

CHOOSING BETWEEN PERSISTENT AND STATIONARY VOLATILITY

BY ILIAS CHRONOPOULOS^{1,a}, LIUDAS GIRAITIS^{2,b} AND GEORGE KAPETANIOS^{3,c}

¹Essex Business School, University of Essex, ^ailias.chronopoulos@essex.ac.uk

²School of Economics and Finance, Queen Mary University of London, ^bl.giraitis@qmul.ac.uk

³King's Business School, King's College London, ^cgeorge.kapetanios@kcl.ac.uk

This paper suggests a multiplicative volatility model where volatility is decomposed into a stationary and a nonstationary persistent part. We provide a testing procedure to determine which type of volatility is prevalent in the data. The persistent part of volatility is associated with a nonstationary persistent process satisfying some smoothness and moment conditions. The stationary part is related to stationary conditional heteroskedasticity. We outline theory and conditions that allow the extraction of the persistent part from the data and enable standard conditional heteroskedasticity tests to detect stationary volatility after persistent volatility is taken into account. Monte Carlo results support the testing strategy in small samples. The empirical application of the theory supports the persistent volatility paradigm, suggesting that stationary conditional heteroskedasticity is considerably less pronounced than previously thought.

1. Introduction. Two important issues widely discussed in the statistical and finance literature, over the last 25 years, are structural change and volatility modelling. Starting with the seminal work of [16], volatility modelling has developed into a large topic of study. Most work has produced volatility models that are stationary and allow for time variation in the conditional variance. There are two important groups of parametric models used to model volatility. The first group represents the conditional variance as a function of observables and includes autoregressive conditional heteroskedasticity (ARCH) and generalised ARCH (GARCH) models. The second group, where the conditional variance is treated as a latent variable and may depend on more than one innovation process, includes stochastic volatility models.

Empirical work, though, has repeatedly concluded that volatility can exhibit extreme persistence. Such persistence is not easily accommodated by stationary volatility models. The challenge is revealed through the *integrated GARCH effect*; see, for example, [30], when parameter estimates are observed to lie near the boundary of stationarity. This effect can be caused by smooth or abrupt structural change in the unconditional variance over time. So, it is possible that once allowed for, volatility can be best characterised by persistent, and possibly nonstationary processes. There is a growing literature that tries to characterise volatility using parameter processes that allow for gradual change in the unconditional variance. First, we succinctly summarise the main ways this problem is addressed in the literature, and then present our main contributions.

The first line of recent research on structural change has focused on paradigms coming from the statistical literature, such as the work of [32] and [9], where parameter processes are smooth deterministic functions of time. Dahlhaus and Rao [11] proposed the locally stationary time-varying ARCH model, that is globally nonstationary. Along the same lines, [36] proposed another model with deterministic smoothly varying parameters, where volatility is

Received November 2021; revised September 2022.

MSC2020 subject classifications. Primary 62M10, 62G08; secondary 62G05, 62F05.

Key words and phrases. ARCH effect, persistence, volatility, time-varying coefficient models, nonparametric estimation.

multiplicatively decomposed into a stationary and nonstationary part. The assumption that a nonstationary part could drive volatility has recently permeated in standard ARCH (stationary) models as well. Specifically, [17] and [6] propose using splines, [28] use the Fourier Flexible form of [19] in their periodic volatility model, [26] consider the generalised GARCH model and, in a series of papers [1] and [2], suggest to model volatility as a linear combination of logistic functions.

While the above characterisations provide an avenue to describe and estimate stationary and nonstationary volatility processes, there is no clear way to separate the two kinds of volatility. Further, the above characterisations are either tied to parametric forms or assumed to be smooth deterministic functions of time.

This paper makes a number of contributions. We suggest a multiplicative specification of a volatility process that allows a stationary and a nonstationary part. The latter can potentially account for the extreme persistence of volatility observed in data. We use ideas from the recent literature on structural change to show how persistence, perhaps surprisingly, allows kernel estimation of the unobserved stochastic persistent part of volatility without strong parametric assumptions and the requirement to be a smooth, deterministic function of time; see, for example, [9, 27, 32] and [36], among others. While smooth deterministic functions for the persistent part of volatility are still allowed, including stochastic elements in persistent volatility modelling can provide a richer representation of volatility.

Recent work by [21] shows that as long as a parameter process satisfies some smoothness and moment or boundedness conditions, it can be stochastic but still estimable using a kernel estimator. Such processes may adequately fit the observed behaviour of volatility, as they are clearly more persistent than stationary processes. In fact, persistence is their most distinctive characteristic. Giraitis et al. [22] essentially ask the following question: Assuming a decomposition of the form $y_t = h_t u_t$, for some observed process y_t and unobserved stationary process u_t , what properties should h_t have, so that h_t^2 can be consistently estimated by, essentially, a rolling window form, mean estimate of y_t^2 ? They show that h_t has to change slowly, in the sense that $|h_t - h_s|$ has to be small when t and s are close, and thus, stationary processes do not qualify. A normalised random walk provides a canonical example for the sort of processes we have in mind.

We demonstrate in this paper how the uniform consistent estimation of h_t leads to a strategy of separation between the stationary and persistent parts of volatility. Basically, if the persistent part can be uniformly estimated, then the rescaled series of residuals $|\widehat{u_t}|^\gamma = (h_t^\gamma)^{-1} |y_t|^\gamma$, $\gamma > 0$, can be used to test for ARCH effects (conditional heteroskedasticity) or the presence of a stationary volatility in u_t . If only persistent volatility is present, standard ARCH tests will not detect ARCH effects in residuals. If the persistent part h_t is absent, the normalisation by $\widehat{h_t}^\gamma$ will not distort the residuals and stationary volatility in u_t will be detected. Our specification allows for both the persistent and stationary parts of volatility to coexist. Moreover, they can be extracted from the data. After testing for ARCH effects in u_t , is performed, in a second step, a stationary volatility model can be fitted to u_t . This extension is beyond the scope of the current paper.

In this paper, we discuss, in detail, conditions needed for consistent estimation of the persistent part, h_t , of volatility and further conditions that enable the use of standard ARCH tests to separate persistent volatility from stationary volatility. We provide illustrative Monte Carlo results that support our approach on testing in small samples. We proceed and present extensive empirical evidence clearly supporting the persistent volatility paradigm, suggesting that stationary time varying conditional volatility is less pronounced than previously thought and, further, conditional second moments of asset returns are very persistent and change slowly.

The remainder of this paper is organised as follows. Section 2 presents our statistical procedure and theoretical results. Section 3 contains the simulation study. Section 4 reports the

empirical results from implementing our testing strategy in financial data. Section 5 concludes. Proofs are relegated to the Supplementary Material [7]. Below \rightarrow_D , \rightarrow_P stand for convergence in distribution and probability.

2. Theoretical considerations. We consider the following white noise model for a series of uncorrelated random variables:

$$(1) \quad y_t = h_t u_t, \quad t = 1, \dots, T,$$

where $\{u_t\}$ is a stationary sequence of uncorrelated random variables with $E u_t = 0$, $E u_t^2 = 1$ and $h_t > 0$ is the persistent part of volatility (stochastic or deterministic). Formally, $y_t = y_{t_n}$, $t = 1, \dots, T$ and $h_t = h_{t_n}$, $t = 1, \dots, T$ are triangular arrays but it is unnecessary to add the additional index in what follows. We assume that sequences $\{u_t\}$ and $\{h_t\}$ are mutually independent. Then

$$\text{cov}(y_t, y_s) = E[h_t h_s] E[u_t u_s] = 0 \quad \text{for } t \neq s.$$

First, we establish notions of persistent and stationary volatility. We will assume that h_t is measurable with respect to the information set \mathcal{F}_{t-1} at time $t - 1$ and $E(u_t | \mathcal{F}_{t-1}) = 0$. Then the conditional variance of y_t is defined by

$$(2) \quad \text{var}(y_t | \mathcal{F}_{t-1}) = h_t^2 E(u_t^2 | \mathcal{F}_{t-1}) = h_t^2 \sigma_t^2.$$

To specify the properties of persistence for h_t , we introduce below Assumption H. Overall, the notion of persistence of h_t simplifies to

$$(3) \quad \lim_{T \rightarrow \infty} (h_{t,T}^2 - h_{t-1,T}^2) =_P 0,$$

for any $1 \leq t = t_T \leq T$. For example, if h_t is a deterministic function, the property (3) will imply that the unconditional variance $\text{var}(y_t) = h_t^2 E u_t^2$ changes smoothly when T increases. Other processes, such as locally stationary and stochastic unit root processes, h_t , satisfy this property as well, as we discuss below. We refer to h_t^2 as the persistent part of volatility.

There is a vast body of literature on modelling stationary volatility. We define the stationary part of volatility as a conditional variance

$$\sigma_t^2 = \text{var}(u_t | \mathcal{F}_{t-1}),$$

with respect to the information set \mathcal{F}_{t-1} at time $t - 1$. Here, both $\{\sigma_t^2\}$ and $\{\sigma_t^2 - \sigma_{t-1}^2\}$ are stationary processes, and thus, the persistence property $\sigma_t^2 - \sigma_{t-1}^2 = o_P(1)$ does not hold. The main two classes of stationary volatility models for σ_t^2 are the autoregressive conditional heteroskedastic ARCH and GARCH models and stochastic volatility models.

Our objective is to test whether the conditional variance $\text{var}(y_t | \mathcal{F}_{t-1})$ contains a stationary component σ_t^2 . A simple general specification of such a hypothesis is given by

$$(4) \quad H_0 : \text{var}(y_t | \mathcal{F}_{t-1}) = h_t^2 \quad \text{vs} \quad H_1 : \text{var}(y_t | \mathcal{F}_{t-1}) = h_t^2 \sigma_t^2,$$

where $\{\sigma_t^2\}$ is a stationary sequence of dependent random variables. To construct a feasible testing procedure, we further assume that, under H_0 , $\{u_t\}$ is a sequence of independent identically distributed (i.i.d.) random variables and y_t is generated according to the following processes:

$$(5) \quad \begin{aligned} y_t &= h_t \varepsilon_t \quad \text{under } H_0, \\ y_t &= h_t u_t, \quad u_t = \sigma_t \varepsilon_t \quad \text{under } H_1, \end{aligned}$$

where $\{\varepsilon_t\}$ is an i.i.d. sequence with $E \varepsilon_t = 0$, $\text{var}(\varepsilon_t) = 1$ and such that $E(\varepsilon_t | \mathcal{F}_{t-1}) = 0$, $E(\varepsilon_t^2 | \mathcal{F}_{t-1}) = 1$. This model specification implies (4).

According to the specification (5), $\{u_t\}$ is a white noise process. Essentially, a standard test for ARCH effect is a white noise test for the squared time series $\{u_t^2\}$ and requires the existence of six finite moments for u_t . To relax this assumption, in this paper we consider tests for ARCH effects, based on $\{|u_t|^\gamma\}$, $\gamma > 0$, which is basically a white noise test for the time series $\{|u_t|^\gamma\}$. The analysis covers the case $\gamma = 2$.

Since $\{u_t\}$ is not observed, we base our test for ARCH effects on the power transform of residuals

$$(6) \quad \widehat{|u_t|}^\gamma \quad \text{where} \quad \widehat{|u_t|}^\gamma = (\widehat{h_t}^\gamma)^{-1} |y_t|^\gamma$$

and $\widehat{h_t}^\gamma > 0$ is an estimate of h_t^γ . Such testing requires uniformly consistent estimation of h_t^γ by $\widehat{h_t}^\gamma$, and thus, stronger conditions on $\{h_t, u_t\}$ are needed, than for consistent point estimation of h_t at time t . This is reflected in the Assumptions M and H that we make for u_t and h_t . In particular, we impose the assumption of mutual independence between $\{h_t^2\}$ and $\{u_t^2\}$, which clearly holds for a deterministic volatility factor h_t .

ASSUMPTION M (α -mixing).

1. $\{u_t\}$ is a stationary white-noise ergodic sequence with $Eu_t = 0$, $Eu_t^2 = 1$, $Eu_t u_s = 0$ for $t \neq s$.
2. $\{u_t\}$ is α -mixing with mixing coefficients $\alpha_k \leq c\phi^k$, $k \geq 1$, for some $0 < \phi < 1$ and $c > 0$.

ASSUMPTION H (Smoothness).

1. For some $\nu \in (1/2, 1]$,

$$(7) \quad \begin{aligned} |h_t - h_j| &\leq C(|t - j|/T)^\nu, \quad t, j = 1, \dots, T \quad \text{or} \\ |h_t - h_j| &\leq (|t - j|/T)^\nu \xi_{tj}, \end{aligned}$$

where $C > 0$ does not depend on t, T , and for some $0 < \alpha < \infty$, $c > 0$,

$$(8) \quad \max_{t, j=1, \dots, T} E[\exp(c|\xi_{tj}|^\alpha)] \leq C < \infty, \quad \max_{t=1, \dots, T} E[\exp(c|h_t|^\alpha)] \leq C < \infty.$$

2. There exists $a > 0$ such that $h_t \geq a > 0$ a.s. for all $t \geq 1$.
3. $\{h_t\}$ and $\{u_t\}$ are mutually independent.

The model specification (1) abstracts from the general case of a time series with a specified conditional mean. It is possible to generalise our test for ARCH effects to a time series with a nonzero conditional mean

$$y_t = \mu_t + h_t u_t, \quad \mu_t = E(y_t | \mathcal{F}_{t-1}), \quad t = 1, \dots, T.$$

The smoothness condition (7) in Assumption H implies that the persistent component of volatility, h_t , drifts slowly in time, which essentially rules out abrupt or explosive behaviour for h_t . This assumption is widely used in the statistical and econometric literature. It allows the use of both deterministic and stochastic time-varying processes h_t and implies the persistence property (3).

The deterministic specification $h_t = g(t/T)$, $t = 1, \dots, T$, where $g(\cdot)$ is a Lipschitz smooth function with parameter $1/2 < \nu \leq 1$, that is, $|g(x) - g(y)| \leq C|x - y|^\nu$ is a standard assumption in the work of Dahlhaus on locally stationary processes (see, e.g., [9] or [10]). It implies $|h_t - h_s| \leq C(|t - s|/T)^\nu$.

The stochastic specification $h_t = T^{-\nu} |\sum_{j=1}^t v_j|$, $t = 1, \dots, T$ of h_t , where $\{v_j\}$ is a stationary sequence with zero mean, was proposed by [20–22] to allow for a persistent process h_t that can be presented as nonstationary random walks; see Example 1 below.

A combination of the two, satisfying (7) with parameter ν , can be summarised as

$$(9) \quad h_t = |T^{-\nu}(v_1 + \cdots + v_t) + g(t/T)| + a, \quad t = 1, \dots, T \quad (a > 0).$$

Our testing procedures will still work for the case of y_t with a nonzero conditional mean; in this case, a first step estimator for the mean will be required; see, for example, [8].

EXAMPLE 1. Let $\{v_j\}$ be a stationary Gaussian ARFIMA(p, d, q) sequence with parameter $d \in (0, 1/2)$ and zero mean; see, for example, Chapter 7 in [23]. Then $h_t = T^{-\nu}|\sum_{j=1}^t v_j|$, $t = 1, \dots, T$ satisfies (7) of Assumption H with $\nu = 1/2 + d$ and $\alpha = 2$. Indeed, for $t > s$,

$$\begin{aligned} |h_t - h_s| &= \left| T^{-\nu} \left| \sum_{j=1}^t v_j \right| - T^{-\nu} \left| \sum_{j=1}^s v_j \right| \right| \\ &\leq T^{-\nu} \left| \sum_{j=s+1}^t v_j \right| \leq (|t-s|/T)^\nu |\xi_{ts}|, \quad \xi_{ts} = (t-s)^{-\nu} \sum_{j=s+1}^t v_j. \end{aligned}$$

Here, ξ_{ts} is a Gaussian variable, and by Proposition 3.3.1 in [23] and stationarity of $\{v_j\}$, the variance $\text{var}(\xi_{ts}) = \text{var}(\xi_{t-s,0}) \rightarrow \nu^2 < \infty$ as $t-j \rightarrow \infty$. Hence, ξ_{ts} and h_t satisfy (8) with $\alpha = 2$.

2.1. *Volatility estimation.* To extract residuals $|\widehat{u}_t|^\gamma = (\widehat{h}_t^\gamma)^{-1}|y_t|^\gamma$, required for the testing of ARCH effects, we need an estimate for h_t^γ in $|y_t|^\gamma = h_t^\gamma |u_t|^\gamma$. Without loss of generality, we assume that h_t is rescaled so that $E|u_t|^\gamma = 1$. We will show that in model (1), under Assumptions H and M, h_t^γ can be consistently estimated by a kernel-type estimate

$$(10) \quad \widehat{h}_t^\gamma = K_t^{-1} \sum_{j=1}^T b_{H,|t-j|} |y_j|^\gamma, \quad K_t = \sum_{j=1}^T b_{H,|t-j|}, \quad t = 1, \dots, T,$$

where $b_{H,|t-j|} = K(|t-j|/H)$ are kernel weights. $K(\cdot)$ is assumed to be a nonnegative and bounded function, with a piecewise bounded derivative, and H is a bandwidth parameter that satisfies $H = o(T)$, as $T \rightarrow \infty$. Commonly used examples of $K(x)$ include:

$$K(x) = (1/2)I(|x| \leq 1), \quad \text{flat kernel,}$$

$$K(x) = (3/4)(1-x^2)I(|x| \leq 1), \quad \text{Epanechnikov kernel,}$$

$$K(x) = (1/\sqrt{2\pi})e^{-x^2/2}, \quad \text{Gaussian kernel.}$$

The first two kernel functions have finite support, whereas the Gaussian kernel has infinite support. We further assume that on its support,

$$(11) \quad K(x) \leq C(1+x^g)^{-1}, \quad |(d/dx)K(x)| \leq C(1+x^g)^{-1}, \quad x \geq 0 \text{ with } g > 4, C > 0.$$

Under this setup, in Lemma 7.2 of the Supplementary Material [7] we show the pointwise consistency of this estimate:

$$(12) \quad |\widehat{h}_t^\gamma - h_t^\gamma| = O_p((H/T)^\nu + H^{-1/2}).$$

Similar results for vector autoregressive models were derived in [22]. In Lemma 7.2 of the Supplementary Material [7], using the results of [14], we establish the uniform convergence

$$(13) \quad \max_{t=1, \dots, T} |\widehat{h}_t^\gamma - h_t^\gamma| = o_P(1).$$

This uniform convergence result will prove useful in our testing procedure for the distinction between the persistent and stationary parts of volatility that follows.

2.2. Testing. In this subsection, we outline how our strategy to discriminate between the persistent and stationary components of volatility works under (1), where u_t is not observed. Basically, it will be seen that the test of hypothesis (4) reduces to a white noise test for $\{|u_t|^\gamma\}$. First, we briefly summarise the basic tests for ARCH effects when a white noise time series $\{u_t\}$ is observed.

The Lagrange Multiplier (LM) test by [16] is the most commonly used standard ARCH test to detect autoregressive conditional heteroskedasticity (or stationary volatility) in $\{u_t\}$. We simply fit to u_t^2 an $AR(p)$, $p \geq 1$, model

$$(14) \quad u_t^2 = \beta_0 + \beta_1 u_{t-1}^2 + \cdots + \beta_p u_{t-p}^2 + \eta_t,$$

where $\beta_0 > 0$ and test the following null hypothesis:

$$(15) \quad H_0: \beta_1 = \beta_2 = \cdots = \beta_p = 0 \quad \text{vs} \quad H_1: \beta_j \neq 0 \quad \text{for some } j = 1, \dots, p.$$

The null hypothesis H_0 implies absence of correlation in the first p lags of the series $\{u_t^2\}$ and *vice versa*. Basically, this ARCH LM test is equivalent to testing for absence of autocorrelation in $\{u_t^2\}$.

The test statistic of the ARCH LM test by [16] is defined as TR^2 , where T is the sample size and R^2 is the coefficient of determination of the AR regression (14). Under H_0 , when $\{u_t\}$ is an i.i.d. sequence with finite fourth moment, the LM statistic follows asymptotically a χ_p^2 distribution. Further tests, such as the Wald and likelihood ratio, have been shown to be asymptotically equivalent to the LM test.

Through testing, the literature mainly addresses two distinct problems: the misspecification of the conditional mean (see, e.g., the discussion in [3]), and the correct specification of the volatility process. Our work naturally falls in the second category by addressing the question of whether allowing for a persistent component h_t in (1), provides a better specification for the volatility process. By ARCH effects in $\{u_t\}$, we mean the presence of correlation in a sequence $\{|u_t|^\gamma\}$, where $\gamma > 0$ is selected in advance. We will test for ARCH effects in the unobserved component u_t of $y_t = h_t u_t$ in (1) by fitting to $|u_t|^\gamma$ an $AR(p)$, $p \geq 1$ model

$$(16) \quad |u_t|^\gamma = \beta_0 + \beta_1 |u_{t-1}|^\gamma + \cdots + \beta_p |u_{t-p}|^\gamma + \eta_t$$

and then testing the hypothesis (15) on β_1, \dots, β_p . We replace the unobserved variables $|u_t|^\gamma$ by residuals

$$(17) \quad \widehat{|u_t|^\gamma} = (\widehat{h_t^\gamma})^{-1} |y_t|^\gamma,$$

where $\widehat{h_t^\gamma}$ is the kernel estimate (10) of the γ -power h_t^γ of the persistent factor h_t of volatility. Our aim is to show that asymptotically it is equivalent to test for ARCH effects using the residuals $\widehat{\mathbf{u}}^\gamma = [\widehat{|u_1|^\gamma}, \widehat{|u_2|^\gamma}, \dots, \widehat{|u_T|^\gamma}]'$, instead of $\mathbf{u}^\gamma = [|u_1|^\gamma, |u_2|^\gamma, \dots, |u_T|^\gamma]'$.

In addition, such an equivalence implies that the residuals, $\widehat{u}_t = \widehat{h_t}^{-1} y_t$ obtained using $\widehat{h_t}$ with $\gamma = 1$, should behave as a white noise. In Theorems 2.2 and 2.3, we show that both the ARCH LM test based on regression on u_t and the correlogram of $\widehat{u}_1, \dots, \widehat{u}_T$ can be used to test for absence of correlation in $\{u_t\}$.

The ARCH test using regression (14) for powers, $|u_t|^\gamma$ obtained in Theorem 2.1, shows that if a stationary process σ_t^2 (co-)drives the volatility via u_t , then the normalisation by $\widehat{h_t^\gamma}$ in (17) will not corrupt the properties of testing.

In our setup and for $p \geq 1$, we consider $TS(\mathbf{u}^\gamma) = TR^2$, the test statistic where R^2 is the coefficient of determination of the AR regression (16) based on \mathbf{u}^γ , and $TS(\widehat{\mathbf{u}}^\gamma) = T\widehat{R}^2$ where \widehat{R}^2 is based on the residuals, $\widehat{\mathbf{u}}^\gamma$, as described above. The formulas of $S(\mathbf{u}^\gamma)$ and $S(\widehat{\mathbf{u}}^\gamma)$ are given in (6.6) and (6.7) of the Supplementary Material [7].

The following theorem, proven in the Supplementary Material [7], gives a sufficient condition for LM test for ARCH effects to be asymptotically valid when applied to $\widehat{\mathbf{u}}^\gamma$, instead of \mathbf{u}^γ .

Denote by $\gamma_k = \text{cov}(|u_k|^\gamma, |u_0|^\gamma)$, $k \geq 0$ the covariance function of a stationary sequence $\{|u_t|^\gamma\}$, $\gamma > 0$, and define the $p \times p$ matrix $\mathbf{\Gamma}_p$ and $p \times 1$ vector $\boldsymbol{\gamma}_p$ by setting

$$\mathbf{\Gamma}_p = (\gamma_{|j-k|})_{j,k=1,\dots,p}, \quad \boldsymbol{\gamma}_p = (\gamma_1, \dots, \gamma_p)'.$$

Denote by $\boldsymbol{\beta}_p$ the $p \times 1$ vector of parameters $(\beta_1, \dots, \beta_p)$, which will appear in testing under the alternative hypothesis H_1 and set

$$\boldsymbol{\beta}_p = (\beta_1, \dots, \beta_p)' = \mathbf{\Gamma}_p^{-1} \boldsymbol{\gamma}_p, \quad \sigma_p^2 = \text{var}(|u_{p+1}|^\gamma - \beta_1 |u_p|^\gamma - \dots - \beta_p |u_1|^\gamma).$$

Notice that the existence of $\mathbf{\Gamma}_p^{-1}$ follows from Lemma 3.1(i) in [13] because the stationary sequence $\{u_t^2\}$ has a bounded spectral density. The latter follows from the absolute summability of the covariance function γ_k ; see (7.8) in the Supplementary Material [7].

Recall notation ν of the smoothness parameter of h_t appearing in (7).

THEOREM 2.1. *Let $\{y_t, t = 1, \dots, T\}$, follow (1), Assumptions H and M hold and H satisfies*

$$(18) \quad T^{1/2+a} \leq H = o(T^{1-(1/4\nu)}) \quad (\text{for some } a > 0).$$

Assume that $\gamma > 0$ and $E|u_t|^{3\gamma} < \infty$. Then the ARCH LM test statistic based on regression (16) on $|u_t|^\gamma$ has the following properties. As $T \rightarrow \infty$, for any $p \geq 1$,

$$(19) \quad S(\widehat{\mathbf{u}}^\gamma) = S(\mathbf{u}^\gamma) + o_P(1) = \sigma_p^{-2} \boldsymbol{\beta}_p' \mathbf{\Gamma}_p \boldsymbol{\beta}_p + o_P(1).$$

In addition, if $\{u_t\}$ is an i.i.d. sequence, then $\boldsymbol{\beta}_p = 0$, and

$$(20) \quad TS(\widehat{\mathbf{u}}^\gamma) = TS(\mathbf{u}^\gamma) + o_P(1) \rightarrow_D \chi_p^2.$$

Result (19) implies that testing for ARCH effects based on regression (16) is equivalent to testing for the white noise for the series $\{|u_t|^\gamma\}$. Indeed, the matrix $\mathbf{\Gamma}_p$ is positive definite and, therefore, its smallest and largest eigenvalues obey $0 < \lambda_{\min} \leq \lambda_{\max} < \infty$. Notice that $\boldsymbol{\beta}_p' \mathbf{\Gamma}_p \boldsymbol{\beta}_p = \boldsymbol{\gamma}_p' \mathbf{\Gamma}_p^{-1} \boldsymbol{\gamma}_p$, and

$$\boldsymbol{\beta}_p' \mathbf{\Gamma}_p \boldsymbol{\beta}_p \geq \|\boldsymbol{\beta}_p\|^2 \lambda_{\min}, \quad \boldsymbol{\gamma}_p' \mathbf{\Gamma}_p^{-1} \boldsymbol{\gamma}_p \geq \|\boldsymbol{\gamma}_p\|^2 \lambda_{\max}^{-1},$$

where $\|\boldsymbol{\beta}_p\|$ denotes the Euclidean norm of $\boldsymbol{\beta}_p$. Hence, $\|\boldsymbol{\beta}_p\| = 0$ implies $\|\boldsymbol{\gamma}_p\| = 0$ and vice versa, which proves the above claim.

Theorem 2.1 implies that a test for ARCH effects in $\{u_t\}$ based on statistic $TS(\widehat{\mathbf{u}}^\gamma)$ has the same asymptotic size and power properties as a test based on $TS(\mathbf{u}^\gamma)$ applied on unobserved \mathbf{u}^γ . If the hypothesis H_0 is not rejected, then this implies the absence of correlation in $\{|u_t|^\gamma\}$ up to lag p . Conversely, the alternative H_1 is detected with a rate T .

Notice that the value $\boldsymbol{\beta}_p = \mathbf{\Gamma}_p^{-1} \boldsymbol{\gamma}_p$ appearing in (19) is the same as the “true” value of the parameter $\boldsymbol{\beta}_p$ estimated by the OLS method in regression (14) when η_t is a white noise sequence.

REMARK 1. If Assumption H is satisfied with $\nu = 1$, for example, which is the case for deterministic weights $h_t = g(t/T)$, where g is a continuous piecewise differentiable function with a bounded derivative, then assumption (18) on the bandwidth H becomes

$$(21) \quad T^{1/2+a} \leq H = o(T^{3/4}).$$

Note that lower values of ν imply narrower permitted intervals for bandwidth H . Given that ν is unknown, a theoretical recommendation for the choice of H is that it should be larger than $T^{1/2}$ and less than $T^{3/4}$. In practice, we recommend trying different values for H , within the interval $T^{1/2}$ to $T^{3/4}$, to establish test robustness. This matter is further discussed in the Monte Carlo study.

REMARK 2. The moment condition $E|u_t|^{3\gamma} < \infty$ of Theorem 2.1 implies that testing for ARCH effects based on regression (14) using u_t^2 , $|u_t|$ or $|u_t|^{1/2}$, which corresponds to $\gamma = 2, 1, 1/2$, requires the existence of $E|u_t|^6$, $E|u_t|^3$ and $E|u_t|^{3/2}$, respectively. If $E|u_t|^{3\gamma+\epsilon} < \infty$ for some $\epsilon > 0$, then (18) can be replaced by

$$T^{1/2} = o(H), \quad H = o(T^{1-(1/4\nu)}).$$

We also suggest two methods to examine whether the unobserved time series $\{u_t\}$ in the model (1) is a white noise sequence. The first approach is similar to testing for ARCH effects. It amounts to fitting to u_t an $AR(p)$, $p \geq 1$ model

$$(22) \quad u_t = \beta_0 + \beta_1 u_{t-1} + \cdots + \beta_p u_{t-p} + \eta_t$$

and testing the hypothesis H_0 and H_1 on $(\beta_1, \dots, \beta_p)$ in (14). The test statistics $TS(\mathbf{u})$ and $TS(\hat{\mathbf{u}})$ in Theorem 2.2 correspond to the AR regression (22) on u_t and $\hat{u}_t = (\hat{h}_t)^{-1} y_t$. They satisfy the asymptotic results of Theorem 2.1 with and β_p , Γ_p defined as in (19) using

$$(23) \quad \gamma_k = \text{cov}(u_k, u_0), \quad \sigma_p^2 = \text{var}(u_{p+1} - \beta_1 u_p - \cdots - \beta_p u_1).$$

THEOREM 2.2. Let $\{y_t, t = 1, \dots, T\}$ be as in (1) and $E|u_t|^3 < \infty$. Suppose that Assumptions M, H and (18) hold. Then statistics $TS(\mathbf{u})$ and $TS(\hat{\mathbf{u}})$ based on the AR regression (22), have the following properties.

For any $p \geq 1$, as $T \rightarrow \infty$,

$$(24) \quad S(\hat{\mathbf{u}}) = S(\mathbf{u}) + o_P(1) = \sigma_p^{-2} \beta_p' \Gamma_p \beta_p + o_P(1).$$

In addition, if $\{u_t\}$ is an i.i.d. sequence, then $\beta_p = 0$, and

$$(25) \quad TS(\hat{\mathbf{u}}) = TS(\mathbf{u}) + o_P(1) \rightarrow_D \chi_p^2.$$

The same argument as used below Theorem 2.1, implies that the hypothesis $H_0 : \beta_1 = \cdots = \beta_p = 0$ is equivalent to the absence of correlation up to lag p in $\{u_t\}$.

Alternatively, absence of correlation in $\{u_t\}$ can be tested using the correlogram of residuals $\hat{u}_t = (\hat{h}_t)^{-1} y_t$. This important step of data analysis allows one to verify the model specification (1) for y_t since standard tests for white noise based on y_t might not be applicable.

For $k = 0, 1, \dots$, denote

$$(26) \quad \begin{aligned} \hat{r}_k &= T^{-1} \sum_{t=k+1}^T (\hat{u}_t - \bar{\hat{u}})(\hat{u}_{t-k} - \bar{\hat{u}}), \\ \tilde{r}_k &= T^{-1} \sum_{t=k+1}^T (u_t - Eu_t)(u_{t-k} - Eu_{t-k}). \end{aligned}$$

THEOREM 2.3. Suppose that assumptions of Theorem 2.2 are satisfied. Then, as $T \rightarrow \infty$,

$$(27) \quad \hat{r}_k = \tilde{r}_k + o_P(1) = \text{cov}(u_k, u_0) + o_P(1), \quad k \geq 0.$$

In addition, if $\{u_t\}$ is an i.i.d. sequence, then

$$(28) \quad T^{1/2} \hat{r}_k = T^{1/2} \tilde{r}_k + o_P(1) \rightarrow \mathcal{N}(0, (Eu_1^2)^2), \quad k \geq 1.$$

Denote $\widehat{\rho}_k = \widehat{r}_k / \widehat{r}_0$, $k = 0, 1, 2, \dots$. If $\{u_t\}$ is an i.i.d. sequence then (28) of Theorem 2.3 implies that

$$(29) \quad T^{1/2}(\widehat{\rho}_1, \dots, \widehat{\rho}_m) \rightarrow_D \mathcal{N}(0, I_m), \quad m \geq 1.$$

This shows that using residuals \widehat{u}_t , we can perform standard tests for the absence of correlation in $\{u_t\}$ at individual lag k and Ljung–Box tests, for the cumulative lag m as if variables $\{u_t\}$ were observed.

Notice that in our setting no ARCH effects in $\{u_t\}$ imply no correlation in $\{|u_t|^\gamma\}$. This is a slightly weaker property than the i.i.d. assumption on $\{u_t\}$ under H_0 . The latter property leads to standard approximations (20) and (29) for the test statistics, which are not guaranteed for a non-i.i.d. white noise $\{u_t\}$.

REMARK 3. In financial and economic applications, the common choice for volatility modelling is a stationary GARCH-type model. Given this, it is also relevant to test the hypothesis

$$H_0 : y_t = u_t \quad \text{vs} \quad H_1 : y_t = h_t u_t,$$

for absence of persistent component h_t ($h_t = 1$) in the model (1), where $\{u_t\}$ is a stationary sequence of uncorrelated random variables. In general, this is equivalent to testing the mean stability of series $y_t^2 = h_t^2 E u_t^2 + h_t^2 (u_t^2 - E u_t^2)$. Tests for the detection of alternatives with deterministic h_t were developed in [12].

REMARK 4. Jumps or abrupt changes in volatility are often observed in real data and have been widely discussed in the literature; see, for example, [29] and [18]. The smoothness Assumption H, of our setting excludes intermittently occurring changes with magnitude, that is bounded away from zero in the persistent part, h_t^2 of volatility. Discrimination between persistent and stationary volatility requires extraction of h_t^2 , which is achieved by using kernel estimation and Assumption H. In our multiplicative volatility model, jumps can appear in the stationary volatility component σ_t^2 . The presence of jumps in that component will lead to rejection of the null hypothesis. Simulation evidence supports this finding and is available upon request.

3. Simulation study. In this section, we use simulations to verify the theoretical properties of the test statistics $TS(\widehat{u}^\gamma)$ for ARCH effects in $\{u_t\}$, and explore its finite-sample size and power performance. In particular, we examine the impact of the three types of persistent volatility h_t (constant, deterministic, stochastic) and the choice of the bandwidth parameter H on the size and power of the test, and how crucial the moment condition $E|u_t|^{3\gamma} < \infty$ is.

We generate an array of samples

$$(30) \quad y_t = h_t u_t, \quad u_t = \sigma_t \varepsilon_t, \quad t = 1, \dots, T,$$

where $\{\varepsilon_t\}$ is an i.i.d. $N(0, 1)$ noise. For σ_t^2 , we use stationary ARCH(1) and GARCH(1, 1) models:

$$(31) \quad \begin{aligned} \sigma_t^2 &= 1 + \beta u_{t-1}^2, \quad \beta = 0, 0.2, 0.4, \quad \text{ARCH(1) model;} \\ \sigma_t^2 &= 1 + 0.2 u_{t-1}^2 + 0.7 \sigma_{t-1}^2, \quad \text{GARCH(1, 1) model.} \end{aligned}$$

The case $\beta = 0$ leads to $\sigma_t^2 = 1$ or H_0 , and

$$(32) \quad y_t = h_t \varepsilon_t.$$

We use (32) to study the empirical size of our test for ARCH effects in unobserved u_t at lag p , in particular, to determine whether the size of this test is robust to the choice of H . For a persistent volatility component h_t , with parameter $\nu = 1$ by Theorem 2.1 under (32),

$$TS(\widehat{\mathbf{u}}^\gamma) \sim TS(\mathbf{u}^\gamma) \sim \chi_p^2 \quad \text{if } T^{0.5} = o(H), \quad H = o(T^{0.75}).$$

Ideally, we expect both the test statistic $TS(\widehat{\mathbf{u}}^\gamma)$ based on residuals, and $TS(\mathbf{u}^\gamma)$ based on u_t , to exhibit similar empirical size. Further, for $\beta > 0$ in ARCH(1) model or GARCH(1, 1) model in (30)–(31), a stationary volatility component is present, and then for the above choice of H , we expect them to achieve similar empirical power, as shown in Theorem 2.1.

We set the significance level to $\alpha = 5\%$ and conduct 5000 replications. We consider $T = 200, 400, 800, 1600$, $\gamma = 2, 1, 1/2$ and we conduct testing for ARCH effects on $|u_t|^\gamma$ at lags 1, 5 and 10 for $\gamma = 2, 1, 1/2$.

In the tables below, we examine the impact of bandwidth H and parameter γ on testing. The shaded grey range of $H = [T^{0.60}, \dots, T^{0.70}]$ denotes theoretically permissible values of H for smoothness parameter $\nu = 1$. As a reference point, we include a third shaded line reporting the size and power for statistics $TS(\mathbf{u}^\gamma)$. Theoretically, as the sample size increases, size and power in the shaded area should approach the benchmark.

We use three models for $\{y_t\}$.

MODEL 1. We set $h_t = 1$ and σ_t^2 follows (31).

This model does not involve a persistent time-varying volatility component. It includes a stationary volatility component except for the case $\beta = 0$, where $\sigma_t = 1$ and $y_t = \varepsilon_t$.

The left panel of Table 1 confirms that both tests $TS(\widehat{\mathbf{u}}^2)$ and $TS(\mathbf{u}^2)$ for ARCH effects at the lag $p = 5$ based on squares ($\gamma = 2$) achieve similar empirical size and power when H is chosen from the permitted range $[T^{0.6}, T^{0.7}]$.

MODEL 2. We set $h_t = \sin(2\pi t/T) + 2$, $t = 1, \dots, T$, and σ_t^2 follows (31).

This model contains a persistent deterministic volatility component h_t , which satisfies Assumption H with parameter $\nu = 1$. For $\beta = 0$, the model $y_t = h_t \varepsilon_t$ includes no stationary volatility (ARCH effects) in u_t . From the right panel of Table 1, we conclude that overall the empirical size and power of the test $TS(\widehat{\mathbf{u}}^2)$ are comparable to those of $TS(\mathbf{u}^2)$ as long as $H \in [T^{0.6}, T^{0.7}]$, but the use of a permissible bandwidth H plays an essential role here. In real data, u_t^2 is not observed and the test $TS(\mathbf{u}^2)$ cannot be performed. Not surprisingly, in the model $y_t = h_t \varepsilon_t$ ($\beta = 0$), the empirical size of the test $TS(\mathbf{y}^2)$ applied on squares y_t^2 of the data is close to 100%, that is, it falsely suggests the presence of a stationary volatility in y_t .

MODEL 3. We set $h_t = T^{-(d-1/2)}|I_{d,t}| + 1$, $t = 1, \dots, T$, where $I_{d,t}$ is a nonstationary ARFIMA(0, d , 0) process. We consider the values $d = 1.2, 1.4, 1.5$ and σ_t^2 follows (31).

Here, we assume that $\{h_t\}$ and $\{\varepsilon_t\}$ are mutually independent. Such a stochastic persistent process h_t satisfies Assumption H with parameter $\nu = d - 1/2 > 1/2$; see Example 1. Table 2 shows that overall the performance of the test for ARCH effects on squares of residuals at the lag $p = 5$ exhibits similar patterns as for Model 2, although the lower degree of persistence of h_t in Model 3 results in a somewhat smaller rate of detection of spurious presence of stationary volatility in y_t by $TS(\mathbf{y}^2)$.

In sum, to test for ARCH effects in $\{u_t\}$, we have used the statistics $S(\widehat{\mathbf{u}}^\gamma)$ based on residuals $|u_t|^\gamma$. We report testing results for the value $\gamma = 2$, while the results for $\gamma = 1, 1/2$ can

TABLE 1

Testing for ARCH effects on squares at the lag $p = 5$ in Model 1 and Model 2. Rejection frequencies (in %) at the 5% significance level ($\beta = 0$ size, $\beta > 0$ power)

T	H	Data	Model 1 (with $h_t = 1$)				Model 2 (with deterministic h_t)			
			$\beta = 0$	$\beta = 0.2$	$\beta = 0.4$	GARCH	$\beta = 0$	$\beta = 0.2$	$\beta = 0.4$	GARCH
200	$T^{0.50}$	$\widehat{u_t^2}$	12.98	19.04	48.82	7.22	10.50	17.96	49.14	8.02
	$T^{0.60}$		4.90	23.30	60.82	28.66	3.20	24.42	62.84	34.96
	$T^{0.70}$		3.72	30.14	69.60	52.42	7.64	40.68	74.92	66.44
	$T^{0.80}$		3.60	34.64	74.32	65.52	25.10	60.12	84.08	83.34
		u_t^2	4.34	40.30	78.94	77.12	4.34	40.30	78.96	77.12
		y_t^2	4.34	40.30	78.94	77.12	73.66	86.78	93.62	95.22
400	$T^{0.50}$	$\widehat{u_t^2}$	13.18	42.80	89.26	33.86	12.10	42.62	89.30	35.12
	$T^{0.60}$		4.80	52.00	94.12	76.14	3.58	53.36	94.34	79.00
	$T^{0.70}$		3.62	60.04	96.08	90.56	7.90	68.78	96.86	95.10
	$T^{0.80}$		3.80	64.70	97.14	94.94	40.86	86.58	98.64	98.60
		u_t^2	4.68	69.06	97.60	97.22	4.68	69.04	97.60	97.22
		y_t^2	4.68	69.06	97.60	97.22	95.40	99.14	99.80	99.92
800	$T^{0.50}$	$\widehat{u_t^2}$	15.32	80.58	99.92	89.00	14.42	80.62	99.92	89.58
	$T^{0.60}$		5.60	87.42	100.00	99.38	4.74	88.16	100.00	99.52
	$T^{0.70}$		4.38	91.38	100.00	99.88	7.78	93.88	100.00	99.96
	$T^{0.80}$		4.18	92.96	100.00	99.92	60.80	98.98	100.00	100.00
		u_t^2	4.74	94.14	100.00	99.98	4.74	94.14	100.00	99.98
		y_t^2	4.74	94.14	100.00	99.98	99.98	100.00	100.00	100.00
1600	$T^{0.50}$	$\widehat{u_t^2}$	16.44	99.04	100.00	99.96	15.96	99.08	100.00	99.96
	$T^{0.60}$		6.20	99.72	100.00	100.00	5.68	99.72	100.00	100.00
	$T^{0.70}$		5.08	99.82	100.00	100.00	7.06	99.84	100.00	100.00
	$T^{0.80}$		5.16	99.84	100.00	100.00	80.06	100.00	100.00	100.00
		u_t^2	5.44	99.86	100.00	100.00	5.44	99.86	100.00	100.00
		y_t^2	5.44	99.86	100.00	100.00	100.00	100.00	100.00	100.00

be found in Section 8 of the Supplementary Material [7]. Testing results for ARCH effects at lags $p = 1, 10$ produce similar patterns as for $p = 5$, and are available upon request. We report additional testing results for errors $\varepsilon_t \sim t(4)$, in Section 8 of the Supplementary Material [7]. They show that the lack of finite $E|u_t|^{3\gamma}$ moment has an impact on testing results.

Finally, in Figure 1 we explore the impact of bandwidth H on the size of the test by plotting the MC average of the $TS(\widehat{u}^\gamma)$ test statistic for $p = 5$ and various values of H under the null hypothesis. We consider data $y_t = h_t \varepsilon_t$ produced by Model 2 for $T = 1600$. From theory, the Monte Carlo average of a well-behaved test statistic should approach the number of degrees of freedom $p = E\chi_p^2$ marked by the light grey line. We see that this is indeed the case for bandwidths $H \in [T^{0.6}, T^{0.7}]$, suggesting that such bandwidth values perform well in small samples and meet the requirements of our theoretical analysis.

Since the absence of a stationary volatility part can be detected, it is intriguing to conduct a comparison of volatility forecasts for a specific case of the model (1):

(33)

$$y_t = h_t \varepsilon_t, \quad t = 1, \dots, T,$$

with a persistent nonparametric volatility $\text{var}(y_t|\mathcal{F}_{t-1}) = h_t^2$ and an i.i.d. $(0, 1)$ noise $\{\varepsilon_t\}$. This model does not includes a stationary volatility component σ_t^2 ; see (2). Our primary interest is to verify whether the kernel forecasting method $\widehat{h}_{T+1|T}^2 := \widehat{h}_T^2$ of the volatility h_{T+1}^2

TABLE 2
Testing for ARCH effects on squares at the lag $p = 5$ in Model 3. Rejection frequencies (in %) at the 5% significance level ($\beta = 0$ size, $\beta > 0$ power)

Model 3 (with stochastic h_t)											
T	H	Data	$\beta = 0$			$\beta = 0.2$			$\beta = 0.4$		
			$d = 1.2$	$d = 1.4$	$d = 1.5$	$d = 1.2$	$d = 1.4$	$d = 1.5$	$d = 1.2$	$d = 1.4$	$d = 1.5$
200	$T^{0.50}$	$\widehat{u_t^2}$	11.20	11.80	12.14	19.12	19.48	19.38	49.26	49.00	49.14
	$T^{0.60}$		3.90	4.06	4.30	23.40	23.30	23.36	62.06	61.40	61.52
	$T^{0.70}$		3.92	3.82	3.68	31.40	30.00	29.70	71.22	70.86	70.70
	$T^{0.80}$		6.32	5.04	4.78	38.68	37.14	36.54	77.26	76.54	76.38
		u_t^2	4.52	4.52	4.52	39.76	39.76	39.76	79.92	79.92	79.92
		y_t^2	22.58	20.24	19.36	56.08	54.18	53.76	84.12	83.46	83.34
400	$T^{0.50}$	$\widehat{u_t^2}$	12.24	13.02	13.10	43.94	43.90	43.86	89.96	90.02	89.88
	$T^{0.60}$		4.08	4.58	4.68	54.08	53.34	53.50	94.64	94.54	94.42
	$T^{0.70}$		4.52	4.20	4.04	63.70	62.56	62.22	96.80	96.60	96.56
	$T^{0.80}$		8.34	6.62	6.00	71.02	69.42	68.80	97.92	97.70	97.68
		u_t^2	5.12	5.12	5.12	70.10	70.10	70.10	97.84	97.84	97.84
		y_t^2	36.40	32.48	30.64	83.80	81.96	81.70	98.78	98.68	98.52
800	$T^{0.50}$	$\widehat{u_t^2}$	14.22	15.06	15.30	81.62	81.34	81.34	99.92	99.92	99.92
	$T^{0.60}$		4.86	5.22	5.38	88.68	88.30	88.22	99.98	99.98	99.98
	$T^{0.70}$		4.66	4.32	4.52	92.58	92.10	91.94	99.98	99.98	99.98
	$T^{0.80}$		9.46	6.28	5.74	95.22	94.10	93.82	99.98	99.98	99.98
		u_t^2	5.34	5.34	5.34	94.46	94.46	94.46	99.98	99.98	99.98
		y_t^2	51.20	45.04	43.28	98.02	97.68	97.46	99.98	99.98	99.98
1600	$T^{0.50}$	$\widehat{u_t^2}$	15.44	15.94	16.10	99.28	99.26	99.24	100.00	100.00	100.00
	$T^{0.60}$		5.32	5.64	5.92	99.72	99.68	99.66	100.00	100.00	100.00
	$T^{0.70}$		5.76	5.14	5.06	99.86	99.84	99.80	100.00	100.00	100.00
	$T^{0.80}$		11.98	7.16	6.22	99.92	99.90	99.88	100.00	100.00	100.00
		u_t^2	5.28	5.28	5.28	99.88	99.88	99.88	100.00	100.00	100.00
		y_t^2	65.84	57.46	55.22	99.98	99.96	99.96	100.00	100.00	100.00




FIG. 1. Average of $TS(\widehat{u}^\gamma)$ test statistic for $p = 5$ and various values of H under H_0 , $y_t = h_t \varepsilon_t$.

outperforms the 1-step ahead forecast $\widehat{h}_{j;T+1|T}^2$ formed from the following five parametric volatility models denoted by “ j ”.

We consider the following stationary volatility models for σ_t^2 in $y_t = \sigma_t \varepsilon_t$. We set $\varepsilon_t \sim$ i.i.d. $\mathcal{N}(0, 1)$ in all models except Model 4:

1. The stationary GARCH(1, 1) model of [4],

$$(34) \quad \sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2, \quad t = 1, \dots, T.$$

This model is the benchmark in our volatility forecasting.

2. The GJR-GARCH(1, 1) model of [24]:

$$\sigma_t^2 = \omega + \alpha u_{t-1}^2 + \gamma u_{t-1}^2 I(u_{t-1} < 0) + \beta \sigma_{t-1}^2, \quad t = 1, \dots, T.$$

This model enables simulation of the leverage effect, which can be important in forecasting volatility.

3. The APARCH(1, 1, 1) model of [15]:

$$\sigma_t^\delta = \omega + \alpha(|u_{t-1}| - \gamma u_{t-1})^\delta + \beta \sigma_{t-1}^\delta, \quad t = 1, \dots, T.$$

It adds flexibility capturing volatility dynamics and asymmetries via parameter δ .

4. The GARCH- $t(1, 1)$ of [5]. It uses (34) and assumes that ε_t follows Student $t(v)$ distribution with $v > 2$ (unknown) degrees of freedom. It allows us to assess whether the choice of an i.i.d. ε_t noise has impact on forecasting.

5. The stochastic volatility (SV) model of [33]. The presence of two separate generating noise provides to this model extra flexibility.

Parameters of these models are estimated using the quasi-maximum likelihood method.

In Monte Carlo simulations, h_t^2 is known. We use two volatility proxies, $p_t = h_t^2$ and y_t^2 . For the given volatility proxy p_t , the best forecasting method j minimises the average quadratic loss

$$(35) \quad \text{MSFE}_j = (T - T_0)^{-1} \sum_{t=T_0+1}^T (p_t - \widehat{h}_{j;t|t-1}^2)^2,$$

over $t \in (T_0, T]$. We set $T = 1000$ and $T_0 = 200$. Forecasting of h_t^2 with the kernel predictor $\widehat{h}_{t|t-1}^2$ is performed with fixed bandwidths $H = t^{0.60}, t^{0.65}, t^{0.70}$ and with a cross-validated bandwidth $H_{\text{CV},t}$, which minimises

$$(36) \quad \sum_{s=t-t_0}^t (p_s - \widehat{h}_{s|s-1}^2)^2$$

over $H = t^{0.55}, t^{0.60}, \dots, t^{0.75}$ where $t_0 = 50$.

To make comparisons across different forecasting methods, we use the benchmark $\text{MSFE}_{\text{GARCH}}$ of the parametric GARCH(1, 1) volatility model and calculate the relative root quadratic loss, $\text{RMSFE}_j = \frac{(\text{MSFE}_j)^{1/2}}{(\text{MSFE}_{\text{GARCH}})^{1/2}}$.

Tables 3 and 4 report the average value of the relative RMSFE_j over 1000 replications for data generating Models 2 and 3 of y_t , and two proxies, $p_t = h_t^2$ and y_t^2 . The smaller that the entry (< 1) is, the better the forecast.

Table 3 reports RMSFE results where the forecasting performance is evaluated using the “optimal” proxy $p_t = h_t^2$. The kernel forecasting methods produce the smallest values of RMSFE and clearly outperform the parametric forecasting methods of volatility in both Models 2 and 3, and the cross-validated bandwidth $H_{\text{CV},t}$ outperforms forecasts with a fixed

TABLE 3
Comparison of forecasting methods. Table reports RMSFE with the (true) volatility proxy $p_t = h_t^2$

Model 2 (with $h_t = \text{deterministic}$)						Model 3 (with $h_t = \text{persistent}$, $d = 1.4$)					
H	Kernel	GJR	SV	GARCH-t	APARCH	H	Kernel	GJR	SV	GARCH-t	APARCH
$t^{0.55}$	0.776	1.009	1.309	0.992	1.084	$t^{0.55}$	1.021	1.041	1.626	1.008	1.095
$t^{0.60}$	0.700					$t^{0.60}$	0.887				
$t^{0.65}$	0.660					$t^{0.65}$	0.790				
$t^{0.70}$	0.668					$t^{0.70}$	0.738				
$t^{0.75}$	0.738					$t^{0.75}$	0.742				
H_{CV}	0.544					H_{CV}	0.651				
H_{CV}^*	0.664					H_{CV}^*	0.894				

bandwidth. For comparison, we report RMSE for the kernel forecast with a feasible cross-validated bandwidth $H_{CV,t}^*$ that minimises (36) with the commonly used volatility proxy $p_t^* = y_t^2$; see [34]. It works well for Model 2 but slightly less well for Model 3.

Table 4 reports results of the same experiments as Table 3, using the imperfect, but observed, proxy $p_t = y_t^2$; see [31]. It is still noticeable that the kernel forecasts outperform the parametric forecasting methods, but the difference becomes marginal. It is clear that the choice of proxy for the cross-validation in empirical analysis is crucial. We use realised variance as a feasible volatility proxy in our empirical forecasting example below.

4. An empirical example. In this section, we illustrate the practical applicability of our testing methodology for the detection of stationary volatility.

We use weekly stock returns for a group of 254 companies in the S&P 500 over the period January 1994 to December 2019, obtained from *Bloomberg*. After data cleaning, it contains $T = 1340$ observations. In particular, we split the historical weekly returns into three subperiods: the pre-crisis period (January 1994–December 2007), the period covering the global financial crisis (January 2005–December 2012), and the post-crisis period (January 2011–December 2019).

In line with the finance literature, we assume that weekly returns of an individual company stock follow the model:

$$r_t = \mu_t + y_t \quad \text{where } y_t = r_t - \mu_t = h_t u_t,$$

where $\mu_t = E[r_t | \mathcal{F}_{t-1}]$ is the conditional mean and y_t is a white noise process. We test for the presence or absence of stationary volatility $\sigma_t^2 = \text{var}[r_t | \mathcal{F}_{t-1}]$ in u_t . If ARCH effects in u_t are not detected, then $r_t = \mu_t + h_t \varepsilon_t$, where $\{\varepsilon_t\}$ is an i.i.d. noise.

TABLE 4
Comparison of forecasting methods. Table reports RMSFE with the volatility proxy $p_t = y_t^2$

Model 2 (with $h_t = \text{deterministic}$)						Model 3 (with $h_t = \text{persistent}$, $d = 1.4$)					
H	Kernel	GJR	SV	GARCH-t	APARCH	H	Kernel	GJR	SV	GARCH-t	APARCH
$t^{0.55}$	0.988	1.001	1.026	1.000	1.005	$t^{0.55}$	1.000	1.001	1.023	1.000	1.003
$t^{0.60}$	0.985					$t^{0.60}$	0.997				
$t^{0.65}$	0.983					$t^{0.65}$	0.995				
$t^{0.70}$	0.984					$t^{0.70}$	0.995				
$t^{0.75}$	0.986					$t^{0.75}$	0.995				
H_{CV}	0.984					H_{CV}	0.997				

TABLE 5

Proportion of stock returns with no ARCH effects. Testing based on squares ($\gamma = 2$) and absolutes ($\gamma = 1$) for different bandwidths H , lags p and subperiods as defined in the main text. Testing at the 5% significance level. μ_t estimated using the single index model

H	Data	$p = 1$			$p = 5$			$p = 10$		
		1994–2007	2005–2012	2011–2019	1994–2007	2005–2012	2011–2019	1994–2007	2005–2012	2011–2019
Testing on squares ($\gamma = 2$)										
$T^{0.50} \widehat{u_t^2}$		93.31	97.24	97.64	97.64	96.06	96.46	94.88	91.73	92.91
$T^{0.55}$		89.76	94.88	96.85	94.49	96.85	97.24	96.85	96.06	98.03
$T^{0.60}$		86.61	90.55	95.67	90.55	92.13	96.06	93.31	93.70	98.43
$T^{0.65}$		81.10	82.28	93.31	86.61	85.43	95.28	87.80	85.43	97.24
$T^{0.70}$		76.77	74.02	92.52	77.95	72.44	92.13	77.17	71.26	93.31
$T^{0.75}$		68.50	65.75	92.52	64.17	57.09	90.16	62.21	54.72	91.73
	y_t^2	39.37	42.91	78.35	22.05	33.47	80.32	20.87	29.53	77.17
Testing on absolutes ($\gamma = 1$)										
$T^{0.50} \widehat{ u_t }$		89.76	94.49	94.49	88.58	84.25	86.22	80.32	72.44	69.69
$T^{0.55}$		86.61	93.70	93.31	90.95	91.34	90.16	89.76	87.80	83.07
$T^{0.60}$		81.50	87.80	94.09	87.40	89.37	93.70	86.61	87.80	89.37
$T^{0.65}$		73.62	80.71	91.73	80.71	83.47	94.09	84.25	82.68	92.91
$T^{0.70}$		62.60	70.87	90.16	68.11	70.87	89.76	70.47	70.87	92.52
$T^{0.75}$		53.15	59.06	89.37	51.18	56.30	88.58	52.76	53.54	90.55
	$ y_t $	12.60	30.71	72.05	2.36	24.41	71.65	1.97	18.11	72.05

In Table 5, we report the proportion of stock returns (in %) exhibiting no ARCH effects, according to our test, in u_t (among 254 stocks). Testing for ARCH effects is based on residuals $|\widehat{u_t}|^\gamma = (\widehat{h_t^\gamma})^{-1}|y_t|^\gamma$ where $y_t = r_t - \widehat{\mu}_t$, and conducted at the 5% significance level. To obtain an estimate for μ_t , we use the single index model $\mu_t = r_{f,t} + \beta_1(R_{m,t} - r_{f,t})$, where $R_{m,t}$ is the market factor and $r_{f,t}$ is the risk-free rate. Specifically, we would like to understand the impact of the choice for the bandwidth H , lag p , subperiods and values of γ on testing. The shaded grey range of $H = [T^{0.60}, \dots, T^{0.70}]$ denotes theoretically permissible values of H for smoothness parameter $\nu = 1$.

The empirical testing results can be summarised as follows: Across lags 1, 5, 10, $\gamma = 2, 1$ and three subperiods, in the recommended (shaded grey) range of $H = [T^{0.60}, \dots, T^{0.70}]$, the vast majority ($\sim 80\%$) of the stock returns have no ARCH effects. Further, in the more volatile subperiod 2005–2012, the proportion of stock returns with no ARCH effects falls slightly. However, when persistent volatility is not taken into account and testing is performed directly on the powers of y_t , the number of stock returns with no ARCH effects drops sharply.

Table 5 shows clear robustness of the testing results across values $\gamma = 1, 2$, which is a reassuring finding. Recall that estimation with $\gamma = 2$ requires at least sixth moments of u_t , whereas estimation with $\gamma = 1$ requires at least three. We include $\gamma = 1$, since by following the methodology of [35], we find that the majority of the stock returns have finite fifth moments, but not sixth. More details on this, and empirical testing results for $\gamma = 1/2$ can be found in Section 9 of the Supplementary Material [7].

Overall, these empirical results are in line with the Monte Carlo experiments and provide clear evidence that stationary volatility might be considerably less pronounced in the data, than previously thought.

Next, we consider the problem of forecasting of persistent volatility of the weekly log returns of some major stock indices and exchange rate series using data from the database “Oxford-Man Institute’s realised library,” Version 0.1, produced by [25] (see also [34]), over the period from January 3, 1999, to December 23, 2007 (469 observations).

TABLE 6
Comparison of the forecasting methods. Table reports RMSFE with the realised variance proxy p_t

Asset	Kernel	GJR	SV	GARCH-t	APARCH
Dow Jones	0.654	0.904	0.665	1.001	0.966
Nasdaq 100	0.696	1.238	1.086	1.137	1.311
Nikkei 250	0.831	1.008	0.971	0.996	1.117
S&P 500	0.705	0.901	0.934	1.085	0.851
USD British pound	0.989	1.038	1.116	1.002	1.160
USD Euro	0.915	1.013	1.043	1.000	1.042
USD Swiss franc	0.925	0.986	0.886	1.006	0.998

We consider the same forecasting methods of volatility as in the simulation study section. In order for these forecasting methods to be applied on returns with persistent volatility as in the model (33), we employ our test for ARCH effects from Section 2 on demeaned returns r_t of assets and select four stock indices and three exchange rates, which returns do not exhibit ARCH effects in squares. Instead of using a noisy proxy of squared returns for the volatility, we use the proxy p_t of “realised variance”; see [31].

Table 6 presents the values of relative RMSFE _{j} introduced in Section 3. Kernel prediction uses cross-validated bandwidth $H_{CV,t}$. The quadratic loss and the cross-validated bandwidth $H_{CV,t}$ are derived using the proxy of realised variance in (35) and (36). We set $T_0 = 200$, $t_0 = 50$. The results suggest that for the majority of asset returns, the kernel forecasting method with cross-validated bandwidth $H_{CV,t}$, significantly outperforms stationary alternatives.

5. Discussion. This paper contributes to the literature in three ways. First, we introduce a setup for persistent processes that can provide a general approximation to the volatility process of a time series. Second, we develop a consistent uniform estimation theory for the unobserved volatility processes, without strong parametric assumptions, and third, we suggest a testing strategy that enables the separation of stationary volatility from its persistent counterpart. To prove our main results, the uniform bounds for kernel-type estimates obtained in [14] and based on Bernstein inequalities for dependent random variables were used.

Testing results on U.S. stock returns provide extensive support for the persistent volatility paradigm, suggesting that the role of stationary conditional heteroskedasticity is not as outstanding in the data, as was previously thought. In addition, forecasting results on persistent volatility of log returns of stock indices and exchange rates provide evidence in favour of using kernel forecasting methods.

There are a number of interesting avenues for future work, in particular, the extension of our testing procedure to multivariate time series.

This is a distinct problem rather than a simple generalisation of the univariate case. In general, a multivariate model for a $p \times 1$ process

$$y_t = H_t u_t, \quad t = 1, \dots, T$$

can include a $p \times p$ matrix, H_t , of persistent volatility and a stationary $p \times 1$ white noise process u_t . The objective would be to test whether components of this white noise exhibit stationary conditional heteroskedasticity (ARCH effects). The estimation of H_t that would enable such testing could be undertaken using the work by [14].

Acknowledgements. We thank the Editor, the Associate Editor and the referees for constructive comments and valuable suggestions, which led to significant improvements in the paper.

Funding. The first author was supported by ESRC Grant ES/P000703/1.

SUPPLEMENTARY MATERIAL

Supplement to: “Choosing between persistent and stationary volatility” (DOI: [10.1214/22-AOS2236SUPP](https://doi.org/10.1214/22-AOS2236SUPP); .pdf). The Supplementary Material [7] provides proofs of all results given in the main paper and additional material for simulation study and empirical exercise.

REFERENCES

- [1] AMADO, C. and TERÄSVIRTA, T. (2013). Modelling volatility by variance decomposition. *J. Econometrics* **175** 142–153. [MR3061936 https://doi.org/10.1016/j.jeconom.2013.03.006](https://doi.org/10.1016/j.jeconom.2013.03.006)
- [2] AMADO, C. and TERÄSVIRTA, T. (2017). Specification and testing of multiplicative time-varying GARCH models with applications. *Econometric Rev.* **36** 421–446. [MR3608383 https://doi.org/10.1080/07474938.2014.977064](https://doi.org/10.1080/07474938.2014.977064)
- [3] BERA, A. K. and HIGGINS, M. L. (1993). ARCH models: Properties, estimation and testing. *J. Econ. Surv.* **7** 305–366.
- [4] BOLLERSLEV, T. (1986). Generalized autoregressive conditional heteroskedasticity. *J. Econometrics* **31** 307–327. [MR0853051 https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- [5] BOLLERSLEV, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *Rev. Econ. Stat.* **69** 542–547.
- [6] BROWNLEES, C. T. and GALLO, G. M. (2009). Comparison of volatility measures: A risk management perspective. *J. Financ. Econom.* **8** 29–56.
- [7] CHRONOPOULOS, I., GIRAITIS, L. and KAPETANIOS, G. (2022). Supplement to “Choosing between persistent and stationary volatility.” <https://doi.org/10.1214/22-AOS2236SUPP>
- [8] CHRONOPOULOS, I., KAPETANIOS, G. and PETROVA, K. (2021). Kernel-based volatility generalised least squares. *Econom. Stat.* **20** 2–11. [MR4302583 https://doi.org/10.1016/j.ecosta.2019.11.001](https://doi.org/10.1016/j.ecosta.2019.11.001)
- [9] DAHLHAUS, R. (2000). A likelihood approximation for locally stationary processes. *Ann. Statist.* **28** 1762–1794. [MR1835040 https://doi.org/10.1214/aos/1015957480](https://doi.org/10.1214/aos/1015957480)
- [10] DAHLHAUS, R. and POLONIK, W. (2006). Nonparametric quasi-maximum likelihood estimation for Gaussian locally stationary processes. *Ann. Statist.* **34** 2790–2824. [MR2329468 https://doi.org/10.1214/0090536060000000867](https://doi.org/10.1214/0090536060000000867)
- [11] DAHLHAUS, R. and SUBBA RAO, S. (2006). Statistical inference for time-varying ARCH processes. *Ann. Statist.* **34** 1075–1114. [MR2278352 https://doi.org/10.1214/0090536060000000227](https://doi.org/10.1214/0090536060000000227)
- [12] DALLA, V., GIRAITIS, L. and PHILLIPS, P. C. B. (2015). Testing mean stability of heteroskedastic time series. Cowles Foundation Discussion Paper No. 2006.
- [13] DALLA, V., GIRAITIS, L. and PHILLIPS, P. C. B. (2022). Robust tests for white noise and cross-correlation. *Econometric Theory* **38** 913–941. [MR4513134 https://doi.org/10.1017/S0266466620000341](https://doi.org/10.1017/S0266466620000341)
- [14] DENDRAMIS, Y., GIRAITIS, L. and KAPETANIOS, G. (2021). Estimation of time-varying covariance matrices for large datasets. *Econometric Theory* **37** 1100–1134. [MR4348398 https://doi.org/10.1017/S0266466620000535](https://doi.org/10.1017/S0266466620000535)
- [15] DING, Z., GRANGER, C. W. and ENGLE, R. F. (1993). A long memory property of stock market returns and a new model. *J. Empir. Finance* **1** 83–106.
- [16] ENGLE, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **50** 987–1007. [MR0666121 https://doi.org/10.2307/1912773](https://doi.org/10.2307/1912773)
- [17] ENGLE, R. F. and RANGEL, J. G. (2008). The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Rev. Financ. Stud.* **21** 1187–1222.
- [18] ERAKER, B., JOHANNES, M. and POLSON, N. (2003). The impact of jumps in volatility and returns. *J. Finance* **58** 1269–1300.
- [19] GALLANT, A. R. (1984). The Fourier flexible form. *Amer. J. Agric. Econ.* **66** 204–208.
- [20] GIRAITIS, L., KAPETANIOS, G., WETHERILT, A. and ŽIKEŠ, F. (2016). Estimating the dynamics and persistence of financial networks, with an application to the sterling money market. *J. Appl. Econometrics* **31** 58–84. [MR3464851 https://doi.org/10.1002/jae.2457](https://doi.org/10.1002/jae.2457)
- [21] GIRAITIS, L., KAPETANIOS, G. and YATES, T. (2014). Inference on stochastic time-varying coefficient models. *J. Econometrics* **179** 46–65. [MR3153648 https://doi.org/10.1016/j.jeconom.2013.10.009](https://doi.org/10.1016/j.jeconom.2013.10.009)
- [22] GIRAITIS, L., KAPETANIOS, G. and YATES, T. (2018). Inference on multivariate heteroscedastic time varying random coefficient models. *J. Time Series Anal.* **39** 129–149. [MR3767057 https://doi.org/10.1111/jtsa.12271](https://doi.org/10.1111/jtsa.12271)

- [23] GIRAITIS, L., KOUL, H. L. and SURGAILIS, D. (2012). *Large Sample Inference for Long Memory Processes*. Imperial College Press, London. [MR2977317](#) <https://doi.org/10.1142/p591>
- [24] GLOSTEN, L. R., JAGANNATHAN, R. and RUNKLE, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Finance* **48** 1779–1801.
- [25] HEBER, G., LUNDE, A., SHEPHARD, N. and SHEPPARD, K. (2009). Oxford-Man Institute's realized library, version 0.1. Available at <https://realized.oxford-man.ox.ac.uk/data/download>.
- [26] JIANG, F., LI, D. and ZHU, K. (2021). Adaptive inference for a semiparametric generalized autoregressive conditional heteroskedasticity model. *J. Econometrics* **224** 306–329. [MR4303783](#) <https://doi.org/10.1016/j.jeconom.2020.10.007>
- [27] KAPETANIOS, G. and YATES, T. (2014). Evolving UK and US macroeconomic dynamics through the lens of a model of deterministic structural change. *Empir. Econ.* **47** 305–345.
- [28] MAZUR, B. and PIPIEN, M. (2012). On the empirical importance of periodicity in the volatility of financial returns-time varying GARCH as a second order APC (2) process. *Cent. Eur. J. Econ. Model. Econom.* **4** 95–116.
- [29] MERCURIO, D. and SPOKOINY, V. (2004). Statistical inference for time-inhomogeneous volatility models. *Ann. Statist.* **32** 577–602. [MR2060170](#) <https://doi.org/10.1214/009053604000000102>
- [30] MIKOSCH, T. and STĂRICĂ, C. (2004). Nonstationarities in financial time series, the long-range dependence, and the IGARCH effects. *Rev. Econ. Stat.* **86** 378–390.
- [31] PATTON, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *J. Econometrics* **160** 246–256. [MR2745881](#) <https://doi.org/10.1016/j.jeconom.2010.03.034>
- [32] PRIESTLEY, M. B. (1965). Evolutionary spectra and non-stationary processes. *J. Roy. Statist. Soc. Ser. B* **27** 204–237. [MR0199886](#)
- [33] RUIZ, E. (1994). Quasi-maximum likelihood estimation of stochastic volatility models. *J. Econometrics* **63** 289–306.
- [34] SHEPHARD, N. and SHEPPARD, K. (2010). Realising the future: Forecasting with high-frequency-based volatility (heavy) models. *J. Appl. Econometrics* **25** 197–231. [MR2758633](#) <https://doi.org/10.1002/jae.1158>
- [35] TRAPANI, L. (2016). Testing for (in)finite moments. *J. Econometrics* **191** 57–68. [MR3434435](#) <https://doi.org/10.1016/j.jeconom.2015.08.006>
- [36] VAN BELLEGEM, S. and VON SACHS, R. (2004). Forecasting economic time series with unconditional time-varying variance. *Int. J. Forecast.* **20** 611–627.