MODELING FOREIGN EXCHANGE VOLATILITY WITH INTRADAY DATA

by

Alexandre Borges Sugiyama

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ECONOMICS

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 9 8

MODELING FOREIGN EXCHANGE VOLATILITY

WITH INTRADAY DATA

by

Alexandre Borges Sugiyama

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF ECONOMICS

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 9 8

UMI Number: 9829336

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

THE UNIVERSITY OF ARIZONA ®
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read the dissertation prepared by Alexandre Borges Sugiyama entitled Modeling Foreign Exchange Volatility with Intraday Data

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Satheesh Aradhyula — Date 4/6/98

William Horrace — Date 4/6/98

Ronald Oaxaca — Date 4/6/98

Stanley Reynolds — Date 4/6/98

Date

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copy of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Ronald L. Oaxaca
Dissertation Director Ronald Oaxaca — Date 4/6/98

Statement by Author

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgement of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: Alexandre Borges Fagundes

Acknowledgements

The author thanks his dissertation committe of Satheesh Aradhyula, William Horrace, Ronald Oaxaca, and Stan Reynolds. Also appreciated are: Devajyoti Ghose, Kenneth Kroner, Nori Sasaki, Lester Taylor, and the participants at the Econometrics Workshop at the University of Arizona for their comments and suggestions. The author thanks Jonathan Kent and Richard Kiser for their computer assistance. Finally, the support of Maria Teresa Vélez and the Graduate College is greatly appreciated.

Table of Contents

Table of Contents—Continued

Table of Contents—Continued

List of Figures

List of Figures—Continued

List of Figures—Continued

List of Figures—Continued

List of Tables

List of Tables—Continued

List of Tables—Continued

List of Tables—Continued

Abstract

This dissertation studies intraday and daily foreign exchange market volatility. First, we address how best to model the intraday seasonality and the serial correlation in return volatility. We find there is no gain from smoothing the intraday seasonal volatility pattern. A model that jointly estimates the intraday seasonal pattern and conditional heteroskedasticity underperforms models that remove seasonal variance through deseasonalization and then model conditional heteroskedasticity with a GARCH model. Secondly, we show how intraday data can be used to create daily volatility estimates. Results show intraday data allow for daily volatility estimates which are independent of a volatility dynamics specification. Lastly, we show that intraday data improve the performance of one-step ahead forecasts based on a one year sample and show that the results are consistent with Monte Carlo simulations.

Chapter 1
Conditional Heteroskedasticity and Intraday Seasonal Volatility in the Foreign Exchange Market

1.1 Introduction

There is a growing area of research in finance which deals with modeling volatility of asset markets. With the construction of high frequency data. it is now possible to model volatility in a more precise manner. Previously. one could only address volatility on a week by week or day by day time frame. While this may be fine for some applications. like options pricing. it may not be for other applications since information about the nature of the market is unused.

New research is now possible due to the collection of data which is of a higher frequency than previously used. Whereas previous studies used weekly or daily data. data sets with prices every few minutes are now available which allow researchers to look at new market dynamics. For example. in high frequency exchange rate data. it has been observed that price changes vary a great deal per day. That is. depending on the time of day. the volatility of the exchange rate market fluctuates. This pattern is completely unobservable when working with daily data. The effect is very similar when working with sales data for firms. There is a big end of the year effect on sales. and that is known as the Christmas/Holiday shopping season. In a regression analysis we need to take holiday sales into account. so too do we have to adjust our high frequency exchange rate models for a periodic intraday pattern. In this case. we have a daily pattern in the variance of the exchange rates.

One of the recent advances in modeling financial time series has been the work of Engle (1982) and Bollerslev (1986) for the ARCH model. This class of models estimates parameters more efficiently by directly modeling the second moment of the model (or variance, or volatility), in addition to the first moment.[1] In financial time series data volatility appears to be time varying; there are periods of high and low volatility. Since financial models often make assumptions on the nature of variance it is important that one uses a class of models which allow for time-varying variance. Assuming that the data has constant variance when it does not could lead one to make incorrect statistical inference. The ARCH framework is appealing for the direct nature of its modeling of time-varying variances.

This chapter tries to address significance of intraday volatility seasonality when we are modeling exchange rates with an ARCH framework. We address whether intraday seasonality matters, and if it does how best to handle it. To show which is the best way to model this seasonality, we estimate a variety of models and evaluate their effectiveness in forecasting out-of-sample.

1.1.1 Literature Review

The approach of this paper is not completely ad hoc. For one, it has been established that high frequency returns do have an intraday seasonal variance pattern.[2] For this reason, it is important to model this seasonality. In addition, there is a well established literature which looks at volatility transmission.[3] This tries to explain why there are periods of high and low volatility in financial markets. A popular

[1] A good survey article is Bollerslev, Chou, and Kroner (1992).
[2] See Dacorogna et. al. (1993).
[3] See Engle, Ito, and Lin (1988).

explanation is that there are periods of high and low 'news' activity. As news arrives to a market, agents will have differing interpretations of the news and it takes time for a new equilibrium to result. The news process is the cause of changes in prices and price expectations and ARCH models are often used to model the arrival of news to the market. Therefore, it seems appropriate to decompose volatility into two distinct components, an intraday seasonal component and a news arrival component.

Secondly, but not of less importance, is that current research is yielding ambiguous results. It has been shown by Drost and Nijman (1993) that there is a theoretical relationship between parameter estimates from a GARCH model as the sampling rate for data becomes less frequent. That is, the parameter estimates using daily returns will be related to the parameter estimates when two-day returns are used. Andersen and Bollerslev (1994) work with five minute Deutschemark exchange rate data and show that this theoretical relationship does not always hold.

When GARCH models are estimated with returns calculated on small intervals (minutes), Andersen and Bollerslev (1994) find the GARCH parameter estimates do not follow the Drost and Nijman (1993) prediction. Also, the parameter estimates yield unreasonable volatility persistence. (Volatility persistence is a measure of how long a shock in the variance takes to dissipate. It can be measured as a half-life or as a median half-life. In either case, given the GARCH parameters, there is a theoretical relationship which yields these half-lives.) Interestingly, when longer time periods (daily and above) are used to calculate returns, the Drost and Nijman (1993) theoretical predictions hold and volatility persistence estimates are more reasonable. It is also worth noting that Andersen and Bollerslev (1994) attempt to correct for intraday seasonality, but still obtain results which they call "unpredictable, occasionally imprecise, and generally difficult to interpret."

1.1.2 Effects of Deseasonalization

The most common way to address the intraday seasonality in foreign exchange returns is to remove it by deseasonalization. This is the approach taken in Andersen and Bollerslev (1994) and more recently by Ghose and Kroner (1997). It is not clear that this approach is the best way to model volatility dynamics. When we deseasonalize intraday returns, we are assuming that what is important is to compare a return with the average return for that time period. It is a relative comparison that is important; all indications of absolute magnitude are lost.

To illustrate this point, consider the days Friday, February 11, 1994 and Monday, December 6, 1993. In both cases, if one divides the fifteen minute returns of those days, by the corresponding sample standard deviation for each return, the difference between the raw return and the deseasonalized returns varies greatly. When the intradaily returns for December 6, 1994 are deseasoned and squared, they appear to be very similar to the raw squared returns (see Figures A.1 and A.2). In fact, the correlation of squared return and the deseasonalized square return is 0.92. For February 11, 1994, deseasonalization causes the returns to appear to be different. When deseaonalized, the early morning returns and late night returns are much larger relative to the afternoon returns (see Figures A.3 and A.4). The correlation of the squared returns and the deseasoned returns is 0.45.

The days just discussed were chosen as extreme examples to illustrate the point that days will look different depending on whether one deseasonalizes or not. The question then becomes whether, if one is interested in forecasting, is it best to (1) deseasonalize returns, model volatility, create forecasts, seasonalize the forecasts or to (2) model seasonality and volatility jointly? This paper investigates whether this sequential modeling approach performs better than a joint modeling approach where the intraday seasonality and the conditional heteroskedasticity are dealt with simultaneously.

1.1.3 Outline of Chapter

This paper's goal is to compare various methods of modeling intraday seasonality. Some models have already been used in the context of exchange rates (Andersen and Bollerslev (1994)), while other models will be extended to allow for a intraday seasonal volatility component. Regardless, there has yet to be a systematic analysis of which models perform better. Such a study is useful to guide theoretical research in this area and further points to new undiscovered dynamics.

The comparison of various models will require some thought. Since the true variance process for returns is unobserved, it is not possible to calculate an R^2 as is often done with estimates from a standard regression. Instead, we will evaluate models statistically with likelihood based criteria and on out-of-sample forecast ability.

This paper analyzes various models and determines which method for handling the intraday seasonality and conditional heteroskedasticity works best for forecasting of volatility.

The rest of the paper is organized as follows: Section 1.2 presents various models which incorporate intraday seasonality and conditional heteroskedasticity. Section 1.3 discusses the data, Section 1.4 discusses how the models will be ranked or compared, Section 1.5 presents the empirical results, and Section 1.6 concludes and summarizes.

1.2 Models

1.2.1 Seasonality

There have been two general methods for dealing with the seasonality in intraday returns. The first approach is to use the sample variance while the second is to create a smoothed estimate. Both are discussed below; two variants of the smoothed approach are presented.

Model 1: Sample Variance

The most straight forward measure of the intraday seasonality is to calculate the sample variance for each intraday time interval. Suppose that our returns are defined to be $R_{t,i}$ where $t = 1 \dots T$ represents the day, and $i = 1 \dots n$ represents the time of day. For example, with fifteen minute data, $n = 96$ and the 12AM return for day 1 of the sample would be denoted $R_{1,1}$. Similarly, $R_{1,2}$ represents the 12:15AM return for day 1. Defining $\gamma(i)$ to be the variance of the returns at time i, we sum over all days T and obtain

$$\gamma(i) = \sum_{t=1}^{T} \left(R_{t,i} - \frac{1}{T} \sum_{t}^{T} R_{t,i} \right)^2 .$$

This approach was used in Ghose and Kroner (1997) to deseasonalize their returns. With their approach, to standardize the 9AM returns, one just divides the returns by the standard deviation of the 9AM returns.

Model 2: Fourier Series

The second method is to use parametric smoothing. One fits the time of day variance to a fourier expansion (a similar approach has been used by Andersen and Bollerslev (1996) to model intraday seasonality). In this case we have $\gamma(i)$ being a function of sines and cosines

$$\gamma_m(i) \quad = \quad \frac{a_0}{2} + \sum_{k=1}^{m} a_k \cos(ki) + \sum_{k=1}^{m-1} b_k \sin(ki)$$

The number of sine and cosines is determined by m, which is left to the discretion of the researcher. The series is defined when the weights (a_k and b_k) are estimated.

There are a couple of approaches to fit a fourier series to the intraday variance. One is to use the sample variances (estimated from Model 1) and then fit a fourier series to the ninety-six variances. The problem with this approach is that it assigns equal importance to each variance estimate. Unfortunately, the confidence intervals for the variance estimates are much larger when the US market is open than when it is closed. To allow all variances to be equally important is inefficient. We can addressed this inefficiency with appropriate weighting of the variances.

Another approach to fitting the fourier series is to use the individual returns rather than the ninety-six sample variances. In this case, given that the mean return is zero for each cross section[4], we can estimate the intraday variance by using squared returns. The fourier series solution for a_k is now extended to take

[4]Results are not reported.

care of the T observations per cross section interval, i.

$$
\begin{aligned}
a_k &= \frac{1}{48 \times T} \sum_{i=1}^{n} \sum_{t}^{T} R_{t,i}^2 \cos(ki) \\
b_k &= \frac{1}{48 \times T} \sum_{i=1}^{n} \sum_{t}^{T} R_{t,i}^2 \sin(ki)
\end{aligned}
$$

One of the benefits of fourier series is that all the terms are orthogonal to one another. This prevents problems of multicollinearity as the number of terms (i.e. m) can be increased with ease.

Model 3: Exponential Seasonality

A variation on the fourier series method is to use polynomials.

$$\gamma(i) = \exp(c_0 + c_1 i + \cdots + c_n i^n)$$

The approach here is similar to that used by Dacorogna et. al. (1993), although there are two distinct changes. First, the polynomial is continuous throughout the entire day. Dacorogna et. al. (1993) use a pieced polynomial, similar in nature to a spline polynomial regression, with restrictions on the coefficients. The second difference is that the polynomial is exponentiated to guarantee positive variances.

In addition we require that $\gamma(i = 1) = \gamma(i = 96)$. This forces $\gamma(i)$ to have the same value at 0 GMT and 24 GMT and is guaranteed by imposing the restrictions on the c_n coefficients.[5]

[5]If the time index i cycles from 0 to 1 (rather than 1 to 96), then the restriction is satisfied when $c_0 = c_0 + \cdots + c_n$. A similar linear restriction is created if the time index is from -1 to 1.

One of the problems with estimating a high order polynomial is that of multicollinearity. The time variables become more and more collinear as the order of the polynomial is increased. To see this we can write

$$i = \begin{pmatrix} 1 \\ 2 \\ \vdots \\ 96 \end{pmatrix}$$

As i cycles through 1 to 96, for example i^9 and i^8 while not identical, appear to be so in a numerical (or computational) environment. Our vector i^n becomes more correlated with i^{n-1} as n increases. A way to reduce this multicollinearity is to use a time index which cycles from -1 to 1 instead of 1 to 96. While this method is not perfect, it has the effect of making i^9 and i^8 more orthogonal to each other.

In order to estimate seasonal exponential, the same approach for the ARCH class of models is used. As a result, one can think of this seasonality as being conditionally seasonal. The model and its corresponding log likelihood is

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \epsilon_{t,i} &\sim N(0, \sigma^2_{t,i}) \\ \sigma^2_{t,i} &= \text{Exp}(c_0 + c_1 i + \cdots + c_n i^n) \\ \log(L) &= -\frac{1}{2}\log(\pi) - \sum_t^T \log(\sigma^2_{t,i}) - \sum_t^T \epsilon_t^2/\sigma^2_{t,i} \end{aligned}$$

which obtains consistent estimates for the c_k's.

1.2.2 Conditional Heteroskedasticity

The serial correlation in foreign exchange volatility has been captured by the ARCH (or GARCH) class of models which were developed by Engle (1982) and Bollerslev (1986). This class of models can be thought of as an Autoregressive Moving Average (ARMA) model in variance and are presented in brief form below:

Model 4: GARCH(1,1)

$$\begin{aligned} y_t &= \epsilon_t \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

Model 5: ARCH(10)

$$\begin{aligned} y_t &= \epsilon_t \\ \sigma_t^2 &= \omega + \alpha_0\epsilon_{t-1}^2 + \cdots + \alpha_{10}\epsilon_{t-10}^2 \end{aligned}$$

1.2.3 Sequential Modeling: Deseasonalization

The most common method to handle the intraday volatility seasonality has been to deasonalize the returns before modeling. This procedure is a sequential one, and involves three steps:

1. Deseasonalize the data with an estimate for the time specific variance. Using our notation of $\gamma(i)$ for our estimate of intraday variance, then our deseasoned returns are
$$\tilde{R}_{t,i} = \frac{R_{t,i}}{\sqrt{\gamma(i)}}.$$

In this step, one can use either the time-of-day sample variance or a smoothed variance estimate.

2. Model the deseasoned returns, $\tilde{R}_{t,i}$, with a GARCH or ARCH model, and obtain variance estimates $\tilde{\sigma}^2_{t,i}$.

3. We are not interested in a deseasoned volatility forecast, but a true volatility forecast. As a result, one needs to take the forecasts from the deseasoned analysis and reseasonalize them. This is done by multiplying the variance forecast by a the time specific sample standard deviation. Specifically,

$$\sigma^2_{t,i} = \tilde{\sigma}^2_{t,i} \times \sqrt{\gamma(i)}.$$

1.2.4 Simultaneous Modeling: Seasonality and Conditional Heteroskedasticity

An alternative approach from the one discussed above is to jointly (or simultaneously) model the seasonality and conditional heteroskedasticity. To investigate this approach, we will estimate a number of joint seasonal/heteroskedastic models.

Model 6: PGARCH(1,1) and PARCH(10)

The first model is the Periodic GARCH (P-GARCH) model (and similarly P-ARCH) from Bollerslev and Ghysels (1996). This model is a GARCH model where the variance parameters vary depending on the time of day. In our case, we will limit the seasonality to two terms for simplicity. This model has two sets of GARCH parameters which are estimated with two dummy variables. One dummy captures the effects of the US market being open, while the other dummy isolates the effects of the US market being closed.

$$\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_0 d_0 + \alpha_1 d_1)\,\epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\,\sigma^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York is open} \\ 0 & \text{New York is closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York is open} \\ 1 & \text{New York is closed} \end{cases}
\end{aligned}$$

The US market is assumed to be open from 1300 to 2300 GMT. Melvin (1995) describes this time period as the time period from which the New York financial center opens to when the San Francisco financial center closes.[6] While the currencies are traded 24 hours a day, this particular window seems to capture the major activity of American traders. This divides a day's ninety-six observations into forty returns for when the US market is open, and fifty-six observations for when the US market is closed. In Bollerslev and Ghysels (1996), the foreign exchange market returns were partitioned into two returns per day: a morning return and an afternoon return.

The PARCH model is similar to the PGARCH, with the variance equation changed from a GARCH to ARCH framework. Otherwise, there is no difference between the two models.

[6]We do not correct for daylight savings. In the summer, the US markets open one hour earlier relative to GMT. The focus here is to approximate the effect of the US market on currency volatility, not to exhaustively study it.

Model 7: Conditionally Periodic GARCH (CP-GARCH) and Conditionally Periodic ARCH (CP-ARCH)

The second seasonal/heteroskedastic model is designed as a GARCH model where the "ω" term varies with the time of day. In fact, we assume that "ω" varies proportionally to the intraday seasonal variance. This model consists of two steps. First, the seasonal intraday variance is estimated, either using the sample variance for each time of day or the exponential seasonality. Secondly, the conditional heteroskedasticity is modeled with the seasonal intraday variance estimates obtained from the first step being placed in the variance equation.

The second stage of the model is:

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i|\gamma(i)} &= \gamma(i)\omega + \alpha\epsilon^2_{t,i-1} + \beta\sigma^2_{t,i-1}. \end{aligned}$$

while in ARCH terms, we get:

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i|\gamma(i)} &= \gamma(i)\omega + \alpha_0\epsilon^2_{t,i-1} + \cdots + \alpha_n\epsilon^2_{t,i-n}. \end{aligned}$$

This model does not estimate the seasonality when modeling the volatility dynamics, but takes it as given (and hence the name conditionally periodic). The ω term is necessary to deflate the γ_t in the variance equation. In a GARCH model, the relationship between the sample variance and the parameter estimates is given by $\sigma = \frac{\omega}{1-\alpha-\beta}$ and therefore the ω term in this model allows the α and β parameters to vary freely.

Model 8: Seasonality & ARCH(10) (S-ARCH)

The last seasonal/heteroskedastic model is just the addition of the seasonal exponential model to the ARCH(10) model, and is the univariate equivalent to the model in Aradhyula and Tronstad (1994). This model jointly estimates intraday seasonality and conditional heteroskedasticity and is:

$$\begin{aligned} y_t &= \epsilon_t \\ \sigma_t^2 &= \exp(c_0 + c_1 i + \cdots + c_n i^n) + \\ &\quad + \alpha_0 \epsilon_{t-1}^2 + \cdots + \alpha_{10} \epsilon_{t-10}^2 \end{aligned}$$

In this model, the seasonal effects are independent of the dynamic effects. As a result, there are no feedback effects of the lagged terms on the seasonal terms. A GARCH specification would allow for feedback between seasonality and conditional heteroskedasticity, and such it is much more difficult to interpret the seasonal terms (see LaFrance and Burt 1983).

1.3 Data

The data consists of 262 daily observations for the Dollar-Deutschemark, Dollar-Swiss Franc, and Dollar-Yen, from September 15, 1993 to September 16, 1994. The sample consists of all weekdays with the exception of December 24 (December 25 and January 1 occurred on weekends.) Figures A.5, A.6, and A.7 show the exchange rates used in this study.

For each day, there are returns for every fifteen minutes, yielding a total of ninety-six observations per day. The returns are calculated as log returns and are multiplied by 100. The notation we will use will be

$$R_{t,i} = \log(P_{t,i}) - \log(P_{t,(i-1)})$$

where $t = 1, 2, \ldots, T$, and $i = 1, 2, \ldots, 96$. The t variable corresponds to day in the sample, while i represents the intraday cycle (there is no correction for daylight savings time). For example, the first fifteen minute return for day 3 at 12AM is calculated as $R_{3,1} = \log(P_{3,1}) - \log(P_{2,96})$, while the fifteen minute return for day 3 at 12:15AM is $R_{3,2} = \log(P_{3,2}) - \log(P_{3,1})$.

1.4 Model Selection

There are many ways of evaluating competing models. Statistically, one can choose the model with the highest mean likelihood. Alternatively, one may view the true value of a model in terms of its ability to perform out of sample. Below we discuss the two methods which will be used to rank models statistically.

1.4.1 Likelihood Ratio Tests

Where one model nests another, we can just use a likelihood ratio test to determine model selection. The formula for a likelihood ratio test is given by

$$\xi_{LR} = -2 \ln \lambda$$

where $\lambda = L_1/L_0$, where L_0 is the likelihood function for the unrestricted model, L_1 is the likelihood function for the restricted model and ξ_{LR} is distributed asymptotically as χ^2_k (where k is the number of restrictions). In the results section, we have the mean log-likelihood for the various models. Simple algebra yields the test statistic in terms of the mean log-likelihood.

$$\begin{aligned} \xi_{LR} &= -2\ln\lambda \\ &= -2\ln\frac{L_1}{L_0} \\ &= 2n(\text{Mean}\ln L_0 - \text{Mean}\ln L_1) \end{aligned}$$

where n is the number of observations.

1.4.2 Likelihood Dominance Criterion

We are also interested in comparing models which are not nested. The likelihood dominance criterion proposed by Pollack and Wales (1991) provides a method for model selection when this occurs. A quick summary is provided below.

Let hypothesis H_1 have n_1 parameters and let hypothesis H_2 have n_2 parameters where $n_1 < n_2$. Then, the likelihood dominance criterion:

- Prefers H_1 to H_2 if $L_2 - L_1 < \frac{C(n_2+1)-C(n_1+1)}{2}$.
- Indecisive between H_1 to H_2 if $\frac{C(n_2-n_1+1)-C(1)}{2} > L_2 - L_1 > \frac{C(n_2+1)-C(n_1+1)}{2}$.
- Prefers H_2 to H_1 if $L_2 - L_1 > \frac{C(n_2-n_1+1)-C(1)}{2}$.

where L_j is the likelihood for model j and $C(m)$ is the critical value of a chi-square with m degrees of freedom.

1.4.3 Heteroskedastic Mean Square Error

The last method to compare rival models is a variant of the standard mean squared error (MSE) measure. Since we are dealing with time series with non-constant variance, it is appropriate to alter the MSE calculation so that large returns do not dominate the calculation. Bollerslev and Ghysels (1996) first proposed using a heteroskedastic consistent MSE measure:

$$\mathrm{HMSE} = \frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}}\left[\frac{\epsilon_{\mathrm{t}}^2}{\hat{\sigma}_{\mathrm{t}}^2} - 1\right]^2$$

where ϵ_t^2 is the realized volatility and $\hat{\sigma}_t^2$ is the forecast of the volatility. We will use HMSE as an indicator of how well the models perform out of sample.

1.5 Empirical Results

The data are divided into two sets; the first for estimation, the second for out of sample forecasting. The estimation (or modeling) data set consists of 202 days, leaving 60 days for evaluating forecasts. Thus roughly 30% of the data is set aside for out-of-sample forecasting. Tables A.1 and A.2 provide the descriptive statistics for the modeling and forecasting data sets. The parameter estimates for all the models discussed below are provided in Tables A.3 to A.25. The mean log likelihood estimates for each model is provided in Table A.26; the sample kurtosis resulting from using the estimated variances are provided in Table A.27.

1.5.1 Seasonality

The intraday returns clearly exhibit seasonality. As evident in Figures A.8, A.9, and A.10, the variance during the day rises after 12 GMT (8 EDT) and starts to fall at 18 GMT (14 EDT) for the three currencies. A likelihood ratio test of intradaily homoskedasticity versus intradaily heteroskedasticity rejects homoskedasticity. The test is a chi-square test with 190 degrees of freedom. DM test statistic is 6,519, for JY it is 2,948, and for the SF it is 7,051; the 1% critical value is 238.

Similarly, intradaily homoskedasticity is rejected against the alternative of the exponential seasonality model. The estimated parameters are presented in Table A.3. Graphs of the smoothed variance estimates are shown in Figures A.17, A.19, and A.21.

As for the the merits of a smoothed variance estimate in comparison with the sample variance estimates, the smoothed estimates fare poorly. Using the likelihood dominance criteria for non-nested hypothesis, we reject the smoothed estimates in favor of the sample variance estimates for all three currencies (see Table A.28).

To get a better understanding of why a likelihood ratio test rejects the smoothed estimates in favor of the sample variance estimates, one just needs to compare the smoothed estimate with the sample variance estimates. While to the eye, the smooth estimates are "close" to the sample variance estimates, they are not. In fact a large number of smoothed variance estimates are not within the 99% confidence interval for the sample variance. Between one-third to half of the exponential estimates are outside the 99% confidence interval. Given so many estimates are so far away from the sample variance estimates, it is understandable we reject exponential seasonality in favor of sample variance seasonality.

It turns out that using another smoothing technique or increased number of parameters does not improve the smooth intraday variance estimates. We also estimated another smoothed estimator, a fourier series expansion for comparison. The fourier series expansion with about the same number of degrees of freedom as the exponential model estimated ($m = 4$) (see Figures A.11, A.13 and A.15) was estimated along with ($m = 9$) (see Figures A.12, A.14 and A.16) which has more than twice the number of parameter estimates. It is clear to the eye that having a greater number of sines and cosines (with $m = 9$) provides a better fit. However, we still find that with $m = 9$, there are still roughly 40% of the time specific variance estimates outside the 99% confidence interval of the sample variance estimates (see Table A.29).

We conclude that the parametric form of the seasonality is not an issue since two parameterizations do so poorly. Therefore, we proceed with exponential seasonality since it can be more easily nested in larger models. Given that higher order number of fourier terms only slightly improve the fit, and since higher ordered polynomials suffer from multicollinearity, we limit the order of the exponential seasonal model to degree nine.(We do not investigate the benefits of changing the degree of the polynomial.)

1.5.2 Deseasonalization

The results of deseasonalization are encouraging. The ARCH, GARCH, P-ARCH, and P-GARCH models all improve statistically when the returns are deseasonalized first, the conditional heteroskedasticity is modeled, and the variance forecasts are reseasoned. The likelihood dominance criterion model selection procedure favors the three step deseasonalizing the returns over modeling the raw returns (see Table A.30) and is consistent with the results obtained in Andersen and Bollerslev (1994).

Secondly, using the sample variance to deseasonalize the returns again outperforms using the smoothed variance estimates from the exponential seasonality model. In the four models mentioned above: ARCH, GARCH, P-ARCH, and P-GARCH all perform better when the sample variance is used (see Table A.31).

Lastly, as seen in the tables which report coefficient estimates, there appears to be little variation in the parameter estimates when deseasonalizing with exponential seasonality or sample standard deviation. While no formal tests are presented here, the variations between parameter estimates are well within the heteroskedastic adjusted standard errors for the parameter estimates.

1.5.3 Seasonality & Heteroskedasticity: Joint Modeling

There are three attempts to incorporate the intraday seasonality and the conditional heteroskedasticity. In all cases, the addition of the seasonality term in the ARCH/GARCH framework improves the models. We discuss each below.

Addition of Seasonality

As is to be expected, the addition of a seasonal term in a conditional heteroskedastic model improves modeling. The P-GARCH, P-ARCH, CP-GARCH, CP-ARCH, and S-ARCH models all fail to be rejected versus the alternative of GARCH or ARCH on raw returns (see Table A.32).

S-ARCH vs. CP-GARCH/P-GARCH

As is expected, S-ARCH is favored over P-GARCH framework for all three currencies (see Table A.33). Since there is so much variation in the intraday seasonality. a model which captures it well will outperform a model which is more restrictive in its seasonality. Since P-GARCH has only two "states". the flexibility of the S-ARCH framework is dominant. However. the case is more ambiguous in comparing S-ARCH and CP-ARCH/CP-GARCH. When one uses time of day variance for the conditional periodic variance, there is no clear winner (see Table A.33).

S-ARCH vs. Deseasonalization

Interestingly, at this time, there appears to be little gain in modeling the intraday seasonality and the conditional heteroskedasticity jointly in comparison with deseasonalization. All three classes of models (P-ARCH. CP-ARCH, and S-ARCH) are rejected in favor of the three step procedure: deseasonalization. modeling, and reseasonalization when sample variance is used as the variance to deseasonalize (see Table A.33). The gains from using the sample variance to deseasonalize are greater than the benefits of simultaneous modeling with exponential seasonality.

One interesting observation from the S-ARCH estimation is that of the estimated seasonality. When adjusted for the conditional heteroskedasticity, the seasonal pattern estimated is much more volatile than that of the seasonal exponential model. As shown in Figures A.18, A.20 and A.22, the expected seasonality conditional on time of day is much more peaked.

1.5.4 ARCH vs. GARCH

For many empirical studies, financial returns are modeled with a GARCH model. It is of interest to see if one would be better off with an ARCH specification instead. In our case here, for the majority of the model specifications, the likelihood dominance criterion favors the ARCH(10) over the GARCH(1,1) specification (see Table A.36). With the exception of the JY, the returns are better modeled when deseasonalized with sample variance with ARCH than with GARCH. The conditionally periodic GARCH models (with or without NY dummies) are also rejected, for all three currencies, in favor of a conditionally periodic ARCH specification.

1.5.5 Does Deseasonalization Capture Everything?

The final question we ask is whether we should treat all returns equally. When we deasonalize returns and then estimate a GARCH model, we are assuming that the volatility dynamics are the same within the day. We know that when the NY market is open, we have much greater volatility. Perhaps the dynamics of volatility vary when the NY market is open as well.

Overwhelmingly, the models are improved when a dummy is used as an indicator of when the NY market is open (see Table A.37). We would expect that the P-GARCH model is favored over the GARCH model on raw returns. After all, we know that the volatility is much larger when the NY market is open and a GARCH model on raw returns does not address this point. However, it is not expected that after deseasonalizing the returns that there is still a NY effect. In the deseasoned GARCH, deseasoned ARCH, CP-GARCH, and CP-ARCH models all fail to reject that dummies for NY should be included. The likelihood dominance criterion favors

NY dummies being included for all four classes of models on all three currencies. We are left to conclude that the NY market behaves differently than the London and Tokyo markets.

1.5.6 Forecasting Performance

The results obtained with the modeling data set hold up well out of sample (see Table A.39 for all the results discussed in this section). Using HMSE as the measure for model comparison, all the results, with the exception of the NY dummies being significant remain valid.

We find that out of sample the three step deseasonalizing procedure out performs the joint models. Still, the sample variance seasonality outperforms the exponential seasonality. The sample variance appears to be quite stable as a predictor of future variance. One would think that perhaps a smoothed variance estimate would perform better out of sample, but it does not.

Again, the joint seasonal/heteroskedastic models outperform those which incorporate either seasonality or heteroskedasticity but not both.

Lastly, the NY effect appears to be less pronounced. Ignoring the raw ARCH and GARCH models, the addition of the NY dummies improves the forecast error, even after deseasonalization. For the DM, both the ARCH and GARCH models deseasonalized with sample variance and exponential seasonality have reduced HMSE. For the JY, the GARCH models improve with the NY dummies regardless of method of deseasonalization. However, the addition of the NY dummies does not improve the ARCH models when the returns are deseasoned, again regardless of deseasonalization method. For the SF, the ARCH and GARCH models with deseasoned returns are improved for sample variance deseasonalization, but only the GARCH model does worse with the NY dummies when the returns had been deseasonalized with exponential seasonality.

1.6 Conclusions

To summarize, this paper asks three questions: (1) what is the best way to model intraday volatility, either deseasonalize returns or jointly model seasonality and conditional volatility; (2) which method for deseasonalizing returns is best; and (3) are volatility dynamics constant throughout the day?

We find that:

- Smoothed seasonal volatility does worse than expected: sample variance is best both in sample and out of sample.

- Joint modeling of seasonality and conditional heteroskedasticity is dominated by the three step procedure of deseasonalizing, modeling, and reseasonalization.

- Volatility dynamics does depend on the time of day: the NY foreign exchange market is different and should not be ignored.

To summarize, what we have discovered is that of the two factors in intraday volatility–seasonality and serial correlation, that getting the seasonality correct has the biggest gain in improving forecast error. Care should be taken in modeling seasonality as it is the dominant factor. The joint models fare so poorly because they are handicapped by the limited seasonality in the models. While the gains from modeling the NY effect, at least in sample, are statistically significant, the largest reduction in HMSE comes from using the best seasonality to deseason the returns. While the academic may chose to focus on the dynamics of volatility when the NY market is open, the practitioner should worry about how the intraday seasonality is affected by government announcements and national holidays.

The results are promising and provide guidance for future research. Clearly, volatility dynamics depend upon the time of day. When the New York market is open, volatility is different. Given the recent work of Andersen and Bollerslev, one can conjecture that macroeconomic announcements are playing a role. It seems plausible that announcements made during US business hours are processed by the foreign exchange markets differently than when other nations make announcements. Future work in this area, is of course, necessary.

Chapter 2

Daily Foreign Exchange Volatility: Creating Estimates with Intraday Data

2.1 Introduction and Motivation

There is a growing area of research in finance which deals with modeling volatility of asset markets. With the construction of high frequency data, it is now possible to model volatility in a more precise manner. Previously, one could only address volatility on a week by week or day by day time frame. The initial studies on daily volatility have used prices spaced twenty-four hours apart. From these observations. a return series was calculated, and the market dynamics were analyzed.

Now due to better recording of prices, we have a much richer data set. In some cases, we have transaction by transaction information. Rather than using one price per day, there is the possibility of using thousands of prices per day. One route researchers are taking is to look at the short term dynamics of asset markets. This is understandable since these issues can now be studied.

Many previous daily volatility studies have focused on the issue by using daily data. In order to estimate a day's volatility, GARCH models use previous information. This is feasible since volatility is serially correlated. High volatility days tend to follow each other while low volatility days tend to follow each other.

One problem with such an approach is that small returns for a day do not necessarily mean the underlying asset did not oscillate wildly during that period. All it does mean is that the price levels ended up close to each other at the end of 24 hours. There may or may not have been large movements throughout the day.

With large returns we do not face this problem. Since a large return means that the end point prices have changed a great deal, then one must have large price changes during the day. The larger the return, the more easily one can say the volatility for the day was high.

Fortunately, the availability of intraday data now facilitates the issue. It is now possible to distinguish between a day with very little price change during the trading period, and a day with large price swings which happen to yield a net small price change at the end of the trading period.

For example, for the Deutshemark on Friday, March 11, 1994 and Monday, May 30, 1994 (Figures B.24 and B.23 respectively), the net change in price level is small in comparison with the other daily returns in the sample year. However, both days are not equal in terms of their volatility if we are able to see the price change during the day. In the first case, the price does not change much during the day. In the second case, the price fluctuates much more wildly and ends up close to the opening price. If one were to focus on prices once a day (and thus look at daily returns only), we would be unable to distinguish between the two days.

Similar results hold with days that have large price changes. For the Deutshemark, Thursday, July 21, 1994 and Tuesday, May 10, 1994 both have large price changes during the day (see Figures B.25 and B.26). The price level for May 10th slowly increases as the day progresses, while for July 21st, the price change is much more rapid, perhaps as result of a macroeconomic announcement. Intraday data allows researchers to distinguish days of larger currency appreciation or depreciation.

Previous volatility which looked to estimate daily volatility needed to make inferences regarding the dynamics of the variance process in order to create a daily estimate. In a GARCH model, the variance forecast is conditional on information up to and including the day before.

Recent work with intraday data has focused understanding the dynamics of volatility during the day. The volatility is now understood to have a regular pattern throughout the day, which repeats every twenty-four hours. Andersen and Bollerslev (1997) model this intraday volatility seasonality and control for macroeconomic announcements and holidays. In addition, Andersen and Bollerslev (1997) appeal to variance estimates from intraday returns to show that GARCH models are good predictors of volatility. Zhou (1996) creates a variance estimator using tick-by-tick data, but to do so requires a huge number of observations to create a daily estimate.

The goal of this paper is to show that one can use intraday returns to create reasonable variance estimates, but unlike Zhou (1996) who requires thousands of observations to create daily variance estimates, we find that fifteen minute returns (ninety-six observations) per day are sufficient for good estimates. In addition, we analyze two variance estimates that have been used by Andersen and Bollerslev and show that one is severely biased, as is therefore inappropriate.

This chapter is organized as follows: Section 2.2 describes various ways to evaluate daily variance measures and discuss the limitations of some which have been used in the literature, Section 2.3 describes and compares two intraday volatility estimators, Section 2.4 discusses the empirical results from a one year sample of three currencies, and Section 2.5 concludes.

2.2 Evaluating Daily Volatility Estimates

For exchange rates we assume that the process of exchange rate determination is

$$\begin{aligned} R_t &= \mu + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \epsilon_t^2 &= \sigma_t^2 v_t^2 \\ v_t &\sim N(0, 1) \end{aligned}$$

thus ϵ_t^2 has a conditional mean of σ_t^2 since the expectation of v_t^2 which is χ_1^2 is 1. However, as pointed out by Lopez (1995) the median of a χ_1^2 is 0.455 so we expect that using ϵ_t^2 as an unbiased estimate of σ_t^2 will lead to an underestimate approximately 54.5% of the time. In addition, the distribution of a χ_1^2 is such that for a single observation, the estimate of σ_t using ϵ_t will be within 50 percent of σ_t^2 only a quarter of the time.

Although the variance process is latent, we can still create estimates of this process. However, we must be confident that our estimates are reasonable. Given the assumptions of our model, we expect that our estimates of σ_t^2 should normalize the data[1], i.e. $R_t/\sqrt{\hat{\sigma}_t^2} \sim N(0, 1)$ and similarly, the squared return divided by our variance estimates should be chi-squared, i.e. $R_t^2/\hat{\sigma}_t^2 \sim \chi_1^2$. Thus we can use statistical tests which are based on the normal distribution or the chi-square distribution. Two standard tests are the Bera-Jarque normality tests, and the Kolmogorov-Smirnov cumulative distribution tests.

[1]For exchange rates, theory indicates that we would expect that $E[R_t] = 0$. Formal statistical tests are not reported.

2.2.1 Normality Test

The Bera-Jarque test uses the skewness and kurtosis of a sample to test whether the sample appears to be normal. Under the assumption of normality, the skewness is zero, and the kurtosis is three. Large deviations from either lead one to conclude that the data on hand is not normal. Formally the test is:

$$\mathrm{BJ} = \mathrm{T} \times \left[\frac{\text{skewness}^2}{6} + \frac{(\text{kurtosis} - 3)^2}{24}\right] \sim \chi^2_2$$

2.2.2 Kolmogorov-Smirnov Test

The Kolmogorov-Smirnov test compares the empirical cumulative distribution function to a theoretical cumulative distribution function (CDF). If the null hypothesis that the sample is distributed as the theoretical distribution, then the empirical CDF and theoretical CDF should be statistically close. The test statistic is based on the largest difference between the empirical and theoretical distribution functions (denoted $F_E(x) - F_T(x)$ respectively.)

$$D = \sup_x |F_E(x) - F_T(x)|$$

We can use this test to test both that standardized returns are normal, and that squared returns divided by daily variance are chi-square with one degree of freedom.

2.2.3 R^2 Measures

One measure which researchers have used to evaluate volatility forecasts is looking at the R^2 from the following regression:

$$\epsilon_t^2 = a + b\hat{\sigma}_t^2 + u_t$$

where ϵ_t^2 is the squared return for day t and $\hat{\sigma}_t^2$ is a forecast for day t's volatility, from a GARCH model.

Andersen and Bollerslev (1997) describe why using such a regression is inappropriate. They show that if the variance process is a GARCH process, the expected R^2 is low.

2.2.4 Error Measures

There are also several measures which are frequently used to determine the quality of variance estimates. They are defined below.

Mean Squared Error

$$\mathrm{MSE} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left(\epsilon_{\mathrm{t}}^2 - \hat{\sigma}_{\mathrm{t}}^2\right)^2$$

Mean Absolute Error

$$\mathrm{MAE} = \frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \left| \epsilon_t^2 - \hat{\sigma}_t^2 \right|$$

Logarithmic Loss

$$\mathrm{LL} = \frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \left[\log\left(\epsilon_t^2\right) - \log\left(\hat{\sigma}_t^2\right) \right]^2$$

Heteroskedasticity-adjusted Mean Square Error

$$\mathrm{HMSE} = \frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \left[\frac{\epsilon_t^2}{\hat{\sigma}_t^2} - 1 \right]^2$$

Gaussian Quasi-Maximum Likelihood

$$\mathrm{GMLE} = \frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \left[\log\left(\hat{\sigma}_t^2\right) + \frac{\epsilon_t^2}{\hat{\sigma}_t^2} \right]$$

A couple of issues are worth noting. First, each of the above measures of fit are minimized when one uses the squared return as the estimate for the day's variance. As a result, each measure will prefer variance estimates that are highly correlated with the squared return to those which are less correlated with the squared return. Given the lose assumptions for the variance process of returns, this feature of these statistical measures of fit is undesirable.

Secondly, if we knew the true variance for each day, none of the statistical measures would ever equal zero. That is because each of the measures involves a random variable. For example, suppose we knew the true variance for each day t, then the absolute error for the day would be

$$\begin{aligned} \text{AE} &= |\epsilon_t^2 - \sigma_t^2| \\ &= |v_t^2\sigma_t^2 - \sigma_t^2| \\ &= |\sigma_t^2(v_t^2 - 1)| \end{aligned}$$

so the expectation of $|\sigma_t^2(v_t^2 - 1)|$ is not equal to zero. Theoretically, with perfect estimates for daily variance we are not guaranteed to minimize absolute error, or any of the other measures. Thus variance estimates which are based on these criterion are unwise.

2.3 Creating Variance Estimates with Intraday Data

2.3.1 Modeling Volatility with GARCH

One of the ways to model asset returns has been with a GARCH model. This class of models estimates both a returns process and a variance process for the returns. The usual version used is the GARCH(1,1) model, where the return process is given by

$$\begin{aligned} y_t &= \mu + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

where y_t are the returns for some asset class. Thus we create variance estimates that are conditional on information from time $t-1$ for day t.

Now with intraday data, it is possible to create variance estimates for the day which take into account information from the day, unlike the estimates from GARCH models which use information from previous days only. Andersen and Bollerslev have used two estimators variance estimators which use intraday data. One is a cumulative absolute return estimator, the second is a cumulative square return estimator. We will now discuss both estimators and show that the cumulative absolute returns estimator is biased.

2.3.2 Cumulative Absolute Returns (CAR)

Andersen and Bollerslev (1996) suggest modeling daily volatility by using absolute moments of returns to create variance estimators. We present their approach below.

Assume that our daily return is the sum of intraday returns

$$X = X_1 + \cdots + X_n$$

where $X_i \sim N(0, \sigma_i^2)$, and the X_i's are uncorrelated with each other. Then we have $X \sim N(0, \sigma^2)$ where $\sigma^2 = \sum_{i=1}^{n} \sigma_i^2$.

Then to use absolute returns as measures of volatility as Andersen and Bollerslev suggest yields

$$E(|X|) = \sqrt{\frac{2}{\pi}}\sigma$$

for the daily return. Using intraday returns, we can calculate the cumulative absolute return as

$$
\begin{aligned}
E\left(\sum_{i=1}^{n}|X_i|\right) &= \sqrt{\frac{2}{\pi}}(\sigma_1 + \cdots + \sigma_n) \\
&= \sqrt{\frac{2}{\pi}}\sum_{i=1}^{n}\sigma_i
\end{aligned}
$$

Constant Intraday Variance

Now if all the σ_i are constant, then we have the variance of each X_i as being equal to σ^2/n since

$$
\begin{aligned}
\sigma^2 &= \sum_{i=1}^{n}\sigma_i^2 \\
&= \sum_{i=1}^{n}\sigma^2/n \\
&= \sigma^2.
\end{aligned}
$$

With our assumption of normality, since $X_i \sim N(0, \sigma^2/n)$,

$$
E(|X_i|) = \sqrt{\frac{2}{\pi}}\left(\frac{\sigma}{\sqrt{n}}\right)
$$

Then we would have

$$
\begin{aligned}
E\left(\sum_{i=1}^{n}|X_i|\right) &= \sqrt{\frac{2}{\pi}}n\left(\frac{\sigma}{\sqrt{n}}\right) \\
&= (2/\pi)^{1/2}n^{1/2}\sigma
\end{aligned}
$$

and so, we can create estimates of the daily volatility based on daily data and based on intraday data.

$$\breve{\sigma}_{n=1} = (\pi/2)^{1/2}|X|$$

and

$$\breve{\sigma}_{n>1} = n^{-1/2}(\pi/2)^{1/2}\sum_{i=1}^{n}|X_i|$$

As for the performance of the estimator, having intraday returns is preferred since the variance of the estimator falls as n increases.

$$\begin{aligned}
\mathrm{Var}(\breve{\sigma}_{n=1}) &= \mathrm{Var}\left((\pi/2)^{1/2}|X|\right) \\
&= (\pi/2)\mathrm{Var}(|X|) \\
&= (\pi/2)(n)\mathrm{Var}(|X_i|)
\end{aligned}$$

since we have

$$n\mathrm{Var}(|X_i|) = \mathrm{Var}(|X|)$$

and so in comparison, the variance of the intraday estimator is

$$\begin{aligned}
\mathrm{Var}(\breve{\sigma}_{n>1}) &= \mathrm{Var}\left(n^{-1/2}(\pi/2)^{1/2}\sum_{i=1}^{n}|X_i|\right) \\
&= n^{-1}(\pi/2)\mathrm{Var}\left(\sum_{i=1}^{n}|X_i|\right) \\
&= n^{-1}(\pi/2)\left(\sum_{i=1}^{n}\mathrm{Var}\left(|X_i|\right) + 2\sum_{i=1}^{n-1}\mathrm{Cov}(|X_i|, |X_{i-1}|)\right) \\
&= n^{-1}(\pi/2)\sum_{i=1}^{n}\mathrm{Var}\left(|X_i|\right) \\
&= n^{-1}(\pi/2)n\mathrm{Var}\left(|X_i|\right) \\
&= (\pi/2)\mathrm{Var}\left(|X_i|\right)
\end{aligned}$$

so we can conclude that

$$\mathrm{Var}(\breve{\sigma}_{i=n}) = (1/n)\mathrm{Var}(\breve{\sigma}_{i=1})$$

and so we should use cumulative absolute returns as an estimator for the daily variance rather than the absolute return as an estimate of daily variance.

Bias Under Heteroskedasticity

Now suppose the σ_i's are not constant, that is, the intradaily return volatility is different throughout the day. This is a reasonable interpretation of intraday returns since it has been clearly demonstrated there is a seasonal pattern to intraday returns.[2] It then turns out that the estimator for which it was assumed that the intraday variances were constant, $\breve{\sigma}_{n>1}$, is biased in the presence of non-constant intraday variance.

To see this, recall that our estimator is:

$$\breve{\sigma}_{n>1} = \sqrt{\frac{\pi}{2n}} \sum_{i=1}^{n} |X_i|.$$

Additionally,

$$\begin{aligned}
E[\breve{\sigma}_{n>1} - \sigma] &= E[\sqrt{\frac{\pi}{2n}} \sum_{i=1}^{n} |X_i|] - \sigma \\
&= \sqrt{\frac{\pi}{2n}} \left(\frac{2}{\pi}\sigma_1 + \cdots + \frac{2}{\pi}\sigma_n \right) \\
&= \frac{\sigma_1 + \cdots + \sigma_n}{\sqrt{n}} - \sigma
\end{aligned}$$

[2] Müller et. al.(1990), Baille and Bollerslev (1991), Andersen and Bollerslev(1994).

now $E[\breve{\sigma}_{n>1}] - \sigma < 0$ since

$$\begin{aligned}
\frac{\sigma_1 + \cdots + \sigma_n}{\sqrt{n}} &< \sigma \\
\frac{\sigma_1 + \cdots + \sigma_n}{\sqrt{n}} &< \sqrt{\sigma_1^2 + \cdots + \sigma_n^2} \\
\sigma_1 + \cdots + \sigma_n &< \sqrt{n}\sqrt{\sigma_1^2 + \cdots + \sigma_n^2} \\
(\sigma_1 + \cdots + \sigma_n)^2 &< n\left(\sigma_1^2 + \cdots + \sigma_n^2\right) \\
\sigma_1^2 + \cdots + \sigma_n^2 + 2\sum_{i<j} \sigma_i\sigma_j &< n\sigma_1^2 + \cdots + n\sigma_n^2 \\
0 &< (n-1)\sigma_1^2 + \cdots + (n-1)\sigma_n^2 - 2\sum_{i<j} \sigma_i\sigma_j \\
0 &< \sum_{i<j} \sigma_i^2 + \sigma_j^2 - 2\sum_{i<j} \sigma_i\sigma_j \\
0 &< \sum_{i<j} \sigma_i^2 - 2\sigma_i\sigma_j + \sigma_j^2 \\
0 &< \sum_{i<j} (\sigma_i - \sigma_j)^2
\end{aligned}$$

Hence the cumulative absolute return standard deviation estimator is biased downwards under intraday heteroskedasticity. The greater the difference between the variances for the intraday periods, the greater the bias in the estimator.

2.3.3 Cumulative Squared Returns (CSR)

There is a problem with cumulative absolute returns. Their appeal requires that one assume volatility is constant throughout the day. However, from the intraday volatility research, we know that is not true. We will now show how using cumulative squared returns as the basis for an estimate for a day's volatility is robust to intraday time varying volatility. As it a result, it produces superior estimation of daily volatility.

Assume that our daily return is the sum of intraday returns

$$X = X_1 + \cdots + X_n$$

where $X_i \sim N(0, \sigma_i^2)$, and the X_i's are uncorrelated with each other. Then we have $X \sim N(0, \sigma^2)$ where $\sigma^2 = \sum_{i=1}^{n} \sigma_i^2$.

Then squared returns are used as measures of volatility since

$$E(X^2) = \sigma^2$$

for the daily return. Using intraday returns, we can use the cumulative squared returns to yield another variance estimator

$$\begin{aligned} E\left(\sum_{i=1}^{n} X_i^2\right) &= \sum_{i=1}^{n} E(X_i^2) \\ &= \sigma_1^2 + \cdots + \sigma_n^2 \\ &= \sigma^2 \end{aligned}$$

and so, we can create estimates of the daily volatility based on daily data and based on intraday data.

$$\hat{\sigma}_{n=1}^2 = X^2$$

and

$$\hat{\sigma}_{n>1}^2 = \sum_{i=1}^{n} X_i^2$$

As for the performance of two estimators, the intraday variance estimator is superior since it has a smaller variance.

To show this, notice that we can write the daily squared return in terms of the intraday returns,

$$\begin{aligned} X &= (X_1 + \cdots + X_n)^2 \\ &= \sum_{i=1}^{n} X_i^2 + \sum_{i<j}^{n} X_i X_j \end{aligned}$$

as a result. the variance of $\hat{\sigma}^2_{n=1}$ is greater than the variance of $\hat{\sigma}^2_{n>1}$ since

$$\begin{aligned} \mathrm{Var}(\hat{\sigma}^2_{n>1}) &< \mathrm{Var}(\hat{\sigma}^2_{n=1}) \\ \mathrm{Var}\left(\sum_{i=1}^{n} X_i^2\right) &< \mathrm{Var}\left(\sum_{i=1}^{n} X_i^2 + \sum_{i<j}^{n} X_i X_j\right) \end{aligned}$$

Thus we should use cumulative squared returns as an estimator for the daily variance rather than the absolute return as an estimate of daily variance since it is more efficient.

The final issue is an empirical one. We need to (1) make sure that our variance estimates appear to be reasonable, and (2) know with what frequency intraday day returns are needed in order to obtain the benefits of the intraday estimators. On the first point, we would like to demonstrate that we can create practical variance estimates with an intraday estimator. On the second point, we need to know if, for example, hourly returns sufficient to obtain a good daily variance estimate or are minute returns required, or do we need something in between.

2.4 Empirical Results

2.4.1 Data

The data consists of 262 daily observations for the Dollar-Deutschemark, Dollar-Swiss Franc, and Dollar-Yen, from September 15, 1993 to September 16, 1994. The sample consists of all weekdays with the exception of December 24 (December 25 and January 1 occurred on weekends.)

For each day, there are returns for every fifteen minutes, yielding a total of ninety-six observations per day. The returns are calculated as log returns and are multiplied by 100. The notation we will use will be

$$R_{t,i} = \log(P_{t,i}) - \log(P_{t,(i-1)})$$

where $t = 1, 2, \ldots, T$, and $i = 1, 2, \ldots, n$. We can define R_t to be the daily realized return, and $R_{t,i}$ to be a return during part of the day, where the time interval is $1/N$. The time variable t corresponds to days, while n represents the intraday cycle. For example, if $n = 96$, then $R_{t,i}$ is a fifteen minute return. Each day's returns go from 12AM GMT to 11:45PM GMT, and are not corrected for daylight savings time.

From the intraday returns, the daily return series can be computed as the sum of ninety-six returns,

$$R_t = \sum_{i=1}^{96} R_{t,i}$$

Tables B.40 and C.62 provide the descriptive statistics for the fifteen-minute and daily returns. See Figure C.44 for a plot of the Deutschemark daily returns.

2.4.2 Daily Variance Estimates

Initial Impressions of CSR Variance Estimates

The cumulative squared returns (CSR) variance estimates using fifteen minute returns are shown in Figure C.46. Compared to the daily squared return (see Figure C.45), the cumulative squared return volatility estimates are much less volatile. As shown in Table C.63, the variance of the squared returns are 0.3886, 1.0289, and 0.5331, respectively for the DM, JY and the SF. The variance of the CSR variance estimates are 0.1281, 0.2104, and 0.0740 respectively (see Table C.64). Similarly the range of the cumulative squared return variance estimates is smaller than the range for daily squared returns for the three currencies.

The returns for each currency are fat tailed (i.e. have excess kurtosis). For example, Figures B.31 and B.32 which shows a histogram for DM returns and the gaussian kernel density estimate with the normal distribution with the same mean and variance as the sample.

As for the distribution of the CSR volatility estimates, they are distributed differently than the daily squared returns. For example, the DM histogram for squared returns has many of the returns in the left most bin. In comparison, the CSR variance estimates are shifted to the right. The median squared return for the DM, JY and SF is 0.1483, 0.1624, and 0.1704, respectively. The median CSR variance estimate is 0.2954, 0.4014, and 0.4406 respectively.

Joint Distribution of Squared Return and CSR Variance Estimates

We have thus seen that the CSR variance estimate appears to be consistent with our theoretic predictions for a variance estimator. We now discuss how the squared return compares with the CSR variance estimate.

Table B.51 provides a summary of the joint distribution of the squared daily return and the corresponding CSR variance estimates. The table partitions the joint density into four parts, by using the median squared return and the median CSR estimate as divisors. We have already seen that on average the unconditional variance calculated from the returns is equal to the unconditional variance calculated using the average CSR variance estimate. However, for each individual day, predicting the daily variance using CSR will yield a different estimate than using the squared daily return.

As in Table B.51, roughly one third of the time, both the squared return is below the median squared return (is "low") and the corresponding CSR estimate for the day is also below the median CSR estimate (is "low".) Similarly, one third of the time, the squared return is above the median squared return (is "high") and the CSR estimate for the day is above the median CSR estimate (is "high".)

When both the squared return and CSR are low, the CSR estimate is almost always higher than the squared return. When the squared return is low and the CSR is high, then the CSR estimate is always larger than the squared return.

When the squared return is high and the CSR is high, then it is roughly equally likely that the CSR will be larger than the squared return, and vice-versa. (For the SF it is 55 of the 85 days have the CSR larger than the squared return.) When the squared return is high and the CSR is low, then the squared return is larger than the CSR about 2 to 1.

When we look at the variance estimates for the four "samples", we see that when the squared return is low and the CSR is low, the mean CSR is much greater than the mean squared return. Similarly, we obtain the same results when the squared return is low and the CSR is high. We conclude that squared return for those days most likely understates the true variance that day. When we look at the high squared return days, the average squared return is much larger than the average CSR. Thus we conclude that on those days the true variance for the day is smaller than the squared return suggests.

It is also clearer why Andersen and Bollerslev (1996) use CAR variance estimates to model DM volatility. They feel it is a good daily estimate because it is more highly correlated with GARCH forecasts. Unfortunately, this is due to GARCH forecasts having a very small variance when compared to CSR estimates (see Table B.53.) Since CAR is less volatile than CSR, it is going to be more strongly correlated with GARCH forecasts (see Table C.61.)

Correlations

While numerous studies have shown that squared returns are positively correlated, for two of the currencies that is not true. Both the DM and SF have a negative correlation. While this result is surprising, it is not unexpected since we only have one year of data to analyze. However, when we look at the serial correlation of CSR variance estimates (see Table C.61), all three currencies exhibit a positive correlation: 0.1569 for the DM, 0.3765 for the JY, and 0.2080 for the SF. (The serial correlation for CAR turns out to be higher than of CSR, being 0.4011, 0.5366, and 0.3543 for the DM, JY and SF respectively. This is misleading since the CAR has a smaller variance than CSR, and is very strongly correlated with CSR, it is to be expected that it exhibits higher serial correlation.)

The CSR variance estimates also exhibit higher correlation at longer lags. In exchange rate studies there is a significant long decay in the autocorrelogram for squared returns which has lead some researchers to model exchange rate volatility with fractionally integrated GARCH models.[3] With one year of data, there is no such pattern for the three currencies. Again, the CSR variance estimates exhibit the pattern found in longer samples when using squared returns but with much less data. While we are not modeling volatility dynamics, we can conclude that the CSR variance estimates appear to have the same underlying dynamic structure that daily squared returns possess.

Bias in CAR Variance Estimates

It is clear that cumulative absolute return (CAR) variance estimates using fifteen minute returns are biased downward. This is evident by comparing the mean CAR estimate (see Table B.43) with those of the CSR estimates. In addition, the variance of the CAR estimates in each case is lower than the variance of the CSR estimates. Visually, the graph of CAR volatility (see Figure B.30) is much more dampened when compared to that of the CSR (see Figure C.46.)

As previous mentioned, the sample means of the CAR variance estimates are lower than those of the CSR variance estimates. Therefore, using the mean of CAR as an estimate of the unconditional variance produces a lower estimate for the unconditional variance than when using CSR. Table B.45 provides a list of

[3]Andersen and Bollerslev (1996), Baille, Bollerslev and Mikkelsen (1996), and Bollerslev and Mikkelsen (1996).

four sample estimates for the sample of the three currencies along with confidence intervals for the estimates.

The sample variance estimates are

$$\sigma^2 = \frac{1}{T}\sum(R_t - \bar{R})^2$$

If we use the absolute value of the return to create an estimate of the standard deviation, then we can use the square of the estimator and use the average as an estimate of the unconditional variance.

$$\sigma^2 = \left(\frac{1}{T}\sqrt{\frac{2}{\pi}}(\sum|R_t|)\right)^2$$

The sample averages using CAR and CSR are just,

$$\sigma^2 = \frac{2}{\pi N}\frac{1}{T}\sum_t^T\left(\sum_{i=1}^n|R_{t,i}|\right)^2$$

and

$$\sigma^2 = \left(\frac{1}{T}\sum_t^T(\sum_{i=1}^n R_{t,i}^2)\right)^2.$$

From the four sets of estimates, we see that the absolute return unconditional variance estimates are lower than that from using realized returns. However, the difference does not appear to be significant since the confidence intervals overlap. The CSR unconditional variance estimates are a little higher than that of the realized returns unconditional variance estimates. Again, the confidence intervals overlap. For the CAR unconditional variance estimates, at least for the DM and the SF, a different story emerges. The point estimates, given the confidence intervals, are low enough to be statistically significant at the 5% level when compared to the realized return unconditional variance estimates. When comparing the CSR and CAR unconditional estimates, we reject that the unconditional variance estimates are equal.[4]

The theory predicts that the CAR daily standard deviation estimate will be biased downwards. While we do not know the true volatility for the day, we can compare the CAR estimates to that of the CSR estimates. Tables B.49 and B.50 give the descriptive statistics in the percentage difference of CAR standard deviation estimates to that of the CSR standard deviation estimates, and likewise with the variance estimates. As presented in the tables, the CAR estimates are almost always below that of the CSR estimates. Only for three days with the JY is the CAR estimate higher than the corresponding CSR estimate; for the SF and DM the CAR estimate is always lower than the CSR estimates. On average the CAR understates the daily variance by 25.86 percent for the DM, 22.12 percent for the JY, and 29.58 for the SF in comparison with CSR variance estimate. For the standard deviation estimates, CAR understates by 14.30 percent for the DM, 12.04 percent for the JY, and 16.42 for the SF.

Finally, while CAR is a biased estimator for daily variance, it sometimes appears to be an adequate measure of a day's volatility. This is a result of the high correlation between CSR and CAR estimates (0.9063, 0.9390, and 0.877 respectively for DM, JY and SF.) For example, standardizing returns with CSR estimates yield returns that look normally distributed, and as a result of the high correlation, standardizing with CAR does also. Thus one may be mislead into using CAR as a variance estimate.

[4] In GARCH models, the parameter estimates are constrained under estimation to be related to the unconditional variance, since $\sigma^2 = \omega/(1-\alpha-\beta)$.

2.4.3 Normality Tests

We see (see Tables B.47) that the Bera-Jarque normality test rejects normality for the DM and JY returns at the 1% level, and rejects normality for the SF at 5%. If we use the CSR as a variance estimate for the day, and standardize the returns, that for all three currencies, we reduce the Bera-Jarque statistic significantly. (Again, the CAR also seems to standardize the returns since it is so highly correlated with CSR.) In comparison, using the conditional variance estimates obtained from a GARCH model improves the Bera-Jarque statistic, but only slightly, and insignificantly when compared to the reduction when using CSR. We conclude that the CSR variance estimates effectively "explains" the fat-tailed nature of daily returns.

2.4.4 R^2 Measures

While one may not be interested in using the regression, $\epsilon_t^2 = a + b\hat{\sigma}_t^2 + u_t$, as a measure to evaluate volatility forecasts, it is useful to compare how the CSR variance estimates perform in relation to GARCH variance estimates. This is of interest since many papers report the R^2 from the above regression.

Table B.55 reports the R^2 for a number of volatility estimates, $\hat{\sigma}_t^2$. We see that when using a GARCH(1,1), the in-sample estimates for volatility yield an R^2 of 0.001, 0.003, and 0.000, for the DM, JY and SF respectively. These results, are in the low range compared to those in the reported in the literature, which range from 0.001 to 0.106, but have more than one year of data and use daily returns.

Finally, we see that the CSR estimates have values of R^2 of 0.216, 0.379, and 0.189. While we expect a higher R^2 from CSR, the results help to provide a picture of what the best case of what the R^2 could look like from a forecasting model. In any case, the results here are higher than any reported from using GARCH models. (Notice that the CAR estimates dominate the CSR estimates for the DM and JY. This is understandable since for those two currencies the correlation between the squared daily returns and the CAR variance estimates is higher for the squared daily returns and the CSR variance estimates. Since CAR variance estimates are biased we still do not want to use them.)

2.4.5 Error Measures

In our case, mean square error and their variants are more useful in providing a comparison of the estimates rather than as a method to rank intraday estimators (see Table B.54.) The reason for that, is that each uses the squared return as the "best" estimate for the volatility for the day. As a result, estimators which are correlated with the squared return will outperform those which are less correlated. But, it is apparent that CSR (and sometimes CAR) both do a much better job as a variance estimate than a GARCH(1,1) model does since on all five measures the CSR estimate yield a lower value than the GARCH(1,1) estimates.

However, it is clear that the LL and MAE measures are inappropriate to use when evaluating volatility estimates. Both favor CAR over the CSR estimates. The GMLE measure also has trouble since it incorrectly favors CAR once out of the three currencies. Only the MSE and HMSE measures favor the CSR variance estimates. These results suggest that researchers be careful in how the evaluate volatility estimates or forecasts. Basing model selection on LL, MAE, or GMLE may produce erroneous results and may thus be inappropriate.

2.4.6 Frequency of Returns

The final issue to address is one of what frequency of returns should researchers analyze. For example, much of Andersen and Bollerslev's recent work on announcement effects and seasonality has used 5-minute returns, while Müller et. al.(1990) have used hourly returns to model daily seasonality.

Current research has shown that there are short run events which take place during macroeconomic announcements. These studies have shown that it takes from two to four hours for the new information to be processed and a new equilibrium to be established. Thus, one would expect that intraday data would have to be of a high enough frequency to capture the market's reaction to news as it is happening. For this reason, we would expect four hour returns (or longer) to look very much like daily returns.

Researchers who have looked at ultra-high frequency data, which are transactions data, have found considerable noise due to several effects, namely the bid-ask bounce, and from a Zhou's (1996) fighting-screen effect where institutions update their quotes so that they are visible on computer screens. Clearly, if one is interested in daily volatility, we wish to avoid the noise from individual transactions and focus on time periods where one could reasonably conclude that price changes are unlikely to be the result of a bid-ask spread, but from a real price change.

To address which frequency of returns should be used, we can turn to some of the same measures we have already used.

Table B.56 shows the Bera-Jarque normality test statistic when standardizing with different frequencies of intraday returns. Except for the JY for which it does not seem to matter which frequency is used, the DM and the SF test statistic increase as the sampling frequency for returns falls.

As for testing the empirical distribution for standardized returns (see Table B.59), we obtain different results as compared to the the Bera-Jarque test. We find that up to (an including) 180 minute returns appear to be satisfactory for creating variance estimates. We fail to reject that standardized returns are normal for all three currencies.

On the other hand, when we test if the distribution of the squared daily returns divided by CSR variance estimates appear to be chi-square with one degree of freedom, we obtain different results (see Table B.60) For the DM and JY, the CSR variance estimates fail to be rejected with 15 to 30 minute returns only. For the SF, we reject that 15 minute returns produce valid variance estimates, but we fail to reject that 30 minute to 180 minute returns are adequate. The result for the SF is puzzling, but we can conclude that 30 minute and 15 minute returns are reasonable to create CSR variance estimates.

Unfortunately, the mean square error measures, and their variants are useless in determining which frequency to choose. The problem is that with lower sampling of returns, the correlation between the squared return and the CSR variance estimates increases. The MSE measures all improve the higher the correlation between the variance estimate and the squared return. In fact, they are completely minimized when one uses the squared return as an estimate for the variance. Unfortunately, as discussed previously, this is a poor measure of variance.

Lastly, we can look at the serial correlation in the CSR estimates as the returns are aggregated (see Table B.58). For the SF, the serial correlation is highest for the 15 minute return CSR. For the DM, the serial correlation is high until 90 minute returns are used; for the JY until 180 minute returns are used. Thus if one is interested in volatility dynamics, it appears that 15 minute returns are the safest level of sampling.

Thus we can conclude that while this analysis has been ad-hoc, the results appear to be consistent. If one uses 15 minute returns, then one can create good variance estimates for the day.

2.5 Conclusions

Previous studies which studied the time varying nature of volatility have had to rely on GARCH or stochastic variance models to both estimate a daily variance and to model the volatility dynamics. We have shown that in order to have a variance estimate for the day, it is no longer necessary to use a GARCH model. In fact, GARCH(1,1) appears to be rather poor at estimating daily volatility. We can now use intraday data to create a good estimate for the variance of the day using cumulative squared returns on fifteen minute returns (or 30 minute returns.) Unlike a GARCH model, this estimate is now independent of how we think volatility today is related to volatility yesterday.

Secondly, cumulative squared return volatility estimates will provide a better measure to evaluate volatility forecasting models. Previous models have relied on squared returns as a measure for the "true" variance for the day. As shown, squared returns tend to understate the variance on some days and overstate the variance on other days. As a result, we are in a position to better rank competing forecasting models.

Finally, it appears that using cumulative squared returns as the "true" volatility for the day in place of squared returns will improve volatility forecasting models.

Since the cumulative squared return gives a better picture of a day's volatility, we should get a more precise picture of the association between today's volatility and yesterday's volatility. Thus the GARCH model's use becomes strictly one of forecasting volatility instead of estimating and forecasting simultaneously. We leave such issues to future research.

Chapter 3

Daily Foreign Exchange Volatility: Predicting with Intraday Data

3.1 Introduction and Motivation

This paper analyzes how volatility forecasts in the foreign exchange market can be improved with intraday data. Intraday data now allows us to create variance estimates which are independent of how we model volatility dynamics. Chapter 2 shows that using fifteen minute returns, one can create good variance estimates using cumulative squared returns.

$$\hat{\sigma}_t^2 = \sum_{i=1}^{96} R_{t,i}^2$$

where $R_{t,i}$ is the fifteen minute return for day t for interval i during the day.

Without intraday data, researchers have typically used GARCH models to estimate the time-varying nature of volatility. This involves having to specify volatility dynamics in order to estimate volatility for a given day. In GARCH models, variance estimates are conditional estimates on lagged information.

The benefit to using intraday variance estimates is seen in Figures C.42 and C.43. Figure C.42 shows how the squared daily return and the previous squared daily return are related graphically. We would expect that there be a positive correlation between squared returns and its own lag. It is the positive correlation which is utilized by the GARCH class of models. However, in Figure C.42 we fail to see a positive relationship at all. In fact it appears as if the relationship is a reciprocal one. Figure C.43, on the other hand, provides us with a much greater association between the cumulative squared return series, than the daily squared return series. It yields the relationship we expect to see. Consequently, one expects that using cumulative squared returns as the basis for forecasting variance would yield better forecasts, and the results are consistent with that intuition.

This paper uses intraday data variance estimates to do two things, (1) better model volatility dynamics for the purposes of forecasting and (2) better evaluate volatility estimates since we have less noisy volatility estimates.

We find that (1) the usual GARCH models are improved when we use intraday data for calculating the variance, (2) forecasting models which fit intraday variance estimates rather than squared returns yield better forecasts, (3) forecasts from models using intraday data yield much greater variance in the forecasts, and (4) intraday volatility models are less persistent than GARCH models: forecasts place a greater weight previous period.

The chapter is organized as follows. Section 3.2 presents various models which utilize intraday variance estimates. Section 3.3 provides the empirical results. Section 3.4 concludes.

3.2 Modeling Volatility Dynamics

From Zhou (1996) and Chapter 2 we know that intraday data can be used to create variance estimates which are good measures of the volatility for the day. This is done by using the returns from the day to calculate the variance of the day. Unfortunately, for forecasting one still needs to specify a dynamic model. However, the intraday variance measures can be used as a proxy for the true latent variance for the day.

One of the ways to model asset returns and volatility has been with a GARCH model. This class of models estimates both a returns process and a variance process for the returns. The variance process is worth adding because it has been observed that a large return will tend to be followed by a large return, and a low return will tend to follow a low return, and as a result volatility is predictable.

The usual version which is used is the GARCH(1,1) model, where the return process is given by

$$\begin{aligned} y_t &= \mu + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

where y_t are the returns for some asset class. One thing worth mentioning is that the GARCH model can be written as an ARMA model in squared residuals. This can be shown by noting that

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

is equivalent to

$$\begin{aligned} \epsilon_t^2 &= \omega + (\alpha + \beta)\epsilon_{t-1}^2 + v_t - \beta v_{t-1} \\ v_t &= \epsilon_t^2 - \sigma_t^2 \\ &= (\eta_t^2 - 1)\sigma_t^2 \end{aligned}$$

where

$$\eta_t \overset{\text{iid}}{\sim} N(0, 1)$$

and so $E(v_t|I_{t-1}) = 0$, where I_{t-1} represents all information available to the modeler at time $t-1$. In this framework, the v_t is the "error" and can also be thought of as the difference between the one step ahead forecast of the squared error term and the realized squared error. Thus the "shocks" are seen as unexpected differences in the forecasted squared error and the realized square error. As a result, GARCH models are dynamic in their specification of the volatility process and are used to study the time varying nature of volatility of returns.

Notice that in a GARCH model, that the information from which to make inferences from the volatility returns come from the returns themselves (in this case, lagged squared returns). If one had additional information, then one would not have to use the returns to both describe its movement over time and how its own volatility changes over time.

3.2.1 Value of Intraday Variance Estimates

The most straight forward way to use intraday data is to replace the lagged squared error term with an intraday variance estimate. Now the estimation of the GARCH(1,1) model can be improved with intraday data. In the variance equation, in place of the ϵ_t^2, we can use a variance estimator which is more efficient. We will be using cumulative squared returns, $\hat{\sigma}_t^2 = \sum_i^n R_{t,i}^2$, in its place. Thus, we transform a GARCH model into a GARCH-CSR model. One can then compare how a GARCH(1,1) model compares to a GARCH-CSR model (and likewise a constant variance model).

Model 1: Constant Volatility

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma^2) \\
\sigma_t^2 &= \sigma_0^2 \\
\log(L) &= -\frac{1}{2}\log(\pi) - T\log(\sigma_0^2) - \sum_t^T \epsilon_t^2/\sigma_0^2
\end{aligned}
$$

Model 2: GARCH

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_{t|t-1}^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1|t-2}^2 \\
\log(L) &= -\frac{1}{2}\log(\pi) - \sum_t^T \log(\sigma_{t|t-1}^2) - \sum_t^T \epsilon_t^2/\sigma_{t|t-1}^2
\end{aligned}
$$

Model 3: GARCH-CSR

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_{t|t-1}^2 &= \omega + \alpha\left(\sum_{i=1}^n R_{t-1,i}^2\right) + \beta\sigma_{t-1|t-2}^2 \\
\log(L) &= -\frac{1}{2}\log(\pi) - \sum_t^T \log(\sigma_{t|t-1}^2) - \sum_t^T \epsilon_t^2/\sigma_{t|t-2}^2
\end{aligned}
$$

In our case, it happens that we can create a model that nests both the GARCH and the GARCH-CSR models. Since $\epsilon_t^2 = \sum_{i=1}^n R_{t-1,i}^2 + \sum_{i<j}^n 2R_{t-1,i}R_{t-1,j}$, we can

estimate the following model,

$$
\begin{aligned}
y_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_t^2 &= \omega + \alpha_0 \sum_{i=1}^{n} R_{t-1,i}^2 + \alpha_1 \sum_{i<j}^{n} 2R_{t-1,i}R_{t-1,j} + \beta\sigma_{t-1}^2
\end{aligned}
$$

when α_1 and α_0 are restricted to be equal, we estimate a GARCH model. When we restrict α_1 to be equal to 0, we estimate a GARCH-CSR model. We can therefore test whether $\alpha_1 = \alpha_0$ and is positive, or if $\alpha_1 = 0$. That is, do the cross terms help us forecast volatility.

3.2.2 Models: Creating Better Volatility Forecasts

The problem with forecasting variance is evaluating the success of a model. For instance, researchers have used mean squared error measures, to evaluate how close a variance forecast is to the squared error. The problem is that while the squared error is an unbiased estimate of the volatility, it is a very noisy one. Therefore, estimation procedures which focus on squared errors in the likelihood function are try to estimate parameters to a very noisy signal. Rather than focusing on noisy variance estimates, we can instead focus on better variance estimates from cumulative squared returns.

Since we now have better estimates of volatility during the day, we should use those estimates in model building. Andersen and Bollerslev (1997) use intraday variance estimates to show that GARCH forecasts are reasonable by comparing them to intraday variance estimates. They limit the use of intraday variance estimate to evaluating the performance of GARCH estimates. However, we can also use intraday variance estimates in our likelihood functions.

This bring us to two different models. The first is an ARMA model on the cumulative squared return (Model 4). This model attempts to minimize the sum of squared errors between estimates of a forecast variance and the "observed" estimate.

The second model (Model 5) is a variant of the GARCH model, where the forecast of the variance at time t is conditional on information available at time $t-1$, but unlike the GARCH model, the likelihood function replaces the squared error from the mean process with the cumulative squared return.

Model 4: ARMA

$$
\begin{aligned}
\sigma_t^2 &= \sum_{i=1}^{n} R_{t,i}^2 \\
\sigma_t^2 &= \psi_0 + \psi_1 \sigma_{t-1}^2 + \epsilon_t + \psi_2 \epsilon_{t-1} \\
\epsilon_t &\sim N(0, \sigma_\epsilon^2) \\
\log(L) &= -\frac{1}{2}\log(\pi) - T\log(\sigma_\epsilon^2) - \sum_{t}^{T} \epsilon_t^2/\sigma_\epsilon^2
\end{aligned}
$$

Model 5: Conditionally Heteroskedastic CSR Variance

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_t^2 &= \left(\sum_{i=1}^{n} R_{t,i}^2\right) \\
\sigma_{t|t-1}^2 &= \omega + \alpha\sigma_{t-1}^2 + \beta\sigma_{t-1|t-2}^2 \\
\log(L) &= -\frac{1}{2}\log(\pi) - \sum_{t}^{T}\log(\sigma_{t|t-1}^2) - \sum_{t}^{T}\sigma_t^2/\sigma_{t|t-1}^2
\end{aligned}
$$

We can think of these two models as trying to fit volatility. The ARMA model fits the CSR variance estimates directly. The Conditionally Heteroskedastic model (Hetero-CSR), is similar to a GARCH model but with two differences. Overall, the GARCH model uses daily information, while the Hetero-CSR model uses intraday information. As a result, the variance equation for the Hetero-CSR model use the lagged cumulative squared return rather than the lagged squared error. In the log-likelihood, the Hetero-CSR model substitutes the CSR variance estimate for the lagged squared error. Both terms have the same expectation, i.e. $E[\epsilon_{t-1}^2] = E[\sum_{i=1}^{n} R_{t,i}^2] = \sigma_{t-1}^2$, but the cumulative squared return is a better estimate, then we would expect more efficient estimates and better forecasts. Consistency should not be an issue since we are substituting a more efficient estimate.

3.2.3 Model Evaluation

Normality Test

The Bera-Jarque test uses the skewness and kurtosis of a sample to test whether the sample appears to be normal. Under the assumption of normality, the skewness is zero, and the kurtosis is three. Large deviations from either lead one to conclude that the data on hand is not normal. Formally the test is:

$$\text{BJ} = \text{T} \times \left[\frac{\text{skewness}^2}{6} + \frac{(\text{kurtosis} - 3)^2}{24}\right] \sim \chi_2^2$$

Error Measures

There are also several measures to determine the quality of variance estimates (see Lopez (1995) for a good discussion). We recall their definitions from Chapter 2:

Mean Squared Error

$$\mathrm{MSE} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left(\epsilon_{\mathrm{t}}^2 - \hat{\sigma}_{\mathrm{t}}^2\right)^2$$

Mean Absolute Error

$$\mathrm{MAE} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left|\epsilon_{\mathrm{t}}^2 - \hat{\sigma}_{\mathrm{t}}^2\right|$$

Logarithmic Loss

$$\mathrm{LL} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left[\log\left(\epsilon_{\mathrm{t}}^2\right) - \log\left(\hat{\sigma}_{\mathrm{t}}^2\right)\right]^2$$

Heteroskedasticity-adjusted Mean Square Error

$$\mathrm{HMSE} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left[\frac{\epsilon_{\mathrm{t}}^2}{\hat{\sigma}_{\mathrm{t}}^2} - 1\right]^2$$

Gaussian Quasi-Maximum Likelihood

$$\mathrm{GMLE} = \frac{1}{\mathrm{T}} \sum_{\mathrm{t}=1}^{\mathrm{T}} \left[\log\left(\hat{\mathrm{h}}_{\mathrm{t}}^2\right) + \frac{\epsilon_{\mathrm{t}}^2}{\hat{\sigma}_{\mathrm{t}}^2}\right]$$

However, we have seen in Chapter 2 that some of these measures do not do a good job at discriminating against various variance estimates. The MAE, LL, and GMLE all favored a biased variance estimate on occasion. In addition, since we know we have heteroskedasticity, it makes sense to favor the HMSE measure instead of the MSE measure. For these reasons, we will present error measures produced by all five error measures but only discuss the results from the HMSE.

There will be two versions of HMSE that will be analyzed. The first is based on the squared daily return. The second is based on the cumulative squared return. Given our assumptions about volatility, the expectation at time t is that both are equivalent. That is, $E[R_t^2] = E[\mathrm{CSR}_t]$. However, as shown in Chapter 2, a volatility estimate based on cumulative squared returns is more efficient. Since the daily volatility literature has used the first volatility measure, it is therefore useful to use both to evaluate forecasts.

R^2 Measures

We can also determine what the coefficient of determination is from the following regression:

$$\hat{\sigma}_t^2 = a + b\hat{h}_t + \epsilon_t$$

and see how close the variance forecasts $\hat{h}_t$ are to the estimated variances, $\hat{\sigma}_t^2$.

Persistence Measures

From a GARCH model the persistence of "shocks" is given by various "half life" measures. It is of interest to see if using variance estimates based on intraday data changes our estimates of the decay of shocks to volatility when compared to the results discussed in the literature.

Two commonly used measures are:

Half Life of Lag

$$-\frac{\log(2)}{\log(\alpha + \beta)}$$

Mean Lag

$$\frac{\alpha}{1-\alpha-2\beta+\alpha\beta+\beta^2}$$

For Models 2, 3, and 5, the calculations are straight forward. To obtain the equivalent persistence measures for Model 4, we solve for the α and β parameters in terms of ψ_1 and ψ_2, using the ARMA representation of a GARCH model. This yields $\psi_1 = \alpha + \beta$ and $\psi_2 = -\beta$.

3.3 Dynamic Models Empirical Results

3.3.1 Data

The data used in this chapter is the same as that used in Chapter 2. It consists of 262 daily observations for the Dollar-Deutschemark, Dollar-Swiss Franc, and Dollar-Yen, from September 15, 1993 to September 16, 1994. The sample consists of all weekdays with the exception of December 24 (December 25 and January 1 occurred on weekends.)

Table C.62 provides the descriptive statistics for the daily returns. See Figure C.44 for a plot of the Deutschemark daily returns.

3.3.2 GARCH & GARCH-CSR

For the three currencies, the parameter estimates for GARCH (Table C.65) and the parameter estimates for GARCH-CSR model (Table C.66) are very similar. In all three cases, the constant term in the mean equation is indistinguishable from zero. In the variance equation, the major difference in the two parameter estimates is in the α estimates. In the GARCH-CSR models, the estimates increase to 0.0369, 0.1014, and 0.0208 for the DM, JY and SF, from the GARCH estimates of 0.0166, 0.0393, and 0.0063. While the difference is large in terms of a magnitude change, it is well within the estimated standard errors. In fact, the GARCH α estimates are quite low compared to Jorion (1995). He reports α estimates for the DM, JY and SF, from roughly 1985 to 1992, of 0.0757, 0.0965, and 0.0521. The smaller estimates for α in our case is most likely due to the small sample on hand. In this context, it appears as if the GARCH-CSR is consistent with GARCH models.[1]

As for the nested hypothesis that the variance equation is GARCH-CSR rather than GARCH, unfortunately, the difference in maximized log-likelihoods is too small to be statistically significant. A constrained maximum likelihood estimation places zero weight on the cross terms from the intraday returns ($\sum_{i<j}^{n} 2R_{t,i}R_{t,j}$).

HMSE Measures

The GARCH and GARCH-CSR one-step ahead in sample forecasts are quite similar for the DM (Figure C.47) and the JY (Figure C.51) and less so for the SF (Figure C.55). The correlation of the estimates is 0.792, 0.777, and 0.253 respectfully (Table C.61). The GARCH-CSR estimates themselves vary much more than the GARCH estimates, and have a larger range. Thus they sometimes yield higher or lower estimates than the GARCH model does. (See Figures C.47, C.51 and C.55).

With respect to the HMSE results, the GARCH-CSR model reduces HMSE for all three currencies, and regardless of how HMSE is calculated. For the HMSE

[1]The β parameters in Jorion (1995) are 0.8828, 0.7870, and 0.9042 for the DM, JY and SF, which is lower than those reported here. However the standard errors of our estimates are roughly within one standard error of Jorion's β estimates.

based on squared returns (see Table C.71. GARCH-CSR has an HMSE of 2.6873. 4.3256. and 2.6094 for the DM. JY and SF respectfully. While GARCH model produces a higher HMSE of 2.9247. 5.0515. and 2.7526 respectfully. The results do not change when we use an HMSE based on CSR (see Table C.72. Again. the GARCH-CSR model reduces HMSE when compared to the GARCH model. For the DM. the reduction is from 0.9049 to 0.8107. for the JY from 0.9891 to 0.9290. and for the SF the reduction is from 0.4112 to 0.3504.

3.3.3 ARMA & Hetero CSR

The variance forecasts from the ARMA model and the Hetero CSR models are quite similar. The correlation of the two forecasts is 0.993. 0.989. and 0.998 for the DM. JY and SF. See Figure C.48 for graphs of the estimates.

HMSE Measures

The HMSE measures provide. on the surface. ambiguous results–but they are consistent. For the HMSE measure based on squared returns. the Hetero-CSR model is favored by the DM and JY. For the SF. the ARMA has a lower HMSE. but not by much (see Table C.71. For the HMSE measure based on the cumulative squared returns (see Table C.72). the ARMA model is favored by the DM and the SF. Since both models produce forecasts so similar to each other. we conclude there is not much difference between the two. We recommend the use of the Hetero-CSR model as it does the best reducing HMSE based on squared return.

3.3.4 Choosing Between Models

Finally, in determining whether one should use a GARCH, GARCH-CSR or an Hetero-CSR model, the results are very clear if we focus on the results from the HMSE measures. First we discuss other measures.

3.3.5 Normality, R^2 and Persistence

In terms of explaining the excess kurtosis of daily returns, model which is the best at explaining the excess kurtosis is the GARCH-CSR model (see Table C.73.) This makes sense since the this is one of the by products of maximizing the log-likelihood function. The ARMA and Hetero-CSR models do focus on trying to be close to the cumulative squared return for the day, and thus one should not expect them to explain excess kurtosis. In hindsight, explaining the excess kurtosis is not a useful method to evaluate forecasts.

In terms of forecasting next day variance as measured by the cumulative squared return, moving from a GARCH model to one with cumulative squared returns (GARCH-CSR) increases the R^2, but not dramatically. In terms of an R^2, all four models produce similar degrees of "fit" (see Table C.73.)

As for the persistence shocks to the variance process, all four models are close to the estimates one obtains from using Jorion's (1995) parameter estimates. While there is some variation in the estimates, overall no discernible effect is evident. It appears as if all models have similar lag effects. This should be reassuring since we would expect the gain from using intraday data to be one of efficiency. As such, it is unlikely that there would be any changes in persistence for each currency. The results confirm this view.

Statistical Measures

Comparing the HMSE (based on squared returns) of the various models: GARCH, GARCH-CSR and Hetero-CSR, we see that the Hetero-CSR model performs the best (see Table C.75.) If we use the GARCH model forecasts as our baseline HMSE, then when we use the Hetero-CSR model we reduce HMSE by 38% for the DM, 29% for the JY, and 15% for the SF. However, given that we have shown in Chapter 2 that the cumulative squared return is a good measure of volatility, we can ask how close we are (in HMSE terms) to reaching the best volatility forecasts we have–those that are based on the day's data rather than a forecast. Using the fact that the HMSE for the DM, JY and SF is 1.3393, 0.8313, and 1.2023 respectively, we can calculate the how close we have gotten to those lower limits. It turns out that we do quite well; for the DM we reach 71% of the lower limit, for the JY we reach 35% and for the SF we reach 25%. Considering we have no macroeconomic or structural variables in the model, the results are very encouraging. Unlike the results obtained when we run a regression, we find that there has been much more explained than previously thought.

3.4 Conclusions

While the short data set used in this study occasions the need for further research, the results are promising. We find that using intraday variance estimates improves variance forecasts. The amount of volatility that is predicted is quite high, from 75% to 25%. It is not clear whether the results will hold with longer data sets–however we expect them to do so. The gains from creating a richer model will also have to be determined later. But given the stylized fact that announcements play a major role in volatility we should expect that incorporating them will be useful.

Chapter 4

A Monte Carlo Investigation of the Heteroskedastic Cumulative Squared Returns Volatility Model

4.1 Introduction to Simulations

In order to determine that the results obtained in Chapter 3 are significant and not a result of the data analyzed, we do some simulations. Recall that the analysis found that there were gains from using intraday data to forecast daily volatility. The Heteroskedastic Cumulative Squared Returns volatility model (Hetero-CSR) reduced Heteroskedastic Mean Squared Error for the Deutschemark, Japanese Yen, and Swiss Franc.

The simulations of the Hetero-CSR model support the empirical results obtained and previously discussed in Chapter 3.

The rest of this chapter is organized as follows: Section 4.2 discusses simulations in Bollerslev and Ghysels (1996), Section 4.3 reviews the empirical model. Section 4.4 discusses the simulation assumptions, Section 4.5 discusses the evaluation process for the simulations, Section 4.6 discusses the results, Section 4.7 concludes.

4.2 Simulations in Bollerslev and Ghysels (1996)

To help us with our simulations, we turn to the GARCH literature to provide us an illustration. In this case, we turn to Bollerslev and Ghysels (1996). In their paper, they present the P-GARCH model and also do simulations. Using their approach on a GARCH model rather than the P-GARCH model, the authors simulated

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i} &= \omega_0 + \alpha_0 \epsilon^2_{t,i-1} + \beta_0 \sigma^2_{t,i-1} \end{aligned}$$

by establishing the data generating process this way

1. Initial Conditions
   (a) Set $\sigma_0^2 = \omega/(1 - \alpha - \beta)$
   (b) Set $\sigma_1 = \sigma_0$
   (c) Draw η_1, a normal variate
   (d) Set $\epsilon_1 = \eta_1 \sigma_1$
2. Do Loop
   (a) Set $\sigma_2^2 = \omega + \alpha\epsilon_1^2 + \beta\sigma_1^2$
   (b) Draw η_2, a normal variate
   (c) Set $\epsilon_2 = \eta_2 \sigma_2$
3. Continue
   (a) Repeat Step 2 until end

4.3 Empirical Model

The model estimated that we want to simulate is

Model 5: Conditionally Heteroskedastic CSR Variance

$$\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_t^2 &= \left(\sum_{i=1}^{n} R_{t,i}^2\right) \\
\sigma_{t|t-1}^2 &= \omega + \alpha\sigma_{t-1}^2 + \beta\sigma_{t-1|t-2}^2 \\
\log(L) &= -\frac{1}{2}\log(\pi) - \sum_t^T \log(\sigma_{t|t-1}^2) - \sum_t^T \sigma_t^2/\sigma_{t|t-1}^2
\end{aligned}$$

Here the daily returns for day t are composed of i intraday returns. With the assumption that the expectation of an intraday return being zero, i.e. $E[R_{t,i}] = 0$, and that the covariance between any intraday return being zero, then we can use intraday returns to compute a daily variance estimate. Both of these assumptions are reasonable given the efficiency of the foreign exchange markets.

4.4 Assumptions for Simulations

The following assumptions are used in the simulations:

1. We use the Deutschemark intraday seasonality as the seasonal pattern for the intraday returns. The Deutschemark's intraday seasonal volatility pattern is similar to that of the Japanese Yen and the Swiss Franc, and so with no loss of generality it is used as the basis. No attempt is made to study other seasonal patterns.

2. There will be two sets of simulations, each run with two different sample sizes:

Simulations 1: The parameter are set to be the following

$$
\begin{aligned}
\mu &= 0 \\
\omega &= 0.0153 \\
\alpha &= 0.1513 \\
\beta &= 0.8086
\end{aligned}
$$

and recall that with a GARCH model we have an unconditional variance of $\sigma_0^2 = \omega/(1-\alpha-\beta)$. In our case it is approximately 0.4. These parameters are similar to empirical estimates from the Deutschemark empirical results. The α parameter has been raised by 0.03 and the β parameter has been lowered by 0.03. This was done to make the α parameter farther from 0. but does not change the unconditional variance.

Simulations 2: The parameter estimates are

$$
\begin{aligned}
\mu &= 0 \\
\omega &= 0.16 \\
\alpha &= 0.20 \\
\beta &= 0.40
\end{aligned}
$$

Again, the unconditional variance is 0.4. These parameters are chosen to be far away from I-GARCH model. (Recall that in the I-GARCH model, $\alpha + \beta = 1$.) Since the parameters from Simulations 1 have the condition that

$\alpha + \beta = 0.9599$, which may be close to I-GARCH, we choose our parameters so that $\alpha + \beta << 1$. This allows us to evaluate how the models perform under conditions far from I-GARCH. Again we have α less than β since this is what is commonly found in empirical work.

3. Replications are 1000. This number is chosen to be large enough to obtain asymptotic results, but small enough to be manageable.

4. Simulations 1 and 2 will be run with 262 observations and 262×5 observations: A total of four treatments will be conducted. The first sample size is exactly the same number of observations in the empirical analysis in Chapter 3. The second sample size is chosen to be consistent with the sample sizes typically found in empirical work with daily returns.

5. Each replication resets the random number generator seed. The random number seed sequence for each treatment will be the same.

6. The log-likelihood function used does not estimate a μ parameter as it is set to zero in the simulations and it is not of interest. As a result, from each replication there will three parameter estimates, $\hat{\omega}$, $\hat{\alpha}$, and $\hat{\beta}$.

4.4.1 Simulating Data

Variables

The following variables are needed:

1. Let $v_1(i)$ and $v_2(i)$, both vectors (96×1) in length.

2. Let $\gamma(i)$, a vector (96×1) which stores intraday seasonality. (In our case, this stores the Deutschemark seasonal pattern of standard deviations for each fifteen minute part of the day.)

3. Let Ret_t stores the daily return, CSR_t stores the cumulative squared return, and Sigma_t^2 store daily variance; both are (1040×1) vectors.

4. Let k be the ratio of the day's variance over the historical variance, σ_0^2.

Simulation Process

The simulations are then created in the following manner:

1. Initial Conditions $(t = 1)$

   (a) Draw 96 normal variates and fill $v_1(i)$. This creates the randomness for the day.

   (b) Set $v_2(i) = v_1(i)\gamma(i)$ for $i = 1$ to 96. This scales intraday returns with the seasonal pattern.

   (c) Set $\text{Ret}_{t=1} = \sum_i v_2(i)$. We sum up the intraday returns to get a daily return.

   (d) Set $\text{CSR}_{t=1} = \sum_i v_2^2(i)$. We sum up the squared intraday returns to get the CSR variance estimate for the day.

   (e) Set $\text{Sigma}_{t=1}^2 = \sigma_0^2$. We set the first day's true variance to the unconditional variance.

2. Remaining Data $(t > 1)$

   (a) Draw 96 normal variates and re-fill $v_1(i)$.

(b) Reset $\text{Sigma}^2_{t=j} = \omega + \alpha\text{CSR}_{j-1} + \beta\text{Sigma}^2_{j-1}$.

(c) Set $k = \text{Sigma}^2{}_{t=j}/\sigma_0^2$.

(d) Reset $v_2(i) = \sqrt{k}\,(v_1(i)\gamma(i))$ for $i = 1$ to 96.

(e) Set $\text{Ret}_{t=j} = \sum_i v_2(i)$.

(f) Set $\text{CSR}_{t=j} = \sum_i v_2^2(i)$.

3. Repeat Until $j = T$, where T is the sample size and is either 262 or 262×5.

4.5 Evaluating Simulations

The data were simulated and then three models were estimated on the data. The first model was a standard GARCH(1,1) model. The second model was the GARCH-CSR model. The last model was the Heteroskedastic CSR model (Hetero-CSR). Note that the GARCH-CSR model is very similar to the Hetero-CSR model. The difference in the two models is in the log-likelihood function. Recall the GARCH-CSR model is:

Model 3: GARCH-CSR Variance

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_t^2 &= \left(\sum_{i=1}^{n} R_{t.i}^2\right) \\
\sigma^2_{t|t-1} &= \omega + \alpha\sigma^2_{t-1} + \beta\sigma^2_{t-1|t-2} \\
\log(L) &= -\frac{1}{2}\log(\pi) - \sum_t^T \log(\sigma^2_{t|t-1}) - \sum_t^T \epsilon_t^2/\sigma^2_{t|t-1}
\end{aligned}
$$

In terms of evaluating the advantage of one model over the other, we will use a mean square error measure adjusted for heteroskedasticity (HMSE). We will use two variants, one that is based on squared daily return, the other is based on cumulative squared returns.

Recall that mean square error is given by the following formula.

$$\mathrm{MSE} = \frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}}\left[\epsilon_{\mathrm{t}}^{2} - \hat{\sigma}_{\mathrm{t}}^{2}\right]^{2}$$

For the HMSE version of MSE, we are interested how close or volatility forecast is to our actual volatility in terms of a percentage. Thus we modify the MSE formula to get:

$$\mathrm{HMSE} = \frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}}\left[\frac{\epsilon_{\mathrm{t}}^{2}}{\hat{\sigma}_{\mathrm{t}}^{2}} - 1\right]^{2}$$

Notice that the basic HMSE is actually based on the squared returns since $R_t^2 = \epsilon_t^2$. Therefore, we can call the basic HMSE, HMSE-Squared Return.

The second version of HMSE is called HMSE-CSR. The reasoning behind this version of the HMSE is that while on average $E[\epsilon_t^2] = \sigma_t^2$, it is not an efficient estimate of daily volatility when we have intraday data (see Chapter 2). As a result, given we have good estimates of σ_t^2 from CSR, we create a HMSE based on our CSR estimates of volatility. We just substitute CSR_t^2 for ϵ_t^2.

$$\mathrm{HMSE} - \mathrm{CSR} = \frac{1}{\mathrm{T}}\sum_{\mathrm{t}=1}^{\mathrm{T}}\left[\frac{\mathrm{CSR}_{\mathrm{t}}^{2}}{\hat{\sigma}_{\mathrm{t}}^{2}} - 1\right]^{2}$$

4.6 Results

The results of the simulations are consistent with expectations. We find the following:

1. Simulated data lacks excess kurtosis. (Results neither reported nor tabulated.) Since the daily simulated returns appear to be normal, then a constant variance model is going to do well. We would not, for example, expect that an excess kurtosis model like a GARCH model, to be favored statistically. Therefore, statistical measures which are related to log-likelihood approach, such as the HMSE-Sq Ret, will not be useful as a measure of forecast performance. However, the HMSE-Sq Ret results are presented since there was no reason a priori to know that the simulations would not produce daily returns without excess kurtosis.

2. GARCH model estimates on both Simulation 1 and Simulation 2 fail a large number of times. The parameter estimates often have α equal to zero, or α and β both equal to zero. (The frequency of such occurances is neither reported nor tabulated.) Thus the data generating process does not appear to be consistent with daily empirical data.

3. Estimating a Hetero-CSR model on the simulated data recovers the parameters better than the GARCH or GARCH-CSR models do (See Tables D.76 and D.77.)

4. A larger sample produces better forecasts for GARCH and GARCH-CSR. For Simulation 1, HMSE-CSR falls from 0.0589 to 0.0506 for GARCH forecasts, and for GARCH-CSR forecasts it falls from 0.0521 to 0.0418 (see Table D.76.)

For Simulation 2, HMSE-CSR falls from 0.0533 to 0.0429 for GARCH forecasts, and for GARCH-CSR forecasts it falls from 0.0514 to 0.0414 (see Table D.77.)

5. Larger sample size does not lower average HMSE-CSR for Hetero-CSR model. For Simulations 1, HMSE-CSR is 0.0374 for a sample size of 262, while it is 0.0378 for a sample size 262×5 (see Table D.76). For Simulations 2, HMSE-CSR is 0.0372 for sample size of 262, while it is 0.0377 for a sample size of 262×5 (see Table D.77.) In both cases, the HMSE-CSR increases, but only slightly.

6. A larger sample makes it more likely that GARCH-CSR model produces better forecasts than a GARCH model (see Table D.78.) For Simulations 1, 775 times out of 1000, the GARCH-CSR model, with a sample size of 262, produced better forecasts than the GARCH forecasts when evaluated with HMSE-CSR. When the sample size increased by five times, the GARCH-CSR had better forecasts 981 times. For Simulations 2, the number of times the GARCH-CSR produced better forecasts increased from 594 to 782.

7. A larger sample makes it more likely that Hetero-CSR model produces better forecasts than a GARCH model (see Table D.79.) For Simulations 1, 996 times out of 1000, the GARCH-CSR model, with a sample size of 262, produced better forecasts than the GARCH forecasts when evaluated with HMSE-CSR. When the sample size increased by five times, the GARCH-CSR had better forecasts 1000 out of the 1000 times. For Simulations 2, the number of times the GARCH-CSR produced better forecasts increased from

982 to 1000. Clearly the Hetero-CSR model does a much better job than the GARCH model at forecasting.

8. Now for comparing GARCH-CSR and Hetero-CSR forecasts. We find that the majority of the time the Hetero-CSR model produced better forecasts (again using the HMSE-CSR metric.) As shown in Table D.80, Simulations 1, 928 times out of 1000, the Hetero-CSR model, with a sample size of 262, produced better forecasts than the GARCH-CSR forecasts. When the sample size increased by five times, the Hetero-CSR model had better forecasts 828 out of the 1000 times. For Simulations 2, the number of times the Hetero-CSR produced better forecasts decreased from 941 to 819.

9. Although the result that the Hetero-CSR model appears to do less well at forecasting when compared to GARCH-CSR model is unexpected, it is not explainable. As denoted on Tables D.76 and D.77, the HMSE-CSR for GARCH-CSR model falls when sample size is increased, but it stays the same for the Hetero-CSR forecasts. As a result, it is not that the Hetero-CSR forecasts are worse, it is just that the GARCH-CSR forecasts get closer to the Hetero-CSR forecasts.

10. We can also look at percentage improvement in forecasts (see Table D.81.) If we assume that the GARCH model forecast is our baseline forecast, then the percentage reduction in forecast HMSE-CSR from using the GARCH-CSR model increases with sample size. For Simulations 1, the percentage improvement increases from 9.44 to 17.00. For Simulations 2, it increases from 1.25 to 3.38.

11. Again, given that the Hetero-CSR forecasts do not improve with increased sample size but the GARCH and GARCH-CSR forecasts do, then using either the GARCH or GARCH-CSR model forecasts as a baseline yields leads us to conclude that benefit of Hetero-CSR forecasts decrease with sample size.

12. In all four treatments, the Hetero-CSR estimates are better than the GARCH-CSR estimates (see Figures D.59 to D.70.) The Gaussian kernel density for every treatment is more concentrated on the actual parameter for the Hetero-CSR estimates.

13. Finally, in both simulations, the Gaussian kernel density estimates of the parameters appear to converge to the true parameter estimates as the number of observations in the sample increases. See Figures D.71 to D.82.

4.6.1 Possible Source of Lack of Excess Kurtosis in Simulations

The likely explanation for lack of excess kurtosis in the simulations is due to the large number of intraday returns that are used to create the CSR daily variance estimate. The difference between the CSR variable and the Sigma variable is probably small due to a law of large numbers effect from summing up 96 returns.

One approach would be to create a richer model. However, at this point this is unnecessary for the given problem at hand. We are interested in evaluating the model previously estimated in Chapter 3. The results here suggest that a more complex model may provide better forecasts. We leave this question for future research.

4.6.2 Final Remarks

Although the simulations fail in creating data with excess kurtosis, we can conclude a great deal. First, we know that the Hetero-CSR model fails to capture excess kurtosis in daily returns, and thus it should be possible to extend the model to capture this feature of the empirical data. This should lead to even better forecasts. Additionally, the benefit of intraday data is clearly demonstrated with the HMSE-CSR forecasts over GARCH forecasts. In small samples, the benefits of intraday data yield better forecasts since we get both better parameter estimates, but also better estimates of daily volatility. While in larger samples, GARCH model forecasts improve in quality, they are still worse than the Hetero-CSR forecasts. As a result of these simulations, we conclude that the results obtained with empirical data are significant.

Finally, it appears that intraday data may allow researchers to investigate whether volatility dynamics change over time. With daily data this has not been possible since estimation typically required all observations be grouped. As a result, sample sizes are five years or more. However, we may suspect that institutional or investor behaviors change. With more data than currently used here, it may be possible to study whether volatility regimes exist as it appears as if one year's worth of intraday data is superior to five years of daily data. The applications to investigate policy questions seems very promising. We leave such investigations to future researchers.

Appendix A: Tables and Figures for Chapter 1

Note: In the following tables, numbers is parentheses are heteroskedastic-consistent standard errors and numbers in brackets are t-ratios. In the cases where a (−) or a [−] is present, standard error estimates were not provided by the statistical software used for estimating the parameters.

Modeling Returns	DM	JY	SF
Min	-0.6495	-0.8571	-0.6863
Max	1.1646	0.9123	0.7814
Range	1.8141	1.7695	1.4677
Mean	-0.0002	-0.0003	-0.0004
Median	0	0	0
Variance	0.0037	0.0056	0.0050
Standard Deviation	0.0611	0.0745	0.0707
Mean Deviation	0.0391	0.0483	0.0456
Median Deviation	0.0265	0.0323	0.0338
Skewness	0.3700	0.1743	0.0570
Excess Kurtosis	19.190	12.264	8.787

Table A.1: Descriptive statistics for fifteen minute return modeling data set ($T \times n = 19,392$).

Forecast Returns	DM	JY	SF
Min	-0.8200	-0.5749	-0.9632
Max	1.2845	0.4623	0.6066
Range	2.1046	1.0373	1.5698
Mean	-0.0001	0.0002	-0.0003
Median	0	0	0
Variance	0.0057	0.0053	0.0063
Standard Deviation	0.0757	0.0726	0.0796
Mean Deviation	0.0462	0.0497	0.0501
Median Deviation	0.0291	0.0351	0.0373
Skewness	0.6135	0.0131	-0.3290
Excess Kurtosis	27.366	4.817	10.225

Table A.2: Descriptive statistics for fifteen minute return forecasting data set ($T \times n = 5,760$).

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i} &= \gamma(i) \\ \gamma(i) &= \mathrm{Exp}(c_0 + c_1 i + \cdots + c_9 i^9) \end{aligned}$$

Estimate	DM	JY	SF
c_0	-5.5210 [-93.409]	-5.1338 [-82.448]	-5.2099 [-105.627]
c_1	3.8641 [12.362]	2.4813 [6.835]	3.5814 [13.856]
c_2	7.9712 [6.145]	5.7983 [5.316]	7.9524 [8.206]
c_3	-22.0738 [-5.936]	-14.9576 [-3.831]	-22.0369 [-7.384]
c_4	-49.5225 [-7.670]	-34.1290 [-6.680]	-50.1452 [-10.205]
c_5	57.9574 [4.148]	43.7188 [3.391]	70.9155 [6.618]
c_6	74.3920 [6.797]	50.9557 [6.071]	74.9036 [8.799]
c_7	-68.2269 [-3.428]	-58.9793 [-3.577]	-98.9381 [-6.716]
c_8	-33.6061 [-5.727]	-23.0231 [-5.239]	-33.5627 [-7.203]
c_9	28.4792 [3.041]	27.7369 [3.871]	46.4780 [6.875]

Table A.3: Parameter estimates for exponential seasonality, ($T = 202, n = 96$).

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i} &= \omega + \alpha\epsilon^2_{t,i-1} + \beta\sigma^2_{t,i-1} \end{aligned}$$

Estimate	DM	JY	SF
ω	1.7831×10^{-4}	2.7745×10^{-4}	2.9422×10^{-4}
	(-)	(-)	(-)
	[-]	[-]	[-]
α	0.2521	0.1654	0.2548
	(0.0233)	(0.0157)	(0.0183)
	[10.816]	[10.519]	[13.947]
β	0.7385	0.7955	0.7199
	(0.0238)	(0.0186)	(0.0159)
	[31.041]	[42.872]	[45.264]

Table A.4: Parameter estimates for GARCH on raw data ($T \times n = 19,392$).

$$\begin{aligned} \tilde{R}_{t,i} &= \epsilon_{t,i} \\ \tilde{\sigma}^2_{t,i} &= \omega + \alpha\epsilon^2_{t,i-1} + \beta\tilde{\sigma}^2_{t,i-1} \end{aligned}$$

Estimate	DM	JY	SF
ω	0.07130	0.02985	0.09475
	(0.0164)	(0.0069)	(0.0947)
	[4.350]	[4.338]	[5.232]
α	0.1324	0.09029	0.1186
	(0.0171)	(0.0120)	(0.0157)
	[7.734]	[7.528]	[7.552]
β	0.7967	0.8826	0.7877
	(0.0327)	(0.0178)	(0.0312)
	[24.400]	[49.610]	[25.234]

Table A.5: Parameter estimates for GARCH on deseasoned by sample standard deviation. ($T \times n = 19,392$).

$$\tilde{R}_{t,i} = \epsilon_{t,i}$$
$$\tilde{\sigma}^2_{t,i} = \omega + \alpha\epsilon^2_{t,i-1} + \beta\tilde{\sigma}^2_{t,i-1}$$

Estimate	DM	JY	SF
ω	0.07135	0.02898	0.1132
	(0.0188)	(0.0068)	(0.0261)
	[3.802]	[4.237]	[4.346]
α	0.1282	0.09351	0.1234
	(0.0186)	(0.0131)	(0.0183)
	[6.883]	[7.142]	[6.749]
β	0.8062	0.8803	0.7669
	(0.0354)	(0.0185)	(0.0415)
	[22.778]	[47.471]	[18.458]

Table A.6: Parameter estimates for GARCH on deseasoned by exponential seasonality. ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= \omega + \alpha_1 \epsilon^2_{t,i-1} + \cdots + \alpha_{10} \epsilon^2_{t,i-10}
\end{aligned}
$$

Estimate	DM	JY	SF
ω	7.0739×10^{-4} [-]	1.5892×10^{-4} [14.633]	1.0923×10^{-4} [-]
α_1	0.3117 [9.278]	0.1934 [10.557]	0.2825 [11.650]
α_2	0.1541 [8.197]	0.1493 [7.429]	0.1542 [8.597]
α_3	0.0895 [4.414]	0.0984 [5.049]	0.1466 [5.016]
α_4	0.1139 [4.601]	0.0656 [4.483]	0.0782 [4.079]
α_5	0.0703 [3.377]	0.0735 [3.681]	0.0699 [3.819]
α_6	0.0606 [3.581]	0.0410 [3.219]	0.0391 [2.912]
α_7	0.0522 [3.247]	0.0399 [2.669]	0.0409 [3.384]
α_8	0.0173 [1.850]	0.0317 [2.664]	0.0323 [2.902]
α_9	0.0384 [1.763]	0.0235 [1.967]	0.0177 [1.586]
α_{10}	0.0520 [2.373]	0.0498 [3.410]	0.0428 [1.689]

Table A.7: Parameter estimates for ARCH(10) on raw data, ($T \times n = 19,392$).

$$\begin{aligned} \tilde{R}_{t,i} &= \epsilon_{t,i} \\ \tilde{\sigma}^2_{t,i} &= \omega + \alpha_1 \epsilon^2_{t,i-1} + \cdots + \alpha_{10} \epsilon^2_{t,i-10} \end{aligned}$$

Estimate	DM	JY	SF
ω	0.4111 [19.555]	0.3743 [17.373]	0.5023 [20.468]
α_1	0.1990 [10.373]	0.1581 [10.813]	0.1619 [10.712]
α_2	0.0946 [7.186]	0.1064 [7.443]	0.0997 [4.825]
α_3	0.0489 [4.139]	0.0695 [4.540]	0.0609 [4.579]
α_4	0.0489 [3.444]	0.0521 [4.270]	0.0368 [2.846]
α_5	0.0498 [3.952]	0.0557 [3.794]	0.0412 [3.889]
α_6	0.0361 [3.414]	0.0339 [3.210]	0.0251 [2.518]
α_7	0.0302 [2.812]	0.0346 [2.821]	0.0187 [2.425]
α_8	0.0143 [1.678]	0.0317 [2.905]	0.0132 [1.582]
α_9	0.0482 [2.336]	0.0254 [2.076]	0.0273 [2.205]
α_{10}	0.0408 [2.249]	0.0717 [3.558]	0.0298 [2.207]

Table A.8: Parameter estimates for ARCH(10) on returns deseasoned by sample standard deviation, ($T \times n = 19,392$).

$$\bar{R}_{t,i} = \epsilon_{t,i}$$
$$\bar{\sigma}^2_{t,i} = \omega + \alpha_1\epsilon^2_{t,i-1} + \cdots + \alpha_{10}\epsilon^2_{t,i-10}$$

Estimate	DM	JY	SF
ω	0.4071	0.3544	0.5139
	[16.081]	[17.023]	[19.831]
α_1	0.2114	0.1638	0.1753
	[9.132]	[10.205]	[10.333]
α_2	0.0911	0.1130	0.0953
	[6.797]	[7.049]	[6.059]
α_3	0.0469	0.0697	0.0605
	[3.571]	[4.163]	[4.173]
α_4	0.0485	0.0583	0.0307
	[3.273]	[4.255]	[2.325]
α_5	0.0590	0.0573	0.0297
	[3.022]	[3.609]	[2.881]
α_6	0.0280	0.0305	0.0224
	[2.453]	[2.967]	[2.256]
α_7	0.0315	0.0364	0.0203
	[2.671]	[2.541]	[2.392]
α_8	0.0101	0.0325	0.0124
	[1.310]	[2.596]	[1.272]
α_9	0.0532	0.0262	0.0221
	[2.055]	[2.069]	[1.897]
α_{10}	0.0475	0.0807	0.0360
	[2.475]	[4.212]	[2.350]

Table A.9: Parameter estimates for ARCH(10) on returns deseasoned exponential seasonality. ($T \times n = 19.392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_0 d_0 + \alpha_1 d_1)\,\epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\,\sigma^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York open} \\ 1 & \text{New York closed} \end{cases}
\end{aligned}
$$

Estimate	DM	JY	SF
ω_0	1.7239×10^{-4} (-) [-]	2.2439×10^{-4} (-) [-]	3.8425×10^{-4} (-) [-]
ω_1	2.4887×10^{-3} (0.0006) [3.964]	1.2407×10^{-3} (0.0003) [3.679]	3.5489×10^{-3} (0.0008) [4.592]
α_0	0.2208 (0.0355) [6.225]	0.1483 (0.0184) [8.068]	0.2462 (0.0242) [10.166]
α_1	0.2407 (0.0481) [5.008]	0.2076 (0.0287) [7.227]	0.2042 (0.0314) [6.493]
β_0	0.7352 (0.0552) [14.071]	0.8119 (0.0269) [30.162]	0.6520 (0.0317) [20.597]
β_1	0.4454 (0.0893) [4.985]	0.6459 (0.0559) [11.552]	0.4835 (0.0797) [6.067]

Table A.10: Parameter estimates for P-GARCH ($T \times n = 19,392$).

$$
\begin{aligned}
\tilde{R}_{t,i} &= \epsilon_{t,i} \\
\tilde{\sigma}^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_0 d_0 + \alpha_1 d_1)\, \epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\, \sigma^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York open} \\ 1 & \text{New York closed} \end{cases}
\end{aligned}
$$

Estimate	DM	JY	SF
ω_0	0.07403	0.03188	0.09599
	(0.0199)	(0.0064)	(0.0231)
	[3.714]	[4.951]	[4.150]
ω_1	0.06925	0.01854	0.09509
	(0.0225)	(0.0096)	(0.0257)
	[3.073]	[1.938]	[3.693]
α_0	0.1272	0.08001	0.1138
	(0.0208)	(0.0106)	(0.0199)
	[6.110]	[7.532]	[5.723]
α_1	0.1359	0.1153	0.1348
	(0.0207)	(0.0185)	(0.0201)
	[6.551]	[6.215]	[6.717]
β_0	0.8024	0.8902	0.7925
	(0.0387)	(0.0160)	(0.0400)
	[20.708]	[55.635]	[19.816]
β_1	0.8004	0.8681	0.7771
	(0.0379)	(0.0238)	(0.0366)
	[21.118]	[36.551]	[21.252]

Table A.11: Parameter estimates for P-GARCH on returns deseasoned by sample standard deviation, $(T \times n = 19,392)$.

$$\begin{aligned}
\tilde{R}_{t,i} &= \epsilon_{t,i} \\
\tilde{\sigma}^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_0 d_0 + \alpha_1 d_1)\,\epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\,\sigma^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York open} \\ 1 & \text{New York closed} \end{cases}
\end{aligned}$$

Estimate	DM	JY	SF
ω_0	0.06820	0.03023	0.10694
	(0.0221)	(0.0064)	(0.0313)
	[3.092]	[4.757]	[3.414]
ω_1	0.06676	0.02139	0.10741
	(0.0275)	(0.0121)	(0.0335)
	[2.426]	[1.760]	[3.206]
α_0	0.1176	0.08354	0.1141
	(0.0230)	(0.0120)	(0.0237)
	[5.113]	[6.966]	[4.818]
α_1	0.1486	0.1173	0.1412
	(0.0266)	(0.0217)	(0.0219)
	[5.584]	[5.397]	[6.462]
β_0	0.8197	0.8896	0.7817
	(0.0428)	(0.0166)	(0.0522)
	[19.152]	[53.736]	[14.964]
β_1	0.7906	0.8614	0.7565
	(0.0458)	(0.0306)	(0.0461)
	[17.252]	[28.162]	[16.417]

Table A.12: Parameter estimates for P-GARCH on returns deseasoned by exponential seasonality, ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_{1,0} d_0 + \alpha_{1,1} d_1)\,\epsilon^2_{t,i-1} + \\
&\quad \cdots + (\alpha_{10,0} d_0 + \alpha_{10,1} d_1)\,\epsilon^2_{t,i-10} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= 1 - d_0
\end{aligned}
$$

Estimate	DM		JY		SF	
	$j=0$	$j=1$	$j=0$	$j=1$	$j=0$	$j=1$
ω_j	0.0007	0.0035	0.0015	0.0035	0.0011	0.0057
	[-]	[11.400]	[14.545]	[10.280]	[-]	[14.123]
$\alpha_{1,j}$	0.2854	0.2458	0.1767	0.1953	0.2518	0.1711
	[7.491]	[4.314]	[8.915]	[6.118]	[9.804]	[4.581]
$\alpha_{2,j}$	0.1333	0.0745	0.1108	0.1569	0.1226	0.1110
	[7.152]	[2.950]	[5.609]	[4.801]	[6.323]	[4.247]
$\alpha_{3,j}$	0.0455	0.0714	0.0879	0.0644	0.1220	0.0821
	[2.256]	[2.648]	[4.297]	[2.013]	[3.234]	[3.541]
$\alpha_{4,j}$	0.0911	0.0301	0.0732	0.0384	0.0475	0.0252
	[4.253]	[1.379]	[4.113]	[2.161]	[2.512]	[1.192]
$\alpha_{5,j}$	0.0637	0.0395	0.0698	0.0524	0.0616	0.0413
	[2.675]	[2.141]	[2.957]	[2.020]	[1.531]	[2.344]
$\alpha_{6,j}$	0.0420	0.0225	0.0329	0.0182	0.0183	0.0315
	[2.342]	[1.458]	[2.713]	[1.176]	[1.907]	[1.961]
$\alpha_{7,j}$	0.0353	0.0244	0.0444	0.0065	0.0264	0.0299
	[2.191]	[2.062]	[2.522]	[0.502]	[1.292]	[1.965]
$\alpha_{8,j}$	0.0027	0.0212	0.0243	0.0353	0.0176	0.0065
	[0.488]	[1.764]	[1.925]	[1.887]	[1.292]	[0.615]
$\alpha_{9,j}$	0.0716	0.0003	0.0065	0.0407	0.0116	0.0043
	[2.072]	[0.041]	[0.075]	[1.628]	[0.862]	[0.577]
$\alpha_{10,j}$	0.0578	0.0472	0.0792	0.0182	0.0515	0.0136
	[1.764]	[1.989]	[3.465]	[1.012]	[1.252]	[1.184]

Table A.13: Parameter estimates for PARCH(10), ($T \times n = 19,392$).

$$
\begin{aligned}
\bar{R}_{t,i} &= \epsilon_{t,i} \\
\bar{\sigma}^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_{1,0} d_0 + \alpha_{1,1} d_1)\,\epsilon^2_{t,i-1} + \\
&\quad \cdots + (\alpha_{10,0} d_0 + \alpha_{10,1} d_1)\,\epsilon^2_{t,i-10} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= 1 - d_0
\end{aligned}
$$

Estimate	DM		JY		SF	
	$j=0$	$j=1$	$j=0$	$j=1$	$j=0$	$j=1$
ω_j	0.4081	0.4212	0.3979	0.2968	0.4987	0.5146
	[17.827]	[11.980]	[16.156]	[9.752]	[18.086]	[13.519]
$\alpha_{1,j}$	0.2019	0.2037	0.1447	0.1919	0.1659	0.1459
	[9.220]	[5.250]	[9.208]	[6.240]	[9.809]	[4.654]
$\alpha_{2,j}$	0.1058	0.0619	0.0889	0.1562	0.1011	0.0940
	[6.678]	[2.880]	[5.932]	[4.994]	[3.892]	[3.695]
$\alpha_{3,j}$	0.0439	0.0579	0.0768	0.0470	0.0625	0.0515
	[3.477]	[2.302]	[4.420]	[1.544]	[3.968]	[2.545]
$\alpha_{4,j}$	0.0535	0.0346	0.0555	0.0403	0.0425	0.0119
	[3.440]	[1.220]	[3.890]	[2.082]	[2.766]	[0.643]
$\alpha_{5,j}$	0.0515	0.0410	0.0564	0.0625	0.0375	0.0543
	[3.470]	[2.230]	[3.451]	[2.094]	[3.082]	[2.712]
$\alpha_{6,j}$	0.0349	0.0388	0.0354	0.0319	0.0246	0.0239
	[2.961]	[1.799]	[2.983]	[1.506]	[2.254]	[1.187]
$\alpha_{7,j}$	0.0305	0.0308	0.0418	0.0136	0.0144	0.0353
	[2.446]	[1.753]	[2.947]	[0.817]	[1.853]	[1.949]
$\alpha_{8,j}$	0.0095	0.0343	0.0269	0.0598	0.0105	0.0221
	[0.965]	[1.947]	[2.293]	[2.383]	[1.112]	[1.370]
$\alpha_{9,j}$	0.0549	0.0212	0.0149	0.0629	0.0275	0.0326
	[2.268]	[1.238]	[1.288]	[2.205]	[1.845]	[1.584]
$\alpha_{10,j}$	0.0266	0.0821	0.0762	0.0462	0.0278	0.0420
	[1.478]	[2.005]	[3.233]	[1.502]	[1.765]	[1.916]

Table A.14: Parameter estimates for PARCH(10) deseasoned with sample variance seasonality, ($T \times n = 19,392$).

$$\begin{aligned}
\tilde{R}_{t,i} &= \epsilon_{t,i} \\
\tilde{\sigma}^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1) + (\alpha_{1,0} d_0 + \alpha_{1,1} d_1)\, \epsilon^2_{t,i-1} + \\
&\quad \cdots + (\alpha_{10,0} d_0 + \alpha_{10,1} d_1)\, \epsilon^2_{t,i-10} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= 1 - d_0
\end{aligned}$$

Estimate	DM		JY		SF	
	$j=0$	$j=1$	$j=0$	$j=1$	$j=0$	$j=1$
ω_j	0.4081 [14.458]	0.4007 [10.735]	0.3636 [15.753]	0.3184 [9.880]	0.5091 [17.488]	0.5296 [13.739]
$\alpha_{1,j}$	0.2140 [8.776]	0.2154 [4.052]	0.1578 [8.682]	0.1801 [6.091]	0.1834 [9.503]	0.1465 [4.309]
$\alpha_{2,j}$	0.0999 [6.342]	0.0640 [2.639]	0.0985 [5.573]	0.1459 [4.764]	0.0942 [4.862]	0.0993 [3.932]
$\alpha_{3,j}$	0.0362 [2.804]	0.0738 [2.483]	0.0786 [4.022]	0.0417 [1.320]	0.0610 [3.602]	0.0571 [2.620]
$\alpha_{4,j}$	0.0587 [3.433]	0.0167 [0.958]	0.0655 [4.033]	0.0389 [2.010]	0.0380 [2.342]	0.0007 [0.047]
$\alpha_{5,j}$	0.0595 [2.688]	0.0491 [2.077]	0.0601 [3.313]	0.0538 [1.986]	0.0260 [2.200]	0.0438 [2.323]
$\alpha_{6,j}$	0.0291 [2.281]	0.0202 [1.075]	0.0334 [2.835]	0.0229 [1.251]	0.0209 [1.994]	0.0272 [1.200]
$\alpha_{7,j}$	0.0326 [2.370]	0.0302 [1.586]	0.0447 [2.708]	0.0123 [0.767]	0.0181 [2.009]	0.0279 [1.572]
$\alpha_{8,j}$	0.0048 [0.708]	0.0453 [1.828]	0.0294 [2.055]	0.0523 [2.251]	0.0099 [0.840]	0.0191 [1.162]
$\alpha_{9,j}$	0.0610 [1.964]	0.0205 [1.276]	0.0138 [1.349]	0.0708 [2.078]	0.0221 [1.564]	0.0267 [1.353]
$\alpha_{10,j}$	0.0375 [1.652]	0.0785 [2.241]	0.0895 [3.824]	0.0503 [1.647]	0.0356 [1.951]	0.0424 [1.775]

Table A.15: Parameter estimates for PARCH(10) deseasoned with exponential seasonality, $(T \times n = 19,392)$.

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= \mathrm{Exp}(c_0 + c_1 i + \cdots + c_n i^n) + \\
&\quad + \alpha_0 \epsilon^2_{t,i-1} + \cdots + \alpha_{10} \epsilon^2_{t,i-10}
\end{aligned}
$$

Estimate	DM	JY	SF
c_0	-6.3233 [-58.099]	-6.0113 [-63.564]	-5.7920 [-58.843]
c_1	5.9681 [8.287]	4.0708 [7.667]	4.9549 [10.520]
c_2	9.7621 [4.822]	6.8001 [3.815]	9.2353 [5.657]
c_3	-53.1047 [-5.273]	-40.9258 [-6.544]	-41.9588 [-6.840]
c_4	-68.0510 [-6.565]	-47.0271 [-5.271]	-59.7734 [-7.273]
c_5	162.2698 [4.219]	131.4393 [5.803]	139.4167 [5.993]
c_6	107.1158 [6.149]	72.6594 [4.913]	88.5169 [6.258]
c_7	-197.6123 [-3.655]	-165.8812 [-5.317]	-186.3538 [-5.598]
c_8	-49.7033 [-5.535]	-32.5576 [-4.240]	-38.7135 [-5.021]
c_9	82.4120 [3.298]	71.0907 [4.943]	84.0627 [5.277]

Table A.16: Seasonality parameter estimates for the joint exponential seasonality and conditional heteroskedasticity model. ($T = 202, n = 96$).

$$\begin{aligned} R_{t,i} &= \epsilon_{t,i} \\ \sigma^2_{t,i} &= \mathrm{Exp}(c_0 + c_1 i + \cdots + c_n i^n) + \\ & \quad +\alpha_0 \epsilon^2_{t,i-1} + \cdots + \alpha_{10}\epsilon^2_{t,i-10} \end{aligned}$$

Estimate	DM	JY	SF
α_1	0.2029 [9.444]	0.1610 [10.216]	0.1691 [10.567]
α_2	0.0881 [6.626]	0.1168 [7.321]	0.0962 [6.003]
α_3	0.0418 [3.814]	0.0691 [4.305]	0.0562 [4.275]
α_4	0.0416 [3.131]	0.0565 [4.522]	0.0356 [2.689]
α_5	0.0597 [3.301]	0.0549 [3.935]	0.0317 [3.149]
α_6	0.0298 [2.630]	0.0295 [3.156]	0.0226 [2.324]
α_7	0.0299 [3.217]	0.0326 [2.346]	0.0161 [2.085]
α_8	0.0161 [2.139]	0.0341 [2.731]	0.0113 [1.403]
α_9	0.0768 [2.259]	0.0252 [2.400]	0.0248 [2.629]
α_{10}	0.0469 [2.711]	0.0625 [4.192]	0.0197 [1.924]

Table A.17: Conditional heteroskedasticity parameter estimates for the joint exponential seasonality and conditional heteroskedasticity model, ($T = 202, n = 96$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= \omega\gamma(i) + \alpha\epsilon^2_{t,i-1} + \beta\sigma^2_{t,i-1} \\
\gamma(i) &= \text{Sample Variance}
\end{aligned}
$$

Estimate	DM	JY	SF
ω	0.3244	0.1187	0.4497
	(0.0641)	(0.0295)	(0.0667)
	[5.065]	[4.031]	[6.737]
α	0.2428	0.1832	0.2225
	(0.0202)	(0.0234)	(0.0144)
	[12.018]	[7.827]	[15.455]
β	0.4339	0.7030	0.3284
	(0.0722)	(0.0507)	(0.0656)
	[6.012]	[13.862]	[5.007]

Table A.18: Parameter estimates for conditionally periodic GARCH with sample variance seasonality, ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= \omega\gamma(i) + \alpha\epsilon^2_{t,i-1} + \beta\sigma^2_{t,i-1} \\
\gamma(i) &= \text{Exp}(c_0 + c_1 i + \cdots + c_n i^n)
\end{aligned}
$$

Estimate	DM	JY	SF
ω	0.2167	0.1376	0.3627
	(0.0289)	(0.0327)	(0.0360)
	[7.492]	[4.210]	[10.065]
α	0.2094	0.1716	0.1974
	(0.0185)	(0.0198)	(0.0138)
	[11.323]	[8.667]	[14.302]
β	0.5802	0.7602	0.5282
	(0.0383)	(0.0357)	(0.0289)
	[15.163]	[21.318]	[18.304]

Table A.19: Parameter estimates for conditionally periodic GARCH with exponential seasonality, ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1)\,\gamma(i) + (\alpha_0 d_0 + \alpha_1 d_1)\,\epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\,\beta^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York open} \\ 1 & \text{New York closed} \end{cases} \\
\gamma(i) &= \text{Sample Variance}
\end{aligned}
$$

Estimate	DM	JY	SF
ω_0	0.3308	0.1041	0.4608
	(0.0701)	(0.0436)	(0.0690)
	[4.717]	[2.388]	[6.683]
ω_1	0.2847	0.1216	0.3347
	(0.0937)	(0.0297)	(0.0996)
	[3.038]	[4.098]	[3.359]
α_0	0.2504	0.1667	0.2264
	(0.0209)	(0.0330)	(0.0166)
	[12.004]	[5.045]	[13.638]
α_1	0.2147	0.2085	0.1949
	(0.0511)	(0.0288)	(0.0359)
	[4.197]	[7.246]	[5.425]
β_0	0.4282	0.7397	0.3155
	(0.0774)	(0.0755)	(0.0655)
	[5.534]	[9.803]	[4.814]
β_1	0.4769	0.6514	0.4645
	(0.1177)	(0.0493)	(0.1195)
	[4.052]	[13.208]	[3.888]

Table A.20: Parameter estimates for conditionally periodic GARCH with sample variance seasonality. ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1)\,\gamma(i) + (\alpha_0 d_0 + \alpha_1 d_1)\,\epsilon^2_{t,i-1} + (\beta_0 d_0 + \beta_1 d_1)\,\beta^2_{t,i-1} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= \begin{cases} 0 & \text{New York open} \\ 1 & \text{New York closed} \end{cases} \\
\gamma(i) &= \mathrm{Exp}(c_0 + c_1 i + \cdots + c_n i^n)
\end{aligned}
$$

Estimate	DM	JY	SF
ω_0	0.1913	0.08555	0.3114
	(0.0362)	(0.0237)	(0.0377)
	[5.284]	[3.610]	[8.255]
ω_1	0.2705	0.1263	0.3366
	(0.0538)	(0.0345)	(0.0705)
	[5.030]	[3.657]	[4.777]
α_0	0.1991	0.1569	0.2023
	(0.0214)	(0.0232)	(0.0161)
	[9.314]	[6.760]	[12.582]
α_1	0.2275	0.2119	0.2014
	(0.0462)	(0.0283)	(0.0329)
	[4.929]	[7.494]	[6.115]
β_0	0.6236	0.7694	0.4928
	(0.0525)	(0.0444)	(0.0384)
	[11.877]	[17.347]	[12.841]
β_1	0.4793	0.6416	0.4609
	(0.0658)	(0.0547)	(0.0820)
	[7.281]	[11.729]	[5.622]

Table A.21: Parameter estimates for conditionally periodic GARCH with exponential seasonality. ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= \omega\gamma(i) + \alpha_1\epsilon^2_{t,i-1} + \cdots + \alpha_{10}\epsilon^2_{t,i-10} \\
\gamma(i) &= \text{Sample Variance}
\end{aligned}
$$

Estimate	DM	JY	SF
ω	0.4851 [21.530]	0.4184 [19.334]	0.5803 [22.853]
α_1	0.1920 [9.556]	0.1607 [10.599]	0.1664 [10.637]
α_2	0.0900 [7.152]	0.1067 [7.101]	0.0902 [5.039]
α_3	0.0372 [3.492]	0.0604 [4.046]	0.0477 [3.779]
α_4	0.0343 [2.953]	0.0507 [4.382]	0.0200 [1.976]
α_5	0.0400 [3.578]	0.0464 [3.759]	0.0319 [3.353]
α_6	0.0250 [2.344]	0.0244 [2.633]	0.0126 [1.436]
α_7	0.0143 [1.736]	0.0248 [1.889]	0.0049 [0.771]
α_8	0.0051 [1.005]	0.0220 [2.199]	0.0000 [-]
α_9	0.0275 [1.593]	0.0128 [1.555]	0.0185 [2.078]
α_{10}	0.0247 [1.693]	0.0511 [3.167]	0.0056 [0.679]

Table A.22: Parameter estimates for conditionally periodic ARCH(10) with sample variance seasonality, $(T \times n = 19,392)$.

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma_{t,i}^2 &= \omega\gamma(i) + \alpha_1\epsilon_{t,i-1}^2 + \cdots + \alpha_{10}\epsilon_{t,i-10}^2 \\
\gamma(i) &= \text{Exp}(c_0 + c_1 i + \cdots + c_n i^n)
\end{aligned}
$$

Estimate	DM	JY	SF
ω	0.4523	0.3926	0.5597
	[18.161]	[18.537]	[20.861]
α_1	0.2102	0.1660	0.1746
	[9.038]	[10.323]	[10.196]
α_2	0.0912	0.1156	0.0915
	[6.536]	[6.939]	[6.063]
α_3	0.0366	0.0701	0.0547
	[3.036]	[4.094]	[3.938]
α_4	0.0406	0.0551	0.0264
	[3.004]	[4.231]	[2.242]
α_5	0.0559	0.0518	0.0293
	[2.823]	[3.456]	[2.945]
α_6	0.0216	0.0241	0.0154
	[1.920]	[2.639]	[1.670]
α_7	0.0195	0.0291	0.0099
	[2.081]	[2.068]	[1.346]
α_8	0.0035	0.0240	0.0031
	[0.680]	[2.097]	[0.371]
α_9	0.0372	0.0138	0.0161
	[1.604]	[1.600]	[1.783]
α_{10}	0.0342	0.0563	0.0131
	[1.885]	[3.606]	[1.277]

Table A.23: Parameter estimates for conditionally periodic ARCH(10) with exponential seasonality, $(T \times n = 19,392)$.

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1)\,\gamma(i) + (\alpha_{1,0} d_0 + \alpha_{1,1} d_1)\,\epsilon^2_{t,i-1} + \\
&\quad \cdots + (\alpha_{10,0} d_0 + \alpha_{10,1} d_1)\,\epsilon^2_{t,i-10} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= 1 - d_0 \\
\gamma(i) &= \text{Sample Variance}
\end{aligned}
$$

	DM		JY		SF	
Estimate	$j=0$	$j=1$	$j=0$	$j=1$	$j=0$	$j=1$
ω_j	0.4692	0.5066	0.4255	0.3746	0.5729	0.6074
	[18.404]	[14.209]	[17.394]	[11.525]	[19.262]	[16.377]
$\alpha_{1,j}$	0.1924	0.2046	0.1478	0.1985	0.1694	0.1620
	[8.909]	[4.417]	[8.987]	[6.302]	[9.764]	[4.661]
$\alpha_{2,j}$	0.1005	0.0538	0.0890	0.1470	0.0909	0.0873
	[6.539]	[3.041]	[5.677]	[4.685]	[3.993]	[3.775]
$\alpha_{3,j}$	0.0319	0.0468	0.0664	0.0479	0.0501	0.0359
	[2.851]	[2.065]	[3.870]	[1.573]	[3.274]	[2.010]
$\alpha_{4,j}$	0.0424	0.0166	0.0564	0.0368	0.0245	0.0049
	[3.066]	[1.165]	[4.102]	[2.142]	[1.997]	[0.382]
$\alpha_{5,j}$	0.0413	0.0328	0.0491	0.0488	0.0287	0.0400
	[2.997]	[2.162]	[3.423]	[2.057]	[2.608]	[2.231]
$\alpha_{6,j}$	0.0263	0.0173	0.0279	0.0129	0.0107	0.0137
	[2.051]	[1.211]	[2.569]	[1.053]	[1.006]	[0.965]
$\alpha_{7,j}$	0.0178	0.0128	0.0351	0.0031	0.0002	0.0134
	[1.619]	[1.253]	[2.198]	[0.262]	[0.024]	[1.140]
$\alpha_{8,j}$	0.0036	0.0140	0.0220	0.0293	0.0000	0.0040
	[0.700]	[1.260]	[1.801]	[1.776]	[-]	[0.521]
$\alpha_{9,j}$	0.0476	0.0052	0.0061	0.0292	0.0286	0.0069
	[1.736]	[0.736]	[0.793]	[1.528]	[2.145]	[0.848]
$\alpha_{10,j}$	0.0176	0.0341	0.0679	0.0211	0.0008	0.0093
	[1.007]	[1.545]	[2.810]	[1.199]	[0.074]	[0.948]

Table A.24: Parameter estimates for condtionally periodic ARCH(10) with with sample variance seasonality and NY dummies, ($T \times n = 19,392$).

$$
\begin{aligned}
R_{t,i} &= \epsilon_{t,i} \\
\sigma^2_{t,i} &= (\omega_0 d_0 + \omega_1 d_1)\,\gamma(i) + (\alpha_{1,0} d_0 + \alpha_{1,1} d_1)\,\epsilon^2_{t,i-1} + \\
&\quad \cdots + (\alpha_{10,0} d_0 + \alpha_{10,1} d_1)\,\epsilon^2_{t,i-10} \\
d_0 &= \begin{cases} 1 & \text{New York open} \\ 0 & \text{New York closed} \end{cases} \\
d_1 &= 1 - d_0 \\
\gamma(i) &= \text{Exponential Seasonality}
\end{aligned}
$$

Estimate	DM		JY		SF	
	$j=0$	$j=1$	$j=0$	$j=1$	$j=0$	$j=1$
ω_j	0.4275 [14.765]	0.4750 [12.974]	0.3869 [16.490]	0.3838 [11.562]	0.5438 [17.572]	0.6074 [16.136]
$\alpha_{1,j}$	0.2125 [8.733]	0.2268 [4.074]	0.1579 [8.739]	0.1908 [6.216]	0.1788 [9.433]	0.1612 [4.381]
$\alpha_{2,j}$	0.1015 [6.203]	0.0554 [2.578]	0.0992 [5.541]	0.1434 [4.666]	0.0918 [4.905]	0.0917 [3.773]
$\alpha_{3,j}$	0.0278 [2.367]	0.0581 [2.250]	0.0787 [3.942]	0.0451 [1.485]	0.0596 [3.519]	0.0437 [2.253]
$\alpha_{4,j}$	0.0543 [3.359]	0.0098 [0.767]	0.0642 [4.060]	0.0336 [1.958]	0.0347 [2.351]	0.0004 [0.031]
$\alpha_{5,j}$	0.0580 [2.598]	0.0342 [2.124]	0.0572 [3.201]	0.0429 [1.833]	0.0262 [2.248]	0.0361 [2.038]
$\alpha_{6,j}$	0.0241 [1.773]	0.0133 [0.985]	0.0279 [2.572]	0.0142 [1.095]	0.0147 [1.399]	0.0186 [1.138]
$\alpha_{7,j}$	0.0250 [1.978]	0.0176 [1.623]	0.0402 [2.387]	0.0026 [0.215]	0.0081 [0.937]	0.0156 [1.254]
$\alpha_{8,j}$	0.0018 [0.361]	0.0159 [1.349]	0.0256 [1.820]	0.0297 [1.710]	0.0009 [0.055]	0.0025 [0.305]
$\alpha_{9,j}$	0.0654 [1.869]	0.0041 [0.550]	0.0074 [0.972]	0.0337 [1.489]	0.0245 [1.762]	0.0057 [0.706]
$\alpha_{10,j}$	0.0363 [1.336]	0.0380 [1.757]	0.0773 [3.272]	0.0240 [1.318]	0.0155 [0.883]	0.0100 [0.953]

Table A.25: Parameter estimates for condtionally periodic ARCH(10) with with exponential variance seasonality and NY dummies, $(T \times n = 19,392)$.

Model	DM	JY	SF
Constant Variance	2.29471	2.09645	2.14999
Sample Variance Seasonality	2.46281	2.17245	2.33182
Exponential Seasonality	2.43974	2.15404	2.31079
GARCH	2.46480	2.23215	2.29638
GARCH: TOD	2.52887	2.26983	2.37662
GARCH: EXP	2.51358	2.25951	2.35621
ARCH	2.46834	2.22995	2.29629
ARCH: TOD	2.53483	2.26374	2.37832
ARCH: EXP	2.51535	2.25386	2.35662
P-GARCH	2.48152	2.23721	2.31740
P-GARCH: TOD	2.53408	2.27146	2.37883
P-GARCH: EXP	2.51384	2.26066	2.35632
P-ARCH	2.48865	2.23868	2.32035
P-ARCH: TOD	2.53554	2.26507	2.37876
P-ARCH: EXP	2.51618	2.25485	2.35712
S-ARCH	2.51650	2.25645	2.35815
Cond. Periodic GARCH: TOD	2.51498	2.24473	2.36377
Cond. Periodic GARCH: EXP	2.50264	2.23714	2.34897
Cond. Periodic GARCH: TOD with dummies	2.51513	2.24580	2.36397
Cond. Periodic GARCH: EXP with dummies	2.50342	2.24371	2.34991
Cond. Periodic ARCH: TOD	2.52746	2.25698	2.37276
Cond. Periodic ARCH: EXP	2.51115	2.24948	2.35442
Cond. Periodic ARCH: TOD with dummies	2.52817	2.25818	2.37323
Cond. Periodic ARCH: EXP with dummies	2.51266	2.25077	2.35501

Table A.26: Mean log-likelihood for all models, $(T \times n = 19,392)$.

Model	DM	JY	SF
Constant Variance	22.193	15.266	11.788
Sample Variance Seasonality	10.310	9.711	7.902
Exponential Seasonality	13.764	12.110	8.413
GARCH	15.441	9.467	12.628
GARCH: TOD	7.700	7.408	7.387
GARCH: EXP	10.904	8.110	8.142
ARCH	14.079	9.566	11.932
ARCH: TOD	7.390	8.256	7.327
ARCH: EXP	10.250	8.614	7.990
P-GARCH	15.405	9.140	13.367
P-GARCH: TOD	7.685	7.294	7.357
P-GARCH: EXP	10.714	8.035	8.052
P-ARCH	13.453	9.036	11.672
P-ARCH: TOD	7.233	8.048	7.348
P-ARCH: EXP	10.105	8.425	8.007
S-ARCH	10.984	8.506	8.010
Cond. Periodic GARCH: TOD	11.512	10.062	8.332
Cond. Periodic GARCH: EXP	14.540	9.461	8.979
Cond. Periodic GARCH: TOD with dummies	11.483	9.574	8.310
Cond. Periodic GARCH: EXP with dummies	14.066	9.094	8.985
Cond. Periodic ARCH: TOD	9.003	8.660	7.400
Cond. Periodic ARCH: EXP	11.516	8.826	8.395
Cond. Periodic ARCH: TOD with dummies	8.671	8.394	7.426
Cond. Periodic ARCH: EXP with dummies	10.867	8.588	8.398

Table A.27: In-sample kurtosis of returns standardized by sample variance estimates for all models, $(T \times n = 19,392)$.

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
Constant Variance	Sample Variance	H_2	H_2	H_2	LR
Constant Variance	Exponential Seasonality	H_2	H_2	H_2	LR
Exponential Seasonality	Sample Variance	H_2	H_2	H_2	LDC

Table A.28: Model selection: Seasonality Likelihood Ratio Tests ($T \times n = 19,392$).

Currency	Rejections (out of 96)	
	$m = 4$	$m = 9$
DM	64	45
JY	55	43
SF	71	52

Table A.29: Number of Fourier Series variance estimates which are outside 99% sample variance estimates confidence invervals ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
ARCH	ARCH:TOD	H_2	H_2	H_2	LDC
ARCH	ARCH:EXP	H_2	H_2	H_2	LDC
GARCH	GARCH:TOD	H_2	H_2	H_2	LDC
GARCH	GARCH:EXP	H_2	H_2	H_2	LDC
P-GARCH	P-GARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH	P-GARCH:EXP	H_2	H_2	H_2	LDC
P-ARCH	P-ARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH	P-ARCH:EXP	H_2	H_2	H_2	LDC

Table A.30: Model selection: Deseasonalization Likelihood Tests ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
ARCH:EXP	ARCH:TOD	H_2	H_2	H_2	LDC
GARCH:EXP	GARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH:EXP	P-GARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH:EXP	P-ARCH:TOD	H_2	H_2	H_2	LDC
CP-ARCH:EXP	CP-ARCH:TOD	H_2	H_2	H_2	LDC
CP-GARCH:EXP	CP-GARCH:TOD	H_2	H_2	H_2	LDC

Table A.31: Model selection: Testing for Exponential Seasonality versus Sample Variance Seasonality ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
ARCH	CP-ARCH:TOD	H_2	H_2	H_2	LDC
ARCH	CP-ARCH:EXP	H_2	H_2	H_2	LDC
GARCH	CP-GARCH:TOD	H_2	H_2	H_2	LDC
GARCH	CP-GARCH:EXP	H_2	H_2	H_2	LDC
ARCH	S-ARCH	H_2	H_2	H_2	LR
GARCH	S-ARCH	H_2	H_2	H_2	LDC
ARCH	P-ARCH	H_2	H_2	H_2	LR
GARCH	P-GARCH	H_2	H_2	H_2	LR

Table A.32: Model selection: Testing Joint Modeling versus Raw Modeling ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
P-ARCH	CP-GARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH	CP-GARCH:EXP	H_2	H_1	H_2	LDC
P-ARCH	CP-ARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH	CP-ARCH:EXP	H_2	H_2	H_2	LDC
P-ARCH	ARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH	ARCH:EXP	H_2	H_2	H_2	LDC
P-ARCH	GARCH:TOD	H_2	H_2	H_2	LDC
P-ARCH	GARCH:EXP	H_2	H_2	H_2	LDC
P-GARCH	CP-GARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH	CP-GARCH:EXP	H_2	Inc.	H_2	LDC
P-GARCH	CP-ARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH	CP-ARCH:EXP	H_2	H_2	H_2	LDC
P-GARCH	ARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH	ARCH:EXP	H_2	H_2	H_2	LDC
P-GARCH	GARCH:TOD	H_2	H_2	H_2	LDC
P-GARCH	GARCH:EXP	H_2	H_2	H_2	LDC

Table A.33: Model selection: Testing Joint Modeling versus Sequential Modeling ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected			Test
		DM	JY	SF	
GARCH:TOD	CP-GARCH:EXP	H_1	H_1	H_1	LDC
GARCH:TOD	CP-ARCH:TOD	H_1	H_1	H_1	LDC
GARCH:TOD	CP-ARCH:EXP	H_1	H_1	H_1	LDC
GARCH:EXP	CP-GARCH:TOD	Inc.	H_1	H_2	LDC
GARCH:EXP	CP-ARCH:TOD	H_2	Inc.	H_2	LDC
GARCH:EXP	CP-ARCH:EXP	H_1	H_1	H_1	LDC
ARCH:TOD	CP-GARCH:TOD	H_1	H_1	H_1	LDC
ARCH:TOD	CP-GARCH:EXP	H_1	H_1	H_1	LDC
ARCH:TOD	CP-ARCH:EXP	H_1	H_1	H_1	LDC
ARCH:EXP	CP-GARCH:TOD	Inc.	H_1	H_2	LDC
ARCH:EXP	CP-GARCH:EXP	H_1	H_1	H_1	LDC
ARCH:EXP	CP-ARCH:TOD	H_2	H_2	H_2	LDC

Table A.34: Model selection: Testing Joint Modeling versus Sequential Modeling ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected DM	JY	SF	Test
S-ARCH	GARCH:TOD	H_2	H_2	H_2	LDC
S-ARCH	ARCH:TOD	H_2	H_2	H_2	LDC
S-ARCH	GARCH:EXP	H_1	H_2	H_1	LDC
S-ARCH	CP-ARCH:TOD	H_2	Inc.	H_2	LDC
S-ARCH	CP-GARCH:TOD	Inc.	H_1	H_2	LDC
S-ARCH	CP-GARCH:EXP	H_1	H_1	H_1	LDC
S-ARCH	P-GARCH	H_1	H_1	H_1	LDC
S-ARCH	P-ARCH	H_1	H_1	H_1	LDC

Table A.35: Model selection: Testing Joint Modeling versus Sequential Modeling ($T \times n = 19,392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected DM	JY	SF	Test
GARCH	ARCH	H_2	H_1	Inc.	LDC
GARCH:TOD	ARCH:TOD	H_2	H_1	H_2	LDC
GARCH:EXP	ARCH:EXP	H_2	H_1	H_2	LDC
P-GARCH	P-ARCH	H_2	H_2	H_2	LDC
P-GARCH:TOD	P-ARCH:TOD	H_2	H_1	Inc.	LDC
P-GARCH:EXP	P-ARCH:EXP	H_2	H_1	H_2	LDC
CP-GARCH:TOD	CP-ARCH:TOD	H_2	H_2	H_2	LDC
CP-GARCH:EXP	CP-ARCH:EXP	H_2	H_2	H_2	LDC
CP-GARCH:TOD + dum	CP-ARCH:TOD + dum	H_2	H_2	H_2	LDC
CP-GARCH:EXP + dum	CP-ARCH:EXP + dum	H_2	H_2	H_2	LDC

Table A.36: Model selection: Testing ARCH versus GARCH ($T \times n = 19.392$).

Hypothesis: H_1	Hypothesis: H_2	Model Selected DM	JY	SF	Test
GARCH	P-GARCH	H_2	H_2	H_2	LR
GARCH:TOD	P-GARCH:TOD	H_2	H_2	H_2	LR
GARCH:EXP	P-GARCH:EXP	H_1	H_2	H_1	LR
ARCH	P-ARCH	H_2	H_2	H_2	LR
ARCH:TOD	P-ARCH:TOD	H_2	H_2	H_2	LR
ARCH:EXP	P-ARCH:EXP	H_2	H_2	H_2	LR
CP-GARCH:TOD	CP-GARCH:TOD + dum	H_1	H_2	H_1	LR
CP-GARCH:EXP	CP-GARCH:EXP + dum	H_2	H_2	H_2	LR
CP-ARCH:TOD	CP-ARCH:TOD + dum	H_2	H_2	H_2	LR
CP-ARCH:EXP	CP-ARCH:EXP + dum	H_2	H_2	H_2	LR

Table A.37: Model selection: Testing for NY Effect ($T \times n = 19.392$).

Model	DM	JY	SF
Constant Variance	21.186	14.262	10.786
Sample Variance Seasonality	9.321	8.704	6.917
Exponential Seasonality	12.762	11.108	7.414
GARCH	14.437	8.465	11.629
GARCH: TOD	7.050	6.308	6.551
GARCH: EXP	9.901	7.106	7.140
ARCH	13.082	8.570	10.921
ARCH: TOD	6.392	7.259	6.321
ARCH: EXP	9.254	7.619	6.986
P-GARCH	14.396	8.129	12.363
P-GARCH: TOD	6.686	6.291	6.354
P-GARCH: EXP	9.710	7.032	7.057
P-ARCH	12.444	8.028	10.664
P-ARCH: TOD	6.230	7.042	6.346
P-ARCH: EXP	9.102	7.418	7.006
S-ARCH	9.986	7.510	7.005
Cond. Periodic GARCH: TOD	10.511	9.062	7.330
Cond. Periodic GARCH: EXP	13.535	8.455	8.045
Cond. Periodic GARCH: TOD with dummies	10.483	8.572	7.310
Cond. Periodic GARCH: EXP with dummies	13.060	8.087	7.984
Cond. Periodic ARCH: TOD	8.006	7.663	6.395
Cond. Periodic ARCH: EXP	10.517	7.831	7.388
Cond. Periodic ARCH: TOD with dummies	7.665	7.386	6.420
Cond. Periodic ARCH: EXP with dummies	9.885	7.581	7.394

Table A.38: In-sample heteroskedastic mean square error (HMSE) measures for all models, ($T \times n = 19,392$).

Model	DM	JY	SF
Constant Variance	69.296	6.143	19.821
Sample Variance Seasonality	46.796	4.824	9.598
Exponential Seasonality	41.980	4.826	11.674
GARCH	24.667	6.233	9.766
GARCH: TOD	12.963	5.351	8.122
GARCH: EXP	16.928	5.355	8.543
ARCH	28.435	6.023	9.253
ARCH: TOD	16.694	4.963	8.306
ARCH: EXP	20.273	5.146	8.380
P-GARCH	28.884	5.541	11.078
P-GARCH: TOD	12.162	5.331	7.909
P-GARCH: EXP	16.883	5.244	8.627
P-ARCH	34.852	5.141	9.625
P-ARCH: TOD	15.336	5.060	8.229
P-ARCH: EXP	19.313	5.195	8.269
S-ARCH	18.216	5.156	8.319
Cond. Periodic GARCH: TOD	20.599	5.196	8.537
Cond. Periodic GARCH: EXP	20.052	5.352	8.343
Cond. Periodic GARCH: TOD with dummies	20.499	5.212	8.604
Cond. Periodic GARCH: EXP with dummies	20.018	5.313	8.231
Cond. Periodic ARCH: TOD	16.749	4.938	8.377
Cond. Periodic ARCH: EXP	20.871	5.186	8.401
Cond. Periodic ARCH: TOD with dummies	15.704	4.985	8.124
Cond. Periodic ARCH: EXP with dummies	20.574	5.151	8.244

Table A.39: Out-of-sample heteroskedastic mean square error (HMSE) measures for all models, $(T \times n = 5,760)$.

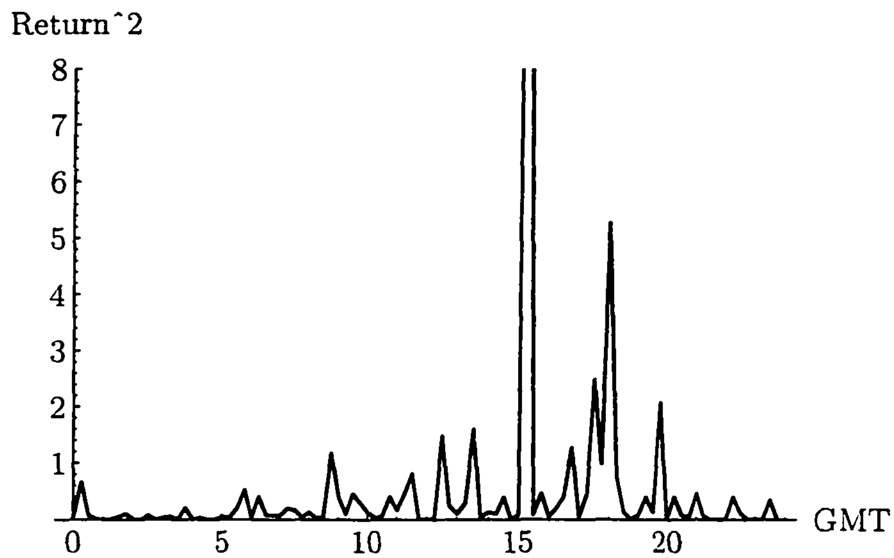

Figure A.1: Raw DM returns for December 6, 1993 (standardized by sample variance).

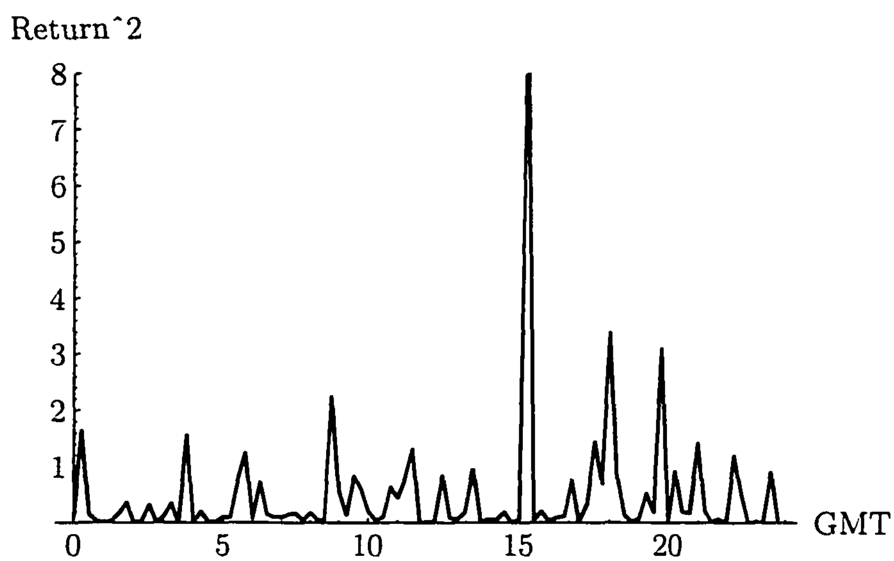

Figure A.2: Deseasoned DM returns for December 6, 1993 (standardized by intraday sample variance).

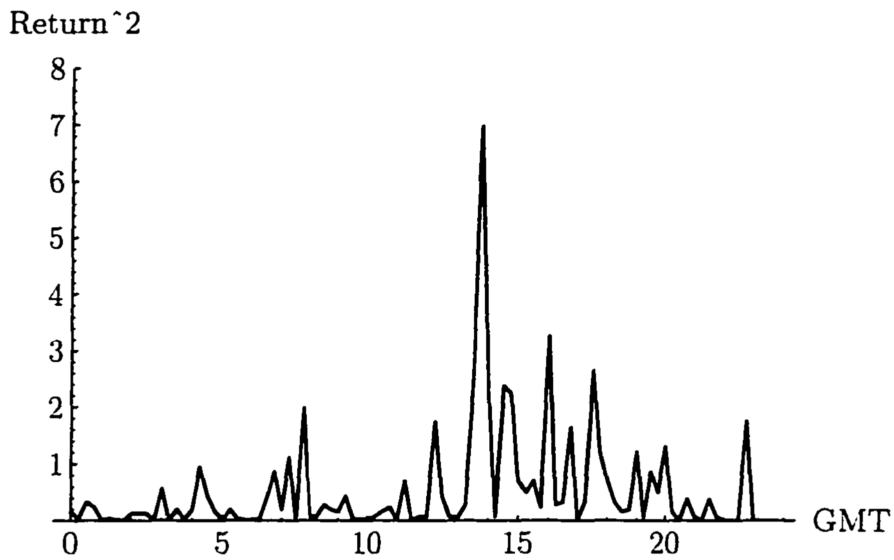

Figure A.3: Raw DM returns for February 11, 1994 (standardized by sample variance).

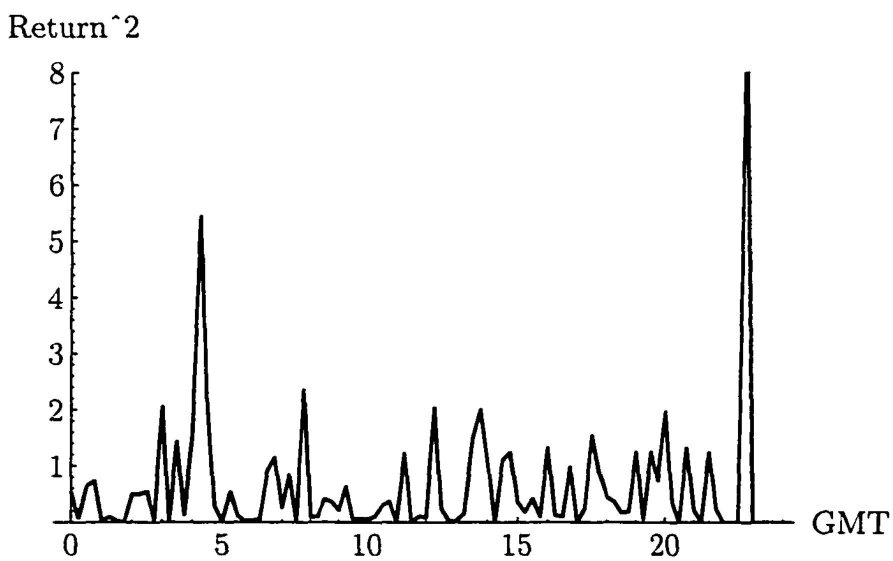

Figure A.4: Deseasoned DM returns for February 11, 1994 (standardized by intraday sample variance).

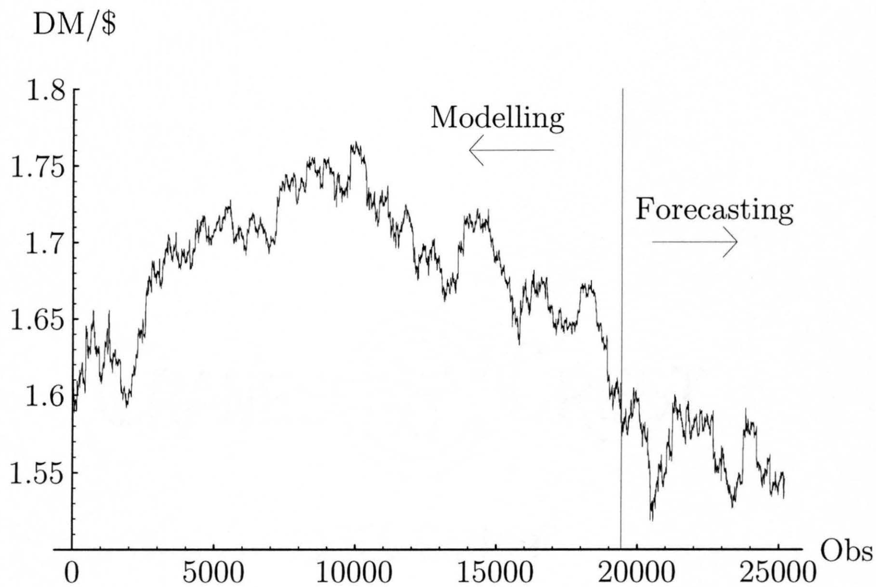

Figure A.5: DM exchange rate sample.

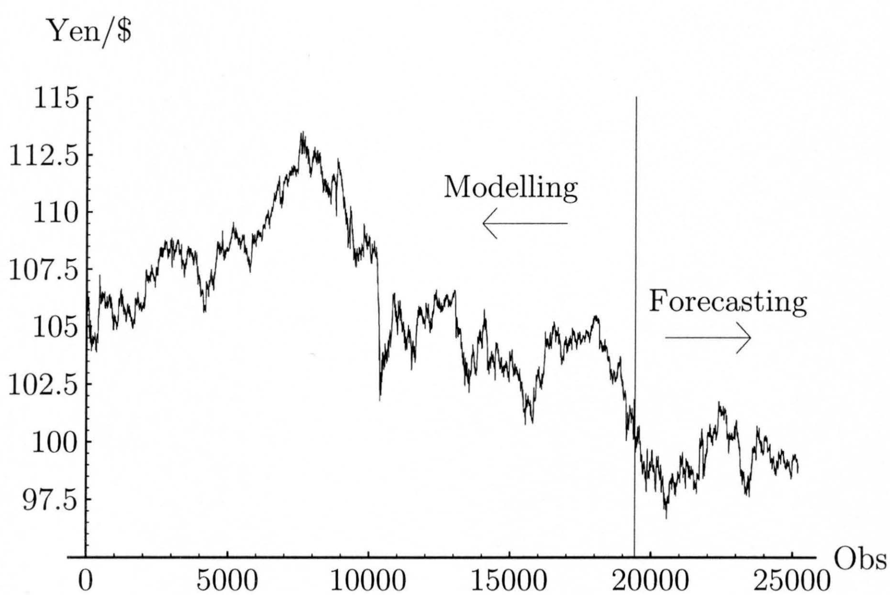

Figure A.6: JY exchange rate sample.

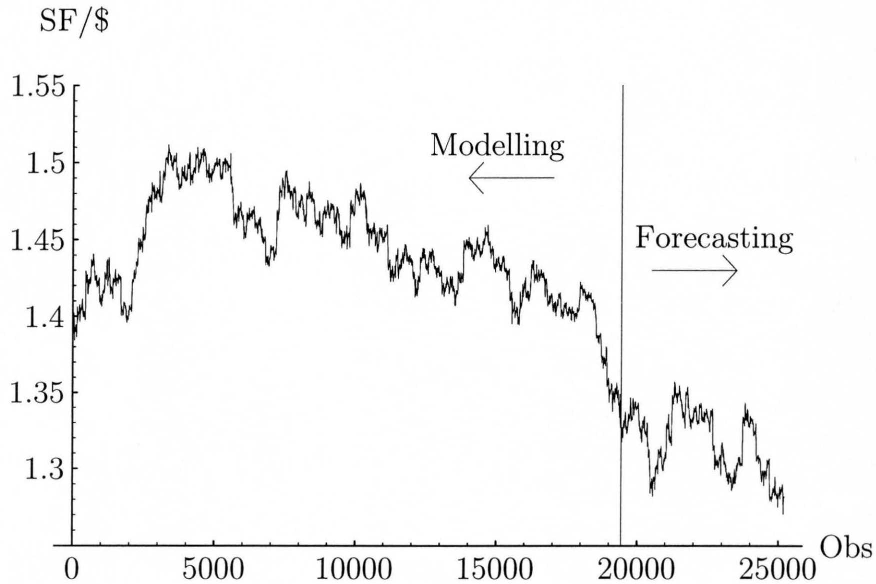

Figure A.7: SF exchange rate sample.

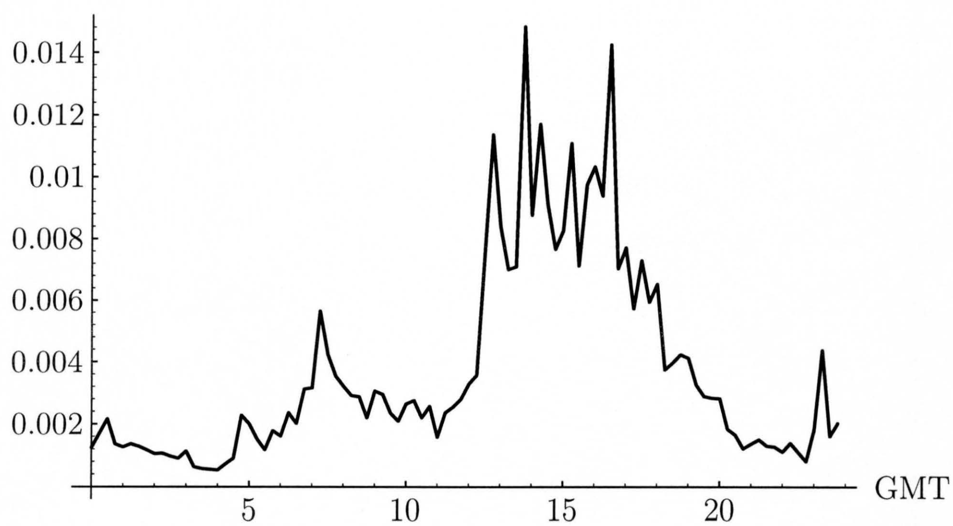

Figure A.8: DM time of day variance.

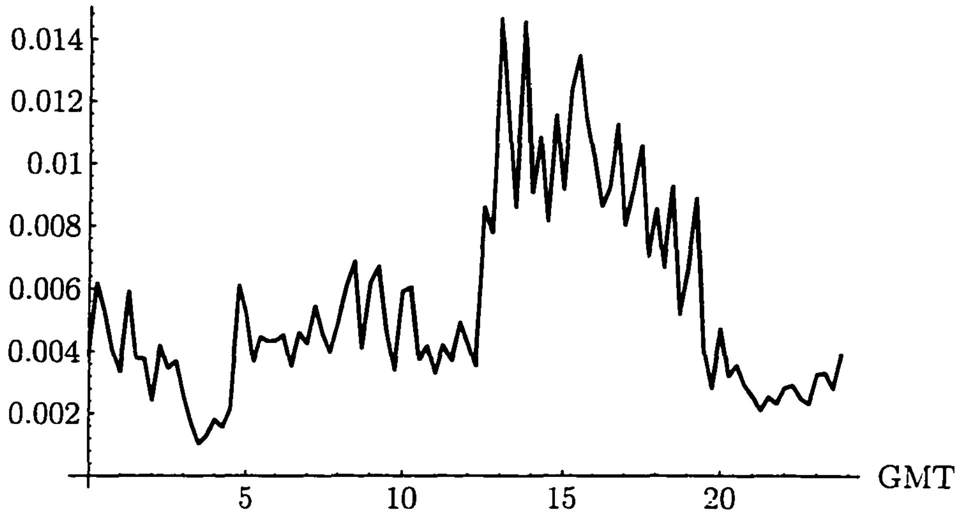

Figure A.9: JY time of day variance.

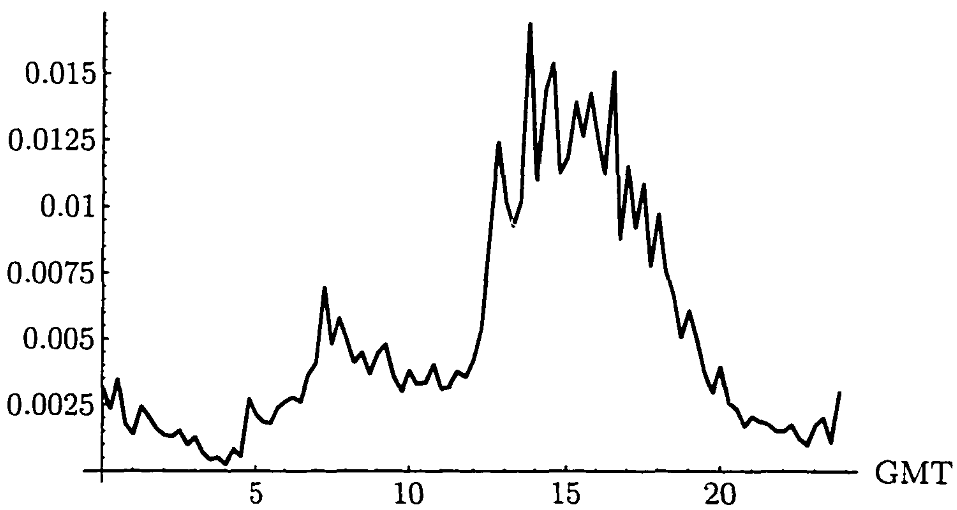

Figure A.10: SF time of day variance.

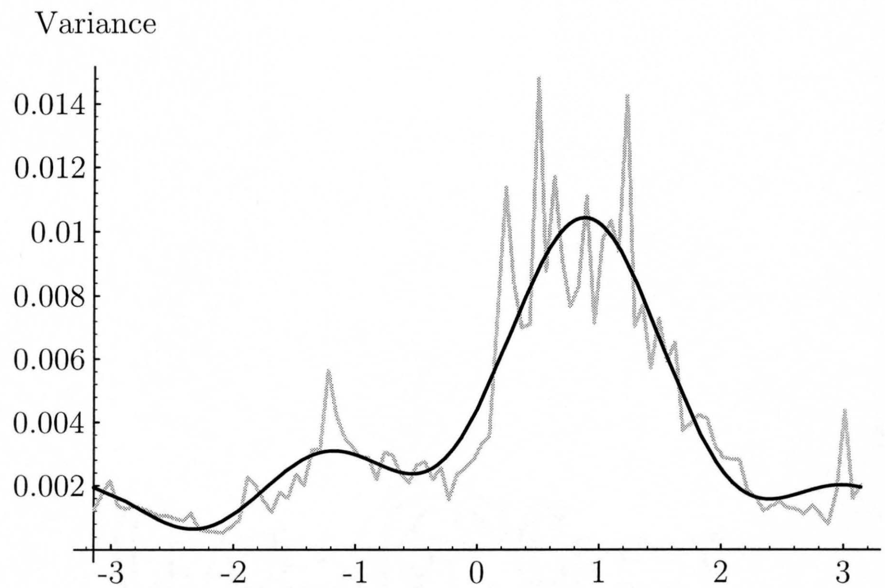

Figure A.11: DM fourier series fit for time of day variance $m = 4$.

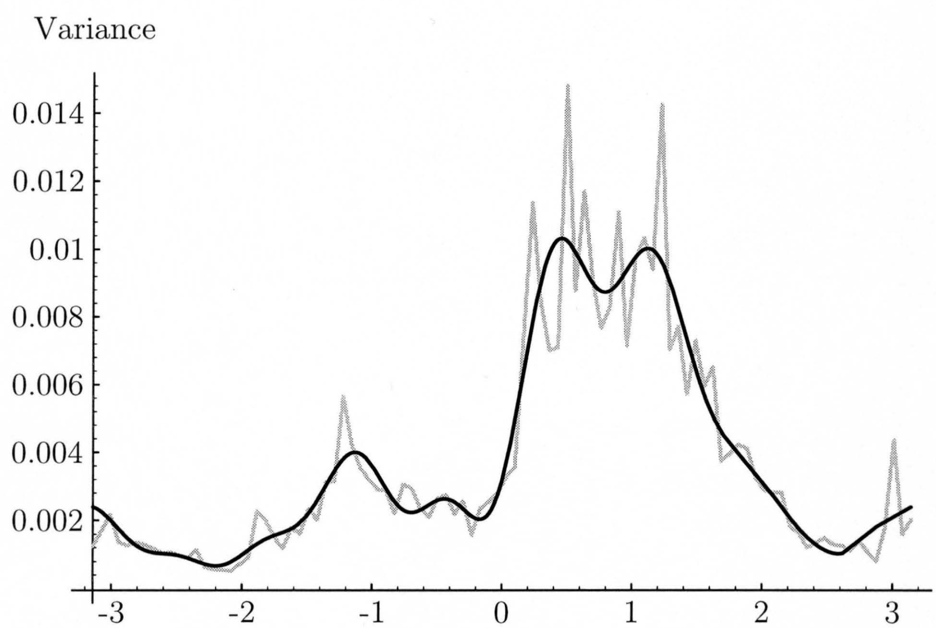

Figure A.12: DM fourier series fit for time of day variance $m = 9$.

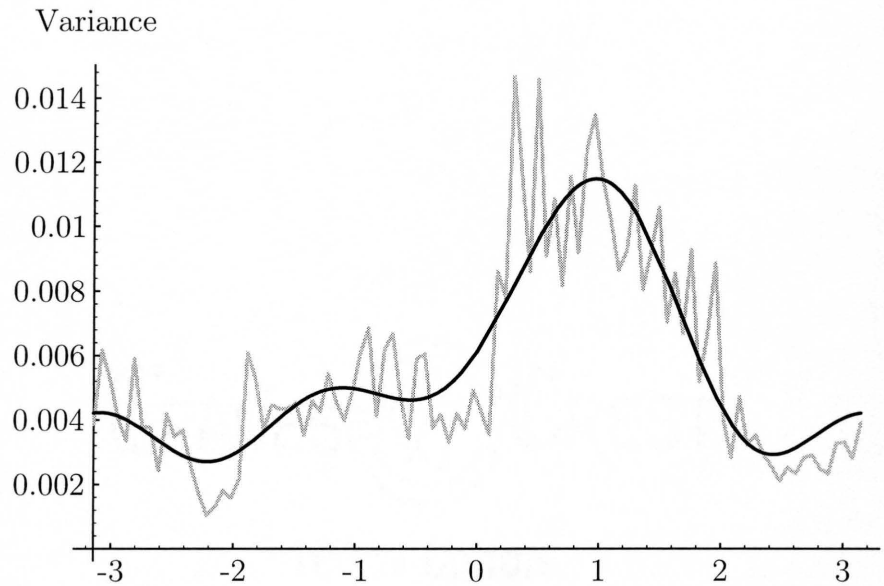

Figure A.13: JY fourier series fit for time of day variance $m = 4$.

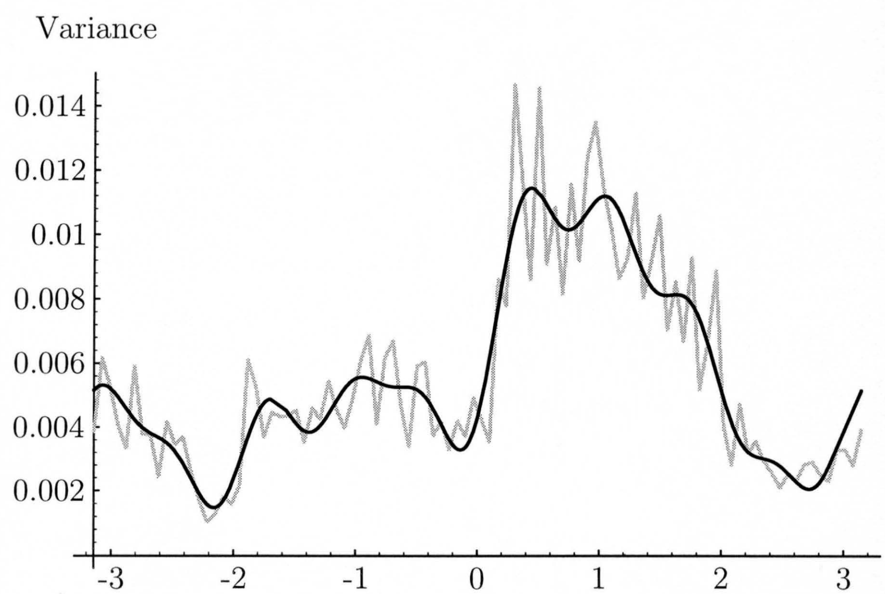

Figure A.14: JY fourier series fit for time of day variance $m = 9$.

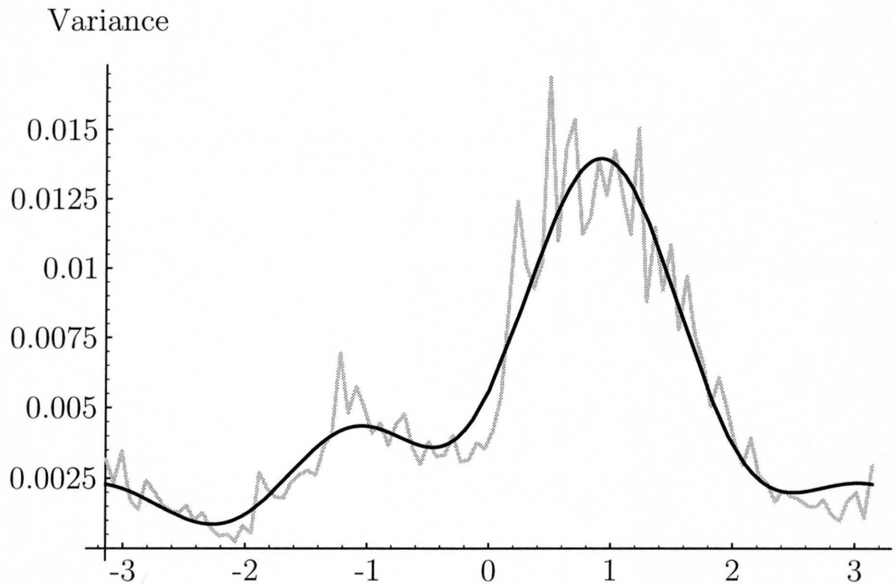

Figure A.15: SF fourier series fit for time of day variance $m = 4$.

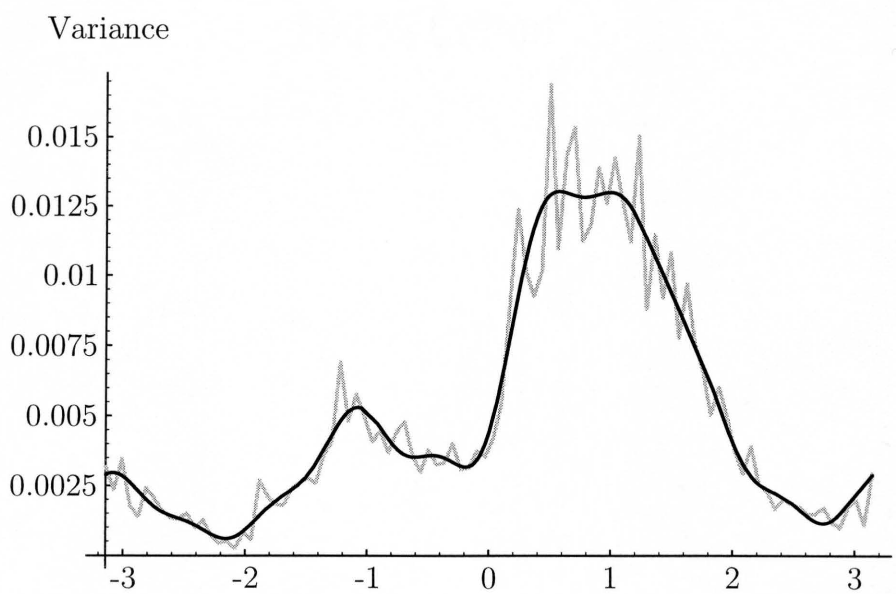

Figure A.16: SF fourier series fit for time of day variance $m = 9$.

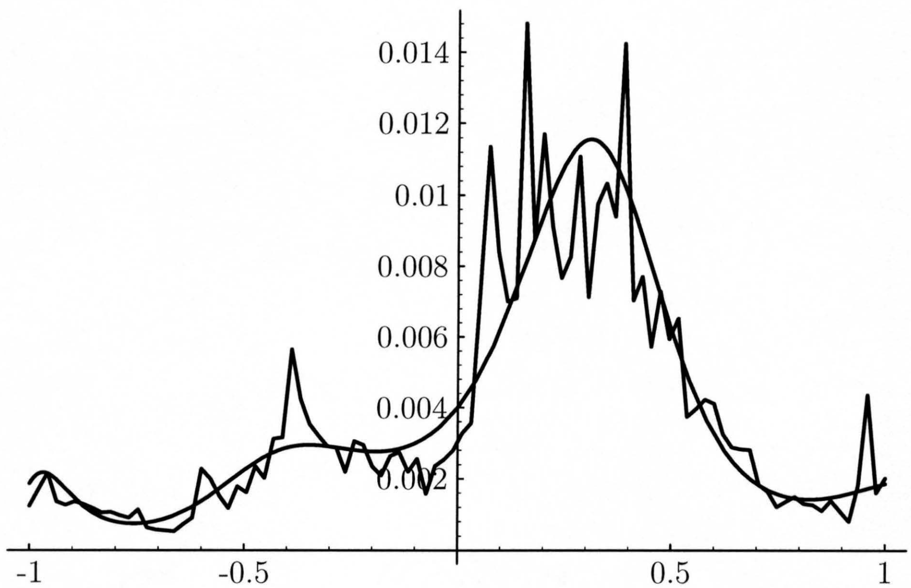

Figure A.17: DM exponential seasonality time of day variance.

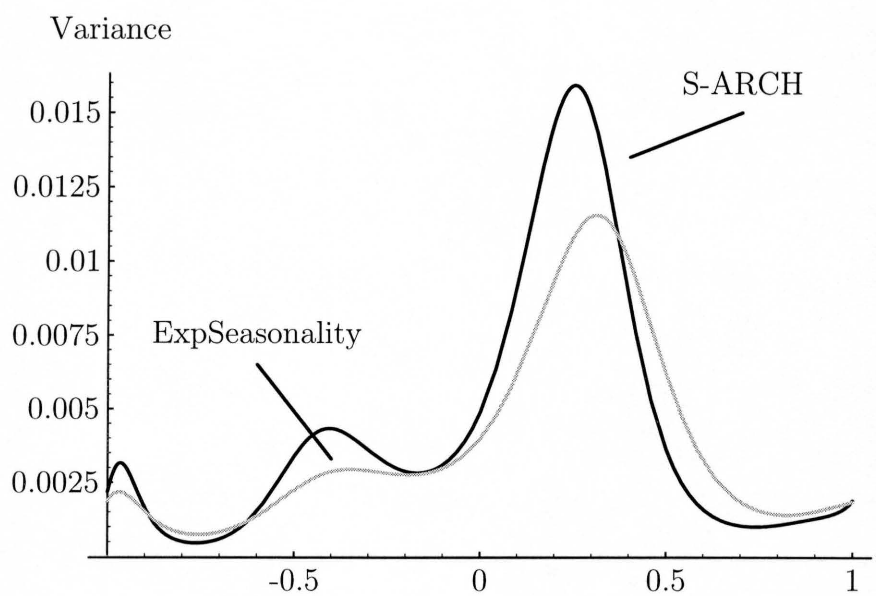

Figure A.18: Comparison of unconditional time of day variance for DM.

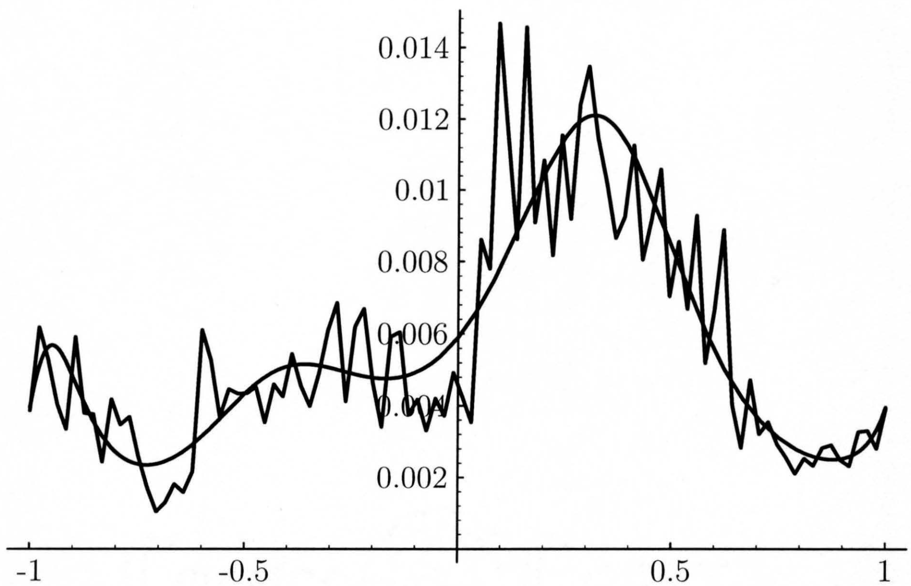

Figure A.19: JY exponential seasonality time of day variance.

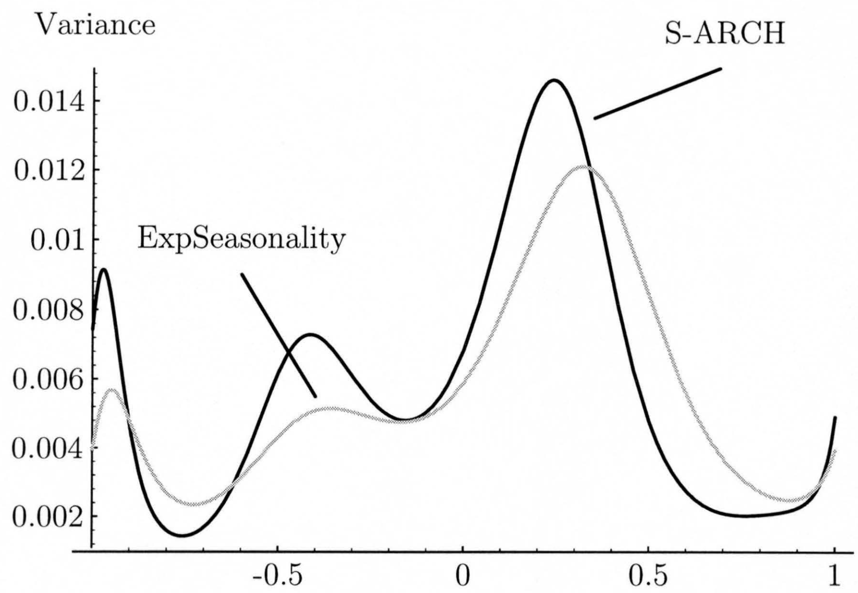

Figure A.20: Comparison of unconditional time of day variance for JY.

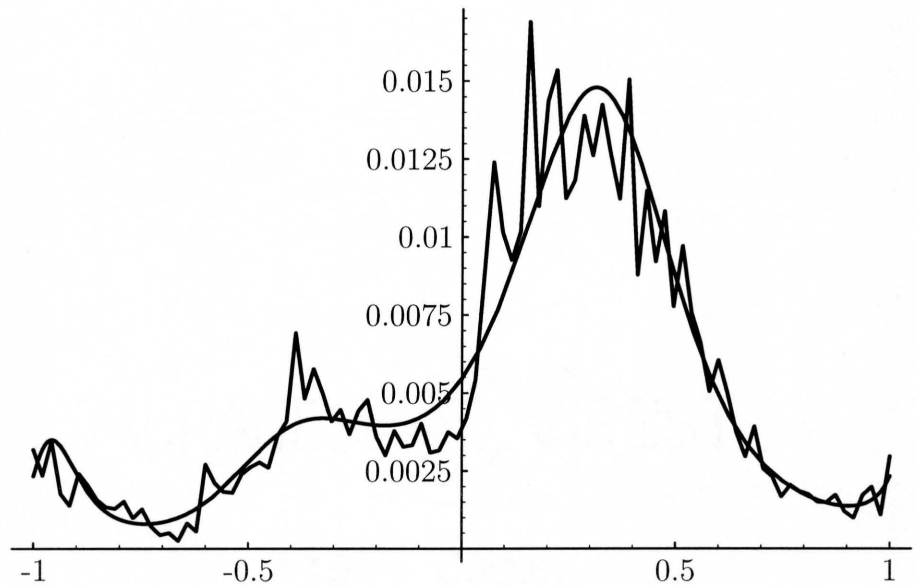

Figure A.21: SF exponential seasonality time of day variance.

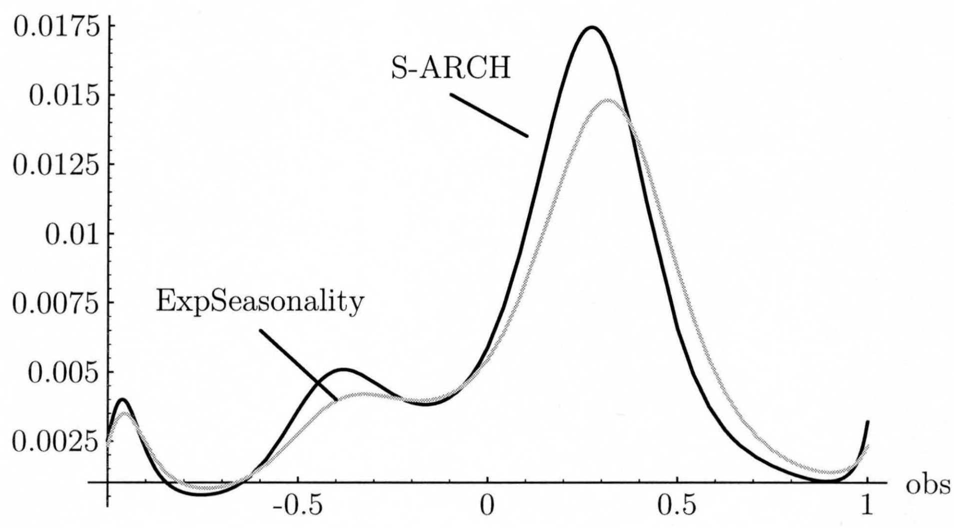

Figure A.22: Comparison of unconditional time of day variance for SF.

Appendix B: Tables and Figures for Chapter 2

Note: In the following tables, numbers is parentheses are heteroskedastic-consistent standard errors and numbers in brackets are t-ratios. See table and figure notes at the end of the tables and figures respectfully for more details.

Fifteen Minute Returns	DM	JY	SF
Equal to Zero	2292	3002	4879
Min	-0.8200	-0.8571	-0.9632
Max	1.2845	0.9123	0.7814
Range	2.1046	1.7695	1.7446
Mean	-0.0002	-0.0002	-0.0004
Median	0	0	0
Variance	0.0042	0.0055	0.0053
Standard Deviation	0.0648	0.0741	0.0728
Mean Deviation	0.0407	0.0486	0.0466
Median Deviation	0.0269	0.0326	0.0340
Skewness	0.4665	0.1394	-0.0582
Excess Kurtosis	23.573	10.699	9.394

Table B.40: Descriptive statistics for fifteen minute returns ($T \times n = 25,152$).

Daily Realized Returns	DM	JY	SF
Min	-1.7194	-3.6904	-1.8069
Max	1.9924	1.9046	2.1870
Range	3.7118	5.5949	3.9939
Mean	-0.0152	-0.0154	-0.0373
Median	-0.0299	-0.0587	-0.0354
Variance	0.3721	0.4466	0.4370
Standard Deviation	0.6100	0.6683	0.6610
Mean Deviation	0.4669	0.5060	0.5026
Median Deviation	0.3860	0.4107	0.4124
Skewness	0.2220	-0.4589	0.2080
Excess Kurtosis	0.8079	3.0872	0.8078

Table B.41: Descriptive statistics for realized returns (T=262).

Daily Squared Daily Returns	DM	JY	SF
Min	0	0	0
Max	3.9697	13.6187	4.7832
Range	3.9697	13.6187	4.7832
Mean	0.3709	0.4451	0.4367
Median	0.1483	0.1624	0.1704
Variance	0.3886	1.0289	0.5331
Standard Deviation	0.6234	1.0144	0.7301
Mean Deviation	0.3894	0.4638	0.4577
Median Deviation	0.1309	0.1480	0.1572
Skewness	3.0579	8.9118	3.0306
Excess Kurtosis	10.3977	107.084	10.6857

Table B.42: Descriptive statistics for squared daily returns (T=262).

CAR Variance Estimates	DM	JY	SF
Min	0.0381	0.0502	0.0811
Max	1.4103	3.4045	1.0992
Range	1.3722	3.3543	1.0181
Mean	0.2684	0.3962	0.3427
Median	0.2287	0.3124	0.3132
Variance	0.0258	0.1114	0.0246
Standard Deviation	0.1607	0.3338	0.1568
Mean Deviation	0.1116	0.1979	0.1184
Median Deviation	0.0748	0.1079	0.0910
Skewness	2.5168	4.4776	1.3304
Excess Kurtosis	11.460	32.523	3.0910

Table B.43: Descriptive statistics for cumulative absolute returns daily variance estimates (n=96).

CSR Variance Estimates	DM	JY	SF
Min	0.0443	0.0826	0.0983
Max	3.5016	4.2520	1.7748
Range	3.4573	4.1694	1.6764
Mean	0.4026	0.5271	0.5089
Median	0.2954	0.4014	0.4406
Variance	0.1281	0.2104	0.0740
Standard Deviation	0.3579	0.4587	0.2720
Mean Deviation	0.2210	0.2798	0.2055
Median Deviation	0.1079	0.1519	0.1466
Skewness	4.4001	3.6680	1.5432
Excess Kurtosis	30.021	19.942	3.4130

Table B.44: Descriptive statistics for cumulative squared returns daily variance estimates (n=96).

Unconditional Daily Variance Estimates	DM	JY	SF
realized returns $\hat{\sigma}^2 = \frac{1}{T}\sum(R_t - \bar{R})^2$	0.3709 (0.32, 0.45)	0.4451 (0.38, 0.53)	0.4367 (0.37, 0.52)
absolute returns $\hat{\sigma}^2 = \left(\frac{1}{T}\sqrt{\frac{2}{\pi}}(\sum\|R_t\|)\right)^2$	0.3433 (0.28, 0.42)	0.4041 (0.32, 0.49)	0.3991 (0.32, 0.49)
cumulative absolute returns $\hat{\sigma}^2 = \frac{2}{\pi N}\frac{1}{T}\sum_t^T\left(\sum_{i=1}^n \|R_{t,i}\|\right)^2$	0.2684 (0.25, 0.29)	0.3962 (0.36, 0.44)	0.3427 (0.32, 0.36)
cumulative squared returns $\hat{\sigma}^2 = \left(\frac{1}{T}\sum_t^T(\sum_{i=1}^n R_{t,i}^2)\right)^2$	0.4026 (0.36, 0.45)	0.5271 (0.47, 0.58)	0.5089 (0.48, 0.54)

Table B.45: Different estimates of variance for daily realized returns with 95% confidence intervals.

Daily CSR Standardized Returns	DM	JY	SF
Min	-2.4046	-2.0579	-2.1800
Max	2.8303	2.1051	2.5093
Range	5.2348	4.1630	4.6893
Mean	-0.0300	-0.0140	-0.0513
Median	-0.0586	-0.0795	-0.0648
Variance	0.8735	0.7752	0.7780
Standard Deviation	0.9346	0.8804	0.8821
Mean Deviation	0.7609	0.7291	0.7035
Median Deviation	0.6543	0.6541	0.7035
Skewness	0.1614	0.1886	0.1720
Excess Kurtosis	-0.2574	-0.6993	-0.0767

Table B.46: Descriptive statistics for returns standardized by dividing by square root cumulative squared returns (n=96).

Bera Jarque Normality Test	DM	JY	SF
Fifteen minute returns $R_{t,i}$	583247	120045	92503
Realized returns R_t	9.27	113.24	9.01
CSR Standardized returns $R_t/\sqrt{\sum_{i=1}^{n} R_{t,i}^2}$	1.86	6.89	1.36
CAR Standardized returns $R_t/\sqrt{2\pi 96}\sum_{i=1}^{n}\lvert R_{t,i}\rvert$	1.71	7.31	1.58
GARCH(1,1) Standardized returns $R_t/\sqrt{\text{GARCH}}$	8.06	88.47	8.66

Table B.47: Normality tests (5% critical value is 5.99; the 1% critical value is 9.21).

Correlations	DM	JY	SF
R_t and R_{t-1}	-0.0612	-0.1213	-0.0888
$\|R_t\|$ and $\|R_{t-1}\|$	-0.1257	0.0363	-0.1423
R_t^2 and R_{t-1}^2	-0.1278	0.0752	-0.1155
CSR_t and CSR_{t-1}	0.1569	0.3765	0.2080
CAR_t and CAR_{t-1}	0.4011	0.5366	0.3543
R_t^2 and CSR_t	0.4649	0.6156	0.4353
R_t^2 and CAR_t	0.4656	0.6510	0.3904
CSR_t and CAR_t	0.9063	0.9390	0.8777
GARCH and R_t^2	0.0349	0.0501	-0.0029
GARCH and CSR_t	0.2237	0.3520	0.1666
GARCH and CAR_t	0.2588	0.4552	0.1594

Table B.48: Serial Correlations and contemporaneous correlations (T=262).

Percentage Bias in CAR	DM	JY	SF
Min	0.22	-3.49	2.92
Max	47.12	39.55	51.92
Range	46.90	43.05	49.00
Mean	14.30	12.04	16.42
Median	13.05	10.73	15.25
Variance	69.44	50.84	56.78
Standard Deviation	8.33	7.13	7.54
Mean Deviation	6.25	5.39	5.59
Median Deviation	4.63	3.88	4.24
Skewness	1.27	1.10	1.21
Excess Kurtosis	2.04	1.63	2.36

Table B.49: Descriptive statistics for $100 \times \left(\frac{\mathrm{CSR}-\mathrm{CAR}}{\mathrm{CSR}}\right)$, the percentage difference in CAR relative to CSR as a daily standard deviation estimator (n=96).

Percentage Bias in CAR	DM	JY	SF
Min	0.43	-7.11	5.76
Max	72.03	63.46	76.88
Range	71.60	70.57	71.13
Mean	25.86	22.12	29.58
Median	24.39	20.30	28.17
Variance	180.63	144.22	142.74
Standard Deviation	13.44	12.01	11.95
Mean Deviation	10.40	9.25	9.06
Median Deviation	7.97	7.08	7.30
Skewness	0.92	0.80	0.82
Excess Kurtosis	0.87	0.81	0.99

Table B.50: Descriptive statistics for $100 \times \left(\frac{\text{CSR}-\text{CAR}}{\text{CSR}}\right)$, the percentage difference in CAR relative to CSR as a daily variance estimator (n=96).

Currency		Low Ret^2 Low CSR^2	Low Ret^2 High CSR^2	High Ret^2 Low CSR^2	High Ret^2 High CSR^2
DM	Variance Estimate from Returns	0.044	0.049	0.389	0.846
	Variance Estimate from CSR	0.157	0.333	0.174	0.395
	Total Number	88	43	43	88
	Number where $Ret^2 > CSR$	1	0	30	45
	Number where $Ret^2 < CSR$	87	43	13	43
JY	Variance Estimate from Returns	0.054	0.057	0.387	1.086
	Variance Estimate from CSR	0.215	0.542	0.222	0.592
	Total Number	84	47	47	84
	Number where $Ret^2 > CSR$	1	0	33	39
	Number where $Ret^2 < CSR$	83	47	14	45
SF	Variance Estimate from Returns	0.048	0.053	0.615	0.937
	Variance Estimate from CSR	0.221	0.429	0.253	0.466
	Total Number	85	46	46	85
	Number where $Ret^2 > CSR$	1	0	33	30
	Number where $Ret^2 < CSR$	84	46	13	55

Table B.51: Joint distribution comparison of daily CSR variance estimates and daily squared returns (T=262).

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2
\end{aligned}
$$

Estimate	DM	JY	SF
μ	-0.0156	-0.0038	-0.0360
	(0.0367)	(0.412)	(0.0406)
	[-0.425]	[-0.093]	[-0.886]
ω	0.0089	0.0392	0.0119
	(0.0097)	(0.0319)	(0.0193)
	[0.923]	[1.229]	[0.615]
α	0.0166	0.0393	0.0063
	(0.0148)	(0.0367)	(0.0147)
	[1.124]	[1.070]	[0.430]
β	0.9587	0.8696	0.9666
	(0.0309)	(0.0841)	(0.0461)
	[31.068]	[10.339]	[20.981]

Table B.52: GARCH parameter estimates (T=262).

GARCH Variance Estimates	DM	JY	SF
Min	0.2993	0.3411	0.4081
Max	0.4617	0.9725	0.4827
Range	0.1624	0.6314	0.0746
Mean	0.3611	0.4371	0.4352
Median	0.3464	0.4201	0.4333
Variance	0.0016	0.0077	0.0003
Standard Deviation	0.0400	0.0876	0.0168
Mean Deviation	0.0346	0.0545	0.0132
Median Deviation	0.0237	0.0368	0.0101
Skewness	0.5991	3.0856	0.7587
Excess Kurtosis	-0.8666	12.975	0.0169

Table B.53: Descriptive statistics for GARCH(1,1) conditional variance estimates.

Currency	Model	MSE	MAE	LL	HMSE	GMLE
DM	Constant	0.3871	0.3894	9.8666	2.8138	0.0082
DM	CSR	0.3091	0.3449	8.8677	1.3393	-0.2666
DM	CAR	0.3304	0.3116	8.1489	3.1040	-0.2015
DM	GARCH	0.3870	0.3862	9.7135	2.9247	0.0075
JY	Constant	1.0250	0.4638	9.5624	5.1732	0.1906
JY	CSR	0.6706	0.4350	8.8256	0.8313	-0.0947
JY	CAR	0.6992	0.3880	8.2585	1.4045	-0.1009
JY	GARCH	1.0238	0.4597	9.4724	5.0515	0.1832
SF	Constant	0.5310	0.4577	24.628	2.7847	0.1715
SF	CSR	0.4377	0.4418	23.903	1.2023	-0.0241
SF	CAR	0.4753	0.3896	22.809	2.7526	-0.0067
SF	GARCH	0.5314	0.4569	24.618	2.7526	0.1724

Table B.54: In-sample statistical evaluation of estimates for daily volatility.

Coefficient of Determination	DM	JY	SF
CSR $\hat{\sigma}_t^2 = \sum_{i=1}^{96} R_{t.i}^2$	0.216	0.379	0.189
CAR $\hat{\sigma}_t^2 = (\pi/192)(\sum_{i=1}^{96} \|R_{t.i}\|)^2$	0.217	0.424	0.152
GARCH $\hat{\sigma}_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\hat{h}_{t-1}$	0.001	0.003	0.000

Table B.55: Coefficient of determination (R^2) from regressing daily squared return on daily variance estimates for various variance estimators, $R_t^2 = a + b\hat{\sigma}_t + u_t$.

Bera-Jarque Statistic	DM	JY	SF
15 minute returns (n=96) $R_t/\sqrt{\text{CSR}}$	1.86	6.89	1.35
30 minute returns (n=48) $R_t/\sqrt{\text{CSR}}$	1.94	6.91	0.86
60 minute returns (n=24) $R_t/\sqrt{\text{CSR}}$	3.85	6.59	2.91
90 minute returns (n=18) $R_t/\sqrt{\text{CSR}}$	7.35	6.03	4.20
120 minute returns (n=12) $R_t/\sqrt{\text{CSR}}$	8.73	7.52	6.81
180 minute returns (n=8) $R_t/\sqrt{\text{CSR}}$	10.23	8.56	8.68
Daily return. R_t	9.29	113.24	9.01

Table B.56: Bera-Jarque normality test when standardizing daily returns with cumulative squared return daily standard deviation estimates with differing intraday returns (5% c.v is 5.99: 1% c.v. is 9.21).

Correlation between $\sum R_{t,i}^2$ and R_t^2	DM	JY	SF
15 minute returns (n=96)	0.465	0.616	0.435
30 minute returns (n=48)	0.507	0.599	0.465
60 minute returns (n=24)	0.582	0.571	0.573
90 minute returns (n=18)	0.596	0.658	0.571
120 minute returns (n=12)	0.648	0.621	0.623
180 minute returns (n=8)	0.716	0.617	0.701

Table B.57: Correlation between CSR and daily squared return for different levels of aggregation.

Correlation between $\sum R^2_{t,i}$ and $\sum R^2_{t-1,i}$	DM	JY	SF
15 minute returns (n=96)	0.157	0.376	0.208
30 minute returns (n=48)	0.178	0.377	0.092
60 minute returns (n=24)	0.147	0.323	0.062
90 minute returns (n=18)	0.075	0.208	0.066
120 minute returns (n=12)	0.086	0.229	0.052
180 minute returns (n=8)	0.059	0.137	0.042

Table B.58: Serial correlation between CSR and its previous value for different levels of aggregation.

Frequency	DM	JY	SF
15 minute returns (n=96)	0.0483	0.0521	0.0720
30 minute returns (n=48)	0.0428	0.0534	0.0517
60 minute returns (n=24)	0.0649	0.0581	0.0475
90 minute returns (n=18)	0.0705	0.0655	0.0519
120 minute returns (n=12)	0.0623	0.0633	0.0578
180 minute returns (n=8)	0.0757	0.0822	0.0664
daily returns (n=1)	0.0659	0.0472	0.0619

Table B.59: Kolmogorov-Smirnov test of normal distribution for $R_t/\sqrt{\mathrm{CSR_t}}$ (5% critical value is 0.0840; 1% critical value is 0.1007.)

Frequency	DM	JY	SF
15 minute returns (n=96)	0.0581	0.0688	0.0907
30 minute returns (n=48)	0.0573	0.0601	0.0432
60 minute returns (n=24)	0.0878	0.0757	0.0494
90 minute returns (n=18)	0.0958	0.0876	0.0587
120 minute returns (n=12)	0.0948	0.1020	0.0716
180 minute returns (n=8)	0.1204	0.1260	0.0742
daily returns (n=1)	0.0689	0.0681	0.0819

Table B.60: Kolmogorov-Smirnov test of chi-square distribution for $R_t^2/\mathrm{CSR_t}$ (5% critical value is 0.0840; 1% critical value is 0.1007).

Table Notes

Table B.45: Unconditional Variance Estimates. (1) Realized returns is the usual sample variance estimate. (2) Absolute returns variance estimate is the average of the squared of the standard deviation estimate for each day using the absolute value of the return. (3) The cumulative absolute return variance estimate ($n = 96$) is the average of the square of the cumulative absolute standard deviation estimates. (4) The cumulative square return variance estimate ($n = 96$) is the average over all days.

Table B.46: Daily standardized returns. The daily return is standardized by dividing by the square root of the CSR variance estimate for the day.

Table C.61: CAR and CSR estimates are variance estimates.

Table B.49: Descriptive statistics for bias in CAR when comparing the CAR and CSR standard deviation estimates.

Table B.50: Descriptive statistics for bias in CAR when comparing the CAR and CSR variance estimates.

Table B.51: Joint distribution of CSR variance estimates and daily squared returns. Variance estimate from returns is the variance estimate when using the days which satisfy the restrictions for the column. The variance estimate using CSR daily estimates is the average CSR for the days which satisfy the restrictions for the column.

Table B.52: GARCH estimates. The GARCH estimates were calculated using a likelihood function which assumed conditional normality.

Table B.55: Coefficient of determination results. The parameter estimates for the regression are available from the author.

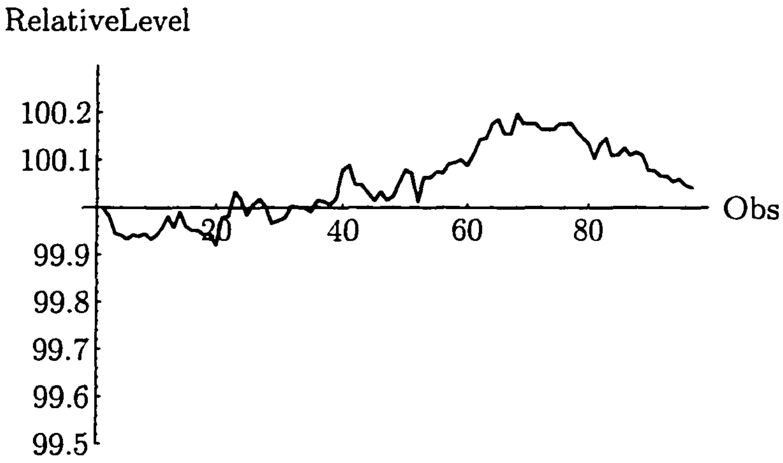

Figure B.23: DM relative price level changes for May 30, 1994. A low volatility day, with a small realized return for the day.

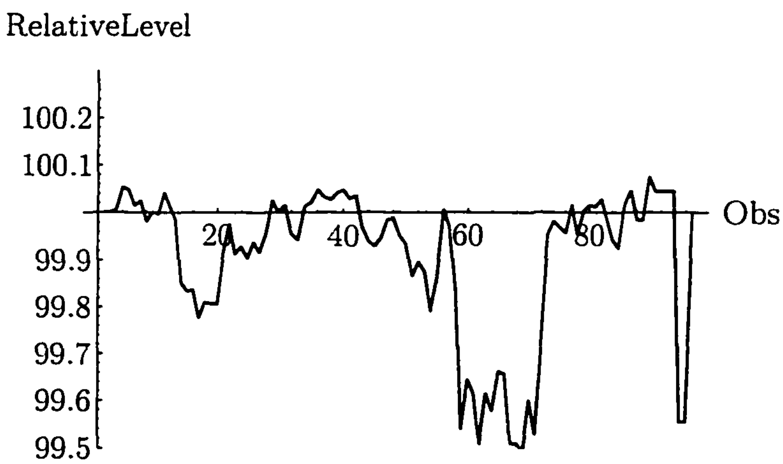

Figure B.24: DM relative price level changes for March 11, 1994. A high volatility day, but small realized return for the day.

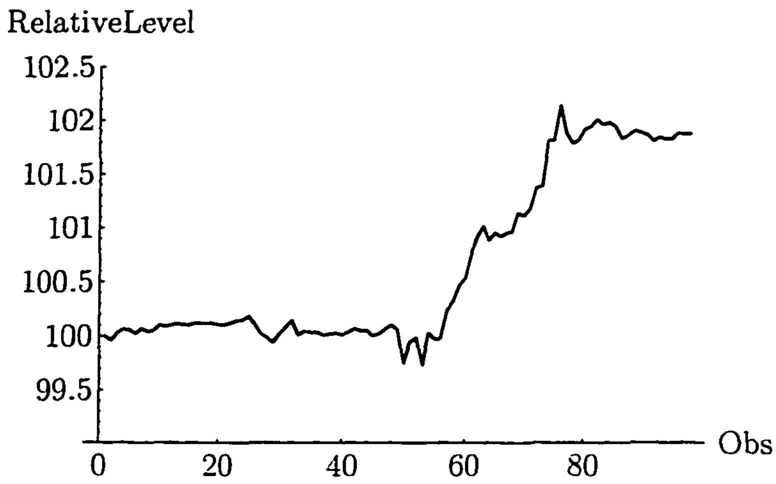

Figure B.25: DM relative price level changes for July 21, 1994. A high volatility day, and a large realized return for the day.

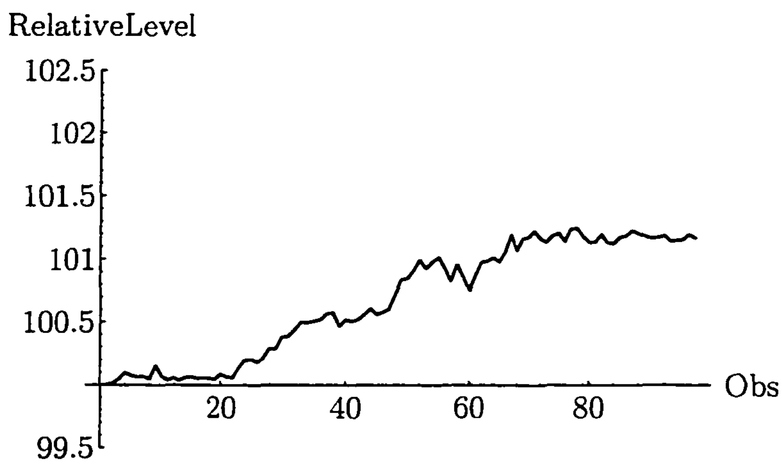

Figure B.26: DM relative price level changes for May 10, 1994. A high volatility day, and a large realized return for the day.

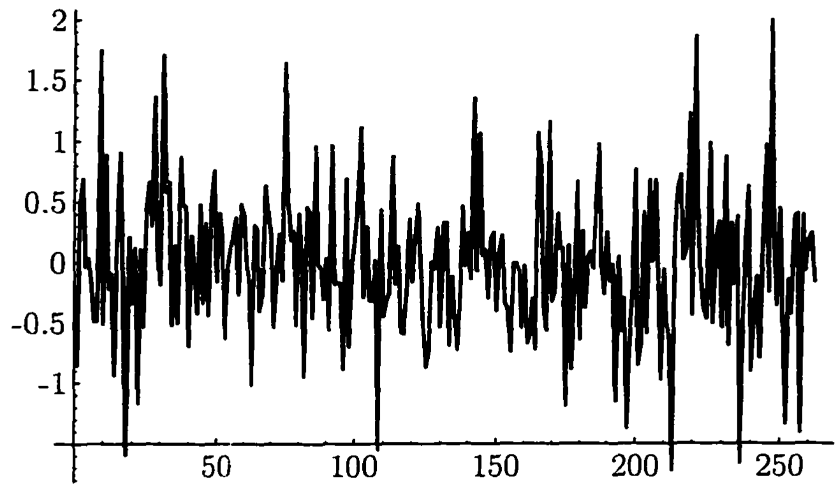

Figure B.27: DM daily realized returns.

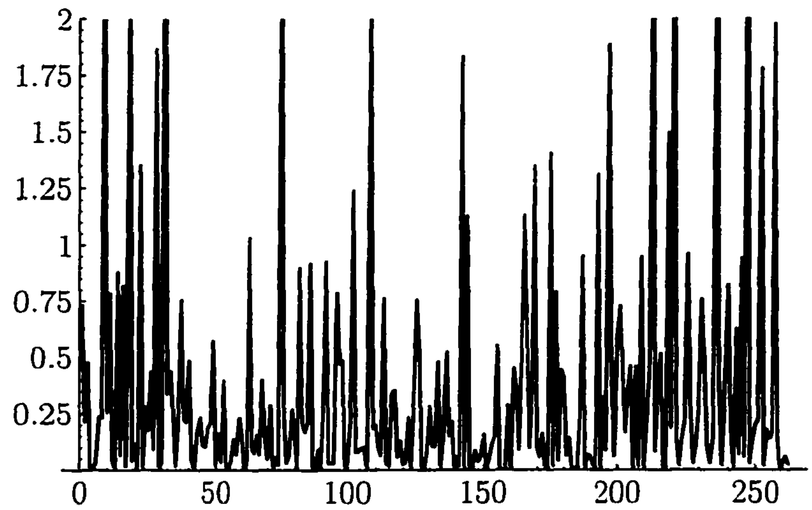

Figure B.28: DM daily squared returns.

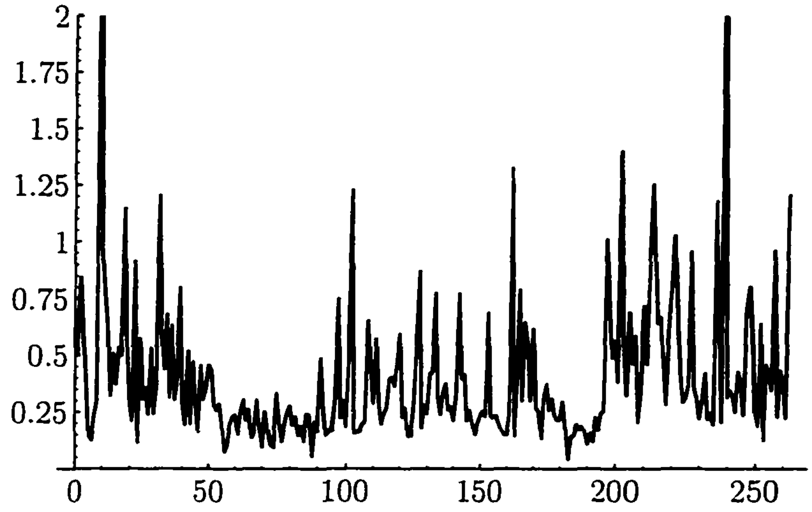

Figure B.29: DM daily cumulative squared return variance estimates.

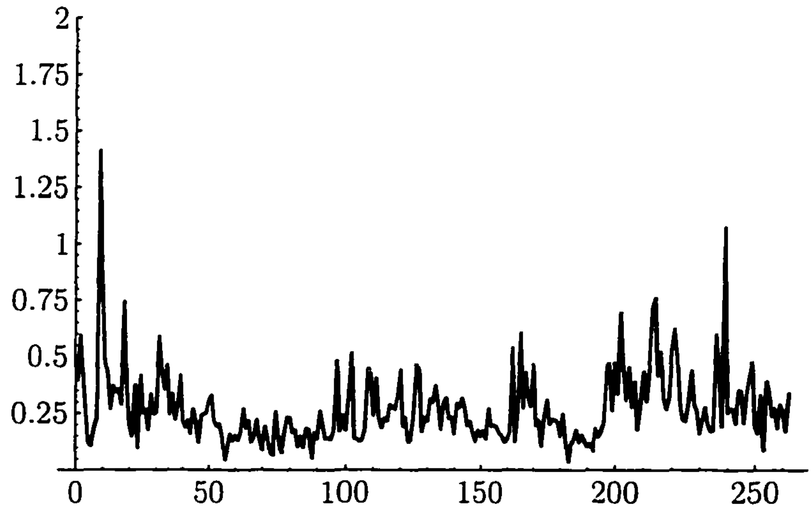

Figure B.30: DM daily cumulative absolute return variance estimates.

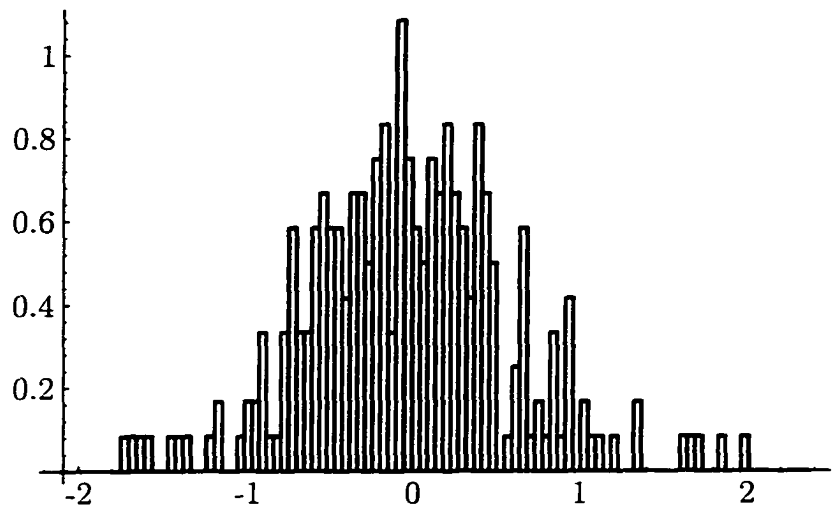

Figure B.31: DM daily realized return histogram.

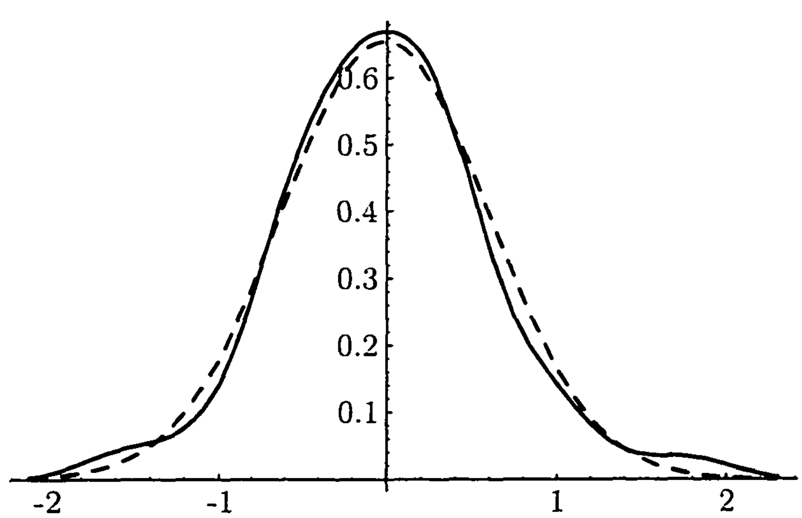

Figure B.32: DM daily realized return density and normal distribution with same mean and variance.

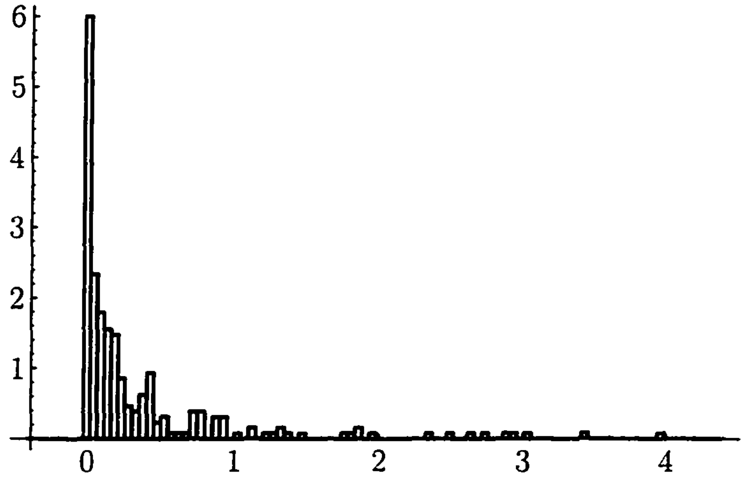

Figure B.33: DM daily squared returns.

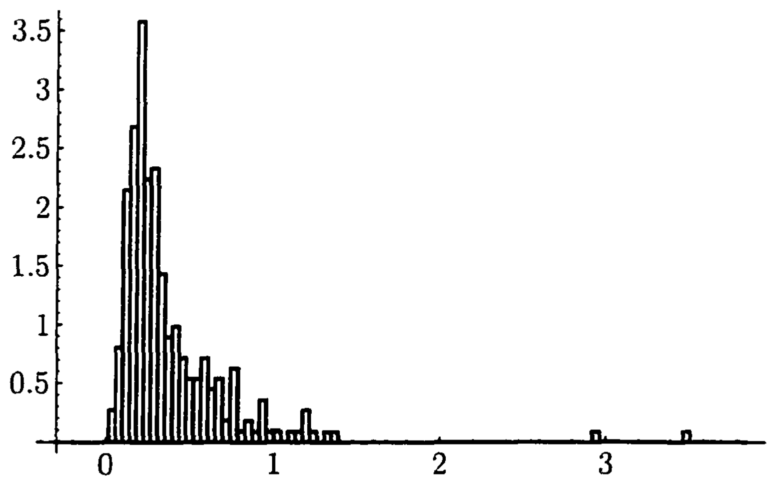

Figure B.34: DM cumulative squared return variance estimator histogram.

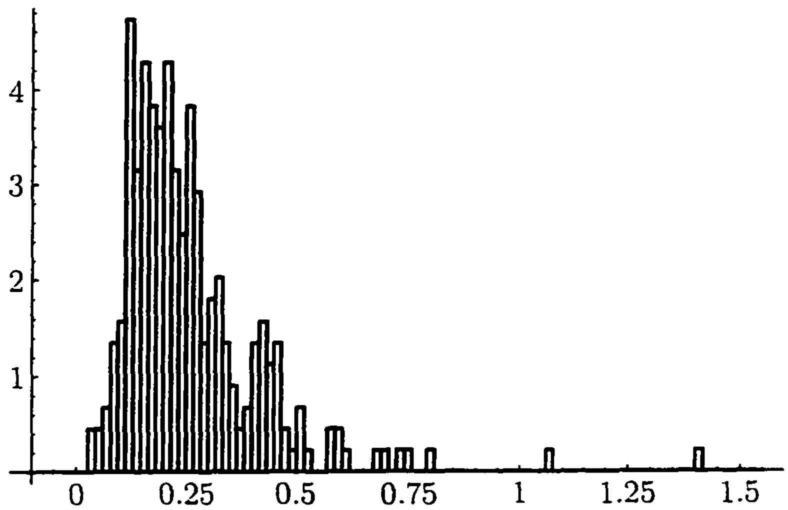

Figure B.35: DM cumulative absolute return variance estimator histogram.

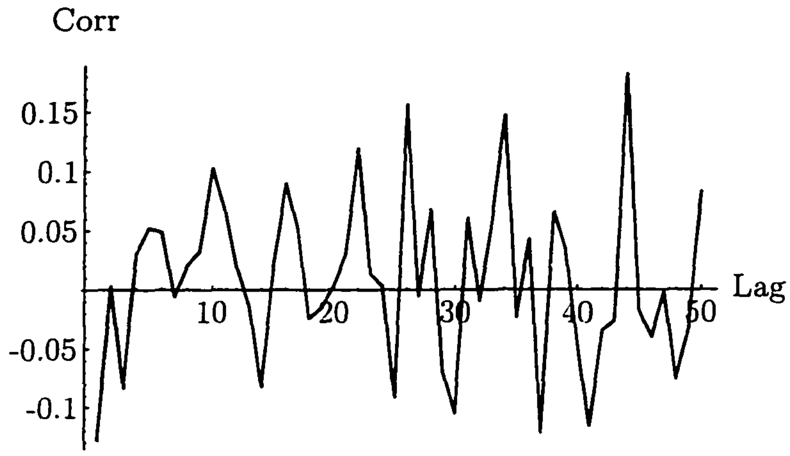

Figure B.36: DM realized squared returns correlogram.

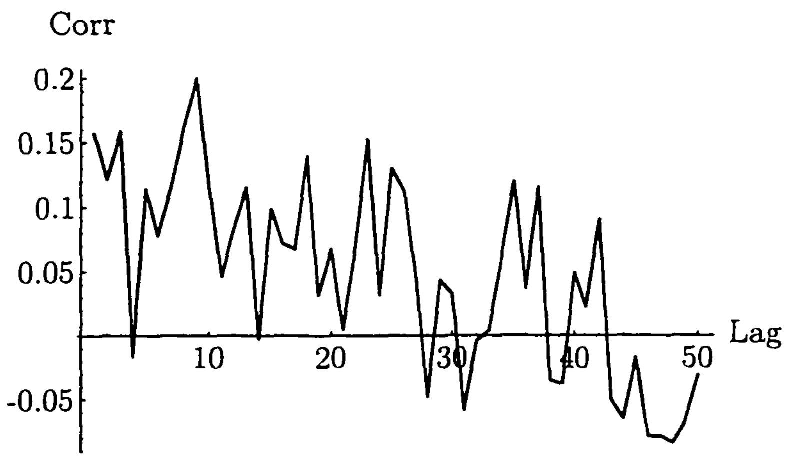

Figure B.37: DM cumulative absolute returns correlogram.

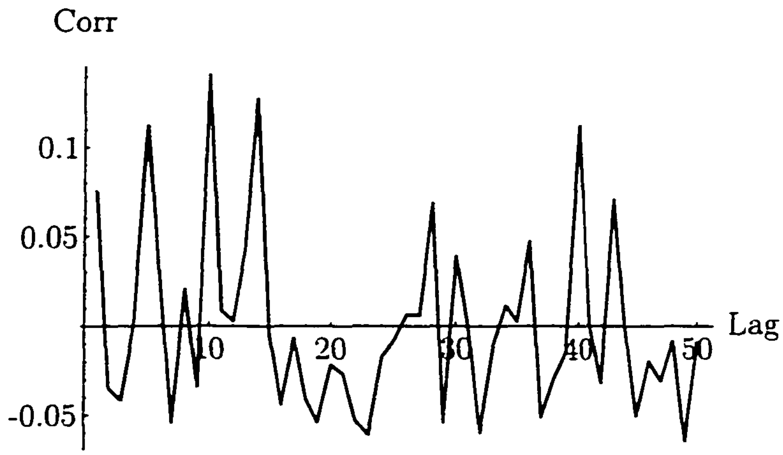

Figure B.38: JY realized squared returns correlogram.

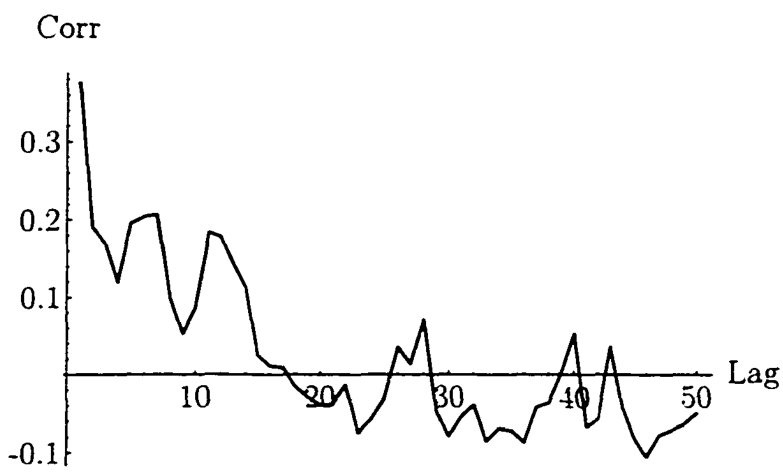

Figure B.39: JY cumulative absolute returns correlogram.

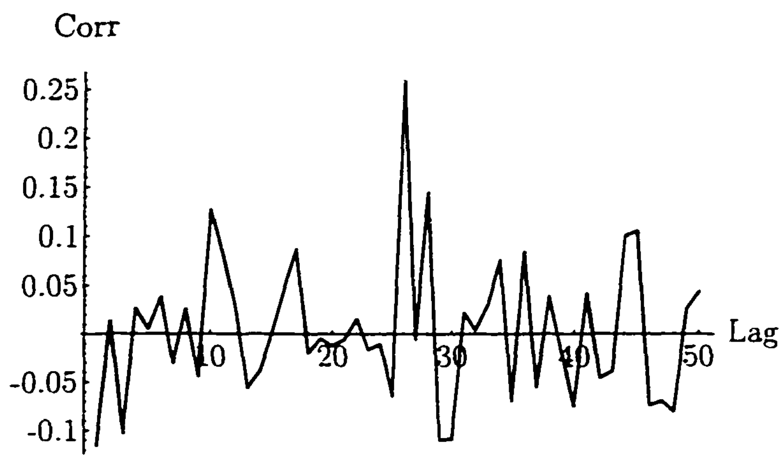

Figure B.40: SF realized squared returns correlogram.

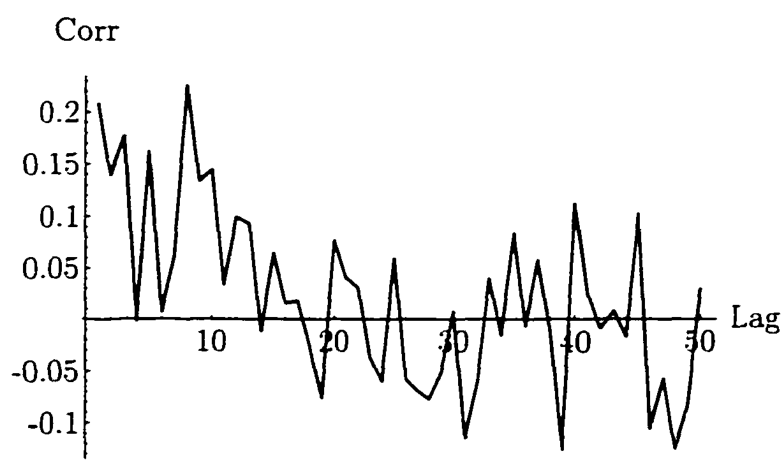

Figure B.41: SF cumulative absolute returns correlogram.

Figure Notes

Figure 1–4: The starting level for the day for each graph has been normalized to 100, thus the difference between the starting and ending value represents the net appreciation or depreciation for the day in terms of logarithmic percentage.

Appendix C: Tables and Figures for Chapter 3

Note: In the following tables, numbers is parentheses are heteroskedastic-consistent standard errors and numbers in brackets are t-ratios.

Correlations	DM	JY	SF
R_t and R_{t-1}	-0.0612	-0.1213	-0.0888
$\|R_t\|$ and $\|R_{t-1}\|$	-0.1257	0.0363	-0.1423
R_t^2 and R_{t-1}^2	-0.1278	0.0752	-0.1155
CSR_t and CSR_{t-1}	0.1569	0.3765	0.2080
R_t^2 and CSR_t	0.4649	0.6156	0.4353

Table C.61: Serial Correlations and contemporaneous correlations (T=262).

Daily Realized Returns	DM	JY	SF
Min	-1.7194	-3.6904	-1.8069
Max	1.9924	1.9046	2.1870
Range	3.7118	5.5949	3.9939
Mean	-0.0152	-0.0154	-0.0373
Median	-0.0299	-0.0587	-0.0354
Variance	0.3721	0.4466	0.4370
Standard Deviation	0.6100	0.6683	0.6610
Mean Deviation	0.4669	0.5060	0.5026
Median Deviation	0.3860	0.4107	0.4124
Skewness	0.2220	-0.4589	0.2080
Excess Kurtosis	0.8079	3.0872	0.8078

Table C.62: Descriptive statistics for realized returns (T=262).

Daily Squared Daily Returns	DM	JY	SF
Min	0	0	0
Max	3.9697	13.6187	4.7832
Range	3.9697	13.6187	4.7832
Mean	0.3709	0.4451	0.4367
Median	0.1483	0.1624	0.1704
Variance	0.3886	1.0289	0.5331
Standard Deviation	0.6234	1.0144	0.7301
Mean Deviation	0.3894	0.4638	0.4577
Median Deviation	0.1309	0.1480	0.1572
Skewness	3.0579	8.9118	3.0306
Excess Kurtosis	10.3977	107.084	10.6857

Table C.63: Descriptive statistics for squared daily returns (T=262).

CSR Variance Estimates	DM	JY	SF
Min	0.0443	0.0826	0.0983
Max	3.5016	4.2520	1.7748
Range	3.4573	4.1694	1.6764
Mean	0.4026	0.5271	0.5089
Median	0.2954	0.4014	0.4406
Variance	0.1281	0.2104	0.0740
Standard Deviation	0.3579	0.4587	0.2720
Mean Deviation	0.2210	0.2798	0.2055
Median Deviation	0.1079	0.1519	0.1466
Skewness	4.4001	3.6680	1.5432
Excess Kurtosis	30.021	19.942	3.4130

Table C.64: Descriptive statistics for cumulative squared returns daily variance estimates (n=96).

$$\begin{aligned} R_t &= \mu + \epsilon_t \\ \sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \end{aligned}$$

Estimate	DM	JY	SF
μ	-0.0156	-0.0038	-0.0360
	(0.0367)	(0.412)	(0.0406)
	[-0.425]	[-0.093]	[-0.886]
ω	0.0089	0.0392	0.0119
	(0.0097)	(0.0319)	(0.0193)
	[0.923]	[1.229]	[0.615]
α	0.0166	0.0393	0.0063
	(0.0148)	(0.0367)	(0.0147)
	[1.124]	[1.070]	[0.430]
β	0.9587	0.8696	0.9666
	(0.0309)	(0.0841)	(0.0461)
	[31.068]	[10.339]	[20.981]

Table C.65: GARCH estimates (using realized returns, T=262).

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\sigma_t^2 &= \omega + \alpha \left(\sum_{i=1}^{n} R_{t-1,i}^2 \right) + \beta \sigma_{t-1}^2
\end{aligned}
$$

Estimate	DM	JY	SF
μ	-0.0235 (0.0367) [-0.640]	-0.0181 (0.0390) [-0.464]	-0.0438 (0.0413) [-1.060]
ω	0.0116 (0.0135) [0.864]	0.0104 (0.0181) [0.575]	0.0107 (0.0221) [0.484]
α	0.0369 (0.0268) [1.377]	0.1014 (0.0468) [2.166]	0.0208 (0.0253) [0.819]
β	0.9279 (0.0557) [16.645]	0.8555 (0.0630) [13.586]	0.9516 (0.0573) [16.619]

Table C.66: GARCH estimates using lagged cumulative squared returns instead of lagged squared error (T=262).

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\sigma_t^2 &= \omega + \alpha_0 \left(\sum_{i=1}^{n} R_{t-1,i}^2 \right) + \alpha_1 \left(\sum_{i<j}^{n} 2R_{t-1,i} R_{t-1,j} \right) + \beta \sigma_{t-1}^2
\end{aligned}
$$

Estimate	DM	JY	SF
μ	-0.0235 (0.0367) [-0.640]	-0.0181 (0.0390) [-0.464]	-0.0438 (0.0413) [-1.060]
ω	0.0116 (0.0135) [0.864]	0.0104 (0.0181) [0.575]	0.0107 (0.0221) [0.484]
α_0	0.0369 (0.0268) [1.377]	0.1014 (0.0468) [2.166]	0.0208 (0.0253) [0.819]
α_1	0 - -	0 - -	0 - -
β	0.9279 (0.0557) [16.645]	0.8555 (0.0630) [13.586]	0.9516 (0.0573) [16.619]

Table C.67: GARCH estimates using lagged cumulative squared returns and cross terms (T=262).

$$
\begin{aligned}
R_t &= \mu + \epsilon_t \\
\epsilon_t &\sim N(0, \sigma_t^2) \\
\sigma_t^2 &= \left(\sum_{i=1}^{n} R_{t,i}^2 \right) \\
\sigma_{t|t-1}^2 &= \omega + \alpha \sigma_{t-1}^2 + \beta \sigma_{t-1|t-2}^2 .
\end{aligned}
$$

Estimate	DM	JY	SF
ω	0.0153	0.0388	0.0352
	(0.0198)	(0.0458)	(0.0585)
	[0.772]	[0.847]	[0.602]
α	0.1213	0.1741	0.1071
	(0.0907)	(0.1050)	(0.1230)
	[1.337]	[1.659]	[0.871]
β	0.8386	0.7479	0.8206
	(0.1168)	(0.1611)	(0.2016)
	[7.183]	[4.642]	[4.071]

Table C.68: Conditional heteroskedastic CSR volatility model (T=262).

$$
\begin{aligned}
\hat{\sigma}_t^2 &= \sum_{i=1}^{n} R_{t,i}^2 \\
\hat{\sigma}_t^2 &= \psi_0 + \psi_1 \hat{\sigma}_{t-1}^2 + \epsilon_t + \psi_2 \epsilon_{t-1}
\end{aligned}
$$

Estimate	DM	JY	SF
ψ_0	0.0162	0.0691	0.0343
	(0.0120)	(0.0460)	(0.0185)
	[1.354]	[1.503]	[1.851]
ψ_1	0.9564	0.8628	0.0185
	(0.0302)	(0.0867)	(0.0369)
	[-31.678]	[9.957]	[25.203]
ψ_2	-0.8749	-0.6627	-0.8379
	(0.0520)	(0.1402)	(0.0589)
	[-16.841]	[-4.725]	[-14.229]

Table C.69: ARMA estimates on cumulative squared returns (T=262).

Correlations	DM	JY	SF
R_t^2 and CSR	0.465	0.616	0.435
R_t^2 and GARCH	0.035	0.050	-0.003
R_t^2 and GARCH-CSR	0.106	0.099	0.002
R_t^2 and ARMA	0.074	0.095	0.031
R_t^2 and Hetero-CSR	0.089	0.097	0.025
CSR and R_t^2	0.465	0.616	0.435
CSR and GARCH	0.224	0.352	0.167
CSR and GARCH-CSR	0.263	0.354	0.203
CSR and ARMA	0.261	0.389	0.275
CSR and Hetero-CSR	0.254	0.383	0.277
GARCH and R_t^2	0.035	0.050	-0.003
GARCH and CSR	0.224	0.352	0.167
GARCH and GARCH-CSR	0.792	0.777	0.253
GARCH and ARMA	0.745	0.789	0.413
GARCH and Hetero-CSR	0.704	0.805	0.416
GARCH-CSR and R_t^2	0.106	0.099	0.002
GARCH-CSR and CSR	0.263	0.354	0.203
GARCH-CSR and GARCH	0.792	0.777	0.253
GARCH-CSR and ARMA	0.961	0.911	0.775
GARCH-CSR and Hetero-CSR	0.924	0.957	0.740
ARMA and R_t^2	0.074	0.095	0.031
ARMA and CSR	0.261	0.389	0.275
ARMA and GARCH	0.745	0.789	0.413
ARMA and GARCH-CSR	0.961	0.911	0.775
ARMA and Hetero-CSR	0.993	0.989	0.998
Hetero-CSR and R_t^2	0.089	0.097	0.025
Hetero-CSR and CSR	0.254	0.383	0.277
Hetero-CSR and GARCH	0.704	0.805	0.416
Hetero-CSR and GARCH-CSR	0.924	0.957	0.740
Hetero-CSR and ARMA	0.993	0.989	0.998

Table C.70: Correlations of in-sample variance forecasts (T=262).

Currency	Model	MSE	MAE	LL	HMSE	GMLE
DM	Constant	0.3871	0.3894	9.8666	2.8138	0.0082
DM	GARCH	0.3870	0.3862	9.7135	2.9247	0.0075
DM	GARCH-CSR	0.3827	0.3869	9.7237	2.6873	-0.0066
DM	ARMA	0.3861	0.3960	9.8529	2.6095	0.0029
DM	Hetero-CSR	0.3906	0.3989	9.8504	1.8056	0.0146
JY	Constant	1.0250	0.4638	9.5624	5.1732	0.1906
JY	GARCH	1.0238	0.4597	9.4724	5.0515	0.1832
JY	GARCH-CSR	1.0197	0.4681	9.3993	4.3256	0.1512
JY	ARMA	1.0281	0.3946	9.8393	3.8105	0.1702
JY	Hetero-CSR	1.0298	0.4972	9.8246	3.5841	0.1659
SF	Constant	0.5310	0.4577	24.628	2.7847	0.1715
SF	GARCH	0.5314	0.4569	24.618	2.7526	0.1724
SF	GARCH-CSR	0.5382	0.4779	24.752	2.6094	0.1814
SF	ARMA	0.5444	0.4997	25.174	2.3212	0.1939
SF	Hetero-CSR	0.5462	0.5008	25.188	2.3602	0.1970

Table C.71: In-sample statistical evaluation of volatility forecasts using squared return as measure of true volatility.

Currency	Model	MSE	MAE	LL	HMSE	GMLE
DM	Constant	0.1276	0.2210	0.4594	0.7875	0.0901
DM	GARCH	0.1246	0.1995	0.3881	0.9049	0.0794
DM	GARCH-CSR	0.1204	0.1935	0.3652	0.8107	0.0646
DM	ARMA	0.1190	0.1941	0.3642	0.7597	0.0579
DM	Hetero-CSR	0.1201	0.1933	0.3609	0.7835	0.0579
JY	Constant	0.2096	0.2798	0.4555	0.7545	0.3597
JY	GARCH	0.1972	0.2372	0.3438	0.9891	0.3451
JY	GARCH-CSR	0.1882	0.2276	0.3034	0.9290	0.3260
JY	ARMA	0.1779	0.2397	0.3195	0.7053	0.3089
JY	Hetero-CSR	0.1790	0.2395	0.3164	0.6965	0.3080
SF	Constant	0.0737	0.2055	0.2677	0.2845	0.3245
SF	GARCH	0.0779	0.1953	0.2463	0.4112	0.3343
SF	GARCH-CSR	0.0733	0.1942	0.2408	0.3504	0.3249
SF	ARMA	0.0681	0.1928	0.2372	0.2727	0.3131
SF	Hetero-CSR	0.0681	0.1925	0.2368	0.2731	0.3129

Table C.72: In-sample statistical evaluation of volatility forecasts using CSR as measure of true volatility.

Currency	Model	BJ	R^2 (CSR)
DM	Constant	9.29	-
DM	GARCH	8.06	0.050
DM	GARCH-CSR	6.22	0.069
DM	ARMA	8.17	0.067
DM	Hetero-CSR	10.67	0.064
JY	Constant	113.24	-
JY	GARCH	88.47	0.124
JY	GARCH-CSR	64.09	0.125
JY	ARMA	106.64	0.151
JY	Hetero-CSR	87.73	0.147
SF	Constant	9.01	-
SF	GARCH	8.66	0.028
SF	GARCH-CSR	9.01	0.041
SF	ARMA	10.23	0.076
SF	Hetero-CSR	10.82	0.077

Table C.73: More in-sample statistical evaluation of volatility forecasts.

	Half Life			Mean Lag		
Model	DM	JY	SF	DM	JY	SF
GARCH	27.71	7.26	25.23	16.27	3.31	6.96
GARCH-CSR	19.34	15.73	24.77	14.54	16.28	15.57
ARMA	15.55	4.70	9.45	14.94	4.32	7.98
Hetero-CSR	16.94	8.54	9.24	18.74	8.85	8.26
GARCH (Jorion)	16.35	5.59	15.51	15.56	3.89	12.44

Table C.74: Half life, mean lag, and median lag.

Currency	DM	JY	SF
Constant Variance	2.8138	5.1732	2.7847
GARCH	2.9247	5.0515	2.7526
GARCH-CSR	2.6873	4.3256	2.6094
Hetero-CSR	1.8056	3.5841	2.3602
Actual CSR	1.3393	0.8313	1.2023

Table C.75: HMSE measures (based on squared return) for differing models.

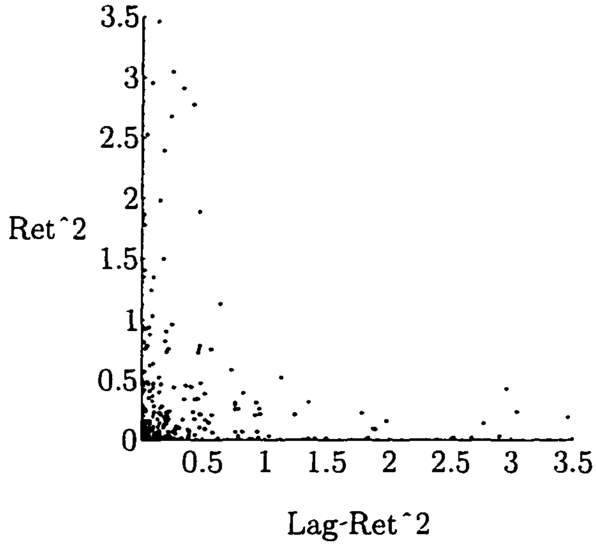

Figure C.42: DM squared realized returns versus lagged squared realized returns.

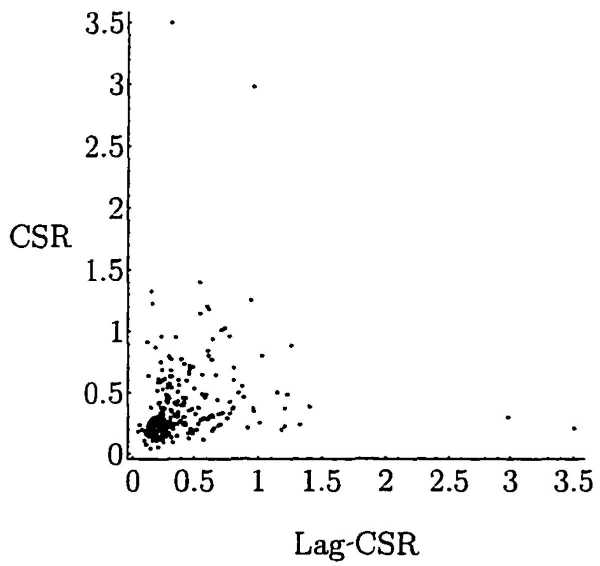

Figure C.43: DM cumulative squared returns versus lagged cumulative squared returns.

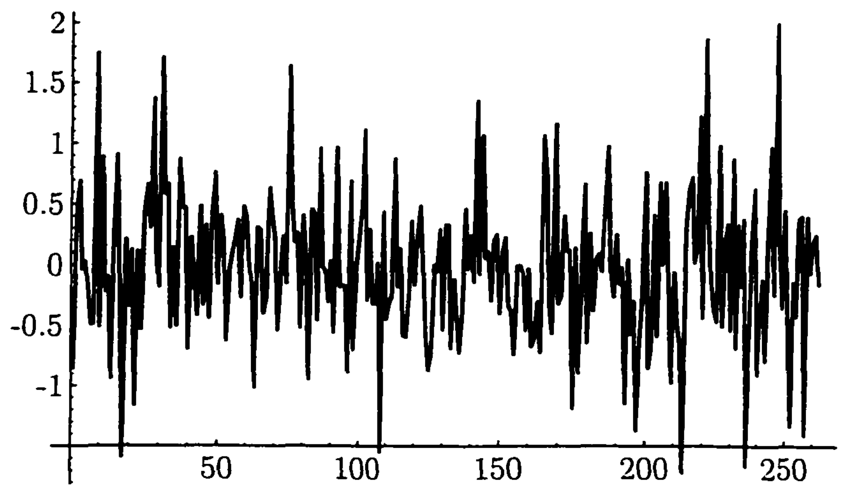

Figure C.44: DM daily realized returns.

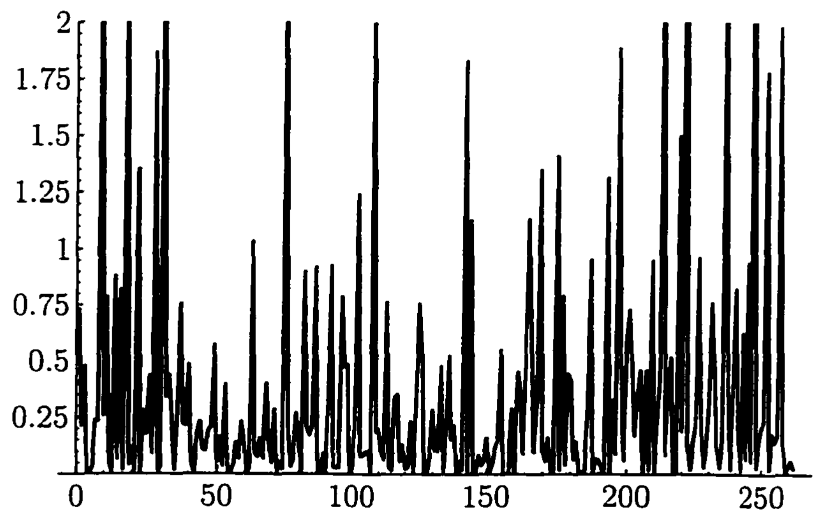

Figure C.45: DM daily squared returns.

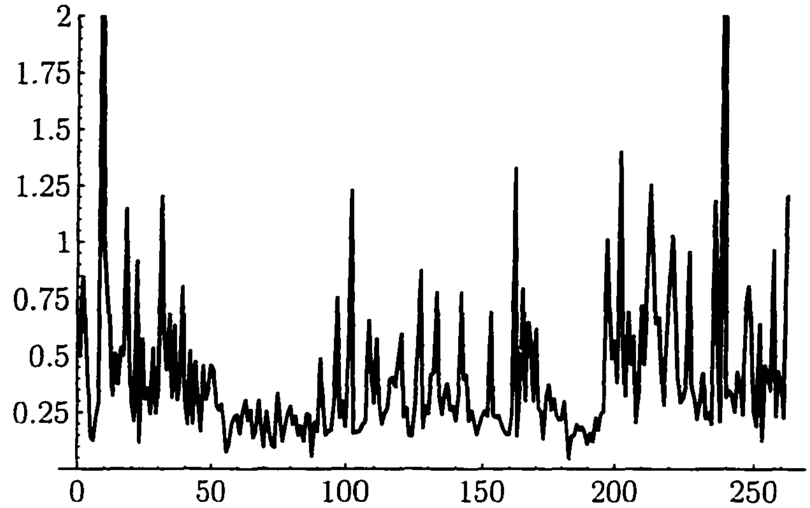

Figure C.46: DM daily cumulative squared return variance estimates.

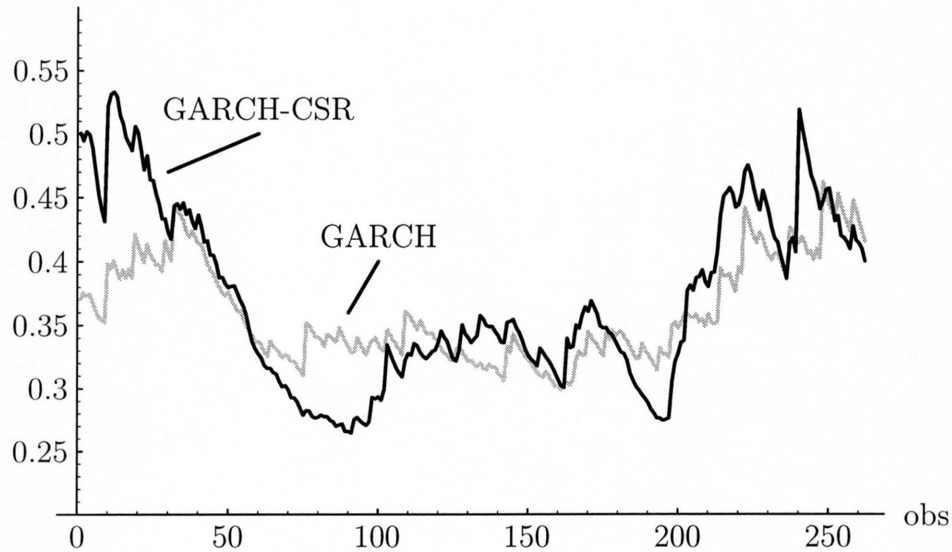

Figure C.47: DM GARCH and GARCH-CSR forecast variances.

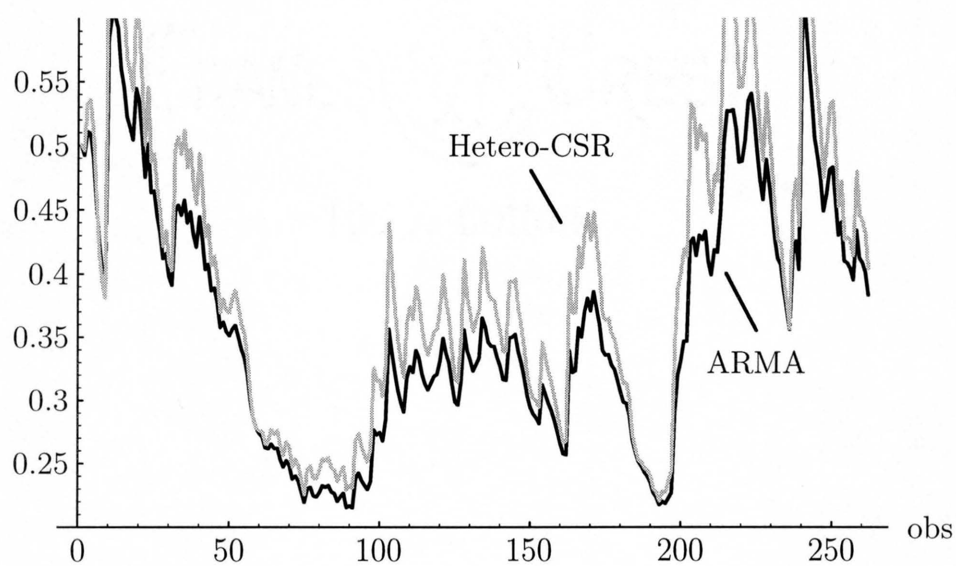

Figure C.48: DM ARMA and Hetero-CSR forecast variances.

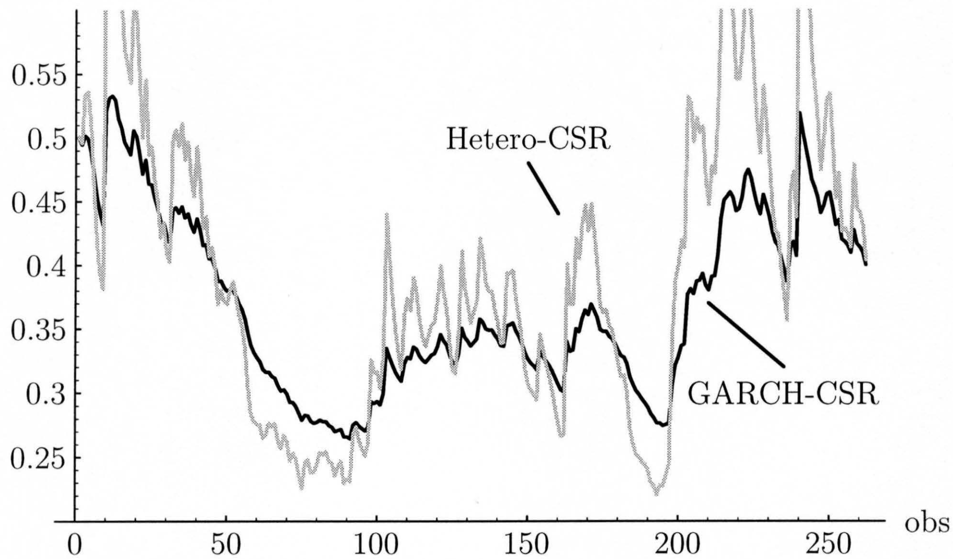

Figure C.49: DM GARCH-CSR and Hetero-CSR forecast variances.

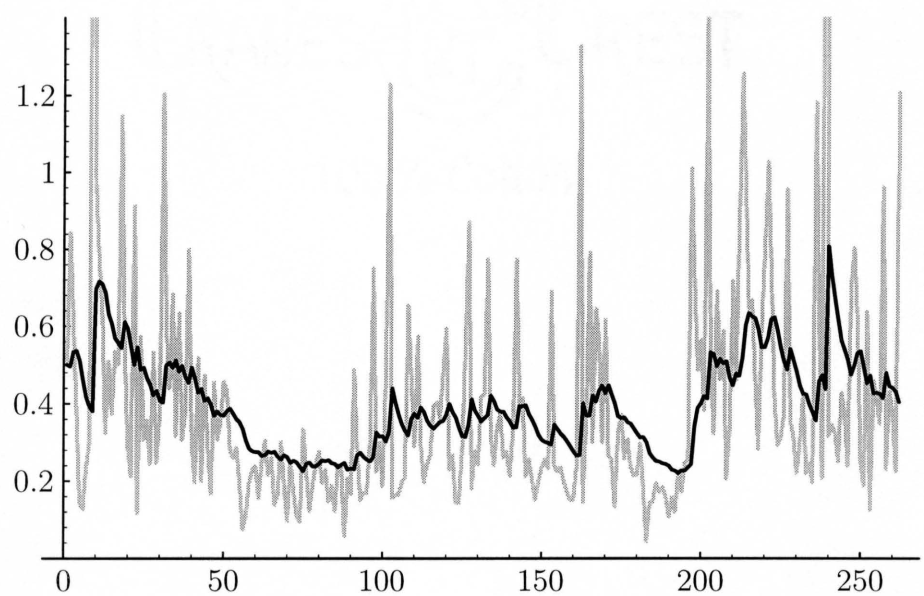

Figure C.50: DM CSR variance estimates and Hetero-CSR forecast variances.

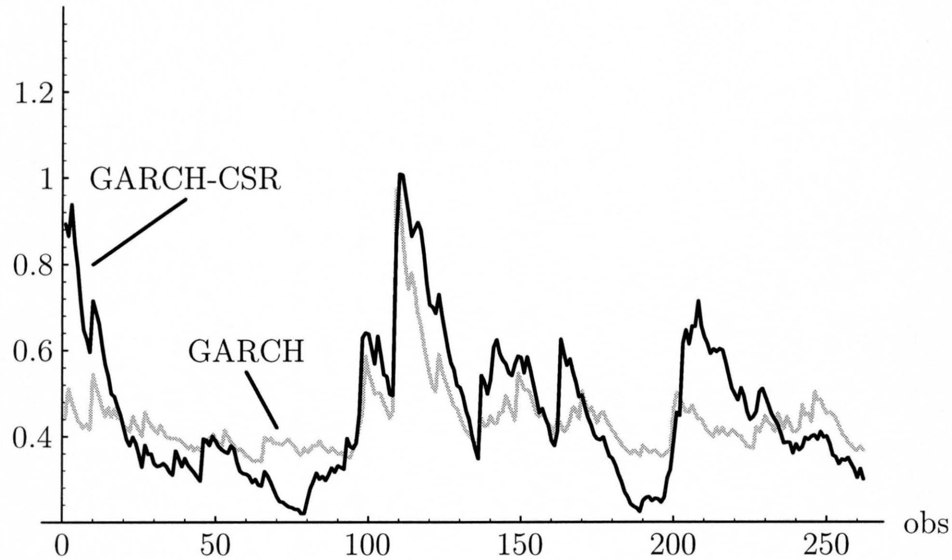

Figure C.51: JY GARCH and GARCH-CSR forecast variances.

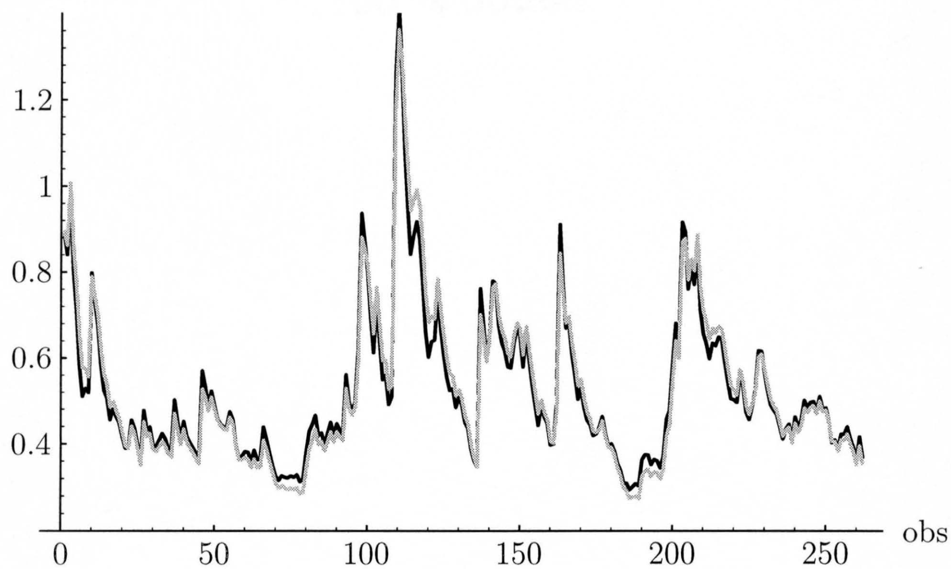

Figure C.52: JY ARMA and Hetero-CSR forecast variances.

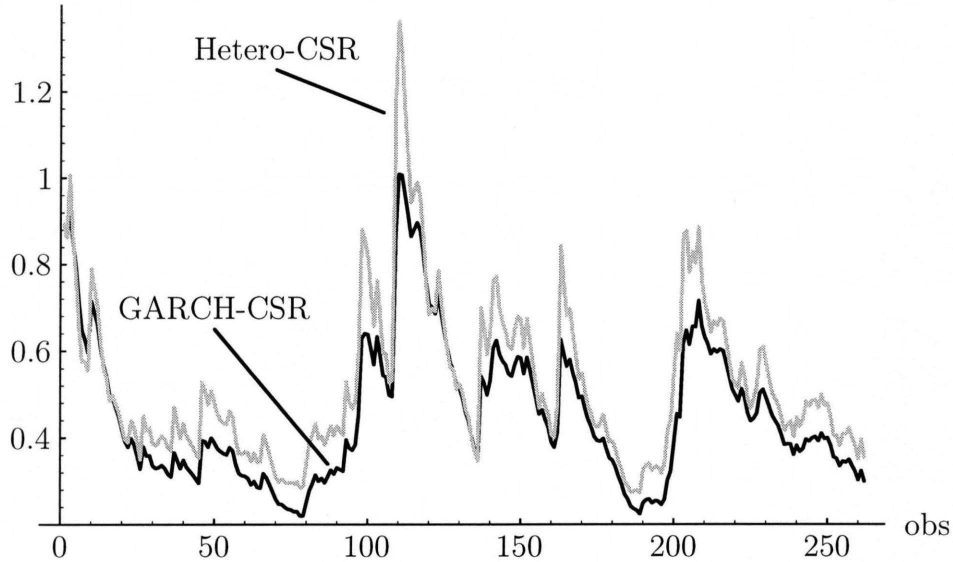

Figure C.53: JY GARCH-CSR and Hetero-CSR forecast variances.

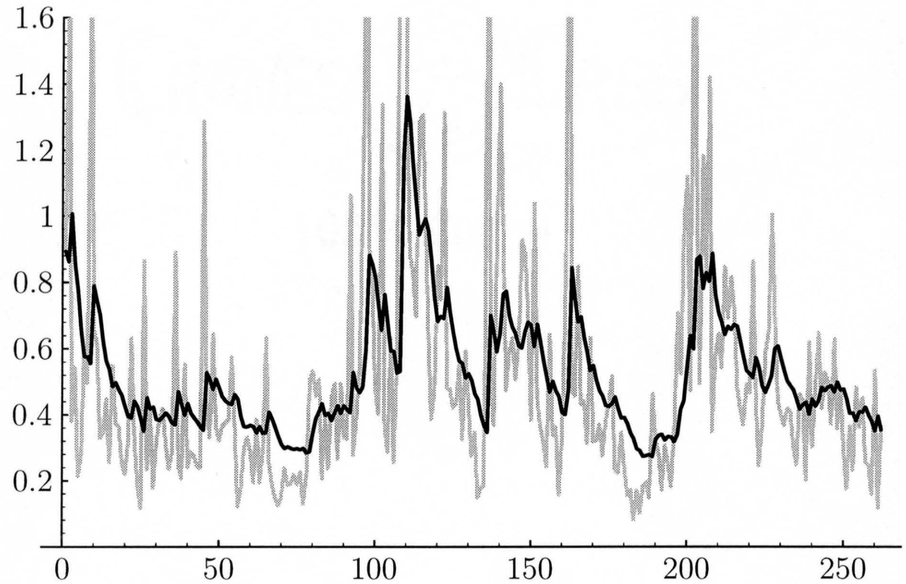

Figure C.54: JY CSR variance estimates and Hetero-CSR forecast variances.

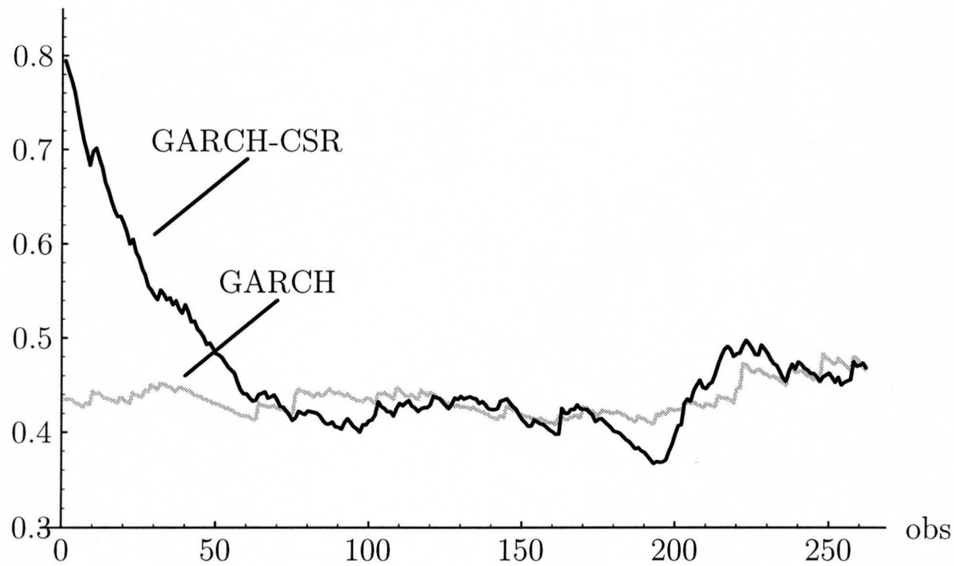

Figure C.55: SF GARCH and GARCH-CSR forecast variances.

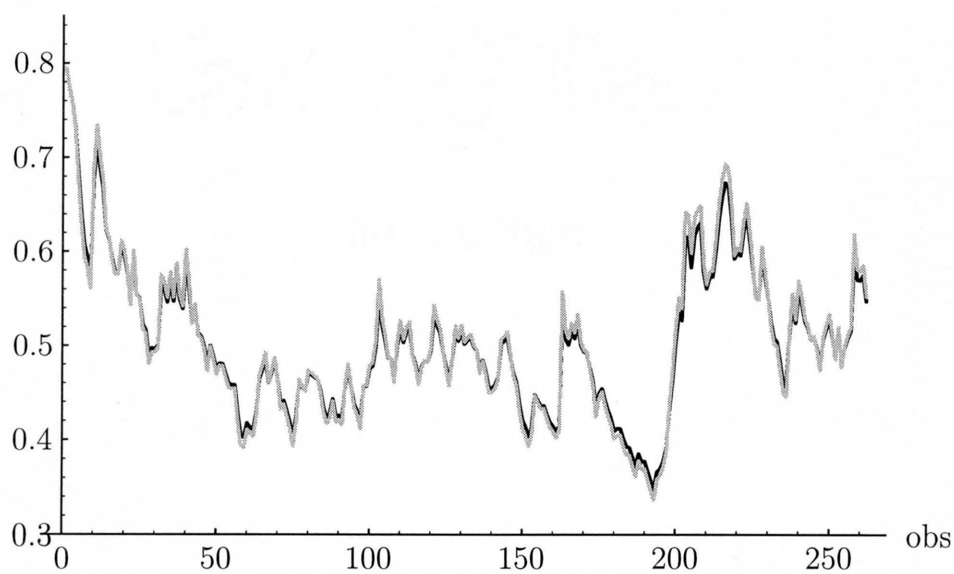

Figure C.56: SF ARMA and Hetero-CSR forecast variances.

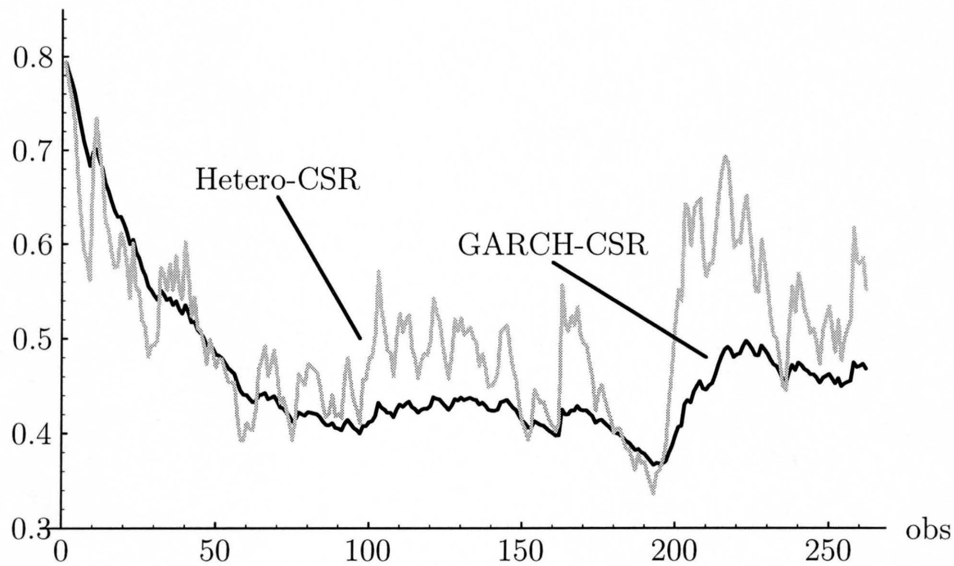

Figure C.57: SF GARCH-CSR and Hetero-CSR forecast variances.

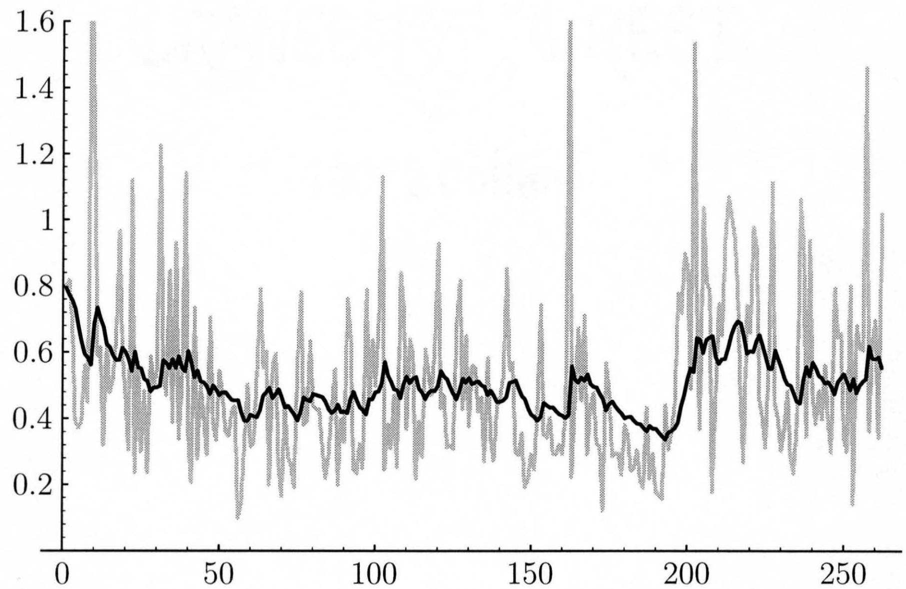

Figure C.58: SF CSR variance estimates and Hetero-CSR forecast variances.

Appendix D: Tables and Figures for Chapter 4

Model	Average Estimates			Average HMSE	
	ω	α	β	Sq Ret	CSR
True	0.0153	0.1513	0.8086	-	-
GARCH					
$n = 262$	0.2372	0.0299	0.3141	1.9790	0.0589
$n = 262 \times 5$	0.2657	0.0171	0.2452	2.0158	0.0506
GARCH-CSR					
$n = 262$	0.0398	0.2370	0.6536	1.9444	0.0521
$n = 262 \times 5$	0.0297	0.1873	0.7300	1.9849	0.0418
Hetero-CSR					
$n = 262$	0.0274	0.1484	0.7755	1.9916	0.0374
$n = 262 \times 5$	0.0167	0.1491	0.8046	1.9950	0.0378

Table D.76: Simulations 1. Mean parameter estimates and average HMSE measures from simulations.

Model	Average Parameter ω	α	β	Average HMSE Sq Ret	CSR
True	0.1600	0.2000	0.4000	-	-
GARCH					
$n = 262$	0.2124	0.0278	0.4302	1.9632	0.0533
$n = 262 \times 5$	0.2313	0.0152	0.3950	1.9949	0.0429
GARCH-CSR					
$n = 262$	0.1220	0.2898	0.3941	1.9433	0.0514
$n = 262 \times 5$	0.1525	0.2269	0.3846	1.9850	0.0414
Hetero-CSR					
$n = 262$	0.1781	0.1940	0.3520	1.9884	0.0372
$n = 262 \times 5$	0.1620	0.1968	0.3905	1.9940	0.0377

Table D.77: Simulations 2. Mean parameter estimates and average HMSE measures from simulations.

Number of Better Forecasts	HMSE-Sq Ret GARCH	GARCH-CSR	HMSE-CSR GARCH	GARCH-CSR
Simulation 1: $(\omega, \alpha, \beta) = (0.02, 0.15, 0.81)$				
$n = 262$	356	638	207	775
$n = 262 \times 5$	185	815	19	981
Simulation 2: $(\omega, \alpha, \beta) = (0.16, 0.20, 0.40)$				
$n = 262$	393	602	392	594
$n = 262 \times 5$	340	660	218	782

Table D.78: Pairwise comparisons between GARCH and GARCH-CSR forecasts. The number of times each model produced better forecasts. (Note: Cases do not always add up to 1000 due to forecasts occasionally being equal.)

Number of Better Forecasts	HMSE-Sq Ret		HMSE-CSR	
	GARCH	Hetero-CSR	GARCH	Hetero-CSR
Simulation 1: $(\omega, \alpha, \beta) = (0.02, 0.15, 0.81)$				
$n = 262$	470	530	4	996
$n = 262 \times 5$	447	553	0	1000
Simulation 2: $(\omega, \alpha, \beta) = (0.16, 0.20, 0.40)$				
$n = 262$	487	513	18	982
$n = 262 \times 5$	488	512	7	993

Table D.79: Pairwise comparisons between GARCH and Hetero-CSR forecasts. The number of times each model produced better forecasts.

Number of Better Forecasts	HMSE-Sq Ret		HMSE-CSR	
	GARCH-CSR	Hetero-CSR	GARCH-CSR	Hetero-CSR
Simulation 1: $(\omega, \alpha, \beta) = (0.02, 0.15, 0.81)$				
$n = 262$	518	482	72	928
$n = 262 \times 5$	519	481	172	828
Simulation 2: $(\omega, \alpha, \beta) = (0.16, 0.20, 0.40)$				
$n = 262$	508	492	59	941
$n = 262 \times 5$	506	494	181	819

Table D.80: Pairwise comparisons between GARCH-CSR and Hetero-CSR forecasts. The number of times each model produced better forecasts.

Treatment	Alternative Forecast	Base Forecast	Alternate Model's Average Percent Reduction in HMSE-Ret Sq	HMSE-CSR
Simulation 1:				
$n = 262$	GARCH-CSR	GARCH	1.61	9.44
$n = 262 \times 5$			1.50	17.00
Simulation 2:				
$n = 262$	GARCH-CSR	GARCH	0.89	1.25
$n = 262 \times 5$			0.48	3.38
Simulation 1:				
$n = 262$	Hetero-CSR	GARCH	-0.73	30.85
$n = 262 \times 5$			0.98	24.41
Simulation 2:				
$n = 262$	Hetero-CSR	GARCH	-1.37	23.97
$n = 262 \times 5$			0.01	11.20
Simulation 1:				
$n = 262$	Hetero-CSR	GARCH-CSR	-2.48	22.24
$n = 262 \times 5$			-0.54	8.62
Simulation 2:				
$n = 262$	Hetero-CSR	GARCH-CSR	-2.32	21.82
$n = 262 \times 5$			-0.48	7.90

Table D.81: Pairwise comparisons between forecasts. The average percentage improvement in forecasts.

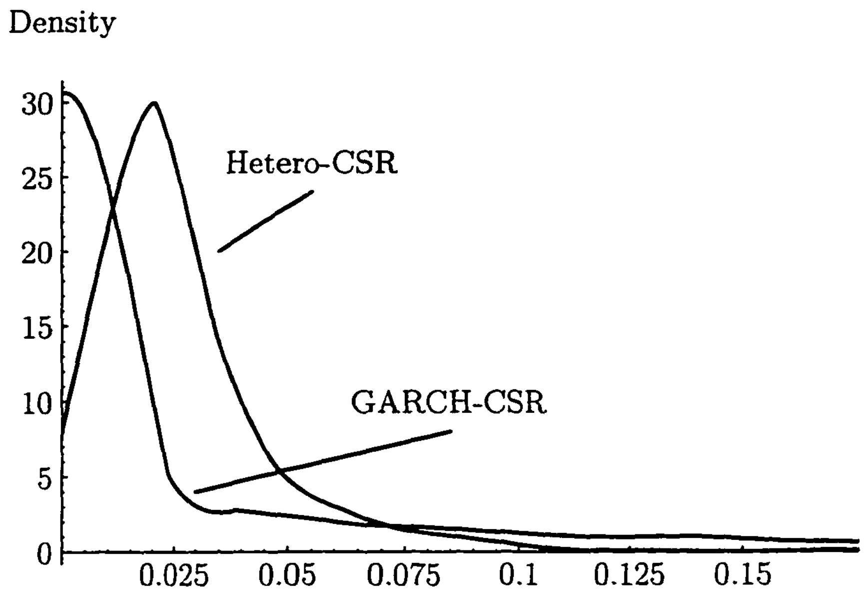

Figure D.59: Simulations 1 ($T = 262$), Gaussian kernel density estimates for ω for Hetero-CSR and GARCH-CSR, (true value is 0.0153).

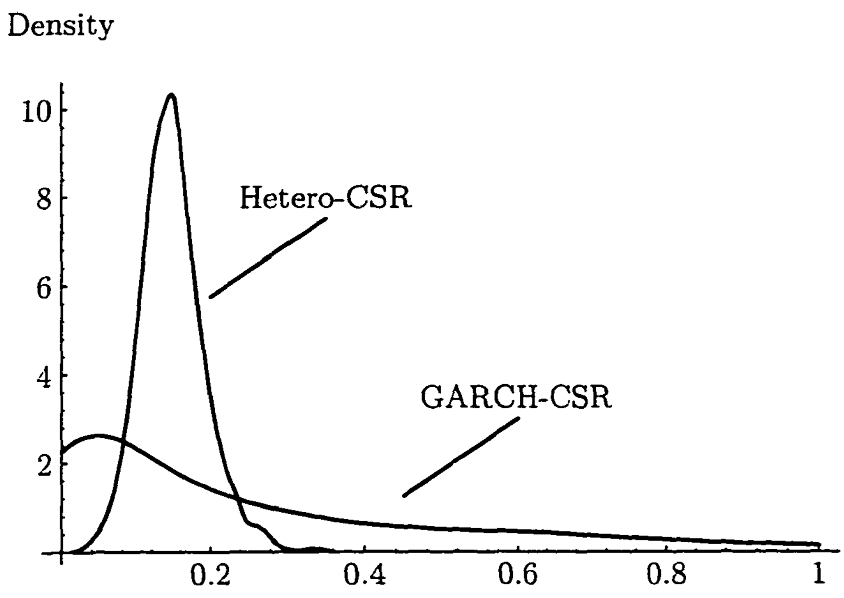

Figure D.60: Simulations 1 ($T = 262$), Gaussian kernel density estimates for α for Hetero-CSR and GARCH-CSR, (true value is 0.1513).

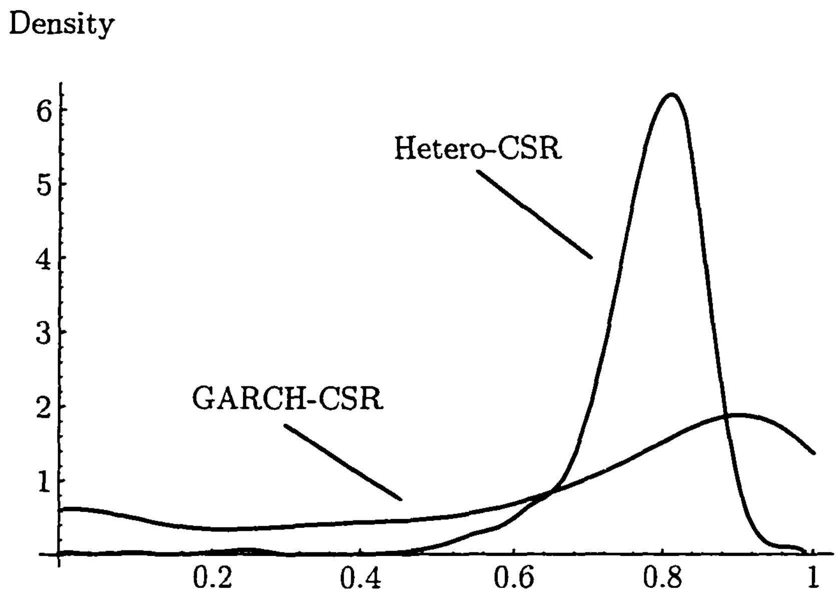

Figure D.61: Simulations 1 ($T = 262$), Gaussian kernel density estimates for β for Hetero-CSR and GARCH-CSR, (true value is 0.8083).

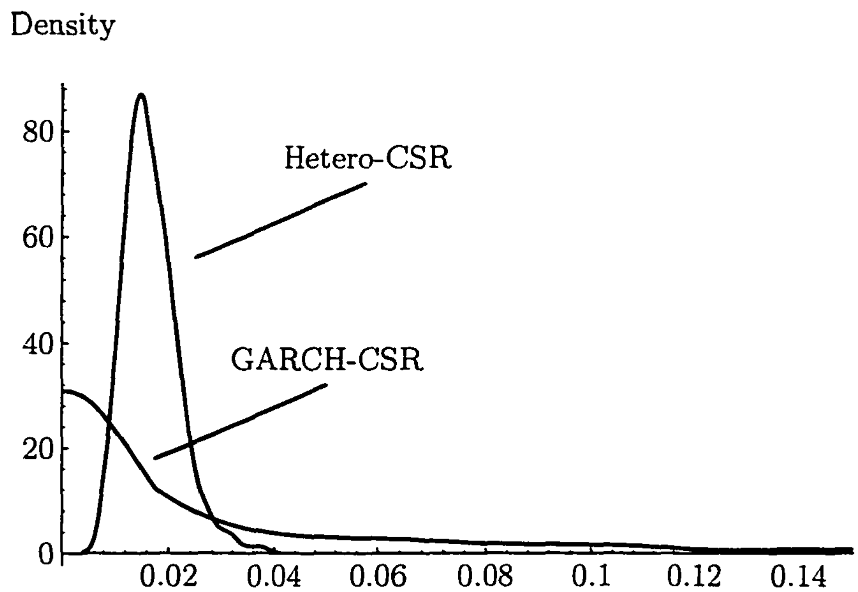

Figure D.62: Simulations 1 ($T = 262 \times 5$), Gaussian kernel density estimates for ω for Hetero-CSR and GARCH-CSR, (true value is 0.0153).

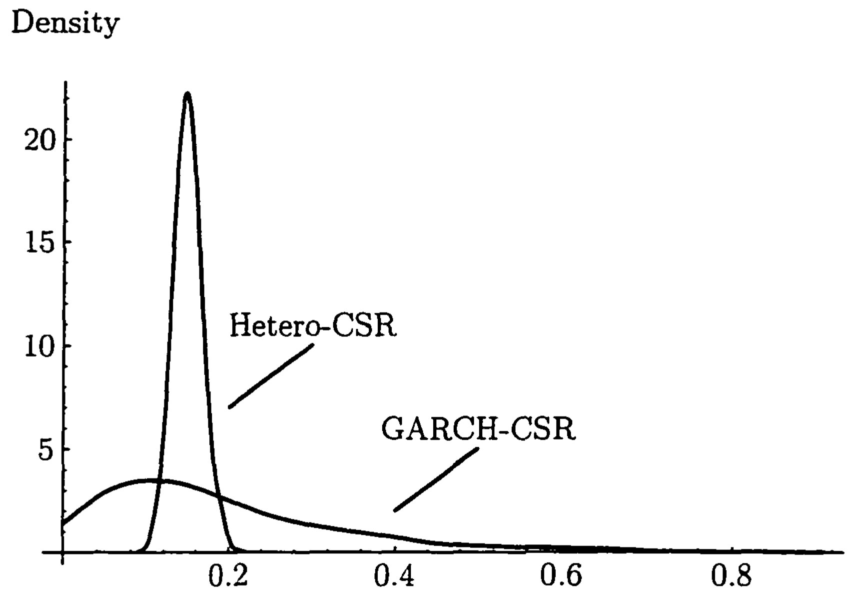

Figure D.63: Simulations 1 ($T = 262 \times 5$), Gaussian kernel density estimates for α for Hetero-CSR and GARCH-CSR, (true value is 0.1513).

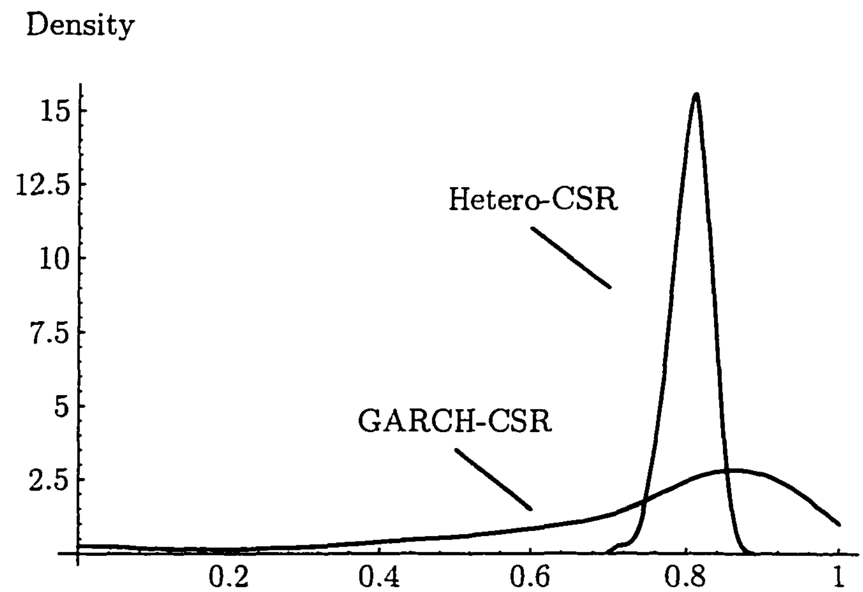

Figure D.64: Simulations 1 ($T = 262 \times 5$), Gaussian kernel density estimates for β for Hetero-CSR and GARCH-CSR, (true value is 0.8083).

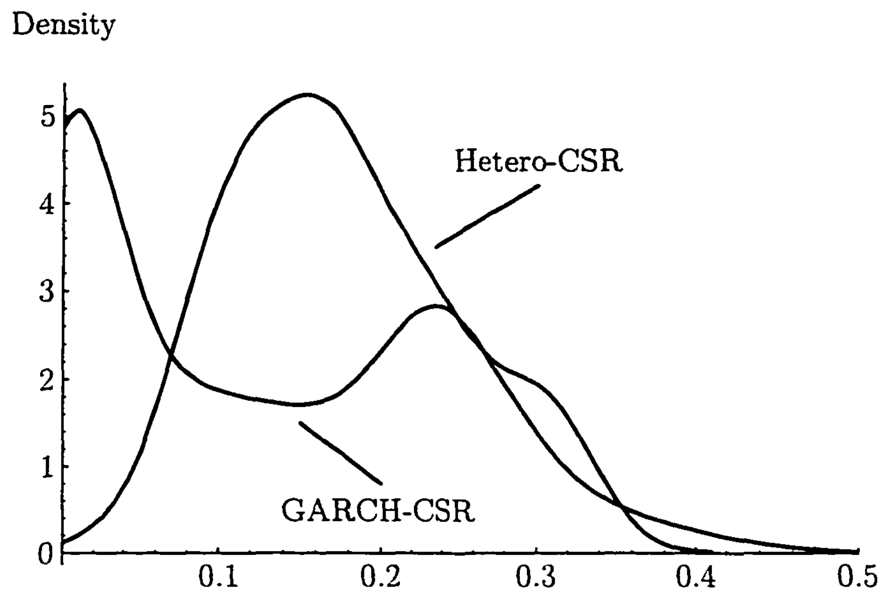

Figure D.65: Simulations 2 ($T = 262$), Gaussian kernel density estimates for ω for Hetero-CSR and GARCH-CSR, (true value is 0.02).

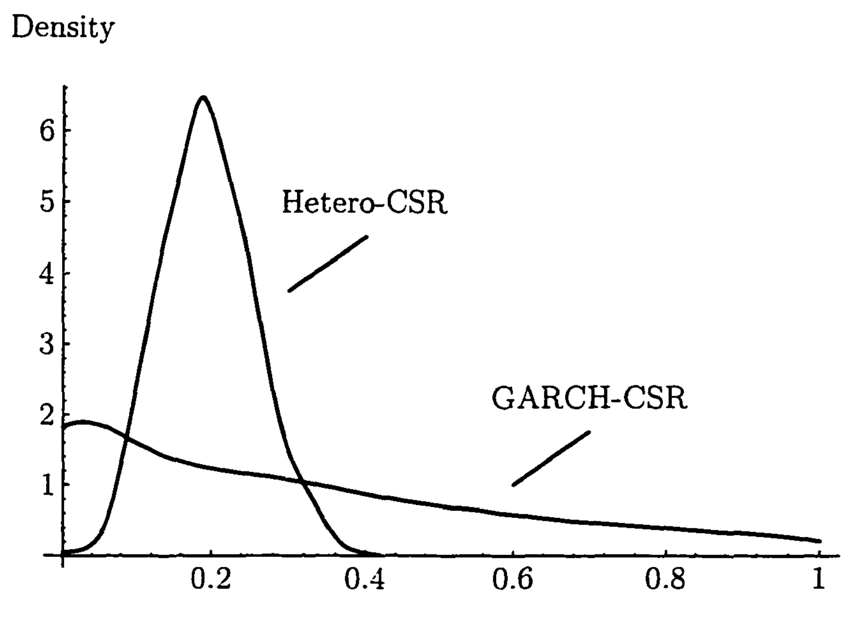

Figure D.66: Simulations 2 ($T = 262$), Gaussian kernel density estimates for α for Hetero-CSR and GARCH-CSR, (true value is 0.20).

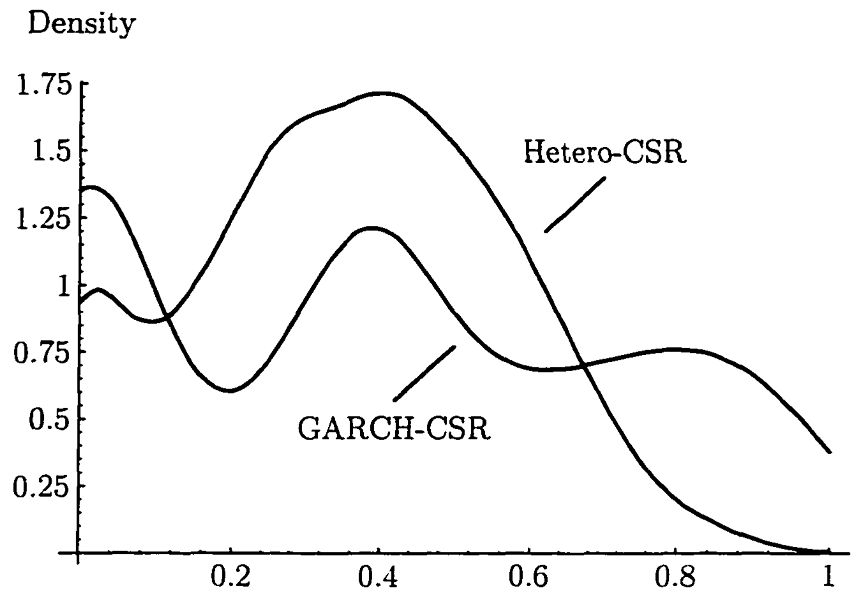

Figure D.67: Simulations 2 ($T = 262$), Gaussian kernel density estimates for β for Hetero-CSR and GARCH-CSR, (true value is 0.40).

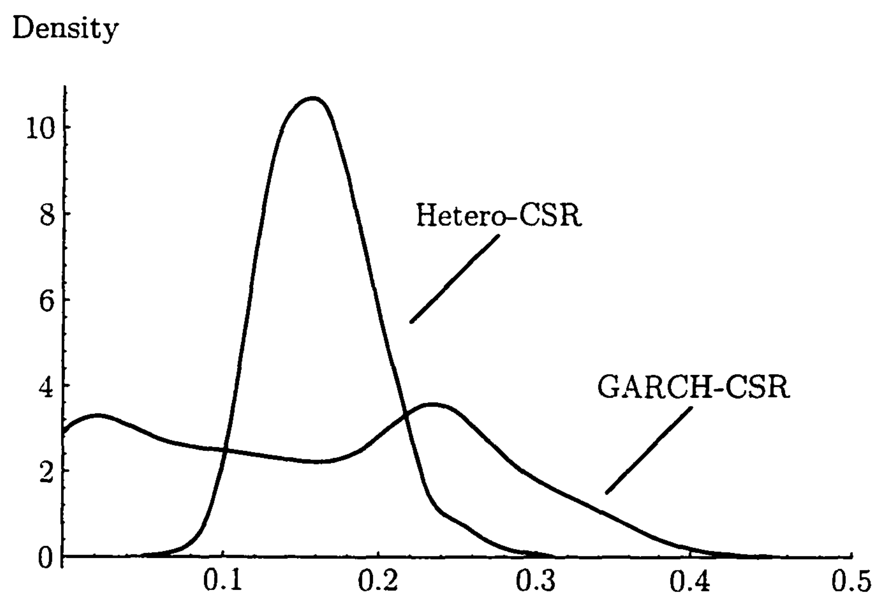

Figure D.68: Simulations 2 ($T = 262 \times 5$), Gaussian kernel density estimates for ω for Hetero-CSR and GARCH-CSR, (true value is 0.02).

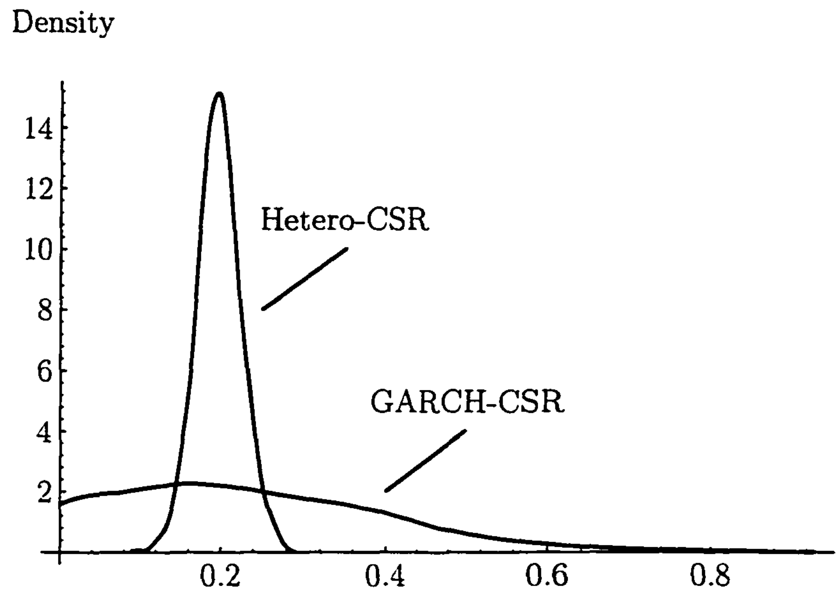

Figure D.69: Simulations 2 ($T = 262 \times 5$), Gaussian kernel density estimates for α for Hetero-CSR and GARCH-CSR, (true value is 0.20).

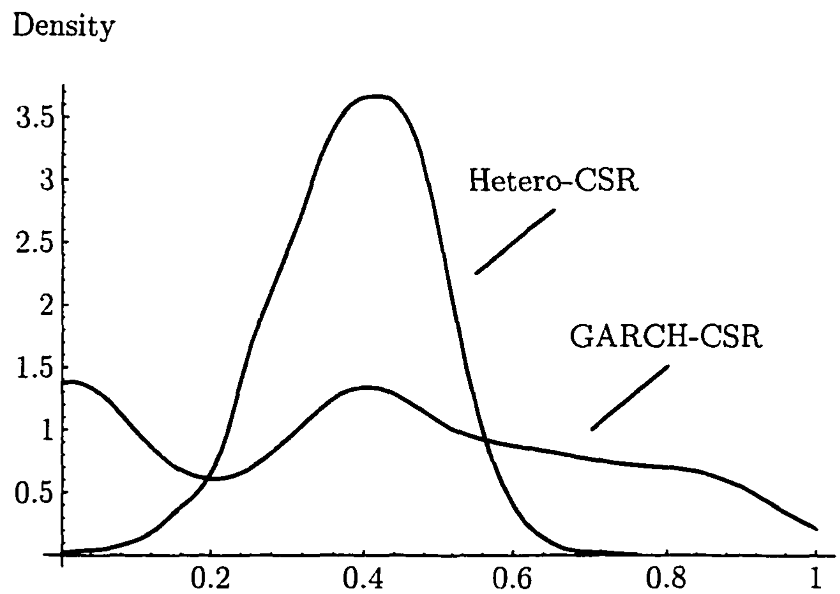

Figure D.70: Simulations 2 ($T = 262 \times 5$), Gaussian kernel density estimates for β for Hetero-CSR and GARCH-CSR, (true value is 0.40).

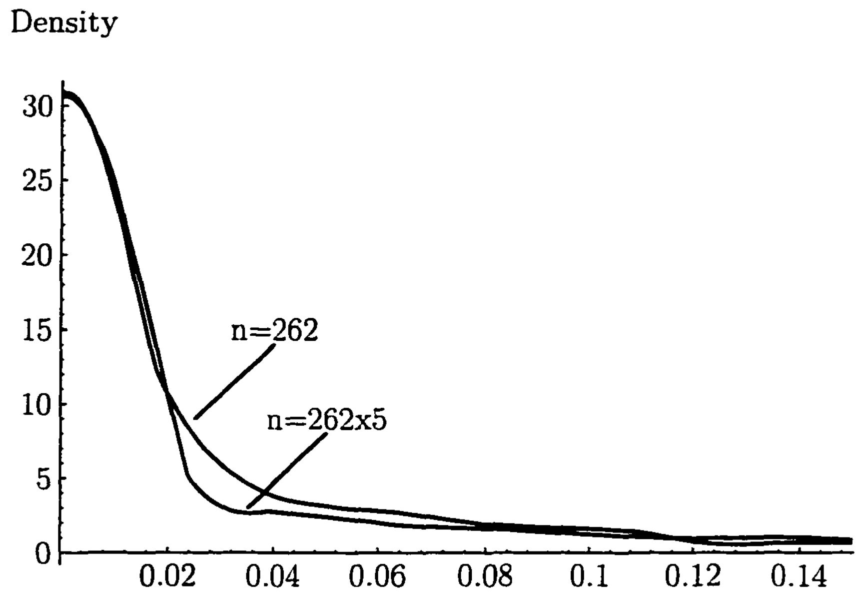

Figure D.71: Simulations 1, Gaussian kernel density for GARCH-CSR ω estimates. (true value is 0.0153).

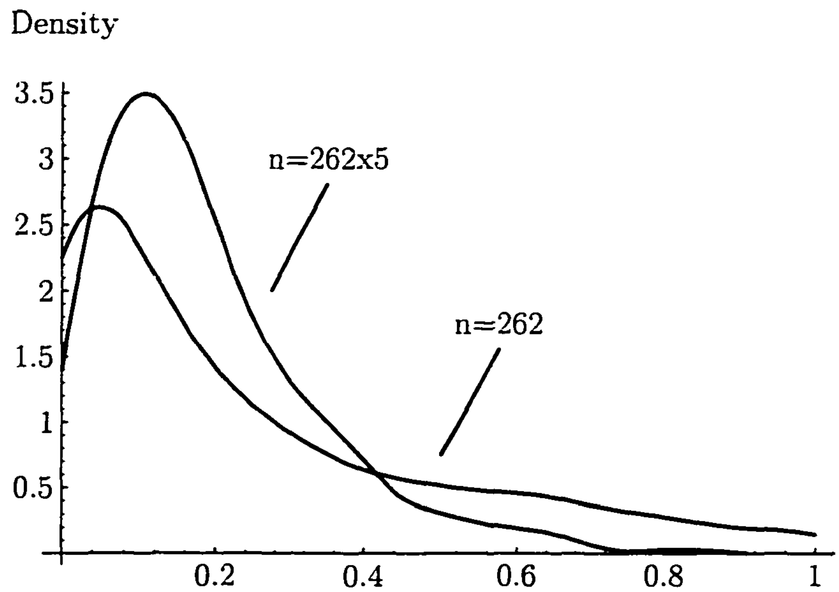

Figure D.72: Simulations 1, Gaussian kernel density for GARCH-CSR α estimates. (true value is 0.1513).

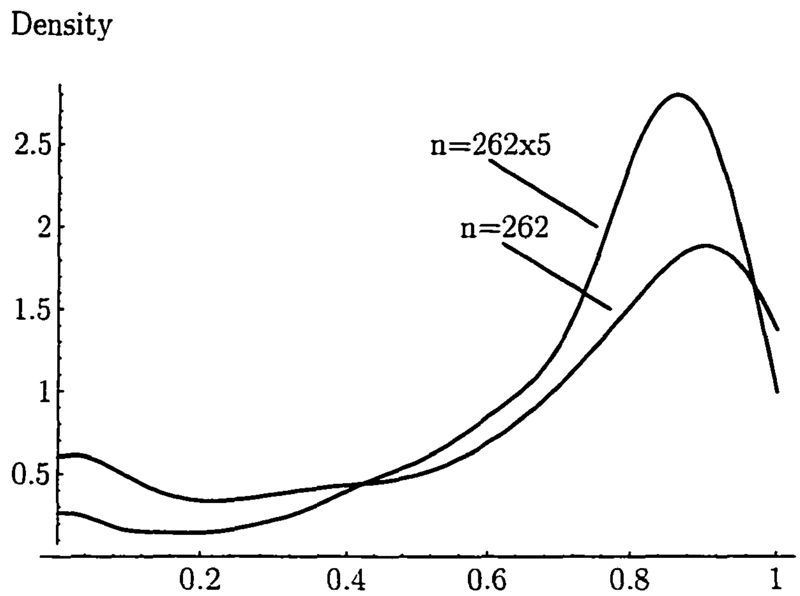

Figure D.73: Simulations 1, Gaussian kernel density for GARCH-CSR β estimates, (true value is 0.8083).

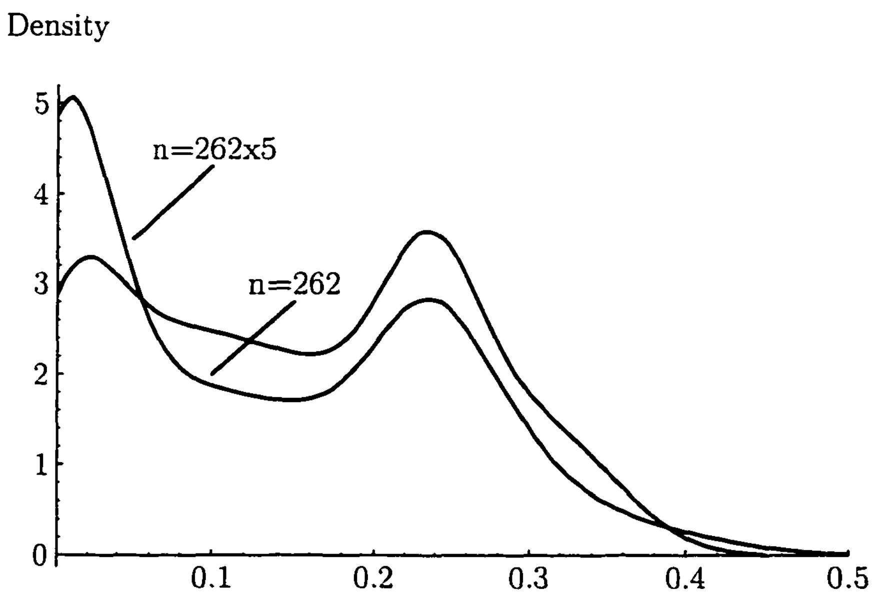

Figure D.74: Simulations 2, Gaussian kernel density for GARCH-CSR ω estimates, (true value is 0.02).

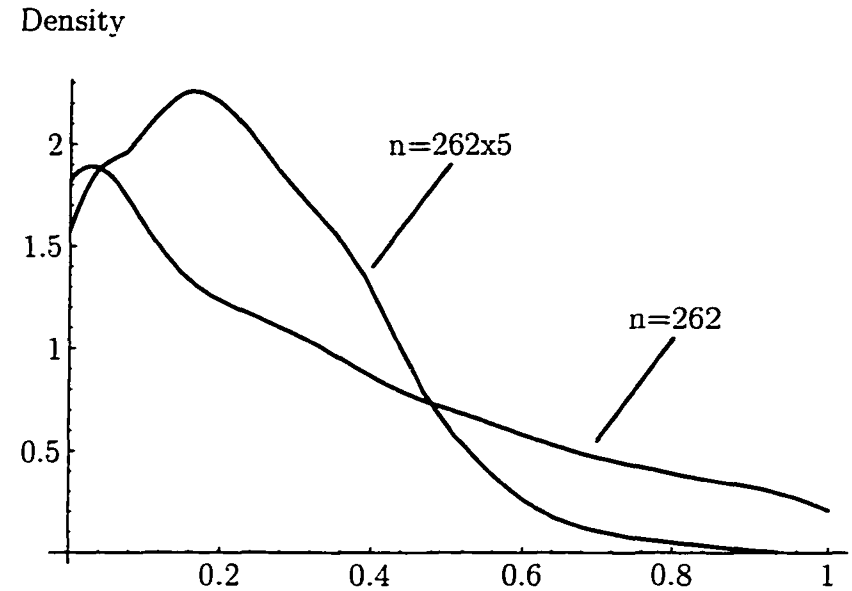

Figure D.75: Simulations 2, Gaussian kernel density for GARCH-CSR α estimates, (true value is 0.20).

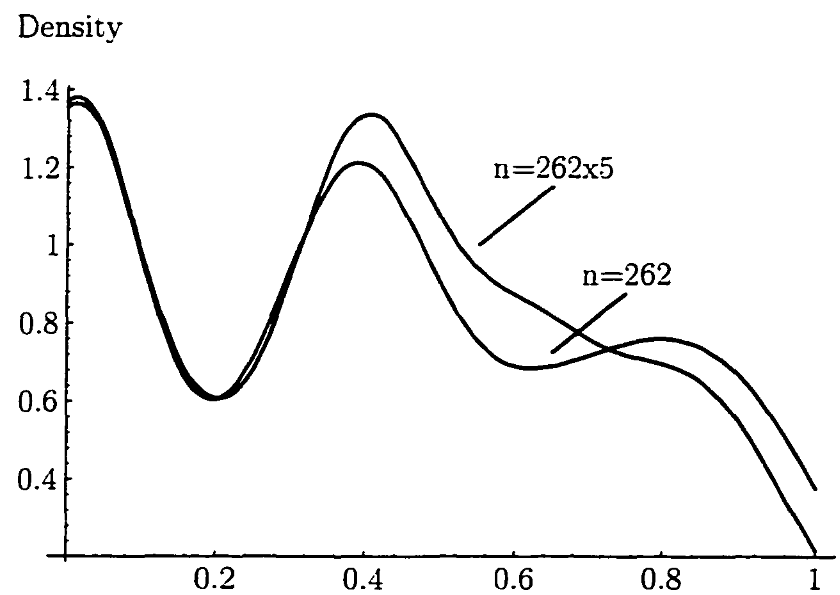

Figure D.76: Simulations 2, Gaussian kernel density for GARCH-CSR β estimates, (true value is 0.40).

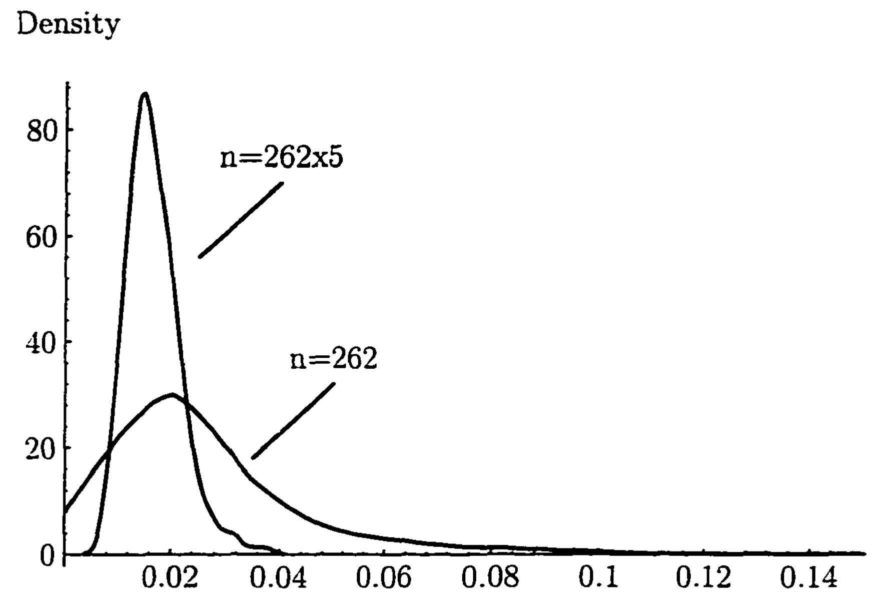

Figure D.77: Simulations 1. Gaussian kernel density for Hetero-CSR ω estimates. (true value is 0.0153).

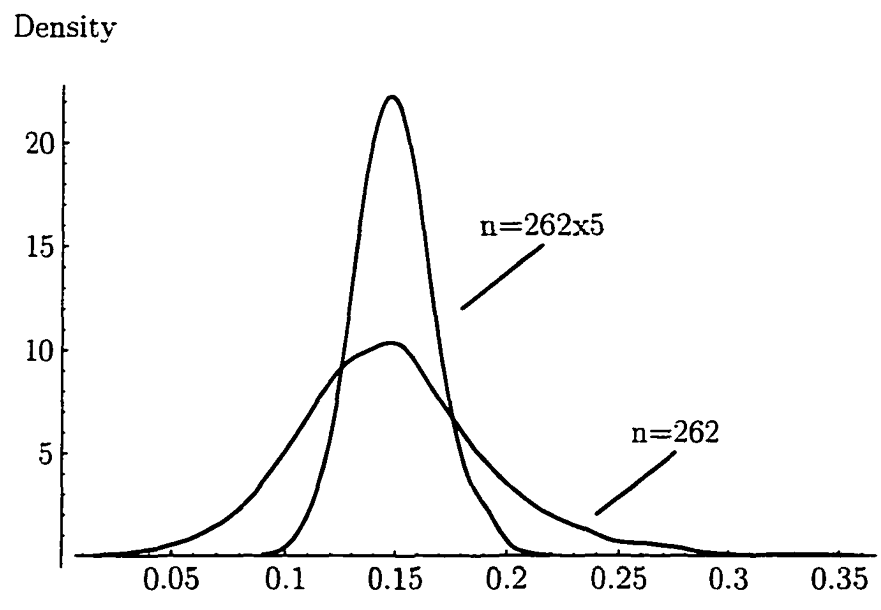

Figure D.78: Simulations 1. Gaussian kernel density for Hetero-CSR α estimates. (true value is 0.1513).

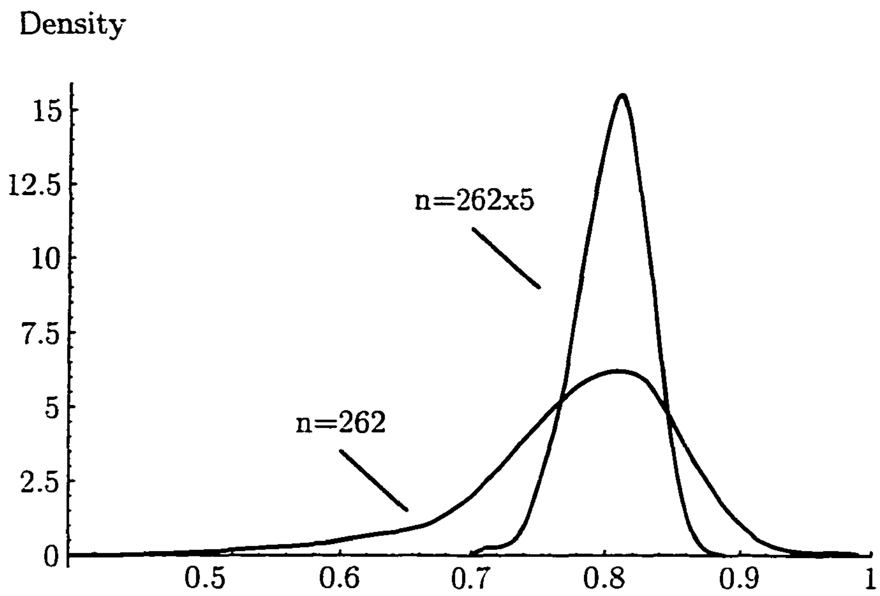

Figure D.79: Simulations 1. Gaussian kernel density for Hetero-CSR β estimates. (true value is 0.8083).

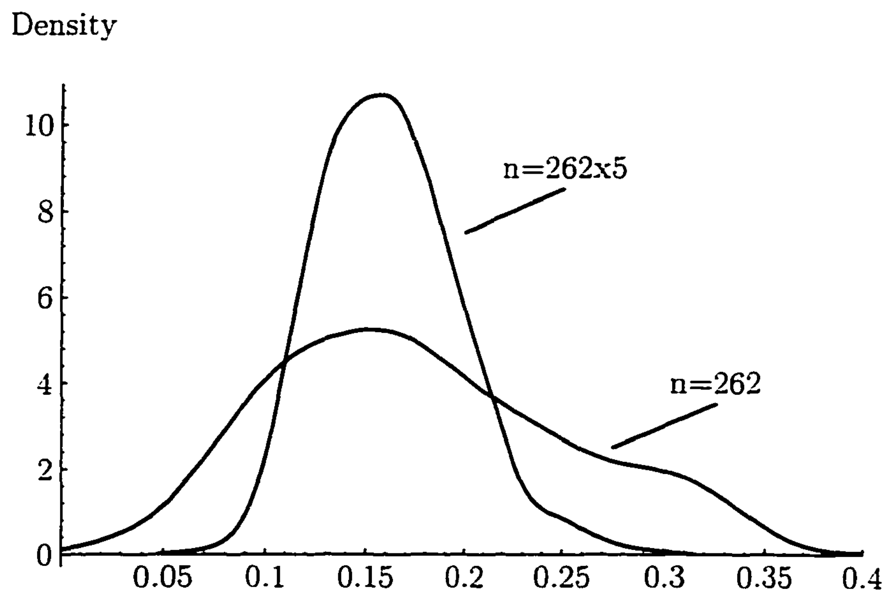

Figure D.80: Simulations 2. Gaussian kernel density for Hetero-CSR ω estimates. (true value is 0.02).

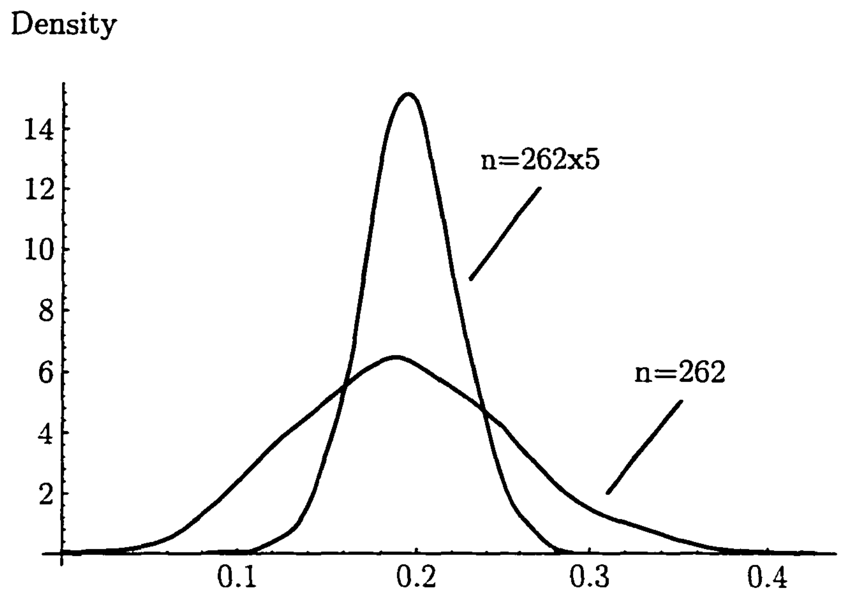

Figure D.81: Simulations 2, Gaussian kernel density for Hetero-CSR α estimates, (true value is 0.20).

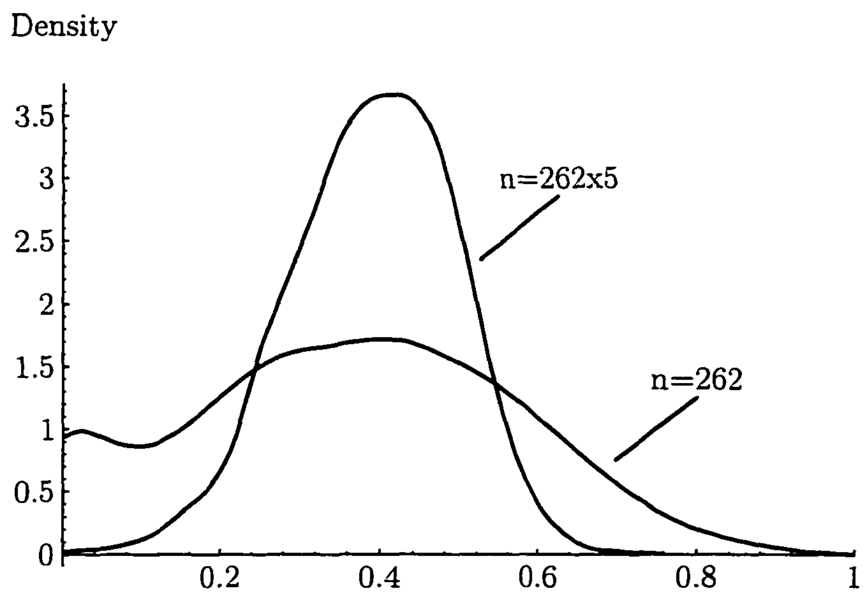

Figure D.82: Simulations 2, Gaussian kernel density for Hetero-CSR β estimates, (true value is 0.40).

References

Andersen, Torben G. and Tim Bollerslev (1997), "Answering the Critics: Yes, ARCH Models do Provide Good Volatility Forecasts", Manuscript, Kellogg Graduate School of Management, Northwestern University.

Andersen, Torben G. and Tim Bollerslev (1996), "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies", Working Paper 217, Kellogg Graduate School of Management, Northwestern University.

Andersen, Torben G. and Tim Bollerslev, "Intraday seasonality and volatility persistence in financial markets", (1994) Working Paper 193, Kellogg Graduate School of Management, Northwestern University.

Aradhyula, Satheesh and Russell Tronstad, "Price seasonality and volatility of different alfalfa hay qualities", Working paper (1994), Department of Agricultural and Resource Economics, University of Arizona.

Baille R. T. and Tim Bollerslev (1991), "The Message in Daily Exchange Rates: A Conditional Variance Tale," *Journal of Business and Economic Statistics*, 7, 297-305.

Baille R. T., Tim Bollerslev and H. O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity" *Journal of Econometrics*, 74, 3-30.

Bollerslev, Tim, "Generalized autoregressive conditional heteroskedasticity", *Journal of Econometrics*, 31, (1986) 307–327.

Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner, "ARCH modeling in finance: a review of the theory and empirical evidence", *Journal of Econometrics*, 52, (1992) 5–59.

Bollerslev, Tim and Eric Ghysels, "Periodic Autoregressive Conditional Heteroskedasticity", *Journal of Business and Economic Statistics*, 14, (1996) 139–151.

Bollerslev, Tim and H. O. Mikkelsen (1996), "Modeling and Pricing Long-Memory in Stock Market Volatility" *Journal of Econometrics*, 73, 151-184.

Dacorogna, Michel M, Ulrich A Müller, Robert J Nagler, Richard B Olsen, and Oliver V. Pictet, "A geographical model for the daily and weekly seasonal volatility in the foreign exchange market", *Journal of International Money and Finance*. 12, (1993) 413–438.

Drost, F. C. and T. E. Nijman, "Temporal aggregation of GARCH processes", *Econometrica*, 61, (1993) 909–927.

Engle, Robert F., "Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, 50, (1982) 987–1007.

Engle, Robert F., Takatoshi Ito, and Wen-Ling Lin, "Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market", University of San Diego Working Paper 88-19R (1988).

Ghose, Devajyoti and Kenneth F. Kroner, "Components of volatility in foreign exchange markets: an empirical analysis of high frequency data", University of Arizona Manuscript, (1997).

Jorion, P. (1995), "Predicting Volatility in the Foreign Exchange Market", *Journal of Finance* 50, 507-528.

LaFrance, Jeffrey T., and Oscar R. Burt, "A modified partial adjustment model of aggregate U.S. agricultural supply", *Western Journal of Agricultural Economics*, 8, (1983) 1–12.

Lopez, Jose A. (1995), "Evaluating the Predictive Accuracy of Volatility Models", Federal Reserve Bank of New York Research Paper No. 9524.

Melvin, Michael, *International Money and Finance*, Fourth Edition, HarperCollins, New York, 1995.

Müller, U. A, M. M. Dacorogna, R. B. Olsen, O. V. Pictet, M. Schwarz and C. Morgenegg. (1990), "Statistical Study if Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis", *Journal of Banking and Finance*, 14, 1189-1208.

Pollak, Robert A., Terence J. Wales "The likelihood dominance criterion: a new approach to model selection", *Journal of Econometrics*, 47 (1991) 227-242.

Zhou, Bin. (1996), "High-Frequency Data and Volatility in Foreign Exchange Rates", *Journal of Business and Economic Statistics*, 14, 1, 45-52.

IMAGE EVALUATION TEST TARGET (QA-3)

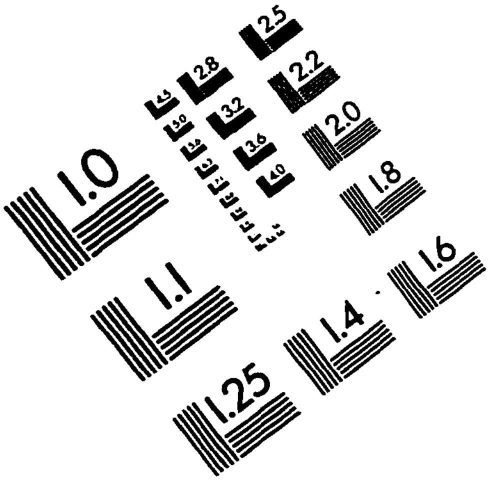
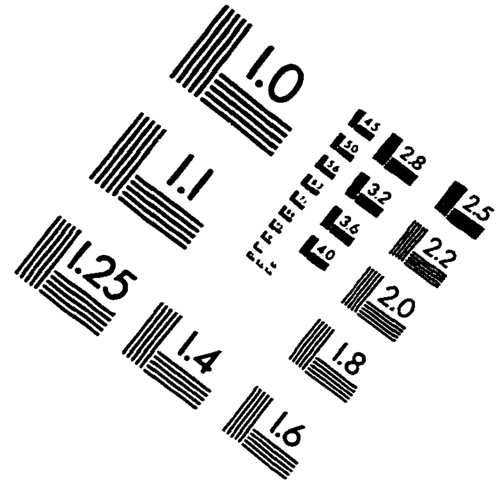

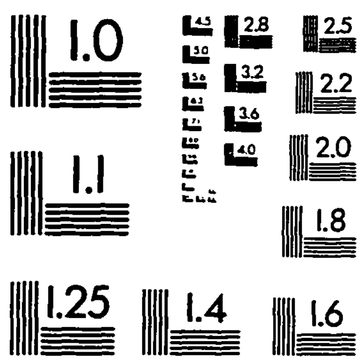

150mm

6″

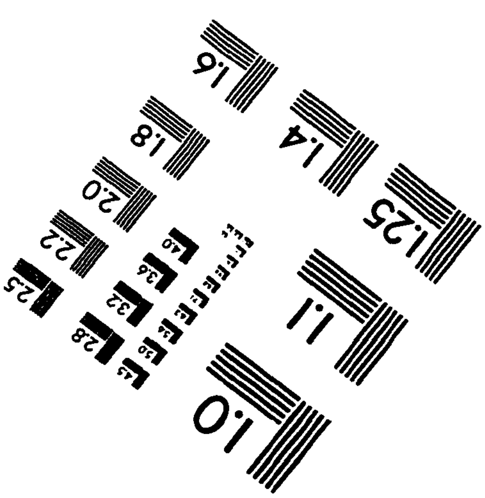
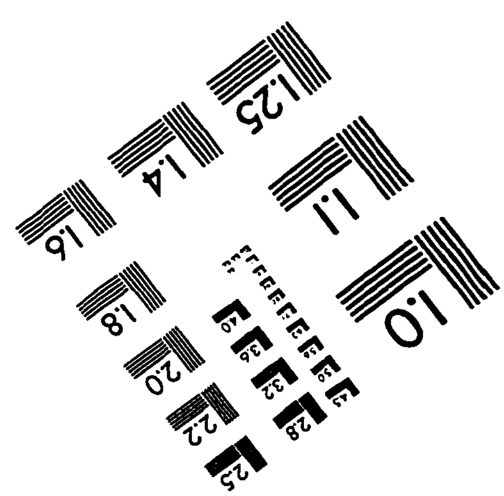

APPLIED IMAGE . Inc
1653 East Main Street
Rochester, NY 14609 USA
Phone: 716/482-0300
Fax: 716/288-5989