University of Illinois — Department of Economics
Spring 2012 — Econ 574 — Roger Koenker

Final Exam "Answers"

(1) This was admittedly a bit silly, a warm-up question as I said in class. Since, e.g.

$$\begin{aligned} V(X_1) &= V(1 - X_2 - X_3) \\ &= V(X_2 + X_3) \\ &= V(X_2) + V(X_3) + 2Cov(X_2, X_3) \end{aligned}$$

$Cov(X_2, X_3) = -1/2$, and the same for X_2 and X_3. It is also clear that the expectations of the X_i need to add to one, contrary to my original posting of the question.

(2) This was more challenging. The Hotelling tube idea seems to be a very attractive strategy for extending LR tests in a wide variety of non-regular cases.

(a) The basic problem here is that under the null the Box-Cox parameter, λ, isn't identified so we need some sort of special trick to evaluate the distribution of the likelihood ratio statistic. Provided we know σ^2, in our simple model the usual twice log likelihood ratio statistic would be

$$LR = \ell_n(\hat{\beta}(\lambda)) - \ell_n(0)$$

but here we find it convenient to divide through by $\ell_n(0)$ and ignore the superfluous -1. It is worth stressing at this point that any monotone transformation of the LR statistic like the one we have just made has the same rejection region as the original one so we are free to choose a transformation that makes computing the null distribution convenient. Hotelling's strategy is appealing in this way. When there are other covariates we can simply follow the usual Gauss-Frisch-Waugh approach and project the response and the Box-Cox'd covariate onto the space orthogonal to these covariates and proceed as before.

(b) Following the recipe in my "Notes" I ran a small simulation experiment with $\lambda \in \{-0.5, 0, 0.5\}$ and $n \in \{20, 50, 100, 500, 1000\}$ and local alternatives $\beta_n = \beta_0/\sqrt{n}$. Results of the experiment can be seen in the table below. (R code will be posted also.) It will be seen that the nominal 0.05 level of the test based on the Hotelling critical values is quite well maintained, and power is also quite respectible, somewhat better for alternatives with $\lambda = \pm 0.5$ than for $\lambda = 0$. I used standard Gaussian errors and standard lognormal x_i's, but it might be interesting to know how results would differ with other choices.

Table 1. Rejection frequencies for the Hotelling likelihood ratio test for a simple Box-Cox example. Tests are nominal level $\alpha = 0.05$. Local alternatives are employed of the form: $\beta_n = \beta_0/\sqrt{n}$.

	$\beta_0 = \mathbf{0}$			$\beta_0 = \mathbf{1}$			$\beta_0 = \mathbf{2}$		
	λ = -0.5	$\lambda = 0$	$\lambda = 0.5$	λ = -0.5	$\lambda = 0$	$\lambda = 0.5$	λ = -0.5	$\lambda = 0$	$\lambda = 0.5$
n = 20	0.056	0.058	0.049	0.313	0.193	0.182	0.781	0.459	0.380
n = 50	0.049	0.051	0.057	0.275	0.225	0.342	0.639	0.577	0.782
n = 100	0.063	0.048	0.056	0.350	0.261	0.281	0.840	0.637	0.704
n = 500	0.048	0.052	0.055	0.298	0.243	0.288	0.747	0.612	0.735
n = 1000	0.063	0.046	0.047	0.299	0.218	0.250	0.724	0.549	0.667

(3) I'll resist making a big commentary on the LeCam paper and simply observe that one-steps with a judicious choice of starting values are sometimes an attractive option. LeCam's innovation of reducing the problem to an approximately normal formulation of a pseudo likelihood based on robust estimates of location and scale is characteristically clever.