Model Based Tests for Asset Price Bubbles

Robert Taylor

University of Essex

Based on joint (very preliminary!) work with Sam Astill, Dave Harvey and Steve Leybourne

ACE Opening Conference, Comwell Kellers Park
May 26–28 2025

Outline of Presentation

Table of Contents

- We develop tests for detecting asset price bubbles that are derived from the standard stock pricing equation commonly employed in the finance literature whereby stock prices are determined by expectations of the future price, future dividends and an unobserved component.
- The general solution to this equation shows that the price of an asset is the sum of an $I(1)$ fundamental price component and an explosive component that only takes non-zero values during bubble periods.
- While the current financial econometrics literature focusses almost exclusively on modelling asset prices as a single autoregressive process and testing the null of a unit root against the alternative of explosivity, this is not the underlying theory model for asset prices.
- We, instead, use the solution to the asset pricing equation to construct Locally Best Invariant [LBI] motivated tests designed to test the null that the innovations to the bubble component of the asset pricing equation have zero variance (and hence no bubble exists), against the alternative that these innovations have non-zero variance (and hence an asset price bubble component exists).

Introduction and Motivation

- Asset price bubbles, defined as a large upward price swing additional to the fundamental price of an asset, represent a misallocation of resources during the upwards bubble phase and a potentially catastrophic loss of value during the inevitable crash phase.
- Given the widespread damage caused by the collapse of asset price bubbles, their detection is a key element of macroprudential policy.
- Most developments in the literature have focussed on using right-tailed univariate unit root tests to test a single series for explosivity, with the price-dividend ratio of individual stocks or stock indices being a common focus.

Introduction and Motivation

- The earliest contribution in this area was made by Phillips, Wu and Yu (2011) [PWY] who developed a test for the null of no explosive behaviour against the alternative of explosivity based on a sequence of forward recursive right-tailed augmented Dickey-Fuller [ADF] test statistics.
- This methodology was further developed by Phillips, Shi and Yu (2015) [PSY] who propose tests based on either a sequence of backward recursive right-tailed ADF test statistics, or a test based on a double recursion across all possible start and end-dates (subject to a minumum window size).
- The PSY tests have subsequently become the industry standard for detecting historical bubbles or ongoing bubbles.

- In all of the aforementioned tests the focus is on treating the price series as a single autoregressive [AR] process and testing the null that the leading AR coefficient is equal to unity (no bubble) in all periods against the alternative that this coefficient is greater than unity (bubble) in some periods.
- The univariate TVAR model employed is, however, not the model implied by finance theory for asset prices during a rational bubble episode, where the asset price series is the sum of a fundamental $I(1)$ component and a separate bubble component.
- While extant tests have been demonstrated to exhibit excellent power when prices are assumed to follow a univariate TVAR process, their efficacy in detecting bubbles generated according to the aforementioned additive model may well be severely diminished.
- While the rational bubble model is very often used to motivate the need for existing tests for asset price bubbles, to our knowledge there has been no attempt in either the econometrics or empirical finance literature to develop tests that are specifically designed to test for bubbles in prices that take the form implied by the finance theory model.

Table of Contents

The Component Model

- The fundamental price of an asset is derived from the standard no arbitrage condition:
$$P_t = \frac{E_t\left[P_{t+1} + D_{t+1}\right]}{1+R} \tag{1}$$
where P_t denotes the price of the asset at time t, $R > 0$ is the constant risk-free rate, D_{t+1} denotes the dividend at time t and $E_t[.]$ denotes the expectation of its argument conditional on information available at time t.
- Forward iteration of (1) leads to the following equation determining the fundamental price of the asset,
$$P_t^f = \sum_{i=1}^{\infty} \frac{1}{(1+R)^i} E_t\left[D_{t+i}\right]. \tag{2}$$
- The fundamental price of the asset is therefore seen to be the present value of all expected future dividends.

- If we impose the transversality condition

$$\lim_{k \to \infty} E_t \left[\frac{1}{(1+R)^k} P_{t+k} \right] = 0 \tag{3}$$

this implies that $P_t = P_t^f$ is the unique solution to (1) and the asset is therefore free of bubbles.

- If Equation (3) does not hold, however, then P_t^f is not the only price process that solves equation (1).

The Component Model

- It can be verified that if a process $\{B_t\}_{t=1}^{\infty}$ is generated according to the submartingale process

$$E_t[B_{t+1}] = (1 + R)B_t \tag{4}$$

then adding B_t to P_t^f will yield another solution to equation (1). Homm and Breitung (2012) note that there are, in fact, infinitely many solutions that take the form

$$P_t = P_t^f + B_t \tag{5}$$

where $\{B_t\}_{t=1}^{\infty}$ is a process that satisfies equation (4).

- Equation (5) implies that the price of an asset can be decomposed into a fundamental component, P_t^f and a "bubble" component, B_t, that is explosive in expectation. If a bubble is present in the price then (4) implies that a rational investor will only be willing to hold the stock if they expect the bubble component to grow at rate R, as the investor is then compensated for the price paid for the stock in addition to its fundamental value.

The Component Model

- We need to specify an econometric model based on the theoretical model. To that end, for an asset price series, P_t, observed over the period $t = 1, ..., T$, we consider an *unobserved components model* of the form

$$P_t = \mu + F_t + B_t.$$

  Here μ is a constant term, while F_t and B_t represent, respectively, the unobserved fundamental and bubble components of P_t.

- Following the finance theory model, the fundamental component is assumed to follow a (martingale) random walk process,

$$F_t = F_{t-1} + \varepsilon_t$$

  across all t with $F_0 = 0$ and ε_t a zero mean white noise process with variance σ^2.

The Component Model

Assuming that the "bubble" component is active between two unknown dates $1 \leq t_{b1} < t_{b2} \leq T$, we can model it using an explosive AR(1) process

$$B_t = \begin{cases} \rho B_{t-1} + \eta_t & t = t_{b1}, ..., t_{b2} \\ 0 & \text{otherwise} \end{cases}$$

where $\rho > 1$ and η_t is also a zero mean white noise process with variance ω^2, independent of ε_t.

Table of Contents

A Locally Best Invariant Test

- We wish to test for the presence of the unobserved bubble component, B_t, in the asset price series, P_t. We can then represent the null hypothesis that no bubble is present as

$$H_0 : \omega^2 = 0$$

- The alternative of a price bubble over $t = t_{b1}, ..., t_{b2}$ is given by

$$H_1 : \omega^2 > 0$$

- In the infeasible case where ρ, t_{b1} and t_{b2} are assumed known, we can derive a Locally Best Invariant (LBI) test of H_0 against H_1, using the general testing approach of King and Hillier (1985).

A Locally Best Invariant Test

- If we assume that $\eta_t \sim NIID(0, \omega^2)$ and $\varepsilon_t \sim NIID(0, \sigma^2)$ we can express the model in obvious Tx1 vector notation as

$$P = \mu\mathbf{1} + F + B$$

  where $\mathbf{1}$ is a vector of 1s.

- Here, $P - \mu\mathbf{1} \sim N(0, Q(\omega^2))$ where $Q(\omega^2) = \sigma^2 D + \omega^2 A$, where

$$D := \begin{bmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 2 & 2 & \cdots & 2 \\ 1 & 2 & 3 & 3 & \cdots & 3 \\ 1 & 2 & 3 & 4 & \cdots & 4 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & 4 & \cdots & T \end{bmatrix}$$

A Locally Best Invariant Test

The matrix A has the block structure

$$A := \begin{bmatrix} 0_{(t_{b1}-1)\times(t_{b1}-1)} & 0_{(t_{b1}-1)\times(t_{b2}-t_{b1}+1)} & 0_{(t_{b1}-1)\times(T-t_{b2})} \\ 0_{(t_{b2}-t_{b1}+1)\times(t_{b1}-1)} & A_{t_{b1},t_{b2}} & 0_{(t_{b2}-t_{b1}+1)\times(T-t_{b2})} \\ 0_{(T-t_{b2})\times(t_{b1}-1)} & 0_{(T-t_{b2})\times(t_{b2}-t_{b1}+1)} & 0_{(T-t_{b2})\times(T-t_{b2})} \end{bmatrix}$$

where

$$A_{t_{b1},t_{b2}} := \begin{bmatrix} 1 & \rho & \cdots & \rho^{t_{b2}-t_{b1}} \\ \rho & 1+\rho^2 & \cdots & \rho^{t_{b2}-t_{b1}+1}+\rho^{t_{b2}-t_{b1}-1} \\ \vdots & \vdots & \ddots & \vdots \\ \rho^{t_{b2}-t_{b1}} & \rho^{t_{b2}-t_{b1}+1}+\rho^{t_{b2}-t_{b1}-1} & \cdots & 1+\rho^2(1+\rho^2+\rho^4+\ldots+\rho^{2(t_{b2}-t_{b1}-1)}) \end{bmatrix}.$$

A Locally Best Invariant Test

- In order to derive the LBI test using King and Hillier (1985), we need the variance of $P - \mu\mathbf{1}$ under the null to be of the form $\sigma^2 I_T$.
- Consider then $M(P - \mu\mathbf{1}) \sim N(0, MQ(\omega^2)M')$ where $M'M = D^{-1}$, i.e.

$$M := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & -1 & 1 \end{bmatrix}.$$

- Then, $MQ(\omega^2)M' = \sigma^2 MDM' + \omega^2 MAM' = \sigma^2 I_T + \omega^2 PAP'$, and it follows that $PQ(0)P' = \sigma^2 I_T$.

A Locally Best Invariant Test

Using King and Hillier, and defining $t_{b1} := \lfloor \tau_{b1} T \rfloor$ and $t_{b2} := \lfloor \tau_{b2} T \rfloor$, the LBI test of H_0 against H_1 rejects for large values of the statistic,

$$S_\rho(\tau_{b1}, \tau_{b2}) := \frac{\{M(P-\mu\mathbf{1})\}' MAM' \{M(P-\mu\mathbf{1})\}}{\{M(P-\mu\mathbf{1})\}'\{M(P-\mu\mathbf{1})\}} = \frac{\{M(P-\mu\mathbf{1})\}' HH' \{M(P-\mu\mathbf{1})\}}{\{M(P-\mu\mathbf{1})\}'\{M(P-\mu\mathbf{1})\}}$$

with

$$H := \begin{bmatrix} 0_{(t_{b1}-1)\times(t_{b1}-1)} & 0_{(t_{b1}-1)\times(t_{b2}-t_{b1}+1)} & 0_{(t_{b1}-1)\times(T-t_{b2})} \\ 0_{(t_{b2}-t_{b1}+1)\times(t_{b1}-1)} & H_{t_{b1},t_{b2}} & 0_{(t_{b2}-t_{b1}+1)\times(T-t_{b2})} \\ 0_{(T-t_{b2})\times(t_{b1}-1)} & 0_{(T-t_{b2})\times(t_{b2}-t_{b1}+1)} & 0_{(T-t_{b2})\times(T-t_{b2})} \end{bmatrix},$$

$$H_{t_{b1},t_{b2}} := \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ \rho-1 & 1 & 0 & \cdots & 0 \\ \rho(\rho-1) & \rho-1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{t_{b2}-t_{b1}-1}(\rho-1) & \rho^{t_{b2}-t_{b1}-2}(\rho-1) & \rho^{t_{b2}-t_{b1}-3}(\rho-1) & \cdots & 1 \end{bmatrix}$$

A Locally Best Invariant Test

… and where we note that

$$M(P - \mu \mathbf{1}) = \begin{bmatrix} P_1 - \mu \\ \Delta P_2 \\ \Delta P_3 \\ \vdots \\ \Delta P_T \end{bmatrix}.$$

A Locally Best Invariant Test

We can re-write $S_\rho(\tau_{b1}, \tau_{b2})$ in scalar summation notation as

$$S_\rho(\tau_{b1}, \tau_{b2}) = \frac{\left((P_1 - \mu) + (\rho - 1)\sum_{j=t_{b1}+1}^{t_{b2}} \rho^{j-t_{b1}-1}\Delta P_j\right)^2 + \sum_{t=t_{b1}+1}^{t_{b2}} \left(\Delta P_t + (\rho - 1)\sum_{j=t+1}^{t_{b2}} \rho^{j-t-1}\Delta P_j\right)^2}{(P_1 - \mu)^2 + \sum_{t=2}^{T}(\Delta P_t)^2}$$

when $t_b = 1$, and

$$S_\rho(\tau_{b1}, \tau_{b2}) = \frac{\sum_{t=t_{b1}}^{t_{b2}} \left(\Delta P_t + (\rho - 1)\sum_{j=t+1}^{t_{b2}} \rho^{j-t-1}\Delta P_j\right)^2}{(P_1 - \mu)^2 + \sum_{t=2}^{T}(\Delta P_t)^2}$$

when $t_{b1} > 1$.

A Locally Best Invariant Test

- For fixed ρ, in the numerator parts of $S_\rho(\tau_{b1}, \tau_{b2})$ the $(P_1 - \mu)$ and ΔP_t terms are of a lower stochastic order of magnitude than the $(\rho - 1) \sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j$ term.

- Further, in the denominator parts the $(P_1 - \mu)^2$ is dominated stochastically by $\sum_{t=2}^{T} (\Delta P_t)^2$.

A Locally Best Invariant Test

Omitting these asymptotically dominated terms we obtain a simplified variant of $S_\rho(\tau_{b1}, \tau_{b2})$, given by

$$S^*_\rho(\tau_{b1}, \tau_{b2}) := \begin{cases} \frac{(\rho-1)^2 \left(\sum_{j=t_{b1}+1}^{t_{b2}} \rho^{j-t_{b1}-1} \Delta P_j\right)^2 + (\rho-1)^2 \sum_{t=t_{b1}+1}^{t_{b2}-1} \left(\sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j\right)^2}{\sum_{t=2}^{T} (\Delta P_t)^2} & t_{b1} = 1 \\ \frac{(\rho-1)^2 \sum_{t=t_{b1}}^{t_{b2}-1} \left(\sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j\right)^2}{\sum_{t=2}^{T} (\Delta P_t)^2} & t_{b1} > 1 \end{cases}.$$

which can be written as

$$\begin{aligned} S^*_\rho(\tau_{b1}, \tau_{b2}) &= \frac{(\rho-1)^2 \sum_{t=t_{b1}}^{t_{b2}-1} \left(\sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j\right)^2}{\sum_{t=2}^{T} (\Delta P_t)^2} \\ &= \frac{(\rho-1)^2 \sum_{t=t_{b1}+1}^{t_{b2}} \left(\sum_{j=t}^{t_{b2}} \rho^{j-t} \Delta P_j\right)^2}{\sum_{t=2}^{T} (\Delta P_t)^2} \qquad (6) \end{aligned}$$

for all t_{b1}.

A Feasible Test

- In practice the value of ρ is unknown, as are the bubble start and end dates, t_{b1} and t_{b2}. We therefore next develop a feasible version of $S^*_\rho(\tau_{b1}, \tau_{b2})$.

- First we introduce notation for arbitrary bubble start and end dates, $t_1 := \lfloor \tau_1 T \rfloor$ and $t_2 := \lfloor \tau_2 T \rfloor$ which may or may not coincide with the true dates $t_{b1} = \lfloor \tau_{b1} T \rfloor$ and $t_{b2} = \lfloor \tau_{b2} T \rfloor$.

- The analogue of the statistic in (6) computed over this arbitrary sub-sample is then given by:

$$S^*_\rho(\tau_1, \tau_2) = \frac{(\rho - 1)^2 \sum_{t=t_1+1}^{t_2} \left(\sum_{j=t}^{t_2} \rho^{j-t} \Delta P_j \right)^2}{\sum_{t=2}^{T} (\Delta P_t)^2}.$$

A Feasible Test

- To be operational we must specify a value for ρ; we denote this by $\bar{\rho} > 1$.
- In order to obtain a statistic with a tractable limiting distribution, we set $\bar{\rho}$ local-to-unity, with the scaling appropriate to the sub-sample size that the numerator of the statistic is based upon. Specifically,

$$\bar{\rho} = 1 + \bar{c}(t_2 - t_1)^{-1}$$

  where $\bar{c} > 0$ denotes a user specified constant.
- The denominator of the statistic is $O_p(T)$ under the null, and so is scaled by T^{-1}. Together, this gives

$$S^*_{\bar{\rho}}(\tau_1, \tau_2) = \frac{(\bar{\rho}-1)^2 \sum_{t=t_1+1}^{t_2} \left(\sum_{j=t}^{t_2} \bar{\rho}^{j-t} \Delta P_j \right)^2}{T^{-1} \sum_{t=2}^{T} (\Delta P_t)^2}$$

  or, substituting for $\bar{\rho}$,

$$S^*_{\bar{c}}(\tau_1, \tau_2) = \frac{\bar{c}^2 (t_2 - t_1)^{-2} \sum_{t=t_1+1}^{t_2} \left(\sum_{j=t}^{t_2} \{1 + \bar{c}(t_2 - t_1)^{-1}\}^{j-t} \Delta P_j \right)^2}{T^{-1} \sum_{t=2}^{T} (\Delta P_t)^2}$$

A Feasible Test

- Finally, since the true bubble start and end dates are unknown, we then propose a feasible statistic that takes the maximum value of $S_{\bar{c}}^*(\tau_1, \tau_2)$ across all possible bubble timings, subject to a constraint on the minimum permitted bubble regime length.
- Given the exponential nature of the statistic (since $\bar{c} > 0$), we also apply a natural log transformation to $S_{\bar{c}}^*(\tau_1, \tau_2)$. The final statistic we propose is then given by:

$$S_{\bar{c}}^* = \sup_{\tau_1 \in [1/T, 1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi, 1]} \ln S_{\bar{c}}^*(\tau_1, \tau_2)$$

  where π represents then minimum permitted value for the window width fraction, $\tau_2 - \tau_1$.
- We therefore have a continuum of possible tests, indexed by $\bar{c}$.

Table of Contents

- Next consider the large sample behaviour of $S^*_{\bar{c}}$ under $H_1 : \omega^2 > 0$, where ρ is local-to-unity so that a well-defined asymptotic distribution for $S^*_{\bar{c}}$ can be obtained.

- Specifically, we set $\rho = 1 + c(t_{b2} - t_{b1})^{-1}$ with $c > 0$. We can set $\mu = 0$ without loss of generality. We make the following assumption regarding η_t and ε_t.

  Assumption 1. ε_t and η_t are independent MDSs with $E(\varepsilon_t^2) = \sigma^2$ and $E(\eta_t^2) = \omega^2$ and finite fourth order moments.

We then obtain the following large sample result,

Theorem 1. *Under H_1 and Assumption 1,*

$$S_{\bar{c}}^* \Rightarrow \sup_{\tau_1 \in [0,1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi,1]} \ln H_{c,\bar{c}}(\tau_1, \tau_2, \omega/\sigma, \tau_{b1}, \tau_{b2})$$

where

$$H_{c,\bar{c}}(\tau_1, \tau_2, \omega/\sigma, \tau_{b1}, \tau_{b2}) := \frac{\bar{c}^2(\tau_2 - \tau_1)^{-2} \int_{\tau_1}^{\tau_2} \left\{ \int_r^{\tau_2} e^{\bar{c}(s-r)(\tau_2-\tau_1)^{-1}} dK_c(s, \omega/\sigma, \tau_{b1}, \tau_{b2}) \right\}^2 dr}{1 + \left(\frac{\omega}{\sigma}\right)^2 (\tau_2 - \tau_1) + \left(\frac{\omega}{\sigma}\right)^2 \left\{ \int_{\tau_{b1}}^{\tau_{b2}} e^{c(\tau_{b2}-s)(\tau_{b2}-\tau_{b1})^{-1}} dW_\eta(s) \right\}^2}$$

with

$$K_c(r, \omega/\sigma, \tau_{b1}, \tau_{b2}) := W_\varepsilon(r) + \mathbb{I}(\tau_{b1} \le r \le \tau_{b2}) \frac{\omega}{\sigma} \int_{\tau_{b1}}^{r} e^{c(r-s)(\tau_{b2}-\tau_{b1})^{-1}} dW_\eta(s)$$

and where $W_\varepsilon(r)$ and $W_\eta(r)$ are independent standard Brownian motions.

Limiting Distribution

- The limiting null distribution of $S_{\bar{c}}^*$ follows on setting $\omega^2 = 0$ in the result in Theorem 1. That is,

$$S_{\bar{c}}^* \Rightarrow \sup_{\tau_1 \in [0,1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi,1]} \ln L_{\bar{c}}(\tau_1, \tau_2)$$

where

$$L_{\bar{c}}(\tau_1, \tau_2) := \bar{c}^2(\tau_2 - \tau_1)^{-2} \int_{\tau_1}^{\tau_2} \left\{ \int_r^{\tau_2} e^{\bar{c}(s-r)(\tau_2-\tau_1)^{-1}} dW_\varepsilon(s) \right\}^2 dr.$$

- Notice that the limiting null distribution of $S_{\bar{c}}^*$ depends on $\bar{c}$.
- Local power is seen to depend on $\bar{c}$, c, ω/σ (the signal-to-noise ratio), and the bubble start and end fractions, τ_{b1} and τ_{b2}.

Table of Contents

A Modified Statistic

- It is clear from the DGP that, with the exception of a bubble which is still on-going at the end of the sample, when the bubble terminates, a level shift occurs as the series returns from $P_t = \mu + B_t + F_t$ to $P_t = \mu + F_t$.

- This induces a one-time outlier in the first differences of P_t, with $\Delta P_{t_{b2}+1} = O_p(T^{1/2})$ and is responsible for the third term in the denominator of $H_{c,\bar{c}}(\tau_1, \tau_2, \omega/\sigma, \tau_{b1}, \tau_{b2})$.

- As this term is positive, it is to be expected that the power of the test would be increased by removing this outlier from the variance calculation in the denominator of the $S^*_{\bar{c}}(\tau_1, \tau_2)$ statistic.

A Modified Statistic

- We note that, for large T, $\max_{t\in[2,\ldots,T]} |\Delta P_t| = |\Delta P_{t_{b2}+1}|$ almost surely, since all other $|\Delta P_t|$ are $O_p(1)$. We therefore also consider a modified version of $S^*_{\bar{c}}(\tau_1, \tau_2)$ where we remove the largest absolute value of ΔP_t from the $T^{-1}\sum_{t=2}^{T}(\Delta P_t)^2$ calculation, i.e. we replace $T^{-1}\sum_{t=2}^{T}(\Delta P_t)^2$ with

$$T^{-1}\left\{\sum_{t=2}^{T}(\Delta P_t)^2 - \max_{t\in[2,\ldots,T]} |\Delta P_t|^2\right\}$$

- It is straightforward to show that

$$T^{-1}\left\{\sum_{t=2}^{T}(\Delta P_t)^2 - \max_{t\in[2,\ldots,T]} |\Delta P_t|^2\right\} \xrightarrow{p} \sigma^2 + \omega^2(\tau_{b2} - \tau_{b1})$$

A Modified Statistic

The modified statistic then becomes

$$S_{\bar{c}}^{\dagger} = \sup_{\tau_1 \in [1/T, 1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi, 1]} \ln S_{\bar{c}}^{\dagger}(\tau_1, \tau_2)$$

where

$$S_{\bar{c}}^{\dagger}(\tau_1, \tau_2) := \frac{\bar{c}^2 (t_2 - t_1)^{-2} \sum_{t=t_1+1}^{t_2} \left(\sum_{j=t}^{t_2} \{1 + \bar{c}(t_2 - t_1)^{-1}\}^{j-t} \Delta P_j \right)^2}{T^{-1} \left\{ \sum_{t=2}^{T} (\Delta P_t)^2 - \max_{t \in [2, \dots, T]} |\Delta P_t|^2 \right\}}$$

It follows that

$$\begin{aligned} S_{\bar{c}}^{\dagger}(\tau_1, \tau_2) \quad &\Rightarrow \quad \frac{\bar{c}^2 (\tau_2 - \tau_1)^{-2} \int_{\tau_1}^{\tau_2} \left\{ \int_r^{\tau_2} e^{\bar{c}(s-r)(\tau_2 - \tau_1)^{-1}} dK_c(s, \omega/\sigma, \tau_{b1}, \tau_{b2}) \right\}^2 dr}{1 + \left(\frac{\omega}{\sigma}\right)^2 (\tau_{b2} - \tau_{b1})} \\ &=: \quad G_{c,\bar{c}}(\tau_1, \tau_2, \omega/\sigma, \tau_{b1}, \tau_{b2}) \end{aligned}$$

and hence

$$S_{\bar{c}}^{\dagger} \Rightarrow \sup_{\tau_1 \in [0, 1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi, 1]} \ln G_{c,\bar{c}}(\tau_1, \tau_2, \omega/\sigma, \tau_{b1}, \tau_{b2})$$

Table of Contents

Monte Carlo Setup

Simulation data were generated according to

$$\begin{aligned} P_t &= F_t + B_t, \qquad t = 1, ..., T \\ F_t &= F_{t-1} + \varepsilon_t \end{aligned}$$

with $F_0 = 0$ and where

$$B_t = \begin{cases} \rho B_{t-1} + \eta_t & t = \lfloor \tau_{b1} T \rfloor, ..., \lfloor \tau_{b2} T \rfloor \\ 0 & \text{otherwise} \end{cases}$$

with $T = 200$, $\rho = 1 + c/T$, $\varepsilon_t \sim NIID(0, 1)$ and $\eta_t \sim NIID(0, \omega^2)$.

We compare the performance of the tests based on the $S_{\bar{c}}^{\dagger}$ statistic (denoted $LBI_{\bar{c}}$ in the graphs) for each of $\bar{c} \in \{2, 4, 6, 8, 10\}$ with the GSADF test of PSY (allowing for a mean with no lag augmentation) under $H_0 : \omega^2 = 0$ and $H_1 : \omega^2 > 0$ (ω^2 is on the horizontal axis).

Power $c = 1$. $\tau_{b1} = 0.3$, $\tau_{b2} = 0.7$

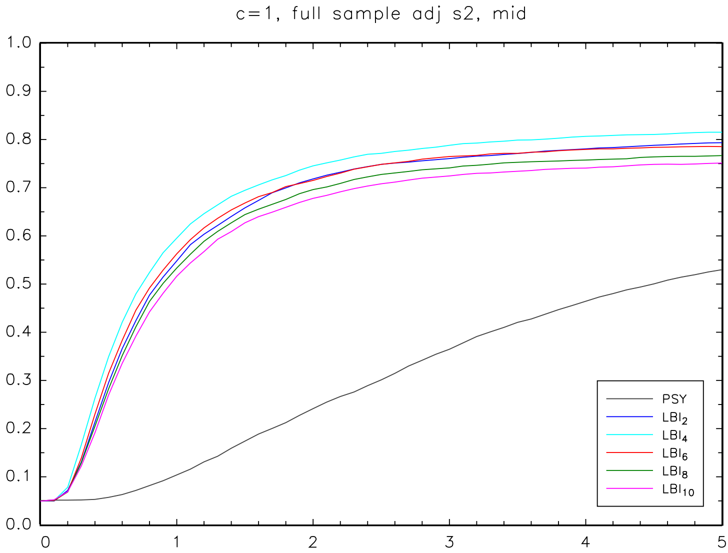

Power $c = 4$. $\tau_{b1} = 0.3$, $\tau_{b2} = 0.7$

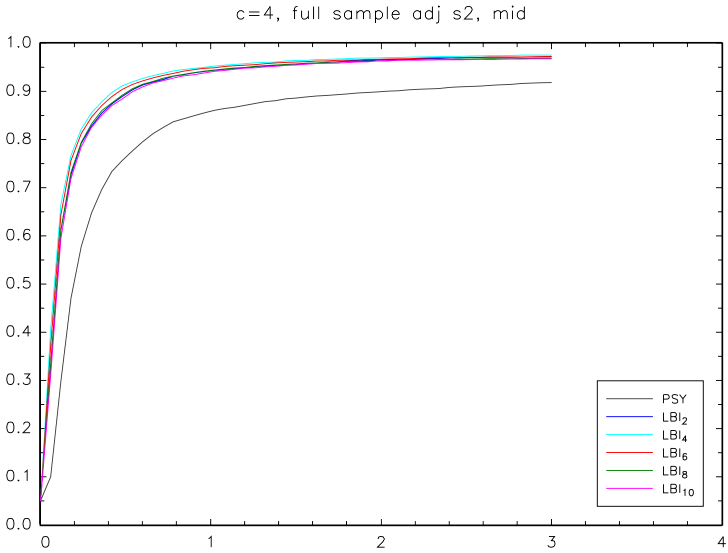

Power $c = 0.5$. $\tau_{b1} = 0.2$, $\tau_{b2} = 0.4$

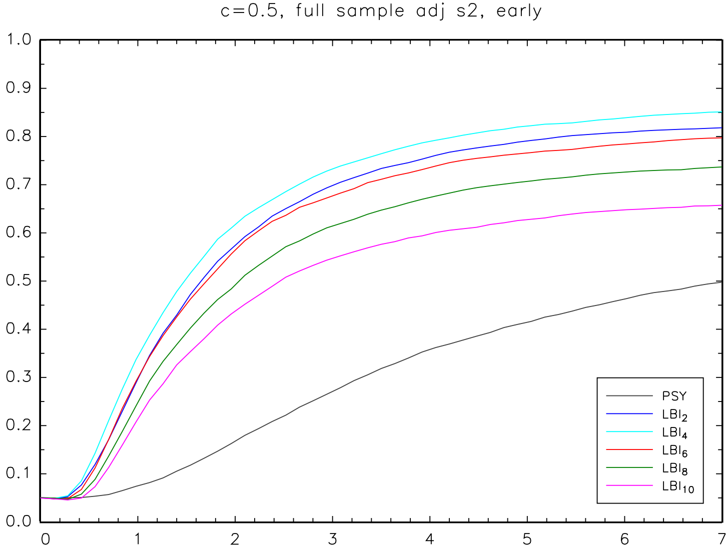

Power $c = 2$. $\tau_{b1} = 0.2$, $\tau_{b2} = 0.4$

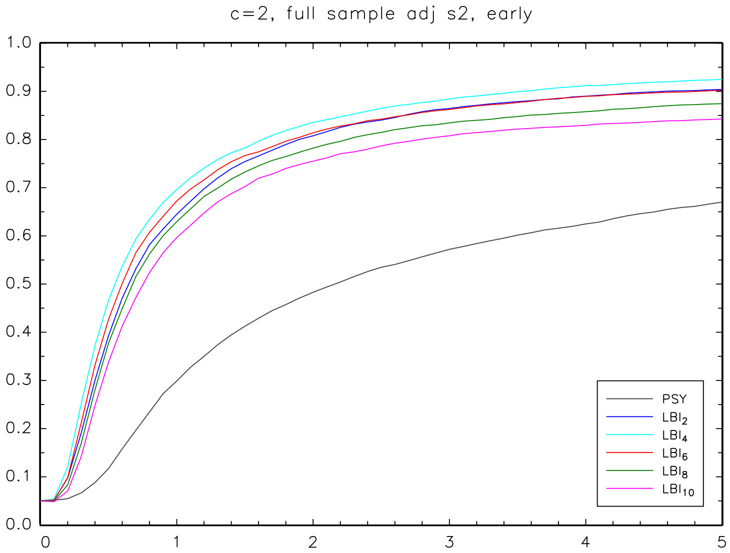

Monte Carlo Discussion

- We see that there is some difference in the power of the tests based on $S_{\bar{c}}^{\dagger}$ tests across $\bar{c}$, with the best overall power arguably offered by the test based on the $S_4^{\dagger}$ statistic.

- In all cases, however, the power of even the worst performing variant of $S_{\bar{c}}^{\dagger}$ is far ahead of that of the *GSADF* test.

- Further simulations show that the tests based on $S_{\bar{c}}^{\dagger}$ retain decent power, close to that of the *GSADF* test, for the $TVAR(1)$ model typically assumed for asset price bubble testing in the rest of the literature.

Table of Contents

Allowing for Heteroskedasticity

- We can introduce unconditional heteroskedasticity into ε_t and η_t. Specifically, replace Assumption 1 with the following

  Assumption 2. $\varepsilon_t = \sigma_t z_{1t}$ and $\eta_t = \omega_t z_{2t}$ where z_{1t} and z_{2t} are independent MDSs with unit variances and finite fourth order moments. The volatility terms σ_t and ω_t satisfy $\sigma_t = \sigma(t/T)$ and $\omega_t = \omega(t/T)$ where $\sigma(\cdot) \in \mathcal{D}$ and $\omega(\cdot) \in \mathcal{D}$ are non-stochastic and strictly positive.

Theorem 2. *Under H_1 and Assumption 2,*

$$S_{\bar{c}}^{*} \Rightarrow \sup_{\tau_1 \in [0,1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi,1]} \ln H_{c,\bar{c}}(\tau_1, \tau_2, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2})$$

where

$$H_{c,\bar{c}}(\tau_1, \tau_2, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) := \frac{\bar{c}^2(\tau_2 - \tau_1)^{-2} \int_{\tau_1}^{\tau_2} \left\{ \int_r^{\tau_2} e^{\bar{c}(s-r)(\tau_2-\tau_1)^{-1}} dK_c(s, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) \right\}^2 dr}{\int_0^1 \sigma(r)^2 dr + \int_{\tau_{b1}}^{\tau_{b2}} \omega(r)^2 dr + \left\{ \int_{\tau_{b1}}^{\tau_{b2}} e^{c(\tau_{b2}-s)(\tau_{b2}-\tau_{b1})^{-1}} \omega(s) dW_\eta(s) \right\}^2}$$

with

$$K_c(r, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) := \int_0^r \sigma(s) dW_\varepsilon(s) + \mathbb{I}(\tau_{b1} \leq r \leq \tau_{b2}) \int_{\tau_{b1}}^r e^{c(r-s)(\tau_{b2}-\tau_{b1})^{-1}} \omega(s) dW_\eta(s)$$

and $W_\varepsilon(r)$ and $W_\eta(r)$ are independent standard Brownian motions.

Limiting Distribution under Heteroskedasticity

- Note that the limiting null distribution in the heteroskedastic case is given by $H_{c,\bar{c}}(\tau_1, \tau_2, 0, \sigma(\cdot), \tau_{b1}, \tau_{b2})$ which we can simply denote as $H_{\bar{c}}(\tau_1, \tau_2, \sigma(\cdot))$ since it doesn't depend on c, τ_{b1} or τ_{b2}.
- Clearly then, heteroskedastcity in ε_t renders the limiting null distribution of $S_{\bar{c}}^*$ non-pivotal.
- In the homoskedastic null case this reduces to $L_{\bar{c}}(\tau_1, \tau_2)$, as previously.

Wild Bootstrap Implementation

- The limiting null distribution of the $S_{\bar{c}}^*$ statistic can therefore be seen to depend on the pattern of the heteroskedasticity in the data.
- We therefore propose obtaining critical values from a wild bootstrap algorithm. Specifically we generate the bootstrap data P_t^b, $t = 1, ..., T$, with $\Delta P_1^b = 0$ and $\Delta P_t^b = w_t \Delta P_t$, $t = 2, ..., T$, where w_t denotes an $NIID(0, 1)$ sequence.
- We then compute the $S_{\bar{c}}^*$ statistic on the bootstrap series; denote this $S_{\bar{c}}^{*b}$. Taking the $1 - \alpha$ quantile of the B such bootstrap statistics gives a critical value appropriate for α level testing.

Theorem 3. *Under H_1 and Assumption 2,*

$$S_{\bar{c}}^{*b} \xrightarrow{w}_p \sup_{\tau_1 \in [0,1-\pi]} \sup_{\tau_2 \in [\tau_1+\pi,1]} \ln H_{0,\bar{c}}(\tau_1, \tau_2, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2})$$

where

$$H_{0,\bar{c}}(\tau_1, \tau_2, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) = \frac{\bar{c}^2(\tau_2 - \tau_1)^{-2} \int_{\tau_1}^{\tau_2} \left\{ \int_r^{\tau_2} e^{\bar{c}(s-r)(\tau_2-\tau_1)^{-1}} dK_0(s, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) \right\}^2 dr}{\int_0^1 \sigma(r)^2 dr + \int_{\tau_{b1}}^{\tau_{b2}} \omega(r)^2 dr + \left\{ \int_{\tau_{b1}}^{\tau_{b2}} e^{c(\tau_{b2}-s)(\tau_{b2}-\tau_{b1})^{-1}} \omega(s) dW_\eta(s) \right\}^2}$$

with

$$K_0(r, \omega(\cdot), \sigma(\cdot), \tau_{b1}, \tau_{b2}) = \int_0^r \sigma(s) dW_\varepsilon(s) + \mathbb{I}(\tau_{b1} \leq r \leq \tau_{b2}) \int_{\tau_{b1}}^r \omega(s) dW_\eta(s)$$

and $W_\varepsilon(r)$ and $W_\eta(r)$ are independent standard Brownian motions.

Wild Bootstrap Implementation

- The heteroskedastic null case is given by $H_{0,\bar{c}}(\tau_1, \tau_2, 0, \sigma(\cdot), \tau_{b1}, \tau_{b2})$ which is $H_{\bar{c}}(\tau_1, \tau_2, \sigma(\cdot))$.
- The wild bootstrap statistic therefore correctly replicates the first order limiting null distribution of $S_{\bar{c}}^*$ under heteroskedasticity of the form specified in Assumption 2.
- The wild bootstrap statistic does not, however, converge to this same limit under H_1, and so the local power of the bootstrap tests will not coincide with that of a (infeasible) size-corrected version of the original test.
- The homoskedastic null case is given by $H_{\bar{c}}(\tau_1, \tau_2, 1)$ and this is equal to $L_{\bar{c}}(\tau_1, \tau_2)$, again replicating the limiting null distribution of $S_{\bar{c}}^*$ under Assumption 1.

Monte Carlo Setup - Heteroskedasticity

Data were generated according to

$$\begin{aligned} P_t &= F_t + B_t, \qquad t = 1, ..., T \\ F_t &= F_{t-1} + \varepsilon_t \end{aligned}$$

with $F_0 = 0$ and where

$$B_t = \begin{cases} \rho B_{t-1} + \eta_t & t = \lfloor \tau_{b1} T \rfloor, ..., \lfloor \tau_{b2} T \rfloor \\ 0 & \text{otherwise} \end{cases}$$

with $T = 200$ and $\rho = 1 + c/T$ with $\varepsilon_t \sim NIID(0, \sigma_t^2)$ and $\eta_t \sim NIID(0, \omega^2)$.

We set $\sigma_t = \sigma_1$ for $t = \lfloor 0.3T \rfloor, ..., \lfloor 0.7T \rfloor$ and $\sigma_t = 1$ otherwise.

Monte Carlo Setup - Heteroskedasticity

- In the simulations that follow, it is the statistic $S^{\dagger}_{\bar{c}}$ that is compared to the bootstrap distribution (obtained from B replications of $S^{\dagger b}_{\bar{c}}$).
- We replace $T^{-1}\sum_{t=2}^{T}(\Delta y_t^b)^2$ in the denominator of $S^{\dagger}_{\bar{c}}(\tau_1, \tau_2)^b$ with $T^{-1}\sum_{t=2}^{T}(\Delta y_t)^2$ as this improves finite sample size
- Under the alternative, we want $S^{\dagger b}_{\bar{c}}$ to be as small as we can make it - that puts more distance between it and the original statistic. So we (a): Don't replace $T^{-1}\sum_{t=2}^{T}(\Delta y_t)^2$ in the modified bootstrap statistic denominator with $T^{-1}\left\{\sum_{t=2}^{T}(\Delta y_t)^2 - \max_{t\in[2,\ldots,T]}|\Delta y_t|^2\right\}$; and (b): Calculate the bootstrap statistic numerator using the sequence $\left\{\Delta y_t - \max_{t\in[2,\ldots,T]}|\Delta y_t|\right\}$ in place of $\{\Delta y_t\}$. Both improve finite sample power.

Monte Carlo Results - Heteroskedasticity

ω	σ_1	$GSADF$	$S_4^{\dagger}$	$GSADF^b$	$S_4^{\dagger b}$
0	1	0.051	0.049	0.013	0.044
	3	0.255	0.374	0.040	0.094
	$1/3$	0.281	0.177	0.051	0.067
2	1	0.216	0.694	0.036	0.611
	3	0.343	0.676	0.054	0.473
	$1/3$	0.210	0.703	0.041	0.648

Table of Contents

Empirical Example

- We apply both the *GSADF* and $S_4^\dagger$ tests to a range of major stock indices.
- We use weekly data (in logs) from 01/01/1995 to 30/12/2001, with this period encompassing the well known dot-com bubble.
- The *GSADF* test is performed with regressions augmented by a single lag of ΔP_t and the $S_4^\dagger$ test is computed using a long-run variance estimator using the QS kernel with automatic bandwidth selection.
- Bootstrap p-values for GSADF are computed using the wild bootstrap algorithm of Harvey *et al.* (2016), and for $S_4^\dagger$ as outlined previously, in each case using $B = 499$ bootstrap replications.
- We see that, with the exception of the Nikkei index, the p-values for the $S_4^\dagger$ test are lower than for the *GSADF* test. If testing at a 10% level of significance the $S_4^\dagger$ test rejects for 9 series, and the *GSADF* test for only 1. These results reinforce the MC results which show a strong power advantage for $S_4^\dagger$ over *GSADF*.

Empirical Example

Index	$p(S_4^\dagger)$	$p(GSADF)$
FTSE 100	0.084	0.468
DAX	0.012	0.114
CAC 40	0.014	0.314
Nikkei	0.174	0.110
NYSE Composite	0.060	0.658
S&P 500	0.076	0.704
Dow Jones Industrial Average	0.018	0.810
Nasdaq 100	0.076	0.262
Nasdaq Composite	0.136	0.346
Nasdaq Computer	0.128	0.370
Nasdaq Biotechnology	0.026	0.024
Nasdaq Telecommunications	0.058	0.300

Table of Contents

Conclusions

- We propose tests for asset price bubbles based on the components model implied from financial theory.
- We show that our proposed tests have superior power to the *GSADF* test of PSY for a components based DGP when the innovations are i.i.d. and have reasonable size and power properties under unconditionally heteroskedastic innovations if a wild bootstrap implementation is used (but there is room for improvement and this is on-going!)
- An empirical application to major stock market indices for data spanning the dot-com bubble show that our proposed tests reject more often, and more strongly, than the *GSADF* test.