THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Heteroscedasticity and interval effects in estimating beta

Citation for published version:

Armitage, S & Brzezczynski, J 2011, 'Heteroscedasticity and interval effects in estimating beta: UK evidence', *Applied Financial Economics*, vol. 21, no. 20, pp. 1525-1538.
<https://doi.org/10.1080/09603107.2011.581208>

Digital Object Identifier (DOI):

[10.1080/09603107.2011.581208](https://doi.org/10.1080/09603107.2011.581208)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Applied Financial Economics

Publisher Rights Statement:

This is an Accepted Manuscript of an article published in Applied Financial Economics on 18 Jul 2011, available online: <http://www.tandfonline.com/10.1080/09603107.2011.581208>.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Heteroscedasticity and interval effects in estimating beta:

UK evidence

Seth Armitage *

University of Edinburgh

Janusz Brzeszczynski **

Heriot-Watt University, Edinburgh

April 2011

Applied Financial Economics 21, 2011, pp. 1525-38

Acknowledgements: We would like to thank Tony Naughton, Abu Mollik, and Bruce Vanstone, participants at the Australasian Finance and Banking Conference 2009, and seminar participants at Heriot-Watt University, for comments on this paper. We are grateful to James Hamilton for a very helpful discussion regarding the estimation issues.

* *Corresponding author:* University of Edinburgh Business School, 50 George Square, Edinburgh, EH8 9JY, UK, Tel.: +44 131 650 9663, Fax: +44 131 650 8077, E-mail: seth.armitage@ed.ac.uk

** Department of Accountancy, Economics and Finance, Heriot-Watt University, Edinburgh, EH14 4AS, UK, Tel.: +44 131 451 3294, Fax: +44 131 451 3296, E-mail: j.brzeszczynski@hw.ac.uk

Heteroscedasticity and interval effects in estimating beta:

UK evidence

Abstract

The paper compares beta estimates obtained from OLS regression with estimates corrected for heteroscedasticity of the error term using ARCH models, for 145 UK shares. The differences are mainly less than 0.10, for betas calculated using daily returns, but even such small differences can matter in practice. OLS tends to overestimate the beta coefficients compared with ARCH models, and selecting an ARCH-type estimate makes most difference for large-cap shares. Regarding the measurement interval, the downward bias in betas from daily returns is associated not only with thin trading but also with the volatility of the share's daily returns. We infer that the idiosyncratic component in daily returns, as well as lack of trading, is responsible for low daily betas.

Keywords: beta estimation, heteroscedasticity, ARCH models, interval effect

JEL codes: G15, C51

1. Introduction

The question of how best to estimate a share's beta is of obvious practical importance. Despite the doubtful performance of beta in explaining differences in observed share returns, recent surveys of company executives show that a large majority of US listed companies use the capital asset pricing model (CAPM) to estimate their cost of equity, and therefore require an estimate of their share's beta (e.g. Graham & Harvey, 2001). The CAPM is also the most common method used in the UK, especially amongst larger listed companies (e.g. McLaney, Pointon, Thomas & Tucker, 2004). It is the primary method recommended in all the leading finance textbooks for estimating a company's cost of equity. In addition, asset and portfolio betas are fundamental in textbook explanations of portfolio theory. Beta is the benchmark measure of systematic risk, and estimation of beta is necessary to implement any model of expected returns which includes market risk as one of the risk factors. All of the well-known multifactor alternatives to the simple CAPM incorporate market risk. Thus, applications of beta remain very widespread in the practice of finance and in academic research, and the estimation of beta is an issue of direct relevance to many finance practitioners, students and researchers.

Both the method of beta estimation and the interval over which returns are measured affect beta estimates. The current paper presents new and straightforward evidence on these questions from a sample of UK shares of all sizes. Most beta estimates used in practice are obtained from ordinary least squares (OLS) regression of the returns

on the share against the monthly, weekly or daily returns on a market index.¹ A common problem in the modelling of financial data is variation over time in the volatility of the OLS error term, especially in higher-frequency data. This problem of heteroscedasticity causes the OLS assumption that the error term is normally distributed with a constant variance to be violated. According to standard statistical theory, OLS estimation of the model's parameters remains unbiased and consistent, but is not efficient, in which case the standard error of the beta coefficient can be either overestimated or underestimated. However, Hamilton (2010) argues that failure to account for heteroscedasticity can also result in incorrect estimates of the regression model's *parameters*. He presents evidence that the estimated parameters can differ substantially between OLS and ARCH models, which supports his theoretical argument.

The current paper investigates estimates from the ARCH (autoregressive conditional heteroscedasticity) class of models.² An ARCH-type model can be used to estimate beta and its standard error whenever there is significant heteroscedasticity in the share returns. The paper compares OLS betas for shares with the corresponding betas from the most appropriate model in the ARCH category. The evidence presented helps to answer the question of how much difference it would make if, in cases with

¹ Some estimation services offer forecast betas with a simple adjustment for the known tendency of betas to regress to the mean. This adjustment gives a two-thirds weight to the estimated beta and a one-third weight to the mean, which is assumed to equal one. London Business School's Risk Measurement Service calculates beta with a Bayesian adjustment.

² Other alternatives to OLS include Bayesian approaches to estimation, regression methods that adjust for infrequent trading, and the conditional methods discussed in note 3.

heteroscedastic errors, a more correct model than OLS, from the statistical point of view, were chosen to estimate beta. The evidence in most previous papers on alternative approaches is from portfolios. Only two papers offer explicit comparisons of betas for shares (Lie, Brooks & Faff, 2000; Brooks, Faff & McKenzie, 2000). Both these papers use small samples of shares, and both calculate alternative betas from the GARCH(1,1) specification only, rather than from the most appropriate ARCH-type model. Our evidence makes clearer what the difference is to beta from using the most correct model on the menu of ARCH extensions to OLS. It turns out that the differences are mostly fairly small. However, the difference in beta using daily data is at least 0.05 for 45% of the sample, and at least 0.1 for 17%. It is not apparent in previous research that these more substantial differences can arise. Perhaps a surprising result is that the difference between OLS and ARCH betas, positive or negative, is largest for the largest, most frequently traded shares. So the estimation technique tends to matter most in cases in which the valuation or risk of large chunks of equity are being estimated.

A difference in beta of 0.05 might look small but it will have a material impact on the valuation of many shares. For example, suppose that the risk-free rate is 3% per year, the equity risk premium is 5% per year, the OLS beta for a given share is 1.00 and the ARCH beta is 0.95. Then, according to the capital asset pricing model, the cost of equity is 8.00% per year using the OLS beta and 7.75% using the ARCH beta. If the year-end dividend is 10p per share and the long-term growth rate of the dividend is 5% per year, then the share price using the dividend-growth model is $10p/(8.00\% - 5.00\%) = 333.3p$ with the OLS beta, and $10p/(7.75\% - 5.00\%) = 363.6p$ with the ARCH beta. The difference in value is substantial: 9.1%.

Another example in which a small change in beta can make a large difference to an outcome is in estimation of the benchmark return for a managed fund. Suppose that a fund j is assessed by comparing its return over a given year t with the expected return from a model that includes the term $\beta_j(R_{M,t} - R_{F,t})$, where β_j is the estimated beta of the fund, and $R_{M,t} - R_{F,t}$ is the return on the stock market minus the risk-free return during year t . The impact of β_j depends on the size of the excess return on the market, but values of at least plus-or-minus 10% are common. With an excess return of 10%, a difference in β_j of 0.05 implies a difference in the benchmark return of 0.5%, and therefore a difference in the performance of the fund in relation to its benchmark of 0.5%. This is a substantial change, in the context of a business in which beating the benchmark by 0.5% or more is seen as quite an achievement.

A problem with betas calculated from daily returns is downward bias in the estimates for less liquid shares. We find that models in the ARCH category do not alleviate the downward bias in daily betas. In fact the mean and median betas are slightly higher using OLS than they are using an ARCH model. We also find that the use of ARCH-type models is not usually warranted when working with monthly returns.

The second question we examine is the impact of the return interval on betas. It is known from previous research that thin or infrequent trading in shares causes severe downward bias in OLS beta estimates, and that the bias is worse in higher-frequency data. While we confirm the impact of thin trading, we find in addition that the difference between betas estimated from daily and monthly data is strongly related to daily volatility. The most volatile shares see very large increases in beta as one moves from daily to monthly betas. In fact there is a stronger link between differences in daily and

monthly betas and daily volatility than there is between the beta differences and frequency of trading. This is a new finding. It suggests that the idiosyncratic portion of daily returns is positively related to their volatility. This appears to be a reason for low daily betas, one that is both separate from thin trading, and at least as significant in terms of its impact on daily betas. We argue that our findings uncover a new explanation for low beta estimates from daily data.

The next section reviews previous research on interval effects and the use of ARCH models. Section three presents our data and method, Section four presents the results and Section five concludes.

2. Previous research

2.1 OLS and ARCH betas

There is a large literature that explores applications of ARCH models to financial market data. Here we concentrate on research involving comparisons of beta estimates from OLS and ARCH-type models. Many papers have reported that the volatility of stock returns, and the volatility of the error term in the OLS model, do vary over time, causing heteroscedasticity. Such evidence provides a clear motivation for exploring the use of ARCH methodology.

Most papers comparing OLS and alternative betas test the accuracy of betas from the different estimation methods, including OLS, ARCH, the Schwert-Seguin (1990) regression model, and the Kalman filter technique.³ The method of testing is to compare

³ The Schwert-Seguin model is a way of allowing for time-varying variance in index returns. The Kalman filter is a recursive algorithm. It is a fundamentally different method of estimation, in that

the forecast returns with the subsequent observed returns, where the forecast returns result from insertion of betas from different estimation models into the market model. Most studies present evidence from portfolios. Faff, Hillier & Hillier (2000) present beta estimates and the results of forecasting tests using ten estimation models. They use daily data for 32 UK industry sectors. The OLS betas and the estimation-period means of three time-varying ARCH betas are highly correlated across the sectors; the difference between the highest and lowest of the OLS and ARCH betas for each sector is less than 0.10 for all but four sectors. The three ARCH betas are much more highly correlated with each other than are the ARCH and Schwert-Seguin betas or the ARCH and Kalman filter betas. The method with the most accurate forecasts is Kalman filter; the ARCH models do not consistently perform better than OLS. Reyes (1999), using monthly returns, also reports that Schwert-Seguin betas for portfolios are substantially different from OLS betas. Braun, Nelson & Sunier (1995) compare forecasts using time-varying US industry betas from an EGARCH model with rolling OLS betas estimated over the previous 60 months. The mean errors in forecast returns are smaller from EGARCH for all industries, though not much smaller. Mergner & Bulla (2008) expand the range of estimation techniques tested to include stochastic-volatility and Markov-switching models.

Fraser, Hamelink, Hoesli & McGregor (2004) compare OLS and ARCH-type models using the Fama-MacBeth approach of estimating portfolio betas and then

it is not based on classical regression. A series of time-varying betas can be estimated via Kalman filter by assuming a particular process for beta. For example, if beta is assumed to follow a random walk, the process is $\beta_{i,t} = \beta_{i,t-1} + \varepsilon_{i,t}$.

estimating the cross-sectional relation between returns and betas in a subsequent period. They use monthly data for seven UK industry sectors. None of the estimation methods produces betas which have a significant relation with returns.

Overall, the results suggest that neither ARCH, Schwert-Seguin nor Kalman betas provide much more accurate forecast returns than do OLS betas, though Kalman betas may have the edge. There is, though, a problem with the tests: they assume implicitly that there is a linear relation between returns and betas, when the latter are accurately estimated. If the relationship between returns and accurate betas is weak or non-existent, as may be the case, then a more accurate estimate of beta will not necessarily result in a more accurate forecast of returns.

The two papers closest to ours are Lie *et al.* (2000) and Brooks *et al.* (2000). Lie *et al.* examine alternative beta estimates for 15 Australian banking shares, using daily data. The OLS betas and estimation-period means of the GARCH(1,1) betas are highly correlated across the 15 shares, with a correlation coefficient of 0.99. Fourteen of the 15 differences are smaller than 0.10. The cross-sectional correlation between the OLS and Kalman betas is lower at 0.93. The GARCH and Kalman betas are not highly correlated over time; the average of the correlation coefficients for each of the shares is only 0.13. Brooks *et al.* (2000) present OLS betas and means of GARCH(1,1) betas for 11 estimation periods for 18 US banking shares. They also report a correlation between the OLS and mean GARCH (1,1) betas across the shares of 0.99, estimated over the 11 periods combined. Both these studies have small samples and examine the GARCH(1,1) specification only. In addition, neither paper compares for each share a single OLS beta with a single beta from each of the other estimation methods, using the same estimation

period for all the shares and methods. Thus, the difference to beta that alternative methods make compared with OLS is not transparent.

2.2 *Measurement interval and thin trading*

The interval over which returns are measured makes a large difference to beta estimates. In principle daily returns are preferable. They contain more information about the variation in the relevant share and market, and they provide a much larger number of observations for a given estimation period, compared with weekly or monthly returns. This results in much greater precision of beta estimates (Daves, Ehrhardt & Kunkel 2000). The number of observations can be increased by lengthening the estimation period, but it causes many beta estimates to change, because betas tend not to be stable over time. Commercial estimates are based on monthly or weekly returns, whereas many academic studies, especially event studies, use betas estimated from daily returns. The use of daily betas in practical applications has also been prominently advocated. For example, Wright, Mason & Miles (2003) is an influential study on estimation of the cost of capital of UK regulated utilities. The authors examine the estimation of beta in detail, especially the question of the return interval, and they argue for daily data to estimate beta, at least for companies whose shares are frequently traded. It is now the norm to estimate the betas of UK utilities from daily returns; see, for example, the report by the consulting firm NERA (2009). Thus, estimation of beta from daily data is an issue of practical as well as academic relevance.

However, the use of daily data is compromised in practice by the effect of thin trading on beta estimates, especially in non-US markets. Thin trading in a given share

causes the share price to change less frequently than the value of the market index. The nonsynchronous changes in the price and the index result in downward bias in an OLS beta estimate. Many shares in markets outside the USA trade less than once a day, and the price changes much less frequently than once a day. This causes severe downward bias in beta estimates. Even in the USA, there is a strong positive relation between beta and frequency of trading.

It has been found that the problem of thin trading is worse the shorter the interval over which returns are measured. The bias is severe using daily returns to calculate beta, but mild or absent using monthly returns. For example, Davies, Unni, Draper & Paudyal (1999, p. 42) present the average beta across samples of 100 large UK companies by market capitalisation, 100 medium-sized and 100 small. The results for a five-year estimation period are:

Average beta using	Full sample	Large co's	Medium co's	Small co's
Daily returns	0.65	1.07	0.49	0.39
Monthly returns	1.01	1.08	1.04	0.88

Techniques proposed to correct for thin trading, including the well-known Dimson (1979) and Scholes-Williams (1979) methods, do not result in much increase in the mean betas of small shares. Further discussion of interval effects and thin trading is contained in Armitage (2005, section 13.3). It is not known to date whether ARCH estimation alleviates the bias from thin trading.

3. Data and method for ARCH-type estimates

We extract from Datastream daily share prices, adjusted for capital changes, of UK-registered companies quoted on the London Stock Exchange. Investment trusts and other investment companies are excluded. The sample period is the five years from 1 January 2002 to 31 December 2006. The sampling procedure is to select the largest company by market capitalisation as at 31 December 2006, then the fifth largest, the tenth and so on. If a share x does not have a complete list of prices for the full five years, the share one rank above x is chosen, then the share one rank below x , then two above then two below. If none of these five shares have five years of data, none is chosen and the process is begun again starting with the share five places below x . The selection stops with the share ranked 1,000. The initial sample is 198 shares.

The starting point of our analysis is estimation of beta from OLS regression:

$$R_{i,t}^{share} = \alpha_i^{OLS} + \beta_i^{OLS} \cdot R_t^{index} + \varepsilon_t \quad (1)$$

where $R_{i,t}^{share}$ is the return on share i on day or month t , R_t^{index} is the return on the FTSE Allshare Index, and ε_t is the error term. We test for heteroscedasticity in the error term share by share, using the Lagrange multiplier test for the ARCH effects in Engle (1982). If heteroscedasticity exists, we start with a basic specification of GARCH(1,1) and test for any remaining ARCH effects. If necessary, we add more lags and estimate higher-order equations, but in 22% of cases the GARCH (S,Q) specification is not applicable, and we use a more parsimonious ARCH(S) variant. We also test for asymmetries in the error term and for in-mean effects by specifying, respectively, asymmetric ARCH and ARCH-in-mean models. If asymmetries or in-mean effects are present, the relevant

variant of the ARCH model is the one that provides the beta estimate used in the paper. The ARCH-type models are estimated in EViews via maximum likelihood.⁴

We select the model that fits the data best by comparison across standard OLS and alternative specifications of ARCH, to identify the most appropriate specification for the share in question. Selection from OLS and ARCH models can be done in a relatively straightforward manner, whereas it can not be done easily if other types of model, such as Kalman filter, are included in the list to select from.

The linear ARCH model we employ has the following form (Engle, 1982):

$$R_{i,t}^{share} = \alpha_i^{ARCH} + \beta_i^{ARCH} \cdot R_t^{index} + \xi_t \quad (2a)$$

⁴ We choose ARCH models to correct for heteroscedasticity rather than alternative methods such as Newey-West (1987) or White (1980), because these methods suffer from shortcomings regarding inference when heteroscedasticity is present. Hamilton (2010) demonstrates that inference about the conditional mean can be inappropriately influenced by high-variance episodes in time-series data, if the conditional variance has not been incorporated directly into the estimation of the mean (or the beta coefficient, in our case). High-variance episodes often occur in stock market data. Hamilton (2010) argues that even when the primary interest is in estimating the conditional mean, having a correct description of the conditional variance is still important. First, hypothesis tests about the mean in a model in which the variance is misspecified will be invalid. Second, by incorporating observed features of the heteroscedasticity into the estimation of the conditional mean, substantially more efficient estimates of the conditional mean can be obtained. Hamilton shows the magnitude of this problem and provides evidence from macroeconomic data on why both the Newey-West (1987) and White (1980) techniques do not fully correct for the above inference problems, which is more successfully done through ARCH methodology.

$$\xi_t = \mathcal{G}_t \sqrt{h_t} \quad (2b)$$

$$h_t = \gamma_0 + \sum_{s=1}^S \gamma_s \varepsilon_{t-s}^2 \quad (2c)$$

where $\mathcal{G}_t : \text{IID}(0,1)$ ($\mathcal{G}_t$ is identically and independently distributed with a mean of zero and a variance of one), $\xi_t : \text{IID}(0, \sigma_\xi^2)$, ε_t is the error term from the OLS regression (1), and $S > 0$, $\gamma_0 > 0$, $\gamma_s \geq 0$. If there is heteroscedasticity in the OLS model, the information contained in the conditional variance function h_t about the time-varying variance of the error term is used in estimating α_i^{ARCH} and β_i^{ARCH} . In a GARCH(S, Q) model we have:

$$h_t = \gamma_0 + \sum_{s=1}^S \gamma_s \varepsilon_{t-s}^2 + \sum_{q=1}^Q \phi_q h_{t-q} \quad (3)$$

where an autoregression of h_t is added in the conditional variance function (2c) (Bollerslev, 1986).⁵ We investigate estimation via several other GARCH-type models, including exponential GARCH (Nelson, 1991), threshold ARCH (Zakoian, 1994), GJR (Glosten, Jagannathan & Runkle, 1993) and in-mean models (Engle, Lilien & Robins,

⁵ A different approach to estimating a GARCH beta is to use GARCH to model the conditional variance of the index $h_{m,t}$ and of the share, $h_{i,t}$. The covariance is then $\rho_{im} \sqrt{h_{i,t} h_{m,t}}$; the correlation coefficient ρ_{im} is assumed to be constant in some studies (e.g. Lie *et al.*, 2000). We view this as an indirect method of estimation compared with the model in (2a) to (2c), which corrects directly for heteroscedasticity in the error term.

1987). However, with the exception of one estimate using EGARCH, these specifications do not produce better-fitting models than simpler alternatives.

The model chosen to estimate the beta of a given share is the one which tests show to be the most appropriate, given the data. For example, if the errors in GARCH(1,1) show no remaining heteroscedasticity, asymmetry or in-mean effects, then GARCH(1,1) is taken to be the best-fitting model. If there is heteroscedasticity remaining and a higher order of GARCH (S, Q) has to be determined, we do that by testing the combinations of S and Q lags in the sequence: first $(1,1)$, then $(2,1)$, $(1,2)$, $(2,2)$ etc. We compare all variants with parameters $S=1,2,3,4 \dots$ etc. and $Q=1,2,3,4 \dots$ etc., and we choose the best one from the point of view of statistical significance and the maximum likelihood criterion.

Once the best ARCH-type model has been chosen for a given share, we test for autocorrelation in the error term using the Ljung-Box Q-statistic. If autocorrelation is present we impose a relevant structure of autoregressive (AR) and/or moving average (MA) terms, in order to eliminate it. Eight of the shares exhibit one or more extreme daily returns, defined as a return with an absolute value greater than 30%. These shares are all very infrequently traded. When this situation occurs it causes persistent autocorrelation which cannot be eliminated by imposing any structure of AR and/or MA terms. We treat those extreme returns as outlying observations and deal with them by adding a relevant dummy variable, which allows us to eliminate the autocorrelation. This method is akin to a robust regression technique to reduce the impact of outliers.

4. Results

4.1 *Daily data*

Table 1 around here

Table 1 shows a breakdown of the best-fitting specifications of the OLS and ARCH-type models. It shows that the different values of the ARCH betas and their standard errors, compared with OLS, result in a different sample of shares with betas significantly different from zero. There are 28 shares (14% of the total) with an OLS beta not significantly different from zero at the 10% level, and no ARCH estimate is calculated for all but one of these, because there is no significant heteroscedasticity, or because none of the ARCH-type betas are significantly different from zero. The one remaining case demonstrates that there can be no relationship between a share and an index using OLS, when one exists using an ARCH model. Twenty-five further shares have a significant OLS beta but no ARCH beta. In all, 52 shares have no ARCH beta, 26% of the initial sample. For 21 shares the reason is that no significant heteroscedasticity was detected in the 1,304 daily returns, in which case no ARCH model could be specified and estimated. The remaining 31 shares have returns that display significant heteroscedasticity, but, despite this, no ARCH beta significantly different from zero could be calculated. 145 shares, 73% of the initial sample, have both an OLS and an ARCH beta that is statistically significant. Of these, 112 betas are obtained from the GARCH(1,1) model, 32 from the ARCH(1) model and one from the asymmetric

EGARCH(1,1) model. The dominance of GARCH(1,1), with 77% of the 145 ARCH-type betas, is consistent with the popularity of this specification in previous research.

The absence of a significant beta estimate is related to lack of trading in the share. As a measure of trading frequency, we calculate the proportion of trading days in the sample period on which the price changed. For the 145 shares with both OLS and ARCH-type betas, the trading frequency is 70% of trading days, with a range between 96% and 12%. In contrast, the mean trading frequency is 34% for shares without a significant OLS beta and 45% for shares with an OLS beta but without an ARCH beta. The lack of significance of the OLS beta for 14% of the sample is a problem for beta estimation with daily data. The problem is similar for ARCH models, with 18% lacking a significant ARCH-type beta (of the 177 shares that exhibit significant heteroscedasticity).

Figure 1 and Table 2 around here

Figure 1 shows the differences between OLS and ARCH-type betas for the 145 shares for which the comparison can be made. The correlation between the two estimates is 0.98. The difference, ARCH less OLS, is negative for 64% of the sample. The mean of the absolute differences is 0.060 (median 0.041), with the largest absolute difference being 0.271. In percentage terms, the mean difference is 15.5% (12.0%) and the largest is 74.6%. Thus, use of ARCH instead of OLS often makes a modest difference, which is in line with the impression gained from the results for the small samples in Lie *et al.* (2000) and Brooks *et al.* (2000). But the change can be substantial: for 65 (45%) of the sample the difference exceeds 0.05, and for 25 (17%) it exceeds 0.10. This finding is not

apparent in previous research. In percentage terms, the difference exceeds 10% for more than half (57%) of the shares and it exceeds 30% for 17% of them.

Tables 2a and 2b divide the sample into categories given by the size of the absolute difference and percentage difference between the ARCH and OLS betas. For all the size categories, more ARCH betas are smaller than larger. This asymmetric result suggests that OLS tends to *overestimate* the beta coefficients compared with ARCH models, which take into account heteroscedasticity of the error term. This finding is not reported earlier in the literature.

The mean (median) OLS beta is 0.491 (0.336), compared with a mean ARCH beta of 0.469 (0.317). The OLS result is consistent with previous findings that mean daily betas for samples dominated by smaller-cap shares are well below one for non-US markets. The main reason is thought to be thin trading, although we offer an additional explanation in Section 4.3.⁶ We find that the downward bias is not alleviated by the ARCH method. There is a strong positive correlation of both OLS and ARCH betas with trading frequency; the correlation coefficient is 0.65 for OLS and 0.66 for ARCH. In addition, the difference ARCH minus OLS beta tends to be more negative for smaller shares. The mean difference in the smallest fifth of the sample is -0.039 compared with -0.005 for the largest fifth (not tabulated).

Figure 2 around here

⁶ In addition, the mean beta can differ substantially from one when a value-weighted index is used, such as the FTSE Allshare Index in the current study. The mean beta would be expected to be closer to one if an equally weighted index were used.

A noteworthy finding is the relationship between the absolute difference between the ARCH and OLS betas and trading frequency, shown in Figure 2. The shares are grouped in quintiles by size of the difference; shares with the largest positive difference are in the first quintile and shares with the largest negative difference are in the fifth. It can be seen that the biggest differences between ARCH and OLS, whether positive or negative, tend to be in the most frequently traded shares. This visual impression is confirmed by the following regression results (*t*-stats in brackets):

$$Abs_diff_i = -0.30 + 0.51 \cdot Trad_freq_i \quad (4)$$

(−2.45) (2.95) $R^2 = 0.15$

where Abs_diff_i is the absolute difference between the ARCH and OLS beta estimates for share i , and $Trad_freq_i$ is the trading frequency measured as proportion of days in the sample period on which the price changed. The slope coefficient is significant at the 1% level. Thus the decision about whether to use an OLS or ARCH-type estimate is likely to make most difference if the share is frequently traded. This is an important finding of our study.

As a further experiment and robustness check, we calculate betas from OLS with adjustment for autocorrelation only. 103 of the 145 shares in the final sample (71%) displayed autocorrelation. In 49 cases (34%) the autocorrelation was of the first order, so the inclusion of an AR(1) or MA(1) term was sufficient to eliminate it completely. In the remaining cases, a more complex AR and/or MA structure was needed to remove the autocorrelation. The estimates of beta from OLS adjusted for autocorrelation are very similar to the standard OLS betas and are not reported. The mean (median) of the absolute value of the difference is only 0.008 (0.004) with daily data. We conclude from

this that adjustment for heteroscedasticity of the error term makes much more difference to the estimated beta than does adjustment for autocorrelation.

4.2 *Monthly data*

When we switch to monthly returns, significant heteroscedasticity is detected for only nine shares. Thus it is usually not possible to estimate an ARCH-type beta with 60 monthly returns, a data specification often chosen for commercial estimates.⁷ 140 of the 145 shares have significant OLS betas from monthly returns. The monthly beta is higher than the daily beta for 129 of the 140 shares, 92% of the sample, as shown in Figure 3 below. The mean (median) difference is huge; it is 0.718 (0.618), with a range between 3.20 and -0.37 . The mean beta increases from 0.478 (0.317) using daily returns to 1.196 (1.039) using monthly returns. In percentage terms the mean difference is 287% (202%), with a range between 1,273% and -52% . The increase tends to be more positive for less frequently traded shares; the correlation between the difference in betas and trading frequency is -0.40 , or -0.77 if the difference is in per cent. Furthermore, the correlation between trading frequency and beta drops from 0.69 for daily data to 0.07 for monthly data. These results confirm that monthly returns greatly alleviate the downward bias caused by thin trading. The larger betas from monthly data for the UK are also noted in Davies *et al.* (1999).

⁷ McClain, Humphreys and Boscan (1996) find that statistically significant ARCH parameters can be obtained for most shares when there are at least 500 monthly returns. Most firms do not have an estimation period of this length, and if they do, the nature and risk of their business is likely to have changed substantially over such a lengthy period.

Sixty-three of the shares have monthly errors that exhibit autocorrelation. We find in unreported results that adjustment for autocorrelation makes very little difference to the OLS betas, as in the case of daily returns.

4.3 *Volatility and the difference between daily and monthly betas*

Although the estimates using monthly data alleviate the bias caused by thin trading, a lot of information about the variation in the share price is lost compared with the information in daily data. The effect of this loss of information on the relationship between share and market returns is uncertain *a priori*. The volatility of a share's returns is a proxy for the impact on returns of new information affecting the share price, and so volatility could help explain differences between daily and monthly betas. For example, suppose that the daily changes in share prices are mainly due to company-specific reasons, but longer term movements tend to be due to reasons affecting the stock market as a whole. Then it could be that, cross-sectionally, the changes from daily to monthly betas are positively related to the shares' volatilities.

Handa, Kothari & Wasley (1989) present a model in which the beta estimate is predicted to rise for betas above one, and fall for betas below one, as the return interval is lengthened. The model assumes that a given return interval consists of n periods, and that the expected one-period returns on share and market, and the expected one-period covariance and market variance, are all constant. The model's prediction arises because, as n is increased, the expected covariance of the share and the market measured over the interval does not change in proportion with the expected variance of the market. Handa *et al.* (1989) present supporting evidence, using US data. However, the construct, that an

interval consists of n identical periods, may not accurately reflect how the covariance between share and market returns changes with the length of the interval.

To investigate the difference between monthly and daily betas, we run the following cross-sectional regression. The dependent variable is the beta estimate from monthly returns minus the estimate from daily returns ($Diff_M-D_i$). The explanatory variables are the standard deviation of the daily returns ($Daily_volatility_i$), and our proxy for trading frequency ($Trad_freq_i$). The daily and monthly betas used in the cross-sectional regression are the ones that are best from a statistical and econometric point of view. The betas are either from the OLS model (where no heteroscedasticity is detected) or from an ARCH-class model (where heteroscedasticity is present), with correction for autocorrelation as appropriate. The results are as follows (t -stats in brackets):

$$Diff_M-D_i = 1.07 + 6.50 \cdot Daily_volatility_i - 1.03 \cdot Trad_freq_i \quad (5)$$

(8.19) (10.01) (-6.01) $R^2 = 0.51$

Both slope coefficients are significant at the 1% level. The negative sign on trading frequency indicates that less trading is associated with a larger increase in the monthly beta, as expected. The positive sign on intra-month daily volatility indicates that high daily volatility is associated with a *low* daily beta compared with the monthly beta. This suggests that, for more volatile shares, much of the short-term volatility is idiosyncratic and is not associated with changes in the market index.

We check whether the value of the beta from daily returns affects the results, as predicted by Handa *et al.* (1989). We include as an explanatory variable the daily beta multiplied by a dummy variable that equals 1 if the daily beta is one or more, and -1 if the daily beta is below one. The predicted sign of the coefficient is positive. The estimated coefficient is negative but not significantly different from zero. This suggests

that the processes by which daily and monthly returns arise in our sample are not adequately represented by the model in Handa et al (1989).

When the regression is run for the two explanatory variables separately, $R^2 = 0.15$ in the model with $Trad_freq_i$ in its own and $R^2 = 0.38$ in the model with $Daily_volatility_i$ on its own. The correlation between the two variables is only -0.05 . Our results imply that daily volatility explains more of the difference between the monthly and daily estimates of beta than does trading frequency. This is a new finding.

Figures 3a and 3b around here

Figure 3a shows the mean differences between monthly and daily beta for ten groups of shares ranked by size of the difference, together with the mean daily volatility. The figure illustrates the strong positive relation that exists between the difference in beta and volatility, especially for the many shares with a large positive difference between monthly and daily beta. Figure 3b compares the difference in beta with trading frequency. It is clear that this relationship is less strong than that between the difference in beta and daily volatility. Once the proportion of days on which the price changes falls below about 75%, there is no relation between the change in beta and trading frequency, though there are big differences in the changes in beta for these thinly traded shares.

We investigate further whether the strength of the relationship that exists between the $Diff_M-D_i$ and $Daily_volatility_i$ variables, captured by equation (5), differs across sub-samples of shares of various sizes. We rank the data according to the shares' market capitalisations, divide the sample into three equal sub-samples of large, medium and

small companies, and ran again regression (5) for each of the sub-samples. The estimates of the slope coefficient of the $Daily_volatility_i$ variable for the largest, medium and smallest samples are 1.54, 8.85 and 5.74; the corresponding t -statistics are 1.59, 9.21 and 5.53, and the R^2 values are 0.24, 0.68 and 0.42. So the relationship between $Diff_M-D_i$ and $Daily_volatility_i$ is most visible in the group of medium-sized shares. This is also the group with the largest mean daily variance, 0.065, compared with a mean variance of 0.048 for large shares and 0.052 for small shares.

The results imply that the idiosyncratic element of share returns forms a larger proportion of the daily returns than of the monthly returns, and that the idiosyncratic proportion is positively related to daily volatility. This effect is separate from the impact of thin-trading bias in explaining the differences between monthly and daily betas. The effect is substantial – it is more important in our sample than the impact of thin trading – and it is most pronounced among shares with high volatility of daily returns. Our findings imply that the low daily betas found in small- and medium-cap shares are not only a result of thin trading, but also arise because of a larger idiosyncratic component in daily returns than in monthly returns of these shares.

4.4 Accuracy of OLS and ARCH betas

Finally, as a check of which beta estimates are more accurate, we perform an out-of-sample comparison of forecast accuracy using the market model, in which we substitute beta estimates from OLS and ARCH models for all 145 shares obtained in the in-sample period of years 2002-06. The forecasting horizons are the next three, six, nine

and 12 months in the year 2007. Of course, an out-of-sample test is less likely to find a difference in forecast accuracy than an in-sample test.

We calculate the returns implied by the market model using the daily betas from OLS and ARCH models, and compare them with the actual market returns in the out-of-sample period. Forecast accuracy is assessed by the mean absolute error and the mean square error. The results are presented in Table 3, and they clearly show that the forecasting errors are smaller for ARCH betas, especially over horizons of six to twelve months, although the improvements on average are small. ARCH betas produce more accurate forecasts in 59% of cases, over a six-month horizon. Tests of this nature assume that there is a linear relation between return and beta correctly estimated, as noted in Section 2.1. To the extent that this assumption is true, the results in Table 3 indicate that ARCH-type betas are not only more correct statistically, but they tend to produce more accurate forecasts of returns.

5. Conclusion

We compare betas estimated from OLS with betas from the best-fitting model in the ARCH category, for a sample of 145 UK shares. The mean absolute difference in the betas from daily data is fairly small at 0.06, although it exceeds 0.10 for 17% of the sample. Even a difference of 0.05 can have a substantial impact on the estimate of a company's market value, or on the measurement of the performance of a managed fund. We also find that correction for heteroscedasticity in the error term, *via* an ARCH model, makes much more difference to beta estimates than does correction for autocorrelation. The absolute value of the difference in betas is positively related to frequency of trading:

choice of model matters more for more liquid shares. Unfortunately, use of an ARCH model does not alleviate the severe downward bias in the betas of illiquid shares using daily data.

We find that it is not always possible to estimate a share's beta using an ARCH model. The standard estimation period in practice is five years, and the reason it is not longer is that the underlying risk of the company is more likely to change as the length of the estimation period is lengthened. But with daily data and a five-year estimation period, no ARCH model is estimated for 26% of the sample, either because no heteroscedasticity is identified in the daily returns – so no ARCH model can be specified – or because the ARCH beta is not significantly different from zero. In these cases the simple OLS beta can be used. With monthly data and a five-year estimation period, ARCH effects rarely exist.

An ARCH-type model is of most potential value in beta estimation for large, frequently traded shares. Daily data can be used for such shares without causing bias, and we have found that the difference between OLS and ARCH betas tends to be greatest amongst the most frequently traded shares. Hence, the potential benefit from using ARCH is greatest for large-cap shares. Our tests of forecast accuracy confirm that ARCH betas are an improvement over OLS, in that they tend to produce more accurate forecasts of returns.

We also examine the difference between beta calculated from monthly returns and from daily returns. The difference is larger for less frequently traded shares, as others have reported. But we also find that the difference is positively related to the daily volatility of the share. In fact volatility explains more of the difference between monthly

and daily betas than does trading frequency. The returns of volatile shares are less sensitive to market returns than they appear to be from monthly betas, at least when the returns are measured over short intervals.

Overall, our results show that adjustment for heteroscedasticity in returns via an ARCH-type model makes a modest difference to the betas of individual shares calculated from daily data – though modest differences in beta can matter in practice. The adjustment should be made in theory, since an ARCH-type model provides a more accurate estimate of the beta's standard error when significant heteroscedasticity is present, as it is for a large majority of shares. In addition, Hamilton (2010) argues that OLS estimates of the parameters are incorrect in the presence of heteroscedasticity. If monthly data are being used, we find that it is rarely possible to use an ARCH model.

A bigger question from a practical perspective is the choice of interval over which returns are measured. The mean increase in monthly over daily beta is 0.72 in our sample, which is very large. The established explanation for this increase is that daily betas suffer from downward bias due to infrequent trading. But we find that daily volatility explains more of the cross-sectional differences in the betas than does trading frequency. When the sample is categorised by size into three sub-samples, the link between the differences and daily volatility is strongest in the middle sub-sample, which has the largest average daily volatility.

These new results deserve further study. Our interpretation is that daily volatility is a measure of the information that is lost when the share prices and their returns are aggregated from daily to monthly frequency. Company-specific (idiosyncratic) information accounts for a larger proportion of daily returns than of monthly returns. The

company-specific proportion of daily returns is positively related to their volatility, whereas this relationship is weaker or absent in monthly returns. So the daily betas of shares with high daily volatility are most biased downwards compared with the monthly betas. This effect is distinct from the effect of thin trading on daily betas.

Our results suggest that the true beta of a share is sensitive to the interval over which returns are measured, and that the impact of the measurement interval is not merely an unwanted by-product of thin trading. An implication is that low daily betas are less spurious than may previously have been thought. Unfortunately there is no obvious way of disentangling the effect of thin trading from the effect of the larger idiosyncratic component in daily returns than monthly returns.

References

Armitage, S. (2005) *The Cost of Capital: Intermediate Theory*, Cambridge University Press, Cambridge.

Bollerslev, T. (1986) Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics*, **31**, 307-327.

Braun, P.A., Nelson, D.B. and Sunier, A.M. (1995) Good news, bad news, volatility and betas, *Journal of Finance*, **50**, 1575-1603.

Brooks, R.D., Faff, R.W. and McKenzie, M.D. (2000) US banking sector risk in an era of regulatory change: a bivariate GARCH approach, *Review of Quantitative Finance and Accounting*, **14**, 17-43.

Daves, P.R., Ehrhardt, M.C. and Kunkel, R.A. (2000), Estimating systematic risk: the choice of return interval and estimation period, *Journal of Financial and Strategic Decisions*, **13**, 7-13.

Davies, R., Unni, S., Draper, P. and Paudyal, K. (1999) *The Cost of Equity Capital*, The Chartered Institute of Management Accountants, London.

Dimson, E. (1979) Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics*, **7**, 197-226.

Engle, R.F. (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica*, **50**, 987-1007.

Engle, R.F., Lilien, D.M. and Robins, R.P. (1987) Estimating time varying risk premia in the term structure: the ARCH-M model, *Econometrica*, **55**, 391-407.

Faff, R.W., Hillier, D. and Hillier, J. (2000) Time varying beta risk: an analysis of alternative modeling techniques, *Journal of Business Finance and Accounting*, **27**, 523-254.

Fraser, P., Hamelink, F., Hoesli, M. and MacGregor, B. (2004) Time-varying betas and the cross-sectional return-risk relation: evidence from the UK, *European Journal of Finance*, **10**, 255-276.

Glosten, L.R., Jagannathan, R. and Runkle, D. (1993) On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance*, **48**, 1779-1801.

Graham, J.R. and Harvey, C.R. (2001) The theory and practice of corporate finance: evidence from the field, *Journal of Financial Economics*, **60**, 187-243.

Hamilton, J.D. (2010) Macroeconomics and ARCH, in *Festschrift in Honor of Robert F. Engle*, (Eds) T. Bollerslev, J. Russel and M. Watson, Oxford University Press, Oxford, pp. 79-96.

Handa, P., Kothari, S.P. and Wasley, C. (1989) The relation between the return interval and betas: implications for the size effect, *Journal of Financial Economics*, **23**, 79-100.

Lie, F., Brooks, R. and Faff, R. (2000) Modelling the equity beta risk of Australian financial sector companies, *Australian Economic Papers*, September, 301-311.

McClain, K.T., Humphreys, H.B. and Boscan, A. (1996) Measuring risk in the mining sector with ARCH models with important observations on sample size, *Journal of Empirical Finance*, **3**, 369-391.

McLaney, E., Pointon, J., Thomas, M. and Tucker, J. (2004) Practitioners' perspectives on the UK cost of capital, *European Journal of Finance*, **10**, 123-138.

Mergner, S. and Bulla, J. (2008) Time-varying beta risk of pan-European industry portfolios: a comparison of alternative modeling techniques, *European Journal of Finance*, **14**, 771-802.

Nelson, D. (1991) Conditional heteroskedasticity in asset returns: a new approach, *Econometrica*, **59**, 347-370.

NERA Economic Consulting (2009) *Cost of Capital for Price Review 09: Final Report for Water UK*, 188 pp, available at <http://www.water.org.uk/home/policy/publications/archive/finance/revised-cost-of-capital-nera-eport/090206-nera-water-uk-wacc-final--2---2-.pdf>

Newey, W.K., and West, K.D. (1987) A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica*, **55**, 703-708.

Reyes, M.G. (1999) Size, time-varying beta and conditional heteroskedasticity in UK stock returns, *Review of Financial Economics*, **8**, 1-10.

Scholes, M.S. and Williams, J. (1977) Estimating beats from non-synchronous data, *Journal of Financial Economics*, **5**, 309-327.

Schwert, G.W. and Seguin, P.J. (1990) Heteroskedasticity in stock returns, *Journal of Finance*, **45**, 1129-1255.

White, H. (1980) A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica*, **48**, 817-838.

Wright, S., Mason, R. and Miles, D. (2003) *A Study into Certain Aspects of the Cost of Capital for Regulated Utilities in the UK*, Smithers & Co. Ltd, London, 144 pp, available at: <http://www.ofgem.gov.uk/Networks/Policy/Documents1/2198-jointregscoc.pdf>

Zakoian, J.-M. (1994) Threshold heteroskedastic models, *Journal of Economic Dynamics and Control*, **18**, 931-955.

Table 1
Breakdown of model specifications selected for betas from daily returns

Initial sample of shares	198
Less	
OLS beta not significantly different from zero	28
<i>Of which:</i>	
<i>No significant heteroscedasticity</i>	<i>6</i>
<i>No ARCH-type beta significant</i>	<i>21</i>
<i>ARCH-type beta significant</i>	<i>1</i>
OLS beta significant but no ARCH-type beta	25
<i>Of which:</i>	
<i>No significant heteroscedasticity</i>	<i>15</i>
<i>No ARCH-type beta significant</i>	<i>10</i>
Sample with significant OLS and ARCH beta	145
<i>In which ARCH model selected is:</i>	
<i>ARCH</i>	<i>32</i>
<i>GARCH</i>	<i>112</i>
<i>EGARCH</i>	<i>1</i>

Table 2a
Differences between ARCH and OLS betas

Absolute value of differences	% of the sample	Differences (ARCH less OLS)	% of the sample
> 0.25	2.07%	> +0.25 < -0.25	0.69% 1.38%
> 0.20	4.14%	> +0.20 < -0.20	0.69% 3.45%
> 0.15	8.28%	> +0.15 < -0.15	2.76% 5.52%
> 0.10	17.24%	> +0.10 < -0.10	5.52% 11.72%
> 0.05	44.83%	> +0.05 < -0.05	13.79% 31.04%

Table 2b
Percentage differences between ARCH and OLS betas

Absolute value of percentage differences	% of the sample	Percentage differences (ARCH less OLS)	% of the sample
> 50%	1.38%	> +50% < -50%	0.69% 0.69%
> 40%	5.52%	> +40% < -40%	2.76% 2.76%
> 30%	16.55%	> +30% < -30%	5.52% 11.03%
> 20%	28.28%	> +20% < -20%	8.28% 20.00%
> 10%	56.55%	> +10% < -10%	17.93% 38.62%
> 5%	75.86%	> +5% < -5%	24.83% 51.03%

Table 3
Accuracy of forecast returns from the market model over four out-of-sample forecasting horizons

Measure	Model	Forecasting horizon			
		3 months	6 months	9 months	12 months
Mean absolute error	OLS	10.555%	17.414%	20.028%	26.418%
	ARCH	10.522%	17.323%	19.979%	26.374%
Mean square error	OLS	2.024%	5.089%	6.547%	11.233%
	ARCH	2.015%	5.064%	6.524%	11.191%
	Prop'n ARCH < OLS	0.497 <i>0.00</i>	0.586 <i>1.93</i>	0.522 <i>0.43</i>	0.571 <i>1.56</i>
N		145	142	138	133

Note

The table shows summary statistics for differences between observed returns and returns predicted from the market model using α and β parameters estimated either via OLS or via the best-fitting ARCH-type model. The out-of-sample periods start on 1 January 2007, so for the three-month forecasting horizon, for example, the forecast error for a share i is given by $R_{i,t} - (\alpha_i + \beta_i R_{M,t})$, where $R_{i,t}$ and $R_{M,t}$ are the return on the share and on the FT-Allshare Index, respectively, during the first three months of 2007, and the α_i and β_i parameters are estimated using data from 2002-06. 'Prop'n ARCH < OLS' is the proportion of observations in which the prediction error from ARCH parameters is smaller in absolute size than the error from OLS parameters; the number below in italics is the z -statistic in a sign test. The sample size diminishes as the forecasting horizon lengthens because of company delistings.

Figure 1

Note

The figure shows the beta from an ARCH-type model less the OLS beta for 145 shares for which both betas are significantly different from zero. The estimates are based on daily returns during 1 January 2002 to 31 December 2006.

Figure 2

Note

Shares are grouped in quintiles by size of difference ARCH beta less OLS beta, most positive (quintile 1) to most negative (quintile 5). The bars show the mean trading frequencies for the shares in each quintile, measured by the percentage of days in the sample on which the price changed (left-hand axis). The diamonds show the means of the absolute values of the differences between ARCH and OLS betas (right-hand axis).

Figure 3a

Figure 3b

Note

The bars in both graphs show, for each group of shares, the mean of the monthly beta less the daily beta across the shares in the group (left-hand scale). The shares are ranked by the size of the difference in their betas, and the ten groups are formed from the ranked shares. The diamonds in Fig. 3a show the mean for each group of the variance of the daily returns. The diamonds in Fig. 3b show the mean for each group of the trading intensities, proxied by the proportion of trading days in the sample on which the price of the relevant share changes.