

COVERAGE BOUND FOR FIXED- b SUBSAMPLING AND GENERALIZED SUBSAMPLING FOR TIME SERIES

Yinxiao Huang and Xiaofeng Shao

Boston University and University of Illinois, Urbana-Champaign

Abstract: We investigate the upper bounds on coverage probabilities of the subsampling-based confidence sets in the time series setting. Under the fixed- b asymptotic framework, where b is the ratio of block size to sample size, we derive the limiting coverage bound, and obtain the finite sample coverage bound by simulations. Our findings suggest that the coverage bound is strictly below 1 for positive b , it can be far away from 1, and the fixed- b subsampling method in Shao and Politis (2013) can exhibit serious undercoverage when the dimension of the parameter is large, the time series dependence is (positively) strong, or b is large. To alleviate the problem, we propose a generalized subsampling method that combines useful features of fixed- b subsampling and self-normalization, and demonstrate its effectiveness in terms of delivering more accurate coverage via numerical studies.

Key words and phrases: Coverage bound, pivot, self-normalization, subsampling.

1. Introduction

Since the seminal work of Politis and Romano (1994), subsampling has become an important and widely applicable tool in various inference problems for time series and dependent data of other types; see Politis, Romano, and Wolf (1999a). The theoretical treatment in most subsampling-related work adopts the traditional small- b asymptotics, where b is the fraction of the subsampling block size (or window width, bandwidth) relative to total sample size, and subsampling can be shown to be consistent under mild conditions. Recently, Shao and Politis (2013) introduced the fixed- b asymptotics (Kiefer and Vogelsang (2005)) into the subsampling-based inference and proposed a p -value based calibration approach to alleviate the inconsistency of subsampling under the fixed- b framework. This new approach allows the effect of block size on the subsampling approximation to be captured to the first order, and is shown to deliver more accurate coverage as compared to its small- b counterpart in simulation studies.

Although the fixed- b subsampling based confidence set improves its small- b counterpart in terms of coverage accuracy, it can be seen from the numerical results of Shao and Politis (2013) that when the dependence of the time series is positively strong and sample size is small, the empirical coverage level can still be

far below the nominal level, resulting in inaccurate inference. This severe undercoverage may be explained by the fact that the theoretical justification in fixed- b subsampling hinges on the Continuous Mapping Theorem and the functional central limit theorem (CLT), and the latter approximation (based on functional CLT) tends to get worse when the time series dependence gets positively stronger and sample size gets smaller. One can quantify the approximation error by developing an Edgeworth expansion for the distribution of the p -value under the fixed- b asymptotics, but that seems very involved; see Zhang and Shao (2013) for a recent attempt in a related context.

We offer a new perspective on the serious undercoverage for fixed- b subsampling based confidence sets, and we discover an intrinsic coverage bound problem associated with the subsampling approach that seems largely unnoticed in the literature. Specifically, there is a nontrivial finite sample (least) upper bound on the coverage probability of the subsampling-based confidence set regardless of its confidence level and this bound is determined by such factors as block size, sample size, dimension of the parameter, the strength and the sign of the time series dependence etc. To gain more insight, we conducted numerical simulations and tabulated the finite sample bounds for a few combinations of sample size, block size, and time series models with varying dependence. We derive the limiting bound as $n \rightarrow \infty$ as a function of b , for both finite and infinite-dimensional parameters.

To alleviate the severe undercoverage associated with the fixed- b subsampling, we propose the generalized subsampling (GS) method as an alternative. The GS still uses values of the statistic computed over blocks of the data, but the blocks can be of different size and a scaling parameter is introduced to allow the finite sample bound and the limiting bound be close to 1. The GS inherits two main ingredients from Shao (2010) and Shao and Politis (2013). It uses the idea of prepivoting by looking at the p -value in Shao and Politis (2013) instead of using a direct studentization, which seems not applicable to the inference of infinite-dimensional parameter. It also uses the recursive subsample estimates, as used in the self-normalized (SN) approach of Shao (2010), which leads to inconsistent subsampling approximation but can be calibrated to yield asymptotically valid inference. We further investigate the coverage bound for GS for both finite and infinite-dimensional parameters. Our numerical results indicate that the bound can be very close to 1 as long as the scaling parameter is in a certain range. We show in simulation studies that the finite sample coverage of the GS approach can be comparable or favorable in some settings in comparison with the SN method and fixed- b subsampling method.

The rest of the paper is organized as follows. In Section 2, we point out the problem of coverage bound for scalar, vector, and infinite-dimensional parameters, and investigate both finite sample and asymptotic bounds by simulation. In

Section 3, we propose a generalized subsampling method and study its coverage bound. In Section 4, we compare GS method to the SN and fixed- b subsampling methods via numerical simulations. Section 5 provides some concluding remarks.

2. Coverage Bounds for Subsampling Based Confidence Sets

We first describe subsampling in the context of a simple inference problem: inference for the mean $\mu = E(X_1)$ of a univariate stationary time series based on the observations $\{X_t\}_{t=1}^n$. Here the subsampling method approximates the sampling distribution of $\sqrt{n}(\bar{X}_n - \mu)$, where $\bar{X}_n = n^{-1} \sum_{t=1}^n X_t$ is the sample mean, with the empirical distribution generated by its subsample counterpart $\sqrt{l}(\bar{X}_{j,j+l-1} - \bar{X}_n)$, where $\bar{X}_{j,j+l-1} = l^{-1} \sum_{i=j}^{j+l-1} X_i$, $j = 1, \dots, N = n - l + 1$. To construct a symmetric two sided confidence interval for μ , consider $\tilde{L}_{n,l}(x) = N^{-1} \sum_{j=1}^N \mathbf{1}(\sqrt{l}|\bar{X}_{j,j+l-1} - \bar{X}_n| \leq x)$, where $\mathbf{1}(A)$ denotes the indicator function of the set A . In the fixed- b framework, the ratio $b = l/n$ is kept constant as $n \rightarrow \infty$.

For a given $\alpha \in [0, 1)$, take the subsampling-based critical values as $\tilde{\pi}_{n,l}(1 - \alpha) = \inf\{x : \tilde{L}_{n,l}(x) \geq 1 - \alpha\}$. Then, under the small- b asymptotic theory, the $100(1 - \alpha)\%$ symmetric confidence interval for μ is

$$\left(\bar{X}_n - n^{-1/2}\tilde{\pi}_{n,l}(1 - \alpha), \bar{X}_n + n^{-1/2}\tilde{\pi}_{n,l}(1 - \alpha)\right). \tag{2.1}$$

In the context of hypothesis testing, if the alternative hypothesis is $H_1 : \mu \neq \mu_0$, then we reject the null hypothesis at the significance level α if p -value is less than or equal to α , where the p -value is

$$\widetilde{pval}_{n,l}^{Fb} = \widetilde{pval}_{n,l}^{Fb}(\mu_0) = \frac{1}{N} \sum_{j=1}^N \mathbf{1}\left(\sqrt{n}|\bar{X}_n - \mu_0| \leq \sqrt{l}|\bar{X}_{j,j+l-1} - \bar{X}_n|\right). \tag{2.2}$$

We often omit the dependence of $\widetilde{pval}_{n,l}^{Fb}$ on μ_0 for notational simplicity. By duality of confidence interval and hypothesis testing, the confidence interval is

$$\left\{\mu : \widetilde{pval}_{n,l}^{Fb}(\mu) \text{ in (2.2)} > \alpha\right\}.$$

The above subsampling-based interval is constructed implicitly assuming the limiting null distribution of p -value is $U(0, 1)$, which is no longer true under the fixed- b framework, see Lahiri (2001). As a remedy, Shao and Politis (2013) calibrated the nominal coverage level on the basis of the pivoting limiting null distribution of the p -value under the fixed- b asymptotics, and obtained more accurate confidence intervals by taking into account the influence of b . Specifically, under the fixed- b asymptotics, the limiting null distribution of $\widetilde{pval}_{n,l}^{Fb}$ is the distribution of \tilde{G}_b , where

$$\tilde{G}_b = (1-b)^{-1} \int_0^{1-b} \mathbf{1} \left(|W(1)| \leq \frac{|W(b+t) - W(t) - bW(1)|}{\sqrt{b}} \right) dt,$$

and $W(t)$ is a standard Brownian motion. If $\tilde{G}_b(\alpha)$ denotes the $100\alpha\%$ quantile of the distribution of \tilde{G}_b , the p -value calibrated fixed- b based $100(1-\alpha)\%$ symmetric confidence interval is $\{\mu : \widetilde{pval}_{n,l}^{Fb}$ in (2.2) $> \tilde{G}_b(\alpha)\}$, namely,

$$\left(\bar{X}_n - n^{-1/2} \tilde{\pi}_{n,l} \left(1 - \tilde{G}_b(\alpha) \right), \bar{X}_n + n^{-1/2} \tilde{\pi}_{n,l} \left(1 - \tilde{G}_b(\alpha) \right) \right). \quad (2.3)$$

In Shao and Politis (2013), the value $\tilde{G}_b(\alpha)$ was obtained by Monte Carlo simulation.

Take

$$\beta_n(b) := P \left(\max_{j=1, \dots, N} |\sqrt{l}(\bar{X}_{j,j+l-1} - \bar{X}_n)| < \sqrt{n} |\bar{X}_n - \mu_0| \right), \quad (2.4)$$

where μ_0 is the true value of μ , with its limit

$$\beta(b) := P \left(\sup_{t \in [0, 1-b]} \frac{|(W(b+t) - W(t) - bW(1))|}{\sqrt{b}} < |W(1)| \right),$$

where $W(t)$ is a standard Brownian motion. Then the coverage probability of (2.1) and (2.3) have a nontrivial upper bound,

$$\begin{aligned} P(\mu_0 \in CI) &\leq P \left(\sqrt{n} |\bar{X}_n - \mu_0| \leq \max_{j=1, \dots, N} |\sqrt{l}(\bar{X}_{j,j+l-1} - \bar{X}_n)| \right) \\ &= 1 - P \left(\sqrt{n} |\bar{X}_n - \mu_0| > \max_{j=1, \dots, N} |\sqrt{l}(\bar{X}_{j,j+l-1} - \bar{X}_n)| \right) \\ &= 1 - \beta_n(b) \rightarrow 1 - \beta(b), \end{aligned} \quad (2.5)$$

where $b = l/n$. We call $1 - \beta_n(b)$ the finite sample coverage bound and $1 - \beta(b)$ the limiting bound. If $1 - \beta_n(b) < 1 - \alpha$, then the confidence set is bound to undercover and the amount of undercoverage gets more severe as $\beta_n(b)$ gets farther from zero. For large n and under the small b asymptotics, this problem does not occur because $\beta_n(b) \approx \beta(0) = 0$. But under the fixed- b asymptotics or for small sample size, the bound is strictly less than 1 in both finite sample or limit. This seems the first time such a phenomenon has been brought up for discussion in the literature.

Since $\beta(b) := P(\tilde{G}_b = 0)$, two cases can occur. If $\beta(b) > \alpha$, then $\tilde{G}_b(\alpha) = 0$, and the inequality in (2.5) becomes an equality. In this case, it is impossible to construct a confidence interval with asymptotically correct coverage. If $\beta(b) \leq \alpha$, then a CI of asymptotically valid coverage can be constructed, but whether the finite sample coverage bound reaches the target level is unknown for a given sample size. The quantity $\beta_n(b)$ depends on the joint distribution of time series,

the form of the parameter, block size, and sample size, so is in general difficult to calculate. Table 1 provides numerical values of $\beta_n(b)$ and $\beta(b)$ for a few combinations of b 's and n 's under different levels of dependence strength. See Section 3 for more discussions.

2.1. Finite dimensional parameter

The issue of coverage bound also exists when we deal with general finite dimensional parameters. Following Politis, Romano, and Wolf (1999b), we assume that the parameter of interest is $\theta(P) \in \mathbf{R}^d$, where P is the joint probability law that governs the p -dimensional stationary sequence $\{X_t\}_{t \in \mathbf{Z}}$. Let $\hat{\theta}_n = \hat{\theta}_n(X_1, \dots, X_n)$ be an estimator of $\theta = \theta(P)$ based on the observations (X_1, \dots, X_n) . We define the subsampling estimator of $\theta(P)$ by $\hat{\theta}_{j,j+l-1} = \hat{\theta}_l(X_j, \dots, X_{j+l-1})$ on the basis of the subsample (X_j, \dots, X_{j+l-1}) , $j = 1, \dots, N$. Let $\|\cdot\|$ be a norm in \mathbf{R}^d . The subsampling-based distribution estimator of $\|\sqrt{n}(\hat{\theta}_n - \theta)\|$ is denoted by $\tilde{L}_{n,l}(x) = N^{-1} \sum_{j=1}^N \mathbf{1}(\|\sqrt{l}(\hat{\theta}_{j,j+l-1} - \hat{\theta}_n)\| \leq x)$. In the testing context (say $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$), we define the subsampling based p -value as

$$\widetilde{pval}_{n,l}^{Fb} = N^{-1} \sum_{j=1}^N \mathbf{1}(\|\sqrt{n}(\hat{\theta}_n - \theta)\| \leq \|\sqrt{l}(\hat{\theta}_{j,j+l-1} - \hat{\theta}_n)\|); \tag{2.6}$$

here we do not distinguish θ and θ_0 for convenience, they are the same under the null.

Suppose that $\theta(P) = T(F)$, where F is the marginal distribution of $X_1 \in \mathbf{R}^p$, and T is a functional that takes value in \mathbf{R}^d . Then a natural estimator of $T(F)$ is $\hat{\theta}_n = T(\rho_{1,n})$, where $\rho_{1,n} = n^{-1} \sum_{t=1}^n \delta_{X_t}$ is the empirical distribution and δ_x stands for the point mass at x . Similarly, $\hat{\theta}_{j,j+l-1} = T(\rho_{j,j+l-1})$, where $\rho_{j,j+l-1} = l^{-1} \sum_{h=j}^{j+l-1} \delta_{X_h}$. Under some regularity conditions on T and moment and weak dependence assumptions on the time series, Theorem 1 in Shao and Politis (2013) showed that the limiting null distribution of $\widetilde{pval}_{n,l}^{Fb}$ is the distribution of $\tilde{G}_{b,d}$, where

$$\tilde{G}_{b,d} = \frac{1}{1-b} \int_0^{1-b} \mathbf{1}\left(\|\Sigma^{1/2}W_d(1)\| \leq \frac{\|\Sigma^{1/2}(W_d(b+r) - W_d(r) - bW_d(1))\|}{\sqrt{b}}\right) dr,$$

with $W_d(\cdot)$ representing the d -dimensional vector of independent Brownian motions and $\Sigma = \Sigma(P)$ denoting the long run variance matrix corresponding to $T(F)$. Specifically,

$$\Sigma = \sum_{j=-\infty}^{\infty} \text{cov}(IF(X_0, P), IF(X_j, P)), \tag{2.7}$$

where $IF(X; P)$ stands for the influence function corresponding to T ; see Section 3.1 of Shao and Politis (2013) for the definition.

In the special case $d = 1$, $\tilde{G}_{b,1} = \tilde{G}_b$, and is pivotal. However, for $d \geq 2$, $\tilde{G}_{b,d}$ is no longer pivotal; it critically depends on the unknown covariance matrix Σ . One way out is to approximate the limiting null distribution $\tilde{G}_{b,d}$ further by subsampling, see Shao and Politis (2013). Since this procedure mimics the idea of a double-bootstrap, it was termed double subsampling procedure in the latter paper.

Let $\tilde{G}_{b,d}(\alpha)$ be the α th-quantile of $\tilde{G}_{b,d}$ that can be consistently approximated by subsampling. Then the calibrated $100(1 - \alpha)\%$ subsampling-based confidence region for θ contains all points θ_0 for which the test of $H_0 : \theta = \theta_0$ fails to reject the null hypothesis,

$$\left\{ \theta \in \mathbf{R}^d : \widetilde{pval}_{n,l}^{Fb} \text{ in (2.6)} > \tilde{G}_{b,d}(\alpha) \right\}, \tag{2.8}$$

whereas the traditional subsampling-based confidence region is $\{ \theta \in \mathbf{R}^d : \widetilde{pval}_{n,l}^{Fb} \text{ in (2.6)} > \alpha \}$. Define

$$\beta_n(b; d; \Sigma) = P \left(\max_{j=1, \dots, N} \|\sqrt{l}(\hat{\theta}_{j,j+l-1} - \hat{\theta}_n)\| < \sqrt{n}\|\hat{\theta}_n - \theta\| \right). \tag{2.9}$$

Let $F(x-) := \lim_{y \uparrow x} F(y)$. The coverage probability of (2.8) is

$$\begin{aligned} P(\theta \in (2.8)) &= P \left(\left\{ \theta \in \mathbf{R}^d : 1 - \tilde{L}_{n,l}(\sqrt{n}\|\hat{\theta}_n - \theta\|^-) > \tilde{G}_{b,d}(\alpha) \right\} \right) \\ &= P \left(\left\{ \theta \in \mathbf{R}^d : \tilde{L}_{n,l}(\sqrt{n}\|\hat{\theta}_n - \theta\|^-) < 1 - \tilde{G}_{b,d}(\alpha) \right\} \right) \\ &= P \left(\left\{ \theta \in \mathbf{R}^d : \sqrt{n}\|\hat{\theta}_n - \theta\| \leq (1 - \tilde{G}_{b,d}(\alpha))\text{-th quantile of } \tilde{L}_{n,l} \right\} \right) \\ &\leq P \left(\sqrt{n}\|\hat{\theta}_n - \theta\| \leq \max_{j=1, \dots, N} \|\sqrt{l}(\hat{\theta}_{j,j+l-1} - \hat{\theta}_n)\| \right) \\ &= 1 - \beta_n(b; d; \Sigma). \end{aligned} \tag{2.10}$$

Let $\beta(b; d; \Sigma)$ be the limit of $\beta_n(b; d; \Sigma)$ as $n \rightarrow \infty$. If $d = 1$, $\beta(b; d; \Sigma) = \beta(b)$ that does not depend on the nuisance parameter Σ . For $d \geq 2$, let

$$\beta(b; d; \Sigma) = P \left(\sup_{r \in [0, 1-b]} \frac{\|\Sigma^{1/2}(W_d(b+r) - W_d(r) - bW_d(1))\|}{\sqrt{b}} < \|\Sigma^{1/2}W_d(1)\| \right).$$

If $\beta(b; d; \Sigma) \geq \alpha$, then $P(\tilde{G}_{b,d} = 0) \geq \alpha$, $\tilde{G}_{b,d}(\alpha) = 0$, and the inequality in (2.10) becomes equality. Again we run into the issue of undercoverage if $\beta_n(b; d; \Sigma) > \alpha$.

2.2. Infinite-dimensional parameter

In the time series setting, subsampling methods have been used to provide an approximation of the nonpivotal limiting distribution when the parameter of interest is of infinite dimension, such as the marginal distribution function and the spectral distribution function of a stationary time series; see Politis, Romano, and Wolf (1999b). We use $\|F - G\|_\infty$ to denote $\sup_{x \in \mathbf{R}} |F(x) - G(x)|$ and focus on the confidence band construction for the marginal distribution function.

Consider a stationary sequence $\{X_t\}_{t \in \mathbf{Z}}$ and let $m(s) = P(X_0 \leq s)$ be its marginal cumulative distribution function (cdf). Given the observations $\{X_t\}_{t=1}^n$, the empirical distribution function is $m_n(s) = n^{-1} \sum_{t=1}^n \mathbf{1}(X_t \leq s)$. To construct a confidence band for $m(\cdot)$, it is known from Berkes, Hörmann, and Schauer (2009) that $\sqrt{n}(m_n(s) - m(s)) \Rightarrow K(s, 1)$, where $\{K(s, t), (s, r) \in [-\infty, \infty] \times [0, 1]\}$ is a two-parameter mean zero Gaussian process with

$$\text{cov}(K(s, r), K(s', r')) = (r \wedge r')\Gamma(s, s'), \tag{2.11}$$

and $\Gamma(s, s') = \sum_{k=-\infty}^\infty \text{cov}(\mathbf{1}(X_0 \leq s), \mathbf{1}(X_k \leq s'))$. Then by the Continuous Mapping Theorem, $\sqrt{n}\|m_n - m\|_\infty \rightarrow_D \sup_{s \in \mathbf{R}} |K(s, 1)|$, whose distribution is not pivotal since the covariance kernel $\Gamma(\cdot, \cdot)$ depends on unknown nuisance parameters. To describe the fixed- b subsampling method, let $\sqrt{l}(m_{t,t+l-1}(s) - m_n(s))$, $t = 1, \dots, N = n - l + 1$, be the subsampling counterpart of $\sqrt{n}(m_n(s) - m(s))$, where $m_{t,t+l-1}(s) = l^{-1} \sum_{h=t}^{t+l-1} \mathbf{1}(X_h \leq s)$. Define the p -value

$$\widetilde{pval}_{n,l}^{Fb} = N^{-1} \sum_{t=1}^N \mathbf{1}\left(\sqrt{l}\|m_{t,t+l-1} - m_n\|_\infty \geq \sqrt{n}\|m_n - m\|_\infty\right). \tag{2.12}$$

Let $b = l/n$. Under fixed- b asymptotics, the limiting null distribution of the p -value is the distribution of \mathcal{G}_b , where

$$\mathcal{G}_b := \frac{1}{1-b} \int_0^{1-b} \mathbf{1}\left(\sup_{s \in \mathbf{R}} \frac{|K(s, r+b) - K(s, r) - bK(s, 1)|}{\sqrt{b}} \geq \sup_{s \in \mathbf{R}} |K(s, 1)|\right) dr.$$

The distribution of \mathcal{G}_b is not pivotal for a given b , because it depends on the Gaussian process $K(s, t)$, whose covariance structure is tied to the unknown dependence structure of X_t . So subsampling at the first stage is insufficient under the fixed- b asymptotic framework.

To make the inference feasible, a double-sampling procedure was employed in Shao and Politis (2013) to approximate the sampling distribution of the p -value or its limiting null distribution; see also Section 3.1. Let $\mathcal{G}_b(\alpha)$ be the α -th quantile of \mathcal{G}_b and $\widehat{\mathcal{G}}_b(\alpha)$ a consistent estimator by subsampling in the second stage. For a given $\alpha \in (0, 1)$, the $100(1 - \alpha)\%$ calibrated subsampling-based confidence band for $m(\cdot)$ is

$$\left\{ m : m \text{ is a distribution function and } \widetilde{pval}_{n,l}^{Fb} \text{ in (2.12)} > \widehat{\mathcal{G}}_b(\alpha) \right\}. \tag{2.13}$$

Take

$$\beta_n(b, K) := P\left(\max_{j=1, \dots, N} \sqrt{l} \|m_{j, j+l-1} - m_n\|_\infty < \sqrt{n} \|m_n - m\|_\infty\right). \quad (2.14)$$

Following the argument in Section 2.1, The coverage bound of (2.13) can be derived as

$$P\left(\max_{j=1, \dots, N} \sqrt{l} \|m_{j, j+l-1} - m_n\|_\infty \geq \sqrt{n} \|m_n - m\|_\infty\right) = 1 - \beta_n(b, K),$$

with the limit of $\beta_n(b, K)$ being

$$\beta(b, K) := P\left(\sup_{r \in (0, 1-b)} \sup_{s \in \mathbf{R}} \frac{|K(s, r+b) - K(s, r) - bK(s, 1)|}{\sqrt{b}} < \sup_{s \in \mathbf{R}} |K(s, 1)|\right).$$

2.3. Finite sample coverage bound for fixed- b subsampling

To investigate the severity of the coverage bound issue, we present the finite sample coverage bounds of the subsampling based confidence set for time series data through numerical simulations. Suppose X_t is generated from a vector autoregression (VAR) model $X_t = \rho I_d X_{t-1} + \epsilon_t$ with $\epsilon_t \in \mathbf{R}^d$, where I_d is a $d \times d$ identity matrix. Assuming that $\epsilon_t \sim i.i.d. N(0, I_d)$, there is no cross-sectional correlation for the time series, we present the values of coverage bounds $\beta_n(b; d; \Sigma_d)$ in Table 1 with different choices of b and sample size n for $d = 1, 2, 3$. The sample size n is 50, 100, 500, 3,000, ∞ . We approximated the asymptotic coverage bound $\beta(b; d; \Sigma_d)$ by simulating independent Wiener processes 5,000 times, where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables.

Tables 1 and 2 summarize the coverage bound for the mean and median respectively by fixed- b subsampling. The coverage bound decreases as the positive dependence strengthens, whereas for negative ρ the coverage bound tends to inflate as compared to the i.i.d. case. In general there is a decreasing trend in the coverage bound as the dimension d increases. The bound much depends on b and is moderately sensitive to b especially in the case of high positive correlation. Although no universal pattern can be found to dictate the selection of b based on the coverage bound, the simulation results do suggest that b in the range of $(0, 0.2)$ tends to deliver higher coverage probability.

Tables 1 and 2 reveal that even after p -value calibration, an asymptotically valid confidence set can have a potential undercoverage issue in finite samples. In that situation, if a high confidence level is desired, it is recommended that the bound be computed to see if such a high confidence level is attainable. From Shao and Politis (2013), “undercoverage occurs (for fixed- b subsampling methods) and it becomes more severe as the dependence positively strengthens”, see Figures 1-3 therein. This phenomenon can be partially explained and is echoed by the coverage bound presented in Table 1.

Table 1. Bounds on the coverage probability $1 - \beta_n(b; d; I_d)$ for the mean by fixed- b subsampling. The number of replications was 5,000 each. The values for $\beta(b; d; I_d)$ in the first column were based on simulating 5,000 independent Wiener processes for $d = 1$ ($d = 2, 3$, respectively), where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables.

$b \setminus n$	$\rho = 0$				$\rho = 0.5$				$\rho = 0.8$				$\rho = -0.5$				
	∞	50	100	500	3,000	50	100	500	3,000	50	100	500	3,000	50	100	500	3,000
$d = 1$																	
0.02	99.56	97.62	98.84	99.66	99.67	82.67	92.78	98.88	99.39	54.00	70.04	96.04	99.31	99.94	99.86	99.72	99.86
0.04	99.36	97.20	98.20	99.22	99.27	87.52	94.86	98.34	98.83	64.58	78.96	97.02	98.88	99.32	99.50	99.44	99.55
0.06	98.93	96.56	97.76	98.68	98.83	89.45	95.02	97.70	98.28	69.70	82.46	96.62	98.56	99.10	98.92	98.88	99.20
0.08	98.43	95.80	97.00	98.22	98.45	89.43	94.74	97.08	97.66	72.30	84.06	95.96	97.95	98.34	98.48	98.58	98.66
0.1	97.82	95.20	96.24	97.62	97.85	89.27	94.06	96.12	96.94	73.74	84.52	95.24	97.18	97.94	97.78	97.92	98.08
0.2	94.68	90.66	91.90	93.80	94.41	84.81	89.68	92.40	93.29	73.28	81.14	90.18	93.41	94.24	94.64	94.14	94.64
0.3	90.46	85.90	87.34	89.08	90.05	79.68	84.66	87.10	88.55	68.92	75.18	84.36	88.75	90.02	90.24	90.36	90.33
0.4	85.77	80.56	81.92	84.34	85.27	73.52	79.08	82.02	83.67	63.22	69.90	79.58	84.02	85.26	85.92	85.56	85.66
0.5	79.71	74.58	76.00	79.30	80.17	67.74	72.74	76.66	78.46	57.22	63.84	73.12	79.24	80.66	81.06	80.58	80.64
$d = 2$																	
0.02	99.81	97.76	98.44	99.26	99.76	73.12	89.32	98.76	99.53	33.36	56.36	94.68	99.09	99.99	99.90	99.86	99.76
0.04	99.64	97.30	98.40	99.10	99.52	82.28	93.12	98.34	99.15	48.00	70.70	96.28	98.87	99.54	99.58	99.58	99.55
0.06	99.26	96.82	97.00	98.56	99.00	85.84	94.02	97.94	98.73	56.14	76.74	95.84	97.95	99.34	99.36	99.24	99.00
0.08	98.53	96.24	96.94	97.84	98.32	86.62	93.80	97.30	98.02	61.26	79.62	94.92	97.60	98.98	98.66	98.72	98.47
0.1	98.04	95.10	96.10	97.26	97.82	86.94	92.94	96.38	97.65	64.16	80.50	94.38	96.89	98.28	98.26	98.26	98.07
0.2	94.25	90.24	90.96	93.42	93.54	82.52	87.62	91.70	93.29	66.28	77.54	88.92	92.25	94.88	94.40	94.58	94.25
0.3	89.22	83.48	84.50	87.70	88.34	75.58	80.18	85.66	87.47	62.12	70.90	82.48	85.93	89.16	88.68	89.52	88.62
0.4	82.10	76.46	78.28	80.28	81.48	67.52	73.06	77.84	81.24	55.00	62.36	74.58	78.80	83.02	82.24	82.32	82.40
0.5	75.17	67.64	69.68	73.38	73.98	59.00	64.38	70.30	73.13	46.76	54.04	66.78	71.33	76.68	75.66	75.42	75.20
$d = 3$																	
0.02	99.58	98.08	98.82	99.60	99.70	66.74	86.78	98.70	99.62	19.98	43.76	92.72	99.20	99.99	99.86	99.86	99.80
0.04	99.28	97.62	98.58	99.06	99.52	79.64	92.38	98.28	99.22	35.48	62.90	95.12	98.69	99.52	99.48	99.60	99.51
0.06	98.81	96.94	97.92	98.54	98.91	84.14	93.10	97.50	98.73	45.00	71.50	95.22	98.16	99.48	98.92	99.28	99.15
0.08	98.37	96.18	97.32	98.04	98.38	85.72	92.80	96.92	98.15	51.52	75.76	94.58	97.55	98.64	98.46	98.80	98.65
0.1	97.87	95.48	95.72	97.34	97.66	85.92	92.14	96.08	97.60	56.02	77.46	93.52	96.65	98.24	97.70	98.16	98.07
0.2	93.28	89.58	90.84	92.56	93.36	81.32	87.02	90.96	92.87	61.16	75.40	87.40	91.96	93.80	93.18	94.42	92.93
0.3	87.38	81.40	85.06	85.52	86.94	73.58	78.68	84.38	85.85	56.30	67.56	79.88	85.29	87.62	87.22	88.04	86.36
0.4	78.76	73.60	74.90	77.18	78.69	64.46	69.64	75.92	77.04	47.74	57.52	70.62	75.85	79.34	79.38	80.28	78.45
0.5	68.79	63.02	65.66	67.68	69.56	54.76	59.86	66.30	68.33	38.36	48.26	61.34	66.44	71.18	70.20	71.24	68.80

Table 2. Bounds on the coverage probability $1 - \beta_n(b; d; I_d)$ for the median by fixed- b subsampling. The number of replications was 5,000 each. The values for $\beta(b; d; I_d)$ in the first column were based on simulating 5,000 independent Wiener processes, where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables.

$b \setminus n$	∞	$\rho = 0$			$\rho = 0.5$			$\rho = 0.8$			$\rho = -0.5$						
		50	100	500	3,000	50	100	500	3,000	50	100	500	3,000				
$d=1$																	
0.02	99.77	95.02	96.70	99.48	99.61	79.16	90.36	98.44	99.26	52.82	68.92	95.30	99.33	99.12	96.42	99.62	99.73
0.04	99.40	94.48	97.28	99.30	99.29	85.60	93.40	98.26	98.97	64.00	78.78	96.34	98.82	94.32	98.06	99.20	99.44
0.06	99.00	95.68	96.64	98.80	98.89	88.79	94.04	97.64	98.54	70.12	82.64	96.18	98.41	99.48	98.28	98.80	98.94
0.08	98.26	94.44	96.44	98.36	98.41	88.95	94.06	97.18	97.92	72.88	83.94	95.86	97.75	96.22	98.18	98.34	98.50
0.1	97.82	94.90	96.04	97.72	97.84	89.59	93.64	96.66	97.18	74.40	84.80	95.20	97.32	98.50	97.94	97.76	97.78
0.2	94.61	90.70	92.54	94.40	94.65	85.56	89.78	92.76	93.64	75.18	81.92	90.28	93.55	93.76	94.36	94.32	94.61
0.3	90.09	86.82	87.74	89.58	89.96	81.24	84.92	88.22	89.11	71.40	76.54	85.68	89.01	90.18	89.70	90.24	90.33
0.4	85.87	80.16	82.38	84.26	85.53	74.94	79.22	82.56	83.91	64.80	70.42	80.62	84.18	83.48	84.96	84.90	85.29
0.5	81.79	74.24	76.36	78.70	80.03	68.41	73.18	76.92	78.51	58.14	65.14	74.14	79.14	77.12	78.12	79.70	80.20
$d=2$																	
0.02	99.83	93.56	95.26	99.08	99.74	69.46	85.62	98.54	99.58	32.22	54.40	93.88	99.25	99.18	94.54	99.58	99.82
0.04	99.66	92.94	96.84	98.78	99.52	79.70	91.62	98.42	99.38	47.36	70.60	96.08	98.65	92.36	97.68	99.46	99.62
0.06	99.38	96.58	95.96	98.82	99.16	85.68	93.08	97.90	98.87	57.52	76.42	95.98	98.33	99.62	97.64	99.22	99.16
0.08	99.05	94.50	95.82	97.96	98.26	85.70	93.22	97.56	98.47	62.48	79.48	95.50	97.91	96.18	97.84	98.80	98.75
0.1	98.71	94.96	96.08	97.58	97.82	87.92	92.64	97.08	97.91	66.82	81.02	94.56	97.29	99.10	97.64	98.46	98.02
0.2	94.74	90.84	92.50	94.22	94.92	83.78	88.64	92.50	93.58	69.88	79.62	89.98	93.13	94.60	94.62	94.92	94.60
0.3	90.15	87.80	87.50	88.38	89.40	79.44	83.06	87.14	88.40	66.00	73.44	83.38	86.80	91.84	90.24	90.16	89.71
0.4	82.65	79.26	81.58	82.24	82.86	72.18	76.92	80.28	82.65	59.26	66.58	76.14	79.80	83.66	84.08	83.50	83.44
0.5	76.61	73.22	73.36	75.34	75.32	65.50	68.20	72.90	75.58	51.46	59.02	68.38	73.04	77.42	76.32	76.46	76.55
$d=3$																	
0.02	99.72	92.60	94.68	99.02	99.82	61.46	81.14	98.00	99.69	19.64	41.58	92.38	99.31	99.00	93.00	99.66	99.80
0.04	99.54	92.26	96.40	99.04	99.48	74.94	90.82	98.00	99.29	34.02	62.82	95.46	99.05	90.42	97.00	99.46	99.44
0.06	99.13	96.00	95.94	98.64	98.54	84.28	92.64	97.90	98.73	46.82	72.18	95.36	98.49	99.58	97.54	99.16	99.15
0.08	98.65	93.50	96.54	98.12	98.58	84.40	92.64	97.46	98.44	52.42	72.22	95.02	97.78	94.84	97.80	99.12	98.58
0.1	98.01	96.08	95.74	97.90	98.12	87.22	92.64	96.90	97.91	58.50	78.94	93.84	96.84	99.22	97.34	98.52	98.18
0.2	92.85	90.50	92.64	94.52	94.40	83.32	88.82	92.46	94.09	64.50	78.84	88.80	92.20	93.66	94.42	94.94	94.16
0.3	86.66	86.04	87.30	87.78	87.92	78.82	82.30	86.60	87.47	62.12	71.48	81.54	85.82	91.46	89.50	89.96	88.18
0.4	78.09	78.02	79.82	80.72	80.74	69.94	74.68	78.46	79.22	53.58	63.58	72.78	77.65	83.50	82.84	82.64	81.25
0.5	69.75	71.86	70.76	71.80	71.20	62.56	66.16	70.04	70.13	45.78	54.18	63.44	68.35	77.76	74.20	74.02	72.38




Figure 1. Bounds on the coverage probabilities for fixed- b subsampling for the mean at nominal level 95% (dashed line). The data were generated from multivariate standard normal distribution with $n = 5,000$ and the number of Monte Carlo replications was 5,000.

We investigated the impact of the dimensionality d of the parameter on the limiting coverage bound. Figure 1 presents the coverage bound for vector mean at $d = 2, 3, 5, 10, 15, 20, 50$ for a range of b 's in $(0, 0.5)$, where the data were generated from i.i.d. multivariate standard normal distribution with sample size $n = 5,000$. As the dimension d increases, even for i.i.d. Gaussian data, the finite sample coverage bound deviates from 1. The upper bound can be close to zero when both b and d are large. We expect a larger deviation from 1 on the coverage bound if the data exhibit positive dependence, or if the sample size is small, as seen from Tables 1 and 2. These findings call for special caution from the practitioners when dealing with a confidence set in one of the following situations if the nominal level is close to 1, say 99%, the dimension of the parameter d is large, the (positive) dependence is strong, b is large, or the sample size n is small.

For the infinite-dimensional case, we examine the confidence band for the marginal distribution function of an AR(1) model: $X_t = \rho X_{t-1} + \epsilon_t$, where $\epsilon_t \sim i.i.d. N(0, 1 - \rho^2)$; the theoretical marginal distribution of X_t is standard normal. Table 3 shows the coverage bound for the confidence band at different levels of dependence. Not surprisingly, the coverage bound is significantly smaller than 1 if sample size is small ($n = 50$) and the dependence is strong ($\rho = 0.8$).

Table 3. Bounds on the coverage probability $1 - \beta_n(b)$ for the marginal cdf by fixed- b subsampling. The number of replications was 5,000 each. The values for $\beta(b)$ in the first column were based on simulating 5,000 independent Wiener processes, where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables. The theoretical marginal distribution in each case is $N(0, 1)$.

$b \setminus n$	$\rho = 0$			$\rho = 0.5$			$\rho = 0.8$			$\rho = -0.5$				
	∞	50	100	50	100	500	50	100	500	3,000	50	100	500	3,000
$d = \infty$														
0.02	99.76	73.66	94.14	99.30	99.62	48.96	78.72	98.48	99.50	22.46	50.56	93.60	99.13	85.92
0.04	99.23	92.84	96.50	99.06	99.14	77.28	92.54	98.20	99.14	47.20	73.70	96.12	98.82	93.94
0.06	98.80	94.02	96.20	98.34	98.62	86.20	93.72	97.42	98.52	59.60	81.84	96.14	98.12	93.94
0.08	98.20	94.02	95.46	97.28	98.24	88.50	93.50	96.82	97.90	67.84	84.96	95.24	97.76	92.92
0.1	97.88	92.92	95.02	96.90	97.30	89.02	92.82	96.10	97.30	71.60	86.36	94.54	96.79	93.26
0.2	93.46	86.80	89.24	92.38	92.78	84.20	87.30	91.24	93.02	73.88	82.86	89.24	92.34	86.08
0.3	87.44	77.64	81.42	85.04	85.64	75.52	80.00	84.88	86.80	67.88	75.78	83.24	86.77	76.78
0.4	77.65	66.32	69.72	75.44	75.84	66.16	71.20	77.12	79.54	59.36	67.50	75.60	80.09	63.82
0.5	67.25	52.46	56.80	63.44	64.76	54.88	60.40	67.46	70.74	48.54	56.90	66.64	71.37	48.08

Table 4. Bounds on the coverage probability $1 - \beta_n^{GS}(g)$ for the marginal cdf by GS. The number of replications was 10,000 each. The values for $\beta^{GS}(g)$ in the first column are based on simulating 5,000 independent Wiener processes, where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables. The theoretical marginal distribution in each case is $N(0, 1)$.

$g \setminus n$	$\rho = 0$			$\rho = 0.5$			$\rho = 0.8$			$\rho = -0.5$				
	∞	50	100	50	100	500	50	100	500	3,000	50	100	500	3,000
$d = \infty$														
1	19.71	13.91	15.60	18.64	19.53	22.27	25.10	28.44	30.93	23.99	28.93	35.56	39.75	8.89
2.5	95.97	93.25	93.83	95.23	95.45	89.74	91.74	93.42	94.27	83.51	88.06	92.12	93.71	94.33
3	98.89	97.79	98.19	98.75	98.77	95.39	96.67	97.43	98.07	90.90	94.13	96.56	97.48	98.70
3.5	99.75	99.23	99.55	99.66	99.73	98.26	98.77	99.02	99.42	95.01	97.06	98.53	99.06	99.66
4	99.94	99.76	99.87	99.91	99.94	99.28	99.48	99.67	99.76	97.28	98.53	99.33	99.59	99.92
5	100.00	99.99	99.99	100.00	100.00	99.90	99.95	99.94	99.98	99.34	99.56	99.85	99.92	99.99

The choice of b plays an important role on the coverage bound with larger b corresponding to lower bound in general, but we notice that for some very small b 's, the bound can be far below 1. The negative correlation $\rho = -0.5$ leads to a different pattern for an infinite-dimensional parameter. Specifically, except for $d = 0.02$, the coverage bound for $\rho = -0.5$ is either comparable to, or smaller than that for $\rho = 0$, indicating that a potential undercoverage can be caused by not only positive, but also negative value of ρ .

3. Generalized Subsampling

In the time series setting, for given observations from a time series $\{X_i\}_{i=1}^n$, a statistic is recomputed over the $n-l+1$ subsets of size l of the form $\{X_i, X_{i+1}, \dots, X_{i+l-1}\}$. The traditional way of performing subsampling in this context has two important features: (i) blocks of the consecutive observations are used and blocks are of the same length to retain the time series dependence; (ii) the block size satisfies $l/n + 1/l = o(1)$ as $n \rightarrow \infty$ to show the consistency of the subsampling based approximation in distribution approximation and in variance estimation; see Politis and Romano (1994).

We propose a new GS method to relax these requirements for the use of the subsampling method. It still uses values of the statistic computed over blocks of the data (subsamples), but the blocks can be of different size and a scaling parameter g is introduced to alleviate the coverage bound problem. The applicability of the GS is limited to approximately linear statistics that are asymptotically normal, so its scope of applicability is narrower than that for subsampling. On the other hand, as we demonstrate later, its coverage bound can be fairly close to 1 and can outperform fixed- b subsampling in terms of coverage accuracy, especially when the coverage bound for fixed- b subsampling is substantially below 1. It is closely related to the SN approach (Shao (2010)), where the use of an inconsistent normalizer leads to a nonstandard limiting distribution of the studentized statistic, but it can be used in the inference of infinite-dimensional parameters, to which the SN approach seems not directly applicable. It is also intimately related to the work of Shao and Politis (2013), in which the subsampling is inconsistent under a fixed- b asymptotic framework and a p -value based calibration is proposed to yield asymptotically valid confidence sets. We follow the same p -value based calibration approach and our development for the GS method is parallel to that in Shao and Politis (2013).

3.1. GS methodology details

Let $\hat{\theta}_t = \hat{\theta}_{1,t}$ and $m_t = m_{1,t}$ be the estimate based on subsample (X_1, \dots, X_t) . For a vector parameter $\theta \in \mathbf{R}^d$, we approximate the sampling distribution of $\sqrt{n}\|\hat{\theta}_n - \theta\|$ by

$$M_{n,g}(x) = n^{-1} \sum_{t=1}^n \mathbf{1} \left(gn^{-1/2} \|t(\hat{\theta}_t - \hat{\theta}_n)\| \leq x \right),$$

where g is a scaling parameter introduced to alleviate the coverage bound problem for $g = 1$. We leave the choice of g and its impact on finite sample performance to later sections. Here the use of recursive estimators $\{\hat{\theta}_t\}_{t=1}^n$ is motivated by the good finite sample performance of the self-normalized approach in Shao (2010); see also Nordman, Bunzel, and Lahiri (2013) for a related formulation in the blockwise empirical likelihood context. Similar to the subsampling under the fixed- b framework, the distribution estimator $M_{n,g}(x)$ is inconsistent, but we can use p -value calibration to construct a confidence set with asymptotically correct coverage. In the testing context (say $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$), we take the GS-based p -value as

$$\widetilde{pval}_{n,g}^{GS} = n^{-1} \sum_{t=1}^n \mathbf{1} \left(g \|n^{-1/2} t(\hat{\theta}_t - \hat{\theta}_n)\| \geq \|\sqrt{n}(\hat{\theta}_n - \theta)\| \right), \quad (3.1)$$

where we do not distinguish θ and θ_0 because they are the same under the null. Its limiting null distribution is

$$H_{g,d} = \int_0^1 \mathbf{1} \left(g \|\Sigma(P)^{1/2} (W_d(r) - rW_d(1))\| \geq \|\Sigma(P)^{1/2} W_d(1)\| \right) dr, \quad (3.2)$$

with $\Sigma(P)$ as in (2.7).

When $d = 1$, $H_{g,1} := H_g$ is pivotal since $\Sigma(P)$ in (3.2) gets canceled, and the quantiles of H_g can be obtained by monte carlo simulation. Let $H_g(\alpha)$ be the α -th quantile of the distribution H_g . Then the corresponding $100(1 - \alpha)\%$ confidence interval for μ is

$$\begin{aligned} & \left\{ \mu : \widetilde{pval}_{n,g}^{GS} > H_g(\alpha) \right\} \\ &= \left\{ \mu : n^{-1} \sum_{t=1}^n \mathbf{1} \left(gn^{-1/2} |t(\bar{X}_t - \bar{X}_n)| \geq |\sqrt{n}(\bar{X}_n - \mu)| \right) > H_g(\alpha) \right\} \\ &= \left(\bar{X}_n - n^{-1/2} c_{n,g}(1 - H_g(\alpha)), \bar{X}_n + n^{-1/2} c_{n,g}(1 - H_g(\alpha)) \right), \end{aligned} \quad (3.3)$$

where $c_{n,g}(1 - \alpha) = \inf\{x : M_{n,g}(x) \geq 1 - \alpha\}$ for $\alpha \in [0, 1)$.

Take $\beta_n^{GS}(g) := P(\max_{t=1, \dots, n} n^{-1/2} g |t(\bar{X}_t - \bar{X}_n)| < \sqrt{n} |\bar{X}_n - \mu|)$ and $\beta^{GS}(g) := P(\sup_{r \in (0,1)} g |W(r) - rW(1)| < |W(1)|)$. The coverage bound of the interval in (3.3) is

$$\begin{aligned} P(\mu \in (3.3)) &= P(\sqrt{n}(|\bar{X}_n - \mu| < c_{n,g}(1 - \alpha)) \\ &\leq P\left(\sqrt{n}|\bar{X}_n - \mu| \leq \max_{t=1, \dots, n} gn^{-1/2} |t(\bar{X}_t - \bar{X}_n)|\right) = 1 - \beta_n^{GS}(g). \end{aligned}$$

For $d > 1$, $H_{g,d}$ depends on the unknown long-run variance matrix $\Sigma(P)$ and is not pivotal; we opt to bypass this difficulty by subsampling approximation at a second stage, see Section 3.2 in Shao and Politis (2013). For completeness, we present the details here. Denote by n' the subsampling width at the second stage. For each subsample $\{X_j, \dots, X_{j+n'-1}\}$, we take the subsampling counterpart of $\widetilde{pval}_{n,l}^{GS}$ as

$$q_{n',g}^{(j)} = (n')^{-1} \sum_{t=1}^{n'} \mathbf{1} \left(g \|n'^{-1/2} t (\hat{\theta}_{j,j+t-1} - \hat{\theta}_{j,j+n'-1})\| \geq \|n'^{1/2} (\hat{\theta}_{j,j+n'-1} - \hat{\theta}_n)\| \right)$$

for $j = 1, \dots, n-n'+1$. Denote the empirical distribution function of $\{q_{n',g}^{(j)}\}_{j=1}^{n-n'+1}$ by $Q_{n,n',g}(x) = (n - n' + 1)^{-1} \sum_{j=1}^{n-n'+1} \mathbf{1}(q_{n',g}^{(j)} \leq x)$, which can be used to approximate the sampling distribution or the limiting null distribution of $\widetilde{pval}_{n,l}^{GS}$. Let $c_{n,n',g}(1 - \alpha) = \inf\{x : Q_{n,n',g}(x) \geq 1 - \alpha\}$. Then the calibrated $100(1 - \alpha)\%$ subsampling-based confidence region for θ is

$$\left\{ \theta \in \mathbf{R}^d : \widetilde{pval}_{n,l}^{GS} \text{ in (3.1)} > c_{n,n',g}(\alpha) \right\}, \tag{3.4}$$

which is equivalent to

$$\left\{ \theta \in \mathbf{R}^d : \sqrt{n} \|\hat{\theta}_n - \theta\| < (1 - c_{n,n',g}(\alpha))\text{-th quantile of } M_{n,g} \right\}. \tag{3.5}$$

Take $\beta_n^{GS}(g; d; \Sigma) := P(\max_{t=1, \dots, n} n^{-1/2} g \|t(\hat{\theta}_t - \hat{\theta}_n)\| < \sqrt{n} \|\hat{\theta}_n - \theta\|)$ and $\beta^{GS}(g; d; \Sigma) := P(\sup_{r \in (0,1)} g \|\Sigma^{1/2}(W_d(r) - rW_d(1))\| < \|\Sigma^{1/2}W_d(1)\|)$. We can show that the coverage bound of (3.4) is $1 - \beta_n^{GS}(g; d; \Sigma)$. The value of $1 - \beta_n^{GS}(g; d; I_d)$ is tabulated in Tables 5 and 6, from which we can see an improvement of the finite as well as asymptotic bound over the fixed- b subsampling results in Tables 1 and 2; see more comments in Section 3.2.

In estimating the marginal cdf of a stationary process, let $g_n(t, s) = n^{-1/2} t (m_t(s) - m_n(s))$, $t = 1, \dots, n$ be the GS counterpart of $\sqrt{n}(m_n(s) - m(s))$; we approximate the distribution of $\|\sqrt{n}(m_n - m)\|_\infty$ by

$$M_{n,g}(x) = n^{-1} \sum_{t=1}^n \mathbf{1} \left(gn^{-1/2} \|t(m_t - m_n)\|_\infty \leq x \right).$$

The p -value

$$\widetilde{pval}_{n,g}^{GS} = n^{-1} \sum_{t=1}^n \mathbf{1} \left(gn^{-1/2} t \|m_t - m_n\|_\infty \geq \sqrt{n} \|m_n - m\|_\infty \right) \tag{3.6}$$

has a limiting null distribution of

$$\mathcal{H}_g := \int_0^1 \mathbf{1} \left(g \sup_{s \in \mathbf{R}} |K(s, r) - rK(s, 1)| \geq \sup_{s \in \mathbf{R}} |K(s, 1)| \right) dr,$$

where the Gaussian process $K(s, r)$ has a covariance structure determined by the unknown dependence structure of X_t , see (2.11). Therefore a second-stage subsampling needs to be employed. Specifically, let n' be the subsampling window size at the second stage. For each subsample $\{X_t, \dots, X_{t+n'-1}\}$, the subsampling counterpart of $\widetilde{pval}_{n,g}^{GS}$ is

$$\begin{aligned} h_{n',g}^{(t)} &= (n')^{-1} \sum_{j=1}^{n'} \mathbf{1}(g(n')^{-1/2} j \|m_{t,t+j-1} - m_{t,t+n'-1}\|_\infty \\ &\geq (n')^{1/2} \|m_{t,t+n'-1} - m_n\|_\infty) \end{aligned}$$

for $t = 1, \dots, n - n' + 1$. Then we approximate the sampling distribution of $\widetilde{pval}_{n,g}^{GS}$, or its limiting null distribution \mathcal{H}_g , by the empirical distribution of $\{h_{n',g}^{(t)}\}_{t=1}^{n-n'+1}$, which is denoted as

$$J_{n,n',g}(x) = (n - n' + 1)^{-1} \sum_{t=1}^{n-n'+1} \mathbf{1}(h_{n',g}^{(t)} \leq x).$$

Let $\mathcal{H}_g(\alpha)$ be the α -th quantile of \mathcal{H}_g and $\widehat{\mathcal{H}}_g(\alpha)$ be the subsampling-based estimator, $\widehat{\mathcal{H}}_g(\alpha) = \inf\{x : J_{n,n',g}(x) \geq \alpha\}$. Then the confidence band for $m(\cdot)$ is

$$\left\{ m : m \text{ is a distribution function and } \widetilde{pval}_{n,g}^{GS} \text{ in (3.6)} > \widehat{\mathcal{H}}_g(\alpha) \right\}. \quad (3.7)$$

The following proposition states the consistency of subsampling in the second stage, which implies that the coverage for the calibrated confidence band is asymptotically correct. Let

$$V_g(r, \epsilon) = P\left\{ \left| g \sup_{s \in \mathbf{R}} |K(s, r) - rK(s, 1)| - \sup_{s \in \mathbf{R}} |K(s, 1)| \right| = \epsilon \right\}.$$

Proposition 1. *Suppose that $1/n' + n'/n = o(1)$, and (4) in Shao and Politis (2013) holds.*

- (a) *The limiting null distribution of the p-value $\widetilde{pval}_{n,g}^{GS}$ is the distribution of \mathcal{H}_g provided $V_g(r, 0) = 0$ for every $r \in [0, 1]$.*
- (b) *Suppose that the process X_t is α -mixing, \mathcal{H}_g is a continuous random variable and $V_g(r, \epsilon) = 0$ for every $r \in [0, 1]$ and $\epsilon \geq 0$. Then*

$$\sup_{x \in \mathbf{R}} |J_{n,n',g}(x) - P(\mathcal{H}_g \leq x)| = o_p(1).$$

This result is analogous to Theorem 3 in Shao and Politis (2013) in terms of technical conditions and results, and the arguments for the proof of Theorems 3 in

Shao and Politis (2013) can be extended in a straightforward fashion to prove it. The conditions on \mathcal{H}_g and $V_g(r, \epsilon)$ are not easy to verify and this may be related to the regularity of the distribution of the maximum of Gaussian processes; see Azaïs and Wschebor (2001) and the references therein. The assumption that $1/n' + n'/n = o(1)$ implies that the consistency of second-stage subsampling holds under the small- b asymptotics, where $b = n'/n$. If we view $b = n'/n$ as fixed, then the asymptotic coverage of the calibrated confidence set is still different from the nominal level. One can perform further calibration by subsampling, but the selection of the subsampling window size at each stage usually requires expensive computation and the finite sample improvement in coverage accuracy is not guaranteed by doing iterative subsampling. Our GS is a form of prepivoting (Beran (1987, 1988)), but the limiting null distribution of the p -value is not $U(0, 1)$ in our setting; it depends on g and possibly some aspects of the data generating process (for vector and infinitely dimensional parameters).

3.2. Coverage bound for GS

Here we adopt the set-up in Section 2.3 and investigate the finite sample coverage bound of GS.

Tables 5 and 6 give the coverage bound for the vector mean and median respectively, by GS. Without the scaling parameter g (i.e., $g = 1$), the coverage bound is at most 60%, indicating the ‘crude’ confidence interval is not meaningful at usual confidence levels (say, 95%). There is an increasing trend in the coverage bound as g increases, and the bound is close to 1 on the range $g \in [3, \infty)$. We refrain from using large values of g since they lead to a wider CI while preserving similar coverage probability, see Section 4 for more details. Simulation study shows that $g \in [3, 5]$ is a sensible choice for most data generating processes. Although strong positive dependence makes the coverage bound lower, it seems GS quickly adapts to dependence, and brings the coverage bound close to 1 by increasing g .

Figure 2 has the same format as Figure 1, it depicts how the dimensionality affects the coverage bound of GS. We find that the coverage bound quickly stabilizes and gets up to 95% for $g \in [3, 5]$ regardless of d , when the crude CI bound ($g = 1$) reaches zero for large d . Compared to Figure 1, we can see that although high coverage bound can be achieved by shrinking the fraction b in a high-dimensional case, it leads to instability as the subsample shrinks.

Table 4 summarizes the coverage bound for the marginal cdf by GS. The bound for GS at small samples ($n = 50, 100$) is quite close to 1 as long as $g \geq 3.5$. For large samples, the coverage bound is quite robust to the choice of g (other than the $g = 1$ case), and it is comparable to the best coverage bound for the fixed- b method.

Table 5. Bounds on the coverage probability $1 - \beta_n^{GS}(g; d; I_d)$ for the mean by generalized subsampling. The number of replications was 5,000 each. The values for $\beta_n^{GS}(g; d; I_d)$ in the first column are based on simulating 5000 independent Wiener processes for $d = 1$ ($d = 2, 3$ respectively), where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables.

$g \setminus n$	$\rho = 0$			$\rho = 0.5$			$\rho = 0.8$			$\rho = -0.5$			
	∞	50	100	500	3,000	50	100	500	3,000	50	100	500	3,000
$d = 1$													
1	60.16	54.92	56.22	58.06	58.44	49.08	53.30	56.00	59.44	40.48	46.06	53.44	58.90
2.5	93.86	91.36	92.08	92.94	93.16	86.24	89.08	91.96	93.30	77.40	83.60	89.92	92.94
3	96.46	95.02	95.48	96.62	96.48	91.30	93.10	95.46	96.84	83.60	89.18	94.56	96.36
3.5	98.18	97.18	97.78	98.34	98.16	94.82	96.08	97.50	98.58	88.04	92.78	97.02	98.00
4	99.08	98.52	98.72	99.28	99.26	96.96	97.60	98.72	99.14	90.96	95.30	98.26	98.82
5	99.76	99.44	99.60	99.80	99.76	98.80	99.28	99.50	99.70	95.06	97.82	99.36	99.66
$d = 2$													
1	43.95	39.46	41.11	41.46	44.23	34.70	37.03	40.26	43.52	25.22	31.43	39.36	42.02
2.5	94.74	91.32	92.38	93.62	93.86	86.74	89.55	92.02	93.29	76.36	83.65	91.16	92.86
3	97.86	95.82	96.67	97.06	97.58	93.28	94.81	96.32	97.18	84.88	90.88	96.34	96.86
3.5	99.13	98.30	98.62	98.74	99.08	96.50	97.41	98.42	98.95	90.56	94.85	98.16	98.73
4	99.71	99.04	99.45	99.58	99.64	98.14	98.85	99.28	99.67	94.06	97.18	99.08	99.47
5	99.99	99.88	99.90	100.00	99.95	99.54	99.75	99.90	99.92	97.90	99.18	99.84	99.90
$d = 3$													
1	34.80	29.92	29.92	32.62	32.30	23.76	27.44	31.10	32.10	17.28	21.46	28.84	30.62
2.5	94.80	91.94	92.82	93.72	94.30	87.52	89.28	92.88	93.42	75.36	84.58	90.84	93.36
3	98.48	96.82	97.60	97.46	98.06	94.06	95.56	97.28	97.68	86.06	92.20	96.22	97.58
3.5	99.42	98.76	99.10	99.16	99.40	97.30	98.36	99.12	99.16	92.10	96.14	98.54	99.18
4	99.83	99.54	99.62	99.74	99.88	98.82	99.30	99.68	99.70	95.40	97.82	99.26	99.72
5	100.00	99.98	99.98	100.00	99.98	99.86	99.90	99.96	99.98	98.74	99.68	99.90	99.98

Table 6. Bounds on the coverage probability $1 - \beta_n^{GS}(g; d; I_d)$ for the median by generalized subsampling. The number of replications was 5,000 each. The values for $\beta_n^{GS}(g; d; I_d)$ in the first column were based on simulating 5000 independent Wiener processes for $d = 1$ ($d = 2, 3$, respectively), where each Wiener process was approximated by a normalized partial sum of 50,000 i.i.d. standard normal random variables.

$g \setminus n$	$\rho = 0$				$\rho = 0.5$				$\rho = 0.8$				$\rho = -0.5$				
	∞	50	100	500	3,000	50	100	500	3,000	50	100	500	3,000	50	100	500	3,000
$d = 1$																	
1	60.36	58.10	58.68	59.56	60.92	52.56	55.92	58.60	59.44	43.14	49.30	55.38	59.30	59.62	59.22	60.78	60.78
2.5	93.76	89.56	91.54	93.02	93.76	87.74	90.04	92.26	93.96	79.10	84.80	90.52	93.18	90.38	91.98	92.90	94.04
3	96.86	93.38	95.02	95.90	96.78	91.76	93.58	95.52	97.02	84.86	90.20	94.86	96.30	93.54	94.96	96.20	96.96
3.5	98.58	95.46	96.78	97.76	98.40	94.18	95.90	97.38	98.34	88.76	93.42	96.96	98.16	95.58	96.98	97.76	98.32
4	99.30	96.62	97.80	98.72	99.06	96.24	97.30	98.52	99.14	91.34	95.50	98.28	99.08	97.08	98.00	98.74	99.08
5	99.82	98.26	99.04	99.60	99.74	98.24	98.88	99.64	99.76	95.18	97.90	99.42	99.66	98.42	98.98	99.54	99.72
$d = 2$																	
1	45.96	47.00	46.58	45.54	45.43	40.48	42.35	43.76	44.91	30.08	35.25	41.60	43.76	51.72	50.07	48.70	45.51
2.5	94.55	92.70	93.48	94.74	94.67	89.90	91.58	93.56	94.08	80.16	86.34	92.08	93.64	94.20	94.34	94.68	94.65
3	97.91	96.24	97.01	97.58	97.85	94.40	95.83	97.18	97.54	87.40	92.41	96.70	96.99	96.98	97.26	97.78	97.86
3.5	99.39	97.96	98.51	99.02	99.12	96.96	97.94	98.62	99.03	92.28	95.77	98.64	98.73	98.48	98.74	99.06	99.07
4	99.80	98.80	99.26	99.58	99.68	98.32	99.01	99.52	99.79	95.16	97.67	99.24	99.55	99.04	99.36	99.50	99.59
5	99.99	99.56	99.82	99.92	99.96	99.48	99.76	99.96	99.99	98.02	99.25	99.92	99.93	99.68	99.82	99.96	99.91
$d = 3$																	
1	35.69	36.88	37.54	36.98	34.46	31.54	32.88	34.30	34.94	21.62	25.58	32.00	32.16	43.44	42.42	40.16	36.12
2.5	95.07	93.62	94.38	95.30	95.46	90.78	92.72	93.82	94.72	80.42	87.02	92.68	93.68	94.88	95.52	96.16	95.40
3	98.37	97.06	97.84	98.36	98.36	96.14	97.22	97.74	98.30	89.02	93.58	96.96	97.80	98.00	98.34	98.62	98.48
3.5	99.43	98.86	99.06	99.34	99.52	98.36	98.74	99.36	99.32	93.86	96.86	98.94	99.24	99.06	99.36	99.50	99.60
4	99.81	99.60	99.62	99.78	99.86	99.34	99.44	99.76	99.80	96.54	98.36	99.74	99.76	99.62	99.76	99.74	99.84
5	100.00	99.84	99.96	100.00	99.98	99.88	99.90	99.98	100.00	98.86	99.62	99.96	99.98	99.92	99.96	99.96	99.98

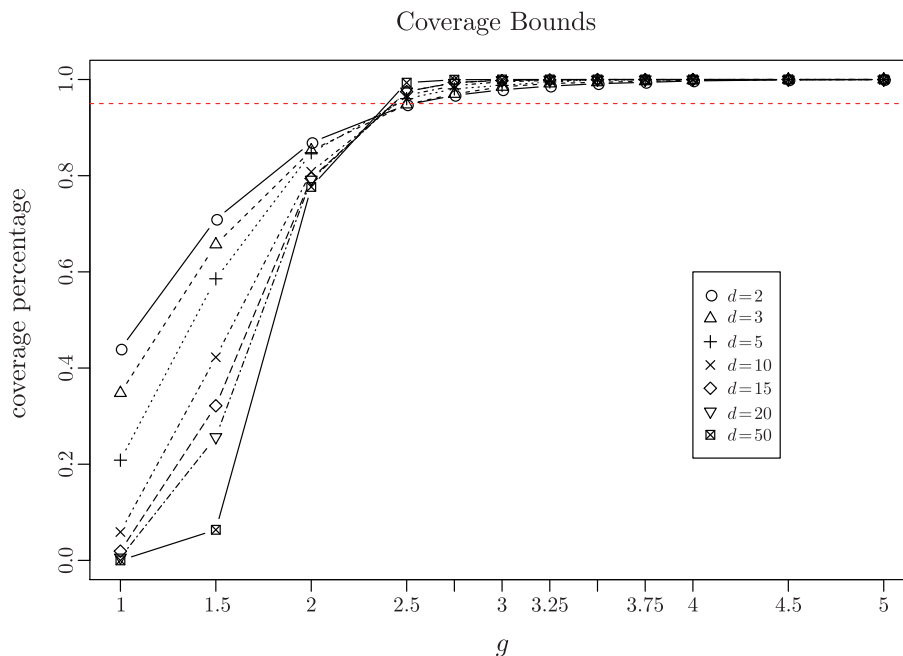


Figure 2. Bounds on the coverage probabilities of GS for the mean at nominal level 95% (dashed line). The data were generated from multivariate standard normal distribution with $n = 5,000$ and the number of Monte Carlo replications was 5,000.

4. Numerical Studies

We have illustrated the improvement on the least upper bound on coverage probability by GS over the fixed- b subsampling. Here we compare the empirical coverage probability of the confidence set by GS, fixed- b , and SN methods, and pay particular attention to the case where the finite sample bound for the fixed- b subsampling is substantially below 1.

We simulated Gaussian (vector) AR(1) model and constructed the confidence set for the finite-dimensional parameter (univariate mean, vector mean) or confidence band for the marginal cdf. In the univariate mean case, the limiting null distribution of the p -value is pivotal for both fixed- b and GS, and the quantile $H_g(\alpha)$ can be simulated, while $\tilde{G}_b(\alpha)$ has been tabulated in Shao and Politis (2013). In vector and infinite-dimensional parameter cases, we employed a second stage subsampling method to approximate the limiting null distribution of the p -value for both fixed- b method and GS, as described in Section 3.1. Following the proposal by Bickel and Sakov (2008), a data-driven bandwidth selection procedure in the second stage subsampling was carried out as follows.

- Step 1. For a predetermined interval $[K_1, K_2]$ and $\gamma \in (0, 1)$, take a sequence of n_j 's of the form $n_j = \lfloor \gamma^{j-1} K_2 \rfloor$, for $j = 1, 2, \dots, \lfloor \log(K_2/K_1)/\log(1/\gamma) \rfloor$.
- Step 2. For each n_j , find J_{n,n_j} , where J_{n,n_j} is the subsampling-based distribution estimator for the sampling distribution of the p -value.
- Step 3. Set $j_0 = \operatorname{argmin}_{j=1, \dots, \lfloor \log(K_2/K_1)/\{-\log(\gamma)\}} \sup_{x \in \mathbf{R}} |J_{n,n_j}(x) - J_{n,n_{j+1}}(x)|$. Then the optimal block size is $\gamma^{j_0} K_2$. If the difference is minimized for a few values of j , then pick the largest among them.

In the simulation, we set $(K_1, K_2, \gamma) = (5, 0.3n, 0.75)$, which corresponds to a sequence of block lengths of $(30, 22, 16, 12, 9, 7, 5)$ when $n = 100$.

Figure 3 compares the empirical coverage probability for the CI of the univariate mean constructed by GS, fixed- b and SN methods. Following Shao and Politis (2013), the range of b is $[0.01, 0.16]$ since the formula given therein for the critical value $\tilde{G}_b(\alpha)$ may become negative for $b > 0.16$ and $\alpha = 0.05$. The range of g is $[2.5, 10]$. It shows that, for $\rho = 0, 0.5, 0.8$, GS delivers more accurate coverage probability than fixed- b subsampling, and is quite comparable to SN for a range of g 's. As $\rho > 0$ becomes larger, fixed- b subsampling encounters more severe undercoverage, whereas GS manages to bring the coverage probability toward the nominal level. The fixed- b subsampling relies heavily on the choice of b , while GS performs quite stably for $g \in [3, 10]$ and is comparable to SN method in terms of both coverage accuracy and interval length. When $\rho = -0.5$, all three methods exhibit over-coverage phenomena, which is consistent with the high coverage bound presented in Tables 1 and 5. The coverage probability delivered by the GS confidence interval is closer to the nominal level and, as a trade off, the length of CI is moderately longer compared to fixed- b approach. The pattern is the same for median and quantiles, and is not shown.

Figure 4 shows the empirical coverage probability of the confidence region for a 3-dimensional vector mean of a VAR(1) model. It can be seen that when the dependence is moderate, say, $\rho = 0.5$, the coverage probability of the fixed- b method is not far from the nominal level and, in this case, GS can barely improve the coverage. As dependence grows (positively) stronger, the coverage probability of the fixed- b subsampling deviates from the 95% level by a larger amount, and GS brings the coverage probability much closer to the nominal level. When $\rho < 0$, GS is still much better than fixed- b subsampling except for $b = 0.01$. We also compare the volumes of the confidence regions. Notice that the confidence region constructed by fixed- b subsampling and GS are d -dimensional balls with radius determined by the p -value calibrated critical value, while SN method results in a confidence ellipsoid. For an ellipsoid with semi-principal axes of length a, b, c , take the *equivalent radius* to be $R = (abc)^{1/3}$. Again we see a wider confidence region radius by GS as a reasonable trade-off for more accurate

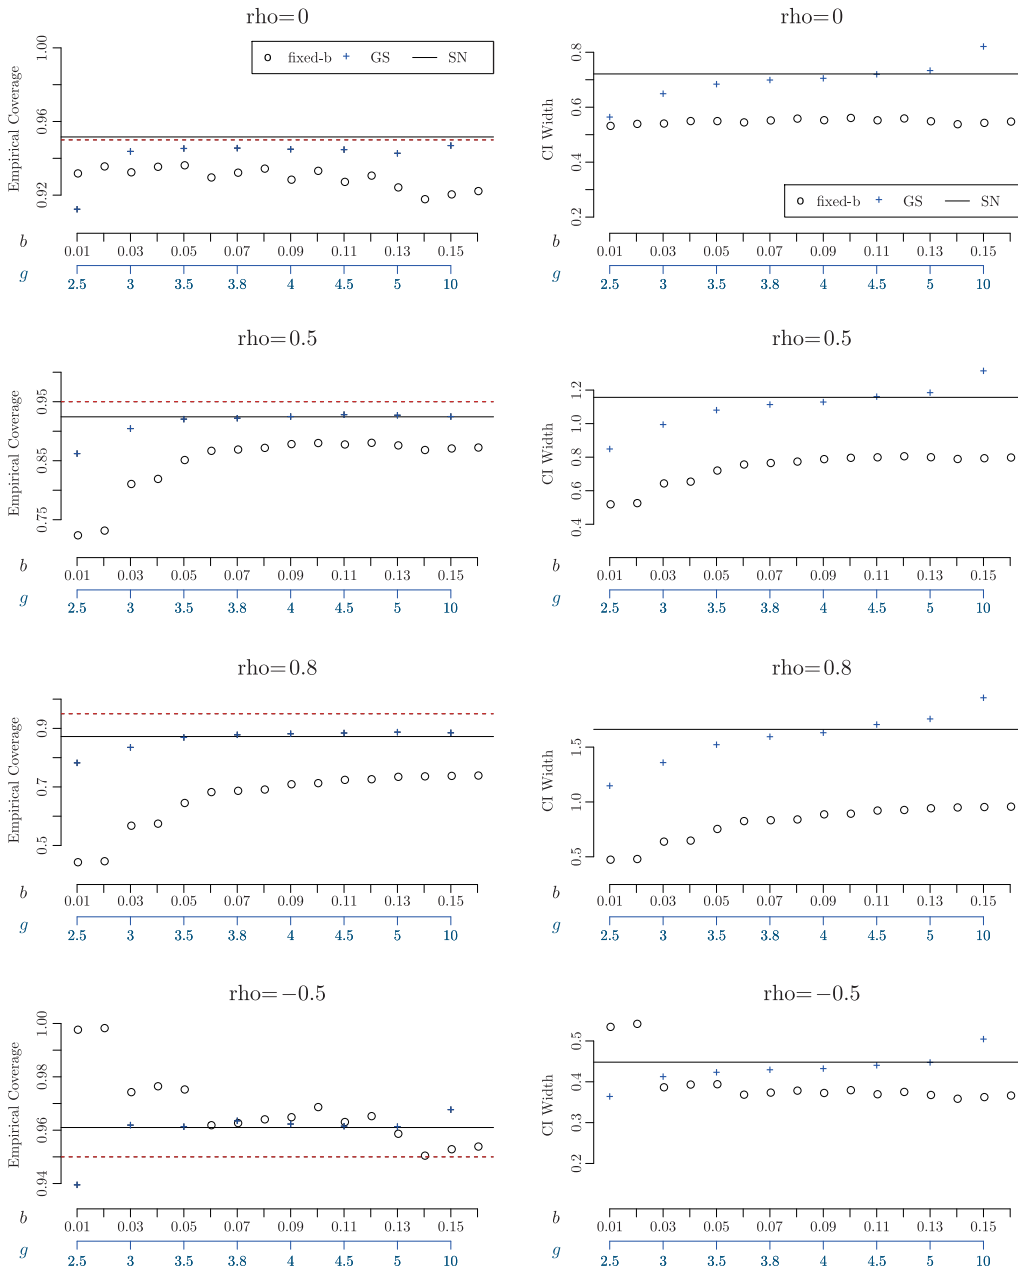


Figure 3. Empirical coverage probabilities and CI width for the mean by GS ('+'), fixed-b ('o') and SN (solid line) at nominal level 95% (dashed line). The data were generated from Gaussian AR(1) models with $n = 50$, the number of Monte Carlo replications was 5,000.

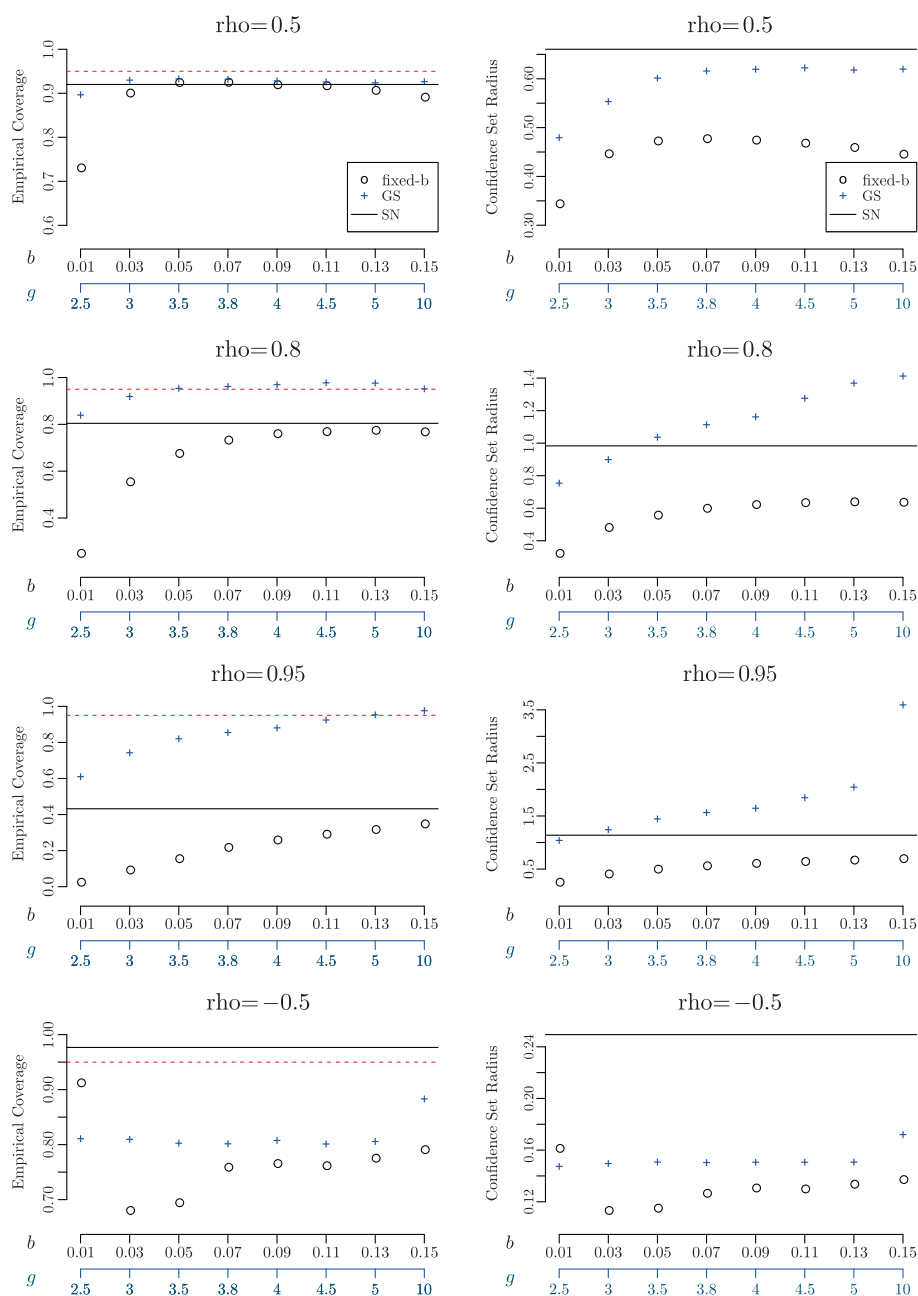


Figure 4. Empirical coverage probabilities and (equivalent) radius of confidence region for 3-dimensional vector mean by GS ('+'), fixed- b ('o') and SN (solid line) at nominal level 95% (dashed line). The data were generated from vector Gaussian AR(1) models with $n = 100$, the number of Monte Carlo replications was 5,000.

coverage probability. Overall, the performance of GS is stable for a wide range of g , and its performance is comparable to the SN method in terms of coverage and confidence region radius when $\rho = 0.5$, and can vastly outperform SN method in coverage when $\rho = 0.8$ or 0.95 ; when $\rho = -0.5$, SN does better in coverage.

Figure 5 shows the empirical coverage probability of the confidence band for the marginal cdf of AR(1) model by GS and the fixed- b method. The SN method is not applicable for infinite-dimensional parameters. It appears that GS is slightly worse than fixed- b subsampling when dependence is weak or moderate, but GS's advantage shows up as the dependence strengthens, for either positive or negative ρ , and the coverage probability for GS is satisfactorily close to the nominal level and quite stable across a wide range. As expected, the GS-based confidence set is typically wider, as a price we pay for more accurate coverage.

Overall, GS's performance is competitive, and it displays distinct advantages over fixed- b subsampling when the dependence is (positively) strong; it is comparable to SN for finite-dimensional parameters.

5. Conclusion

In this paper, we study the coverage upper bound on the coverage probabilities of fixed- b subsampling based confidence sets. We derive the formulae for both finite sample bound and the limiting bound, and tabulate them for several combinations of (n, b) and time series models with strong/weak dependence. This seems to be the first time that the coverage bound problem is recognized for subsampling methods. Our numerical results show that the bound can be far from 1 and the finite sample coverage can be far below the nominal level when the dimension of the parameters is large, the dependence of the time series is positively strong, or b is large. This finding suggests that caution be taken when applying subsampling methods to time series inference with a high-dimensional parameter or strong dependence.

Our proposed GS methodology combines the recursive subsample idea in the self-normalization of Shao (2010) and the p -value calibration idea in Shao and Politis (2013), and introduces a scaling parameter g to alleviate the bound problem when $g = 1$. The numerical comparison with SN and fixed- b subsampling shows that the GS can deliver comparable or sometimes favorable coverage accuracy for a range of g 's. It is interesting to ask about optimal g and whether there is a data-driven algorithm for its choice. We leave this for future research. There are a few variants of the GS method. For example, we can use different weights for the values computed over blocks of possibly different size, generalizing the scaling idea. Or we can use all possible blocks of consecutive observations in the calculation of values of the subsample statistic. We expect the coverage bound problem also occurs for a block bootstrap based confidence set, in view of the

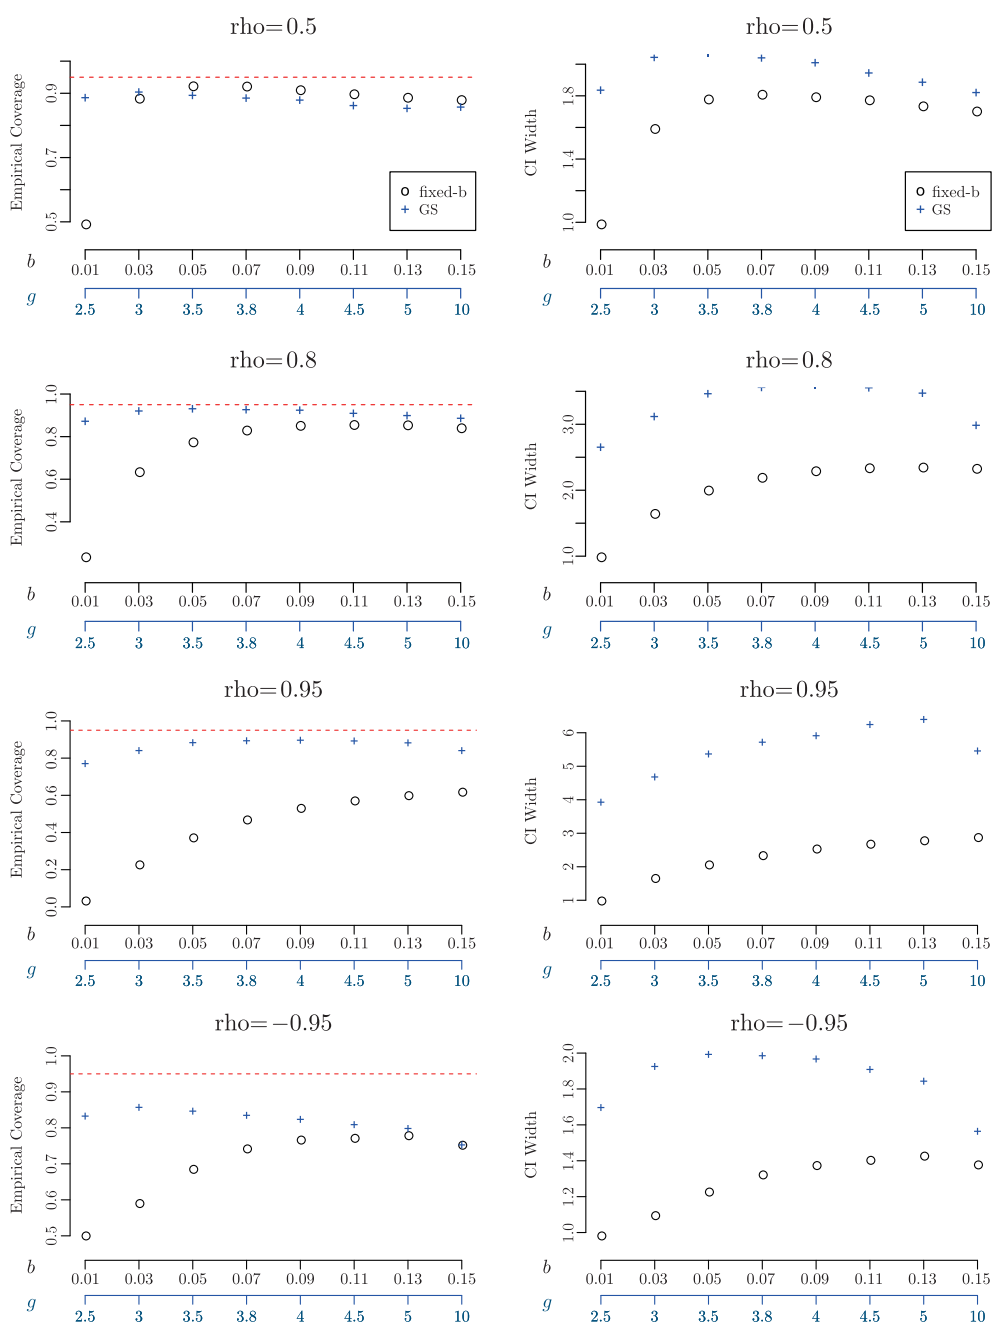


Figure 5. Coverage probabilities and widths of confidence bands for the marginal cdf by fixed- b and generalized subsampling. The data were generated from AR(1) models with $n = 100$ and the number of Monte Carlo replications was 5,000.

developments on fixed- b block bootstrap in Shao and Politis (2013). All these topics are worthy of further investigation.

Acknowledgement

The work of the second author is supported in part by National Science Foundation Grant DMS-1104545. We are grateful to two reviewers for their helpful comments, which led to improvements.

References

- Azaïs, J. M. and Wschebor, M. (2001). On the regularity of the distribution of the maximum of one-parameter Gaussian processes. *Probab. Theory Related Fields* **119**, 70-98.
- Berkes, I., Hörmann, S. and Schauer, J. (2009). Asymptotic results for the empirical process of stationary sequences. *Stochastic Process. Appl.* **119**, 1298-1324.
- Beran, R. (1987). Previoting to reduce level error of confidence sets. *Biometrika* **74**, 457-468.
- Beran, R. (1988). Prepivotting test statistics: A bootstrap view of asymptotic refinements. *J. Amer. Statist. Assoc.* **83**, 687-697.
- Bickel, P. and Sakov, A. (2008). On the choice of m in the m out of n bootstrap and confidence bounds for extrema. *Statist. Sinica* **18**, 967-985.
- Kiefer, N. M. and Vogelsang, T. J. (2005). A new asymptotic theory for heteroskedasticity-autocorrelation robust tests. *Econom. Theory* **21**, 1130-1164.
- Lahiri, S. N. (2001). Effects of block lengths on the validity of block resampling methods. *Probab. Theory Related Fields* **121**, 73-97.
- Nordman, D. J., Bunzel, H. and Lahiri, S. N. (2013). A non-standard empirical likelihood for time series. *Ann. Statist.* **4**, 3050-3073.
- Politis, D. N. and Romano, J. P. (1994). Large sample confidence regions based on subsamples under minimal assumptions. *Ann. Statist.* **22**, 2031-2050.
- Politis, D. N., Romano, J. P. and Wolf, M. (1999a). *Subsampling*. Springer-Verlag, New York.
- Politis, D. N., Romano, J. P. and Wolf, M. (1999b). Weak convergence of dependent empirical measures with application to subsampling in function spaces. *J. Statist. Plann. Inference* **79**, 179-190.
- Shao, X. (2010). A self-normalized approach to confidence interval construction in time series. *J. Roy. Statist. Soc. Ser. B* **72**, 343-366. Corrigendum: 2010, **72**, 695-696.
- Shao, X. and Politis, D. N. (2013). Fixed- b subsampling and block bootstrap: improved confidence sets based on p -value calibration. *J. Roy. Statist. Soc. Ser. B* **75**, 161-184.
- Zhang, X. and Shao, X. (2013). Fixed-smoothing asymptotics for time series. *Ann. Statist.* **41**, 1329-1349.

Department of Mathematics and Statistics, Boston University, Boston, MA 02135, USA.

E-mail: yxhuang@bu.edu

Department of Statistics, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA.

E-mail: xshao@illinois.edu

(Received May 2014; accepted December 2014)