

Analysis of a longitudinal ordinal response clinical trial using dynamic models.

P.J. Lindsey[†]

J. Kaufmann[‡]

May 8, 2002

Abstract: In many areas of pharmaceutical research, there has been increasing use of categorical data and more specifically ordinal responses. In many cases, complex models are required to account for different types of dependencies present among the responses.

The clinical trial considered here involved patients who were required to remain in a particular state in order to enable the doctors to examine their heart. The aim of this trial was to study the relationship between the dose of the drug administered and the time spent by the patient in the state permitting examination. The patient's state was measured every second by a continuous Doppler signal which was categorised by the doctors into one of four ordered categories. Hence, the response consisted of repeated ordinal series of different lengths because the drug effect did not last the same length of time for all patients.

A general method for generating new ordinal distributions is presented, which is flexible enough to handle unbalanced ordinal repeated measurements. It consists in obtaining a cumulative mixture distribution from a Laplace transform and introducing into it the integrated intensity of a binary logistic, continuation-ratio, or proportional odds model. Then, a multivariate distribution is constructed by a procedure similar to the updating process of the Kalman filter. Several types of history dependencies are proposed.

Keywords: Categorical data, compartment models, gamma distribution, history dependencies, Laplace transform, likelihood, logistic distribution, mixture distribution, non-linear regression models, repeated measurements.

[†]Biostatistics, Limburgs Universitair Centrum, Diepenbeek (Email: plindsey@luc.ac.be)

[‡]Schering AG, Diagnostics-Biometry, Berlin (Email: joerg.kaufmann@schering.de)

1 Introduction

To a certain extent, collecting observations always leads to discretisation. In most cases, a measurement scale must be determined to be able to collect data. This may be arbitrarily chosen or determined by the precision of the instrument used to perform the measurement. The lack of precision during the recording process is usually disregarded if the discretisation leads to many ordered categories.

Often, either a response variable is assumed to be continuous or the number of categories is reduced to two, leading to a binary response. But until recently, it has been rare to analyse series of ordinal responses, with more than two categories (Lang and Agresti, 1994; Glonek and McCullagh, 1995).

During the 1970's, the introduction of generalised linear models by Nelder and Wedderburn (1972) led to a wider range of models for continuous as well as discrete responses. But these models were not meant to take into account the dependence among related observations. Thus, many extensions have since been proposed to introduce different types of dependencies corresponding to various relationships among observations.

Recent research has led to various other methods of introducing dependencies among related observations, for example, dynamic models for time series (Harvey, 1989), mixture distributions in survival analysis (Hougaard, 1986; Yue and Chan, 1997) to introduce a frailty dependence, copulas (Joe, 1997; Nelsen, 1999), and so on.

McCullagh and Nelder (1983) showed how independent ordinal responses could be modelled using standard generalised linear model software. This was possible by specific re-parameterisations of the logistic distribution such as the continuation-ratio and proportional odds models. A few binary dependency models have been extended to allow more than two categories. For example, the binary model of Conaway (1990) has been extended by Crouchley (1995) and Ten Have (1996) to allow modelling of ordinal response; the Rasch model has been generalised to ordinal responses by Conaway (1989) and Agresti and Lang (1993).

This particular case of dependencies among ordinal responses will be considered here. A general family of dynamic models is derived by adapting methods for continuous non-normal responses, especially survival techniques. This enables unbalanced repeated measurements to be modelled with a wide range of possible dependencies.

Section 2 describes a large longitudinal ordinal clinical trial. Section 3 develops a general method to obtain new categorical distributions and various types of dependencies. Section 4 illustrates the use of one particular new distribution with several different types of dependencies through the analysis of this longitudinal ordinal clinical trial. Finally, Section 5 discusses this new method and its use in practice.

2 The clinical trial

The clinical trial considered here involved patients who required a heart examination. Doctors had to be able to assess the left ventricular volume and ejection fraction for each patient by intermittent harmonic colour Doppler. This is only possible if an enhancement of colour Doppler images in the heart of each patient is induced by a drug administered continuously over time (intravenous infusion). For confidentiality reasons, the drug under study in this clinical trial and any revealing characteristics cannot be mentioned.

2.1 Design

The trial involved 85 patients randomised into three groups: 17, 30, and 38 patients respectively receiving a 2.5, 4, and 8 gram dose over time of the same drug all at the same constant concentration of 300 mg per ml.

The intravenous infusion started at time zero with rate of 1ml per minute. Depending on the patient's response to the drug, this infusion rate could either be decreased to 0.5ml per minute, or increased to 2 or 4ml per minute at prespecified times. This was necessary for ten patients (considered below) in order to reach an appropriate Doppler signal and maintain them at it. Because the infusion rate is endogenous, an appropriate method would be to specify a bivariate model for this rate and the Doppler signal but this is not considered here.

Although the Doppler signal could have been recorded as a continuous variable, the doctors discretized it into one of four ordered categories: no signal, an insufficient signal, an appropriate signal, and an excess signal.

The signal was observed frequently and regularly over time resulting in a total of 19690 observations being collected. Patients have series of different lengths due to the different rates of drug intake and the effects wearing off differently among the patients. The series have a maximum length of 781, 2089, and 1643 seconds (ie. about 13, 35, and 27 minutes) respectively for the 2.5, 4, and 8 grams dose groups.

The responses are also unbalanced with respect to the chosen observation time points. These were chosen very close together in time (a few seconds apart) at the start of the trial, slightly further apart (up to thirty seconds) during the period where an appropriate signal level was observed, and finally far apart (up to sixty seconds or more) when the drug effect started to wear off.

2.2 Aim

The purpose of the trial was to investigate the relationship between the dose of the drug administered and the amount of time the patients spent within the required signal range.

It had already been established that the drug under investigation provides at least 2.5 minutes of usable signal enhancement for an intravenous intake dose of 2.5 grams. For larger doses of the drug, the duration of diagnostically usable signal enhancement was assumed to behave proportionally

to the dose administered. In other words, the doctors were hoping that the amount of time spent by patients in the appropriate range of Doppler signal would be greater than 4 and 8 minutes for respective doses of 4 and 8 grams.

Hence, this led to at least the following three questions.

- Is there a difference between the dose groups?
- Is the difference between the dose groups proportional to the time spent in the appropriate range of Doppler signal?
- Is there an interaction between dose group and time (or are the three dose groups parallel over time)?

Each of these points is investigated and answered during the analysis presented in Section 4.

2.3 Visualisation

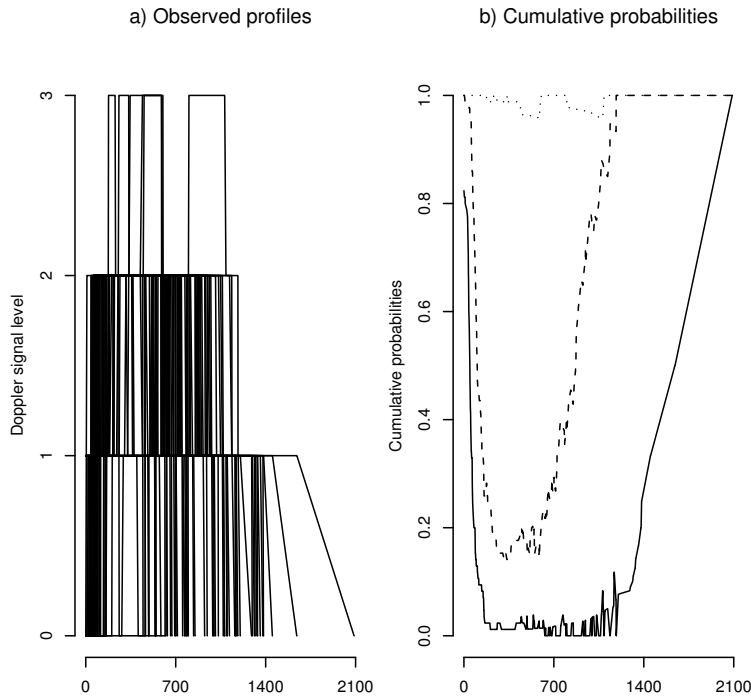
In contrast to continuous longitudinal responses, an individual profile plot of longitudinal ordinal data is not very informative. Indeed, due to the small number of different possible levels of outcome, most profiles overlap as can be seen in Figure 1a. Hence, a different type of graph such as a plot of the cumulative probabilities over time can be more informative. The data are shown in this way in the remaining plots of Figure 1. Such plots are quite different from individual profile plots because they do not represent the data at an individual level. Indeed, the outcomes of a particular individual cannot be followed over time. On the other hand, they give a general idea of the patients' distribution over the possible outcomes.

Figure 1b represents the change over time of the probability of being in a particular category. Most patients have no Doppler signal during the first 50 seconds of the trial. Then, patients have a rapid increase in signal level. During the next 50 seconds patients have an insufficient signal level. After 100 seconds, the majority of the patients remain at an appropriate signal level for almost 13 minutes. From then on, a patient's signal level slowly disappears.

The cumulative probabilities over time are also plotted separately for the three dose groups. Patients who had a larger dose of the drug administered (Figures 1d and 1e) have longer observed series and also spend a longer time at the appropriate signal level compared to the reference dose group of 2.5g (Figure 1c).

Finally, it can also be seen that more variability appears in the two smaller dose groups (2.5g and 4g) once the drug effect starts wearing off (roughly between 500 and 1200 seconds after study initiation). This can most certainly be attributed to the constant unit drug infusion rate imposed in the 2.5 grams dose group and the reduced number of patients who had their drug infusion rate changed over time in the 4 grams dose group compared to the 8 grams dose group. Indeed, it is sensible to think that tuning the infusion rate would be more often necessary when greater doses of a drug are administered because patients are expected to react faster and for longer periods of

All patients



Patients sorted by volume group




Figure 1: a) Observed profiles for all patients. Cumulative probability plots for b) all patients, c) the 2.5g dose group, d) the 4g dose group, and e) the 8g dose group. The area below the solid line indicates no signal, that between the solid and the dashed line, an insufficient signal, between the dashed and dotted line, an appropriate signal, and above the dotted line, an excess signal.

time. Out of the 85 patients, only ten (two and eight respectively in the 4g and 8g dose groups) had a non-unit or non-constant drug infusion rate over time as shown in Figure 2.

3 Model concepts

3.1 A gamma mixture distribution

Mixture distributions have been used in many areas (such as survival analysis, times to event, repeated measurements) for continuous responses or count data. This method has proved to be a flexible way to introduce frailty, especially in the case of non-Gaussian distributions. However, with the exception of the references cited above, little has been done in the case of ordinal data.

A mixture distribution, $f_m(y)$, can be obtained by integration of the product of a density and a conditional density.

$$f_m(y) = \int_0^\infty f(y|\lambda)f(\lambda)d\lambda \quad (1)$$

The density $f(\lambda)$ is called the *mixing distribution* and the density $f(y|\lambda)$ is conditional on the random parameter of this mixing distribution. This can be interpreted as the parameter λ of $f(y|\lambda)$ varying randomly in a population according to $f(\lambda)$.

In the case of discrete data, the densities do not always have a closed form. This can be overcome by considering instead the cumulative distribution function.

Under regularity conditions, Equation (1) can be rewritten in terms of the conditional cumulative distribution, $F(y|\lambda)$, and cumulative mixing distribution, $F(\lambda)$, which will yield the mixture cumulative distribution function, $F_m(y)$.

$$F_m(y) = \int_0^\infty F(y|\lambda)dF(\lambda) \quad (2)$$

A tractable form arises if a further step is made by using the relationship between a cumulative distribution function and its survival function,

$$\begin{aligned} F(y|\lambda) &= 1 - S(y|\lambda) \\ &= 1 - e^{-H(y|\lambda)} \end{aligned}$$

where $S(y|\lambda)$ is the conditional survival function and $H(y|\lambda)$ is the conditional integrated hazard function. Suppose that the parameter acts multiplicatively on the integrated hazard

$$H(y|\lambda) = \lambda H(y) \quad (3)$$

as in Hougaard (1984), where $H(y)$ is the population integrated hazard.

Now, Equation (2) can be written in terms of the conditional survival function which produces the mixture survival function,

$$S_m(y) = \int_0^\infty S(y|\lambda)dF(\lambda)$$

Patients with non-unit or non-constant infusion rate over time.

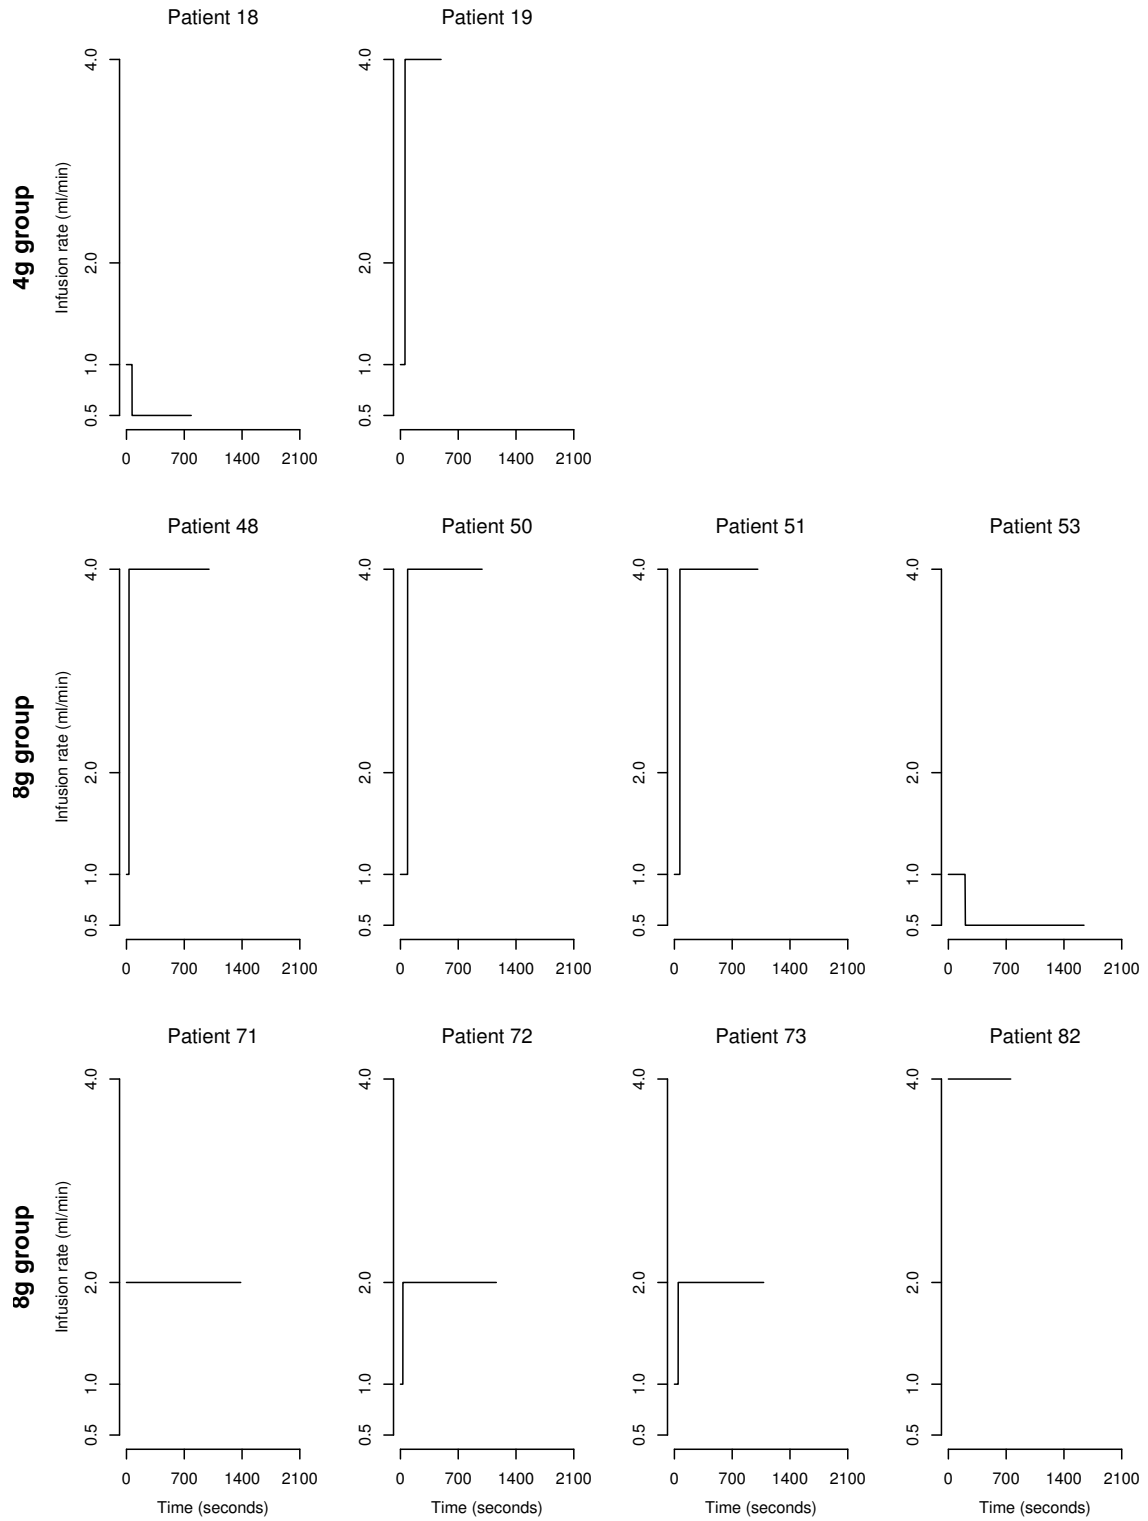


Figure 2: Drug infusion rate profiles for the ten patients who either required changes over time or received the drug at a different infusion rate than 1ml per minute.

$$\begin{aligned}
&= \int_0^\infty e^{-\lambda H(y)} dF(\lambda) \\
&= \phi[H(y)]
\end{aligned} \tag{4}$$

where $\phi[H(y)]$ is a Laplace transform (Hougaard, 1984; Joe, 1997, pp. 85–86).

Equation (4) is a useful alternative method to obtain a mixture for categorical as well as for continuous responses because the integrated intensity can be obtained by

$$H(y) = -\log[1 - F(y)]$$

The gamma Laplace transform or gamma mixture survival function corresponding to the two parameter gamma distribution function, $f(z) = \Gamma^{-1}(\alpha)\beta^\alpha y^{\alpha-1} e^{-y\beta}$, is

$$S_m(y) = \left(\frac{\beta}{\beta + H(y)} \right)^\alpha \tag{5}$$

This corresponds to the survival function of the Pareto distribution (Cox and Oakes, 1984, p. 20) for the transformation $H(y)$. This relationship as well as additional details in the case of continuous responses are described in Lindsey (2000).

Finally, Equation (5) can be written in terms of the cumulative distribution function,

$$F_m(y) = 1 - \left(\frac{\beta}{\beta - \log[1 - F(y)]} \right)^\alpha \tag{6}$$

where $F(y)$ will be chosen for this paper to be an ordinal parameterisation of the logistic cumulative distribution function, the proportional odds model which contrasts lower to higher categories around various cut points.

The cumulative probabilities corresponding to outcome i are

$$p_{ik} = \frac{e^{\gamma_k + g(\boldsymbol{\theta}, \mathbf{x})}}{1 + e^{\gamma_k + g(\boldsymbol{\theta}, \mathbf{x})}}$$

where γ_k is the intercept coefficient of level k , $g(\cdot)$ is some linear or non-linear regression function as in Subsection 4.1, and the vectors $\boldsymbol{\theta}$ and \mathbf{x} respectively correspond to the coefficients and covariates for this regression equation.

The probability of observation i being in the response category k is then

$$\begin{aligned}
\pi_{i1} &= p_{i1} \\
\pi_{ik} &= p_{ik} - p_{i,k-1} && \text{if } 1 < k < K \\
\pi_{iK} &= 1 - p_{i,K-1}
\end{aligned}$$

for the proportional odds model, where K indicates the last response category.

As already stated, the mixture density does not have a closed form in the case of categorical responses. Therefore, the probabilities corresponding to the cut-off point of each category must be obtained by taking the appropriate differences among the mixture cumulative probabilities. For example, in the proportional odds case, if an individual's response i is observed at level k , then

the original cumulative distribution function corresponds to

$$F(y_{ik}) = \sum_{l=1}^k \pi_{il} = p_{ik}$$

The mixture cumulative probabilities can then be written as

$$F_m(y_{ik}) = 1 - \left(\frac{\beta}{\beta - \log[1 - p_{ik}]} \right)^\alpha$$

Finally, the mixture probability of observation i being in the response category k is

$$\begin{aligned} \pi_{m_{i1}} &= 1 - \left(\frac{\beta}{\beta - \log[1 - p_{i1}]} \right)^\alpha \\ \pi_{m_{ik}} &= \left[1 - \left(\frac{\beta}{\beta - \log[1 - p_{ik}]} \right)^\alpha \right] - \left[1 - \left(\frac{\beta}{\beta - \log[1 - p_{i,k-1}]} \right)^\alpha \right] \\ &= \left(\frac{\beta}{\beta - \log[1 - p_{i,k-1}]} \right)^\alpha - \left(\frac{\beta}{\beta - \log[1 - p_{ik}]} \right)^\alpha && \text{if } 1 < k < K \\ \pi_{m_{iK}} &= \left(\frac{\beta}{\beta - \log[1 - p_{i,K-1}]} \right)^\alpha \end{aligned}$$

In Gaussian theory, the mixing distribution has an additive effect and its location is held to zero. In the present case, from Equation (3), the mixing distribution has a multiplicative effect which requires its location to be held to unity ($\frac{\alpha}{\beta} = 1$).

Thus, Equation (6) can be re-parameterised, by letting $\alpha = \beta = \frac{1}{\delta}$,

$$F_m(y) = 1 - \left(\frac{1}{1 - \delta \log[1 - F(y)]} \right)^{\frac{1}{\delta}} \quad (7)$$

This yields the original cumulative distribution function $F(y)$ as δ tends to zero.

Equation (7) can now be used to model independent observations. The parameter δ introduced by the gamma mixing distribution will act as an over-dispersion parameter.

3.2 History dependencies

Dependence among succeeding observations of an individual can be introduced to obtain a multivariate distribution. Different types of history or auto-regression dependencies are available (Harvey, 1989, pp. 350–358; Lindsey, 1999, pp. 69–72). A particular structure can therefore be chosen to allow the individual variability to be taken into account in the model in a desired way.

We shall require indices for the two parameters, α and β , of the gamma mixing distribution of Equation (6):

$$F_M(y_{ijk} | y_{1jk}, \dots, y_{i-1,jk}) = 1 - \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - F(y_{ijk})]} \right)^{\alpha_{i-1,j}}$$

where i indicates a particular observation of individual j at level k . We can now dynamically-update the process, with the parameter α_{ij} counting the number of observations and β_{ij} summing up the amount of information brought to the likelihood by every new observation on an individual.

Three of these dependencies are presented. The first will be called a *cumulative update* and is obtained by

$$\begin{aligned}\alpha_{ij} &= \alpha_{i-1,j} + 1 \\ \beta_{ij} &= \beta_{i-1,j} - \log[1 - F(y_{ijk})]\end{aligned}$$

The initial state of α_{ij} and β_{ij} is also required. It is defined as $\alpha_0 = \beta_0 = \frac{1}{\delta}$ to ensure that the mean of the mixing distribution is held to unity in the initial state.

More complex dependencies among successive observations can be achieved by allowing the two parameters, α and β , to depend on previous observation with a fading effect as distance increases between following observation.

Hence, a *serial update* can be obtained by

$$\begin{aligned}\alpha_{ij} &= (1 - \omega^{z_{ij}})\delta + \omega^{z_{ij}}\alpha_{i-1,j} + 1 \\ \beta_{ij} &= (1 - \omega^{z_{ij}})\delta + \omega^{z_{ij}}\beta_{i-1,j} - \log[1 - F(y_{ijk})]\end{aligned}$$

where ω can be called the discount parameter because it reduces the influence of previous observations, z_{ij} represents the elapsed time between observation i and $i - 1$ of individual j , and δ is the initiation parameter ($\alpha_0 = \beta_0 = \frac{1}{\delta}$).

Similarly, a *Markov update* can be obtained by

$$\begin{aligned}\alpha_{ij} &= (1 - \omega^{z_{ij}})\delta + \omega^{z_{ij}}\alpha_{i-1,j} + 1 \\ \beta_{ij} &= \delta - \omega^{z_{ij}} \log[1 - F(y_{i-1,jk})] - \log[1 - F(y_{ijk})]\end{aligned}$$

In this case, the parameter β depends only on its previous value discounted as a function of the distance between the two observations, rather than an accumulation of all the individual's previous values.

Independence is obtained for both of these dependencies by setting $\omega = 0$. The cumulative dependence is obtained for the serial update when $\omega \rightarrow 1$.

Now, the cumulative probabilities corresponding to outcome i from individual j are

$$p_{ijk} = \frac{e^{\gamma_k + g(\boldsymbol{\theta}, \boldsymbol{x})}}{1 + e^{\gamma_k + g(\boldsymbol{\theta}, \boldsymbol{x})}} \quad (8)$$

where γ_k is the intercept coefficient of level k , $g(\cdot)$ is some linear or non-linear regression function as in Subsection 4.1, and the vectors $\boldsymbol{\theta}$ and \boldsymbol{x} respectively correspond to the coefficients and covariates for this regression equation.

Hence for a proportional odds model, if response i from individual j is observed at level k then the original cumulative distribution function corresponds to

$$F(y_{ijk}) = \sum_{l=1}^k \pi_{ijl} = p_{ijk}$$

Then the multivariate probability of observation i from individual j being in response category k can be obtained by taking the appropriate differences, as for the cumulative mixture.

$$\begin{aligned}
\pi_{M_{ij1}} &= 1 - \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ij1}]} \right)^{\alpha_{i-1,j}} \\
\pi_{M_{ijk}} &= \left[1 - \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ijk}]} \right)^{\alpha_{i-1,j}} \right] - \left[1 - \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ij,k-1}]} \right)^{\alpha_{i-1,j}} \right] \\
&= \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ij,k-1}]} \right)^{\alpha_{i-1,j}} - \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ijk}]} \right)^{\alpha_{i-1,j}} \quad \text{if } 1 < k < K \\
\pi_{M_{ijK}} &= \left(\frac{\beta_{i-1,j}}{\beta_{i-1,j} - \log[1 - p_{ij,K-1}]} \right)^{\alpha_{i-1,j}}
\end{aligned}$$

Finally, the likelihood is obtained by multiplying together the appropriate multivariate probabilities where p_{ijk} is given by Equation (8) which contains the (non-)linear regression equation.

4 Analysis

Because the modelling process is exploratory, the inference criterion used for comparing the models under consideration is their ability to fit the observed data, that is how probable they make the data. In other words, models are compared directly through their minimised minus log likelihood. When the numbers of parameters in models differ, they are penalised by adding the number of estimated parameters, a form of the Akaike information criterion (AIC, see Akaike, 1973; Lindsey and Jones, 1998). Smaller values indicate more preferable models. This criterion allows direct comparisons among models, that are not required to be nested.

4.1 Model building

To begin, an independent multinomial regression was fitted. This null model just contains the three intercept parameters and has an AIC of 19720. This provides us with a reference point for comparison with further fitted models.

From the plots in Figure 1, it is clear that the cumulative probabilities have a somewhat parabolic shape over time. To start, a reasonable regression curve might therefore be a second-degree polynomial in time. This five-parameter proportional odds model, still with independence among observations, lowers the AIC to 16591.

However, it is also clear from Figure 1 that the curve is not symmetric. Hence, non-linear regression models must now be considered. Among several regression equations fitted, a sum of exponentials

$$g(\boldsymbol{\theta}, \boldsymbol{x}) = \theta_1 \times \left(e^{1 - \frac{\text{time}}{\theta_2}} + e^{\frac{\text{time}}{\theta_3}} \right)$$

fits better than the second-degree polynomial, where \boldsymbol{x} refers to whatever covariates are in the model. Indeed, this six parameter model lowers the AIC to 14811. This regression curve is similar to a first-order one-compartment model used in pharmacokinetics where θ_1 would correspond to

a volume parameter, θ_2 to the absorption rate, and θ_3 to the elimination rate. (All non-linear regression curves are presented without the intercepts; the appropriate intercept for each given ordinal category must be added.)

A model where the current infusion rate is linearly added to the regression equation has seven parameters and an AIC of 14669. This model is further improved by also linearly adding the dose groups, lowering the AIC to 14620 with nine parameters.

Next, the various dose groups can be allowed to evolve differently over time. This is introduced into the model by creating interactions between the dose groups and time. Each dose group requires a different curve. This fourteen-parameter model further lowers the AIC to 14448:

$$g(\boldsymbol{\theta}, \mathbf{x}) = \theta_1 \times \left(e^{1 - \frac{\text{time}}{\theta_2}} + e^{\frac{\text{time}}{\theta_3}} \right) + \theta_4 \times \text{rate} + \theta_5 \times \text{dose4} + \theta_6 \times \text{dose8} \\ + \theta_7 \times \text{dose4} \times \left(e^{1 - \frac{\text{time}}{\theta_8}} + e^{\frac{\text{time}}{\theta_9}} \right) + \theta_{10} \times \text{dose8} \times e^{1 - \frac{\text{time}}{\theta_{11}}}$$

where θ_1 to θ_3 describes the change over time for the reference (2.5g) dose group, θ_4 is the infusion rate parameter, θ_5 is the 4g dose group, θ_6 is the 8g dose group, parameters θ_7 to θ_9 describe the change over time for the 4g dose group, and θ_{10} and θ_{11} describe the change over time for the 8g dose group. Note that for patients in the 8g dose group, the drug wears off no differently than for patients in the 2.5g dose group.

All these models assumed independence among the response observations. An over-dispersion model introduces an additional parameter, but on the other hand, the interaction involving the 4g dose group can now be simplified. Indeed, it appears that the drug actually wears off at a same speed for all patients regardless of the dose group they were randomised to. This results in a thirteen-parameter model with a slightly lower AIC of 14441:

$$g(\boldsymbol{\theta}, \mathbf{x}) = \theta_1 \times \left(e^{1 - \frac{\text{time}}{\theta_2}} + e^{\frac{\text{time}}{\theta_3}} \right) + \theta_4 \times \text{rate} + \theta_5 \times \text{dose4} + \theta_6 \times \text{dose8} \\ + \theta_7 \times \text{dose4} \times e^{1 - \frac{\text{time}}{\theta_2}} + \theta_8 \times \text{dose8} \times e^{1 - \frac{\text{time}}{\theta_3}} \quad (9)$$

Now, we introduce different types of dependencies among successive observations. In each case, the regression is identical to that of Equation (9), the over-dispersion model. The cumulative model yields an AIC of 12858 with thirteen parameters, the serial dependence yields an AIC of 10618 with fourteen parameters, and the Markov dependence yields an AIC of 13426 with fourteen parameters, as can be seen in Table 1.

4.2 Results

The lowest AIC is obtained using the serial update to introduce dependence among the succeeding observations. The dependence parameter is estimated to be 0.95, indicating a very strong dependence on the previous responses. Unfortunately, this dependence close to 1 combined with a particular feature of this trial creates an undesirable effect.

Table 1: AICs for the different steps of the model building.

	AIC	Parameters
Multinomial	19720	3
Prop. Odds	14448	14
Over-dispersion	14441	13
Serial	<i>10618</i>	14
Markov	13426	14
Cumulative	12858	13

Prop. Odds: proportional odds with covariates.

Looking closely at the individual series, we notice that, due to the high frequency of observation, the series consist of long sequences of values at a particular signal level with very few changes. Due to this lack of variability, the high dependence is enough to predict the outcome at a given time point from the previous observation without any regression model. Because the model is dynamic, a wrong prediction is rapidly overcome by automatically readjusting the model at the following prediction. Hence, the underlying population mean curve is unnecessary and can simply be described by an intercept model.

Due to the presence of a strong dependence and the lack of variability among the responses in this particular trial, we decided that the serial update is not suitable to answer the doctors' questions. Among the remaining models, the cumulative model has the lowest AIC (12858).

The parameter estimates from this model, and their standard errors, are presented in Table 2. The infusion rate coefficient (θ_4) is positive implying that an increased rate leads to higher response levels. The sum of exponentials used to induce non-linear regression curves has the time covariate measured in seconds explaining the relatively large values for coefficients θ_2 , θ_3 , and θ_9 which involve times up to 2089 seconds. The coefficients for the 4g (θ_5) and 8g (θ_6) dose groups main effects are negative but the interaction coefficients are positive. This implies that the two higher dose groups will have intercepts slightly lower than the reference (2.5g) dose group. But they will also increase faster to an appropriate signal level. Finally, the dependence introduced by the cumulative update is estimated to be 0.17. No standard errors are presented in Table 2 for this dependence parameter because a transformation was actually estimated, along with corresponding standard errors, in order to ensure that this parameter would remain in the unit interval.

Because a dynamic dependence has been created among the repeated observation, two types of curves can be obtained: individual or recursive fitted values and population or average predictions. Indeed, a recursive curve is fundamentally different from an average curve because it is not representative of the entire population but of a particular patient. This is due to the dynamic part of the model which readjusts the recursive curve at each time point accordingly to the previous response of the patient under consideration. Hence, recursive curves are necessarily individual specific.

Table 2: Estimates for the proportional-odds model with cumulative dependence and non-linear regression specified by Equation (9).

	Estimates	Standard Errors
Main effects (Intercepts)		
Level 1	-12.88	0.161
Level 2	-10.00	0.149
Level 3	-5.50	0.152
Regression		
θ_1	3.94	0.059
θ_2	87.55	1.363
θ_3	1626.00	31.331
θ_4	1.25	0.054
θ_5	-0.90	0.070
θ_6	-1.56	0.132
θ_7	0.51	0.062
θ_8	0.91	0.049
θ_9	443.87	51.003
Dependence		
δ	0.17	—

Both types of curves can be represented by cumulative probabilities or by the highest probability categories. A recursive mean curve can also be computed by taking the mean of the numbered ordinal categories, but care must be taken with such curves. Indeed, they assume that the ordinal response is following an underlying continuous variable. Hence, the range of this underlying scale corresponding to each of the response categories must be known in order to obtain a recursive mean curve which is not distorted. Unfortunately, this is in most cases unknown and hard to determine.

Recursive fitted response levels are presented in Figure 3 for three individuals in each dose group who remained at an infusion rate of 1ml per minute during the entire trial. The computed recursive means (of the ordinal responses) yield smooth curves typical of a cumulative dependence. The dynamic process can clearly be observed from these curves. The recursive highest probabilities can be very useful because they represent a specific response category. These curves follow each patient's observations reasonably well.

The population predicted cumulative frequencies, for combinations of dose groups (2.5g, 4g, and 8g) and infusion rates (0.5ml/min, 1ml/min, 2ml/min, and 4ml/min) which occurred, are presented in Figure 4. The probabilities of being in the highest category (level 3: excess Doppler signal, above the dotted line) are significantly increased compared to Figures 1c to 1e. Although, such deviations could be thought as implying a poor fit, this is not necessarily the case. Indeed, the

Recursive fitted levels of Doppler signal

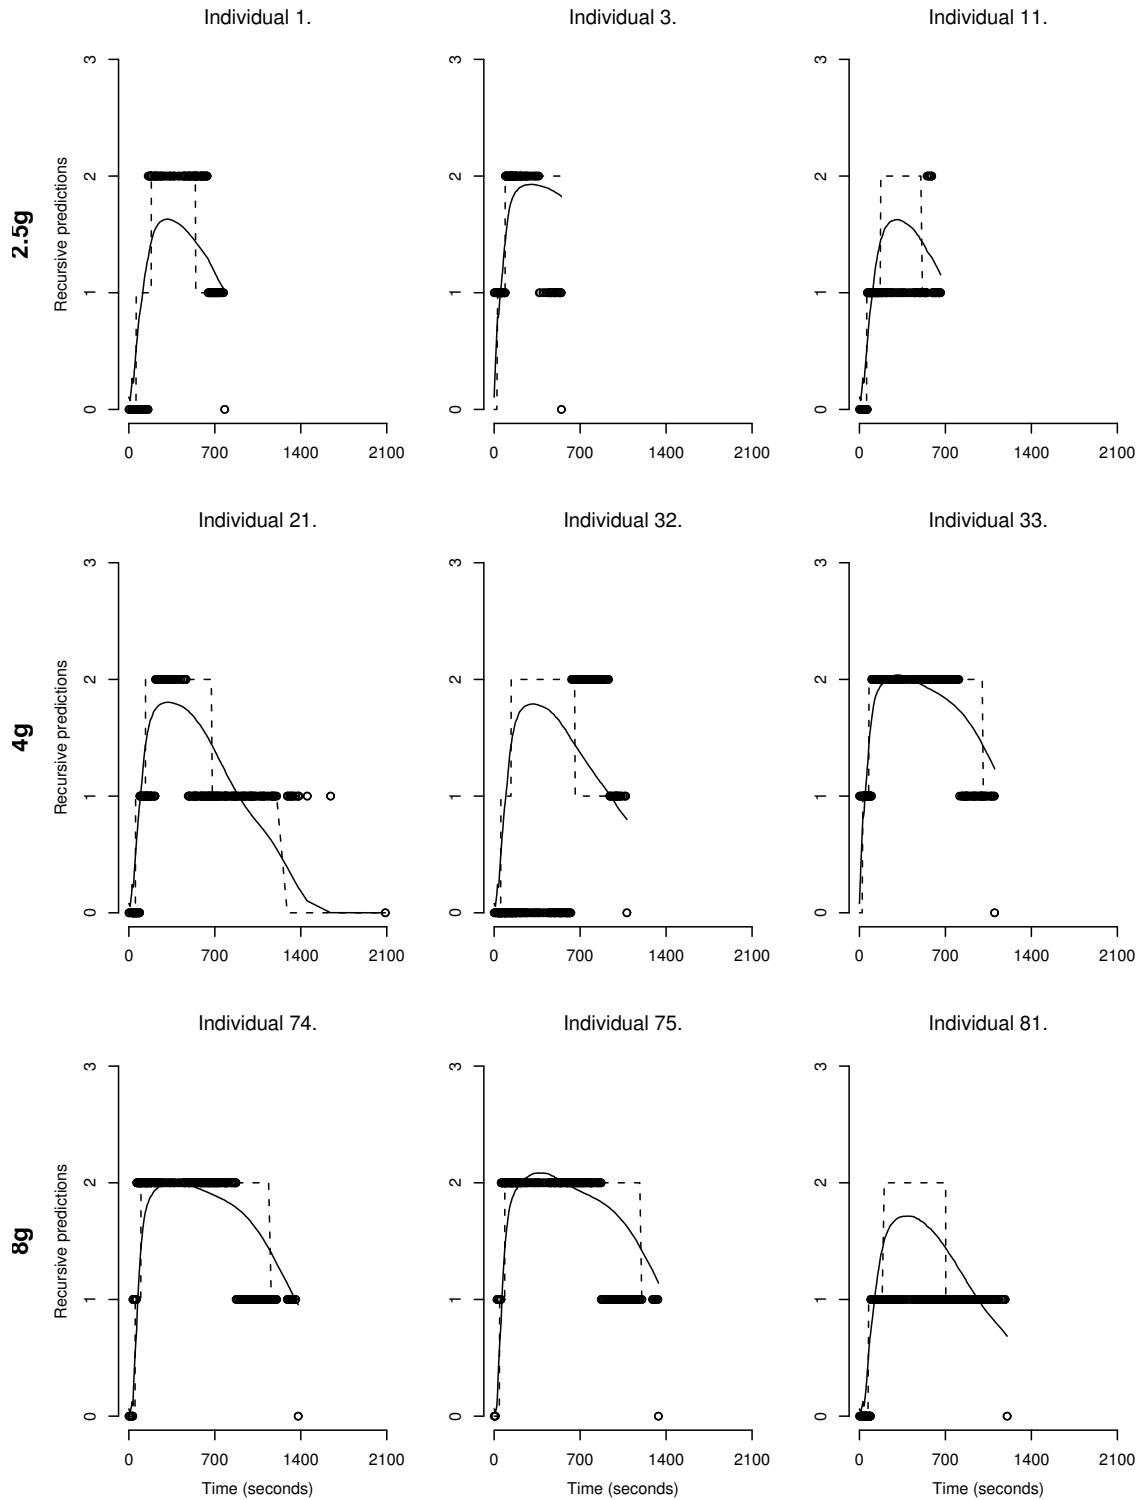


Figure 3: Recursive fitted response levels for the three dose groups (2.5g, 4g, and 8g). Solid line: mean fitted response levels. Dashed line: highest probability fitted response levels. Circles: observations.

Population cumulative frequencies

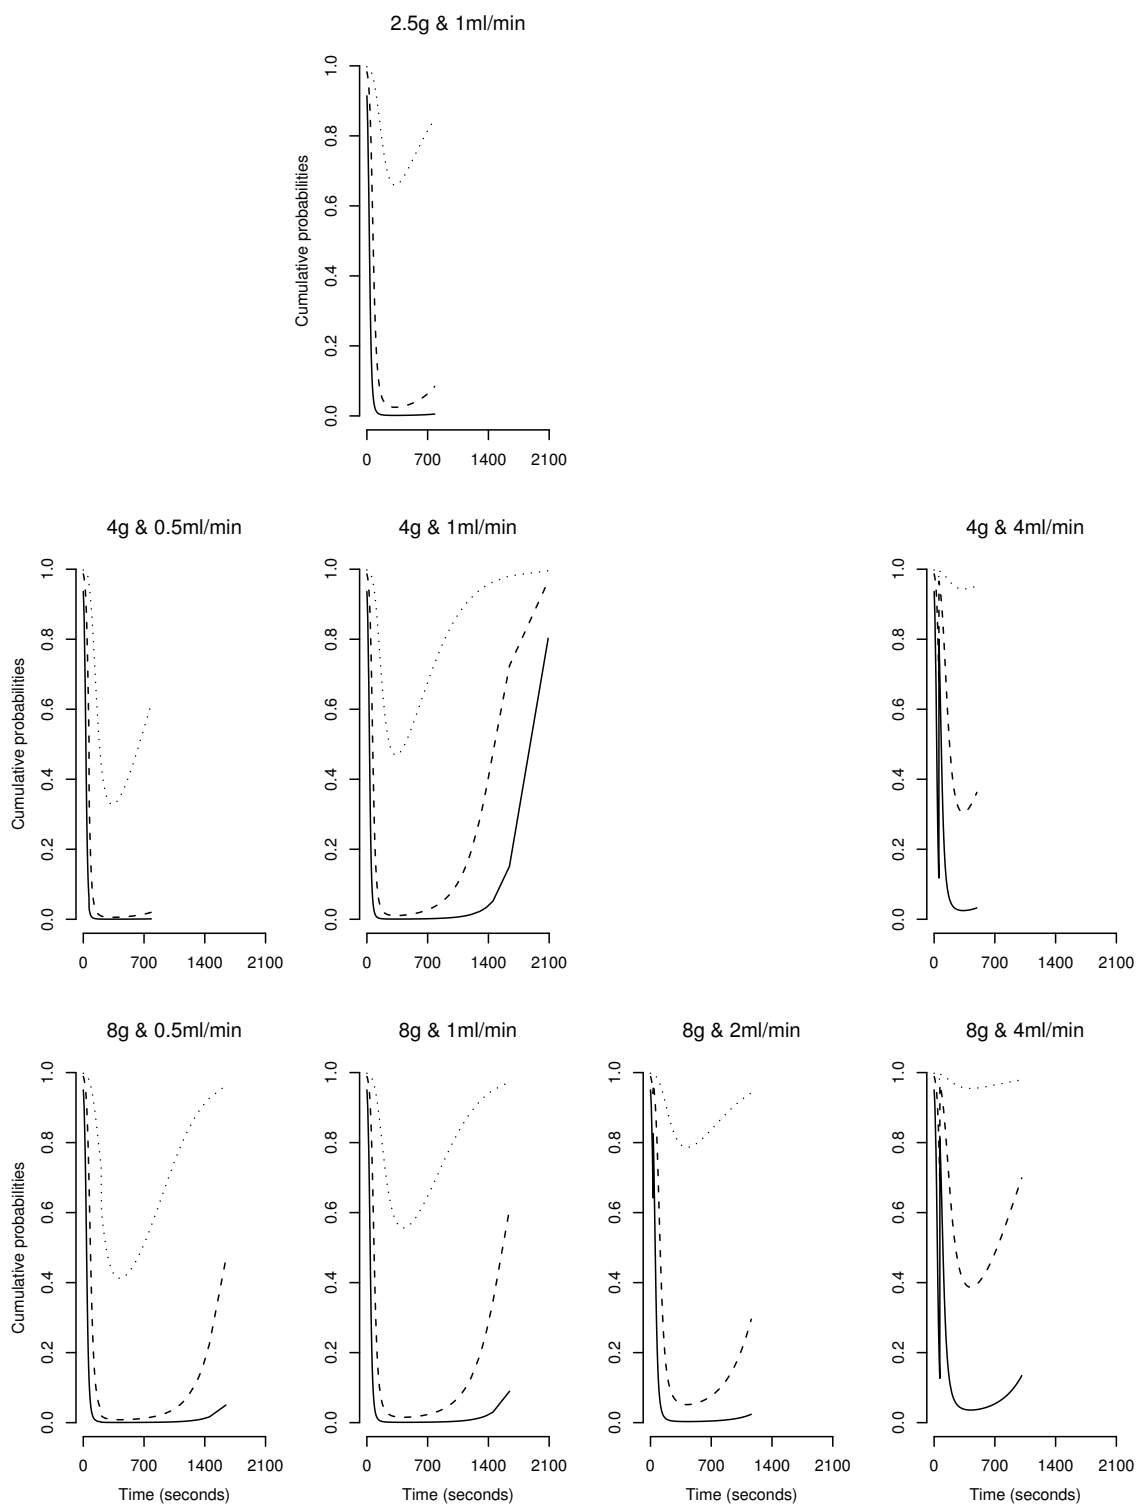


Figure 4: Mean predicted cumulative frequencies for the three drug dose (2.5g, 4g, and 8g), and by the four infusion rates (0.5ml/min, 1ml/min, 2ml/min, and 4ml/min).

Table 3: Estimated time spent at an appropriate Doppler signal level.

Doses	Infusion rates			
	0.5ml/min	1ml/min	2ml/min	4ml/min
2.5g	—	11 min, 46 sec	—	—
4.0g	04 min, 09 sec*	19 min, 57 sec**	—	5 min, 17 sec
8.0g	18 min, 47 sec*	22 min, 49 sec	17 min, 31 sec	8 min, 24 sec

*: in addition, approximately 7 minutes and 30 seconds were spent at an excess signal level.

**: in addition, approximately 3 minutes were spent at an excess signal level.

regression curve constraints are imposed on the individual curves (a conditional model is fitted), not on the population curve (as would be the case with a marginal model). Hence, the individual fitted curves are much more important in assessing the fit of conditional models and these were found, in the present case, to be following each patient’s observations relatively closely.

The estimated time a patient would spend at an appropriate Doppler signal level, can be obtained from these mean predictions. Table 3 summarises the time spent at the required level depending on the dose of drug administered and whether the patient requires a particular adjustment of its infusion rate during the trial.

Finally, it is clear from this analysis that there is a significant difference among the three dose groups. Additionally, the level of Doppler signal observed over time behaves differently for each dose group (ie. the interaction between dose group and time as also significant). From Table 3, it also can be seen that increasing the dose from 2.5g to 4g increases by a factor a little greater than 1.6 the amount of time spent at an appropriate Doppler signal level. Unfortunately, further increasing the dose administered to 8g no longer corresponds to a proportional increase of the estimated time spent at an appropriate Doppler signal level. One should also notice that several patients receiving drug doses of 4g and 8g experienced considerable amounts of time at an excess Doppler signal level and that more changes in the infusion rates are required as the dose administered increases. Hence, some additional studies on high drug doses at lower infusion rates might be of interest.

5 Discussion

The method described in Section 3 is very flexible in many respects. First, different mixture distributions can be obtained with specific properties inherited from the Laplace transform used.

For example, the gamma Laplace transform chosen in this paper does not have the reflection symmetry property. Hence, although changing the ordering direction does not change the proportional odds model, once over-dispersion or any type of dependence is introduced, this symmetry property no longer holds. For the data at hand, the ordering direction was fairly straightforward because the no Doppler signal category clearly corresponds to the lowest possible level of signal.

However, a different Laplace transform might be desirable in the case of some other response variable.

Secondly, many different types of updates can be introduced in order to obtain the desired dependence structure. The serial and Markov updates were not appropriate here to model the relationship among successive responses due to the lack of variability among these. But in the case of responses frequently varying from one response level to another, these types of dependencies prove to be very useful.

Finally, these two features provide a general family of models to handle any type of non-normally distributed repeated measurements. Indeed, such models were already available for continuous responses, count data, and times to event. Now, we have extended them to handle ordinal outcomes. But most of all, this has provided a wide set of models in the area of repeated ordinal measurements where few were until now available (Crouchley, 1995; Ten Have, 1996; Conaway, 1989; Agresti and Lang, 1993).

Acknowledgements The analysis presented in this paper were performed using a special library in R (Ihaka and Gentleman, 1996). R is a fast S-Plus clone freely available under the GNU license, which we thank Robert Gentleman, Ross Ihaka, and the R core group for developing. The R ordinal public library available at www.luc.ac.be/~plindsey/rlibs.html includes the function `kalord` which can fit user-specified non-linear regression models.

We would like to thank Alan Agresti, Robert Crouchley, Paul Janssen, Philippe Lambert, Jim Lindsey, and Geert Molenberghs, who provided many helpful comments. Finally, we would also like to thank the referees for their valuable help in clarifying the different issues discussed.

References

- [1] Agresti, A., and Lang, J.B. (1993) A proportional odds model with subject-specific effects for repeated ordered categorical responses. *Biometrika* **80**, 527–534.
- [2] Akaike, H. (1973) Information theory and an extension of the maximum likelihood principle. in Petrov, B.N. and Csàki, F., eds *Second International Symposium on Inference Theory*, Budapest: Akadémiai Kiadó, pp. 267–281.
- [3] Conaway, M.R. (1989) Analysis of repeated categorical measurements with conditional likelihood methods. *Journal of the American Statistical Association* **84**, 53–62.
- [4] Conaway, M.R. (1990) A random effects model for binary data. *Biometrics* **46**, 317–328.
- [5] Cox, D.R. and Oakes, D. (1984) *Analysis of Survival Data*. London: Chapman and Hall.
- [6] Crouchley, R. (1995) A random-effects model for ordered categorical data. *Journal of the American Statistical Association* **90**, 489–498.

- [7] Glonek, G.F.V. and McCullagh, P. (1995) Multivariate logistic models. *Journal of the Royal Statistical Society B* **57**, 533–546.
- [8] Harvey, A.C. (1989) *Forecasting, structural time series models and the Kalman filter*. Cambridge: Cambridge university press.
- [9] Hougaard, P. (1984) Life time methods for heterogeneous populations: Distributions describing the heterogeneity. *Biometrika* **71**, 75–83.
- [10] Hougaard, P. (1986) Survival models for heterogeneous populations derived from stable distributions. *Biometrika* **73**, 387–396.
- [11] Ihaka, R. and Gentleman, R. (1996) R: a language for data analysis and graphics. *Journal of Computational Graphics and Statistics*, **5**, 299–314.
- [12] Joe, H. (1997) *Multivariate Models and Dependence Concepts*. London: Chapman and Hall.
- [13] Lang, J.B. and Agresti, A. (1994) Simultaneously modeling joint and marginal distributions of multivariate categorical responses. *Journal of the American Statistical Association* **89**, 625–632.
- [14] Lindsey, J.K. (1999) *Models for repeated measurements, 2nd edition*. Oxford: Oxford University Press.
- [15] Lindsey, J.K. (2000) A family of models for uniform and serial dependence in repeated measurements studies. *Journal of the Royal Statistical Society C* **49**, 343–357.
- [16] Lindsey, J.K. and Jones, B. (1998) Choosing among generalized linear models applied to medical data. *Statistics in Medicine* **16**, 59–68.
- [17] McCullagh, P., and Nelder, J.A. (1983) *Generalized Linear Models*. London: Chapman and Hall.
- [18] Nelder, J.A. and Wedderburn, R.W.M. (1972) Generalized linear models. *Journal of the Royal Statistical Society A* **135**, 370–383.
- [19] Nelsen, R.B. (1999) *An Introduction to Copulas*. New-York: Springer.
- [20] Ten Have, T.R. (1996) A mixed effects model for multivariate ordinal response data including correlated discrete failure times with ordinal responses. *Biometrics* **52**, 473–491.
- [21] Yue, H. and Chan, K.S. (1997) A dynamic frailty model for multivariate survival data. *Biometrics* **53**, 785–793.