Time Series Modeling of Daily Log-Price Ranges for CHF/USD and USD/GBP

CELSO BRUNETTI AND PETER M. LILDHOLDT

CELSO BRUNETTI
is an assistant professor in the Department of Finance at Johns Hopkins University in Baltimore, MD.
celsob@jhu.edu

PETER M. LILDHOLDT
is a junior fellow in CREATES, School of Economics and Management at the University of Aarhus and senior analyst at Danske Bank in Copenhagen, Denmark.
plildholdt@econ.au.dk

The aim of this article is to model the dynamic evolution of daily price ranges for two foreign exchange rates, CHF/USD and USD/GBP. Following Engle [2002] and Chou [2005], we adopt the Multiplicative Error Model (MEM). The range is a highly efficient measure of daily volatility, and our empirical results provide insights into the volatility dynamics for CHF/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 MEM. Semi-parametric and parametric models are estimated. In a simple forecasting exercise we show that (FI)MEMs are able to predict daily ranges rather well. Moreover, using straddles and strangles trading strategies, we show that volatility forecasts from (FI)MEMs may have significant economic value.

Time series modeling 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 many fields (e.g., risk management and option pricing). A complication in volatility modeling is that volatility is not observable. However, several proxies for volatility exist: squared returns, absolute returns, realized volatility measures, and so forth. 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, modeling the volatility proxy directly (e.g., Andersen, Bollerslev, Diebold, and Labys [2003] model realized volatility measures; and Beran and Ocker [2001] model power transformations of absolute returns). We follow the latter approach in this article. In particular, we employ a nonstandard volatility proxy: the daily price range.[1]

Several reasons motivate our interest in the range. First, the range is a very efficient estimator of volatility. Parkinson [1980], and Brunetti and Lildholdt [2004] show that the daily range is more efficient than return-based volatility estimators. Second, historical data on high and low financial asset prices, required for the computation of the range, are available for several decades; this is not the case for high frequency data which, for example, are necessary for the computation of realized volatility. Third, the range contains valuable information even when compared to realized volatility. Engle and Gallo [2006] analyze the three most common measures of volatility: daily squared return, daily range, and realized volatility. They find that the range carries additional information to realized volatility. Fourth, realized volatility—the sum of (intraday) squared returns sampled at high frequencies—is contaminated by market microstructure noise. Several solutions have been proposed in the literature to overcome this problem. The issue of which of these solutions should be used is still open. Perhaps more importantly, Brandt and Diebold [2006] provide evidence that the range is not affected by market microstructure noise.

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 CHF/USD. Following Engle [2002], and Chou [2005], we adopt the Multiplicative Error Model (MEM) introduced in the literature by Engle and Russell [1998].[2] This family of models is very flexible and is a natural choice given that they model all variables defined on the real positive line.

The main goal of the article is to study the dynamic properties of the range. 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 MEM to allow for fractional integration (see Jasiak [1999]), the dynamic properties of daily price ranges for USD/GBP seem to be well described by the Fractionally Integrated MEM, or FIMEM(1,d,1). In contrast, the short memory MEM(1,2) appears to be the appropriate specification for the CHF/USD price range. We also find that the Quasi-maximum likelihood estimation (QMLE) procedure is able to produce accurate parameter estimates and robust inference.

Finally, we check the forecasting performance of (FI)MEMs. We adopt two different approaches. The first is one-step-ahead forecasting using the selected models. We find that (FI)MEMs are indeed able to forecast the range. The second forecasting exercise involves straddles and strangles trading strategies. The volatility forecasts are plugged in the Black-Scholes option pricing formula to compute straddles and strangles prices. Those strategies are bought if the market price is lower than the forecasted price and sold if the market price is higher than the forecasted price. The same forecasting exercise is also conducted employing the widely used GARCH(1,1) model. We demonstrate that (FI)MEM volatility forecasts are indeed profitable. This is not the case for the GARCH(1,1) forecasts.

The structure of the article is as follows: The second section motivates the approach of the article, and the third section describes the data. In the fourth section, we present the MEM of Engle [2002] and introduce the FIMEM. In the fifth section, we present semi-parametric and parametric estimates of the (FI)MEM. The sixth section is devoted to forecasting exercises, and the last section concludes.

MOTIVATION

In this section, we explain our interest in modeling 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 $P_{t,s}^{CHF/USD}$. Further, assume $P_{t,s}^{CHF/USD}$ evolves according to a driftless Brownian Motion with diffusion coefficient σ_s

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

The daily price range on day s is defined by

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

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)

$$(P_{T,s}^{CHF/USD} - P_{0,s}^{CHF/USD})^2$$

absolute returns,

$$|P_{T,s}^{CHF/USD} - P_{0,s}^{CHF/USD}|$$

or variations on the estimators above based on intraday returns (e.g., Andersen, Bollerslev, Diebold, and Labys [2001], and Barndorff-Nielsen and Shephard [2003]).

Parkinson [1980] shows that the unbiased range-based estimator of the squared diffusion coefficient,

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

is approximately 5 times more efficient than the unbiased estimator based on the daily squared returns. Brunetti and Lildholdt [2004] show that the unbiased estimator of $\sigma_s^{CHF/USD}$ based on the range,

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

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

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

The normalization constant in Equation (2) rescales the range to produce an unbiased estimator of the standard deviation. The derivation comes directly from the distribution of the range developed by Feller [1951].[3]

In Equation (1), the diffusion coefficient stays constant during the day. Andersen and Bollerslev [1998, footnote 20] adopt a stochastic volatility model where the diffusion coefficient evolves continuously. 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 also, the daily price range provides a highly efficient estimator of daily volatility—integrated volatility (see also Christensen and Podolski [2005]).

In an empirical study, Li and 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, and so forth.

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. However, return-based estimates of the volatility process are based on the opening and closing prices only—no intraday information is used.

Motivated by these appealing features, we would like to scrutinize the dynamic properties of scaled (log of the price) ranges for the CHF/USD and USD/GBP FX rates,

$$R_s^{CHF/USD} = \sqrt{\frac{\pi}{8}} l_s^{CHF/USD} \tag{3}$$

and likewise

$$R_s^{USD/GBP} = \sqrt{\frac{\pi}{8}} l_s^{USD/GBP} \tag{4}$$

DATA

We analyze data on the foreign exchange rates, CHF/USD and USD/GBP, over the period January 3, 1991 to May 31, 2002. We divide the sample into two subsamples. The first subsample, January 3, 1991 to September 20, 2001, is used for in-sample estimation, while the second subsample is used for forecasting.

Data are from Bloomberg and correspond to intraday 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. 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 obtain high/low transaction prices spanning an eleven-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, therefore, might 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 U.K. was part of the EMS and experienced high volatility in exchange rates as a consequence of severe speculative attacks.

Exhibits 1 and 2 show time series plots of daily price ranges (see Equations (3) and (4)) for the two rates. Both series appear to be stationary, and the September 1992 speculative attack 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 greater 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 U.K. left the EMS. In Exhibit 3, we report summary statistics for daily price ranges. The mean and standard deviation of USD/GBP price ranges are lower than those of CHF/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. Exhibit 4 reports Ljung-Box (LB) statistics for serial correlation up to lag(s), 1, 15, 30, 50, 100 and 250. All statistics are significant at the 5% significance level,[4] and it is evident that the series are highly autocorrelated. The values of the LB-test are particularly high for the USD/GBP range. This finding is also supported by Exhibits 5 and 6, where autocorrelation functions (ACF) for the two series are shown. It is evident that the

EXHIBIT 1
Time Series Plot for $R^{CHF/USD}$

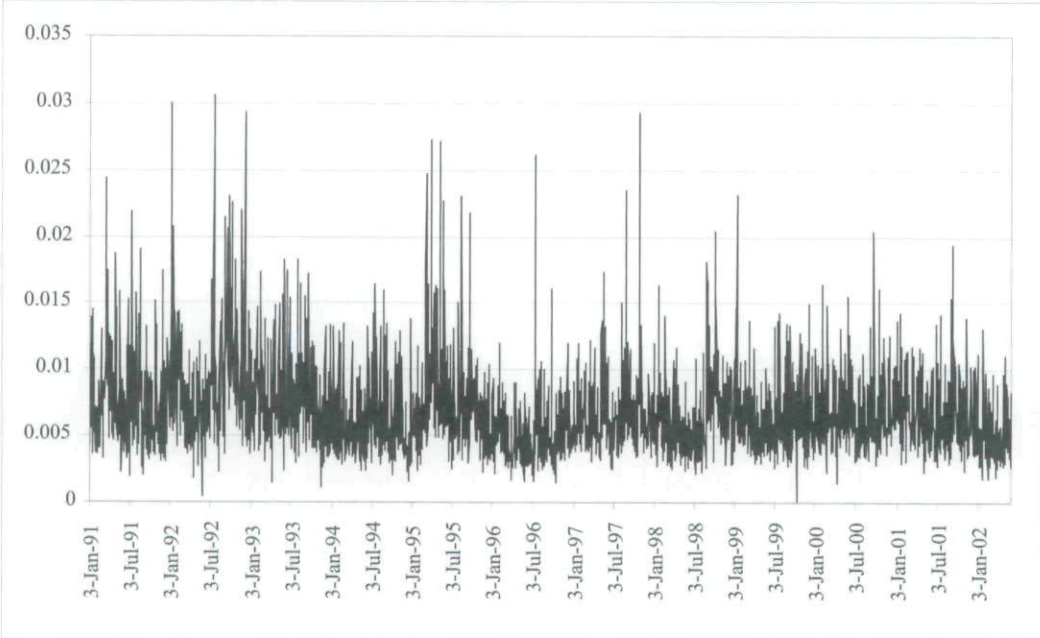

EXHIBIT 2
Time Series Plot for $R^{USD/GBP}$

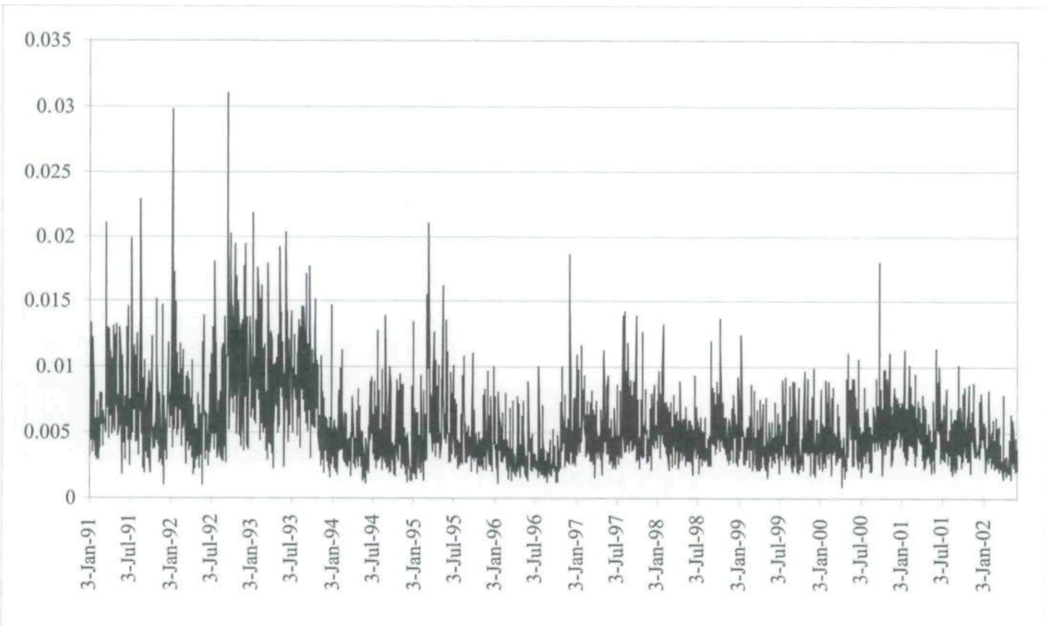

EXHIBIT 3
Summary Statistics for Daily Price Ranges

	Mean	St.Dev.	Skewness	Kurtosis
$R^{CHF/USD}$	0.0069	0.0034	1.9320	9.6612
$R^{USD/GBP}$	0.0055	0.0030	1.9614	9.6585

EXHIBIT 4
Ljung-Box Statistics for Daily Price Ranges

	$R^{CHF/USD}$	$R^{USD/GBP}$
LB (1)	237.3*	623.4*
LB (15)	1688.7*	5294.8*
LB (30)	2431.5*	8725.4*
LB (50)	3093.1*	12376.1*
LB (100)	3566.1*	18804.2*
LB (250)	4000.6*	27529.5*

*Note: Rejection of the null of no serial correlation, at a 5% significance level, is denoted by *.*

USD/GBP range is very persistent. The ACF is always positive and remains significantly different from zero for high lags. In contrast, the CHF/USD range is less persistent and the ACF is statistically insignificant at lag 71.

MULTIPLICATIVE ERROR MODEL

Following Engle [2002], the Multiplicative Error Model (MEM) might be defined by

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

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

where R_s is the daily price range and $a(L) = a_1 L + a_2 L^2 + \cdots + a_p L^p$ and $b(L) = b_1 L + b_2 L^2 + \cdots + b_q L^q$ are lag polynomials of order p and q, respectively. The variable λ_s is the conditional mean of the range. The model is identical to the Autoregressive Conditional Duration (ACD) model of Engle and Russell [1998], and CARR model of Chou [2005]. This is a very appealing model because, as shown by Engle and 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 $\varepsilon_s = R_s - \lambda_s$, the model may be expressed as an ARMA process in R_s

EXHIBIT 5
ACF for $R^{CHF/USD}$

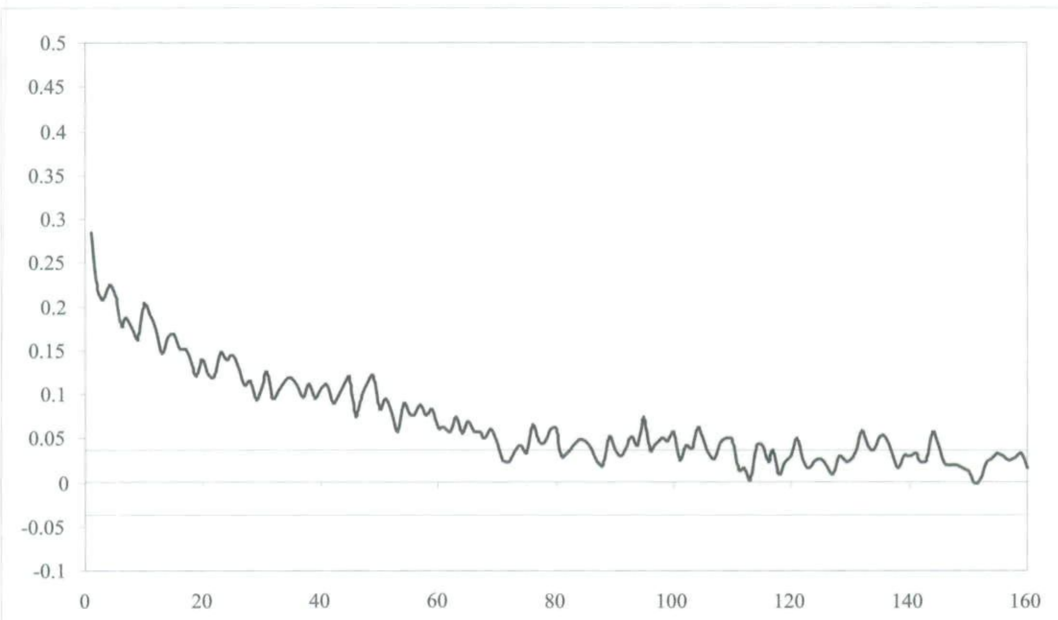

EXHIBIT 6
ACF for $R^{USD/GBP}$

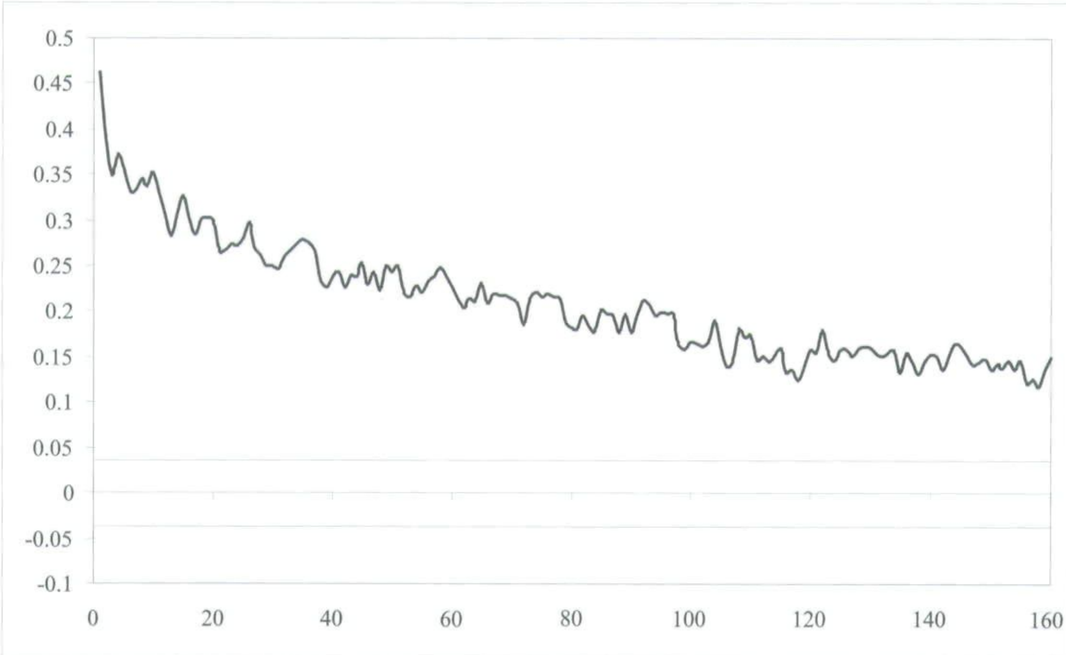

$$(1 - a(L) - b(L))R_s = \omega + (1 - b(L))\varepsilon_s \tag{8}$$

Using the results from the vast GARCH literature, it follows that the MEM 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 (e.g., Baillie, Bollerslev, and Mikkelsen [1996]; Baillie [1996], and Andersen et al. [2003]). For this reason, we introduce the long memory extension of the MEM (see Jasiak [1999]). To create more persistent impulse responses for the model, a fractional root in the autoregressive polynomial for R_s is introduced in Equation (8),

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

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

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

or, in a more parsimonious form

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

where $\gamma(L) = \frac{[1-b(L) - \phi(L)(1 - L)^d]}{1-b(L)}$, 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 MEM (FIMEM). 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.[5] The FIMEM is similar to the FIGARCH of Baillie et al. [1996]; parallel to the FIGARCH specification, the unconditional mean of R_s is infinite. This implies that FIMEM is not weakly stationary. However, following Nelson [1990], the FIMEM is stationary and ergodic for $0 \le d \le 1$.

We have not specified the distribution of the innovations, ε_s, in Equation (5) yet. The ACD literature concentrates on the family of Gamma densities,

$$f(\varepsilon_s|\Lambda_{s-1}) = \frac{1}{\Gamma(\alpha)\beta^{\alpha}}\varepsilon_s^{\alpha-1}\exp\left(-\frac{\varepsilon_s}{\beta}\right) \quad (13)$$

From Equation (5), we know that $E(\varepsilon_s) = 1$ which implies that $\beta = \alpha^{-1}$, therefore, Equation (13) becomes

$$f(\varepsilon_s|\Lambda_{s-1}) = \frac{1}{\Gamma(\alpha)}\alpha^{\alpha}R_s^{\alpha-1}\lambda_s^{1-\alpha}\exp\left(-\alpha\frac{R_s}{\lambda_s}\right) \quad (14)$$

Engle and Gallo [2006] derive the log-likelihood of the above conditional density and show two important results. First, the maximization of the log-likelihood from Equation (14) with respect to the parameters of interest does not depend on α. Second, the robust variance-covariance matrix (see Lee and Hansen [1994]) also does not depend on α. These findings imply that we can use the simple exponential density ($\alpha = 1$) to estimate the parameters of interest. The log-likelihood for the exponential density is given by

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

Moreover, this is a quasi-maximum likelihood estimation (QMLE) method that provides consistent and asymptotically normal estimates with a covariance matrix given by the familiar robust standard errors from Lee and Hansen [1994] (see Engle and Russell [1998]).

Estimation of the FIMEM involves truncating the infinite lag-polynomial in Equation (12) at lag 1,000 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,

$$\hat{\varepsilon}_s = \frac{R_s}{\hat{\lambda}_s}$$

where $\hat{\lambda}$ denotes the conditional expectation of the range computed via Equation (6) or Equation (12) by use of estimated coefficients.

EMPIRICAL RESULTS

Semi-parametric MEMs

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 MEMs: MEM(2,2), MEM(2,1), MEM(1,2), and MEM(1,1).

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

- general-to-specific (e.g., Hendry [2000]); testing down from a general model by t-statistics of individual parameters;
- log-likelihood; for nested models, LR-tests are adopted, and for non-nested models, Akaike's (AIC) and Schwartz's (SIC) information criteria are used; and
- ability of the model to fit the assumption of *iid* innovations; Ljung-Box statistics for serial correlation in the standardized innovations are used.

Exhibit 7 shows the estimation results for the CHF/USD range. The most general model, the MEM(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. The restrictions in Equation (7) are sufficient, but not necessary, conditions for the conditional range to be positive (see Nelson and Cao [1992]). Using the general-to-specific approach, the preferred model is the MEM(1,2). However, the LR-test, the AIC, and the SIC select the simple MEM(1,1). The dispute between the MEM(1,2) and the MEM(1,1) is resolved by looking at the LB statistics. It is evident that the values of the Ljung-Box statistics are much lower for the MEM(1,2) than for the MEM(1,1). A fundamental assumption in the MEM (see Equation (5)) is that the innovation term is *iid*. Exhibit 7 shows that significant autocorrelation at lag 1 remains in the standardized innovations for the MEM(1,1), and hence the MEM(1,2) is preferred.

For each specification we estimate two models: with and without a dummy variable for the September 1992 speculative attack. It is evident that the speculative attack of September 1992 against the GBP rates increased volatility dramatically even for the CHF/USD rate.

EXHIBIT 7
Semi-parametric MEMs for $R^{CHF/USD}$

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

	MEM(2,2)	MEM(2,1)	MEM(1,2)	MEM(1,1)
$\widehat{\omega}$	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
$\widehat{SEP92}$	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

*Note: Robust standard errors are in parenthesis. "AIC" and "SIC" refer to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a 5% significance level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean of standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation of standardized innovations and the theoretical value of unity.*

However, the introduction of the dummy does not materially modify the parameter estimates.[6] A common characteristic of all the estimated models is that volatility is very persistent: the sum of the coefficients in the lag polynomials $a(L)$ and $b(L)$ ranges from 0.96 to 0.98. This result is in line with the findings of the GARCH literature.

The last two rows in Exhibit 7 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.

Exhibit 8 shows the MEM estimation results for USD/GBP. There are remarkable similarities between the results for USD/GBP and CHF/USD (see Exhibits 7 and 8). In the MEM(2,2), diagnostics support the restriction of $b_2 = 0$. The preferred model using the general-to-specific approach and the Ljung-Box statistic is the MEM(1,2). If we use the LR test, the AIC, and the SIC, the selected model is the MEM(1,1). However, the MEM(1,2) accounts for serial correlation in data much better than MEM(1,1). Therefore, the MEM(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 CHF/USD rate. We also estimate the same models omitting the dummy variable (results not reported) and obtain very similar parameter estimates to those in Exhibit 8.

The general impression from Exhibits 7 and 8 is that persistency is strong and ranges from 0.99 for the MEM(2,2) to 0.978 for the MEM(1,1).

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

Semi-parametric FIMEMs

For the FIMEM we also estimate four different specifications: FIMEM(2,d,2), FIMEM(2,d,1), FIMEM(1,d,2), and FIMEM(1,d,1).[7]

Exhibit 8

Semi-Parametric MEMs for $R^{USD/GBP}$

	$R_s = \lambda_s \varepsilon_s, \quad \varepsilon_s \sim \exp(1)$			
	$\lambda_s = \omega + a_1 R_{s-1} + a_2 R_{s-2} + b_1 \lambda_{s-1} + b_2 \lambda_{s-2} + SEP92 \cdot DUMs$			
	MEM(2,2)	MEM(2,1)	MEM(1,2)	MEM(1,1)
$\widehat{\omega}$	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
$\widehat{SEP92}$	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

*Note: Robust standard errors are in parenthesis. "AIC" and "SIC" refer to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a 5% significance level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean of standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation of standardized innovations and the theoretical value of unity.*

Exhibit 9 shows the results of the FIMEMs for the CHF/USD range. The long memory parameter, d, is insignificant except for FIMEM(1,d,1). Testing down from the general FIMEM(2,d,2), 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 FIMEM(1,d,1). 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 FIMEM provides a suitable model for CHF/USD.

The results of the FIMEM estimates for USD/GBP are shown in Exhibit 10. The results are unambiguous: regardless of the selection criteria used, the FIMEM(1,d,1) is always the preferred model. In fact, the diagnostics favor the restrictions $\phi_2 = b_2 = 0$.[8] 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 a 10% significance level). We also estimate the same FIMEM without the September 1992 dummy variable (results not reported). In line with the findings obtained for the MEM, the introduction of the dummy does not materially modify the parameter estimates.

In the following section, we compare the preferred MEMs to the preferred FIMEMs.

Comparison of Semi-parametric MEMs and FIMEMs

We first analyze the USD/GBP series. For the general-to-specific approach, it is important to note that the MEM(2,2) is nested in the FIMEM(2,d,2)—see Equations (8) and (9). Therefore, using the general-to-specific approach, the preferred model is the FIMEM(1,d,1). Both models, MEM(1,2) and FIMEM(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 FIMEM(1,d,1). This

EXHIBIT 9
Semi-parametric FIMEMs for $R^{CHF/USD}$

$$R_s = \lambda_s \varepsilon_s, \qquad \varepsilon_s \sim \exp(1)$$
$$\lambda_s = \omega + [1 - b_1 L - b_2 L^2 - (1 + \phi_1 L + \phi_2 L^2)(1 - L)^d] R_s + b_1 \lambda_{s-1} + b_2 \lambda_{s-2} + SEP92 \cdot DUM_s$$

	FIMEM(2,d,2)	FIMEM(2,d,1)	FIMEM(1,d,2)	FIMEM(1,d,1)
$\widehat{\omega}$	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
$\widehat{SEP92}$	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

*Note: Robust standard errors are in parenthesis. "AIC" and "SIC" refer to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where * denotes rejection (at a 5% significance level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean of standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation of standardized innovations and the theoretical value of unity.*

is particularly true for long lags. It seems that the FIMEM(1,d,1) is able to account for both short and the long range dependence. It is quite interesting that the preferred MEM is of order (1,2). Note that a MEM(1,2) may be rewritten as an ARMA(2,1) model in daily price ranges—see Equation (9). Gallant, Hsu, and 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 FIMEM(1,d,1) for the USD/GBP range reduce dramatically when compared to those of the MEMs. 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 CHF/USD, the preferred model is the MEM(1,2). The FIMEM(1,d,1) for the CHF/USD range (last column in Exhibit 9), 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 MEM(1,2) (see Exhibit 7). When $d = 0$, the FIMEM collapses into the MEM with $\phi(L) = 1 - a(L) - b(L)$. The parameters of the FIMEM(1,d,2) are not individually significantly different from the corresponding parameters from the MEM(1,2),[9] which provides further evidence in favor of the MEM(1,2).

The results above are consistent with the empirical autocorrelation functions. The ACF for CHF/USD range

EXHIBIT 10

Semi-parametric FIMEMs for $R^{USD/GBP}$

$$R_s = \lambda_s \varepsilon_s, \qquad \varepsilon_s \sim \exp(1)$$

$$\lambda_s = \omega + [1 - b_1L - b_2L^2 - (1 + \phi_1 L + \phi_2 L^2)(1-L)^d] R_s + b_1\lambda_{s-1} + b_2\lambda_{s-2} + SEP92 \cdot DUM_s$$

	FIMEM(2,d,2)	FIMEM(2,d,1)	FIMEM(1,d,2)	FIMEM(1,d,1)
$\widehat{\omega}$	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
$\widehat{SEP92}$	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

Note: Robust standard errors are in parenthesis. "AIC" and "SIC" refer to Akaike's and Schwartz's information criterion, respectively. LB(n) are Ljung-Box statistics, where ∗ denotes rejection (at a 5% significance level) of the null of no serial correlation from lag 1 to n. "Mean-th.mean" refers to the difference between the empirical mean of standardized innovations and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation of standardized innovations and the theoretical value of unity.

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.[10]

Parametric (FI)MEMs

The focus of this section is to fit parametric (FI)MEMs 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: MEM(1,2) and FIMEM(1,d,1). If the models are correctly specified, the standardized innovations are distributed according to an Exponential distribution. In Exhibit 11, two densities are depicted. The density of the standardized innovations from the MEM(1,2) for $R_s^{CHF/USD}$ (solid line) is shown in addition to a unit Exponential density (dotted line). We estimate density functions by the Gamma kernel proposed in Chen [2000].[11] This procedure overcomes the traditional boundary bias of standard kernel density estimators associated with estimating probability density functions with bounded support. It is quite clear from Exhibit 11 that the model is mis-specified: the two density functions are far from being identical. The same conclusion is valid for the FIMEM(1,d,1) for $R^{USD/GBP}$ (see Exhibit 12). The mis-specification is clear also from the last rows of Exhibits 7 to 10; in fact, the standard deviation of the standardized innovations are far in excess of the theoretical value of unity.

Weibull (FI)MEMs. A good candidate for the density of the standardized innovations is the Weibull distribution,[12] because it allows for the hump-shape of the price ranges. Hence, we estimate Weibull (FI)MEMs, abbreviated by WE-(FI)MEM.[13] The use of the Weibull

EXHIBIT 11
Density of Standardized Innovations, MEM(1,2), $R^{CHF/USD}$

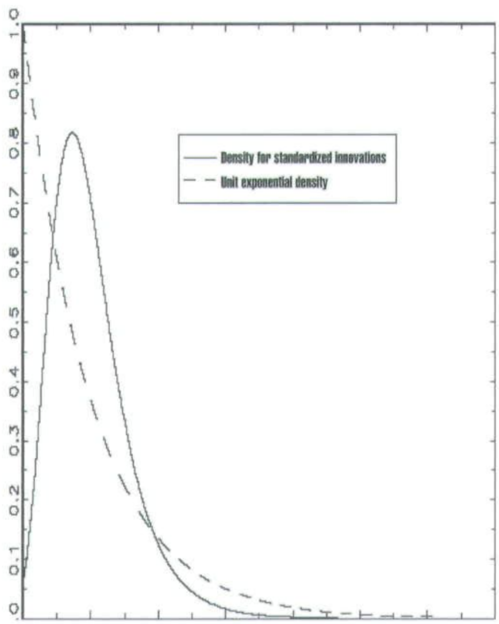

EXHIBIT 12
Density of Standardized Innovations, FIMEM(1,d,1), $R^{USD/GBP}$

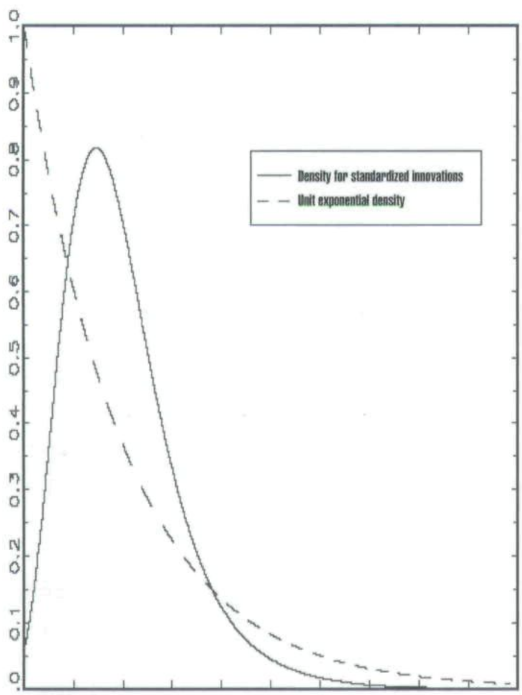

distribution in the ACD framework was proposed by Engle and Russell [1998]. Following Lunde [1999], we compute Cox-Snell residuals (see Cox and Snell [1968]) for the parametric (FI)MEMs. These residuals are exponentially distributed, which allows for a direct comparison of the fit of various parametric models. They are defined by

$$u_s = \int_0^{\varepsilon_s} h_\varepsilon(x)dx \tag{15}$$

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

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

where $f_\varepsilon(\cdot)$ denotes the density function for ε_s, and $F_\varepsilon(\varepsilon_s)$ denotes the cumulative distribution. Rewriting Equation (15),

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

it follows that the Cox-Snell residuals follow an Exponential distribution.[14] Cox-Snell residuals for WE-(FI)MEM are computed according to

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

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

The estimated parameters and standard errors[15] are in line with those obtained with the QMLE in Exhibits 7 to 10. Exhibits 13 and 14 show empirical densities of Cox-Snell residuals from the WE-(FI)MEMs. The fit improves relative to Exhibits 11 and 12, however, the standard deviation of the Cox-Snell residuals is far in excess of the theoretical value of unity for all models.

Gamma (FI)MEMs. The Weibull distribution of the former subsection seems to improve the fit of the models,

EXHIBIT 13
Density of Cox-Snell Residuals, WE-MEM(1,2), $R^{CHF/USD}$

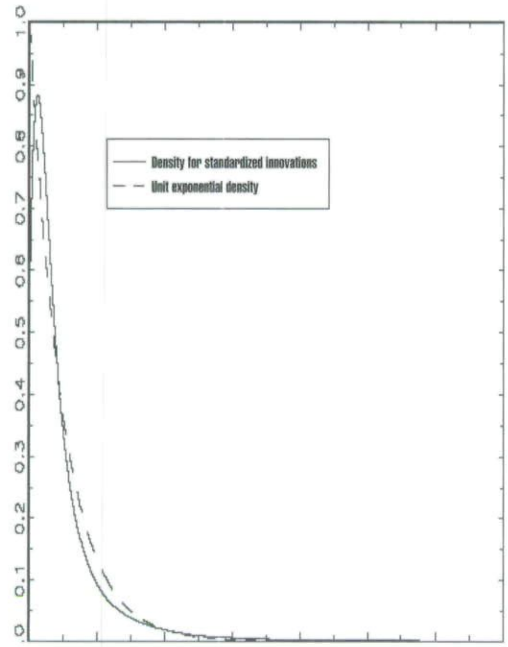

EXHIBIT 14
Density of Cox-Snell Residuals, WE-FIMEM(1,d,1), $R^{USD/GBP}$

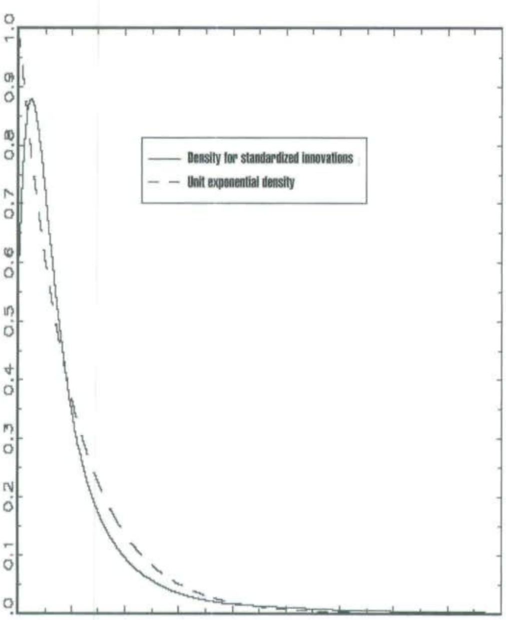

although they are still mis-specified. We therefore adopt the Gamma distribution—see Equation (14).[16]

Computation of Cox-Snell residuals for this model is rather 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

$$\hat{u}_s = \int_0^{\hat{\varepsilon}_s} \hat{h}_\varepsilon(\hat{\varepsilon}_s) = \int_0^{\hat{\varepsilon}_s} \frac{\frac{(\hat{\varepsilon}_s)^{\hat{\alpha}-1}}{\Gamma(\hat{\alpha})}\exp(-\hat{\varepsilon}_s)}{1 - \int_0^{\hat{\varepsilon}_s} \frac{(\hat{\varepsilon}_s)^{\hat{\alpha}-1}}{\Gamma(\hat{\alpha})}\exp(-\hat{\varepsilon}_s)},$$

$$\hat{\varepsilon}_s = \frac{\hat{\alpha} R_s}{\hat{\lambda}_s} \tag{16}$$

where the integral is computed numerically.[17]

The coefficient estimates and the associated standard errors (results not reported) are exactly the same as those obtained with the Exponential distribution in the QMLE procedure.[18] This is not surprising given the results in Engle and Gallo [2006]. 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 Exhibit 15). Exhibits 16 and 17 show the densities of Cox-Snell residuals from the estimated Gamma (FI)MEMs. The densities seem to fit the unit Exponential rather well. In contrast, Exhibit 15, middle section, shows the mean, standard deviation, skewness, and kurtosis of the Cox-Snell residuals from the Gamma (FI)MEMs. It is evident that they do not seem to follow a unit Exponential distribution.[19] Moreover, the excess dispersion test of Engle and 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 Exhibit 15, we compute the same sample moments, excluding the 20 (for $R^{CHF/USD}$) 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 fails to reject the null of the standard deviation being equal to unity. The general message from Exhibits 15, 16, and 17 is that the source of the mis-specification is a relatively small number of extreme observations.

EXHIBIT 15

(FI)MEMs based on the Gamma Distribution

"Mean-th.mean" refers to the difference between the empirical mean of Cox-Snell residuals and the theoretical value of unity. "Std-th.std" refers to the difference between the empirical standard deviation of Cox-Snell residuals and the theoretical value of unity. The middle panel shows summary statistics for Cox-Snell residuals from the GA-(FI)MEM 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.

	$R^{CHF/USD}$		$R^{USD/GBP}$	
	GA-MEM(1,2)	GA-FIMEM(1,d,1)	GA-MEM(1,2)	GA-FIMEM(1,d,1)
Mean-th.mean	0.01584	0.01605	0.01949	0.01562
Std-th.std	0.27559	0.27436	0.28052	0.27211
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

EXHIBIT 16

Density of Cox-Snell Residuals, GA-MEM(1,2), $R^{CHF/USD}$

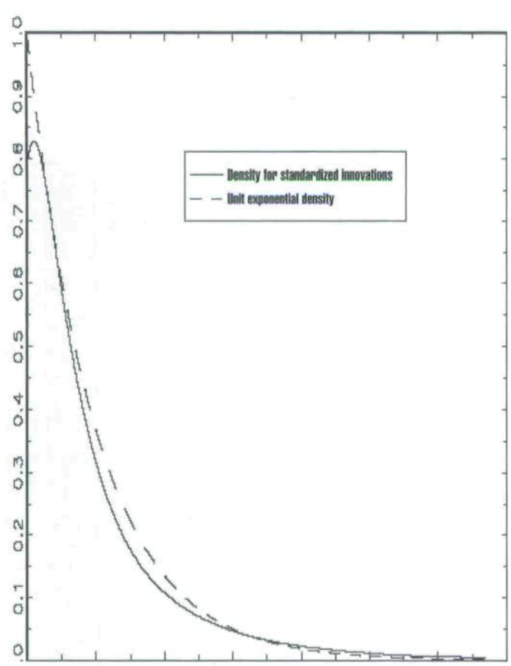

EXHIBIT 17

Density of Cox-Snell Residuals, GA-FIMEM(1,d,1), $R^{USD/GBP}$

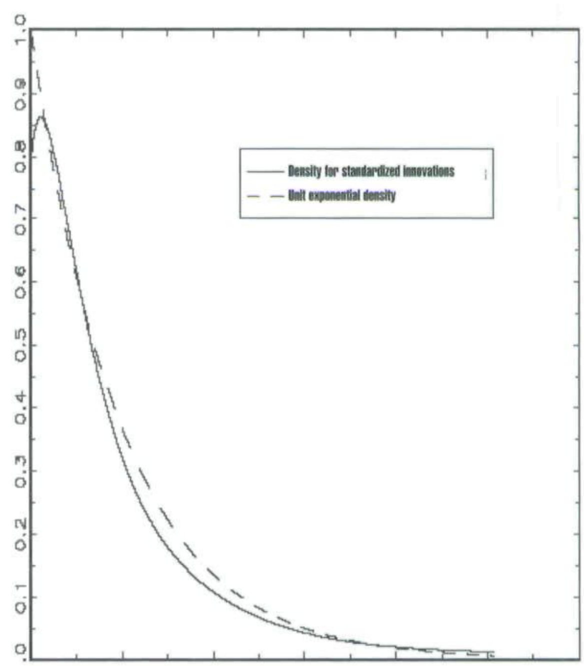

We further analyzed the largest Cox-Snell residuals for the four models. For the CHF/USD, the largest Cox-Snell residuals for the Gamma-MEM(1,2) coincide with the largest Cox-Snell residuals for the GA-FIMEM(1,d,1). Moreover, the largest Cox-Snell residuals of the two models match the largest price ranges. Press reports on the days corresponding to the largest ranges (very high volatility) reveal that unusual events took place (e.g., Russian coup, Asian crisis, Russian crisis, and so forth). The same analysis holds for the USD/GBP rate. It seems that the GA-(FI)MEM works well in standard market conditions, but is not able to account for abnormal events.

Summary of Empirical Results

By a semi-parametric procedure (QMLE), we estimated several short memory MEMs and found evidence in favor of the MEM(1,2). The results show clear evidence of a strong degree of persistence. For this reason, we extend the model to allow for long memory. The FIMEM(1,d,1) is able to capture the dynamic properties of the USD/GBP range, but the MEM(1,2) is the preferred model for the CHF/USD range.

For the parametric estimates of the (FI)MEMs, we use three distributions: the Exponential, the Weibull and the ordinary Gamma. We find evidence that the Gamma distribution provides a fairly good approximation of the distribution of the standardized innovations, except in the very far end of the tail. However, the results in Engle and Gallo [2006] suggest use of the simple Exponential distribution which has the additional advantage of a QMLE interpretation.

FORECASTING EXERCISES

Forecasting the volatility process of asset returns is not an easy task. In fact, the volatility process cannot be directly observed. There are, however, several volatility proxies such as realized volatility (see Andersen, Bollerslev, Diebold, and Labys [2000]). A possible strategy to assess the forecasting performance of the selected (FI)MEMs is to consider these volatility proxies as the "true" volatility. Hansen and Lunde [2006] show that to overcome the bias of market microstructure noise, realized volatility measures should be constructed using all available transactions or all available quotes. A problem with FX data is that transactions are OTC and, therefore, not available. Data on intraday quotes on FX markets are available to us at designated time intervals (i.e., we do not have data on all quotes). It is thus questionable whether the realized volatility measures that we are able to construct represent the "true" variance. Consequently, we decided to adopt different approaches.

Forecasting the Range: One-step-ahead Forecasts

The (FI)MEM estimates cover the period from January 3, 1991 to September 20, 2001. Using the selected models, MEM(1,2) for the CHF/USD rate and FIMEM (1,d,1) for the USD/GBP rate, we forecast one-day ranges for the period from September 21, 2001 to May 31, 2002 for a total of 176 trading days. To assess the forecasting performance of the (FI)MEMs we then run the following regression

$$R_s^i = \theta_0 + \theta_1 \tilde{R}_s^i + \upsilon_s \qquad (17)$$

where $\tilde{R}_s^i$ is the one-step-ahead forecast of the range, R_s^i is the actual range, and $i = CHF/USD, USD/GBP$.

If the (FI)MEMs are able to forecast the range well, we should expect $\theta_0 = 0$ and $\theta_1 = 1$. Exhibit 18 shows that when considered separately, we fail to reject the two hypotheses at standard significance levels (5%). The test for the joint linear restrictions, $\theta_0 = 0$ and $\theta_1 = 1$, fails to reject for the USD/GBP range only. The last row of Exhibit 18 reports r^2 for Equation (17). Results are in line with previous forecasts of the daily range (see Chou [2005]). Overall, Exhibit 18 suggests that (FI)MEMs are able to forecast the daily range rather well.

Trading Straddles and Strangles

Assume that a trader reads the above results on the (FI)MEMs and would like to use these models for trading. Then the issue becomes: Can the trader make money?

Let us call this trader Ihmm (I hope to make money). Ihmm buys and sells volatility using straddle and strangle strategies.[20] A straddle strategy consists of buying a call and a put (at the money) with the same maturity and strike price on the same underlying asset. The buyer of the straddle strategy makes a profit if the price of the underlying asset moves, regardless of which direction. But if the price of the underlying asset at maturity is equal to the strike price, the owner of the straddle will face a loss

EXHIBIT 18
One-step-ahead Forecasts

	CHF/USD - MEM(1,2)	USD/GBP FIMEM(1,d,1)
$\theta_0 = 0$ (p-values)	13.6%	53.7%
$\theta_1 = 1$ (p-values)	7.08%	34.8%
$\theta_0 = 0$ & $\theta_1 = 1$ (p-values)	1.28%	9.29%
r^2	14.7%	13.8%

equal to the options' premia. The strangle strategy is similar to the straddle, but the call has a higher strike price than the put (see Hull [2003]).

Ihmm operates as follows: On day s, Ihmm forecasts the volatility until maturity τ using the preferred (FI)MEM. Then, using the simple Black-Scholes option pricing model, Ihmm computes the price of the straddle (strangle) strategy. On $s + 1$, Ihmm compares the price of the straddle (strangle) obtained at s using the (FI)MEM volatility forecasts to the price quoted in the market. If the market price is higher than the forecasted price, Ihmm sells the straddle (strangle). But if the market price is lower than the forecasted price, Ihmm buys the straddle (strangle). When the forecasted price and the market price are the same, Ihmm does not trade. In other words, Ihmm is comparing the forecasted volatility from the (FI)MEM to implied volatility. The general wisdom about implied volatility is that it represents the market expectations of the volatility of the underlying asset until maturity, and it is the volatility measure most used in the finance industry.

Ihmm's forecasts. The call and put prices are derived by the Black-Scholes formula adjusted for foreign currencies. The call and put price for next day, $s + 1$, conditional on the information set available in s, can be written as

$$\tilde{c}^{\tau}_{s+1}\big|\Lambda_s = S^{A/B}_s \exp(-\rho_s\tau)N(d_1) - X^{A/B}\exp(-r_s\tau)N(d_2) \tag{18}$$

$$\tilde{p}_{s+1}\big|\Lambda_s = X^{A/B}\exp(-r_s\tau)N(d_2) - S^{A/B}_s \exp(-\rho_s\tau)N(d_1) \tag{19}$$

$$d_1 = \frac{\ln\left(S^{A/B}_s / X^{A/B}\right) + \left[r_s - \rho_s + (1/2)(\sigma^{\tau}_{s+1}\big|\Lambda_s)^2\right]}{(\sigma^{\tau}_{s+1} \mid \Lambda_s)\sqrt{\tau}} \tag{20}$$

$$d_2 = \frac{\ln\left(S^{A/B}_s / X^{A/B}\right) + \left[r_s - \rho_s - (1/2)(\sigma^{\tau}_{s+1}\big|\Lambda_s)^2\right]}{(\sigma^{\tau}_{s+1} \mid \Lambda_s)\sqrt{\tau}} \tag{21}$$

where τ is the maturity of the option, $S^{A/B}_s$ is the spot price in terms of currency A of one unit of currency B in day s, $X^{A/B}$ is the strike price, r_s and ρ_s are the annualized risk free rates for maturity τ for the domestic and the foreign currency, respectively, and the volatility is given by

$$\sigma^{\tau}_{s+1}\big|\Lambda_s = \sqrt{\frac{1}{\tau}\sum\nolimits_{h=2}^{\tau} \tilde{R}^2_{s+h}} \tag{22}$$

Therefore, $\sigma^{\tau}_{s+1} \mid \Lambda_s$ is the forecast made at the origin s of the average volatility from $s + 1$ until the option's maturity date, τ.

Volatility forecasts are made using the MEM(1,2) for the CHF/USD rate and the FIMEM(1,d,1) for the USD/GBP rate. The forecasting period covers September 21, 2001 to May 31, 2002. The starting values of the parameters for each forecast are those obtained in Exhibit 7, column 3, for the CHF/USD rate, and Exhibit 10, column 4, for the USD/GBP rate. These refer to the period January 3, 1991 to September 20, 2001. For each forecast, convergence is achieved with few (less than ten) iterations. Forecasts are made using all the available information until time s, the day before the start of the forecasting period. The estimated parameters and relative standard errors are consistently stable across forecasts and in line with results reported in the previous section.

The data for the options are taken from the Philadelphia Stock Exchange (PHLX). We concentrate on European options and analyze transaction data. Trading hours for the currency options are from 2:30 AM EST to 2:30 PM EST. The idea is to allow Ihmm to trade straddles or strangles at transaction prices using the first available transactions in a trading day.[21]

To trade a straddle Ihmm needs both a put and a call on the same underlying asset, with the same maturity and strike. For a strangle, the strikes are different and, in particular, the call's strike is higher than the put's strike. Unfortunately, our data show that for the 176 days in the forecasting period, only on 63 days Ihmm could trade straddles or strangles in the CHF/USD and USD/GBP rates (FX transactions are mainly OTC). However, on the 63 days that Ihmm could trade, our data show that trades in puts and calls take place in pairs with exactly the same volume. This might suggest that traders effectively use straddles and strangles to buy and sell volatility.

From January 2002, the volume of traded options on the CHF/USD FX dropped dramatically. We conjecture that this might be due to the inception of the Euro which probably attracted most of the trading in options on European currencies. For the whole forecasting period analyzed, the volume of transactions on the USD/GBP was never particularly high. In total, there are 50 days for the CHF/USD and only 13 days for the GBP/USD on which it is possible to trade straddles and/or strangles.[22]

The data for the risk-free interest rates are taken from Datastream and have maturities of one day; one and two weeks; and one, two, three, six, nine, and twelve months.[23]

We are aware that our approach has many limitations. First, we are using the Black-Scholes option pricing formula which assumes constant volatility. More sophisticated pricing formulas with stochastic volatility are available (e.g., Hull and White [1987]). However, the use of European options might make Black-Scholes sufficiently accurate. Moreover, the goal of this section is to perform a very simple forecasting exercise which makes use of standard tools such as Black-Scholes. Second, there are only 63 days on which Ihmm can trade. This may not provide sufficient evidence of the forecasting performance of (FI)MEMs. The analysis of stock market data will provide more dates where straddle and strangles can be traded. We leave this to future research.

Computing the rate of returns. Each day, $s+1$, Ihmm goes to the market to trade straddles and strangles using the price forecasted on s. Ihmm does not have any capital. In fact, when Ihmm buys the straddle (strangle), he borrows the money at the risk-free rate. When Ihmm sells the straddle (strangle), he invests the proceeds at the risk-free rate. The minimum contract sizes for options quoted on the PHLX is 62,500 and 31,250 for the CHF/USD rate and the GBP/USD rate, respectively. The initial amount of money, on $s+1$, involved in the transaction is given by

$$y_{s+1} = \Theta_i(c_{s+1} + p_{s+1}) \tag{23}$$

where c_{s+1} and p_{s+1} are the call and put market prices on day $s+1$, while $\Theta_i = 2 \times 62{,}500$ when $i = CHF/USD$, and $\Theta_i = 2 \times 31{,}250$ when $i = GBP/USD$. The variable y_{s+1} is positive when selling the straddle (strangle) and negative when buying the straddle (strangle).

Noh et al. [1994] computed returns by looking at the difference between the straddle price on $s+1$ and $s+2$. We cannot adopt the same approach because the options markets we are analyzing are not very liquid; in other words, we are not sure whether we will be able to observe a market price for the straddle (strangle) the following day. Therefore, we use a different strategy. Once Ihmm opens a position, he keeps the position until the options' maturity. At maturity, if Ihmm bought the straddle (strangle), he compares the strike price of the options to the market (close) price of the underlying asset and determines whether he should exercise the options or not. If Ihmm sold the straddle (strangle), the buyer will decide whether to exercise. On $s+\tau$ (maturity), profits for buying the straddle are computed as follows:

$$\pi_{s+\tau} = \begin{cases} -\Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad +\Theta_i\left(S_{s+\tau}^{A/B} - X^{A/B}\right) \text{ if the call is exercized} \\ -\Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad +\Theta_i\left(X^{A/B} - S_{s+\tau}^{A/B}\right) \text{ if the put is exercized} \\ -\Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad \text{if neither the put nor the call are exercized} \end{cases}$$

while those for selling the straddle (strangle) are

$$\pi_{s+\tau} = \begin{cases} \Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad -\Theta_i\left(S_{s+\tau}^{A/B} - X^{A/B}\right) \text{ if the call is exercized} \\ \Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad -\Theta_i\left(X^{A/B} - S_{s+\tau}^{A/B}\right) \text{ if the put is exercized} \\ \Theta_i(c_{s+1} + p_{s+1})\exp[r_{s+1}\tau] \\ \quad \text{if neither the put nor the call are exercized} \end{cases}$$

Returns at maturity are given by the ratio $\pi_{s+\tau}/\gamma_{s+1}$. We also allow for transaction costs which we set equal to \$0.05 per option. Therefore, the transaction costs for buying a straddle on the CHF/USD rate are ($2 \times 62{,}500 \times$ \$0.05). The introduction of transaction costs will affect the profits from trading straddles and strangles. However, it does not have any implication for the forecasting ability of (FI)MEMs. Throughout this exercise we adopt the usual convention of 252 trading days per year.

Ihmm's profits. As previously mentioned, there are 63 days when Ihmm can trade: 41 times he sells the straddle (strangle), 19 times he buys the straddle (strangle), and 3 times he does not trade because the forecasted straddle (strangle) price is exactly equal to the market price. Therefore, 65% of the time Ihmm is selling which indicates that the straddle (strangle) market price is higher than the forecasted price. In turn, this implies that (FI)MEM volatility forecasts are lower than implied volatilities. Maturities range from one day to a maximum of 60 days, with a median of 21 days which is a trading month.

Exhibit 19 shows the annualized rate of return of trading straddles and strangles.

For the CHF/USD rate, there are 49 transactions which generate an average rate of return equal to 31.2% with a standard deviation of 65.9%. Ihmm makes profits 71% of the time. When introducing transaction costs, the average rate of returns drops to 27.5% and Ihmm makes a profit in 65% of his trades. For the GBP/USD rate, the average rate of return is very high and very volatile. The t-ratio[24] indicates that the average rate of return is only significant at the 20% significance level. When trading straddles (strangles) on the GBP/USD rate, Ihmm makes money 64% of the time. As expected, when introducing transaction costs, the average rate of return decreases to 60%. Given the small number of data points available, it is perhaps more appropriate to analyze the aggregate results. Overall, Ihmm earns an annual rate of 40.1% and his trades are profitable 70% of the time. Transaction costs reduce the average rate of return to 33.6%. The t-ratio indicates that these profits are significantly greater than zero.

Selling a straddle (strangle) bears a much higher risk than buying it. In the former case the potential loss is unbounded, while in the latter case it is bounded by the options' premia. A riskier trading strategy requires higher returns. To check whether our results are affected by the fact that Ihmm sells more than he buys, we compute the profits from a simple strategy which consists of selling the straddle (strangle) at all times.[25] The last row of the first part of Exhibit 19 shows that this strategy is less profitable than Ihmm's strategy, and therefore confirms that Ihmm's profits are robust against this critique. When Ihmm buys the straddle (strangle) he makes an average profit of 10.9% and, as expected, when he sells the straddle (strangle) his profits average 53.6%.

GARCH(1,1) forecasts. To corroborate our results, we conduct the same forecasting exercise using the famous and widely used GARCH(1,1) model. Hansen and Lunde [2005] analyze the forecasting performances of 330 ARCH-types models and find that the simple GARCH(1,1) model outperforms more sophisticated models in forecasting exchange rate volatility. We,

EXHIBIT 19
Ihmm's Average Annualized Profits

		Range Forecasts							
		No Transaction Costs				Transaction Costs			
	# Obs.	Mean	St. Dev.	t-ratio	% π > 0	Mean	St. Dev.	t-ratio	% π > 0
CHF/USD	49	31.2%	65.9%	3.61	71%	27.5%	61.4%	3.42	65%
GBP/USD	11	79.8%	176%	1.50	64%	60.3%	137.2%	1.46	64%
Total	60	40.1%	93.8%	3.31	70%	33.6%	79.3%	3.28	70%
Always sell	60	33.2%	96.5%	2.67	68%	30.2%	79.9%	3.13	68%
		GARCH(1,1) Forecasts							
		No Transaction Costs				Transaction Costs			
	# Obs.	Mean	St. Dev.	t-ratio	% π > 0	Mean	St. Dev.	t-ratio	% π > 0
CHF/USD	49	-19.5%	72.9%	-2.03	44.9%	-20.8%	67.9%	-2.33	42.9%
GBP/USD	11	61.6%	184%	1.11	54.5%	42.1%	149%	0.96	54.5%
Total	60	-4.60%	102%	-0.35	45.0%	-9.26%	86.9%	-0.82	43.3%

therefore, compare the forecasting performance of (FI)MEMs against that of the GARCH(1,1). The set-up is exactly the same as the preceding example. The bottom part of Exhibit 19 reports profits from trading straddles and strangles using the GARCH(1,1) forecasts. The results show a net profit for the GBP/USD rate, but a loss for the CHF/USD rate. The aggregate result is an annualized loss of 4.6% without transaction costs, and of 9.3% with transaction costs. However, in both cases, the t-ratio indicates that these values are not significantly different from zero. These results are very different from those obtained using the (FI)MEMs.

Our findings should be interpreted with caution given that straddles and strangles trading strategies could be applied only to a few days in the forecasting sample. (FI)MEMs positive annualized returns are probably exaggerated by lack of observations. In the same way, the poor performance of the GARCH(1,1) volatility forecasts might be due to the few data points used in the forecasting exercise. We are, nevertheless, confident that volatility forecasts made via (FI)MEMs might have significant economic value. Even if our forecasting exercise is limited to 63 data points, there is evidence that the range, as a volatility proxy, may contain valuable information, and that the (FI)MEM may provide a useful tool for modeling and forecasting exchange rate volatility.

CONCLUSION

The aim of the article 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 Engle [2002] and Chou [2005], we model daily price ranges for two major exchange rates, CHF/USD and USD/GBP, by the MEM.

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. However, we also find that the CHF/USD price range is best described by a short memory MEM(1,2). In a simple forecasting exercise, we show that (FI)MEMs are able to predict daily ranges rather well. Moreover, using straddles and strangles trading strategies, we show that volatility forecasts from (FI)MEMs may have significant economic value.

ENDNOTES

We thank Niels Haldrup, Alessio Caldarera, Stephen Figlewski, the referee, and seminar participants at JHU for useful comments. The usual disclaimer applies.

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

[2]Note that Engle and Russell [1998] apply the model to intertrade durations and, therefore, name the model Autoregressive Conditional Duration (ACD). Chou [2005] uses the same specification to model the range and, therefore, names the model Conditional AutoRegressive Range (CARR).

[3]See Parkinson [1980].

[4]LB(n) statistics follow a chi-squared distribution with n degrees of freedom under the null of no serial correlation from lag 1 to n. Critical values for the LB tests in Exhibit 4 (at a significance level of 5%) are {3.84,25.00,43.77,67.50,124.34, 287.88}.

[5]See Baillie et al. [1996].

[6]To conserve space, we do not report the results of the MEM estimates without the September 1992 dummy.

[7]Although the FIMEM is not weakly stationary with an unbounded, unconditional mean when $\omega > 0$, the main goal of using the FIMEM is to model the conditional range. In fact, we are interested in the dynamic properties of the range.

[8]We also investigate the FIMEM(1,d,0) and find strong evidence in favor of the FIMEM(1,d,1).

[9]The parameter $\hat{\phi}_1$ is not significantly different from $\hat{a}_1 + \hat{b}_1$ for the MEM(1,2), $\hat{\phi}_2$ is not significantly different from $\hat{a}_2 + \hat{b}_2$, and $\hat{b}_1$ (for FIMEM(1,d,2)) is not significantly different from $\hat{b}_1$ (for MEM(1,2)).

[10]Strictly speaking, the FIMEM(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.

[11]We use a bandwith of 0.15.

[12]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.

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

[14]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, and Peacock [1993, p.61]).

[15]Note that the standard errors are computed from the hessian (i.e., we did not use Lee and Hansen [1994]). To conserve space we do not report the estimation results for the WE-(FI)MEM.

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

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

[18]We also computed standard errors from the hessian and found that they were, in general, smaller than the robust standard errors of the QMLE procedure.

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

[20]This approach is similar to that developed in Noh, Engle, Kane [1994].

[21]Alternatively, we could have used close prices. However, data on close prices were not available from the Philadelphia Stock Exchange. The idea is to allow Ihmm to trade straddles and strangles taking advantage of the first available opportunity. The transaction data also show, see below, that puts and calls on the same underlying, maturity, and strike (only for straddles) were traded almost at the same time.

For a few days, straddles and strangles strategies could have been implemented twice or three times in a day. In this case, hoping to consolidate our results, we opt for the conservative choice and trade the straddle (strangle) that produces either the lowest profits or the highest losses.

[22]We also analyzed option data from June 2002 to March 2003, and we could not find a single day where straddles and/or strangles could be traded.

[23]We use linear interpolation for intermediate maturities.

[24]The t-ratio is based on the assumption that rates of return from straddles and strangles are independent and normally distributed. This is not the case with our data. Therefore, caution should be exercised in interpreting these results.

[25]We thank Jeffrey Russell for this suggestion.

REFERENCES

Andersen, T.G., and T. Bollerslev. "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts." *International Economic Review*, 39 (1998), pp. 885–905.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys, "Great Realizations." *Risk*, 13 (2000), pp. 105–108.

——. "The Distribution of Exchange Rate Volatility." *Journal of the American Statistical Association*, 96 (2001), pp. 42–55.

——. "Modeling and Forecasting Realized Volatility." *Econometrica*, 71 (2003), pp. 529–626.

Baillie, R.T. "Long Memory Processes and Fractional Integration in Econometrics." *Journal of Econometrics*, 73 (1996), pp. 5–59.

Baillie, R.T., T. Bollerslev, and H. Mikkelsen. "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics*, 74 (1996), pp. 3–30.

Barndorff-Nielsen, O., and N. Shephard. "Realized Power Variation and Stochastic Volatility Models." *Bernoulli*, 9 (2003), pp. 243–265.

Beran, J., and D. Ocker. "Volatility of Stock-Market Indexes—An Analysis Based on SEMIFAR Models." *Journal of Business and Economic Statistics*, 19 (2001), pp. 103–116.

Bierens, H. "Higher-Order Sample Autocorrelations and the Unit Root Hypothesis." *Journal of Econometrics*, 57 (1993), pp. 137–160.

Brandt, M., and F.X. Diebold. "A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations." *Journal of Business*, 79 (2006), pp. 61–74.

Brunetti, C., and P.M. Lildholdt. "Relative Efficiency of Return—and Range-Based Volatility Estimators." Manuscript, (2004).

Chen, S.X. "Probability Density Function Estimation Using Gamma Kernels." *Annals of the Institute Statistical Mathematics*, 52 (2000), pp. 471–480.

Chou, R. "Forecasting Financial Volatilities with Extreme Values: The Conditional Autoregressive Range (CARR) Model." *Journal of Money, Credit, and Banking*, 37 (2005), pp. 561–582.

Christensen, K., and M. Podolski. "Asymptotic Theory of Range-Based Estimation of Integrated Variance of a Continuous Semi-Martingale." Manuscript, (2005).

Cox, D., and E. Snell. "A General Definition of Residuals." *Journal of the Royal Statistical Society*, Series B, 30 (1968), pp. 248–275.

Engle, R. "New Frontiers for ARCH Models." *Journal of Applied Econometrics*, 17 (2002), pp. 425–446.

Engle, R., and G. Gallo. "A Multiple-Indicators Model for Volatility Using Intra-Daily Data." *Journal of Econometrics*, 131 (2006), pp. 3–27.

Engle, R., and J. Russell. "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data." *Econometrica*, 66 (1998), pp. 1127–1162.

Evans, M., N. Hastings, and B. Peacock. *Statistical Distributions*. New York, NY: John Wiley and Sons, 1993.

Feller, W. "The Asymptotic Distribution of the Range of Sums of Independent Random Variables." *Annals of Mathematical Statistics*, 22 (1951), pp. 427–432.

Gallant, A., C. Hsu, and G. Tauchen. "Using Daily Range Data to Calibrate Volatility Diffusions And Extract the Forward Integrated Variance." *The Review of Economic, and Statistics*, 81 (1999), pp. 617–631.

Hansen, P., and A. Lunde. "A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)." *Journal of Applied Econometrics*, 20 (2005), pp. 873–889.

——. "Realized Variance and Market Microstructure Noise." *Journal of Business and Economic Statistics*, 24 (2006), pp. 127–218.

Hasza, D. "The Asymptotic Distribution of the Sample Autocorrelations for the Integrated ARMA Process." *Journal of the American Statistical Association*, 75 (1980), pp. 349–352.

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

Hull, J.C. *Options, Futures and Other Derivatives*. Prentice Hall, 2003.

Hull, J.C., and A. White. "The Pricing of Options on Assets With Stochastic Volatility." *Journal of Finance*, 42 (1987), pp. 281–300.

Jasiak, J. "Persistence in Intertrade Durations." Manuscript, (1999).

Lee, S., and B. Hansen. "Asymptotic Theory for the GARCH(1,1) Quasi Maximum Likelihood Estimator." *Econometric Theory*, 10 (1994), pp. 29–52.

Li, K., and D. Weinbaum. "The Empirical Performance of Alternative Extreme Value Volatility Estimators." Manuscript, (2000).

Lunde, A. "A Generalized Gamma Autoregressive Conditional Duration Model." Manuscript, (1999).

Nelson, D. "Stationary and Persistence in the GARCH(1,1) Model." *Econometric Theory*, 6 (1990), pp. 318–334.

Nelson, D., and Cao, C. "Inequality Constraints in the Univariate GARCH Model." *Journal of Business and Economic Statistics*, 10 (1992), pp. 229–235.

Noh, J., R.F., Engle, and A. Kane. "Forecasting Volatility and Option Prices of the S&P 500 Index." *Journal of Derivatives*, 1(1994) pp.17–30.

Parkinson, M. "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business*, 53 (1980), pp. 61–65.