

RESEARCH ARTICLE

Multivariate fractional integration tests allowing for conditional heteroskedasticity with an application to return volatility and trading volume

Marina Balboa¹ | Paulo M. M. Rodrigues² | Antonio Rubia¹ | A. M. Robert Taylor³

¹Department of Financial Economics and Accounting, University of Alicante, Alicante, Spain

²Banco de Portugal and Nova SBE, Universidade Nova de Lisboa, Lisbon, Portugal

³Essex Business School, University of Essex, Essex, UK

Correspondence

A. M. Robert Taylor, Essex Business School, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK.
 Email: robert.taylor@essex.ac.uk

Funding information

Spanish Ministry of Economy, Grant/Award Number: ECO2017-87069-P; Portuguese Science Foundation (FCT), Grant/Award Number: PTDC/EGE-ECO/28924/2017, UID/ECO/00124/2013; Social Sciences DataLab, Project, Grant/Award Number: 22209; POR Lisboa, Grant/Award Number: LISBOA-01-0145-FEDER-007722; POR Norte

Summary

We introduce a new joint test for the order of fractional integration of a multivariate fractionally integrated vector autoregressive (FIVAR) time series based on applying the Lagrange multiplier principle to a feasible generalised least squares estimate of the FIVAR model obtained under the null hypothesis. A key feature of the test we propose is that it is constructed using a heteroskedasticity-robust estimate of the variance matrix. As a result, the test has a standard χ^2 limiting null distribution under considerably weaker conditions on the innovations than are permitted in the extant literature. Specifically, we allow the innovations driving the FIVAR model to follow a vector martingale difference sequence allowing for both serial and cross-sectional dependence in the conditional second-order moments. We also do not constrain the order of fractional integration of each element of the series to lie in a particular region, thereby allowing for both stationary and non-stationary dynamics, nor do we assume any particular distribution for the innovations. A Monte Carlo study demonstrates that our proposed tests avoid the large oversizing problems seen with extant tests when conditional heteroskedasticity is present in the data. We report an empirical case study for a sample of major US stocks investigating the order of fractional integration in trading volume and different measures of volatility in returns, including realised variance. Our results suggest that both return volatility and trading volume are fractionally integrated, but with the former generally found to be more persistent (having a higher fractional exponent) than the latter, when more reliable proxies for volatility such as the range or realised variance are used.

1 | INTRODUCTION

Long memory models have been used to model time series data in a wide range of fields of application. The class of (multivariate) fractionally integrated autoregressive moving average (ARFIMA) models provides a parsimonious means of simultaneously modelling the patterns of long and short range dependence typically seen in many macroeconomic and financial data sets; see, for example, the surveys in Baillie (1996) and Robinson (2003). In the context of the ARFIMA class of models, the long memory parameter, or fractional exponent (vector of exponents in the multivariate case), is the key parameter driving the behaviour of the series. Where this is zero, a weakly dependent (short memory) ARMA series obtains. If it is less than one half, the series is weakly stationary; otherwise, it is non-stationary, the familiar autoregressive

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. Journal of Applied Econometrics published by John Wiley & Sons Ltd.

unit root case occurring where the exponent is unity. Consequently, considerable interest has been paid to developing methods of inference on the fractional exponent both as a parameter of interest in its own right and for preliminary data analysis. A leading example is a test of the null hypothesis of weak dependence (against fractional alternatives); here a non-rejection would allow for the use of standard methods for conducting, among other things, causality, structural vector autoregression or impulse response analyses. More generally, such tests could usefully be employed to indicate what order of differencing of the data is required for such methods to be suitably employed.

In the univariate setting, a number of hypothesis tests on the fractional exponent have been proposed; see, among others, Robinson (1994), Tanaka (1999), Breitung and Hassler (2002), Nielsen (2004b), Demetrescu et al. (2008), Hassler et al. (2009, 2016) and Cavaliere et al. (2017). In the context of a vector series, one could perform such univariate fractional integration tests separately on each element of the vector. However, the overall size of such a testing procedure would be hard to control. Moreover, multivariate testing can improve efficiency relative to univariate testing because it explicitly acknowledges and exploits the existence of any endogenous cross dependencies in the vector series, which can reduce the variability in the estimation errors and, hence, improve efficiency in estimation and testing. In this paper, we develop multivariate fractional integration tests designed to test joint null hypotheses concerning the values of the long memory parameters of the elements of a fractionally integrated vector autoregressive (FIVAR) model. Specifically, we propose a parametric multivariate Lagrange multiplier (LM)-type test in the time domain, which generalises the univariate regression-based LM-type test of Demetrescu et al. (2008) to the multivariate case. The method we propose can also be used to construct confidence sets, at a given asymptotic coverage level, for the true values of the long memory coefficients.

Our testing procedure is implemented in a regression-based context, based on feasible generalised least squares (FGLS) estimation of the multivariate FIVAR model under the null hypothesis, coupled with a heteroskedasticity-robust variance matrix estimate. A key advantage of, and motivation for, this approach is that it allows us to significantly weaken the technical conditions needed on the innovations, relative to existing multivariate fractional tests including, among others, Robinson (1995), Lobato and Robinson (1998), Lobato (1999), Lobato and Velasco (2000), Marinucci and Robinson (2001), Breitung and Hassler (2002), Shimotsu (2007) and (Nielsen, 2004b, 2005, 2011). In particular, we allow the driving innovations in the data generating process (DGP) to follow a vector martingale difference sequence (MDS), which is permitted to exhibit time-varying conditional heteroskedasticity. This therefore allows for both serial and cross-sectional dependence in the conditional second-order moments, which is of particular empirical relevance when modelling financial data and is not, to the best of our knowledge, allowed by any extant multivariate fractional integration test.

Like Nielsen (2004a, 2005), we work within the context of a multivariate FIVAR model. This model allows each series within the vector process to have different fractional exponents irrespective of the parameters of the short-run component of the model. This property is not guaranteed when using the class of vector autoregressive fractionally integrated (VARFI) models where the orders of integration of the elements of the vector series are not constant throughout the parameter space of the model; for further details see Nielsen (2005, pp. 381–382). This is important for the empirical case study we consider in this paper with respect to trading volume and return volatility where we aim to explicitly investigate whether the data support the hypothesis that these series admit a common fractional exponent or not. For a further recent empirical application using FIVAR models, investigating the effects of monetary policy on the economy, where it is important to allow the elements of the vector time series to have potentially different fractional exponents, see Lovcha and Perez-Laborda (2018). An implication of the FIVAR model, however, is that fractional cointegration is not possible between the elements of the vector time series. In common with the tests in Nielsen (2005), we do not restrict the fractional exponents to lie within a particular region, thereby allowing for both stationary and non-stationary dynamics. We also do not impose any particular distributional law on the innovations.

Under the conditionally heteroskedastic setting outlined above, our proposed test retains a standard χ^2 limiting null distribution (irrespective of the null values of the long memory parameters being tested) and exhibits non-trivial power against a sequence of local (Pitman drift) alternatives. Moreover, where the errors are independent and identically distributed (i.i.d.) our test statistic is asymptotically equivalent to the multivariate LM statistic discussed in Nielsen (2004a, 2005). As a consequence, the LM-type test we propose is asymptotically (locally) efficient when the errors are i.i.d. Gaussian. Monte Carlo simulation experiments show that our proposed multivariate fractional integration test displays good finite sample size control and power performance in the presence of empirically relevant data features, such as short-run dependence and time-varying GARCH-type conditional variances for both Gaussian and non-Gaussian innovations. In contrast, extant tests are shown to display quite poor finite sample size control in the presence of such features.

Multivariate testing is naturally intended to address joint hypotheses involving the degree of persistence of a set of variables. This has important practical applications. As a leading example, there has been considerable interest in both

the theoretical and empirical finance literatures on understanding the link between trading volume and return volatility. A number of papers have analysed if the long-run dynamics of these variables share a common order of fractional integration, with mixed evidence; see, for example, Bollerslev and Jubinski (1999), Lobato and Velasco (2000), Luu and Martens (2003), Fleming and Kirby (2011) and Rossi and de Magistris (2013).

In our empirical analysis, we apply our new approach to conduct joint inference on the order of fractional integration of trading volumes and different measures of return volatility focusing on 30 major US stocks from the Dow Jones Industrial Average Index (DJI). We also investigate the existence of a common order of fractional integration between volume and these measures of return volatility. Because our tests do not require a particular distribution and, more importantly, allow for time-varying conditional second-order moments, our results are likely to be more robust than those reported in previous studies, which are based on estimation techniques which neglect these empirically relevant data features (e.g., Bollerslev & Jubinski, 1999; Lobato & Velasco, 2000; Fleming & Kirby, 2011).

Together with daily log volume, we consider the log transformations of three alternative measures of return volatility with increasing degrees of efficiency, namely, absolute-valued returns, the range-based estimator of Garman and Klass (1980) and a measure of realised variance constructed from 5-min returns. An important aspect of this analysis is to investigate the influence that measurement errors have on the conclusions drawn from the data. Our empirical findings suggest that a common fractional exponent cannot in general be rejected when return volatility is proxied by absolute-valued returns but can be rejected when it is proxied by more accurate estimates such as the range or realised variance. These findings are consistent with the previous literature and help us to understand the disparity between empirical results where different proxies for volatility are used. Our empirical results indicate that return volatility tends to exhibit a larger fractional integration exponent than trading volume, with long-term behaviour possibly driven by non-stationary dynamics. Heterogeneous degrees of fractional integration, such that return volatility tends to be more persistent than trading volume, could originate in certain types of trading strategies associated with imitation and herding in investors and market microstructure environmental conditions; see, for example, LeBaron and Yamamoto (2008) and Yamamoto (2011).

The remainder of the paper is organised as follows. Section 2 introduces the DGP and the main assumptions underlying our theoretical analysis. In Section 3, we detail our new LM-type multivariate fractional integration test and derive its asymptotic distribution under both the null hypothesis and a sequence of local alternatives. Section 4 discusses the results of our finite sample Monte Carlo study. Section 5 analyses the empirical relationship between trading volume and return volatility for stocks from the DJI. Section 6 concludes. Supporting information appendix contains mathematical proofs of the large sample results given in Section 3 together with additional material relating to the Monte Carlo analysis in Section 4 and to the empirical application in Section 5.

In what follows, $\Rightarrow$ and $\xrightarrow{p}$ denote weak convergence and convergence in probability, respectively, as $T \rightarrow \infty$. $I(\cdot)$ is an indicator function that equals 1 if the condition in parenthesis is fulfilled and equals 0 otherwise. The operators $\otimes$ and $\odot$ correspond to the Kronecker and Hadamard products, respectively. The quantities $\mathbf{I}_n$ and $\mathbf{0}_{n \times m}$ denote an n -dimensional identity matrix and an $n \times m$ zero matrix, respectively. The notation $\mathbf{A} = \{a_{ij}\}$ denotes that the (i, j) th element of the matrix $\mathbf{A}$ is given by a_{ij} .

2 | A FIVAR MODEL WITH HETEROSKEDASTICITY

We consider the observable k -dimensional time series vector $\{\mathbf{y}_t\}_{t=1}^T$, where $\mathbf{y}_t \equiv (y_{1,t}, \dots, y_{k,t})'$ is generated according to the DGP:

$$\Delta^{d+\theta}(L)\mathbf{y}_t = \varepsilon_t I(t \geq 1) \quad (1)$$

where $\Delta^{d+\theta}(L)$ is a $k \times k$ diagonal matrix polynomial in the conventional lag operator, L , with characteristic element given by $(1 - L)^{d_i+\theta_i}$, $i \in \{1, \dots, k\}$. The real-valued fractional exponent, $d_i + \theta_i$, is commonly referred to as the long memory or fractional integration parameter, such that $\mathbf{d} + \boldsymbol{\theta} \equiv (d_1 + \theta_1, \dots, d_k + \theta_k)'$. The k -dimensional vector $\varepsilon_t \equiv (\varepsilon_{1,t}, \dots, \varepsilon_{k,t})'$ is a weakly dependent (short memory or $I(0)$) noise process with bounded spectral density that is bounded away from zero at the origin. Our focus is on developing tests of the null hypothesis that $\mathbf{d}$ is the true order of integration of $\{\mathbf{y}_t\}$; that is, $H_0 : \boldsymbol{\theta} = \mathbf{0}$, against the alternative hypothesis that at least one element of $\boldsymbol{\theta}$ is non-zero.

Assumption 1 details the formal properties, which we will assume to hold on $\{\varepsilon_t\}$ in (1).

Assumption 1. $\{\varepsilon_t\}$ in (1) is generated as $\Pi(L)\varepsilon_t = \mathbf{e}_t \equiv (e_{1,t}, \dots, e_{k,t})'$, with $\Pi(L) := \mathbf{I}_p - \sum_{j=1}^p \Pi_j L^j$, where Π_j are $k \times k$ parameter matrices such that $\Pi(L)$ has all of its roots lying outside the unit circle and $\{\mathbf{e}_t\}$ satisfies the following conditions:

- (A1) $E(\mathbf{e}_t) = 0$ and $E(\mathbf{e}_t \mathbf{e}_t') = \Sigma$, with Σ positive definite.
 (A2) $\sup_t E(|\mathbf{e}_t|^{4+\eta}) < \infty$ for some $\eta > 0$.
 (A3) $\{\mathbf{e}_t, \mathcal{F}_t\}_{t=-\infty}^{\infty}$ is a strictly stationary and ergodic vector MDS, with respect to the natural filtration $\mathcal{F}_t$, the σ -field generated by $\{\mathbf{e}_s : s \leq t\}$.
 (A4) $\sum_{i=1}^{\infty} \sum_{j=1, i \neq j}^{\infty} E|e_{h,t} e_{s,t} e_{r,t-i} e_{u,t-j}| < \infty$, for any $1 \leq h, s, r, u \leq k$.

Remark 1. Assumption 1 allows the short memory component of $\{\mathbf{y}_t\}$ to be driven by a stationary VAR(p) process. Accordingly, (1) is a FIVAR model in which each component $\{y_{i,t}\}$, $i = 1, \dots, k$, follows a Type-II ARFIMA ($p, d_i + \theta_i, 0$) process. The choice of Type-II fractional integration in our setting has the desirable feature that the same definition is valid for an arbitrarily large range of admissible values of the fractional parameters, $d_i + \theta_i$, $i = 1, \dots, k$; in particular, these are not restricted to lie in the interval $(-0.5, 0.5)$, a necessary condition for stationarity and invertibility.

Remark 2. (A1) and (A2) are standard moment conditions. Unlike the existing multivariate fractional integration tests discussed in Section 1, (A3) allows the innovations to exhibit time-varying conditional variances. The absolute summability condition (A4) limits the amount of temporal and cross-sectional dependence in the second-order moments and is equivalent to requiring absolutely summable fourth-order joint cumulants. Our conditions are weaker than requiring $\{\mathbf{e}_t, \mathcal{F}_t\}$ to be either conditionally homoskedastic or independent, both of which imply (A4). Finally, (A1) and (A3) imply that $E(e_{i,t} e_{j,t+h}) = 0$ whenever $h \neq 0$, but allows $E(e_{i,t} e_{j,t}) \neq 0$ when $i \neq j$ because Σ is not restricted to be diagonal.

Remark 3. Assumption 1 imposes that the unconditional variance matrix of $\mathbf{e}_t$, Σ , is constant. However, the pivotal χ^2 limiting null distribution of our proposed FGLS statistic, LM_d^{FGLS} , defined below in (8), in Theorem 2 under the conditions of Assumption 1 remains valid in so-called non-stationary volatility cases where $E(\mathbf{e}_t \mathbf{e}_t') = \Sigma_t = \Sigma_t \Sigma_t'$, provided the unconditional volatility matrix, Σ_t , satisfies the regularity conditions detailed in, for example, Assumption 2(a) of Boswijk et al. (2016).¹ In particular, these entail that $\Sigma_t := \Sigma(t/T)$, for all $t = 1, \dots, T$, where $\Sigma(\cdot)$ is a non-stochastic element of the space of $k \times k$ matrices of càdlàg functions on $[0, 1]$ equipped with the Skorokhod metric, and is such that $\Sigma(u) := \Sigma(u) \Sigma(u)'$ is positive definite for all $u \in [0, 1]$. For further discussion, including a number of examples satisfying these conditions, see Boswijk et al. (2016, p. 66).

Remark 4. Under Assumption 1, the model in (1) can be re-written as $\Pi(L) \Delta^{d+\theta}(L) \mathbf{y}_t = \mathbf{e}_t$. Given the stationarity restriction imposed under Assumption 1, for a sufficiently large value of p , the FIVAR representation could be viewed as an approximation to the more general class of FIVARMA models, although we treat p as fixed (independent of the sample size) in this paper. We conjecture that it should be possible to extend our analysis to allow p to increase with the sample size but this would considerably complicate the theoretical analysis and is beyond the scope of this paper.

Remark 5. The FIVAR model in (1) under Assumption 1 rules out the possibility of fractional cointegration between the elements of $\{\mathbf{y}_t\}$; for further discussion, see, among others, Sela and Hurvich (2009) and Nielsen (2005, pp. 381–382). The maintained assumption of no fractional cointegration is also made in all of the extant multivariate fractional integration tests cited in the Introduction. However, noting from Remark 10 below that the feasible GLS multivariate fractional integration test we propose in Section 3.2 is asymptotically equivalent to the multivariate LM fractional test in Nielsen (2005), then for the same reasons as are discussed in Nielsen (2005, pp. 378–379), the LM-type test, LM_d^{FGLS} , developed in Section 3, is also implicitly a test of the null of no fractional cointegration (in the sense defined in Nielsen, 2005, p. 378) and will diverge at rate $O_p(T)$ under fractional cointegration.² It therefore seems advisable to consider the tests proposed in this paper alongside tests for fractional cointegration. We

¹Numerical experiments investigating the properties of the LM_d^{FGLS} test for data with a one-time break in unconditional variance are reported in the supporting information appendix. These results suggest that even quite large variance breaks have very little impact on the finite sample size of the FGLS-based tests.

²Numerical experiments investigating rejection rates of the LM_d^{FGLS} test and the tests of Nielsen (2005) and Breitung and Hassler (2002) in a fractionally cointegrated model are reported in the supporting information appendix. These show that, as expected, all three tests display empirical rejection frequencies in excess of the nominal level, which are larger, other things equal, the larger is T or the strength of cointegration. Of the three tests, our FGLS test tends to reject with slightly lower frequency than the other two tests.

adopt this approach in the empirical application in Section 5 by also considering the procedures developed in Nielsen and Shimotsu (2007).

3 | A MULTIVARIATE LM-TYPE FRACTIONAL INTEGRATION TEST

3.1 | Preliminaries

Given the observable time series vector $\{\mathbf{y}_t\}$ generated as in (1) and an arbitrary real-valued vector $\mathbf{g} \equiv (g_1, \dots, g_k)'$, define the k -dimensional stochastic processes

$$\boldsymbol{\varepsilon}_{t,\mathbf{g}} \equiv (\varepsilon_{1,t,\mathbf{g}_1}, \dots, \varepsilon_{k,t,\mathbf{g}_k})' := (1 - L)_+^{\mathbf{g}} \mathbf{y}_t = \sum_{j=0}^{t-1} \Lambda_j(\mathbf{g}) \mathbf{y}_{t-j}, \quad (2)$$

where $(1 - L)_+^{\mathbf{g}} := \sum_{j=0}^{t-1} \Lambda_j(\mathbf{g}) L^j$, and

$$\mathbf{z}_{t-1,\mathbf{g}}^* \equiv (z_{1,t-1,\mathbf{g}_1}^*, \dots, z_{k,t-1,\mathbf{g}_k}^*)' := \sum_{j=1}^{t-1} j^{-1} \boldsymbol{\varepsilon}_{t-j,\mathbf{g}}, \quad t = 2, \dots, T \quad (3)$$

with $\{\Lambda_j(\mathbf{g})\}_{j=0}^{t-1}$ denoting a sequence of $k \times k$ diagonal matrices with i th diagonal element

$$\lambda_0(g_i) := 1, \text{ and } \lambda_j(g_i) := \frac{j-1-g_i}{j} \lambda_{j-1}(g_i), \quad j \geq 1, \quad (4)$$

corresponding to the truncated series of polynomial coefficients in the binomial expansion $(1 - L)^{g_i} := \sum_{j=0}^{\infty} \lambda_j(g_i) L^j$. These variables are straightforward generalisations of the corresponding univariate processes in Breitung and Hassler (2002) to the multivariate context, with the characteristic harmonic weighting in (3) arising from the derivative of a (Gaussian) score function. Remark 10 below gives further insight into the key role played by these variables in the construction of our proposed LM-type test statistic. Let Φ denote a $k \times k$ diagonal matrix with i th diagonal element ϕ_{ii} , $i = 1, \dots, k$. Under Assumption 1, testing the null hypothesis that $\mathbf{d}$ is the true order of integration of $\{\mathbf{y}_t\}$, $H_0 : \boldsymbol{\theta} = \mathbf{0}$, is equivalent to testing $H_0 : \Phi = \mathbf{0}$ in the multivariate linear regression model

$$\boldsymbol{\varepsilon}_{t,\mathbf{d}} = \Phi \mathbf{z}_{t-1,\mathbf{d}}^* + \sum_{j=1}^p \Pi_j \boldsymbol{\varepsilon}_{t-j,\mathbf{d}} + \mathbf{v}_t, \quad t = p^* + 1, \dots, T \quad (5)$$

where $p^* := \max(1, p)$. This equivalence holds because, under $H_0 : \boldsymbol{\theta} = \mathbf{0}$, (5) and (1) are bijective with $\phi_{ii} = 0$ for all $i = 1, \dots, k$ and $\mathbf{v}_t = \mathbf{e}_t$ in (5); see also Breitung and Hassler (2002), Demetrescu et al. (2008), and Hassler et al. (2009).

It will prove convenient to re-write (5) in matrix notation. First, corresponding to the time series of observations for each element of $\mathbf{y}_t$, we have the equivalent representation:

$$\mathbf{Y}_{i,d_i} = \mathbf{X}_{i,-1,\mathbf{d}}^* \boldsymbol{\beta}_i + \mathbf{u}_i, \quad 1 \leq i \leq k \quad (6)$$

where $\mathbf{Y}_{i,d_i} := (\varepsilon_{i,p^*+1,d_i}, \dots, \varepsilon_{i,T,d_i})'$ is a $(T - p^*) \times 1$ vector, $\boldsymbol{\beta}_i := (\phi_{ii}, \Pi_{i1}, \dots, \Pi_{ip})'$ is a k' -dimensional parameter vector, with $k' := pk + 1$, and Π_{ij} denotes the i -th row of Π_j , $j = 1, \dots, p$, $\mathbf{u}_i := (v_{i,p^*+1}, \dots, v_{i,T})'$ is a $(T - p^*) \times 1$ vector of innovations and $\mathbf{X}_{i,-1,\mathbf{d}}^*$ is the $(T - p^*) \times k'$ matrix of observations of the (lagged) right-hand side variables $\mathbf{x}_{i,t-1,\mathbf{d}}^* := (z_{i,t-1,d_i}^*, \varepsilon'_{t-1,\mathbf{d}}, \dots, \varepsilon'_{t-p,\mathbf{d}})'$. With the exception of the first regressor, all other right-hand side variables that characterise the i th equation (6) are the same, because these always correspond to lagged values of $\boldsymbol{\varepsilon}_{t,\mathbf{d}}$. Then, given $T' := k(T - p^*)$,

we can write the system of equations (6) compactly as $\mathbf{Y}_d = \mathbf{X}_{-1,d}^* \boldsymbol{\beta} + \mathbf{u}$, with these terms defined implicitly as

$$\begin{bmatrix} \mathbf{Y}_{1,d_1} \\ \mathbf{Y}_{2,d_2} \\ \vdots \\ \mathbf{Y}_{k,d_k} \end{bmatrix}_{T' \times 1} = \begin{bmatrix} \mathbf{X}_{1,-1,d}^* & \mathbf{0}_{(T-p^*) \times k'} & \cdots & \mathbf{0}_{(T-p^*) \times k'} \\ \mathbf{0}_{(T-p^*) \times k'} & \mathbf{X}_{2,-1,d}^* & \cdots & \mathbf{0}_{(T-p^*) \times k'} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{(T-p^*) \times k'} & \mathbf{0}_{(T-p^*) \times k'} & \cdots & \mathbf{X}_{k,-1,d}^* \end{bmatrix}_{T' \times kk'} \begin{bmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \\ \vdots \\ \boldsymbol{\beta}_k \end{bmatrix}_{kk' \times 1} + \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_k \end{bmatrix}_{T' \times 1}.$$

3.2 | A heteroskedasticity-robust LM test

Under Assumption 1 and $H_0 : \boldsymbol{\theta} = \mathbf{0}$, it follows that $E(\mathbf{u}\mathbf{u}') = \Sigma \otimes \mathbf{I}_{T-p^*}$. Equation (5) defines a seemingly unrelated regression equation (SURE) system. Although equation-by-equation ordinary least squares [OLS] estimation will deliver consistent estimates of $\boldsymbol{\beta}$, these estimates will not be efficient unless Σ is diagonal (recalling that the regressors differ across the equations in the system). We will therefore consider a FGLS estimator of $\boldsymbol{\beta}$ based on a preliminary consistent estimate of Σ (obtained using OLS residuals estimated on an equation-by-equation basis). The resulting FGLS estimator of $\boldsymbol{\beta}$ is defined as follows:

$$\hat{\boldsymbol{\beta}} := \left(\mathbf{X}_{-1,d}^{*'} \left[\tilde{\Sigma}^{-1} \otimes \mathbf{I}_{T-p^*} \right] \mathbf{X}_{-1,d}^* \right)^{-1} \left(\mathbf{X}_{-1,d}^{*'} \left[\tilde{\Sigma}^{-1} \otimes \mathbf{I}_{T-p^*} \right] \mathbf{Y}_d \right) \quad (7)$$

where $\tilde{\Sigma} = \{\tilde{\sigma}_{ij}\}$ is estimated as $\tilde{\sigma}_{ij} := T^{-1} \tilde{\mathbf{u}}_i' \tilde{\mathbf{u}}_j$, with $\tilde{\mathbf{u}}_s := \mathbf{Y}_{s,d_s} - \mathbf{X}_{s,-1,d}^* \tilde{\boldsymbol{\beta}}_s$ and $\tilde{\boldsymbol{\beta}}_s$ denotes the equation-by-equation OLS estimate of $\boldsymbol{\beta}_s$, $s = 1, \dots, k$ in (6).³

In Theorem 1, we now characterise the asymptotic distribution of the FGLS estimate, $\hat{\boldsymbol{\beta}}$, under Assumption 1 and $H_0 : \boldsymbol{\theta} = \mathbf{0}$.

Theorem 1. Let $\mathbf{y}_t$ be generated according to (1) and let $\hat{\boldsymbol{\beta}}$ be the vector of FGLS estimates defined in (7). Under Assumption 1 and $H_0 : \boldsymbol{\theta} = \mathbf{0}$, $\sqrt{T}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \Rightarrow \mathcal{N}(\mathbf{0}, \Omega_{\boldsymbol{\beta}})$ where $\boldsymbol{\beta}_0 \equiv (\boldsymbol{\beta}'_{01}, \dots, \boldsymbol{\beta}'_{0k})'$ with $\boldsymbol{\beta}_{0s} := (0, \Pi_{s1}, \dots, \Pi_{sp})'$, $s = 1, \dots, k$, and $\Omega_{\boldsymbol{\beta}} := \mathbf{A}_{\boldsymbol{\beta}}^{-1} \mathbf{B}_{\boldsymbol{\beta}} \mathbf{A}_{\boldsymbol{\beta}}^{-1}$, with $\mathbf{A}_{\boldsymbol{\beta}} := \text{plim}_{T \rightarrow \infty} \left(\frac{1}{T} \mathbf{X}_{-1,d}^{*'} \left[\Sigma^{-1} \otimes \mathbf{I}_{T-p^*} \right] \mathbf{X}_{-1,d}^* \right)$, $\mathbf{B}_{\boldsymbol{\beta}} := \text{plim}_{T \rightarrow \infty} \left(\frac{1}{T} \mathbf{w}_{-1,d}^* \mathbf{w}_{-1,d}^{*'} \right)$, and $\mathbf{w}_{-1,d}^* := \mathbf{X}_{-1,d}^{*'} \left[\Sigma^{-1} \otimes \mathbf{I}_{T-p^*} \right] \mathbf{u}$.

The dependence of the asymptotic variance of the FGLS estimator on nuisance parameters arising from any weak dependence and/or cross sectional correlation in ϵ_t implies that asymptotically pivotal inference on the long memory parameters will need to be based on a heteroskedasticity-robust statistic formed using a consistent estimate of $\Omega_{\boldsymbol{\beta}}$. This can be achieved by using the familiar Eicker–Huber–White approach building on the preliminary OLS estimate $\tilde{\Sigma}$ and the FGLS residuals $\hat{\mathbf{u}} := \mathbf{Y}_d - \mathbf{X}_{-1,d}^* \hat{\boldsymbol{\beta}}$. In particular, a heteroskedasticity-robust estimate of the variance matrix $\Omega_{\boldsymbol{\beta}}$ is given by $\hat{\Omega}_{\boldsymbol{\beta}} := \mathbf{A}_T^{*-1} \mathbf{B}_T^* \mathbf{A}_T^{*-1}$, where $\mathbf{A}_T^* := \mathbf{X}_{-1,d}^{*'} \left[\tilde{\Sigma}^{-1/2} \otimes \mathbf{I}_{T-p^*} \right] \mathbf{X}_{-1,d}^* / T$ and $\mathbf{B}_T^* := \hat{\mathbf{w}}_{-1,d}^* \hat{\mathbf{w}}_{-1,d}^{*'} / T$, with $\hat{\mathbf{w}}_{-1,d}^* := \mathbf{X}_{-1,d}^{*'} \left[\tilde{\Sigma}^{-1} \otimes \mathbf{I}_{T-p^*} \right] \hat{\mathbf{u}}$. It is shown in the supporting information appendix that $\hat{\Omega}_{\boldsymbol{\beta}}$ is a consistent estimate of $\Omega_{\boldsymbol{\beta}}$ under the conditions given in Assumption 1.

Based on the heteroskedasticity-robust estimate, $\hat{\Omega}_{\boldsymbol{\beta}}$, it is then straightforward to construct a test statistic for the joint hypothesis $H_0 : \boldsymbol{\theta} = \mathbf{0}$ using the LM testing principle. Specifically, we can form a heteroskedasticity-robust LM-type test which rejects $H_0 : \boldsymbol{\theta} = \mathbf{0}$ for large values of the statistic

$$LM_d^{FGLS} := T \left[\mathbf{R} \hat{\boldsymbol{\beta}} \right]' \left[\mathbf{R} \hat{\Omega}_{\boldsymbol{\beta}} \mathbf{R}' \right]^{-1} \left[\mathbf{R} \hat{\boldsymbol{\beta}} \right] \quad (8)$$

where $\mathbf{R} = \{r_{ij}\}$ is a $k \times kk'$ indicator matrix taking a value equal to one when $j = (i-1)k' + 1, i = 1, \dots, k$, and zero otherwise. In Theorem 2, we next derive the large sample behaviour of LM_d^{FGLS} under both the null hypothesis, $H_0 : \boldsymbol{\theta} = \mathbf{0}$, and under the sequence of local alternatives $H_c : \boldsymbol{\theta} = \mathbf{c}/\sqrt{T}$, where $\mathbf{c} \equiv (c_1, \dots, c_k)'$ is a k vector of finite constants (Pitman drifts) at least one of which is non-zero.

³Some numerical experiments comparing the finite sample properties of the equation-by-equation OLS estimate and the FGLS estimate in (7) of $\boldsymbol{\beta}$ are given in the supporting information appendix. These clearly demonstrate the efficiency gains that can be obtained by FGLS over OLS.

Theorem 2. Let $\mathbf{y}_t$ be generated according to (1), and let Assumption 1 hold. Let LM_d^{FGLS} be as defined in (8). Then, (i) under the null hypothesis, $H_0 : \boldsymbol{\theta} = \mathbf{0}$, $LM_d^{FGLS} \Rightarrow \chi^2_{(k)}$, and (ii) under the sequence of local alternatives, $H_c : \boldsymbol{\theta} = \mathbf{c}/\sqrt{T}$, with at least one element of $\mathbf{c}$ non-zero, $LM_d^{FGLS} \Rightarrow \chi^2_{(k,\xi)}$, where $\chi^2_{(k)}$ and $\chi^2_{(k,\xi)}$ denote a standard χ^2 distribution with k degrees of freedom, and a non-central χ^2 distribution with k degrees of freedom and non-centrality parameter $\xi := (\mathbf{L}'^{-1}\mathbf{c})'(\mathbf{L}'^{-1}\mathbf{c})$, respectively, with $\mathbf{L}$ denoting an upper triangular matrix such that $\mathbf{L}'\mathbf{L} = \mathbf{R}\boldsymbol{\Omega}_\beta\mathbf{R}'$.

Remark 6. The result in Part (i) of Theorem 2 shows that the limiting null distribution of LM_d^{FGLS} is pivotal and that a test of $H_0 : \boldsymbol{\theta} = \mathbf{0}$ can be run using standard critical values from the $\chi^2_{(k)}$ distribution, where k is the dimension of $\mathbf{y}_t$. Part (ii) of Theorem 2 establishes that the asymptotic distribution of LM_d^{FGLS} displays a non-trivial positive offset under the local alternative, $H_c : \boldsymbol{\theta} = \mathbf{c}/\sqrt{T}$, vis-à-vis the null, H_0 , but that its asymptotic local power function will, in general, depend on nuisance parameters arising from any weak dependence or cross sectional correlation present in $\boldsymbol{\varepsilon}_t$. The same is also true of the extant multivariate fractional integration tests discussed in Section 1, except that these do not, in general, have pivotal limiting null distributions when conditional heteroskedasticity is present in $\mathbf{e}_t$, as a consequence of the fact that they are not based around a heteroskedasticity-robust estimate of the variance matrix $\boldsymbol{\Omega}_\beta$.

Remark 7. Theorem 2 provides a theoretical basis for the construction of confidence sets. This can be achieved by inverting the non-rejection region of the test statistic; see Hassler et al. (2009). More specifically, let LM_g denote the value of the LM statistic when testing $H_0 : \boldsymbol{\theta} = \mathbf{0}$ for an arbitrary $\mathbf{g} \in \mathbb{R}^k$, and let Ψ be an arbitrary compact set in $\mathbb{R}^k$. Define $D_\lambda := \left\{ \mathbf{g} \in \Psi : \Pr \left[\chi^2_{(k)} > LM_g \right] \leq 1 - \lambda \right\}$ with $\lambda \in (0, 1)$, that is, the subset of Ψ for which H_0 cannot be rejected at the λ significance level. From Theorem 2, it follows that if Ψ is large enough so as to contain the true values of the long memory parameter vector, then the probability of the true order of integration lying within D_λ is at least $(1 - \lambda)$.

Remark 8. Our proposed test procedure can be generalised to account for non-zero means following the approach in Robinson (1994). To that end, consider the extended form of the DGP in (1) given by $\mathbf{y}_t = \boldsymbol{\mu} + \Delta(L)^{-d-\boldsymbol{\theta}}\boldsymbol{\varepsilon}_t I(t \geq 1)$, where $\boldsymbol{\mu} \equiv (\mu_1, \dots, \mu_k)'$ is a fixed vector. Under $H_0 : \boldsymbol{\theta} = \mathbf{0}$, $(1-L)_+^{d_i} y_{it} = (1-L)_+^{d_i} \mu_i + \varepsilon_t I(t \geq 1)$, $1 \leq i \leq k$. Following Robinson (1994), we regress the differences $(1-L)_+^{d_i} y_{it} := \sum_{j=0}^{t-1} \lambda_j(d_i) y_{it-j}$ on $h_{t,d_i} := \sum_{j=0}^{t-1} \lambda_j(d_i)$, $t = 2, \dots, T$, with $\{\lambda_j(d_i)\}$ as defined in (4). Denote the resulting estimates as $\tilde{\mu}_i$, $i = 1, \dots, k$, and the corresponding residuals as $\tilde{\varepsilon}_{it,d_i} := (1-L)_+^{d_i} y_{it} - \tilde{\mu}_i h_{t,d_i}$. One then simply redefines the i th element of the vector $\boldsymbol{\varepsilon}_{t,d}$ from (2) to be $\tilde{\varepsilon}_{it,d_i}$, $i = 1, \dots, k$, and then proceeds as before. Let $\tilde{\boldsymbol{\beta}}$ denote the FGLS estimator obtained in this way. Then, following the approach taken in Proposition 4 of Demetrescu et al. (2008), it can be shown that Theorem 1 holds with $\hat{\boldsymbol{\beta}}$ replaced by $\tilde{\boldsymbol{\beta}}$ because $\|\tilde{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}\| = o_p(T^{-1/2})$ under the restrictions considered and the additional condition that $d > 0$. More generally, the results can be extended to account for, among other things, deterministic polynomial time trends and deterministic seasonal effects; see also Nielsen (2005) and Demetrescu et al. (2008). The large sample results given in this section are not affected by accounting for such deterministics.⁴

Remark 9. In practical applications of the tests, the lag order p will typically be unknown and so could be selected using a standard consistent information criterion such as the Bayes information criterion (BIC). Demetrescu et al. (2008) argue that these can lead to substantial finite-sample biases in the context of the tests considered here. As an alternative, Demetrescu et al. (2008) advocate the use of a deterministic lag selection rule, such as the popular Schwert (1989) rule, which sets $p = \lfloor K(T/100)^{1/4} \rfloor$, where $\lfloor \cdot \rfloor$ denotes the integer part of its argument and K is a finite positive constant. Provided that the true lag order p is finite, as we assume in this paper, then the limiting distribution theory given in this section will remain apposite for tests based on a lag length determined according to such deterministic rules. We will implement Schwert's rule in the empirical application considered in Section 5.

Remark 10. It is useful to compare the large sample properties of our proposed test with the Gaussian LM test of Nielsen (2005) in comparable settings. To this end, consider the case where, as required by the conditions imposed in Theorem 3 of Nielsen (2005, p. 381), Assumption 1 is restricted such that $p = 0$ and $\mathbf{e}_t$ is an i.i.d. innovation sequence.

⁴Numerical experiments investigating the finite sample rejection rates of our tests when a non-zero mean is allowed for are given in section B.3 of the supporting information appendix. These confirm the (exact) invariance of such tests and the lack of invariance of tests which do not allow for a non-zero mean. The loss of finite sample power from allowing for a non-zero mean appears very modest.

It is straightforward to show that under these additional restrictions

$$LM_d^{FGLS} = \left(X_{-1,d}' \left[\tilde{\Sigma}^{-1} \otimes I_{T-1} \right] Y_d \right)' \left(X_{-1,d}' \left[\tilde{\Sigma}^{-1} \otimes I_{T-1} \right] X_{-1,d}^* \right)^{-1} \left(X_{-1,d}' \left[\tilde{\Sigma}^{-1} \otimes I_{T-1} \right] Y_d \right) + o_p(1)$$

where $X_{-1,d}^* := \text{diag} \left\{ X_{1,-1,d}^*, \dots, X_{k,-1,d}^* \right\}$ in which $X_{i,-1,d}^* = Z_{-1,d}^* a_i$, where $Z_{-1,d}^* := (z_{1,d}^*, \dots, z_{T-1,d}^*)'$ and a_i denotes the i -th unit k -dimensional vector. Noting, moreover, that $z_{t-1,d}^* := \sum_{j=1}^{t-1} j^{-1} e_{t-j,d} = -\ln(1-L)_+ e_{t-j,d}$ under the null hypothesis, the vector $X_{-1,d}' \left[\tilde{\Sigma}^{-1} \otimes I_{T-1} \right] Y_d$ corresponds to the Gaussian score vector $S_T := J_k' \text{vec} \left(\tilde{\Sigma}^{-1} S_{10}' \right)$ given in eq. (11) of Nielsen (2005), where $S_{10} := \sum_{t=2}^T e_{t-1}^* e_t'$ with $e_{t-1}^* := \sum_{j=1}^{t-1} j^{-1} e_{t-j}$, and $J_k := (\text{vec}(A_{11}), \dots, \text{vec}(A_{kk}))$ with $A_{ii} := a_i a_i'$. Because $\frac{1}{T} X_{-1,d}' \left[\tilde{\Sigma}^{-1} \otimes I_{T-1} \right] X_{-1,d}^* \xrightarrow{p} A_\beta$, where $A_\beta = \frac{\pi^2}{6} \Sigma \otimes \Sigma^{-1}$ under the additional restrictions outlined above, it can be seen that LM_d^{FGLS} is asymptotically equivalent to the Gaussian LM test proposed in Nielsen (2005). Consequently, LM_d^{FGLS} is asymptotically locally efficient when e_t is a Gaussian i.i.d. sequence; see Nielsen (2004a). Where $p > 0$, the two tests differ crucially on how the short-run autocorrelation is handled. While LM_d^{FGLS} uses p th-order augmentation in (5), Nielsen's (2005) test relies on prewhitening using the residuals from a VAR(p) model in a two-stage procedure. Augmentation and prewhitening are asymptotically equivalent strategies but will differ in finite samples, as will be explored in the next section.

4 | MONTE CARLO SIMULATIONS

We consider the simulation DGP

$$\begin{bmatrix} (1-L)^{1+\theta_1} & 0 \\ 0 & (1-L)^{1+\theta_2} \end{bmatrix} y_t = \varepsilon_t I(t \geq 1), \quad t = 1, \dots, T, \quad (9)$$

where $y_t \equiv (y_{1t}, y_{2t})'$, $\Pi(L)\varepsilon_t = e_t$ with $\Pi(L) = \text{diag}\{1 - \pi_1 L, 1 - \pi_2 L\}$, and $(\pi_1, \pi_2) \in \{(0, 0), (0.4, 0.4)\}$; such that the former corresponds to white noise, while the latter yields weakly stationary VAR(1) errors. As the particular values of the long memory coefficients play no role in our context, we set $d_1 = d_2 = 1$. We report results for $T \in \{500, 1000\}$.

The innovations $\{e_t\}$ are generated to exhibit time-varying conditional second-order moments according to the design:

$$e_t = \begin{bmatrix} \sigma_{1t} & 0 \\ 0 & \sigma_{2t} \end{bmatrix} \eta_t; \quad E(\eta_t) = 0; \quad E(\eta_t \eta_t') =: \Omega_\rho = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

where $\eta_t := (\eta_{1t}, \eta_{2t})'$ is an i.i.d. vector drawn from either a multivariate Gaussian distribution or a (heavy-tailed) multivariate Student- t distribution with 5 degrees of freedom. The covariance matrix Ω_ρ depends on the contemporaneous correlation coefficient ρ , $\rho \in \{0, 0.2, 0.4, 0.6, 0.8\}$. The conditional variances $\{\sigma_{it}^2\}$ are driven by (normalised) stationary GARCH(1,1) processes $\sigma_{it}^2 = (1 - \alpha - \beta) + \alpha e_{it-1}^2 + \beta \sigma_{it-1}^2$, $i = 1, 2$ with $\alpha, \beta \geq 0$ and $\alpha + \beta < 1$, such that $E(e_{it}^2) = 1$. We consider $(\alpha, \beta) \in \{(0, 0), (0.1, 0.5), (0.1, 0.7), (0.1, 0.8), (0.1, 0.85)\}$. The case $\alpha = \beta = 0$ corresponds to conditional homoskedasticity.

To simplify our discussion, we fix $\theta_2 = 0$ in all of the reported simulations and vary θ_1 among $\{-0.3, -0.25, \dots, 0, \dots, 0.25, 0.3\}$. Consequently, while the true order of integration of $\{y_{2t}\}$ is always one, the true order of integration of $\{y_{1t}\}$ is $1 + \theta_1$. The case where $\theta_1 = 0$ allows us to investigate the empirical size properties of LM_d^{FGLS} , while the cases where $\theta_1 \neq 0$ allow us to investigate its finite sample power against an alternative where one of the long memory parameters deviates from the null hypothesis. For each of the parameter configurations $(\alpha, \beta, \rho, \pi_1, \pi_2, \theta_1)$, the two sample lengths, and the two conditional distributions, we compute LM_d^{FGLS} and determine the empirical rejection frequencies [ERFs] at the 5% nominal (asymptotic) level over 5000 replications.

We also benchmark the performance of LM_d^{FGLS} against two alternative (but related) tests. The first is the multivariate LM test of Nielsen (2004a, 2005) discussed in Remark 10 above, denoted LM_d^{MLE} in what follows, and the second is the multivariate trace test of Breitung and Hassler (2002), which we denote BH_d . While LM_d^{FGLS} corrects for stationary serial correlation in ε_t via lag augmentation in (5), both LM_d^{MLE} and BH_d use a prewhitening approach. Both LM_d^{MLE} and BH_d require that $\{e_t\}$ is i.i.d., and so neither allows for the presence of conditional heteroskedasticity in $\{e_t\}$. Under these conditions, LM_d^{MLE} has a limiting $\chi_{(k)}^2$ null distribution, while BH_d has a limiting $\chi_{(k^2)}^2$ null distribution. Nielsen's LM_d^{MLE} is

designed to test the same null hypotheses as LM_d^{FGLS} and so is the most natural candidate to benchmark our test against. In contrast, the BH_d test is for the null hypothesis of a common order of integration between the elements of the vector time series. Our simulation DGP is such that this condition holds under the null hypothesis, but not under the alternative, so a comparison with this test is appropriate.

4.1 | ERFs with no augmentation/prewhitening

Table 1 reports the empirical size properties ($\theta_1 = \theta_2 = 0$), for LM_d^{FGLS} , LM_d^{MLE} and BH_d where no short-run dynamics are present ($\pi_1 = \pi_2 = 0$), and where, accordingly, no lag augmentation or prewhitening is needed. This allows us to first investigate the impact of GARCH effects, contemporaneous correlations, and the conditional distribution of the innovations on each test.

The results show that LM_d^{FGLS} displays ERFs close to the nominal asymptotic 5% level in almost all cases. Some mild oversizing is seen for the smaller sample size considered when the innovations are conditionally Student- t distributed with relatively high GARCH persistence, $\alpha = 0.1$ and $\beta \geq 0.80$, and significant levels of endogeneity, $\rho \geq 0.4$. These distortions are largely ameliorated as the sample size increases. Where the innovations are i.i.d. ($\alpha = \beta = 0$), both LM_d^{MLE} and BH_d display good finite sample size control regardless of the conditional distribution or the degree of endogeneity. However, where the innovations exhibit conditional heteroskedasticity, a very different pattern emerges for both LM_d^{MLE} and BH_d . These tests display a tendency to strong oversizing, with these distortions being larger (other things equal): The stronger the degree of persistence of the GARCH process, the larger the degree of endogenous correlation $|\rho|$, and for innovations drawn from a heavy-tailed distribution. Moreover, these size distortions are not ameliorated as the sample size increases. To illustrate, for $T = 500$ and $\rho = 0.8$, the ERFs of LM_d^{MLE} and BH_d with GARCH errors driven by $(\alpha, \beta) = (0.10, 0.85)$ and Student- t innovations are 36% and 38.8%, respectively. In contrast, the LM_d^{FGLS} test is only slightly oversized at 7.1%. For $T = 1000$, the corresponding ERFs of LM_d^{MLE} and BH_d increase significantly to 48.3% and 52.8%, respectively, while that of LM_d^{FGLS} reduces to 6.1%.

4.2 | ERFs with augmentation/prewhitening

We now analyse the finite sample size and power properties of LM_d^{FGLS} , LM_d^{MLE} and BH_d in the case where the errors, ε_t , can display first-order stationary VAR dynamics. Accordingly, we set $p = 1$ in (5) in relation to the LM_d^{FGLS} test, while analogously, we use a VAR(1) for prewhitening in connection with the LM_d^{MLE} and BH_d tests. For ε_t we consider (i) $\pi_1 = \pi_2 = 0$, so that augmentation/prewhitening is in fact unnecessary, and (ii) $\pi_1 = \pi_2 = 0.4$, so that the correct order of augmentation/prewhitening is employed.

Table 2 reports ERFs of the three tests in the Gaussian homoskedastic case ($\alpha = \beta = 0$). Results for the Student- t case are not reported as these are almost identical to the results reported in Table 2. Also, to keep the size of the subsequent tables to manageable proportions, we will only report results for two values of the correlation coefficient, namely, $\rho = 0$ and $\rho = 0.8$. Corresponding results for $\rho \in \{0.2, 0.4, 0.6\}$ can be obtained on request.

Consider first the results for the case where $\theta_1 = 0$ so that the null hypothesis holds. Here, we see that the ERFs of the augmented LM_d^{FGLS} test lie close to the nominal asymptotic level throughout, even where the lag augmentation is unnecessary. Prewhitening also appears to be effective for the LM_d^{MLE} and BH_d tests, with the exception of the case where $\rho = 0.8$ where these tests are somewhat oversized for $T = 1000$. Turning next to the empirical power results for $\theta_1 \neq 0$, we see that LM_d^{FGLS} displays good finite sample power properties with power increasing, other things equal, both as $|\theta_1|$ increases and as T increases, as would be expected. Power is also larger, other things equal, for $\rho = 0.8$ than for $\rho = 0$, illustrating the efficiency benefits gained from multivariate modelling when the variables are cross-correlated. In terms of a comparison between the three tests, overall the finite sample power properties of LM_d^{FGLS} and LM_d^{MLE} seen in Table 2 are very similar for alternatives where $\theta_1 < 0$, as might be expected given the asymptotic equivalence of these tests when the innovations are i.i.d.; compare Remark 10. For alternatives where $\theta_1 > 0$ (i.e., when the process is more persistent than posited under the null) LM_d^{FGLS} can display somewhat higher power than LM_d^{MLE} , particularly in the case where the errors are first-order autocorrelated, $\pi_1 = \pi_2 = 0.4$; for example, for $\pi_1 = \pi_2 = 0.4$, $\rho = 0$, $T = 500$ and $\theta_1 = 0.3$ the power of LM_d^{FGLS} and LM_d^{MLE} are 57.4% and 32.3%, respectively. These differences are likely attributable to the use of lag augmentation rather than prewhitening in the construction of LM_d^{FGLS} . Both LM_d^{FGLS} and LM_d^{MLE} clearly dominate BH_d on power; in the previous example, the power of BH_d is only 26.3%. The power functions of all of the tests are asymmetric in the sign of θ_1 , for a given DGP, such that a false null hypothesis which leads to an over-differenced series ($\theta_1 < 0$) is seen to be more easily rejected than an incorrect null, which leads to an under-differenced series ($\theta_1 > 0$) where the magnitude of the under/over difference is the same. To illustrate, for $\pi_1 = \pi_2 = 0.4$, $\rho = 0$ and $T = 500$, the power of LM_d^{FGLS} to detect

TABLE 1 Empirical rejection frequencies (empirical sizes) for $\theta_1 = \theta_2 = 0$ at the 5% nominal asymptotic significance level of the LM_d^{FGLS} , LM_d^{MLE} and BH_d tests, for different values for the contemporaneous correlation parameter ρ , the GARCH parameters (α, β) and sample lengths T

GARCH	ρ	Gaussian innovations				Student- t innovations			
		$T = 500$		$T = 1000$		$T = 500$		$T = 1000$	
		LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}
$\alpha = 0, \beta = 0$	0.0	0.051	0.045	0.050	0.048	0.048	0.048	0.053	0.049
	0.2	0.050	0.047	0.050	0.047	0.048	0.048	0.056	0.050
	0.4	0.053	0.050	0.050	0.044	0.045	0.048	0.056	0.050
	0.6	0.055	0.051	0.050	0.047	0.045	0.048	0.055	0.052
	0.8	0.055	0.050	0.050	0.048	0.049	0.048	0.056	0.047
$\alpha = 0.1, \beta = 0.5$	0.0	0.054	0.071	0.065	0.049	0.070	0.067	0.060	0.180
	0.2	0.053	0.072	0.065	0.048	0.068	0.066	0.060	0.179
	0.4	0.054	0.072	0.064	0.049	0.068	0.066	0.061	0.181
	0.6	0.055	0.074	0.065	0.049	0.071	0.067	0.059	0.189
	0.8	0.056	0.075	0.066	0.052	0.075	0.067	0.064	0.199
$\alpha = 0.1, \beta = 0.7$	0.0	0.053	0.080	0.070	0.048	0.080	0.071	0.060	0.219
	0.2	0.054	0.080	0.071	0.047	0.077	0.072	0.060	0.218
	0.4	0.054	0.080	0.070	0.048	0.079	0.073	0.058	0.221
	0.6	0.055	0.080	0.073	0.050	0.081	0.075	0.059	0.233
	0.8	0.056	0.080	0.074	0.054	0.084	0.077	0.063	0.242
$\alpha = 0.1, \beta = 0.8$	0.0	0.052	0.088	0.076	0.048	0.092	0.081	0.062	0.269
	0.2	0.052	0.088	0.079	0.047	0.088	0.079	0.059	0.275
	0.4	0.052	0.089	0.079	0.045	0.089	0.080	0.061	0.281
	0.6	0.056	0.091	0.082	0.049	0.095	0.085	0.065	0.296
	0.8	0.055	0.093	0.086	0.051	0.097	0.089	0.069	0.306
$\alpha = 0.1, \beta = 0.85$	0	0.053	0.103	0.089	0.050	0.113	0.093	0.064	0.337
	0.2	0.050	0.104	0.089	0.049	0.114	0.094	0.063	0.346
	0.4	0.054	0.103	0.090	0.048	0.117	0.096	0.066	0.350
	0.6	0.056	0.107	0.096	0.049	0.124	0.104	0.068	0.353
	0.8	0.058	0.114	0.105	0.052	0.121	0.111	0.071	0.360

Note: The innovations are drawn from either a multivariate normal distribution or a multivariate Student- t distribution with 5 degrees of freedom.

TABLE 2 Empirical rejection frequencies at the 5% nominal asymptotic level for $\theta_2 = 0$ and the values $\theta_1 = 0$ (empirical size) and $|\theta_1| > 0$ (empirical power) of the LM_d^{FGLS} , LM_d^{MLE} and BH_d tests for the correlation coefficient ρ and sample length T

θ_1	ρ	$\pi_1 = \pi_2 = 0$						$\pi_1 = \pi_2 = 0.4$					
		$T = 500$			$T = 1000$			$T = 500$			$T = 1000$		
		LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d
-0.30	0.0	1.000	0.999	0.994	1.000	1.000	1.000	0.875	0.876	0.777	0.997	0.998	0.990
-0.25	0.0	0.987	0.982	0.946	1.000	1.000	0.999	0.720	0.718	0.591	0.967	0.972	0.932
-0.20	0.0	0.899	0.879	0.781	0.999	0.999	0.990	0.496	0.486	0.374	0.835	0.855	0.748
-0.15	0.0	0.644	0.626	0.493	0.941	0.934	0.868	0.289	0.288	0.214	0.558	0.590	0.465
-0.10	0.0	0.338	0.321	0.239	0.609	0.601	0.478	0.154	0.147	0.123	0.271	0.284	0.213
-0.05	0.0	0.118	0.112	0.091	0.183	0.179	0.134	0.074	0.070	0.067	0.099	0.098	0.085
0	0.0	0.057	0.051	0.054	0.052	0.050	0.049	0.053	0.045	0.053	0.050	0.048	0.046
0.05	0.0	0.121	0.112	0.091	0.179	0.174	0.142	0.080	0.079	0.069	0.100	0.097	0.083
0.10	0.0	0.321	0.311	0.243	0.584	0.569	0.472	0.145	0.143	0.119	0.267	0.270	0.205
0.15	0.0	0.600	0.581	0.498	0.888	0.876	0.814	0.272	0.254	0.198	0.479	0.454	0.360
0.20	0.0	0.812	0.788	0.701	0.983	0.978	0.957	0.392	0.321	0.265	0.674	0.597	0.507
0.25	0.0	0.926	0.895	0.831	0.998	0.996	0.988	0.496	0.358	0.296	0.793	0.652	0.576
0.30	0.0	0.969	0.942	0.902	1.000	0.999	0.997	0.574	0.323	0.263	0.859	0.613	0.546
-0.30	0.8	1.000	1.000	1.000	1.000	1.000	1.000	0.998	0.999	0.997	1.000	1.000	1.000
-0.25	0.8	1.000	1.000	1.000	1.000	1.000	1.000	0.985	0.990	0.969	1.000	1.000	1.000
-0.20	0.8	1.000	1.000	0.999	1.000	1.000	1.000	0.908	0.925	0.848	1.000	1.000	1.000
-0.15	0.8	0.983	0.981	0.945	1.000	1.000	0.999	0.678	0.705	0.572	1.000	1.000	0.999
-0.10	0.8	0.757	0.738	0.621	0.920	0.954	0.909	0.350	0.360	0.271	0.920	0.954	0.909
-0.05	0.8	0.243	0.234	0.179	0.369	0.463	0.369	0.124	0.124	0.098	0.369	0.463	0.369
0	0.8	0.058	0.050	0.054	0.053	0.088	0.080	0.057	0.056	0.053	0.053	0.088	0.080
0.05	0.8	0.245	0.233	0.180	0.372	0.448	0.370	0.117	0.125	0.100	0.372	0.448	0.370
0.10	0.8	0.712	0.700	0.596	0.906	0.929	0.884	0.324	0.339	0.253	0.906	0.929	0.884
0.15	0.8	0.951	0.943	0.904	0.996	0.997	0.995	0.567	0.559	0.464	0.996	0.997	0.995
0.20	0.8	0.994	0.993	0.982	1.000	1.000	1.000	0.743	0.692	0.594	1.000	1.000	1.000
0.25	0.8	0.999	0.998	0.997	1.000	1.000	1.000	0.845	0.774	0.691	1.000	1.000	1.000
0.30	0.8	1.000	1.000	0.999	1.000	1.000	1.000	0.898	0.756	0.679	1.000	1.000	1.000

Note: The short-run errors in the DGP obey VAR(1) dynamics with on-diagonal coefficients π_1 and π_2 and off-diagonal coefficients $\pi_{12} = \pi_{21} = 0$. LM_d^{FGLS} statistic is computed from an augmented auxiliary regression with one lag of the dependent variable. LM_d^{MLE} and BH_d are computed from VAR(1) residuals. The innovations are drawn from a multivariate Gaussian distribution.

$\theta_1 = 0.25$ and $\theta_1 = -0.25$ is 49.6% and 72.0%, respectively. Breitung and Hassler (2002) report a similar asymmetry in the power properties of their univariate tests.

Finally, we turn to the case where the innovations may display GARCH effects and excess kurtosis. Table 3 ($T = 500$) and Table 4 ($T = 1000$) report the ERFs for LM_d^{FGLS} , LM_d^{MLE} and BH_d for both Gaussian and Student- t innovations for $\pi_1 = \pi_2 = 0.4$, $\rho \in \{0, 0.8\}$, $(\alpha, \beta) \in \{(0.10, 0.80), (0.10, 0.85)\}$. The results for $\theta_1 = 0$ show that the empirical size properties of the tests in the presence of GARCH are similar to the corresponding results reported previously for the serially uncorrelated case with no augmentation/prewhitening in Table 1. In particular, while the empirical size of LM_d^{FGLS} is reasonably close to the nominal asymptotic 5% level throughout (size departures are not greater than 1.6% for $T = 500$ and not greater than 0.8% for $T = 1000$), incorrectly assuming conditional homoskedasticity causes significant oversizing in both LM_d^{MLE} and BH_d which is not ameliorated by increasing the sample size. To illustrate, for $\rho = 0.8$, and $(\alpha, \beta) = (0.10, 0.85)$, LM_d^{MLE} and BH_d , respectively, display ERFs of 9.2% and 8.7% for $T = 500$ and 10.4% and 9.9% for $T = 1000$ with Gaussian innovations, increasing to 28.2% and 32.4% for $T = 500$ and 40.4% and 46.2% for $T = 1000$ with Student- t innovations.

For non-zero values of θ_1 , we observe qualitatively similar patterns in relation to the power properties of LM_d^{FGLS} as were reported in Table 2 in the homoskedastic case, albeit persistent GARCH-type behaviour in the innovations can be seen to clearly lower the finite sample power of LM_d^{FGLS} relative to the i.i.d. case, and particularly so when the conditional distribution of the innovations is heavy tailed. This is of course consistent with Theorem 2 where it was shown that the asymptotic local power function of the LM_d^{FGLS} test depends on any nuisance parameters arising from conditional heteroskedasticity in the innovations. To illustrate, from Table 2 for $\pi_1 = \pi_2 = 0.4$, $\rho = 0$, $T = 500$, the power of LM_d^{FGLS} to detect $\theta_1 = 0.3$ ($\theta_1 = -0.3$) in the i.i.d. case is 57.4% (87.5%). However, from Table B.1, under GARCH dependence with $(\alpha, \beta) = (0.10, 0.85)$ the respective probabilities are 52.7% (75.7%) in the Gaussian case, and 37.3% (49.3%) in the Student- t case. Similarly, for $T = 1000$ in the previous example power is seen from Table 3 to be 98.7% (99.9%) in the Gaussian case and 71.6% (81.1%) in the Student- t case. A comparison between the finite sample power of LM_d^{FGLS} and that of LM_d^{MLE} and BH_d is somewhat uninformative here because of the poor size control of the latter two tests under conditional heteroskedasticity.

5 | LONG-RUN DYNAMICS IN VOLUME AND VOLATILITY

Understanding the linkages between return volatility, liquidity and trading activity has been an area of considerable research interest in the finance literature. We apply the multivariate testing approach developed in this paper to perform joint inference on the order of fractional integration of trading volume and return volatility for a sample of major stocks traded in the US market. As part of this, we also investigate the hypothesis that these variables exhibit the same order of fractional integration.

A number of previous studies have investigated this hypothesis in trading volume and return volatility within a multivariate ARFIMA framework. No strong consensus has emerged across these studies, which are based on a variety of methods of estimation and inference and employ a number of different observable variables to proxy the latent return volatility process. Bollerslev and Jubinski (1999) and Lobato and Velasco (2000) use semiparametric multivariate periodogram-based estimators in the frequency domain, proxying return volatility by absolute-valued returns. They conclude that, for most of the stocks analysed, the hypothesis that trading volume and return volatility share the same order of fractional integration cannot be rejected. However, Fleming and Kirby (2011) argue that the slow rate of convergence of periodogram-based estimators raises concerns about estimation efficiency. Consequently, they implement a parametric Gaussian quasi-maximum likelihood (QML) approach as in Nielsen (2004a) to estimate a bivariate FIVAR model, allowing for short-run dependencies, but under the assumption of conditional homoskedasticity. Moreover, Fleming and Kirby (2011) proxy return volatility using intra-day data with the aim of improving accuracy over the use of absolute-valued returns and reject the hypothesis of a common long memory coefficient in most cases.

Our testing procedure is expected to be useful here for two key reasons. First, as shown in Theorem 1, the FGLS-based test achieves the usual $\sqrt{T}$ rate of convergence in parametric testing and is therefore expected to yield improved finite-sample power performance relative to periodogram-based estimators; see, for example, Tanaka (1999). This consideration addresses concerns surrounding efficiency raised by Fleming and Kirby (2011). Second, and arguably most importantly, our testing approach is valid in the presence of stationary conditionally time-varying second-order moments and heavy-tailed innovations, unlike the QML approach of Nielsen (2004a) used by Fleming and Kirby (2011).

TABLE 3 Empirical rejection frequencies at the 5% nominal asymptotic level for $\theta_2 = 0$ and $\theta_1 = 0$ (empirical size) and $|\theta_1| > 0$ (empirical power) of the LM_d^{FGLS} , LM_d^{MLE} and BH_d tests for the correlation coefficient ρ

$T = 500$		GARCH: $\alpha = 0.1, \beta = 0.8$						GARCH: $\alpha = 0.1, \beta = 0.85$					
θ_1	ρ	Gaussian			Student-t			Gaussian			Student-t		
		LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d
-0.30	0.0	0.795	0.859	0.766	0.568	0.840	0.776	0.757	0.855	0.768	0.493	0.836	0.784
-0.25	0.0	0.625	0.718	0.599	0.419	0.724	0.644	0.591	0.716	0.612	0.365	0.730	0.673
-0.20	0.0	0.419	0.512	0.398	0.295	0.579	0.508	0.398	0.521	0.417	0.259	0.615	0.557
-0.15	0.0	0.250	0.328	0.251	0.192	0.434	0.384	0.240	0.341	0.268	0.165	0.490	0.450
-0.10	0.0	0.142	0.182	0.152	0.124	0.310	0.282	0.136	0.199	0.162	0.113	0.385	0.362
-0.05	0.0	0.076	0.102	0.084	0.073	0.218	0.207	0.074	0.121	0.098	0.074	0.294	0.291
0	0.0	0.051	0.072	0.066	0.056	0.186	0.180	0.048	0.083	0.076	0.057	0.252	0.257
0.05	0.0	0.074	0.102	0.085	0.074	0.202	0.192	0.074	0.117	0.097	0.070	0.265	0.272
0.10	0.0	0.136	0.164	0.127	0.118	0.241	0.232	0.131	0.173	0.137	0.105	0.306	0.310
0.15	0.0	0.251	0.266	0.213	0.192	0.321	0.291	0.241	0.277	0.222	0.162	0.360	0.351
0.20	0.0	0.358	0.334	0.278	0.284	0.391	0.350	0.345	0.344	0.281	0.242	0.420	0.409
0.25	0.0	0.461	0.377	0.312	0.377	0.432	0.395	0.449	0.381	0.318	0.310	0.456	0.436
0.30	0.0	0.543	0.346	0.284	0.463	0.418	0.388	0.527	0.353	0.294	0.373	0.447	0.430
-0.30	0.8	0.992	0.998	0.994	0.839	0.977	0.971	0.984	0.997	0.992	0.744	0.961	0.957
-0.25	0.8	0.954	0.981	0.961	0.729	0.931	0.916	0.928	0.973	0.950	0.614	0.902	0.896
-0.20	0.8	0.836	0.901	0.839	0.576	0.837	0.804	0.794	0.892	0.831	0.466	0.810	0.786
-0.15	0.8	0.588	0.693	0.583	0.371	0.658	0.610	0.545	0.686	0.587	0.294	0.642	0.632
-0.10	0.8	0.300	0.378	0.295	0.212	0.461	0.426	0.279	0.390	0.307	0.172	0.478	0.493
-0.05	0.8	0.121	0.159	0.128	0.096	0.271	0.276	0.112	0.176	0.150	0.086	0.325	0.367
0	0.8	0.054	0.077	0.0670	0.061	0.202	0.221	0.055	0.092	0.087	0.066	0.282	0.324
0.05	0.8	0.103	0.152	0.120	0.102	0.249	0.256	0.098	0.165	0.137	0.085	0.302	0.339
0.10	0.8	0.296	0.351	0.277	0.214	0.377	0.371	0.278	0.355	0.288	0.169	0.410	0.424
0.15	0.8	0.519	0.550	0.464	0.363	0.549	0.513	0.492	0.546	0.469	0.282	0.544	0.542
0.20	0.8	0.696	0.689	0.604	0.521	0.658	0.619	0.662	0.680	0.604	0.410	0.615	0.608
0.25	0.8	0.809	0.763	0.689	0.640	0.714	0.694	0.775	0.751	0.682	0.514	0.677	0.682
0.30	0.8	0.870	0.735	0.676	0.719	0.705	0.703	0.846	0.725	0.674	0.586	0.660	0.687

Note: The short-run errors follow a VAR(1) model with $\pi_1 = \pi_2 = 0.4$ and GARCH innovations with parameters α and β . The LM_d^{FGLS} test statistic is computed from an augmented auxiliary regression with one lag of the dependent variable. LM_d^{MLE} and BH_d are computed from VAR(1) residuals. The sample length is $T = 500$.

TABLE 4 Empirical rejection frequencies at the 5% nominal asymptotic level for $\theta_2 = 0$ and $\theta_1 = 0$ (empirical size) and $|\theta_1| > 0$ (empirical power) of the LM_d^{FGLS} , LM_d^{MLE} and BH_d tests for the correlation coefficient ρ

$T = 1000$		GARCH: $\alpha = 0.1, \beta = 0.8$						GARCH: $\alpha = 0.1, \beta = 0.85$					
θ_1	ρ	Gaussian			Student-t			Gaussian			Student-t		
		LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d	LM_d^{FGLS}	LM_d^{MLE}	BH_d
-0.30	0.0	0.983	0.993	0.983	0.753	0.961	0.949	0.967	0.991	0.979	0.603	0.947	0.930
-0.25	0.0	0.916	0.959	0.920	0.632	0.915	0.884	0.891	0.956	0.914	0.481	0.904	0.879
-0.20	0.0	0.744	0.841	0.744	0.457	0.811	0.754	0.702	0.836	0.742	0.324	0.817	0.774
-0.15	0.0	0.471	0.597	0.481	0.287	0.645	0.592	0.437	0.603	0.489	0.219	0.688	0.655
-0.10	0.0	0.237	0.316	0.248	0.155	0.457	0.415	0.220	0.332	0.269	0.135	0.550	0.524
-0.05	0.0	0.093	0.132	0.107	0.080	0.302	0.286	0.093	0.150	0.124	0.073	0.446	0.424
0	0.0	0.051	0.072	0.066	0.058	0.243	0.237	0.053	0.089	0.083	0.056	0.397	0.387
0.05	0.0	0.091	0.122	0.098	0.077	0.262	0.247	0.084	0.139	0.112	0.073	0.398	0.391
0.10	0.0	0.246	0.284	0.220	0.176	0.385	0.353	0.226	0.293	0.228	0.131	0.472	0.456
0.15	0.0	0.437	0.461	0.376	0.288	0.508	0.470	0.414	0.464	0.384	0.213	0.558	0.539
0.20	0.0	0.635	0.603	0.523	0.434	0.607	0.557	0.607	0.607	0.531	0.322	0.631	0.596
0.25	0.0	0.756	0.653	0.576	0.588	0.670	0.625	0.729	0.652	0.578	0.433	0.668	0.639
0.30	0.0	0.830	0.620	0.559	0.673	0.664	0.623	0.811	0.621	0.555	0.516	0.659	0.649
-0.30	0.8	1.000	1.000	1.000	0.929	0.998	0.999	1.000	1.000	1.000	0.811	0.991	0.994
-0.25	0.8	1.000	1.000	1.000	0.877	0.999	0.992	0.997	1.000	1.000	0.715	0.974	0.980
-0.20	0.8	0.988	0.998	0.993	0.752	0.965	0.965	0.977	0.995	0.992	0.545	0.926	0.929
-0.15	0.8	0.893	0.951	0.906	0.555	0.876	0.860	0.846	0.940	0.898	0.387	0.833	0.838
-0.10	0.8	0.535	0.658	0.554	0.295	0.670	0.635	0.491	0.655	0.560	0.203	0.665	0.679
-0.05	0.8	0.164	0.228	0.180	0.112	0.386	0.391	0.155	0.254	0.207	0.091	0.488	0.535
0	0.8	0.058	0.083	0.080	0.057	0.271	0.296	0.058	0.104	0.099	0.058	0.404	0.462
0.05	0.8	0.175	0.228	0.180	0.113	0.349	0.352	0.162	0.246	0.196	0.088	0.441	0.487
0.10	0.8	0.510	0.578	0.477	0.307	0.591	0.569	0.471	0.574	0.485	0.206	0.584	0.611
0.15	0.8	0.822	0.848	0.780	0.542	0.771	0.760	0.781	0.839	0.768	0.366	0.711	0.725
0.20	0.8	0.940	0.940	0.908	0.714	0.873	0.860	0.916	0.928	0.897	0.504	0.806	0.816
0.25	0.8	0.977	0.969	0.949	0.819	0.911	0.903	0.964	0.961	0.943	0.632	0.833	0.856
0.30	0.8	0.993	0.958	0.943	0.885	0.900	0.908	0.987	0.946	0.930	0.716	0.840	0.867

Note: The short-run errors follow a VAR(1) model with $\pi_1 = \pi_2 = 0.4$ and GARCH innovations with parameters α and β . The LM_d^{FGLS} test statistic is computed from an augmented auxiliary regression with one lag of the dependent variable. LM_d^{MLE} and BH_d are computed from VAR(1) residuals. The sample length is $T = 1000$.

5.1 | Data

Our analysis focuses on 30 major US stocks from the DJI. We analyse data sampled from 2 January 2003 to 31 December 2014. Unlike trading volume, return volatility cannot be directly observed. The literature has suggested a number of different estimation methods in increasing degree of accuracy, which we implement. The simplest approach uses absolute-valued returns computed from close-to-close daily prices. Unfortunately, this measure is known to be highly inefficient and subject to large estimation errors. More accurate estimates can be constructed building on intra-day information. Following Garman and Klass (1980), we also proxy daily return variability as $u_t^2/2 - (2 \ln 2 - 1)c_t^2$, where u_t and c_t are the differences in the natural logarithms of the high and low and of the closing and opening prices, respectively. Such range-based estimators produce more efficient estimates than absolute-valued returns computed from close-to-close prices (Parkinson, 1980) and, as discussed in Andersen and Bollerslev (1998), can be as efficient a measure of return volatility as realised volatility computed on the basis of 3- to 4-h returns. The last estimator we consider is a realised variance measure computed from aggregating 5-min squared continuously compounded returns over the trading session. Daily share volumes and high, low, opening and closing prices are obtained from CRSP. High-frequency prices necessary to compute realised variances are obtained from the NYSE Trade and Quote (TAQ) database. As is customary in this literature, we implement log transforms in both trading volume and return volatility variables. Standard descriptive statistics for the aforementioned variables as well as a statistical analysis highlighting statistically significant evidence for the presence of time-varying second order moments in the data are presented in Tables C1 and C2 of the supporting information appendix.

5.2 | Implementation issues

In conducting our analysis of the long memory properties of log-trading volume and log-return volatility, hereafter denoted as $(d(vlm), d(\sigma))'$, a number of key implementation issues arise, which we now detail.

First, we construct 99%, 95% and 90% confidence sets for $(d(vlm), d(\sigma))'$ by inverting the non-rejection regions of the multivariate test in a discrete grid search over the support $\Psi = [-0.2, 1.2] \times [-0.2, 1.2]$ (see Remark 7). More specifically, we evaluate LM_d^{FGLS} for any pair of values d_1 and d_2 in the grid sequence $\{-0.2, -0.1, \dots, 1.1, 1.2\}$. Point estimates of the long memory parameter vector can also be obtained by minimising the value of LM_d^{FGLS} over Ψ ; notice that this estimate does not depend on the confidence level used. This method of point estimation has been used in the univariate context; see, for example, Gil-Alaña and Robinson (1997). We denote the resulting point estimates of the long memory parameter for log trading volume and log volatility as $\hat{d}_{\min}(vlm)$ and $\hat{d}_{\min}(\sigma)$, respectively.⁵

Second, to account for deterministic effects in the level of these series, we apply the OLS-based demeaning procedure described in Remark 8. While most papers do not consider deterministic trends as a stylised feature of return volatility, trading volume is widely accepted to exhibit trending paths conformable with increasing growth in the number of traders and trading activity; see Fleming and Kirby (2011) and references therein. Consequently, for the log volatility measures, our main analysis is carried out by including a constant to capture a non-zero drift, as in Hassler et al. (2016), while in the case of log volume we allow for a quadratic time trend polynomial of the form $\mu_t = \mu_0 + \mu_1(t/T) + \mu_2(t/T)^2$, as advocated by, among others, Luu and Martens (2003) and Fleming and Kirby (2011). Parameters in these functions are estimated through univariate OLS (see Remark 8), with the multivariate fractional integration test then computed on the resultant residuals.

Third, as discussed in Remark 9, we determine the lag length according to Schwert's rule, $p = \lfloor 4(T/100)^{1/4} \rfloor$. Given the large sample size involved, Schwert's rule ensures a relatively long lag length, so that the short-run component of log volume and log realised variance should be well captured in the auxiliary regression. Andersen et al. (2003) also adopt a

⁵Numerical experiments investigating the finite sample accuracy (bias and MSE) of these estimates in the context of a bivariate model are reported in Section B.2 of the supporting information appendix.

relatively long lag length in estimating their FIVAR model for the realised volatility of exchange rates in order to maintain a conservative approach.⁶

5.3 | Main results

For each stock and for each volatility measure, Table 5 reports the resulting point estimates $\hat{d}_{\min}(vlm)$ and $\hat{d}_{\min}(\sigma)$. Table 5 also gives the upper and lower bounds of the corresponding 95% confidence ellipsoids formed as the vertical and horizontal projections of the confidence set onto the log-trading volume and log-volatility axes, respectively.⁷ The columns headed 'Common d ' in Table 5 report the range of values $\bar{d}$ for which the null hypothesis $H_0 : d(vlm) = d(\sigma) = \bar{d}$ cannot be rejected at the asymptotic 5% nominal significance level. If this region is non-empty, it shows the set of values along the 45-degree line contained within the 95% confidence ellipsoid; that is, those values of a common order of fractional integration for which the null cannot be rejected. Notice that, by construction, the resulting interval contains the true value of a common long memory parameter with an (asymptotic) probability not smaller than 95%. In addition, given $\hat{d}_{\min}(vlm)$ and $\hat{d}_{\min}(\sigma)$, we can compute the residuals from the multivariate FGLS regression and use these to estimate the contemporaneous correlation between the innovations to log-volume and a given return volatility measure; this estimate is denoted by $\hat{\rho}_e$ in Table 5. Large values of $\hat{\rho}_e$ are supportive of the usefulness of the multivariate approach we advocate. And, indeed, we see from Table 5 that this estimated correlation is generally quite large and positive.

Let us first discuss the results from the analysis of the joint dynamics of log volume and log absolute returns. Consistent with previous literature, we observe that for most stocks considered our multivariate test rejects both the null hypothesis that the order of integration of the bivariate series is $I(0)$ (such that both variables are weakly dependent) and the null hypothesis that it is $I(1)$ (such that both series admit an autoregressive unit root). The only exceptions are INTC (Intel) and MSFT (Microsoft), for which the multivariate test cannot reject the null hypothesis that log-volume is $I(0)$ at the 5% level. Taking a simple average of the estimates $\hat{d}_{\min}(vlm)$ and $\hat{d}_{\min}(\sigma)$ across all stocks considered yields 0.41 and 0.39, respectively, essentially matching the 'characteristic' value of 0.40 typically found in literature; see, Andersen et al. (2003). While for many of the stocks considered the point estimates of the vector of fractional exponents are below the non-stationary threshold, we note that for most of the stocks the respective confidence sets cover both the stationary and non-stationary regions of the parameter space, preventing us from drawing clear conclusions on the stationarity of the underlying series. This is a common finding in the realised-volatility modelling literature; see, for example, Kellard and Sarantis (2010).

Reflecting the strong similarities seen between the estimates of the two long memory parameters in the bivariate system, the hypothesis that trading volume and return volatility are driven by a FIVAR model with the same fractional exponent can be rejected for only five of the stocks considered at the 5% level, which constitutes about 20% of the stocks in our sample. This is, however, considerably higher than the corresponding frequency found by Bollerslev and Jubinski (1999) who only reject for 8% of the series they considered but is the same as Lobato and Velasco (2000) who also reject the null hypothesis of a common long memory parameter for 20% of the series they consider.⁸ In their study of log-volume and log absolute returns, Fleming and Kirby (2011) reject the common long memory parameter null for 100% of the series they analyse. They attribute this to estimation bias in the QML-based inference they use yielding systematically larger parameter estimates for the long memory coefficient for trading volume, and conjecture that departures from normality in log absolute returns may be causing a pervasive effect on QML estimation; see Fleming and Kirby (2011, pp. 1721–1722). Our Monte Carlo simulations in Section 4 accord with this conjecture suggesting that the combination of persistent time-varying volatility and non-Gaussian features in the data can introduce sizeable biases into the QML-based methods of Nielsen (2004a, 2005) used by Fleming and Kirby (2011).

⁶We also investigated the robustness of our main conclusions to the lag augmentation order used in the FGLS regression and to the inclusion of a deterministic time trend in connection with the return volatility measures. To that end, as in Fleming and Kirby (2011), we also looked at the case where a linear time trend was allowed for in return volatility and a low-order VAR(p) was fitted. Table C.3 in the supporting information appendix reports the main results from this analysis, focusing directly on log volume and log realised variance, with $p = 2$ and both with and without a linear time trend in return volatility. Here, we also report the related results when p is chosen according to Schwert's rule and return volatility includes a deterministic time trend. While the results show some sensitivity to these variations in the estimated model, the main qualitative picture that emerges is essentially very similar to that discussed below.

⁷These bounds (projections) define a rectangular approximation to the true confidence interval ellipsoids, whose area cannot be smaller than that of the true ellipsoid. However, they have the advantage that they provide a summary measure which can easily be tabulated. The full set of confidence ellipsoids for each stock considered can be found in Sections C.2.2–C.2.4 of the supporting information appendix.

⁸In making such comparisons, it is important to note, however, that these authors use different sample data than we do involving different stocks and different time periods. In particular, the sample lengths considered in Lobato and Velasco (2000) are more than double those we consider and our findings of the same frequency of rejections of a common exponent as they do may reflect the greater efficiency of the methods used here.

TABLE 5 Results from the implementation of the FGLS multivariate test on log-volume and log-volatility (proxied by log-absolute returns, log-range estimator and log-realised variance) for each stock in the sample

Stock	Log absolute returns						Log range						Log RV					
	$\hat{d}_{\min}(vlm)$	95% CIB	$\hat{d}_{\min}(\sigma)$	95% CIB	Common d	$\hat{\rho}_e$	$\hat{d}_{\min}(vlm)$	95% CIB	Common d	$\hat{\rho}_e$	$\hat{d}_{\min}(vlm)$	95% CIB	$\hat{d}_{\min}(\sigma)$	95% CIB	Common d	$\hat{\rho}_e$	$\hat{d}_{\min}(vlm)$	95% CIB
AAPL	0.62	[0.52,0.71]	0.38	[0.31,0.45]	-	0.41	0.65	[0.56,0.73]	0.59	[0.52,0.67]	0.63	0.66	[0.57,0.74]	0.64	[0.55,0.72]	0.62	0.62	[0.58,0.71]
AXP	0.49	[0.39,0.58]	0.46	[0.37,0.56]	[0.40,0.55]	0.37	0.53	[0.45,0.62]	0.61	[0.53,0.69]	0.53	0.53	[0.45,0.60]	0.61	[0.51,0.70]	0.49	0.49	[0.54,0.56]
BA	0.34	[0.23,0.44]	0.30	[0.22,0.39]	[0.25,0.39]	0.39	0.41	[0.30,0.50]	0.57	[0.48,0.66]	-	0.52	0.40	[0.30,0.51]	0.58	[0.46,0.69]	-	0.52
CAT	0.42	[0.29,0.54]	0.41	[0.32,0.50]	[0.32,0.50]	0.43	0.45	[0.33,0.56]	0.56	[0.46,0.65]	[0.49,0.55]	0.53	0.48	[0.37,0.59]	0.61	[0.50,0.70]	-	0.58
CSCO	0.35	[0.23,0.45]	0.38	[0.29,0.49]	[0.29,0.44]	0.36	0.39	[0.27,0.49]	0.57	[0.48,0.65]	-	0.48	0.41	[0.31,0.51]	0.58	[0.49,0.66]	-	0.53
CVX	0.52	[0.40,0.63]	0.37	[0.29,0.45]	-	0.30	0.50	[0.41,0.59]	0.53	[0.43,0.64]	[0.44,0.59]	0.47	0.50	[0.41,0.59]	0.57	[0.47,0.68]	[0.47,0.58]	0.43
DD	0.36	[0.22,0.48]	0.40	[0.30,0.50]	[0.30,0.47]	0.32	0.38	[0.26,0.50]	0.55	[0.44,0.64]	-	0.47	0.39	[0.28,0.50]	0.57	[0.46,0.68]	-	0.46
DIS	0.32	[0.18,0.45]	0.40	[0.30,0.52]	[0.31,0.45]	0.37	0.38	[0.26,0.50]	0.62	[0.52,0.71]	-	0.47	0.38	[0.27,0.49]	0.59	[0.49,0.68]	-	0.42
GE	0.47	[0.37,0.58]	0.40	[0.32,0.49]	[0.38,0.48]	0.42	0.52	[0.43,0.61]	0.58	[0.49,0.67]	[0.49,0.61]	0.54	0.52	[0.43,0.61]	0.59	[0.48,0.68]	[0.48,0.59]	0.51
GS	0.52	[0.39,0.65]	0.40	[0.31,0.48]	[0.40,0.44]	0.38	0.57	[0.47,0.68]	0.62	[0.53,0.71]	[0.53,0.68]	0.58	0.55	[0.46,0.65]	0.58	[0.48,0.68]	[0.48,0.65]	0.58
HD	0.54	[0.44,0.63]	0.43	[0.34,0.53]	[0.45,0.52]	0.35	0.57	[0.47,0.65]	0.61	[0.54,0.69]	[0.54,0.65]	0.46	0.56	[0.47,0.63]	0.62	[0.52,0.70]	[0.52,0.63]	0.47
IBM	0.36	[0.24,0.47]	0.30	[0.20,0.41]	[0.24,0.41]	0.38	0.39	[0.29,0.48]	0.54	[0.42,0.64]	-	0.47	0.41	[0.31,0.50]	0.55	[0.44,0.65]	-	0.47
INTC	0.18	[-0.03,0.34]	0.41	[0.30,0.54]	-	0.36	0.28	[0.10,0.42]	0.64	[0.53,0.73]	-	0.46	0.29	[0.12,0.42]	0.60	[0.51,0.69]	-	0.50
JNJ	0.44	[0.32,0.54]	0.35	[0.22,0.51]	[0.32,0.51]	0.36	0.46	[0.37,0.56]	0.58	[0.48,0.67]	-	0.49	0.46	[0.37,0.55]	0.62	[0.50,0.72]	-	0.44
JPM	0.56	[0.45,0.65]	0.47	[0.37,0.56]	[0.47,0.56]	0.38	0.59	[0.52,0.67]	0.62	[0.55,0.69]	[0.55,0.67]	0.56	0.58	[0.51,0.66]	0.61	[0.53,0.70]	[0.53,0.66]	0.56
KO	0.44	[0.30,0.57]	0.38	[0.29,0.47]	[0.32,0.47]	0.38	0.44	[0.30,0.56]	0.58	[0.50,0.66]	-	0.51	0.46	[0.33,0.58]	0.61	[0.52,0.69]	-	0.47
MCD	0.38	[0.25,0.52]	0.44	[0.33,0.56]	[0.33,0.51]	0.33	0.40	[0.29,0.52]	0.63	[0.55,0.71]	-	0.49	0.42	[0.30,0.54]	0.62	[0.53,0.70]	-	0.46
MMM	0.30	[0.11,0.47]	0.36	[0.25,0.46]	[0.28,0.45]	0.37	0.35	[0.20,0.48]	0.52	[0.42,0.61]	-	0.49	0.37	[0.24,0.50]	0.55	[0.45,0.65]	-	0.47
MRK	0.35	[0.17,0.50]	0.35	[0.27,0.44]	[0.27,0.44]	0.39	0.43	[0.24,0.56]	0.55	[0.46,0.64]	[0.48,0.56]	0.50	0.43	[0.28,0.55]	0.55	[0.45,0.64]	[0.48,0.54]	0.50
MSFT	0.16	[-0.11,0.37]	0.43	[0.35,0.54]	-	0.36	0.25	[0.03,0.44]	0.58	[0.47,0.68]	-	0.49	0.25	[0.06,0.43]	0.52	[0.41,0.62]	-	0.50
NKE	0.41	[0.27,0.53]	0.37	[0.28,0.47]	[0.30,0.47]	0.35	0.38	[0.26,0.49]	0.53	[0.44,0.62]	-	0.46	0.39	[0.28,0.50]	0.55	[0.46,0.63]	-	0.44
PFE	0.32	[0.04,0.52]	0.35	[0.26,0.46]	[0.26,0.46]	0.36	0.39	[0.21,0.55]	0.53	[0.42,0.64]	[0.48,0.54]	0.50	0.36	[0.19,0.52]	0.54	[0.42,0.66]	-	0.45
PG	0.38	[0.20,0.53]	0.36	[0.27,0.46]	[0.27,0.45]	0.34	0.40	[0.24,0.53]	0.47	[0.36,0.57]	[0.36,0.53]	0.45	0.39	[0.24,0.52]	0.50	[0.40,0.61]	[0.43,0.52]	0.43
TRV	0.40	[0.24,0.54]	0.45	[0.33,0.59]	[0.34,0.53]	0.27	0.47	[0.34,0.59]	0.66	[0.52,0.80]	-	0.42	0.47	[0.36,0.59]	0.62	[0.50,0.76]	-	0.43
UNH	0.26	[0.06,0.42]	0.35	[0.26,0.44]	[0.27,0.41]	0.37	0.39	[0.20,0.54]	0.59	[0.51,0.67]	-	0.54	0.39	[0.22,0.53]	0.61	[0.52,0.70]	-	0.52
UTX	0.43	[0.30,0.54]	0.37	[0.28,0.46]	[0.31,0.46]	0.30	0.45	[0.34,0.55]	0.57	[0.47,0.65]	-	0.46	0.44	[0.34,0.53]	0.57	[0.48,0.67]	-	0.44
V	0.73	[0.62,0.83]	0.43	[0.33,0.53]	-	0.30	0.71	[0.61,0.81]	0.62	[0.54,0.71]	[0.62,0.69]	0.42	0.73	[0.63,0.83]	0.63	[0.55,0.72]	[0.64,0.70]	0.44
VZ	0.40	[0.24,0.54]	0.41	[0.30,0.53]	[0.31,0.50]	0.32	0.44	[0.27,0.58]	0.60	[0.51,0.69]	-	0.48	0.46	[0.31,0.58]	0.61	[0.50,0.70]	-	0.42
WMT	0.38	[0.24,0.51]	0.37	[0.29,0.47]	[0.29,0.46]	0.40	0.44	[0.32,0.55]	0.54	[0.45,0.63]	[0.49,0.53]	0.50	0.44	[0.30,0.55]	0.57	[0.47,0.66]	-	0.51
XOM	0.48	[0.33,0.60]	0.30	[0.21,0.40]	-	0.30	0.50	[0.38,0.60]	0.53	[0.41,0.64]	[0.41,0.60]	0.47	0.49	[0.37,0.59]	0.52	[0.39,0.64]	[0.39,0.59]	0.46
Average	0.41		0.39		20.00%	0.36	0.45		0.58		0.49	0.45			0.58			
Rejection rate																		63.33%

Note: The columns headed: ' $\hat{d}_{\min}(vlm)$ ' and ' $\hat{d}_{\min}(\sigma)$ ' denote the values which minimise the test statistic; '95% CIB' denote the bounds of the 95% confidence ellipsoid for the long memory parameter; when non-empty, 'Common d ' reports the values along the 45° line contained in the 95% confidence ellipsoid (if empty, no common value is detected); ' $\hat{\rho}_e$ ' reports the sample cross-correlation coefficient of the residuals from the FGLS auxiliary regression given the values $\hat{d}_{\min}(vlm)$ and $\hat{d}_{\min}(\sigma)$.

We now move to a discussion of the results relating to the use of the log range and log realised proxies for return volatility. The overall picture that emerges here is remarkably similar in both cases. Multivariate estimation provides even stronger support for fractional integration in this context, with the $I(0)$ and $I(1)$ null hypotheses both being rejected at the 5% level for all of the stocks considered. For log trading volume, some changes are seen relative to the results discussed previously relating to the use of log absolute returns.⁹ Overall, the average value of $\hat{d}_{\min}(\text{vlm})$ taken across all of the stocks considered for the log range and log realised variance estimators is 0.45, reasonably similar to the 0.41 value reported above in connection with the use of log absolute returns. In contrast, the estimates of the order of integration of return volatility based on either the log-range or log realised variance measures show a marked increase compared to log absolute returns. In both cases, the average value of $\hat{d}_{\min}(\sigma)$ is 0.58, and the overall evidence is strongly suggestive that return volatility displays non-stationary dynamics over the period because the lower bounds of the confidence ellipsoids are not smaller than the 0.5 threshold for many of the stocks considered. Evidence of non-stationary fractionally integrated dynamics in realised volatility over this period (which includes the financial crisis) is consistent with the results reported by Hassler et al. (2016); see also Bandi and Perron (2006). Consequently, and because the persistence of log realised variance tends to be greater than that of log volume, the hypothesis that both variables share a common fractional exponent is rejected at the 5% level for a significantly larger proportion of the stocks considered, namely, 53.33% when using log range and 63.33% when using log realised variance.

It is well understood in the financial econometrics literature that measurement errors in absolute returns can cause bias in (univariate) long memory parameter estimation. Essentially, log absolute returns are subject to noisy additive measurement errors with large variability, which will make the underlying process appear less persistent than it really is, leading to downward-biased estimates of the true order of fractional integration; see, among others, Bollerslev and Wright (2000), Haldrup and Nielsen (2007) and Dalla (2015). This provides a straightforward and plausible explanation for the systematic differences seen in the long memory estimation results for the different return volatility measures reported in Table 5.¹⁰ According to our results, the more efficient the estimate of return volatility used the higher the percentage of the stocks for which the null hypothesis of a common order of integration can be rejected. Essentially, downward biases in the estimation of the long memory parameter on absolute returns biases the tests to non-rejection of a common order of integration. Using more accurate return volatility measures reduces this estimation bias and leads to increased evidence that return volatility is more persistent than volume.

5.4 | Fractional cointegration

As noted in Remark 5, our FGLS-based LM test, like the LM test of Nielsen (2005), assumes the absence of fractional cointegration between the variables, and diverges if fractional cointegration is present. Given that we reject the null hypothesis of a common order of integration for trading volume and return volatility for most of the stocks considered, we now also investigate the order of fractional integration of the series using the semiparametric approach of Nielsen and Shimotsu (2007) (NS henceforth), detailed in the supporting information appendix. NS's procedure allows us (under certain regularity conditions) to consistently estimate the cointegration rank of the series and, using the approach of Robinson and Yajima (2002), to test the null hypothesis that the elements of long memory vector, $\mathbf{d}$, are equal (although it is important to note that this is not a multivariate test as it is based on the univariate estimates of the fractional exponents). Denoting the statistic for the latter as T_0 , NS show that $T_0 \xrightarrow{p} 0$ when the cointegration rank, r , is greater than zero (i.e., where the variables are cointegrated), whereas $T_0 \Rightarrow \chi^2_{(1)}$ when $r = 0$ (where the variables are not cointegrated) and the null of an equal order of integration holds on $\mathbf{d}$. NS argue that large values of the test statistic provide evidence against the hypothesis of a common order of integration, regardless of whether the underlying series are fractionally cointegrated or not.

Given that the highest frequency of rejections of a common fractional exponent occurred when using log realised variances, we only report that case here. Consistent with the analysis in Table 5, we account for a deterministic drift in log volatility and a polynomial time trend in log volume by prior detrending of the data, using the two-stage exact local Whittle estimator in Shimotsu (2010). Following the empirical analysis in NS, we estimated $\mathbf{d}$ by setting $m_T = \lfloor T^{0.6} \rfloor$

⁹Because we conduct joint estimation, and the innovations to the short-term component of volume and return volatility are strongly positively correlated, as reported in the column $\hat{\rho}_e$ in Table 5, the estimates of the long memory parameter of log volume would be expected to be somewhat sensitive to changes in the variable used as a proxy for return volatility.

¹⁰An alternative explanation, put forward by a referee, is that the VAR dynamics may be misspecified and, as a result, some of the high-frequency measurement error is picked up in the estimate of the fractional exponent. However, the relatively long lag length used should mitigate against this and, moreover, as discussed in footnote 6, the results appear relatively robust to the lag order specified.

TABLE 6 Results from the the Nielsen and Shimotsu (2007)-based approach

Stock	Nielsen–Shimotsu testing approach									
	$\hat{d}(vlm)$	95% CI	$\hat{d}(\sigma)$	95% CI	T_0	p value	$L(0)$	$L(1)$	$\hat{r}_T$	$\hat{r}_T^*$
AAPL	0.62	[0.53,0.71]	0.48	[0.40,0.57]	5.37	0.02	−1.465	−1.337	0	−
AXP	0.52	[0.43,0.61]	0.63	[0.54,0.72]	3.90	0.05	−1.465	−1.444	0	−
BA	0.37	[0.28,0.46]	0.57	[0.48,0.66]	11.41	0.00	−1.465	−1.366	0	−
CAT	0.40	[0.31,0.48]	0.56	[0.47,0.65]	8.74	0.00	−1.465	−1.419	0	−
CSCO	0.34	[0.25,0.42]	0.52	[0.44,0.61]	10.46	0.00	−1.465	−1.416	0	−
CVX	0.52	[0.43,0.61]	0.61	[0.52,0.70]	2.78	0.10	−1.465	−1.455	0	0
DD	0.35	[0.26,0.44]	0.55	[0.46,0.64]	12.08	0.00	−1.465	−1.378	0	−
DIS	0.38	[0.29,0.47]	0.61	[0.52,0.70]	14.55	0.00	−1.465	−1.408	0	−
GE	0.55	[0.46,0.64]	0.60	[0.51,0.69]	0.90	0.34	−1.465	−1.468	1	1
GS	0.53	[0.44,0.62]	0.56	[0.47,0.65]	0.35	0.56	−1.465	−1.482	1	1
HD	0.60	[0.51,0.69]	0.57	[0.48,0.66]	0.33	0.57	−1.465	−1.380	0	0
IBM	0.39	[0.31,0.48]	0.58	[0.49,0.66]	10.83	0.00	−1.465	−1.460	0	−
INTC	0.19	[0.10,0.28]	0.54	[0.45,0.63]	33.64	0.00	−1.465	−1.301	0	−
JNJ	0.41	[0.33,0.50]	0.60	[0.51,0.69]	10.86	0.00	−1.465	−1.452	0	−
JPM	0.59	[0.50,0.68]	0.62	[0.53,0.71]	0.31	0.58	−1.465	−1.530	1	1
KO	0.40	[0.31,0.49]	0.61	[0.52,0.70]	13.03	0.00	−1.465	−1.350	0	−
MCD	0.35	[0.26,0.44]	0.60	[0.51,0.69]	16.74	0.00	−1.465	−1.258	0	−
MMM	0.36	[0.27,0.45]	0.54	[0.45,0.63]	10.87	0.00	−1.465	−1.447	0	−
MRK	0.33	[0.24,0.41]	0.52	[0.43,0.61]	11.35	0.00	−1.465	−1.368	0	−
MSFT	0.22	[0.14,0.31]	0.48	[0.39,0.56]	16.50	0.00	−1.465	−1.351	0	−
NKE	0.39	[0.30,0.48]	0.52	[0.43,0.60]	4.45	0.03	−1.465	−1.326	0	−
PFE	0.37	[0.28,0.46]	0.53	[0.44,0.62]	7.52	0.01	−1.465	−1.332	0	−
PG	0.36	[0.27,0.45]	0.48	[0.39,0.57]	3.89	0.05	−1.465	−1.323	0	−
TRV	0.41	[0.30,0.51]	0.59	[0.49,0.69]	7.34	0.01	−1.426	−1.399	0	−
UNH	0.33	[0.24,0.41]	0.56	[0.47,0.65]	16.39	0.00	−1.465	−1.338	0	−
UTX	0.41	[0.32,0.50]	0.58	[0.49,0.67]	9.01	0.00	−1.465	−1.420	0	−
VZ	0.39	[0.30,0.48]	0.57	[0.48,0.66]	8.22	0.00	−1.465	−1.248	0	−
V	0.64	[0.55,0.74]	0.62	[0.53,0.72]	0.11	0.73	−1.452	−1.305	0	0
WMT	0.40	[0.32,0.49]	0.57	[0.48,0.66]	7.82	0.01	−1.465	−1.392	0	−
XOM	0.47	[0.38,0.56]	0.53	[0.44,0.62]	1.06	0.30	−1.465	−1.493	1	1
Average	0.42		0.56						13.33%	13.33%
Rejection 95%							76.67%			

Note: The columns headed: ' $\hat{d}(vlm)$ ' and ' $\hat{d}(\sigma)$ ' report point estimates from the two-stage univariate exact local Whittle estimator in Shimotsu and Phillips (2005); ' T_0 ' and ' p value' report the test statistic for the null of a common order of integration and related p values; ' $L(u)$ ', $u = 0, 1$, report the objective function used to infer the cointegration rank in a model selection procedure; finally, ' $\hat{r}_T$ ' reports the estimated cointegration rank resulting from this criterion, and ' $\hat{r}_T^*$ ' reports the conditional estimates of r for the cases in which the null hypothesis of a common order of integration cannot be rejected at 5% level.

and compute T_0 with $s_T = \log T$. Following NS, we also use $m_{1T} = \lfloor T^{0.55} \rfloor$ and $v_T = m_{1T}^{-0.3}$ in the estimation of the cointegration rank. Table 6 reports the point estimates and 95% asymptotic confidence level estimates of $\mathbf{d}$, the T_0 test statistic and related p values, the values of the function $L(u)$ used in the model selection procedure, and the estimates of the cointegration rank, $\hat{r}_T$. The column $\hat{r}_T^*$ reports the conditional estimates of r for the cases in which the hypothesis of a common order of integration cannot be rejected at the 5% nominal size level.

Three key features arise from this analysis. First, the results based on the NS test provide the same qualitative evidence as the tests based on FGLS estimation. There exists strong evidence of fractional integration in both series, which is again suggestive that realised volatility is more persistent than trading volume. Accordingly, the hypothesis of an equal order of integration is rejected at any of the usual significance levels for the majority of the stocks in our sample. Second, in most cases, the FGLS and the NS tests agree on whether to reject or not the null hypothesis of a common order of fractional integration. In particular, all of the cases in which the FGLS test rejects at the 5% level correspond to stocks for which the NS test also rejects at this level. There are, however, five stocks for which NS rejects the null but the FGLS test does not, so the average rejection rate of the NS test taken across all of the stocks considered is slightly higher at 76.67%. Crucially, the p values of the T_0 test in three of these five cases are only slightly below the 5% threshold, suggesting that the differences with the FGLS test are caused by only marginal differences in significance. Finally, the estimates of the cointegration

rank suggests that volume and realised volatility are not in general cointegrated, supporting the suitability of FIVAR-type modelling. In particular, fractional cointegration, indicated by $\hat{r}_T = 1$ or $\hat{r}_T^* = 1$, is found for only 13.33% of the stocks, but crucially, these are all stocks for which neither the NS T_0 test nor our FGLS-based tests reject the null hypothesis of a common order of integration; that is, none of the rejections of a common order of integration seen with the FGLS-based test in Table 5 are associated with a non-zero estimate of the cointegration rank.

It is worth pointing out in conclusion that the assumptions on which the NS approach are based include the requirement of conditional homoskedasticity. To check how sensitive the NS test is to violations of this assumption, we conducted a small Monte Carlo experiment using the sample simulation DGP as in Section 4. For samples of size $T = 500$ and $T = 1000$, $\rho = 0.8$ and persistent GARCH processes with $(\alpha, \beta) = (0.1, 0.85)$, the ERFs of the T_0 test under Gaussian innovations at the 5% nominal asymptotic level were 3.10% and 4.30%, respectively, suggesting approximately correct size. However, under Student- t innovations with 5 degrees of freedom, the respective ERFs were 15.5% and 24.5%. Although clearly oversized, these distortions are considerably smaller than those that were seen in the corresponding results in Section 4 for the QML-based test of Nielsen (2005). These simulation results might help explain the differences seen between the results for the FGLS and T_0 tests in our empirical study whereby slightly more rejections of the null of a common order of integration are obtained when using the T_0 test, given that this test has a tendency to be somewhat oversized when the data display conditional heteroskedasticity and heavy-tailed behaviour, as is the case with the data in our empirical study.

6 | CONCLUSIONS

We have proposed that a new test for fractional integration in the context of a quite general FIVAR model, which allows for conditional heteroskedasticity in the innovations, does not require the order of integration of the elements of the vector time series to coincide or to lie in a certain region (thereby allowing for both stationary and non-stationary dynamics) and does not assume a particular distribution for the innovations. To the best of our knowledge, none of the methods in the extant literature has achieved this degree of flexibility. Our approach is based on an LM-type test statistic using a heteroskedasticity-robust estimate of the variance matrix and can be readily implemented using FGLS estimation in a regression-based context. We have demonstrated that our proposed test statistic has a standard χ^2 asymptotic null distribution, that the test exhibits non-trivial power to reject against a sequence of local alternatives and that in the case of i.i.d. Gaussian errors the test is asymptotically locally efficient. Monte Carlo analysis was used to show that while our test is approximately correctly sized in finite samples of data exhibiting conditional heteroskedasticity and heavy-tailed features, extant tests in the literature that neglect conditional heteroskedasticity can be severely oversized even for very large samples.

In an empirical case study, we have used our proposed testing procedure to jointly infer the order of fractional integration of trading volume and return volatility in a sample of major stocks traded in the US market. Return volatility was proxied by three different measures with increasing degrees of accuracy: absolute returns, a range-based estimator and a realised variance computed over 5-min returns. The evidence from the analysis based on the realised variance and range-based estimates delivered similar conclusions, namely, that for many stocks in the sample return volatility is more persistent than trading volume. On the other hand, the analysis based on log absolute returns showed that volume and return volatility share the same order of fractional integration. Because long memory estimation in absolute returns is known to be downward biased, measurement errors in the data would seem to be a plausible explanation for the evidence from log absolute returns.

For applied work, it is of interest that our conclusions based on the realised variance and range-based estimators of return volatility were very similar. While the former is a more efficient estimate of conditional variability, the latter seems to provide a reasonable enough level of accuracy such that the conclusions drawn from the data are not markedly different. This might be a useful observation in practice because for many applications the high-frequency intra-day data needed to construct realised variance and related measures is often not available, for example, when considering small or illiquid markets. In the absence of intra-day data, but when information on high and low prices are available, inference based on range-based volatility estimates may still lead to reliable conclusions.

We finish with a suggestion for further research. Here, we have proposed parametric FGLS-based multivariate fractional integration tests which, unlike other extant parametric tests, allow for conditionally heteroskedastic innovations. There are relatively few semiparametric multivariate fractional integration tests in the literature, most notably Lobato and Robinson (1998), Lobato (1999), Marinucci and Robinson (2001) and Shimotsu (2007), all of which assume conditionally homoskedastic innovations. Investigating whether or not these tests remain asymptotically valid under conditional het-

eroskedasticity and, if so, comparing their finite sample performance with the tests developed in this paper is beyond the scope of the present paper but would constitute an interesting topic for further research.

ACKNOWLEDGMENTS

We thank the Co-Editor, Eric Ghysels, and three anonymous referees for their helpful and constructive comments on an earlier version of this paper. Balboa and Rubia gratefully acknowledge financial support from the Spanish Ministry of Economy through project ECO2017-87069-P. Rodrigues gratefully acknowledges financial support from the Portuguese Science Foundation (FCT) through project PTDC/EGE-ECO/28924/2017, and (UID/ECO/00124/2013 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

OPEN RESEARCH BADGES



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available at [[https://urldefense.com/v3/__http://qed.econ.queensu.ca/jae/datasets/balboa001/_;!!N11eV2iwtfs!7QfEpR5YsUNsoUTjHTRM9PI1lx78tJKDTbd1jOZReF1VGtY0u3XnXY0-pKST2zF3\\$](https://urldefense.com/v3/__http://qed.econ.queensu.ca/jae/datasets/balboa001/_;!!N11eV2iwtfs!7QfEpR5YsUNsoUTjHTRM9PI1lx78tJKDTbd1jOZReF1VGtY0u3XnXY0-pKST2zF3$)].

REFERENCES

- Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4), 885–905.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71, 579–625.
- Baillie, R. T. (1996). Long memory processes and fractional integration in econometrics. *Journal of Econometrics*, 73, 5–59.
- Bandi, F. M., & Perron, B. (2006). Long memory and the relation between implied and realized volatility. *Journal of Financial Econometrics*, 4, 636–670.
- Bollerslev, T., & Jubinski, D. (1999). Equity trading volume and volatility: Latent information arrivals and common long-run dependencies. *Journal of Business & Economic Statistics*, 17(1), 9–21.
- Bollerslev, T., & Wright, J. (2000). Semiparametric estimation of long-memory volatility dependencies: The role of high-frequency data. *Journal of Econometrics*, 98, 81–106.
- Boswijk, H. P., Cavaliere, G., Rahbek, A., & Taylor, A. M. R. (2016). Inference on co-integration parameters in heteroskedastic vector autoregressions. *Journal of Econometrics*, 192(1), 64–85.
- Breitung, J., & Hassler, U. (2002). Inference on the cointegration rank in fractionally integrated processes. *Journal of Econometrics*, 110, 167–185.
- Cavaliere, G., Nielsen, M. O., & Taylor, A. M. R. (2017). Quasi-maximum likelihood estimation and bootstrap inference in fractional time series models with heteroskedasticity of unknown form. *Journal of Econometrics*, 198(1), 165–188.
- Dalla, V. (2015). Power transformations of absolute returns and long memory estimation. *Journal of Empirical Finance*, 33, 1–18.
- Demetrescu, M., Kuzin, V., & Hassler, U. (2008). Long memory testing in the time domain. *Econometric Theory*, 24(1), 176–215.
- Fleming, J., & Kirby, C. (2011). Long memory in volatility and trading volume. *Journal of Banking and Finance*, 35(7), 1714–1726.
- Garman, M. B., & Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *The Journal of Business*, 53(1), 67–78.
- Gil-Alaña, L. A., & Robinson, P. M. (1997). Testing of unit root and other nonstationary hypotheses in macroeconomic time series. *Journal of Econometrics*, 80(2), 241–268.
- Haldrup, N., & Nielsen, M. O. (2007). Estimation of fractional integration in the presence of data noise. *Computational Statistics & Data Analysis*, 51(6), 3100–3114.
- Hassler, U., Rodrigues, P. M. M., & Rubia, A. (2009). Testing for the general unit root hypothesis in the time domain. *Econometric Theory*, 25(6), 1793–1828.
- Hassler, U., Rodrigues, P. M. M., & Rubia, A. (2016). Quantile regression for long memory testing: A case of realized volatility. *Journal of Financial Econometrics*, 25(4), 693–724.
- Kellard, C. D. N., & Sarantis, N. (2010). Foreign exchange, fractional cointegration and the implied–realized volatility relation. *Journal of Banking and Finance*, 34(4), 882–891.
- LeBaron, B., & Yamamoto, R. (2008). The impact of imitation on long memory in an order-driven market. *Eastern Economic Journal*, 34, 504–517.
- Lobato, I. N. (1999). A semiparametric two-step estimator in a multivariate long memory model. *Journal of Econometrics*, 90, 129–153.
- Lobato, I. N., & Robinson, P. M. (1998). A nonparametric test for $i(0)$. *The Review of Economic Studies*, 65(3), 475–495.
- Lobato, I. N., & Velasco, C. (2000). Long memory in stock-market trading volume. *Journal of Business & Economic Statistics*, 18(4), 410–427.
- Lovcha, Y., & Perez-Laborda, P. A. (2018). Monetary policy shocks, inflation persistence, and long memory. *Journal of Macroeconomics*, 55, 117–127.

- Luu, J., & Martens, M. (2003). Testing the mixture of distributions hypothesis using “realized” volatility. *The Journal of Futures Markets*, 23, 61–69.
- Marinucci, D., & Robinson, P. M. (2001). Semiparametric fractional cointegration analysis. *Journal of Econometrics*, 105(1), 225–247.
- Nielsen, F. S. (2011). Local whittle estimation of multi-variate fractionally integrated processes. *Journal of Time Series Analysis*, 32, 317–335.
- Nielsen, M. O. (2004a). Efficient inference in multivariate fractionally integrated time series models. *Econometrics Journal*, 7(1), 63–97.
- Nielsen, M. O. (2004b). Efficient likelihood inference in nonstationary univariate models. *Econometric Theory*, 20(1), 116–146.
- Nielsen, M. O. (2005). Multivariate lagrange multiplier tests for fractional integration. *Journal of Financial Econometrics*, 3, 372–398.
- Nielsen, M. O., & Shimotsu, K. (2007). Determining the cointegrating rank in nonstationary fractional systems by the exact local whittle approach. *Journal of Econometrics*, 141(2), 574–596.
- Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *The Journal of Business*, 53(1), 61–65.
- Robinson, P. M. (1994). Efficient tests of nonstationary hypotheses. *Journal of the American Statistical Association*, 89, 1420–1437.
- Robinson, P. M. (1995). Log-periodogram regression of time series with long range dependence. *The Annals of Statistics*, 23(3), 1048–1072.
- Robinson, P. M. (2003). Long memory time series. In Robinson, P. M. (Ed.), *Time series with long memory* (pp. 4–32). Oxford: Oxford University Press.
- Robinson, P. M., & Yajima, Y. (2002). Determination of cointegrating rank in fractional systems. *Journal of Econometrics*, 106(2), 217–241.
- Rossi, E., & de Magistris, P. S. (2013). Long memory and tail dependence in trading volume and volatility. *Journal of Empirical Finance*, 22, 94–112.
- Schwert, G. W. (1989). Tests for unit roots: A Monte Carlo investigation. *Journal of Business & Economic Statistics*, 7(2), 147–160.
- Sela, R. J., & Hurvich, C. M. (2009). Computationally efficient methods for two multivariate fractionally integrated models. *Journal of Time Series Analysis*, 30, 631–651.
- Shimotsu, K. (2007). Gaussian semiparametric estimation of multivariate fractionally integrated processes. *Journal of Econometrics*, 137, 277–310.
- Shimotsu, K. (2010). Exact local whittle estimation of fractional integration with unknown mean and time trend. *Econometric Theory*, 26, 501–540.
- Shimotsu, K., & Phillips, P. C. B. (2005). Exact local whittle estimation of fractional integration. *The Annals of Statistics*, 33(4), 1890–1933.
- Tanaka, K. (1999). The nonstationary fractional unit root. *Econometric Theory*, 15, 549–582.
- Yamamoto, R. (2011). Order aggressiveness, pre-trade transparency, and long memory in an order-driven market. *Journal of Economic Dynamics and Control*, 35, 1938–1963.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Balboa, M., Rodrigues, P. M. M., Rubia, A., & Taylor, A. M. R. (2021) Multivariate fractional integration tests allowing for conditional heteroskedasticity with an application to return volatility and trading volume. *Journal of Applied Econometrics*. 2021;36:544–565. <https://doi.org/10.1002/jae.2829>