13 *Sequential Decisions: Trees and the Value of Information*

After this chapter you will be able to:

1. distinguish a decision node (square) from a chance node (circle) in a decision tree
2. compute the expected monetary value $\text{EMV} = \sum_k p_k V_k$ at a chance node
3. build a two-stage tree and roll it back, taking the maximum at each square and the expected value at each circle
4. compute the expected value of perfect information (EVPI), and recognize it as a ceiling on the value of any test
5. convert a prior probability and a test's reliability into posterior probabilities with Bayes' rule
6. compute the expected value of sample information (EVSI), and decide whether to buy the test at a stated price

chapter 12 treated an uncertain project as a single lottery: attach probabilities to the outcomes, take the expected present worth, and decide once. That is the right picture when the whole plan is committed up front. But many real decisions are not made all at once. A firm builds a pilot before it builds a plant, drills one hole before it drills the field, signs a first customer before it scales the factory. Each early step is cheap, and its purpose is to *learn* something before the expensive step is committed. The value of that learning is the subject of this chapter. We build the decision *tree*, a diagram that separates the moves the decision maker controls from the moves nature controls, and we evaluate it by *rollback*: work from the tips back to the root, taking expectations where nature acts and the best choice where the decision maker acts. Two numbers fall out that were invisible in chapter 12: the value of perfect information, which caps what any study could ever be worth, and the value of an imperfect test, which is what a real, fallible pilot or survey is worth once its reliability is priced through Bayes' rule. This chapter connects backward to the expected-present-worth machinery of chapter 12 and forward to the flexibility valuation of chapter 15, where the same trees acquire a market discipline.

Which node takes the maximum.

Q1. A decision tree has two kinds of node: square nodes, where the decision maker chooses, and round nodes, where nature chooses. When you roll the tree back, at which node do you take the *maximum* of the branches, and at which do you take the *expected value*?

Q2. The expected value of perfect information (EVPI) is the most a firm should pay to learn the true state of the world before deciding. Can the EVPI ever be negative?

Q3. A geological survey is not perfectly reliable: it sometimes says ''favorable'' when the ore is poor. Is such an imperfect survey worth more or less than perfect information?

Answers.

Q1. Take the **maximum** at a **square** (decision) node — the decision maker will pick the best branch. Take the **expected value** at a **round** (chance) node — nature picks, and the decision maker gets a probability-weighted average of the branches. Reversing these, taking a maximum where nature acts, is the most common tree error, and it values a project as if the decision maker could choose which way the dice land.

Q2. *No.* Perfect information can only help: in the worst case the decision maker ignores it and does exactly what they would have done anyway, so the value of information is at least zero. The EVPI is therefore never negative, and it is a *ceiling* — no test, however good, can be worth more than knowing the truth outright.

Q3. *Less.* A fallible test leaves some uncertainty unresolved, so it cannot be worth more than eliminating the uncertainty entirely. Its value, the expected value of sample information (EVSI), sits between zero and the EVPI. The closer the test's reliability is to perfect, the closer its EVSI creeps toward the EVPI ceiling.

13.1 *Decision Nodes, Chance Nodes, and Rollback*

A decision tree is a drawing of a sequential decision, read left to right in time. It has exactly two kinds of node, and the whole method depends on never confusing them.

The branches leaving a decision node are the *alternatives*; the decision maker will take whichever leads to the highest worth, so a decision node is worth the *maximum* over its branches. The branches leaving a chance node are the *states of nature*; the decision maker cannot choose among them, so a chance node is worth the probability-weighted *expected value* of its branches. *Rollback* applies these two rules from the tips inward: evaluate the rightmost nodes first, replace each with its worth, and repeat until only the root remains.

Reference. Park, 3rd ed., Chapter 11 (handling project uncertainty); decision trees extend the expected-value analysis of chapter 12.

Decision node. Drawn as a **square**. A point where the decision maker chooses one branch. Evaluated by taking the *best* (maximum worth) branch.

Chance node. Drawn as a **circle**. A point where nature chooses, with each branch carrying a probability. Evaluated by taking the *expected value* over the branches.

Rollback. Evaluating a tree from its tips back to its root: expected value at each chance node, maximum at each decision node. The root value is the worth of the whole sequential decision.

Expected monetary value (EMV). The probability-weighted average worth at a chance node: EMV $= \sum_k p_k V_k$. The same expected present worth of chapter 12, now sitting at a node in a tree.

Example 13.1 — A one-stage tree: build or decline

Najd Copper can build a full leaching plant on a new deposit, or decline. If it builds, the ore grade is uncertain: with probability 0.55 the grade is high and the plant's present worth is $+80$ million SAR; with probability 0.45 the grade is low and the present worth is -40 million SAR. Declining is worth 0. Draw the tree and decide.

Solution.

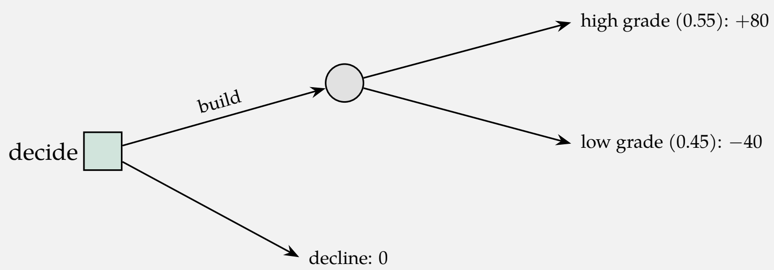

The chance node (circle) is worth its expected value:

$$\text{EMV(build)} = 0.55\,(80) + 0.45\,(-40) = 44 - 18 = 26 \text{ million SAR}.$$

The decision node (square) takes the better of build (26) and decline (0):

$$\text{root value} = \max(26,\ 0) = \boxed{\text{26 million SAR}} \ \Rightarrow\ \text{build now.}$$

Building looks good on this one-shot analysis. But notice the low-grade branch loses 40 million SAR, and the firm commits before it knows the grade. The rest of the chapter asks what it would be worth to learn the grade *first*.

Insight. The two node rules encode who is in control. At a square, the decision maker is in control and rational, so they take the maximum. At a circle, nature is in control and indifferent to the decision maker's wishes, so the best the decision maker can do is average over what nature might do. Every tree in this chapter is just these two operations, applied from the tips back to the root. Get the node shapes right and the arithmetic is mechanical.

13.2 *Rollback: A Two-Stage Pilot Decision*

The interesting trees are the ones where an early, cheap decision buys information for a later, expensive one. The canonical case is a pilot plant: a small trial that reveals the ore grade before the full plant is committed.

Example 13.2 — Pilot, then plant: a two-stage tree

Extend Example 13.1. Before committing to the full plant, *Najd Copper* can build a pilot plant for 5 million SAR that reveals the ore grade with certainty. After the pilot reports, the firm decides whether to build the full plant (present worth $+80$ if high grade, -40 if low) or abandon (0). The prior probabilities are unchanged: high grade 0.55, low grade 0.45. The firm's three root alternatives are: build now, pilot first, or decline. Roll the tree back.

Solution.

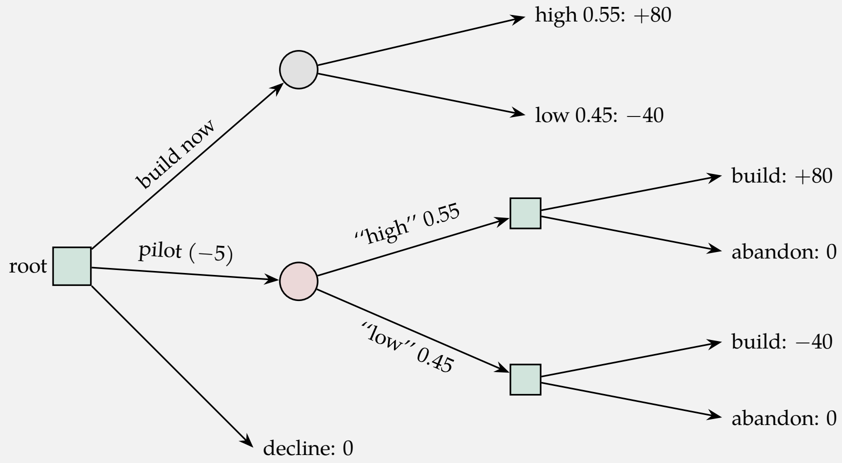

Step 1: roll back the second-stage decision nodes.

After a pilot report, the firm chooses build or abandon, knowing the grade:

$$\text{report ``high'':}\ \max(80,\ 0) = 80, \qquad \text{report ``low'':}\ \max(-40,\ 0) = 0.$$

The firm builds on a high report and abandons on a low one: the pilot lets the firm decline the loser.

Step 2: roll back the pilot's chance node.

The pilot reveals the grade, which is high with probability 0.55 and low with 0.45, so the expected worth *after* the pilot report (before subtracting the pilot cost) is

$$0.55\,(80) + 0.45\,(0) = 44 \text{ million SAR.}$$

Subtract the 5 million SAR pilot cost:

$$\mathrm{EMV}(\text{pilot}) = 44 - 5 = 39 \text{ million SAR.}$$

Step 3: roll back the root decision.

Compare the three alternatives:

$$\text{build now } 26, \quad \text{pilot } 39, \quad \text{decline } 0 \ \Rightarrow\ \text{root} = \max = \boxed{39 \text{ million SAR}}.$$

Run the pilot. It is worth 39 million SAR against 26 for building blind, a gain of 13 million SAR, because it converts the 40 million SAR loss on the low-grade branch into a costless walk-away. The pilot's whole value is that it lets the firm avoid building the loser.

Insight. Read the rollback as a conversation between the two node types. The second-stage *squares* say: once I know the grade, I will build only the winner. The pilot's *circle* says: I do not know which report I will get, so average over them. The root *square* says: given all that, is paying for the pilot better than building blind or walking away? Each node does one job; the tree chains them. The pilot's 13-million-SAR gain is not magic — it is exactly the low-grade loss the firm no longer has to take, weighted by how often it would have occurred.

In practice. Whether a pilot "reveals the grade with certainty" is an engineering claim, and usually an optimistic one. A pilot plant samples one part of a deposit under trial conditions; it estimates the grade, it does not photograph it. Treating the pilot as perfect, as this example does, gives an *upper* bound on its value. The imperfect-information section replaces that idealization with a stated reliability, which is where the real assay physics enters the economics.

Reference. Park, 3rd ed., Chapter 11; the EVPI is the standard decision-analysis bound on the worth of any test.

Expected value of perfect information (EVPI). The gain from learning the true state before acting: EVPI $=$ $\mathbb{E}[\text{best act under perfect information}] - \mathbb{E}[\text{best act now}]$. Always ≥ 0; a ceiling on the value of any test.

13.3 *The Value of Perfect Information*

The pilot of Example 13.2 was a *perfect* test: it revealed the grade with certainty. Its net value was 13 million SAR, but that already nets out the 5 million SAR cost. Strip the cost away and ask a cleaner question: what is the raw value of *knowing the grade* before deciding, whatever a test to learn it might cost? That number is the expected value of perfect information, and it is the ceiling on what any study, pilot, or survey could ever be worth.

The recipe is two expectations and a subtraction. First, the value of acting *now*, under uncertainty: that is the best root alternative from Example 13.1, the EMV of building blind, 26 million SAR. Second, the value of acting *with perfect information*: for each state, take the best action *given* that state, then weight by the state's prior probability.

Example 13.3 — EVPI for the grade decision

Find the EVPI of learning the ore grade before deciding, using the data of Example 13.1 (high grade 0.55, present worth $+80$; low grade 0.45, present worth -40).

Solution.

Step 1: the best act now.

Without information, the firm builds, worth

$$\mathbb{E}[\text{best act now}] = \max(26,\ 0) = 26 \text{ million SAR}.$$

Step 2: the best act under perfect information.

If a clairvoyant announced the grade first, the firm would build on high grade (worth 80) and decline on low grade (worth 0), each with its prior probability:

$$\mathbb{E}[\text{perfect info}] = 0.55\,(80) + 0.45\,(0) = 44 \text{ million SAR}.$$

Step 3: subtract.

$$\text{EVPI} = 44 - 26 = \boxed{18 \text{ million SAR}}.$$

Perfect knowledge of the grade is worth 18 million SAR. Notice the check against Example 13.2: the perfect pilot cost 5 million SAR and returned $44 - 5 = 39$, a net gain over building-blind of $39 - 26 = 13$ million SAR — which is exactly $\text{EVPI} - \text{cost} = 18 - 5$. A perfect test is worth the EVPI, gross; its net value is the EVPI minus its price. Any test priced above 18 million SAR is not worth buying, no matter how good.

Insight. The EVPI answers "how much uncertainty is even at stake here?" before any specific test is on the table. It is set entirely by the payoffs and the priors, not by any survey's design. That makes it the first number to compute: if the EVPI is small, no study is worth commissioning, and the firm should just decide. If the EVPI is large, information is valuable and the next question — what a *real*, imperfect test is worth — becomes worth asking. The EVPI screens the whole question of whether to gather data.

Reference. Park, 3rd ed., Chapter 11; Bayes' rule links a test's reliability to the revision of the state probabilities.

13.4 Imperfect Information and Bayes' Rule

Reliability (likelihoods). The conditional probabilities of each test result given each true state, e.g. $P(\text{favorable} \mid \text{high}) = 0.80$. The test's sensitivity and specificity.

Posterior probability. The revised probability of a state *after* seeing the test result, $P(\text{state} \mid \text{result})$, computed from the prior and the reliability by Bayes' rule.

Real tests are fallible. An assay says ''favorable'' or ''unfavorable,'' and it is right most of the time but not always. To value such a test we need two things: its *reliability*, stated as conditional probabilities, and Bayes' rule to turn the reliability and the prior into revised (*posterior*) probabilities of the grade given each test result.

The mechanics are a clean three-column march: prior, then joint, then posterior. Multiply each prior by the matching reliability to get the *joint* probability of (state and result); sum the joints down each result column to get the *marginal* probability of that result; divide each joint by its column's marginal to get the *posterior*. The following example carries the arithmetic in one table.

Example 13.4 — Revising the grade with an assay survey

Najd Copper can commission an assay survey that reports ''favorable'' (F) or ''unfavorable'' (U). The survey is reliable but not perfect: $P(F \mid \text{high}) = 0.80$ and $P(U \mid \text{low}) = 0.75$ (so $P(U \mid \text{high}) = 0.20$ and $P(F \mid \text{low}) = 0.25$). The priors are high grade 0.55, low grade 0.45. Build the prior–joint–posterior table.

Solution.

State	Prior	$P(F \mid \cdot)$	Joint $(\cdot, F)$	Joint $(\cdot, U)$
High	0.55	0.80	$0.55(0.80) = 0.4400$	$0.55(0.20) = 0.1100$
Low	0.45	0.25	$0.45(0.25) = 0.1125$	$0.45(0.75) = 0.3375$
Marginal			$P(F) = 0.5525$	$P(U) = 0.4475$

Table 13.1. The joint table. Each joint is prior × reliability; the column sums are the marginal probabilities of each survey result. The four joints sum to 1.

The posteriors divide each joint by its column marginal:

	After ''favorable'' (F)	After ''unfavorable'' (U)
$P(\text{high} \mid \cdot)$	$0.4400/0.5525 = 0.7964$	$0.1100/0.4475 = 0.2458$
$P(\text{low} \mid \cdot)$	$0.1125/0.5525 = 0.2036$	$0.3375/0.4475 = 0.7542$

Table 13.2. The posterior probabilities. A ''favorable'' report lifts the high-grade probability from the prior 0.55 to 0.80; an ''unfavorable'' report drops it to 0.25. Each column sums to 1.

Read what the survey does. Before it, the high-grade probability is 0.55. A favorable report raises it to 0.7964; an unfavorable report cuts it to 0.2458. The survey does not settle the question — it is not perfect — but it moves the odds, and moved odds change decisions.

In practice. The two reliability numbers, $P(F \mid \text{high}) = 0.80$ and $P(U \mid \text{low}) = 0.75$, are not free parameters — they come from the assay method's physics. The sensitivity $P(F \mid \text{high})$ is how often the assay flags truly rich ore; the specificity $P(U \mid \text{low})$ is how often it clears truly poor ore. Both depend on sample size, the mineralogy, the detection threshold, and the laboratory protocol. A geologist sets these by choosing the survey design; the economist only turns them, through Bayes' rule, into a value. Better physics means likelihoods closer to 1, posteriors closer to certainty, and information closer to the EVPI ceiling.

Expected value of sample information (EVSI). The gain from a fallible test: EVSI $= \mathbb{E}[\text{best act with the test}] - \mathbb{E}[\text{best act now}]$. Lies between 0 and the EVPI.

13.5 *The Value of a Fallible Test: EVSI*

The posteriors of Example 13.4 let us value the survey. The logic is the rollback of section 13.2, but the second-stage decision now uses the *posterior* probabilities, not the prior, because after the report the firm knows more.

Example 13.5— EVSI of the assay survey, and whether to buy it

Using the posteriors of Example 13.4, find the expected value of the survey (EVSI), compare it with the EVPI of 18 million SAR, and decide whether to commission the survey if it costs (a) 3 million SAR or (b) 6 million SAR.

Solution.

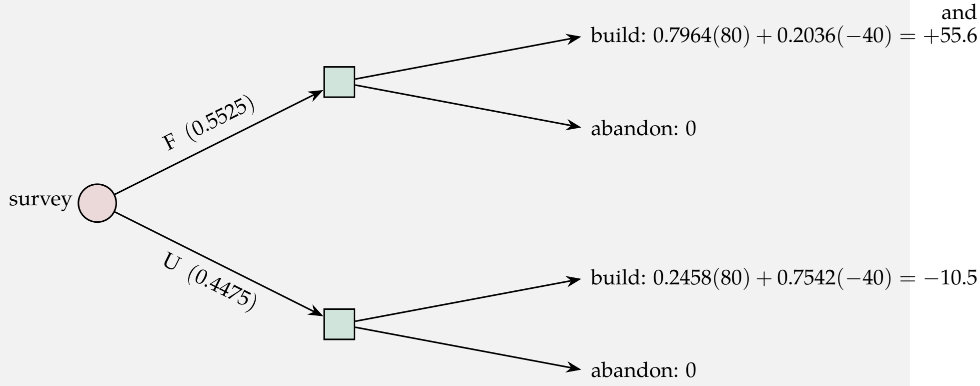

Step 1: the second-stage decision under each report.

Using the posteriors, the expected present worth of building is

$$\text{after F: } 0.7964\,(80) + 0.2036\,(-40) = 63.71 - 8.14 = +55.57,$$

$$\text{after U: } 0.2458\,(80) + 0.7542\,(-40) = 19.66 - 30.17 = -10.50.$$

So the firm **builds after a favorable report** (55.57 beats 0) and **abandons after an unfavorable report** (0 beats -10.50).

Step 2: average over the reports.

The firm does not know in advance which report it will get, so weight the second-stage values by the marginal report probabilities:

$$\mathbb{E}[\text{with survey}] = 0.5525\,(55.57) + 0.4475\,(0) = 30.7 \text{ million SAR.}$$

(Equivalently, straight from the joint table: $0.44\,(80) + 0.1125\,(-40) = 35.2 - 4.5 = 30.7$, building only on the F branch.)

Step 3: the EVSI.

Subtract the best act now, 26 million SAR:

$$\text{EVSI} = 30.7 - 26 = \boxed{4.7 \text{ million SAR}}.$$

The survey is worth 4.7 million SAR, gross of its cost. Compare with the EVPI:

$$\frac{\text{EVSI}}{\text{EVPI}} = \frac{4.7}{18} = 0.26,$$

so this fallible survey captures about 26% of what perfect information would be worth — sensible, since it moves the odds but does not settle them.

Step 4: the pay-for-the-survey verdict.

Buy the survey only if its value exceeds its price:

$$\text{(a) cost } 3 < 4.7 \;\Rightarrow\; \text{commission it (net } +1.7),$$

$$\text{(b) cost } 6 > 4.7 \;\Rightarrow\; \text{skip it (net } -1.3).$$

The break-even price is the EVSI itself, 4.7 million SAR: a survey priced below that is worth buying, one priced above it is not.

Figure 13.1 places the survey among the tests the firm might have bought instead. Read it for two things: the flat stretch on the left, where a test is too unreliable to change any decision and is therefore worth nothing, and the distance from the survey's 4.7 up to the 18-million-SAR ceiling, which is what a sharper assay could still buy.

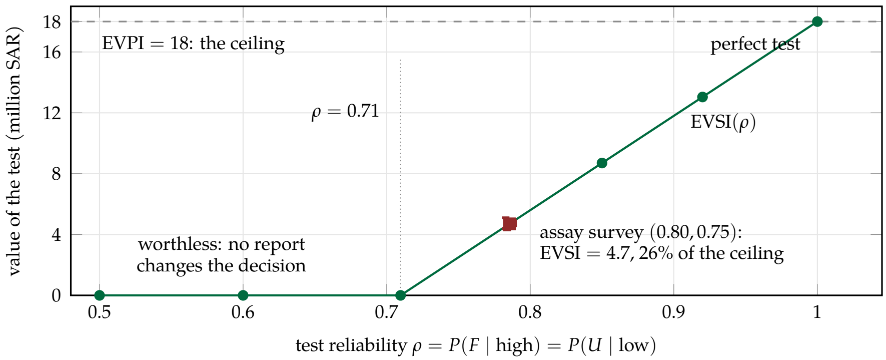

Figure 13.1. The value of information rises with the quality of the test and stops at the EVPI: the curve is this chapter's own EVSI arithmetic re-run over a family of tests of symmetric reliability ρ. A test too weak to change any decision is worth nothing, the assay survey captures 4.7 of the 18 million SAR ceiling, and only a clairvoyant collects all of it.

Insight. The EVSI sits inside a firm ordering: $0 \leq \text{EVSI} \leq \text{EVPI}$. The lower bound says a test can never hurt — worst case, you ignore it. The upper bound says a test can never beat clairvoyance. A perfect test hits the EVPI exactly; a useless test (one whose result is independent of the state) has EVSI zero. Every real survey lands between, and its position is set by its reliability: the 0.80 and 0.75 likelihoods put this one 26% of the way up. Improving the assay raises the likelihoods, the posteriors sharpen, and the EVSI climbs toward the 18-million-SAR ceiling.

13.6 *Case Study from the Literature*

Value-of-information analysis is used to decide how much a country, a utility, or a developer should spend to resolve a specific uncertainty before committing capital. Patankar, de Queiroz, DeCarolis, Bazilian, and Chattopadhyay (2019) apply the idea to electricity planning in South Sudan, where the uncertainty is not ore grade but future armed *conflict* that could damage newly built generators. They ask what planners should be willing to pay to eliminate that uncertainty — the expected value of perfect information — and how it depends on the modelling assumptions. Their amounts are stated as percentages of a baseline system cost and in US dollars per kilowatt-hour, the units of the paper, and we keep them.

Reference. Patankar et al. (2019), arXiv:1911.04652. Open access; full reference at the end of the chapter.

Example 13.6 — Case study: the value of resolving conflict uncertainty

In Patankar et al.'s stochastic planning model, the near-term decision is which generators to build, and the chance node is whether future conflict damages them (their damage-probability parameter ranges over $P(\text{no damage}) = 0.10$ to 0.50). They compute the EVPI as the difference between the cost of a plan made with foreknowledge of the conflict outcome (the ''wait-and-see'' solution) and the cost of the best plan made under uncertainty (the stochastic solution). Interpret their headline numbers with the tree vocabulary of this chapter.

Solution.

Step 1: match the model to the tree.

The build choice is a decision node; conflict-or-not is a chance node; the EVPI is $\mathbb{E}[\text{best plan under perfect information}] - \mathbb{E}[\text{best plan now}]$, exactly as in Example 13.3, but with cost minimized rather than worth maximized, so the roles of ''high'' and ''low'' payoff swap sign.

Step 2: read the value of information.

The authors report that the EVPI *peaks* at about **8% of the baseline system cost**, under moderate damage probability and a curtailment cost near 0.20 USD per kWh. In tree terms: when conflict is a near-certainty or a remote chance, the best plan is obvious and information is nearly worthless; it is at *intermediate* odds, where the decision is in real doubt, that resolving the uncertainty is worth the most. This is the same shape as our grade example, where the EVPI is largest when the prior does not already point clearly to build or decline.

Step 3: read the companion metric.

They also report the value of the stochastic solution — the gain from planning *with* the uncertainty rather than ignoring it — peaking near **4% of baseline cost**, again at intermediate odds. And ignoring conflict entirely (a naive least-cost plan) can cost up to about **60% of baseline** in expensive recourse when a hydro-heavy build is later damaged. The lesson mirrors this chapter: the value of information is real, it is largest when the decision is truly uncertain, and pretending the chance node is not there is the most expensive mistake of all.

What did the researchers conclude? That explicitly modelling conflict as a chance node changes near-term investment, favouring solar photovoltaics as a hedge, and that the value of resolving the uncertainty is highest at moderate conflict odds (Patankar et al., 2019). The finding that information is worth most when the decision is in real doubt is exactly the EVPI intuition of Example 13.3, measured in a real infrastructure model.

Insight. The case study verifies the method, not the specific numbers we could recompute. Given the model's stated damage probabilities and cost structure, the EVPI follows from the same subtraction this chapter uses: cost with foreknowledge minus cost under uncertainty. What the tree cannot check from outside are the inputs — the damage probabilities, the curtailment cost, the capital costs of each generator. The value-of-information arithmetic is exact; its inputs are engineering and political judgment. That division is the theme of every case study in this book.

13.7 *Chapter Summary*

- A **decision tree** draws a sequential decision left to right, with two node types: **decision nodes** (squares, where the decision maker chooses) and **chance nodes** (circles, where nature chooses, with probabilities).
- **Rollback** evaluates the tree from tips to root: take the **expected value** at each chance node and the **maximum** at each decision node. Taking a maximum at a chance node is the classic error — it values the project as if the decision maker controlled nature.
- The **EMV** at a chance node is $\sum_k p_k V_k$: the expected present worth of chapter 12, sitting at a node.
- The **EVPI** is $\mathbb{E}[\text{best act under perfect information}] - \mathbb{E}[\text{best act now}]$. It is never negative and is a **ceiling** on the value of any test. Compute it first: a small EVPI means no study is worth commissioning.
- A **perfect test** is worth the EVPI, gross; its net value is the EVPI minus its price.
- For an **imperfect test**, state its **reliability** (likelihoods $P(\text{result} \mid \text{state})$) and use **Bayes' rule**: prior $\times$ likelihood $=$ joint; sum joints for the marginal; joint / marginal $=$ **posterior**. The posterior drives the second-stage decision.

- The **EVSI** is $\mathbb{E}$[best act with the test] $-$ $\mathbb{E}$[best act now], and it satisfies $0 \leq \text{EVSI} \leq \text{EVPI}$. Buy the test only if its price is below the EVSI; the break-even price *is* the EVSI.
- The value of information is largest when the decision is **in doubt**. At extreme priors the best act is obvious and information is nearly worthless.
- These trees connect back to the expected-present-worth analysis of chapter 12 and forward to chapter 15, where the same branching is priced with a market discipline.

13.8 Exercises

These problems follow the chapter's build: rollback on one- and two-stage trees, then the EVPI ceiling, then Bayes-revised posteriors and the EVSI of a fallible test.

Rollback

Exercise 13.1. A firm can launch a product (present worth $+300$ if the market is strong, -200 if weak) or not launch (0). The market is strong with probability 0.6. Draw the one-stage tree, compute the EMV of launching, and state the decision.

Exercise 13.2. Reconsider the launch of Exercise 13.1. Before launching, the firm can run a market test for 20 that reveals the market state with certainty. Roll back the two-stage tree (test, then launch or not) and decide whether to test.

EVPI

Exercise 13.3. For the launch of Exercise 13.1, compute the EVPI of learning the market state before deciding, and use it to explain why the perfect test of Exercise 13.2 was worth running.

Exercise 13.4. A project's best act now is worth 40. A clairvoyant would let the firm earn an expected 40 as well, because the same action is best in every state. What is the EVPI, and what does it say about commissioning any study?

Imperfect information and Bayes' rule

Exercise 13.5. A deposit is high grade with prior probability 0.5. An assay reports "favorable" or "unfavorable" with reliability $P(F \mid \text{high}) = 0.9$ and $P(U \mid \text{low}) = 0.6$. Build the prior–joint–posterior table and find $P(\text{high} \mid F)$ and $P(\text{high} \mid U)$.

Exercise 13.6. Using the revision of Exercise 13.5, suppose building is worth $+100$ if high grade and -60 if low grade, and abandoning is 0. Find the EVSI of the assay, and state the most the firm should pay for it.

Exam-style problems

The problems below follow the format, length, and point weights of the final exam. Draw each tree, mark squares and circles, roll back carefully, and close with the decision sentence.

Exercise 13.7. *Wadi Lithium* can build a brine-processing plant now (present worth $+120$ million SAR if the brine is rich, -70 million SAR if lean; the brine is rich with prior probability 0.5), or decline (0). It can instead drill a test well for 8 million SAR that reveals the brine quality with certainty before the build decision. (a) (*6 pts*) Compute the EMV of building now and state the act-now decision. (b) (*8 pts*) Roll back the two-stage tree (drill, then build or decline) and state whether to drill. (c) (*6 pts*) Compute the EVPI and use it to check the drilling decision. (d) (*4 pts*) At what test-well price would the firm be indifferent between drilling and building blind?

Exercise 13.8. *Najd Copper* faces the build decision of Example 13.1 (high grade 0.55, $+80$; low grade 0.45, -40; build now EMV 26; EVPI 18 million SAR). A quick geophysical scan reports "F" or "U" with reliability $P(F \mid \text{high}) = 0.70$ and $P(U \mid \text{low}) = 0.80$. (a) (*10 pts*) Build the prior–joint–posterior table and give $P(\text{high} \mid F)$ and $P(\text{high} \mid U)$. (b) (*8 pts*) Roll back the survey tree and compute the EVSI. (c) (*4 pts*) The scan costs 3 million SAR. Should the firm buy it? What fraction of the EVPI does it capture?

Assessing outside recommendations

Each problem below quotes a finished analysis produced by someone else. The work looks professional and its arithmetic is internally consistent, but each contains a genuine method error. Decide whether to accept or reject, name the flaw precisely, and produce the corrected number and decision.

Exercise 13.9. An analyst evaluates *Najd Copper*'s pilot decision of Example 13.2 (build now EMV 26; perfect pilot costs 5; high 0.55, $+80$; low 0.45, -40). Their memo reads:

> "After the pilot, the firm gets the better of the two grade outcomes at the chance node, so the pilot branch is worth $\max\big(0.55(80),\ 0.45(-40)\big) = \max(44,\ -18) = 44$, minus the 5 million SAR cost, giving 39. But the same maximum logic applied to building now gives $\max(0.55 \cdot 80,\ 0.45 \cdot (-40)) = 44$ as well, so building now is worth 44 and dominates the pilot's 39. Recommendation: build now, do not pilot."

Accept or reject the memo? Name the flaw and give the correct root decision.

Exercise 13.10. A consultant computes the EVPI for the brine decision of Exercise 13.7 (rich 0.5, $+120$; lean 0.5, -70). Their memo reads:

> "With perfect information the firm earns $0.5(120) + 0.5(-70) = 25$. So the value of perfect information is $25 - 25 = 0$, and no test well is ever worth drilling."

Accept or reject the memo? Name the flaw and give the correct EVPI.

References

Park, C. S. (2012). *Fundamentals of Engineering Economics* (3rd ed.). Pearson.

Patankar, N., de Queiroz, A. R., DeCarolis, J. F., Bazilian, M. D., & Chattopadhyay, D. (2019). Building conflict uncertainty into electricity planning: A South Sudan case study. *arXiv preprint*, arXiv:1911.04652. `https://arxiv.org/abs/1911.04652`