

The bootstrap approach to autoregressive time series analysis

by

Felicia Henriette de Koster

Submitted in fulfilment of part of the
requirements for the degree of

Master of Mathematical Statistics

in the Faculty of Science

University of Pretoria

Pretoria

October 1999

The financial assistance of the Centre for Science Development (HSRC South Africa) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the Centre for Science Development.

Acknowledgements

I would like to express my appreciation towards my supervisor Dr H Boraine. Her suggestions, contributions and enthusiasm were of great value throughout my research.

I am grateful to my father for his assistance with the computer application.

This work is completed with immense gratitude to loved ones. I thank my parents, family and dear friends for their continuous support and encouragement.

To the Creator for the blessed ability to study and advance in knowledge.

Summary

The bootstrap approach to autoregressive
time series analysis

by

Felicia Henriette de Koster

Supervisor: Dr H Borraine

Department of Statistics

Submitted in fulfilment of part of the requirements
for the degree of Master of Mathematical Statistics

The bootstrap is a non-parametric computer-intensive statistical technique that uses a unique finite sample to describe the variability of a statistic without making any distributional assumptions about the data. The bootstrap is especially useful in situations where finite sample theory is difficult or even impossible to derive or when only asymptotic theory is available. The latter being the case with time series data.

Formally, the bootstrap consists of a methodology for estimating standard errors by repeatedly resampling with replacement from the original finite sample, which is believed to be a sample of independent and identically distributed (i.i.d.) observations from an unknown probability distribution. Until recently however, it

was not possible to use the bootstrap in studies of time series data. The reason lies in the assumption of i.i.d. random variables which is violated when observations are serially correlated. Different approaches to this problem of preserving the correlation in the data are considered.

A brief introduction to the standard bootstrap principle is provided and the different bootstrap approaches to the estimation of autoregressive time series parameters are discussed. The resampling of residuals, the moving blocks bootstrap and the stationary bootstrap methodologies are all examined.

A comparison is made of the different bootstrap methods in terms of the sampling distributions of the parameters. The standard errors obtained by the methods are also examined. Finally, these methods are applied to the construction of prediction intervals.

The results of a simulation study are included in order to compare the different bootstrap approaches with the conventional estimation methods for a specific second order autoregressive process.

A discussion of the SAS program used for the bootstrap computations of the preferred method is given and the use of the application included on the diskette is also explained.

Samevatting

Die bootstrap benadering tot outoregressiewe
tydreeksanalise
deur

Felicia Henriette de Koster

Studieleier: Dr H Borraine

Departement Statistiek

Voorgelê ter vervulling van 'n deel van die vereistes
vir die graad Magister in Wiskundige Statistiek

Die bootstrap is 'n nie-parametriese rekenaar intensiewe statistiese tegniek wat gebruik maak van 'n unieke eindige steekproef om die variasie van 'n grootheid te beskryf sonder om enige aannames te maak oor die verdeling van die data. Die bootstrap is veral nuttig wanneer dit moeilik of selfs onmoontlik is om eindige steekproef teorie af te lei of as slegs asimptotiese teorie beskikbaar is, soos wat die geval is met tydreeksdata.

Formeel kan die bootstrap beskryf word as 'n metode om standaardfoute te beraam deur herhaaldelike steekproefneming met teruglegging uit die oorspronklike eindige steekproef. Daar word aanvaar dat die steekproef bestaan uit onafhanklike en identies verdeelde waarnemings vanuit 'n onbekende

waarskynlikheidsverdeling. Tot onlangs was dit egter nie moontlik om die bootstrap toe te pas op studies van tydreeksdata nie. Die rede hiervoor lê in die aanname van onafhanklike en identies verdeelde stogastiese veranderlikes wat nie meer geld wanneer opeenvolgende waarnemings gekorreleerd is nie. Verskillende benaderings tot hierdie probleem om die korrelasie in die data te behou word oorweeg.

'n Kort inleiding tot die basiese bootstrap beginsel word gegee en die verskillende bootstrap benaderings tot die beraming van outoregressiewe tydreeks parameters word bespreek. Die hertrekking van residue, die bewegende blokke bootstrap en die stationêre bootstrap metodiek word beskou

'n Vergelyking van die verskillende bootstrap metodes word getref in terme van die steekproefverdelings van die parameters. Die standaardfoute verkry deur die metodes word ook beskou. Uiteindelik word hierdie metodes ook toegepas om voorspellingsintervalle te bereken.

Die resultate van 'n simulasiestudie word ingesluit sodat die verskillende bootstrap benaderings vergelyk kan word met die tradisionele beramingsmetodes vir 'n tweede orde outoregressiewe proses.

Die SAS program waarmee die bootstrap berekenings uitgevoer is, word bespreek en die gebruik van die toepassing wat ingesluit is op die disket, word verduidelik.

Notation

The following provides a list consisting of the notation adopted.

Symbol	Commentary
$\{Z_t, t = 1, \dots, n\}$	Observed time series
Z_t	Time series observation at time t
$\{Z_t^*\}$	Bootstrap time series
n	Length of time series
p	Order of time series
$\phi_i, i = 1, \dots, p$	Time series parameters
$\Phi = (\phi_1, \dots, \phi_p)$	Vector of parameters
$\hat{\Phi}$	Estimate of Φ based on the original series
$\hat{\Phi}_j^*$	Estimate of Φ based on the j^{th} bootstrap series
B	Number of bootstrap replications
*	Star notation indicates a bootstrap replication
l	Block length
m	Number of blocks resampled
n^*	Length of pseudo-time series
a_t	Residual at time t
\hat{a}_t	Estimated residual at time t
a_t^*	Resampled residual at time t
\bar{a}	Mean of the residuals
$F(\cdot)$	Distribution of the error terms
$\hat{F}(\cdot)$	Estimate of $F(\cdot)$
σ_a^2	Error term variance
Z_{n+h}	Future value of Z at time $n+h$
$\hat{Z}_n(h)$	Predicted value of Z at time $n+h$
$\{X_i\}$	Set of random i.i.d. variables
$[x]$	Integer part of x
$j(\text{mod } N)$	Integer remainder of j/N

Table of Contents

Chapter 1	Introduction	1
Chapter 2	The Standard Bootstrap Principle	3
Chapter 3	The Bootstrap Approach to Parameter Estimation	8
Chapter 4	Resampling the Residuals	13
	- Simulation Results	19
Chapter 5	The Moving Blocks Bootstrap	23
	- Choosing the Optimal Block Length	26
	- Simulation Results	30
Chapter 6	The Stationary Bootstrap	33
	- Simulation Results	39
Chapter 7	Comparing the Methods	41
	- Comparison of the Sampling Distributions	41
	- Comparison of the Standard Errors	48

Chapter 8	Bootstrap Prediction Intervals	55
	- An Improved Estimator	62
	- Simulation Results	64
Chapter 9	Discussion of Bootstrap Computations	70
Chapter 10	Conclusion	76
Appendix A	The Infinite Order Moving Average Representation of an Autoregressive Model	79
References		80

Chapter 1

Introduction

The bootstrap is a non-parametric computer-intensive statistical technique that was first introduced by Efron in 1979. This technique uses a unique finite sample to describe the variability of a statistic without making any distributional assumptions about the data. The bootstrap is especially useful in situations where finite sample theory is difficult or even impossible to derive or when only asymptotic theory is available. The latter being the case with time series data.

Formally, the bootstrap consists of a methodology for estimating standard errors by repeatedly resampling with replacement from the original finite sample, which is believed to be a sample of independent and identically distributed (i.i.d.) observations from an unknown probability distribution. The resamples obtained in this manner are called *pseudo-data* or *bootstrap samples* and are used to estimate the statistics of interest. These estimates are then referred to as the *bootstrap estimates*.

Until recently however, it was not possible to use the bootstrap in studies of time series data. The reason lies in the assumption of i.i.d. random variables which is violated when observations are serially correlated. Different approaches to this problem have

been suggested and a few will be considered in the following chapters.

Chapter 2 provides a brief introduction to the standard bootstrap principle, while chapter 3 discusses the different bootstrap approaches to parameter estimation in the context of autoregressive time series. The resampling of residuals, the moving blocks bootstrap and the stationary bootstrap methodologies are discussed in chapters 4, 5 and 6. Each one of these chapters points out relevant factors to consider for the different procedures. Chapter 7 compares the different bootstrap methods in terms of the sampling distributions of the parameters. The standard errors obtained by the methods are also examined. Finally in chapter 8, these methods are applied to the construction of prediction intervals. The results of a simulation study are included in order to compare the different bootstrap approaches with the conventional estimation methods for a specific autoregressive process of order 2. A discussion of the SAS program used for the bootstrap computations of the preferred method, is given in Chapter 9. This chapter also explains the use of the application included on the diskette. The last chapter offers conclusions and recommendations for further study.



Chapter 2

The Standard Bootstrap Principle

Let $\mathbf{X}=(X_1, \dots, X_n)$ represent the original finite sample of n independent identically distributed (i.i.d.) observations, obtained from a common unknown probability distribution $F(\cdot)$ and let $T_n(\mathbf{X})$ be some statistic of interest which is under consideration.

Denote the empirical distribution by $\hat{F}(\cdot)$ which assigns probability mass n^{-1} to each sample element. The bootstrap is used to approximate the sampling distribution of $T_n(\mathbf{X})$ under $F(\cdot)$ by the bootstrap distribution of $T_n(\mathbf{X}^*)$ under $\hat{F}(\cdot)$, where $\mathbf{X}^*=(X_1^*, \dots, X_n^*)$ ¹ is a bootstrap sample of size n obtained by randomly sampling *with replacement* from the original sample \mathbf{X} . This approximation can be done according to the Monte Carlo method by repetitively resampling from the original data and recalculating the value of the statistic. The desired accuracy can be obtained by increasing the number of repetitions sufficiently.


The bootstrap algorithm begins by generating a large number, B , of independent bootstrap samples denoted by \mathbf{X}_i^* , $i=1, \dots, B$, each of size n . These samples are in fact, drawn from the empirical

¹ The star notation indicates that \mathbf{X}^* is not the original data set \mathbf{X} , but rather a resampled version of \mathbf{X}

distribution $\hat{F}(\cdot)$. Corresponding to each bootstrap sample \mathbf{X}_i^* is a bootstrap replication of T_n , namely $T_n(\mathbf{X}_i^*)$, the value of the statistic evaluated for \mathbf{X}_i^* . This bootstrap resampling scheme is illustrated in Figure 2.1.

Figure 2.1

Schematic representation of the bootstrap process for estimating the sampling distribution of a statistic $T_n(\mathbf{X})$.



The set of bootstrap estimates $\{T_n(\mathbf{X}_i^*), i = 1, \dots, B\}$ forms an approximation to the true sampling distribution of the statistic



$T_n(\mathbf{X})$. In particular, estimates of location, spread, confidence intervals and bias of the statistic can be evaluated from the set $\{T_n(\mathbf{X}_i^*)\}$.

The bootstrap estimate of the parameter, denoted by T_n^* , is defined by the mean of the bootstrap replications

$$T_n^* = \frac{1}{B} \sum_{i=1}^B T_n(\mathbf{X}_i^*) \quad (2.1)$$

and the variance is estimated by the empirical variance of the replications which is denoted by $V(T_n^*)$ and defined as

$$V(T_n^*) = \frac{1}{(B-1)} \sum_{i=1}^B [T_n(\mathbf{X}_i^*) - T_n^*]^2 \quad (2.2)$$

To illustrate the practical use of the bootstrap for estimating standard errors and probabilities, the following two examples are included. The second example is given with acknowledgement to Swanepoel (1990).

Example 2.1

Consider the estimation of the standard error of the sample mean

$$T_n(\mathbf{X}) = \bar{X} = n^{-1} \sum_{i=1}^n X_i.$$



The observed sample is $X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$ and $F(\cdot)$ is estimated by the empirical distribution $\hat{F}(\cdot)$,

$$\hat{F} : \text{Probability} = n^{-1} \text{ at each } x_i, i = 1, 2, \dots, n.$$

A bootstrap random sample $\mathbf{X}^* = (X_1^*, X_2^*, \dots, X_n^*)$ is drawn from $\hat{F}(\cdot)$. This means that each X_i^* is drawn independently with replacement and with equal probability n^{-1} from the observed set $\{x_1, x_2, \dots, x_n\}$.

Calculate the sample mean for this bootstrap sample,

$$T_n(\mathbf{X}^*) = \bar{X}^* = n^{-1} \sum_{i=1}^n X_i^*$$

The bootstrap sampling and estimate calculation as explained above is repeated a large number, B , times to obtain the bootstrap replications $T_n(\mathbf{X}_1^*) = \bar{X}_1^*, \dots, T_n(\mathbf{X}_B^*) = \bar{X}_B^*$.

Finally, the standard error of the sample mean is estimated by the sample standard deviation of the B bootstrap replications,

$$s = (B - 1)^{-1} \sum_{i=1}^B [T_n(\mathbf{X}_i^*) - \bar{T}_n]^2$$

where $\bar{T}_n = B^{-1} \sum_{i=1}^B T_n(\mathbf{X}_i^*)$ is the sample mean of the B bootstrap replications.

Example 2.2

Consider the estimation of the probability

$$p = P(\sqrt{n}(\bar{X} - \mu) \leq 0.5)$$

where $\mu = E(X_i)$ and $\bar{X} = n^{-1} \sum_{i=1}^n X_i$.

Define the statistic under consideration as $T_n(\mathbf{X}) = \sqrt{n}(\bar{X} - \mu)$.

Once again, the observed sample is

$$X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$$

and $F(\cdot)$ is estimated by the empirical distribution $\hat{F}(\cdot)$,

$$\hat{F} : \text{Probability} = n^{-1} \text{ at each } x_i, i = 1, 2, \dots, n.$$

Draw a bootstrap random sample $\mathbf{X}^* = (X_1^*, X_2^*, \dots, X_n^*)$

independently from $\hat{F}(\cdot)$ as described in Example 2.1.

Calculate the relevant statistic for this bootstrap sample by

$$T_n(\mathbf{X}^*) = \sqrt{n}(\bar{X}^* - \bar{X})$$

where $\bar{X}^* = n^{-1} \sum_{i=1}^n X_i^*$

Repeat the above two steps a large number, B , times to obtain the

bootstrap replications $T_n(\mathbf{X}_1^*), \dots, T_n(\mathbf{X}_B^*)$.

Finally, calculate

$$\hat{p} = B^{-1} \sum_{i=1}^B I(T_n(\mathbf{X}_i^*) \leq 0.5)$$

where $I(\cdot)$ is the indicator function.



Chapter 3

The Bootstrap Approach to Parameter Estimation

Efron (1979) introduced the bootstrap procedure for estimating sampling distributions based on a finite sample of i.i.d. observations. After two decades it is already well known that, in the i.i.d. setup, the bootstrap often provides more accurate approximations than classical large sample approximations (Singh 1981; Babu 1986). However, Singh (1981) showed that the original bootstrap no longer succeeds when the observations are not necessarily independent. Originally this prevented the use of the bootstrap in studies of time series data, because the assumption of i.i.d. random variables is violated when observations are serially correlated. Several solutions to this problem have now been suggested, although most extensions in the literature so far only apply to the stationary case. They can roughly be divided into resampling and subsampling methods.

There are broadly speaking, two approaches to the use of resampling methods for strictly stationary dependent data. The first is to apply Efron's bootstrap to an approximate i.i.d. setting by focusing on the residuals of some general regression model. Such examples include linear regression, autoregressive time series




as well as other applications. In these situations the residuals are resampled instead of the original observations. In addition to being restricted to relatively simple contexts where structural models are both plausible and tractable, little is known about how this approach would perform for heteroskedastic observations. The fitted residuals will, in general, no longer behave like i.i.d. observations, but exhibit some form of heteroskedasticity. Fortunately it has been shown that Efron's bootstrap performs reasonably well even when the data are independent but not identically distributed. (Freedman 1981; Liu 1988; Liu and Singh 1992). Therefore, one might hope for a certain degree of robustness to heteroskedasticity as well.

As a second approach resampling methods for less restrictive contexts have more recently been suggested. They are based on *blocking* techniques in which the data are divided into blocks of adjacent observations and these blocks are resampled rather than the single original observations or the estimated residuals. This is done in order to capture the dependence in consecutive observations. Carlstein (1986) proposed non-overlapping blocks, whereas Künsch (1989) and Liu and Singh (1992) independently introduced the *moving blocks* method which utilises overlapping blocks. These overlapping blocks use the data more efficiently and consequently reduce variability. Figure 3.1 compares these two blocking methods schematically.


Figure 3.1

Schematic comparison of Carlstein's non-overlapping blocks and Künsch's moving blocks. The black circles represent the original observations and are divided into blocks of length 3.

Carlstein's non-overlapping blocks:



Künsch's moving blocks:



Hall et al (1995) analyze the asymptotic behaviour for estimates based on a scalar sample mean within a stationary environment and find that the moving blocks bootstrap enjoys an asymptotic superiority relative to Carlstein's method. Carlstein's method is therefore not discussed further in this study as the asymptotic quality is considered relevant for typical sample sizes found in economic situations. Furthermore, Carlstein's rule also neglects dependency between the different blocks of observations in the sample, which in particular might be of great importance for a nonstationary process.



Politis and Romano (1992) also consider a *blocks of blocks* scheme to obtain valid inference of parameters of the infinite-dimensional joint distribution of the process, such as the spectrum. It turns out that Künsch's bootstrap enjoys some robustness property to heteroskedasticity as was pointed out by Lahiri (1992) in the case of the sample mean.

In both Carlstein's and Künsch's bootstrap blocks of fixed length are resampled which means that the newly generated pseudo-time series is no longer stationary. To fix this shortcoming, Politis and Romano (1994a) suggested the stationary bootstrap which is based on resampling blocks of random length, where the length of each block has a geometric distribution.

Building new pseudo-time series by joining randomly selected independent blocks together induces a different probability mechanism. Dependency will be reduced and for both Carlstein's and Künsch's bootstrap stationarity will be lost. However in typical applications the underlying dependence is sufficiently weak. Therefore the main contributions come from the short lags which are well approximated by the blocking methods when an appropriate block size is chosen. This ensures that these methods work nevertheless.

As an alternative to resampling methods, Politis and Romano (1994b) proposed the subsampling approach. Rather than



resampling blocks from the original series in order to construct a new pseudo-time series, each individual subblock or subseries of observations is viewed as a valid *sub-time series* in its own right. The motivation is that each block as a part of the original series, was generated according to the true underlying probability mechanism. It then seems reasonable to hope that one can gain information about the sampling distribution of a statistic by evaluating it on all of these subseries.

Another attractive feature of the subsampling method is that it has been shown to be valid under very weak assumptions. Apart from regularity and dependency conditions the only requirement in the stationary setup is that the sampling distribution of the properly normalized statistic of interest has a nondegenerate limiting distribution. Politis, Romano and Wolf (1997) consider subsampling for heteroskedastic time series. The subject of heteroskedasticity is beyond the scope of this study, therefore their methods will not be considered in greater depth.

In the following chapters some of these proposed methods are explored in the context of autoregressive time series. Results of a simulation study is also given in an attempt to measure the performance of these methods and to compare them with the traditional Box-Jenkins methodology (1994).

Chapter 4

Resampling the Residuals

Let $\{Z_t, t = 1, \dots, n\}$ represent n observations of a stationary autoregressive time series of order p . The AR(p) structure proposed by Box, Jenkins and Reinsel (1994) to model Z_t is given by

$$Z_t = \phi_1 Z_{t-1} + \phi_2 Z_{t-2} + \dots + \phi_p Z_{t-p} + a_t \quad (4.1)$$

where the a_t are the i.i.d. residuals from an unknown distribution $F(\cdot)$ with mean zero and finite variance σ_a^2 . For the purpose of explaining the procedures, assume that the mean of the series has been subtracted from each observation in order that Z_t has a zero mean. Furthermore, for stationarity all p roots of the characteristic equation $1 - \phi_1 x - \phi_2 x^2 - \dots - \phi_p x^p = 0$ lie outside of the unit circle.

The vector of unknown parameters $\Phi = (\phi_1, \dots, \phi_p)$ is obtained by the minimisation of the unconditional sum of squares of the residuals

$$S(\Phi) = \sum_{t=1}^n [a_t | \Phi, \mathbf{Z}]^2 \quad (4.2)$$

where $\mathbf{Z} = (Z_1, \dots, Z_t)$

Due to the non-linear nature of $S(\Phi)$ the minimisation is carried out by an iterative process using a non-linear algorithm such as that proposed by Marquardt (1963). It is of interest to recall that if the residuals, a_t , in (4.1) are assumed to possess a Gaussian distribution, then the least squares estimates obtained by the minimisation of $S(\Phi)$ in equation (4.2) are equivalent to the maximum likelihood estimates for large samples (Box, Jenkins and Reinsel, 1994). The above equivalence is used in the evaluation of the variance of the estimates. It should be kept in mind that the variance estimator is an approximation of the true variance as it only holds asymptotically, i.e. it is only valid for large samples.

Given the stationary time series $\{Z_t\}$, the conventional Box and Jenkins approach is used to fit an AR(p) model. The corresponding vector of parameter estimates $\hat{\Phi} = (\hat{\phi}_1, \dots, \hat{\phi}_p)$ is then obtained by the traditional parametric estimation (4.2). This enables the evaluation of the model residuals, \hat{a}_t by

$$\hat{a}_t = Z_t - \hat{\phi}_1 Z_{t-1} - \dots - \hat{\phi}_p Z_{t-p} \quad (4.3)$$

In this way, the required random sample $\{\hat{a}_t, t = 1, \dots, n\}$ of size n is obtained and B bootstrap samples can be generated by sampling with replacement from the empirical distribution of the estimated residuals, $\hat{F}(\cdot)$, as described in Chapter 2. Note that the a_t^* are

independent bootstrap random variables sampled from $\hat{F}(\cdot)$, even though the \hat{a}_t are not independent random variables in the usual sense.

Chatterjee (1986), Holbert and Son (1986) and Efron and Tibshirani (1993) used the procedure described above to resample directly from the estimated residuals. This approach is justified when the mean of $\hat{F}(\cdot)$ is close to zero. If it is not, Efron and Tibshirani (1993) suggested that the definition $E_F(a_t) = 0$ could be honoured by centering the residuals around their mean. This means that the bootstrap residual samples are generated from the empirical distribution that puts mass n^{-1} on $(\hat{a}_t - \bar{a})$ where

$$\bar{a} = \frac{1}{n} \sum_{t=1}^n \hat{a}_t.$$

Shao and Tu (1995) also resampled from the empirical distribution of \hat{a}_t centered at \bar{a} and Freedman (1981) concluded that without centering, the bootstrap will usually fail.

In their simulation study, Souza and Neto (1996) inflated the estimated residuals by the factor $[n/(n-k)]^{1/2}$, where k is the number of free parameters of the model. This is done due to the degrees of freedom lost in the fitting process. McCullough (1994) used the same scaling factor in his study of bootstrapping forecast intervals.

Finally, combining these two ideas, Shao (1996) generated i.i.d. a_1^*, \dots, a_n^* from the empirical distribution putting mass n^{-1} on $(\hat{a}_t - \bar{a})/\sqrt{1 - p/n}$, $t = 1, \dots, n$ where \bar{a} is the average of the \hat{a}_t and p is the order of the model.

Let $\{a_{it}^*, t = 1, \dots, n\}$ represent the i^{th} bootstrap residual sample of size n , generated according to one of the proposed rules, where $i = 1, \dots, B$. The corresponding bootstrap series $\{Z_{it}^*\}$ can now be evaluated by

$$Z_{it}^* = \hat{\phi}_1 Z_{i,t-1}^* + \hat{\phi}_2 Z_{i,t-2}^* + \dots + \hat{\phi}_p Z_{i,t-p}^* + a_{it}^* \quad (4.4)$$

where the model parameters $\hat{\Phi} = (\hat{\phi}_1, \dots, \hat{\phi}_p)$ used in the generation of the pseudo-time series are those estimated according to the original series $\{Z_t\}$. By repeating the above procedure, a set of B bootstrap series each of size n , can be generated and each series can be modelled by the same AR(p) structure adopted for the original series $\{Z_t\}$, i.e.

$$\{Z_{it}^*, i = 1, \dots, B; t = 1, \dots, n\} \sim \text{AR}(p) \quad (4.5)$$

With the bootstrap series generated according to the above description, the traditional parametric estimation procedure can be

applied to each series $\{Z_{it}^*\}$ in order to obtain a set of bootstrap estimates $\{\Phi_j^*, j = 1, \dots, B\}$. Using this sequence of estimates an approximation of the true sampling distribution of each $\phi_i, i = 1, \dots, p$ can be calculated. In particular, bootstrap standard errors and confidence intervals for each parameter ϕ_i of Φ can be constructed.

The above procedure was proposed by Chatterjee (1986), in a pioneer time series simulation study to bootstrap ARMA models. However in the case of the existence of autoregressive terms in the model for the series (i.e. $p \geq 1$) problems arise in the calculation of the bootstrap series $\{Z_{it}^*\}$ in equation (4.4) due to the starting values $Z_{i,0}^*, Z_{i,-1}^*, \dots, Z_{i,1-p}^*$ required in the evaluation of $Z_{i,1}^*$.

One possible solution to this problem is to assume these unknowns fixed for each bootstrap sample. These values can thus be set equal to those used in the original series, i.e.

$$Z_{i,0}^* = Z_0 \quad , \quad Z_{i,-1}^* = Z_{-1} \quad , \quad \dots \quad , \quad Z_{i,1-p}^* = Z_{1-p} \quad (4.6)$$

This approach has been adopted by Efron and Tibshirani (1986) for an AR(1) structure as well as by Chatterjee (1986) for AR(2) and ARMA(1,1) structures.

Souza and Neto (1996) approached this problem differently by generating the starting values $Z_{i,0}^*$ and $Z_{i,-1}^*$ (i.e. for the AR(1), AR(2) and ARMA(1,1) structures) for each bootstrap series. These starting values were generated according to their corresponding Gaussian marginal and conditional probability distributions. They have therefore, in each bootstrap series $\{Z_{it}^*\}$ generated according to equation (4.4), used the corresponding random starting observations instead of considering them fixed. This random selection process was also adopted by Holbert and Son (1986) for the AR(2) case. They found that the bootstrap estimates of the parameters' standard errors improved considerably when this random selection was used.

Masarotto (1990) suggested that in order to obtain stationary data Z_{it}^* could be generated from $t = -k$ where k is a sufficiently large integer e.g. $k = 50$ or 100 . This is also the approach followed in the simulation study.

Simulation Results

In order to test the procedure described in the previous section, a simulation study was conducted. The following AR(2) structure is considered

$$Z_t = 10 + 1.3Z_{t-1} - 0.4Z_{t-2} + a_t \quad (4.7)$$

where the a_t are independent random observations from a normal distribution with zero mean and variance $\sigma_a^2 = 2$. The mean of this series is $\mu = 100$. Another important factor to consider is the influence of the number of observations and therefore two series lengths, $n = 30$ and $n = 100$, were investigated.

In order to simulate the population properties, 1000 Monte Carlo replications were generated. One series was then generated on which the bootstrap methods were performed. The parameters for this series were estimated by means of maximum likelihood estimation and the order was considered fixed as $p = 2$. The bootstrap estimates were computed from $B = 1000$ pseudo-time series, constructed according to (4.4) and generated from $t = -50$.

Two approaches to the resampling of the residuals were followed. Firstly, the centered residuals, $(\hat{a}_t - \bar{a})$ were resampled and

secondly, the resamples were taken from the empirical distribution of $(\hat{a}_t - \bar{a})/\sqrt{1-p/n}$.

The results of these tests are summarised in the Tables 4.1 and 4.2. For each parameter, the estimated value as well as the standard deviation (SD) is given.

Based on these results it seems that the bootstrap standard error estimates compare very well with the maximum likelihood estimates and are in accordance with the simulation results. Only in the case of the series mean, it seems as if all three methods underestimate the standard error. The bootstrap estimates may follow the trend of the maximum likelihood estimate, because the bootstrap series generation is dependent on the values of the maximum likelihood estimates.

Even though the difference is quite small, resampling from $(\hat{a}_t - \bar{a})/\sqrt{1-p/n}$ appears to offer an improvement on the standard errors, compared to resampling from the centered residuals.

One advantage of the bootstrap is that an estimate of the standard error of σ_a^2 can be calculated, which is not the case with the traditional estimation methods.

For a shorter series, where asymptotic validity is often questioned, the bootstrap method also performs quite well. The results obtained in this study show an improvement on the maximum likelihood standard error estimates. Once again it does seem useful to multiply the centered residuals by the factor $(1 - p/n)^{-1/2}$.

Table 4.1

A comparison of the parameter estimates obtained by maximum likelihood estimation and two residual resampling schemes for a series of length $n = 100$

Parameter		Simulation results	Maximum likelihood estimation	Resampling from $(\hat{a}_t - \bar{a})$	Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$
σ_a^2	Estimate	2.00495	2.08226	2.09515	2.13720
	SD	0.28200		0.28722	0.28514
μ	Estimate	100.01815	99.26085	99.25040	99.27405
	SD	1.32986	0.91234	0.96615	0.94508
ϕ_1	Estimate	1.28383	1.30623	1.29090	1.28869
	SD	0.08912	0.08915	0.09009	0.08867
ϕ_2	Estimate	-0.40876	-0.45893	-0.46496	-0.46340
	SD	0.08592	0.08917	0.08632	0.08771



Table 4.2

A comparison of the parameter estimates obtained by maximum likelihood estimation and two residual resampling schemes for a series of length $n = 30$

Parameter		Simulation results	Maximum likelihood estimation	Resampling from $(\hat{a}_t - \bar{a})$	Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$
σ_a^2	Estimate	1.98757	1.85625	1.66753	1.80883
	SD	0.52908		0.41403	0.45126
μ	Estimate	99.81101	101.07308	101.12170	101.08721
	SD	2.34217	1.60515	1.60978	1.61796
ϕ_1	Estimate	1.23587	1.25915	1.19295	1.18531
	SD	0.18291	0.17031	0.18373	0.18394
ϕ_2	Estimate	-0.42543	-0.39604	-0.42401	-0.41350
	SD	0.16729	0.17200	0.16593	0.16441

Chapter 5

The Moving Blocks Bootstrap

The moving blocks bootstrap is presented as an extension of the standard bootstrap principle, to take into account weak dependency in the data. Rather than fitting a model and then sampling from the residuals, this method takes an approach closer to that used for independent data problems. The idea is to generate a bootstrap realization of the time series by choosing a suitable block length and considering all possible contiguous blocks of this length. These blocks are sampled with replacement and pasted together to form the bootstrap time series. Just enough blocks are sampled to obtain a series of roughly the same length as the original series. The standard bootstrap is the special case when the block size equals unity. Figure 5.1 provides a schematic representation of the moving blocks bootstrap applied to a time series.


For a description of the method, consider the moving blocks bootstrap resampling for a statistic of interest. Given a time series $\{Z_t\}$ and the statistic $T_n(\mathbf{Z}) = T_n(Z_1, \dots, Z_n)$, Fitzenberger (1997) defined B_j as the block of l consecutive observations starting from Z_j , that is

$$B_j = (Z_j, \dots, Z_{j+l-1}) \text{ for } j = 1, \dots, n-l+1 \quad (5.1)$$

Figure 5.1

Schematic diagram of the moving blocks bootstrap for time series. The black circles are the original time series observations and the white circles denote a bootstrap realization of the time series.

This figure represents blocks of length 3.



Let $q = n - l + 1$ and $n^* = ml$, with m and n^* positive integers, such that n^* is the smallest integer multiple of l which is greater than or equal to n . Now m blocks Y_1, \dots, Y_m can be resampled independently from $\{B_1, \dots, B_q\}$ with equal probability q^{-1} where each Y_i , $i = 1, \dots, m$ is a block of size l with $Y_i = (y_{i1}, \dots, y_{il})$. The bootstrap distribution of Y_i , conditional on the series $\{Z_t\}$ can be denoted by P^* . Thus, given $\{Z_t\}$ the m random blocks Y_1, \dots, Y_m are i.i.d. according to the conditional distribution P^* . The moving blocks bootstrap resample of size n^* , denoted by $\{Z_1^*, \dots, Z_{n^*}^*\}$ is then formed by joining the blocks Y_1, \dots, Y_m into

one series i.e. $Z_i^* = y_{\tau v}$ for $\tau = \left\lfloor \frac{i-1}{l} \right\rfloor + 1$, where $\lfloor x \rfloor$ denotes the integer part of x and $v = i - l(\tau - 1)$. The resample statistic is then calculated according to $T_{n^*}^*(Z_1^*, \dots, Z_{n^*}^*)$.

Fitzenberger (1997) used the actual resample size as $n^* = ml$ but other authors (Künsch 1989; Liu and Singh 1992) have proposed deleting the observations Z_j^* for $j > n$. If n is not exactly divisible by l , the bootstrap standard errors need to be multiplied by $\sqrt{\frac{ml}{n}}$ to adjust for the difference in lengths of the series.

The justification for the moving blocks bootstrap lies in the correlation structure of the data. Simply resampling from the original observations will destroy this correlation that needs to be captured and modelled. With the moving blocks bootstrap, the idea is to choose a block length l large enough so that the observations more than l time units apart will be nearly independent. By sampling the blocks of length l , the correlation present in observations less than l time units apart, is retained. However, the correct choice of the block length is quite important and requires careful consideration. The next section is devoted to this subject.

The moving blocks bootstrap also has the advantage of being less model dependent than the bootstrapping of residuals approach as the latter method depends on the model that is fit to the original time series. Although the same model structure is used to estimate the parameters for the bootstrap time series as for the original series, it is not used in the generation of the bootstrap realizations of the time series. Note however that in the moving blocks bootstrap sampling, the first and last few observations in the series do not have the same chance of being drawn as the observations in the middle part of the series. Politis and Romano (1992a) and Shao and Yu (1993) proposed a *circular block* method by wrapping the observations Z_1, \dots, Z_n around in a circle and then generating consecutive blocks of the bootstrap data from the circle. This circular method is also used in the *stationary bootstrap* of Politis and Romano (1994) and will be considered again in Chapter 6.

Choosing the Optimal Block Length

The purpose of this section is to address the problem of optimal choice of block length when the block bootstrap is used in a variety of different contexts. According to Hall, Horowitz and Jing (1995), it turns out that optimal block size depends very much on the context. They identify three different settings of practical importance, namely estimation of bias or variance, estimation of a

one-side distribution function, e.g. $F_1(x) = P\left\{\frac{(\hat{\theta} - \theta)}{s} \leq x\right\}$, and

estimation of a two-sided distribution function, e.g.

$F_2(x) = P\left\{\frac{|\hat{\theta} - \theta|}{s} \leq x\right\}$. The two-sided distribution function is

used to construct symmetrical confidence intervals for an unknown parameter and these intervals were shown to enjoy enhanced coverage accuracy among two-sided confidence regions (Hall 1988). Hall, Horowitz and Jing (1995) show that optimal block lengths in the three problems are of different order sizes, being $n^{1/3}$, $n^{1/4}$ and $n^{1/5}$ respectively, where n is the length of the time series. At first this disparity might be confusing but they offer a simple explanation for it. Optimal block length is achieved by balancing error-about-the-mean against bias to minimise mean square error. Bias terms in all three problems are of similar sizes, but variances are quite different for the three cases as they are essentially the variances of standardised second, third and fourth cumulants in the respective problems. Elementary calculations based on this observation lead Hall, Horowitz and Jing (1995) to their claims about orders of block length in the different cases.

The optimal asymptotic formula for block length l is $l \cong Cn^{1/k}$ where $k = 3, 4$ or 5 . As indicated above, the value of k is known, being determined by the specific context. This result is of practical benefit as well as theoretical interest, since it may be used as the

basis for a simple rule for choosing block size. The rule operates by using empirical methods to choose the block size for a sub-series of the original data set of length $m < n$. This quantity may be re-calibrated so that it applies to the original larger sample size, by multiplying by the factor $(n/m)^{1/k}$.

Without considering the detailed theoretical results, the empirical rule for estimating optimal block size for a time series of smaller length than the original, say $m < n$ is first investigated (Hall, Horowitz and Jing, 1995). Let \hat{l}_m denote this block size for a time series of length m . Once \hat{l}_m has been determined, the optimal block size \hat{l}_n for the original time series of length n may be estimated by the formula $\hat{l}_n = (n/m)^{1/k} \hat{l}_m$.

Let ζ denote the set of all $n - m + 1$ runs of length m , obtainable from the original time series and apply the moving blocks bootstrap method to each subseries from ζ . Let l' denote the block size used here. Each application of the bootstrap produces a point estimate of the statistic of interest. Let $T_m(\mathbf{Z}_i^*)$, for $1 \leq i \leq n - m + 1$ denote the bootstrap estimates of T_n computed from the $n - m + 1$ runs of length m in ζ and let $T_n(\mathbf{Z})$ be the estimate computed from the entire data set of length n , using a plausible pilot block size l . An estimate of mean squared error in a sample of size m , using block length l , is the average of the squares

of the differences $T_m(\mathbf{Z}_i^*) - T_n(\mathbf{Z})$. Select that value of l' , denoted by \hat{l}_m , which minimises this quantity and then revise the choice of the pilot block size l , to $(n/m)^{1/k} \hat{l}_m$. This procedure may be iterated if desired, replacing the original pilot choice of l with an updated version.

In their simulation study, Hall, Horowitz and Jing (1995) chose the length of the subseries to be $m = 25$ for a time series of 100 observations. They considered a moving average model for which they found that the number of iterations required for this procedure to converge was one for 93% of their variance estimation simulations as well as for 92% of their distribution estimation simulations. For the remaining cases, convergence was always achieved after two iterations. In the case of the distribution estimation, the theoretical optimum could not be reached. They do however state that larger sample sizes give better performance which suggests that this empirical rule might not work as effectively for short series.

They also studied block bootstrap estimation of the one-sided distribution function for two first order autoregressive models. In both these cases the medians of the distributions of empirically chosen block length were equal to the optimal block sizes.

Simulation Results

The simulation study was also extended to include the moving blocks bootstrap. It was executed for two cases. In the first, the bootstrap time series were of length ml and the corresponding standard errors were multiplied by $\sqrt{\frac{ml}{n}}$. In the second case, the bootstrap time series lengths were all taken equal to n .

In both of these cases, the block sizes were chosen to be $l = 13$. This is the optimal value based on the minimisation of the mean squared error (MSE) of the variance of the sample mean, with respect to the block length l and not a value determined by the empirical rule of Hall, Horowitz and Jing (1995). For each block length, the bootstrap variance of the mean of the series was computed and averaged over 500 simulations. From this average and the true value of the variance of the mean of an AR(2) process, the bias was calculated. This bias was then used together with the variance of the 500 simulations to obtain the MSE. A minimum MSE was reached at the block length $l = 13$.

A comparison of the results is set out in Table 5.1 and Table 5.2. The moving blocks bootstrap doesn't seem to perform as well as the resampling of residuals. For the series of length 100 the bootstrap standard errors tend to overestimate, except in the case of the mean, where the bootstrap estimate is considerably less than

the simulated standard error. For the series of length 30, it appears as if the bootstrap generally underestimates the standard errors, except in the case of the error variance. This can be ascribed to the fact that an AR(2) structure is fit to the pseudo-time series, which are not necessarily AR(2) processes. The estimation of the error variance is largely dependent on the fit of the model. In the residual resampling case the AR(2) structure is used to reconstruct the bootstrap time series and therefore the error variance can be estimated more accurately.

Table 5.1

A comparison of the parameter estimates obtained by maximum likelihood estimation and two moving blocks bootstrap resampling schemes for a series of length $n = 100$

Parameter		Simulation results	Maximum likelihood estimation	Bootstrap series length m	Bootstrap series length n
σ_a^2	Estimate	2.00495	2.08226	3.77201	3.81258
	SD	0.28200		1.13101	1.15876
μ	Estimate	100.01815	99.26085	99.09187	99.06240
	SD	1.32986	0.91234	0.81098	0.78419
ϕ_1	Estimate	1.28383	1.30623	1.05678	1.05350
	SD	0.08912	0.08915	0.12013	0.12815
ϕ_2	Estimate	-0.40876	-0.45893	-0.28663	-0.28718
	SD	0.08592	0.08917	0.10028	0.10358



Furthermore, for both series the bootstrap estimates of the parameters are generally unacceptably far from the true values.

Table 5.2

A comparison of the parameter estimates obtained by maximum likelihood estimation and two moving blocks bootstrap resampling schemes for a series of length $n = 30$

Parameter		Simulation results	Maximum likelihood estimation	Bootstrap series length m	Bootstrap series length n
σ_a^2	Estimate	1.98757	1.85625	3.50937	3.78603
	SD	0.52908		1.72439	1.92458
μ	Estimate	99.81101	101.07308	101.91317	102.28872
	SD	2.34217	1.60515	1.27075	1.24002
ϕ_1	Estimate	1.23587	1.25915	1.03640	0.94938
	SD	0.18291	0.17031	0.18068	0.20487
ϕ_2	Estimate	-0.42543	-0.39604	-0.25881	-0.23774
	SD	0.16729	0.17200	0.14758	0.14704

Chapter 6

The Stationary Bootstrap

In the previous chapter, techniques for resampling blocks were considered. These methods share the construction of pseudo-time series by resampling blocks of observations, so that the statistic of interest may be recalculated based on the resampled data set. In the context of applying this method to stationary data, it is natural to require that, conditional on the original data, the resampled pseudo-time series should also be stationary.

In this chapter another resampling method, introduced by Politis and Romano (1994) is discussed. This stationary bootstrap is generally applicable for stationary weakly dependent time series. Similar to the block resampling techniques, the stationary bootstrap involves resampling the original data to form a pseudo-time series from which the statistic of interest may be recalculated, thereby building up an approximation to the sampling distribution of the statistic. In contrast to the aforementioned moving blocks bootstrap though, the pseudo-time series generated by the stationary bootstrap method is actually a stationary time series. This means that, conditional on the original series, Z_1, \dots, Z_n , a stationary pseudo-time series, Z_1^*, \dots, Z_n^* , is generated by an appropriate resampling scheme. Hence, this procedure attempts

to mimic the original model by retaining the stationarity property of the original series in the resampled pseudo-time series.

Basically, the pseudo-time series is generated by resampling blocks of random size, where the length of each block has a geometric distribution.

Suppose that $\{Z_t, t = 1, \dots, n\}$ is a strictly stationary and weakly dependent time series. Typically, an estimate of the sampling distribution of some statistic $T_n(\mathbf{Z}) = T_n(Z_1, \dots, Z_n)$, is required and the stationary bootstrap was developed for this purpose. Let $B_{jl} = (Z_j, \dots, Z_{j+l-1})$ be the block consisting of l consecutive observations starting from Z_j . To ensure that all observations have the same probability of being drawn a circular block scheme is introduced. In the case of $i > n$, Z_i is defined as Z_j where $j = i(\text{mod } n)$ and $Z_0 = Z_n$. Let p be a fixed number in the interval $[0, 1]$. Independent of Z_1, \dots, Z_n let L_1, L_2, \dots be a sequence of i.i.d. random variables having the geometric distribution, such that

$$P(L_i = m) = (1 - p)^{m-1} p \quad \text{for } m = 1, 2, \dots \quad (6.1)$$

Furthermore, independent of the Z_i and the L_i let I_1, I_2, \dots be a sequence of i.i.d. variables that have a discrete uniform distribution on $\{1, \dots, n\}$. Now a pseudo-time series Z_1^*, \dots, Z_n^* is generated in

the following manner. Sample a sequence of blocks of random length by the prescription $B_{I_1, L_1}, B_{I_2, L_2}, \dots$. The first L_1 observations of the pseudo-time series are determined by the first block, B_{I_1, L_1} of observations $Z_{I_1}, \dots, Z_{I_1+L_1-1}$ and the following L_2 observations of the pseudo-time series are the observations in the second sampled block B_{I_2, L_2} , namely $Z_{I_2}, \dots, Z_{I_2+L_2-1}$. Of course this process is stopped as soon as n observations have been generated for the pseudo-time series, although it is clear that this resampling method allows for time series of arbitrary length.

Once Z_1^*, \dots, Z_n^* have been generated $T_n(Z_1^*, \dots, Z_n^*)$ can be computed. By simulating a large number, B , of pseudo-time series according to the same algorithm, the true distribution of $T_n(\mathbf{Z})$ can be approximated by the empirical distribution of the B values of $T_n(\mathbf{Z}^*)$.

An alternative description of this resampling algorithm can also be given. Let Z_1^* be picked at random from the original n observations such that $Z_1^* = Z_{I_1}$. Now, let Z_2^* be picked at random from the original n observations with probability p and with probability $(1-p)$, let $Z_2^* = Z_{I_1+1}$ so that Z_2^* will be the observation in the original time series following directly after Z_{I_1} .

This means that

$$P(Z_2^* = Z_{I_2} \mid Z_1^* = Z_{I_1}) = p$$

and
$$P(Z_2^* = Z_{I_1+1} \mid Z_1^* = Z_{I_1}) = 1 - p$$

where $I_1, I_2 \sim \text{DU}(1, 2, \dots, n)$ (6.2)

In general, given that Z_i^* is determined by the J^{th} observation Z_J in the original time series, let Z_{i+1}^* be equal to Z_{J+1} with probability $(1 - p)$ and picked at random from the original n observations with probability p , i.e.

$$P(Z_{i+1}^* = Z_I \mid Z_i^* = Z_J) = p$$

and
$$P(Z_{i+1}^* = Z_{J+1} \mid Z_i^* = Z_J) = 1 - p$$

where $I \sim \text{DU}(1, 2, \dots, n)$ (6.3)

Politis and Romano (1994) proposed that by following this procedure, Z_1^*, \dots, Z_n^* is stationary, conditional on Z_1, \dots, Z_n .

Some of the similarities and differences between the stationary bootstrap and the moving blocks bootstrap are pointed out by Politis and Romano (1994). To begin with, the pseudo-time series generated by the moving blocks bootstrap is not stationary. Both methods involve sampling blocks of observations. For the moving blocks technique, the number of observations in each block is a

fixed number l , whereas in the stationary bootstrap method, the number of observations in each block is random and has a geometric distribution. The methods also differ in how they deal with end effects. For example, because there is no data after n , the moving blocks method does not define a block of length l beginning at Z_n (if $l > 1$). To achieve stationarity for the resampled time series, the stationary bootstrap method wraps the data around in a circle so that Z_1 follows Z_n .

Other variants on the stationary bootstrap based on resampling blocks of random length are also possible. Instead of assuming a geometric distribution for the L_i , other distributions can be considered. Alternative distributions for the I_i can be used as well. In this sense, the moving blocks bootstrap may be seen as a special case. The choice of the L_i having a geometric distribution and the I_i as the discrete uniform distribution was made by Politis and Romano (1994) so that the resampled series is stationary.

There is also another way to think about the difference between the moving blocks bootstrap and the stationary bootstrap. For each fixed block size l , one can compute a bootstrap distribution or an estimate of standard error of a given statistic. The stationary bootstrap proposed here is essentially a weighted average of these moving blocks bootstrap distributions or estimates of standard

error, where the weights are determined by a geometric distribution.

It is important to keep in mind that a difficult aspect in applying these methods is how to choose l in the moving blocks scheme and how to choose the geometric distribution parameter p for the stationary bootstrap. Politis and Romano (1994) showed that the stationary bootstrap estimate of variance and the moving blocks bootstrap estimate of variance are quite close, provided that p^{-1} is approximately equal to l . However, from a simulation study, they concluded that the stationary bootstrap estimate of variance is much less variable, that is it is less sensitive to the choice of p than the moving blocks bootstrap is to the choice of l . This means that the choice of p in the stationary bootstrap is less crucial than the choice of l in the moving blocks scheme.

Finally, they acknowledge that in practice, a data-based choice of p would be inevitable, but state that as long as p satisfies $p \rightarrow 0$ and $np \rightarrow \infty$, the choice of p will not enter into first-order properties, such as coverage error of the stationary bootstrap. Further work, focusing on the optimal choice of p is of vital importance and it should be investigated.

Simulation Results

In order to investigate the performance of the stationary bootstrap, it was also applied to the generated AR(2) time series.

The value for the geometric distribution parameter was taken as

$p = \frac{1}{13}$. This choice is in accordance with the optimal value of the

block length determined in the previous chapter.

Table 6.1

A comparison of the parameter estimates obtained by maximum likelihood estimation and the stationary bootstrap for a series of length $n = 100$

Parameter		Simulation results	Maximum likelihood estimation	Stationary bootstrap
σ_a^2	Estimate	2.00495	2.08226	4.12246
	SD	0.28200		1.29009
μ	Estimate	100.01815	99.26085	99.24669
	SD	1.32986	0.91234	0.71283
ϕ_1	Estimate	1.28383	1.30623	1.01484
	SD	0.08912	0.08915	0.14706
ϕ_2	Estimate	-0.40876	-0.45893	-0.25860
	SD	0.08592	0.08917	0.11333

Tables 6.1 and 6.2 contain the results found for this procedure. As shown by Politis and Romano (1994) the standard error estimates are relatively close to those obtained with the moving blocks bootstrap. For the series of length 100, the stationary bootstrap overestimates the standard errors, except in the case of the mean. The same is true for the series of length 30, but the standard error for the second autoregressive parameter ϕ_2 , is underestimated. These results may also be due to the fact that the pseudo-time series are assumed to be AR(2) processes. The parameter estimates are also disappointingly far from the simulated values.

Table 6.2

A comparison of the parameter estimates obtained by maximum likelihood estimation and the stationary bootstrap for a series of length $n = 30$

Parameter		Simulation results	Maximum likelihood estimation	Stationary bootstrap
σ_a^2	Estimate	1.98757	1.85625	3.73788
	SD	0.52908		1.65648
μ	Estimate	99.81101	101.07308	101.40458
	SD	2.34217	1.60515	1.08577
ϕ_1	Estimate	1.23587	1.25915	0.98168
	SD	0.18291	0.17031	0.20978
ϕ_2	Estimate	-0.42543	-0.39604	-0.24906
	SD	0.16729	0.17200	0.14851

Chapter 7

Comparing the Methods

This chapter offers an overall comparison of all the methods discussed so far in terms of the sampling distributions of the parameters of the time series. They are judged both on location as well as spread. The standard errors of the parameters as obtained by the different methods are also compared in order to determine the success of the methods to estimate accuracy measures.

Based on these comparisons, the performance of the different methods may be evaluated and the use of the most reliable methods might be suggested. The comparisons are presented in the form of graphs that clearly show the difference in the methods. Apart from the graphs, only a short discussion is offered to point out the most important findings.

Comparison of the Sampling Distributions

The following abbreviations are used to distinguish between the methods:

- sim - The simulated distribution
- res (c) - Resampling from the centered residuals

- res (cs) - Resampling from the centered residuals, scaled by the factor $1/\sqrt{1-p/n}$
- mbb (ml) - The moving blocks bootstrap with $n^* = ml$
- mbb (n) - The moving blocks bootstrap with $n^* = n$
- stat - The stationary bootstrap

Figures 7.1 and 7.2 represent the sampling distributions of σ_a^2 for the two series lengths. It is clear that the residual resampling methods closely resemble the simulation results in both location and spread. In contrast, the moving blocks methods and the stationary bootstrap show a much wider spread, with a peak at a higher value than that obtained in the simulation. This is true for both the longer and the shorter series.

In order to compare the sampling distributions of the series mean μ , Figures 7.3 and 7.4 are considered. In this case, all the bootstrap methods seem to provide similar sampling distributions. For both the series, the bootstrap spread is narrower than the spread of the simulated values. This is to be expected as the bootstrap time series are reconstructed from one original series and are dependent on the mean of that series. The means of the simulated series vary more because they were generated independently. Therefore the spread of the bootstrap series means should be narrower than that of the different simulated series.

Figure 7.1

Comparison of the sampling distributions of σ_a^2 for a series of length $n = 100$ obtained by the different methods




Figure 7.2

Comparison of the sampling distributions of σ_a^2 for a series of length $n = 30$ obtained by the different methods




Figure 7.3

Comparison of the sampling distributions of μ for a series of length $n = 100$ obtained by the different methods





Figure 7.4

Comparison of the sampling distributions of μ for a series of length $n = 30$ obtained by the different methods



For the longer series, the simulated peak is at a higher value than the peaks reached by the bootstrap methods. The converse is true in the case of the shorter series, where the simulated peak is at a lower value than those of the bootstrap methods.

Figure 7.5 and Figure 7.6, serve to depict the sampling distributions of the first autoregressive parameter ϕ_1 . Once again, the residual resampling methods resemble the simulated distribution very closely.

For the series of length 100, the moving blocks methods and the stationary bootstrap show a peak at a lower value and a slightly wider spread than the simulated distribution. Although the spread seems more comparable, the moving blocks methods also peak at a lower value for the series of length 30.

Lastly, the sampling distributions obtained for the second autoregressive parameter ϕ_2 are examined by Figures 7.7 and 7.8. When the longer series is considered, it is seen that the residual resampling methods peak at a lower value and the moving blocks and stationary bootstrap methods peak at a higher value than the simulated distribution.

The spread of the residual resampling methods are however closer to the simulated spread, while the other three methods tend to be spread a bit more widely. In the case of the shorter series, the

Figure 7.5

Comparison of the sampling distributions of ϕ_1 for a series of length $n = 100$ obtained by the different methods




Figure 7.6

Comparison of the sampling distributions of ϕ_1 for a series of length $n = 30$ obtained by the different methods

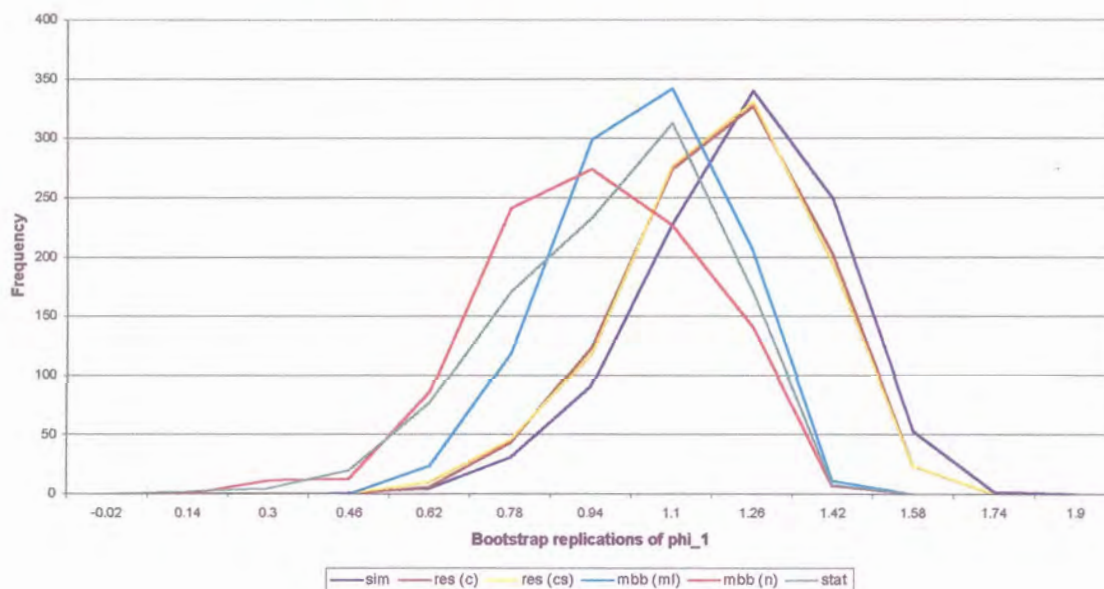


Figure 7.7

Comparison of the sampling distributions of ϕ_2 for a series of length $n = 100$ obtained by the different methods





Figure 7.8

Comparison of the sampling distributions of ϕ_2 for a series of length $n = 30$ obtained by the different methods



sampling distributions obtained by the residual resampling schemes once again correspond to the simulated distribution, while the moving blocks methods and the stationary bootstrap peak at a slightly higher value.

Comparison of the Standard Errors

The bootstrap methodology is especially useful for the estimation of standard errors. Therefore the standard errors of the parameters as obtained by the different methods are of crucial importance. The rest of the figures presented in this chapter are thus concerned with this comparison. In the case of the error variance σ_a^2 , maximum likelihood estimation offers no estimate of the standard error but in all the other cases the maximum likelihood estimates are also included in the discussion.

For both the short and the long series the standard errors of σ_a^2 obtained by the residual resampling methods compare very well with the simulated values, while the moving blocks and stationary bootstrap methods tend to greatly overestimate this quantity. This can be seen in the graphs in Figure 7.9 and Figure 7.10.

Figures 7.11 and 7.12 show that in the case of estimating the standard error of the series mean μ , all the estimation methods

Figure 7.9

Comparison of the standard errors of σ_a^2 for a series of length $n = 100$ obtained by the different methods




Figure 7.10

Comparison of the standard errors of σ_a^2 for a series of length $n = 30$ obtained by the different methods




Figure 7.11
Comparison of the standard errors of μ for a series of length
 $n = 100$ obtained by the different methods





Figure 7.12
Comparison of the standard errors of μ for a series of length
 $n = 30$ obtained by the different methods



provide a value that is lower than the simulated standard error. The estimates obtained by the residual resampling methods correspond better to the maximum likelihood estimate, while the moving blocks and the stationary bootstrap methods result in even lower standard error estimates.

One has to keep in mind that the residual resampling methods are dependent on the maximum likelihood estimates. This is because of the fact that the maximum likelihood estimates are used in the generation of the bootstrap time series. It is therefore reasonable to expect the estimates of the residual resampling methods to follow the behaviour of the maximum likelihood estimates.

The standard error estimates of the first autoregressive parameter ϕ_1 , are considered in Figure 7.13 and Figure 7.14. It once again seems that the methods of resampling the residuals outperform the rest. This is particularly true in the case of the series of length 100 where the standard error estimates calculated according to the moving blocks and stationary bootstrap methods are higher than the simulated values and also do not compare well with the maximum likelihood estimate. When looking at the shorter series, the difference is not that clear with only the last two methods resulting in a slightly higher estimate of standard error.

The last parameter to be considered in Figure 7.15 and Figure 7.16 is the second autoregressive parameter ϕ_2 . Following the general

Figure 7.13

Comparison of the standard errors of ϕ_1 for a series of length $n = 100$ obtained by the different methods




Figure 7.14

Comparison of the standard errors of ϕ_1 for a series of length $n = 30$ obtained by the different methods




Figure 7.15

Comparison of the standard errors of ϕ_2 for a series of length $n = 100$ obtained by the different methods





Figure 7.16

Comparison of the standard errors of ϕ_2 for a series of length $n = 30$ obtained by the different methods



trend, the estimates of standard error provided by the moving blocks and stationary bootstrap methods were once again the least satisfactory. For the long series these methods overestimated and for the short series they underestimated the standard errors as compared to the simulated values.

After studying these comparative sampling distributions and estimates of standard error, it becomes apparent that in the context of this study the methods based on the resampling of the residuals are superior to the other bootstrap methods tried. In addition, they compare well with the simulated values and those obtained by maximum likelihood estimation.

Of course one has to remember the dependency on the observed series' representation of the underlying mechanism being modelled as well as the accuracy of the maximum likelihood estimation used in the construction of the bootstrap series. When these assumptions are handled with caution, the residual resampling methods may be used with confidence in situations similar to those considered in this study.

Chapter 8

Bootstrap Prediction Intervals

When studying a time series, one of the main goals is to forecast its future values. However, point estimates are seldom really helpful unless their reliability is assessed. Therefore, interval forecasts are often preferable because they take into account the predictor variability. The traditional method of forecasting for ARMA and ARIMA models is due to Box, Jenkins and Reinsel (1994). However, their method assumes a Gaussian distribution for the errors and does not consider the variability caused by the estimation of the model parameters.

This chapter describes the application of the bootstrap technique to the construction of prediction intervals. Masarotto (1990) highlighted that the resulting method seems interesting, since

- i. it is distribution-free
- ii. it explicitly takes into account that the parameters and the order of the model is unknown, and
- iii. the suggested probability limits can be computed using available computer software.

Consider the stationary autoregressive time series $\{Z_t\}$ which satisfies

$$Z_t = \phi_1 Z_{t-1} + \phi_2 Z_{t-2} + \dots + \phi_p Z_{t-p} + a_t \quad (8.1)$$

where the a_t are i.i.d. random variables and $\Phi = (\phi_1, \dots, \phi_p)$ are parameters such that the roots of the characteristic equation $1 - \phi_1 x - \dots - \phi_p x^p = 0$ are greater than one in absolute value. In addition assume that the a_t have zero mean, finite variance σ_a^2 and common distribution function $F(\cdot)$. The problem to be considered is that of predicting the future value Z_{n+h} , having observed (Z_1, \dots, Z_n) . It is well known that the minimum squared error linear predictor, $Z_n(h)$ can be obtained, when p and Φ are known, by the recursion

$$Z_n(h) = \phi_1 Z_n(h-1) + \dots + \phi_p Z_n(h-p) \quad (8.2)$$

where it is understood that $Z_n(h) = Z_{n+h}$ when $h \leq 0$.

When the model is unknown, equation (8.2) is used with p and Φ replaced by suitable estimates \hat{p} and $\hat{\Phi}$. Denote the resulting predictions by $\hat{Z}_n(h)$.

The usual $(1 - 2\alpha)$ -level prediction interval for Z_{n+h} , due to Box, Jenkins and Reinsel (1994), is given by

$$\left[\hat{Z}_n(h) - z_{1-\alpha} \hat{\sigma}_a \hat{v}_b, \hat{Z}_n(h) + z_{1-\alpha} \hat{\sigma}_a \hat{v}_b \right] \quad (8.3)$$

Here z_α is the α^{th} percentage point of a standard normal variable, $\hat{\sigma}_a$ is an estimate of σ_a and $\hat{v}_b^2 = 1 + \hat{\psi}_1^2 + \dots + \hat{\psi}_{b-1}^2$, where the ψ_j^2 are the coefficients of the infinite order moving average representation of the fitted model.¹

When using (3), the following assumptions are made:

- i. $p = \hat{p}$, i.e. the true model order is \hat{p}
- ii. $\Phi = \hat{\Phi}$, i.e. the true model parameters are $\hat{\Phi} = (\hat{\phi}_1, \dots, \hat{\phi}_p)$
and
- iii. $F(x) = \Phi\left(\frac{x}{\sigma_a}\right)$, i.e. the a_t 's distribution $F(\cdot)$ is Gaussian with zero mean and variance σ_a^2 , where $\Phi(x)$ is the distribution function of a standard Gaussian variable.

Hence (8.3) is based on a parametric assumption about $F(\cdot)$ and in addition, it does not account for the sampling variability of the

¹ See Appendix A on the infinite order moving average representation of autoregressive models

parameter estimates $\hat{\Phi}$ and \hat{p} . In order to set up exact probability limits for Z_{n+b} , the distribution of the prediction errors

$$e_n(b) = Z_{n+b} - \hat{Z}_n(b) \quad (8.4)$$

must be known. Denote this distribution as $H_b(\cdot; p, \Phi, F)$.

Alternatively, the distribution, $G_b(\cdot; p, \Phi, F)$ of the standardised prediction errors

$$r_n(b) = \frac{e_n(b)}{\sigma_a \nu_b} \quad (8.5)$$

can be used.

However, these distributions are generally unknown. This chapter aims to discuss an estimate for $H_b(\cdot; p, \Phi, F)$ or $G_b(\cdot; p, \Phi, F)$, which is not based on parametric assumptions about $F(\cdot)$ (Masarotto, 1990). The focus is on the second distribution, $G_b(\cdot; p, \Phi, F)$, because the forecast intervals based on the standardised prediction error have in general coverage probabilities closer to the nominal level (Li and Maddala, 1996; Kabaila, 1993), but the method described here can easily be modified to estimate $H_b(\cdot; p, \Phi, F)$.

Firstly, suppose that the order of the model, p , is given. Since Φ and $F(\cdot)$ are not known, $G_b(\cdot; p, \Phi, F)$ cannot be used to set up probability limits for Z_{n+h} . However, consistent estimators can be found for Φ and $F(\cdot)$ from the available data. In particular, Φ can be estimated using some variant of the method of least squares while a non-parametric estimate of $F(\cdot)$ can be obtained from the residuals

$$\hat{a}_t = Z_t - \hat{\phi}_1 Z_{t-1} - \dots - \hat{\phi}_p Z_{t-p} \quad t = 1, \dots, n \quad (8.6)$$

where $\hat{\Phi} = (\hat{\phi}_1, \dots, \hat{\phi}_p)$ denotes the estimate of Φ .

For example, a possible estimate of $F(\cdot)$ is the empirical distribution of the residuals i.e. the distribution which assigns probability n^{-1} to each point \hat{a}_t , $t = 1, \dots, n$. Other estimates of $F(\cdot)$ can also be used, as discussed in Chapter 4 on the methods that utilise the resampling of residuals.

Let $\hat{F}(\cdot)$ be an estimate of $F(\cdot)$. Then $G_b(\cdot; p, \hat{\Phi}, \hat{F})$ seems a natural estimator of the standardised prediction error distribution.

This is equivalent to the approximation of the unknown distribution of $r_n(b)$ with the distribution of

$$r_n^*(b) = \frac{Z_{n+b}^* - \hat{Z}_n^*(b)}{\sigma_a^* \upsilon_b^*} \quad (8.7)$$

where $\{Z_t^*\}$ are random variables which satisfy

$$Z_t^* = \hat{\phi}_1 Z_{t-1}^* + \dots + \hat{\phi}_p Z_{t-p}^* + a_t^* \quad (8.8)$$

and $\hat{Z}_n^*(b)$, σ_a^* and υ_b^* are obtained from (Z_1^*, \dots, Z_n^*) with the same method used to calculate $\hat{Z}_n(b)$, $\hat{\sigma}_a$ and $\hat{\upsilon}_b$ from (Z_1, \dots, Z_n) . In this case, the a_t^* are independent random variables with distribution $\hat{F}(\cdot)$.

Exact computation of $G_b(\cdot; p, \hat{\Phi}, \hat{F})$ is not possible. However, the distribution can be approximated by means of the Monte Carlo method in the following way. Firstly, obtain $\hat{\Phi}$, $\hat{F}(\cdot)$ and $\hat{Z}_n(b)$ from the observed data. Generate $n+b$ pseudo-data $(Z_1^*, \dots, Z_{n+b}^*)$ from model (8.8) and, using the same estimation procedure as before, compute $\hat{Z}_n^*(b)$, σ_a^* and υ_b^* from the first n pseudo-data and then calculate the pseudo-error according to (8.7). This generation of pseudo-time series and estimation of the pseudo-error is repeated B times, where B is some large number.

Finally, take the $(1 - 2\alpha)$ -level prediction interval for Z_{n+b} to be

$$I_b = \left[\hat{Z}_n(b) + \hat{\sigma}_a \hat{\nu}_b r_n^*(b)^{(j)} \quad , \quad \hat{Z}_n(b) + \hat{\sigma}_a \hat{\nu}_b r_n^*(b)^{(B-j)} \right] \quad (8.9)$$

where $j = \lfloor B\alpha \rfloor$, with $\lfloor x \rfloor$ denoting the integer part of x and $(r_n^*(b)^{(1)}, \dots, r_n^*(b)^{(B)})$ are the ordered bootstrap replications of the standardised prediction errors.

Masarotto (1990) also showed that, under reasonable requirements for a_t , $\hat{\Phi}$ and $\hat{F}(\cdot)$, the obtained prediction intervals are asymptotically consistent in the sense that

$$\lim_{n,m} P\{Z_{n+b} \in I_b\} = 1 - 2\alpha$$

In practice, it cannot be assumed that p is known. However, it is easy to adopt the given approach if p is estimated by means of some model selection criteria. To do that, let \hat{p} be an estimate of p . Then $(Z_1^*, \dots, Z_{n+b}^*)$ can be found by substituting \hat{p} for p in (8.8) and from this pseudo-series, $r_n^*(b)$ can be computed.

Masarotto (1990) notes that in order to account for the variability in the estimate of p , the point forecast $\hat{Z}_n(b)$ and the related

quantities σ_a^* and υ_b^* , must be computed with reference to a model of order p^* , where p^* is obtained from the first n pseudo-data with the same method used to compute p from the observed series.

Masarotto (1990) also showed that the resulting prediction intervals are consistent for the usual model selection criteria in particular the Akaike AIC criterion (1974), the Hannan and Quinn criterion (1979) and the Schwartz criterion (1978).

An Improved Estimator

Grigoletto (1998) showed that better probability limits than those introduced by Masarotto (1990), can be found. The bootstrap forecast may be expressed as

$$\hat{Z}_n^*(b) = \sum_{i=1}^P d_{bi}(\Phi^*) Z_{n+1-i}^* \quad (8.10)$$

where $d_{bi}(\cdot)$, $i = 1, \dots, P$ are continuous functions and P is the maximum lag tried in the selection of the order.

Then

$$\begin{aligned}
 r_n^*(b) &= \frac{Z_{n+b}^* - \hat{Z}_n^*(b)}{\sigma_a^* \nu_b^*} \\
 &= \frac{\left\{ Z_{n+b}^* - \sum_{i=1}^P d_{bi}(\hat{\Phi}) Z_{n+1-i}^* \right\} + \left\{ \sum_{i=1}^P [d_{bi}(\hat{\Phi}) - d_{bi}(\Phi^*)] Z_{n+1-i}^* \right\}}{\sigma_a^* \nu_b^*} \\
 &= \frac{A + B}{C}
 \end{aligned} \tag{8.11}$$

The estimator of the α -percentile of the standardised prediction error bootstrap distribution can be seen as the x value which satisfies

$$U_B = \frac{1}{B} \sum_{i=1}^B I \left\{ \frac{A_i + B_i}{C_i} \leq x \right\} \approx \alpha \tag{8.12}$$

where A_i , B_i and C_i , $i = 1, \dots, B$ are the outcomes of the i^{th} bootstrap replication and $I(\cdot)$ is the indicator function.

Refer to Grigoletto (1998) for a more detailed explanation, as well as the most efficient use of this proposed method. He has claimed that this new estimating procedure leads to a substantial reduction in the variances of the predictive distribution percentile estimators.

Simulation Results

Prediction intervals for Z_{n+1}, \dots, Z_{n+5} were constructed using all of the methods discussed in this study. The resulting intervals are compiled in Tables 8.1 to 8.4.

For a series with 100 observations, Table 8.1 summarises the prediction intervals obtained for lead time one, while Table 8.2 provides the results for lead time five.

Both the residual resampling methods compare very well with the intervals obtained by maximum likelihood estimation. These intervals are all in line with the simulated intervals.

The moving blocks bootstrap results seem less appropriate. Although not that far off, the intervals obtained by the moving blocks method with $n^* = n$ are generally narrower than the simulated intervals. It is however the method with $n^* = m$ that is troubling. The very wide intervals raise suspicion and therefore the procedure was repeated. This time the interval for lead time one was forecast as (90.1557 ; 104.4018) and for lead time five it was (88.8079 ; 105.8884). Using this moving blocks method for prediction interval calculation is thus not recommended, although the reason for this unexpected result should be investigated.

Finally, the stationary bootstrap also results in intervals that are in good comparison with the simulated intervals.

Table 8.1

Summary of 95% prediction intervals for Z_{n+1} obtained by the respective methods for a series of length $n = 100$

Method	Lower bound	Upper bound	Length of interval
Simulation	93.9812	100.1760	6.1948
Maximum likelihood	94.1337	99.7901	5.6564
Resampling from $(\hat{a}_t - \bar{a})$	93.7537	99.5997	5.8460
Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$	93.9130	99.5568	5.6438
Moving blocks with $n^* = ml$	89.6360	104.3145	14.6785
Moving blocks with $n^* = n$	94.1409	99.0750	4.9341
Stationary bootstrap	93.9002	99.8602	5.9600

Table 8.2

Summary of 95% prediction intervals for Z_{n+5} obtained by the respective methods for a series of length $n = 100$

Method	Lower bound	Upper bound	Length of interval
Simulation	90.8837	105.8718	14.9881
Maximum likelihood	91.6101	105.3725	13.7624
Resampling from $(\hat{a}_t - \bar{a})$	90.7183	105.5992	14.8809
Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$	90.6544	105.9882	15.3338
Moving blocks with $n^* = ml$	89.2167	105.6776	16.4609
Moving blocks with $n^* = n$	89.6324	103.6326	14.0002
Stationary bootstrap	88.6778	105.5382	16.8604

The same prediction intervals were constructed for the shorter series with 30 observations. Those results are given in Table 8.3 and Table 8.4.

For lead time one, the residual resampling methods provide intervals that are very close to the simulated interval. In this case, the maximum likelihood interval is slightly narrower.

Both the moving blocks methods provide rather narrow intervals, with the method using $n^* = m$ once again being the worst performer.

The stationary bootstrap doesn't compare badly, providing a slightly wider interval than the simulated one.

The least consistent results were obtained for the short series with a lead time five. Resampling the scaled, centered residuals offered the interval comparing closest to the simulated values, while the boundaries obtained by maximum likelihood estimation leads to an interval that might prove to be too narrow.

The other bootstrap methods don't appear to be faring badly, except for the moving blocks method with $n^* = n$ that resulted in a wide interval of which especially the lower bound is less than expected.

Table 8.3

Summary of 95% prediction intervals for Z_{n+1} obtained by the respective methods for a series of length $n = 30$

Method	Lower bound	Upper bound	Length of interval
Simulation	95.6075	101.4563	5.8488
Maximum likelihood	95.7513	101.0920	5.3407
Resampling from $(\hat{a}_t - \bar{a})$	95.0348	100.8797	5.8449
Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$	95.0761	100.9874	5.9113
Moving blocks with $n^* = ml$	97.7511	100.9532	3.2021
Moving blocks with $n^* = n$	95.3052	100.1493	4.8441
Stationary bootstrap	95.2070	101.5447	6.3377

Table 8.4

Summary of 95% prediction intervals for Z_{n+5} obtained by the respective methods for a series of length $n = 30$

Method	Lower bound	Upper bound	Length of interval
Simulation	90.7074	108.9629	18.2555
Maximum likelihood	93.7688	106.4265	12.6577
Resampling from $(\hat{a}_t - \bar{a})$	91.2115	108.7503	17.5388
Resampling from $\frac{(\hat{a}_t - \bar{a})}{\sqrt{1 - p/n}}$	90.5343	108.6828	18.1485
Moving blocks with $n^* = ml$	89.6484	106.1413	16.4929
Moving blocks with $n^* = n$	82.6086	105.4858	22.8772
Stationary bootstrap	89.3027	105.8956	16.5929



Chapter 9

Discussion of Bootstrap Computations

In order to apply the different bootstrap methods, programming was done in SAS. A program was written for each bootstrapping procedure in order to estimate the parameters for the second order autoregressive time series model. These programs are set out according to the different algorithms as described in the relevant chapters.

After executing the programs, it was found that the residual resampling methods are the most reliable. This is evident from the results of the simulation study and is clearly seen in the comparison between the methods in Chapter 7. For this reason, only the program that performs the resampling from the centered, scaled residuals $(\hat{a}_t - \bar{a})/\sqrt{1 - p/n}$, will be discussed here.

A copy of this SAS program is given on the enclosed diskette under the filename "BOOTRES.SAS". Although the simulation study was only performed on an AR(2) model, this program was adjusted to accommodate autoregressive models of any order. To facilitate the use of the SAS bootstrap application, another program "BTS01.EXE" is included. This program serves to capture the necessary information for the application, which includes the



length of the time series and the autoregressive order to be modelled. As this application does not include model identification, the user first has to examine the data in order to determine the correct autoregressive order to be applied to the observed time series. When the file “BTS01.EXE” is run from the diskette, the necessary files are created on the diskette. The SAS program reads these files from the A drive. If the user’s stiffy is not located on the A drive, the path specified in the main part of the SAS program “BOOTRES.SAS”, needs to be changed to the appropriate drive.

The observed series can be captured in two ways. Firstly, if the data is already available on an existing file, this file need only be copied onto the diskette under the name that is referenced by the SAS application which is “BTSDATA.DAT”. Figure 9.1 shows an example of the correct format to be used for the file containing the observations.

The second option is to capture each observation individually as prompted. In this case, the data file is automatically created and the number of observations is counted. The user will thus not be asked to specify the length of the time series.

As soon as the necessary data has been captured, the SAS application can be executed. This is done by submitting the program statements of “BOOTAPPL.SAS” to the SAS System.

Figure 9.1

Example of the format to be used for the file containing the observed data

```
101.314  
103.351  
105.058  
105.492  
103.563  
101.691  
99.349  
97.522  
93.797  
92.486  
91.894  
88.999  
87.944  
88.905  
93.516  
96.550  
100.763  
100.702  
100.039  
100.333  
...  
...  
...
```

The SAS program starts by reading the captured data from the correct files. An autoregressive model is fit to the observed series in order to estimate the vector of parameters $\Phi = (\phi_1, \dots, \phi_p)$, which is used to construct the bootstrap time series. Then the program executes a macro that performs the bootstrap computations. Within this macro, the process is repeated for $B = 1000$ replications.

For each replication, the macro involves centering and multiplying the residuals with the factor $1/\sqrt{1-p/n}$ and then draws a random sample of size n from these adjusted residuals. Using these resampled residuals, a pseudo-time series is constructed and the parameters $\phi_1^*, \dots, \phi_p^*$ are estimated for this bootstrap series.



Finally, the standardised prediction errors are calculated which are used in the construction of the forecast intervals.

When the macro ends, data sets containing 1000 bootstrap replicates of the parameter estimates and the pseudo-errors are available. The sample mean and standard deviation of the parameter values are calculated to obtain the bootstrap estimates. Finally, the pseudo-errors are sorted and the forecast intervals are calculated accordingly.

A brief description of the given SAS output is as follows. Firstly the output prints the autoregressive order, the number of observations as well as the observed series. Furthermore, it provides bootstrap estimates of the parameters as well as their standard deviations. The variable ERR_VAR denotes the error variance σ_a^2 ; the variable MU is the series mean μ and the rest of the variables are the autoregressive parameters. This means that COL3 is equivalent to ϕ_1 , COL4 to ϕ_2 and so forth. Then 95% forecast intervals for the first five step forecasts are also displayed. The output is also written to files on the diskette. These files are "ESTIM.OUT" containing the estimates and their standard errors and "INTERVAL.OUT" containing the bounds of the forecast intervals.

An example of the output is given in Figure 9.2 where the bootstrap was applied to a stationary AR(2) model with 100



observations. The one step forecast interval is (102.975; 108.462) and the five step forecast interval is (93.664; 108.462) .

Figure 9.2
Example of the output obtained by the SAS application

```
The SAS System
Autoregressive order
  OBS      P
   1       2

The SAS System
Number of observations
  OBS      N
   1     100

The SAS System
The observed series
  OBS      YT
   1    102.097
   2    102.404
   3    102.675
   4    102.723
   5    102.202
   6    102.139
   7    103.160
   8    102.353
   9    103.813
  10    105.845
  11    105.373
  12    106.024
  13    105.323
      ...
      ...
      ...
```



Figure 9.2 (continued)
Example of the output obtained by the SAS application

The SAS System

Bootstrap estimates of the parameters
and their standard deviations

Variable	Mean	Std Dev
ERR_VAR	2.1331355	0.2799856
MU	99.2800188	0.8939345
COL3	1.2896609	0.0879916
COL4	-0.4606597	0.0875169

The SAS System

95% forecast intervals for the first five step forecasts

OBS	LOWER	UPPER
1	102.975	108.462
2	100.181	109.686
3	97.286	109.226
4	95.240	108.980
5	93.664	108.462

It was found however, that performing the bootstrap procedures in SAS is not particularly efficient. The execution times of the programs are unfavourably long and they would not be fit for general use. Better approaches surely exist and should be investigated and developed further.

Chapter 10

Conclusion

In this study, the methodology and performance of bootstrapping procedures were studied. Different methods were discussed in the context of autoregressive time series and tested on a specific AR(2) model.

For this specific application the resampling of residuals method seems to perform better than the moving blocks, or stationary bootstrap schemes. The standard error estimates obtained by the residual resampling procedure compare very well to the simulated values, which suggest that the bootstrap competes well with conventional methods. The value of the two blocking methods should however, not be underestimated. Further study in the field of the optimal block choice and the choice of the geometric distribution parameter, might prove to be extremely valuable.

One would hope that the bootstrap performs well for a short series in comparison to traditional estimation methods that are based on asymptotic results. Such a statement cannot be made based on the results of this study. This might be due to the fact that the residual resampling methods that were found to work most satisfactory, are still dependent on the maximum likelihood estimates.

One aspect of the formulation of time series models was not considered in this study. This is actually the first stage of the process, namely the model identification. For the purpose of this study, the order of the models were considered to be fixed, because of the fact that an AR(2) model was generated. In practical situations, such an assumption would of course not be possible. Therefore, this first step is of great importance, since an incorrectly identified model would likely lead to poor forecasts.

Aczel and Josephy (1992) propose that a bootstrap procedure can make a difference in model identification by improving accuracy and thereby leading to a correct model being identified more often and resulting in better forecasts. They show how the bootstrap method can be used as an alternative to the normal theory in estimating the distribution of sample autocorrelations and setting confidence bounds for these parameters.

Shao (1996) developed a bootstrap model selection procedure in the context of linear models which can easily be extended to more complicated problems, such as autoregressive time series. This procedure is based on the minimisation of the average conditional expected loss in prediction. This means that the selected model is the one with the best prediction ability. He points out that it might be preferable to use the same method both in model selection and in the subsequent inference based on the selected model. The bootstrap observations generated for the model

selection can then also be used in the inference, which means that there is no extra cost in terms of generating bootstrap observations.

A complete application of the bootstrap methods will definitely contain a bootstrap model identification procedure. This is thus another field that needs to be explored in greater detail.

The conclusions reached in this study, are based on one application of the bootstrap methods. In order to make strong statements, this experiment should ideally be repeated a number of times. Only then can the comparative performance of the bootstrap methods be judged.

An extension of these methods, to be applied to higher order models as well as moving average and mixed models would make a valuable contribution to the statistical community.

The SAS programs used in this study were a basic attempt to apply the methods and compare the results. They were not found to be sufficiently effective and other options should be researched. Ultimately, the development of a software package that includes these bootstrapping methods is a goal worth striving for and an exciting prospect.

Appendix A

The Infinite Order Moving Average Representation of an Autoregressive Model

A *moving average process* of order q , can be represented as a weighted linear combination of $q + 1$ independent and identically distributed random variables, i.e

$$Z_t = a_t + \psi_1 a_{t-1} + \psi_2 a_{t-2} + \dots + \psi_q a_{t-q}$$

Any autoregressive series can be rewritten in the form of an infinite order moving average process. In particular, the coefficients ψ_i in the MA(∞) representation for an AR(2) series are given by

$$\psi_0 = 1$$

$$\psi_1 = \phi_1$$

$$\psi_k = \phi_1 \psi_{k-1} + \phi_2 \psi_{k-2} \quad \text{for } k = 2, 3, \dots$$

References

- Aczel A.D. and Josephy N.H. (1992) Using the Bootstrap for Improved ARIMA Model Identification, *Journal of Forecasting*, **11**, 71-80
- Akaike, H. (1974) A new look at the statistical model identification, *IEEE Transactions on Automatic Control*, **AC-19**, 716-723
- Babu G.J. (1986) A note on bootstrapping the variance of sample quantiles, *Ann. Inst. Statist. Math.*, **38**, 439-443
- Box G.E.P. , Jenkins G.M. and Reinsel G.C. (1994) *Time Series Analysis: Forecasting and Control* Third Edition. Prentice-Hall, New Jersey
- Carlstein E. (1986) The use of subseries values for estimating the variance of a general statistic from a stationary sequence, *Annals of Statistics*, **14**, 1171-1194
- Chatterjee S. (1986) Bootstrapping ARMA Models: Some Simulations, *IEEE Transactions on Systems, Man and Cybernetics*, **SMC-16**, 294-299
- Cryer J.D. (1986) *Time Series Analysis* Duxbury Press, Boston
- Efron B. (1979) Bootstrap Methods: Another look at the jackknife, *Annals of Statistics*, **7**, 1-26
- Efron B. and Tibshirani R.J. (1986) Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy, *Statistical Science*, **1**, 54-77
- Efron B. and Tibshirani R.J. (1993) *An Introduction to the Bootstrap* Chapman & Hall, New York

- Fitzenberger B. (1997) The moving blocks bootstrap and robust inference for linear least squares and quantile regressions, *Journal of Econometrics*, **82**, 235-287
- Freedman D.A. (1981) Bootstrapping Regression Models, *The Annals of Statistics*, **9**, 1218-1228
- Grigoletto M (1998) Bootstrap prediction intervals for autoregressions: some alternatives, *International Journal of Forecasting*, **14**, 447-456
- Hall P. (1988a) Theoretical comparison of bootstrap confidence intervals, *The Annals of Statistics*, **16**, 927-985
- Hall P. (1988b) On symmetric bootstrap confidence intervals, *Journal of the Royal Statistical Society, Series B*, **50**, 35-45.
- Hall P., Horowitz J.L. and Jing B. (1995) On blocking rules for the bootstrap with dependent data, *Biometrika*, **82**, 561-574
- Hamilton J.D. (1994) *Time Series Analysis* Princeton University Press, New Jersey
- Hannan E.J. and Quinn B.G. (1979) The determination of the order of an autoregression, *Journal of the Royal Statistical Society, Series B*, **41**, 190-195
- Holbert D. and Son M.S. (1986) Bootstrapping a time series model: Some empirical results, *Communications in Statistics*, **15**, 3669-3691
- Kabaila P. (1993) On bootstrap predictive inference for autoregressive processes *Journal of Time Series Analysis*, **14**, 473-484
- Künsch H.R. (1989) The Jackknife and the Bootstrap for general Stationary Observations *The Annals of Statistics*, **17**, 1217-1241

- Lahiri S. (1992) Edgeworth Correction by Moving Blocks Bootstrap for Stationary and Nonstationary Data, *Exploring the Limits of Bootstrap* eds. R. LePage and L. Billard, John Wiley, New York
- Li H and Maddala G.S. (1996) Bootstrapping time series models *Econometric reviews*, 15, 115-158
- Liu R.Y. (1988) Bootstrap procedures under some non-i.i.d. models *Annals of Statistics*, 16, 1697-1708
- Liu R.Y. and Singh K. (1992a) Efficiency and robustness in resampling *Annals of Statistics*, 20, 370-384
- Liu R.Y. and Singh K. (1992b) Moving Blocks Jackknife and Bootstrap Capture Weak Dependence, *Exploring the Limits of Bootstrap* eds. R. LePage and L. Billard, John Wiley, New York
- Marquardt D.W. (1963) An algorithm for least square estimation of nonlinear parameters, *SIAM Journal*, 11, 2, 431-441
- Masarotto G. (1990) Bootstrap prediction intervals for autoregressions, *International Journal of Forecasting*, 6, 229-239
- McCullough B.D. (1994) Bootstrapping Forecast Intervals: An Application to AR(p) Models, *Journal of Forecasting*, 13, 51-66
- McCullough B.D. (1996) Consistent Forecast Intervals when the Forecast-period Exogenous Variables are Stochastic, *Journal of Forecasting*, 15, 293-304
- Politis D.N. and Romano J.P. (1992a) A Circular Block Resampling Procedure for Stationary Data, *Exploring the Limits of Bootstrap*, eds. R. LePage and L. Billard, John Wiley, New York

- Politis D.N. and Romano J.P. (1992b) A general resampling scheme for triangular arrays of α -mixing random variables with application to the problem of spectral density estimation, *Annals of Statistics*, **20**, 1985-2007
- Politis D.N. and Romano J.P. (1994) The Stationary Bootstrap, *Journal of the American Statistical Association*, **89**, 1303-1313
- Politis D.N., Romano J.P. and Wolf M. (1997) Subsampling for heteroskedastic time series, *Journal of Econometrics*, **81**, 281-317
- Schwartz, G. (1978) Estimating of dimension of a model, *Annals of Statistics*, **6**, 461-464
- Shao J. (1996) Bootstrap Model Selection, *Journal of the American Statistical Association*, **91**, 655-665
- Shao J. and Tu D. (1995) *The Jackknife and Bootstrap* Springer-Verlag, New York
- Shao J. and Yu H. (1993) Bootstrapping the sample means for stationary mixing sequences, *Stochastic Processes and Their Applications*, **48**, 175-190
- Singh K. (1981) On the asymptotic accuracy of Efron's bootstrap, *Annals of Statistics*, **9**, 1187-1195
- Souza R.C. and Neto A.C. (1996) A Bootstrap Simulation Study in ARMA(p,q) Structures, *Journal of Forecasting*, **15**, 343-353
- Swanepoel J.W.H. (1990) A Review of Bootstrap Methods, *South African Statistical Journal*, **24**, 1-34