



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/194958/>

Version: Published Version

Article:

Whitehouse, E.J., Harvey, D.I. and Leybourne, S.J. (2023) Real-time monitoring of bubbles and crashes. *Oxford Bulletin of Economics and Statistics*, 85 (3). pp. 482-513. ISSN: 0305-9049

<https://doi.org/10.1111/obes.12540>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Real-Time Monitoring of Bubbles and Crashes

EMILY J. WHITEHOUSE,[†]  DAVID I. HARVEY[‡]  and STEPHEN J. LEYBOURNE[‡]

[†]*Department of Economics, University of Sheffield, Sheffield, S1 4DT, UK (e-mail: e.whitehouse@sheffield.ac.uk)*

[‡]*School of Economics, University of Nottingham, Nottingham, NG7 2RD, UK*

Abstract

Given the financial and economic damage that can be caused by the collapse of an asset price bubble, it is of critical importance to rapidly detect the onset of a crash once a bubble has been identified. We develop a real-time monitoring procedure for detecting a crash episode in a time series. We adopt an autoregressive framework, with bubble and crash regimes modelled by explosive and stationary dynamics, respectively. The first stage of our approach is to monitor for a bubble; conditional on which, we monitor for a crash in real time as new data emerges. Our crash detection procedure employs a statistic based on the different signs of the means of the first differences associated with explosive and stationary regimes, and critical values are obtained using a training period of data. We show that the procedure has desirable asymptotic properties in terms of its ability to rapidly detect a crash while never indicating a crash earlier than one occurs. Monte Carlo simulations further demonstrate that our procedure can offer a well-controlled false positive rate during a bubble regime. Application to the US housing market demonstrates the efficacy of our procedure in rapidly detecting the house price crash of 2006.

I. Introduction

Asset price bubbles and crashes are a prevalent feature in economic and financial markets, with notable examples including the Dot-com bubble in technology stock prices in the late 1990s, the sub-prime mortgage bubble in the US housing market in the mid-2000s, and, more recently, the presence of bubbles in cryptocurrencies. The collapse of the sub-prime mortgage bubble, in particular, illustrated how devastating the emergence and collapse of asset price bubbles can be, not just for the asset market in which the bubble occurs, but for the global economy as a whole. The quicker that policy makers are alerted to the onset of a crash, the quicker they can respond to mitigate the effects of that crash. Developing early detection tools that can provide warning signals of such an event are therefore of critical importance. In this paper we propose a real-time monitoring procedure for the crash of an asset price bubble in order to provide this fast detection.

JEL Classification numbers: C12, C22, G01.

Much of the econometric literature concerning the identification of bubbles and crashes has focused on historical detection, where a bubble episode has emerged and then terminated within a sample of observed data. The focus of this literature has been primarily on rational bubbles, where the price of an asset diverges from the underlying fundamental value of that asset, yet investors continue to purchase the asset due to an expectation that prices will grow beyond the price paid. In a rational bubble framework, an asset price bubble is present if explosive behaviour is observed in the price series of an asset but not in the corresponding fundamental values. Modelling asset price bubbles as explosive autoregressive processes, Phillips, Wu, and Yu (2011) propose the use of sub-sample right-tailed augmented Dickey–Fuller unit root tests, implemented recursively, to distinguish between a process which is unit root across a full sample period or instead exhibits a single episode of temporary explosiveness at some point in that period. Phillips, Shi, and Yu (2015) subsequently extend this recursive unit root testing approach to consider the detection of multiple bubble episodes. Further developments in the econometric detection of explosive bubbles have considered, *inter alia*, CUSUM-based procedures (Homm and Breitung, 2012), bootstrap implementations of recursive unit root test procedures (Harvey *et al.*, 2016; Phillips and Shi, 2020) and the use of Generalised Least Squares-based recursive unit root testing (Whitehouse, 2019). These techniques have been employed in the empirical literature to detect past bubble episodes in a wide range of asset markets such as stocks (Caspi and Graham, 2018; Hu and Oxley, 2018; Basse *et al.*, 2021), housing (Anundsen *et al.*, 2016; Pavlidis, Paya, and Peel, 2018; Anundsen, 2019), commodities (Etienne, Irwin, and Garcia, 2014; Etienne, Irwin, and Garcia, 2015; Figuerola-Ferretti and McCrorie, 2016), and cryptocurrencies (Corbet, Lucey, and Yarovaya, 2018; Gronwald, 2021).

Determining the presence and timing of historical asset bubbles provides useful information for empirical researchers, allowing for a more rigorous analysis of the timeline and determinants of bubble behaviour. However, from a policy perspective, the detection of ongoing bubbles in a real-time monitoring exercise is of clear interest. Astill *et al.* (2017) examine the possibility of detecting an end-of-sample bubble. That is, a bubble which has emerged close to the end of an observed sample period of data and is still ongoing at the end of the sample. Their procedure requires the application of a test statistic to a finite number of end-of-sample observations, where critical values are obtained through sub-sampling of the test statistic throughout the remainder of the sample. The test statistic used is motivated by a Taylor series expansion of the first differences of an explosive process. The advantage of this methodology is that as the critical values are obtained from the data itself, the test procedure is robust to serial correlation and conditional heteroskedasticity in the data.

The Astill *et al.* (2017) approach provides a method for conducting a one-shot test for an end-of-sample bubble. If no bubble is detected through such a procedure, it may then be of interest to repeatedly test for the presence of a bubble as new data observations are released. We refer to this repeated testing every time a new data point is observed as real-time monitoring. Inherent in real-time monitoring is the multiple testing problem, in which repeated application of the same test statistic as each new observation is realised can lead to an empirical size for the procedure far beyond the theoretical size of a one-off application of the test statistic. Astill *et al.* (2018) (AHLST) therefore propose a real-time

monitoring procedure, based on the test statistic of Astill *et al.* (2017), but implemented in such a way that the false positive rate (FPR) for bubble detection (i.e. the probability of false detection at each point of monitoring) can be determined at any point in the monitoring horizon. Specifically, test statistics are computed over rolling sub-samples of fixed length within a training period, with the maximum test statistic within this training period forming the critical value to which monitoring statistics, computed using the most recent data available, are compared. Such an approach also allows the practitioner to set a maximum monitoring horizon to ensure that the FPR never exceeds some specified level. Monte Carlo simulation results demonstrate that the AHLST 2017 real-time monitoring procedure delivers FPRs close to their theoretical level in finite samples, and offers rapid detection of explosive bubbles.

An equally important issue to real-time detection of the emergence of a bubble is real-time monitoring for the subsequent termination of that bubble in the form of a crash. In the context of detecting historical crash episodes, Harvey, Leybourne, and Sollis (2017) and Phillips and Shi (2018) model the crash regime as a stationary autoregressive process that immediately follows the explosive autoregressive bubble phase. In this paper, we propose a real-time monitoring procedure for stationary crashes, with our procedure conditional on having first detected the presence of a bubble. Our procedure relies on the sequential application of a new test statistic for the detection of a crash, motivated by the differing signs of the means of the first differences of explosive and stationary processes. The test statistic is constructed in such a way that a user-chosen parameter offers practitioners the choice to potentially introduce a degree of detection delay in order to reduce the FPR of crash detection. We follow AHLST 2017 in computing sub-sample test statistics over a training period, although it is now the minimum of these training sample statistics which forms our critical value to which monitoring statistics for crash detection are compared. We rely on AHLST 2017 for bubble detection, such that we begin monitoring for a crash conditional on first detecting an explosive bubble. We show that our crash monitoring procedure has desirable asymptotic properties in terms of its ability to rapidly detect a crash regime while never indicating a crash earlier than one occurs. Monte Carlo simulations demonstrate that our recommended crash monitoring procedure can offer a well-controlled FPR in finite samples, while also allowing rapid detection of a crash.

Once a crash has been detected by our procedure, it may be the case that, rather than ending the monitoring exercise, a practitioner wishes to continue monitoring for the possible emergence of subsequent bubble episodes. We therefore also propose an extension of our monitoring procedure to the multiple bubble case, where, after detection of a crash regime, bubble monitoring then resumes. Simulation results show that our procedure is effective also in this multiple regime context, delivering good FPR control and rapid detection of multiple explosive bubble and stationary collapse regimes.

The usefulness of our proposed crash monitoring procedure is demonstrated through an empirical application of the procedure to the US house price to rent ratio. In a pseudo real-time monitoring exercise, we begin our monitoring in 1998:Q1, with detection of an explosive bubble occurring in 2000:Q1 and subsequent detection of a stationary crash occurring in 2006:Q2, thus providing timely detection of changes in the dynamics of the US housing market.

In the next section we present a bubble and crash model and introduce the hypotheses of interest. Section III describes the AHLST 2017 real-time monitoring procedure for bubbles. In section IV we outline our crash monitoring procedure and establish asymptotic results for its behaviour under the alternative hypothesis of interest. Monte Carlo simulation results are presented in section V. Section VI discusses the extension of our proposed procedure to multiple bubble and crash episodes. In section VII we provide an application of our real-time monitoring procedure to the US housing market. Section VIII concludes. Proofs are contained in Appendices. We adopt the following notation: ‘ $\lfloor \cdot \rfloor$ ’ denotes the integer part, ‘ $\mathbb{I}(\cdot)$ ’ denotes the indicator function, and we use the order notation $O_p^+(\cdot)$ to imply that the term concerned is positive.

II. The model and real-time monitoring framework

We consider the following DGP for a time series y_t , $t = 1, \dots, T$, which represents prices (or prices relative to fundamentals):

$$y_t = \mu + u_t, \quad (1)$$

$$u_t = \begin{cases} u_{t-1} + \varepsilon_t & t = 2, \dots, \lfloor \tau_1 T \rfloor \\ (1 + \delta_1)u_{t-1} + \varepsilon_t & t = \lfloor \tau_1 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor \\ (1 - \delta_2)u_{t-1} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_3 T \rfloor \\ u_{t-1} + \varepsilon_t & t = \lfloor \tau_3 T \rfloor + 1, \dots, T, \end{cases} \quad (2)$$

with $u_1 = O_p(1)$, $\delta_1 > 0$ and $1 \geq \delta_2 > 0$. We assume the error term ε_t is a strictly stationary, possibly conditionally heteroskedastic, process with zero mean.

In the context of (1) and (2), if $\tau_1 = 1$ then y_t admits a unit autoregressive root throughout the sample period. If $\tau_1 < 1$, then the y_t process changes at time $\lfloor \tau_1 T \rfloor$ from unit root to explosive autoregressive dynamics (with explosive offset δ_1), providing a model of bubble behaviour. In this case we assume that $u_{\lfloor \tau_1 T \rfloor} > 0$, such that $u_{\lfloor \tau_1 T \rfloor} = O_p^+(T^{1/2})$, ensuring that the explosive regime has an upwards (rather than downwards) trajectory, in line with typical bubble behaviour. If $\tau_2 = 1$ the explosive regime is ongoing at the end of the sample, while if $\tau_2 < 1$ the explosive behaviour terminates at time $\lfloor \tau_2 T \rfloor$. After the explosive regime terminates, the process switches into a stationary collapse regime (with stationary offset δ_2), which acts as a model for a post-bubble crash. The stationary collapse regime runs to time $\lfloor \tau_3 T \rfloor$, at which point, provided $\tau_3 < 1$, unit root behaviour resumes.¹

Our focus is on real-time monitoring first for an explosive regime, and then, conditional on detecting such explosive behaviour, monitoring for a stationary collapse. In the first stage, we therefore wish to distinguish between the hypotheses:

$$\begin{aligned} H_0^{(1)} : \tau_1 &= 1 && \text{(unit root)} \\ H_1^{(1)} : \tau_1 &< \tau_2 \leq 1 && \text{(unit root then explosive, with or without stationary collapse)} \end{aligned}$$

¹An additional possibility at the end of the explosive regime is for the process to return directly to unit root behaviour without collapse, i.e. $\tau_2 = \tau_3 < 1$. However, given that bubbles almost invariably terminate in collapse, we focus our main attention on the $\tau_2 < \tau_3$ case, such that a stationary collapse regime follows the termination of explosive behaviour.

Conditional on having rejected $H_0^{(1)}$ in favour of $H_1^{(1)}$, in the second stage we then assume $H_1^{(1)}$ is true and proceed to distinguish between the following sub-hypotheses of $H_1^{(1)}$:

$$\begin{aligned} H_0^{(2)} &: \tau_1 < \tau_2 = 1 \quad (\text{unit root then explosive without stationary collapse}) \\ H_1^{(2)} &: \tau_1 < \tau_2 < 1 \quad (\text{unit root then explosive with stationary collapse}) \end{aligned}$$

The real-time monitoring framework we adopt follows AHLST and considers y_1, \dots, y_{T^*} , $T^* = \lfloor \lambda T \rfloor \leq \lfloor \tau_1 T \rfloor$ for some $\lambda \in (0, 1)$, as a training period, during which it is assumed that no explosive behaviour is present, i.e. $T^* \leq \lfloor \tau_1 T \rfloor$. We will subsequently consider monitoring from some time period T^\dagger onwards, employing the training period data in a calibration role. The first stage involves monitoring for a change from $H_0^{(1)}$ to $H_1^{(1)}$, so as to detect the onset of an explosive regime. Once an explosive regime has been detected, the second stage involves subsequent monitoring for a change from $H_0^{(2)}$ to $H_1^{(2)}$, in order to detect the termination of explosive behaviour and the start of a stationary collapse regime.

III. Monitoring for an explosive regime

Several approaches to testing between $H_0^{(1)}$ and $H_1^{(1)}$, i.e. testing for a period of temporary explosiveness, have been developed in the literature. While the majority of these approaches focus on historical detection, some recent developments have emerged to deal with real-time monitoring for explosive episodes. In particular, AHLST develop a real-time monitoring procedure based on the Astill *et al.* (2017) test for an end-of-sample explosive regime, which in turn is based on the instability tests of Andrews (2003) and Andrews and Kim (2006). In Astill *et al.* (2017), a test statistic designed to detect explosive behaviour is computed over a finite-sized window of observations at the end of the sample period, and compared to a critical value obtained from computing the same statistic repeatedly over sub-samples of the previous observations. This sub-sampling approach has the desirable feature that the test is robust to conditional heteroskedasticity and serial correlation, while robustness to unconditional heteroskedasticity is delivered through a White-type correction in the test statistic. AHLST adapt this approach to the real-time monitoring context by comparing test statistics computed from rolling finite-sized windows in the monitoring period with the maximum of the statistics computed over sub-samples in the training period.

The test statistic that Astill *et al.* (2017) employ is motivated by a Taylor series expansion of the first differences Δy_t during the explosive regime, and essentially amounts to testing for the presence of an upward trend in Δy_t . Letting k denote the chosen window width over which the statistic is computed, and e the last observation used in the statistic's calculation, the statistic is given by

$$A_{e,k} = \frac{\sum_{t=e-k+1}^e (t - e + k) \Delta y_t}{\sqrt{\sum_{t=e-k+1}^e \{(t - e + k) \Delta y_t\}^2}}$$

The real-time monitoring procedure of AHLST then proceeds as follows. Suppose that we wish to start monitoring for a bubble at the present time period, $t = T^\dagger$. We let

$t = 1, \dots, T^*$ form an initial training sample, where $T^* = T^\dagger - k$. The $A_{e,k}$ statistic is computed over rolling sub-samples of length k within this training sample, producing a set of training sample statistics. The maximum training sample statistic, which we denote by $A_{\max}^* = \max_{e \in [k+1, T^*]} A_{e,k}$, forms the critical value for the monitoring procedure. Beginning at time $t = T^\dagger$, the first monitoring statistic is computed using data from time periods $t = T^\dagger - k + 1, \dots, T^\dagger = T^* + 1, \dots, T^* + k$, then subsequent monitoring statistics are computed as each new observation occurs, rolling forwards the window of k observations (e.g. the second monitoring statistic is computed at time $t = T^\dagger + 1$ using data from $t = T^\dagger - k + 2, \dots, T^\dagger + 1$). Detection of an explosive regime is triggered at the first point where a monitoring statistic $A_{e,k}$, $e = T^\dagger, T^\dagger + 1, \dots$, exceeds the critical value A_{\max}^* . At an arbitrary point in the monitoring period, $t = e$, we can then write the monitoring decision rule as:

$$\text{Detect } H_1^{(1)} \text{ at time } e \text{ if } A_{e,k} > A_{\max}^*. \quad (3)$$

The time period at which an explosive regime is detected is then denoted $t = T^\diamond$. AHLST discuss how the FPR of such a procedure can be controlled, and we formalise this in Theorem 1 below. Hereafter, we refer to this explosive regime monitoring procedure as $A_{\text{MAX}}(k)$.

We now establish the theoretical FPR of the $A_{\text{MAX}}(k)$ procedure under $H_0^{(1)}$ as $T \rightarrow \infty$, where we assume that monitoring has been run to some point T' , and that T^* and T' are such that $T^* = \lfloor \lambda_1 T \rfloor$ and $T' = \lfloor \lambda_2 T \rfloor$, where $0 < \lambda_1 < \lambda_2 \leq 1$. This is done by observing that the decision rule in (3) is equivalent to determining whether the maximum of the monitoring statistics $A_{e,k}$, $e = T^\dagger, T^\dagger + 1, \dots, T'$, exceeds the corresponding maximum statistic over the training period A_{\max}^* . Evaluating the limiting probabilities of these exceedances under $H_0^{(1)}$ gives the result of the following theorem.

Theorem 1. Under $H_0^{(1)}$ and assuming that $\{\varepsilon_t\}$ satisfies the mixing conditions of Ferreira and Scotto (2002, p. 476), then as $T \rightarrow \infty$,

$$\lim_{T \rightarrow \infty} P \left(\max_{e \in [T^*+k, T']} A_{e,k} > \max_{e \in [k+1, T^*]} A_{e,k} \right) = \alpha,$$

where

$$\alpha = \lim_{T \rightarrow \infty} \left(\frac{T' - T^* - k + 1}{T' - 2k + 1} \right) = \lim_{T \rightarrow \infty} \left(\frac{T' - T^*}{T'} \right).$$

For given values of T^* and k , we can use the result in Theorem 1 to approximate the empirical FPR that would be obtained in practice for any monitoring point T' , i.e.

$$\alpha \approx \frac{T' - T^* - k + 1}{T' - 2k + 1}. \quad (4)$$

We can also rearrange (4) to identify the monitoring time period T' at which the FPR of the procedure will (approximately) reach the level α , allowing us to determine how far one can monitor into the future whilst maintaining a chosen FPR:

$$T' \approx \frac{T^* + k - 1 - \alpha(2k - 1)}{1 - \alpha}.$$

IV. Monitoring for a stationary collapse regime

Our main focus in this paper is on the second stage monitoring where, given prior detection of an explosive regime, the aim is to detect the termination of explosive behaviour and the onset of a stationary collapse regime. We now motivate a test statistic for distinguishing between $H_0^{(2)}$ and $H_1^{(2)}$, with the aim of using this statistic to monitor for a collapse using a similar approach to the $A_{MAX}(k)$ procedure for monitoring explosive behaviour. Consider first the model (1) and (2) expressed in first differences:

$$\Delta y_t = \begin{cases} \varepsilon_t & t = 2, \dots, \lfloor \tau_1 T \rfloor \\ \delta_1 u_{t-1} + \varepsilon_t & t = \lfloor \tau_1 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor \\ -\delta_2 u_{t-1} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_3 T \rfloor \\ \varepsilon_t & t = \lfloor \tau_3 T \rfloor + 1, \dots, T \end{cases}$$

Next consider the observations in the immediate neighbourhood of the explosive regime endpoint $\lfloor \tau_2 T \rfloor$. Specifically, for a finite number m of observations on Δy_t up to $\lfloor \tau_2 T \rfloor$, and for a finite number n of observations on Δy_t immediately after $\lfloor \tau_2 T \rfloor$, we can use the autoregressive recursion to write

$$\Delta y_t = \begin{cases} \delta_1(1 + \delta_1)^{t - (\lfloor \tau_2 T \rfloor - m) - 1} u_{\lfloor \tau_2 T \rfloor - m} + \sum_{i=0}^{t - (\lfloor \tau_2 T \rfloor - m) - 1} \delta_1(1 + \delta_1)^{i-1} \varepsilon_{t-i} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor - m + 1, \dots, \lfloor \tau_2 T \rfloor \\ -\delta_2(1 - \delta_2)^{t - \lfloor \tau_2 T \rfloor - 1} u_{\lfloor \tau_2 T \rfloor} - \sum_{i=1}^{t - \lfloor \tau_2 T \rfloor - 1} \delta_2(1 - \delta_2)^{i-1} \varepsilon_{t-i} + \varepsilon_t & t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor + n \end{cases} \quad (5)$$

For $t = \lfloor \tau_2 T \rfloor - m + 1, \dots, \lfloor \tau_2 T \rfloor$, for finite m , $\delta_1(1 + \delta_1)^{t - (\lfloor \tau_2 T \rfloor - m) - 1} u_{\lfloor \tau_2 T \rfloor - m}$ is of the same order as $u_{\lfloor \tau_2 T \rfloor - m}$, which is also of the same order as $u_{\lfloor \tau_2 T \rfloor}$; then, from lemma 1(i) of Harvey *et al.* (2017), it follows that $u_{\lfloor \tau_2 T \rfloor} = O_p(S_T^{1/2})$, where $S_T = \lfloor \tau_1 T \rfloor (1 + \delta_1)^{2(\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor)}$, and hence

$$\delta_1(1 + \delta_1)^{t - (\lfloor \tau_2 T \rfloor - m) - 1} u_{\lfloor \tau_2 T \rfloor - m} = O_p(S_T^{1/2}). \quad (6)$$

Next, $\sum_{i=0}^{t - (\lfloor \tau_2 T \rfloor - m) - 1} \delta_1(1 + \delta_1)^{i-1} \varepsilon_{t-i} + \varepsilon_t = O_p(1)$ since we have a finite sum of $O_p(1)$ variates, and so the term in (6) dominates. Similarly, for $t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor + n$ and finite n , lemma 1(i) of Harvey *et al.* (2017) gives

$$\delta_2(1 - \delta_2)^{t - \lfloor \tau_2 T \rfloor - 1} u_{\lfloor \tau_2 T \rfloor} = O_p(S_T^{1/2}), \quad (7)$$

and $\sum_{i=1}^{t - \lfloor \tau_2 T \rfloor - 1} \delta_2(1 - \delta_2)^{i-1} \varepsilon_{t-i} + \varepsilon_t = O_p(1)$, so that (7) dominates over this second time period. In each regime, therefore, the first terms in (5) dominate the stochastic behaviour of Δy_t . Finally, for δ_1 and δ_2 close to zero, taking the first term of a Taylor expansion with respect to δ_1 and δ_2 , respectively, yields the following approximations:

$$\begin{aligned} \delta_1(1 + \delta_1)^{t - (\lfloor \tau_2 T \rfloor - m) - 1} u_{\lfloor \tau_2 T \rfloor - m} &\approx \delta_1 u_{\lfloor \tau_2 T \rfloor - m} & t = \lfloor \tau_2 T \rfloor - m + 1, \dots, \lfloor \tau_2 T \rfloor \\ \delta_2(1 - \delta_2)^{t - \lfloor \tau_2 T \rfloor - 1} u_{\lfloor \tau_2 T \rfloor} &\approx \delta_2 u_{\lfloor \tau_2 T \rfloor} & t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor + n \end{aligned} \quad (8)$$

This allows Δy_t in the neighbourhood of $\lfloor \tau_2 T \rfloor$ to be expressed as

$$\Delta y_t = \begin{cases} \beta_1 + \eta_t & t = \lfloor \tau_2 T \rfloor - m + 1, \dots, \lfloor \tau_2 T \rfloor \\ \beta_2 + \eta_t & t = \lfloor \tau_2 T \rfloor + 1, \dots, \lfloor \tau_2 T \rfloor + n \end{cases}, \quad (9)$$

where $\beta_1 = \delta_1 u_{\lfloor \tau_2 T \rfloor - m}$, $\beta_2 = -\delta_2 u_{\lfloor \tau_2 T \rfloor}$ and η_t generically denotes an error containing the approximation errors in (8) and the lower order $O_p(1)$ terms in (5).

Given the presence of an explosive regime ($\delta_1 > 0$), the onset of a stationary collapse ($\delta_2 > 0$) implies a change in the mean of Δy_t from a positive value β_1 to a negative value β_2 at time $\lfloor \tau_2 T \rfloor + 1$. If we consider simple OLS estimators of β_1 and β_2 over the two respective sub-samples given in (9), i.e.:

$$\hat{\beta}_1 = m^{-1} \sum_{t=\lfloor \tau_2 T \rfloor - m + 1}^{\lfloor \tau_2 T \rfloor} \Delta y_t,$$

$$\hat{\beta}_2 = n^{-1} \sum_{t=\lfloor \tau_2 T \rfloor + 1}^{\lfloor \tau_2 T \rfloor + n} \Delta y_t,$$

then we can motivate a statistic for detecting a change from explosive to stationary collapse behaviour at time $\lfloor \tau_2 T \rfloor + 1$ as one based on the sign of the product $\hat{\beta}_1 \hat{\beta}_2$.

In the monitoring context, where the putative collapse change point is unknown, we can consider the following statistic indexed by the last observation used in the statistic's calculation, again denoted e :

$$N_{e,m,n} = \left(m^{-1} \sum_{t=e-n-m+1}^{e-n} \Delta y_t \right) \left(n^{-1} \sum_{t=e-n+1}^e \Delta y_t \right),$$

i.e. $\hat{\beta}_1 \hat{\beta}_2$ with $e - n$ replacing $\lfloor \tau_2 T \rfloor$. This statistic is therefore suitable for monitoring for a collapse that begins at time $t = e - n + 1$.

We next consider introducing a variance standardization to imbue $N_{e,m,n}$ with a degree of robustness to possible changes in the unconditional variance of the errors ε_t ; such changes are formally excluded from our assumptions but may of course occur in practice. We propose standardizing the components of $N_{e,m,n}$ by error variance estimates obtained over the respective sub-samples $t = e - n - m + 1, \dots, e - n$ and $t = e - n + 1, \dots, e$. For the sub-sample $t = e - n - m + 1, \dots, e - n$, given the first-order autoregressive structure of the model given in (1) and (2), we consider the error variance estimator $m^{-1} \sum_{t=e-n-m+1}^{e-n} \hat{\varepsilon}_t^2$ where the $\hat{\varepsilon}_t$ are OLS residuals obtained from a regression of Δy_t on a constant and y_{t-1} , over the sub-sample $t = e - n - m + 1, \dots, e - n$. For the sub-sample $t = e - n + 1, \dots, e$, we will below advocate use of very small values of n in the monitoring procedure, hence a regression-based error variance estimator is impractical. Instead, we propose the simpler $H_0^{(1)}$ -based error variance estimator $n^{-1} \sum_{t=e-n+1}^e (\Delta y_t)^2$. Standardizing $N_{e,m,n}$ by these two error variance estimators results in the following statistic:

$$\frac{\left(m^{-1} \sum_{t=e-n-m+1}^{e-n} \Delta y_t \right) \left(n^{-1} \sum_{t=e-n+1}^e \Delta y_t \right)}{\sqrt{m^{-1} \sum_{t=e-n-m+1}^{e-n} \hat{\varepsilon}_t^2} \sqrt{n^{-1} \sum_{t=e-n+1}^e (\Delta y_t)^2}}.$$

Note that the m^{-1} and n^{-1} constants are not needed in the monitoring procedure that follows, since such constant scalings apply for all e and become redundant when comparing training and monitoring period statistics. Consequently, our proposed statistic for monitoring for a stationary collapse that begins at time $t = e - n + 1$ is given by:

$$S_{e,m,n} = \frac{\sum_{t=e-n-m+1}^{e-n} \Delta y_t \sum_{t=e-n+1}^e \Delta y_t}{\sqrt{\sum_{t=e-n-m+1}^{e-n} \hat{\epsilon}_t^2 \sum_{t=e-n+1}^e (\Delta y_t)^2}} \tag{10}$$

The real-time monitoring for collapse procedure we propose proceeds as follows, mirroring the AHLST procedure described above. First, the $S_{e,m,n}$ statistic is computed over the training sample $t = 1, \dots, T^*$ for all possible rolling sub-samples of length $m + n$, and the critical value to be used in monitoring is set to the minimum of these training sample statistics, denoted S_{\min}^* . As we are only concerned with monitoring for a collapse following detection of a prior explosive regime, we now consider the situation where monitoring using the $A_{\text{MAX}}(k)$ procedure has signalled the presence of an explosive regime at some time period $t = T^\diamond$. Conditional on this finding, we then switch to monitoring for a stationary collapse regime by computing the $S_{e,m,n}$ statistic on a rolling basis for $e = T^\diamond + 1, T^\diamond + 2, \dots$. Detection of a collapse regime is triggered at the first point where a monitoring statistic $S_{e,m,n}$ falls below the critical value S_{\min}^* . At an arbitrary point in the collapse monitoring period, $t = e$, we can then write the monitoring decision rule as:

$$\text{Detect } H_1^{(2)} \text{ at time } e \text{ if } S_{e,m,n} < S_{\min}^* \tag{11}$$

We refer to this stationary collapse monitoring procedure as $S_{\text{MIN}}(m, n)$.

The FPR of the $S_{\text{MIN}}(m, n)$ procedure is more difficult to establish compared to the $A_{\text{MAX}}(k)$ procedure for explosive regime detection, since monitoring using $S_{\text{MIN}}(m, n)$ is only performed following detection of an explosive regime by $A_{\text{MAX}}(k)$. When $S_{\text{MIN}}(m, n)$ is conducted and $H_0^{(1)}$ is true, it follows that the explosive regime detection signalled by $A_{\text{MAX}}(k)$ was false, and hence the FPR of $S_{\text{MIN}}(m, n)$ is bounded by the FPR of $A_{\text{MAX}}(k)$ at the point $t = T^\diamond$ (where explosive behaviour was erroneously detected). When $S_{\text{MIN}}(m, n)$ is conducted and $H_1^{(1)}$ is true, which is the main case of interest, the $A_{\text{MAX}}(k)$ explosive regime detection at time $t = T^\diamond$ was correct, and the FPR of the subsequently implemented $S_{\text{MIN}}(m, n)$ procedure is bounded by the true positive rate of $A_{\text{MAX}}(k)$ at time $t = T^\diamond$. Under $H_0^{(2)}$, when monitoring for a collapse at a point e that is $O(T)$ observations into the explosive regime, i.e. $e - \lfloor \tau_1 T \rfloor = O(T)$, then the explosive properties of the process are dominant and a result similar to Theorem 2(a) below can be obtained to show that $\lim_{T \rightarrow \infty} \Pr(S_{e,m,n} < S_{\min}^*) = 0$, i.e. the limit probability of spurious collapse detection by $S_{e,m,n}$ is zero at this point. However, given that an explosive regime can be detected by the $A_{\text{MAX}}(k)$ procedure after only a *finite* number of monitoring periods, we cannot assume that monitoring for a collapse will only be done from a point that can be considered $O(T)$ observations into the explosive regime. As such, establishing a result for $\Pr(S_{e,m,n} < S_{\min}^*)$ in this case, and consequently a result for the FPR associated with $S_{\text{MIN}}(m, n)$, is not tractable, since it will inevitably depend on DGP parameters such as the magnitude of the explosive autoregressive parameter and the duration of the explosive

regime. Consequently, it is not possible to quantify analytically the FPR of the $S_{\text{MIN}}(m, n)$ procedure under $H_0^{(2)}$ for an arbitrary monitoring point e . Instead we examine the FPR of $S_{\text{MIN}}(m, n)$ by simulation in the next section for a range of DGP settings.

In the following theorem, we establish the large sample properties of the $S_{\text{MIN}}(m, n)$ procedure under the stationary collapse hypothesis $H_1^{(2)}$, for different points e in the region of the explosive regime endpoint $\lfloor \tau_2 T \rfloor$.

Theorem 2. Under $H_1^{(2)}$ and assuming that $\{\varepsilon_t\}$ satisfies the mixing conditions of Ferreira and Scotto (2002, p. 476), then as $T \rightarrow \infty$:

(a) If $e = \lfloor \tau_2 T \rfloor - j$ with $j = 0, \dots, c$ and c finite,

$$\lim_{T \rightarrow \infty} \Pr(S_{e,m,n} < S_{\text{min}}^*) = 0.$$

(b) If $e = \lfloor \tau_2 T \rfloor + j$ with $j = 1, \dots, n - 1$,

$$\lim_{T \rightarrow \infty} \Pr(S_{e,m,n} < S_{\text{min}}^*) = \begin{cases} 0 & (1 - \delta_2)^j > (1 + \delta_1)^{j-n} \\ \in \{0, 1\} & (1 - \delta_2)^j = (1 + \delta_1)^{j-n} \\ 1 & (1 - \delta_2)^j < (1 + \delta_1)^{j-n} \end{cases}.$$

(c) If $e = \lfloor \tau_2 T \rfloor + n$,

$$\lim_{T \rightarrow \infty} \Pr(S_{e,m,n} < S_{\text{min}}^*) = 1.$$

Part (a) of Theorem 2 shows that, asymptotically, $S_{e,m,n}$ will never fall below S_{min}^* within a finite region prior to the first observation in the stationary collapse regime ($\lfloor \tau_2 T \rfloor + 1$). Consequently, the $S_{\text{MIN}}(m, n)$ monitoring procedure will not indicate a collapse in this region earlier than it occurs. At the other extreme, part (c) demonstrates that, asymptotically under $H_1^{(2)}$, $S_{\text{MIN}}(m, n)$ will always signal a collapse by time period $\lfloor \tau_2 T \rfloor + n$, thereby ensuring detection of the collapse regime with a delay of no more than $n - 1$ periods. When $n = 1$, therefore, $S_{\text{MIN}}(m, n)$ guarantees collapse detection at exactly the first point of this regime. When $n > 1$, part (b) of Theorem 2 becomes relevant, and the result shows that $S_{\text{MIN}}(m, n)$ may indicate a collapse with a delay of less than $n - 1$ periods, depending on the magnitudes of δ_1 and δ_2 . For example, when $n = 2$, a collapse will be indicated with no delay (i.e. at time period $\lfloor \tau_2 T \rfloor + j$ with $j = 1$) when $(1 - \delta_2) < (1 + \delta_1)^{-1}$, but a one period delay will be induced when $(1 - \delta_2) > (1 + \delta_1)^{-1}$, while either outcome is possible when $(1 - \delta_2) = (1 + \delta_1)^{-1}$.² When $n = 3$, a collapse will be indicated with no delay ($j = 1$) if $(1 - \delta_2) < (1 + \delta_1)^{-2}$ and a one period delay ($j = 2$) if $(1 - \delta_2)^2 < (1 + \delta_1)^{-1}$.³

In practice choices must be made for m and n . Setting $n = 1$ affords the opportunity of detecting a collapse most quickly, since (11) can signal the presence of a stationary collapse regime when the monitoring date is the very first observation of the collapse regime, while setting $n > 1$ would represent a more risk-averse strategy, reducing the chance of an outlying downward movement in y_t during an ongoing explosive regime

²A sufficient condition for $(1 - \delta_2) < (1 + \delta_1)^{-1}$ is that $\delta_2 > \delta_1$, i.e. that the stationary collapse offset is larger in magnitude than the explosive offset.

³A sufficient condition for $(1 - \delta_2) < (1 + \delta_1)^{-2}$ is that $\delta_2 > 2\delta_1$.

spuriously triggering detection of a collapse. Such considerations are explored in the finite sample simulations of the next section.

Note that the real-time monitoring methodology employed ensures that the $S_{\text{MIN}}(m, n)$ procedure is robust to conditional heteroskedasticity and serial correlation in ε_t , in line with the robustness properties of $A_{\text{MAX}}(k)$. The procedure will also be asymptotically robust to a finite number of volatility shifts that occur over the training or monitoring sample periods, since only a finite number of sub-sample statistics will have a variance standardization contaminated by variance changes, hence the effect on the procedure becomes asymptotically negligible.

V. Simulation results

To examine the finite sample performance of the $S_{\text{MIN}}(m, n)$ crash monitoring procedures, we consider a Monte Carlo simulation exercise with data generated by (1)–(2) which allows for a single explosive and collapse regime, with $\mu = 0$ (without loss of generality) and $\varepsilon_t \sim \text{IIDN}(0, 1)$. We set $u_1 = 100$ to ensure that, under $H_1^{(1)}$, we generate only positive explosive regimes. We evaluate the performance of $S_{\text{MIN}}(m, n)$ statistics computed for $m = \{5, 10, 15\}$ and $n = \{1, 2, 3\}$, in each case using $A_{\text{MAX}}(k)$ test statistics for prior bubble detection, with $k = \{5, 10, 15\}$. We set the beginning of the monitoring period to $T^\dagger = 200$ such that the training sample end date is $T^* = 200 - k$ throughout. Monte Carlo simulations are conducted using 10,000 replications in Gauss 20.

First, we consider the empirical FPR of the $S_{\text{MIN}}(m, n)$ monitoring procedures when no stationary collapse regime is present. Figure 1 displays the cumulative rejection frequencies of the $A_{\text{MAX}}(k)$ and $S_{\text{MIN}}(m, n)$ procedures under no collapse scenarios. The cumulative rejection frequency of $A_{\text{MAX}}(k)$ at T' is defined as the proportion of replications where we detect $H_1^{(1)}$ at any point in the monitoring period up to and including T' . The cumulative rejection frequency of $S_{\text{MIN}}(m, n)$ at T' is defined as the proportion of replications where we detect $H_1^{(2)}$ at any point in the monitoring period up to and including T' .⁴ The cumulative rejection frequencies of $A_{\text{MAX}}(k)$ under $H_0^{(1)}$ and of $S_{\text{MIN}}(m, n)$ under $H_0^{(1)}$ or $H_0^{(2)}$ represent the empirical FPRs for the procedures. Figure 1a–c considers the case of $H_0^{(1)}$ where there is neither an explosive bubble nor stationary collapse such that $\tau_1 = 1$, while Figure 1d–f examines $H_0^{(2)}$ where an explosive regime occurs and continues until the end of the monitoring period without stationary collapse. Beginning with the $H_0^{(1)}$ case, it is apparent from Figure 1a–c that the empirical FPR of $A_{\text{MAX}}(k)$ tracks its theoretical FPR very closely, confirming the results of AHLST, while the empirical FPRs of the $S_{\text{MIN}}(m, n)$ procedures are always somewhat lower than this theoretical FPR for all combinations of m and n considered here. As discussed in section IV, under $H_0^{(1)}$ the FPR of the $S_{\text{MIN}}(m, n)$ procedures should be bounded by the theoretical FPR for $A_{\text{MAX}}(k)$ and our simulation results demonstrate that this holds in finite samples.

⁴Another interpretation of the cumulative rejection frequency of $S_{\text{MIN}}(m, n)$ is the joint probability of $A_{\text{MAX}}(k)$ and $S_{\text{MIN}}(m, n)$ rejecting by a certain point in the monitoring period. Note though that the $S_{\text{MIN}}(m, n)$ procedure itself is conditional, in the sense that a rejection of $H_0^{(1)}$ using $A_{\text{MAX}}(k)$ is required before the $S_{\text{MIN}}(m, n)$ procedure is ever applied.

Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$ under $H_0^{(1)}$ and $H_0^{(2)}$

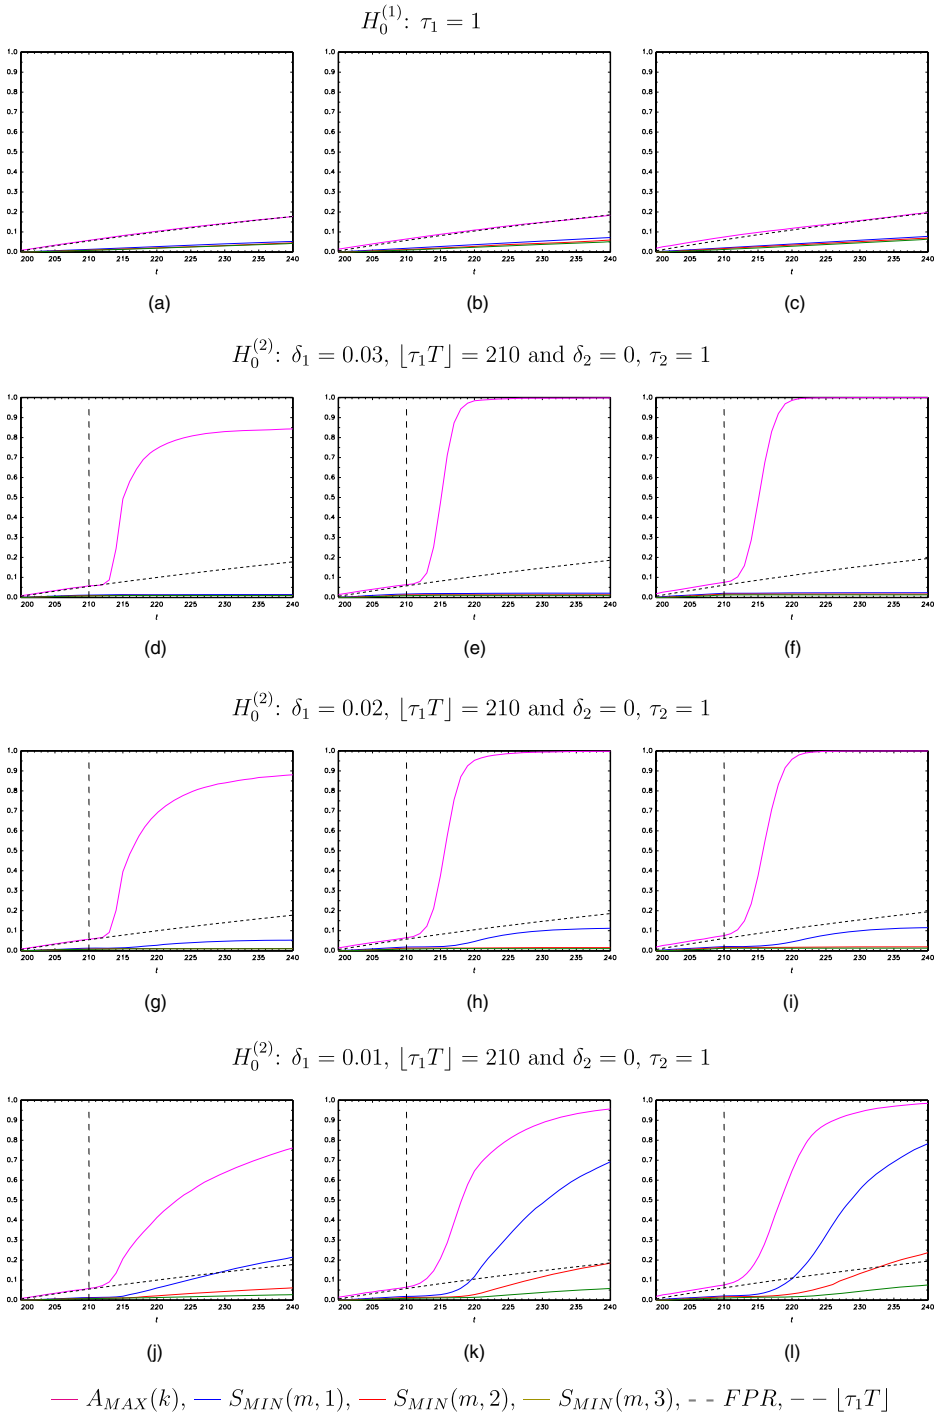


FIGURE 1. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$ under $H_0^{(1)}$ and $H_0^{(2)}$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $m = k = 5$; (e) $m = k = 10$; (f) $m = k = 15$; (g) $m = k = 5$; (h) $m = k = 10$; (i) $m = k = 15$; (j) $m = k = 5$; (k) $m = k = 10$; (l) $m = k = 15$

In Figure 1d–f, we now consider the case of an explosive regime which begins at $[\tau_1 T] = 210$ and continues until the end of the monitoring period. The magnitude of the explosive regime is set to $\delta_1 = 0.03$. The empirical FPR of $A_{\text{MAX}}(k)$ now tracks the theoretical FPR closely until the beginning of the explosive regime, at which point empirical rejection frequencies increase substantially due to the detection of this regime. The empirical FPRs of the $S_{\text{MIN}}(m, n)$ procedures, however, lie close to zero throughout the monitoring period for all settings of m and n considered, suggesting a reassuring degree of FPR control for this bubble magnitude setting. A similar set of results is obtained in Figure 1g–i where we now consider a smaller magnitude explosive regime, setting $\delta_1 = 0.02$. In this case, the empirical FPR of the $S_{\text{MIN}}(m, 1)$ procedure increases slightly at the beginning of the explosive regime, before levelling off to a point below 0.06 in the case of $m = 5$ and below 0.12 in the case of $m = \{10, 15\}$. As in the $\delta_1 = 0.03$ case, the empirical FPRs of $S_{\text{MIN}}(m, 2)$ and $S_{\text{MIN}}(m, 3)$ lie close to zero throughout the monitoring period. Finally, turning our attention to Figure 1j–l, we now set $\delta_1 = 0.01$ such that the magnitude of the explosive regime is smaller still. In this case, the empirical FPR of $S_{\text{MIN}}(m, 1)$ is seen to increase beyond the level seen previously, with this being particularly apparent for $m = \{10, 15\}$ where high FPR levels are ultimately obtained. $S_{\text{MIN}}(m, 2)$ and $S_{\text{MIN}}(m, 3)$, however, display much lower empirical FPRs throughout the monitoring period. These latter procedures therefore offer a greater degree of FPR control relative to $S_{\text{MIN}}(m, 1)$.

It is clear that the empirical FPRs of the $S_{\text{MIN}}(m, n)$ procedures are influenced by the magnitude of the explosive parameter. The closer the explosive offset is to zero, the higher the likelihood of negative changes in y_t occurring during the explosive regime, and thus the higher the probability of rejection by $S_{\text{MIN}}(m, n)$. The relationship between the empirical FPR of $S_{\text{MIN}}(m, n)$ and δ_1 is non-monotone, in the sense that the empirical FPR will be smaller in the case of $\delta_1 = 0$ and in the case of higher values of $\delta_1 \neq 0$, than for lower values of $\delta_1 \neq 0$. With reference to the statistics' motivation in section IV, during the bubble regime, $\hat{\beta}_1$ is typically large and positive, hence even a modest negative $\hat{\beta}_2$ can generate a large negative statistic and trigger crash detection. The potential for a high FPR in the presence of a small magnitude bubble is particularly acute for $S_{\text{MIN}}(m, 1)$ since only a one period decrease is needed to make $\hat{\beta}_2$ negative, an outcome that is potentially sufficient to trigger spurious detection of a crash. In contrast, the potential for spurious detection of a crash regime is substantially reduced by using $S_{\text{MIN}}(m, 2)$ or $S_{\text{MIN}}(m, 3)$, since, other things equal, $\hat{\beta}_2$ is less likely to be negative in the presence of only a one period decrease. The FPRs of these procedures are consequently less sensitive to the magnitude of the explosive regime.

We now turn to examining the finite sample performance of our monitoring procedures under the alternative hypothesis, $H_1^{(2)}$, where the explosive regime is followed by a stationary collapse. Figures 2–6 report the cumulative rejection frequencies for the $A_{\text{MAX}}(k)$ and $S_{\text{MIN}}(m, n)$ procedures under this hypothesis, as well as histograms of crash detection dates obtained by $S_{\text{MIN}}(m, n)$, to allow us to evaluate both the power and crash timing properties of our procedures.⁵

⁵As noted in section II, it is possible that when an explosive regime terminates it is followed by unit root behaviour rather than a collapse. While such a scenario is less common in practice, and as such is not the focus of this paper,

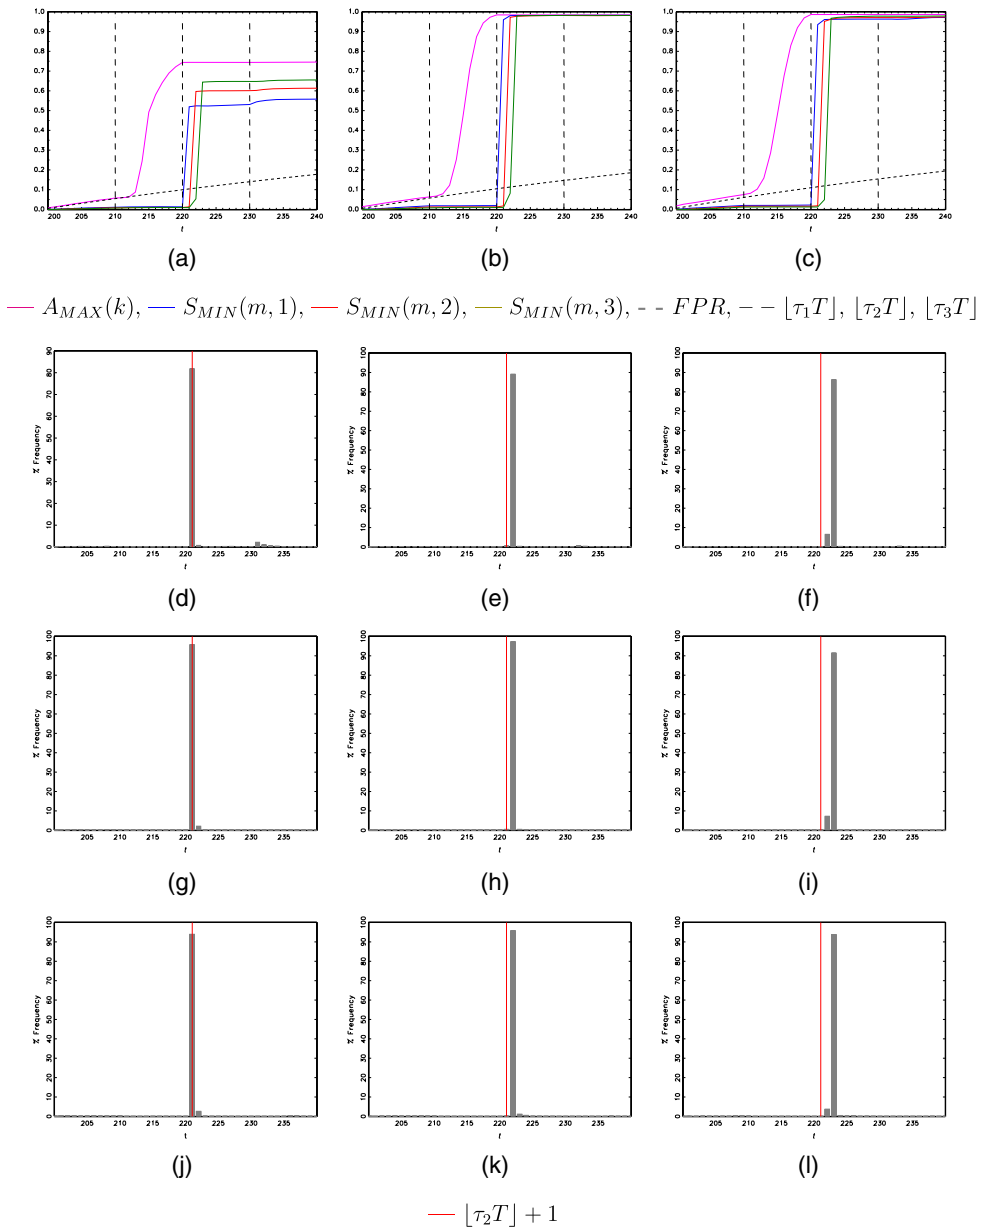


FIGURE 2. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$, and histograms of $S_{MIN}(m, n)$ detection dates: $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 220$, $\lfloor \tau_3 T \rfloor = 230$, $\delta_1 = 0.03$ and $\delta_2 = 0.015$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $S_{MIN}(5, 1)$; (e) $S_{MIN}(5, 2)$; (f) $S_{MIN}(5, 3)$; (g) $S_{MIN}(10, 1)$; (h) $S_{MIN}(10, 2)$; (i) $S_{MIN}(10, 3)$; (j) $S_{MIN}(15, 1)$; (k) $S_{MIN}(15, 2)$; (l) $S_{MIN}(15, 3)$

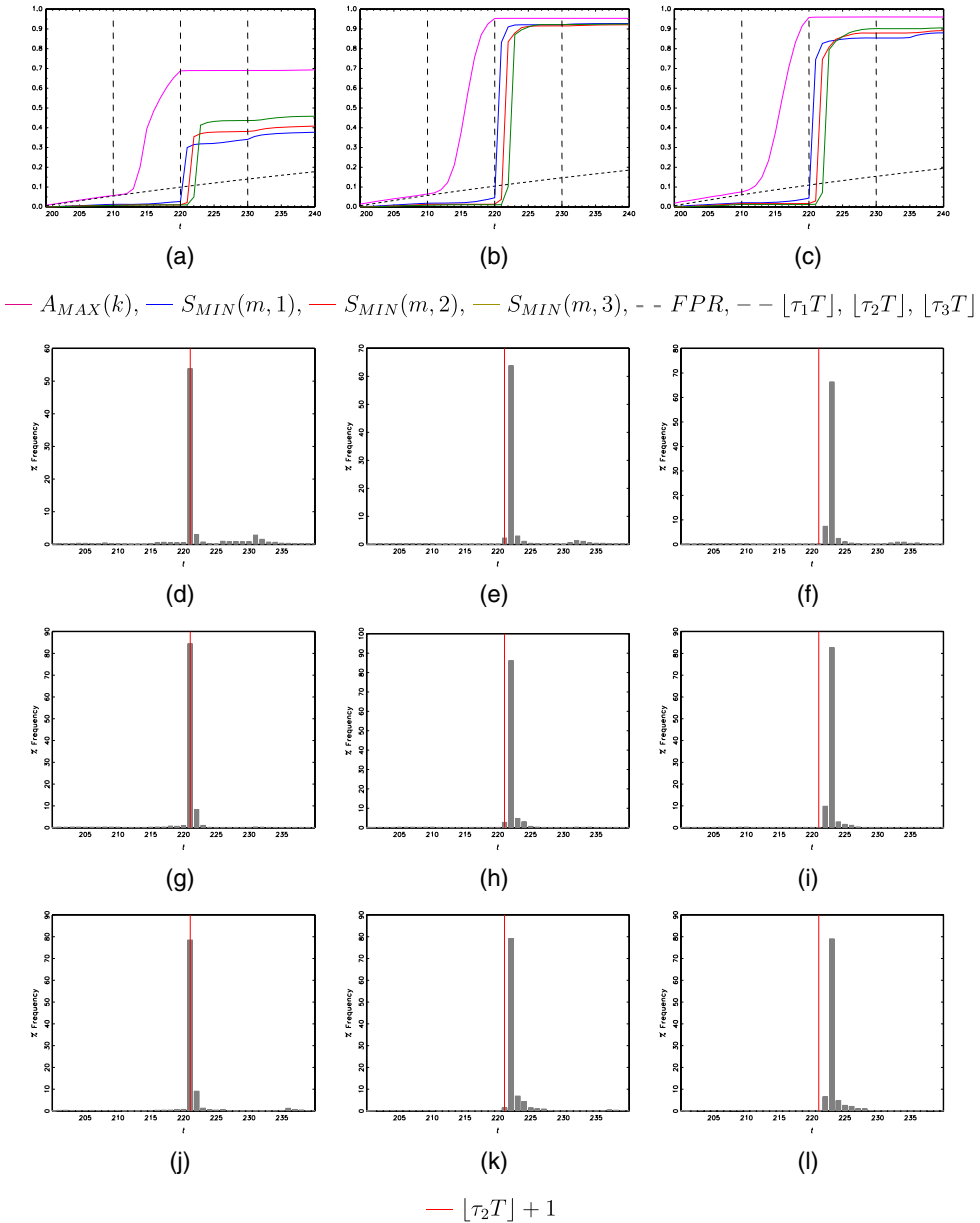


FIGURE 3. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$, and histograms of $S_{MIN}(m, n)$ detection dates: $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 220$, $\lfloor \tau_3 T \rfloor = 230$, $\delta_1 = 0.02$ and $\delta_2 = 0.01$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $S_{MIN}(5, 1)$; (e) $S_{MIN}(5, 2)$; (f) $S_{MIN}(5, 3)$; (g) $S_{MIN}(10, 1)$; (h) $S_{MIN}(10, 2)$; (i) $S_{MIN}(10, 3)$; (j) $S_{MIN}(15, 1)$; (k) $S_{MIN}(15, 2)$; (l) $S_{MIN}(15, 3)$

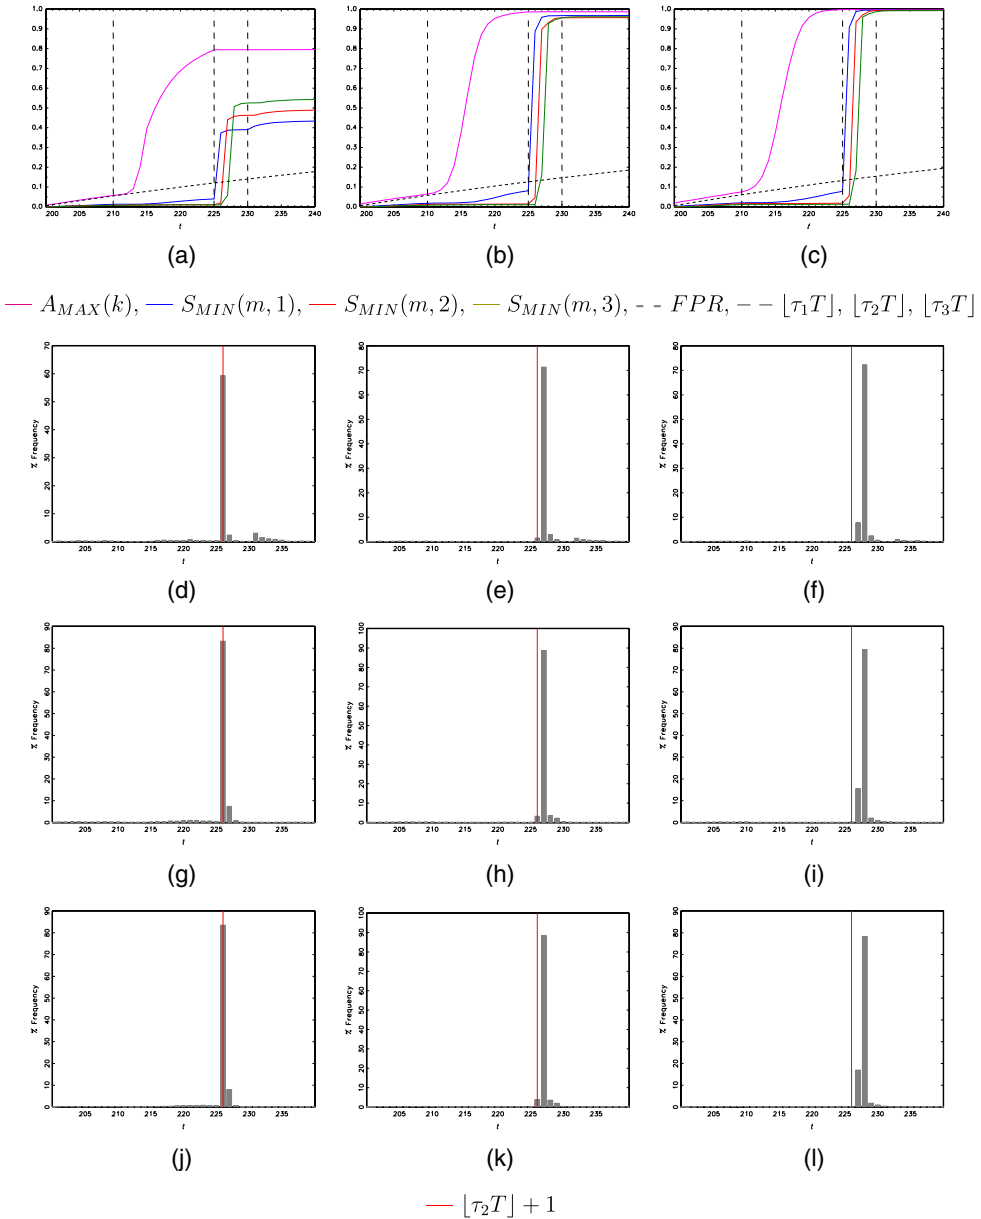


FIGURE 4. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$, and histograms of $S_{MIN}(m, n)$ detection dates: $[\tau_1 T] = 210$, $[\tau_2 T] = 225$, $[\tau_3 T] = 230$, $\delta_1 = 0.02$ and $\delta_2 = 0.01$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $S_{MIN}(5, 1)$; (e) $S_{MIN}(5, 2)$; (f) $S_{MIN}(5, 3)$; (g) $S_{MIN}(10, 1)$; (h) $S_{MIN}(10, 2)$; (i) $S_{MIN}(10, 3)$; (j) $S_{MIN}(15, 1)$; (k) $S_{MIN}(15, 2)$; (l) $S_{MIN}(15, 3)$

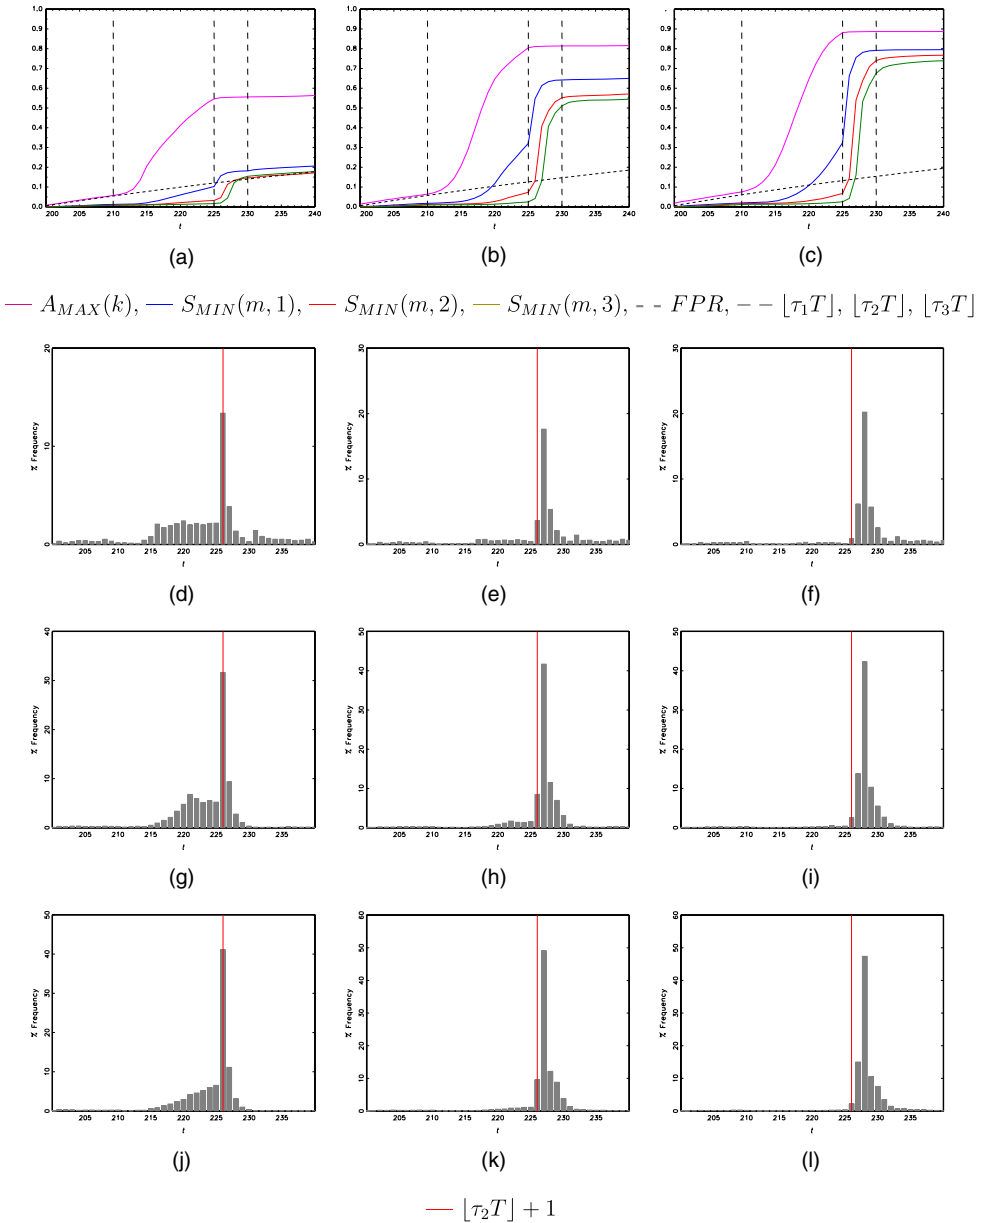


FIGURE 5. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$, and histograms of $S_{MIN}(m, n)$ detection dates: $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 225$, $\lfloor \tau_3 T \rfloor = 230$, $\delta_1 = 0.01$ and $\delta_2 = 0.005$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $S_{MIN}(5, 1)$; (e) $S_{MIN}(5, 2)$; (f) $S_{MIN}(5, 3)$; (g) $S_{MIN}(10, 1)$; (h) $S_{MIN}(10, 2)$; (i) $S_{MIN}(10, 3)$; (j) $S_{MIN}(15, 1)$; (k) $S_{MIN}(15, 2)$; (l) $S_{MIN}(15, 3)$

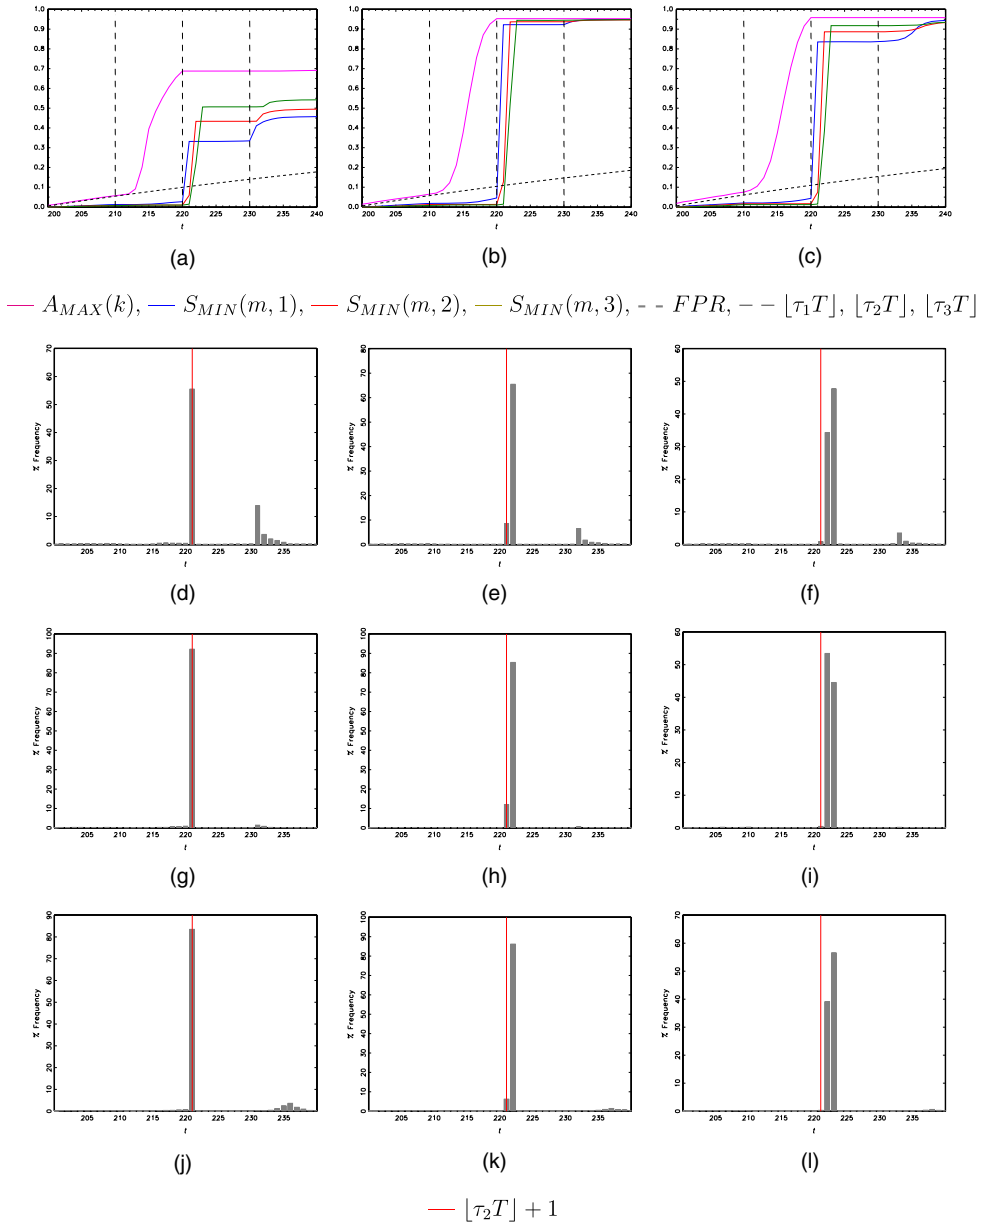


FIGURE 6. Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m, n)$, and histograms of $S_{MIN}(m, n)$ detection dates: $\lceil \tau_1 T \rceil = 210$, $\lceil \tau_2 T \rceil = 220$, $\lceil \tau_3 T \rceil = 230$, $\delta_1 = 0.02$ and $\delta_2 = 0.03$. (a) $m = k = 5$; (b) $m = k = 10$; (c) $m = k = 15$; (d) $S_{MIN}(5, 1)$; (e) $S_{MIN}(5, 2)$; (f) $S_{MIN}(5, 3)$; (g) $S_{MIN}(10, 1)$; (h) $S_{MIN}(10, 2)$; (i) $S_{MIN}(10, 3)$; (j) $S_{MIN}(15, 1)$; (k) $S_{MIN}(15, 2)$; (l) $S_{MIN}(15, 3)$

In Figures 2 and 3 we set $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 220$ and $\lfloor \tau_3 T \rfloor = 230$ such that both the explosive and stationary collapse regimes have a duration of 10 observations. We set the collapse magnitude to be half of the explosive magnitude such that $\delta_2 = \delta_1/2$ and consider ‘high’ magnitude settings of $\{\delta_1, \delta_2\} = \{0.03, 0.015\}$ in Figure 2 and ‘low’ magnitude settings of $\{\delta_1, \delta_2\} = \{0.02, 0.01\}$ in Figure 3. Examining first Figure 2, we note that the $S_{\text{MIN}}(m, n)$ procedure achieves rejection frequencies of between approximately 0.53 and 0.65 by the end of the collapse regime in the case of $m = 5$, whilst the rejection frequencies are very close to 1 in the case of $m = 10$ and $m = 15$, across all n . These rejection frequencies are close to or the same as $A_{\text{MAX}}(k)$, particularly in the case of $m = 10$ and $m = 15$, so that the collapse is detected in almost all cases where a bubble is found.

While the rejection frequencies for bubble detection by $A_{\text{MAX}}(k)$ are increasing over the duration of the explosive regime, one key feature of $S_{\text{MIN}}(m, n)$ is that the majority of rejections occur at one point in time. In the case of $S_{\text{MIN}}(m, 1)$, this is the first observation of the collapse regime, whereas for $S_{\text{MIN}}(m, 2)$ and $S_{\text{MIN}}(m, 3)$ a one and two observation lag is introduced, respectively. To examine this in more detail, Figure 2d–l displays histograms of the date at which the $S_{\text{MIN}}(m, n)$ procedures detected the crash in each Monte Carlo replication where one was detected during the monitoring period. For example, comparing Figure 2d–f, it is clear that when a rejection is found by the $S_{\text{MIN}}(5, 1)$ procedure it is on the true crash date, $\lfloor \tau_2 T \rfloor + 1$, in almost all replications. This suggests that, in line with the theoretical results of section IV, our procedure is highly capable of real-time detection of crashes. For $S_{\text{MIN}}(5, 2)$ almost all rejections occur at $\lfloor \tau_2 T \rfloor + 2$, i.e. one observation after the true crash date, and for $S_{\text{MIN}}(5, 3)$ almost all occur at $\lfloor \tau_2 T \rfloor + 3$, two observations after the true crash date. Hence the key asymptotic result that use of $S_{\text{MIN}}(m, n)$ gives rise to crash detection with a delay of at most $n - 1$ periods is borne out in finite samples. Given that under $H_0^{(2)}$, use of $n > 1$ offers a greater degree of FPR control relative to using $n = 1$, a trade-off exists between the speed of possible crash detection and the probability of spuriously finding a crash in an ongoing bubble.

Examining now the ‘low’ magnitude settings in Figure 3, our results are qualitatively similar to those discussed in the ‘high’ magnitude case. The lower magnitude explosive and stationary collapse regimes lead to a reduction in the cumulative rejection frequencies, as we would expect, with these reductions more noticeable in the case of $m = 5$. For $m = 10$ and $m = 15$ the rejection frequencies of $S_{\text{MIN}}(m, n)$ range from 0.85 to 0.92 under these settings, therefore still maintaining excellent power to detect a crash. The histograms in Figure 3d–l show that the immediate detection of a crash by $S_{\text{MIN}}(m, 1)$ and one or two observation lagged detection by $S_{\text{MIN}}(m, 2)$ and $S_{\text{MIN}}(m, 3)$ discussed previously holds for this lower magnitude setting.

In Figures 4 and 5, we consider settings of $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 225$ and $\lfloor \tau_3 T \rfloor = 230$ such that the explosive regime now lasts 15 observations and the collapse regime only five observations to examine the impact on our procedures of varying regime lengths. Again, we consider ‘high’ and ‘low’ magnitude settings for the explosive and stationary

we note that in unreported simulations our $S_{\text{MIN}}(m, n)$ procedures can detect the termination of an explosive regime even when it is not followed by a collapse.

regimes. As the explosive regime is now longer in duration, our ‘high’ magnitude settings (reported in Figure 4) are the ‘low’ settings that we used previously to consider a 10 observation explosive regime, namely $\{\delta_1, \delta_2\} = \{0.02, 0.01\}$. Figure 5 considers ‘low’ magnitude settings of $\{\delta_1, \delta_2\} = \{0.01, 0.005\}$.

Examining Figure 4, the results obtained are very similar to those discussed above, with moderate cumulative rejection frequencies in the case of $m = 5$ and very high frequencies in the cases of $m = 10$ and $m = 15$. As before, we see that setting $n = 1$ yields detection dates equal to the true crash date in the majority of replications, whereas setting $n > 1$ leads to a $n - 1$ detection delay.

Turning our attention now to Figure 5, we note that the cumulative rejection frequencies for $S_{\text{MIN}}(5, n)$ are low, while $S_{\text{MIN}}(10, n)$ and $S_{\text{MIN}}(15, n)$ offer much more promising crash detection levels with all settings of n offering reasonable rejection frequencies. We note, however, that under these settings the cumulative rejection frequencies of $S_{\text{MIN}}(m, 1)$ are quite high during the explosive regime, giving a further indication of a potential lack of good FPR control when $n = 1$. This behaviour mirrors that of the empirical FPR in Figure 1j–l that we discussed previously. Crucially, such behaviour is not observed for $S_{\text{MIN}}(m, 2)$ and $S_{\text{MIN}}(m, 3)$, with these procedures maintaining low spurious rejection frequencies throughout the explosive phase. To examine this behaviour further, consider the histogram of detected crash dates for $S_{\text{MIN}}(10, 1)$ in Figure 5g. While the majority of detections occur for this procedure at the true crash date, we observe that smaller numbers of false rejections arise at dates before the crash occurs. If we now consider the histogram of detected crash dates for $S_{\text{MIN}}(10, 2)$ in Figure 5h, the majority of these false rejections before the true crash date have been eliminated. Of course, the trade-off here is that detection of the crash occurs at $\lfloor \tau_2 T \rfloor + 2$ in the majority of replications such that we have introduced a one observation delay in detection.

Finally, in Figure 6 we examine the case where the crash magnitude is a little larger than the explosive magnitude. We set $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 220$ and $\lfloor \tau_3 T \rfloor = 230$ as before, and now consider $\{\delta_1, \delta_2\} = \{0.02, 0.03\}$ such that $\delta_2 > \delta_1$. As discussed in section IV, an implication of Theorem 2 is that the $S_{\text{MIN}}(m, 2)$ procedure will asymptotically indicate a crash with no delay in this scenario. Figure 6e,h,k shows that the crash is detected at the true date in a number of cases, but the majority of detection dates remain one period late. This is perhaps to be expected in finite samples, since for these settings of δ_1 and δ_2 , the relevant inequality driving the theoretical result, $(1 - \delta_2) < (1 + \delta_1)^{-1}$, is only just satisfied ($0.97 < 0.98$). For $S_{\text{MIN}}(m, 3)$, the results of Theorem 2 indicate that crash detection should occur with a one period delay asymptotically, since $(1 - \delta_2) > (1 + \delta_1)^{-2}$ but $(1 - \delta_2)^2 < (1 + \delta_1)^{-1}$. Figure 6f,i,l demonstrates that $S_{\text{MIN}}(m, 3)$ does detect the crash with a one period delay in a substantial number of cases, with the remainder of detections obtained with a two-period delay.

The results of this section demonstrate that our $S_{\text{MIN}}(m, n)$ monitoring procedure is able to detect a crash with a high degree of power either at or very close to the date at which the crash occurs for the vast majority of settings considered here. In general, setting $m = 10$ or $m = 15$ yields higher rejection frequencies than $m = 5$. Comparing $m = 10$ and $m = 15$, we observe a slight advantage for the $S_{\text{MIN}}(m, n)$ procedure where the value of m matches the length of the explosive regime. This is more evident in the ‘low’ magnitude settings for the explosive and stationary regimes than the ‘high’ magnitude settings. In practice, we

will not know the length of any explosive regime at the point that we begin monitoring for its crash. We suggest, based on the simulation results shown here, that selecting $m = 10$ will be suitable for many scenarios and we use this setting in what follows. We have also demonstrated how the flexibility in the procedure's construction allows a practitioner to prioritise their monitoring preferences. Immediate detection of a crash is clearly important for policy makers, allowing them to react to changes in market conditions as they occur. Our results demonstrate that setting $n = 1$ allows for the immediate detection of a crash. However, as observed in Figure 5, when the magnitude of an explosive process is small, downwards movement in the series during the explosive phase could potentially trigger false crash detection. In some contexts, any decline in prices will be of interest, but to others this feature may be less desirable. Our results demonstrate that setting $n > 1$ reduces the probability of these pre-emptive crash detections occurring, with the obvious trade-off that when crashes do occur they may be detected with a delay of up to $n - 1$ observations. Of course, depending on the practitioner's motivations, even more risk-averse approaches to crash detection than the $n = \{2, 3\}$ procedures we considered in this section could be implemented. However our results show that introducing just a one observation delay is sufficient to eliminate the majority of false detections that arise in situations where the explosive bubble is small in magnitude. Given this, we suggest that setting $n = 2$ will be suitable for many scenarios, but with practitioners retaining the flexibility to adjust this parameter to suit their preferences and the particular monitoring context.

VI. Monitoring for multiple bubble and crash regimes

The testing approach outlined in sections III and IV concerns monitoring for a single bubble and crash episode. However, it may be the case that once a crash has been detected, rather than ending the monitoring exercise, a practitioner wishes to continue monitoring for subsequent bubble episodes. In this section we consider how the testing approach outlined in this paper could be extended to deal with monitoring for multiple bubbles and crashes. Consider the following multiple bubble and crash DGP in which we allow for $j = 1, \dots, N$ explosive bubble and stationary collapse regimes. For $t = 1, \dots, T$:

$$y_t = \mu + x_t + u_t, \quad (12)$$

$$u_t = (1 + \delta_t)u_{t-1} + \varepsilon_t + u_1 \mathbb{I}(t = \lfloor \tau_{j,3} T \rfloor + 1) \quad (13)$$

$$\delta_t = \sum_{j=1}^N \left\{ \delta_{j,1} \mathbb{I}(\lfloor \tau_{j,1} T \rfloor < t \leq \lfloor \tau_{j,2} T \rfloor) - \delta_{j,2} \mathbb{I}(\lfloor \tau_{j,2} T \rfloor < t \leq \lfloor \tau_{j,3} T \rfloor) - \mathbb{I}(t = \lfloor \tau_{j,3} T \rfloor + 1) \right\}$$

$$x_t = \sum_{j=1}^N (u_{\lfloor \tau_{j,3} T \rfloor} - u_1) \mathbb{I}(t > \lfloor \tau_{j,3} T \rfloor),$$

with, as before, $u_1 = O_p(1)$, $\delta_{j,1} > 0$ and $\delta_{j,2} > 0$ for all j . The inclusion of x_t in (12) and the indicator function term in (13) prevents the magnitude of one explosive regime from

entering the dynamics of subsequent explosive regimes; this DGP represents a simple modification of that adopted in Harvey, Leybourne, and Whitehouse (2020), here allowing for non-zero u_1 . Under this specification, y_t can undergo N explosive bubble phases with start and end dates $\lfloor \tau_{j,1}T \rfloor + 1$ and $\lfloor \tau_{j,2}T \rfloor$, respectively, and N stationary crash regimes with start and end dates $\lfloor \tau_{j,2}T \rfloor + 1$ and $\lfloor \tau_{j,3}T \rfloor$, respectively, for $j = 1, \dots, N$.

In order to distinguish between explosive bubble and stationary collapse regimes in this multiple bubble context, our hypotheses of interest become:

$$\begin{aligned} H_{0,j}^{(1)} : \tau_{j,1} &= 1 && \text{(unit root)} \\ H_{1,j}^{(1)} : \tau_{j,1} < \tau_{j,2} &\leq 1 && \text{(unit root then explosive, with or without stationary collapse),} \end{aligned}$$

and

$$\begin{aligned} H_{0,j}^{(2)} : \tau_{j,1} < \tau_{j,2} &= 1 && \text{(unit root then explosive without stationary collapse)} \\ H_{1,j}^{(2)} : \tau_{j,1} < \tau_{j,2} &< 1 && \text{(unit root then explosive with stationary collapse).} \end{aligned}$$

Consider the simple case of $N = 2$ where $\tau_{j,1} < \tau_{j,2} < \tau_{j,3} < 1$ for $j = 1, 2$ such that there exist two explosive bubble regimes, each followed by a stationary collapse. Following the detection of the first bubble and crash regime through rejection of $H_{0,1}^{(1)}$ and subsequently $H_{0,1}^{(2)}$ in favour of $H_{1,1}^{(2)}$, we should then switch back into monitoring for the second explosive bubble regime by considering $H_{0,2}^{(1)}$.

A question that arises in this situation is whether we want to begin monitoring for the next explosive bubble immediately upon detection of a crash, given that the stationary collapse is likely to be still ongoing. To mitigate against possible problems associated with bubble monitoring resuming during a crash regime, we impose a minimum window width gap between the detection of a stationary collapse regime and the start of monitoring for a subsequent explosive bubble. A natural candidate for this minimum window here would be k .

Our multiple bubble and crash monitoring decision rules can therefore be written as follows. First, explosive bubble monitoring is undertaken using $A_{\text{MAX}}(k)$:

$$\text{Detect } H_{1,j}^{(1)} \text{ at time } e \text{ if } A_{e,k} > A_{\text{max}}^*.$$

If an explosive bubble is detected, we denote the detection date as T_j^\diamond . Next, stationary collapse monitoring is undertaken over $e \in [T_j^\diamond + 1, T]$ using $S_{\text{MIN}}(m, n)$:

$$\text{Detect } H_{1,j}^{(2)} \text{ at time } e \text{ if } S_{e,m,n} < S_{\text{min}}^*.$$

The date of stationary collapse detection is denoted $T_j^{\diamond\diamond}$. Monitoring for a subsequent explosive bubble is then undertaken over $e \in [T_j^{\diamond\diamond} + k, T]$ using $A_{\text{MAX}}(k)$:

$$\text{Detect } H_{1,j+1}^{(1)} \text{ at time } e \text{ if } A_{e,k} > A_{\text{max}}^*.$$

Monitoring continues in this manner, switching between $A_{\text{MAX}}(k)$ and $S_{\text{MIN}}(m, n)$ for as long as desired.




FIGURE 7. Rejection frequencies of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for multiple bubble and crash episodes: $\tau_{j,1} = 1$

We consider the performance of our proposed multiple bubble and crash monitoring procedure through Monte Carlo simulation of (12), where we again use $\varepsilon_t \sim IIDN(0, 1)$ and set $\mu = 0$ and $u_1 = 100$. Given the finite sample performance of our monitoring procedure in the single bubble context displayed in section V, we provide results for the $A_{MAX}(10)$ and $S_{MIN}(10,2)$ procedures here. We set $T^* = 200 - k$.

Figure 7 displays the empirical FPRs of the $A_{MAX}(10)$ and $S_{MIN}(10,2)$ procedures under $H_{0,1}^{(1)}$, i.e. with $\tau_{1,1} = 1$. The empirical FPRs of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for one explosive and collapse regime will be identical to those discussed in the single bubble case, displayed in Figure 1. However, in a multiple bubble context, we now wish to consider the empirical FPRs displayed for subsequent bubble regimes. That is, we wish to consider the probability of our procedures falsely detecting more than one bubble and crash regime. We display the empirical FPRs for detecting one, two and three bubble and crash regimes. It is clear from Figure 7 that the empirical FPRs of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for two bubbles/crashes remain small throughout the monitoring period, while the empirical FPRs for three bubbles/crashes are near zero throughout the monitoring period. Given the conditional nature of our multiple bubble procedure, i.e. that we do not monitor for a $(j + 1)^{th}$ bubble unless we have detected a collapse of bubble j , this should not be surprising.

Figure 8 displays the cumulative rejection frequencies of the $A_{MAX}(10)$ and $S_{MIN}(10,2)$ procedures for multiple bubble and collapse regimes, where $\{\lfloor \tau_{1,1}T \rfloor, \lfloor \tau_{1,2}T \rfloor, \lfloor \tau_{1,3}T \rfloor\} = \{215, 225, 235\}$, $\{\lfloor \tau_{2,1}T \rfloor, \lfloor \tau_{2,2}T \rfloor, \lfloor \tau_{2,3}T \rfloor\} = \{255, 265, 275\}$, and $\{\lfloor \tau_{3,1}T \rfloor, \lfloor \tau_{3,2}T \rfloor, \lfloor \tau_{3,3}T \rfloor\} = \{295, 305, 315\}$. That is, the DGP contains three explosive bubble and stationary crash regimes, with the bubble and crash phases for each regime lasting 10 observations. In Figure 8a we set $\delta_{j,1} = 0.03$ and $\delta_{j,2} = 0.015$ for $j = 1, \dots, 3$, such that each explosive bubble (and each stationary crash) is of the same magnitude. It is clear that our proposed procedures have excellent power to detect multiple bubble and crash regimes under these settings, with cumulative rejection frequencies close to 1 for each regime. The speed of detection matches that observed in the single bubble case, with the $S_{MIN}(10,2)$ detection date being equal to the




FIGURE 8. Rejection frequencies of $A_{MAX}(10)$ and $S_{MIN}(10, 2)$ for multiple bubble and crash episodes: $\{[\tau_{1,1}T], [\tau_{1,2}T], [\tau_{1,3}T]\} = \{215, 225, 235\}$, $\{[\tau_{2,1}T], [\tau_{2,2}T], [\tau_{2,3}T]\} = \{255, 265, 275\}$, $\{[\tau_{3,1}T], [\tau_{3,2}T], [\tau_{3,3}T]\} = \{295, 305, 315\}$. (a) $\delta_{j,1} = 0.03$ and $\delta_{j,2} = 0.015$; (b) $\delta_{j,1} = 0.02$ and $\delta_{j,2} = 0.01$

second observation of the stationary collapse regime in most replications. In Figure 8b we consider smaller settings of $\delta_{j,1} = 0.02$ and $\delta_{j,2} = 0.01$ for $j = 1, \dots, 3$. Under these settings, the conditional nature of the testing approach becomes more obvious, as we observe slightly lower cumulative rejection frequencies for later bubble/crash regimes relative to earlier regimes. However, our proposed procedures still obtain very good levels of power for multiple regimes under these settings. We have therefore demonstrated in this section that the procedures proposed in this paper for real-time monitoring of a crash extend simply to a multiple bubble monitoring context.

VII. Empirical application

To demonstrate the effectiveness of our crash monitoring procedure, we consider an empirical application of the $A_{MAX}(10)$ and $S_{MIN}(10, n)$ procedures to the US housing market. The sub-prime mortgage crisis and subsequent financial distress of the late 2000s has led to increased scrutiny of the dynamics of house prices. Several recent studies

have investigated historical bubble behaviour in housing markets (see, *inter alia*, Pavlidis *et al.*, 2016; Anundsen *et al.*, 2016; Anundsen, 2019; Fabozzi *et al.*, 2020). In a recent study, Harvey *et al.* (2020) propose a method of date-stamping multiple bubble and crash regimes based on Bayesian Information Criterion model selection and use this technique to investigate the dynamics of the housing market in 20 OECD countries, finding substantial evidence of both bubbles and crashes across many countries, including the US. While there is now a consensus that the US housing market underwent a bubble during the 2000s, at the time the issue was contested. Addressing the Joint Economic Committee of the US Congress in 2002, Federal Reserve Chairman Alan Greenspan remarked that a comparison of house prices to the bubble and crash behaviour observed in stock markets was not appropriate due to the high transaction costs and limited arbitrage opportunities in housing. He also stated that, instead of a national market, US housing could be seen as a collection of local markets, such that ‘even if a bubble were to develop in a local market, it would not necessarily have implications for the nation as a whole’ (Monetary Policy and the Economic Outlook, 2002). Real-time monitoring techniques such as the one proposed in this paper could have provided evidence of changes in the dynamics of the housing market and allowed policy makers to respond quickly to these events.

We consider a pseudo-real-time monitoring exercise of the US housing market. As discussed in section I, a rational bubble manifests itself as the presence of explosive behaviour in asset prices with the absence of explosive behaviour in the corresponding fundamental values. We therefore examine a price to fundamental ratio for housing, using rent as our proxy of housing fundamentals, as is common in the literature (see Pavlidis *et al.*, 2016, for example). A quarterly house price to rent ratio is obtained from the OECD (OECD, 2021) for the period 1975:Q4 - 2021:Q1, yielding $T = 182$ observations. We begin monitoring for an explosive bubble in 1998:Q1. We select window widths of length $k = 10$ and $m = 10$, such that our preferred $A_{MAX}(10)$ and $S_{MIN}(10, n)$ tests are used. This provides us with a training sample of $T^* = 80$ observations over which our training statistics are computed from $t = 1, \dots, T^*$.

Figure 9 displays the US house price to rent ratio in the first panel and the computed $S_{MIN}(10, 1)$, $S_{MIN}(10, 2)$, and $S_{MIN}(10, 3)$ test statistics in the second panel. We examine the monitoring performance of all three crash procedures here to demonstrate the trade-offs of choosing $n > 1$ in terms of speed of detection. We begin by monitoring for an explosive bubble using $A_{MAX}(10)$, which detects the presence of a bubble in 2000:Q1, pre-dating Greenspan’s remarks by 2 years. The theoretical FPR of $A_{MAX}(10)$ is 0.11 at the point of detection. At this point, we switch into crash monitoring using the $S_{MIN}(10, 1)$, $S_{MIN}(10, 2)$ and $S_{MIN}(10, 3)$ procedures. $S_{MIN}(10, 1)$ detects a crash in 2006:Q2, whilst both $S_{MIN}(10, 2)$ and $S_{MIN}(10, 3)$ detect a crash in 2006:Q3. The effect of increasing n in the $S_{MIN}(10, n)$ monitoring procedure mirrors our theoretical and simulation results, with $n > 1$ introducing a delay in detection of $n - 1$ observations or fewer. We note that the detection delay is one observation for both $S_{MIN}(10, 2)$ and $S_{MIN}(10, 3)$ here. Visual inspection of the full sample of the US house price to rent ratio (which would, of course, not have been possible were we doing this in real time) shows that the crash date indicated by $S_{MIN}(10, 1)$ corresponds to the first observation after the explosive bubble where the ratio begins to decline before it substantially decreases throughout 2007 and 2008, therefore suggesting that the monitoring procedure has worked very well. In October 2007, over a




FIGURE 9. US house price to rent ratio and $S_{MIN}(m, n)$ test statistics, 1975–2021

year after the detection of a crash by $S_{MIN}(10, 1)$, Treasury Secretary Henry Paulson stated that the continuing decline in house prices marked the ‘most significant current risk to [the US] economy’ (U.S. Department of the Treasury, 2007). Our application demonstrates that our proposed monitoring procedures can be used to detect bubble and crash behaviour in macroeconomic or financial data in real time, and this in turn can allow policy makers to respond quickly to such events.

VIII. Conclusion

In this paper, we have developed a real-time monitoring procedure for detecting a crash episode in a time series, conditional on having first detected a bubble regime. Our proposed procedure makes use of a training period, over which no bubble or crash occurs, to calibrate critical values, and then proceeds to monitor for significant evidence of a crash in a real-time environment as new data emerges. The new statistic we use for crash detection is based on an autoregressive modelling framework, with bubble and crash regimes modelled by explosive and stationary autoregressive dynamics, respectively. The statistic exploits the different signs of the means of the first differences associated with explosive and stationary regimes. A user-chosen parameter allows practitioners to trade off the speed of possible crash detection with the probability of spurious crash detection during a bubble regime. Asymptotic results establish that our procedure has desirable properties in terms of its ability to rapidly indicate the onset of a crash. Our Monte Carlo simulations suggest that, in finite samples, the recommended crash monitoring procedure has a well-controlled FPR during a bubble phase, while also allowing rapid detection of

a crash when one occurs. We have also considered how our procedure can be extended to a multiple bubble and crash environment, and demonstrate through simulation that the extended procedure performs well in this more general context. An application to the US housing market demonstrated the efficacy of our procedure in rapidly detecting the housing price crash which occurred in 2006.

Appendix A: Proof of Theorem 1

We first note that $A_{e,k}$ is a measurable function of a *finite* number of observations on Δy_t , and that under $H_0^{(1)}$, $\Delta y_t = \varepsilon_t$. It then follows that $\{A_{e,k}\}$ is a strictly stationary sequence and, from lemma 2.1 of White and Domowitz (1984), it is mixing of the same size as $\{\varepsilon_t\}$. Theorem 1 assumes that the mixing conditions of $\{\varepsilon_t\}$, and hence $\{A_{e,k}\}$, satisfy the mixing (long range dependence) conditions of Ferreira and Scotto (2002) (see Definition on p. 476). Hence the conditions underpinning the result in theorem 2.1 of Ferreira and Scotto (2002) are satisfied for the sequence $\{A_{e,k}\}$. Theorem 2.1 of Ferreira and Scotto (2002) for the case $r = s = 1$ (in their notation) then implies that, for two disjoint subintervals $I_{T,a}$ and $I_{T,b}$ of $[1, T]$, with respective lengths a_T and b_T such that $a_T/T \rightarrow a$ and $b_T/T \rightarrow b$,

$$\lim_{T \rightarrow \infty} P\left(\max_{e \in I_{T,a}} A_{e,k} \leq \max_{e \in I_{T,b}} A_{e,k}\right) = \frac{b}{a+b},$$

or

$$\lim_{T \rightarrow \infty} P\left(\max_{e \in I_{T,a}} A_{e,k} > \max_{e \in I_{T,b}} A_{e,k}\right) = \frac{a}{a+b}.$$

Setting $I_{T,a} = [T^* + k, T']$ and $I_{T,b} = [k + 1, T^*]$ we obtain

$$\begin{aligned} \lim_{T \rightarrow \infty} P\left(\max_{e \in [T^*+k, T']} A_{e,k} > \max_{e \in [k+1, T^*]} A_{e,k}\right) &= \lim_{T \rightarrow \infty} \left(\frac{T' - T^* - k + 1}{T' - 2k + 1}\right) \\ &= \lim_{T \rightarrow \infty} \left(\frac{T' - T^*}{T'}\right) \\ &= \alpha. \end{aligned}$$

Appendix B: Proof of Theorem 2

First we define $D_T = \lfloor \tau_1 T \rfloor (1 + \delta_1)^{2(\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor)}$. In the proof we will make use of the following results:

- (i) For any process x_t , an application of the Cauchy–Schwarz inequality shows that

$$\frac{(\sum_{t=a}^b x_t)^2}{\sum_{t=a}^b x_t^2} \leq b - a,$$

and hence, with finite $b - a$,

$$\frac{\sum_{t=a}^b x_t}{\sqrt{\sum_{t=a}^b x_t^2}} = O_p(1). \tag{B1}$$

(ii) For an AR(1) process x_s with generic parameter $1 + \delta$ and innovation ε_s , for $s, k > 0$, we can write

$$x_{s+k} = (1 + \delta)^k x_s + \sum_{j=0}^{s+k-1} (1 + \delta)^j \varepsilon_{s+k-j}, \tag{B2}$$

$$x_{s-k} = (1 + \delta)^{-k} x_s - \sum_{j=0}^{k-1} (1 + \delta)^{-k+j} \varepsilon_{s-j}. \tag{B3}$$

We next examine the behaviour of u_t in the region of $t = \lfloor \tau_2 T \rfloor$. First, for $t = \lfloor \tau_2 T \rfloor$ we use (B2) with $s + k = \lfloor \tau_2 T \rfloor$ and $k = \lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor$ to give

$$\begin{aligned} u_{\lfloor \tau_2 T \rfloor} &= (1 + \delta_1)^{\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor} u_{\lfloor \tau_1 T \rfloor} + \sum_{j=0}^{\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor - 1} (1 + \delta_1)^j \varepsilon_{\lfloor \tau_2 T \rfloor - j}, \\ D_T^{-1/2} u_{\lfloor \tau_2 T \rfloor} &= D_T^{-1/2} (1 + \delta_1)^{\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor} u_{\lfloor \tau_1 T \rfloor} + O_p(T^{-1/2}) \\ &= O_p^+(1), \end{aligned}$$

since $u_{\lfloor \tau_1 T \rfloor} = O_p^+(T^{1/2})$. Hence, $u_{\lfloor \tau_2 T \rfloor} = O_p^+(D_T^{1/2})$. Second, consider $t < \lfloor \tau_2 T \rfloor$. Using (B3) we have

$$\begin{aligned} u_{\lfloor \tau_2 T \rfloor - k} &= (1 + \delta_1)^{-k} u_{\lfloor \tau_2 T \rfloor} + O_p(1) \\ &= (1 + \delta_1)^{-k} O_p^+(D_T^{1/2}). \end{aligned}$$

Finally, for $t > \lfloor \tau_2 T \rfloor$, using (B2) we find

$$\begin{aligned} u_{\lfloor \tau_2 T \rfloor + k} &= (1 - \delta_2)^k u_{\lfloor \tau_2 T \rfloor} + O_p(1) \\ &= (1 + \delta_2)^k O_p^+(D_T^{1/2}). \end{aligned}$$

Now consider the behaviour of $S_{e,m,n}$ in the region of $t = \lfloor \tau_2 T \rfloor$, for the three cases of theorem 2.

(a) Consider the case of $j = 0$, i.e. $e = \lfloor \tau_2 T \rfloor$. Then

$$\begin{aligned} \sum_{t=e-n-m+1}^{e-n} \Delta y_t &= \sum_{t=\lfloor \tau_2 T \rfloor - n - m + 1}^{\lfloor \tau_2 T \rfloor - n} \Delta y_t \\ &= u_{\lfloor \tau_2 T \rfloor - n} - u_{\lfloor \tau_2 T \rfloor - n - m} \\ &= ((1 + \delta_1)^{-n} - (1 + \delta_1)^{-n-m}) u_{\lfloor \tau_2 T \rfloor} + O_p(1), \end{aligned}$$

using (B3). Here $(1 + \delta_1)^{-n} - (1 + \delta_1)^{-n-m} > 0$ so

$$\sum_{t=e-n-m+1}^{e-n} \Delta y_t = O_p^+(D_T^{1/2}).$$

Similarly,

$$\begin{aligned} \sum_{t=e-n+1}^e \Delta y_t &= \sum_{t=\lfloor \tau_2 T \rfloor - n + 1}^{\lfloor \tau_2 T \rfloor} \Delta y_t \\ &= u_{\lfloor \tau_2 T \rfloor} - u_{\lfloor \tau_2 T \rfloor - n} \\ &= (1 - (1 + \delta_1)^{-n})u_{\lfloor \tau_2 T \rfloor} + O_p(1). \end{aligned}$$

Here $1 - (1 + \delta_1)^{-n} > 0$ so

$$\sum_{t=e-n+1}^e \Delta y_t = O_p^+(D_T^{1/2}).$$

Next, since the data is generated according to $\Delta y_t = \delta_1(y_{t-1} - \mu) + \varepsilon_t = -\delta_1\mu + \delta_1 y_{t-1} + \varepsilon_t$, a linear regression of Δy_t on a constant and y_{t-1} represents a correctly specified estimating model. Standard OLS results then imply that $\sum_{t=\lfloor \tau_2 T \rfloor - n - m + 1}^{\lfloor \tau_2 T \rfloor - n} \hat{\varepsilon}_t^2 = \sum_{t=\lfloor \tau_2 T \rfloor - n - m + 1}^{\lfloor \tau_2 T \rfloor - n} \varepsilon_t^2 + O_p(1) = O_p(1)$. Hence we find

$$\begin{aligned} S_{e,m,n} &= \frac{\sum_{t=e-n-m+1}^{e-n} \Delta y_t}{\sqrt{\sum_{t=e-n-m+1}^{e-n} \hat{\varepsilon}_t^2}} \times \frac{\sum_{t=e-n+1}^e \Delta y_t}{\sqrt{\sum_{t=e-n+1}^e \Delta y_t^2}} \\ &= \frac{O_p^+(D_T^{1/2})}{\sqrt{O_p(1)}} \times O_p^+(1) \\ &= O_p^+(D_T^{1/2}), \end{aligned}$$

using (B1). Therefore we see that $S_{e,m,n}$ is diverging to $+\infty$ at the rate $D_T^{1/2}$. For $j = 1, 2, \dots, c$, the Δy_t involved in $S_{e,m,n}$ have the same order properties as when $j = 0$. Hence it also follows that $S_{e,m,n} = O_p^+(D_T^{1/2})$ for $e = \lfloor \tau_2 T \rfloor - 1, \lfloor \tau_2 T \rfloor - 2, \dots, \lfloor \tau_2 T \rfloor - c$. The result of Theorem 2(a) then follows since S_{\min}^* is the minimum of a sequence of $O_p(1)$ variates (given that $S_{e,m,n} = O_p(1)$ over the unit root training period), and therefore $S_{\min}^* = o_p(D_T^{1/2})$.

(b) Consider $e = \lfloor \tau_2 T \rfloor + j$ with $j = 1, \dots, n - 1$. It remains true that

$$\begin{aligned} \sum_{t=e-n-m+1}^{e-n} \Delta y_t &= O_p^+(D_T^{1/2}), \\ \sum_{t=e-n-m+1}^{e-n} \hat{\varepsilon}_t^2 &= O_p(1). \end{aligned}$$

Now,

$$\begin{aligned} \sum_{t=e-n+1}^e \Delta y_t &= \sum_{t=\lfloor \tau_2 T \rfloor + j - n + 1}^{\lfloor \tau_2 T \rfloor + j} \Delta y_t \\ &= u_{\lfloor \tau_2 T \rfloor + j} - u_{\lfloor \tau_2 T \rfloor + j - n} \\ &= (1 - \delta_2)^j u_{\lfloor \tau_2 T \rfloor} - (1 + \delta_1)^{j-n} u_{\lfloor \tau_2 T \rfloor} + O_p(1) \\ &= ((1 - \delta_2)^j - (1 + \delta_1)^{j-n}) u_{\lfloor \tau_2 T \rfloor} + O_p(1), \end{aligned}$$

so

$$\sum_{t=e-n+1}^e \Delta y_t = \begin{cases} ((1 - \delta_2)^j - (1 + \delta_1)^{j-n}) O_p^+(D_T^{1/2}) & (1 - \delta_2)^j \neq (1 + \delta_1)^{j-n} \\ O_p(1) & (1 - \delta_2)^j = (1 + \delta_1)^{j-n} \end{cases}$$

and therefore

$$\begin{aligned} S_{e,m,n} &= \frac{O_p^+(D_T^{1/2})}{\sqrt{O_p(1)}} \times \begin{cases} O_p^+(1) & (1 - \delta_2)^j > (1 + \delta_1)^{j-n} \\ O_p(1) & (1 - \delta_2)^j = (1 + \delta_1)^{j-n} \\ O_p^-(1) & (1 - \delta_2)^j < (1 + \delta_1)^{j-n} \end{cases} \\ &= \begin{cases} O_p^+(D_T^{1/2}) & (1 - \delta_2)^j > (1 + \delta_1)^{j-n} \\ O_p(D_T^{1/2}) & (1 - \delta_2)^j = (1 + \delta_1)^{j-n} \\ O_p^-(D_T^{1/2}) & (1 - \delta_2)^j < (1 + \delta_1)^{j-n} \end{cases}. \end{aligned}$$

Given $S_{\min}^* = o_p(D_T^{1/2})$, the result of Theorem 2(b) then follows directly.

(c) Consider $e = \lfloor \tau_2 T \rfloor + n$. Once again,

$$\begin{aligned} \sum_{t=e-n-m+1}^{e-n} \Delta y_t &= O_p^+(D_T^{1/2}), \\ \sum_{t=e-n-m+1}^{e-n} \hat{\varepsilon}_t^2 &= O_p(1). \end{aligned}$$

Now,

$$\begin{aligned} \sum_{t=e-n+1}^e \Delta y_t &= \sum_{t=\lfloor \tau_2 T \rfloor + 1}^{\lfloor \tau_2 T \rfloor + n} \Delta y_t \\ &= u_{\lfloor \tau_2 T \rfloor + n} - u_{\lfloor \tau_2 T \rfloor} \\ &= (1 - \delta_2)^n u_{\lfloor \tau_2 T \rfloor} - u_{\lfloor \tau_2 T \rfloor} + O_p(1) \\ &= ((1 - \delta_2)^n - 1) u_{\lfloor \tau_2 T \rfloor} + O_p(1). \end{aligned}$$

Here, $(1 - \delta_2)^n - 1 < 0$ so

$$\begin{aligned} S_{e,m,n} &= \frac{O_p^+(D_T^{1/2})}{\sqrt{O_p(1)}} \times O_p^-(1) \\ &= O_p^-(D_T^{1/2}) \end{aligned}$$

and the result in Theorem 2(c) follows.

Final Manuscript Received: July 2021

References

- Andrews, D. W. K. (2003). 'End-of-sample instability tests', *Econometrica*, Vol. 71, pp. 1661–1694.
- Andrews, D. W. K. and Kim, J.-Y. (2006). 'Tests for cointegration breakdown over a short time period', *Journal of Business and Economic Statistics*, Vol. 24, pp. 379–394.
- Anundsen, A. K. (2019). 'Detecting imbalances in house prices: what goes up must come down?', *The Scandinavian Journal of Economics*, Vol. 121, pp. 1587–1619.
- Anundsen, A. K., Gerdrup, K., Hansen, F. and Kragh-Sorensen, K. (2016). 'Bubbles and crises: the role of house prices and credit', *Journal of Applied Econometrics*, Vol. 31, pp. 1291–1311.
- Astill, S., Harvey, D. I., Leybourne, S. J. and Taylor, A. M. R. (2017). 'Tests for an end-of-sample bubble in financial time series', *Econometric Reviews*, Vol. 36, pp. 651–666.
- Astill, S., Harvey, D. I., Leybourne, S. J., Sollis, R. and Taylor, A. M. R. (2018). 'Real-time monitoring for explosive financial bubbles', *Journal of Time Series Analysis*, Vol. 39, pp. 863–891.
- Basse, T., Klein, T., Vigne, S. A. and Wegener, C. (2021). 'U.S. stock prices and the dot-com bubble: can dividend policy rescue the efficient market hypothesis?', *Journal of Corporate Finance*, Vol. 67, 101892.
- Caspi, I. and Graham, M. (2018). 'Testing for bubbles in stock markets with irregular dividend distribution', *Finance Research Letters*, Vol. 26, pp. 89–94.
- Corbet, S., Lucey, B. and Yarovaya, L. (2018). 'Datestamping the Bitcoin and Ethereum bubbles', *Finance Research Letters*, Vol. 26, pp. 81–88.
- Etienne, X. L., Irwin, S. H. and Garcia, P. (2014). 'Bubbles in food commodity markets: four decades of evidence', *Journal of International Money and Finance*, Vol. 42, pp. 129–155.
- Etienne, X. L., Irwin, S. H. and Garcia, P. (2015). 'Price explosiveness, speculation, and grain futures prices', *American Journal of Agricultural Economics*, Vol. 97, pp. 65–87.
- Fabozzi, F. J., Kynigakis, I., Panopoulou, E. and Tunaru, R. S. (2020). 'Detecting bubbles in the US and UK real estate markets', *The Journal of Real Estate Finance and Economics*, Vol. 60, pp. 469–513.
- Ferreira, H. and Scotto, M. (2002). 'On the asymptotic location of high values of a stationary sequence', *Statistics and Probability Letters*, Vol. 60, pp. 475–482.
- Figuerola-Ferretti, I. and McCrorie, J. (2016). 'The shine of precious metals around the global financial crisis', *Journal of Empirical Finance*, Vol. 38, pp. 717–738.
- Gronwald, M. (2021). 'How explosive are cryptocurrency prices?', *Finance Research Letters*, Vol. 38, 101603.
- Harvey, D. I., Leybourne, S. J., Sollis, R. and Taylor, A. R. (2016). 'Tests for explosive financial bubbles in the presence of non-stationary volatility', *Journal of Empirical Finance*, Vol. 38, pp. 548–574.
- Harvey, D. I., Leybourne, S. J. and Sollis, R. (2017). 'Improving the accuracy of asset price bubble start and end date estimators', *Journal of Empirical Finance*, Vol. 40, pp. 121–138.
- Harvey, D., Leybourne, S. and Whitehouse, E. (2020). 'Date-stamping multiple bubble regimes', *Journal of Empirical Finance*, Vol. 58, pp. 226–246.
- Homm, U. and Breitung, J. (2012). 'Testing for speculative bubbles in stock markets: a comparison of alternative methods', *Journal of Financial Econometrics*, Vol. 10, pp. 198–231.
- Hu, Y. and Oxley, L. (2018). 'Do 18th century 'bubbles' survive the scrutiny of 21st century time series econometrics?', *Economics Letters*, Vol. 162, pp. 131–134.

- Monetary Policy and the Economic Outlook. (2002). *Hearing before the Joint Economic Committee, Congress of the United States (testimony of Alan Greenspan)*. 107th Cong, April 17.
- OECD. (2021). *Housing Prices (indicator)*. <https://doi.org/10.1787/63008438-en>.
- Pavlidis, E., Yusupova, A., Paya, I., Peel, D., Martinez-Garcia, E., Mack, A. and Grossman, A. (2016). 'Episodes of exuberance in housing markets: in search of the smoking gun', *Journal of Real Estate Finance and Economics*, Vol. 53, pp. 419–449.
- Pavlidis, E. G., Paya, I. and Peel, D. A. (2018). 'Using market expectations to test for speculative bubbles in the crude oil market', *Journal of Money, Credit and Banking*, Vol. 50, pp. 833–856.
- Phillips, P. C. B. and Shi, S. (2018). 'Financial bubble implosion and reverse regression', *Econometric Theory*, Vol. 34, pp. 705–753.
- Phillips, P. C. B. and Shi, S. (2020). 'Real time monitoring of asset markets: Bubbles and crises', in Vinod H. D. and Rao C. (eds), *Handbook of Statistics*, Vol. 42, Elsevier, Amsterdam, pp. 61–80.
- Phillips, P. C. B., Wu, Y. and Yu, J. (2011). 'Explosive behavior in the 1990s NASDAQ: when did exuberance escalate asset values?', *International Economic Review*, Vol. 52, pp. 201–226.
- Phillips, P. C. B., Shi, S. and Yu, J. (2015). 'Testing for multiple bubbles: historical episodes of exuberance and collapse in the S&P 500', *International Economic Review*, Vol. 56, pp. 1043–1078.
- U.S. Department of the Treasury. (2007). *Remarks by Secretary Henry M. Paulson, Jr. on Current Housing and Mortgage Market Developments*, Georgetown University Law Center. <https://home.treasury.gov/news/press-releases/hp612>.
- White, H. and Domowitz, I. (1984). 'Nonlinear regression with dependent observations', *Econometrica*, Vol. 52, pp. 143–162.
- Whitehouse, E. J. (2019). 'Explosive asset price bubble detection with unknown bubble length and initial condition', *Oxford Bulletin of Economics and Statistics*, Vol. 81, pp. 20–41.