A latent-class mixture model for incomplete longitudinal Gaussian data

Peer-reviewed author version

BEUNCKENS, Caroline; MOLENBERGHS, Geert; VERBEKE, Geert & Mallinckrodt, Craig (2008) A latent-class mixture model for incomplete longitudinal Gaussian data. In: BIOMETRICS, 64(1). p. 96-105.

DOI: 10.1111/j.1541-0420.2007.00837.x

Handle: <http://hdl.handle.net/1942/9518>

A Latent-Class Mixture Model For Incomplete Longitudinal Gaussian Data

Caroline Beunckens,^{1,*} Geert Molenberghs,¹

Geert Verbeke,² and Craig Mallinckrodt³

¹ Center for Statistics, Hasselt University, Agoralaan 1, 3590 Diepenbeek, Belgium.

² Biostatistical Centre, Catholic University of Leuven,
Kapucijnenvoer 35, 3000 Leuven, Belgium.

³ Eli Lilly & Company, Lilly Corporate Center, Indianapolis, IN 46285, U.S.A.

* *email:* caroline.beunckens@uhasselt.be

SUMMARY. In the analyses of incomplete longitudinal clinical trial data, there has been a shift, away from simple methods that are valid only if the data are missing completely at random (MCAR), to more principled ignorable analyses, which are valid under the less restrictive missing at random (MAR) assumption. The availability of the necessary standard statistical software nowadays allows for such analyses in practice. While the possibility of data missing not at random (MNAR) cannot be ruled out, it is argued that analyses valid under MNAR are not well suited for the primary analysis in clinical trials. Rather than either forgetting about or blindly shifting to an MNAR framework, the optimal place for MNAR analyses is within a sensitivity analysis context. One such route for sensitivity analysis is to consider, next to selection models, pattern-mixture models or shared-parameter models. The latter can also be extended to a latent-class mixture model, the route taken in this paper. The so-obtained flexible model is submitted to the test in simulations and applied to data from a depression trial.

KEY WORDS: Latent class, Nonrandom missingness, Random effect, Shared parameter

1 Introduction

Repeated measures are often prone to incompleteness, often taking the dropout form. The nature of the dropout mechanism can affect inference. Since one can never be certain about the dropout mechanism, certain assumptions have to be made. We will use the terminology introduced by Rubin (1976). A non-response process is *missing completely at random* (MCAR) if the missingness is independent of both unobserved and observed data, and *missing at random* (MAR) if, conditional on the observed data, the missingness is independent of the unobserved measurements. A process that is neither MCAR nor MAR is *non-random* (MNAR). For likelihood inference, and when the parameters describing the measurement process are functionally independent of the parameters describing the missingness process, MCAR and MAR are *ignorable*, in which case the missingness process can be ignored when interest is in inference for the longitudinal process only. For frequentist inference, the stronger condition of MCAR is required for ignorability.

Recently, there has been a move away from simple methods, such as complete case analysis and last observation carried forward analysis, towards such likelihood-based methods as mixed models (Molenberghs et al., 2004; Jansen et al., 2006). Of course, since non-random methods allow the missingness to depend on the unobserved or missing values, it is clear that the MNAR assumption is not verifiable (Laird, 1994; Molenberghs, Kenward and Lesaffre, 1997). Therefore, fitting a single MNAR model will not be trustworthy, which does not mean we should ignore such models but rather frame them within a sensitivity analysis. So far, we have used selection-model concepts. Alternatively, pattern-mixture and shared-parameter models can be used. They are introduced in the next section.

We propose a *latent-class mixture model*, bringing together features of the selection, pattern-mixture, and shared-parameter model frameworks. Information from the location and evolution of the response profiles, a selection model concept, and from the dropout patterns, a pattern-mixture idea, is used simultaneously to define latent groups and variables, a shared-

parameter feature. This brings several appealing features. First, one uses information in a more symmetric, elegant way. Second, apart from providing a more flexible modeling tool, there is room for use as a sensitivity analysis instrument. Third, a strong advantage over existing methods is that we are able to classify subjects into latent groups. If done with due caution, it can enhance substantive knowledge and generate hypotheses. Fourth, while computational burden increases, fitting the proposed method is remarkably stable and acceptable in terms of computation time for the settings considered here. Clearly, neither the proposed model nor any other alternative can be seen as a tool to definitively test for MAR *versus* MNAR, as amply documented in the sensitivity analysis literature. This is why the method's use predominantly lies within the sensitivity analysis context. Such a sensitivity analysis is of use both when it modifies the results of a simpler analysis, for further scrutiny, as well as when it confirms these.

The general latent-class mixture model is presented in Section 2. A simulation study is reported in Section 5. In Section 6, data from a depression clinical trial are analyzed using a latent-class mixture model within a sensitivity analysis.

2 Latent-Class Mixture Models

Let Y_{ij} denote the response for the i th individual, designed to be measured at time t_{ij} , $i = 1, \dots, N$, $j = 1, \dots, n_i$. Group the outcomes into a vector $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{in_i})'$. Define a dropout indicator D_i for the occasion at which dropout occurs, with the convention that $D_i = n_i + 1$ for a complete sequence. Split the vector $\mathbf{Y}_i$ into observed ($\mathbf{Y}_i^o$) and missing ($\mathbf{Y}_i^m$) components, respectively.

In principle, one would like to consider the density of the full data $f(\mathbf{y}_i, d_i | \boldsymbol{\theta}, \boldsymbol{\psi})$, where the parameter vectors $\boldsymbol{\theta}$ and $\boldsymbol{\psi}$ describe the measurement and missingness processes, respectively. Covariates are allowed but suppressed from notation. This full density function can be factorized in different ways. The *selection model* framework is based on (Rubin, 1976; Little

and Rubin, 1987): $f(\mathbf{y}_i, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}) = f(\mathbf{y}_i | \boldsymbol{\theta})f(d_i | \mathbf{y}_i, \boldsymbol{\psi})$. The first factor is the marginal density of the measurement process and the second one is the density of the missingness process, conditional on the outcomes. Alternatively, one can consider *pattern-mixture models* (Little, 1993, 1994) using the reversed factorization $f(\mathbf{y}_i, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}) = f(\mathbf{y}_i | d_i, \boldsymbol{\theta})f(d_i | \boldsymbol{\psi})$. This can be seen as a mixture of different populations, characterized by the observed pattern of missingness. Instead of using the selection modelling or pattern-mixture modelling framework, the measurement and the dropout process can be jointly modeled by using a *shared-parameter model* (Wu and Carrol, 1988; Wu and Bailey, 1989; Ten Have et al., 1998). These methods assume there exists a vector of random effects $\mathbf{b}_i$, conditional upon which the measurement and dropout processes are independent: $f(\mathbf{y}_i, d_i | \mathbf{b}_i, \boldsymbol{\theta}, \boldsymbol{\psi}) = f(\mathbf{y}_i | \mathbf{b}_i, \boldsymbol{\theta})f(d_i | \mathbf{b}_i, \boldsymbol{\psi})$. We propose an extension, capturing possible heterogeneity between the subjects not measured through covariates but rather through a latent variable. We call this model a *latent-class mixture model*. Next to one or more so-called shared parameters, $\mathbf{b}_i$, the model contains a latent variable, $\mathbf{Q}_i$, dividing the population into g subgroups. This latent variable is a vector of group indicators $\mathbf{Q}_i = (Q_{i1}, \dots, Q_{ig})$, defined as $Q_{ik} = 1$, if subject i belongs to group k , and 0 otherwise. The measurement process as well as the dropout process depend on this latent variable, directly and through the subject-specific effects $\mathbf{b}_i$. The distribution of $\mathbf{Q}_i$ is multinomial and defined by $P(Q_{ik} = 1) = \pi_k$ ($k = 1, \dots, g$), where π_k denotes the group or component probability, also termed prior probabilities of the components. These are restricted through $\sum_{k=1}^g \pi_k = 1$.

The measurement process will be modeled by a heterogeneity linear mixed model (Verbeke and Lesaffre, 1996; Verbeke and Molenberghs, 2000): $\mathbf{Y}_i | Q_{ik} = 1, \mathbf{b}_i \sim N(\mathbf{X}_i \boldsymbol{\beta}_k + \mathbf{Z}_i \mathbf{b}_i, \boldsymbol{\Sigma}_i^{(k)})$, where $\mathbf{X}_i$ and $\mathbf{Z}_i$ are design matrices, $\boldsymbol{\beta}_k$ component-dependent fixed effects, $\mathbf{b}_i$ denote the shared parameters, following a mixture of g normal distributions with mean vectors $\boldsymbol{\mu}_k$ and covariance matrices $\mathbf{D}_k$, i.e., $\mathbf{b}_i | Q_{ik} = 1 \sim N(\boldsymbol{\mu}_k, \mathbf{D}_k)$ and thus $\mathbf{b}_i \sim \sum_{k=1}^g \pi_k N(\boldsymbol{\mu}_k, \mathbf{D}_k)$. The measurement error terms $\boldsymbol{\varepsilon}_i$ follow a normal distribution with mean zero and covariance

matrix $\Sigma_i^{(k)}$ and are independent of the shared parameters. The mean and the variance of $\mathbf{Y}_i$ are:

$$E(\mathbf{Y}_i) = \mathbf{X}_i \sum_{k=1}^g \pi_k \boldsymbol{\beta}_k + \mathbf{Z}_i \sum_{k=1}^g \pi_k \boldsymbol{\mu}_k, \quad (1)$$

$$\text{Var}(\mathbf{Y}_i) = \mathbf{Z}_i' \left[\sum_{k=1}^g \pi_k \boldsymbol{\mu}_k^2 - \left(\sum_{k=1}^g \pi_k \boldsymbol{\mu}_k \right)^2 + \sum_{k=1}^g \pi_k \mathbf{D}_k \right] \mathbf{Z}_i + \sum_{k=1}^g \pi_k \Sigma_i^{(k)}. \quad (2)$$

Assuming that the shared effects are ‘calibrated’, i.e., $\sum_{k=1}^g \pi_k \boldsymbol{\mu}_k = \mathbf{0}$, (1) and (2) simplify to:

$$E(\mathbf{Y}_i) = \mathbf{X}_i \sum_{k=1}^g \pi_k \boldsymbol{\beta}_k,$$

$$\text{Var}(\mathbf{Y}_i) = \mathbf{Z}_i' \left[\sum_{k=1}^g \pi_k \boldsymbol{\mu}_k^2 + \sum_{k=1}^g \pi_k \mathbf{D}_k \right] \mathbf{Z}_i + \sum_{k=1}^g \pi_k \Sigma_i^{(k)}.$$

Assuming that the first measurement Y_{i1} is obtained for every subject in the study, the model for the dropout process is based on a logistic regression for the probability of dropout at occasion j , given: (1) the subject was still in the study up to occasion j , (2) $\mathbf{b}_i$, and (3) that the subject belongs to the k th component. Denote this probability by $g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})$, in which $\mathbf{w}_{ij}$ is a vector containing all relevant covariates: $g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik}) = P(D_i = j | D_i \geq j, \mathbf{w}_{ij}, \mathbf{b}_i, q_{ik} = 1)$. We then assume that $g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})$ satisfies $\text{logit}[g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] = \mathbf{w}_{ij} \boldsymbol{\gamma}_k + \boldsymbol{\lambda} \mathbf{b}_i$. The joint likelihood of the measurement and dropout processes will take the form:

$$\begin{aligned} f(\mathbf{y}_i, d_i) &= \sum_{k=1}^g P(q_{ik} = 1) f(\mathbf{y}_i, d_i | q_{ik} = 1) \\ &= \sum_{k=1}^g \pi_k \int f(\mathbf{y}_i, d_i | q_{ik} = 1, \mathbf{b}_i) f_k(\mathbf{b}_i) d\mathbf{b}_i \\ &= \sum_{k=1}^g \pi_k \int f(\mathbf{y}_i | q_{ik} = 1, \mathbf{b}_i, \mathbf{X}_i, \mathbf{Z}_i) f(d_i | q_{ik} = 1, \mathbf{b}_i, \mathbf{w}_i) f_k(\mathbf{b}_i) d\mathbf{b}_i, \end{aligned} \quad (3)$$

with $f(\mathbf{y}_i | q_{ik} = 1, \mathbf{b}_i, \mathbf{X}_i, \mathbf{Z}_i)$ the density function of the normal distribution $N(\mathbf{X}_i \boldsymbol{\beta}_k +$

$\mathbf{Z}_i \mathbf{b}_i, \Sigma_i^{(k)}$, $f_k(\mathbf{b}_i)$ the density function of $N(\boldsymbol{\mu}_k, \mathbf{D}_k)$, and

$$f(d_i | q_{ik} = 1, \mathbf{b}_i, \mathbf{w}_i) = \begin{cases} g_{id_i}(\mathbf{w}_{id_i}, \mathbf{b}_i, q_{ik}) \times \prod_{j=2}^{d_i-1} [1 - g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] & \text{if incomplete,} \\ \prod_{j=2}^{n_i} [1 - g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] & \text{if complete.} \end{cases}$$

Note that the dropout model can depend, not only on the outcomes, but also on relevant covariates such as treatment allocation, time, gender, age, etc. Not all models that can be formulated in this way are identified, so restrictions might be needed. We return to this point in Section 3.

The latent-class mixture models are based on assuming a latent structure additional to the measurement and dropout processes and grouping the subjects by means of a latent variable, thereby accounting for inter-group differences both in terms of their dropout pattern as well as their measurement profiles.

3 Likelihood Function and Estimation

Estimation of the unknown parameters in the latent-class mixture model described in the previous section will be based on maximum likelihood. Since it would be very cumbersome to maximize this likelihood function analytically, the EM algorithm (Dempster, Laird and Rubin, 1977) is proposed as it is a practical tool for maximum likelihood estimation in the case of finite mixtures (Redner and Walker, 1984).

Let $\boldsymbol{\pi}$ be the vector of component probabilities $\boldsymbol{\pi}' = (\pi_1, \dots, \pi_g)$ and group all other unknown parameters of the measurement process in the vector $\boldsymbol{\theta}$, of the dropout process in $\boldsymbol{\psi}$, and of the mixture distribution in $\boldsymbol{\alpha}$. If $\boldsymbol{\sigma}$ denotes the vector of covariance parameters of all $\Sigma_i^{(k)}$, $\boldsymbol{\delta}$ the covariance parameters of all $\mathbf{D}_k$, $\boldsymbol{\mu}' = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_g)$, and $\boldsymbol{\gamma}' = (\gamma_1, \dots, \gamma_g)$, then $\boldsymbol{\theta} = (\boldsymbol{\beta}, \boldsymbol{\sigma})$, $\boldsymbol{\psi} = (\boldsymbol{\gamma}, \boldsymbol{\lambda})$ and $\boldsymbol{\alpha} = (\boldsymbol{\mu}, \boldsymbol{\delta})$. Now, the vector $\boldsymbol{\Omega}$ will be the vector containing all unknown parameters in the model, i.e., $\boldsymbol{\Omega}' = (\boldsymbol{\pi}', \boldsymbol{\theta}', \boldsymbol{\psi}', \boldsymbol{\alpha}')$.

The observed-data likelihood, $L(\boldsymbol{\Omega}|\mathbf{y}^o, \mathbf{d})$, is obtained as follows:

$$\begin{aligned}
L(\boldsymbol{\Omega}|\mathbf{y}^o, \mathbf{d}) &= \prod_{i=1}^N f(\mathbf{y}_i^o, d_i|\boldsymbol{\Omega}) = \prod_{i=1}^N \int f(\mathbf{y}_i, d_i|\boldsymbol{\Omega}) d\mathbf{y}_i^m \\
&= \prod_{i=1}^N \int \left\{ \sum_{k=1}^g \pi_k \int f(\mathbf{y}_i|\boldsymbol{\theta}, \mathbf{b}_i, q_{ik} = 1) f(d_i|\boldsymbol{\psi}, \mathbf{b}_i, q_{ik} = 1) f_k(\mathbf{b}_i|\boldsymbol{\alpha}) d\mathbf{b}_i \right\} d\mathbf{y}_i^m \\
&= \prod_{i=1}^N \sum_{k=1}^g \pi_k \int f(\mathbf{y}_i^o|\boldsymbol{\theta}, \mathbf{b}_i, q_{ik} = 1) f(d_i|\boldsymbol{\psi}, \mathbf{b}_i, q_{ik} = 1) f_k(\mathbf{b}_i|\boldsymbol{\alpha}) d\mathbf{b}_i, \tag{4}
\end{aligned}$$

where $\mathbf{y}^{o'} = (\mathbf{y}_1^o, \dots, \mathbf{y}_N^o)$ is the vector containing all observed response values and $\mathbf{d} = (d_1, \dots, d_N)$ is the vector of all values of the dropout indicator. This likelihood function is invariant under the $g!$ possible permutations of the parameters corresponding to each of the g mixture components. To overcome this, the constraint suggested by Aitkin and Rubin (1985), $\pi_1 \geq \pi_2 \geq \dots \geq \pi_g$, are imposed.

Identifiability is a delicate issue. Böhning (1999) shows that a mixture of two normals with simultaneously different means and different variances is not identifiable. Such problems arise from latency, now occurring through latent classes, random effects, and missingness. In line with Böhning (1999) and McLachlan and Peel (2000), one could consider several variations to the target model. The likelihood values, parameter estimates, and information matrices can then be studied in view of identifiability. All of this cautions against certain uncritical uses of the model. For example, one should not place a blind belief in the number of components resulting when applying the model; this would be a license for mischief. Rather, the model can play a role as a sensitivity analysis tool. The allocation of subjects to latent groups can help formulate hypotheses, which then have to be checked against substantive scientific knowledge and/or in follow up studies. On the other hand, when focusing on certain inferences, such as testing for treatment effect, the number of components, for example, may be less essential.

To maximize the log-likelihood function, corresponding to (4), with respect to $\boldsymbol{\Omega}$, the EM algorithm is used (Dempster, Laird and Rubin, 1977). The underlying latent variable $\mathbf{Q}_i$, representing component membership, will be considered missing. Thus, the response vector

$\mathbf{Y}_i^o$, the dropout indicator D_i , and the (unobserved) population indicators $\mathbf{Q}_i$, are the so-called augmented data even though only vectors $\mathbf{Y}_i^o$ and D_i are observed. Since the joint density of $\mathbf{Y}_i^o$, D_i and $\mathbf{Q}_i$ equals

$$\begin{aligned} & f_i(\mathbf{y}_i^o, d_i, Q_{i1} = q_{i1}, \dots, Q_{ig} = q_{ig}) \\ &= f_i(\mathbf{y}_i^o, d_i | Q_{i1} = q_{i1}, \dots, Q_{ig} = q_{ig}) \times P(Q_{i1} = q_{i1}, \dots, Q_{ig} = q_{ig}) \\ &= \prod_{k=1}^g [\pi_k f_{ik}(\mathbf{y}_i^o, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\alpha})]^{q_{ik}}, \end{aligned}$$

the joint likelihood $L(\boldsymbol{\Omega} | \mathbf{y}^o, \mathbf{d}, \mathbf{q})$ of the augmented data, is

$$L(\boldsymbol{\Omega} | \mathbf{y}^o, \mathbf{d}, \mathbf{q}) = \prod_{i=1}^N \prod_{k=1}^g [\pi_k f_{ik}(\mathbf{y}_i^o, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\alpha})]^{q_{ik}}, \quad (5)$$

with $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_n)'$ the vector of all hypothetically observed population indicators.

Maximizing $\ell(\boldsymbol{\Omega} | \mathbf{y}^o, \mathbf{d}, \mathbf{q})$, the corresponding log-likelihood, is easier than maximizing the log-likelihood $\ell(\boldsymbol{\Omega} | \mathbf{y}^o, \mathbf{d})$. Precisely, the E-step of the EM algorithm computes $E[\ell(\boldsymbol{\Omega} | \mathbf{y}^o, \mathbf{d}, \mathbf{Q}) | \mathbf{y}, \mathbf{d}]$, where the expectation is over $\mathbf{y}^o$ and $\mathbf{d}$. In the maximization (M) step, this function is maximized. Denote the expected log-likelihood function, the so-called objective function, by $\mathcal{O}$. The EM algorithm is initiated by means of an initial value $\boldsymbol{\Omega}^{(0)}$, after which one oscillates between the E- and M-steps, until convergence. More details on the EM algorithm can be found in the electronically available Appendix (see Supplementary Materials). With an eye on speed of convergence, one might start with the EM algorithm and then switch to direct likelihood maximization. Of course, the marginal likelihood may take a somewhat cumbersome form.

4 Classification

One can also classify the subjects into the different latent subgroups of the fitted model. Through the structure of the latent-class mixture model, the subdivision of the population in latent groups depends on the number of observed measurements, i.e., on the dropout

indicator or pattern, as well as on the values of the observed response measurements and hence can be useful to assess the coherence between the dropout process and the measurement process. In certain cases such latent groups can have substantive meaning.

The decision to which component of the mixture a specific subject is most likely to belong will be based on *posterior probabilities*. We have that $P(Q_{ik} = 1) = \pi_k$, thus the component probabilities π_k , express how likely the i th subject is to belong to group k , without using information from outcomes and dropout pattern. For this reason, the component probabilities are often called *prior* probabilities. The *posterior* probability for subject i to belong to the k th group is given by

$$\pi_{ik} = P(Q_{ik} = 1 | \mathbf{y}_i^o, d_i) = \frac{f_i(\mathbf{y}_i^o, d_i | Q_{ik} = 1) P(Q_{ik} = 1)}{f_i(\mathbf{y}_i^o, d_i)} \bigg|_{\hat{\Omega}} = \frac{\pi_k f_{ik}(\mathbf{y}_i^o, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\alpha})}{\sum_{k=1}^g \pi_k f_{ik}(\mathbf{y}_i^o, d_i | \boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\alpha})} \bigg|_{\hat{\Omega}}.$$

This expresses how likely the i th subject is to belong to group k , taking into account the observed response $\mathbf{y}_i$ as well as the dropout indicator d_i of that subject. Using these posterior probabilities, we classify subject i into component k if and only if $\pi_{ik} = \max_j \{\pi_{ij}\}$.

Clearly, we should be cautious with the resulting classification, since for a particular subject i , the vector of posterior probabilities is given by $\boldsymbol{\pi}_i = (\pi_{i1}, \dots, \pi_{ig})$ with $\sum_{k=1}^g \pi_{ik} = 1$. For a good comfort level, one of these posterior probabilities for subject i would lie close to 1. However, a scenario would be that two or more posterior probabilities are almost equal, of which one is the maximum of all posterior probabilities for that particular subject. This makes classification nearly random and misclassification is likely to occur. Therefore, rather than merely considering the classification of subjects into the latent subgroups, it is instructive to inspect the posterior probabilities in full. Furthermore, we can vary the number of latent groups g and explore the sensitivity of the classification to the number of latent subgroups considered.

5 Simulation Study

An advantage of the latent-class mixture model is its flexible structure, making the model a helpful tool for analyzing incomplete longitudinal data. However, as already seen in Section 3, the estimation of the model parameters is based on a doubly iterative method, which we might expect to be computationally intensive. To assess whether this disadvantage counterbalances the advantage of model flexibility, and to assess performance, we conduct a simulation study. Section 5.1 describes a simplification of the latent-class mixture model used in the simulations and in the application in Section 6. The design and results of the simulation study are given in Sections 5.2 and 5.3, respectively.

5.1 A Simplification of the Latent-Class Mixture Model

We assume equal covariance matrices for the different mixture components, i.e., $\mathbf{D}_1 = \dots = \mathbf{D}_g = \mathbf{D}$, and equal residual covariance matrices, i.e., $\Sigma_i^{(1)} = \dots = \Sigma_i^{(g)} = \Sigma_i$, which leads to $\mathbf{Y}_i | q_{ik} = 1, \mathbf{b}_i \sim N(X_i \boldsymbol{\beta} + Z_i \mathbf{b}_i, \Sigma_i)$, with $\mathbf{b}_i \sim \sum_{k=1}^g \pi_k N(\boldsymbol{\mu}_k, \mathbf{D})$. The matrices $\Sigma_i^{(k)}$ usually depend on i only through the dimension of the response vector for subject i , while parameters are common to all. We further simplify the model in two steps. First, it is assumed that there is only one subject-specific effect b_i , a shared intercept, influencing the measurement process, not the dropout process. Second, the measurement process is assumed to depend on the latent variable, not in a direct way, but only through the shared intercept.

5.2 Design of the Simulation Study

We simulated 250 datasets, each containing measurements and covariate information of 100 subjects. The latent variable in the model is assumed to split the subjects into two latent subgroups with component probabilities $\pi_1 = 0.6$ and $\pi_2 = 1 - \pi_1 = 0.4$. There are five measurement occasions and the outcome follows a linear trend over time with intercept $\beta_0 = 9.4$ and slope $\beta_1 = 2.25$. The shared intercept follows a mixture of two normal distributions with means $\mu_1 = -4.4$ and $\mu_2 = -\frac{\pi_1 \mu_1}{\pi_2} = 6.6$. In line with Section 5.1, the

variances of these two normal distributions are assumed to be equal and are denoted by d^2 . The measurement error variance is σ^2 .

Four settings will be considered, based on varying d^2 and σ^2 . In the first setting both variance parameters are relatively small, $d = 2.0$ and $\sigma = 0.25$. While only the measurement error variance is increased in the second setting, $\sigma = 0.75$, both variance parameters are increased in the third setting, $d = 3.5$ and $\sigma = 1.00$. Up to the third setting, the chosen parameters result in a bimodal, well-separated mixture distribution. Since this might improve estimation of the parameters, we consider a fourth, unimodal setting with $d = 6$ and $\sigma = 2$.

Finally, in the dropout model, the logistic regression is based on an intercept only, which differs for both latent classes, namely, $\gamma_1 = -2.5$ and $\gamma_2 = -1.25$, respectively, with corresponding probabilities 0.73 and 0.45 of completing the study.

The latent-class mixture model can now be formulated as follows. For a subject $i = 1, \dots, 100$, belonging to latent group $k = 1, 2$, the measurement at time $j = 1, \dots, 5$ is modeled by $Y_{ij} = \beta_0 + \beta_1 \text{time}_j + b_i + \varepsilon_{ij}^{(k)}$, with $b_i \sim \pi_1 N(\mu_1, d^2) + \pi_2 N(\mu_2, d^2)$ and $\varepsilon_i^{(k)} \sim N(0, \sigma^2 \mathbf{I}_5)$. Further, the dropout model is expressed as $\text{logit}[g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] = \gamma_k$.

5.3 Results of the Simulation Study

Table 1 contains the results of the simulation study. Besides comparing the mean estimates and true values of the parameters through the bias, we also consider the mean squared error (MSE), simultaneously involving bias and precision.

Let us discuss the four simulation settings in turn. For the first one, there is a clear distinction between both groups, which owes to the small variance, d^2 , of the mixture distribution, relative to the systematic difference between the mean of both groups, $\mu_1 - \mu_2$. Further, the small measurement error variance, σ^2 , ensures the within-subject variability to be small, resulting in almost straight individual profiles. From Table 1, the mean estimates of the parameters are close to the true values, with biases of the order 10^{-2} or less. Together

with small mean squared error values, of which the magnitude does not exceed 10^{-4} , this indicates the fit of the latent-class mixture model is very close to the simulated data. This was expected due to earlier observations.

Increasing the measurement error variance in the second simulation setting leads to an increased within-subject variability. The discrepancy between both latent groups is still present in the second setting. The bias increases slightly, but remains of the same order. For the MSE values, we observe a small increase, but its magnitude does not exceed 10^{-3} . So, we can conclude the model fits the data well, even with a larger within-subject variability.

In the penultimate simulation setting, not only the measurement error variance is increased, but also the variance in the mixture components. In the third setting, there is no gap in between both latent groups. The discrepancy between the groups seems to have vanished, and profiles appear to be homogeneous. Let us look at the results in Table 1 to see whether this has an influence on the model fit. For some of the parameters, the mean estimates deviates little from the true value. However, bias and MSE values remain small, the order of magnitude not exceeding 10^{-1} and 10^{-3} , respectively.

Finally, in the last simulation setting, in which even larger values for both variance parameters result in simulated data following an unimodal mixture distribution, profiles again seem to be homogeneous. Remarkably, even in this setting, bias and MSE values remain small, both with order of magnitude below 10^{-1} .

In all four simulation settings, we ascertain small bias and MSE values. This is true, not only for the model parameters, but also for derived quantities, such as the treatment effect at the last time and the area under the curve. Thus the latent-class mixture model does fit the simulated data well. Moreover, from the four simulation settings we conclude that, whenever the model is correctly specified, it fits very well; so, this applies even when the mixture distribution is unimodal. This suggests that, for a real application, the fit is likely to be good in cases where the researcher has decent insight into the true mean structure.

Computation time increased from about 30 minutes for fitting the latent-class mixture model to a simulated dataset of the first setting, to a bit over two hours for fitting one of the later settings. Thus, fitting the latent-class mixture model is not unreasonable in terms of computation time, perhaps against initial expectation.

6 Analysis of Depression Trial Data

We apply the latent-class mixture model to a depression trial, arising from a randomized, double-blind psychiatric clinical trial, conducted in the United States. The primary objective of this trial was to compare the efficacy of an experimental anti-depressant with a non-experimental one. In these retrospective analyses, data from 170 patients are considered. The Hamilton Depression Rating Scale ($HAMD_{17}$) is used to measure the depression status of the patients. For each patient, a baseline assessment is available, as well as 5 post-baseline visits going from visit 4 to 8.

In the two subsequent sections, a latent-class mixture model is fitted to the depression trial and a sensitivity analysis performed. The latter will establish the latent-class mixture model as a viable sensitivity tool.

6.1 Formulating a Latent-Class Mixture Model

The latent-class mixture model framework is used to analyze the depression trial, assuming the patients can be split into g latent subgroups.

The mean structure is determined based on an exploratory analysis. As a result, the heterogeneity linear mixed model for the change in $HAMD_{17}$ score includes as fixed effects an intercept, the treatment variable, the baseline $HAMD_{17}$ score, the linear and quadratic time variable, and the interaction between treatment and time. In the latent-class mixture model with two group, the parameter values for these fixed effects are assumed to be equal for both latent subgroups. The measurement error terms are assumed to be independent and to follow a normal distribution with mean 0 and variance σ^2 .

A shared intercept, b_i , is included in the measurement model, which follows a mixture of g normal distributions with different means, $\mu_1, \dots, \mu_g$ respectively, but with equal variance d^2 . Further, the dropout process is modeled based on a logistic regression, including an intercept and time variable, which, in case of the two-group latent-class model, can differ between both latent subgroups ($\gamma_{0,1}, \dots, \gamma_{0,g}$ corresponding to the intercept, and $\gamma_{1,1}, \dots, \gamma_{1,g}$ corresponding to the slope). While dependence on covariates was allowed, none were retained during model building.

At first, the same simplification as used in the simulation study in Section 5 is considered, i.e., the shared intercept is only included in the measurement model, not in the dropout model. Afterwards, we extend the model by adding the shared intercept to the dropout model as well, meaning the dropout model changes from

$$\text{logit}[g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] = \gamma_{0,k} + \gamma_{1,k} t_j \quad (6)$$

to

$$\text{logit}[g_{ij}(\mathbf{w}_{ij}, \mathbf{b}_i, q_{ik})] = \gamma_{0,k} + \gamma_{1,k} t_j + \lambda b_i, \quad (7)$$

where t_j is the j th visit.

An overview of the models considered is given in Table 2. Since assessing the number of components by a classical likelihood ratio test is not valid in the mixture model framework (McLachlan and Peel, 2000), we calculated the Akaike's Information Criterion (AIC) and Bayesian Information Criterion (BIC) for all models.

A model building exercise is performed starting with fitting a one-component latent-class mixture model, which comes down to a classical shared parameter model, as well as a two-component latent-class mixture model. Next, we compare these models using the AIC and BIC criteria, and depending on the choice made by both criteria, we decide whether we fit a latent-class mixture model with three latent subgroups.

Table 2 shows that when assuming dropout model (6), AIC opts for the model with two

latent subgroups (Model 2), whereas BIC gives preference to the shared-parameter model (Model 1). Further, in case of dropout model (7) however, both information criteria select the shared-parameter model (Model 4). Note that, since the dropout model in Model 1 does not depend on the shared intercept, the dropout model and the measurement model are independent, resulting in the MCAR assumption, whereas in Model 2, the dropout model is linked to the measurement model through the latent classes (MNAR). Overall, the AIC criterion prefers Model 2, the 2-component latent-class mixture model with no random effect in the dropout model, whereas BIC picks Model 4, the classical shared parameter model. Since both criteria select a different model, we will take a more detailed look at the latent-class mixture model with two components, indicated by AIC, whereas we will consider the classical shared-parameter model in a sensitivity analysis in the next section.

Parameter estimates with corresponding standard errors and p -values of the two-component latent-class mixture model are shown in Table 3. Once this latent-class mixture model has been fitted to the depression trial data, the posterior probabilities can be used to classify the patients into two subgroups as shown in Section 4. The 170 patients split into 79 and 91 patients classified into the first and second group, respectively. In Figure 1, the left panel represents the individual profiles of patients classified into the first latent group, and the right one represents the individual profiles of patients classified into the second group. Clearly, the first group corresponds to patients with lower $HAMD_{17}$ scores, that continue to decrease over time. This means these are the patients getting better. On the other hand, the second group contains patients with a higher change versus baseline compared to the patients from the first group. Their changes of $HAMD_{17}$ score fluctuate around 0, more specifically somewhere in the region between -10 and 10 . In addition, without taking into account the within-subject variability, their profiles appear more or less time-constant. A more formal comparison of both latent groups regarding their change of $HAMD_{17}$ score versus baseline confirms this association between the classification and the profile over time.

Furthermore, a formal test for association of baseline values and group classification is not significant, indicating similar baseline $HAMD_{17}$ scores for patients in both groups.

Based on this difference in location of the profiles between both groups, this classification of subjects can be interpreted as being a split into acute versus chronic depression. Patients in both the acute and chronic groups enter the study with a baseline value indicating depression. However, the profiles of the patients in the acute group show recovery during the trial, whereas the depression score of patients in the chronic group remains more or less level.

Further, this difference between both latent groups is not due to treatment, since the classification of subjects in latent subgroups is independent of their treatment allocation. Indeed, the estimated odds ratio between the latent classification variable and the treatment allocation is 0.75, which was expected since the observed treatment groups are included in the mean structure of the measurement model. Moreover, when the treatment variable would be included in the dropout model, this independence would even increase.

Regarding the incompleteness of the patients in both latent groups, we notice a clear difference, which is confirmed by chi-square tests for independence, implying a significant association between the dropout pattern and the latent classification. The first latent group mainly contains patients who complete the study, 62 in total. Of the 17 patients who drop out, merely 2 drop out at visit 6, 3 more at visit 7, and 12 patients missed the last visit only. The dropout percentage in the second latent group is larger, 48.4% compared to 21.5% in the first group, or 44 out of 91 patients. Of these incompleters, 17 drop out after the first visit, 10 more at visit 6, 11 at the penultimate visit, and 6 more at the last visit.

Finally, the latent groups can also be compared by focusing on demographic characteristics such as age, gender, and origin, yielding no association between the latent classification with either gender or origin, but a significant association with age. Patients in the acute group are younger than the patients in the chronic group, with a mean age of 38.5 and 42.4, and corresponding 95% confidence intervals [36.1, 41.0] and [40.0, 44.7], respectively.

This is important insight, even though it may be hard to disentangle the causal relationship between age and chronicity. Even though age explains a part of the latent-class structure, it is relevant to further entertain the connection with chronicity, since this may have important implications for differentiated, more effective therapy.

However, as mentioned in Section 4, using this classification rule does not render insight into how sure the classification is in one of the two groups. This will depend on the magnitude of the maximal posterior probability. Since the latent-class mixture model considered here only contains two latent groups, we merely need to look at one of the posterior probabilities, e.g., at the posterior probability that the subject belongs to group 1, π_{i1} . Based on this π_{i1} , the subjects can be classified following the guidelines of Table 4. If the posterior probability π_{i1} lies between 0.45 and 0.55, it is uncertain to which group the subject can be classified. Only 8 out of 170 patients in the depression trial are in this situation. For most patients, 152 or 89.4%, it is clear into which group they can be classified, since their maximal posterior probability is above 0.60. Furthermore, aforementioned association of the latent classification with the location of profiles, the dropout pattern, and patient’s age as well as independence of baseline values and patient’s origin and gender, is confirmed by testing the independence of these variables with the posterior probabilities, which can be viewed as continuous variables ranging from 0 to 1.

6.2 A Sensitivity Analysis

In this section, we apply latent-class mixture models as a sensitivity analysis tool. In addition to the two-component latent-class mixture model shown in Section 6.1, a classical shared-parameter model will be fitted to the depression trial, as well as a pattern-mixture model, and two selection models, based on the selection models introduced by Diggle and Kenward (1994). All models contain the same fixed effects as in the two-component latent-class mixture model, i.e., intercept, treatment, time, baseline, time², and treatment-by-time interaction.

The classical shared-parameter model, selected by the BIC criterion in Section 6.1, includes a shared intercept $b_i \sim N(0, d^2)$, conditional upon which the measurement model follows a normal distribution $\mathbf{Y}_i | b_i \sim N(\mathbf{X}_i \beta + b_i, \sigma^2 \mathbf{I}_{n_i})$, and the dropout process is based on (7).

Next, the Diggle-Kenward (DK) model combines a multivariate normal model for the measurement process with a logistic regression model for the dropout process. More specifically, the measurement model assumes that the vector $\mathbf{Y}_i$ of repeated measurements for the i th subject satisfies the linear regression model $\mathbf{Y}_i \sim N(\mathbf{X}_i \beta, \mathbf{V}_i)$, $i = 1, \dots, N$. The matrix $\mathbf{V}_i$ can be left unstructured or assumed of a specific form. For the depression trial, the linear mixed model (Verbeke and Molenberghs, 2000) is used to model the measurement process, with an unstructured covariance matrix. Further, let $\mathbf{h}_{ij} = (y_{i1}, \dots, y_{i,j-1})$ denote the observed history of subject i up to time $t_{i,j-1}$. The DK model for the dropout process allows the conditional probability for dropout at occasion j , given that the subject was still observed at the previous occasion, to depend on the history $\mathbf{h}_{ij}$ and the possibly unobserved current outcome y_{ij} , but not on future outcomes y_{ik} , $k > j$. In the two models considered for the depression trial, the logistic dropout model will take the form

$$\text{logit}[P(D_i = j \mid D_i \geq j, \mathbf{h}_{ij}, y_{ij}, \boldsymbol{\Omega})] = \psi_0 + \psi_1 y_{i,j-1} + \psi_2 y_{ij}. \quad (8)$$

Regarding the missingness mechanism, the first selection model assumes the MAR assumption to hold, yielding $\psi_2 = 0$, whereas the second one assumes MNAR.

Finally, a pattern-mixture model is fitted, taking the same form as the DK models, except that there are separate intercepts and slopes for each of the dropout patterns occurring. Notice that the classification function in the latent-class mixture model is a data driven approach to define groups, whereas pattern-mixture models use the assumption to define groups in function of dropout patterns.

Since the main interest of the depression trial was in the treatment effect at the last visit, Table 5 shows the estimates, standard errors, and p -values for this effect under the five fitted

models. Clearly, the p -values resulting from all five models are very similar and between around 0.07 and 0.11, yielding the same conclusion for the treatment effect at visit 8. Thus, the significance results are not sensitive to the model used, and hence more trust can be put into the conclusion. This is because a deflated estimate is combined with a reduced standard error. However, note that using both the two-component latent-class mixture model and the classical shared-parameter model, the standard error is reduced by 0.3 units, compared to either selection model, or pattern-mixture model, resulting in a more accurate confidence interval for the treatment effect at the last visit.

Furthermore, we explore the sensitivity of the treatment-by-time interaction by comparing the estimates, standard errors and p -values under the five fitted models in Table 5. The p -values are clearly moving around the significance level of 0.05. Whereas under the latent-class mixture model and the shared-parameter model the p -value is about 0.03, the p -value under both selection models and the pattern-mixture model is around 0.07. While one should be cautious with over-interpretation of p -values, there are contexts, such as regulated clinical trials, where strict decision rules are implemented. In such a case and when in addition the treatment by time interaction is the primary effect, the latent-class mixture model and the shared-parameter model would lead to a claim of significance, whereas this would not be justified with neither the selection models nor the pattern-mixture model.

7 Concluding Remarks

We have proposed latent-class mixture models to analyze incomplete longitudinal data in which incompleteness is due to dropout of subjects. Through its structure, the model captures unobserved heterogeneity between latent subgroups of the population. It is an extension of the shared-parameter model, in the sense that both the measurement and dropout processes are allowed to share a set of random effects, conditional upon which both processes are assumed to be independent. It can, at the same time, be seen as an extension of the

pattern-mixture model, now with latent rather than explicitly observed groups. As shown in the simulation study, the flexibility of such latent-class mixture models outweighs the expected modelling complexity.

Our proposal can be used for flexible modeling, as a sensitivity analysis instrument, and for further exploration of the latent class membership. Of course, care has to be taken when interpreting latent classes, since in some applications they may merely be artifacts, without any substantive grounds. In others, there may be more basis for their existence. We believe, together with mental health scientists, the two-component classification in our example, refers to the natural split of the patients, regardless of which treatment they were allocated to, into the more chronic and the more acute ones. An additional word of caution is needed regarding the number of latent classes to be considered. This is a tricky but well documented problem (McLachlan and Peel 2000). A practical way out is to consider several choices for the number of components, pick the most reasonable one, and assess whether alternative choices would substantially alter the conclusions.

Evidently, the computational burden of the LCMM increases over non-latent-class models, but is still reasonable. For example, whereas the MNAR version of the Diggle and Kenward model takes around one hour and the one-component mixture needs about the same amount of time, the two-component mixture increases needs around one order of magnitude more. Furthermore, the performance of the algorithm is remarkably computationally stable, given sensible starting values (e.g., built from non-mixture classical models). Details on starting value selection are embedded in an electronically available companion manual.

Supplementary Materials

A Web appendix, referenced in Section 3, and a software manual, referenced in the Concluding Remarks, are available under the Paper Information link at the Biometrics website <http://www.tibs.org/biometrics>.

Acknowledgements

Caroline Beunckens, Geert Molenberghs and Geert Verbeke gratefully acknowledge the financial support from the IAP research Network P6/03 of the Belgian Government (Belgian Science Policy).

References

- Aitkin, M. and Rubin, D.B. (1985). Estimation and Hypothesis Testing in Finite Mixture Models. *Journal of the Royal Statistical Society, Series B*, **47**, 67–75.
- Böhning, D. (1999). *Computer Assited Analysis of Mixtures (C.A.M.A.N.)*. Marcel and Decer Inc., New York.
- Dempster, A.P., Laird, N.M., and Rubin, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society, Series B*, **39**, 1–38.
- Diggle, P.D., Kenward, M.G. (1994). Informative dropout in longitudinal data analysis (with discussion). *Applied Statistics*, **43**, 49–93.
- Jansen, I., Beunckens, C., Molenberghs, G., Verbeke, G. and Mallinckrodt, C. (2006). Analyzing incomplete binary longitudinal clinical trial data. *Statistical Science*, **21**, 52–69.
- Laird, N.M. (1994). Discussion to Diggle, P.J. and Kenward, M.G.: Informative dropout in longitudinal data analysis. *Applied Statistics*, **43**, 84.
- Little, R.J.A. (1993). Pattern-mixture models for multivariate incomplete data. *Journal of the American Statistical Association*, **88**, 125–134.
- Little, R.J.A. (1994). A class of pattern-mixture models for normal incomplete data. *Biometrika*, **81**, 471–483.

- Little, R.J.A., Rubin, D.B. (1987) *Statistical Analysis with Missing Data*. New York: Wiley.
- McLachlan, G.J. and Peel, D. (2000) *Finite mixture models*. New York: Wiley.
- Molenberghs, G., Kenward, M.G., and Lesaffre, E. (1997). The analysis of longitudinal ordinal data with non-random dropout. *Biometrika*, **84**, 33–44.
- Molenberghs, G., Thijs, H., Jansen, I., Beunckens, C., Kenward, M.G., Mallinckrodt, C., and Carroll, R.J. (2004). Analyzing incomplete longitudinal clinical trial data. *Biostatistics*, **5**, 445–464.
- Redner, R.A. and Walker, H.F. (1984). Mixture densities, maximum likelihood, and the EM algorithm. *SIAM Review*, **26**, **2**, 195–239.
- Rubin, D.B. (1976). Inference and missing data. *Biometrika*, **63**, 581–592.
- Ten Have, T.R., Kunselman, A.R., Pulkstenis, E.P., and Landis, J.R. (1998). Mixed effects logistic regression models for longitudinal binary response data with informative drop-out. *Biometrics*, **54**, 367–383.
- Verbeke, G. and Lesaffre, E. (1996). A linear mixed-effects model with heterogeneity in the random-effects population. *Journal of the American Statistical Association*, **91**, 217–222.
- Verbeke, G. and Molenberghs, G. (2000) *Linear Mixed Models for Longitudinal Data*. New York: Springer-Verlag.
- Wu, M.C. and Bailey, K.R. (1989). Estimation and comparison of changes in the presence of informative right censoring: conditional linear model. *Biometrics*, **45**, 939–955.
- Wu, M.C. and Carroll, R.J. (1988). Estimation and comparison of changes in the presence of informative right censoring by modelling the censoring process. *Biometrics*, **44**, 175–188.

Table 1: *Simulation Study. Results of the simulation study: mean and true value, bias, and mean squared error (MSE) of the parameters, under the four simulations settings.*

Setting 1					Setting 2				
Effect	Mean	True	Bias	MSE	Effect	Mean	True	Bias	MSE
<i>Measurement Model</i>					<i>Measurement Model</i>				
β_0	9.37	9.40	-2.84×10^{-2}	8.07×10^{-4}	β_0	9.34	9.40	-5.75×10^{-2}	3.31×10^{-3}
β_1	2.25	2.25	1.30×10^{-4}	1.68×10^{-8}	β_1	2.25	2.25	7.56×10^{-4}	5.72×10^{-7}
σ	0.25	0.25	-2.49×10^{-4}	6.18×10^{-8}	σ	0.75	0.75	6.27×10^{-4}	3.93×10^{-7}
μ_1	-4.39	-4.40	1.31×10^{-2}	1.73×10^{-8}	μ_1	-4.36	-4.40	4.48×10^{-2}	2.00×10^{-3}
d	1.98	2.00	-1.70×10^{-2}	2.89×10^{-4}	d	1.97	2.00	-2.53×10^{-2}	6.38×10^{-4}
π_1	0.60	0.60	4.60×10^{-4}	2.12×10^{-7}	π_1	0.60	0.60	4.22×10^{-3}	1.79×10^{-5}
<i>Dropout Model</i>					<i>Dropout Model</i>				
γ_1	-2.52	-2.50	-2.28×10^{-2}	5.19×10^{-4}	γ_1	-2.51	-2.50	-1.26×10^{-2}	1.58×10^{-4}
γ_2	-1.26	-1.25	-1.23×10^{-2}	1.53×10^{-4}	γ_2	-1.27	-1.25	-2.30×10^{-2}	5.27×10^{-4}
Setting 3					Setting 4				
Effect	Mean	True	Bias	MSE	Effect	Mean	True	Bias	MSE
<i>Measurement Model</i>					<i>Measurement Model</i>				
β_0	9.44	9.40	3.83×10^{-2}	1.46×10^{-3}	β_0	9.59	9.40	1.92×10^{-1}	3.70×10^{-3}
β_1	2.25	2.25	1.91×10^{-4}	3.66×10^{-8}	β_1	2.24	2.25	-1.44×10^{-2}	2.06×10^{-4}
σ	0.99	1.00	-5.45×10^{-3}	2.06×10^{-5}	σ	2.01	2.00	6.07×10^{-3}	3.69×10^{-5}
μ_1	-4.69	-4.40	-2.86×10^{-1}	8.18×10^{-2}	μ_1	-4.84	-4.40	-4.39×10^{-1}	1.93×10^{-1}
d	3.43	3.50	-7.00×10^{-2}	4.90×10^{-3}	d	6.02	6.00	2.03×10^{-2}	4.10×10^{-4}
π_1	0.57	0.60	3.36×10^{-2}	1.13×10^{-3}	π_1	0.52	0.60	-8.06×10^{-2}	6.50×10^{-3}
<i>Dropout Model</i>					<i>Dropout Model</i>				
γ_1	-2.61	-2.50	-1.07×10^{-1}	1.14×10^{-2}	γ_1	-2.97	-2.50	-4.73×10^{-1}	2.23×10^{-1}
γ_2	-1.27	-1.25	-2.04×10^{-2}	4.17×10^{-4}	γ_2	-1.29	-1.25	-3.89×10^{-2}	1.51×10^{-3}

Table 2: *Depression Trial. Information criteria AIC and BIC, for models with dropout model (6) or (7), and $g = 1, 2, 3$.*

Model	Dropout Model	g	# Par	-2ℓ	AIC	BIC
1	$\gamma_{0,k} + \gamma_{1,k} t_j$	1	10	4676.07	4696.08	4727.44
2	$\gamma_{0,k} + \gamma_{1,k} t_j$	2	14	4662.37	4690.37	4734.27
3	$\gamma_{0,k} + \gamma_{1,k} t_j$	3	18	4662.03	4698.03	4754.48
4	$\gamma_{0,k} + \gamma_{1,k} t_j + \lambda b_i$	1	11	4669.12	4691.12	4725.61
5	$\gamma_{0,k} + \gamma_{1,k} t_j + \lambda b_i$	2	15	4662.02	4692.02	4739.06

Table 3: *Depression Trial. Parameter estimates, standard errors, and p-values for the latent-class mixture model applied to the depression trial.*

Effect	Estimate	s.e.	p-value
<u>Measurement Model</u>			
Intercept : β_0	23.17	3.75	< 0.0001
Treatment : β_1	2.69	1.49	0.072
Time : β_2	-6.18	1.18	< 0.0001
Time $\times$ Treatment : β_3	-0.52	0.24	0.028
Baseline : β_4	-0.42	0.07	< 0.0001
Time $\times$ Time : β_5	0.41	0.10	< 0.0001
Measurement Error : σ	4.24	0.13	< 0.0001
<u>Dropout Model</u>			
Intercept Group 1 : $\gamma_{0,1}$	-8.58	3.57	0.009
Time Group 1 : $\gamma_{1,1}$	0.83	0.44	0.056
Intercept Group 2 : $\gamma_{0,2}$	-1.35	1.28	0.292
Time Group 1 : $\gamma_{1,2}$	-0.05	0.20	0.793
<u>Shared Effects</u>			
Mean Shared Intercept Group 1 : μ_1	-3.64	0.43	< 0.0001
Variance Shared Intercept : d	2.67	0.50	< 0.0001
Prior probability Group 1 : $\pi_1 = \pi$	0.48	0.10	< 0.0001
Loglikelihood	-2331.18		

Table 4: *Depression Trial. Classification of subjects based on the magnitude of posterior probabilities π_{i1} .*

π_{i1}	Classification	# Patients
$0.80 \rightarrow 1.00$	Clearly Group 1	61
$0.60 \rightarrow 0.80$	Group 1	8
$0.55 \rightarrow 0.60$	Doubtful, more likely Group 1	5
$0.45 \rightarrow 0.55$	Uncertain	8
$0.40 \rightarrow 0.45$	Doubtful, more likely Group 2	5
$0.20 \rightarrow 0.40$	Group 2	19
$0.00 \rightarrow 0.20$	Clearly Group 2	64

Table 5: *Depression Trial. Estimates, standard errors, and p-values for the treatment effect at visit 8, as well as the treatment-by-time interaction, for the latent-class mixture model (Model 4 in Table 2), the shared-parameter model (Model 2 in Table 2), the pattern-mixture model, and both selection models, assuming either MAR (which equals MCAR) or MNAR.*

Model	Treatment at Endpoint			Treatment $\times$ Time		
	Estimate	s.e.	p-value	Estimate	s.e.	p-value
Latent-Class Mixture Model	-1.44	0.91	0.114	-0.52	0.23	0.028
Shared-Parameter Model	-1.69	0.93	0.069	-0.50	0.24	0.035
Pattern-Mixture Model	-2.01	1.20	0.096	-0.55	0.31	0.077
MCAR $\equiv$ MAR Selection Model	-2.17	1.25	0.082	-0.58	0.32	0.068
MNAR Selection Model	-2.16	1.24	0.081	-0.57	0.31	0.068

Figure 1: *Depression Trial. Classification of the subjects of the depression trial based on a latent-class mixture model. The left panel corresponds to patients classified into first group, the right panel to patients classified into second one.*