the interval based on the t distribution; the latter was inflated by the contributions of the large observations to the sample variance. ■

We close this section with an illustration of the use of the bootstrap in a two-sample problem. As before, suppose that $X_1, X_2, \dots, X_n$ and $Y_1, Y_2, \dots, Y_m$ are two independent samples from distributions F and G, respectively, and that $\pi = P(X < Y)$ is estimated by $\hat{\pi}$. How can the standard error of $\hat{\pi}$ be estimated and how can an approximate confidence interval for π be constructed? (Note that the calculations of Theorem A are not directly relevant, since they are done under the assumption that $F = G$.)

The problem can be approached in the following way: First suppose for the moment that F and G were known. Then the sampling distribution of $\hat{\pi}$ and its standard error could be estimated by simulation. A sample of size n would be generated from F, an independent sample of size m would be generated from G, and the resulting value of $\hat{\pi}$ would be computed. This procedure would be repeated many times, say B times, producing $\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_B$. A histogram of these values would be an indication of the sampling distribution of $\hat{\pi}$ and their standard deviation would be an estimate of the standard error of $\hat{\pi}$.

Of course, this procedure cannot be implemented, because F and G are not known. But as in the previous chapter, an approximation can be obtained by using the empirical distributions F_n and G_n in their places. This means that a bootstrap value of $\hat{\pi}$ is generated by randomly selecting n values from $X_1, X_2, \dots, X_n$ with replacement, m values from $Y_1, Y_2, \dots, Y_m$ with replacement and calculating the resulting value of $\hat{\pi}$. In this way, a bootstrap sample $\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_B$ is generated.

11.2.4 Bayesian Approach

We consider a Bayesian approach to the model, which stipulates that the X_i are i.i.d. normal with mean μ_X and precision ξ; and the Y_j are i.i.d. normal with mean μ_Y, precision ξ, and independent of the X_i. In general, a prior joint distribution assigned to (μ_X, μ_Y, ξ) would be multiplied by the likelihood and normalized to integrate to 1 to produce a three-dimensional joint posterior distribution for (μ_X, μ_Y, ξ). The marginal joint distribution of (μ_X, μ_Y) could be obtained by integrating out ξ. The marginal distribution of $\mu_X - \mu_Y$ could then be obtained by another integration as in Section 3.6.1. Several integrations would thus have to be done, either analytically or numerically. Special Monte Carlo methods have been devised for high dimensional Bayesian problems, but we will not consider them here.

An approximate result can be obtained using improper priors. We take (μ_X, μ_Y, ξ) to be independent. The means μ_X and μ_Y are given improper priors that are constant on $(-\infty, \infty)$, and ξ is given the improper prior $f_\Xi(\xi) = \xi^{-1}$. The posterior is thus proportional to the likelihood multiplied by ξ^{-1}:

$$f_{\text{post}}(\mu_X, \mu_Y, \xi) \propto \xi^{\frac{n+m}{2}-1} \exp\left(-\frac{\xi^{m+n}}{2}\left[\sum_{i=1}^{n}(x_i - \mu_X)^2 + \sum_{j=1}^{m}(y_j - \mu_Y)^2\right]\right)$$