Chapter 7

Maximum likelihood for GLMMs

7.1 Introduction

As noted in Chapter 1, creation of a GLMM by incorporating random factors in the linear predictor of a GLM leads to difficult-to-handle likelihoods. This is first laid out more carefully in a simple example and then general approaches to maximum likelihood are described.

7.2 A simple example

To fix ideas consider the following logit-normal example:

$$
\begin{aligned}
Y_{ij}|\mathbf{u} &\sim \text{indep. Bernoulli}(p_{ij}), \qquad i = 1, 2, \ldots, q;\ j = 1, 2, \ldots, n, \\
\text{logit}(p_{ij}) &= \beta x_{ij} + u_i, \\
u_i &\sim \text{indep. } \mathcal{N}(0, \sigma^2).
\end{aligned}
\tag{7.1}
$$

In this scenario there are q clusters, each with n observations, a logit link and a single random and single fixed factor. The random effects, u_i, are assumed to be i.i.d. normally distributed.

The example is so simplified it is a stretch to come up with a realistic situation it might reflect, but here is an attempt. Suppose we record $Y_{ij} = 1$ if a subject's blood pressure decreases on day j of treatment with a blood pressure drug at dose x_{ij}, and is 0 otherwise. There are q individuals tested, each at n different doses. Since the intercept is zero, when the dose is 0 and for u_i equal to its mean of zero, the probability of a decrease is 0.5. The interpretation of u_i is the person-specific propensity to decrease or increase blood pressure in response to treatment (a type of individual-specific placebo effect).

Since the model is specified conditionally, it is easiest to derive the likelihood