Fundamental Theory of Statistical Inference

G. Alastair Young

Department of Mathematics
Imperial College London

LTCC, 2017

Fundamental characteristic

Explicit optimality criteria.

- Hypothesis testing: seek test which maximises power.
- Point estimation: seek estimator which minimises risk.

Formulation, hypothesis testing

We have parameter space Ω_θ, and consider hypotheses of the form

$$H_0 : \theta \in \Theta_0 \quad \text{vs.} \quad H_1 : \theta \in \Theta_1$$

where Θ_0 and Θ_1 are two disjoint subsets of Ω_θ, possibly satisfying $\Theta_0 \cup \Theta_1 = \Omega_\theta$.

Simple/composite hypotheses

If a hypothesis consists of a single member of Ω_θ, $H_0 : \theta = \theta_0$, then it is a simple hypothesis. Otherwise it is composite.

Nuisance parameters

Beware of nuisance parameters!

$Y_1, \ldots, Y_n$ IID $N(\mu, \sigma^2)$, μ and σ^2 unknown. $H_0 : \mu = 0$ is composite, because of nuisance parameter σ^2.

Classical approach

Adopt the following criterion: fix a small probability α (known as the size) and seek a test for which

$$\mathbb{P}_\theta\{\text{Reject } H_0\} \le \alpha \quad \text{for all } \theta \in \Theta_0. \qquad (\dagger)$$

H_0 and H_1 are treated asymmetrically. Usually H_0 is called the null hypothesis and H_1 the alternative hypothesis.

Test functions

Conventional formulation: choose a test statistic $t(Y)$ with distribution depending on θ and a critical region C_α, reject H_0 based on $Y = y$ iff $t(y) \in C_\alpha$. Critical region chosen to satisfy (†).

Slight reformulation: define the test function $\phi(y)$ by

$$\phi(y) = \begin{cases} 1 & \text{if } t(y) \in C_\alpha, \\ 0 & \text{otherwise.} \end{cases}$$

If we observe $\phi(y) = 1$, we reject H_0, while if $\phi(y) = 0$, we accept.

Randomised Tests

To cope with the case when $t(Y)$ has a discrete distribution, generalise concept of a test function to allow $\phi(y)$ to take on any value in the interval $[0, 1]$.

Having observed data y and evaluated $\phi(y)$, we use some independent randomisation device to draw a random number $W \in \{0, 1\}$ which takes value 1 with probability $\phi(y)$ and 0 otherwise. We then reject H_0 if and only if $W = 1$.

We interpret $\phi(y)$ to be the probability that H_0 is rejected given $Y = y$.

Power

Criterion for deciding whether one test is better than another: power.

The power function of a test ϕ is defined to be

$$w(\theta) = \mathbb{P}_\theta\{\text{Reject } H_0\} = \mathbb{E}_\theta\{\phi(Y)\},$$

defined for all $\theta \in \Omega_\theta$.

Idea

A good test of a given size α is one which makes $w(\theta)$ as large as possible for $\theta \in \Theta_1$ while satisfying the constraint $w(\theta) \leq \alpha$ for all $\theta \in \Theta_0$.

Classes of problem

- (i) Simple H_0 vs. simple H_1. Complete theory, given by the Neyman-Pearson Theorem.
- (ii) Simple H_0 vs. composite H_1. In some cases, but not all, there is a uniformly most powerful test, with $w(\theta)$ largest over all tests, uniformly for all $\theta \in \Theta_1$. When the family of distributions has the property of Monotone Likelihood ratio.
- (iii) Composite H_0 vs. composite H_1. How do we cope with nuisance parameters? With certain distributions, we can use conditional tests.

The Neyman-Pearson Theorem

Test of simple null hypothesis $H_0: \ \theta = \theta_0$ against simple alternative hypothesis $H_1: \ \theta = \theta_1$, where θ_0 and θ_1 are specified.

Let the pdf of Y be $f(y;\theta)$ specialised to $f_0(y) = f(y;\theta_0)$ and $f_1(y) = f(y;\theta_1)$.

Define the likelihood ratio $\Lambda(y)$ by

$$\Lambda(y) = \frac{f_1(y)}{f_0(y)}.$$

Best test of size α in terms of power is of the form: reject H_0 when $\Lambda(Y) > k_\alpha$ where k_α is chosen to guarantee test has size α.

Randomised tests

In a generalised form of Neyman-Pearson Theorem, we allow for the possibility of randomised tests.

The (randomised) test with test function ϕ_0 is said to be a likelihood ratio test (LRT for short) if it is of the form

$$\phi_0(y) = \begin{cases} 1 & \text{if } f_1(y) > Kf_0(y), \\ \gamma(y) & \text{if } f_1(y) = Kf_0(y), \\ 0 & \text{if } f_1(y) < Kf_0(y), \end{cases}$$

where $K \geq 0$ is a constant and $\gamma(y)$ an arbitrary function satisfying $0 \leq \gamma(y) \leq 1$ for all y.

Theorem

(a) (Optimality). For any K and $\gamma(y)$, the test ϕ_0 has maximum power among all tests whose size is no greater than the size of ϕ_0.

(b) (Existence). Given α $(0 < \alpha < 1)$ there exist constants K and γ_0 such that the LRT defined by this K and $\gamma(y) = \gamma_0$ for all y has size exactly α.

(c) (Uniqueness). If the test ϕ has size α, and is of maximum power amongst all possible tests of size α, then ϕ is necessarily a LRT, except possibly on a set of values of y which has probability 0 under both H_0 and H_1.

Monotone Likelihood Ratio and UMP Tests

A uniformly most powerful or UMP test of size α is a test $\phi_0(\cdot)$ for which

- (i) $\mathbb{E}_\theta \phi_0(Y) \leq \alpha$ for all $\theta \in \Theta_0$;
- (ii) Given any other test $\phi(\cdot)$ for which $\mathbb{E}_\theta \phi(Y) \leq \alpha$ for all $\theta \in \Theta_0$, we have $\mathbb{E}_\theta \phi_0(Y) \geq \mathbb{E}_\theta \phi(Y)$ for all $\theta \in \Theta_1$.

Existence of UMP test

Asking that the Neyman-Pearson test for simple vs. simple hypotheses should be the same for every pair of simple hypotheses contained within H_0 and H_1.

For one-sided testing problems involving just a single parameter ($\Omega_\theta \subseteq \mathbb{R}$), there is a wide class of parametric families for which such a property holds. Such families are said to have monotone likelihood ratio or MLR.

Definition

The family of densities $\{f(y;\theta),\ \theta \in \Omega_\theta \subseteq \mathbb{R}\}$ with real scalar parameter θ is said to be of monotone likelihood ratio (MLR) if there exists a function $t(y)$ such that the likelihood ratio

$$\frac{f(y;\theta_2)}{f(y;\theta_1)}$$

is a non-decreasing function of $t(y)$ whenever $\theta_1 \leq \theta_2$.

Note that non-increasing as a function of $t(y) \equiv$ non-decreasing in $-t(y)$.

The main result

Suppose Y has a distribution from a family which is MLR with respect to a statistic $t(Y)$, and that we wish to test $H_0: \theta \leq \theta_0$ against $H_1: \theta > \theta_0$. Suppose the distribution function of $t(Y)$ is continuous.

(a) The test

$$\phi_0(y) = \begin{cases} 1 & \text{if } \; t(y) > t_0, \\ 0 & \text{if } \; t(y) < t_0, \end{cases}$$

is UMP among all tests of size $\leq \mathbb{E}_{\theta_0}\{\phi_0(Y)\}$.

(b) Given some α, where $0 < \alpha < 1$, there exists some t_0 such that the test in (a) has size exactly α.

Two-sided tests and conditional inference

Revisit hypothesis testing to consider:

- two-sided hypotheses of the form $H_0: \theta \in [\theta_1, \theta_2]$ (with $\theta_1 < \theta_2$) or $H_0: \theta = \theta_0$ with the alternative H_1 including all θ not part of H_0. Cannot find a UMP test of a given size α. Introduce concept of unbiasedness, define uniformly most powerful unbiased, or UMPU, tests. Focus on exponential families.
- Extension to multiparameter exponential families, using notion of conditional tests. Discuss two situations where conditional tests arise, when there are ancillary statistics, and where conditional procedures are used to construct similar tests.

Two-sided hypotheses and two-sided tests

Consider a one-dimensional parameter $\theta \in \Omega_\theta \subseteq \mathbb{R}$. Consider the case where the null hypothesis is $H_0 : \ \theta \in \Theta_0$ where Θ_0 is either the interval $[\theta_1, \theta_2]$ for some $\theta_1 < \theta_2$, or else the single point $\Theta_0 = \{\theta_0\}$, and $\Theta_1 = \Omega_\theta - \Theta_0$.

We cannot in general expect to find a UMP test. If we construct a best test of say $\theta = \theta_0$ against $\theta = \theta_1$ for some $\theta_1 \neq \theta_0$, the test takes quite a different form when $\theta_1 > \theta_0$ from when $\theta_1 < \theta_0$.

Two-sided tests

For an exponential or MLR family with natural statistic $S = s(Y)$, might expect tests of the form

$$\phi(y) = \begin{cases} 1 & \text{if } s(y) > t_2 \text{ or } s(y) < t_1, \\ \gamma(y) & \text{if } s(y) = t_2 \text{ or } s(y) = t_1, \\ 0 & \text{if } t_1 < s(y) < t_2, \end{cases}$$

where $t_1 < t_2$ and $0 \le \gamma(y) \le 1$, to have good properties.

Such tests are called two-sided tests.

Unbiasedness and UMPU tests

A test ϕ of $H_0: \theta \in \Theta_0$ against $H_1: \theta \in \Theta_1$ is unbiased of size α if

$$\sup_{\theta \in \Theta_0} \mathbb{E}_\theta\{\phi(Y)\} = \alpha$$

and

$$\mathbb{E}_\theta\{\phi(Y)\} \geq \alpha \text{ for all } \theta \in \Theta_1.$$

A test which is uniformly most powerful amongst the class of all unbiased tests is uniformly most powerful unbiased, abbreviated UMPU.

Discussion

Unbiasedness is not by itself an optimality criterion.

There is no reason why the optimal decision procedure should turn out to be unbiased.

The principal role of unbiasedness is to restrict the class of possible decision procedures and hence to make the problem of determining an optimal procedure more manageable.

UMPU tests for one-parameter exponential families

Consider an exponential family for a random variable Y, with real-valued parameter $\theta \in \mathbb{R}$ and density of form

$$f(y;\theta) = c(\theta)h(y)e^{\theta s(y)},$$

where $S = s(Y)$ is a real-valued natural statistic.

This implies that S also has an exponential family distribution, with density of form

$$f_S(s;\theta) = c(\theta)h_S(t)e^{\theta s}.$$

Assume that S is a continuous random variable with $h_S(s) > 0$ on the open set which defines the range of S, to avoid the need for randomised tests.

The set-up

Consider the case

$$\Theta_0 = [\theta_1, \theta_2], \quad \Theta_1 = (-\infty, \theta_1) \cup (\theta_2, \infty),$$

where $\theta_1 < \theta_2$.

Theorem

Let ϕ be any test function. Then there exists a unique two-sided test ϕ' which is a function of S such that

$$\mathbb{E}_{\theta_j}\phi'(Y) = \mathbb{E}_{\theta_j}\phi(Y), \quad j = 1, 2.$$

Also,

$$\mathbb{E}_{\theta}\phi'(Y) - \mathbb{E}_{\theta}\phi(Y) \begin{cases} \leq 0 & \text{for } \theta_1 < \theta < \theta_2, \\ \geq 0 & \text{for } \theta < \theta_1 \text{ or } \theta > \theta_2. \end{cases}$$

Corollary

For any $0 < \alpha < 1$, there exists a UMPU test of size α, which is of two-sided form in S.

Testing a point null hypothesis

Consider the case $H_0: \ \theta = \theta_0$ against $H_1: \ \theta \neq \theta_0$ for a given value of θ_0. By the previous case, letting $\theta_2 - \theta_1 \to 0$, there exists a two-sided test ϕ' for which

$$\mathbb{E}_{\theta_0}\{\phi'(Y)\} = \alpha, \quad \frac{d}{d\theta}\mathbb{E}_{\theta}\{\phi'(Y)\}\bigg|_{\theta=\theta_0} = 0.$$

Existence of derivative follows from assumption of exponential family.

Such a test is UMPU.

Conditional inference: a story

An experiment is conducted to measure the carbon monoxide level in the exhaust of a car. A sample of exhaust gas is collected, and taken to the laboratory for analysis. Inside the laboratory are two machines, one of which is expensive and very accurate, the other an older model which is much less accurate. We will use the accurate machine if we can, but this may be out of service or already in use for another analysis. We do not have time to wait for this machine to become available, so if we cannot use the more accurate machine we use the other one instead (which is always available). Before arriving at the laboratory we have no idea whether the accurate machine will be available, but we do know that the probability that it is available is $\frac{1}{2}$ (independently from one visit to the next).

Formulation

We observe (δ, Y), where δ (=1 or 2) represents the machine used and Y the subsequent observation. The distributions are $\mathbb{P}\{\delta = 1\} = \mathbb{P}\{\delta = 2\} = \frac{1}{2}$ and, given δ, $Y \sim N(\theta, \sigma_\delta^2)$ where θ is unknown and σ_1, σ_2 are known, with $\sigma_1 < \sigma_2$.

We want to test $H_0: \ \theta \leq \theta_0$ against $H_1: \ \theta > \theta_0$.

Two possible tests

Consider the following tests:

- Procedure 1. Reject H_0 if $Y > c$, where c is chosen so that the test has prescribed size α, i.e.

$$\frac{1}{2}\left\{1-\Phi\left(\frac{c-\theta_0}{\sigma_1}\right)\right\}+\frac{1}{2}\left\{1-\Phi\left(\frac{c-\theta_0}{\sigma_2}\right)\right\}=\alpha.$$

- Procedure 2. Reject H_0 if $Y > z_\alpha \sigma_\delta + \theta_0$, z_α is upper α-quantile of $N(0,1)$.

Comparison

Procedure 1 sets a single critical level c, regardless of which machine is used, while Procedure 2 determines its critical level solely on the standard deviation for the machine that was actually used without taking the other machine into account at all.

Procedure 2 is a conditional test because it conditions on the observed value of δ. The distribution of δ itself does not depend on the unknown parameter θ, so we are not losing any information by doing this.

Power comparison

We might expect Procedure 2 to be more reasonable, but if we compare the two in terms of power it is not so clear-cut. The diagram shows the power curves of the two tests in the case $\sigma_1 = 1,\ \sigma_2 = 3,\ \alpha = 0.05$, for which $z_\alpha = 1.645$ and $c = 3.8457 + \theta_0$. When the difference in means, $\theta_1 - \theta_0$, is small, procedure 2 is much more powerful, but for larger values when $\theta_1 > \theta_0 + 4.9$, procedure 1 is better.

The power functions

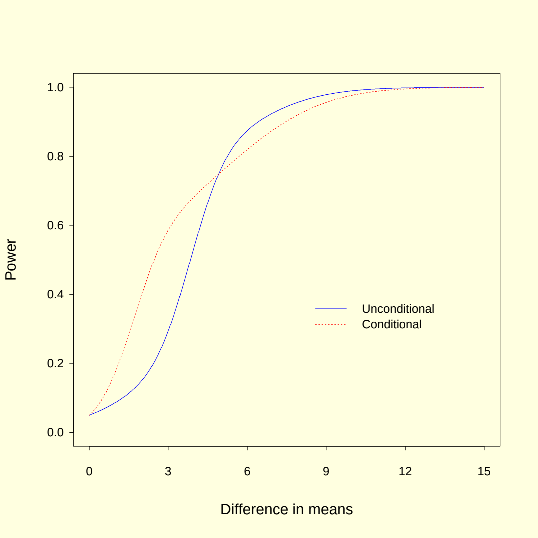

Smith and Jones

Smith and Jones are two statisticians. Smith works for the environmental health department of Cambridge City Council and Jones is retained as a consultant by a large haulage firm which operates in the Cambridge area.

Smith carries out a test of the exhaust fumes emitted by one of the lorries belonging to the haulage firm. He has to use machine 2 and the observation is $X = \theta_0 + 4.0$, where θ_0 is the permitted standard.

It has been agreed in advance that all statistical tests will be carried out at the 5% level and therefore, following procedure 1, he reports that the company is in violation of the standard.

Send for Jones!

The company is not satisfied with the conclusion and sends the results to Jones for comment. The information available to Jones is that a test was conducted on a machine for which the standard deviation of all measurements is 3 units, that the observed measurement exceeded the standard by 4 units, and that therefore the null hypothesis (that the lorry is meeting the standard) is rejected at the 5% level.

Jones calculates that the critical level should be $\theta_0 + 3z_{0.05} = \theta_0 + 3 \times 1.645 = \theta_0 + 4.935$ and therefore queries why the null hypothesis was rejected.

Back to Smith

The query is referred back to Smith who now describes the details of the test including the existence of the other machine and Smith's preference for procedure 1 over procedure 2 on the grounds that procedure 1 is of higher power when $|\theta_1 - \theta_0|$ is large.

This however is all news to Jones who was not previously aware that the other machine even existed.

Call in the lawyers!

The question facing Jones now is: should she revise her opinion on the basis of the new information provided by Smith?

She does not see why she should. There is no new information about either the sample that was collected or the way that it was analysed. All that is new is that there was another machine which might have been used for the test, but which in the event was not. Jones cannot see why this is relevant and therefore advises the company to challenge the test in court.

The Conditionality Principle, revisited

The minimal sufficient statistic for θ is (Y, δ), and δ has a distribution not depending on θ.

The Conditionality Principle argues that inference about θ should be based on the conditional distribution of Y given δ.

Similar tests

Suppose we have $\theta = (\psi, \lambda)$, with ψ the interest parameter and λ a nuisance parameter. Suppose the minimal sufficient statistic $T = (S, C)$, where the conditional distribution of S given $C = c$ depends on ψ, but not λ, for each c.

Test using the conditional distribution of S given C. This eliminates the nuisance parameter: the test is similar.

Similarity: definition

Suppose $\theta = (\psi, \lambda)$ and the parameter space is of the form $\Theta = \Psi \times \Lambda$. Suppose we wish to test the null hypothesis $H_0 : \ \psi = \psi_0$ against the alternative $H_1 : \ \psi \neq \psi_0$, with λ treated as a nuisance parameter.

Suppose $\phi(y), \ y \in \mathcal{Y}$ is a test of size α for which

$$\mathbb{E}_{\psi_0,\lambda}\{\phi(X)\} = \alpha \text{ for all } \lambda \in \Lambda.$$

Then ϕ is called a similar test of size α.

Some discussion

More generally, if the parameter space is $\theta \in \Omega_\theta$ and the null hypothesis is of the form $\theta \in \Theta_0$, where Θ_0 is a subset of Ω_θ, then a similar test is one for which $\mathbb{E}_\theta\{\phi(X)\} = \alpha$ on the boundary of Θ_0.

By analogy with UMPU tests, if a test is uniformly most powerful among the class of all similar tests, we call it UMP similar.

If the power function is continuous in θ then any unbiased test of size α must have power exactly α on the boundary between Θ_0 and Θ_1, i.e. it is similar.

In such cases, if we can find a UMP similar test, and if this test turns out also to be unbiased, then it is necessarily UMPU.

Some further discussion

In many cases we can show that a test which is UMP among all tests based on the conditional distribution of S given C, is UMP amongst all similar tests. In particular, this is valid when C is a complete sufficient statistic for λ.

In summary, there are many cases when a test which is UMP (one-sided) or UMPU (two-sided), based on the conditional distribution of S given C, is in fact UMP similar or UMPU among the class of all tests.

Reasons for conditioning

So, there are two quite distinct arguments for conditioning.

- When the conditioning statistic is ancillary, failure to condition may lead to paradoxical situations in which two analysts may form completely different viewpoints of the same data, even though conditioning may run counter to the strict viewpoint of maximising power.
- Under certain circumstances a conditional test may satisfy the conditions needed to be UMP similar or UMPU. This argument is explicitly based on power considerations.

Multiparameter exponential families

Consider a full exponential family model in its natural parametrisation,

$$f(y;\theta) = c(\theta)h(y)\exp\left(\sum_{i=1}^{m} t_i(y)\theta^i\right),$$

where y represents the value of a data vector Y and $t_i(Y),\ i = 1, \ldots, m$ are the natural statistics. Write T_i in place of $t_i(Y)$.

Suppose interest is in one particular parameter, θ^1. Consider the test $H_0 : \theta^1 \leq \theta^{1*}$ against $H_1 : \theta^1 > \theta^{1*}$, where θ^{1*} is prescribed.

Take $S = T_1$ and $C = (T_2, ..., T_m)$. Then the conditional distribution of S given C is also of exponential family form and does not depend on $\theta^2, ..., \theta^m$. Therefore, C is sufficient for $\lambda = (\theta^2, ..., \theta^m)$ and since it is also complete, arguments concerning similar tests suggest that we ought to construct tests for θ^1 based on the conditional distribution of S given C.

Such tests do turn out to be UMPU.

Sometimes C is an ancillary statistic for θ^1: then there is the stronger argument based on the Conditionality Principle for conditioning on C.

The form of test: one-sided case

If the distribution of T_1 is continuous, the optimal one-sided test will then be of the following form. Suppose we observe $T_1 = t_1, ..., T_m = t_m$. Then we reject H_0 if and only if $t_1 > t_1^*$, where t_1^* is calculated from

$$\mathbb{P}_{\theta^{1*}}\{T_1 > t_1^* | T_2 = t_2, ..., T_m = t_m\} = \alpha.$$

It can be shown that this test is UMPU of size α.

The form of test: two-sided case

If we want to construct a two-sided test $H_0: \ \theta^{1*} \leq \theta^1 \leq \theta^{1**}$ against the alternative, $H_1: \ \theta^1 < \theta^{1*}$ or $\theta^1 > \theta^{1**}$, where $\theta^{1*} < \theta^{1**}$ are given, we proceed by defining the conditional power function based on T_1 as

$$w_{\theta^1}(\phi; t_2, ..., t_m) = \mathbb{E}_{\theta^1}\{\phi(T_1) | T_2 = t_2, ..., T_m = t_m\}.$$

This quantity depends only on θ^1 and not on $\theta^2, .., \theta^m$.

We can then consider a two-sided conditional test of the form

$$\phi'(t_1) = \begin{cases} 1 & \text{if } \; t_1 < t_1^* \text{ or } t_1 > t_1^{**}, \\ 0 & \text{if } \; t_1^* \leq t_1 \leq t_1^{**}, \end{cases}$$

where t_1^* and t_1^{**} are chosen such that

$$w_{\theta^1}(\phi'; t_2, ..., t_m) = \alpha \quad \text{when } \theta^1 = \theta^{1*} \text{ and } \theta^1 = \theta^{1**}. \qquad (P)$$

A Useful Result

Suppose $C = (T_2, \ldots, T_m)$, and suppose that $V \equiv V(T_1, C)$ is a statistic independent of C, with $V(t_1, c)$ increasing in t_1 for each c.

The UMPU test is equivalent to that based on the marginal distribution of V: the conditional test is the same as that obtained by testing H_0 against H_1 using V as test statistic.

An Example

Normal distribution $N(\mu, \sigma^2)$: given an independent sample $X_1, \ldots, X_n$, to test a hypothesis about σ^2, the conditional test is based on the conditional distribution of $T_1 \equiv \sum_{i=1}^n X_i^2$, given the observed value of $C \equiv \bar{X}$.

Let $V = T_1 - nC^2 \equiv \sum_{i=1}^{n}(X_i - \bar{X})^2$.

We know that V is independent of C (from general properties of the normal distribution), so the optimal conditional test is equivalent to that based on the marginal distribution of V.

We have that V/σ^2 is chi-squared, χ^2_{n-1}.

The form of test: two-sided case, point hypothesis

If the hypotheses are of the form $H_0:\ \theta^1 = \theta^{1*}$ against $H_1:\ \theta^1 \neq \theta^{1*}$, then the test is of the same form but with (P) replaced by

$$w_{\theta^{1*}}(\phi'; t_2, ..., t_m) = \alpha,$$

$$\frac{d}{d\theta^1}\left\{w_{\theta^1}(\phi'; t_2, ..., t_m)\right\}\Big|_{\theta^1=\theta^{1*}} = 0.$$

Optimal point estimation

Optimal point estimation of (scalar) parameter θ.

- Minimum variance unbiased estimator.
- More generally, unbiased estimator minimising convex loss function.

Jensen's inequality

If $g : \mathbb{R} \to \mathbb{R}$ is a convex function and X is a real-valued random variable, then $\mathbb{E}\{g(X)\} \geq g(\mathbb{E}\{X\})$.

Convex: $g(\alpha x_1 + (1-\alpha)x_2) \leq \alpha g(x_1) + (1-\alpha)g(x_2)$, for any x_1, x_2, $\alpha \in [0,1]$

Theorem

To estimate a real-valued parameter θ with an estimator $d(Y)$ say. The loss function $L(\theta, d)$ is a convex function of d for each θ.

Let $d_1(Y)$ be an unbiased estimator for θ and suppose T is a sufficient statistic. Then the estimator

$$\chi(T) = \mathbb{E}\{d_1(Y)|T\}$$

is also unbiased and has risk not exceeding that of d_1.

Remarks

- The inequality above will be strict unless L is a linear function of d, or the conditional distribution of $d_1(Y)$ given T is degenerate. In all other cases, $\chi(T)$ strictly dominates $d_1(Y)$.
- If T is also complete, then $\chi(T)$ is the unique unbiased estimator minimising the risk.
- If $L(\theta, d) = (\theta - d)^2$ then this is the Rao-Blackwell Theorem. Now risk of an unbiased estimator $\equiv$ variance, so there is a unique minimum variance unbiased estimator which is a function of the complete sufficient statistic.