Decision Problems

Christos Dimitrakakis

Chalmers

1/11/2013

- Decisions $d \in \mathcal{D}$
- Experiments with outcomes in Ω.
- Reward $r \in \mathcal{R}$ depending on experiment and outcome.
- Utility $U : \mathcal{R} \to \mathbb{R}$.

Example 1 (Taking the umbrella)

- There is some probability of rain.
- We don't like carrying an umbrella.
- We really don't like getting wet.

- Random outcome $\omega \sim P$.
- Decision $d \in D$

Definition 2 (Reward function)

When we take decision d, then ω is randomly chosen, and we obtain a reward:

$$r = \rho(\omega, d). \tag{2.1}$$

For every $d \in \mathcal{D}$, the function $\rho : \Omega \times D \to \mathcal{R}$ induces a probability distribution P_d on $\mathcal{R}$.

$$P_d(B) \triangleq P(\{\omega \mid \rho(\omega, d) \in B\}). \tag{2.2}$$

Thus, instead of directly choosing some distribution of rewards, we choose a decision d, which corresponds to a particular distribution P_d.

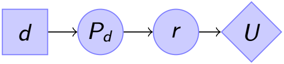

(a) The combined decision problem

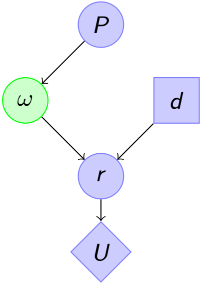

(b) The separated decision problem

Expected utility

$$\mathbb{E}_{P_d}(U) = \sum_{r \in \mathcal{R}} U(r) P_d(r) = \sum_{\omega \in \Omega} U[\rho(\omega, d)] P(\omega). \tag{2.3}$$

Example 3

You are going to work, and it might rain. The forecast said that the probability of rain (ω_1) was 20%. What do you do?

- d_1: Take the umbrella.
- d_2: Risk it!

$\rho(\omega, d)$	d_1	d_2
ω_1	dry, carrying umbrella	wet
ω_2	dry, carrying umbrella	dry

$U[\rho(\omega, d)]$	d_1	d_2
ω_1	0	-10
ω_2	0	1
$\mathbb{E}_P(U \mid d)$	0	-1.2

Table: Rewards, utilities, expected utility for 20% probability of rain.

.

Application to statistical estimation

Example 4 (Voting)

Let us say for example that you wish to estimate the number of votes for different candidates in an election. The *unknown parameters* of the problem mainly include: the percentage of likely voters in the population, the probability that a likely voter is going to vote for each candidate. One simple way to estimate this is by polling.

- The unknown outcome of the experiment ω is called a parameter.

Application to statistical estimation

Example 4 (Voting)

Let us say for example that you wish to estimate the number of votes for different candidates in an election. The *unknown parameters* of the problem mainly include: the percentage of likely voters in the population, the probability that a likely voter is going to vote for each candidate. One simple way to estimate this is by polling.

- The unknown outcome of the experiment ω is called a parameter.
- The set of outcomes Ω is called the parameter space.

Application to statistical estimation

Example 4 (Voting)

Let us say for example that you wish to estimate the number of votes for different candidates in an election. The *unknown parameters* of the problem mainly include: the percentage of likely voters in the population, the probability that a likely voter is going to vote for each candidate. One simple way to estimate this is by polling.

- The unknown outcome of the experiment ω is called a parameter.
- The set of outcomes Ω is called the parameter space.
- We wish to guess a particular value $d \in D = \Omega$ for the parameter.

Application to statistical estimation

Example 4 (Voting)

Let us say for example that you wish to estimate the number of votes for different candidates in an election. The *unknown parameters* of the problem mainly include: the percentage of likely voters in the population, the probability that a likely voter is going to vote for each candidate. One simple way to estimate this is by polling.

- The unknown outcome of the experiment ω is called a parameter.
- The set of outcomes Ω is called the parameter space.
- We wish to guess a particular value $d \in D = \Omega$ for the parameter.
- $\rho(\omega, d)$ measures how close our guess is to the parameter.

Application to statistical estimation

- The unknown outcome of the experiment ω is called a parameter.
- The set of outcomes Ω is called the parameter space.
- We wish to guess a particular value $d \in D = \Omega$ for the parameter.
- $\rho(\omega, d)$ measures how close our guess is to the parameter.

Definition 4 (Simplified expected utility of a given decision)

$$U(P, d) \triangleq \sum_{\omega \in \Omega} U[\rho(\omega, d)] P(\omega). \tag{2.4}$$

Definition 5 (Bayes-optimal utility)

$$U^*(P) \triangleq \max_d U(P, d) \tag{2.5}$$

Voting example

- Consider a nation with k political parties.
- Let $\omega = (\omega_1, \ldots, \omega_k) \in [0, 1]^k$ be the voting percentages for each party.
- We wish to make a guess $d \in [0, 1]^k$.
- How should we guess, given a distribution $P(\omega)$?
- How should we select U and ρ?

Voting example

- Consider a nation with k political parties.
- Let $\omega = (\omega_1, \ldots, \omega_k) \in [0,1]^k$ be the voting percentages for each party.
- We wish to make a guess $d \in [0,1]^k$.
- How should we guess, given a distribution $P(\omega)$?
- How should we select U and ρ?

Squared error

We can set $\rho(\omega, d) = (\omega_1 - d_1, \ldots, \omega_k - d_k)$, our error vector $r \in [0,1]^k$. Then we set $U(r) \triangleq -\|r\|^2$, where $\|r\|^2 = \sum_i |x_i|^2$.

Voting example

- Consider a nation with k political parties.
- Let $\omega = (\omega_1, \ldots, \omega_k) \in [0,1]^k$ be the voting percentages for each party.
- We wish to make a guess $d \in [0,1]^k$.
- How should we guess, given a distribution $P(\omega)$?
- How should we select U and ρ?

Squared error

We can set $\rho(\omega, d) = (\omega_1 - d_1, \ldots, \omega_k - d_k)$, our error vector $r \in [0,1]^k$. Then we set $U(r) \triangleq -\|r\|^2$, where $\|r\|^2 = \sum_i |x_i|^2$.

Predicting the winner

In that case $\rho(\omega, d) = 1$ if $\arg\max_i \omega_i = \arg\max_i d_i$ and 0 otherwise, and $U(r) = r$.

Definition 6 (The set of maximising arguments)

The set of all maximising values of a function f is denoted by $\arg\max_x f(x)$. More formally,

$$\arg\max_x f(x) = \{x \mid f(x) \geq g(y) \forall y\}.$$

If there is no maximising value, then the set is empty.

Example 7 (Squared error)

Consider the case $\Omega = D = \mathbb{R}$. Our problem is:

$$\max_d U(P, d), \qquad U(P, d) \triangleq -\sum_\omega |\omega - d|^2 P(\omega).$$

Example 7 (Squared error)

Consider the case $\Omega = D = \mathbb{R}$. Our problem is:

$$\max_d U(P, d), \qquad U(P, d) \triangleq -\sum_\omega |\omega - d|^2 P(\omega).$$

By taking the differential inside the sum, we have

$$\frac{\partial}{\partial d} \sum_\omega |\omega - d|^2 P(\omega) = \sum_\omega \frac{\partial}{\partial d} |\omega - d|^2 P(\omega) \tag{2.9}$$

Example 7 (Squared error)

Consider the case $\Omega = D = \mathbb{R}$. Our problem is:

$$\max_d U(P, d), \qquad U(P, d) \triangleq -\sum_\omega |\omega - d|^2 P(\omega).$$

By taking the differential inside the sum, we have

$$\begin{aligned} \frac{\partial}{\partial d} \sum_\omega |\omega - d|^2 P(\omega) &= \sum_\omega \frac{\partial}{\partial d} |\omega - d|^2 P(\omega) \qquad (2.6) \\ &= 2 \sum_\omega (d - \omega) P(\omega) \end{aligned}$$

$$(2.9)$$

Example 7 (Squared error)

Consider the case $\Omega = D = \mathbb{R}$. Our problem is:

$$\max_d U(P, d), \qquad U(P, d) \triangleq -\sum_\omega |\omega - d|^2 P(\omega).$$

By taking the differential inside the sum, we have

$$\begin{aligned}
\frac{\partial}{\partial d} \sum_\omega |\omega - d|^2 P(\omega) &= \sum_\omega \frac{\partial}{\partial d} |\omega - d|^2 P(\omega) && (2.6)\\
&= 2 \sum_\omega (d - \omega) P(\omega) && (2.7)\\
&= 2 \sum_\omega d P(\omega) - 2 \sum_\omega \omega P(\omega) &&
\end{aligned}$$

$$(2.9)$$

Example 7 (Squared error)

Consider the case $\Omega = D = \mathbb{R}$. Our problem is:

$$\max_d U(P, d), \qquad U(P, d) \triangleq -\sum_\omega |\omega - d|^2 P(\omega).$$

By taking the differential inside the sum, we have

$$\frac{\partial}{\partial d} \sum_\omega |\omega - d|^2 P(\omega) = \sum_\omega \frac{\partial}{\partial d} |\omega - d|^2 P(\omega) \tag{2.6}$$

$$= 2 \sum_\omega (d - \omega) P(\omega) \tag{2.7}$$

$$= 2 \sum_\omega d P(\omega) - 2 \sum_\omega \omega P(\omega) \tag{2.8}$$

$$= 2d - 2\,\mathbb{E}_P(\omega), \tag{2.9}$$

so the optimal decision is $d = \mathbb{E}(\omega)$.

The utility for quadratic loss

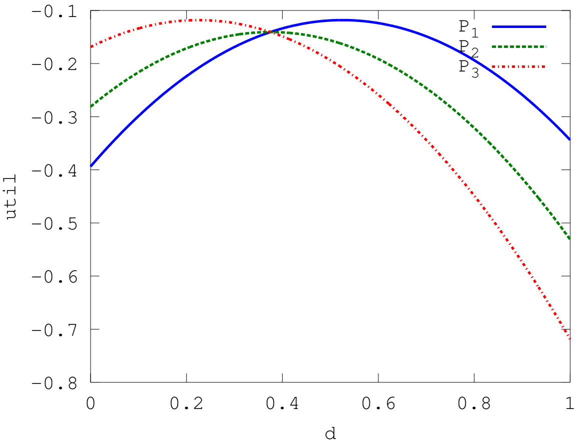

Figure: Fixed distribution, varying decision. The decision utility under three different distributions.

Only prior information

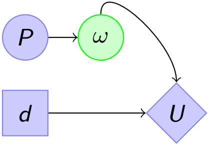

Figure: Statistical decision problem without observations. The green circles are unknown quantities. Squares indicate decisions. Diamonds indicate utilities.

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.

Only prior information

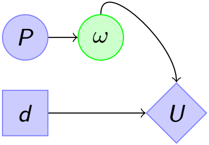

Figure: Statistical decision problem without observations. The green circles are unknown quantities. Squares indicate decisions. Diamonds indicate utilities.

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Our utility is $U : \Omega \times D \to \mathbb{R}$.

Only prior information

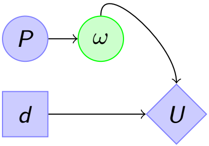

Figure: Statistical decision problem without observations. The green circles are unknown quantities. Squares indicate decisions. Diamonds indicate utilities.

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Our utility is $U : \Omega \times D \to \mathbb{R}$.
3. We want to choose $d \in \mathcal{D}$, taking into account P:

$$\max_d U(P, d) = \max_d \sum_{\omega \in \Omega} U(\omega, d) P(\omega).$$

Obtaining information

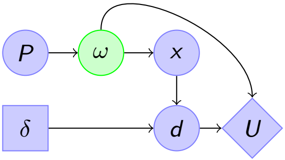

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.

Obtaining information

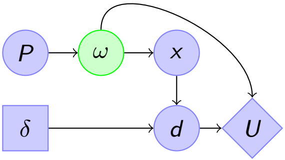

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Now consider a family of conditional probabilities on the observation set $\mathcal{S}$:

$$\mathcal{F} = \{P(\cdot \mid \omega) \mid \omega \in \Omega\},$$

such that $P(x \mid \omega)$ is the probability of observing $x \in \mathcal{S}$ under parameter ω.

Obtaining information

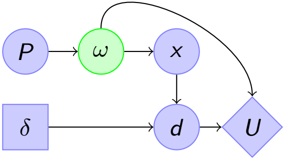

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Now consider a family of conditional probabilities on the observation set $\mathcal{S}$:

$$\mathcal{F} = \{P(\cdot \mid \omega) \mid \omega \in \Omega\},$$

   such that $P(x \mid \omega)$ is the probability of observing $x \in \mathcal{S}$ under parameter ω.
3. Let $x \in \mathcal{S}$ be a random variable with distribution $P(x \mid \omega)$ for some (unknown) ω.

Obtaining information

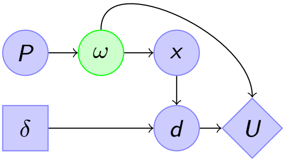

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Now consider a family of conditional probabilities on the observation set $\mathcal{S}$:

$$\mathcal{F} = \{P(\cdot \mid \omega) \mid \omega \in \Omega\},$$

   such that $P(x \mid \omega)$ is the probability of observing $x \in \mathcal{S}$ under parameter ω.
3. Let $x \in \mathcal{S}$ be a random variable with distribution $P(x \mid \omega)$ for some (unknown) ω.
4. Our utility is $U : \Omega \times D \to \mathbb{R}$.

Obtaining information

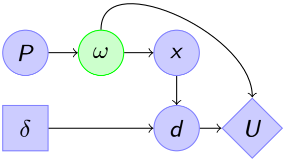

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Now consider a family of conditional probabilities on the observation set $\mathcal{S}$:

$$\mathcal{F} = \{P(\cdot \mid \omega) \mid \omega \in \Omega\},$$

   such that $P(x \mid \omega)$ is the probability of observing $x \in \mathcal{S}$ under parameter ω.
3. Let $x \in \mathcal{S}$ be a random variable with distribution $P(x \mid \omega)$ for some (unknown) ω.
4. Our utility is $U : \Omega \times D \to \mathbb{R}$.
5. We want to choose $d \in \mathcal{D}$, taking into account both P and the evidence x.

Obtaining information

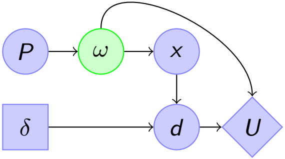

Figure: Statistical decision problem with observations

1. There is an unknown parameter $\omega \in \Omega$ with $\omega \sim P$.
2. Now consider a family of conditional probabilities on the observation set $\mathcal{S}$:

$$\mathcal{F} = \{P(\cdot \mid \omega) \mid \omega \in \Omega\},$$

   such that $P(x \mid \omega)$ is the probability of observing $x \in \mathcal{S}$ under parameter ω.
3. Let $x \in \mathcal{S}$ be a random variable with distribution $P(x \mid \omega)$ for some (unknown) ω.
4. Our utility is $U : \Omega \times D \to \mathbb{R}$.
5. We want to choose $d \in \mathcal{D}$, taking into account both P and the evidence x.
6. We want to find a decision function $\delta : \mathcal{S} \to D$ maximising expected utility

Maximising expected utility a posteriori

- Prior probability $P(\omega)$

Maximising expected utility a posteriori

- Prior probability $P(\omega)$
- Observation x.

Maximising expected utility a posteriori

- Prior probability $P(\omega)$
- Observation x.
- Posterior probability

$$P(\omega \mid x) = \frac{P(x \mid \omega)P(\omega)}{\sum_{\omega'} P(x \mid \omega')P(\omega')}$$

Maximising expected utility a posteriori

- Prior probability $P(\omega)$
- Observation x.
- Posterior probability

$$P(\omega \mid x) = \frac{P(x \mid \omega)P(\omega)}{\sum_{\omega'} P(x \mid \omega')P(\omega')}$$

- Expected utility of decision d under the posterior

$$\mathbb{E}_P(U \mid d, x) = \sum_{\omega \in \Omega} U(\omega, d)P(\omega \mid x)$$

Maximising expected utility a posteriori

- Prior probability $P(\omega)$
- Observation x.
- Posterior probability

$$P(\omega \mid x) = \frac{P(x \mid \omega)P(\omega)}{\sum_{\omega'} P(x \mid \omega')P(\omega')}$$

- Expected utility of decision d under the posterior

$$\mathbb{E}_P(U \mid d, x) = \sum_{\omega \in \Omega} U(\omega, d)P(\omega \mid x)$$

- Bayes decision rule:

$$\delta^*(x) \in \arg\max_{d \in D} \mathbb{E}_P(U \mid d, x).$$

Exercise

Abdul Alhazred claims that he is psychic and can always predict a coin toss. Let $P(A) = 2^{-16}$ be your prior belief that AA is a psychic.

- Abdul bets you 100 CU that he can predict the next four coin tosses. How much are you willing to bet against that (assuming that you are using a fair coin).
- You throw the coin 4 times, and AA guesses correctly all four times. Abdul now bets you another 100 CU that he can predict the next four coin tosses. Up to how much would you bet now?

Assumption 1

- *You use a fair coin, such that the probability of it coming heads is* $1/2$.
- *Your utility for money is linear, i.e.* $U(x) = x$ *for any amount of money* x.

Quick summary

- We want to make a decision against an unknown parameter ω.
- Our knowledge is represented by a distribution $P(\omega)$.
- The Bayes utility is the maximum expected utility under the distribution of ω.
- Our decisions can depend on observations, via a decision function.
- We can construct a complete decision function by computing the optimal decision for every possible observation.
- We can instead wait until we observe x and then:
  1. Compute the posterior distribution $P(\omega \mid x)$.
  2. Compute the expected utility under the distribution $P(\omega \mid x)$ of ω, for all decisions d.
  3. Choose the decision with the highest expected utility *a posteriori*.

  Thus, decision making under prior and posterior distributions can be handled in the same exact framework.