An Unobserved Components Based Test for Asset Price Bubbles

Robert Taylor

University of Essex

(joint work with Sam Astill, Dave Harvey and Steve Leybourne)

Special Invited Session: A Tribute to Hashem Pesaran I
CFE, December 15th 2025

Outline of Presentation

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

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, most recently after the Global Financial Crisis, their detection is now a crucial element of macroprudential policy.
- Most developments in the literature have focussed on using right-tailed univariate unit root statistics to test for explosivity, with the price-dividend ratio of individual stocks or stock indices being a common focus.
- The earliest meaningful 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] 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 statistics, or a test based on a double recursion across all possible start and end-dates (subject to a minimum window size).
- The PSY tests have subsequently become the industry standard for detecting and date-stamping bubble episodes.
- The aforementioned tests focus on treating the price series as a univariate time-varying autoregressive [TVAR] process and testing the null that the leading AR coefficient is fixed and 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 consistent with the general solution of the standard stock pricing equation commonly employed in the finance literature which decomposes the price of an asset into the sum of a fundamental price component and a separate (unrelated) bubble component which is explosive in (conditional) expectation.

- 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 components model may well be severely diminished.
- Although the rational bubble model is very often used to motivate the need for tests for asset price bubbles, to our knowledge there has been no attempt in either the econometrics or empirical finance literature to develop tests for asset price bubbles in context of the additive components structure implied by the finance theory model.
- To that end, we consider an unobserved components model motivated by the general solution to the asset pricing equation. In the context of this model we construct tests based on the *locally best invariant* [LBI] principle, for the null that the innovations to the bubble component 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).

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

- The standard *no-arbitrage* condition for the determination of the price, P_t, of an asset at time t, implies that it can be written as the conditional expectation of the future price, P_{t+1}, plus dividend (or other fundamentals, such as rental value with property assets), D_{t+1}, discounted by the constant risk-free rate, $R > 0$; that is,

$$P_t = \frac{E_t\left[P_{t+1} + D_{t+1}\right]}{1 + R} \tag{1}$$

where $E_t[.]$ denotes expectation conditional on information available at time t.

- Forward iteration of (1) leads to the general solution to (1),

$$P_t = \lim_{k\to\infty} \frac{E_t[P_{t+k}]}{(1+R)^k} + \sum_{i=1}^{\infty} \frac{1}{(1+R)^i} E_t\left[D_{t+i}\right]. \tag{2}$$

- The *fundamental* price, F_t, of the asset is defined as the (unique) particular solution for P_t in (2) where the *transversality condition*,

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

holds.

- Hence,

$$F_t = \sum_{i=1}^{\infty} \frac{1}{(1+R)^i} E_t \left[D_{t+i} \right]. \tag{4}$$

- The fundamental price of the asset is therefore seen to be the present value of all expected future dividends.

- Under (3), $P_t = F_t$, and the asset price is seen to be free of bubbles.

- Where the condition in (3) does not hold, F_t is not the only price process that solves equation (1). Here the general solution to equation (1) is given by

$$P_t = F_t + B_t \tag{5}$$

where $\{B_t\}_{t=1}^{\infty}$ satisfies the condition that

$$E_t[B_{t+1}] = (1 + R)B_t. \tag{6}$$

- B_t seen to be explosive in (conditional) expectation since, by the law of iterated expectations, (6) implies that $E_t[B_{t+\kappa}] = (1 + R)^{\kappa} B_t$, recalling that $R > 0$.
- Where B_t is constrained to be non-negative (which rules out negative bubbles) the condition in (6) entails that B_t is a *submartingale*.
- Where B_t is constrained to be positive it constitutes a *rational bubble*.

The Unobserved Components Model

- Homm and Breitung (2012) note that there are infinitely many solutions of the form given in (5).
- Equation (5) implies that the price of an asset can be decomposed into the sum of a fundamental component, F_t and an unrelated "bubble" component, B_t. If a bubble is present in the price then (6) 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.
- We next specify an econometric model based on the general solution to the stock pricing equation.

- To that end, for an asset price series, P_t, observed over the period $t = 1, ..., T$, we consider an *unobserved components* [UC] model of the form

$$P_t = \mu + F_t + B_t. \quad (7)$$

  Here μ is a constant term allowing for an arbitrary level in prices, while F_t and B_t represent, respectively, the fundamental and bubble components of P_t.

- To complete the specification of the model we need to specify stochastic processes corresponding to F_t and B_t in (7).

- Consider F_t first. From the general solution in (2), the behaviour of F_t is determined by the character of the dividend series. A standard assumption, dating back to Bachelier (1900), is that the fundamental price is a martingale process, such that $E_t(F_{t+1}) = F_t$, with some justifications for this assumption discussed in e.g. Breitung and Kruse (2013).

- Accordingly, we specify the component for F_t to be generated by the first-order unit root autoregression,

$$F_t = F_{t-1} + \varepsilon_t, \quad t = 1, ..., T \tag{8}$$

with $F_0 = o_p(T^{1/2})$ and where ε_t is a martingale difference sequence [MDS] with constant (unconditional) variance σ^2 which is bounded and bounded away from zero.

- Turning to B_t, this needs to satisfy the condition in (6), that $E_t[B_{t+1}] = (1+R)B_t$, $R > 0$. This allows for a wide range of possible stochastic processes, including the deterministic bubble model of Blanchard and Watson (1982), the randomly starting bubble model, and the periodically collapsing bubble model of Evans (1991).

- A first-order autoregression whose lag coefficient is strictly greater than one also satisfies (9) (though it is not a submartingale), with empirical studies finding that any explosive behaviour found in prices tends to be episodic. Accordingly, we adopt a simple model of episodic explosive behaviour for B_t, whereby B_t is zero other than between the two unknown dates $1 \leq t_{b1} < t_{b2} \leq T$, where it follows an explosive first-order autoregression; that is,

$$B_t = \begin{cases} \rho B_{t-1} + \eta_t, & t = t_{b1}, \ldots, t_{b2} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where $\rho > 1$ and η_t is a MDS, independent of ε_t, with bounded variance, ω^2.

The Unobserved Components Model

- Clearly, under the specification for B_t in (9), if $\omega^2 > 0$, there is a single explosive episode occurring between two deterministic time points, t_{b1} (the bubble inception date) and t_{b2} (the crash date). In contrast, if $\omega^2 = 0$, $B_t = 0$ for all t and the bubble component is absent from P_t.
- In the context of the UC model (7)-(8)-(9), we will use the LBI principle to derive a test for the null hypothesis that $\omega^2 = 0$ against $\omega^2 > 0$ in the case where the break dates t_{b1} and t_{b2} and the autoregressive parameter, ρ, are all assumed known.
- Subsequently, we will develop a feasible version of this test which does not require knowledge of these parameters.
- Although, by design, the power function of these tests is directed against the specific alternative of a single explosive episode, the tests will also have power against many other more complicated alternatives, including, for example, models with multiple bubble episodes.

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

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 in the context of the UC model in (7)-(8)-(9). The null hypothesis of no bubble is given by

$$H_0 : \omega^2/\sigma^2 = 0.$$

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

$$H_1 : \omega^2/\sigma^2 > 0.$$

- Given the specification in (8), P_t therefore follows a random walk under H_0. Notice, therefore, that our null model coincides with the null model considered by PWY and PSY.
- 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

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

where $\boldsymbol{P} := (P_1, ..., P_T)'$, $\boldsymbol{F} := (F_1, ..., F_T)'$, $\boldsymbol{B} := (B_1, ..., B_T)'$, and $\mathbf{1}$ is a $(T \times 1)$ vector of 1s.

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

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

The matrix $\boldsymbol{A}$ has the block structure

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

where

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

A Locally Best Invariant Test

- We can use equation (6) of King and Hillier (1985,p.99) to obtain the critical region of the LBI test for H_0 against H_1 if $\boldsymbol{Q}(0) = \sigma^2 \boldsymbol{I}_T$, but as we have seen that is not the case, due to the unit root present in F_t.
- We need to apply a transformation to the data, such that this holds on the resulting variance matrix. To that end, consider the data transformation

$$\boldsymbol{M}(\boldsymbol{P} - \mu \mathbf{1}) \sim N(0, \boldsymbol{M}\boldsymbol{Q}(\omega^2)\boldsymbol{M}')$$

where $\boldsymbol{M}'\boldsymbol{M} = \boldsymbol{D}^{-1}$, i.e.,

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

- Aside from end effects, this transformation essentially corresponds to first differencing the data.

- We then have that $\boldsymbol{M}(\boldsymbol{P} - \mu\mathbf{1}) = [P_1 - \mu, \Delta P_2, ..., \Delta P_T]'$ whose covariance matrix is given by
$$\boldsymbol{M}\boldsymbol{Q}(\omega^2)\boldsymbol{M}' = \sigma^2 \boldsymbol{I}_T + \omega^2 \boldsymbol{M}\boldsymbol{A}\boldsymbol{M}'$$
and, hence, it follows that $\boldsymbol{M}\boldsymbol{Q}(0)\boldsymbol{M}' = \sigma^2 \boldsymbol{I}_T$.

- Applying King and Hillier (1985), we obtain that a LBI test of H_0 against H_1 rejects for large positive values of the statistic
$$S_\rho(\tau_{b1}, \tau_{b2}) := \frac{\boldsymbol{r}'\boldsymbol{M}\boldsymbol{A}\boldsymbol{M}'\boldsymbol{r}}{T^{-1}\boldsymbol{r}'\boldsymbol{r}} \qquad (10)$$
where $\tau_{b1} := t_{b1}/T$ and $\tau_{b2} := t_{b2}/T$, $\boldsymbol{r} := [0, \Delta P_2, ..., \Delta P_T]'$ is the vector of residuals from the OLS (null) regression of $[P_1, \Delta P_2, ..., \Delta P_T]'$ on $[1, 0, ..., 0]'$.

- $S_\rho(\tau_{b1}, \tau_{b2})$ can be equivalently written in scalar notation as

$$S_\rho(\tau_{b1}, \tau_{b2}) = \frac{\sum_{t=t_{b1}+\mathbb{I}(t_{b1}=1)}^{t_{b2}} (\Delta P_t + (\rho - 1) \sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j)^2 + \mathbb{I}(t_{b1} = 1)((\rho - 1) \sum_{j=t_{b1}+1}^{t_{b2}} \rho^{j-t_{b1}-1} \Delta P_j)}{\hat{\sigma}^2}$$

  where $\hat{\sigma}^2 := T^{-1} \sum_{t=2}^{T} (\Delta P_t)^2$ is a scale estimate, '$\mathbb{I}(t_{b1} = 1)$' is the indicator function taking the value 1 when $t_{b1} = 1$, zero otherwise.

- ΔP_t in the first term in the numerator is of lower stochastic order of magnitude than $(\rho - 1) \sum_{j=t+1}^{t_{b2}} \rho^{j-t-1} \Delta P_j$. We therefore use a simplified statistic which omits this term. After re-arranging the remaining terms, this yields the statistic,

$$S_\rho^*(\tau_{b1}, \tau_{b2}) := \frac{(\rho - 1)^2 \sum_{t=t_{b1}+1}^{t_{b2}} \left(\sum_{j=t}^{t_{b2}} \rho^{j-t} \Delta P_j \right)^2}{\hat{\sigma}^2}. \tag{11}$$

- Unreported simulation evidence suggests that there is essentially no loss in finite sample power from basing our proposed feasible tests which follow on this simplified version of the LBI statistic.

- 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 (11) computed over this arbitrary sub-sample is then given by:

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

A Feasible Test

- To be operational we must specify a value for ρ; we denote this by $\bar{\rho} > 1$, yielding

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

- 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.
- Hence, substituting for $\bar{\rho}$,

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

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

1 Introduction

2 The Unobserved Components Model

3 A Locally Best Invariant Test

4 Asymptotics

5 A Modified Statistic

6 Finite Sample Simulations

7 Allowing for Weak Dependence

8 Allowing for Unconditional Heteroskedasticity

9 Empirical Application

10 Conclusions

Limiting Distribution

- We next consider the large sample behaviour of $S^*_{\bar{c}}$ under $H_1 : \omega^2/\sigma^2 > 0$, where the true autoregressive coefficient 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.

- Assumption 1 does not impose Gaussianity on either ε_t or η_t. Moreover, it allows for a wide class of models of conditional heteroskedasticty in both η_t and ε_t, including ARCH and GARCH models.

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

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

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 $\bar{c}$.

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

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 $\hat{\sigma}^2 = T^{-1}\sum_{t=2}^{T}(\Delta P_t)^2$ with

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

- It is straightforward to show that

$$\hat{\sigma}^2_m \xrightarrow{p} \sigma^2 + \omega^2(\tau_{b2} - \tau_{b1})$$

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

$$S_{\bar{c}}^{\dagger}(\tau_1, \tau_2) := \frac{\bar{c}^2 (t_2 - t_1)^{-2} \sum_{t=t_1+1}^{t_2} \left(\sum_{j=t}^{t_2} \{1 + \bar{c}(t_2 - t_1)^{-1}\}^{j-t} \Delta P_j \right)^2}{\hat{\sigma}_m^2}$$

It is easily shown 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

1 Introduction
2 The Unobserved Components Model
3 A Locally Best Invariant Test
4 Asymptotics
5 A Modified Statistic
6 Finite Sample Simulations
7 Allowing for Weak Dependence
8 Allowing for Unconditional Heteroskedasticity
9 Empirical Application
10 Conclusions

Single Bubble Episode

- We compare the performance of the tests based on the $S_{\bar{c}}^{\dagger}$ statistic 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 is on the horizontal axis in all plots) for a variety of DGPs.

- The 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 $F_0 = 0$.

- In the first set of experiments a single bubble episode occurs between $t = \lfloor \tau_{b1} T \rfloor$ and $t = \lfloor \tau_{b2} T \rfloor$:

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

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

Power $c = 1$. $\tau_{b1} = 0.3$, $\tau_{b2} = 0.7$ [mid-sample bubble]

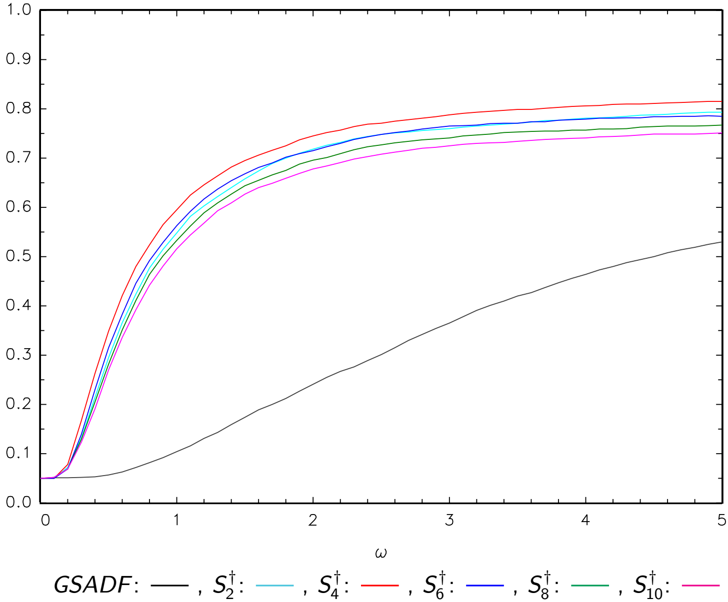

Power $c = 2$. $\tau_{b1} = 0.3$, $\tau_{b2} = 0.7$ [mid-sample bubble]

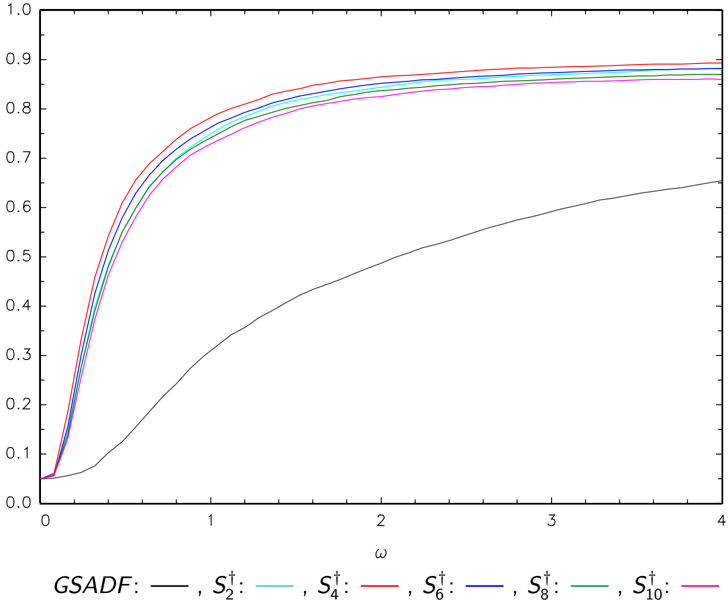

Power $c = 4$. $\tau_{b1} = 0.3$, $\tau_{b2} = 0.7$ [mid-sample bubble]

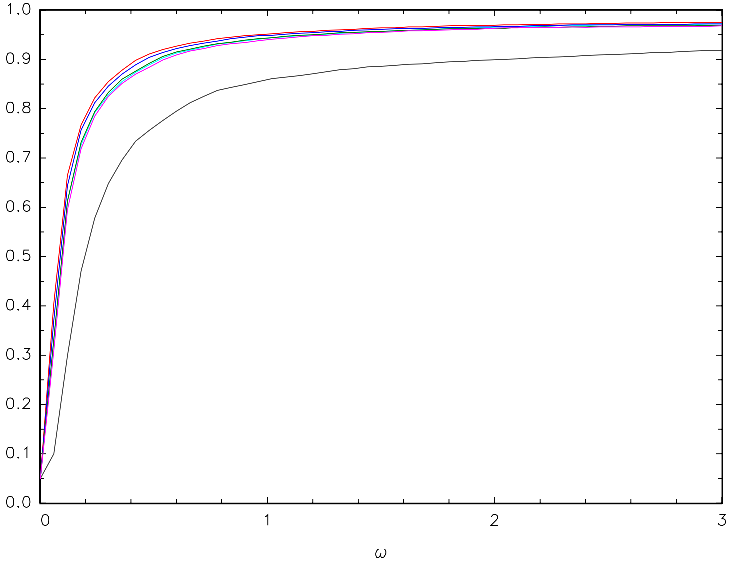

$GSADF$: —— , $S_2^\dagger$: —— , $S_4^\dagger$: —— , $S_6^\dagger$: —— , $S_8^\dagger$: —— , $S_{10}^\dagger$: ——

Power $c = 0.5$. $\tau_{b1} = 0.2$, $\tau_{b2} = 0.4$ [early bubble]

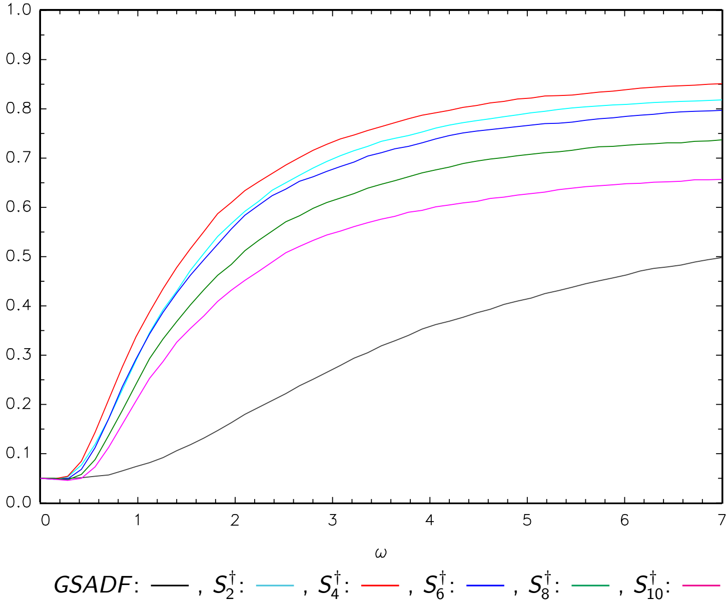

Power $c = 1$. $\tau_{b1} = 0.2$, $\tau_{b2} = 0.4$ [early bubble]

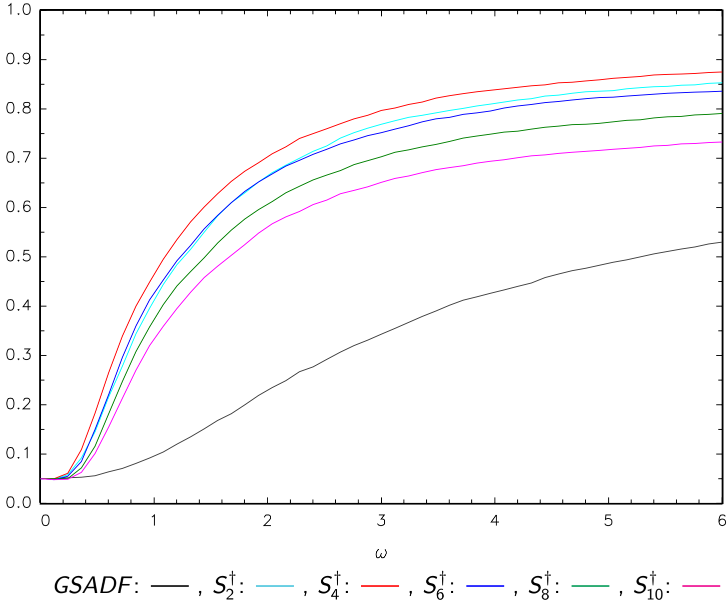

Power $c = 2$. $\tau_{b1} = 0.2$, $\tau_{b2} = 0.4$ [early bubble]

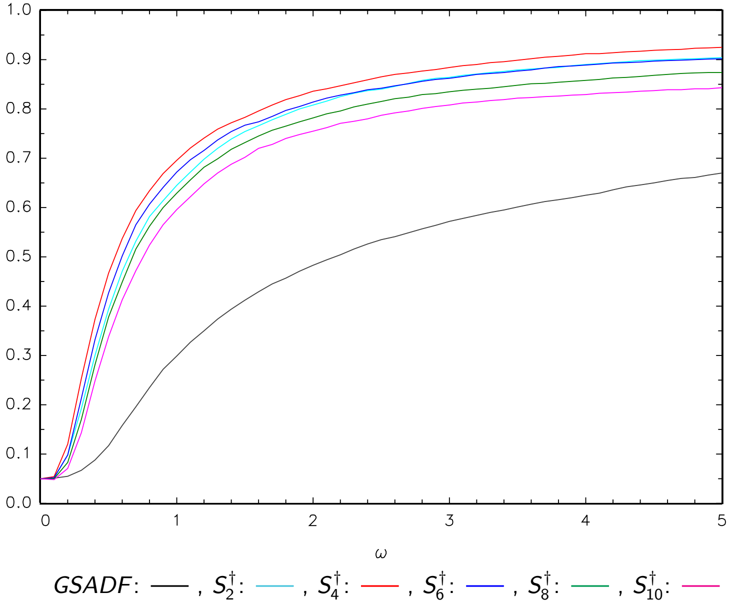

Power $c = 0.5$. $\tau_{b1} = 0.8$, $\tau_{b2} = 1$ [late bubble]

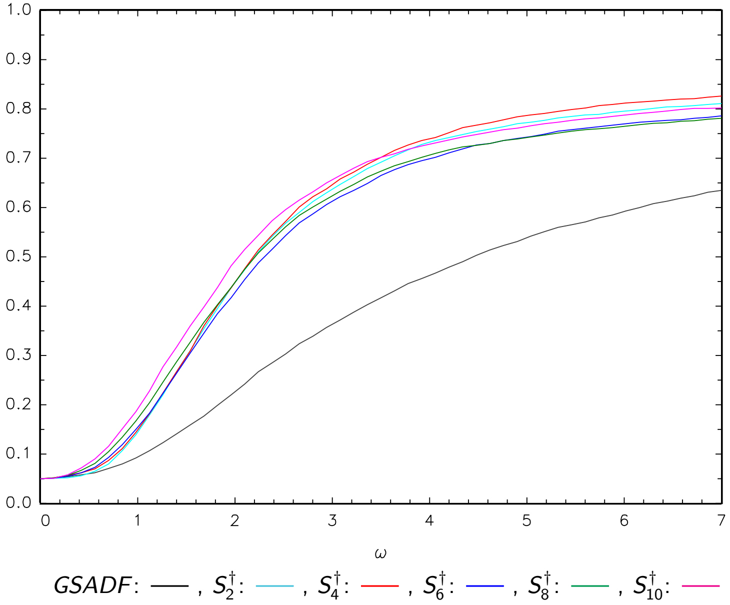

Power $c = 1$. $\tau_{b1} = 0.8$, $\tau_{b2} = 1$ [late bubble]

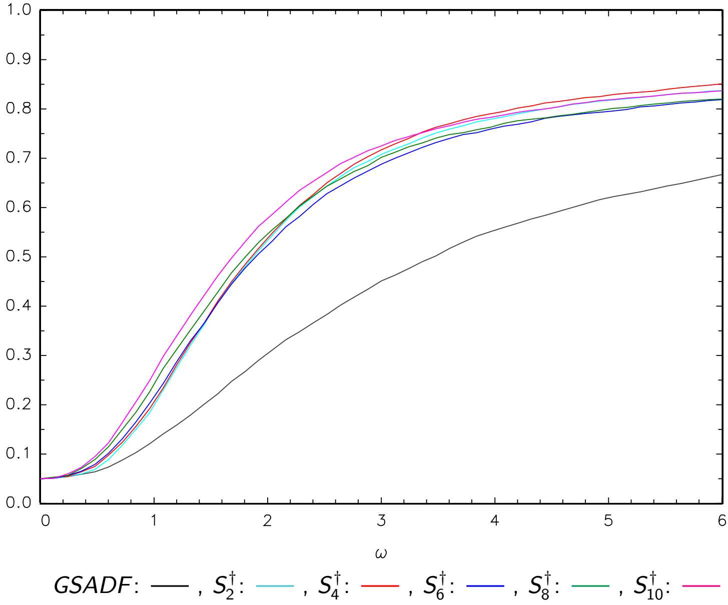

Power $c = 2$. $\tau_{b1} = 0.8$, $\tau_{b2} = 1$ [late bubble]

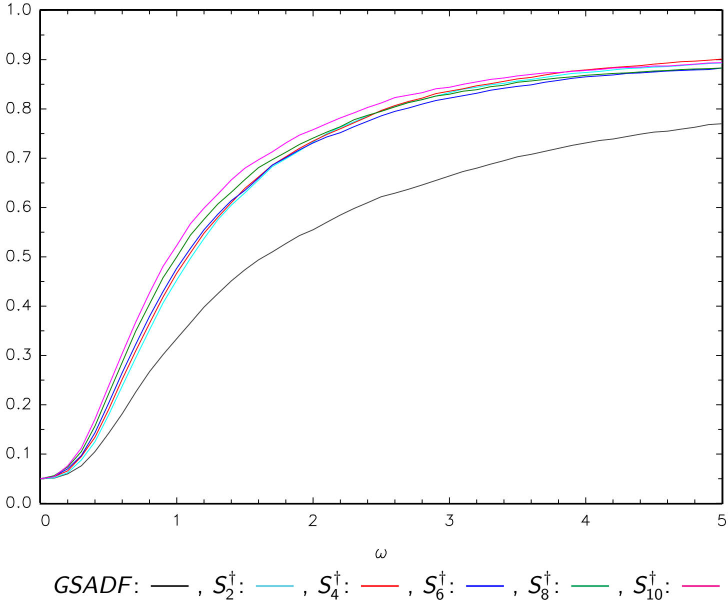

- The results for a single bubble episode highlight some differences in the power profiles of the tests based on $S_{\bar{c}}^{\dagger}$ tests across $\bar{c}$, with the overall best power profile offered by the test based on the $S_{4}^{\dagger}$ statistic.

- 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.

- We next consider results for a DGP with two distinct bubble episodes. Here the DGP is as before but now

$$B_t = \begin{cases} \rho B_{t-1} + \eta_t & t = t_{b11}, \ldots, t_{b21} \text{ and } t = t_{b12}, \ldots, t_{b22} \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

where $t_{bij} = \lfloor \tau_{i,j} T \rfloor \; i, j = 1, 2$ with $0 < \tau_{1,1} < \tau_{2,1} < \tau_{1,2} < \tau_{2,2} < 1$

Power $c = 0.5$. $\tau_{1,1} = 0.2$, $\tau_{2,1} = 0.4$ and $\tau_{1,2} = 0.6$, $\tau_{2,2} = 0.8$ [2 bubbles]

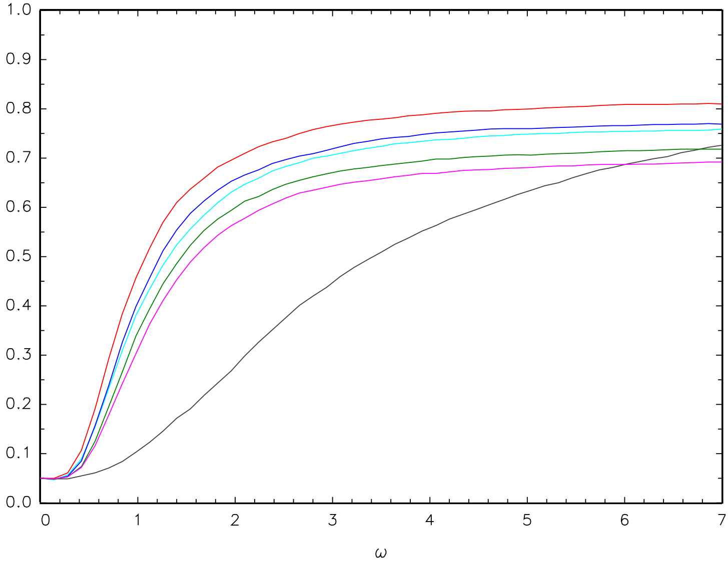

$GSADF$: —— , $S_2^\dagger$: —— , $S_4^\dagger$: —— , $S_6^\dagger$: —— , $S_8^\dagger$: —— , $S_{10}^\dagger$: ——

Power $c = 1$. $\tau_{1,1} = 0.2$, $\tau_{2,1} = 0.4$ and $\tau_{1,2} = 0.6$, $\tau_{2,2} = 0.8$ [2 bubbles]

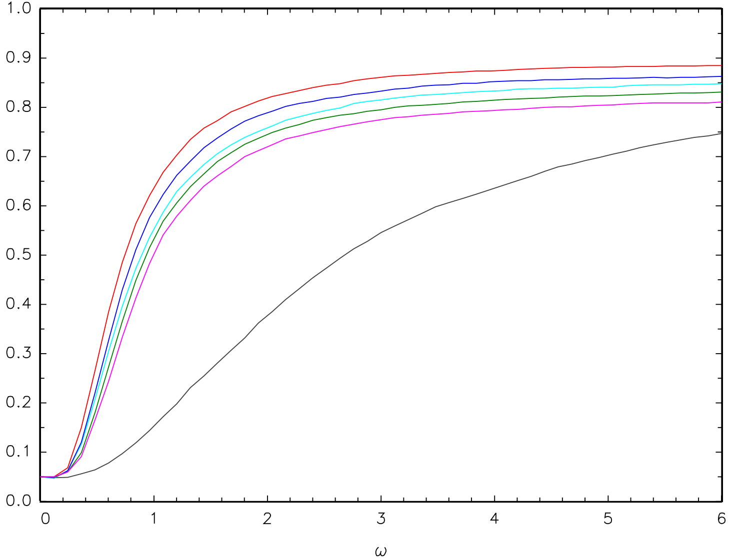

$GSADF$: —— , $S_2^\dagger$: —— , $S_4^\dagger$: —— , $S_6^\dagger$: —— , $S_8^\dagger$: —— , $S_{10}^\dagger$: ——

Power $c = 2$. $\tau_{1,1} = 0.2$, $\tau_{2,1} = 0.4$ and $\tau_{1,2} = 0.6$, $\tau_{2,2} = 0.8$ [2 bubbles]

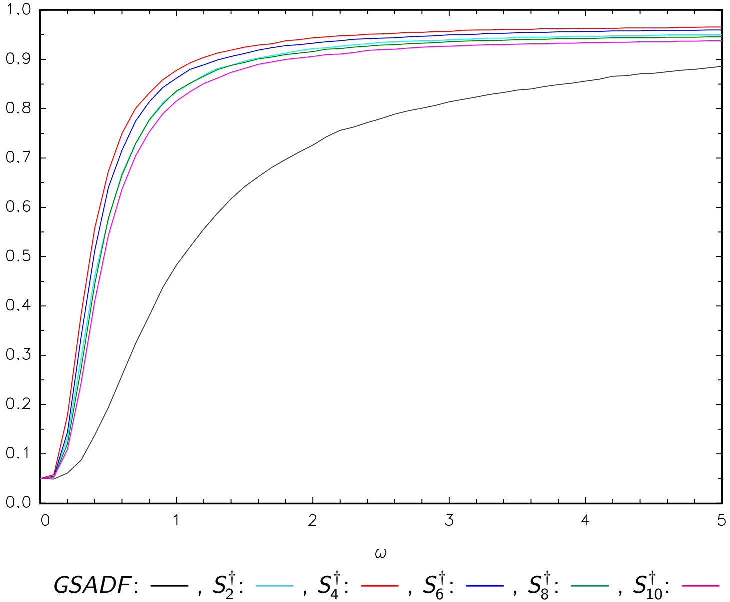

Monte Carlo Discussion - Two Bubble Episodes

- As in the single bubble case, among the $S_{\bar{c}}^{\dagger}$ tests the best power profile is again displayed by the $S_4^{\dagger}$ test.
- The $S_{\bar{c}}^{\dagger}$ tests continue for the most part to significantly outperform the *GSADF* test.
- We next consider the case where the price data are generated by a $TVAR(1)$ model, as is typically assumed for asset price bubble testing in the rest of the literature.

$TVAR(1)$ DGP

- We consider three TVAR(1) DGPs, varying in how the bubble phase ends.

- In the first DGP P_t is a unit root process until time $t = \lfloor \tau_1 T \rfloor$ when a bubble occurs running until $t = \lfloor \tau_2 T \rfloor$, after which point P_t reverts to a unit root process until the end of the sample:

$$P_t = \begin{cases} P_{t-1} + \varepsilon_t & t = 1, \ldots, \lfloor \tau_1 T \rfloor \\ \phi_1 P_{t-1} + \varepsilon_t & t = \lfloor \tau_1 T \rfloor + 1, \ldots, \lfloor \tau_2 T \rfloor \\ P_{t-1} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1, \ldots, T. \end{cases}$$

- In the second DGP P_t is a unit root process followed by a bubble regime, and then subsequently follows a stationary collapse regime from time $\lfloor \tau_2 T \rfloor + 1, ..., \lfloor \tau_3 T \rfloor$, before the series reverts to a unit root process from $t = \lfloor \tau_3 T \rfloor + 1$ until the end of the sample:

$$P_t = \begin{cases} P_{t-1} + \varepsilon_t & t = 1, \ldots, \lfloor \tau_1 T \rfloor \\ \phi_1 P_{t-1} + \varepsilon_t & t = \lfloor \tau_1 T \rfloor + 1, \ldots, \lfloor \tau_2 T \rfloor \\ \phi_2 P_{t-1} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1, \ldots, \lfloor \tau_3 T \rfloor \\ P_{t-1} + \varepsilon_t & t = \lfloor \tau_3 T \rfloor + 1, \ldots, T. \end{cases}$$

- In the third DGP P_t begins as a unit root process followed by a bubble regime, with the price series then collapsing to its pre-bubble level, before reverting to unit root behaviour thereafter, i.e.

$$P_t = \begin{cases} P_{t-1} + \varepsilon_t & t = 1, \ldots, \lfloor \tau_1 T \rfloor \\ \phi_1 P_{t-1} + \varepsilon_t & t = \lfloor \tau_1 T \rfloor + 1, \ldots, \lfloor \tau_2 T \rfloor \\ P_{\lfloor \tau_1 T \rfloor} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1 \\ P_{t-1} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 2, \ldots, T. \end{cases}$$

- We set $\tau_1 = 0.3$, $\tau_2 = 0.5$ and $\tau_3 = 0.7$. We generate the innovations as $\varepsilon_t \sim NIID(0, 1)$ and set $\phi_1 = 1 + c^*/T$, and for the stationary collapse model set $\phi_2 = 1 - c^*/T$. We examine the power of the tests for a grid of values of $c^* \in (0, 30]$. Again, we set $T = 200$.

- Given that the best overall performance for the $S_{\bar{c}}^{\dagger}$ tests in the previous results was given setting $\bar{c} = 4$, we report results for the $S_4^{\dagger}$ and $GSADF$ tests.

First DGP. $\tau_1 = 0.3$, $\tau_2 = 0.5$ [$TVAR(1)$ - no crash]

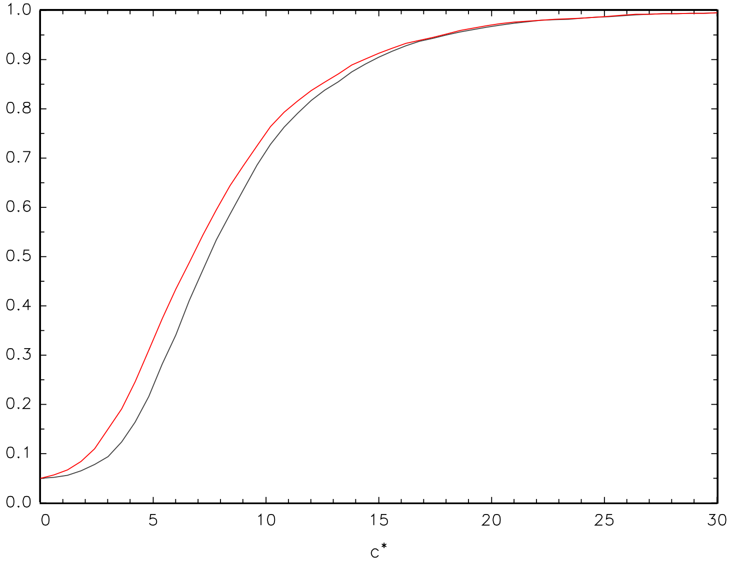

$GSADF$: —— , $S_4^{\dagger}$: ——

Second DGP. $\tau_1 = 0.3$, $\tau_2 = 0.5$, $\tau_3 = 0.7$ [$TVAR(1)$ - stationary collapse]

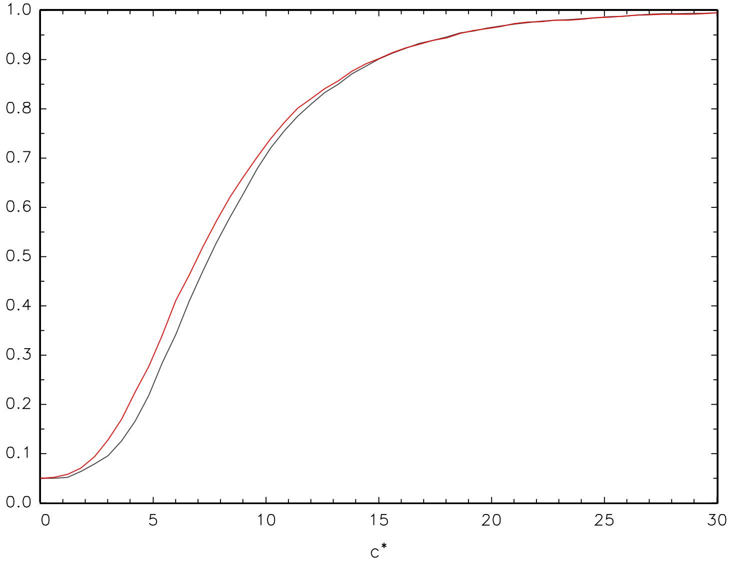

$GSADF$: —— , $S_4^{\dagger}$: ——

Third DGP. $\tau_1 = 0.3$, $\tau_2 = 0.5$ [$TVAR(1)$ - instantaneous crash]

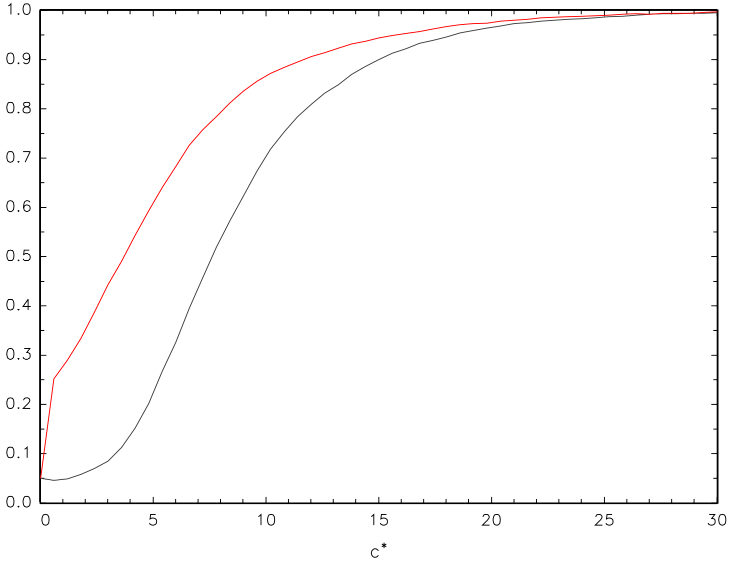

$GSADF$: —— , $S_4^{\dagger}$: ——

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

- Thus far we have assumed that F_t in (8) follows a random walk, so that ΔP_t is a MDS under H_0. This assumption is standard, but in practice weak autocorrelation could be present in ΔP_t under the null. This could arise from e.g. the presence of weak dependence in the shocks, ε_t, in (8), or from observed transaction prices being subject to an additive weakly autocorrelated [$I(0)$] measurement error (microstructure noise) arising from imperfections in the trading process; see, e.g., Aït-Sahalia and Yu (2009).

- In both examples above, applying first differences to (5) we obtain that

$$\Delta P_t = \Delta B_t + u_t, \; t = 2, ..., T \tag{13}$$

  where u_t is a weakly dependent [$I(0)$] series.

- Assume that u_t satisfies the mixing conditions given in, for example, Assumption $\mathcal{E}$ of Cavaliere and Taylor (2005,p.1114). As discussed in Cavaliere and Taylor (2005,p.1115), this allows u_t to belong to a wide class of weakly dependent (and conditionally heteroskedastic) stationary processes, including most stationary and invertible *ARMA* processes.

- In this case the limiting null distribution of $S^*_{\bar{c}}$ will depend on both the long run variance, say λ^2, and the short-run variance of u_t when u_t is weakly dependent.
- To recover the limiting null distribution of $S^*_{\bar{c}}$, given previously, in cases where u_t is weakly dependent, we can use the same solution as outlined in the context of the LBI test of Nyblom and Mäkeläinen (1983) by Kwiatkowski *et al.* (1992) [KPSS]. This entails replacing the short run variance estimator, $\hat{\sigma}^2$, with an estimate of the long run variance which has the property of being consistent under H_0.
- Following KPSS, we will use the familiar kernel-based estimate

$$\hat{\lambda}^2 := \sum_{j=-T+1}^{T-1} k\left(\frac{j}{q_T}\right)\hat{\gamma}(j), \quad \hat{\gamma}(j) := T^{-1}\sum_{t=|j|+2}^{T} \Delta P_t \Delta P_{t-|j|} \tag{14}$$

  where q_T and $k(\cdot)$ are the bandwidth and kernel function, respectively, assumed to satisfy standard regularity conditions, such as Assumption $\mathcal{K}$ of Cavaliere and Taylor (2005,p.1115).
- As demonstrated in Cavaliere and Taylor (2005), these conditions ensure that $\hat{\lambda}^2 \xrightarrow{p} \lambda^2$ under H_0.

Allowing for Weak Dependence

- Consequently, a modified statistic which replaces $\hat{\sigma}^2$ by $\hat{\lambda}^2$ in the calculation of $S^*_{\bar{c}}(\tau_1, \tau_2)$, will have the same limiting null distribution as that given previously for $S^*_{\bar{c}}$ in the case where u_t is serially uncorrelated.
- Correcting for weak dependence in the errors in the context of the modification outlined previously to allow for outliers arising from bubble terminations follows in an obvious fashion, and we denote the long-run variance estimate in this case as $\hat{\lambda}^2_m$.
- We investigated the finite sample efficacy of this correction for the $S_4^\dagger$ test using the simulation DGP $P_t = F_t$, $t = 1, ..., T$, with $F_t = F_{t-1} + \varepsilon_t$ with $F_0 = 0$ and $\varepsilon_t = \varphi\varepsilon_{t-1} + e_t$, where $e_t \sim NIID(0, 1)$. We consider $T \in \{100, 200, 400\}$ and $\varphi \in \{0, 0.3, 0.5, 0.7, 0.9\}$.
- The long run variance estimate was obtained using the quadratic spectral kernel with the automatic bandwidth selection method of Andrews (1991).

Table 1: Empirical Size of Nominal 5% Tests.

Panel A: Short-Run Variance Estimate			
	$T = 100$	$T = 200$	$T = 400$
φ	$S_4^{\dagger}$	$S_4^{\dagger}$	$S_4^{\dagger}$
0.0	0.050	0.050	0.050
0.3	0.313	0.403	0.499
0.5	0.636	0.788	0.894
0.7	0.914	0.982	0.998
0.9	0.998	1.000	1.000
Panel B: Long-Run Variance Estimate			
	$T = 100$	$T = 200$	$T = 400$
φ	$S_4^{\dagger}$	$S_4^{\dagger}$	$S_4^{\dagger}$
0.0	0.048	0.047	0.047
0.3	0.074	0.056	0.051
0.5	0.073	0.050	0.044
0.7	0.074	0.047	0.030
0.9	0.123	0.058	0.035

Allowing for Weak Dependence

- The results in Table 1 show that, as would be expected, the original $S_4^{\dagger}$ tests display severe size distortions where $\varphi \neq 0$.
- However, the modification to allow for weak dependence performs remarkably well (especially considering how poorly it is known to work in the context of the KPSS test) in restoring empirical size close to the nominal level.
- The only exception of note is for $\varphi = 0.9$ when $T = 100$ where the tests are still rather oversized (around 12%). However, $\varphi = 0.9$ implies the changes in prices are close to being $I(1)$, arguably well beyond the level of dependence that might be anticipated in practice.

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

- Thus far we have assumed that ε_t and η_t are unconditionally homoskedastic MDSs.
- Unconditional heteroskedasticity in η_t will affect the asymptotic local power of our proposed tests but not their limiting null distributions, while unconditional heteroskedasticity in ε_t will affect both.
- We now explore the impact of unconditional heteroskedasticity in ε_t and η_t on our proposed tests. To that end, we 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.

- 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.

- 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

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

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 homoskedastic null case is given by $H_{\bar{c}}(\tau_1, \tau_2, 1)$ and this is equal to $L_{\bar{c}}(\tau_1, \tau_2)$.
- We replace $T^{-1}\sum_{t=2}^{T}(\Delta y_t^b)^2$ in the denominator of $S_{\bar{c}}^{*}(\tau_1, \tau_2)^b$ with $T^{-1}\sum_{t=2}^{T}(\Delta y_t)^2$ as this improves finite sample size
- Under the alternative, we want $S_{\bar{c}}^{*b}$ 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 bootstrap statistic denominator with the modification $T^{-1}\left\{\sum_{t=2}^{T}(\Delta y_t)^2 - \max_{t\in[2,\ldots,T]}|\Delta y_t|^2\right\}$; 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.
- It is the statistic $S_{\bar{c}}^{\dagger}$ that is compared to be bootstrap distribution.

Monte Carlo Setup - Unconditional 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

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

where $t_{b1} = \lfloor \tau_{b1} T \rfloor$ and $t_{b2} = \lfloor \tau_{b2} T \rfloor$.

- We set $\tau_{b1} = 0.3$ and $\tau_{b2} = 0.7$, $T = 200$, $\rho = 1 + c(t_{b2} - t_{b1})^{-1}$ with $c = 2$, $\varepsilon_t \sim NIID(0, \sigma_t^2)$, and $\eta_t \sim NIID(0, \omega^2)$.
- We set $\sigma_t = \sigma$ for $t = \lfloor 0.3T \rfloor, ..., \lfloor 0.7T \rfloor$ and $\sigma_t = 1$ otherwise. We consider $\sigma \in \{1/3, 1/2, 1, 2, 3\}$.
- $B = 500$ bootstrap replications were used.

Table 2: Empirical Size and Power of Nominal 5% Tests with Non-Stationary Volatility

	$\omega = 0$		$\omega = 1$	
σ	$S_4^{\dagger}$	Bootstrap $S_4^{\dagger}$	$S_4^{\dagger}$	Bootstrap $S_4^{\dagger}$
1	0.049	0.044	0.762	0.757
1/2	0.124	0.064	0.782	0.792
1/3	0.177	0.067	0.787	0.801
2	0.252	0.082	0.747	0.668
3	0.374	0.094	0.729	0.582

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

- We apply both the *GSADF* and S_4 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 so-called *dot-com bubble*.
- The *GSADF* test is performed with regressions which include an intercept and one lag of ΔP_t. The $S_4^{\dagger}$ test is computed using a long-run variance estimator again using the QS kernel with automatic bandwidth selection. For all tests we set $\pi = 0.1$ yielding a minimum window width of 36.
- The $\mathcal{H}_{AD}$ stationary volatility test of Cavaliere and Taylor (2008) is also calculated, using their recommended settings.
- Bootstrap critical values are computed using either a parametric bootstrap or the wild bootstrap algorithm of Harvey *et al.* (2016), in each case using $B = 1000$ bootstrap replications.
- We see that, with the exception of the Nikkei and Nasdaq Biotechnology indices, the heteroskedasticity-robust 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 apparently reinforce the MC results which show a strong power advantage for $S_4^{\dagger}$ over *GSADF*.

Table 3: Empirical Application: Stationary Volatility and Bubble Detection Test Results

		Bootstrap p-values			
		Parametric		Wild	
Series	$\mathcal{H}_{AD}$	$S_4^\dagger$	*GSADF*	$S_4^\dagger$	*GSADF*
FTSE 100	3.517**	0.019	0.206	0.084	0.468
DAX	3.950***	0.001	0.002	0.012	0.114
CAC 40	1.953*	0.003	0.110	0.014	0.314
Nikkei	0.187	0.200	0.043	0.174	0.110
NYSE Composite	3.626**	0.023	0.453	0.060	0.658
S&P 500	4.697***	0.025	0.417	0.076	0.704
Dow Jones Industrial Average	3.754**	0.004	0.674	0.018	0.810
Nasdaq 100	8.926***	0.004	0.055	0.076	0.262
Nasdaq Composite	8.799***	0.023	0.061	0.136	0.346
Nasdaq Computer	7.652***	0.027	0.128	0.128	0.370
Nasdaq Biotechnology	9.300***	0.000	0.000	0.026	0.024
Nasdaq Telecommunications	10.558***	0.003	0.048	0.058	0.300

*Note: In the context of the stationary volatility tests, *, ** and *** denote rejection at the 10%, 5% and 1% nominal significance levels, respectively, based on the critical values from Table 5 of Shorack and Wellner (1987, p.148).*

Table of Contents

1. Introduction
2. The Unobserved Components Model
3. A Locally Best Invariant Test
4. Asymptotics
5. A Modified Statistic
6. Finite Sample Simulations
7. Allowing for Weak Dependence
8. Allowing for Unconditional Heteroskedasticity
9. Empirical Application
10. Conclusions

Conclusions

- We propose tests for asset price bubbles derived from an unobserved components time series model motivated by the general solution of the standard stock pricing equation.

- We present Monte Carlo simulations which 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 retain decent size and power properties under unconditionally heteroskedastic innovations provided a wild bootstrap implementation is used.

- An empirical application to major stock market indices for data spanning the purported dot-com bubble shows that our proposed tests reject more often, and more strongly, than the *GSADF* test, consistent with the Monte Carlo findings in the paper.