Journal of Banking & Finance 35 (2011) 3055–3064

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

A cyclical model of exchange rate volatility

Richard D.F. Harris [a,*], Evarist Stoja [b], Fatih Yilmaz [c]

[a] Xfi Centre for Finance and Investment, University of Exeter, The Queen's Drive, Exeter EX4 4QJ, UK
[b] School of Economics, Finance and Management, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK
[c] SLJ Macro Partners, 52 Brook Street, 3F, London W1K 5DS, UK

ARTICLE INFO

Article history:
Received 26 April 2010
Accepted 19 April 2011
Available online 23 April 2011

JEL classification:
G11
G17

Keywords:
Conditional volatility
Intraday range
Non-parametric filter

ABSTRACT

In this paper, we investigate the long run dynamics of the intraday range of the GBP/USD, JPY/USD and CHF/USD exchange rates. We use a non-parametric filter to extract the low frequency component of the intraday range, and model the cyclical deviation of the range from the long run trend as a stationary autoregressive process. We use the cyclical volatility model to generate out-of-sample forecasts of exchange rate volatility for horizons of up to 1 year under the assumption that the long run trend is fully persistent. As a benchmark, we compare the forecasts of the cyclical volatility model with those of the range-based EGARCH and FIEGARCH models of Brandt and Jones (2006). Not only does the cyclical volatility model provide a very substantial computational advantage over the EGARCH and FIEGARCH models, but it also offers an improvement in out-of-sample forecast performance.

1. Introduction

There is now substantial evidence that financial market volatility is both time-varying and highly predictable (see, for example, Andersen et al., 2004). This has important implications for many applications in finance, including portfolio optimisation, risk measurement and option pricing, and has given rise to a large literature on volatility measurement and forecasting. As noted by Brandt and Jones (2006), the efficacy of a volatility model depends on a number of factors. The first is the adequacy of the proxy for unobserved volatility that is employed in the model. Traditional volatility proxies based on the squared demeaned return are unbiased estimators of the latent integrated variance because the integrated volatility is, by construction, the expectation of the squared demeaned return. However, measures based on the squared return are inefficient owing to the fact that they employ only a single measurement of the price each period and hence contain no information about the intra-period trajectory of the price. An improvement in efficiency can be obtained by using intraday data. Indeed, Andersen et al. (2004) show that under very general assumptions, the sum of squared intraday returns converges to the unobserved integrated volatility as the intraday interval goes to zero. However, the construction of realized volatility relies on high frequency data, which is often not readily available over extended periods. Moreover, as the intraday frequency increases, market microstructure effects distort the measurement of returns, leading to an upward bias in realized volatility.

An alternative volatility proxy, and one that has recently experienced renewed attention, is the intraday range, which is defined as the scaled difference between the intraday high and low prices. Building on the earlier results of Parkinson (1980), Garman and Klass (1980) and others, Alizadeh et al. (2002) show that in addition to being significantly more efficient than the squared return, the intraday range is more robust than realized volatility to market microstructure noise. Moreover, an important practical advantage of the intraday range is that in contrast with the high frequency data that are required for the construction of realized volatility, intraday high and low prices are readily available for almost all financial assets over extended periods of time.[1] The intraday range has been employed in a number of conditional volatility models, including Chou (2005), who develops a conditional autoregressive range (CARR) estimator that is based on the conditional duration model of Engle and Russell (1998) and Brandt and Jones (2006), who extend the EGARCH model of Nelson (1991) using the intraday range in place of the absolute return. In both cases, the range-based GARCH estimators generate more accurate volatility forecasts than equivalent models based on squared returns.

The second factor that determines the efficacy of a volatility model is the specification of the process that governs volatility

* Corresponding author. Tel.: +44 1392 263215; fax: +44 1392 263242.
E-mail addresses: R.D.F.Harris@exeter.ac.uk (R.D.F. Harris), e.stoja@bristol.ac.uk (E. Stoja), faylmz@yahoo.co.uk (F. Yilmaz).

[1] For example, Datastream records the intraday range for most securities, including equities, currencies and commodities, going back to about 1985.

0378-4266/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2011.04.007

dynamics. Increasingly, evidence suggests that volatility is characterised by a multi-factor structure, with different dynamic processes governing the long run and short run dynamics of volatility. Engle and Lee (1999) introduce a component GARCH model, which decomposes volatility into a permanent long run trend component and a transitory short-run component that is mean-reverting towards the long run trend.[2] Empirical evidence suggests that the two-factor GARCH model provides a better fit to the data than an equivalent one-factor model. Alizadeh et al. (2002) estimate both one-factor and two-factor range-based stochastic volatility models for the daily returns of a number of exchange rates and find that the evidence strongly supports a two-factor model with one highly persistent factor and one rapidly mean-reverting factor. Similarly, Brandt and Jones (2006) estimate one-factor and two-factor range-based EGARCH models for daily returns on the S&P 500 index. They too show that volatility is well characterised by a two-factor model with one highly persistent factor and one strongly stationary factor.[3]

The finding that volatility has both a highly persistent factor and a strongly stationary factor has important implications for modelling and forecasting volatility over both short and long horizons. In particular, using the range-based two-factor EGARCH model, Brandt and Jones (2006) show that there is substantial predictability in volatility at horizons of up to 1 year. This is in contrast with earlier studies, such as West and Cho (1995) and Christoffersen and Diebold (2000), both of which conclude that volatility predictability is essentially a short horizon phenomenon. It is clear that the success of two-factor models in forecasting volatility rests on their ability to correctly identify the current long run level of volatility, and to exploit the dynamics of the short-run component to forecast reversion of volatility to the current trend. To the extent that the long-run component is close to being non-stationary, its dynamics are only relevant, if at all, over much longer forecasting horizons.

Motivated by the above interpretation of two-factor volatility models, we explore an alternative, very simple approach to modelling and forecasting volatility over both short and long horizons. In particular, we estimate the long run trend in measured volatility using a non-parametric filter, and then model the dynamics of the short-run component as a stationary autoregressive process around this long run trend. Our measure of volatility is based on the intraday range in order to exploit the improvement in efficiency that it offers over the squared return. Rather than apply the non-parametric filter directly to the intraday range, we separately extract the long-run components of intraday high and low prices, and then use these to construct a range-based estimate of the cyclical component of volatility. This is motivated by the fact that intraday high and low prices are more likely to satisfy the assumptions of the non-parametric filter and hence give reliable estimates of the underlying long run trend in volatility. To extract the long run trends in intraday high and low prices we use the low-pass filter of Hodrick and Prescott (1997) and the band-pass filter of Christiano and Fitzgerald (2003).

We implement the cyclical volatility model using GBP/USD, JPY/USD and CHF/USD daily exchange rates over the period 1 January, 1987 to 28 April, 2008. Consistent with the findings of Engle and Lee (1999), Alizadeh et al. (2002) and Brandt and Jones (2006), we show that the long-run component of volatility is characterised by a time-varying but highly persistent trend, while the short-run component is strongly mean-reverting to this trend. We use the model to generate out-of-sample forecasts of exchange rate volatility. We assume that over the forecast horizon, the long-run component of volatility follows a random walk and use the estimated parameters of the autoregressive model to forecast the deviation of volatility from the long-run component. As a benchmark, we compare the forecasts of our model with those of the one-factor and two-factor range-based EGARCH models and the range-based FIEGARCH model of Brandt and Jones (2006). Following the approach of Brandt and Jones (2006), we forecast volatility up to 1 year ahead, and use a range-based volatility proxy to evaluate the forecasts from each model. In almost all cases, the cyclical volatility model using the Hodrick–Prescott filter provides an improvement in forecast performance over the EGARCH and FIEGARCH models, in terms of both accuracy and informational content. The improvement in performance is particularly evident over shorter horizons where the random walk assumption for the long trend is most likely to be a good approximation. A significant practical advantage of the cyclical volatility model is the ease with which it can be estimated relative to more sophisticated models. In particular, while the recursive estimation of the two-factor EGARCH model and the FIEGARCH model over our sample is measured in tens of hours, the cyclical volatility model using the Hodrick–Prescott filter can be estimated in a matter of seconds. In practical situations that require either timely forecasts (such as for intraday options trading) or recursive estimation (such as for back testing risk management systems), the cyclical volatility model provides a feasible way of incorporating long memory in volatility forecasting in way that the EGARCH and FIEGARCH models do not.

The outline of the remainder of this paper is as follows. In Section 2 we present the theoretical framework of the new model. Section 3 describes the data used in the empirical analysis and the forecast evaluation criteria. Section 4 presents the empirical results. Section 5 provides a summary and offers some concluding remarks.

2. Theoretical background

Suppose that the log price of an asset, $p(t)$, follows a continuous-time diffusion given by

$$dp(t) = \sigma^2(t)dW(t) \tag{1}$$

where $dW(t)$ is the increment of a Wiener process and $\sigma^2(t)$ is the instantaneous variance, which is strictly stationary and independent of $dW(t)$.[4] Suppose that the price is recorded at daily intervals $t = 1, \ldots, T$. Then conditional on the sample path of $\sigma^2(t)$, the daily logarithmic return, $r_t = p_t - p_{t-1}$, is normally distributed with integrated variance $\sigma_t^2 = \int_{t-1}^{t} \sigma^2(s)ds$. The integrated variance, σ_t^2, is unobserved, but in principle can be estimated arbitrarily accurately using a measure of realized volatility based on intraday returns. In particular, Andersen et al. (2004) show that under very general conditions, the sum of squared intraday returns converges to the unobserved integrated volatility as the intraday interval goes to zero. However, the construction of realized volatility relies on high frequency intraday data, which are often not readily available over extended periods. Moreover, the accuracy of realized volatility as a proxy for integrated volatility is limited by the fact that as the intraday measurement frequency increases, market microstructure effects distort the measurement of returns, leading to an upward bias in estimated volatility (see, for example, Aït-Sahalia et al., 2005). An alternative approach is to construct an estimate of volatility based on the intraday range, given by

$$\sigma_{R,t}^2 = \frac{1}{4\ln 2}(p_t^H - p_t^L)^2 \tag{2}$$

[2] See also Christoffersen et al. (2008), who derive a two-component version of the GARCH model developed by Heston and Nandi (2000), which permits a closed-form solution for option valuation.

[3] See also Gallant et al. (1999), Chernov et al. (2003), Barndorff-Nielsen and Shephard (2001), Bollerslev and Zhou (2002), Maheu (2005) and Li (2011).

[4] For convenience, we assume that the drift of the log price process is zero, which is a common assumption when dealing with short horizon returns. However, it is straightforward to relax this assumption.