

Final Examination (Fall 1996)
Economics 573
Suggested Answer Sheets

In what follows are the final exam questions (in bold) and suggested answers:

1. (20 points) Consider the linear regression model, with:

$$y = X\mathbf{b} + \mathbf{e} \tag{1}$$

where $E[\mathbf{e}|X] = 0$ and $E[\mathbf{e}\mathbf{e}'|X] = \mathbf{S}^2\mathbf{\Omega}$. Let:

$$\hat{\mathbf{b}} = (X'\mathbf{\Omega}^{-1}X)^{-1}X'\mathbf{\Omega}^{-1}y. \tag{2}$$

denote the GLS estimator of \mathbf{b} . Prove that any other linear unbiased estimator of \mathbf{b} , say \mathbf{b}^* , will have a variance which exceeds that of $\hat{\mathbf{b}}$ by a positive definite matrix.

You could approach this proof in two ways. First, since \mathbf{b}^* is linear, we can write it in the form:

$$\begin{aligned} \mathbf{b}^* &= Ay \\ &= [(X'\mathbf{\Omega}^{-1}X)^{-1}X'\mathbf{\Omega}^{-1} + C]y \end{aligned} \tag{A.1}$$

where

$$C \equiv [(X'\mathbf{\Omega}^{-1}X)^{-1}X'\mathbf{\Omega}^{-1} - A] \tag{A.2}$$

Notice that the GLS estimator corresponds to $C = \mathbf{0}_{K \times N}$, so that:

$$\begin{aligned} \mathbf{b}^* &= \hat{\mathbf{b}} + Cy \\ &= \hat{\mathbf{b}} + CX\mathbf{b} + C\mathbf{e} \end{aligned} \tag{A.3}$$

Since $\hat{\mathbf{b}}$ is unbiased, \mathbf{b}^* will be unbiased only if:

$$CX = 0. \tag{A.4}$$

Thus

$$\mathbf{b}^* - \mathbf{b} = [(X'\mathbf{\Omega}^{-1}X)^{-1}X'\mathbf{\Omega}^{-1} + C]\mathbf{e} \tag{A.5}$$

Finally, then, the covariance matrix of \mathbf{b}^* is then given by:

$$\begin{aligned}
\text{Var}[\mathbf{b}^*] &= E\left\{ \left[\mathbf{b}^* - \mathbf{b} \right] \left[\mathbf{b}^* - \mathbf{b} \right]' \right\} \\
&= E\left\{ \left[\mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{C} \right] \mathbf{e}\mathbf{e}' \left[\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{C}' \right] \right\} \\
&= \left[\mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{C} \right] \mathbf{s}^2 \Omega \left[\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{C}' \right] \\
&= \mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1}\mathbf{s}^2\Omega\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{C}\mathbf{s}^2\Omega\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1} \\
&\quad + \mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1}\mathbf{s}^2\Omega\mathbf{C}' + \mathbf{C}\mathbf{s}^2\Omega\mathbf{C}' \\
&= \mathbf{s}^2\mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1} + \mathbf{s}^2\mathbf{C}X\mathbb{H}^{-1}X'\Omega^{-1} + \mathbf{s}^2\mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\mathbf{C}' + \mathbf{s}^2\mathbf{C}\Omega\mathbf{C}' \\
&= \mathbf{s}^2\mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1} + \mathbf{s}^2\mathbf{C}\Omega\mathbf{C}' \\
&= \text{Var}[\hat{\mathbf{b}}] + \mathbf{s}^2\mathbf{C}\Omega\mathbf{C}'
\end{aligned} \tag{A.5}$$

where the second to the last equality follows from (A.3).

An alternative approach to this problem would have been to note that we can define P such that:

$$P'P = \Omega^{-1}. \tag{A.6}$$

Then premultiplying both sides of equation (1) by P yields:

$$\begin{aligned}
Py &= PX\mathbf{b} + P\mathbf{e} \\
\text{or} & \tag{A.7}
\end{aligned}$$

$$\tilde{y} = \tilde{X}\mathbf{b} + \tilde{\mathbf{e}}$$

where

$$\begin{aligned}
E[\tilde{\mathbf{e}}] &= PE[\mathbf{e}] = 0 \\
\text{Var}[\tilde{\mathbf{e}}\tilde{\mathbf{e}}'] &= \mathbf{s}^2 I
\end{aligned} \tag{A.8}$$

so that transformed model in satisfies the classical linear regression model assumptions used in the Gauss-Markov Theorem. In this case, the OLS estimator of \mathbf{b} using the transformed model is *best* (i.e., minimum variance) among the class of unbiased estimators. But,

$$\begin{aligned}
\tilde{\mathbf{b}} &= \left(\tilde{X}'\tilde{X} \right)^{-1} \tilde{X}'\tilde{y} \\
&= \left(X'P'PX \right)^{-1} X'P'Py \\
&= \mathbf{C}X'\Omega^{-1}X\mathbb{H}^{-1}X'\Omega^{-1}y \\
&= \hat{\mathbf{b}}
\end{aligned} \tag{A.9}$$

Thus, the GLS estimator of \mathbf{b} is also *best*.

2. (25 points) Consider the linear regression model, with:

$$y_i = \mathbf{b}x_i + \mathbf{e}_i \quad i = 1, \dots, N \quad (3)$$

where x_i is a scalar and $E[\mathbf{e}_i|x_i] = 0$, $E[\mathbf{e}_i^2|x_i] = \mathbf{s}_i^2$, and $E[\mathbf{e}_i\mathbf{e}_j|X] = 0$. Under what conditions (i.e., specifications for \mathbf{s}_i^2) will the GLS estimator of \mathbf{b} be given by:

$$\text{i. } \hat{\mathbf{b}} = \frac{\sum_{i=1}^N x_i^2 y_i}{\sum_{i=1}^N x_i^3}$$

$$\text{ii. } \hat{\mathbf{b}} = \frac{1}{N} \sum_{i=1}^N \frac{y_i}{x_i}$$

$$\text{iii. } \hat{\mathbf{b}} = \frac{\sum_{i=1}^N y_i}{\sum_{i=1}^N x_i}$$

$$\text{iv. } \hat{\mathbf{b}} = \frac{\sum_{i=1}^N x_i y_i}{\sum_{i=1}^N x_i^2}$$

The first thing to note about this problem is that the model is one with heteroskedasticity. As we discussed in class, the GLS estimator in this case is given by:

$$\begin{aligned} \hat{\mathbf{b}} &= (X' \Omega^{-1} X)^{-1} X' \Omega^{-1} y \\ &= \left(\sum_{i=1}^N \frac{x_i x_i'}{\mathbf{s}_i^2} \right)^{-1} \sum_{i=1}^N \frac{x_i y_i}{\mathbf{s}_i^2} \end{aligned} \quad (\text{A.10})$$

For the case in which x_i is a scalar, this reduces to:

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^N x_i y_i}{\sum_{i=1}^N \frac{x_i^2}{\mathbf{s}_i^2}} \quad (\text{A.11})$$

We can now answer the questions posed above.

- i. For this question, if $\mathbf{s}_i^2 = \frac{\mathbf{S}^2}{x_i}$, then (A.11) yields

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^N \frac{x_i y_i}{\mathbf{s}_i^2 x_i^{-1}}}{\sum_{i=1}^N \frac{x_i^2}{\mathbf{s}_i^2 x_i^{-1}}} = \frac{\sum_{i=1}^N x_i^2 y_i}{\sum_{i=1}^N x_i^3} \quad (\text{A.12})$$

- ii. For this question, if $\mathbf{s}_i^2 = \mathbf{S}^2 x_i^2$, then (A.11) yields

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^N \frac{x_i y_i}{\mathbf{S}^2 x_i^2}}{\sum_{i=1}^N \frac{x_i^2}{\mathbf{S}^2 x_i^2}} = \frac{\sum_{i=1}^N \frac{y_i}{x_i}}{\sum_{i=1}^N 1} = \frac{1}{N} \sum_{i=1}^N \frac{y_i}{x_i} \quad (\text{A.13})$$

- iii. For this question, if $\mathbf{s}_i^2 = \mathbf{S}^2 x_i$, then (A.11) yields

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^N \frac{x_i y_i}{\mathbf{S}^2 x_i}}{\sum_{i=1}^N \frac{x_i^2}{\mathbf{S}^2 x_i}} = \frac{\sum_{i=1}^N y_i}{\sum_{i=1}^N x_i} \quad (\text{A.14})$$

- iv. For this question, if $\mathbf{s}_i^2 = \mathbf{S}^2$, then (A.11) yields

$$\hat{\mathbf{b}} = \frac{\sum_{i=1}^N \frac{x_i y_i}{\mathbf{S}^2}}{\sum_{i=1}^N \frac{x_i^2}{\mathbf{S}^2}} = \frac{\sum_{i=1}^N x_i y_i}{\sum_{i=1}^N x_i^2} \quad (\text{A.15})$$

Notice that in each of these cases, \mathbf{S}^2 cancels out of the numerator and denominator of each estimator.

3. (20 points) You are given the following two models of an individual's income. In the first model,

$$y_i = \mathbf{a} + \mathbf{b}_E E_i + \mathbf{e}_i \quad (4)$$

where y_i denotes the individual's annual income and

$$E_i = \begin{cases} 3 & \text{individual } i\text{'s highest degree is a Ph.D. or M.D.} \\ 2 & \text{individual } i\text{'s highest degree is a Masters degree} \\ 1 & \text{individual } i\text{'s highest degree is a Bachelors.} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

indicates the highest degree achieved by the individual. In the second model, we have:

$$y_i = \mathbf{a} + \mathbf{b}_{BS} D_{BS,i} + \mathbf{b}_{MS} D_{MS,i} + \mathbf{b}_{PhD} D_{PhD,i} + \mathbf{e}_i \quad (6)$$

where

$$D_{BS,i} = \begin{cases} 1 & \text{individual } i\text{'s highest degree is a Bachelors} \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

$$D_{MS,i} = \begin{cases} 1 & \text{individual } i\text{'s highest degree is a Masters} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

$$D_{PhD,i} = \begin{cases} 1 & \text{individual } i\text{'s highest degree is a Ph.D. or M.D.} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

How would you test model (4) as a restriction on model (6)? In particular, specify the hypothesis that you would be making and the test statistic you would employ.

The hypothesis of interest corresponds to:

$$H_0: \mathbf{b}_{PhD} = 3\mathbf{b}_{BS} \text{ and } \mathbf{b}_{MS} = 2\mathbf{b}_{BS} \quad (A.16)$$

or, equivalently:

$$H_0: \frac{1}{3}\mathbf{b}_{PhD} = \frac{1}{2}\mathbf{b}_{MS} = \mathbf{b}_{BS} \quad (A.17)$$

Substituting this restriction into model (6) yields:

$$\begin{aligned} y_i &= \mathbf{a} + \mathbf{b}_{BS} D_{BS,i} + 2\mathbf{b}_{BS} D_{MS,i} + 3\mathbf{b}_{BS} D_{PhD,i} + \mathbf{e}_i \\ &= \mathbf{a} + \mathbf{b}_{BS} (D_{BS,i} + 2D_{MS,i} + 3D_{PhD,i}) + \mathbf{e}_i \\ &= \mathbf{a} + \mathbf{b}_{BS} E_i + \mathbf{e}_i \end{aligned} \quad (A.18)$$

where \mathbf{b}_{BS} would correspond to \mathbf{b}_E in equation (4). Note that there are only two restrictions being imposed, as (A.16) makes clear. We could test this restriction using

- the standard F-test, distributed F with [2,N-4] degrees of freedom, or

- a Wald test, distributed χ^2_2 .

4. (35 points) You are interested in modeling the electricity usage of households in Iowa. The local utility can provide you with two data sets to help you in your analysis. The first data set consists of 500 households randomly selected from their *low usage* customers (i.e., with annual usage less than 6,000 kWh). The second data set consists of 500 households randomly selected from their *high usage* customers (i.e., with annual usage greater than 20,000 kWh). For each household, you are provided with data on their:

- annual usage (K_i),
- annual household income (I_i), and
- the number of household members (M_i).

You hypothesize that, in the population as a whole, annual household usage follows the linear regression model:

$$K_i = a + bI_i + gM_i + e_i \quad (10)$$

a. Detail how you would estimate your model, including a complete specification of the associated likelihood function you would employ and any additional assumptions you would need.

There are two problems with the available data sets. First, we have no observations from the middle of the distribution of usage. Instead, the two sets combined represent a truncated data set, with the middle of the sample (i.e., $6,000 \leq K_i \leq 20,000$) missing. There is a second problem with this data set, namely that we are not given a random sample from the combined truncated regions (i.e., $K_i \leq 6,000$ and $20,000 \leq K_i$). Instead we are given 500 observations from each of the upper and lower tails of the usage distribution, respectively. For now, let's ignore this second problem and assume that the low and high usage groups are roughly the same size in the population. What we have then is a truncated sample, with observations only in the tails of the distribution. In general, we can write the truncated distribution as:

$$\begin{aligned} f_K(K_i | K_i \leq 6,000 \text{ or } K_i \geq 20,000) &= f_e(e_i | e_i \leq [6,000 - a - bI_i - gM_i] \text{ or } e_i \geq [20,000 - a - bI_i - gM_i]) \\ &= \frac{f_e(e_i)}{\Pr[e_i \leq [6,000 - a - bI_i - gM_i] \text{ or } e_i \geq [20,000 - a - bI_i - gM_i]]} \quad (A.19) \\ &= \frac{f_e(e_i)}{1 - \Pr([6,000 - a - bI_i - gM_i] < e_i < [20,000 - a - bI_i - gM_i])} \end{aligned}$$

The first equality comes from the fact that K_i and e_i are linearly related through equation (10). In order to estimate the model in question, we will have to make some assumption about the error distribution. If we assume that $e_i \sim iidN(0, \mathbf{S}^2)$, then

$$f_{e_i|e_i} = \frac{1}{s} f\left(\frac{e_i}{s}\right) \quad (\text{A.20})$$

where $f(\cdot)$ denotes the standard normal pdf and

$$\begin{aligned} \Pr\left[6,000 - a - bI_i - gM_i < e_i < [20,000 - a - bI_i - gM_i]\right] \\ = \Phi\left(\frac{20,000 - a - bI_i - gM_i}{s}\right) - \Phi\left(\frac{6,000 - a - bI_i - gM_i}{s}\right) \end{aligned} \quad (\text{A.21})$$

Substituting these into equation (A.19) yields:

$$\begin{aligned} f_K(K_i | K_i \leq 6,000 \text{ or } K_i \geq 20,000) \\ = \frac{\frac{1}{s} f\left(\frac{K_i - a - bI_i - gM_i}{s}\right)}{1 - \Phi\left(\frac{20,000 - a - bI_i - gM_i}{s}\right) + \Phi\left(\frac{6,000 - a - bI_i - gM_i}{s}\right)} \end{aligned} \quad (\text{A.22})$$

The corresponding log-likelihood function becomes:

$$\begin{aligned} l(K, I, M; \mathbf{a}, \mathbf{b}, \mathbf{g}, \mathbf{s}) = -\frac{N}{2} \ln(2\pi) - N \ln(\mathbf{s}) - \frac{1}{2\mathbf{s}^2} \sum_{i=1}^N (K_i - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i)^2 \\ + \sum_{i=1}^N \ln\left[1 - \Phi\left(\frac{20,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}}\right) + \Phi\left(\frac{6,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}}\right)\right] \end{aligned} \quad (\text{A.23})$$

where the last term in the likelihood function controls for the truncation in the data base. This last term is analogous to what we saw in class when the tails of the distribution were truncated. In this case, it is the center of the distribution that has been truncated.

The second problem that I mentioned early on is that we do not really have a random sample from the tails of the distribution of K_i . Instead, we have 500 observations from the lower tail and 500 observations from the upper tail. To see that these are not the same, suppose that in the original (untruncated) population, there were 10,000 households with annual usage below 6,000 kWh and only 1,000 households with annual usage above 20,000. Thus, there would be 11,000 households in the population meeting the conditions: $K_i \leq 6,000$ and $20,000 \leq K_i$. A random sample from these 11,000 households would, on average, yield 91 households from high usage group and 909 from the lower usage households. This is because the chance of randomly drawing a high usage household from the truncated sample should be:

$$P_H = \frac{\Pr[K_i \geq 20,000]}{\Pr[K_i \leq 6,000] + \Pr[K_i \geq 20,000]} = \frac{1,000}{11,000} = .091 \quad (\text{A.24})$$

We can account for non-random nature of the sampling to reflect the population proportions. Suppose that the first N_H observations consist of the high usage observations. The adjusted log-likelihood function would then become:

$$\begin{aligned}
 l_{adj}(K, I, M; \mathbf{a}, \mathbf{b}, \mathbf{g}, \mathbf{s}) = & \frac{.5}{P_H} \left[\frac{N_H}{2} \ln(2\mathbf{p}) - N_H \ln(\mathbf{s}) - \frac{1}{2\mathbf{s}^2} \sum_{i=1}^{N_H} (K_i - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i)^2 \right. \\
 & + \sum_{i=1}^{N_H} \ln \left[1 - \Phi \left(\frac{20,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}} \right) - \Phi \left(\frac{6,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}} \right) \right] \\
 & + \frac{.5}{1 - P_H} \left[\frac{N_L}{2} \ln(2\mathbf{p}) - N_L \ln(\mathbf{s}) - \frac{1}{2\mathbf{s}^2} \sum_{i=N_H+1}^N (K_i - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i)^2 \right. \\
 & \left. + \sum_{i=N_H+1}^N \ln \left[1 - \Phi \left(\frac{20,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}} \right) - \Phi \left(\frac{6,000 - \mathbf{a} - \mathbf{b}I_i - \mathbf{g}M_i}{\mathbf{s}} \right) \right] \right]
 \end{aligned} \tag{A.23}$$

where $N_L = N - N_H$.

Note: In grading the exam, I did not deduct points for failing to specify the above likelihood function, but I did take a couple of points off if you failed to notice the second problem at all.

b. How would you assess the fit of your model. Be specific about your measure of fit, including its advantages and limitations.

There is no perfect answer in this case. One obvious suggestion would be to use a pseudo- R^2 statistic, similar to McFadden's statistic in the case of the probit model. That is:

$$LRI = 1 - \frac{l}{l_0} \tag{10.30}$$

where l_0 denotes the log-likelihood value if we constrain $\mathbf{b} = \mathbf{g} = 0$ in our model. If the model fits perfectly, then all of the probabilities would become 1 and $l = 0$, with $LRI=1$, whereas if the model did no better than one with only the intercept included, then $\frac{l}{l_0} = 1$,

so that $LRI=0$. The problem with this measure is that it is difficult to interpret the intermediate values of LRI. What is a "good fit?" and how much worse is $LRI = .5$ from an $LRI = .6$?