QUANTITATIVE FIELD EXAM

SECTION ONE:
Choose two out of the following three questions.

1. Let X be a random variable with probability function

$$f_X(x;\alpha) = \frac{\exp(\alpha x)\exp(-e^{\alpha})}{x},$$

for $x = 0, 1, \ldots$, and zero elsewhere.

(a) Calculate the mean and variance of X.

(b) Calculate the Cramer-Rao bound for α.

(c) Find, if possible, an unbiased estimator for α.

(d) Let $x_1, x_2, \ldots, x_N$ be a random sample from this distribution. Find a one dimensional sufficient statistic.

(e) Let $N = 100, \sum x_i = 200$, and $\sum x_i^2 = 610$. Find the maximum likelihood estimate for α.

(f) What is the standard error of the maximum likelihood estimator and give an estimate for this standard error.

(g) Construct an approximate 95% confidence interval for α.

2. Let $X_1, X_2, \ldots, X_N$ be independent random variables with normal distributions with mean μ and variance σ^2 (both mean and variance unknown).

(a) Find the moment generating function for a random variable with a normal distribution with mean μ and variance σ^2.

(b) Find the minimum variance unbiased estimator for σ^2.

(c) Find the maximum likelihood estimator for σ^2.

(d) Calculate the mean squared error for both the mvue and mle.

(e) Show that the difference between the mvue and mle converges to zero in probability as the sample size goes to infinity.

(f) Let $N = 10$, $\sum x_i = 30$ and $\sum x_i^2 = 100$. Calculate the mle.

(g) Estimate the standard error of the mle.

3. Let X be a random variable with probability density function

$$f_X(x;\lambda) = \lambda \exp(-x\lambda),$$

for $x > 0$ and zero elsewhere.

(a) Calculate the mean and variance of X.

(b) Calculate the mean and variance of X conditional on $X > 7$.

(c) Let $x_1, x_2, \ldots, x_N$ be a random sample from this distribution, with $N = 20$, $\sum x_i = 100$, and $\sum x_i^2 = 605$. Calculate the maximum likelihood estimate.

(d) Test the hypothesis that $\lambda = 1/4$ at the 5% level using a Wald test.

(e) Test the same hypothesis using a likelihood ratio test.

(f) Test the same hypothesis using a Lagrange multiplier (score) test.

Section Two

Answer two of the following three questions. Use a separate bluebook *for this section of the examination.*

I. Consider the variable y_i which, conditional on $\mathbf{x}_i$ follows an *iid* t-distribution

$$f(y_i|\mathbf{X}) = c^{-1}\left(1 + \frac{1}{\nu}\left(\frac{y_i - x_i'\beta}{\sigma}\right)^2\right)^{-(\nu+1)/2}.$$

A. What is the log likelihood function for y_i?

B. Derive the first-order conditions for the maximum likelihood estimator of β.

C. For this t-distributed model, under what conditions is the least squares estimator $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ BLUE?

D. Derive the first order conditions for the maximum likelihood estimator of the multiple regression model, with normal, *iid*, but heteroskedastic errors σ_i.

E. Compare these two specifications: the student-t, and the normal-heteroskedastic model. In what sense might we think of the regression model with student-t errors as implying heteroskedasticity?

F. Contrast the student-t specification with an ARCH(1) model. Most importantly, note the difference in the conditioning sets defining the conditional variances.

G. Both the student-t innovation model, and the ARCH(1) model display errors that are leptokurtotic — that is marginal distributions of the errors have heavy tails. What feature of the data distinguishes these models?

H. Why is the heteroskedasticity interpretation incorrect? Specifically, evaluate the heteroskedasticity interpretation in terms of the efficiency conditions (of least squares) that you described above. Hint: is the conditional variance well specified?

I. Compare the least squares estimator $\hat{\beta}$ with the maximum likelihood estimator β^* for the correctly specified (t-distributed) model. Which estimator is preferred, and why? Specifically comment on the efficiency of the two estimators.

II. Given $\mathbf{x}' = (\mathbf{x}_1', \mathbf{x}_2')$ jointly multivariate normal, the conditional distribution of $\mathbf{x}_1$ given $\mathbf{x}_2$ is given by

$$f(\mathbf{x}_1|\mathbf{x}_2) = \mathcal{N}\left(\boldsymbol{\mu}_1 + \Sigma_{12}\Sigma_{22}^{-1}(\mathbf{x}_2 - \boldsymbol{\mu}_2), \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}\right).$$

Use this result to analyze the following problems.

Suppose that three variables x, y, z are multivariate normally distributed, with means μ_x, μ_y, μ_z respectively. Denote the variances and covariances of these variables as σ_{ij}, for $i, j \in \{x, y, z\}$. Let the capital letters X, Y, Z denote n observations on each of the variables.

A. Denote the conditional mean of y as

$$E(y|x,z) = \alpha + \beta x + \gamma z.$$

Determine the constants α, β, γ as functions of the underlying mean and variance parameters μ_i, σ_{ij}, for $i, j \in \{x, y, z\}$.

B. Let $\varepsilon = y - E(y|x,z)$. Derive the (conditional) variance of ε, σ_ε^2.

C. What is the variance of the least squares estimator $\hat{\beta}$? Express your answer in terms of σ_ε^2 and matrices involving X and Z.

Next consider the following simple regression model

$$y_t = a + bx_t + u_t.$$

Maintain the assumption that (x, y, z) are trivariate normal as above.

D. What is the conditional distribution $f(y_t|X)$? You may express your answer in terms of the parameters α, β, and γ (as well as other parameters).

E. What is the conditional mean of the simple least squares estimator $E(\hat{b}|X)$?

F. What is the (conditional) variance of u_t, σ_u^2?

G. Show that $\sigma_\varepsilon^2 = V(\varepsilon|X, Z) \leq \sigma_u^2 = V(u|X)$.

H. Give two conditions under which $\sigma_\varepsilon^2 = \sigma_u^2$.

I. What is the conditional distribution of $\hat{b}$?

J. Show that $V(\hat{b})/\sigma_u^2 \leq V(\hat{\beta})/\sigma_\varepsilon^2$.

K. What is the conditional (given X) distribution of the t-statistic for the hypotheses that $b = \beta$?

L. For each of the following conditions, describe the likely effect on the hypothesis $b > \beta$. That is, discuss the size and the power of the t-test you analyzed in (K) relative to this hypothesis.

  i. $\gamma > 0$, $\sigma_{xz} < 0$.
  ii. $\gamma < 0$, $\sigma_{xz} < 0$.
  iii. $\gamma = 0$, $\sigma_{xz} < 0$.

III. Consider the conditional moment specifications

$$\begin{aligned} E(y_t|X, Z) &= X_t'\beta_1, \\ \mathrm{Var}(y_t|X, Z) &= Z_t'\beta_2. \end{aligned}$$

Here

$$X = \begin{pmatrix} X_1' \\ X_2' \\ \vdots \\ X_T' \end{pmatrix}, \text{ and } Z = \begin{pmatrix} Z_1' \\ Z_2' \\ \vdots \\ Z_T' \end{pmatrix},$$

are $(T \times k_1)$ and $(T \times k_2)$ matrices of observations.

A. Provide consistent estimators of β_1 and β_2.

B. Suppose that $x_{j1} = z_{j1} = 1, j = 1, \ldots, T$. Let $\beta_1^* = (\beta_{12}, \ldots, \beta_{1k})'$, and $\beta_2^* = (\beta_{22}, \ldots, \beta_{2k})'$, so that $\beta_1 = (\beta_{11}, \beta_1^{*\prime})'$ and $\beta_2 = (\beta_{21}, \beta_2^{*\prime})'$. Suppose we want to test the hypothesis that

$$H_2 : \beta_2^* = 0.$$

Describe the construction of a test statistic, and give its distribution under the null hypothesis.

C. Suppose that $k_2 = 2$. Can you provide an alternative test of H_2 that may be more efficient? Be explicit in describing the construction of the test and its distribution.

D. Suppose instead that Z_t consists of X_t and its interaction and quadratic terms. Provide an alternative interpretation of the test you are proposing in (B).

E. Based on part (B), suppose you reject the null H_2. Can you use this information to improve your estimator of β_1? If so, how are you improving it? Suppose instead that you accept the null hypothesis H_2. What does this tell you about the properties of your estimator for β_1? Again explain.

F. Next you want to test

$$H_1 : \beta_1^* = 0.$$

Construct a test of this hypothesis, based on your results from part (E). Specifically, explain how you construct a test statistic for the case that you accept H_2, and for the case that you reject H_2. Specify the limiting distribution of your test statistics for both cases.

G. What is a potential problem with the procedure you described in parts (E) and (F)? How does this affect your testing results in part (F)?

H. Explain why you can, or cannot reverse the procedures above. Can you first test H_1, and base your tests of H_2 on this outcome?

I. Suppose that $X_t = Z_t$. How can you test the hypothesis that $\beta_1 = \beta_2$?

Section Three — Use a separate bluebook when answering this Section.

Answer any <u>two</u> of the following <u>three</u> questions.

1. Consider the following model:

$$y_1 = y^* + \varepsilon_1 \tag{1}$$
$$y_2 = \beta y^* + \varepsilon_2 \tag{2}$$
$$y^* = x'\pi + u \tag{3}$$

where: y_1 is a household's measured annual income; y_2 is its measured annual consumption expenditures; y^* is the household's permanent income; x is a $K \times 1$ vector of exogenous determinants of permanent income; the scalar β and the $K \times 1$ vector, π, are unknown parameters; and ε_1, ε_2, and u are independently distributed disturbances with zero expectations and uncorrelated with x. You are given a random sample of N observations consisting of data on y_1, y_2, and x.

(a) Find the reduced form system of equations for the above model.

(b) Will OLS and GLS estimator of the unknown reduced form parameters from part (a) be the same? If so, why? If not, why not? Be precise in stating your answer.

(c) Are the unknown parameters, β and π, identified from the reduced form representation of the model? If so, why? If not, why not?

(d) Propose an Indirect Feasible GLS estimator of β and π, carefully defining what you mean by such an estimator, and characterize its asymptotic distribution based on the information provided for the above model and its sampling process. Be precise in characterizing your answer and what you are assuming that you know about the sampling process and stochastic properties of the model.

(e) Would your estimator of β and π in part (d) be more efficient if you knew that the variances of ε_1 and ε_2 were each *homoskedastic* and *proportional* to one another? Explain your answer.

2. Suppose you are trying to model whether or not consumers purchased a group of M goods. For the t^{th} consumer, let

$$y_{mt} = \begin{cases} 1, & \text{if consumer } t \text{ purchased the good } m \\ 0, & \text{if consumer } t \text{ did not purchase good } m \end{cases} \tag{4}$$

for $m = 1,\ldots,M$. The following expression determines, y^*_{mt}, the t^{th} consumer's "utility" for the good m:

$$y^*_{mt} = x'_{mt}\beta + \varepsilon_{mt}, \tag{5}$$

where x_{mt} is a $K \times 1$ vector of observed characteristics associated with good m, β is an unknown vector of parameters, and ε_{mt} is an unobserved disturbance, $m = 1,\ldots,M$. Assume that a consumer purchases the m^{th} good if and only if $y^*_{mt} > 0$. The following assumptions hold for ε_{mt}:

$$E(\varepsilon_{mt} | x_{mt}) = 0$$

$$\varepsilon_t = [\varepsilon_{1t}, \varepsilon_{2t}, \ldots, \varepsilon_{Mt}]' \sim N(0, \Sigma)$$

where Σ is a full covariance matrix, i.e., its off-diagonal elements cannot be assumed to be zero. (Assume throughout this problem that the all of the diagonal elements of Σ are equal to 1.) Finally, assume that you have data on $[(y_{1t}, x'_{1t}), (y_{2t}, x'_{2t}), \ldots, (y_{mt}, x'_{mt}), \ldots, (y_{Mt}, x'_{Mt})]$ for a random sample of $t = 1,\ldots,N$ consumers.

(a) Characterize the likelihood function for the above model and describe how you would use this function to obtain the maximum likelihood (ML) estimator of β and the unrestricted elements of Σ. [HINT: Let $\phi(\xi | \mu, \Omega)$ denote the multivariate normal density function for a $M \times 1$ disturbance vector, ξ, with mean, μ, and $M \times M$ variance matrix, Ω. What is the probability of observing a particular set of purchases for the M goods by the t^{th} consumer? Use this probability to characterize the likelihood for the purchases that the typical consumer is observed to make and "replicate" this for all of the N consumers in your sample.]

(b) Based on the above model, what is the expression for $E(y_{mt} | x_{mt})$?

(c) Using your answer to part (b), define a GMM estimator for β. Be precise in your characterization of the form of the GMM criterion function and how you would use it to obtain the GMM estimator of β. Use an arbitrary weighting matrix, W, in your formulation of the GMM criterion function.

(d) Given the assumptions made above, will the GMM estimator of β that you describe in part (c) be consistent? Given an explanation for your conclusion on this part.

(e) Would your GMM estimator of β be as efficient, asymptotically, as the ML estimator of β? Give an explanation for your conclusion.

3. This question has 3 parts. Be sure to answer all three parts.

(a) An analyst, the first one, is interested in estimating the following relationship,

$$y_t = \beta_1 x_{1t} + \beta_2 x_{2t} + \varepsilon_t, \tag{6}$$

for a random sample of $t = 1,\ldots,N$ observations, where she is quite sure that $E(\varepsilon_t | x_{2t}) = 0$. At the same time, she is initially worried that x_{1t} is endogenous with respect to the specification in (6). She estimates the coefficients in (6) using ordinary least squares (OLS) and finds that the resulting estimates of β_1 and β_2 concur with her theoretical priors as to what the sign and magnitude of these coefficients should be. She concludes that x_{1t} must be exogenous in the model characterized by equation (6). Evaluate her conclusion.

(b) A second analyst, certain that x_{1t} is endogenous, writes down the following second relationship

$$x_{1t} = \gamma_1 y_t + \gamma_2 x_{2t} + v_t \tag{7}$$

where he is certain that $E(v_t | x_{2t}) = 0$. The second analyst estimates (7) by OLS and finds that the resulting estimate of γ_1 is significantly different from zero. Armed with this result, the second analyst attacks the conclusions of the first analyst (who has published her findings in *Econometrica*) that x_{1t} is exogenous for the relationship given in (6). Evaluate the soundness of the second analyst's critique.

(c) As a newly minted assistant professor of economics, you decide to assess the "debate" between the first and second analyst. Your new "angle" on the problem is that you *know* that the following assumptions are true:

$$\beta_2 \neq 0, \gamma_2 = 0, \text{ and } Cov(\varepsilon_t, v_t) = 0, \forall t.$$

How would you use these assumptions to shed light (i.e., test) whether it is legitimate to treat x_1 as exogenous in the relationship given in (6)? Is the assumption that $Cov(\varepsilon_t, v_t) = 0$ crucial to your ability to make progress on resolving the dispute between the first and second analysts? Why or why not?

End of Section Three