2

Accuracy of Estimators, Confidence Intervals and Tests

To determine the accuracy of the estimates that we made in Chapter 1, we demonstrate how to calculate confidence intervals and perform tests. The specific concerns of each particular experiment are described herein. We then introduce the methodology, first describing classical asymptotic procedures and two asymptotic tests, the Wald test and the likelihood ratio test, and then we present procedures and tests based on a resampling method, the bootstrap. As before, we conclude the chapter by applying these methods to the examples.

2.1 Examples

Let us assume that in Example 1.1.1 we are interested in the maximum yield, θ_1. We have calculated one estimate of θ_1, $\widehat{\theta}_1 = 69.95$; however, if we do another experiment under the same experimental conditions, the observed values of Y and the estimates of the parameters will be different. Thus, knowing one estimate is not entirely satisfactory; we need to quantify its accuracy.

In Example 1.1.2, we calculated an estimate of the calibration curve. Suppose we now want to estimate the dose of hormone D contained in a preparation that has the expected response $\mu = 2000$ c.p.m. To do this, we must use Equation (1.6), replacing the parameters with their estimates. We find $\widehat{X} = -1.1033$ and $\widehat{D} = \exp \widehat{X} \log 10 = 0.3318$ ng/.1 ml, but we now need to calculate how much confidence we can place in this estimate.

Let us consider Example 1.1.3, in which we estimated two ELISA response curves. Our concern in this experiment was to estimate the relative potency of the two different sera. In order to do this, however, we must verify whether the condition of *parallelism*, as expressed in Equation (1.7), is true. If it does exist for all values of x, then the parameters satisfy that $\theta_1^{\mathrm{May}} = \theta_1^{\mathrm{June}}$, $\theta_2^{\mathrm{May}} = \theta_2^{\mathrm{June}}$, $\theta_3^{\mathrm{May}} = \theta_3^{\mathrm{June}}$. In this case, $\beta = \theta_4^{\mathrm{May}} - \theta_4^{\mathrm{June}}$ is the horizontal distance between the two curves at the inflection point. To determine parallelism we will first test whether these relations between the parameters are true or, more exactly, if they do not contradict the data. If the test does not reject the hypothesis of parallelism, we will be able to estimate β and test if it is significantly different from zero.

In Example 1.1.4 we were interested in comparing parameters P_w and P_s. Because water permeability in cells is higher than propane-diol permeability, water flows out of the cells more rapidly than the propane-diol flows in, resulting in high cell shrinkage. Thus we are interested in the speed of intracellular propane-diol penetration; the cryoprotectant must permeate the ovocytes in a short time. To this end, we will compare the values of V_s at times $T_1 = 1$ mn, $T_2 = 5$ mn, and so on.

Using Example 1.1.5, we will compare different methods for calculating confidence intervals for the parameters.

In sum, what we want to do in all of these examples is to determine if we estimated the function of the parameters denoted $\lambda(\theta)$ accurately. In the first example, $\lambda = \theta_1$, and in the second example, $\lambda = \exp(X \log 10)$, where X is defined by

$$X = \frac{1}{\theta_4}\left\{\log\left[\exp\frac{1}{\theta_5}\log\frac{\theta_2-\theta_1}{\mu-\theta_1} - 1\right] - \theta_3\right\}.$$

In the third example, if the hypothesis of parallelism is not rejected, we are interested in $\lambda = \beta$ or $\lambda = \exp -\beta$. In the fourth example we are interested in the pairs (P_w, P_s) for each curve, in $\lambda = V_s(T_1)$, and, in $\lambda = V_s(T_2)$.

2.2 Problem Formulation

The nonlinear regression model was defined in Section 1.2. Let $\widehat{\theta}$ be the least squares estimator of θ when we have homogeneous variances, and let it be the weighted least squares estimator of θ when we have heterogeneous[1] variances (see Section 1.3).

Let λ be a function of the parameters and $\widehat{\lambda}$ an estimator of λ: $\widehat{\lambda} = \lambda(\widehat{\theta})$. In this chapter, we will describe how to calculate a confidence interval for λ, and how to do a test.

The function λ must satisfy some *regularity* assumptions. It must be a continuous function of θ with continuous partial derivatives with respect to θ.

2.3 Solutions

2.3.1 Classical Asymptotic Results

$\widehat{\theta}$ is a function of the Y_{ij}, and when the number of observations tends to infinity, its distribution is known: $\widehat{\theta} - \theta$ tends to 0, and the *limiting distribution*

---

[1] See Chapter 3 for a complete treatment when the variance of errors is not constant.

of $V_{\widehat{\theta}}^{-1/2}(\widehat{\theta}-\theta)$ is a standard p-dimensional normal (Gaussian) distribution $\mathcal{N}(0, \mathrm{I}_p)$ with expectation 0 and variance I_p, where I_p is the $p \times p$ identity matrix and $V_{\widehat{\theta}}$ is the estimated asymptotic *covariance matrix* of $\widehat{\theta}$. Thus, for sufficiently large n, the distribution of $\widehat{\theta}$ may be approximated by the normal distribution $\mathcal{N}(\theta, V_{\widehat{\theta}})$.

We need a result for $\widehat{\lambda} = \lambda(\widehat{\theta})$. The limiting distribution (when n tends to infinity) of

$$\widehat{T} = \frac{\widehat{\lambda} - \lambda}{\widehat{S}}$$

will be a centered normal distribution, $\mathcal{N}(0,1)$, where $\widehat{S}$ is an estimate of the standard error of $\widehat{\lambda}$.

Notations and Formulas: f_i is for $f(x_i, \theta)$. The p vector of derivatives of f with respect to θ calculated in x_i is denoted by $\partial f_i/\partial\theta$. The components of $\partial f_i/\partial\theta$ are $[(\partial f/\partial\theta_a)(x_i,\theta)]$, $a = 1, \ldots p$.

Let $\mathit{\Gamma}_\theta$ be the $p \times p$ matrix defined as follows:

$$\mathit{\Gamma}_\theta = \frac{1}{\sigma^2}\frac{1}{n}\sum_{i=1}^{k} n_i \frac{\partial f_i}{\partial\theta}\left(\frac{\partial f_i}{\partial\theta}\right)^T,$$

where the exponent T means that the vector is transposed. The elements (a,b) of $\mathit{\Gamma}_\theta$ are

$$\mathit{\Gamma}_{\theta,ab} = \frac{1}{\sigma^2}\frac{1}{n}\sum_{i=1}^{k} n_i \frac{\partial f_i}{\partial\theta_a}\frac{\partial f_i}{\partial\theta_b}.$$

Let $\mathit{\Delta}_\theta$ be the $p \times p$ matrix

$$\frac{1}{n}\sum_{i=1}^{k} \frac{n_i}{\sigma_i^2} \frac{\partial f_i}{\partial\theta}\left(\frac{\partial f_i}{\partial\theta}\right)^T.$$

Let $\widehat{f_i} = f(x_i, \widehat{\theta})$, $\partial\widehat{f_i}/\partial\theta$ be the vector with components $\partial f/\partial\theta_a(x_i,\widehat{\theta})$, and $\mathit{\Gamma}_{\widehat{\theta}}$ and $\mathit{\Delta}_{\widehat{\theta}}$ be the matrices $\mathit{\Gamma}_\theta$ and $\mathit{\Delta}_\theta$, where the unknown parameters are replaced by their estimators:

$$\mathit{\Gamma}_{\widehat{\theta}} = \frac{1}{\widehat{\sigma}^2}\frac{1}{n}\sum_{i=1}^{k} n_i \frac{\partial\widehat{f_i}}{\partial\theta}\left(\frac{\partial\widehat{f_i}}{\partial\theta}\right)^T, \quad \mathit{\Delta}_{\widehat{\theta}} = \frac{1}{n}\sum_{i=1}^{k} \frac{n_i}{\widehat{\sigma}_i^2} \frac{\partial\widehat{f_i}}{\partial\theta}\left(\frac{\partial\widehat{f_i}}{\partial\theta}\right)^T,$$

and $\widehat{\sigma}^2 = C(\widehat{\theta})/n$, $\widehat{\sigma}_i^2 = s_i^2$.

$V_{\widehat{\theta}}$ is the estimate of V_θ, the $p \times p$ asymptotic covariance matrix of $\widehat{\theta}$:

In the case $\mathrm{Var}(\varepsilon_{ij}) = \sigma^2$, $V_{\widehat{\theta}} = \dfrac{1}{n}\mathit{\Gamma}_{\widehat{\theta}}^{-1}$,

In the case $\mathrm{Var}(\varepsilon_{ij}) = \sigma_i^2$, $V_{\widehat{\theta}} = \dfrac{1}{n}\mathit{\Delta}_{\widehat{\theta}}^{-1}$.

Because $\widehat{\theta}-\theta$ is small, we get the limiting distribution of $\widehat{\lambda}$ by approximating $\lambda(\widehat{\theta}) - \lambda(\theta)$ by a linear function of $\widehat{\theta} - \theta$:

$$\left(\frac{\partial\lambda}{\partial\theta}\right)^T (\widehat{\theta}-\theta) = \sum_{a=1}^{p} \frac{\partial\lambda}{\partial\theta_a}(\widehat{\theta}_a - \theta_a).$$

Thus $(\widehat{\lambda}-\lambda)/S_{\widehat{\theta}}$ is distributed, when n tends to infinity, as an $\mathcal{N}(0,1)$ (a Gaussian centered variate with variance 1), where

$$S_\theta^2 = \left(\frac{\partial\lambda}{\partial\theta}\right)^T V_\theta \frac{\partial\lambda}{\partial\theta} = \sum_{a=1}^{p}\sum_{b=1}^{p} \frac{\partial\lambda}{\partial\theta_a}\frac{\partial\lambda}{\partial\theta_b} V_{\theta,ab} \tag{2.1}$$

and $\widehat{S} = S_{\widehat{\theta}}$ is the asymptotic estimate of the standard error.

2.3.2 Asymptotic Confidence Intervals for $\boldsymbol{\lambda}$

If the distribution of $\widehat{T}$ were known, say $F(u) = \Pr(\widehat{T} \leq u)$, we would calculate the $\alpha/2$ and $1-\alpha/2$ percentiles[2] of $\widehat{T}$, say u_α, $u_{1-\alpha/2}$. The interval

$$\widehat{I} = \left[\widehat{\lambda} - u_{1-\alpha/2}\widehat{S}\ ;\ \widehat{\lambda} - u_{\alpha/2}\widehat{S}\right]$$

would be a confidence interval for λ, with level $1-\alpha$. In this case, the *coverage probability* of $\widehat{I}$, the probability that $\widehat{I}$ covers λ, would be $1-\alpha$.

However, as we have seen in the preceding paragraph, we can only approximate, when n is sufficiently large, the distribution of $\widehat{T}$. Thus we use this approximation to calculate confidence intervals with coverage probability close to $1-\alpha$.

Let $\mathcal{N}$ be a variate distributed as an $\mathcal{N}(0,1)$. Let ν_α be the α percentile of $\mathcal{N}$. From the result of Section 2.3.1, we can deduce a confidence interval for λ:

$$\widehat{I}_{\mathcal{N}} = \left[\widehat{\lambda} - \nu_{1-\alpha/2}\widehat{S}; \widehat{\lambda} + \nu_{1-\alpha/2}\widehat{S}\right]. \tag{2.2}$$

This interval is symmetric around $\widehat{\lambda}$ for $\nu_\alpha = -\nu_{1-\alpha}$.

The probability that $\widehat{T}$ is less than ν_α tends to α when n tends to infinity; the probability for λ to lie in $\widehat{I}_{\mathcal{N}}$ tends to $1-\alpha$ when n tends to infinity. We say that $\widehat{I}_{\mathcal{N}}$ has *asymptotic level* $1-\alpha$.

Remarks

1. By analogy to the Gaussian linear regression case, in the nonlinear regression model with homogeneous variance, we define an alternative confidence interval for λ. In Equation (2.2), we replace ν_α with $\sqrt{n/(n-p)}t_\alpha$, where t_α is the α percentile of a Student variate with $n-p$ degrees of freedom:

$$\widehat{I}_{\mathcal{T}} = \left[\widehat{\lambda} - \sqrt{\frac{n}{n-p}}t_{1-\alpha/2}\widehat{S};\ \widehat{\lambda} + \sqrt{\frac{n}{n-p}}t_{1-\alpha/2}\widehat{S}\right]. \tag{2.3}$$

---

[2] The α percentile of a variate with distribution function F is the value of u, say u_α, such that $F(u_\alpha) = \alpha$ and $0 < \alpha < 1$.

$\widehat{I}_{\mathcal{T}}$ has the same asymptotic level as $\widehat{I}_{\mathcal{N}}$, but $\widehat{I}_{\mathcal{T}}$ is wider than $\widehat{I}_{\mathcal{N}}$ and its coverage probability will be greater. Some studies [HJM89] have shown that $\widehat{I}_{\mathcal{T}}$ has a coverage probability closer to $1-\alpha$ than $\widehat{I}_{\mathcal{N}}$.

2. The intervals $\widehat{I}_{\mathcal{N}}$ and $\widehat{I}_{\mathcal{T}}$ are symmetric around $\widehat{\lambda}$. In some applications, a part of the symmetric confidence interval might not coincide with the set of variations of the parameter λ. For example, consider $\lambda = \exp \theta_3$ in the pasture regrowth example. If the estimate of the standard error of $\widehat{\lambda}$, say $\widehat{S}$, is bigger than $\widehat{\lambda}/\nu_{1-\alpha/2}$, then the lower bound of $\widehat{I}_{\mathcal{N}}$ is negative even though λ is strictly positive. In that case, it is easy to see that it is more appropriate to calculate a confidence interval for θ_3 and then to transform this interval taking the exponential of its limits to find a confidence interval for λ. More generally, let $\widehat{S}_3$ be the estimate of the standard error of $\widehat{\theta}_3$, and let g be a strictly increasing function of θ_3. If θ_3 lies in
$$\left[\widehat{\theta}_3 - \nu_{1-\alpha/2}\widehat{S}_3; \widehat{\theta}_3 + \nu_{1-\alpha/2}\widehat{S}_3\right],$$
then $\lambda = g(\theta_3)$ lies in
$$\left[g(\widehat{\theta}_3 - \nu_{1-\alpha/2}\widehat{S}_3); g(\widehat{\theta}_3 + \nu_{1-\alpha/2}\widehat{S}_3)\right].$$

2.3.3 Asymptotic Tests of $\boldsymbol{\lambda = \lambda_0}$ against $\boldsymbol{\lambda \neq \lambda_0}$

Let λ_0 be a fixed value of λ and let the hypothesis of interest be H: $\lambda = \lambda_0$, against the alternative A: $\lambda \neq \lambda_0$.

Wald Test When H is true the limiting distribution of $(\widehat{\lambda} - \lambda_0)/\widehat{S}$ is an $\mathcal{N}(0,1)$. Thus, the limiting distribution of the test statistic
$$\mathcal{S}_{\mathrm{W}} = \frac{(\widehat{\lambda} - \lambda_0)^2}{\widehat{S}^2}$$
is a χ^2 with one degree of freedom. Hypothesis H will be rejected for large values of $\mathcal{S}_{\mathrm{W}}$, say $\mathcal{S}_{\mathrm{W}} > C$, where C is chosen such that $\Pr(Z_1 \leq C) = 1 - \alpha$, where Z_1 is distributed as a χ^2 with one degree of freedom.

This is the *Wald test.* When H is true, the probability for $\mathcal{S}_{\mathrm{W}}$ to be greater than C (in other words, the probability that hypothesis H is rejected when it should be accepted) tends to α when n tends to infinity. We say that this test has an *asymptotic error of the first kind*, equal to α. Assume now that H is false. Then the power of the test defined as the probability for $\mathcal{S}_{\mathrm{W}}$ to be greater than C (in other words, the probability to reject the hypothesis H when H is false) tends to 1 when n tends to infinity. We say that this test is consistent.

Remark As in Section 2.3.2, homogeneous variances can be considered separately; hypothesis H will be rejected if

$$\frac{n-p}{n}\mathcal{S}_{\mathrm{W}} > C,$$

where C is chosen such that $\Pr(F_{1,n-p} \leq C) = 1-\alpha$, where $F_{1,n-p}$ is distributed as a Fisher variable with one and $n-p$ degrees of freedom.

Likelihood Ratio Test Another idea is to estimate the parameters under the constraint $\lambda = \lambda_0$, say $\widehat{\theta}_{\mathrm{H}}$; then to estimate them without the constraint, say $\widehat{\theta}_{\mathrm{A}}$; and then to compare the estimation criteria (1.10) $C(\widehat{\theta}_{\mathrm{H}})$ and $C(\widehat{\theta}_{\mathrm{A}})$ in the case of homogeneous variances. If H is true, the difference between $C(\widehat{\theta}_{\mathrm{H}})$ and $C(\widehat{\theta}_{\mathrm{A}})$ will be small. Let

$$\mathcal{S}_{\mathrm{L}} = n\log C(\widehat{\theta}_{\mathrm{H}}) - n\log C(\widehat{\theta}_{\mathrm{A}})$$

be the test statistic. When n tends to infinity, it can be shown that the limiting distribution of $\mathcal{S}_{\mathrm{L}}$ is a χ^2 with one degree of freedom. Hypothesis H will be rejected when $\mathcal{S}_{\mathrm{L}} > C$, where C is chosen such that $\Pr(Z_1 \leq C) = 1-\alpha$.

This test based on $\mathcal{S}_{\mathrm{L}}$ is called a likelihood ratio test. It has the same asymptotic properties as the Wald test. Although the Wald test is easier to calculate, some theoretical arguments favor the likelihood ratio test.

2.3.4 Asymptotic Tests of $\Lambda\theta = L_0$ against $\Lambda\theta \neq L_0$

Let us return to Example 1.1.5 and assume that we want to test whether the parameters θ_2, θ_3, and θ_4 are identical. The hypothesis of interest is H: $\theta_2 = \theta_3 = \theta_4$, against the alternative that at least two of these parameters are different. H can be written as $\Lambda\theta = 0$, where Λ is the following 2×4 matrix:

$$\Lambda = \begin{pmatrix} 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix}. \tag{2.4}$$

The problems just defined can be solved by returning to the general case, with θ of dimension p. We aim to test the hypothesis H: $\Lambda\theta = L_0$ against A: $\Lambda\theta \neq L_0$, where Λ is a $q\times p$ matrix of rank q, $q<p$, and L_0 is a vector of dimension q. The model defined by hypothesis H is a model nested in the more general one defined by hypothesis A.

The Wald Test When H is true, the limiting distribution of

$$\mathcal{S}' = (\Lambda V_{\widehat{\theta}}\Lambda^T)^{-1/2}(\Lambda\widehat{\theta} - L_0)$$

is a q-dimensional Gaussian variable with mean 0 and covariance matrix equal to the $q\times q$ identity matrix. In these cases, the limiting distribution of the test statistic $\mathcal{S}_{\mathrm{W}} = \sum_{a=1}^{q}\mathcal{S}_a'^2$ is a χ^2 with q degrees of freedom. The Wald test is defined by the rejection of H when $\mathcal{S}_{\mathrm{W}} > C$, where $\Pr(Z_q \leq C) = 1-\alpha$ and Z_q is distributed as a χ^2 with q degrees of freedom.

Remark In the case of homogeneous variances, the test is defined by the rejection of H when

$$\frac{n-p}{n}\frac{\mathcal{S}_{\mathrm{W}}}{q} > C, \tag{2.5}$$

where C is chosen such that $\Pr(F_{q,n-p} \le C) = 1 - \alpha$, where $F_{q,n-p}$ is distributed as a Fisher with q and $n-p$ degrees of freedom.

The Likelihood Ratio Test Let $\widehat{\theta}_{\mathrm{H}}$ be the estimation of θ under the constraint $\Lambda\theta = L_0$; then, in the case of homogeneous variances, the limiting distribution of the test statistic $\mathcal{S}_{\mathrm{L}} = n \log C(\widehat{\theta}_{\mathrm{H}}) - n \log C(\widehat{\theta}_{\mathrm{A}})$ is a χ^2 with q degrees of freedom. This result provides the likelihood ratio test.

Curve Comparison Let us return to Example 1.1.3, where we needed to compare two curves. The hypothesis of interest is H: $\theta_1^{\mathrm{May}} = \theta_1^{\mathrm{June}}$, $\theta_2^{\mathrm{May}} = \theta_2^{\mathrm{June}}$, $\theta_3^{\mathrm{May}} = \theta_3^{\mathrm{June}}$ against the alternative that at least one of these equalities is false. We create a data set by joining the data observed in May and June. We define the vector of parameters by joining θ^{May} and θ^{June}: Let

$$\theta = \begin{pmatrix} \theta^{\mathrm{May}} \\ \theta^{\mathrm{June}} \end{pmatrix}$$

be the $2p$ vector of parameters for the two curves. Then hypothesis H can be written as earlier $\Lambda\theta = 0$, where Λ is the following $3 \times 2p$ matrix:

$$\Lambda = \begin{pmatrix} 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 \end{pmatrix}. \tag{2.6}$$

As before, we define a test using the statistic $\mathcal{S}_{\mathrm{W}}$ or $\mathcal{S}_{\mathrm{L}}$.

2.3.5 Bootstrap Estimations

Resampling methods like the jackknife and the *bootstrap* are especially useful for estimating the accuracy of an estimator. We observe $Y_1, Y_2, \ldots Y_n$; we choose a parametric nonlinear regression model with parameters θ, and we find an estimation procedure to estimate a function of θ, say $\lambda(\theta)$. We get $\widehat{\lambda} = \lambda(\widehat{\theta})$, but we are interested in calculating the accuracy of $\widehat{\lambda}$ or, more generally, in knowing its distribution (or some characteristics of it). If we were able to repeat the experiment under exactly the same conditions, we would observe $Y_1^1, Y_2^1, \ldots Y_n^1$, and in the same way as for $\widehat{\lambda}$ we would calculate $\widehat{\lambda}^1$. We could repeat it again and calculate $\widehat{\lambda}^2$ with $Y_1^2, Y_2^2, \ldots Y_n^2$. $\widehat{\lambda}^1, \widehat{\lambda}^2, \ldots$ would be a sample of random variables distributed as $\widehat{\lambda}$. This sample would approximate the distribution of $\widehat{\lambda}$. In short, resampling methods are a way to mimic the repetition of the experiment.

Bootstrap estimations are based on estimates $\widehat{\lambda}^\star = \lambda(\widehat{\theta}^\star)$ calculated from artificial bootstrap samples $(x_i, Y_{ij}^\star)$, $j = 1, \ldots, n_i$, $i = 1, \ldots k$, where

$$Y^{\star}_{ij} = f(x_i, \widehat{\theta}) + \varepsilon^{\star}_{ij}.$$

The errors $\varepsilon^{\star}_{ij}$ are simulated in the following way: Let $\widehat{\varepsilon}_{ij} = Y_{ij} - f(x_i, \widehat{\theta})$ be the residuals, and let $\widetilde{\varepsilon}_{ij} = \widehat{\varepsilon}_{ij} - \widehat{\varepsilon}_{\bullet}$, be the centered residuals, where $\widehat{\varepsilon}_{\bullet}$ is the sample mean, and $\widehat{\varepsilon}_{\bullet} = \sum_{i,j} \widehat{\varepsilon}_{ij}/n$. The set of $\varepsilon^{\star}_{ij}$, for $j = 1, \dots n_i$, and $i = 1, \dots k$ is a random sample from the empirical distribution function based on the $\widetilde{\varepsilon}_{ij}$ (n $\widetilde{\varepsilon}_{ij}$ are drawn with replacement, each with probability $1/n$). There are n^n such different samples.

$\widehat{\theta}^{\star}$ will be the value of θ that minimizes

$$C^{\star}(\theta) = \sum_{i=1}^{k} \sum_{j=1}^{n_i} (Y^{\star}_{ij} - f(x_i, \theta))^2.$$

The bootstrap estimate of λ is $\widehat{\lambda}^{\star} = \lambda(\widehat{\theta}^{\star})$.

Let B be the number of bootstrap simulations. $(\widehat{\lambda}^{\star,b} = \lambda(\widehat{\theta}^{\star,b})$, $b = 1, \dots, B)$ is a B *sample* of bootstrap estimates of λ. The choice of B will be discussed at the end of this section. The important result is that the distribution of $\widehat{\lambda}^{\star}$, estimated by the empirical[3] distribution function of the $(\widehat{\lambda}^{\star,b}$, $b = 1, \dots B)$, approximates the distribution of $\widehat{\lambda}$. Let

$$\widehat{T}^{\star} = \frac{\widehat{\lambda}^{\star} - \widehat{\lambda}}{S_{\widehat{\theta}^{\star}}}.$$

Roughly speaking, the difference between the distribution functions of $\widehat{T}$ and $\widehat{T}^{\star}$ tends to 0 when the number of observations n is large; thus, we can use the quantiles of $\widehat{T}^{\star}$ instead of those of $\widehat{T}$ to construct confidence intervals or tests.

Let us emphasize that the bootstrap distribution for approximating the distribution of $\widehat{T}$ is theoretically justified when n is large and is an alternative to the centered normal distribution presented in Section 2.3.1. In real data sets, the number of observations is fixed and may be small. No theoretical result is known about the superiority of one of these approximations over the others.

Bootstrap Confidence Interval for λ

Let $(\widehat{T}^{\star,b}$, $b = 1, \dots B)$ be a B sample of $\widehat{T}^{\star}$; $\widehat{T}^{\star}$ is calculated in the same way as $\widehat{T}$, replacing Y_{ij} with $Y^{\star}_{ij}$. Let b_{α} be the α percentile of the $\widehat{T}^{\star,b}$ (the way of calculating b_{α} is detailed in Section 2.4.1). It can be shown that $\Pr(\widehat{T} \leq b_{\alpha})$ tends to α when n tends to infinity. This gives a *bootstrap confidence interval* for λ:

$$\widehat{I}_B = \left[\widehat{\lambda} - b_{1-\alpha/2}\widehat{S}; \widehat{\lambda} - b_{\alpha/2}\widehat{S}\right]. \tag{2.7}$$

For large n and B, the coverage probability of $\widehat{I}_B$ is close to $1 - \alpha$.

---

[3] Obviously, B must be large enough that the empirical distribution function is a good approximation of the distribution of $\widehat{\lambda}^{\star}$. If $B = n^n$, and if we draw all of the possible samples, we get the exact distribution of $\widehat{\lambda}^{\star}$.

Bootstrap Estimation of the Accuracy of $\widehat{\lambda}$

The variance, and even the bias, of $\widehat{\lambda}$ may be infinite or undefined. Nevertheless, their estimates (using the asymptotic results of Section 2.3.2 or the bootstrap) measure the localization and dispersion of the distribution of $\widehat{\lambda}$.

Variance The *bootstrap estimation of the variance* is calculated using the empirical variance of the B sample $(\widehat{\lambda}^{\star,b},\ b = 1, \ldots, B)$:

$$\widehat{S}^{\star 2} = \sum_{b=1}^{B} \frac{1}{B-1} \left(\widehat{\lambda}^{\star,b} - \widehat{\lambda}^{\star,\bullet} \right)^2 , \tag{2.8}$$

where $\widehat{\lambda}^{\star,\bullet}$ is the sample mean $\widehat{\lambda}^{\star,\bullet} = \sum_{b=1}^{B} \widehat{\lambda}^{\star,b}/B$.

Bias As we noted in Section 2.3.2, the expectation of $\widehat{\lambda}$, $\mathrm{E}(\widehat{\lambda})$, is close to λ when we have large values of n. In other words, the bias of $\widehat{\lambda}$, $\mathrm{BIAS} = \mathrm{E}(\widehat{\lambda}) - \lambda$, is close to 0. We can use the bootstrap sample to estimate this bias:

$$\widehat{\mathrm{BIAS}}^{\star} = \widehat{\lambda}^{\star,\bullet} - \widehat{\lambda}. \tag{2.9}$$

Mean Square Error We can estimate the mean square error (MSE) in a similar way: $\mathrm{MSE} = \mathrm{E}(\widehat{\lambda} - \lambda)^2 = S^2 + \mathrm{BIAS}^2$, where $S^2 = \mathrm{E}(\widehat{\lambda} - \mathrm{E}(\widehat{\lambda}))^2$ is the variance of $\widehat{\lambda}$; it is estimated by

$$\widehat{\mathrm{MSE}}^{\star} = \widehat{S}^{\star 2} + \widehat{\mathrm{BIAS}}^{\star 2}.$$

Median Because it is always defined, the median error, the median of $\widehat{\lambda} - \lambda$, is of special interest. Its bootstrap estimate, $\widehat{\mathrm{MED}}^{\star}$, is the median of the B values $|\widehat{\lambda}^{\star,b} - \widehat{\lambda}|$.

Remarks

1. We have seen that the number of different bootstrap samples equals n^n. Obviously, we never choose for B a value that rapidly becomes unusable ($8^8 = 16,777,216$!). In practice, however, a moderate number usually suffices: If B is around 50, we can estimate the accuracy characteristic, and if B is around 200, we can calculate a confidence interval, for example.
2. Other resampling methods, like the jackknife, are available especially to estimate the accuracy characteristics (see [Wu86] and [Bun90] for details). These methods are less reliable than the bootstrap, however.
3. The bootstrap method in the case of heterogeneous variances is discussed in Section 3.4.3.

2.4 Applications

2.4.1 Pasture Regrowth: Calculation of a Confidence Interval for the Maximum Yield

Model The regression function is

$$f(x, \theta) = \theta_1 - \theta_2 \exp\left(-\exp(\theta_3 + \theta_4 \log x)\right),$$

and the variances are homogeneous: $\mathrm{Var}(\varepsilon_i) = \sigma^2$.

Results

Parameters	Estimated Values	Asymptotic Covariance Matrix			
θ_1	69.95	3.09			
θ_2	61.68	3.87	6.66		
θ_3	-9.209	0.76	1.25	0.37	
θ_4	2.378	-0.22	-0.35	-0.09	0.027
σ^2	0.9306				

The parameter of interest is $\lambda(\theta) = \theta_1$.

Calculation of Confidence Intervals with Asymptotic Level 95%, Using Results of Section 2.3.2 , (d.f. is for degree of freedom):

$\widehat{\lambda}$	$\widehat{S}$	$\nu_{0.975}$	$\widehat{I}_{\mathcal{N}}$	$t_{0.975}$ (5 d.f.)	$\widehat{I}_{\mathcal{T}}$
69.95	1.76	1.96	[66.5 , 73.4]	2.57	[63.9 , 76.0]

d.f. is for degree of freedom

Calculation of Confidence Intervals with Asymptotic Level 95%, Using the Bootstrap Method Table 2.1 gives the estimated values of f, $\widehat{f}_i$, and the centered residuals $\widetilde{\varepsilon}_i$. For two bootstrap simulations, the table gives the bootstrap errors $\varepsilon_i^\star$, the bootstrap observations $Y_i^\star$, the bootstrap estimate of θ_1, and the corresponding asymptotic variance $S_{\widehat{\theta}^\star}$.

B, the number of bootstrap simulations, equals 199. The histogram of the $\widehat{T}^{\star,b}$, $b = 1, \ldots B$, is shown in Figure 2.1.

Calculation of the Percentiles of ($\widehat{T}^{\star,b}$, $b = 1, \ldots B$) We calculate the 0.025 and 0.975 percentiles of the $\widehat{T}^{\star,b}$ as follows: Let $\widehat{T}^{\star,(b)}$ be the ordered values of $\widehat{T}^{\star,b}$ so that $\widehat{T}^{\star,(1)} \leq \widehat{T}^{\star,(2)} \leq \ldots \leq \widehat{T}^{\star,(199)}$; b_α is $\widehat{T}^{\star,(q_\alpha)}$, where q_α is the smallest integer such that q_α/B is greater than or equal to α. When $B = 199$, we find $b_{0.025} = \widehat{T}^{\star,(5)}$ and $b_{0.975} = \widehat{T}^{\star,(195)}$:

$\widehat{\lambda}$	$\widehat{S}$	$b_{0.025}$	$b_{0.975}$	$\widehat{I}_B$
69.95	1.76	-4.19	3.73	[63.4, 77.3]

We will see in Section 2.6 that in practice we use the function `quantile` of S-Plus.

Table 2.1. Results for two bootstrap simulations

$\widehat{f}_i$	$\widetilde{\varepsilon}_i$	$\varepsilon_i^{\star,1}$	$Y_i^{\star,1}$	$\varepsilon_i^{\star,2}$	$Y_i^{\star,2}$
9.411	-0.481	-0.481	8.93	-0.734	8.677
11.47	-0.669	-0.067	11.4	0.025	11.49
16.30	2.284	-0.734	15.57	-0.5	15.80
23.17	-0.843	0.025	23.20	-0.5	22.67
40.08	-0.734	-0.669	39.41	-0.481	39.60
56.18	-0.067	2.284	58.46	-0.669	55.51
60.74	0.986	0.025	60.77	-0.843	59.9
64.59	0.025	-0.734	63.86	-0.734	63.86
67.58	-0.5	2.284	69.86	2.284	69.86
$\widehat{\theta}_1 = 69.95$		$\widehat{\theta}_1^{\star,1} = 71.65$		$\widehat{\theta}_1^{\star,2} = 74.92$	
$\widehat{S} = 1.76$		$S_{\widehat{\theta}^{\star,1}} = 1.91$		$S_{\widehat{\theta}^{\star,2}} = 1.59$	

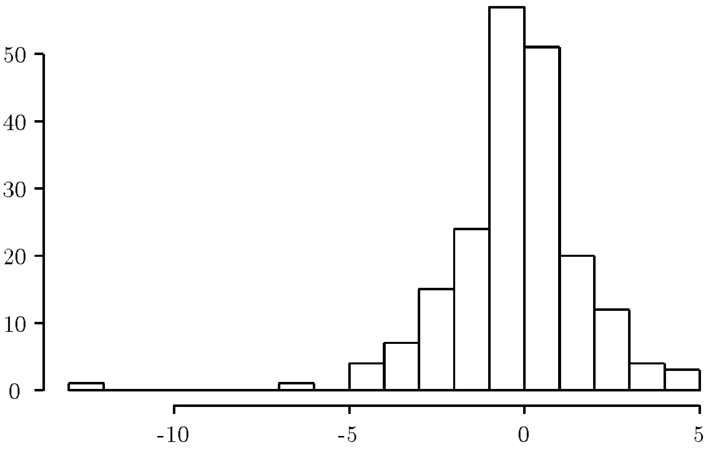

Figure 2.1. Pasture regrowth example: Histogram of ($\widehat{T}^{\star,b}$, $b = 1, \dots B$)

Bootstrap Estimate of the Accuracy Characteristics:

$\widehat{\mathrm{BIAS}}^\star$	$\widehat{S}^\star$	$\widehat{\mathrm{MSE}}^\star$	$\widehat{\mathrm{MED}}^\star$
0.378 (0.5% of $\widehat{\theta}^1$)	2.30	5.42	69.78

2.4.2 Cortisol Assay: Estimation of the Accuracy of the Estimated Dose $\widehat{D}$

Model The regression function is

$$f(x, \theta) = \theta_1 + \frac{\theta_2 - \theta_1}{(1 + \exp(\theta_3 + \theta_4 x))^{\theta_5}},$$

and the variances are heterogeneous: $\mathrm{Var}(\varepsilon_i) = \sigma_i^2$.

Results

	Estimates	Asymptotic Covariance Matrix				
θ_1	133.30	0.727				
θ_2	2759.8	0.264	801			
θ_3	3.0057	-0.0137	-2.34	0.0338		
θ_4	3.1497	-0.00723	-2.33	0.0241	0.01845	
θ_5	0.64309	0.00341	0.568	-0.00714	-0.00516	0.00152

The parameter of interest is $D = \lambda(\theta) = 10^{f^{-1}(\mu,\theta)}$; see Equation (1.6), with $\mu = 2000$.

Calculation of Confidence Intervals with Asymptotic Level 95%, Using Results of Section 2.3.2

$\widehat{D}$	$\widehat{S}$	$\nu_{0.975}$	$\widehat{I}_{\mathcal{N}}$
0.0856	0.00175	1.96	[0.0822 , 0.0891]

We will discuss other methods for calculating the accuracy of $\widehat{D}$ in Chapter 5.

2.4.3 ELISA Test: Comparison of Curves

We want to test the *parallelism* of the response curves in order to estimate the difference $\beta = \theta_4^{\text{May}} - \theta_4^{\text{June}}$. We can do this by testing hypothesis H: $\Lambda\theta = 0$ against the alternative A: $\Lambda\theta \neq 0$, where Λ is the matrix defined by Equation (2.6),

Model The regression function is

$$f(x,\theta) = \theta_1 + \frac{\theta_2 - \theta_1}{1 + \exp\theta_3\,(x - \theta_4)},$$

and the variances are homogeneous: $\text{Var}(\varepsilon_i) = \sigma^2$.

Results

Estimated Values

θ_1^{May}	θ_2^{May}	θ_3^{May}	θ_4^{May}	θ_1^{June}	θ_2^{June}	θ_3^{June}	θ_4^{June}
0.04279	1.936	2.568	3.467	0.0581	1.909	2.836	3.251
Asymptotic Covariance Matrix ($\times 10^4$)							
4.51							
-1.01	1.81						
15.2	-8.22	95.7					
-2.09	-0.502	-4.10	2.43				
0	0	0	0	2.51			
0	0	0	0	-0.647	1.92		
0	0	0	0	10.5	-8.74	106	
0	0	0	0	-1.05	-0.727	-0.988	1.83

$\widehat{\sigma}^2 = 0.0005602$.

Wald Test of H: $\Lambda\theta = 0$ against A: $\Lambda\theta \neq 0$. See Section2.3.3.

The value of the Wald test is $\mathcal{S}_{\mathrm{W}} = 4.53$. This number must be compared to 7.8, which is the 0.95 quantile of a χ^2 with three degrees of freedom. The hypothesis of *parallelism* is not rejected.

Because the variance of errors is constant, we can compare

$$\frac{n-p}{n}\frac{\mathcal{S}_{\mathrm{W}}}{q} = \frac{32-8}{32}\frac{7.8}{3} = 1.133$$

to 3, which is the 0.95 quantile of a Fisher with 3 and 24 degrees of freedom.

Likelihood ratio tests See Section2.3.3.

The two first columns of Table 2.2 show the estimated parameters under A and under the constraints defined by H and the corresponding values of the sum of squares C.

Table 2.2. ELISA test example: Estimated parameters under hypothesis A, the *parallelism* is not verified; hypothesis H, the curves are *parallel*; and the hypothesis that $\beta = 0$

	Under A	Under H	Under $\beta = 0$
θ_1^{May}	0.0428	0.0501	0.0504
θ_2^{May}	1.936	1.924	1.926
θ_3^{May}	2.568	2.688	2.635
θ_4^{May}	3.467	3.470	3.356
θ_1^{June}	0.058	0.0501	0.0504
θ_2^{June}	1.909	1.924	1.926
θ_3^{June}	2.836	2.688	2.635
θ_4^{June}	3.251	3.247	3.356
$\mathcal{C}(\theta)$	0.0179	0.0206	0.183

The test statistic $\mathcal{S}_{\mathrm{L}} = 32 * (\log(0.0206) - \log(0.0179))$ equals 4.5. Thus, hypothesis H is not rejected.

The estimated value of $\beta = \theta_4^{\mathrm{May}} - \theta_4^{\mathrm{June}}$ is $\widehat{\beta} = 0.223$. We can carry out a likelihood ratio test by comparing $C(\widehat{\theta}_{\mathrm{H}})$ with $C(\widehat{\theta}_{\beta=0})$, which is the sum of squares when the parameters are estimated under the constraint $\beta = 0$; see the third column of Table 2.2. $\mathcal{S}_{\mathrm{L}} = 32 * (\log(0.183) - \log(0.0206)) = 69.9$. This number must be compared to 3.8, the 0.95 quantile of a χ^2 with one degree of freedom. The hypothesis $\beta = 0$ is rejected.

Conclusion of the Test In this experiment, we conclude that the *potency* ρ of the serum taken in June *relative* to the serum taken in May is estimated by $\widehat{\rho} = 10^{-\widehat{\beta}} = 0.59$ and that ρ is significantly different from 1.

Calculation of a Confidence Interval for ρ The parameter of interest is $\rho = \lambda(\theta) = 10^{\theta_4^{\mathrm{June}} - \theta_4^{\mathrm{May}}}$.

Calculation of Confidence Intervals with Asymptotic Level 95%, Using Results of Section 2.3.2 (d.f. is for degree of freedom):

$\widehat{\rho}$	$\widehat{S}$	$\nu_{0.975}$	$\widehat{I}_{\mathcal{N}}$	$t_{0.975}$ (27 d.f.)	$\widehat{I}_{\mathcal{T}}$
0.59	0.0192	1.96	[0.561 , 0.636]	2.05	[0.555 , 0.642]

$\widehat{I}_{\mathcal{T}}$ is not much different from $\widehat{I}_{\mathcal{N}}$ because $n-p$ is large, $n/(n-p)$ is close to 1, and the difference between ν_α and t_α is small.

Calculation of Confidence Intervals with Asymptotic Level 95%, Using the Bootstrap Method The number of bootstrap simulations is $B = 199$. The histogram of the $\widehat{T}^{\star,b}$, $b = 1, \dots B$, is shown in Figure 2.2.

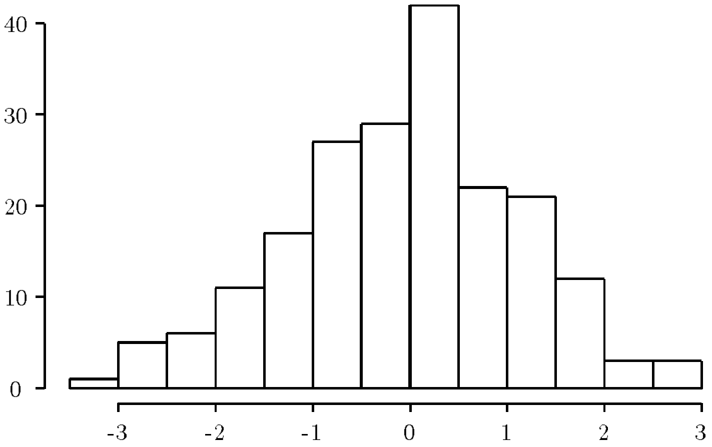

Figure 2.2. ELISA test example: Histogram of ($\widehat{T}^{\star,b}$, $b = 1, \dots B$)

The results follow:

$\widehat{\rho}$	$\widehat{S}$	$b_{0.025}$	$b_{0.975}$	$\widehat{I}_B$
0.59	0.0192	-2.61	2.17	[0.557 , 0.649]

In this example, the bootstrap shows that $\widehat{I}_B$ is longer than $\widehat{I}_{\mathcal{N}}$ but is nearly equal to $\widehat{I}_{\mathcal{T}}$. Moreover, the bootstrap estimation of the standard error of $\widehat{\rho}$ is $\widehat{S}^\star = 0.0199$. The differences between the methods are not very important from a practical point of view.

2.4.4 Ovocytes: Calculation of Confidence Regions

Although this experiment tested several cryoprotectants in different experimental conditions (with or without treatment, at several temperatures) and yielded 15 curves to be estimated and compared, we present here the results for only two curves. We compute two types of confidence regions: Confidence ellipsoids and likelihood contours.

Confidence Ellipsoids Let θ be the pair (P_w, P_s). When the number of observations tends to infinity, the limiting distribution of

$$\mathcal{S}'(\theta) = V_{\widehat{\theta}}^{-1/2}(\widehat{\theta} - \theta)$$

is a standard two dimensional normal distribution $\mathcal{N}(0, I_2)$, or the limiting distribution of

$$\mathcal{S}_{\mathrm{W}}(\theta) = \mathcal{S}_1'^2 + \mathcal{S}_2'^2$$

is a χ^2 with two degrees of freedom. Let $r_\alpha(2)$ be the α percentile of a χ^2 with two degrees of freedom, and let $\mathcal{R}_{\mathrm{W}}$ be the set of θ such that $\mathcal{S}_{\mathrm{W}}(\theta) \leq r_{1-\alpha}(2)$. $\mathcal{R}_{\mathrm{W}}$ is an ellipse that covers θ with probability close to $1 - \alpha$.

Likelihood Contours Constructing confidence ellipses is based on the limiting distribution of $\widehat{\theta} - \theta$. Another way to calculate confidence regions for θ is to consider the limiting distribution of the statistic

$$\mathcal{S}_{\mathrm{L}}(\theta) = n \log C(\theta) - n \log C(\widehat{\theta}).$$

$\mathcal{S}_{\mathrm{L}}(\theta)$ is a χ^2 with two degrees of freedom. Let $\mathcal{R}_{\mathrm{L}}$ be the set of θ such that $\mathcal{S}_{\mathrm{L}}(\theta) \leq r_{1-\alpha}(2)$. $\mathcal{R}_{\mathrm{L}}$ is a region of the plane that covers θ with probability close to $1 - \alpha$.

Figure 2.3 illustrates the confidence ellipses and the likelihood contours with level 95% for the parameters (P_w, P_s) in Example 1.1.4. In this example, the likelihood contours are close to the ellipses, but we will see that this is not always the case. In fact, the discrepancy between these regions is related to the discrepancy between the distribution of $\widehat{\theta} - \theta$ and its approximation by a centered Gaussian variable.

Remarks When p, the dimension of θ, is greater than 2, the confidence ellipsoids for θ are defined by the set of θ such that $\mathcal{S}_{\mathrm{W}}(\theta) \leq r_{1-\alpha}(p)$, where $r_\alpha(p)$ is the α percentile of a χ^2 with p degrees of freedom. Usually the sections of the regions are drawn in two dimensions, and they give conditional information. Consider the case $p = 3$. The sections of the confidence regions in the plane (θ_1, θ_2) are the sets of (θ_1, θ_2) such that $\mathcal{S}_{\mathrm{W}}(\theta_1, \theta_2, \widehat{\theta}_3) \leq r_{1-\alpha}(3)$.

2.4.5 Isomerization: An Awkward Example

Model The regression function is

$$f(x, \theta) = \frac{\theta_1 \theta_3 (P - I/1.632)}{1 + \theta_2 H + \theta_3 P + \theta_4 I},$$

and the variances are homogeneous: $\mathrm{Var}(\varepsilon_i) = \sigma^2$.

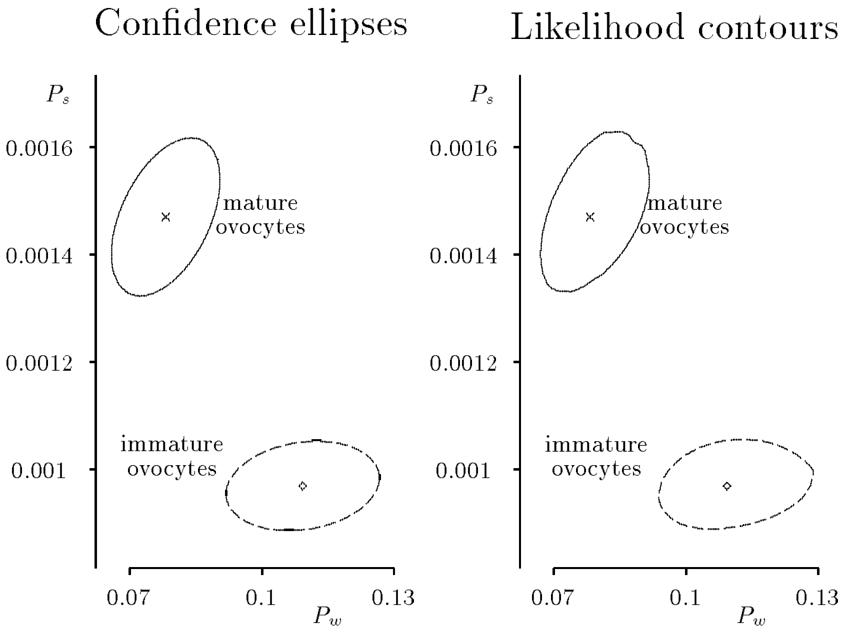

Figure 2.3. Ovocytes example: 95% confidence ellipses and likelihood contours for the parameters (P_w, P_s)

Table 2.3. Isomerization example: Estimated parameters and standard errors, normal confidence intervals

	Estimates	Standard Errors ($\widehat{S}$)	95% Confidence Interval $\widehat{I}_{\mathcal{N}}$
θ_1	35.9193	7.49	[21.20 50.60]
θ_2	0.0708583	0.163	[−0.249 0.391]
θ_3	0.0377385	0.0913	[−0.141 0.217]
θ_4	0.167166	0.379	[−0.577 0.911]
σ^2	0.13477		

Calculation of Confidence Intervals Using the Percentiles of a Gaussian Distribution for Each Parameter The estimated values of the parameters and their standard errors and of the confidence intervals calculated using Equation (2.2) are given in Table 2.3.

The standard errors are so large for parameters θ_2, θ_3, and θ_4 that the value 0 is inside the confidence intervals. Obviously, the hypothesis $\theta_2 = \theta_3 = \theta_4 = 0$ is meaningless. Figure 2.4 illustrates the discrepancy between the distribution of $\widehat{\theta} - \theta$ and its approximation by the distribution $\mathcal{N}(0, V_\theta)$. In this figure, the sections of the confidence ellipsoids and likelihood contours in the plane (θ_1, θ_2) clearly differ significantly in their appearance. Thus, we cannot use the percentiles of a Gaussian distribution to calculate the confidence intervals. Instead, let us try the bootstrap method.

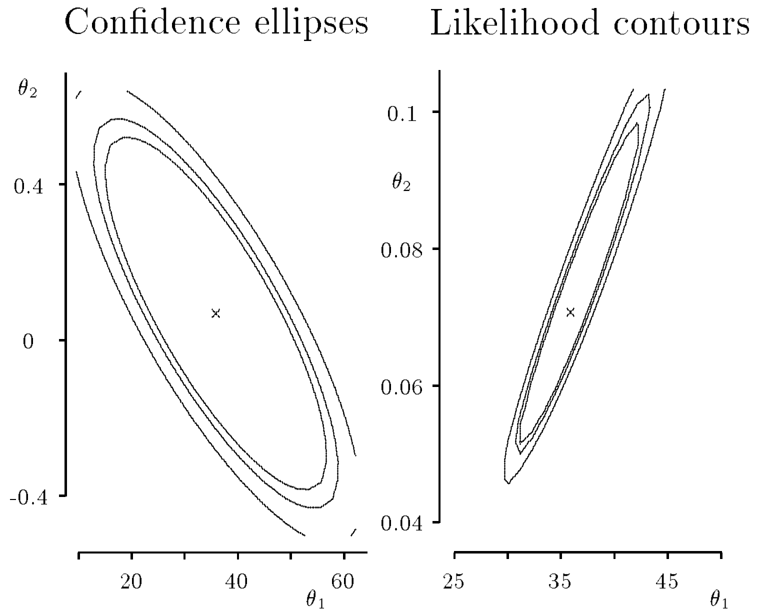

Figure 2.4. Isomerization example: The confidence ellipses and likelihood contours for the parameters (θ_1, θ_2) are drawn at levels 90%, 95%, and 99%

Calculation of Confidence Intervals Using the Bootstrap Method $B = 199$ bootstrap simulations have been done to calculate another approximation of the distribution of $\widehat{\theta}$. Let

$$\widehat{T}_a = \frac{\widehat{\theta}_a - \theta_a}{\widehat{S}_a},$$

where θ_a is the component a of θ and $\widehat{S}_a$ is the estimation of the standard error of $\widehat{\theta}_a$. For each parameter, the bootstrap estimation of the distribution of $\widehat{T}_a$ is shown in Figure 2.5. Note that the bootstrap distributions are very far from the Gaussian distribution, except for the first parameter.

Bootstrap Estimations of Standard Error and Bias for Each Parameter Estimator The results are Table 2.4.

The bootstrap estimations of the standard errors, $\widehat{S}_a^{\star}$ (see Equation (2.8)), are of the same magnitude as $\widehat{S}_a$. The bootstrap yields a high value of the bias (see Equation (2.9)) for θ_1, but the bias is small for the other parameters. The 0.025 and 0.975 percentiles of the $\widehat{T}_a^{\star,b}$ ($\widehat{T}_a^{*}$ is the bootstrap version of $\widehat{T}_a$) are calculated as in Section 2.4.1. They show the asymmetry of the estimators' distribution. By comparison, the 0.025 percentile of a Gaussian $\mathcal{N}(0,1)$

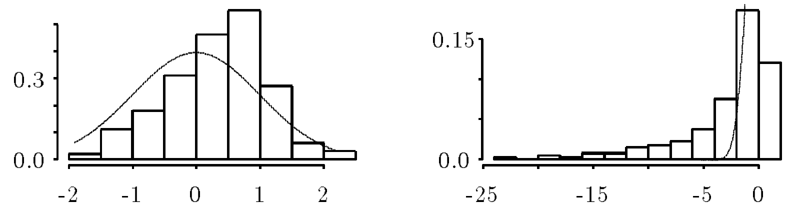

Histogram of T_1 Histogram of T_2

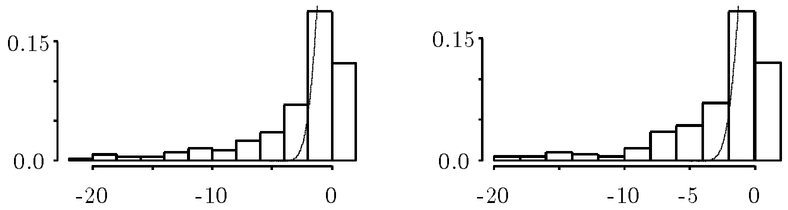

Histogram of T_3 Histogram of T_4

Figure 2.5. Isomerization example: Histogram of ($\widehat{T}^{\star,b}$, $b = 1, \ldots B$) for each parameter; the line is the probability density of an $\mathcal{N}(0,1)$

Table 2.4. Isomerization example: Bootstrap estimation of standard errors and bias for the parameters θ; 2.5% and 97.5% percentiles of the bootstrap distribution of $\widehat{T}$; bootstrap confidence intervals for the θ

	$\widehat{S}^\star$	$\widehat{\text{BIAS}}^\star$ (% of bias)	$b_{0.025}$	$b_{0.975}$	$\widehat{I}_B$
θ_1	9.83	5.47 (15)	-1.38	1.69	[23.2 46.2]
θ_2	0.133	0.002 (3)	-15.2	0.137	[0.048 2.56]
θ_3	0.080	0.002 (6)	-16.4	0.151	[0.024 1.54]
θ_4	0.322	0.008 (5)	-15.4	0.144	[0.112 6.05]

distribution is $\nu_{0.975} = 1.96$. The last column gives the bootstrap confidence intervals (see Equation 2.7). For the three last parameters the lower bound of the intervals is positive; this condition is more realistic than the negative bounds obtained with $\widehat{I}_{\mathcal{N}}$. The confidence intervals are not symmetric around the estimated values $\widehat{\theta}_a$.

Calculation of Confidence Intervals Using a New Parameterization of the Function f An alternative to the bootstrap is to find another parameterization of the function $f(x, \theta)$ that reduces the discrepancy between the distribution of $\widehat{T}$ and the approximation of $\widehat{T}$ by a Gaussian distribution.

Model A new parameterization, suggested by several authors (see [BW88], for example) is defined by considering a new set of parameters, say $(\beta_1, \beta_2, \beta_3, \beta_4)$.

These are obtained by eliminating the product $\theta_1\theta_3$ in Equation (1.9). The model function is now defined by

$$f(x,\beta) = \frac{P - I/1.632}{\beta_1 + \beta_2 H + \beta_3 P + \beta_4 I}, \tag{2.10}$$

and $\mathrm{Var}(\varepsilon_i) = \sigma^2$.

Table 2.5. Isomerization example with the new parameterization: Estimated parameters and standard errors; normal confidence intervals

	Estimates	Standard Errors ($\widehat{S}$)	95% Confidence Interval $\widehat{I}_{\mathcal{N}}$
β_1	0.73738	1.66	[−2.51 3.98]
β_2	0.052274	0.00418	[0.0441 0.0605]
β_3	0.027841	0.00581	[0.0164 0.0392]
β_4	0.12331	0.0161	[0.0917 0.155]
σ^2	0.13477		

Parameter and Standard Error Estimations The results are given in Table 2.5. The estimated accuracy of the parameters is reasonable, and the discrepancy between confidence ellipses and likelihood contours for the pair (β_1, β_2) is not as big as for the pair (θ_1, θ_2); see Figure 2.6.

Let us assume that we are interested in calculating confidence intervals for the parameters θ. We want to calculate confidence intervals for each θ_a using the confidence regions calculated for β. In our example, the relations between θ and β are easy to write:

$$\begin{aligned}
\theta_1 &= 1/\beta_3, \\
\theta_2 &= \beta_2/\beta_1, \\
\theta_3 &= \beta_3/\beta_1, \\
\theta_4 &= \beta_4/\beta_1.
\end{aligned}$$

If β_3 lies in the interval $[0.0164, 0.0392]$, then θ_1 lies in $[25.5, 60.8]$; however, we see that the same reasoning cannot be applied to θ_2, because 0 lies in the confidence interval for β_1. Thus, in this example, the calculation of a confidence interval for the parameters θ_a, $a = 1, \ldots p$, is only possible for θ_1. Only the bootstrap method allows us to calculate confidence intervals for each parameter.

2.4.6 Pasture Regrowth: Calculation of a Confidence Interval for $\boldsymbol{\lambda} = \exp \boldsymbol{\theta_3}$

Let us return to our pasture regrowth example to illustrate how to calculate a confidence interval that respects the constraint that parameter λ must be positive.

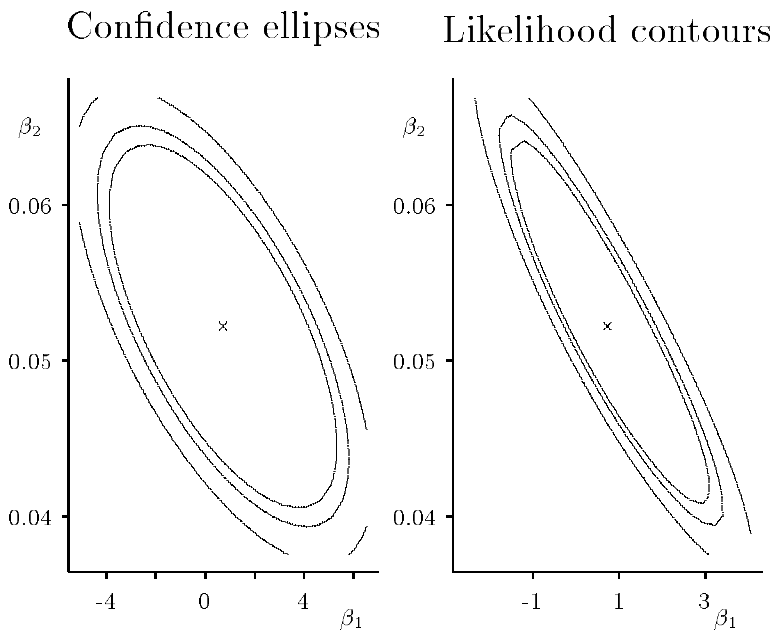

Figure 2.6. Isomerization example: The confidence ellipses and likelihood contours for the parameters (β_1, β_2) are drawn at levels 90%, 95%, and 99%

Model The regression function is

$$f(x, \theta) = \theta_1 - \theta_2 \exp\left(-\exp(\theta_3 + \theta_4 \log x)\right),$$

and the variances are homogeneous: $\mathrm{Var}(\varepsilon_i) = \sigma^2$.

Results

Parameters	Estimated Values	Asymptotic Covariance Matrix			
θ_1	69.95	3.09			
θ_2	61.68	3.87	6.66		
θ_3	-9.209	0.76	1.25	0.37	
θ_4	2.378	-0.22	-0.35	-0.09	0.027
σ^2	0.9306				

The parameter of interest is $\lambda(\theta) = \exp \theta_3$.

Calculation of Confidence Intervals with Asymptotic Level 95%, Using Results of Section 2.3.2

$\widehat{\lambda}$	$\widehat{S}$	$\nu_{0.975}$	$\widehat{I}_{\mathcal{N}}$
0.0001001	0.0000609	1.96	$[-.0000194, .000219]$

It is immediately apparent that this confidence interval is unusable because λ cannot be negative. The result would have been the same if we had estimated λ by replacing the model function (1.1) with the following one: $f(x, \theta_1, \theta_2, \lambda, \theta_4) = \theta_1 - \theta_2 \exp(-\lambda x^{\theta_4})$.

Another way to calculate a confidence interval for λ is to transform the confidence interval calculated for θ_3. If θ_3 lies in $[-10.4\ -8.01]$, then $\lambda = \exp \theta_3$ lies in $[0.0000303\ ,\ 0.000331]$. By construction, this confidence interval is adapted to the set of variation of λ.

2.5 Conclusion

We have proposed several methods to estimate confidence intervals based on approximating the distribution of $\widehat{\lambda}$ by either the Gaussian distribution or the bootstrap distribution. In some cases (see Example 2.4.1), it is more convenient to consider a monotone transformation of the parameter of interest in place of the parameter itself, because the distribution of its estimator is better approximated by a Gaussian distribution. In other cases, the bootstrap method is more appropriate (see Example 2.4.5). The examples treated in this chapter show that there is no rule to decide in advance which is the correct method. Nevertheless, in each case, the final choice was based on the adequacy of the result and the nature of the parameter of interest.

2.6 Using nls2

This section reproduces the commands and files used in this chapter to analyze the examples using **nls2**. It is assumed that the commands introduced in Section 1.6 have already been executed.

Pasture Regrowth Example: Confidence Interval for the Maximum Yield

The results of estimation have been stored in the structure `pasture.nl1` (see Section 1.6, page 19). Now we want to calculate a confidence interval for the maximum yield, the parameter θ_1.

Confidence Interval for $\lambda = \theta_1$ *with Asymptotic Level 95%*

We use the function `confidence`. This function calculates the confidence interval $\widehat{I}_{\mathcal{N}}$ defined by Equation (2.2), page 32, and the confidence interval $\widehat{I}_{\mathcal{T}}$ defined by Equation (2.3), page 32. By default, the asymptotic level is 95%.

```
> pasture.conf.par <- confidence(pasture.nl1)
```

We display the values of $\widehat{\lambda}$, $\widehat{S}$, $\nu_{0.975}$, $\widehat{I}_{\mathcal{N}}$, $t_{0.975}$ (five degrees of freedom), and $\widehat{I}_{\mathcal{T}}$:

```
> cat("Estimated value of lambda:", pasture.conf.par$psi[1],"\n" )
> cat("Estimated value of S:",pasture.conf.par$std.error[1],"\n" )
> cat("nu_(0.975):", qnorm(0.975),"\n" )
> cat("Estimated value of In:",
      pasture.conf.par$normal.conf.int[1,],"\n" )
> cat("t_(0.975, 5):", qt(0.975,5),"\n" )
> cat("Estimated value of It:",
      pasture.conf.par$student.conf.int[1,],"\n" )
```

Confidence Interval for $\lambda = \theta_1$ *Using Bootstrap with Asymptotic Level 95%*

To calculate the confidence interval $\widehat{I}_B$ defined by Equation (2.7), page 36, we use the function `bootstrap`. Several methods of bootstrap simulation are possible. Here, we choose `residuals`, which means that pseudoerrors are randomly generated among the centered residuals.

To initialize the iterative process of `bootstrap`, `nls2` must be called with the option `renls2`. We also set the option `control` so that intermediary results are not printed. Finally, we call the function `delnls2` to destroy any internal structures:

```
> pasture.nl1 <- nls2(pasture, "pasture.mod1",
            list(theta.start= c(70, 60, 0, 1), max.iters=100),
            control=list(freq=0),
            renls2=T)
> pasture.boot <- bootstrap(pasture.nl1,
             method="residuals",
             n.loops=199)
> delnls2()
```

We calculate the values of $\widehat{T}^{\star,b}$, $b = 1, \dots 199$ (see Section 2.3.5) and illustrate their distribution function by plotting them in a histogram (see Figure 2.1, page 39):

```
> P1.B <- pasture.boot$pStar[,1]
> SE.P1.B <- sqrt(pasture.boot$var.pStar[,1])
> T.B <- (P1.B-pasture.nl1$theta[1])/SE.P1.B
> hist(T.B,nclass=12,
      title="Pasture regrowth example",
      sub="Histogram of bootstrap estimations for T")
```

We calculate the 0.0025 and 0.975 percentiles of the $\widehat{T}^{\star,b}$, $b_{0.025}$ and $b_{0.975}$ (see Section 2.4.1) using the function `quantile` of S-Plus, and we display the values of $\widehat{\lambda}$, $\widehat{S}$, $b_{0.025}$, $b_{0.975}$, and $\widehat{I}_B$:

```
>    # Print the results:
> cat("Estimated value of lambda:", pasture.nl1$theta[1],"\n")
> cat("Estimated value of S:",coef(pasture.nl1)$std.error[1],"\n")
> qu  <- quantile(T.B,probs=c(0.975,0.025))
> cat("b_(0.025):", qu[2],"\n" )
```

```
> cat("b_(0.975):", qu[1],"\n" )
> cat("Estimated value of Ib:",
   pasture.nl1$theta[1]+qu[2]*coef(pasture.nl1)$std.error[1],
   pasture.nl1$theta[1]+qu[1]*coef(pasture.nl1)$std.error[1],"\n")
```

Finally, we calculate the accuracy characteristics of the bootstrap estimation: the bias ($\widehat{\mathrm{BIAS}}^{\star}$; Equation (2.9)), the variance ($\widehat{S}^{\star}$; Equation (2.8)), the mean square error ($\widehat{\mathrm{MSE}}^{\star}$), and the median ($\widehat{\mathrm{MED}}^{\star}$):

```
> cat("BIAS:" , (mean(P1.B)-pasture.nl1$theta[1]),"\n" )
> cat("S:" ,sqrt(var(P1.B)),"\n" )
> cat("MSE:" ,
    var(P1.B)+(mean(P1.B)-pasture.nl1$theta[1])^2 ,"\n" )
> cat("MED:" ,median(P1.B) ,"\n" )
```

Note: The bootstrap method generates different numbers on each execution. Thus, results of these commands may vary slightly from those displayed in Section 2.4.1, page 38.

Cortisol Assay Example: Confidence Interval for D

We calculated an estimate of the calibration curve (see Section 1.6, page 21) in structure `corti.nl1`, and now we want to calculate a confidence interval for the estimation of the dose of hormone D contained in a preparation that has the expected response $\mu = 2000$ c.p.m.

Confidence Interval for D

We describe the function D in a file called `corti.D`. The expected response μ is introduced by the key word `pbispsi`:

```
psi D;
ppsi n,d,a,b,g;
pbispsi mu;
aux X1, X2, X;
subroutine;
begin
X1 = log((d-n)/(mu-n));
X2 = exp(X1/g);
X = (log(X2-1)-a)/b;
D = 10**X;
end
```

To calculate a confidence interval for D, we apply the `confidence` function. Then we display the results of interest, $\widehat{D}$, $\widehat{S}$, $\nu_{0.975}$, and $\widehat{I}_{\mathcal{N}}$:

```
> loadnls2(psi="")
> corti.conf.D <- confidence(corti.nl1,file="corti.D",pbispsi=2000)
>    # Print the results:
```

```
> cat("Estimated value of D:", corti.conf.D$psi,"\n" )
> cat("Estimated value of S:",corti.conf.D$std.error,"\n" )
> cat("nu_(0.975):", qnorm(0.975),"\n" )
> cat("Estimated value of In:",corti.conf.D$normal.conf.int,"\n" )
```

(The results are given in Section 2.4.2, page 39.)

ELISA Test Example: Comparison of Curves

The results of estimation have been stored in the structure `elisa.nl1` (see Section 1.6, page 22). Now we want to test the *parallelism* of the May and June curves.

Wald Test with Asymptotic Level 95%

To test the parallelism of the response curves with a Wald test (see Section 2.3.3), we use the function `wald`.

First we describe the functions to be tested in a file called `elisa.wald`:

```
psi d1,d2,d3;
ppsi p1_c1, p2_c1, p3_c1, p1_c2, p2_c2, p3_c2;
subroutine;
begin
d1=p1_c1-p1_c2;
d2=p2_c1-p2_c2;
d3=p3_c1-p3_c2;
end
```

We apply the function `wald` and display the value of the statistic $\mathcal{S}_{\mathrm{W}}$ and the 0.95 quantile of a χ^2 with three degrees of freedom from which it should be compared:

```
> elisa.wald <- wald(elisa.nl1,file="elisa.wald")
>    # Print the results:
> cat("SW:",elisa.wald$statistic,"\n" )
> cat("X2(3):", qchisq(0.95, 3),"\n" )
```

Because the variances are homogeneous, we calculate the test statistic defined in Equation (2.5), page 35:

```
> SF <- (elisa.wald$statistic*(32-8))/(32*3)
> cat("SF:", SF,"\n" )
> cat("F(3,24):", qf(0.95, 3,24), "\n" )
```

Likelihood Ratio Tests

To test the parallelism of the curves by likelihood ratio tests (see Section 2.3.3), we have to estimate the parameters under hypothesis A (*the parallelism of the curves is not verified*), under hypothesis H (*the curves are parallel*), and under the last hypothesis (*the curves are identical*).

Estimation under hypothesis A has already been calculated: A is the hypothesis under which structure `elisa.nl1` has been built.

Estimation under hypothesis H is done by setting equality constraints on all of the parameters except for the last one. Equality constraints are given with the option `eqp.theta`:

```
> elisa.nlH <- nls2(elisa,
    list(file="elisa.mod1", eqp.theta=c(1,2,3,4,1,2,3,5)),
    rep(c(2,0,1,0),2))
```

Estimation under the last hypothesis is done by setting equality constraints on all of the parameters:

```
> elisa.nlb <- nls2(elisa,
    list(file="elisa.mod1", eqp.theta=c(1,2,3,4,1,2,3,4)),
    rep(c(2,0,1,0),2))
```

We display the estimated values of the parameters and the sums of squares for the three hypothesis:

```
> cat("Estimated values of the parameters for the 3 hypothesis:\n")
> print(elisa.nl1$theta)
> print(elisa.nlH$theta)
> print(elisa.nlb$theta)
> cat("Estimated sums of squares for the 3 hypothesis:\n",
   elisa.nl1$rss, "\n", elisa.nlH$rss,"\n",
   elisa.nlb$rss,"\n")
```

Now, we calculate the test statistic $\mathcal{S}_{\mathrm{L}}$ and display the 0.95 quantile of a χ^2 with one degree of freedom from which they should be compared.

We also print the estimated value of $\beta = (\theta_4^{\mathrm{May}} - \theta_4^{\mathrm{June}})$:

```
> cat("Sl:",
    32*(log(elisa.nlH$rss) - log(elisa.nl1$rss)),
    32*(log(elisa.nlb$rss) - log(elisa.nlH$rss)),"\n" )
> cat("X2(0.95,1):", qchisq(0.95,1),"\n" )
> cat("Estimated value of beta:",
     elisa.nlH$theta["p4_c1"] -  elisa.nlH$theta["p4_c2"],"\n" )
```

Confidence Interval for ρ with Asymptotic Level 95%

Now we want to calculate a confidence interval for a function of the parameters: $\rho = \lambda(\theta) = 10^{(\theta_4^{\mathrm{June}} - \theta_4^{\mathrm{May}})}$.

We describe ρ in a file called `elisa.ro`:

```
psi ro;
ppsi p4_c1, p4_c2;
subroutine;
begin
ro=10**(p4_c2-p4_c1);
end
```

The function `confidence` is applied to the structure `elisa.nlH`, which contains the results of estimation under hypothesis H (*the curves are parallel*). We display the values of $\widehat{\rho}$, the standard error ($\widehat{S}$), $\nu_{0.975}$, $\widehat{I}_{\mathcal{N}}$, $t_{0.975}$ (27 degrees of freedom), and $\widehat{I}_{\mathcal{T}}$:

```
>  elisa.ro <- confidence(elisa.nlH, file="elisa.ro")
>    # Print the results:
> cat("Estimated value of rho:", elisa.ro$psi,"\n" )
> cat("Estimated value of S:",elisa.ro$std.error,"\n" )
> cat("nu_(0.975):", qnorm(0.975),"\n" )
> cat("Estimated value of In:",elisa.ro$normal.conf.int,"\n" )
> cat("t_(0.975, 27):", qt(0.975,27),"\n" )
> cat("Estimated value of It:",elisa.ro$student.conf.int,"\n" )
```

Confidence Interval for ρ Using Bootstrap with Asymptotic Level 95%

To calculate confidence intervals for ρ with bootstrap simulations (see Section 2.3.5, page 35) we apply the function `bootstrap`.

To initialize the iterative bootstrap process, `nls2` is first called with the option `renls2`, and, finally, the function `delnls2` cleans the internal structures:

```
> elisa.nlH <- nls2(elisa,
    list(file="elisa.mod1", eqp.theta=c(1,2,3,4,1,2,3,5)),
    rep(c(2,0,1,0),2),
    control=list(freq=0),
    renls2=T)
> elisa.boot.ro <- bootstrap(elisa.nlH,method="residuals",
                  file="elisa.ro", n.loops=199)
> delnls2()
```

We display the values of $\widehat{\rho}$, $\widehat{S}$, $b_{0.025}$, $b_{0.975}$, and $\widehat{I}_B$:

```
> cat("Estimated value of rho:",elisa.ro$psi,"\n" )
> cat("Estimated value of S:", elisa.ro$std.error,"\n" )
> qu  <- quantile((elisa.boot.ro$tStar,probs=c(0.975,0.025))
> cat("b_(0.025):", qu[2],"\n" )
> cat("b_(0.975):", qu[1],"\n" )
> cat("Estimated value of Ib:", elisa.boot.ro$conf.int ,"\n" )
> cat("Bootstrap standard error:", sqrt(var(elisa.boot.ro$psiStar))
```

To illustrate the distribution function of $\widehat{T}^{\star}$, we plot a histogram of their values (see Figure 2.2, page 42):

```
> hist(elisa.boot.ro$tStar, nclass=9,
       title="ELISA example",
       sub="Histogram of bootstrap estimations for T")
```

Note: The bootstrap method generates different numbers on each execution. Thus, results of these commands may vary slightly from those displayed in Section 2.4.3, page 40.

Ovocytes Example

Confidence Ellipsoids and Likelihood Contours for the Parameters (P_w, P_s)

The results of estimation by `nls2` have been stored in the structure `ovo.nl1` (see Section 1.6, page 25). We want to compare the parameters P_w and P_s by calculating confidence ellipsoids and likelihood contours in the space of these parameters.

The functions `ellips` and `iso` are used. `ellips` returns what is necessary to plot confidence ellipsoids, and `iso` returns what is necessary to define confidence regions in a two-dimensional space of parameters. The plots themselves are drawn by the graphical functions of S-Plus:

```
> ovo.ell1 <- ellips(ovo.nl1, axis=c("Pw_c1","Ps_c1"))
> ovo.ell2 <- ellips(ovo.nl1, axis=c("Pw_c2","Ps_c2"))
> ovo.iso1 <- iso(ovo.nl1, axis=c("Pw_c1","Ps_c1"))
> ovo.iso2 <- iso(ovo.nl1, axis=c("Pw_c2","Ps_c2"))
>    # Graphical functions of Splus
> par(mfrow=c(1,2))
> plot(x=c(.06,.13),y=c(0.0008,.0017),type="n",xlab="Pw",ylab="Ps")
> contour(ovo.ell1,levels=qchisq(0.95,2),add=T,labex=0)
> contour(ovo.ell2,levels=qchisq(0.95,2),add=T,labex=0)
> text(0.1,0.0015,"mature ovocytes")
> text(0.08,0.001,"immature ovocytes")
> title("Confidence ellipses")
> plot(x=c(.06,.13),y=c(0.0008,.0017),type="n",xlab="Pw",ylab="Ps")
> contour(ovo.iso1,levels=qchisq(0.95,2),add=T,labex=0)
> contour(ovo.iso2,levels=qchisq(0.95,2),add=T,labex=0)
> text(0.095,0.0015,"mature ovocytes")
> text(0.08,0.001,"immature ovocytes")
> title("Likelihood contours")
```

(See Figure 2.3, page 44.)

Isomerization Example

We have calculated one estimate of the parameters (see Section 1.6, page 26) in the structure `isomer.nl1`, and now we want to calculate confidence intervals for each parameter.

Confidence Intervals for Each Parameter with Asymptotic Level 95%

We use the function `confidence` to calculate the confidence interval $\widehat{I}_{\mathcal{N}}$ defined by Equation (2.2), page 32, for each parameter:

```
> isomer.conf.par <- confidence(isomer.nl1)
```

We display the estimated values of the parameters, their standard errors ($\widehat{S}$), and the confidence interval $\widehat{I}_{\mathcal{N}}$:

```
> print(matrix(c(isomer.conf.par$psi,
                 isomer.conf.par$std.error,
                 isomer.conf.par$normal.conf.int[,"lower"],
                 isomer.conf.par$normal.conf.int[,"upper"]),
        ncol=4,
        dimnames=list(names(isomer.conf.par$psi),
          c("parameters","std","lower bound","upper bound" ))))
```

(Results are shown Table 2.3, page 44.)

Confidence Regions for Parameters

We use the functions `ellips` and `iso` and the graphical functions of S-Plus to plot confidence ellipsoids and likelihood contours in the space of the parameters (θ_1, θ_2):

```
> isomer.ell <- ellips(isomer.nl1, axis=c(1,2))
> isomer.cont <- iso(isomer.nl1, axis=c(1,2),
          bounds=matrix(c(25,50,0.03,0.11),nrow=2))
>    # Graphical functions of Splus
> par(mfrow=c(1,2))
> contour(isomer.ell,levels=qchisq(c(0.90,0.95,0.99),4),labex=0)
> points(x=isomer.nl1$theta[1], y=isomer.nl1$theta[2])
> title("Confidence ellipses")
> contour(isomer.cont,levels=qchisq(c(0.90,0.95,0.99),4),labex=0)
> title("Likelihood contours")
> points(x=isomer.nl1$theta[1], y=isomer.nl1$theta[2])
```

(See Figure 2.4, page 45.)

Calculation of Confidence Intervals Using Bootstrap

Confidence intervals using the bootstrap method (see Section 2.3.5, page 35) are calculated using the function `bootstrap`.

Here, to reduce the execution time, which may be long because the model must be calculated several times at each loop, we choose to evaluate by the C program rather than by syntaxical trees (see the Ovocytes Example, Section 1.6, page 24).

To generate the program that calculates the model, we type the operating system command:

```
$  analDer isomer.mod1
```

We then load the program into our S-Plus session:

```
> loadnls2("isomer.mod1.c")
```

To initialize the iterative bootstrap process, `nls2` is first called with the option `renls2`, and, finally, the function `delnls2` cleans the internal structures:

```
> isomer.nl1 <- nls2(isomer,"isomer.mod1", c(36,.07,.04,.2),
          control=list(freq=0), renls2=T)
> isomer.boot <- bootstrap(isomer.nl1,
          method="residuals", n.loops=199)
> delnls2()
```

Histograms of ($\widehat{T}^{\star,b}$, $b = 1, \ldots 199$) for Each Parameter

Histograms of $\widehat{T}^{\star}$ for each parameter illustrate the boostrap estimation of their distribution. Only the results corresponding to correct estimations, i.e., when `isomer.boot$code=0`, are taken into account:

```
> pStar_ isomer.boot$pStar[isomer.boot$code==0,]
> var.pStar_ isomer.boot$var.pStar[isomer.boot$code==0,]
> theta<-matrix(rep(isomer.nl1$theta,isomer.boot$n.loops),
          ncol=4, byrow=T)
> TT <-(pStar - theta)/  sqrt(var.pStar)
> par(mfrow=c(2,2))
> for (a in 1:4)
>  {
>  hist(TT[,a],probability=T,main="Isomerization example",
    sub=paste("Histogram of bootstrap estimations for T",a),xlab="")
>  qx<-seq(from=min(TT[,a]),to=max(TT[,a]),length=75)
>  lines(qx,dnorm(qx))
>  }
```

(See Figure 2.5, page 46.)

Bootstrap Estimations of Standard Error and Bias for Each Parameter Estimator

We calculate the accuracy characteristics of the bootstrap estimation and display the values of the standard error ($\widehat{S}^{\star}$), the bias ($\widehat{\mathrm{BIAS}}^{\star}$), the percentage of bias, the 0.0025 and 0.975 percentiles ($b_{0.025}$ and $b_{0.975}$), and the confidence interval $\widehat{I}_B$:

```
>  SE.boot <- sqrt(diag(var(pStar)))
>  bias.boot <-  apply(pStar,2,mean)-isomer.nl1$theta
>  Pbias.boot <- 100*bias.boot/isomer.nl1$theta
>  b0.025.boot <- apply(TT,2,quantile,probs=0.025)
```

```
>  b0.975.boot <- apply(TT,2,quantile,probs=0.975)
>  binf.boot <- isomer.nl1$theta -
          b0.975.boot*coef(isomer.nl1)$std.error
>  bsup.boot <- isomer.nl1$theta -
          b0.025.boot*coef(isomer.nl1)$std.error
>    # Print the results:
> print(matrix(c(SE.boot, bias.boot,Pbias.boot,
                b0.025.boot,b0.975.boot,
                binf.boot, bsup.boot), ncol=7,
        dimnames=list(names(isomer.nl1$theta),
                c("S","BIAS","% of BIAS","b0.025","b0.975",
                   "lower bound","upper bound" ))))
```

Note: The bootstrap method generates different numbers on each execution. Thus, results of these commands may vary slightly from those displayed in Section 2.4, page 46.

Confidence Intervals Using a New Parameterization of the Function f

A new parameterization of the regression function f is considered. The model of Equation (2.10), page 47, is defined in a file called `isomer.mod2`:

```
resp r;
varind H,P,I;
aux a1, a2;
parresp b1,b2,b3,b4;
subroutine;
begin
a1= P - I/1.632;
a2= b1 + b2*H + b3*P + b4*I;
r=a1/a2;
end
```

Before calling `nls2` to estimate the parameters, we have to call the function `loadnls2`. If we do not do this, the program `isomer.mod1.c`, previously loaded into the S-Plus session, will still be current. `loadnls2` is called without any argument to reset the default action; the default is to calculate the model by syntaxical trees:

```
> loadnls2(psi="")
>  isomer.nl2<-nls2(isomer,"isomer.mod2",rep(1,4))
```

Confidence intervals are calculated using the function `confidence`.

We display the estimated values of the parameters, their standard errors ($\widehat{S}$), and the 95% confidence interval $\widehat{I}_{\mathcal{N}}$:

```
> isomer.conf.par2 <- confidence(isomer.nl2)
> print(matrix(c(isomer.conf.par2$psi,
                isomer.conf.par2$std.error,
                isomer.conf.par2$normal.conf.int[,"lower"],
```

```
              isomer.conf.par2$normal.conf.int[,"upper"]),
        ncol=4,
        dimnames=list(names(isomer.conf.par2$psi),
                c("parameters","S",
                   "lower bound","upper bound" ))))
```

(Results are given in Table 2.5, page 47.)

Confidence Regions with the New Parameterization

We plot confidence ellipses and likelihood contours in the space of the parameters (β_1, β_2) using the functions `ellips` and `iso` and graphical functions of S-Plus:

```
> isomer.ell2 <- ellips(isomer.nl2, axis=c(1,2))
> isomer.iso2 <- iso(isomer.nl2, axis=c(1,2))
>    # Graphical functions of Splus
> par(mfrow=c(1,2))
> plot(x=c(-5,7),y=c(0.03,.07),type="n",xlab="b1",ylab="b2")
> contour(isomer.ell2,levels=qchisq(c(0.90,0.95,0.99),4),
        add=T,labex=0)
> points(x=isomer.nl2$theta[1], y=isomer.nl2$theta[2])
> title("Confidence ellipses")
> plot(x=c(-5,7),y=c(0.03,.07),type="n",xlab="b1",ylab="b2")
> contour(isomer.iso2,levels=qchisq(c(0.90,0.95,0.99),4),
        add=T,labex=0)
> title("Likelihood contours")
```

(See Figure 2.6, page 48.)

Pasture Regrowth Example: Confidence Interval for $\lambda = \exp\theta_3$

Let us return to the pasture regrowth example to calculate a confidence interval for $\lambda = \exp\theta_3$.

We define the function λ in a file called `pasture.lambda`:

```
psi lambda;
ppsi p3;
subroutine;
begin
lambda = exp(p3);
end
```

A confidence interval for λ is calculated using the `confidence` function. We display the values of $\widehat{\lambda}$, $\widehat{S}$, $\nu_{0.975}$, and $\widehat{I_{\mathcal{N}}}$:

```
> loadnls2(psi="")
> pasture.conf.expP3 <- confidence(pasture.nl1,
            file="pasture.lambda")
```

```
>    # Print the results:
> cat("Estimated value of lambda:", pasture.conf.expP3$psi,"\n" )
> cat("Estimated value of S:",pasture.conf.expP3$std.error,"\n" )
> cat("nu_(0.975):", qnorm(0.975),"\n" )
> cat("Estimated value of In for exp(p3):",
    pasture.conf.expP3$normal.conf.int,"\n" )
> cat("Estimated value of In for p3:",
    pasture.conf.par$normal.conf.int[3,],"\n" )
> cat("Exponential transformation of the preceding interval:",
    exp(pasture.conf.par$normal.conf.int[3,]),"\n" )
```

(Results for this example are given in Section 2.4.6, page 47.)