

## ST 762, HOMEWORK 6 SOLUTIONS, FALL 2009

1. (a) Here, with the subject-specific model (2), rather than modeling the probabilities of having good ( $Y = 1$ ) respiratory status in the population of subjects over time on the two treatments, the probabilities that an individual subject would have good respiratory status over time are being modeled. Thus, focusing on  $\beta_3$ , whereas in the population-average model (1)  $\beta_3$  represents the difference in log odds post-baseline in the two groups,  $\beta_3$  in this model represents the typical value of the difference in log odds post-baseline were the *same* subject to receive placebo vs. treatment. The interpretation is definitely different, with  $\beta$  in the population-average model thus representing the parameter values giving the probabilities in the population and  $\beta$  in the subject-specific model being the typical values of parameters dictating individual-level probabilities.

(b)–(d) See the programs and output.

(e) The results of all the model fits are similar, but with some definite difference; most pronounced are the difference in estimation of the random effect variance  $D$ . From the documentation for `proc glimmix`, with the `mspl` option, this procedure should be doing something similar to what the `glimmix` macro with `method=ml` does, and with the `quad` option should be doing something similar to what `proc nlmixed` with the default adaptive quadrature is doing.

Because the interpretation of  $\beta$  is different for models (1) and (2), there is no reason to expect the estimates to be similar, as different quantities are being estimated in each model. Comparing the results to those from fitting the population-average model (1), we see that the estimates of  $\beta_3$  are around 1.0 for the population-average fits and greater (between 1.4 and 2.0) for the subject-specific fits, and, in each case, the evidence to reject the hypothesis that  $\beta_3 = 0$  is overwhelming. Clearly, the treatment has an effect post-baseline both on the population-level probability of good respiratory status and on the within-individual probability.

2. See the programs, log files (parts (a) and (c)), and output for all three fits. Note from the log file for the fit using `proc nlmixed` that, although convergence was declared, it is not entirely satisfactory. This is not surprising – `proc nlmixed` is carrying out 3-dimensional numerical integration.

The results are for the most part very similar for estimation of all parameters. The estimates of  $\theta$  in (b) and (c) suggest that there is intra-individual constant coefficient of variation.