Journal of Banking & Finance 34 (2010) 882–891

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Foreign exchange, fractional cointegration and the implied–realized volatility relation

Neil Kellard [a,*], Christian Dunis [b], Nicholas Sarantis [c]

[a] Essex Finance Centre, Essex Business School, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom
[b] Centre for International Banking, Economics and Finance (CIBEF), Liverpool Business School, Liverpool John Moores University, John Foster Building, 98, Mount Pleasant, Liverpool L3 5UZ, United Kingdom
[c] Centre for International Capital Markets, London Metropolitan Business School, London Metropolitan University, 84, Moorgate, London EC2M 6SQ, United Kingdom

ARTICLE INFO

Article history:
Received 1 May 2008
Accepted 24 September 2009
Available online 30 September 2009

JEL classification:
F31
G14
C14
C22

Keywords:
Market efficiency
Options markets
Fractional cointegration
Narrow band least squares
Bootstrap
Traded volatility
Intra-day data

ABSTRACT

Almost all relevant literature has characterized implied volatility as a biased predictor of realized volatility. In this paper we provide new time series techniques to investigate the validity of this finding in several foreign exchange options markets, including the Euro market. First, we develop a new fractional cointegration test that is shown to be robust to both stationary and non-stationary regions. Second, we employ both intra-day and daily data to measure realized volatility in order to assess the relevance of data frequency in resolving the bias. Third, we use data on implied volatility traded on the market. In contrast to previous studies, we show that the frequency of data used for measuring realized volatility within a fractionally cointegrating framework is important for the results of unbiasedness tests. Significantly, for many popular exchange rates, the use of intra-day rather than daily data affects the emergence of a different bias, as the possibility of a fractionally integrated risk premium admits itself!

1. Introduction

Market efficiency in options markets is typically examined by estimating the following regression:

$$\sigma^{\mathrm{RV}}_{t+\tau} = \alpha + \beta\sigma^{\mathrm{IV}}_{t} + u_{t+\tau}, \tag{1}$$

where σ^{IV}_{t} is the implied volatility (IV) over a period of time τ and $\sigma^{\mathrm{RV}}_{t+\tau}$ is the realized volatility (RV). Unbiasedness holds in (1) when $\alpha = 0$, $\beta = 1$ and $u_{t+\tau}$ is serially uncorrelated. Of course, unbiasedness is a sufficient condition for market efficiency but is not necessary in the presence of either a constant or a time-varying option market risk premium.

Conventional tests in the previous literature have generally led to the conclusion that IV is a biased forecast of RV in the sense that the slope parameter in (1) is not equal to unity (see, *inter alia*, Christensen and Prabhala, 1998; Poteshman, 2000). This conclusion is found to be robust across a variety of asset markets (see Neely, 2009) and has thus provided the motivation for several attempted explanations of this common finding. Popular suggestions include computing RV with low frequency data (Poteshman, 2000); that the standard estimation with overlapping observations produces inconsistent parameter estimates (Dunis and Keller, 1995; Christensen et al., 2001); and that volatility risk is not priced (Poteshman, 2000; Chernov, 2007). However, Neely (2009) evaluates these possible solutions and finds that the bias in IV is not removed.

Of course, the optimality of the estimation procedure applied to (1) depends critically on the order of integration of the component variables. Given the acknowledged persistence in individual volatility series, the recent literature suggests they are well represented as fractionally integrated processes (see, *inter alia*, Andersen et al., 2001a,b). Notably Bandi and Perron (2006), Christensen and Nielson (2006), and Nielsen (2007) have begun to examine the consequences of this approach for regression (1).

Employing stock market data, Bandi and Perron (2006), Christensen and Nielson (2006), and Nielsen (2007) suggest that

* Corresponding author. Tel.: +44 1206 874153; fax: +44 1206 873429.
E-mail addresses: nkellard@essex.ac.uk (N. Kellard), c.dunis@ljmu.ac.uk (C. Dunis), n.sarantis@londonmet.ac.uk (N. Sarantis).

0378-4266/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2009.09.017

IV and RV are fractionally cointegrated series.[1] Interestingly, Bandi and Perron (2006) stress the fractional order of volatility is found in the non-stationary region whereas Christensen and Nielson (2006) and Nielsen (2007) indicate the stationary region. However, allowing for 95% confidence intervals, the estimates could plausibly lie in either region. In any case, Marinucci and Robinson (2001) stress that it is typically difficult to determine the integration order of fractional variables because a smooth transition exists between stationary and non-stationary regions. Christensen and Nielson (2006) and Nielsen (2007) note that when the fractional nature of the data is accounted for a slope parameter of unity in Eq. (1) cannot be rejected. Bandi and Perron (2006), noting the non-standard asymptotic distribution of conventional estimators in the non-stationary region, cannot formally test the relevant null hypothesis. However, subsampling shows their results also give support to the unbiasedness hypothesis.

This paper builds on the empirical work of Bandi and Perron (2006), Christensen and Nielson (2006), and Nielsen (2007) in five steps. Firstly, we employ data for several foreign exchange markets including the relatively new Euro market. Importantly, the IV data collected is traded on the market (and hence is directly observable). Since these data are directly quoted from brokers, they avoid the potential measurement errors associated with the more common approach (see, *inter alia,* Christensen and Prabhala, 1998) of backing out implied volatilities from a specific option-pricing model.

Secondly, it is important to note that in the recent literature, RV is constructed either from (i) high frequency intra-day return data (see, for example, Nielsen, 2007) or (ii) daily return data (see Bandi and Perron, 2006). Neely (2009) suggests that, at least in the context of least squares regression, the use of intra-day instead of daily data, does not resolve the biased slope coefficient. However, to our knowledge, this comparison has not been formally drawn in a fractionally cointegrated setting. Additionally, given that RV constructed from intra-day data is likely to be a less noisy proxy[2] for the unobserved but true volatility, the key to detecting (small) time-varying risk premia might be the use of such high frequency data. For example, consider augmenting regression (1) with a time-varying risk premium term rp_t

$$\sigma^{\mathrm{RV}}_{t+\tau} = \alpha + \beta\sigma^{\mathrm{IV}}_{t} + \delta rp_t + u_{t+\tau}. \tag{2}$$

Bivariate fractional cointegration between RV and IV implies any risk premium will be of a lower order of (fractional) integration than the original regressors. As a result, and as noted by Bandi and Perron (2006), the use of spectral methods like narrow band least squares will estimate regression (1) consistently, even in the presence of the risk premium. Re-arranging (2) leads to

$$\sigma^{\mathrm{RV}}_{t+\tau} - \alpha - \beta\sigma^{\mathrm{IV}}_{t} = \delta rp_t + u_{t+\tau}. \tag{3}$$

Given that daily data is relatively noisy, it might be that any long memory behaviour of the risk premium[3] is swamped[4] by $u_{t+\tau}$ in finite samples. In other words, a potential pitfall of employing daily data to construct RV is that it might render the risk premium undetectable. On the other hand, the use of a less noisy intra-day derived RV may lead to a smaller $u_{t+\tau}$ and therefore the revealing of a time-varying risk premium. Following Bandi and Perron (2006), we deliberately eschew modelling a specific functional form for a risk premium, simply suggesting that fractionally integrated behaviour in the residual of (1) provides prima facie evidence for latent risk premia. To examine these issues, we construct two RV series from intra-day[5] and daily data.

Thirdly, the possibility of fractional cointegration is examined formally using a new adaptation of the recently developed semi-parametric technique of Hassler et al. (2006) [hereafter HMV]. Under certain assumptions HMV prove that a residual-based log periodogram estimator, where the first few harmonic frequencies have been trimmed, has a limiting normality property. In particular, this methodology provides an asymptotically reliable testing procedure for fractional cointegration when the fractional order of regressors presents a particular type of non-stationarity. However, given the noted empirical uncertainty, (foreign exchange) volatility may present an integration order that violates the assumptions for the HMV test, as well as other fractional cointegration tests. To circumvent this uncertainty, we suggest, examine and apply an adapted fractional cointegration test robust to both stationary and non-stationary regions.

Fourthly, given the non-standard asymptotic distribution of conventional estimators when using fractionally integrated data, we employ a wild bootstrap procedure as suggested by Gerolimetto (2006) to compute appropriate confidence intervals in (1). Again, this specifically overcomes the difficulties encountered when estimators are applied in the non-stationary region.

Fifthly, we stress that the existence of fractional cointegration and that $\alpha = 0$ and $\beta = 1$ in (1) are only necessary conditions for unbiasedness. The important condition, that $u_{t+\tau}$ in (1) is serially uncorrelated, is required but such tests have been neglected by the recent extant literature. For completeness therefore, we employ an appropriate portmanteau test to the fractionally cointegrating residual.

The paper is divided into five sections: Section 2 presents the empirical methodology; Section 3 describes the data; Section 4 analyses the empirical results and, finally, Section 5 concludes.

2. Empirical methodology

2.1. Fractional integration

Many in the literature (see, *inter alia*, Bandi and Perron, 2006; Vilasuso, 2002; Andersen et al., 2001a; Baillie et al., 1996) have suggested that asset price volatility is neither an I(1) nor an I(0) process but rather a fractionally integrated or I(d) process. The introduction of the autoregressive fractionally integrated moving average (ARFIMA) model by Granger and Joyeux (1980) and Hosking (1981) allows the modelling of persistence or long memory where $0 < d < 1$. A time series y_t follows an ARFIMA[6] (p, d, q) process if

$$\Phi(L)(1-L)^d y_t = \mu + \Theta(L)\varepsilon_t, \quad \varepsilon_t \sim iid(0,\sigma^2), \tag{4}$$

where $\Phi(L) = 1 - \varphi_1 L - \cdots - \varphi_p L^p$ and $\Theta(L) = 1 - \theta_1 L - \cdots - \theta_q L^q$. Such models may be better able to describe the long-run behaviour of certain variables. For example, when $0 < d < 1/2$, y_t is stationary but contains long memory, possessing shocks that disappear hyperbolically not geometrically. Contrastingly, for $1/2 < d < 1$, the

[1] Although this recent work predominantly investigates stock markets, Bandi and Perron (2006) also analyse options on Deutsche Mark/US Dollar futures. Finding similar results to those for stock markets they suggest that fractional cointegration in the implied–realized relation is a stylised fact.

[2] Andersen and Bollerslev (1998) suggest that daily squared returns are noisy estimators for daily volatility and show that the sum of squared intra-day returns is a less noisy proxy. Employing the theory of quadratic variation, Andersen et al. (2001a) provide theoretical rationale for the intra-day approach.

[3] Kellard and Sarantis (2008) provide evidence for a fractionally integrated risk premium in forward foreign exchange markets. For discussion of volatility risk premia in other markets (see Almeida and Vicente, 2009; Doran and Ronn, 2008).

[4] For further discussions of swamping in time series see Maynard and Phillips (2001), Kellard (2006), and Kellard and Sarantis (2008).

[5] We thank an anonymous referee for enquiring about the use of intra-day data for constructing RV.

[6] ARFIMA models have been often used to model and forecast volatility (see, *inter alios*, Konstantinidi et al., 2008; Becker et al., 2007).