Statistical Decision Theory and Bayesian Analysis

CH2. UTILITY AND LOSS

§ 2.2 Utility Theory

- Reward $r \in \mathcal{R}$ "a consequence"

  P: a probability distribution over r caused by an action a.

  expected utility

  $U(r)$: a utility function (satisfying $E^P[U(r)]$ exists)

- "Rationality axioms"

  P_2 is preferred | P_1, P_2 equivalent

Axiom 1. If $P_1, P_2 \in \mathcal{P}$, then either $P_1 \prec P_2$, $P_1 \approx P_2$, $P_2 \prec P_1$

Axiom 2. If $P_1 \precsim P_2$ and $P_2 \precsim P_3$, then $P_1 \precsim P_3$

Axiom 3. If $P_1 \prec P_2$ then $\alpha P_1 + (1-\alpha) P_3 \prec \alpha P_2 + (1-\alpha) P_3$

$\forall\ 0 < \alpha < 1$ and $P_3 \in \mathcal{P}$

conditionality principle

Axiom 4. If $P_1 \prec P_2 \prec P_3$, there are numbers $0 < \alpha < 1$, and $0 < \beta < 1$ such that

$$\alpha P_1 + (1-\alpha) P_3 \prec P_2$$

$$\text{and } P_2 \prec \beta P_1 + (1-\beta) P_3$$

"No heaven or hell"

If P_1 is infinitely bad (i.e. $E^{P_1}[U(r)] = -\infty$)

$\Rightarrow \nexists\ \beta > 0$ s.t. $P_2 \prec \beta P_1 + (1-\beta) P_3$.

That is to say, U must be bounded

Axioms $1 \sim 5$ $\Rightarrow$ a unique utility function (for a particular scale)

However, determining U is difficult.

- Values of consequences may not have obvious scale (e.g. prestige, reputation, etc)
- Utility $\not\propto$ "True value"

- Construction of U

<u>STEP 1:</u> Consider $r_1 \not\approx r_2$. Assume $r_1 \prec r_2$. Then let

$$U(r_1)=0 \qquad U(r_2)=1$$

Choose r_2 as the best reward and r_1 worst for convenience. But any choice is acceptable

<u>STEP 2.</u> For r_3 s.t. $r_1 \prec r_3 \prec r_2$ find the α $(0<\alpha<1)$ s.t.

one-point distribution

$$r_3 \approx p \triangleq \alpha\langle r_1\rangle + (1-\alpha)\langle r_2\rangle$$

$$U(r_3) = E^P[U(r)] = \alpha\, U(r_1) + (1-\alpha)\, U(r_2) = 1-\alpha$$

<u>STEP 3.</u> For r_3 s.t. $r_3 \prec r_1 \prec r_2$. find the α $(0<\alpha<1)$. s.t.

$$r_1 \approx P \triangleq \alpha\langle r_3\rangle + (1-\alpha)\langle r_2\rangle$$

$$U(r_3) = -\frac{(1-\alpha)}{\alpha}$$

<u>STEP 4:</u> For r_3 s.t. $r_1 \prec r_2 \prec r_3$. find the α $(0<\alpha<1)$ s.t.

$$r_2 \approx P \triangleq \alpha\langle r_1\rangle + (1-\alpha)\langle r_3\rangle$$

Then

$$1 = U(r_2) = E^P[U(r)] = \alpha\, U(r_1) + (1-\alpha)\, U(r_3) = (1-\alpha)\, U(r_3)$$

$$\Rightarrow \quad U(r_3) = \frac{1}{1-\alpha}$$

<u>STEP 5:</u> Periodically check the construction process for consistency by comparing new combinations of rewards.

Assume $r_3 \prec r_4 \prec r_5$ found by STEP 1~4. Then find α.

$$r_4 \approx P \triangleq \alpha\langle r_3\rangle + (1-\alpha)\, U(r_5)$$

We find define α s.t. $U(r_4) = \alpha U(r_3) + (1-\alpha) U(r_5)$. But that may not reflect real preference over r_3, r_4, r_5.

make sure

$$U(r_4) = \alpha\, U(r_3) + (1-\alpha)\, U(r_5)$$

Notes:
1. This process of comparing and recomparing is often how the best judgement can be made.
2. People do not intuitively tend to act in accordance with a utility function. Thus we are, in essence, defining rational behavior for an individual, and suggesting that such behavior is good.

Example

	a_1: football	a_2: movie
rain (θ_1)	$r_1 = (a_1, \theta_1)$ 0	$r_3 = (a_2, \theta_1)$ 0.75
no rain (θ_2)	$r_2 = (a_1, \theta_2)$ 1	$r_4 = (a_2, \theta_2)$ 0.6

$\pi(\theta_1) = 0.4$

$\pi(\theta_2) = 0.6$

Determine r_1, r_2, r_3, r_4 and take an action.

- It is obvious that $r_1 \prec r_4 \prec r_3 \prec r_2$ assume we prefer football

- We assign $U(r_1) = 0$ $U(r_2) = 1$

- Gamble r_4 against $\alpha \langle r_1 \rangle + (1-\alpha) \langle r_2 \rangle$ until we feel equally happy.

  After some soul searching, we decide $\alpha = 0.4$. i.e.

  $r_4 \approx 0.4 \langle r_1 \rangle + (1-\alpha) \langle r_2 \rangle$

  $U(r_4) = 0.6$

- Gamble r_3 againt $\alpha \langle r_1 \rangle + (1-\alpha) \langle r_2 \rangle$, we decide $\alpha = 0.3$

  $r_3 \approx 0.3 \langle r_1 \rangle + 0.7 \langle r_2 \rangle$. $U(r_3) = 0.7$

- Check consistency:

  $r_3 \approx \alpha \langle r_4 \rangle + (1-\alpha) \langle r_2 \rangle$, what is α?

  Then we feel α should be 0.6

  $0.6 = U(r_3) = \alpha U(r_4) + (1-\alpha) U(r_2) = 0.6 \cdot 0.6 + (1-0.6) \cdot 1 = 0.76$

  ↯ contradiction

- Re-examine by introspection again until utilities are consistent.

  Say $U(r_4) = 0.6$ $U(r_3) = 0.75$

- Watching foot ball $E^{\pi, a_1}[U(r)] = \pi(\theta_1) \cdot U(r_1) + \pi(\theta_2) \cdot U(r_2) = 0.4 \cdot 0 + 0.6 \cdot 1 = 0.6$

  Watching a movie $E^{\pi, a_2}[U(r)] = \pi(\theta_1) \cdot U(r_3) + \pi(\theta_2) \cdot U(r_4) = 0.4 \cdot 0.75 + 0.6 \cdot 0.6 = 0.66$

  ⇒ The optimal action is go to the movie.

- Dealing with large $\mathcal{R}$
  - Estimate at a few points and then interpolate
  - High dimensional $\mathcal{R}$: regression.

E.g. $U(r) = \sum_{i=1}^{m} K_i U_i(r_i)$ (Naive Bayes assumption)

$U(r) = \sum_{i=1}^{m} K_i U_i(r_i) \sum_{i=1}^{m} \sum_{j=1}^{m} K_{ij} U_i(r_i) U_j(r_j)$

(Two side effects of a drug might be acceptable separately, but very dangerous if they occur together)

§ 2.3

- The Utility of Money

"backward binary construction"

1° Choose $r_1 < r_2$ and set $U(r_1)=0$ and $U(r_2)=1$

2° Find $r_3 \approx \frac{1}{2} U\langle r_1\rangle + \frac{1}{2} U\langle r_2 \rangle$, indicating

$U(r_3) = \frac{1}{2} U(r_1) + \frac{1}{2} U(r_2) = \frac{1}{2}$.

3° Find $r_4 \approx \frac{1}{2}\langle r_1\rangle + \frac{1}{2}\langle r_3\rangle$ and $r_5 \approx \frac{1}{2}\langle r_3\rangle + \frac{1}{2}\langle r_2\rangle$

$U(r_4) = \frac{1}{4}$ $\quad U(r_5) = \frac{3}{4}$

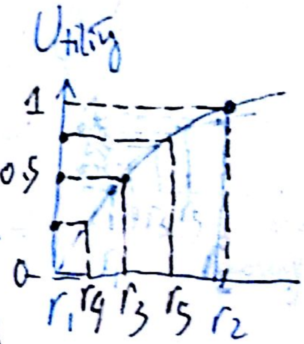

4° Is $r_3 = \frac{1}{2}\langle r_4\rangle + \frac{1}{2}\langle r_5\rangle$ acceptable?

Petersburg paradox

Gamble at cost c. A fair coin is flipped until a tail first appears.

Reward: $\$2^n$ where n is the number of flips it takes until a tail appears.

Expected gain (in money): $\sum_{n=1}^{\infty} 2^n P(n \text{ flips}) = \sum_{n=1}^{\infty} 2^n \cdot \frac{1}{2^n} = \infty$

Expected utility (for change): $\sum_{n=1}^{\infty} U(2^n - c)\cdot 2^{-n}$

Either positive or negative depending on c.

"Quasi-utilities"

§ 2.4 The Loss Function

- Loss functions should be defined on utilities. However, utilities are usually not linear

<u>Example:</u> A coin is flipped n times.

$p(\text{Head}) = 0.6 \Rightarrow$ win \$10,000

$p(\text{Tail}) = 0.4 \Rightarrow$ Lose \$1,000

Let Z_i denot the amount won <u>on</u> the i-th flip.

Assume $U(r) = \begin{cases} r^{1/3} & \text{if } r \geq 0 \\ 2r^{2/3} & \text{if } r < 0 \end{cases}$

$$E[U(Z_i)] = 0.6 \cdot U(1000) + 0.4\, U(-1000)$$
$$= 0.6 \cdot 10 + 0.4 \cdot (-20)$$
$$= -2$$

However $\sum_{i=1}^{n} E[U(Z_i)] \neq -2n$ if $i \neq 1$.

In fact the true outcome $Z = \sum_{i=1}^{n} Z_i$ can be positive

- Standard Loss Function: <u>Squared-Error Loss</u>

✗ Cons: Unbounded, convex.

✓ Pros: Assume $Z = h(\theta - a, Y)$ — an estimator of θ; Some randomness other than θ. (also independent of θ)

$g(\theta - a, Y) \triangleq U(h(\theta - a, Y))$, assume $\theta - a$ is small

Tayler expansion, derivatives wrt. $(\theta - a)$

$$g(\theta - a, Y) \simeq g(0, Y) + (\theta - a)\, g'(0, Y) + \frac{1}{2}(\theta - a)^2 g''(0, Y)$$

Y is the only thing random in Z.

Denote $K_1 = -E^Y[g(0, Y)]$, $K_2 = -E[g'(0, Y)]$ and $K_3 = \frac{1}{2} E[g''(0, Y)]$

Then $L(\theta, a)$

$$= -E[U(Z)] \simeq K_1 + K_2(\theta - a) + K_3(\theta - a)^2$$

✓ Another merit: For an estimator $R(\theta, \hat{\theta}) = E_\theta^X L(\theta, \hat{\theta}) = E_\theta^X[\theta - \hat{\theta}(X)]^2$. R happens to be the variance of $\hat{\theta}$, if $\hat{\theta}$ unbiased

$$\Rightarrow \quad L(\theta-a) \simeq k_3\left(\theta-a+\frac{k_2}{2k_3}\right)^2+\left(k_1-\frac{k_2^2}{4k_3}\right)$$

Provided $k_3 > 0$, $k_2 = 0$ (if $g(0,y)$ symmetric in y and Y is symmetric distribution)

and $\theta \approx a$, ~~then~~ the loss is then squared error.

Variants: – $L(\theta, a) = (\theta - a + c)^2$

– Weighted squared error.

$$L(\theta, a) = w(\theta)\cdot(\theta-a)^2$$

– Quadratic loss

$$L(\theta, a) = (\theta-a)^T Q \, (\theta-a)$$

If diganal

$$L(\theta, a) = \sum_{i=1}^{p} q_i(\theta_i - a_i)^2$$

- Standard Loss Function: Linear Loss

$$L(\theta, a) = \begin{cases} k_0(\theta-a) & \text{if } \theta - a \geq 0 \\ k_1(a-\theta) & \text{if } \theta - a < 0 \end{cases}$$

or $L(\theta, a) = |\theta - a|$

or weighted linear loss

- Standard Loss Function: "0-1" Loss

$$L(\theta, a_i) = \begin{cases} 0, & \text{if } \theta \in \Theta_i \\ 1, & \text{if } \theta \in \Theta_j \quad j \neq i \end{cases}$$

Risk: $R(\theta, \delta) = E_\theta^X[L(\theta, \delta(X))] = P_\theta\{\delta(X) \text{ incorrect}\}$ Either type I & II error, depending on $\theta \in \Theta_0$ or $\theta \in \Theta_1$

Bayesian expected loss: $\rho(\pi^*, a_i) = \int L(\theta, a_i)\, dF^{\pi^*}(\theta) = 1 - P^{\pi^*}(\theta \in \Theta_i)$

More realistic approximations

$$L(\theta, a_i) = \begin{cases} 0, & \text{if } \theta \in \Theta_i \\ K_i, & \text{if } \theta \in \Theta_j \quad (i \neq j) \end{cases} \quad \text{or} \quad L(\theta, a_i) = \begin{cases} 0 & \text{if } \theta \in \Theta_i \\ K_i(\theta) & \text{if } \theta \in \Theta_j \ (i \neq j) \end{cases}$$

- For Inference Problems

<u>Example.</u> Let C denotes a confidence rule, $C(x)$ being a confidence set.

Define $L(\theta, C(x)) = 1 - I_{C(x)}(\theta) = \begin{cases} 1 & \text{if } \theta \notin C(x) \\ 0 & \text{if } \theta \in C(x) \end{cases}$

Then $R(\theta, C) = E_\theta[1 - I_{C(X)}(\theta)] = 1 - P_\theta\{C(X) \text{ contains } \theta\}$

$$\rho(\pi^*, C(x)) = E^{\pi^*}[1 - I_{C(x)}(\theta)] = 1 - P^{\pi^*}(\theta \in C(x))$$

<u>Example:</u> Measuring the "communication quantity"

Consider the example on p. 11. (SDT/BA-CH1).

where $X = (X_1, X_2)$ X_i iid. $P_\theta(X_i = \theta - 1) = P_\theta(X_i = \theta + 1) = \frac{1}{2}$

Frequentist confidence $\alpha_1(x) \equiv .75$

Conditionalist confidence $\alpha_2(x) = \begin{cases} 1 & \text{if } x_1 \neq x_2 \\ 0.5 & \text{if } x_1 = x_2 \end{cases}$

Define $L_c(\theta, \alpha(x)) = (I_{C(x)}(\theta) - \alpha(x))^2$

$$R_c(\theta, \alpha_1) = E_\theta^X L_c(\theta, \alpha_1(X)) = \frac{3}{16}$$

$$R_c(\theta, \alpha_2) = E_\theta^X L_c(\theta, \alpha_2(X)) = \frac{1}{8}$$

- For Predictive Problems

To predict $Z \sim g(z|\theta)$

Then $L(\theta, a) = E_\theta^Z[L^*(Z,a)] = \int L^*(z,a) g(z|\theta) dz$

Example: $Z \sim \mathcal{N}(\theta, \sigma^2)$ Let $L^*(z,a) = (z-a)^2$

$$L(\theta,a) = E_\theta^Z[Z-a]^2 = E_\theta^Z[Z-\theta+\theta-a]^2$$

$$= E_\theta^Z[Z-\theta]^2 + E_\theta^Z[\theta-a]^2 = \sigma^2 + (\theta-a)^2$$

$\Rightarrow$ Working with $L(\theta,a)$ $\Leftrightarrow$ working with squared-error loss for θ.