Econometric Theory, **42**, 2026, 470–500.
doi:10.1017/S0266466624000331

PARAMETERS ON THE BOUNDARY IN PREDICTIVE REGRESSION

GIUSEPPE CAVALIERE
University of Bologna, University of Exeter

ILIYAN GEORGIEV
University of Bologna

EDOARDO ZANELLI
University of Bologna

We consider bootstrap inference in predictive (or Granger-causality) regressions when the parameter of interest may lie on the boundary of the parameter space, here defined by means of a smooth inequality constraint. For instance, this situation occurs when the definition of the parameter space allows for the cases of either no predictability or sign-restricted predictability. We show that in this context constrained estimation gives rise to bootstrap statistics whose limit distribution is, in general, random, and thus distinct from the limit null distribution of the original statistics of interest. This is due to both (i) the possible location of the true parameter vector on the boundary of the parameter space and (ii) the possible non-stationarity of the posited predicting (resp. Granger-causing) variable. We discuss a modification of the standard fixed-regressor wild bootstrap scheme where the bootstrap parameter space is shifted by a data-dependent function in order to eliminate the portion of limiting bootstrap randomness attributable to the boundary and prove validity of the associated bootstrap inference under non-stationarity of the predicting variable as the only remaining source of limiting bootstrap randomness. Our approach, which is initially presented in a simple location model, has bearing on inference in parameter-on-the-boundary situations beyond the predictive regression problem.

1. INTRODUCTION

In this paper, we revisit the well-known problem of bootstrap inference in regressions with parameter space defined by means of smooth inequality constraints. For instance, consider the setup of a regression $y_t = \alpha + \beta x_{t-1} + \varepsilon_t$, where the parameter space for (α, β) is defined by the constraint $\beta \geq 0$. This framework arises when only the possibilities $\beta = 0$ of no predictability (or no first-order Granger causality, generalizable to higher orders), and $\beta > 0$ of sign-restricted predictability, are entertained, and the model is estimated under the constraint

We thank Yixiao Sun (co-editor), two anonymous referees, Rasmus Søndergaard Pedersen and Mervyn Silvapulle for comments. Financial support from the Italian Ministry of University and Research (PRIN 2020 Grant 2020B2AKFW) is gratefully acknowledged. Address correspondence to: Giuseppe Cavaliere, Department of Economics, University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy; email: giuseppe.cavaliere@unibo.it.

© The Author(s), 2025. Published by Cambridge University Press.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

$\beta \in [0, \infty)$. In applications, economic theory is often informative about the direction of predictability, and such information could be used to improve the efficiency of estimators and increase the power of hypotheses tests. A prominent example is provided by predictive regressions for financial returns; see, e.g., Phillips (2014) and the references therein. Interest can then be in testing the hypothesis of no predictability (i.e., $\beta = 0$) by means of a one-sided test, or a special case of this hypothesis (e.g., $\alpha = \beta = 0$), or a hypothesis where the parameter vector may but need not lie on the boundary of the parameter space (e.g., $\alpha + \beta = 0$).

Though in this context the bootstrap is potentially useful, its application is not straightforward if the parameter vector may lie on the boundary of the parameter space; see Andrews (2000). In particular, as we discuss in the following, even in a simple location model where the parameter space is a closed half-line, the cumulative distribution function [cdf] of the parametric bootstrap t-statistic, conditional on the original data, converges weakly to a random cdf, rather than to the target asymptotic distribution of the t-statistic computed from the original data.

Our first contribution is to show that in predictive regressions with parameter values on the boundary, the distribution of fixed-regressor[1] bootstrap statistics, like the t-statistic for $\beta = 0$ in the regression above, may be random in the limit. Limiting randomness may arise in two ways. A first possible source of randomness in the limit bootstrap measure is in the non-stationarity of the regressor, which operates through the random limits of sample product moments. This is hardly surprising, see, e.g., Georgiev et al. (2019). A second potential source of randomness is the location of the parameter vector on the boundary of the parameter space. Invalidity of standard bootstrap schemes when a parameter is on the boundary was initially discussed in Andrews (2000), where a simple location-model example was given; see also Chatterjee and Lahiri (2011). In the context of hypotheses tests in predictive regressions, we revisit Andrews' result and show that, for a general bootstrap scheme, the occurrence or non-occurrence of limiting bootstrap randomness due to the possible location of a parameter on the boundary of the parameter space depends on how well the bootstrap scheme approximates the mutual position of three objects: (i) the boundary, (ii) the parameter set identified by the null hypothesis, and (iii) the true parameter value. Standard bootstrap approximations of this mutual position may not be sufficiently precise, giving rise to complex conditioning in the limit bootstrap distribution, with ensuing bootstrap validity only for special types of statistics.

Our second contribution is to show that certain nonstandard bootstrap schemes, designed to provide a better match with the geometric configuration in the original parameter space, give rise to limit bootstrap distributions where randomness, if present, is not attributable to the boundary value of the parameter vector. This fact allows us to establish bootstrap validity in an "unconditional" sense; see Cavaliere and Georgiev (2020). That is, although randomness of the limiting bootstrap cdf

[1] We focus on "fixed-regressor" bootstrap schemes as they do not require knowledge on the regressor generating process. For instance, and in contrast to recursive-based schemes, they can be applied to both I(0) and I(1) settings.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

prevents the possibility that the bootstrap could mimic the asymptotic distribution of the original statistic, we can show that, in large samples, bootstrap tests and asymptotic tests are correctly sized for essentially the same set of nominal sizes.

Formally, we make use of the following definition, which generalizes the definition of unconditional bootstrap validity given in Cavaliere and Georgiev (2020, p.2555). Let p_n and p_n^* be, respectively, the p-value of an asymptotic test and of its bootstrap analog. Let also

$$C := \{q \in (0,1) : \lim_{n\to\infty} P(p_n \leq q) = q | \mathsf{H}_0\},$$

such that a test rejecting for $p_n \leq q$ (or for $p_n > q$) is correctly sized for nominal significance levels q (resp. $1-q$) with $q \in C$, as $n \to \infty$.[2] If, under the null hypothesis H_0,

$$P(p_n^* \leq q) \to q \text{ for all } q \in \text{int } C, \tag{1.1}$$

where int C denotes the interior of the set C, we say that the bootstrap test based on p_n^* is valid for H_0.[3] The meaning is that the bootstrap test and the asymptotic test are first-order asymptotically equivalent in terms of correct size control. In particular, bootstrap validity for simple hypotheses H_0 characterizes pointwise size control.

Notice that bootstrap validity as in (1.1) is implied by the classic definition of bootstrap consistency, namely that $\sup_{x\in\mathbb{R}} |F_n^*(x) - F(x)| \to_p 0$ for a bootstrap statistic with cdf F_n^* conditionally on the data and an original test statistic with continuous asymptotic cdf F. The converse does not hold; that is, (1.1) does not imply classic bootstrap consistency, see the discussion in Cavaliere and Georgiev (2020).

For test statistics whose asymptotic distribution is continuous, it holds that int $C = (0,1)$ and hence condition (1.1) should hold for all $q \in (0,1)$ for the bootstrap to be valid. Unfortunately, parameter values on the boundary of the parameter space may induce discontinuities in the limiting cdf's, such that not even the exact p-values of the associated tests are asymptotically standard uniform on $[0,1]$. This makes the above weaker version of the validity definition unavoidable.

Finally, we turn to the special case of one-sided tests for the null hypothesis that the parameter vector lies on the boundary of the parameter space, such that the boundary coincides with the parameter set identified by the null hypothesis. This case provides a transparent example of a limit bootstrap cdf which is random only on a subset of its domain. Then, if bootstrap validity is defined as in (1.1), in this case also some standard bootstrap schemes can be proved to be valid.

This paper is related to recent work by Fang and Santos (2019) and Hong and Li (2020). The latter two papers propose nonstandard bootstrap schemes—involving a tuning tool—which correct the inconsistency of "classic" bootstrap methods in settings that cover parameters on the boundary as a special case.

[2] For $q \in$ int C it holds that $P(p_n = q) \to 0$ and rejections for $p_n \leq q$ (or $p_n > q$) are asymptotically equivalent to rejections for $p_n < q$ (or $p_n \geq q$).

[3] Bootstrap unconditional validity as in Cavaliere and Georgiev (2020) is obtained as the special case int $C = (0,1)$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

The main difference from the present contribution is that our theory applies to random limit bootstrap measures. Thus, Fang and Santos (2019) consider bootstrap inference in settings where the target asymptotic distribution, say that of a random element τ, can be thought of as a transformation φ of another random element τ', and both the distribution of τ' and the transformation φ need to be estimated; see also the related works by Dümbgen (1993), Hirano and Porter (2012), Fang (2014) and Chen and Fang (2019). Although Fang and Santos (2019) consider deterministic φ and the unconditional distribution of τ', such that their results are not directly applicable here, their way of conceptualizing the problem remains fruitful also in the case of random φ and random conditional distributions $\tau'|\tau''$(for some random element τ'').We discuss this in Section 5.2.

Our contribution is also related to Hong and Li (2020), who propose a "numerical bootstrap," which is valid in settings where a parameter space can be approximated locally by a cone with vertex at the true value of the parameters; see Geyer (1994) for a detailed discussion of the approximation. Both the approaches in this paper and that by Hong and Li (2020) are connected to the large body of literature considering estimation and inference for constrained M-estimators; see, among others, Geyer (1994), Andrews (1999, 2000), and the references therein. In Section 5.2, we argue that, when applied to a restricted predictive regression, the "numerical bootstrap" of Hong and Li (2020) performs a geometric approximation of the kind we propose, though at the cost of a slower-than-standard convergence rate for the resulting bootstrap estimator.

We present our main idea using first a simple location model for i.i.d. scalar data whose location parameter is constrained to be positive. This is done in Section 2. The predictive regression framework is presented in Section 3; in this section, we also show that the bootstrap limit measure associated with standard fixed-regressor wild bootstrap schemes is random. A new family of bootstrap algorithms and their validity are discussed in Section 4. Results on the validity of one-sided tests, connections to the previous literature, and uniform size control for the bootstrap tests are discussed in Section 5. Section 6 provides simulation evidence, whereas Section 7 concludes. Proofs are collected in the Appendix.

1.1. Notation and definitions

We use the following notation throughout. The spaces of càdlàg functions $[0,1] \to \mathbb{R}^n$, $[0,1] \to \mathbb{R}^{m\times n}$ and $\mathbb{R} \to \mathbb{R}$, all equipped with the respective Skorokhod J_1-topologies, are denoted by $\mathscr{D}_n$, $\mathscr{D}_{m\times n}$ and $\mathscr{D}(\mathbb{R})$, respectively; see Kallenberg (1997, Appendix A2). For $n = 1$, the subscript in $\mathscr{D}_n$ is suppressed. $\mathscr{C}_n(\mathbb{R}^n)$ is the space of continuous functions from $\mathbb{R}^n$ to $\mathbb{R}^n$ equipped with the topology of uniform convergence on compacts. Integrals are over $[0,1]$ unless otherwise stated, Φ is the standard Gaussian cdf, $U_{[0,1]}$ is the uniform distribution on $[0,1]$ and $\mathbb{I}_{\{\cdot\}}$ is the indicator function. If F is a cdf, possibly random, F^{-1} stands for the right-continuous generalized inverse, i.e., $F^{-1}(u) := \sup\{v \in \mathbb{R} : F(v) \le u\}$, $u \in \mathbb{R}$. Unless differently specified, limits are for $n \to \infty$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

With (Z_n, Y_n) and (Z, Y) being random elements of the metric spaces $S_Z \times S_{Y_n}$ and $S_Z \times S_Y$ ($n \in \mathbb{N}$), and defined on a common probability space, we denote by "$Z_n|Y_n \xrightarrow{w}_p Z|Y$" (resp. "$Z_n|Y_n \xrightarrow{w}_{a.s.} Z|Y$") the fact that $E\{g(Z_n)|Y_n\} \to E\{g(Z)|Y\}$ in probability (resp. a.s.) for all bounded continuous functions $g : S_Z \to \mathbb{R}$. When Z_n is a bootstrap statistic and Y_n denotes the original data, we write "$Z_n \xrightarrow{w^*}_p Z|Y$" (resp. "$Z_n \xrightarrow{w^*}_{a.s.} Z|Y$"). Finally, with (Z_n, Y_n) and (Z, Y) possibly defined on different probability spaces, "$Z_n|Y_n \xrightarrow{w}_w Z|Y$" means that $E\{g(Z_n)|Y_n\} \xrightarrow{w} E\{g(Z)|Y\}$ for all bounded continuous functions $g : S_Z \to \mathbb{R}$, see Kallenberg (1997, 2017); we label this fact "weak convergence in distribution". For the special case of scalar random variables Z_n and Z, if the conditional distribution $Z|Y$ is diffuse (non-atomic), weak convergence in distribution is equivalent to the following weak convergence in $\mathscr{D}(\mathbb{R})$:

$$F_n(\cdot|Y_n) := P(Z_n \leq \cdot|Y_n) \xrightarrow{w} P(Z \leq \cdot|Y) =: F(\cdot|Y). \tag{1.2}$$

When Z_n is a bootstrap statistic and conditioning is on the original data, we use the notation "$\xrightarrow{w^*}_w$". For multivariate generalizations, we refer to Cavaliere and Georgiev (2020, Appendix A).

2. PREVIEW OF THE RESULTS IN A LOCATION MODEL

To illustrate the main arguments that will be proposed in the predictive regression framework later, consider as in Andrews (2000) and Cavaliere et al. (2017) the location model

$$y_t = \theta + \varepsilon_t \ (t = 1, ..., n),$$

where the ε_t's are i.i.d.$(0, 1)$ and the parameter space is $\Theta := \{\theta \in \mathbb{R} : \theta \geq 0\}$. Interest is in inference on the true value θ_0 of θ by using the Gaussian QMLE, $\hat{\theta}$. With $l_n(\theta) := -\frac{1}{2}\sum_{t=1}^{n}(y_t - \theta)^2$, we find $\hat{\theta} := \arg\max_{\theta \in \Theta} l_n(\theta) = \max\{0, \bar{y}_n\}$, $\bar{y}_n := n^{-1}\sum_{t=1}^{n} y_t$. If θ_0 is an interior point of Θ, i.e., $\theta_0 > 0$, then $n^{1/2}(\hat{\theta} - \theta_0) \xrightarrow{w} \xi$, $\xi \sim N(0, 1)$. In contrast, if θ_0 is on the boundary of Θ, i.e., $\theta_0 = 0$, the asymptotic distribution of $\hat{\theta}$ is

$$n^{1/2}(\hat{\theta} - \theta_0) = n^{1/2}\hat{\theta} \xrightarrow{w} \ell := \max\{0, \xi\}, \tag{2.1}$$

again with $\xi \sim N(0, 1)$.

The first takeaway of this section is the fact that the location of a parameter on the boundary of the parameter space may induce limiting bootstrap randomness of a kind that invalidates bootstrap inference. To see this, consider in the context of the location model a standard Gaussian parametric bootstrap based on the bootstrap sample

$$y_t^* = \hat{\theta} + \varepsilon_t^*,$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

where the ε_t^*'s are i.i.d.$N(0,1)$ independent of the original data. The bootstrap counterpart of $l_n(\theta)$ is $l_n^*(\theta) := -\frac{1}{2}\sum_{t=1}^n (y_t^* - \theta)^2$, and the usual bootstrap QMLE is $\hat{\theta}^* := \arg\max_{\theta\in\Theta} l_n^*(\theta) = \max\{0, \bar{y}_n^*\}$, $\bar{y}_n^* := \hat{\theta} + \bar{\varepsilon}_n^*$, $\bar{\varepsilon}_n^* := n^{-1}\sum_{t=1}^n \varepsilon_t^*$. Conditionally on the original sample, $\hat{\theta}^*$'s exact distribution is

$$n^{1/2}(\hat{\theta}^* - \hat{\theta}) = n^{1/2}\max\{-\hat{\theta}, \bar{\varepsilon}_n^*\} \sim \max\{-n^{1/2}\hat{\theta}, \xi^*\}|\hat{\theta}, \xi^*|\hat{\theta} \sim \xi \sim N(0,1), \tag{2.2}$$

with associated conditional cdf given by

$$P^*(n^{1/2}(\hat{\theta}^* - \hat{\theta}) \le x) = \Phi(x)\,\mathbb{I}_{\{x \ge -n^{1/2}\hat{\theta}\}},\ x \in \mathbb{R}. \tag{2.3}$$

Now, when θ_0 is an interior point of Θ, $-n^{1/2}\hat{\theta}$ diverges to $-\infty$ in probability and the distribution of $n^{1/2}(\hat{\theta}^* - \hat{\theta})$ given the data converges weakly in probability to the non-random distribution of ξ^*; the bootstrap therefore mimics the $N(0,1)$ asymptotic distribution of the original statistic, the bootstrap distributional approximation is consistent and bootstrap inference is valid in the sense of (1.1), with int $C = (0,1)$. Conversely, when θ_0 is on the boundary of the parameter space, the cdf in (2.3) converges weakly in $\mathscr{D}(\mathbb{R})$ to the random cdf $\Phi(x)\,\mathbb{I}_{\{x \ge -\ell\}}$. In terms of weak convergence in distribution,

$$n^{1/2}(\hat{\theta}^* - \hat{\theta}) \xrightarrow{w^*}_w \ell^*|\ell,\ \ell^* := \max\{-\ell, \xi^*\}, \tag{2.4}$$

where ℓ is distributed as in (2.1) and is independent of ξ^*. The limit distribution in (2.4) is random since its cdf is a stochastic process depending on the conditioning random variable ℓ. Thus, it is distinct from the limit distribution in (2.1), which is unconditional and hence characterized by a non-random cdf. Because the bootstrap limit distribution is random, the bootstrap approximation is not consistent for the limit in (2.1).

As we shall see in Section 4, limiting bootstrap randomness could be of two kinds: "benign," thus not compromising the validity of bootstrap inference in the sense of (1.1), or "malignant," thus invalidating bootstrap inference. In this example, a bootstrap test employing a bootstrap statistic $\tau_n^* := \phi(n^{1/2}(\hat{\theta}^* - \hat{\theta}))$ as the analog of a statistic $\tau_n := \phi(n^{1/2}\hat{\theta})$, where ϕ is a real function, may not be valid in the sense of (1.1) under the null hypothesis $\mathsf{H}_0 : \theta_0 = 0$ even if the function ϕ is continuous, thus implying "malignant" randomness.

To get some further insight into the source of limiting bootstrap randomness, which will be exploited in the next sections, it is useful to notice that the asymptotic distributions in (2.1) and (2.2) can be written as

$$\begin{aligned}
\ell &= \max\{0,\xi\} = \arg\min_{\lambda\in\Lambda} |\lambda - \xi|,\ \Lambda := \{\lambda \in \mathbb{R} : \lambda \ge 0\} \\
\ell^*|\ell &= \max\{-\ell, \xi^*\}|\ell = \left(\arg\min_{\lambda\in\Lambda(\ell)} |\lambda - \xi^*|\right)\Big|\ell,\ \Lambda(\ell) := \{\lambda \in \mathbb{R} : \lambda \ge -\ell\} \\
&= \Lambda - \ell,
\end{aligned}$$

respectively. Hence, bootstrap randomness, and the implied bootstrap invalidity, can be attributed to the fact that in the bootstrap world the limit constraint set for

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

the objective function $|\lambda - \xi^*|$ is the *random* half line $\Lambda(\ell)$ rather than the original fixed half line $\Lambda = \Lambda(0)$. That is, the chosen bootstrap scheme shifts the constraint set by the random variable $-\ell$, which is non-zero with probability $1/2$.

The second takeaway of this section is the fact that bootstrap validity could be restored by offsetting properly the previous shift of the limit constraint set. Specifically, this requires an ad hoc construction of a bootstrap parameter space intended to approximate well the mutual position of the true parameter value and the boundary of the original parameter space.

Consider a bootstrap scheme where the boundary of the bootstrap parameter space Θ^* is chosen in a data-driven way such that the mutual position of θ_0 and the boundary of Θ is well approximated irrespective of whether θ_0 belongs to $\partial\Theta$ or not. To this aim, introduce the half line $\Theta^* := \{\theta : \theta \geq g^*(\hat{\theta})\}$, where $g^*(\theta) := \theta - |\theta|^{1+\kappa}$, $\kappa > 0$, and the associated $\hat{\theta}^* := \arg\max_{\theta \in \Theta^*} l_n^*(\theta) = \max\{g^*(\hat{\theta}), \bar{y}_n^*\}$. The bootstrap QMLE statistic is then given by $n^{1/2}(\hat{\theta}^* - \hat{\theta}) = n^{1/2}\max\{g^*(\hat{\theta}) - \hat{\theta}, \bar{\varepsilon}_n^*\}$. Conditionally on the data, it is distributed as $\max\{n^{1/2}(g^*(\hat{\theta}) - \hat{\theta}), \xi^*\}|\hat{\theta}$, with $\xi^*|\hat{\theta} \sim N(0,1)$. If $\theta_0 = 0$, it then follows that $n^{1/2}(g^*(\hat{\theta}) - \hat{\theta}) = -n^{1/2}\hat{\theta}^{1+\kappa} \xrightarrow{p} 0$, and the bootstrap statistic conditionally on the data converges weakly in probability to ℓ of (2.1). Conversely, if $\theta_0 > 0$ then $n^{1/2}(g^*(\hat{\theta}) - \hat{\theta}) = -n^{1/2}\hat{\theta}^{1+\kappa} \xrightarrow{p} -\infty$ and the bootstrap statistic conditionally on the data converges weakly in probability to the $N(0,1)$ distribution. In both cases, the bootstrap mimics the asymptotic distribution of $n^{1/2}(\hat{\theta} - \theta_0)$ and bootstrap validity in the sense of (1.1) can be seen to be successfully restored.

Remark. In the location model, an appropriate choice of Θ^* simultaneously restores bootstrap validity and removes all the randomness from the limit bootstrap distribution. In the predictive regression framework, we shall conclude that, in order to achieve bootstrap validity, it is essential to remove only the portion of limiting bootstrap randomness that is due to the location of the parameter vector on the boundary of the parameter space. As no other sources of limiting bootstrap randomness exist in the context of the location model, in this section, the previous conclusion simplifies to eliminating all the limiting bootstrap randomness.

Before moving on to predictive regressions, we notice that when a test of $\mathsf{H}_0 : \theta_0 = 0$ against $\mathsf{H}_1 : \theta_0 > 0$ is performed, employing $\tau_n^* := n^{1/2}(\hat{\theta}^* - \hat{\theta})$ as the bootstrap analog of $\tau_n := n^{1/2}\hat{\theta}$, the standard parametric bootstrap with $\Theta^* = \Theta$ is valid in the sense of (1.1); see also Andrews (2000). Specifically, the bootstrap test rejects H_0 when the bootstrap p-value $\tilde{p}_n^* = 1 - p_n^*$ is small, with the following convergence satisfied under the null hypothesis:

$$p_n^* = P^*(\tau_n^* \leq \tau_n) = P^*(\xi^* \leq \tau_n) \xrightarrow{w} \Phi(\ell).$$

A similar convergence is satisfied by the p-value $\tilde{p}_n = 1 - p_n$ of the asymptotic test, with

$$p_n = P(\ell \leq u)|_{u=\tau_n} = \tfrac{1}{2}\mathbb{I}_{\{\tau_n = 0\}} + \Phi(\tau_n)\mathbb{I}_{\{\tau_n > 0\}} = \Phi(\tau_n) \xrightarrow{w} \Phi(\ell).$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

As ℓ is distributed like $\Phi^{-1}(U)\mathbb{I}_{\{U>1/2\}}$, $U \sim U_{[0,1]}$, it follows that $\Phi(\ell)$ is distributed like $\Phi(\Phi^{-1}(U)\mathbb{I}_{\{U>1/2\}})$. As a result, both the bootstrap and the asymptotic test are correctly sized for nominal levels below $1/2$. This phenomenon, whose extensions to predictive regression are discussed in Section 5.1, does not generalize to hypotheses where one-sided tests are not appropriate or straightforward. Therefore, a remedy is necessary for the inference-invalidating limiting bootstrap randomness induced by the location of a parameter on the boundary.

3. THE PREDICTIVE REGRESSION SETUP

Consider the following predictive regression in a triangular array setup:

$$y_t = \theta_1 + \theta_2 x_{n,t-1} + \varepsilon_t, \quad (t = 1, ..., n;\ n = 1, 2, ...), \tag{3.1}$$

where ε_t is a martingale difference sequence [mds] and $x_{n,t}$ is a non-stationary posited predicting variable satisfying the following assumption; see, e.g. Müller and Watson (2008) for references to primitive conditions.

Assumption 1. Let $z_{n,t} := n^{-1/2}\sum_{s=1}^{t}\varepsilon_s$. Then:

(a) $\{\varepsilon_t\}$ is an mds w.r.t. some filtration to which $(x_{n,t}, z_{n,t})$ is adapted, with $E\varepsilon_t^2 = \omega_{zz} \in (0, \infty)$.

(b) a law of large numbers holds as $n \to \infty$:

$$\sum_{t=1}^{n}\begin{pmatrix}\Delta x_{n,t}\\ \Delta z_{n,t}\end{pmatrix}\begin{pmatrix}\Delta x_{n,t} & \Delta z_{n,t}\end{pmatrix} \xrightarrow{p} \Omega := \begin{pmatrix}\omega_{xx} & \omega_{xz}\\ \omega_{xz} & \omega_{zz}\end{pmatrix} > 0.$$

(c) an invariance principle holds in $\mathscr{D}_2$ as $n \to \infty$:

$$(x_{n,\lfloor n\cdot\rfloor}, z_{n,\lfloor n\cdot\rfloor})' \xrightarrow{w} (X, Z)' \sim BM(0, \Omega),$$

a bivariate Brownian motion on $[0, 1]$.

Assumption 1 covers the specification $x_{n,t} = n^{-1/2}x_t$ for an $I(1)$ process x_t driven by an mds that could be contemporaneously correlated with ε_t.[4,5]

Assumption 1 implies that $\sum_{t=1}^{n} x_{n,t-1}\Delta z_{n,t} \xrightarrow{w} \int XdZ$, which need not have a mixed Gaussian distribution because X and Z need not be independent. Nevertheless, it holds that $\sum_{t=1}^{n} x_{n,t-1}(\Delta z_{n,t} - \omega_{xz}\omega_{xx}^{-1}\Delta x_{n,t}) \xrightarrow{w} \int Xd(Z - \omega_{xz}\omega_{xx}^{-1}X)$, which is zero-mean mixed Gaussian with conditional variance $\sigma_e^2\int X^2$, where $\sigma_e^2 := \omega_{zz} - \omega_{xz}^2\omega_{xx}^{-1}$ is the variance of ε_t corrected for $\Delta x_{n,t}$. The bootstrap schemes discussed below all rely on the independence of the processes X and $Z - \omega_{xz}\omega_{xx}^{-1}X$.

[4] As the bootstrap p-values discussed in the paper are invariant to rescaling of the regressor, the normalization of x_t by $n^{-1/2}$ has no practical implication. It is equivalent to specifying a local-to-zero regression coefficient, as is frequent in applications where y_t is a financial return and x_t is non-stationary.

[5] Results under two alternative stochastic specifications of $x_{n,t}$, as a near-unit root and as a stationary process, are given in the accompanying supplement, Section S.1.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Further, Assumption 1 imposes unconditional homoskedasticity for simplicity. As all the bootstrap approximations below are based on "wild" bootstrap schemes, unconditional heteroskedasticity can be accommodated at only a notational cost.

The next assumption specifies the parameter space, say Θ, by means of a smooth inequality constraint.

Assumption 2. The parameter space is $\Theta := \{\theta = (\theta_1, \theta_2)' \in \mathbb{R}^2 : g(\theta) \geq 0\}$, with non-empty boundary $\partial\Theta := \{\theta \in \mathbb{R}^2 : g(\theta) = 0\}$, where $g : \mathbb{R}^2 \to \mathbb{R}$ is continuously differentiable in some neighborhood of the true parameter value $\theta_0 := (\theta_{1,0}, \theta_{2,0})'$ with gradient $\frac{\partial}{\partial \theta'} g(\theta) \neq 0$ in that neighborhood.

In the following, $\dot{g}$ will denote the gradient of the function g evaluated at θ_0.

Assumption 2 generalizes the leading example of the parameter space $\Theta = \mathbb{R} \times [0, \infty)$ obtained by setting $g(\theta) = (0, 1)\theta = \theta_2$. The boundary of Θ then corresponds to the case $\theta_2 = 0$ of no predictability of y_t by $x_{n,t-1}$, whereas the interior of Θ corresponds to the case of sign-restricted predictability.

Interest is in bootstrap inference on a null hypothesis H_0 identifying a set of parameter values that has a nonempty intersection with the boundary of the parameter space. In particular, we consider the following mutual positions of the boundary, the parameter set identified by H_0 and the true value θ_0:

$\mathscr{G}_1$. H_0 is the hypothesis that θ_0 belongs to the boundary: $\mathsf{H}_0 : g(\theta_0) = 0$;
$\mathscr{G}_2$. H_0 is a simple null hypothesis on the boundary: $\mathsf{H}_0 : \theta_0 = \bar{\theta}$, $g(\bar{\theta}) = 0$;
$\mathscr{G}_3$. $\mathsf{H}_0 : h(\theta_0) = 0$, where $\{\theta \in \mathbb{R}^2 : h(\theta) = 0\}$ is not a subset of the boundary $\partial\Theta$, but meets $\partial\Theta$ at a singleton set.

For example, let again $g(\theta) = \theta_2$, such that the parameter space is $\mathbb{R} \times [0, \infty)$ with boundary $\partial\Theta = \mathbb{R} \times \{0\}$. Then the hypothesis of no predictability $\mathsf{H}_0 : \theta_{2,0} = 0$ falls under $\mathscr{G}_1$. The hypothesis $\mathsf{H}_0 : \theta_0 = \bar{\theta} = (0, 0)'$ that y_t is unpredictable with zero mean falls under $\mathscr{G}_2$. Finally, the hypothesis $\mathsf{H}_0 : (1, 1)\theta_0 = \theta_{1,0} + \theta_{2,0} = 0$ falls under $\mathscr{G}_3$ by setting $h(\theta) := (1, 1)\theta$; in this case, the intersection point of the boundary and the parameter set identified by H_0 is $(0, 0)'$, which might, but need not, be the true value under H_0.

3.1. Asymptotic distributions

Let $\hat{\theta}$ be the OLS estimator of $(\theta_1, \theta_2)'$ in the equation

$$y_t = \theta_1 + \theta_2 x_{n,t-1} + \delta \Delta x_{n,t} + e_t, \tag{3.2}$$

subject to the constraint $\hat{\theta} \in \Theta$, i.e., $g(\hat{\theta}) \geq 0$, and where the role of the regressor $\Delta x_{n,t}$ is to ensure that the residuals are asymptotically uncorrelated with the innovations driving $x_{n,t}$, a convenient prerequisite for the bootstrap implementations. The existence, with probability approaching one, of a measurable minimizer of the residual sum of squares in (3.2) over the set Θ can be established in a similar but simpler way than that of its bootstrap counterpart in our detailed proof of Theorem 4.1. Moreover, any two such minimizers are first-order asymptotically

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

equivalent, explaining our usage of "the" associated with the constrained OLS estimator. Specifically, any such minimizer $\hat{\theta}$ satisfies $n^{1/2}(\hat{\theta}-\theta_0) \xrightarrow{w} \ell(\theta_0)$, with $\ell(\theta_0)$ depending on the position of θ_0 relative to the boundary $\partial\Theta$. Thus, $\ell(\theta_0) = \tilde{\ell} := M^{-1/2}\xi$ if $\theta_0 \in \text{int}\Theta := \Theta \setminus \partial\Theta$, where $M := \int \tilde{X}\tilde{X}'$, $\tilde{X} := (1, X)'$, $\xi \sim N\left(0, \sigma_e^2 I_2\right)$ is independent of X, and $\sigma_e^2 > 0$ is the variance of ε_t corrected for $\Delta x_{n,t}$, whereas

$$n^{1/2}(\hat{\theta}-\theta_0) \xrightarrow{w} \ell(\theta_0) = \ell := \underset{\lambda\in\Lambda}{\arg\min}||\lambda - M^{-1/2}\xi||_M,\ \Lambda := \{\lambda \in \mathbb{R}^2 : \dot{g}'\lambda \geq 0\}, \tag{3.3}$$

if $g(\theta_0) = 0$, with $||x||_M := (x'Mx)^{1/2}$ for $x \in \mathbb{R}^2$; see Section 12 in the working paper version of Andrews (1999) or the proof of Theorem 4.1 for the bootstrap counterpart.

The previous asymptotic result is sufficient in order to see that the possibility of having θ_0 at the boundary of the parameter space Θ induces a dichotomy in the limit distribution of $n^{1/2}(\hat{\theta}-\theta_0)$ similar to the dichotomy established in the introductory location-model example. Replicating the constraint set in the limit distribution by means of a bootstrap scheme will be our main concern in what follows.

3.2. Standard bootstrap invalidity

Consider first a fixed-regressor wild bootstrap sample generated as

$$y_t^* = \hat{\theta}_1 + \hat{\theta}_2 x_{n,t-1} + \varepsilon_t^*, \tag{3.4}$$

where $\varepsilon_t^* = \hat{e}_t w_t^*$, $t = 1, ..., n$, with $\hat{e}_t$ the residuals of (3.2) and w_t^* i.i.d. $N(0,1)$, independent of the original data.[6] Then the distribution of $n^{1/2}(\hat{\theta}-\theta_0)$ could be tentatively approximated by the distribution of $n^{1/2}(\hat{\theta}^*-\hat{\theta})$ conditional on the original data, where $\hat{\theta}^*$ is obtained by regressing y_t^* on $(1, x_{n,t-1})'$ under the constraint $\hat{\theta}^* \in \Theta^* = \Theta$, i.e., $g(\hat{\theta}^*) \geq 0$ as for the original estimator; see Andrews (2000)[7].

To motivate the analysis in the next section, it is useful to anticipate some asymptotic properties of $\hat{\theta}^*$ which obtain by specializing Theorem 4.1 below to the fixed-regressor wild bootstrap scheme. For $\theta_0 \in \text{int}\Theta$, it turns out that the bootstrap distribution converges to a conditional version of the limit distribution of $n^{1/2}(\hat{\theta}-\theta_0)$ found earlier:

$$n^{1/2}(\hat{\theta}^*-\hat{\theta}) = n^{1/2}(\tilde{\theta}^*-\hat{\theta}) + o_p(1) \xrightarrow{w^*}_w \tilde{\ell}|M, \tag{3.5}$$

where $\tilde{\theta}^*$ denotes the unconstrained OLS estimator from the bootstrap sample. The limit bootstrap distribution is, therefore, random. The vehicle of limiting bootstrap randomness is the random matrix M, such that limiting bootstrap randomness is

[6]The conclusions do not change if another zero-mean unit-variance distribution with a finite fourth moment is used instead of the standard Gaussian distribution.

[7]Note that the term $\Delta x_{n,t}$ is no longer necessary because $x_{n,t-1}$ and ε_t^* are independent conditionally on the data.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

fully attributable to the stochastic properties of the regressor. Due to the fact that the bootstrap replicates a conditional version of the limit distribution of the original estimator $\hat{\theta}$, bootstrap inference is not invalidated. Rigorous statements in this sense will be provided in Corollary 4.2.

On the other hand, if $\theta_0 \in \partial\Theta$ the bootstrap statistic converges as follows:

$$n^{1/2}(\hat{\theta}^* - \hat{\theta}) \xrightarrow{w^*}_w \ell^* | (M, \ell) \tag{3.6}$$
$$\ell^* := \arg\min_{\lambda \in \Lambda^*_\ell} ||\lambda - M^{-1/2}\xi^*||_M,\ \Lambda^*_\ell := \{\lambda \in \mathbb{R}^2 : \dot{g}'\lambda \geq -\dot{g}'\ell\},$$

where $\xi^* \sim N(0, \sigma_e^2 I_2)$ is independent of (M, ℓ). In contrast with the case $\theta_0 \in \text{int}\Theta$ and additionally to the random matrix M, in (3.6) also the random vector ℓ appears as a vehicle of limiting bootstrap randomness. Moreover, the limit in (3.6) is not a conditional version of the limit of $n^{1/2}(\hat{\theta} - \theta_0)$, inasmuch as Λ^*_ℓ in (3.6) is a random half-plane, rather than the original admissible set Λ of (3.3). The kind of limiting bootstrap randomness introduced by ℓ is similar to the one established in the introductory location model and, in general, it invalidates bootstrap inference. The reason for the discrepancy between Λ and Λ^*_ℓ is that the parameter space of the standard fixed-regressor wild bootstrap does not approximate well the original mutual position of the true value θ_0 and the boundary, unless $g(\hat{\theta}) = 0$. Other, non-standard bootstrap schemes may be designed in order to provide better approximations, at least under the null hypothesis. Under these schemes the possible boundary position of θ_0 is no longer a vehicle of limiting bootstrap randomness, while the role of the random matrix M in the limit bootstrap distribution is maintained. This topic is analyzed in the next section.

4. ASYMPTOTICALLY VALID BOOTSTRAP SCHEMES

In order to unify the discussion of several bootstrap schemes for inference on H_0 under the three cases $\mathscr{G}_1$, $\mathscr{G}_2$, and $\mathscr{G}_3$, consider a bootstrap sample generated as in (3.4) and, more generally than before, a bootstrap OLS estimator $\hat{\theta}^*$ constrained to belong to a bootstrap parameter space Θ^* satisfying the following assumption.

Assumption 3. The bootstrap parameter space is $\Theta^* := \{\theta \in \mathbb{R}^2 : g(\theta) \geq g^*(\hat{\theta})\}$ for some function $g^* : \mathbb{R}^2 \to \mathbb{R}$, which is continuously differentiable in a neighborhood of θ_0 and satisfies $g^*(\theta) \leq g(\theta)$ for $\theta \in \Theta$.

The standard bootstrap considered in Section 3 obtains by setting $g^* = 0$, such that $\Theta^* = \Theta$, the original parameter space. Alternatively, setting $g^* = g$ restricts the bootstrap true value $\hat{\theta}$ to lie on the boundary of the bootstrap parameter space Θ^*.[8] Finally, setting $g^* = g - |g|^{1+\kappa}$ for some $\kappa > 0$ introduces a correction, in the spirit of an alternative to the standard bootstrap mentioned in Andrews (2000, p.403,

[8] As $\hat{\theta} \xrightarrow{p} \theta_0$ and $\dot{g}(\theta_0) \neq 0$, it follows by continuity that $P(\dot{g}(\hat{\theta}) \neq 0) \to 1$, such that, with probability approaching one, $\hat{\theta}$ is not a stationary point of g. In particular, with probability approaching one, $\hat{\theta}$ is not a local minimizer of g, implying that $\hat{\theta} \in \partial\Theta^*$ under $\Theta^* = \{\theta \in R^2 : g(\theta) \geq g(\hat{\theta})\}$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Method two), Fang and Santos (2019, Example 2.1) and Cavaliere et al. (2022), where the bootstrap true value either shrinks to the boundary of the bootstrap parameter space at a proper rate or remains bounded away from this boundary, according to whether θ_0 belongs to the original boundary $\partial\Theta$ or not. Other choices of g^* with the same implication are discussed in Sections 5.2 and 6.3.

To formulate the next theorem, recall M and $\ell(\theta_0)$ introduced in Section 3.1, and let $\xi^*|(M,\ell(\theta_0)) \sim N(0,\sigma_e^2 I_2)$ as in Section 3.2. Let also $D_n = \{y_t, x_{n,t-1}\}_{t=1}^n$ denote the original data. Finally, call a convergence in distribution $Z_n \xrightarrow{w} Z$ and a weak convergence of random distributions $Z_n^*|D_n \xrightarrow{w}_w Z^*|Y$ joint, denoted as $(Z_n,(Z_n^*|D_n)) \xrightarrow{w}_w (Z,(Z^*|Y))$, if $(Z_n, E\{g(Z_n^*)|D_n\}) \xrightarrow{w} (Z, E\{g(Z^*)|Y\})$ for all continuous and bounded real functions g with matching domain.

THEOREM 4.1. *Under a null hypothesis* H_0 *as in* $\mathscr{G}_1$–$\mathscr{G}_3$ *and under Assumptions 1-3, the bootstrap estimator* $\hat\theta^*$ *obtained by regressing* y_t^* *of (3.4) on* $(1, x_{n,t-1})'$ *under the constraint* $\hat\theta^* \in \Theta^*$*, satisfies*

$$(n^{1/2}(\hat\theta-\theta_0),(n^{1/2}(\hat\theta^*-\hat\theta)|D_n)) \xrightarrow{w}_w \left(\ell(\theta_0),(\ell^*(\theta_0)|(M,\ell(\theta_0)))\right),$$

where in the case $g^*(\theta_0) < g(\theta_0)$,

$$\ell^*(\theta_0) = \tilde\ell^* := M^{-1/2}\xi^* \text{ with } \tilde\ell^*|(M,\ell(\theta_0)) = \tilde\ell|M, \tag{4.1}$$

in the sense of a.s. equality of conditional distributions, whereas in the case $g^*(\theta_0) = g(\theta_0)$,

$$\ell^*(\theta_0) = \ell^* := \underset{\lambda\in\Lambda_\ell^*}{\arg\min}||\lambda - M^{-1/2}\xi^*||_M,\ \Lambda_\ell^* := \{\lambda\in\mathbb{R}^2 : \dot g'\lambda \geq (\dot g^* - \dot g)'\ell(\theta_0)\}. \tag{4.2}$$

The following conclusions could be drawn.

(i) Consider first configurations $\mathscr{G}_1$ and $\mathscr{G}_2$ under H_0, such that $g(\theta_0) = 0$. Consider the magnitude order, in probability, of the distance between the bootstrap "true" value $\hat\theta$ and the bootstrap boundary $\partial\Theta^*$ as a measure of how precisely the bootstrap approximates the geometry of $\mathscr{G}_1$ and $\mathscr{G}_2$. As seen previously, the standard bootstrap corresponds to $g^* = 0$ and approximates the geometry up to an exact magnitude order of $n^{-1/2}$, resulting in a situation where the belonging of θ_0 to the boundary contributes to the randomness of the limit bootstrap distribution given by (3.6) and (4.2) via conditioning on the r.v. $\ell(\theta_0) = \ell$. Conversely, bootstrap schemes employing $g^*(\theta_0) = g(\theta_0)$ and $\dot g^* = \dot g$, such that the bootstrap boundary is tangent to the original boundary at θ_0, give rise to approximations of order $o_p(n^{-1/2})$ and all the randomness in the bootstrap limit is due to the properties of the stochastic regressor via the random variable M, as now $\ell^*|(M,\ell) = \ell|M$ in the sense of a.s. equality of random distributions; see (3.3) and (4.2). Moreover, for such schemes the bootstrap mimics a conditional version of the asymptotic distribution of the original estimator: $n^{1/2}(\hat\theta^*-\hat\theta) \xrightarrow{w^*}_w \ell|M$. Examples are the

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

"restricted" bootstrap based on $g^* = g$, which replicates the geometry of the original data under H_0 by putting $\hat{\theta}$ on the bootstrap boundary, and the choice $g^* = g - |g|^{1+\kappa}$ for some $\kappa > 0$. For general g^*, the limit distribution of the resulting bootstrap estimator is random, with randomness depending on both the stochastic regressor and the position of θ_0 relative to the boundary.

(ii) Consider now the case in $\mathscr{G}_3$, such that $g(\theta_0) = 0$ need not, but may hold under H_0. Among the bootstraps considered in (i), the standard one would fail to mimic a conditional version of the original limit distribution if $g(\theta_0) = 0$, while the "restricted" one would fail if $g(\theta_0) > 0$. As an alternative, consider the bootstrap based on $g^* = g - |g|^{1+\kappa}$ for some $\kappa > 0$. If $\theta_0 \in \partial\Theta$, then this choice puts the bootstrap true value $\hat{\theta}$ at the asymptotically negligible distance of $o_p(n^{-1/2})$ from the bootstrap boundary, whereas if $\theta_0 \in \text{int}\Theta$, then $\hat{\theta}$ is bounded away from the bootstrap boundary, in probability. This guarantees bootstrap validity under some regularity conditions, see (ii) in Corollary 4.2 below.

In general, bootstrap validity in the sense of (1.1) can be evaluated through the following corollary of Theorem 4.1 above.

COROLLARY 4.2. *Under the assumptions of Theorem 4.1, a necessary and sufficient condition for the convergence*

$$(n^{1/2}(\hat{\theta} - \theta_0), (n^{1/2}(\hat{\theta}^* - \hat{\theta})|D_n)) \xrightarrow{w}_w (\ell(\theta_0), (\ell(\theta_0)|M)), \tag{4.3}$$

is that: (i) under $\mathscr{G}_1$ and $\mathscr{G}_2$, $g(\theta_0) = g^*(\theta_0)$ and $\dot{g} = \dot{g}^*$; (ii) under $\mathscr{G}_3$, either $g(\theta_0) = g^*(\theta_0)$ and $\dot{g} = \dot{g}^*$, or $g(\theta_0) > \max\{0, g^*(\theta_0)\}$.

Moreover, under (4.3) the bootstrap is valid in the sense of (1.1) for any pair of statistics τ_n, τ_n^* such that, under H_0, $\tau_n = \phi(n^{1/2}(\hat{\theta} - \theta_0)) + o_p(1)$ and $\tau_n^* = \phi(n^{1/2}(\hat{\theta}^* - \hat{\theta})) + o_p(1)$ for some continuous real function ϕ such that $\phi(\ell(\theta_0))$ is well defined a.s.

The class of functions $g^* = g - |g|^{1+\kappa}$ for $\kappa > 0$ satisfies both conditions (i) and (ii) of Corollary 4.2; hence, the ensuing bootstrap inference is valid under all of $\mathscr{G}_1$-$\mathscr{G}_3$. In contrast, the standard bootstrap violates condition (i) and, in general, is asymptotically invalid if $g(\theta_0) = 0$. An exception is when the discrepancy between the original and the bootstrap geometry is offset by the use of a test statistic that takes into account the geometric position of the null hypothesis in the original parameter space. Section 5.1 focuses on this setup.

Remark. The practical implications of Corollary 4.2 depend on the choice of the statistic τ_n and the respective function ϕ, which will typically be a linear $\phi(l) = l' \frac{\partial r}{\partial \theta'}(\theta_0)$ arising from the delta method, with $l \in \mathbb{R}^2$, $\frac{\partial r}{\partial \theta'}(\theta_0) \neq 0$. For instance, if $g(\theta_0) = 0$ and $\phi(\ell)$ depends on ℓ only through $\dot{g}'\ell = \max\{0, \dot{g}' M^{-1/2} \xi\}$, then the cdf of $\phi(\ell)$ will not be continuous. Still, the bootstrap will be valid in the sense of (1.1), meaning that the largest open subset of $[0, 1]$ on which the bootstrap test is correctly sized as $n \to \infty$ coincides with the analogous set for the asymptotic test. This set will be smaller than $(0, 1)$, however. An example is $\tau_n = n^{1/2} g(\hat{\theta})$,

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

$\tau_n^* = n^{1/2}(g(\hat{\theta}^*) - g(\hat{\theta}))$ with $\phi(\ell) = \dot{g}'\ell$, corresponding to a right-sided test of $\mathsf{H}_0 : g(\theta_0) = 0$.

Remark. Bootstrap validity extends readily to statistics where $n^{1/2}(\hat{\theta} - \theta_0)$ is normalized by some $\hat{\Sigma} = \Sigma(M_n) + o_p(1)$, with M_n defined after (5.2) below, for a function $\Sigma : \mathbb{R}^{2\times 2} \to \mathbb{R}^{2\times 2}$ which is continuous on the set of positive definite matrices. Specifically, bootstrap validity holds if, under H_0, $\tau_n = \phi(n^{1/2}\hat{\Sigma}(\hat{\theta} - \theta_0)) + o_p(1)$ and $\tau_n^* = \phi(n^{1/2}\hat{\Sigma}(\hat{\theta}^* - \hat{\theta})) + o_p(1)$, where ϕ is a continuous real function such that $\phi(\Sigma(M)\ell(\theta_0))$ is a.s. well defined.

5. DISCUSSION AND EXTENSIONS

In this section we address the following three issues: (i) the validity of one-sided bootstrap tests; (ii) a discussion of the bootstrap schemes from Corollary 4.2 within the paradigm of some previous works – specifically, Fang and Santos (2019) and Hong and Li (2020); and (iii) uniform bootstrap validity.

5.1. Validity of one-sided standard bootstrap tests

Under case $\mathscr{G}_1$, consider testing $\mathsf{H}_0 : g(\theta_0) = 0$ against the alternative $H_1 : g(\theta_0) > 0$ using a one-sided test and the standard bootstrap, i.e., with $g^* = 0$. For a test statistic of the form $\tau_n := n^{1/2}g(\hat{\theta})$, a bootstrap counterpart is given by $\tau_n^* := n^{1/2}(g(\hat{\theta}^*) - g(\hat{\theta}))$ and the associated one-sided bootstrap test rejects for large values of the bootstrap p-value $p_n^* := P^*(\tau_n^* \le \tau_n)$; equivalently, for small values of $\tilde{p}_n^* := 1 - p_n^*$. As for $\hat{\theta}^*$, also τ_n^* is affected in the limit by extra randomness due to θ_0 being on the boundary. From (4.1), which reduces to (3.3) and (3.6), it follows by the delta method that

$$(\tau_n, (\tau_n^*|D_n)) \xrightarrow{w}_w \left(\dot{g}'\ell, \left(\dot{g}'\ell^*|(M,\ell)\right)\right) = (\dot{g}'\ell, (\max\{-\dot{g}'\ell, \dot{g}'\tilde{\ell}^*\}|(M,\ell))), \tag{5.1}$$

with ℓ, ℓ^*, and $\tilde{\ell}^*$ as previously defined. For τ_n^*, however, the randomness induced by conditioning on ℓ affects the sample paths of the associated random cdf on the negative half-line alone, because $\dot{g}'\ell \ge 0$, and is thus irrelevant for bootstrap tests with nominal size in $(0, \frac{1}{2})$. Put differently, the bootstrap p-values $\tilde{p}_n^*$ are asymptotically uniformly distributed below $\frac{1}{2}$. This follows rigorously from the next generalization of Theorem 3.1 in Cavaliere and Georgiev (2020), the proof being analogous, where conditions for bootstrap validity restricted to a subset of nominal testing levels are formulated.

THEOREM 5.1. *Let there exist a random variable τ and a random element X, both defined on the same probability space, such that the support of τ_n is contained in a finite or infinite closed interval T, and $(\tau_n, F_n^*) \xrightarrow{w} (\tau, F)$ in $\mathbb{R} \times \mathscr{D}(T)$ for $F_n^*(u) := P(\tau_n^* \le u|D_n)$ and $F(u) := P(\tau \le u|X)$, $u \in T$. If the possibly random cdf*

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

F is sample-path continuous on T, then the bootstrap p-value $p_n^* := F_n^*(\tau_n)$ satisfies

$$P(p_n^* \leq q) \to q,$$

for q such that $q \in F(T)$ a.s.

By Theorem 5.1 with $T = [0, \infty)$, which corresponds to the support of τ_n and $\tau := \dot{g}'\ell$, it follows that the standard bootstrap applied to the one-sided statistic τ_n is asymptotically correctly sized for nominal test sizes in $(0, \frac{1}{2})$.

5.2. Fang and Santos (2019) and Hong and Li (2020)

In this section we put the geometric considerations of Section 4 in the perspective of Fang and Santos (2019), and of the numerical bootstrap of Hong and Li (2020). The discussion is often specialized to the case of an affine constraint.

Consider the constrained OLS estimator $\hat{\theta}$ of Section 3.1. Its limit distribution, see (3.3), is the distribution of $\ell(\theta_0) = \varphi_{\theta_0}(M^{-1/2}\xi)$ with

$$\varphi_{\theta_0}(u) = \begin{cases} u & \text{if } g(\theta_0) > 0 \\ \dot{g}_\perp(\dot{g}_\perp' M \dot{g}_\perp)^{-1}\dot{g}_\perp' M u + M^{-1}\dot{g}(\dot{g}' M^{-1}\dot{g})^{-1}\max\{0, \dot{g}'u\} & \text{if } g(\theta_0) = 0 \end{cases},$$

with $u \in \mathbb{R}^2$. By a projection identity, the expression in the second line of the previous display collapses to u whenever $\dot{g}'u \geq 0$. Note that the distribution of $M^{-1/2}\xi$ conditional on M can be estimated consistently by the distribution of the unconstrained bootstrap OLS estimator conditional on the data; that is,

$$n^{1/2}(\tilde{\theta}^* - \hat{\theta}) \xrightarrow{w^*}_w M^{-1/2}\xi | M.$$

One can then ask what properties of an estimator $\hat{\varphi}_n$ of φ_{θ_0} are sufficient for $\hat{\varphi}_n(n^{1/2}(\tilde{\theta}^* - \hat{\theta})) \xrightarrow{w^*}_w \varphi_{\theta_0}(M^{-1/2}\xi)|M$ to hold. Fang and Santos (2019) address this question in the setup of deterministic transformations of non-random limit distributions, instead of the random transformation φ_{θ_0} of the random distribution $M^{-1/2}\xi|M$. Although not directly applicable here, Theorem 3.2 of Fang and Santos (2019) provides the key insight: there should be sufficient uniformity in the convergence of $\hat{\varphi}_n$ to φ_{θ_0}. Consider for instance

$$\begin{aligned} \hat{\varphi}_n(u) = {} & \hat{\dot{g}}_\perp(\hat{\dot{g}}_\perp' M_n \hat{\dot{g}}_\perp)^{-1}\hat{\dot{g}}_\perp' M_n u \\ & + M_n^{-1}\hat{\dot{g}}(\hat{\dot{g}}' M_n^{-1}\hat{\dot{g}})^{-1}\max\{-n^{1/2}|g(\hat{\theta})|^{1+\kappa}, \hat{\dot{g}}'u\},\ u \in \mathbb{R}^2, \end{aligned} \tag{5.2}$$

where $\hat{\dot{g}} = \frac{\partial}{\partial\theta'}g(\hat{\theta})$, $M_n = n^{-1}\sum_{t=1}^n \tilde{x}_t\tilde{x}_t'$ with $\tilde{x}_t = (1, x_{n,t-1})'$, and $\kappa > 0$. Given that $M_n \xrightarrow{w} M$, $\hat{\dot{g}} \xrightarrow{p} \dot{g}$ and $n^{1/2}|g(\hat{\theta})|^{1+\kappa} \xrightarrow{p} \infty\mathbb{I}_{\{g(\theta_0)>0\}}$, it is easily checked that $\tilde{\varphi}_n \xrightarrow{w} \varphi_{\theta_0}$ on $\mathscr{C}_2(\mathbb{R}^2)$, and the convergence of $\hat{\varphi}_n$ is joint with that of $n^{1/2}(\hat{\theta} - \theta_0)$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

and $n^{1/2}(\tilde{\theta}^* - \hat{\theta})$, the latter one given the data. These facts are sufficient to ensure that

$$(n^{1/2}(\hat{\theta} - \theta_0), (\hat{\varphi}_n(n^{1/2}(\tilde{\theta}^* - \hat{\theta}))|D_n)) \xrightarrow{w}_w (\ell(\theta_0), (\ell(\theta_0)|M)),$$

on $\mathbb{R}^4$, essentially as a consequence of the continuous mapping theorem (CMT) and the continuity of the evaluation map from $\mathscr{C}_2(\mathbb{R}^2) \times \mathbb{R}^2$ to $\mathbb{R}^2$. As the previous limit is the same as in Corollary 4.2, it follows that bootstrap inference based on the distribution of $\hat{\varphi}_n(n^{1/2}(\tilde{\theta}^* - \hat{\theta}))$ conditional on the data is valid. Moreover, for the valid bootstrap schemes obtained from Corollary 4.2 with $g^* = g - |g|^{1+\kappa}$, $\kappa > 0$, the bootstrap estimator $\hat{\theta}^*$ satisfies $n^{1/2}(\hat{\theta}^* - \hat{\theta}) = \hat{\varphi}_n(n^{1/2}(\tilde{\theta}^* - \hat{\theta}))$ for affine functions g. It can be concluded that $\hat{\varphi}_n$ of (5.2) implicitly performs the geometric approximation proposed in Section 4, and so does any other estimator of φ_{θ_0} that converges like $\hat{\varphi}_n$.

We now argue that such an estimator of φ_{θ_0} is embedded in the numerical bootstrap of Hong and Li (2020). This ensures the validity of the numerical bootstrap for the predictive regression of interest here, though at the cost of a slower consistency rate of the bootstrap estimator than in Corollary 4.2. Let $s_n \to \infty$ be a sequence such that $n^{-1/2}s_n \to 0$. Hong and Li (2020) propose in their eq. (4.9) a bootstrap estimator $\hat{\theta}^*_{nb}$ where the constraint set of our $\ell(\theta_0)$ (i.e., $\mathbb{R}^2$ if $\theta_0 \in \text{int}\Theta$ and the half-plane Λ if $\theta_0 \in \partial\Theta$), would be estimated by $\Lambda^*_{nb} = \{\lambda \in \mathbb{R}^2 : g(\hat{\theta} + s_n^{-1}\lambda) \geq 0\}$, the implied bootstrap parameter space being $\Theta^*_{nb} = \hat{\theta} + s_n^{-1}\Lambda^*_{nb} = \Theta$. The bootstrap estimator itself, adapted to our setup, could be written as

$$\hat{\theta}^*_{nb} = \underset{g(\theta)\geq 0}{\arg\min} ||s_n(\theta - \hat{\theta}) - M_n^{-1/2}\xi^*_n||_{M_n},$$

where ξ^*_n is a bootstrap variable such that $\xi^*_n \xrightarrow{w^*}_p N(0, I_2)$; e.g., $\xi^*_n = n^{1/2}M_n^{1/2}(\tilde{\theta}^* - \hat{\theta})$. In the simple case of an affine g we find the explicit expression

$$s_n(\hat{\theta}^*_{nb} - \hat{\theta}) = \bar{\varphi}_n(M_n^{-1/2}\xi^*_n),$$

for $\bar{\varphi}_n$ defined similarly to $\hat{\varphi}_n$, with the only difference that in (5.2) the term $n^{1/2}|g(\hat{\theta})|^{1+\kappa}$ is replaced by $s_n g(\hat{\theta})$. As $s_n g(\hat{\theta}) \xrightarrow{p} \infty \mathbb{I}_{\{g(\theta_0)>0\}}$ similarly to $n^{1/2}|g(\hat{\theta})|^{1+\kappa}$, $\kappa > 0$, it follows that $\bar{\varphi}_n$ converges similarly to $\hat{\varphi}_n$. As a result,

$$(n^{1/2}(\hat{\theta} - \theta_0), (s_n(\hat{\theta}^*_{nb} - \hat{\theta})|D_n)) \xrightarrow{w}_w (\ell(\theta_0), (\ell(\theta_0)|M)),$$

ensuring the validity of the numerical bootstrap, though the consistency rate of $\hat{\theta}^*_{nb}$ is $s_n = o(n^{1/2})$ instead of $n^{1/2}$. In contrast, the rate of $n^{1/2}$ would be achieved by our proposed bootstrap estimator, with $n^{1/2}(\hat{\theta}^* - \hat{\theta}) = \bar{\varphi}_n(M_n^{-1/2}\xi^*_n)$, if $\Theta^* = \{\theta \in \mathbb{R}^2 : g(\theta) \geq g^*_n(\hat{\theta})\}$ with $g^*_n = g - n^{-1/2}s_n|g|$ is specified in Assumption 3.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

5.3. Uniformity considerations

In agreement with Chatterjee and Lahiri (2011), Remark 3, the focus in this paper is on pointwise bootstrap validity. For situations where uniform bootstrap validity is of interest, our key takeaways are similar to the literature on non-random limiting bootstrap measures. First, for the null hypothesis $\mathscr{G}_1$ that the true parameter value lies on the boundary of the parameter space, the pointwise-valid bootstrap schemes outlined in Corollary 4.2 display asymptotic rejection probabilities matching the local power of the bootstrap test whenever the true parameter value varies along a sequence that is local to the boundary at the $n^{-1/2}$ rate. This fact is associated with rejection frequencies above the nominal test size along local-to-the-boundary parameter sequences (cf. Fang and Santos, 2019, Remark 3.6). Second, if conservative bootstrap inference along such parameter sequences is desired, it can be achieved for hypotheses $\mathscr{G}_1$– $\mathscr{G}_3$ by adapting the approach of Doko Tchatoka and Wang (2021), and Cavaliere et al. (2024), at the cost of a potential decrease in power.

To illustrate these points, consider a sequence of true parameter values $\theta_n = \theta_0 + n^{-1/2}\vartheta$ such that $g(\theta_0) = 0$ and $\dot{g}'\vartheta = c > 0$ with $g(\theta_n) = n^{-1/2}c + o(n^{-1/2})$. Moreover, let

$$\ell(\vartheta, c) := \vartheta + \underset{\lambda \in \Lambda^c}{\arg\min} ||\lambda - M^{-1/2}\xi||_M,\ \Lambda^c = \{\lambda \in \mathbb{R}^2 : \dot{g}'\lambda + c \geq 0\}, \tag{5.3}$$

and $\ell(0,0) = \ell$ of eq. (3.3). Then, the joint convergence result

$$(n^{1/2}(\hat{\theta} - \theta_0), (n^{1/2}(\hat{\theta}^* - \hat{\theta})|D_n)) \xrightarrow{w}_w (\ell(\vartheta, c), (\ell(0,0)|M)), \tag{5.4}$$

holds for the bootstrap schemes satisfying conditions (i) and (ii) of Corollary 4.2. For a function $r : \mathbb{R}^2 \to \mathbb{R}$, which is continuously differentiable close to θ_0, consider the statistics $\tau_n = n^{1/2}r(\hat{\theta})$ and $\tau_n^* = n^{1/2}(r(\hat{\theta}^*) - r(\hat{\theta}))$, and distinguish among the extreme possibilities $\dot{r} = \alpha\dot{g}$ with $\alpha > 0$, and $\dot{r} = \alpha\dot{g}_\perp$ with $\alpha \neq 0$, where $\dot{r} = \frac{\partial r}{\partial \theta'}(\theta_0)$. The former possibility arises in testing the null hypothesis that θ_0 lies on the boundary (e.g., with $r = g$), whereas the latter one arises when the null is orthogonal to the boundary (e.g., with $r(\theta) = \theta_1$, $\mathsf{H}_0 : \theta_1 = 0$ and $\Theta = \mathbb{R} \times [0, \infty)$). If $\dot{r} = \dot{g}$ and, without loss of generality, $\alpha = 1$, the delta method yields

$$(\tau_n, (\tau_n^*|D_n)) \xrightarrow{w}_w (\max\{0, \dot{g}'M^{-1/2}\xi + c\}, (\max\{0, \dot{g}'M^{-1/2}\xi\}|M)).$$

With $\gamma_M := (\dot{g}'M^{-1}\dot{g})^{-1/2}$, it follows that

$$\begin{aligned} P^*(\tau_n^* \leq \tau_n) &\xrightarrow{w} \pi(c; M, \xi) \\ &:= \tfrac{1}{2}\mathbb{I}_{\{\dot{g}'M^{-1/2}\xi + c < 0\}} + \Phi(\gamma_M(\dot{g}'M^{-1/2}\xi + c))\mathbb{I}_{\{\dot{g}'M^{-1/2}\xi + c \geq 0\}} \\ &> \pi(0; M, \xi), \end{aligned}$$

where $\pi(0; M, \xi) \overset{d}{=} \frac{1}{2}\mathbb{I}_{\{U < 0.5\}} + U\mathbb{I}_{\{U \geq 0.5\}}$, $U \sim U_{[0,1]}$, represents the limit distribution of the bootstrap p-value under the null. The inequality above implies

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

that bootstrap tests rejecting for large bootstrap p -values will exhibit rejection frequencies above the nominal test size.

On the other hand, if $\dot{r} = \alpha \dot{g}_\perp$, it holds that

$$(\tau_n, (\tau_n^* | D_n)) \xrightarrow{w}_w (\dot{r}' \gamma_M^\perp M^{1/2} \xi + \dot{r}' \gamma_M^\perp M \vartheta, (\dot{r}' \gamma_M^\perp M^{1/2} \xi | M)),$$

with $\gamma_M^\perp := \dot{g}_\perp \dot{g}_\perp' (\dot{g}_\perp' M \dot{g}_\perp)^{-1}$, such that the boundary is asymptotically irrelevant. Bootstrap tests of the null that $r(\theta_0) = 0$ could be conservative or liberal according to the sign of $\dot{r}' \dot{g}_\perp \dot{g}_\perp' M \vartheta$. Similar considerations apply whenever $\dot{r}' \dot{g}_\perp \neq 0$.

For situations where liberal tests are not desirable, a possible remedy is suggested next. It involves a continuum of boundaries for the bootstrap parameter space and its implementation requires a discretization of that continuum.

Let $\tilde{\theta}$ be the unrestricted OLS estimator of θ in regression (3.2). For every $s \in I_n := [-|g(\tilde{\theta})|^{1-\mu}, g(\hat{\theta})]$, let $\hat{\theta}_s^*$ be the bootstrap estimator over the parameter space $\Theta_s^* := \{\theta \in \mathbb{R}^2 : g(\theta) \geq s - g(\hat{\theta})^{1+\kappa}\}$, where $\mu \in (0,1)$ and $\kappa > 0$ are fixed. For a continuously differentiable function r, let $p_n^*(s)$ be the p-value of a test based on $\tau_n = n^{1/2} r(\hat{\theta})$ and $\tau_n^* = n^{1/2}(r(\hat{\theta}^*) - r(\hat{\theta}))$. Then

$$\limsup_{n \to \infty} P(\sup_{s \in I_n} p_n^*(s) \leq q) \leq q,$$

for all $q \in \text{int } C$, where C is the set from display (1.1) for the benchmark asymptotic test based on the unfeasible statistic $n^{1/2}(r(\hat{\theta}) - r(\theta_n))$ and the simple null hypothesis that θ_n is the true parameter value. This conservative generalization of the validity property (1.1) holds irrespective of the values of the drift parameter c. Specifically, the role of $-|g(\tilde{\theta})|^{1-\mu}$ in the definition of I_n is to guarantee that $g(\hat{\theta}) - cn^{-1/2} \in I_n$ with probability approaching one. Conservative size control then follows from the fact that $\hat{\theta}_s^*$ with $s = g(\hat{\theta}) - cn^{-1/2}$ satisfies

$$(n^{1/2}(\hat{\theta} - \theta_n), (n^{1/2}(\hat{\theta}_s^* - \hat{\theta}) | D_n)) \xrightarrow{w}_w (\ell(0,c), (\ell(0,c) | M));$$

see eqs. (5.3)–(5.4).

6. NUMERICAL RESULTS AND CHOICE OF THE TUNING PARAMETERS

In this section we analyze the finite sample performance of the proposed bootstrap methodology by means of numerical simulations. The purpose is twofold: first, to investigate the practical advantage of our methodology over standard bootstrap methods; second, to provide some practical guidance on how to choose the functions g^* and the tuning parameter κ in the definition of the bootstrap parameter space. Simulations are based on setup $\mathcal{G}_3$ of Section 3, as it covers the general case of a true parameter value that could, but need not, lie on the boundary of the parameter space under the null hypothesis. This section is organized as follows. In Section 6.1, we describe the data generating processes (DGPs), the null hypotheses and the adopted bootstrap schemes. In Section 6.2, we discuss the performance of the tests both under the null and under local alternatives. Section 6.3 deals with the

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

choice of g^* and κ. Additional numerical results are provided in the accompanying supplement, Section S.2.

6.1. Monte Carlo design

We consider the same DGP as in (3.1), where $x_{n,t} = n^{-1/2}x_t$, $x_t := \sum_{i=1}^{t} \varepsilon_{x,i}$, $\varepsilon_{x,t} \sim$ i.i.d. $N(0,1)$, with the following specifications of ε_t:

1. $\varepsilon_t \sim$ i.i.d. $N(0,1)$;
2. $\varepsilon_t = \sigma_t \nu_t$, where $\sigma_t^2 = 0.7 + 0.3\varepsilon_{t-1}^2$ and $\nu_t \sim$ i.i.d. $N(0,1)$;
3. $\varepsilon_t = \sqrt{0.5}\varepsilon_{x,t} + \sqrt{0.5}\eta_t$, where $\eta_t \sim$ i.i.d. $N(0,1)$.

In each case, $\{\varepsilon_{x,t}\}$ is independent of, respectively, $\{\varepsilon_t\}, \{\nu_t\}$ and $\{\eta_t\}$. In Case 1, the regression errors are independent and Gaussian, while in Case 2, they exhibit ARCH-type conditional heteroskedasticity. Case 3 allows for correlation between ε_t and the regressor's innovation $\varepsilon_{x,t}$.

The parameter space is specified as $\Theta := \{\theta \in \mathbb{R}^2 : g(\theta) \geq 0\}$ where $g(\theta) = \theta_2$. That is, $\Theta := \mathbb{R} \times [0, \infty)$ – such that its boundary is given by $\partial\Theta = \mathbb{R} \times \{0\}$. For all parameter values, we test the null hypothesis $\mathsf{H}_0 : h(\theta_0) = 0$, with $h(\theta) = \theta_1 + \theta_2$, against the two-sided alternative $h(\theta_0) \neq 0$. To do so, we employ the test statistics $\tau_n = \phi(\sqrt{n}h(\hat{\theta}))$ and $\tau_n^* = \phi(\sqrt{n}(h(\hat{\theta}^*) - h(\hat{\theta})))$, where $\phi(x) = x^2$, while $\hat{\theta}$ and $\hat{\theta}^*$ denote the original and bootstrap constrained LS estimators, respectively. In order to analyze size control and power of the proposed tests, we consider both empirical rejection probabilities [ERPs] under the null and under local alternatives. For tests performed under the null, we consider three different choices of the true value θ_0, one located on $\partial\Theta$ and two located on $\Theta \backslash \partial\Theta$; specifically, $\theta_0 \in \{(0,0)', (-0.75, 0.75)', (-1.5, 1.5)'\}$. Under H_1, we employ a local alternative of the form $\theta_0 = a_0 n^{-1/2}$, $a_0 \in \mathbb{R}^2$, such that $h(\theta_0) \neq 0$ unless $a = (0,0)'$.

Tests are based on p-values obtained using a 'standard' – i.e., with $\Theta^* = \Theta$ – fixed-regressor Gaussian wild bootstrap and the proposed "corrected" bootstrap scheme. For the latter, the bootstrap parameter space is set to $\Theta^* = \mathbb{R} \times [g^*(\hat{\theta}_2), \infty)$, where the function g^* satisfies the assumptions of Corollary 4.2, see also Section 6.3. In order to assess the impact of the tuning parameter κ, we consider a grid of possible values for κ. Numerical results are based on 50,000 Monte Carlo simulations, each involving $B = 999$ bootstrap repetitions. Sample sizes are set to $n \in \{100, 200, 400, 800, 1600\}$.

6.2. Empirical rejection probabilities

We now discuss the ERPs of the bootstrap tests. Specifically, the Monte Carlo results in Table 1 and 2 refer to the case in which data are generated under the null and under local alternatives, respectively. The proposed modified bootstrap parameter space is based on the function $g^* = g - |g|^{1+\kappa}$ for several values of $\kappa > 0$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

TABLE 1. Empirical rejection probabilities (ERPs) of bootstrap tests under the null

Nominal level: 0.05																
		$\theta_0 = (0,0)'$					$\theta_0 = (-0.75, 0.75)'$					$\theta_0 = (-1.50, 1.50)'$				
		b_1	b_2				b_1	b_2				b_1	b_2			
			κ					κ					κ			
dist.	n		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0
ξ_1	100	4.2	4.7	5.0	5.3	5.4	6.9	7.0	7.2	7.3	7.5	6.3	6.3	6.3	6.4	6.5
	400	3.9	4.8	5.1	5.3	5.3	5.5	5.6	5.8	6.2	6.7	5.3	5.3	5.3	5.3	5.3
	800	3.7	4.8	5.1	5.2	5.2	5.2	5.3	5.4	5.6	6.2	5.2	5.2	5.2	5.2	5.2
ξ_2	100	4.2	4.7	5.0	5.3	5.5	7.1	7.3	7.4	7.6	7.8	6.2	6.3	6.3	6.4	6.5
	400	3.8	4.7	5.0	5.1	5.2	5.7	5.9	6.1	6.4	6.9	5.3	5.3	5.3	5.3	5.3
	800	3.6	4.6	4.8	4.9	4.9	5.1	5.2	5.3	5.5	6.0	5.1	5.1	5.1	5.1	5.1
ξ_3	100	4.3	4.7	5.0	5.3	5.5	7.1	7.2	7.3	7.5	7.7	6.4	6.4	6.4	6.5	6.6
	400	3.7	4.6	4.9	5.1	5.1	5.5	5.7	5.9	6.2	6.7	5.2	5.2	5.2	5.2	5.2
	800	3.7	4.8	5.0	5.1	5.2	5.1	5.2	5.3	5.5	6.0	5.1	5.1	5.1	5.1	5.1
Nominal level: 0.10																
		$\theta_0 = (0,0)'$					$\theta_0 = (-0.75, 0.75)'$					$\theta_0 = (-1.50, 1.50)'$				
		b_1	b_2				b_1	b_2				b_1	b_2			
			κ					κ					κ			
dist.	n		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0
ξ_1	100	8.0	9.0	9.7	10.3	10.6	13.0	13.3	13.6	14.1	14.6	11.5	11.6	11.6	11.7	11.8
	400	7.7	9.5	10.1	10.4	10.5	10.4	10.5	10.8	11.3	12.4	10.3	10.3	10.3	10.3	10.3
	800	7.4	9.4	9.9	10.1	10.1	10.4	10.4	10.5	10.7	11.5	10.1	10.1	10.1	10.1	10.1
ξ_2	100	8.1	9.0	9.7	10.3	10.5	13.2	13.5	13.8	14.3	14.7	11.3	11.3	11.4	11.5	11.6
	400	7.5	9.2	9.9	10.2	10.3	10.7	10.9	11.1	11.6	12.5	10.2	10.2	10.2	10.2	10.3
	800	7.2	9.2	9.8	10.0	10.0	10.2	10.3	10.3	10.5	11.3	10.3	10.3	10.3	10.3	10.3
ξ_3	100	8.3	9.2	9.9	10.5	10.8	13.3	13.7	14.0	14.5	15.0	11.7	11.7	11.8	11.9	12.0
	400	7.6	9.4	10.0	10.3	10.3	10.4	10.5	10.8	11.3	12.4	10.2	10.2	10.2	10.2	10.2
	800	7.4	9.3	9.9	10.1	10.1	10.1	10.1	10.2	10.4	11.2	10.0	10.0	10.0	10.0	10.0

Note: bootstrap tests are based on a standard fixed-regressor wild bootstrap (b_1) and on the proposed corrected wild bootstrap method (b_2) of Section 4, using $g^* = g - |g|^{1+\kappa}$. ERPs are estimated using 50,000 Monte Carlo replications and 999 bootstrap repetitions. The column "dist." shows the distributions of ε_t: $\xi_1 \sim iidN(0,1)$, $\xi_2 \sim ARCH(1)$ and $\xi_3 = \sqrt{0.5}v_t + \sqrt{0.5}\varepsilon_{x,t}$, where $v_t \sim iidN(0,1)$ and $\varepsilon_{x,t}$ is the error term of the predictive variable $x_{n,t}$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

TABLE 2. Empirical rejection probabilities (ERPs) of bootstrap tests under local alternatives

Nominal level: 0.05																
		$a_0=(-3,0)'$					$a_0=(3,0)'$					$a_0=(5,0)'$				
		b_1	b_2				b_1	b_2				b_1	b_2			
			κ					κ					κ			
dist.	n		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0
ξ_1	100	21.0	21.0	21.1	21.2	21.3	40.6	40.9	41.0	41.0	41.0	68.0	68.0	68.0	68.0	68.0
	400	18.9	19.1	19.3	19.4	19.5	38.5	38.8	38.8	38.8	38.8	64.9	64.9	64.9	64.9	64.9
	800	18.6	18.8	19.0	19.1	19.1	37.6	37.9	37.9	37.9	38.0	64.0	64.0	64.0	64.0	64.0
ξ_2	100	21.7	21.8	21.9	22.0	22.1	41.9	42.1	42.2	42.2	42.2	68.5	68.5	68.5	68.5	68.5
	400	19.2	19.4	19.5	19.7	19.8	38.3	38.7	38.7	38.7	38.7	64.7	64.8	64.8	64.8	64.8
	800	18.6	18.8	19.0	19.1	19.1	37.8	38.1	38.1	38.1	38.1	64.2	64.2	64.2	64.2	64.2
ξ_3	100	20.6	20.7	20.8	20.8	21.0	40.8	41.0	41.1	41.1	41.1	67.3	67.3	67.3	67.3	67.3
	400	19.0	19.1	19.3	19.4	19.4	38.1	38.4	38.5	38.5	38.5	65.0	65.0	65.0	65.0	65.0
	800	18.3	18.5	18.8	18.9	18.9	37.7	38.0	38.1	38.1	38.1	63.5	63.5	63.5	63.5	63.5
Nominal level: 0.10																
		$a_0=(-3,0)'$					$a_0=(3,0)'$					$a_0=(5,0)'$				
		b_1	b_2				b_1	b_2				b_1	b_2			
			κ					κ					κ			
dist.	n		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0		0.25	0.50	1.0	2.0
ξ_1	100	29.6	29.8	29.9	30.1	30.3	54.7	55.0	55.1	55.2	55.2	81.7	81.7	81.8	81.8	81.8
	400	27.0	27.3	27.7	28.1	28.2	52.2	52.6	52.7	52.7	52.7	79.6	79.6	79.6	79.6	79.6
	800	26.4	26.9	27.3	27.6	27.6	51.7	52.1	52.2	52.2	52.2	78.7	78.7	78.7	78.7	78.7
ξ_2	100	30.2	30.4	30.6	30.8	31.0	55.7	55.9	56.0	56.0	56.0	82.0	82.0	82.0	82.0	82.0
	400	27.1	27.4	27.9	28.2	28.3	51.8	52.1	52.2	52.2	52.2	79.3	79.3	79.3	79.3	79.3
	800	26.6	27.0	27.5	27.7	27.7	51.5	51.9	51.9	51.9	51.9	78.6	78.6	78.6	78.6	78.6
ξ_3	100	29.1	29.3	29.4	29.7	29.9	54.2	54.5	54.6	54.6	54.6	80.9	80.9	80.9	80.9	80.9
	400	26.7	27.0	27.4	27.7	27.8	51.7	52.1	52.2	52.2	52.2	79.4	79.4	79.4	79.4	79.4
	800	26.2	26.6	27.1	27.3	27.3	51.3	51.7	51.7	51.7	51.7	78.5	78.5	78.5	78.5	78.5

Note: bootstrap tests are based on a standard fixed-regressor wild bootstrap (b_1) and on the proposed corrected wild bootstrap method (b_2) of Section 4, using $g^* = g - |g|^{1+\kappa}$. ERPs are estimated using 50,000 Monte Carlo replications and 999 bootstrap repetitions. The column "dist." shows the distributions of ε_t: $\xi_1 \sim iidN(0,1)$, $\xi_2 \sim ARCH(1)$ and $\xi_3 = \sqrt{0.5}v_t + \sqrt{0.5}\varepsilon_{x,t}$, where $v_t \sim iidN(0,1)$ and $\varepsilon_{x,t}$ is the error term of the predictive variable $x_{n,t}$.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Table 1 shows that the "standard" bootstrap scheme typically under-rejects the true null hypothesis when the parameter lies on the boundary of the parameter space Θ whereas, as expected, its ERPs are closer to the nominal level when θ_0 is in the interior of Θ. Our proposed bootstrap performs similarly to the "standard" bootstrap for very small values of κ, with the impact of the correction becoming more relevant as κ increases. If the parameter is on the boundary of the parameter space ($\theta_0 \in \partial\Theta$), our proposed bootstrap scheme gives rise to smaller absolute size distortions than the "standard" bootstrap, for all the considered DGPs and all values of κ. When $\theta_0 \in \text{int}\Theta$, we observe very little variability in the ERPs across the different bootstrap methods, at least for reasonably small values of κ.

Table 2 reports the ERPs of the tests when data are generated under local alternatives $\theta_0 = a_0 n^{-1/2}$, $a_0 \in \{(-3,0)', (3,0)', (5,0)'\}$, such that the true parameter values lie on the boundary of the parameter space. Results show that both bootstrap schemes have power under local alternatives, with the "corrected" bootstrap generally showing higher ERPs than the "standard" bootstrap, in line with the results obtained under the null. Finally, we notice that the sign of the deviations from the null hypothesis matters, with positive deviations showing higher ERPs. This finding can be explained by the fact that the limit distribution of $n^{1/2}(h(\hat{\theta}) - h(\theta_0))$ is asymmetric when θ_0 lie on the boundary of Θ. Results about local alternatives such that θ_0 are $n^{-1/2}$-local to the boundary are substantially similar and are reported in Section S.2 of the supplement.

6.3. Choice of g^* and κ

We now consider the practical issue of choosing the function g^* and the tuning parameter κ used to construct the modified bootstrap parameter space Θ^*.

Regarding g^*, in Section 4 we discussed the functions $g^*_{(1)} := g - |g|^{1+\kappa}$, $\kappa > 0$, which satisfy the assumptions of Corollary 4.2 and were employed in the simulations so far, whereas in Section 5.2 we considered also $g^*_{(2)} := g - n^{-\kappa}|g|$, $\kappa \in (0, 1/2)$, corresponding to $s_n = n^{1/2-\kappa}$ in the concluding paragraph of Section 5.2. Numerical results in Table 1 and 2 and in the accompanying Supplement, Section S.2, show that both choices of g^* deliver good test performance, both under the null and under local alternatives. The most salient difference between $g^*_{(1)}$ and $g^*_{(2)}$ is that tests based on $g^*_{(1)}$ tend to be more robust to the choice of κ when $g(\theta_0) \geq 1$.

Concerning the choice of the tuning parameter κ, we focus on $g^* = g^*_{(1)}$. Preliminary considerations point at a possible trade-off between the cases of a boundary and an interior location of the true parameter θ_0. Thus, for $\theta_0 \in \partial\Theta$, larger values of κ accelerate the convergence of $g(\hat{\theta})^{1+\kappa}$ to zero, which can be expected to favor bootstrap performance as the bootstrap true value $\hat{\theta}$ is put at a smaller distance from the bootstrap boundary. On the other hand, if $\theta_0 \in \text{int}(\Theta)$ and $g(\theta_0) \in (0, 1)$, in small samples large values of κ may put $\hat{\theta}$ too close to the bootstrap boundary, yielding inferior bootstrap performance.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Our Monte Carlo study indeed confirms that small values of κ are preferable when $\theta_0 \in \text{int}(\Theta)$ and $g(\theta_0) \in (0,1)$; however, it also shows that the proposed correction quickly provides (as n increases) satisfactory size control for small values of κ even when $\theta_0 \in \partial\Theta$. Finally, we notice that when $\theta_0 \in \text{int}(\Theta)$ and $g(\theta_0) \geq 1$ the choice of κ has negligible impact on the ERPs. Overall, our numerical analysis suggests that choices of κ close to 0.5 provide quite satisfactory size control across all the considered scenarios.

Remark. The above guideline about the choice of κ is based on numerical evidence; it delivers a reasonable simple choice which can be easily implemented. It is not optimal in any sense, and indeed alternative methods could be employed to obtain data-driven choices of κ. For instance, the unrestricted parameter estimates could be used to assess how far the true parameter value θ_0 is from the boundary of the parameter space, and then calibrate the choice of κ accordingly. This approach would be in the spirit of Romano, Shaikh and Wolf (2014), who suggest to improve the power of tests of moment inequalities by introducing a first step, where a confidence region for the moments is constructed using their unrestricted estimates. Although this approach may improve the finite sample properties of our tests, it would require a preliminary choice of further tuning parameters, such as β in Romano, Shaikh and Wolf (2014), hence introducing an extra layer of complexity.

7. CONCLUSIONS

In this paper, we analyzed the problem of bootstrap hypotheses tests on the parameters (α, β) of a predictive regression $y_t = \alpha + \beta x_{t-1} + \varepsilon_t$, generalizable to higher dimensions, when the parameter space is defined by means of a smooth constraint $g(\alpha, \beta) \geq 0$ and the true parameter vector under the null hypothesis may lie on the boundary of the parameter space. In the framework of constrained parameter estimation, implementation of the bootstrap is not straightforward, as the presence of a parameter on the boundary of the parameter space makes the bootstrap measure random in the limit.

We discussed possible solutions to this inference problem. Specifically, we presented some modifications of standard bootstrap schemes where the bootstrap parameter space is shifted by a data-dependent function, thus allowing us to regain control over the boundary as a source of limiting bootstrap randomness. We also proved validity of the associated bootstrap inference in the cases where the posited predicting variable is I(1).

Our contribution is novel in the framework of predictive regression, in that the existing literature has not analyzed the bootstrap in contexts combining non-stationarity of the posited predictor with a priori knowledge about the possible form of predictability, represented by a restricted parameter space. The value of our work is to provide valid bootstrap implementations in this setting.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

APPENDIX

A. Mathematical Appendix

A.1. Proof of Theorem 4.1

Introduce $\tilde{x}_t := (1, x_{n,t-1})'$. Let $\mu_n := n^{1/2}(\hat{\theta} - \theta_0)$, $M_n := n^{-1}\sum_{t=1}^n \tilde{x}_t\tilde{x}_t'$ and $N_n^* := n^{-1/2}\sum_{t=1}^n \varepsilon_t^*\tilde{x}_t$. Moreover, let the normalized bootstrap estimator be denoted by $\mu_n^* := n^{1/2}(\hat{\theta}^* - \hat{\theta})$; similarly, $\tilde{\mu}_n^* := n^{1/2}(\tilde{\theta}^* - \hat{\theta})$, where $\tilde{\theta}^*$ is the unrestricted (OLS) bootstrap estimator. On the event $\{\det(M_n) > 0\}$ with $P(\det(M_n) > 0) \to 1$, the estimator $\tilde{\theta}^*$ is well defined and unique. As we are interested in distributional convergence results, without loss of generality we proceed as if $P(\det(M_n) > 0) = 1$.

By arguments similar to the proof of Theorem 4.1 in Cavaliere and Georgiev (2020), it can be concluded that $(\mu_n, M_n, N_n^*) \xrightarrow{w^*}_w (\ell(\theta_0), M, M^{1/2}\xi^*)|(M, \ell(\theta_0)))$ in $\mathbb{R}^{2\times 4}$, where M is of full rank with probability one, $\xi^*|(M, \ell(\theta_0)) \sim N(0, \sigma_e^2 I_2)$ and σ_e^2 denotes the variance of ε_t corrected for $\Delta x_{n,t}$. To derive the result (4.1), we analyze the properties of μ_n^* on a special probability space where (μ_n, M_n, N_n^*) given the data converge weakly a.s. rather than weakly in distribution. Specifically, by Lemma A.2(a) in Cavaliere and Georgiev (2020), we can consider a probability space (where $\ell(\theta_0), M$ and, for every $n \in N$, also the original data and the bootstrap sample can be redefined, maintaining their distribution), such that

$$\mu_n \xrightarrow{a.s.} \ell(\theta_0),\ M_n \xrightarrow{a.s.} M,\ N_n^* \xrightarrow{w^*}_{a.s.} M^{1/2}\xi^*|(M, \ell(\theta_0)) = M^{1/2}\xi^*|M, \tag{A.1}$$

the last equality being an a.s. equality of conditional distributions.

Let $q_n^*(\theta) := n^{-1}\sum_{t=1}^n (y_t^* - \theta'\tilde{x}_t)^2$ with $\tilde{\theta}^* := \arg\min_{\theta\in\mathbb{R}^2} q_n^*(\theta)$ being well defined and unique for outcomes in the event $\{\det(M_n) > 0\}$. On the special probability space, the asymptotic distribution of $\tilde{\mu}_n^* = n^{1/2}(\tilde{\theta}^* - \hat{\theta}) = M_n^{-1}N_n^*$ follows from (A.1) and a CMT (Theorem 10 of Sweeting, 1989):

$$\tilde{\mu}_n^* \xrightarrow{w^*}_{a.s.} \tilde{\ell}^*|(M, \ell(\theta_0)) = \tilde{\ell}^*|M,\ \tilde{\ell}^* := M^{-1/2}\xi^*. \tag{A.2}$$

Let us turn now to the bootstrap estimator $\hat{\theta}^*$. If $g(\theta_0) > g^*(\theta_0)$, then the consistency facts $\hat{\theta} \xrightarrow{a.s.} \theta_0$ (from (A.1)) and $\tilde{\theta}^* \xrightarrow{w^*}_{a.s.} \theta_0$ (from (A.2)), jointly with the continuity of g, g^* at θ_0, imply that $P^*(g(\tilde{\theta}^*) \geq g^*(\hat{\theta})) \xrightarrow{a.s.} 1$. Hence, $\tilde{\theta}^*$ uniquely minimizes q_n^* on Θ^* with P^*-probability approaching one a.s. This establishes the existence of $\hat{\theta}^*$ with P^*-probability approaching one a.s., as well as the facts $P^*(\hat{\theta}^* = \tilde{\theta}^*) \xrightarrow{a.s.} 1$ and $P^*(\mu_n^* = \tilde{\mu}_n^*) \xrightarrow{a.s.} 1$. Using also (A.2), it follows that $\mu_n^* \xrightarrow{w^*}_{a.s.} \tilde{\ell}^*|M$ on the special probability space, and since $\mu_n \xrightarrow{a.s.} \ell(\theta_0)$ on this space, it follows further that $(\mu_n, (\mu_n^*|D_n)) \xrightarrow{w}_w (\ell(\theta_0), (\tilde{\ell}^*|M))$ on a general probability space, as asserted in (4.1).

In the case where $g^*(\theta_0) = g(\theta_0)$, it still holds that $\tilde{\theta}^*$ uniquely minimizes q_n^* on Θ^* whenever $g(\tilde{\theta}^*) \geq g^*(\hat{\theta})$, such that $\hat{\theta}^*$ exists and equals $\tilde{\theta}^*$ on the event $\{g(\tilde{\theta}^*) \geq g^*(\hat{\theta})\}$. However, the probability of this event no longer tends to one. Whenever $g(\tilde{\theta}^*) < g^*(\hat{\theta})$, a minimizer of q_n^* on Θ^* exists if and only if a minimizer, say $\breve{\theta}^*$, of q_n^* on $\partial\Theta^*$ exists and minimizes q_n^* over the entire Θ^* (this claim is due to the fact that, for outcomes in the event $\{\det(M_n) > 0\}$, the function $q_n^*(\theta)$ is locally minimized uniquely at $\tilde{\theta}^*$). Let $\mathbb{I}_n^* := \mathbb{I}_{\{b(\tilde{\theta}^*)\geq 0\}}$ with $b(\theta) := g(\theta) - g^*(\hat{\theta})$. We show in Section A.2 below that $\breve{\theta}^*(1 - \mathbb{I}_n^*)$, with a measurable

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

$\check{\theta}^*$, is well defined with P^*-probability approaching one a.s. and $(q_n^*(\check{\theta}^*) - q_n^*(\theta))(1 - \mathbb{I}_n^*) \leq 0$ for all $\theta \in \Theta^*$, with P^*-probability approaching one a.s. This establishes the possibility to define the bootstrap estimator $\hat{\theta}^*$ as

$$\hat{\theta}^* = \tilde{\theta}^* \mathbb{I}_n^* + \check{\theta}^*(1 - \mathbb{I}_n^*), \tag{A.3}$$

and, therefore, the existence of $\hat{\theta}^*$ with P^* -probability approaching one a.s. The existence result carries over to a general probability space with P^*-probability approaching one in probability.

In Section A.2, we also show that $\|\check{\theta}^* - \hat{\theta}\|(1 - \mathbb{I}_n^*) = O_{p^*}(n^{-1/2})$ a.s., and as a result, $\|\hat{\theta}^* - \hat{\theta}\| = O_{p^*}(n^{-1/2})$ a.s., using also (A.2). We do not discuss the uniqueness of $\check{\theta}^*$ but instead we argue next that the measurable minimizers of q_n^* over the bootstrap boundary are asymptotically equivalent, as they give rise to the same asymptotic distribution of $\hat{\theta}^*$.

To accomplish this, we use the result of Section A.2 that $\check{\theta}^*$ satisfies a first-order condition [foc] with P^* -probability approaching one a.s. Let dots over function names denote differentiation w.r.t. θ (e.g., $\dot{q}_n^*(\theta) := (\partial \dot{q}_n^* / \partial \theta')(\theta)$, a column vector). Then the foc takes the form

$$\{\dot{q}_n^*(\check{\theta}^*) + \check{\delta}_n \dot{b}(\check{\theta}^*)\}(1 - \mathbb{I}_n^*) = \{\dot{q}_n^*(\check{\theta}^*) + \check{\delta}_n \dot{g}(\check{\theta}^*)\}(1 - \mathbb{I}_n^*) = 0,\ b(\check{\theta}^*)(1 - \mathbb{I}_n^*) = 0,$$

where $\check{\delta}_n \in \mathbb{R}$ is a Lagrange multiplier. The foc implies, by means of a standard argument, the existence of a measurable $\bar{\theta}^*$ between $\check{\theta}^*$ and $\hat{\theta}$ such that

$$\{n^{1/2}(\check{\theta}^* - \hat{\theta}) - (I_2 - A_n^* \dot{g}(\bar{\theta}^*)')\tilde{\mu}_n^* + A_n^* n^{1/2} b(\hat{\theta})\}(1 - \mathbb{I}_n^*) = 0,$$

where $A_n^* := M_n^{-1} \dot{g}(\check{\theta}^*)[\dot{g}(\bar{\theta}^*)' M_n^{-1} \dot{g}(\check{\theta}^*)]^{-1}$ is well defined with P^*-probability approaching one a.s. As further $\|\check{\theta}^* - \hat{\theta}\|(1 - \mathbb{I}_n^*) = O_{p^*}(n^{-1/2})$ a.s., $\|\bar{\theta}^* - \hat{\theta}\|(1 - \mathbb{I}_n^*) = O_{p^*}(n^{-1/2})$ a.s. and $\hat{\theta} - \theta_0 = O(n^{-1/2})$ a.s., using the continuity of $\dot{g}(\theta)$ at θ_0 it follows that

$$\{n^{1/2}(\check{\theta}^* - \hat{\theta}) - [(I_2 - A^* \dot{g}')\tilde{\mu}_n^* - A^*(\dot{g} - \dot{g}^*)' n^{1/2}(\hat{\theta} - \theta_0)]\}(1 - \mathbb{I}_n^*) = o_{p^*}(1) \text{ a.s.},$$

where $A^* := M^{-1} \dot{g}[\dot{g}' M^{-1} \dot{g}]^{-1}$ and $P^*(|o_{p^*}(1)| > \eta) \xrightarrow{a.s.} 1$ for all $\eta > 0$.

Returning to (A.3), we conclude that

$$n^{1/2}(\hat{\theta}^* - \hat{\theta}) = \tilde{\mu}_n^* \mathbb{I}_n^* + \{(I_2 - A^* \dot{g}')\tilde{\mu}_n^* - A^*(\dot{g} - \dot{g}^*)' n^{1/2}(\hat{\theta} - \theta_0)\}(1 - \mathbb{I}_n^*) + o_{p^*}(1) \text{ a.s.} \tag{A.4}$$

Consider the event indicated by $\mathbb{I}_n^*$. As $\|\tilde{\theta}^* - \hat{\theta}\| = O_{p^*}(n^{-1/2})$ a.s. and $\hat{\theta} - \theta_0 = O(n^{-1/2})$ a.s., by the mean value theorem and the continuous differentiability of g, g^* it holds that

$$n^{1/2} b(\tilde{\theta}^*) = \dot{g}' \tilde{\mu}_n^* + (\dot{g} - \dot{g}^*)' \mu_n + o_{p^*}(1) \text{ a.s.}$$

Then $\mathbb{I}_n^* \xrightarrow{w^*}_{a.s.} \mathbb{I}_\infty | (M, \ell(\theta_0))$ with $\mathbb{I}_\infty := \mathbb{I}_{\{\dot{g}' \tilde{\ell}^* \geq (\dot{g}^* - \dot{g})' \ell(\theta_0)\}}$, by (A.1)-(A.2) and the CMT for weak a.s. convergence (Theorem 10 of Sweeting, 1989), as the probability of the limiting discontinuities is 0: $P(\dot{g}' \tilde{\ell}^* = (\dot{g}^* - \dot{g})' \ell(\theta_0) | (M, \ell(\theta_0))) = 0$ a.s. By exactly the same facts, passage to the limit directly in (A.4) yields

$$n^{1/2}(\hat{\theta}^* - \hat{\theta}) \xrightarrow{w^*}_{a.s.} \{\tilde{\ell}^* \mathbb{I}_\infty + \check{\ell}^*(1 - \mathbb{I}_\infty)\} | (M, \ell(\theta_0)),\ \check{\ell}^* := (I_2 - A^* \dot{g}')\tilde{\ell}^* - A^*(\dot{g} - \dot{g}^*)' \ell$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

on the special probability space, where also $\mu_n \overset{a.s.}{\to} \ell(\theta_0)$ by (A.1). Therefore, on a general probability space it holds that

$$(\mu_n, (n^{1/2}(\hat{\theta}^* - \hat{\theta})|D_n)) \overset{w}{\to}_w (\ell(\theta_0), [\{\tilde{\ell}^* \mathbb{I}_\infty + \check{\ell}^*(1 - \mathbb{I}_\infty)\}|(M, \ell(\theta_0))]).$$

As $I_2 - A^* \dot{g}' = \dot{g}_\perp (\dot{g}_\perp' M \dot{g}_\perp)^{-1} \dot{g}_\perp' M$ and $\tilde{\ell}^* = M^{-1/2} \xi^*$, it follows that

$$\begin{aligned}\tilde{\ell}^* \mathbb{I}_\infty + \check{\ell}^*(1 - \mathbb{I}_\infty) = {} & \dot{g}_\perp (\dot{g}_\perp' M \dot{g}_\perp)^{-1} \dot{g}_\perp' M^{1/2} \xi^* \\ & + M^{-1} \dot{g} (\dot{g}' M^{-1} \dot{g})^{-1} \max\{(\dot{g}^* - \dot{g})' \ell, \dot{g}' M^{-1/2} \xi^*\},\end{aligned}$$

which is $\arg\min_{\{\dot{g}'\lambda \geq (\dot{g}^* - \dot{g})'\ell\}} ||\lambda - M^{-1/2}\xi^*||_M$ a.s. as asserted in (4.2).

For use in the proof of Corollary 4.2, we notice here a useful consequence of the previous argument. Return to the special probability space where

$$(\mu_n, (n^{1/2}(\hat{\theta}^* - \hat{\theta})|D_n)) \overset{w}{\to}_{a.s.} (\ell(\theta_0), [\{\tilde{\ell}^* \mathbb{I}_\infty + \check{\ell}^*(1 - \mathbb{I}_\infty)\}|(M, \ell(\theta_0))]).$$

Let $\tau_n := \phi(\mu_n)$, $\tau_n^* := \phi(n^{1/2}(\hat{\theta}^* - \hat{\theta}))$, $\tau := \phi(\ell(\theta_0))$ and $\tau^* := \phi(\tilde{\ell}^* \mathbb{I}_\infty + \check{\ell}^*(1 - \mathbb{I}_\infty))$ for a continuous $\phi : \mathbb{R} \to \mathbb{R}$. Then

$$(\tau_n, (\tau_n^*|D_n)) \overset{w}{\to}_{a.s.} (\tau, \tau^*|(M, \ell(\theta_0)))$$

by the CMT of Sweeting (1989). Furthermore, the regular conditional distributions $\tau_n^*|D_n$ converge weakly to the regular conditional distribution $\tau^*|(M, \ell(\theta_0))$ for almost all outcomes; see Theorem 2.2 of Berti, et al. (2006). For any fixed outcome such that the previous convergence holds, also $F_n^{*-1}(q_i) \to F_{M,\ell}^{-1}(q_i)$, $i = 1, 2$, hold for the sample paths of the respective conditional quantile functions, provided that q_1, q_2 are continuity points of the sample path of $F_{M,\ell}^{-1}$. If q_1, q_2 are continuity points of almost all sample paths of $F_{M,\ell}^{-1}$, it follows that $F_n^{*-1}(q_i) \to_{a.s.} F_{M,\ell}^{-1}(q_i)$, $i = 1, 2$. Therefore, on a general probability space,

$$(\tau_n, F_n^{*-1}(q_1), F_n^{*-1}(q_2), (\tau_n^*|D_n)) \overset{w}{\to}_w (\tau, F_{M,\ell}^{-1}(q_1), F_{M,\ell}^{-1}(q_2), \tau^*|(M, \ell(\theta_0))), \tag{A.5}$$

provided that $F_{M,\ell}^{-1}$ is a.s. continuous at q_1, q_2.

A.2. Details of the proof of Theorem 4.1

Let $g^*(\theta_0) = g(\theta_0)$ throughout this subsection. For outcomes such that $\tilde{\theta}^* \notin \Theta^*$ and $\lambda_{\min}(M_n) > 0$, the quadratic function q_n^* is not minimized over Θ^* at any interior point of Θ^* (for otherwise this point would have to be the stationary point $\tilde{\theta}^* \notin \Theta^*$ of q_n^*, a contradiction). For such outcomes, if q_n^* is at all minimized over Θ^*, then this has to occur at a boundary point of Θ^*. Since $\partial\Theta^* \subseteq \{\theta \in \mathbb{R}^2 : g(\theta) = g^*(\hat{\theta})\} =: \tilde{\partial}\Theta^*$, we proceed by constructing a minimizer of q_n^* over the latter set and by showing that this minimizer is in fact a global one over Θ^*. This (and some added measurability considerations) establishes the well-definition of $\check{\theta}^*$ in (A.3). Then we establish the $n^{-1/2}$ consistency rate of $\check{\theta}^*$ in the sense that $\|\check{\theta}^* - \hat{\theta}\|(1 - \mathbb{I}_n^*) = O_{p^*}(n^{-1/2})$ a.s.

Step 1. Existence of a minimizer of q_n^* over a portion of $\tilde{\partial}\Theta^*$close to θ_0. The point $(\theta', c)' = (\theta_0', g(\theta_0))' \in \mathbb{R}^3$ trivially satisfies the equation $g(\theta) = c$. Since g is continuously differentiable in a neighborhood of θ_0 and $\dot{g} = (\dot{g}_1(\theta_0), \dot{g}_2(\theta_0))' \neq 0$ (say that $\dot{g}_1(\theta_0) \neq 0$, with the subscript denoting partial differentiation), by the implicit function theorem there exist an $r > 0$ and a unique function $\gamma : [\theta_{2,0} - r, \theta_{2,0} + r] \times [g(\theta_0) - r, g(\theta_0) + r] \to [\theta_{1,0} -$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

$r, \theta_{1,0}+r]$ such that $\gamma(\theta_{2,0}, g(\theta_0)) = \theta_{1,0}$, $g(\gamma(\theta_2, c), \theta_2) = c$; moreover, γ is continuously differentiable. For outcomes such that $|g^*(\hat{\theta}) - g(\theta_0)| \le r$, the (non-empty) portion of the curve $\tilde{\partial}\Theta^* = \{\theta \in \mathbb{R}^2 : g(\theta) = g^*(\hat{\theta})\}$ contained in the square $\Pi := [\theta_{1,0}-r, \theta_{1,0}+r] \times [\theta_{2,0}-r, \theta_{2,0}+r]$ can be parameterized as $\theta_1 = \gamma(\theta_2, g^*(\hat{\theta}))$, $\theta_2 \in [\theta_{2,0}-r, \theta_{2,0}+r]$. Define $\breve{\theta}^* := (\gamma(\breve{\theta}_2^*, g^*(\hat{\theta}^r)), \breve{\theta}_2^*)'$, where $\breve{\theta}_2^*$ is a measurable minimizer of the continuous function $q_n^*(\gamma(\theta_2, g^*(\hat{\theta}^r)), \theta_2)$ over $\theta_2 \in [\theta_{2,0}-r, \theta_{2,0}+r]$, with $\hat{\theta}^r := \hat{\theta}\mathbb{I}_{\{|g^*(\hat{\theta})-g(\theta_0)|\le r\}} + \theta_0 \mathbb{I}_{\{|g^*(\hat{\theta})-g(\theta_0)|>r\}}$. Since $\mathbb{I}_{\{|g^*(\hat{\theta})-g(\theta_0)|\le r\}} \overset{a.s.}{\to} 1$ under $g^*(\theta_0) = g(\theta_0)$, it follows that $\breve{\theta}^*$ minimizes q_n^* over $\tilde{\partial}\Theta^* \cap \Pi$ with P^*-probability approaching one a.s.

Step 2. Minimization of q_n^* over the entire bootstrap parameter space. For outcomes in

$$\mathcal{A}_n := \{|g^*(\hat{\theta}) - g(\theta_0)| \le r\} \cap \{g(\tilde{\theta}^*) < g^*(\hat{\theta})\} \cap \{\|\hat{\theta} - \theta_0\| + \|\tilde{\theta}^* - \hat{\theta}\| \le \tfrac{r}{2}\},$$

the minimum of q_n^* over the entire bootstrap parameter space Θ^* exists and is attained only in Π (e.g., at $\breve{\theta}^*$ defined in Step 1), provided that

$$\alpha_n := \lambda_{\min}(M_n)\tfrac{r^2}{4} - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \hat{\theta}\|^2 > 0.$$

To see this, consider $\theta^c := c\hat{\theta} + (1-c)\tilde{\theta}^*$ where $c := \inf\{a \in [0,1] : b(a\hat{\theta} + (1-a)\tilde{\theta}^*) = 0\}$; θ^c is well defined whenever $g(\tilde{\theta}^*) < g^*(\hat{\theta})$ because $g(\hat{\theta}) \ge g^*(\hat{\theta})$ and b is continuous. Moreover, $\theta^c \in \Pi$ for outcomes in $\mathcal{A}_n$ because $\|\theta^c - \theta_0\| \le \|\hat{\theta} - \theta_0\| + \|\tilde{\theta}^* - \hat{\theta}\| \le \frac{r}{2}$ and, hence, $q_n^*(\theta^c) \ge q_n^*(\breve{\theta}^*)$ for outcomes in $\mathcal{A}_n$, by the minimizing property of $\breve{\theta}^*$ on $\tilde{\partial}\Theta^* \cap \Pi$ and the fact that $b(\theta^c) = 0$. For any $\theta \notin \Pi$ and outcomes in $\mathcal{A}_n$, we therefore find that

$$\begin{aligned}
q_n^*(\theta) - q_n^*(\breve{\theta}^*) &\ge q_n^*(\theta) - q_n^*(\theta^c) = q_n^*(\theta) - q_n^*(\tilde{\theta}^*) + q_n^*(\tilde{\theta}^*) - q_n^*(\theta^c) \\
&\ge \lambda_{\min}(M_n)\|\theta - \tilde{\theta}^*\|^2 - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \theta^c\|^2 \\
&\ge \lambda_{\min}(M_n)\{\|\theta - \theta_0\| - \|\tilde{\theta}^* - \theta_0\|\}^2 - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \hat{\theta}\|^2 \\
&\ge \lambda_{\min}(M_n)\{r - \|\tilde{\theta}^* - \hat{\theta}\| - \|\hat{\theta} - \theta_0\|\}^2 - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \hat{\theta}\|^2 \\
&\ge \lambda_{\min}(M_n)\tfrac{r^2}{4} - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \hat{\theta}\|^2 = \alpha_n.
\end{aligned}$$

Thus, for outcomes in $\mathcal{A}_n \cap \{\alpha_n > 0\}$, q_n^* out of Π is larger than $\min_{\theta \in \tilde{\partial}\Theta^* \cap \Pi} q_n^*(\theta)$. As $\tilde{\partial}\Theta^* \subseteq \Theta^*$, it follows that $\min_{\theta \in \Theta^* \cap \Pi} q_n^*(\theta)$ (which exists) for such outcomes is actually $\min_{\theta \in \Theta^*} q_n^*(\theta)$. Moreover,

$$\min_{\theta \in \Theta^*} q_n^*(\theta) = \min_{\theta \in \Theta^* \cap \Pi} q_n^*(\theta) = \min_{\theta \in \tilde{\partial}\Theta^* \cap \Pi} q_n^*(\theta) = \min_{\theta \in \partial\Theta^* \cap \Pi} q_n^*(\theta),$$

for if $\min_{\theta \in \Theta^* \cap \Pi} q_n^*(\theta) < \min_{\theta \in \partial\Theta^* \cap \Pi} q_n^*(\theta)$, then $\min_{\theta \in \Theta^* \cap \Pi} q_n^*(\theta)$ (and thus, $\min_{\theta \in \Theta^*} q_n^*(\theta)$) is achieved at an interior point of Θ^*, which can only be $\tilde{\theta}^*$, a contradiction with $\tilde{\theta}^* \notin \Theta^*$ (i.e., with $g(\tilde{\theta}^*) < g^*(\hat{\theta})$). To summarize, for outcomes in $\mathcal{A}_n \cap \{\alpha_n > 0\}$, $\breve{\theta}^*$ minimizes q_n^* over Θ^* and is at the boundary of Θ^*.

We find the associated probability

$$\begin{aligned}
&P^*\left((1-\mathbb{I}_n^*)q_n^*(\breve{\theta}^*) < (1-\mathbb{I}_n^*)q_n(\theta) \ \ \forall \theta \in \Theta^* \setminus \Pi\right) \\
&\quad \ge P^*\left(|g^*(\hat{\theta}) - g(\theta_0)| \le r, \|\hat{\theta} - \theta_0\| \le \tfrac{r}{4}, \|\tilde{\theta}^* - \hat{\theta}\| \le \frac{r}{4}, \alpha_n > 0\right) \\
&\quad = \mathbb{I}_{\{|g^*(\hat{\theta})-g(\theta_0)|\le r\} \cap \{\|\hat{\theta}-\theta_0\| \le r/4\}} P^*\left(\|\tilde{\theta}^* - \hat{\theta}\| \le \tfrac{r}{4}, \alpha_n > 0\right) \overset{a.s.}{\to} 1
\end{aligned}$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

because $g(\hat{\theta}) \xrightarrow{a.s.} g(\theta_0)$, $\lambda_{\min}(M_n) \to \lambda_{\min}(M) > 0$ a.s., $\lambda_{\max}(M_n) \to \lambda_{\max}(M) < \infty$ a.s. and $\|\tilde{\theta}^* - \hat{\theta}\| \xrightarrow{w^*}_{a.s.} 0$. This establishes the fact that $\hat{\theta}^*$ of (A.3), with $\check{\theta}^*$ as defined in Step 1, minimizes q_n^* over the bootstrap parameter space Θ^* with P^*-probability approaching one a.s.

Step 3. Consistency rate of $\check{\theta}^*$. Similarly to Step 2, for outcomes in $\mathcal{A}_n$,

$$0 \geq q_n^*(\check{\theta}^*) - q_n^*(\theta^c) \geq \lambda_{\min}(M_n)\|\check{\theta}^* - \tilde{\theta}^*\|^2 - \lambda_{\max}(M_n)\|\tilde{\theta}^* - \hat{\theta}\|^2,$$

the first inequality by the minimizing property of $\check{\theta}^*$ over $\tilde{\partial}\Theta^* \cap \Pi$. Therefore,

$$\begin{aligned} &P^*\left((1-\mathbb{I}_n^*)\|\check{\theta}^* - \tilde{\theta}^*\|^2 \leq (1-\mathbb{I}_n^*)\frac{\lambda_{\max}(M_n)}{\lambda_{\min}(M_n)}\|\tilde{\theta}^* - \hat{\theta}\|^2\right) \\ &\geq P^*\left(|g^*(\hat{\theta}) - g(\theta_0)| \leq r, \|\hat{\theta} - \theta_0\| \leq \tfrac{r}{4}, \|\tilde{\theta}^* - \hat{\theta}\| \leq \tfrac{r}{4}\right) \\ &= \mathbb{I}_{\{|g^*(\hat{\theta})-g(\theta_0)|\leq r\}\cap\{\|\hat{\theta}-\theta_0\|\leq r/4\}} P^*\left(\|\tilde{\theta}^* - \hat{\theta}\| \leq \tfrac{r}{4}\right) \xrightarrow{a.s.} 1. \end{aligned}$$

As $\lambda_{\max}(M_n)/\lambda_{\min}(M_n) \xrightarrow{a.s.} \lambda_{\max}(M)/\lambda_{\min}(M)$ and $\|\tilde{\theta}^* - \hat{\theta}\| = O_{p^*}(n^{-1/2})P$-a.s. (the latter, by (A.2)), it follows that $(1-\mathbb{I}_n^*)\|\check{\theta}^* - \tilde{\theta}^*\| = O_{p^*}(n^{-1/2})P$-a.s. and $\|\hat{\theta}^* - \tilde{\theta}^*\| = O_{p^*}(n^{-1/2})P$-a.s. for $\hat{\theta}^*$ of (A.3). Thus, $\hat{\theta}^*$ has the same consistency rate as $\tilde{\theta}^*$. This argument applies to any measurable $\check{\theta}^*$ which minimizes q_n^* over $\tilde{\partial}\Theta^* \cap \Pi$ for outcomes in $\mathcal{A}_n$. This completes Step 3.

Finally, consider the first-order condition [foc] for minimization of q_n^* on $\tilde{\partial}\Theta^*$. As $\|\check{\theta}^* - \theta_0\|(1-\mathbb{I}_n^*) \leq \{\|\check{\theta}^* - \tilde{\theta}^*\| + \|\check{\theta}^* - \theta_0\|\}(1-\mathbb{I}_n^*) \xrightarrow{w^*}_{a.s.} 0$, it follows that $\mathbb{I}_{\{\check{\theta}^* \in \mathrm{int}(\Pi)\}}(1-\mathbb{I}_n^*) + \mathbb{I}_n^* \xrightarrow{w}_{a.s.} 1$. As additionally $\dot{g}(\theta_0) \neq 0$, by continuity of $\dot{g}(\theta) := (\partial g/\partial\theta')(\theta)$, the foc takes the form

$$P^*\left(\{\dot{q}_n(\check{\theta}^*) + \check{\delta}_n \dot{g}(\check{\theta}^*)\}(1-\mathbb{I}_n^*) = 0\right) \xrightarrow{a.s.} 1,$$

where $\check{\delta}_n \in \mathbb{R}$ are measurable Lagrange multipliers that can be determined, for outcomes in the event $\mathbb{I}_n^* = 1$, by involving also the constraint $b(\check{\theta}^*)(1-\mathbb{I}_n^*) = 0$.

A.3. Proof of Corollary 4.2

We only discuss the bootstrap validity part of the corollary, as the convergence part (4.3) was explained in the main text.

Let $\tau_n := \phi(n^{1/2}(\hat{\theta} - \theta_0)) + o_p(1)$, $\tau_n^* := \phi(n^{1/2}(\hat{\theta}^* - \hat{\theta})) + o_p(1)$ and $\tau := \phi(\ell(\theta_0))$. Convergence (4.3) and the continuity of ϕ imply that

$$(\tau_n, (\tau_n^*|D_n)) \xrightarrow{w}_w (\tau, (\tau|M)).$$

If the (random) cdf of $\tau|M$ is sample-path continuous, bootstrap validity follows from Theorem 3.1 and Lemma A.2(b) of Cavaliere and Georgiev (2020). We reduce the general case to the globally continuous case by a local argument for the cdf's $F(\cdot) := P(\tau \leq \cdot)$ and $F_M(\cdot) := P(\tau \leq \cdot|M)$. For concreteness, we focus on the technically more involved possibility $g(\theta_0) = 0$, such that $\theta_0 \in \partial\Theta$ given the assumption $\dot{g} \neq 0$. With

$$l(B) := \dot{g}_\perp(\dot{g}_\perp' B \dot{g}_\perp)^{-1}\dot{g}_\perp' B^{1/2}\xi + B^{-1}\dot{g}(\dot{g}'B^{-1}\dot{g})^{-1}\max\{0, \dot{g}'B^{-1/2}\xi\},$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

for positive definite $B \in \mathbb{R}^{2\times 2}$ and with $\ell = l(M)$, notice the following. If B is a fixed positive definite matrix such that

$$P(\phi(l(B)) = a) > 0, \tag{A.6}$$

for some $a \in \mathbb{R}$, then by equivalence (i.e., mutual absolute continuity) considerations for nonsingular Gaussian distributions, also

$$P(\phi(l(D)) = a) > 0,$$

for any positive definite $D \in \mathbb{R}^{2\times 2}$. In fact, let $\psi : \mathbb{R} \to \mathbb{R}$ be defined as $\psi(\cdot) := \phi(\dot{g}_\perp(\cdot))$ and let $\phi^{\leftarrow}(\cdot), \psi^{\leftarrow}(\cdot)$ denote inverse images. Then the probability in (A.6) equals

$$\begin{aligned}
&P\left(l(B) \in \phi^{\leftarrow}(\{a\}) \cap \partial\Lambda\right) + P(l(B) \in \phi^{\leftarrow}(\{a\}) \cap \mathrm{int}\Lambda) \\
&\quad = P(\{\dot{g}'B^{-1/2}\xi \le 0\} \cap \{(\dot{g}'_\perp B\dot{g}_\perp)^{-1}\dot{g}'_\perp B^{1/2}\xi \in \psi^{\leftarrow}(\{a\})\}) \\
&\qquad + P(\{\dot{g}'B^{-1/2}\xi > 0\} \cap \{B^{-1/2}\xi \in \phi^{\leftarrow}(\{a\})\}) \\
&\quad = P(\dot{g}'B^{-1/2}\xi \le 0)P((\dot{g}'_\perp B\dot{g}_\perp)^{-1}\dot{g}'_\perp B^{1/2}\xi \in \psi^{\leftarrow}(\{a\})\}) \\
&\qquad + P(B^{-1/2}\xi \in \phi^{\leftarrow}(\{a\}) \cap \mathrm{int}\Lambda),
\end{aligned}$$

the equality because $Cov(\dot{g}'B^{-1/2}\xi, (\dot{g}'_\perp B\dot{g}_\perp)^{-1}\dot{g}'_\perp B^{1/2}\xi) = 0$ and ξ is Gaussian. In the previous display, $P(\dot{g}'B^{-1/2}\xi \le 0) = P(N(0, \dot{g}'B^{-1}\dot{g}) \le 0) > 0$ for all positive definite B,

$$P((\dot{g}'_\perp B\dot{g}_\perp)^{-1}\dot{g}'_\perp B^{1/2}\xi \in \psi^{\leftarrow}(\{a\})\}) = P(N(0, (\dot{g}'_\perp B\dot{g}_\perp)^{-1}) \in \psi^{\leftarrow}(\{a\})),$$

is either 0 for all positive definite B or positive for all positive definite B, and the same applies to $P(B^{-1/2}\xi \in \phi^{\leftarrow}(\{a\}) \cap \mathrm{int}\Lambda)$. Therefore, the sign of the probability in (A.6) is the same (zero or positive) for all positive definite B.

The cdf F_M is a measurable transformation of M determined a.s. uniquely by the distribution of (M, ξ); it can be identified (up to a set of measure zero) as

$$F_M(\cdot) = P(\phi(l(B)) \le \cdot)|_{B=M},$$

by virtue of the independence of M and ξ. Since M is positive definite a.s., from the argument in the previous paragraph, we can conclude that every point on the line is either a discontinuity point of almost all sample paths of F_M, or a continuity point of almost all sample paths of F_M. By averaging, a point on the line is a discontinuity point of F if and only if it is a discontinuity point of almost all sample paths of F_M.

Let now q_0 be an interior point of the set

$$C = \{q \in (0,1) : \lim_{n\to\infty} P(F(\tau_n) \le q) \to q | \mathsf{H}_0\},$$

such that the asymptotic test is correctly sized for $q \in (q_0 - 2\epsilon, q_0 + 2\epsilon) \subset (0,1)$ for some $\epsilon > 0$. As $\tau_n \xrightarrow{w} \tau \sim F$, this implies that F and (by the discussion in previous paragraph) F_M skip no values from the interval $(q_0 - 2\epsilon, q_0 + 2\epsilon)$ (for F_M, a.s.). In particular, almost all sample paths of F_M are continuous on the (random) open superset $\left(F_M^{-1}(q_0 - \frac{3}{2}\epsilon), F_M^{-1}(q_0 + \frac{3}{2}\epsilon)\right)$ of $I_\epsilon := [F_M^{-1}(q_0 - \epsilon), F_M^{-1}(q_0 + \epsilon)]$, with

$$F_M^{-1}\left(q_0 - \frac{3}{2}\epsilon\right) < F_M^{-1}(q_0 - \epsilon) < F_M^{-1}(q_0 + \epsilon) < F_M^{-1}\left(q_0 + \frac{3}{2}\epsilon\right) \quad \text{a.s.} \tag{A.7}$$

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Without loss of generality, ϵ can be considered such that $q_0 \pm \epsilon$ are continuity points of F_M^{-1} a.s. (because F_M^{-1} is chosen to be càdlàg and its discontinuity points on, say $[\frac{q_0}{2}, \frac{q_0+1}{2}]$ are countably many). Let $\Psi^-(a,x)$ and $\Psi^+(a,x)$ be generalized inverses of the cdf's of a standard Gaussian variable conditioned to take values respectively in $(-\infty, a]$ and $[a, \infty)$. On extensions of the probability spaces where the data and (τ, M) are defined, consider a $U_{[0,1]}$ variable υ. Define $F_n^*(\cdot) := P^*(\tau_n^* \leq \cdot)$, $I_{n,\epsilon} := [F_n^{*-1}(q_0-\epsilon), F_n^{*-1}(q_0+\epsilon)]$ and

$$\begin{aligned}
\tilde{\tau}_n &= \tau_n \mathbb{I}_{\{\tau_n \in I_{n,\epsilon}\}} + \Psi^-(F_n^{*-1}(q_0-\epsilon), \upsilon)\mathbb{I}_{\{\tau_n < F_n^{*-1}(q_0-\epsilon)\}} \\
&\quad + \Psi^+(F_n^{*-1}(q_0+\epsilon), \upsilon)\mathbb{I}_{\{\tau_n > F_n^{*-1}(q_0+\epsilon)\}}, \\
\tilde{\tau}_n^* &= \tau_n^* \mathbb{I}_{\{\tau_n^* \in I_{n,\epsilon}\}} + \Psi^-(F_n^{*-1}(q_0-\epsilon), \upsilon)\mathbb{I}_{\{\tau_n^* < F_n^{*-1}(q_0-\epsilon)\}} \\
&\quad + \Psi^+(F_n^{*-1}(q_0+\epsilon), \upsilon)\mathbb{I}_{\{\tau_n^* > F_n^{*-1}(q_0+\epsilon)\}}, \\
\tilde{\tau} &= \tau \mathbb{I}_{\{\tau \in I_\epsilon\}} + \Psi^-(F_M^{-1}(q_0-\epsilon), \upsilon)\mathbb{I}_{\{\tau < F_M^{-1}(q_0-\epsilon)\}} \\
&\quad + \Psi^+(F_M^{-1}(q_0+\epsilon), \upsilon)\mathbb{I}_{\{\tau > F_M^{-1}(q_0+\epsilon)\}}.
\end{aligned}$$

Then

$$(\tilde{\tau}_n, (\tilde{\tau}_n^* | D_n)) \xrightarrow{w}_w (\tilde{\tau}, (\tilde{\tau} | M)),$$

because

$$(f_1(\tilde{\tau}_n), E\{f_2(\tilde{\tau}_n^*) | D_n\}) \xrightarrow{w} (f_1(\tilde{\tau}_n), E\{f_2(\tilde{\tau}) | M\}),$$

for any continuous and bounded real functions f_1, f_2, as a result of (A.5) and the fact that $P(\tau = F_M^{-1}(q_0 \pm \epsilon)|M) = 0$ a.s. by sample-path continuity of F_M on an open superset of I_ϵ. As the cdf of $\tilde{\tau}|M$ is a.s. sample-path continuous by construction, it follows that $P^*(\tilde{\tau}_n^* \leq \tilde{\tau}_n) \xrightarrow{w} U_{[0,1]}$, by Theorem 3.1 and Lemma A.2(b) of Cavaliere and Georgiev (2020).

Let $\tilde{F}_n^*(\cdot) := P^*(\tilde{\tau}_n^* \leq \cdot)$. We now return to the original variables. By considerations of equalities of events, it holds that

$$P(F_n^*(\tau_n) \leq q_0) = P(F_n^*(\tilde{\tau}_n) \leq q_0) = P(\tilde{F}_n^*(\tilde{\tau}_n) \leq q_0) = P(P^*(\tilde{\tau}_n^* \leq \tilde{\tau}_n) \leq q_0) \to q_0,$$

using the fact that $P^*(\tilde{\tau}_n^* \leq \tilde{\tau}_n) \xrightarrow{w} U_{[0,1]}$. This completes the proof.

SUPPLEMENTAL MATERIALS

The supplementary material for this article can be found at https://doi.org/10.1017/S0266466624000331.

REFERENCES

Andrews, D. W. K. (1999). Estimation when a parameter is on a boundary. *Econometrica*, 67, 1341–1383.

Andrews, D. W. K. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica*, 68, 399–405.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press

Berti, P., Pratelli, L., & Rigo, P. (2006). Almost sure weak convergence of random probability measures. *Stochastics: An International Journal of Probability and Stochastic Processes*, 78, 91–97.

Cavaliere, G., Nielsen, H. B., Pedersen, R. S. & Rahbek, A. (2022). Bootstrap inference on the boundary of the parameter space, with application to conditional volatility models. *Journal of Econometrics*, 227, 241–263.

Cavaliere, G., Nielsen, H. B., & Rahbek, A. (2017). On the consistency of bootstrap testing for a parameter on the boundary of the parameter space. *Journal of Time Series Analysis*, 38, 513–534.

Cavaliere, G., Perera, I., & Rahbek, A. (2024). Specification tests for GARCH processes with nuisance parameters on the boundary. *Journal of Business & Economic Statistics*, 42, 197–214.

Cavaliere, G., & Georgiev, I. (2020). Inference under random limit bootstrap measures. *Econometrica*, 88, 2547–2574.

Chatterjee, A., & Lahiri, S. N. (2011). Bootstrapping Lasso estimators. *Journal of the American Statistical Association*, 106, 608–25.

Chen, Q., & Fang, Z. (2019). Inference on functionals under first order degeneracy. *Journal of Econometrics*, 210, 459–481.

Demetrescu, M., Georgiev, I., Taylor, A. M. R. & Rodrigues, P. M. M. (2023). Extensions to IVX methods of inference for return predictability. *Journal of Econometrics*, 237 (2, Part C).

Doko Tchatoka, F., & Wang, W. (2021). Uniform inference after pretesting for exogeneity with heteroskedastic data, MPRA paper 106408, University Library of Munich, Germany.

Dümbgen, L. (1993). On nondifferentiable functions and the bootstrap. *Probability Theory and Related Fields*, 95, 125–140.

Fang, Z. (2014). Optimal plug-in estimators of directionally differentiable functionals. Unpublished manuscript.

Fang, Z., & Santos, A. (2019). Inference on directionally differentiable functions. *Review of Economic Studies*, 86, 377–412.

Georgiev, I., Harvey, D., Leybourne, S., & Taylor, A. M. R. (2019). A bootstrap stationarity test for predictive regression invalidity. *Journal of Business & Economic Statistics*, 37, 528–541.

Geyer, C. J. (1994). On the asymptotics of constrained M-estimation. *Annals of Statistics*, 22, 1993–2010.

Hirano, K., & Porter, J. R. (2012). Impossibility results for nondifferentiable functionals. *Econometrica*, 80, 1769–1790.

Hong, H., & Li, J. (2020). The numerical bootstrap. *Annals of Statistics*, 48, 397–412.

Kallenberg, O. (1997). *Foundations of modern probability*. Springer-Verlag.

Kallenberg, O. (2017). *Random measures: Theory and applications*. Springer-Verlag.

Müller, U. K., & Watson, M. (2008). Testing models of low frequency variability. *Econometrica*, 76, 979–1016.

Phillips, P. C. B. (2014). On confidence intervals for autoregressive roots and predictive regression. *Econometrica*, 82, 1177–1195.

Romano, J. P., Shaikh, A. M., & Wolf, M. (2014). A practical two-step method for testing moment inequalities. *Econometrica*, 82, 1979–2002.

Sweeting, T. J. (1989). On conditional weak convergence. *Journal of Theoretical Probability*, 2, 461–474.

https://doi.org/10.1017/S0266466624000331 Published online by Cambridge University Press