University of Illinois	**Final Exam**	Department of Economics
Spring 2017	**Economics 574**	Roger Koenker

The exam is due Friday, May 12 at 5pm. Feel free to email or stop by my office if you feel that there are points that need clarification; I'll post corrections or clarifications on the course web page, if necessary. **Please answer both of the first two questions, and one of the last two questions.**

(1) In the 2002 574 exam I asked: "Suppose $\{y_1, \ldots, y_n\}$ are iid random variables, each normally distributed with mean μ and variance μ^2. Find the mle of μ and argue its consistency. Compare the asymptotic efficiency of the mle in this problem with that of the sample mean. This problem is related to estimating models of heteroscedasticity in linear regression which have parameters in common with the model for the conditional mean." A concise answer to this question is available from the archive of prior 574 exams. Suppose instead of considering the MLE, you decided that you wanted to investigate the GMM estimator of μ in these circumstances, formulate the GMM estimator and contrast it with the MLE.

(2) As a further example of the futility of moments, consider the problem of estimating the raw second moment $S \equiv \mathbb{E}X^2$ for X lognormal, i.e. $\log X \sim \mathcal{N}(\mu, \sigma^2)$. The sample analogue $\hat{S} = n^{-1} \sum X_i^2$ has mean $S = \exp\{2\mu + 2\sigma^2\}$, but note that, for $Z \sim \mathcal{N}(0, 1)$,

$$\begin{aligned} \mathbb{P}(X^2 > S) &= \mathbb{P}(\log(X^2) > 2\mu + 2\sigma^2) \\ &= \mathbb{P}(\mu + \sigma Z > \mu + \sigma^2) \\ &= \mathbb{P}(Z > \sigma) \\ &= 1 - \Phi(\sigma) \end{aligned}$$

Thus, for $\sigma = 4$ even when n is quite large we *almost never* see X's in the extremely long right tail of the distribution, and consequently the sample analogue estimator is quite poor. Provide some simulation evidence for this, and contemplate its implications for the law of large numbers. Hint: what would Pafnuty Chebyshev have to say?

(3) In the 2012 574 final exam there was a rather challenging question asking to evaluate the Hotelling tube approach to testing in a simple Box-Cox transformation model. In the exam archive you can find the original exam question, a brief "answer" and some R code used to construct the table appearing in the answer. Motivation for such methods is provided by Hansen (1996) Inference When a Nuisance Parameter Is Not Identified Under the Null Hypothesis, Econometrica, 64, 413-430, who considers several alternative approaches, based on simulated critical values. Suppose that you wanted to compare Hansen's SupLM with the Hotelling approach for the simple Box-Cox problem, design a small monte-carlo experiment to evaluate this and provide a concise description of the results. Most importantly, explain how the SupLM test is computed and any advantages or disadvantages you might see compared to the Hotelling form of the test.

(4) The empirical Bayes estimator for the Gaussian compound decision (aka sequence) model proposed by Martin and Walker (EJS, 8, 2188-2206) can be computed by MCMC Gibbs sampling as described in their (8a-b). An implementation in R is available from my Empirical Bayes web page to accompany the "Bakeoff" paper.

(a) Provide a exposition/tutorial detailing how/why this Gibbs sampler “works,” why does the posterior mean of the Markov chain provide a reasonable estimate of the vector of unknown θ’s? Try to provide some graphical evidence for this.

(b) Martin and Walker briefly mention the simulation results reported in the “Bakeoff” paper, but rather dismissively suggest that they shouldn’t be taken seriously since the Kiefer-Wolfowitz NPMLE isn’t “provably minimax.” This view seems to imply that there is a least favorable setting for which the Martin and Walker estimator would out-perform the NPMLE Bayes rule. Try to cook-up such a case.