

Question 2:

The evaluation of new compounds for the treatment of schizophrenia often involve longitudinal measurements of disease symptoms. The Positive and Negative Syndrome Scale (PANSS) is an instrument devised to measure changes in psychopathology of psychotic (usually schizophrenic) conditions over time. It is now the predominant instrument used to measure change in pharmacological trials of antipsychotic agents.

The data consist of repeated measurements on 517 subjects, with roughly 87 subjects in each of 6 treatment groups. In this analysis we will be comparing 3 treatment groups; placebo, haloperidol and risperidone. The risperidone treatment group is formed by combining 2 of the 6 original treatment groups. This results in a total of 344 subjects that consider here. The PANSS score was measured on subjects over a three month period, resulting in possible total of 7 measurements at -1, 0, 1, 2, 4, 6, and 8 weeks. A larger score indicates a poorer response.

As we shall see in part (b) there is considerable dropout among the subjects. Throughout the analysis we assume that missingness in the data is monotone. That is, we assume that once a subject drops out (i.e. they miss a scheduled visit) then they remain off study. Among the 344 subjects that we consider, three of them had a subsequent visit after having missed a previous visit. To maintain the assumption of monotonicity the data on subsequent visits was treated as missing. In addition, 2 subjects had their final visit during week -1 (pre-randomisation) and are consequently not included in the analysis. This leaves 342 subjects.

- (a) Figure 2 provides a plot of the scores for the 342 subjects in the PANSS data set. Included on the plot are estimates of the mean trends for the three treatment groups that we consider here.




Figure 2: PANSS scores over time (weeks), along with mean trends for each of the three treatment groups.

The empirical variance-covariance and correlation matrices are provided below. We can see from the variance-covariance matrix that the variance of the scores seems to increase with time. This suggests that a random slopes component to the error structure may be appropriate. From the empirical correlation matrix we see that the correlation between observations seem to decrease the further apart we take them. Although the correlation is decreasing, they do not appear to go to zero suggesting that they reach an asymptote of roughly 0.4. This suggests that a combination of random intercepts

and a serial component may be appropriate. The empirical variogram given in Figure 3 confirms this.

Empirical Variance-Covariance Matrix:

	week -1	week 0	week 1	week 2	week 4	week 6	week 8
week -1	247	215	181	177	171	161	155
week 0	0	378	259	259	273	263	203
week 1	0	0	405	378	354	338	224
week 2	0	0	0	506	454	435	278
week 4	0	0	0	0	570	538	422
week 6	0	0	0	0	0	612	508
week 8	0	0	0	0	0	0	551

Empirical Correlation Matrix:

	week -1	week 0	week 1	week 2	week 4	week 6	week 8
week -1	1	0.70	0.57	0.50	0.46	0.41	0.42
week 0	0	1	0.66	0.59	0.59	0.55	0.44
week 1	0	0	1	0.84	0.74	0.68	0.48
week 2	0	0	0	1	0.85	0.78	0.53
week 4	0	0	0	0	1	0.91	0.75
week 6	0	0	0	0	0	1	0.87
week 8	0	0	0	0	0	0	1




Figure 3: Variogram for the PANSS Data.

From Figure 2, there seems to be convincing evidence that the three treatment regimes are different with respect to mean PANSS score over the course of 8 weeks. However, as we can see from the subsequent EDA attention will need to be paid to the error structure to ensure that standard error estimates are valid.

- (b) As Figure 2 indicates there is a considerable amount of dropout in the data. In particular, we can clearly see that the data at week 8 is more sparse than the data at week 0. Table 1 provides numerical summaries to describe the dropout by treatment group. In particular, for each week and each treatment arm the number of subjects who have their last visit during that week is provided. We see that in the placebo arm a total of 29 subjects had their last visit during the 8th week (i.e. they completed the study). In addition we see that 21 subjects in the risperidone arm had their last visit

Week	0	1	2	4	6	8
Number						
Placebo	2	15	14	16	11	29
Haloperidol	2	11	10	18	5	41
Risperidone	1	16	13	21	14	103
Percent						
Placebo	2.3	17.2	16.1	18.4	12.6	33.3
Haloperidol	2.3	12.6	11.5	20.7	5.7	47.1
Risperidone	0.6	9.5	7.7	12.5	8.3	61.3
Cumulative Percent						
Placebo	2.3	19.5	35.6	54.0	66.6	100.0
Haloperidol	2.3	14.9	26.4	47.1	52.8	100.0
Risperidone	0.6	10.1	17.8	30.3	38.6	100.0

Table 1: Number/Percent/Cumulative Percent of subjects for whom the last visit was during a specified week, by treatment group.

during week 4, and therefore dropped out during the 6th week.

As the group sizes are not equal, Table 1 also provides the distribution (via percentages) of the dropout for each treatment group. We can see that fewer subjects complete the study in the placebo group (33%) than either of the other two arms (47% and 61%). We can also see from the cumulative percentages that there is more dropout earlier on for the haloperidol group and, especially, the placebo group than the risperidone group.

- (c) To investigate the missingness mechanism we can use a continuation ratio model. For this model we consider the categorical outcome of dropout time. The CRM models the conditional probability of dropout, given that the subject did not dropout before hand. The conditional probability is modelled as a function of *level* (current score) and *trend* (change in score since previous visit), where we assume that the effect of level and trend is independent of the current state. The output below provides a series of logistic regressions, modelling the conditional probabilities independently, as well as the full conditional probability. From the individual logistic regressions we can see that the assumption of constant 'level' effect is not unreasonable across the times. With the exception of the impact at week 0, the assumption of constant 'trend' effect seems reasonable.

```

Time: 0
      Value Std. Error t value
(Intercept) -4.844      2.621  -1.848
level      0.007      0.029   0.262
trend     -0.014      0.044  -0.318

Time: 1
      Value Std. Error t value
(Intercept) -6.314      1.106  -5.710
level      0.046      0.011   4.218
trend      0.052      0.016   3.343

Time: 2
      Value Std. Error t value
(Intercept) -5.235      0.980  -5.343
level      0.037      0.010   3.537
trend      0.018      0.016   1.147

```

Time: 4

	Value	Std. Error	t value
(Intercept)	-3.499	0.743	-4.712
level	0.024	0.009	2.817
trend	0.046	0.014	3.240

Time: 6

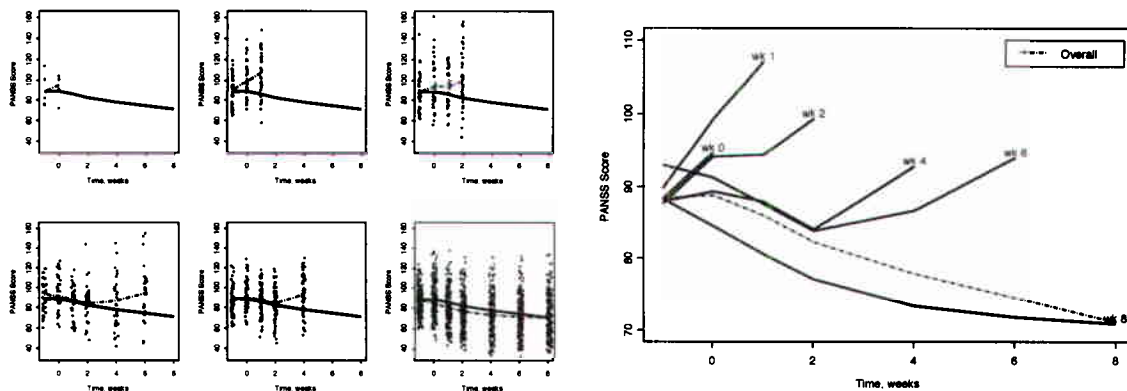
	Value	Std. Error	t value
(Intercept)	-4.766	0.921	-5.173
level	0.034	0.010	3.287
trend	0.046	0.019	2.420

Full Continuation Ratio Model

	Value	Std. Error	t value
(Intercept)	-7.886	0.666	-11.832
time.f1	2.977	0.495	6.019
time.f2	3.024	0.500	6.048
time.f4	3.704	0.492	7.530
time.f6	3.320	0.511	6.495
level	0.032	0.005	6.712
trend	0.039	0.008	5.034

The results from the CRM indicate that there is a significant association between dropout time and both level and trend. In particular, we find that, at each week, the odds of dropping out, given that the subject has not dropped out so far, increase with both level and trend. We can interpret the level coefficient as follows. If we consider two individuals whose current score (level) differ by 10 units, holding trend constant, then the individual with the higher score will have odds of dropping out that are approximately 38% ($\exp^{10 \times 0.032} \approx 1.377$) higher than the odds for the individual with the lower score.

Both of the positive coefficients suggest that subjects who are doing worse, either having a higher current score or having had an increase in score from the previous visit, tend to have higher odds of dropping out.



(a) Longitudinal response patterns given dropout times; 0,1,2,4,6,8. Each plot also contains the overall mean trend.

(b) Mean response until dropout for each dropout time, as well as the overall mean trend.

Figure 4: Pattern Mixture Model for dropout.

- (d) An alternative to the CRM is to adopt a pattern mixture model for the dropout. Figure 4 provides plots that examine the longitudinal response patterns for each dropout time. In particular, Figure 4(b) shows the mean PANSS score trends that are specific to individuals that dropout at each visit. We see that for the subjects that complete the study, i.e. drop out after week 8, that the score pattern resembles the overall score pattern. For each of the previous visits, we see the mean trend for those that dropout on the subsequent visit. For each visit prior to the end of study we see that the score trends are increasing indicating that subjects who drop out were doing progressively worse on previous visit.
- (e) If we assume that the missing data in the PANSS study is MAR (Missing-at-Random), then we can implement likelihood-based methods (as long as we are confident in the assumptions of the model). In particular, we can use mixed effects models to model the repeated outcomes for each subject. Interest lies in assessing differences at 8 weeks between the three treatment regimes. Assuming a linear time effect (centered at week 8) and a three level categorical treatment we can use an interaction between time and treatment to answer the primary question of interest. In each of the following models we can interpret the main effects for treatment (given by tx1 for haloperidol and tx2 for risperidone) as the impact of each treatment relative to placebo at 8 weeks. In addition, due to randomization of the treatment regimes, all data that is pre-randomization (i.e. week -1) has been removed for these analyses.

Below we examine the impact of making various assumptions regarding the form of the error structure. Decisions can, in part, be guided by the EDA of part (a). Each model is fitted via maximum likelihood.

Random Intercepts

```
Random effects:
Formula: ~ 1 | id
          (Intercept) Residual
StdDev:    17.62421 12.30737

Fixed effects: score ~ time8 * tx
              Value Std.Error DF   t-value p-value
(Intercept)  89.15333  2.624716 1263   33.96685 <.0001
time8       -0.24754  0.285833 1263   -0.86603  0.3866
tx1         -9.48725  3.604728  339   -2.63189  0.0089
tx2        -18.25608  3.115421  339   -5.85991 <.0001
time8tx1     -1.21417  0.384459 1263   -3.15813  0.0016
time8tx2     -1.84880  0.330510 1263   -5.59377 <.0001
```

Random Intercepts + Random Slopes

```
Random effects:
Formula: ~ time8 | id
Structure: General positive-definite
              StdDev  Corr
(Intercept) 25.142194 (Intercept)
time8       2.228351 0.739
Residual   10.745259

Fixed effects: score ~ time8 * tx
              Value Std.Error DF   t-value p-value
(Intercept)  93.69646  3.465718 1263   27.03522 <.0001
time8        0.39910  0.401963 1263    0.99289  0.3210
tx1        -11.54986  4.776369  339   -2.41813  0.0161
tx2        -21.66569  4.122133  339   -5.25594 <.0001
time8tx1     -1.50970  0.547602 1263   -2.75694  0.0059
time8tx2     -2.33377  0.470513 1263   -4.96004 <.0001
```

Random Intercepts + AR(1)

Random effects:

Formula: ~ 1 | id
(Intercept) Residual

StdDev: 15.80355 14.82967

Correlation Structure: Exponential spatial correlation

Formula: ~ time8 | id

Parameter estimate(s):

range

1.961153

Fixed effects: score ~ time8 * tx

	Value	Std.Error	DF	t-value	p-value
(Intercept)	89.50163	2.941718	1263	30.42495	<.0001
time8	-0.27638	0.381871	1263	-0.72374	0.4694
tx1	-10.14650	4.009237	339	-2.53078	0.0118
tx2	-18.93131	3.454840	339	-5.47965	<.0001
time8tx1	-1.31771	0.517540	1263	-2.54611	0.0110
time8tx2	-2.07011	0.445345	1263	-4.64832	<.0001

Random Intercepts + AR(1) + Measurement Error

Random effects:

Formula: ~ 1 | id
(Intercept) Residual

StdDev: 0.9878355 21.47729

Correlation Structure: Exponential spatial correlation

Formula: ~ time8 | id

Parameter estimate(s):

range nugget

11.41318 0.1329292

Fixed effects: score ~ time8 * tx

	Value	Std.Error	DF	t-value	p-value
(Intercept)	90.65820	3.025782	1263	29.96191	<.0001
time8	-0.11122	0.405437	1263	-0.27431	0.7839
tx1	-10.96340	4.122918	339	-2.65914	0.0082
tx2	-19.08215	3.549864	339	-5.37546	<.0001
time8tx1	-1.46082	0.552177	1263	-2.64556	0.0083
time8tx2	-2.15802	0.475421	1263	-4.53918	<.0001

Random Intercepts + Random Slopes + AR(1) + Measurement Error

Random effects:

Formula: ~ time8 | id
Structure: General positive-definite
StdDev Corr

(Intercept)	19.976803	(Inter
time8	1.158248	1
Residual	16.234720	

Correlation Structure: Exponential spatial correlation

Formula: ~ time8 | id

Parameter estimate(s):

range nugget

5.27029 0.2199205

Fixed effects: score ~ time8 * tx

	Value	Std.Error	DF	t-value	p-value
(Intercept)	92.69179	3.445997	1263	26.89840	<.0001
time8	0.17616	0.417354	1263	0.42208	0.6730
tx1	-11.04393	4.723844	339	-2.33791	0.0200
tx2	-20.69583	4.072425	339	-5.08194	<.0001
time8tx1	-1.45692	0.567228	1263	-2.56849	0.0103
time8tx2	-2.33094	0.487618	1263	-4.78025	<.0001

From the output from each of the models we see that the estimation of the treatment effects for both haloperidol and risperidone (relative to placebo) do not depend to a great extent on the choice of the error structure. The following output provide comparisons of the above five models, in terms of their ability to fit the data. Both of the AIC and BIC criteria indicate that the best overall fit of the data is provided by model 5, which is the most complex of the error structures.

Model	df	AIC	BIC	logLik	Test	L.Ratio	p-value
1. Random Intercepts + ME	8	13441.34	13484.40	-6712.668			
2. Random Intercepts/Slopes + ME	10	13321.88	13375.71	-6650.941	1 vs 2	123.4535	<.0001
3. Random Intercepts + AR(1)	9	13317.16	13365.61	-6649.582	2 vs 3	2.7188	0.0992
4. Random Intercepts + AR(1) + ME	10	13292.95	13346.78	-6636.476	3 vs 4	26.2119	<.0001
5. Random Intercepts/Slopes + AR(1) + ME	12	13273.98	13338.57	-6624.988	4 vs 5	22.9755	<.0001

- (f) We can examine the residuals (population level) to assess assumptions of the model. In particular, we can assess the linearity of the time effect within each of the treatment groups by examining the residuals plots in Figure 5.

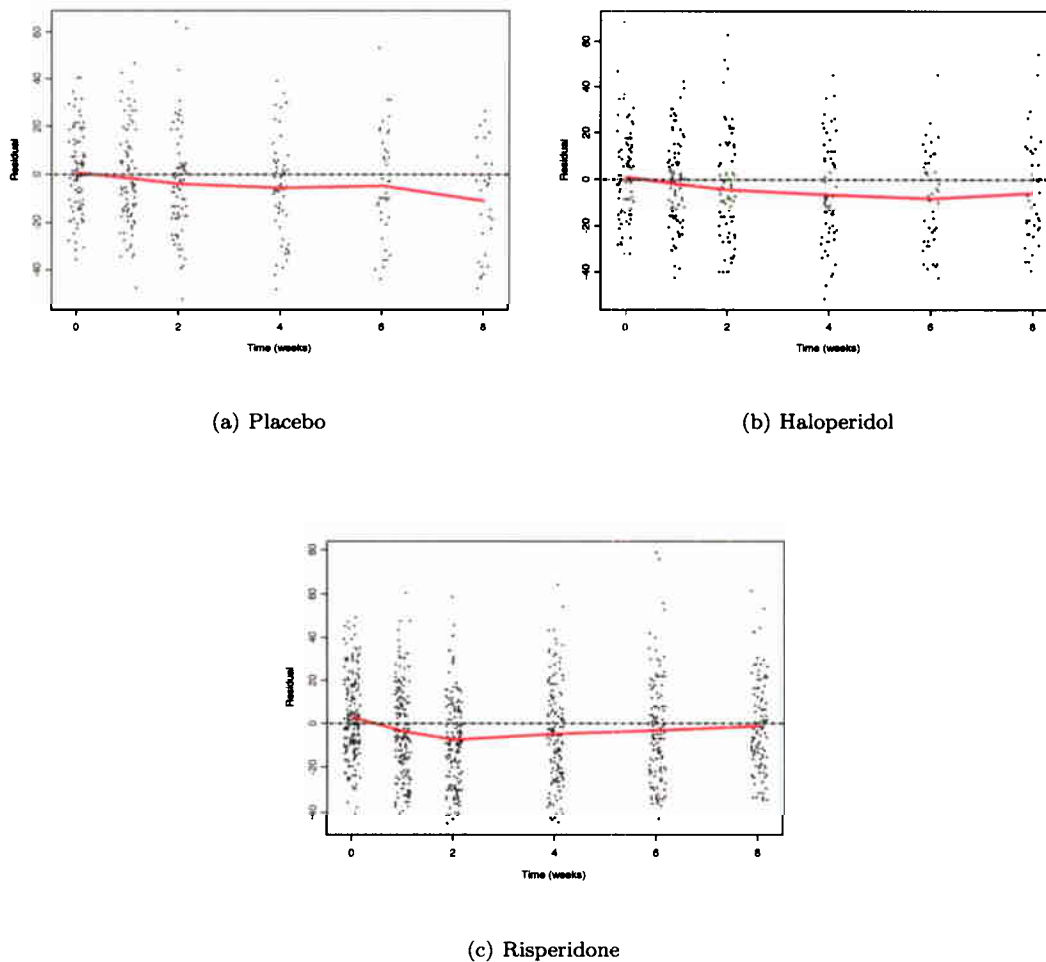


Figure 5: Plots of (population-level) residuals versus time (weeks), by treatment group.

From Figure 5 there seems to be an indication, especially for the risperidone group, that the linearity assumption may not be adequate and that a quadratic term may be needed. Assuming the error structure identified in part (e), we can assess this directly by adding in a quadratic term into the model. The following provides the resulting output and likelihood ratio test (of the quadratic time components) based on ML fits of the mixed effects model.

```
Random effects:
Formula: ~ time8 | id
Structure: General positive-definite
          StdDev  Corr
(Intercept) 21.322210 (Intercept)
time8      1.499273 1
Residual    16.574307

Correlation Structure: Exponential spatial correlation
Formula: ~ time8 | id
Parameter estimate(s):
range  nugget
7.730288 0.2373044
Fixed effects: score ~ (time8 + time8^2) * tx
              Value Std.Error  DF  t-value p-value
(Intercept)  96.14975   3.891712 1260   24.70629  <.0001
time8       1.94723   1.144486 1260    1.70140  0.0891
I(time8^2)   0.17518   0.115329 1260    1.51893  0.1290
tx1        -10.14795   5.271703  339   -1.92498  0.0551
tx2        -18.79110   4.548034  339   -4.13170  <.0001
time8tx1     -0.29647   1.524215 1260   -0.19451  0.8458
time8tx2     0.60625   1.314868 1260    0.46107  0.6448
I(time8^2)tx1 0.14260   0.155786 1260    0.91538  0.3602
I(time8^2)tx2 0.36736   0.134712 1260    2.72703  0.0065
```

Likelihood Ratio Test:

Model	df	AIC	BIC	logLik	Test	L.Ratio	p-value
1. Linear	12	13273.98	13338.57	-6624.988			
2. Quadratic	15	13212.92	13293.66	-6591.462	1 vs 2	67.05266	<.0001

From the results of the likelihood ratio test we see that there is evidence that the quadratic term is needed in the model. Due to the centering of the time variable, the inclusion of the quadratic terms do not impact the interpretations of the treatment main effects from those that are desirable for the scientific question of interest.

The final model that we adopt incorporates time (centered at week 8) via a quadratic model and treatment via a three level factor variable. To address the primary question of interest we can include an interaction between time and treatment, and concentrate on estimation/inference regarding the main effects for treatment. We adopt a fairly complex error structure for the repeated measure, which includes both random intercepts and slopes, as well as a serial autocorrelation component and finally measurement error.

Before interpreting the individual coefficients, we can examine the treatment effects by performing a likelihood ratio test comparing the above model (Alternative) to one that assumes no treatment effect (Null):

Model	df	AIC	BIC	logLik	Test	L.Ratio	p-value
Null	9	13242.69	13291.13	-6612.344			
Alternative	15	13212.92	13293.66	-6591.462	1 vs 2	41.7652	<.0001

The results of the global test of the treatment effect we see that is strong evidence to indicate a difference between the three treatment regimes.

From the output of the model above, we find that if we compare two individuals where one was randomised to placebo and the other to haloperidol (tx1), then the subject on haloperidol is estimated to have a expected PANSS score at week 8 which is approximately 10.15 units lower than the subject on placebo. Also, if we compare two individuals where one was randomised to placebo and the other to risperidone (tx2), then the subject on risperidone is estimated to have an expected PANSS score at week 8 which is approximately 18.79 units lower than the subject on placebo.

The following table provides the estimated expected week 8 PANSS scores for each of the treatment groups, along with approximate 95% confidence intervals.

Treatment Group	Expected PANSS score at week 8	Approximate 95% CI
Placebo	96.15	(88.52, 103.78)
Haloperidol	86.00	(79.03, 92.97)
Risperidone	77.36	(72.75, 81.97)

Any conclusions that we draw from this analysis are based on the two fairly strong assumptions. The first is that the missingness is MAR. From parts (c) and (d) we saw that we can identify fairly strong upward trends in patient PANSS scores just before dropout. This suggests that patients are dropping out (or perhaps being taken off the study) due to a worsening in their symptoms. Since we are able to identify such a reasonable mechanism from the data it is not unreasonable to assume that contributions from other mechanisms may be minimal. Of course, and unfortunately, there is no way for us to assess this explicitly. Given that we are happy with the MAR assumption, we still have to assume that we have specified the likelihood fully and correctly. This translates into correct specification of the mean and covariance structure. Given our specific mean structure, and although not shown here, the results (estimates/inference) do not vary greatly with the assumptions of the covariance structure.

S-Plus Code

```
#
## Schizophrenia trial data: PANSS scores for patients in six treatment groups.
##   Treatment codes: 1 = haloperidol    2 = placebo
##                   3 = risperidone10  4 = risperidone16
##                   5 = risperidone2    6 = risperidone6
##   Variables in column order:
##       group      = treatment group
##       score.m1    = score at time = -1
##       score.0     = score at time = 0
##       score.1     = score at time = 1
##       score.2     = score at time = 2
##       score.4     = score at time = 4
##       score.6     = score at time = 6
##       score.8     = score at time = 8
## Clinical References: Chouinard et al. (1993)
##                     Marder and Meibach (1994) Am J Psychiatry 151(6):825-35
#
options( contrasts = c("contr.treatment", "contr.poly") )
source( "stacked2wide.q" )
source( "variogram.q" )
library( nlme3 )
#
panss.full <- read.table( "C:\\TA\\571_W03\\homework\\exercise7\\code\\panss_data.txt" )
names( panss.full ) <- c( "group", "score.m1", "score.0", "score.1", "score.2", "score.4", "score.6", "score.8" )
panss.full$id <- 1:517
```

```

#
#### Only consider those in the following 3 treatment groups: placebo, haloperidol, risperidone (risperidone 6 and 10)
#
panss.full$tx <- rep( 9, 517 )
panss.full$tx[panss.full$group == 2] <- 0          # 0 = placebo
panss.full$tx[panss.full$group == 1] <- 1          # 1 = haloperidol
panss.full$tx[panss.full$group == 3 | panss.full$group == 6] <- 2      # 2 = risperidone
panss.full <- panss.full[panss.full$tx != 9,]
#
#### Investigate the monotonicity of the missingness; i.e. if a subject drops out they don't come back.
#
for( i in 1:344 ){
  nonmono <- F
  for( j in 3:8 ){ if( is.na(panss.full[i,(j-1)]) == T && is.na(panss.full[i,j]) == F ) nonmono <- T }
  if( nonmono == T ) print( panss.full[i,] )
}
#
#### Subjects 44, 68, 220 have non-monotone missingness
#
panss.full$score.4[44] <- panss.full$score.4[68] <- panss.full$score.1[220] <- NA
#
#### Convert dataset into a long form and remove observations that have the outcome missing (i.e. post-dropout):
#
panss.long <- matrix( 0, 344*7, 4 )
for( i in 1:344 ){
  panss.long[(((i-1)*7) + 1:7),] <- cbind( rep(i,7), rep(panss.full$tx[i],7), c(-1,0,1,2,4,6,8), unlist(as.vector(panss.full[i,2:8])) )
}
panss.long <- as.data.frame( panss.long )
names( panss.long ) <- c( "id", "tx", "time", "score" )
panss <- na.omit( panss.long )
#
#####
#### Part (a) ####
#####
#
plot( jitter(panss$time), panss$score, xlab = "Time, weeks", ylab = "PANSS Score", pch = "." )
lines( loess( panss$time[panss$tx == 0], panss$score[panss$tx == 0] ), col = 1, lwd = 3, lty = 1 )
lines( loess( panss$time[panss$tx == 1], panss$score[panss$tx == 1] ), col = 4, lwd = 3, lty = 3 )
lines( loess( panss$time[panss$tx == 2], panss$score[panss$tx == 2] ), col = 8, lwd = 3, lty = 4 )
legend( 5, 160, col = c(1,4,8), lwd = c(3,3,3), lty = c(1,3,4), c("placebo", "haloperidol", "risperidone") )
#
#### Observations per subject
#
n.obs <- unlist( lapply( split( panss$id, panss$id ), length ) )
table( n.obs )
table( panss$time )
#
#### Empirical correlation matrix
#
fit.sat <- lm( score ~ ns(time, knots = c(1, 4)) * as.factor(tx), data = panss )
resids <- panss$score - fitted( fit.sat )
rmat <- stacked2wide( panss$id, resids, panss$time, c(-1,0,1,2,4,6,8), 5 )
cmat <- matrix( 0, 7, 7 )
nmat <- matrix( 0, 7, 7 )
for( j in 1:7 ){
  for( k in j:7 ){
    njk <- sum( !is.na( rmat[ ,j]*rmat[ ,k] ) )
    sjk <- sum( rmat[ ,j]*rmat[ ,k], na.rm=T ) / njk
    cmatrix[j, k] <- sjk
    nmatrix[j, k] <- njk
  }
}
#
vvec <- diag( cmatrix )
cormat <- cmatrix / ( outer( sqrt( vvec ), sqrt( vvec ) ) )
dimnames( cmatrix ) <- dimnames( cormat ) <- dimnames( nmatrix ) <- list( paste("week ", c(-1,0,1,2,4,6,8), sep = "" ),
paste(" week ", c(-1,0,1,2,4,6,8), sep = "" ) )
print( round( cmatrix ) )
print( round( cormat, 2 ) )
print( nmatrix )
#
#### Variogram
#
panss.vario <- lda.variogram( panss$id, resids, panss$time )
var.est <- var( resids )
plot( jitter(panss.vario$delta.x), panss.vario$delta.y,
pch = ".", ylim = c( 0, 1.2*var.est ), xlab = "Change in time", ylab = "Change in residual over time squared" )
lines( smooth.spline( panss.vario$delta.x, panss.vario$delta.y, df = 5 ), lwd = 3 )
abline( h = var.est, lty = 3, lwd = 3 )
#
#
#

```

```
#####
### Part (b) ###
#####

### Visits completed by treatment group
#
final.visit <- unlist( lapply( split( panss$time, panss$id ), max ) )
tx.group <- unlist( lapply( split( panss$tx, panss$id ), max ) )
table( tx.group, final.visit )[, -1]
#
dropout.percent <- table( tx.group, final.visit )[, -1]
for( i in 1:3 ){ dropout.percent[i,] <- round( dropout.percent[i,] / sum(dropout.percent[i,]), 3 ) * 100 }
dropout.percent
#
dropout.cum <- dropout.percent
for( i in 1:3 ){ dropout.cum[i,] <- cumsum( dropout.cum[i,] ) }
dropout.cum
#
#
#####
### Part (c) ###
#####

### Construct data set with failure indicator; y
#
panss.crm <- panss
names( panss.crm ) <- c( "id", "tx", "time", "level" )
panss.crm$trend <- rep( NA, 1952 )
for( i in 1:1952 ){ if( panss.crm$time[i] != -1 ) panss.crm$trend[i] <- panss.crm$level[i] - panss.crm$level[(i-1)] }
panss.crm$y <- rep( 0, length(panss.crm[,1]) )
for( i in 1:1951 ){ if(panss.crm$id[i] != panss.crm$id[(i+1)]) panss.crm$y[i] <- 1 }
#
### Construct level-specific intercepts:
#
ncuts <- length( unique(panss.crm$time) )
panss.crm <- cbind( panss.crm, matrix( 0, length(panss.crm[,1]), ncuts ) )
names( panss.crm ) [7:13] <- c( paste( "Int", unique(panss.crm$time), sep = " " ) )
for( i in 7:13 ){ panss.crm[,i][panss.crm$time == unique(panss.crm$time)[(i-6)]] <- 1 }
#
panss.crm$time.f <- as.factor( panss.crm$time )
panss.crm <- panss.crm[panss.crm$time != -1 & panss.crm$time != 8,]
#
#
### Individual logistic regressions
#
for( i in 1:5 ){
  cat( paste( "\nTime:", unique(panss.crm$time)[i], "\n" ) )
  temp.panss.crm <- panss.crm[is.element(panss.crm$time, unique(panss.crm$time)[i]), ]
  temp.crm <- glm( y ~ level + trend, family = binomial, data = temp.panss.crm )
  print( round( summary(temp.crm)$coef, 3 ) )
}
#
### Full Continuation Ratio Model
#
crm.fit.0 <- glm( y ~ time.f + level + trend, family = binomial, data = panss.crm, na.action = na.omit )
round( summary(crm.fit.0)$coef, 3 )
crm.fit.1 <- glm( y ~ time.f + level + trend, family = binomial, data = panss.crm, na.action = na.omit )
round( summary(crm.fit.1)$coef, 3 )
#
#####
### Part (d) ###
#####

#
panss$dropout.group <- rep( final.visit, as.vector( table(panss$id) ) )
panss <- panss[(panss$dropout.group != -1),]
panss$dropout.group <- as.factor( panss$dropout.group )
PANSSgrouped <- groupedData( score ~ time | id, outer = ~ dropout.group,
                           data = panss,
                           labels = list( x = "Time, weeks", y = "PANSS Score" ) )
plot( PANSSgrouped, outer = ~ dropout.group, aspect = 1 )
#
dg <- unique( panss$dropout.group )
plot( panss$time, panss$score, xlab = "Time, weeks", ylab = "PANSS Score", type = "n", ylim = c(70,110) )
for( i in 1:length(dg) ){
  value <- loess(panss$time[panss$dropout.group == dg[i]], panss$score[panss$dropout.group == dg[i]])
  lines( value, lwd = 3 )
  text( x = value$x[length(value$x)], y = value$y[length(value$y)] + 1, paste("wk", as.character(dg[i])) )
}
lines( loess(panss$time, panss$score), lty = 3, lwd = 3 )
legend( 5.8, 110, lty = 3, lwd = 3, c("Overall") )
```

```

#
par( mfrow = c(2,3) )
for( i in length(dg):1 ){
  plot( panss$time, panss$score, xlab = "Time, weeks", ylab = "PANSS Score", type = "n" )
  points( jitter(panss$time[panss$dropout.group == dg[i]]), panss$score[panss$dropout.group == dg[i]], pch = "." )
  lines( lowess(panss$time, panss$score), lwd = 3 )
  lines( lowess(panss$time[panss$dropout.group == dg[i]], panss$score[panss$dropout.group == dg[i]]), lty = 3, lwd = 3 )
}
#
#####
#### Part (e) ####
#####
#
#### Create a new variable time8 = time - 8, so that the intercept in the model refers to the week 8 measurement
#
panss$time8 <- panss$time - 8
panss$tx <- as.factor( panss$tx )
panss.lme <- panss[panss$time != -1,]
#
#### Random intercepts:
#
fit.0 <- lme( fixed = score ~ time8 * tx,
             method = "ML",
             random = reStruct( ~ 1 | id, pdClass = "pdSymm", REML = F ),
             data = panss.lme )
summary( fit.0 )
#
#### Random intercepts and random slopes:
#
fit.1 <- lme( fixed = score ~ time8 * tx,
             method = "ML",
             random = reStruct( ~ time8 | id, pdClass = "pdSymm", REML = F ),
             data = panss.lme )
summary( fit.1 )
#
#### Random intercept and AR(1):
#
fit.2 <- lme( fixed = score ~ time8 * tx,
             method = "ML",
             random = reStruct( ~ 1 | id, pdClass = "pdSymm", REML = F ),
             correlation = corExp( form = ~ time8 | id, nugget = F ),
             data = panss.lme )
summary( fit.2 )
#
#### Random intercept and AR(1) and measurement error:
#
fit.3 <- lme( fixed = score ~ time8 * tx,
             method = "ML",
             random = reStruct( ~ 1 | id, pdClass = "pdSymm", REML = F ),
             correlation = corExp( form = ~ time8 | id, nugget = T ),
             data = panss.lme )
summary( fit.3 )
#
#### Random intercept/slopes and AR(1) and measurement error:
#
fit.4 <- lme( fixed = score ~ time8 * tx,
             method = "ML",
             random = reStruct( ~ time8 | id, pdClass = "pdSymm", REML = F ),
             correlation = corExp( form = ~ time8 | id, nugget = T ),
             data = panss.lme )
summary( fit.4 )
#
#### Comparison of models:
#
summary( fit.0 )
summary( fit.1 )
summary( fit.2 )
summary( fit.3 )
summary( fit.4 )
anova( fit.0, fit.1, fit.2, fit.3, fit.4 )
#
#####
#### Part (f) ####
#####
#
#### Residuals from fit.4: Population-level residuals
#
pop.res <- resid( fit.4, level = 0 )
#
plot( jitter(panss.lme$time[panss.lme$tx == 0]), sub.res[panss.lme$tx == 0], pch = ".",
      xlab = "Time (weeks)", ylab = "Residual" )

```

```

lines( smooth.spline( panss.lme$time[panss.lme$tx == 0], sub.res[panss.lme$tx == 0], df = 5 ), col = 8, lwd = 3 )
abline( h = 0, lwd = 3, lty = 3 )
#
plot( jitter(panss.lme$time[panss.lme$tx == 1]), sub.res[panss.lme$tx == 1], pch = ".",
      xlab = "Time (weeks)", ylab = "Residual" )
lines( smooth.spline( panss.lme$time[panss.lme$tx == 1], sub.res[panss.lme$tx == 1], df = 5 ), col = 8, lwd = 3 )
abline( h = 0, lwd = 3, lty = 3 )
#
plot( jitter(panss.lme$time[panss.lme$tx == 2]), sub.res[panss.lme$tx == 2], pch = ".",
      xlab = "Time (weeks)", ylab = "Residual" )
lines( smooth.spline( panss.lme$time[panss.lme$tx == 2], sub.res[panss.lme$tx == 2], df = 5 ), col = 8, lwd = 3 )
abline( h = 0, lwd = 3, lty = 3 )
#
#### Do we need the quadratic term ?
#
fit.5 <- lme( fixed = score ~ (time8 + time8^2) * tx,
             method = "ML",
             random = reStruct( ~ time8 | id, pdClass = "pdSymm", REML = F ),
             correlation = corExp( form = ~ time8 | id, nugget = T ),
             data = panss.lme )

summary( fit.5 )
#
anova( fit.4, fit.5 )
#
#### Assess the treatment effects:
#
fit.6 <- lme( fixed = score ~ time8 + time8^2,
             method = "ML",
             random = reStruct( ~ time8 | id, pdClass = "pdSymm", REML = F ),
             correlation = corExp( form = ~ time8 | id, nugget = T ),
             data = panss.lme )

summary( fit.6 )
#
anova( fit.6, fit.5 )
#
#### Compute expected PANSS scores at week 8 based on fit.5 (final model)
#
LinCom <- function( fit, x = 0, alpha = 0.05 ){
  coef.est <- fixef( fit )
  varcov.est <- fit$varFix * (fit$dim$N) / (fit$dim$N - length(coef.est))      ## Adjustment for ML estimation
  point <- x %*% coef.est
  point.se <- sqrt( t(x) %*% varcov.est %*% x )
  lower <- point - (qnorm(1 - (alpha/2)) * (point.se))
  upper <- point + (qnorm(1 - (alpha/2)) * (point.se))
  return( round( c(point, lower, upper), 2 ) )
}
LinCom(fit.5, x = c(1,0,0,0,0,0,0,0))
LinCom(fit.5, x = c(1,0,0,1,0,0,0,0))
LinCom(fit.5, x = c(1,0,0,0,1,0,0,0))

```