

Making Trade-offs Under Risk Explicit

SRAI Executive Pathway Module 2 Canonical Executive Lesson

Control field	Value
Module	EP-M02
Asset	EP-M02-A01
Version	v0.1
Status	G2 content prototype
Primary source	Former Book 6 Chapter 2 Utility and Value
Companion	EP-M02-A08 Companion Notebook v0.1
Terminal artifact	Value Consequence and Risk Profile
Decision	ACCEPT, REVISE or ESCALATE
Date	16 September 2026

Purpose

This module helps executives make defensible choices when alternatives have uncertain consequences, stakeholder impacts and possible severe losses. It preserves the technical foundations of utility and value while translating them into a governed institutional decision process. The aim is not to replace judgment with a formula. The aim is to make value judgments, risk attitudes, information needs and non-negotiable constraints visible enough to inspect, challenge and authorize.

The module culminates in a Value Consequence and Risk Profile. That record documents what outcomes matter, whose values are represented, how uncertainty is treated, which consequences cannot be traded away, whether further information is worth acquiring, and whether the proposed value-and-risk model should be accepted, revised or escalated.

1 Executive learning contract

1.1 Central question

When outcomes, risks, rights and stakeholder priorities differ, how should an institution decide what it values and how much uncertainty or downside it is prepared to accept?

1.2 Learning outcomes

1. Separate outcomes, monetary amounts, institutional value and utility.
2. Identify whose values are represented, who bears consequences and who can authorize the model.
3. Explain why the alternative with the highest expected money is not automatically the rational choice.
4. Use expected utility to compare uncertain prospects under an explicit risk posture.
5. Interpret a certainty equivalent and risk premium.
6. Keep legal rights, mandatory safety limits and unacceptable losses outside compensable scoring.
7. Estimate whether additional information is worth its cost and delay.
8. Test whether a recommendation survives plausible changes in assumptions.
9. Complete and defend a Value Consequence and Risk Profile.

1.3 Prerequisite

Participants should begin with an EP-M01 decision frame that identifies the decision owner, alternatives, stakeholders, constraints, uncertainties, evidence gaps and conditions for PROCEED, REVISE or STOP. Module 2 does not reopen whether AI is necessary. It governs how consequences and risk are valued once the decision has been framed.

2 Opening case

2.1 A choice that cannot be reduced to average money

A public agency must decide whether to deploy an AI-supported inspection system. Alternative A promises larger average savings but carries a small probability of a severe failure that could deny essential services to a vulnerable group. Alternative B produces lower average savings but reduces the severe-loss probability. Alternative C delays deployment for six months to collect more evidence. The decision cannot be defended by pointing only to expected savings.

The agency must decide how it values savings, continuity of service, avoidable harm, equity, reversibility, public trust and delay. It must also determine whether any safety or rights condition is a hard gate. If a severe consequence is legally or ethically unacceptable, no amount of aggregate savings can compensate for it.

2.2 First executive test

Question	Why it matters
What happens under each alternative?	Defines consequences before assigning value.
Who experiences each consequence?	Prevents aggregate results from hiding distributional harm.
Which outcomes can be traded?	Separates compensable preferences from hard constraints.
How uncertain are the estimates?	Prevents point estimates from being treated as facts.
Can new evidence change the action?	Determines whether information has decision value.
Who can approve the risk posture?	Links the model to legitimate authority.

3 Outcomes value and utility

3.1 Three distinct ideas

Concept	Meaning	Example
Outcome	What occurs under an alternative and state of the world.	Cost, lives affected, service delay or error rate.
Value	Why and how much an outcome matters to the institution or affected people.	Continuity of essential services may matter more than equal monetary savings.
Utility	A numerical representation of preferences over outcomes under declared conditions.	A concave function assigns diminishing marginal utility and represents risk aversion over a stated range.

Confusing these concepts creates false objectivity. A model may report an outcome in dollars, but dollars do not automatically encode irreversibility, public obligations, distributional harm or risk attitude. Utility is not an intrinsic property of an outcome. It represents a preference model whose owner, range and purpose must be declared.

3.2 Measurement discipline

- Define the outcome unit and time horizon.
- State whether higher or lower values are preferable.
- Declare the range over which the value or utility function is intended to apply.
- Do not extrapolate a locally elicited utility curve to catastrophic outcomes without justification.
- Keep subgroup consequences visible before aggregation.

4 Expected value as a benchmark

4.1 Formal definition

For an alternative a with possible consequences x_i and probabilities p_i , the expected value is $E[X | a] = \sum_i p_i x_i$. The probabilities must be non-negative and sum to one. The calculation is an average across repeated or comparable uncertainty; it is not a promise that the average will occur in a single decision.

4.2 When expected value is adequate

Expected value is a defensible objective when the numerical outcome scale already represents the decision maker's value over the relevant range, risk attitude is effectively neutral for the stakes involved, distributional and rights constraints are separately satisfied, and the decision is sufficiently repeatable or diversified. These are substantive conditions, not mathematical conveniences.

4.3 Why the benchmark can fail

- A rare catastrophic loss can be hidden by a favorable average.
- Money may not represent nonfinancial objectives or public obligations.
- The same aggregate result can distribute gains and losses very differently.
- An irreversible decision may justify a different risk posture than a reversible pilot.
- A legal right or safety requirement cannot be converted into an ordinary weighted benefit.

5 Utility under uncertainty

5.1 Expected utility

Expected utility evaluates an alternative by $EU(a) = \sum_i p_i u(x_i)$, where $u(x_i)$ is the utility assigned to consequence x_i . A rational choice under the stated model selects the feasible alternative with the greatest expected utility. The result is conditional on the alternatives, probabilities, consequence model, utility function and hard constraints.

5.2 Linear and concave utility

A linear utility function over a declared range represents risk neutrality: equal increments in the outcome produce equal increments in utility. A concave utility function represents diminishing marginal utility and risk aversion over that range. Concavity does not prove that an institution is risk averse in every domain. Risk posture can differ by mandate, consequence type, affected group and reversibility.

Model	Illustrative form	Interpretation
Linear	$u(x) = (x - L) / (U - L)$	Risk-neutral benchmark between lower bound L and upper bound U.
Exponential	$u(x) = (1 - \exp(-r(x-L))) / (1 - \exp(-r(U-L)))$	Constant absolute risk-aversion form after normalization; r controls curvature.
Hard constraint	$g_j(a) \leq \text{limit}_j$	Feasibility gate, not a utility term.

6 Eliciting utility responsibly

6.1 Authority before arithmetic

Analysts may facilitate elicitation, test internal consistency and calculate implications, but they do not own institutional values by default. The decision record must state who authorized the utility model, which affected groups informed it, which evidence supports the parameters and how disagreements were handled.

6.2 Practical elicitation methods

- Certain-equivalent questions: identify the guaranteed outcome considered equivalent to a stated risky prospect.
- Probability-equivalent questions: identify the probability of a best outcome that makes a gamble equivalent to a sure intermediate outcome.
- Pairwise comparisons: compare carefully designed prospects to reveal preference reversals and inconsistencies.
- Policy calibration: compare the implied trade-offs with previous authorized decisions and statutory obligations.
- Sensitivity envelope: use a range of plausible functions when a single curve cannot be legitimately fixed.

6.3 Bias and process risks

Framing, anchoring, scale choice, group power and misunderstanding of probability can distort elicitation. A technically fitted curve does not cure an illegitimate process. Record who participated, which questions were used, whether comprehension was checked, and which disagreements remain unresolved.

7 Worked example

7.1 Prospect definition

Consider outcomes 0, 50 and 100 with probabilities 0.20, 0.50 and 0.30. Expected value is 55. Under normalized linear utility on 0 to 100, utilities are 0, 0.5 and 1, so expected linear utility is 0.55. Now introduce normalized exponential utility with risk parameter $r = 0.02$. The utility of 50 becomes greater than 0.5 because the function is concave, while the marginal gain from 50 to 100 is smaller than the gain from 0 to 50.

Outcome x	Probability p	p x	Linear u(x)	p u(x)
0	.20	0	0	.000
50	.50	25	.5	.250
100	.30	30	1	.300
Total	1.00	55		.550

7.2 Interpretation

The expected value is a benchmark, not a complete authorization. Under a concave utility model, the same prospect has a certainty equivalent below 55 and therefore a positive risk premium. The companion notebook calculates the exact values and varies r to show how the implied risk posture changes.

8 Certainty equivalent and risk premium

8.1 Definitions

The certainty equivalent CE is the guaranteed outcome that provides the same utility as the uncertain prospect: $u(CE) = EU(a)$. If u is invertible, $CE = u^{-1}(EU(a))$. The risk premium is $\pi = E[X] - CE$. For a risk-averse model, CE is normally below expected value and π is positive.

8.2 Executive interpretation

The risk premium is not an insurance price automatically payable in the market. It is a model-based expression of how much expected outcome the represented decision maker would surrender to remove uncertainty. Its meaning depends on the outcome scale, utility curve, valid range and decision context.

8.3 Control questions

- Is the utility function monotonic and invertible over the relevant range?
- Does the certainty equivalent remain inside the valid outcome range?
- Was the risk parameter elicited or merely assumed?
- Does the risk premium remain stable under plausible parameter changes?
- Would a severe subgroup harm be hidden by the aggregate certainty equivalent?

9 Rights safety and non-compensatory constraints

9.1 What must remain outside utility

Legal rights, mandatory safety requirements, prohibited discrimination, minimum service obligations and other unacceptable losses must remain feasibility conditions. They should not be converted into low utility values that can be offset by enough benefit elsewhere. An alternative that violates a hard constraint is infeasible even if its expected utility is high.

9.2 Aggregation and distribution

Institutional or social utility can conceal who benefits and who bears harm. Before aggregation, display consequences by affected group, geography, vulnerability, time horizon and relevant protected status. After aggregation, retain the disaggregated table as part of the evidence record.

Condition	Treatment	Decision consequence
Legal or rights violation	Hard constraint	Reject or escalate; do not score as a trade-off.
Mandatory safety threshold	Hard constraint	Alternative is infeasible until corrected.
Material subgroup harm	Disaggregated consequence plus governance review	Revise safeguards, compensate where lawful, or escalate.
Ordinary preference trade-off	Utility or multi-criteria treatment	Compare transparently and test sensitivity.

10 Value of information

10.1 Decision value rather than curiosity

Information has value when it can change the selected action or improve how the action is implemented. Gross value of information compares the optimized expected result with the information to the optimized result without it. If the decision maker may ignore unhelpful evidence, gross value cannot be negative before acquisition cost.

10.2 Net value

Net information value subtracts the monetary and nonmonetary cost of acquisition, analysis, governance and delay. A study with positive gross value can still be a poor choice if it arrives after the decision window or exposes people to unacceptable risk while evidence is collected.

Component	Question
Decision sensitivity	Could the evidence change the preferred feasible alternative?
Gross improvement	How much does optimized expected utility improve?
Acquisition cost	What money, staff time and operational burden are required?
Delay cost	What is lost or harmed while waiting?
Implementation value	Can the institution act on the result?
Net decision	Acquire now, acquire later, or act without more evidence?

11 Sensitivity analysis

11.1 The recommendation is conditional

A single expected-utility result can create false confidence. Sensitivity analysis asks whether the recommendation changes when contested inputs move across plausible ranges. The aim is not to produce every imaginable scenario. It is to expose the assumptions that control the decision.

11.2 Required tests

- Probability sensitivity, especially for severe-loss states.
- Utility-shape and risk-parameter sensitivity.
- Outcome-range and normalization sensitivity.
- Catastrophe-consequence sensitivity.
- Information-cost and acquisition-delay sensitivity.
- Subgroup-distribution sensitivity.
- Constraint-threshold sensitivity.

11.3 Decision regions

Report ranges in which each alternative is preferred, points at which rankings reverse, and assumptions that do not affect the decision. A robust recommendation remains acceptable across the plausible envelope. A fragile recommendation requires revision, further information or escalation.

12 Failure modes

Failure mode	Why it fails	Control
Highest expected money wins	Ignores risk attitude, rights and nonfinancial consequences.	Compare expected value with utility and constraints.
Utility curve chosen by analyst	Confuses calculation authority with value authority.	Name the approving authority and elicitation process.
Hard limits entered as low scores	Allows prohibited harms to be compensated.	Model them as feasibility gates.
One parameter reported	Hides fragility.	Publish sensitivity ranges and reversal points.
Aggregate utility only	Conceals distributional harm.	Retain subgroup consequence tables.
Information valued without delay	Overstates research value.	Calculate net value after cost and delay.
Notebook result treated as authorization	Automates a governance decision.	Require named human approval and decision record.

12.1 AI Reasoning Clinic

Claim: “The alternative with the highest expected money is rational.” Evaluation: the claim is defensible only if money correctly represents value over the relevant range, risk neutrality is authorized, hard constraints are satisfied, distributional consequences are acceptable and uncertainty is credible. Without those conditions, the statement is an unsupported simplification.

13 Sector applications

Sector	Decision example	Protected concern
Agriculture	Allocate scarce inspection or extension resources under uncertain crop risk.	Food security, smallholder distribution and seasonal irreversibility.
Education	Choose an early-warning intervention under uncertain benefit.	False positives, stigma, access and student rights.
Health	Select a screening or triage policy.	Patient safety, subgroup performance and clinical authority.
Habitat	Prioritize settlements for infrastructure or risk mitigation.	Displacement, environmental justice and long-horizon consequences.
Public finance	Sequence controls or audits under uncertain recovery.	Due process, service continuity and institutional mandate.

13.1 Transfer rule

The mathematics can transfer across sectors, but the value authority, consequence units, hard constraints, uncertainty sources and acceptable risk posture do not transfer automatically. Each application requires a sector-specific Value Consequence and Risk Profile.

14 Value Consequence and Risk Profile

14.1 Controlled fields

Field	Required content
Decision identity	Owner, date, version and EP-M01 frame reference.
Alternatives	Feasible choices and status quo.
Consequences	Units, time horizons, affected groups and sources.
Value authority	Whose values, elicitation method and approving authority.
Risk posture	Utility form, range, parameter and justification.
Hard constraints	Rights, safety, legal and non-compensatory limits.
Calculations	Expected value, expected utility, CE and risk premium.
Information	Evidence gaps, gross value, cost, delay and net value.
Sensitivity	Plausible ranges, ranking reversals and fragile assumptions.
Decision	ACCEPT, REVISE or ESCALATE; approver, review date and triggers.

14.2 Decision meanings

- **ACCEPT:** the value authority, model, constraints, information decision and sensitivity evidence are adequate for the stated use.
- **REVISE:** correctable weaknesses remain in elicitation, parameters, evidence, safeguards or documentation.
- **ESCALATE:** authority is unclear, stakeholder conflict is unresolved, or a rights, safety or legitimacy issue cannot be aggregated or corrected at the current level.

15 Questions executives should ask technical teams

1. What exact outcome is being modeled, in what unit and over what time horizon?
2. Whose preferences does the utility function represent?
3. How was the utility function elicited, fitted or selected?
4. Over what range is it valid, and what happens outside that range?
5. Which consequences are hard constraints rather than tradeable objectives?
6. How were probabilities estimated and calibrated?
7. Which assumptions cause the recommendation to reverse?
8. How are severe losses and subgroup harms displayed?
9. What is the certainty equivalent and how should it be interpreted?
10. What information could change the decision, and what would it cost or delay?
11. Can the notebook be rerun from a clean environment?
12. Who remains accountable for the final decision?

16 Companion notebook guide

The companion notebook is an advisory analytical instrument. It implements the equations, reproduces the worked example, compares risk-neutral and risk-averse choices, calculates certainty equivalents and risk premiums, estimates information value, runs sensitivity tests and exports an executive summary. It does not issue an institutional authorization.

Notebook section	Decision purpose
Environment and invariants	Establish reproducibility and validate probabilities.
Alternatives and consequences	Make the decision data visible.
Expected value	Provide a neutral benchmark.
Utility functions	Represent declared risk posture.
Expected utility comparison	Compare feasible alternatives.
Certainty equivalent and risk premium	Translate risk posture into interpretable quantities.
Information value	Test whether evidence is worth cost and delay.
Sensitivity	Expose recommendation reversals.
Distribution and constraints	Prevent invalid aggregation.
Executive export and gate	Summarize evidence for human review.

17 Practice exercises

1. Concept distinction. For a proposed AI inspection system, list three outcomes, three value judgments and one possible utility representation. Explain why they are not interchangeable.
2. Expected value. Calculate the expected value of outcomes 0, 50 and 100 with probabilities 0.20, 0.50 and 0.30. State two reasons the result is insufficient for authorization.
3. Expected utility. Using normalized linear utility on 0 to 100, calculate expected utility. Then explain qualitatively how a concave function changes the interpretation.
4. Certainty equivalent. Given an expected value of 55 and a certainty equivalent of 44, calculate the risk premium and explain what it means.
5. Constraints. Classify each as tradeable or non-compensatory: budget overrun, prohibited discrimination, small implementation delay, mandatory safety threshold and reputational discomfort.
6. Information value. A study can improve optimized decision value from 76 to 84, costs 5 units and creates a delay cost of 2. Calculate gross and net value. Recommend whether to acquire the information.
7. Sensitivity. Identify four assumptions to vary before accepting a recommendation based on exponential utility.
8. Executive decision. Complete a Value Consequence and Risk Profile and choose ACCEPT, REVISE or ESCALATE.

18 Knowledge check

1. Why is expected money not always the correct decision objective?
2. What does concave utility represent, and over what scope?
3. How is a certainty equivalent defined?
4. Why is gross value of information non-negative when evidence may be ignored?
5. Which outcomes should remain hard constraints?
6. What does a positive risk premium indicate?
7. Why must distributional effects remain visible after aggregation?
8. What distinguishes ACCEPT from ESCALATE in the terminal decision?

18.1 Answer standard

A satisfactory answer must state the result, show the reasoning, identify assumptions and uncertainty, and describe the governance consequence. Numerical accuracy without a defensible interpretation is incomplete.

19 Synthesis

Utility analysis is valuable because it makes consequence trade-offs and risk posture explicit. Its legitimacy depends on what it refuses to hide: value authority, affected groups, valid ranges, severe losses, hard constraints, information gaps and sensitivity. A precise calculation built on an unauthorized value model is not a defensible decision.

19.1 SRAI Principle 142

Utility makes contextual value and risk judgments explicit but cannot convert rights or mandatory safety limits into compensable scores.

19.2 Decision chain

Stage	Required evidence
Frame	Decision, owner, alternatives, affected groups and constraints.
Represent	Outcomes, values, utility and risk posture.
Calculate	Expected value, expected utility, CE, risk premium and information value.
Challenge	Sensitivity, severe-loss and distributional tests.
Govern	Hard constraints, value authority and accountability.
Decide	ACCEPT, REVISE or ESCALATE with review triggers.

20 Glossary

Term	Definition
Certainty equivalent	Guaranteed consequence with utility equal to the expected utility of a prospect.
Expected utility	Probability-weighted average of utilities under an alternative.
Expected value	Probability-weighted average of numerical consequences.
Hard constraint	Requirement that determines feasibility and cannot be compensated by gains elsewhere.
Net value of information	Gross decision improvement minus acquisition cost, burden and delay.
Risk aversion	Preference represented by concave utility over a declared range.
Risk premium	Expected value minus certainty equivalent.
Utility	Numerical representation of preferences under stated conditions.
Value authority	Person or body legitimately empowered to approve represented values and risk posture.
Value Consequence and Risk Profile	Controlled executive record supporting the Module 2 decision.

21 References and provenance

- Primary SRAI source: Decision Intelligence Complete Edition v1.0, Chapter 2 Utility and Value.
- Primary notebook source: V6_N02_utility_value.ipynb.
- von Neumann, J. and Morgenstern, O. Theory of Games and Economic Behavior.
- Savage, L. J. The Foundations of Statistics.
- Raiffa, H. Decision Analysis Introductory Lectures on Choices Under Uncertainty.
- Keeney, R. L. and Raiffa, H. Decisions with Multiple Objectives Preferences and Value Tradeoffs.
- Howard, R. A. Information Value Theory.
- SRAI Executive Pathway Module 1 accepted structural baseline and EP-M02 G1 architecture v0.1.

21.1 Technical preservation statement

This lesson retains and expands every technical concept in former Book 6 Chapter 2: outcomes versus utility, expected value limitations, linear and concave utility, expected utility, certainty equivalent, risk premium, value authority, valid range, distributional effects, hard rights and safety constraints, value of information, catastrophe sensitivity, practice and assessment. The companion notebook makes calculations and sensitivity visible and reproducible.

Looking ahead

Module 3, Deciding When Evidence Is Incomplete, develops Bayesian updating, posterior beliefs, asymmetric loss and formal value-of-information analysis. Module 2 introduces the information decision but does not pre-empt that treatment.