

Oomen, Roel C. A.

**Working Paper**

## Modelling realized variance when returns are serially correlated

WZB Discussion Paper, No. SP II 2004-11

**Provided in Cooperation with:**

WZB Berlin Social Science Center

*Suggested Citation:* Oomen, Roel C. A. (2004) : Modelling realized variance when returns are serially correlated, WZB Discussion Paper, No. SP II 2004-11, Wissenschaftszentrum Berlin für Sozialforschung (WZB), Berlin

This Version is available at:

<https://hdl.handle.net/10419/51076>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



WISSENSCHAFTSZENTRUM BERLIN  
FÜR SOZIALFORSCHUNG

SOCIAL SCIENCE RESEARCH  
CENTER BERLIN

Roel C. A. Oomen \*

**Modelling Realized Variance when Returns are  
Serially Correlated**

\* Warwick Business School, University of Warwick

**SP II 2004 – 11**

May 2004

ISSN Nr. 0722 – 6748

**Research Area**  
**Markets and Political Economy**

**Research Unit**  
**Market Processes and Governance**

**Forschungsschwerpunkt**  
**Markt und politische Ökonomie**

**Abteilung**  
**Marktprozesse und Steuerung**

Zitierweise/Citation:

Roel C. A. Oomen, **Modelling Realized Variance when Returns are Serially Correlated**, Discussion Paper SP II 2004 – 11, Wissenschaftszentrum Berlin, 2004.

Wissenschaftszentrum Berlin für Sozialforschung gGmbH,  
Reichpietschufer 50, 10785 Berlin, Germany, Tel. (030) 2 54 91 – 0  
Internet: [www.wz-berlin.de](http://www.wz-berlin.de)

## ABSTRACT

### **Modelling Realized Variance when Returns are Serially Correlated**

by Roel C. A. Oomen\*

This article examines the impact of serial correlation in high frequency returns on the realized variance measure. In particular, it is shown that the realized variance measure yields a biased estimate of the conditional return variance when returns are serially correlated. Using 10 years of FTSE-100 minute by minute data we demonstrate that a careful choice of sampling frequency is crucial in avoiding substantial biases. Moreover, we find that the autocovariance structure (magnitude and rate of decay) of FTSE-100 returns at different sampling frequencies is consistent with that of an ARMA process under temporal aggregation. A simple autocovariance function based method is proposed for choosing the “optimal” sampling frequency, that is, the highest available frequency at which the serial correlation of returns has a negligible impact on the realized variance measure. We find that the logarithmic realized variance series of the FTSE-100 index, constructed using an optimal sampling frequency of 25 minutes, can be modelled as an ARFIMA process. Exogenous variables such as lagged returns and contemporaneous trading volume appear to be highly significant regressors and are able to explain a large portion of the variation in daily realized variance.

*Keywords: High frequency data, realized return variance, market microstructure, temporal aggregation, long memory, bootstrap*

---

\* The paper has greatly benefitted from discussions with Abhay Abhyankar, Frank Diebold, Marcelo Fernandes, George Jiang, Mark Lauer, Bent Sørensen, Juan Toro and Ilias Tsiakas. I am particularly grateful to Søren Johansen for many detailed comments and suggestions. I thank Logical Information Machines, Inc. who kindly provided the data needed for the analysis. The research was supported by the Netherlands Organization for Scientific Research (NWO) under reference number SIR124010 and R46447.

## ZUSAMMENFASSUNG

### **Modellierung realisierter Varianz bei autokorrelierten Erträgen**

Dieser Artikel untersucht die Auswirkungen von autokorrelierten Erträgen auf das Maß der realisierten Varianz bei hochfrequenten Daten über die Erträge. Es wird gezeigt, dass die realisierte Varianz ein verzerrter Schätzer für die bedingte Varianz der Erträge bei Vorliegen von Autokorrelation ist. Unter Verwendung eines zehnjährigen Datensatzes von Minutendaten des FTSE-100 wird dargestellt, dass eine sorgfältige Auswahl der Stichprobenfrequenz unabdingbar zur Vermeidung von Verzerrungen ist. Eine einfache Methode zur Bestimmung der optimalen Stichprobenfrequenz, basierend auf der Autokovarianzfunktion, wird vorgeschlagen. Diese ergibt sich als die höchste Frequenz, bei der die vorhandene Autokorrelation noch einen vernachlässigbaren Einfluss auf das Maß der realisierten Varianz hat. Für den betrachteten Datensatz ergibt sich eine optimale Frequenz von 25 Minuten. Unter Verwendung dieser Frequenz können die logarithmierten Erträge des FTSE-100 als ARFIMA Prozess modelliert werden.

# 1 Introduction

A crucial element in the theory and practice of derivative pricing, asset allocation and financial risk management is the modelling of asset return variance. The Stochastic Volatility and the Autoregressive Conditional Heteroskedasticity class of models have become widely established and successful approaches to the modelling of the return variance process in both the theoretical and the empirical literature (see for example Bollerslev, Engle, and Nelson (1994) and Ghysels, Harvey, and Renault (1996)). Despite the enormous amount of research on return variance modelling carried out over the past two decades, complemented with overwhelming empirical evidence on the presence of heteroskedastic effects in virtually all financial time series, the variety of competing variance models highlights the disagreement on what the correct model specification should be. An alternative route to identifying the dynamics of the return variance process is to utilize the information contained in option prices. Yet, also here, several studies have documented a severe degree of model misspecification even for the more general option pricing formulas that incorporate stochastic volatility, interest rates and jumps (see for example Bakshi, Cao, and Chen (1997)). It is therefore not surprising that a growing number of researchers have turned their attention to the use of high frequency data which, under certain conditions, allow for an essentially non-parametric or model-free approach to the measurement of return variance. The objective of this paper is twofold. First, explore the extent to which the now widely available intra-day data on financial asset prices can be used to improve and facilitate the estimation and modelling of return variance. Special attention is given to the impact that market microstructure-induced serial correlations, present in returns sampled at high frequency, have on the resulting variance estimates. Second, analyze and model the time series of estimated (daily) return variance. Here the focus is on identifying a suitable model plus a set of exogenous variables that is able to characterize and explain variation in the return variance.

The idea of inferring the unobserved return variance from high frequency data is not new. In fact, it can be traced back to Merton (1980) who notes that the variance of a time-invariant Gaussian diffusion process (over a fixed time-interval) can be estimated arbitrarily accurately as the sum of squared realizations, provided that the data are available at a sufficiently high sampling frequency. Empirical studies making use of this insight include French, Schwert, and Stambaugh (1987), who estimate monthly return variance as the sum of squared daily returns and Andersen and Bollerslev (1998), Hsieh (1991), and Taylor and Xu (1997) who estimate daily return variance as the sum of squared intra-day returns. More recent studies that apply and develop this idea further include Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001, 2003), Areal and Taylor (2002), Barndorff-Nielsen and Shephard (2002, 2003), Blair, Poon, and Taylor (2001), Maheu and McCurdy (2002), and Martens (2002).

One of the main attractions that has been put forward of estimating return variance by the sum of squared intra-period returns, a measure commonly referred to as “realized variance” (or “realized volatility” being the square root of realized variance), is that this approach does not require the specification of a potentially misspecified parametric model. In addition, when constructing the realized variance measure there is no need to take the widely documented and pronounced intra-day variance pattern of the return process into account. This feature contrasts sharply with parametric variance models which generally require the explicit modelling of intra-day regularities in return variance (see for example Engle (2000)). Finally, calculating realized variance is straightforward and can be expected to yield accurate variance estimates as it relies on large amounts of intra-day data. The theoretical justification for using the realized variance measure has been provided in a series of recent papers by Andersen, Bollerslev, Diebold, and Labys (2001, 2003, ABDL hereafter). In particular, ABDL have shown that when the return process follows a special semi-martingale, the Quadratic Variation (QV) process is the dominant determinant of the conditional return variance. By definition, QV can be approximated by the sum of squared returns at high sampling frequency, or in other words realized variance. Moreover, under certain restrictions on the conditional mean of the process, QV is the single determinant of the conditional return variance, thereby underlining the importance of the realized variance measure. In related work, Barndorff-Nielsen and Shephard (2003) derive the limiting distribution of realized power variation, that is the sum of absolute powers of increments (i.e. returns) of a process, for a wide class of SV models. It is important to note that, in contrast to conventional asymptotic theory, here, the limit distribution results rely on the concept of “in-fill” or “continuous-record” asymptotics, i.e. letting the number of observations tends to infinity while keeping the time interval fixed. In the context of (realized) variance estimation, this translates into cutting up, say, the daily return into a sequence of intra-day returns sampled at an increasingly high frequency (see for example Foster and Nelson (1996)).

The recently derived consistency and asymptotic normality of the realized variance measure greatly contribute to a better understanding of its properties and, in addition, provide a formal justification for its use in high frequency data based variance measurement. However, a major concern that has largely been ignored in the literature so far, is that in practice the applicability of these asymptotic results is severely limited for two reasons. First, the amount of data available over a fixed time interval is bounded by the number of transactions recorded. Second, the presence of market microstructure effects in high frequency data potentially invalidate the asymptotic results.

This paper studies the properties of the realized variance in the presence of market microstructure-induced serial correlation. In particular, we show that the realized variance measure is a biased estimator of the conditional return variance when returns are serially correlated. The return dependence at high sampling frequencies is analyzed using a decade of minute by minute FTSE-100 index returns.

We find that the autocovariance structure (magnitude and rate of decay) of returns at different sampling frequencies is consistent with that of an ARMA process under temporal aggregation. Based on this finding, an autocovariance based method is proposed to determine the “optimal” sampling frequency of returns, that is, the highest available frequency at which the market microstructure-induced serial correlations have a negligible impact on the realized variance measure<sup>1</sup>.

Following the methodology outlined above, we find that the optimal sampling frequency for the FTSE-100 data set lies around 25 minutes. We construct a time series of daily realized variance, confirm several stylized facts reported in earlier studies, and find that the logarithmic realized variance series can be modelled well using an ARFIMA specification. Exogenous variables such as lagged returns and contemporaneous trading volume appear to be highly significant regressors, explaining a large portion of the variation in daily realized variance. While the regression coefficients of lagged returns indicate the presence of Black’s leverage effect, there is no indication of reduced persistence in the return variance process upon inclusion of contemporaneous trading volume. This latter finding is in sharp contrast with the study by Lamoureux and Lastrapes (1990).

The remainder of this paper is organized as follows. Section 2 investigates the impact of serial correlation in returns on the realized variance measure. Here, results on temporal aggregation of an ARMA process are used to characterize the bias of the realized variance measure at different sampling frequencies. Section 3 reports the empirical findings based on the FTSE-100 data set while Section 4 concludes.

## 2 Realized Variance

The notion of realized variance, as introduced by ABDL, is typically discussed in a continuous time framework where logarithmic prices are characterized by a semi-martingale. More restrictive specifications have been considered by Barndorff-Nielsen and Shephard (2002, 2003). In this setting, the quadratic variation (QV) of the return process can be consistently estimated as the sum of squared intra-period returns. It is this measure that is commonly referred to as realized variance. Importantly, ABDL show that QV is the crucial determinant of the conditional return (co-) variance thereby establishing the relevance of the realized variance measure. In particular, when the conditional mean of the return process is deterministic or a function of variables contained in the information set, the QV is in fact equal to the conditional return variance which can thus be estimated consistently as the sum of squared returns. Notice that this case precludes random *intra*-period evolution of the instantaneous

---

<sup>1</sup>Independent work by Andersen, Bollerslev, Diebold, and Labys (2000b), Corsi, Zumbach, Müller, and Dacorogna (2001) have proposed a similar approach to determine the optimal sampling frequency. Other related studies include Aït-Sahalia and Mykland (2003), Andreou and Ghysels (2001), Bai, Russell, and Tiao (2001).



mean. However, it is argued by ABDL that such effects are likely to be trivial in magnitude and that the QV therefore remains the dominant determinant of the conditional return variance.

Below we analyze the impact of serial correlation in returns on the realized variance measure. As opposed to ABDL and Barndorff-Nielsen and Shephard (2002, 2003), a simple discrete time model for returns is used for the sole reason that it is sufficient to illustrate the main ideas. In what follows, the period of interest is set to one day.

Let  $S_{t,j}$  ( $j = 1, \dots, N$ ) denote the  $j^{\text{th}}$  intra day- $t$  logarithmic price of security  $S$ . At sampling frequency  $f$ , assuming equi-time spaced<sup>2</sup> observations,  $N_f = \frac{N}{f}$  intra-day returns can be constructed as  $R_{f,t,i} = S_{t,if} - S_{t,(i-1)f}$ , for  $i = 1, \dots, N_f$  and  $S_{t,0} = S_{t-1,N}$ . By the additive property of returns, it follows that the day- $t$  return is given by:

$$R_t = \sum_{i=1}^{N_f} R_{t,f,i}.$$

We assume that the (excess) return follows a *martingale difference sequence* and that its conditional distribution, i.e.  $R_{t,f,i} | \mathcal{F}_{t,f,(i-1)}$  where  $\mathcal{F}_{t,f,j}$  denotes the information set available up to the  $j^{\text{th}}$  period of day  $t$ , is symmetric. The need for this symmetry assumption will become clear later on. While this specification allows for deterministic and stochastic fluctuations in the return variance, it also implies that returns are necessarily uncorrelated. Let  $V_1 \equiv R_t^2$ , i.e. the squared day- $t$  return, and  $V_2 \equiv \sum_{i=1}^{N_f} R_{t,f,i}^2$ , i.e. the sum of squared *intra-day- $t$*  returns sampled at frequency  $f$ . In the current context,  $V_2$  is referred to as the realized variance measure. Since returns are serially uncorrelated at any given frequency  $f$ , it follows that:

$$V [R_t | \mathcal{F}_t] = E [R_t^2 | \mathcal{F}_t] = E \left[ \sum_{i=1}^{N_f} R_{t,f,i}^2 | \mathcal{F}_t \right], \quad (1)$$

where  $\mathcal{F}_t$  denotes the information set available prior to the start of day  $t$ . Realized variance, like squared daily return, is therefore an *unbiased estimator of the conditional return variance*. However, it turns out that the variance of  $V_2$  is strictly smaller than the variance of  $V_1$  and is therefore the preferred estimator. To see this, it is sufficient to show that  $E [V_2^2 | \mathcal{F}_t] < E [V_1^2 | \mathcal{F}_t]$ :

$$E [V_1^2 | \mathcal{F}_t] = E \left[ \sum_i \sum_j \sum_k \sum_m R_{t,f,i} R_{t,f,j} R_{t,f,k} R_{t,f,m} | \mathcal{F}_t \right] = E \left[ \sum_i R_{t,f,i}^4 + 3 \sum_i \sum_{j \neq i} R_{t,f,i}^2 R_{t,f,j}^2 | \mathcal{F}_t \right],$$

because the cross product of returns is zero except when (i)  $i = j = k = m$ , (ii)  $i = j \neq k = m$ , (iii)  $i = k \neq j = m$ , (iv)  $i = m \neq j = k$ . Notice that  $E [R_{t,f,i} R_{t,f,j}^3 | \mathcal{F}_t] = 0$  for  $i > j$  by the martingale

---

<sup>2</sup>This can straightforwardly be generalized to irregularly time spaced returns.

difference assumption and  $E [R_{t,f,i}R_{t,f,j}^3|\mathcal{F}_t] = 0$  for  $i < j$  by symmetry of the conditional distribution of returns. On the other hand

$$E [V_2^2|\mathcal{F}_t] = E \left[ \sum_i R_{t,f,i}^4 + \sum_i \sum_{j \neq i} R_{t,f,i}^2 R_{t,f,j}^2 |\mathcal{F}_t \right],$$

from which it directly follows that

$$V [V_2|\mathcal{F}_t] < V [V_1|\mathcal{F}_t].$$

The conditional return variance over a fixed period can thus be estimated arbitrarily accurate by summing up squared intra-period returns sampled at increasingly high frequency. While this result does not depend on the choice of period (i.e. one day), it does crucially rely on the property that returns are serially uncorrelated at any sampling frequency. The additional symmetry assumption rules out any feedback effects from returns into the conditional third moment of returns but allows for skewness in the unconditional return distribution. Other than that, weak conditions are imposed on the return process. As mentioned above, the specification of the return dynamics is sufficiently general so as to allow for deterministic and stochastic fluctuations in the return variance and, as a result, encompasses a wide class of variance models.

## 2.1 Realized Variance in Practice

The results above suggest that straightforward use of high frequency returns can reduce the measurement error in the return variance estimates provided that the return series is a martingale difference sequence (with a symmetric conditional return distribution). This section focuses on the implementation and potential pitfalls that may be encountered in practice. In particular, minute by minute FTSE-100 index level data<sup>3</sup> are used to investigate whether the method of calculating the daily realized variance measure will yield satisfactory results. The additive property of returns allows us to decompose the squared daily return as:

$$R_t^2 = \left[ \sum_{i=1}^{N_f} R_{f,t,i} \right]^2 = \sum_{i=1}^{N_f} R_{f,t,i}^2 + 2 \sum_{i=1}^{N_f-1} \sum_{j=i+1}^{N_f} R_{f,t,i} R_{f,t,j}. \quad (2)$$

It is clear that when the returns are serially uncorrelated at sampling frequency  $f$ , the second term on the right hand side of expression (2) is zero in expectation and the realized variance measure constitutes an unbiased estimator of the conditional return variance. However, when returns are serially correlated

---

<sup>3</sup>I thank *Logical Information Machines, Inc.* who kindly provided the data needed for the analysis. The data set contains minute by minute data on the FTSE-100 index level, starting May 1, 1990 and ending January 11, 2000. For each day, the data is available from 8:35 until 16:10 (except for the period from July 17, 1998 until September 17, 1999 during which the data is available from 9:00 until 16:10). The total number of observations exceeds one million.

the cross product of returns may not vanish in expectation which, in turn, introduces a bias into the realized variance measure. In particular, when returns are positively (negatively) correlated<sup>4</sup>, the sum of squared intra-day returns will *under-estimate* (*over-estimate*) daily conditional return variance as the cross multiplication of returns will be positive (negative) in expectation.

At first sight, the practical relevance of this finding seems to be challenged by the efficient markets hypothesis which claims that the presence of significant serial correlation in returns, if any, is unlikely to persist for extended periods of time. It is important to note, however, that the efficient markets hypothesis concerns *economic* and not statistical significance of serial correlation. Therefore, due to the presence of market microstructure<sup>5</sup> effects and transaction costs, a certain degree of serial correlation in returns does not necessarily conflict with market efficiency.

In the market microstructure literature, a prominent hypothesis that is able to rationalize serial correlation in stock index returns is non-synchronous trading. The basic idea is that when individual securities in an index do not trade simultaneously, the contemporaneous correlation among returns induces serial correlation in the index returns. Intuitively, when the index components non-synchronously incorporate shocks to a common factor that is driving their price, this will result in a sequence of correlated price changes at the aggregate or index price level. As discussed by Lo and MacKinlay (1990), non-synchronous trading induces positive serial correlation in the index returns. On the other hand, the Roll (1984) bid/ask bounce hypothesis often applies to single asset returns which are typically found to exhibit negative serial correlation. Here the argument is as follows: when at a given point in time no new information arrives in a (dealer) market, the stock price is expected to bounce between the bid and the ask price whenever a trade occurs. Although this phenomenon may not be apparent at a daily or weekly frequency, it is likely to have a discernible impact on returns sampled at high (intra-day) frequency. Finally, transaction costs and feedback trading, in addition to non-synchronous trading and the bid-ask bounce, may also induce serial correlation in returns. For an empirical investigation of these issues see for example Säfvenblad (2000). Although this paper does not aim to analyze the various market microstructure effects in specific, we do want to highlight the presence of such effects and study their impact on the realized variance measure.

Several studies have encountered the impact of serial correlation in returns on the estimates of return variance. For example, French and Roll (1986) find that stock return variance is much lower when estimated using hourly instead of daily data, indicating the presence of positive serial correlation in their data set. Recognizing the presence of serial correlation, French, Schwert, and Stambaugh

---

<sup>4</sup>When returns exhibit both positive and negative serial correlation, the effect is not clear. The realized variance measure may be biased or unbiased depending on the relative magnitudes of the return autocovariance at different orders.

<sup>5</sup>For an in depth discussion of the relation between market microstructure and price dynamics see for instance Campbell, Lo, and MacKinlay (1997), Lequeux (1999), Madhavan (2000), O'Hara (1995), Wood (2000) and references therein.

(1987) estimate monthly return variance as the sum of squared daily returns *plus* twice the sum of the products of adjacent returns. Froot and Perold (1995) also find significant positive serial correlation in 15 minute returns on S&P500 cash index from 1983-1989 and show that the annualized return variance estimates based on weekly data are significantly higher (about 20%) than the variance estimates based on 15-minute data. More recently, Andersen, Bollerslev, Diebold, and Labys (2000b) document the dependence of the realized variance measure on return serial correlation.

These findings offer an early recognition of the central idea of this paper: the results derived in the previous section, and the consistency and asymptotic results derived in ABDL and Barndorff-Nielsen and Shephard (2003), are not applicable to return data that exhibit a substantial degree of serial dependence. In particular, the conditional mean specification used in these studies does typically not allow for the random intra-day evolution of the conditional mean<sup>6</sup>. It is commonly argued that this flexibility is not required at low, say daily or weekly, frequencies. However, when moving to higher intra-day sampling frequencies, the characteristics of the data may change dramatically due to the presence of market microstructure which in turn, leads to substantial dependence in the conditional mean of the return process.

Because market microstructure effects are present in virtually all financial return series, the issue outlined above is central to the discussion of high frequency data based variance measurement. This is emphasized in the empirical analysis which is based on minute by minute returns on the FTSE-100 stock market index. Specifically, the 10 year average (1990-2000) of the two terms on the right hand side of expression (2) is computed for sampling frequencies between 1 and 45 minutes and the results are displayed in Figure 1. The implicit assumption we make here is that the return process is weakly stationary<sup>7</sup> so that the averaging (over time) is justified and the estimates can be interpreted as (co)variance estimates.

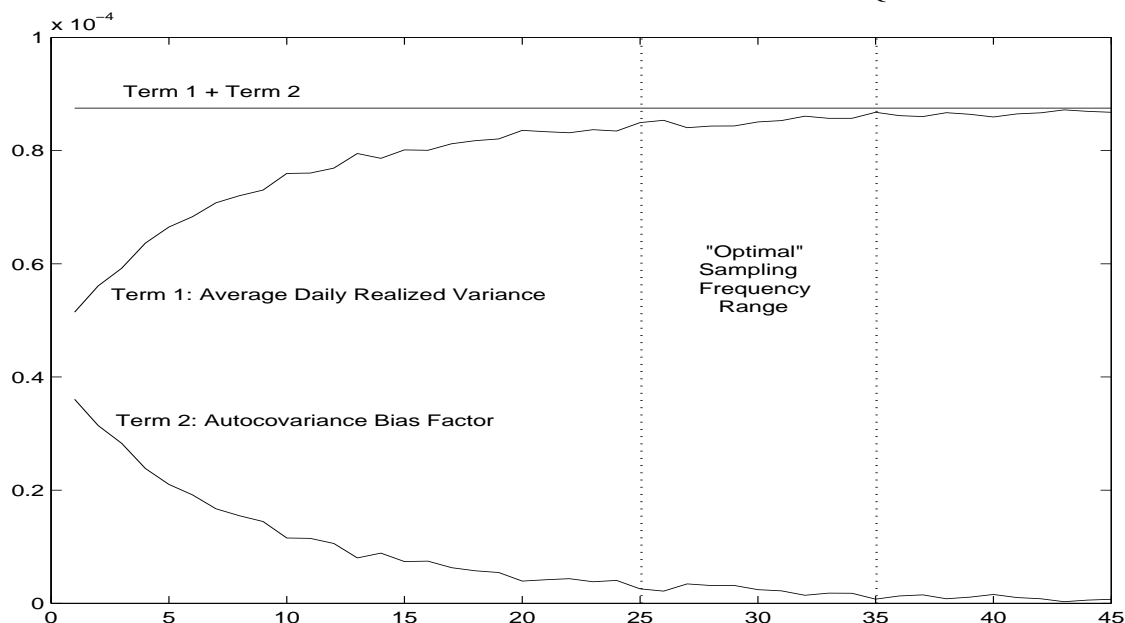
It is clear that for FTSE-100 data the first term, the realized variance measure, increases with a decrease in sampling frequency while the second term, the summation of cross multiplied returns, decreases. The positivity of the second term indicates that the FTSE-100 returns are positively correlated, introducing a *downward bias* into the realized variance measure, while its decreasing pattern demonstrates that this dependence, and consequently the bias, diminishes when sampling is done less frequently. This term, which measures the bias that is introduced by the serial dependence of returns, is referred to as the “*autocovariance bias factor*” in the remainder of this paper. Figure 1 illustrates that an ad hoc choice of sampling frequency can lead to a substantial (downward) bias in the realized variance measure. In fact, at the highest available sampling frequency of 1 minute, the bias in the

---

<sup>6</sup>An exception is the general model covered by Theorem 1 in Andersen, Bollerslev, Diebold, and Labys (2003) from which it is also clear that the realized variance measure yields a biased estimate of the conditional return variance.

<sup>7</sup>For the bootstrap analysis of Section 2.2 we need to impose strict stationarity and weak dependence on the return process.

FIGURE 1: REALIZED VARIANCE VERSUS SAMPLING FREQUENCY



Notes: The 1990-2000 average daily FTSE-100 realized variance (i.e.  $T^{-1} \sum_{t=1}^T \sum_{i=1}^{N_f} R_{f,t,i}^2$ ) and autocovariance bias factor (i.e.  $2T^{-1} \sum_{t=1}^T \sum_{i=1}^{N_f-1} \sum_{j=i+1}^{N_f} R_{f,t,i} R_{f,t,j}$ ) for sampling frequencies between 1 and 45 minutes.

variance estimate is more than 35%! To stress the economic significance of this finding, we notice that in a Black Scholes world, a mere 10% under-estimation of the return variance leads to a 14.5% underpricing of a 3 month, 15% out of the money option. Also the option's delta is 8.2% lower than its true value. Indeed, Figlewski (1998) finds that an accurate return variance estimate is of single most importance when hedging derivatives. When the return variance is stochastic, Jiang and Oomen (2002) also find that for the hedging of derivatives accurate estimation of the *level* of return variance is far more important than accurate estimation of the dynamic parameters of the variance process. Pricing and hedging options aside, it is easy to think of a number of other situations where accurate return variance estimates are of crucial importance. Risk managers often derive Value at Risk figures from the estimated return variance of a position. Also, in a multivariate setting, the covariance matrix of returns is the primary input for portfolio choice and asset allocation.

The above discussion naturally leads to the important question at which frequency the data should be sampled. Figure 1 plays a central role in answering this question by providing a graphical depiction of the trade-off one faces when constructing the realized variance measure: an increase in the sampling frequency yields a greater amount of data, thereby attaining higher levels of efficiency (in theory), while at the same time a decrease in the sampling frequency mitigates the biases due to market microstructure effects surfacing at the highest sampling frequencies. A balance must be struck between these oppos-

ing effects and it is argued here that an autocovariance based method, such as the autocovariance bias factor of Figure 1, can be used to determine *the “optimal” sampling frequency as the highest available sampling frequency for which the autocovariance bias term is negligible*<sup>8</sup>. Clearly, deciding whether the bias term is “negligible”, and whether the sampling frequency is therefore “optimal”, may prove a difficult issue for at least two reasons. First, even though it may be possible to bootstrap confidence bounds around the autocovariance bias factor in order to determine the frequency at which the bias is statistically indistinguishable from zero (see Table 1 for some related results), for many applications economic significance, as opposed to statistical significance, may be the relevant metric with which to measure “negligibility”. The optimal sampling frequency may therefore very well depend on the particular application at hand. Second, when aggregating returns, a reduction in bias should generally be weighed against the loss in efficiency. In practice, however, both the loss or gain in bias and efficiency will often be difficult to quantify which, in turn, complicates the choice of optimal sampling frequency. It should be noted that for a general SV model, Barndorff-Nielsen and Shephard (2003) have shown that the realized variance measure converges to integrated variance at rate  $\sqrt{N}$  where  $N$  is the number of intra-period observations. Also, Oomen (2003) has derived an explicit characterization of the bias term as a function of the sampling frequency when the price process follows a compound Poisson process with correlated innovations. While the results in these studies may yield some valuable insights into the bias-efficiency trade-off, it is important to keep in mind that they are derived under potentially restrictive parametric specifications for the price process. As such, they should be interpreted cautiously when applied to high frequency data which, as we show below, are often contaminated by market microstructure effects. Without further going into this, it seems reasonable to expect that for the FTSE-100 data the optimal sampling frequency lies somewhere between 25 and 35 minutes, i.e. the range indicated in Figure 1.

## 2.2 Serial Correlation, Time Aggregation & Sampling Frequency

We now take a closer look at the autocovariance bias term and show how its shape is intimately related to the dynamic properties of intra-day returns at different sampling frequencies.

Table 1 reports some standard descriptive statistics for the FTSE-100 return data. Because it is well known that financial returns, and in particular high frequency returns, are not independently and identically distributed we bootstrap the confidence bounds around the statistics instead of deriving them from the well known asymptotic distributions that are valid under the iid null hypothesis. For the return volatility and the skewness and kurtosis coefficients we use the stationary bootstrap of Politis and Romano (1994) who show that this procedure is valid for strictly stationary, weakly dependent data. Let

---

<sup>8</sup>Independent work by Andersen, Bollerslev, Diebold, and Labys (2000b), Corsi, Zumbach, Müller, and Dacorogna (2001) have proposed a similar approach to determine the optimal sampling frequency.

$\mathbf{x} = (x_1, \dots, x_N)$  denote the original data set (i.e. time series of returns at a given sampling frequency) and let  $X_{i,k} \equiv (x_i, \dots, x_{i+k-1})$  where  $i = 1, \dots, N$ ,  $k = 1, 2, \dots$ , and  $x_j = x_{j \bmod N}$  for  $j > N$ . A bootstrap sample is constructed as  $\mathbf{x}^* = (X_{i_1, k_1}, \dots, X_{i_b, k_b})$  where  $\sum_{j=1}^b k_j = N$ ,  $i$  has a discrete uniform distribution on  $\{1, \dots, N\}$ , and  $k$  has a geometric distribution, i.e.  $P(k = m) = p(1-p)^{m-1}$  for  $m = 1, 2, \dots$ . Based on this re-sampled time series, we then compute the relevant test statistics. By simulating a large number  $B$  of bootstrap samples we can approximate the true distribution of the test statistics by the empirical distribution of the  $B$  values of the associated statistics. The idea behind sampling blocks instead of single entries is that, when the block length is sufficiently large, the dependence structure of the original series will be preserved in the re-sampled series to a certain extent. Evidently, the correspondence between the distribution of the original and the re-sampled series will be closer the weaker the dependence and the longer the block length. To choose  $p$ , or equivalently the expected block length  $E[k] = 1/p$ , we have experimented with a number of different values but find, in line with several other studies (Horowitz, Lobato, Nankervis, and Savin 2002, Romano and Thombs 1996), that the results are rather insensitive to the choice of  $p$ . The results reported in Table 1 are based on  $p = 1/15$  (i.e.  $E[k] = 15$ ) and  $B = 5,000$ .

The confidence intervals for the correlation coefficients and the critical value the Box-Ljung test statistic are obtained by the “blocks-of-blocks” bootstrap. Instead of sampling a  $1 \times k$  block, as is done in the stationary bootstrap, we now sample an  $h \times k$  block  $X_{i,k,h} = (x_i^h, \dots, x_{i+k-1}^h)$  where  $x_i^h = (x_i, \dots, x_{i+h-1})'$  and  $h - 1$  matches the maximum order of correlation coefficient to be computed. Analogous to the procedure described above, an  $h \times N$  bootstrap sample is constructed as  $\mathbf{x}^* = (X_{i_1, k_1, h}, \dots, X_{i_b, k_b, h})$  from which the  $k^{\text{th}}$  order correlation coefficients can be computed as

$$\hat{\rho}_k = \frac{\sum_{i=1}^N (\mathbf{x}_{1,i}^* - \bar{\mathbf{x}}_{1,\cdot}^*) (\mathbf{x}_{k+1,i}^* - \bar{\mathbf{x}}_{k+1,\cdot}^*)}{\left[ \sum_{i=1}^N (\mathbf{x}_{1,i}^* - \bar{\mathbf{x}}_{1,\cdot}^*)^2 \sum_{i=1}^N (\mathbf{x}_{k+1,i}^* - \bar{\mathbf{x}}_{k+1,\cdot}^*)^2 \right]^{1/2}}$$

where  $\bar{\mathbf{x}}_{i,\cdot}^* = N^{-1} \sum_{j=1}^N \mathbf{x}_{i,j}^*$ . Because the null-hypothesis for the Box-Ljung statistic is uncorrelatedness, we first pre-whitened the data using an AR(15)<sup>9</sup> and implement the bootstrap procedure on the residuals. As above, the geometric parameter and the number of bootstrap replications are set as  $p = 1/15$  and  $B = 5,000$ . For more details on how to approximate the sampling distribution of the correlation coefficients and the Box-Ljung statistics using the (blocks-of-blocks) bootstrap see for example Davison and Hinkley (1997), Horowitz, Lobato, Nankervis, and Savin (2002) and Romano and Thombs (1996).

Based on the above bootstrap procedures we construct 95% confidence bounds for the descriptive

---

<sup>9</sup>We note that the choice of AR-order is relatively ad hoc, and could arguably be lowered with a decrease in sampling frequency. However, with the amount of data we work with here, it can be expected that the efficiency loss associated with the potentially redundant AR-terms is minimal. Hence, for simplicity, we keep the AR-order fixed across the different sampling frequencies.

TABLE 1: DESCRIPTIVE STATISTICS OF FTSE-100 RETURNS.


Frequency	No. Obs.	Volatility	Skewness	Kurtosis	BL[15]		
1 Min	1,046,862	11.5 (10.9;12.2)	1.61* (-14.9;17.0)	<b>3305*</b> (1209;5333)	<b>26874*</b> (312.4)		
5 Min	209,372	13.2 (12.6;14.0)	-0.43* (-6.82;5.18)	<b>508.3*</b> (177.5;886.2)	<b>4844.2*</b> (223.6)		
10 Min	104,686	14.2 (13.4;15.1)	-1.89* (-8.50;3.04)	<b>344.8*</b> (78.5;697.2)	<b>1206.9*</b> (138.0)		
30 Min	34,895	15.1 (14.2;16.4)	-1.11* (-4.61;1.51)	<b>115.2*</b> (27.1;234.8)	<b>221.1*</b> (181.6)		
60 Min	17,447	15.6 (14.3;17.3)	-0.32* (-2.71;1.86)	<b>84.3*</b> (14.3;161.5)	117.2* (146.2)		
1 Day	2,407	15.2 (14.0;16.6)	0.05 (-0.21;0.32)	<b>5.43*</b> (4.37;6.56)	<b>44.34*</b> (37.3)		
		$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_5$	$\rho_{10}$ $\rho_{15}$
1 Min	<b>10.4*</b> (9.06;11.7)	<b>7.39*</b> (6.22;8.55)	<b>6.09*</b> (5.08;7.12)	<b>5.13*</b> (3.89;6.68)	<b>2.75*</b> (1.77;3.68)	0.59* (-0.43;1.34)	<b>0.95*</b> (0.32;1.69)
5 Min	<b>14.6*</b> (12.4;16.7)	<b>3.35*</b> (1.92;4.75)	<b>1.85*</b> (0.70;3.05)	0.78* (-0.29;2.00)	<b>0.96*</b> (0.15;1.81)	-0.10 (-1.23;1.32)	-0.60* (-1.57;0.35)
10 Min	<b>10.2*</b> (8.40;12.1)	<b>2.35*</b> (0.94;3.83)	<b>1.06*</b> (0.11;2.07)	0.04 (-1.45;1.65)	-0.55 (-0.20;0.12)	0.17 (-0.71;1.13)	-0.68* (-2.72;0.82)
30 Min	<b>6.39*</b> (4.25;8.68)	-0.60 (-2.48;1.24)	1.28* (-0.58;3.19)	-0.88 (-2.21;0.50)	-0.10 (-3.57;2.58)	2.02* (-0.10;4.46)	-1.83* (-6.35;2.00)
60 Min	<b>3.28*</b> (0.12;6.41)	-0.79 (-4.59;2.58)	<b>3.66*</b> (1.16;6.04)	-0.00 (-2.40;2.41)	2.84* (-0.14;6.26)	-1.55* (-3.03;0.07)	-0.26 (-3.27;2.64)
1 Day	<b>7.56*</b> (3.64;11.2)	-3.62 (-8.45;1.64)	-3.70 (-8.76;1.63)	-0.46 (-5.69;4.73)	-2.44 (-8.20;3.20)	3.33 (-2.26;8.73)	1.84 (-3.23;6.75)

Notes: The upper panel reports the *annualized* return volatility in percentage points (“Volatility”), the skewness coefficient (“Skewness”), kurtosis coefficient (“Kurtosis”), and the Box-Ljung test statistic on the first 15 autocorrelations (“BL[15]”) for FTSE-100 returns sampled at frequencies between 1 minute and 1 day over the period 1990-2000. The lower panel reports the serial correlation coefficients in percentage points ( $\rho_k$  denotes the  $k^{th}$  order correlation coefficient). Bootstrapped 95% confidence bounds (and critical values for the Box-Ljung test) are reported in parentheses below. An asterisk indicates significance at 95% confidence level under the null hypothesis that returns are iid distributed;  $\sqrt{T}$ Skewness  $\xrightarrow{d} \mathcal{N}(0, 6)$ ,  $\sqrt{T}$ Kurtosis  $\xrightarrow{d} \mathcal{N}(0, 24)$ ,  $\sqrt{T}\rho_k \xrightarrow{d} \mathcal{N}(0, 1)$ , and  $BL[K] \xrightarrow{d} \chi_K^2$ .

statistics under the null that returns are weakly dependent and report them in parentheses in Table 1. The statistics that are significant are printed in bold. For comparison purposes, an asterisk indicates 95% significance under the alternative null hypothesis that returns are independently and identically distributed. For this case, it is well known that the square root of the sample size times the  $k^{th}$  order serial correlation, skewness, and kurtosis coefficients of returns are asymptotically distributed as normal with variance 1, 6, and 24 respectively. The Box-Ljung statistic on the first  $K$  autocorrelations,  $BL[K]$ , is asymptotically distributed as chi-square with  $K$  degrees of freedom. Turning to the results in Table 1, we find that there is substantial excess kurtosis and serial correlation in high frequency returns. At the minute frequency, most of the serial correlation coefficients up to order 15 are significant and



FIGURE 2: CORRELOGRAM OF 1 MINUTE AND 5 MINUTE FTSE-100 RETURNS



Notes: Correlogram of 1 minute (left panel) and 5 minute (right panel) FTSE-100 index returns for the period 1990-2000

the kurtosis coefficient indicates the presence of an extremely fat tailed marginal return distribution. However, aggregation of returns brings the distribution of returns closer to normal and reduces both the order and magnitude of the serial correlation (see also Figure 2). At the daily frequency, the excess kurtosis has come down from around 3000 to about 2.5, and the serial correlation coefficients of order higher than one are all insignificantly different from zero. Consistent with the autocovariance bias term above (Figure 1), we also see that the (annualized) return volatility increases with a decrease of the sampling frequency. Interestingly, the 95% confidence bounds for frequencies lower than 30 minutes (i.e. 1, 5, and 10 minutes) do not include the point estimate of the annualized return variance based on daily data. This suggests that the autocovariance bias term at these frequencies is statistically different from zero which, in turn, corroborates our choice of “optimal” sampling frequency range on statistical grounds.

It is also clear from Table 1 that the bootstrapped confidence bounds deviate substantially from their iid-asymptotic counterparts. As a result, a number of statistics that are significant under the (invalid) iid null hypothesis, turn out to be insignificant based on the bootstrapped confidence bounds which allows for weak dependence in the return data. For example, while the skewness of intra-daily returns is significant under the iid hypothesis, none of the skewness coefficients are significant under the alternative null-hypothesis. Also, the maximum order of the significant correlation coefficients is generally lower for the bootstrapped critical values than for the iid-asymptotic values. For example, at frequencies between 5 and 30 minutes,  $\rho_{15}$  is found significant under the iid-hypothesis but insignificant under the alternative hypothesis. These findings emphasize the inadequacy of the “iid-” asymptotic distributions for this data and illustrate the value of the bootstrap method.

Turning to the specification of the return process, we notice that the overwhelming significance of

the serial correlation coefficients reported in Table 1 and Figure 2 suggests that the characteristics of *intra*-day returns are not consistent with those of a martingale difference sequence. Instead, modelling *intra*-day returns as an ARMA<sup>10</sup> process is a natural and, as it turns out, successful approach for it is well suited to account for the serial dependence of returns at various sampling frequencies. From a market microstructure point of view, the AR part will arguably be able to capture any autocorrelation induced by non-synchronous trading while the MA part will account for potential negative first order autocorrelation induced by the bid-ask bounce. Further, the decreasing order and magnitude of serial correlation with the sampling frequency is, as it turns out, a consequence of temporal aggregation of the return process.

Suppose that returns at the highest sampling frequency,  $R_1$  (the  $t$  subscript is momentarily dropped for notational convenience), can be described as an ARMA(p,q) process:

$$\alpha(L) R_{1,i} = \beta(L) \varepsilon_{1,i},$$

where  $\alpha(L)$  and  $\beta(L)$  are lag polynomials of lengths  $p$  and  $q$  respectively. As before, we also assume that the return process is weakly stationary which justifies expression 4 and the analysis below. Consider the case where all the reciprocals of the roots of  $\alpha(L) = 0$ , denoted by  $\theta_1, \dots, \theta_p$ , lie inside the unit circle. The model through which the returns at an arbitrary sampling (or aggregation) frequency can be represented is derived using the results of Wei (1981) on temporal aggregation<sup>11</sup>. In particular, if  $R_1$  follows an ARMA(p,q) process, the returns sampled at frequency  $f$ , denoted by  $R_f$ , can be represented by an ARMA(p,r) process:

$$\prod_{j=1}^p (1 - \theta_j^f L^f) R_{f,i} = \prod_{j=1}^p \frac{1 - \theta_j^f L^f}{1 - \theta_j L} \frac{1 - L^f}{1 - L} \beta(L) \varepsilon_{f,i},$$

where  $r$  equals the integer part of  $p + \frac{q-p}{f}$  and  $\varepsilon_{f,i} = \sum_{j=0}^{f-1} \varepsilon_{1,fi-j}$ . Due to the invertibility of the AR polynomial, the above model can be rewritten in MA( $\infty$ ) form with parameters  $\{\psi_j\}_{j=0}^{\infty}$  and  $\psi_0 = 1$ . Let  $\varphi_h^f$  denote the  $h^{\text{th}}$  autocovariance of the temporally aggregated returns at frequency  $f$ :

$$\varphi_h^f = E[R_{f,i} R_{f,i-h}] \propto \sum_{j=0}^{\infty} \left[ \left( \sum_{i=\max(0, j-f+1)}^j \psi_i \right) \left( \sum_{i=j+1+f(h-1)}^{j+fh} \psi_i \right) \right]. \quad (3)$$

It can be shown that the  $\psi_j$  coefficients decay exponentially fast in terms of  $j$  and, as a result, the autocovariances disappear under temporal aggregation. To see this, let  $|\psi_j| < w\delta^j$  for  $|\delta| < 1$  and  $w$

<sup>10</sup>More generally, one could specify an ARFIMA model for returns, thereby allowing for a hyperbolic decay of serial correlation. However, market microstructure and efficiency considerations aside, casual inspection of Table 1 and Figure 2 suggests that an ARMA process is sufficiently flexible to capture the dynamics of the returns process at high frequency.

<sup>11</sup>Temporal aggregation for ARMA models is discussed in Brewer (1973), Tiao (1972), Wei (1981), Weiss (1984) and the VARFIMA in Marcellino (1999).

some positive constant and notice that:

$$\varphi_h^f \propto \sum_{j=0}^{\infty} \left[ \sum_{i=0}^j w \delta^i \sum_{i=j+f(h-1)}^{j+fh} w \delta^i \right] < \frac{w^2}{(1-\delta)^3} \delta^{f(h-1)},$$

from which it can be seen that the autocovariances of order higher than two disappear when either the sampling frequency,  $f$ , or the displacement,  $h$ , increases. While it does not follow from the above that the first order autocovariance term also disappears, Wei (1981) has shown that the limit model of an ARMA(p,q) process under temporal aggregation is indeed an ARMA(0,0) or equivalently white noise.

It is important to emphasize that these theoretical properties of the ARMA process appear very much in accordance with the empirical properties of the return process as reported in Table 1. In particular, at high sampling frequencies the ARMA model can account for the observed serial dependence while at lower sampling frequencies these dependencies die off as a consequence of temporal aggregation of the return process. In addition, as the ARMA(p,q) model converges to an ARMA(0,0) under temporal aggregation, the model specification for returns at high frequency does not necessarily conflict with a model for returns at low frequency.

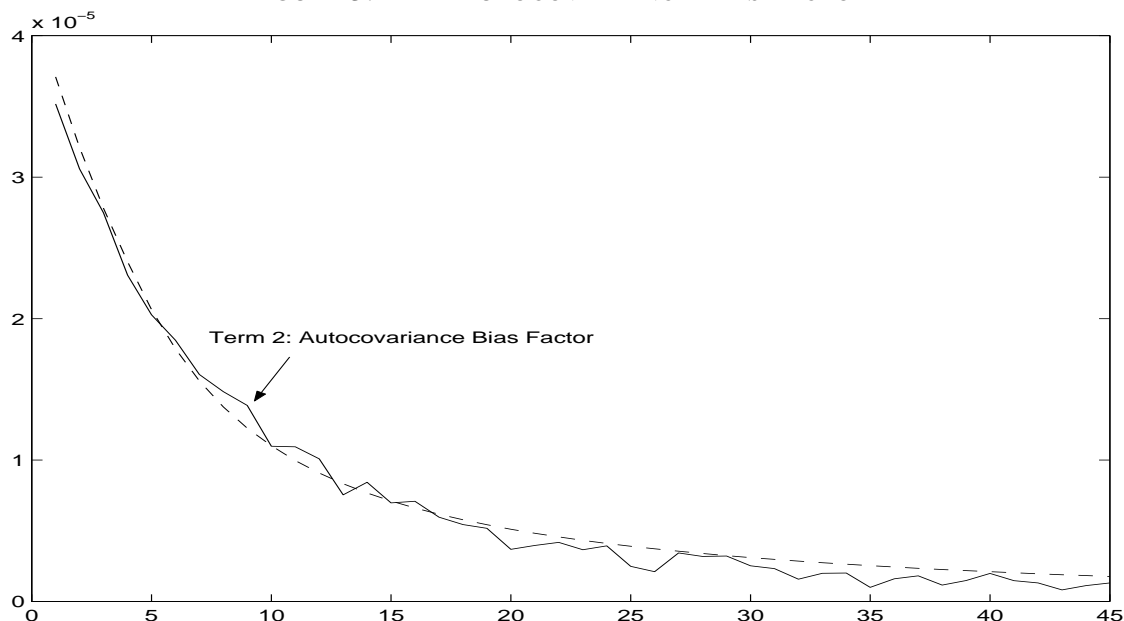
Relating the above aggregation results to the discussion of the previous section, we note that the expression for the autocovariance function of the ARMA process can be used to check the consistency of the model with the properties of the data by comparing the temporal aggregation implied decay of the autocovariance bias term with the empirically observed one. To this end, we estimate various ARMA models using the minute by minute returns on the FTSE-100 index and find that an ARMA(6,0) model yields satisfactory results<sup>12</sup>. Although the residuals are highly heteroskedastic, the OLS parameter estimates remain consistent (Amemiya 1985). Moreover, the efficiency loss due to the non-normality of the errors is unimportant given the large amount of data. Based on the *single* set of ARMA(6,0) parameters associated with the 1-minute data, the autocovariances for the estimated return process at various sampling frequencies can be *deduced* using expression (3). It is noted that:

$$E \left[ \sum_{i=1}^{N_f-1} \sum_{j=i+1}^{N_f} R_{f,t,i} R_{f,t,j} \right] = \sum_{h=1}^{N_f-1} (N_f - h) \varphi_h^f. \quad (4)$$

Hence, the “aggregation implied” autocovariance estimates can be used to calculate the “aggregation implied” autocovariance bias term as in expression (4). In particular, a single set of ARMA(6,0) parameters for the 1-minute data are used to imply the autocovariance bias factor at sampling frequencies between 1 and 45 minutes. Figure 3 demonstrates that the empirical and theoretically implied curves are remarkably close.

<sup>12</sup>Some of the higher order AR terms could arguably be replaced by low order MA terms. However, the AR specification has the advantage that inference is straightforward from a numerical point of view, as opposed to an MA specification. Since the AR and MA specification are largely equivalent preference is given here to the AR specification.

FIGURE 3: THE “AUTOCOVARANCE BIAS FACTOR”



Notes: The empirical autocovariance bias factor (solid line, see also Figure 1) and the superimposed aggregation implied autocovariance bias factor (dotted line, see also expressions (3) and (4)) for sampling frequencies between 1 and 45 minutes.

The above results illustrate that the ARMA model is a good description of the return data at different sampling frequencies. In fact, the decay of the (market microstructure-induced) serial dependencies in high frequency returns is consistent with the decay of an ARMA process under temporal aggregation. Also, it can be shown, based on expression 4, that the autocovariance bias term decays at a hyperbolic rate under temporal aggregation (i.e.  $\sum_{h=1}^{N_f-1} (N_f - h) \varphi_h^f < \frac{N_f^2}{f^2}$ ). Finally, we notice that it is possible to trace out the entire autocovariance bias factor curve, and hence determine the optimal frequency, using solely a *single* set of ARMA parameters.

In summary, we have shown that the conditional return variance can be estimated consistently by the realized variance measure, provided that the intra-day returns are serially uncorrelated. When the intra-day returns are serially correlated, realized variance will either overestimate (with negative correlation) or underestimate (with positive correlation) the conditional return variance. Correcting for the bias term by adding up the cross products of intra-day returns, they are known after all, is not desirable as this is equivalent to using the squared daily return to estimate daily realized variance. Here we suggest that when the available high frequency return data are serially correlated, one approach<sup>13</sup> is to aggregate the returns down to a frequency at which the correlation has disappeared, thereby avoiding

<sup>13</sup>An alternative approach would be to utilize all of the observations by explicitly modelling the high-frequency market microstructure. However, as noted by Andersen, Bollerslev, Diebold, and Ebens (2001), that approach is much more complicated and subject to numerous pitfalls of its own.

(potentially) large biases in the realized variance measure. Plotting the sum of squared intra-day returns or the autocovariance bias factor versus the sampling frequency, as is done in Figure 1 proves a very helpful and easily implementable strategy to determine the frequency at which the correlation has died off. Further analysis suggests that the decay of the autocovariance bias factor is consistent with an ARMA process under temporal aggregation. This finding provides an alternative, yet closely related, parametric approach to determining the optimal sampling frequency.

### 3 Modelling Realized Variance

A number of studies<sup>14</sup> have analyzed high frequency data for a variety of financial securities. Regarding the properties of the realized variance measure, several studies find that (i) the marginal distribution of realized variance is distinctly non-normal and extremely right skewed, whereas the marginal distribution of logarithmic realized variance is close to Gaussian, (ii) logarithmic realized variance displays a high degree of (positive) serial correlation which dies out very slowly (iii) logarithmic realized variance does not seem to have a unit root, but there is clear evidence of fractional integration<sup>15</sup>, roughly of order 0.40 and (iv) daily returns standardized by realized volatility<sup>16</sup>, i.e. the square root of realized variance, are close to Gaussian.

Based on the analysis in Sections 2.1 and 2.2, which indicates that the daily conditional return variance of the FTSE-100 can be estimated unbiasedly as the sum of squared intra-day returns sampled at a frequency of 25 minutes, a time series of (logarithmic) realized variance is constructed and is displayed in the left panel of Figure 4. Table 2 reports some descriptive statistics of the time series of realized variance and returns.

We find that our results are very much in line with the findings described above. In particular, the unconditional distribution of the realized variance appears significantly skewed and exhibits severe kurtosis, while the unconditional distribution of logarithmic realized variance is much less skewed and displays significantly reduced kurtosis (Table 2). Furthermore, the correlogram for the realized variance measure decays only very slowly but the Augmented Dickey Fuller test strongly rejects the null hypothesis of a unit root (Table 2 and right panel of Figure 4). This finding indicates that the (logarithmic) realized variance series may exhibit long memory, a feature that will be discussed below.

---

<sup>14</sup>See for example Andersen, Bollerslev, Diebold, and Labys (2000b, 2000b), Blair, Poon, and Taylor (2001), Dacorogna, Gençay, Müller, Olsen, and Pictet (2001), Froot and Perold (1995), Goodhart and O'Hara (1997), Hsieh (1991), Lequeux (1999), Stoll and Whaley (1990), Zhou (1996).

<sup>15</sup>See for example Baillie (1996), Baillie, Bollerslev, and Mikkelsen (1996), Breidt, Crato, and de Lima (1998), Comte and Renault (1998), Henry and Payne (1998), Liu (2000), Lo (1991).


<sup>16</sup>In a multivariate setting it is found that the distribution of correlations between realized variance is close to normal with positive mean, and that the autocorrelations of realized correlation decays extremely slow.

TABLE 2: DESCRIPTIVE STATISTICS OF REALIZED VARIANCE AND RETURNS.

	Mean	Volatility	Skewness	Kurtosis	ADF[5]
Realized Variance	8.5e-5	2.6e-4	21.21	596	-16.2
Log Realized Variance	-9.98	0.962	0.558	4.11	-8.83
Daily Returns	4.6e-4	0.009	0.063	5.29	-21.8
Standardized Daily Return	0.091	1.091	0.036	2.23	-22.3

Notes: Descriptive statistics based on the FTSE-100 data set for the period 1990-2000. The augmented Dickey Fuller test (“ADF[5]”) includes a constant and 5 lags and has a 5% (1%) critical value of -2.865 (-3.439).

FIGURE 4: LOGARITHMIC REALIZED VOLATILITY



Notes: Time series (left panel) and correlogram (right panel) of FTSE-100 daily logarithmic realized variance constructed at a sampling frequency of 25 minutes over the period 1990-2000. The superimposed dotted lines in the right panel represent the correlogram of a fractional process for values of  $d$  equal to 0.30, 0.40, and 0.45.

Finally, daily returns standardized by realized variance are close to normal (Table 2). This indicates that the empirical findings obtained by Andersen, Bollerslev, Diebold, and Labys (2000a) on exchange rate data can be extended to the FTSE-100 stock market index data.

### 3.1 Fractional Integration & Realized Variance

A time series,  $X_t$  is said to be fractionally integrated of order  $d$  if after applying the difference operator  $(1 - L)^d$  it follows a stationary ARMA( $p, q$ ) process where  $p$  and  $q$  are finite nonnegative integers. This concept has been developed by Granger (1980), Granger (1981), and Granger and Joyeux (1980). For

values of  $d$  between 0 and 0.5, the fractionally integrated process<sup>17</sup> exhibits “long memory” which has the property that the effect of a shock to the process is highly persistent but decays over time. This is in contrast to  $I(1)$  processes, where a shock has infinite persistence, or at the other extreme  $I(0)$  processes, where the effect of a shock decays exponentially fast. The ARFIMA( $p,d,q$ ) model can be written as

$$\alpha(L)(1 - L)^d X_t = \beta(L)\varepsilon_t, \quad (5)$$

where  $\alpha(L)$  and  $\beta(L)$  are lag polynomials of order  $p$  and  $q$  respectively. For  $d < \frac{1}{2}$  and  $d \neq 0$ , it can be shown that the decay of the correlogram is *hyperbolic*, i.e.

$$\varphi_h = \text{corr}(X_t, X_{t-h}) = \frac{\Gamma(1-d)}{\Gamma(d)} \frac{\Gamma(h+d)}{\Gamma(h+1-d)} \underset{h \text{ large}}{\propto} h^{2d-1}. \quad (6)$$

Regarding the estimation of  $d$ , Geweke and Porter-Hudak (1983, GPH hereafter) propose the use of a log periodogram regression. In particular, for given  $\{X_t\}_{t=1}^T$ , the fractional parameter  $d$  can be estimated as the slope coefficient in a linear regression of  $I(\lambda_j) = \frac{1}{2\pi T} \left| \sum_{t=1}^T X_t e^{i\lambda_j t} \right|^2$ , the log periodogram at harmonic frequency  $\lambda_j = \frac{2\pi j}{T}$ , on a constant and  $\ln [4 \sin^2(\lambda_j/2)]$  for  $j = 1, \dots, m \ll T$ . The “bandwidth” parameter  $m$  is required to increase at a slower rate than the sample size  $T$  and in many applications  $m$  is set to equal to the square root of the sample size  $T$ . Robinson (1995a, 1995b) derives an alternative estimator for  $d$ , which is shown to be asymptotically more efficient than the GPH estimator, and is given by the value of  $d$  that minimizes the following objective function:

$$Q(c, d) = \frac{1}{m} \sum_{j=1}^m \left[ \ln(c\lambda_j^{-2d}) + \frac{\lambda_j^{2d}}{c} I(\lambda_j) \right],$$


where  $c > 0$  and  $-\frac{1}{2} < d < \frac{1}{2}$ .

Turning to the FTSE-100 realized variance series, it is clear that long memory features are very much present. The right panel of Figure 4 displays the correlogram of the log realized variance series while the right panel of Figure 5 displays the correlogram of the *fractionally differenced* log realized variance series based on an ad hoc parameter value<sup>18</sup> of  $d = 0.40$ . The serial correlations of the log realized variance series decay at a hyperbolic rate and the resemblance between the sample correlogram and the superimposed correlograms of a fractionally integrated process for various values of  $d$  is remarkable. In sharp contrast, the fractionally differenced series is virtually uncorrelated. A supplementary diagnostic check for the presence of long memory is based on expression (6) above. In particular, when the realized variance series exhibits long memory, its log autocorrelation function should yield a linear relationship in terms of log displacement, i.e.  $\ln \varphi_h \propto (2d - 1) \ln h$ . Figure 6 (left panel) indicates the

<sup>17</sup>The process is stationary with long memory for  $0 < d < 0.5$  but stationary with intermediate memory for  $-0.5 < d < 0$ . For  $d \geq 0.5$ , the process is non-stationary.

<sup>18</sup>Various values between 0.35 and 0.45 have been used but the results appear robust to the specific choice of  $d$ .

FIGURE 5: FRACTIONALLY DIFFERENCED LOGARITHMIC REALIZED VOLATILITY



Notes: Time series (left panel) and correlogram (right panel) of FTSE-100 *fractionally differenced* daily logarithmic realized variance constructed at a sampling frequency of 25 minutes over the period 1990-2000. The dotted lines in the right panel are the 95% confidence bounds calculated as  $\pm 2N^{-1/2}$  where  $N$  denotes the number of observations.

required linear relationship between  $\ln \varphi_h$  and  $\ln h$  for values of  $h$  up to 100. An OLS regression can be used to determine the slope. Based on the entire sample ( $h = 250$ ) the results suggest a value for  $d$  of around 0.37. Ignoring the last 150 autocorrelations ( $h = 100$ ) raises  $d$  to about 0.43. Finally, the GPH and Robinson estimators, described above, are implemented. The bandwidth parameter  $m$  (controlling the range of periodic frequencies used), is set equal to a range of values between<sup>19</sup> 25 and 275. The results of this estimation are summarized in Figure 6 (right panel) where the GPH and Robinson estimates are plotted as a function of  $m$ . For small  $m$ , the two alternative estimates both fall into the non-stationary region while for large  $m$  (above 150) they are both below 0.5. Although it is clear from this that the value for  $d$  will be close to 0.5, it is difficult to judge on the stationarity of the process as the choice of  $m$  is relatively arbitrary. In summary, all of the test results reported above suggest that the FTSE-100 log realized variance series is fractionally integrated and appear roughly consistent with Andersen, Bollerslev, Diebold, and Ebens (2001) who find that for their data set  $d$  is around 0.40.

### 3.2 Empirical Results


Motivated by the preliminary tests discussed above, the focus of our modelling approach will center around the ARFIMA specification. We consider the following model:

$$\alpha(L)(1 - L)^d [\ln \hat{\sigma}_{25,t}^2 - \pi' X_t] = \beta(L)\varepsilon_t, \quad (7)$$

<sup>19</sup>The sample size is 2445 and hence the range of  $m$  is between  $T^{0.40}$  and  $T^{0.70}$ . This is in line with e.g. Bollerslev, Cai, and Song (2000) which set  $m = T^{0.50}$  or Dittmann and Granger (2002) which set  $m = T^{0.8}$ .



FIGURE 6: TWO TESTS FOR FRACTIONAL INTEGRATION



Notes: Two tests for fractional integration. Linearity of  $\ln \varphi_h$  versus  $\ln h$  (left panel) and the Geweke Porter-Hudak and Robinson estimate for  $d$  as a function of the bandwidth  $m$  (right panel).

where  $\hat{\sigma}_{25,t}^2$  denotes the day- $t$  realized variance measure constructed based on 25 minute intra-day returns,  $\alpha(L)$  is a lag polynomial of order  $p$ ,  $\beta(L)$  a lag polynomial of order  $q$ , and  $\varepsilon_t$  is a residual error term. The  $k \times 1$  vector  $X_t$  allows for the inclusion of exogenous variables and deterministic terms such as a constant and time trend. Here, we consider the following specification:

$$\pi' X_t = \omega + \sum_{j=1}^k (\zeta_j R_{t-j} + \bar{\zeta}_j |R_{t-j}|) + \sum_{j=0}^m \lambda_j \ln VOL_{t-j} + \sum_{j=0}^n \delta_j (IR_{t-j} - IR_{t-j-1}) \quad (8)$$

where  $IR_t$  and  $VOL_t$  denote the day- $t$  short term interest rate (1 month UK Interbank rate) and daily trading volume respectively. The inclusion of lagged returns and lagged absolute returns mirrors the EGARCH specification of Nelson (1991) and is, in part, motivated by the well documented Black's leverage effect or the asymmetric impact that lagged returns have on the return variance. In particular, Black (1976) argues that one should expect negative returns to have a larger impact on future variance than positive returns. In the above specification we can test whether such a leverage effect is present at horizon  $h$  by testing whether  $\zeta_h$  is significantly less than zero<sup>20</sup>.

Next, trading volume is included because it is often argued that it is intimately related to the return variance. A model which can rationalize such a relationship has been proposed by Clark (1973) where prices follow a subordinated process with information flow (proxied by trading volume or number of

<sup>20</sup>Suppressing subscripts momentarily, define  $R^+ = R$  when  $R > 0$  and  $R^+ = 0$  when  $R \leq 0$ . Similarly, define  $R^- = -R$  when  $R \leq 0$  and  $R^- = 0$  when  $R > 0$ . Hence,  $R = R^+ - R^-$  and  $|R| = R^+ + R^-$ . It is now straightforward to show that  $\zeta R + \bar{\zeta} |R| = \zeta^+ R^+ + \zeta^- R^-$  where  $\zeta^+ = \zeta + \bar{\zeta}$  and  $\zeta^- = \bar{\zeta} - \zeta$ . For the leverage effect to be present, it is required that  $\zeta^- > \zeta^+ \Leftrightarrow \zeta^- - \zeta^+ > 0 \Leftrightarrow \zeta < 0$ .

trades) being the subordinator. A number of papers have addressed the relationship from an empirical point of view (e.g. Karpoff (1987), Gallant, Rossi, and Tauchen (1992) and more recently Ané and Geman (2000)) and invariably report positive correlation between return variance and trading volume. In addition, an influential paper by Lamoureux and Lastrapes (1990) finds that the *persistence* of return variance decreases (or even disappears) when trading volume is accounted for. Finally, the inclusion of (changes in) the short term interest rate is motivated by Glosten, Jaganathan, and Runkle (1993) who find that it has a significant positive effect on stock market volatility.

Before moving on to the estimation results, we point out that the above specification does not allow us to study the *causal* relation between return volatility and trading volume. In particular, it could well be that, in addition to trading volume causing return volatility, return volatility also has a feedback effect onto subsequent trade activity. Whether such dynamics can be identified at a daily frequency is questionable but are clearly of interest. The theoretical market microstructure has studied such relationships extensively. However, the primary focus has been on the impact of trade duration on the price process and results are mixed (see for example, Admati and Pfleiderer (1988), Diamond and Verrecchia (1987), Easley and O'Hara (1992), and Glosten and Milgrom (1985)). Engle (2000) has also focussed on the impact of trade durations on the price process. Using IBM high frequency data, he finds that low trading activity leads to a reduction in future return volatility (supporting the implications of the Easley and O'Hara (1992) model). A related study by Renault and Werker (2002) investigates the *instantaneous* causality relation between transaction durations and prices and finds that about two-thirds of return volatility can be attributed to instantaneous durations - in other words - transaction times cause transaction prices. Under the assumption that trade durations are inversely proportional to trade volume, the model we have specified in (7) and (8) is directly in line with the above mentioned work, although it should be kept in mind that we work with data at a daily frequency as opposed to transaction level data. The feedback effect of return volatility on trade durations - or trade volume - is, although of interest, not studied here.

Under the assumptions that (i) the roots of  $\alpha(L)$  are simple and lie outside the unit circle, (ii) the residuals are i.i.d. Gaussian, and (iii)  $d < \frac{1}{2}$ , the ARFIMA model, specified by (7) and (8) above, can be estimated<sup>21</sup> using the maximum likelihood procedure of Sowell (1992). Alternatively, the model could have been estimated using a two-step procedure in which the fractional parameter is estimated in the first step (e.g. with the GPH or Robinson estimator), while the remaining ARMA coefficients are estimated in the second step based on the fractionally differenced data using ordinary least squares. However, as documented by Smith, Taylor, and Yadav (1997), such an approach may well lead to inaccurate or biased ARMA coefficient estimates. The Sowell procedure, allowing for the simultaneous estimation of the model parameters, is therefore preferred.

<sup>21</sup>We have used the ARFIMA package in PcGive version 10.0. See Doornik and Ooms (1999) and Doornik (2002) for documentation.

TABLE 3: ARFIMA ESTIMATION RESULTS

Par	Full Sample (1990-2000)				Sub Sample (1990-1997)			
	ARFIMA	+ Returns	+ Volume	+ Interest	ARFIMA	+ Returns	+ Volume	+ Interest
$d$	0.483 (22.6)	0.476 (16.6)	0.484 (24.0)	0.484 (23.8)	0.441 (8.74)	0.391 (6.73)	0.476 (16.5)	0.475 (16.2)
$\alpha_1$	0.356 (5.76)	0.337 (7.27)	0.337 (7.94)	0.335 (7.87)	0.437 (5.70)	0.325 (4.69)	0.321 (6.42)	0.318 (6.30)
$\beta_1$	-0.602 (10.3)	-0.678 (15.2)	-0.695 (18.9)	-0.695 (18.8)	-0.635 (7.75)	-0.599 (6.99)	-0.680 (14.8)	-0.678 (14.5)
$\zeta_1$	-	-3.377 (5.03)	-3.863 (5.99)	-3.999 (6.18)	-	-2.362 (2.60)	-3.708 (4.25)	-3.958 (4.51)
$\zeta_2$	-	-2.652 (3.96)	-3.086 (4.79)	-3.107 (4.81)	-	-1.906 (2.10)	-2.834 (3.25)	-2.827 (3.22)
$\zeta_3$	-	-1.439 (2.14)	-1.472 (2.28)	-1.449 (2.25)	-	-2.448 (2.69)	-2.312 (2.65)	-2.290 (2.63)
$\zeta_4$	-	-1.206 (1.79)	-1.498 (2.32)	-1.527 (2.37)	-	-0.925 (1.02)	-1.393 (1.61)	-1.404 (1.62)
$\bar{\zeta}_1$	-	30.36 (28.0)	27.01 (25.3)	27.15 (25.4)	-	37.22 (25.4)	32.73 (22.8)	32.99 (23.0)
$\bar{\zeta}_2$	-	12.85 (11.7)	11.30 (10.5)	11.29 (10.5)	-	16.99 (11.5)	14.89 (10.3)	14.80 (10.2)
$\bar{\zeta}_3$	-	6.468 (5.91)	5.817 (5.52)	5.729 (5.44)	-	6.288 (4.27)	5.812 (4.11)	5.610 (3.97)
$\bar{\zeta}_4$	-	4.541 (4.19)	4.413 (4.24)	4.355 (4.19)	-	3.953 (2.70)	3.926 (2.81)	3.851 (2.76)
$\lambda_0$	-	-	0.338 (14.0)	0.335 (13.9)	-	-	0.370 (13.1)	0.365 (12.9)
$\lambda_1$	-	-	-0.007 (0.28)	-0.007 (0.30)	-	-	-0.011 (0.39)	-0.010 (0.37)
$\delta_0$	-	-	-	-0.179 (2.29)	-	-	-	-0.198 (2.41)
$\delta_1$	-	-	-	0.031 (0.39)	-	-	-	0.068 (0.83)
-LogL	977.3	607.7	504.5	501.6	710.5	414.0	328.8	325.4
AIC/T	0.805	0.509	0.426	0.425	0.795	0.475	0.382	0.381
No. Par	5	13	15	17	5	13	15	17
Skew	0.675	0.371	0.385	0.380	0.721	0.330	0.345	0.345
Kurt	5.680	4.291	4.184	4.144	5.850	3.924	3.785	3.734
PM[5]	3.888 (0.143)	2.822 (0.244)	1.945 (0.378)	2.095 (0.351)	5.391 (0.068)	4.906 (0.086)	4.365 (0.113)	4.516 (0.105)
ARCH[5]	4.443 (0.001)	4.143 (0.001)	2.662 (0.021)	2.028 (0.072)	3.402 (0.005)	5.389 (0.000)	4.718 (0.000)	3.406 (0.005)

Notes: ARFIMA(1,d,1) estimation results for the full sample (2 May 1990 - 11 January 2000; 2445 observations) and the sub sample (2 May 1990 - 15 June 1997; 1803 observations). The full model specification is given by expressions (7) and (8). The table reports all parameter estimates (except  $\omega$ ) with absolute t-statistics in parenthesis below. The residual test statistics include skewness (“Skew”), kurtosis (“Kurt”), and the Portmanteau (“PM[5]”,  $\mathcal{X}_2^2$ ) and ARCH (“ARCH[5]”,  $F(5, 1775)$  for sub-sample and  $F(5, 2419)$  for full sample) statistics including 5 lags. p-values are reported in parenthesis below PM[5] and ARCH[5].

We first estimate the model without any exogenous variables and then subsequently add returns, volume, and the short rate. To address the concern that long memory may be induced by infrequent structural breaks<sup>22</sup>, we re-estimate the model on various subsamples of the data set. Table 3, summarizes the estimation results<sup>23</sup> for two different samples and  $p = q = 1$ ,  $k = 4$ , and  $m = n = 2$ . The first sample is the full sample while the second sample covers the period May 1, 1990 until June 15, 1997. As the point estimates for the fractional parameter remain within a tight range (with one exception, all estimates are between 0.44 and 0.48) and turn out to be highly significant irrespective of the sample period or the model specification, we argue that the realized variance series clearly exhibits a long memory feature that is not caused by structural breaks. Based on the  $t$ -statistic<sup>24</sup>, however, we cannot reject that  $d > 0.5$  at a 95% confidence level, i.e. the realized variance series is potentially non-stationary. Turning to the exogenous variables, we notice a dramatic increase in log likelihood - accompanied by a substantial decrease in AIC criterion - upon inclusion of lagged (absolute) returns. In particular, for  $k = 4$ , the number of parameters increases by 8 to a total of 13 while the log likelihood increases by almost 370! As a result, the AIC criterion drops from 0.80 to 0.50. Further, the sign and significance of the  $\zeta$  parameters suggest that Black's leverage effect is present at horizons up to 3 or 4 days. This finding provides support for the GJR-GARCH (Glosten, Jaganathan, and Runkle 1993) and EGARCH (Nelson 1991) specifications which explicitly account for this asymmetric effect that returns have on future variance. Regarding trading volume, we find that contemporaneous values further improve the fit of the model. Consistent with Clark's model, we find that the sign of  $\lambda_0$  is positive and highly significant. However, in contrast to the findings of Lamoureux and Lastrapes, it appears that the *persistence* of the variance process (as measured by  $d$ ) remains largely unchanged when trading volume is conditioned upon. Finally, the estimate for  $\delta_0$  suggests that an interest rate cut is accompanied by higher volatility than an interest rate hike. It must be said, however, that this effect is marginally significant and that the associated likelihood increase minimal. As for trading volume, lagged changes in the interest rate are found to have an insignificant impact. A similar pattern is observed for the sub-sample.

<sup>22</sup>See for example Diebold and Inoue (2001), Engle and Smith (1999), Granger (1999), and Granger and Hyung (1999). A simple and representative model that can cause long memory is the stochastic break model, which takes the form:  $y_t = u_t + \varepsilon_t$ , where  $u_t = u_{t-1} + q_{t-1}\eta_t$ ,  $\varepsilon_t \sim iid\mathcal{N}(0, \sigma_y^2)$ ,  $\eta_t \sim iid\mathcal{N}(0, \sigma_u^2)$  and  $q_t$  equals 0 with probability  $p$  and 1 with probability  $1 - p$ . Diebold and Inoue (2001) note that in order to achieve a slowly declining autocorrelation function, whatever the model may be, the key idea is to let  $p$  decrease with the sample size so that regardless of the sample size, realizations of the process tend to have just a few breaks.

<sup>23</sup>Based on the likelihood ratio test and the AIC criterion we find that an ARFIMA(1,d,1) model provides a parsimonious specification. The choice of  $k$ ,  $m$ , and  $n$  is guided by the significance of the parameters.  $\zeta_4$ ,  $\bar{\zeta}_4$ ,  $\lambda_1$ , and  $\delta_1$  are included for completeness.

<sup>24</sup>The validity of the  $t$ -statistics crucially relies on whether the residuals are IID Gaussian. The diagnostic tests reported in Table 3 indicate that even though the residuals appear uncorrelated some skewness, kurtosis and heteroskedasticity is present. Fortunately, these effects diminish to some extent when lagged returns and trading volume are included and we will therefore work under the assumption that the  $t$ -statistics - in particular for the full model - are reasonably accurate.

## 4 Conclusion

Under certain assumptions on the return process, a number of recent papers have shown that realized variance is a consistent and virtually measurement error-free estimator of the conditional return variance. In this paper we show that realized variance measure constitutes a *biased* estimate of the return variance when (excess) returns are serially correlated. 10 years of FTSE-100 minute by minute data are used to illustrate that a careful choice of sampling frequency is crucial in avoiding a substantial bias. The relation between the sampling frequency and the presence of serial correlation is analyzed in detail and demonstrates that serial correlation in returns disappears under temporal aggregation at a rate of decay that is consistent with that one of an ARMA process. An autocovariance function based method is proposed for choosing the optimal sampling frequency, that is, the highest available sampling frequency for which the autocovariance bias term is negligible. Many alternative approaches to deal with this issue can be considered though. One route is to use all available data by explicitly modelling the market microstructure effects. Another is to “correct” for the bias by dividing the biased realized variance estimate by an appropriate constant (or any sort of function that achieves unbiasedness of the estimator). A third approach, which we may explore in future research, is to use a Newey-West type covariance estimator in order to take into account the serial correlation in the data. The advantage here is that it is potentially more efficient than the aggregation approach outlined in this paper as it makes use of all available data while the non-parametric nature of the estimator avoids the need to explicitly model the market microstructure.

Regarding the FTSE-100 data set, we find that the realized variance series can be modelled as an ARFIMA process. Exogenous variables such as lagged returns and contemporaneous trading volume appear to be highly significant regressors and are able to explain a large portion of the variation in realized variance. Also, statistical tests suggest that Black’s leverage effect is significant at three or four days. Regarding contemporaneous trading volume we find that, despite its significance, the persistence of the variance process remains largely unchanged.

## References

- Admati, A., and P. Pfleiderer, 1988, “A Theory of Intraday Patterns: Volume and Price Variability,” *Review of Financial Studies*, 1 (1), 3–40.
- Ait-Sahalia, Y., and P. A. Mykland, 2003, “How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise,” NBER working paper 9611.
- Amemiya, T., 1985, *Advanced Econometrics*. Harvard University Press, Boston, MA.
- Andersen, T. G., and T. Bollerslev, 1998, “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts,” *International Economic Review*, 39(4), 885–905.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens, 2001, “The Distribution of Stock Return Volatility,” *Journal of Financial Economics*, 61, 43–76.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys, 2000a, “Exchange Rate Returns Standardized by Realized Volatility are (nearly) Gaussian,” *Multinational Finance Journal*, 4, 159–179.
- , 2000b, “Great Realizations,” *Risk*, pp. 105–108.
- , 2001, “The Distribution of Realized Exchange Rate Volatility,” *Journal of the American Statistical Association*, 96, 42–55.
- , 2003, “Modeling and Forecasting Realized Volatility,” *Econometrica*, 71 (2), 579–625.
- Andreou, E., and E. Ghysels, 2001, “When Does Microstructure Noise Affect the Estimation of Asset Pricing Models?,” Manuscript University of Manchester and North Carolina.
- Ané, T., and H. Geman, 2000, “Order Flow, Transaction Clock, and Normality of Asset Returns,” *Journal of Finance*, 55(5), 2259–2284.
- Areal, N. M., and S. J. Taylor, 2002, “The Realized Volatility of FTSE-100 Futures Prices,” *Journal of Futures Markets*, 22 (7), 627–648.
- Bai, X., J. R. Russell, and G. C. Tiao, 2001, “Beyond Merton’s Utopia (I): Effects of Non-Normality and Dependence on the Precision of Variance Estimates using High Frequency Data,” Manuscript University of Chicago.
- Baillie, R. T., 1996, “Long Memory Processes and Fractional Integration in Econometrics,” *Journal of Econometrics*, 73, 5–59.

- Baillie, R. T., T. Bollerslev, and H. O. Mikkelsen, 1996, "Modelling and Pricing Long Memory in Stock Market Volatility," *Journal of Econometrics*, 73, 151–184.
- Bakshi, G., C. Cao, and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Model," *Journal of Finance*, 52(5), 2003–2049.
- Barndorff-Nielsen, O. E., and N. Shephard, 2002, "Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society, Series B*, 64, 252–280.
- Barndorff-Nielsen, O. E., and N. Shephard, 2003, "Realised Power Variation and Stochastic Volatility Models," forthcoming *Bernoulli* 9.
- Black, F., 1976, "Studies of Stock Price Volatility Changes," Proceeding from the American Statistical Association, Business and Economic Statistics Section.
- Blair, B. J., S.-H. Poon, and S. J. Taylor, 2001, "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns," *Journal of Econometrics*, 105, 5–26.
- Bollerslev, T., J. Cai, and F. M. Song, 2000, "Intra-day Periodicity, Long Memory Volatility and Macroeconomic announcement effects in the US Treasury Bond Market," *Journal of Empirical Finance*, 7, 37–55.
- Bollerslev, T., R. F. Engle, and D. B. Nelson, 1994, "ARCH Models," In: Engle, R.F., McFadden, D. (Eds.), *Handbook of Econometrics*, vol. 4, North-Holland, Amsterdam, pp. 2959–3038.
- Breidt, F. J., N. Crato, and P. de Lima, 1998, "The Detection and Estimation of Long Memory in Stochastic Volatility," *Journal of Econometrics*, 83, 325–348.
- Brewer, K., 1973, "Some Consequences of Temporal Aggregation and Systematic Sampling for ARMA and ARMAX Models," *Journal of Econometrics*, 1, 133–154.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.
- Clark, P. K., 1973, "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41(1), 135–155.
- Comte, F., and E. Renault, 1998, "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291–323.

- Corsi, F., G. Zumbach, U. A. Müller, and M. Dacorogna, 2001, “Consistent High-Precision Volatility from High-Frequency Data,” Olsen Group Working Paper.
- Dacorogna, M. M., R. Gençay, U. Müller, R. B. Olsen, and O. V. Pictet, 2001, *An Introduction to High-Frequency Finance*. Academic Press, London, UK.
- Davison, A., and D. Hinkley, 1997, *Bootstrap Methods and their Applications*. Cambridge University Press, Cambridge, UK.
- Diamond, D., and R. Verrecchia, 1987, “Constraints on Short-Selling and Asset Price Adjustments to Private Information,” *Journal of Financial Economics*, 18, 277–311.
- Diebold, F. X., and A. Inoue, 2001, “Long Memory and Regime Switching,” *Journal of Econometrics*, 105, 131–159.
- Dittmann, I., and C. W. Granger, 2002, “Properties of Nonlinear Transformations of Fractionally Integrated Processes,” *Journal of Econometrics*, 110 (2), 113–133.
- Doornik, J. A., 2002, *Object-Oriented Matrix Programming Using Ox*. Timberlake Consultants Press, London, UK, 3 edn.
- Doornik, J. A., and M. Ooms, 1999, “A Package for Estimating, Forecasting and Simulating ARFIMA Models,” manuscript Nuffield College, Oxford ([www.nuff.ox.ac.uk/Users/Doornik](http://www.nuff.ox.ac.uk/Users/Doornik)).
- Easley, D., and M. O’Hara, 1992, “Time and the Process of Security Price Adjustment,” *Journal of Finance*, 47 (2), 577–605.
- Engle, R. F., 2000, “The Econometrics of Ultra-High Frequency Data,” *Econometrica*, 68(1), 1–22.
- Engle, R. F., and D. Smith, 1999, “Stochastic Permanent Breaks,” *Review of Economics & Statistics*, 81 (4), 553–574.
- Figlewski, S., 1998, “Derivatives Risks, Old and New,” *Brookings Wharton Papers on Financial Services*, pp. 159–238.
- Foster, D. P., and D. B. Nelson, 1996, “Continuous Record Asymptotics for Rolling Sample Variance Estimators,” *Econometrica*, 64 (1), 139–174.
- French, K., and R. Roll, 1986, “Stock Return Variances: The Arrival of Information and the Reaction of Traders,” *Journal of Financial Economics*, 17, 5–26.
- French, K. R., G. W. Schwert, and R. Stambaugh, 1987, “Expected Stock Returns and Volatility,” *Journal of Financial Economics*, 19, 3–29.



- Froot, K. A., and A. F. Perold, 1995, "New Trading Practices and Short-Run Market Efficiency," *Journal of Futures Markets*, 15(7), 731–765.
- Gallant, A. R., P. E. Rossi, and G. Tauchen, 1992, "Stock Prices and Volume," *Review of Financial Studies*, 5, 199–242.
- Geweke, J. F., and S. Porter-Hudak, 1983, "The Estimation and Application of Long Memory Time Series Models," *Journal of Time Series Analysis*, 4(4), 221–238.
- Ghysels, E., A. Harvey, and E. Renault, 1996, "Stochastic Volatility," In *Handbook of Statistics (14) Statistical Methods in Finance*, eds. G.S. Maddala and C.R. Rao, Elsevier Amsterdam, 119-191.
- Glosten, L., and P. R. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.
- Glosten, L. R., R. Jagannathan, and D. E. Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779–1801.
- Goodhart, C. A., and M. O'Hara, 1997, "High Frequency Data in Financial Markets: Issues and Applications," *Journal of Empirical Finance*, 4, 73–114.
- Granger, C. W., 1980, "Long Memory Relationships and the Aggregation of Dynamic Models," *Journal of Econometrics*, 14, 227–238.
- , 1981, "Some Properties of Time Series Data and Their Use in Econometric Model Specification," *Journal of Econometrics*, 16, 121–130.
- , 1999, "Aspects of Research Strategies for Time Series Analysis," Outline of presentation to the conference on New Developments in Time Series Economics at the Cowless Foundation, New Haven.
- Granger, C. W., and N. Hyung, 1999, "Occasional Structural Breaks and Long Memory," UCSD discussion paper 99-14.
- Granger, C. W., and R. Joyeux, 1980, "An Introduction to Long-Memory Time Series and Fractional Differencing," *Journal of Time Series Analysis*, 1, 15–30.
- Henry, M., and R. Payne, 1998, "An Investigation of Long Range Dependence in Intra-Day Foreign Exchange Rate Volatility," Financial Markets Group discussion paper No. 264.
- Horowitz, J. L., I. Lobato, J. C. Nankervis, and N. Savin, 2002, "Bootstrapping the Box-Pierce Q Test: A Robust Test for Uncorrelatedness," manuscript Northwestern University, Department of Economics.

- Hsieh, D. A., 1991, "Chaos and Nonlinear Dynamics: Application to Financial Markets," *Journal of Finance*, 46(5), 1839–1877.
- Jiang, G. J., and R. C. Oomen, 2002, "Hedging Derivatives Risks," Manuscript University of Arizona and Warwick Business School.
- Karpoff, J., 1987, "The Relation Between Price Change and Trading Volume: A Survey," *Journal of Financial & Quantitative Analysis*, 22, 109–126.
- Lamoureux, C. G., and W. D. Lastrapes, 1990, "Heteroskedasticity in Stock Return Data: Volume versus GARCH effects," *Journal of Finance*, 45(1), 221–229.
- Lequeux, P., 1999, *Financial Markets Tick by Tick: Insights in Financial Market Micro Structure*. John Wiley & Sons Ltd, England.
- Liu, M., 2000, "Modelling Long Memory in Stock Market Volatility," *Journal of Econometrics*, 99, 139–171.
- Lo, A. W., 1991, "Long Memory in Stock Market Prices," *Econometrica*, 59(5), 1279–1313.
- Lo, A. W., and A. C. MacKinlay, 1990, "An Econometric Analysis of Nonsynchronous-Trading," *Journal of Econometrics*, 45, 181–212.
- Madhavan, A. N., 2000, "Market Microstructure: A Survey," *Journal of Financial Markets*, 3, 205–258.
- Maheu, J. M., and T. H. McCurdy, 2002, "Nonlinear Features of Realized FX Volatility," *Review of Economics & Statistics*, 84 (4), 668–681.
- Marcellino, M., 1999, "Some Consequences of Temporal Aggregation in Empirical Analysis," *Journal of Business & Economic Statistics*, 17(1), 129–136.
- Martens, M., 2002, "Measuring and Forecasting S&P 500 Index-Futures Volatility Using High-Frequency Data," *Journal of Futures Markets*, 22 (6), 497–518.
- Merton, R. C., 1980, "On Estimating the Expected Return on the Market; An Exploratory Investigation," *Journal of Financial Economics*, 8, 323–361.
- Nelson, D. B., 1991, "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59 (2), 347–370.
- O'Hara, M., 1995, *Market Microstructure Theory*. Blackwell Publishers Ltd, Oxford, UK.

- Oomen, R. C., 2003, "Statistical Models for High Frequency Security Prices," Manuscript Warwick Business School.
- Politis, D. N., and J. P. Romano, 1994, "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89 (428), 1303–1313.
- Renault, E., and B. J. Werker, 2002, "Stochastic Volatility Models with Transaction Time Risk," manuscript Université de Montréal and Tilburg University.
- Robinson, P. M., 1995a, "Gaussian Semiparametric Estimation of Long Range Dependence," *Annals of Statistics*, 23, 1630–1661.
- , 1995b, "Log-Periodogram Regression of Time Series with Long Range Dependence," *Annals of Statistics*, 23(3), 1048–1072.
- Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 39(4), 1127–1139.
- Romano, J. P., and L. A. Thombs, 1996, "Inference for Autocorrelations under Weak Assumptions," *Journal of the American Statistical Association*, 91 (434), 590–600.
- Säfvenblad, P., 2000, "Trading Volume and Autocorrelation: Empirical Evidence from the Stockholm Stock Exchange," *Journal of Banking & Finance*, 24, 1275–1287.
- Smith, J., N. Taylor, and S. Yadav, 1997, "Comparing the Bias and Misspecification in ARFIMA models," *Journal of Time Series Analysis*, 18(5), 507–527.
- Sowell, F., 1992, "Maximum Likelihood Estimation of Stationary Univariate Fractionally Integrated Time Series Models," *Journal of Econometrics*, 53, 165–188.
- Stoll, H. R., and R. E. Whaley, 1990, "The Dynamics of Stock Index and Stock Index Futures Returns," *Journal of Financial & Quantitative Analysis*, 25(4), 441–468.
- Taylor, S. J., and X. Xu, 1997, "The Incremental Volatility Information in One Million Foreign Exchange Quotations," *Journal of Empirical Finance*, 4, 317–340.
- Tiao, G., 1972, "Asymptotic Behavior of Temporal Aggregates of Time Series," *Biometrika*, 59, 525–531.
- Wei, W., 1981, "Effects of Systematic Sampling on ARIMA Models," *Communications in Statistical-Theoretical Mathematics*, A10, 2389–2398.

- Weiss, A., 1984, "Systematic Sampling and Temporal Aggregation in Time Series Models," *Journal of Econometrics*, pp. 271–281.
- Wood, R., 2000, "Market Microstructure Research Databases: History and Projections," *Journal of Business & Economic Statistics*, 18(2), 140–145.
- Zhou, B., 1996, "High Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business & Economic Statistics*, 14(1), 45–52.