

This is the final peer-reviewed accepted manuscript of:

Cavaliere, G., Nielsen, M. Ø., & Robert Taylor, A. M. (2022). Adaptive inference in heteroscedastic fractional time series models. *Journal of Business & Economic Statistics*, 40(1), 50-65.

The final published version is available online at:

<https://doi.org/10.1080/07350015.2020.1773275>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

ADAPTIVE INFERENCE IN HETEROSKEDASTIC FRACTIONAL TIME SERIES MODELS

Giuseppe Cavaliere^a, Morten Ørregaard Nielsen^b and A.M. Robert Taylor^{c,*}

^a Department of Economics, University of Bologna and Department of Economics, Exeter Business School

^b Department of Economics, Queen's University and CREATES

^c Essex Business School, University of Essex

May 7, 2020

Abstract

We consider estimation and inference in fractionally integrated time series models driven by shocks which can display conditional and unconditional heteroskedasticity of unknown form. Although the standard conditional sum-of-squares (CSS) estimator remains consistent and asymptotically normal in such cases, unconditional heteroskedasticity inflates its variance matrix by a scalar quantity, $\lambda > 1$, thereby inducing a loss in efficiency relative to the unconditionally homoskedastic case, $\lambda = 1$. We propose an adaptive version of the CSS estimator, based on non-parametric kernel-based estimation of the unconditional volatility process. We show that adaptive estimation eliminates the factor λ from the variance matrix, thereby delivering the same asymptotic efficiency as that attained by the standard CSS estimator in the unconditionally homoskedastic case and, hence, asymptotic efficiency under Gaussianity. Importantly, the asymptotic analysis is based on a novel proof strategy, which does not require consistent estimation (in the sup norm) of the volatility process. Consequently, we are able to work under a weaker set of assumptions than those employed in the extant literature. The asymptotic variance matrices of both the standard and adaptive CSS estimators depend on any weak parametric autocorrelation present in the fractional model and any conditional heteroskedasticity in the shocks. Consequently, asymptotically pivotal inference can be achieved through the development of confidence regions or hypothesis tests using either heteroskedasticity-robust standard errors and/or a wild bootstrap. Monte Carlo simulations and empirical applications illustrate the practical usefulness of the methods proposed.

Keywords: adaptive estimation; conditional sum-of-squares; fractional integration; heteroskedasticity; quasi-maximum likelihood estimation; wild bootstrap.

JEL classification: C12, C13, C22.

*Correspondence to: Robert Taylor, Essex Business School, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom. Email: robert.taylor@essex.ac.uk

1 Introduction

Fractionally integrated time series models have proved highly effective in a wide range of fields of application including economics, finance, internet modeling, hydrology, climate studies, linguistics, opinion polling and DNA sequencing to name but a few; see, for example, the survey by Samorodnitsky (2007), the books by Giraitis *et al.* (2012) and Beran *et al.* (2013), and references therein.

As a result, a great many papers have appeared in the literature addressing estimation of and inference on the parameters of fractional time series models. The earliest work developed semi-parametric estimation methods; e.g., Geweke and Porter-Hudak (1983) and Robinson (1995). More relevant to this paper, parametric maximum likelihood and quasi-maximum likelihood (QML) based estimation and inference methods have been developed in, among others, Fox and Taqqu (1986), Robinson (1994), Tanaka (1999), Johansen and Nielsen (2010), Hualde and Robinson (2011), and Nielsen (2015). Of these estimation methods, the (conditional) QML estimator - equivalently the conditional sum-of-squares (CSS) estimator - was shown by Hualde and Robinson (2011) and Nielsen (2015) to be consistent over an arbitrarily large set of admissible values of the long memory parameter, thereby solving a long-standing problem arising from the non-uniform convergence of the objective function when the range of values the long memory parameter may take is large.

All of the papers cited above assume that the shocks driving the model are unconditionally homoskedastic (and in many cases the even stronger assumption that the shocks are conditionally homoskedastic or even independent and identically distributed, IID). That is, they do not allow for shocks which display permanent changes in their unconditional variance over time, what is often referred to as non-stationary volatility. Non-stationary volatility appears to be a relevant data phenomenon in a range of applied subject areas. For example, Sensier and van Dijk (2004) report that a large variety of both real and nominal economic and financial variables reject the null of constant unconditional variance. Many empirical studies also report a substantial decline, often referred to as the Great Moderation, in the unconditional volatility of the shocks driving macroeconomic series in the twenty years or so leading up to the Great Recession that started in late 2007, with a subsequent sharp increase again in volatility observed after 2007; see, *inter alia*, McConnell and Perez-Quiros (2000), Stock and Watson (2012), and the references therein. Fractional integration models have been very successfully applied in financial volatility modeling, particularly for the realised variance [RV]; see, e.g., Andersen *et al.* (2007) and references therein.

Moreover, the variance of RV is known to be time-varying; see Barndorff-Nielsen and Shephard (2002) for theoretical arguments and Corsi *et al.* (2008) for empirical evidence on this point.

In spite of this there are still relatively few papers which consider estimation and inference methods for fractionally integrated models driven by unconditionally heteroskedastic shocks. Harris and Kew (2017) extend the score-based tests of Robinson (1994) to develop LM-type tests on the value of the long memory parameter in the context of an unconditionally heteroskedastic, but conditionally homoskedastic, fractionally integrated model. Cavaliere, Nielsen and Taylor (2017, henceforth CNT) demonstrate that the QML estimator of Hualde and Robinson (2011), originally developed under the assumption of conditional homoskedasticity, retains its global consistency and asymptotic normality properties in cases where the innovations display non-stationary volatility and/or conditional heteroskedasticity. However, CNT also demonstrate that non-stationary volatility inflates the limiting covariance matrix of the QML estimator by a scalar quantity $\lambda > 1$ relative to the unconditionally homoskedastic case, thereby implying a loss in asymptotic efficiency. As a result, standard hypothesis tests based on the QML estimate lose asymptotic efficiency relative to the unconditionally homoskedastic case when non-stationary volatility is present.

The contribution of this paper to the literature is to develop asymptotically efficient QML-based methods of estimation and inference for the long and/or short memory parameters of univariate fractionally integrated time series models driven by shocks which display non-stationary volatility. To the best of our knowledge the methods developed in this paper are the first to provide a full theory of adaptive estimation and inference, including bootstrap-based inferential theory, in a heteroskedastic fractional time series model. This also allows us to develop methods of model selection and model-based inference. Our focus is on the QML estimation principle of Hualde and Robinson (2011). We will develop two-step QML-based estimation and inference methods for fractional time series models which non-parametrically adapt to unconditional heteroskedasticity of unknown form and we demonstrate that these recover the (asymptotic) efficiency losses experienced by the standard QML-based methods estimator under non-stationary volatility identified by CNT. In the first step, the unconditional variance process is estimated using a kernel-based non-parametric regression on the squares of the residuals which obtain on fitting the model using the standard QML estimator. In the second step, the sum of squares criterion, which is minimised to deliver the standard QML estimator, is scaled by the estimated volatility process, and is then subsequently minimised. We will refer to this estimator as the adaptive CSS [ACSS] estimator.

A very large and well-established literature exists relating to the use of kernel-based methods for estimating local variance functions; see, among many others, Fan and Yao (1998), Hall and Tajividi (2000), Yu and Jones (2004) and Xu and Phillips (2011). Adaptive inference based on using such an estimate of the unconditional variance was proposed in the context of inference on the parameters of finite-order unconditionally heteroskedastic but conditionally homoskedastic autoregressive models by Xu and Phillips (2008). More generally, adaptive estimators and tests designed to account for non-parametric heteroskedasticity have been widely used in the time series literature; see among others, Carroll (1982), Robinson (1987), Harvey and Robinson (1988), Hansen (1995), Xu and Phillips (2011) and Xu and Yang (2015). Adaptive methods have also been used in the context of testing for autoregressive conditional heteroskedasticity [ARCH] effects in unconditionally heteroskedastic autoregressive models by Patilea and Raïssi (2014), and for the adaptive estimation of VAR models in Patilea and Raïssi (2012, 2013). Gou and Koul (2008) consider adaptive estimation of the unconditional variance in the context of heteroskedastic regression models which display long memory in the regressors and regression disturbances. Of most relevance to the present paper is Harris and Kew (2017) whose LM-type tests on the value of the long memory parameter are based around an adaptive estimate of the unconditional volatility process. Importantly, however, Harris and Kew (2017) develop only LM tests and do not consider estimation. Moreover, they do not allow for the presence of conditional heteroskedasticity and impose considerably stronger conditions on the unconditional volatility process than we can allow for here; see Remark 3.4 for further details.

Under suitable conditions, we demonstrate that our proposed ACSS estimator is asymptotically equivalent to an infeasible estimator obtained by minimising the sum of squares criterion when divided by the true (unknown) volatility process. As a consequence, the estimator is asymptotically efficient under Gaussianity. We further demonstrate the global consistency of the adaptive estimator and show that it attains the same limiting distribution as would be attained by the standard QML estimator under unconditional homoskedasticity, other things being equal. A key consequence of this result is that, while there is no loss of asymptotic efficiency from using the ACSS estimator rather than the standard QML estimator in the conditionally homoskedastic case, efficiency gains relative to the QML estimator will be obtained where the innovations display non-stationary volatility. Although the limiting distribution of the ACSS estimator does not depend on any non-stationary volatility present in the shocks, our results show that it does, like the standard QML estimator, depend on any conditional heteroskedasticity present in the shocks and on

any weak parametric autocorrelation present in the model. Interval estimation for (functions of) the long run and/or short run parameters of the fractional time series model based on the ACSS estimator must therefore either be based on the use of heteroskedasticity robust standard errors or by using an asymptotically valid bootstrap method. We investigate both approaches, and for the latter we propose a wild bootstrap implementation of the ACSS estimator and show that this is asymptotically valid. We also develop associated heteroskedasticity-robust t and Wald statistics for testing hypotheses on these parameters. We demonstrate that these have standard limiting null distributions under our assumptions and may also be validly bootstrapped. The finite sample performance of our proposed methods are explored using Monte Carlo simulation.

In an important related literature, Baillie, Chung, and Tieslau (1996), Ling and Li (1997), Li, Ling, and McAleer (2002) and Ling (2003), among others, consider efficient maximum likelihood estimation of an ARFIMA model in the presence of parametric generalised ARCH [GARCH] models under Gaussianity, in each case assuming unconditional homoskedasticity. The aim of these authors is different from ours here which is to develop adaptive methods of inference valid under general forms of both conditional and unconditional heteroskedasticity without requiring the practitioner to specify a parametric model for either form of heteroskedasticity. As our results show, adaptive methods can deliver asymptotically efficient inference under unconditional heteroskedasticity. To develop efficient inference in the presence of conditional heteroskedasticity, a parametric model would have to be fitted, similarly to the approach adopted in the papers cited above. This will only be asymptotically efficient if the correct parametric model for the conditional heteroskedasticity is chosen and has the potential to behave very poorly where a misspecified model is chosen.

The remainder of the paper is structured as follows. Section 2 outlines our heteroskedastic fractional time series model and our main assumptions. Section 3 outlines the properties of the QML estimator under our set-up, details our ACSS estimator and establishes its large sample properties. Section 4 develops methods of inference based on confidence intervals, heteroskedasticity-robust t and Wald statistics, and wild bootstrap implementations thereof. Monte Carlo simulation results are given in Section 5. Data examples relating to realised variance and government bond yields are reported in Section 6. Section 7 concludes. All mathematical proofs and additional data analysis are reported in the Supplementary Appendix; Cavaliere, Nielsen and Taylor (2019). Equations, lemmas, etc., numbered S.x can be found in the Supplementary Appendix.

NOTATION. We use c to denote a generic, finite constant, and $\|\cdot\|, \|\cdot\|_r$ to denote the Euclidean and

L_r -norms, respectively. A function $f(x) : \mathbb{R}^q \rightarrow \mathbb{R}$ satisfies a Lipschitz condition of order α , or is in $\text{Lip}(\alpha)$, if there exists a $c > 0$ such that $|f(x_1) - f(x_2)| \leq c\|x_1 - x_2\|^\alpha$ for all $x_1, x_2 \in \mathbb{R}^q$. We use \xrightarrow{w} , \xrightarrow{p} and $\xrightarrow{L_r}$ to denote convergence in distribution, in probability, and in L_r -norm, respectively, as the sample size $T \rightarrow \infty$. The probability and expectation conditional on the realisation of the original sample are denoted P^* and E^* , respectively. For a sequence X_T^* computed on the bootstrap data, $X_T^* \xrightarrow{p^*} 0$ or $X_T^* = o_p^*(1)$, in probability, denote that $P^*(|X_T^*| > \epsilon) \rightarrow 0$ in probability for any $\epsilon > 0$, $X_T^* = O_p^*(1)$, in probability, denotes that there exists a $c > 0$ such that $P^*(|X_T^*| > c) \rightarrow 0$ in probability, and $\xrightarrow{w^*}$ denotes weak convergence in probability, in each case as $T \rightarrow \infty$.

2 The Heteroskedastic Fractional Model and Assumptions

We consider the (type II) fractional time series model

$$X_t = \Delta_+^{-d} u_t \text{ with } u_t = a(L, \psi) \varepsilon_t, \quad (1)$$

where L is the usual lag operator, and the operator Δ_+^{-d} is given, for a generic variable x_t , by $\Delta_+^{-d} x_t := \Delta^{-d} x_t \mathbb{I}(t \geq 1) = \sum_{n=0}^{t-1} \pi_n(d) x_{t-n}$, where $\mathbb{I}(\cdot)$ denotes the indicator function, and with $\pi_n(d) := \frac{\Gamma(d+n)}{\Gamma(d)\Gamma(1+n)} = \frac{d(d+1)\dots(d+n-1)}{n!}$ denoting the coefficients in the usual binomial expansion of $(1-z)^{-d}$, and where ψ is a p -dimensional parameter vector and $a(z, \psi) := \sum_{n=0}^{\infty} a_n(\psi) z^n$. We let $\theta := (d, \psi)'$ denote the full parameter vector. The parametric form of the function $a(z, \psi)$ will be assumed known, so that, specifically, u_t is assumed to be a linear process governed by an underlying unknown p -dimensional parameter vector, ψ . For example, any process that can be written as a finite order ARMA model is permitted, as is the exponential spectrum model of Bloomfield (1973). Further discussion on $a(z, \psi)$ can be found in Hualde and Robinson (2011). Thus, our focus is model-based inference (on either the long memory parameter, d , or the short memory parameter, ψ , or jointly on both). As such we assume a statistical model characterised by a finite-dimensional vector of parameters and the objective is one of estimation and inference on those parameters.

We now outline the assumptions that we will place on the model in (1). It is important to note that none of the assumptions which follow impose Gaussianity on (1).

Assumption 1. *The innovations $\{\varepsilon_t\}_{t \in \mathbb{Z}}$ are such that $\varepsilon_t = \sigma_t z_t$, where $\{z_t\}_{t \in \mathbb{Z}}$ and $\{\sigma_t\}_{t \in \mathbb{Z}}$ satisfy the conditions in parts (a) and (b), respectively, below:*

(a) $\{z_t\}_{t \in \mathbb{Z}}$ is a (conditionally heteroskedastic) martingale difference sequence with respect to

the natural filtration \mathcal{F}_t , the sigma-field generated by $\{z_s\}_{s \leq t}$, such that $\mathcal{F}_{t-1} \subseteq \mathcal{F}_t$ for $t = \dots, -1, 0, 1, 2, \dots$, and satisfies

- (i) $E(z_t^2) = 1$,
 - (ii) $\tau_{r,s} := E(z_t^2 z_{t-r} z_{t-s})$ is uniformly bounded for all $r \geq 0, s \geq 0$,
 - (iii) For all integers q such that $3 \leq q \leq 8$ and for all integers $r_1, \dots, r_{q-2} \geq 1$, the q 'th order cumulants $\kappa_q(t, t, t - r_1, \dots, t - r_{q-2})$ of $(z_t, z_t, z_{t-r_1}, \dots, z_{t-r_{q-2}})$ satisfy the condition that $\sup_t \sum_{r_1, \dots, r_{q-2}=1}^{\infty} |\kappa_q(t, t, t - r_1, \dots, t - r_{q-2})| < \infty$.
- (b) $\{\sigma_t\}_{t \in \mathbb{Z}}$ is a non-stochastic sequence satisfying
- (i) $\sup_{t \in \mathbb{Z}} \sigma_t < \infty$,
 - (ii) for all $t = 1, \dots, T$, $\sigma_t = \sigma(t/T)$, where $\sigma(\cdot) \in \mathcal{D}([0, 1])$, the space of càdlàg functions on $[0, 1]$, satisfies $\inf_{0 \leq u \leq 1} \sigma(u) > 0$.

Assumption 2. It holds that $\theta_0 = (d_0, \psi_0)' \in D \times \Psi =: \Theta$, where $D := [d_1, d_2]$ with $-\infty < d_1 \leq d_2 < \infty$ and the set $\Psi \subset \mathbb{R}^p$ is convex and compact.

Assumption 3. It holds that:

- (i) For all z in the complex unit disk $\{z \in \mathbb{C} : |z| \leq 1\}$ and for all $\psi \in \Psi$, $a(z, \psi)$ is bounded and bounded away from zero and $a_0(\psi) = 1$.
- (ii) For all $\psi \in \Psi$, $a(e^{i\omega}, \psi)$ is twice differentiable in ω with second derivative in $\text{Lip}(\zeta)$ for $\zeta > 0$.
- (iii) For all λ , $a(e^{i\omega}, \psi)$ is differentiable in ψ on $\psi \in \Psi$, and for all $\psi \in \Psi$, $\dot{a}(e^{i\omega}, \psi) := \frac{\partial a(e^{i\omega}, \psi)}{\partial \psi}$ is twice differentiable in ω with derivative in $\text{Lip}(\zeta)$ for $\zeta > 0$.
- (iv) For all ω , $a(e^{i\omega}, \psi)$ is thrice differentiable in ψ on the closed neighborhood $\mathcal{N}_\delta(\psi_0) := \{\psi \in \Psi : \|\psi - \psi_0\| \leq \delta\}$ for some $\delta > 0$, and for all $\psi \in \mathcal{N}_\delta(\psi_0)$, the second and third derivatives of $a(e^{i\omega}, \psi)$ with respect to ψ are twice differentiable in ω with derivative in $\text{Lip}(\zeta)$ for $\zeta > 0$.

Assumption 4. For all $\psi \in \Psi \setminus \{\psi_0\}$ it holds that $a(z, \psi) \neq a(z, \psi_0)$ on a subset of $\{z \in \mathbb{C} : |z| = 1\}$ of positive Lebesgue measure.

Remark 2.1. Assumption 1(a), in particular part (iii), rules out long memory in z_t^2 because, for $q = 4$, it implies that $\sum_{r=1}^{\infty} |\text{Cov}(z_t^2, z_{t-r}^2)| < \infty$. However, volatility clustering, such as GARCH, is permitted by the fact that the quantity $\tau_{r,r}$ is not necessarily equal to $E(z_t^2)E(z_{t-r}^2) = 1$. Asymmetric volatility clustering is allowed for by non-zero $\tau_{r,s}$ for $r \neq s$. Statistical leverage is also permitted, which occurs when the quantity $E(z_t^2 z_{t-i})$ is non-zero for some $i \geq 1$, noting that $E(z_t^2 z_{t-i}) = E(h_t z_{t-i})$, where $h_t := E(z_t^2 | \mathcal{F}_{t-1})$ is the conditional variance function. The stated

conditions, including the summability condition on the eighth-order cumulants of ε_t , are typical, but rather weaker than, those used in the fractional literature; see, for example, Robinson (1991), Hassler *et al.* (2009), and Harris and Kew (2017). Notably, these authors impose further conditions which rule out, among other things, statistical leverage and asymmetric volatility clustering. Harris and Kew (2017) additionally impose conditional homoskedasticity on z_t . \diamond

Remark 2.2. Assumption 1 implies that the time-varying scale factor σ_t^2 corresponds to the unconditional variance of ε_t . Thus, both the conditional and the unconditional variance of ε_t are allowed to display time-varying behaviour under Assumption 1. \diamond

Remark 2.3. Assumption 1(b) imposes relatively mild conditions on the sequence $\{\sigma_t\}$. In particular, the càdlàg assumption on $\sigma(\cdot)$ appears much weaker than those usually applied in the literature. For example, Assumption S of Harris and Kew (2017) and Assumption (i) of Xu and Phillips (2008, p. 267) require $\sigma(\cdot)$ to satisfy a uniform first-order Lipschitz condition with at most a finite number of discontinuities. In contrast, our assumption allows a countable number of jumps, which admits an extremely wide class of potential models for the unconditional variance of ε_t (formally, $\varepsilon_t = \varepsilon_{T,t} = \sigma_{T,t}z_t$ is a triangular array, but for simplicity the subscript T is suppressed). Models of single or multiple variance shifts satisfy part (b) of Assumption 1 with $\sigma(\cdot)$ piecewise constant; for example, a one-time break in variance from σ_0^2 to σ_1^2 at time $\lfloor \tau T \rfloor$, $0 < \tau < 1$, corresponds to $\sigma(u) := \sigma_0 + (\sigma_1 - \sigma_0)\mathbb{I}(u > \tau)$. Piecewise affine functions are also permitted. \diamond

Remark 2.4. Zhang and Wu (2012) consider shocks which are ‘locally stationary’; that is, of the form $u_t = G(t/T; F_t)$, where $F_t = (\dots, \varepsilon_{t-1}, \varepsilon_t)$ is shift process (see, for example, Rosenblatt, 1959) of IID random variables $\{\varepsilon_t\}$ and G is a sufficiently smooth (Lipschitz) measurable function. It is not difficult to see that our u_t process in (1) satisfies local stationarity under Assumption 1, up to an $O_p(T^{-1})$ term, provided $\sigma(\cdot)$ is Lipschitz measurable. The main difference between the ‘locally stationary’ set-up and ours is that the former places smoothness (and moment) restrictions on the function G , whereas we impose a linear dependence structure on u_t through the function $a(z, \psi)$, leaving the scaling function $\sigma(\cdot)$ essentially unrestricted. \diamond

Remark 2.5. Assumption 2 permits the length of the interval of admissible values of the parameter d to be arbitrarily large such that the model in (1) is sufficiently general to simultaneously accommodate both non-stationary, (asymptotically) stationary, and over-differenced processes. \diamond

Remark 2.6. Assumption 3 relates to the coefficients of the linear filter $a(z, \psi)$ and is easily satisfied, for example, by stationary and invertible finite order ARMA processes. In particular, Assumptions 3(i)-(ii) ensure that u_t in (1) is an invertible short-memory process with power transfer function (scale-free spectral density) that is bounded and bounded away from zero at all frequencies. Under Assumption 3(i) the function $b(z, \psi) := \sum_{n=0}^{\infty} b_n(\psi)z^n = a(z, \psi)^{-1}$ is well-defined by its power series expansion for $|z| \leq 1 + \epsilon$ for some $\epsilon > 0$, and is also bounded and bounded away from zero on the complex unit disk and $b_0(\psi) = 1$. Under Assumption 3 the coefficients $a_n(\psi)$, $b_n(\psi)$, $\dot{a}_n(\psi) := \partial a_n(\psi)/\partial \psi$, and $\dot{b}_n(\psi) := \partial b_n(\psi)/\partial \psi$ satisfy

$$|a_n(\psi)| = O(n^{-2-\zeta}), |b_n(\psi)| = O(n^{-2-\zeta}), \|\dot{a}_n(\psi)\| = O(n^{-2-\zeta}), \|\dot{b}_n(\psi)\| = O(n^{-2-\zeta}) \quad (2)$$

uniformly in $\psi \in \Psi$; see Zygmund (2003, pp. 46 and 71). The second and third derivatives with respect to ψ satisfy the same bounds uniformly over the neighborhood $\mathcal{N}_\delta(\psi_0)$. \diamond

Remark 2.7. Assumption 3(i) coincides with Assumption A1(iv) of Hualde and Robinson (2011), while Assumption 3(ii) strengthens their Assumption A1(ii) from once differentiable in ω with derivative in $\text{Lip}(\zeta)$ for $\zeta > 1/2$, and Assumption 3(iii) strengthens their Assumption A1(iii) from continuity in ψ to differentiability. Assumption 3(iv) requires $a(z, \psi)$ to be thrice differentiable in ψ rather than the corresponding twice differentiable condition in Assumption A3(ii) of Hualde and Robinson (2011) with associated Lipschitz conditions in ω . The latter are used to obtain the bounds in (2), and also appear to be needed to obtain the corresponding bounds in Hualde and Robinson (2011, p. 3169). \diamond

Remark 2.8. Assumption 3 is assumed to apply for all ψ in the user-chosen optimizing set Ψ . For example, if u_t is an ARMA model, the set Ψ can then be chosen as any compact and convex subset of the (open) set for which the roots of the AR and MA polynomials are strictly outside the unit circle. Specifically, if u_t is modeled as a first-order AR model then Assumption 3 is clearly satisfied for all $\psi \in (-1, 1)$, and the optimizing set Ψ can be chosen as any compact and convex subset of $(-1, 1)$. \diamond

Remark 2.9. The identification condition in Assumption 4 is identical to Assumption A1(i) in Hualde and Robinson (2011) and is satisfied, for example, by all stationary and invertible finite order ARMA processes whose AR and MA polynomials do not admit any common factors. \diamond

To conclude this section we need to set up some additional notation and an associated final

assumption that will be required in the next section when stating the large sample properties of the QML estimator and of our proposed adaptive estimator. To that end, define

$$A_0 := \sum_{n,m=1}^{\infty} \tau_{n,m} \begin{bmatrix} n^{-1}m^{-1} & -\gamma_n(\psi_0)'/m \\ -\gamma_n(\psi_0)/m & \gamma_n(\psi_0)\gamma_m(\psi_0)' \end{bmatrix} \text{ and } B_0 := \sum_{n=1}^{\infty} \begin{bmatrix} n^{-2} & -\gamma_n(\psi_0)'/n \\ -\gamma_n(\psi_0)/n & \gamma_n(\psi_0)\gamma_n(\psi_0)' \end{bmatrix},$$

where $\tau_{n,m}$ is defined in Assumption 1(a)(ii) and $\gamma_n(\psi) := \sum_{m=0}^{n-1} a_m(\psi)\dot{b}_{n-m}(\psi)$. Observe that A_0 (and hence B_0) is finite because $\sum_{n=0}^{\infty} \|\gamma_n(\psi)\| < \infty$ under Assumption 3 and $\sum_{n,m=1}^{\infty} |\tau_{n,m}| < \infty$ by Assumption 1(a)(iii). The matrix B_0 coincides with the matrix A in Hualde and Robinson (2011) and derives from the autocorrelation present in the process through $a(z, \psi)$. The matrix A_0 also includes the effects of any conditional heteroskedasticity present in ε_t ; if none is present $A_0 = B_0$ because here $\tau_{n,m} = \mathbb{I}(n = m)$. Neither A_0 nor B_0 are affected by any unconditional heteroskedasticity arising from Assumption 1(b). As in Hualde and Robinson (2011), to state the limiting distribution of the QML and adaptive estimators we will require B_0 to be invertible:

Assumption 5. *The matrix B_0 is non-singular.*

This condition is satisfied by, for example, stationary and invertible ARMA processes.

3 Adaptive Estimation

Adaptive estimation requires a preliminary consistent estimator. To this end, we first we review the standard QML estimator in Section 3.1. In Section 3.2, we then detail our adaptive estimator.

3.1 Standard QML Estimation

Define the residuals

$$\varepsilon_t(\theta) := \sum_{n=0}^{t-1} b_n(\psi) \Delta_+^d X_{t-n}. \quad (3)$$

Then the conditional (on initial values) Gaussian QML estimator of θ is identical to the classical least squares or CSS estimator, which is found by minimizing the sum of squared residuals; that is,

$$\tilde{\theta} := \arg \min_{\theta \in \Theta} Q_T(\theta), \quad Q_T(\theta) := T^{-1} \sum_{t=1}^T \varepsilon_t(\theta)^2. \quad (4)$$

CNT show that if X_t is generated according to (1) under Assumptions 1–5, then it holds that

$$\sqrt{T}(\tilde{\theta} - \theta_0) \xrightarrow{w} N(0, \lambda C_0), \quad (5)$$

where $C_0 := B_0^{-1}A_0B_0^{-1}$ and $\lambda := \int_0^1 \sigma^4(s)ds / (\int_0^1 \sigma^2(s)ds)^2$.

Remark 3.1. A fairly standard conditionally homoskedastic alternative to Assumption 1 (see, for example, Hannan, 1973, and Hualde and Robinson, 2011) is one where the innovations $\{\varepsilon_t\}$ are assumed to form a conditionally homoskedastic martingale difference sequence with respect to the filtration \mathcal{F}_t , i.e., where $E(\varepsilon_t^2 | \mathcal{F}_{t-1}) = \sigma^2$ almost surely and $\sup_t E(|\varepsilon_t|^q) \leq c < \infty$ for some $q \geq 4$. Under these conditions, $A_0 = B_0$ and $\lambda = 1$ and, hence, the result in (5) reduces to the result in Theorem 2.2 of Hualde and Robinson (2011). In the case where $\lambda = 1$ and z_t in Assumption 1(a) is Gaussian (and hence i.i.d.), the QML estimator $\tilde{\theta}$ of (4) is asymptotically efficient. \diamond

Remark 3.2. If heteroskedasticity arises only through part (a) of Assumption 1 then $\lambda = 1$ and the variance matrix in (5) reduces to C_0 whose form is determined only by the weak dependence (conditional heteroskedasticity and/or weak parametric autocorrelation) present in $\{u_t\}$. Alternatively, if heteroskedasticity arises only through part (b) of Assumption 1, then C_0 reduces to B_0^{-1} which is determined solely by the form of weak parametric autocorrelation present in $\{u_t\}$. \diamond

Remark 3.3. As (5) shows, the variance of the asymptotic distribution of the standard QML estimator, $\tilde{\theta}$, depends on the scalar parameter λ , which is a measure of the degree of unconditional heteroskedasticity (non-stationary volatility) present in $\{\varepsilon_t\}$ (and, hence, in $\{u_t\}$). For an unconditionally homoskedastic process, where $\sigma(\cdot)$ is constant, $\lambda = 1$, whereas when $\sigma(\cdot)$ is non-constant, $\lambda > 1$ by the Cauchy-Schwarz inequality. Other things being equal, the variance of the asymptotic distribution of the QML estimator is therefore inflated by unconditional heteroskedasticity in $\{\varepsilon_t\}$. \diamond

Suppose for the present that $\{\sigma_t^2\}$ was known. In such circumstances, a (infeasible) weighted CSS estimate of θ could be formed as

$$\bar{\theta} := \arg \min_{\theta \in \Theta} \bar{Q}_T(\theta), \text{ where } \bar{Q}_T(\theta) := T^{-1} \sum_{t=1}^T \left(\frac{\varepsilon_t(\theta)}{\sigma_t} \right)^2. \quad (6)$$

In Theorem 1 below it is shown that if X_t is generated by (1) under Assumptions 1–5, then $\sqrt{T}(\bar{\theta} - \theta_0) \xrightarrow{w} N(0, C_0)$. The asymptotic variance matrix of the QML estimator, $\tilde{\theta}$, therefore

differs from that of the infeasible weighted CSS estimator, $\bar{\theta}$, by the factor λ . As C_0 is invariant to the function $\sigma(\cdot)$, the inefficiency of the standard QML estimator relative to the weighted CSS estimator is solely determined by λ in large samples. Where this factor is large, QML will be highly inefficient relative to the weighted CSS estimator, whereas if it is close to unity QML will lose little in efficiency and would be close to being asymptotically optimal under Gaussianity.

3.2 Adaptive QML Estimation

The weighted CSS estimator $\bar{\theta}$ in (6) is infeasible, because the true values of σ_t^2 are unknown in practice. However, it is possible to implement a feasible version of $\bar{\theta}$ that has the same asymptotic distribution as $\bar{\theta}$ following the approach used in, e.g., Xu and Phillips (2008, p. 271) using a kernel-based nonparametric estimate of σ_t^2 . To that end, we first need a preliminary estimator of θ which, although based on an assumption of homoskedasticity, is nonetheless (root- T) consistent under heteroskedasticity. The standard QML estimator, $\tilde{\theta}$ of (4), satisfies this requirement, see (5).

Defining $\tilde{\varepsilon}_t := \varepsilon_t(\tilde{\theta})$ as the standard QML residuals, our proposed feasible CSS estimator of θ is then defined as

$$\hat{\theta} := \arg \min_{\theta \in \Theta} \hat{Q}_T(\theta), \quad \hat{Q}_T(\theta) := T^{-1} \sum_{t=1}^T \frac{\varepsilon_t(\theta)^2}{\hat{\sigma}_t^2}, \quad (7)$$

where

$$\hat{\sigma}_t^2 := \sum_{i=1}^T k_{ti} \tilde{\varepsilon}_i^2 \quad \text{with} \quad k_{ti} := \frac{K\left(\frac{t-i}{Tb}\right)}{\sum_{i=1}^T K\left(\frac{t-i}{Tb}\right)}. \quad (8)$$

Here $K(\cdot)$ is a bounded, nonnegative, continuous kernel function and $b := b(T)$ is a bandwidth parameter which depends on the sample size, T . We will refer to $\hat{\theta}$ as the adaptive CSS [ACSS] estimator of θ ; this in the sense that it is a feasible version of $\bar{\theta}$ based on adaptive estimation of σ_t^2 .

To establish the large sample properties of the adaptive ACSS estimator $\hat{\theta}$ in (7), we need to impose some additional conditions on the kernel function, $K(\cdot)$, and on the bandwidth, b .

Assumption 6. *The kernel function $K(\cdot) : [-\infty, \infty] \rightarrow \mathbb{R}^+ \cup \{0\}$ is continuous and satisfies $\sup_{-\infty \leq x \leq \infty} K(x) < \infty$ and $\int_{-\infty}^{\infty} K(u) du \in (0, \infty)$.*

Assumption 7. *The bandwidth $b := b(T)$ satisfies $b + T^{-1}b^{-2} \rightarrow 0$.*

We now detail the asymptotic distribution of the ACSS estimator, $\hat{\theta}$ of (7).

Theorem 1. *Let X_t be generated according to (1) and let Assumptions 1–7 be satisfied. Then,*

$$\sqrt{T}(\bar{\theta} - \hat{\theta}) \xrightarrow{P} 0, \quad (9)$$

$$\sqrt{T}(\hat{\theta} - \theta_0) \xrightarrow{w} N(0, C_0). \quad (10)$$

Remark 3.4. Notice that Theorem 1 holds without the need to strengthen the càdlàg condition in part (b) of Assumption 1. In contrast, additional smoothness conditions on $\sigma(\cdot)$ are routinely needed in the adaptive inference literature; see Remark 2.3. This generalization with respect to previous works in the field is made possible because our method of proof for Theorem 1, which is novel in the literature, rests on showing that $T^{-1} \sum_{t=1}^T (\hat{\sigma}_t^2 - \sigma_t^2)^2 \xrightarrow{P} 0$ (see Lemma S.1), rather than establishing the stronger result that $\sup_t |\hat{\sigma}_t^2 - \sigma_t^2| \xrightarrow{P} 0$. The latter necessitates imposing additional smoothness conditions on the kernel function and, in particular, on $\sigma(\cdot)$. For instance, in addition to our Assumption 1(b), Xu and Phillips (2008) follow Cavaliere (2004) and impose that $\sigma(\cdot)$ also satisfies a Lipschitz condition except at a finite number of points of discontinuity; the same condition is also employed in Harris and Kew (2017) and in a series of papers by Patilea and Raïssi (2012, 2013, 2014). In the case of non-stationary time series, Chang and Park (2010) and Beare (2018) require twice differentiability of $\sigma(\cdot)$, while Boswijk and Zu (2018) assume continuity of $\sigma(\cdot)$. In addition, all of these papers (with the exception of Beare, 2018) also need to assume that z_t is conditionally homoskedastic or even that z_t is i.i.d. Conversely, our proof of Theorem 1 allows for conditional heteroskedasticity of a very general form. \diamond

Remark 3.5. Adaptive estimation of a function, such as $\sigma(\cdot)$ in our framework, has been pursued in various contexts in the statistics literature on long memory processes, although it is usually assumed to be smooth whereas our $\sigma(\cdot)$ is only assumed càdlàg. For example, Beran and Weiershäuser (2011) estimate a smooth (spline) time-varying location model where the errors are described by a subordinated Gaussian long memory process (for textbook discussions of Gaussian subordination methods, including key references, see, *inter alia*, Giraitis *et al.*, 2012, and Beran *et al.*, 2013). Although these authors focus on a (time-varying) mean, their approach could be implemented in order to estimate, under appropriate conditions, a time-varying volatility process. It is, however, important to stress that our approach differs from theirs, as well as from other approaches based on subordination of a Gaussian long memory process, in that we filter out the time dependence of the series, and in particular the long memory component, prior to adaptive estimation, rather

than performing adaptive estimation on the subordinated series itself. \diamond

Remark 3.6. Implementation of $\hat{\sigma}_t^2$ depends on the choice of kernel function, $K(\cdot)$, and the bandwidth, b . Commonly used kernels which satisfy Assumption 6 include the uniform, Epanechnikov, biweight and Gaussian functions. The bandwidth condition in Assumption 7 implies that $b \rightarrow 0$ but at a slower rate than $T^{-1/2}$. Noticeably, the condition does not depend on the long memory parameter, d , because adaptive estimation is applied to squared residuals, $\tilde{\varepsilon}_t^2$, in (8). As shown in Lemma S.1 and the proof of Theorem 1, these are close to the squared errors, ε_t^2 , which do not display long memory; see also Remark 2.1. In practice bandwidth selection is crucial to performance and here the data-driven method of Wong (1983), which uses cross-validation on the average squared error, could be employed. This cross-validators choice of b is the value b^* which minimises $\widehat{CV}(b) := T^{-1} \sum_{t=1}^T (\tilde{\varepsilon}_t^2 - \hat{\sigma}_t^2)^2$, but where a leave-one-out procedure is used in (8) such that the observation $\tilde{\varepsilon}_t^2$ is omitted which is done by defining $K(\frac{t-i}{Tb}) := 0$ for $t = i$ in (8). We note that, although Wong (1983) assumes independent data, there is some evidence in, e.g., Xu and Phillips (2008) and Harris and Kew (2017) that the method works well under assumptions similar to ours. \diamond

Remark 3.7. A comparison of the result in Theorem 1 with that given at the end of Section 3.1 for $\bar{\theta}$ of (6) shows that the asymptotic distribution of the ACSS estimator coincides with that of the infeasible weighted CSS estimator. The limiting distribution of the ACSS estimator is therefore seen not to depend on the non-stationary volatility process, $\sigma(\cdot)$. \diamond

Remark 3.8. Recalling that $\lambda \geq 1$, a comparison of (5) and (10) shows that there can never be a loss of asymptotic efficiency from using the ACSS estimator rather than the QML estimator (the relative efficiency being given by λ) even where the shocks are unconditionally homoskedastic. Thus, the ACSS estimator is asymptotically more efficient than the QML estimator. \diamond

Remark 3.9. Where z_t in Assumption 1(a) is conditionally homoskedastic but ε_t is unconditionally heteroskedastic the ACSS estimator recovers the asymptotic distribution attained by the QML estimator in the purely homoskedastic case. \diamond

Remark 3.10. Consider the case where z_t in Assumption 1(a) is Gaussian. Then the adaptive estimator $\hat{\theta}$ is asymptotically efficient in the sense that it has the same asymptotic variance, $C_0 = B_0^{-1}$, as obtains under homoskedasticity for the standard QML estimator $\tilde{\theta}$ in (4). As a consequence, the standard QML estimator is not efficient, even when z_t is Gaussian, except in the special case where $\lambda = 1$, due to the factor $\lambda \geq 1$ appearing in its asymptotic variance. \diamond

4 Adaptive Inference

For inference purposes, a consistent estimator of C_0 is required. In Section 4.1 we discuss such an estimator and show that it can be used to obtain asymptotically valid confidence regions for (functions of) θ . Section 4.2 discusses asymptotically pivotal hypothesis tests on the elements of θ based on this estimator, and details the asymptotic power functions of these tests under Pitman drift. Bootstrap implementations of these methods are explored in Section 4.3.

4.1 Confidence Regions

To construct adaptive confidence regions for (functions of) the elements of θ based on the limiting result in (10) we will require a consistent estimate of C_0 . Following Eicker (1967), Huber (1967), and White (1982), we consider the familiar sandwich-type estimator of C_0 ,

$$\hat{C} := \left(\left(\frac{\partial^2 Q_T(\theta)}{\partial \theta \partial \theta'} \right)^{-1} \left(T^{-1} \sum_{t=1}^T \frac{\partial q_t(\theta)}{\partial \theta} \frac{\partial q_t(\theta)}{\partial \theta'} \right) \left(\frac{\partial^2 Q_T(\theta)}{\partial \theta \partial \theta'} \right)^{-1} \right) \Bigg|_{\theta=\hat{\theta}} \quad (11)$$

with $q_t(\theta) := \varepsilon_t(\theta)^2 / \hat{\sigma}_t^2$. The consistency of \hat{C} for C_0 is formalised in Theorem 2.

Theorem 2. *Let the conditions of Theorem 1 hold. It then follows that $\hat{C} \xrightarrow{P} C_0$.*

Theorem 2, taken together with the result in (10), implies that we can construct asymptotically pivotal feasible adaptive confidence regions for the elements of θ , using \hat{C} , in the usual way. Let $f : R^{p+1} \rightarrow R^q$ be a (possibly non-linear) function which is continuously differentiable at θ_0 and let $\xi := f(\theta)$ denote the parameter of interest. The adaptive estimator of ξ is given by $\hat{\xi} := f(\hat{\theta})$. Given previous results, the asymptotic distribution of $\hat{\xi}$ can straightforwardly be shown, using the delta method, to be $\sqrt{T}(\hat{\xi} - \xi_0) \xrightarrow{w} N(0, F(\theta_0)C_0F(\theta_0)')$, where $F(\theta) := \frac{\partial}{\partial \theta} f(\theta)$ is the Jacobian of the function $f(\theta)$. Confidence regions for ξ can then be formed using a consistent estimator of $F(\theta_0)C_0F(\theta_0)'$, an obvious candidate for which is $F(\hat{\theta})\hat{C}F(\hat{\theta})'$.

4.2 Hypothesis Testing and Local Power Considerations

To complement the material on feasible adaptive confidence regions, we now discuss adaptive tests of general hypotheses on the elements of θ . To that end, suppose we wish to test the null hypothesis,

$$H_0 : \xi := f(\theta) = 0, \quad (12)$$

where $f(\theta)$ is as defined below Theorem 2, and consider the sequence of local (Pitman) alternatives,

$$H_{1,T} : \xi_T = \delta/\sqrt{T}, \quad (13)$$

where δ is a fixed q -vector. In the case where we wish to test for linear restrictions on θ , we would set $f(\theta) = M'\theta - m$ where M is a $(p+1) \times q$ full-rank matrix of constants defining q (linearly independent) restrictions on the parameter vector θ and m is a q -vector of constants. An obvious example involves testing hypotheses on the long memory parameter d , important cases thereof are $d = 0$ (short memory), $d = 0.5$ (the series being weakly stationary, in the absence of unconditional heteroskedasticity, if $d < 0.5$) and $d = 1$ (unit root). As a second example, testing hypotheses on the elements of ψ could be used for order determination for the short memory dynamics, such as establishing an autoregressive order. Finally, joint hypotheses involving both d and ψ can be tested; for example, $d = 1 \cap \psi = 0$ corresponds to the pure (possibly heteroskedastic) random walk hypothesis, while $d = 0 \cap \psi = 0$ yields a martingale difference sequence.

The null hypothesis in (12) can be tested using the familiar Wald statistic,

$$W_T := T f(\hat{\theta})' (F(\hat{\theta}) \hat{C} F(\hat{\theta})')^{-1} f(\hat{\theta}), \quad (14)$$

rejecting H_0 in favour of $H_1 : \xi \neq 0$ for large values of W_T . Where only a single restriction is being tested, so that $q = 1$, one can also use the t -type statistic

$$t_T := \frac{\sqrt{T} f(\hat{\theta})}{\sqrt{F(\hat{\theta}) \hat{C} F(\hat{\theta})'}} \quad (15)$$

with H_0 rejected in favour of H_1 for large absolute values of this statistic. The statistic in (15) can also be used to test H_0 against one-sided alternatives of the form $H_{1,L} : \xi < 0$ and $H_{1,U} : \xi > 0$ by rejecting for large negative and large positive values of t_T , respectively. A familiar special case of t_T in (15) which obtains for testing the simple null hypothesis that the i 'th element of θ is equal to some hypothesised value m_i , $H_{0,i} : \theta_i - m_i = 0$ say, is given by $t_{i,T} := T^{1/2}(\hat{\theta}_i - m_i)/(\hat{C}_{ii})^{1/2}$, which again can be performed as either a one-sided or two-sided test.

The asymptotic distributions of W_T and t_T under H_0 follow immediately from Theorems 1 and 2, which we state as a corollary.

Corollary 1. *Let the conditions of Theorem 1 hold. Then, under H_0 of (12) and provided $F(\theta_0)$ is of full row rank, $W_T \xrightarrow{w} \chi^2(q)$ and $t_T \xrightarrow{w} N(0, 1)$.*

As an obvious consequence of Corollary 1, critical regions for the tests are found from standard tables, and hence the tests are easily implemented in practice.

We proceed to discuss asymptotic local power and optimality of the tests under the assumption that z_t is Gaussian; c.f. Remark 3.10 relating to efficiency of the estimator. The following corollary is implied by Theorems 1 and 2 and Le Cam's Third Lemma.

Corollary 2. *Let the conditions of Theorem 1 hold and assume also that z_t is Gaussian. Then, under $H_{1,T}$ of (13) and provided $F(\theta_0)$ is of full row rank, it holds that a one-sided test based on the t_T statistic is asymptotically uniformly most powerful (UMP) while a two-sided test based on t_T will be an asymptotically UMP unbiased. Specifically, $W_T \xrightarrow{w} \chi_q^2(\delta'(F(\theta_0)B_0^{-1}F(\theta_0)')^{-1}\delta)$ and $t_T \xrightarrow{w} N(\delta(F(\theta_0)B_0^{-1}F(\theta_0)')^{-1/2}, 1)$, where $\chi_q^2(g)$ indicates a noncentral χ_q^2 distribution with noncentrality parameter g .*

Remark 4.1. As with Remark 3.10, Corollary 2 imposes Gaussianity, and hence conditional homoskedasticity, on z_t , so that $A_0 = B_0$ and $C_0 = B_0^{-1}$. Consequently, the limiting distributions given in Corollary 2 for the W_T and t_T statistics, and as a result the asymptotic local power functions of the tests based on these statistics, coincide with those which would obtain for the corresponding statistics in the homoskedastic Gaussian case, and the optimality statements follow. That is, even in the presence of heteroskedasticity of the form in Assumption 1(b), and regardless of the value of λ , the tests based on W_T and t_T achieve the same asymptotic local power as in the homoskedastic Gaussian case. Examples of these asymptotic local power functions, showing the impact of λ , are graphed in Cavaliere *et al.* (2015). \diamond

Remark 4.2. Asymptotically pivotal test statistics can also be constructed based on the QML estimator, $\tilde{\theta}$, using the sandwich estimator \tilde{C} defined as in (11), but with $Q_T(\theta)$ given by (4) and $q_t(\theta) = \varepsilon_t(\theta)^2$ and evaluated at $\theta = \tilde{\theta}$. It is shown in Theorem 2 of CNT that if X_t is generated according to (1) under Assumptions 1–5 then $\tilde{C} - \lambda C_0 \xrightarrow{p} 0$. Consequently, defining \tilde{W}_T and \tilde{t}_T as in (14) and (15) but now based on $\tilde{\theta}$ and \tilde{C} , it follows that, under the additional assumption that z_t is Gaussian, $\tilde{W}_T \xrightarrow{w} \chi_q^2(\lambda^{-1}\delta'(F(\theta_0)B_0^{-1}F(\theta_0)')^{-1}\delta)$ and $\tilde{t}_T \xrightarrow{w} N(\lambda^{-1/2}\delta(F(\theta_0)B_0^{-1}F(\theta_0)')^{-1/2}, 1)$. It is seen from a comparison with the results in Corollary 2 that the noncentrality parameters of

tests based on our ACSS estimator are larger than those based on the QML estimator by a factor of $\lambda \geq 1$ (for W_T) or $\lambda^{1/2} \geq 1$ (for t_T) and so asymptotic local power is correspondingly higher. \diamond

Remark 4.3. Although possessing standard limiting null distributions, cf Corollary 1, it is seen from the results in Corollary 2 that tests based on W_T and t_T will have asymptotic local power functions that depend on any weak dependence present in u_t . \diamond

4.3 Bootstrap Methods

As an alternative to the asymptotic approach to forming confidence regions for θ outlined in Section 4.1, we now consider bootstrap-based methods of constructing confidence regions for θ . We will subsequently also explore bootstrap implementations of the robust Wald and t tests from Section 4.2. To that end, we first outline our proposed bootstrap algorithm. Because we allow for the presence of conditional heteroskedasticity under Assumption 1, we use a wild bootstrap-based approach (Wu, 1986). Specifically, with $\hat{\varepsilon}_t := \varepsilon_t(\hat{\theta})$ denoting the residuals based on the ACSS estimate, $\hat{\theta}$, we construct the bootstrap innovations $\varepsilon_t^* := \hat{\varepsilon}_t w_t$, where w_t , $t = 1, \dots, T$, is an i.i.d. sequence with $E(w_t) = 0$, $E(w_t^2) = 1$ and $E(w_t^{16}) < \infty$, setting $\varepsilon_t^* = 0$ for $t \leq 0$. Then the bootstrap sample $\{X_t^*\}$ is generated from the recursion

$$X_t^* := \Delta_+^{-d} u_t^* \text{ with } u_t^* := a(L, \hat{\psi}) \varepsilon_t^*, \quad t = 1, \dots, T, \quad (16)$$

and the bootstrap ACSS estimator is given by

$$\hat{\theta}^* := \arg \min_{\theta \in \Theta} \hat{Q}_T^*(\theta), \quad \hat{Q}_T^*(\theta) := T^{-1} \sum_{t=1}^T \frac{\varepsilon_t^*(\theta)^2}{\hat{\sigma}_t^{*2}}, \quad (17)$$

where

$$\varepsilon_t^*(\theta) := \sum_{n=0}^{t-1} b_n(\psi) \Delta_+^d X_{t-n}^* \quad (18)$$

and $\hat{\sigma}_t^{*2}$ is defined as the estimator (8) computed from the residuals $\tilde{\varepsilon}_t^* := \varepsilon_t^*(\tilde{\theta}^*)$ with $\tilde{\theta}^*$ denoting the preliminary standard QML estimator computed on the bootstrap sample.

Remark 4.4. The assumption that $E(w_t^{16}) < \infty$ is not restrictive in practice as it is satisfied by all common choices of the distribution of w_t , e.g., Gaussian, Rademacher, and other two-point distributions (Mammen, 1993, or Liu, 1988). \diamond

Remark 4.5. Notice that $\hat{\theta}^*$ employs an unrestricted estimate of θ in constructing the bootstrap data. Because the bootstrap data generating process is then based on $\hat{\theta}$, it is the distribution of $\sqrt{T}(\hat{\theta}^* - \hat{\theta})$, conditional on the original data, that will be used to approximate that of $\sqrt{T}(\hat{\theta} - \theta_0)$. As is standard, the former can be approximated numerically to any desired degree of accuracy. \diamond

Remark 4.6. An alternative to $\hat{\theta}^*$ is to calculate the following bootstrap estimator,

$$\check{\theta}^* := \arg \min_{\theta \in \Theta} \check{Q}_T^*(\theta), \quad \check{Q}_T^*(\theta) := T^{-1} \sum_{t=1}^T \frac{\varepsilon_t^*(\theta)^2}{\hat{\sigma}_t^2}, \quad (19)$$

where $\hat{\sigma}_t^2$ is the kernel-based estimator of σ_t^2 computed on the original data. In contrast, $\hat{\theta}^*$ is based on the bootstrap analogue of the adaptive estimator, $\hat{\sigma}_t^{*2}$. The estimator in (19) has the advantage of being computationally significantly less intensive than $\hat{\theta}^*$ because it eliminates the need to calculate the preliminary QML estimator on each bootstrap sample, but would be expected to display inferior finite sample performance. All of the large sample results given below for $\hat{\theta}^*$ also hold for $\check{\theta}^*$. \diamond

We are now in a position to establish the large-sample distribution theory for $\hat{\theta}^*$ and its computationally simpler analogue, $\check{\theta}^*$. By analogy to the use of \hat{C} to estimate C_0 in the case of the original data, we also consider the bootstrap analogue of \hat{C} , defined as

$$\hat{C}^* := \left(\left(\frac{\partial^2 \hat{Q}_T^*(\theta)}{\partial \theta \partial \theta'} \right)^{-1} \left(T^{-1} \sum_{t=1}^T \frac{\partial q_t^*(\theta)}{\partial \theta} \frac{\partial q_t^*(\theta)}{\partial \theta'} \right) \left(\frac{\partial^2 \hat{Q}_T^*(\theta)}{\partial \theta \partial \theta'} \right)^{-1} \right) \Bigg|_{\theta = \hat{\theta}^*} \quad (20)$$

with $q_t^*(\theta) := \frac{\varepsilon_t^*(\theta)^2}{\hat{\sigma}_t^{*2}}$. The variance estimator corresponding to $\check{\theta}^*$, denoted \check{C}^* , is defined accordingly.

Theorem 3. *Let Assumptions 1–7 be satisfied and assume that $\theta_0 \in \text{int}(\Theta)$. Then,*

$$\sqrt{T}(\hat{\theta}^* - \hat{\theta}) \xrightarrow{w_p^*} N(0, C_0^\dagger) \text{ and } \hat{C}^* \xrightarrow{p_p^*} C_0^\dagger, \quad (21)$$

where $C_0^\dagger := B_0^{-1} A_0^\dagger B_0^{-1}$ with $A_0^\dagger := \sum_{n=1}^{\infty} \tau_{n,n}$ $\begin{bmatrix} n^{-2} & -\gamma_n(\psi_0)' / n \\ -\gamma_n(\psi_0) / n & \gamma_n(\psi_0) \gamma_n(\psi_0)' \end{bmatrix}$. Furthermore,

$$\sqrt{T}(\check{\theta}^* - \hat{\theta}^*) \xrightarrow{p_p^*} 0 \text{ and } \check{C}^* - \hat{C}^* \xrightarrow{p_p^*} 0. \quad (22)$$

The large sample result in (21) can be used as a basis for developing asymptotically valid

bootstrap confidence regions and hypothesis tests for θ . We describe these for $\hat{\theta}^*$ in the following remarks; corresponding results for $\check{\theta}^*$ follow entirely analogously.

Remark 4.7. It is immediately seen from a comparison of the limiting covariance matrices which appear in (10) and (21) that bootstrap confidence regions for θ based on non-studentized quantities will be (asymptotically) valid provided $C_0^\dagger = C_0$; that is, when $\tau_{r,s} = 0$ for $r \neq s$, so that $A_0^\dagger = A_0$. This additional condition rules out certain asymmetries in the fourth-order moments of z_t , but importantly does not place any restrictions on the third-order moments of z_t and hence does not restrict leverage, see also Remark 2.1. As an example, in the case where θ is a scalar parameter, letting $\hat{\theta}_{(\alpha)}^*$ denote the α percent quantile of the bootstrap distribution of $\hat{\theta}^*$, the asymptotic $(1 - \alpha)\%$ -level naïve (or basic) and percentile bootstrap confidence intervals for θ are given by $[2\hat{\theta} - \hat{\theta}_{(1-\alpha/2)}^*; 2\hat{\theta} - \hat{\theta}_{(\alpha/2)}^*]$ and $[\hat{\theta}_{(\alpha/2)}^*, \hat{\theta}_{(1-\alpha/2)}^*]$, respectively. \diamond

Remark 4.8. The additional condition in Remark 4.7 can be avoided by bootstrapping pivotal statistics such as the studentized quantities W_T of (14) and t_T of (15), as will be considered in Remark 4.9 to follow, or studentized bootstrap confidence intervals. Because \hat{C}^* converges to the correct limiting variance as shown in (21), the fact that $C_0 \neq C_0^\dagger$ is inconsequential for the validity of bootstrap procedures as long as these are properly studentized. For example, letting $t_{(\alpha)}^*$ and $|t^*|_{(\alpha)}$ denote the α percent quantiles of the bootstrap distributions of $t_{i,T}^* = T^{1/2}(\hat{\theta}_i^* - \hat{\theta}_i)/(\hat{C}_{ii}^*)^{1/2}$ and $|t_{i,T}^*|$, respectively, the asymptotic $(1 - \alpha)\%$ -level equal-tailed and symmetric studentized (or percentile- t) bootstrap confidence intervals for θ_i are $[\hat{\theta}_i - t_{(1-\alpha/2)}^* T^{-1/2}(\hat{C}_{ii})^{1/2}; \hat{\theta}_i - t_{(\alpha/2)}^* T^{-1/2}(\hat{C}_{ii})^{1/2}]$ and $[\hat{\theta}_i - |t^*|_{(1-\alpha)} T^{-1/2}(\hat{C}_{ii})^{1/2}; \hat{\theta}_i + |t^*|_{(1-\alpha)} T^{-1/2}(\hat{C}_{ii})^{1/2}]$, respectively. As these intervals are based on studentized quantities, they do not require any additional conditions, and their (asymptotic) validity follows immediately from Theorem 3 under the conditions stated there. \diamond

Remark 4.9. Wild bootstrap analogues of the robust W_T and t_T statistics of (14) and (15), respectively, are given by

$$W_T^* := T(f(\hat{\theta}^*) - f(\hat{\theta}))'(F(\hat{\theta}^*)\hat{C}^*F(\hat{\theta}^*))^{-1}(f(\hat{\theta}^*) - f(\hat{\theta})), \quad (23)$$

$$t_T^* := \frac{\sqrt{T}(f(\hat{\theta}^*) - f(\hat{\theta}))}{\sqrt{F(\hat{\theta}^*)\hat{C}^*F(\hat{\theta}^*)}}. \quad (24)$$

It is immediate from Theorem 3 that these statistics attain the same first order limiting null distributions as their non-bootstrap counterparts. This implies that the wild bootstrap tests based

on W_T^* and t_T^* will have correct asymptotic size regardless of any conditional or unconditional heteroskedasticity (satisfying Assumption 1) present in ε_t and hence establishes their (asymptotic) validity. Again this result does not require the additional condition in Remark 4.7. \diamond

5 Monte Carlo Simulations

We report results from a simulation study comparing the finite sample properties of confidence intervals based on the asymptotic and bootstrap theory described above, for a fractionally integrated process allowing for both weak parametric autocorrelation in the model and heteroskedastic errors.

5.1 Monte Carlo Setup

The Monte Carlo data are simulated from the model in (1) with u_t generated according to either an AR(1) or an MA(1) process; that is, the errors, u_t , will satisfy

$$(1 - a_1L)u_t = \varepsilon_t \quad \text{or} \quad u_t = (1 + a_2L)\varepsilon_t, \quad (25)$$

where in each case the structure of the innovations $\varepsilon_t = \sigma_t z_t$ will be defined below. Our focus is on the finite sample behaviour of confidence intervals for the long memory parameter, d . We set $d_0 = 0$ in what follows with no loss of generality. Additional simulation results for estimators of d are reported in Section S.4 of the Supplementary Appendix (Cavaliere, Nielsen and Taylor, 2019).

We report results for asymptotic confidence intervals based on the QML estimator (reported under \tilde{d}) and the corresponding ACSS estimator (reported under \hat{d}). For the latter, $\hat{\sigma}_t$ was estimated using the Gaussian kernel and with the bandwidth parameter chosen by cross-validation as outlined in Remark 3.6. In each case the confidence intervals were based on robust standard errors, using \tilde{C} and \hat{C} (as defined in Remark 4.2 and (11)) for the QML and ACSS estimators, respectively. We also report results for symmetric studentized (or percentile- t) wild bootstrap intervals (see Remark 4.8 or Gonçalves and Kilian, 2004, p. 100) based on the QML estimator (reported under \tilde{d}^*) and based on the ACSS estimator. In the latter case, we report results based on (19) (reported under \check{d}^*) as well as results based on (17). For the results based on (17), the bandwidth is either re-determined by cross-validation for each bootstrap replication (reported under \hat{d}_1^*) or the bandwidth is simply chosen to be the same as that used for the original sample (reported under \hat{d}_2^*). For each method, we report the coverage percentage and the median length (across the Monte Carlo replications) of the confidence interval for d based on 10,000 Monte Carlo replications. All of these methods of

interval estimation are asymptotically pivotal for each of the models we will consider here.

Results are reported for $T = 100, 250$ and 500 . All confidence intervals are nominal 90% intervals. The variance estimators required in, for example, (11) were implemented using numerical derivatives. For the bootstrap implementations, 999 bootstrap replications were used and the w_t were generated according to the Rademacher distribution.

5.2 Results With Heteroskedastic, Uncorrelated Errors

We consider first the case where the errors, u_t , are not autocorrelated; i.e. it is known that $a_1 = a_2 = 0$ such that $u_t = \varepsilon_t$. This allows us to analyse the impact of heteroskedasticity on the confidence intervals, uncontaminated by the influence of autocorrelation in u_t .

The unconditional volatility process is generated according to the deterministic one-shift volatility process, $\sigma_t = v_1 + (v_2 - v_1)\mathbb{I}(t \geq \tau T)$; i.e., there is an abrupt single shift in the variance from v_1^2 to v_2^2 at time τT , for some $\tau \in (0, 1)$. In this example, $\lambda = (\tau + (1 - \tau)(v_2/v_1)^2)^{-2}(\tau + (1 - \tau)(v_2/v_1)^4)$. This function is graphed in Xu and Phillips (2008, p. 270), whereby it is seen that the variance of the QML estimator will be least inflated by either early positive ($v_2/v_1 > 1$) or late negative ($v_2/v_1 < 1$) breaks, but most inflated by either early negative or late positive breaks. Without loss of generality we normalise $v_1^2 = 1$. We let the break date vary among $\tau \in \{1/4, 3/4\}$ and the ratio $v := v_2/v_1$ among $v \in \{1/3, 1, 3\}$. Note that $v = 1$ corresponds to homoskedastic errors. These values of τ and v are motivated by the so-called Great Moderation and the recent Great Recession, as mentioned in the introduction, suggesting a decline in volatility early in the sample and an increase in volatility late in the sample, respectively.

The results with uncorrelated errors are presented in Table 1. We consider the following three models for $\{z_t\}$, in each case with $\{e_t\}$ forming an i.i.d. standard normal sequence:

$$\begin{aligned} \text{Panel A : } z_t &= e_t, & \text{Panel B : } z_t &= h_t^{1/2} e_t, h_t = 0.1 + 0.2z_{t-1}^2 + 0.79h_{t-1}, \\ \text{Panel C : } z_t &= e_t \exp(h_t), h_t = 0.936h_{t-1} + 0.424v_t, (v_t, e_t) \sim N(0, I_2). \end{aligned}$$

Panel A relates to the case where z_t is conditionally homoskedastic, while Panels B and C contain results for conditionally heteroskedastic GARCH(1,1) and first-order autoregressive stochastic volatility [ARSV] specifications for z_t , respectively. These generating mechanisms and parameter values are taken from Gonçalves and Kilian (2004), where empirical evidence documenting their

practical relevance is presented; see also Gonçalves and Kilian (2007). When z_t follows either a GARCH or an ARSV process, we simulate $T + 100$ values and discard the first 100 as initialization.

Consider first the results in the first three rows of Panels A, B and C where ε_t is unconditionally homoskedastic ($\nu = 1$). In this case, all of the reported confidence intervals have coverage rates which lie reasonably close to the nominal 90% level, albeit the standard QML estimator with asymptotic standard errors (\tilde{d}) has a coverage rate somewhat below the nominal level for $T = 100$, but the wild bootstrap-based confidence interval (\tilde{d}^*) rectifies this. Here there is also little to choose between the median confidence interval lengths, each of which decreases as the sample size increases, as would be expected given the consistency of both the QML and ACSS estimates. Where conditional heteroskedasticity is present the confidence intervals based on asymptotic standard errors do not perform as well, most notably where ε_t displays ARSV, with coverage rates consistently below the nominal level, increasingly so the smaller the sample size and for QML *vis-à-vis* ACSS. With one exception, the corresponding bootstrap confidence intervals do a good job in correcting the coverage rates. The exception is the wild bootstrap confidence interval for the ACSS estimate which uses the same bandwidth in the bootstrap samples as was estimated on the original data (\hat{d}_2^*) which has a coverage rate considerably in excess of the nominal level even for $T = 500$. For all methods, as expected from the results in (5) and Theorem 1, the median length of the confidence intervals is larger under conditional heteroskedasticity than under homoskedasticity. Comparing across methods, it is clear that the ACSS-based \tilde{d}^* and \hat{d}_1^* intervals perform similarly to one another and are clearly superior to both the wild bootstrap interval based on the QML estimator (\tilde{d}^*) and also to \hat{d}_2^* , in that they deliver approximately correct coverage rates and the smallest width intervals.

Consider next the results where unconditional heteroskedasticity is present in ε_t . Relative to the unconditionally homoskedastic results, we see a clear deterioration in finite sample coverage rates for the standard QML estimator with asymptotic standard errors. Other things equal, it performs worst where λ is largest. Its performance is particularly poor in the case where z_t is also an ARSV process (Panel C); here, even for $T = 500$, the coverage rate is still only around 84%. The ACSS estimator with asymptotic standard errors displays significantly better finite sample coverage, albeit coverage rates under ARSV z_t are also significantly below the nominal level. For both the QML and ACSS estimators, these effects are considerably ameliorated when implemented with a wild bootstrap. Amongst the wild bootstrap confidence intervals, the ACSS-based interval where the bandwidth is determined in each bootstrap sample using cross validation (\hat{d}_1^*) performs best with

Table 1: Simulation results with uncorrelated errors, 90% nominal intervals

τ	ν	T	Empirical coverage rate (%)						Median length					
			\tilde{d}	\tilde{d}^*	\hat{d}	\check{d}^*	\hat{d}_1^*	\hat{d}_2^*	\tilde{d}	\tilde{d}^*	\hat{d}	\check{d}^*	\hat{d}_1^*	\hat{d}_2^*
Panel A: IID errors														
	1	100	88.7	89.8	89.1	90.2	89.8	90.5	0.265	0.274	0.269	0.278	0.276	0.281
	1	250	89.0	89.4	89.2	89.6	89.5	89.8	0.166	0.167	0.167	0.168	0.168	0.169
	1	500	90.2	90.5	90.3	90.4	90.4	90.5	0.116	0.117	0.117	0.117	0.117	0.118
1/4	1/3	100	85.9	89.4	88.6	90.3	89.8	92.5	0.358	0.405	0.282	0.295	0.293	0.319
1/4	1/3	250	88.0	89.9	89.6	90.0	90.0	91.8	0.239	0.253	0.172	0.174	0.174	0.183
1/4	1/3	500	88.8	89.8	89.5	89.6	89.6	91.0	0.173	0.178	0.119	0.120	0.120	0.124
1/4	3	100	88.0	89.3	88.6	89.9	89.5	91.2	0.295	0.310	0.292	0.304	0.301	0.317
1/4	3	250	88.7	89.2	88.7	89.1	89.0	90.4	0.185	0.188	0.176	0.179	0.178	0.184
1/4	3	500	90.2	90.6	90.0	90.2	90.1	90.8	0.130	0.131	0.121	0.122	0.122	0.125
3/4	1/3	100	88.3	89.6	88.8	90.4	89.8	91.3	0.288	0.302	0.280	0.293	0.289	0.305
3/4	1/3	250	89.4	90.0	89.9	90.2	90.1	91.1	0.183	0.186	0.173	0.175	0.175	0.180
3/4	1/3	500	89.4	89.7	89.5	89.6	89.5	90.3	0.129	0.129	0.120	0.120	0.120	0.123
3/4	3	100	85.5	89.0	88.7	90.4	90.0	92.8	0.378	0.430	0.295	0.307	0.306	0.333
3/4	3	250	87.8	89.5	89.7	90.1	89.9	91.9	0.246	0.261	0.176	0.178	0.178	0.188
3/4	3	500	88.4	89.3	89.3	89.5	89.4	90.8	0.175	0.181	0.121	0.121	0.121	0.126
Panel B: GARCH errors														
	1	100	87.3	89.8	88.8	90.2	90.3	92.8	0.310	0.333	0.290	0.303	0.302	0.328
	1	250	88.2	89.5	89.8	90.3	90.3	93.8	0.216	0.226	0.180	0.183	0.183	0.205
	1	500	88.7	89.6	89.4	89.7	89.7	94.0	0.168	0.174	0.125	0.126	0.126	0.144
1/4	1/3	100	85.2	89.0	87.6	89.7	89.6	92.8	0.366	0.416	0.284	0.304	0.302	0.342
1/4	1/3	250	87.4	89.6	89.5	90.3	90.1	93.9	0.259	0.281	0.179	0.182	0.182	0.206
1/4	1/3	500	88.3	90.0	89.6	89.9	89.9	94.0	0.201	0.211	0.124	0.125	0.125	0.142
1/4	3	100	87.2	89.6	88.0	89.7	89.6	92.7	0.336	0.365	0.302	0.318	0.317	0.354
1/4	3	250	87.9	89.5	89.3	89.8	89.9	93.3	0.230	0.242	0.183	0.186	0.186	0.210
1/4	3	500	88.4	89.6	89.4	89.5	89.3	94.1	0.180	0.186	0.127	0.127	0.128	0.145
3/4	1/3	100	86.3	89.0	87.3	89.6	89.4	92.4	0.324	0.352	0.284	0.302	0.300	0.333
3/4	1/3	250	88.2	89.8	89.5	90.4	90.3	93.8	0.227	0.239	0.177	0.180	0.180	0.203
3/4	1/3	500	88.8	89.9	89.2	89.5	89.6	93.9	0.177	0.183	0.124	0.124	0.125	0.142
3/4	3	100	85.2	89.1	87.8	89.7	89.5	93.1	0.391	0.448	0.302	0.319	0.319	0.363
3/4	3	250	87.3	89.7	88.8	89.3	89.5	93.4	0.268	0.290	0.184	0.188	0.188	0.214
3/4	3	500	87.9	89.5	89.6	89.8	90.1	94.2	0.206	0.217	0.127	0.128	0.128	0.147
Panel C: SV errors														
	1	100	82.8	88.2	84.0	88.7	88.9	92.3	0.410	0.489	0.306	0.345	0.348	0.406
	1	250	83.6	88.2	86.1	88.9	89.5	94.1	0.325	0.375	0.181	0.195	0.199	0.251
	1	500	84.9	89.1	87.2	88.7	89.7	95.2	0.270	0.305	0.117	0.122	0.125	0.167
1/4	1/3	100	82.2	88.0	83.8	88.9	89.6	92.5	0.421	0.517	0.280	0.324	0.328	0.383
1/4	1/3	250	83.3	88.5	85.7	88.6	89.5	94.1	0.339	0.400	0.168	0.182	0.189	0.239
1/4	1/3	500	84.1	88.8	87.1	88.7	90.0	95.6	0.289	0.334	0.111	0.116	0.121	0.168
1/4	3	100	83.3	88.2	84.3	88.7	88.8	92.5	0.422	0.507	0.317	0.356	0.360	0.424
1/4	3	250	84.4	88.7	86.0	88.5	89.3	94.1	0.334	0.387	0.188	0.200	0.205	0.258
1/4	3	500	84.7	89.1	87.0	88.6	89.7	95.5	0.282	0.322	0.122	0.127	0.131	0.175
3/4	1/3	100	82.5	87.9	83.6	88.6	89.0	92.4	0.408	0.493	0.279	0.323	0.327	0.381
3/4	1/3	250	84.0	89.0	85.5	88.7	89.7	94.2	0.327	0.382	0.167	0.181	0.185	0.236
3/4	1/3	500	84.3	88.6	86.6	88.4	89.4	95.1	0.279	0.319	0.108	0.114	0.118	0.161
3/4	3	100	82.0	88.3	84.6	88.8	89.1	92.9	0.440	0.534	0.325	0.367	0.373	0.435
3/4	3	250	84.0	88.8	86.6	89.1	90.0	94.4	0.348	0.412	0.191	0.203	0.209	0.264
3/4	3	500	84.0	88.7	87.1	88.9	89.9	95.3	0.294	0.340	0.124	0.129	0.133	0.180

Notes: The table reports empirical coverage percentage and median length of confidence intervals for d based on 10,000 replications. The reported intervals are based on the QML estimator with robust standard errors (\tilde{d}), the wild bootstrap equivalent (\tilde{d}^*), the ACSS estimator (\hat{d}), the wild bootstrap ACSS in (19) (\check{d}^*), the wild bootstrap ACSS estimator in (17), where the bandwidth is re-determined for each bootstrap sample (\hat{d}_1^*), and the wild bootstrap ACSS in (17), using the bandwidth from the original sample on each bootstrap sample (\hat{d}_2^*). The bootstrap intervals are symmetric studentized bootstrap confidence intervals (see Remark 4.8) based on 999 bootstrap replications.

Table 2: Simulation results with AR or MA errors, 90% nominal intervals

a_1	a_2	T	Empirical coverage rate (%)						Median length					
			\tilde{d}	\tilde{d}^*	\hat{d}	\check{d}^*	\hat{d}_1^*	\hat{d}_2^*	\tilde{d}	\tilde{d}^*	\hat{d}	\check{d}^*	\hat{d}_1^*	\hat{d}_2^*
Panel A: No break, $v = 1$														
0.0		100	83.7	91.3	83.7	91.5	91.2	91.4	0.432	0.622	0.438	0.637	0.626	0.630
0.0		250	88.8	91.4	88.8	91.6	91.4	91.7	0.270	0.282	0.272	0.284	0.283	0.285
0.0		500	89.2	89.5	89.2	89.5	89.6	89.8	0.189	0.191	0.190	0.191	0.190	0.192
-0.8		100	88.3	89.8	88.8	90.2	89.8	90.4	0.284	0.300	0.289	0.304	0.301	0.306
-0.8		250	89.0	89.5	89.2	89.7	89.6	90.1	0.176	0.180	0.177	0.181	0.180	0.182
-0.8		500	89.7	89.9	89.8	89.9	90.0	90.0	0.124	0.125	0.124	0.125	0.125	0.126
0.8		100	76.4	96.8	76.8	96.9	96.9	97.1	0.461	0.921	0.466	0.941	0.922	0.941
0.8		250	79.3	96.4	79.5	96.3	96.4	96.5	0.349	0.619	0.350	0.623	0.617	0.625
0.8		500	82.8	95.6	82.8	95.7	95.7	95.8	0.282	0.414	0.283	0.416	0.415	0.416
	0.0	100	87.3	90.2	87.5	90.6	90.2	90.7	0.445	0.488	0.453	0.501	0.492	0.503
	0.0	250	89.8	90.4	90.0	90.5	90.5	90.6	0.272	0.274	0.274	0.275	0.275	0.278
	0.0	500	89.7	89.6	89.7	89.6	89.7	89.7	0.189	0.189	0.190	0.189	0.189	0.190
	-0.8	100	69.9	94.1	70.3	94.2	93.9	94.2	0.447	1.472	0.451	1.467	1.456	1.492
	-0.8	250	74.1	93.7	74.3	93.9	93.7	93.9	0.353	0.942	0.355	0.949	0.945	0.947
	-0.8	500	79.3	93.8	79.4	93.9	93.8	93.8	0.288	0.547	0.290	0.547	0.546	0.550
	0.8	100	88.9	90.0	89.2	90.2	89.8	90.4	0.290	0.300	0.295	0.305	0.300	0.307
	0.8	250	89.1	89.5	89.2	89.5	89.3	89.7	0.178	0.179	0.179	0.181	0.180	0.182
	0.8	500	90.2	90.0	90.3	90.3	90.3	90.3	0.124	0.125	0.125	0.125	0.125	0.125
Panel B: Early break, $\tau = 1/4, v = 1/3$														
0.0		100	77.7	89.5	82.5	92.4	91.9	93.3	0.512	0.791	0.427	0.629	0.610	0.645
0.0		250	84.7	92.6	88.8	91.6	91.3	92.9	0.370	0.457	0.276	0.293	0.291	0.306
0.0		500	87.9	91.2	89.4	89.7	89.4	90.8	0.272	0.294	0.192	0.194	0.193	0.201
-0.8		100	84.0	89.2	87.6	90.1	89.4	92.1	0.375	0.441	0.298	0.319	0.316	0.344
-0.8		250	87.5	89.8	89.3	90.1	89.9	91.7	0.253	0.273	0.183	0.187	0.186	0.197
-0.8		500	88.6	90.0	89.6	90.0	89.8	91.1	0.184	0.191	0.127	0.128	0.128	0.133
0.8		100	75.9	96.1	78.6	97.9	97.5	98.1	0.545	1.007	0.471	0.950	0.907	0.960
0.8		250	78.2	96.5	82.2	97.1	96.9	97.3	0.432	0.785	0.351	0.590	0.580	0.610
0.8		500	80.5	96.0	84.0	96.3	96.2	96.6	0.354	0.602	0.278	0.394	0.392	0.407
	0.0	100	84.2	89.2	87.2	90.5	90.0	91.9	0.577	0.704	0.448	0.511	0.499	0.541
	0.0	250	87.9	89.7	90.0	90.5	90.3	92.1	0.386	0.414	0.278	0.280	0.279	0.295
	0.0	500	89.1	89.8	90.1	90.2	89.9	91.2	0.277	0.282	0.192	0.192	0.190	0.199
	-0.8	100	70.8	93.2	69.8	93.4	93.2	93.8	0.518	1.812	0.440	1.353	1.333	1.507
	-0.8	250	71.4	93.3	74.3	93.1	92.8	93.2	0.410	1.269	0.351	0.873	0.866	0.915
	-0.8	500	73.4	92.4	79.7	93.5	93.5	93.8	0.344	0.869	0.285	0.506	0.502	0.525
	0.8	100	86.7	89.3	89.0	90.3	89.7	92.5	0.391	0.431	0.305	0.318	0.313	0.343
	0.8	250	87.8	89.2	89.4	89.5	89.4	91.4	0.257	0.269	0.185	0.186	0.185	0.196
	0.8	500	88.6	89.2	90.2	90.2	90.1	91.3	0.184	0.188	0.128	0.128	0.128	0.133
Panel C: Late break, $\tau = 3/4, v = 3$														
0.0		100	75.4	88.9	81.9	90.7	90.4	92.1	0.567	0.945	0.475	0.741	0.723	0.766
0.0		250	84.9	93.1	88.9	92.3	92.2	93.5	0.390	0.495	0.287	0.306	0.304	0.320
0.0		500	87.3	90.9	89.7	89.9	89.9	91.2	0.282	0.304	0.197	0.199	0.198	0.207
-0.8		100	84.1	88.7	87.9	89.7	89.8	92.3	0.401	0.472	0.316	0.337	0.333	0.365
-0.8		250	87.3	89.7	89.1	89.7	89.6	91.6	0.260	0.281	0.188	0.191	0.191	0.202
-0.8		500	88.2	89.5	89.2	89.7	89.6	91.1	0.186	0.194	0.129	0.130	0.130	0.135
0.8		100	73.8	97.0	75.2	97.1	97.0	97.5	0.572	1.188	0.491	1.053	1.029	1.119
0.8		250	76.2	96.4	78.5	96.1	96.0	96.5	0.437	0.854	0.363	0.677	0.669	0.707
0.8		500	78.1	95.7	82.0	95.6	95.7	96.0	0.362	0.648	0.290	0.446	0.442	0.463
	0.0	100	81.2	89.4	86.0	90.6	90.2	92.4	0.627	0.886	0.499	0.592	0.579	0.633
	0.0	250	87.9	90.3	90.0	91.0	90.7	92.6	0.407	0.445	0.291	0.295	0.292	0.310
	0.0	500	89.0	90.0	90.0	89.9	89.7	91.1	0.287	0.294	0.197	0.197	0.196	0.205
	-0.8	100	70.7	95.7	70.7	95.5	95.0	95.8	0.553	1.861	0.476	1.544	1.519	1.719
	-0.8	250	71.0	94.9	73.7	94.1	94.1	94.5	0.428	1.400	0.365	1.026	1.014	1.088
	-0.8	500	73.6	93.1	78.8	94.0	93.9	94.2	0.359	0.940	0.294	0.590	0.585	0.613
	0.8	100	87.1	89.7	88.9	90.2	90.0	92.7	0.417	0.468	0.322	0.335	0.333	0.366
	0.8	250	88.0	89.8	89.1	89.4	89.5	91.5	0.267	0.280	0.189	0.191	0.190	0.202
	0.8	500	89.5	90.2	90.6	90.4	90.6	91.8	0.189	0.193	0.129	0.130	0.130	0.135

Notes: See notes to Table 1.

an empirical coverage rate very close to the nominal level throughout. The median length of the QML-based confidence intervals are, other things equal, inflated (often considerably so) the larger the value of λ , consistent with the impact of unconditional heteroskedasticity on the asymptotic variance matrix of the QML estimate; see (5). To illustrate, in the case where z_t is IID (Panel A) the median length of the wild bootstrap confidence interval based on \tilde{d} ($T = 100$) increases from 0.274 under unconditional homoskedasticity to 0.430 when a late positive break in variance occurs. In contrast, the median length of the confidence intervals based on the ACSS estimator appear relatively unaffected by unconditional heteroskedasticity, as anticipated by Theorems 1 and 3. Notice also that, consistent with the large sample theory, the ratio of the median length of the wild bootstrap confidence intervals based on \tilde{d}^* (QML) and \hat{d}_1^* (ACSS) is approximately unity, regardless of T , when $v = 1$ (such that $\sqrt{\lambda} = 1$), and is 1.41, 1.47 and 1.50 for $T = 100, 250$ and 500, respectively, in the late positive break case (where $\sqrt{\lambda} \approx 1.53$).

5.3 Results With Heteroskedastic, Autocorrelated Errors

Table 2 reports results for cases where u_t can display both weak parametric autocorrelation and heteroskedasticity. Specifically, u_t is generated according to one of the models in (25) with $a_1, a_2 \in \{-0.8, 0, 0.8\}$. The reported cases where either $a_1 = 0$ or $a_2 = 0$ correspond to the situation where an AR(1) or MA(1) specification is estimated, respectively, even though it is not present in the data generating process. Results are given for where ε_t is either IID (Panel A) or displays an early negative (Panel B) or late positive (Panel C) break in unconditional variance.

Relative to the results in Table 1, autocorrelation can be seen to have a significant impact on both the coverage rate and median length of the confidence intervals. For both the QML and ACSS estimators, coverage rates based on asymptotic standard errors are not as accurate (relative to the corresponding results in Table 1) when autocorrelation is either present and modelled or not present but allowed for in the estimated model. For example, where $a_1 = 0$ (so that an AR(1) is modelled but not actually present in the data) the coverage rates for \tilde{d} and \hat{d} are both 83.7% when $T = 100$ in the homoskedastic case compared to 88.7% and 89.1%, respectively, in the corresponding case in Table 1. Where autocorrelation is present, the empirical coverage rates of \tilde{d} and \hat{d} can lie significantly below the nominal level. This is seen most obviously for the cases where u_t is either positively autocorrelated ($a_1 = 0.8$) or follows a negative moving average ($a_2 = -0.8$). To illustrate, when $a_2 = -0.8$ the coverage rates for \tilde{d} and \hat{d} are around 70% when $T = 100$, regardless

of whether ε_t is homoskedastic or contains a break in variance. As with the impact of conditional heteroskedasticity, the wild bootstrap considerably improves the coverage rates of the confidence intervals, albeit in the most problematic cases above ($a_1 = 0.8$ and $a_2 = -0.8$) the wild bootstrap tends to rather over correct such that the resulting confidence intervals are somewhat too liberal.

As with the presence of conditional heteroskedasticity, the median lengths of the confidence intervals are seen to vary with a_1 and a_2 , again as expected given the results for the QML and ACSS estimates in (5) and Theorem 1, respectively. The confidence intervals are considerably wider when either $a_1 = 0.8$ or $a_2 = -0.8$ than when $a_1 = -0.8$ or $a_2 = 0.8$. Controlling for the impact of autocorrelation, however, the results in Table 2 reveal qualitatively similar conclusions to those drawn from the results in Table 1; that is, the efficiency gains from basing confidence intervals around the ACSS estimator, \hat{d} , rather than the QML estimator, \tilde{d} are clearly visible when unconditional heteroskedasticity is present in ε_t . As with the conclusions drawn from Table 1, the ACSS-based interval with the bandwidth determined in each bootstrap sample using cross validation (\hat{d}_1^*) appears to deliver the best overall performance.

6 Data Examples

We now apply the methods discussed in this paper to a variety of data sets. All bootstrap confidence intervals were based on 999 replications using the Rademacher distribution for w_t . For each data set we report the QML and ACSS estimates of d along with their (robust or sandwich) standard errors and 95% confidence intervals based on (5) and (10), respectively, together with 95% wild bootstrap percentile- t confidence intervals (for the ACSS these re-estimate σ_t and estimate the bandwidth parameter in each bootstrap replication using cross-validation). In each case an ARFIMA($p,d,0$) model was fitted to the (de-meaned) data with p chosen by the BIC applied to the adaptive estimation model (choosing over $p = 0, \dots, 10$). Plots of all data series and additional graphical analysis of the residuals from the chosen models for each data set along with formal statistical tests for heteroskedasticity are reported in Section S.5 in the Supplementary Appendix (Cavaliere, Nielsen and Taylor, 2019). Both the conditional and unconditional heteroskedasticity tests reported display highly significant rejections for most of the series, providing strong evidence of heteroskedastic behaviour in these data sets.

Table 3: Andersen *et al.* (2007) Realised Variance Data

	p	QML						ACSS					
		Est. d	$se(d)$	ACI _L	ACI _U	WBCI _L	WBCI _U	Est. d	$se(d)$	ACI _L	ACI _U	WBCI _L	WBCI _U
S&P500	0	0.321	0.046	0.231	0.412	0.211	0.431	0.360	0.039	0.284	0.436	0.276	0.444
DM/\$	0	0.382	0.032	0.320	0.444	0.319	0.445	0.390	0.017	0.357	0.423	0.358	0.422
T-bond	0	0.186	0.015	0.156	0.215	0.155	0.216	0.198	0.015	0.168	0.228	0.168	0.228

Notes: Daily data from January 2, 1990 to December 31, 2002 (SP500 and Tbond) and December 2, 1986 to June 30, 1999 (DM/\$). The sample sizes are $T = 3262$ and $T = 3045$, respectively. An ARFIMA($p, d, 0$) model was estimated on the de-measured data. $se(d)$ is the (robust) standard error, ACI are asymptotic confidence intervals and WBCI are wild bootstrap confidence intervals (for ACSS re-estimating σ_t and estimating the bandwidth parameter in each bootstrap replication using cross-validation). All intervals are 95% nominal level.

6.1 Realised Variance Data

We first revisit a classic data set in finance from Andersen *et al.* (2007). The data set consists of daily realised variance [RV], based on intra-daily five-minute return data for the S&P500 index, DM/\$ exchange rate futures, and 30-year US Treasury bond futures. The sample periods are January 2, 1990 to December 31, 2002 for the S&P500 and T-bond data ($T = 3262$), and December 2, 1986 to June 30, 1999 for the DM/\$ data ($T = 3045$). Detailed data descriptions can be found in Andersen *et al.* (2007, Section III.A).

The results are reported in Table 3. Both the QML and ACSS estimates of d are in line with the typical 0.3–0.4 found in the literature, although the estimates for the T-bond RV data are slightly lower. For the S&P500 RV data and the DM/\$ RV data, the standard errors for ACSS are much smaller than those for QML, which translates into commensurately smaller confidence interval widths for d when based on ACSS rather than QML estimation. However, for the T-bond RV data they are the same. This is exactly as expected from the theoretical results in (5) and Theorem 1, given that the unconditional heteroskedasticity tests reject the null of unconditional homoskedasticity for the S&P500 and DM/\$ data, but do not reject for the T-bond data. Thus, for the S&P500 and DM/\$ data, we can draw much more accurate inferences from the ACSS estimator compared with the QML estimator, in the sense of smaller standard errors and narrower confidence intervals. Finally, the wild bootstrap confidence intervals are very similar to the asymptotic confidence intervals, although the former are slightly wider for the S&P500 data.

We next analyse a similar, but more recent, data set consisting of the RV for 28 of the 30 stocks in the U.S. market which comprise the Dow Jones Industrial Average (tickers TRV and V have missing data issues and, hence, are omitted). The RVs are calculated from continuously compounded

Table 4: Dow Jones Realised Variance Data

	p	QML						ACSS					
		Est. d	se(d)	ACI _L	ACI _U	WBCI _L	WBCI _U	Est. d	se(d)	ACI _L	ACI _U	WBCI _L	WBCI _U
AAPL	0	0.322	0.020	0.283	0.360	0.281	0.362	0.319	0.016	0.288	0.351	0.286	0.353
AXP	0	0.405	0.025	0.356	0.453	0.352	0.457	0.374	0.013	0.349	0.398	0.347	0.400
BA	0	0.345	0.020	0.305	0.385	0.307	0.383	0.316	0.014	0.287	0.344	0.285	0.346
CAT	0	0.359	0.023	0.315	0.404	0.316	0.403	0.336	0.015	0.306	0.365	0.305	0.366
CSCO	0	0.339	0.022	0.296	0.381	0.295	0.383	0.320	0.015	0.290	0.349	0.291	0.349
CVX	0	0.423	0.038	0.348	0.499	0.344	0.503	0.376	0.024	0.329	0.422	0.328	0.423
DD	0	0.356	0.022	0.313	0.399	0.309	0.403	0.321	0.016	0.290	0.351	0.290	0.351
DIS	0	0.382	0.024	0.335	0.430	0.336	0.428	0.350	0.018	0.315	0.385	0.315	0.385
GE	0	0.410	0.036	0.340	0.479	0.335	0.485	0.355	0.016	0.324	0.386	0.324	0.387
GS	1	0.477	0.074	0.332	0.622	0.324	0.631	0.433	0.027	0.380	0.486	0.380	0.486
HD	0	0.358	0.027	0.304	0.412	0.303	0.413	0.330	0.015	0.300	0.359	0.301	0.359
IBM	0	0.380	0.022	0.337	0.424	0.336	0.425	0.338	0.016	0.307	0.368	0.308	0.367
INTC	0	0.330	0.018	0.296	0.365	0.296	0.365	0.310	0.013	0.284	0.335	0.286	0.333
JNJ	0	0.376	0.033	0.313	0.442	0.313	0.441	0.351	0.026	0.301	0.401	0.302	0.401
JPM	0	0.459	0.039	0.382	0.536	0.383	0.535	0.396	0.017	0.363	0.428	0.364	0.428
KO	0	0.369	0.029	0.313	0.426	0.307	0.431	0.331	0.020	0.292	0.369	0.291	0.370
MCD	0	0.314	0.032	0.252	0.376	0.242	0.386	0.325	0.022	0.282	0.367	0.280	0.370
MMM	0	0.317	0.036	0.246	0.388	0.236	0.397	0.313	0.019	0.275	0.351	0.275	0.352
MRK	0	0.337	0.028	0.281	0.392	0.283	0.390	0.322	0.018	0.286	0.357	0.286	0.357
MSFT	0	0.312	0.020	0.272	0.352	0.270	0.353	0.280	0.016	0.249	0.311	0.249	0.312
NKE	0	0.349	0.024	0.303	0.396	0.302	0.397	0.319	0.016	0.288	0.350	0.288	0.350
PFE	0	0.295	0.024	0.249	0.342	0.247	0.344	0.291	0.020	0.253	0.329	0.249	0.333
PG	0	0.238	0.057	0.125	0.350	0.072	0.404	0.267	0.026	0.216	0.319	0.211	0.323
UNH	0	0.348	0.039	0.272	0.424	0.270	0.426	0.318	0.019	0.281	0.355	0.281	0.355
UTX	0	0.360	0.027	0.308	0.412	0.306	0.414	0.323	0.017	0.289	0.356	0.291	0.354
VZ	0	0.357	0.028	0.301	0.413	0.298	0.415	0.324	0.018	0.289	0.360	0.287	0.361
WMT	0	0.350	0.036	0.280	0.420	0.276	0.424	0.312	0.021	0.271	0.353	0.271	0.352
XOM	0	0.436	0.042	0.354	0.517	0.351	0.520	0.381	0.025	0.332	0.430	0.332	0.430

Notes: Daily data from January 2, 2003 to December 31, 2014 ($T = 31021$)y. An ARFIMA($p, d, 0$) model was estimated on the de-measured data. se(d) is the (robust) standard error, ACI are asymptotic confidence intervals and WBCI are wild bootstrap confidence intervals (for ACSS re-estimating σ_t and estimating the bandwidth parameter in each bootstrap replication using cross-validation). All intervals are 95% nominal level.

intra-daily five-minute returns based on prices obtained from the NYSE Trades and Quotes (TAQ) database. The sample period is January 2, 2003 to December 31, 2014 for a sample size of $T = 3021$.

The results for this data set are presented in Table 4. The overall conclusions are in line with those for the classic data set in Table 3. Specifically, the estimates of d are all about 0.3–0.4, the ACSS standard errors are smaller than the corresponding QML standard errors in every case, and the wild bootstrap and asymptotic confidence intervals are very similar with few exceptions (most notably for ticker PG). Most interestingly, for tickers CVX, GS, JPM, and XOM, the QML confidence intervals include the value $d = 1/2$, which defines the boundary between stationarity and nonstationarity in standard fractional time series models, and hence $d \geq 1/2$ cannot be excluded based on QML inference. However, for each of those series, due to a combination of slightly lower estimates of d and much narrower confidence intervals, the ACSS confidence intervals do not include

Table 5: Government Bond Yield Data

	p	QML						ACSS					
		Est. d	$se(d)$	ACI_L	ACI_U	$WBCI_L$	$WBCI_U$	Est. d	$se(d)$	ACI_L	ACI_U	$WBCI_L$	$WBCI_U$
Belgium	0	1.126	0.040	1.047	1.206	1.049	1.204	1.093	0.026	1.043	1.144	1.043	1.144
Finland	1	0.961	0.026	0.911	1.011	0.909	1.014	0.998	0.022	0.956	1.041	0.957	1.039
France	0	1.005	0.019	0.968	1.042	0.968	1.042	1.008	0.017	0.975	1.041	0.974	1.042
Germany	1	0.971	0.025	0.922	1.020	0.923	1.019	0.992	0.022	0.948	1.036	0.948	1.036
Ireland	1	1.052	0.071	0.912	1.192	0.890	1.213	1.008	0.028	0.952	1.064	0.951	1.065
Italy	1	0.896	0.042	0.814	0.979	0.814	0.979	0.928	0.026	0.877	0.979	0.877	0.979
Portugal	1	0.965	0.062	0.843	1.087	0.834	1.096	0.970	0.022	0.928	1.013	0.926	1.014
Spain	1	0.872	0.038	0.797	0.947	0.794	0.950	0.935	0.025	0.886	0.984	0.884	0.986

Notes: Daily data on 10-year government bond yields from September 5, 2005, to February 2, 2016. Sample size is $T = 2715$. An ARFIMA($p, d, 0$) model was estimated on the de-meaned data. $se(d)$ is the (robust) standard error, ACI are asymptotic confidence intervals and WBCI are wild bootstrap confidence intervals (for ACSS re-estimating σ_t and estimating the bandwidth parameter in each bootstrap replication using cross-validation). All intervals are 95% nominal level.

$d = 1/2$, so based on the ACSS we can reject that $d \geq 1/2$ at the 5% level.

6.2 Government Bond Yield Data

We next analyse the sovereign debt data series for the following eight Eurozone countries considered in Martins and Amado (2016): Belgium, Finland, France, Germany, Ireland, Italy, Portugal, and Spain. These are daily data on 10-year government bond yields from September 5, 2005, to February 2, 2016, with $T = 2715$; see Martins and Amado (2016, Section 3.1) for detailed data descriptions.

The ACSS and QML estimates of d , reported in Table 5 are similar for each country considered with estimated values for all countries lying relatively close to 1. As for the RV data sets, the ACSS standard errors are smaller, in many cases substantially so, for each country than the corresponding QML ones. To illustrate, comparing the widths of wild bootstrap confidence intervals which do not re-estimate σ_t for each bootstrap replication, we see that for Portugal the width of the ACSS-based interval is only about a third of that of the QML-based interval, while for Ireland the ACSS-based interval is about half the width of the QML-based interval. Interestingly, the differences in the widths of the ACSS- and QML-based confidence intervals do not lead to a different outcome when using these intervals to test the null hypothesis that $d = 1$ for any of the series. In particular, the unit root null hypothesis cannot be rejected at the 95% level for all series except Spain and Ireland. The former (latter) appears to display lower (higher) persistence than a unit root process.

7 Conclusions

In this paper we have discussed estimation and inference on the parameters of fractionally integrated time series models driven by shocks which can display conditional and/or unconditional heteroskedasticity. The asymptotic variance matrix of the limiting distribution of the standard QML estimator is inflated under unconditional heteroskedasticity relative to the unconditionally homoskedastic case. We have shown that an adaptive version of the QML estimator, based on a non-parametric kernel-based estimator of the unconditional variance process, attains the same asymptotic variance matrix when unconditional heteroskedasticity is present as the standard QML estimator would achieve under unconditional homoskedasticity and, hence, achieves (asymptotic) efficiency gains over the QML estimator. Under Gaussianity, the adaptive estimator is asymptotically efficient. We have shown that asymptotically pivotal inference based on the adaptive estimator can be achieved through the development of confidence regions or hypothesis tests using either heteroskedasticity robust standard errors and/or a wild bootstrap, and that these deliver gains in asymptotic efficiency/local power over their counterparts based on the standard QML estimator. Monte Carlo simulation results reported suggest that the large sample advantages of basing confidence intervals on adaptive methods carry over into finite samples. Empirical applications to realised variance and government bond yield data were provided to illustrate the improvements in inference that the adaptive methods discussed in this paper can deliver when applied to economic and financial data displaying non-constant variance.

Finally, although we have chosen to focus attention on developing inference methods for the elements of θ , our results also facilitate running standard diagnostic (misspecification) tests on the residuals from the estimated model, such as tests for unmodelled serial correlation, provided they are implemented with either heteroskedasticity-robust standard errors or the wild bootstrap.

Acknowledgements

We are grateful to the Editor, Jianqing Fan, an Associate Editor and three referees for their helpful and constructive comments on an earlier version of this paper. We thank Tim Bollerslev, Marina Balboa and Antonio Rubia for providing the realised variance data sets analysed in Section 6.1 and Luca De Angelis for providing the data set analysed in Section 6.2. Cavaliere thanks the University of Bologna, Alma Idea grants, for financial support. Nielsen thanks the Canada Research Chairs

program, the Social Sciences and Humanities Research Council of Canada (SSHRC), and the Center for Research in Econometric Analysis of Time Series (CREATES, funded by the Danish National Research Foundation, DNRF78) for financial support. Replication files for Tables 1–5 are available from the authors.

References

- Andersen, T.G., T. Bollerslev and F.X. Diebold (2007), Roughing it up: including jump components in the measurement, modeling, and forecasting of return volatility, *Review of Economics and Statistics* 89, 701–720.
- Baillie, R.T., C.-F. Chung and M.A. Tieslau (1996), Analyzing inflation by the fractionally integrated ARFIMA-GARCH model, *Journal of Applied Econometrics* 11, 23–40.
- Barndorff-Nielsen, O.E. and N. Shephard (2002), Econometric analysis of realised volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society Series B* 64, 253–280.
- Beare, B. (2018), Unit root testing with unstable volatility, *Journal of Time Series Analysis* 39, 816–835.
- Beran, J., Y. Feng, S. Ghosh and R. Kulik (2013), *Long-Memory Processes: Probabilistic Properties and Statistical Methods*, Springer, Berlin.
- Beran, J. and A. Weiershäuser (2011), On spline regression under Gaussian subordination with long memory, *Journal of Multivariate Analysis* 102, 315–335.
- Bloomfield, P. (1973), An exponential model for the spectrum of a scalar time series, *Biometrika* 60, 217–226.
- Boswijk, H.P. and Y. Zu (2018), Adaptive wild bootstrap tests for a unit root with non-stationary volatility, *Econometrics Journal* 21, 87–113.
- Carroll, R.J. (1982), Adapting for heteroscedasticity in linear models, *Annals of Statistics* 10, 1224–1233.
- Cavaliere, G. (2004), Unit root tests under time-varying variances, *Econometric Reviews* 23, 259–292.
- Cavaliere, G., M.Ø. Nielsen and A.M.R. Taylor (2015), Bootstrap score tests for fractional integration in heteroskedastic ARFIMA models, with an application to price dynamics in commodity spot and futures markets, *Journal of Econometrics* 187, 557–579.

- Cavaliere, G., M.Ø. Nielsen and A.M.R. Taylor (2017), Quasi-maximum likelihood estimation and bootstrap inference in fractional time series models with heteroskedasticity of unknown form, *Journal of Econometrics* 198, 165–188.
- Cavaliere, G., M.Ø. Nielsen and A.M.R. Taylor (2019), Supplementary Appendix: Adaptive inference in heteroskedastic fractional time series models, QED working paper 1390, Queen’s University.
- Chang, S.K. and J.Y. Park (2010), Cointegrating regressions with time heterogeneity, *Econometric Reviews* 29, 397–438.
- Corsi, F., S. Mittnik, C. Pigorsch and U. Pigorsch (2008), The volatility of realized volatility, *Econometric Reviews* 27, 46–78.
- Eicker, F. (1967), Limit theorems for regression with unequal and dependent errors, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 1, 59–82, University of California Press, Berkeley.
- Fan, J. and Q. Yao (1998), Efficient estimation of conditional variance functions in stochastic regression, *Biometrika* 85, 645–660.
- Fox, R. and M.S. Taqqu (1986), Large-sample properties of parameter estimates for strongly dependent stationary Gaussian time series, *Annals of Statistics* 14, 517–532.
- Geweke, J. and S. Porter-Hudak (1983), The estimation and application of long memory time series models. *Journal of Time Series Analysis* 4, 221–238.
- Giraitis, L., H.L. Koul and D. Surgailis (2012), *Large Sample Inference for Long Memory Processes*, Imperial College Press, London.
- Gonçalves, S. and L. Kilian (2004), Bootstrapping autoregressions with conditional heteroskedasticity of unknown form, *Journal of Econometrics* 123, 89–120.
- Gonçalves, S. and L. Kilian (2007), Asymptotic and bootstrap inference for $AR(\infty)$ processes with conditional heteroskedasticity, *Econometric Reviews* 26, 609–641.
- Guo, H. and H.L. Koul (2008), Asymptotic inference for some regression models under heteroscedasticity and long memory design and errors, *Annals of Statistics* 36, 458–487.
- Hall, P. and N. Tajvidi (2000), Nonparametric analysis of temporal trend when fitting parametric models to extreme-value data, *Statistical Science* 15, 153–167.
- Hannan, E.J. (1973), The asymptotic theory of linear time-series models, *Journal of Applied Probability* 10, 130–145.

- Hansen, B.E. (1995), Regression with nonstationary volatility, *Econometrica* 63, 1113–1132.
- Harris, D. and H. Kew (2017), Adaptive long memory testing under heteroskedasticity, *Econometric Theory* 33, 755–778.
- Harvey, A.C. and P.M. Robinson (1988), Efficient estimation of nonstationary time series regression, *Journal of Time Series Analysis* 9, 201–204.
- Hassler, U., P.M.M. Rodrigues and A. Rubia (2009), Testing for general fractional integration in the time domain, *Econometric Theory* 25, 1793–1828.
- Hualde, J. and P.M. Robinson (2011), Gaussian pseudo-maximum likelihood estimation of fractional time series models, *Annals of Statistics* 39, 3152–3181.
- Huber, P.J. (1967), The behavior of the maximum likelihood estimator under nonstandard conditions, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 1, 221–233, University of California Press, Berkeley.
- Johansen, S. and M.Ø. Nielsen (2010), Likelihood inference for a nonstationary fractional autoregressive model, *Journal of Econometrics* 158, 51–66.
- Li, W.K., S. Ling and M. McAleer (2002), Recent theoretical results for time series models with GARCH errors, *Journal of Economic Surveys* 16, 245–249.
- Ling, S. (2003), Adaptive estimators and tests of stationary and nonstationary short- and long-memory ARFIMA-GARCH models, *Journal of the American Statistical Association* 98, 955–967.
- Ling, S. and W.K. Li (1997), On fractionally integrated autoregressive moving-average time series models with conditional heteroscedasticity, *Journal of the American Statistical Association* 92, 1184–1194.
- Liu, R.Y. (1988), Bootstrap procedures under some non-I.I.D. models, *Annals of Statistics* 16, 1696–1708.
- Mammen, E. (1993), Bootstrap and wild bootstrap for high dimensional linear models, *Annals of Statistics* 21, 255–285.
- Martins, S. and C. Amado (2016), Modelling sovereign debt contagion: a smooth transition approach, University of Minho working paper.
- McConnell, M.M. and G. Perez-Quiros (2000), Output fluctuations in the United States: what has changed since the early 1980s?, *American Economic Review* 90, 1464–1476.
- Nielsen, M.Ø. (2015), Asymptotics for the conditional-sum-of-squares estimator in multivariate fractional time-series models, *Journal of Time Series Analysis* 36, 154–188.

- Patilea, V. and H. Raïssi (2012), Adaptive estimation of vector autoregressive models with time varying variance: application to testing linear causality in mean, *Journal of Statistical Planning and Inference* 142, 2891–2912.
- Patilea, V. and H. Raïssi (2013), Corrected portmanteau tests for VAR models with time-varying variance, *Journal of Multivariate Analysis* 116, 190–207.
- Patilea, V. and H. Raïssi (2014), Testing second-order dynamics for autoregressive processes in presence of time-varying variance, *Journal of the American Statistical Association* 109, 1099–1111.
- Robinson, P.M. (1987), Asymptotically efficient estimation in the presence of heteroskedasticity of unknown form, *Econometrica* 55, 875–891.
- Robinson, P.M. (1991), Testing for strong serial correlation and dynamic conditional heteroskedasticity in multiple regression, *Journal of Econometrics* 47, 67–84.
- Robinson, P.M. (1994), Efficient tests of nonstationary hypotheses, *Journal of the American Statistical Association* 89, 1420–1437.
- Robinson, P.M. (1995), Gaussian semiparametric estimation of long range dependence, *Annals of Statistics* 23, 1630–1661.
- Rosenblatt, M. (1959), Stationary processes as shifts of functions of independent random variables, *Journal of Mathematics and Mechanics* 8, 665–681.
- Samorodnitsky, G. (2007), Long range dependence, *Foundations and Trends in Stochastic Systems* 1, 163–257.
- Sensier, M. and D. van Dijk (2004), Testing for volatility changes in U.S. macroeconomic time series, *Review of Economics and Statistics* 86, 833–839.
- Stock, J.H. and M.W. Watson (2012), Disentangling the channels of the 2007-09 recession, *Brookings Papers on Economic Activity* Spring Volume, 81–156.
- Tanaka, K. (1999), The nonstationary fractional unit root, *Econometric Theory* 15, 549–582.
- White, H. (1982), Maximum likelihood estimation of misspecified models, *Econometrica* 50, 1–25.
- Wong, W.H. (1983), On the consistency of cross validation in kernel nonparametric regression, *Annals of Statistics* 11, 1136–1141.
- Wu, C.F.J. (1986), Jackknife bootstrap and other resampling methods in regression analysis, *Annals of Statistics* 14, 1261–1295.
- Xu, K.-L. and P.C.B. Phillips (2008), Adaptive estimation of autoregressive models with time-varying variances, *Journal of Econometrics* 142, 265–280.

- Xu, K.-L. and P.C.B. Phillips (2011), Tilted nonparametric estimation of volatility functions with empirical applications, *Journal of Business and Economic Statistics* 29, 518–528.
- Xu, K.-L. and J.-C. Yang (2015), Towards uniformly efficient trend estimation under weak/strong correlation and non-stationary volatility, *Scandinavian Journal of Statistics* 42, 63–86.
- Yu, K. and M.C. Jones (2004), Likelihood-based local linear estimation of the conditional variance function, *Journal of the American Statistical Association* 99, 139–144.
- Zhang, T. and W.B. Wu (2012), Inference of time-varying regression models, *Annals of Statistics* 40, 1376–1402.
- Zygmund, A. (2003), *Trigonometric Series*, third edition, Cambridge University Press, Cambridge.