A Helicopter Tour on Random Limit Bootstrap Measures

Giuseppe Cavaliere

Department of Economics, University of Bologna
and Exeter Business School

SIdE Webinars Series
8 October 2020

Aim of this talk

- Classic (first-order) validity of the bootstrap usually refers to the case where the distribution of the bootstrap statistic (conditional on the original sample) replicates the asymptotic distribution of an estimator or a test statistic.
- In many situations, however, the (conditional) distribution of the bootstrap statistic is *random in the limit* (as $n \to \infty$) ...
  ... and bootstrap inference is commonly regarded as invalid.
- Tha aim of this talk is to give a helicopter tour on inference under *random limit bootstrap measures*.

Aim of this talk

- Main takeaways from this seminar:
  – Lack of classic bootstrap validity does not imply that the bootstrap cannot provide reliable inference;
  – In fact, bootstrap validity can be proved even in cases where the limit bootstrap measure is random;
  – Need a twist in what we mean by bootstrap validity – focus on *bootstrap p-values*;
  – New (old) machinery needed: theory of random measures (Kallenberg)
- Presentation based on 2013–2020 joint works with Iliyan Georgiev,[2013,16,18,20] Rob Taylor,[2013,16,18] Anders Rahbek[2015,20] and Peter Boswijk[2020]

Background

- Starting point: *the bootstrap is usually considered a device to estimate the asymptotic distribution of a statistic of interest*.
- Consider a sample $\{y_1, ..., y_n\}$ and a statistic

$$\tau_n := f(y_1, ..., y_n)$$

with asymptotic distribution

$$\tau_n \xrightarrow{w} \tau_\infty$$
$$G_n(x) := P(\tau_n \leq x) \to P(\tau_\infty \leq x) =: G_\infty(x)$$

- Let $\{y_1^*, ..., y_n^*\}$ be the bootstrap sample, which depends on $\{y_1, ..., y_n\}$ and a set of bootstrap innovations $\{\eta_1^*, ..., \eta_n^*\}$, independent of the original sample. The bootstrap analog of τ_n is

$$\tau_n^* := f(y_1^*, ..., y_n^*) = f(y_1, ..., y_n, \eta_1^*, ..., \eta_n^*)$$

Its distribution, given the original data, is

$$G_n^*(x) := P^*(\tau_n^* \leq x) = P(\tau_n^* \leq x | \{y_1, ..., y_n\})$$

Background

- 'Bootstrap consistency': the bootstrap estimates the asymptotic distribution of τ_n, that is
$$G_n^*(x) := P^*(\tau_n \le x) \to_p G_\infty(x)$$
(often denoted as $\tau_n^* \overset{w^*}{\to}_p \tau_\infty$, 'weak convergence in probability')
- Continuity of G_∞ also implies the stronger result (Polya's theorem):
$$\sup_{x \in \mathbb{R}} |P^*(\tau_n^* \le x) - P(\tau_n \le x)| \to_p 0$$
- Key implication: the bootstrap p-value satisfies
$$p_n^* := P^*(\tau_n^* \le x)|_{x=\tau_n} = G_n^*(\tau_n) \to_w U[0,1]$$

Background

- Interestingly, there are many cases where $G_n^*(\cdot)$ converges (in some sense) to a random process;
- Examples are application of the bootstrap to:
  - first-order autoregressions with a unit root (Basawa et al, 1991[AoS]);
  - when a parameter is on the boundary of the parameter space (Andrews, 2000[ECMA]);
  - models with infinite variance innovations (Knight, 1988[AoS]);
  - tests on the co-integrating relations in VAR models (Cavaliere et al, 2015[ECMA]);
  - 2SLS estimators under weak instruments;
  - cube-root consistent estimators (Cattaneo et al., 2020[ECMA]);
  - inference after model selection;
  - fixed b asymptotics (Lahiri, 2010; Shao and Politis, 2013[JRSSb]);
  - Hodges-LeCam superefficient estimators (Beran, 1997[AISM])
  - proxy SVARs with weak instruments (Fanelli et al. 2020)
- This automatically implies that the bootstrap is not consistent for the asymptotic distribution of the statistic of interest:

$$G_n^*(x) := P(\tau_n^* \leq x | y_1, ..., y_n) \nrightarrow_p G_\infty(x)$$

Background - Unit root test

- First order autoregression with a unit root (Basawa et al, 1991^{AoS}):

$$y_t = \alpha y_{t-1} + \varepsilon_t,\ \alpha = 1,\ \varepsilon_t \text{ i.i.d.}(0, \sigma^2)$$

  - J_c: Ornstein-Uhlenbeck process with mean reversion parameter c (Brownian Motion for $c = 0$)

$$H_c(x) := P\left(\int J_c dJ_c \Big/ \int J_c^2 \leq x\right)\ [H_0(\cdot) = \text{cdf of DF distribution}]$$

  - With $\hat{\alpha} := OLS(y_t | y_{t-1})$,

$$\tau_n := n(\hat{\alpha} - 1) \rightarrow_w \tau_\infty := \int J_0 dJ_0 \Big/ \int J_0^2 du, \text{ i.e. } P(\tau_n \leq x) \rightarrow H_0(x)$$

- (Fully parametric) Bootstrap:
  - $y_t^* = \hat{\alpha} y_{t-1}^* + \varepsilon_t^*$, ε_t^* i.i.d. $N(0, \sigma^2)$
  - bootstrap estimator $\hat{\alpha}^* := \mathrm{OLS}(y_t^* | y_{t-1}^*)$ and test statistic $\tau_n^* := n(\hat{\alpha}^* - \hat{\alpha})$.
  - As $n \rightarrow \infty$, the limit bootstrap measure is *random*:

$$P^*(\tau_n^* \leq x) \rightarrow_w H_{\tau_\infty}(x) = P\left(\int J_{\tau_\infty} dJ_{\tau_\infty} \Big/ \int J_{\tau_\infty}^2 du \leq x \,\Big|\, \tau_\infty\right)$$

- Wayout: a different bootstrap scheme (see several papers on bootstrap UR testing); m/n bootstrap

Background - Parameters on the boundary

- Simple location model (Andrews, 2000^{ECMA});
  - $y_t = \theta + \varepsilon_t$, ε_t i.i.d.$(0,1)$, $\theta \in [0, \infty)$.
  - Gaussian (Q)MLE: $\hat{\theta}_n = \max\{0, \bar{y}_n\}$, with asymptotic distribution

$$\tau_n := \sqrt{n}(\hat{\theta}_n - \theta) \to_w \begin{cases} Z & \text{if } \theta > 0 \\ Z^+ := \max\{0, Z\} & \text{if } \theta = 0 \end{cases} \quad (Z \sim N(0,1))$$

- Bootstrap:
  - $y_t^* = \hat{\theta}_n + \varepsilon_t^*$, ε_t^* i.i.d. $N(0,1)$ (independent of the original data) $\hat{\theta}_n^* = \max\{0, \bar{y}_n^*\}$ with test statistic $\tau_n^* = \sqrt{n}(\hat{\theta}_n^* - \hat{\theta}_n)$ with

$$P^*(\tau_n^* \leq x) = P^*(\sqrt{n}(\hat{\theta}_n^* - \hat{\theta}_n) \leq x) = \Phi(x) I(x \geq -\sqrt{n}\hat{\theta}_n)$$

  - When $\theta > 0$,

$$P^*(\tau_n^* \leq x) \to_p \Phi(x) = P(Z \leq x)$$

  - However, when $\theta = 0$,

$$P^*(\sqrt{n}(\hat{\theta}_n^* - \hat{\theta}_n) \leq x) \to_w \Phi(x)(x \geq -Z^+)$$

    or, equivalently, with Z^+, Z independent,

$$\tau_n^* \overset{w^*}{\to}_w \max\{-Z^+, Z\} | Z^+$$

- Solution: a different bootstrap scheme (see Cavaliere-Nielsen-Rahbek, 2016^{JTSA})

Background - Infinite variance

- $y_t = \mu + \varepsilon_t$, ε_t *symmetric* with $E\varepsilon_t^2 = +\infty$ and in the domain of attraction of a stable law with index $\alpha \in (0, 2)$, $\varepsilon_t \in \mathcal{D}(\alpha)$
  - With $\hat{\mu}_n := \bar{y}_n = n^{-1} \sum_{t=1}^n y_t$ it holds that

$$\tau_n := a_n^{-1} n(\hat{\mu}_n - \mu) \xrightarrow{w} S(\alpha)$$

- Bootstrap: $y_t^* = \hat{\mu}_n + \varepsilon_t^*$, where ε_t^* i.i.d. from the centred residuals $\hat{\varepsilon}_t := y_t - \hat{\mu}_n$
  - the bootstrap statistic is

$$\tau_n^* := a_n^{-1} n(\hat{\mu}_n^* - \hat{\mu}_n),\ \hat{\mu}_n^* := \bar{y}_n^* = n^{-1} \sum\nolimits_{t=1}^n y_t^*.$$

  - Then (Knight, 1989^{AoS})

$$P^*(\tau_n^* \leq x) \to_w P\left(\sum\nolimits_{k=1}^\infty \delta_k Z_k (M_k^* - 1) \leq x | Z_1, Z_2, ...\right)$$

    where: M_k^* iid Poisson(1); δ_k iid with $P(\delta_k = 1) = P(\delta_k = -1) = \frac{1}{2}$; $\{Z_k^{-\alpha}\}_{k=1}^\infty$ arrival times of a standard Poisson process.
  - $Z_\infty := (Z_1, Z_2, ...)'$: limit of the normalized order statistic of the $|\varepsilon_t|$'s

- Wayout: m out of n bootstrap

Background - Non-stationary stochastic vol [NSV]

- Location model with NSV (Hansen, 1995[ECMA]; Boswijk et al., 2020)

$$y_t = \mu + \sigma_t \varepsilon_t,\ E_{t-1}(\varepsilon_t) = 0,\ E_{t-1}(\varepsilon_t^2) = 1$$

- The assumption of non-stationary stochastic vol:
  - With $\sigma_n(u)$ the $D[0,1]$ approximant of $\{\sigma_t\}$

$$\sigma_n(u) := \sigma_{\lfloor nu \rfloor},\ u \in [0,1]$$

  it holds that $\sigma_n(\cdot) \to_w \sigma(\cdot) \in \mathcal{C}[0,1]$.

- Examples:
  - Hansen's NSV: $\sigma_t := h(c_0 + \frac{c_1}{n^{1/2}} \sum_{i=1}^t e_i)$, $\sigma_n(\cdot) \to_w \sigma(\cdot) := h(c_0 + c_1 B(\cdot))$
  - near integrated GARCH: $\omega = a_0 n^{-1/2}$, $\alpha_1 = a_1 n^{-1/2}$, $\beta = 1 - \alpha + a_2 n^{-1}$ (Nelson, 1991[JE]); $\sigma(\cdot)$ is a diffusion

- The standardized estimator of the mean is *mixed normal* as $n \to \infty$:

$$\tau_n := \sqrt{n}(\hat{\mu}_n - \mu) = n^{-1/2} \sum_{t=1}^n \sigma_t \varepsilon_t \to_w \tau_\infty = \int_0^1 \sigma(u)\, \mathrm{d}B(u)$$

$$\equiv MN(0, V),\ \text{with } V = \int_0^1 \sigma^2(u) \mathrm{d}u$$

Background - Non-stationary stochastic vol [NSV]

- (Gaussian) Wild Bootstrap (Wu, Liu, Mammen, Hansen, etc.):

$$y_t^* := \varepsilon_t \eta_t^*,\ \eta_t^* \text{ iid } N(0,1)\text{, independent of } \{y_t\}$$

with bootstrap statistic

$$\tau_n^* := n^{-1/2} \sum_{t=1}^n y_t^* = n^{-1/2} \sum_{t=1}^n \varepsilon_t \eta_t^*$$

- Conditional on the original data, with $\hat{V}_n := n^{-1} \sum_{t=1}^n \varepsilon_t^2$,

$$\tau_n^* | \{y_t\} \sim N(0, \hat{V}_n),$$

or equivalently (with Φ denoting the $N(0,1)$ cdf)

$$P^*(\tau_n^* \leq x) = P^*(N(0, \hat{V}_n) \leq x) = \Phi(\hat{V}_n^{-1/2} x),\ x \in \mathbb{R}$$

- Since $\hat{V}_n \to_w V := \int \sigma^2(u)\,\mathrm{d}u$, we have a random limit:

$$P^*(\tau_n^* \leq x) \to_w \Phi(V^{-1/2} x) = P(N(0,V) \leq x | V),\ x \in \mathbb{R}$$
$$\tau_n^* \overset{w^*}{\to}_w N(0,V)\,|V$$

The bootstrap does not replicate the asymptotic mixed normality of τ_n.

Background and key insights

- In all the previous examples, the presence of a *random* limiting distribution for the bootstrap statistic invalidates classic bootstrap inference (the bootstrap does not estimate the unconditional distribution of the statistic of interest).
- However, I will illustrate why this fact does not automatically imply that bootstrap inference is not valid:
  - the bootstrap can still deliver confidence intervals (or hypothesis tests) with the desired coverage probability (or size) when $n \to \infty$;
  - bootstrap inference may also have the appealing asymptotic interpretation of a *conditional inferential procedure ...*

    ...hence delivering efficiency (or power) gains over standard unconditional inference.

Sketch of what follows

- Random bootstrap measures in a Gaussian regression model and new definitions of bootstrap validity
- A first result on bootstrap validity under random bootstrap measures
  - Revisiting infinite variance case
- A second result on bootstrap validity
  - The case of co-integrating regressions
- Application: tests for parameter constancy in a linear model with non-stationarity regressors

Random bootstrap measures in a Gaussian linear model

- Consider a simple linear model

$$y_t = \beta x_t + \varepsilon_t, \qquad t = 1, 2, ..., n$$

where ε_t's are i.i.d. $\mathrm{N}(0,1)$ and x_t's are observable random variables, independent of the unobservable ε_t's. We assume further that

$$M_n := \sum_{t=1}^{n} x_t^2 > 0, \ a.s.$$

- Let $\hat{\beta} = \mathrm{OLS}(y_t|x_t)$ and $\tau_n := \hat{\beta} - \beta$ (its distribution is unknown)
- Important fact (see later): *conditional on the x_t's,*

$$\tau_n|\{x_t\} = M_n^{-1} \sum_{t=1}^{n} x_t \varepsilon_t \Big| M_n \sim N\left(0, M_n^{-1}\right) \Big| M_n = M_n^{-1/2} N\left(0,1\right) \Big| M_n$$

- In terms of CDFs,

$$P(\tau_n \leq x|\{x_t\}) = \Phi(M_n^{1/2} x), \ x \in \mathbb{R}$$

Random bootstrap measures in a Gaussian linear model

- Classic (parametric) fixed-design bootstrap sample (e.g. Hall, 1992):

$$y_t^* = \hat{\beta}' x_t + \varepsilon_t^* \qquad (t = 1, 2, ..., n)$$

where ε_t^* is i.i.d. N$(0, 1)$, independent of the original data.

- With $\hat{\beta}^* = \mathrm{OLS}(y_t^* | x_t)$ and $\tau_n^* := \hat{\beta}^* - \hat{\beta}$, *conditionally on the original data*,

$$\tau_n^* | \{y_t, x_t\} = M_n^{-1} \sum\nolimits_{t=1}^{n} x_t \varepsilon_t^* | \{y_t, x_t\} \sim N\left(0, M_n^{-1}\right) | M_n,$$

- That is,

$$P^*(\tau_n^* \le x) = P(\tau_n^* \le x | \{x_t, y_t\}) = \Phi(M_n^{1/2} x),\ x \in \mathbb{R}$$

- Important fact: the distribution the bootstrap statistic τ_n^*, *conditionally on the data*, coincides with the distribution of the original statistic τ_n, *conditionally on* $\{x_t\}$:

$$G_n^*(x) := P^*(\tau_n^* \le x) = P(\tau_n \le x | \{x_t\}) = \Phi(M_n^{1/2} x)$$

- Another important fact: even in this very simple example, *the limit bootstrap measure can be random*.

Random bootstrap measures in a Gaussian linear model: Fact 1

- *Case 1 (stationary&ergodic regressor)*:

$$M_n \to_p M > 0 \ (M \text{ non-stochastic})$$

Then,

$$G_n^*(x) = \Phi(M_n^{1/2}x) \to_p \Phi(M^{1/2}x),\ x \in \mathbb{R}$$

$$\tau_n^* \xrightarrow{w^*}_p N\left(0, M^{-1}\right)$$

- Notice that in this case, $\tau_n \to_w N\left(0, M^{-1}\right)$.
- Hence, classic bootstrap consistency holds.

Random bootstrap measures in a Gaussian linear model: Fact 1

- *Case 2 (nonstationary regressor)*:

$$M_n \to_w M > 0 \ (M \ stochastic)$$

For example, $x_t = T^{-1/2} \sum_{i=1}^{t} \varepsilon_i$. Then,

$$G_n^* (x) = \Phi(M_n^{1/2} x) \to_w \Phi(M^{1/2} x),\ x \in \mathbb{R}$$

$$\tau_n^* \overset{w^*}{\to}_w N\left(0, M^{-1}\right) |M$$

(*weak convergence in distribution*)

- Notice that in Case 2, $\tau_n \to_w MN\left(0, M^{-1}\right)$
- Hence the limit bootstrap measure does not replicate the asymptotic distribution of τ_n.
- Is the bootstrap really doing badly here?

Random bootstrap measures in a Gaussian linear model: Fact 2

- Twist the focus on the bootstrap p-value,

$$p_n^* = P^*(\tau_n^* \leq x)|_{x=\tau_n} = G_n^*(\tau_n)$$

- Using $G_n^*(x) = \Phi(M_n^{1/2}x)$ and $\tau_n|\{x_n\} \sim [M_n^{-1/2}N(0,1)]|M_n$,

$$p_n^* = G_n^*(\tau_n) = \Phi(M_n^{1/2}\tau_n) \overset{d}{=} \Phi(M_n^{1/2}M_n^{-1/2}N(0,1)) = \Phi(N(0,1)) \sim U[0,1]$$

- This same result holds even *conditionally* on M_n:

$$p_n^*|M_n \sim U[0,1]$$

- Notice that the usual bootstrap consistency,

$$\sup_{x\in\mathbb{R}} |P^*(\tau_n^* \leq x) - P(\tau_n \leq x)| \to_p 0$$

does not hold here.

Random bootstrap measures in a Gaussian linear model: Fact 3

- Why do we obtain uniformity of the p-values?
  - The key fact (to be exploited later for general applications) is that

$$\begin{pmatrix} P(\tau_n \le x|\{x_t\}) \\ P^*(\tau_n^* \le x) \end{pmatrix} = \begin{pmatrix} \Phi(M_n^{1/2}x) \\ \Phi(M_n^{1/2}x) \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \Phi(M_n^{1/2}x)$$

    such that the bootstrap replicates a particular conditional distribution of τ_n (and not its asymptotic/unconditional distribution)
  - Going to the limit

$$\begin{pmatrix} P(\tau_n \le x|\{x_t\}) \\ P^*(\tau_n^* \le x) \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \Phi(M_n^{1/2}x) \to_w \begin{pmatrix} 1 \\ 1 \end{pmatrix} \Phi(M^{1/2}x)$$

  - Equivalently ('weak convergence in distribution')

$$\begin{pmatrix} \tau_n|\{x_t\} \\ \tau_n^*|\{x_t, y_t\} \end{pmatrix} \overset{w}{\rightarrow}_w \begin{pmatrix} 1 \\ 1 \end{pmatrix} N\left(0, M^{-1}\right)|M$$

- Thus, the bootstrap is consistent for a limiting *conditional* distribution of τ_n (but not for its asymptotic distribution).

(Re-)defining bootstrap validity

- Let $\tau_n = \tau_n(D_n)$ be the original statistic and $\tau_n^* = \tau_n^*(D_n, W_n^*)$ the bootstrap statistic:
  - D_n denotes the data;
  - W_n^* are auxiliary bootstrap variates.
- For instance,
  - in the regression model case, $D_n := \{y_t, x_t\}_{t=1}^n$ and $W_n^* = (\varepsilon_1^*, ..., \varepsilon_n^*)'$ (iid $N(0,1)$ bootstrap shocks);
  - in the NSV case, $D_n := \{y_t\}_{t=1}^n$, $W_n^* = (w_1^*, ..., w_n^*)'$ (Rademacher bootstrap shocks).
- We also consider a (possibly unobservable) random element X_n (on the same probability space of D_n and W_n^*):
  - in the regression model case $X_n = (x_1, ..., x_n)'$;
  - in the NSV case, $X_n := \{\sigma_t\}_{t=1}^n$.

(Re-)defining bootstrap validity

- Let $p_n^* := P^*(\tau_n^* \leq \tau_n)$ denote the bootstrap p-value. Let X_n be a random element (e.g. a function of the data).

Definition

(i) Bootstrap inference is asymptotically *unconditionally* valid if

$$P\left(p_n^* \leq q\right) \rightarrow q \text{ (all } q \in (0,1))$$

so that p_n^* is asymptotically $U(0,1)$.
(ii) Bootstrap inference is asymptotically valid *conditionally on* X_n if

$$P\left(p_n^* \leq q \middle| X_n\right) \xrightarrow{p} q \text{ (all } q \in (0,1))$$

so that p_n^* is asymptotically $U(0,1)$ distributed *conditionally on* X_n.

Assessing bootstrap (conditional) validity

- Based on Cavaliere and Georgiev (2020[ECMA]) [GC]
- We start by providing sufficient conditions for the bootstrap to be valid conditionally.
- Such conditions only require (joint) weak convergence of the distribution of the bootstrap statistic, conditional on the original data, and of a conditional distribution of the original statistic.

Theorem

If, as $n \to \infty$, τ_n and τ_n^* satisfy

$$\left[\begin{array}{c} P\left(\tau_n \leq u \middle| X_n\right) \\ P\left(\tau_n^* \leq u \middle| D_n\right) \end{array} \right] \xrightarrow{w} \left[\begin{array}{c} 1 \\ 1 \end{array} \right] F(u)$$

in $D(\mathbb{R}) \times D(\mathbb{R})$, where F is a random cdf with a.s. continuous sample paths, then:

$$\sup_{u \in \mathbb{R}} \left| P\left(\tau_n^* \leq u \middle| D_n\right) - P\left(\tau_n \leq u \middle| X_n\right) \right| \xrightarrow{p} 0.$$

Assessing bootstrap (conditional) validity

- A consequence of the previous result is the following.

Theorem (cont'd)

Moreover,

$$p_n^* | X_n \xrightarrow{w}_p U(0,1),$$

so the bootstrap is asymptotically valid conditionally on X_n, and hence unconditionally:

$$p_n^* \xrightarrow{w} U(0,1).$$

- If $P(\tau_n \leq u | X_n)$ and $P^*(\tau_n^* \leq u)$ do not converge to the same limit, the bootstrap cannot be valid conditionally on X_n (it could be valid conditionally on some other random element)

Infinite variance (reprise): bootstrap validity

- Consider the following wild bootstrap (Cavaliere, Georgiev and Taylor, 2013ER, 2016^{AoS}):

$$\tilde{\tau}_n^* := a_n^{-1} \textstyle\sum_{t=1}^n \varepsilon_t^*, \text{ where } \varepsilon_t^* = \hat{\varepsilon}_t w_t^*,\ w_t^* \text{ i.i.d.}\,(0,1)\ \text{ Rademacher}$$

  (wild bootstrap with Rademacher bootstrap shocks)

- CGT^{03} show that

$$P^* \left(\tilde{\tau}_n^* \leq x \right) = P \left(\frac{1}{a_n} \sum_{t=1}^n \hat{\varepsilon}_t w_t^* \leq x \middle| \{\varepsilon_t\} \right) \xrightarrow{w} P \left(\sum_{t=1}^{\infty} \delta_t Z_t \leq x \middle| Z \right)$$

- Again, a random limiting distribution. Similar to Aue et al. (2008^{BER}) for permutation CUSUM tests.

Infinite variance (reprise): bootstrap validity

- Why is this bootstrap better than the i.i.d. bootstrap?
- In the symmetric case ($P(\delta_t = 1) = P(\delta_t = -1) = \frac{1}{2}$), as in Knight (1989),

$$P(\tau_n \leq x \mid |\varepsilon_1|, ..., |\varepsilon_n|) \xrightarrow{w} P\left(\sum_{t=1}^{\infty} \delta_t Z_t \leq x \middle| Z \right)$$

  which is a random distribution....
- ... identical to the weak limit of the Rademacher-wild bootstrap statistic

$$P^*(\tilde{\tau}_n^* \leq x) \xrightarrow{w} P\left(\sum_{t=1}^{\infty} \delta_t Z_t \leq x \middle| Z \right)$$

  where the limit is a continuous random cdf (a.s.)
- The two convergences are joint:

$$\begin{pmatrix} P(\tau_n \leq x | \{|\varepsilon_t|\}) \\ P^*(\tilde{\tau}_n^* \leq x) \end{pmatrix} \rightarrow_w \begin{pmatrix} 1 \\ 1 \end{pmatrix} P\left(\sum_{t=1}^{\infty} \delta_t Z_t \leq x \middle| Z \right)$$

- According to the previous theorem:
  - the bootstrap *mimics a particular conditional distribution of the original statistic*, i.e. the distribution of S_n *conditional* on $\{|\varepsilon_t|\}$;
  - bootstrap *p*-values are uniformly distributed, conditional on $\{|\varepsilon_t|\}$.

Infinite variance (reprise): coverage in finite samples

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage		
α	n	asympt.	iid BS	Wild BS
0.75	20	0.959	0.887	
0.75	100	0.959	0.904	
0.75	500	0.958	0.906	
1.00	20	0.950	0.899	
1.00	100	0.951	0.913	
1.00	500	0.951	0.915	
1.25	20	0.957	0.908	
1.25	100	0.956	0.925	
1.25	500	0.957	0.925	
1.50	20	0.950	0.917	
1.50	100	0.949	0.932	
1.50	500	0.950	0.934	
2.00	20	0.950	0.929	
2.00	100	0.950	0.945	
2.00	500	0.950	0.949	

Infinite variance (reprise): coverage in finite samples

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage		
α	n	asympt.	iid BS	Wild BS
0.75	20	0.959	0.887	0.955
0.75	100	0.959	0.904	0.964
0.75	500	0.958	0.906	0.960
1.00	20	0.950	0.899	0.943
1.00	100	0.951	0.913	0.953
1.00	500	0.951	0.915	0.952
1.25	20	0.957	0.908	0.936
1.25	100	0.956	0.925	0.950
1.25	500	0.957	0.925	0.950
1.50	20	0.950	0.917	0.931
1.50	100	0.949	0.932	0.948
1.50	500	0.950	0.934	0.948
2.00	20	0.950	0.929	0.927
2.00	100	0.950	0.945	0.945
2.00	500	0.950	0.949	0.950

Conditional bootstrap validity and conditional inference

- Under conditional validity,

$$p_n^* | X_n \rightarrow_w U(0, 1).$$

- This implies that the bootstrap can be seen as a tool for conditional inference.
- Hence, we expect efficiency (or power) gains over standard unconditional inference.
- Price to pay (if any): power functions/length of confidence sets are random (they depend on X_n)
  - in the InfV example, they depend on $\{|\varepsilon_t|\}$
  - in the linear model example, they depend on M_n

  also in the limit $n \rightarrow \infty$.

Infinite variance (reprise): efficiency gains

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage			COND vs UNC CI widths		
α	n	asympt.	iid BS	Wild BS	25%	median	75%
0.75	20	0.959	0.887	0.955			
0.75	100	0.959	0.904	0.964			
0.75	500	0.958	0.906	0.960			
1.00	20	0.950	0.899	0.943			
1.00	100	0.951	0.913	0.953			
1.00	500	0.951	0.915	0.952			
1.25	20	0.957	0.908	0.936			
1.25	100	0.956	0.925	0.950			
1.25	500	0.957	0.925	0.950			
1.50	20	0.950	0.917	0.931			
1.50	100	0.949	0.932	0.948			
1.50	500	0.950	0.934	0.948			
2.00	20	0.950	0.929	0.927			
2.00	100	0.950	0.945	0.945			
2.00	500	0.950	0.949	0.950			

Infinite variance (reprise): efficiency gains

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage			COND vs UNC CI widths		
α	n	asympt.	iid BS	Wild BS	25%	median	75%
0.75	20	0.959	0.887	0.955			
0.75	100	0.959	0.904	0.964			
0.75	500	0.958	0.906	0.960			
1.00	20	0.950	0.899	0.943			
1.00	100	0.951	0.913	0.953			
1.00	500	0.951	0.915	0.952			
1.25	20	0.957	0.908	0.936			
1.25	100	0.956	0.925	0.950			
1.25	500	0.957	0.925	0.950			
1.50	20	0.950	0.917	0.931			
1.50	100	0.949	0.932	0.948			
1.50	500	0.950	0.934	0.948			
2.00	20	0.950	0.929	0.927	0.843	0.948	1.057
2.00	100	0.950	0.945	0.945	0.940	0.991	1.044
2.00	500	0.950	0.949	0.950	0.972	1.000	1.031

Infinite variance (reprise): efficiency gains

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage			COND vs UNC CI widths		
α	n	asympt.	iid BS	Wild BS	25%	median	75%
0.75	20	0.959	0.887	0.955			
0.75	100	0.959	0.904	0.964			
0.75	500	0.958	0.906	0.960			
1.00	20	0.950	0.899	0.943			
1.00	100	0.951	0.913	0.953			
1.00	500	0.951	0.915	0.952			
1.25	20	0.957	0.908	0.936			
1.25	100	0.956	0.925	0.950			
1.25	500	0.957	0.925	0.950			
1.50	20	0.950	0.917	0.931	0.397	0.503	0.687
1.50	100	0.949	0.932	0.948	0.418	0.514	0.694
1.50	500	0.950	0.934	0.948	0.422	0.518	0.691
2.00	20	0.950	0.929	0.927	0.843	0.948	1.057
2.00	100	0.950	0.945	0.945	0.940	0.991	1.044
2.00	500	0.950	0.949	0.950	0.972	1.000	1.031

Infinite variance (reprise): efficiency gains

($N = 50,000$ MC simulations, $B = 999$ bootstrap repetitions; ε_t symmetric)

		coverage			COND vs UNC CI widths		
α	n	asympt.	iid BS	Wild BS	25%	median	75%
0.75	20	0.959	0.887	0.955	0.020	0.043	0.118
0.75	100	0.959	0.904	0.964	0.021	0.043	0.113
0.75	500	0.958	0.906	0.960	0.021	0.043	0.111
1.00	20	0.950	0.899	0.943	0.096	0.159	0.306
1.00	100	0.951	0.913	0.953	0.100	0.162	0.308
1.00	500	0.951	0.915	0.952	0.101	0.161	0.303
1.25	20	0.957	0.908	0.936	0.196	0.275	0.439
1.25	100	0.956	0.925	0.950	0.205	0.283	0.440
1.25	500	0.957	0.925	0.950	0.207	0.285	0.439
1.50	20	0.950	0.917	0.931	0.397	0.503	0.687
1.50	100	0.949	0.932	0.948	0.418	0.514	0.694
1.50	500	0.950	0.934	0.948	0.422	0.518	0.691
2.00	20	0.950	0.929	0.927	0.843	0.948	1.057
2.00	100	0.950	0.945	0.945	0.940	0.991	1.044
2.00	500	0.950	0.949	0.950	0.972	1.000	1.031

Bootstrap 'unconditional' validity

- Bootstrap validity 'unconditionally': $p_n^* \rightarrow_w U(0,1)$ (weaker than $p_n^*|X_n \rightarrow_w U(0,1)$)
- Consider the co-integrating regression

$$y_t = \beta x_t + \varepsilon_t$$

where:

  - $x_t = x_{t-1} + \eta_t$ $(x_0 = 0)$
  - $e_t := (\varepsilon_t, \eta_t)'$ stationary and ergodic MDS, $\Omega := Ee_t e_t' = \mathrm{diag}\{\omega_{\varepsilon\varepsilon}, \omega_{\eta\eta}\} > 0$.

- Object: bootstrap inference on β based on $\hat{\beta} = \mathrm{OLS}(y_t|x_t)$
- As is known (e.g. Chan and Wei, 1988^{AoS}),

$$n^{-1/2} \sum_{t=1}^{\lfloor n\cdot \rfloor} e_t \xrightarrow{w} (B_\varepsilon, B_\eta)'$$

$$\tau_n := n(\hat{\beta} - \beta) \xrightarrow{w} \left(\int B_\eta^2 \right)^{-1} \int B_\eta dB_\varepsilon \stackrel{d}{=} N(0, \omega_{\varepsilon\varepsilon} M^{-1}) \text{ (mixed Gaussian)}$$

due to the independence of B_η and B_ε (here $M := \int B_\eta^2$).

Bootstrap 'unconditional' validity

- *(Fixed regressor) bootstrap*. With $\hat{\omega}_{\varepsilon\varepsilon} := n^{-1}\sum_{t=1}^{n}(y_t - \hat{\beta}x_t)^2$,

$$y_t^* = \hat{\beta}x_t + \hat{\omega}_{\varepsilon\varepsilon}^{1/2}\varepsilon_t^*, \ \varepsilon_t^* \text{ i.i.d.} N(0,1) \text{ (independent of the data)}$$

  The bootstrap analogue of $\hat{\beta}$ is $\hat{\beta}^* = \mathrm{OLS}(y_t^*|x_t)$

- For $\tau_n^* := n(\hat{\beta}^* - \hat{\beta})$, as in the linear model example,

$$\tau_n^*|\{x_t, y_t\} \sim N(0, n\hat{\omega}_{\varepsilon\varepsilon}M_n^{-1})\big|\,(M_n, \hat{\omega}_{\varepsilon\varepsilon})$$

- Since $n^{-2}M_n := n^{-2}\sum_{t=1}^{n} x_t^2 \to_w M := \int B_\eta^2$, if $\hat{\omega}_{\varepsilon\varepsilon} \xrightarrow{p} \omega_{\varepsilon\varepsilon}$, by the CMT

$$P^*(\tau_n^* \leq x) = \Phi(\hat{\omega}_{\varepsilon\varepsilon}^{-1/2}(n^{-1}M_n^{1/2})x) \xrightarrow{w} \Phi(\omega_{\varepsilon\varepsilon}^{-1}M^{1/2}x),\ x \in \mathbb{R},$$

  or, equivalently,

$$\tau_n^* \xrightarrow{w^*}_w N(0, \omega_{\varepsilon\varepsilon}M^{-1})\big|\,M$$

Bootstrap 'unconditional' validity

- In summary,

$$\tau_n := n(\hat{\beta} - \beta) \to_w N(0, \omega_{\varepsilon\varepsilon} M^{-1}) \text{ (mixed Gaussian)}$$
$$\tau_n^* := n(\hat{\beta}^* - \hat{\beta}) \overset{w^*}{\to}_w N(0, \omega_{\varepsilon\varepsilon} M^{-1})\big| M \text{ (a component of the mixture)}$$

- The unconditional limit of $\tau_n := n(\hat{\beta} - \beta)$ obtains by integrating over M the conditional limit of $\tau_n^* := n(\hat{\beta}^* - \hat{\beta})$ given the data.
- This implies asymptotic *unconditional* validity of the bootstrap, see next slide. Notice that by direct evaluation

$$p_n^* = P^*(\tau_n^* \leq u)|_{u=\tau_n} = \Phi(\hat{\omega}_{\varepsilon\varepsilon}^{-1/2} M_n^{1/2}(\hat{\beta} - \beta))$$
$$\overset{w}{\to} \Phi((\omega_{\varepsilon\varepsilon} \textstyle\int B_\eta^2)^{-1/2} \int B_\eta dB_\varepsilon) = \Phi(N(0,1)) \overset{d}{=} U(0,1)$$

- Notice that without additional assumptions, the same result does not hold conditionally on $\{x_t\}$

Bootstrap 'unconditional' validity: a general result

- Let
  - τ_n original statistic; τ_n^* bootstrap statistic;
  - $G_n^*(x) := P^*(\tau_n^* \leq x)$ bootstrap (conditional) distribution.
- Sufficient conditions for bootstrap unconditional validity (Theorem 3.1 in CG):

Theorem

Suppose that for some random element X

$$\tau_n \rightarrow_w \tau_\infty, \qquad G_n^*(\cdot) \rightarrow_w G(\cdot) = P(\tau_\infty \leq \cdot | X) \qquad \text{(jointly)}$$

Then, if G has a.s. continuous sample paths, the bootstrap is valid unconditionally:

$$p_n^* := G_n^*(\tau_n) \rightarrow_w U[0,1]$$

- Notice that $P(\tau_\infty \leq x) = \int P(\tau_\infty \leq x | X) \, dP(X)$;
- It applies to the co-integrating regression case;
- This theorem does not imply conditional validity of the bootstrap.

Bootstrap 'unconditional' validity

- Bootstrap validity conditionally on $\{x_t\}$ requires strengthening of the assumptions.
- Assume that ε_t is MDS with respect to $\mathcal{G}_t = \sigma(\{\varepsilon\}_{s=-\infty}^{t} \cup \{\eta_s\}_{s\in\mathbb{Z}})$, and that $n^{-1}\sum_{t=1}^{n} E(\varepsilon_t^2|\{\eta_s\}_{s\in\mathbb{Z}}) \overset{a.s.}{\to} \omega_{\varepsilon\varepsilon}$. Then,

$$\begin{pmatrix} \tau_n|\{x_t\} \\ \tau_n^*|\{x_t, y_t\} \end{pmatrix} \overset{w}{\to}_w \begin{pmatrix} 1 \\ 1 \end{pmatrix} \left(\int_0^1 B_\eta^2\right)^{-1} \int_0^1 B_\eta dB_\varepsilon \bigg| B_\eta$$
$$= \begin{pmatrix} 1 \\ 1 \end{pmatrix} N(0, \omega_{\varepsilon\varepsilon} M^{-1}) \bigg| M$$

or, equivalently,

$$\begin{pmatrix} P(\tau_n \le \cdot|\{x_t\}) \\ P^*(\tau_n^* \le \cdot) \end{pmatrix} \overset{w}{\to} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \Phi(\omega_{\varepsilon\varepsilon}^{-1/2} M^{1/2} \cdot), \ \cdot \in \mathbb{R}$$

- Hence, the bootstrap is consistent for the limiting *conditional* distribution of $\tau_n|\{x_t\}$ and

$$p_n^*|\{x_t\} \overset{w}{\to}_p U(0,1)$$

‘Conditional’ vs ‘unconditional’ bootstrap validity

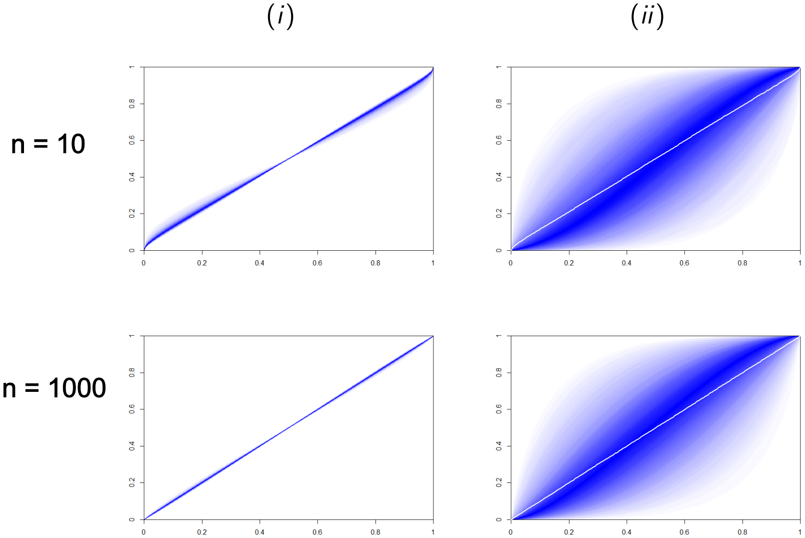

Figure: Fan chart of the simulated cdfs (conditional on X_n) of the bootstrap p-values.

Application: bootstrap tests of parameter constancy

- Consider the classical problem of parameter constancy testing in regression models (Chow, 1960[ECMA]; Quandt, 1960[JASA]; Nyblom, 1989[JASA]; Andrews, 1993[ECMA]; Andrews and Ploberger, 1994[ECMA]; Elliott and Müller, 2006[RES]; Perron and Qu, 2006[JoE],...).
- Specifically, we deal with bootstrap implementations when the moments of the regressors may be unstable over time (Hansen, 2000[JoE]).
- Model:

$$y_{nt} = \beta_t' x_{nt} + \varepsilon_{nt} \qquad (t = 1, 2, ..., n).$$

with null hypothesis $\mathrm{H}_0 : \beta_t = \beta_1 \, (t = 2, ..., n)$

- In order to test H_0 against H_1, we consider the 'sup F' test statistic

$$\mathcal{F}_n := \max_{r \in [\underline{r}, \overline{r}]} F_{\lfloor nr \rfloor}$$

Here $F_{\lfloor nr \rfloor}$ is the usual F statistic for testing the auxiliary null $\theta = 0$ in the regression

$$y_{nt} = \beta' x_{nt} + \theta' x_{nt} \mathbb{I}_{\{t \geq \lfloor rn \rfloor\}} + \varepsilon_{nt}$$

Bootstrap tests of parameter constancy

- A high level assumption on x_t and ε_t (e.g. allowing for unit root regressors or for infrequent random level shifts) is the following (Hansen, 2000[JoE])
- *Assumption $\mathcal{H}$*
  - *(mda) ε_{nt} is a martingale difference array (mda) with respect to the current value of x_{nt} and the lagged values of $(x_{nt}, \varepsilon_{nt})$;*
  - *(wlln) ε_{nt}^2 satisfies the law of large numbers $n^{-1} \sum_{t=1}^{\lfloor nr \rfloor} \varepsilon_{nt}^2 \xrightarrow{p} r(E\varepsilon_{nt}^2) = r\sigma^2 > 0$ as $n \to \infty$, for all $r \in (0, 1]$;*
  - *(non-stationarity) in $D_{m\times m} \times D_{m\times m} \times D_m$:*

$$\left(\frac{1}{n} \sum_{t=1}^{\lfloor n\cdot \rfloor} x_{nt} x_{nt}', \frac{1}{n\sigma^2} \sum_{t=1}^{\lfloor n\cdot \rfloor} x_{nt} x_{nt}' \varepsilon_t^2, \frac{1}{n^{1/2}\sigma} \sum_{t=1}^{\lfloor n\cdot \rfloor} x_{nt} \varepsilon_{nt} \right) \xrightarrow{w} (M, V, N)$$

    as $n \to \infty$, where M and V are a.s. continuous and (except at 0) strictly positive-definite valued processes, whereas N, conditionally on $\{V, M\}$, is a zero-mean Gaussian process with covariance kernel $E\{N(r_1) N(r_2)'\} = V(r_1)$ $(0 \le r_1 \le r_2 \le 1)$

Bootstrap tests of parameter constancy

- Asymptotic null distribution of $\mathcal{F}_n$:

$$\mathcal{F}_n \xrightarrow{w} \tau_\infty := \sup_{r\in[\underline{r},\overline{r}]} \left\{ \tilde{N}(r)' \tilde{M}(r)^{-1} \tilde{N}(r) \right\}$$

$$\begin{aligned} \tilde{N}(u) &:= N(u) - M(u) M(1)^{-1} N(1) \\ \tilde{M}(r) &:= M(r) - M(r) M(1)^{-1} M(r) \end{aligned}$$

- With $G_n(x) := P(\mathcal{F}_n \leq x)$ and $G_\infty(x) := P(\tau_\infty \leq x)$, we have that

$$G_n(x) \to G_\infty(x), \; x \in \mathbb{R}$$

- For (asymptotically) stationary regressors, it corresponds to the supremum of a squared tied-down Bessell process (Andrews, 1993[ECMA]).
- Quite impossible to simulate critical values for this asymptotic distribution.

Bootstrap tests of parameter constancy

- Fixed design/wild Bootstrap (Hansen, 2000^{JoE}; Goncalves and Kilian, 2004^{JoE}):
- Used to accommodate possible conditional heteroskedasticity of ε_{nt}
- Based on the OLS residuals $\tilde{e}_{nt}$ from the regression of y_{nt} on x_{nt} and $x_{nt}\mathbb{I}_{\{t \geq \lfloor \tilde{r}n \rfloor\}}$, where $\tilde{r} := \arg\max_{r \in [\underline{r}, \overline{r}]} F_{\lfloor nr \rfloor}$ is the estimated break fraction.
- Bootstrap sample:

$$y_t^* := \tilde{e}_{nt} w_t^*, \; w_t^* \text{ i.i.d. } N(0,1)$$

- Bootstrap statistic:

$$\mathcal{F}_n^* := \max_{r \in [\underline{r}, \overline{r}]} F_{\lfloor nr \rfloor}^*,$$

where $F_{\lfloor nr \rfloor}^*$ is the F statistic for the auxiliary null that $\theta^* = 0$ in the regression

$$y_t^* = \beta^{*\prime} x_{nt} + \theta^{*\prime} x_{nt} \mathbb{I}_{\{t \geq \lfloor rn \rfloor\}} + \text{error}_{nt}^*.$$

- The associated bootstrap p-value is

$$p_n^* := P(\mathcal{F}_n^* \geq \mathcal{F}_n | \{y_1, x_1, ..., y_n, x_n\}) =: P^*(\mathcal{F}_n^* \geq \mathcal{F}_n)$$

Bootstrap tests of parameter constancy

- The usual bootstrap validity argument requires a 'weak convergence in probability' statement like the following

$$\mathcal{F}_n^* \xrightarrow{w^*}_p \tau_\infty := \sup_{r \in [\underline{r}, \overline{r}]} \left\{ \tilde{N}(r)' \tilde{M}(r)^{-1} \tilde{N}(r) \right\}$$

- That is,

$$G_n^*(x) := P(\mathcal{F}_n^* \leq x | \{y_1, x_1, ..., y_n, x_n\}) =: P^*(\mathcal{F}_n^* \leq x) \rightarrow_p G_\infty(x)$$

  However, this case is much more involved.
- Results such as the previous one do not hold.

Bootstrap tests of parameter constancy

- Asymptotic properties of the bootstrap statistic $\mathcal{F}_n^*$:

Theorem

Under Assumption $\mathcal{H}$, it holds that

$$\mathcal{F}_n^* \xrightarrow{w^*}_w \sup_{r\in[\underline{r},\overline{r}]} \left\{ \tilde{N}(r)' \tilde{M}(r)^{-1} \tilde{N}(r) \right\} \Big| M, V$$

where $\tilde{M}(r) = M(r) - M(r)M(1)^{-1}M(r)$, $\tilde{N}(r) = N(r) - M(r)M(1)^{-1}N(1)$.

- This contradicts the claim in Hansen (2000^{JoE}):

$$\mathcal{F}_n^* \xrightarrow{w}_p \sup_{r\in[\underline{r},\overline{r}]} \left\{ \tilde{N}(r)' \tilde{M}(r)^{-1} \tilde{N}(r) \right\}$$

which is not correct as

  - the limiting distribution of the bootstrap statistic is random; hence the bootstrap does not estimate the unconditional distribution of the test statistic
  - convergence is weak in distribution, not weak in probability.

Bootstrap tests of parameter constancy

- Simulations in Hansen (2000) suggest that $p_n^* \xrightarrow{w} U(0,1)$. Can we (at least) save this result?
- The stated assumption are indeed sufficient for unconditional bootstrap validity (details in CG):

Theorem

Let the parameter constancy hypothesis H_0 *hold under assumption* $\mathcal{H}$. *Then, the bootstrap based on* $\tau_n = \mathcal{F}_n$ *and* $\tau_n^* = \mathcal{F}_n^*$ *is asymptotically valid unconditionally.*

- Hence, although Hansen's distributional result on the bootstrap statistic is not correct, the final claim on bootstrap unconditional validity holds.

Bootstrap tests of parameter constancy

- What about bootstrap validity conditional on x_{nt}?
- Not surprisingly, we need further assumptions
- **Assumption** $\mathcal{C}$. *Joint with the convergence facts in Assumption* $\mathcal{H}$,

$$\left(\frac{1}{n} \sum_{t=1}^{\lfloor n \cdot \rfloor} x_{nt} x_{nt}', \frac{1}{n\sigma^2} \sum_{t=1}^{\lfloor n \cdot \rfloor} x_{nt} x_{nt}' \varepsilon_{nt}^2, \frac{1}{n^{1/2}\sigma} \sum_{t=1}^{\lfloor n \cdot \rfloor} x_{nt} \varepsilon_{nt} \right) \Bigg| \{x_{nt}\} \xrightarrow{w}_w (M, V, N) \,|(M, V)$$

- Then (see CG):

Theorem

Let the parameter constancy hypothesis H_0 *hold under assumption* $\mathcal{H}$ *and* $\mathcal{C}$. *Then, the bootstrap based on* $\tau_n = \mathcal{F}_n$ *and* $\tau_n^* = \mathcal{F}_n^*$ *is asymptotically valid conditionally on* x_{nt}.

That is, $p_n^* | \{x_{nt}\} \xrightarrow{w}_w U(0,1)$.

Concluding remarks

- As is known, there are many cases where bootstrap statistics have random limiting distributions
- This feature invalidates the bootstrap as a mean of estimating the unconditional distribution of a statistic of interest.
- Don't panic – the bootstrap can still be (very) useful:
  - as a device to estimate a particular conditional distribution of the statistic of interest (hence, to draw *conditional* inferences);
  - as a device to obtain p-values and confidence intervals.
- The analysis of bootstrap validity requires non-standard tools (weak convergence of random measures).
- Currently dealing with:
  - proxy SVARs when the proxies are weak (with L. Fanelli and G. Angelini);
  - solution to the NSV example (if you would like to see it, please join SIdE's 9th ICEEE in Cagliari!)