C H A P T E R 5

FITTING MODELS TO THE ECOLOGICAL DATA SET

5.1 Background to Ecological Data Set

The ecological data set originates from a study conducted in order to determine if the suppression of water hyacinth (*Eichhornia crassipes*) – a potent water weed – under biological control can be aided by spraying the plants with a sublethal dosage of herbicide during a single spraying event, and if this is affected by the nutrients (phosphates and nitrates) in the water (Kirton, 2005). A sublethal dosage of 0.8% of a glyphosate herbicide was applied to water hyacinth plants grown at three different nutrient levels. Control plants were grown in identical conditions, but without herbicide. Two pairs of water hyacinth weevils, *Neochetina eichhorniae*, were added to each plant. Plant growth parameters and insect feeding intensity were measured weekly over an eight week period, resulting in nine measurements of each variable per plant.

The purpose of spraying the plants with the sublethal dosage of herbicide was to stunt the growth of the plants, but not to negatively affect the biological control agents. If the dosage was too high, it could result in total loss of the plants, resulting in loss of the biological control agents as well, which would then allow new plants to grow without suppression from the biological control agents. The growth rate of unsprayed plants is dependent on the amount of nutrients in the water. Therefore, if the growth rate of sprayed plants also increases with increasing nutrient levels in the water, the success of this integrated method of control could be dependent on the amount of nutrients in the water.

Findings of the study were that sprayed plants stopped growing under the influence of a sublethal herbicide dose and nutrients had no effect on sprayed plant growth. For all nutrient levels, treated plants underwent very little asexual reproduction, and leaf production halted.

At the onset of the study, four water hyacinth plants, removed from the same pool, were placed into circular 50ℓ plastic tubs containing 42ℓ of water and a mix of nutrients, on a laboratory roof top at the University of the Witwatersrand. As the plants had been grown in the same conditions over a long period of time, the growth forms (which depend on the amount of space, nutrients and light available to the plants) of the plants were all the same, and the size of the plants were very similar. Three nutrient levels were used: low, 0.5mg Nitrate/ℓ and 0.08mg Phosphate/ℓ; medium, 1.5mg Nitrate /ℓ and 0.22mg Phosphate /ℓ; and high, 3mg Nitrate /ℓ and 0.43mg Phosphate /ℓ. These levels were chosen using country-wide water quality analyses from the South African Institute for Water Quality Service (IWQS) recording stations. The ratio of nitrogen to phosphorus was approximately 7:1 in each case, which has been suggested to yield optimum growth of water hyacinth plants (Wilson, 2002). Thirty tubs were used in total. These tubs were then divided, by means of a random allocation process, into the three nutrient groups. Half of the tubs, selected at random, in each nutrient group were sprayed with the sublethal dosage of herbicide at the beginning of the experiment, and the other half were sprayed with water. The experimental tubs were moved to a different location on the roof for spraying, and a screen was placed up to avoid spraying the control tubs. Five experimental tubs and five control tubs were used for each nutrient level. Water and nutrients were replaced weekly during the course of the experiment. The plastic tubs were individually enclosed in a net canopy to ensure that the weevils remained on the plants.

Four plants were placed into each tub to ensure that growth of the plants was still possible. Two *N. eichhorniae* weevil pairs were released onto each plant in all thirty tubs, resulting in an initial weevil density of 4 weevils/plant or 16 weevils/tub, matching field infestation rates. Two days before the herbicide was applied, two water hyacinth rosettes were randomly chosen from each of the tubs and tagged, and weekly measurements were made on these plants. The length of the petiole of the second youngest leaf, which is the distance from the point of attachment at the rhizome to the base of the lamina was one of the plant performance measurements taken. This is an important measure as it reveals information about the state of health and growth form of the plant.

5.2 Data Exploration

The ecological data set contains nine measurements on each of the sixty plants, resulting in 540 measurements in total, which is larger compared to the PR data set, which has a total of 27 subjects measured at four occasions each, resulting in a total of 108 measurements.

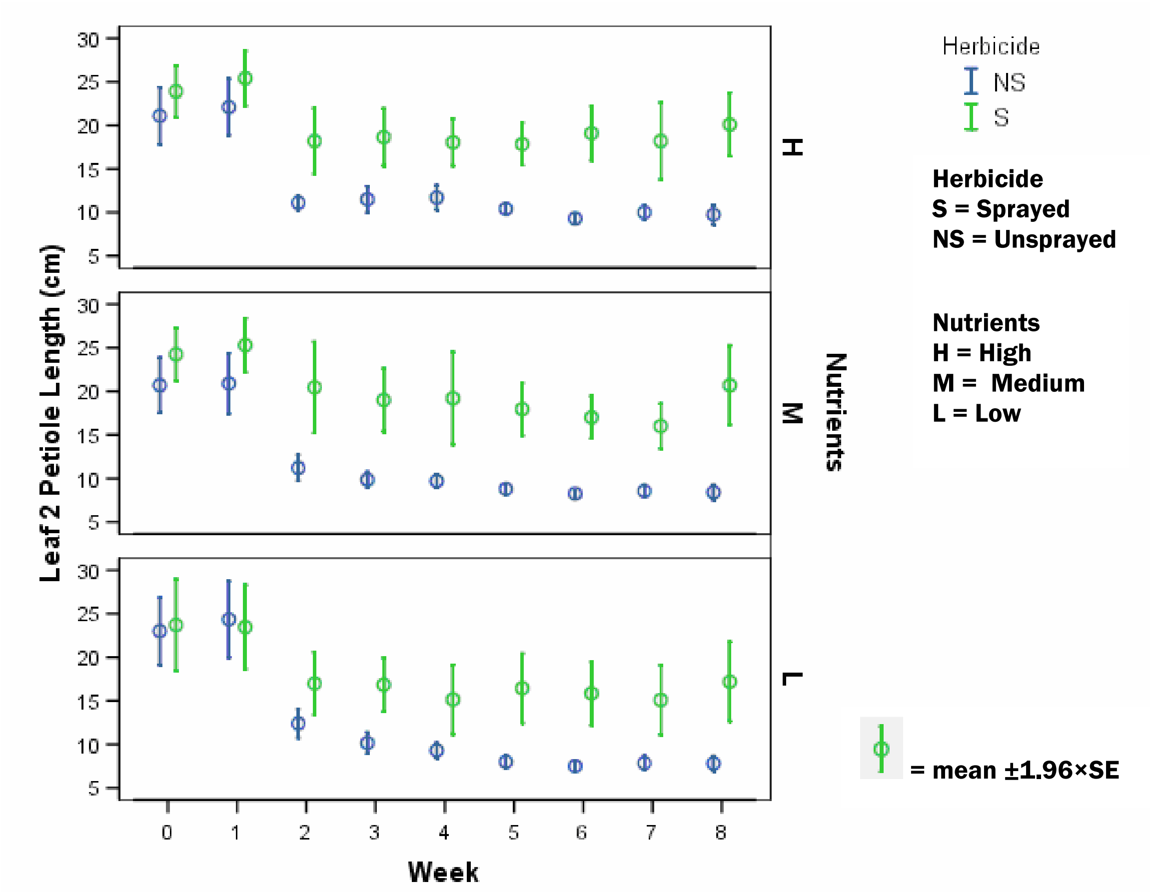

Fig. 5.1: Means plot of second petiole (leaf 2) length, separated by nutrient level and herbicide application.

The response analysed was the length of the second petiole. The plot of these data over time, separated by nutrient level and herbicide application, appears in Fig. 5.1. The petiole length of a healthy plant is expected change from week to week as new leaves are formed. The plots show that the second petiole length of the sprayed (S) plants started off being similar to the length of the unsprayed (NS) plants, but over time the unsprayed plants showed a bigger decrease in petiole length compared to the sprayed plants, before reaching a constant value. This same pattern is displayed for all three nutrient levels (H = high, M = medium, L = low), and a comparison of the values at each week, compared between nutrient levels, shows that they appear to be very similar. The large change in leaf 2 petiole length observed for the unsprayed plants is thought to be due to the generation of new leaves, thereby replacing the leaf two position with a smaller leaf. The small decrease in leaf 2 petiole length in the sprayed plants is thought to be due to the shrinking back of the plant due to its compromised state of health (Kirton, 2005).

The analysis which follows on this data set considers two different approaches to modelling the mean response. The first is a typical ANOVA approach to modelling the mean. This is the type of approach which may have been used if a researcher simply wished to obtain p-values for the different fixed effects in the model (an approach I have observed among ecology students). Therefore a simplistic linear model is used. This analysis assesses the effect of an incorrect mean model choice on the conclusions of the study.

Close examination of the mean response over time (Fig. 5.1) shows that a simple linear relationship between the length of the second petiole and time does not adequately represent the observed relationship. At least two turning points are apparent from this figure. Therefore a more appropriate mean structure would include at least a quadratic time term. Additional adjustment parameters could also be included in order to accommodate for the change in process which appears to take place after the first week in the experiment. Analysis of this more complex quadratic mean structure allows investigation of the effect of different variance models when the mean structure used is adequate.

The purpose of carrying out this analysis was (i) to corroborate the findings of the simulation study by determining the appropriateness of the random effects models with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN and with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN, and the no random effects model with TOEP errors; (ii) to determine how covariance structures compared, particularly covariance structures which performed well in the simulation study, the OLS model, and more complex structures; and (iii) to interpret the results of the linear mixed effects model fitting exercise in the context of an ecological study, both under a simplistic model and under a more complex mean model.

5.3 Simplistic Linear Model Fitting and Analysis

5.3.1 Simplistic linear mean models

Models considered in this section are the linear mixed effects models, along with the simplest case of these models: the OLS ordinary regression model. Preliminary fitting included the nutrients effect and the herbicide effect under various covariance structures. This analysis confirmed that nutrients was not a significant predictor (results for two-way interaction models shown for the random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN in Table 5.1).

The fullest parameterisation of the model considered was

$$\begin{aligned}
y_i = {} & \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Nutrients1) + \beta_4(Nutrients2) \\
& + \beta_5(Herbicide)\times(Week) + \beta_6(Nutrients1)\times(Week) + \beta_7(Nutrients2)\times(Week) \\
& + \beta_8(Nutrients1)\times(Herbicide) + \beta_9(Nutrients2)\times(Herbicide) \\
& + \beta_{10}(Nutrients1)\times(Herbicide)\times(Week) + \beta_{11}(Nutrients2)\times(Herbicide)\times(Week) \\
& + b_{1i} + b_{2i}(Week) + \varepsilon_i
\end{aligned}$$

$$\mathbf{b}_i \sim N(\mathbf{0},\Sigma), \boldsymbol{\varepsilon}_i \sim N(\mathbf{0},\boldsymbol{\omega}_i)$$

where y_i is a response from the i^{th} plant, *Herbicide* is a binary variable indicating sprayed or not sprayed (unsprayed = 1), *Nutrients*1 (high = 1) and *Nutrients*2 (low =

1) are binary indicator variables for the nutrient level (medium is coded as *Nutrients1* = 0 and *Nutrients2* = 0), and *Week* is a quantitative variable for the week number, $\boldsymbol{\beta}_i$ are the parameters of the fixed effects, $\mathbf{b}_i$ are the parameters of the random effects, and $\boldsymbol{\varepsilon}_i$ is the error vector of the i^{th} subject. The matrices $\boldsymbol{\Sigma}$ and $\boldsymbol{\omega}_i$ are the random effects and error covariance matrices respectively. Up to two random effects were included in the model, with *Week* used as the variable for the random slope parameter.

Models were considered with all two-way and three-way interactions. The three-way interaction was non-significant for all models (three-way interaction terms had p-values = 0.78 and 0.35 for the random effects model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN), and was therefore removed from the model. Non-significant two-way interaction terms were removed next, one at a time, followed by non-significant main effect terms that were not contained within any significant interaction terms. All terms containing the nutrients effect were found to be non-significant. The confidence intervals of the terms containing nutrients fluctuated the most between models with different mean structures. Therefore the nutrients effect was excluded from the final model, which contained fixed effects for herbicide and week. These models were fit using SAS PROC MIXED (ver. 9.1).

Table 5.1: Fixed effects results for simplistic linear model with the nutrients and herbicide effects, including all two-way effects, under independent error covariance and random intercept and slope with unstructured covariance.

Effect	**Estimate**	**p-value**	**Lower 95% CI Limit**	**Upper 95% CI Limit**
Model with all two-way interactions				
Intercept	22.9142	<0.0001	20.3152	25.5133
Week	-0.7431	0.0001	-1.1082	-0.3779
Herbicide	-4.4785	0.0084	-7.7713	-1.1856
Nutrient1	-0.7180	0.6826	-4.2057	2.7698
Nutrient2	-0.9347	0.5947	-4.4225	2.5530
Herbicide×Week	-0.9014	<0.0001	-1.2665	-0.5363
Nutrient1×Week	0.1754	0.4353	-0.2718	0.6226
Nutrient2×Week	-0.2729	0.2266	-0.7201	0.1743
Nutrient1×Herbicide	1.1659	0.5690	-2.9129	5.2447
Nutrient2×Herbicide	2.3750	0.2482	-1.7037	6.4538
Model with two-way interactions including week				
Intercept	22.2341	<0.0001	19.9814	24.6668
Week	-0.7431	0.0001	-1.1082	-0.3779
Herbicide	-3.2981	0.0066	-5.6408	-0.9555
Nutrient1	-0.1350	0.9252	-3.0042	2.7342
Nutrient2	0.2528	0.8605	-2.6164	3.1220
Herbicide×Week	-0.9014	<0.0001	-1.2665	-0.5362
Nutrient1×Week	0.1754	0.4353	-0.2718	0.6226
Nutrient2×Week	-0.2729	0.2266	-0.7201	0.1743
Model with all main effects and the interaction between week and herbicide effect				
Intercept	22.4717	<0.0001	20.4417	24.5018
Week	-0.7756	<0.0001	-1.0382	-0.5129
Herbicide	-3.2981	0.0071	-5.6624	-0.9339
Nutrient1	0.6620	0.5153	-1.3636	2.6877
Nutrient2	-0.9873	0.3331	-3.0129	1.0384
Herbicide×Week	-0.9014	<0.0001	-1.2728	-0.5300

5.3.2 Model fits under various covariance structures

The proposed model for the data was:

$$y_i = \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Herbicide) \times (Week) + b_{1i} + b_{2i}(Week) + \varepsilon_i$$
$$\mathbf{b}_i \sim N(\mathbf{0}, \Sigma), \boldsymbol{\varepsilon}_i \sim N(\mathbf{0}, \boldsymbol{\omega}_i) \qquad \ldots(1)$$

where y_i is a response from the i^{th} plant, *Herbicide* is a binary variable indicating sprayed or not sprayed (unsprayed = 1), *Week* is a quantitative variable for the week number, $\boldsymbol{\beta}_i$ are the parameters of the fixed effects, $\mathbf{b}_i$ are the parameters of the random effects, and $\boldsymbol{\varepsilon}_i$ is the error vector of the i^{th} subject. The matrices $\boldsymbol{\Sigma}$ and $\boldsymbol{\omega}_i$ are the random effects and error covariance matrices respectively.

All of the parameter estimates, including the covariance parameter estimates, and information criteria for each model successfully fitted to the data appear in Appendix B1. Due to the large quantity, the outlier and influence diagnostics for selected models, have been included in Appendix B2. A summary of model fit results appears in Table 5.2 and are discussed below. The parameter estimates appear in the order presented in equation (1). This table also shows which models were not fitted to the data and the reason for failure to fit these models.

As opposed to the simulation study, certain models with UN or heterogeneous error covariance structures performed well. Models fitted with UN and heterogeneous error covariance structures obtained the lowest AIC, BIC and AICc values. The most likely explanation is that the study was balanced and that there were more subjects per treatment compared to the simulation study data set, accommodating more parameters in the model. Models which performed particularly badly, obtaining invalid estimates or resulting in non-convergence, were random intercept and slope models fitted with an AR(1) or ARH(1) error structure. None of the models in this category obtained valid estimates. As observed in the simulation study, when covariance structures were fitted to the random effects which forced the variance of the random intercept and slope to be the same, such as CS or TOEP structure, the resulting covariance structure was not positive definite. This was demonstrated through the calculation of negative eigenvalues for these estimated covariance matrices.

Table 5.2: Summary of model fitting results.

Successfully fitted models						
$\boldsymbol{\omega}_i$	**Σ**	$\boldsymbol{B}_i$	**Standard Error**	**AIC**	**BIC**	**AICc**
UN	None	$\begin{pmatrix} 19.50 \\ -4.54 \\ -0.21 \\ -0.65 \end{pmatrix}$	$\begin{pmatrix} 0.66 \\ 0.93 \\ 0.09 \\ 0.13 \end{pmatrix}$	2929.0	3023.2	2937.4
CSH	CSH/ARH(1)	$\begin{pmatrix} 17.87 \\ -6.32 \\ -0.28 \\ -0.50 \end{pmatrix}$	$\begin{pmatrix} 0.85 \\ 1.20 \\ 0.11 \\ 0.16 \end{pmatrix}$	3092.7	3120.0	3093.4
CSH	UN	$\begin{pmatrix} 17.87 \\ -6.32 \\ -0.28 \\ -0.50 \end{pmatrix}$	$\begin{pmatrix} 0.85 \\ 1.20 \\ 0.11 \\ 0.16 \end{pmatrix}$	3092.7	3120.0	3093.4
ARH(1)	None	$\begin{pmatrix} 21.15 \\ -4.20 \\ -0.60 \\ -0.80 \end{pmatrix}$	$\begin{pmatrix} 0.93 \\ 1.32 \\ 0.17 \\ 0.25 \end{pmatrix}$	3097.6	3118.6	3098.0
CSH	None	$\begin{pmatrix} 17.14 \\ -6.65 \\ -0.31 \\ -0.54 \end{pmatrix}$	$\begin{pmatrix} 0.83 \\ 1.18 \\ 0.10 \\ 0.15 \end{pmatrix}$	3124.2	3145.2	3124.7
TOEP	None	$\begin{pmatrix} 22.72 \\ -2.93 \\ -0.81 \\ -1.01 \end{pmatrix}$	$\begin{pmatrix} 0.82 \\ 1.15 \\ 0.14 \\ 0.20 \end{pmatrix}$	3173.1	3192.0	3173.5

Table 5.2 (cont): Summary of model fitting results

Successfully fitted models						
ω_i	**Σ**	**B_i**	**Standard Error**	**AIC**	**BIC**	**AICc**
AR(1)	Intercept only	$\begin{pmatrix} 22.43 \\ -3.02 \\ -0.70 \\ -0.96 \end{pmatrix}$	$\begin{pmatrix} 0.84 \\ 1.19 \\ 0.15 \\ 0.21 \end{pmatrix}$	3177.4	3183.7	3177.4
AR(1)	None	$\begin{pmatrix} 22.47 \\ -2.94 \\ -0.68 \\ -0.97 \end{pmatrix}$	$\begin{pmatrix} 0.85 \\ 1.21 \\ 0.17 \\ 0.23 \end{pmatrix}$	3180.4	3184.6	3180.4
VC	CSH/ARH(1)	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.81 \\ 1.15 \\ 0.13 \\ 0.19 \end{pmatrix}$	3243.8	3252.1	3243.8
VC	UN	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.81 \\ 1.15 \\ 0.13 \\ 0.19 \end{pmatrix}$	3243.8	3252.1	3243.8
VC	VC	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.70 \\ 1.00 \\ 0.11 \\ 0.16 \end{pmatrix}$	3247.4	3253.7	3247.5
CS	None	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.72 \\ 1.02 \\ 0.11 \\ 0.15 \end{pmatrix}$	3247.6	3251.8	3247.7
VC	Intercept only	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.72 \\ 1.02 \\ 0.11 \\ 0.15 \end{pmatrix}$	3247.7	3251.8	3247.7
VC	None	$\begin{pmatrix} 22.36 \\ -3.30 \\ -0.78 \\ -0.90 \end{pmatrix}$	$\begin{pmatrix} 0.59 \\ 0.84 \\ 0.13 \\ 0.17 \end{pmatrix}$	3334.9	3339.2	3334.9

Table 5.2 (cont): Summary of model fitting results

Models which resulted in estimates of ∑ that were not positive definite					
$\boldsymbol{\omega}_i$	**Σ**				
VC	CS				
CS	TOEP				
AR(1)	VC				
AR(1)	UN				
ARH(1)	VC				
ARH(1)	CSH/ARH(1)				
ARH(1)	UN				
TOEP	CS				
TOEP	TOEP				
Models which had a final Hessian matrix which was not positive definite					
$\boldsymbol{\omega}_i$	**Σ**				
CS	CSH				
Models which did not converge					
$\boldsymbol{\omega}_i$	**Σ**				
AR(1)	CSH/ARH(1)				

The complexity of models ranges. The simplest model, the OLS model, has four fixed effects parameters and one variance parameter, totalling five parameters, and resulting in an observation to parameter ratio of 108:1. The no random effects models with a maximum of 45 covariance parameters, which totals to 49 parameters in total, results in an observation to parameter ratio of 11.0:1. The random intercept models have an additional sixty random effects parameters that need to be calculated and one random effects variance, resulting in a maximum of 110 model parameters, and an observation to parameter ratio of 4.9:1. The random intercept and slope models have 120 random effects parameters that need to be estimated and up to three random effects covariance parameters, resulting in a maximum of 172 parameters, and leading to an observation to parameter ratio of 3.1:1. The lowest observation to parameter ratio of the ecological data set is twice as large compared to the lowest observation to parameter ratio of the

PR data set. Models successfully fitted to the data had varied observation to parameter ratios, ranging between 3.1:1 and 108:1.

Models which were expected to perform well from the results of the simulation study include the no random effects model with TOEP error structure, and the random intercept and slope models with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN, and with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN. These models obtained information criteria that were close to the minimum relatively consistently, and obtained the best coverage probabilities for the fixed effects estimates. Of these three models, only two were successfully fitted to the data, with the random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN obtaining an estimated variance for the slope of zero in the random effects covariance matrix. Models which obtained low mean values for the information criteria in the simulation study, namely the no random effects model with $\boldsymbol{\omega}_i$ = CS model, the random intercept model with independent errors, and the random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = VC, fitted the data successfully, but obtained relatively high AIC and BIC values. The random intercept model with $\boldsymbol{\omega}_i$ = AR(1), which performed relatively well in the simulation study with respect to the information criteria analysis and obtained reasonable coverage probabilities, was fitted successfully and obtained relatively low values for AIC, BIC and AICc. The no random effects model with $\boldsymbol{\omega}_i$ = AR(1), which obtained coverage probabilities that were too high in the simulation study, had a similar fit to the random intercept model and obtained similar fixed effects parameter estimates. The OLS model, which also had coverage probabilities that were too high in the simulation study, obtained the worst fit.

5.3.3 Goodness-of-fit analysis of the linear model under various covariance structures

A closer examination will now be carried out on the fit of individual models.

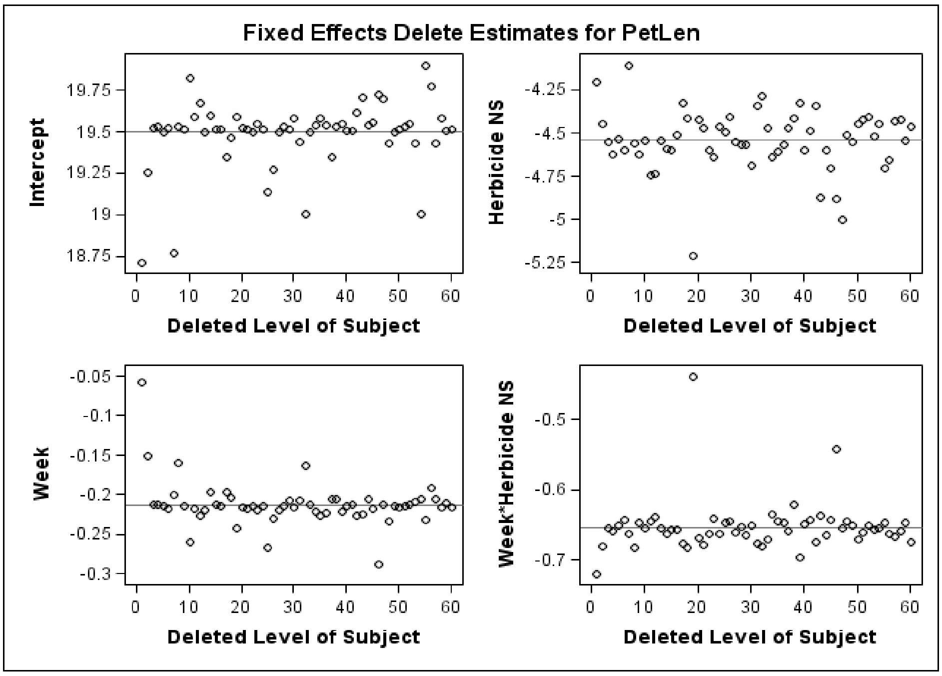

Fig. 5.2: Estimates of fixed parameters obtained after the removal of each data point for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = UN.

5.3.3.1 No random effects model with $\boldsymbol{\omega}_i$ = UN

The no random effects model with $\boldsymbol{\omega}_i$ = UN is the best fitting model according to the information criteria, as this model obtained the minimum AIC of 2929, a minimum BIC of 3023 and a minimum AICc of 2937 (Table 5.2). Table 5.2 shows that the estimated fixed effects were all significant at the 5% level. Fig. 5.2 is a plot of the fixed effects parameter estimates after the removal of each subject's observations in the study, where the x-axis indicates which subject has been removed. The horizontal line in each plot indicates the fixed effect parameter estimate obtained for the model fitted to the full data set. The purpose of these plots is to demonstrate the effect an individual subject's observations have on the estimate of the fixed effects parameters. This plot shows that the estimates for the intercept coefficients of both sprayed (figure with y-axis labelled "Intercept") and unsprayed plants (figure with y-axis labelled "Herbicide NS") are quite stable, with a maximum difference of less than one unit from the actual estimate for subject 1. The estimates for the slopes also show relatively little change (of difference of less than 0.1) for deletion of most subjects, except in the case of subject 1, for the estimate of the slope of the sprayed plants, and subject 19, for the estimate of the difference in slopes between the unsprayed and sprayed plants, which result in differences between the deleted estimate and the actual estimate of close to 0.2. Both subjects 1 and 19 were in the sprayed treatment.

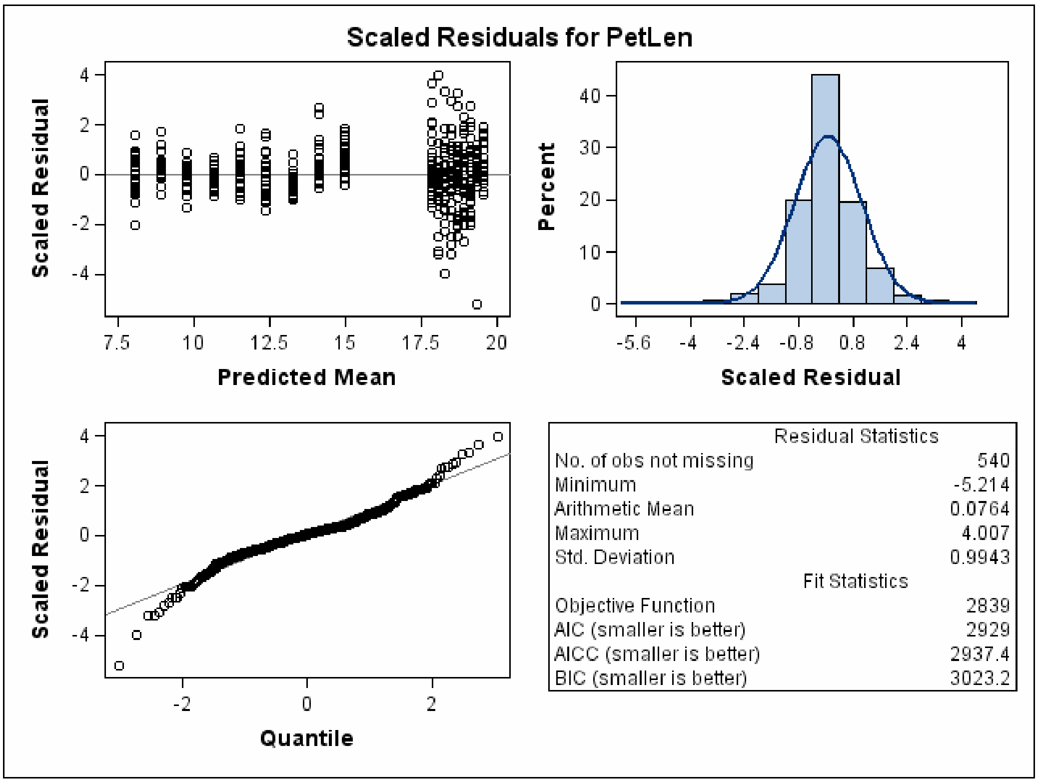

Fig. 5.3: Plots of scaled residuals for the no random effects simplistic linear model with ω_i = UN.

Fig. 5.3 contains plots of the transformed residuals. The transformation of these residuals is discussed in Chapter two. The transformed residuals can be treated in the same way as residuals from an ordinary regression analysis. These residuals should be approximately normally distributed and have a mean of zero and a variance of one. The first plot is of the predicted values against the transformed residuals. This plot shows that the transformed residuals vary between 4 and -4. Since the time variable is discrete, the separate bands of residuals that can be seen are the residuals for predicted values at a particular week. The bands of residuals on the left of the plot are for the predicted values of the unsprayed plants, and those on the right are for the predicted values of the sprayed plants. The spread of the residuals for the unsprayed plants appear to deviate around the zero line, except those in the last band of residuals. The slope of the predicted equation for unsprayed plants is negative, but the plot of the observed data (Fig. 5.1) reveals that although there is a decreasing trend in the data, the mean of the response for week 8 is increasing relative to week 7, and the mean for week 1 is smaller compared to the mean for week 2. Since the equation predicted is linear, this information cannot be captured by the model, and therefore the estimates for week 8 will be lower compared to the observed data, and the estimates for week 1 will be higher compared to the observed data. The variability in the residuals of the sprayed plants is much higher compared to those of the unsprayed plants. Therefore it does not seem that the variability within the sprayed and unsprayed plants is the same. This is also suggested by the plot of the observed data (Fig. 5.1). The next two plots in Fig. 5.3 are a histogram and a q-q plot of the transformed residuals, and are intended to show the normality of the transformed residuals. The histogram of the residuals is close to symmetrical, but with slight skewness to the right. The q-q plot shows that most of the points fall on the one-to-one line, but that there is deviation from normality in the tails of the distribution of the transformed residuals.

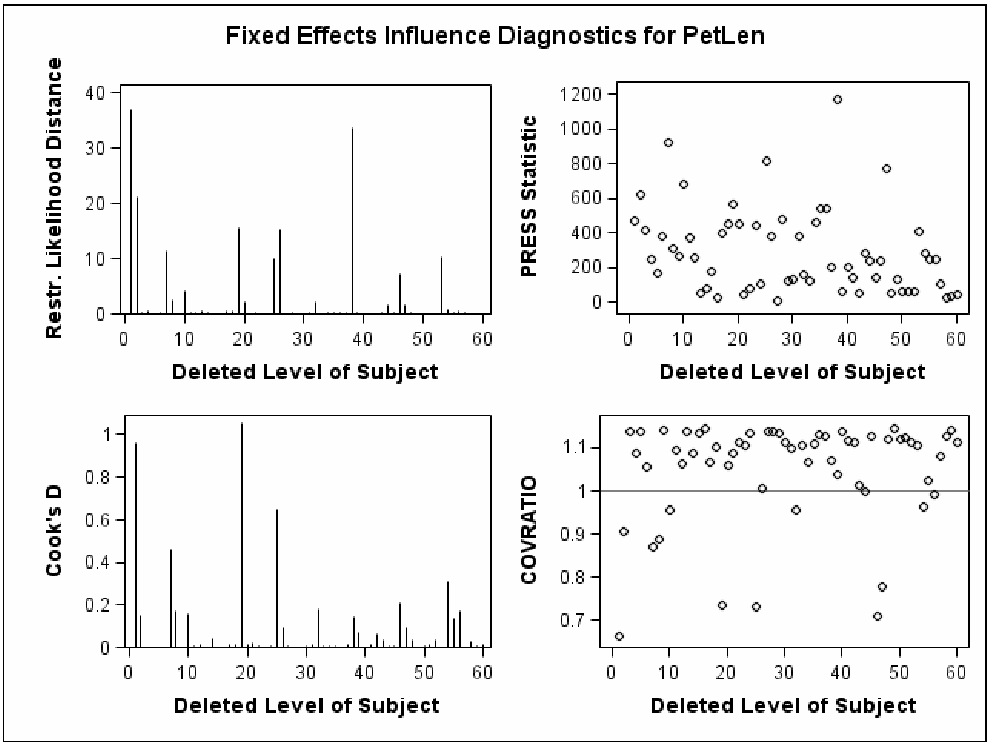

Fig. 5.4: Plots of the influence diagnostics for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = UN.

Fig 5.4 contains plots of four different influence diagnostics. These diagnostics are the restricted likelihood distance, the PRESS statistic, Cook's distance (Cook's D), and covariance ratio (COVRATIO). Each of these diagnostics examines a different aspect of the influence an observation has on the estimated model. The restricted likelihood distance measures the overall influence of an observation, Cook's distance measures the influence of an observation on all predicted values, the covariance ratio measures the influence of an observation on the precision of the estimates, and the PRESS residuals measure the influence of an observation on its own predicted value. The deletion takes place at the level of a subject, and so these diagnostics show the influence that a plant's observations have on the estimated model. The full set of diagnostic values for each subject is available in Appendix B2. Influential subjects are identified as those subjects with absolute diagnostic values relatively larger compared to other subjects. Subjects 1, 19, and 38 were highlighted as being potentially influential, as the diagnostic values for these subjects are higher relative to all other subjects. In particular, the values for Cook's distance and the restricted likelihood distance highlight these three subjects, and so these subjects influence the predicted values of all other subjects. These three subjects belonged to the sprayed group of plants, so therefore the model may be predicting the unsprayed plants better than the sprayed plants.

5.3.3.2 Random intercept and slope model with $\boldsymbol{\omega}_i$ = CSH and $\sum$ = CSH

The random intercept and slope model with $\boldsymbol{\omega}_i$ = CSH and $\sum$ = CSH performed only slightly worse. This model obtained an AIC of 3093, a BIC of 3120 and an AICc of 3093 (Table 5.2). Again all of the fixed parameters were significant. The parameter estimates for the intercepts are about two units lower compared to those estimated for the previous model, but the slope estimates are very similar. The stability of the fixed effects parameter estimates with removal of subjects from the data set was similar compared to the previous model, but showed smaller extreme deviations in estimates for the slopes (Fig. 5.5). The transformed residuals varied between 4 and -4 (Fig. 5.6). The plot of the predicted values against the residuals show the same pattern as for the previous model, but show a smaller range in the residuals for the unsprayed plants, but with residuals of large predicted value suggesting negative bias in the predicted values. The q-q plot shows that more points fall on the one-to-one line compared to the previous model, but the histogram is slightly skewed to the left. The influence diagnostics only highlight subject 38 as an influential point (Fig. 5.7). The size of the diagnostics is smaller in comparison to the previous model.

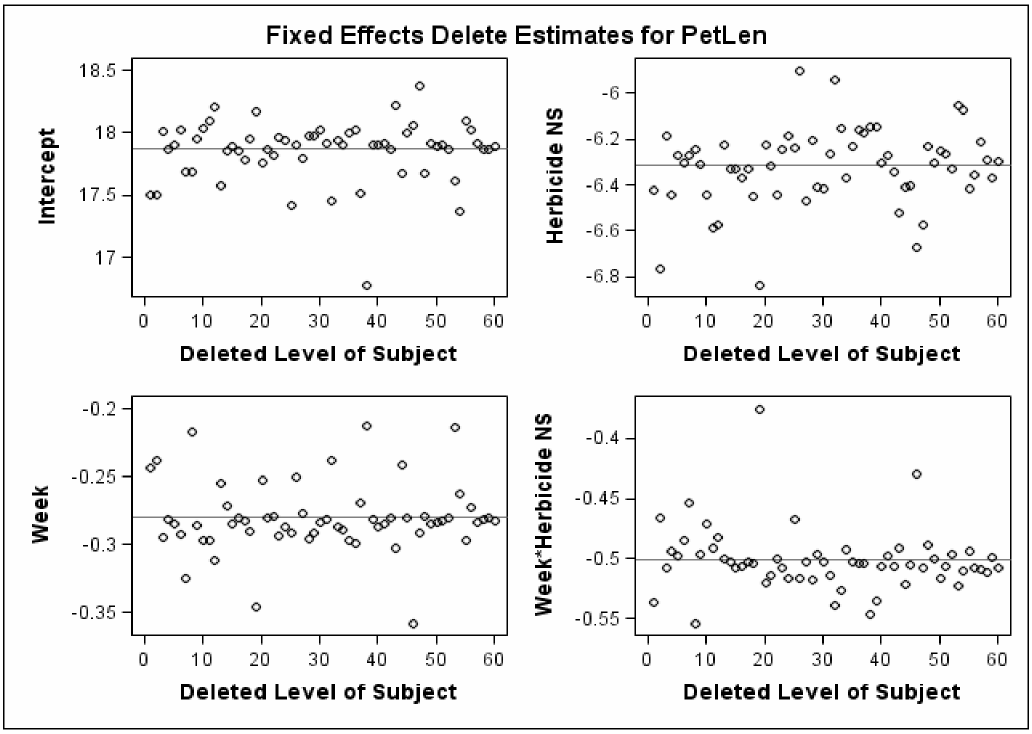

Fig. 5.5: Estimates of fixed parameters obtained after the removal of each data point for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH.

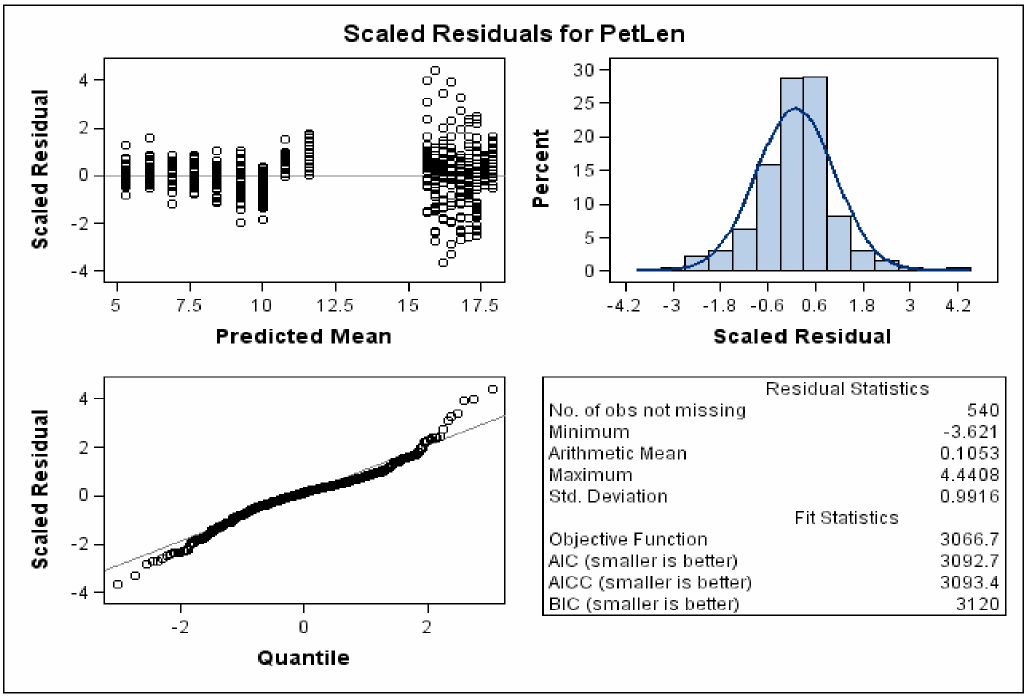

<u>Fig. 5.6: Plots of scaled residuals for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH.</u>

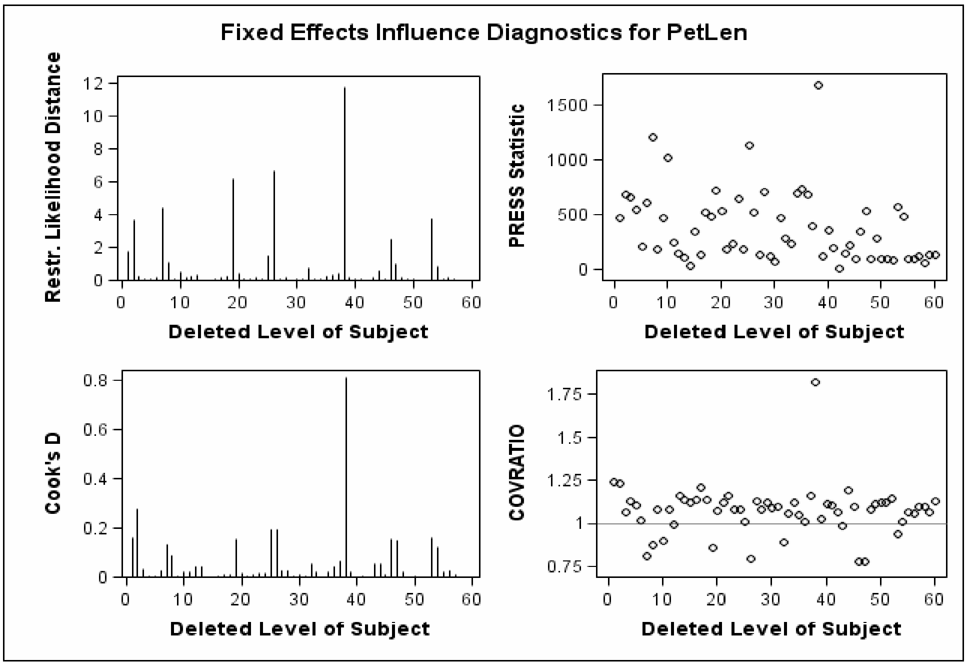

<u>Fig. 5.7: Plots of the influence diagnostics for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH.</u>

5.3.3.3 Random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN

The random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN was shown to be a relatively robust model in the previous simulation study. This model obtained an AIC value of 3244, a BIC value of 3252 and an AICc value of 3244 (Table 5.2). The estimated parameters of the fixed effects were all significant at the 5% level and showed only small deviations, relative to those of the best fitting model, when subjects were removed (Fig. 5.8). The estimates of the fixed effects intercept parameters were approximately two units higher compared to those obtained for the best fitting model, and had steeper downward slopes. The transformed residuals varied between 4 and -3 and were close to normality, but some pattern was apparent in the plot of the predicted values against the transformed residuals for those residuals of the unsprayed plants and was more exaggerated compared to previous models (Fig. 5.9). Along with subject 38, the influence diagnostics highlight at least five other subjects as potential outliers, including observations 2, 7, 25, and 47 (Fig. 5.10). The restricted likelihood distances are much smaller compared to all three previous models.

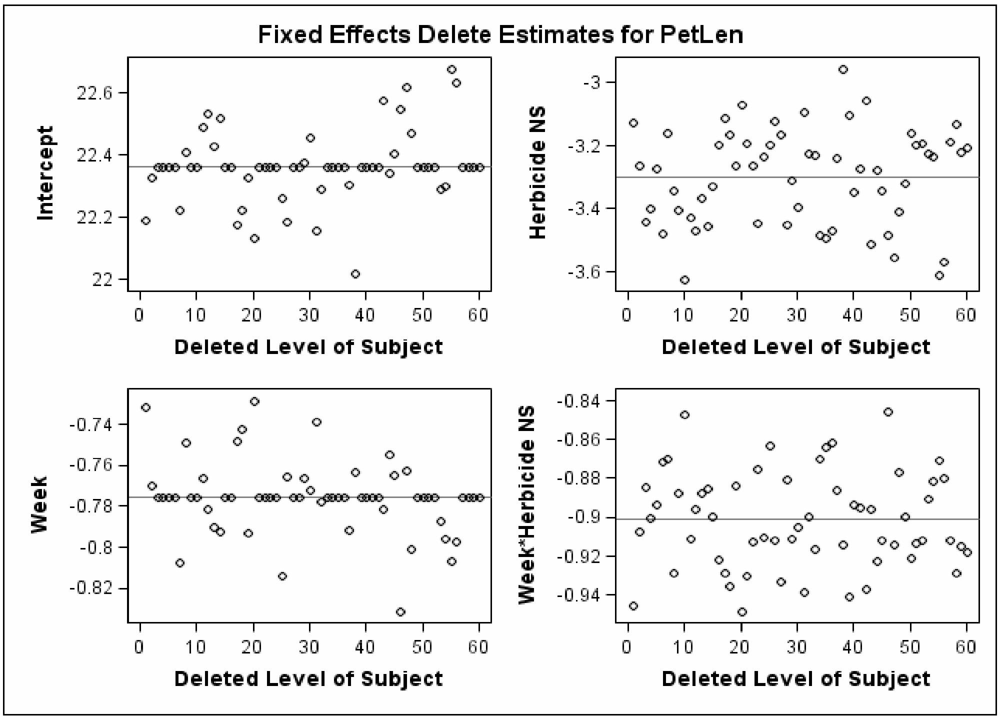

Fig. 5.8: Estimates of fixed parameters obtained after the removal of each data point for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

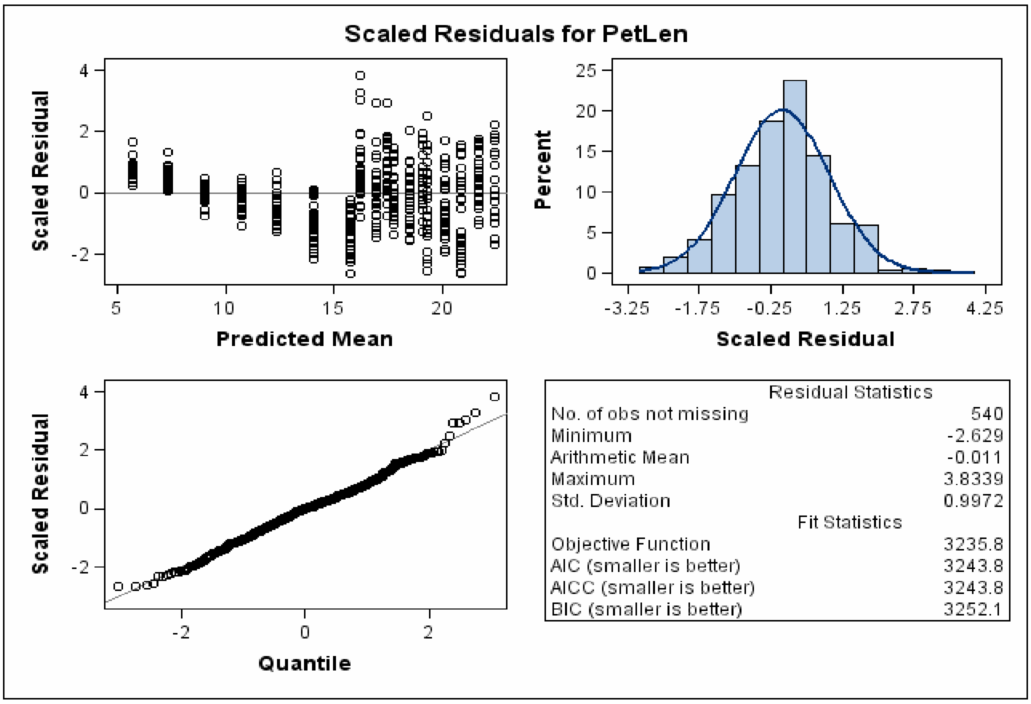

Fig. 5.9: Plots of scaled residuals for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

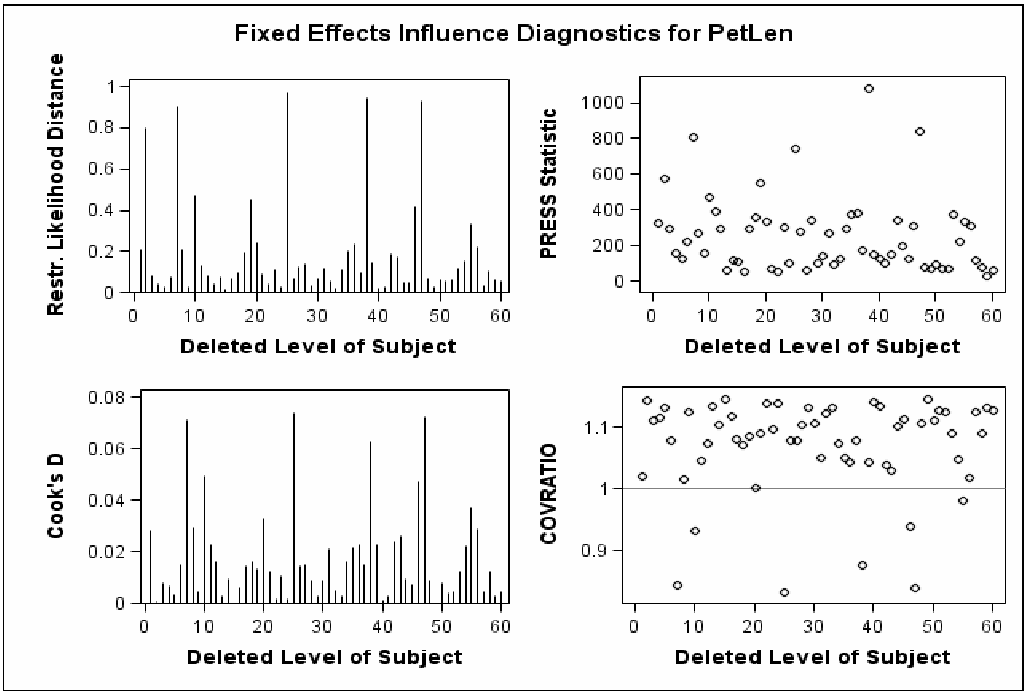

Fig. 5.10: Plots of the influence diagnostics for the random intercept and slope simplistic linear model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

5.3.3.4 Random intercept model with $\boldsymbol{\omega}_i$ = AR(1)

The random intercept model with $\boldsymbol{\omega}_i$ = AR(1) was another model shown to be relatively robust by the simulation study. This model obtained AIC, BIC and AICc values of 3177, 3184 and 3177 respectively. These values were more or less in the middle of the range of values obtained for the information criteria by models in this study. The estimates of the fixed effects (Table 5.2) were very close to those obtained in the previous random effects model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN. The deleted estimates for the fixed effects show that the estimated fixed effects estimates showed relatively small deviation in estimated values when subjects were deleted compared to the best fitting model (Fig. 5.11). The transformed residuals varied between 4 and -3 (Fig. 5.12). The plot of the predicted values against the residuals again shows a

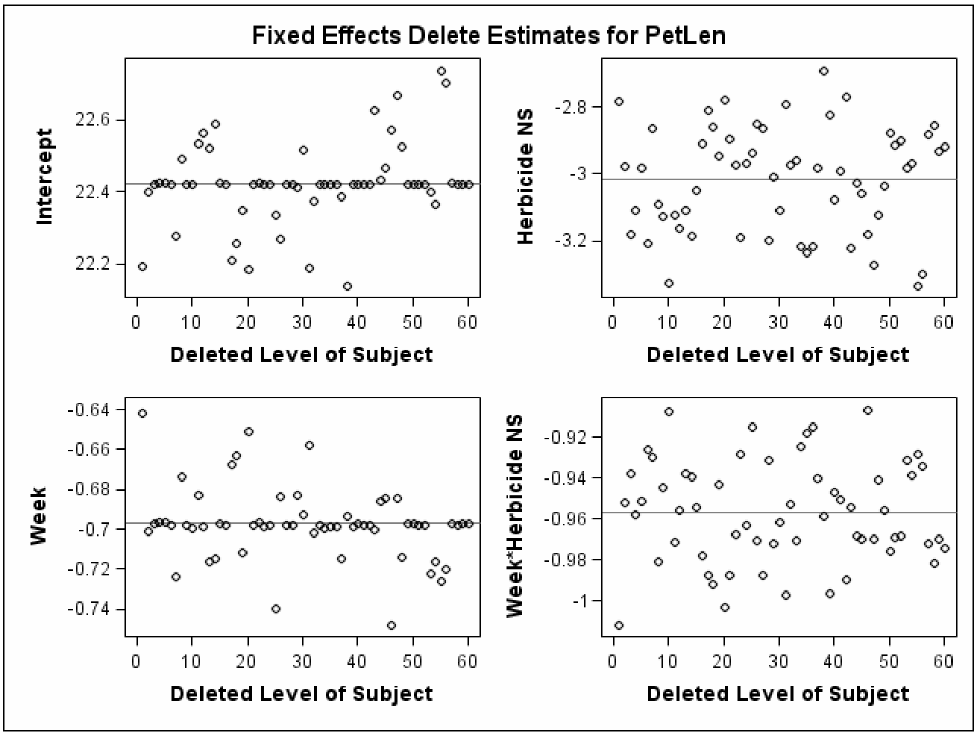

Fig. 5.11: Estimates of fixed parameters obtained after the removal of each data point for the random intercept simplistic linear model with $\boldsymbol{\omega}_i$ = AR(1).

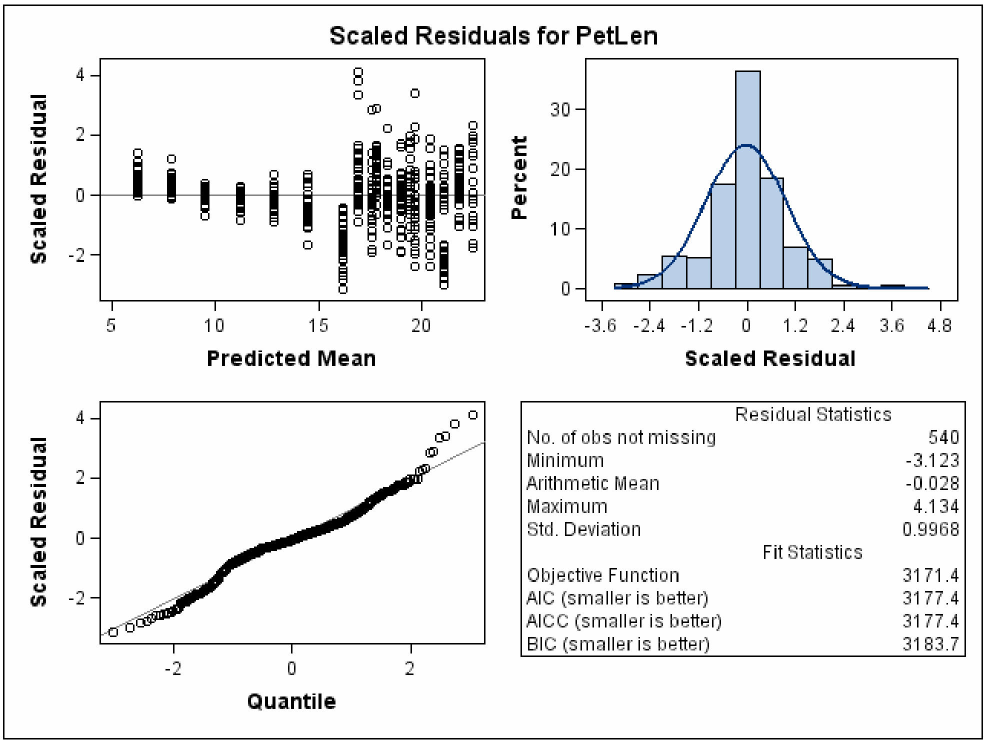

Fig. 5.12: Plots of scaled residuals for the random intercept simplistic linear model with $\boldsymbol{\omega}_i$ = AR(1).

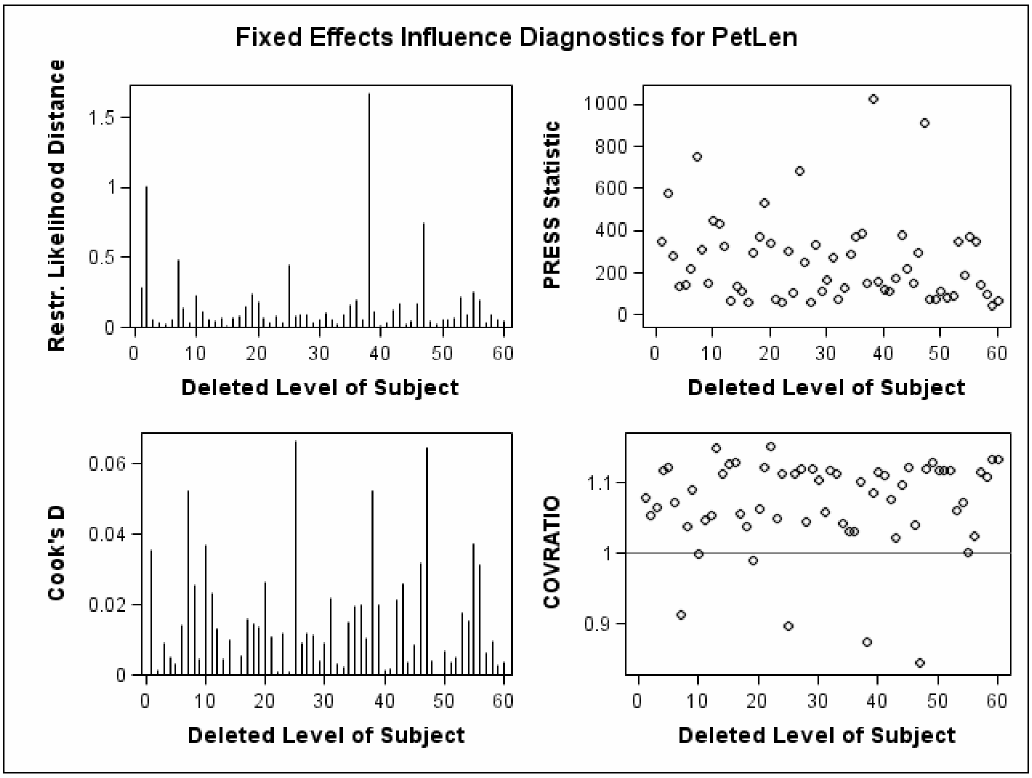

Fig. 5.13: Plots of the influence diagnostics for the random intercept simplistic linear model with $\boldsymbol{\omega}_i$ = AR(1).

pattern for the unsprayed plants, where small predicted values tend to be negatively biased and large predicted values tend to be positively biased, except in the case of the largest predicted values, where the bias changes to negative again. The influence and outlier diagnostics show that observations 2 and 38 are potentially influential points (Fig. 5.13). The size of the restricted likelihood distances and Cook's distances is much smaller compared to those obtained for the best fitting model.

5.3.3.5 No random effects model with $\boldsymbol{\omega}_i$ = TOEP

The no random effects model with $\boldsymbol{\omega}_i$ = TOEP, predicted to be a robust model from the simulation study, resulted in an AIC, BIC and AICc values of 3174, 3192 and 3174 respectively (Table 5.2). The fixed parameter estimates were all significant at the 5% level and did not show large deviations in size when subjects were excluded (Fig. 5.14). Comparing the parameter estimates to the best fitting model, the no random effects model with $\boldsymbol{\omega}_i$ = TOEP obtained similar estimates compared to the random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN and the random intercept model with $\boldsymbol{\omega}_i$ = AR(1) (Table 5.2). The transformed residuals of this model were between 4 and -3, and only deviated from normality for large values of scaled residuals (Fig. 5.15). The plot of the predicted values against the transformed residuals is the best compared to all previous models. This plot shows that the bands of residuals vary around the zero line, and that no trend in the residuals is evident indicating that the predicted values are not biased. But as in the case of the previous models, the homogeneity of the residual variance is poor. Subject 38 was once again highlighted as an influential subject, along with subjects 2, 25, 47 and 54 (Fig. 5.16). The restricted likelihood distances and Cook's distances were slightly larger compared to the previous model, but not as large as those of the no random effects model with $\boldsymbol{\omega}_i$ = UN.

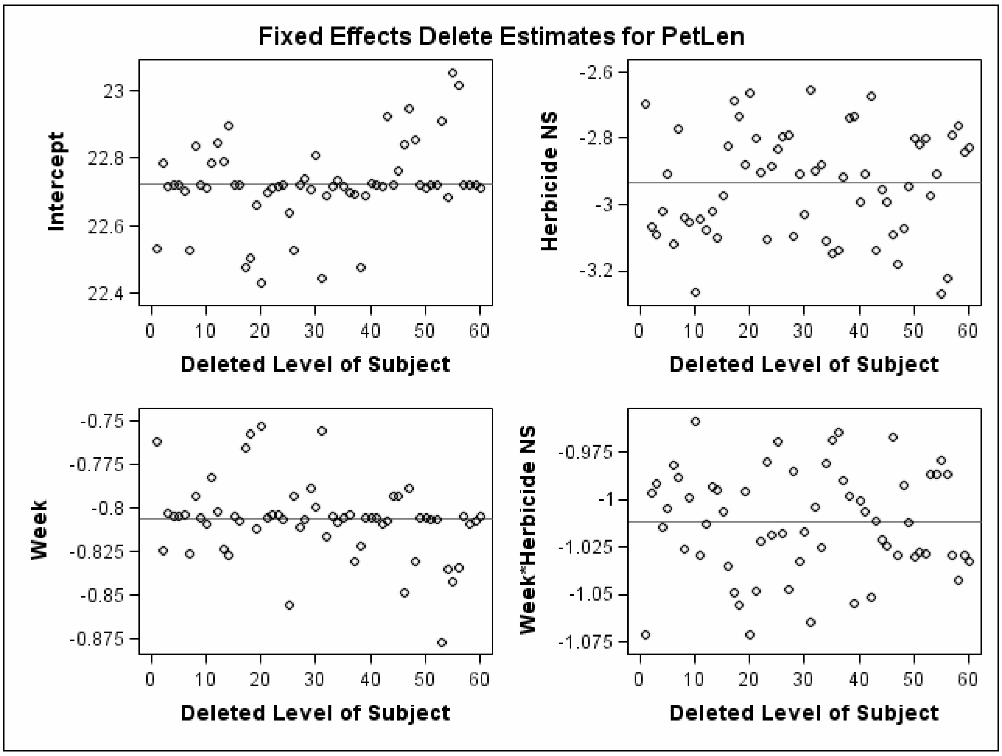

<u>Fig. 5.14: Estimates of fixed parameters obtained after the removal of each data point for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = TOEP.</u>

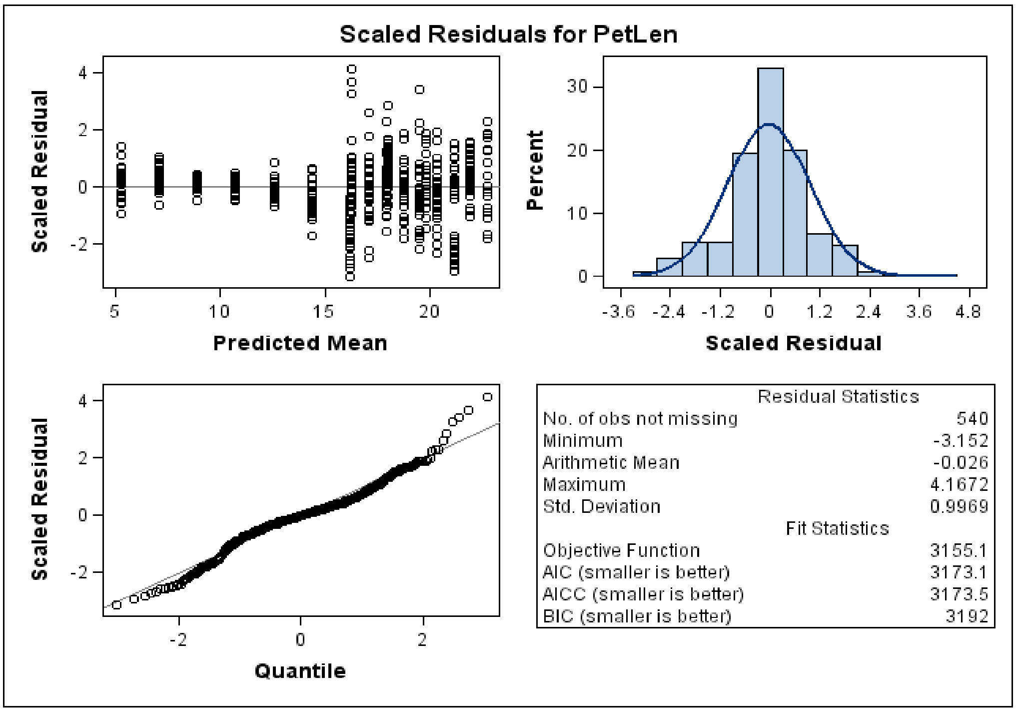

<u>Fig. 5.15: Plots of scaled residuals for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = TOEP.</u>

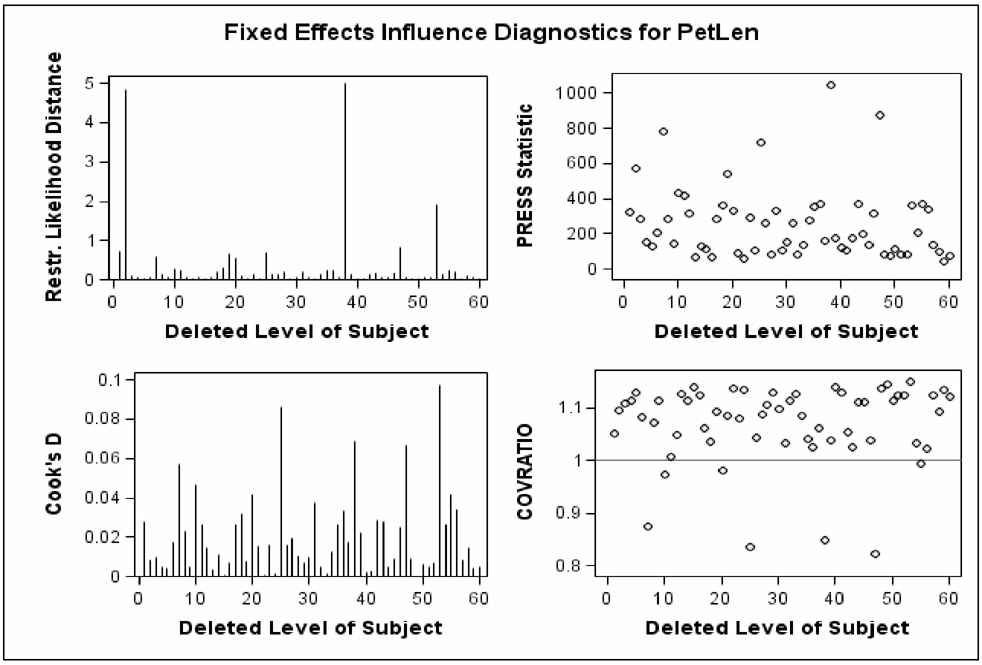

Fig. 5.16: Plots of the influence diagnostics for the no random effects model simplistic linear with $\boldsymbol{\omega}_i$ = TOEP.

5.3.3.6 OLS model

The final model investigated was the OLS model with independent error covariance structure. This model had the highest AIC, BIC and AICc values of 3335, 3339 and 3335 respectively (Table 5.2). The fixed parameter estimates were very similar to those for the previous model. These estimates were all significant at the 5% significance level and did not show high sensitivity to the removal of subjects (Fig. 5.17). The transformed residuals varied between 4 and -3, and when plotted against the predicted values showed a trend in the bias of the predictions. The plots for normality showed only small departures from normality, but skewness to the right (Fig. 5.18). The influence diagnostics highlighted subjects 7, 25, 38 and 47. The Cook's distances were larger compared to previous models, but other diagnostics were similar in size (Fig. 5.19).

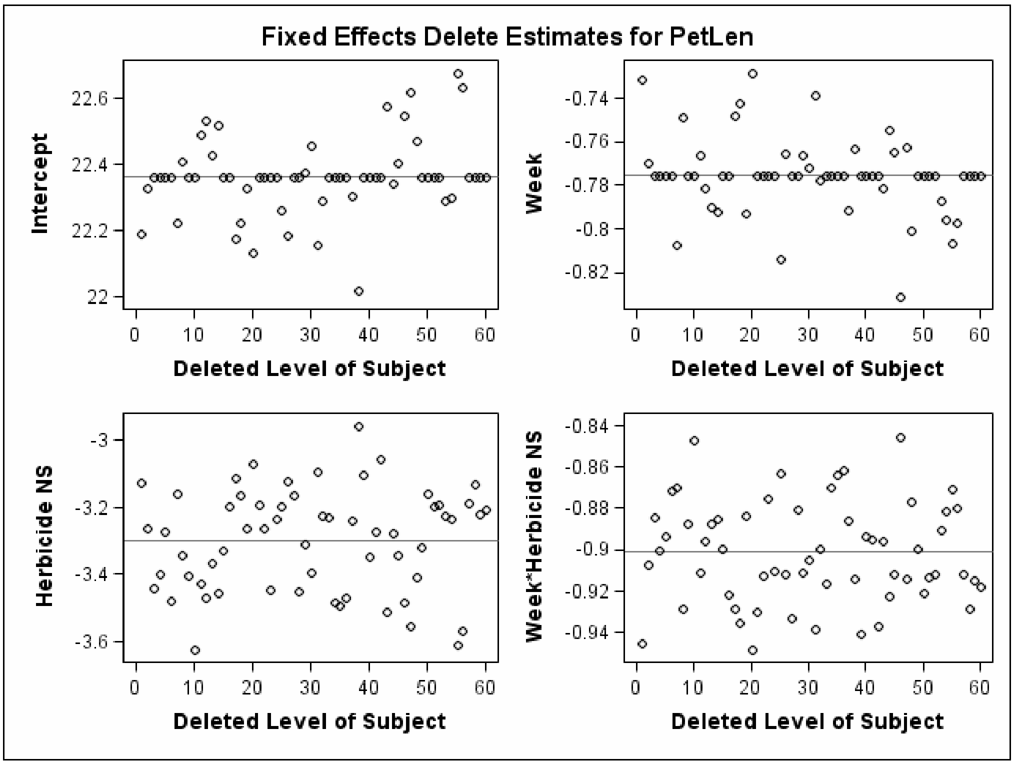

Fig. 5.17: Estimates of fixed parameters obtained after the removal of each data point for the simplistic linear OLS model with independent error covariance structure.

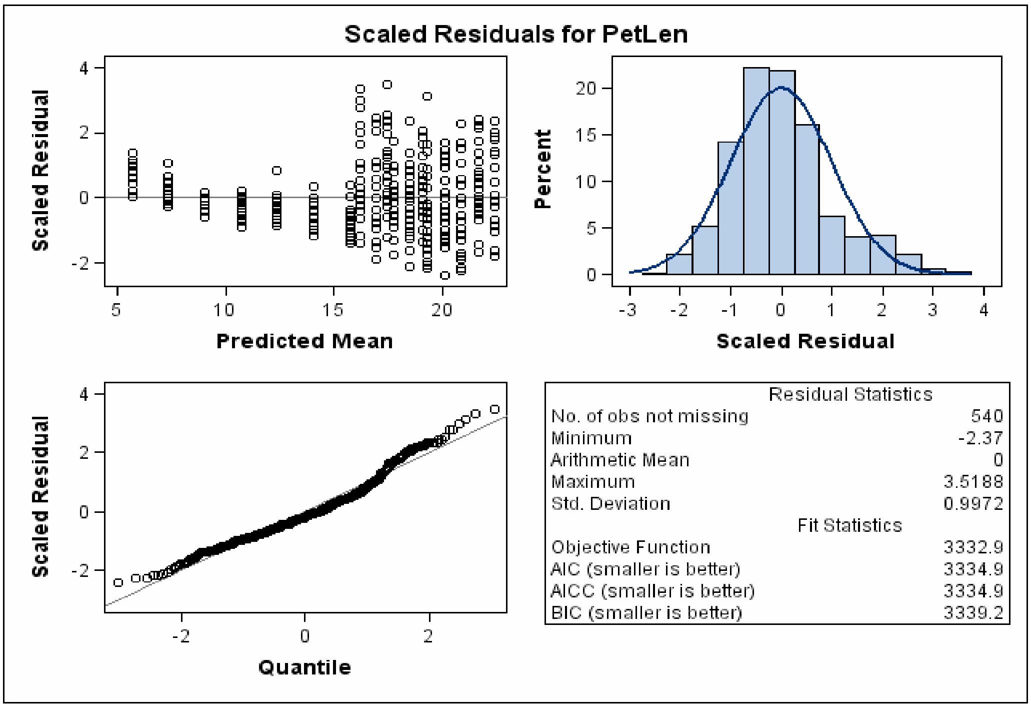

Fig. 5.18: Plots of scaled residuals for the simplistic linear OLS model with independent error covariance structure.

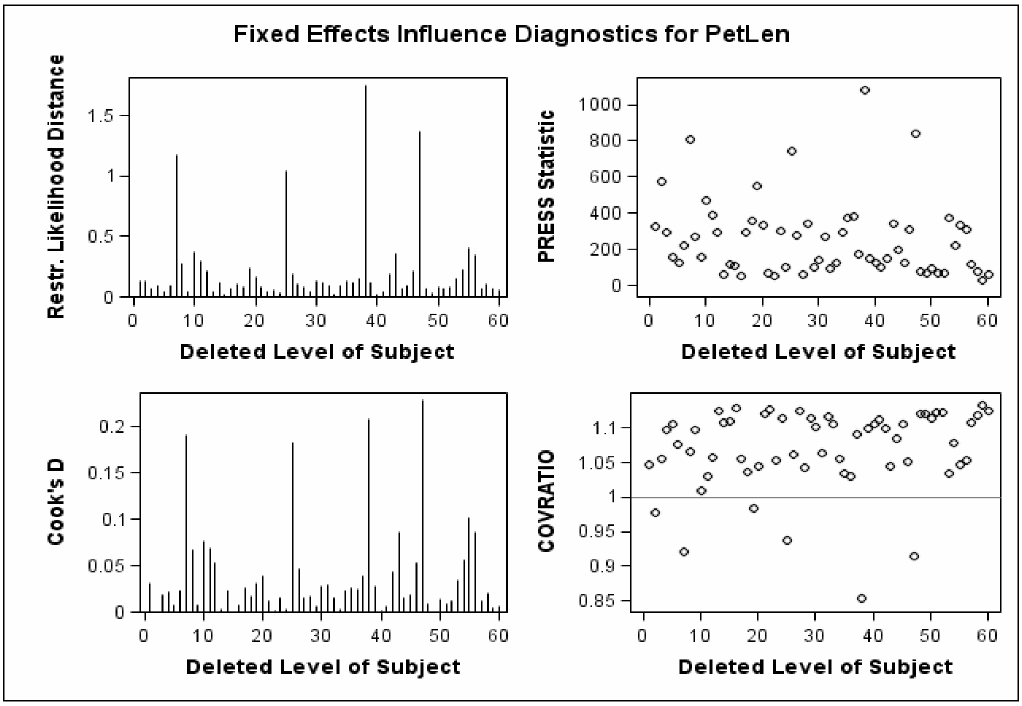

Fig. 5.19: Plots of the influence diagnostics for the simplistic linear OLS model with independent error covariance structure.

5.3.3.7 Summary

The model fit analysis reveals that for all models, the residuals show that the variability of the responses within the sprayed and unsprayed plants is not equal. The plot of the predicted values against the transformed residuals shows very different spreads for the sprayed and unsprayed plants. Most models also showed a trend in the residuals when plotted against predicted values. The size of the deviations of deleted fixed effects estimates from actual estimates was similar between models. The size of the influence diagnostics varied between models with the restricted likelihood distance varying the most and the random effects model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN having the largest values for influence diagnostics. Subject 38, followed by subject 1, was shown to be a potential outlier in all models. Both these subjects were sprayed plants. The no random effects model with $\boldsymbol{\omega}_i$ = TOEP obtained the best transformed residuals compared to all other models considered. The residuals varied about the zero line and showed very little trend, but as in all the cases considered, the homogeneity of the residual variance was poor. The influence diagnostics were not the smallest compared to other models, but were smaller compared to the no random effects model with $\boldsymbol{\omega}_i$ = UN.

5.3.4 Fit of the covariance structures under the simplistic linear mean model

In order to determine how well the models performed in terms of the covariance structure fitted, graphical methods were employed. The methods include a semi-variogram-type approach and a plot of the covariances against the lags in time, as discussed in Chapter two. SAS proc mixed (ver. 9.1) does not include an option to display the empirical semi-variogram for the fitted linear mixed effect model. Applying the available generic function for the calculation of the empirical semi-variogram (PROC VARIOGRAM (SAS ver. 9.1)) is non-trivial, as it is specifically for spatial data and requires two distance variables, namely latitude and longitude. Longitudinal data only have one dimension, which is the length in time between observations. The equation of the empirical semi-variogram for longitudinal data is:

$$\gamma(h_{ijk}) = \tfrac{1}{2}\operatorname{var}(r_{ij}) + \tfrac{1}{2}\operatorname{var}(r_{ik}) - \operatorname{cov}(r_{ij,}r_{ik})$$

where i represents the subject, j and k are measurement occasions, and r_{ij} is the residual on subject i at time j, and h_{ijk} is the distance in time between times j and k. The transformed residuals (r_{ij}^*) were used in place of the ordinary residuals, therefore, if the covariance is correctly specified, then the plot of the semi-variogram should show a horizontal scatter around the value of one. In order to obtain an estimate of var(r_{ij}^*), the variance of the residuals obtained between subjects at time j was calculated. Similarly, the covariances of the residuals between times j and k were calculated for all subjects in order to obtain an estimate of cov(r_{ij}^*, r_{ik}^*). These values were then used to calculate $\hat{\gamma}(h_{ijk}^*)$ and were plotted as a function of the distance in time between measurements. The plot of the semi-variogram can be used to determine if there is misfit of the covariance model.

To obtain a plot of the covariances, the estimated covariance matrices for the random errors and random effects (as shown in Appendix B1) were used to calculate the estimated covariance matrix of the responses, $\hat{\mathbf{V}}_i = \mathbf{Z}_i\hat{\mathbf{\Sigma}}\mathbf{Z}_i' + \hat{\boldsymbol{\omega}}_i$. These covariances were also plotted as a function of the lags between observations. This plot gives a visual representation of the covariance matrix, and is useful for comparing the variance and covariance estimates between different models. The values plotted at lag zero are the estimated variances of the response at each week.

5.3.4.1 No random effects model with $\boldsymbol{\omega}_i$ = UN

From the analysis of the information criteria, the model which performed best was the no random effects model with $\boldsymbol{\omega}_i$ = UN. A plot of the approximate semi-variogram of the transformed residuals (Fig. 5.20) shows that that the points in general are close to one, with only a small number occurring below 0.8. The plot of the covariances (Fig. 5.20) gives a visual description of the covariance structure. It shows that the values of the variances (those at lag 0) are quite high relative to the values of the covariances, with week 0's variance being the highest. At one weeks lag, a very large covariance estimate was obtained relative to other covariance at this lag, suggesting that this model may be over-fitting. The covariances show a decreasing trend towards five weeks lag, but then increases again from five weeks to seven weeks lag, with a slight drop at eight weeks lag. The standard errors of this model (Table 5.2) are some of the lowest values compared to all other models.

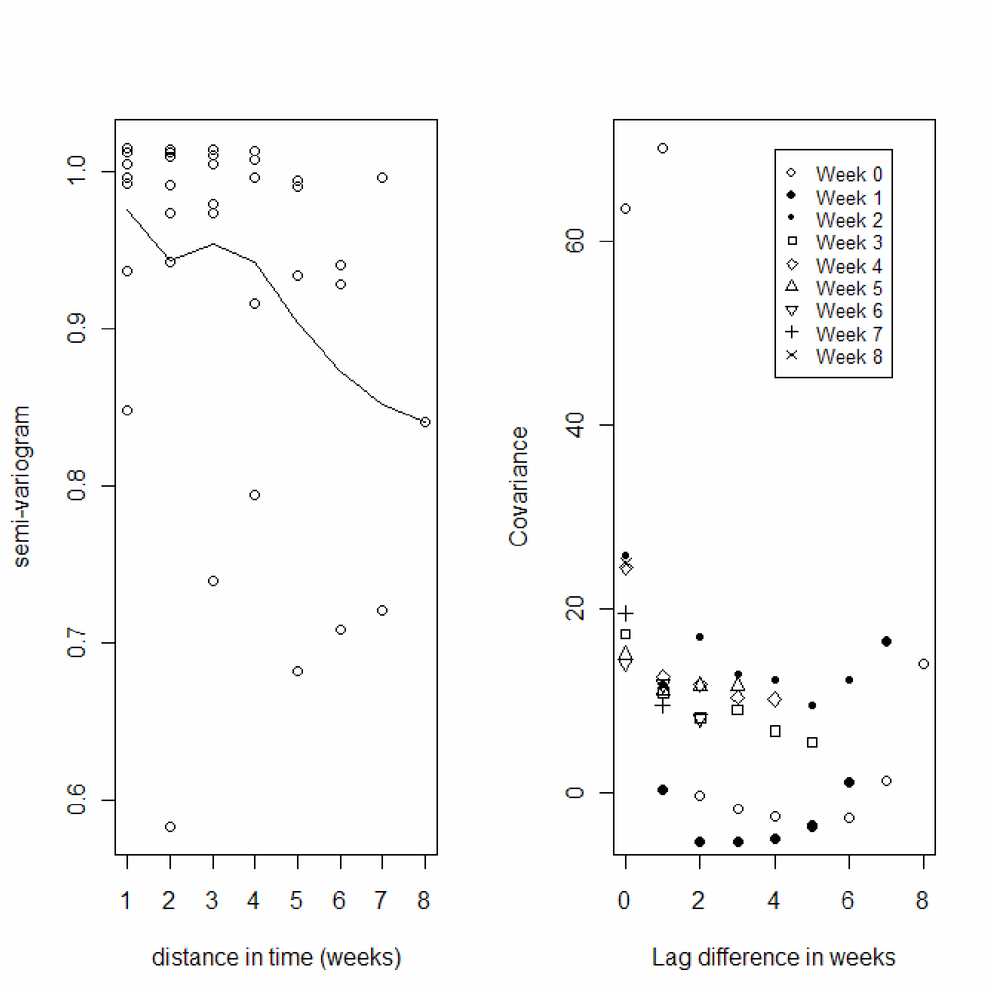

Fig. 5.20: Plot of semi-variogram (left) and the covariances (right) as function of lag in weeks between observations for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = UN.

5.3.4.2 Random intercept and slope model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH

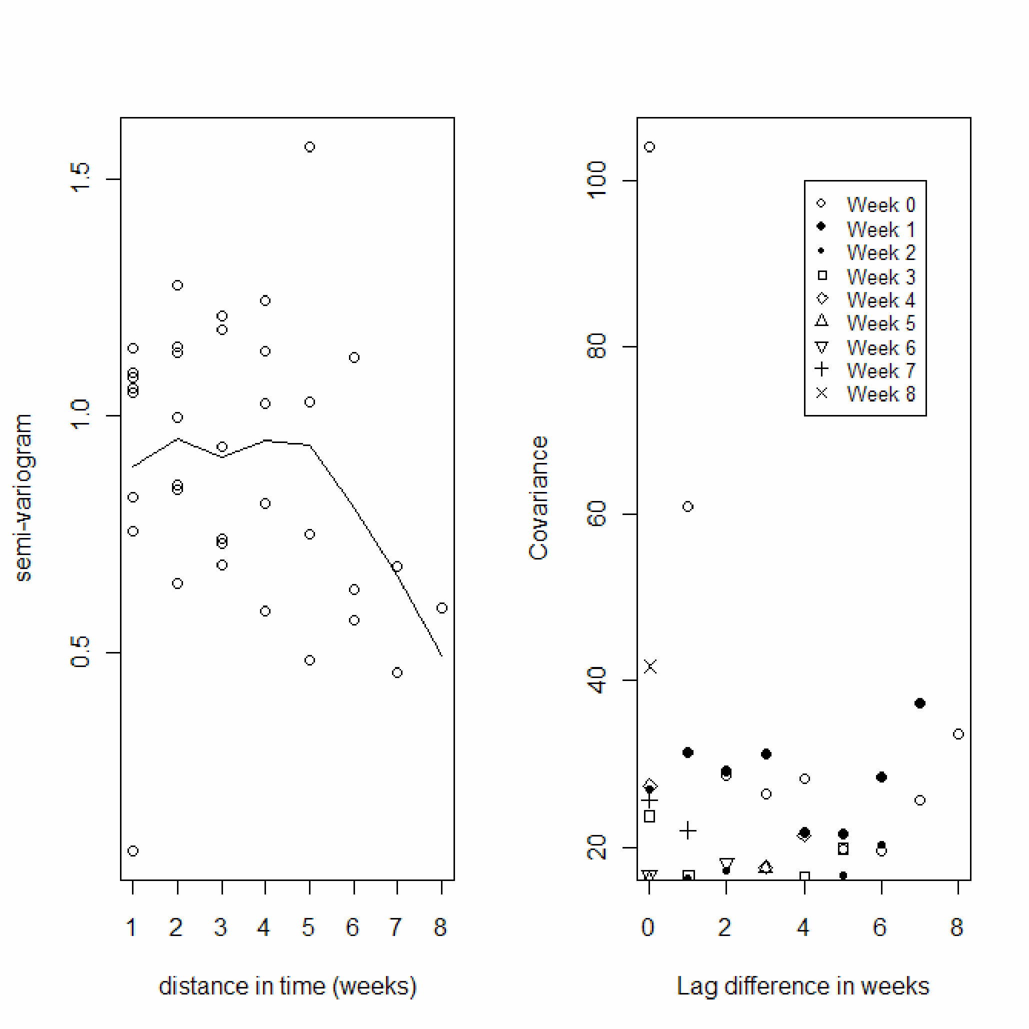

<u>Fig. 5.21: Plot of semi-variogram (left) and the covariances (right) as function of lag in weeks between observations for the simplistic linear model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH.</u>

The next best performing model was the random intercept and slope model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH. The semi-variogram of this model (Fig. 5.21) shows a relatively higher spread in values compared to the no random effects model with $\boldsymbol{\omega}_i$ = UN, but does seem to vary close to the value of one. The plot of the covariances (Fig. 5.21) indicates that not all the variances in this case have been estimated to be higher than the covariances, as was seen in the previous model. The maximum covariance value is much higher, with a value of over 100 estimated for the variance of week 0. This plot also shows an increase in the values of the covariances starting from around six weeks lag. The standard errors estimated for this model are slightly higher compared to the no random effects model with $\boldsymbol{\omega}_i$ = UN (Table 5.2). The plots for the random intercept and slope model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = UN were found to be identical to these covariance plots, as expected from the estimated covariance matrices obtained for these two models (Appendix B1).

5.3.4.3 Random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN

The random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN did not perform as well according to the goodness-of-fit analysis as the previous three models discussed. The approximate semi-variogram (Fig. 5.22) shows that most of the values are well below the value one, suggesting that the covariance structure may not be well specified. The plot of the covariances (Fig. 5.22) shows a decreasing trend in the covariance estimates as the lag in weeks increases. The size of the variances estimated at lag zero are relatively smaller compared to the previous two models, and the variances are more similar in size. Therefore this structure is more restrictive compared to the no random effects model with $\boldsymbol{\omega}_i$ = UN. The standard errors estimated for this model were some of the highest estimated among all the analysed models (Table 5.2).

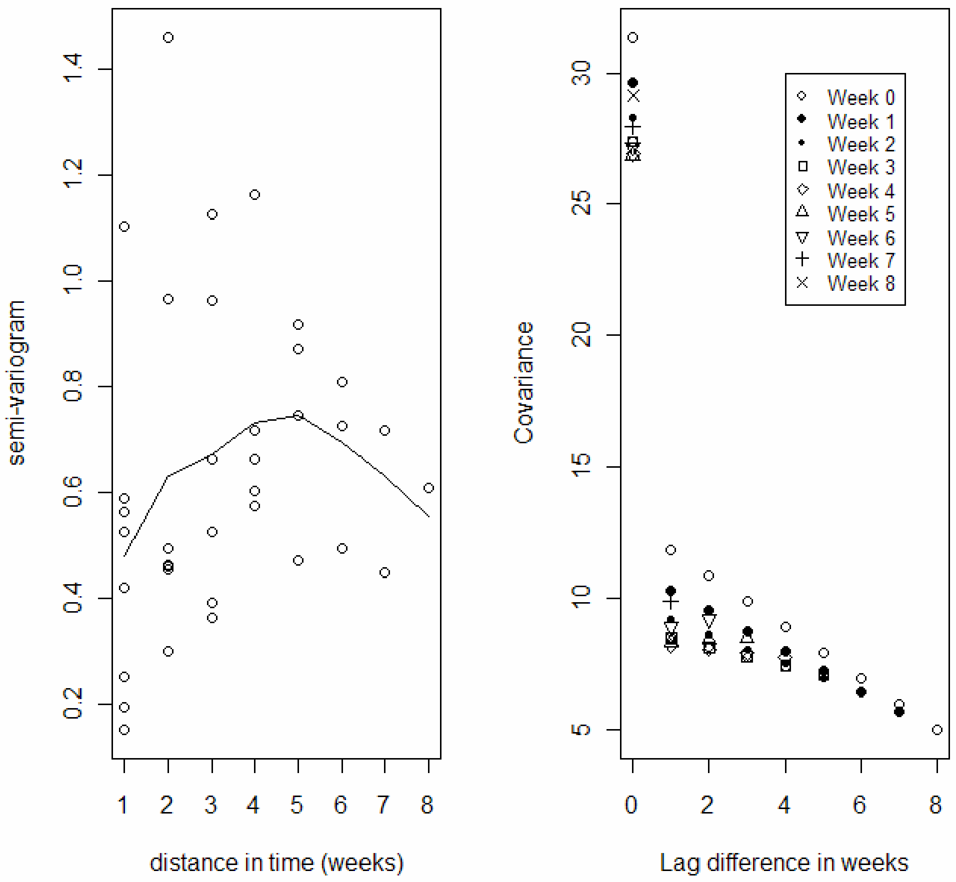

Fig. 5.22: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the simplistic linear model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

5.3.4.4 Random intercept model with $\boldsymbol{\omega}_i$ = AR(1)

The random intercept model with $\boldsymbol{\omega}_i$ = AR(1) performed slightly better than the previous model, obtaining slightly lower AIC, BIC and AICc values. An analysis of the approximate semi-variogram (Fig. 5.23) shows that the values are close to the value one, but mostly falling below one. The spread of semi-variogram values is not as great compared to the previous model. The covariance plot (Fig. 5.23) shows estimates of the variances being around 30 and then the covariance estimates starting from values around 15 and decreasing as the lag in weeks increases. The standard errors for this model are larger compared to the best fitting model, and close to the estimated standard errors obtained for the previous model and for the random effects model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH (Table 5.2).

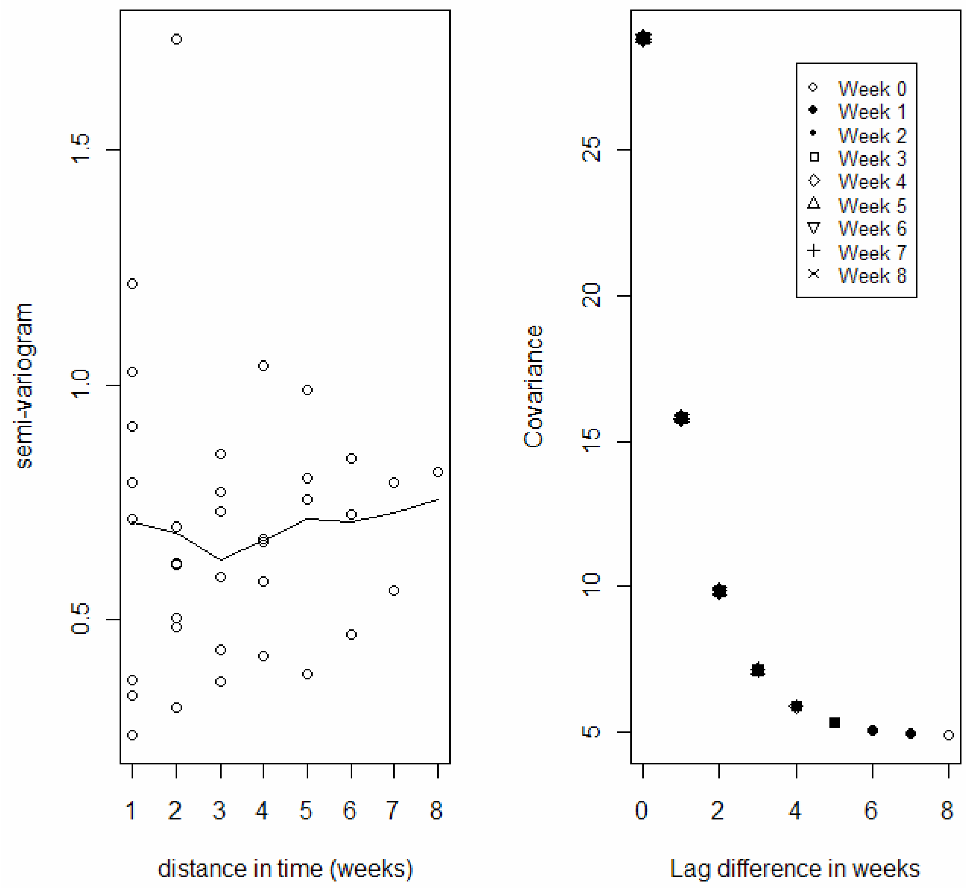

Fig. 5.23: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the random intercept simplistic linear model with $\boldsymbol{\omega}_i$ = AR(1).

5.3.4.5 No random effects model with $\boldsymbol{\omega}_i$ = TOEP

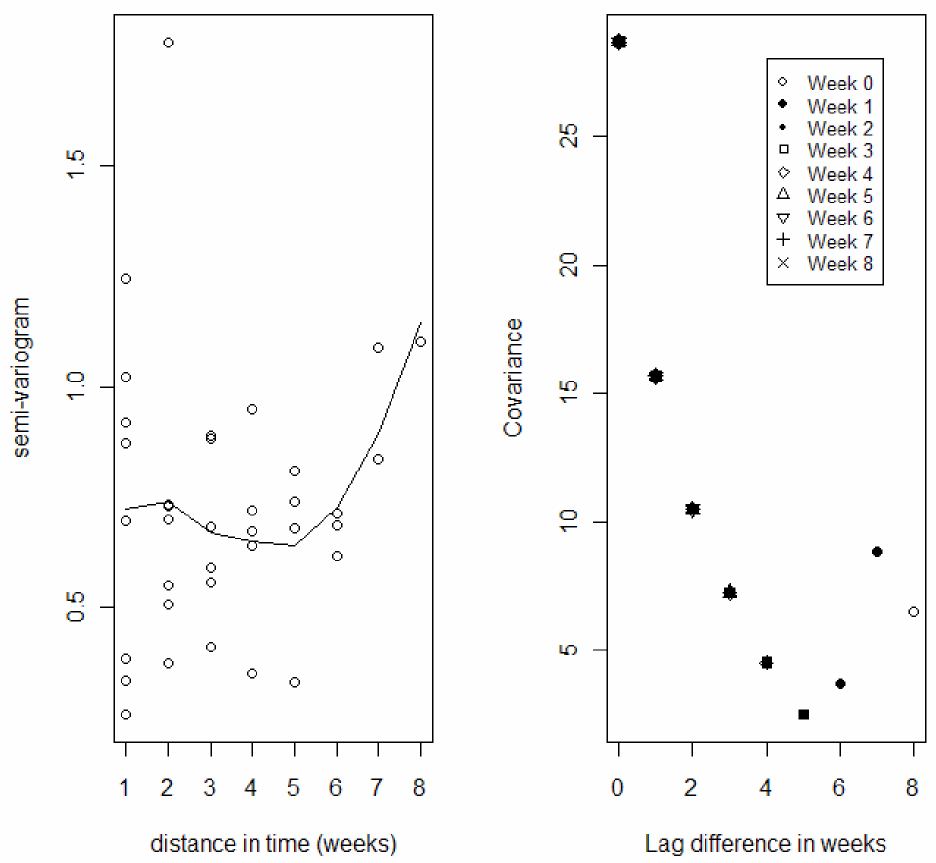

Fig. 5.24: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = TOEP.

The no random effects model with $\boldsymbol{\omega}_i$ = TOEP obtained slightly lower AIC, BIC and AICc values compared to the previous AR(1) model. The plot of the approximate semi-variogram (Fig. 5.24) shows that the values are close to one, but most values are below one. The range of semi-variogram values is similar to that observed for the previous model, with most values between 0.5 and 1.0. The plot of the covariance shows that the TOEP structure was able to obtain the same basic structure inherent in the no random effects model with $\boldsymbol{\omega}_i$ = UN. The covariances initially show a decline in value as the lag in weeks increases, but then from week 5 the covariances begin to increase, with a dip at week 8 (Fig. 5.24). The standard errors estimated are similar to those obtained for the previous AR(1) model and for the model with $\boldsymbol{\omega}_i$ = CSH and $\boldsymbol{\Sigma}$ = CSH (Table 5.2).

5.3.4.6 OLS model

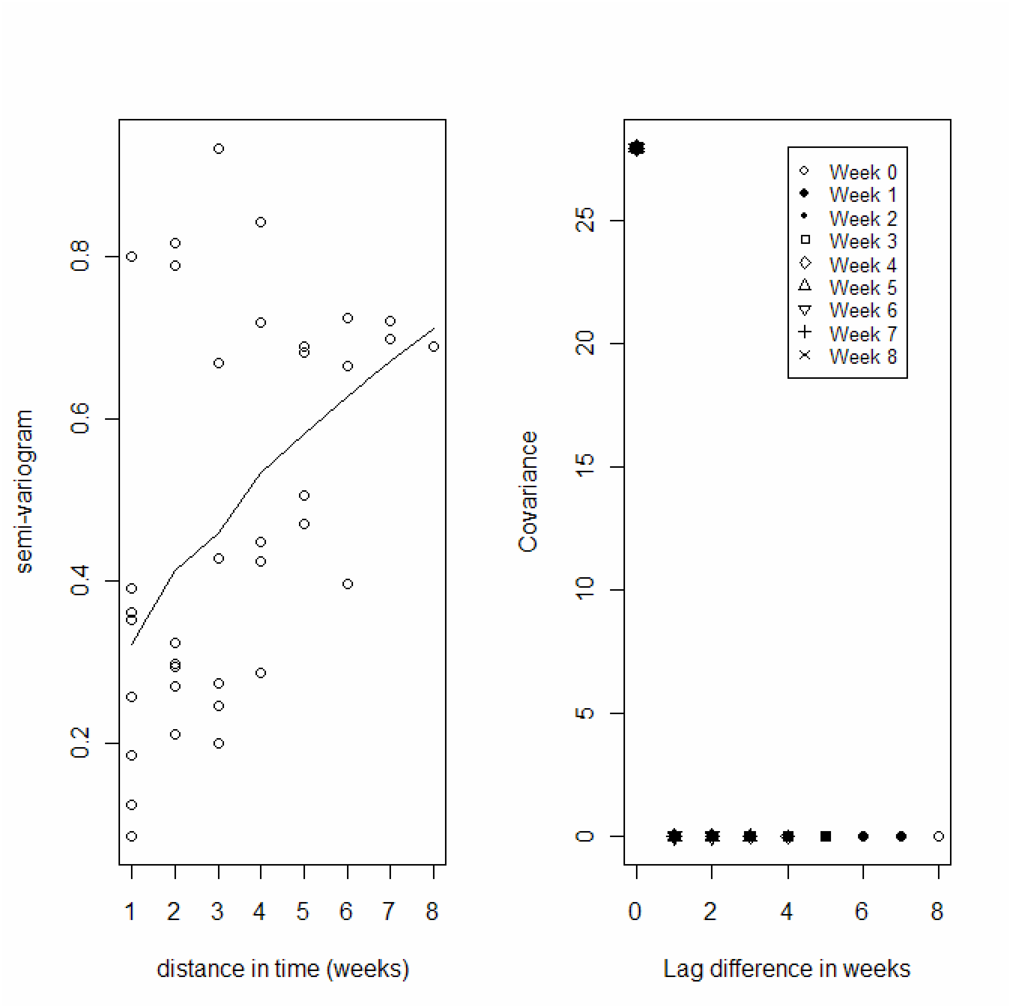

Fig. 5.25: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the no random effects simplistic linear model with $\boldsymbol{\omega}_i$ = VC.

The covariances of the OLS model were also analysed. This model performed the worst in terms of the information cirtieria out of all models under consideration. The plot of the semi-variogram (Fig. 5.25) shows values well below the value one, indicating that the covariance is not correctly specified. The plot of the covariances displays the zero estimates for the covariance terms, as assumed for the OLS model, and an estimate of 28 for the variance, which is assumed to be the same at all weeks (Fig. 5.25). As opposed to what would be expected from the literature (e.g. Weiss, 2005), the estimates of the standard errors for the OLS models were relative small compared to other models, which would result in narrower confidence intervals compared to other models (Table 5.2). The estimates of the standard errors are close to those obtained for the no random effects model with $\boldsymbol{\omega}_i$ = UN. This may be due to the absence of the large variance estimate observed in the best fitting models. The variance estimate for the OLS model was just below thirty, similar to the estimate of the variance from the random intercept model with AR(1) errors and the no random effects model with TOEP error structure.

5.3.4.7 Summary

The semi-variogram plots help to identify models where the covariance model has been misspecified. The plots for the models considered indicate that the OLS model, which assumed independence of observations, had semi-variogram values all below zero, and had the most incorrect specification for the covariance structure compared to all other models. Comparing the covariance plots of the OLS model to the best fitting model indicates that covariance matrix under the OLS model is very different from that specified by the best fitting model. The no random effects model with $\boldsymbol{\omega}_i$ = TOEP obtained reasonable values for the semi-variogram, and obtained a very similar covariance plot to that of the best fitting model, even though the random effects model with $\boldsymbol{\omega}_i$ = TOEP only has nine parameters, compared to the 45 parameters of the random effects model with $\boldsymbol{\omega}_i$ = UN.

5.3.5 Discussion on the simplistic linear mean model

Interpretability of linear mixed effects models is an important consideration when deciding upon a model. The fixed effect parameter estimates can be considered as the average effect of the predictor on the response. The mean structure in this case is described by equation (1), which uses sprayed plants as the base group. All of the fitted models predict a positive value for the intercept (Table 5.2). This value can be considered as the average petiole length of the sprayed plants at the beginning of the experiment. The estimate for β_2 is the average difference in lengths between the unsprayed and sprayed plants at the beginning of the experiment. All models predict this value to be small but negative. The estimate for β_1 is the average increase or decrease in petiole length of the sprayed plants over each week. There are large differences between models for this estimate. It is predicted to be negative and around three by all the models, except the models with UN or heterogeneous error covariance structures, which predicted it to be around six. The estimate for β_3 is the average difference in changes in the petiole length of the unsprayed plants compared to the sprayed plants. This was estimated to be negative and around one for most models, expect again the UN and heterogeneous error covariance models which predicted it to be around 0.5. Overall, the models shown to be more robust predicted bigger differences in the change in the second petiole length of sprayed plants versus unsprayed plants, with the second petioles of the unsprayed plants becoming shorter.

The parameters of the covariance structures can also be interpreted, but this is easier done for less complex covariance structures (estimates available in Appendix B1). The σ^2 term of the VC error covariance structure is the predicted variance in observations of a subject. In this example, the models fitted with VC error covariance structures have estimated σ^2 values around 20. This value changes depending on how much of the overall variance of the model is attributed to differences between subjects, and how much is attributed to random error. The AR(1) covariance structure has two parameters, σ^2 and ρ. The σ^2 term can be interpreted in the same way as for the VC error structure. The ρ term is the estimated correlation between observations from a subject one time unit apart. The autoregressive models in the example had estimated variances close to 25 and correlation coefficients around 0.5, indicating quite high amounts of correlation between observations close in time. The TOEP structure is a little more flexible in that it assumes that observations that are the same number of time units apart have the same correlation, but this does not depend on the correlation of units that are different time units apart. This means that there are many more parameters that need to be interpreted, increasing as the number of measurement occasions increases. In this example, the estimates of the TOEP covariance structure show that the within individual variance of the observations is close to 30, and that the covariance between observations decreases with increasing number of lags until lag five is reached, and then increases from lags five to seven. From lags seven to eight a decrease in the covariance is estimated.

The covariance structures which have been discussed so far all assume that the variance is constant across all measurement occasions. The UN and heterogeneous covariance structures do not make this assumption, increasing the number of estimated parameters and the complexity allowed in the covariance matrix, and increasing the difficulty in interpreting these values. The values for these covariance parameters are very dependent on the analysed data set, and will show much larger changes if a new data set were selected from the same population compared to the previous covariances. These structures are more appropriate for the random effects which should show flexibility as different subjects are included in a population.

From the goodness-of-fit analysis, it appears that the error structure inherent in these data is quite complex as the more flexible models consistently performed better. Of the simpler models, the no random effects model with $\boldsymbol{\omega}_i$ = TOEP performed best. This is most likely due to the more flexible structures available through this specification compared to simpler models, such as the AR(1) model which makes very restrictive assumptions about the covariance structure at different time lags. In this particular case, the AR(1) structure performed particularly poorly as most covariance structures involving an AR(1) specification failed to obtain valid estimates, indicating that this is not an appropriate covariance structure for these data. By exploring the raw data (Fig. 5.1), it does appear that the mean length of the second petiole increases towards the end of the experiment. This would mean that the values for second petiole length would be closer to the values obtained at the beginning of the experiment than to those obtained after the initial drop off in second petiole length. This phenomenon can be captured in the TOEP structure, as it is possible to obtain higher covariances at later lags in time compared to earlier lags. A reason for this increase in petiole length at the end of the experiment could be due to the increased density of plants in the tubs (as asexual reproduction of these would have been taking place throughout the eight week period), resulting in the growth form of the plant adapting to produce longer petioles in an attempt to compete for light. Therefore the more complex covariance structures reveal a particular biological phenomenon in the data.

As concluded in the simulation study, the OLS model performed worst compared to all other models considered. Therefore it is inappropriate to obtain estimates via OLS estimation if the data is longitudinal and may potentially have a complex covariance structure, particularly if there is interest in the variability within the data.

Residual analysis of each of the model structures considered in this simplistic analysis clearly indicates that the variance of the residuals for the unsprayed and spray plants is not the same. Therefore methods for accounting for this difference in variances should be considered.

A second obvious problem is that the linear mean structure considered is not appropriate for the data, as can be observed in Fig. 5.1. The shape of the curve seems to be somewhat parabolic. To correct the linear model, additional terms could be added into the mean structure, such as a quadratic in time, or adjustment parameters. More advanced models will be considered in the next section to improve the model fit and appropriateness.

5.4 More Complex Quadratic Mean Structure Analysis

Two methods were considered in order to improve the homogeneity of the residual variance. The first method accounted for difference in variance between the sprayed and unsprayed plants by adding an additional random effect for the categorical herbicide variable. The idea behind this method was that an additional variance term would be added to the random effects covariance structure for herbicide, and this would result in differing covariance matrices, $\mathbf{V}_i$, for the sprayed and unsprayed plants. This method was not successfully implemented, as the addition of herbicide as a random effect resulted in non-positive-definiteness of the final Hessian matrix.

The second method implemented was to log the responses. This method was successfully implemented, and improved the homogeneity of the variances of the residuals.

To improve the appropriateness of the mean structure, a quadratic term for *Week* and adjustment terms for the start of the experiment and the first week were added. The adjustment terms were added as it appears from Fig. 5.1 that a change in the growth process occurred after the first week. These additional indicator variables would result in estimates of the deviation away from the curve predicted over the full time period, accounting for this change in process. These terms were added into the mean structure (now modelling the logged length of the second petiole), and were found to be highly significant. The results for this analysis are presented in the following sections.

5.4.1 Information criteria obtained for three different quadratic mean models

As new terms were added into the mean structure, it was necessary to include the nutrient level in the mean structure once again. The fullest mean structure considered was:

$$
\begin{aligned}
\log(y_i) = {} & \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Nutrients1) + \beta_4(Nutrients2) \\
& + \beta_5(Week0) + \beta_6(Week1) + \beta_7(Week^2) \\
& + \beta_8(Herbicide)\times(Week) + \beta_9(Nutrients1)\times(Week) + \beta_{10}(Nutrients2)\times(Week) \\
& + \beta_{11}(Herbicide)\times(Week0) + \beta_{12}(Nutrients1)\times(Week0) + \beta_{13}(Nutrients2)\times(Week0) \\
& + \beta_{14}(Herbicide)\times(Week1) + \beta_{15}(Nutrients1)\times(Week1) + \beta_{16}(Nutrients2)\times(Week1) \\
& + \beta_{17}(Herbicide)\times(Week^2) + \beta_{18}(Nutrients1)\times(Week^2) + \beta_{19}(Nutrients2)\times(Week^2) \\
& + \beta_{20}(Nutrients1)\times(Herbicide) + \beta_{21}(Nutrients2)\times(Herbicide) \\
& + b_{1i} + b_{2i}(Week) + \varepsilon_i
\end{aligned}
$$

where all terms are as previously described, $Week^2$ is a quadratic term for *Week*, and *Week0* and *Week1* are the adjustment parameters for the start of the experiment and the first week respectively. An analysis of this model under all covariance structures showed that the nutrient effect was not significant – on it's own or in interaction terms. The details of this analysis are given in the next section.

The second mean structure considered was the fullest structure excluding the nutrient effect:

$$
\begin{aligned}
\log(y_i) = {} & \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Week0) + \beta_4(Week1) + \beta_5(Week^2) \\
& + \beta_6(Herbicide)\times(Week) + \beta_7(Herbicide)\times(Week0) + \beta_8(Herbicide)\times(Week1) \\
& + \beta_9(Herbicide)\times(Week^2) + b_{1i} + b_{2i}(Week) + \varepsilon_i
\end{aligned}
$$

where the terms are as described for the previous model. An analysis of this mean structure under various covariance structures showed that the interaction term between *Week* and *Herbicide* was non-significant. Therefore this term was dropped from the mean structure, resulting in the final mean structure considered:

$$
\begin{aligned}
\log(y_i) = {} & \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Week0) + \beta_4(Week1) + \beta_5(Week^2) \\
& + \beta_6(Herbicide)\times(Week) + \beta_7(Herbicide)\times(Week0) + \beta_8(Herbicide)\times(Week1) \\
& + b_{1i} + b_{2i}(Week) + \varepsilon_i .
\end{aligned}
$$

The information criteria obtained for each mean model, under each covariance structure considered, are presented in Table 5.3. As it is clear from the simulation study, and from the previous analysis, that selecting a random effects covariance matrix which assumes equal elements along the diagonal is inappropriate (resulting in non-positive definite random effects covariance structures), these covariance structures (i.e. CS, TOEP and AR(1)) were not considered for the random effects.

The information criteria from the three different mean models are presented in Table 5.3. Models under all covariance structures obtained the lowest information criteria for the third mean structure. The covariance structure which obtained the lowest AIC and AICc values of -167.9 and -159.3 for the third mean model was the no random effects model with unstructured errors. The lowest BIC value of -98.0 was obtained for the random intercept and slope model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN/CSH/ARH(1). Selecting UN, CSH or ARH(1) for the random effects covariance structure results in equivalent model results, therefore the results for these random effects covariance structures appear together.

In contrast to the simplistic linear model, the models with error covariance structures of AR(1) or ARH(1) performed well, obtaining some of the best values for the AIC, BIC and AICc, whereas for the previous simpler linear model, these error structure resulted in invalid covariance estimates when included with random effects. In the case of the quadratic model, the ARH(1) error structure performed better compared the CSH structure, which resulted in invalid parameter estimates.

Table 5.3: Information criteria obtained under three different quadratic mean models for each of the covariance structures specified.

Successfully fitted models				
$\log(y_i) = \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Nutrients1) + \beta_4(Nutrients2) + \beta_5(Week0) + \beta_6(Week1) + \beta_7(Week^2) + \beta_8(Herbicide)\times(Week) + \beta_9(Nutrients1)\times(Week) + \beta_{10}(Nutrients2)\times(Week) + \beta_{11}(Herbicide)\times(Week0) + \beta_{12}(Nutrients1)\times(Week0) + \beta_{13}(Nutrients2)\times(Week0) + \beta_{14}(Herbicide)\times(Week1) + \beta_{15}(Nutrients1)\times(Week1) + \beta_{16}(Nutrients2)\times(Week1) + \beta_{17}(Herbicide)\times(Week^2) + \beta_{18}(Nutrients1)\times(Week^2) + \beta_{19}(Nutrients2)\times(Week^2) + \beta_{20}(Nutrients1)\times(Herbicide) + \beta_{21}(Nutrients2)\times(Herbicide) + b_{1i} + b_{2i}(Week) + \varepsilon_i$				
ω_i	**Σ**	**AIC**	**BIC**	**AICc**
UN	None	-125.4	-31.2	-116.6
ARH(1)	UN/CSH/ARH(1)	-74.0	-46.8	-73.3
ARH(1)	VC	-48.2	-23.1	-47.6
ARH(1)	Intercept only	-46.9	-23.9	-46.4
AR(1)	Intercept only	-33.2	-26.9	-33.2
AR(1)	VC	-32.8	-24.4	-32.7
AR(1)	UN/CSH/ARH(1)	-31.0	-20.6	-30.9
TOEP	None	-28.6	-9.7	-28.2
ARH(1)	None	-28.0	-7.1	-27.6
CSH	VC	-25.5	-0.3	-24.8
VC	UN/CSH/ARH(1)	-22.6	-6.2	-14.5
AR(1)	None	-12.8	-8.6	-12.7
CSH	Intercept only	-12.5	10.5	-12.0
VC	VC	-11.8	-5.5	-11.7
CSH	None	1.6	22.5	2.0
CS	None	2.9	7.1	2.9
VC	Intercept only	2.9	7.1	2.9
VC	None	158.6	162.8	158.6
$\log(y_i) = \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Week0) + \beta_4(Week1) + \beta_5(Week^2) + \beta_6(Herbicide)\times(Week) + \beta_7(Herbicide)\times(Week0) + \beta_8(Herbicide)\times(Week1) + \beta_9(Herbicide)\times(Week^2) + b_{1i} + b_{2i}(Week) + \varepsilon_i$				
ω_i	**Σ**	**AIC**	**BIC**	**AICc**
UN	None	-160.6	-66.4	-152.1
ARH(1)	UN/CSH/ARH(1)	-117.2	-89.9	-116.5
ARH(1)	VC	-91.3	-66.2	-90.7
ARH(1)	Intercept only	-86.9	-63.8	-86.4
AR(1)	VC	-73.0	-64.6	-72.9

Table 5.3 (cont.): Information criteria obtained under three different mean models for each of the covariance structures specified.

ω_i	Σ	AIC	BIC	AICc
AR(1)	UN/CSH/ARH(1)	-71.2	-60.7	-71.1
AR(1)	Intercept only	-71.0	-64.7	-70.9
ARH(1)	None	-69.3	-48.3	-68.9
TOEP	None	-68.7	-49.8	-68.3
CSH	VC	-65.6	-40.5	-65.0
VC	UN/CSH/ARH(1)	-54.0	-45.6	-53.9
AR(1)	None	-52.4	-48.2	-52.4
VC	VC	-51.3	-45.0	-51.2
CSH	Intercept only	-42.2	-19.2	-41.7
CS	None	-27.0	-22.8	-27.0
VC	Intercept only	-27.0	-22.8	-27.0
CSH	None	-25.3	-4.3	-24.8
VC	None	140.3	144.6	140.3

$$\log(y_i) = \beta_0 + \beta_1(Herbicide) + \beta_2(Week) + \beta_3(Week0) + \beta_4(Week1) + \beta_5(Week^2) + \beta_6(Herbicide)\times(Week) + \beta_7(Herbicide)\times(Week0) + \beta_8(Herbicide)\times(Week1) + b_{1i} + b_{2i}(Week) + \varepsilon_i$$

ω_i	Σ	AIC	BIC	AICc
UN	None	-167.9	-73.6	-159.3
ARH(1)	UN/CSH/ARH(1)	-125.3	-98.0	-124.6
ARH(1)	VC	-99.6	-74.5	-99.0
ARH(1)	Intercept only	-95.0	-72.0	-94.5
AR(1)	VC	-81.4	-73.1	-81.4
AR(1)	UN/CSH/ARH(1)	-79.7	-62.2	-75.5
AR(1)	Intercept only	-79.3	-73.0	-79.3
ARH(1)	None	-77.3	-56.3	-76.8
TOEP	None	-77.1	-58.3	-76.8
CSH	VC	-73.8	-48.7	-73.2
VC	UN/CSH/ARH(1)	-62.6	-54.3	-62.6
AR(1)	None	-60.5	-56.3	-60.4
VC	VC	-59.9	-53.6	-59.9
CSH	Intercept only	-50.4	-27.4	-49.9
CS	None	-35.5	-31.3	-35.4
VC	Intercept only	-35.4	-31.3	-35.4
CSH	None	-33.6	-12.7	-33.2
VC	None	132.3	136.6	132.3

Table 5.3 (cont.): Information criteria obtained under three different mean models for each of the covariance structures specified.

Models which had a final Hessian matrix which was not positive definite				
$\boldsymbol{\omega}_i$	$\boldsymbol{\Sigma}$			
CS	UN/CSH/ARH(1)			
CS	VC			
CS	Intercept only			
TOEP	UN/CSH/ARH(1)			
TOEP	VC			
TOEP	Intercept only			
UN	UN			
Models which did not converge				
$\boldsymbol{\omega}_i$	$\boldsymbol{\Sigma}$			
TOEP	VC			
CSH	UN/CSH/ARH(1)			
UN	UN/CSH/ARH(1)			
UN	VC			
UN	Intercept only			

The models predicted to perform well from the simulation study, i.e. the no random effects model with TOEP error structure, the random intercept and slope model with AR(1) errors and unstructured random effects, and the random intercept and slope model with VC errors and unstructured random effects, obtained valid estimates and obtained AIC and AICc values below -60, and BIC values below -50. The random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN obtained the best information criteria of -79.7, -75.5 and -62.2 for the AIC, BIC and AICc respectively, followed by the no random effects model with $\boldsymbol{\omega}_i$ = TOEP and the random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

The OLS model obtained the highest information criteria, with an AIC and AICc of 132.3 each, and a BIC of 136.6. Therefore the information criteria show that the OLS model is inferior to models where the covariance structure had been modelled.

The observation to parameter ratio of the successfully fitted models ranged between 54:1 to 3.8:1. The lowest observation to parameter ratio was obtained for the random effects models with ARH(1) errors and random effects covariance structures of UN, CSH, and ARH(1). For the unsuccessfully fitted models, the highest observation to parameter ratio was 7.5:1, which was obtained for the random intercept model with a CS error structure. The majority of unsuccessfully fitted models had observation to parameter ratios of less than 5:1. A large proportion of the successfully fitted models also had a parameter ratio of less than 5:1, and this is mainly due to the large number of random effects (120 in the case of the random intercept and slope models) which needed to be estimated. The observation to parameter ratio if only taking into the random effects of the random intercept and slope models is 4.5:1.

5.4.2 Analysis of quadratic mean models

The mean structure including all the effects, including nutrient level, as well as the two-way interactions between these effects, had very few significant parameters (Table 5.4). If the nutrient effect is excluded from the model, the significance of the other mean parameters improved (Table 5.4). The model with all two-way effects excluding nutrients still had non-significant parameters (Table 5.4). Although the interaction term with *Herbicide* and *Week* has a larger p-value than the interaction term with *Herbicide* and $Week^2$, the interaction with the quadratic was excluded first, as it does not make sense to have a model where there is an interaction term with the quadratic, but not with the linear term. Excluding the interaction term with $Week^2$ resulted in an improved model where all the remaining interaction terms were significant.

Table 5.4: Fixed effects results for advanced mean structures, including effects for herbicide, nutrients and the adjustment variables. These estimates were obtained under a no random effects model with unstructured error covariance.

Effect	**Estimate**	**p-value**	**Lower 95% CI Limit**	**Upper 95% CI Limit**
Model with all two-way interactions				
Intercept	3.0320	<0.0001	2.7786	3.2854
Week	-0.0683	0.1287	-0.1569	0.0204
Week^2	0.0067	0.1140	-0.0017	0.0152
Week 0	0.1407	0.2511	-0.1024	0.3838
Week 1	0.2246	0.0177	0.0404	0.4088
Herbicide	-0.3226	0.0214	-0.5961	-0.0491
Nutrient1	-0.2266	0.1605	-0.5453	0.0921
Nutrient2	0.0210	0.8958	-0.2977	0.3397
Herbicide×Week	-0.1105	0.0155	-0.1992	-0.0218
Nutrient1×Week	0.0963	0.0811	-0.0123	0.2049
Nutrient2×Week	-0.0592	0.2794	-0.1678	0.0494
Herbicide×Week^2	0.0061	0.1511	-0.0023	0.0145
Nutrient1×Week^2	-0.0079	0.1287	-0.0183	0.0024
Nutrient2×Week^2	0.0041	0.4273	-0.0062	0.0144
Herbicide×Week 0	0.1507	0.2195	-0.0924	0.3938
Nutrient1×Week 0	0.1744	0.2457	-0.1234	0.4721
Nutrient2×Week 0	-0.0757	0.6126	-0.3734	0.2221
Herbicide×Week 1	0.2676	0.0052	0.0834	0.4518
Nutrient1×Week 1	0.1247	0.2729	-0.1009	0.3503
Nutrient2×Week 1	-0.0246	0.8281	-0.2502	0.2010
Nutrient1×Herbicide	0.1085	0.3104	-0.1039	0.3209
Nutrient2×Herbicide	0.1364	0.2035	-0.0760	0.3488

Table 5.4 (cont.): Fixed effects results for advanced mean structures, including effects for herbicide, nutrients and the adjustment variables. These estimates were obtained under a no random effects model with unstructured error covariance.

Fullest model with two-way interactions, excluding the nutrient effect				
Effect	**Estimate**	**p-value**	**Lower 95% CI Limit**	**Upper 95% CI Limit**
Intercept	2.9624	<0.0001	2.7853	3.1394
Week	-0.0558	0.0954	-0.1217	0.0101
Week^2	0.0055	0.0797	-0.0007	0.01159
Week 0	0.1741	0.0488	0.0009	0.3472
Week 1	0.2586	0.0002	0.1285	0.3887
Herbicide	-0.2359	0.0644	-0.4863	0.0145
Herbicide×Week	-0.1140	0.0174	-0.2072	-0.0208
Herbicide×Week^2	0.0064	0.1452	-0.0023	0.0151
Herbicide×Week0	0.1449	0.2409	-0.0999	0.3897
Herbicide×Week1	0.2640	0.0057	0.0799	0.4480
Fullest model with two-way interactions, excluding the nutrient effect and the interaction between Week^2 and Herbicide				
Effect	**Estimate**	**p-value**	**Lower 95% CI Limit**	**Upper 95% CI Limit**
Intercept	3.0400	<0.0001	2.8957	3.1843
Week	-0.0903	0.0004	-0.1383	-0.0424
Week^2	0.0088	0.0002	0.0044	0.0131
Week 0	0.1092	0.1517	-0.0410	0.2594
Week 1	0.2260	0.0005	0.1033	0.3487
Herbicide	-0.3867	<0.0001	-0.5313	-0.2420
Herbicide×Week	-0.0470	<0.0001	-0.0678	-0.0261
Herbicide×Week0	0.2709	0.0031	0.0955	0.4463
Herbicide×Week1	0.3272	0.0002	0.1644	0.4900

This final mean model has significant interactions between *Herbicide* and the time variables *Week*, *Week 0* and *Week 1* (Table 5.4). This can be interpreted to mean that the slope parameter differs between sprayed and unsprayed plants, as well as the adjustments required at the beginning of the experiment and at the first week to accommodate for a change in process which occurred after this time. Since the interaction term with the quadratic term is non-significant, it means that the curvature parameter does not differ significantly between sprayed and unsprayed plants. The derivative with respect to *Week* for both the sprayed plants and unsprayed plants' curve is increasing, and therefore, by definition, the curve will be concave up (Stewart, 1998). The adjustment terms, as well as the adjustment terms in interaction with *Herbicide*, are positive (Table 5.4), indicating that the values at weeks 0 and 1 need to be increased compared to the estimated quadratic over the full time period, and more so for the unsprayed plants. A plot of the predicted line superimposed over the observed mean line appears in Fig. 5.26. This plot shows that the observed mean values fit within the confidence limits of the predicted values, as they should, showing that the model is compatible with the data.

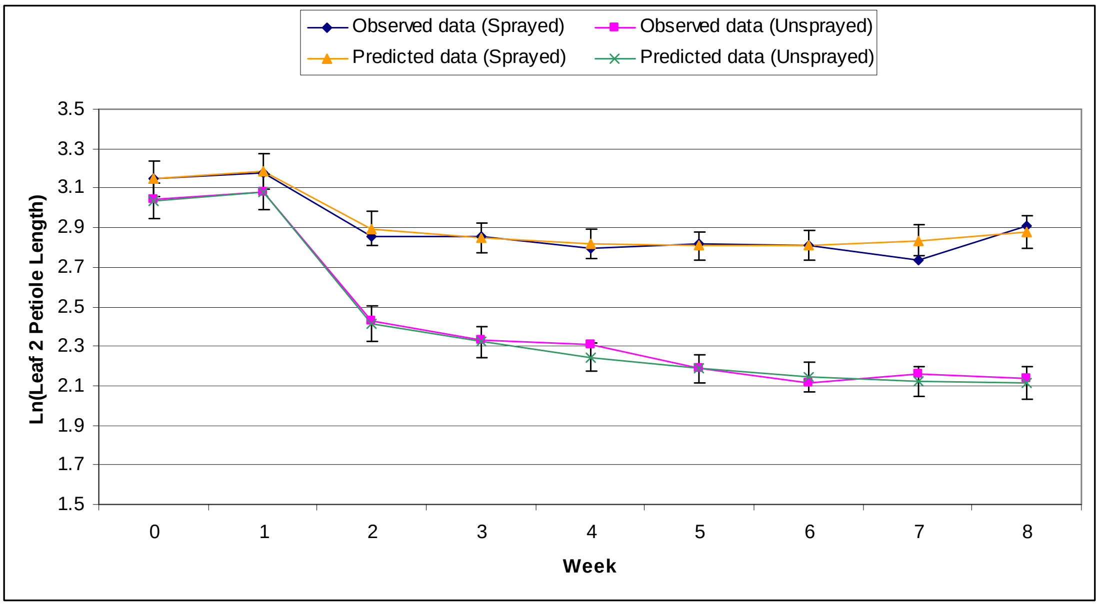

Fig. 5.26: Plot of the predicted mean log response over week superimposed over the plot for the observed mean log response. The error bars represent the 95% confidence interval of the predicted mean values.

5.4.3 Goodness-of-fit analysis of the quadratic model, with adjustment parameters, under various covariance structures

5.4.3.1 No random effects model with $\boldsymbol{\omega}_i$ *= UN*

As for the linear model, the no random effects model with unstructured error covariance matrix obtained the best fit compared to the models with other covariance structures according to the AIC and AICc values (Table 5.3). For the goodness-of-fit analysis of the quadratic model, the deleted estimates are not shown due to the large quantity of plots resulting from the additional fixed effects parameters. These plots are included in Appendix B2 which shows the full influence analysis of each of the models considered. The residual analysis plots shown in Fig. 5.27 show that the

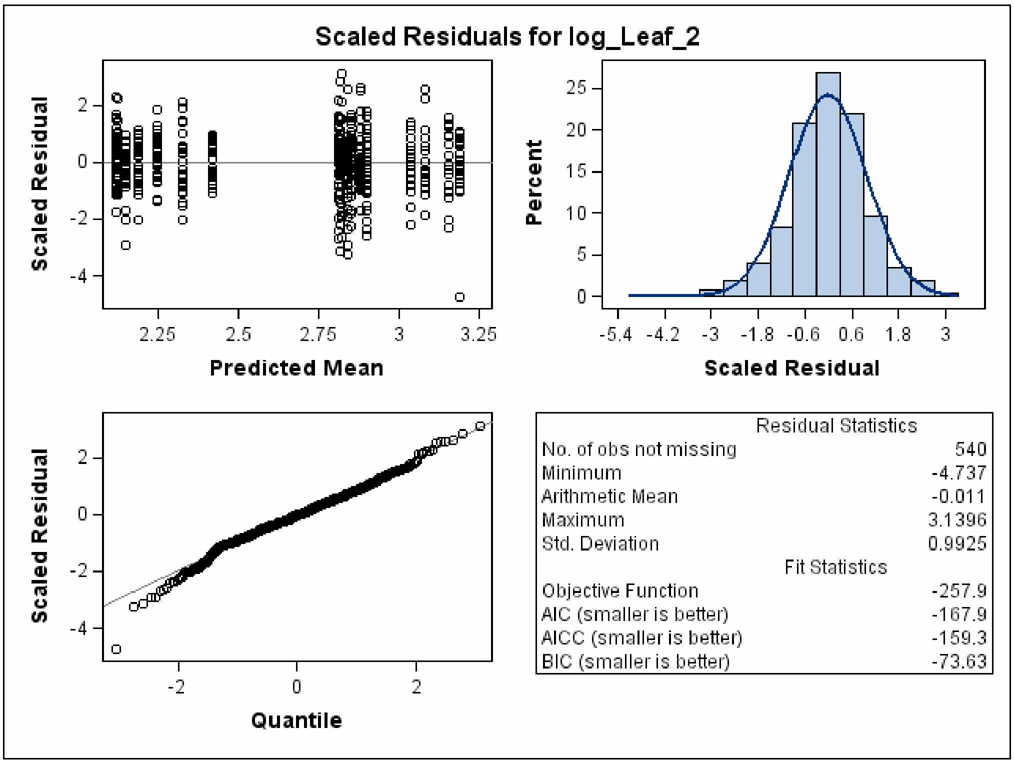

Fig. 5.27: Plots of scaled residuals for the quadratic model with no random effects and $\boldsymbol{\omega}_i$ = UN.

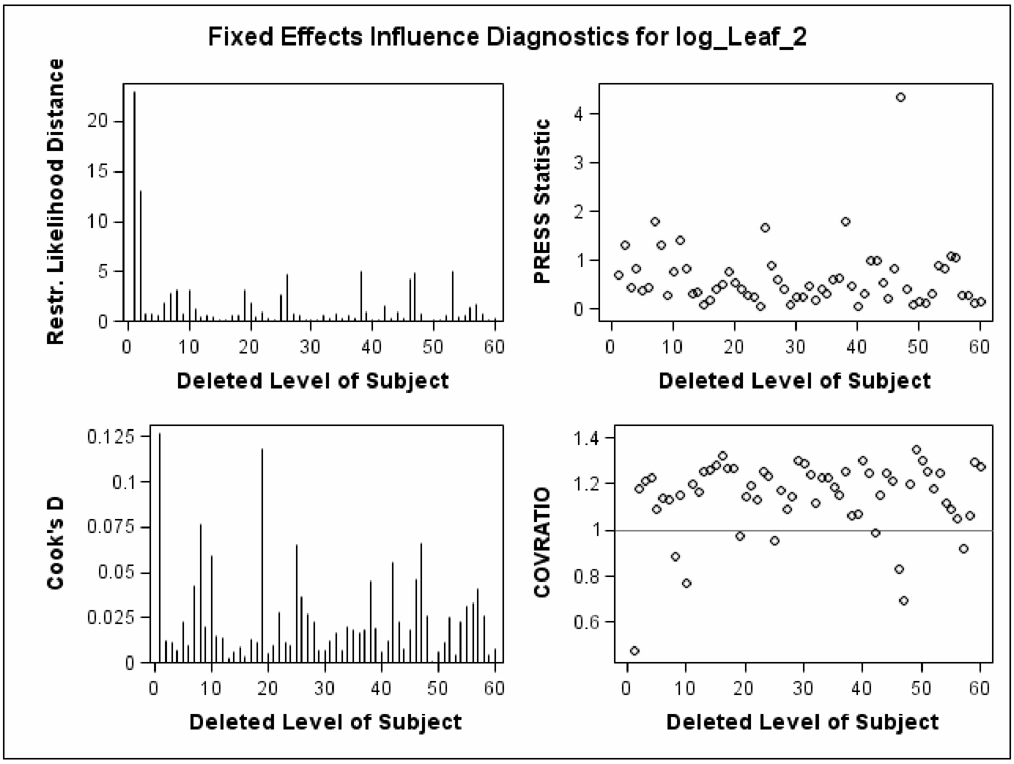

Fig. 5.28: Plots of the influence diagnostics for the quadratic model with no random effects and $\boldsymbol{\omega}_i$ = UN.

homogeneity of the residual variance is improved compared to the linear model, and the mean of the residuals is close to zero. This analysis also shows that the residuals are close to normality. The restricted likelihood distance highlights subjects 1 and 2 as potential outliers, Cook's D highlights subjects 1 and 19, the PRESS statistic highlights subject 47, and the COVRATIO highlights subject 1 (Fig. 5.28). All these subjects highlighted are sprayed plants. Therefore although the goodness-of-fit plots show that the quadratic model, with additional adjustment parameters, modelling the logged response is an improvement of the linear model, there appears to be some bias in the estimates for the sprayed plants.

5.4.3.2 Random intercept and slope model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN

The random intercept and slope model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN was the second best model, according the AIC and AICc values, and the best fitting model according to the BIC value. The variance of the residuals is relatively homogeneous, the mean of the residuals is close to zero, and the residuals are close to normality (Fig. 5.29). The restricted maximum likelihood values (Fig. 5.30) again highlight subjects 1 and 2, but are smaller in comparison to those of the previous model (Fig. 5.28). Cook's D is also smaller relative to the results from the previous model, and highlights subjects 1, 42 and 47 (Fig. 5.28). The PRESS statistics and COVRATIO are similar in size to those of the previous model, and highlight the same subjects (Fig. 5.28). Subject 42 is an unsprayed plant, but all the other highlighted observations are sprayed plants.

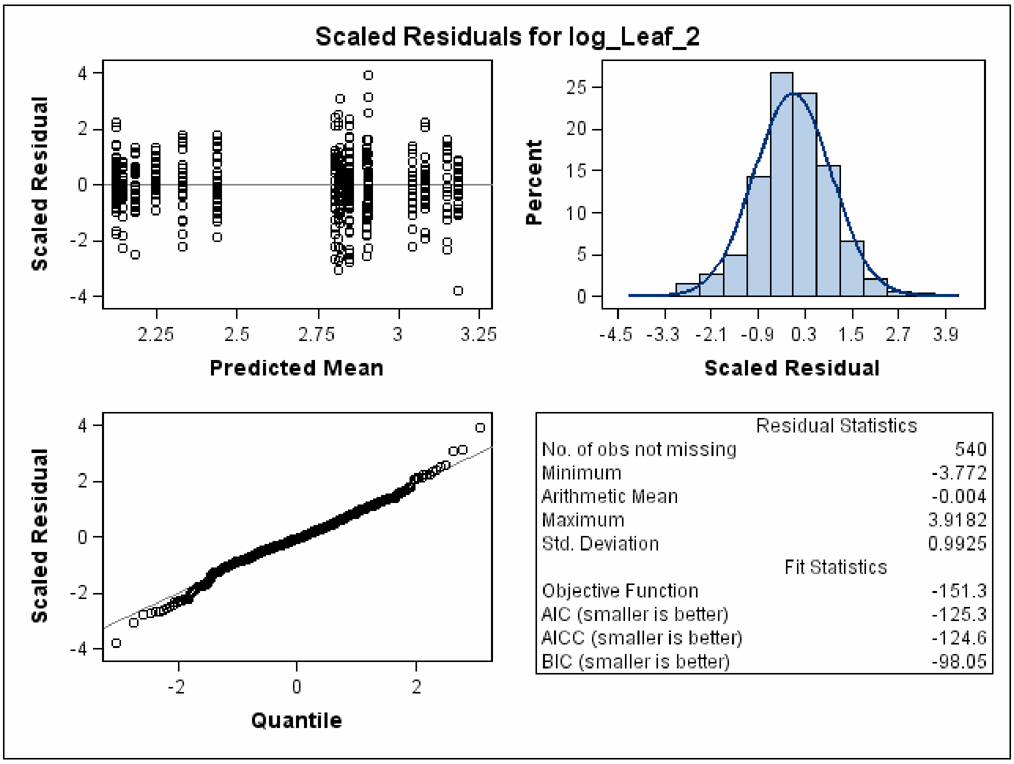

Fig. 5.29: Plots of scaled residuals for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN.

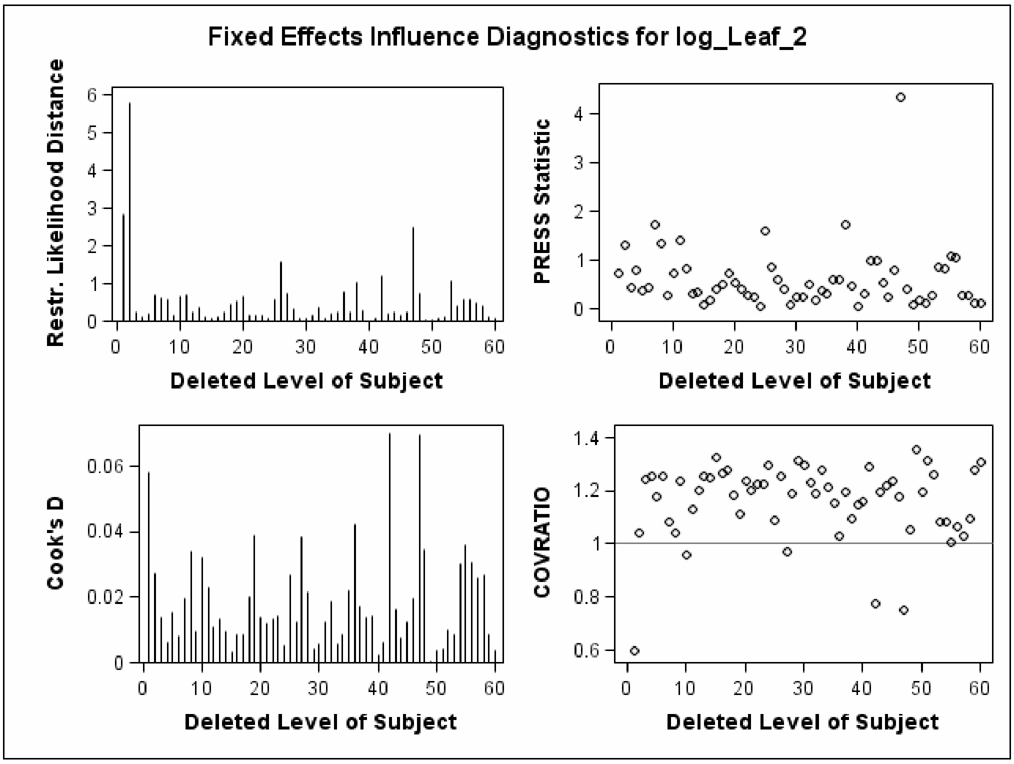

Fig. 5.30: Plots of influence diagnostics for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN.

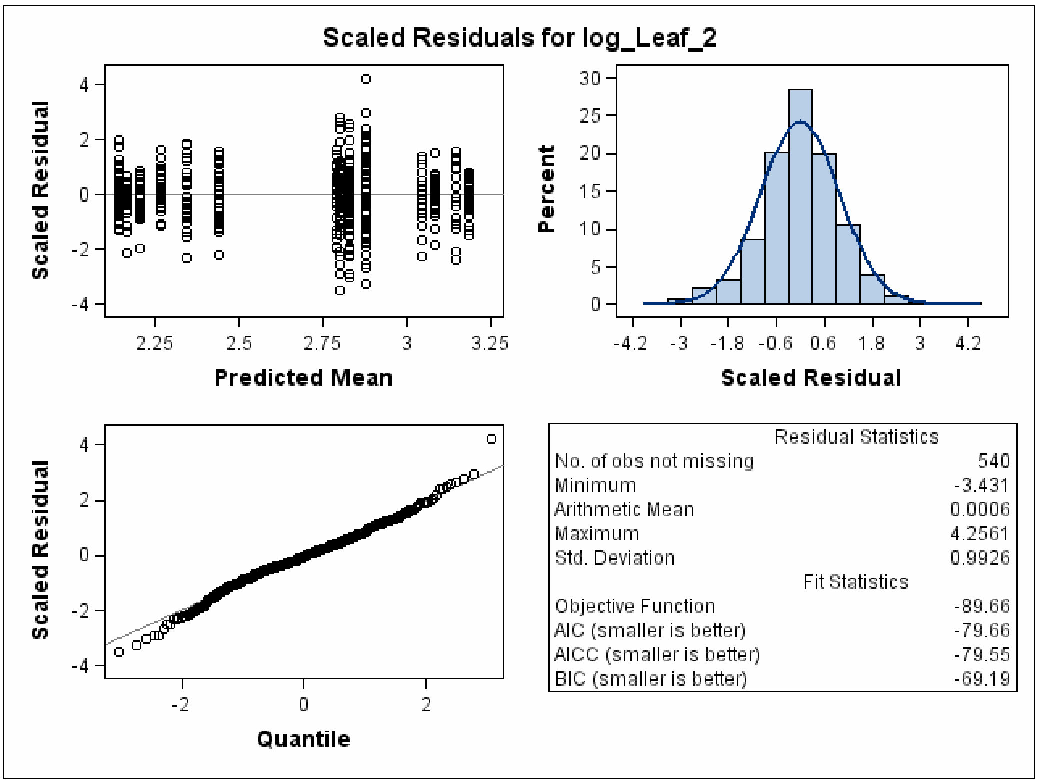

Fig. 5.31: Plots of scaled residuals for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = AR (1) and $\boldsymbol{\Sigma}$ = UN.

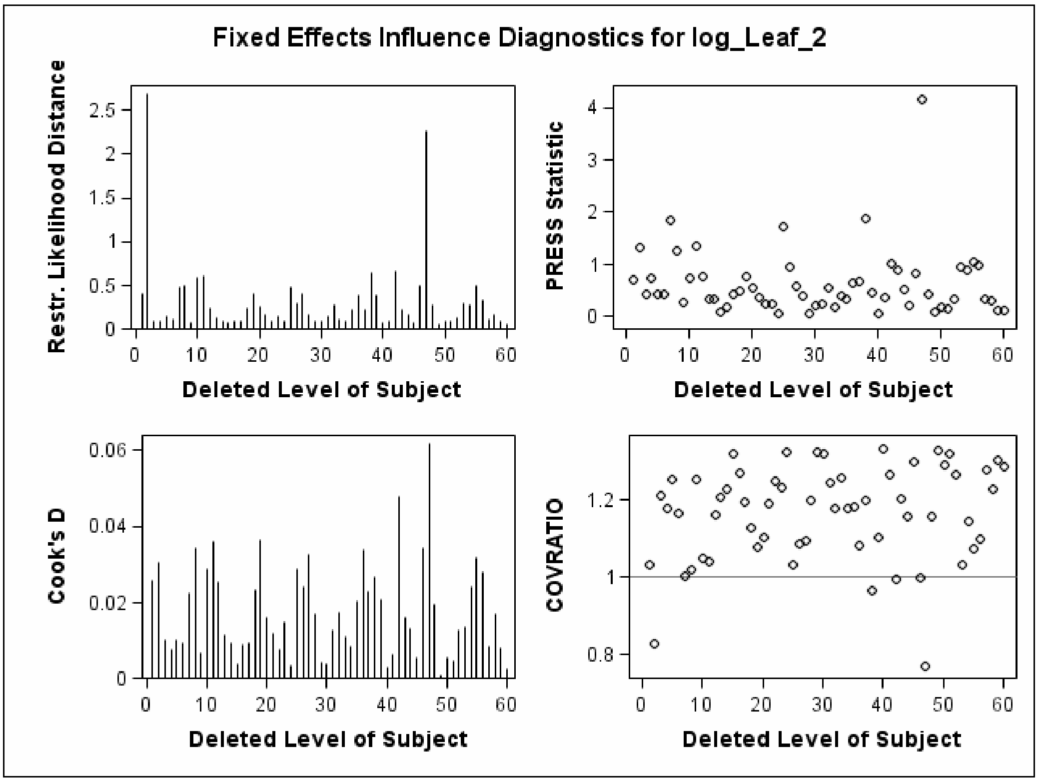

Fig. 5.32: Plots of influence diagnostics for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = AR (1) and $\boldsymbol{\Sigma}$ = UN.

5.4.3.3 Random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN

Of the three models expected to perform well from the simulation study, the random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN obtained the lowest information criteria values. The residual plots for this model (Fig. 5.31) show that the residuals are closely distributed as normal, the mean of the residuals is close to zero, and that the homogeneity of the variance is similar to that of the previous two models. The restricted likelihood distance is smaller compared to the previous two models, and highlights subjects 2 and 47 as potential outliers (Fig. 5.32). Subject 47 is also highlighted by Cook's D, the PRESS Statistic and the COVRATIO (Fig. 5.32).

5.4.3.4 No random effects model with $\boldsymbol{\omega}_i$ = TOEP

The no random effects model with TOEP error structure was another of the models shown to be robust under misspecification. This model obtained only slightly larger values for the information criteria compared to the previous model. The residuals of this model have a mean of exactly zero, and the variance of the residuals is similar to that of the previous models (Fig. 5.33). The residuals are closely distributed as normal (Fig. 5.33). The plots of the influence diagnostics (Fig. 5.34) highlight the same observations mentioned for the previous models, in particular Subject 2, which obtained a restricted likelihood distance much larger in comparison with the other subjects.

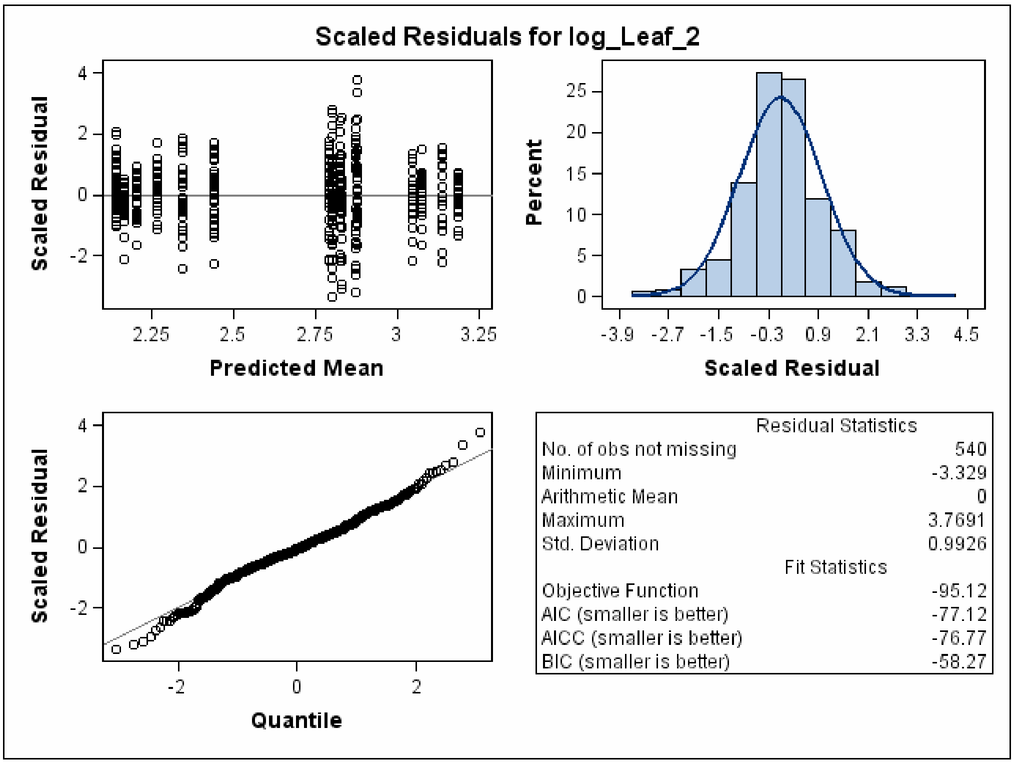

Fig. 5.33: Plots of scaled residuals for the quadratic model with no random effects, and $\boldsymbol{\omega}_i$ = TOEP.

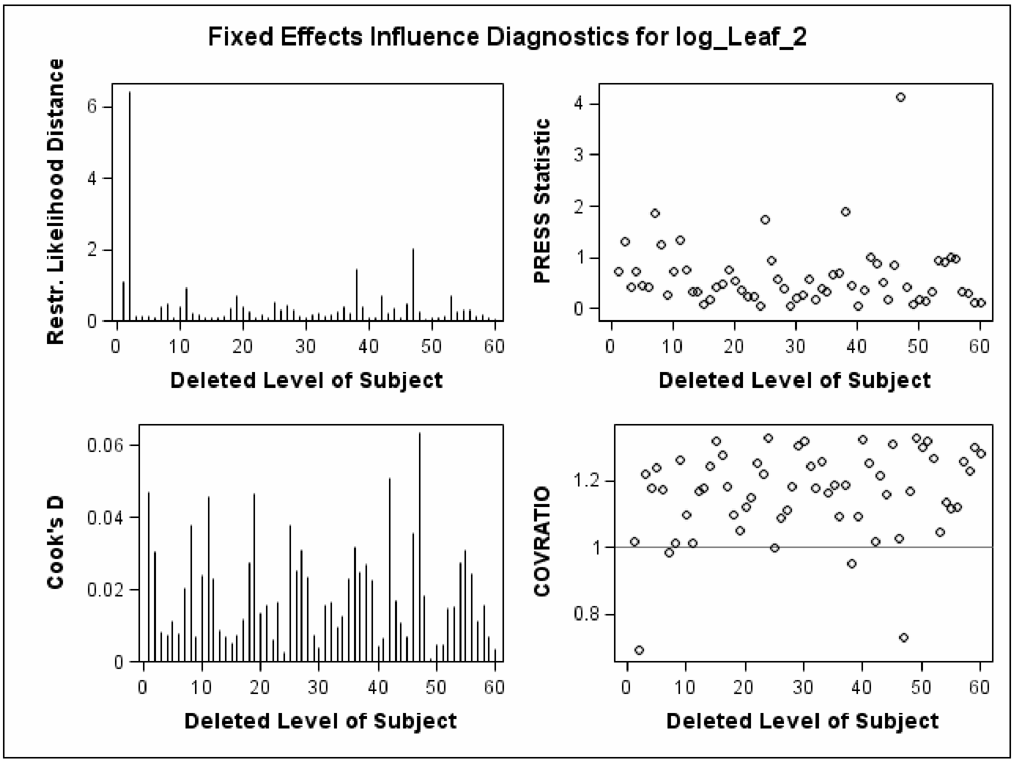

Fig. 5.34: Plots of influence diagnostics for the quadratic model with no random effects, and $\boldsymbol{\omega}_i$ = TOEP.

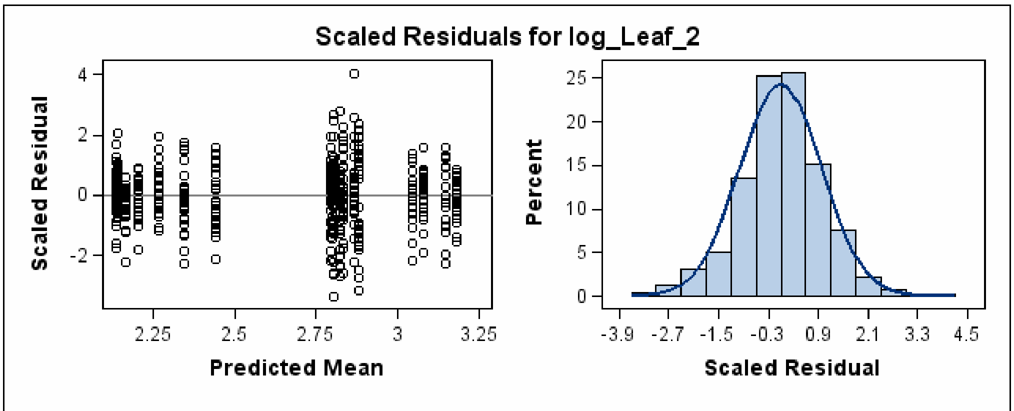

Residual Statistics	
No. of obs not missing	540
Minimum	-3.369
Arithmetic Mean	0.0003
Maximum	4.0664
Std. Deviation	0.9926
Fit Statistics	
Objective Function	-70.64
AIC (smaller is better)	-62.64
AICC (smaller is better)	-62.57
BIC (smaller is better)	-54.26

Fig. 5.35: Plots of scaled residuals for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

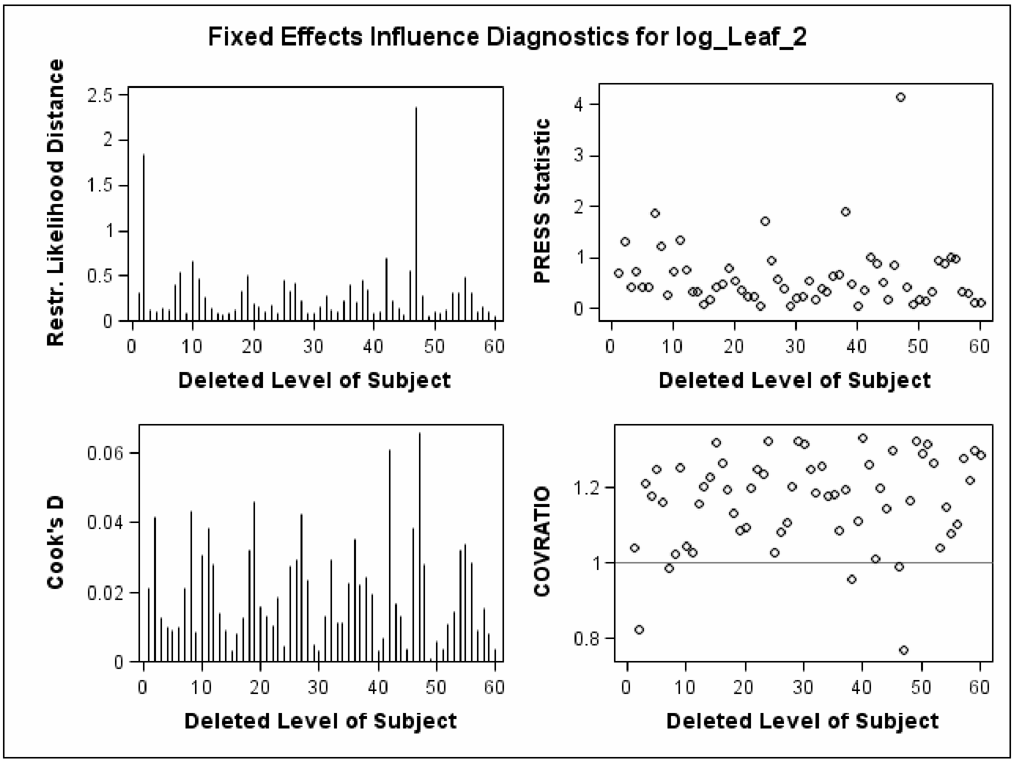

Fig. 5.36: Plots of influence diagnostics for the quadratic model with random intercept and slope, and $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

5.4.3.5 Random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN

The random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN is the third model expected to perform well from the simulation study, and obtained the highest information criteria of the three models. The plots of the scaled residuals (Fig. 5.35) indicate, as for the previous models, that the mean of the residuals is close to zero, the variance of the residuals is relatively homogeneous, and are closely distributed as normal. The influence diagnostics are similar to those obtained for the random intercept and slope model with $\boldsymbol{\omega}_i$ = AR (1) and $\boldsymbol{\Sigma}$ = UN (Fig. 5.36). Subject 47 is identified by all the diagnostics as a potential outlier.

5.4.3.6 Random intercept model with $\boldsymbol{\omega}_i$ *= AR(1)*

In order to include a random intercept model for comparison, the random intercept model with $\boldsymbol{\omega}_i$ = AR(1) was selected. This model performed relatively well in the simulation study, and obtained low values for the information criteria (Table 5.3). The residual plots of the model (Fig. 5.37) are very similar to those obtained for the previous model. The distribution of the residuals is close to normal, the mean of the residuals is very close to zero, and the homogeneity of the residual variance is similar as for previous models. The plots of the influence diagnostics (Fig. 5.38) are also very similar to those of the previous model, once again highlighting subjects 2 and 47.

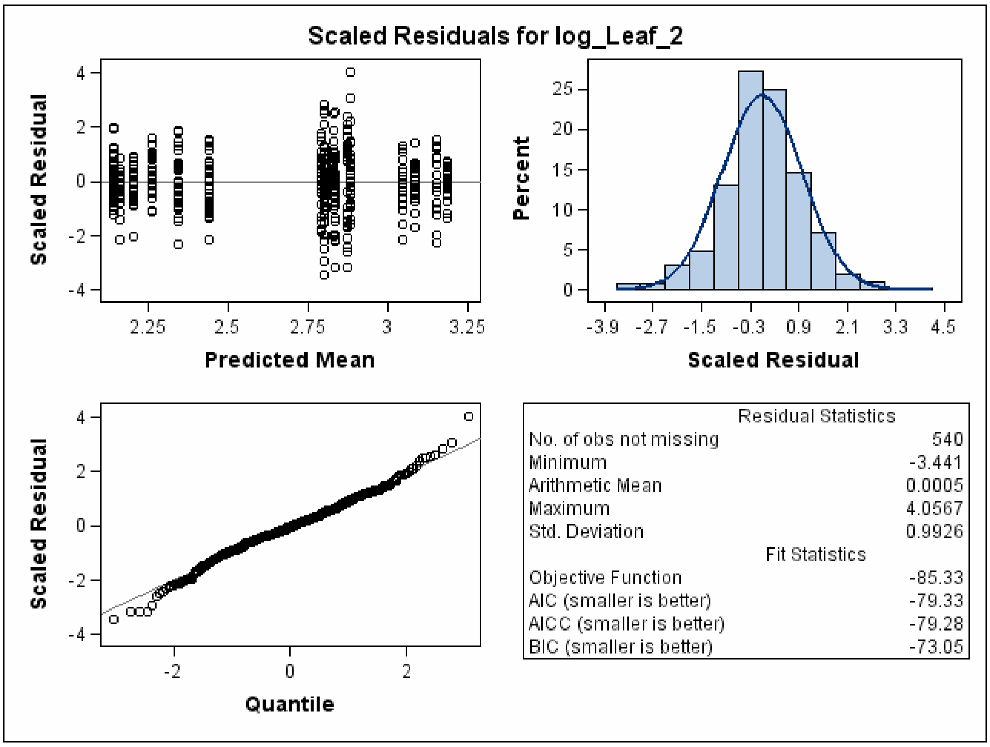

Fig. 5.37: Plots of scaled residuals for the quadratic model with random intercept, and $\boldsymbol{\omega}_i$ = AR(1).

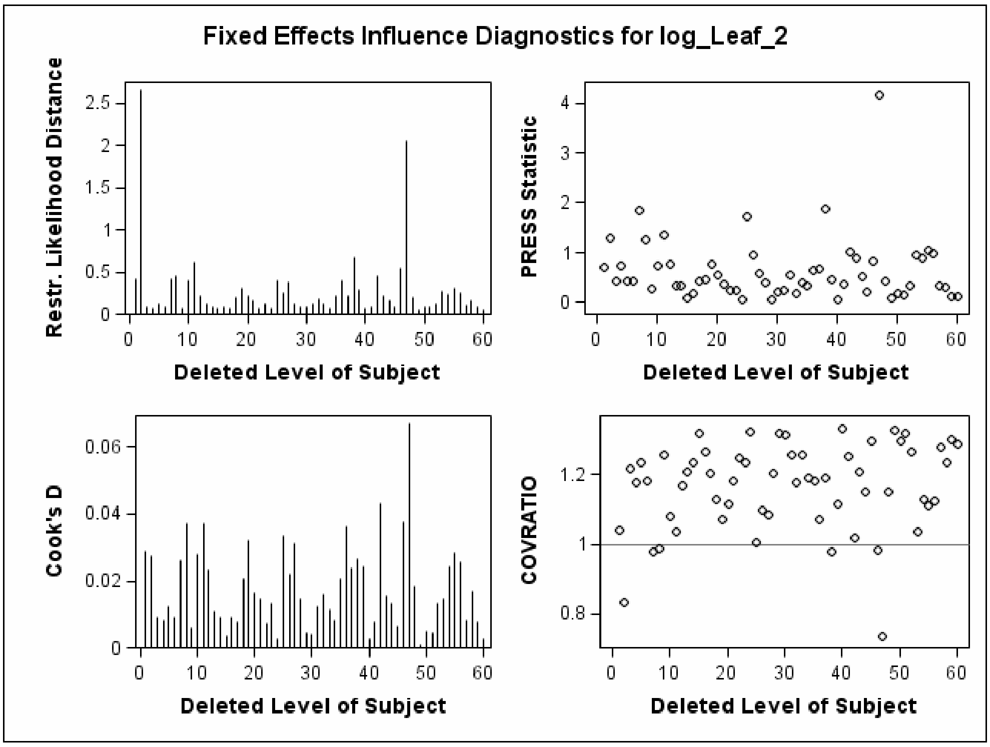

Fig. 5.38: Plots of influence diagnostics for the quadratic model with random intercept, and $\boldsymbol{\omega}_i$ = AR(1).

5.4.3.7 OLS model

The residual diagnostics of the OLS model (Fig. 5.39) indicate that the residual variance is relatively homogeneous, but that there appears to be slight skewness to the right, thereby deviating more from normality than the previous models discussed. The influence diagnostics are similar in size to the previous models (Fig. 5.40). Only subject 47 is clearly highlighted as a potentially influential point.

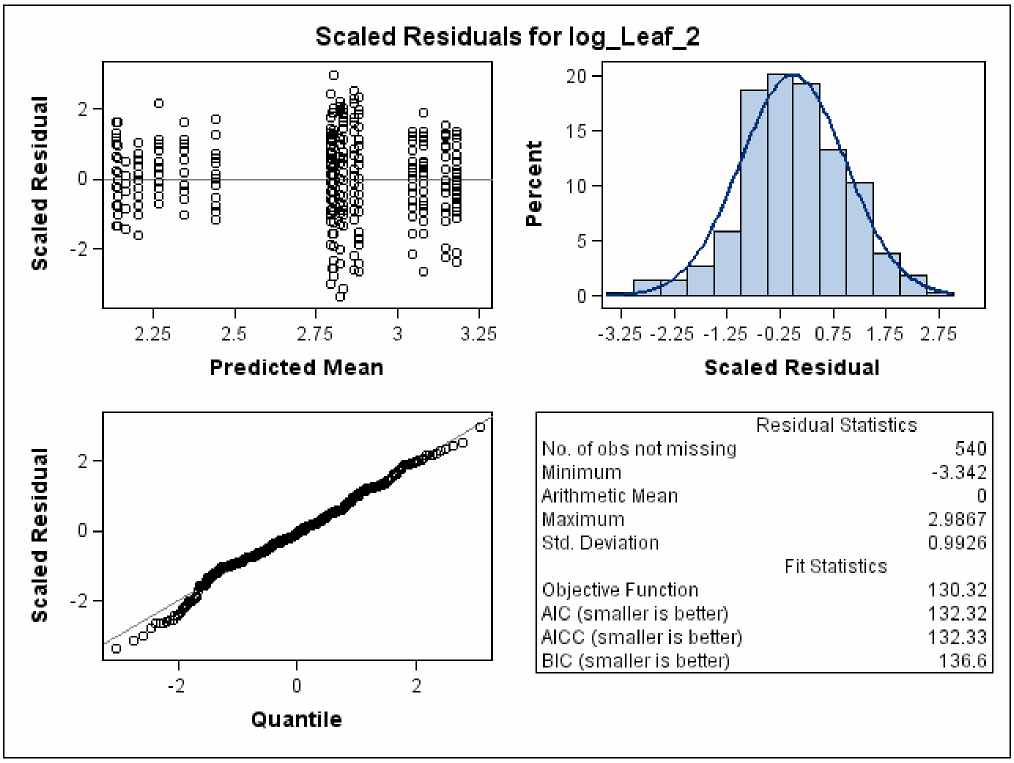

Fig. 5.39: Plots of scaled residuals for the quadratic model with no random effects, and $\boldsymbol{\omega}_i$ = VC (OLS Model).

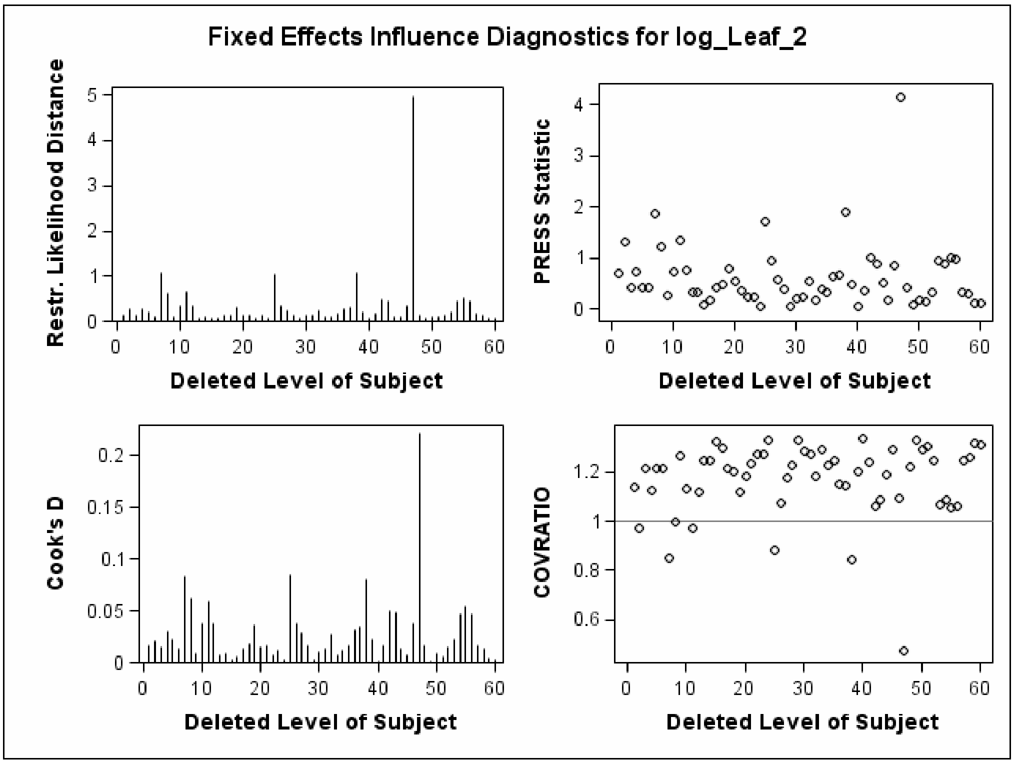

Fig. 5.40: Plots of influence diagnostics for the quadratic model with no random effects, and $\boldsymbol{\omega}_i$ = VC (OLS Model).

5.4.3.8 Summary

In general the residual diagnostics were very similar between all the quadratic models considered. These analyses show an improvement in the homogeneity of the variance of the residuals compared to the simpler linear models. The influence diagnostics tended to highlight the same subjects, in particular, subjects 2 and 47. Although the no random effects model with unstructured error covariance obtained the lowest AIC and AICc values, the residuals are not as close to normality compared to those of other models considered, nor are the influence diagnostics the smallest in size. The random intercept and slope models with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN, with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN, and with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN, and the no random effects model with $\boldsymbol{\omega}_i$ = TOEP obtained very similar residual plots. The size of the influence diagnostics were also similar, except in the case of the restricted likelihood diagnostic of the no random effects model with $\boldsymbol{\omega}_i$ = TOEP, which was relatively large for one observation.

5.4.4 Fit of the covariance structures under quadratic mean model

5.4.4.1 No random effects model with $\boldsymbol{\omega}_i$ *= UN*

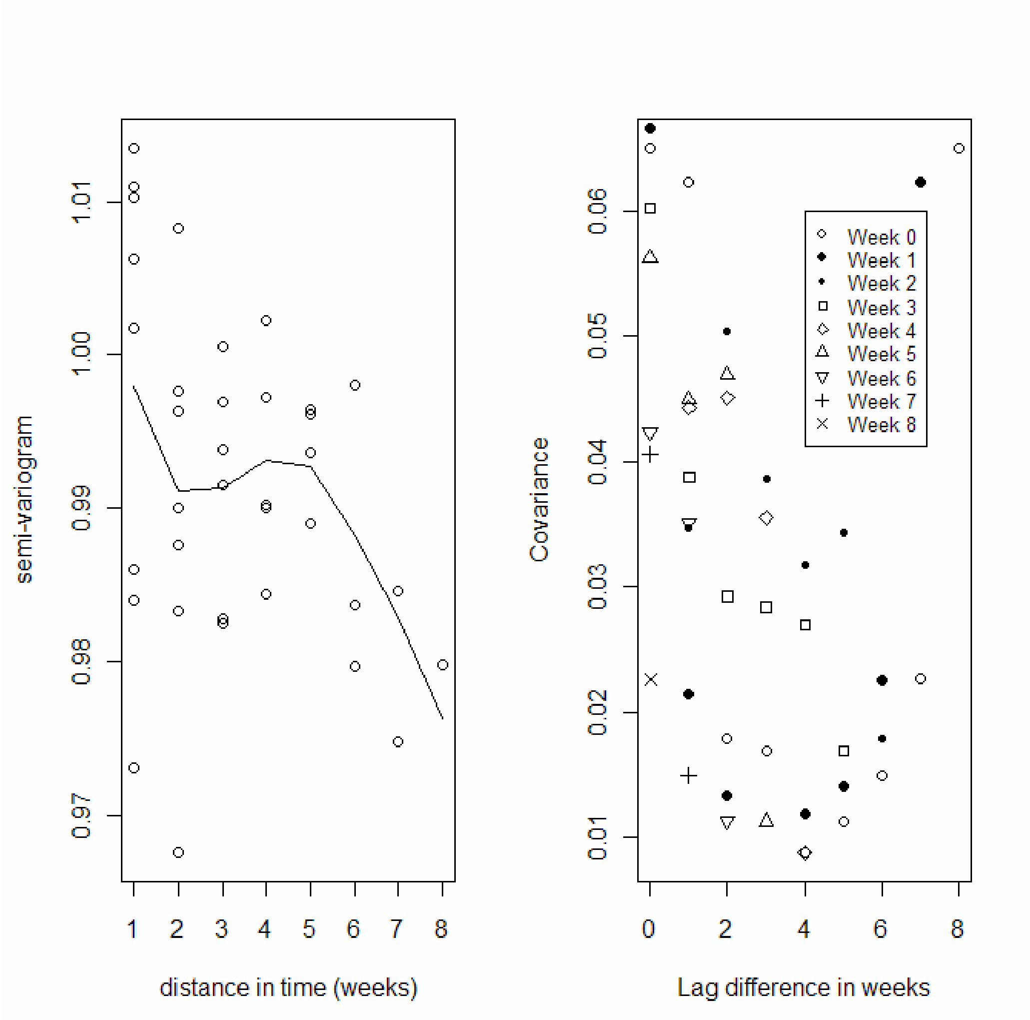

Fig. 5.41: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the no random effects quadratic model with $\boldsymbol{\omega}_i$ = UN.

The plot of the semi-variogram of the best fitting model (in terms of the AIC and AICc), the no random effects model with $\boldsymbol{\omega}_i$ = UN, shows values very close to one, ranging between 0.96 and 1.02 (Fig. 5.41). The plot of the covariances shows that the size of the covariances decreases with increasing lag until a lag 5, where after the covariances again start to increase in size (Fig. 5.41). Therefore the pattern of the covariances over increasing lag forms a U-shape.

5.4.4.2 Random intercept and slope model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN

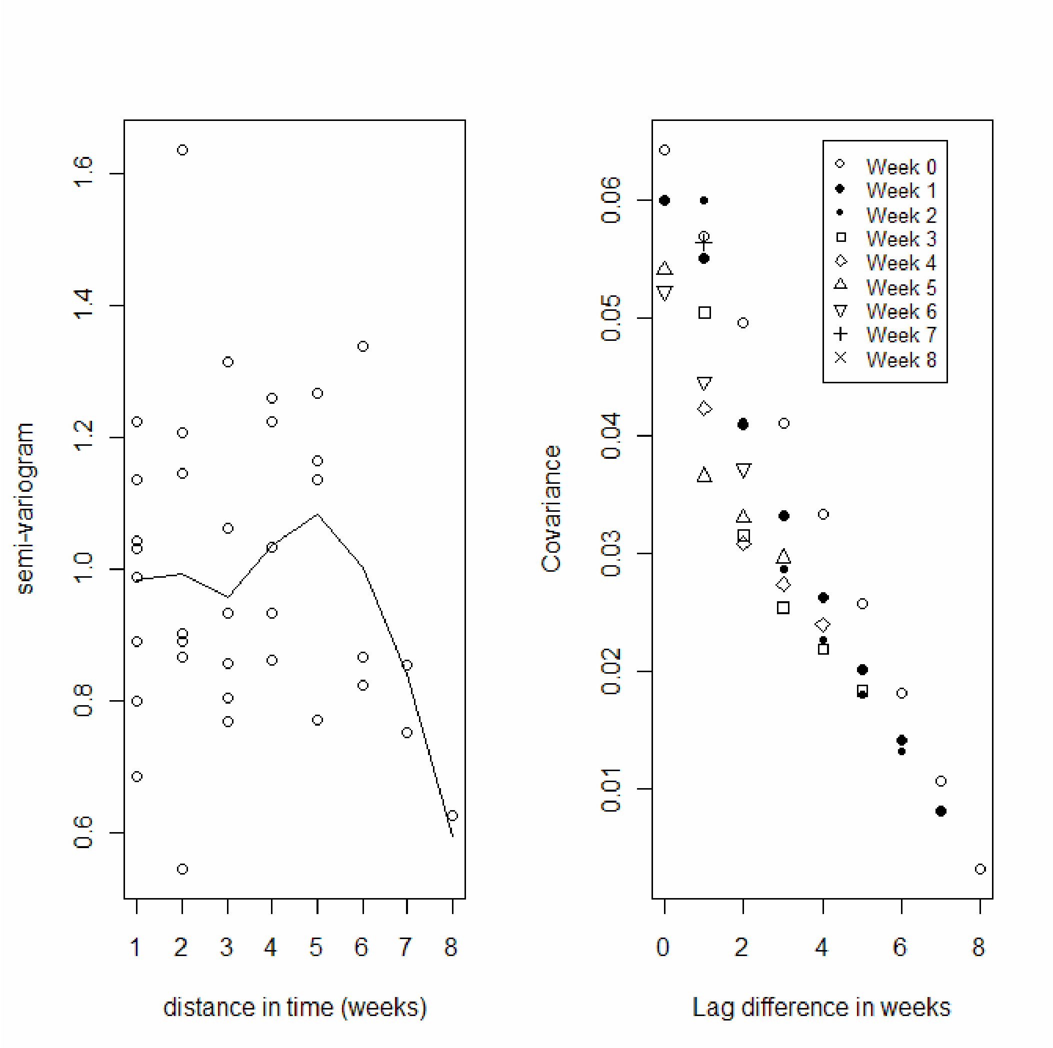

Fig. 5.42: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the random intercept and slope quadratic model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN.

Although the random intercept and slope model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN was the second best fitting model in terms of the AIC and AICc values, and the best fitting model in terms of the BIC value, the covariance plots appear very different to those of the previous model, which obtained the best AIC and AICc values (Fig. 5.42). The range of values for the semi-variogram is much wider, ranging between 0.5 and 1.7. Although the size of the covariances is similar to that of the previous model, the pattern of the covariances is very different, showing a sharply decreasing trend in the size of the covariances as the lag in weeks increases.

5.4.4.3 Random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN

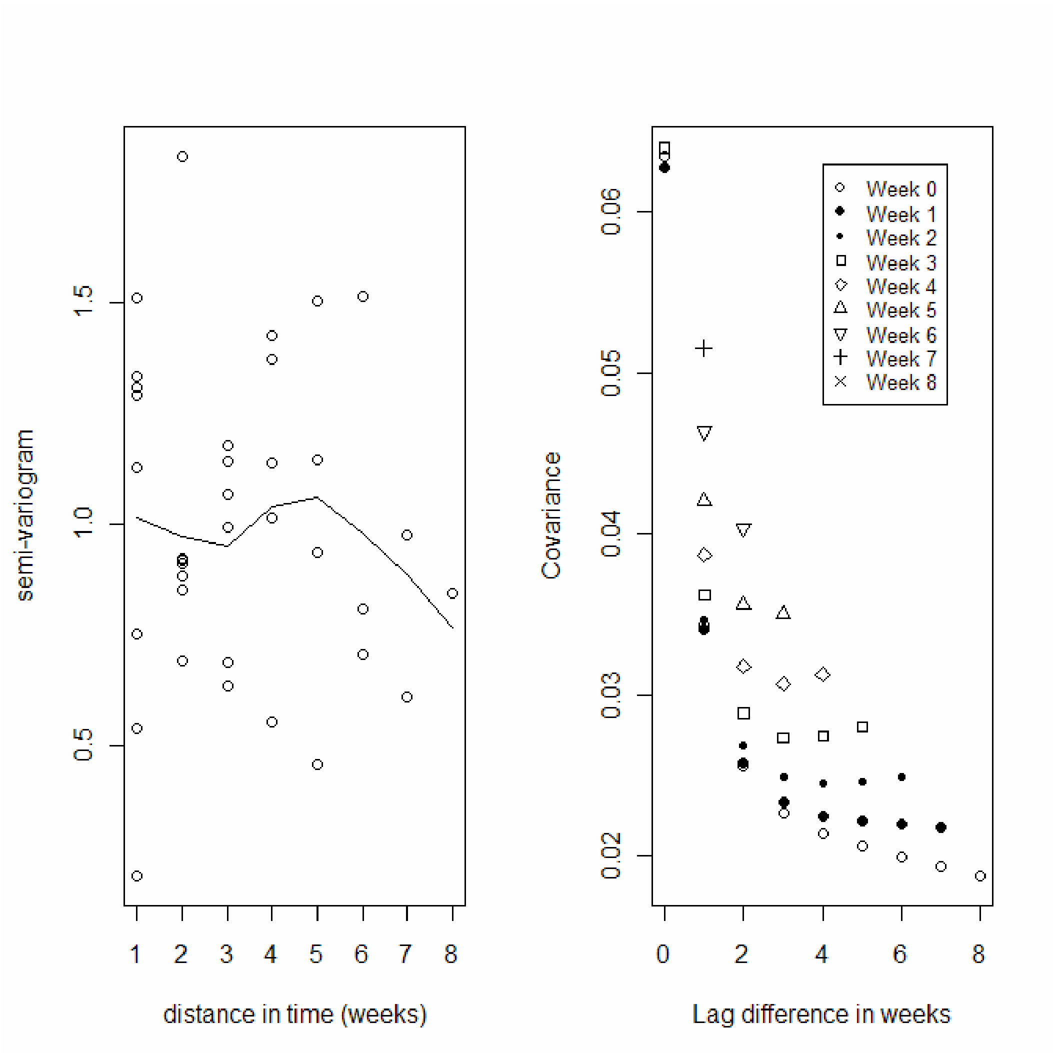

Fig. 5.43: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the random intercept and slope quadratic model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN.

The random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN obtained values for the semi-variogram similar to those for the random effects model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN (Fig. 5.43). The size and pattern of the covariances is also similar (Fig. 5.43).

5.4.4.4 No random effects model with $\boldsymbol{\omega}_i$ = TOEP

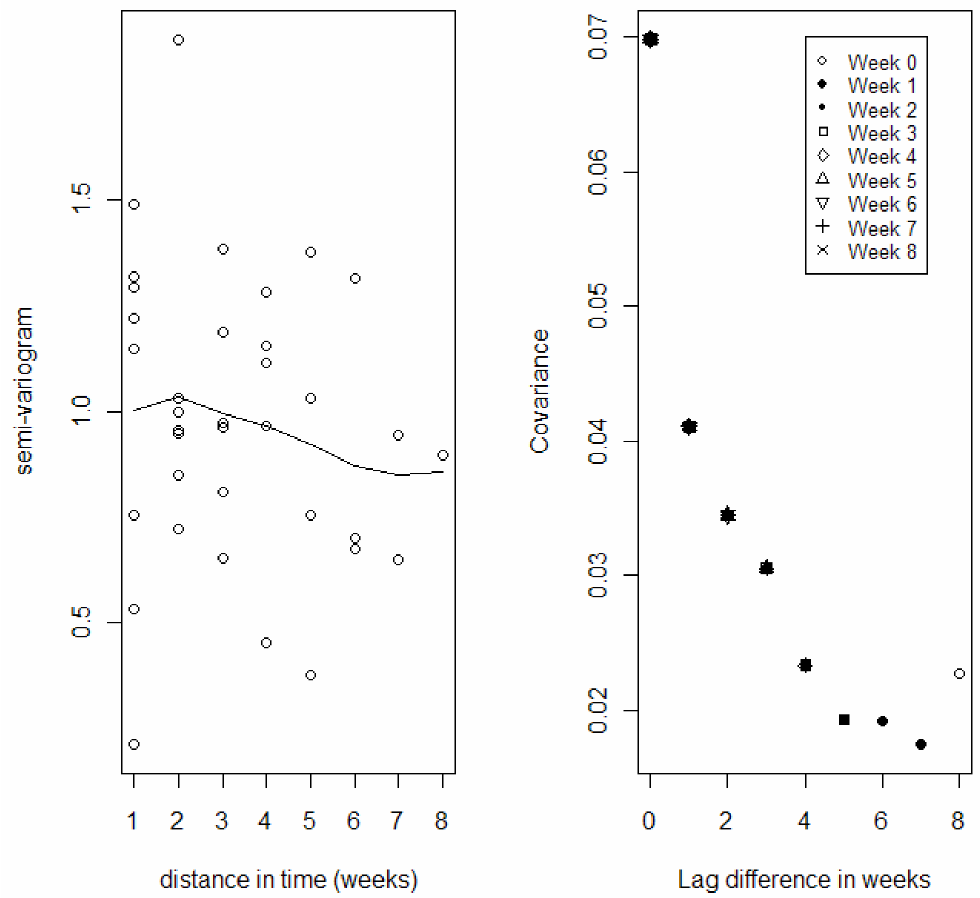

Fig. 5.44: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the no random effects quadratic model with $\boldsymbol{\omega}_i$ = TOEP.

The values of the semi-variogram for the no random effects model with TOEP error structure have an average of approximately one, and range between zero and two. The plots of the covariances show that there is a steep decline in the covariances from lag 0 to lag 5, after which the decline becomes less steep until lag 8 when the covariances increase.

5.4.4.5 Random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN

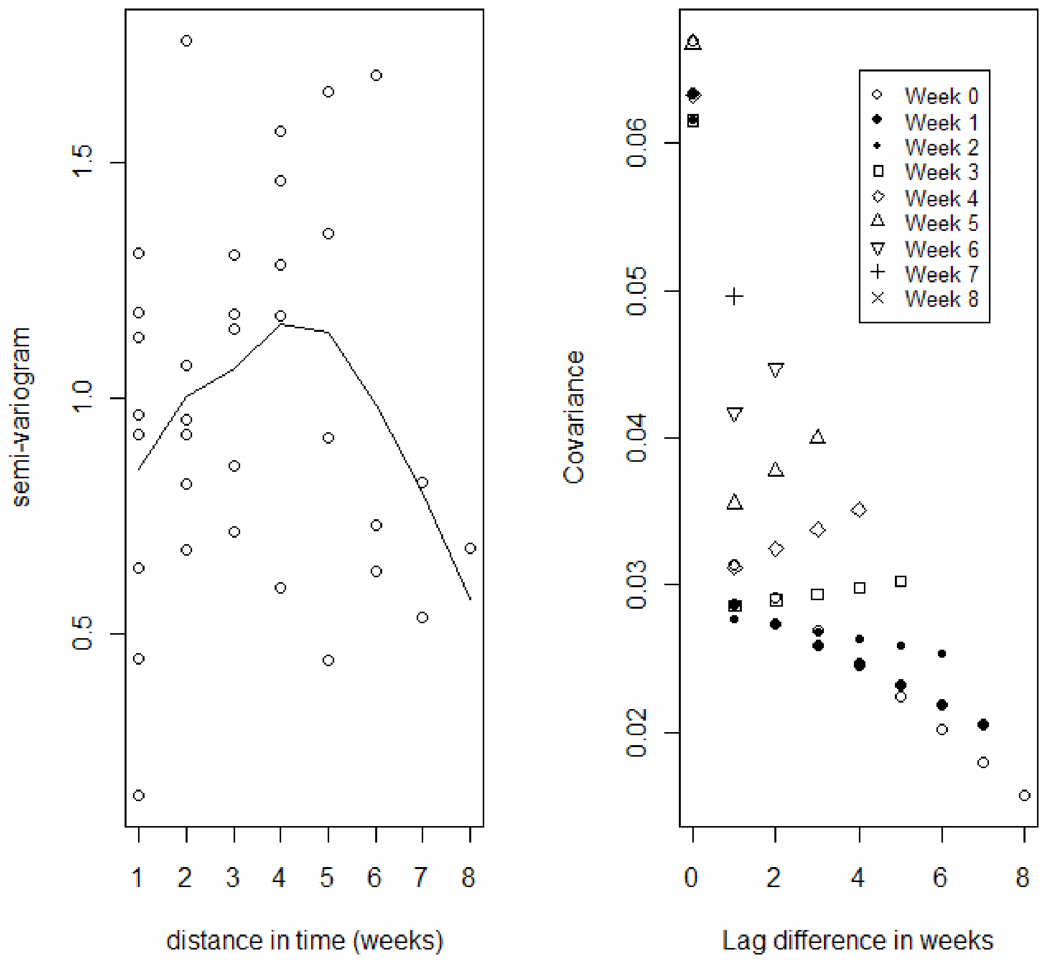

Fig. 5.45: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the random intercept and slope quadratic model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN.

The semi-variogram and the covariance plots of the random intercept and slope model with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN is very similar to that of the random intercept and slope model with $\boldsymbol{\omega}_i$ = AR(1) and $\boldsymbol{\Sigma}$ = UN. The range of values for the semi-variogram is almost identical, and the pattern of the covariances is very similar. In the no random effects cases of these models, i.e. the OLS model and the AR(1) no random effects model, the estimated covariance matrices would have been quite different. Therefore the covariance matrix of the random effects is playing an important role in determining the pattern of the overall model covariances, resulting in two models with different error covariance structures obtaining very similar overall covariance matrices.

5.4.4.6 Random intercept model with $\boldsymbol{\omega}_i$ = AR(1)

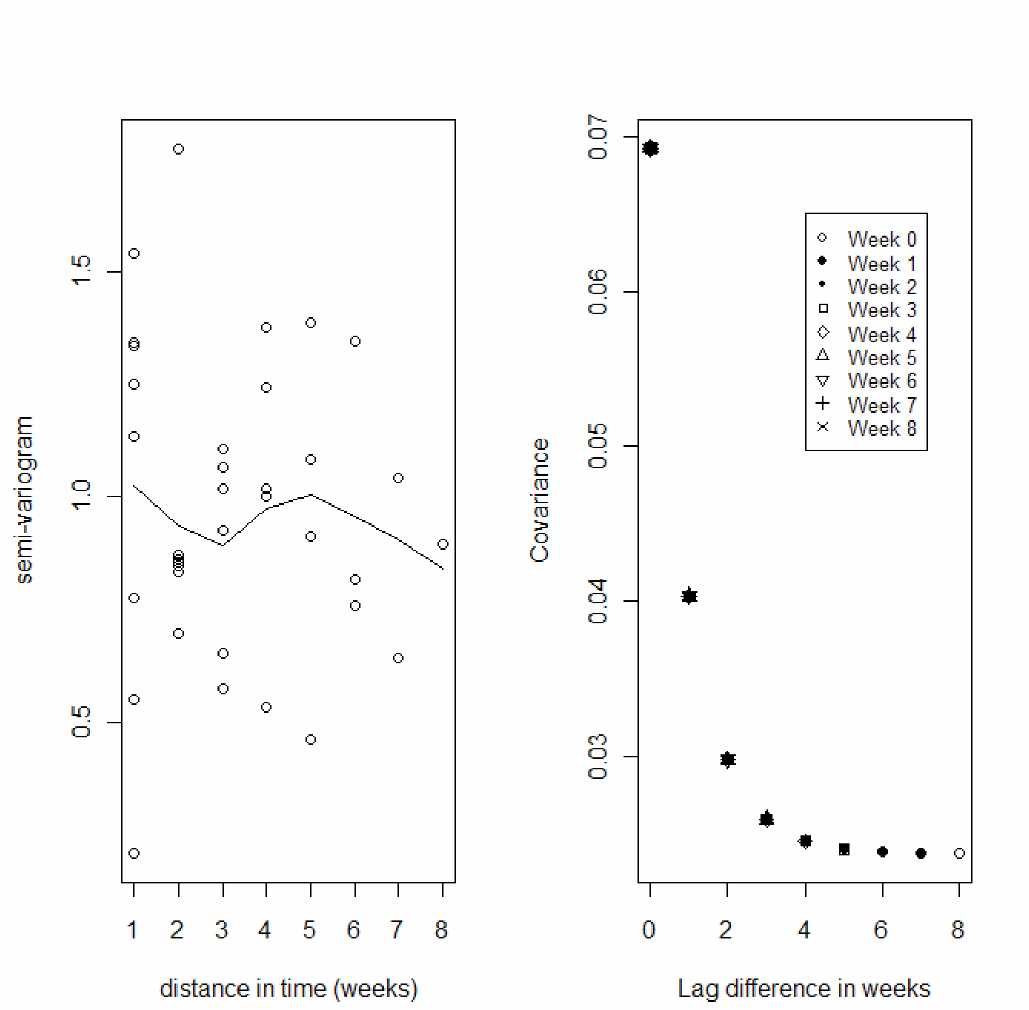

Fig. 5.46: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the random intercept quadratic model with $\boldsymbol{\omega}_i$ = AR(1).

The random intercept model with $\boldsymbol{\omega}_i$ = AR(1) has a semi-variogram with values ranging between zero and two (Fig. 5.46). The covariance plots shows covariances which steeply decrease with increasing lag until lag 4, after which covariances remain constant at a value of approximately two.

5.4.4.7 OLS model

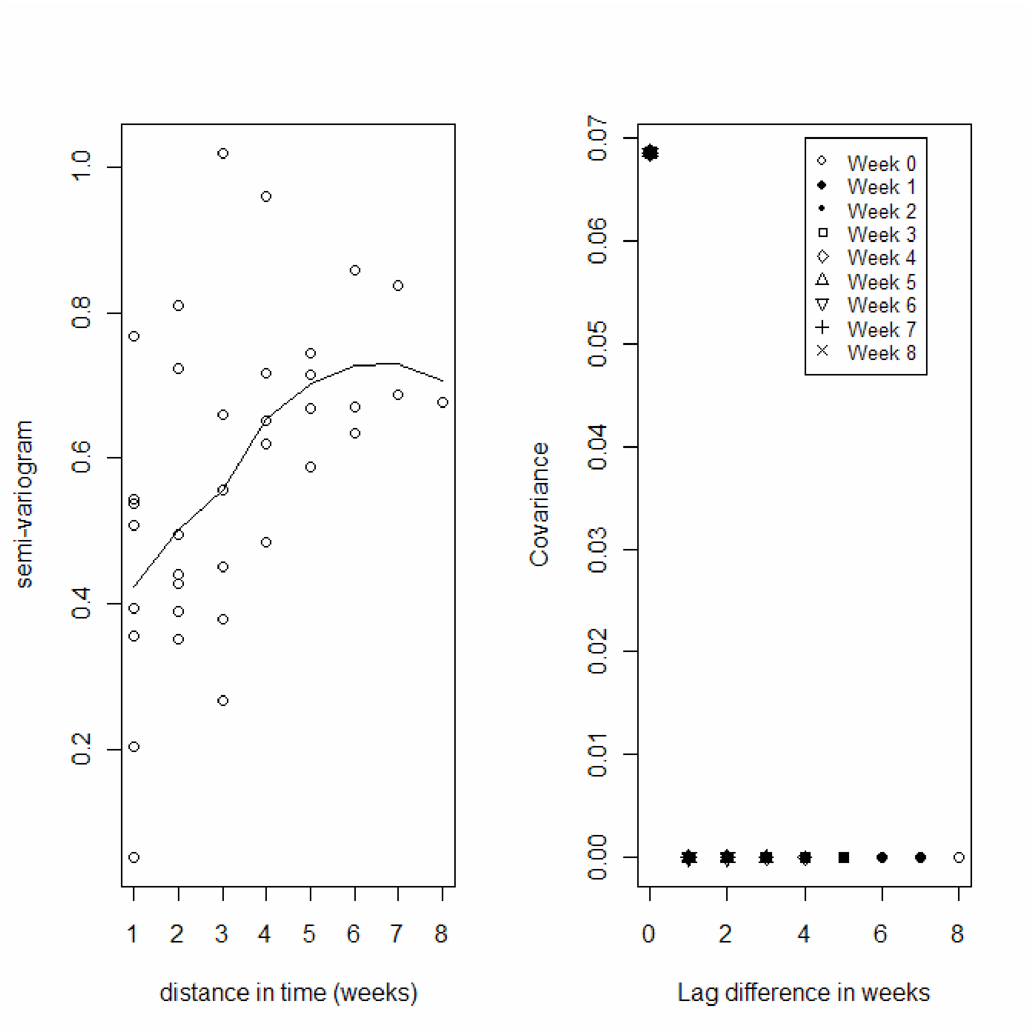

Fig. 5.47: Plot of semi-variogram (left) and the covariances (right) as a function of lag in weeks between observations for the quadratic model with VC errors (OLS model).

The OLS model obtained a semi-variogram with values mostly below one, averaging approximately 0.6. The variance estimate obtained was close to 0.07, similar to the variances obtained by the previous models.

5.4.4.8 Summary

With the exception of the no random effects model with $\boldsymbol{\omega}_i$ = UN and the OLS model, the plots of the semi-variogram and the covariances were very similar between the models considered, i.e. the random intercept and slope models with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN, with $\boldsymbol{\omega}_i$ = AR (1) and $\boldsymbol{\Sigma}$ = UN, and with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN, the random intercept model with $\boldsymbol{\omega}_i$ = AR (1), and the no random effects model with $\boldsymbol{\omega}_i$ = TOEP, obtained very similar overall covariance matrices. The no random effects model with $\boldsymbol{\omega}_i$ = UN obtained the best semi-variogram as all the values were close to one. The pattern of the covariances was very different to other models, which generally predicted a decline the size of the covariances as the lag between weeks increased. Therefore, in this particular example, choosing the model preferred by the AIC, or choosing the model preferred by the BIC would have resulted in fairly different covariance matrices for the responses, thereby affecting the standard errors of the estimated fixed effects. The OLS model obtained the worst values for the semi-variogram as the majority of the values were below one.

5.4.5 Discussion on the quadratic model

The quadratic mean model based on the logged length of the second petiole, which includes adjustment parameters, is an improvement over the linear model based on the actual length of the second petiole. This is revealed by the improvement of the distribution of the residuals, and the increase in homogeneity of the residual variance. The plot of the predicted mean response over time superimposed over the observed mean response over time (Fig. 5.26) clearly shows that the estimated mean model very accurately describes the mean response of the logged observations.

The residual and influence diagnostics were very similar between the different covariance structures considered. The covariance analysis showed that most models had very similar estimated covariance matrices, except in the case of the no random effects model with UN errors. The model under this covariance structure resulted in very large estimates for the covariance at eight weeks lag relative to the estimates obtained by other models. The semi-variogram values of this model were close to one, more so compared to the other models considered.

In this example, if the choice of the best fitting model had been based on the AIC or AICc values, the no random effects model with UN errors would have been the clear choice. But if the choice were based on the BIC, the random effects model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN would have been chosen. Therefore this may be an example where the BIC over penalises a model for having a large number of covariance parameters.

The three covariance structures expected to perform well based on the simulation study (i.e. random intercept and slope models with $\boldsymbol{\omega}_i$ = AR (1) and $\boldsymbol{\Sigma}$ = UN, and with $\boldsymbol{\omega}_i$ = VC and $\boldsymbol{\Sigma}$ = UN, and the no random effects model with $\boldsymbol{\omega}_i$ = TOEP) obtained valid model estimates, and obtained information criteria close that of the random effects model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN, the model selected by the BIC. The residual and influence diagnostics, as well as the covariance analysis, showed that these three models obtained estimates for both the mean model and the covariance model that were very close to the random effects model with $\boldsymbol{\omega}_i$ = ARH(1) and $\boldsymbol{\Sigma}$ = UN.

As for the linear model, the OLS model performed the worst compared to all other models. The information criteria obtained was substantially higher compared to other models, the residuals deviated from normality more compared to other models, and the simple covariance structure obtained for this model did not adequately describe the complexity of the covariance matrix evident from the estimates obtained for other models.

5.5 Comparison between the Simplistic Linear Model and the More Complex Quadratic Model

As concluded in the previous discussion section, the quadratic mean model more adequately describes the relationship of the response over time compared to the simplistic linear model. By modelling the logged length of the second petiole, the normality of the residuals and the homogeneity of the residual variance improved compared to the linear model based on the actual length of the second petiole.

Considering the performance of the three covariance structures predicted to perform well from the simulation study, compared to the simpler linear model, these three covariance structures performed better. Therefore, if a mean structure is selected which adequately describes the mean process, it's possible to fit simpler covariance structures to the data, and successfully model the covariance of the data. In the same breath, the more complex covariance structures under the linear mean model gave more interesting estimates for the covariance matrix of the response compared to the same covariance structures under the quadratic linear model. For example, the no random effects model with TOEP error structure gave estimates for the covariances that resulted in a spike in the size of the covariances from approximately six weeks lag under the linear mean model. The same covariance structure for the quadratic model also resulted in a spike, but comparatively smaller, and only at eight weeks lag. Therefore the covariance estimates are accommodating for the misspecification of the mean model for the linear case, by inflating the covariances for those lags where the error of the mean structure would have been most apparent.

The quadratic model is slightly more difficult to interpret, compared to the simple linear model. The researcher analysing the model output needs insight into the effect that each coefficient has on the shape of the response curve over time in order to be able to accurately assess the estimated response curve, as well as an understanding of the effect that the coefficients of the indicator variable for *Herbicide* would have on the two separate curves for the two levels of this categorical variable. In this example both models predict a separation in the curves of sprayed and unsprayed plants, with the curve of the sprayed plants sitting above the curve of the unsprayed plants. Therefore the conclusion that the growth process of the sprayed and unsprayed plants is different could be concluded from both the simplistic linear model and the more complex quadratic model. On the other hand, the linear model predicts a steady decline in growth of the plants over time, whereas the quadratic model predicts a decelerating decline until approximately week 7 (for the sprayed plants) and week 8 (for the unsprayed plants) when the growth then begins to increase. The description of the growth available from the quadratic model is more accurate, and potentially reveals important characteristics about the growth curves of the sprayed and unsprayed plants not available from the linear model.

The disadvantage of the quadratic model compared to the simplistic linear mean model is the further limitation of the observation to parameter ratio, which is already constrained in the case of the random intercept and slope models. Therefore any additional parameters added to the model will result in greater sensitivity of the model to individual observations.