

Bootstrapping two-stage quasi-maximum likelihood estimators of time series models*

Silvia Gonçalves[†], Ulrich Hounyo[‡], Andrew J. Patton[§] and Kevin Sheppard[¶]

January 15, 2022

Abstract

This paper provides results on the validity of bootstrap inference methods for two-stage quasi-maximum likelihood estimation involving time series data, such as those used for multivariate volatility models or copula-based models. Existing approaches require the researcher to compute and combine many first- and second-order derivatives, which can be difficult to do and is susceptible to error. Bootstrap methods are simpler to apply, allowing the substitution of capital (CPU cycles) for labor (keeping track of derivatives). We show the consistency of the bootstrap distribution and consistency of bootstrap variance estimators, thereby justifying the use of bootstrap percentile intervals and bootstrap standard errors.

*A MATLAB toolbox for this article is available at <https://sites.google.com/site/ulrichounyo>.

[†]Department of Economics, McGill University, Leacock Building, Room 506, 855 Sherbrooke Street West, Montreal, Quebec, H3A 2T7. Email: silvia.goncalves@mcgill.ca.

[‡]Department of Economics, University at Albany, SUNY, Buiding 25, room 231, 1400 Washington Ave, Albany, New York, 12222, United States. Email: khounyo@albany.edu.

[§]Department of Economics, Duke University, Durham, NC 27708, United States. Email: andrew.patton@duke.edu.

[¶]Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, United Kingdom. Email: kevin.sheppard@economics.ox.ac.uk.

Keywords: inference, multivariate GARCH, copulas

J.E.L. Codes: C15, C32, C58

1 Introduction

Many econometric models are estimated in stages in order to make them more computationally tractable. For example, copula-based models often have the marginal distribution parameters estimated in the first N stages, and then the copula parameters estimated in a final stage (see Joe (1995) and Patton (2006), for example). Similarly, the famous Dynamic Conditional Correlation (DCC) GARCH model of Engle (2002) first estimates univariate GARCH models for each variable, and then the correlation matrix component is estimated. In empirical finance, the ubiquitous Fama-MacBeth (1973) procedure involves estimating assets' exposures to risk factors in a first stage, before estimating risk premia associated with those factors in a second stage. In all cases, valid inference on parameters in the final stage must account for estimation error accumulated in earlier stages.

Standard methods to account for estimation error from earlier estimation stages requires computing many first- and second-order derivatives of the overall objective function.¹ In most cases, it is difficult or tedious to obtain these analytically, and many authors instead rely on numerical derivatives. Both of these approaches are susceptible to error: human, in the first case; numerical, in the second. Inference based on the bootstrap is an attractive alternative, allowing the substitution of capital (CPU cycles) for labor (keeping track of derivatives). See, e.g. Cochrane (2001, Chapter 15.2) for bootstrap inference in two-pass regressions and Patton (2012) for bootstrap inference in copula-based models.

Some related results are available in the literature. For example, Chen et al. (2003) and Armstrong et al. (2014) consider the bootstrap for two-step nonlinear parametric

¹For a textbook discussion of inference in multi-stage estimation see White (1994) and Newey and MacFadden (1994).

and semiparametric models, and Cattaneo et al. (2019) considers bootstrap inference for two-step GMM estimators, but all of these papers consider only i.i.d. data. Bootstrap methods for *one*-step quasi-maximum likelihood estimators (QMLE) have been considered (e.g., Gonçalves and White (2004)) for both i.i.d. and time series data, but no general results are available for multistep QMLE. We address this gap and show bootstrap validity for such applications under very general time series dependence and heterogeneity.

We consider two approaches for bootstrap inference. The first jointly resamples the contributions to the quasi log-likelihood functions in the various stages of estimation. Since model misspecification at any stage can induce time series dependence in the scores of each model, we rely on the moving blocks bootstrap (MBB) of Künsch (1989) and Liu and Singh (1992). This approach involves re-estimating the model on each bootstrapped objective function, and this “fully optimized” approach is thus computationally demanding.

The second approach is a fast resampling method that avoids any optimization problem in the bootstrap world. In particular, our proposal is to resample the score function underlying the asymptotic linear representation of the final-step QMLE, evaluated at the parameters estimated in the real data. This is related to the resampling method proposed in Armstrong et al. (2014), however in their approach the first-stage model is re-estimated on each bootstrap sample (which has advantages in the semiparametric settings considered in that paper) while our proposed method requires no estimation in the bootstrap world².

We prove two sets of results for both bootstrap methods. For simplicity, we provide formal results only for the two-step QMLE³. First, we show the consistency of the bootstrap distribution of the final-stage parameter estimator for the usual limit distribution of that parameter, justifying the construction of bootstrap percentile intervals. We consider

²Other papers that have proposed fast resampling methods include Davidson and MacKinnon (1999), Andrews (2002), Gonçalves and White (2004), Hong and Scaillet (2006), and La Vecchia et al. (2020), among others. However, all these papers consider one-step M or GMM estimators.

³These results can easily be generalized to multistage QMLE. In our Monte Carlo simulations, we estimate a copula model involving a three-stage QMLE.

regularity conditions that allow for time series dependence and heterogeneity of unknown form, extending the work of Gonçalves and White (2004) who show the asymptotic validity of the MBB for inference on one-step QMLE. Second, we prove the consistency of bootstrap standard errors. Although these are very popular in applied work, their consistency does not automatically follow from the consistency of the bootstrap distribution. This was recently emphasized by Hahn and Liao (2020), who proved that bootstrap standard errors for smooth functions of i.i.d. data may lead to conservative inference. Here, we show that the bootstrap variance estimator of the two-step QMLE is consistent by verifying a certain uniform integrability condition. We follow the approach of Kato (2011) and Cheng (2015) and rely on empirical process theory to prove our results. The results of those two papers apply to one-step M-estimators with i.i.d. data and do not cover time series applications. They also do not cover two-step QMLE, even for the i.i.d. data. Similarly, although the results of Gonçalves and White (2005) allow for time series dependence, they are specific to the one-step least squares estimator. Thus, no results appear to be available regarding the consistency of bootstrap variance estimators of one or multi-stage QMLE with general time series dependence.

We present Monte Carlo simulations to illustrate the usefulness of our results. We consider the estimation of a bivariate copula model, where estimation is done in stages and the parameter of interest is the copula parameter, which is estimated in the final (third) stage. In addition to the standard confidence interval that relies on analytical standard errors, we consider two types of bootstrap-based intervals: intervals that rely on bootstrap standard errors, but use the normal critical value, and bootstrap percentile intervals. We find that all methods provide similarly good coverage probabilities, even for the smaller sample sizes. This is in agreement with the theory of the bootstrap since none of the methods promises asymptotic refinements. However, we find that the competing inference methods differ in their confidence interval lengths. In particular, the intervals based on the

fully optimized bootstrap method tend to be narrower than the intervals based on either the fast resampling method or the usual asymptotic approach using analytical standard errors. We show that this is due to the fully optimized bootstrap standard error having a smaller mean squared error than the competing methods. Thus, although more computationally intensive than the fast resampling method, the fully optimized bootstrap intervals have better finite sample properties.

The rest of the paper is organized as follows. In Section 2, we present the framework and provide an example of a two-step QMLE based on the bivariate copula model. In Section 3, we describe our two bootstrap methods and prove their consistency in Section 4. Section 5 contains the simulation results and Section 6 concludes. Assumptions are collected in Appendix A. An online supplementary appendix contains the proofs of the bootstrap distribution consistency and the proof of the bootstrap variance consistency results. This online appendix also contains two general bootstrap consistency theorems for two-step M and GMM bootstrap estimators, which could be of independent interest as their high level conditions can be verified for any bootstrap scheme. Our focus on the two-stage QMLE estimator based on the moving blocks bootstrap in the main text is justified by the fact that this covers the main applications in finance we have in mind.

2 Framework

2.1 Two-stage quasi maximum likelihood estimation

Suppose $\{X_t : \Omega \rightarrow \mathbb{R}^l, t \in \mathbb{N}\}$ denotes a sequence of \mathbb{R}^l -valued random vectors defined on some probability space (Ω, \mathcal{F}, P) . Let $\Theta = \mathcal{A} \times \mathcal{B}$, where \mathcal{A} and \mathcal{B} are compact subsets of finite dimensional Euclidean spaces. Given data $\{X_t : t = 1, \dots, n\}$, our goal is to estimate a parameter vector $\beta_0 \in \mathcal{B} \subset \mathbb{R}^p$ by a two-stage quasi-maximum likelihood (2QMLE) estimator. We focus on two-stage QMLE for ease of presentation, but our results generalize

directly to multi-stage QMLEs. (E.g., our simulation study considers a three-stage QMLE problem.)

In the first step, we estimate $\alpha_0 \in \mathcal{A} \subset \mathbb{R}^k$ with

$$\hat{\alpha}_n = \arg \max_{\alpha \in \mathcal{A}} Q_{1n}(\alpha),$$

where $Q_{1n}(\alpha) \equiv n^{-1} \sum_{t=1}^n \log f_{1t}(X^t, \alpha)$, with $X^t \equiv (X_1, \dots, X_{t-1}, X_t)$, for some quasi-likelihood function $f_{1t}(X^t, \alpha) : \mathbb{R}^{lt} \times \mathcal{A} \rightarrow \mathbb{R}^+$. To simplify the notation, we sometimes write $f_{1t}(\alpha) \equiv f_{1t}(X^t, \alpha)$. A similar notation is used for any other function of X^t throughout.

In the second step, we estimate β with

$$\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} Q_{2n}(\hat{\alpha}_n, \beta)$$

where $Q_{2n}(\alpha, \beta) \equiv n^{-1} \sum_{t=1}^n \log f_{2t}(X^t, \alpha, \beta)$, for a conditional quasi-likelihood function $f_{2t}(\alpha, \beta) \equiv f_{2t}(X^t, \alpha, \beta) : \mathbb{R}^{lt} \times \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}^+$. We allow for time heterogeneity in $f_{1t}(\alpha)$ and $f_{2t}(\alpha, \beta)$ (i.e. the functional forms may depend on t) and we also allow for the possibility that these functions depend on the past information up to time t (i.e., X^t is a vector of possibly growing dimension).

2.2 An example: copula models

An example of time series models that are often estimated in multiple stages are copula-based multivariate models. These models combine separately estimated marginal distributions via a copula function to form a joint distribution. When the parameters that characterize the marginal distributions are different from those that characterize the copula density function, estimation and inference can be done in stages. Our results can be useful in this context.

To illustrate, let $X_t \equiv (X_{1t}, X_{2t})'$ denote a random vector whose joint conditional density

we would like to model. By the usual decomposition, we can write

$$\log g(X_1, \dots, X_n, \theta) = \sum_{t=1}^n \log g_t(X_t | \mathcal{F}^{t-1}, \theta),$$

where $g_t(X_t | \mathcal{F}^{t-1}, \theta)$ is the conditional density function of X_t given \mathcal{F}^{t-1} . Suppose $X_{it} | \mathcal{F}^{t-1} \sim G_{it}(\alpha_i)$, some distribution function parametrized by a set of parameters α_i with density function $g_{it}(\alpha_i)$. The joint (conditional) pdf of X_t is then given by

$$g_t(X_t | \mathcal{F}^{t-1}, \theta) = g_{1t}(X_{1t}, \alpha_1) g_{2t}(X_{2t}, \alpha_2) c_t(G_{1t}(X_{1t}, \alpha_1), G_{2t}(X_{2t}, \alpha_2), \beta),$$

where $c_t(\cdot, \cdot, \beta)$ is a copula density function parametrized by β , and $\theta = (\alpha_1, \alpha_2, \beta)'$ denotes the full set of parameters. It follows that the joint log-likelihood function can be written as

$$\begin{aligned} \log g(X_1, \dots, X_n, \theta) &= \sum_{t=1}^n \log g_{1t}(X_{1t} | \mathcal{F}^{t-1}, \alpha_1) + \sum_{t=1}^n \log g_{2t}(X_{2t} | \mathcal{F}^{t-1}, \alpha_2) \\ &\quad + \sum_{t=1}^n \log c_t(G_{1t}(X_{1t}, \alpha_1), G_{2t}(X_{2t}, \alpha_2) | \mathcal{F}^{t-1}, \beta). \end{aligned}$$

When the parameters characterizing the marginals and the copula function are separable (i.e. the parameters that enter one marginal do not enter another marginal nor the copula function and there are no cross equation restrictions), these parameters can be estimated in stages. In particular, first estimate α_i by QMLE:

$$\hat{\alpha}_{in} = \arg \max_{\alpha_i \in \mathcal{A}} \sum_{t=1}^n \log g_{it}(X_{it} | \mathcal{F}^{t-1}, \alpha_i), \text{ for } i = 1, 2,$$

and then estimate the copula parameters β in a second stage by

$$\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \sum_{t=1}^n \log c_t(G_{1t}(X_{1t}, \hat{\alpha}_{1n}), G_{2t}(X_{2t}, \hat{\alpha}_{2n}) | \mathcal{F}^{t-1}, \beta).$$

Thus, in our previous notation,

$$Q_{2n}(\hat{\alpha}_n, \beta) = \sum_{t=1}^n \log f_{2t}(X^t, \hat{\alpha}_n, \beta), \text{ where } \hat{\alpha}_n = (\hat{\alpha}_{1n}, \hat{\alpha}_{2n})', \text{ and}$$

$$f_{2t}(X^t, \hat{\alpha}_n, \beta) \equiv c_t(G_{1t}(X_{1t}, \hat{\alpha}_{1n}), G_{2t}(X_{2t}, \hat{\alpha}_{2n}) | \mathcal{F}^{t-1}, \beta).$$

The contributions to this quasi-log-likelihood function depend on the sample of $X_t = (X_{1t}, X_{2t})'$ up to time t through the integral probability transforms $G_{it}(X_{it}, \hat{\alpha}_{in})$.

2.3 Asymptotic properties of two-stage QMLE

We now briefly review the asymptotic properties of the two-stage QMLE to assist in later explaining the properties that the bootstrap needs to have in order to be asymptotically valid. These properties are well known in the literature, see e.g. White (1994), Newey and McFadden (1994) and Wooldridge (1994).

Let α_0 be the unique maximizer of $\bar{Q}_1(\alpha) \equiv \lim_{n \rightarrow \infty} E(Q_{1n}(\alpha))$ on \mathcal{A} and let β_0 be the unique maximizer of $\bar{Q}_2(\alpha_0, \beta) \equiv \lim_{n \rightarrow \infty} E(Q_{2n}(\alpha_0, \beta))$ on \mathcal{B} .⁴ Then, under Assumption A in Appendix A, we can show that $\hat{\beta}_n \xrightarrow{P} \beta_0$ and

$$\sqrt{n}(\hat{\beta}_n - \beta_0) \xrightarrow{d} N(0, C_0),$$

where the asymptotic covariance matrix $C_0 \equiv H_0^{-1} J_0 H_0^{-1}$, with

$$H_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \beta'} s_{2t}(\alpha_0, \beta_0) \right), \text{ such that } s_{2t}(\alpha_0, \beta_0) \equiv \frac{\partial}{\partial \beta} \log f_{2t}(\alpha_0, \beta_0),$$

and

$$J_0 \equiv \lim_{n \rightarrow \infty} \text{Var} \left(n^{-\frac{1}{2}} \sum_{t=1}^n (s_{2t}(\alpha_0, \beta_0) - F_0 A_0^{-1} s_{1t}(\alpha_0)) \right),$$

⁴Under general heterogeneity and time series dependence, α_0 and β_0 could depend on n . We omit the index n to simplify the notation.

where

$$s_{1t}(\alpha_0) \equiv \frac{\partial}{\partial \alpha} \log f_{1t}(\alpha_0), \quad A_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{1t}(\alpha_0) \right), \quad \text{and}$$

$$F_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{2t}(\alpha_0, \beta_0) \right).$$

As this result shows, the impact of the first stage estimation of α_0 is not negligible asymptotically except when $F_0 = 0$. This implies that we need to adjust the standard errors of $\hat{\beta}_n$ for the added estimation uncertainty of $\hat{\alpha}_n$. Although a consistent estimator of J_0 can be obtained by applying a HAC (heteroskedasticity and autocorrelation covariance) estimator to $\left\{ s_{2t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1t}(\hat{\alpha}_n) \right\}$ (where \hat{F}_n and \hat{A}_n are consistent estimators of F_0 and A_0) in practice the bootstrap is often used. Our goal is to provide a set of conditions that justify this practice in time series applications.

3 Bootstrap methods

The asymptotic validity of the bootstrap depends on its ability to mimic the asymptotic variance-covariance matrix of $\hat{\beta}_n$. The form of J_0 suggests that the bootstrap should replicate the time series dependence and the heterogeneity properties of the score vector $\left\{ s_t(\alpha_0, \beta_0) \equiv s_{2t}(\alpha_0, \beta_0) - F_0 A_0^{-1} s_{1t}(\alpha_0) \right\}$. Model misspecification at any stage can induce time series dependence and our approach is to use a block bootstrap. In particular, we rely on the moving blocks bootstrap (MBB) of Künsch (1989) and Liu and Singh (1992). See also Gonçalves and White (2002, 2004, 2005) for the validity of the MBB under general time series dependence and heterogeneity. Although we focus on the MBB, other bootstrap methods which are robust to serial dependence could be used. Examples include the tapered block bootstrap of Paparoditis and Politis (2002) and the kernel block bootstrap of Parente and Smith (2021). The validity of these alternative bootstrap schemes could be

established by verifying the high level conditions given in Theorems S4.3 and S4.4 in the online Appendix.

We consider two different methods. One is based on resampling the contributions to the likelihood functions $\{f_{1t}(\alpha)\}$ (which yields a bootstrap QMLE $\hat{\alpha}_n^*$) and $\{f_{2t}(\hat{\alpha}_n^*, \beta)\}$ (which is optimized over β to yield $\hat{\beta}_n^*$). The same bootstrap indices obtained with the MBB are used across the two stages, ensuring that this method mimics the time series dependence of the extended score. Because it requires two (or multiple) sets of maximization, this method may be computationally intensive. For this reason, we also propose another bootstrap method which resamples directly the estimated score $s_t(\hat{\alpha}_n, \hat{\beta}_n) \equiv s_{2t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1t}(\hat{\alpha}_n)$. Our simulations show that this method yields wider confidence intervals than the fully optimized bootstrap method. In particular, the fast resampling standard errors have larger mean squared errors compared to the fully optimized standard errors, especially for the smaller sample sizes. This translates into wider confidence intervals for the parameter of interest.

Both methods involve resampling certain functions of the data using the MBB to obtain new indices, which can be described as follows. For a generic time series $\{Z_t : t = 1, \dots, n\}$, let $\ell = \ell_n \in \mathbb{N}$ ($1 \leq \ell < n$) be a block length. Define $B_{t,\ell} = \{Z_t, Z_{t+1}, \dots, Z_{t+\ell-1}\}$ as the block of ℓ consecutive observations starting at Z_t ($\ell = 1$ corresponds to the standard i.i.d. bootstrap). For simplicity take $n = k\ell$. The MBB draws $k = n/\ell$ blocks randomly with replacement from the set of overlapping blocks $\{B_{1,\ell}, \dots, B_{n-\ell+1,\ell}\}$. Letting I_1, \dots, I_k be i.i.d. random variables distributed uniformly on $\{0, \dots, n - \ell\}$, we have $\{Z_t^* = Z_{\tau_t}, t = 1, \dots, n\}$, where τ_t is defined as $\{\tau_t\} \equiv \{I_1 + 1, \dots, I_1 + \ell, \dots, I_k + 1, \dots, I_k + \ell\}$.

3.1 The fully optimized bootstrap method

The first method we consider requires resampling the contributions to the two (or more) likelihood functions f_{1t} and f_{2t} and then computing $\hat{\alpha}_n^*$ and $\hat{\beta}_n^*$ using these resampled

likelihood functions. More specifically, the bootstrap analogue of $\hat{\alpha}_n$ is given by

$$\hat{\alpha}_n^* = \arg \max_{\alpha \in \mathcal{A}} Q_{1n}^*(\alpha),$$

where $Q_{1n}^*(\alpha) \equiv n^{-1} \sum_{t=1}^n \log f_{1t}^*(\alpha)$, and $f_{1t}^*(\alpha) = f_{1,\tau_t}(\alpha) \equiv f_{1,\tau_t}(X^{\tau_t}, \alpha)$ is a resampled version of $f_{1t}(\alpha) \equiv f_{1t}(X^t, \alpha)$, where the indices τ_t are chosen by the bootstrap. Thus, we resample the functions $f_{1t}(\alpha)$ rather than the data directly. However, when $f_{1t}(\alpha) = f_1(Z_t, \alpha)$ where the function f_{1t} does not depend on t and is a function of $Z_t \equiv (X_t, X_{t-1}, \dots, X_{t-k})'$ for some finite $k \geq 0$, resampling $f_{1t}(\alpha)$ is equivalent to resampling the vector Z_t , i.e. $f_{1t}^*(\alpha) \equiv f_{1t}^*(X^{*\tau_t}, \alpha) = f_1(Z_{\tau_t}^*, \alpha) = f_1(Z_{\tau_t}, \alpha)$.

The second-step bootstrap estimator $\hat{\beta}_n^*$ is obtained as

$$\hat{\beta}_n^* = \arg \max_{\beta \in \mathcal{B}} Q_{2n}^*(\hat{\alpha}_n^*, \beta),$$

where $Q_{2n}^*(\hat{\alpha}_n^*, \beta) \equiv n^{-1} \sum_{t=1}^n \log f_{2t}^*(\hat{\alpha}_n^*, \beta)$, with $f_{2t}^*(\hat{\alpha}_n^*, \beta) = f_{2,\tau_t}(\hat{\alpha}_n^*, \beta) \equiv f_{2,\tau_t}(X^{\tau_t}, \hat{\alpha}_n^*, \beta)$. Thus, we resample the functions $f_{2t}(\alpha, \beta) \equiv f_{2t}(X^t, \alpha, \beta)$ evaluated at $\alpha = \hat{\alpha}_n^*$ using the *same* indices τ_t used in computing $\hat{\alpha}_n^*$. Resampling both functions f_{1t} and f_{2t} with the same set of indices is important because this will preserve the form of dependence between the two functions. In particular, this will guarantee that the bootstrap is able to mimic the dependence in the score vector $s_t(\hat{\alpha}_n, \hat{\beta}_n)$. If instead we used two different sets of indices, say τ_{1t} and τ_{2t} , generated independently of each other, this would induce an independence between f_{1t}^* and f_{2t}^* which would not necessarily exist for the original functions.

3.2 A fast resampling method

Bootstrapping multi-stage extremum estimators can be computationally intensive as this may require solving multiple optimization problems for each resample. For this reason, we also consider a fast resampling method for bootstrapping two-step QMLE which has a

closed form expression and avoids any numerical optimization. To describe this estimator, let

$$\hat{H}_n = n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \beta'} s_{2t}(\hat{\alpha}_n, \hat{\beta}_n), \quad \hat{A}_n = n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{1t}(\hat{\alpha}_n), \quad \text{and} \quad \hat{F}_n = n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{2t}(\hat{\alpha}_n, \hat{\beta}_n).$$

The fast resample two-step QMLE is given by

$$\hat{\beta}_{1,n}^* = \hat{\beta}_n - \hat{H}_n^{-1} n^{-1} \sum_{t=1}^n s_t^*(\hat{\alpha}_n, \hat{\beta}_n),$$

where $s_t^*(\hat{\alpha}_n, \hat{\beta}_n)$ is a resampled version of

$$s_t(\hat{\alpha}_n, \hat{\beta}_n) \equiv s_{2t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1t}(\hat{\alpha}_n),$$

i.e. $s_t^*(\hat{\alpha}_n, \hat{\beta}_n) = s_{\tau_t}(\hat{\alpha}_n, \hat{\beta}_n) \equiv s_{2\tau_t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1\tau_t}(\hat{\alpha}_n)$. Thus, $\hat{\beta}_{1,n}^*$ is a one-step bootstrap QMLE which updates $\hat{\beta}_n$ using the estimated Hessian \hat{H}_n and the bootstrap scores $s_t^*(\hat{\alpha}_n, \hat{\beta}_n)$, evaluated at $\hat{\alpha}_n$ and $\hat{\beta}_n$.

A special case of $\hat{\beta}_{1,n}^*$ is a version of the one-step bootstrap QMLE considered in Gonçalves and White (2004) (henceforth GW(2004)) in the context of one-stage QMLE. In that paper, $\hat{\beta}_n$ does not depend on a first stage estimator $\hat{\alpha}_n$, implying that $s_{2t}(\hat{\alpha}_n, \hat{\beta}_n) = s_{2t}(\hat{\beta}_n)$ and $s_t(\hat{\alpha}_n, \hat{\beta}_n) = s_{2t}(\hat{\beta}_n)$. The only difference with respect to GW(2004) in this case is that our proposal only resamples the estimated scores and does not involve resampling the contributions to the Hessian \hat{H}_n (instead, their one-step bootstrap QMLE involves resampling both; see also Davidson and MacKinnon (2004) and Andrews (2002) who proposed k-step bootstrap methods that resample the contributions to the Hessian and the score vector at each iteration, starting from the original estimators).

$\hat{\beta}_{1,n}^*$ is also related to a fast resampling approach proposed by Armstrong et al. (2014) in the context of a two-step GMM estimator with i.i.d. data, where the first step is

a potentially nonparametric estimator (see also Chen et al. (2003) and Chen and Liao (2015)). In our context, it amounts to

$$\tilde{\beta}_{1,n}^* = \hat{\beta}_n - \hat{H}_n^{-1} n^{-1} \sum_{t=1}^n s_{2t}^* \left(\hat{\alpha}_n^*, \hat{\beta}_n \right).$$

There are two main differences between $\hat{\beta}_{1,n}^*$ and $\tilde{\beta}_{1,n}^*$. First, $\tilde{\beta}_{1,n}^*$ requires computing $\hat{\alpha}_n^*$ whereas $\hat{\beta}_{1,n}^*$ does not. Hence, $\tilde{\beta}_{1,n}^*$ only avoids the computational burden of the second step and not of the first step. Instead, our method avoids computing $\hat{\alpha}_n^*$ for each resample and therefore is computationally more attractive. Second, $\tilde{\beta}_{1,n}^*$ resamples the scores of the second-stage model (evaluated at $(\hat{\alpha}_n^*, \hat{\beta}_n)$), whereas our method involves resampling $s_t(\hat{\alpha}_n, \hat{\beta}_n) = s_{2t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1t}(\hat{\alpha}_n)$. We can think of this vector as an “extended” version of the scores for the second-stage, extended by the term $-\hat{F}_n \hat{A}_n^{-1} s_{1t}(\hat{\alpha}_n)$. This term corrects for the added uncertainty due to the first step. We note that it would not be valid to resample $s_{2t}(\hat{\alpha}_n, \hat{\beta}_n)$ unless $\hat{F}_n = 0$.

4 Bootstrap theory

We discuss two uses of the bootstrap for inference on β using $\hat{\beta}_n$. In Section 4.1 we consider using the bootstrap distribution of $\sqrt{n}(\hat{\beta}_n^* - \hat{\beta}_n)$ (or $\sqrt{n}(\hat{\beta}_{1,n}^* - \hat{\beta}_n)$) to approximate the quantiles of the distribution of $\sqrt{n}(\hat{\beta}_n - \beta_0)$. This approach underlies the construction of percentile bootstrap intervals for β . Even though it does not promise asymptotic refinements, it is empirically attractive as it does not require computing standard errors for $\hat{\beta}_n$. An alternative is to use the bootstrap to estimate standard errors, which we consider in Section 4.2.

4.1 Bootstrap distribution consistency for non-studentized statistics

The first order asymptotic validity of the MBB based on the fully optimized bootstrap two-step QMLE $\hat{\beta}_n^*$ follows by showing that the bootstrap distribution of $\sqrt{n}(\hat{\beta}_n^* - \hat{\beta}_n)$ is consistent for the distribution of $\sqrt{n}(\hat{\beta}_n - \beta_0)$. This result requires we strengthen Assumption A as follows.

Assumption B For some $r > 2$ and some $\delta > 0$,

B.1: (i) $\{s_{1t}(\alpha_0)\}$ is $r + \delta$ -dominated on \mathcal{A} uniformly in t .

(ii) $\{s_{2t}(\alpha_0, \beta_0)\}$ is $r + \delta$ -dominated on $\mathcal{A} \times \mathcal{B}$ uniformly in t .

B.2: $\{V_t\}$ is an α -mixing sequence of size $-\frac{(2+\delta)(r+\delta)}{r-2}$.

B.3: (i) The elements of $\{s_{1t}(\alpha)\}$ are $L_{2+\delta}$ -NED on $\{V_t\}$ of size -1 , uniformly on (\mathcal{A}, ρ) .

(ii) The elements of $\{s_{2t}(\alpha, \beta)\}$ are $L_{2+\delta}$ -NED on $\{V_t\}$ of size -1 , uniformly on $(\mathcal{A} \times \mathcal{B}, \rho)$.

B.4: (i) $n^{-1} \sum_{t=1}^n E(s_{1t}(\alpha_0)) E(s_{1t}(\alpha_0))' = o(\ell_n^{-1})$, where $\ell_n = o(n)$ and $\ell_n \rightarrow \infty$.

(ii) $n^{-1} \sum_{t=1}^n E(s_{2t}(\alpha_0, \beta_0)) E(s_{2t}(\alpha_0, \beta_0))' = o(\ell_n^{-1})$, where $\ell_n = o(n)$ and $\ell_n \rightarrow \infty$.

These assumptions are weaker than those used by GW (2004) (see their Assumption 2.1) and are sufficient to show that a bootstrap CLT applies to $\{s_{2t}^*(\alpha_0, \beta_0) - F_0 A_0^{-1} s_{1t}^*(\alpha_0)\}$, as shown by Gonçalves and de Jong (2003). Assumption B.4 is a restatement of Assumption 2.2 of Gonçalves and White (2002) and is satisfied when the models are correctly specified or when the scores $\{s_{1t}(X^t, \alpha_0)\}$ and $\{s_{2t}(X^t, \alpha_0, \beta_0)\}$ are stationary (this follows if $\{X_t\}$ is a strictly stationary process, the likelihood functions $\{f_{1t}(\alpha)\}$ and $\{f_{2t}(\alpha, \beta)\}$ depend only on a finite number of lags of X_t and there is no time heterogeneity on $\{f_{1t}\}$ and $\{f_{2t}\}$). Under this assumption, the bootstrap covariance matrix of the scaled average of

$\{s_{2t}^*(\alpha_0, \beta_0) - F_0 A_0^{-1} s_{1t}^*(\alpha_0)\}$ converges to J_0 , the correct asymptotic covariance matrix of $\sqrt{n}(\hat{\beta}_n - \beta_0)$.

In the following theorem, and throughout, we let E^* , Var^* and P^* denote the bootstrap expectation, variance and probability measure induced by the resampling, conditional on the original sample.

Theorem 4.1. *Let Assumption A as strengthened by Assumption B hold. If $\ell_n \rightarrow \infty$ and $\ell_n = o(n^{1/2})$, then*

$$\sup_{x \in \mathbb{R}^p} \left| P^* \left(\sqrt{n} \left(\hat{\beta}_n^* - \hat{\beta}_n \right) \leq x \right) - P \left(\sqrt{n} \left(\hat{\beta}_n - \beta_0 \right) \leq x \right) \right| = o_P(1). \quad (1)$$

To prove Theorem 4.1, we verify the conditions of Theorem S2.4 in the online appendix. This result shows the consistency of the bootstrap distribution of a general two-step M-estimator $\hat{\beta}_n^*$ (based on an asymptotically linear first-step estimator $\hat{\alpha}_n^*$) under a set of bootstrap high level conditions (Assumption \mathcal{B}^*). We show that Assumption A strengthened by Assumption B verifies Assumption \mathcal{B}^* .

The first-order asymptotic validity of the fast resampling method is given in the next result. Its proof is a by-product of the proof of Theorem 4.1 and is omitted.

Theorem 4.2. *Under the same assumptions as in Theorem 4.1,*

$$\sup_{x \in \mathbb{R}^p} \left| P^* \left(\sqrt{n} \left(\hat{\beta}_{1,n}^* - \hat{\beta}_n \right) \leq x \right) - P \left(\sqrt{n} \left(\hat{\beta}_n - \beta_0 \right) \leq x \right) \right| = o_P(1).$$

4.2 Bootstrap distribution consistency for studentized statistics

Here, we focus on testing hypotheses about β_0 based on Wald statistics⁵. Let $r : \mathcal{B} \rightarrow \mathbb{R}^q$, where $\mathcal{B} \subset \mathbb{R}^p$, $q \leq p$, be a continuously differentiable function on \mathcal{B} such that $R_0 \equiv \frac{\partial}{\partial \beta'} r(\beta_0)$

⁵The validity of the bootstrap for LM statistics follows from similar arguments. To conserve space and because Wald tests are more popular than LM tests, we only report results for Wald statistics.

has full row rank q . The Wald statistic for testing $\mathcal{H}_0 : r(\beta_0) = 0$ is

$$\mathcal{W}_n = n\hat{r}'_n \left(\hat{R}_n \hat{C}_n \hat{R}'_n \right)^{-1} \hat{r}_n,$$

where $\hat{r}_n = r(\hat{\beta}_n)$, $\hat{R}_n = \frac{\partial}{\partial \beta'} r(\hat{\beta}_n)$ and $\hat{C}_n = \hat{H}_n^{-1} \hat{J}_n \hat{H}_n^{-1}$ is consistent for C_0 . In particular, $\hat{H}_n \equiv n^{-1} \sum_{t=1}^n \frac{\partial^2}{\partial \beta' \partial \beta} \log f_{2t}(\hat{\alpha}_n, \hat{\beta}_n)$ is a consistent estimator of H_0 , and \hat{J}_n is such that $\hat{J}_n - J_0 = o_P(1)$. We define the bootstrap Wald statistic as follows

$$\mathcal{W}_n^* = n(\hat{r}_n^* - \hat{r}_n)' \left(\hat{R}_n^* \hat{C}_n^* \hat{R}_n^{*'} \right)^{-1} (\hat{r}_n^* - \hat{r}_n).$$

For the fully optimized bootstrap method, we set $\hat{r}_n^* = r(\hat{\beta}_n^*)$, $\hat{R}_n^* = \frac{\partial}{\partial \beta'} r(\hat{\beta}_n^*)$ and

$$\hat{C}_n^* = \hat{H}_n^{*-1} \hat{J}_n^* \hat{H}_n^{*-1}, \text{ with} \quad (2)$$

$$\begin{aligned} \hat{H}_n^* &= n^{-1} \sum_{t=1}^n \frac{\partial^2}{\partial \beta' \partial \beta} \log f_{2,\tau_t} \left(X^{\tau_t}, \hat{\alpha}_n^*, \hat{\beta}_n^* \right), \text{ and} \\ \hat{J}_n^* &= k^{-1} \sum_{i=1}^k \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}} \left(\hat{\alpha}_n^*, \hat{\beta}_n^* \right) \right) \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}} \left(\hat{\alpha}_n^*, \hat{\beta}_n^* \right) \right)', \end{aligned}$$

such that $s_{\tau_t}(\hat{\alpha}_n^*, \hat{\beta}_n^*) = s_{2\tau_t}(\hat{\alpha}_n^*, \hat{\beta}_n^*) - \hat{F}_n^* \hat{A}_n^{*-1} s_{1\tau_t}(\hat{\alpha}_n^*)$, where $\hat{A}_n^* = n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{1\tau_t}(\hat{\alpha}_n^*)$ and $\hat{F}_n^* = n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{2\tau_t}(\hat{\alpha}_n^*, \hat{\beta}_n^*)$.

For the fast resampling method, we set $\hat{r}_n^* = r(\hat{\beta}_{1,n}^*)$, $\hat{R}_n^* = \frac{\partial}{\partial \beta'} r(\hat{\beta}_{1,n}^*)$ and

$$\hat{C}_n^* = \hat{H}_n^{-1} \hat{J}_{1,n}^* \hat{H}_n^{-1}, \quad (3)$$

with $\hat{J}_{1,n}^* = k^{-1} \sum_{i=1}^k \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}}(\hat{\alpha}_n, \hat{\beta}_n) \right) \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}}(\hat{\alpha}_n, \hat{\beta}_n) \right)'$, such that $s_{\tau_t}(\hat{\alpha}_n, \hat{\beta}_n) = s_{2\tau_t}(\hat{\alpha}_n, \hat{\beta}_n) - \hat{F}_n \hat{A}_n^{-1} s_{1\tau_t}(\hat{\alpha}_n)$.

For the fast resampling approach proposed by Armstrong et al. (2014), we set $\hat{r}_n^* =$

$r\left(\tilde{\beta}_{1,n}^*\right), \hat{R}_n^* = \frac{\partial}{\partial \beta'} r\left(\tilde{\beta}_{1,n}^*\right)$ and

$$\hat{C}_n^* = \hat{H}_n^{-1} \tilde{J}_{1,n}^* \hat{H}_n^{-1}, \quad (4)$$

where $\tilde{J}_{1,n}^* = k^{-1} \sum_{i=1}^k \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}} \left(\hat{\alpha}_n^*, \hat{\beta}_n \right) \right) \left(\ell^{-1/2} \sum_{t=1}^{\ell} s_{I_{i+t}} \left(\hat{\alpha}_n^*, \hat{\beta}_n \right) \right)'$, such that $s_{\tau_t} \left(\hat{\alpha}_n^*, \hat{\beta}_n \right) = s_{2\tau_t} \left(\hat{\alpha}_n^*, \hat{\beta}_n \right) - \hat{F}_n \hat{A}_n^{-1} s_{1\tau_t} \left(\hat{\alpha}_n^* \right)$.

Next we show that the sampling distribution of \mathcal{W}_n is well approximated by the bootstrap distribution of \mathcal{W}_n^* . For this, we strengthen Assumption B4 as follows.

Assumption B.4' $E\left(s_{1t}\left(\alpha_0\right)\right) = 0$ and $E\left(s_{2t}\left(\alpha_0, \beta_0\right)\right) = 0$ for all $t = 1, \dots, n$.

This is a mild strengthening of Assumption B.4, which is satisfied whenever the score functions are not heterogeneous and/or the data $\{X_t\}$ are stationary.

Theorem 4.3. *Let the assumptions of Theorem 4.1 as strengthened by Assumption B.4' hold. Suppose further that \mathcal{W}_n uses a consistent estimator of J_0 . Then, under \mathcal{H}_0 , if $\ell_n \rightarrow \infty$ and $\ell_n = o\left(n^{1/2}\right)$,*

$$\sup_{x \in \mathbb{R}} \left| P^* \left(\mathcal{W}_n^* \leq x \right) - P \left(\mathcal{W}_n \leq x \right) \right| = o_P(1).$$

The proof of Theorem 4.3 is in the online Appendix; it builds upon results in Gonçalves and White (2004, cf. Theorem 3.1).

4.3 Bootstrap variance consistency

Bootstrap standard errors are often used in applied work as they are easy to compute, avoiding the need to look up complicated formulas. This is especially true in multistage estimation, where these formulas become quickly involved due to the need to keep track of the added uncertainty caused by each estimation stage. Instead, bootstrap standard errors are easily computed by Monte Carlo simulation. For instance, we can approximate the

bootstrap variance estimator of the parameter $\hat{\beta}_{n,j}^*$, $Var^* \left(\hat{\beta}_{n,j}^* \right)$, with the sample variance obtained across B replications of $\hat{\beta}_{n,j}^*$,

$$\frac{1}{B} \sum_{k=1}^B \left(\hat{\beta}_{n,j}^{*(k)} - \overline{\hat{\beta}_{n,j}^{*(k)}} \right)^2, \text{ where } \overline{\hat{\beta}_{n,j}^{*(k)}} = \frac{1}{B} \sum_{k=1}^B \hat{\beta}_{n,j}^{*(k)}.$$

The corresponding bootstrap standard error is the square root of this expression.

The previous results (Theorems 4.1 and 4.2) do not justify by themselves the consistency of bootstrap standard errors based on $\hat{\beta}_n^*$ or $\hat{\beta}_{1,n}^*$. The reason is that convergence in distribution of a random sequence does not imply convergence of moments. For instance, Ghosh et al. (1984) and Shao (1992) give examples of the inconsistency of bootstrap variance estimators for the sample median and smooth functions of sample means, respectively, in the i.i.d. context. Despite this, applied researchers routinely apply the bootstrap when computing standard errors. For a recent paper emphasizing this point, see Hahn and Liao (2020), who show that bootstrap standard errors can lead to conservative inference.

The main goal of this section is to provide a theoretical justification for computing bootstrap standard errors in the context of two-step QMLE with time series data. The current bootstrap literature does not cover this case as it has either assumed i.i.d. data (as in Kato (2011) and Cheng (2015), who prove the consistency of bootstrap variance estimators for one-step M-estimators) or has considered time series least squares estimators, as in Gonçalves and White (2005). No results appear to be available for multistage QMLE, even for i.i.d. data.

Given our previous bootstrap distribution consistency results, a sufficient condition for showing the consistency of the corresponding bootstrap standard errors is to show that a uniform integrability condition holds. In particular, to show that $Var^* \left(\sqrt{n} \hat{\beta}_n^* \right)$ is consistent, it suffices to show that $E^* \left| \sqrt{n} \left(\hat{\beta}_n^* - \hat{\beta}_n \right) \right|^{2+\delta} = O_P(1)$ for some small $\delta > 0$. Because $\hat{\beta}_{1,n}^*$ has a closed form expression, it is substantially easier to verify this condition for the fast resampling method than for the fully optimized bootstrap two-stage QMLE.

For this reason, we focus on this estimator first.

We impose a smoothness condition on the vector of scores $\{s_{1t}\}$ and $\{s_{2t}\}$ which we did not need for bootstrap distribution consistency.

Assumption B.5 (i) $\{s_{1t}(\alpha)\}$ is Lipschitz continuous on \mathcal{A} , a.s.- P with Lipschitz functions $\{L_{1t}\}$ that satisfy the condition $n^{-1} \sum_{t=1}^n E(L_{1t})^{2+\delta} = O(1)$.

(ii) $\{s_{2t}(\alpha, \beta)\}$ is Lipschitz continuous on $\mathcal{A} \times \mathcal{B}$, a.s.- P with Lipschitz functions $\{L_{2t}\}$ satisfying the condition $n^{-1} \sum_{t=1}^n E(L_{2t})^{2+\delta} = O(1)$.

Theorem 4.4. *Under Assumptions A and B strengthened by B.4' and B.5, $Var^* \left(\sqrt{n} \hat{\beta}_{1,n}^* \right) \xrightarrow{P} C_0$.*

Next, we consider the fully optimized bootstrap estimator $\hat{\beta}_n^*$. Similarly to Theorem 4.4, we prove the consistency of $Var^* \left(\sqrt{n} \hat{\beta}_n^* \right)$ by relying on Theorem 4.1 and showing that $E^* \left| \sqrt{n} \left(\hat{\beta}_n^* - \hat{\beta}_n \right) \right|^{2+\delta} = O_P(1)$ for some small $\delta > 0$. Because $\hat{\beta}_n^*$ does not have a closed form expression, this condition is much harder to verify than for $\hat{\beta}_{1,n}^*$ and requires a different method of proof and a different set of assumptions.

Our proof and regularity conditions are inspired by Kato (2011) and Cheng (2015). Kato (2011) shows the consistency of bootstrap moment estimators for M-estimators of parametric models, whereas Cheng (2015) allows for semiparametric models, where the parameter of interest is a finite dimensional parameter, but the model also contains a nuisance parameter that is potentially infinite dimensional. Both papers focus on one-step M-estimators and give sufficient conditions for bootstrap variance consistency that only cover i.i.d. data. Our contribution is to extend those results to two-stage M-estimation with time series data.

To present our regularity conditions, we need to introduce more notation. First, because our proof is based on showing that the unconditional moment of $\left| \sqrt{n} \left(\hat{\beta}_n^* - \hat{\beta}_n \right) \right|^{2+\delta}$ is finite, we need to introduce the joint probability measure $\mathbb{P} = P \times P^*$ that accounts for the two

sources of randomness in $\hat{\beta}_n^*$: the randomness that comes from the original data (and which is described by P) and the randomness that comes from the resampling, conditional on the original sample (described by P^*). In the following, we write \mathbb{E} to denote expected value with respect to \mathbb{P} . Second, to prove that $\mathbb{E} \left| \sqrt{n} \left(\hat{\beta}_n^* - \hat{\beta}_n \right) \right|^{2+\delta} < \infty$, we assume the uniform square integrability of the original two-step QMLE estimator (i.e. we assume that $E \left| \sqrt{n} \left(\hat{\beta}_n - \beta_0 \right) \right|^{2+\delta} < \infty$) and provide regularity conditions that allows us to show that $\mathbb{E} \left| \sqrt{n} \left(\hat{\beta}_n^* - \beta_0 \right) \right|^{2+\delta} < \infty$. We follow Kato (2011) and Cheng (2015) and use an argument that entails bounding the tail probability $\mathbb{P} \left(\left| \sqrt{n} \left(\hat{\beta}_n^* - \beta_0 \right) \right| > u \right)$ for large u . This requires empirical process theory and maximal inequalities. In particular, we impose bounds on the L_p -moments (with $p > 2 + \delta$) of the supremum of certain empirical processes which we describe next.

For any class of functions $\mathcal{F} = \{f_t\}$, define the empirical process $\mathbb{G}_n f = n^{-1/2} \sum_{t=1}^n (f_t - E f_t)$ and let its norm be given by $\|\mathbb{G}_n\|_{\mathcal{F}} = \sup_{f \in \mathcal{F}} |\mathbb{G}_n f|$.

Our assumptions are as follows.

Assumption B.6

- (i) For any $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$, the log-likelihood function $\log f_2(\cdot, \alpha, \beta)$ and its expectation $\bar{Q}_2(\alpha, \beta) \equiv E(\log f_2(X^t, \alpha, \beta))$ are time invariant.
- (ii) There exists a positive constant K independent of β for which for all $\beta \in \mathcal{B}$, $\bar{Q}_2(\alpha_0, \beta) - \bar{Q}_2(\alpha_0, \beta_0) \leq -K \|\beta - \beta_0\|^2$.
- (iii) Given $\eta > 0$, define the class of functions

$$\mathcal{N}_\eta = \{\log f_2(\alpha, \beta) - \log f_2(\alpha, \beta_0) : \|\beta - \beta_0\| \leq \eta, (\alpha, \beta) \in \mathcal{A} \times \mathcal{B}\}.$$

Then, for some $p > 2 + \delta$, and every $\eta > 0$, there exists a positive constant K such

that

$$\left[E \left(\|\mathbb{G}_n\|_{\mathcal{N}_\eta}^p \right) \right]^{1/p} \leq K\eta. \quad (5)$$

(iv) The functions $\{\log f_2(\alpha, \beta)\}$, $\{\frac{\partial}{\partial \alpha'} \log f_2(\alpha, \beta)\}$ and $\{\frac{\partial}{\partial \alpha \partial \alpha'} \log f_2(\alpha, \beta)\}$ satisfy a Lipschitz continuity condition on $\mathcal{A} \times \mathcal{B}$, a.s.- P with Lipschitz functions $\{L_t\}$, $\{L_{1t}\}$ and $\{L_{2t}\}$ such that $E|L_t|^p < \infty$, $E\left(|L_{1t}|^{\frac{\varepsilon}{\varepsilon-1}p}\right) < \infty$ and $E\left(|L_{2t}|^{\frac{\varepsilon}{\varepsilon-1}p}\right) < \infty$, respectively, for $p > 2 + \delta$ as in (iii) and for some $\varepsilon > 1$.

(v) The first-step estimator $\hat{\alpha}_n$ and its bootstrap analog $\hat{\alpha}_n^*$ are such that

$$E \left| \sqrt{n} (\hat{\alpha}_n - \alpha_0) \right|^{3\varepsilon p} = O(1) \text{ and } \mathbb{E} \left| \sqrt{n} (\hat{\alpha}_n^* - \hat{\alpha}_n) \right|^{3\varepsilon p} = O(1), \quad (6)$$

where $\varepsilon > 1$ and $p > 2 + \delta$ are as defined in (iv).

(vi) $\sup_n E \left| \sqrt{n} (\hat{\beta}_n - \beta_0) \right|^{2+\delta} < \infty$.

Assumption B.6 (i) assumes that the log-likelihood functions f_{2t} and the population criterion function $\bar{Q}_2(\alpha, \beta) \equiv E(\log f_2(X^t, \alpha, \beta))$ are time invariant (the latter will follow from the first under stationarity of $\{X_t\}$).

To understand Assumptions B.6(ii)-(iv), suppose that β is a scalar and consider for example the generated regressor problem. In particular, suppose the following regression model $y_t = \beta q_t + e_t$, where the regressor $q_t = w_t \alpha$ is latent because we do not observe α . If $x_t = w_t \alpha + \eta_t$, where x_t and w_t are observed, we can obtain a consistent estimator of α by running an OLS regression of x_t on w_t . This yields a generated regressor $\hat{q}_t = w_t \hat{\alpha}$, which we can then use to obtain a consistent estimator of β . In terms of our notations, and letting $e_t | w_t, x_t \sim \text{i.i.d. } N(0, \sigma^2)$, the log-likelihood function underlying the second step is given by $\log f_{2t}(\alpha, \beta) = -\log(2\pi\sigma^2) - \frac{1}{2\sigma^2} (y_t - w_t \alpha \beta)^2$ and $\bar{Q}_2(\alpha_0, \beta) = -\log(2\pi\sigma^2) - \frac{1}{2\sigma^2} E(y_t - w_t \alpha_0 \beta)^2$. This verifies condition B.6(i). To verify Condition B.6(ii), note that by a second-order Taylor expansion of $\bar{Q}_2(\alpha_0, \beta)$ around β_0 , we get $\bar{Q}_2(\alpha_0, \beta) = \bar{Q}_2(\alpha_0, \beta_0) +$

$\frac{\partial}{\partial \beta} \bar{Q}_2(\alpha_0, \beta_0) (\beta - \beta_0) + \frac{1}{2} \frac{\partial^2}{\partial \beta^2} \bar{Q}_2(\alpha_0, \ddot{\beta}) (\beta - \beta_0)^2$, where $\ddot{\beta}$ lies between β and β_0 . Since (α_0, β_0) maximizes $\bar{Q}_2(\alpha, \beta)$, $\frac{\partial}{\partial \beta} \bar{Q}_2(\alpha_0, \beta_0) = 0$, implying that $\bar{Q}_2(\alpha_0, \beta) = \bar{Q}_2(\alpha_0, \beta_0) + \frac{1}{2} \frac{\partial^2}{\partial \beta^2} \bar{Q}_2(\alpha_0, \ddot{\beta}) (\beta - \beta_0)^2$. So, the condition will be satisfied if we can bound $\frac{\partial^2}{\partial \beta^2} \bar{Q}_2(\alpha_0, \beta)$ by a negative constant $-K$, for any value of β . This is true if $\bar{Q}_2(\alpha_0, \beta)$ is a quadratic function of β , as in the generated regressor problem. This is a strong condition since it imposes a global restriction on $\bar{Q}_2(\alpha_0, \beta)$, but it is crucial for controlling the tail probability $\mathbb{P}\left(\left|\sqrt{n}\left(\hat{\beta}_n^* - \beta_0\right)\right| > u\right)$, as Kato (2011) and Cheng (2015) note. A similar condition is also used by Nishiyama (2010) to prove the moment convergence of the original M-estimator.

Assumption B.6(iii) (cf. equation (5)) is a high level condition on the empirical process \mathbb{G}_n . Cheng (2015) relies on a similar assumption to show the consistency of bootstrap one-step moment estimators of any order $p \geq 1$ for i.i.d. data. This so-called L_p -maximal inequality condition can be verified under more primitive conditions involving the structure of the function class \mathcal{N}_η , e.g. Cheng (2015) shows that (for one-step M-estimators with i.i.d. data) condition (5) is implied by a finite uniform entropy integral condition, which is verified when the functions in \mathcal{N}_η are Lipschitz continuous. Our Assumption B.6(iii) adapts Cheng (2015)'s high level condition to the two-step QMLE context, where we use the function class \mathcal{N}_η .

To verify Condition B.6(iii) in the context of the generated regressor problem, note that $\log f_{2t}(\alpha, \beta) - \log f_{2t}(\alpha, \beta_0) = (w_t \alpha (\beta - \beta_0)) (2y_t - w_t \alpha (\beta + \beta_0))$. Given that β and $\beta_0 \in B$, a compact subset of \mathbb{R} , it follows that $|\log f_{2t}(\alpha, \beta) - \log f_{2t}(\alpha, \beta_0)| \leq C \|w_t\| (|y_t| + |w_t|) |\beta - \beta_0|$, for some sufficiently large constant C . Hence, the condition follows by Remark 2.3 of Kato (2011) (relying for instance on Lemma 2.14.1 of van der Vaart and Wellner (1996)) provided that $E \|w_t\| (|y_t| + |w_t|)^p < \infty < \infty$, for some $p > 2 + \delta$.

Assumptions B.6(iv) and (v) are new to the two-step estimators we treat here. Part (iv) imposes a Lipschitz continuity condition on the score and the Hessian of $\log f_{2t}(\alpha, \beta)$ with respect to α . Note that in the example of OLS with generated regressor, condition

(iv) follows provided that $E\left(|w_t^2 + y_t^2|^{\frac{\varepsilon}{\varepsilon-1}p}\right) < \infty$, for $p > 2 + \delta$ as in Assumption B.6(iii) and for some $\varepsilon > 1$.

Part (v) imposes uniform integrability conditions on the first-step estimator $\hat{\alpha}_n$ and its bootstrap analog $\hat{\alpha}_n^*$. Similarly, part (vi) assumes the uniform integrability condition on $\hat{\beta}_n$. These high level conditions could be derived from more primitive conditions such as the ones used by Cheng (2015) or Kato (2011), but we prefer to state them as high level conditions since our focus is on the second-step bootstrap estimator $\hat{\beta}_n^*$. It is nevertheless interesting to note that stronger than usual uniform square integrability conditions on the first-step estimators are imposed in order to verify the uniform square integrability condition on the second stage bootstrap estimator $\hat{\beta}_n^*$. In particular, we require the existence of a bit more than six moments for $\sqrt{n}(\hat{\alpha}_n - \alpha_0)$ and its bootstrap analogue. This is three times more than the number of moments for the second-step estimators $\hat{\beta}_n$ and $\hat{\beta}_n^*$. When the log-likelihood function $\log f_{2t}$ is quadratic in α and β , Assumption B.6 (v) can be weakened to $E|\sqrt{n}(\hat{\alpha}_n - \alpha_0)|^{2\varepsilon p} = O(1)$ and $\mathbb{E}|\sqrt{n}(\hat{\alpha}_n^* - \hat{\alpha}_n)|^{2\varepsilon p} = O(1)$. Under these assumptions, we can prove the following theorem.

Theorem 4.5. *Suppose Assumptions A and B strengthened by Assumption B6 holds. Then, for some $\delta > 0$, $\sup_n \mathbb{E}\left|\sqrt{n}\left(\hat{\beta}_n^* - \hat{\beta}_n\right)\right|^{2+\delta} < \infty$, implying that $\text{Var}^*\left(\sqrt{n}\hat{\beta}_n^*\right) \xrightarrow{P} C_0$.*

5 Monte Carlo simulations

We here assess the properties of the bootstrap approximation proposed in Sections 3 and 4. We do so via detailed and realistic Monte Carlo simulations and we start by describing the design of the study. We consider a bivariate copula-based model. Each variable's marginal distribution is an AR(1)-GARCH(1,1) with standardized Student's t errors:

$$\begin{aligned} X_{it} &= \phi_{0,i} + \phi_{1,i}X_{i,t-1} + \varepsilon_{it}, \quad \text{where } \varepsilon_{it} = \sigma_{it}\eta_{it}, \\ \sigma_{it}^2 &= \tilde{\omega}_i + \tilde{\alpha}_i\varepsilon_{i,t-1}^2 + \tilde{\beta}_i\sigma_{i,t-1}^2 \quad \text{and } \eta_{it} \sim iid t(0, 1, \nu_i). \end{aligned}$$

We examine the case where the amount of dependence between the two variables X_1 and X_2 is related to the Clayton copula, with parameter $\beta = 1$, which roughly implies linear correlation of 0.5 (see Nelsen (1999) for more on this copula). We use parameters similar to those found in applied work (see, e.g., Oh and Patton (2012)): for $i = 1, 2$, $[\phi_{0,i}, \phi_{1,i}] = [0, 0.4]$, $[\tilde{\omega}_i, \tilde{\alpha}_i, \tilde{\beta}_i] = [0.05, 0.05, 0.9]$, and $\nu_1 = \nu_2 = \nu$, such that $\nu \in \{6, 10\}$.

Thus, we have two DGPs, which differ only in the value of the Student's t parameter ν , which control the thickness of the tail of the distributions. Note that when $\nu \rightarrow \infty$, this implies that $\eta \sim N(0, 1)$. We generate repeated trials of length $n \in \{200, 500, 2500\}$ from these processes and conduct inference based on a misspecified model that assumes $\phi_{1,i} = 0$ (i.e. we estimate a constant mean-GARCH(1,1)-Student's t -Clayton copula model for each trial).

It is easy to see that our bivariate density models constructed using copulas can be partitioned into elements relating only to a marginal distribution and elements that relate only to the copula. As pointed out by Joe (2005) and Patton (2006), when such a partition is not possible, the familiar one-stage maximum likelihood estimator is the natural estimator to employ. However, when this partitioning is possible as in our simulation setting, great computational savings may be achieved by employing a multi-stage estimator. Therefore, in the following we consider the multi-stage maximum likelihood estimator (MSMLE).

Our estimation steps are:

1. Estimate the conditional mean of X_i , $i = 1, 2$, using the sample mean $n^{-1} \sum_{t=1}^n X_{it}$. Obtain the estimated residuals $\hat{\epsilon}_{it} = X_{it} - n^{-1} \sum_{t=1}^n X_{it}$.
2. Estimate the conditional variance parameter using QML (with normal log-likelihood) and the residuals from step 1, conditioning on the realizations for $t = 1$. Obtain the estimated standardized residuals $\hat{\eta}_{it}$.
3. Estimate ν using ML and the standardized residuals from step 2. Obtain the esti-

mated probability integral transforms (PITs) \hat{G}_{it} .

4. Estimate β using ML and the estimated PITs from step 3.

More generally in a multivariate d -dimensional application (with $d \geq 2$) there are a total of $3d + 1$ estimation steps: three steps for each marginal distribution, and 1 step for the copula. With misspecification as in our context, the scores are generally correlated, hence justifying the use of blockwise bootstrap methods. Furthermore, for misspecified models, the QMLE is generally inconsistent for the copula true parameter β , instead confidence intervals for the copula pseudo-true parameter β_0 are obtained. We evaluate by simulation, the pseudo-true parameter using 1 million observations.

To generate the bootstrap data, we use the MBB. The number of Monte Carlo trials is 1,000 with $B = 999$ bootstrap replications each. We implement three resampling methods: the fully optimized bootstrap procedure B1, the fast resampling approach B2 and the fast resampling approach proposed by Armstrong et al. (2014) B3. To select the block size, we rely on the asymptotic equivalence between the MBB and the Bartlett kernel variance estimators, and choose ℓ equal to the bandwidth chosen by Andrews's automatic procedure for the Bartlett kernel.

We consider three types of confidence intervals for the copula pseudo-true parameter β_0 : asymptotic normal theory-based confidence intervals, computed by using the quantile of the standard normal distribution, bootstrap percentile confidence intervals, and bootstrap percentile- t confidence intervals, which use the bootstrap methods (B1, B2 and B3) to compute critical values for the non-studentized and studentized statistics based on $\hat{\beta}_n$, respectively. The asymptotic normal theory-based confidence interval for β_0 is given by $\hat{\beta}_n \pm 1.96 \cdot \widehat{SE}(\hat{\beta}_n)$, where $\widehat{SE}(\hat{\beta}_n)$ is a consistent estimator of $SE(\hat{\beta}_n) = \sqrt{Var(\hat{\beta}_n)}$. Four choices are used to compute $\widehat{SE}(\hat{\beta}_n)$. For our first choice of $\widehat{SE}(\hat{\beta}_n)$, we use the multi-stage maximum likelihood (MSML) standard errors estimator as described in detail in Section 3.1.1 in Patton (2012) (cf. equation (41)). Then, we construct a MSMLE

variance, asymptotic normal theory-based confidence interval for β_0 by using $\widehat{SE}(\hat{\beta}_n) = \widehat{SE}^{\text{MSML}}(\hat{\beta}_n)$, where $\widehat{SE}^{\text{MSML}}(\hat{\beta}_n)$ is the estimated MSML standard error of $\hat{\beta}_n$ (which has a sandwich form).

In our second, third and fourth choices of $\widehat{SE}(\hat{\beta}_n)$, we use the bootstrap approaches B1, B2 and B3. In particular, a fully optimized bootstrap procedure variance, asymptotic normal theory-based confidence interval for β_0 can be obtained with $\widehat{SE}(\hat{\beta}_n) = \widehat{SE}^{\text{B1}}(\hat{\beta}_n)$, where $\widehat{SE}^{\text{B1}}(\hat{\beta}_n)$ is the estimated standard error of $\hat{\beta}_n$ based on B1. Similarly, fast resampling procedure variance, asymptotic normal theory-based confidence intervals for β are obtained by using $\widehat{SE}(\hat{\beta}_n) = \widehat{SE}^{\text{B2}}(\hat{\beta}_n)$ and $\widehat{SE}(\hat{\beta}_n) = \widehat{SE}^{\text{B3}}(\hat{\beta}_n)$, where $\widehat{SE}^{\text{B2}}(\hat{\beta}_n)$ and $\widehat{SE}^{\text{B3}}(\hat{\beta}_n)$ are the estimated standard error of $\hat{\beta}_n$ based on B2 and B3, respectively. Note that, the standard errors $\widehat{SE}^{\text{B1}}(\hat{\beta}_n)$, $\widehat{SE}^{\text{B2}}(\hat{\beta}_n)$ and $\widehat{SE}^{\text{B3}}(\hat{\beta}_n)$ are obtained by computing the statistics $\sqrt{\frac{1}{B} \sum_{i=1}^B (\hat{\beta}_n^{*(i)} - \bar{\hat{\beta}}_n^*)^2}$ with $\bar{\hat{\beta}}_n^* = \frac{1}{B} \sum_{j=1}^B \hat{\beta}_n^{*(j)}$, where B is the number of bootstrap replications.

The second set of intervals we consider are bootstrap percentile confidence intervals, which are very simple to compute since they avoid the need to explicitly compute standard errors. For each resampling approach (B1, B2 and B3), a symmetric bootstrap percentile confidence intervals for β_0 is given by $\hat{\beta}_n \pm p_{95}^*$, where p_{95}^* is the 95% quantile of the bootstrap distribution of $|\hat{\beta}_n^* - \hat{\beta}_n|$.

Finally, we also consider bootstrap percentile- t confidence intervals. For each resampling approach, a symmetric bootstrap percentile- t confidence intervals for β_0 is given by $\hat{\beta}_n \pm q_{95}^* \widehat{SE}^{\text{MSML}}(\hat{\beta}_n)$, where q_{95}^* is the 95% quantile of the bootstrap distribution of $\sqrt{n}|\hat{\beta}_n^* - \hat{\beta}_n|/\sqrt{\hat{C}_n^*}$, with \hat{C}_n^* as in (2), (3) and (4), for B1, B2 and B3, respectively.

Table 1 reports the coverage rates and lengths of 95% confidence intervals of the copula pseudo-true parameter β_0 for the two DGPs, respectively. Results in Table 1 are not too sensitive to the value of the Students t parameter ν . For a given sample size, all intervals have approximately the same coverage rate, with only small differences among them. Both

fast resampling procedures B2 and B3, and the MSML approach perform well even for the small sample size $n = 200$. Indeed, the evidence of presence of serial correlation in the scores is confirmed by the average value of the block sizes chosen by Andrews (1991) method, which is equal to 3.90 in our simulations. However, as Table 1 suggests, there are notable differences among the different methods when considering their confidence interval lengths. This table clearly shows that for all DGP's, the intervals based on B1 (either using the CLT-based or the bootstrap percentile and/or percentile- t approach) tend to display shorter intervals for the smaller sample sizes compared to CLT-MSML, B2 and the B3 intervals.

Table 1: Coverage Rates and length confidence intervals of nominal 95% intervals for β_0

	CLT						Percentile			Percentile- t			
	CLT-MSML	CLT-B1	CLT-B2	CLT-B3	B1	B2	B3	B1	B2	B3	B1	B2	B3
Coverage Rates													
				$\nu = 6$									
$n = 200$	90.70	91.40	90.60	90.70	90.80	90.00	90.30	91.70	91.80	91.90	91.70	91.80	91.90
$n = 500$	92.50	92.70	92.40	92.90	92.50	92.80	92.80	93.60	93.00	93.40	93.60	93.00	93.40
$n = 2500$	93.90	94.20	94.00	94.00	94.40	94.20	94.10	94.90	94.40	94.80	94.90	94.40	94.80
				$\nu = 10$									
$n = 200$	89.50	89.30	90.10	90.20	89.60	89.50	89.40	90.20	90.30	90.40	90.20	90.30	90.40
$n = 500$	92.20	92.60	92.50	92.70	92.30	92.70	92.50	93.00	93.00	93.10	93.00	93.00	93.10
$n = 2500$	94.20	94.90	94.50	94.20	94.60	94.70	94.80	95.50	94.70	95.60	95.50	94.70	95.60
Length confidence intervals													
				$\nu = 6$									
$n = 200$	0.66	0.61	0.65	0.63	0.61	0.68	0.64	0.60	0.68	0.64	0.61	0.68	0.64
$n = 500$	0.42	0.42	0.41	0.41	0.42	0.43	0.42	0.41	0.44	0.49	0.42	0.44	0.49
$n = 2500$	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19
				$\nu = 10$									
$n = 200$	0.70	0.61	0.70	0.66	0.62	0.73	0.63	0.61	0.73	0.68	0.62	0.73	0.68
$n = 500$	0.43	0.41	0.43	0.43	0.42	0.45	0.42	0.42	0.45	0.44	0.42	0.45	0.44
$n = 2500$	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.20	0.19	0.19	0.20

Notes: CLT-MSML, CLT-B1, CLT-B2, and CLT-B3 -intervals based on the normal using estimated standard error based on the MSML, the B1, the B2, and the B3, respectively; B1 bootstrap intervals based on the fully optimized procedure, B2 bootstrap intervals based on the fast resampling procedure, B3 bootstrap intervals based on the fast resampling approach proposed by Armstrong et al. (2014). 1,000 Monte Carlo trials with 999 bootstrap replications each. Pseudo-true parameters were calculated using 1 million observations. For $\nu = 6$, $\beta_0 = 0.99$, and $\beta_0 = 1.07$ when $\nu = 10$.

Table 2: Comparison of standard errors estimation of $\hat{\beta}_n$

	$n = 200$			$n = 500$			$n = 2500$					
	MSML	B1	B2	B3	MSML	B1	B2	B3	MSML	B1	B2	B3
$\nu = 6$												
(MSE of estimated SE) · 10 ³	1.31	0.82	1.26	0.93	0.28	0.16	0.28	0.15	0.05	0.01	0.05	0.02
Ratio of SE over the true value	1.12	1.01	1.14	0.99	1.00	1.00	0.99	0.99	1.00	0.99	1.00	1.01
$\nu = 10$												
(MSE of estimated SE) · 10 ³	7.03	0.96	6.31	3.28	0.42	0.16	0.31	0.19	0.17	0.02	0.15	0.03
Ratio of SE over the true value	1.11	1.01	1.12	1.02	1.00	0.99	1.00	1.00	0.99	1.00	1.00	0.99

Notes: This table provides the MSE of estimated standard errors and the ratio of estimated standard errors over the true value (simulation based) of standard errors. We compute the ratio of standard error of an estimator $\hat{\theta}$ over the true value as: $S^{-1} \sum_{i=1}^S \frac{\hat{SE}_i^j(\hat{\theta})}{SE(\hat{\theta})}$, where S is the number of Monte Carlo replications, $i = 1, \dots, S$, $j = \text{MSML, B1, B2, B3}$ thus $\hat{SE}_i^j(\hat{\theta})$ is the estimated value of the standard error of an estimator $\hat{\theta}$ on the i th Monte Carlo replication obtained by using the method j . $SE(\hat{\theta})$ is defined as $SE(\hat{\theta}) = \sqrt{S^{-1} \sum_{i=1}^S (\hat{\theta}_i - S^{-1} \sum_{s=1}^S \hat{\theta}_s)^2}$, with $\hat{\theta}_i$ the estimated value of the parameter θ on the i th Monte Carlo replication. Similarly, we compute the MSE as $S^{-1} \sum_{i=1}^S (SE_i^j(\hat{\theta}) - SE(\hat{\theta}))^2$. Simulations were done with 1,000 Monte Carlo trials with 999 bootstrap replications each.

Note that all three asymptotic normal theory-based confidence intervals differ only by the way that the estimated standard errors of $\hat{\beta}_n$ have been computed. In order to gain further insight into the "relatively" good performance of these asymptotic normal theory-based confidence intervals in finite samples, we compute the ratio of the estimated standard error over the true value and the mean-square error (MSE) of the estimated standard errors. The results are presented in Table 2. For small sample sizes, on average MSML and B2 overestimate the standard errors, with the ratio of estimated standard error over the true value above 1. For instance, when $n = 200$ and $\nu = 6$, the ratio of estimated standard error over the true value based on MSML and B2 are 1.12 and 1.14 for MSML and B2, respectively, whereas this ratio is 1.01 for B1 and 0.99 for B3. Consequently, the length of confidence intervals based on estimated standard error from MSML and B2 are larger than the one based on B1 and/or B3.

The gains associated with the bootstrap methods can be quite substantial when the main goal of the researcher and/or practitioner is to estimate the standard errors. The results in Table 2 are in favor of the bootstrap particularly for small sample sizes. More specifically, the full resampling method B1 is better than using MSML and/or B2 standard errors. For small samples, the bootstrap method B1 estimates the standard error of the copulas parameter estimator $\hat{\beta}_n$ more precisely than the MSML, B2 and B3 approaches. For large sample sizes, we have approximately the same performance for all three methods. For instance, when $n = 200$ and $\nu = 10$, the MSE of the estimated standard errors of $\hat{\beta}_n$ based on MSML, B1, B2 and B3 are $7.03 \cdot 10^{-3}$, $0.96 \cdot 10^{-3}$, $6.31 \cdot 10^{-3}$ and $3.28 \cdot 10^{-3}$, respectively. Whereas, for $n = 2,500$ and $\nu = 10$, the MSE become $0.17 \cdot 10^{-3}$, $0.02 \cdot 10^{-3}$, $0.15 \cdot 10^{-3}$ and $0.03 \cdot 10^{-3}$, respectively. Thus we see that although all four methods MSML, B1, B2 and B3 are asymptotically equivalent, and the full resampling method B1 may be computationally much more demanding, in small samples, the improved estimates of the standard errors based on B1 may outweigh its computational cost. Overall, the performance of B2 is

comparable to that of MSML, B3 outperforms B2, whereas B1 outperforms B2, B3 and MSML and provides more accurate estimators of the standard errors, especially when the sample size is small.

6 Conclusion

This paper proposes and theoretically justifies bootstrap methods for inference on nonlinear dynamic models that are estimated in two (or more) stages of quasi-maximum likelihood estimation (QMLE). In particular, we show the consistency of the bootstrap distribution of the two-step QML estimator using dependence and heterogeneity conditions similar to those used by Gonçalves and White (2004) for the one-step case. In addition, we also prove the consistency of bootstrap standard errors for two-step QMLE, a result that does not seem to be available even for i.i.d. data. This justifies the standard practice of computing bootstrap standard errors instead of computing analytical standard errors, which quickly becomes cumbersome in the multistage QMLE context. Our simulation results show that intervals based on bootstrap standard errors or bootstrap percentile intervals obtained with the fully optimized method that resamples the log-likelihood functions jointly are shorter on average than intervals based only on asymptotic theory or on the fast resampling method we propose. Thus, although more computationally demanding, the fully optimized bootstrap method has better finite sample properties than the other methods we consider.

Appendix A

In this appendix, we provide a set of primitive assumptions under which the asymptotic theory of the two-step QMLE (consistency and asymptotic distribution) follows. These assumptions extends the assumptions of GW (2004) to the two-step QMLE context and are used to prove our bootstrap results.

Assumption A

A.1: Let (Ω, \mathcal{F}, P) be a complete probability space. The observed data are a realization of a stochastic process $\{X_t : \Omega \rightarrow \mathbb{R}^l, t \in \mathbb{N}\}$, with

$$X_t(\omega) = W_t(\dots, V_{t-1}(\omega), V_t(\omega), V_{t+1}(\omega), \dots),$$

$V_t : \Omega \rightarrow \mathbb{R}^v$, and $W_t : \times_{\tau=-\infty}^{\infty} \mathbb{R}^v \rightarrow \mathbb{R}^l$ is such that X_t is measurable for t .

A.2: (i) The functions $\{f_{1t}(X^t, \alpha)\}$ are such that $f_{1t}(\cdot, \alpha)$ is measurable for each $\alpha \in \mathcal{A}$, where \mathcal{A} is a compact subset of \mathbb{R}^k , $f_{1t}(X^t, \cdot)$ is continuous on \mathcal{A} , a.s.- P , and $f_{1t}(X^t, \cdot)$ is twice continuously differentiable on $\text{int}(\mathcal{A})$, a.s.- P .

(ii) The functions $\{f_{2t}(X^t, \alpha, \beta)\}$ are such that $f_{2t}(\cdot, \alpha, \beta)$ is measurable for each $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$, where \mathcal{B} is a compact subset of \mathbb{R}^p , $f_{2t}(X^t, \cdot, \cdot)$ is continuous on $\Theta = \mathcal{A} \times \mathcal{B}$, a.s.- P , and $f_{2t}(X^t, \cdot, \cdot)$ is twice continuously differentiable on $\text{int}(\Theta)$, a.s.- P .

A.3: (i) α_0 is the unique maximizer of $\bar{Q}_1(\alpha) \equiv \lim_{n \rightarrow \infty} E(Q_{1n}(\alpha))$ on \mathcal{A} .

(ii) β_0 is the unique maximizer of $\bar{Q}_2(\alpha_0, \beta) \equiv \lim_{n \rightarrow \infty} E(Q_{2n}(\alpha_0, \beta))$ on \mathcal{B} .

(iii) $\theta_0 = (\alpha_0, \beta_0)$ is interior to $\Theta = \mathcal{A} \times \mathcal{B}$.

A.4: (i) The functions $\{\log f_{1t}(X^t, \alpha)\}$ and $\{\frac{\partial}{\partial \alpha'} s_{1t}(X^t, \alpha)\}$ are Lipschitz continuous on \mathcal{A} , a.s.- P , where $s_{1t}(X^t, \alpha) \equiv \frac{\partial}{\partial \alpha} \log f_{1t}(X^t, \alpha)$.

(ii) The functions $\{\log f_{2t}(X^t, \alpha, \beta)\}$, $\{\frac{\partial}{\partial \beta'} s_{2t}(X^t, \alpha, \beta)\}$ and $\{\frac{\partial}{\partial \alpha'} s_{2t}(X^t, \alpha, \beta)\}$ are Lipschitz continuous on $\mathcal{A} \times \mathcal{B}$, a.s.- P , where $s_{2t}(X^t, \alpha, \beta) \equiv \frac{\partial}{\partial \beta} \log f_{2t}(X^t, \alpha, \beta)$.

A.5: For some $r > 2$,

(i) The functions $\{\log f_{1t}(X^t, \alpha)\}$, $\{s_{1t}(X^t, \alpha)\}$ and $\{\frac{\partial}{\partial \alpha'} s_{1t}(X^t, \alpha)\}$ are r -dominated on \mathcal{A} uniformly in t .

(ii) The functions $\{\log f_{2t}(X^t, \alpha, \beta)\}$, $\{s_{2t}(X^t, \alpha, \beta)\}$, $\left\{\frac{\partial}{\partial \beta'} s_{2t}(X^t, \alpha, \beta)\right\}$ and $\left\{\frac{\partial}{\partial \alpha'} s_{2t}(X^t, \alpha, \beta)\right\}$ are r -dominated on $\Theta = \mathcal{A} \times \mathcal{B}$ uniformly in t .

A.6: $\{V_t\}$ is an α -mixing sequence of size $-\frac{2r}{r-2}$, with $r > 2$.

A.7: The elements of (i) $\{\log f_{1t}(X^t, \alpha)\}$ and $\left\{\frac{\partial}{\partial \alpha'} s_{1t}(X^t, \alpha)\right\}$ are L_2 -NED on $\{V_t\}$ of size $-\frac{1}{2}$, and those of $\{s_{1t}(X^t, \alpha)\}$ are L_2 -NED on $\{V_t\}$ of size -1 , uniformly on (\mathcal{A}, ρ) , where ρ is a metric on \mathbb{R}^k ;

(ii) $\{\log f_{2t}(X^t, \alpha, \beta)\}$, $\left\{\frac{\partial}{\partial \beta'} s_{2t}(X^t, \alpha, \beta)\right\}$ and $\left\{\frac{\partial}{\partial \alpha'} s_{2t}(X^t, \alpha, \beta)\right\}$ are L_2 -NED on $\{V_t\}$ of size $-\frac{1}{2}$, and those of $\{s_{2t}(X^t, \alpha, \beta)\}$ are L_2 -NED on $\{V_t\}$ of size -1 , uniformly on $(\mathcal{A} \times \mathcal{B}, \rho)$, where ρ is a metric on $\mathbb{R}^k \times \mathbb{R}^p$.

A.8: (i) $A_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{1t}(X^t, \alpha_0) \right)$ is nonsingular and

$B_0 \equiv \lim_{n \rightarrow \infty} Var \left(n^{-\frac{1}{2}} \sum_{t=1}^n s_{1t}(X^t, \alpha_0) \right)$ is positive definite.

(ii) $H_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \beta'} s_{2t}(X^t, \alpha_0, \beta_0) \right)$ is nonsingular,

$$J_0 \equiv \lim_{n \rightarrow \infty} Var \left(n^{-\frac{1}{2}} \sum_{t=1}^n (s_{2t}(X^t, \alpha_0, \beta_0) - F_0 A_0^{-1} s_{1t}(X^t, \alpha_0)) \right)$$

is positive definite, and $F_0 \equiv \lim_{n \rightarrow \infty} E \left(n^{-1} \sum_{t=1}^n \frac{\partial}{\partial \alpha'} s_{2t}(X^t, \alpha_0, \beta_0) \right) < \infty$.

References

- [1] Andrews, D. (2002). Higher-order improvements of a computationally attractive k -step bootstrap for extremum estimators. *Econometrica* 70, 119-162.
- [2] Armstrong, T. B., M. Bertanha and H. Hong (2014). A fast resample method for parametric and semiparametric models. *Journal of Econometrics*, 179, 128–133.

- [3] Cattaneo, M.D., M. Jansson and X. Ma (2019). Two-step estimation and inference with possibly many included covariates. *Review of Economic Studies*, 86, 1095-1122.
- [4] Cochrane, J. H. (2001). *Asset Pricing*. Princeton University Press.
- [5] Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. In: *Handbook of Econometrics*, VI, Elsevier Science.
- [6] Chen, X., and Z. Liao (2015). Sieve semiparametric two-step GMM with weakly dependent data. *Journal of Econometrics*, 189, 163-186.
- [7] Chen, X., O. Linton and I. van Keilegom. Estimation of Semiparametric Models When the Criterion Function is not Smooth, *Econometrica*, 2003, 71, 1591-1608.
- [8] Cheng, G. (2015). Moment Consistency of the Exchangeably Weighted Bootstrap for Semiparametric M estimation. *Scandinavian Journal of Statistics*, 42, 665-684.
- [9] Davidson, R. and J. MacKinnon (1999). Bootstrap testing in nonlinear models. *International Economic Review*, 40, 487-508.
- [10] Engle, R. and K. Sheppard (2001). Theoretical and Empirical Properties of Dynamic Conditional Correlation Multivariate GARCH, mimeo.
- [11] Gallant, A.R. and H. White (1988) *A unified theory of estimation and inference for nonlinear dynamic models*. New York: Basil Blackwell.
- [12] Ghosh, M., Parr, W.C., Singh, K. and G.J. Babu (1984). A note on bootstrapping the sample median. *Annals of Statistics*, 12, 1130-1135.
- [13] Gonçalves, S. and R. de Jong. (2003). Consistency of the stationary bootstrap under weak moment conditions. *Economics Letters* , 81, 273-278.

- [14] Gonçalves, S. and H. White (2002). The Bootstrap of the mean for dependent heterogeneous arrays, *Econometric Theory*, 2002, 18, 1367-1384.
- [15] Gonçalves, S. and H. White (2004). Maximum likelihood and the bootstrap for nonlinear dynamic models. *Journal of Econometrics*, 119, 199-219.
- [16] Gonçalves, S. and H. White (2005). Bootstrap standard error estimates for linear regressions, *Journal of the American Statistical Association*, 100, 471, 970-979.
- [17] Hahn, J. and Z. Liao (2020). Bootstrap standard error estimates and inference, forthcoming in *Econometrica*.
- [18] Hong, H. and O. Scaillet (2004). A fast subsampling method for nonlinear dynamic models, *Journal of Econometrics*, 133, 557-578.
- [19] Kato, K. (2011). A note on moment convergence of bootstrap M estimators. *Statistics & Decisions* 28, 51-61.
- [20] Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17, 1217-1241.
- [21] Joe, H (1995). Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, 94, 401-19.
- [22] La Vecchia, D., A. Moor and O. Scaillet (2020). A higher-order correct fast moving-average bootstrap for dependent data, University of Geneva, mimeo.
- [23] Liu, R. Y. and K. Singh (1992). Moving blocks jackknife and bootstrap capture weak dependence. In: LePage, R., Billiard, L. (Eds.), *Exploring the Limits of the Bootstrap*, Wiley, New York, 224-248.
- [24] Nishiyama, Y. (2010). Moment convergence of M estimators. *Stat. Neerlandica* 64, 505-507.

- [25] Nelsen, R. B. (1999). *An Introduction to Copulas*. Lecture notes in statistics, Springer.
- [26] Newey, W.K., McFadden, D. (1994). Large sample estimation and hypothesis testing. In: Engle, R. McFadden, D. (Eds.), *Handbook of Econometrics, vol IV*, North Holland.
- [27] Oh, D.H. and A.J. Patton (2013). Simulated method of moments estimation for copula-based multivariate models. *Journal of the American Statistical Association*, 108, 689-700.
- [28] Parente, P. and R.J. Smith (2021). Quasi-maximum likelihood and the kernel block bootstrap for nonlinear dynamic models. *Journal of Time Series Analysis*, 42, 377-405.
- [29] Patton, A. J., (2006). Estimation of Multivariate Models for Time Series of Possibly Different Lengths, *Journal of Applied Econometrics* , 21(2), 14-173 .
- [30] Patton, A. J., (2012). Copula Methods for Forecasting Multivariate Time Series, *in Handbook of Economic Forecasting*, eds. G. Elliott and A. Timmermann, (Vol. 2), Oxford: Elsevier.
- [31] Paparoditis, E. and D. Politis (2002). The tapered block bootstrap for general statistics from stationary sequences. *Econometrics Journal*, 5, 131-148.
- [32] Shao, J. (1992). Bootstrap variance estimators with truncation. *Statistics and Probability Letters*, 15, 95-101.
- [33] White, H. (1994). *Estimation, Inference and Specification Analysis*. Econometric Society Monographs 22, Cambridge University Press: Cambridge.
- [34] Wooldridge, J. (1994). Estimation and inference for dependent processes. In: Engle, R. McFadden, D. (Eds.), *Handbook of Econometrics, vol IV*, North Holland.