Journal of Econometrics 188 (2015) 94–110

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Higher-order improvements of the sieve bootstrap for fractionally integrated processes☆

D.S. Poskitt, Simone D. Grose, Gael M. Martin*

Department of Econometrics & Business Statistics, Monash University, Australia

ARTICLE INFO

Article history:
Received 5 June 2012
Received in revised form
15 October 2013
Accepted 31 March 2015
Available online 21 April 2015

JEL classification:
C18
C22
C52

Keywords:
Long memory
ARFIMA
Sieve bootstrap
Bootstrap-based inference
Edgeworth expansion
Sampling distribution

ABSTRACT

This paper investigates the accuracy of bootstrap-based inference in the case of long memory fractionally integrated processes. The re-sampling method is based on the semi-parametric sieve, whereby the dynamics in the process used to produce the bootstrap draws are captured by an autoregressive approximation. Application of the sieve method to data pre-filtered by a semi-parametric estimate of the long memory parameter is also explored. Higher-order improvements yielded by both forms of re-sampling are demonstrated using Edgeworth expansions for a broad class of statistics that includes first- and second-order moments, the discrete Fourier transform and regression coefficients. The methods are then applied to the problem of estimating the sampling distributions of the sample mean and of selected sample autocorrelation coefficients, in experimental settings. In the case of the sample mean, the pre-filtered version of the bootstrap is shown to avoid the distinct underestimation of the sampling variance of the mean which the raw sieve method demonstrates in finite samples, higher-order accuracy of the latter notwithstanding. Pre-filtering also produces gains in terms of the accuracy with which the sampling distributions of the sample autocorrelations are reproduced, most notably in the part of the parameter space in which asymptotic normality does not obtain.

1. Introduction

Many empirical time series have been found to exhibit behaviour characteristic of long memory, or long-range dependent, processes, and the class of fractionally integrated ($I(d)$) processes introduced by Granger and Joyeux (1980) and Hosking (1981) is perhaps the most popular model used to describe the features of such processes. $I(d)$ processes can be characterized by the specification

$$y(t) = \sum_{j=0}^{\infty} k(j)\varepsilon(t-j) = \frac{\kappa(z)}{(1-z)^d}\,\varepsilon(t), \tag{1.1}$$

where $\varepsilon(t), t \in \mathbb{Z}$, is a zero mean white noise process with variance σ^2, z is here interpreted as the lag operator ($z^j y(t) = y(t-j)$), and $\kappa(z) = \sum_{j\geq 0} \kappa(j) z^j$. The behaviour of this process naturally depends on the fractional integration parameter d; for instance, if the "non-fractional" component $\kappa(z)$ is the transfer function of a stable, invertible autoregressive moving-average (ARMA) and $|d| < 0.5$, then the coefficients of $k(z)$ are square-summable, $\sum_{j\geq 0} |k(j)|^2 < \infty$, and $y(t)$ is well-defined as the limit in mean square of a covariance-stationary process. More pertinently, for any $d > 0$ the impulse response coefficients of $k(z)$ in the representation (1.1) are not absolutely summable and the autocovariances decline at a hyperbolic rate, $\gamma(\tau) \sim C\tau^{2d-1}$, rather than the exponential rate typical of an ARMA process. For a detailed description of the properties of long memory processes see Beran (1994).

Statistical procedures for analysing fractional processes are discussed in Hosking (1996), and techniques for estimating fractional models have ranged from the likelihood-based methods studied in Fox and Taqqu (1986), Dahlhaus (1989), Sowell (1992) and Beran (1995), to the semi-parametric methods advanced by Geweke and Porter-Hudak (1983) and Robinson (1995a,b), among others. These techniques typically focus on obtaining an accurate estimate of the parameter governing the long-term behaviour of the process, and the asymptotic theory for these estimators is well established. In particular, we have consistency, asymptotic efficiency, and asymptotic normality for the maximum likelihood estimator (MLE), and the semi-parametric estimators

☆ This research has been supported by Australian Research Council Discovery Grants DP0985234 and DP120102344, and Australian Research Council Future Fellowship FT0991045. The authors would like to thank the Editor and two referees for very detailed and constructive comments on earlier drafts of the paper.

* Correspondence to: Department of Econometrics and Business Statistics, Monash University, Clayton, Victoria 3800, Australia. Tel.: +61 3 9905 1189; fax: +61 3 9905 5474.

E-mail address: gael.martin@monash.edu (G.M. Martin).

http://dx.doi.org/10.1016/j.jeconom.2015.03.045
0304-4076/

are consistent and asymptotically pivotal with particularly simple asymptotic normal distributions.

Concurrent with the development of the asymptotic theory associated with the estimation of long memory models, focus has also been directed at the production of more accurate estimates of finite sample distributions in this setting. An explicit form for the Edgeworth expansion for the sample autocorrelation function of a stationary Gaussian long memory process is derived in Lieberman et al. (2001), and Lieberman et al. (2003) establish the validity of an Edgeworth expansion for the distribution of the MLE of the parameters of such a process, with a zero mean assumed. The unknown mean case is covered in Andrews and Lieberman (2005), with the estimator defined by maximizing the log-likelihood with the unknown mean replaced by the sample mean (referred to as the "plug-in" MLE, or PML). Andrews and Lieberman (2005) also derive results for the Whittle MLE (WML) and for the plug-in version (PWML). Giraitis and Robinson (2003) derive an Edgeworth expansion for the semi-parametric local Whittle estimator of the long memory parameter (Robinson, 1995a) (SPLW), whilst Lieberman and Phillips (2004) derive an explicit form for the first-order expansion for the MLE of the long memory parameter in the fractional noise case.

From the point of view of practical implementation, evaluation of the terms in such expansions, for general long memory models, is no trivial task and typically requires knowledge of the values of population ensemble parameters. These expansions are also usually only valid under more restrictive assumptions than are required for first-order asymptotic approximations; see, for example, Lieberman et al. (2001) and Giraitis and Robinson (2003). Accordingly, much attention has also been given to the application of bootstrap-based inference in these models. Building on the Edgeworth results of Lieberman et al. (2003), Andrews and Lieberman (2005) and Andrews et al. (2006) derive the error rate for the parametric bootstrap for the PML and PWML estimators in Gaussian autoregressive fractionally integrated moving average (ARFIMA) models. In contrast, Poskitt (2008) proposes a semi-parametric approach, based on the sieve bootstrap, and provides both theoretical and simulation-based results regarding the accuracy with which the method estimates the true sampling distribution of suitably continuous linear statistics. To the authors' knowledge Andrews et al. (2006) and Poskitt (2008) are amongst the earliest papers in the literature to have examined the theoretical properties of bootstrap methods in the context of fractionally integrated (long memory) processes.

The current paper builds upon certain developments in Poskitt (2008) and produces new results regarding error rates for sieve-based bootstrap techniques in the context of fractionally integrated processes. Using Edgeworth expansions, it is shown that the procedure we here refer to as the "raw" sieve bootstrap can achieve an error rate of $O_p(T^{-(1-d')+\beta})$ for all $\beta > 0$ where $d' = \max\{0, d\}$, for a class of statistics that includes the sample mean, the sample autocovariance and autocorrelation functions, the discrete Fourier transform and ordinary least squares (OLS) regression coefficients. We also present a new methodology based on a modified form of the sieve bootstrap. The modification uses a consistent semi-parametric estimator of the long memory parameter to pre-filter the raw data, prior to the application of a long autoregressive approximation which acts as the "sieve" from which bootstrap samples are produced. We refer to this as the pre-filtered sieve bootstrap. We establish that, subject to appropriate regularity, for any fractionally integrated processes with $|d| < 0.5$ the error rate of the pre-filtered sieve bootstrap is $O_p(T^{-1+\beta})$ for all $\beta > 0$. These results generalize those of Choi and Hall (2000) who show that, for linear statistics characterized by polynomial products, double sieve bootstrap calibrated percentile methods and sieve bootstrap percentile t confidence intervals evaluated in the short memory case converge at a rate arbitrarily close to that obtained with simple random samples, namely $O_p(T^{-1+\beta})$ for all $\beta > 0$.

Choi and Hall (2000) argue that for short memory processes the sieve bootstrap is to be preferred over the block bootstrap (Künsch, 1989). In particular they note that although the block bootstrap accurately replicates the first-order dependence structure of the original times series it fails to reproduce second-order effects, because these are corrupted by the blocking process. Use of an adjusted variance estimate to correct for the failure to approximate second-order effects results, in turn, in an error rate of only $O_p(T^{-2/3+\beta})$ for the block bootstrap. In contrast, the second-order structure is shown to be preserved by the sieve. Choi and Hall (2000) demonstrate that the performance of the sieve is robust to the selected order for the autoregressive approximation, whilst noting that the choice of block length and other tuning parameters can be crucial to the performance of the block bootstrap. Moreover, as these authors also remark, the use of an automated method such as Akaike's information criterion (AIC) to determine the autoregressive order offers obvious practical advantages, again in contrast with the situation that prevails for the block bootstrap, whereby generic selection rules for the block length are unavailable. These deficiencies identified in the block bootstrap technique are likely to be manifest with long-range dependent data *a-fortiori*, suggesting that the sieve bootstrap is likely to be even more favoured for fractionally integrated processes. For a review of block and sieve bootstrap methods and further discussion of their associated features see Politis (2003).

We illustrate our proposed methods by means of a simulation study, in which we examine the sieve bootstrap approximation to the sampling distribution of two types of statistic that satisfy the relevant conditions for the convergence results to hold. Firstly, we compare and contrast the performance of the raw and the pre-filtered sieve bootstrap in correctly characterizing the known finite sample properties of the sample mean under long memory. In particular, we investigate the previously noted tendency of bootstrap techniques to underestimate the true variance of the sample mean in this setting (Hesterberg, 1997). The pre-filtering is shown to correct for the distinct underestimation of the sampling variance still produced by the raw sieve, the higher-order accuracy of the latter notwithstanding. Secondly, we document the performance of the two bootstrap methods in estimating the (unknown) sampling distributions of selected autocorrelation coefficients. We undertake two exercises here. We begin by comparing the estimates of the sampling distributions produced by the (raw) sieve bootstrap with those produced via an Edgeworth approximation, in the region of the parameter space where such an approximation is valid (see Lieberman et al., 2001). The bootstrap method is shown to produce distributions that are visually indistinguishable from those produced by the second-order Edgeworth expansion which, in turn, replicate the Monte Carlo estimates. Encouraged by the accuracy of the bootstrap method in the case in which an analytical finite sample comparator is available, we then assess the relative performance of the two alternative sieve bootstrap methods – raw and pre-filtered – in the part of the parameter space in which it is not. The pre-filtered method (in particular) is shown to produce particularly accurate estimates of the "true" (Monte Carlo) distributions in this region, auguring well for its general usefulness in empirical settings.

The paper proceeds as follows. Section 2 briefly outlines the statistical properties of autoregressive approximations to fractionally integrated processes, and summarizes the properties of the raw sieve bootstrap in this context. In Section 3 we present relevant Edgeworth expansions for a given class of statistics, and exploit these representations to establish the stated error rates for the raw sieve bootstrap technique. Section 4 outlines the methodology underlying the pre-filtered sieve bootstrap