Filtering With Confidence
In-sample Confidence Bands For GARCH Filters

Marcin Zamojski

UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

PDF @ zmks.co/fwc
slides @ zmks.co/fwc/slides
www @ zamojski.net

Local In TimE (LITE) bootstrap for observation-driven filters

- I propose i) **a novel bootstrap** procedure and ii) **a new resampling scheme**, which in combination, allow for **bootstrapping conditionally on observed data** (vs. the typical unconditional procedures).
- The resampler constructs series which mimic dependence patterns in the data by drawing with replacement normalized residuals in a moving window of a relatively short length.
- LITE produces **in-sample confidence bands around time-varying parameters in observation-driven methods** which account for both parameter and filtering uncertainty.
- Simulation evidence suggests that the obtained confidence bands have **good coverage properties**.
- LITE is both an easily implementable (in less than 20 lines of code; see slides) and a flexible method (applies to a wide range of observation-driven models).

GARCH(1,1) as an example observation-driven filter

The paper focuses on obtaining confidence bands for time-varying volatility in the standard GARCH(1,1) model originated by Engle (1982) and Bollerslev (1986):

$$y_t = \varepsilon_t,$$
$$\varepsilon_t = \eta_t \sqrt{\sigma_t^2}, \qquad \eta_t \sim \mathrm{N}(0,1),$$
$$\sigma_{t+1}^2 = \omega + \beta\sigma_t^2 + \alpha y_t^2.$$

However, the method can be adapted to other observation-driven filters. The underlying assumption in the paper is that the GARCH(1,1) is only a misspecified filter rather than a true data-generating process. In particular, in the true and unknown data-generating process the time-varying parameter (here, volatility) can evolve stochastically or in a deterministic fashion rather than being observation-driven.

Traditional bootstrap procedure

(Hall and Yao, 2003; Gonçalves and Kilian, 2004; Pascual et al., 2006)

1. Use GARCH(1,1) to get $\hat{\sigma}_t^2$ and fitted residuals $\hat{\varepsilon}_t$, get raw residuals: $\hat{\eta}_t = \hat{\varepsilon}_t(\hat{\sigma}_t)^{-1}$.
2. Resample raw residuals treating them as *i.i.d.* to get B bootstrap samples.
3. Use the GARCH(1,1) transition equation to get new paths of volatility, $\sigma_{t+1}^{2\,*} = \hat{\omega} + \hat{\beta}\sigma_t^{2*} + \hat{\alpha}y_t^{2*}$.
   - Estimated path with two bootstrap draws:

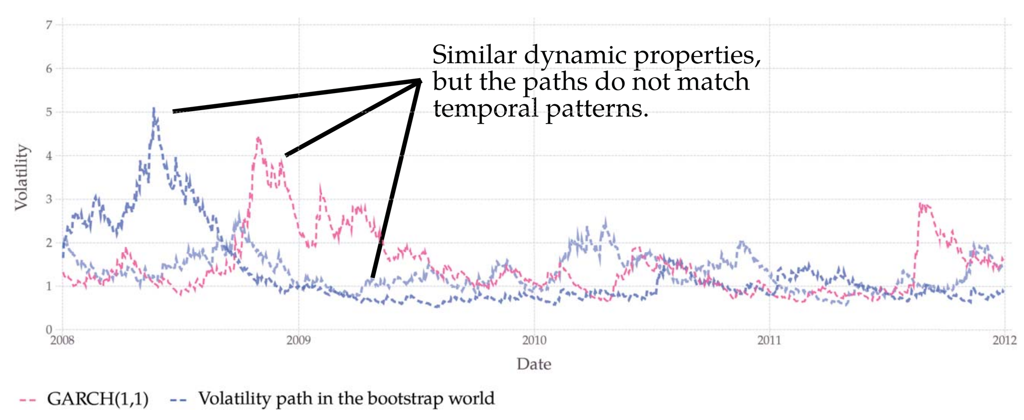

4. Fit GARCH(1,1) to all bootstrap samples.
5. Compute confidence intervals for... ω, α, and β.
   - Naïve bands:

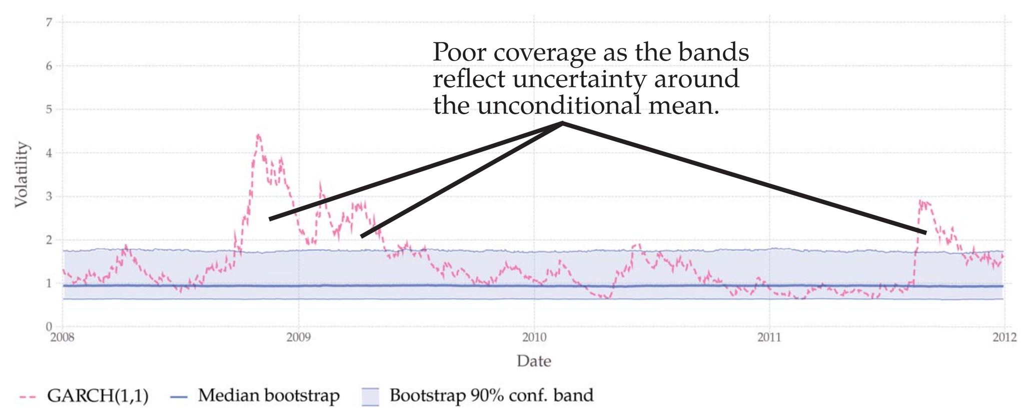

LITE bootstrap procedure

1. Use GARCH(1,1) to get $\hat{\sigma}_t^2$ and fitted residuals $\hat{\varepsilon}_t$, get raw residuals: $\hat{\eta}_t = \hat{\varepsilon}_t(\hat{\sigma}_t)^{-1}$.
2. Resample raw residuals with the LITE resampler to get B bootstrap samples.
3. Re-scale the raw bootstrap residuals with $\hat{\sigma}_t^2$, from (1.), $y_t^* = \hat{\sigma}_t\hat{\eta}_t^*$.
   - Estimated path with two bootstrap draws:

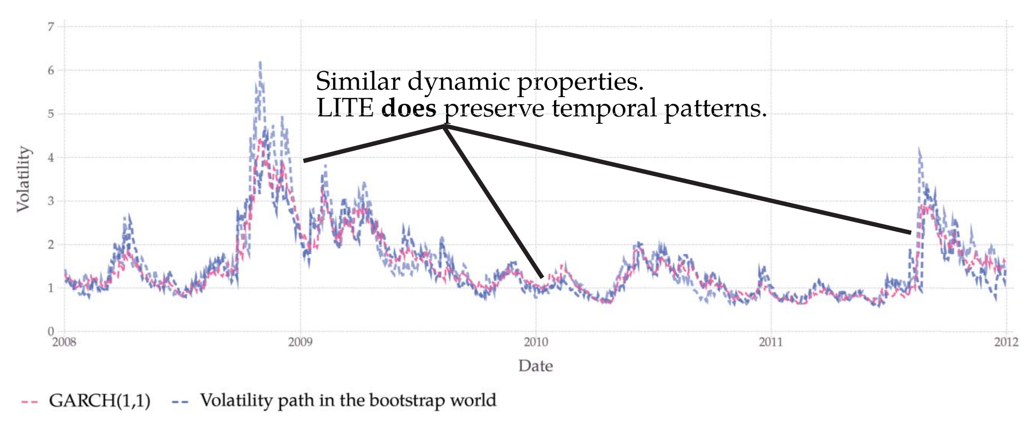

4. Fit GARCH(1,1) to all bootstrap samples.
5. Compute confidence intervals for... σ, ω, α, and β.
   - LITE bands:

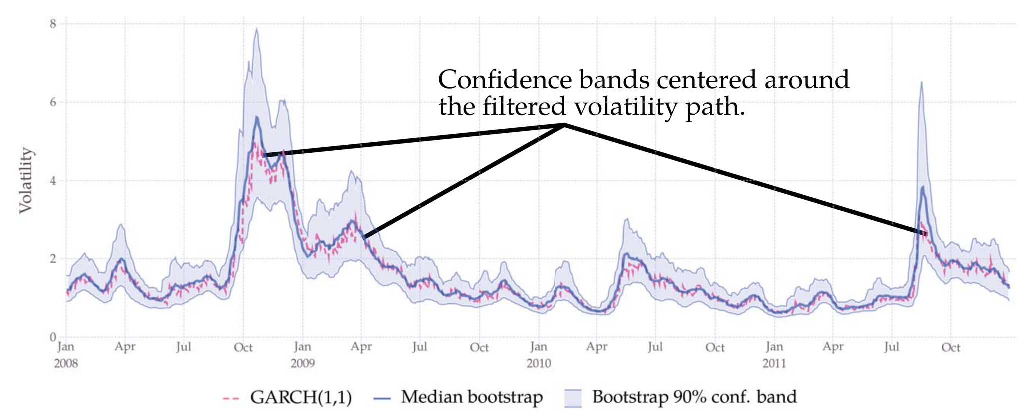

LITE resampler

The LITE resampler operates in a moving window and draws normalized residuals with replacement from adjacent points in time.

The figure shows tracking of the true volatility by the GARCH(1,1) filter and deviations in absolute raw residuals, $|\hat{\eta}_t| - |\eta_t|$ as a proxy for under-/over-scaling.

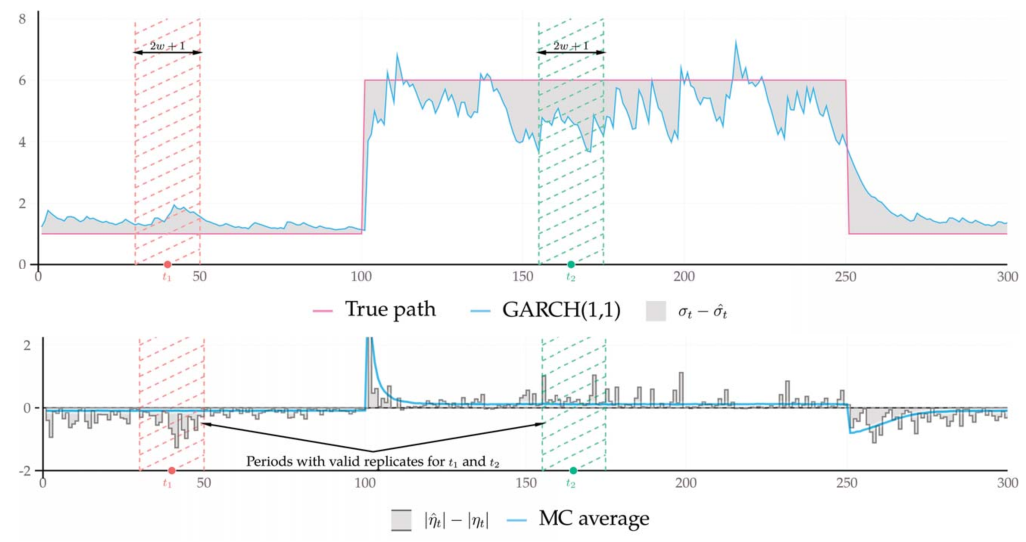

The shaded regions around points t_1 and t_2 show periods with valid replicates based on the resampling bandwidth of w. Note that within these two periods observations share three characteristics:

- the unobserved, true volatility values are close,
- the estimated values are close,
- in effect the estimated volatility path is locally affected by similar levels of misfit.

The raw residuals from the period around t_2 are not valid replicates for the residual at t_1 as they are under-scaled and, in effect, their distribution has higher standard deviation than is required to imitate the conditional distribution around t_1.

Coverage properties based on simulations

For a series of deterministic 'true' volatility paths I conduct a simulation study to measure real coverage of a 90% confidence interval. Results for an S&P 500 inspired example (in the paper) are shown below. The 'true path' is common across the replications.

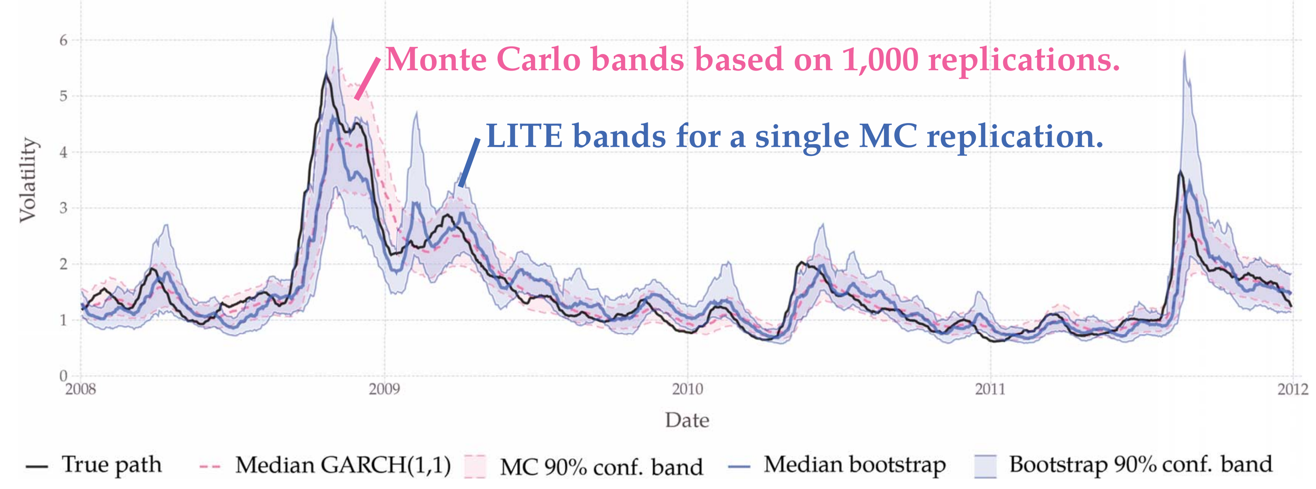

Coverage statistics:

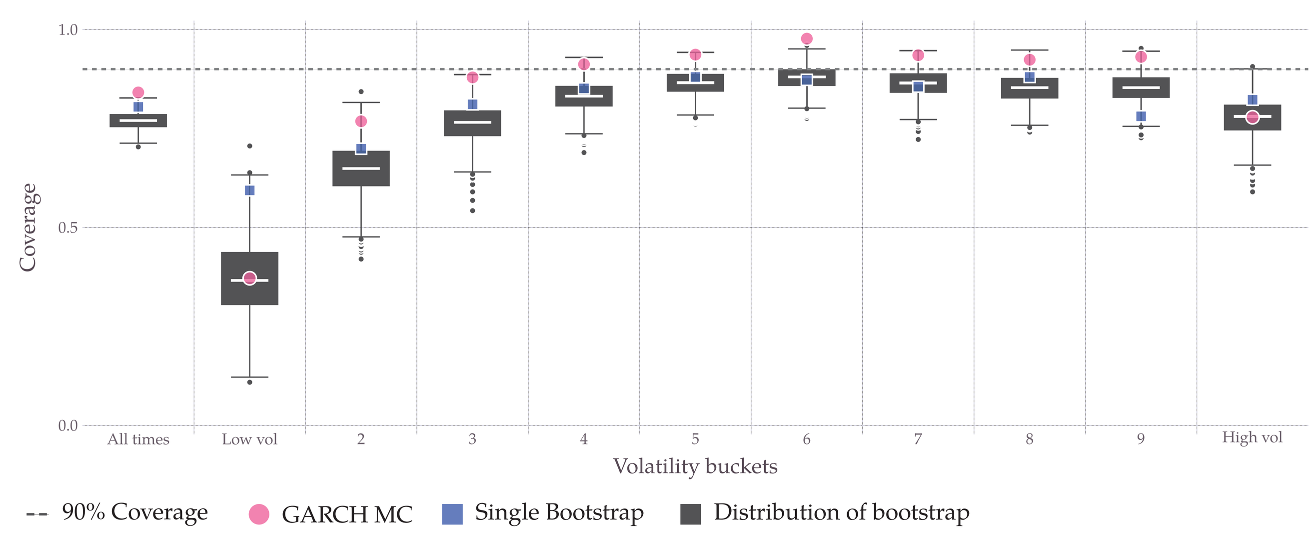

LITE bands have unconditional coverage of approx. 77% while the benchmark, GARCH MC, is at 84%. The coverage depends on the ability of the used filter to produce an unbiased estimate of the true volatility which (as in this case) may depend on the value of volatility itself.

Coverage convergence

Nelson (1992) shows that even if the GARCH filter is misspecified, the tracking quality improves as the sampling frequency increases. In this simulation study I increase the sampling frequency seven-fold from approx. 7,000 observations to 50,000.

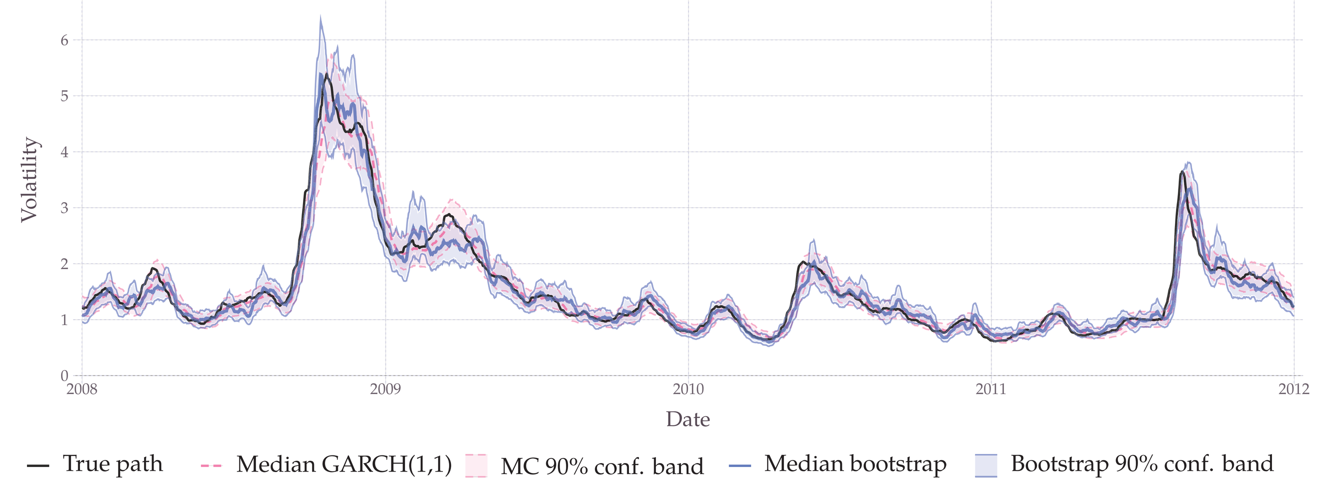

Coverage statistics:

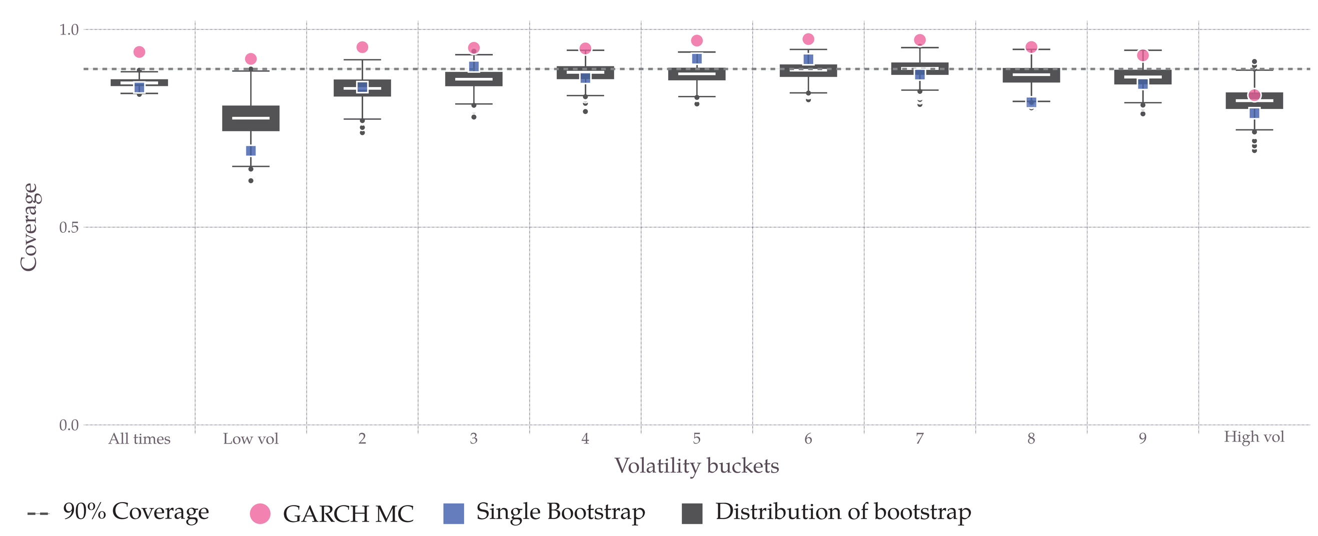

As sampling frequency is increased, the average path coverage of the LITE confidence bands converges to the nominal level. Here, the unconditional path coverage is on average 86%. The increase is mostly driven by vast improvements of coverage in extremely low and high volatility regimes.