UCLA, Department of Economics, Patrik Guggenberger
Economics 203C (Spring, 2009), May 14, 2009

Problem Set VII (due Wed 5/27/09)

1. (Empirical coverage probabilities of delta method and bootstrap confidence intervals) Let X be distributed as normal with unknown mean $\mu = 0$ and known variance $\sigma^2 = 6$ and assume we are given a sample of n iid observations (X_i). The parameter of interest is $\theta_0 = \exp(\mu) = 1$ for which the estimator $\widehat{\theta}_n = \exp(\overline{X})$ is used, where $\overline{X}$ is the sample average.

For $n = 10$ simulate empirical coverage probabilities and average lengths of nominal 95%, two–sided symmetric and one–sided confidence intervals for θ_0 based on $R = 1000$ repetitions using

(i) the delta method

(ii) the nonparametric iid bootstrap based on $B = 250$ bootstrap resamples

(iii) method (ii) using $B = 1000$.

Compare and discuss your results. (In the one–sided case define "length" as the right endpoint of the confidence interval.)

2. True, questionable or false?

(i) You want to test $H_0 : \theta = 0$ against $H_1 : \theta > 0$. The test for H_0 is to reject if $T_n = \widehat{\theta}/s(\widehat{\theta}) > c$, where c is picked so that the type I error is α (the data are iid and as usually $\widehat{\theta}$ is a root$-n$ consistent estimator of θ and $s(\widehat{\theta})$ is a standard deviation estimator). You do this as follows. Using the non-parametric bootstrap, you generate B bootstrap samples, calculate the estimates $\widehat{\theta}^*$ on these samples and then calculate $T_n^* := \widehat{\theta}^*/s(\widehat{\theta}^*)$ B times. Let $q_n^*(1-\alpha)$ denote the $100(1-\alpha)\%$ quantile of the empirical distribution of T_n^*. You replace c with $q_n^*(1-\alpha)$ and thus reject H_0 if $T_n = \widehat{\theta}/s(\widehat{\theta}) > q_n^*(1-\alpha)$. Claim: For this test the null rejection probability converges to α.

(ii) Take a random sample $\{y_1, ..., y_n\}$ with $\mu = Ey_i$ and $\sigma^2 = var(y_i)$. The statistic of interest is the sample mean $T_n = n^{-1}\sum_{i=1}^n y_i$. Calculate the population moments ET_n and $var(T_n)$. Let $\{y_1^*, ..., y_n^*\}$ be a random sample from the empirical distribution function and let $T_n^* = n^{-1}\sum_{i=1}^n y_i^*$ be its sample mean. The first two bootstrap moments of T_n^* coincide with ET_n and $var(T_n)$, respectively.

(iii) In general, the optimal linear forecast is not a linear function.

(iv) A HAC covariance matrix estimator for a scalar estimator based on the truncated kernel can be negative.

(v) In HAC covariance matrix estimation we need the bandwidth S_n grow to infinity with sample size n growing to infinity in order for the variance of the HAC estimator to converge to zero.

3. (i) Prove that for Gaussian processes the optimal (with respect to squared loss) unrestricted forecast is given by the linear projection as long as a constant term is included among the variables on which the forecast is based. (See Hamilton, chapter 4.6)

(ii) Let $\{X_t\}$ be a stationary process with mean μ. Show that the optimal linear forecast of X_{n+h} based on $1, X_1, ..., X_n$ equals μ plus the optimal linear forecast of Y_{n+h} based on $Y_1, ..., Y_n$, where the zero mean process $\{Y_t\}$ is defined by $Y_t := X_t - \mu$ and where h is a fixed integer ≥ 1.

4. (HAC estimation) Consider the regression model

$$y_t = \mu + \delta x_t + u_t,$$

where

$$\begin{aligned} u_t &= \rho_1 u_{t-1} + p_2 u_{t-2} + e_t, \\ x_t &= .5x_{t-1} + \eta_t \in R, \end{aligned}$$

e_t and η_t i.i.d. $N(0,1)$ with $cov(e_t, \eta_t) = 0$, and $x_0 = u_0 = u_{-1} = 0$. For the true μ and δ being 0, do a Monte Carlo study to investigate the finite-sample size properties of a nominal 5% one-sided t-test of the null $H_0 : \delta = 0$ versus $H_1 : \delta > 0$. The regression is estimated by OLS and to estimate the standard deviation of $\widehat{\delta}_{OLS}$ (that is needed to compute the t-statistic), we use the Bartlett kernel HAC estimator, where the bandwidth S_T is chosen as the square root of the sample size. Report finite-sample rejection probabilities under the null for the 8 parameter combinations $T = 100$ and (ρ_1, p_2) equal to $(-.5, 0)$, $(0, 0)$, $(.5, 0)$, $(.9, 0)$, $(.99, 0)$, $(1.5, -.75)$, $(1.9, -.95)$, and $(0.8, .1)$ using $R = 2000$ samples each time. Compare your results to a t-test that uses White (1980, Ecta) standard errors. Discuss your findings. For further reading, consult Andrews, D.W.K. (1991): "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," Econometrica, 59, 817—858.