University of Illinois	**Final Exam**	Department of Economics
Spring 2012	**Economics 574**	Professor Roger Koenker

The ultimate objective of field courses is (presumably) to prepare you to digest new material in future research. Thus, the exam is structured to allow you to demonstrate some facility for this, and in the process to learn some (marginally) fun, new stuff. The exam is due friday, May 11 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 webpage, if necessary.

1. Suppose you have random variables (X_1, X_2, X_3) with joint distribution F, with $\mathbb{V}X_i = 1$ for $i = 1, 2, 3$, and in addition satisfying $X_1 + X_2 + X_3 = 1$. What can be said about the possible values that can be taken by their three covariances?

2. In my "Some Notes on Hotelling Tubes" available at `http://bit.ly/IAuBZa` I've tried to sketch an elementary exposition of the 1939 Hotelling approach to a rather general class of irregular likelihood ratio testing problems. In the simplest parametric setting illustrating this approach we have a model of the form,

$$y_i = \beta \frac{x_i^\lambda - 1}{\lambda} + u_i \quad i = 1, \cdots, n,$$

   with $\{u_i\}$ iid $\mathcal{N}(0, \sigma^2)$. Suppose we have reason to restrict the Box Cox parameter $\lambda \in [-1, 1]$. We would like to test the hypothesis that β is zero, versus the general alternative that it is non-zero. The Hotelling formulation is based on the likelihood ratio statistic, assuming known σ,

$$T_n = \inf_{\lambda \in [-1,1]} \left\{ \sum_{i=1}^n (y_i - \hat{\beta}(\lambda)(x_i^\lambda - 1)/\lambda)^2 / \sum_{i=1}^n y_i^2 \right\}.$$

   (a) Explain the form of T_n, and explain how you would reduce the testing problem to this T_n even if you had additional linear covariates in the model for the conditional mean. (Note that $\hat{\beta}(\lambda)$ in the above expression is just the usual least squares estimate, which obviously depends on λ.)

   (b) Following the recipe in the "Notes" implement the test in R and conduct a small Monte Carlo experiment to validate its performance. (For the sake of comparison, use $m = 500$ for the discretization of the grid for $\|\gamma\|$.)

3. Lucien LeCam is generally recognized as one of the most profound statistical minds of the 20th century. His 1986 monograph *Asymptotic Methods in Statistical Decision Theory* is on a par with Wittgenstein's *Tractatus* for its depth and opacity. A much more readable acccount of some of his ideas is available in his paper "Maximum Likelihood: An Introduction" published in the International Statistical Review, but available at the class website in its preprint form at `http://bit.ly/IAwvcb`. Read the paper, reflect on the examples, ponder his principles – especially number zero, and finally write a brief comment/explanation on his proposed solution to the final problem discussed in the paper.