

STATISTICAL INFERENCE FOR MINIMUM INHIBITORY CONCENTRATION DATA

by

Huanhuan Wu

Bachelor of Science, Simon Fraser University, 2006

Bachelor of Business Administration, Jinan University, 2002

A PROJECT SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the Department
of
Statistics and Actuarial Science

© Huanhuan Wu 2008
SIMON FRASER UNIVERSITY
Spring 2008

All rights reserved. This work may not be
reproduced in whole or in part, by photocopy
or other means, without the permission of the author.

APPROVAL

Name: Huanhuan Wu
Degree: Master of Science
Title of project: STATISTICAL INFERENCE FOR MINIMUM INHIBITORY
CONCENTRATION DATA

Examining Committee: Dr. Brad McNeney
Chair

Dr. Carl J. Schwarz
Senior Supervisor
Simon Fraser University

Dr. Thomas M. Loughin
Simon Fraser University

Dr. Tim Swartz
External Examiner
Simon Fraser University

Date Approved: April 7, 2008

Abstract

The Minimum Inhibitory Concentration (MIC) is the smallest concentration of an antimicrobial agent that inhibits the growth of bacteria. The value is obtained in a highly mechanized fashion, but this procedure only provides interval censored reading. It is often of interest to use data collected from complex experiments to see how the mean MIC is affected by different factors.

Because the MIC value is interval censored, ordinary least squares cannot be used. For models containing only fixed effects, maximum likelihood estimates (MLE) can be obtained. For models containing random effects, MLE methods are infeasible and Bayesian approaches are required. Model building, selection and diagnostic procedures are presented for selecting the appropriate model. In cases where several models seem to fit the data equally well, model averaging is also performed to get model averaged estimates. Four real data sets are analyzed using the methodology we developed.

Keywords: Bayesian inference; Interval censoring; Maximum likelihood estimate; Minimum inhibitory concentration; Model average

Executive Summary

The current summary reports of the analysis of MIC data all have some limitation. All current reports focus on selected percentiles of the anti-microbial resistance pattern, but do not use the entire resistance pattern to look for changes. In the analysis procedure, none of the collected covariates is used to investigate if the MIC values vary by region of origin, producer of origin, feeding regime, size of herd/flock, temporally, etc. Additionally, none of the current summary reports have any indication of the precision of the estimates.

The main goals of this project are to use statistical methods to relate the logMIC (because it is assumed that the unknown MIC has a log-normal distribution) to covariates; to get general estimates that explicitly account for the sampling design; and to obtain the precision of estimates.

A methodology was developed to deal with interval censored response variables. We start by building possible models, then we select a model and finally perform model diagnostics. When several models fit the data equally well, we model average. In the last step, we calculate estimates from the appropriate model and use them to make inferences.

The analysis of the four real data sets support different “best” model. When a model contains fixed effects only, the maximum likelihood method (implemented by SAS) is used. When a model contains random effects, a Bayesian analysis using Markov Chain Monte Carlo (MCMC) techniques (implemented by WinBUGS) is performed.

The analysis of the four real data sets all indicate that the mean logMIC is related to some covariates. This information is useful in designing future experiments.

If the variance of the random effects (such attributed to farms) is small relative to the variance of random noise; the estimates and precision derived using maximum likelihood ignoring the random effects and derived using Bayesian method including the random effects are quite similar.

Dedication

To Wei and Xinxin.

Acknowledgments

I spent four years for my second Bachelor degree and my M.Sc. degree at Simon Fraser University. My journey in academic is full of challenge and difficulty. I want to thank many people who involved in my life at SFU. Without their help I can not be working in academic today.

First, I would like to thank my supervisor, Dr. Carl Schwarz whose support, patience, guidance, encouragement and understanding helped me through the “hard” time. Carl give me a chance to achieve my academic goal.

I would like to thank my committee members, Dr. Tom loughin and Dr. Tim Swartz for their time and suggestions.

I would like to thank my undergraduate adviser Dr. Larry Weldon for his kind and insight suggestion, Dr. Boxin Tang for his help, Dr. Charmaine Dean and Robin Insley for suggestion on being new parent.

Thanks to my fellow graduate students, Carolyn, Cindy, Jean, Luyao, Monica, Simon, Suli and Wendell for their help.

I would like to thank my parents, my grand parents for their education and support.

Last but not least, I give my special thanks to my husband Wei and my daughter Xinxin, for their love.

Abbreviation

AIC	Akaike's Information Criterion
APE	Bayesian Average Posterior Estimate
BCMAL	The British Columbia Ministry of Agriculture and Lands
BIC	Bayesian Information Criterion
BF	Bayes Factor
BGR	Gelman-Rubin statistic
DIC	Deviance Information Criterion
MCMC	Markov Chain Monte Carlo
MIC	Minimum Inhibitory Concentration
MLE	Maximum Likelihood Estimate
OLM	Ordinary Linear Model
SAS	Statistics Analysis System
SD	Standard Deviation
SE	Standard Error
WinBUGS	Windows version of the Bayesian inference Using Gibbs Sampling

Contents

Approval	ii
Abstract	iii
Executive Summary	iv
Dedication	v
Acknowledgments	vi
Abbreviation	vii
Contents	viii
List of Tables	x
List of Figures	xii
1 Introduction	1
2 Models for the data	3
2.1 Model for data with no censoring and with no random effects acting upon the mean response	4
2.2 Model for data with no censoring, but with random effects acting upon the mean response	4
2.3 Model for data with censoring but with no random effects acting upon the mean response	4

2.4	Model for data with censoring and with random effects acting upon the mean response	5
2.5	Model notation	6
3	Illustrative analyses	8
3.1	VISIT linear, no TYPE and no FARM effects	9
3.2	VISIT linear, TYPE classification, no FARM effect	11
3.3	VISIT classification, TYPE classification, no FARM effect	13
3.4	VISIT linear, TYPE classification, FARM random effect (big)	16
3.5	VISIT linear, TYPE classification, FARM random effect (small)	19
3.6	Summary of simulation results	22
4	Model building, selection and diagnostics	23
4.1	Model building	23
4.2	Model selection	27
4.2.1	Likelihood Model	27
4.2.2	Bayesian Model	28
4.3	Model checking	30
4.3.1	Likelihood model	30
4.3.2	Bayesian model	31
5	Real Examples	33
5.1	Avian data set (Avian plate)	33
5.2	Avian data set (NARMS plate)	41
5.3	Dairy data set (<i>Campylobacter</i>)	48
5.4	Dairy data set (<i>E.coli</i>)	61
6	Conclusions and Suggestions	70
6.1	Future work	71
	Bibliography	72

List of Tables

3.1	True values and estimates of parameters from model $[Y]=\text{VISIT}$	11
3.2	True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT}$. .	13
3.3	True values of parameter of model $[Y]=\text{TYPE}(C) \text{ VISIT}(C)$	14
3.4	True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT}(C)$.	16
3.5	True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT}$ FARM(R)(big).	19
3.6	True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT}$ FARM(R)(small).	20
4.1	Example of model averaging	29
5.1	Results of fitting different models to the Avian-CEFTIF data using MLE . .	36
5.2	Estimates for the Avian-CEFTIF data from the MLE and two Bayesian anal- yses.	39
5.3	DIC value for model $[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE FARM}(R)$ with and without farm random effect	40
5.4	DIC value for all possible models	40
5.5	Results of fitting different models to the Avian (NARMS plate)-CEFTIF data using MLE	44
5.6	Estimates for the Avian (NARMS plate)-CEFTIF data from the MLE and two Bayesian analyses.	47
5.7	DIC values for model $[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE FARM}(R)$ with and without farm random effect.	48
5.8	Results of fitting different models to the Dairy (<i>Camylobacter</i>)-Azithromycin data using MLE	53

5.9	Estimates for the Dairy (<i>Camylobacter</i>)-Azithromycin data (full data) from the MLE and two Bayesian analyse.	56
5.10	DIC values for model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$ with and without random farm effect.	57
5.11	Results of fitting different models after outliers were deleted use MLE method	57
5.12	Estimates for the Dairy (<i>Camylobacter</i>)-Azithromycin data (reduced data) from the MLE and two Bayesian analyse.	60
5.13	DIC values for model $[Y] = \text{VISIT}(C) \text{ FARM}(R)$ with and without farm random effect.	61
5.14	Results of fitting different models to the Dairy (<i>E.coli</i>)-CEFTIF data set using MLE method	65
5.15	Estimates for the Dairy (<i>E.coli</i>)-CEFTIF data from the MLE and two Bayesian analyses.	69
5.16	DIC values for model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$ with and without farm random effect.	69

List of Figures

3.1	Plot of the lower bounds of simulated data from the model $[Y]=\text{VISIT}$ with fitted regression line.	10
3.2	Plot of the lower bounds of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT with fitted regression lines.	12
3.3	Plot of the lower bounds of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT(C) with fitted regression lines.	15
3.4	Plot of the lower bounds of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT FARM(R)(big) with fitted regression lines.	18
3.5	Plot of the lower bounds of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT FARM(R)(small) with fitted regression lines.	21
4.1	Portion of the Avian data set	23
4.2	Plot of the lower bounds of logMIC vs. visit	25
4.3	Plot of connecting median lower bounds of logMIC of each visit of each farm vs. visit for different farms	26
5.1	Part of the raw data for the Avian data set.	34
5.2	Plot of the lower bounds of logMIC_CEFITIF vs. visit for Avian data set.	35
5.3	Residual plot of model $[Y] = \text{VISIT}(C) \text{TYPE}(C) \text{VISIT} \times \text{TYPE}$	37
5.4	Part of raw data for the Avian (NARMS plate) data set.	41
5.5	Plot of the lower bounds of logMIC_CEFITIF vs. visit for Avian (NARMS plate) data set	42
5.6	Plot of connecting median lower bounds of logMIC_CEFITIF of each visit of each farm vs. visit for different farm for Avian (NARMS plate) data set	43
5.7	Residual plot of model $[Y] = \text{VISIT}(C) \text{TYPE}(C) \text{VISIT} \times \text{TYPE}$	45

5.8	Part of the raw data of the Dairy (<i>Campylobacter</i>) data set.	49
5.9	Plot of the lower bounds of logMIC-Azithromycin vs. visit with fitted lines for different group	50
5.10	Plot of connecting median lower bounds of logMIC_CEFOTIF of each visit of each farm vs. visit for different farm	52
5.11	Residual plot of model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$	54
5.12	Residual plot of model $[Y] = \text{VISIT}(C)$	58
5.13	Part of raw data of Dairy (<i>E.coli</i>) data set.	62
5.14	Plot of the lower bounds of logMIC-CEFTIF vs. visit with fitted loess lines for Dairy (<i>E.coli</i>) data set	63
5.15	Plot of connecting median lower bound of logMIC_CEFOTIF of each visit of each farm vs. visit for different farm.	64
5.16	Residual plot of model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$	66

Chapter 1

Introduction

The British Columbia Ministry of Agriculture and Lands (BCMAL) has initiated several projects to measure the anti-microbial resistance of bacteria found in chickens, swine, cattle, and other livestock. Data have been collected for the past several years on the bacteria *Escherichia coli*, *Salmonella* sp., *Campylobacter* sp. and *Enterococcus* sp.

The sampling design is complex and depends upon the animal being tested. For example, 51 dairy calves on 26 farms were enrolled in the Dairy project. Four visits were made to each farm to collect samples at 1-7 days of age; 14-28 days of age; 90-120 days of age; and 6-8 months of age. At each visit, one individual sample was taken from the enrolled calf, one pooled sample from five calves of the same age, and one pooled sample from five cows. Additionally, one sample was taken from the mother at the first visit, and one liquid calf feed specimen was taken at the first, second, and third visits. A one-page questionnaire was completed at each of the four visits and calf antibiotic treatments were recorded.

Samples were cultured and the resulting bacterial isolates were tested against an array of anti-microbial agents at a variety of concentrations. For example, *E.coli* and *Salmonella* bacterial cultures were tested in a highly mechanized fashion against an array of 17 anti-microbial agents at a variety of concentrations (a total of 96 combinations of 17 anti-microbials x 3-9 concentrations per agent) in a microtitre plate. The output from testing one sample is a vector of MIC values for each anti-microbial agent, as determined by the broth microdilution method (NCCLS/CLSI - M7-A5).

For example, the array may contain the anti-microbial agent Ceftiofur at concentrations of 0.12, 0.25, 0.50, 1, 2, 4 and 8 mg/L. A particular bacterial isolate may show inhibition of bacterial growth at 2, 4 and 8 mg/L but growth at lower concentrations. The reported

MIC value would then be 2 mg/L. This value means that a concentration of 2 mg/L showed inhibition, but a concentration of 1 mg/L did not show inhibition. Consequently, the true inhibition is between 1 and 2 mg/L. Similarly, if the reported MIC value is less than 0.12 mg/L (the lowest reading on the plate), all that is known is that the true inhibitory concentration is between 0 and 0.12 mg/L. Equally, if the reported MIC value is greater than 8 mg/L, this indicates that no concentrations on the plate showed inhibition of growth so the actual concentration that inhibited bacterial growth is greater than 8 mg/L. In all cases, the MIC value reported is a censored value: all that is known is the MIC is either below the minimum concentration tested, or between two concentrations, or above the maximum concentration tested in the array for that anti-microbial agent.

Using the E-Test, samples of *Campylobacter* sp. were tested against 8 anti-microbials. This process is less mechanized, but also provides MIC data.

Chapter 2

Models for the data

The analysis of the MIC data has several goals. What is the relationship between the MIC and other explanatory variables such as the type of farm operation, visit, or the age of animal? How can the precision for the estimated parameters be found? How can trend lines be fitted with censored data?

Several assumptions, which are common for these types of data (Lee and Whitmore, 1999), will be used in the analyses.

1. Tests are carried out without measurement error.
2. Tests are monotonic in the sense that if inhibition is detected at concentration C , then it is detected at any C' , where $C' > C$.
3. The unknown MIC value of the anti-microbial agent, when tested against a bacteria, has a log-normal distribution.

The logarithm of the MIC value, then, gives a normally distributed random variable that we denote by $Y = \log(\text{MIC})$. Hence Y has the normal distribution $Y \sim N(\mu, \sigma^2)$ where μ and $\sigma^2 > 0$ are the unknown mean and variance parameters. The mean μ may depend upon covariates. These parameters are the quantities that we wish to estimate.

2.1 Model for data with no censoring and with no random effects acting upon the mean response

If the exact MIC value is measured, and if all the explanatory variables are fixed effects, the analysis would proceed using the Ordinary Linear Model (OLM):

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$$

where \mathbf{Y} is an $n \times 1$ response vector, \mathbf{X} is an $n \times p$ predictor matrix, $\boldsymbol{\beta}$ is an unknown $p \times 1$ regression coefficient vector, p is the number of regression coefficients, and $\boldsymbol{\epsilon}$ is a noise term for the model which is multivariate normal $\text{MVN}(0, \sigma^2 \mathbf{I}_n)$, where \mathbf{I}_n is the $n \times n$ identity matrix. Estimates are available in closed form, and are discussed in standard text books (Kutner et al., 2004) on regression. Standard software (e.g., SAS (SAS Institute 2008) (PROC REG, PROC GLM), R (R Development Core Team. 2008) (package LM, GLM)) can be used to obtain estimates.

2.2 Model for data with no censoring, but with random effects acting upon the mean response

If the MIC value is non-censored, but some of the explanatory variables are random effects, this leads to the standard linear mixed model

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\epsilon}$$

where $\boldsymbol{\gamma}$ is a vector of random effects that affect the mean response through the design matrix \mathbf{Z} , and γ_i is distributed as $N(0, \sigma_\gamma^2)$ where γ_i is the i^{th} element of $\boldsymbol{\gamma}$. Again, the analysis is standard (Dobson, 2002). Standard software (e.g., SAS (PROC MIXED), R (package lme4)) can be used to obtain estimates.

2.3 Model for data with censoring but with no random effects acting upon the mean response

When the MIC value is censored (as in these experiments), the exact value of the MIC is unknown; only an interval containing the MIC is known. Ordinary least squares cannot be

used; instead tobit regression (Long, 1997) is used and estimation is done using maximum likelihood methods.

Let T_1, T_2, \dots, T_m represent the concentrations available in the array for the anti-microbial agents, with $T_0 = 0, T_{m+1} = \infty, T_i < T_{i+1}, \forall i \in \{0, 1, 2, \dots, m\}$. The likelihood function is given by

$$L = \prod_{i=1}^{m+1} [S(T_{i-1}) - S(T_i)]^{l_i}$$

where l_i is the number of samples with the MIC value between T_{i-1} and T_i , $S(T_i) = Pr(T > T_i) = 1 - \Phi\left(\frac{\log T_i - \mu_i}{\sigma}\right)$, $S(T_0) = 1, S(T_{m+1}) = 0$, $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp\left(-\frac{u^2}{2}\right) du$ is cumulative distribution function for the standard normal distribution with $\mu = 0$ and $\sigma = 1$, and

$$\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\beta}$$

Estimates are not available in closed form and standard numerical methods must be used. Standard software (e.g., SAS (PROC LIFEREG), R (package KMsurv)) can be used.

2.4 Model for data with censoring and with random effects acting upon the mean response

When the MIC value is censored and some of the explanatory variables are random effects, the analysis becomes more complicated.

A random-effect tobit model (Long, 1997) could be used and it has a likelihood of

$$L = \prod_{i=1}^{m+1} \left[\int_{\gamma} (S(T_{i-1}) - S(T_i)) f(\gamma) \right]^{l_i}$$

where

$$S(T_i) = 1 - \Phi\left(\frac{\log T_i - \mu_i}{\sigma}\right)$$

but now

$$\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\gamma}$$

where $\boldsymbol{\gamma}$ is a vector of random effects, γ_i is distributed as $N(0, \sigma_\gamma^2)$, and γ_i is the i^{th} element of $\boldsymbol{\gamma}$.

Likelihood estimation is intractable because of the need to integrate over the random effects in μ .

What happens if the random effects are ignored in the modeling process? The model, the analysis, and the conclusions may not be correct. For example, in many situations, if random effects are ignored, the variance of the response variable is understated. Estimates can be unbiased, but the estimated standard errors will typically be too small. Test statistics based on the wrong standard error will typically lead to false positive results, i.e., claiming an effect when none really exists.

In these situations, a Bayesian analysis is often performed and Markov chain Monte Carlo (MCMC) methods are used to “integrate” over the random effects. The Bayesian model has the same likelihood of

$$L = \prod_{i=1}^{m+1} \left[\int_{\gamma} (S(T_{i-1}) - S(T_i)) f(\gamma) \right]^{l_i}$$

and

$$\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\gamma}$$

but adds the prior distributions for the parameters

$$\beta_i \sim N(0, \sigma_{\beta_i}^2) \quad i = 1, \dots, p \text{ and}$$

$$\gamma \sim N(0, \sigma_{\gamma}^2)$$

$$\sigma_{\gamma}^2 \sim \text{GAMMA}(a, b)$$

where p is the number of regression coefficients, and $\sigma_{\beta_i}^2$ and a, b are assumed to be known. Priors are usually chosen based on previous knowledge (e.g., the results of earlier studies). In our study, this form of the priors is chosen as above because these are standard priors for these types of data.

WinBUGS (Lunn et al. 2000) software, an interactive Windows version of the BUGS (Bayesian inference Using Gibbs Sampling) program for Bayesian analysis of complex statistical models using MCMC techniques, can be used for inference in these models.

2.5 Model notation

Despite having different types of models, we can define a standard notation to describe models in terms of how they fit the data. This notation is similar to that used in linear

models.

Let

Y represent the logMIC if known exactly

[Y] represent the logMIC if censored with the censoring bounds implicitly known

X represent a linear fixed effect (such as a straight line)

X(C) represent a fixed classification effect (such as different groups)

X(R) represent a random linear effect (rarely used)

X(C)(R) represent a random classification effect

Nesting and crossing of effects are denoted in the usual way. For example, consider an experiment where there are two TYPES of farms (e.g., different husbandry practices). At each of several VISITS, equally spaced apart, a sample of fecal matter is obtained, bacteria isolated, and the MIC obtained as outlined in the introduction. A potential model for this experiment would be

$$[Y] = \text{TYPE}(C) \text{ FARM}(\text{TYPE})(R) \text{ VISIT } \text{TYPE}(C) \times \text{VISIT}$$

This model means the dependent variable is logMIC, which is censored. The explanatory variables are TYPE (classification), FARM (random effect, nested within TYPE), VISIT (linear), and an interaction term between TYPE and VISIT (e.g., the linear effect is different for each type of farm).

Chapter 3

Illustrative analyses

Simulated data will be used to demonstrate the use of various computer packages for the different models discussed in Section 2. For each model type, the parameter values used to generate the data and the estimates from two types of analysis are presented.

Why do we use simulated data? The first reason is that the real data is limited in size; for example, actual sample sizes range from 323 to 519. The simulated data have sample sizes ranging from 1000 to 2000. With large sample sizes, the results are sufficiently precise that any approximate bias in the estimates can be determined.

The second reason is that a simulation can demonstrate the robustness of the algorithm used in this analysis. If the methods used are robust, then the result from a simulated data set should be similar or consistent with the true parameter values.

The analysis procedure is as follows. Different models will be used to generate simulated data and estimates will be obtained using maximum likelihood (implemented with SAS PROC LIFEREG) and a Bayesian approach (implemented with WinBUGS). The advantage of SAS PROC LIFEREG is that it is the only procedure in SAS that can handle interval censoring data; the disadvantage is that it is not designed to work with random effects models. The Bayesian approach can deal with complicated models that include mixed effects model, but priors need to be specified and the implementation is not as simple as it is in SAS PROC LIFEREG.

What type of data is simulated? The general scenario is that multiple measurements are taken over several VISITS from FARMS of two TYPES, VISIT is the number of times that farms were being visited. We start with the simplest model $[Y] = \text{VISIT}$, where logMIC is only related to VISIT. Then we add a variable for the TYPE of farm, and the model

becomes $[Y] = \text{VISIT} \text{ TYPE}(C)$, i.e., $\log\text{MIC}$ is related to VISIT and TYPE . When VISIT is a classification variable, we have another model. Last, we add a random FARM effect, and the model becomes $[Y] = \text{VISIT} \text{ TYPE}(C) \text{ FARM}(\text{TYPE})(R)$, where $\log\text{MIC}$ is related to VISIT , TYPE and FARM .

The following sections are arranged in the order of the models used. For each model, we will describe the simulation experiment, present the fitted line and raw data on a graph, and summarize the simulation results.

3.1 VISIT linear, no TYPE and no FARM effects

In this simulation experiment, separate farms of the same type are measured at each of five VISITS. Consequently, the farm effects are confounded with measurement error. The total sample size is 1000. The concentrations of the anti-microbial are predetermined at 0.25, 0.5, 1, 2, 4 mg/L, which imply that the censoring intervals are: $(\log(0)=-\infty, \log(0.25))$; $(\log(0.25), \log(0.5))$; $(\log(0.5), \log(1))$; $(\log(1), \log(2))$; and $(\log(4), \log(\infty)=\infty)$. Once the observed true MIC value is simulated, the interval that contains the true value is then used to represent the $\log\text{MIC}$ value obtained.

The simplified syntax for the model of this experiment is

$$[Y] = \text{VISIT}$$

The true model is

$$y_i = 1.3 + 0.86\text{visit}_i + \epsilon_i$$

$$\epsilon_i \sim N(0, 2^2), \text{visit}_i \text{ is the number of visit to } i^{\text{th}} \text{ observation.}$$

Figure 3.1 shows the raw data and the fitted line. The line doesn't seem to fit the data in the figure, because (a) only the lower bound of each $\log\text{MIC}$ was plotted, and (b) the largest concentration of anti-microbial used was 4, so all the $\log\text{MIC}$ s that exceed 4 are represented as $> \log 4$.




Figure 3.1: Plot of the lower bounds of simulated data from the model $[Y]=\text{VISIT}$ with fitted regression line. Values are jittered to avoid overplotting. Anti-microbial concentrations used were 0.25, 0.5, 1, 2, 4 mg/L.

Table 3.1 summarizes the simulation results. All the estimates are close to the true parameter values, indicating that censored data provide sufficient information to estimate the parameters reliably if the sample size is sufficiently large.

	Intercept	Effect of VISITS	Standard deviation of random noise
	β_0 (SE)	β_{visit} (SE)	σ (SE)
True Value	1.30	0.86	2.00
MLE	1.32 (0.21)	0.80 (0.09)	2.05 (0.15)
APE	1.32 (0.21)	0.78 (0.08)	2.01 (0.12)

Table 3.1: True values and estimates of parameters from model $[Y]=\text{VISIT}$. Estimates are Maximum Likelihood Estimates (MLE) and Bayesian Average Posterior Estimates (APE). Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs.

3.2 VISIT linear, TYPE classification, no FARM effect

In this simulation experiment, separate farms of two TYPES are measured at each of five VISITS. The farm effects are again confounded with measurement error. The total sample size is 1000. The same anti-microbial concentrations as in Section 3.1 are used.

The simplified syntax for the model of this experiment is

$$[Y] = \text{TYPE}(C) \text{ VISIT}$$

The true model is

$$y_i = 0.3 + 0.8\text{visit}_i + \epsilon_i \quad \text{for farms of type}=1$$

$$y_i = 0.8 + 0.8\text{visit}_i + \epsilon_i \quad \text{for farms of type}=2$$

$$\epsilon_i \sim N(0, 2^2)$$

Figure 3.2 shows the raw data and the fitted lines. The two fitted regression lines are parallel. This is consistent with the true relationship between the two types.




Figure 3.2: Plot of the lower bounds (star for farms of TYPE=1, dot for farms of TYPE=2) of simulated data from the model $[Y]=\text{TYPE}(C) \text{ VISIT}$ with fitted regression lines (dashed line for farms of TYPE=1, solid line for farms of TYPE=2). Values are jittered to avoid overplotting. Values of TYPE=1 were jittered to left of visit. Values of TYPE=2 were jittered to right of visit. Anti-microbial concentrations used were 0.25, 0.5, 1, 2, 4 mg/L.

Table 3.2 summarizes the true values and the estimates of the parameters. All the estimates are close to the true parameter values, indicating that censored data provides sufficient information to estimate the parameters reliably, if the sample size is sufficiently large.

	Intercept of type=2	Effect of TYPE of FARM	Effect of VISITS	Standard deviation of random noise
	β_0 (SE)	$\beta_{type1} - \beta_{type2}$ (SE)	β_{visit} (SE)	σ (SE)
True Value	0.80	-0.50	0.80	2.00
MLE	0.78 (0.21)	-0.45 (0.18)	0.79 (0.08)	2.09 (0.12)
APE	0.79 (0.21)	-0.46 (0.18)	0.79 (0.08)	2.10 (0.12)

Table 3.2: True values and estimates of parameters from model $[Y]=TYPE(C) VISIT$. Estimates are Maximum Likelihood Estimates (MLE) and Bayesian Average Posterior Estimates (APE). Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs.

3.3 VISIT classification, TYPE classification, no FARM effect

In this simulation experiment, separate farms of two TYPES are measured at each of five VISITS. VISIT is treated as a classification effect. The farm effects are confounded with measurement error. The total sample size is 2000. The same anti-microbial concentrations as in Section 3.1 are used.

The simplified syntax for the model of this experiment is

$$[Y] = TYPE(C) VISIT(C)$$

Table 3.3 gives the parameters for the true models of each type and visit.

Figure 3.3 shows the raw data and the fitted lines. The two fitted regression lines are parallel. This is consistent with the true relationship between the two types.

	Type1	Type2
Visit1	$y_i=3.0+\epsilon_i$	$y_i=3.5+\epsilon_i$
Visit2	$y_i=2.8+\epsilon_i$	$y_i=3.3+\epsilon_i$
Visit3	$y_i=2.4+\epsilon_i$	$y_i=2.9+\epsilon_i$
Visit4	$y_i=2.1+\epsilon_i$	$y_i=2.6+\epsilon_i$
Visit5	$y_i=3.1+\epsilon_i$	$y_i=3.6+\epsilon_i$
$\epsilon_i \sim N(0, 2^2)$		

Table 3.3: True model for the simulation of the model $[Y]=\text{TYPE}(C) \text{ VISIT}(C)$.




Figure 3.3: Plot of the lower bounds (star for farms of TYPE=1, dot for farms of TYPE=2) of simulated data from the model $[Y]=\text{TYPE}(C) \text{ VISIT}(C)$ with fitted regression lines (dashed line for farms of TYPE=1, solid line for farms of TYPE=2). Values are jittered to avoid overplotting. Values of TYPE=1 were jittered to left of visit. Values of TYPE=2 were jittered to right of visit. Anti-microbial concentrations used were 0.25, 0.5, 1, 2, 4 mg/L.

Table 3.4 summarizes the true values and the estimates of the parameters. Some estimates are not close to the true parameter values, but the 95% confidence intervals still cover the true values.

	Intercept of type=2 visit=5 β_0 (SE)	Effect of TYPE of FARM $\beta_{type_1} - \beta_{type_2}$ (SE)	Effect of VISITS $\beta_{visit_1} - \beta_{visit_5}$ (SE) $\beta_{visit_2} - \beta_{visit_5}$ (SE) $\beta_{visit_3} - \beta_{visit_5}$ (SE) $\beta_{visit_4} - \beta_{visit_5}$ (SE)	Standard deviation of random noise σ (SE)
true value	3.60	-0.50	-0.10 -0.30 -0.70 -1.00	2.00
MLE	3.61 (0.19)	-0.46 (0.12)	-0.18 (0.20) -0.41 (0.20) -0.74 (0.20) -1.15 (0.20)	2.04 (0.08)
APE	3.62 (0.20)	-0.46 (0.12)	-0.17 (0.21) -0.40 (0.21) -0.73 (0.20) -1.14 (0.20)	2.06 (0.08)

Table 3.4: True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT}(C)$. Estimates are Maximum Likelihood Estimates (MLE) and Bayesian Average Posterior Estimates (APE). Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs.

3.4 VISIT linear, TYPE classification, FARM random effect (big)

In this simulation experiment, separate FARMS of two TYPES are repeatedly measured at each of five VISITS. For each farm type, there are 20 farms, and each farm has 10 replications. The total sample size is 2000. We assume the farm effect is normally distributed with mean zero and is treated as a random effect. In this simulation, the farm effect has a variance much larger than the variance of the random noise. The same anti-microbial concentrations as in Section 3.1 are used.

The simplified syntax for the model of this experiment is

$$[Y] = \text{TYPE}(C) \text{ VISIT} \text{ FARM}(R)$$

The true model is

$$\begin{aligned}
 y_i &= 1.5 + 0.5\text{visit}_i + \text{farm.effect}_i + \epsilon_i \text{ for farms of type=1} \\
 y_i &= 0.5 + 0.5\text{visit}_i + \text{farm.effect}_i + \epsilon_i \text{ for farms of type=2} \\
 \text{farm.effect}_i &\sim N(0, 1.5^2) \\
 \epsilon_i &\sim N(0, 1^2)
 \end{aligned}$$

The MLE and Bayesian are different in this case because the SAS implementation of the MLE model cannot easily include the farm random effect.

The MLE model is

$$[Y] = \text{TYPE}(C) \text{ VISIT}$$

The model for the Bayesian analysis is

$$[Y] = \text{TYPE}(C) \text{ VISIT FARM}(R)$$

The prior distributions are: $\beta_0(\text{intercept}) \sim N(0, 10)$, $\beta.\text{type} \sim N(0, 10)$, $\beta.\text{visit} \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$, $\tau_{\text{farm}} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{\text{farm}} = 1/\sigma_{\text{farm}}^2$.

Figure 3.4 shows the raw data and the fitted regression lines from the MLE and Bayesian analyses. The two analysis procedures give similar fitted lines and the fitted regression lines for the different type of farms are parallel. This is consistent with the true relationship between the two types.




Figure 3.4: Plot of the lower bounds (star for farms of TYPE=1, dot for farms of TYPE=2) of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT FARM(R) with fitted regression lines from the MLE method (solid lines, upper line for farms of TYPE=1, bottom line for farms of TYPE=2) and the Bayesian analysis (dashed lines, upper line for farms of TYPE=1, bottom line for farms of TYPE=2). Values are jittered to avoid overplotting. Values of TYPE=1 were jittered to left of visit. Values of TYPE=2 were jittered to right of visit. Anti-microbial concentrations used were 0.25, 0.5, 1, 2, 4 mg/L.

Table 3.5 summarizes the true values and the estimates of the parameters. Because a simple MLE analysis (SAS PROC LIFEREG) ignores random effects, the farm effect is “confounded” with error, so the estimate of σ is inflated. The Bayesian analysis (WinBUGS) gives a closer estimate of the farm effect and a closer estimate of σ than the Maximum Likelihood method.

The standard errors of the MLEs are generally smaller than the Bayesian posterior standard errors, which is reasonable because Bayesian inference incorporates prior uncertainty and other sources of uncertainty.

The estimate of σ from the Bayesian analysis is still a poor estimate. When the simulated true logMIC is $\geq \log(4)$, the censoring implies that the lower bound is $\log(4)$ but the upper bound is ∞ . For example, from Figure 3.4 we can see when *visit* is 5, most lower bounds gather around $\log(4)$. Information about the data is lost. Consequently, it is difficult to get a precise estimate of σ .

	Intercept of type=2	Effect of TYPE of FARM	Effect of VISITS	Standard deviation of random FARM effect	Standard deviation of random noise
	β_0 (SE)	$\beta_{type_1} - \beta_{type_2}$ (SE)	β_{visit} (SE)	σ_{farm} (SE)	σ (SE)
True value	0.50	1.00	0.50	1.50	1.00
MLE	0.52 (0.11)	1.17 (0.10)	0.47 (0.04)	—	1.61 (0.06)
APE	0.60 (0.28)	1.14 (0.10)	0.49 (0.04)	1.21 (0.23)	1.34 (0.05)

Table 3.5: True values and estimates of parameters from model $[Y]=TYPE(C) VISIT FARM(R)$. Estimates are Maximum Likelihood Estimates (MLE) and Bayesian Average Posterior Estimates (APE). Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs.

3.5 VISIT linear, TYPE classification, FARM random effect (small)

In this simulation experiment, separate FARMS of two TYPES are repeatedly measured at each of five VISITS. For each farm type, there are 20 farms, and each farm has 10 replications. The total sample size is 2000. The farm effects are again random effects but the farm effects are now small relative to random noise. The same anti-microbial concentrations as in Section 3.1 are used.

The simplified syntax for the model of this experiment is

$$[Y] = \text{TYPE}(C) \text{ VISIT FARM}(R)$$

The true model is

$$y_i = 1.5 + 0.5\text{visit}_i + \text{farm.effect}_i + \epsilon_i \quad \text{for farms of type}=1$$

$$y_i = 0.5 + 0.5\text{visit}_i + \text{farm.effect}_i + \epsilon_i \quad \text{for farms of type}=2$$

$$\text{farm.effect}_i \sim N(0, 0.1^2)$$

$$\epsilon_i \sim N(0, 1^2)$$

The models used for the analysis are the same as in Section 3.4.

Figure 3.5 shows the raw data and the fitted regression lines from the MLE and Bayesian analyses. The two analysis procedures give almost the same fitted lines and the fitted regression lines for the different types of farms are parallel. This is consistent with the true relationship between the two types.

Table 3.6 summarizes the true values and the estimates of the parameters. When the farm random effect is small, as in this case, the MLE and Bayesian analyses give similar estimates. The Bayesian approach also gives a close estimate of the random effects. This is not unexpected as Bayesian models for data with small random effects should behave similarly to MLE models where the small random effects are ignored.

	Intercept of type=2	Effect of TYPE of FARM	Effect of VISITS	Standard deviation of random FARM effect	Standard deviation of random noise
	β_0 (SE)	$\beta_{\text{type}1} - \beta_{\text{type}2}$ (SE)	β_{visit} (SE)	σ_{farm} (SE)	σ (SE)
True value	0.50	1.00	0.50	0.10	1.00
MLE	0.56 (0.08)	0.99 (0.08)	0.48 (0.03)	—	1.01 (0.04)
APE	0.56 (0.09)	1.00 (0.08)	0.49 (0.03)	0.12 (0.04)	1.03 (0.04)

Table 3.6: True values and estimates of parameters from model $[Y]=\text{TYPE}(C) \text{ VISIT FARM}(R)$. Estimates are Maximum Likelihood Estimates (MLE) and Bayesian Average Posterior Estimates (APE). Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs.




Figure 3.5: Plot of the lower bounds (star for farms of TYPE=1, dot for farms of TYPE=2) of simulated data from the model $[Y]=\text{TYPE}(C)$ VISIT FARM(R) with fitted regression lines from the MLE method (solid lines, upper line for farms of TYPE=1, bottom line for farms of TYPE=2) and the Bayesian analysis (dashed lines, upper line for farms of TYPE=1, bottom line for farms of TYPE=2). Values are jittered to avoid overplotting. Values of TYPE=1 were jittered to left of visit. Values of TYPE=2 were jittered to right of visit. Anti-microbial concentrations used were 0.25, 0.5, 1, 2, 4 mg/L.

3.6 Summary of simulation results

The simulation results support our supposition that when the sample size is large, the maximum likelihood method and Bayesian analysis give similar estimates. The analysis methods also appear to be comparable for a variety of models. This implies that for different types of data, the two methods used are robust.

The simulation results also demonstrate that the Bayesian method can be used in the presence of random effects. When random effects are small relative to random noise, the MLE method also works well, because a model with small random effects will behave similarly to a model that ignores the random effects.

The simulation study still has some limitations. For example, it does not simulate a model that contains interaction terms between effects, nor does it simulate data with small sample sizes.

Chapter 4

Model building, selection and diagnostics

Any statistical model should reflect the experimental structure of the data. The procedure of finding a “good” model is as follows: model building → model selection → model assessment.

4.1 Model building

Model building starts with a set of candidate effects for the experiment. First, we need to determine which effects and types of relationships should be included in the model. This depends upon how the data are collected, information about likely explanatory variables obtained, and can be aided by some preliminary plots.

IdNumber	Farm	Visit	Strain	Type	AMOXIC
1-1-1-1	1	1	1	1	>16
1-1-2-1	1	1	1	2	>16
1-1-3-1	1	1	1	3	>16
1-1-4-1	1	1	1	4	>16
1-1-5-1	1	1	1	5	>16
1-1-6-1	1	1	1	6	>16
1-1-7-1	1	1	1	7	>16
1-1-8-1	1	1	1	8	>16
1-1-9-1	1	1	1	9	>16
1-1-10-1	1	1	1	10	>16
1-2-1-1	1	2	1	1	>16
1-2-2-1	1	2	2	2	>16
1-2-3-1	1	2	3	3	>16

Figure 4.1: Portion of the Avian data set

As an example, the Avian data set will be used to illustrate how to build a starting model. A portion of the raw data is shown in Figure 4.1.

The sample protocol for the Avian project is that five visits are made to each farm of two diet types; nine farms used a medicated commercial ration, while two other farms used a vegetarian, non-medicated diet. From each sample collected, 5 or 10 bacterial colonies or “strains” were selected for inoculation onto anti-microbial plates. The sample Id-number, Farm, Visit, Strain, Type, and MIC values were recorded.

Based on the sample protocol, *farm*, *visit*, *strain* and *type* may be potential explanatory variables. *Strain1* to *strain10* are replicated readings from same sample (i.e., pseudo-replicates). *Visit* refers to the sampling time, that is relative to the age of the animals. *Type* refers to the type of diet. *Farm* refers to where these samples come from. The last three variables are considered to be potential explanatory variables for the MIC value so they should be included in the model.

In this experiment we are only interested in the two specific diet types, so *type* is treated as a fixed effect. Similarly, *visit* is also pre-determined; it is also a fixed effect. *Farm* is generated from a large population of farms, the level of farms cannot be repeatedly chosen. If the experiment were to be repeated, the same levels of *visit* (age of animals) and *type* could be chosen while a new set of levels of farm would be chosen, so *farm* effect is treated as a random effect.

Figure 4.2 is a plot of the lower bound of logMIC of *CEFTIF* (Ceftiofur, one of the anti-microbial agents used in the Avian project) vs. *visit*, with a simple loess line (connecting the mean lower bound value of every visit) for different *types* of farms.

The plot reveals some useful things. The two simple loess lines do not appear to be parallel, which indicates that an interaction between *type* and *visit* may exist.

Figure 4.3 is a plot of the farm effects. It is drawn by connecting the median lower bound of the logMIC of each visit of each farm. From the plot, we can see some farms have higher median lower bound of logMIC than other farms. This implies that a farm effect may exist and may be large.

After examining the raw data and the plots of the data, it is clear that the initial model needs to include terms for the effects of *farm*, *visit* and *type*.




Figure 4.2: Plot of the lower bounds of $\log\text{MIC}_{\text{CEFTIF}}$ vs. visit (star for farms of $\text{TYPE}=0$, dot for farms of $\text{TYPE}=1$) with fitted loess lines (dashed line for farms of $\text{TYPE}=0$, solid line for farms of $\text{TYPE}=1$). Values are jittered to avoid overplotting. Values of $\text{TYPE}=1$ were jittered to left of visit; values of $\text{TYPE}=0$ were jittered to right of visit.




Figure 4.3: Plot of connecting median lower bounds of $\log\text{MIC.CEFTIF}$ of each visit of each farm vs. visit for different farms (each line represents a farm). Values are jittered to avoid overplotting.

4.2 Model selection

Once the base model is determined, there is still the question of whether other models are more suitable. For example, is there evidence that interaction terms are large? Is there evidence that the slope of the relationship between variables is different than zero? Are random effects large?

A natural way to compare models is to use a criterion based on a trade-off between the fit of the data to the model and the corresponding complexity of the model. Likelihood models and Bayesian models have different but similar criteria based on this principle.

4.2.1 Likelihood Model

Akaike's information criterion (AIC) is a model selection criterion for likelihood models (Burnham et al. 2002). This statistic trades off goodness of fit (measured by the maximized log-likelihood) against model complexity (measured by p , which is the number of parameters in each model under consideration). Comparisons between a number of possible models can be made on the basis of the statistic

$$AIC = -2 \times \log(\text{maximum likelihood}) + 2p$$

The AIC penalizes models for adding parameters. The model with the smallest AIC value is chosen to be the best model.

AICc is AIC with a second order correction for small sample sizes n . It is defined as:

$$AICc = AIC + \frac{2p(p+1)}{n-p-1}$$

Since AICc converges to AIC as n gets large, AICc should be employed regardless of sample size.

In practice, it is impossible to explain all relationships between variables using a single model. For example, consider a case where two models are fitted to some biological data. Model M1 may be statistically superior to model M2, but model M2 is more biologically meaningful than model M1. In this case, model averaging is introduced to generate model averaged estimates (Burnham and Anderson, 2002). We use $\Delta AICc$ to represent the difference in AICc between model and the model with the smallest AICc value.

Table 4.1 is an example of model averaging procedures. Suppose model M1, M2 and M3 are possible models of an experiment, and $n = 100$ is the sample size. Let's use M1, M2

and M3 to demonstrate the details of the calculations.

The process is as follows: fit model \rightarrow get $\log(\text{Maximum likelihood}) \rightarrow$ calculate AIC and AICc \rightarrow calculate Weight and Normalize weight \rightarrow calculate model averaged estimates.

$\log(\text{maximum likelihood})$ is obtained by using the MLE method. For example, if SAS PROC LIFEREG is used to analyze the data set, the program will automatically give the $\log(\text{maximum likelihood})$.

p is the number of parameters in the model.

The AIC is calculated by $AIC = -2 \times \log(\text{maximum likelihood}) + 2p$. For example, the AIC for M3 is $-2 \times (-210) + 2 \times 4 = 428.00$.

The AICc is calculated by $AICc = AIC + \frac{2p(p+1)}{n-p-1}$. The AICc for M3 is $428 + \frac{2 \times 4 \times (4+1)}{100-4-1} = 428.41$; The AICc for M1 and M2 were obtained similarly. Of the three models, M3 has the minimum AICc 428.41.

The ΔAIC is calculated by $AICc - \text{minimum} AICc$, which is $444.12 - 428.41 = 15.71$ for M1, $436.24 - 428.41 = 7.83$ for M2, and $428.41 - 428.41 = 0$ for M3.

The model averaged weight is calculated by $weight = \exp(-0.5 \times \Delta AIC) = \exp(-0.5 \times 15.71) = 0$ for M1, $\exp(-0.5 \times 7.83) = 0.02$ for M2, and $\exp(-0.5 \times 0) = 1$ for M3.

The normalized weight is calculated by $normalized\ weight = \frac{weight}{\sum(weight)} = \frac{0}{0.00+0.02+1.00} = 0$ for M1, $\frac{0.02}{0.00+0.02+1.00} = 0.02$ for M2, and $\frac{1.00}{0.00+0.02+1.00} = 0.98$ for M3.

From the output of the computer program, we can get the estimates for parameter S and the corresponding standard errors.

The model averaged estimate of S , which is \bar{S} , is calculated by $\sum estimate \times normalize\ weight$. For example, it would be $= 7 \times 0.98 + 6 \times 0.02 + 5 \times 0 = 6.98$.

The model averaged estimate of the standard error for S is the estimate of the unconditional standard error. It is defined as $\sum normalize\ weight \times \sqrt{var(S) + (S - \bar{S})^2}$. For this example, the calculation would be $= 0 \times \sqrt{1.00^2 + (5.00 - 6.98)^2} + 0.02 \times \sqrt{1.50^2 + (6.00 - 6.98)^2} + 0.98 \times \sqrt{2.00^2 + (7.00 - 6.98)^2} = 0 \times 2.22 + 0.02 \times 1.79 + 0.98 \times 2.00 = 0 + 0.04 + 1.96 = 2.00$.

4.2.2 Bayesian Model

Spiegelhalter et al. (2002) proposed a Bayesian model comparison criterion. The Deviance Information Criterion (DIC) is defined as

$$DIC = \text{“goodness of fit”} + \text{“complexity”}$$

Model	M1	M2	M3
Log(maximum likelihood)	-220.00	-215.00	-210.00
P	2	3	4
AIC	444.00	436.00	428.00
AICc	444.12	436.24	428.41
$\Delta\text{AIC}=\text{AICc}-\text{smallest AICc}$	15.71	7.83	0.00
Weight= $\exp(-0.5\Delta\text{AIC})$	0	0.02	1.00
Normalize Weight= $\frac{\text{weight}}{\text{sum}(\text{weight})}$	0	0.02	0.98
Estimate of S	5.00	6.00	7.00
SE of estimate	1.00	1.50	2.00
Model averaged estimate of S	S=6.98		
Unconditional standard error	SE=2.00		

Table 4.1: Example of model averaging

“Goodness of fit” is measured by the deviance, where deviance is defined as $D(\theta) = -2\log L(\text{data}|\theta)$, that is $-2\log(\text{likelihood})$.

Complexity is measured by an estimate of the effective number of parameters defined as $P_D = \bar{D} - D(\bar{\theta})$, i.e., the posterior mean deviance minus the deviance evaluated at the posterior mean of the parameters.

Sometimes P_D is a fractional number, because \bar{D} and $D(\bar{\theta})$ are not always integers. In a random effect model, P_D can be a negative number. This happens when there is either a substantial conflict between the priors and the data or when the posterior distribution for a parameter is extremely asymmetric.

The DIC is defined analogously to AIC as

$$DIC = \bar{D} + P_D$$

The model with the smallest DIC is chosen as the model that would best predict a replicate dataset with the same structure as that currently observed. The model that has the minimum DIC will make the best short-term predictions. The DIC has the same spirit as Akaike’s criterion.

It is difficult to say what would constitute an important difference in DIC. Very roughly, differences of more than 10 might definitely rule out the model with the higher DIC, differences between 5 and 10 are substantial, but if the difference in DIC is, say, less than 5, and

the models make very different inferences, then it could be misleading just to report the model with the lowest DIC. In this case, estimates from both models should be reported.

Bayes factor (BF) is often used for comparing two models. It is defined as the ratio of the marginal likelihood under one model to the marginal likelihood under a second model. If we observe Y_{obs} and have 2 models, M_1 and M_2 , then the Bayes Factor is

$$BF = \frac{p(Y_{obs}|M_1)}{p(Y_{obs}|M_2)} = \frac{\int p(\theta_1|M_1)p(Y_{obs}|\theta_1, M_1)d\theta_1}{\int p(\theta_2|M_2)p(Y_{obs}|\theta_2, M_2)d\theta_2}$$

where $p(Y_{obs}|M_i)$ is called the marginal likelihood for model i . This provides the relative weight of evidence for model M_1 compared to model M_2 . Roughly, a BF of more than 150 gives strong evidence to keep M_1 , a BF of between 12 and 150 is substantial, a BF of between 3 and 12 gives some evidence for M_1 against M_2 , and a BF of between 1 and 3 gives weak evidence for M_1 against M_2 .

4.3 Model checking

Once we have accomplished the first two steps of a data analysis — constructing models and selecting among models, we still need to assess the fit of the model to the data. It is difficult to include all of one's knowledge about a problem in one model, so it is wise to investigate what aspects of reality are not captured by the model.

4.3.1 Likelihood model

When searching for a parametric model that fits the data well, there are many methods of model checking. Graphical displays are the most useful method to check a model's appropriateness.

An important assumption in regression analysis is that the residual errors, the deviations of the observed values of the response from their expectations, are independent and identically distributed with a mean of zero. This assumption can be verified by viewing residual plots. Independence among observations is often checked by plotting residuals vs. the order of observation. If observations are correlated, we expect to find residuals exhibiting dependence over time.

Residual plots for censored data are constructed similarly as for non-censored data. The only difference is that we plot a residual interval instead of a residual. A residual interval is

an interval with a lower bound equal to the difference between the predicted data and data's lower bound, and an upper bound equal to the difference between the predicted data and the data's upper bound. For example, in our experiment the lower bound of the residual interval equals the lower bound of the logMIC minus prediction, and the upper bound of the residual interval equals the upper bound of the logMIC minus prediction. We know the true value is within an interval, so the true residual is also within an interval.

4.3.2 Bayesian model

There are several diagnostics used for Bayesian model checking: chain convergence, Bayesian p-values and posterior predictive checking.

Chain convergence refers to the idea that eventually the Gibbs Sampler or other MCMC technique will reach a stationary distribution. After the model has converged, samples from the conditional distributions are used to summarize the posterior distribution of parameters of interest.

Trace plots of parameter values against iteration numbers can be very useful in assessing convergence. A trace plot indicates whether the chain has converged to its stationary distribution. If convergence has not been reached, a longer burn-in period is often advisable. Those simulated values which are generated before the chain reaches its stationary distribution should be discarded. A trace plot can also indicate whether the chain is mixing well. A chain that mixes well traverses its posterior space rapidly, and it can jump from one remote region of the posterior space to another in relatively few steps.

The Brooks-Gelman-Rubin (BGR) statistic (Brooks and Gelman, 1998) is another convergence diagnostic. It uses parallel chains with dispersed initial values to test whether they all converge to the same target distribution. Failure could indicate the presence of a multi-mode posterior distribution (different chains converging to different local modes) or the need to run a longer chain (burn-in is yet to be completed).

The BGR is computed as follows. The MCMC algorithm is run from J separate, equal-length chains ($J \geq 2$) with starting values dispersed over the support of the target density. Let L denote the length of each chain after discarding D burn-in iterates. For example, suppose that the variable of interest is X , and its value at the t th iteration of the j th chain is $x_j^{(t)}$. Thus, for the j th chain, the D values $x_j^{(0)}, \dots, x_j^{(D-1)}$ are discarded and the L values

$x_j^{(D)}, \dots, x_j^{(D+L-1)}$ are retained. Then

$$BGR = \frac{\frac{L-1}{L}W + \frac{1}{L}B}{W}$$

where W represent the mean of the within-chain estimated variances, B represent the between-chain variance. If all the chains are stationary, then both the numerator and the denominator should estimate the marginal variance of X . As $L \rightarrow \infty$, $\sqrt{BGR} \rightarrow 1$. In practice, when $\sqrt{BGR} \leq 1.2$ we think D and L are acceptable.

In a Bayesian context, a posterior p-value is the probability, given the data, that a future observation is more extreme (as measured by some test statistics) than the data. If the model fits, then replicated data generated under the model should look similar to observed data. Our basic technique for checking the fit of a model to the data is to draw simulated values from the posterior predictive distribution of replicated data and compare these samples to the observed data. Any systematic differences between the simulations and the data indicates potential failings of the model. A discrepancy function for the data is computed for each simulated set and for the real data at the amount values of the parameters, and then these two values are plotted against each other. If the model fits well, the points will tend to fall along a reference line that has an intercept of 0 and a slope equal to 1.

Chapter 5

Real Examples

5.1 Avian data set (Avian plate)

In this project, 11 barns on 9 farms were visited. Fecal samples were collected on Day 0 (day the chicks arrived), and on approximately Day 10, Day 25, and Day 40 (just prior to slaughter). In addition, a visit was made to the slaughter plant at the time the flock (barn floor) of interest was being slaughtered, and carcass swab samples were collected. Two barns used a vegetarian, non-medicated diet, while the 9 remaining barns used a medicated commercial ration. From each farm visit, two pooled samples of 5-10 fresh fecal specimens were collected from one floor of one barn on the farm (2 barns on farms where a vegetarian diet was used). From each slaughter plant visit, two pooled samples of 10 carcass swabs were collected. The bacteria colonies from each sample were inoculated onto Avian and National Antimicrobial Resistance Monitoring System (NARMS) sensititre plates. The Avian plate data set is the outcome of samples inoculated into Avian sensititre plates; the NARMS plate data set is the outcome of samples inoculated into NARMS sensititre plates.

After samples were collected, the bacteria colonies from each sample were tested against an array of 18 anti-microbial agents on Avian sensititre plates (Amoxicillin, Ceftiofur, Clindamycin, Enrofloxacin, Erythromycin, Gentamycin, Neomycin, Novobiocin, Oxytetracycline, Penicillin, Sarafloxacin, Spectinomycin, Streptomycin, Sulfadimethoxine, Sulfathiazole, Tetracycline, Trimethoprim/Sulfamethoxazole, Tylosin Tartrate) and reported MIC values were given. We will use the CEFTIF (Abbreviation of Ceftiofur) value as an example to show how to analyse the data. The dilution series is 0.25, 0.5, 1, 2, 4 mg/L, representing the different concentrations of the anti-microbial agent CEFTIF. Figure 5.1 presents some

sample data from this experiment.

IdNumber	Farm	Visit	Strain	Type	AMOXIC	CEFTIF	CLINDA	ENROFL	ERYTH
1-1-1-1	1	1	1	1	>16	0.5	>4	<=0.12	>4
1-1-2-1	1	1	2	1	>16	0.5	>4	<=0.12	>4
1-1-3-1	1	1	3	1	>16	0.5	>4	<=0.12	>4
1-1-4-1	1	1	4	1	>16	0.5	>4	<=0.12	>4
1-1-5-1	1	1	5	1	>16	>4	>4	<=0.12	>4
1-1-6-1	1	1	6	1	>16	0.5	>4	<=0.12	1
1-1-7-1	1	1	7	1	>16	<=0.25	>4	<=0.12	1
1-1-8-1	1	1	8	1	>16	>4	>4	<=0.12	>4
1-1-9-1	1	1	9	1	>16	>4	>4	<=0.12	>4
1-1-10-1	1	1	10	1	>16	<=0.25	>4	<=0.12	>4

Figure 5.1: Part of the raw data for the Avian data set. Column CEFTIF lists the reported MIC values of anti-microbial CEFTIF.

Each of the reported MIC values are converted to a pair of numbers of either the form $(\log(0), xx)$, $(xxx/2, xxx)$, or $(xxxx, \infty)$, representing the intervals that contain the \log MIC values.

As outlined in Chapter 4, we decided that *type*, *visit* and *farm* should be included in model.

Figure 5.2 presents a plot of \log MIC (lower bound) vs. visit for different farm types with simple loess lines.




Figure 5.2: Plot of the lower bounds of $\log\text{MIC_CEFTIF}$ vs. visit (star for farms of $\text{TYPE}=0$, which used the vegetarian non-medicated diet; dot for farms of $\text{TYPE}=1$, which used the medicated commercial ration), with fitted loess lines (dashed line for farms of $\text{TYPE}=0$, solid line for farms of $\text{TYPE}=1$). Values are jittered to avoid overplotting. Values of $\text{TYPE}=1$ were jittered to left of visit. Values of $\text{TYPE}=0$ were jittered to right of visit.

When the likelihood methods are used, we cannot include farm as a random effect in our model. We will ignore the farm effect first, and then include it in our model when we do a Bayesian model analysis.

Several models, named as M1 to M7, are fitted as summarized in Table 5.1. In terms of the AICc criterion, support for model $[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE}$ overwhelm other models, and this model has dominant weight. Model M1 will be used to do further analysis.

Model	Log(maximum likelihood)	P	AICc	ΔAIC	Normalized Weight
M1:[Y]=Visit(C) Type(C) Visit \times Type	-803.43	11	1629.44	0.00	0.97
M2:[Y]=Visit Type(C)	-815.09	4	1638.27	8.83	0.01
M3:[Y]=Visit(C) Type(C)	-812.12	7	1638.49	9.05	0.01
M4:[Y]=Visit Type(C) Visit \times Type	-814.70	5	1639.53	10.10	0.01
M5:[Y]=Type(C)	-817.29	3	1640.63	11.20	0.00
M6:[Y]=Visit(C)	-815.51	6	1643.20	13.76	0.00
M7:[Y]=Visit	-818.76	3	1643.57	14.13	0.00

Table 5.1: Results of fitting different models to the Avian-CEFTIF data using maximum likelihood estimation from PROC LIFEREG (SAS).

Figure 5.3 is the residual plot of M1. Residuals are scattered around 0 and no pattern is apparent in the plot. This indicates that there is no evidence of lack of fit of model M1.




Figure 5.3: Residual plot of model $[Y] = \text{VISIT}(C) \text{TYPE}(C) \text{VISIT} \times \text{TYPE}$. Symbol “X” on the graph indicates truncated residuals which occur when the upper bound of $\log\text{MIC}$ is $\log(\infty)$. Similarly, when the lower bound of $\log\text{MIC}$ is $\log(0)$, the lower bound of the residual interval = $-\infty$. This is also represented by a “X”. Values are jittered to avoid overplotting.

Two Bayesian models with and without farm as a random effect are also used to examine the data.

- Bayesian model without farm effect

$$[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE}$$

The prior distributions are: $\beta_0 \sim N(0, 10)$, $\beta.type \sim N(0, 10)$, $\beta.visit \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$.

- Bayesian model with farm as a random effect

$$[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE} \text{ FARM}(R)$$

The prior for farm is: $\tau_{farm} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{farm} = 1/\sigma_{farm}^2$. Other prior distributions are the same as above.

Table 5.2 presents the estimates for the MLE and Bayesian models. Estimates are similar even when different analysis methods were used because the variance of random effects is small relative to the random error. However, standard errors are somewhat higher with the random effect in the model.

From Table 5.3, the DIC is smaller when farm is included in the model than when it doesn't, indicating that

$$[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE} \text{ FARM}(R)$$

is the more “useful” model. Also from Table 5.4, this model has the smallest DIC value among all 14 possible models and BGR statistic close to 1 means different chains converge to same target distribution. Summary statistics from this model will be used to draw inferences for the quantities of interest.

The summary statistics from the Bayesian model that includes farm random effects are as follows: The difference in mean logMIC between the two diet types (farms of $type=1$ used a medicated commercial ration, while farms of $type=0$ used a vegetarian, non-medicated diet) is not statistically significant. The difference of -0.18 is small compared to the SD of 0.79. In anti-log scale, the ratio of the mean logMIC between the two diet types is $\exp(-0.18)=0.84$. The MLE estimate of -0.14 is small compared to the SE of 0.62. The same conclusion is

	MLE	Bayesian excluding farm effect	Bayesian including random farm effect
β_0 (intercept) (SE)	-0.62 (0.24)	-0.58 (0.24)	-0.56 (0.34)
$\beta_{type_0} - \beta_{type_1}$ (SE)	-0.14 (0.62)	-0.30 (0.59)	-0.18 (0.79)
$\beta_{visit_1} - \beta_{visit_5}$ (SE)	0.39 (0.38)	0.34 (0.38)	0.35 (0.38)
$\beta_{visit_2} - \beta_{visit_5}$ (SE)	0.88 (0.33)	0.85 (0.33)	0.83 (0.31)
$\beta_{visit_3} - \beta_{visit_5}$ (SE)	0.81 (0.33)	0.76 (0.33)	0.72 (0.32)
$\beta_{visit_4} - \beta_{visit_5}$ (SE)	0.21 (0.33)	0.16 (0.32)	0.14 (0.31)
$\mu_{visit_1type_1}$ (SE)	-0.22 (0.30)	-0.23 (0.29)	-0.21 (0.38)
$\mu_{visit_2type_1}$ (SE)	0.26 (0.22)	0.27 (0.22)	0.27 (0.32)
$\mu_{visit_3type_1}$ (SE)	0.20 (0.22)	0.18 (0.22)	0.16 (0.33)
$\mu_{visit_4type_1}$ (SE)	-0.41 (0.22)	-0.42 (0.22)	-0.42 (0.32)
$\mu_{visit_5type_1}$ (SE)	-0.62 (0.24)	-0.58 (0.24)	-0.56 (0.34)
$\mu_{visit_1type_0}$ (SE)	-1.01 (0.53)	-1.01 (0.53)	-1.13 (0.73)
$\mu_{visit_2type_0}$ (SE)	0.18 (0.47)	0.19 (0.50)	0.18 (0.69)
$\mu_{visit_3type_0}$ (SE)	-2.90 (0.66)	-2.85 (0.64)	-2.85 (0.82)
$\mu_{visit_4type_0}$ (SE)	-0.40 (0.45)	-0.38 (0.45)	-0.41 (0.68)
$\mu_{visit_5type_0}$ (SE)	-0.76 (0.57)	-0.88 (0.54)	-0.74 (0.73)
$\sigma_{residual}$ (SE)	1.94 (0.10)	2.00 (0.11)	1.88 (0.10)
σ_{farm} (SE)			0.69 (0.21)

Table 5.2: Estimates for the Avian data from the MLE and two Bayesian analyses. Numbers in brackets are the standard errors for MLE or standard deviations for Bayesian APE. $\beta_{type_0} - \beta_{type_1}$ is the difference in mean logMIC between different types of operation. $\beta_{visit_i} - \beta_{visit_5}$ is the difference in mean logMIC between $visit_i$ and $visit_5$. $\mu_{visit_i type_j}$ is the mean logMIC at the i^{th} visit to a farm of j^{th} diet type.

drawn: there is no strong evidence to say that the mean logMIC values of the two diet types are different when averaged over visits.

The difference in mean logMIC between visit 2 and visit 5 is statistically significant, as is the difference in mean logMIC between visit 3 and visit 5. This means there is evidence that the mean logMIC value of visit 2 and visit 5 are different; the mean logMIC value of visit 3 and visit 5 are different. However, there is no strong evidence that the mean logMIC value of visit 1 and visit 5 are different, and that the mean logMIC value of visit 4 and visit 5 are different.

When we compare the mean logMIC of visit 2 and visit 3, that is $\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_3} = (\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_5}) - (\widehat{\beta}_{visit_3} - \widehat{\beta}_{visit_5}) = 0.11$, the standard deviation is too complex to compute by hand because of the correlation between estimates, but is easily calculated by the computer, which gives an SD of 0.30. From the MLE output, the $\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_3} = 0.07$, the $SE = \left((se_{\beta_{visit_2} - \beta_{visit_5}})^2 + (se_{\beta_{visit_3} - \beta_{visit_5}})^2 + 2cov(\beta_{visit_2} - \beta_{visit_5}, \beta_{visit_3} - \beta_{visit_5}) \right)^{\frac{1}{2}} =$

Model	P_D	DIC	Δ DIC
Model without farm effect	10.60	1628.50	31.70
Model with farm effect	18.40	1596.80	0.00

Table 5.3: DIC value for model $[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE FARM}(R)$ with and without farm random effect

Model	DIC	P_D	Max(BGR)	Δ DIC
MF1: $[Y]=\text{Visit}(C) \text{ Type}(C) \text{ Visit}*\text{Type Farm}(R)$	1596.84	18.40	1.03	0.00
MF3: $[Y]=\text{Visit}(C) \text{ Type}(C) \text{ Farm}(R)$	1608.24	14.25	1.02	11.40
MF6: $[Y]=\text{Visit}(C) \text{ Farm}(R)$	1608.69	14.42	1.02	11.85
MF7: $[Y]=\text{Visit Farm}(R)$	1609.32	11.40	1.02	12.48
MF4: $[Y]=\text{Visit Type}(C) \text{ Visit}*\text{Type Farm}(R)$	1609.38	12.19	1.03	12.54
MF2: $[Y]=\text{Visit Type}(C) \text{ Farm}(R)$	1609.54	11.54	1.04	12.70
MF5: $[Y]=\text{Type}(C) \text{ Farm}(R)$	1612.05	10.68	1.03	15.21
M1: $[Y]=\text{Visit}(C) \text{ Type}(C) \text{ Visit}*\text{Type}$	1628.45	10.63	1.04	31.61
M2: $[Y]=\text{Visit Type}(C)$	1638.17	3.98	1.01	41.33
M3: $[Y]=\text{Visit}(C) \text{ Type}(C)$	1638.46	7.06	1.02	41.62
M4: $[Y]=\text{Visit Type}(C) \text{ Visit}*\text{Type}$	1639.29	4.92	1.03	42.45
M5: $[Y]=\text{Type}(C)$	1640.76	3.08	1.02	43.92
M6: $[Y]=\text{Visit}(C)$	1643.17	6.04	1.02	46.33
M7: $[Y]=\text{Visit}$	1643.24	2.85	1.02	46.40

Table 5.4: DIC value for all possible models

$\left((0.33)^2 + (0.33)^2 + 2 \times 0.06 \right)^{\frac{1}{2}} = 0.58$. The estimate of 0.07 is small compared to the SE of 0.58. The same conclusion is drawn: there is no strong evidence to say that the mean logMIC values of visit 2 and visit 3 are different from each other.

The variance of the farm effect 0.52 (0.35) is small compared to the variance of random noise 3.56 (0.39).

The estimates from the MLE method have reported SE for some fixed effects that are smaller relative to the standard deviation of the Bayesian estimates. This is because the estimates of some fixed effects, such as type, are obtained by comparison between farms. When estimating the fixed effects, the farm random effects will be subtracted, and the variance of farm random effects will not be counted in the calculation of standard error. The mean values, such as $\mu_{\text{visit}_i, \text{type}_j}$, are obtained by averaging among farms. In this case,

the variance of random effects is small compared to the random noise, but the Bayesian method still needs to incorporate prior uncertainty and other sources of uncertainty.

5.2 Avian data set (NARMS plate)

The sample protocol is same as the Avian plate data set. Samples are tested against an array of 16 anti-microbial agents (Amikacin, Amoxicillin/Clavulanic Acid, Ampicillin, Cefoxitin, Ceftiofur, Ceftriaxone, Cephalothin, Chloramphenicol, Ciprofloxacin, Gentamycin, Kanamycin, Nalidixic Acid, Streptomycin, Sulfamethoxazole, Tetracycline, Trimethoprim/Sulfamethoxazole) and the reported MIC values are given. We will use the CEFTIF (Abbreviation of Ceftiofur) values to show how to analyse the data. The dilution series is 0.12, 0.25, 0.5, 1, 2, 4, 8 mg/L, representing the different concentrations of the anti-microbial agent CEFTIF. Figure 5.4 presents a portion of the raw data.

IdNumber	Type	Farm	Visit	Strain	AMIKAC	AMOCLA	AMPICI	CEFOXI	CEFTIF
1-1-1-1	1	1	1	1	<=0.5	2	>32	2	1
1-1-2-1	1	1	1	2	<=0.5	16	>32	>16	0.5
1-1-3-1	1	1	1	3	<=0.5	2	>32	4	0.5
1-1-4-1	1	1	1	4	<=0.5	2	>32	2	0.5
1-1-5-1	1	1	1	5	1	32	>32	>16	8
1-1-6-1	1	1	1	6	1	16	>32	>16	0.5
1-1-7-1	1	1	1	7	<=0.5	16	>32	16	0.25
1-1-8-1	1	1	1	8	1	>32	>32	>16	8
1-1-9-1	1	1	1	9	<=0.5	>32	>32	>16	8
1-1-10-1	1	1	1	10	<=0.5	16	32	>16	0.25

Figure 5.4: Part of raw data for the Avian (NARMS plate) data set. Column CEFTIF lists the reported MIC values.

Each of the reported MIC values are converted to a pair of numbers of either the form $(\log(0), xx)$, $(xxx/2, xxx)$, or $(xxxx, \infty)$, representing the intervals that contain the $\log\text{MIC}$ values.

Figure 5.5 presents a plot of $\log\text{MIC}$ (lower bound) vs. *visit* for different farm *types* with simple loess lines. The two simple loess lines do not appear to be parallel, which indicates that an interaction between *type* and *visit* may exist. Farm effects are shown in Figure 5.6; we can see some farms have higher median lower bounds of $\log\text{MIC}$ than other farms. This means that a large farm effect may exist.

After examining the raw data and the plot, the initial model needs to include terms for the effects of *farm*, *visit* and *type*.




Figure 5.5: Plot of the lower bounds of $\log\text{MIC}_{\text{CEFTIF}}$ vs. visit (star for farms of $\text{TYPE}=0$, which used a vegetarian non-medicated diet; dot for farms of $\text{TYPE}=1$, which used a medicated commercial ration), with fitted loess lines (dashed line for farms of $\text{TYPE}=0$, solid line for farms of $\text{TYPE}=1$). Values are jittered to avoid overplotting. Values of $\text{TYPE}=1$ were jittered to left of visit. Values of $\text{TYPE}=0$ were jittered to right of visit.




Figure 5.6: Plot of connecting median lower bounds of logMIC.CEFTIF of each visit of each farm vs. visit for different farm (each line represents a farm). Values are jittered to avoid overplotting.

We follow the same procedure as before. Based on Section 4, we decide that *type* and *visit* should be included in the model. Several models are fitted as summarized in Table 5.5. In terms of the AICc criterion, support for model $[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE}$ overwhelm other models, and this model has the dominant weight. It will be used to do further analysis.

Model	Log(maximum likelihood)	P	AICc	ΔAIC	Normalized Weight
M1:[Y]=Visit(C) Type(C) Visit \times Type	-979.88	11	1982.34	0.00	0.95
M2:[Y]=Visit(C) Type(C)	-988.14	7	1990.53	8.19	0.02
M3:[Y]=Visit Type(C) Visit \times Type	-990.22	5	1990.56	8.23	0.02
M4:[Y]=Visit Type(C)	-991.44	4	1990.96	8.62	0.01
M5:[Y]=Visit(C)	-990.60	6	1993.37	11.04	0.00
M6:[Y]=Type(C)	-993.68	3	1993.40	11.07	0.00
M7:[Y]=Visit	-994.03	3	1994.11	11.78	0.00

Table 5.5: Results of fitting different models to the Avian-CEFTIF data using maximum likelihood estimation from PROC LIFEREG (SAS).

Figure 5.7 is the residual plot of model M1. Residuals are scattered around 0 and no pattern is apparent in the plot. This indicates that there is no evidence of lack of fit of model M1.




Figure 5.7: Residual plot of model $[Y] = \text{VISIT}(C) \text{TYPE}(C) \text{VISIT} \times \text{TYPE}$. The symbol “X” on the graph indicates truncated residuals as explained in Figure 5.3. Values are jittered to avoid overplotting.

Two Bayesian models with and without farm as a random effect are also used to examine the data.

- Bayesian model without farm effect

$$[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE}$$

The prior distributions are: $\beta_0 \sim N(0, 10)$, $\beta.type \sim N(0, 10)$, $\beta.visit \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$.

- Bayesian model with farm as a random effect

$$[Y] = \text{VISIT}(C) \text{ TYPE}(C) \text{ VISIT} \times \text{TYPE} \text{ FARM}(\text{TYPE})(R)$$

The prior for farm is: $\tau_{farm} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{farm} = 1/\sigma_{farm}^2$. Other prior distributions are the same as above.

Table 5.6 presents the estimates for these models. Estimates are similar even when different analysis methods were used because the variance of the farm random effects ($0.54^2 = 0.29$) is small relative to the residual variance ($1.46^2 = 2.13$).

From Table 5.7, the DIC is smaller when farm is included in the model than when it doesn't, indicating that

$$[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE} \text{ FARM}(\text{TYPE})(R)$$

is the more “useful” model. Summary statistics from this model will be used to draw inferences for the quantities of interest.

The summary statistics from the Bayesian model that includes random farm effects are as follows. The difference in mean logMIC between the two diet types (farms of $type=1$ used a medicated commercial ration, farms of $type=0$ used a vegetarian, non-medicated diet) is not statistically significant. There is no evidence to say the mean logMIC values of the two diet types are different. The estimate of 0.41 is small compared to the SE of 0.61. The MLE analysis leads to same conclusion because the estimate of 0.44 is small compared to the SE of 0.46.

The difference in mean logMIC between visit 2 and visit 5 is statistically significant, as is the difference in mean logMIC between visit 3 and visit 5. This means there is evidence

	MLE	Bayesian excluding random farm effect	Bayesian including random farm effect
β_0 (intercept) (SE)	-0.46 (0.18)	-0.44 (0.17)	-0.43 (0.25)
$\beta_{type_0} - \beta_{type_1}$ (SE)	0.44 (0.46)	0.36 (0.42)	0.41 (0.61)
$\beta_{visit_1} - \beta_{visit_5}$ (SE)	0.35 (0.28)	0.33 (0.28)	0.33 (0.27)
$\beta_{visit_2} - \beta_{visit_5}$ (SE)	0.70 (0.24)	0.68 (0.24)	0.68 (0.23)
$\beta_{visit_3} - \beta_{visit_5}$ (SE)	0.53 (0.24)	0.51 (0.24)	0.47 (0.24)
$\beta_{visit_4} - \beta_{visit_5}$ (SE)	0.10 (0.24)	0.09 (0.24)	0.07 (0.23)
$\mu_{visit_1type_1}$ (SE)	-0.10 (0.22)	-0.11 (0.22)	-0.10 (0.29)
$\mu_{visit_2type_1}$ (SE)	0.24 (0.16)	0.24 (0.16)	0.25 (0.25)
$\mu_{visit_3type_1}$ (SE)	0.08 (0.17)	0.08 (0.17)	0.04 (0.25)
$\mu_{visit_4type_1}$ (SE)	-0.35 (0.16)	-0.35 (0.16)	-0.36 (0.24)
$\mu_{visit_5type_1}$ (SE)	-0.46 (0.18)	-0.44 (0.17)	-0.43 (0.25)
$\mu_{visit_1type_0}$ (SE)	-0.81 (0.39)	-0.78 (0.39)	-0.84 (0.54)
$\mu_{visit_2type_0}$ (SE)	0.20 (0.35)	0.22 (0.35)	0.21 (0.52)
$\mu_{visit_3type_0}$ (SE)	-1.80 (0.41)	-1.78 (0.40)	-1.79 (0.55)
$\mu_{visit_4type_0}$ (SE)	-0.45 (0.34)	-0.45 (0.34)	-0.43 (0.52)
$\mu_{visit_5type_0}$ (SE)	-0.02 (0.42)	-0.08 (0.40)	-0.03 (0.56)
$\sigma_{residual}$ (SE)	1.51 (0.06)	1.53 (0.06)	1.46 (0.06)
σ_{farm} (SE)			0.54 (0.17)

Table 5.6: Estimates for the Narms data from the MLE and two Bayesian analyses. Numbers in brackets are the standard errors for MLE or standard deviations for Bayesian APE. $\beta_{type_0} - \beta_{type_1}$ is the difference in mean logMIC between different types of operation. $\beta_{visit_i} - \beta_{visit_5}$ is the difference in mean logMIC between $visit_i$ and $visit_5$. $\mu_{visit_i type_j}$ is the mean logMIC at the i^{th} visit to a farm of j^{th} diet type.

that the mean logMIC value of visit 2 and visit 5 are different; the mean logMIC value of visit 3 and visit 5 are different.

The difference in mean logMIC between visit 1 and visit 5, and between visit 4 and visit 5 are not statistically significant. This means there is no strong evidence that the mean logMIC value of visit 1 and visit 5 are different; and that the mean logMIC value of visit 4 and visit 5 are different.

When we compare the mean logMIC values of visit 2 and visit 3, that is $\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_3} = (\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_5}) - (\widehat{\beta}_{visit_3} - \widehat{\beta}_{visit_5}) = 0.21$, the standard error is too complex to compute by hand because of correlations between estimates, but can be calculated by the computer program and the reported SE is 0.23. From the MLE output, the $\widehat{\beta}_{visit_2} - \widehat{\beta}_{visit_3} = 0.17$, the $SE = \left((se_{\beta_{visit_2} - \beta_{visit_5}})^2 + (se_{\beta_{visit_3} - \beta_{visit_5}})^2 + 2cov(\beta_{visit_2} - \beta_{visit_5}, \beta_{visit_3} - \beta_{visit_5}) \right)^{\frac{1}{2}} = \left((0.24)^2 + (0.24)^2 + 2 \times 0.032 \right)^{\frac{1}{2}} = 0.42$. The estimate of 0.17 is small compared to the SE

Model	P_D	DIC
Model without farm effect	10.50	1981.00
Model with farm effect	18.70	1946.00

Table 5.7: DIC values for model $[Y] = \text{TYPE}(C) \text{ VISIT}(C) \text{ VISIT} \times \text{TYPE FARM}(R)$ with and without farm random effect.

of 0.42. The same conclusion is drawn: there is no strong evidence to say that the mean logMIC values of visit 2 and visit 3 are different from each other.

The mean logMIC values from the MLE method have reported SEs that are small relative to standard deviations of the Bayesian model when the random farm effect is present. The reason is same as in Avian data set.

5.3 Dairy data set (*Campylobacter*)

The sample protocol for this project was that 51 dairy calves on 26 farms were enrolled in the Dairy project. Six veterinarians made four visits to the farms to collect samples when each enrolled calf was 1-7 days of age, 14-28 days of age, 90-120 days of age, and 6-8 months of age. At each visit, one individual fecal sample was taken from the enrolled calf, one pooled fecal sample from five calves of the same age, and one pooled fecal sample from five cows. Additionally, one sample was taken from the mother of enrolled calf at the first visit, and one liquid calf feed specimen was taken at the first, second, and third visits. A one-page questionnaire was completed at each of the four visits.

Bacteria were cultured from the fecal and feed samples. *Campylobacter* sp. were tested using AMR E-Test strips against 8 anti-microbial agents (Azithromycin, Chloramphenicol, Ciprofloxacin, Erythromycin, Gentamycin, Nalidixic Acid, Tetracycline, Clindamycin) and MIC values were given. We will use Azithromycin to show how to analyse the data. The concentrations of the anti-microbial agent Azithromycin are 0.016, 0.023, 0.032, 0.047, 0.064, 0.094, 0.125, 0.19, 0.25, 0.38, 0.50, 0.75, 1.0, 1.5, 2, 3, 4, 6, 8, 12, 16, 24, 32, 48, 64, 96, 128, 192, 256 mg/L. Figure 5.8 presents a portion of the raw data.

Each of the reported MIC values are converted to a pair of numbers that represents the lower bound and upper bound of the true MIC values.

Based on the sample protocol, *vet*, *visit*, *date*, *farm*, and *group* may be potential

Vet	visit	date	farm	sample	coli/jejuni	Azithromycin-2
SC	1	20-Mar-06	F4	C412-C1	jejuni	0.064
DC	1	6-Feb-06	F2	C45-C1	jejuni	0.125
DC	1	6-Feb-06	F2	C45-C2	jejuni	0.25
BG	1	9-Mar-06	F4	C616-C1	jejuni	0.25
BG	1	9-Mar-06	F4	C616-C2	jejuni	0.38
KL	1	31-Jan-06	F1	Co722/723-C1	jejuni	0.094
KL	1	31-Jan-06	F1	Co722/723-C2	jejuni	0.25
KL	1	10-Feb-06	F3	D1201-C1	jejuni	0.125
KL	1	10-Feb-06	F3	D1201-C2	jejuni	0.25

Figure 5.8: Part of the raw data of the Dairy (*Campylobacter*) data set. *Vet* refers to the names of vets, *visit* refers to the time of a visit, *date* refers to the exact date of visit, *farm* refers to where a sample comes from, *sample* indicates from which animal group or pooled animals group the sample is taken. Column Azithromycin-2 lists the reported MIC values. In order to give a clear explanation, *sample* will be called as *group* and *sampling* will be called as *visit* in the following part.

explanatory variables. *Vet* refers to the veterinarian. *Visit* is the number of the visit. *Date* is the exact date of each visit. *Farm* is the name of the farm visited. Because farm F1 visited by veterinarian 1 is not same as farm F1 visited by veterinarian 2, this implies that farms are nested with veterinarians. The variable *group* is the type and ID-number of animal sampled; this code is complex because it is not in the form of a consistent encoding scheme. Based on *group*, animals are grouped into five types: enrolled calf, dam (mother of enrolled calf), pooled calf (pooled sample from five calves of the same age), pooled cow (pooled sample from five cows), and liquid calf feed. Unfortunately, there are only 2 observations for liquid calf feed, so this set of data was deleted. *Visit*, *group* and *farm* are considered most likely to be related to the MIC value so they should be included in the model. The variable *date* has a similar measure to *visit* so it is not included in the model. With only six veterinarians, *vet* is not considered to have a big effect. So *Vet* and *date* are dropped from the model.

In this experiment, we are interested in the specific five samples from the animals, so *group* is treated as a fixed effect. Similarly, *visit* is also pre-determined; it is a fixed effect too. *Farm* is generated from a large population of farms, and the level of farms cannot be repeatedly chosen. If the experiment were to be repeated, the same levels of *visit* and *group* could be chosen while a new set of levels of farm would be chosen, so *farm* effect is treated as a random effect.

Figure 5.9 presents a plot of logMIC (lower bound) vs. *visit* for different *groups* with simple loess lines.




Figure 5.9: Plot of the lower bounds of logMIC-Azithromycin vs. visit (star for group of calf, dot for sample from mother of enrolled calf, circle for sample of pooled calf, triangle for sample of pooled cow) with fitted lines for different samples. Values are jittered to avoid overplotting.

The plot reveals some useful things. The simple loess line of $group=2$ (dam) seems to diverge from the others, and the sample protocol indicates that $group=2$ is only measured at the first visit, so we will treat the four samples from $group=2$ and $visit=2$ as outliers. There are two observations have lower bounds equal to 256 when $group=4$ $visit=1$; these may be outliers. We will discuss the six total outliers later. The simple loess lines of the other three groups (calf, pooled calf, pooled cow) are almost parallel, which indicates that an interaction between $group$ and $visit$ may not exist.

Farm effect can also be shown by plotting a line connecting the median lower bounds of the logMIC of every visit of a farm. From Figure 5.10, we can see that all the lines are mixed together, which implies that the farm effect is not large.

After examining the raw data and the plots of the data, it is apparent that the initial model needs to include terms for the effects of $farm$, $visit$ and $group$.




Figure 5.10: Plot of connecting median lower bounds of $\log\text{MIC}_{\text{CEFTIF}}$ of each visit of each farm vs. visit for different farm (each line represents a farm). Values are jittered to avoid overplotting. The farm which is represented by the upper left line seems to have a big effect; this is because the farm has two observations whose lower bounds are 256 (outliers).

The analysis procedure is similar to that for the Avian data set. We will ignore the farm effect at first, and then include it in our model when we do the Bayesian model analysis. Several models were fitted using the MLE method. The two observations have lower bounds equal to 256 made the WinBUGS program not converge, so we will drop the two observations at this stage. Results are summarized in Table 5.8. In terms of the AICc criterion, support for model

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$$

overwhelm other models, and this model has the dominant weight. This will be used to do further analysis.

Model	Log(maximum likelihood)	P	AICc	Δ AIC	Normalized Weight
M1:[Y]=Visit Group(C) Visit \times Group	-694.74	9	1408.06	0.00	0.85
M2:[Y]=Visit(C) Group(C) Visit \times Group	-689.96	15	1411.53	3.47	0.15
M3:[Y]=Visit(C)	-708.09	5	1426.37	18.31	0.00
M4:[Y]=Visit(C) Group(C)	-706.71	8	1429.89	21.83	0.00
M5:[Y]=Visit	-718.53	3	1443.14	35.08	0.00
M6:[Y]=Visit Group(C)	-718.35	6	1448.98	40.91	0.00
M7:[Y]=Group(C)	-729.25	5	1468.69	60.63	0.00

Table 5.8: Results of fitting different models to the Dairy (*Camylobacter*)-Azithromycin data using maximum likelihood estimation from PROC LIFEREG (SAS).

Figure 5.11 is the residual plot of model M1. Residuals are scattered around 0 and no pattern is apparent in the plot. This indicates that there is no evidence of a lack of fit of model M1.




Figure 5.11: Residual plot of model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$. Values are jittered to avoid overplotting.

Two Bayesian models with and without farm as a random effect are also used to examine the data.

- Bayesian model without farm effect

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$$

The prior distributions are: $\beta_0 \sim N(0, 10)$, $\beta_{.group} \sim N(0, 10)$, $\beta_{.visit} \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$

- Bayesian model with farm as a random effect

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$$

The prior for farm is: $\tau_{farm} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{farm} = 1/\sigma_{farm}^2$. Other prior distributions are the same as above.

Table 5.9 presents the estimates for those models. Estimates are similar even when different analysis methods are used because the variance of the farm random effect ($0.26^2 = 0.07$) is relatively small compared to the variance of the residual error ($0.70^2 = 0.49$).

From Table 5.10, the DIC is smaller when farm is included in model than when it doesn't, indicating that

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$$

is the more "useful" model. Summary statistics from this model will be used to draw inferences for the quantities of interest.

The summary statistics from the Bayesian model that includes farm random effects are as follows. The differences in mean logMIC between calf and pooled cow (estimate 0.59 compared to the SD 0.35 is not big enough), pooled calf and pooled cow (estimate 0.06 is small compared to the SD 0.29), calf and pooled calf (estimate 0.53 compared to the SD 0.40 is not big enough) are not statistically significant. $\widehat{\beta}_{group1} - \widehat{\beta}_{group3} = (\widehat{\beta}_{group1} - \widehat{\beta}_{group4}) - (\widehat{\beta}_{group3} - \widehat{\beta}_{group4}) = 0.53$, the standard error is calculated by the computer program and the reported SE is 0.40. From the MLE output, the $\widehat{\beta}_{group1} - \widehat{\beta}_{group3} = 0.45$, the SE = $\left(\left(se_{\beta_{group1} - \beta_{group4}} \right)^2 + \left(se_{\beta_{group1} - \beta_{group4}} \right)^2 + 2cov(\beta_{group1} - \beta_{group4}, \beta_{group3} - \beta_{group4}) \right)^{\frac{1}{2}} =$

	MLE	Bayesian excluding random farm effect	Bayesian including random farm effect
β_0 (intercept) (SE)	-1.66 (0.16)	-1.67 (0.17)	-1.71 (0.17)
$\beta_{group_1} - \beta_{group_4}$ (SE)	0.56 (0.35)	0.55 (0.35)	0.59 (0.35)
$\beta_{group_2} - \beta_{group_4}$ (SE)	-2.74 (0.45)	-2.66 (0.44)	-2.64 (0.47)
$\beta_{group_3} - \beta_{group_4}$ (SE)	0.11 (0.29)	0.11 (0.28)	0.06 (0.29)
β_{visit} (SE)	-0.21 (0.07)	-0.21 (0.07)	-0.20 (0.06)
$\mu_{visit_1group_1}$ (SE)	-1.51 (0.22)	-1.52 (0.22)	-1.51 (0.22)
$\mu_{visit_2group_1}$ (SE)	-1.91 (0.14)	-1.92 (0.14)	-1.90 (0.15)
$\mu_{visit_3group_1}$ (SE)	-2.32 (0.08)	-2.32 (0.08)	-2.29 (0.10)
$\mu_{visit_4group_1}$ (SE)	-2.72 (0.11)	-2.72 (0.12)	-2.68 (0.13)
$\mu_{visit_1group_2}$ (SE)	-2.01 (0.11)	-2.01 (0.11)	-2.05 (0.11)
$\mu_{visit_2group_2}$ (SE)	0.38 (0.37)	0.31 (0.37)	0.25 (0.40)
$\mu_{visit_1group_3}$ (SE)	-1.79 (0.16)	-1.80 (0.16)	-1.86 (0.18)
$\mu_{visit_2group_3}$ (SE)	-2.04 (0.11)	-2.04 (0.10)	-2.06 (0.12)
$\mu_{visit_3group_3}$ (SE)	-2.28 (0.10)	-2.28 (0.10)	-2.27 (0.11)
$\mu_{visit_4group_3}$ (SE)	-2.52 (0.15)	-2.52 (0.15)	-2.48 (0.15)
$\mu_{visit_1group_4}$ (SE)	-1.88 (0.11)	-1.88 (0.11)	-1.91 (0.12)
$\mu_{visit_2group_4}$ (SE)	-2.09 (0.07)	-2.09 (0.07)	-2.10 (0.09)
$\mu_{visit_3group_4}$ (SE)	-2.30 (0.08)	-2.30 (0.08)	-2.30 (0.09)
$\mu_{visit_4group_4}$ (SE)	-2.52 (0.13)	-2.51 (0.13)	-2.49 (0.13)
$\sigma_{residual}$ (SE)	0.73 (0.03)	0.74 (0.03)	0.70 (0.03)
σ_{farm} (SE)			0.26 (0.07)

Table 5.9: Estimates for the Dairy (*Camylobacter*) Project from MLE and Bayesian. Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs. Group=1 is the sample of the enrolled calf, group=2 is the sample from the mother of the enrolled calf, group=3 is the sample of pooled calves, and group=4 is the sample of pooled cows. $\beta_{group_i} - \beta_{group_4}$ is the difference in mean logMIC between i^{th} group and group of pooled cows. β_{visit} is the difference in mean logMIC between the i^{th} and the $(i + 1)^{th}$ visit. $\mu_{visit_i group_j}$ is the mean logMIC of the i^{th} visit and j^{th} group.

Model	P_D	DIC
Model without farm effect	8.80	1407.10
Model with farm effect	24.40	1388.20

Table 5.10: DIC values for model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$ with and without random farm effect.

$\left((0.35)^2 + (0.29)^2 + 2 \times 0.027 \right)^{\frac{1}{2}} = 0.51$. This means there is no strong evidence that the mean logMIC values of calf and pooled cow, pooled calf and pooled cow, calf and pooled calf are different. The difference in mean logMIC between dam and pooled cow (estimate -2.64 is large compared to the SD 0.47) is statistically significant. This means there is evidence that the mean logMIC values of dam and pooled cow are different. The MLE analysis also leads to the same conclusion.

The differences in mean logMIC between successive visits is statistically significant (the estimate of -0.20 is large compared to the SD of 0.06). This means the logMIC is decreasing with visits for the calf, dam, pooled cow and pooled calf groups. The MLE analysis reaches the same conclusion.

Previous plots revealed some potential outliers. We will repeat the procedure, dropping the outliers. Table 5.11 gives the results of re-fitting the models without the four outliers (observations from $group=2$ and $visit=2$). In terms of the AICc criterion, support for model $[Y] = \text{VISIT}(C)$ overwhelm other models, and this model has the dominant weight. Model M1 will be used to do further analysis.

Model	Log(maximum likelihood)	P	AICc	ΔAIC	Normalized Weight
M1:[Y]=Visit(C)	-684.52	5	1379.23	0.00	0.90
M2:[Y]=Visit(C) Group(C)	-684.13	8	1384.74	5.51	0.06
M3:[Y]=Visit	-689.79	3	1385.65	6.42	0.04
M4:[Y]=Visit Group(C)	-688.39	6	1389.05	9.82	0.00
M5:[Y]=Visit Group(C) Visit \times Group	-686.99	8	1390.45	11.22	0.00
M6:[Y]=Visit(C) Group(C) Visit \times Group	-682.27	14	1393.95	14.71	0.00
M7:[Y]=Group(C)	-704.59	5	1419.38	40.15	0.00

Table 5.11: Results of re-fitting different models to the Dairy (*Camylobacter*)-Azithromycin data (without outliers) using maximum likelihood estimation from PROC LIFEREG (SAS).




Figure 5.12: Residual plot of model $[Y] = \text{VISIT}(C)$ after outliers are deleted. Values are jittered to avoid overplotting.

Figure 5.12 is the residual plot of model M1. Residuals are scattered around 0 and no pattern is apparent in the plot. This indicates that there is no evidence of lack of fit of model M1.

Two Bayesian models with and without farm as a random effect are also used to examine the data.

- Bayesian model without farm effect

$$[Y] = \text{VISIT}(C)$$

The prior distributions are: $\beta_0 \sim N(0, 10)$, $\beta.visit \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$

- Bayesian model with farm as a random effect

$$[Y] = \text{VISIT}(C)$$

The prior for farm is: $\tau_{farm} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{farm} = 1/\sigma_{farm}^2$. Other prior distributions are the same as above.

Table 5.12 presents the estimates for those models. Estimates are similar even when different analysis methods are used because the variance of the farm random effect ($0.22^2 = 0.05$) is relatively small compared to the variance of the residual error ($0.71^2 = 0.50$).

From Table 5.13, the DIC is smaller when the farm effect is included in the model than when it doesn't, indicating that

$$[Y] = \text{VISIT}(C) \text{ FARM}(R)$$

is the more "useful" model. Summary statistics from this model will be used to draw inferences for the quantities of interest.

The DIC results indicate that the model $[Y] = \text{VISIT}(C) \text{ FARM}(R)$ has the smallest DIC value among the four Bayesian models. It is still hard to say which model is better than the others because of the influence of the six outliers. We will leave this question to the biologists. They can make a reasonable choice between the two models.

The summary statistics from the Bayesian model $[Y] = \text{VISIT}(C) \text{ FARM}(R)$ is as follows. For the calf, dam, pooled cow and pooled calf groups, the difference in mean logMIC between

	MLE	Bayesian excluding random farm effect	Bayesian including random farm effect
β_0 (intercept) (SE)	-2.65 (0.09)	-2.65 (0.09)	-2.62 (0.11)
$\beta_{visit_1} - \beta_{visit_4}$ (SE)	0.68 (0.12)	0.68 (0.11)	0.60 (0.11)
$\beta_{visit_2} - \beta_{visit_4}$ (SE)	0.85 (0.14)	0.84 (0.14)	0.79 (0.14)
$\beta_{visit_3} - \beta_{visit_4}$ (SE)	0.32 (0.12)	0.31 (0.12)	0.31 (0.11)
$\beta_{visit_1} - \beta_{visit_2}$ (SE)	-0.17 (0.22)	-0.16 (0.13)	-0.19 (0.13)
$\beta_{visit_2} - \beta_{visit_3}$ (SE)	0.53 (0.22)	0.53 (0.13)	0.50 (0.13)
$\mu_{visit_1group_1}$ (SE)	-1.97 (0.07)	-1.97 (0.07)	-2.02 (0.09)
$\mu_{visit_2group_1}$ (SE)	-1.80 (0.11)	-1.81 (0.11)	-1.83 (0.12)
$\mu_{visit_3group_1}$ (SE)	-2.34 (0.07)	-2.34 (0.08)	-2.31 (0.09)
$\mu_{visit_4group_1}$ (SE)	-2.65 (0.09)	-2.65 (0.09)	-2.62 (0.11)
$\mu_{visit_1group_2}$ (SE)	-1.97 (0.07)	-1.97 (0.07)	-2.02 (0.09)
$\mu_{visit_1group_3}$ (SE)	-1.97 (0.07)	-1.97 (0.07)	-2.02 (0.09)
$\mu_{visit_2group_3}$ (SE)	-1.80 (0.11)	-1.81 (0.11)	-1.83 (0.12)
$\mu_{visit_3group_3}$ (SE)	-2.34 (0.07)	-2.34 (0.08)	-2.31 (0.09)
$\mu_{visit_4group_3}$ (SE)	-2.65 (0.09)	-2.65 (0.09)	-2.62 (0.11)
$\mu_{visit_1group_4}$ (SE)	-1.97 (0.07)	-1.97 (0.07)	-2.02 (0.09)
$\mu_{visit_2group_4}$ (SE)	-1.80 (0.11)	-1.81 (0.11)	-1.83 (0.12)
$\mu_{visit_3group_4}$ (SE)	-2.34 (0.07)	-2.34 (0.08)	-2.31 (0.09)
$\mu_{visit_4group_4}$ (SE)	-2.65 (0.09)	-2.65 (0.09)	-2.62 (0.11)
$\sigma_{residual}$ (SE)	0.72 (0.03)	0.74 (0.03)	0.69 (0.03)
σ_{farm} (SE)			0.25 (0.07)

Table 5.12: Estimates for re-fit the Dairy data set without outliers from MLE and Bayesian. Numbers in brackets state the standard errors for the MLEs or the standard deviations for the APEs. Group=1 is sample of calf, group=2 is sample from mother of enrolled calf, group=3 is sample of pooled calf, and group=4 is sample of pooled cow. $\beta_{visit_i} - \beta_{visit_j}$ is the difference in mean logMIC between the i^{th} and j^{th} visit. $\mu_{visit_i group_j}$ is the mean logMIC of the i^{th} visit and j^{th} group.

Model	P_D	DIC
Model without farm effect	4.90	1378.90
Model with farm effect	20.00	1361.40

Table 5.13: DIC values for model $[Y] = \text{VISIT}(C) \text{ FARM}(R)$ with and without farm random effect.

visit occasions 1 and 4, 2 and 4, 3 and 4 are statistically significant. While the difference in mean logMIC between visit occasions 1 and 2 is not statistically significant. This means there is no evidence that the mean logMIC of visit 1 and visit 2 are different. The estimate of -0.19 is not big compared to the SD 0.13. There is evidence that the mean logMIC of visit 1 and visit 4 are different. The estimate of 0.60 is big compared to the SD 0.11. There is evidence that the mean logMIC of visit 2 and visit 4 are different. The estimate of 0.79 is big compared to the SD 0.14. There is evidence that the mean logMIC of visit 3 and visit 4 are different. The estimate of 0.31 is big compared to the SD 0.11. There is evidence that the mean logMIC of visit 2 and visit 3 are different. The estimate 0.50 is big compared to the SE of 0.13. The MLE analysis also leads to the same conclusion. The mean logMIC value is decreasing with visit except when $visit=2$.

5.4 Dairy data set (*E.coli*)

The sample protocol is the same as the Dairy (*Campylobacter*) data set. In this case, colonies of *E.coli* bacteria from each sample were inoculated onto NARMS sensititre plates. Samples were tested against an array of 15 anti-microbial agents (Amikacin, Amoxicillin/Clavulanic Acid, Ampicillin, Cefoxitin, Ceftiofur, Ceftriaxone, Chloramphenicol, Ciprofloxacin, Gentamycin, Kanamycin, Nalidixic Acid, Streptomycin, Sulfamethoxazole, Tetracycline, Trimethoprim/Sulfamethoxazole) and the reported MIC values were given. We will use the CEFTIF (abbreviation of Ceftiofur) value as an example to show how to analyse the data. A portion of the raw data is shown in Figure 5.13. The dilution series is 0.12, 0.25, 0.5, 1, 2, 4, 8 mg/L, representing the different concentrations of anti-microbial agent CEFTIF.

Each of the reported MIC values are converted to a pair of numbers that represent the lower bound and upper bound of the true MIC value.

Based on similar reasoning to that outlined in Section 5.3, the initial model needs to

project	sample	isolate	sex	speices	collected	CEFTIF/8/4
Dairy Calf Project	F2C2975SD2	1 E.CNH	F	BOV	12/02/2006 8:49	= 0.25 NOINTP
Dairy Calf Project	F2PCA2975SD2	1 E.CNH	F	BOV	12/02/2006 8:50	= 0.25 NOINTP
Dairy Calf Project	F2PCA2974SD2	1 E.CNH	F	BOV	12/02/2006 8:44	= 0.25 NOINTP
Dairy Calf Project	F2D800KL	1 E.CNH	F	BOV	12/02/2006 9:09	= 0.25 NOINTP
Dairy Calf Project	F2PCO800KL	1 E.CNH	F	BOV	12/02/2006 9:07	= 0.25 NOINTP
Dairy Calf Project	F2PCA800KL	1 E.CNH	F	BOV	12/02/2006 9:06	= 0.25 NOINTP
Dairy Calf Project	F2C800KL	1 E.CNH	F	BOV	12/02/2006 9:04	= 0.25 NOINTP
Dairy Calf Project	F2LF800KL	1 E.CNH	F	BOV	12/02/2006 9:03	= 8 NOINTP
Dairy Calf Project	F3C31SC	1 E.CNH	F	BOV	12/02/2006 9:01	= 0.25 NOINTP
Dairy Calf Project	F3D31SC	1 E.CNH	F	BOV	12/02/2006 9:10	= 0.25 NOINTP

Figure 5.13: Part of raw data of Dairy (*E.coli*) data set. *Sample* are the combined animal ID and information of *Vet* (names of veterinarians), *visit* (time of visit), *farm* (where those sample come from), *group* (sample from which animal group or pooled animal group). *Isolate* is the bacteria isolated, *sex* is sex of animal, *species* is the bacteria species, *collected* is the date and time sample was analyzed. Column CEFTIF lists the reported MIC values.

include terms for the effects of *farm*, *visit* and *group*.

Figure 5.14 presents a plot of logMIC (lower bound) vs. *visit* for different groups with simple loess lines.

The plot reveals some useful things. The five simple loess lines are not parallel, which indicates that an interaction between *group* and *visit* may exist. Farm effect can also be shown by plotting a line connecting the median lower bounds of the logMIC of every visit of a farm. From Figure 5.15, we can see that most lines are mixed together; this indicates that the farm effect is not large.

After examining the raw data and the plot, it is clear that the initial model needs to include terms for the effects of *farm*, *visit* and *group*.




Figure 5.14: Plot of the lower bounds of logMIC-CEFTIF vs. visit (star for sample of calf, dot for sample from mother of enrolled calf, circle for sample of pooled calf, triangle for sample of pooled cow, heart is liquid calf feed sample), with fitted loess lines for different groups. Values are jittered to avoid overplotting.




Figure 5.15: Plot of connecting median lower bound of $\log\text{MIC_CEFTIF}$ of each visit of each farm vs. visit for different farms. Each line represents a farm. Values are jittered to avoid overplotting.

The analysis procedure is the same as before. We will ignore the farm effect at first, and then include it in our model when we do the Bayesian model analysis.

Several models were fitted using the MLE method, results are summarized in Table 5.14. In terms of the AICc criterion, support for model

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$$

overwhelms other models, and this model has the dominant weight. It will be used to do further analysis.

Model	Log(maximum likelihood)	P	AICc	Δ AIC	Normalized Weight
M1:[Y]=Visit Group(C) Visit \times Group	-724.32	11	1471.15	0.00	0.97
M2:[Y]=Visit	-736.56	3	1479.16	8.01	0.02
M3:[Y]=Visit(C) Group(C) Visit \times Group	-722.08	18	1481.52	10.37	0.01
M4:[Y]=Visit(C)	-735.99	5	1482.10	10.95	0.00
M5:[Y]=Group(C)	-735.21	6	1482.59	11.43	0.00
M6:[Y]=Visit Group(C)	-734.94	7	1484.09	12.94	0.00
M7:[Y]=Visit(C) Group(C)	-734.33	9	1487.01	15.86	0.00

Table 5.14: Results of fitting different models to the Dairy (*E.coli*)-CEFTIF data using maximum likelihood estimation from PROC LIFEREG (SAS).

Figure 5.16 is the residual plot of model M1. Residuals are scattered around 0 and no pattern is apparent in the plot. This indicates that there is no evidence of a lack of fit of the model M1.




Figure 5.16: Residual plot of model $[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$. The symbol “X” on the graph indicates truncated residuals as explained in Figure 5.3. Values are jittered to avoid overplotting.

Two Bayesian models with and without farm as a random effect are also used to examine the data.

- Bayesian model without farm effect

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP}$$

The prior distributions are: $\beta_0 \sim N(0, 10)$, $\beta_{\text{group}} \sim N(0, 10)$, $\beta_{\text{visit}} \sim N(0, 10)$, $\tau_Y \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_Y = 1/\sigma^2$

- Bayesian model with farm as a random effect

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$$

The prior for farm is: $\tau_{\text{farm}} \sim \text{GAMMA}(0.1, 0.01)$ and $\tau_{\text{farm}} = 1/\sigma_{\text{farm}}^2$. Other prior distributions are the same as above.

Table 5.15 presents the estimates for these models. Most estimates are similar even when different analysis methods were used because the variance of the farm random effect ($0.17^2 = 0.03$) is relatively small compared to the variance of the residual error ($0.70^2 = 0.49$). A few estimates are quite different; this may be because of the small sample size. For example, when $\text{group}=2$ and $\text{visit}=2$, the sample size is 2; when $\text{group}=5$ and $\text{visit}=1$, the sample size is 1; and when $\text{group}=5$ and $\text{visit}=3$, the sample size is 1.

From Table 5.16, the DIC is smaller when farm is included in the model than when it doesn't, indicating that

$$[Y] = \text{VISIT GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$$

is the more "useful" model. Summary statistics from this model will be used to draw inferences for the quantities of interest.

The summary statistics from the Bayesian model that includes farm random effects are as follows. The differences in mean logMIC between calf and pooled cow (estimate of -0.04 is small compared to the SD of 0.20), between dam and pooled cow (estimate of 0.26 is small compared to the SD of 0.58), between pooled calf and pooled cow (estimate of -0.09 is small compared to the SD of 0.21), and between calf and pooled calf (estimate 0.05 is small compared to the se 0.19) are not statistically significant. $\widehat{\beta}_{\text{group}_1} - \widehat{\beta}_{\text{group}_3} = (\widehat{\beta}_{\text{group}_1} - \widehat{\beta}_{\text{group}_4}) - (\widehat{\beta}_{\text{group}_3} - \widehat{\beta}_{\text{group}_4}) = 0.05$, the standard error is calculated by the computer program and the

reported SE is 0.19. This means there is no strong evidence that the mean logMIC values of calf and pooled cow, dam and pooled cow, calf and pooled calf, or pooled calf and pooled cow are different. The difference in mean logMIC between liquid feed calf and pooled cow (estimate of 4.11 is big compared to the SD of 1.11) is statistically significant. This means that there is evidence that the mean logMIC values of liquid feed calf and pooled cow are different. The MLE analysis also indicates the same conclusion.

The difference in mean logMIC between successive visit (estimate of -0.02 is small compared to the SD of 0.06) is not statistically significant. This means there is no strong evidence that the mean logMIC changes with visit for the calf, dam, pooled cow, pooled calf and liquid feed calf. The MLE also demonstrates the same conclusion.

	MLE	Bayesian excluding random farm effect	Bayesian including random farm effect
β_0 (intercept) (SE)	-1.54 (0.16)	-1.52 (0.15)	-1.51 (0.16)
$\beta_{group_1} - \beta_{group_4}$ (SE)	0.00 (0.20)	-0.03 (0.20)	-0.04 (0.20)
$\beta_{group_2} - \beta_{group_4}$ (SE)	0.03 (0.57)	0.00 (0.56)	0.26 (0.58)
$\beta_{group_3} - \beta_{group_4}$ (SE)	-0.05 (0.22)	-0.07 (0.23)	-0.09 (0.21)
$\beta_{group_5} - \beta_{group_4}$ (SE)	6.46 (1.49)	5.08 (1.27)	4.11 (1.11)
β_{visit} (SE)	-0.01 (0.06)	-0.02 (0.06)	-0.02 (0.06)
$\mu_{visit_1group_1}$ (SE)	-1.57 (0.09)	-1.58 (0.08)	-1.58 (0.09)
$\mu_{visit_2group_1}$ (SE)	-1.60 (0.06)	-1.61 (0.05)	-1.61 (0.06)
$\mu_{visit_3group_1}$ (SE)	-1.64 (0.06)	-1.64 (0.06)	-1.64 (0.06)
$\mu_{visit_4group_1}$ (SE)	-1.67 (0.09)	-1.67 (0.08)	-1.67 (0.09)
$\mu_{visit_1group_2}$ (SE)	-1.62 (0.10)	-1.63 (0.10)	-1.63 (0.11)
$\mu_{visit_2group_2}$ (SE)	-1.73 (0.51)	-1.73 (0.52)	-2.01 (0.51)
$\mu_{visit_1group_3}$ (SE)	-1.59 (0.11)	-1.59 (0.11)	-1.60 (0.11)
$\mu_{visit_2group_3}$ (SE)	-1.59 (0.07)	-1.59 (0.07)	-1.59 (0.08)
$\mu_{visit_3group_3}$ (SE)	-1.58 (0.08)	-1.59 (0.08)	-1.58 (0.08)
$\mu_{visit_4group_3}$ (SE)	-1.58 (0.12)	-1.59 (0.13)	-1.57 (0.12)
$\mu_{visit_1group_4}$ (SE)	-1.55 (0.10)	-1.54 (0.10)	-1.53 (0.11)
$\mu_{visit_2group_4}$ (SE)	-1.56 (0.07)	-1.56 (0.07)	-1.55 (0.08)
$\mu_{visit_3group_4}$ (SE)	-1.57 (0.07)	-1.57 (0.07)	-1.57 (0.08)
$\mu_{visit_4group_4}$ (SE)	-1.58 (0.11)	-1.59 (0.11)	-1.59 (0.11)
$\mu_{visit_1group_5}$ (SE)	1.73 (0.72)	1.08 (0.63)	0.61 (0.58)
$\mu_{visit_2group_5}$ (SE)	-1.46 (0.32)	-1.40 (0.32)	-1.38 (0.31)
$\mu_{visit_3group_5}$ (SE)	-4.64 (0.97)	-3.87 (0.84)	-3.36 (0.72)
$\sigma_{residual}$ (SE)	0.69 (0.02)	0.70 (0.02)	0.69 (0.02)
σ_{farm} (SE)			0.17 (0.05)

Table 5.15: Estimates for the Dairy (*E.coli*)-CEFTIF data from MLE and Bayesian analyses. Numbers in bracket state the standard errors for the MLEs or standard deviations for the APEs. Group=1 is sample of calf, group=2 is sample from dam, group=3 is sample of pooled calf, group=4 is sample of pooled cow, and group=5 is liquid calf feed sample. $\beta_{group_i} - \beta_{group_4}$ is the difference in mean logMIC between the i^{th} group with group of pooled cows. β_{visit} is the difference in mean logMIC between the i^{th} and $(i + 1)^{th}$ visit. $\mu_{visit_i group_j}$ is the mean logMIC of the i^{th} visit and j^{th} group.

Model	P_D	DIC
Model without farm effect	10.41	1470.33
Model with farm effect	23.85	1461.11

Table 5.16: DIC values for model $[Y] = VISIT \text{ GROUP}(C) \text{ VISIT} \times \text{GROUP FARM}(R)$ with and without farm random effect.

Chapter 6

Conclusions and Suggestions

We assumed a log-normal distribution of MIC data, which are interval censored, and developed different models for which the unknown mean value of logMIC is related to some explanatory variables. Four real data sets were analyzed. The MLE and Bayesian analyses were performed and estimates from the two analysis procedures were compared with each other. When the variance of random effects is small compared to the variance of random noise, the two methods give similar results. The result of the simulation study also confirmed that when the variance of the random effect is small relative to the variance of random noise, estimates from two analysis procedures are similar. This implies that MLE analysis (using SAS PROC LIFEREG) may be sufficient for many datasets.

If the MLE method is used, ignoring farm random effects, what is the approximate increase in the SE that should be seen to account for farm effects? In our examples, the increase were $1.40\times$, $1.43\times$, $1.11\times$ ($1.23\times$ for model M1 when fitting reduced data), and $1.01\times$ respectively. The number is obtained by calculating the ratios $\frac{se(\text{from Bayesian analysis})}{se(\text{MLE})}$ for estimates of mean logMIC, obtaining the average of the ratios; the average number is the approximate increase in the SE. Those numbers are small, which indicates again a small effect for ignoring farm random effect.

The suggestion for increasing the precision of the estimates is as follows. (1) Increase sample size. Animals in the same farm are usually correlated with each other; 1000 correlated samples may give similar information as 500 independent samples. (2) Increasing the number of concentrations of anti-microbial agents in the microtitre plate may reduce the size of the interval that contains the true MIC. If, from previous knowledge, we can estimate the MIC value of an anti-microbial agent as 0.50 mg/L, we can have concentrations of 0.80, 0.90,

1, 1.30, 1.50 and 2 mg/L instead of concentrations of 0.40, 0.50, 1 and 2 mg/L on the microtitre plate. (3) Change the number of concentrations of some anti-microbial agents in the microtitre plate. For example, in the Avian data set, the MIC values results for anti-microbial CLINDA are all ≥ 4 , more appropriate concentrations may be 4, 5, and 6 mg/L instead of concentrations 1, 2 and 4 mg/L on the microtitre plate.

Using a simple and consistent coding system for recording may make future work easier. For example, in the Dairy (*E.coli*) data set, the variable sample contain information on farm, vet, visit, and animal (group).

6.1 Future work

For future work, we suggest to plan studies early. Existing estimates have relatively poor precision. Hence, we need to do a power analysis to see what size of sample is required to get accurate parameter estimates. Also it is hard to know how much precision is lost by censoring, a simulation study may help us to investigate it.

Bibliography

- Brooks, S. and Gelman, A. (1998) Assessing convergence of Markov Chain Monte Carlo algorithms. *Statistics and Computing*, 8, 319-335.
- Burnham, K.P. and Anderson, D.R. (2002) *Model selection and multi-model inference : a practical information-theoretic approach*, Springer, New York.
- Dobson, A.J. (2002). *An Introduction to Generalized Linear Models*, Second Edition. Chapman and Hall, London.
- Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B. (2003) *Bayesian Data Analysis*, Chapman and Hall/CRC.
- Givens, G.H. and Hoeting, J.A. (2005) *Computational Statistics*, John Wiley & Sons, Inc., Hoboken, New Jersey.
- Kutner, M.H., Nachtstein, C.J., Neter, J. and Li, W. (2004) *Applied Linear Statistical Models*, 5th Ed. McGraw-Hill/Irwin, Boston.
- Lee, M.T. and Whitmore, G.A. (1999). Statistical Inference for Serial Dilution Assay Data. *Biometrics*, 55, 1215-1220.
- Long, J.S. (1997). *Regression Models for Categorical and Limited Dependent Variables*, SAGE Publications, Thousand Oaks.
- Lunn, D.J., Thomas, A., Best, N., and Spiegelhalter, D. (2000). WinBUGS - a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, 10:325-337.
- R Development Core Team (2008). R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org>

SAS Institute Inc. (2008). SAS/STAT 9.2 Users Guide. Cary, NC: SAS Institute Inc.

Spiegelhalter, D.J., Best, N.G., Carlin, B.P. and Van der Linde, A. (2002). Bayesian Measures of Model Complexity and Fit, *Journal of the Royal Statistical Society, Series B*, 64, 583-639.