

NOTES

BOOTSTRAPPING MULTIVARIATE SPECTRA

Jeremy Berkowitz and Francis X. Diebold*

Abstract—We generalize the Franke-Härdle (1992) spectral-density bootstrap to the multivariate case. The extension is nontrivial and facilitates use of the Franke-Härdle bootstrap in frequency-domain econometric work, which often centers on crossvariable dynamic interactions. We document the bootstrap's good finite-sample performance in a small Monte Carlo experiment, and we conclude by highlighting key directions for future research.

I. Introduction

Inference for spectra is central to both theoretical and applied-time series econometrics, hence the interest in Franke and Härdle's (1992) method for bootstrap inference on estimated univariate spectral-density functions. In econometric work, however, interest often centers on crossvariable dynamic interactions, and hence *multivariate* spectra, as for example in Diebold et al. (1998) and Wright (1997). Unfortunately, the Franke-Härdle bootstrap has not yet been generalized to the multivariate case. In this paper we do so. We provide a brief summary of the univariate procedure in Section II to fix ideas and establish notation, and we provide the multivariate generalization in Section III. In Section IV, we perform a small Monte Carlo experiment, which suggests that the performance of the multivariate Franke-Härdle procedure is superior to conventional, first-order asymptotics for sample sizes relevant in econometrics. We offer concluding remarks in Section V.

II. The Univariate Bootstrap

The Franke-Härdle (1992) bootstrap is based directly on frequency-domain considerations. We briefly discuss the univariate case. Let $f(\omega_j)$ be the population spectral density of a covariance stationary stochastic process at frequency ω_j , let $I(\omega_j)$ be the unsmoothed (inconsistent) sample spectral density, and let $\hat{f}(\omega_j)$ be the smoothed (consistent) sample spectral density.

It is well known that in large samples the asymptotically pivotal statistic $2I(\omega_j)/f(\omega_j) \equiv \epsilon_j$ is independently and identically distributed as χ_2^2 , $\omega_j = 2\pi j/T, j = 1, \dots, T/2 - 1$. Equivalently,

$$I(\omega_j) = \frac{1}{2} f(\omega_j) \chi_2^2,$$

$\omega_j = 2\pi j/T, j = 1, \dots, T/2 - 1$. Franke and Härdle suggest replacing the population spectral density on the right side of the equation with a consistent estimator, drawing from the distribution of the ϵ_j 's to form a drawing of $I(\omega_j)$'s which are then smoothed, and repeating many times to build up an approximation to the distribution of the $\hat{f}(\omega_j)$'s.

Two variants of the bootstrap, both nonparametric in the sense that no distributional assumptions are made for the process whose spec-

trum is bootstrapped, may be performed. The first proceeds by making use of the known asymptotic distribution of the normalized sample spectral density. At bootstrap replication i , we draw $\{\epsilon_j^{(i)}\}_{j=1}^{(T/2-1)}$ from a χ_2^2 distribution and convert them into $\{I^{(i)}(\omega_j)\}_{j=1}^{(T/2-1)}$ via

$$I^{(i)}(\omega_j) = \frac{1}{2} \hat{f}(\omega_j) \epsilon_j^{(i)}.$$

Then, we smooth to obtain $\hat{f}^{(i)}(\omega_j), j = 1, \dots, T/2 - 1$. At the end we have $\hat{f}^{(i)}(\omega_j), j = 1, \dots, T/2 - 1, i = 1, \dots, R$, where R is the number of bootstrap replications. Finally, we form confidence intervals for $f(\omega_j), j = 1, \dots, T/2 - 1$, using the empirical percentiles of the $\hat{f}^{(i)}(\omega_j)$'s.

The second variant of the bootstrap is motivated by the fact that the χ_2^2 distribution for ϵ_j obtains only asymptotically, so that it may be preferable to sample with replacement directly from the empirical distribution of the "observed" $\epsilon_j^* = 2I(\omega_j)/\hat{f}(\omega_j), j = 1, \dots, T/2 - 1$.¹ We proceed as follows. At bootstrap replication i we draw $\{\epsilon_j^{(i)}\}_{j=1}^{(T/2-1)}$ by sampling with replacement from $\{\epsilon_j^*\}_{j=1}^{(T/2-1)}$ and convert them into $\{I^{(i)}(\omega_j)\}_{j=1}^{(T/2-1)}$ via

$$I^{(i)}(\omega_j) = \frac{1}{2} \hat{f}(\omega_j) \epsilon_j^{(i)},$$

$j = 1, \dots, T/2 - 1$. Then we smooth to obtain $\hat{f}^{(i)}(\omega_j), j = 1, \dots, T/2 - 1$. At the end we have $\hat{f}^{(i)}(\omega_j), i = 1, \dots, R$, from which we form confidence intervals for $f(\omega_j)$ using the empirical percentiles of the $\hat{f}^{(i)}(\omega_j)$ values, $j = 1, \dots, T/2 - 1$.

III. The Multivariate Bootstrap

Now we generalize the Franke-Härdle procedure to the N -dimensional multivariate case. In obvious parallel to the univariate discussion, let $F(\omega_j)$ be the population spectral-density matrix at frequency ω_j , let $I(\omega_j)$ be the unsmoothed sample spectral-density matrix, and let $\hat{F}(\omega_j)$ be the smoothed sample spectral-density matrix.

We proceed by exploiting the result that $I(\omega_j) \xrightarrow{d} W_N^c(1, F(\omega_j))$, an N -dimensional complex Wishart, $j = 1, \dots, T/2 - 1$.² Thus, in parallel with our univariate discussion, we define $F(\omega_j)^{-1/2} I(\omega_j) F(\omega_j)^{-1/2} \equiv \epsilon_j$, and we have that, in large samples,

$$\epsilon_j \stackrel{iid}{\sim} W_N^c(1, I),$$

or equivalently,

$$I(\omega_j) = F(\omega_j)^{1/2} W_N^c(1, I) F(\omega_j)^{1/2},$$

$$\omega_j = \frac{2\pi j}{T}, j = 1, \dots, \frac{T}{2} - 1.$$

¹ We first scale the ϵ_j^* 's so that their sample mean equals their known asymptotic population mean, 2, using the transformation $\epsilon_j^* = 2\epsilon_j^*/\bar{\epsilon}^*$.

² See Brillinger (1981).

Received for publication October 28, 1996. Revision accepted for publication November 26, 1997.

* Federal Reserve Board; and University of Pennsylvania and NBER, respectively.

We gratefully acknowledge helpful comments from the co-editor (Jim Stock), two referees, Bill Brown, Jin Hahn, Danny Quah, and Jonathan Wright. All remaining flaws are ours. We thank the National Science Foundation, the Sloan Foundation, and the University of Pennsylvania Research Foundation for support.

The above result suggests our first multivariate spectral-density bootstrap. At bootstrap replication i , we draw a sample spectral density as

$$I^{(i)}(\omega_j) = \hat{F}(\omega_j)^{1/2} W_N^{c(i)}(1, I) \hat{F}(\omega_j)^{1/2},$$

$j = 1, \dots, T/2 - 1, i = 1, \dots, R$, where $W_N^{c(i)}(1, I)$ is a draw from an N -dimensional complex Wishart distribution. We then smooth to obtain $\hat{F}^{(i)}(\omega_j), j = 1, \dots, T/2 - 1, i = 1, \dots, R$, from which we compute confidence intervals.

The second variant of the multivariate bootstrap proceeds by sampling $\{\epsilon_j^{(i)}\}_{j=1}^{(T/2-1)}$ with replacement from $\{\epsilon_j^*\}_{j=1}^{(T/2-1)}$, where $\epsilon_j^* = \hat{F}(\omega_j)^{-1/2} I(\omega_j) \hat{F}(\omega_j)^{-1/2}, j = 1, \dots, T/2 - 1$. We then form

$$I^{(i)}(\omega_j) = \hat{F}(\omega_j)^{1/2} \epsilon_j^{(i)} \hat{F}(\omega_j)^{1/2},$$

and we smooth to obtain $\hat{F}^{(i)}(\omega_j), j = 1, \dots, T/2 - 1, i = 1, \dots, R$, from which we compute confidence intervals.

IV. Monte Carlo Evaluation of the Multivariate Spectral Bootstrap

We perform a Monte Carlo experiment to help assess the performance of our multivariate spectral bootstrap. The data-generating process is a bivariate first-order Gaussian vector autoregression,

$$\begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \phi & 0 \\ 0.5 & 0.5 \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix}$$

$$\begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix} \sim N(0, I).$$


We vary persistence by setting $\phi = 0, 0.5, 0.7, 0.9, 0.95, 0.97$, and 0.99 , and for each value of ϕ we consider six sample sizes, $T = 140, 200, 400, 600, 800$, and $1,000$. At each design point and at each Monte Carlo replication i , we generate a realization, $\{y_t^{(i)}\}_{t=1}^T$, and calculate a consistent estimate of the spectral-density matrix, $\hat{F}(\omega)$, at $\omega = \pi/2$.³ We perform 2,000 Monte Carlo replications, which yield the exact finite-sample distributions of the elements of $\hat{F}(\omega)$ (that is, estimates of the two spectra, the cospectrum, and the quadrature spectrum, $\hat{f}_1(\omega), \hat{f}_2(\omega), \hat{c}(\omega)$, and $\hat{q}(\omega)$). Then, for each of the four estimated elements of the spectrum, we compare the actual finite-sample distribution to two approximations: one delivered by conventional Gaussian asymptotics and the other by the multivariate spectral bootstrap.⁴

The Gaussian approximation is simple and follows from the well-known result that under regularity conditions the standardized elements of the estimated spectral-density matrix are asymptotically Gaussian (e.g., Brockwell and Davis, 1992, 447–448). The bootstrap approximation is also simple, when used in one-time applications to real data; one merely follows the recipe we provided in Section 3. The computational burden of Monte Carlo evaluation of the bootstrap, however, quickly becomes prohibitive. Thus, we do not implement the bootstrap at each Monte Carlo replication; instead, precisely following Franke and Härdle (1992, p. 128), we adopt an alternative and clever,

³ Following Franke and Härdle, we compute the estimate by smoothing the periodogram with a Bartlett-Priestley window. We select the bandwidth automatically using the method of Beltrao and Bloomfield (1987).

⁴ We use the variant of the bootstrap based on sampling from the complex Wishart.

FIGURE 1.—EXACT SAMPLING DISTRIBUTION OF ESTIMATED SPECTRAL DENSITY, AND TWO APPROXIMATIONS $\phi = 0.95$



Notes: $f_1(\omega)$ and $f_2(\omega)$ are the two spectra, $c(\omega)$ is the cospectrum, and $q(\omega)$ is the quadrature spectrum. Dashed lines are actual empirical CDFs based on 2,000 Monte Carlo iterations. Solid lines are asymptotic normal approximations, and dotted lines are bootstrap approximations based on 2,000 bootstrap iterations from a representative sample. See text for details.

if second-best, strategy. For each sample size, we obtain a typical realization on which to perform bootstrap replications, which we do only once. We obtain the typical realization by generating nine samples and choosing the one for which the integrated squared error (ISE) of the estimated spectrum with respect to the true spectrum,

$$ISE = \sqrt{\sum_{\omega_j} \left(\frac{\hat{f}_1 - f_1}{f_1} \right)^2 + \sum_{\omega_j} \left(\frac{\hat{f}_2 - f_2}{f_2} \right)^2 + \sum_{\omega_j} \left(\frac{\hat{c} - c}{c} \right)^2 + \sum_{\omega_j} \left(\frac{\hat{q} - q}{q} \right)^2},$$

achieves its median, where $\omega_j = 2\pi j/T, j = 1, \dots, T/2 - 1$. We then implement the bootstrap on that realization, from which we obtain our bootstrap approximations to the actual finite-sample distributions of the elements of $\hat{F}(\omega)$.

We display the results for a realistic value of persistence ($\phi = 0.95$) in figure 1. Evidently, the Gaussian approximations are comparatively poor guides to the finite-sample distributions in small samples. The bootstrap approximations, in contrast, fare quite well for the most part. And, of course, as the sample size increases, all distributions become closer.

One can only go so far with visual assessments of the sort contained in figure 1. Response surfaces are preferable for a number of reasons. First, they enable us to summarize the results of the entire experiment, rather than just a single design point. This helps us understand how the comparative performance of the bootstrap approximation varies with (ϕ, T) . Second, response surfaces help us to reduce the specificity of the Monte Carlo experiment by facilitating generalization to other (ϕ, T) values. Finally, response surfaces force us to adopt an explicit measure of divergence between the true sampling distribution and the various approximations, which is helpful for making precise assessments.

TABLE 1.—RESPONSE SURFACE ESTIMATES DEPENDENT VARIABLE: (ISE^G–ISE^B)


Coefficient on:	$T^{-1/4}$	$\phi T^{-1/4}$	$T^{-1/2}$	$\phi T^{-1/2}$	R ²
Component of spectral density:					
$f_1(\omega)$	-0.166 (0.157)	-0.116 (0.288)	0.447 (0.767)	1.56 (1.26)	0.49
$f_2(\omega)$	-0.002 (0.085)	-0.310 (0.157)	-0.059 (0.429)	1.95 (0.779)	0.56
$c(\omega)$	0.018 (0.092)	-0.212 (0.210)	-0.216 (0.383)	1.07 (0.991)	0.14
$q(\omega)$	-0.109 (0.038)	-0.038 (0.058)	0.524 (0.154)	0.135 (0.261)	0.48

Notes: $f_1(\omega)$ and $f_2(\omega)$ are the two spectra, $c(\omega)$ is the cospectrum, and $q(\omega)$ is the quadrature spectrum. We report ordinary least-squares parameter estimates, based on an ordinary least-squares response surface regression with sample size 42, which is the number of design points. Newey-West standard errors, computed using a Bartlett window, appear in parentheses.

We model the difference between the integrated squared error of the Gaussian and bootstrap approximations (relative to the true sampling distribution) as a function of ϕ and T . We have seen that as $T \rightarrow \infty$, the difference in approximation accuracy of the Gaussian and bootstrap approximations appears to vanish, which suggests expansion in negative powers of T . Some experimentation revealed that expansion in powers of $T^{-1/4}$ provides a good approximation. Thus, for each of the four elements of the spectral-density matrix, we estimate a response surface of the form

$$(ISE^G - ISE^B) = T^{-1/4}(\beta_0 + \beta_1\phi + \beta_2T^{-1/4} + \beta_3\phi T^{-1/4}),$$

where ISE^G and ISE^B denote the ISE of the Gaussian and bootstrap approximations with respect to the true finite-sample distribution. We report the least squares estimates in table 1, together with R² values and robust standard errors. Many of the estimated parameters are statistically significant, and three of the four R² values are in the neighborhood of 50%.

FIGURE 2.—RESPONSE SURFACES DEPENDENT VARIABLE: (ISE^G – ISE^B)

Notes: $f_1(\omega)$ and $f_2(\omega)$ are the two spectra, $c(\omega)$ is the cospectrum, and $q(\omega)$ is the quadrature spectrum. Response surfaces are fitted values from ordinary least-squares regressions with sample size 42, the number of design points. The lefthand-side variable is the integrated squared error of the Gaussian CDF minus the integrated squared error of the bootstrapped CDF. The righthand-side variables are sample size and ϕ .

We display the four estimated response surfaces in figure 2. The surfaces are usually positive, indicating that, throughout much of the range of the experimental design, the bootstrap approximation is better than the Gaussian approximation. The graphs also reveal the nature of the dependence of the ISE difference on T . It is interesting to note that the comparative accuracy of the bootstrap approximation is best in the region of most importance for applications such as macroeconomics: large ϕ , small T . Although it is true that the *absolute* performance of both the Gaussian and bootstrap approximations deteriorates as ϕ approaches 1, the performance of the bootstrap approximation deteriorates more slowly.

V. Conclusions and Directions for Future Research

We have generalized the Franke-Härdle (1992) spectral-density bootstrap to the multivariate case and provided some evidence of good finite-sample performance in a Monte Carlo experiment. As with its univariate counterpart, the multivariate spectral bootstrap is intuitive and appealing, because it is based directly upon the well-known asymptotic distribution of the periodogram ordinates. The approach complements recent work by Politis and Romano (1992), who generalize the Künsch (1989) and Liu and Singh (1992) moving-blocks bootstrap so that it delivers valid inference for spectra, in contrast to the original Künsch-Liu-Singh procedure.

The obvious item at the top of our agenda for future work is proof of first-order asymptotic validity of the multivariate spectral bootstrap. We conjecture that it is first-order valid, because it is a natural extension of the Franke-Härdle univariate bootstrap, which is first-order valid. Our Monte Carlo analysis supports that conjecture, as the Gaussian and bootstrap approximations appear to get closer as sample size grows. We conjecture as well that the multivariate spectral bootstrap may offer second-order refinements, because it is based on an asymptotically pivotal statistic. Our Monte Carlo analysis also supports that conjecture, as the bootstrap approximation appear superior in small samples. Our Monte Carlo analysis is not exhaustive, however, and could be improved in several ways. Formal proofs of the asymptotic properties of the bootstrap, moreover, are not at all trivial and remain elusive, as the univariate proof methods do not readily generalize to the multivariate case.

REFERENCES

- Beltrao, K. I., and Bloomfield, P., (1987), "Determining the Bandwidth of a Kernel Spectrum Estimate," *Journal of Time Series Analysis* 8, 21–38.
- Brillinger, D. R., *Time Series: Data Analysis and Theory* (Expanded Edition), (New York: McGraw-Hill, 1981).
- Brockwell, P. J., and Davis, R. A., *Time Series: Theory and Methods*, Second Edition, (New York: Springer-Verlag, 1992).
- Diebold, F. X., Ohanian, L. E., and Berkowitz, J., "Dynamic Equilibrium Economies: A Framework for Comparing Models and Data," *Review of Economic Studies* 65 (1998), 433–452.
- Franke, J., and Härdle, W. (1992), "On Bootstrapping Kernel Spectral Estimates," *Annals of Statistics* 20, 121–145.
- Künsch, H. R., "The Jackknife and the Bootstrap for General Stationary Observations," *Annals of Statistics* 17 (1989), 1217–1241.
- Liu, R. Y., and Singh, K., "Moving Blocks Jackknife and Bootstrap Capture Weak Dependence," in R. LePage and L. Billard (Eds.), *Exploring the Limits of the Bootstrap*, (New York: John Wiley, 1992).
- Politis, D. N., and Romano, J. P., "A General Resampling Scheme for Triangular Arrays of α -Mixing Random Variables with an Application to the Problem of Spectral Density Estimation," *Annals of Statistics* 20 (1992), 1985–2007.
- Wright, J. H. (1997), "Frequency Domain Inference for Impulse Responses," Manuscript, Department of Economics, University of Virginia.