Real-time monitoring of bubbles and crashes

Whitehouse, E. J., Harvey, D. I., Leybourne, S. J.

Real-time monitoring of bubbles and crashes[*]

Whitehouse, E. J.[a], Harvey, D. I.[b] and Leybourne, S. J.[b]

[a] Department of Economics, University of Sheffield
[b] School of Economics, University of Nottingham

April 6, 2022

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 the bubble and crash regimes modelled by explosive and stationary dynamics respectively. The first stage of our approach is to monitor for the presence of a bubble; conditional on having detected a bubble, 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, over which no bubble or crash is assumed to occur. Monte Carlo simulations suggest that our recommended procedure has a well-controlled false positive rate during a bubble regime, while also allowing very rapid detection of a crash when one occurs. Application to the US housing market demonstrates the efficacy of our procedure in rapidly detecting the house price crash of 2006.

Keywords: Real-time monitoring; Bubble; Crash; Explosive autoregression; Stationary autoregression.

JEL Classification: C12, C22, G01.

---

[*]Correspondence to: Emily Whitehouse, Department of Economics, University of Sheffield, 9 Mappin Street, Sheffield, S1 4DT, UK; email: e.whitehouse@sheffield.ac.uk

1 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.

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 et al. (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 et al. (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; Anundsen, 2019; Pavlidis et al., 2018), commodities (Etienne et al., 2014, 2015; Figuerola-Ferretti and McCrorie, 2016) and cryptocurrencies (Corbet et al., 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 Astill et al. (2018) 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 et al. (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 introduce a degree of detection delay in order to trade off speed of detection versus reducing the FPR of crash detection. We follow Astill et al. (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 Astill et al. (2017) for bubble detection, such that we begin monitoring for a crash conditional on first detecting an explosive bubble. Monte Carlo simulations demonstrate that our recommended crash monitoring procedure offers 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 3 describes the Astill et al. (2017) real-time monitoring procedure for bubbles. In Section 4 we outline our crash monitoring procedure. Monte Carlo simulation results are presented in 5. Section 6 discusses the extension of our proposed procedure to multiple bubble and crash episodes. In Section 7 we provide an application of our real-time monitoring procedure to the US housing market. Section 8 concludes.

2 The model and real-time monitoring framework

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

$$y_t = \mu + u_t \tag{1}$$

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

with $u_0 = \psi$ where ψ is a finite positive constant, $\delta_1 > 0$ and $\delta_2 < 0$, and where $\lfloor . \rfloor$ denotes the integer part of the argument. We assume the error term ε_t is a strictly stationary, possibly conditionally heteroskedastic, process with zero mean.

In the context of (1)-(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, providing a model of 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, 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. We therefore wish to distinguish between the following hypotheses:

$$\begin{array}{lll} H_0 : \tau_1 = 1 & \text{(unit root)} \\ H_{1,1} : \tau_1 < 1 & \text{(unit root then explosive)} \\ H_{1,2} : \tau_1 < \tau_2 < 1 & \text{(unit root then explosive then stationary collapse)} \end{array}$$

The real-time monitoring framework we adopt follows AHLST and considers $y_1, ...,$ 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

---

[1]An 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.

data in a calibration role. The first stage involves monitoring for a change from H_0 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 $H_{1,1}$ to $H_{1,2}$, in order to detect the termination of explosive behaviour and the start of a stationary collapse regime.

3 Monitoring for an explosive regime

Several approaches to testing between H_0 and $H_{1,1}$, i.e. testing for a period of temporary explosiveness, have been developed in the literature. Whilst 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

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

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 $t = 1, ..., T^*$ 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 $t = T^\dagger - k + 1, ..., T^\dagger = T^* + 1, ..., T^* + k$, 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 $t = T^\dagger - k + 2, ..., T^\dagger + 1$). Detection of an explosive regime is triggered at the first point where a monitoring statistic $A_{e,k}$, $e = T^\dagger, T^\dagger + 1, ...$, exceeds the critical value $A_{\max}^*$. At an arbitrary point in the monitoring period, $t = T'$, we can then write the monitoring decision rule as:

$$\text{Detect } H_{1,1} \text{ at time } t = T' \text{ if } A_{T',k} > A_{\max}^*. \tag{3}$$

4

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 $A_{MAX}(k)$.

We now establish the theoretical FPR of the $A_{MAX}(k)$ procedure under H_0 as $T \to \infty$, 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 $A_{e,k}$, $e = T^\dagger, T^\dagger + 1, ..., T'$, exceeds the corresponding maximum statistic over the training period $A^*_{\max}$. Evaluating the limiting probabilities of these exceedances under H_0 gives the result of the following theorem.

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

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

where

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

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.:

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

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:

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

4 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 $H_{1,1}$ 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)-(2) expressed in first differences:

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

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

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

In each regime, the first term dominates the stochastic behaviour of Δy_t. For $t = \lfloor \tau_2 T \rfloor - m + 1, ..., \lfloor \tau_2 T \rfloor$, this follows since it can be shown that, for finite m, $\delta_1(1+\delta_1)^{t-(\lfloor \tau_2 T \rfloor - m)-1}u_{\lfloor \tau_2 T \rfloor - m} = O_p(S_T^{1/2})$, where $S_T = \lfloor \tau_1 T \rfloor (1+\delta_1)^{2(\lfloor \tau_2 T \rfloor - \lfloor \tau_1 T \rfloor)}$ (see Harvey et al. (2017) for details), and $\sum_{i=0}^{t-(\lfloor \tau_2 T \rfloor - m)-1}\delta_1(1+\delta_1)^{i-1}\varepsilon_{t-i} + \varepsilon_t = O_p(1)$. Similarly, for $t = \lfloor \tau_2 T \rfloor + 1, ..., \lfloor \tau_2 T \rfloor + n$ and finite n, $\delta_2(1+\delta_2)^{t-\lfloor \tau_2 T \rfloor - 1}u_{\lfloor \tau_2 T \rfloor} = O_p(S_T^{1/2})$ and $\sum_{i=1}^{t-\lfloor \tau_2 T \rfloor - 1}\delta_2(1+\delta_2)^{i-1}\varepsilon_{t-i} + \varepsilon_t = O_p(1)$. Furthermore, for small m and n, and δ_1 and δ_2 close to zero, we can note the following approximations:

$$
\begin{array}{ll}
\delta_1(1+\delta_1)^{t-(\lfloor \tau_2 T \rfloor - m)-1}u_{\lfloor \tau_2 T \rfloor - m} \approx \delta_1 u_{\lfloor \tau_2 T \rfloor - m} & t = \lfloor \tau_2 T \rfloor - m + 1, ..., \lfloor \tau_2 T \rfloor \\
\delta_2(1+\delta_2)^{t-\lfloor \tau_2 T \rfloor - 1}u_{\lfloor \tau_2 T \rfloor} \approx \delta_2 u_{\lfloor \tau_2 T \rfloor} & t = \lfloor \tau_2 T \rfloor + 1, ..., \lfloor \tau_2 T \rfloor + n
\end{array} . \quad (6)
$$

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

$$
\Delta y_t = \begin{cases}
\beta_1 + \eta_t & t = \lfloor \tau_2 T \rfloor - m + 1, ..., \lfloor \tau_2 T \rfloor \\
\beta_2 + \eta_t & t = \lfloor \tau_2 T \rfloor + 1, ..., \lfloor \tau_2 T \rfloor + n
\end{cases} \quad (7)
$$

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 (6) 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 (7), 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:

$$
S_{e,m,n} = \frac{\sum_{t=e-n-m+1}^{e-n}\Delta y_t \sum_{t=e-n+1}^{e}\Delta y_t}{\sqrt{\sum_{t=e-n-m+1}^{e-n}(\Delta y_t)^2 \sum_{t=e-n+1}^{e}(\Delta y_t)^2}}. \quad (8)
$$

The numerator of this statistic is essentially $\hat{\beta}_1\hat{\beta}_2$ with $e - n$ replacing $\lfloor \tau_2 T \rfloor$, and is therefore suitable for monitoring for a collapse that begins at time $t = e - n + 1$. Note

that the m^{-1} and n^{-1} constants in $\hat{\beta}_1$ and $\hat{\beta}_2$ 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. The denominator of (8) is a variance standardization, introduced to imbue the statistic with a degree of robustness to possible changes in unconditional variance; such changes are formally excluded from our assumptions but may of course occur in practice.

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 $t = 1, ..., T^*$ 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 $A_{MAX}(k)$ 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 $e = T^\diamond + 1, T^\diamond + 2, ...$ 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, $t = T'$, we can then write the monitoring decision rule as:

$$\text{Detect } H_{1,2} \text{ at time } t = T' \text{ if } S_{T',m,n} < S_{\min}^*. \tag{9}$$

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

The FPR of the $S_{MIN}(m, n)$ procedure is more difficult to establish compared to the $A_{MAX}(k)$ procedure for explosive regime detection, since monitoring using $S_{MIN}(m, n)$ is only performed following detection of an explosive regime by $A_{MAX}(k)$. When $S_{MIN}(m, n)$ is conducted and H_0 is true, it follows that the explosive regime detection signalled by $A_{MAX}(k)$ was false, and hence the FPR of $S_{MIN}(m, n)$ is bounded by the FPR of $A_{MAX}(k)$ at the point $t = T^\diamond$ (where explosive behaviour was erroneously detected). When $S_{MIN}(m, n)$ is conducted and $H_{1,1}$ is true, which is the main case of interest, the $A_{MAX}(k)$ explosive regime detection at time $t = T^\diamond$ was correct, and the FPR of the subsequently implemented $S_{MIN}(m, n)$ procedure is bounded by the true positive rate of $A_{MAX}(k)$ at time $t = T^\diamond$. Further theoretical analysis of the FPR of the $S_{MIN}(m, n)$ procedure in this case is not possible, since it will depend on nuisance parameters such as the magnitude of the explosive autoregressive parameter and the duration of the explosive regime. Instead we examine the FPR of $S_{MIN}(m, n)$ by simulation in the next section for a range of DGP settings.

In practice choices must be made for m and n. Setting $n = 1$ affords the opportunity of detecting a collapse most quickly, since (9) 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 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 $S_{MIN}(m, n)$ procedure is robust to conditional heteroskedasticity and serial correlation in ε_t, in line with the robustness properties of $A_{MAX}(k)$. 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

7

becomes asymptotically negligible.

5 Simulation results

To examine the finite sample performance of the $S_{MIN}(m, n)$ 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 $\varepsilon_t \sim IIDN(0, 1)$. We set $\psi = 100$ to ensure that, under $H_{1,1}$, we generate only positive explosive regimes. We evaluate the performance of $S_{MIN}(m, n)$ statistics computed for $m = \{5, 10, 15\}$ and $n = \{1, 2, 3\}$, in each case using $A_{MAX}(k)$ 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 $S_{MIN}(m, n)$ monitoring procedures when no stationary collapse regime is present. Figure 1 displays the cumulative rejection frequencies of the $A_{MAX}(k)$ and $S_{MIN}(m, n)$ procedures under no collapse scenarios. The cumulative rejection frequencies of $A_{MAX}(k)$ under H_0 and of $S_{MIN}(m, n)$ under H_0 and $H_{1,1}$ represent the empirical FPRs for the procedures. Figures 1(a)-1(c) consider the case of H_0 where there is neither an explosive bubble nor stationary collapse such that $\tau_1 = 1$, whilst Figures 1(d)-1(l) examine $H_{1,1}$ where an explosive regime occurs and continues until the end of the monitoring period. Beginning with the H_0 case, it is apparent from Figures 1(a)-1(c) that the empirical FPR of $A_{MAX}(k)$ tracks its theoretical FPR very closely, confirming the results of AHLST, whilst the $S_{MIN}(m, n)$ procedures are always somewhat lower than this theoretical FPR for all combinations of m and n considered here. As discussed in Section 4, under H_0 the FPR of the $S_{MIN}(m, n)$ procedures should be bounded by the theoretical FPR for $A_{MAX}(k)$ and our simulation results demonstrate that this holds in finite samples.

In Figures 1(d)-1(f), we now consider the case of an explosive regime which begins at $\lfloor \tau_1 T \rfloor = 210$ 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 $A_{MAX}(k)$ 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 $S_{MIN}(m, n)$ 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 Figures 1(g)-1(i) where we now consider a smaller magnitude explosive regime, setting $\delta_1 = 0.02$. In this case, the empirical FPR of the $S_{MIN}(m, 1)$ procedure increases slightly at the beginning of the explosive regime, before levelling off to a point below 0.08 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 $S_{MIN}(m, 2)$ and $S_{MIN}(m, 3)$ lie close to zero throughout the monitoring period. Finally, turning our attention to Figures 1(j)-1(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 $S_{MIN}(m, 1)$ 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. $S_{MIN}(m, 2)$ and $S_{MIN}(m, 3)$, however, display much lower empirical FPRs throughout the monitoring period, and never exceed 0.20 even under this small explosive setting. These latter procedures therefore offer a greater degree of FPR control relative

to $S_{MIN}(m,1)$.

It is clear that the empirical FPRs of the $S_{MIN}(m,n)$ procedures are clearly influenced by the magnitude of the explosive parameter. The closer the explosive offset is to zero, the higher the likelihood of negative changes in the series occurring during the explosive regime, and thus the higher the probability of rejection by $S_{MIN}(m,n)$. With reference to the statistics' motivation in Section 4, 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 $S_{MIN}(m,1)$ 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 $S_{MIN}(m,2)$ or $S_{MIN}(m,3)$, 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 much less sensitive to the magnitude of the explosive regime, and good FPR control is displayed for all m settings considered.

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 2-5 report the cumulative rejection frequencies for the $A_{MAX}(k)$ and $S_{MIN}(m,n)$ procedures under this hypothesis, as well as histograms of crash detection dates obtained by $S_{MIN}(m,n)$, to allow us to evaluate both the power and dating accuracy of our procedures.[2]

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. Due to the conditional testing approach we employ, the $S_{MIN}(m,n)$ rejection frequencies are bounded by those of $A_{MAX}(k)$. Examining first Figure 2, we note that the $S_{MIN}(m,n)$ procedure achieves rejection frequencies close to or the same as $A_{MAX}(k)$ across m and n, so that the collapse is detected in almost all cases where a bubble was found. In particular, for $m = 5$, by the end of the collapse regime, the $S_{MIN}(5,n)$ procedures have rejection frequencies between 0.09 and 0.15 lower than those of $A_{MAX}(5)$; for $m = 10$ and $m = 15$, these differences are no more than 0.06. Indeed, for $S_{MIN}(10,n)$ and $S_{MIN}(15,n)$, the procedures have rejection frequencies very close to 1 by the end of the collapse regime.

While the rejection frequencies for bubble detection by $A_{MAX}(k)$ are increasing over the duration of the explosive regime, one key feature of $S_{MIN}(m,n)$ is that the majority of rejections occur at one point in time. In the case of $S_{MIN}(m,1)$, this is the first observation of the collapse regime, whereas for $S_{MIN}(m,2)$ and $S_{MIN}(m,3)$ a one and two observation lag is introduced respectively. To examine this in more detail, Figures 2(d)-2(l) display histograms of the date at which the $S_{MIN}(m,n)$ procedures detected a crash in each Monte Carlo replication where a crash was detected during the monitoring period. For example, comparing Figures 2(d), 2(e) and 2(f), it is clear that when a rejection is found by the $S_{MIN}(5,1)$ procedure it is on the true crash date, $\lfloor \tau_2 T \rfloor + 1$, in almost all replications. This suggests that our procedure is highly capable of real-time

---

[2]As noted in Section 2, it is possible that when an explosive regime terminates it is followed by unit root behaviour rather than a collapse. Whilst such a scenario is less common in practice, and as such is not the focus of this paper, we note that in unreported simulations our $S_{MIN}(m,n)$ procedures can detect the termination of an explosive regime even when it is not followed by a collapse.

detection of crashes. For $S_{MIN}(5,2)$ almost all rejections occur at $\lfloor \tau_2 T \rfloor + 2$, i.e. one observation after the true crash date, and for $S_{MIN}(5,3)$ almost all rejections occur at at $\lfloor \tau_2 T \rfloor + 3$, two observations after the true crash date. Using $n = 1$ in $S_{MIN}(m,n)$ thus offers immediate detection of a crash, whilst using $n > 1$ can be seen as invoking an $n - 1$ observation delay in detection. Given that under $H_{1,1}$, 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 small 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 $S_{MIN}(m,n)$ range from 0.82 to 0.93 under these settings, therefore still maintaining excellent power to detect a crash. The histograms displayed in Figures 3(d)-3(l) show that the immediate detection of a crash by $S_{MIN}(m,1)$ and one or two observation lagged detection by $S_{MIN}(m,2)$ and $S_{MIN}(m,3)$ 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 5 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 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 $S_{MIN}(5,n)$ are relatively low, while $S_{MIN}(10,n)$ and $S_{MIN}(15,n)$ 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 $S_{MIN}(m,1)$ 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 Figures 1(j)-1(l) that we discussed previously. Crucially, such behaviour is not observed for $S_{MIN}(m,2)$ and $S_{MIN}(m,3)$, 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 $S_{MIN}(10,1)$ in Figure 5(g). Whilst 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 $S_{MIN}(10,2)$ in Figure 5(h), 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.

The results of this section demonstrate that our $S_{MIN}(m,n)$ 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, settings of $m = 10$ and $m = 15$ yielded similarly high rejection frequencies, with lower frequencies obtained for $m = 5$. We suggest that selecting $m = 10$ will be suitable for most scenarios. 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 $n = 1$ in our procedure 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 will be detected with a chosen (i.e. $n - 1$) delay. 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.

6   Monitoring for multiple bubble and crash regimes

The testing approach outlined in Sections 3-4 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 $j = 1, ..., N$ explosive bubble and stationary collapse regimes. For $t = 1, ..., T$:

$$y_t = \mu + x_t + u_t \tag{10}$$
$$u_t = (1 + \delta_t)u_{t-1} + \varepsilon_t + \psi\mathbb{I}(t = \lfloor \tau_{j,3}T \rfloor + 1) \tag{11}$$
$$\delta_t = \sum_{j=1}^{N} \{\delta_{j,1}\mathbb{I}(\lfloor \tau_{j,1}T \rfloor < t \leqslant \lfloor \tau_{j,2}T \rfloor) + \delta_{j,2}\mathbb{I}(\lfloor \tau_{j,2}T \rfloor < t \leqslant \lfloor \tau_{j,3}T \rfloor) - \mathbb{I}(t = \lfloor \tau_{j,3}T \rfloor + 1)\}$$
$$x_t = \sum_{j=1}^{N} (u_{\lfloor \tau_{j,3}T \rfloor} - \psi)\mathbb{I}(t > \lfloor \tau_{j,3}T \rfloor)$$

with, as before, $u_0 = \psi$ where ψ is a finite positive constant, $\delta_{j,1} > 0$ and $\delta_{j,2} < 0$ for all j. The inclusion of x_t in (10) and the indicator function term in (11) 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 et al. (2020), here allowing for non-zero ψ. 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 $j = 1, ..., N$.

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

$$H_0^{(j)} : \tau_{j,1} = 1 \qquad \text{(unit root)}$$
$$H_{1,1}^{(j)} : \tau_{j,1} < 1 \qquad \text{(unit root then explosive)}$$
$$H_{1,2}^{(j)} : \tau_{j,1} < \tau_{j,2} < 1 \quad \text{(unit root then explosive then stationary collapse)}$$

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 $H_0^{(1)}$ in favour of $H_{1,2}^{(1)}$, we should then switch back into monitoring for the second explosive bubble regime by considering $H_0^{(2)}$.

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 $A_{MAX}(k)$:

$$\text{Detect } H_{1,1}^{(j)} \text{ at time } t = T' \text{ if } A_{T',m} > A_{\max}^*.$$

If an explosive bubble is detected, we denote the detection date as $T_j^{\diamond}$. Next, stationary collapse monitoring is undertaken using $S_{MIN}(m, n)$:

$$\text{Detect } H_{1,2}^{(j)} \text{ at time } t = T' \text{ if } S_{T',m,n} < S_{\min}^*$$

where $T' \in [T_j^{\diamond} + 1, T]$. The date of stationary collapse detection is denoted $T_j^{\diamond\diamond}$. Monitoring for a subsequent explosive bubble then begins at $T_j^{\diamond\diamond} + k$ using $A_{MAX}(k)$:

$$\text{Detect } H_{1,1}^{(j+1)} \text{ at time } t = T' \text{ if } A_{T',m} > A_{\max}^*$$

where $T' \in [T_j^{\diamond\diamond} + k, T]$. Monitoring continues in this manner, switching between $A_{MAX}(k)$ and $S_{MIN}(m, n)$ for as long as desired.

We consider the performance of our proposed multiple bubble and crash monitoring procedure through Monte Carlo simulation of (10), where we again use $\varepsilon_t \sim IIDN(0, 1)$ and set $\mu = 0$ and $\psi = 100$. Given the finite sample performance of our monitoring procedure in the single bubble context displayed in Section 5, we provide results for the $A_{MAX}(10)$ and $S_{MIN}(10, 2)$ procedures here. We set $T^* = 200$.

Figure 6 displays the empirical FPRs of the $A_{MAX}(10)$ and $S_{MIN}(10, 2)$ procedures under H_0^j, 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 6 that the empirical FPRs of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for two bubbles/crashes remain small throughout the monitoring period, whilst 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 7 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 7(a) we set $\delta_{j,1} = 0.03$ and $\delta_{j,2} = 0.015$ for $j = 1, .., 3$, 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 second observation of the stationary collapse regime in most replications. In Figure 7(b) we consider smaller settings of $\delta_{j,1} = 0.02$ and $\delta_{j,2} = 0.01$ for $j = 1, .., 3$. 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.

7 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 United States 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*, Anundsen et al. (2016), Anundsen (2019), Pavlidis et al. (2016), 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. Whilst 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 1, 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 $t = 1, ..., T^*$.

Figure 8 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 two 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, $S_{MIN}(10, 2)$ in 2006:Q3, and $S_{MIN}(10, 3)$ in 2006:Q4. Increasing n in the $S_{MIN}(10, n)$ monitoring procedure has worked exactly as the simulations demonstrated, with $n > 1$ introducing a $n - 1$ observation delay in detection. 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 Oct 2007, over a 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.

8   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. Our Monte Carlo simulations suggest that 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 also 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.

References

Andrews, D. W. K. (2003), 'End-of-sample instability tests', *Econometrica* **71**, 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* **24**, 379–394.

Anundsen, A. K. (2019), 'Detecting imbalances in house prices: What goes up must come down?', *The Scandinavian Journal of Economics.* **121**, 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* **31**, 1291–1311.

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* **39**, 863–891.

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* **36**, 651–666.

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* **67**, 101892.

Caspi, I. and Graham, M. (2018), 'Testing for bubbles in stock markets with irregular dividend distribution', *Finance Research Letters* **26**, 89–94.

Corbet, S., Lucey, B. and Yarovaya, L. (2018), 'Datestamping the Bitcoin and Ethereum bubbles', *Finance Research Letters* **26**, 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* **42**, 129–155.

Etienne, X. L., Irwin, S. H. and Garcia, P. (2015), 'Price explosiveness, speculation, and grain futures prices', *American Journal of Agricultural Economics* **97**, 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* **60**, 469–513.

Ferreira, H. and Scotto, M. (2002), 'On the asymptotic location of high values of a stationary sequence', *Statistics and Probability Letters* **60**, 475–482.

Figuerola-Ferretti, I. and McCrorie, J. (2016), 'The shine of precious metals around the global financial crisis', *Journal of Empirical Finance* **38**, 717–738.

Gronwald, M. (2021), 'How explosive are cryptocurrency prices?', *Finance Research Letters* **38**, 101603.

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* **40**, 121–138.

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* **38**, 548–574.

Harvey, D., Leybourne, S. and Whitehouse, E. (2020), 'Date-stamping multiple bubble regimes', *Journal of Empirical Finance* **58**, 226–246.

Homm, U. and Breitung, J. (2012), 'Testing for speculative bubbles in stock markets: A comparison of alternative methods', *Journal of Financial Econometrics* **10**, 198–231.

Hu, Y. and Oxley, L. (2018), 'Do 18th century 'bubbles' survive the scrutiny of 21st century time series econometrics?', *Economics Letters* **162**, 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)'. doi: 10.1787/63008438-en.

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* **50**, 833–856.

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* **53**, 419–449.

Phillips, P. C. B. and Shi, S. (2018), 'Financial bubble implosion and reverse regression', *Econometric Theory* **34**, 705–753.

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* **56**, 1043–1078.

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* **52,1**, 201–226.

Phillips, P. C. and Shi, S. (2020), Real time monitoring of asset markets: Bubbles and crises, *in* H. D. Vinod and C. Rao, eds, 'Handbook of Statistics', Vol. 42, Elsevier, pp. 61–80.

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. October 16.

White, H. and Domowitz, I. (1984), 'Nonlinear regression with dependent observations', *Econometrica* **52**, 143–162.

Whitehouse, E. J. (2019), 'Explosive asset price bubble detection with unknown bubble length and initial condition', *Oxford Bulletin of Economics and Statistics* **81**, 20–41.

A   Proof of Theorem 1

We first note that $A_{e,m}$ is a measureable function of a *finite* number of observations on Δy_t, and that under H_0, $\Delta y_t = \varepsilon_t$. It then follows that $\{A_{e,m}\}$ is a strictly stationary sequence and, from Lemma 2.1 of White and Domowitz (1984), it is mixing of the same size as $\{\varepsilon_t\}$. These conditions then allow us to appeal to the result in Theorem 2.1 of Ferreira and Scotto (2002) which concerns the limiting probability that the rth ($r \geq 1$) largest value in one of two disjoint subsequences exceeds the sth ($s \geq 1$) largest value in the other subsequence. Setting $r = s = 1$ gives the limiting probability that the maximum of one disjoint subsequence exceeds the maximum of another as being equal to the limiting ratio of the length of the former subsequence to the total length of the two subintervals. This result can then be used to establish the result for $A_{e,m}$ in Theorem 1.

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

H_0: $\tau_1 = 1$
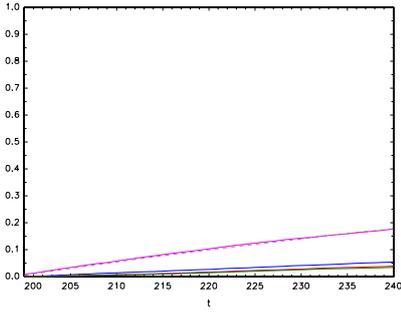
(a) $m = k = 5$
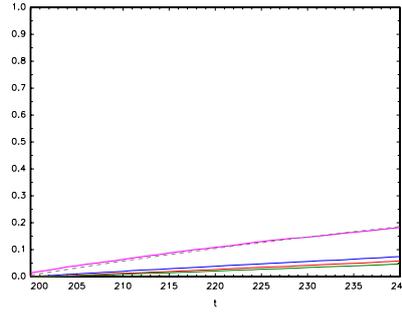
(b) $m = k = 10$
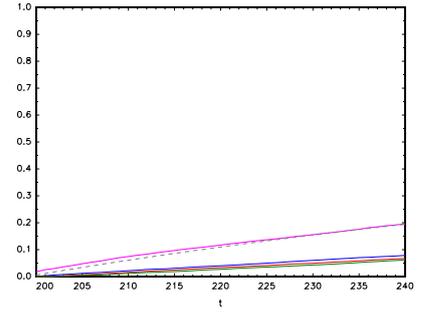
(c) $m = k = 15$

$H_{1,1}$: $\delta_1 = 0.03$, $\lfloor \tau_1 T \rfloor = 210$ and $\delta_2 = 0$, $\tau_2 = 1$
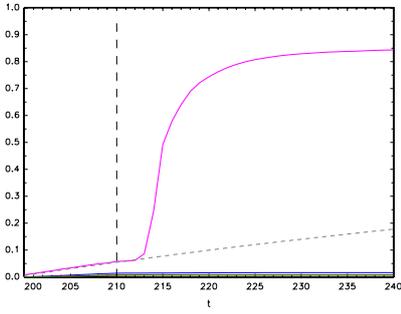
(d) $m = k = 5$
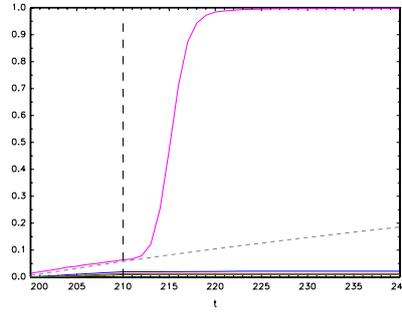
(e) $m = k = 10$
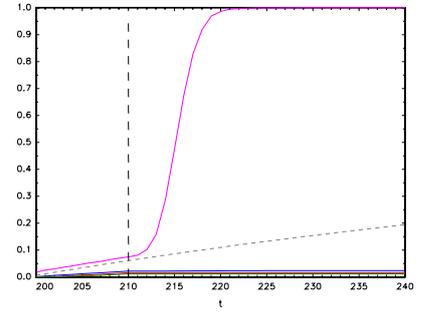
(f) $m = k = 15$

$H_{1,1}$: $\delta_1 = 0.02$, $\lfloor \tau_1 T \rfloor = 210$ and $\delta_2 = 0$, $\tau_2 = 1$
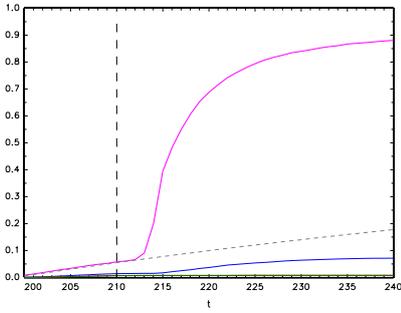
(g) $m = k = 5$
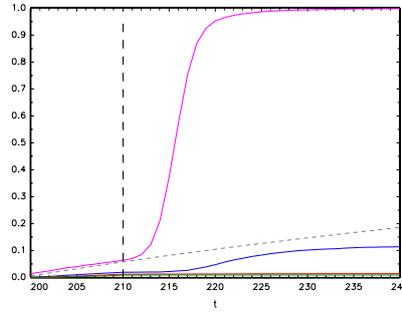
(h) $m = k = 10$
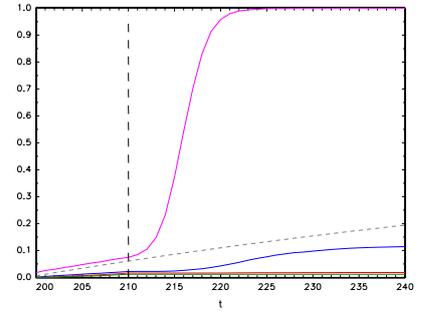
(i) $m = k = 15$

$H_{1,1}$: $\delta_1 = 0.01$, $\lfloor \tau_1 T \rfloor = 210$ and $\delta_2 = 0$, $\tau_2 = 1$
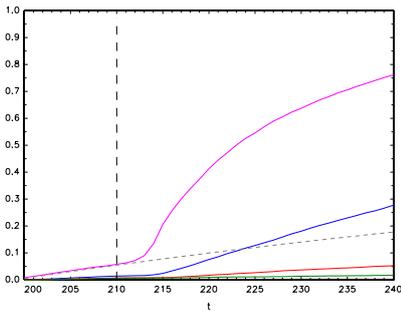
(j) $m = k = 5$
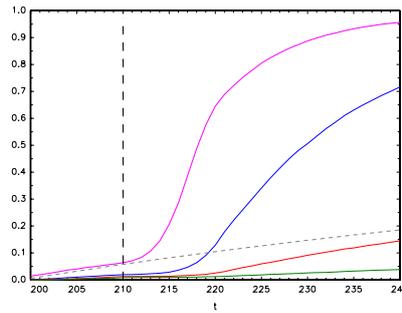
(k) $m = k = 10$
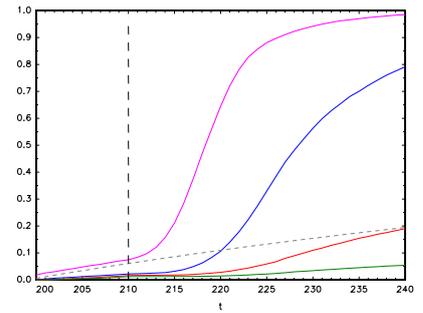
(l) $m = k = 15$

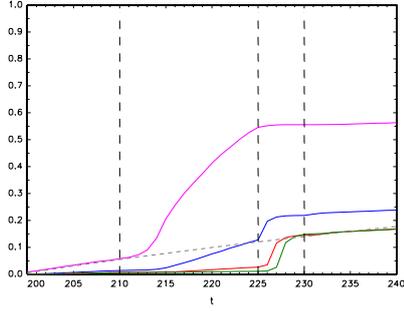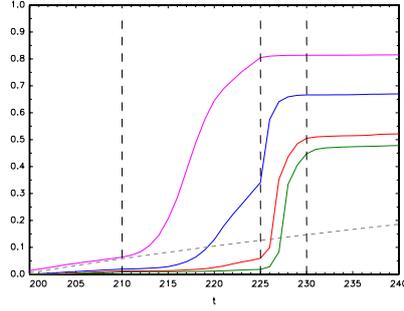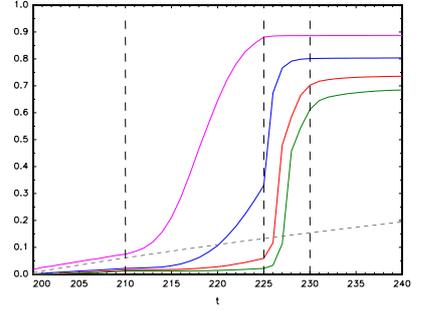$A_{MAX}(k)$, —— $S_{MIN}(m,1)$, —— $S_{MIN}(m,2)$, —— $S_{MIN}(m,3)$, - - FPR, — — $\lfloor \tau_1 T \rfloor$

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 $\delta_2 = 0.015$
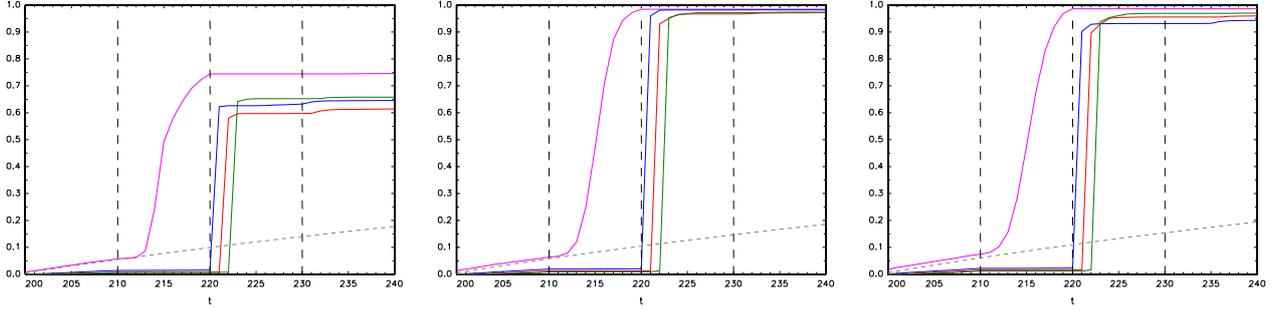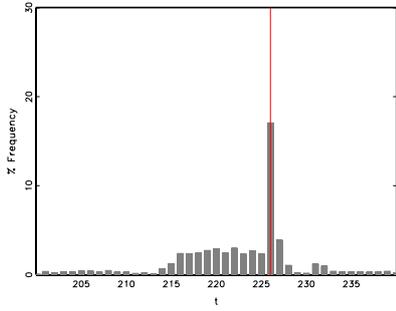

(a) $m = k = 5$ (b) $m = k = 10$ (c) $m = k = 15$
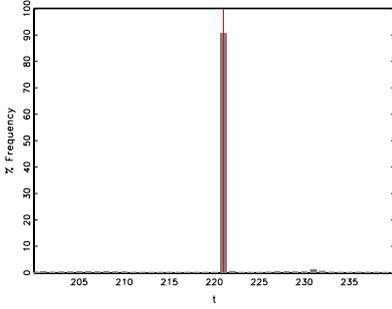
—— $A_{MAX}(k)$, —— $S_{MIN}(m,1)$, —— $S_{MIN}(m,2)$, —— $S_{MIN}(m,3)$, - - FPR, -- $\lfloor\tau_1 T\rfloor$, $\lfloor\tau_2 T\rfloor$, $\lfloor\tau_3 T\rfloor$
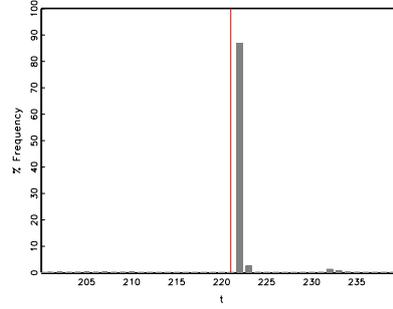
(d) $S_{MIN}(5,1)$ (e) $S_{MIN}(5,2)$ (f) $S_{MIN}(5,3)$
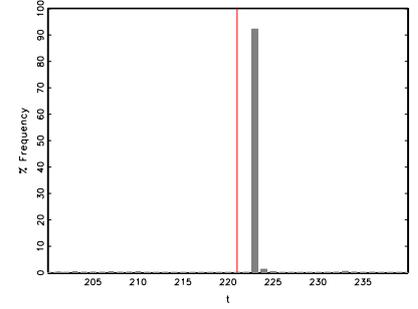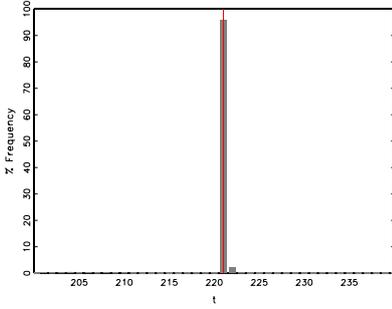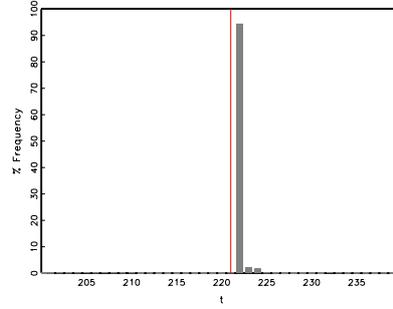
(g) $S_{MIN}(10,1)$ (h) $S_{MIN}(10,2)$ (i) $S_{MIN}(10,3)$
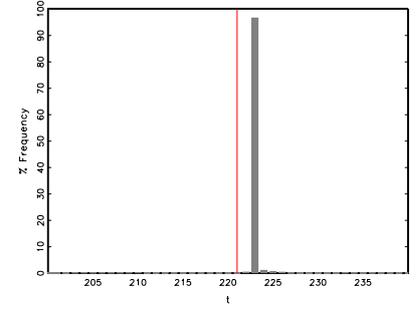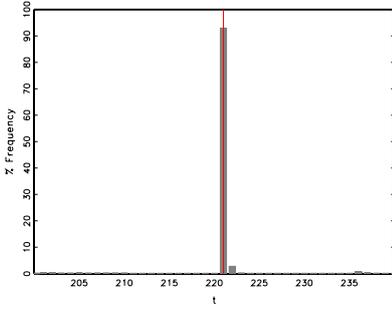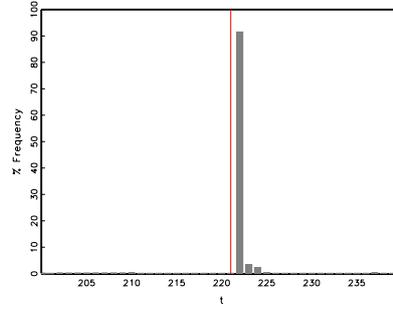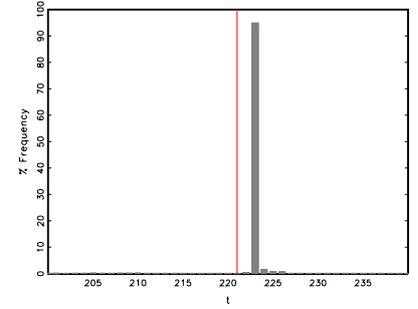
(j) $S_{MIN}(15,1)$ (k) $S_{MIN}(15,2)$ (l) $S_{MIN}(15,3)$

—— $\lfloor\tau_2 T\rfloor + 1$

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 $\delta_2 = 0.01$
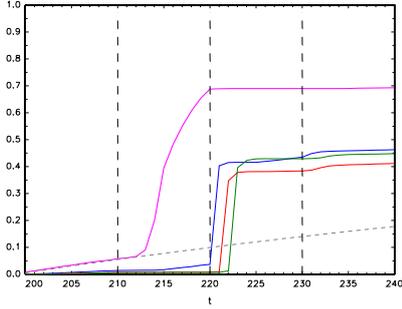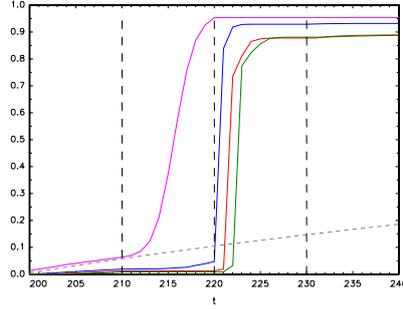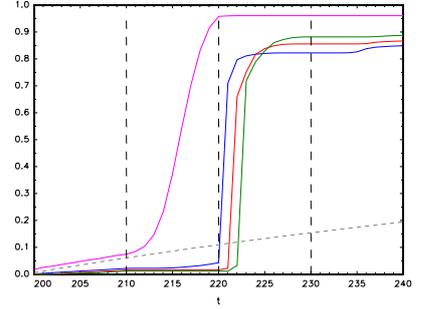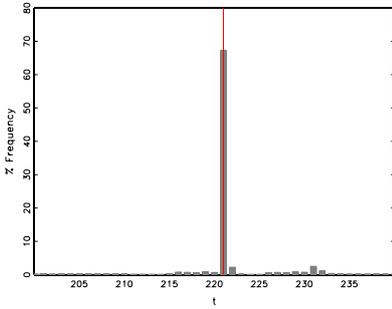
(a) $m = k = 5$ (b) $m = k = 10$ (c) $m = k = 15$

—— $A_{MAX}(k)$, —— $S_{MIN}(m,1)$, —— $S_{MIN}(m,2)$, —— $S_{MIN}(m,3)$, - - FPR, - - $\lfloor \tau_1 T \rfloor$, $\lfloor \tau_2 T \rfloor$, $\lfloor \tau_3 T \rfloor$
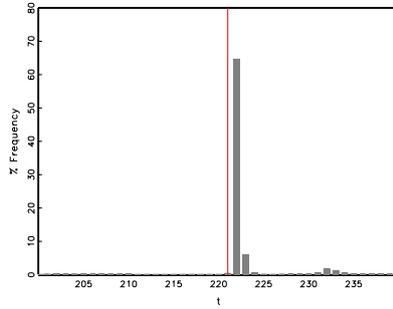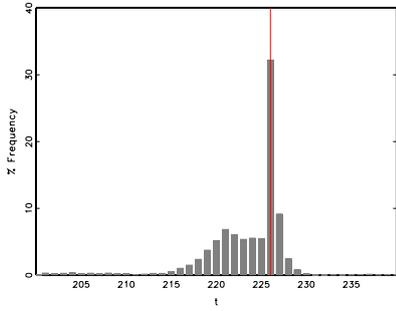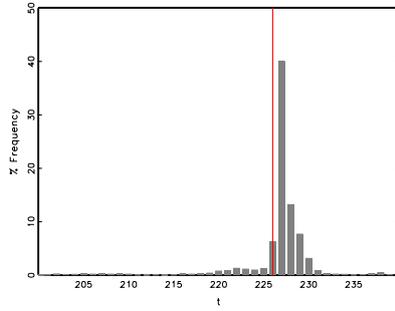
(d) $S_{MIN}(5,1)$ (e) $S_{MIN}(5,2)$ (f) $S_{MIN}(5,3)$
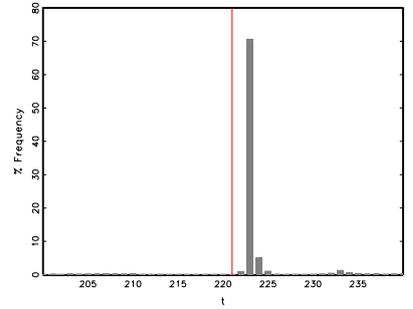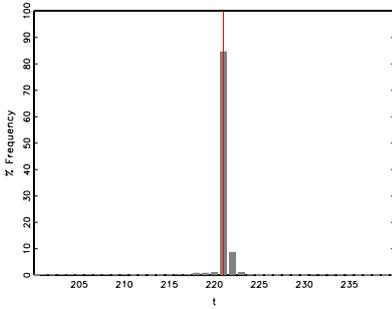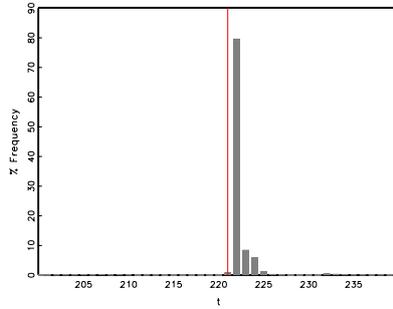
(g) $S_{MIN}(10,1)$ (h) $S_{MIN}(10,2)$ (i) $S_{MIN}(10,3)$
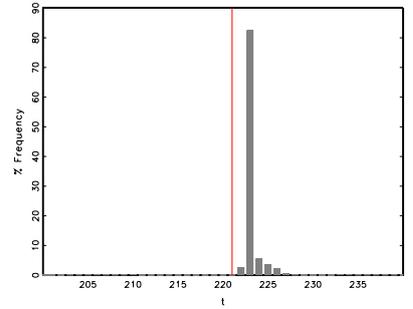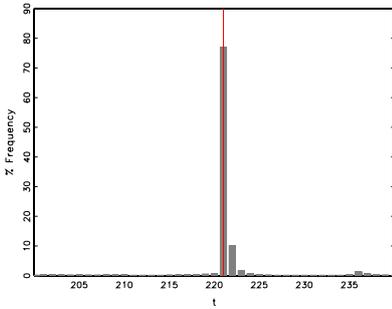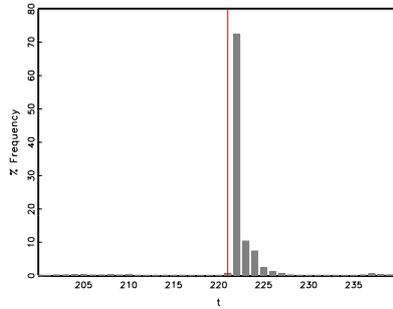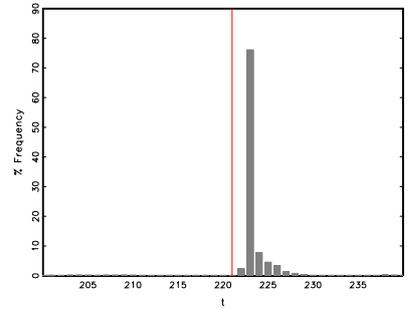
(j) $S_{MIN}(15,1)$ (k) $S_{MIN}(15,2)$ (l) $S_{MIN}(15,3)$

—— $\lfloor \tau_2 T \rfloor + 1$

Figure 4: Rejection frequencies of $A_{MAX}(k)$ and $S_{MIN}(m,n)$, and histograms of $S_{MIN}(m,n)$ detection dates: $\lfloor \tau_1 T \rfloor = 210$, $\lfloor \tau_2 T \rfloor = 225$, $\lfloor \tau_3 T \rfloor = 230$, $\delta_1 = 0.02$ and $\delta_2 = 0.01$
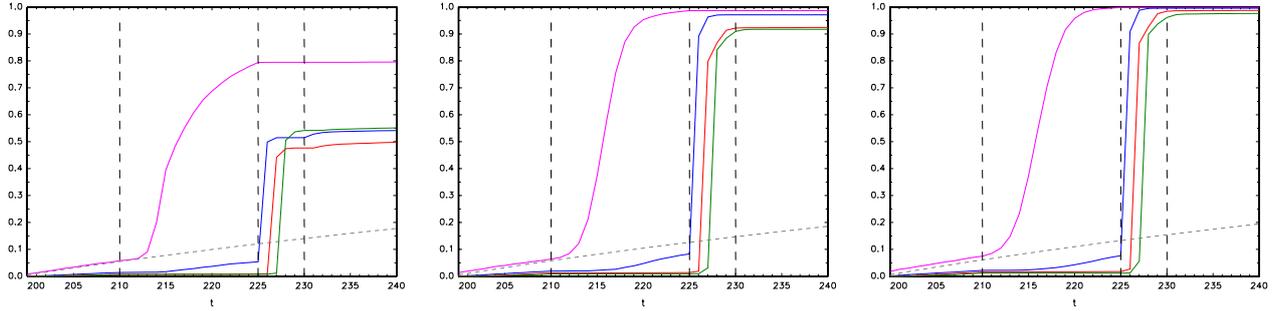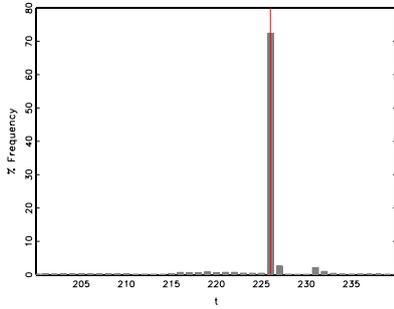
(a) $m = k = 5$ (b) $m = k = 10$ (c) $m = k = 15$

— $A_{MAX}(k)$, — $S_{MIN}(m,1)$, — $S_{MIN}(m,2)$, — $S_{MIN}(m,3)$, - - FPR, - - $\lfloor \tau_1 T \rfloor$, $\lfloor \tau_2 T \rfloor$, $\lfloor \tau_3 T \rfloor$
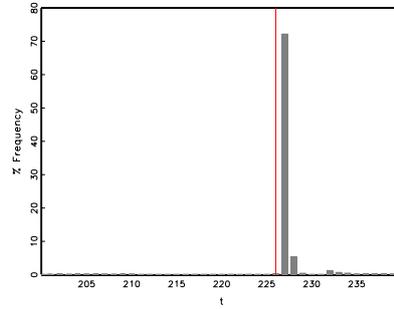
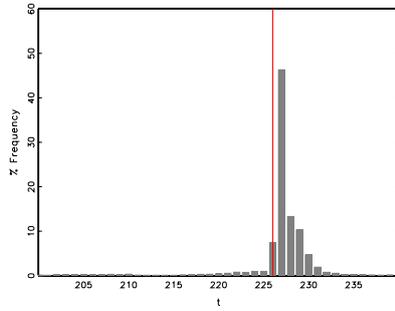(d) $S_{MIN}(5,1)$ (e) $S_{MIN}(5,2)$ (f) $S_{MIN}(5,3)$
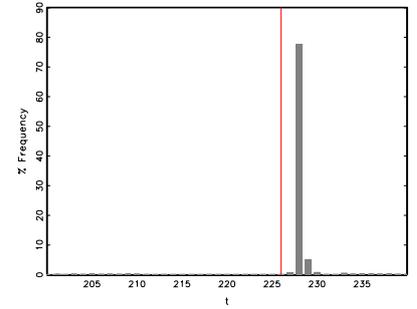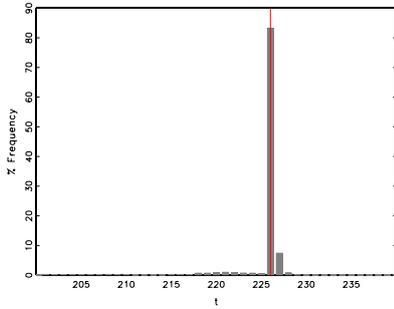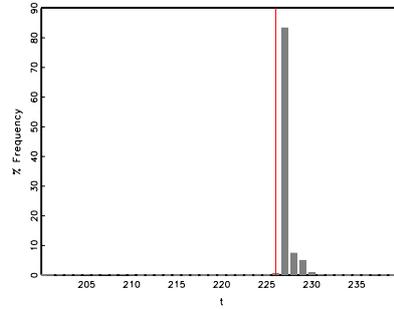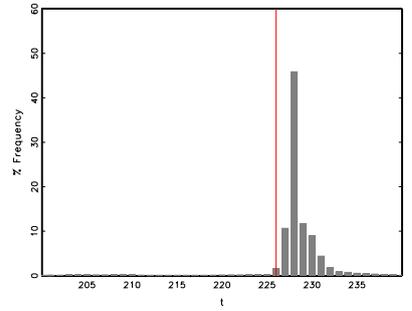
(g) $S_{MIN}(10,1)$ (h) $S_{MIN}(10,2)$ (i) $S_{MIN}(10,3)$
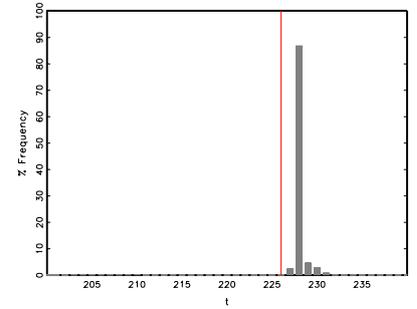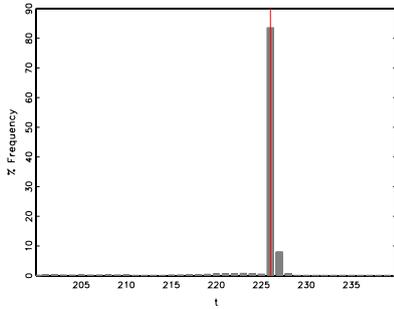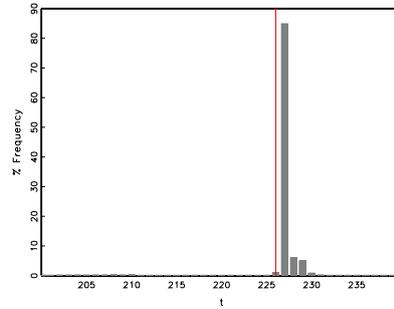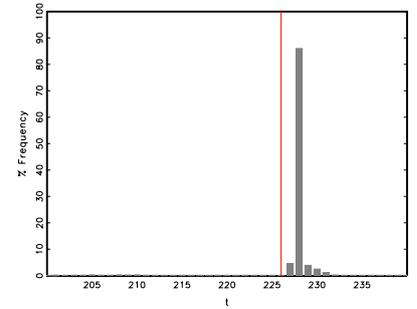
(j) $S_{MIN}(15,1)$ (k) $S_{MIN}(15,2)$ (l) $S_{MIN}(15,3)$

— $\lfloor \tau_2 T \rfloor + 1$

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) $m = k = 15$

—— $A_{MAX}(k)$, —— $S_{MIN}(m,1)$, —— $S_{MIN}(m,2)$, —— $S_{MIN}(m,3)$, - - FPR, - - $\lfloor \tau_1 T \rfloor$, $\lfloor \tau_2 T \rfloor$, $\lfloor \tau_3 T \rfloor$
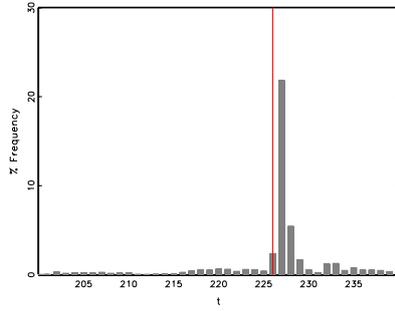
(d) $S_{MIN}(5,1)$ (e) $S_{MIN}(5,2)$ (f) $S_{MIN}(5,3)$
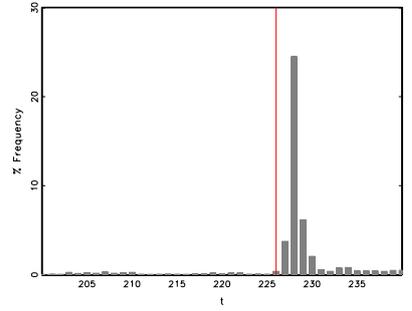
(g) $S_{MIN}(10,1)$ (h) $S_{MIN}(10,2)$ (i) $S_{MIN}(10,3)$
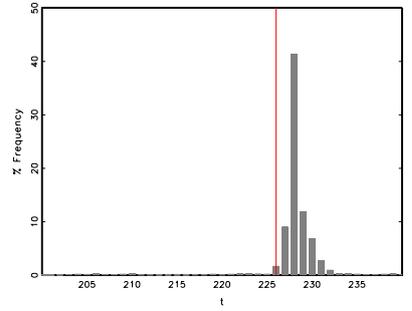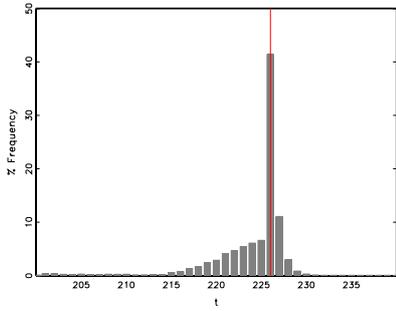
(j) $S_{MIN}(15,1)$ (k) $S_{MIN}(15,2)$ (l) $S_{MIN}(15,3)$

—— $\lfloor \tau_2 T \rfloor + 1$

Figure 6: Rejection frequencies of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for multiple bubble and crash episodes: $\tau_{j,1} = 1$
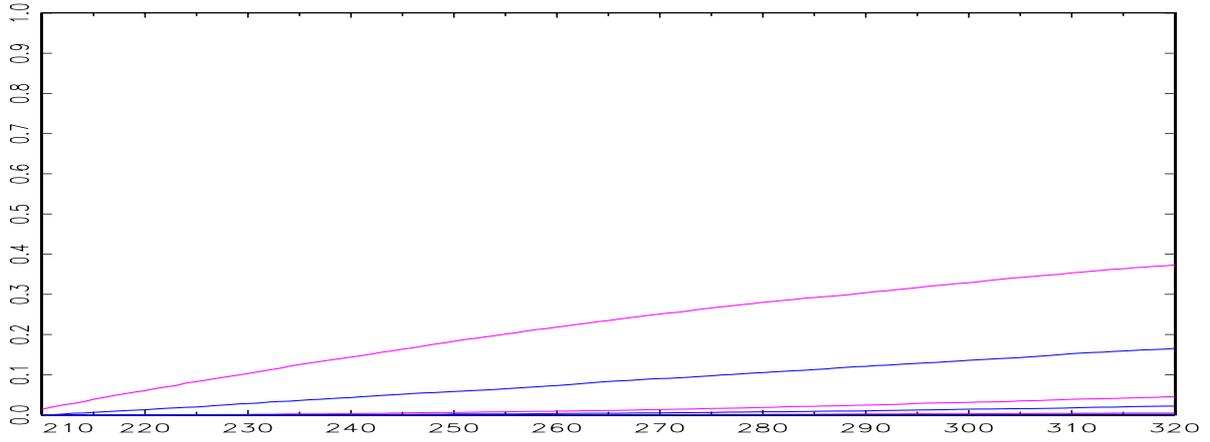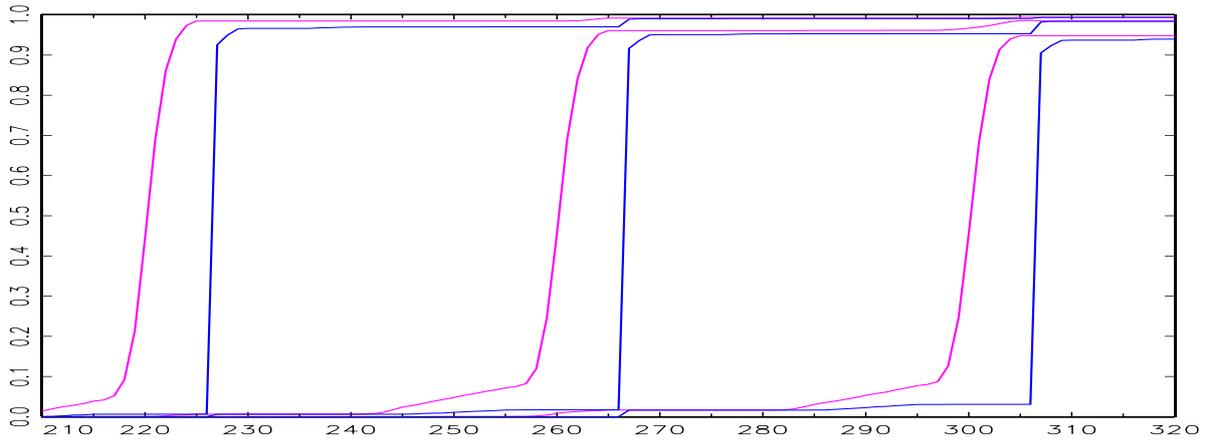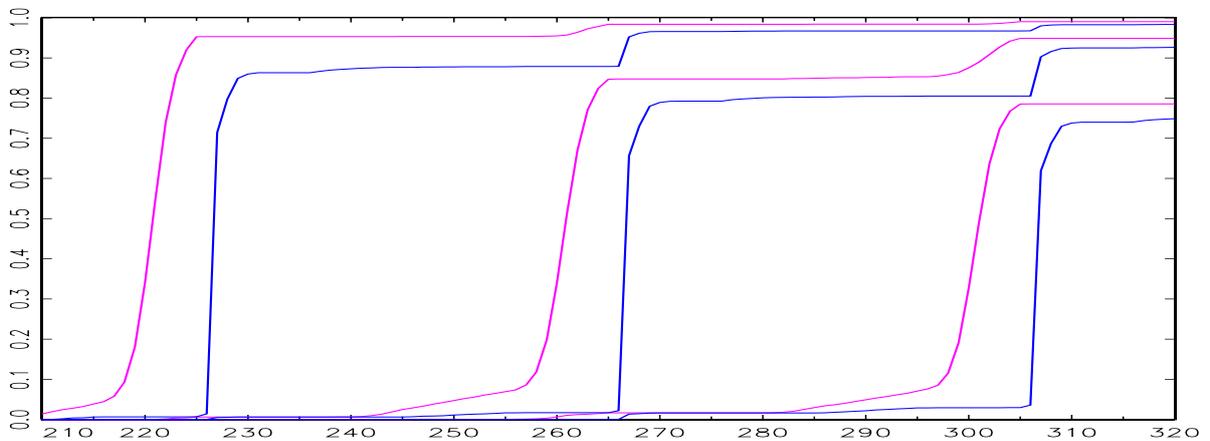

Figure 7: Rejection frequencies of $A_{MAX}(10)$ and $S_{MIN}(10,2)$ for multiple bubble and crash episodes: $\{\lfloor\tau_{1,1}T\rfloor, \lfloor\tau_{1,2}T\rfloor, \lfloor\tau_{1,3}T\rfloor\} = \{215, 225, 235\}$, $\{\lfloor\tau_{2,1}T\rfloor, \lfloor\tau_{2,2}T\rfloor, \lfloor\tau_{2,3}T\rfloor\} = \{255, 265, 275\}$, $\{\lfloor\tau_{3,1}T\rfloor, \lfloor\tau_{3,2}T\rfloor, \lfloor\tau_{3,3}T\rfloor\} = \{295, 305, 315\}$
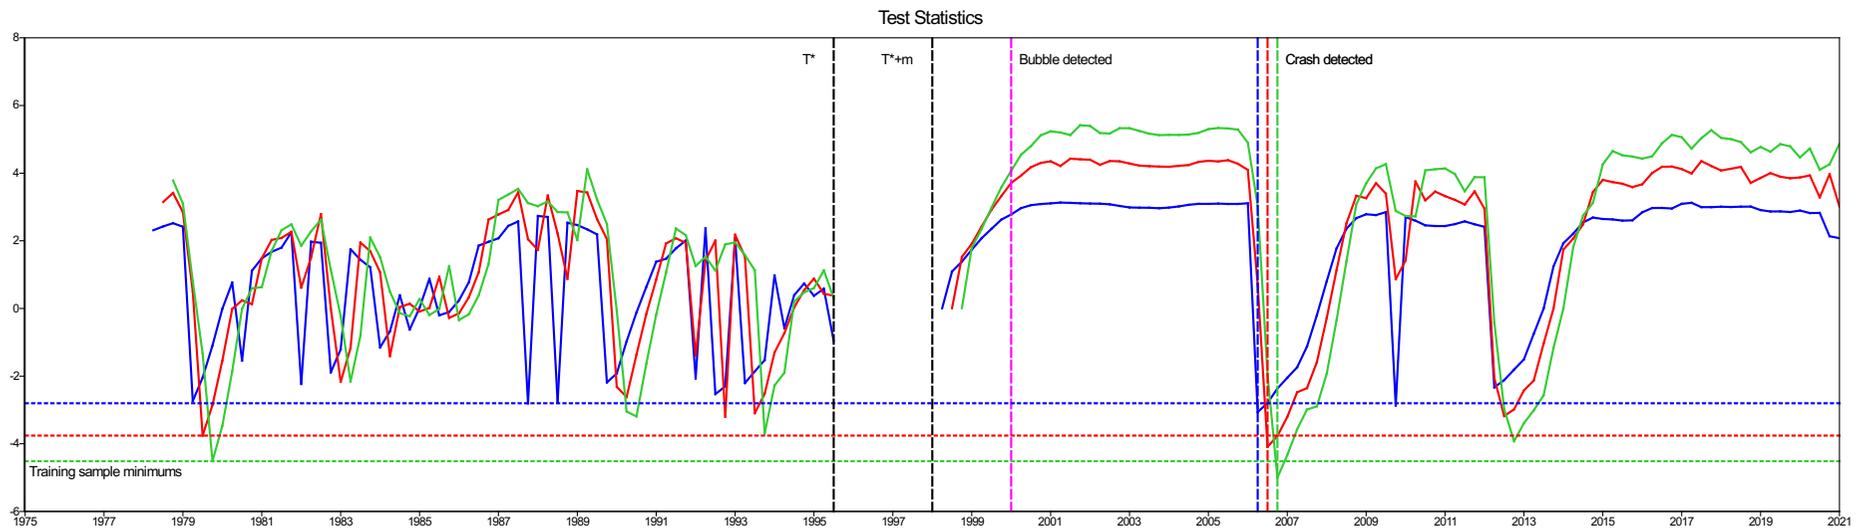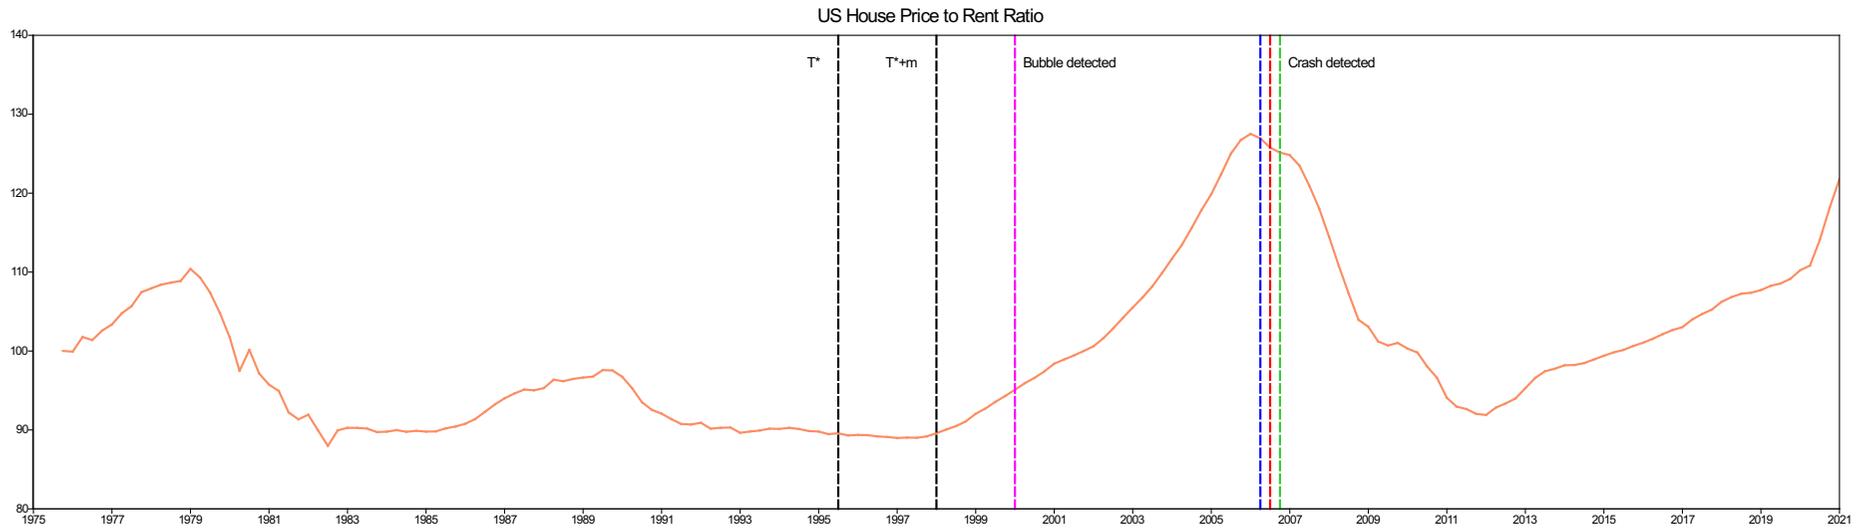

(a) $\delta_{j,1} = 0.03$ and $\delta_{j,2} = 0.015$

(b) $\delta_{j,1} = 0.02$ and $\delta_{j,2} = 0.01$

$\qquad$ $A_{MAX}(10)$, $\qquad$ $S_{MIN}(10,2)$, $--$ $\lfloor\tau_{j,1}T\rfloor$, $\lfloor\tau_{j,2}T\rfloor$, $\lfloor\tau_{j,3}T\rfloor$

Figure 8: US house price to rent ratio and $S_{MIN}(m,n)$ test statistics, 1975-2021