

ST 762, HOMEWORK 3, FALL 2009

These problems are to be turned in on the due date.

1. *Using standard nonlinear regression software to implement GLS and normal theory ML with estimation of θ .*

For drugs intended to be administered orally, such as solid tablets, it is important that the drug is appropriately dissolved within the gastrointestinal (GI) tract so that it may be absorbed into the system. Moreover, it is critical that the way in which this dissolution occurs is well-understood and consistent over batches of product manufactured and over a range of conditions, so that product quality and reliability can be ensured. Thus, it is standard in the pharmaceutical industry, and required by regulatory authorities (the Food and Drug Administration), for companies to carry out so-called dissolution studies, which are meant to assist them in evaluating and monitoring the dissolution properties of a drug. These dissolution studies use a sophisticated apparatus that is meant to simulate the conditions of the GI tract. Specifically, a tablet is placed in a vessel that contains a dissolution medium, which is a liquid whose properties are meant to mimic those existing in the GI tract. The tablet will dissolve in the medium over time. The medium is stirred throughout the study by an automated paddle, and a number of conditions, such as volume of liquid, temperature, stirring speed, and so on, under which this is to take place may be specified. Time zero (0) is the time at which the tablet is placed in the medium, and the study continues for some prespecified length of time (usually 1–2 hours). At various intermittent points in time following time 0, a sample is taken from the vessel, and the percent of the tablet dissolved (the so-called “dissolution value”) is determined.

The file `dissolution.dat`, available on the class web page, contains data from a dissolution study from which the scientists wished to understand the effects of *vibration* on the dissolution properties of an oral medication. The no vibration condition was meant to simulate how a tablet might be dissolved in the GI tract of a patient who remains in a resting position for the 2 hours following ingestion of the tablet, while the vibration condition was meant to simulate roughly how a tablet would be dissolved in a person who begins engaging in physical activity shortly after ingesting a tablet.

Twelve (12) tablets were taken at random from a recently manufactured lot of tablets. Each was placed at time 0 minutes in a vessel containing 500 ml of medium held at 37 degrees Celsius, with the paddle set to a stirring speed of 50 revolutions per minute. Dissolution values were ascertained over the next 2 hours (120 minutes). The 12 tablets were randomly assigned to either dissolution with no vibration during the entire 120 minutes (6 tablets) or dissolution with vibration introduced after 20 minutes (6 tablets). For each of the 6 tablets assigned to the no vibration condition, the vessel was kept at the above temperature, volume, and paddle speed over the next 2 hours (120 minutes), and no vibration was introduced. Dissolution values were ascertained on each at 2.5, 5, 7.5, 15, 20, 30, 45, 60, 75, 90, and 120 minutes. For each of the 6 tablets assigned to the vibration condition, the vessel was kept at this temperature, volume, and paddle speed over the next 2 hours; however, at 20 minutes, the vessel was in addition subjected to vibration. That is, starting at 20 minutes after introduction of the tablet into the vessel, the table on which apparatus was mounted was shaken for the rest of the 2 hours. As for the no vibration tablets, dissolution values were ascertained at 2.5, 5, 7.5, 15, 20, 30, 45, 60, 75, 90, and 120 minutes.

The data were then as follows. At each time point, the dissolution values from the 6 tablets under the no vibration condition were averaged, and similarly for the vibration condition.

Thus, for each condition, there are 11 observations, where each consists of the time point and the average of the 6 dissolution values. The columns of the data set are (1) vibration condition (0=no vibration, 1=vibration); (2) time, x , in minutes; and Y , average of 6 dissolution values (percent), the response.

A standard model for the dissolution-time relationship under no vibration is the so-called *Weibull model*

$$W_0(x) = W_{\max}[1 - \exp\{-\log(2)(x/\tau)^\beta\}], \quad x \geq 0, \quad (1)$$

where $W_0(x)$ is dissolution value at time x , depending on parameters W_{\max} , the “amplitude,” representing the asymptote of dissolution as time gets large; τ , the “scale,” representing the time at which dissolution value is equal to half of W_{\max} ; and β , the “shape” parameter, which governs the shape of the curve ($\beta=1$ leads to an exponential curve and $\beta > 1$ to an S-shaped curve). When vibration is introduced, at $t_0 > 0$ minutes (so $t_0 = 20$ in the current experiment), the effect is to “shift” the Weibull relationship. They have thus proposed the following model to capture this phenomenon. Let $W_1(x)$ be the dissolution value at any time $x \geq 0$ when vibration is introduced at $t_0 > 0$. Then the model is

$$\begin{aligned} W_1(x) &= W_0(x), & 0 \leq x \leq t_0 \\ &= W_0(t_0) + W_{\max}^* \left(1 - \exp[-\log(2)\{(x - t_0)/\tau^*\}^{\beta^*}]\right), & x > t_0. \end{aligned} \quad (2)$$

In (2), then, the dissolution profile follows relationship (1), that expected under no vibration, up until time t_0 . (Note that the data up to 20 minutes from *both* experimental conditions have information in them about the values of W_{\max}, τ, β .) After t_0 , it is as if time is “reset” so that t_0 is the time origin for the dissolution under vibration. The value of dissolution at t_0 , $W_0(t_0)$, is where dissolution “starts out” once the vibration is introduced. W_{\max}^* is then the asymptote of dissolution under vibration measured from $W_0(t_0)$; τ^* is the “scale” after t_0 , i.e., the time it takes after t_0 for dissolution to rise halfway between $W_0(t_0)$ and W_{\max}^* ; and β^* is the “shape” parameter describing the shape after vibration is introduced, which could very well be different from the shape under no vibration.

In carrying out the analyses below, make the following assumptions: (i) the magnitude and pattern of variation in average dissolution values is the same under no vibration and vibration; and (ii) serial correlation in average dissolution values is negligible. Thus, you may assume that dissolution values Y_j at times x_j , $j = 1, \dots, n$ under both vibration conditions are independent and follow a model $E(Y_j|x_j) = f(x_j, \boldsymbol{\beta})$ and $\text{var}(Y_j|x_j) = \sigma^2 f^{2\theta}(x_j, \boldsymbol{\beta})$, where σ and θ are the same for both vibration conditions. The model $f(x_j, \boldsymbol{\beta})$ should pertain to both conditions (vibration and no vibration); it is up to you to determine what $f(x_j, \boldsymbol{\beta})$ should be based on (1) and (2) above.

(a) Make a plot of the data using your favorite software, using a different symbol for each vibration condition. This will give you a sense of the shape of the response-time relationship for both conditions.

(b) With the examples in Section 6.8 as a guide, write a program to implement the 3-step GLS algorithm using *either* SAS `proc nlin` or R function `nls()` for fitting your model to data from both vibration conditions simultaneously under the above assumptions, estimating θ at step (ii) of the algorithm by pseudolikelihood (PL) implemented by “The Trick.” You will need to

- Iterate the 3-step GLS algorithm at least $C = 10$ times

- Determine suitable starting values by some *ad hoc* procedure
- Estimate σ^2 using the final estimates of β and θ from the 3-step procedure using the usual bias-adjusted estimator

(c) Re-do (b), except replace PL estimation of θ in step (ii) by estimation using the *identity transformation*. To do this, you will need to implement “The Trick” using the result of part (c) of Extra Problem 5 with $\lambda = 1$. You may use the same (quadratic) estimator for σ^2 as you did in (b), obtaining the estimate by substituting the final estimates of β and θ from this procedure.

(d) Re-do (b) yet again, this time replacing PL estimation of θ in step (ii) by using the log transformation estimator described in Example 6.3 on page 136 of the notes. You may use the same estimator for σ^2 as you did in (a), using the final estimates of β and θ from this procedure. Compare the square root of this estimate to the antilogarithm of the final “intercept” estimate you get for estimating θ (i.e. $\exp(\eta)$) – how do they differ and what do you think may be the reason?

(e) An alternative to using GLS methods, which rely on a linear estimating equation, is to use quadratic estimating equations to define estimators for β , θ , and σ . As discussed in the notes, one such set of equations arises from differentiation of the loglikelihood for β , θ , and σ found by assuming that $Y_j|\mathbf{x}_j$ is normally distributed; these equations are given in (6.9) on page 125 of the notes and correspond to what we have called *normal theory maximum likelihood*. Solving these equations can be implemented by using “The Trick” with both β and θ treated as “regression parameters,” as discussed on page 132 of the notes. Using this method, obtain still another set of estimates for β , θ , and σ^2 . Compare the results to those you obtained in part (b) using PL – are the point estimates similar?

(f) Inspecting the point estimates of θ from all four fits, do you think there is convincing evidence that the data exhibit variance that increases with the mean response?

(g) Make a plot of the data again, superimposing the fit from part (b) on the data. Does the model seem to give a good description of the pattern of the relationship?

2. *Diagnostics for nonconstant variance.* You will now get to try your hand at constructing and interpreting the various plots discussed in the notes. Consider again the dissolution data in Problem 1.

(a) Find the OLS estimate of β . Also find the estimate of the assumed common variance under your OLS fit, $\hat{\sigma}_{OLS}^2$. (You should have done this in Problem 1.)

(b) Write a program (as always, in your favorite programming language) to compute, given the OLS estimate of β , (approximate) studentized OLS residuals

$$b_j = \frac{r_j}{\hat{\sigma}_{OLS}(1 - h_{jj})^{1/2}},$$

where r_j are the usual OLS residuals and h_{jj} are the diagonal elements of the approximate “hat” matrix. Note that `proc nlin` in SAS has an option to compute these; *do not* use this option, but rather, write a program to compute the hat matrix yourself. Make the following plots:

- usual OLS residuals vs. log predicted values (use log predicted values to “stretch out” the horizontal axis)

- (i) studentized OLS residuals vs. log predicted values
- (iii) log absolute studentized residuals vs. log predicted values
- (iv) 2/3-root (absolute) studentized residuals vs. log predicted values

Comment on the evidence in the plots regarding nonconstant variance in general and the relevance of the power variance model. Do the plots seem to support the suitability of this variance model? Is there another variance model you might consider?

(c) Now write a program to compute weighted (approximate) studentized residuals as described in Section 7.3 of the notes (again, write it yourself). Using your results from Problem 1(b) for the 3-step GLS fit using PL estimation of θ in (a) make the same plots you did in (b) using the weighted studentized residuals. Did it “work”?

(d) Using the “folklore” result, compute standard errors for the GLS-PL estimator for β . Compare the standard errors you compute to the final standard errors output by either `proc nlin` or `nls()` in your implementation of the three-step algorithm. Hopefully, they are the same!

(e) Compute “robust sandwich” standard errors for the GLS estimator as described in Section 9.4. How do these alternative standard errors compare to the usual values based on the folklore result?