Chapter 9

Interval Estimation

9.1 Introduction

Definition 9.1.1 *An interval estimate of a real-values parameter θ is any pair of functions, $L(x_1, \ldots, x_n)$ and $U(x_1, \ldots, x_n)$, of a sample that satisfy $L(\boldsymbol{x}) \leq U(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathcal{X}$. If $\boldsymbol{X} = \boldsymbol{x}$ is observed, the inference $L(\boldsymbol{x}) \leq \theta \leq U(\boldsymbol{x})$ is made. The random interval $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ is called an interval estimator.*

Although in the majority of cases we will work with finite values for L and U, there is sometimes interest in one-sided interval estimates. For instance, if $L(\boldsymbol{x}) = -\infty$, then we have the one-sided interval $(-\infty, U(x)]$ and the assertion is that $\theta \leq U(\boldsymbol{x})$. We could similarly take $U(\boldsymbol{x}) = \infty$ and have a one-sided interval $[L(\boldsymbol{x}), \infty)$. Although the definition mentions a closed interval $[L(\boldsymbol{x}), U(\boldsymbol{x})]$, it will sometimes be more natural to use an open interval $([L(\boldsymbol{x}), U(\boldsymbol{x}))$ or even a half-open and half-closed interval. We will use whichever seems most appropriate for the particular problem at hand.

Example 9.1.1 (**Interval estimator**) *Let $X_1, \ldots, X_4$ be a sample from $N(\mu, 1)$. When we estimate μ by $\bar{X}$, the probability that we are exactly correct, that is, $P(\bar{X} = \mu) = 0$. However, with an interval estimator, for example, $[\bar{X}-1, \bar{X}+1]$, we have a positive probability of being correct. The probability that μ is covered by the interval is*

$$\begin{aligned} P(\mu \in [\bar{X} - 1, \bar{X} + 1]) &= P(\bar{X} - 1 \leq \mu \leq \bar{X} + 1) \\ &= P(-2 \leq (\bar{X} - \mu)/\sqrt{1/4} \leq 2) = 0.9544. \end{aligned}$$

Sacrificing some precision in our estimate, in moving from a point to an interval, has resulted in increased confidence that our assertion is correct.

Definition 9.1.2 *For an interval estimator* $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ *of a parameter* θ*, the coverage probability of* $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ *is the probability that the random interval* $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ *covers the true parameter,* θ*. In symbols, it is denoted by either* $P_\theta(\theta \in [L(\boldsymbol{X}), U(\boldsymbol{X})])$ *or* $P_\theta(\theta \in [L(\boldsymbol{X}), U(\boldsymbol{X})]|\theta)$.

Definition 9.1.3 *For an interval estimator* $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ *of a parameter* θ*, the confidence coefficient of* $[L(\boldsymbol{X}), U(\boldsymbol{X})]$ *is the infimum of the coverage probabilities,* $\inf_\theta P_\theta(\theta \in [L(\boldsymbol{X}), U(\boldsymbol{X})])$.

Note that the interval is the random quantity, not the parameter. Interval estimators, together with a measure of confidence (usually a confidence coefficient), are sometimes known as confidence intervals. We will often use this term interchangeably with interval estimator. Another point is concerned with coverage probabilities and confidence coefficients. Since we do not know the true value of θ, we can only guarantee a coverage probability equal to the infimum, the confidence coefficient. In some cases this does not matter because the coverage probability will be a constant function of θ. In other cases, however, the coverage probability can be a fairly variable function of θ.

Example 9.1.2 *(***Scale uniform interval estimator***) Let $X_1, \ldots, X_n$ be a random sample from a uniform$(0, \theta)$ population and let $Y = \max\{X_1, \ldots, X_n\}$. We are interested in an interval estimator of θ. We consider two candidate estimators: $[aY, bY]$, $1 \le a < b$, and $[Y + c, Y + d]$, $0 \le c < d$, where a, b, c and d are specified constants. For the first interval we have*

$$P_\theta(\theta \in [aY, bY]) = P_\theta(aY \le \theta \le bY) = P_\theta(\frac{1}{b} \le \frac{Y}{b} \le \frac{1}{a}) = P_\theta(\frac{1}{b} \le T \le \frac{1}{a}$$

by defining $T = Y/\theta$. The pdf of T is $f_T(t) = nt^{n-1}$, $0 \le t \le 1$. We therefore have

$$P_\theta(\frac{1}{b} \le T \le \frac{1}{a}) = \int_{1/b}^{1/a} nt^{n-1} dt = (\frac{1}{a})^n - (\frac{1}{b})^n.$$

Thus, the coverage probability of the interval is independent of the value of θ, $(\frac{1}{a})^n - (\frac{1}{b})^n$ is the confidence coefficient of the interval.

For the other interval, for $\theta \ge d$ a similar calculation yields

$$P_\theta(\theta \in [Y + c, Y + d]) = P_\theta(1 - \frac{d}{\theta} \le T \le 1 - \frac{c}{\theta}) = \int_{1-d/\theta}^{1-c/\theta} nt^{n-1} dt$$
$$= (1 - \frac{c}{\theta})^n - (1 - \frac{d}{\theta})^n.$$

In this case, the coverage probability depends on θ. It is easy to *see that for any constants c and d,*

$$\lim_{\theta\to\infty}(1-\frac{c}{\theta})^n-(1-\frac{d}{\theta})^n=0,$$

showing that the confidence coefficient of this interval estimator is 0.

9.2 Methods of Finding Interval Estimators

9.2.1 Inverting a test Statistic

There is a strong correspondence between hypothesis testing and interval estimation. In fact, we can say in general that every confidence set corresponding to a test and vice versa. Consider the following example.

Example 9.2.1 *(***Inverting a normal test***) Let $X_1, \ldots, X_n$ be iid $N(\mu, \sigma^2)$ and consider testing $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$. For a fixed α level, a reasonable test has rejection region $\{\boldsymbol{x} : |\bar{x} - \mu_0| > z_{\alpha/2}\sigma/\sqrt{n}\}$. That is, H_0 is accepted with*

$$\bar{x} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \leq \mu_0 \leq \bar{x} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}.$$

Since the test has size α, this means that $P(H_0$ is accepted$|\mu = \mu_0) = 1 - \alpha$, i.e.,

$$P(\bar{X} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \leq \mu_0 \leq \bar{X} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}|\mu = \mu_0) = 1 - \alpha.$$

But this probability statement is true for every μ_0. Hence, the statement

$$P_\mu(\bar{X} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{X} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}) = 1 - \alpha$$

is true. The interval $[\bar{x}-z_{\alpha/2}\frac{\sigma}{\sqrt{n}}, \bar{x}+z_{\alpha/2}\frac{\sigma}{\sqrt{n}}]$*, obtained by inverting the acceptance region of the level* α *test, is a* $1-\alpha$ *confidence interval.*

In this example, the acceptance region, the set in sample space for which $H_0 : \mu = \mu_0$ is accepted, is given by

$$A(\mu_0) = \{(x_1, \dots, x_n) : \mu_0 - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \le \bar{x} \le \mu_0 + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\},$$

and the confidence interval, the set in parameter space with plausible values of μ, is given by

$$C(x_1, \dots, x_n) = \{\mu : \bar{x} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{x} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\}.$$

These sets are connected to each other by the tautology

$$(x_1, \dots, x_n) \in A(\mu_0) \Leftrightarrow \mu_0 \in C(x_1, \dots, x_n).$$

Hypothesis test and confidence set both look for consistency between sample statistics and population parameters. The hypothesis test fixes the parameter and asks what sample values (acceptance region) are consistent with that fixed value. The confidence set fixes the sample value and asks what parameter values (the confidence interval) make this sample value most plausible. The next theorem gives a formal version of this correspondence.

Theorem 9.2.1 *For each $\theta_0 \in \Theta$, let $A(\theta_0)$ be the acceptance region of a level α test of $H_0 : \theta = \theta_0$. For each $\boldsymbol{x} \in \mathcal{X}$, define a set $C(\boldsymbol{x})$ in the parameter space by*

$$C(\boldsymbol{x}) = \{\theta_0 : \boldsymbol{x} \in A(\theta_0)\}. \tag{9.1}$$

Then the random set $C(\boldsymbol{X})$ is a $1-\alpha$ confidence set. Conversely, let $C(\boldsymbol{X})$ be a $1-\alpha$ confidence set. For any $\theta_0 \in \Theta$, define

$$A(\theta_0) = \{\boldsymbol{x} : \theta_0 \in C(\boldsymbol{x})\}.$$

Then $A(\theta_0)$ is the acceptance region of a level α test of $H_0 : \theta = \theta_0$.

Proof: For the first part, since $A(\theta_0)$ is the acceptance region of a level α test,

$$P_{\theta_0}(\boldsymbol{X} \notin A(\theta_0)) \leq \alpha \quad \text{and hence} \quad P_{\theta_0}(\boldsymbol{X} \in A(\theta_0)) \geq 1 - \alpha.$$

Since θ_0 is arbitrary, write θ instead of θ_0. The above inequality, together with (9.1), showing that the coverage probability of the set $C(\boldsymbol{X})$ is given by

$$P_\theta(\theta \in C(\boldsymbol{X})) = P_\theta(\boldsymbol{X} \in A(\theta)) \geq 1 - \alpha,$$

showing that $C(\boldsymbol{X})$ is a $1 - \alpha$ confidence set.

For the second part, the type I error probability for the test of $H_0 : \theta = \theta_0$ with acceptance region $A(\theta_0)$ is

$$P_{\theta_0}(\boldsymbol{X} \notin A(\theta_0)) = P_{\theta_0}(\theta_0 \notin C(\boldsymbol{X})) \leq \alpha.$$

So this is a level α test. $\square$

In Theorem 9.2.1, we stated only the null hypothesis $H_0 : \theta = \theta_0$. All that is required of the acceptance region is

$$P_{\theta_0}(\boldsymbol{X} \in A(\theta_0)) \geq 1 - \alpha.$$

In practice, when constructing a confidence set by test inversion, we will also have in mind an alternative hypothesis such as $H_1 : \theta \neq \theta_0$ or $H_0 : \theta > \theta_0$. The alternative will dictate the form of $A(\theta_0)$ that is reasonable, and the form of $A(\theta_0)$ will determine the shape of $C(\boldsymbol{x})$. Note, however, that we carefully the word *set* rather than *interval*. This is because there is no guarantee that the confidence set obtained by test inversion will be an interval. In most cases, however, one-sided tests give one-sided intervals, two-sided tests give two sided intervals, strange-shaped acceptance regions give strange-shaped confidence sets. Later examples will exhibit this.

The properties of the inverted test also carry over to the confidence set. For example, unbiased tests, when inverted, will produce unbiased confidence sets. Also, and more important, since we know that we can confine attention to sufficient statistics when looking for a good test, it follows that we can confine attention to sufficient statistics when looking for a good confidence sets.

Example 9.2.2 *(***Inverting an LRT***) Suppose that we want a confidence interval for the mean, λ, of an exponential(λ) population. We can obtain such an interval by inverting a level α test of $H_0 : \lambda = \lambda_0$ versus $H_1 : \lambda \neq \lambda_0$.*

Take a random sample $X_1, \dots, X_n$, the LRT statistic is given by

$$\frac{\frac{1}{\lambda_0^n} e^{-\sum x_i/\lambda_0}}{\sup_\lambda \frac{1}{\lambda^n} e^{-\sum x_i/\lambda}} = \frac{\frac{1}{\lambda_0^n} e^{-\sum x_i/\lambda_0}}{\frac{1}{(\sum x_i/n)^n} e^{-n}} = \Big(\frac{\sum x_i}{n\lambda_0}\Big)^n e^n e^{-\sum x_i/\lambda_0}.$$

For fixed λ_0, the acceptance region is given by

$$A(\lambda_0) = \Big\{\boldsymbol{x} : \frac{\sum x_i}{\lambda_0}\Big)^n e^{-\sum x_i/\lambda_0} \geq k^*\Big\}, \tag{9.2}$$

where k^* is a constant chosen to satisfy $P_{\lambda_0}(\boldsymbol{X} \in A(\lambda_0)) = 1-\alpha$. (The constant e^n/n^n has been absorbed into k^*.) Inverting this *acceptance region gives the* $1-\alpha$ *confidence set*

$$C(\boldsymbol{x}) = \big\{\lambda : \frac{\sum x_i}{\lambda}\big)^n e^{-\sum x_i/\lambda} \geq k^*\big\}.$$

So the confidence interval can be expressed in the form

$$C(\sum x_i) = \{\lambda : L(\sum x_i) \leq \lambda \leq U(\sum x_i)\},$$

where L *and* U *are functions determined by the constraint that the set (9.2) has probability* $1-\alpha$ *and*

$$\big(\frac{\sum x_i}{L(\sum x_i)}\big)^n e^{-\sum x_i/L(\sum x_i)} = \big(\frac{\sum x_i}{U(\sum x_i)}\big)^n e^{-\sum x_i/U(\sum x_i)}. \tag{9.3}$$

If we set

$$\frac{\sum x_i}{L(\sum x_i)} = a \quad \text{and} \quad \frac{\sum x_i}{U(\sum x_i)} = b, \tag{9.4}$$

where $a > b$, *then (9.3) becomes*

$$a^n e^{-a} = b^n e^{-b}.$$

To work out some details, let $n = 2$ *and note that* $\sum x_i \sim \mathrm{Gamma}(2,\lambda)$ *and* $\sum x_i/\lambda \sim \mathrm{Gamma}(2,1)$. *Hence, from (9.4), the confidence interval becomes* $\{\lambda : \frac{1}{a}\sum x_i \leq \lambda \leq \frac{1}{b}\sum x_i\}$, *where* a *and* b *satisfy*

$$P_\lambda(\frac{1}{a}\sum x_i \leq \lambda \leq \frac{1}{b}\sum x_i) = P(b \leq \frac{\sum x_i}{\lambda} \leq a)$$
$$= \int_b^a te^{-t}dt = e^{-b}(b+1) - e^{-a}(a+1) = 1-\alpha.$$

To get a $1-\alpha$ *confidence interval, we can solve the system*

$$e^{-b}(b+1) - e^{-a}(a+1) = 1-\alpha.$$

$$a^2e^{-a} = b^2e^{-b}.$$

For example, let $\alpha = 0.1$*, we get* $a = 5.480$*,* $b = 0.441$*, with a confidence coefficient of 0.90006. Thus,*

$$P_\lambda(\frac{1}{5.480}\sum X_i \le \lambda \le \frac{1}{0.441}\sum X_i) = 0.90006.$$

The region obtained by inverting the LRT of $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$ is of the form

$$\text{accept } H_0 \text{ if } \tfrac{L(\theta_0|\boldsymbol{x})}{L(\hat{\theta}|\boldsymbol{x})} \le k(\theta_0),$$

with the resulting confidence region

$$\{\theta : L(\theta|\boldsymbol{x}) \ge k'(\boldsymbol{x}, \theta)\},$$

for some function k' that gives $1-\alpha$ confidence.

The test inversion method is completely general in that we can invert any test and obtain a confidence set. In the preceding example, we inverted LRTs, but we could have used a test constructed by any method.

Example 9.2.3 *(***Normal one-sided confidence bound***) Let $X_1, \ldots, X_n$ be a random sample from a $N(\mu, \sigma^2)$ population. Consider constructing a $1-\alpha$ upper confidence bound for μ. That is, we want a confidence interval of the form $C(\boldsymbol{x}) = (-\infty, U(\boldsymbol{x})]$. To obtain such an interval using Theorem 9.2.1, we will invert one-sided tests of $H_0 : \mu = \mu_0$ versus $H_1 : \mu < \mu_0$. The size α LRT of H_0 versus H_1 rejects H_0 if*

$$\frac{\bar{X} - \mu_0}{S/\sqrt{n}} < -t_{n-1,\alpha}.$$

Thus the acceptance region for this test is

$$A(\mu_0) = \{\boldsymbol{x} : \bar{x} \geq \mu_0 - t_{n-1,\alpha}\frac{s}{\sqrt{n}}\}$$

and $\boldsymbol{x} \in A(\mu_0) \Leftrightarrow \bar{x} + t_{n-1,\alpha}s/\sqrt{n} \geq \mu_0$. We define

$$C(\boldsymbol{x}) = \{\mu_0 : \boldsymbol{x} \in A(\mu_0)\} = \{\mu_0 : \bar{x} + t_{n-1,\alpha}\frac{s}{\sqrt{n}} \geq \mu_0\}.$$

By Theorem 9.2.1, the random set $C(\boldsymbol{X}) = (-\infty, \bar{X} + t_{n-1,\alpha}S/\sqrt{n}]$ is a $1-\alpha$ confidence set for μ. We see that, inverting the one-sided test gave a one-sided confidence interval.

9.2.2 Pivotal Quantities

The two confidence intervals that we saw in Example 9.1.2 differed in many respects. One important difference was that the coverage probability of the interval $\{AY, bY\}$ did not depend on the value of the parameter θ, while that of $\{Y + c, Y + d\}$ did. This happened because the coverage probability of $\{aY, bY\}$ could be expresses in terms of the quantity Y/θ, a random variable whose distribution does not depend on the parameter, a quantity known as a *pivotal quantity*, or *pivot*.

Definition 9.2.1 *A random variable* $Q(\boldsymbol{X}, \theta) = Q(X_1, \ldots, X_n, \theta)$ *is a pivotal quantity (or pivot) if the distribution of* $Q(\boldsymbol{X}, \theta)$ *is independent of all parameters. That is, if* $\boldsymbol{X} \sim F(\boldsymbol{x}|theta)$*, then* $Q(\boldsymbol{X}, \theta)$ *has the same distribution for all values of* θ*.*

The function $Q(\boldsymbol{x}, \theta)$ will usually explicitly contain both parameters and statistics, but for any set $\mathcal{A}$, $P_\theta(Q(\boldsymbol{X}, \theta) \in \mathcal{A})$ cannot depend on θ. The technique of constructing confidence sets from pivots relies on being able to find a pivot and a set $\mathcal{A}$ so that the set $\{\theta : Q(\boldsymbol{X}, \theta) \in \mathcal{A}\}$ is a set estimate of θ.

Table 9.1: Location-scale pivots

Form of pdf	Type of pdf	Pivotal quantity
$f(x-\mu)$	location	$\bar{X}-\mu$
$\frac{1}{\sigma}f(\frac{x}{\sigma})$	scale	$\frac{\bar{X}}{\sigma}$
$\frac{1}{\sigma}f(\frac{x-\mu}{\sigma})$ location-scale	$\frac{\bar{X}-\mu}{S}$	

Example 9.2.4 *(***Location-scale pivots***) In location and scale cases there are lots of pivots. Let $X_1, \ldots, X_n$ be a random sample from the indicated pdfs, and let $\bar{X}$ and S be the sample mean and standard deviation. In particular if $X_1, \ldots, X_n$ is from $N(\mu, \sigma^2)$, then the t statistic $\frac{\bar{X}-\mu}{S}$ is pivot because the t distribution does not depend on the parameters μ and σ^2.*

In general, differences are pivotal for location problems, while ratios are pivotal for scale problems.

Example 9.2.5 (**Gamma pivot**) *Suppose that* $X_1, \ldots, X_n$ *are iid exponential*(λ)*. Then* $T = \sum X_i$ *is a sufficient statistic for* λ *and* $T \sim \text{gamma}(n, \lambda)$*. In the gamma pdf* t *and* λ *appear together as* t/λ *and, in fact the gamma*(n, λ) *pdf* $(\Gamma(n)\lambda^n)^{-1} t^{n-1} e^{-t/\lambda}$ *is a scale family. Thus, if* $Q(T, \lambda) = 2T/\lambda$*, then*

$$Q(T, \lambda) \sim \text{gamma}(n, \lambda(2/\lambda)) = \text{gamma}(n, 2) \quad \textit{or} \quad \chi^2_{2n},$$

which does nor depend on λ*. The quantity* $Q(T, \lambda) = 2T/\lambda$ *is a pivot.*

In the above example, the quantity t/λ appeared in the pdf and this turned out to be a pivot. In the normal pdf, the quantity $(\bar{x} - \mu)/\sigma$ appears and this quantity is also a pivot. In general, suppose the pdf of a statistic T, $f(t|\theta)$, can be expressed in the form

$$f(t|\theta) = g(Q(t, \theta)) \left| \frac{\partial}{\partial t} Q(t, \theta) \right| \tag{9.5}$$

for some function and some monotone function Q (monotone in t for each θ). Then $Q(T, \theta)$ is a pivot.

Once we have a pivot, it is easy to construct a confidence set. We will use the pivot to specify the specific form of our acceptance region, and use the test inversion method to construct the confidence set. If $Q(\boldsymbol{x}, t)$ is a pivot, then for a specified value of α we can find numbers a and b, which do not depend on θ, to satisfy

$$P_\theta(a \le Q(\boldsymbol{x}, \theta) \le b) \ge 1 - \alpha.$$

Then, for each $\theta_0 \in \Theta$,

$$A(\theta_0) = \{\boldsymbol{x} : a \le Q(\boldsymbol{x}, \theta_0) \le b\}$$

is the accept region for a level α test of $H_0 : \theta = \theta_0$. Using Theorem 9.2.1, we invert these tests to obtain

$$C(\boldsymbol{x}) = \{\theta_0 : a \le Q(\boldsymbol{x}, \theta_0) \le b\},$$

and $C(\boldsymbol{X})$ is a $1 - \alpha$ confidence set for θ. If θ is a real-valued parameter and if, for each $\boldsymbol{x} \in \mathcal{X}$, $Q(\boldsymbol{x}, \theta)$ is a monotone function of θ, then $C(\boldsymbol{x})$ will be an interval. In fact, if $Q(\boldsymbol{x}, \theta)$ is an increasing function of θ, then $C(\boldsymbol{x})$ has the form $L(\boldsymbol{x}, a) \le \theta \le U(\boldsymbol{x}, b)$. If $Q(\boldsymbol{x}, \theta)$ is a decreasing function of θ, then $C(\boldsymbol{x})$ has the form $L(\boldsymbol{x}, b) \le \theta \le U(\boldsymbol{x}, a)$.

Example 9.2.6 *(***Continuation of Example 9.2.5***) Consider the test $H_0 : \lambda = \lambda_0$ versus $\lambda \neq \lambda_0$, if we choose constants a and b to satisfy $P(a \leq \chi^2_{2n} \leq b) = 1 - \alpha$, then*

$$P_\lambda(a \leq \frac{2T}{\lambda} \leq b) = P_\lambda(a \leq Q(T, \lambda) \leq b) = P(a \leq \chi^2_{2n} \leq b) = 1-\alpha.$$

Inverting the set $A(\lambda) = \{t : a \leq \frac{2t}{\lambda} \leq b\}$ gives $C(t) = \{\lambda : \frac{2t}{b} \leq \lambda \leq \frac{2t}{a}\}$, which is a $1 - \alpha$ confidence interval. Note here $Q(\boldsymbol{x}, \lambda)$ is decreasing in λ. For example, if $n = 10$, then consulting a table of χ^2 cutoffs shows that a 95% confidence interval is given by $\{\lambda : \frac{2T}{34.17} \leq \lambda \leq \frac{2T}{9.59}\}$.

Example 9.2.7 *(***Normal pivotal interval***) If $X_1, \ldots, X_n$ are iid $N(\mu, \sigma^2)$, then $(\bar{X} - \mu)/(\sigma/\sqrt{n})$ is a pivot. If σ^2 is known, we can use this pivot to construct a confidence interval for μ. For any constant a,*

$$P(-a \leq \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \leq a) = P(-a \leq Z \leq a),$$

where Z is standard normal, and the resulting confidence interval is

$$\{\mu : \bar{x} - a\frac{\sigma}{\sqrt{n}} \leq \mu \leq \mu + a\frac{\sigma}{\sqrt{n}}\}.$$

If σ^2 is unknown, we can use the location-scale pivot $\frac{\bar{X}-\mu}{S/\sqrt{n}}$. Since

$\frac{\bar{X}-\mu}{S/\sqrt{n}}$ *has Student's t distribution,*

$$P(-a \le \frac{\bar{X}-\mu}{S/\sqrt{n}} \le a) = P(-a \le T_{n-1} \le a).$$

Thus, for any given α, *if we take* $a = t_{n-1,\alpha/2}$, *we find that a* $1-\alpha$ *confidence interval is given by*

$$\{\mu : \bar{x} - t_{n-1,\alpha/2}\frac{s}{\sqrt{n}} \le \mu \le \bar{x} + t_{n-1,\alpha/2}\frac{s}{\sqrt{n}}\},$$

which is the classic $1-\alpha$ *confidence interval for* μ *based on Student's t distribution.*

Because $(n-1)S^2/\sigma^2 \sim \chi^2_{n-1}$, $(n-1)S^2/\sigma^2$ *is also a pivot. A confidence interval for* σ^2 *can be constructed as follows. Choose a and b to satisfy*

$$P(a \le (n-1)S^2/\sigma^2 \le b) = P(a \le \chi^2_{n-1} \le b) = 1-\alpha.$$

Invert this set to obtain the $1-\alpha$ *confidence interval*

$$\{\sigma^2 : (n-1)S^2/b \le \sigma^2 \le (n-1)S^2/a\}$$

or, equivalently,

$$\{\sigma : \sqrt{(n-1)S^2/b} \le \sigma \le \sqrt{(n-1)S^2/a}\}.$$

One choice of a and b that will produce the required interval is $a = \chi^2_{n-1,1-\alpha/2}$ *and* $b = \chi^2_{n-1,\alpha/2}$.

9.2.3 Pivoting the CDF

First we consider the case where T is continuous. The situation where T is discrete is similar but has a few additional technique details to consider. We, therefore, state the discrete in a separate theorem.

First, recall theorem 2.1.10 :

Let X have continuous cdf $F_X(x)$ and define the random variable Y as $Y = F_X(X)$. Then Y is uniformly distributed on $(0,1)$, that is, $P(Y \leq y) = y$, $0 < y < 1$.

This theorem tells us that the random variable $F_T(T|\theta)$ is uniform(0,1), a pivot. Thus, if $\alpha_1 + \alpha_2 = \alpha$, an α-level acceptance region of the hypothesis $H_0 : \theta = \theta_0$ is

$$\{t : \alpha_1 \leq F_T(t|\theta) \leq 1 - \alpha_2\},$$

with associated confidence set

$$\{\theta : \alpha_1 \leq F_T(t|\theta) \leq 1 - \alpha_2\}.$$

Now to guarantee that the confidence set is an interval, we need to have $F_T(t|\theta)$ to be monotone in θ. But we have seen this already, in the definitions of stochastically increasing and stochastically decreasing.

Theorem 9.2.2 *(***Pivoting a continuous cdf***) Let T be a statistic with continuous cdf $F_t(t|\theta)$. Let $\alpha_1 + \alpha_2 = \alpha$ with $0 < \alpha < 1$ be fixed values. Suppose that for each $t \in \mathcal{T}$, the functions $\theta_L(t)$ and $\theta_U(t)$ can be defined as follows.*

i. If $F_T(t|\theta)$ is a decreasing function of θ for each t, define $\theta_L(t)$ and $\theta_U(t)$ by

$$F_T(t|\theta_U(t)) = \alpha_1, \quad F_T(t|\theta_L(t)) = 1 - \alpha_2.$$

ii. If $F_T(t|\theta)$ is an increasing function of θ for each t, define $\theta_L(t)$ and $\theta_U(t)$ by

$$F_T(t|\theta_U(t)) = 1 - \alpha_2, \quad F_T(t|\theta_L(t)) = \alpha_1.$$

Then the random interval $[\theta_L(T), \theta_U(T)]$ is a $1 - \alpha$ confidence interval for θ.

Proof: We will prove only part (i). The proof of part (ii) is similar. Assume that we have constructed the $1 - \alpha$ acceptance region

$$\{t : \alpha_1 \le F_T(t|\theta_0) \le 1 - \alpha_2\}.$$

Since $F_T(t|\theta)$ is a decreasing function of θ for each t and $1 - \alpha_2 > \alpha_1$,

$\theta_L(t) < \theta_U(t)$, and the values $\theta_L(t)$ and $\theta_U(t)$ are unique. Also,

$$F_T(t|\theta) < \alpha_1 \Leftrightarrow \theta > \theta_U(t)$$

$$F_T(t|\theta) > 1 - \alpha_2 \Leftrightarrow \theta < \theta_L(t)$$

and hence $\{\theta : \alpha_1 \leq F_T(t|\theta) \leq 1 - \alpha_2\} = \{\theta : \theta_L(T) \leq \theta \leq \theta_U(T)\}$. $\square$

We note that, in the absence of additional information, it is common to choose $\alpha_1 = \alpha_2 = \alpha/2$. Although this may not be optimal, it is certainly a reasonable strategy in most situations. If a one-sided interval is desired, however, this can easily be achieved by choosing either α_1 or α_2 equal to 0.

The equations for case (i) can also expressed in terms of the pdf of the statistic T. The functions $\theta_U(t)$ and $\theta_L(t)$ can be defined to satisfy

$$\int_{-\infty}^{t} f_T(u|\theta_U(t))du = \alpha_1 \quad \text{and} \quad \int_{t}^{\infty} f_T(u|\theta_L(t))du = \alpha_2. \quad (9.6)$$

A similar set of equations holds for case (ii).

Example 9.2.8 *(***Location exponential interval***) If $X_1, \ldots, X_n$ are iid with pdf $f(x|\mu) = e^{-(x-\mu)} I_{[\mu,\infty)}(x)$, then $Y = \min\{X_1, \ldots, X_n\}$ is sufficient for μ with pdf*

$$f_y(y|\mu) = ne^{-n(y-\mu)} I_{[\mu,\infty)}(y).$$

Fix α and define $\mu_L(y)$ and $\mu_U(y)$ to satisfy

$$\int_{\mu_U(y)}^{y} ne^{-n(u-\mu_U(y))}du = \frac{\alpha}{2}, \quad \int_{y}^{\infty} ne^{-n(u-\mu_L(y))}du = \frac{\alpha}{2}.$$

These integrals can be evaluated to give the equations

$$1 - e^{-n(y-\mu_U(y))} = \frac{\alpha}{2}, \quad e^{-n(y-\mu_L(y))} = \frac{\alpha}{2},$$

which give us the solutions

$$\mu_U(y) = y + \frac{1}{n}\log(1-\frac{\alpha}{2}), \quad \mu_L(y) = y + \frac{1}{n}\log(\frac{\alpha}{2}).$$

Hence, the random interval

$$C(Y) = \int \{\mu : Y + \frac{1}{n}\log(\frac{\alpha}{2}) \leq \mu \leq Y + \frac{1}{n}\log(1-\frac{\alpha}{2})\},$$

is a -1α confidence interval for μ.

Note two things about the use of this method. First, the equations (9.6) need to be solved only for the values of the statistics actually observed. If $T = t_0$ is observed, then the realized confidence interval on θ will be $[\theta_L(t_0), \theta_U(t_0)]$. Thus, we need to solve only the two equations

$$\int_{-\infty}^{t_0} f_T(u|\theta_U(t_0))du = \alpha_1 \quad \text{and} \quad \int_{t_0}^{\infty} f_T(u|\theta_L(t_0))du = \alpha_2$$

for $\theta_L(t_0)$ and $\theta_U(t_0)$. Second, the two equations can be solved numerically if a analytical solution is not available.

Theorem 9.2.3 (**Pivoting a discrete cdf**) *Let T be a discrete statistic with cdf $F_T(t\theta) = P(T \leq t|\theta)$. Let $\alpha_1 + \alpha_2 = \alpha$ with $0 < \alpha < 1$ be fixed values. Suppose that for each $t \in \mathcal{T}$, $\theta_L(t)$ and $\theta_U(t)$ can be defined as follows.*

i. If $F_T(t\theta)$ is a decreasing function of θ for each t, define $\theta_L(t)$ and $\theta_U(t)$ by

$$P(T \leq t|\theta_U(t)) = \alpha_1, \quad P(T \geq t|\theta_L(t)) = \alpha_2.$$

ii. If $F_T(t\theta)$ is an increasing function of θ for each t, define $\theta_L(t)$ and $\theta_U(t)$ by

$$P(T \geq t|\theta_U(t)) = \alpha_1, \quad P(T \leq t|\theta_L(t)) = \alpha_2.$$

Then the random interval $[\theta_L(t), \theta_U(t)]$ is $1-\alpha$ confidence interval for θ.

Proof: We will only sketch the proof of part (i). First recall Exercise 2.10, where it was shown that $F_T(t\theta)$ is stochastically greater than a uniform random variable, that is, $P_\theta(F_T(T|\theta) \leq x) \leq x$. This implies that the set

$$\{\theta : F_t(T \leq t|\theta) \geq \alpha_1 \quad \text{and} \quad 1 - F_T(T|\theta) \geq \alpha_2\}$$

is a $1-\alpha$ confidence set.

The fact that $F_T(t|\theta)$ is a decreasing function of θ for each t implies that $1 - F_T(t|\theta)$ is a nondecreasing function of θ for each t. It therefore follows that

$$\theta > \theta_U(t) \Rightarrow F_T(t|\theta) < \alpha_1, \quad \theta < \theta_L(t) \Rightarrow 1 - F_T(t|\theta) < \alpha_2,$$

and hence $\{\theta : F_t(T \leq t|\theta) \geq \alpha_1 \quad \text{and} \quad 1 - F_T(T|\theta) \geq \alpha_2\} = \{\theta : \theta_L(T) \leq \theta \leq \theta_U(T)\}$. $\square$

Example 9.2.9 *(***Poisson interval estimator***) Let $X_1, \ldots, X_n$ be a random sample from a Poisson population with parameter λ and define $Y = \sum X_i$. Y is sufficient for λ and $Y \sim$ Poisson$(n\lambda)$. Applying the above method with $\alpha_1 = \alpha_2 = \alpha/2$, if $Y = y_0$ is observed, we are led to solve for λ in the equations*

$$\sum_{k=0}^{y_0} e^{-n\lambda}\frac{(n\lambda)^k}{k!} = \frac{\alpha}{2} \quad \textit{and} \quad \sum_{k=y_0}^{\infty} e^{-n\lambda}\frac{(n\lambda)^k}{k!} = \frac{\alpha}{2}.$$

Recall the identity, from Example 3.3.1, linking the Poisson and gamma families. Apply that identity to the above sums, we can write

$$\frac{\alpha}{2} = \sum_{k=0}^{y_0} e^{-n\lambda}\frac{(n\lambda)^\lambda}{k!} = P(Y \leq y_0|\lambda) = P(\chi^2_{2(y_0+1)} > sn\lambda),$$

where $\chi^2_{2(y_0+1)}$ *is a chi squared random variable with* $2(y_0+1)$ *degrees of freedom. Thus, the solution to the above equation is to take*

$$\lambda = \frac{1}{2n}\chi^2_{2(y_0+1),\alpha/2}.$$

Similarly, applying the identity to the other equation yields

$$\frac{\alpha}{2} = \sum_{k=y_0}^{\infty} e^{-n\lambda}\frac{(n\lambda)^k}{k!} = P(Y \geq y_0|\lambda) = P(\chi^2_{2y_0} < 2n\lambda).$$

Thus, we obtain the $1-\alpha$ *confidence interval for* λ *as*

$$\{\lambda : \frac{1}{2n}\chi^2_{2y_0,1-\alpha/2} \leq \lambda \leq \frac{1}{2n}\chi^2_{2(y_0+1),\alpha/2}\}.$$

(At $y_0 = 0$ *we define* $\chi^2_{2y_0,1-\alpha/2} = 0$.*)*

For a numerical example, consider $d = 10$ *and observe* $y_0 = \sum x_i = 6$*. A* 90% *confidence interval for* λ *is given by*

$$\frac{1}{20}\chi^2_{12,0.95} \leq \lambda \leq \frac{1}{20}\chi^2_{14,0.05},$$

which is $0.262 \leq \lambda \leq 1.184$.

9.2.4 Bayesian Intervals

To keep the distinction between Bayesian and classical sets clear, the Bayesian set estimates are referred to as *credible sets* rather than confidence sets.

If $\pi(\theta|\boldsymbol{x})$ is the posterior distribution of θ given $\boldsymbol{X} = \boldsymbol{x}$, then for any set $A \subset \Theta$, the credible probability of A is

$$P(\theta \in A|\boldsymbol{x}) = \int_A \pi(\theta|\boldsymbol{x}) d\theta, \tag{9.7}$$

and A is a credible set for θ. If $\pi(\theta|\boldsymbol{x})$ is a pmf, we replace integrals with sums in the above expressions.

Example 9.2.10 *(***Poisson credible set***) Let $X_1, \dots, X_n$ be iid Poisson(λ) and assume that λ has a gamma prior pdf, $\lambda \sim$* gamma(a, b). *The posterior pdf of λ is*

$$pi(\lambda|\sum X_i = \sum x_i) = \text{gamma}(a + \sum x_i, [n + 1/b]^{-1}).$$

We can form a credible set for λ in many different ways, as any set A satisfying (9.7) will do. One simple way is to split the α equally between the upper and lower endpoints. Thus, a $1-\alpha$ credible interval is

$$\{\lambda : \frac{b}{2(nb+1)} \chi^2_{2(\sum x + a), 1-\alpha/2} \leq \lambda \leq \frac{b}{2(nb+1)} \chi^2_{2(\sum x + a), \alpha/2}\}.$$

If we take $a = b = 1$, the posterior distribution of λ given $\sum X = \sum x$ can then be expressed as $2(n+1)\lambda \sim \chi^2_{2(\sum x + 1)}$. As in Example 9.2.9, assume $n = 10$ and $\sum x = 6$, a 90% credible set for λ is given by $[0.299, 1.077]$.

Example 9.2.11 *(***Coverage of a normal credible set***) Let $X_1, \ldots, X_n$ be iid $N(\theta, \sigma^2)$, and let θ have the prior pdf $N(\mu, \tau^2)$, where μ, σ and τ are all known. Thus, we have*

$$\pi(\theta|\bar{x}) \sim N(\delta^B(\bar{x}), Var(\theta|\bar{x})),$$

where $\delta^B(\bar{x}) = \frac{\sigma^2}{\sigma^2+n\tau^2}\mu + \frac{n\tau^2}{\sigma^2+n\tau^2}\bar{x}$ and $Var(\theta|\bar{x}) = \frac{\sigma^2\tau^2}{\sigma^2+n\tau^2}$. It therefore follows that under the posterior distribution,

$$\frac{\theta - \delta^B(\bar{x})}{\sqrt{Var(\theta|\bar{x})}} \sim N(0,1),$$

and a $1-\alpha$ credible set for θ is given by

$$\delta^B(\bar{x}) - z_{\alpha/2}\sqrt{Var(\theta|\bar{x})} \leq \theta \leq \delta^B(\bar{x}) + z_{\alpha/2}\sqrt{Var(\theta|\bar{x})}.$$

9.3 Methods of Evaluating Interval Estimators

We now have seen many methods for deriving confidence sets and, in fact, we can derive different confidence sets for the same problem. In such situations we would, of course, want to choose a best one. Therefore, we now examine some methods and criteria for evaluating set estimators. In set estimation two quantities vie against each other, size and coverage probability. Naturally, we want our set to have small size and large coverage probability, but such sets are usually difficult to construct.

The coverage probability of a confidence set will, except in special cases, be a function of the parameter, so there is not one value to consider but an infinite number of values. For the most part, however, we will measure coverage probability performance by the confidence coefficient, the infimum of the coverage probabilities.

When we speak of the size of a confidence set we will usually mean the length of the confidence set, if the set is an interval. If the set is not an interval, or if we are dealing with a multidimensional set, then length will usually become volume.

9.3.1 Size and Coverage Probability

Example 9.3.1 *(***Optimizing length***) Let $X_1, \ldots, X_n$ be iid $N(\mu, \sigma^2)$, where σ is known. From the method of section 9.2.2 and the fact that*

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$$

is a pivot with a standard normal distribution, any a and b that satisfy

$$P(a \leq Z \leq b) = 1 - \alpha$$

will give the $1-\alpha$ *confidence interval*

$$\{\mu : \bar{x} - b\frac{\sigma}{\sqrt{n}} \le \mu \le \bar{x} - a\frac{\sigma}{\sqrt{n}}\}.$$

The length of the interval is $(b-a)\frac{\sigma}{\sqrt{n}}$ *and it can be minimized by choosing* $a = -z_{\alpha/2}$ *and* $b = z_{\alpha/2}$ *according to the following theorem.*

Theorem 9.3.1 *If* $f(x)$ *be a unimodal pdf. If the interval* $[a,b]$ *satisfies*

i. $\int_a^b f(x)dx = 1-\alpha$,

ii. $f(a) = f(b) > 0$,

iii. $a \le x^* \le b$, *where* x^* *is a mode of* $f(x)$,

then $[a,b]$ *is the shortest among all intervals that satisfy (i).*

Proof: First, let's recall the definition of unimodal: A pdf $f(x)$ is unimodal if there exists x^* such that $f(x)$ is nondecreasing for $x \le x^*$ and $f(x)$ is non-increasing for $x \ge x^*$.

Let $[a', b']$ be any interval with $b'-a' < b-a$. We will show that this implies $\int_{a'}^{b'} f(x)dx < 1-\alpha$. The result will be proved only for $a' \le a$, the proof being similar if $a' > a$. Also, two cases need to be considered, $b' \le a$ and $b' > a$.

If$b' \le a$, then $a' \le b' \le a \le x^*$ and

$$\int_{a'}^{b'} f(x)dx \le f(b')(b'-a') \le f(a)(b'-a')$$
$$< f(a)(b-a) \le int_a^b f(x)dx = 1-\alpha.$$

If $b' > a$, then $a' \le a < b' < b$. In this case we have

$$\int_{a'}^{b'} f(x)dx = \int_a^b f(x)dx + [\int_{a'}^{a} f(x)dx - \int_{b'}^{b} f(x)dx]$$
$$= (1-\alpha) + [\int_{a'}^{a} f(x)dx - \int_{b'}^{b} f(x)dx]$$

and the theorem will be proved if we show $\int_{a'}^{a} f(x)dx - \int_{b'}^{b} f(x)dx < 0$. Since

$$\int_{a'}^{a} f(x)dx \le f(a)(a-a'), \quad \text{and} \quad \int_{b'}^{b} f(x)dx \ge f(b)(b-b').$$

Thus,

$$\int_{a'}^{a} f(x)dx - \int_{b'}^{b} f(x)dx \le f(a)(a-a') - f(b)(b-b')$$
$$= f(a)[(a-a')-(b-b')] = f(a)[(b'-a')-(b-a)]$$

which is negative if $b'-a' < b-a$ and $f(a) > 0$. □

Example 9.3.2 *(***Shortest pivotal interval***) Suppose $X \sim rmgamma(k, \beta)$. The quantity $Y = X/\beta$ is a pivot, with $Y \sim$* gamma$(k, 1)$, *so we can get a confidence interval by finding con-*

stants a and b to satisfy

$$P(a \leq Y \leq b) = 1 - \alpha.$$

However, a blind application of Theorem 9.3.1 will not give the shortest confidence interval. That is because, the interval on β is of the form

$$\{\beta : \frac{x}{b} \leq \beta \leq \frac{x}{a}\},$$

so the length of the interval is $(1/a - 1/b)x$; that is, it is proportional to $(1/a - 1/b$ and not to $b - a$.

The shortest pivotal interval can be found by solving the following constrained minimization problem:

Minimize, with respect to a: $\frac{1}{a} - \frac{1}{b(a)}$ *subject to:* $\int_a^{b(a)} f_Y(y)dy = 1 - \alpha$

where b is defined as a function of a.

9.3.2 Test-Related Optimality

Since there is a one-to-one correspondence between confidence sets and tests of hypotheses, there is some correspondence between optimality of tests and optimality of confidence sets. Usually, test-related optimality properties of confidence sets do not directly relate to the size of the set but rather to the probability of the set covering false values, or the probability of false coverage.

We first consider the general situation, where $\boldsymbol{X} \sim f(\boldsymbol{x}|\theta)$, and we construct a $1-\alpha$ confidence set for θ, $C(\boldsymbol{x})$, by inverting an acceptance region, $A(\theta)$. The probability of coverage of $C(\boldsymbol{x})$, that is, the probability of true coverage, is the function of θ given by $P_\theta(\theta \in C(\boldsymbol{X}))$. The probability of the false coverage is the function of θ and θ' defined by

$$\begin{aligned} P_\theta(\theta' \in C(\boldsymbol{X})), \theta \neq \theta', &\text{if} \quad C(\boldsymbol{X}) = [L(bX), U(\boldsymbol{X})] \\ P_\theta(\theta' \in C(\boldsymbol{X})), \theta' < \theta, &\text{if} \quad C(\boldsymbol{X}) = [L(bX), \infty) \\ P_\theta(\theta' \in C(\boldsymbol{X})), \theta' > \theta, &\text{if} \quad C(\boldsymbol{X}) = (-\infty, U(\boldsymbol{X})] \end{aligned}$$

the probability of covering θ' when θ is the true parameter.

A $1-\alpha$ confidence set that minimizes the probability of false coverage over a class of $1-\alpha$ confidence sets is called a *uniformly most accurate* (UMA) confidence set.

Theorem 9.3.2 *Let $\boldsymbol{X} \sim f(\boldsymbol{x}|\theta)$, where θ is a real-valued parameter. For each $\theta_0 \in \Theta$, let $A^*(\theta_0)$ be the UMP level α acceptance region of a test of $H_0 : \theta = \theta_0$ versus $H_1 : \theta > \theta_0$. Let $C^*(\boldsymbol{x})$ be the $1-\alpha$ confidence set formed by inverting the UMP* *acceptance regions. Then for any other* $1 - \alpha$ *confidence set* C*,*

$$P_\theta(\theta' \in C^*(\boldsymbol{X})) \leq P_\theta(\theta' \in C(\boldsymbol{X})) \quad \text{for all} \quad \theta' < \theta.$$

PROOF: Let θ' be any value less than θ. Let $A(\theta')$ be the acceptance region of the level α test of $H_0 : \theta = \theta'$ obtained by inverting C. Since $A^*(\theta')$ is the UMP acceptance region for testing $H_0 : \theta = \theta'$ versus $H_1 : \theta > \theta'$, and since $\theta > \theta'$, we have

$$P_\theta(\theta' \in C^*(\boldsymbol{X})) = P_\theta(\boldsymbol{X} \in A^*(\theta)) \leq P_\theta(\boldsymbol{X} \in A(\theta')) = P_\theta(\theta' \in C(\boldsymbol{X})).$$

Note that the above inequality is "$\leq$" because we are working with probabilities of acceptance regions. This is 1-power, so UMP tests will minimize these acceptance region probabilities. □

Note that in Theorem 9.3.2, the alternative hypothesis leads to lower confidence bounds; that is, if the sets are intervals, they are of the form $[L(\boldsymbol{X}), \infty)$.

Example 9.3.3 *(***UMA confidence bound***) let* $X_1, \ldots, X_n$ *be iid* $N(\mu, \sigma^2)$*, where* σ^2 *is known. The interval*

$$C(\bar{x}) = \{\mu : \mu \geq \bar{x} - z_\alpha \frac{\sigma}{\sqrt{n}}\}$$

is a $1 - \alpha$ *UMA lower confidence bound since it can be obtained by inverting the UMP test of* $H_0 : \mu = \mu_0$ *versus* $H_1 : \mu > \mu_0$*.*

The more common two-sided interval,

$$C(\bar{x}) = \{\mu : \bar{x} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{x} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\},$$

is not UMA, since it is obtained by inverting the two-sided acceptance region from the test of $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$, hypotheses for which no UMP test exists.

In the testing problem, when considering two-sided tests, we found the property of unbiasedness to be both compelling and useful. In the confidence interval problem, similar ides apply. When we deal with two-sided confidence intervals, it is reasonable to restrict consideration to unbiased confidence sets.

Definition 9.3.1 *A $1 - \alpha$ confidence set $C(\boldsymbol{x})$ is unbiased if $P_\theta(\theta' \in C(\boldsymbol{X})) \leq 1 - \alpha$ for all $\theta \neq \theta'$.*

Thus, for an unbiased confidence set, the probability of false coverage is never more than the minimum probability of true coverage. Unbiased confidence sets can be obtained by inverting unbiased test. That is, if $A(\theta_0)$ is an unbiased level α acceptance region of a test of $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$ and $C(\boldsymbol{x})$ is the $1 - \alpha$ confidence set formed by inverting the acceptance regions, then $C(\boldsymbol{x})$ is an unbiased $1 - \alpha$ confidence set.

Example 9.3.4 *(***Continuation of Example 9.3.3***) The two-sided normal interval*

$$C(\bar{x}) = \{\mu : \bar{x} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}} \leq \mu \leq \bar{x} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\},$$

is an unbiased interval. It can be obtained by inverting the unbiased test of $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$.

9.3.3 Bayesian Optimality

If we have a posterior distribution $\pi(\theta|\boldsymbol{x})$, the posterior distribution of θ given $\boldsymbol{X} = \boldsymbol{x}$, we would like to find the set $C(\boldsymbol{x})$ that satisfies

$$(i) \quad \int_{C(\boldsymbol{x})} \pi(\theta|\boldsymbol{x})d\boldsymbol{x} = 1 - \alpha$$

$$(ii) \quad \text{Size}(C(\boldsymbol{x})) \leq \text{Size}(C'(\boldsymbol{x}))$$

for any set $C'(\boldsymbol{x})$ satisfying $\int_{C'(\boldsymbol{x})} \pi(\theta|\boldsymbol{x})d\boldsymbol{x} \geq 1 - \alpha$.

If we take our measure of size to be length, then we can apply Theorem 9.3.1 and obtain the following result.

Corollary 9.3.1 *If the posterior density $\pi(\theta|\boldsymbol{x})$ is unimodal, then for a given value of α, the shortest credible interval for θ is given by*

$$\{\theta : \pi(\theta|\boldsymbol{x}) \geq k\} \quad \textit{where} \quad \int_{\{\theta:\pi(\theta|\boldsymbol{x})\geq k\}} \pi(\theta|\boldsymbol{x})d\theta = 1 - \alpha.$$

Example 9.3.5 *(***Poisson HPD region***) In Example 9.2.10 we derived a $1-\alpha$ credible set for a Poisson parameter. We now construct an HPD region. By Corollary 9.3.1, this region is given by $\{\lambda : \pi(\lambda|\sum x) \geq k\}$, where k is chosen so that*

$$1-\alpha = \int_{\{\lambda:\pi(\lambda|\sum x)\geq k\}} \pi(\lambda|\sum x)d\lambda.$$

Recall that the posterior pdf of λ is gamma($a+\sum x, [n+(1/b)]^{-1}$), so we need to find λ_L and λ_U such that

$$\pi(\lambda_L|\sum x) = \pi(\lambda_U|\sum x) \quad \textit{and} \quad \int_{\lambda_L}^{\lambda_U} \pi(\lambda|\sum x)d\lambda = 1-\alpha.$$

If we take $a=b=1$, the posterior distribution of λ given $\sum X = \sum x$ can be expressed as $2(n+1)\lambda \sim \chi^2_{2(\sum x+1)}$ and, if $n=10$ and $\sum x = 6$, the 90% HPD credible set for λ is given by $[0.253, 1.005]$.

Example 9.3.6 *(***Normal HPD region***) The equal-tailed credible set derived in Example 9.2.11 is, in fact, an HPD region. Since the posterior distribution of θ is normal with mean δ^B, it follows that $\{\theta : \pi(\theta|\bar{x}) \geq k\} = \{\theta : \theta \in \delta^B \pm k'\}$ for some k'. So the HPD region is symmetric about the mean $\delta^B(\bar{x})$.*