

Kansas State University Libraries

New Prairie Press

Conference on Applied Statistics in Agriculture

2005 - 17th Annual Conference Proceedings

A BAYESIAN RANDOM COEFFICIENT NONLINEAR REGRESSION MODEL FOR A SPLIT-PLOT EXPERIMENT

Reid D. Landes

Trey Spencer

Ian A. Zelaya

Follow this and additional works at: <https://newprairiepress.org/agstatconference>



Part of the [Agriculture Commons](#), and the [Applied Statistics Commons](#)



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License](#).

Recommended Citation

Landes, Reid D.; Spencer, Trey; and Zelaya, Ian A. (2005). "A BAYESIAN RANDOM COEFFICIENT NONLINEAR REGRESSION MODEL FOR A SPLIT-PLOT EXPERIMENT," *Conference on Applied Statistics in Agriculture*. <https://doi.org/10.4148/2475-7772.1140>

This is brought to you for free and open access by the Conferences at New Prairie Press. It has been accepted for inclusion in Conference on Applied Statistics in Agriculture by an authorized administrator of New Prairie Press. For more information, please contact cads@k-state.edu.

A BAYESIAN RANDOM COEFFICIENT NONLINEAR REGRESSION MODEL FOR A SPLIT-PLOT EXPERIMENT

Reid D. Landes[†], Trey Spencer[†], and Ian A. Zelaya[‡]

[†] Department of Biostatistics, University of Arkansas for Medical Sciences

[‡] Department of Agronomy, Iowa State University

Abstract

In random coefficients regression, we are often interested in the mean of a certain parameter particular to the experimental unit (EU). When the mean depends on some treatment regimen, we are then interested in comparing the means among the different treatments. When the EUs are repeatedly measured on a variable containing information about the EU parameter, a standard procedure is to estimate each EU parameter and treat the estimates as the response variables. This is especially true when the regression model for an EU is nonlinear. Often, for designed experiments with a factorial treatment structure, the estimated EU parameters are then modeled with an appropriate linear (mixed) model. Here, we consider a split-plot experiment conducted to detect differences in the half-life of a compound between different treatment regimens of the compound, namely compound preparation and temperature (whole-plot factors) and initial compound amount (split-plot factor). Initially, we provide a standard (classical) analysis plan, and then present a Bayes random coefficients regression model to address the researcher's questions of interest. We finally compare the results from the standard and Bayes analyses.

Keywords: *hierarchical model, MCMC*

1 Introduction

For certain types of planned longitudinal experiments, we cannot directly observe the desired characteristic, say g , of an experimental unit (EU). However, longitudinal data particular to the EU may be modeled with a regression function of g , time, and possibly other EU-specific parameters. We may then estimate g for each EU. To compare means of g among different treatments, we may proceed in one of two ways:

1. fit the estimated g s with a linear (mixed) model and compare different treatment means with appropriate contrasts, or
2. fit the longitudinal data with a random coefficients regression (RCR) model and compare the treatment effects with appropriate contrasts of coefficients.

In this paper, we compare a Bayesian version of the latter to a classical version of the former for a longitudinal study conducted in a split-plot experiment with subsampling.

We organize the paper starting with a description of the experiment in Section 2. The model for the longitudinal data along with the analysis plans following from a classical linear mixed model and a Bayes RCR model are presented in Section 3. We present and compare the results from the two analyses in Section 4, and discuss strengths and weaknesses of the two methods in the final section.

2 Background and Description of the Experiment

Shikimic acid (SA) is often used in plant sciences as an indicator of efficacy of herbicides containing glyphosate (e.g. Roundup[®] and Touchdown[®]). For most plants, endogenous SA accumulates in response to glyphosate applications onto the plants. However, glyphosate-resistant plants, such as Roundup Ready[®] crops and certain weeds (e.g. Conya, Amaranthus, Lolium, etc.), do not accumulate SA when treated with field rates of glyphosate. To measure the amount of SA in a sample of plant tissue, SA must first be separated from the tissue. When the tissue is left undisturbed, SA remains stable, but once SA is extracted from the tissue, SA starts to degrade. It is important for laboratory technicians to know how quickly SA degrades. If it decays slowly, then the technician may choose to process a large number of samples at the same time. If it degrades quickly, then the technician should work with smaller numbers of samples at a time.

Typically, when preparing a tissue sample for SA determination via spectrophotometry, a technician will use one of two compound preparation solutions, say P1 or P2. In addition, the samples are typically prepared in either cold ($\sim 4^\circ\text{C}$) or room temperature ($\sim 23^\circ\text{C}$) conditions. We will call a preparation at a given temperature a treatment. Treatments 1 through 4 will correspond with P1 at 4° , P1 at 23° , P2 at 4° , and P2 at 23° , respectively. Wanting to know which treatment, if any, prolongs SA half-life the most, the researcher decided to run an experiment that had the four treatments at two different initial amounts of SA, say 1 and 2 units.

This 4×2 factorial treatment structure was carried out in a split-plot experiment with subsampling. Treatments were the whole-plots and were replicated by two runs of the experiment. The two levels of initial amount were randomly assigned to halves of a whole-plot, thus constituting the split-plots. From each split-plot, five tubes of the treatment-amount mixture were prepared. From each tube, the optical density (OD) of SA was recorded at times 0, 1, ..., 14, 28, and 50 hours. (See Zelaya, Owen, and Landes (manuscript) for more details.)

We note that inferences on the split-plot treatment effects and interactions with them are irrelevant since half-life is not dependent on amount. The different amounts were included in this experiment to simulate both high and low levels of SA that would be expected to occur in samples with unknown amounts of SA.

3 Statistical Methods

We now describe how the experimental data can be analyzed by fitting a classical linear mixed model and a Bayesian RCR model. Common to both approaches, we let y_{hijkl} be observed OD at time l on tube k with initial amount j prepared in run i of treatment h . With x_l being the time at observation l and

$$EY_{hijkl} = a_{hijk} + b_{hijk} \exp \left\{ - \frac{(\ln 2) x_l}{g_{hijk}} \right\}, \quad (1)$$

we model

$$y_{hijkl} = EY_{hijkl} + \varepsilon_{hijkl}, \quad (2)$$

where

$$\varepsilon_{hijkl} \sim^{iid} N(0, \tau_\varepsilon^{-1}) \quad (3)$$

(Note that τ_ε is the precision of ε_{hijk} .) From (1) we see that $a_{hijk} \in \mathfrak{R}$ is the lower asymptote, $b_{hijk} \in \mathfrak{R}^+$ is the range between the lower asymptote and initial OD reading, and $g_{hijk} \in \mathfrak{R}^+$ is

the half-life for a tube.

3.1 A classical approach

In this experiment, there was no missing data, and all tubes were measured at the same time intervals. Hence, we use the idea of derived variables and fit these with a linear mixed model. (See Section 6.3 of Diggle, Liang, and Zeger (1994).)

According to (1)-(3), we regress \mathbf{y}_{hijk} , the longitudinal OD data from a tube, on time $\mathbf{x} = (0, 1, \dots, 14, 28, 50)$. (We did this with the NLIN procedure in SAS[®].) We use the resulting iteratively reweighted least squares estimate of g_{hijk} , denoted \hat{g}_{hijk} , as the (derived) response variable for tube k with initial amount j in run i of treatment h .

Following from the design of the experiment, we model

$$\hat{g}_{hijk} = \text{TRT}_h + \text{RUN}_i + \omega_{hi} + \text{AMT}_j + \text{TRT*AMT}_{ij} + \epsilon_{hij} + \delta_{hijk}, \quad (4)$$

where

TRT_h , RUN_i , AMT_j , and TRT*AMT_{hj} are, respectively, the fixed effects for treatment, run of the experiment, initial amount, and the interaction between treatment and initial amount,

ω_{hi} is the whole-plot error of TRT h in RUN i ,

ϵ_{hij} is the split-plot error of TRT*AMT (h, j) in RUN i , and

δ_{hijk} is the sampling error of TUBE k from TRT*AMT (h, j) in RUN i .

The ω s, ϵ s, and δ s are taken to be mutually independent and distributed normal with 0 means and respective variances σ_ω^2 , σ_ϵ^2 , and σ_δ^2 . The ANOVA for model (4) is as follows.

| Source | DF | F to test |
|-------------------|----|-----------------------------|
| TRT | 3 | MS(TRT)/MS(TRT*RUN) |
| RUN | 1 | – |
| TRT*RUN | 3 | – |
| AMT | 1 | MS(AMT)/MS(TRT*RUN*AMT) |
| TRT*AMT | 3 | MS(TRT*AMT)/MS(TRT*RUN*AMT) |
| TRT*RUN*AMT | 4 | – |
| TUBE(TRT RUN AMT) | 64 | – |
| Total | 79 | |

The effects involving amount (AMT) are not of direct interest, since we assume that the amount of SA has no effect on SA half-life. However, we leave these effects in the model for two reasons: (1) they handle potential imbalance of split-plot and sampling units in such experiments; and (2) they serve as a check of the half-life assumption.

The comparisons of interest are among the means of the whole-plot treatments. We are particularly interested in identifying a “best” treatment or set of treatments. Using contrasts of the four treatment means, we may test the main effects of preparation and temperature and the preparation by temperature interaction. Alternatively, we may use mean separation techniques to potentially identify a best treatment. We use Tukey’s adjustment method for all pairwise comparisons of the four treatment means.

3.2 Bayesian Approach

Displays (1)-(3) describe a model for the data from a sampling unit, which is nested within a split-plot that is nested within a whole-plot. Considering the sampling unit parameters a_{hijk} , b_{hijk} , and g_{hijk} in this context, a random coefficients regression (RCR) model for all the OD data follows naturally. RCRs are hierarchical models and may be approached with a classical or Bayesian philosophy. Here, we describe a Bayesian hierarchical model and how inferences on the four treatment means are conducted from it.

3.2.1 The Bayesian hierarchical model

The first level of the hierarchy (**the OD model**) is described in (1) – (3). In the second level – **the tube model**, we assign τ_ε in (3) the prior $\text{Gam}(0.001, 0.001)$. In view of Equation (2), this prior is conjugate. The hyperparameters taking value 0.001 spread the mass of the prior over a very large range, which, in practicality, makes the prior noninformative. We model the tube parameters a_{hijk} , b_{hijk} , and g_{hijk} as

$$\begin{aligned} a_{hijk} &\sim N(A_{hij}, \tau_a^{-1}), \\ b_{hijk} &\sim \text{Gam}(B_{hij}^2 \tau_b, B_{hij} \tau_b), \text{ and} \\ g_{hijk} &\sim \text{Gam}(G_{hij}^2 \tau_g, G_{hij} \tau_g), \end{aligned}$$

with a_{hijk} , b_{hijk} , and g_{hijk} being mutually independent. Parameters A_{hij} , B_{hij} , and G_{hij} are specific to split-plot (h, i, j) . We conveniently chose the normal and gamma distributions here since their domains coincide with the ranges of the respective tube parameters. Parameterizing the gamma distribution as $\text{Gam}(m^2 p, mp)$, where m and p are the mean and precision, respectively, we can model b_{hijk} and g_{hijk} as having constant precisions (i.e. constant variances) and means that are specific to split-plot (h, i, j) . Constant precision and split-plot-specific means are also assumed in the normal prior for a_{hijk} . We use this rationale for choosing models and priors in subsequent levels of the hierarchy.

The model for the split-plot parameters, A_{hij} , B_{hij} , and G_{hij} , and the priors on precisions τ_a , τ_b , and τ_g comprise the third level – **the split-plot model**. We assign $\tau_a \sim \text{Gam}(0.001, 0.001)$, $\tau_b \sim \text{Gam}(0.03, 0.03)$, and $\tau_g \sim \text{Gam}(0.03, 0.03)$, choosing the hyperparameters so that the priors are relatively noninformative. We model

$$\begin{aligned} A_{hij} &\sim N(\alpha_{hi}, \tau_A^{-1}), \\ B_{hij} &\sim \text{Gam}((\beta_{hi} \text{AMT}_j)^2 \tau_B, \beta_{hi} \text{AMT}_j \tau_B), \text{ and} \\ G_{hij} &\sim \text{Gam}(\gamma_{hi}^2 \tau_G, \gamma_{hi} \tau_G), \end{aligned}$$

with A_{hij} , B_{hij} , and G_{hij} being mutually independent. Note that B_{hij} is the range between OD at time zero and A_{hij} for split-plot (h, i, j) . OD at time zero is a linear function of initial amount (AMT_j for $j = 1, 2$), and the mean of B_{hij} , $\beta_{hi} \text{AMT}_j$, accommodates this. Parameters α_{hi} , β_{hi} , and γ_{hi} are specific to whole-plot (h, i) .

In the fourth level of the model – **the whole-plot model** – we assign $\tau_A \sim \text{Gam}(0.001, 0.001)$, $\tau_B \sim \text{Gam}(0.0075, 0.0075)$, and $\tau_G \sim \text{Gam}(0.0075, 0.0075)$, which are relatively noninformative. We model

$$\alpha_{hi} \sim N(A_h, \tau_\alpha^{-1}),$$

$$\begin{aligned}\beta_{hi} &\sim \text{Gam}(\mathbf{B}_h^2\tau_\beta, \mathbf{B}_h\tau_\beta), \text{ and} \\ \gamma_{hi} &\sim \text{Gam}(\Gamma_h^2\tau_\gamma, \Gamma_h\tau_\gamma),\end{aligned}$$

with α_{hi} , β_{hi} , and γ_{hi} being mutually independent. \mathbf{A}_h , \mathbf{B}_h , and Γ_h are parameters specific to treatment h .

The treatment model is the final level of the hierarchical model and is made up of the following relatively noninformative priors: $\mathbf{A}_h \sim \text{N}(0, 0.001^{-1})$, $\mathbf{B}_h \sim \text{Gam}(0.001, 0.001)$, $\Gamma_h \sim \text{Gam}(0.001, 0.001)$, $\tau_\alpha \sim \text{Gam}(0.005, 0.005)$, $\tau_\beta \sim \text{Gam}(0.005, 0.005)$, and $\tau_\gamma \sim \text{Gam}(0.005, 0.005)$.

We note that the researcher had no historical data from which we could inform our prior distributions. We thus strove to choose hyperparameters that would spread the mass of the prior over a very large range. For priors of the precision parameters (e.g. τ_g), this often led to numerical difficulties when fitting the hierarchical model; hence, we chose these priors to be as diffuse as possible as allowed by the software (described below).

3.2.2 Inferences from the hierarchical model

In the Bayesian hierarchical model presented above, the quantities of interest are the individual elements of $\mathbf{\Gamma} = (\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4)$ and functions of them. The particular functions of $\mathbf{\Gamma}$ in which we are interested are

$$\begin{aligned}\phi_{INT} &\equiv (\Gamma_1 - \Gamma_2 - \Gamma_3 + \Gamma_4), \\ \phi_{PREP} &\equiv 0.5(\Gamma_1 + \Gamma_2 - \Gamma_3 - \Gamma_4), \\ \phi_{TMP} &\equiv 0.5(\Gamma_1 - \Gamma_2 + \Gamma_3 - \Gamma_4), \\ \phi_{h-h'} &\equiv \Gamma_h - \Gamma_{h'} \text{ for } 1 \leq h < h' \leq 4, \text{ and} \\ \phi_{ORD} &\equiv \text{order}(\mathbf{\Gamma}).\end{aligned}$$

Here, ϕ_{INT} is the interaction effect of preparation and temperature, ϕ_{PREP} is the main effect of preparation, ϕ_{TMP} is the main effect of temperature, $\phi_{h-h'}$ is the mean difference between treatments h and h' , and ϕ_{ORD} is the ordered vector of $\mathbf{\Gamma}$ from least to greatest.

Integration over the nuisance parameters to obtain marginal posterior densities is intractable for $p(\Gamma_h|\mathbf{y})$ with $h = 1, 2, 3, 4$ and $p(\phi(\mathbf{\Gamma})|\mathbf{y})$ with $\phi(\mathbf{\Gamma})$ being a function defined above. However, empirical distributions of these marginal posterior densities are readily obtained using Markov Chain Monte Carlo (MCMC) simulation in the hierarchical model. MCMC simulation may be conducted by employing Metropolis-Hasting algorithms in successive substitution sampling of $p(\theta_{(m)}|\boldsymbol{\theta}_{(\sim m)}, \mathbf{y})$, where $\theta_{(m)}$ is one of the parameters in $(\mathbf{a}, \mathbf{b}, \mathbf{g}, \mathbf{A}, \mathbf{B}, \mathbf{G}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{A}, \mathbf{B}, \mathbf{\Gamma}, \boldsymbol{\tau})$ and $\boldsymbol{\theta}_{(\sim m)}$ are all of the remaining parameters. (See Chapter 11 of Gelman, Carlin, Stern, and Rubin (1995) and Chapter 5 of Carlin and Louis (2000).) A large collection of iterations of the quantity of interest from the MCMC simulation serves as the empirical distribution of that quantity given the data \mathbf{y} .

From the univariate empirical distributions we may obtain point and interval estimates for the quantities of interest. We use medians for point estimates and quantiles $\lambda/2$ and $1 - \lambda/2$ for $(1 - \lambda)100\%$ credible intervals. When constructing simultaneous $(1 - \lambda)100\%$ credible intervals for all pairwise differences, we use Bonferroni's adjustment. Specifically, we construct individual $(1 - \lambda/6)100\%$ credible intervals for the 6 possible pairwise differences.

We note that the distribution of $\phi_{ORD}|\mathbf{y}$ is a multivariate distribution. For this distribution, we determine (i) the probability of Γ_h having order r given \mathbf{y} for $h = 1, 2, 3, 4$ and $r = 1, 2, 3, 4$,

and (ii) the most probable ordered sequence. To determine (i) we obtain the relative frequency of times that Γ_h has order r in the empirical posterior distribution. For (ii), we obtain the relative frequency of the mode ϕ_{ORD} vector.

Implementing MCMC simulation in such models as the one described here has been greatly simplified with software such as WinBUGS. We used the `R2WinBUGS` package by Sturtz, Liggs, and Gelman (2004) to call WinBUGS version 1.4 from R version 2.0.1. WinBUGS was used to carry out the MCMC simulations, and R was used to compute statistics of interest from the empirical posterior distributions generated in WinBUGS.

When running the MCMC, we ran three chains simultaneously. This allowed us to monitor the convergence of the three chains to a common posterior distribution for any of the parameters of interest. Based on plots of the Brooks-Gelman-Rubin statistic versus iterate, all of the monitored posterior distributions reached a stable state. All posterior distributions were stable after 6000 iterations of the MCMC. We collected every 50th iterate after the burn-in of 6000 to bring autocorrelation between iterate values within a chain to negligible levels. We used approximately 6000 thinned iterates to provide the empirical distributions. Starting values of the parameters for the three MCMC chains were chosen to cover a wide range of values.

4 Results

In Figure 1, we present the estimated half-life parameters Γ_h for treatments $h = 1, 2, 3, 4$, and 95% confidence/credible intervals (henceforth CIs) resulting from the two different models: the linear mixed model (LMM) described in Section 3.1 and the Bayesian hierarchical model (BHM) described in Section 3.2.1. The point estimates are means from the LMM and medians from the BHM. The maximum absolute difference between the LMM and BHM point estimates was 0.16. The widths of the BHM CIs were, at most, 0.61 of those of the LMM CIs.

To test for main effects of preparation and temperature and for their interaction effect, we use the CI approach to hypothesis testing: if zero is contained in the CI, there is no evidence of an effect. The point and interval estimates of these effects are presented in Figure 2. From the figure, we see that the conclusions on the main effect of temperature and the interaction effect coincide for the LMM and the BHM (based on a significance level of 0.05). The BHM allows us to detect a preparation effect, whereas the LMM does not.

The maximum absolute difference between the LMM and BHM point estimates of these effects was 0.19. The widths of the BHM CIs for these effects were no more than 0.58 of those of the LMM CIs.

To determine if there were any detectable differences between Γ_h for $h = 1, 2, 3, 4$, we performed all pairwise comparisons between them. In Figure 3, we present point estimates and simultaneous 95% CIs for all 6 differences. The maximum absolute difference between the LMM and BHM point estimates was 0.28. The widths of the BHM CIs were no more than 0.66 of those of the LMM CIs. Both the LMM and BHM detected a difference between treatments 1 and 4. Using the BHM, we were also able to detect differences between treatments 1 and 2 and treatments 3 and 4.

Using the BHM, we were able to produce posterior probabilities of Γ_h having order r for $h = 1, 2, 3, 4$ and $r = 1, 2, 3, 4$. These probabilities are presented in Table 1. Analogous probabilities using the LMM cannot be computed. Table 1 may also be used to find, for example, the posterior probability that the temperature 4° (i.e. $0.5(\Gamma_1 + \Gamma_3)$) has order r^* , with $r^* = 1, 2$. By way of ex-

ample, we compute $P(\text{order}(0.5(\Gamma_1 + \Gamma_3)) = 2 | \mathbf{y}) = 0.5 [(0.0012 + 0.9790) + (0.9322 + 0.0010)] = 0.9567$.

Additionally, the BHM allows us to find the most probable ordering of the vector $\mathbf{\Gamma}$. It is, from least to greatest, $(\Gamma_4, \Gamma_2, \Gamma_3, \Gamma_1,)$ with posterior probability 0.9167. Again, there is no analogue to this using the LMM.

5 Discussion

For parameters that could be estimated using either the LMM or BHM, resulting point estimates differed very little. However, the gain in precision of point estimation using the BHM was substantial when compared to that using the LMM. The CIs produced from the BHM were less than two thirds the width of analogous CIs produced from the LMM. As such, we were able to detect certain differences under the BHM that were not detectable under the LMM. There are at least two reasons for the BHM having more precision than the LMM: the BHM constrains the estimates of Γ_h to be positive, whereas the LMM does not; and the BHM “borrows strength” from other EUs in the estimation of parameters for a particular EU, which in turn allows more precise estimates of the treatment parameters, e.g. Γ_h .

To better understand the effects of hyperparameters on (Bayesian) point and interval estimates of Γ_h , we made all priors increasingly more concentrated than the ones reported above, but left the means unchanged. This is equivalent to saying that we know more than when we use diffuse priors. By making the priors more concentrated, we are essentially saying we are more sure that the means of the posteriors will be closer to the means of the priors. Accordingly, the point estimates for Γ_h were shifted downward toward 1 – the mean of the prior on Γ_h (results not shown). We thought it was interesting that the interval estimates arising from these more concentrated priors were wider than those arising from the diffuse priors (results not shown). Since the researcher did not have any prior expectation of what the half-life values might be, among the Bayesian results, we are more comfortable with the results from the diffuse priors.

Some concerns that we have not addressed in this article are the possible violations of the assumptions (common to both models) implied in Display (3) and, from the BHM, the independence assumptions among the elements in (a, b, g) , (A, B, G) , (α, β, γ) , and (A, B, Γ) . We have also not addressed the large discrepancy in modeling assumptions between the LMM and BHM. How these issues affect the results we have presented here is unknown. A better comparison would be between the classical and Bayesian hierarchical models; we are currently investigating this.

We recall the objective of the experiment: determine which treatment has the longest half-life. Using the BHM, we could directly address this objective by obtaining posterior probabilities of Γ_h having order r for $h = 1, 2, 3, 4$ and $r = 1, 2, 3, 4$. (The half-life of SA using P1 at 4°C (i.e. Γ_1) had the longest half-life with posterior probability 0.9790 (see Table 1).) With the LMM, we are left to mean separation methods, such as all pairwise comparisons with an adjustment for the multiple comparisons. If the treatment having the highest estimated mean is not significantly different from the next lower mean, then no single treatment can be selected as the best. (This is often the case in other experiments of this nature and is the case here.) More importantly, we cannot compute a probability that a treatment is the best using the LMM. When using the LMM, the best that can be said is that 4°C is better than 23°C (see Figure 2).

The classical approach we have presented for this problem is often used, primarily for its ease of implementation and fit to standard methods often used in subject-specific publications.

However, this approach fails if the time intervals differ between any of the measured units. (Note that missing data on any of the units would result in differing time intervals.) Such would violate the assumption of homogeneity of variance among the (derived) responses. We believe that these types of longitudinal experiments free of this problem are a minority. Additionally, when fitting the nonlinear model to the data from each unit, it is a good idea to have at least 5 observations for each parameter to be estimated in the nonlinear model. Again, very few experiments like the one presented here have the luxury of taking so many observations from one unit. So, though the classical approach is appealing due to its simplicity, we feel that it can seldom be applied correctly due to limitations in the data or the experiment itself.

The random coefficients nonlinear regression is a viable approach when the time intervals differ between the units and when there are relatively few measurements from each unit (Chapter 4 and 8 of Davidian and Giltinan (1995)). This is true for either a classical or Bayesian approach to the regression. Using this sort of regression model allows us to more appropriately analyze the data from these types of longitudinal experiments. However, implementation can be difficult. Besides the actual coding of statistical software, other difficulties that often occur are selecting good starting values for the parameters in the model, modeling correlations between parameters within an EU, and obtaining convergence. These problems are common to both frequentists and Bayesians, and we experienced all of these in our analysis.

Though difficult to implement, random coefficients nonlinear regression (i.e. hierarchical modeling) is more often the better way to approach such longitudinal experiments than using derived response variables in a linear (mixed) model. The beauty of a Bayesian hierarchical model is that we can easily produce direct answers to the researchers' questions. Whereas with classical methods, we must often rely on a series of hypothesis tests to arrive at answers in a "round-about" way.

In this paper, we have exhibited a classical and Bayesian analysis of a longitudinal experiment with differently sized EU. We found that the Bayesian method produced not only more precise estimates of desired quantities, but was also able to directly answer the researcher's main question: "Which treatment is the best?". Though much easier to implement than the Bayesian method, the classical approach that used derived response variables in a linear mixed model produced similar point estimates as the Bayesian method, but could not cleanly answer the researcher's main question.

References

- [1] Carlin, B.P. and Louis, T.A. (2000) *Bayes and empirical Bayes methods for data analysis, 2nd edition*, Boca Raton: Chapman & Hall/CRC.
- [2] Davidian, M. and Giltinan, D.M. (1995) *Nonlinear models for repeated measurement data*, Boca Raton: Chapman & Hall/CRC.
- [3] Diggle, P.J., Liang, K.-Y., and Zeger, S.L. (1994) *Analysis of longitudinal data*, New York: Oxford University Press, Inc.
- [4] Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. (1995) *Bayesian data analysis*, London; New York: Chapman & Hall Ltd.

- [5] Sturtz, S., Ligges, U., and Gelman, A. (2004) “R package R2WinBUGS: running WinBUGS from R (version 0.2-1).” Available at <http://cran.R-project.org>; accessed 11 May 2004.
- [6] Zelaya, I.A., Owen, M.D.K., and Landes, R.D. (manuscript) “Evaluation of spectrophotometric and HPLC methods for shikimic acid determination: models in glyphosate resistant and susceptible soybeans (*Glycine max* L. Merr) and maize (*Zea mays* L.)”

Appendix

Figure 1: Point and interval estimates of treatment half-life parameters Γ_h for $h = 1, 2, 3, 4$. LMM and BHM point estimates are noted with \blacksquare and \bullet , respectively.




Table 1: Posterior probabilities of Γ_h having order r for $h = 1, 2, 3, 4$ and $r = 1, 2, 3, 4$.

| \mathbf{p} | $\mathbf{r = 1}$ | $\mathbf{r = 2}$ | $\mathbf{r = 3}$ | $\mathbf{r = 4}$ |
|--------------|------------------|------------------|------------------|------------------|
| Γ_1 | 0.0008 | 0.0190 | 0.0012 | 0.9790 |
| Γ_2 | 0.0017 | 0.9606 | 0.0178 | 0.0199 |
| Γ_3 | 0.0471 | 0.0197 | 0.9322 | 0.0010 |
| Γ_4 | 0.9504 | 0.0007 | 0.0488 | 0.0001 |

Figure 2: Point and interval estimates of effects of preparation, temperature, and their interaction. LMM and BHM point estimates are noted with ■ and ●, respectively.




Figure 3: Point and interval estimates of all pairwise differences between Γ_h for $h = 1, 2, 3, 4$. LMM and BHM point estimates are noted with ■ and ●, respectively.

