Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://www.econ.upenn.edu/pier

PIER Working Paper 02-017

"Time series modelling of daily log-price ranges for SF/USD and USD/GBP"

by

Celso Brunetti and Peter M. Lildholdt

Time series modelling of daily log-price ranges for SF/USD and USD/GBP

Celso Brunetti*
Department of Economics, University of Pennsylvania

Peter M. Lildholdt†
Department of Economics, University of Aarhus,
and Centre for Analytical Finance

First draft: June 1, 2001
This draft: June 29, 2002

Abstract

The aim of this paper is to model the dynamic evolution of daily log-price ranges for two foreign exchange rates, SF/USD and USD/GBP. Following Chou (2001), we adopt the CARR model, which is identical to the ACD model of Engle & Russell (1998). Log-price ranges are highly efficient measures of daily volatilities and hence our empirical results provide insights into the volatility dynamics for SF/USD and USD/GBP. We find that both series are highly persistent, and in particular, USD/GBP calls for a long memory specification in the form of a fractionally integrated CARR model. Semi-parametric and parametric models are estimated, and the parametric (fractionally integrated) CARR with a Gamma distribution is the preferred model. However, the estimation results of the simple semi-parametric procedure (QMLE) are virtually identical to the results of the preferred parametric models.

Keywords: Daily price range, foreign exchange rate volatility, persistency, CARR, FICARR

JEL Classification: C22, G15

Acknowledgments: We thank Niels Haldrup for useful comments. The usual disclaimer applies.

*E-mail: celsob@econ.upenn.edu

†E-mail: plildholdt@econ.au.dk

1 Introduction

Time series modelling of financial market volatility has been a very productive area of research in empirical finance. Accurate modeling and forecasting of financial asset volatility is of paramount importance in e.g. risk management and option pricing. A complication in volatility-modelling is that volatility is not observable. However, several proxies for volatility exist: Squared returns, absolute returns, realized volatility measures, etc. In the framework of Stochastic Volatility/GARCH, the underlying volatility is extracted from the volatility proxy, thereby explicitly recognizing that volatility is non-observable. Another approach is to treat volatility as an observable variable, modelling the volatility proxy directly, e.g. Andersen, Bollerslev, Diebold & Labys (2001*b*) model realized volatility measures, and Beran & Ocker (2001) model power transformations of absolute returns. We follow the latter approach in this paper. In particular, we employ a nonstandard volatility proxy: The daily price range,[1] which is a very efficient estimate of the daily volatility, see Parkinson (1980) and Brunetti & Lildholdt (2002). Treating the daily price range as an observable, highly efficient estimator of daily volatility, we model two time series of daily price ranges for foreign exchange rates: USD/GBP and SF/USD. Following Chou (2001), we adopt the Conditional AutoRegressive Range model (CARR), which is identical to the Autoregressive Conditional Duration (ACD) model of Engle & Russell (1998). The difference is that Engle & Russell (1998) applies the model to inter-trade durations and Chou (2001) applies the model to price ranges. This family of models is a natural choice, because durations and price ranges are both defined on the real positive line.

The main goal of the paper is to study the dynamic properties of the log-price ranges. Our findings confirm that volatility is a highly persistent process and, paralleling the results from the empirical GARCH literature, we find evidence of autoregressive roots close to unity. Further, extending the CARR model to allow for fractional integration, see Jasiak (1999) and Brunetti & Gilbert (2001), the dynamic properties of daily price ranges for USD/GBP seem to be well described by the FICARR(1,d,1) model. On the other hand, the short memory CARR(1,2) model appears to be the appropriate specification for the SF/USD price range. The distributional assumption for the innovation term in the CARR model plays an important role in the estimation procedure. However, comparing semi-parametric estimates and parametric estimates, we find that the former procedure is able to produce accurate parameter estimates and robust inference.

The structure of the paper is as follows: Section 2 motivates the approach of this paper and section 3 describes the data. In section 4, we present the CARR model of Chou (2001) and we introduce the FICARR model. In section 5.1 - 5.2, we present semi-parametric estimates of the (FI)CARR models and section 5.4 proceeds with parametric estimation results. Section 6 concludes.

[1] The daily price range is equal to the difference between the intradaily high log-price and the intradaily low log-price.

2 Motivation

In the following, we will motivate our interest in modelling the daily price range by illustrating its statistical properties as a volatility estimator. Assume the log-price (measured in log-US Dollars) on day s, at time t, of one Swiss Franc is denoted by $P_{t,s}^{SF/USD}$. Further, assume $P_{t,s}^{SF/USD}$ evolves according to a driftless Brownian Motion with diffusion coefficient σ_s

$$P_{t,s}^{SF/USD} = \sigma_s^{SF/USD} W_{t,s} \tag{1}$$

The daily price range on day s is defined by

$$l_s^{SF/USD} = \sup_{0 \leq t \leq T} P_{t,s}^{SF/USD} - \inf_{0 \leq t \leq T} P_{t,s}^{SF/USD}$$

where T denotes the length of one trading day. As mentioned in the introduction, conventional volatility proxies are based on returns, e.g. daily squared returns,

$$\left(P_{T,s}^{SF/USD} - P_{0,s}^{SF/USD} \right)^2$$

absolute returns,

$$\left| P_{T,s}^{SF/USD} - P_{0,s}^{SF/USD} \right|$$

or variations on the estimators above based on intradaily returns, see e.g. Andersen, Bollerslev, Diebold & Labys (2001*a*) and Barndorff-Nielsen & Shephard (2001).

Parkinson (1980) showed that the unbiased range-based estimator of the squared diffusion coefficient,

$$\left(\widehat{\sigma_s^{SF/USD}} \right)^2 = \frac{\left(l_s^{SF/USD} \right)^2}{4 \ln(2)},$$

is approximately 5 times more efficient than the unbiased estimator based on the daily squared returns. Brunetti & Lildholdt (2002) showed that the unbiased estimator of $\sigma_s^{SF/USD}$ based on the range

$$\widehat{\sigma_s^{SF/USD}} = \sqrt{\frac{\pi}{8}} l_s^{SF/USD} \tag{2}$$

is approximately 6.5 times more efficient than the unbiased estimator based on absolute returns,

$$\widetilde{\sigma_s^{SF/USD}} = \sqrt{\frac{\pi}{2}} \left| P_{T,s}^{SF/USD} - P_{0,s}^{SF/USD} \right|$$

In (1), the diffusion coefficient stays constant during the day. Andersen & Bollerslev (1998, footnote 20) adopt a stochastic volatility model where the diffusion coefficient evolves continuously, and they report that the MSE of the range-based estimator of integrated volatility is approximately equal to the MSE of a realized volatility estimator based on two- or three-hour returns. Hence, in this setting too, the daily price range provides a highly efficient estimator of daily volatility (integrated volatility).

In an empirical study, Li & Weinbaum (2000) compare the empirical performance of various return- and range-based volatility estimators against the benchmark of a realized volatility measure. The daily price range provides an estimator which is clearly superior to the return-based estimator, measured by various criteria such as MSE, variance, MAD etc.

The intuition underlying the superior performance of the range may be explained by the fact that computation of the range requires the full sample path of the price process, even though only two prices are used. On the other hand, return-based estimates of the volatility process are based on the opening and closing prices only - no intradaily information is used.

Motivated by these properties, we would like to scrutinize the dynamic properties of scaled log-price ranges. In particular, we will refer to the daily log-price range as a scaled version of the true price range from equation (2):

$$R_s^{SF/USD} = \sqrt{\frac{\pi}{8}} l_s^{SF/USD} = \sqrt{\frac{\pi}{8}} \left[\max_{0 \le t \le T} \ln \left(p_{t,s}^{SF/USD} \right) - \min_{0 \le t \le T} \ln \left(p_{t,s}^{SF/USD} \right) \right] \tag{3}$$

and likewise

$$R_s^{USD/GBP} = \sqrt{\frac{\pi}{8}} l_s^{USD/GBP} = \sqrt{\frac{\pi}{8}} \left[\max_{0 \le t \le T} \ln \left(p_{t,s}^{USD/GBP} \right) - \min_{0 \le t \le T} \ln \left(p_{t,s}^{USD/GBP} \right) \right] \tag{4}$$

3 Data

We analyze data on foreign exchange rates of SF/USD and USD/GBP over the period 3 January 1991 - 20 September 2001. Data are from Bloomberg and correspond to intradaily high and low prices. The foreign exchange market is open on a 24-hour basis and in the data set a new day starts at 06:00 PM GMT. The data set contains 2734 observations for each rate[2]. Prices refer to quotes recorded over the 24 hour period. Quotes might not be truly representative of market conditions, but we have not been able to get high/low transaction prices spanning a ten year period from the foreign exchange market. A closely related issue is that high and low prices might be recorded at times when markets are not very active, and the price range might therefore not reflect the true price range. However, in the absence of transactions data, it is almost impossible to correct for these potential flaws.

It is important to note that the sample period includes the September 1992 collapse of the European Monetary System (EMS). The UK was part of the EMS and experienced high volatility in the exchange rates as a consequence of severe speculative attacks.

Figure 1 shows time series plots of daily price ranges, see equations (3) and (4), for the two rates. Both series are stationary, and the September 1992 speculative attack

[2]For further information about the data set, we refer to Brunetti & Lildholdt (2002) where the same data set is used.

against the GBP is evident in both series. This is to be expected, given that, starting from September 1992, many European currencies experienced a very difficult time due to the uncertainty surrounding the future of the EMS. The effect of the speculative attack is stronger in the USD/GBP range. In both series, volatility seems itself more volatile at the beginning of the period analyzed. For the USD/GBP range, the volatility process stabilized at the end of 1993. It is quite surprising that the USD/GBP rate is less volatile after the UK left the EMS.

In Table 1, we report summary statistics for daily price ranges. The mean and the standard deviation of USD/GBP price ranges are lower than those of SF/USD ranges. Skewness and kurtosis of the two series are very similar and indicate that the ranges are not symmetrically distributed and exhibit excess kurtosis. These results are confirmed by looking at Figure 2, which shows the empirical distribution (histograms) for daily price ranges.

Table 2 reports Ljung-Box (LB) statistics for serial correlation up to lags, 1, 15, 30, 50, 100 and 250. All statistics are significant[3] at 5% significance level (denoted by *), and it is evident that the series are highly correlated. The values of the LB-test are particularly high for the USD/GBP range. This finding is also supported by Figure 3, where autocorrelation functions (ACF) for the two series are shown. It is evident that the USD/GBP range is very persistent. The ACF is always positive and remains significantly different from zero for high lags. On the other hand, the SF/USD range is less persistent and the ACF is statistically insignificant at lag 70.

4 CARR models

Following Chou (2001), the CARR model is defined by

$$R_s = \lambda_s \varepsilon_s, \qquad \varepsilon_s \sim iid, \qquad \varepsilon_s \geq 0, \qquad E(\varepsilon_s) = 1 \tag{5}$$

$$\lambda_s = \omega + a(L) R_s + b(L) \lambda_s \tag{6}$$

where R_s is the daily price range, $a(L) = a_1 L + a_2 L^2 + ... + a_p L^p$ and $b(L) = b_1 L + b_2 L^2 + ... + b_q L^q$ are lag polynomials of order p and q, respectively. λ_s is the conditional mean of the range. The model is identical to the Autoregressive Conditional Duration (ACD)[4] model of Engle & Russell (1998). This is a very appealing model because, as shown by Engle & Russell (1998), the specification is similar to the well-known and widely used GARCH model. Sufficient conditions for λ_s to be positive are

$$\omega > 0, a_i \geq 0, b_i \geq 0 \tag{7}$$

By defining $\epsilon_s = R_s - \lambda_s$, the model may be expressed as an ARMA process in R_s

$$(1 - a(L) - b(L))\, R_s = \omega + (1 - b(L))\, \epsilon_s \tag{8}$$

[3] LB(n) statistics follow a chi-squared distribution with n degress of freedom under the null of no serial correlation from lag 1 to n. Critical values for the LB tests in Table 2 (at a significance level of 5%) are: $\{3.84, 25.00, 43.77, 67.50, 124.34, 287.88\}$.

[4] We do not adopt this name for the model because durations are not involved in our setting.

Using the results from the vast GARCH literature, it follows that the CARR model is stationary when the roots of $(1 - a(L) - b(L))$ lie outside the unit circle. A stylized fact in finance is that the volatility of returns seems to be highly persistent in the sense that autocorrelations of volatility proxies decay very slowly. This feature is called the Long Memory property, see e.g. Baillie, Bollerslev & Mikkelsen (1996), Baillie (1996) and Andersen et al. (2001*b*). For this reason, we introduce the long memory extension of the CARR model, which is identical to the FIACD models of Jasiak (1999) and Brunetti & Gilbert (2001). The motivation for the FIACD model in intertrade duration analysis is that in empirical applications, $(1 - a(L) - b(L))$ from the ACD model has roots very close to unity. However, the presence of a unit root in $(1 - a(L) - b(L))$ is associated with the inappropriate property that shocks to the model have ever-lasting effects, see Jasiak (1999). It is conjectured in Jasiak (1999) that the FIACD model overcomes this feature in the sense that shocks have persistent effects that eventually die out over time. Hence, to create more persistent impulse responses for the model, a fractional root in the autoregressive polynomial for R_s is introduced in (8):

$$\phi(L)(1-L)^d R_s = \omega + (1 - b(L))\, \epsilon_s \tag{9}$$

The fractional differencing operator $(1-L)^d$ in equation (9) is defined by the hypergeometric function

$$(1-L)^d = \sum_{j=0}^{\infty} \frac{\Gamma(j-d)}{\Gamma(j+1)\,\Gamma(-d)} L^j \tag{10}$$

Rewriting equation (9) in terms of the conditional expectation of the range yields

$$\lambda_s = \omega + \left[1 - b(L) - \phi(L)(1-L)^d\right] R_s + b(L)\lambda_s \tag{11}$$

or, in a more parsimonious form

$$\lambda_s = \frac{\omega}{[1 - b(1)]} + \gamma(L) R_s \tag{12}$$

where $\gamma(L) = \frac{\left[1-b(L)-\phi(L)(1-L)^d\right]}{1-b(L)}$, and therefore it is a lag polynomial of infinite order. Equation (12) may be interpreted as an infinite MA process. We refer to this model as the Fractionally Integrated CARR model (FICARR). Sufficient conditions for $\lambda_s > 0$ in equation (11) are difficult to establish for the general case, but it is possible to find conditions for specific parametrizations, see Baillie et al. (1996). The FICARR is similar to the FIGARCH specification of Baillie et al. (1996). Paralleling the FIGARCH model, the unconditional mean of the FICARR model is infinite. The FICARR process is, therefore, not weakly stationary. Despite this, the FICARR process is strictly stationary and ergodic for $0 \leq d \leq 1$, see Baillie et al. (1996).[5]

We have not specified the distribution of the innovations, ε_s, in (5) yet. Empirical work on CARR models is limited to Chou (2001), and hence, our prior on the distribution of the

[5] A straightforward extension of the ACD specification is to model the natural logarithm of the expected range, see Engle & Lunde (1999). We abstain from that specification in the current paper.

innovation term is very diffuse. We adopt a semi-parametric method for estimating the models, and the estimation proceeds by quasi maximum likelihood with an Exponential density for the innovation term. The log-likelihood function is

$$\log L = -\sum_{j=1}^{S}\left[\ln(\lambda_s) + \frac{R_s}{\lambda_s}\right]$$

Engle & Russell (1998) (for the ACD model) and Jasiak (1999) (for the FIACD model) show that this Quasi Maximum Likelihood Method (QMLE) provides consistent and asymptotically normal estimates with a covariance matrix given by the familiar robust standard errors from Lee & Hansen (1994).

Estimation of the FICARR model involves truncating the infinite lag-polynomial in (12) at lag 1000 and setting pre-sample values of λ_s equal to the unconditional mean of the daily price ranges R_s.

In order to measure the goodness of fit of the models estimated by QMLE, we compute standardized innovations:

$$\widehat{u}_s = \frac{R_s}{\widehat{\lambda}_s}$$

$\widehat{\lambda}_s$ denotes the conditional expectation of the range computed via (6) or (12) by use of estimated coefficients. For a well-specified model, $\widehat{u_s}$'s should be independent and we compute Ljung-Box statistics to test this hypothesis.

5 Empirical results

As a first step in the empirical analysis, we estimated AR(FI)MA models for $\ln(R_s^{SF/USD})$ and $\ln(R_s^{USD/GBP})$.[6] However, it was virtually impossible to obtain well-specified models due to indications of strong ARCH effects and departures from normality for the innovation term. Hence, we adopted the CARR model.

5.1 Semi-parametric CARR models

In the vast GARCH literature, a common finding is that a persistent GARCH(1,1) specification is able to describe the volatility of financial returns. However, little is known about the dynamic properties of the range. For this reason, we estimate four CARR specifications: CARR(2,2), CARR(2,1), CARR(1,2) and CARR(1,1)

$$R_s = \lambda_s \varepsilon_s, \qquad \varepsilon_s \sim Exp(1) \tag{13}$$
$$\lambda_s = \omega + a_1 R_{s-1} + a_2 R_{s-2} + b_1 \lambda_{s-1} + b_2 \lambda_{s-2} + SEP92 \cdot DUM_s \tag{14}$$

The choice of the preferred model is based on three approaches:

[6] Log-transformations were adopted because the support of an AR(FI)MA process is the entire real line.

- General-to-specific approach, see e.g. Hendry (2000). Testing down from a general model by t-statistics of individual parameters.
- Log-likelihood based approach. For nested models, LR-tests are adopted and for non-nested models, Akaike's (AIC) and Schwartz's (SIC) information criteria are used.
- Model selection based on the ability of the model to fit the assumption of *iid* innovations. Ljung-Box-statistics for serial correlation in the standardized innovations are used.

Table 3 shows the estimation results for the SF/USD range. The most general model, the CARR(2,2), tells us that the parameter b_2 is not statistically significant. It is important to note that some of the estimated parameters are negative. Nevertheless, the conditional ranges, λ_s, are strictly positive in the sample. The restrictions in (7) are sufficient but not necessary conditions for the conditional range to be positive, see Nelson & Cao (1992). Using the general-to-specific approach, the preferred model is the CARR(1,2). However, the LR-test, the AIC and the SIC select the simple CARR(1,1) specification. The dispute between the CARR(1,2) and the CARR(1,1) is resolved by looking at the Ljung-Box statistics. It is evident that the values of the Ljung-Box statistics are much lower for the CARR(1,2) than for the CARR(1,1). A fundamental assumption in the CARR model - see equation (5) - is that the innovation term is *i.i.d.* Table 3 shows that significant autocorrelation at lag 1 remains in the standardized innovations for the CARR(1,1) model, and hence the CARR(1,2) specification is preferred.

For each specification we estimated two models: with and without a dummy variable for the September 1992 speculative attack.[7] It is evident that the speculative attack of September 1992 against the GBP rates increased volatility dramatically even for the SF/USD rate. However, the introduction of the dummy[8] does not modify the parameter estimates. A common characteristic of all the estimated models is that volatility is very persistent: $\widehat{a}(1) + \widehat{b}(1)$ ranges from 0.98 to 0.96. This result is in line with the findings of the GARCH literature.

The last two rows in Table 3 refer to the difference between the mean and the standard deviation of the standardized innovations and their theoretical values of unity. While the mean of the standardized innovations is very close to one (by first-order conditions), the standard deviation is not. Nevertheless, the QMLE procedure allows us to make asymptotically valid inference.

Table 4 shows the CARR model estimation results for USD/GBP. There are remarkable similarities between the results in Tables 3 and 4. In the CARR(2,2) model, diagnostics support the restriction of $b_2 = 0$. The preferred model using the general-to-specific approach and the Ljung-Box statistic is the CARR(1,2). If we use the LR test,

[7] We introduced the dummy variable in equation (14) following the standard practice in the GARCH literature, e.g. Lamoureux & Lastrapes (1990).

We use four different dummy variables. The reported results refer to the dummy variable which covers the period 16/Sep/92 - 28/Sep/92. For the other dummies the results are similar to those reported in Table 3.

[8] To conserve space we do not report the results of the ACD estimates without the September 1992 dummy.

the AIC and the SIC, the selected model is the CARR(1,1). However, the CARR(1,2) is accounting for serial correlation in data much better than CARR(1,1). Therefore, the CARR(1,2) is the preferred model.

As expected, the impact of the September 1992 dummy is significant (at the 10% level) and is stronger for the volatility of the USD/GBP rate than for the volatility of the SF/USD rate. We also estimated the same models omitting the dummy variable (results not reported) and obtained very similar parameter estimates to those in Table 4.

The general impression from Tables 3 and 4 is that persistency, measured by $\widehat{a}(1) + \widehat{b}(1)$, is strong and ranges from 0.99 for the CARR(2,2) to 0.978 for the CARR(1,1).

In the next section, we turn our attention to the semi-parametric FICARR estimates.

5.2 Semi-parametric FICARR models

For the FICARR model we also estimated four different specifications: FICARR(2,d,2), FICARR(2,d,1), FICARR(1,d,2), and FICARR(1,d,1), nested within the general model

$$R_s = \lambda_s \varepsilon_s, \qquad \varepsilon_s \sim Exp(1) \tag{15}$$

$$\begin{aligned}\lambda_s = \omega &+ \left[1 - b_1 L - b_2 L^2 - (1 + \phi_1 L + \phi_2 L^2)(1-L)^d\right] R_s \\ &+ b_1 \lambda_{s-1} + b_2 \lambda_{s-2} + SEP92 \cdot DUM_s\end{aligned} \tag{16}$$

Table 5 shows the results of the FICARR models for the SF/USD range. The long memory parameter, d, is insignificant except for FICARR(1,d,1). Testing down from the general FICARR(2,d,2) model, the final model specification depends on how (sequentially) we discard parameters. All models are able to produce uncorrelated innovations and the information criteria are in favor of the FICARR(1,d,1) model. However, this model is not completely satisfactory due to the fact that d is insignificant for the other specifications. It might be the case that d is spuriously capturing some short memory dynamics, and it is questionable whether a FICARR specification provides a suitable model for SF/USD.

The results of the FICARR estimates for USD/GBP are shown in Table 6. The results are unambiguous: Regardless of the selection criteria used, the FICARR(1,d,1) is always the preferred model. In fact, the diagnostics favor the restrictions $\phi_2 = b_2 = 0$.[9] The long memory parameter d is always significantly different from zero, indicating that the USD/GBP range is a long memory process.

The September 1992 dummy variable, is statistically important (at 10% significance level). We also estimated the same FICARR specification without the September 1992 dummy variable (results not reported). In line with the findings obtained for the CARR models, the introduction of the dummy does not modify the parameter estimates.

In the following section, we compare the preferred CARR models to the preferred FICARR models.

[9] We also investigated the FICARR(1,d,0) specification and found strong evidence in favor of the FICARR(1,d,1) model.

5.3 Comparison of semi-parametric CARR and FICARR models

We first analyze the USD/GBP series. For the general-to-specific approach, it is important to note that the CARR(2,2) is nested in the FICARR(2,d,2) - see equations (8) and (9). Therefore, using the general-to-specific approach, the preferred model is the FICARR(1,d,1). Both models, CARR(1,2) and FICARR(1,d,1) have very similar log-likelihood values, and the LB-statistic may help us in selecting the right specification. The values of the LB-statistics are lower for the FICARR(1,d,1) model. This is particularly true for long lags. It seems that the FICARR(1,d,1) model is able to account both for the short and the long range dependence. It is quite interesting that the preferred CARR model is of order (1,2). Note that a CARR(1,2) model may be rewritten as an ARMA(2,1) model in daily price ranges, see equation (9). Gallant, Hsu & Tauchen (1999) showed that the sum of two independent AR(1)-processes may approximate a fractionally integrated process. The sum of two independent AR(1) processes produces an ARMA(2,1) process, and hence the ARMA(2,1) model is able to mimic the long memory property of the series. This feature generally points to the fact that, to some degree, higher order ARMA processes are capable of capturing slowly decaying autocorrelations.

The estimates of b_1 in the FICARR(1,d,1) model for the USD/GBP range reduce dramatically when compared to those of the CARR models. This is in line with the results of Baillie et al. (1996). They show that if the underlying process is long memory and a short memory GARCH model is used in the estimation procedure, the estimates of the parameter b_1 are upward biased. The simple GARCH model is spuriously capturing the long memory effect through the b_1 parameter.

Turning to SF/USD, the preferred model is the CARR(1,2) specification. The FICARR(1,d,1) specification for the SF/USD range (last column in Table 5), shows a statistically significant value of the fractional integration parameter. This might be due to mis-specification of the model. The long memory parameter is spuriously capturing the short run dynamic of the a_2 parameter in the CARR(1,2) specification (see Table 3). When $d = 0$, the FICARR model collapses into the CARR specification with $\phi(L) = 1 - a(L) - b(L)$. The parameters of the FICARR(1,d,2) model are not individually significantly different from the corresponding parameters from the CARR(1,2) model,[10] which provides further evidence in favor of the CARR(1,2) model.

The results above are consistent with the empirical autocorrelation functions. The ACF for SF/USD range is not highly persistent, which indicates a short memory specification. In contrast, the ACF for the USD/GBP range is highly persistent and hence a long memory specification is needed to accommodate this feature.[11]

[10] $\widehat{\phi}_1$ is not significantly different from $\widehat{a}_1 + \widehat{b}_1$ for the CARR(1,2), $\widehat{\phi}_2$ is not significantly different from $\widehat{a}_2 + \widehat{b}_2$, and $\widehat{b}_1$ (for FICARR(1,d,2)) is not significantly different from $\widehat{b}_1$ (for CARR(1,2)).

[11] Stricly speaking, the FIACD(1,d,1) model is not weakly stationary and hence autocorrelations from that model are not well-defined. However, the intuition from Hasza (1980) and Bierens (1993) suggests that the above reasoning may be appropriate.

5.4 Parametric (FI)CARR models

The models from the former section are semi-parametric in the sense that the density of the innovations is left unspecified. The focus of this section is to fit parametric (FI)CARR models for the two time series of daily price ranges. In other words, we specify a parametric distribution for the innovations, ε_s, in equation (5).

A natural starting point is to adopt the Exponential assumption for the innovation term. We concentrate our analysis on the preferred models: CARR(1,2) and FI-CARR(1,d,1) specifications. If the models are correctly specified, the standardized innovations are distributed according to an Exponential distribution. In Figure 4, two densities are depicted. The density of the standardized innovations from the CARR(1,2) model for $R_s^{SF/USD}$ (see Table 3, third column) is shown in addition to a unit Exponential density. We estimate density functions by the Gamma kernel proposed by Chen (2000). This procedure overcomes the traditional boundary bias of standard kernel density estimators associated with estimating probability density functions with bounded support.[12] It is quite clear from the figure that the model is mis-specified: The two density functions are far from being identical. The same conclusion is valid for the CARR(1,2) model for $R^{USD/GBP}$ and the FICARR(1,d,1) models for both series, see Figures 5, 6 and 7, respectively. The mis-specification is clear also from the last rows of Tables 3 - 6: In fact, the standard deviation of the standardized innovations are far in excess of the theoretical value of unity.

5.4.1 Weibull (FI)CARR

A good candidate for the density of the standardized innovations is the Weibull[13] distribution, because it allows for the hump-shape of the price ranges in Figure 2. Hence, we estimate Weibull (FI)CARR models, abbreviated by WE-(FI)CARR and defined by

$$R_s = \lambda_s \frac{\varepsilon_s}{\Gamma\left(1+\frac{1}{\delta}\right)}, \qquad \varepsilon_s \overset{iid}{\sim} Weibull(1,\delta) \tag{17}$$

The use of the Weibull distribution in the ACD framework was proposed by Engle & Russell (1998). Following e.g. Lunde (1999), we compute Cox-Snell residuals, see Cox & Snell (1968), for the parametric (FI)CARR models. These residuals have the property that they are exponentially distributed, which allows for a direct comparison of the fit of various parametric models. They are defined by

$$u_s = \int_0^{\varepsilon_s} h_\varepsilon(x)dx \tag{18}$$

where $h_\varepsilon(\cdot)$ denotes the hazard function for ε_s. The hazard function is defined by

$$h_\varepsilon(\varepsilon_s) = \frac{f_\varepsilon(\varepsilon_s)}{1-F_\varepsilon(\varepsilon_s)}$$

[12] The standardized innovations are bounded from below by zero.

[13] If X^δ is distributed according to a unit exponential distribution, then we say that X follows a Weibull distribution with shape parameter δ. The unit exponential is a special case ($\delta = 1$) of the Weibull distribution.

where $f_\varepsilon(\cdot)$ denotes the density function for ε_s, and $F_\varepsilon(\varepsilon_s)$ denotes the cumulative distribution. Rewriting (18),

$$u_s = \int_0^{\varepsilon_s} h_\varepsilon(x)dx = \int_0^{\varepsilon_s} \frac{f_\varepsilon(x)}{1-F_\varepsilon(x)}dx = \left[-\ln\left(1-F_\varepsilon(x)\right)\right]_0^{\varepsilon_s} = -\ln\left(1-F_\varepsilon(\varepsilon_s)\right)$$

it follows that the Cox-Snell residuals follow an Exponential distribution. In fact, $F_\varepsilon(\varepsilon_s)$ follows a uniform distribution, and minus the log of a uniform random variable generates an exponentially distributed random variable, see e.g. Evans, Hastings & Peacock (1993, p. 61). Cox-Snell residuals for WE-(FI)CARR models are computed according to

$$\widehat{u}_s = \left(\widehat{\varepsilon}_s\right)^{\widehat{\delta}} = \left(\frac{R_s}{\widehat{\lambda}_s}\Gamma\left(1+\frac{1}{\widehat{\delta}}\right)\right)^{\widehat{\delta}}$$

If the model is well specified, $\widehat{u}_s$ is distributed according to a unit Exponential distribution. In the following, we also refer to Cox-Snell residuals as standardized innovations.

Estimates[14] of WE-(FI)CARR are reported in Table 7. The estimated parameters and standard errors are in line with those obtained with the QMLE in Tables 3 - 6. It is interesting to note that there is less persistency in the range when estimating the CARR model with the Weibull distribution, $\widehat{a}(1)+\widehat{b}(1)$ is equal to 0.95 for the SF/USD and 0.97 for the USD/GBP when using the Weibull and 0.98 for the SF/USD and 0.99 for the USD/GBP when using QMLE. The Weibull assumption for the error term produces values for the LB statistics that are much lower than those of the QMLE.

Figures 8 - 11 show empirical densities of Cox-Snell residuals from the WE-(FI)CARR models of Table 7. The fit has improved relative to Figures 4 - 7, but as indicated by the last row of Table 7, the standard deviation of the Cox-Snell residuals is far in excess of the theoretical value of unity for all models.

5.4.2 Gamma (FI)CARR

The Weibull distribution of the former section seemed to improve the fit of the models, although they are still mis-specified. Following Lunde (1999) for the ACD model, we adopt the Generalized Gamma distribution for the innovations in equation (5). The Generalized Gamma distribution has two shape parameters (α, δ), and it nests the Weibull distribution as a special case when $\alpha = 1$. For all estimated models, the restriction $\delta = 1$ cannot be rejected and hence the distribution of the innovations reduces to the Ordinary Gamma distribution. The Gamma (FI)CARR model is defined by

$$R_s = \lambda_s \frac{\varepsilon_s}{\alpha}, \qquad \varepsilon_s \overset{iid}{\sim} Gamma(1,\alpha)$$

where $Gamma(1,\alpha)$ denotes the Gamma distribution with the scale parameter of unity and the shape parameter α. Computation of Cox-Snell residuals for this model is rather

[14] The log-likelihood is given by: $\ln L(\delta,\omega;R_s|R_0) = \sum_{s=1}^{S} l_s(\delta,\omega)$ where $l_s(\delta,\omega) = \ln(\delta) + (\delta-1)\ln(R_s) - \delta\ln\left(\frac{\lambda_s}{\Gamma(1+\frac{1}{\delta})}\right) - \left(\frac{\Gamma(1+\frac{1}{\delta})R_s}{\lambda_s}\right)^{\delta}$.

tedious. Note that the hazard function for ε_s is

$$h_\varepsilon(\varepsilon) = \frac{f_\varepsilon(\varepsilon)}{1 - F_\varepsilon(\varepsilon)} = \frac{\frac{(\varepsilon_s)^{\alpha-1}}{\Gamma(\alpha)} \exp(-\varepsilon_s)}{1 - \int_0^{\varepsilon_s} \frac{(\varepsilon_s)^{\alpha-1}}{\Gamma(\alpha)} \exp(-\varepsilon_s)}$$

and hence Cox-Snell residuals are defined by

$$\widehat{u}_s = \int_0^{\widehat{\varepsilon}_s} \widehat{h}_\varepsilon(\widehat{\varepsilon}_s) = \int_0^{\widehat{\varepsilon}_s} \frac{\frac{(\widehat{\varepsilon}_s)^{\widehat{\alpha}-1}}{\Gamma(\widehat{\alpha})} \exp(-\widehat{\varepsilon}_s)}{1 - \int_0^{\widehat{\varepsilon}_s} \frac{(\widehat{\varepsilon}_s)^{\widehat{\alpha}-1}}{\Gamma(\widehat{\alpha})} \exp(-\widehat{\varepsilon}_s)}, \qquad \widehat{\varepsilon}_s = \frac{\widehat{\alpha} R_s}{\widehat{\lambda}_s} \tag{19}$$

where the integral is computed numerically.[15]

Table 8 shows estimates[16] of the Gamma (FI)CARR models. The coefficient estimates and the associated robust standard errors are virtually the same as those obtained with the Exponential distribution in the QMLE procedure. Moreover, there is evidence that the assumption of a Gamma distributed innovation term is more reasonable: Both the mean and the standard deviation of the Cox-Snell residuals are close to their theoretical counterparts, see the last two rows of Table 8. Figures 12 - 15 show the densities of Cox-Snell residuals from the estimated Gamma (FI)CARR models. The densities seem to fit the unit Exponential rather well. In contrast, Table 9, upper section shows the mean, standard deviation, skewness and kurtosis of the Cox-Snell residuals from the models of Table 8. It is evident that they do not seem to follow a unit Exponential[17] distribution. Moreover, the excess dispersion test of Engle & Russell (1998, p. 1144) rejects the null of the standard deviation being equal to unity for all conventional significance levels.

In the lower section of Table 9, we compute the same sample moments, excluding the 20 (for $R^{SF/USD}$) and 25 (for $R^{USD/GBP}$) largest Cox-Snell residuals. The empirical means, standard deviations, skewness and kurtosis of the standardized innovations are very close to their theoretical counterparts. The Excess Dispersion test, see Engle & Russell (1998), fails to reject the null of the standard deviation being equal to unity. The general message from Table 9 and Figures 12 - 15 is that the source of the mis-specification is a relatively small number of extreme observations.[18]

We further analyzed the largest Cox-Snell residuals for the four models. For the SF/USD, the largest Cox-Snell residuals for the GA-CARR(1,2) coincide with the largest Cox-Snell residuals for the GA-FICARR(1,d,1). Moreover, the largest Cox-Snell residuals of the two models match the largest price ranges. Press reports on the days corresponding to the largest ranges (very high volatility) reveal that unusual events took place - e.g.

[15] The integral is computed by the numerical procedure INTQUAD1 (using Gauss-Legendre quadrature) in Gauss.

[16] The log-likelihood is given by: $\ln L(\alpha, \omega; R_s|R_0) = \sum_{s=1}^{S} l_s(\alpha, \omega)$ where $l_s(\delta, \omega) = (\alpha - 1)\ln(R_s) - \alpha \ln(\frac{\lambda_s}{\alpha}) - \ln(\Gamma(\alpha)) - \left(\frac{\alpha R_s}{\lambda_s}\right)$.

[17] The unit Exponential distribution is characterized by mean=1, st.dev=1, skewness=2, and kurtosis=9.

[18] The same analysis has been performed also for the Cox-Snell residuals for the WE-(FI)CARR models. The results clearly show that even when deleting the extreme residuals, the assumption that the innovation term follows a Weibull distribution is poor.

Russian coup, Asian crisis, Russian crisis, etc. The same analysis holds for the USD/GBP rate. It seems that the (FI)CARR model works well in standard market conditions but is not able to account for abnormal events.

5.5 Summary of empirical results

By a semi-parametric procedure (QMLE), we estimated several short memory CARR models and found evidence in favor of the CARR(1,2) specifications. The results show clear evidence of a strong degree of persistency. For this reason, we extend the model to allow for long memory. The FICARR(1,d,1) model is able to capture the dynamic properties of the USD/GBP range while the CARR(1,2) is the preferred model for the SF/USD range.

For the parametric estimates of the (FI)CARR models, we use two alternative distributions, the Weibull and the Generalized Gamma. We find evidence that the Ordinary Gamma distribution, which is a special case of the Generalized Gamma, provides a fairly good approximation to the distribution of the standardized innovations, except in the very far end of the tail. We also note that the semi-parametric estimates reported in Tables 3 - 6 produce approximately the same parameter estimates and standard errors as the parametric Gamma (FI)CARR estimates.

For S&P500 index data, Chou (2001) performed volatility-forecasting for CARR and GARCH models and concluded that the CARR was superior. Our results indicate that the semi-parametric procedure (QMLE) is suitable for performing point-forecasting of daily price ranges in the foreign exchange market because it produces parameter estimates (and thereby point-forecasts) almost identical to a parametric model. However, for the purpose of density forecasting of daily price ranges, a (FI)CARR model with a Gamma distribution for the innovation term would be suitable.

Our empirical results support the conclusion that the (FI)CARR models estimated with the simple QMLE method are able to capture the dynamics of exchange rate volatility measured by the daily price range.

6 Conclusion

The aim of the paper is to shed light on the dynamic properties of exchange rate volatility. We adopt a non-standard volatility proxy: The daily price range, which, under certain assumptions, is highly efficient compared to volatility measures based on daily returns. Building on Chou (2001), we model daily price ranges for two major exchange rates, SF/USD and USD/GBP, by the CARR model.

For the foreign exchange market, there is extensive empirical evidence of a very persistent component in volatility using return-based volatility proxies. We provide additional and complementary evidence of this feature using the daily price range. We find that the USD/GBP range is a long memory process, paralleling the results obtained in the GARCH/FIGARCH and realized volatility literature. On the other hand, we also find that the SF/USD price range is best described by a short memory CARR(1,2) model.

The (FI)CARR models provide a simple, yet effective framework for modeling the daily price range. It would be interesting to explore whether alternative volatility proxies, such as absolute returns and realized volatility measures, fit the class of (FI)CARR models.

The CARR specification is a model for daily price ranges and it does not include daily returns. In other words, the model is agnostic about the evolution of returns. A more structural model including the return process might be desirable. However, a full understanding of the empirical properties of daily price ranges is a necessary first step in that development, we believe. Ongoing work, see Lildholdt (2002), provides a model for asset returns utilizing daily high and low prices in addition to daily return data (open and close prices).

The forecasting performance of the CARR model compared to return-based models, like GARCH and Stochastic Volatility models, is an interesting area for future research. Chou (2001) provides evidence of superior forecasting performance of the CARR model for the S&P500 index, and it is highly relevant to look into the forecasting performance for the FICARR model. At short horizons, the CARR and the FICARR are very similar, but the properties of the models differ at long horizons. The empirical results of this paper indicate that the FICARR model is well suited for capturing the strong persistence of daily price ranges, and that may result in superior forecasting performance at longer horizons. Needless to say, superior forecasting of daily price ranges has important applications in fields, such as VaR, option pricing, risk management, etc.

References

Andersen, T. & Bollerslev, T. (1998), 'Answering the skeptics: Yes, standard volatility models do provide accurate forecasts', *International Economic Review* **39**(4), 885–905.

Andersen, T. G., Bollerslev, T., Diebold, F. & Labys, P. (2001*a*), 'The distribution of exchange rate volatility', *Journal of the American Statistical Association* **96**, 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2001*b*), Modelling and forecasting realized volatility. Manuscript.

Baillie, R., Bollerslev, T. & Mikkelsen, H. (1996), 'Fractionally integrated generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **74**, 3–30.

Baillie, R. T. (1996), 'Long memory processes and fractional integration in econometrics', *Journal of Econometrics* **73**, 5–59.

Barndorff-Nielsen, O. & Shephard, N. (2001), 'Realized power variation and stochastic volatility models', *Forthcoming in Bernoulli* .

Beran, J. & Ocker, D. (2001), 'Volatility of stock-market indexes - an analysis based on SEMIFAR models', *Journal of Business and Economic Statistics* **19**(1), 103–116.

Bierens, H. (1993), 'Higher-order sample autocorrelations and the unit root hypothesis', *Journal of Econometrics* **57**, 137–160.

Brunetti, C. & Gilbert, C. (2001), Fractional integration in conditional duration. Manuscript.

Brunetti, C. & Lildholdt, P. (2002), Return-based and range-based (co)variance estimation - with an application to foreign exchange markets. Manuscript.

Chen, S. (2000), 'Probability density function estimation using gamma kernels', *Ann. Inst. Statist. Math.* **52**(3), 471–480.

Chou, R. (2001), Forecasting financial volatilities with extreme values: The conditional autoregressive range (CARR) model. Manuscript.

Cox, D. & Snell, E. (1968), 'A general definition of residuals', *Journal of the Royal Statistical Society. Series B* **30**(2), 248–275.

Engle, R. & Lunde, A. (1999), Trades and quotes: A bivariate point process. Manuscript.

Engle, R. & Russell, J. (1998), 'Autoregressive conditional duration: A new model for irregularly spaced transaction data', *Econometrica* **66**, 1127–1162.

Evans, M., Hastings, N. & Peacock, B. (1993), *Statistical Distributions*, John Wiley and Sons.

Gallant, A., Hsu, C. & Tauchen, G. (1999), 'Using daily range data to calibrate volatility diffusions and extract the forward integrated variance', *The Review of Economic and Statistics* **81**(4), 617–631.

Hasza, D. (1980), 'The asymptotic distribution of the sample autocorrelations for the integrated ARMA process', *Journal of the American Statistical Association* **75**, 349–352.

Hendry, D. (2000), *Econometrics: Alchemy or Science?*, Oxford University Press.

Jasiak, J. (1999), Persistence in intertrade durations. Manuscript.

Lamoureux, C. & Lastrapes, W. (1990), 'Persistence in variance, structural change and the GARCH model', *Journal of Business and Economic Statistics* **8**, 225–235.

Lee, S. & Hansen, B. (1994), 'Asymptotic theory for the GARCH(1,1) quasi maximum likelihood estimator', *Econometric Theory* **10**, 29–52.

Li, K. & Weinbaum, D. (2000), The empirical performance of alternative extreme value volatility estimators. Manuscript.

Lildholdt, P. (2002), Estimation of GARCH models based on open, close, high, and low prices. Manuscript.

Lunde, A. (1999), A generalized gamma autoregressive conditional duration model. Manuscript.

Nelson, D. & Cao, C. (1992), 'Inequality constraints in the univariate GARCH model', *Journal of Business and Economic Statistics* **10**, 229–235.

Parkinson, M. (1980), 'The extreme value method for estimating the variance of the rate of return', *Journal of Business* **53**(1), 61–65.

7 Tables and Figures

	Mean	St.Dev.	Skewness	Kurtosis
$R^{SF/USD}$	0.0069	0.0034	1.9245	9.55
$R^{USD/GBP}$	0.0056	0.0031	1.9222	9.40

Table 1: Summary statistics for daily price ranges.

	$R^{SF/USD}$	$R^{USD/GBP}$
$LB(1)$	222*	573.8*
$LB(15)$	1519.2*	4761.1*
$LB(30)$	2165.3*	7789.7*
$LB(50)$	2737.4*	11036.7*
$LB(100)$	3130.7*	16839.4*
$LB(250)$	3497.9*	24818.1*

Table 2: Ljung-Box statistics for daily price ranges. Rejection of the null of no serial correlation, at a five percent level, is denoted by *.

	CARR(2,2)	CARR(2,1)	CARR(1,2)	CARR(1,1)
$\widehat{w}$	0.00013 (.00004)	0.00027 (.00008)	0.00016 (.00005)	0.00022 (.00007)
$\widehat{a}_1$	0.14178 (.02108)	0.12241 (.01694)	0.14213 (.02119)	0.09553 (.01425)
$\widehat{a}_2$	-0.07886 (.02258)	0	-0.06659 (.02442)	0
$\widehat{b}_1$	1.03939 (.08151)	0.49541 (.12074)	0.90129 (.01813)	0.87295 (.02118)
$\widehat{b}_2$	-0.12139 (.08341)	0.34291 (.12347)	0	0
$\widehat{SEP}92$	0.00043 (.00021)	0.00080 (.00040)	0.00052 (.00025)	0.00061 (.00031)
$logL$	10895.08120	10894.77144	10895.06769	10894.23306
AIC	-21778.16240	-21779.54288	-21780.13537	-21780.46612
SIC	-21742.68128	-21749.97527	-21750.56777	-21756.81203
$LB(1)$	0.01926	1.16843	0.01374	6.02412*
$LB(15)$	12.75953	14.27363	12.55174	19.27550
$LB(30)$	26.34942	27.75567	26.19926	32.54778
$LB(50)$	43.09774	44.98081	43.02890	50.00172
$LB(100)$	79.69487	80.83788	79.54364	85.55877
$LB(250)$	188.40762	191.68716	188.57547	196.52490
Mean-th.mean	0.00024	0.00028	0.00025	0.00030
Std-th.std	-0.54912	-0.54842	-0.54906	-0.54735

Table 3: Semi-parametric CARR models for $R^{SF/USD}$. Robust standard errors in brackets. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for standardized innovations and the theoretical value of unity.

	CARR(2,2)	CARR(2,1)	CARR(1,2)	CARR(1,1)
$\widehat{w}$	0.00005 (.00002)	0.00015 (.00004)	0.00006 (.00002)	0.00012 (.00004)
$\widehat{a}_1$	0.20891 (.01984)	0.15954 (.01916)	0.21005 (.01994)	0.11951 (.01753)
$\widehat{a}_2$	-0.14578 (.02381)	0	-0.13600 (.02485)	0
$\widehat{b}_1$	1.01468 (.07749)	0.42459 (.06973)	0.91461 (.01749)	0.85841 (.02217)
$\widehat{b}_2$	-0.08732 (.07329)	0.38785 (.06951)	0	0
$\widehat{SEP}92$	0.00052 (.00030)	0.00123 (.00069)	0.00059 (.00033)	0.00093 (.00053)
$logL$	11570.61333	11569.03629	11570.56885	11567.52762
AIC	-23129.22666	-23128.07259	-23131.13770	-23127.05523
SIC	-23093.74554	-23098.50498	-23101.57010	-23103.40115
$LB(1)$	0.00536	5.64210*	0.00011	19.25986*
$LB(15)$	25.49626*	35.56088*	25.37068*	48.91251*
$LB(30)$	44.24465*	53.09440*	43.99266	66.77556*
$LB(50)$	68.95864*	84.32155*	68.85469*	100.54906*
$LB(100)$	114.20526	132.86273*	114.60486	149.19119*
$LB(250)$	253.63685	271.28362	255.10771	285.56505
Mean-th.mean	0.00047	0.00050	0.00049	0.00051
Std-th.std	-0.55110	-0.54882	-0.55097	-0.54775

Table 4: Semi-parametric CARR models for $R^{USD/GBP}$. Robust standard errors in parenthesis. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for standardized innovations and the theoretical value of unity.

	FICARR(2,d,2)	FICARR(2,d,1)	FICARR(1,d,2)	FICARR(1,d,1)
$\widehat{w}$	0.00011 (.00007)	0.00011 (.00009)	0.00011 (.00007)	0.00009 (.00005)
$\widehat{d}$	0.06640 (.06284)	0.06312 (.07918)	0.06181 (.06497)	0.09169 (.03066)
$\widehat{\phi}_1$	0.92586 (.06137)	0.97346 (.01022)	0.99955 (.04777)	0.97264 (.01083)
$\widehat{\phi}_2$	0.04626 (.06564)	0	-0.02531 (.04533)	0
$\widehat{b}_1$	0.85058 (.09602)	0.89486 (.08364)	0.92012 (.02936)	0.92725 (.02313)
$\widehat{b}_2$	0.06508 (.08227)	0.02355 (.05108)	0	0
$\widehat{SEP}92$	0.00047 (.00027)	0.00045 (.00028)	0.00045 (.00025)	0.00042 (.00023)
$logL$	10895.13334	10895.13306	10895.08781	10895.06278
AIC	-21776.26667	-21778.26611	-21778.17563	-21780.12557
SIC	-21734.87202	-21742.78499	-21742.69450	-21750.55796
$LB(1)$	0.01851	0.01824	0.02570	0.13293
$LB(15)$	12.88177	12.87652	12.82152	13.39530
$LB(30)$	26.03434	26.04533	25.96626	26.27759
$LB(50)$	42.74174	42.75553	42.74413	43.02246
$LB(100)$	79.37485	79.39527	79.26736	79.56550
$LB(250)$	187.60925	187.62668	187.49193	187.58613
Mean-th.mean	0.00029	0.00028	0.00038	0.00036
Std-th.std	-0.54931	-0.54930	-0.54920	-0.54919

Table 5: Semi-parametric FICARR models for $R^{SF/USD}$. Robust standard errors are in parenthesis. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for standardized innovations and the theoretical value of unity.

	FICARR(2,d,2)	FICARR(2,d,1)	FICARR(1,d,2)	FICARR(1,d,1)
$\widehat{w}$	0.00006 (.00003)	0.00006 (.00003)	0.00006 (.00003)	0.00006 (.00003)
$\widehat{d}$	0.49129 (.06176)	0.49167 (.06072)	0.49072 (.06075)	0.48919 (.05599)
$\widehat{\phi}_1$	0.41239 (.05628)	0.41319 (.10295)	0.40943 (.05350)	0.41166 (.04788)
$\widehat{\phi}_2$	-0.00157 (.03075)	0	-0.00146 (.03018)	0
$\widehat{b}_1$	0.69900 (.07549)	0.70018 (.12669)	0.69629 (.07780)	0.69726 (.05010)
$\widehat{b}_2$	-0.00044 (.00772)	0.00034 (.04827)	0	0
$\widehat{SEP}92$	0.00180 (.00100)	0.00177 (.00102)	0.00180 (.00100)	0.00180 (.00094)
$logL$	11571.07567	11571.07565	11570.98361	11570.98329
AIC	-23128.15134	-23130.15130	-23129.96723	-23131.96658
SIC	-23086.75669	-23094.67017	-23094.48610	-23102.39897
$LB(1)$	0.00541	0.00569	0.01201	0.01523
$LB(15)$	22.67706	22.68616	22.69850	22.70016
$LB(30)$	37.82009	37.84826	37.81927	37.82371
$LB(50)$	64.33296	64.33167	64.36620	64.37237
$LB(100)$	110.69273	110.65718	110.69505	110.65440
$LB(250)$	244.63787	244.60313	244.37025	244.20530
Mean-th.mean	-0.00096	-0.00094	-0.00099	-0.00095
Std-th.std	-0.55307	-0.55307	-0.55293	-0.55294

Table 6: Semi-parametric FICARR models for $R^{USD/GBP}$. Robust standard errors are in parenthesis. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for standardized innovations and the theoretical value of unity.

	$R^{SF/USD}$		$R^{USD/GBP}$	
	WE-CARR(1,2)	WE-FICARR(1,d,1)	WE-CARR(1,2)	WE-FICARR(1,d,1)
$\widehat{w}$	0.00024 (.00006)	0.00011 (.00019)	0.00008 (.00002)	0.00012 (.00004)
$\widehat{d}$		0.13228 (.02831)		0.38358 (.03995)
$\widehat{a}_1$	0.16545 (.01995)		0.20602 (.02085)	
$\widehat{a}_2$	-0.10527 (.02717)		-0.15991 (.02784)	
$\widehat{b}_1$	0.89335 (.02095)	0.92918 (.02985)	0.92038 (.01367)	0.66432 (.05551)
$\widehat{\phi}_1$		0.95608 (.01887)		0.47356 (.05907)
$\widehat{SEP}92$	0.00049 (.00030)	0.00037 (.00024)	0.00071 (.00027)	0.00228 (.00084)
$\widehat{\delta}$	2.30125 (.03012)	2.30229 (.03014)	2.31510 (.03036)	2.32508 (.03055)
$\widehat{\alpha}$	0	0	0	0
$logL$	12101.78723	12101.06548	12789.49972	12796.02099
AIC	-24191.57446	-24190.13097	-25566.99945	-25580.04197
SIC	-24156.09333	-24154.64984	-25531.51832	-25544.56085
$LB(1)$	0.11106	0.31377	0.00427	0.01428
$LB(15)$	6.12108	6.71985	27.62687*	21.51776
$LB(30)$	12.75710	12.66115	40.73496	34.43380
$LB(50)$	26.85175	26.95915	63.41004	56.38221
$LB(100)$	57.16571	58.81398	108.31067	99.40100
$LB(250)$	152.69068	151.87644	263.94935	244.86624
Mean-th.mean	0.00056	0.00067	0.00167	-0.00053
Std-th.std	0.48397	0.48260	0.42934	0.42635

Table 7: Parametric (FI)CARR models based on the Weibull distribution. Standard errors are in parenthesis. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for Cox-Snell residuals and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for Cox-Snell residuals and the theoretical value of unity.

	$R^{SF/USD}$		$R^{USD/GBP}$	
	GA-CARR(1,2)	GA-FICARR(1,d,1)	GA-CARR(1,2)	GA-FICARR(1,d,1)
$\widehat{w}$	0.00016 (.00004)	0.00009 (.00004)	0.00006 (.00002)	0.00006 (.00002)
$\widehat{d}$		0.09193 (.02667)		0.48990 (.04883)
$\widehat{a}_1$	0.14214 (.01864)		0.21005 (.01927)	
$\widehat{a}_2$	-0.06660 (.02154)		-0.13596 (.02252)	
$\widehat{b}_1$	0.90129 (.01504)	0.92751 (.01925)	0.91457 (.01243)	0.69779 (.04529)
$\widehat{\phi}_1$		0.97276 (.00864)		0.41146 (.04497)
$\widehat{SEP}92$	0.00052 (.00027)	0.00041 (.00023)	0.00059 (.00025)	0.00180 (.00071)
$\widehat{\delta}$	0	0	0	0
$\widehat{\alpha}$	5.95417 (.15680)	5.95390 (.15679)	6.07406 (.16004)	6.08544 (.16035)
$logL$	12347.06365	12347.03491	13046.64817	13049.16799
AIC	-24682.12730	-24682.06981	-26081.29634	-26086.33598
SIC	-24646.64617	-24646.58868	-26045.81521	-26050.85485
$LB(1)$	0.88853	1.29818	0.02656	0.03222
$LB(15)$	10.34527	10.74357	26.87987*	23.60701
$LB(30)$	21.67441	21.23823	42.31167	37.17976
$LB(50)$	37.63922	37.48296	66.49537	62.46784
$LB(100)$	70.47017	70.70475	112.01293	109.24246
$LB(250)$	181.59285	181.02023	264.17881	253.19349
Mean-th.mean	0.01584	0.01605	0.01949	0.01562
Std-th.std	0.27559	0.27436	0.28052	0.27211

Table 8: Parametric (FI)CARR models based on Gamma distributions. Standard errors are in parenthesis. "AIC" and "SIC" refers to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a five percent level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean for Cox-Snell residuals and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation for Cox-Snell residuals and the theoretical value of unity.

	$R^{SF/USD}$		$R^{USD/GBP}$	
	GA-CARR(1,2)	GA-FICARR(1,d,1)	GA-CARR(1,2)	GA-FICARR(1,d,1)
Mean	1.01584	1.01605	1.01949	1.01562
Std	1.27559	1.27436	1.28052	1.27211
Skew	4.37086	4.34275	3.75244	3.75421
Kurt	43.24276	42.70847	27.02591	27.57099
ED test	11.57830	11.52042	11.81116	11.41432
# excl. extremes	20	20	25	25
Mean	0.95600	0.95637	0.95051	0.94766
Std	1.03056	1.03086	1.03360	1.02962
Skew	2.17782	2.17284	2.26138	2.24429
Kurt	8.57259	8.52109	9.06255	8.89708
ED test	1.13542	1.14663	1.24982	1.09852

Table 9: Upper panel shows summary statistics for Cox-Snell residuals from the GA-(FI)CARR models. The lower panel shows the same summary statistics, excluding the largest 20/25 Cox-Snell residuals. The Excess Dispersion (ED) test follows a standard normal distribution under the null of the standard deviation being equal to unity.

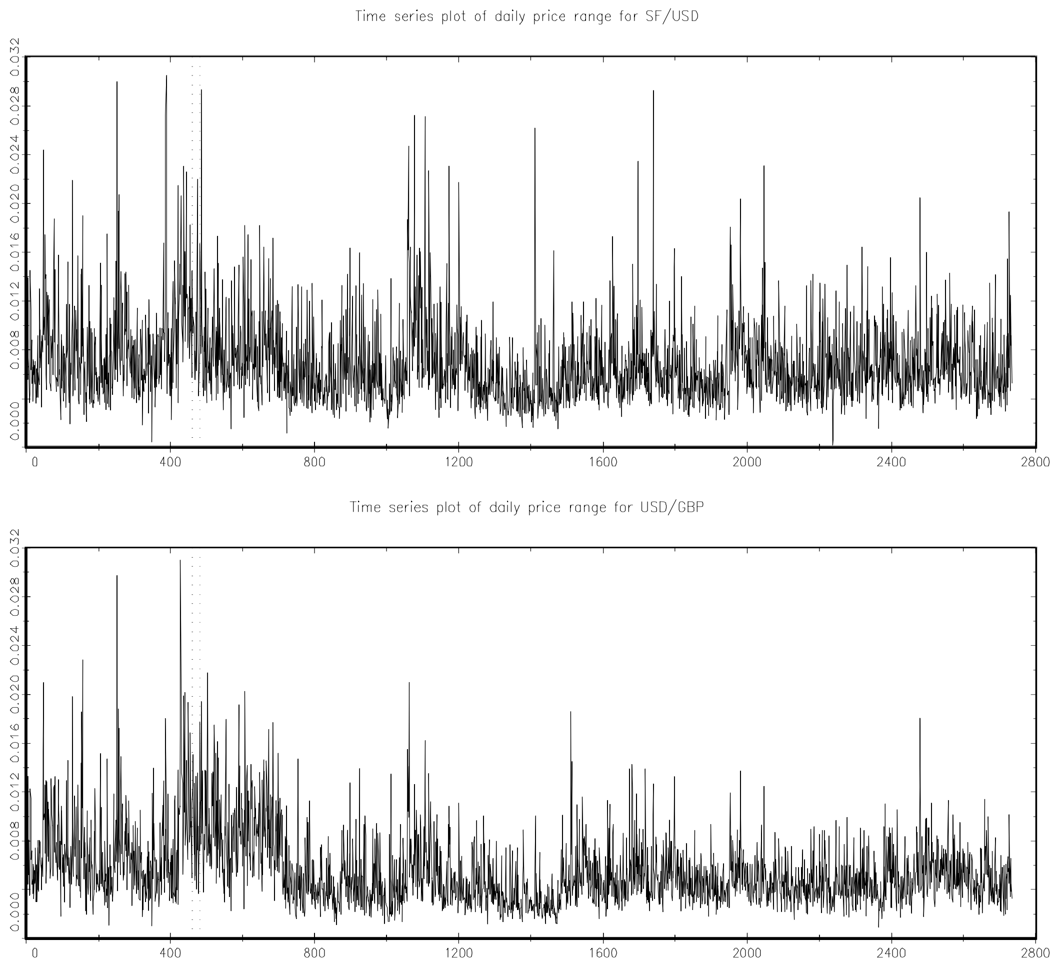

Figure 1: Time series plot of daily price ranges for $R^{SF/USD}$ in the upper panel, and $R^{USD/GBP}$ in the lower panel. The horizontal lines at observation number 461 and 481 denote the beginning and end of September 1992 (UK leaving ERM).

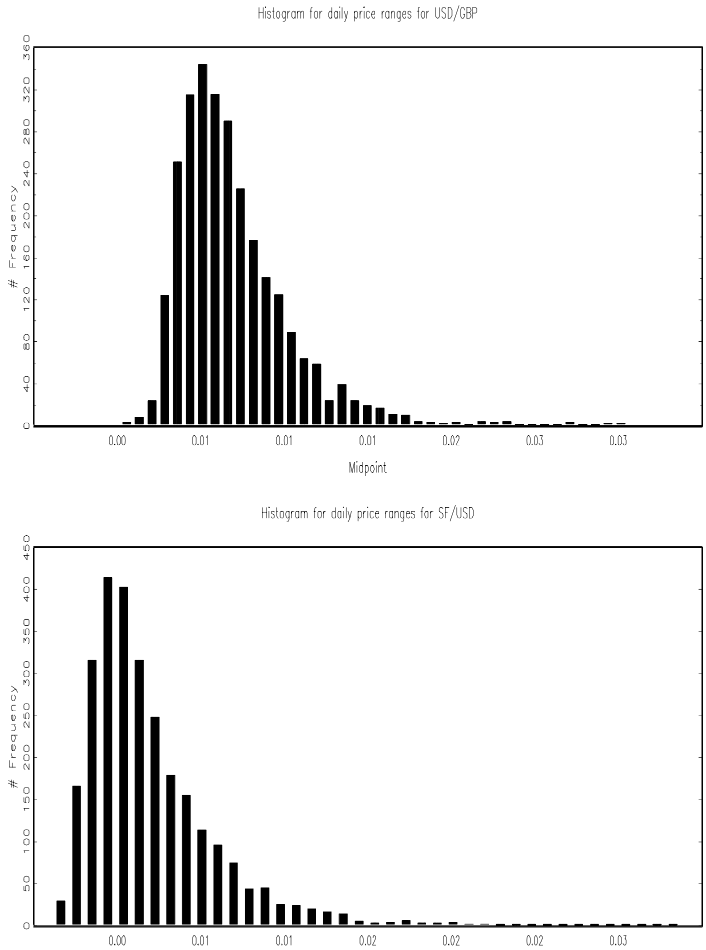

Figure 2: Upper panel shows histogram for $R^{USD/GBP}$ and lower panel shows histogram for $R^{SF/USD}$.

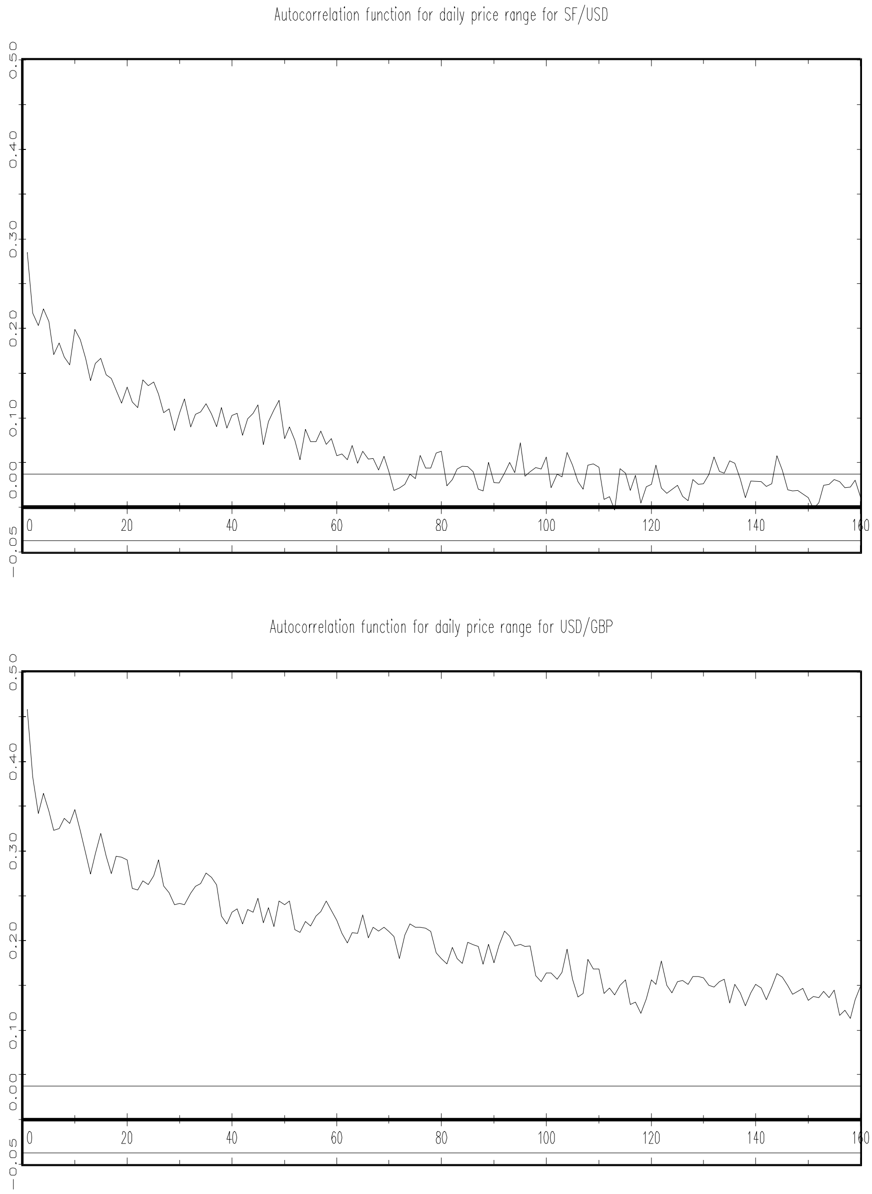

Figure 3: Autocorrelation functions for daily price ranges corresponding to $R^{SF/USD}$ in the upper panel and $R^{USD/GBP}$ in the lower panel. Horizontal lines correspond to critical values (at five percent level) for testing the null of the autocorrelation coefficient equal to zero.

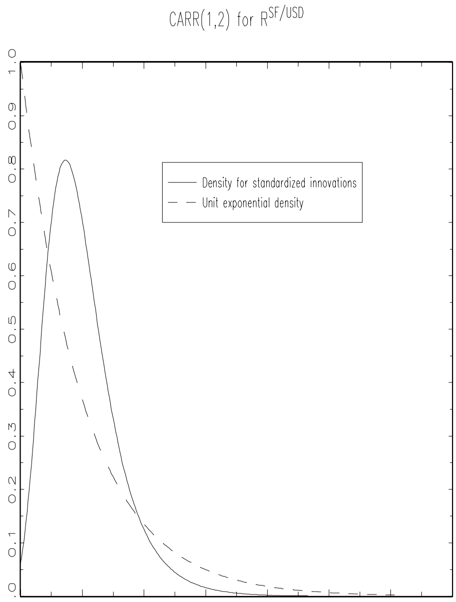

Figure 4: Empirical density of standardized innovations from CARR(1,2) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

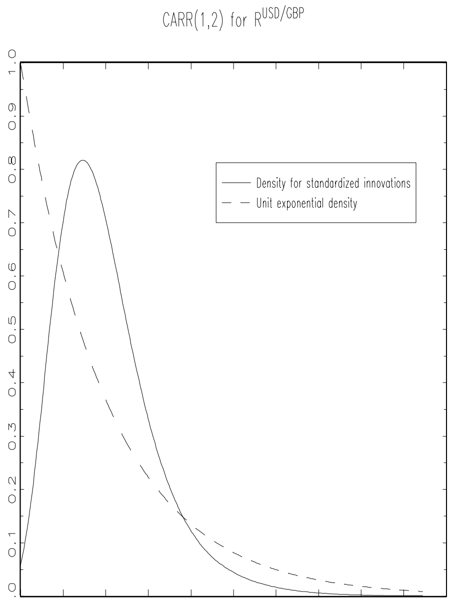

Figure 5: Empirical density of standardized innovations from CARR(1,2) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

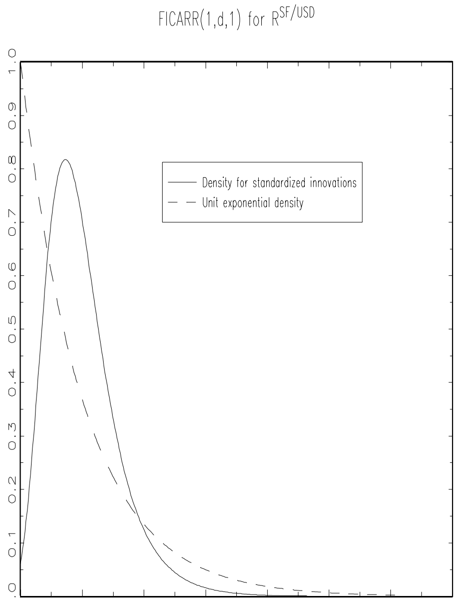

Figure 6: Empirical density of standardized innovations from FICARR(1,d,1) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

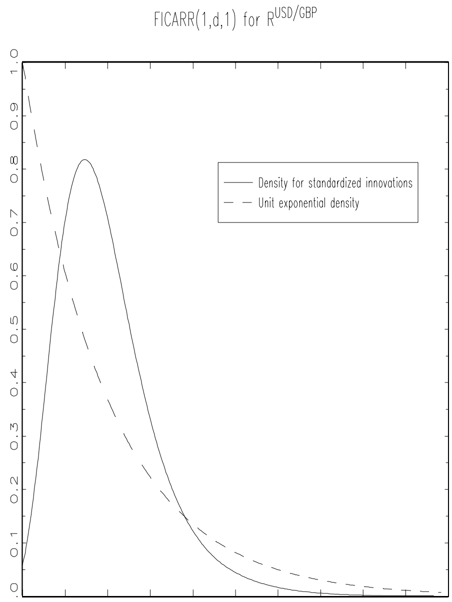

Figure 7: Empirical density of standardized innovations from FICARR(1,d,1) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

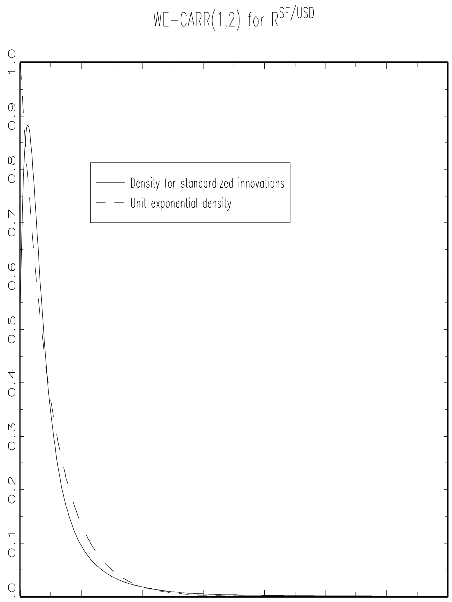

Figure 8: Empirical density of Cox-Snell residuals from WE-CARR(1,2) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

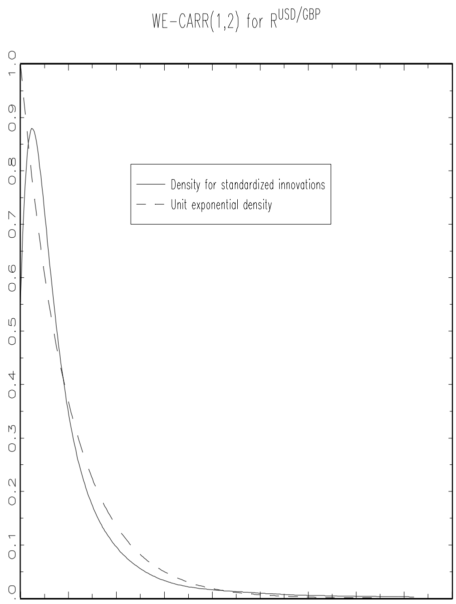

Figure 9: Empirical density of Cox-Snell residuals from WE-CARR(1,2) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

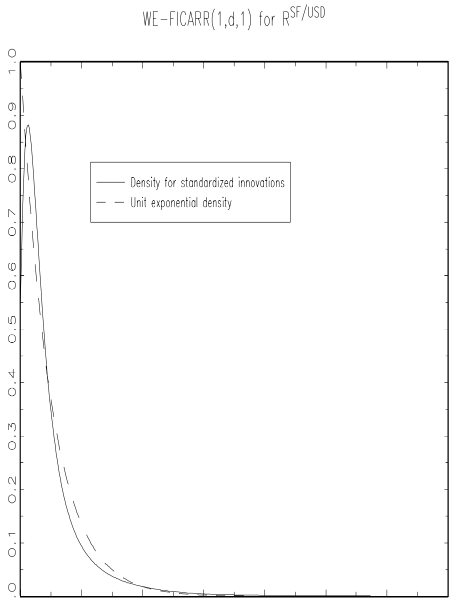

Figure 10: Empirical density of Cox-Snell residuals from WE-FICARR(1,d,1) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

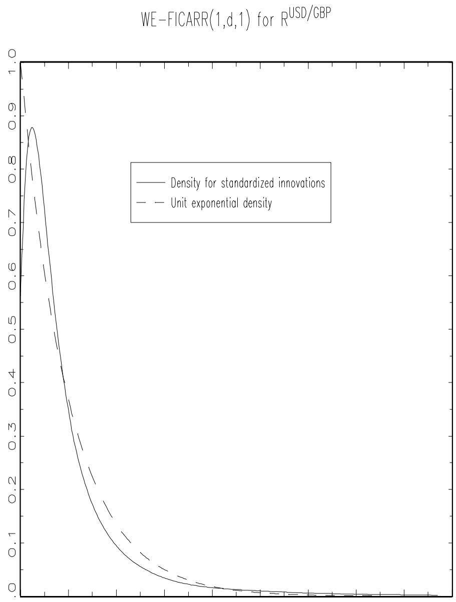

Figure 11: Empirical density of Cox-Snell residuals from WE-FICARR(1,d,1) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

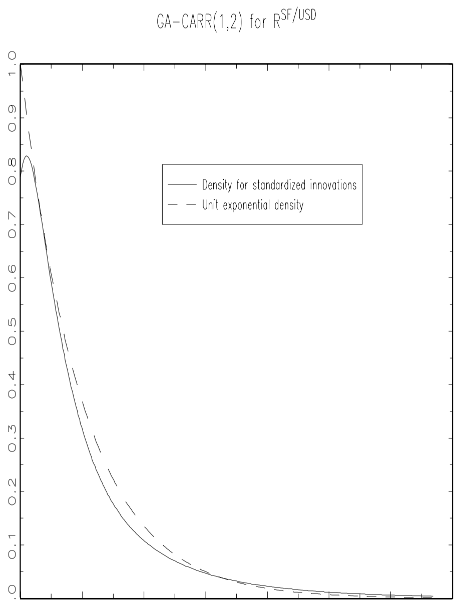

Figure 12: Empirical density of Cox-Snell residuals from GA-CARR(1,2) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

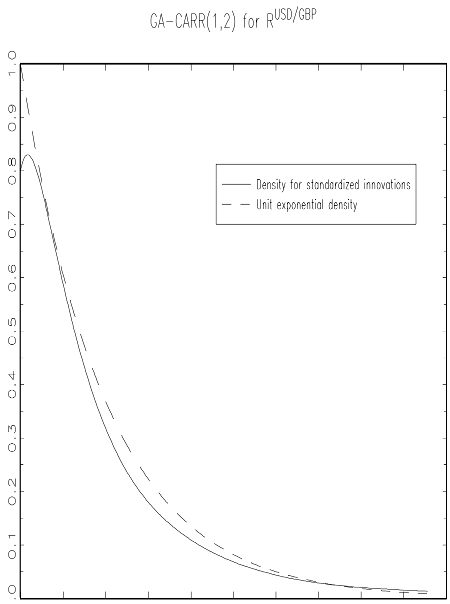

Figure 13: Empirical density of Cox-Snell residuals from GA-CARR(1,2) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

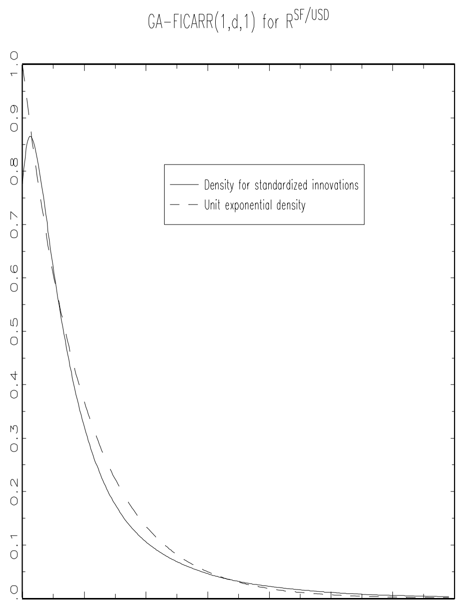

Figure 14: Empirical density of Cox-Snell residuals from GA-FICARR(1,d,1) model for $R^{SF/USD}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.

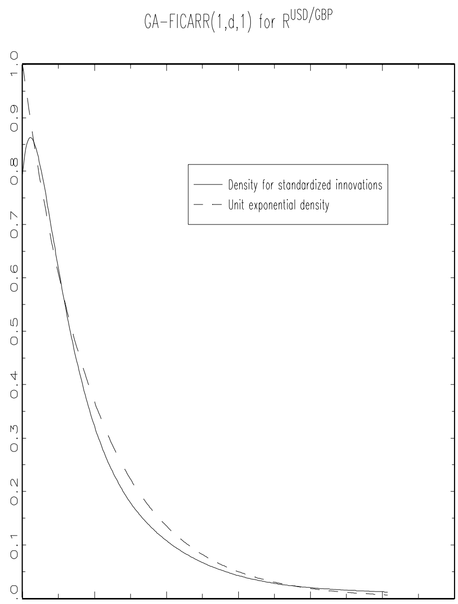

Figure 15: Empirical density of Cox-Snell residuals from GA-FICARR(1,d,1) model for $R^{USD/GBP}$ is depicted by the solid line. The density is estimated by the Gamma kernel of Chen (2000) with a bandwidth of 0.15. The unit Exponential density is depicted by the dotted line.