Testing explosive bubbles with time-varying volatility

David I. Harvey*, Stephen J. Leybourne and Yang Zu
University of Nottingham

July 2018

Abstract

This paper considers the problem of testing for an explosive bubble in financial data in the presence of time-varying volatility. We propose a weighted least squares-based variant of the Phillips, Wu and Yu (2011) test for explosive autoregressive behaviour. We find that such an approach has appealing asymptotic power properties, with the potential to deliver substantially greater power than the established OLS-based approach for many volatility and bubble settings. Given that the OLS-based test can outperform the weighted least squares-based test for other volatility and bubble specifications, we also suggested a union of rejections procedure that succeeds in capturing the better power available from the two constituent tests for a given alternative. Our approach involves a nonparametric kernel-based volatility function estimator for computation of the weighted least squares-based statistic, together with the use of a wild bootstrap procedure applied jointly to both individual tests, delivering a powerful testing procedure that is asymptotically size-robust to a wide range of time-varying volatility specifications.

Keywords: Rational bubble; Explosive autoregression; Time-varying volatility; Weighted least squares; Right-tailed unit root testing.

JEL Classification: C12; C14, C22.

*Corresponding author. Address: School of Economics, University of Nottingham, University Park, Nottingham, NG7 2RD, UK. Email: dave.harvey@nottingham.ac.uk

1 Introduction

Empirical identification of explosive behaviour in financial asset price series is closely related to the study of rational bubbles, with a rational bubble deemed to have occurred if explosive characteristics are manifest in the time path of prices, but not for the dividends. Consequently, methods for testing for explosive time series behaviour have been a focus of much recent research. In a now seminal paper, Phillips et al. (2011) [PWY] model potential bubble behaviour using an explosive autoregressive specification, and suggest testing for such a property using the supremum of a sequence of forward recursive right-tailed Dickey-Fuller unit root tests. The test has been widely applied, and has also prompted the development of related test procedures, such as those of Homm and Breitung (2012) who consider a supremum of backward recursive Chow-type Dickey-Fuller statistics, and Phillips et al. (2015) who consider a double supremum of forward and backward recursive statistics.

While these original papers assumed constant unconditional volatility in the underlying error process, in practice time-varying, typically nonstationary, volatility is a well-known stylised fact observed in empirical financial data. Harvey et al. (2016) [HLST] demonstrate that the asymptotic null distribution of the PWY test depends on the nature of the volatility, and if the test is implemented using critical values derived under a homoskedastic error assumption, size is not controlled for nonstationary volatility patterns. This raises the possibility of misleading inference when using the standard PWY test in the presence of time-varying volatility, with the potential for spurious identification of a bubble. HLST propose a wild bootstrap implementation of the PWY test, which ensures correct asymptotic size in the presence of time-varying volatility. This procedure also retains the same local asymptotic power as the original PWY test, if the latter were to be infeasibly size-corrected to account for the volatility pattern.

The PWY test and its related variants discussed above are all fundamentally based on OLS estimation of the underlying autoregression. In the context of time-varying volatility in the model errors, it is natural to consider whether a GLS-type transformation can deliver a more powerful testing approach. In this paper, we focus on this possibility, developing tests based on Dickey-Fuller unit root statistics derived from a weighted least squares [WLS] transformation of the model, along the lines of the approach considered by Boswijk and Zu (2015) in the context of full-sample testing for a unit root against a left-tailed stationary alternative.[1] Specifically, we propose a PWY-type test, based on the supremum of a sequence of forward recursive statistics, but where the OLS-based Dickey-Fuller statistic is replaced by a WLS-based equivalent.

We begin by treating the volatility path as known, and demonstrate that our WLS-

[1] See also Xu and Phillips (2008) and Xu and Yang (2015) for using similar kernel-type GLS corrections for nonstationary volatility in time series models.

based test can offer substantially greater local asymptotic power than the PWY procedure for many volatility patterns and bubble specifications. This positive result suggests that the WLS-based approach merits development, offering the potential to improve our capacity to detect explosive autoregressive behaviour. Despite these potential large gains in local asymptotic power, we find that for certain volatility and bubble settings, the power rankings of the tests can be reversed, hence the PWY test can still offer a valuable role in bubble detection. In order to capture the relative power advantages of both tests across different volatility patterns, we then proceed to consider a union of rejections approach (cf. Harvey et al. (2009)), whereby the null is rejected in favour of explosive behaviour if either the WLS-based test or the PWY test rejects, subject to a scaling applied to the asymptotic critical values to ensure correct size of the composite procedure. We find that the union of rejections testing strategy performs very well across the range of volatility and bubble specifications that we consider.

Calculation of the WLS-based test statistic requires the volatility at each point in time, which is of course unknown in practice. To render the statistic feasible, we employ a nonparametric kernel-based volatility function estimator, and we show that substitution of this estimator in place of the true volatility path results in a statistic with the same limiting null and local alternative distributions as for the infeasible statistic. In common with the PWY statistic, our feasible WLS-based statistic has a limiting null distribution that depends on the volatility path. Following HLST, we suggest a wild bootstrap implementation of the test to achieve an asymptotically size-controlled procedure. For the union of rejections procedure, we then apply both the wild bootstrap HLST variant of PWY, and the wild bootstrap version of our WLS-based test. In order to ensure that the bootstrap-based union of rejections procedure is asymptotically correctly sized, we implement the wild bootstrap procedure to the two statistics *jointly*, and also calculate the required critical value scaling constant from the wild bootstrap algorithm. We demonstrate the asymptotic validity of the joint procedure, showing that the local asymptotic power profiles coincide with those for the infeasible case where the volatility is treated as known. Note that while we concentrate on introducing the techniques of this paper in the context of the PWY test, which is the prototype for the recent literature on recursive testing for bubbles, it should be noted that the methods we develop in principle apply more widely, e.g. to the double supremum extension by Phillips et al. (2015). We discuss this issue briefly in the final section of the paper.

The rest of the paper is organised as follows. Section 2 outlines the model and introduces the WLS-based approach. Here we establish its limit behaviour in the known volatility case, and compare its local asymptotic power to the PWY test. This section also introduces the union of rejections testing strategy and evaluates its large sample power behaviour. Nonparametric volatility function estimation and a feasible WLS-based statistic are discussed in section 3, while the wild bootstrap algorithm is developed in

section 4, and its asymptotic properties established. Finite sample properties of the tests are explored by Monte Carlo simulation in section 5, and an empirical example is provided in section 6 using FTSE and S&P 500 data. Section 7 discusses possible extensions of the current methodology and concludes the paper. Proofs of the results are given in the Appendix.

2 The bubble model and tests

2.1 The model

We will consider the following model with time-varying volatility for a time series $\{y_t\}$, $t = 1, ..., T$:

$$y_t = \mu + x_t \tag{1}$$
$$x_t = (1 + \rho_t)x_{t-1} + u_t, \quad t = 2, ..., T \tag{2}$$
$$u_t = \sigma_t \varepsilon_t \tag{3}$$

with $x_1 = O_p(1)$. We make the following assumptions regarding the error ε_t in (3) and the error standard deviation at time t, σ_t:

A1 $E(\varepsilon_t|\mathcal{F}_{t-1}) = 0$, $E(\varepsilon_t^2|\mathcal{F}_{t-1}) = 1$ with $\mathcal{F}_t$ the natural filtration generated by $\{u_s\}_{s\geq 1}$, and $E(\varepsilon_t^4) < \infty$.

A2 $\sigma_t = \sigma(t/T)$, where $\sigma(.)$ is a strictly positive function with $\sigma(.) \in D[0, 1]$, the space of right continuous with left limit (càdlàg) processes on $[0, 1]$.

Under Assumption A1, ε_t is a martingale difference sequence, and is hence conditionally first order uncorrelated, cf. Xu and Phillips (2008). Note that for empirical financial data, such an assumption is standard, where any admitted dependence is typically in the second moment (i.e. in the volatility). Assumption A2 implies that the innovation variance is non-stochastic, bounded and displays a countable number of jumps, cf. Cavaliere and Taylor (2007), allowing for a very wide range of volatility dynamics.

For the time-varying autoregressive parameter $(1 + \rho_t)$ in (2) we adopt the following specification:

$$\rho_t = \begin{cases} 0 & t = 2, \ldots, [\tau^* T] \\ c/T & t = [\tau^* T] + 1, \ldots, T \end{cases}$$

where $[.]$ denotes the integer part of its argument. When $c > 0$, y_t follows a unit root process up to time $[\tau^* T]$, after which point it displays locally explosive autoregressive behaviour over the remaining sample period $t = [\tau^* T] + 1, ..., T$. In the context of asset price behaviour, this setup models the case where prices follow the usual market

behaviour with unit root dynamics up to a particular point in the sample, after which a bubble originates and explosive behaviour is manifest. Extensions to the case where the bubble terminates in-sample (with or without some form of collapse) could easily be entertained; see, for example, Harvey et al. (2017).

To test for the presence of a bubble, we consider a null hypothesis $\mathcal{H}_0 : c = 0$ against the alternative $\mathcal{H}_1 : c > 0$. Under the null and local alternative, we can make use of the following invariance principle, which holds under Assumptions A1 and A2:

$$T^{-1/2} \sum_{t=2}^{\lfloor rT \rfloor} u_t \Rightarrow \int_0^r \sigma(s) dW(s)$$

where $\Rightarrow$ denotes weak convergence and $W(r)$ is a standard Brownian motion process.

2.2 A weighted least squares-based test

PWY propose a test for a bubble based on the supremum of recursive right-tailed Dickey-Fuller tests, based on OLS estimation. In view of the heteroskedasticity present in our model, it is natural to consider whether a GLS-type transformation can deliver a more powerful testing approach. Consequently, we now consider Dickey-Fuller t-statistics based on a WLS transformation of the model, initially for the infeasible case where the σ_t are assumed known. Considering first the underlying x_t process in (2), the transformed model can be written as

$$\frac{\Delta x_t}{\sigma_t} = \rho_t \frac{x_{t-1}}{\sigma_t} + \varepsilon_t, \qquad t = 2, \ldots, T. \tag{4}$$

Here, (4) is an infeasible homoskedastic regression model of $\{\Delta x_t/\sigma_t\}$ on $\{x_{t-1}/\sigma_t\}$, with coefficient ρ_t. Assuming knowledge of x_t and σ_t, a bubble test statistic could be constructed using a sequence of WLS-based Dickey-Fuller regressions, analogous to the PWY OLS-based test statistic. In practice we do not observe x_t, but $y_t = \mu + x_t$ as in (1). In order to achieve invariance to μ, we simply replace x_t in (4) by $\tilde{y}_t = y_t - y_1$, which is equivalent to GLS-demeaning of y_t in the sense of Elliott et al. (1996) using $\bar{\alpha} = 1$ in their notation, so that y_1 becomes the estimator of μ. Our infeasible test statistic that assumes knowledge of σ_t can then be written as

$$\text{supBZ} = \sup_{\tau \in [\tau_0, 1]} \text{BZ}_\tau \tag{5}$$

where BZ_τ denotes the Dickey-Fuller statistic calculated over the sub-sample $\{y_1, \ldots, y_{[\tau T]}\}$, that is

$$\text{BZ}_\tau = \frac{\sum_{t=2}^{[\tau T]} \Delta \tilde{y}_t \tilde{y}_{t-1} / \sigma_t^2}{\left(\sum_{t=2}^{[\tau T]} \tilde{y}_{t-1}^2 / \sigma_t^2 \right)^{1/2}}.$$

In (5), the minimum sample length admitted in the sequence of sub-sample regressions is $[\tau_0 T]$. Note that the full-sample statistic BZ_1 coincides with the infeasible test statistic considered in Boswijk and Zu (2015) in the context of left-tailed unit root testing against a stationary alternative.

The following theorem gives the limit distribution of supBZ:

Theorem 1. *Under $\mathcal{H}_1$ and Assumptions A1 and A2,*

$$\mathrm{supBZ} \Rightarrow \sup_{\tau\in[\tau_0,1]} L_c(\tau) = \mathbb{M}_c^{\mathrm{BZ}}$$

where

$$L_c(\tau) = \begin{cases} \frac{\int_0^\tau V_c(r)\mathrm{d}W(r)}{\left(\int_0^\tau V_c(r)^2\mathrm{d}r\right)^{1/2}} & \tau \leqslant \tau^* \\ \frac{\int_0^\tau V_c(r)\mathrm{d}W(r) + c\int_{\tau^*}^\tau V_c(r)^2\mathrm{d}r}{\left(\int_0^\tau V_c(r)^2\mathrm{d}r\right)^{1/2}} & \tau > \tau^* \end{cases}$$

with $V_c(r) = U_c(r)/\sigma(r)$ and

$$U_c(r) = \begin{cases} \int_0^r \sigma(s)\mathrm{d}W(s) & r \leqslant \tau^* \\ e^{c(r-\tau^*)}\int_0^{\tau^*}\sigma(s)\mathrm{d}W(s) + \int_{\tau^*}^r e^{c(r-s)}\sigma(s)\mathrm{d}W(s) & r > \tau^* \end{cases}.$$

Remark 1. The null limit distribution of supBZ is obtained from the result in Theorem 1 simply by setting $c = 0$, so that $\mathrm{supBZ} \Rightarrow \mathbb{M}_0^{\mathrm{BZ}}$. Of course, the limit distribution under both the null and local alternative depends on $\sigma(r)$, hence the critical values and local asymptotic power function will be contingent on the volatility pattern present in the innovations.

We now proceed to evaluate the local asymptotic power of the supBZ test, comparing it to the local power of an infeasibly size-corrected PWY test. The PWY statistic follows a similar form to supBZ, but is based on the supremum of the t-ratios associated with $\hat{\beta}_\tau$ in the fitted OLS regressions

$$\Delta y_t = \hat{\alpha}_\tau + \hat{\beta}_\tau y_{t-1} + \hat{e}_t, \qquad t = 2, \ldots, [\tau T]$$

for $\tau \in [\tau_0, 1]$. Denoting this statistic by supDF, then under $\mathcal{H}_1$ and Assumptions A1 and A2, it can be shown from the results in Harvey et al. (2016) that

$$\mathrm{supDF} \Rightarrow \sup_{\tau\in[\tau_0,1]} J_c(\tau) = \mathbb{M}_c^{\mathrm{DF}}$$

where

$$J_c(\tau) = \begin{cases} \frac{\int_0^\tau \tilde{U}_c(r)\mathrm{d}U_c(r)}{\left(\tau^{-1}\int_0^\tau \sigma(r)^2 dr \int_0^\tau \tilde{U}_c(r)^2\mathrm{d}r\right)^{1/2}} & \tau \leqslant \tau^* \\ \frac{\int_0^\tau \tilde{U}_c(r)\mathrm{d}U_c(r) + c\int_{\tau^*}^\tau \tilde{U}_c(r)^2\mathrm{d}r}{\left(\tau^{-1}\int_0^\tau \sigma(r)^2 dr \int_0^\tau \tilde{U}_c(r)^2\mathrm{d}r\right)^{1/2}} & \tau > \tau^* \end{cases}$$

with $\tilde{U}_c(r) = U_c(r) - \frac{1}{\tau}\int_0^\tau U_c(s)ds$, and where $U_c(r)$ is as defined in Theorem 1. The null limit distribution for supDF is obtained on setting $c = 0$, so that supDF $\Rightarrow \mathbb{M}_0^{\mathrm{DF}}$. As with the result for supBZ, the limit distribution under both the null and local alternative depends on $\sigma(r)$. For the purposes of comparing local power with supBZ, we treat the volatility path as known, and compute infeasibly size-adjusted local asymptotic powers for a given $\sigma(r)$.

For the local power simulations, we consider the following volatility specifications for $\sigma(r)$, where $\mathbb{I}(.)$ denotes the indicator function:

(a) Constant volatility: $\sigma(r) = 1 \quad \forall r$.

(b) Early upward shift: $\sigma(r) = 1 + 5\mathbb{I}(r \geq 0.3)$.

(c) Late upward shift: $\sigma(r) = 1 + 5\mathbb{I}(r \geq 0.8)$.

(d) Early downward shift: $\sigma(r) = 1 + 5\mathbb{I}(r < 0.3)$.

(e) Late downward shift: $\sigma(r) = 1 + 5\mathbb{I}(r < 0.8)$.

(f) Upward trend: $\sigma(r) = 1 + 5r$.

(g) Downward trend: $\sigma(r) = 6 - 5r$.

(h) Double shift: $\sigma(r) = 1 + 5\mathbb{I}(0.4 < r \leqslant 0.6)$.

Here, (a) is the benchmark homoskedastic case, while (b)-(e) specify one-off variance shifts, (h) a double shift, and (f)-(g) linearly trending variances.

In Figures 1-2, we plot the asymptotic local power functions of supBZ and supDF, simulating the limit distributions $\mathbb{M}_c^{\mathrm{BZ}}$ and $\mathbb{M}_c^{\mathrm{DF}}$ using 10,000 Monte Carlo replications, and approximating the Brownian motion processes in the limiting functionals using $NIID(0, 1)$ random variates, with the integrals approximated by normalized sums of 1,000 steps. Here and throughout the paper, we set $\tau_0 = 0.1$, as in PWY. Figures 1 and 2 report results for the bubble timings $\tau^* = 0.6$ and $\tau^* = 0.8$, respectively, under each volatility pattern, using a grid of c values from 0-8 in Figure 1 and 0-20 in Figure 2. For each test, power is evaluated using the 0.05-level null critical value appropriate for each volatility specification (i.e. from $\mathbb{M}_0^{\mathrm{BZ}}$ or $\mathbb{M}_0^{\mathrm{DF}}$), and therefore is infeasibly size-corrected.

Examining first Figure 1 ($\tau^* = 0.6$), we observe that in the homoskedastic case (a), there is relatively little difference in the local power profiles of supBZ and supDF. However, for most time-varying volatility specifications, supBZ demonstrates higher power across c than supDF, with the gains appearing particularly substantial for volatility specifications that incorporate a decrease in variance, i.e. (d), (e), (g) and (h). Large gains are also observed for an early upward shift in volatility (b), while gains are more modest for the late upward shift (c). For the upward trend (f), supBZ is slightly less

powerful than supDF. Turning to the later bubble timing of Figure 2 ($\tau^* = 0.8$), the power gains for supBZ over supDF remain evident in most cases, with the gains again being substantial in cases (b), (d), (e), (g) and (h). As in Figure 1, the tests have broadly similar power levels under homoskedasticity, while for an upward trend (f) and now also a late upward shift (c), we see that supDF outperforms supBZ, albeit to a lesser degree than in the cases where supBZ outperforms supDF.

2.3 A union of rejections testing procedure

As neither test is dominant across all volatility specifications, we can consider employing a *union of rejections* strategy along the lines of Harvey et al. (2009). These authors suggest such an approach in the context of combining inference from two unit root tests, one of which permits a linear trend in its deterministic specification, the other of which excludes the trend, the idea being to harness the better power of the two when the presence of a trend is uncertain. The same principle can be used here, combining inference from supBZ and supDF in the presence of uncertainty over the volatility specification, in an attempt to capitalize on the relative power advantages of each across different volatility patterns.

Specifically, denoting the asymptotic ξ level null critical values of supDF and supBZ (i.e. from $\mathbb{M}_0^{\mathrm{DF}}$and $\mathbb{M}_0^{\mathrm{BZ}}$) by $q_\xi^{\mathrm{DF}}$and q_ξ^{BZ}, respectively, a union of rejections strategy can be written as the decision rule

$$\text{Reject } \mathcal{H}_0 \quad \text{if} \quad \{\text{supDF} > \psi_\xi q_\xi^{\mathrm{DF}} \text{or } \text{supBZ} > \psi_\xi q_\xi^{\mathrm{BZ}}\}$$

where ψ_ξ is a scaling constant chosen so that this decision rule yields an asymptotic size of ξ under $\mathcal{H}_0$. Defining a single statistic $\mathcal{U}$ as

$$\mathcal{U} = \max\left(\text{supDF}, \frac{q_\xi^{\mathrm{DF}}}{q_\xi^{\mathrm{BZ}}}\,\text{supBZ}\right)$$

the decision rule is then equivalent to

$$\text{Reject } H_0 \quad \text{if} \quad \mathcal{U} > \psi_\xi q_\xi^{\mathrm{DF}}.$$

Using the asymptotic results of the previous section, an application of the continuous mapping theorem (CMT) establishes that

$$\mathcal{U} \Rightarrow \max\left(\mathbb{M}_c^{\mathrm{DF}}, \frac{q_\xi^{\mathrm{DF}}}{q_\xi^{\mathrm{BZ}}}\mathbb{M}_c^{\mathrm{BZ}}\right).$$

The scaling constant ψ_ξ can easily be determined from the limit distribution of $\mathcal{U}$ with $c = 0$, but there is actually no need to calculate it explicitly since, for a given ratio

$q_\xi^{\mathrm{DF}}/q_\xi^{\mathrm{BZ}}$, all we require is the critical value $\psi_\xi q_\xi^{\mathrm{DF}}$, which we denote $q_\xi^{\mathcal{U}}$. This can be obtained directly from the null limit distribution of $\mathcal{U}$. Finally, notice that since q_ξ^{DF} and q_ξ^{BZ} depend on the particular form of volatility specified by $\sigma(r)$, so too will the critical value $q_\xi^{\mathcal{U}}$. At this point therefore, $\mathcal{U}$, along with supDF and supBZ, is an infeasible testing procedure.

Along with the infeasible size-adjusted local asymptotic powers of supBZ and supDF, Figures 1 and 2 show the corresponding power of the union of rejections procedure $\mathcal{U}$. We see throughout that the power profile of $\mathcal{U}$ is always very similar to whichever profile of supDF and supBZ obtains the higher power. There is at worst only a small deficit compared with the better of the two, suggesting that the union procedure is performing well here.

Thus far we have considered only the large sample properties of an infeasible variant of supBZ that is based on knowledge of the volatility function σ_t. For any practical implementation, construction of supBZ will require estimation of σ_t. We address this issue in the next section.

3 Nonparametric kernel estimation of the volatility function and the feasible supBZ statistic

For estimation of σ_t^2 we employ a simple nonparametric kernel smoothing estimator of the form

$$\hat{\sigma}_t^2 = \frac{\sum_{i=2}^{T} K_h\left(\frac{i-t}{T}\right)(\Delta y_i)^2}{\sum_{i=2}^{T} K_h\left(\frac{i-t}{T}\right)} \tag{6}$$

where $K_h(s) = K(s/h)/h$ and $K(.)$ is a kernel function with bandwidth parameter h. Based on $\hat{\sigma}_t^2$, a feasible version of BZ_τ is then given by

$$\mathrm{BZ}_\tau = \frac{\sum_{t=2}^{[\tau T]} \Delta\tilde{y}_t\tilde{y}_{t-1}/\hat{\sigma}_t^2}{\left(\sum_{t=2}^{[\tau T]} \tilde{y}_{t-1}^2/\hat{\sigma}_t^2\right)^{1/2}}$$

where, to economize on notation, we have redefined BZ_τ, and we redefine supBZ analogously. To derive the asymptotic distribution of supBZ, in addition to A1 and A2, we make the following further assumptions:

A3 ε_t follows a symmetric distribution, and $E(\varepsilon_t^8) < \infty$.

A4 $\sigma(.)$ is a Lipschitz continuous function on $[0, 1]$.

A5 $K(.)$ is a bounded nonnegative function defined on the real line and $\int_{-\infty}^{\infty} K(r)dr = 1$.

A6 As $T \to \infty$, $h \to 0$ and $Th^2 \to \infty$.

The assumption $E(\varepsilon_t^8) < \infty$ in A3 is also used in Xu and Phillips (2008). The symmetry assumption for ε_t in A3 is made for technical reasons, and is usually easily satisfied for the kind of equity or equity index returns considered in this paper. The continuity assumption in A4 is used for ease of exposition and could be relaxed to allow for a finite number of discontinuities using the strategy in Xu and Phillips (2008), thereby incorporating examples of volatility specifications involving jumps, such as the shifts in volatility cases of section 2.2. From a modelling perspective, large movements in the volatility can be incorporated in the Lipschitz continuous assumption as well; see Boswijk and Zu (2015) for related discussions. The assumption on the volatility function is nonparametric and can allow for a wide range of volatility dynamics such as trending volatility, (multiple) smooth transition (e.g. logistic) changes in variance, or volatility with Fourier-form periodicity. Our assumption A5 on the kernel function is more general than the Xu and Phillips leave-one-out kernel, and is also more general than the truncated kernel considered in Boswijk and Zu (2015), while the rate condition in A6 coincides with Xu and Phillips (2008).

The following theorem gives the limit distribution of the feasible version of supBZ:

Theorem 2. *Under $\mathcal{H}_1$ and Assumptions A1-A6*

$$\text{supBZ} \Rightarrow \mathbb{M}_c^{\text{BZ}}.$$

Remark 2. The feasible statistic supBZ has the same limiting properties as its infeasible counterpart.

The remaining issue that pertains to a full feasible application of supBZ, supDF and hence $\mathcal{U}$, is that the appropriate asymptotic null critical values q_ξ^{BZ} and q_ξ^{DF} arising from $\mathbb{M}_0^{\text{BZ}}$ and $\mathbb{M}_0^{\text{DF}}$ depend on the volatility function $\sigma(s)$. In the context of supDF, Harvey et al. (2016) employ a wild bootstrap procedure to obtain asymptotically valid null critical values. We now show that this same approach can be employed for supBZ and $\mathcal{U}$.

4 A wild bootstrap procedure

Following Harvey et al. (2016), our wild bootstrap algorithm is defined as follows:

1. Generate a wild bootstrap sample $\{y_t^b\}_{t=1}^T$ by setting

$$y_1^b = 0, \quad y_t^b = y_{t-1}^b + \Delta y_t z_t, \quad t = 2, .., T$$

   where the z_t are *iid* standard normal variates.

2. Use the wild bootstrap sample to compute the pair of statistics supDF and supBZ, but using $\hat{\sigma}_t^2$ from the original sample $\{y_t\}_{t=1}^T$.

3. Repeat step 1 and step 2 M times, denoting the resulting pairs of statistics by $\{\text{supDF}_1^b, \text{supBZ}_1^b\}, ..., \{\text{supDF}_M^b, \text{supBZ}_M^b\}$.

Theorem 3. *Under $\mathcal{H}_1$ and Assumptions A1-A6*

$$\begin{pmatrix} \text{supDF}_m^b \\ \text{supBZ}_m^b \end{pmatrix} \xRightarrow{p} \begin{pmatrix} \mathbb{M}_0^{\text{DF}} \\ \mathbb{M}_0^{\text{BZ}} \end{pmatrix}$$

jointly, for any $1 \leq m \leq M$, where $\xRightarrow{p}$ denotes weak convergence in probability.

Remark 3. The results of Theorem 3 shows that the wild bootstrap procedure is first order valid in approximating the asymptotic null distributions of the supDF and supBZ statistics under $\mathcal{H}_1$ (which includes $\mathcal{H}_0$ as a special case). The asymptotic validity of the marginal bootstrap supDF statistic is shown in Harvey et al. (2016). Theorem 3 strengthens their results with the marginal convergence of the bootstrap supBZ statistic and their joint convergence. The joint convergence occurs because both statistics are calculated from the *same* bootstrap sample; this result is needed for the validity of the union test strategy.

The ξ level bootstrap critical values are obtained from the empirical distribution functions of supDF_m^b and supBZ_m^b calculated from M bootstrap replications. Denoting these critical values as $q_\xi^{b,\text{DF}}$ and $q_\xi^{b,\text{BZ}}$, a rejection of $\mathcal{H}_0$ for supDF is obtained if $\text{supDF} > q_\xi^{b,\text{DF}}$ and a rejection of $\mathcal{H}_0$ for supBZ is obtained if $\text{supBZ} > q_\xi^{b,\text{BZ}}$. As $T, N \to \infty$, it follows that $q_\xi^{b,\text{DF}}$ and $q_\xi^{b,\text{BZ}}$ converge in probability to q_ξ^{DF} and q_ξ^{BZ}, so these bootstrap procedures are correctly sized in the limit under $\mathcal{H}_0$, and inherit exactly the same asymptotic local power functions under $\mathcal{H}_1$ as their infeasibly size-corrected counterparts in section 2.2.

The wild bootstrap counterpart of the union statistic $\mathcal{U}$ is given by

$$\mathcal{U}_m^b = \max\left(\text{supDF}_m^b, \frac{q_\xi^{b,\text{DF}}}{q_\xi^{b,\text{BZ}}}\text{supBZ}_m^b\right)$$

for $m = 1, ..., M$. The results in Theorem 3, and an application of the continuous mapping theorem (CMT), verifies that

$$\mathcal{U}_m^b \xRightarrow{p} \max\left(\mathbb{M}_0^{\text{DF}}, \frac{q_\xi^{\text{DF}}}{q_\xi^{\text{BZ}}}\mathbb{M}_0^{\text{BZ}}\right).$$

The ξ level bootstrap critical value for the union is obtained from the empirical distribution function of $\mathcal{U}_m^b$, and denoting this critical value as $q_\xi^{b,\mathcal{U}}$ we reject $\mathcal{H}_0$ when $\mathcal{U} > q_\xi^{b,\mathcal{U}}$, where

$$\mathcal{U} = \max\left(\text{supDF}, \frac{q_\xi^{b,\text{DF}}}{q_\xi^{b,\text{BZ}}}\text{supBZ}\right).$$

Remark 4. Notice that this is a feasible variant of $\mathcal{U}$ which is based on replacing $q_\xi^{\mathrm{DF}}/q_\xi^{\mathrm{BZ}}$ with $q_\xi^{b,\mathrm{DF}}/q_\xi^{b,\mathrm{BZ}}$.

As $T, N \to \infty$, $\mathcal{U}$ is correctly sized in the limit under $\mathcal{H}_0$, since $q_\xi^{b,\mathcal{U}}$ converges in probability to $q_\xi^{\mathcal{U}}$, and also obtains the same asymptotic local power function under $\mathcal{H}_1$ as the infeasibly size-corrected version in section 2.2.

We have therefore established asymptotic validity of bootstrap variants of supDF, supBZ and $\mathcal{U}$ in terms of size control and local power. We now turn to a comparison of the finite sample properties of these procedures.

5 Finite sample properties

Our finite sample simulations are based on (1)-(3) with $T = 200$. Here we set $\mu = 0$ and $x_1 = 0$, with the ε_t are generated as $NIID(0,1)$ random variates. Figures 3 and 4 show 0.05-level finite sample sizes and powers for the same settings of τ^* as used in the asymptotic simulations of Figures 1 and 2, respectively, with the finite sample volatility functions σ_t being the discrete time analogues of those given in cases (a)-(h) of section 2.2. Here we use 1,000 Monte Carlo replications, together with $M = 499$ bootstrap replications.

For the volatility estimates $\hat{\sigma}_t^2$ we employ the Gaussian kernel

$$K(r) = \frac{1}{\sqrt{2\pi}} \exp(-r^2/2).$$

We determine the bandwidth h using a standard leave-one-out cross-validation bandwidth selection procedure. Specifically, for the cross-validation criteria defined by

$$CV(h) = \sum_{t=2}^{T} ((\Delta y_t)^2 - \hat{\sigma}_{t,-}^2)^2$$

where $\hat{\sigma}_{t,-}^2$ is the Gaussian kernel-based variance estimator $\hat{\sigma}_t^2$ that imposes $K(0) = 0$, the bandwidth is chosen as

$$h_{CV} = \arg \min_{h \in [h_l, h_u]} CV(h).$$

We then construct the $\hat{\sigma}_t^2$ in (6) with h_{CV} in place of h.[2]

The finite sample power curves corresponding to Figure 1 ($\tau^* = 0.6$) and Figure 2 ($\tau^* = 0.8$) are given in Figure 3 (a)-(h) and Figure 4 (a)-(h), respectively. The results

[2]In our implementation, we set $h_l = 1/(2T)$ and $h_u = 1/6$, which ensures that the interval of observations for which the kernel weights are non-negligible ranges from 3 to T.

for $c = 0$ (which are of course the same across Figures 3 and 4) show that the feasible version of supBZ displays excellent size control for $T = 200$. The power curves of the bootstrapped tests in Figure 3 ($\tau^* = 0.6$) generally bear close resemblance to their asymptotic counterparts, with the exception of the late upward volatility shift (c) where the powers of supBZ appear lower than in the limit case. For (c) supBZ is now less powerful than supDF.

In Figure 4 ($\tau^* = 0.8$), we observe that supBZ demonstrates non-monotonicity in its power profiles, with power reversals observed for the larger values of c (something not observed in Figure 3). We conjecture that this behaviour is due to large values of c causing the estimates of $\hat{\sigma}_t^2$ to become inflated via their dependence on $\Delta y_t = \Delta x_t = (c/T)x_{t-1} + u_t$, which in finite samples is not necessarily a good proxy for u_t unless T is large relative to c. Despite this tendency for power reversals with supBZ, it is still the case that the power profile of $\mathcal{U}$ is monotonic and again similar to whichever profile of supDF and supBZ obtains the higher power, though obviously for large c, this profile is now typically that of supDF rather than supBZ, in contrast to what was typically observed in our asymptotic simulations. One noteworthy observation is that for the cases where supBZ displays non-monotonic power, $\mathcal{U}$ can actually have power greater than either supDF or supBZ for certain intermediate c values around the intersection of the supDF and supBZ power profiles. It appears, therefore, that $\mathcal{U}$ offers a robust approach to testing, capturing most of the relatively high power that supBZ can offer over supDF for small to moderate magnitude bubbles, while retaining high power across c in cases where the power of supBZ can drop relative to supDF.

Our model assumes a deterministic volatility function, however it is also of interest to evaluate the performance of the tests under stochastic volatility, which is not covered by assumption A2 but is of empirical relevance. The volatility model we simulate for this exercise is the so-called square root process

$$\mathrm{d}\sigma^2(r) = 0.03(0.25 - \sigma^2(r))\mathrm{d}r + 0.1\sqrt{\sigma^2(r)}\mathrm{d}B(r)$$

where $B(r)$ is a standard Brownian motion process, and the parameter settings used are representative of those considered in Bollerslev and Zhou (2002). The volatility model is re-simulated in each replication of the Monte Carlo experiment, using $NIID(0,1)$ drawings to approximate the Brownian motion increments, with these drawings being independent of those for ε_t. Figures 3 (i) and 4 (i) report the empirical rejection frequencies of the tests for this model across c. When $c = 0$, it can be seen that the empirical rejection frequencies of our bootstrap tests are very close to the nominal significance level. However, in unreported results, we found that this does not appear to hold if the Brownian motion increments are correlated with the ε_t, e.g. when leverage effects are

present.[3] As regards the power of the tests, supBZ and supDF tend to perform quite similarly under this stochastic volatility model, with supBZ displaying small power gains over supDF for small c, and a reverse pattern for larger c.

6 Empirical illustrations

In this section, we apply the bootstrap supDF, supBZ and $\mathcal{U}$ procedures to two data sets, with supBZ using the same Gaussian kernel and cross-validation bandwidth selection method discussed in section 5. The data are logarithms of the inflation-adjusted FTSE index from December 1985 to December 1999 and the S&P 500 index from January 1980 to March 2000. For each dataset we consider monthly, weekly and daily frequencies. Figure 5 and Figure 6 show the time series plots of log prices and the first differences (log returns) for the FTSE index and the S&P 500 index, respectively, at the three sampling frequencies. We see that the levels of both series are generally increasing during the period considered (Homm and Breitung, 2012, consider similar sample periods so that the sample endpoints correspond to periods where the prices reach their peaks). From the plots of the first differences (i.e. log returns), the presence of time varying volatility is clearly a plausible phenomenon.

The three procedures supDF, supBZ and $\mathcal{U}$ are applied to each of the six series, with the results given in Table 1, the entries being bootstrap p-values associated with the different procedures. First we observe that the supDF test does not reject the null in favour of explosive behaviour for any of the series considered at conventional significance levels.[4] Turning to supBZ, we find evidence of explosive behaviour, at least at the 0.05-level, for the daily FTSE series and all frequencies of the S&P 500 index. These rejections are preserved when considering the $\mathcal{U}$ procedure, albeit at a slightly weaker significance level for monthly S&P 500. This pattern of results, where rejections are obtained by supBZ and $\mathcal{U}$ but not supDF, fits well with our asymptotic and finite sample simulation findings when an explosive period of small to modest magnitude is present in the data, along with time-varying volatility. Given that the supDF test alone fails to detect these explosive episodes, our application reinforces our earlier findings that the WLS-based supBZ procedure can offer enhanced levels of detectability of explosive behaviour. Moreover, given that the union of rejections procedure does not reduce the number of series for which rejections of the unit root null are found, we again find that there is little cost to adopting this joint test approach, which also provides a degree of insurance for other potential cases where supDF might reject but supBZ not.

[3] This is not a surprising result, as the wild bootstrap method only mimics the heteroskedastic pattern in the data, but is unable to reproduce the dependence between the volatility process increments and the model errors.

[4] We also compared the supDF test statistic with standard critical values obtained under an assumption of constant volatility (as in PWY), and again failed to reject the null in all cases.

7 Discussion and conclusion

In this paper we have proposed a WLS-based variant of the PWY test for explosive autoregressive behaviour in a financial time series. We find that such an approach has appealing asymptotic power properties, with the potential to deliver substantially greater power than the established OLS-based approach for many volatility and explosive settings. Given that the OLS-based test can outperform the WLS-based test for other volatility and explosive specifications, we also suggested a union of rejections procedure that succeeds in capturing the better power available from the two constituent tests for a given alternative. Our approach involves a nonparametric kernel-based volatility function estimator for computation of the WLS-based statistic, together with the use of a wild bootstrap procedure applied jointly to both individual tests, delivering a powerful testing procedure that is asymptotically size-robust to a wide range of time-varying volatility specifications. Finite sample simulations indicate that the procedures should work well in practice, and application of the tests to FTSE and S&P 500 price data supports our premise that the WLS-based test can provide improved ability to detect explosive behaviour compared to extant procedures.

While we have focused our attention on a PWY-type framework, using a single supremum of forward recursively calculated statistics, our WLS approach could of course be used in the context of the double supremum testing approach of Phillips et al. (2015). The WLS variant of their double supremum test statistic is given by

$$\text{supBZ}^* = \sup_{\tau_1 \in [0, 1-\pi], \tau_2 \in [\tau_1+\pi, 1]} \text{BZ}_{\tau_1,\tau_2}$$

where

$$\text{BZ}_{\tau_1,\tau_2} = \frac{\sum_{t=[\tau_1 T]+1}^{[\tau_2 T]} \Delta\tilde{y}_t \tilde{y}_{t-1}/\hat{\sigma}_t^2}{\left(\sum_{t=[\tau_1 T]+1}^{[\tau_2 T]} \tilde{y}_{t-1}^2/\hat{\sigma}_t^2\right)^{1/2}}$$

with $\tilde{y}_t$ and $\hat{\sigma}_t^2$ as defined above. The wild bootstrap and union of rejections methodologies can be applied to supBZ* in an entirely similar way as to supBZ, and we would anticipate power gains to be available for these more general tests also.

Finally, the autoregressive specification that we have adopted in this paper involves a one-time change from unit root to explosive dynamics. One could also consider a more general nonparametric specification for ρ_t, along with an appropriate modification to the testing strategy, e.g. a CUSUM-based test along the lines of Homm and Breitung (2012). As Xu (2015) demonstrates, CUSUM tests in general are not robust to nonstationary volatility, hence such an approach would need to be robustified to the potential heteroskedasticity that we consider in this paper. One approach that could be considered would be to modify the Homm and Breitung (2012) CUSUM test using our volatility estimator $\hat{\sigma}_t^2$. This would be an interesting avenue for future research.

References

Bollerslev, T. and Zhou, H. (2002). Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *Journal of Econometrics* 109, 33–65.

Boswijk, H. P. and Zu, Y. (2018). Adaptive wild bootstrap testing for a unit root with nonstationary volatility. *The Econometrics Journal*, forthcoming, doi:10.1111/ectj.12100.

Cavaliere, G. and Taylor, A.M.R. (2007). Testing for unit roots in time series models with non-stationary volatility. *Journal of Econometrics* 140, 919–947.

Elliott, G., Rothenberg, T.J. and Stock, J.H. (1996). Efficient tests for an autoregressive unit root. *Econometrica* 64, 813–836.

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.M.R. (2016). Tests for explosive financial bubbles in the presence of non-stationary volatility. *Journal of Empirical Finance* 38, 548–574.

Harvey, D.I., Leybourne, S.J. and Taylor, A.M.R. (2009). Unit root testing in practice: dealing with uncertainty over the trend and initial condition (with commentaries and rejoinder). *Econometric Theory* 25, 587-667.

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.

Phillips, P.C.B., Shi, S.-P. and Yu, J. (2015). Testing for multiple bubbles: historical episodes of exuberance and collapse in the S&P 500. *International Economic Review* 56, 1043–1077.

Phillips, P.C.B., Wu, Y. and Yu, J. (2011). Explosive behavior in the 1990s Nasdaq: when did exuberance escalate stock values? *International Economic Review* 52, 201–226.

Robinson, P. M. (1987). Asymptotically efficient estimation in the presence of heteroskedasticity of unknown form. *Econometrica* 55, 875–891.

Xu, K.-L. (2015). Testing for structural change under non-stationary variances. *The Econometrics Journal* 18, 274-305.

Xu, K.-L. and Phillips, P.C.B. (2008). Adaptive estimation of autoregressive models with time-varying variances. *Journal of Econometrics* 142, 265–280.

Xu, K.-L. and Yang, J.C. (2015). Towards uniformly efficient trend estimation under weak/strong correlation and non-stationary volatility. *Scandinavian Journal of Statistics*, 42, 63-86.

Appendix: Proofs of theorems

In what follows, we set $\mu = 0$ and $y_1 = 0$, without loss of generality, so that $y_t = \tilde{y}_t = x_t$.

Proof of Theorem 1

First notice that

$$\mathrm{BZ}_\tau = \frac{\sum_{t=2}^{[\tau T]} \frac{\Delta y_t y_{t-1}}{\sigma_t^2}}{\left(\sum_{t=2}^{[\tau T]} \frac{y_{t-1}^2}{\sigma_t^2}\right)^{1/2}}.$$

Using the result

$$T^{-1/2} y_{[rT]} \Rightarrow U_c(r)$$

which follows straightforwardly from Theorem 1 of Harvey et al. (2016), it follows that

$$T^{-1/2} \frac{y_{[rT]-1}}{\sigma_t} \Rightarrow U_c(r)/\sigma(r) = V_c(r).$$

Notice that this holds for all $r \in (0,1)$ and the U_c (thus also V_c) process is defined to have two regimes. One immediate consequence of this weak convergence result is that

$$T^{-2} \sum_{t=2}^{[\tau T]} \frac{y_{t-1}^2}{\sigma_t^2} \Rightarrow \int_0^\tau V_c(r)^2 \mathrm{d}r.$$

For $\Delta y_t/\sigma_t$, notice that

$$\frac{\Delta y_t}{\sigma_t} = \begin{cases} \varepsilon_t & \tau \leqslant \tau^* \\ (c/T)\frac{y_{t-1}}{\sigma_t} + \varepsilon_t & \tau > \tau^* \end{cases}.$$

Then it follows easily that when $\tau \leqslant \tau^*$,

$$T^{-1} \sum_{t=2}^{[\tau T]} \frac{\Delta y_t y_{t-1}}{\sigma_t^2} \Rightarrow \int_0^\tau V_c(r) \mathrm{d}W(r)$$

and so

$$\mathrm{BZ}_\tau \Rightarrow \frac{\int_0^\tau V_c(r) \mathrm{d}W(r)}{\left(\int_0^\tau V_c(r)^2 \mathrm{d}r\right)^{1/2}}.$$

When $\tau > \tau^*$,

$$\begin{aligned} T^{-1}\sum_{t=2}^{[\tau T]}\frac{\Delta y_t y_{t-1}}{\sigma_t^2} &= T^{-1}\sum_{t=2}^{[\tau T]}\frac{\varepsilon_t y_{t-1}}{\sigma_t} + T^{-1}\sum_{t=[\tau^* T]+1}^{[\tau T]}\frac{(c/T)y_{t-1}^2}{\sigma_t^2} \\ &\Rightarrow \int_0^\tau V_c(r)\mathrm{d}W(r) + c\int_{\tau^*}^\tau V_c(r)^2\mathrm{d}r \end{aligned}$$

and so

$$\mathrm{BZ}_\tau \Rightarrow \frac{\int_0^\tau V_c(r)\mathrm{d}W(r) + c\int_{\tau^*}^\tau V_c(r)^2\mathrm{d}r}{\left(\int_0^\tau V_c(r)^2\mathrm{d}r\right)^{1/2}}.$$

It then follows that, via the continuous mapping theorem,

$$\mathrm{supBZ} \Rightarrow \sup_{\tau\in[\tau_0,1]} L_c(\tau) = \mathbb{M}_c^{\mathrm{BZ}}.$$

Proof of Theorem 2

In order for the *feasible* statistic BZ_τ, in which σ_t^2 is replaced by $\hat{\sigma}_t^2$, to converge to $L_c(\tau)$ (the limit of its infeasible counterpart), we require the following two conditions to hold:

$$T^{-1}\left(\sum_{t=2}^T(\Delta y_t y_{t-1}/\hat{\sigma}_t^2) - \sum_{t=2}^T(\Delta y_t y_{t-1}/\sigma_t^2)\right) = o_p(1) \quad (7)$$

$$T^{-2}\left(\sum_{t=2}^T(y_{t-1}/\hat{\sigma}_t)^2 - \sum_{t=2}^T(y_{t-1}/\sigma_t)^2\right) = o_p(1). \quad (8)$$

To show (7) and (8), we largely adapt the strategy used in Robinson (1987) and Xu and Phillips (2008).

For (7), first notice that the spot variance estimator can be written as

$$\hat{\sigma}_t^2 = \sum_{i=2}^T w_{t,i}\hat{u}_i^2$$

where the weights are defined as $w_{t,i} = K_h\left(\frac{i-t}{T}\right) / \sum_{i=2}^T K_h\left(\frac{i-t}{T}\right)$ and satisfy $\sum_{i=2}^T w_{t,i} = 1$. Defining $\tilde{\sigma}_t^2 = \sum_{i=2}^T w_{t,i}u_i^2$ and $\bar{\sigma}_t^2 = \sum_{i=2}^T w_{t,i}\sigma_i^2$, we can make the following decomposition:

$$\begin{aligned} T^{-1}\sum_{t=2}^T \Delta y_t y_{t-1}(1/\hat{\sigma}_t^2 - 1/\sigma_t^2) &= T^{-1}\sum_{t=2}^T \Delta y_t y_{t-1}(1/\hat{\sigma}_t^2 - 1/\tilde{\sigma}_t^2) + T^{-1}\sum_{t=2}^T \Delta y_t y_{t-1}(1/\tilde{\sigma}_t^2 - 1/\bar{\sigma}_t^2) \\ &\quad + T^{-1}\sum_{t=2}^T \Delta y_t y_{t-1}(1/\bar{\sigma}_t^2 - 1/\sigma_t^2) \end{aligned}$$

where the three right-hand-side terms are denoted as A, B and C respectively. We next

show that all the three terms are $o_p(1)$.

For A,

$$\begin{aligned} |A| &\leqslant \left(\max_t \left|1/(\hat{\sigma}_t^2\tilde{\sigma}_t^2)\right|\right) \sum_{t=2}^{T} \left|(\Delta y_t y_{t-1})(\hat{\sigma}_t^2 - \tilde{\sigma}_t^2)/T\right| \\ &\leqslant \left(\max_t \left|1/(\hat{\sigma}_t^2\tilde{\sigma}_t^2)\right|\right) \left(T^{-2}\sum_{t=2}^{T}(\Delta y_t y_{t-1})^2\right)^{1/2} \left(\sum_{t=2}^{T}(\hat{\sigma}_t^2 - \tilde{\sigma}_t^2)^2\right)^{1/2} \end{aligned}$$

using Cauchy-Schwartz inequality. Notice $\hat{\sigma}_t^2$ and $\tilde{\sigma}_t^2$ are bounded away from 0, $\max_t \left|1/\hat{\sigma}_t^2\tilde{\sigma}_t^2\right|$ will at most be $O_p(1)$; also it is easy to obtain that $T^{-2}\sum_{t=2}^{T}(\Delta y_t y_{t-1})^2 = O_p(1)$. Then to show $A = o_p(1)$, we are left to show $\sum_{t=2}^{T}(\hat{\sigma}_t^2 - \tilde{\sigma}_t^2)^2 = o_p(1)$. Now

$$\begin{aligned} \hat{\sigma}_t^2 - \tilde{\sigma}_t^2 &= \sum_{i=2}^{T} w_{t,i}(\hat{u}_i^2 - u_i^2) \\ &= \sum_{i=2}^{T} w_{t,i}((\Delta y_i)^2 - (\sigma_i\varepsilon_i)^2) \\ &= \sum_{i=2}^{T} w_{t,i}((\rho_i y_{i-1} + \sigma_i\varepsilon_i)^2 - (\sigma_i\varepsilon_i)^2) \\ &= O_p(T^{-1}) \end{aligned} \tag{9}$$

where in the last step we use the definition that $\rho_i = 0$ before τ^* and $\rho_i = c/T$ after τ^*, so $\hat{\sigma}_t^2 - \tilde{\sigma}_t^2$ will be at most $O_p(T^{-1})$. This further implies that $\sum_{t=2}^{T}(\hat{\sigma}_t^2 - \tilde{\sigma}_t^2)^2 = O_p(T^{-1})$, which completes the proof for $A = o_p(1)$.

For B, notice that

$$\begin{aligned} B &= T^{-1}\sum_{t=2}^{T} \Delta y_t y_{t-1}(1/\tilde{\sigma}_t^2 - 1/\bar{\sigma}_t^2) \\ &= T^{-1}\sum_{t=2}^{T} \Delta y_t y_{t-1}(\bar{\sigma}_t^2 - \tilde{\sigma}_t^2)\bar{\sigma}_t^{-4} + T^{-1}\sum_{t=2}^{T} \Delta y_t y_{t-1}(\bar{\sigma}_t^2 - \tilde{\sigma}_t^2)^2\bar{\sigma}_t^{-4}\tilde{\sigma}_t^{-2} \end{aligned} \tag{10}$$

where the equality $p^{-1} - q^{-1} = (q-p)q^2 + (q-p)^2p^{-1}q^{-2}$ is used. We denote the two

terms as $B1$ and $B2$, and we look at them separately. For $B1$,

$$\begin{aligned} B1 &= T^{-1}\sum_{t=2}^{T}\Delta y_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4} \\ &= T^{-1}\sum_{t=2}^{T}\left(\rho_t y_{t-1}+u_t\right)y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4} \\ &= cT^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4}+T^{-1}\sum_{t=2}^{T}u_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4} \\ &= B11+B12 \end{aligned}$$

where $B11$ and $B12$ are defined implicitly. Under the null $B11=0$, while under the alternative we have

$$\begin{aligned} |B11| &= \left|cT^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4}\right| \\ &\leqslant \mathbb{C}\max_t\left|(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\right|T^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2 \end{aligned}$$

where $\mathbb{C}$ generically denotes a positive constant. Notice that

$$\begin{aligned} P(\max_t|\tilde{\sigma}_t^2-\bar{\sigma}_t^2|>\varepsilon) &\leqslant \sum_{t=2}^{T}P(|\tilde{\sigma}_t^2-\bar{\sigma}_t^2|>\varepsilon) \\ &\leqslant \sum_{t=2}^{T}E|\tilde{\sigma}_t^2-\bar{\sigma}_t^2|^4/\varepsilon^4 = O_p\left(1/(Th^2)\right) \end{aligned} \tag{11}$$

and it is straightforward to show that $T^{-2}\sum_{i=2}^{T}y_{t_{i-1}}^2=O_p(1)$, so we have $B11=o_p(1)$ because $Th^2\to\infty$. For the term $B12$, first notice that $\left\{u_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t\right\}$ is a martingale difference sequence with respect to the filtration $\mathcal{F}_t$. This is because

$$\begin{aligned} E\left(u_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4}|\mathcal{F}_{t-1}\right) &= y_{t-1}\bar{\sigma}_t^{-2}E\left(u_t|\mathcal{F}_{t-1}\right)+\bar{\sigma}_t^{-4}E\left(u_t y_{t-1}\sum_{i=2}^{T}w_i u_i^2|\mathcal{F}_{t-1}\right) \\ &= \bar{\sigma}_t^{-4}E\left(u_t y_{t-1}\sum_{i=2}^{t-1}w_i u_i^2|\mathcal{F}_{t-1}\right)+\bar{\sigma}_t^{-4}E\left(u_t y_{t-1}\sum_{i=t+1}^{T}w_i u_i^2|\mathcal{F}_{t-1}\right) \\ &\quad +\bar{\sigma}_t^{-4}E\left(u_t^3 y_{t-1}w_t|\mathcal{F}_{t-1}\right) \end{aligned}$$

where the first term is clearly 0, the second term is also 0 by noticing that, for $i>t$,

$$E\left(u_t u_i^2|\mathcal{F}_{t-1}\right)=E\left(u_t E\left(u_i^2|\mathcal{F}_{i-1}\right)|\mathcal{F}_{t-1}\right)=0$$

while the third term is also 0 by the assumption that $E\left(u_t^3|\mathcal{F}_{t-1}\right)=0$.[5] Using Markov's inequality and the fact that $B12$ is a sum of a martingale difference sequence we have

$$\begin{aligned}
P\left(|B12|>\varepsilon\right) &= P\left(\left|T^{-1}\sum_{t=2}^{T}\Delta y_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4}\right|>\varepsilon\right)\\
&\leqslant E\left|T^{-1}\sum_{t=2}^{T}\Delta y_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)\bar{\sigma}_t^{-4}\right|^2/\varepsilon^2\\
&\leqslant \mathbb{C}T^{-2}\sum_{t=2}^{T}E|\Delta y_t y_{t-1}|^2(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^2/\varepsilon^2\\
&\leqslant \mathbb{C}T^{-2}\varepsilon^{-2}\sum_{t=2}^{T}(E|\Delta y_t y_{t-1}|^4)^{1/2}(E(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^4)^{1/2}\\
&\leqslant \left(\max_t E(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^4\right)^{1/2}\times\mathbb{C}T^{-2}\varepsilon^{-2}\sum_{t=2}^{T}(E|\Delta y_t y_{t-1}|^4)^{1/2}.
\end{aligned}$$

It is easy to see that $T^{-2}\sum_{t=2}^{T}(E|\Delta y_t y_{t-1}|^4)^{1/2}=O_p(1)$, and we next show that $\max_t E(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^4=o_p(1)$:

$$\begin{aligned}
E(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^4 &= E\left(\sum_{i=2}^{T}w_{t,i}(u_i^2-\sigma_i^2)\right)^4\\
&\leqslant E\left(\sum_{i=2}^{T}w_{t,i}^2(u_i^2-\sigma_i^2)^2\right)^2\\
&\leqslant (1/Th)^2\,E\left(\sum_{i=2}^{T}w_{t,i}(u_i^2-\sigma_i^2)^2\right)^2\\
&\leqslant (1/Th)^2\sum_{i=2}^{T}w_{t,i}E(u_i^2-\sigma_i^2)^4\\
&= O_p\left((1/Th)^2\right) \qquad (12)
\end{aligned}$$

where in the second step Burkholder's inequality for a martingale difference sequence is used, in the third step the fact that $\max_i w_{t,i}=O(1/(Th))$ is used, in the fourth step we use Jensen's inequality, and in the last step we use the boundedness of the 8th moment of ε_t from Assumption A3. So, we have shown $B12=o_p\left(1\right)$, which implies $B1=o_p(1)$. For the term $B2$ in (10),

$$\left|T^{-1}\sum_{t=2}^{T}\Delta y_t y_{t-1}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^2\bar{\sigma}_t^{-4}\tilde{\sigma}_t^{-2}\right|\leqslant\mathbb{C}\left(T^{-2}\sum_{t=2}^{T}(\Delta y_t y_{t-1})^2\right)^{1/2}\left(\sum_{t=2}^{T}(\bar{\sigma}_t^2-\tilde{\sigma}_t^2)^4\right)^{1/2}.$$

[5] In Xu and Philips (2008), a leave-one-out estimator for spot volatility is used to make the third term 0. Here our assumption of a symmetric ε_t achieves the same result without the need to use the leave-one-out estimator.

As in the derivation of $A = o_p(1)$, it is easy to see that $T^{-2}\sum_{t=2}^{T}(\Delta y_t y_{t-1})^2 = O_p(1)$. Using the Markov inequality, it follows that

$$\begin{aligned} P\left(\left|\sum_{t=2}^{T}(\bar{\sigma}_t^2 - \tilde{\sigma}_t^2)^4\right| > \varepsilon\right) &\leqslant E\sum_{t=2}^{T}(\bar{\sigma}_t^2 - \tilde{\sigma}_t^2)^4/\varepsilon \\ &= O_p(1/(Th^2)) = o_p(1) \end{aligned}$$

by applying the result in (12), and recalling the assumption that $Th^2 \to \infty$, so we have $B2 = o_p(1)$. This completes the proof for $B = o_p(1)$.

For C, notice that

$$\begin{aligned} C &= T^{-1}\sum_{t=2}^{T}\Delta y_t y_{t-1}(1/\bar{\sigma}_t^2 - 1/\sigma_t^2) \\ &= cT^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2(1/\bar{\sigma}_t^2 - 1/\sigma_t^2) + T^{-1}\sum_{t=2}^{T} u_t y_{t-1}(1/\bar{\sigma}_t^2 - 1/\sigma_t^2) \\ &= C1 + C2 \end{aligned}$$

where $C1$ and $C2$ are defined implicitly. For $C1$,

$$\begin{aligned} C1 &= \left|cT^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2(1/\bar{\sigma}_t^2 - 1/\sigma_t^2)\right| \\ &\leqslant c\max_t\left|\bar{\sigma}_t^2\sigma_t^2\right|\max_t\left|\sigma_t^2 - \bar{\sigma}_t^2\right|T^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2. \end{aligned}$$

Here $T^{-2}\sum_{t=[\tau^* T]+1}^{T} y_{t-1}^2 = O_p(1)$, $\max_t|\bar{\sigma}_t^2\sigma_t^2| = O_p(1)$, and we next show that $\max_t|\sigma_t^2 - \bar{\sigma}_t^2| = o(1)$, as follows

$$\begin{aligned} \bar{\sigma}_t^2 - \sigma_t^2 &= \frac{\sum_{i=2}^{T} K_h\left(\frac{i-t}{T}\right)\sigma_i^2}{\sum_{i=2}^{T} K_h\left(\frac{i-t}{T}\right)} - \sigma_t^2 \\ &= \int_0^1 \frac{1}{h}K\left(\frac{s - t/T}{h}\right)\sigma(s)\mathrm{d}s(1 + o(T^{-1})) - \sigma_t^2 \\ &= \int_{-\infty}^{\infty} K(s)\mathrm{d}s\sigma_t(1 + o(T^{-1})) - \sigma_t^2 = o(1) \end{aligned} \tag{13}$$

uniformly for all t. Then we have $C1 = o_p(1)$. For $C2$, notice that $u_t y_{t-1}$ is a martingale

difference sequence, so

$$\begin{aligned} E\left(C2\right)^2 &= E\left(T^{-1}\sum_{t=2}^{T} u_t y_{t-1}(1/\bar{\sigma}_t^2 - 1/\sigma_t^2)\right)^2 \\ &= T^{-2}\sum_{t=2}^{T} E[E(u_t y_{t-1})^2|\mathcal{F}_{t-1}](1/\bar{\sigma}_t^2 - 1/\sigma_t^2)^2 \\ &= T^{-2}\sum_{t=2}^{T} E(y_{t-1}^2)(\bar{\sigma}_t^2 - \sigma_t^2)^2\bar{\sigma}_t^{-4}\sigma_t^{-2} \\ &\leqslant (\max_t |\bar{\sigma}_t^2 - \sigma_t^2|)^2 \max_t \left|\bar{\sigma}_t^{-4}\sigma_t^{-2}\right| T^{-2}\sum_{t=2}^{T} E(y_{t-1}^2). \end{aligned}$$

Clearly $T^{-2}\sum_{t=2}^{T} E(y_{t-1}^2) = O_p(1)$, and it suffices to show $\max_t |\bar{\sigma}_t^2 - \sigma_t^2| = o_p(1)$, which is already shown in (13). We have thus shown that $C2 = o_p\left(1\right)$ and also $C = o_p\left(1\right)$.

For (8), notice that

$$T^{-2}\left|\sum_{t=2}^{T}(y_{t-1}/\hat{\sigma}_t)^2 - \sum_{i=2}^{T}(y_{t-1}/\sigma_t)^2\right| \leqslant \mathbb{C}T^{-2}\sum_{t=2}^{T} y_{t-1}^2|\hat{\sigma}_t^2 - \sigma_t^2|.$$

Since $T^{-2}\sum_{t=2}^{T} y_{t-1}^2 = O_p(1)$, for (8) to hold it suffices to show

$$\max_t |\hat{\sigma}_t^2 - \sigma_t^2| = o_p(1). \tag{14}$$

First make the decomposition

$$\max_t |\hat{\sigma}_t^2 - \sigma_t^2| \leqslant \max_t |\hat{\sigma}_t^2 - \tilde{\sigma}_t^2| + \max_t |\tilde{\sigma}_t^2 - \bar{\sigma}_t^2| + \max_t |\bar{\sigma}_t^2 - \sigma_t^2| = D + E + F$$

where the three terms are defined implicitly. For the term D, using the result in (9), it easily follows that $D = O_p(T^{-1})$. $E = o_p\left(1\right)$ has already been shown in (11). $F = o(1)$ has already been shown in (13). We have thus shown (14). This completes the whole proof.

Proof of Theorem 3

The result for supDF_m^b follows straightforwardly from Harvey et al. (2016). Also from Harvey et al. (2016), it follows that

$$T^{-1/2}y_{[rT]}^b \overset{p}{\Rightarrow} \int_0^r \sigma(s)\mathrm{d}W(s). \tag{15}$$

It then follows easily by the uniform consistency of the spot volatility estimator in (14), and since $\sigma_{[rT]+1} \to \sigma(r)$, that

$$T^{-1/2} y^b_{[rT]} / \hat{\sigma}_{[rT]+1} \overset{p}{\Rightarrow} V_0(r).$$

Application of the continuous mapping theorem then gives

$$\begin{aligned} \sum_{t=2}^{[\tau T]} \Delta y^b_t y^b_{t-1} / \hat{\sigma}^2_t &\overset{p}{\Rightarrow} \int_0^\tau V_0(r) \mathrm{d}W(r) \\ \sum_{t=2}^{[\tau T]} y^2_{t-1} / \hat{\sigma}^2_t &\overset{p}{\Rightarrow} \int_0^\tau V_0(r)^2 \mathrm{d}r. \end{aligned}$$

Denoting by BZ^b_τ the BZ_τ statistic based on a bootstrap sample, we then obtain

$$\mathrm{BZ}^b_\tau \overset{p}{\Rightarrow} \frac{\int_0^\tau V_0(r) \mathrm{d}W(r)}{\left(\int_0^\tau V_0(r)^2 \mathrm{d}r \right)^{1/2}}$$

and thus

$$\mathrm{sup}\mathrm{BZ}^b_m \overset{p}{\Rightarrow} \mathbb{M}^{\mathrm{BZ}}_0.$$

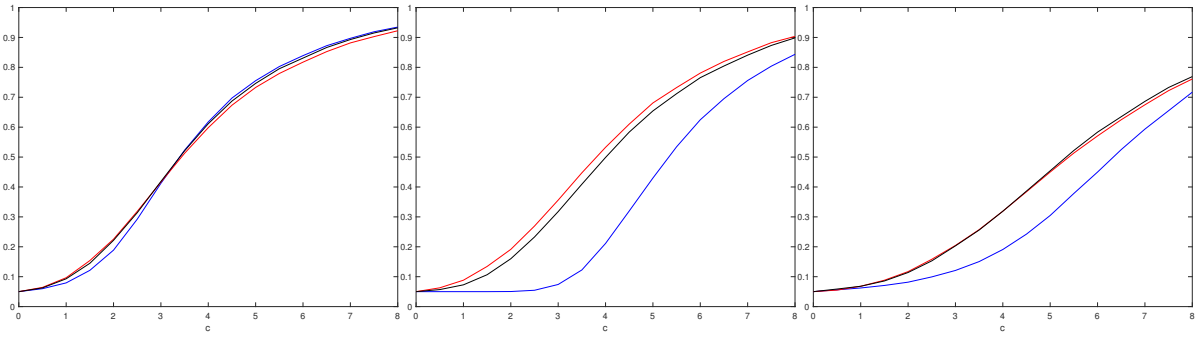

(a) Constant volatility (b) Early upward shift (c) Late upward shift

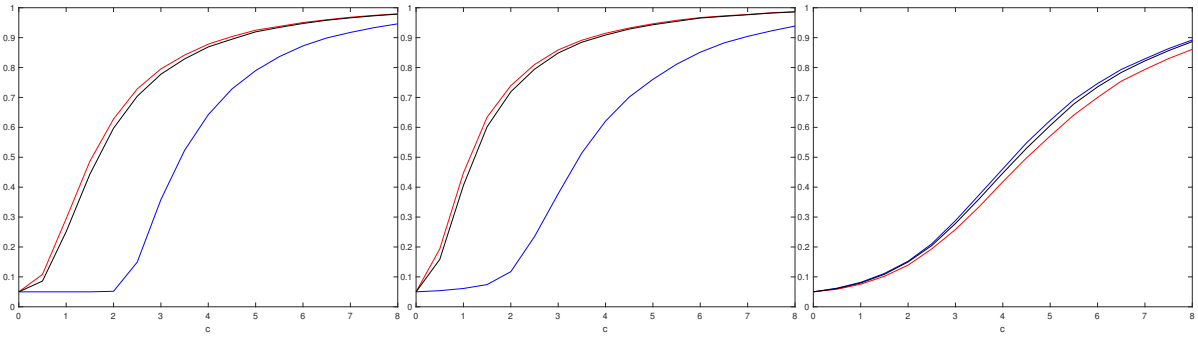

(d) Early downward shift (e) Late downward shift (f) Upward trend

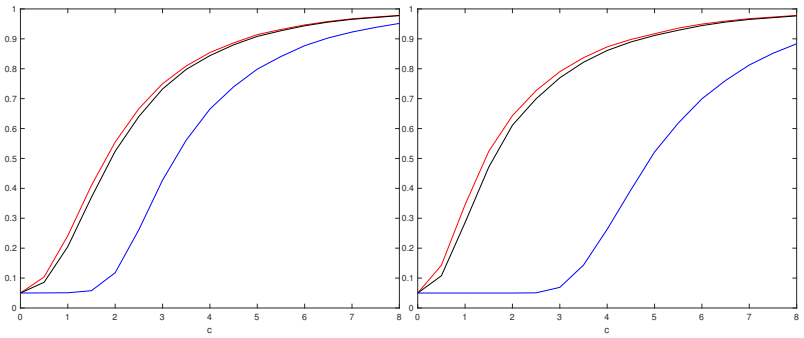

(g) Downward trend (h) Double shift

Figure 1: Limiting local power curves, $\tau^* = 0.6$, supDF: —, supBZ: —, $\mathcal{U}$: —.

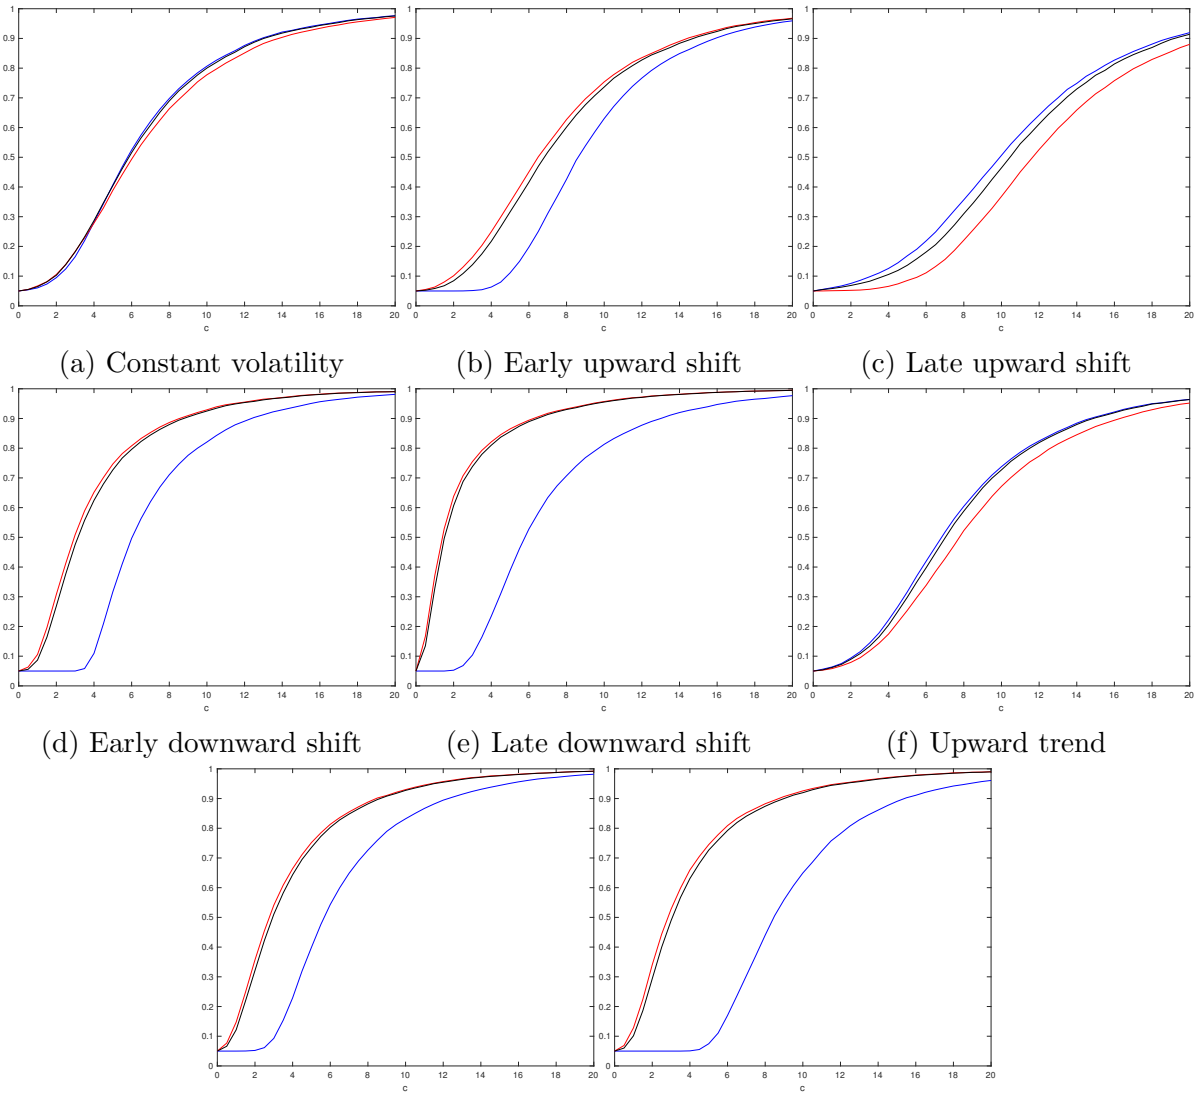

(a) Constant volatility (b) Early upward shift (c) Late upward shift

(d) Early downward shift (e) Late downward shift (f) Upward trend

(g) Downward trend (h) Double shift

Figure 2: Limiting local power curves, $\tau^* = 0.8$, supDF: —, supBZ: —, $\mathcal{U}$: —.

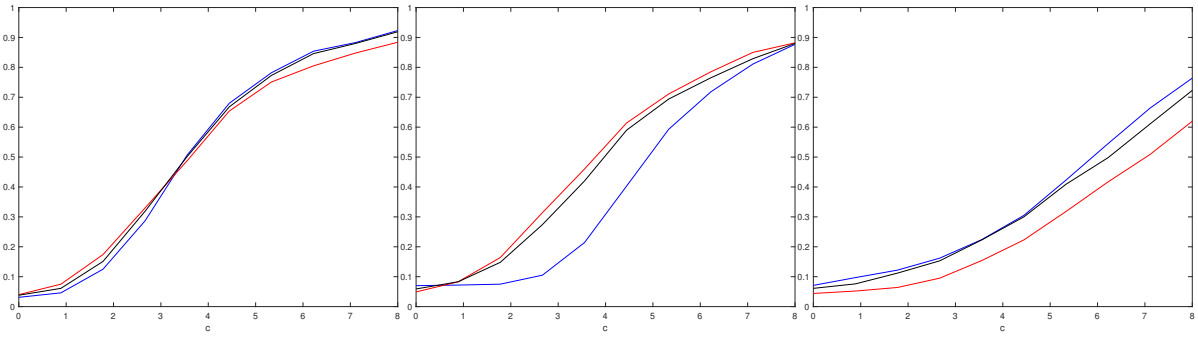

(a) Constant volatility (b) Early upward shift (c) Late upward shift

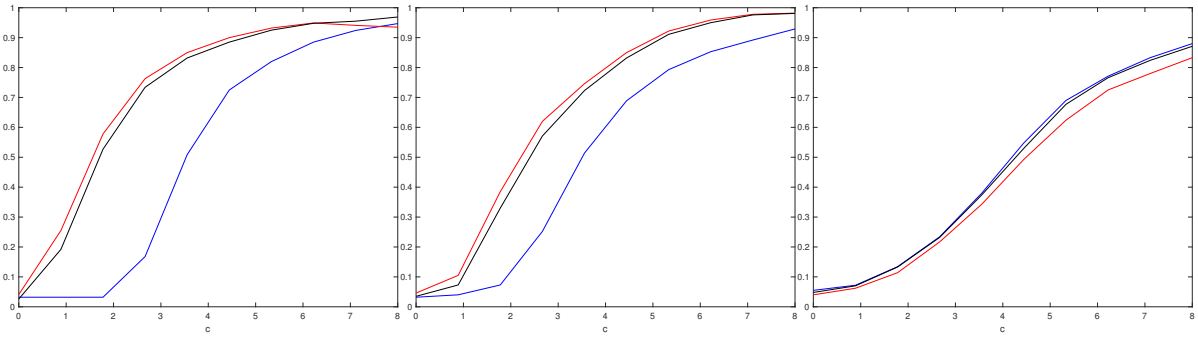

(d) Early downward shift (e) Late downward shift (f) Upward trend

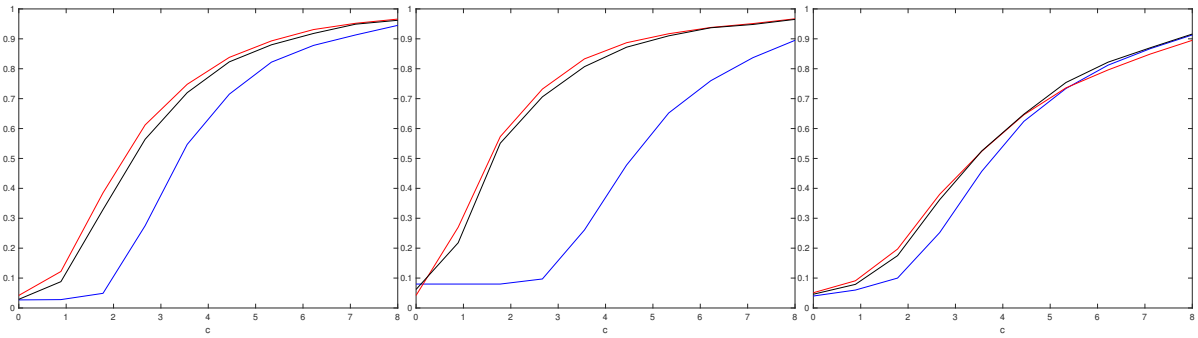

(g) Downward trend (h) Double shift (i) Stochastic volatility

Figure 3: Finite sample local power curves, $\tau^* = 0.6$, supDF: —, supBZ: —, $\mathcal{U}$: —.

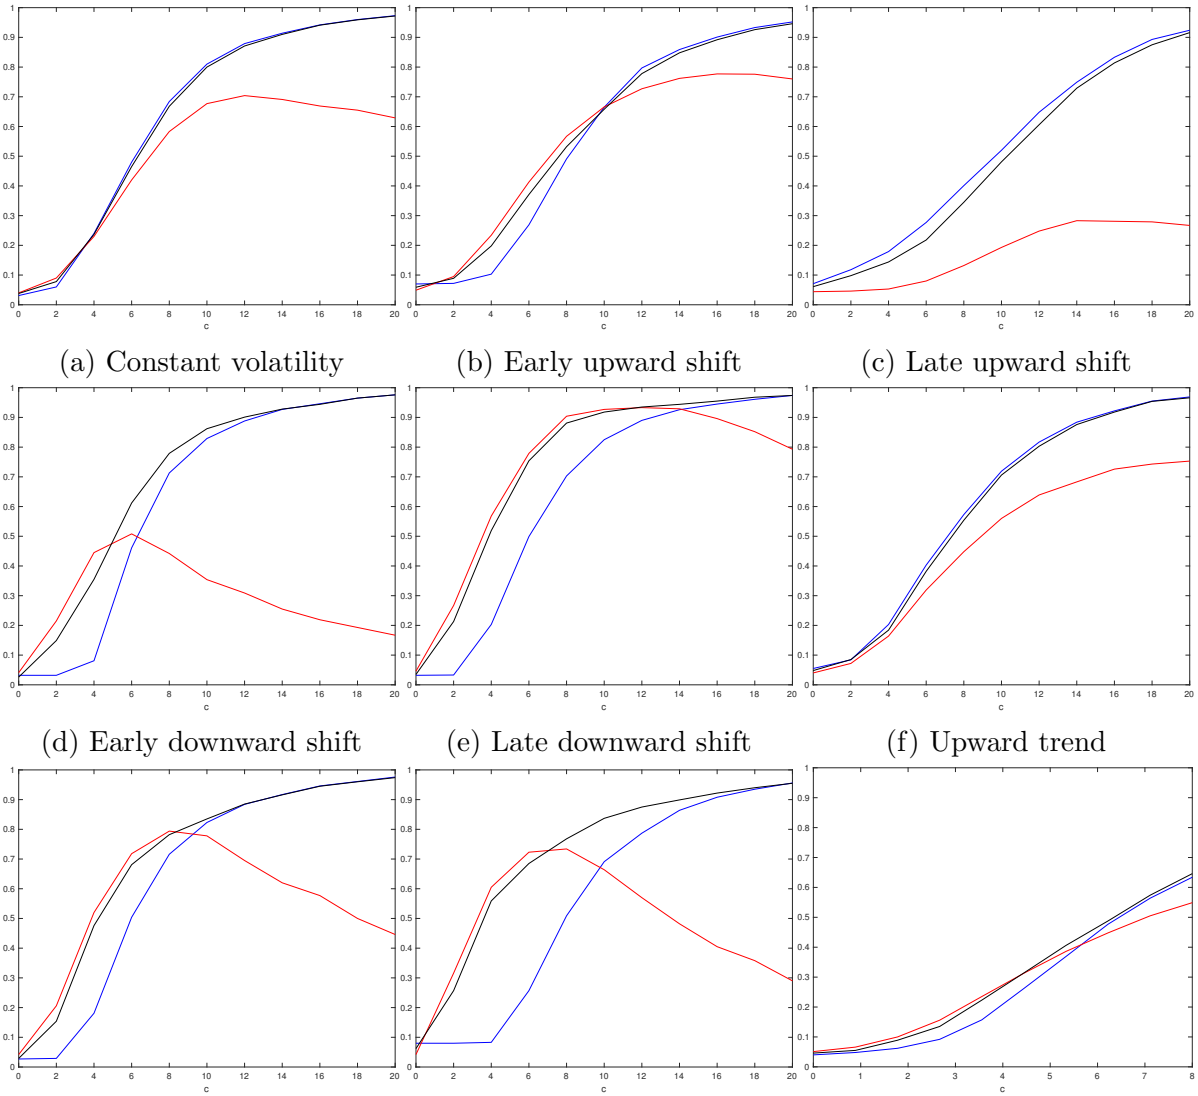

(a) Constant volatility (b) Early upward shift (c) Late upward shift

(d) Early downward shift (e) Late downward shift (f) Upward trend

(g) Downward trend (h) Double shift (i) Stochastic volatility

Figure 4: Finite sample local power curves, $\tau^* = 0.8$, supDF: —, supBZ: —, $\mathcal{U}$: —.

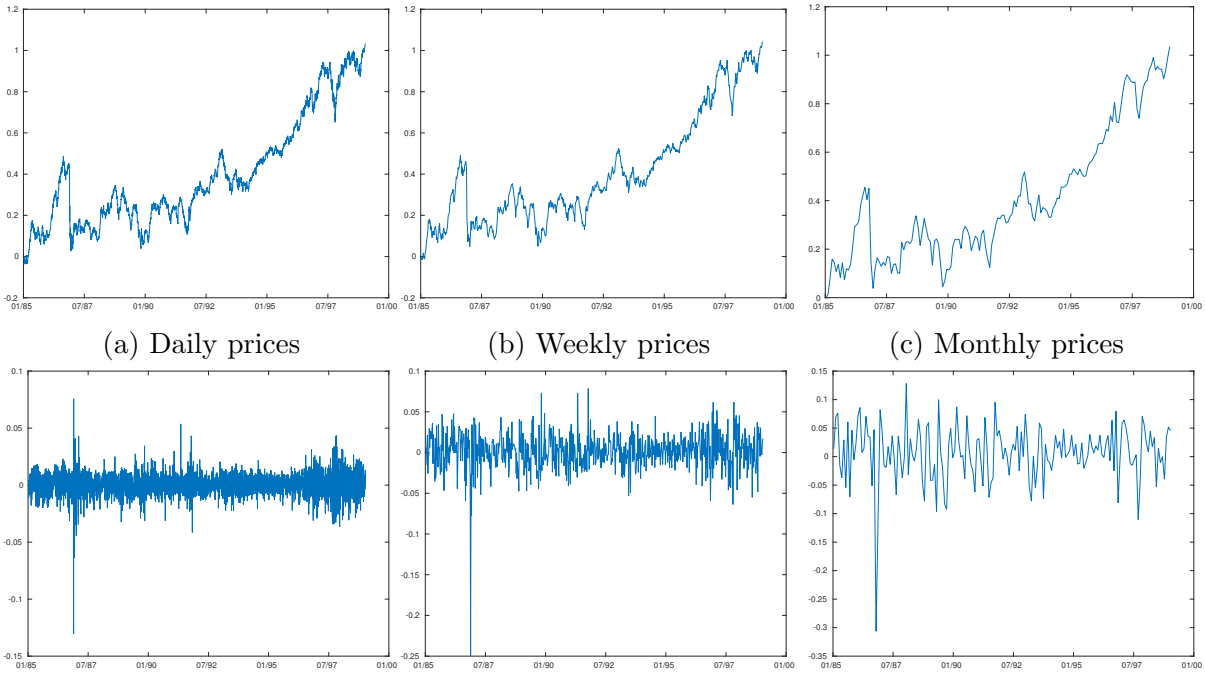

(a) Daily prices (b) Weekly prices (c) Monthly prices

(d) Daily returns (e) Weekly returns (f) Monthly returns

Figure 5: Time series plot of FTSE index prices and returns, December 1985 to December 1999.

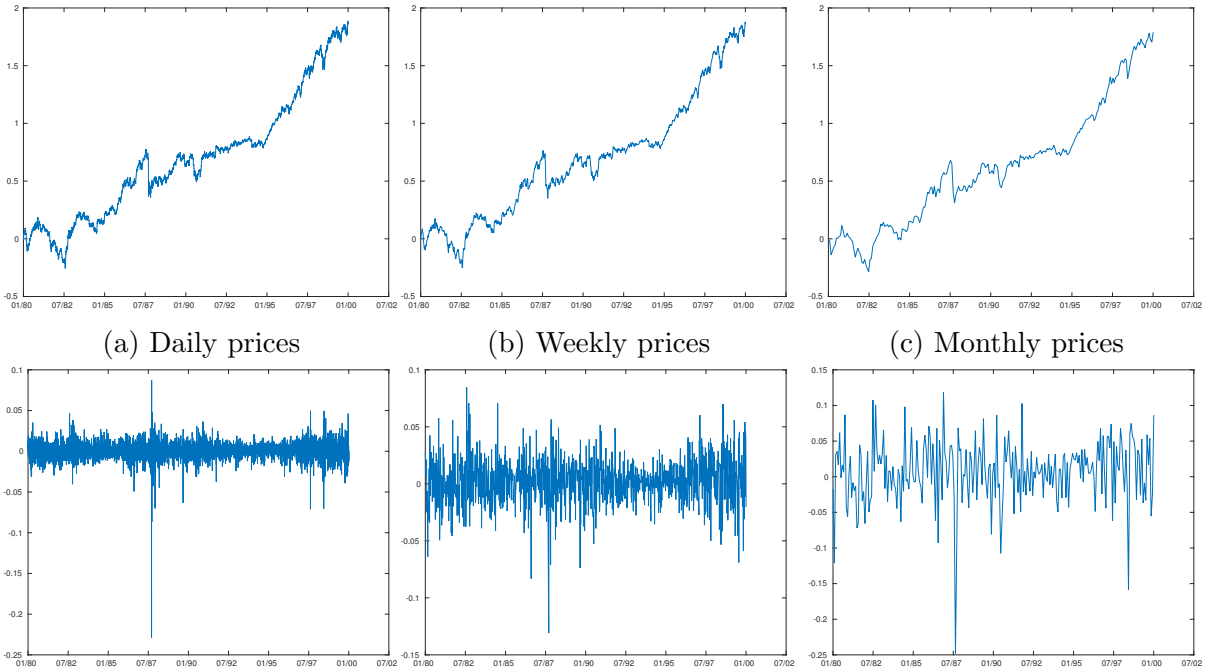

(a) Daily prices (b) Weekly prices (c) Monthly prices

(d) Daily returns (e) Weekly returns (f) Monthly returns

Figure 6: Time Series plot of S&P500 index prices and returns, January 1980 to March 2000.

	supDF	**supBZ**	$\mathcal{U}$
FTSE Daily	0.288	0.016	0.046
FTSE Weekly	0.275	0.146	0.201
FTSE Monthly	0.477	0.279	0.315
SP500 Daily	0.267	0.000	0.003
SP500 Weekly	0.170	0.002	0.011
SP500 Monthly	0.153	0.044	0.071

Table 1: Bootstrap p-values of supDF, supBZ and $\mathcal{U}$.