Decision Problems

Decision Making under Uncertainty, Part III

Christos Dimitrakakis

Chalmers

1/11/2013

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

Example (Taking the umbrella)

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

Assumption (Outcomes)

There exists a probability measure P on $(\Omega, \mathcal{F}_\Omega)$ such that the probability of the random outcome ω being in $A \subset \Omega$ is:

$$\mathbb{P}(\omega \in A) = P(A), \qquad \forall A \in \mathcal{F}_\Omega. \tag{2.1}$$

Assumption (Utilities)

Preferences about rewards in R are transitive, all rewards are comparable and there exists a utility function U, measurable with respect to $\mathcal{F}_R$ such that $U(r') \geq U(r)$ iff $r \succ^* r'$.

Definition (Reward function)

$$r = \rho(\omega, d). \tag{2.2}$$

The probability measure induced by decisions

For every $d \in D$, the function $\rho : \Omega \times D \to R$ induces a probability P_d on R. In fact, for any $B \in \mathcal{F}_R$:

$$P_d(B) \triangleq \mathbb{P}(\rho(\omega, d) \in B) = P(\{\omega \mid \rho(\omega, d) \in B\}). \tag{2.3}$$

Assumption

The sets $\{\omega \mid \rho(\omega, d) \in B\}$ must belong to $\mathcal{F}_\Omega$. In other words, ρ must be $\mathcal{F}_\Omega$-measurable for any d.

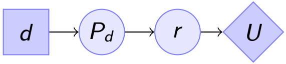

(a) The combined decision problem

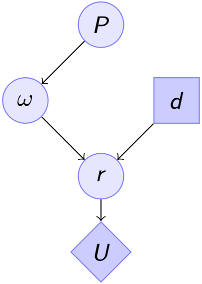

(b) The separated decision problem

Expected utility

$$\mathbb{E}_{P_i}(U) = \int_R U(r)\,\mathrm{d}P_i(r) = \int_\Omega U[\rho(\omega, i)]\,\mathrm{d}P(\omega) = \mathbb{E}_P(U \mid d = i) \tag{2.4}$$

Example

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

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

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.

Definition (Loss)

$$\ell(\omega, d) = -U[\rho(\omega, d)]. \tag{2.5}$$

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.

Definition (Loss)

$$\ell(\omega, d) = -U[\rho(\omega, d)]. \tag{2.5}$$

Definition (Risk)

$$\sigma(P, d) = \int_{\Omega} \ell(\omega, d) \, \mathrm{d}P(\omega). \tag{2.6}$$

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.

Definition (Loss)

$$\ell(\omega, d) = -U[\rho(\omega, d)]. \tag{2.5}$$

Definition (Risk)

$$\sigma(P, d) = \int_{\Omega} \ell(\omega, d) \, \mathrm{d}P(\omega). \tag{2.6}$$

Of course, the optimal decision is d minimising σ.

Bayes risk

Consider parameter space Ω, decision space D, loss function ℓ.

Definition (Bayes risk)

$$\sigma^*(P) = \inf_{d \in D} \sigma(P, d) \tag{3.1}$$

Remark

For any function $f : X \to Y$, where Y is equipped with a complete binary relation $<$, we define, for any $A \subset X$

$$M = \inf_{x \in A} f(x)$$

s.t. $M \leq f(x)$ for any $x \in A$. Furthermore, for any $M' > M$, there exists some $x' \in A$ s.t. $M' > f(x')$.

Example

Let $\Omega = \{0, 1\}$ and $D = [0, 1]$. For an $\alpha \geq 1$, we define the loss $L : \Omega \times D \to \mathbb{R}$ as

$$\ell(\omega, d) = |\omega - d|^{\alpha}. \tag{3.2}$$

Assume that the distribution of outcomes is

$$\mathbb{P}(\omega = 0) = u \qquad \mathbb{P}(\omega = 1) = 1 - u. \tag{3.3}$$

For $\alpha = 1$ we have

$$\sigma(P, d) = \ell(0, d)u + \ell(1, d)(1 - u) = du + (1 - d)(1 - u). \tag{3.4}$$

Hence, if $u > 1/2$ the risk is minimised for $d^* = 0$, otherwise for $d^* = 1$.

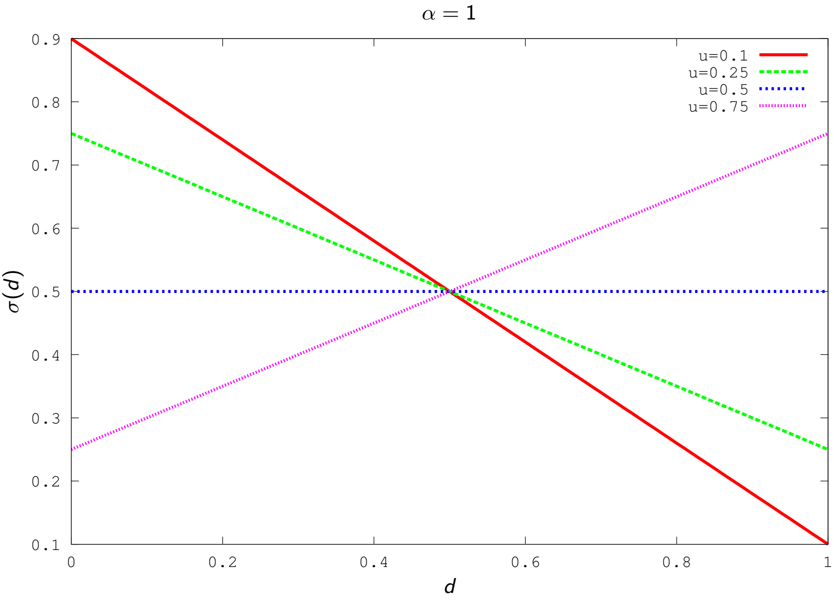

Figure: Risk for four different distributions with absolute loss.

Example

Let $\Omega = \{0, 1\}$ and $D = [0, 1]$. For an $\alpha \geq 1$, we define the loss $L : \Omega \times D \to \mathbb{R}$ as

$$\ell(\omega, d) = |\omega - d|^{\alpha}. \tag{3.2}$$

Assume that the distribution of outcomes is

$$\mathbb{P}(\omega = 0) = u \qquad \mathbb{P}(\omega = 1) = 1 - u. \tag{3.3}$$

For $\alpha > 1$,

$$\sigma(P, d) = d^{\alpha} u + (1 - d)^{\alpha}(1 - u), \tag{3.4}$$

and by differentiating we find that the optimal decision is

$$d^* = \left[1 + \left(\frac{1}{1/u - 1}\right)^{\frac{1}{\alpha - 1}}\right]^{-1}.$$

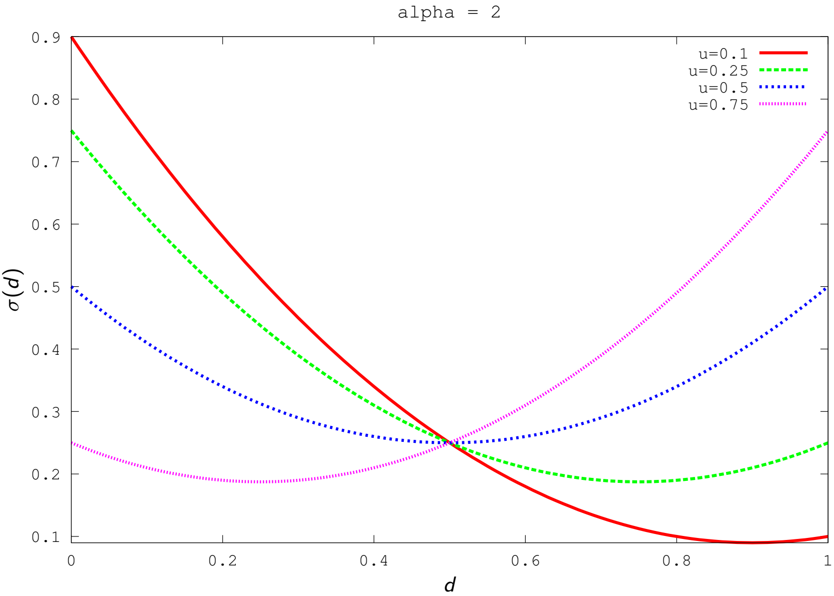

Figure: Risk for four different distributions with quadratic loss.

Example (Quadratic loss)

Now consider $\Omega = \mathbb{R}$ with measure P and $\mathcal{D} = \mathbb{R}$. For any point $\omega \in \mathbb{R}$, the loss is

$$\ell(\omega, d) = |\omega - d|^2. \tag{3.4}$$

The optimal decision minimises

$$\mathbb{E}(\ell \mid d) = \int_{\mathbb{R}} |\omega - d|^2 \, \mathrm{d}P(\omega).$$

Then, as long as $\partial/\partial d |\omega - d|^2$ is measurable with respect to $\mathcal{F}_{\mathbb{R}}$

$$\frac{\partial}{\partial d} \int_{\mathbb{R}} |\omega - d|^2 \, \mathrm{d}P(\omega) = \int_{\mathbb{R}} \frac{\partial}{\partial d} |\omega - d|^2 \, \mathrm{d}P(\omega) \tag{3.5}$$

$$= 2 \int_{\mathbb{R}} (\omega - d) \, \mathrm{d}P(\omega) \tag{3.6}$$

$$= 2 \int_{\mathbb{R}} \omega \, \mathrm{d}P(\omega) - 2 \int_{\mathbb{R}} d \, \mathrm{d}P(\omega) \tag{3.7}$$

$$= 2 \, \mathbb{E}(\omega) - 2d, \tag{3.8}$$

so the cost is minimised for $d = \mathbb{E}(\omega)$.

A mixture of distributions

Consider two probability measures P, Q on $(\Omega, \mathcal{F}_\Omega)$.

A mixture of distributions

Consider two probability measures P, Q on $(\Omega, \mathcal{F}_\Omega)$.
These define two alternative distributions for ω. For any A

A mixture of distributions

Consider two probability measures P, Q on $(\Omega, \mathcal{F}_\Omega)$.
These define two alternative distributions for ω. For any A For any P, Q and $\alpha \in [0, 1]$, we define

$$Z_\alpha = \alpha P + (1 - \alpha) Q$$

to mean the probability measure such that

$$Z_\alpha(A) = \alpha P(A) + (1 - \alpha) Q(A)$$

for any $A \in \mathcal{F}_\Omega$.

Concavity of the Bayes risk

Theorem

For probability measures P, Q on $\mathcal{F}_\Omega$ and any $\alpha \in [0, 1]$

$$\sigma^*[\alpha P + (1-\alpha)Q] \geq \alpha\sigma^*(P) + (1-\alpha)\sigma^*(Q). \tag{3.9}$$

Proof.

From the definition of risk (2.6), for any decision $d \in D$,

$$\sigma[\alpha P + (1-\alpha)Q, d] = \alpha\sigma(P, d) + (1-\alpha)\sigma(Q, d).$$

Hence, by definition (3.1) of the Bayes risk,

$$\begin{aligned}\sigma^*[\alpha P + (1-\alpha)Q] &= \inf_{d \in D} \sigma[\alpha P + (1-\alpha)Q, d] \\ &= \inf_{d \in D} [\alpha\sigma(P, d) + (1-\alpha)\sigma(Q, d)].\end{aligned}$$

□

Concavity of the Bayes risk

Theorem

For probability measures P, Q *on* $\mathcal{F}_\Omega$ *and any* $\alpha \in [0, 1]$

$$\sigma^*[\alpha P + (1-\alpha)Q] \geq \alpha\sigma^*(P) + (1-\alpha)\sigma^*(Q). \tag{3.9}$$

Proof.

$$\sigma^*[\alpha P + (1-\alpha)Q] = \inf_{d \in D} [\alpha\sigma(P, d) + (1-\alpha)\sigma(Q, d)].$$

□

Concavity of the Bayes risk

Theorem

For probability measures P, Q on $\mathcal{F}_\Omega$ and any $\alpha \in [0, 1]$

$$\sigma^*[\alpha P + (1-\alpha)Q] \geq \alpha\sigma^*(P) + (1-\alpha)\sigma^*(Q). \tag{3.9}$$

Proof.

$$\sigma^*[\alpha P + (1-\alpha)Q] = \inf_{d \in D}[\alpha\sigma(P, d) + (1-\alpha)\sigma(Q, d)].$$

Since $\inf_x[f(x) + g(x)] \geq \inf_x f(x) + \inf_x g(x)$,

$$\begin{aligned}\sigma^*[\alpha P + (1-\alpha)Q] &\geq \alpha \inf_{d \in D} \sigma(P, d) + (1-\alpha) \inf_{d \in D} \sigma(Q, d)\\ &= \alpha\sigma^*(P) + (1-\alpha)\sigma^*(Q).\end{aligned}$$

□

The risk function for quadratic loss

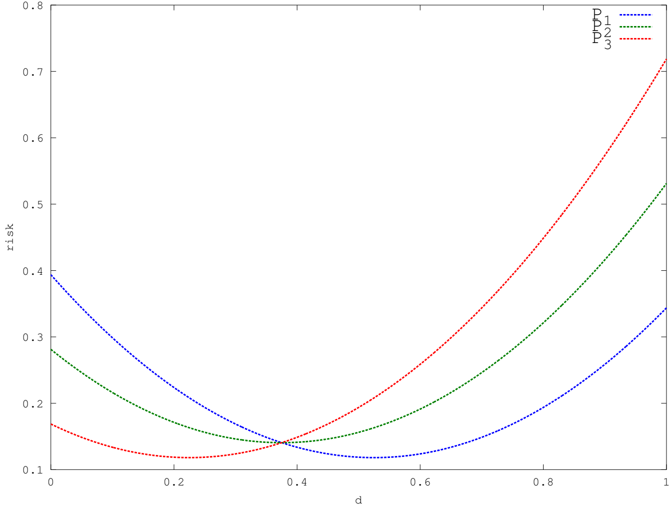

Figure: Fixed distribution, varying decision. The decision risk under three different distributions.

Concavity of the Bayes risk

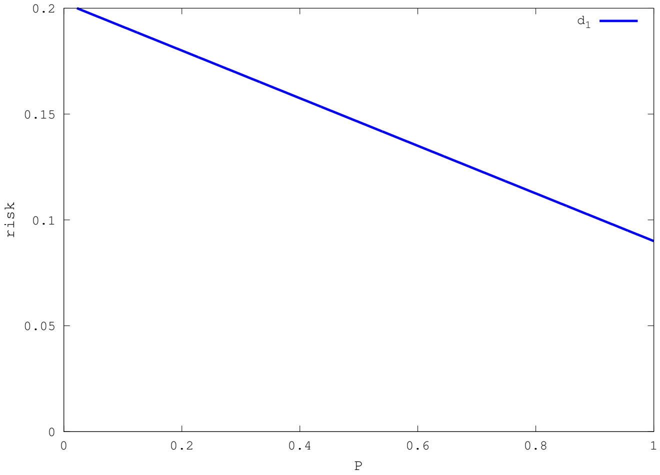

Figure: Fixed decision, varying distribution. The risk of a fixed decision is a linear function of P

Concavity of the Bayes risk

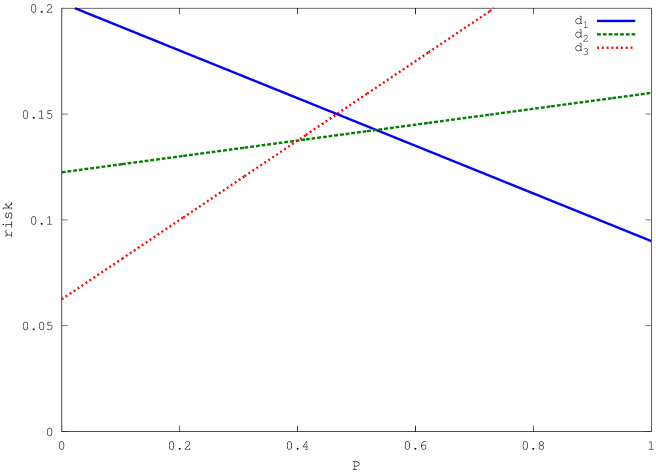

Figure: The risk of a few decisions as P varies. Each decision corresponds to one of these lines.

Concavity of the Bayes risk

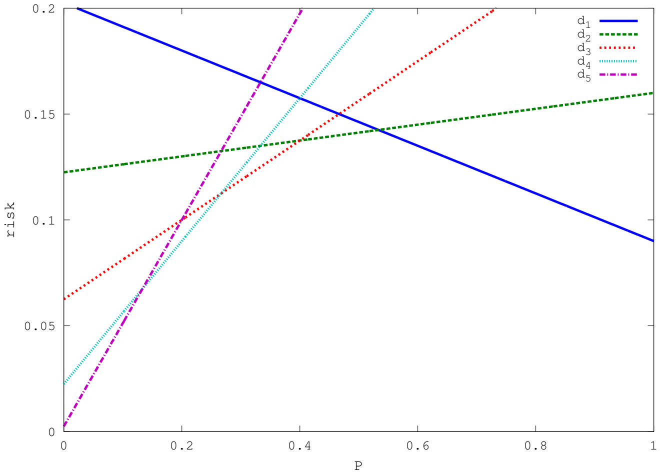

Figure: For each P, there is at least one decision minimising the risk.

Concavity of the Bayes risk

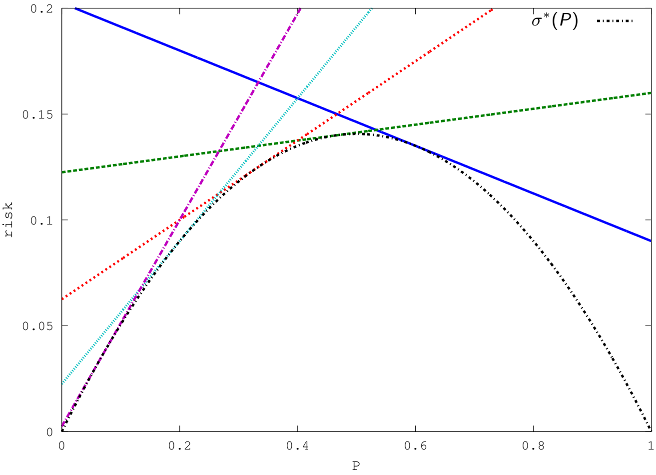

Figure: The Bayes risk is concave and the minimising decision is tangent to it.

Concavity of the Bayes risk

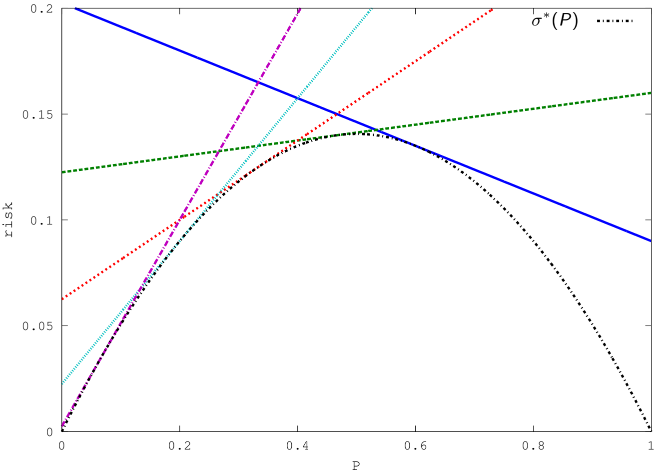

Figure: If we are not very wrong about P, then we are not far from optimal.

Concavity of the Bayes risk

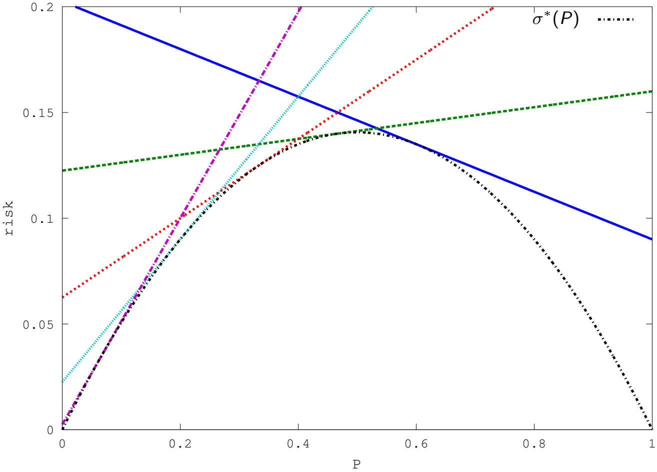

Figure: We can approximate the Bayes risk by taking the minimum of a finite number of decisions.

Mixed decisions

A distribution over decisions

- Consider a probability measure π on $\mathcal{D}$.
- We select decisions according to probability

$$\pi(A) \triangleq \mathbb{P}(d \in A).$$

for any appropriate $A \subset \mathcal{D}$.

Theorem

Consider any statistical decision problem with probability measure P on outcomes Ω and with utility function $U : \Omega \times \mathcal{D} \to \mathbb{R}$. Further let $d^* \in \mathcal{D}$ such that $\mathbb{E}(U \mid d^*) \geq \mathbb{E}(U \mid d)$ for all $d \in \mathcal{D}$. Then for any probability measure π on $\mathcal{D}$,

$$\mathbb{E}(U \mid d^*) \geq \mathbb{E}(U \mid \pi).$$

Mixed decisions

Theorem

Consider any statistical decision problem with probability measure P on outcomes Ω and with utility function $U : \Omega \times \mathcal{D} \to \mathbb{R}$. Further let $d^* \in \mathcal{D}$ such that $\mathbb{E}(U \mid d^*) \geq \mathbb{E}(U \mid d)$ for all $d \in \mathcal{D}$. Then for any probability measure π on $\mathcal{D}$,

$$\mathbb{E}(U \mid d^*) \geq \mathbb{E}(U \mid \pi).$$

Proof.

$$\begin{aligned}
\mathbb{E}(U \mid \pi) &= \int_D \mathbb{E}(U \mid d) \, \mathrm{d}\pi(d) \\
&\leq \int_D \mathbb{E}(U \mid d^*) \, \mathrm{d}\pi(d) \\
&= \mathbb{E}(U \mid d^*) \int_D \mathrm{d}\pi(d) = \mathbb{E}(U \mid d^*)
\end{aligned}$$

□

Alternative decision rules

Maximin rule

Select d maximising $\min_{w \in W} U(w, d)$.

ϵ-optimal rule

For some $\epsilon > 0$, select d maximising

$$P\left(\left\{\omega \;\middle|\; U(\omega, d) > \inf_{d' \in \mathcal{D}} U(\omega, d') + \epsilon\right\}\right). \tag{4.1}$$

Minimax/Maximin values

$$U_* = \max_d \min_\omega U(\omega, d) = \min_\omega U(\omega, d^*) \qquad \text{(maximin)}$$

$$U^* = \min_\omega \max_d U(\omega, d) = \max_d U(\omega^*, d), \qquad \text{(minimax)}$$

Note that by definition

$$U^* \geq U(\omega^*, d^*) \geq U_*. \qquad (4.2)$$

Regret

Consider a problem with two possible outcomes ω_1, ω_2, two possible decisions, d_1, d_2, a utility function $U(\omega, d)$ and a prior distribution $P(\omega_i) = 1/2$.

$U(\omega, d)$	d_1	d_2
ω_1	-1	0
ω_2	10	1
$\mathbb{E}(U \mid P, d)$	4.5	0.5
$\min_\omega U(\omega, d)$	-1	0

Table: Utility function, expected utility and maximin utility.

Regret

Consider a problem with two possible outcomes ω_1, ω_2, two possible decisions, d_1, d_2, a utility function $U(\omega, d)$ and a prior distribution $P(\omega_i) = 1/2$.

$U(\omega, d)$	d_1	d_2
ω_1	-1	0
ω_2	10	1
$\mathbb{E}(U \mid P, d)$	4.5	0.5
$\min_\omega U(\omega, d)$	-1	0

Table: Utility function, expected utility and maximin utility.

Definition (Regret)

$$L(\omega, \pi) \triangleq \max_{\pi'} U(\omega, \pi') - U(\omega, \pi). \tag{4.3}$$

Regret

Consider a problem with two possible outcomes ω_1, ω_2, two possible decisions, d_1, d_2, a utility function $U(\omega, d)$ and a prior distribution $P(\omega_i) = 1/2$.

$U(\omega, d)$	d_1	d_2
ω_1	-1	0
ω_2	10	1
$\mathbb{E}(U \mid P, d)$	4.5	0.5
$\min_\omega U(\omega, d)$	-1	0

Table: Utility function, expected utility and maximin utility.

$L(\omega, d)$	d_1	d_2
ω_1	1	0
ω_2	0	9
$\mathbb{E}(L \mid P, d)$	0.5	4.5
$\max_\omega L(\omega, d)$	1	9

Table: Regret, in expectation and minimax.

Minimax utility, regret and loss

Remark

For each ω, there is some d such that:

$$U(\omega, d) \in \max_{\pi} U(\omega, \pi). \tag{4.3}$$

Remark

$$L(\omega, \pi) = \sum_{d} \pi(d) L(w, d) \geq 0, \tag{4.4}$$

with equality iff π is ω-optimal.

Remark

$$L(\omega, \pi) = \max_{d} U(\omega, d) - U(\omega, \pi). \tag{4.5}$$

Remark

$L(\omega, \pi) = -U(\omega, \pi) = L(\omega, \pi)$ *iff* $\max_d U(\omega, d) = 0$.

Example

(An even-money bet)

U	ω_1	ω_2
d_1	1	-1
d_2	0	0

Table: Even-bet utility

For two distributions π, ξ on D and Ω, define our expected utility to be:

$$U(\xi, \pi) \triangleq \sum_{w \in \Omega} \sum_{d \in D} U(w, d) \xi(w) \pi(d). \tag{4.6}$$

Then we define the maximin policy π^* such that:

$$\min_\xi U(\xi, \pi^*) = U_* \triangleq \max_\pi \min_\xi U(\xi, \pi) \tag{4.7}$$

Then we define the minimax prior ξ^* such that

$$\max_\pi U(\xi^*, \pi) = U^* \triangleq \min_\xi \max_\pi U(\xi, \pi) \tag{4.8}$$

Expected regret

$$\begin{aligned} L(\xi, \pi) &= \max_{\pi'} \sum_w \xi(w) \left\{ U(w, \pi') - U(w, \pi) \right\} \\ &= \max_{\pi'} U(\xi, \pi') - U(\xi, \pi). \end{aligned} \tag{4.9}$$

Theorem

If there exist $\xi^*, \pi^* \in D$ and $C \in \mathbb{R}$ such that

$$U(\xi^*, \pi) \leq C \leq U(\xi, \pi^*)$$

then

$$U^* = U_* = U(\xi^*, \pi^*) = C.$$

Definition

A bilinear game is a tuple (U, Ξ, Π, Ω, D) with $U : \Xi \times \Pi \to \mathbb{R}$ such that all $\xi \in \Xi$ are arbitrary distributions on Ω and all $\pi \in \Pi$ are arbitrary distributions on D:

$$U(\xi, \pi) \triangleq \mathbb{E}(U \mid \xi, \pi) = \sum_{w,d} U(w, d)\pi(d)\xi(w).$$

Theorem

For a bilinear game, $U^* = U_*$. In addition, the following three conditions are equivalent:

1. *π^* is maximin, ξ^* is minimax and $U^* = C$.*
2. *$U(\xi, \pi^*) \geq C \geq U(\xi^*, \pi)$ for all ξ, π.*
3. *$U(w, \pi^*) \geq C \geq U(\xi^*, d)$ for all w, d.*

Linear programming formulation

The problem

$$\max_{\pi} \min_{\xi} U(\xi, \pi),$$

where ξ, π are distributions over finite domains, can be converted to finding π with the greatest lower bound. Using matrix notation,

$$\max \left\{ v_\pi \;\middle|\; (\boldsymbol{U}\boldsymbol{\pi})_j \geq v_\pi \forall j, \;\; \sum_i \pi_i = 1, \;\; \pi_i \geq 0 \forall i \right\},$$

where everything has been written in matrix form. Equivalently, we can find ξ with the least upper bound:

$$\min \left\{ v_\xi \;\middle|\; (\boldsymbol{\xi}^\top \boldsymbol{U})_i \leq v_\xi \forall i, \;\; \sum_j \xi_j = 1, \;\; \xi_j \geq 0 \forall j \right\},$$

where everything has been written in matrix form. In fact, one can show that $v_\xi = v_\pi$, thus obtaining Theorem 2.

Obtaining information

1. We must choose a decision from D.

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.
3. Our loss is $\ell : \Omega \times D \to \mathbb{R}$.

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.
3. Our loss is $\ell : \Omega \times D \to \mathbb{R}$.
4. Now consider a family of probability measures on the observation set $\mathcal{S}$:

$$\{\psi_\omega \mid \omega \in \Omega\}.$$

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.
3. Our loss is $\ell : \Omega \times D \to \mathbb{R}$.
4. Now consider a family of probability measures on the observation set $\mathcal{S}$:

$$\{\psi_\omega \mid \omega \in \Omega\}.$$

5. Let $x \in \mathcal{S}$ be a random variable with distribution ψ_ω.

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.
3. Our loss is $\ell : \Omega \times D \to \mathbb{R}$.
4. Now consider a family of probability measures on the observation set $\mathcal{S}$:

$$\{\psi_\omega \mid \omega \in \Omega\} .$$

5. Let $x \in \mathcal{S}$ be a random variable with distribution ψ_ω.
6. We want to choose $d \in D$, taking into account both ξ and the evidence x.

Obtaining information

1. We must choose a decision from D.
2. There is an unknown parameter $\omega \in \Omega$ with measure ξ.
3. Our loss is $\ell : \Omega \times D \to \mathbb{R}$.
4. Now consider a family of probability measures on the observation set $\mathcal{S}$:

$$\{\psi_\omega \mid \omega \in \Omega\}.$$

5. Let $x \in \mathcal{S}$ be a random variable with distribution ψ_ω.
6. We want to choose $d \in D$, taking into account both ξ and the evidence x.
7. We want to find a decision function $\delta : S \to D$ that minimises the risk

$$\sigma(\xi, \delta) = \mathbb{E}\{\ell[\omega, \delta(X)]\} = \int_\Omega \left(\int_S \ell[\omega, \delta(x)] \, \mathrm{d}\psi_\omega(x) \right) \mathrm{d}\xi(\omega).$$

Minimising the risk

Expected loss of a fixed decision d with $\omega \sim \xi$

$$\sigma(\xi, d) = \int_\Omega L(\omega, d) \, \mathrm{d}\xi(\omega). \tag{5.1}$$

Expected loss of a decision function δ with fixed $\omega \in \Omega$

$$\sigma(\omega, \delta) = \int_S L(\omega, \delta(x)) \, \mathrm{d}\psi_\omega(x). \tag{5.2}$$

Expected loss of a decision function δ with $W \sim \xi$

$$\sigma(\xi, \delta) = \int_\Omega \rho(\omega, \delta) \, \mathrm{d}\xi(\omega), \qquad \sigma^*(\xi) \triangleq \inf_\delta \sigma(\xi, \delta) = \rho(\xi, \delta^*). \tag{5.3}$$

Bayes decision functions

Extensive form

$$\sigma(\xi, \delta) = \int_{\Omega} \int_{S} \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega) \, \mathrm{d}\psi_{\omega}(x) \tag{5.4}$$

$$= \int_{S} \int_{\Omega} \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega \mid x) \, \mathrm{d}f(x), \tag{5.5}$$

where $f(x) = \int_{\Omega} \psi_{\omega}(x) \, \mathrm{d}\xi(\omega)$.

$$\delta^*(x) \triangleq \arg\max_{d \in D} \mathbb{E}_{\xi}(\ell \mid x, d) = \arg\max_{d \in D} \int_{\Omega} \ell(w, d) \, \mathrm{d}\xi(w \mid x).$$

Bayes decision functions

Extensive form

$$\sigma(\xi, \delta) = \int_\Omega \int_S \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega) \, \mathrm{d}\psi_\omega(x) \tag{5.4}$$

$$= \int_S \int_\Omega \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega \mid x) \, \mathrm{d}f(x), \tag{5.5}$$

where $f(x) = \int_\Omega \psi_\omega(x) \, \mathrm{d}\xi(\omega)$.

$$\delta^*(x) \triangleq \arg\max_{d \in D} \mathbb{E}_\xi(\ell \mid x, d) = \arg\max_{d \in D} \int_\Omega \ell(w, d) \, \mathrm{d}\xi(w \mid x).$$

$$\int_S \int_\Omega \ell[w, \delta^*(x)] \, \mathrm{d}\xi(w \mid x) \, \mathrm{d}f(x) = \int_S \left\{ \min_d \int_\Omega \ell[w, d] \, \mathrm{d}\xi(w \mid x) \right\} \mathrm{d}f(x).$$

Bayes decision functions

Extensive form

$$\sigma(\xi, \delta) = \int_\Omega \int_S \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega) \, \mathrm{d}\psi_\omega(x) \tag{5.4}$$

$$= \int_S \int_\Omega \ell[\omega, \delta(x)] \, \mathrm{d}\xi(\omega \mid x) \, \mathrm{d}f(x), \tag{5.5}$$

where $f(x) = \int_\Omega \psi_\omega(x) \, \mathrm{d}\xi(\omega)$.

Definition (Prior distribution)

The distribution ξ is called the prior distribution of ω.

Definition (Marginal distribution)

The distribution f is called the (prior) marginal distribution of x.

Definition (Posterior distribution)

The conditional distribution $\xi(\cdot \mid x)$ is called the posterior distribution of ω.

Minimax worlds with observations

Consider a utility function $U : \Omega \times D \to \mathbb{R}$. There are two players, the statistician and nature, each selecting $d \in D$ and $\omega \in \Omega$ respectively. The statistician's maximin decision without observations is:

$$\max_{d \in D} \min_{\omega \in \Omega} \mathbb{E}(U \mid \omega, d) = \max_{d \in D} \min_{\omega \in \Omega} U(\omega, d).$$

Now consider an observation $x \in S$, with $x \sim \psi(\cdot \mid \omega)$. The statistician now selects a decision function $\delta \in \Delta$. For any δ, the worst-case expected utility is:

$$\min_{\omega \in \Omega} \mathbb{E}(U \mid \omega, \delta) = \min_{\omega \in \Omega} \int_S U[\omega, \delta(x)] \, \mathrm{d}\phi_\omega(x) \tag{5.6}$$

$$= \min_{\omega \in \Omega} \sum_{d \in D} U(\omega, d) \phi_\omega \left(\{x \in S \mid \delta(x) = d\} \right). \tag{5.7}$$

Minimax priors with observations

The maximin problem

$$\max_{\delta \in \Delta} \min_{\xi \in \Xi} \mathbb{E}(U \mid \xi, \delta) = \max_{\delta \in \Delta} \min_{\xi \in \Xi} \int_S \int_\Omega U[\omega, \delta(x)] \, \mathrm{d}\xi(\omega \mid x) \, \mathrm{d}p_\xi(x). \tag{5.8}$$

The minimax problem

$$\min_{\xi \in \Xi} \max_{\delta \in \Delta} \mathbb{E}(U \mid \xi, \delta) = \min_{\xi \in \Xi} \int_S \max_{d \in D} \int_\Omega U[\omega, d] \, \mathrm{d}\xi(\omega \mid x) \, \mathrm{d}p_\xi(x). \tag{5.9}$$

Lemma

If Ξ contains all priors, then

$$\inf_{\xi \in \Xi} U(\xi, \delta) = \inf_{\omega \in \Omega} U(\omega, \delta) \tag{5.10}$$

Decision problems with two points and hypothesis testing

	d_1	d_2
ω_1	0	c_1
ω_2	c_2	0

Table: Cost function of a simple hypothesis testing problem

- We observe the value of some random variable X and then choose decision $\delta(X)$.
- Let $\alpha(\delta)$ be the conditional probability that we choose d_2 when $\omega = \omega_1$.
- Let $\beta(\delta)$ be the conditional probability that we choose d_1 when $\omega = \omega_2$.
- Let $a \triangleq c_1 \, \mathbb{P}(\omega = \omega_1)$ and $b \triangleq c_2 \, \mathbb{P}(\omega = \omega_2)$.

The risk of δ is:

$$a\alpha(\delta) + b\beta(\delta) \tag{5.11}$$

Decision problems with two points and hypothesis testing

	d_1	d_2
ω_1	0	c_1
ω_2	c_2	0

Table: Cost function of a simple hypothesis testing problem

The risk of δ is:

$$a\alpha(\delta) + b\beta(\delta) \tag{5.11}$$

Theorem (Neymann-Pearson lemma)

Let where ψ_w be densities or probabilities on $\mathcal{S}$. For any $a > 0$, $b > 0$, let δ^* be a decision function such that,

$$\delta^*(x) = d_1, \qquad \text{if } a\psi_{\omega_1}(x) > b\psi_{\omega_2}(x) \tag{5.12}$$

$$\delta^*(x) = d_2, \qquad \text{if } a\psi_{\omega_2}(x) < b\psi_{\omega_2}(x), \tag{5.13}$$

and either d_1, d_2 otherwise. Then, for any other δ:

$$a\alpha(\delta^*) + b\beta(\delta^*) \leq a\alpha(\delta) + b\beta(\delta)$$

Posterior distributions for multiple observations

Assume that we observe a value $x^n \triangleq x_1, \ldots, x_n$ of a random variable $X^n \triangleq X_1, \ldots, X_n$. We have a prior ξ on Ω. For the observations, we write:

Observation probability given history x^{n-1} and parameter ω

$$\psi(x_n \mid x^{n-1}, \omega) = \frac{\psi_\omega(x^n)}{\psi_\omega(x^{n-1})}$$

Posterior recursion

$$\xi(\omega \mid x^n) = \frac{\psi_\omega(x^n)\xi(\omega)}{f(x^n)} = \frac{\xi(x_n \mid x^{n-1}, \omega)\xi(\omega \mid x^{n-1})}{f(x_n \mid x^{n-1})}. \tag{5.14}$$

Posterior distributions for multiple independent observations

If $\psi(x_n \mid \omega, x^{n-1}) = \psi_\omega(x_n)$ then $\psi_\omega(x^n) = \prod_{k=1}^n \psi_\omega(x_k)$. Then:

Posterior recursion with conditional independence

$$\xi_n(w) \triangleq \xi_0(\omega \mid x^n) = \frac{\psi_\omega(x^n)\xi_0(w)}{f_0(x_n)} \tag{5.15}$$

$$= \xi_{n-1}(\omega \mid x_n) = \frac{\psi_\omega(x_n)\xi_{n-1}(w)}{f_{n-1}(x_n)} \tag{5.16}$$

where we define ξ_t to be the belief at time t.

Conditional independence allows us to write the posterior update as an identical recursion at each time t.

Observation cost

Expected cost of observation

Let $c : \mathcal{S} \times \Omega \to \mathbb{R}$ be an observation cost function. Then the expected cost is

$$\mathbb{E}_\xi[c(x, \omega)] = \int_\Omega \int_S c(\omega, x) \, \mathrm{d}\psi_\omega(x) \, \mathrm{d}\xi(\omega). \tag{5.17}$$

The total risk of observing x and using a decision function δ

is then given by

$$\sigma(\xi, \delta) + \mathbb{E}_\xi[c(\omega, x)]$$

Fixed cost per observation

- Consider that we can choose the size n of a sample $x_1, \ldots, x_n$.
- The cost of the sample of size n is γn.
- Let δ_n be the (random) Bayes decision function after observing $x_1, \ldots, x_n$:

$$\delta_n \triangleq \arg\min_{d \in D} \sigma[\xi(\cdot \mid x_1, \ldots, x_n), d] \tag{5.18}$$

- Thus, the Bayes risk of n observations is

$$\sigma_t(\xi, \delta_n) = \sigma(\xi, \delta_n) + \gamma n. \tag{5.19}$$

- Now we have another decision problem: How many observations to take?

Exercise

Prove that if the risk is bounded, then there exists an optimal number n of observations.

Quick summary

- We want to make a decision against an unknown parameter W.
- The risk is the negative expected utility.
- The Bayes risk is the minimum risk, and it is concave with respect to the distribution of W.
- Our decisions can depend on observations, via a decision function.
- We can construct a complete decision function by computing $\sigma(\xi, \delta)$ for all decision functions (normal form).
- We can instead wait until we observe x and compute $\sigma[\xi(\cdot \mid x), d]$ for all decisions (extensive form).
- In minimax settings, we can consider a fixed but unknown parameter w or a fixed but unknown prior ξ. This links decision theory to game theory.
- When each observation has cost γ, there is an optimal value n of minimising $\sigma[\xi(\cdot \mid X_n), \delta_n] + \gamma n$, where δ_n is the Bayes decision function after n observations.
- The posterior given multiple observations can be computed recursively using independence.
- Our decision at a certain time, affects the future information available.
- Problems where future decisions must be considered, require planning ahead and are called sequential decision problems.