Lecture 7: Decision Theory and Risk
Mathematical Modeling Lecture

Mathematical Modeling

Kenneth Sok Kin Cheng

Session Route

Launch	Build	Test
Question Motivating context, key vocabulary, and the first modelling decision.	**Method** Worked example, diagram or table, calculation, and parameter meaning.	**Critique** Checkpoint practice, limitation check, and communication of results.

Learning Goals

By the end of the lesson, readers should be able to:

1. distinguish a **decision alternative** from a **state of nature**;
2. compute and interpret **expected value** for repeated risky decisions;
3. build and solve a simple **decision tree** using the fold-back method;
4. use **Bayes' theorem** to update probabilities after new information;
5. compute the **expected value of perfect information**;
6. apply Laplace, Maximin, Maximax, Hurwicz, and Minimax Regret criteria when probabilities are unknown;
7. explain why the mathematically optimal decision may depend on risk tolerance, repetition, and consequences.

Why decision theory?

Core explanation, worked reasoning, and application

Why decision theory?

Previous lectures often assumed that the model inputs were known. In real decisions, some inputs are uncertain: the weather may change, market demand may be high or low, a bid may be accepted or rejected, and a medical test may be correct or misleading. **Decision theory** provides a disciplined way to choose actions when outcomes depend on uncertain future states.
A decision model separates three ingredients:

1. **Alternatives**: actions the decision maker can choose.
2. **States of nature**: uncertain conditions not controlled by the decision maker.
3. **Payoffs**: consequences produced by each alternative-state combination.

Activity: Launch: What should the stall sell?

A school fundraising stall can sell either hot coffee or iced tea. Coffee earns more on a cold day; iced tea earns more on a warm day. Tomorrow's weather is uncertain.
In pairs, discuss:

1. What are the alternatives? What are the states of nature?
2. What information would change the decision?
3. Is this a repeated decision or a one-time decision?
4. Would you rather maximize average profit or avoid a very bad outcome?

Modeling lesson: Before computing anything, identify who chooses, what is uncertain, and what consequence matters.

Illustration: Why decision theory?

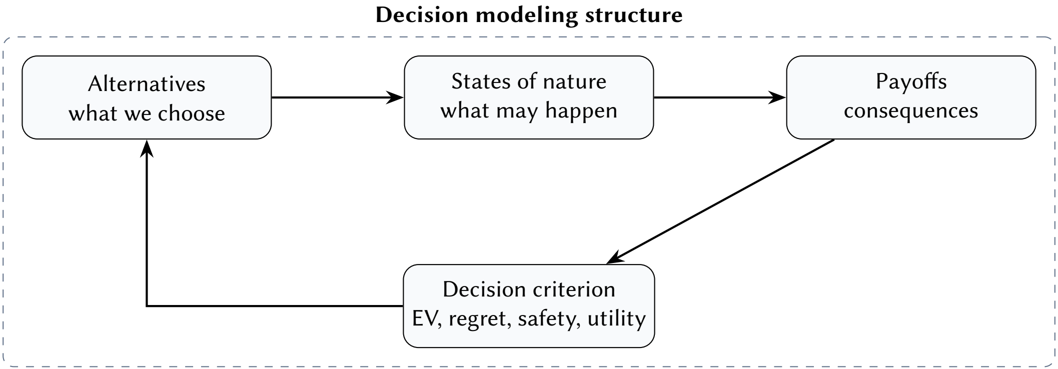

A decision model combines choices, uncertain states, consequences, and a criterion for judging them.

Concept: Repeated decisions versus one-time decisions

Expected value is a long-run average. It is most convincing when the same kind of decision is repeated many times, such as an insurance company selling many policies. For a one-time decision with large downside risk, the highest expected value may not be the most acceptable choice.

This distinction is not a minor philosophical issue. It determines whether we should use expected value, a risk-averse criterion, or a regret-based criterion.

Expected value for repeated risky decisions

Core explanation, worked reasoning, and application

Expected value for repeated risky decisions

If a random outcome has payoffs $w_1, w_2, \ldots, w_n$ with probabilities $p_1, p_2, \ldots, p_n$, where $\sum p_i = 1$, its **expected value** is

$$E = \sum_{i=1}^{n} p_i w_i.$$

The value E is not necessarily a possible outcome. It is the average payoff expected over many repetitions under the same probability model.

Worked Example: Worked Example 1: A dice game

A game costs \$1 to play. If two fair dice sum to 7, the player receives \$6 total; otherwise the player receives nothing. The net payoff is therefore \$5 for a sum of 7 and $-\$1$ otherwise.
There are 36 equally likely dice outcomes and 6 of them sum to 7, so

$$P(\text{sum} = 7) = \frac{6}{36} = \frac{1}{6}.$$

The expected net payoff is

$$E = 5\left(\frac{1}{6}\right) + (-1)\left(\frac{5}{6}\right) = 0.$$

The game is **fair** in the expected-value sense. This does not mean every player breaks even; it means that the average net gain approaches zero over many plays.

Model Critique: Expected value is not a promise

A student might say, "The expected value is zero, so I will neither win nor lose." This is wrong. A single play gives either \$5 or $-\$1$. Expected value describes the centre of the distribution over repeated trials, not the outcome of one trial.

Worked Example: Worked Example 2: Concession decision with known weather probability

A concession firm must choose one product for tomorrow's school sports day.

	Cold day	**Warm day**
Coffee	\$4000	\$1000
Iced tea	\$1500	\$5000

The forecast gives $P(\text{cold}) = 0.30$, so $P(\text{warm}) = 0.70$.

$$E(\text{coffee}) = 0.30(4000) + 0.70(1000) = 1900.$$

$$E(\text{iced tea}) = 0.30(1500) + 0.70(5000) = 3950.$$

The expected-value decision is to sell iced tea.
The **break-even probability** for cold weather satisfies

$$4000p + 1000(1 - p) = 1500p + 5000(1 - p).$$

Solving gives $p = \frac{4000}{6500} \approx 0.615$. Coffee becomes preferable only if the probability of a cold day exceeds about 61.5%.

Checkpoint: Expected value

A bid costs \$2000 to prepare. If accepted, it earns a net profit of \$18,000 after subtracting preparation cost. If rejected, the company loses the \$2000 preparation cost.

1. Write the expected value in terms of the acceptance probability p.
2. Find the break-even value of p.
3. Explain whether the company should bid if $p = 0.15$.

Decision trees and the fold-back method

Core explanation, worked reasoning, and application

Decision trees and the fold-back method

A **decision tree** is useful when a decision is followed by one or more uncertain events. We use square nodes for decisions, circular nodes for uncertainty, and terminal values for payoffs.

Concept: Fold-back method

Solve a decision tree from right to left:

1. At each chance node, compute the expected value of its branches.
2. At each decision node, keep the branch with the best expected value.
3. Replace the solved node by its value and continue moving left.

This method is called **fold-back**, because we simplify the future part of the tree before choosing the earlier action.

Illustration: Decision trees and the fold-back method

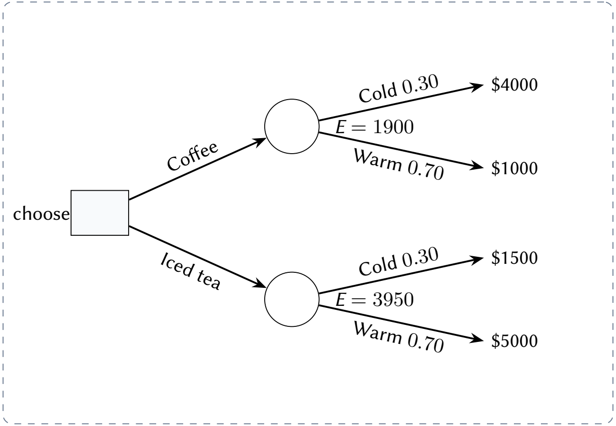

Decision tree for the concession example. Fold-back chooses the branch with the larger expected value.

Worked Example: Worked Example 3: Plant-size decision

A company may build a large plant, build a small plant, or build nothing. Demand may be high, moderate, or low.

Alternative	**High** (0.25)	**Moderate** (0.40)	**Low** (0.35)
Large plant	200000	100000	-120000
Small plant	90000	50000	-20000
No plant	0	0	0

$$E(\text{large}) = 0.25(200000) + 0.40(100000) + 0.35(-120000) = 48000.$$

$$E(\text{small}) = 0.25(90000) + 0.40(50000) + 0.35(-20000) = 35500.$$

$$E(\text{none}) = 0.$$

By expected value, the large plant is best. However, the low-demand loss is $120,000, so the recommendation should mention risk, not only the average.

Model Critique: Risk profile

Two alternatives can have similar expected values but very different downside risk. A complete decision analysis should report:

1. expected value;
2. best-case and worst-case payoff;
3. probability of loss;
4. whether the decision is repeated or one-time;
5. the decision maker's risk tolerance.

Information, Bayes' theorem, and EVPI

Core explanation, worked reasoning, and application

Information, Bayes' theorem, and EVPI

Information is valuable only if it can change the decision. Decision theory allows us to place an upper bound on the value of information and, when the information is imperfect, to compute whether it is worth buying.

Concept: Perfect information

The **expected value of perfect information** is

$$\text{EVPI} = E(\text{best action after knowing the state}) - E(\text{best action now}).$$

It is the most we should ever pay for flawless information. Any imperfect forecast or survey must be worth no more than EVPI.

Worked Example: Worked Example 4: EVPI for the plant-size decision

Without information, the best expected value is \$48,000 from the large plant.
With perfect information, the company would choose the best action after seeing the true demand state:

State	Best action	Payoff
High	Large plant	\$200000
Moderate	Large plant	\$100000
Low	No plant	\$0

Therefore,

$$E(\text{perfect information}) = 0.25(200000) + 0.40(100000) + 0.35(0) = 90000.$$

$$\boxed{\text{EVPI} = 90000 - 48000 = 42000.}$$

No market research report, even a perfect one, is worth more than \$42,000 for this decision.

Information, Bayes' theorem, and EVPI

When information is imperfect, we update probabilities using **Bayes' theorem**. If $H_1, \ldots, H_n$ are possible states and F is evidence, then

$$P(H_i \mid F) = \frac{P(F \mid H_i) P(H_i)}{\sum_j P(F \mid H_j) P(H_j)}.$$

Illustration: Information, Bayes' theorem, and EVPI

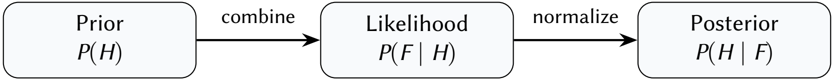

Bayes update: new evidence changes the probability model before the decision is evaluated.

Bayes' theorem updates beliefs from prior probabilities to posterior probabilities using evidence.

Worked Example: Worked Example 5: Market research with Bayes update

For the plant-size problem, suppose a market report is either favorable F or unfavorable U. Its historical accuracy is:

State	$P(F \mid \text{state})$	$P(U \mid \text{state})$
High	0.70	0.30
Moderate	0.50	0.50
Low	0.20	0.80

First compute

$$P(F) = 0.70(0.25) + 0.50(0.40) + 0.20(0.35) = 0.445.$$

The posterior probabilities after a favorable report are

$$P(H \mid F) = \frac{0.70(0.25)}{0.445} \approx 0.393, \quad P(M \mid F) = \frac{0.50(0.40)}{0.445} \approx 0.449, \quad P(L \mid F) = \frac{0.20(0.35)}{0.445} \approx 0.157.$$

Re-evaluating the large plant after a favorable report gives

$$E(\text{large} \mid F) \approx 0.393(200000) + 0.449(100000) + 0.157(-120000) = 104640.$$

Worked Example: Worked Example 5: Market research with Bayes update

Small plant gives about \$54,690, so a favorable report supports the large plant.
For an unfavorable report, $P(U) = 0.555$ and

$$P(H \mid U) \approx 0.135, \quad P(M \mid U) \approx 0.360, \quad P(L \mid U) \approx 0.505.$$

Then

$$E(\text{large} \mid U) \approx 2400, \qquad E(\text{small} \mid U) \approx 20020,$$

so an unfavorable report supports the small plant. The report is useful because it changes the decision in some cases.

Checkpoint: Value of information

In the market-research example, after a favorable report the decision is large plant, and after an unfavorable report the decision is small plant.

1. Compute the expected value if the company uses the report: $0.445E_F + 0.555E_U$.
2. Subtract the best no-report expected value \$48,000 to obtain EVSI.
3. Should the company buy the report if it costs \$10,000? What if it costs \$25,000?

When probabilities are unknown

Core explanation, worked reasoning, and application

When probabilities are unknown

Sometimes we cannot assign credible probabilities. In that case, expected value may give a false impression of precision. We use alternative criteria that represent different attitudes toward uncertainty.

Worked Example: Running Example: Investment payoffs

A student club has funds to invest for one year. Payoffs are in thousand dollars.

Alternative	Fast growth	Normal growth	Slow growth
Stocks	10	6.5	-4
Bonds	8	6	1
Savings	5	5	5

No reliable probabilities are available for the three economic states.

When probabilities are unknown

Criterion	Rule	Decision in example
Laplace	Treat all states as equally likely and average each row.	Bonds/Savings tie at 5.0.
Maximin	For each row, find the worst payoff; choose the best worst payoff.	Savings.
Maximax	For each row, find the best payoff; choose the best best payoff.	Stocks.
Hurwicz	Use $\alpha(\text{best}) + (1 - \alpha)(\text{worst})$.	Depends on α; at $\alpha = 0.6$, Bonds.
Minimax regret	Build a regret matrix and minimize the largest regret.	Bonds.

Illustration: When probabilities are unknown

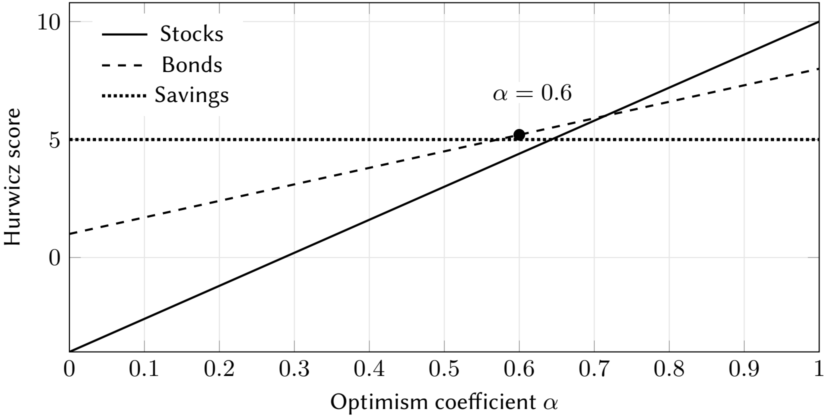

Hurwicz scores change with the optimism coefficient. The chosen alternative depends on the decision maker's attitude toward risk.

Worked Example: Worked Example 6: Minimax regret

For each state, subtract each payoff from the best payoff in that column.

	Fast	Normal	Slow	Max regret
Stocks	$10 - 10 = 0$	$6.5 - 6.5 = 0$	$5 - (-4) = 9$	9
Bonds	$10 - 8 = 2$	$6.5 - 6 = 0.5$	$5 - 1 = 4$	4
Savings	$10 - 5 = 5$	$6.5 - 5 = 1.5$	$5 - 5 = 0$	5

The smallest maximum regret is 4, so the minimax-regret decision is Bonds. This criterion asks: "If the true state is later revealed, which decision limits how much we will regret not having chosen the best action for that state?"

Activity: Decision studio

15 min

Use the investment example.

1. **Core route:** Apply Laplace, Maximin, Maximax, and Minimax Regret. Write one sentence explaining each attitude.
2. **Challenge route:** Find the α values where Hurwicz changes from Savings to Bonds and from Bonds to Stocks.
3. **Extension route:** Suppose the club cannot tolerate any possibility of loss. Which criteria respect this constraint? Which criterion ignores it?

Integrated mini-case: stocking under uncertain demand

Core explanation, worked reasoning, and application

Integrated mini-case: stocking under uncertain demand

Decision theory becomes most useful when the modeler must connect a business rule, uncertain demand, and a recommendation. The following example is intentionally small enough to compute by hand, but it contains the same structure as larger inventory decisions.

Worked Example: Worked Example 7: Seasonal T-shirt stocking

A school club can order 10, 20, or 30 festival T-shirts. Each shirt costs \$50 and sells for \$80. Unsold shirts can be cleared for \$20 each. If demand exceeds stock, each unmet demand unit causes a goodwill loss of \$10. Demand is estimated as:

$$P(D = 10) = 0.30, \qquad P(D = 20) = 0.40, \qquad P(D = 30) = 0.30.$$

For an order quantity Q and demand D, the payoff is

$$80 \min(Q, D) + 20 \max(Q - D, 0) - 50Q - 10 \max(D - Q, 0).$$

This gives the payoff matrix:

Order quantity	$D = 10$	$D = 20$	$D = 30$	**Expected value**
$Q = 10$	300	200	100	$0.30(300) + 0.40(200) + 0.30(100) = 200$
$Q = 20$	0	600	500	$0.30(0) + 0.40(600) + 0.30(500) = 390$
$Q = 30$	-300	300	900	$0.30(-300) + 0.40(300) + 0.30(900) = 300$

By expected value, the best order is $Q = 20$.
Different risk criteria tell a more nuanced story:

Worked Example: Worked Example 7: Seasonal T-shirt stocking

- Maximin chooses $Q = 10$, because its worst payoff is still \$100.
- Maximax chooses $Q = 30$, because it can earn \$900 if demand is high.
- Minimax regret chooses $Q = 20$: its maximum regret is \$400, smaller than the maximum regret for $Q = 10$ or $Q = 30$.

Thus the same payoff matrix can support different recommendations depending on whether the club prioritizes average profit, downside protection, or regret control.

Activity: Sensitivity check

8 min

In the T-shirt example, suppose the probability of high demand $D = 30$ increases while the probability of low demand $D = 10$ decreases. Without recalculating every criterion, predict which order quantity becomes more attractive. Then compute the expected values when

$$P(D = 10) = 0.20, \qquad P(D = 20) = 0.40, \qquad P(D = 30) = 0.40.$$

Explain whether the decision changes.

Model Critique: Hidden assumptions in payoff matrices

A payoff matrix looks precise, but each number depends on assumptions: selling price, clearance price, lost goodwill cost, and the demand probabilities. If these inputs are estimated weakly, the report should include sensitivity checks rather than present the final order quantity as certain.

Writing a defensible decision recommendation

Core explanation, worked reasoning, and application

Writing a defensible decision recommendation

Decision theory does not remove judgment. It makes the judgment explicit. In a modeling report, the recommendation should not simply state the branch with the largest expected value; it should state the assumptions under which the recommendation holds.

Model Critique: Common mistakes in decision models

1. Treating guessed probabilities as if they were measured facts.
2. Using expected value for a one-time catastrophic decision without discussing risk.
3. Forgetting to subtract the cost of information from EVSI.
4. Confusing a decision node with a chance node: the modeler chooses at one; nature decides at the other.
5. Reporting a criterion result without explaining the risk attitude represented by the criterion.

Worked Example: Competition-style conclusion

For the plant-size case, a concise conclusion might be:
Using the prior demand probabilities, the large plant has the highest expected value, $48,000, and is therefore the recommended decision if management is willing to accept a possible $120,000 loss in the low-demand state. Perfect information would be worth at most $42,000, so any market research costing more than this should be rejected immediately. Because the large-plant recommendation is sensitive to the low-demand loss and to the estimated probability of high demand, we also recommend reporting a risk profile and considering the small plant if management is strongly risk-averse.

Summary: Lesson summary

1. A decision model separates alternatives, states of nature, probabilities, and payoffs.
2. Expected value is a long-run average and is most appropriate for repeated decisions.
3. Decision trees are solved by fold-back: expected values at chance nodes, best choices at decision nodes.
4. Bayes' theorem updates probabilities after evidence.
5. EVPI is an upper bound on the value of information; EVSI must be compared with the cost of the information.
6. When probabilities are unknown, different criteria represent different attitudes toward risk and uncertainty.

Checkpoint: Exit ticket

2 min

Complete the sentence:
"Expected value is a useful decision criterion when ______________________________, but it may be misleading when ______________________________."

Practice problems

Core explanation, worked reasoning, and application

Exercises: Core practice

1. A game costs \$2 to play. A player wins \$12 with probability 0.20 and wins nothing otherwise. Compute the expected net payoff. Is the game fair?
2. A food stall must choose noodles or sandwiches. Profits are:

	Rainy $(p = 0.40)$	Sunny $(p = 0.60)$
Noodles	3200	1600
Sandwiches	1800	3600

   Compute both expected values and find the break-even probability of rain.
3. For the plant-size problem in Worked Example 3, compute the probability of losing money under each alternative. Does this change how you would present the recommendation?
4. A medical screening test has $P(+ \mid D) = 0.95$, $P(+ \mid \text{no } D) = 0.08$, and disease prevalence $P(D) = 0.02$. Compute $P(D \mid +)$. Explain why a highly sensitive test can still have a moderate positive predictive value.
5. For the investment payoff matrix, apply Hurwicz with $\alpha = 0.3$ and $\alpha = 0.8$. Explain why the decision changes.
6. For the following payoff matrix, construct the regret matrix and apply minimax regret:

Exercises: Core practice

	S_1	S_2	S_3
A	20	5	8
B	12	12	12
C	4	18	10

Exercises: Modeling challenge

Choose one of the following and prepare a one-page decision analysis brief containing: alternatives, states of nature, payoff table, probability assumptions if any, chosen criterion, recommendation, sensitivity comment, and one limitation.

1. A school club deciding whether to hold an outdoor fundraising event, indoor event, or cancel.
2. A student team deciding whether to submit a safe project, a risky innovative project, or a hybrid project to a modeling competition.
3. A small shop deciding whether to stock 10, 20, or 30 units of a seasonal item.

Your answer should explain why the chosen criterion is appropriate for the situation.

Optional extension: expected utility

Core explanation, worked reasoning, and application

Optional extension: expected utility

Expected value treats each dollar equally. In reality, the pain of losing \$100,000 may be more than twice the pain of losing \$50,000. **Expected utility** replaces payoff w by a utility function $u(w)$ and chooses the alternative with largest

$$E[u(W)] = \sum_i p_i u(w_i).$$

A concave utility function represents risk aversion: gaining an extra dollar matters less when the decision maker is already wealthy, while losses near a critical threshold may be unacceptable. This is why a risk-averse manager may choose the small plant even when the large plant has a larger expected monetary value.

Optional extension: expected value of sample information

Core explanation, worked reasoning, and application

Optional extension: expected value of sample information

If a report can be favorable or unfavorable, the expected value with sample information is

$$E(\text{with report}) = P(F)\max_a E(a \mid F) + P(U)\max_a E(a \mid U).$$

Then

$$\text{EVSI} = E(\text{with report}) - E(\text{best decision without report}).$$

The report should be purchased only when EVSI exceeds its cost. If the cost is included in terminal payoffs, do not subtract it a second time.