Muhammad Farid Ahmed³ / Stephen Satchell^{1,2}

What Proportion of Time is a Particular Market Inefficient? ... A Method for Analysing the Frequency of Market Efficiency when Equity Prices Follow Threshold Autoregressions

¹ University of Sydney, Camperdown, NSW 2006, Australia

² Trinity College, University of Cambridge, Cambridge, UK

³ Faculty of Economics, Selwyn College, University of Cambridge, Sidgwick Avenue Cambridge, Cambridge, CB3 9DD, UK, E-mail: faridahm@gmail.com

Abstract:

We assume that equity returns follow multi-state threshold autoregressions and generalize existing results for threshold autoregressive models presented in Knight and Satchell 2011. "Some new results for threshold AR(1) models," *Journal of Time Series Econometrics* 3(2011):1–42 and Knight, Satchell, and Srivastava (2014) for the existence of a stationary process and the conditions necessary for the existence of a mean and a variance; we also present formulae for these moments. Using a simulation study, we explore what these results entail with respect to the impact they can have on tests for detecting bubbles or market efficiency. We find that bubbles are easier to detect in processes where a stationary distribution does not exist. Furthermore, we explore how threshold autoregressive models with i.i.d trigger variables may enable us to identify how often asset markets are inefficient. We find, unsurprisingly, that the fraction of time spent in an efficient state depends upon the full specification of the model; the notion of how efficient a market is, in this context at least, a model-dependent concept. However, our methodology allows us to compare efficiency across different asset markets.

Keywords: market efficiency, bubbles, threshold auto-regressions

JEL classification: C24, G14

DOI: 10.1515/jtse-2016-0021

1 Introduction

The tendency of asset prices to go through locally explosive and mean reverting states is well documented and has intrigued both theoretical and empirical economists. Regime switching models, such as the ones introduced by Goldfeld and Quandt (1973), Tong (1978), and Hamilton (1989), have often been employed to empirically estimate asset prices with regime changes. These models include hidden Markov-state models as well as threshold autoregressive models.

Knight and Satchell (2011) study the steady-state properties of asset prices that are estimated using threshold auto-regressive models. Their article formalizes necessary and sufficient conditions for the existence of a stationary distribution for regime-switching threshold models. Analytical expressions for the mean, variance, co-variance and the steady-state distribution are also derived. While Knight et al. carry out most of their analysis under the assumption of an independent and identically distributed trigger variable they also consider the case of a threshold auto-regressive model (TAR henceforth) with a Markov trigger.

The current article builds on their results and generalizes the conditions required for a TAR model to have a stationary distribution with any number of finite regimes. We only consider the case where the TAR model is driven by an independent and identically distributed (i.i.d) exogenous variable that triggers regime switches. We also derive analytical expressions for the mean and variance of these models, noting the conditions that need to be satisfied for their existence. The two moments are derived with a switching drift, with a constant drift and with no drift.

This enables us to treat market efficiency as a state, instead of a condition that holds universally whereby an asset market is either efficient or inefficient. Considering efficiency as one of several states, we make a number of

Muhammad Farid Ahmed is the corresponding author.

© 2018 Walter de Gruyter GmbH, Berlin/Boston.

contributions to the literature on the efficient market hypothesis. We contribute to the bubble testing literature by carrying out a simulation study which compares the power of bubble detection tests in situations where a stationary distribution exists against situations where it does not. Here, the explosive or bubble state is one state of a multi-state price process.

Evans (1991) pointed out in his seminal study that bubble detection tests are less useful when an observed series contains multiple instances of collapsing bubbles. His study showed that such tests lose power when the number of bubbles and collapses in a series increases. A number of studies have attempted to address this criticism through alternative methodologies. The most notable ones among them are Hall, Psaradakis, and Sola (1999) and Phillips, Shi, and Yu (2013). Hall et al used a Markov-state regime switching model to estimate the probability of an asset being in an explosive state. Phillips et al. on the other hand have devised a recursive procedure using the Augmented Dickey Fuller test which allows them to not only test for explosiveness but also date these bubbles. The GSADF test as it is now called has proven to be popular with macroeconomists and financial economists. Our simulations show that while this test is statistically powerful in empirical application it has its limitations.

Using both i.i.d and Markov switching-regime simulated series we show that when a time series resembling an asset price fails to satisfy the conditions for a stationary distribution, the GSADF test has high power. On the other hand, the power of the GSADF test falls considerably when the process has a stationary distribution even though locally explosive regimes continue to be present. Thus, our simulation study builds on our theoretical results and further elaborates on observations made in Knight, Satchell, and Srivastava (2014), who outline reasons for the failure of bubble detection tests when a series has a stationary distribution.

Our analysis provides a limiting feature of the GSADF test as the test is premised on a process not having a mean reverting state. Phillips et al. show that the GSADF test has higher power than other alternative bubble detection test; thus, we contend that these results should have external validity for other bubble tests. Furthermore, we note that the power of the GSADF test increases the farther the explosive regime parameter is from unity. This observation is also supported by the formulae we derive in our workings and is discussed in the relevant section.

Finally, we also contribute to the market efficiency literature by providing a methodology that may be used to estimate how often an asset market is efficient and also allows us to compare efficiency across different markets. The efficient market hypothesis is perhaps the most well-known as well as the most divisive hypothesis in economics. While economists were aware of market efficiency for a very long time before him, Fama (1965) was the first to define the market as being efficient in his seminal article on stock prices where he concluded that stock market prices followed a random walk.

Since then a large number of economists have contributed to this literature with both proponents and opponents of the hypothesis contributing. Seminal contributions have been made to this literature by Samuelson (1965), De Bondt and Thaler (1985), Marsh and Merton (1986), and Shiller (2000), among others. The hypothesis itself states that for a given information set I_t , systematic gains cannot be made by trading on the information set alone. In fact, it may be argued that the threshold auto-regressive model literature and the bubble testing literature is a subset of the efficient markets literature as the techniques developed to estimate asset markets have often been discussed in the context of efficient markets.

Note that in this article when we mentioned market efficiency we are referring to the weak form of market efficiency which states that returns cannot be predicted based on prior information, i. e. the impact of prior prices is already reflected in the current price. In econometric terms this suggests that asset price follow a random walk with (or without) drift, i. e. $P_t = \alpha + P_{t-1} + \varepsilon_t$. In the context of this research, thus, a deviation from efficiency should be understood as a deviation from the random walk process or any deviation in slope from the specification above. Instead of arguing for or against the efficient market hypothesis we recognize that although markets may be mostly weak form efficient, they can deviate from efficiency for significant periods of time. Our estimation methodology aims at identifying how long, in terms of what proportion of time, these inefficient periods last.

We outline an estimation methodology for asset markets where market efficiency is one of several states. Our estimation methodology aims at identifying how long periods of efficiency and inefficiency last. To the best of our knowledge this research would be the first to provide a metric for market efficiency using threshold auto-regressions. Historically, research has taken a binary view towards market efficiency and has been based around the presence or absence of market efficiency taken on average over a sample period; our metric provides a more detailed view. We believe that markets can undergo efficient as well as inefficient states and our methodology helps us determine and estimate how long such states last.

We provide an illustrative empirical application of our methodology through estimating a TAR(1) model for the S&P500 and FTSE 100 stock market indices where the parameter switches due to an exogenous trigger variable. We estimate the TAR(1) using a constant drift, a regime-switching drift and no drift. Our results indicate that the inclusion of a drift term, particularly a regime-switching drift term, reduces the impact of the

regime-switching auto-regressive parameter. A switching drift term is able to explain changes in the return process and thus, the high variance observed during explosive periods. With a switching drift term, markets appear to be more efficient than with no drift as we are unable to reject the random walk hypothesis for a number of coefficients. This observation indicates that in the context of regime-switching models, market efficiency is a model driven concept. However, our methodology does allow us to compare efficiency across different markets for the same model specification.

Other measures of efficiency based on trading volumes or number of informed traders may be used to gain estimates of market efficiency. This would require estimating a trading model for the former and a heterogeneous agent model for the latter. We argue that our estimation methodology is much simpler and the data required for estimation (i. e. prices or returns) are much easily available compared to trading volumes and private information of traders. The methodology may be used for forecasting purposes; however, no data testing on forecasting is carried out as it goes beyond the scope of our research which primarily aims at specifying and estimating a metric for market efficiency.

In summary, our main contributions to the literature are a generalization of conditions required for a threshold auto-regressive model to have a mean and variance when a steady-state distribution exists, a simulation study of how bubble tests behave when processes do or do not have steady-state distributions and an empirical methodology for analysing market efficiency through threshold autoregressive models driven by an exogenous trigger variable. Section 2 generalizes the results in Knight and Satchell (2011) to a finite number of regimes and presents results on the mean and variance of TAR(1) models. Section 3 explains the results of the simulation study using both i.i.d and Markov chain threshold exogenous triggers. Section 4 illustrates how the model in Section 2 may be estimated in practice and how it can be used to construct efficiency measures. Section 5 concludes.

2 Conditions for the Existence of a Mean and Variance

Let P_t be the price (or log price) of some asset. We assume that:

$$P_t = \psi_{t-1} + \phi_{t-1}P_{t-1} + \eta_t \text{ where } \eta_t \sim N(0, \sigma_\eta^2) \quad (1)$$

For illustrative purposes we consider the three state case although the results will be applicable to a finite number of 'k' states. For $k = 3$ we specify values of the trigger or driving variable Z_t for which the parameters switch between values.

Thus, for a three-state case we have:

$$\psi_{t-1} = \alpha_1, \phi_{t-1} = \beta_1 \text{ if } -\infty < Z_{t-1} < c_1$$

$$\psi_{t-1} = \alpha_2, \phi_{t-1} = \beta_2 \text{ if } c_1 \leq Z_{t-1} < c_2$$

$$\psi_{t-1} = \alpha_3, \phi_{t-1} = \beta_3 \text{ if } c_2 \leq Z_{t-1} < \infty$$

c_1 and c_2 are threshold levels which trigger the switch between states. The above framework can be generalized to k intervals and $k+1$ constants where $c_0 = -\infty$ and $c_k = \infty$.

We need Z_{t-1} to be i.i.d to derive our results. The probability that Z_{t-1} will take a value between any two constants is assumed to be π_{j+1} , i. e. $P(c_j \leq Z_{t-1} < c_{j+1}) = \pi_{j+1}$ and

$$\sum_{j=1}^k \pi_j = 1$$

As a result there will be '2k' different parameters. We denote the regime specific parameter by α_j, β_j .

We note that conditions for the existence of a stationary distribution will be similar if Z_t followed a Markov process but the moments will be different. (Knight et al. 2011 Theorem 2).

If $E(\ln |\phi_{t-m}|) = \sum_{j=1}^k \pi_j \ln |\beta_j| < 0$ and $\psi_{t-1} = \alpha$ (Knight et al. 2011) then the TAR model given by eq. (1) has the solution

$$P_t = \alpha \left(\sum_{n=0}^{\infty} S_n(t) \right) + \sum_{n=0}^{\infty} S_n(t) \eta_{t-n}, \text{ where } S_n(t) = \prod_{m=1}^n \phi_{t-m} \quad (2)$$

and $S_0(t) = 1$, (Quinn 1982).

When $\psi_{t-1} = \alpha_i$ for $i = 1, 2 \dots k$

$$\begin{aligned} P_t &= \psi_{t-1} + \phi_{t-1}P_{t-1} + \eta_t \\ &= \psi_{t-1} + \phi_{t-1}(\psi_{t-2} + \phi_{t-2}P_{t-2} + \eta_{t-1}) + \eta_t \\ &= (\psi_{t-1} + \phi_{t-1}\psi_{t-2} + \phi_{t-1}\phi_{t-2}\psi_{t-3} + \dots) + \sum_{n=0}^{\infty} S_n(t) \eta_{t-n} \\ P_t &= \sum_{n=0}^{\infty} \psi_{t-1-n} S_n(t) + \sum_{n=0}^{\infty} S_n(t) \eta_{t-n} \end{aligned} \quad (3)$$

It follows that as $n \rightarrow \infty$:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln |S_n(t)| = E[\ln |\phi_{t-m}|] \quad (4)$$

We note that the finiteness of the first term in eq. (3) is governed by the behaviour of $S_n(t)$ but that eq. (4) being satisfied is enough to ensure the existence of eq. (3). We further note that existence of the process does not imply existence of the mean so the first term may not converge to a finite limit in expectation.

In the following sub-sections we derive the mean and variance for the general case $\psi_{t-1} = \alpha_i$ and discuss special cases, i. e. $\psi_{t-1} = \alpha$ and $\psi_{t-1} = 0$.

Mean:

$$\begin{aligned} E(P_t) &= E(\psi_{t-1}) + E(\phi_{t-1}P_{t-1}) \\ &= E(\psi_{t-1}) + E(\phi_{t-1})E(P_{t-1}) \end{aligned} \quad (5)$$

As the ψ 's and ϕ 's switch independently (due to Z_{t-1} being i.i.d), we have:

$$E(\phi_{t-1}) = \sum_{j=1}^k \pi_j \beta_j \quad (6)$$

$$E(\psi_{t-1}) = \sum_{j=1}^k \pi_j \alpha_j \quad (7)$$

It follows from eqs (6) and (7)

$$E(P_t) = \frac{\sum_{j=1}^k \pi_j \alpha_j}{1 - \sum_{j=1}^k \pi_j \beta_j} \quad (8)$$

Thus, the mean will exist if $\sum_{j=1}^k \pi_j \beta_j < 1$

If

$$\psi_{t-1} = \alpha,$$

$$E(P_t) = \frac{\alpha}{1 - \sum_{j=1}^k \pi_j \beta_j} \quad (9)$$

The mean is zero if there is no drift term.

Variance:

For

$$\psi_{t-1} = \alpha_i$$

Under independence of η_t and Z_{t-1} it follows from eq. (1) that:

$$\text{Var}(P_t) = \text{Var}(\psi_{t-1}) + \text{Var}(\eta_t) + \text{Var}(\phi_{t-1}P_{t-1}) + \text{Cov}(\psi_{t-1}, \phi_{t-1}P_{t-1}) \quad (10)$$

We evaluate each term in eq. (10) separately

$$\begin{aligned} \text{var}(\psi_{t-1}) &= E(\psi_{t-1}^2) - [E(\psi_{t-1})]^2 \\ &= \sum_{j=1}^k \pi_j \alpha_j^2 - \left(\sum_{j=1}^k \pi_j \alpha_j \right)^2 \end{aligned} \quad (11)$$

$$\text{Var}(\eta_t) = \sigma_\eta^2 \quad (12)$$

$$\begin{aligned} \text{Var}(\phi_{t-1} P_{t-1}) &= E(\phi_{t-1}^2) E(P_{t-1}^2) - [E(\phi_{t-1}) E(P_{t-1})]^2 \\ &= E(\phi_{t-1}^2) [\text{Var}(P_{t-1}) + [E(P_{t-1})]^2] - [E(\phi_{t-1}) E(P_{t-1})]^2 \\ &= \text{var}(\phi_{t-1}) [E(P_{t-1})]^2 + E(\phi_{t-1}^2) \text{Var}(P_{t-1}) \end{aligned} \quad (13)$$

Equation (13) again makes use of the fact that Z_{t-1} is an i.i.d process which implies independence between ϕ_{t-1} and P_{t-1} .

Using the definition of variance we know that:

$$\begin{aligned} \text{var}(\phi_{t-1}) &= E(\phi_{t-1}^2) - [E(\phi_{t-1})]^2 \\ &= \sum_{j=1}^k \pi_j \beta_j^2 - \left(\sum_{j=1}^k \pi_j \beta_j \right)^2 \end{aligned} \quad (14)$$

Finally, we evaluate $\text{Cov}(\psi_{t-1}, \phi_{t-1} P_{t-1})$

$$\begin{aligned} \text{Cov}(\psi_{t-1}, \phi_{t-1} P_{t-1}) &= E(\psi_{t-1} \phi_{t-1} P_{t-1}) - E(\psi_{t-1}) E(\phi_{t-1}) E(P_{t-1}) \\ &= E(\psi_{t-1} \phi_{t-1}) E(P_{t-1}) - E(\psi_{t-1}) E(\phi_{t-1}) E(P_{t-1}) \\ &= \text{Cov}(\psi_{t-1}, \phi_{t-1}) E(P_{t-1}) \end{aligned} \quad (15)$$

Equation (15) relies on the independence of ψ_{t-1} and ϕ_{t-1} from P_{t-1} .

Thus, eqs (11)–(15) allow us to calculate $\text{Var}(P_t)$

$$\text{Var}(P_t) = [\text{Var}(\psi_{t-1}) + \sigma_\eta^2 + \text{Var}(\phi_{t-1}) [E(P_{t-1})]^2 + E(\phi_{t-1}^2) \text{Var}(P_{t-1}) + \text{Cov}(\psi_{t-1}, \phi_{t-1}) E(P_t)] \quad (16)$$

Rearranging eq. (16) and recognizing that if P_t has a stationary distribution $\text{Var}(P_{t-1}) = \text{Var}(P_t)$ and $E(P_{t-1}) = E(P_t)$

$$\text{Var}(P_t) (1 - E(\phi_{t-1}^2)) = \text{Var}(\psi_{t-1}) + \sigma_\eta^2 + \text{Var}(\phi_{t-1}) [E(P_t)]^2 + \text{Cov}(\psi_{t-1}, \phi_{t-1}) E(P_t)$$

Thus,

$$\text{Var}(P_t) = \frac{[\text{Var}(\psi_{t-1}) + \sigma_\eta^2 + \text{Var}(\phi_{t-1}) [E(P_t)]^2 + \text{Cov}(\psi_{t-1}, \phi_{t-1}) E(P_t)]}{1 - E(\phi_{t-1}^2)} \quad (17)$$

The condition for the existence of a finite variance with switching regimes is thus,

$$E(\phi_{t-1}^2) = \sum_{j=1}^k \pi_j \beta_j^2 < 1 \quad (18)$$

The expression in eq. (17) suggests that the variance of the price series is increasing in the variance of the drift parameter ψ_{t-1} , the variance of the error term σ_η^2 , the variance of the coefficients ϕ_{t-1} , the expected value of P_t (which includes the absolute value of the drift) and the covariance between the drift and coefficient parameters. The latter will be positive since we assume that the same exogenous parameter causes a regime switch in both the drift and the coefficient parameter. If we have regimes with slope coefficients deviating far from unity (a case we will be interested in when considering the efficient market hypothesis), we will get a much higher variance for the price series. Similarly, if there is a large drift term, it can lead to the variance of price being high. This also gives us an early indication that model specification may be an important determinant in analysing asset price series. Given the amount of variation in an observed set of series, different specification will lead to the variation being captured by different parameters. Indeed, this is what we observe in Section 4.

Now we consider the case where the drift is constant, i. e. $\psi_{t-1} = \alpha$

A constant drift implies $Var(\psi_{t-1}) = 0$ and $Cov(\psi_{t-1}, \phi_{t-1}) = 0$. Therefore, the expression for the variance of the asset price with a constant drift term reduces to:

$$Var(P_t) = \frac{[\sigma^2 + Var(\phi_{t-1}) [E(P_{t-1})]^2]}{1 - E(\phi_{t-1}^2)} \quad (19)$$

With a constant drift term the variance of the process relies on the variance of the error term, the variance of the coefficient parameter and the absolute value of the expected price which itself is a function of the drift term; thus, the variance of price is dependent on the absolute value of the drift term.

Finally, we consider the case where $\psi_{t-1} = 0$. With no drift term, eq. (9) implies that $E(P_{t-1}) = 0$. Thus,

$$Var(P_t) = \frac{\sigma^2}{1 - E(\phi_{t-1}^2)} \quad (20)$$

In the vicinity of a unit root, $E(\phi_{t-1}^2)$ is likely to be close to 1, which will lead to a very large variance for the process. Nevertheless, the variance will be finite and will exist as long as the condition in eq. (18) is satisfied. Sections 3 and 4 analyse the implications of the formulae derived above. Section 3 considers a series of simulations to show how violations of the criterion for a stationary steady-state distribution specified above and the variance of the parameter coefficient impact the ability of statistical tests to detect explosive roots or bubbles. Section 3 uses an illustrative empirical study to show how specifications such as eq. (1) may be used to analyse the efficiency of asset markets.

3 Simulation Results

In order to support our theoretical analysis we carry out a simulation study. We simulate series with a switching autoregressive parameter and a standard normal error term. We do not consider a switching drift term for our simulations as a constant drift adequately addresses the issue we wish to highlight. The simulated series takes the following form:

$$y_t = \alpha + \phi_{t-1}y_{t-1} + \epsilon_t \text{ where } \epsilon_t \sim N(0, 1) \quad (21)$$

In this context y_t can be thought of as the logarithm of a price variable so that if $\phi_{t-1} = 1$, the return $y_t - y_{t-1} = \alpha + \epsilon_t$. In the simulations below we consider the case when $\alpha = 0$, $\alpha = 0.01$ and $\alpha = 0.025$.

Our simulation study considers the three-state case for computational ease, although the results will also hold for a finite number of k-states. The switching parameter depends on the pseudo sentiment variable, Z_{t-1} which in our simulations is either a multinomial vector or a Markov chain variable. For the multinomial vector case, we select the probability with which each state occurs. Thus, Z is $[1 \ 0 \ 0]$ when in state 1, $[0 \ 1 \ 0]$ when in state 2 and $[0 \ 0 \ 1]$ when in state three with probabilities π_1, π_2 and π_3 respectively. The value taken by Z_{t-1} depends on the parameter vector specified.

On the other hand if Z_t is a Markov chain variable, it takes on the values 1, 2 or 3 depending on which state the series is in. For the Markov chain simulations we need to specify a transition matrix instead of a probability vector. A Markov chain is more intuitive for the type of series we are concerned with as states tend to be more persistent in this case once the switch occurs. It is also more comparable to the kind of simulations used in the literature related to test for explosiveness or bubbles. In addition to the process followed by the switch inducing variable Z_{t-1} , we also need to specify values for the switching parameters. Together, the switching parameter and the probability of Z_{t-1} enable us to verify if the criterion for a steady-state distribution is satisfied.

In order to illustrate what happens when the criterion is satisfied and when it is not we carry out the GSADF or Phillips, Shi and Yu (PSY henceforth) test on each simulated series to check if the test is able to detect explosiveness in the series which in the context of our simulations may be construed as the process not having a steady-state distribution. This takes the form of a power test. Each simulated series contains periods of explosiveness which is the primary reason for taking the series away from stationarity. This statistic is being used in different areas of economics and PSY have shown it to have high power in detecting explosiveness or bubbles. The test involves estimating recursive regressions of the following form:

$$\Delta y_t = \alpha + \phi y_{t-1} + \sum_{i=1}^k \psi^i \Delta y_{t-i} + \epsilon_t$$

In the above regression the parameter of interest is ϕ which is estimated through expanding, rolling windows with a minimum window size specified by the researcher. For each regression a right-sided unit root statistic is calculated. The supremum(sup) of all right-sided unit root statistics thus calculated is the GSADF or PSY statistic. The sup value can be compared to simulated critical values, allowing the user to comment on whether a bubble may be present in the series under consideration. We refer the interested reader to Phillips, Shi, and Yu (2013) for further details on the test procedure and asymptotic properties of the statistic.

Multinomial Switching variable:

For clarity, we indicate the specific form taken by our simulated series. When Z_{t-1} is a multinomial vector taking values $Z_{t-1,j}$, it takes the following form: $y_t = \alpha + (\sum_{j=1}^3 \beta_j Z_{t-1,j}) y_{t-1} + \epsilon_t$

$$y_t = \alpha + \beta_1 y_{t-1} + \epsilon_t \text{ when } Z_{t-1,1} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \text{ with probability } \pi_1$$

$$y_t = \alpha + \beta_2 y_{t-1} + \epsilon_t \text{ when } Z_{t-1,2} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \text{ with probability } \pi_2$$

$$y_t = \alpha + \beta_3 y_{t-1} + \epsilon_t \text{ when } Z_{t-1,3} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \text{ with probability } (\pi_3 = 1 - \pi_1 - \pi_2)$$

Section 2 considers a similar exogenous trigger. To aid understanding, the first element of the Z_{t-1} vector indicates the mean reverting state, the second element indicates the random walk or efficient state and the third element indicates the explosive or bubble state. For each set of parameter and probability values we generated 500 simulated series and the GSADF test was conducted on each series. Each series is 1000 observations long and the minimum window size was stipulated to be 10 % of the series or 100 observations. Critical values were generated separately using the MATLAB code provided by PSY. The GSADF test was conducted at the 5 % level (critical values for the GSADF test at the 5 % level for series of length 1000 with initial window size of 100 is 2.16 for series without drift and 2.233 with drift). Since each simulated series is explosive 10 % of the time on average, each series exhibits the type of explosive behaviour that the GSADF test seeks to detect. The power is simply calculated by dividing the number of bubbles detected by 500 for each set of 500 simulations.

Table 1 shows the results of our simulations along with the parameter values and the probability vector. Column 3 shows the value of the criterion for a steady-state distribution. The criterion is said to be satisfied whenever $\sum_{j=1}^3 \pi_j \ln |\beta_j| < 0$. We ensured that we chose a range of values so that for some values the criterion was satisfied and for other values it was not. For the multinomial vector case, we note that the power of the GSADF test is much higher when the criterion is not satisfied. We illustrate our results by considering some sets of parameter and probability values. For parameter values [0.96 1 1.05] and a probability vector [0.10 0.80 0.10] we obtain a criterion value of 0.00080 and as per our theoretical results the series should not have a stationary distribution. We see that when the criterion threshold is breached, we get a power of 17.6 % from the GSADF test. Note that the existence of a stationary distribution does not guarantee the existence of moments. With two exceptions (parameter vector = [0.98 1.02 1.05] and [0.96 1 1.03]), all other sets of values do not have a mean or variance even though the distribution may exist (when the criterion is satisfied).

Table 1: Power test for the GSADF statistic using a multinomial vector ($\alpha = 0$).

Parameter vector	Multinomial probability vector	Criterion $\sum_{j=1}^3 \pi_j \ln \beta_j < 0$	Power (%)
[0.98 1 1.02]	[0.10 0.80 0.10]	-0.00004	5.6
[0.99 1 1.02]	[0.10 0.80 0.10]	0.00098	12.6
[0.96 1 1.03]	[0.10 0.80 0.10]	-0.0011	3.6
[0.97 1 1.03]	[0.10 0.80 0.10]	-0.00009	5.8
[0.98 1 1.03]	[0.10 0.80 0.10]	0.00094	11.0
[0.97 1 1.04]	[0.10 0.80 0.10]	0.0009	15.8
[0.95 1 1.05]	[0.10 0.80 0.10]	-0.00025	9.0

[0.96 1 1.05]	[0.10 0.80 0.10]	0.00080	17.6
[0.98 1.02 1.05]	[0.80 0.10 0.10]	-0.0093	1.2
[0.90 1 1.10]	[0.10 0.80 0.10]	-0.0010	23.4

Contrast this with cases when the criterion is satisfied e. g. when the parameter vector is [0.95 1 1.05] and the probability vector is [0.10 0.80 0.10]. The intensity of the bubble or explosive behaviour stays the same, i. e. the bubble increases the value of the series by 10 % each period and the explosive state occurs 10 % of the time in the long run. We note that even with no change in the bubble state parameter value the power of the GSADF test reduces markedly down to 9.0 %. This supports our theoretical results and shows that if bubbles occur in assets which may have a long run steady-state distribution, they may be harder to detect.

In their article PSY do not use a mean reverting state. Their analysis is based on a random walk and a mildly explosive regime which will not satisfy the criterion and is thus, likely to result in a higher power for their test based on what we observe in our simulations. While the test has undoubtedly been useful in many applications, it is important to keep its limitations in mind particularly when it is unable to detect bubbles in an asset which may otherwise be thought to have gone through periods of explosiveness. The types of series considered by PSY are closer to the Markov-chain simulations in the following sub-section, thus, some of the low power detected in this sub-section may be attributed to the choice of our i.i.d exogenous trigger.

We also note that the power of the test increases the farther apart from unity the explosive state is i. e. $\text{var}(\phi_{t-1})$. For instance, when we reduce $\text{var}(\phi_{t-1})$ and consider the parameter vector [0.98 1 1.02] with the same probability vector as before, the power reduces to 5.6 % even though the criterion is smaller than before. We also consider the case with 1 mean reverting and two explosive states with the mean reverting state occurring 80 % of the time. With the same probability vector, this set of values attained a power of only 1.2 %. This is the only set of values for which a mean and variance exists.

Tables 2 and 3 on the other hand report results for simulations which include a constant drift term. Table 2 contains results for a drift of 0.01. Table 3 contains results for a drift of 0.025. The higher drift value is chosen in order to illustrate how the power of the test depends on the drift term. Note that as per the results in Section 2, the size of the drift plays a role in determining the variance and is likely to impact the results.

Table 2: Power test for the GSADF statistic using a multinomial vector with drift ($\alpha = 0.01$).

Parameter vector	Multinomial probability vector	Criterion $\sum_{j=1}^k \pi_j \ln \beta_j < 0$	Power (%)
[0.98 1 1.02]	[0.10 0.80 0.10]	-0.00004	8.4
[0.99 1 1.02]	[0.10 0.80 0.10]	0.00098	13
[0.96 1 1.03]	[0.10 0.80 0.10]	-0.0011	4.2
[0.97 1 1.03]	[0.10 0.80 0.10]	-0.00009	8
[0.98 1 1.03]	[0.10 0.80 0.10]	0.00094	16.6
[0.97 1 1.04]	[0.10 0.80 0.10]	0.0009	20.2
[0.95 1 1.05]	[0.10 0.80 0.10]	-0.00025	11.2
[0.96 1 1.05]	[0.10 0.80 0.10]	0.00080	20.4
[0.98 1.02 1.05]	[0.80 0.10 0.10]	-0.0093	1.3
[0.90 1 1.10]	[0.10 0.80 0.10]	-0.0010	29.6

Table 3: Power test for the GSADF statistic using a multinomial vector with drift ($\alpha = 0.025$).

Parameter vector	Multinomial probability vector	Criterion $\sum_{j=1}^k \pi_j \ln \beta_j < 0$	Power (%)
[0.98 1 1.02]	[0.10 0.80 0.10]	-0.00004	6.6
[0.99 1 1.02]	[0.10 0.80 0.10]	0.00098	12.2
[0.96 1 1.03]	[0.10 0.80 0.10]	-0.0011	4.6
[0.97 1 1.03]	[0.10 0.80 0.10]	-0.00009	8
[0.98 1 1.03]	[0.10 0.80 0.10]	0.00094	18.8
[0.97 1 1.04]	[0.10 0.80 0.10]	0.0009	18.6
[0.95 1 1.05]	[0.10 0.80 0.10]	-0.00025	15
[0.96 1 1.05]	[0.10 0.80 0.10]	0.00080	20.8
[0.98 1.02 1.05]	[0.80 0.10 0.10]	-0.0093	1.4
[0.90 1 1.10]	[0.10 0.80 0.10]	-0.0010	31.6

When a constant drift term is included the power of the GSADF test goes up for all sets of values compared to the case with no drift and our main results continue to hold, i. e. the power is lower if the criterion for a steady-state distribution is satisfied and the variance of the parameter vector is low despite the fact that the mean and variance for the set of values chosen do not exist. As per the expressions in Section 2 for the mean and variance of the simulated series, a higher α implies not just a higher mean but also a higher variance. While it is not clear from our results whether a higher α necessarily leads to a higher power for the GSADF test, the powers attained are higher than the no drift case. We do note, however, that for the case where both a mean and a variance exist a higher α leads to a higher power.

Markov Chain trigger variable:

In this section we consider a Markov-chain trigger variable which influences states. Instead of a three-dimensional multinomial vector, Z_{t-1} is now a Markov-chain variable dependant on a transition matrix. Instead of specifying a vector of probabilities we now specify a transition matrix which determines the value of Z_{t-1} , the pseudo sentiment variable. Thus, the process becomes:

$$y_t = \alpha + \beta_1 y_{t-1} + \epsilon_t \text{ when } Z_{t-1} = 1$$

$$y_t = \alpha + \beta_2 y_{t-1} + \epsilon_t \text{ when } Z_{t-1} = 2$$

$$y_t = \alpha + \beta_3 y_{t-1} + \epsilon_t \text{ when } Z_{t-1} = 3$$

with the transition matrix

$$\begin{pmatrix} p_{11} & p_{12} & p_{13} \\ p_{21} & p_{22} & p_{23} \\ p_{31} & p_{32} & p_{33} \end{pmatrix}$$

The probability that Z_{t-1} takes a value of 1 given that it was 1 in the previous period is p_{11} ; the probability that Z_{t-1} takes a value of 2 given that it took a value of 1 in the previous period is p_{12} and so on. Thus Z_{t-1} is a one period Markov-chain variable. Simulating the series in this way has the desirable property that states tend to be more persistent compared to the multinomial case. While over the long term the duration of each state is similar to the multinomial case, as evident from the steady-state probability vector, each state tends to last longer. Thus, it can be argued that series generated in this way share more properties of actual asset price series.

As in the previous section, we carry out 500 simulations for each set of values. In order to aid comparison we ensure that the transition matrix was such that the steady-state probabilities of states were similar to those used in the multinomial series. One simplification made in selecting values for the transition matrix is that there are no switches from the explosive state to the mean reverting state and vice versa. Thus, whenever there is a switch from either the explosive or the mean reverting state, it is to the random walk state in the first instance.

Simulating and testing Markov-chain series further strengthens the results obtained from the previous subsection. Using Markov-chains instead of multinomial vectors the simulated series exhibit more asset price like properties and due to state persistence we obtain higher powers for each set of values compared to the multinomial vector counterpart. Table 4 reports results for Markov-chain simulations below. When the criterion is set to -0.00004 the multinomial vector series have a power of 5.6 % compared to 25.8 % for the corresponding Markov-chain simulations. A similar pattern is observed for the remaining values. This observation may be attributed to state persistence introduced by the Markov-chain which enables detection via the GSADF test.

Table 4: Power Test for the GSADF statistic using a Markov-Chain trigger ($\alpha = 0$).

Parameters	Transition matrix	Criteria	Steady-state probabilities	Power (%)
[0.98 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00004	[0.10 0.80 0.10]	26.8
[0.99 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00098	[0.10 0.80 0.10]	38.2
[0.96 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.0011	[0.10 0.80 0.10]	41

[0.97 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00009	[0.10 0.80 0.10]	45.4
[0.98 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00094	[0.10 0.80 0.10]	52.6
[0.97 1 1.04]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0009	[0.10 0.80 0.10]	68.4
[0.95 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00025	[0.10 0.80 0.10]	68.2
[0.96 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0008	[0.10 0.80 0.10]	73.8
[0.98 1.02 1.05]	$\begin{pmatrix} .025 & .95 & 0.025 \\ 0.20 & .80 & 0 \\ 0.20 & 0 & .80 \end{pmatrix}$	-0.0093	[0.80 0.10 0.10]	25.8
[0.90 1.00 1.10]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.0010	[0.80 0.10 0.10]	91.4

As noted previously, $\text{var}(\phi_{t-1})$ influences the results. When the criterion is not satisfied and the non-efficient states significantly deviate from 1, we get very high power. For example, when the parameter vector is $\begin{pmatrix} 0.98 \\ 1 \\ 1.02 \end{pmatrix}$ with a transition matrix so chosen to give a steady-state probability vector $\begin{pmatrix} 0.10 \\ 0.80 \\ 0.10 \end{pmatrix}$ we find a power of 26.8 % (criterion value -0.00004). Keeping the explosive state at 1.02, if the mean reverting state is made more persistent (0.99 from 0.98), the criterion is violated (0.00098) and we get a higher power for the GSADF test at 38.2 %. If we increase the deviations from the random-walk state while maintaining the same steady-state probabilities, the power increases further even though the value of the criterion itself does not change significantly.

For example, if the parameter vector is $\begin{pmatrix} 0.96 \\ 1 \\ 1.05 \end{pmatrix}$ with the same steady-state probabilities and a criterion value of 0.0008, the power increases to 73.8 %.

We also consider the case when we have multiple explosive states and a mean reverting state (in this case the mean and variance of the process exist according to the criteria set in Section 2. Among the two explosive states one is more explosive than the other but both explosive states have the same steady-state probability.

The parameter vector is $\begin{pmatrix} 0.98 \\ 1.02 \\ 1.05 \end{pmatrix}$ with steady-state probabilities $\begin{pmatrix} 0.80 \\ 0.10 \\ 0.10 \end{pmatrix}$ yielding a criterion value of -0.0093.

Corresponding to these values we get a power of 25.8 % for the GSADF statistic. With the multinomial regime-switching variable we observed a power of only 1.2 % for the same set of values. Thus, even with two explosive states we note that if the steady-state distribution criterion is satisfied, explosive behaviour may not be discernible using conventional right-sided unit root tests such as the GSADF test.

We also report results for simulations which include drift terms, $\alpha = 0.01$ and 0.025). Results are reported in Tables 5 and 6. In the Markov-switching case we note that the size of the drift term matters significantly. With a small drift term we do not note a significant change in power compared to the case with no drift (note that the critical values for the two tests are different). For the Markov chain simulations we also note that the power of the test increases as the drift term is increased from 0.01 to 0.025 for all sets of values except one. PSY carried out the test for small deviations from random walk while our simulations include much larger deviations which explain why we observe much higher power despite including a mean reverting term. Nevertheless the results are consistent with the previous set of simulations and the same set of factors namely the value of the criterion, the existence of a drift term and the variance of the parameter vector tend to determine the power of the GSADF test.

Table 5: Power Test for the GSADF statistic using a Markov-Chain trigger ($\alpha = 0.01$).

Parameters	Transition matrix	Criteria	Steady state probabilities	Power (%)
------------	-------------------	----------	----------------------------	-----------

[0.98 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00004	[0.10 0.80 0.10]	29
[0.99 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00098	[0.10 0.80 0.10]	38.8
[0.96 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.0011	[0.10 0.80 0.10]	42.2
[0.97 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00009	[0.10 0.80 0.10]	47.8
[0.98 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00094	[0.10 0.80 0.10]	54.8
[0.97 1 1.04]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0009	[0.10 0.80 0.10]	70.8
[0.95 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00025	[0.10 0.80 0.10]	71.6
[0.96 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0008	[0.10 0.80 0.10]	75.6
[0.98 1.02 1.05]	$\begin{pmatrix} .025 & .95 & 0.025 \\ 0.20 & .80 & 0 \\ 0.20 & 0 & .80 \end{pmatrix}$	-0.0093	[0.80 0.10 0.10]	26
[0.90 1.00 1.10]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.0010	[0.80 0.10 0.10]	92.4

Table 6: Power Test for the GSADF statistic using a Markov-Chain trigger ($\alpha = 0.025$).

Parameters	Transition matrix	Criteria	Steady State probabilities	Power (%)
[0.98 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00004	[0.10 0.80 0.10]	32
[0.99 1 1.02]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00098	[0.10 0.80 0.10]	41.2
[0.96 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.0011	[0.10 0.80 0.10]	44.2
[0.97 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00009	[0.10 0.80 0.10]	50.4
[0.98 1 1.03]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.00094	[0.10 0.80 0.10]	61.6
[0.97 1 1.04]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0009	[0.10 0.80 0.10]	69.6
[0.95 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	-0.00025	[0.10 0.80 0.10]	74.8
[0.96 1 1.05]	$\begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix}$	0.0008	[0.10 0.80 0.10]	75.8
[0.98 1.02 1.05]	$\begin{pmatrix} .025 & .95 & 0.025 \\ 0.20 & .80 & 0 \\ 0.20 & 0 & .80 \end{pmatrix}$	-0.0093	[0.80 0.10 0.10]	27.4

$$[0.90 \ 1.00 \ 1.10] \begin{pmatrix} .80 & .20 & 0 \\ .025 & .95 & .025 \\ 0 & 0.20 & .80 \end{pmatrix} -0.0010 \quad [0.80 \ 0.10 \ 0.10] \quad 94.4$$

Thus, our simulations provide evidence for our theoretical results. The power of bubble detection tests will be higher if the estimated parameters and state probabilities are such that the criterion for a stationary distribution is not satisfied. We also note that the power of such tests is higher when $\text{var}(\phi_{t-1})$ is high.

4 A Metric for Market Efficiency

We use the results in Sections 2 and 3 to outline a methodology that may be employed as a measure for market efficiency. Estimating different states in a process will enable us to understand the stationarity properties of an asset price series and whether the process has a stationary distribution. It will also allow us to evaluate how much time the process spends in each of the states.

While we use time as a metric for market efficiency, other potential metrics may also be employed (these could include but not be limited to metrics using trading volumes or market liquidity)²; however, we argue that our metric is the simplest to estimate and understand. For the data that we use in our empirical application, transaction costs are limited (around 1 basis point for ETFs trading S&P500); these ETFs also tend to be highly liquid and we make the argument that transaction costs and illiquidity are not an issue for the asset markets we consider in our empirical application. Nevertheless, it is important to keep in mind that inefficiency can also result from illiquidity or due to high transaction costs. To measure inefficiency from a transaction or liquidity viewpoint, the TAR model will need to be estimated on measures of transactions and liquidity rather than price. We outline our strategy for estimating market efficiency using pricing data below.

4.1 Methodology

Consider eq. (22):

$$\Delta p_t = \psi_{t-1} + (\phi_{t-1} - 1)p_{t-1} + \eta_t \quad \text{where} \quad \eta_t \sim N(0, \sigma_\eta^2) \quad (22)$$

where,

$$\psi_{t-1} = \alpha_1, \phi_{t-1} = \beta_1 \text{ if } -\infty < Z_{t-1} < c_1$$

$$\psi_{t-1} = \alpha_2, \phi_{t-1} = 0 \text{ if } c_1 \leq Z_{t-1} < c_2$$

$$\psi_{t-1} = \alpha_3, \phi_{t-1} = \beta_3 \text{ if } c_2 \leq Z_{t-1} < \infty$$

We define an n -point grid over the extreme values taken by the trigger variable, Z_{t-1} . Thus, our grid takes values $\{m_1, m_2 \dots m_n\}$. Our recursive procedure starts by considering an initial pair of values for c_1 and c_2 (m_1 and m_2 respectively). Using the values of c_1 and c_2 , the sample is divided into three sub-samples, i. e. the first sub-sample contains all values of the asset price that occur when the trigger variable is less than c_1 , the second sub-sample contains all values of the asset price when the trigger variable is between c_1 and c_2 and the third sub-sample contains all values of the asset price when the trigger variable is greater than or equal to c_2 .

The number of grid points is chosen so that we always have more than two observations between two consecutive grid points (as no estimation will be possible without at least two observations). Following similar terminology to Sections 2 and 3, β_2 is restricted to a value of 0 in order to ensure that the second sub-sample is consistent with an efficient market (if we run the regression in levels instead of logs, β_2 would equal 1). Thus, the market is efficient between the thresholds c_1 and c_2 . If $c_1 = c_2$ then the market is not efficient for any amount of time. We estimate parameters $\alpha_1, \beta_1, \alpha_2, \alpha_3$ and β_3 for the three sub-samples using least squares. α_1, β_1 are estimated using all values of p_t that correspond to $Z_{t-1} < c_1$ and α_3, β_3 are estimated using all values of p_t that correspond to $Z_{t-1} > c_2$.

We calculate the sum of squared residuals, $\sum_{t=2}^T \epsilon_t^2 = \sum_{t=2}^T (\Delta p_t - \psi_{t-1} - (\phi_{t-1} - 1)p_{t-1})^2$ for each sub-sample in order to obtain the parameter values and the corresponding diagnostics. In the next iteration of the algorithm we keep c_1 fixed at m_1 and change the value of c_2 to m_3 . The above process is repeated and a new set of parameter

estimates for $\alpha_1, \beta_1, \alpha_2, \alpha_3$ and β_3 is obtained. The procedure is repeated until $c_2 = m_n$, the last point on the grid. Following this first recursion, the recursive procedure is re-started by altering the value taken by c_1 to m_2 and c_2 to m_3 . The above double recursion is repeated until $c_1 = m_{n-1}$ and $c_2 = m_n$, which gives us our final parameter estimates. The values of c_1 and c_2 and the corresponding $\alpha_1, \beta_1, \alpha_2, \alpha_3$ and β_3 that minimize the sum of squared residuals are selected and parameters estimated along with their asymptotic standard errors.

Once we have found the switching parameters and the corresponding thresholds for the trigger variable, we can measure market efficiency by analysing the time the asset price spends in each state. Using the thresholds, we can divide the data into efficient periods (corresponding to periods when $\beta_j = 0$) and inefficient periods (corresponding to periods when statistically $\beta_j \neq 0$). If we have a total sample size n with n_1 efficient and n_2 inefficient periods, we can argue that the proportion of time the asset market is efficient is $\frac{n_1}{n}$ and the proportion of time it is inefficient is $\frac{n_2}{n}$.

Note, that if the c_1 and c_2 that minimize the sum of squared residuals are close, it implies that markets are rarely fully efficient (provided that the auto-regressive parameters for the sub-samples are significantly different from 0). Following convention from Section 2, α_1, β_1 indicate the first state, α_2, β_2 the second state and so on. If the series contains one mean reverting, one efficient and one explosive state we should find that $\beta_1 < 0$, $\beta_2 = 0$ and $\beta_3 > 0$ or that $\beta_1 > 0$ and $\beta_3 < 0$ given how the recursive procedure and the thresholds operate (it also depends on the relationship between Z_t and the parameters as well as the value of the drift parameters).

We neither need to restrict the number of states nor the values that the parameter β_j will take (except the β_j corresponding to the efficient state). Thus, we may observe multiple mean reverting or explosive states if we consider more than three states. We restrict our empirical application to the three state cases for tractability. Increasing the number of states beyond three will increase computational time substantially. We provide an illustrative example in the sub-section below.

We would like to note that the empirical example is illustrative in nature and is primarily aimed at showing how the methodology should be used in empirical work. To be more rigorous any empirical application aiming to derive firm and robust conclusions on market efficiency will need to justify the use of a trigger and also be specific with model selection (i. e. whether to use a specification with or without a switching or constant drift term). Since our example is illustrative, we provide results for all model specifications discussed in Section 2 and comment on how model specification changes the efficiency metric.

While we do not touch upon forecasting and predictability of states in this article, we would like to make some observations. Firstly, if the exogenous trigger is i.i.d, investors or the public will not be able to predict when the switches happen as the switching probabilities will not change period on period, if it is not a leading indicator of a regime change. However, when the exogenous trigger is a Markov chain, stakeholders may be able to estimate the probability of a particular asset market being in a certain state which would enable them to make more informed decisions. In either case the exogenous variable must be specified a prior. Secondly, the exogenous trigger may have predictive power if it is a leading indicator of the regime; investors will be able to anticipate a regime-change with some probability if the exogenous trigger approaches regime-switching thresholds and if its probability distribution is known. As we note below, the exogenous triggers that we use for our empirical application are more likely to be Markovian. We also note that in order to forecast the asset price beyond 1 period, we would also need to forecast the exogenous trigger. Work on unobserved state-switching and corresponding efficiency metrics has also been considered and is currently in progress.

Our procedure allows us to say with a chosen probability (we chose a 95 % confidence interval) whether the asset market under question is inefficient at a particular point in time. It may be argued that the transition from efficiency to inefficiency is smoother and thus, a smooth transition model may be more viable.³ While the argument is valid, a different definition of efficiency will be required if a smooth transition model is considered and we may have to arbitrarily identify regions which we classify as efficient or inefficient, based on the steady-state or dynamic properties of the trigger variable. Here, we have identified a conventional criterion (a random walk process) to classify an asset market as efficient. Secondly, when we refer to inefficiency in the context of our article, we are aware that the asset market is inefficient in a probabilistic sense, i. e. we are 90 % confident that an asset market may be inefficient. Thus, our assertions are not absolute. The empirical application below will clarify how we intend the methodology to be used.

4.2 Empirical Application

For our empirical study we use the S&P 500 and FTSE 100 indices. Data for the two indices were obtained from Google Finance and Bloomberg. The setting in Section 2 suggests that the states are triggered by an exogenous variable. We use the University of Michigan's Index of Consumer Sentiment as our exogenous trigger (MCSI henceforth). We also employ VIX as an additional exogenous trigger and compare the results. The lower bound

on the MCSI is 51 and the upper bound is set at 112 so that the regression in eq. (22) can be estimated; for the VIX the lower bound is 10 and the upper bound is 70. We use a 200 point grid on both exogenous triggers.

The MCSI is a monthly survey collected by the University of Michigan. It asks questions on personal finance and economic trends of individuals and households through telephonic interviews. Survey respondents are representative of the American population and each month more than 500 interviews are conducted. More information on the Survey and sample design is available on the MCSI website. The VIX on the other hand is a measure of implied volatility for the S&P500 derived from a variety of different options available on the S&P500. It seeks to estimate future volatility in the index and is therefore representative of investors' view on the future direction of the stock market index.

Data for the S&P500 index and the MCSI were obtained from January 1978 to June 2015. The underlying assumption required for the MCSI and VIX to be valid trigger variables in this setting is that contemporaneous and one-period lagged values of the two indices do not impact either the MCSI or the VIX. For the VIX, data start from 1st February 1990, but are available at a daily frequency. Thus, for VIX regressions we use daily data from the 1st February 1990 up to 30 June 2015 for both the S&P500 and the FTSE100 Indices.

While data for the MCSI are available from 1964, the survey was initially collected twice a year and the index only becomes a monthly index in 1978. The FTSE 100 index on the other hand is used from its inception in 1984. For the MCSI regressions we use monthly data. The use of MCSI and VIX as trigger variables for FTSE 100 is justified by the strong correlation between the S&P 500 and FTSE 100 indices. The MCSI and VIX do not appear to be independently distributed (first order autocorrelation > 0.9 for MCSI and 0.98 for VIX); thus, they are closer to being Markovian trigger variables.

Although we have not explicitly calculated moments for the case when the trigger variable is Markovian, we refer the interested reader to Knight and Satchell (2011), which discusses these results for the two-state case with a constant drift. We did test with a variety of other potential trigger variables but they failed basic Granger causality tests and thus, cannot be considered exogenous to the indices. Selecting a Markovian variable does not impact our estimators or our estimation strategy; however, we will not be able to use the formulae derived in Section 2 to calculate the mean and variance of our assets. Nevertheless, the formulae offer valuable insights as we will note in our discussion. We first discuss the results using MCSI as the exogenous trigger.

High consumer sentiment, i. e. a positive outlook towards personal finance and general business environment in the country is reflected through a high value of the index. On the other hand low consumer sentiment is reflected as a lower value. We posit that high consumer sentiment that persists for long periods is indicative of explosiveness or bubbles, i. e. if consumers have a very positive outlook they are likely to invest in assets and if a large number of consumers enter the asset markets or in this case the stock market the increase in demand could lead to a switch from efficiency to explosiveness.

Similarly when lower values persist we posit that the market is correcting itself and we get mean reverting behaviour. Figure 1 shows the MCSI and the log of the S&P500 index indicating how the MCSI varies with the log of S&P500 index. We see a spike in consumer sentiment in the run up to the dot com bubble. A similar increase is seen near the 2008–2009 financial crisis. Mean reverting behaviour is observed after the 1979 oil crisis as well as in the aftermath of the financial crisis. While the MCSI may not always respond contemporaneously to movements in the S&P500 index, it nevertheless acts as a valuable trigger variable for our illustrative example.

Figure 1: S&P 500 index and the Michigan consumer sentiment index.

We estimate the autoregressive and drift parameters in eq. (22) using 3 states for the S&P500 & FTSE100 indices respectively and use the MCSI as the trigger variable. The dependant variable in the regression is asset returns instead of log prices in order to ensure consistency of standard errors. We could have used log levels instead of returns but returns are more intuitive; secondly, using levels we would find some non-stationary states. Using the return formulation also ensures that the criterion for the existence of a stationary distribution, specified in Section 2, is satisfied.

We estimate the model with and without the switching drift term ψ_{t-1} . When we do use a drift term, we report results with both a switching drift term, i. e. the drift changes in each state and a constant drift term, i. e. the drift does not change across states. The most commonly employed specification in related literature is that with a constant drift. Our aim is to find thresholds c_1 and c_2 for the MCSI that minimize the residual sum of squares for the threshold auto-regression which in turn also yield the parameter estimates for the 2 inefficient states.

Note, that we do not impose any restrictions on the parameters of the other two states; both states may be mean reverting or explosive. The only restriction imposed is $c_1 > c_2$. We use the procedure outlined above to estimate the threshold, c_1 and c_2 . 51 and 112 represent the extreme values taken by the trigger variable, MCSI in this case. For our assets these values are reported in Table 7 below and include results for both stock market indices with and without a (switching/constant) drift term. The table also notes the time the stock-market index is in each of the 3-states; this will allow us to comment on the proportion of time each of the stock-market indices is efficient or inefficient.

Table 7: Non-linear least squares regression results (with standard errors) – with MCSI as the exogenous trigger.

Index	α_1 (SE)	β_1 (n_1) (SE)	α_2 (SE)	β_2 (n_2) (SE)	α_3 (SE)	β_3 (n_3) (SE)	c_1	c_2
S&P500 w/mov- ing drift	0.1809 (0.127)	-0.037(2 %) (0.0203)	0.0089 (0.002)	0(95 %) N/A	1.6368 (0.4359)	-0.230 (3 %) (0.0606)	59.7	107.5
S&P500 w/const. drift	0.0212 (0.0055)	-0.012 (2 %) (0.0022)	0 N/A	0(11 %) N/A	0 N/A	-0.002 (87 %) (0.0009)	59.6	68.7
S&P500 w'out drift	0 N/A	-0.009 (2 %) (0.0038)	0 N/A	0(2 %) N/A	0 N/A	0.001 (96 %) (0.0003)	59.7	62.2
FTSE100 w/mov- ing drift	1.157 (0.820)	-0.143(2 %) (0.0966)	0.0061 (0.0024)	0(92 %) N/A	2.406 (0.575)	-0.276 (6 %) (0.0661)	59.6	105.7
FTSE100 w/const. drift	0.0302 (0.0088)	-0.010 (2 %) (0.0022)	0 N/A	0(6 %) N/A	0 N/A	-0.003 (92 %) (0.0011)	59.6	68.7
FTSE100 w'out drift	0 N/A	-0.007 (2 %) (0.0020)	0 N/A	0(29 %) N/A	0 N/A	0.001 (69 %) (0.0003)	59.6	82.6

n_1 , n_2 and n_3 are percentages for which the respective index is in that state. SE represents standard errors. Entries in bold are statistically significant at the 5% level.

We postulate that $\beta_1 < 0$ if MCSI is low and $\beta_3 > 0$ when MCSI is high; the postulated relationship will vary based on our choice of trigger variables). Columns (8) and (9) in Table 7 report the thresholds corresponding to the minimum sum of squared residuals. Note that for the specification with drift, the intercept is also switching and as illustrated in Section 2, this changes the mean and variance of the process (if they exist) significantly.

Since we use different model specifications for this example, it is no surprise that the results in Table 7 present a mixed picture. When we consider the case of a shifting drift term in addition to a shifting slope coefficient, a lot of the variance in the series is captured by the shifting drift term. Inclusion of a moving drift substantially reduces the impact of the moving slope terms and we find no more than 32 observations in non-efficient regimes (20 for S&P500 and 32 for FTSE100). This amounts to 5 % inefficiency for the S&P500 and 8 % for the FTSE100. For both indices we also note that the first state is not statistically significant so we may not classify that state as inefficient. We also find little evidence of explosive behaviour due to the slopes.

Thus, explosive and mean reverting episodes under a model with a moving drift are primarily caused by the change in drift. Note that the drift terms are larger in magnitude and appear farther apart which implies that they have a higher variance. As per our formulae in Section 2, a higher variance of the drift parameter leads to a higher variance of the series. We find that the drift term in the explosive regime is statistically significant and greater than the drift term in other regimes. Thus, under this specification explosiveness in the S&P500 and FTSE100 indices is due to a temporary increase in average returns. With a constant drift term, both indices appear stationary and do not have explosive regimes.

One way to systematically beat the market in such a situation will be through predicting when the switches will occur provided that investors are aware of what state they are in as soon as the switch has occurred (and thereby becomes a part of the information set). Thus, we are referring to efficiency in a broader sense. In the conventional auto-regressive sense, a market is said to be efficient if the auto-regressive parameter is 1, i. e. the process is a random walk so that the only change in asset returns is due to unpredictable factors and no gains can be made based on the existing information set. In the threshold auto-regressive case, in addition to the restriction on the auto-regressive parameter we would also require the state switches be unpredictable; although once the switch occurs everyone becomes aware of it. Thus, the information set will also include information about what state the exogenous trigger is in. If markets are weak form efficient all rational investors will find out about the switch at the same time although they may not know when the switch may occur.

For the specification without a drift the results are closer to the behaviour observed in the simulations, i. e. we observe three states although the deviation from efficiency is very small. When $\beta_3 > 0$, i. e. we are in the explosive or bubble regime, we observe additional annualized gains of only 1.5 % in the S&P500 index and 1.1 % in the FTSE100 index. It may be argued that the additional annualized gains being captured by the parameter are in fact accounting for the missing drift term. For models with a switching drift, the criterion for a steady-state stationary distribution is trivially satisfied as for both FTSE100 and S&P500 we do not find an explosive slope coefficient. Figure 2 shows the areas that fall under the different states under this specification for log prices based on the thresholds estimated by the procedure outlined above.

Figure 2: S&P 500 index: Results for non-linear least squares with a switching drift with MCSI as the exogenous trigger.

We note that the first state corresponds to periods of relative slow down, i. e. in the aftermath of the 2nd oil price crisis, in the immediate aftermath of the financial crisis and in late 2011 when fears of a double dip recession abounded. The other non-efficient state occurs in the run up to the East Asian financial crisis and the dot-com bubble when consumer sentiment was at an all-time high. Our grid-search results do not indicate a deviation from a random walk during the financial crisis. The area under the non-random walk states has been shaded (blue for mean-reverting and red for explosive). Figure 3 shows similar results for the FTSE100 index. Note that apart from the dot-com bubble period in early 2000, 1998 is identified as a period of explosiveness for both indices. Both indices attained historical highs in the 1998 which is reflected in consumer sentiment.

Figure 3: FTSE100 index: Results for non-linear least squares with a switching drift with MCSI as the exogenous trigger.

By regressing log prices on their lags instead of returns (i. e. add 1 to each coefficient estimated in Section 3), we can calculate the value of the criterion function specified in Section 2 which allows us to comment on whether the series has a stationary distribution. For the case without a drift, the value of the criterion for the S&P500 and FTSE100 is 0.0010 and 0.0005 respectively (the criterion in this case is calculated as $\sum_{i=1}^3 \hat{\pi}_j(\hat{B}_j+1)$). Neither of the two indices satisfies the criterion for a steady-state stationary distribution under specifications without a drift. This indicates that any test for explosiveness that either assumes a constant drift term without shifting slope coefficients (not reported) or that drop the drift term are more likely to find the criterion violated for the S&P500 and FTSE100 indices.

As mentioned before, if MCSI was an independent and identically distributed variable we would be able to use our formulae from Section 2 and be able to calculate the mean and variance for both the S&P500 and the FTSE100 series when they are estimated using a threshold auto-regression. This will have allowed us to compute metrics such as Sharpe ratios enabling us to comment further on market efficiency and investor behaviour. Since MCSI is closer to a Markovian variable we are unable to use the formulae derived earlier. However, our results do allow us to compare efficiency across the two markets. In the following discussion when we talk about inefficient states we are referring to the proportion of time spent by each index in a state that is statistically significantly different from the random-walk.

When specifications with a drift are considered, the FTSE100 index appears more inefficient than the S&P500 index. The S&P500 index is inefficient for 2 % of the time with the switching drift specification and 86 % of the time with a constant drift. In contrast the FTSE 100 index is inefficient for 6.7 % of the time under the switching drift specification and 93 % of the time under the constant drift specification. On the other hand when no drift is included, the S&P500 appears mostly inefficient (95.5 %) compared to the FTSE 100 (68.7 %). If we compare similar periods, i. e. from 1981 onwards, the results remain robust. The mixed results do not offer a clear answer as to which market appears more inefficient; nevertheless the methodology is applicable to other assets. If an asset appears to spend more time in inefficient states under all specifications compared to another we will be able to conclude that the market for that particular asset is inefficient more often. Our results using VIX as the exogenous trigger shed further light on the efficiency question. Table 8 reports the results for when VIX is used as the exogenous trigger.

Table 8: Non-linear least squares regression results (with standard errors) – with VIX as the exogenous trigger.

Index	$\alpha_1(\text{SE})$	$\beta_1(n_1)$ (SE)	$\alpha_2(\text{SE})$	$\beta_2(n_2)$ (SE)	$\alpha_3(\text{SE})$	$\beta_3(n_3)$ (SE)	c_1	c_2
-------	-----------------------	------------------------	-----------------------	------------------------	-----------------------	------------------------	-------	-------

S&P500 w/mov- ing drift	0.0029 (0.002)	0.0(99 %) (0.0004)	-0.041 (0.019)	0(0.1 %) N/A	1.822 (0.818)	-0.267(1 %) (0.121)	53.72	55.52
S&P500 w/const. drift	-0.0043 (0.001)	0.001 (99 %) (0.0001)	0 N/A	0(0.9 %) N/A	0 N/A	0.006 (0.1 %) (0.0005)	47.39	67.89
S&P500 w/out drift	0 N/A	0.0 (99 %) (0.00002)	0 N/A	0(0.9 %) N/A	0 N/A	0.005 (0.1 %) (0.0021)	47.39	67.89
FTSE100 w/mov- ing drift	0.007 (0.003)	-0.001(99 %) (0.0004)	-0.032 (0.011)	0(0.1 %) N/A	2.194 (1.173)	-0.262(0.9 %) (0.141)	54.6	57.6
FTSE100 w/const. drift	-0.010 (0.002)	0.001 (99 %) (0.0002)	0 N/A	0(0.5 %) N/A	0 N/A	0.003 (0.5 %) (0.0003)	47.7	58.5
FTSE100 w/out drift	0 N/A	0.000 (99 %) (0.0)	0 N/A	0(0.7 %) N/A	0 N/A	0.002(0.3 %) (0.001)	45.9	63.1

n_1 , n_2 and n_3 are percentages for which the respective index is in that state. SE represents standard errors. Entries in bold are statistically significant at the 5% level.

Before we discuss the regression results using VIX in detail, it is important to highlight the differences in the two exogenous triggers. While data for the MCSI are only available at a monthly frequency, VIX is reported daily; this allows us to estimate the regression using daily returns thereby allowing us to measure market efficiency at a higher frequency. Thus, VIX is measuring short-term market efficiency while with MCSI we were able to measure medium-term market efficiency. However, since the data for VIX are only available from 1990, we are unable to include observations from the oil price crises and large stock market downturns in the 1970s and 1980s.

Figure 4 plots the S&P500 and the VIX. The pattern suggests that periods of high volatility, marked by a high value of the VIX, tend to correspond to low values of the S&P500 index. Thus, we detect high volatility in the market whenever the S&P500 index is facing a downturn. No such pattern is immediately visible for the opposite scenario. High values of the S&P500 do not always correspond to low values of the VIX. It appears that while the VIX may be a good potential trigger for the stationary or mean reverting regime, it may not have the necessary variance for the explosive state. Also, as we pointed out initially, the VIX is a measure of investment sentiment while the MCSI is a measure of consumer sentiment. This is another reason why the two results can be different; the information sets of consumers and investors may be quite different. Since efficiency is defined with respect to an information set, it is no surprise that the results vary, although we do note some common trends. The prevailing market price that we observe results from the different demand functions of consumers and professional investors. Both indices contain valuable information and it is advisable to consider results from both sets of regression results.

Figure 4: S&P 500 index and VIX – Daily Data.

While results from the MCSI were mixed, the ones obtained when VIX is used as the exogenous trigger are largely in favour of efficient markets. As before, we restricted state 2 to be the efficient state; however, our estimates with VIX contain significant evidence that there are only 2 as opposed to 3 states. β_1 is always close to 0 (to 3 decimal places or more) or is statistically insignificant, irrespective of the model specification. It is difficult to distinguish this state from the efficient state. β_3 on the other hand does indicate the presence of inefficient periods.

Although the proportion of inefficient observations is small (1 % or less), it accounts for around 50 periods/-days. Figure 5 indicates that these tend to occur in 2008, in the aftermath of the financial crisis when volatility as recorded by VIX was very high. These periods also coincide with the periods of mean reversion detected when MCSI was employed as the exogenous variable. We note, however, that the procedure does not indicate the presence of explosive states in this case. It neither detected the dotcom bubble, nor the period in the run up to the financial crisis. The MCSI on the other hand, did flag the dotcom bubble period as inefficient. As mentioned before, VIX does not appear to be a good predictor of the explosive regime although it does do much better for indicating the onset of mean reversion. The results for FTSE100 looked similar and are reported in Figure 6.

Figure 5: S&P 500 index: Results for non-linear least squares with a switching drift with VIX as the exogenous trigger.

Figure 6: FTSE100 index: Results for non-linear least squares with a switching drift with VIX as the exogenous trigger.

Taking both the MCSI and VIX results together, our empirical results supplement our findings in Sections 2 and 3. Explosiveness is more likely to be detected in asset price series where the criterion function is violated and the variance of the switching slope parameters is large (i. e. there are many regimes or the regimes are much farther apart). Inclusion of a switching drift term may explain most of the explosiveness and may make the price series appear efficient. Another way of analysing results could be through comparing the different series.

The set of results reported above also depend on the selection of the trigger variable as we have noted. Finding an appropriate trigger variable that may indicate switches in regimes is non-trivial in practice and will require rigorous theoretical, empirical or experimental basis so that regime identification criteria can be appropriately set. We also need to consider the time period for which we are measuring market efficiency. As we have seen above, markets may appear to be more efficient when higher frequency data are considered.

While we primarily relied on Granger causality tests to determine appropriate lags for the trigger variables and for determining exogeneity, it could be argued that further testing on exogeneity, including testing for non-linear relationships, may be warranted. This requires judgement on the part of the researcher, the specific asset price being considered and the relationship between the asset price and the trigger variable. We did consider other suitable triggers including the fear and greed index and the economic policy uncertainty index but these were either not exogenous or did not include sufficient data. Furthermore, it is also possible to consider a combination of exogenous triggers which would combine elements of investor and consumer sentiment. We leave this exercise for future research as our example was aimed as an illustration of the outlined methodology. Additionally, the researcher also needs to consider the number of states to be used. A price series could exhibit multiple explosive or mean reverting states.

Our methodology can thus work in practice for different asset markets. We have shown with our example that the methodology may be used to identify the incidence of market efficiency for different assets while also highlighting the importance of model-specification when testing for market efficiency. Model specification is not important just in terms of estimation but completely changes the theoretical meaning of the results. Once a researcher has identified an appropriate specification for an asset price or return based on either technical analysis or through solving a structural model, our methodology will allow her to comment on market efficiency for that asset. However, irrespective of the specification, the results may still be used to compare different markets and identify which markets are more efficient for the given information set.

5 Conclusion

In this article we have extended existing general conditions that need to be satisfied by threshold autoregressive models in order to have a steady-state distribution and for a mean and variance to exist and have provided formulae for them. The results have been extended to include the case where a switching drift is included in

addition to a switching coefficient parameter. We have also considered the case of models with and without drift, specifically considering an i.i.d variable as an exogenous regime switching trigger. We believe that the results can be extended to other types of trigger variables, such as Markovian trigger variables, although we do not evaluate analytical expressions for such cases in this article. We have shown that when a steady-state distribution does exist for a TAR(1) model with a switching drift term, the variance depends on the variance of the error term, the variance of the drift parameters, the variance of the coefficient parameters as well as the covariance between the drift and coefficient parameters.

A simulation study is carried out to evaluate the power of the GSADF bubble detection test under conditions where a steady-state distribution may not exist. Our simulation study has shown that if a series has a steady-state distribution, bubbles may be more difficult to detect. We further note that the power of such tests increases with the variance of the regime parameters, i. e. the farther apart the parameters are from unity, the higher is the power. These results enable us to understand why bubble tests may fail to detect explosiveness even though it may be locally present in a series.

We also provide a methodology that may be used in practice to estimate TAR models with exogenous i.i.d trigger variables or Markovian triggers. Using non-linear least squares we find threshold levels that minimize the sum of squared residuals. Our results indicate that model specification is critical when analysing weak form market efficiency using price series. Series that may appear to exhibit inefficiency when a financial analyst assumes no drift will appear efficient when a regime-switching drift term is used which highlights the need for carefully considering model specification prior to estimation. The methodology allows us to state how often a market is inefficient instead of making a binary classification.

Our empirical results vindicate our theoretical findings, i. e. the variance of a price process depends not only on the regime-switching coefficients but also on the regime-switching drift term. Additionally we also extend the notion of efficiency to include the predictability of state switching, i. e. a market is more efficient if state switching is unpredictable. We believe this methodology is applicable to a variety of different markets including commodity and foreign exchange markets. The methodology also allows us to compare different asset markets and comment on their efficiency relative to one another. We found evidence that the S&P500 and FTSE100 are mostly efficient under the most generalized (switching-drift) specification although there are pockets of data where inefficiency is noted.

Multiple avenues of further research open up as a result of this article. The theoretical results may be further expanded to include a Markovian process as the regime switching variable or consider results for TAR (p) models. In addition, our simulation results highlight one limitation of the GSADF test and also the need for having bubble tests that may be applied locally or on sub-samples as considering the full price process may make detection difficult. Although we have used only non-linear least squares for our illustration, other methodologies including Markov-switching regressions may also be applied.

Our empirical methodology opens up a wide array of possibilities for financial analysts and econometricians alike who may be interested in understanding market efficiency in different markets. In particular, trend-following commodity trading analysts could use such procedures to determine which markets are efficient most of the time and avoid trading dynamically in them unless their mean-variance properties make them intrinsically appealing. This article also raises the question of what exogenous trigger variables may be most appropriate for a particular asset market. Finally, identification of a suitable trigger may have policy implications, i. e. the government may try to influence expectations about these variables in order to move asset markets towards efficiency although we do caution against using tenuous relationships to draw policy conclusions.

Notes

- 1 $Var(\psi_{t-1})$, $Var(\phi_{t-1})$ and $E(\phi_{t-1}^2)$ are as before. $Cov(\psi_{t-1}, \phi_{t-1}) = \sum_{j=1}^k \pi_j \alpha_j \beta_j - \left(\sum_{j=1}^k \pi_j \alpha_j \right) \left(\sum_{j=1}^k \pi_j \beta_j \right)$
- 2 We would like to thank a referee for the discussion in this paragraph.
- 3 A smooth transition Autoregressive model (STAR) is written as:

$$p_t = (\psi_1 + \phi_1 p_{t-1}) F(Z_{t-1}) + (\psi_2 + \phi_2 p_{t-1}) (1 - F(Z_{t-1}))$$

where Z_{t-1} is the exogenous trigger and $F(Z_{t-1})$ is the cumulative distribution of the exogenous trigger. We have outlined the two state cases, but this can potentially have multiple state. $F(Z_{t-1})$ can be a smooth function such as a logistic function. To define efficiency in this context, we will need to define a region of the density function $f(Z_{t-1})$ which we can classify as the efficient region. (Tong 2011)

References

Bloomberg, L.P. S&P500, FTSE100 and VIX daily data series 01/02/1990 to 30/06/2015 2016 from Bloomberg database. Jan.132016. Retrieved on.

- De Bondt, W.F.M., and R. Thaler. 1985. "Does the Stock Market Overreact?" *The Journal of Finance* 40 (3): 793–805.
- Evans, G.W. 1991. "Pitfalls in Testing for Explosive Bubbles in Asset Prices." *American Economic Review* 81: 1189–1214.
- Fama, E.F. 1965. "The Behaviour of Stock-Market Prices." *Journal of Business* 38 (1): 34–105.
- Goldfeld, S.M., and R.E. Quandt. 1973. "A Markov Model for Switching Regressions." *Journal of Econometrics* 1: 3–16.
- "Google Finance". *S&P 500 index*. Retrieved on 03/09/2015 <https://www.google.co.uk/finance>.
- "Google Finance". *FTSE 100 Index*. Retrieved on 03/09/2015 <https://www.google.co.uk/finance>
- Hall, S.G., Z. Psaradakis, and M. Sola. 1999. "Detecting Periodically Collapsing Bubbles: A Markov-Switching Unit Root Test." *Journal of Applied Econometrics* 14 (2): 143–154.
- Hamilton, J.D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica: Journal of the Econometric Society* 57: 319–352.
- Hamilton, J.D. 1990. "Analysis of Time Series Subject to Regime Changes." *Journal of Econometrics* 45: 39–70.
- Knight, J., and S.E. Satchell. 2011. "Some New Results for Threshold AR(1) Models." *Journal of Time Series Econometrics* 3: 1–42.
- Knight, J., S.E. Satchell, and N. Srivastava. 2014. "Steady State Distributions for Models of Locally Explosive Regimes: Existence and Econometric Implications." *Economic Modelling* 41: 281–288.
- Marsh, T.A., and R.C. Merton. 1986. "Divident Variability and Variance Bounds Tests for the Rationality of Stock Market Prices." *The American Economic Review* 76 (3): 483–498.
- Phillips, P.C.B., S. Shi, and J. Yu. 2013. "Testing for Multiple Bubbles: Historical Episodes of Exuberance and Collapse in the S&P 500." In *Cowles Foundation Discussion Papers 1914*. Cowles Foundation for Research in Economics Yale University.
- Quinn, B.G. 1982. "A Note on the Existence of Strictly Stationary Solutions to Bilinear Equations." *Journal of Time Series Analysis* 3: 249–252.
- Samuelson, P.A. 1965. "Proof that Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review* 6 (2): 41–49.
- Sewel, M. 2011. *A History of the Efficient Market Hypothesis*. Research Note: UCL Department of Computer Science.
- Shiller, R.J. 2000. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.
- Tong, H. 1978. "On a Threshold Model." In Chen, C.H. (Ed.) *Pattern Recognition and Signal Processing*. Amsterdam: Sijthoff and Noordhoff.
- Tong, H. 2011. "Threshold Models in Time Series Analysis – 30 Years on (With Discussions by P. Whittle, M. Rosenblatt, B. E. Hansen, P. Brockwell, N. I. Samia and F. Battaglia)." *Statistics and Its Interface* 4: 107–118.
- University of Michigan. *The Index of Consumer Sentiment*. Retrieved on 03/09/2015 <https://data.sca.isr.umich.edu/tables.php>.