University of Illinois
Spring 2010

Final Exam
Economics 574

Department of Economics
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 14 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. (Yet another weird interpretation of 2SLS) Given a conventional linear model specification

$$y = Z\alpha + X\beta + u$$

Suppose we have instrumental variables W as well as exogonous variables X and we consider the 2SLS estimator of the vector α

$$\hat{\alpha}_{2SLS} = (Z^\top P_{M_X W} Z)^{-1} Z^\top P_{M_X W} y$$

(a) Explain the notation, show that this really *is* 2SLS, and explain how you would go about computing $\hat{\alpha}_{2SLS}$ using only OLS steps. (Note that this is just the usual Frisch-Waugh partial residual plot algebra combined with the usual 2SLS algebra.)

(b) Now consider an expanded version of the model,

$$y = Z\alpha + X\beta + W\gamma + v$$

Denote the least-squares estimator of γ in the expanded model, holding α at a fixed value, by $\hat{\gamma}(\alpha) = (W^\top M_X W)^{-1} W^\top M_X (y - Z\alpha)$. Let $\|x\|_A^2 = x^\top A x$. Carefully choosing a matrix A, show that $\mathrm{argmin}_\alpha \parallel \hat{\gamma}(\alpha) \parallel_A^2 = \hat{\alpha}_{2SLS}$.

(c) Interpret the result in (b.), that is, provide some sort of intuitive motivation for why it is (might be) true.[1]

2. (*How to get rich quick*) "Assume that we are hardened and unscrupulous types with an infinitely wealthy friend." Breiman (1961). The friend agrees, *à la* Samuelson at lunch,[2] to flip a coin biased in your favor on which you may bet any amount $B > 0$. The payoff of the bet will be

$$\begin{cases} B & \text{with probability } p > \frac{1}{2} \\ -B & \text{with probability } q = (1-p) \end{cases}$$

Suppose, perhaps because we have little imagination, we adopt a strategy of betting a constant fraction, θ, of our wealth W_n in each period. Thus, if we have initial wealth W_0,

$$W_n = W_0 (1+\theta)^{S_n} (1-\theta)^{F_n}$$

[1] This problem is motivated by the Chernozhukov and Hansen IV estimator for quantile regression.

[2] Samuelson (1963) describes an incident at lunch at MIT in which he asked his department chair, E Cary Brown, if he would be willing to make a 50-50 bet (a coin flip) for which if Brown called the coin correctly Samuelson would pay \$200, while if the coin was called incorrectly Brown would have to pay Samuelson \$100. Brown said "no, but I would be willing to accept 100 such bets." This engendered a small literature because Samuelson felt that Brown's response was deeply irrational. On further reflection Brown's response seems very reasonable, why risk losing \$100 at lunch? But who could resist the temptation to earn/win an expected 5,000, when the probability of losing money was roughly 1/2300. So Samuelson's published diatribe seems a bit unfair. (Typical expected utility computations, except for somewhat pathological utility functions, suggest that if taking 100 bets is good, then taking one should be too.)

where S_n is the number of "successes" in the flipping, and $F_n = n - S_n$ is the number of "failures."
A measure of our average rate of increase in wealth is

$$\log((W_n/W_0)^{1/n}) = \frac{S_n}{n}\log(1+\theta) + \frac{n-S_n}{n}\log(1-\theta)$$

Following Bernoulli, we might consider maximizing the expectation of this quantity:

$$E\log(W_n/W_0)^{1/n} = p\log(1+\theta) + q\log(1-\theta).$$

(a) Show $\theta^* = 2p-1$ is the maximizer.

(b) Show that there is a $\tilde{\theta} \in (\theta^*, 1)$ such that for $\theta > \tilde{\theta}$ we almost surely are ruined, that is we experience the humiliating event $W_n \to 0$, while for $\theta < \tilde{\theta}$ we will have W_n eventually exceeding any fixed bound.

(c) Taking $p = .6$ do some simulations to support your findings in parts (a) and (b). In the process find an approximate value for $\tilde{\theta}$ in part (b).

(d) Generalize the foregoing to a situation in which you are offered several coins with differing p_i's to bet on in each period and you can allocate your bets among these coins and the safe option.

3. Suppose Z is $N(\mu, \sigma^2)$ and $Z \equiv g(X) = \log(X - \alpha)$, then X is said to have a 3-parameter log-normal distribution.

(a) Show that the density of X is

$$f_X(x|\alpha,\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp\{-(\log(x-\alpha)-\mu)^2/2\sigma^2\}[x-\alpha]^{-1}$$

(b) Note that for any fixed α we have a standard normal mle problem for the parameters μ and σ^2 so we may concentrate the likelihood to get

$$L(\alpha|x) = K\hat{\sigma}^{-n}\prod_{i=1}^{n}[x_i - \alpha]^{-1}.$$

Show that for α sufficiently near $x_{(1)}$

$$\hat{\sigma}^2(\alpha) = n^{-1}\sum(\log(x_i - \hat{\alpha}) - \hat{\mu})^2 \leq (\log(x_{(1)} - \alpha))^2$$

so

$$L(\alpha|x) \geq K|\log(x_{(1)} - \alpha)|^{-n}\prod[x_i - \alpha]^{-1}.$$

and since,

$$\lim_{u\to 0}\frac{1}{|\log u|^n u} = \infty.$$

the likelihood becomes unbounded as $\alpha \to x_{(1)}$.

(c) Even though the mle is obviously silly, if literally interpreted in this case, the likelihood can still play a useful role in drawing inferences about the parameter α. In R draw a sample of 200 realizations from the 3 parameter lognormal with parameter $(\alpha, \mu, \sigma^2) = (1, 0, 1)$ and plot the concentrated likelihood as a function of α. Use the asymptotic theory of the the likelihood at the local maximum below the first order statistic to construct a confidence interval for α.

(d) Using the R function `optim` you can easily find the optimizing $\hat{\alpha}$ for this problem. To get an idea of how much it is costing you to estimate α, run a small simulation experiment and compare the bias and mean squared errors for the MLE's of all three parameters with the MLE performance of μ and σ^2 under the assumption that α is known. For simplicity, stick with the $n = 200$, $(\alpha, \mu, \sigma) = (1, 0, 1)$ case.