

Meeting 11:

The value of information

Two-stage decision problem

Suppose we have the following (general) decision matrix for decision problem with r actions and s states:

<i>Actions</i>	<i>States</i>			
	θ_1	θ_2	\dots	θ_s
a_1	U_{11}	U_{12}	\dots	U_{1n}
a_2	U_{21}	U_{22}	\dots	U_{2n}
\vdots	\vdots	\vdots	\ddots	\vdots
a_r	U_{m1}	U_{m2}	\dots	U_{mn}

Using the *EU*-criterion, the optimal action a^* is the action that maximises the expected utility, i.e.

$$a^* = \arg \max_i \left\{ \sum_{j=1}^s U_{ij} \cdot P(\theta_j) \right\}$$

The probabilities $P(s_j)$ are general, but we know that they are prior probabilities when assigned on background information only and posterior probabilities when new data is taken into account.

Hence, we could rather consider...

Optimal action in *prior sense*:
$$a' = \arg \max_i \left\{ \sum_{j=1}^s U_{ij} \cdot f'(\theta_j) \right\}$$

Optimal action in *posterior sense*:
$$a'' = a'' | \mathbf{y}_n = \arg \max_i \left\{ \sum_{j=1}^s U_{ij} \cdot f''(\theta_j | \mathbf{y}_n) \right\}$$

where $\mathbf{y}_n = (y_1, \dots, y_n)$ are the values in a sample from the population characterised by the state θ .

When the set of states, Θ is continuous the corresponding formulas are

$$a' = \arg \max_i \left\{ \int_{\Theta} U_i(\theta) f'(\theta) d\theta \right\} \quad a'' | \mathbf{y}_n = \arg \max_i \left\{ \int_{\Theta} U_i(\theta) f''(\theta | \mathbf{y}_n) d\theta \right\}$$

where $U_i(\theta) = U(a_i, \theta)$ is the utility function associated with action a_i .

We can therefore say that when the decision maker has the possibility to acquire new data to inform their decision, they are faced to

- i. Choose an action based on the current information.
- ii. Await the choice of action until new data has possibly been acquired.

A better formulation is however

1. Decide whether new data should be acquired.
2. Choose an action based on the information available after 1.

which is referred to as a *two-stage* decision problem.

The decision in stage 1 will be based on the *value* of the information that *may* be obtained from acquiring new data and the cost of obtaining these data.

The decision in stage 2 will be based on the expected utilities of the different actions under the distribution of the states assigned upon having conducted stage 1.

Focusing on Stage 1

Once a probability distribution of the states have been (terminally) assigned, choosing an action according to the *EU*-criterion is as previously taken up in the course.

The value of (new) information

Example

Suppose we choose between selling ice cream on the beach “tomorrow” (a Saturday in July).

If the proportion of beach “guests” buying ice cream is at least 30% you deem your minimum payoff to be 50. If lower than 30% your deemed minimum payoff is –200.

Using your deemed minimum payoff as utilities, your decision matrix is

<i>Actions</i>	<i>Proportion of beach guests buying</i>	
	$\geq 30\%$	$< 30\%$
Sell	50	–200
Do not sell	0	0

<i>Actions</i>	<i>Proportion of beach guests buying</i>	
	$\geq 30\%$	$< 30\%$
Sell	50	−200
Do not sell	0	0

Being ignorant about the proportion, q , of beach guests buying you would probably assign a uniform distribution to this proportion.

This means that $P(q \geq 30\%) = 0.7$

$$\Rightarrow E(U(\text{"Sell"})) = 50 \cdot 0.7 - 200 \cdot 0.3 = -25 \text{ and } E(U(\text{"Do not sell"})) = 0$$

$$\Rightarrow a' = \text{"Do not sell"}$$

Now, what would it be worth to you to learn that among 10 beach guests today, 6 bought ice cream?

This information (6 out of 10) can be used to obtain a posterior distribution of q .

The likelihood of q in light of this information is

$$f(6|q, (10)) = \binom{10}{6} q^6 (1 - q)^4$$

The prior distribution of q is a Beta(1,1)-distribution and updating with the likelihood we obtain a Beta(1+6,1+4) = Beta(7,5) posterior distribution.

$$\text{Hence, } P(q \geq 30\% | 6 \text{ out of } 10) = \int_{0.3}^1 \frac{q^6 (1-q)^4}{B(7,5)} dq \approx 0.978$$

$$\Rightarrow E(U(\text{"Sell"} | 6 \text{ out of } 10) | 6 \text{ out of } 10) = 50 \cdot 0.978 - 200 \cdot 0.022 = 44.5$$

and $E(U(\text{"Do not sell"} | 6 \text{ out of } 10)) = 0$

$$\Rightarrow a'' | 6 \text{ out of } 10 = \text{"Sell"}$$

The difference in posterior expected utility between $a'' | 6 \text{ out of } 10$ and a' is $44.5 - 0 = 44.5$, which is then the value (in utility) of the information (VOI) .

Now, assume instead that the information you obtained was that of the 10 beach guests, 4 bought ice cream.

That would give a Beta(5,7) posterior distribution of q .

$$\text{Hence, } P(q \geq 30\% | 6 \text{ out of } 10) = \int_{0.3}^1 \frac{q^6(1-q)^4}{B(7,5)} dq \approx 0.790$$

$$\Rightarrow E(U(\text{"Sell"} | 4 \text{ out of } 10) | 4 \text{ out of } 10) = 50 \cdot 0.79 - 200 \cdot 0.21 = -2.5$$

and $E(U(\text{"Do not sell"} | 4 \text{ out of } 10)) = 0$

$$\Rightarrow a'' | 4 \text{ out of } 10 = \text{"Do not sell"}$$

The difference in posterior expected utility between $a'' | 4 \text{ out of } 10$ and a' is here $0 - 0 = 0$, so the value (in utility) of the information is 0.

In general...

The value of new information, I_N is defined as

$$\text{VOI}(I_N) = E(U(a''|I_N)|I_N) - E(U(a')|I_N)$$

Alternatively, it can be defined in terms of expected loss:

$$\text{VOI}(I_N) = E(L(a')|I_N) - E(L(a''|I_N)|I_N)$$

Different kinds of new information

Perfect information: Information that will remove all uncertainty – the true state is known
(PI)

Sample information: Information that to some extent will reduce the uncertainty about which state is true –also referred to as partial information
(SI)

In any case, at the outset we don't have this information. The question is: Should we endeavour to obtain it? – depends on the value of the information and the cost of acquiring it.

But since we don't know on forehand what the information outcome (data) will be, we don't know its value on forehand.

We must however know which the different outcomes are and with which probability distribution.

⇒ We should be able to calculate the *expected value of information* (EVOI)

- EVSI, EVPI

Expected value of sample information (EVSI)

Assume there is a possibility to obtain a sample of size n from a population for which the true state of nature/world is θ .

Let $\mathbf{y}_n = (y_1, \dots, y_n)$ be the values in such a potential sample with probability density/mass function $f(y_k|\theta)$.

The value of sample information, VSI, is then defined as

$$\text{VSI}(\mathbf{y}_n) = E(U(a''|\mathbf{y}_n)|\mathbf{y}_n) - E(U(a')|\mathbf{y}_n)$$

or

$$\text{VSI}(\mathbf{y}_n) = E(L(a')|\mathbf{y}_n) - E(L(a''|\mathbf{y}_n)|\mathbf{y}_n)$$

The posterior expected utility from taking the posterior optimal action minus the posterior expected utility from taking the prior optimal action

The posterior expected loss from taking the prior optimal action minus the posterior expected loss from taking the posterior optimal action

When utility is linear in money, we could also define VSI as

$$\text{VSI}(\mathbf{y}_n) = E(R(a''|\mathbf{y}_n)|\mathbf{y}_n) - E(R(a')|\mathbf{y}_n)$$

Using the definition of $VSI(y)$ the expected value of sample information (EVSI) is

$$EVSI = \int VSI(\mathbf{y}_n) \cdot f(\mathbf{y}_n) d\mathbf{y}_n \quad y_k \text{ continuous-valued}$$

$$EVSI = \sum VSI(\mathbf{y}_n) \cdot f(\mathbf{y}_n) \quad y_k \text{ discrete-valued}$$

where the integration (or summing) is over the sample space of the y_k s and $f(\mathbf{y}_n)$ is the prior-predictive distribution of \mathbf{y}_n , i.e.

$$f(\mathbf{y}_n) = \int f(\mathbf{y}_n|\theta) f'(\theta) d\theta$$

Most often, however, the sample is a random sample from the same population and thus

$$f(\mathbf{y}_n) = \prod_{k=1}^n f(y_k) = \prod_{k=1}^n \int f(y_k|\theta) f'(\theta) d\theta$$

Exercise 6.15

15. A store must decide whether or not to stock a new item. The decision depends on the reaction of consumers to the item, and the payoff table (in dollars) is as follows.

		PROPORTION OF CONSUMERS PURCHASING				
		0.10	0.20	0.30	0.40	0.50
DECISION	Stock 100	-10	-2	12	22	40
	Stock 50	-4	6	12	16	16
	Do not stock	0	0	0	0	0

If $P(0.10) = 0.2$, $P(0.20) = 0.3$, $P(0.30) = 0.3$, $P(0.40) = 0.1$, and $P(0.50) = 0.1$, what decision maximizes expected payoff? ~~If perfect information is available, find VPI for each of the five possible states of the world and compute EVPI.~~

DECISION	PROPORTION OF CUSTOMERS BUYING (θ)				
	0.10	0.20	0.30	0.40	0.50
Stock 100	-10	-2	12	22	40
Stock 50	-4	6	12	16	16
Do not stock	0	0	0	0	0

θ	$P(\theta)$
0.10	0.2
0.20	0.3
0.30	0.3
0.40	0.1
0.50	0.1

$$a' = \arg \max_a \{E'R(a)\}$$

$$E(R(\text{Stock 100})) = (-10) \cdot 0.2 + (-2) \cdot 0.3 + 12 \cdot 0.3 + 22 \cdot 0.1 + 40 \cdot 0.1 = 7.2$$

$$E(R(\text{Stock 50})) = (-4) \cdot 0.2 + 6 \cdot 0.3 + 12 \cdot 0.3 + 16 \cdot 0.1 + 16 \cdot 0.1 = 7.8$$

$$E(R(\text{Do not stock})) = 0 \cdot 0.2 + 0 \cdot 0.3 + 0 \cdot 0.3 + 0 \cdot 0.1 + 0 \cdot 0.1 = 0$$

$$\Rightarrow a' = \text{Stock 50}$$

Exercise 6.16

16. In Exercise 15, suppose that sample information is available in the form of a random sample of consumers. For a sample of size *one*,
- (a) find the posterior distribution if the one person sampled will purchase the item, and find the value of this sample information;
 - (b) find the posterior distribution if the one person sampled will *not* purchase the item, and find the value of this sample information;
 - (c) find the expected value of sample information.

DECISION	PROPORTION OF CUSTOMERS BUYING (θ)				
	0.10	0.20	0.30	0.40	0.50
Stock 100	-10	-2	12	22	40
Stock 50	-4	6	12	16	16
Do not stock	0	0	0	0	0

θ	$P(\theta)$
0.10	0.2
0.20	0.3
0.30	0.3
0.40	0.1
0.50	0.1

Below we have used the word “BUY” instead of “PURCHASE”

(a) Posterior distribution:

$$P(\theta|\text{BUY}) = \frac{P(\text{BUY}|\theta) \cdot P(\theta)}{P(\text{BUY})} = \frac{P(\text{BUY}|\theta) \cdot P(\theta)}{\sum_{\lambda} P(\text{BUY}|\lambda) \cdot P(\lambda)}$$

$$P(\text{BUY}|\theta) = \theta \quad \text{since } \theta \text{ is the proportion of buying customers}$$

\Rightarrow

$$P(\text{BUY}) = 0.10 \cdot 0.2 + 0.20 \cdot 0.3 + 0.30 \cdot 0.3 + 0.40 \cdot 0.1 + 0.50 \cdot 0.1 = 0.26$$

$$P(0.10|\text{BUY}) = 0.10 \cdot 0.2 / 0.26 \approx 0.0769$$

$$P(0.20|\text{BUY}) = 0.20 \cdot 0.3 / 0.26 \approx 0.2308$$

$$P(0.30|\text{BUY}) = 0.30 \cdot 0.3 / 0.26 \approx 0.3462$$

$$P(0.40|\text{BUY}) = 0.40 \cdot 0.1 / 0.26 \approx 0.1538$$

$$P(0.50|\text{BUY}) = 0.50 \cdot 0.1 / 0.26 \approx 0.1923$$

$$\text{VSI}(\text{BUY}) = E(R(a''|\text{BUY})|\text{BUY}) - E(R(a')|\text{BUY})$$

$$a' = \langle \text{from Exercise 6.15} \rangle = \text{Stock 50}$$

$$a''|\text{BUY} = \arg \max_a \{E(R(a|\text{BUY})|\text{BUY})\}$$

$$E(R(a|\text{BUY})|\text{BUY}) = \sum_a R(a, \theta) \cdot P(\theta|\text{BUY})$$

\Rightarrow

$$\begin{aligned} E(R(\text{Stock 100}|\text{BUY})|\text{BUY}) &= (-10) \cdot 0.0769 \dots + (-2) \cdot 0.2308 \dots + 12 \cdot 0.3462 \dots \\ &\quad + 22 \cdot 0.1538 \dots + 40 \cdot 0.1923 \dots = 14.00 \end{aligned}$$

$$\begin{aligned} E(R(\text{Stock 50}|\text{BUY})|\text{BUY}) &= (-4) \cdot 0.0769 \dots + 6 \cdot 0.2308 \dots + 12 \cdot 0.3462 \dots + \\ &\quad 26 \cdot 0.1538 \dots + 16 \cdot 0.1923 \dots \approx 10.77 \end{aligned}$$

$$\begin{aligned} E(R(\text{Do not stock}|\text{BUY})|\text{BUY}) &= 0 \cdot 0.0769 \dots + 0 \cdot 0.2308 \dots + 0 \cdot 0.3462 \dots + \\ &\quad 0 \cdot 0.1538 \dots + 0 \cdot 0.1923 \dots = 0 \end{aligned}$$

$$\Rightarrow a''|\text{BUY} = \text{Stock 100}$$

$$\begin{aligned} \Rightarrow \text{VSI}(\text{BUY}) &= E(R(\text{Stock 100}|\text{BUY})|\text{BUY}) - E(R(\text{Stock 50})|\text{BUY}) \\ &= 14.00 - 10.77 = 3.23 \end{aligned}$$

(b) Posterior distribution:

$$P(\theta|\text{NOT BUY}) = \frac{P(\text{NOT BUY}|\theta) \cdot P(\theta)}{P(\text{NOT BUY})} = \frac{P(\text{NOT BUY}|\theta) \cdot P(\theta)}{\sum_{\lambda} P(\text{NOT BUY}|\lambda) \cdot P(\lambda)}$$

$$P(\text{NOT BUY}|\theta) = 1 - \theta$$

\Rightarrow

$$P(\text{NOT BUY}) = 0.90 \cdot 0.2 + 0.80 \cdot 0.3 + 0.70 \cdot 0.3 + 0.60 \cdot 0.1 + 0.50 \cdot 0.1 = 0.74$$

$$P(0.10|\text{NOT BUY}) = 0.90 \cdot 0.2 / 0.74 \approx 0.2432$$

$$P(0.20|\text{NOT BUY}) = 0.80 \cdot 0.3 / 0.74 \approx 0.3243$$

$$P(0.30|\text{NOT BUY}) = 0.70 \cdot 0.3 / 0.74 \approx 0.2838$$

$$P(0.40|\text{NOT BUY}) = 0.60 \cdot 0.1 / 0.74 \approx 0.0811$$

$$P(0.50|\text{NOT BUY}) = 0.50 \cdot 0.1 / 0.74 \approx 0.0676$$

$$\text{VSI}(\text{NOT BUY}) = E(R(a''|\text{NOT BUY})|\text{NOT BUY}) - E(R(a')|\text{NOT BUY})$$

$$a' = \text{Stock 50 (Same as in (a))}$$

$$a''|\text{NOT BUY} = \arg \max_a \{E(R(a|\text{NOT BUY})|\text{NOT BUY})\}$$

$$E(R(a|\text{NOT BUY})|\text{NOT BUY}) = \sum_a R(a, \theta) \cdot P(\theta|\text{NOT BUY})$$

\Rightarrow

$$E(R(\text{Stock 100}|\text{NOT BUY})|\text{NOT BUY}) = (-10) \cdot 0.2432 \dots + (-2) \cdot 0.3243 \dots + 12 \cdot 0.2838 \dots + 22 \cdot 0.0811 \dots + 40 \cdot 0.0676 \dots \approx 4.81$$

$$E(R(\text{Stock 50}|\text{NOT BUY})|\text{NOT BUY}) = (-4) \cdot 0.2432 \dots + 6 \cdot 0.3243 \dots + 12 \cdot 0.2838 \dots + 26 \cdot 0.0811 \dots + 16 \cdot 0.0676 \dots \approx 6.76$$

$$E(R(\text{Do not stock}|\text{NOT BUY})|\text{NOT BUY}) = 0 \cdot 0.2432 \dots + 0 \cdot 0.3243 \dots + 0 \cdot 0.2838 \dots + 0 \cdot 0.0811 \dots + 0 \cdot 0.0676 \dots = 0$$

$$\Rightarrow a''|\text{NOT BUY} = \text{Stock 50}$$

$$\begin{aligned} \Rightarrow \text{VSI}(\text{NOT BUY}) &= \\ &= E(R(\text{Stock 50}|\text{NOT BUY})|\text{NOT BUY}) - E(R(\text{Stock 50})|\text{NOT BUY}) = 0 \end{aligned}$$

(c) (c) find the expected value of sample information.

$$\text{EVSI} = \sum_y \text{VSI}(y) \cdot P(y) = 3.2 \cdot 0.26 + 0 \cdot 0.74 = 0.832$$

Expected value of perfect information (EVPI)

Perfect information means that there is no uncertainty left for the decision maker. Hence the true state of the world is known.

The *value of perfect information*, VPI is what this information is worth to the decision maker.

Recall the general formulas for value of information:

$$\text{VOI}(I_N) = E(U(a''|I_N)|I_N) - E(U(a')|I_N)$$

$$\text{VOI}(I_N) = E(L(a')|I_N) - E(L(a''|I_N)|I_N)$$

Perfect information means that $I_N \Leftrightarrow \theta$ is known” – it is not about updating the probabilities of the states and not about expectations with respect to prior-predictive distributions.

The optimal action given θ can be written $a_\theta (= \arg \max_i \{U(a_i, \theta)\})$

Hence the formulas changes into

$$\text{VPI}(\theta) = U(a_\theta, \theta) - U(a', \theta)$$

$$\text{VPI}(\theta) = L(a', \theta) - L(a_\theta, \theta)$$

$$\text{VPI}(\theta) = R(a_\theta, \theta) - R(a', \theta) \quad \text{when utility is linear in money}$$

Using the second form of definition we can deduce a general expression for the *expected value of perfect information* (EVPI) as

$$\begin{aligned} \text{EVPI} &= \int (L(a', \theta) - L(a_\theta, \theta)) f'(\theta) d\theta \\ &= \int L(a', \theta) f'(\theta) d\theta - \int L(a_\theta, \theta) f'(\theta) d\theta = E(L(a')) - \int 0 \cdot f'(\theta) d\theta \\ &= E(L(a')) \end{aligned}$$

since the optimal action for each θ comes with zero loss.

Hence, the expected value of perfect information is equal to the expected loss of the optimal action in prior sense.

Exercise 6.15 (again)

15. A store must decide whether or not to stock a new item. The decision depends on the reaction of consumers to the item, and the payoff table (in dollars) is as follows.

		PROPORTION OF CONSUMERS PURCHASING				
		0.10	0.20	0.30	0.40	0.50
DECISION	Stock 100	-10	-2	12	22	40
	Stock 50	-4	6	12	16	16
	Do not stock	0	0	0	0	0

If $P(0.10) = 0.2$, $P(0.20) = 0.3$, $P(0.30) = 0.3$, $P(0.40) = 0.1$, and $P(0.50) = 0.1$, what decision maximizes expected payoff? If perfect information is available, find VPI for each of the five possible states of the world and compute EVPI.

$$VPI(\theta) = R(a_\theta, \theta) - R(a', \theta)$$

DECISION	PROPORTION OF CUSTOMERS BUYING				
	0.10	0.20	0.30	0.40	0.50
Stock 100	-10	-2	12	22	40
Stock 50	-4	6	12	16	16
Do not stock	0	0	0	0	0

From previously: $a' = \text{Stock 50}$

\Rightarrow

$$VPI(0.10) = 0 - (-4) = 4$$

$$VPI(0.20) = 6 - 6 = 0$$

$$VPI(0.30) = 12 - 12 = 0$$

$$VPI(0.40) = 22 - 16 = 6$$

$$VPI(0.50) = 40 - 16 = 24$$

$$\Rightarrow EVPI = 4 \cdot 0.2 + 0 \cdot 0.3 + 0 \cdot 0.3 + 6 \cdot 0.1 + 24 \cdot 0.1 = 3.8$$