ELSEVIER

Journal of International Money and Finance
20 (2001) 1–23

Journal of International Money and Finance

www.elsevier.nl/locate/econbase

Forecasting daily exchange rate volatility using intraday returns

Martin Martens *

School of Banking and Finance, University of New South Wales, Sydney, NSW 2052, Australia

Abstract

This study investigates whether *intraday* returns contain important information for forecasting *daily* volatility. Whereas in the existing literature volatility models for daily returns are improved by including intraday information such as the daily high and low, volume, the number of trades, and intraday returns, here the volatility of intraday returns is explicitly modelled. Daily volatility forecasts are constructed from multiple volatility forecasts for intraday intervals. It is shown for the DEM/USD and the YEN/USD exchange rates that this results in superior forecasts for daily volatility. © 2001 Elsevier Science Ltd. All rights reserved.

JEL classification: G15

Keywords: Foreign exchange rates; Forecasting volatility; GARCH; Value-at-risk

1. Introduction

Nelson (1992) and Foster and Nelson (1995) show that in theory volatility estimates and forecasts can be made as accurate as required for many diffusion models by using ARCH estimates and sufficiently frequent price measurements. As trading is not continuous, bid/ask spreads exist and intraday volatility has strong seasonal characteristics, there are, of course, practical limits to the benefits obtainable from intraday data. However, it remains an interesting empirical question how one can improve, for example, daily volatility *forecasts* using intraday information.

Andersen and Bollerslev (1998b) show that although the daily squared return is an unbiased estimator of the *realized* daily volatility, it is also extremely noisy. The

* Tel.: +61-2-9385-4291; fax: +61-2-9385-6347.
E-mail address: m.martens@unsw.edu.au (M. Martens).

PII: S0261-5606(00)00047-4

sum of squared intraday returns substantially reduces the noise and should therefore be used for realized volatility. This argument should equally apply to forecasting. For example in the popular GARCH(1,1) model applied to daily returns, the conditional variance is a function of both the conditional variance and the squared (mean-adjusted) return from the previous trading day. If on the previous trading day the return was zero, but within the day prices fluctuated heavily, the lagged squared return (equal to zero) is misleading information. Other measures are then needed to capture the information that the previous trading day was actually quite volatile. This provides a partial explanation for the success of applying the daily high and low (Parkinson, 1980; Garman and Klass, 1980; Beckers, 1983; Taylor, 1987, among others), daily volume (Bessembinder and Seguin, 1993), number of price changes (Laux and Ng, 1993) and the standard deviation of intraday returns (Taylor and Xu, 1997). Including the standard deviation of intraday returns is very similar to replacing the daily squared return with the sum of squared intraday returns for the realized daily volatility as in Andersen and Bollerslev.

Whereas in the above studies the daily model is extended by including the additional intraday information, here it is investigated whether modelling the intraday returns directly provides better out-of-sample forecasts of daily volatility. Multiple period forecasts from intraday volatility models are compared with forecasts from daily volatility models, with or without additional intraday information. These forecasts are evaluated using the new realized volatility measure of Andersen and Bollerslev.

First a simulation experiment is carried out for a continuous time GARCH diffusion process including an intraday seasonal volatility pattern. The results from this experiment show that within this framework the use of higher return frequencies results in superior out-of-sample daily volatility forecasts. This is to be expected since the Data Generating Process (DGP) will be more closely approximated for higher return frequencies. It is also shown that ignoring the seasonal volatility pattern only leads to a marginal loss in efficiency. Extending the daily model with the difference between the daily high and low also improves out-of-sample forecasting, but not as much as the improvement obtained by modelling the highest return frequencies directly. Extending the daily model with the sum of intraday squared returns leads to a similar improvement as modelling the intraday returns directly.

Second, empirical results are obtained from intraday GARCH models for the DEM/USD and YEN/USD exchange rates. The daily GARCH(1,1) model is estimated and subsequently the parameters for the intraday GARCH models are derived using the aggregation results of Drost and Nijman (1993). The results indicate that modelling[1] intraday returns and volatility improves the out-of-sample forecasting of

---

[1] As mentioned above nowhere in the paper are GARCH models for intraday returns actually estimated. Andersen and Bollerslev (1997) show that the aggregation results of Drost and Nijman (1993) break down at the highest frequencies due to market microstructure effects. These inconsistencies with high-frequency GARCH estimates are circumvented by using implied intraday parameters from the estimated daily GARCH(1,1) model. Hence, here and in the following 'modelling intraday returns' refers to this procedure.