Robust Estimation of the Optimal Hedge Ratio

Richard D.F. Harris and Jian Shen

School of Business and Economics
University of Exeter

September 2002

Abstract

When using derivative instruments such as futures in order to hedge a portfolio of risky assets, the primary objective is to estimate the optimal hedge ratio (OHR). When agents have mean-variance utility and the futures price follows a martingale, the OHR is equivalent to the minimum variance hedge ratio, which can be estimated by regressing the spot market return on the futures market return using ordinary least squares. In order to accommodate time-varying volatility in asset returns, estimators based on rolling windows, GARCH or EWMA models are commonly employed. However, all of these approaches are based on the sample variance and covariance estimators of returns, which while consistent irrespective of the underlying distribution of the data, are not in general efficient. In particular, when the distribution of the data is leptokurtic, as is commonly found for short horizon asset returns, these estimators will attach too much weight to extreme observations. This paper proposes an alternative to the standard approach to the estimation of the OHR that is robust to the leptokurtosis of returns. We use the robust OHR to construct a dynamic hedging strategy for daily returns on the FTSE100 index using index futures. We estimate the robust OHR using both the rolling window approach and the EWMA approach, and compare our results to those based on the standard rolling window and EWMA estimators. It is shown that the robust OHR yields a hedged portfolio variance that is marginally lower than that based on the standard estimator. Moreover, the variance of the robust OHR is as much as 70% lower than the variance of the standard OHR, substantially reducing the transaction costs that are associated with dynamic hedging strategies.

KEYWORDS: Optimal hedge ratio; Hedging; Robust estimation; Futures; FTSE100 stock index.

Address for Correspondence: Richard D.F. Harris, School of Business and Economics, University of Exeter, Exeter EX4 4PU, UK. Email: R.D.F.Harris@exeter.ac.uk.

1. Introduction

The rapid expansion of derivatives markets over the last twenty five years has led to a corresponding increase in interest in the theory and practice of hedging. When using derivative instruments such as futures in order to hedge a portfolio of risky assets, the primary objective is to estimate the size of the short position that must be held in the futures market, as a proportion of the long position held in the spot market, that maximises the agent s expected utility, defined over the risk and expected return of the hedged portfolio. This is the problem of estimating the optimal hedge ratio (OHR).

In the mean-variance framework, risk is measured by the standard deviation of the hedged portfolio. Varying the hedge ratio traces out a feasible set for the hedged portfolio in expected return—standard deviation space. The OHR is the hedge ratio that equates the agent s marginal rate of substitution between the expected return and the standard deviation of the hedged portfolio with the slope of this feasible set (see Cecchetti, Cumby and Figlewski, 1988). If the futures price follows a martingale, the expected futures return is zero and so the OHR is simply that which minimises the variance of the hedged portfolio.[1] The minimum variance hedge ratio can be estimated by regressing the spot market return on the futures market return using ordinary least squares.[2] Almost all applications of optimal hedging use the criterion of minimum variance in order to estimate the OHR.

A shortcoming of the conventional OLS approach is that it ignores the fact that the conditional distribution of most financial asset returns tends to vary over time. When the

---

[1] The empirical evidence on whether futures prices follow a martingale process is mixed. For some futures series, the martingale hypothesis cannot be rejected. However, for others, there is evidence of mean reversion. See, for instance, Cargill and Rausser (1975), Taylor (1985), Doukas and Raman (1987) and Goldenberg (1988).

[2] Early studies (such as Ederington, 1979) use price levels in the OLS regression instead of price changes, but if the spot and futures prices are non-stationary and not cointegrated then the estimated OHR will be spurious. Other studies use the change in the price level instead of the return, but this is likely to induce heteroscedasticity in the error term of the OLS regression. See Myers and Thompson (1989) for a discussion of the merits of these different specifications.

conditional distribution of spot and futures returns is predictable, a more efficient estimate of the OHR can be obtained by conditioning on recent information. The most commonly adopted solution is to use a rolling window estimator of the variance-covariance matrix, in which the variance and covariance of spot and futures returns are estimated using an equally weighted moving average of past squared returns and their cross-products. However, rolling window estimators adopt an all-or-nothing approach in which observations have equal weight in the variance-covariance matrix estimator until they reach some arbitrarily defined age, after which they have zero weight. More recently, dynamic hedging strategies based on the GARCH class of models (Engle, 1982; Bollerslev, 1986, Bollerslev, Engle and Wooldridge, 1988; Bollerslev, 1990) have been proposed. A special case of the GARCH model is the exponentially weighted moving average (EWMA) estimator, which has been widely used for estimating conditional variance-covariance matrices for the purpose of calculating Value at Risk. The EWMA estimator imposes the restriction that the variance-covariance matrix is integrated, and is therefore equivalent to an integrated GARCH or IGARCH model. Both GARCH and EWMA models have been widely used for estimation of the optimal hedge ratio.

All of the approaches to estimating the OHR described above are based on the sample variance and covariance estimators of returns. These are consistent estimators of the population variance and covariance, irrespective of the underlying distribution of the data, but they are not in general efficient. In particular, when the distribution of the data is fat-tailed, the sample variance and covariance — and hence also the rolling window, GARCH and EWMA estimators that are based on these — will attach too much weight to extreme observations (see Nelson and Foster, 1996). There is now considerable evidence that short horizon financial asset returns are not normally distributed, even conditionally (see, for instance, Baillie and DeGennaro, 1990; Bollerslev, Chou and Kroner, 1992). As a result, hedging strategies based on the standard estimators of the OHR will yield hedged portfolios that will not generally have minimum variance because the estimated OHR will itself be excessively volatile. Moreover, the excessive volatility of the estimated OHR will increase the transaction costs that are incurred when rebalancing the hedged portfolio in response to changes in the OHR, and which serve to reduce the hedged portfolio s expected return.

This paper proposes an alternative approach to the estimation of the OHR that is robust to the leptokurtosis of the distribution of returns. The robust OHR uses the robust conditional variance estimator of Guermat and Harris (2002), which is based on the maximum likelihood estimator of the variance of the power exponential (PE) distribution. The PE distribution nests the normal distribution as a special case, but also nests other distributions that are fat-tailed. Estimators of the OHR based on these leptokurtic distributions can be expected to be less sensitive to the extreme observations that are commonly found in short horizon asset returns. In order to compute a robust estimate of the conditional covariance between spot and futures returns — which is required for the computation of the OHR — we employ an identity that relates the covariance between two random variables to the variance of their sum and the variance of their difference. Like the standard EWMA estimator, the estimator that we propose can be easily implemented in a spreadsheet package such as Excel™.

We use the robust OHR to construct a dynamic hedging strategy for daily returns on the FTSE100 index, using index futures contracts. We estimate the robust conditional OHR using both the rolling window approach and the EWMA approach, and compare our results to those based on the standard rolling window and EWMA estimators. It is shown that the robust estimator yields a hedged portfolio variance that is marginally lower than that based on the standard estimator. Moreover, the variance of the robust OHR is as much as 70% lower than the variance of the standard OHR, substantially reducing the transaction costs that are associated with dynamic hedging strategies.

The rest of the paper is organised as follows. The following section reviews the theoretical framework for estimating the OHR. Section 3 introduces the robust estimator of the OHR. Section 4 describes the data that we use in the empirical evaluation and outlines the empirical methodology. Section 5 presents the empirical results. Section 6 concludes.

2. Estimation of the Optimal Hedge Ratio

Hedging using futures involves taking a position in the futures market that is opposite to the position held in the spot market. For a long position in the spot market, the return of a hedged portfolio is given by

$$x_t = s_t - hf_t \tag{1}$$

where s_t is the return in the spot market at time *t*, f_t is the return in the futures market at time *t* and *h* is the hedge ratio. The optimal hedge ratio (OHR) is the value of *h* that maximises the investor s expected utility, defined over the expected return and risk of the hedged portfolio. In the mean-variance framework, risk is defined by the variance of the return of the hedged portfolio, which is given by

$$\mathrm{var}(x_t) = \mathrm{var}(s_t - hf_t) \tag{2}$$

When the futures price follows a martingale, the expected futures return is zero and so the futures position will not affect the expected return of the portfolio. Then, the OHR is simply the value of *h* that minimises (2), which is given by

$$\frac{\partial\,\mathrm{var}(x_t)}{\partial h} = 2h\sigma_f^2 - 2\sigma_{sf} = 0 \tag{3}$$

where σ_f^2 is the variance of the futures return and σ_{sf} is the covariance between the spot return and the futures return. This is solved to yield the conventional OHR,

$$h = \frac{\sigma_{sf}}{\sigma_f^2} \tag{4}$$

The OHR given by (4) can be estimated by regressing the spot return on the futures return using OLS. However, a well-established feature of many asset returns, including both financial assets and physical commodities, is that their conditional distribution is time-varying. In particular, short horizon asset returns tend to display significant volatility clustering. When the conditional distribution of spot and futures returns is predictable, a more efficient estimate of the OHR can be obtained by conditioning on recent information. Kroner and Sultan (1993) show that when the variance-covariance matrix of spot and futures returns is time-varying, the OHR is equal to

$$h_t = \frac{\sigma_{sf,t}}{\sigma^2_{f,t}} \tag{5}$$

where $\sigma_{sf,t}$ is the covariance of spot and futures returns at time t and $\sigma^2_{f,t}$ is the variance of futures returns at time t, conditioning on the time $t-1$ information set. To allow for this time-variation in the variance-covariance matrix of returns, a number of approaches have been adopted. The simplest, and most commonly used in practice, is a rolling window approach, in which the variance and covariance of spot and futures returns are estimated using an equally weighted moving average of past squared returns and their cross-products.[3]

$$\sigma^2_{s,t} = \frac{1}{m}\sum_{i=1}^{m} s^2_{t-i} \tag{6}$$

$$\sigma^2_{f,t} = \frac{1}{m}\sum_{i=1}^{m} f^2_{t-i} \tag{7}$$

$$\sigma_{sf,t} = \frac{1}{m}\sum_{i=1}^{m} s_{t-i} f_{t-i} \tag{8}$$

[3] We assume that the mean return is zero, which is a common assumption when dealing with daily returns, as we do in the empirical evaluation below.

where m is the window length typically set to a value between six months and five years. While the rolling window approach allows the OHR to respond to changes in the variance-covariance matrix of returns, they adopt an all-or-nothing approach in which observations have equal weight in the variance-covariance matrix estimator until they reach some arbitrarily defined age (the value of m), after which they have zero weight. This leads to ghost features in the estimated variance-covariance matrix (see Alexander and Leigh, 1997).

More recently, dynamic hedging strategies based on the GARCH class of models (Engle, 1982; Bollerslev, 1986; Bollerslev, Engle and Wooldridge, 1988; Bollerslev, 1990) have been proposed. These models typically employ a multivariate GARCH model to capture the dynamic evolution of the variance-covariance matrix and construct an estimate of the OHR using the conditional covariance of spot and futures returns and the conditional variance of futures returns. There are a number of different specifications of the multivariate GARCH model, each of which imposes different parameter restrictions to make estimation of the model feasible. With an assumption about the conditional distribution of returns, the parameters of the multivariate GARCH model can be estimated using non-linear maximum likelihood. A problem with multivariate GARCH models, however, is that the likelihood function is not globally concave, and so its maximisation is often extremely problematic, relying crucially on arbitrarily chosen starting values of the parameters. Nevertheless, dynamic hedging strategies based on the multivariate GARCH model have been employed using foreign currency futures (Kroner and Sultan, 1993; Lin, Najand and Yung, 1994; Brooks and Chong, 2001), commodity futures (Baillie and Myers, 1991), interest rate futures (Gagnon and Lypny, 1995) and stock index futures (Park and Switzer, 1995; Tong, 1996).

A related approach to estimating the conditional variance-covariance matrix of returns is the exponentially weighted moving average, or EWMA, estimator, given by

$$\sigma_{s,t}^2 = \lambda_1 \sigma_{s,t-1}^2 + (1-\lambda_1) s_{t-1}^2 \tag{9}$$

$$\sigma_{f,t}^2 = \lambda_2 \sigma_{f,t-1}^2 + (1-\lambda_2) f_{t-1}^2 \tag{10}$$

$$\sigma_{sf,t} = \lambda_3 \sigma_{sf,t-1} + (1-\lambda_3) s_{t-1} f_{t-1} \tag{11}$$

where λ_1, λ_2 and λ_3 are the decay factors for the spot and futures variances, and their covariance, respectively. The EWMA estimator is a special case of the multivariate diagonal VECH GARCH(1,1) model of Bollerslev, Engle and Wooldridge (1988). In particular, it imposes the restriction that the variance-covariance matrix is integrated, and consequently the EWMA model is also known as an integrated GARCH, or IGARCH, model. The EWMA model is widely used for estimating conditional variance-covariance matrices for the purpose of calculating Value at Risk following the introduction of Riskmetrics (see J.P. Morgan, 1996).

The EWMA model is characterised by the fact that shocks to the conditional variance-covariance matrix are permanent. Consequently, long run forecasts based on the EWMA model of the variance-covariance matrix do not converge to the unconditional variance-covariance matrix. While this means that the EWMA model clearly doesn t capture the true long run dynamics of most asset returns, it is often found to generate short run forecasts of the variance-covariance matrix that are as good as those of more sophisticated volatility models (see, for instance, Alexander and Leigh, 1997; Boudoukh, Richardson and Whitelaw, 1997). Combined with its ease of implementation — it can be straightforwardly implemented using Excel™ — this has made it extremely popular among practitioners. In the context of hedging, Brooks and Chong (2001) show that estimates of the OHR computed using EWMA estimates of the variance-covariance matrix of returns are superior to those generated by the GARCH model, using several different multivariate GARCH specifications.

3. Robust Estimation of the Optimal Hedge Ratio

All of the approaches to estimating the OHR described above are based on the sample variance and covariance estimators of returns. These are consistent estimators of the population variance and covariance, irrespective of the underlying distribution of the data, but they are not in general efficient, except when the underlying distribution is normal (see Spanos, 1999; Bondesson, 1975). In particular, when the distribution of the data is fat-tailed, the sample variance and covariance — and hence also the GARCH and EWMA estimators that are based on these — will attach too much weight to extreme observations. There is now considerable evidence that short horizon financial asset returns are not normally distributed, even conditionally (see, for instance, Baillie and DeGennaro, 1990; Bollerslev, Chou and Kroner, 1992). As a result, hedging strategies based on the standard estimators of the OHR will yield hedged portfolios that will not generally have minimum variance because the estimated OHR will itself be excessively volatile. Moreover, the excessive volatility of the estimated OHR will increase the transaction costs that are incurred in rebalancing the hedged portfolio in response to changes in the OHR

This paper proposes an alternative approach to the estimation of the OHR that is robust to the leptokurtosis of the distribution of returns. The robust OHR uses the conditional variance estimator of Guermat and Harris (2002), which is based on the maximum likelihood estimator of the variance of the power exponential (PE) distribution, whose density is given by

$$f(z,\sigma,k) = \frac{k}{\theta 2^{\frac{k+1}{2}}\Gamma(1/k)\sigma} e^{-\frac{1}{2}\left|\frac{z}{\theta\sigma}\right|^{k}} \tag{12}$$

where

$$\theta = \left\{\frac{2^{-2/k}\Gamma(1/k)}{\Gamma(3/k)}\right\}^{1/2} \tag{13}$$

$\Gamma(.)$ is the gamma function and k is a parameter that controls the kurtosis of the distribution. The PE distribution nests the normal distribution as a special case (when $k = 2$), but also nests other distributions that are fat-tailed. For instance, when $k = 1$, the PE distribution reduces to the Laplace distribution, which is commonly used in the context of robust estimation. The maximum likelihood estimator of the unconditional variance of members of the PE distribution, defined by different values of k, is given by

$$\hat{\sigma}^k = g(k)\frac{1}{T}\sum_{t=1}^{T}|z_t|^k \tag{14}$$

where

$$g(k) = k\left[\frac{\Gamma(3/k)}{\Gamma(1/k)}\right]^{k/2} \tag{15}$$

The power EWMA estimator of the variance derived by Guermat and Harris (2002) is given by

$$\sigma_t^k = \lambda\sigma_{t-1}^k + (1-\lambda)g(k)|z_{t-1}|^k \tag{16}$$

When $k = 2$, the maximum likelihood estimator of the variance of the PE distribution given by (12) reduces to the sample variance, and the power EWMA estimator given by (16) coincides with the standard EWMA variance estimator given by (9) or (10). However, when $k < 2$, these estimators will be less sensitive to extreme observations than the standard estimators and so could be expected to provide more efficient estimates of the variance when the distribution of returns is leptokurtic.

Computation of the robust OHR requires a robust estimate of the covariance of spot and futures returns as well as their variance. However, derivation of a covariance estimator based on the PE distribution is not straightforward, since a multivariate distribution that is consistent with the marginal PE distribution is not well defined (see Kano, 1994). In order to derive a robust covariance estimator, we instead use a simple identity that relates the covariance between any two square-integrable random variables to the variance of their sum and the variance of their difference. This approach is commonly used in the robust estimation literature when there exists a robust variance estimator, but no corresponding robust covariance estimator (see Huber, 1982). For the covariance between spot and futures returns, we have

$$\operatorname{cov}(s_t, f_t) = \frac{1}{4}\left\{\operatorname{var}(s_t + f_t) - \operatorname{var}(s_t - f_t)\right\} \tag{17}$$

By using the robust estimator to estimate the variance of the sum and difference of spot and futures returns, we can therefore indirectly obtain a robust estimator of the covariance of spot and futures returns. Combined with the robust estimate of the variance of futures returns, we can then compute the robust OHR using the formula given by (5). When $k = 2$, it is straightforward to show that the robust OHR will reduce to the standard OHR, for both the rolling window and EWMA estimators. However, when $k < 2$, the robust OHR will be less sensitive to extreme return observations. Like the standard OHR estimator, the robust OHR estimator can be easily implemented in a spreadsheet package such as Excel$^{\mathrm{TM}}$.

4. Empirical Evaluation

In order to illustrate the procedure outlined in the preceding section, we construct a dynamic hedging strategy for returns on the FTSE100 index, using both the rolling window and EWMA estimators to compute the OHR. We obtained daily closing prices for the FTSE100 index from Datastream and for the FTSE100 index futures contracts from LIFFE, for the period 04.05.84 to 03.05.02. At any one time, there are four futures contracts outstanding. On each day, we use the nearest contract to delivery, but rollover to the next nearest contract on the first day of the delivery month in order to avoid thin trading and expiration effects. Continuously compounded spot and futures returns were computed using the following formulae

$$s_t = \ln\left\{1 + \frac{S_t - S_{t-1}}{S_{t-1}}\right\} \tag{18}$$

$$f_t = \ln\left\{1 + \frac{F_t - F_{t-1}}{S_{t-1}}\right\} \tag{19}$$

where S_t and F_t are the spot and futures prices, respectively, at time t. The formula for the futures return reflects the fact that the return is defined relative to the actual investment in the spot market, not the nominal investment in the futures market (see Solnik, 2000). In order to initialise the EWMA and rolling window estimators, we discard the first 500 observations, and so the full evaluation sample is 29.04.86 to 03.05.02, a total of 4542 observations. As well as the full sample, we consider three equal length sub-samples, 29.04.86 to 23.08.91, 24.08.91 to 18.12.96 and 19.12.96 to 03.05.02, which captures periods of very different market conditions. Table 1, below gives summary statistics for the full sample and each of the three sub-samples.

[Table 1]

For the full sample, and for all three sub-samples, the mean return for both spot and futures is close to zero, as assumed in the empirical methodology. The volatility of futures returns is marginally higher than the volatility of spot returns for the full sample. Spot return volatility is highest for the third sub-sample, while futures return volatility is highest for the first sub-sample. For both spot and futures returns, volatility is lowest in the second sub-sample. The skewness and excess kurtosis coefficients indicate that both spot and futures returns are highly non-normal, but there are considerable differences in the degree of non-normality in different sub-samples. The first sub-sample captures the stock market crash of 1987 (the minimum spot return of 13.0% and the minimum futures return of 17.9%), as well as several other large one-day market movements, and this is reflected in the very high excess kurtosis coefficient for this period. Nevertheless, all three sub-periods display very significant non-normality, implying that the standard approach to calculation of the OHR will be inefficient.

To evaluate our estimator, we compute the OHR using both the unconditional estimator given by (4) and the conditional estimator given by (5). For the unconditional OHR, we use the estimator of the unconditional variance of the PE distribution given by (14) evaluated using the whole sample. For the conditional OHR, we employ two approaches. Firstly, we use a rolling window estimator of the unconditional variance of the PE distribution, based on (14), with window lengths of 125 days, 250 days and 500 days. Secondly, we use the power EWMA variance estimator given by (16), with λ set to 0.94, 0.96 and 0.98. The value λ = 0.94 is used by Riskmetrics to estimate daily volatility for the purpose of calculating Value at Risk. We apply the variance estimators to spot and futures returns, and to their sum and their difference, and compute the covariance between spot and futures returns using the identity given by (17).

In order to implement the variance and covariance estimators, we need to specify the power parameter, k. To establish the sensitivity of our estimator to different values of this parameter, and to compare the robust OHR estimators with the standard OHR estimators, we consider values of k equal to 1.00, 1.25, 1.50, 1.75, 2.00. The value $k = 2$ corresponds to the standard estimator. Each day, we use the conditional OHR in order to construct a hedged portfolio, and compute its return using (1). We report the variance of the hedged portfolio return and the variance of the OHR itself. For reference, we also report the variance of the unhedged portfolio. All of the computations were undertaken using Excel^{TM}. In practice, one could select the values of k and λ on the basis of out-of-sample calibration results similar to those reported here. This is the approach taken by Riskmetrics in their calibration of λ in the standard EWMA model, and is consistent with loss-function estimation of forecasting models (see, for instance, Clements and Hendry, 1998). Alternatively, one could specify a conditional distribution for spot and futures returns, and for their sum and their difference, and to estimate the parameters k and λ for each series using maximum likelihood.[4]

[4] For instance, maximum likelihood estimation for the full sample, specifying a conditional power exponential distribution and using the BHHH algorithm with a convergence criterion of 0.00001 applied to the function value, yields estimates of k and λ equal to 1.420 and 0.935 for spot returns, 1.295 and 0.945 for futures returns, 1.408 and 0.937 for the sum of spot and futures returns, and 1.026 and 0.951 for the difference between spot and futures returns.

5. Empirical Results

We first report results for each of the hedging strategies using the standard unconditional and conditional estimators of the variance-covariance matrix of returns, i.e. with k set to 2. Panel A of Table 2 reports the variance of the hedged portfolio using each of the strategies and the variance of the unhedged portfolio. Panel B reports the percentage reduction in the variance of the hedged portfolio over the unhedged portfolio. Using the standard estimators of the variance-covariance matrix, all of the hedging strategies yield a substantial reduction in the hedged portfolio variance. Note the while the unconditional OHR yields a reduction in the hedged portfolio variance, it could not be exploited by investors in real time since it is based on in-sample estimates of the variance-covariance matrix of returns. Of the conditional models, the rolling window with 500 observations provides the biggest reduction in variance for the full-sample, about 0.75% better than the best performing EWMA model. The rolling window approach also gives a lower hedged portfolio variance for the first and second sub-samples, but for the third sub-sample, the EWMA estimator with a decay factor of 0.98 yields the best variance reduction. Panel C of Table 2 reports the variance of the estimated OHR. As one would expect, the volatility of the OHR is higher for the EWMA estimator than for the rolling window estimator, and decreases with the EWMA decay factor and with the rolling window length.

[Table 2]

We now report results for each of the hedging strategies using the robust estimators of the variance-covariance matrix, using $k = \{1.00, 1.25, 1.50, 1.75\}$. Table 3 reports the percentage reduction in the hedged portfolio variance for each of these estimators over the corresponding standard hedged portfolio variance, reported in Panel A of Table 2. For both the rolling window estimators and the EWMA estimators, the robust approach generates a reduction in the hedged portfolio variance in the full sample across almost all the models and for all values of k, except for $k = 1.75$, where there is a small increase for the EWMA model when the decay factor is 0.94. In general, the reductions are largest for k = 1.00 or k = 1.25. The largest reduction is for the EWMA model with a decay factor of 0.98, yielding a reduction of 4.4%. The robust estimator also generates reductions in the hedged portfolio variance across all models for the first and second sub-samples. For the third sub-sample, hedged portfolio variance is largely unchanged by the use of the robust estimator, with generally small but insignificant increases. The largest reduction occurs in the first sub-sample, reflecting the fact that this is the most leptokurtic period, including as it does the stock market crash of October 1987. The largest reduction in the first sub-period is 6.8% for the EWMA model with a decay factor of 0.98, for k = 1.00. For the unconditional model, the use of the robust estimator has virtually no impact on the variance of the hedged portfolio, with, in most cases, small but insignificant increases in the variance. This is to be expected since the unconditional OHR is estimated using the full sample, where the reduced efficiency of the standard estimator is less likely to be a problem owing to the large sample size. Overall, therefore, it seems that use of the robust estimator generates modest reductions in the variance of the hedged portfolio for all of the hedging strategies considered, particularly during periods of high kurtosis such as those that include stock market crashes.

[Table 3]

Table 4 reports the percentage reduction in the estimated OHR itself for the robust estimators, compared to the standard case reported in Panel C of Table 2. For the full sample, and for all three sub-samples, there are very substantial reductions in the volatility of the estimated hedge ratio in virtually all cases. The reductions are particularly large for the rolling window estimator, but are also significant for the EWMA estimator, particularly for k = 1.25 or k = 1.50. The largest reduction is for the rolling window model with 500 observations, where the variance of the OHR is reduced by 71% for the second sub-sample, when k = 1.00. Even for k = 1.75, the volatility of the hedge ratio is very substantially reduced using all of the estimators. It seems, therefore, that while the reduction in the hedge portfolio variance from using the robust estimator is relatively modest, the reduction in the volatility of the estimated hedge ratio is substantial. Since transaction costs are typically quoted as a percentage of the value traded in the spot and futures markets, the use of the robust estimator should substantially reduce the costs of rebalancing that are associated with changes in the OHR.

[Table 4]

6. Conclusion

Dynamic hedging strategies involve estimating the OHR for a portfolio allowing for time-variation in the variance-covariance matrix of spot and futures returns. However, all of the standard approaches to estimating the OHR are based on the sample variance and covariance estimators of returns. These are consistent estimators of the population variance and covariance, irrespective of the underlying distribution of the data, but they are not in general efficient. As a result, hedging strategies based on the standard estimators of the OHR will yield hedged portfolios that will not generally have minimum variance because the estimated OHR itself will be excessively volatile. Moreover, the excessive volatility of the estimated OHR will increase the transaction costs that are incurred when rebalancing the hedged portfolio in response to changes in the OHR.

This paper proposes an alternative approach to the estimation of the OHR that is robust to the leptokurtosis of the distribution of returns. We use the robust OHR estimation procedure to construct a dynamic hedging strategy for daily returns on the FTSE100 index, using index futures contracts. We estimate robust conditional OHR s using both the rolling window approach and the EWMA approach, and compare our results to those based on the standard rolling window and EWMA estimators. It is shown that the robust estimator yields a hedged portfolio variance that is marginally lower than that based on the standard estimator. Moreover, the variance of the robust OHR is as much as 70% lower than the variance of the standard OHR, substantially reducing the transaction costs that are associated with dynamic hedging strategies. In the mean-variance framework, therefore, the robust OHR can be expected to enhance investor utility relative to the standard OHR both by reducing the risk of the hedged portfolio and by increasing its expected return

References

Alexander, C. and C. Leigh, 1997, On the Covariance Matrices Used in Value at Risk Models , *Journal of Derivatives* 4, 50-62.

Baillie, R., and R. DeGennaro, 1990, Stock Returns and Volatility , *Journal of Financial and Quantitative Analysis* 25, 203-214.

Benninga, S., R. Eldor, and I. Zilcha, 1984, The Optimal Hedge Ratio in Unbiased Futures Markets , *Journal of Futures Markets* 4, 155-159.

Bollerslev, T., 1986, Generalised Autoregressive Conditional Heteroskedasticity , *Journal of Econometrics* 31, 307-327.

Bollersley, T., 1990, Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Approach , *Review of Economics and Statistics* 72, 498-505.

Bollerslev, T., Engle, R. F., & Wooldridge, J. M., 1988, A Capital Asset Pricing Model With Time-Varying Covariances , *Journal of Political Economy* 96, 116-131.

Bondesson, L., 1975, Normal Distribution, Gamma Distribution and Quadratic Polynomial Statistics , *Scandinavian Journal of Statistics* 2, 138-144.

Boudoukh, J., M. Richardson and R. Whitelaw, 1997, Investigation of a Class of Volatility Estimation , *Journal of Derivatives* 4, 63-71.

Brooks, C., and J. Chong, 2001, The Cross-Currency Hedging Performance of Implied Versus Statistical Forecasting Models , *The Journal of Futures Markets* 21, 1043-1069 .

Cargill, T., and T. Rausser, 1975, Temporal Price Behaviour in Commodity Futures Markets , *Journal of Finance* 30, 1043-1053.

Cecchetti, S. G., R. E. Cumby, and S. Figlewski, 1988, Estimation of Optimal Futures Hedge , *Review of Econometrics and Statistic*s 70, 623-630.

Clements, M., and D. Hendry, 1998, *Forecasting Economic Time Series*, Cambridge University Press.

Doukas, J., and Rahman, A., 1987, Unit Roots Tests: Evidence from the Foreign Exchange Futures Market , *Journal of Financial and Quantitative Analysis* 22, 101-108

Ederington, L. H., 1979, The Hedging Performance of the New Futures Market , *Journal of Finance* 34, 157-170.

Engle, R. F., 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. inflation , *Econometrica* 50, 987-1008.

Engle, R. F., and T. Bollerslev, 1986, Modelling the Persistence of Conditional Variances , *Econometric Reviews* 5, 1-50.

Gagnon, L., and Lypny, G., 1995, Hedging Short-Term Interest Risk Under Time-Varying Distributions , *Journal of Futures Markets* 15, 767-783.

Goldenberg, D., 1988, Trading Frictions and Futures Price Movements , *Journal of Financial and Quantitative Analysis* 23, 465-483.

Guermat, C., and R.D.F. Harris, 2002, Robust Conditional Variance Estimation and Value at Risk , *Journal of Risk* 4, 25-41.

JP Morgan, 1996, *Riskmetrics*™ Technical Document, fourth edition, New York.

Kano, Y., 1994, Consistency Property of Elliptical Probability Density Funtions , *Journal of Multivariate Analysis* 51, 139-147.

Kroner, K. F., and Sultan, J., 1991, Exchange Rate Volatility and Time Varying Hedge Ratios , *Pacific-Basin Capital Markets Research* 2, 397-412.

Kroner, K. F., and Sultan, J., 1993, Time Varying Distribution and Dynamic Hedging with Foreign Currency Futures , *Journal of Finance and Quantitative Analysis* 28, 535-551.

Lin, J. W., Najand, M., and Yung, K., 1994, Hedging with Currency Futures: OLS vs. GARCH , *Journal of Multinational Financial Management* 4, 45-67.

Myers, R. J., and S. R. Thompson, 1989, Generalized Optimal Hedge Ratio Estimation , *American Journal of Agricultural Economics* 71, 858-868.

Nelson, D., and D. Foster, 1996, Asymptotic Filtering Theory for Univariate ARCH Models , Econometrica 62, 1-41.

Park, T. H., and Switzer, L. N., 1995, Bivariate GARCH Estimation of the Optimal Hedge Ratios for Stock Index Futures: A note , *Journal of Futures Markets* 15, 61-67.

Richard T. Baillie and Robert J. Myers, 1991, Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge , *Journal of Applied Econometrics* 6, 109-124.

Solnik, B., 2000, *International Investments*, Addison Wesley Longman Inc, USA, 4th ed.

Spanos, A., 1999, *Probability Theory and Statistical Inference, 2nd edition,* Cambridge University Press.

Taylor, S., 1985, The Behaviour of Futures Prices over Time , *Applied Economics* 17, 713-734.

Tong, W. H. S., 1996, An Examination of Dynamic Hedging , *Journal of International Money and Finance* 15, 19-35.

Table 1 Summary Statistics

	Spot Returns				Futures Returns			
	Full Sample	Period 1	Period 2	Period 3	Full Sample	Period 1	Period 2	Period 3
Mean	0.00034	0.00036	0.00031	0.00019	0.00033	0.00035	0.00030	0.00019
Stand Dev	0.01001	0.01060	0.00755	0.01184	0.01143	0.01307	0.00893	0.01270
Minimum	-0.13029	-0.13029	-0.04140	-0.06401	-0.17883	-0.17883	-0.04615	-0.07055
Maximum	0.07597	0.07597	0.05440	0.04345	0.07997	0.07997	0.05042	0.04644
Skewness	-0.70587	-1.81187	0.32883	-0.20241	-1.10043	-2.36129	0.15974	-0.15471
Ex. Kurtosis	9.50900	21.91022	3.54270	1.18687	16.61940	30.69602	1.67282	1.26851
Bera-Jarque	16577.83	26536.48	778.23	33.65	51438.68	52353.43	192.92	55.62
Sample Size	4542	1347	1347	1348	4542	1347	1347	1348

Notes: The table reports summary statistics for the FTSE 100 index spot and futures returns series. The Bera-Jarque statistic is computed as $T(b_1/6+b_2^2/24)$, where b_1 and b_2 are estimated coefficients of skewness and excess kurtosis, respectively, and T is the sample size, and has a chi-squared distribution with two degree of freedom under the null hypothesis.

Table 2 Standard Optimal Hedge Ratio (k = 2.00)

Panel A: Hedged Portfolio Variance

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	1.0307	0.1475	0.1489	0.1466	0.1457	0.1679	0.1602	0.1534
Period 1	1.1221	0.2586	0.2969	0.2889	0.2844	0.3511	0.3294	0.3103
Period 2	0.5703	0.0522	0.0530	0.0527	0.0528	0.0541	0.0538	0.0533
Period 3	1.4009	0.1033	0.0971	0.0986	0.1004	0.0988	0.0978	0.0971

Panel B: Percentage Reduction in Hedged Portfolio Variance

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-85.69	-85.56	-85.77	-85.86	-83.71	-84.46	-85.11
Period 1	-	-76.96	-73.54	-74.25	-74.66	-68.71	-70.64	-72.35
Period 2	-	-90.85	-90.70	-90.76	-90.74	-90.51	-90.56	-90.65
Period 3	-	-92.62	-93.07	-92.96	-92.83	-92.95	-93.02	-93.07

Panel C: Estimated OHR Variance

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-	0.0069	0.0060	0.0048	0.0100	0.0086	0.0072
Period 1	-	-	0.0034	0.0023	0.0035	0.0086	0.0061	0.0029
Period 2	-	-	0.0009	0.0004	0.0005	0.0033	0.0023	0.0012
Period 3	-	-	0.0053	0.0045	0.0026	0.0074	0.0065	0.0056

Notes: Panel A reports the variance (times 10,000) of the hedged portfolio returns, $x_t = s_t - h_t f_t$, for the full sample and for each of the three sub-samples using (a) the unconditional OHR computed using the full-sample, (b) the conditional OHR computed using rolling window estimators of the variance-covariance matrix with window lengths of 125, 250 and 500 days, and (c) the conditional OHR computed using EWMA estimators of the variance-covariance matrix with decay factors of 0.94, 0.96 and 0.98. In all cases, the variance-covariance matrix is estimated using k = 2.00. Panel B reports the percentage variance reduction of portfolio returns over the unhedged portfolio variance. Panel C reports the variance of the estimated OHR for each model.

Table 3 Percentage Reduction in Variance of Hedged Portfolio Over Standard Estimator

***k* = 1.00**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	0.83	**-3.79**	**-2.81**	**-1.06**	**-3.28**	**-3.80**	**-4.40**
Period 1	-	2.71	**-6.23**	**-4.59**	**-1.49**	**-5.25**	**-5.91**	**-6.82**
Period 2	-	0.01	0.02	**-0.25**	**-0.58**	**-0.07**	**-0.37**	**-0.39**
Period 3	-	0.10	1.54	1.00	**-0.09**	1.97	1.42	1.11

***k* = 1.25**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	0.56	**-3.50**	**-2.94**	**-1.65**	**-1.28**	**-1.99**	**-3.18**
Period 1	-	1.67	**-5.56**	**-4.62**	**-2.43**	**-2.11**	**-3.07**	**-4.85**
Period 2	-	0.03	**-0.06**	**-0.23**	**-0.39**	**-0.28**	**-0.43**	**-0.38**
Period 3	-	0.05	0.94	0.54	**-0.11**	1.09	0.79	0.62

***k* = 1.50**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	0.33	**-2.64**	**-2.49**	**-1.73**	**-0.10**	**-0.67**	**-1.83**
Period 1	-	0.85	**-4.12**	**-3.85**	**-2.57**	**-0.26**	**-1.04**	**-2.75**
Period 2	-	0.02	**-0.14**	**-0.20**	**-0.25**	**-0.33**	**-0.40**	**-0.35**
Period 3	-	0.02	0.51	0.24	**-0.11**	0.57	0.40	0.31

***k* = 1.75**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	0.11	**-1.38**	**-1.46**	**-1.20**	0.23	**-0.05**	**-0.71**
Period 1	-	0.25	**-2.13**	**-2.23**	**-1.79**	0.30	**-0.08**	**-1.05**
Period 2	-	0.00	**-0.14**	**-0.14**	**-0.13**	**-0.24**	**-0.27**	**-0.24**
Period 3	-	0.01	0.21	0.07	**-0.08**	0.22	0.16	0.12

Notes: The table reports the percentage reduction of the variance of hedged portfolio returns for k = 1.00, 1.25, 1.50 and 1.75, over the standard OHR estimator (k = 2.00). Reductions in the variance are reported in bold.

Table 4 Percentage Reduction in Variance of OHR Over Standard Estimator

***k* = 1.00**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-	**-23.29**	**-30.74**	**-38.00**	**-6.59**	**-12.13**	**-20.48**
Period 1	-	-	**-14.58**	**-17.76**	**-27.25**	**-3.81**	**-7.21**	1.59
Period 2	-	-	11.73	**-18.66**	**-71.10**	6.19	6.98	10.22
Period 3	-	-	**-14.25**	**-23.01**	**-23.41**	15.31	6.08	**-8.52**

***k* = 1.25**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-	**-20.01**	**-25.85**	**-32.96**	**-8.51**	**-12.15**	**-18.01**
Period 1	-	-	**-18.04**	**-22.03**	**-27.70**	**-8.58**	**-11.02**	**-5.09**
Period 2	-	-	**-0.20**	**-23.86**	**-60.76**	**-3.58**	**-3.44**	**-1.48**
Period 3	-	-	**-11.36**	**-17.42**	**-18.87**	8.19	2.14	**-7.53**

***k* = 1.50**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-	**-14.93**	**-19.29**	**-25.53**	**-7.37**	**-9.67**	**-13.65**
Period 1	-	-	**-16.27**	**-20.68**	**-23.49**	**-8.34**	**-10.02**	**-6.41**
Period 2	-	-	**-6.23**	**-23.72**	**-45.87**	**-6.18**	**-6.76**	**-6.64**
Period 3	-	-	**-7.61**	**-11.39**	**-13.08**	4.09	0.49	**-5.30**

***k* = 1.75**

	Unhedged	Uncond. Hedge	Rolling Window 125 Days	Rolling Window 250 days	Rolling window 500 days	EWMA 0.94	EWMA 0.96	EWMA 0.98
Full sample	-	-	**-8.24**	**-10.77**	**-14.89**	**-4.37**	**-5.52**	**-7.67**
Period 1	-	-	**-10.15**	**-13.57**	**-14.57**	**-5.18**	**-6.09**	**-4.36**
Period 2	-	-	**-6.86**	**-17.17**	**-26.20**	**-4.60**	**-5.40**	**-6.42**
Period 3	-	-	**-3.72**	**-5.48**	**-6.68**	1.55	**-0.06**	**-2.66**

Notes: The table reports the percentage reduction of the variance of the estimated OHR for k = 1.00, 1.25, 1.50 and 1.75, over the standard OHR estimator (k = 2.00). Reductions in the variance are reported in bold.