can conclude that zero is really the best fit. You can also compute a profile (negative log-)likelihood on one particular contribution with values ranging upward from zero and see that the minimum really is at zero. However, going to all this trouble every time you have a parameter or set of parameters that appear to have their best fit on the boundary is quite tedious.

One final issue with parameters on the boundary is that the standard model selection machinery discussed in Chapter 6 (Likelihood Ratio Test, AIC, etc.) always assumes that the null (or nested) values of parameters do not lie on the boundary of their feasible range. This issue is well-known but still problematic in a wide range of statistical applications, for example, in deciding whether to set a variance parameter to zero. For the specific case of linear mixed-effect models (i.e., models with linear responses and normally distributed random variables), the problem is relatively well studied. Pinheiro and Bates (2000) suggest the following approaches (listed in order of increasing sophistication):

- Simply ignore the problem, and treat the parameter as though it were not on the boundary—i.e., use a likelihood ratio test with 1 degree of freedom. Analyses of linear mixed-effect models (Self and Liang, 1987; Stram and Lee, 1994) suggest that this procedure is conservative; it will reject the null hypothesis less often (sometimes much less often) than the nominal type I error rate α.*
- Some analyses of mixed-effect models suggest that the distribution of the log-likelihood ratio under the null hypothesis when n parameters are on the boundary is a mixture of χ^2_n and χ^2_{n-1} distributions rather than a χ^2_n distribution. If you are testing a single parameter, as is most often the case, then $n = 1$ and χ^2_{n-1} is χ^2_0—defined as a spike at zero with area 1. For most models, the distribution is a 50/50 mixture of χ^2_n and χ^2_{n-1}, which Goldman and Whelan (2000) call the $\bar{\chi}^2_n$ distribution. For $n = 1$, $\bar{\chi}^2_1(1-\alpha) = \chi^2_1(1-2\alpha)$. In this case the 95% critical value for the likelihood ratio test would thus be $\chi^2_1(0.95)/2 =$ `qchisq(0.9,1)/2`=1.35 instead of the usual value of 1.92. The `qchibarsq` function in the `emdbook` package will compute critical values for $\bar{\chi}^2_n$.
- The distribution of deviances may *not* be an equal mixture of χ^2_n and χ^2_{n-1} (Pinheiro and Bates, 2000). The "gold standard" is to simulate the null hypothesis and determine the distribution of the log-likelihood ratio under the null hypothesis; see Section 7.6.1 for a worked example.

7.5 Estimating Confidence Limits of Functions of Parameters

Quite often, you estimate a set of parameters from data, but you actually want to say something about a value that is not a parameter (e.g., about the predicted population size some time in the future). It's easy to get the point estimate—you just feed the parameter estimates into the population model and see what comes out. But how do you estimate the confidence limits on that prediction?

* Whether this is a good idea or not, it is the standard approach—as far as I can tell it is *always* what is done in ecological analyses, although some evolutionary analyses are more sophisticated.