



Revisiting the long memory dynamics of the implied–realized volatility relationship: New evidence from the wavelet regression[☆]



Jozef Baruník^{a,b,*}, Michaela Hlínková^a

^aInstitute of Economic Studies, Charles University, Opletalova 26, 110 00 Prague, Czech Republic

^bInstitute of Information Theory and Automation, Czech Academy of Sciences, Pod Vodarenskou Vezi 4, 182 00 Prague, Czech Republic

ARTICLE INFO

Article history:

Accepted 19 January 2016

Available online 22 February 2016

Keywords:

Wavelet band spectrum regression

Corridor implied volatility

Realized volatility

Fractional cointegration

ABSTRACT

The literature studying stock index options confirms severe biases and inefficiencies in using implied volatility as a forecast of future volatility. In this paper, we revisit the implied–realized volatility relationship with wavelet band least squares (WBLS) exploring the long memory of volatility, a possible cause of the bias. Using the S&P 500 and DAX monthly and bi-weekly option prices covering the recent financial crisis, we conclude that the implied–realized volatility relation is driven solely by the lower frequencies of the spectra representing long investment horizons. The findings enable improvement of future volatility forecasts as they support unbiasedness of implied volatility as a good proxy for future volatility in the long run.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Option prices are widely believed to carry information relating to expectations of market participants about the future movement of the underlying asset prices in financial markets, mostly its volatility. The volatility implied by an option's price is the forecast of the future return volatility over the remaining life of the relevant option if the option markets are efficient. Early papers studying the phenomenon of implied–realized volatility relation use volatility implied by option pricing models – most commonly [Black and Scholes \(1973\)](#) or [Hull and White \(1987\)](#) – and come to the conclusion that volatility inferred from option markets is a biased predictor of stock return volatility ([Day and Lewis, 1992](#); [Lamoureux and Lastrapes, 1993](#); [Canina and Figlewski, 1993](#); [Jorion, 1995](#)).

In contrast, [Christensen and Prabhala \(1998\)](#) and [Christensen and Hansen \(2002\)](#) use a wide variety of methods to show that informational content of option implied volatility is superior to that of the past volatility, and it is a less biased (although still biased) predictor of

future realized volatility than what has been previously shown. The authors shed new light on the dubiety about the informational content of option implied volatility by specifying the sources of error in previous research. For example, the choice of particular option contracts for extracting volatility and lower liquidity of the option market than in the underlying asset market. Moreover, [Christensen and Prabhala \(1998\)](#) and [Christensen and Hansen \(2002\)](#) find that overlapping data errors can cause cross-correlation in the volatility series, which stems from the overlapping period between the current implied volatility and future implied volatility. In light of these methodological issues, [Christensen and Hansen \(2002\)](#) conclude that option implied volatility is a more efficient forecast for future realized volatility than historical volatility, but it does not subsume all information contained in historical volatility, and it results in upward biased forecasts.

Unlike the traditional concepts using the work of [Black and Scholes \(1973\)](#) or [Hull and White \(1987\)](#) to extract volatilities from options, model-free implied volatility (MFIV) introduced by [Britten-Jones and Neuberger \(2000\)](#) is not based on any specific option pricing model, and it is derived from no-arbitrage conditions. [Jiang and Tian \(2005\)](#) extended the simple measure of implied volatility to all martingale asset price processes and express the formula in forward rather than spot prices. Most notably, [Jiang and Tian \(2005\)](#) first find that the MFIV subsumes all information contained by historical and [Black and Scholes \(1973\)](#) implied volatility and is a more efficient forecast of future realized volatility. Hence, informational content of option implied volatility in the subsequent research is analyzed using the model-free

[☆] Jozef Baruník gratefully acknowledges the support of the Czech Science Foundation project no. P402/12/G097 DYME - “Dynamic Models in Economics”. The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007–2013) under grant agreement no. FP7-SSH-612955 (FinMaP).

* Corresponding author at: Institute of Economic Studies, Charles University in Prague, Opletalova 26, 110 00, Prague, Czech Republic. Tel.: +420 776 259273.

E-mail address: barunik@fsv.cuni.cz (J. Baruník).

measure.¹ For example, Seo and Kim (2015) find implied volatility to have varying forecasting ability depending on the level of investor sentiment. Although weaker than for stock markets, Chatrath et al. (2015) and Padungsakasawadi and Daigler (2014) confirm the relationship in commodity markets.

In subsequent research, Andersen and Bondarenko (2007) and Andersen et al. (2015) argue that MFIV computation brings serious practical limitations, yielding inaccurate results. The main problem is the lack of liquid options with strike prices covering the entire return distribution, including its tails. The authors advocate using limited strike ranges at a given point in time instead. The concept is called model-free corridor implied volatility (CIV), previously introduced by Carr and Madan (1998a). While different measures can be obtained, depending on the width and positioning of the strike ranges, Andersen et al. (2015) advocate fixing the range of strikes at a level that provides broad coverage but avoids excessive extrapolation of noisy or non-existing quotes for the out-of-the-money options. Recently, Muzzioli (2013) shed more light on the information content of different parts of the risk neutral distribution of the stock price by considering different corridors in CIV.

When assessing the efficiency of implied volatility forecasts, one needs to have return volatility at hand. However, actual volatility has not been a directly observable variable for a long time. In recent years, as a consequence of the increased availability of high-frequency data, another subject has brought new insight into the implied–realized volatility relationship; the concept of realized volatility. Andersen et al. (2003) and Barndorff-Nielsen and Shephard (2004a) have shown that realized variance provides a consistent nonparametric measure of price variability over a given time interval. An immense literature studying the realized volatility emerged in the past decade discussing the impact of noise as well as jumps in the volatility measurement, concluding that realized volatility is unbiased and a consistent measure of quadratic variation only if we assume no market microstructure noise in the process. The literature also argues that it is important to separate the jump process and use the estimator robust to noise to recover a true underlying volatility. For the measurement of realized volatility, we use one of the most recent jump wavelet two-scale realized volatility estimators (JWTSRV) by Baruník and Vacha (2015), which compares to other estimators used in the literature very well. JWTSRV is able to estimate volatility under the jumps and microstructure noise consistently. The forecasting power of the estimator is studied using a Realized GARCH framework in Baruník et al. (2016).

Using increasingly precise measures, the recent literature suggests that the predictive regression between implied volatility and realized volatility is a cointegrating relationship, and OLS estimation should be avoided as it will result in biased estimates (Bandi and Perron, 2006; Nielsen and Frederiksen, 2011). The relationship is driven primarily by the long memory of both implied and realized volatilities, a key stylized fact commonly found in empirical research across a wide variety of asset classes (Baillie, 1996b; Cont, 2001). Using a spectral method, Bandi and Perron (2006) and Nielsen and Frederiksen (2011) confirm that in a long run, implied volatility is an unbiased predictor of future realized volatility. Still, the results do not say anything about short-term unbiasedness, and they rely on Black–Scholes implied volatility only. Generally, band spectrum regression may be a useful tool in the situation where we believe the relationship between variables is dependent on frequency. The concept was introduced to econometrics by

Engle (1974) and further shown to be useful for estimation of cointegrating regressions (Phillips, 1991; Marinucci and Robinson, 2001).

While Bandi and Perron (2006); Nielsen and Frederiksen (2011); Kellard et al. (2010) use a Fourier transform to estimate the relationship in the frequency domain, we contribute to the literature by proposing the band regression on the spectrum estimated by wavelet coefficients. The wavelet transform offers localized frequency decomposition, providing information about frequency components. As a result, wavelets have significant advantages over basic Fourier analysis when the object under study is locally stationary and inhomogeneous – see Gençay et al. (2002); Percival and Walden (2000); Ramsay (2002). This can be a crucial property, as the implied–realized volatility cointegrating relationship may potentially lie in a non-stationary region (Kellard et al., 2010). Wavelets also allow us to study the relationship in the time–frequency domain. We motivate this dynamic by estimating the wavelet coherence measure to study the implied–realized relationship. While wavelet coherence may be used as the “lens” into the relationship that shows the dynamics through time, as well as frequencies at once, a newly proposed wavelet band spectral regression allows us to estimate the relationship.²

The contribution of this paper is twofold. First, we emphasize the importance of the implied volatility measure in studying the implied–realized volatility relationship. We compare MFIV and the recently proposed CIV as measures of option implied volatilities with realized volatility and recently proposed jump-wavelet realized volatility (JWTSRV) capable of separating the continuous part of the volatility from jumps as well as noise. We argue that it is crucial to use proper measures for finding the answer to the question of whether the option implied volatility is an efficient forecast of the future realized volatility. Second, we bring new evidence on the unbiasedness of ex-ante implied volatility as a predictor of ex-post realized volatility by allowing long memory dynamics in the time series. We find that the dependence comes solely from the longer time horizons, and when estimated using wavelet band least squares, the implied volatility forecasts are unbiased forecasts of future volatility. These findings greatly improve the understanding of volatility dynamics and add to previous findings of Li (2002), who stress the importance of long memory in studying the implied–realized volatility relationship, or Kinatader and Wagner (2014) who argue that long memory volatility prediction is influenced by the variance term structure.

The methods are applied to the German DAX and U.S. S&P 500 stock market indices covering the 2008 financial crisis, with abrupt changes in prices. Unlike the previous studies, we use both call and put options, and we use options with monthly as well as bi-weekly maturities.

2. Volatility measurement

Consider a univariate risky logarithmic asset price process p_t defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The price process evolves in continuous time t over the interval $[0, T]$, where T is a finite positive integer according to a jump diffusion process:

$$dp_t = \mu_t dt + \sigma_t dW_t + \xi_t dq_t, \quad (1)$$

where μ_t is a predictable mean, σ_t a strictly positive volatility process, W_t is standard Brownian motion, $\xi_t dq_t$ is a random jump

¹ Many studies use the Chicago Board Option Exchange (CBOE) Volatility Index (VIX) as a proxy for model-free implied volatility of S&P 500. Introduced by the CBOE in 1993, its methodology was revised in 2003 using a new model-free measure of expected volatility and thus can be used conveniently.

² Note that long memory properties in the emerging markets have been largely studied in the recent literature (Yalama and Celik, 2013; Degiannakis and Livada, 2013; Hull and McGroarty, 2014; Charfeddine and Ajmi, 2013)

process allowing for occasional jumps in p_t , and q is a Poisson process uncorrelated with W and governed by the constant jump intensity λ . The magnitude of the jump in the return process is controlled by factor $\xi_t \sim N(\bar{\xi}, \sigma_\xi^2)$.

Generally, we assume the latent logarithmic asset price process is contaminated with microstructure noise. Let y_t be the observed log prices, which will be equal to the latent, so-called “true” log-price process p_t in Eq. (1) and will contain microstructure noise ϵ_t , a zero mean *i.i.d.* noise with variance η^2

$$y_t = p_t + \epsilon_t. \tag{2}$$

The main object of interest in financial econometrics is the estimated integrated variance of the latent price process, $IV_{t,h} = \int_{t-h}^t \sigma_s^2 dt$. Quadratic return variation over the $[t - h, t]$ time interval, $0 \leq h \leq t \leq T$ is

$$QV_{t,h} = \underbrace{\int_{t-h}^t \sigma_s^2 ds}_{IV_{t,h}} + \underbrace{\sum_{t-h \leq s < t} J_s^2}_{JV_{t,h}}. \tag{3}$$

Thus, quadratic variation $QV_{t,h}$ is equal to the integrated volatility of the continuous path and the sum of jumps variation.

2.1. Realized volatility

The recently popularized simple measure of quadratic variation – realized variance – is a consistent and an unbiased estimator of the quadratic variation if the sampling goes to infinity (Andersen and Bollerslev, 1998; Andersen et al., 2003; Barndorff-Nielsen and Shephard, 2001, 2002a,b). The realized variance over $[t - h, t]$, for $0 \leq h \leq t \leq T$, is defined by

$$\hat{RV}_{t,h} = \sum_{i=1}^n r_{t-h+\frac{i}{n}h}^2 \tag{4}$$

where n is the number of observations in $[t - h, t]$. While the realized variance measure is widely used due to its simplicity, it estimates the whole part of the quadratic variation and is subject to the bias from the microstructure noise. Several estimators dealing with microstructure noise and jumps have been introduced recently. For example Zhang et al. (2005) propose the solution to the problem of noise by introducing the two-scale realized volatility (TSRV henceforth) estimator. Another estimator, which is able to address the noise is the realized kernels (RK) introduced by Barndorff-Nielsen et al. (2008). Barndorff-Nielsen and Shephard (2004b, 2006) developed a very powerful and complete way of detecting the presence of jumps in high-frequency data, bipower variation. The basic idea is to compare two measures of the integrated variance, one containing the jump variation and the other being robust to jumps and hence containing only the integrated variation part.

More recently, jump-adjusted wavelet two-scale realized volatility (JWTSRV) has been proposed to measure the integrated variance in the presence of jumps and noise by Baruník and Vacha (2015). JWTSRV is able to consistently estimate jumps using wavelet transforms and is also robust to microstructure noise because of a Zhang et al.’s (2005) framework. Let us introduce the estimator.

Starting with the jump detection, Baruník and Vacha (2015) utilize the methodology proposed by Fan and Wang (2007), who use the wavelet jump detection to the deterministic functions with *i.i.d.* additive noise ϵ_t of Wang (1995). For the estimation of the jump location, the universal threshold of Donoho and Johnstone (1994) on

the 1st level wavelet coefficients of y_t over $[t - h, t]$, $\tilde{W}_{1,k}$ is used.³ If for some $\tilde{W}_{1,k}$, $|\tilde{W}_{1,k}| > d\sqrt{2 \log n}$, then $\hat{\tau}_t = \{k\}$ is the estimated jump location with size $\bar{y}_{\hat{\tau}_t+} - \bar{y}_{\hat{\tau}_t-}$ (averages over $[\hat{\tau}_t, \hat{\tau}_t + \delta_n]$ and $[\hat{\tau}_t, \hat{\tau}_t - \delta_n]$, respectively, with $\delta_n > 0$ being the small neighbourhood of the estimated jump location $\hat{\tau}_t \pm \delta_n$) and where d is the median absolute deviation estimator defined as $(2^{1/2}) \text{median}(|\tilde{W}_{1,k}|, k = 1, \dots, n)/0.6745$, for more details see Percival and Walden (2000).

Using the result of Fan and Wang (2007), the jump variation is then estimated by the sum of the squares of all the estimated jump sizes:

$$\hat{J}V_{t,h}^W = \sum_{l=1}^{N_t} (\bar{y}_{t,h,\hat{\tau}_{t+}} - \bar{y}_{t,h,\hat{\tau}_{t-}})^2. \tag{5}$$

Thus, we are able to estimate the jump variation from the process consistently with the convergence rate $N^{-1/4}$.

Following Baruník and Vacha (2015), we define the jump-adjusted wavelet two-scale realized variance (JWTSRV) estimator over $[t - h, t]$, for $0 \leq h \leq t \leq T$, on the observed jump-adjusted data, $y_{t,h}^{(j)} = y_{t,h} - \sum_{l=1}^{N_t} J_l$ as:

$$\hat{RV}_{t,h}^{(JWTSRV)} = \sum_{j=1}^{m+1} \hat{RV}_{j,t,h}^{(JWTSRV)} = \sum_{j=1}^{m+1} \left(\hat{RV}_{j,t,h}^{(WJ)} - \frac{\bar{N}}{N} \hat{RV}_{j,t,h}^{(WRVJ)} \right), \tag{6}$$

where $\hat{RV}_{j,t,h}^{(WJ)} = \frac{1}{G} \sum_{g=1}^G \sum_{k=1}^N \mathcal{W}_{j,t-h+\frac{k}{N}h}^2$ is obtained from wavelet coefficient estimates on a grid of size $\bar{N} = N/G$ and $\hat{RV}_{j,t,h}^{(WRVJ)} = \sum_{k=1}^N \mathcal{W}_{j,t-h+\frac{k}{N}h}^2$ is the wavelet realized variance estimator at a scale j on the jump-adjusted observed data, $y_{t,h}^{(j)}$.

The JWTSRV estimator decomposes the realized variance into an arbitrarily chosen number of investment horizons and jumps. Baruník and Vacha (2015) discuss that it is a consistent estimator of the integrated variance as it converges in probability to the integrated variance of the process p_t with the speed $N^{-1/6}$ inherited by the TSRV structure of the estimator. The estimator will also have limiting distribution of the TSRV. Baruník and Vacha (2015) test the small sample performance of the estimator in a large Monte Carlo study, and they find that it is able to recover true integrated variance from the noisy process with jumps precisely.

2.2. Model-free implied volatility (MFIV) and corridor implied volatility (CIV)

While realized volatility measures the volatility from the high frequency returns, model-free implied volatility (MFIV) can be used to infer the volatility from the option prices. This approach, derived by Britten-Jones and Neuberger (2000), uses a cross-section of option prices to calculate the volatility as the risk-neutral expected sum of squared returns between two dates. The resulting implied volatility does not depend on any parametric model and provides ex-ante risk-neutral expectations of the future volatilities. The most serious forerunner of MFIV was the volatility inverted from the Black and Scholes (1973) option pricing formula. Nevertheless, it has been proven that Black–Scholes implied that volatility featured a notari-ally known moneyness bias, known as volatility smile or smirk (Macbeth and Merville (1979) which were amongst the first studies to describe this issue).

Britten-Jones and Neuberger (2000) derived the model-free implied volatility under the diffusion assumption. They extended the

³ Not to distract the reader from the main text, we provide the necessary introduction to wavelet analysis in the Appendix 6.

work of [Derman and Kani \(1994\)](#) and [Dupire \(1994\)](#) to infer a forecast of underlying asset's volatility from a continuum of European call options with strikes and maturities ranging from zero to infinity. The complete set of option prices is used to extract a condition that characterizes the set of continuous processes consistent with current option prices.

In the option setting, current time is fixed to $t = 0$, pay-off occurs at a future fixed date T , and time to maturity is denoted as $\tau = T - t$. For $0 \leq t \leq T$, F_t denotes the time t value of the futures contract expiring at date T . Prices of European put and call options with strike K and expiration date T are given by $P_t(K)$ and $C_t(K)$, and the risk-free rate is assumed to be zero.

The authors first derive the risk-neutral probability of the stock price that is fully determined by the initial set of option prices. Consequently, they show that the set of initial option prices also determine the probability of the stock price reaching any two price levels at two consecutive dates. [Jiang and Tian \(2005\)](#) further extend [Britten-Jones and Neuberger \(2000\)](#) volatility measure to all martingale asset price processes and express the formulae in forward rather than spot prices. Because a martingale can be decomposed canonically into the orthogonal sum of a purely continuous martingale and a purely discontinuous martingale ([Jacod and Shirayev, 1987](#); [Protter, 1990](#)), the model-free relationship between the asset return variance and option prices also holds when asset prices contain jumps.

A forecast of integrated variance for the period $[0, T]$ can be determined from observed European call option prices with maturity T as follows:

$$E_0^F \left[\int_0^T \sigma_s^2 ds \right] = 2 \int_0^\infty \frac{C_t(T, K) - \max(0, F_0 - K)}{K^2} dK \quad (7)$$

where E_t^F denotes the time t expectation with respect to the risk-neutral distribution (RND) of the asset price, T denotes the expiration date, F the forward probability measure, K the strike price and F_t and $C_t(T, K)$ are forward asset price and forward option price, respectively. To calculate the integral, the numerical integration can be used, e.g., trapezoidal rule:

$$\int_{K_{\min}}^{K_{\max}} \frac{C_t(T, K) - \max(0, F_0 - K)}{K^2} dK \approx \sum_{i=1}^m [g(T, K_i) + g(T, K_{i-1})] \Delta K, \quad (8)$$

where $\Delta K = (K_{\max} - K_{\min})/m$, $K_i = K_{\min} + i\Delta K$, and $g(T, K_i) = [C_t(T, K_i) - \max(0, F_0 - K_i)]/K_i^2$.

[Jiang and Tian \(2005\)](#) further developed a simple method to implement the MFIV on observed option prices, based on examination of implementation issues. They identify two types of errors associated with implementation: truncation and discretization errors. Truncation errors are present when tails in the RND are ignored (due to the limited availability of the strike prices for listed options). The authors find that truncation errors are negligible if RND is truncated at two standard deviations from F_0 and propose the flat extrapolation scheme for the range of strike prices outside the available set of the prices. In their later work [Jiang and Tian \(2007\)](#), the authors propose to impose smooth pasting condition at the minimum and maximum of available strike prices to avoid kinks in the implied volatility function at the lower and upper price bounds. We discuss both schemes later on in the data section. Discretization errors are minimized using interpolation between listed strike prices. Cubic spline interpolation is not applied directly to option prices, as there is nonlinear relationship between the option prices and option strike prices. Implied volatilities are

obtained from [Black and Scholes \(1973\)](#) formula, a smooth function is fitted to implied volatilities, and using [Black and Scholes \(1973\)](#) formula, volatilities are again translated to option prices at the desired strike prices. The [Black and Scholes \(1973\)](#) model works as a one-to-one mapping between volatilities and option prices and does not impose any model dependency on the calculation of model-free volatility.

[Andersen and Bondarenko \(2007\)](#) and [Andersen et al. \(2015\)](#) argue that the lack of the liquid contract results in non-trivial measurement errors that are amplified by the stochastic nature of the availability of strike prices that vary over time. The authors propose a corridor implied volatility measure following the work of [Carr and Madan \(1998b\)](#), with a cut-off criterion that is determined endogenously by option prices (part of the estimated risk neutral density inferred from option prices). As such, it allows for reflecting the pricing of volatility across an economically equivalent fraction of the strike range. This would ensure inter-temporal coherence of the measure. The authors further show that due to the lack of availability of strike prices, the implementation of MFIV brings the resulting implied volatilities to corridor-implied volatilities, rather than model-free implied volatilities.

By defining two positive barriers, the lower B_1 , the upper B_2 and the following indicator function,

$$I_t(B_1, B_2) = I_t = 1 [B_1 \leq F_t \leq B_2], \quad (9)$$

corridor integrated variance is

$$CIV AR_t(B_1, B_2) = \int_0^T \sigma_s^2 I_s(B_1, B_2) ds. \quad (10)$$

In comparison to the pure integrated variance, now the return variation is accumulated only when the futures price at time t is between the two barrier levels. [Carr and Madan \(1998b\)](#) and [Andersen and Bondarenko \(2007\)](#) demonstrate that risk-neutral expectation of future corridor implied variance (CIV) at time $t = 0$ is

$$E_0^F \left[\int_0^T \sigma_s^2 I_s(B_1, B_2) ds \right] = 2 \int_{B_1}^{B_2} \frac{C_t(T, K) - \max(0, F_0 - K)}{K^2} dK \quad (11)$$

when $B_1 = 0$ and $B_2 = \infty$, the corridor integrated variance is equivalent to model-free implied variance. Expected future volatility can be obtained by taking square roots of CIV as well as MFIV.

3. Data

To calculate option-implied volatility, we use the set of European-style option prices on the German DAX and U.S. S&P 500 stock market indices. The cross-sectional data contain daily option prices for all listed maturities and strike prices and cover the period from July 2006 to October 2010. The data have been provided by the Option-Metrics database. We use settlement mid prices for the analysis. Settlement prices have the advantage over close prices, as they are not plagued by nonsynchronous trading. Following [Bakshi et al., \(1997\)](#), we apply exclusions filters on the datasets to prevent liquidity related bias and mitigate the impact of price discreteness. First, options with less than one week to expiration are excluded. Second, price quotes lower than 0.375 are excluded. Third, the quotes that do not satisfy the no-arbitrage condition $C \geq \max(0, S_t - X_t)$ for calls and $P \geq \max(0, X_t - S_t)$ for puts are dropped as well. We include only out-of-the money options where the strike price is strictly higher than the spot price for call options and vice versa for put options. This approach is common in all comparable studies, as

in-the-money options are less liquid and thus introduce bias into the calculation of implied volatility. Last, we calculate the measures if and only if more than five options for the given maturity and different strikes passed the above-mentioned criteria. Altogether, we discarded 27% of option prices and used, on average, 80 options per day, per maturity, per index.

To address the fact that measures work with the spot price and options are written on index forwards, we follow the literature and translate spot prices into forward rates using zero coupon rates for given currencies and maturities for each day. In case we do not have zero coupons for the given maturity, we interpolate (extrapolate) between the two nearest rates available. Zero coupon rates are again obtained from the OptionMetrics database. To obtain forward prices, spot prices are multiplied by $e^{r_f(T-t)}$, where r_f is the risk-free rate corresponding the maturity of options.

To calculate the model-free measures, we used the datasets with fixed time to maturity; specifically, we use 15 and 28 calendar days for DAX and 15 and 29 for S&P 500. Note that the one-day difference is caused by different listing conditions for the index options. We refer to the different time to maturity as monthly and bi-weekly maturity further in the text. We use non-overlapping datasets using the data from the trading days when options with fixed time to maturity have been available, so we do not introduce any autocorrelation into the forecasting regressions. When dealing with implementation of model-free implied volatility on real option datasets, one has to inevitably resort to an approximation method introducing potential bias. Jiang and Tian (2005) note some practical limitations associated with the implementation of MFIV, which results in truncation errors when the tails of distribution are ignored, and discretization errors due to numerical integration and limited availability of strike prices range. The authors as well propose the procedure that improve the implementation and render more or less negligible errors.

In our calculations, we follow part of this procedure. To limit the discretization errors, we use a step of one unit of the index to calculate the integral numerically. The trapezoidal rule is applied to obtain the variance (correspondingly, square root to obtain volatility). To overcome the limited availability of strike prices, we infer the prices of absent options from the prices of listed options. Due to nonlinearities in the option prices, we apply the cubic spline curve-fitting method to implied volatilities, instead of option prices directly. The implied volatilities are reverted from listed option prices using the Black and Scholes (1973) pricing formula, and a smooth function is fitted to them. Implied volatilities for absent strikes are then extracted, and the same pricing formula is used again to calculate the corresponding option prices. The pricing formula is used only as a one-to-one mapping between option prices and strike prices. Thus, the procedure retains its model-free grounds.

The subsequent procedure differs for MFIV and CIV calculation. We impose flat extrapolation to the prices beyond the available daily price range in MFIV measure implementation. Jiang and Tian (2007) adjust the slope of the extrapolated segment to match the corresponding slope of the interior segment at the extremes. We did not apply this approach, as it actually rendered much worse results. The number of listed strikes that passed all the abovementioned criteria varies between the days. In some cases, the slope at the interior segment would send the extrapolated prices to unrealistic numbers, thus introducing a large bias into the calculated volatility. To avoid truncation errors, we use one standard deviation from forward prices as an integration range. We do not follow Jiang and Tian (2005) the recommended approach to set the truncation point at two standard deviations from F_0 for simple reasons: the lack of available options for some days and the high volatility of index prices through the time period that involves financial crisis. The interval of strike prices that then needs to be extrapolated $[F_0 - 2SD, K_{\min}] \cup [K_{\max}, F_0 + 2SD]$ becomes too large and accounts for the majority of the inputs that

enters the calculation formula, which necessarily introduces a great amount of error into the calculation. We calculate CIV with the corridors covering the range from the 5th to 95th percentile (CIV1) and from the 2.5th to 97.5th percentile (CIV2) of the RND, estimated from the available option prices that passed the exclusion criteria for the given daily maturity.

Corresponding to each of the implied volatilities, we compute the realized volatility (RV) using 5-minute returns and JWTSRV using all the available data. Tick by tick data were provided by the Tick Data. After computing the daily realized measures, we aggregate them into monthly and bi-weekly groups, according to the maturities, in order to avoid introducing the bias from the maturity mismatch.

4. Time-frequency dynamics in volatility

The information content of implied volatility is typically assessed in the literature by estimating the following regression:

$$RV_{t+h} = \alpha + \beta IV_t + \epsilon_t, \quad (12)$$

with ordinary least squares assuming ϵ_t to follow *i.i.d.* errors with zero mean and finite constant variance. RV_{t+h} is the ex-post realized volatility for period $t+h$, and IV_t is the implied volatility at the beginning of period t , being an ex-ante measure of $t+h$ volatility. In case implied volatility is an unbiased forecast of future realized volatility, α should not be significantly different from 0, and β should not be significantly different from 1. In case implied volatility is efficient, the residuals ϵ_t should be zero mean, constant and finite variance, and they should be serially uncorrelated. It is important to note that $RV_{t+h} - IV_t$ is the return to buying variance in a variance swap contract. Therefore, residuals from the regression ϵ_t can be interpreted as conditional variance risk premium (Carr and Wu, 2009). Similar regressions are common to obtain variance risk premium in the literature (Bekaert and Hoerova, 2014).

The initial literature has generally found that implied volatility is a biased forecast of the future realized volatility, while β is significantly different from unity – see, for example, Christensen and Prabhala (1998). Few researchers (Bandi and Perron, 2006; Christensen and Nielsen, 2006) suggest that the implied–realized volatility relation might be a fractional cointegration relationship, as volatility is typically found to be a long memory process. In this case, ϵ_t would not be integrated on order $I(0)$, and standard OLS should not be used. Before proceeding further with the cointegrating relationship, we study the time–frequency dynamics of the relationship using wavelet coherence to determine how the dependence varies over different frequencies. This will provide an important insight for further analysis.

4.1. Dynamic dependence: a wavelet coherence

To better understand the relationship, it is useful to look at it from the point of view of different frequencies. Here, the wavelet analysis may be well utilized, as it allows for studying the time series in the time–frequency domain. As wavelet coefficients estimate the spectrum of the time series, wavelet coherence can be seen as the estimate of the cross-spectra between two series scaled by the spectra of both series. The coherence is analogous to the square of the correlation between two series. Zero coherence suggests that there is no relation; when coherence equals one, we have perfect correlation. The main advantage of this approach is that it provides us with the localized correlation at time–frequency domain. Such “lenses” into dependence between economic variables have been used recently by many researchers in various

fields (Vacha and Baruník, 2012; Aloui and Hkiri, 2014; Aguiar-Conraria and Soares, 2014). This approach is extremely useful as a first step in our analysis, as the period we study includes a turbulent crisis period. Hence, wavelet coherence will reveal possible structural breaks in the dependence, which we will need to focus on later.

The wavelet transform offers localized frequency decomposition, providing information about frequency components. As a result, wavelets have significant advantages over basic Fourier analysis when the object under study is locally stationary and inhomogeneous — see Gençay et al. (2002), Percival and Walden (2000) and Ramsay (2002).

To be able to study the interaction between two time series, we use a bivariate framework of wavelet coherence.⁴ Following Torrence and Compo (1998), we define the cross wavelet transform of two time series RV_{t+h} and IV_t as

$$W_{IV_t RV_{t+h}}(u, s) = W_{IV_t}(u, s) W_{RV_{t+h}}^*(u, s), \quad (13)$$

where $W_{IV_t}(u, s)$ and $W_{RV_{t+h}}(u, s)$ are continuous wavelet transforms of RV_{t+h} and IV_t , respectively, u is a position index, and s denotes the scale, while the symbol $*$ denotes a complex conjugate. The cross wavelet power can easily be computed using the cross wavelet transform as $|W_{IV_t RV_{t+h}}(u, s)|$. The cross wavelet power reveals areas in the time–frequency space where the time series show a high common power, i.e., it represents the local covariance between the time series at each scale.

The wavelet coherence can detect regions in the time–frequency space where the examined time series co-move but do not necessarily have a high common power. Following the approach of Torrence and Webster (1999), we define the squared wavelet coherence coefficient as:

$$R^2(u, s) = \frac{|S(s^{-1} W_{IV_t RV_{t+h}}(u, s))|^2}{S(s^{-1} |W_{IV_t}(u, s)|^2) S(s^{-1} |W_{RV_{t+h}}(u, s)|^2)}, \quad (14)$$

where S is a smoothing operator.⁵ The squared wavelet coherence coefficient is in the range $0 \leq R^2(u, s) \leq 1$. Values close to zero indicate a weak correlation, while values close to one provide evidence of a strong correlation. Hence, the squared wavelet coherence measures the local linear correlation between two stationary time series at each scale and is analogous to the squared correlation coefficient in linear regression. Because the theoretical distribution for the wavelet coherence is not known, the statistical significance of dependence is tested using Monte Carlo methods (Grinsted et al., 2004; Torrence and Compo, 1998).⁶

Finally, wavelet coherence phase differences may be used to assess the details about the delays in the oscillation (cycles) between

the two time series under study (see Torrence and Webster (1999) for the details). The phase is indicated by arrows on the wavelet coherence plots. A zero phase difference means that the examined time series move together. The arrows point to the right (left) when the time series are in-phase (anti-phase) or are positively (negatively) correlated. Arrows pointing up mean that the first time series leads the second by 90° , whereas arrows pointing down indicate that the second time series leads the first by 90° . Usually, we have a mixture of positions; for example, an arrow pointing up and right means that the time series are in phase, with the first time series leading the second.

Fig. A1 brings the wavelet coherence plots of the two variance series. We use CIV1, CIV2 as well as MFIV to measure volatility implied by options and RV and JWTSRV robust to noise and jumps to measure the realized volatility. Moreover, for each index, we use options with monthly and bi-weekly maturity to study the difference. Realized volatility is computed correspondingly.

Dynamic dependence reveals interesting findings. A distinct change in the general pattern can be found at the 2^5 frequency corresponding to 32 days, or approximately 1.5 months when 21 trading days is considered in one month. In the investment horizons less than 1.5 months, no dependence is found, while the wavelet coherence is significant through horizons longer than 1.5 months and all time periods considered. Thus, the effects of market frictions and short-run fluctuations disappear in the long run, and the dynamic relationship between the variances is nearly perfect in the long run for the entire studied period. In the long run (low frequencies), coherence close to one implies that implied volatility is an unbiased forecast of future realized volatility, and no forecast error or premium for bearing volatility risk exists. In the short run, this equilibrium is broken, and zero coherence implies that most of the changes in implied volatility is coming from the risk premium or errors in future expectations.

Interestingly, options with different maturities used to calculate the implied volatility of S&P 500 do not bring any difference into the relationship. The implied–realized volatility relationship is very strong for all investment horizons longer than 1.5 months and time periods without exception. The situation is similar in the DAX. The only difference is that the relation appears to break in the last years of the sample even with long horizons. This may suggest that the German option market is not as efficient as that of the U.S.

An important distinction can be seen from the wavelet coherence plots when we consider the different measures of volatility used. When CIV measures are used to calculate the implied volatility, the long-term relationship is much stronger than in the case when MFIV is used. This may suggest that MFIV provides a biased measure of implied volatility, as the long run dynamics of the relationship is not so pronounced.

A final observation can be made when looking at the phases (arrows in the plot) that point down to the right. This means that implied volatility generally leads the future realized volatility, which is expected in case implied volatility provides an efficient expectation about future volatility.

While wavelet coherence plots provide us the “lenses” into the implied–realized relationship, in the next sections, we will develop a rigorous methodology to estimate the long-term fractional cointegration relationship using wavelets. Our main motivation in doing so is that wavelets are capable of dealing with non-stationary time series, which will become a crucial property for the analysis.

4.2. Fractional cointegration in variances

Fractional integration provides a framework for studying long-run dependencies in economic time series (Baillie, 1996a). A

⁴ In our work, we use continuous wavelet analysis tools. For any interested reader, we include the necessary introduction to wavelet analysis in the Appendix 6

⁵ Without smoothing, the wavelet coherence equals one at all scales. Smoothing is achieved by convolution in both time and scale. The time convolution is performed with a Gaussian window, while the scale convolution is conducted with a rectangular window — see Grinsted et al. (2004).

⁶ The use of wavelets brings with it the difficulty of dealing with boundary conditions on a dataset with finite length. This is a common problem with any transformation relying on filters. In our paper, we address this problem by padding the time series with a sufficient number of zeroes. The area where the errors caused by discontinuities in the wavelet transform cannot be ignored, i.e., where edge effects become important, is called the cone of influence. The cone of influence is highly dependent on the type of wavelet used — see Torrence and Compo (1998). The cone of influence lies under a cone that is bordered by a thin black line.

stationary time series y_t is said to be fractionally integrated of order $d \in (0, 0.5)$, $I(d)$, if

$$\Delta^d y_t = \epsilon_t, \tag{15}$$

where ϵ_t is an $I(0)$ process, and $\Delta^d = (1 - L)^d$ is the fractional difference operator. The empirical evidence suggests that financial market volatilities are well described by the $I(d)$ processes Andersen et al. (2001, 2003) and Christensen and Nielsen (2006), and long memory quickly became one of the key stylized facts about volatility Cont (2001). Naturally, we can expect that the implied and realized variance series will be tied together in the long-run relationship in the form of fractional cointegration, and a linear combination of the two will be integrated into an order lower than $I(d)$. Thus, the difference between the implied and realized variances, the variance risk premium, should be less persistent than the two individual variance series. This result has been documented by Christensen and Nielsen (2006). Interestingly, Bandi and Perron (2006) and Kellard et al. (2010) report a fractional order of volatility in a non-stationary region when $1/2 < d < 1$, although it is difficult to determine the integration order of the fractional variables, as a smooth transition exists between the stationary and non-stationary regions (Marinucci and Robinson, 2001).

When looking at the regression Eq. (12), for $\alpha = 0$ and $\beta = 1$, residuals ϵ_t clearly decreases to a variance risk premium. In case ϵ_t is an $I(d_u)$ with $d_u < d$, we may suspect a fractional cointegrating relationship between the implied and realized variances.

4.3. Band spectral regression approach

Part of the literature proposes to use a band spectral regression to estimate a fractionally cointegrating relationship in implied and realized volatilities, as OLS estimates of β are inconsistent. Robinson and Marinucci (2003); Christensen and Nielsen (2006) have shown that narrow band least squares (NBLS) result in an estimator that is consistent and normally distributed. The basic idea is transforming the time series into the frequency domain using Fourier transforms and estimating β on the narrow band of the spectrum (Fourier coefficients) not far from the zero frequency on the long memory region. Recently, Nielsen and Frederiksen (2011) generalize this idea to a fully modified NBLS (FMNBLS), which is able to address the bias introduced by correlation between regressors and errors. In this paper, we build on these ideas, but use a rather different approach of band least squares on the spectra estimated on wavelet coefficients.

Let us introduce the approach by considering the regression model

$$y_t = x_t \beta + \epsilon_t, \tag{16}$$

where $\{x_t, t = 1, \dots, T\}$, $\{y_t, t = 1, \dots, T\}$ and $\epsilon_t \sim N(0, \sigma^2)$. The OLS estimator of β is

$$\hat{\beta}^{OLS} = (x^T x)^{-1} x^T y = \frac{\text{Cov}(x_t, y_t)}{\text{Var}(x_t)}. \tag{17}$$

Engle (1974) was among the first to consider estimation of β in the frequency domain. In fact, the frequency domain is very intuitive, as the variance and covariance are the spectrum and co-spectrum of the series and can be simply estimated, for example, using Fourier transforms.

Recently, Kellard et al. (2010) finds that realized as well implied volatility series may lie in the non-stationary region when $1/2 <$

$d < 1$. Frequency domain least squares using the Fourier transform are able to accommodate non-stationary fractional cointegration by transforming potentially non-stationary series x_t , which are $I(d)$ with $d > 1/2$ using $\gamma \geq 0$ into the resulting $\Delta^\gamma x_t$, which are $I(d - \gamma)$ (Nielsen and Frederiksen, 2011). The choice of γ affects the estimation procedure, and different choices will lead to different estimators. The authors propose the best choice of $\gamma = d_u$, where d_u is the memory parameter of the residuals which can be estimated.

Frequency domain least squares based on the wavelet estimation of spectra are able to address this problem, as wavelets are generally a very convenient tool in case we are dealing with the non-stationary series (Fan and Whitcher, 2003; Roueff and Sachs, 2011). Although wavelets do not improve the estimation of d in the standard stationary context $d < 1/2$, Faÿ et al. (2009) showed that in the presence of trends, or series with $d \geq 1/2$ and $d \leq -1/2$, they are helpful because they allow differencing implicitly.

4.4. Wavelet band spectral regression (WBLS)

Using wavelet transform, we are able to divide the entire frequency spectrum into frequency bands represented by wavelet scales j . After the transform, the resulting spectrum on the j -th scale has the following form: $f_j^W \in [1/2^{j+1}, 1/2^j]$. The wavelet spectral density function at a scale j can be expressed as $S_{(x)j}(f) = \mathcal{H}_j(f) S_x(f)$ where $\mathcal{H}_j(f)$ is the transfer function of the wavelet filter at a scale j , and $S_x(f)$ denotes the spectrum of x_t . Similarly, the wavelet cross-spectrum at a scale j is defined as $S_{(xy)j}(f) = \mathcal{H}_j(f) S_{xy}(f)$, where $S_{(xy)j}(f)$ represents the cross-spectrum of x_t and y_t . Furthermore, wavelet variance $\nu_x^2(j)$ and wavelet covariance $\gamma_{xy}(j)$ at a scale j reads:

$$\nu_x^2(j) = \int_{-1/2}^{1/2} S_{(x)j}(f) df = \int_{-1/2}^{1/2} \mathcal{H}_j(f) S_x(f) df, \tag{18}$$

$$\gamma_{xy}(j) = \int_{-1/2}^{1/2} S_{(xy)j}(f) df = \int_{-1/2}^{1/2} \mathcal{H}_j(f) S_{xy}(f) df. \tag{19}$$

In case $T \rightarrow \infty$, and therefore, the maximum number of wavelet scales $J \rightarrow \infty$ is available, we can write total variance and covariance as a sum of wavelet variances and covariances at all scales as (Whitcher et al., 1999):

$$\text{Var}(x_t) = (x^T x) = \int_{-1/2}^{1/2} S_{(x)j}(f) df = \sum_{j=1}^{\infty} \nu_x^2(j) \tag{20}$$

$$\text{Cov}(x_t, y_t) = (x^T y) = \int_{-1/2}^{1/2} S_{(xy)j}(f) df = \sum_{j=1}^{\infty} \gamma_{xy}(j). \tag{21}$$

Thus, β can be estimated in the frequency domain using the wavelet least square estimator (WLS) as follows

$$\hat{\beta}^{WLS}(1, \infty) = \left(\sum_{j=1}^{\infty} \nu_x^2(j) \right)^{-1} \left(\sum_{j=1}^{\infty} \gamma_{xy}(j) \right). \tag{22}$$

Asymptotically, $\hat{\beta}^{WLS}$ is equal to $\hat{\beta}^{OLS}$. In many situations, time series carry different information in the low and high part of the spectra. It can be viewed as estimating the simple linear

regression on the different time horizons j obtained from the wavelet transform

$$y_t^{(j)} = x_t^{(j)}\beta + \epsilon_t^{(j)}. \quad (23)$$

As revealed by the wavelet coherence, this is also the case for the volatility relationship. As the relationship comes solely from the long run part of the spectra, we need a tool that will be able to estimate the relation only on this part. Similarly to the NBLS and FMNBLS (Robinson and Marinucci, 2003; Christensen and Nielsen, 2006; Nielsen and Frederiksen, 2011), which we introduce later in the text, as we use it for comparison with our estimator, we can obtain the estimate on the narrow band of the spectrum not far from the zero frequency on the long memory region. More precisely, we can use only scales j , which cover the long memory region.

Our final estimator, the wavelet band least square estimator, simply estimates the β on the band of scales $j \in [k, l]$ in Eq. (22), thus using frequency band $f \in [1/2^{l+1}, 1/2^k]$. For the estimation of spectra, we use modified discrete wavelet transforms (MODWT).⁷ The WBLS estimator is then

$$\hat{\beta}^{WBLS}(k, l) = \left(\sum_{j=k}^l v_x^2(j) \right)^{-1} \left(\sum_{j=k}^l \gamma_{xy}(j) \right) \quad (24)$$

where j -th scale represents the frequency band $f \in [1/2^{j+1}, 1/2^j]$. For example $\hat{\beta}^{WBLS}(3, 4)$ will estimate β over the frequency band of $f \in [1/2^5, 1/2^3]$. The estimator in Eq. (24) can be expressed in terms of the MODWT coefficients $\tilde{W}_x(j, s)$ and $\tilde{W}_y(j, s)$, where j and s denote scale position of the transform for x_t and y_t as

$$\hat{\beta}^{WBLS}(k, l) = \left(\sum_{j=k}^l \left[\frac{1}{T} \sum_{u=1}^T \tilde{W}_x^2(j, s) \right] \right)^{-1} \left(\sum_{j=k}^l \left[\frac{1}{T} \sum_{u=1}^T \tilde{W}_x(j, s) \tilde{W}_y(j, s) \right] \right). \quad (25)$$

By WBLS, we can focus on estimating the long memory part of the spectra using the frequency band near the origin and obtain the long memory relationship. Thus, the information content of the implied volatility can simply be assessed by estimating the relationship on the different bands of spectra as

$$RV_{t+h}^{(j)} = \alpha + \beta IV_t^{(j)} + \epsilon_t^{(j)}, \quad (26)$$

when we use all scales j and as $j \rightarrow \infty$, $\hat{\beta}^{WBLS}$ will be equivalent to $\hat{\beta}^{WLS}$ and will converge on an OLS estimator. The properties of this type of regression in the long memory setting, together with the limiting distributions of the estimates, are studied in Fadili and Bullmore (2002).

4.5. FMNBLS

For comparison with the newly proposed WBLS estimator of the fractional cointegrating implied-realized relationship, we use the frequency domain least squares methods, which is well established in the literature. The basic distinction from the WBLS

is that instead of using wavelet coefficients to estimate the spectra and co-spectra, a Fourier transform is used by the rest of the literature.

The basic idea of the narrow band least squares (NBLS) estimator is to transform the time series into the frequency domain using Fourier transforms and estimating β on the narrow band of the spectrum not far from the zero frequency on the long memory region. Robinson and Marinucci (2003) and Christensen and Nielsen (2006) have shown that narrow band least squares (NBLS) results in an estimator that is consistent and normally distributed. Averaged (co-) cross-periodogram used for the estimation of spectrum is $\hat{F}_{xy}(k, l) = 2\pi/T \sum_{j=k}^l I_{xy}(\lambda_j)$ for any $0 \leq k \leq l \leq T-1$ and for $\lambda_j = 2\pi j/T$, where $I_{xy}(\lambda_j) = 1/2\pi T \sum_{t=1}^T \sum_{s=1}^T x_t y_s e^{-i(t-s)\lambda}$ is cross-periodogram, or estimated cross-spectrum between two series on a specific frequency band $[k, l]$. Analogously, $\hat{F}_x(k, l)$ is estimated spectrum of x_t . Then, the cointegrating relation between two time series $\{x_t\}$ and $\{y_t\}$ can be estimated as

$$\hat{\beta}^{NBLS}(k, l) = \hat{F}_x^{-1}(k, l) \hat{F}_{xy}(k, l), \quad (27)$$

where k and l define the frequency band used for the estimation of β .

By definition, $\hat{\beta}^{NBLS}(1, T-1)$ is algebraically identical to the usual OLS estimator of β and thus identical to the WLS estimator in Eq. (22). If $\frac{1}{T} + \frac{1}{T} \rightarrow 0$ as $T \rightarrow \infty$, $\hat{\beta}^{NBLS}(k, l)$ is an NBLS estimator using only a degenerating band of frequencies near the origin. While l must tend to infinity to have information, it also needs to remain in a neighbourhood of zero where we have assumed knowledge about the spectral density, so l/T must tend to zero.

Nielsen and Frederiksen (2011) based on their previous work (Christensen and Nielsen, 2006) show that the absence of non-coherence between regressors and errors at zero frequency imposes bias on the NBLS estimates, and they propose a fully modified NBLS estimator to eliminate this bias. The FMNBLS estimator is simply NBLS corrected for the asymptotic bias estimated by running an auxiliary NBLS regression of the (differenced) residuals from the initial NBLS on the same regressors. To keep the text under control, we point any interested reader to the work of Nielsen and Frederiksen (2011) for details of the methodology.

5. Results

The main aim of the paper is to revisit the relationship between implied and future realized volatility using the new unbiased measures of volatilities and newly proposed wavelet band spectral regression.

Wavelet coherence suggests that implied volatility may be an efficient forecast of the future volatility in the long run, while the existence of risk premia makes it an inefficient forecast in the short run. Until now, we have been using the daily data as wavelets that are known for their decorrelation properties, and they can address the non-stationary data as well. Christensen and Prabhala (1998) were the first to note that overlapping data may affect the estimation considerably. To overcome this problem, we aggregate the daily data to monthly (and bi-weekly) non-overlapping data for further analysis.

We begin to study the volatilities by estimating their long memory parameter. Table 1 shows the long memory estimates of different implied and realized volatility measures used in our study. For the estimation, we use the popular GPH method (Geweke and Porter-Hudak, 1983; Robinson, 1995), and we report the estimates up to various different frequencies $T^{0.6}$, $T^{0.7}$ and $T^{0.8}$. Interestingly, nearly all estimated volatilities show memory lying

⁷ For a more detailed treatment, see the Appendix.

in between the stationary and non-stationary region of $d = 0.5$. When implied volatility is measured using corridor-implied volatility (CIV1 and CIV2), it displays larger memory than the conventional MFIV measure constituting the popular VIX index. Both realized volatility measures show similar memory, which is uniformly lower than the memory of implied volatility. These results suggest that choice of the different measures may have serious consequences for the estimation of the relationship, in particular, due to the implied volatility measure. When compared to the literature that uses either model-driven implied volatilities or MFIV, our estimated implied volatilities have larger memory possibly crossing the boundary of stationarity. It is worth noting that the period used in our study covers the recent financial crisis, which could possibly introduce this feature.

As a preliminary step in the estimation of the relationship, we estimate it using OLS. The results are reported together with frequency domain estimates separately for each series. Tables 2, 3, 4 and 5 report the OLS estimates for the S&P 500 using monthly and bi-weekly maturities, and the DAX using monthly and bi-weekly maturities in the implied volatility measurement, respectively. The first rows show the coefficients of the Eq. (12), suggesting that generally, IV is a biased predictor of RV. However, when CIV is used to measure the implied volatility, slope coefficients are strikingly closer to unity than MFIV estimates for both markets and both maturities used. This measurement error may have led researchers to find bias more pronounced than it may be.

Memory estimates of the residuals from the OLS confirm this result. When MFIV is used, residuals are $I(d)$ suggesting implied-realized volatility is a cointegrating relationship. When CIV is used as the implied volatility measure, except for the case of S&P 500 with monthly maturity options, residuals do not display as pronounced long memory as reported in the literature. In all the cases, when residuals show a long memory parameter close to zero, or $I(0)$, the slope coefficient is very close to unity.

While estimated volatilities tend to be at the boundary of stationarity, wavelet band least squares may serve as the best tool for the estimation of the relationship. From the wavelet coherence plots, we can see that there is no relationship until the 32nd period, while the relationship for the periods higher than 32 appears to be nearly perfect. Thus, we can utilize the result and conveniently estimate the relationship only on this part of the spectra using wavelet band least squares.

Tables 2, 3, 4 and 5 report the results for the estimated β using WBLS. We use $j = \{5, 6\}$ levels from the 6 level wavelet decomposition to estimate the relationship. All the β 's are much closer to unity, regardless of the measure used. It is interesting to note that when CIV is used, the relationship is not significantly different from unity in most of the cases, while the CIV1 measure implies coefficients closest to unity than does the CIV2. The only exception is the S&P 500 data, with monthly options where the coefficients are significantly lower than unity. MFIV provide the β , which is always lower than unity. The difference between the RV and JWTSRV measures is not so pronounced, although JSTWRV does bring some improvement in the estimates.⁸ This confirms the result of Martin et al. (2009), who assess the robustness of the relative performance of various estimators to the microstructure noise, and they find that

results are invariant to the method of noise correction in the realized volatility.

For comparison, we also estimate the FMNBLS. While in the WBLS, we have motivated the choice of the bands by the wavelet coherence plots, in the FMNBLS, we follow Nielsen and Frederiksen (2011) and use the same bands to estimate the relationship. Namely, we use $[T^{0.4}, T^{0.6}]$, $[T^{0.4}, T^{0.7}]$, $[T^{0.4}, T^{0.8}]$, $[T^{0.5}, T^{0.6}]$, $[T^{0.5}, T^{0.7}]$ and $[T^{0.5}, T^{0.8}]$ frequencies. All of the bands overlap the bands used in the WBLS, although they interfere also with higher frequencies. The results are statistically similar to those obtained by the WBLS estimates. When CIV measures are used, β is closer to unity (or significantly does not differ from unity) when compared to the usual OLS. When MFIV is used, β is closer to unity, but in most cases, they are significantly lower than unity.

Finally, it is worth noting the difference between monthly and bi-weekly regressions. In the case of monthly regressions, when the CIV1 measure is used, spectral band regressions confirm the long-run unbiasedness of the implied volatility forecasts. Residuals do not display significant long memory; thus, fractional cointegration describes the relationship well. However, the unbiasedness of bi-weekly volatility forecasts is confirmed, but residuals from both WBLS as well as FMNBLS display significant long memory. OLS residuals do not show significant long memory, and β s from OLS are very close to unity.

6. Conclusion

In this paper, we study the long run unbiasedness of implied volatility as a predictor of future volatility. While this relationship has been studied previously in the literature, our work contributes to the findings in several ways. First, we propose new spectral techniques to estimate the potential fractional cointegrating relationship of the implied and realized volatilities based on wavelets. The main advantage in comparison to common spectral regression techniques based on Fourier coefficients is that wavelets allow us to work with locally stationary series. Second, we study the fractional cointegration of the implied and realized volatilities using accurate corridor implied volatility (CIV measure) for the first time, as most of the literature uses model based implied volatilities.

When CIV is used to measure implied volatility on options with monthly maturities, implied volatility is found to be an unbiased forecast of the future realized volatility in the long-term horizon over one month. This result holds for the options on S&P 500 as well as DAX indices. Implied and realized volatilities are confirmed to have a fractionally cointegrating relationship, jointly using our newly proposed wavelet band least squares as well as fully modified narrow band least squares. In contrast, when MFIV is used as a measure of implied volatility, all estimates are lower than unity. This result strongly suggests that the measurement of volatility implied by option prices is crucial for the volatility forecasts, as incorrect measurements introduce bias into the forecasts. The result is also important to the literature, as it may suggest that the estimated bias of option implied volatility forecasts might not be that pronounced, as it is caused by an imprecise measurement.

We also question the importance of measures used on the other side of the regression testing the unbiasedness of implied volatility, namely, realized volatility. Interestingly, we find that the results are invariant to the method of noise correction in the realized volatility.

⁸ As JWTSRV is a relatively new measure, we have also used realized measures well-established in the literature, namely bipower variation (BPV), realized kernels (RK) and the two-scale realized variance (TSRV), and the results are very similar to those reported here. To keep minimize the length of the document, we do not report these results but make them available upon request from the authors.

Appendix A

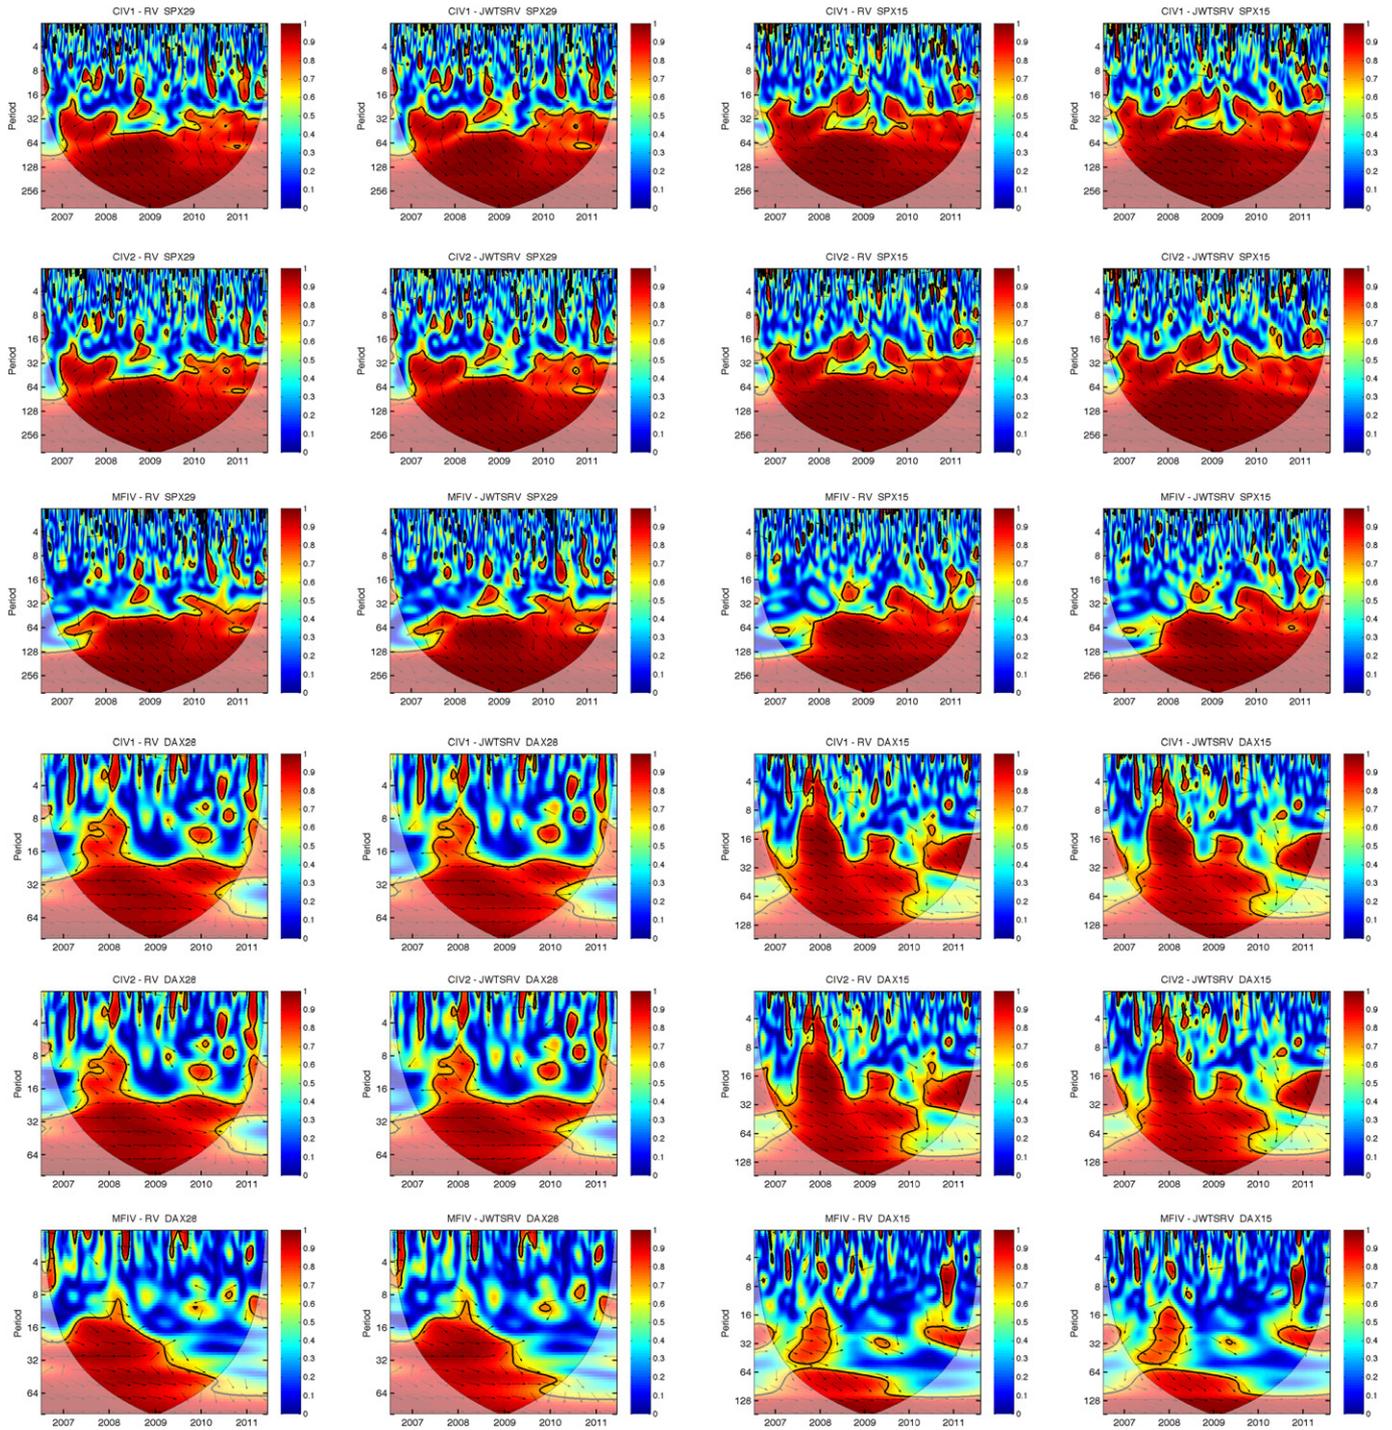


Fig. 1. S&P 500 and DAX monthly and bi-weekly wavelet coherences between implied and realized volatilities measured by CIV1, CIV2, MFIV and RV and JWTSRV, respectively. x-axis shows the studied period extending from July 2006 to October 2010; y-axis shows the periods in days.

Table 1

Long memory estimates of CIV1, CIV2, MFIV measures of implied volatility, RV and JWTSRV measures of realized volatility series. Long memory parameter d is estimated using GPH estimator with different frequency bands $[T^m]$ for $m \in \{0.6, 0.7, 0.8\}$. Standard errors are provided in parentheses.

	CIV1	CIV2	MFIV	RV	JWTSRV
<i>SPX monthly</i>					
$[T^{0.6}]$	0.581 (0.144)	0.573 (0.144)	0.582 (0.144)	0.515 (0.144)	0.534 (0.144)
$[T^{0.7}]$	0.523 (0.115)	0.517 (0.115)	0.546 (0.115)	0.538 (0.115)	0.555 (0.115)
$[T^{0.8}]$	0.641 (0.091)	0.636 (0.091)	0.665 (0.091)	0.776 (0.091)	0.776 (0.091)
<i>SPX bi-weekly</i>					
$[T^{0.6}]$	0.654 (0.144)	0.662 (0.144)	0.494 (0.144)	0.526 (0.144)	0.520 (0.144)
$[T^{0.7}]$	0.782 (0.118)	0.784 (0.118)	0.702 (0.118)	0.693 (0.118)	0.637 (0.118)
$[T^{0.8}]$	0.711 (0.094)	0.683 (0.094)	0.527 (0.094)	0.779 (0.094)	0.752 (0.094)
<i>DAX monthly</i>					
$[T^{0.6}]$	0.626 (0.139)	0.631 (0.139)	0.570 (0.139)	0.448 (0.139)	0.486 (0.139)
$[T^{0.7}]$	0.506 (0.112)	0.509 (0.112)	0.343 (0.112)	0.498 (0.112)	0.531 (0.112)
$[T^{0.8}]$	0.574 (0.091)	0.577 (0.091)	0.300 (0.091)	0.613 (0.091)	0.633 (0.091)
<i>DAX bi-weekly</i>					
$[T^{0.6}]$	0.674 (0.112)	0.680 (0.112)	0.543 (0.112)	0.425 (0.112)	0.441 (0.112)
$[T^{0.7}]$	0.572 (0.087)	0.573 (0.087)	0.371 (0.087)	0.544 (0.087)	0.542 (0.087)
$[T^{0.8}]$	0.700 (0.067)	0.701 (0.067)	0.333 (0.067)	0.689 (0.067)	0.706 (0.067)

Appendix B. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.econmod.2016.01.014>.

References

Aguiar-Conraria, Luís, Soares, Maria,Joana, 2014. The continuous wavelet transform: moving beyond uni-and bivariate analysis. *J. Econ. Surv.* 28 (2), 344–375.

Aloui, Chaker, Hkiri, Besma, 2014. Co-movements of GCC emerging stock markets: new evidence from wavelet coherence analysis. *Econ. Model.* 36, 421–431.

Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *Int. Econ. Rev.* 39 (4), 885–905.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. *J. Financ. Econ.* 61, 43–76.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* (71), 579–625.

Andersen, T.G., Bondarenko, O., 2007. Volatility as an Asset Class. *Risk Books*: London, chap. Constructon and interpretation of model-free implied volatility. 141–181.

M.T.Gonzales-Perez, Andersen, T.G., Bondarenko, O., 2015. Exploring return dynamics via corridor implied volatility. *Rev. Financ. Stud.* 28 (10), 2902–2945.

Baillie, R., 1996. Long memory processes and fractional integration in econometrics. *J. Econ.* 73 (1), 5–59.

Baillie, Richard.T., 1996. Long memory processes and fractional integration in econometrics. *J. Econ.* 73 (1), 5–59.

Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *J. Financ.* 52 (5), 2003–2049.

Bandi, Federico, Perron, Benoit, 2006. Long memory and the relation between implied and realized volatility. *J. Financ. Econ.* 4 (4), 636–670.

Barndorff-Nielsen, O.E., Hansen, P., Lunde, A., Shephard, N., 2008. Designing realized kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica* 76 (6), 1481–1536.

Barndorff-Nielsen, O.E., Shephard, N., 2001. Non-gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *J. R. Stat. Soc. Ser. B* (63), 167–241.

Barndorff-Nielsen, O.E., Shephard, N., 2002. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *J. R. Stat. Soc. Ser. B* (64), 253–280.

Barndorff-Nielsen, O.E., Shephard, N., 2002. Estimating quadratic variation using realized variance. *J. Appl. Econ.* (17), 457–477.

Barndorff-Nielsen, O.E., Shephard, N., 2004. Econometric analysis of realized covariation: high frequency based covariance, regression, and correlation in financial economics. *Econometrica* 72 (3), 885–925.

Barndorff-Nielsen, O.E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. *J. Financ. Econ.* (2), 1–48.

Barndorff-Nielsen, O.E., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. *J. Financ. Econ.* (4), 1–30.

Baruník, J., Krehlik, T., Vacha, L., 2016. Modeling and forecasting exchange rate volatility in time–frequency domain. *Eur. J. Oper. Res.* 251 (1), 329–340.

Baruník, Jozef, Vacha, Lukas, 2015. Realized wavelet-based estimation of integrated variance and jumps in the presence of noise. *Quant. Finan.* 15 (8), 1347–1364.

Bekaert, Geert, Hoerova, Marie, 2014. The VIX, the variance premium and stock market volatility. *J. Econ.* 183 (2), 181–192.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *J. Polit. Econ.* 81 (3), 637–659.

Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. *J. Financ.* 55 (2), 839–866.

Canina, L., Figlewski, S., 1993. The informational content of implied volatility. *Rev. Financ. Stud.* 6, 659–681.

Carr, P., Madan, D.B., 1998. Option valuation using the fast fourier transform. *J. Comput. Financ.* (2), 61–73.

Carr, P., Madan, D., 1998. Towards a theory of volatility trading. *Volatility: New Estimation Techniques for Pricing Derivatives*, R. Jarrow,(ed.), RISK Publications

Carr, Peter, Wu, Liuren, 2009. Variance risk premiums. *Rev. Financ. Stud.* 22 (3), 1311–1341.

Charfeddine, Lanouar, Ajmi, Ahdi.Noomen, 2013. The Tunisian stock market index volatility: long memory vs. switching regime. *Emerg. Mark. Rev.* 16, 170–182.

Chatrath, Arjun, Miao, Hong, Ramchander, Sanjay, Wang, Tianyang, 2015. The forecasting efficacy of risk-neutral moments for crude oil volatility. *J. Forecast.* 34 (3), 177–190.

Christensen, B.J., Hansen, C.H.S., 2002. New evidence on the implied-realized volatility relation. *Eur. J. Financ.* 8 (2), 187–205.

Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. *Journal of Financial Economics* 50 (2), 125–150.

Christensen, B.J., Nielsen, M.O., 2006. Asymptotic normality of narrow-band least squares in the stationary fractional cointegration model and volatility forecasting. *J. Econ.* 133 (1), 343–371. ISSN 0304-4076.

Cont, Rama, 2001. Empirical properties of asset returns: stylized facts and statistical issues.

Day, T., Lewis, C., 1992. Stock market volatility and the information content of stock index options. *J. Econ.* (52), 267–287.

DeGiannakis, Stavros, Livada, Alexandra, 2013. Realized volatility or price range: evidence from a discrete simulation of the continuous time diffusion process. *Ecol. Model.* 30, 212–216.

Derman, E., Kani, I., 1994. Riding on a smile. *RISK* (52), 32–39.

Donoho, D.L., Johnstone, I.M., 1994. Ideal spatial adaptation by wavelet shrinkage. *Biometrica* (81), 425–455.

Dupire, B., 1994. Pricing With a Smile. *RISK* (7), 18–20.

Engle, R.F., 1974. Band spectrum regression. *Int. Econ. Rev.* 15 (1), 1–11.

Fadili, M.J., Bullmore, E.T., 2002. Wavelet-generalized least squares: a new blu estimator of linear regression models with 1/f errors. *NeuroImage* 15 (1), 217–232.

Fan, J., Wang, Y., 2007. Multi-scale jump and volatility analysis for high-frequency financial data. *J. Am. Stat. Assoc.* (102), 1349–1362.

Fan, Y., Whitcher, B., 2003. A wavelet solution to the spurious regression of fractionally differenced processes. *Appl. Stoch. Model. Bus. Ind.* 19 (3), 171–183.

Faÿ, G., Moulines, E., Roueff, F., Taqqu, M.S., 2009. Estimators of long-memory: fourier versus wavelets. *J. Econ.* 151 (2), 159–177.

Gençay, R., Selçuk, F., Whitcher, B., 2002. *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. Academic Press

Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time series models. *J. Time Ser. Anal.* 4, 221–238.

Grinsted, A., Moore, J.C., Jevrejeva, S., 2004. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Process. Geophys.* 11, 561–566.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatility. *J. Financ.* 42, 281–300.

- Hull, Matthew, McGroarty, Frank, 2014. Do emerging markets become more efficient as they develop? Long memory persistence in equity indices. *Emerg. Mark. Rev.* 18, 45–61.
- Jacod, K., Shirayev, A.N., 1987. *Limit theorems for stochastic processes*. Springer-Verlag, Berlin.
- Jiang, George, Tian, Yisong, 2005. The model-free implied volatility and its information content. *Rev. Financ. Stud.* 18 (4), 1305–1342. ISSN 0893-9454.
- Jiang, G.J., Tian, Y.S., 2007. Extracting model-free volatility from option prices: an examination of the vix index. *J. Deriv.* 14, 1–26.
- Jorion, P., 1995. Predicting volatility in the foreign exchange market. *J. Financ.* (50), 507–528.
- Kellard, N., Dunis, Ch, Sarantis, N., 2010. Foreign exchange, fractional cointegration and the implied realized volatility relation. *J. Bank. Financ.* 34 (4), 882–891.
- Kinateder, Harald, Wagner, Niklas, 2014. Multiple-period market risk prediction under long memory: when var is higher than expected. *J. Risk Financ.* 15 (1), 4–32.
- Lamoureux, C.G., Lastrapes, W., 1993. Forecasting stock return variance: towards understanding stochastic implied volatility. *Rev. Financ. Stud.* (6), 293–326.
- Li, Kai, 2002. Long-memory versus option-implied volatility predictions. *J. Deriv.* 9 (3), 9–25.
- Macbeth, J.D., Merville, L.J., 1979. An empirical examination of the Black-Scholes call option pricing model. *J. Financ.* 34 (6), 1173–1186.
- Marinucci, D., Robinson, P.M., 2001. Finite-sample improvements in statistical inference with $I(1)$ processes. *J. Appl. Econ.* 16, 431–444.
- Martin, G.M., Reidy, A., Wright, J., 2009. Does the option market produce superior forecasts of noise-corrected volatility measures? *Journal of Applied Econometrics* 24, 77–104.
- Muzzioli, Silvia, 2013. The forecasting performance of corridor implied volatility in the Italian market. *Comput. Econ.* 41 (3), 359–386.
- Nielsen, M.O., Frederiksen, P., 2011. Fully modified narrow-band least squares estimation of weak fractional cointegration. *Econ. J.* 14, 77–120.
- Padungsaksawasdi, Chaiyuth, Daigler, Robert.T., 2014. The return-implied volatility relation for commodity ETFs. *J. Futur. Mark.* 34 (3), 261–281.
- Percival, D.B., Walden, A.T., 2000. *Wavelet Methods for Time Series Analysis*. Cambridge University Press.
- Phillips, P.C.B., 1991. Spectral regression for cointegrated time series. In: Barnett, W.A., Powell, J., Tauchen, G.E. (Eds.), *Nonparametric and Semiparametric Methods in Econometrics and Statistics: Proceedings of the Fifth International Symposium in Economic Theory and Econometrics*. Cambridge: Cambridge University Press, pp. 413–435.
- Protter, P., 1990. *Stochastic Integration and Differential Equations: A New Approach*. Springer, Berlin.
- Ramsay, J.B., 2002. Wavelets in economics and finance: past and future. *Stud. Nonlinear Dyn. & Econ.* 6 (3).
- Robinson, P.M., Marinucci, D., 2003. *Time Series with Long Memory*. Oxford University Press, Oxford, chap. Semiparametric frequency domain analysis of fractional cointegration.
- Robinson, P.M., 1995. Log-periodogram regression of time series with long range dependence. *Ann. Stat.* 23, 1048–1072.
- Roueff, F., Sachs, R., 2011. Locally stationary long memory estimation. *Stochastic Processes and their Applications* 121 (4), 813–844. ISSN 0304-4149.
- Seo, Sung.Won, Kim, Jun.Sik, 2015. The information content of option-implied information for volatility forecasting with investor sentiment. *J. Bank. Financ.* 50, 106–120.
- Torrence, C., Compo, G.P., 1998. A practical guide to wavelet analysis. *Bull. Am. Meteorol. Soc.* 79 (1), 61–78.
- Torrence, C., Webster, P.J., 1999. Interdecadal changes in the enso-monsoon system. *J. Clim.* 12 (8), 2679–2690.
- Vacha, L., Baruník, J., 2012. Co-movement of energy commodities revisited: evidence from wavelet coherence analysis. *Energy Econ.* 34 (1), 241–247.
- Wang, Y., 1995. Jump and sharp cusp detection via wavelets. *Biometrika* (82), 385–397.
- Whitcher, B., Guttorp, P., Percival, D.B., 1999. Mathematical background for wavelets estimators for cross covariance and cross correlation. *Tech. Rep. 38, Natl. Res. Cent. for stat. and the Environ.*
- Yalama, Abdullah, Celik, Sibel, 2013. Real or spurious long memory characteristics of volatility: empirical evidence from an emerging market. *Econ. Model.* 30, 67–72.
- Zhang, L., Mykland, P.A., Ait-Sahalia, Y., 2005. A tale of two time scales: determining integrated volatility with noisy high frequency data. *J. Am. Stat. Assoc.* (100), 1394–1411.