19
The Epilepsy Data

19.1 Introduction

In this chapter, a marginal and a random-effects approach toward modelling repeated counts will be illustrated based on the Epilepsy data, introduced in Section 2.5. We will fit a marginal GEE model (Section 19.2) as well as a generalized linear mixed model (Section 19.3), and we will extensively compare the results in Section 19.4.

Throughout this chapter, Y_{ij} represents the number of epileptic seizures patient i experiences during week j of the follow-up period. Further, as before, let t_{ij} be the time-point at which Y_{ij} has been measured, $t_{ij} = 1, 2, \ldots$ until at most 27.

19.2 A Marginal GEE Analysis

We will first perform a GEE1 analysis (Section 8.2), assuming a marginal Poisson model, with logarithmic natural link function, and with linear, treatment-specific, time-effects. More specifically, it will be assumed that

$$\begin{aligned} Y_{ij} &\sim \text{Poisson}(\lambda_{ij}), \\ \log(\lambda_{ij}) &= \begin{cases} \beta_0 + \beta_1 t_{ij} & \text{if placebo} \\ \beta_0 + \beta_2 t_{ij} & \text{if treated.} \end{cases} \end{aligned} \tag{19.1}$$

We assume a common intercept for both treatment groups in order to incorporate our prior belief that, due to the randomization, there is no sys-

TABLE 19.1. *Epilepsy Study. Parameter estimates and standard errors (empirically corrected; model-based) for the regression coefficients in Model (19.1), obtained from a GEE1 analysis with AR(1) working correlation matrix.*

Effect	Parameter	Estimate (s.e.)
Common intercept	β_0	1.3140 (0.1435; 0.1601)
Slope placebo	β_1	−0.0142 (0.0234; 0.0185)
Slope treatment	β_2	−0.0192 (0.0178; 0.0174)

tematic difference between both groups at the start of the study. Given the high number of repeated measurements (up to 27), an unstructured working correlation would require estimation of many correlation parameters. Further, the long observation period makes the assumption of an exchangeable correlation structure quite unrealistic. We therefore use the AR(1) working correlation structure, which is meaningful because we have equally spaced time points at which measurements have been taken.

Prior to the fitting of the model in (19.1), an extended model was fitted including quadratic time-evolutions, but these turned out not to be significantly different from zero ($p = 0.5239$). Therefore, from now on, we will restrict to models with linear time-effects. The analysis has been performed using the SAS procedure GENMOD. Without going into any more detail, the SAS program used for the GEE analysis for Model (19.1) is given by

```
proc genmod data=test;
class id timeclss trt;
model nseizw =  trt*time / dist=poisson;
repeated subject=id / withinsubject=timeclss
                      type=AR(1) corrw modelse;
estimate 'diff slopes' trt*time 1 -1 ;
run;
```

and we refer to Section 10.3 for more details on fitting GEE models within the SAS system.

The results of the analysis are shown in Table 19.1. The auto-correlation coefficient has been estimated as 0.5963, i.e., two measurements from the same subject one week apart have correlation equal to 0.5963. For measurements two weeks apart, the correlation is estimated to be $0.5963^2 = 0.3556$, and so on. Note that the small differences between the model-based and the empirically corrected standard errors do not lead to different conclusions with respect to hypothesis testing. None of the average time effects is significantly different from zero (empirically corrected p-values equal to 0.5429 and 0.2795 for the placebo and the treated group, respectively), nor are they significantly different from each other ($p = 0.8721$, obtained by the

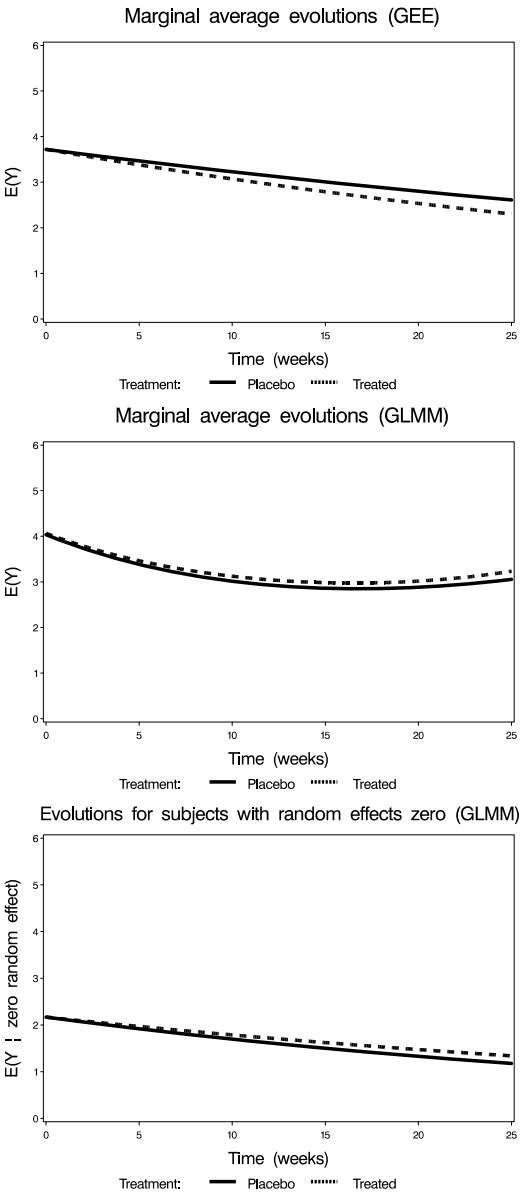

FIGURE 19.1. *Epilepsy Study. Treatment-arm specific evolutions. (a) Marginal evolutions as obtained from a marginal (GEE) model, (b) marginal evolutions as obtained from integrating out a GLMM, and (c) evolutions for an "average" subject from a GLMM, i.e., with* $\boldsymbol{b}_i = 0$.

ESTIMATE statement in the above program). Finally, the fitted average evolutions are shown in panel (a) of Figure 19.1.

19.3 A Generalized Linear Mixed Model

An alternative analysis could be based on a random-effects approach towards modeling the association structure. We then assume that, conditionally on random effects, the Y_{ij} are independent Poisson distributed random variables. As before, a logarithmic link function is used, with linear, treatment-specific, time trends. As random effects, we include random intercepts as well as random time effects. More specific, the model is given by

$$\begin{aligned} Y_{ij}|\boldsymbol{b}_i &\sim \text{Poisson}(\lambda_{ij}), \\ \log(\lambda_{ij}) &= \begin{cases} (\beta_0 + b_{i1}) + (\beta_1 + b_{i2})t_{ij} & \text{if placebo} \\ (\beta_0 + b_{i1}) + (\beta_2 + b_{i2})t_{ij} & \text{if treated,} \end{cases} \end{aligned} \tag{19.2}$$

where the random effects $\boldsymbol{b}_i = (b_{i1}, b_{i2})'$ are assumed to be normally distributed with mean vector $\mathbf{0}$ and 2×2 covariance matrix D. As in the marginal model, we assume the same fixed intercept for the two groups. This reflects our prior belief that, due to the randomization, the initial values are equally distributed in both treatment groups.

The analysis has been performed using the SAS procedures GLIMMIX and NLMIXED. First, PQL and MQL have been applied, with REML estimation for the linear mixed models for the pseudo data (Section 15.2). Afterwards, we used adaptive Gaussian quadrature with 1 and with 10 quadrature points. Note that the adaptive Gaussian quadrature with one quadrature point is equivalent to applying the Laplace approximation to the integrals in the marginal likelihood function (Section 14.5.2).

The SAS programs are given by

```
proc glimmix data=test method=RSPL;
class id trt;
model nseizw =  trt*time / dist=poisson solution;
random intercept time / type=UNR subject=id;
estimate 'diff slopes' trt*time 1 -1;
run;

proc nlmixed data=test qpoints=1;
parms beta0=1 beta1=-0.1 beta2=-0.1
      d11=1 rho=0 d22=0.1;
if (trt = 0) then eta = beta0 + b1
             + beta1*time + b2*time;
else if (trt = 1) then eta = beta0 + b1
```

TABLE 19.2. *Epilepsy Study. Parameter estimates and standard errors for the regression coefficients in Model (19.2), obtained from an MQL and PQL analysis, from an analysis based on the Laplace approximation, and from an analysis based on adaptive Gaussian quadrature with 10 quadrature points (QUAD).*

		MQL	PQL
Effect	Parameter	Estimate (s.e.)	Estimate (s.e.)
Common intercept	β_0	1.3525 (0.1492)	0.8079 (0.1261)
Slope placebo	β_1	−0.0180 (0.0144)	−0.0242 (0.0094)
Slope treatment	β_2	−0.0151 (0.0144)	−0.0191 (0.0094)
Variance of intercepts	d_{11}	1.9017 (0.2986)	1.2510 (0.2155)
Variance of slopes	d_{22}	0.0084 (0.0014)	0.0024 (0.0006)
Correlation rand.eff.	ρ	−0.3268 (0.1039)	−0.3394 (0.1294)
		Laplace	QUAD
Effect	Parameter	Estimate (s.e.)	Estimate (s.e.)
Common intercept	β_0	0.7740 (0.1291)	0.7739 (0.1293)
Slope placebo	β_1	−0.0244 (0.0096)	−0.0245 (0.0096)
Slope treatment	β_2	−0.0193 (0.0096)	−0.0193 (0.0097)
Variance of intercepts	d_{11}	1.2814 (0.2220)	1.2859 (0.2231)
Variance of slopes	d_{22}	0.0024 (0.0006)	0.0024 (0.0006)
Correlation rand.eff.	ρ	−0.3347 (0.1317)	−0.3349 (0.1318)

```
            + beta2*time + b2*time;
lambda = exp(eta);
model nseizw ~ poisson(lambda);
random b1 b2 ~ normal([0, 0],
       [d11, rho*sqrt(d11)*sqrt(d22), d22])
       subject = id;
estimate 'diff slopes' beta1-beta2;
run;
```

We refer to the Sections 15.2 and 15.4 for more details on the SAS procedures GLIMMIX and NLMIXED, respectively.

The results of our analyses are summarized in Table 19.2. We find substantial differences between the MQL and PQL methods. For example, the difference in estimates for the intercepts equals $1.3525 - 0.8079 = 0.5785$, which is large when compared to the estimated standard errors. A similar remark holds for the random-intercepts variance d_{11}. Note also the similarity of the fixed-effects estimates obtained from the MQL method and those reported in Table 19.1, obtained from fitting a marginal GEE method. This phenomenon was already observed earlier in the context of the toenail dataset (Section 16.4). Further, we find very little differences between the results from the Laplace approximation and the results from the adaptive Gaussian quadrature with 10 quadrature points. Hence, in contrast

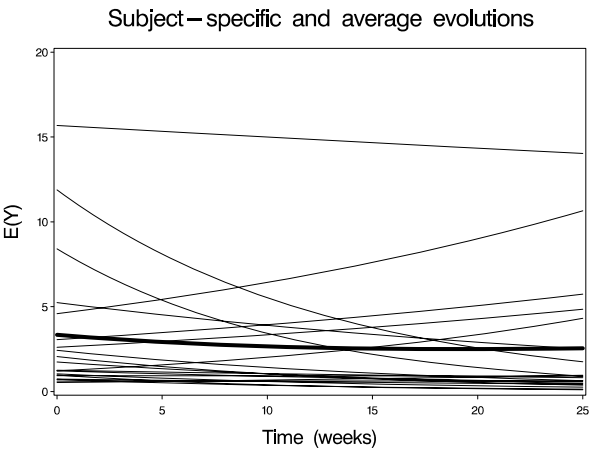

FIGURE 19.2. *Epilepsy Study. Sampled predicted profiles for 20 subjects in the placebo group (thin lines), and the resulting marginal evolution obtained from averaging over the 20 subjects (bold line).*

to earlier examples, the number of quadrature points used in the adaptive Gaussian quadrature approximation has negligable effect on the results. As was indicated in Section 14.3, this will typically be the case in datasets with many repeated measurements per subject, as in the present example. As was also observed earlier (Section 16.4), the results obtained from the PQL approach are closer to those obtained from adaptive Gaussian quadrature than those resulting from the MQL approach. Finally, in contrast to our earlier results based on the marginal GEE model, we now obtain slopes that are significantly different from zero (all p-values smaller than 0.05), unless under the MQL approach, but none of the four analyses revealed a significant difference between the slopes β_1 and β_2 (all p-values larger than 0.6).

19.4 Marginalizing the Mixed Model

As explained in Chapter 16, the regression coefficients in (19.2) need to be interpreted conditionally on the random effects $\boldsymbol{b}_i$, i.e., the parameters have a subject-specific interpretation. In case the population-averaged, marginal, evolutions are of interest, additional computations are needed. For example, the marginal expectation of the outcome Y_{ij}, measured at time-point t_{ij}, in the placebo group, is given by

$$\begin{aligned} E[Y_{ij}] &= E[E[Y_{ij}|\boldsymbol{b}_i]] \\ &= E\left[\exp[(\beta_0 + b_{i1}) + (\beta_1 + b_{i2})t_{ij}]\right] \end{aligned} \tag{19.3}$$

$$\neq \quad \exp[\beta_0 + \beta_1 t_{ij}]$$

with an expression similar to (19.3) for the expected evolution in the treated group. Calculation of (19.3) requires integrating out the random effects over their fitted distribution. As explained in Section 16.3, this can be done based on numerical integration techniques or based on numerical averaging. Here, we will follow the latter procedure, with 1000 draws for each treatment group.

As an example, let us consider the placebo group, and let the model be fitted using adaptive Gaussian quadrature with 10 quadrature points. We start by randomly drawing 1000 realized values for the random effects $\boldsymbol{b}_i$, taken from a bivariate normal distribution with mean vector zero, and with covariance matrix equal to the fitted random-effects covariance matrix (see Table 19.2)

$$D = \begin{pmatrix} 1.2859 & -0.0185 \\ -0.0185 & 0.0024 \end{pmatrix}.$$

The Cholesky decomposition of D, defined as the upper triangular matrix L such that $L'L = D$, and needed in the SAS code for drawing the 1000 random vectors $\boldsymbol{b}_i$ is given by

$$L = \begin{pmatrix} 1.1340 & -0.0163 \\ 0 & 0.0462 \end{pmatrix}.$$

For each of the 1000 realized random vectors $\boldsymbol{b}_i$, and for a fine grid of time points t, the conditional expectation $\exp[(\beta_0+b_{i1})+(\beta_1+b_{i2})t]$ is calculated, with the fixed effects β_0 and β_1 replaced by their fitted values 0.7739 and -0.0245, respectively (see Table 19.2). An estimate for the unconditional mean at a given point t in time is then obtained from averaging the 1000 conditional means, i.e.,

$$\widehat{E}[Y(t)] = \frac{1}{1000} \sum_{i=1}^{1000} \exp[(0.7739 + b_{i1}) + (-0.0245 + b_{i2})t].$$

A graphical representation of the average evolution for the placebo group is then obtained by plotting this estimate for a sufficiently fine grid of t-values. A graphical representation of this procedure is given in Figure 19.2, for 20 placebo subjects randomly drawn from the fitted model (rather than the 1000 actually used in the calculations).

The SAS code needed for the implemenation of the above procedure is given by:

```
data h;
do treat=0 to 1 by 1;
   do subject=1 to 1000 by 1;
      b1=rannor(-1);b2=rannor(-1);
```

```
      ranint=1.1340*b1; ranslope=-0.0163*b1 + 0.0462*b2;
      do t=0 to 27 by 0.1;
         if treat=0 then y=exp(0.7739+ranint
                           +(-0.0245+ranslope)*t);
         else y=exp(0.7739+ranint +(-0.0193+ranslope)*t);
         output;
      end;
   end;
end;

proc sort data=h;
by t treat;
run;

proc means data=h;
var y;
by t treat;
output out=out;
run;

proc gplot data=out;
plot y*t=treat  / haxis=axis1 vaxis=axis2 legend=legend1;
axis1 label=(h=2.5  'Time (weeks)') value=(h=1.5)
      order=(0 to 25 by 5)  minor=none;
axis2 label=(h=2.5 A=90 'E(Y)') value=(h=1.5)
      order=(0 to 6 by 1) minor=none;
legend1 label=(h=2 'Treatment: ')
        value=(h=2 'Placebo' 'Treated');
title h=3 'Marginal average evolutions (GLMM)';
symbol1 c=black i=join w=5 l=1 mode=include;
symbol2 c=black i=join w=5 l=2 mode=include;
where _stat_='MEAN';
run;
```

The result is shown in panel (b) of Figure 19.1. Note the difference between the estimated average profiles obtained from this generalized linear mixed model and those obtained earlier from a marginal GEE analysis [shown in panel (a) of the same figure]. First, the GLMM results show clear curvature in the fitted average profiles and even suggest a small increase in average number of epileptic seizures toward the end of the study. This is completely absent in the GEE profiles. A possible explanation is that the GEE model (19.1) restricts the fitted averages to be monotone functions over time. The GLMM model can accomodate non-monotonicity, through the random effects, even though the linear predictor in (19.2) is linear in time. Further, the GEE approach slightly favors the treatment

group, while the GLMM results tend to favor the placebo group (although none of the differences where found to be statistically significant). A possible explanation can be found in the fact that, as has been explained in Section 2.5, many patients leave the study after week 16. When those patients are compared to those who are still in the study at week 17, one can observe that, in the placebo group, the worst patients continue, while the opposite is true for the treated group. Hence, the two treatment groups are different with respect to the type of subjects that continue past week 16. As will be explained in Section 27.5, the GEE approach does not correct for this which may yield possibly over-optimistic conclusions about the treated group.

Finally, panel (c) in Figure 19.1 also shows the fitted evolution in both treatment groups, for 'average' patients, i.e., patients with random-effects values equal to zero. This again illustrates that the non-linearity of the link function implies that the average evolution cannot be obtained from setting the random effects in the generalized linear mixed model equal to zero, which is in contrast to the linear mixed model (Chapter 4).