WATER RESOURCES RESEARCH, VOL. 32, NO. 6, PAGES 1875–1882, JUNE 1996

The moving blocks bootstrap versus parametric time series models

Richard M. Vogel
Department of Civil and Environmental Engineering, Tufts University, Medford, Massachusetts

Amy L. Shallcross
Camp Dresser & McKee Incorporated, Edison, New Jersey

Abstract. The application of a parametric time series model to a water resources problem involves selecting a model and estimating its parameters, both steps adding uncertainty to the analysis. The moving blocks bootstrap is a simple resampling algorithm which can replace parametric time series models, avoiding model selection and only requiring an estimate of the moving block length. The moving blocks bootstrap resamples the observed time series using approximately independent moving blocks. A Monte Carlo experiment is performed involving the use of a time series model to estimate the storage capacity S of a surface water reservoir. Our results document that the bootstrap always produced storage estimates with lower root-mean-square-error than a parametric alternative, even when no model error is introduced into the parametric scheme. These results suggest that the moving blocks bootstrap can provide a simple and attractive alternative to more complex multivariate ARMA models.

Introduction

Time series models are intended to mimic both the deterministic and the random nature of hydrologic variables in both space and time. The spatial and temporal structure of hydrologic time series is extremely complex, leading to an evolution in parametric time series methods from the relatively simple univariate and multivariate models, introduced by *Fiering* [1967] and others, to the more sophisticated disaggregation models [*Lane*, 1979; *Salas et al*., 1980; *Stedinger and Vogel*, 1984; *Grygier and Stedinger*, 1988]. *Salas et al*. [1980] and *Salas* [1993] review parametric time series methods in water resources. Parametric methods differ significantly from their nonparametric alternatives because parametric methods require assumptions regarding (1) the marginal probability distributions of the variables and (2) the spatial and temporal covariance structure of the variables. Nonparametric methods simply retain the empirical structure of the observed variables. More importantly, parametric methods require estimates of various model parameters which nonparametric methods can either minimize or avoid altogether. *Stedinger and Taylor* [1982b] and *Vogel and Stedinger* [1988] have documented that errors arising from parameter estimation of time series models can easily overwhelm issues of model choice, given the short hydrologic records usually available. Can nonparameteric time series methods such as the bootstrap reduce uncertainties due to parameter estimation and model choice while coincidently preserving important statistical properties of the flow vectors?

The Bootstrap and the Moving Blocks Bootstrap

Nonparametric time series methods are becoming increasingly popular in water resources as evidenced by their coverage in a recent textbook [*Helsel and Hirsch*, 1992] and review article [*Lall*, 1995]. The field of nonparametric statistics is under rapid growth. A few years ago, one approach, known as the bootstrap, was only described in the mathematics research literature. Now there are textbooks describing the bootstrap [*LePage and Billard*, 1992; *Efron and Tibshirani*, 1993; *Hjorth*, 1994].

The bootstrap is the simplest technique for simulating the probability distribution of any statistic, without making any assumptions or estimating any parameters. It is a good example of a new class of nonparametric statistical methods which substitute computer intensive computations for complex mathematical (parametric) models. Bootstrapping amounts to resampling a record, with replacement, to generate B bootstrap samples, from which one can simulate B estimates of a given statistic, leading to an empirical probability distribution of the statistic. Suppose one wishes to estimate the empirical cdf of a statistic $\hat{\theta}_i$ which is estimated from a given sample x_i, $i = 1, \cdots, n$ which we denote X. Each observation x_i is resampled, with replacement, with an equal probability of $1/n$. The sample X continues to be resampled with replacement B times, until B bootstrap samples X_i, $i = 1, \cdots, B$ are obtained. Each bootstrap sample X_i yields a bootstrap estimate of the statistic θ leading to the B bootstrap estimates $\hat{\theta}_i$, $i = 1, \cdots, B$, the desired result. The bootstrap is simple to implement using a spreadsheet [*Willemain*, 1994], an important fact in water resource engineering because many models are impemented using spreadsheets.

The bootstrap can be used as a nonparametric time series model. One simply resamples, with replacement, from the historical record. The challenge is to resample the records, in such a way as to assure that the temporal and spatial covariance structure of the original time series is preserved. For example, suppose a time series of annual streamflow arises from an AR(1) process with serial correlation ρ. For small values of ρ, sequences of 5-year flows will be approximately independent, and hence applying the bootstrap to 5-year flow traces will preserve, approximately, the serial correlation structure of flow

Paper number 96WR00928.
0043-1397/96/96WR-00928$09.00

series. Resampling λ-year blocks, as described above, is known as the moving-blocks bootstrap first introduced by Künsch [1989] and discussed by *LePage and Billard* [1992] and *Efron and Tibshirani* [1993]. In a moving-blocks bootstrap, one chooses a block length $\lambda \approx n/k$, where n is the record length and k is the number of blocks to resample. The idea is to choose a large enough block length λ so that observations more than λ time units apart will be nearly independent.

Again, suppose the original time series is AR(1) with serial correlation ρ. By sampling blocks of length λ, we retain the original correlation ρ among the observations within each block, yet adjacent blocks are uncorrelated. Using a circular definition of serial correlation, the average effective serial correlation of the moving-blocks bootstrap samples will be

$$\rho_{\text{effective}} = \left(\frac{k}{n}\right)0 + \left(\frac{n-k}{n}\right)\rho = \left(\frac{\lambda - 1}{\lambda}\right)\rho \qquad (1)$$

because there are k adjacent block intersections, leaving the remaining $n - k$ observations in their original serial order. As the block length λ increases, $\rho_{\text{effective}}$ approaches ρ; however, increasing λ leads to a smaller number of blocks $k = n/\lambda$ left to bootstrap. Equation (1) also shows that for a given block length λ, bootstrap flow traces will always have lower serial correlation than the parent trace and that downward bias does not disappear with increasing sample size. However, larger sample sizes allow for larger block lengths λ because there are more blocks $k = n/\lambda$ to bootstrap.

Some Applications of the Bootstrap for Simulating Streamflows

This section reviews a few studies which are similar in spirit to this study in the sense that they use nonparametric methods to simulate streamflows. For a more detailed survey of applications of nonparametric methods to time series problems in water resources, see *Lall* [1995] and *Lall and Sharma* [1996]. The bootstrap is not really new to hydrology. *Sudler* [1927] recommended a method which amounts to shuffling and dealing a deck of cards, in which each card contains an observed annual streamflow. His approach produces equally likely sets of independent streamflow traces. Sudler's approach is equivalent to the modern bootstrap; however, the bootstrap is now much easier to implement using a computer. More recently, *Lall and Sharma* [1996] introduce a bootstrap method for resampling monthly streamflows which preserves the dependence in a probabilistic sense. Their bootstrap is termed a nearest neighbor bootstrap because it searches the historical record to find the historic nearest neighbors and subsequently resamples their successors to preserve the empirical dependence of the flow trace. *Zucchini and Adamson* [1988] used a simple bootstrap to estimate confidence intervals associated with annual inflows to a reservoir system. *Hausman* [1990] showed that use of a moving-blocks bootstrap (with $\lambda = 2$) led to about the same storage-reliability-yield relationship as did an AR(1) annual flow model for the Boston water supply system.

Nonparametric resampling algorithms, such as the bootstrap, are much simpler to implement than their parametric alternatives. It is for this reason that the Bureau of Reclamation of the U.S. Department of the Interior often uses a nonparametric resampling scheme termed the indexed sequential method (ISM) instead of their own parametric disaggregation model (*W. Cheney*, personal communication, 1995) termed LAST [*Lane*, 1979]. The ISM method generates equally likely streamflow traces from the historical record, by using short overlapping traces. For example, from a 50-year flow record, one selects the flows in years 1–10 for realization number 1, the flows in years 2–11 for realization number 2, and so on. This method suffers from a number of disadvantages including (1) intercorrelation among the resampled flow records, (2) limited number of potential resampled flow records, and (3) length of resampled flow records must be much shorter than the original record length. In spite of these significant shortcomings, *Labadie et al*. [1987] and *Kendall and Dracup* [1991] conclude that the ISM procedure is comparable to parametric alternatives: it is certainly much simpler. The ISM is not a moving-blocks bootstrap because new sequences are not assembled by randomly piecing together different blocks. Each new ISM sequence contains every historical streamflow, unlike the bootstrap sequences which do not usually contain every historical streamflow.

Limitations of the Bootstrap

A bootstrap trace, unlike a trace from a parametric time series model, is limited to the original historic observations. Since the bootstrap will never generate an observation either larger or smaller than the maximum or minimum historical observation, the bootstrap is not useful for examining the probability distribution of the largest or smallest observation, unless the sample size n is greater than the planning horizon N. Similarly, if one bootstraps annual streamflows in a reservoir application, one never achieves a smaller flow than the minimum historical flow, hence the bootstrap is not useful for systems dominated by within-year storage requirements because the bootstrap could never produce a worse drought than that experienced historically. Yet the bootstrap does have potential for systems either dominated by over-year storage requirements or for systems with a combination of over-year and within-year requirements. For such systems the changes in the serial structure of the annual flows is enough to provide a rich set of alternate drought sequences with which to evaluate reservoir operations. For example, reordering of the historical sequence using the bootstrap can lead to two or more of the lowest flows occurring consecutively, producing a drought which is much larger than the most severe historical drought.

Monte Carlo Experiment

Given the increasing need to evaluate the operations of complex water resource systems, time series methods are likely to find increasing usage. The question remains whether an analyst should consider selecting, estimating, verifying, and validating a parametric time series model using the procedures outlined by *Stedinger and Taylor* [1982a] or apply a nonparametric method where model selection, parameter estimation, and model validation issues are grossly simplified in comparison to the parametric method. We perform a simple experiment to document the advantages and limitations of nonparametric time series methods over their parametric alternatives, for a simple reservoir design problem.

Assume a water resource engineer is faced with the problem of estimating the storage capacity S_p associated with a surface water reservoir. Here S_p is the capacity of a storage reservoir required to supply a constant no-failure annual yield of $\alpha\mu$ over an N-year planning period with reliability p. Here μ is the mean annual inflow to the reservoir and α is the level of development or the fraction of the mean annual flow which is delivered by the reservoir. The traditional solution to this problem has been to route the historical streamflows though a hypothetical reservoir system, releasing the yield $\alpha\mu$ and choosing S as the reservoir capacity which would have just delivered the yield, without failure, over the n-year historical record. This solution was first proposed by *Rippl* [1883] and has been used in the design of almost every reservoir in the United States. Unfortunately, a single historic streamflow trace only leads to a single value of S which has a reliability of 50%, on average [*Stedinger et al.*, 1983]. Time series models are useful for the generation of alternative yet likely streamflow traces and, subsequently, routing them through a storage reservoir, leading to a probability distribution of S.

Vogel and Stedinger [1987] document relationships for estimating S_p when inflows follow an AR(1) lognormal model. Those analytic relationships take the form

$$S_p = f(m,\ C_v,\ \rho,\ \alpha,\ p,\ N) \qquad (2)$$

where

$$m = \frac{(1-\alpha)\mu}{\sigma} = \frac{(1-\alpha)}{C_v};$$

- μ mean annual inflow to reservoir;
- σ standard deviation of annual inflows;
- ρ lag-one correlation of annual inflows;
- C_v coefficient of variation of annual inflows;
- α demand as a fraction of mean annual inflow;
- p yield reliability over N-year period;
- N length of planning period.

Equation (2) is used to obtain population values of S_p to compare with sample estimates of S_p derived from the use of both parametric and nonparametric time series models.

Experimental Design

All of the experiments follow the same general procedure. First 20,000 sets of "historic" flow traces of length $n = 40$ and 80 were generated from an AR(1) lognormal model with $\mu = 1.0$, $\sigma = 0.4$, and $\rho = 0.0$ and 0.3. Each flow trace of length n is assumed to be a flow trace which would be available to a stochastic hydrologist, who would then use that trace to fit a time series model and estimate S_p. Two approaches to estimating S_p are taken: (1) the parametric stochastic hydrologist (PSH) fits an AR(1) lognormal model to the "historic" flow traces and (2) the nonparametric stochastic hydrologist (NSH) bootstraps the "historic" flow traces.

One Monte Carlo Experiment Using the Parametric Time Series Model

The PSH is assumed to use the ordinary product-moment estimators

$$\hat{\mu} = \bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i,$$

$$\hat{\sigma}^2 = s^2 = \frac{1}{(n-1)}\sum_{i=1}^{n} (x_i - \bar{x})^2,$$

$$\hat{\rho} = \sum_{i=1}^{n} (x_{i+1} - \bar{x})(x_i - \bar{x}) \Bigg/ \sum_{i=1}^{n} (x_i - \bar{x})^2$$

to fit an AR(1) lognormal model to each "historical" flow trace. The fitted AR(1) lognormal model is then used to generate 1000 "synthetic" flow traces of length $N = 40$ and 80. Each "synthetic" flow trace is then routed through a reservoir system using *Rippl*'s [1883] mass curve approach, leading to 1000 estimates of required storage S_i, $i = 1, \cdots, 1000$ which are ranked and sorted to obtain a single estimate $\hat{S}_p$.

One Monte Carlo Experiment Using the Bootstrap

The NSH uses the bootstrap to resample, with replacement, from the "historic" flow trace, producing a total of $B = 1000$ "synthetic" flow traces of length $N = 40$ and 80. Similar to the PSH, each "synthetic" flow trace is then routed through a reservoir system using *Rippl*'s [1883] mass curve approach, leading to 1000 estimates of required storage S_i, $i = 1, \cdots, 1000$ which are ranked and sorted to obtain a single estimate $\hat{S}_p$.

Summary of Monte Carlo Experiments

Each Monte Carlo experiment results in a single estimate of $\hat{S}_p$ for the $p = 5$, 25, 50, 75, and 95th percentiles, for the PSH and the NSH. For example, $\hat{S}_{50}$ represents an estimate of the storage capacity that has a 50% reliability of meeting its yield, without failure, in future N-year planning periods. Experiments are performed for levels of development $\alpha = 0.7$ and 0.9. A total of 20,000 Monte Carlo experiments are performed, leading to 20,000 values of $\hat{S}_p$ for each yield, for both the PSH and the NSH. Note that each of those 20,000 Monte Carlo experiments involve generating 1000 synthetic flow traces which are in turn used to estimate S_p. Figures 1–6 summarize the results of these experiments in terms of the percent bias and percent root-mean-square-error (rmse) associated with $\hat{S}_p$. Here percent bias is defined as $(E[\hat{S}_p] - S_p)/S_p$ and percent rmse is defined as $E[(\hat{S}_p - S_p)^2]^{1/2}/S_p$ with the true values of S_p obtained by using the relations given by *Vogel and Stedinger* [1987].

Results

The simple bootstrap ($\lambda = 1$). Figures 1 and 2 illustrate the bias and rmse associated with $\hat{S}_p$ for independent flow traces ($\rho = 0$) with $n = N = 40$ and 80, respectively. The bootstrap always produced estimates of S_p with lower root-mean-square-error than did the parametric model. As expected, both procedures exhibit minimal bias, with any bias nearly disappearing for the larger sample size $n = N = 80$. Figures 1 and 2 document the advantage of the bootstrap over one parametric alternative, when the block length is chosen correctly.

Figures 3 and 4 illustrate the bias and rmse associated with $\hat{S}_p$ for correlated flow traces ($\rho = 0.3$) with $n = N = 40$ and 80, respectively. Now the bootstrap produced estimates of S_p with much greater (downward) bias than the parametric model, yet still, the rmse associated with $\hat{S}_p$ was about equal for both methods, indicating that the bootstrap produces estimates of S_p with much lower variance than the parametric method. Note that mean-square-error of an estimator is equal to the sum of its variance and its bias squared.

The simple bootstrap results in significant downward bias in $\hat{S}_p$, for correlated samples, because it generates "synthetic" flow records which lack serial correlation. *Phatarfod* [1986] has shown that estimates of S_p for correlated flows are inflated by a factor of $(1 + \rho)/(1 - \rho)$ over independent flows, which amounts to an almost doubling for $\rho = 0.3$. Therefore the extraordinary downward bias exhibited in Figures 3 and 4, for

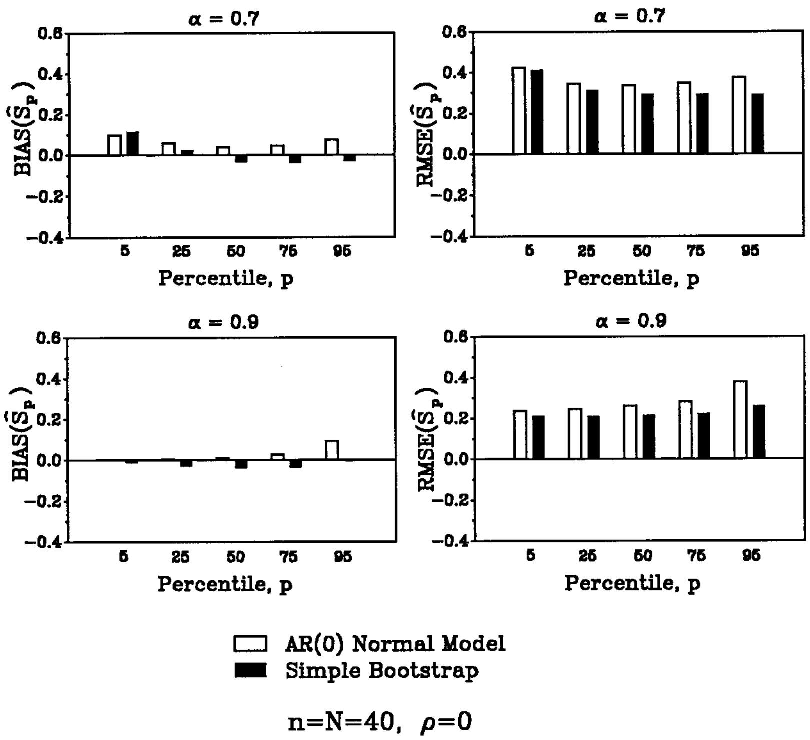

Figure 1. Comparison of bias and root-mean-square-error (rmse) associated with $\hat{S}_p$ using an AR(0) normal model and a simple bootstrap with $\lambda = 1$, when $\rho = 0$ and $n = N = 40$.

the bootstrap, is expected, when the block length is too short to preserve the serial correlation of the flows.

Results

The moving-blocks bootstrap ($\lambda > 1$). Figures 1–4 document that the bootstrap has potential as a time series model if methods were available to introduce the proper serial correlation into the bootstrapped "synthetic" time series. One simple approach employed here is to bootstrap $\lambda = 2$-year flow sequences; that is, rather than resampling the "historic" time-series record in 1-year blocks, resample it in 2-year blocks so as to reproduce half of the year-to-year serial correlation of the flows. Note that for $\lambda = 2$, equation (1) yields $\rho_{\text{effective}} = \rho/2$. This approach was taken, and the experiments which led to Figures 3 and 4 were repeated and summarized in Figures 5 and 6 for $n = N = 40$ and 80, respectively. Figures 5 and 6 reveal that the moving-blocks bootstrap always led to lower rmse associated with $\hat{S}_p$ than the parametric procedure. The overall reductions in rmse result from having bootstrapped 2-year blocks, which reduced the bias in $\hat{S}_p$. Yet still, Figures 5 and 6 reveal that the moving-blocks bootstrap leads to significant downward bias in storage capacity estimates and that bias does not disappear for larger samples. This significant downward bias is to be expected because the effective correlation of the bootstrap sequences is only 0.15 instead of 0.3 as it should be. Increasing the length $\lambda > 2$ of the moving blocks would lead to still lower bias in $\hat{S}_p$. No methods are currently available for choosing an effective bootstrap block length [*Efron and Tibshirani*, 1993]; future research should address this subject.

Discussion of Results

Our results suggest that resampling from the empirical distribution function of lognormal data is more efficient than generating data from a lognormal model fit using method of moments. *Stedinger* [1980] documents that for the cases considered here; $C_v = 0.4$, $n = 40$ and 80, MLE's are slightly more efficient than method-of-moment estimators, hence our results are due in part to our use of method-of-moments estimators instead of MLE's. On the one hand, we give an advantage to the NSH by not allowing the PSH to use the most efficient estimators. On the other hand, we give an advantage to the PSH by choosing the correct model structure because, in practice, the PSH has added uncertainty resulting from errors in model choice.

The increased uncertainty resulting from parameter estimation is even more complex than described above. Suppose the problem were simpler and the PSH draws samples from a normal distribution with estimated mean $\bar{x}$ and variance s^2. The PSH generates samples using $x = \bar{x} + sz$, where z is a standard normal random variable. In this case, *Stedinger* [1983, equation (33)] shows that $\text{Var}(x) = \sigma^2(1 + 1/n)$, so the PSH generates flows with an inflated variance, across all samples. Note that the inflation in the variance of the flows generated by the PSH is likely to be even worse for lognormal flows and for serially correlated flows. The NSH always draws samples from a normal distribution with the correct mean μ and variance σ^2. The inflation in the variance of the generated flows associated with the PSH leads to slightly larger rmse ($\hat{S}_p$) than for the NSH who generates flows with the proper variance.

Other Promising Bootstraps

The moving-blocks bootstrap is only one approach for preserving the serial correlation of the original time series. *Politis and Romano* [1992] suggest using a circular moving-blocks

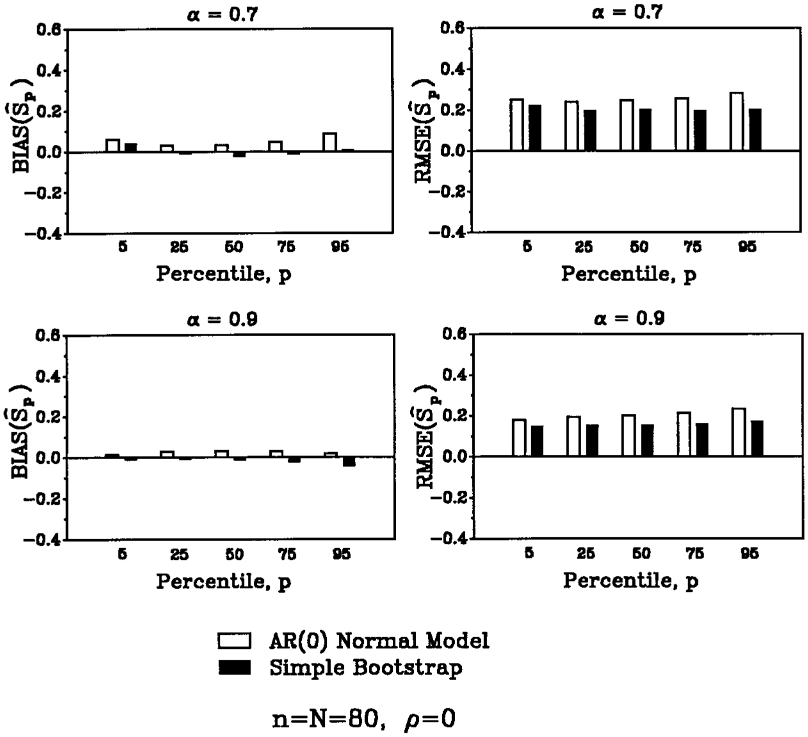

Figure 2. Comparison of bias and rmse associated with $\hat{S}_p$ using an AR(0) normal model and a simple bootstrap with $\lambda = 1$, when $\rho = 0$ and $n = N = 80$.

bootstrap to preserve the sample mean of the original time series. *Efron and Tibshirani* [1992] and *Hjorth* [1994] suggest the use of residual resampling schemes, where one first assumes a model structure, estimates model parameters and model residuals. One then bootstraps the estimated model residuals and, finally, the bootstrapped residuals are used to synthesize a time series. This approach preserves the empirical density function of the original time series; however, in our

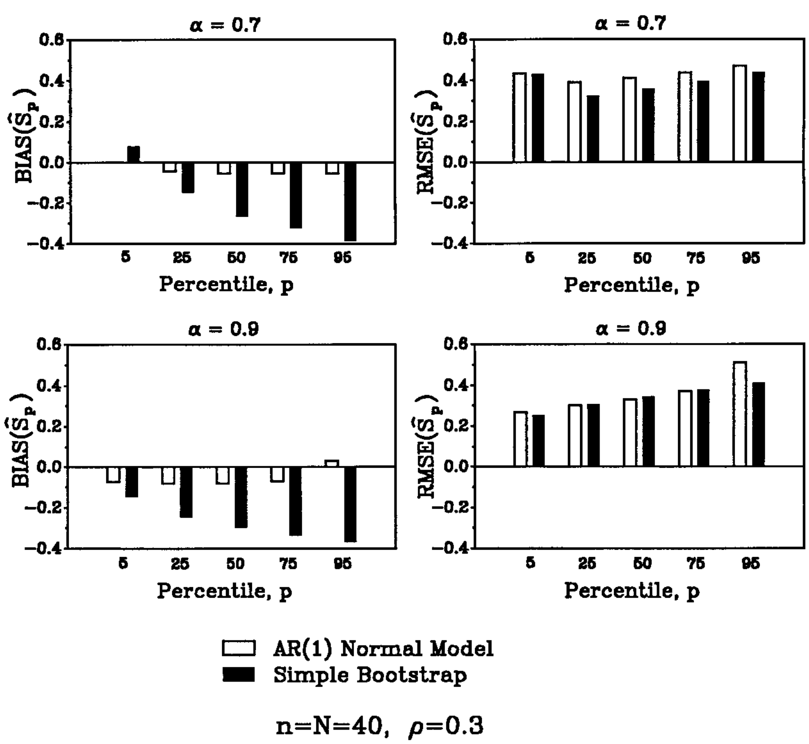

Figure 3. Comparison of bias and rmse associated with $\hat{S}_p$ using an AR(1) normal model and a simple bootstrap with $\lambda = 1$, when $\rho = 0.3$ and $n = N = 40$.

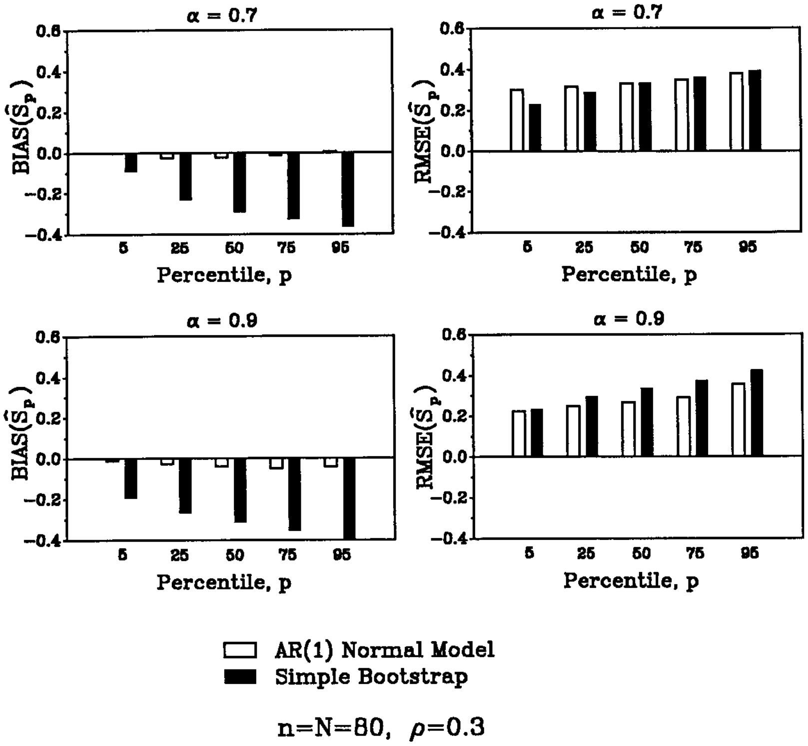

Figure 4. Comparison of bias and rmse associated with $\hat{S}_p$ using an AR(1) normal model and a simple bootstrap with $\lambda = 1$, when $\rho = 0.3$ and $n = N = 80$.

opinion it has little other advantage. *Oliveira et al.* [1988] discuss the use of a nonparametric model of residuals for generating monthly streamflows at many sites in a region. *Hjorth* [1994] also describes a spectral bootstrap in which the temporal correlation of the observations is transformed into independent spectral increments. *Efron and Tibshirani* [1993] also provide citations to numerous other time series applications of the bootstrap.

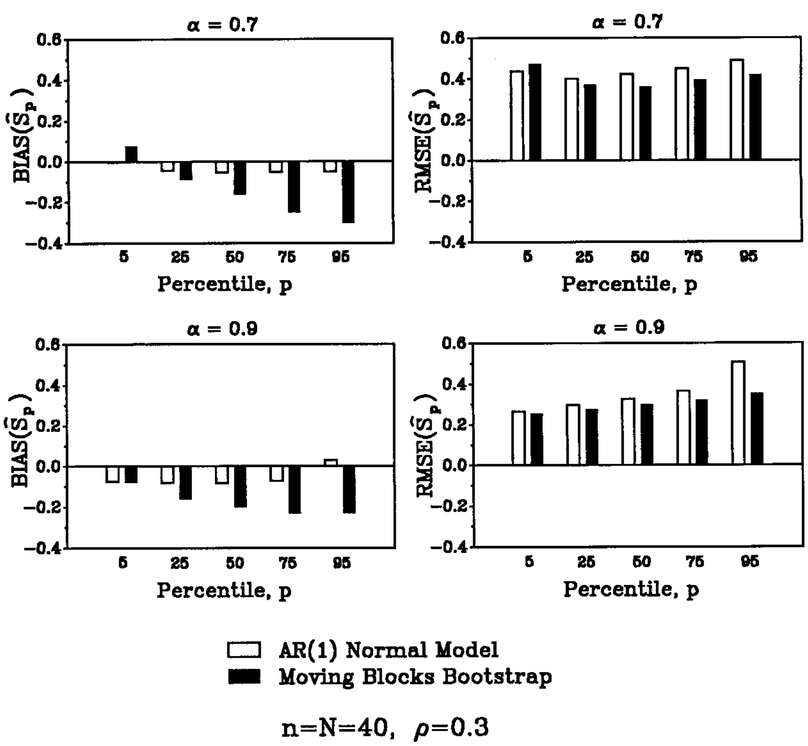

Figure 5. Comparison of bias and rmse associated with $\hat{S}_p$ using an AR(1) normal model and a moving-blocks bootstrap with $\lambda = 2$, when $\rho = 0.3$ and $n = N = 40$.

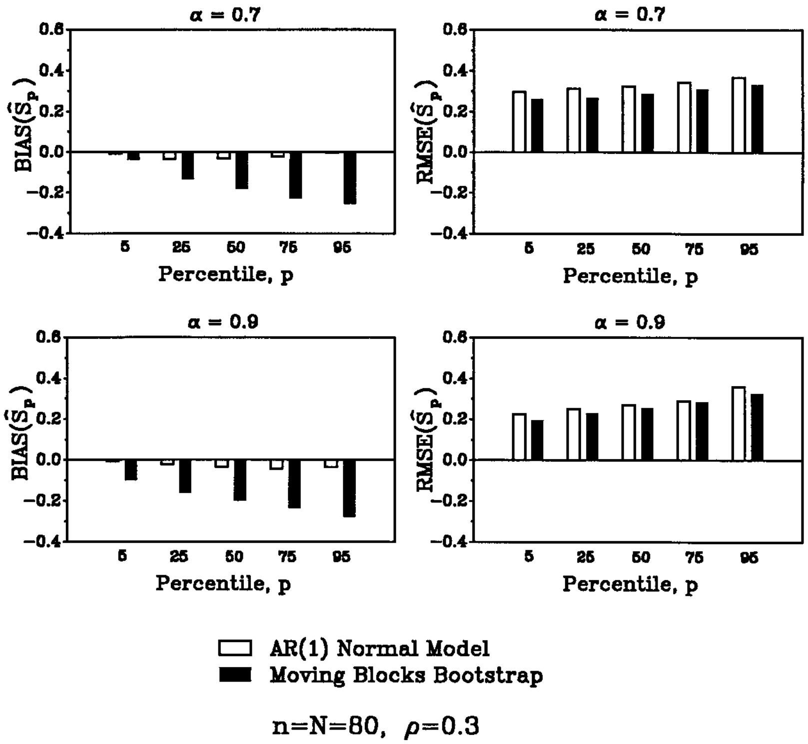

Figure 6. Comparison of bias and rmse associated with $\hat{S}_p$ using an AR(1) normal model and a moving-blocks bootstrap with $\lambda = 2$, when $\rho = 0.3$ and $n = N = 80$.

Conclusions

This study has shown that the moving-blocks bootstrap, which simply resamples the observed time series in approximately independent blocks, can provide a very simple and attractive alternative to parametric time series models. The moving-blocks bootstrap provides "synthetic" time series which preserve the empirical probability distribution of the original observations. If the block length is chosen correctly, the moving-blocks bootstrap can also approximate the empirical covariance structure of the original observations. A Monte Carlo experiment revealed that even though the bootstrap does not introduce any new "synthetic" observations, still, it always produced estimates of the storage capacity with lower root-mean-square-error than estimates based on the correct parametric time series model. The advantage of the bootstrap over the parametric model is due primarily to the fact that the parametric model requires estimates of model parameters which contain uncertainty because of the short records upon which they are based. The advantage of the bootstrap over parametric models would be even greater in practice because of the added uncertainty resulting from the fact that one never knows the correct parametric model to apply. The moving-blocks bootstrap does not require one to select a model and the only parameter required is the block length λ.

This study compares the use of a bootstrap with parametric alternatives for generating univariate (annual) time series. The use of the bootstrap in time series analysis is receiving considerable attention in the statistics literature, as documented by *LePage and Billard* [1992], *Efron and Tibshirani*, [1993], and *Hjorth* [1994]. The moving-blocks bootstrap concept can be easily extended to multivariate time series in both space and time similar to the way in which disaggregation models are applied. In fact, it is the extension of the moving-blocks bootstrap to multivariate time series analysis which motivated this study. We concentrated on the bootstrap of univariate time series first, to set the stage for future work on multivariate time series. Hopefully, future research will address the potential advantages of the bootstrap in a multivariate framework.

Parametric multivariate time series models can be extremely complex. Sometimes the number of parameters in a multivariate disaggregation model is larger than the number of streamflows available for parameter estimation. For example, *Grygier and Stedinger* [1988] discuss a 10-site Valencia-Schaake disaggregation model which has 8580 parameters and requires 72 years of record to have as many data points as there are model parameters. In that example, Grygier and Stedinger point out that 130 years of data would be necessary to have sufficient information to uniquely define some of the model parameters. Clearly, the correct implementation of models of that scale of complexity requires considerable time, money, computational ability, and theoretical knowledge. The moving-blocks bootstrap offers a dramatically simpler approach with potential advantages due to the lack of a required model structure and associated parameter estimation schemes.

Acknowledgments. The authors appreciate the valuable reviews provided by Jery R. Stedinger, Upmanu Lall, and one anonymous reviewer whose suggestions led to many important improvements in our manuscript.

References

Efron, B., and R. J. Tibshirani, *An Introduction to the Bootstrap*, Chapman and Hall, New York, 1993.

Fiering, M. B, *Streamflow Synthesis*, Harvard Univ. Press, Cambridge, Mass., 1967.

Grygier, J. C., and J. R. Stedinger, Condensed disaggregation procedures and conservation corrections for stochastic hydrology, *Water Resour. Res.*, *24*(10), 1574–1584, 1988.

Hausman, E. D., *Analysis of Water Supply and Demand in Eastern Massachusetts*, Masters thesis, Tufts Univ., Medford, Mass., 1990.

Helsel, D. R., and R. M. Hirsch, *Statistical Methods in Water Resources*, Elsevier, New York, 1992.

Hjorth, J. S. U., *Computer Intensive Statistical Methods—Validation Model Selection and Bootstrap*, Chapman and Hall, New York, 1994.

Kendall, D. R., and J. A. Dracup, A comparison of index-sequential and AR(1) generated hydrologic sequences, *J. Hydrol.*, *122*, 335–352, 1991.

Künsch, H. R., The jacknife and the bootstrap for general stationary observations, *Ann. Stat.*, *17*(3), 1217–1241, 1989.

Labadie, J. W., D. G. Fontane, G. Q. Tabios III, and N. F. Chou, Stochastic analysis of dependable hydropower capacity, *J. Water Resour. Plann. Manage. ASCE*, *113*(3), 422–437, 1987.

Lall, U., Recent advances in nonparametric function estimation: Hydrologic applications, *US Natl. Rep. Int. Union Geod. Geophys. 1991–94, Rev. Geophys.*, *33*, 1093–1102, 1995.

Lall, U., and A. Sharma, A nearest neighbor bootstrap for resampling hydrologic time series, *Water Resour. Res.*, *32*, 679, 1996.

Lane, W. L., *Applied Stochastic Techniquies (LAST Computer Package)*: *User Manual*, Div. of Plann. Tech. Serv., U.S. Bur. of Reclam., Denver, Colo., 1979.

LePage, R., and L. Billard, *Exploring the Limits of Bootstrap*, 525 pp., John Wiley, New York, 1992.

Oliveira, G. C., J. Kelman, M. V. F. Pereira, and J. R. Stedinger, A representation of spatial cross correlations in large stochastic seasonal streamflow models, *Water Resour. Res.*, *24*(5), 781–785, 1988.

Phatarfod, R. M., The effect of serial correlation on reservoir size, *Water Resour. Res.*, *22*, 927–934, 1986.

Politis, D. N., and J. P. Romano, A circular block-resampling procedure for stationary data, in *Exploring the Limits of Bootstrap*, edited by Raoul LePage and Lynne Billard, 525 pp., John Wiley, New York, 1992.

Rippl, W., The capacity of storage reservoirs for water supply, *Proc. Inst. Civ. Eng.*, *61*, 270–278, 1883.

Salas, J. D., Analysis and modeling of hydrologic time series, chap. 19, in *Handbook of Hydrology*, edited by D. R. Maidment, 1993.

Salas, J. D., J. W. Delleur, V. Yevjevich, and W. L. Lane, *Applied Modelling of Hydrologic Series*, Water Resour. Pub., Fort Collins, Colo., 1980.

Stedinger, J. R., Fitting log normal distributions to hydrologic data, *Water Resour. Res.*, *16*(3), 481–490, 1980.

Stedinger, J. R., Design events with specified flood risk, *Water Resour. Res.*, *19*(2), 511–522, 1983.

Stedinger, J. R., and M. R. Taylor, Synthetic streamflow generation, 1, Model verification and validation, *Water Resour. Res.*, *18*(4), 909–918, 1982a.

Stedinger, J. R., and M. R. Taylor, Synthetic streamflow generation, 2, Effect of parameter uncertainty, *Water Resour. Res.*, *18*(4), 919–924, 1982b.

Stedinger, J. R., and R. M. Vogel, Disaggregation procedures for generating serially correlated flow vectors, *Water Resour. Res.*, *20*(1), 47–56, 1984.

Stedinger, J. R., B. F. Sule, and D. Pei, Multiple reservoir system screening models, *Water Resour. Res.*, *19*(6), 1383–1393, 1983.

Sudler, C. E., Storage required for the regulation of streamflow, *Trans. Am. Soc. Civ. Eng.*, *91*, 622–660, 1927.

Vogel, R. M., and J. R. Stedinger, Generalized storage-reliability-yield relationships, *J. Hydrol.*, *89*, 303–327, 1987.

Vogel, R. M., and J. R. Stedinger, The value of stochastic streamflow models in overyear reservoir design applications, *Water Resour. Res.*, *24*(9), 1483–1490, 1988.

Willemain, T. R., Bootstrap on a shoestring: Resampling using spreadsheets, *Am. Stat.*, *48*(1), 40–42, 1994.

Zucchini, W., and P. T. Adamson, On the application of the bootstrap to assess the risk of deficient annual inflows to a reservoir, *Water Resour. Manage.*, *2*, 245–254, 1988.

A. Shallcross, Camp Dresser & McKee Inc., Raritan Plaza I, Raritan Center, Edison, NJ 08818. (email: shallcross@cdm.com)

R. Vogel, Department of Civil and Environmental Engineering, Tufts University, Medford, MA 02155. (e-mail: rvogel@tufts.edu)

(Received June 30, 1995; revised December 27, 1995; accepted March 22, 1996.)