Brief Introduction to Decision Theory

Historically, Bayesian statistics has been linked closely with *decision theory*. There are at least three important principles for choosing a decision rule, one of which is *Bayes principle*.

Basic elements of statistical decision problem

1. A parameter space Θ representing all the possible values of an unknown parameter $\boldsymbol{\theta}$.

2. A set $\mathcal{A}$ of actions or decisions that are available to the statistician.

3. A real-valued loss function, $L(\boldsymbol{\theta}, a)$, defined on $\Theta \times \mathcal{A}$.

4. Statistician does an experiment to obtain information about $\boldsymbol{\theta}$. The information is contained in a random vector $\boldsymbol{Y}$, which has distribution $f(\boldsymbol{y}|\boldsymbol{\theta})$ when $\boldsymbol{\theta}$ is the true value of the parameter.

Usually it's assumed that $L(\boldsymbol{\theta}, a) \geq -K > -\infty$, where a negative loss is a gain to the statistician.

Utility

From our point of view, utility is just negative loss.

Utility theory provides a framework for making a wise choice of loss function.

Utility is concerned with the *value* associated with the consequence of an action.

In the decision problem, suppose there is a reward, r, associated with each $(\theta, a) \in \Theta \times \mathcal{A}$.

The utility function U is a real-valued function defined over the set $\mathcal{R}$ of all possible rewards. Given U, our loss function is $L(\theta, a) = -U(r)$, where r is the reward associated with (θ, a).

Determining one's utility function for money

Suppose the possible rewards in a business venture are

$$r_1 = -100, \quad r_2 = 0, \quad r_3 = 200, \quad r_4 = 500.$$

We assume that r_4 is preferred to r_3, which is preferred to r_2, which is preferred to r_1. Then we let

$$U(0) = 0 \quad \text{and} \quad U(500) = 1.$$

Now, consider a game where you get reward r_2 with probability α and you get reward r_4 with probability $1 - \alpha$.

What would α have to be in order for this game to be equivalent to a "game" where you are simply given r_3?

If we can determine α by introspection, then

$$U(200) = \alpha U(0) + (1 - \alpha)U(500).$$

Suppose we decide that $\alpha = 0.10$. This means that a sure 200 bucks is considered to be equivalent to a 90% chance at winning 500 bucks and a 10% chance of winning nothing.

In this case

$$U(200) = 0.90.$$

Now determine α so that being assured of r_2 (0 dollars) is equivalent to losing 100 bucks with probability α and winning 200 bucks with probability $1-\alpha$.

If $\alpha = 0.30$, then

$$U(0) = 0.30U(-100) + 0.70U(200),$$

implying that

$$U(-100) = -0.63/0.30 = -2.1.$$

Reference for utility theory:

Blackwell, D. and Girshick, M.A. *Theory of Games and Statistical Decisions*, Chapter 4

Example 1 *A statistical decision problem*

1. "Nature" chooses θ, the probability of some event, and $\Theta = \{\theta : 0 \leq \theta \leq 1\}$.

2. Statistician estimates θ by a, with $\mathcal{A} = \Theta$.

3. Statistician chooses the loss function $L(\theta, a) = (\theta - a)^2$.

4. Statistician performs a binomial experiment and observes $Y \sim \text{bin}(n, \theta)$.

As in the Bayesian statistical model discussed previously, the statistician might also make use of *prior information* about $\boldsymbol{\theta}$. This prior info may be in the form of a probability distribution π for $\boldsymbol{\theta}$.

Decision rules

Let $\mathcal{Y}$ be the sample space, i.e., the set of all possible data vectors $\boldsymbol{y}$.

Definition A nonrandomized decision rule δ is a function from $\mathcal{Y}$ into $\mathcal{A}$, with the interpretation that $\delta(\boldsymbol{y})$ is the action taken by the statistician when the observed data are $\boldsymbol{y}$.

Example 2 An unknown probability density function belongs to the set of densities $\{f_1, f_2, \ldots, f_k\}$, where $k \geq 2$.

$\Theta = \{1, 2, \ldots, k\}$, where $\theta = i$ corresponds to f_i being the true density. The action space is $\mathcal{A} = \Theta$, and let the loss function be

$$L(i, j) = \int \Big(f_i(t) - f_j(t)\Big)^2 \, dt.$$

The statistician observes n independent observations, $\boldsymbol{Y} = (Y_1, \ldots, Y_n)$, from the true density. Hence,

$$f(\boldsymbol{y}|\theta) = \prod_{i=1}^{n} f_\theta(y_i).$$

A possible decision rule is as follows:

$$\delta(\boldsymbol{y}) = \text{the value of } \theta \text{ that maximizes } f(\boldsymbol{y}|\theta)$$

When the statistician uses a decision rule δ and observes data $\boldsymbol{y}$, the loss at $\boldsymbol{\theta}$ is $L(\boldsymbol{\theta}, \delta(\boldsymbol{y}))$. Obviously this will vary with $\boldsymbol{y}$.

Different decision rules may be compared by *averaging* the loss function with respect to $\boldsymbol{Y}$.

<u>*Definition*</u> The *risk function* of the decision rule δ is

$$R(\boldsymbol{\theta}, \delta) = E_{\boldsymbol{\theta}} \left[L(\boldsymbol{\theta}, \delta(\boldsymbol{Y})) \right].$$

Suppose, for example, that $\boldsymbol{Y}$ has density $f(\boldsymbol{y}|\boldsymbol{\theta})$. Then

$$R(\boldsymbol{\theta}, \delta) = \int L(\boldsymbol{\theta}, \delta(\boldsymbol{y})) f(\boldsymbol{y}|\boldsymbol{\theta}) \, d\boldsymbol{y}.$$

A rule δ_1 is said to be R-*better* than some other rule δ_2 if $R(\boldsymbol{\theta}, \delta_1) \leq R(\boldsymbol{\theta}, \delta_2)$ for all $\boldsymbol{\theta} \in \Theta$ with strict inequality for some $\boldsymbol{\theta}$.

Definition Let Δ be a collection of decision rules. Then $\delta \in \Delta$ is said to be *admissible* in Δ if there exists no R-better rule in Δ. A rule $\delta \in \Delta$ is *inadmissible* if there exists an R-better rule in Δ.

Example 3 Let $\Theta = [0,1]$ and suppose $\Delta = \{\delta_1, \delta_2, \delta_3\}$. The risk functions for the three rules are as follows:

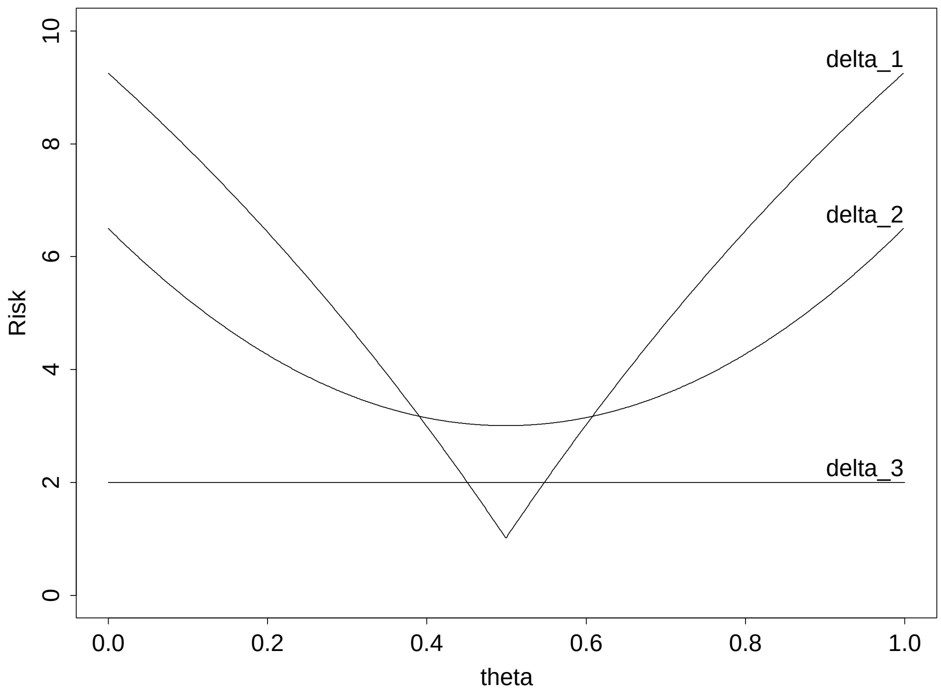

The rule δ_2 is inadmissible since $R(\theta, \delta_3) < R(\theta, \delta_2)$ for each $\theta \in \Theta$. We can discard δ_2 with the assurance that there is a uniformly better rule (at least in terms of risk).

Note that δ_1 is admissible, but not particularly good. It has smaller risk than δ_3 only in a relatively small set of θ values.

<u>Example 4</u> Suppose $\boldsymbol{Y} = (Y_1, Y_2)$, where Y_1 and Y_2 are independent and identically distributed (i.i.d.), each having density

$$f_\theta(y) = \frac{1}{\theta} e^{-y/\theta} I_{(0,\infty)}(y).$$

$$\Theta = \{\theta : \theta > 0\},\ \mathcal{A} = \Theta$$

$$L(\theta, a) = \frac{(\theta - a)^2}{\theta}$$

Let X_1 and X_2 denote the smaller and larger of Y_1 and Y_2, respectively, and consider the following class of decision rules:

$$\delta_p(\boldsymbol{y}) = px_1 + (1-p)x_2, \quad 0 \le p \le 1.$$

$$\begin{aligned} R(\theta, \delta_p) &= E_\theta \left[\frac{1}{\theta} (\delta_p(\boldsymbol{Y}) - \theta)^2 \right] \\ &= \frac{1}{\theta} \left[\text{Var}_\theta (\delta_p(\boldsymbol{Y})) + (E_\theta(\delta_p(\boldsymbol{Y})) - \theta)^2 \right] \\ &= \frac{\theta}{2}(4p^2 - 6p + 3). \end{aligned}$$

The only admissible rule in the class is $\delta_{3/4}$, since

$$R(\theta, \delta_{3/4}) < R(\theta, \delta_p), \quad \forall\ \theta > 0$$

and any $p \in [0,1]$ such that $p \neq 3/4$. In particular, note that $\delta_{1/2}(\boldsymbol{y}) = (y_1 + y_2)/2$ is inadmissible.

Decision principles

How do we choose a decision rule? There usually exists no rule that is R-better than every other rule.

In decision theory, there are three main principles for choosing a decision rule:

I. Bayes principle

II. Minimax principle

III. Invariance principle

I. The Bayes principle requires a prior distribution π.

<u>*Definition*</u> The *Bayes risk* of δ with respect to π is

$$r(\pi, \delta) = E_\pi[R(\boldsymbol{\theta}, \delta)].$$

According to the Bayes principle, δ_1 is preferred to δ_2 if

$$r(\pi, \delta_1) < r(\pi, \delta_2).$$

If there exists a δ^π which minimizes $r(\pi, \delta)$ over all decision rules, then δ^π is called a *Bayes rule*. The number $r(\pi) = r(\pi, \delta^\pi)$ is called the *Bayes risk of* π.

II. The minimax principle says that rule δ_1 is preferred to rule δ_2 if

$$\sup_{\boldsymbol{\theta}} R(\boldsymbol{\theta}, \delta_1) < \sup_{\boldsymbol{\theta}} R(\boldsymbol{\theta}, \delta_2).$$

The rule δ_M is a *minimax rule* if

$$\sup_{\boldsymbol{\theta}} R(\boldsymbol{\theta}, \delta_M) = \inf_{\delta} \sup_{\boldsymbol{\theta}} R(\boldsymbol{\theta}, \delta) = MV.$$

MV is called the *minimax value* for the problem.

III. The invariance principle says that attention should be restricted to a class of decision rules that do not depend on the unit of measurement used, or any other such arbitrary aspects of the problem. By making the class of rules from which to choose smaller, it is easier to find a best rule.

One example of "restricting the class of rules" is in the problem of estimating the mean of a probability distribution. If i.i.d. measurements $Y_1, \ldots, Y_n$ are available, it seems natural to restrict attention to rules satisfying

$$\delta(y_1 + c, \ldots, y_n + c) = \delta(y_1, \ldots, y_n) + c.$$