UCD GEARY INSTITUTE
DISCUSSION PAPER SERIES

Performance of Utility Based Hedges

John Cotter*
UCD School of Business,
University College Dublin

Jim Hanly*
College of Business,
Dublin Institute of Technology

_____

Performance of Utility Based Hedges

John Cotter and Jim Hanly[*]

Abstract

Hedgers as investors are concerned with both risk and return; however the literature has generally neglected the role of both returns and investor risk aversion by its focus on minimum variance hedging. In this paper we address this by using utility based performance metrics to evaluate the hedging effectiveness of utility based hedges for hedgers with both moderate and high risk aversion together with the more traditional minimum variance approach. We apply our approach to two asset classes, equity and energy, for three different hedging horizons, daily, weekly and monthly. We find significant differences between the minimum variance and utility based hedges and their attendant performance in-sample for all frequencies. However out of sample performance differences persist for the monthly frequency only.

Keywords: Hedging Performance; Utility, Energy, Risk Aversion.
JEL classification: G10, G12, G15.

December 2013

1. INTRODUCTION

Recent uncertainty and the accompanying volatility in financial markets have highlighted the need for strategies that can address the risks associated with large and unexpected market movements. This is particularly important for energy consumers given that supply side shocks can arise from any number of sources, manmade or natural. It is also important for all financial market participants, as recent events have shown how even well diversified investors can suffer from large market movements over even short time horizons[1].

Hedging using futures has become an important way for investors to manage the risk of their exposures and a large literature has developed alongside the use of futures, detailing how to estimate an Optimal Hedge Ratio (OHR). This literature has in the main assumed that investors are infinitely risk averse and that therefore the optimal hedge is one that minimizes risk as measured by the variance[2]. A large number of papers (see for example, Lien, 2004) have used a utility maximisation approach but have assumed that futures prices follow a martingale in which case the utility maximising and variance minimising approaches are the same. A more limited number of papers have looked at hedging for other risk measures such as Value at Risk (Harris and Shen, 2006), however, only a relatively small number of papers (see for example, Lien, 2007, 2012) have followed a non-martingale utility based approach which implies that investors are concerned not just with risk, but also with the expected return. Also, where hedging performance has been examined, a variety of performance measures have been applied but in many cases, this has been done in a manner which fails to match the performance metric with the underlying optimization method for generating the hedge ratio[3]. Furthermore, papers that have looked at utility based hedges tend to do so only for a single time horizon.

---

[1] For example, the S&P500 showed daily movements in excess of 5% on over 10 occasions since the onset of the 2008 financial crisis.

[2] One of the first papers to detail this approach is Ederington (1979). Papers have also used a conditional framework using GARCH models. See Brooks et al. (2002) for example.

[3] For example, Kroner and Sultan, 1993 assume that futures prices are a martingale and therefore effectively estimate a minimum variance hedge which they evaluate using average quadratic utility to measure hedging performance

In this paper we contribute by addressing some of these issues. We estimate and compare optimal utility based hedges using performance metrics based on the underlying optimization criteria as well as an economic comparison between the different hedging models using Value at Risk (VaR). Given the potential impact of hedge horizon on optimal hedges (see Juhl, Kawaller and Koch, 2012) we estimate hedge strategies across three time horizons, daily, weekly and monthly. For our analysis we use two of the most broadly used characterisations of investor utility; quadratic and exponential. We also recognise that hedgers are not homogenous in terms of their risk aversion. We therefore estimate utility based hedges allowing for investors who are infinitely risk averse to more realistic approximations including moderate and low risk aversion. This allows us to analyse the impact that different attitudes towards risk have on the OHR. Furthermore, while the literature has generally used a framework based on short hedgers, we apply our analysis to long hedgers who are short the underlying asset, since such hedgers are in many respects more susceptible to unexpected market movements[4]. Finally, we examine how utility based hedges perform in terms of their optimisation criteria in an out-of-sample setting. Our approach allows us to make a wide ranging comparison that encompasses different estimation methods and performance criteria across different time horizons, to see how well utility based hedges perform in a variety of settings as compared with more traditional approaches such as variance minimisation.

Our empirical results for both energy and equity hedgers show that there are significant differences for the quadratic and exponential hedge ratios and also between each of the utility based hedges and the variance minimising hedge. This relates to the non-normal characteristics of both equity and energy returns such as the presence of skewness and excess kurtosis. We use the S&P500 and the West Texas Intermediate (WTI) Light Sweet Crude to proxy for equity and energy markets respectively. Since these characteristics are generally present for such data (and for financial market data in general), it shows the importance of matching the utility function used to generate the hedge to the risk investor's preferences. There are also significant differences between the OHR's across time horizons with a clear implication that optimal hedges are frequency specific. We also find that using

---

[4] Demirer and Lien (2003) found that long hedgers tend to be more active in the futures market than short hedgers which may indicate that they felt more exposed to adverse price movements.

performance methods based on the optimization criteria will yield significant performance differentials for equity hedger's but not for oil hedgers between the different hedging models . The differences are more marked for investors with moderate risk aversion and at the monthly hedging frequency. The implication for oil market participants is that hedging performance is not significantly different enough to warrant using a more complicated utility based approach as compared with the simpler minimum variance based hedge.

The remainder of this paper proceeds as follows. In Section 2 we describe the utility based performance measures together with the associated models used for estimating optimal hedge ratios. Section 3 details the utility framework we employ. Section 4 introduces our data set. Section 5 presents our empirical results. Section 6 concludes.

2.   HEDGING MODELS AND PERFORMANCE

Optimal hedge strategies using futures have generally followed two distinct approaches. The first approach assumes that investors are essentially infinitely risk averse. This means that hedgers will ignore the return component of their investments and focus on minimising the risk via a minimum variance hedge ratio (MVHR). The problem with this approach is that it doesn't explicitly incorporate an investor's risk aversion, or the expected return in the estimation of the optimal hedge. The second approach focuses on maximizing expected utility. This framework incorporates risk, expected return and risk aversion as key elements in hedge ratio estimation. Many papers have avoided making a distinction between the utility maximizing and the variance minimising approaches by assuming that futures prices are unbiased[5]. Under this assumption, the MVHR and the utility maximizing hedge ratios are equivalent, however, there is evidence that oil futures markets (Sadorsky, 2002) and equity futures markets (Chen Lee and Shrestha, 2001, Fung, Jarrett and Leung, 1990) do not follow a martingale.

---

[5] A majority of the literature on optimal hedging takes this approach despite some evidence to the contrary. See deVille deGoyet, Dhaene and Sercu (2008) for evidence and a discussion.

In this paper, we examine the hedging performance of four different hedging strategies with a focus on hedge strategies that are optimised based on attitudes towards risk and return via the investors utility function. This enables us to incorporate different attitudes towards risk within the hedging context and to allow the expected return to play a part. We use a rolling window model based hedge, incorporating three different approaches that reflect an infinitely risk averse investor, an investor that has a quadratic utility function[6] and an investor that has an exponential utility function. Let us first turn to our hedging models and then detail the performance methods used.

2.1 Optimal Hedge Ratio's under the Expected Utility Framework

Assuming a fixed spot position let r_{st} and r_{ft} be logarithmic returns on the spot and futures series respectively, and β be the OHR. The return to the hedged portfolio is constructed as follows:

$$R_p = -r_s + \beta r_f \tag{1}$$

We define the OHR as the weight of the futures asset in the hedged portfolio that is chosen either to minimise risk or to maximize expected utility, depending on the underlying framework that is being applied. Assuming that the agent has a quadratic utility function, R_p, then the OHR can be calculated as:

$$\beta = \frac{-E(r_{ft})}{2\lambda \sigma_{ft}^2} + \frac{\sigma_{sft}}{\sigma_{ft}^2} \tag{2}$$

where $E(r_{ft})$ is the expected return on futures, λ is the risk aversion parameter, σ_{ft}^2 is the futures variance and σ_{sft} is the covariance between spot and futures. Equation (2) establishes the relationship between the risk aversion parameter λ and the OHR. As risk aversion increases, the individual hedges more and speculates less relative to the spot position, such that for extremely large levels of risk aversion, the first term in equation (2) will have a negligible influence on the hedge strategy. Therefore, the OHR for an investor with infinite risk aversion will be the minimum variance hedge which is given by

$$\beta = \frac{\sigma_{sft}}{\sigma_{ft}^2} \tag{3}$$

---

[6] We estimate both quadratic and exponential hedges using moderate and high values for risk aversion.

For an investor with exponential utility, their OHR will approximate the quadratic hedge where returns are normally distributed, however for non-normally distributed returns they will differ. The optimal exponential hedge can be estimated by choosing β to maximize the following expression:

$$\beta = \mu_{pt} - \frac{1}{2}\lambda\sigma_{pt}^2 + \frac{\tau}{6}\lambda^2\sigma_{pt}^3 - \frac{\kappa-3}{24}\lambda^3\sigma_{pt}^4$$

(4)

where μ_{pt} is the expected return on the hedged portfolio, λ is the risk aversion parameter, σ_{pt}^2 is the variance of the portfolio, τ is the skewness of the portfolio and κ is the kurtosis[7].

We use the following hedges.

2.2   No Hedge

This is a hedge ratio (HR) of zero, the exposure is left unhedged. This is included for use as a performance benchmark. This is included as many firms and financial market participants choose not to hedge their exposures since hedging reduces not only risk, but also the expected return to bearing that risk.

2.3   Model Based Hedge

GARCH models have been broadly applied in the hedging literature to estimate time varying OHR's (see for example, Cotter and Hanly, 2006). However, the performance of these models has been mixed and in particular their out-of-sample performance has been average at best (Brooks Cerny and Miffre, 2012). This suggests that there is little to be gained in terms of hedging efficiency from the use of GARCH models despite their additional complexity, as compared with rolling window OLS models (Miffre, 2004). We therefore use a rolling window approach to estimate our model based hedges given its relative simplicity and robust performance. This estimates the variance covariance matrix by conditioning on recent information on a period-by-period basis whereby the most recent observation is added and the oldest observation is removed from the sample, thus

---

[7] See Alexander (2008) for more details.

keeping the number of estimation observations unchanged. The advantage of this method from a hedging perspective is that by updating the information set we still obtain a more efficient estimate of the hedge ratio, which takes time variation in the return distribution into account. However, in common with other methods that require the hedge ratio to be time-varying, it may be expensive as changing the hedge ratio increases transactions costs.

To obtain the OHR's we require estimates of the variance and covariance's of the spot and futures for both the equity and oil contracts. For the utility maximising hedges we also require risk aversion estimates. We use two values of the CRRA to reflect values for the risk aversion parameter that are found in the asset pricing literature[8]. We use a value of five which reflects a moderately risk averse investor and a value of ten to reflect a more risk averse investor[9]. The MVHR is estimated using equation (3), and the quadratic utility maximizing HR (Quad) is estimated using equation (2). The exponential HR is obtained by maximizinq equation (4). Each of the utility maximizing hedges is estimated for a CRRA of 5 and a CRRA of 10. In effect this gives us a total of six hedges to consider. The procedure is outlined in more detail in Section 4.

2.4   Hedging Effectiveness

A key contribution of this paper is that we align the methods used for hedge ratio estimation with the performance methods used to measure hedging effectiveness. We focus on three different hedging models, each designed to optimise different criterion. These are variance minimisation, quadratic utility maximisation and exponential utility maximisation. Additionally, we evaluate each of the different OHR's across all three hedging effectiveness measures. This allows us to build up a profile, not only of how well a given hedge works within its own optimisation setting, but also how well that method transfers for broader application. This is important because hedgers may seek hedging strategies that perform

---

[8] A number of estimates for the CRRA have been used in the literature, but they generally fall within the range 1 - 10. See for example Mehra and Prescott (1985), Ghysels, Santa Clara, and Valkanov (2005), Guo and Whitelaw (2006). See also Cotter and Hanly (2010) for a more comprehensive discussion and derivation of the CRRA in the hedging context.
[9] We also estimated but do not report hedge ratios using a risk aversion value of 1 which reflects an investor with a log utility function. The resulting hedges were in excess of two and considered unrealistic.

well across a broad range of criteria (see Cotter and Hanly, 2006; who found diverse hedging performance using different evaluation methods).

The first performance metric we use to examine hedging performance is the variance. The variance metric (HE_1) measures the percentage reduction in the variance of a hedged portfolio as compared with the variance of an unhedged portfolio[10]. The performance metric is:

$$HE_1 = 1 - \left[\frac{VARIANCE_{HedgedPortfolio}}{VARIANCE_{UnhedgedPortfolio}} \right] \tag{5}$$

This gives us the percentage reduction in the variance of the hedged portfolio as compared with the unhedged portfolio. When the futures contract completely eliminates risk, we obtain $HE_1 = 1$ which indicates a 100% reduction in the variance, whereas we obtain $HE_1 = 0$ when hedging with the futures contract does not reduce risk. Therefore, a larger number indicates better hedging performance. As the standard measure of risk in finance, it is important to include the variance despite its shortcomings, chief of which is that it implicitly caters for investors who are infinitely risk averse. Since this is not a realistic proposition for real world investors we also use utility based measures. Therefore, the second metric we use to examine hedging performance is (HE_2) which is based on quadratic utility. This measures the percentage increase in the quadratic utility of a hedged portfolio as compared with the quadratic utility of an unhedged portfolio. It is worth noting that expected utility doesn't have a simple intuitive interpretation[11] in a way that other statistical measures do. The number of itself does not have an inherent meaning; rather it is the ordering of the utilities that matters. For this reason using the percentage change is important as it enables us to translate the utilities to provide a clear and unambiguous measure to allow us to compare utility based hedging effectiveness. The performance metric is:

$$HE_2 = 1 - \left[\frac{Quadratic\_Utility_{HedgedPortfolio}}{Quadratic\_Utility_{UnhedgedPortfolio}} \right] \tag{6}$$

---

[10] An unhedged portfolio has a hedge ratio of zero.
[11] See Alexander (2008) page 228 for further discussion.

The third metric we use to examine hedging performance is (HE₃) which based on the exponential utility. This measures the percentage increase in the exponential utility of a hedged portfolio as compared with the exponential utility of an unhedged portfolio. The performance metric is:

$$HE_2 = 1 - \left[\frac{Exponential\_Utility_{HedgedPortfolio}}{Exponential\_Utility_{UnhedgedPortfolio}} \right] \qquad (7)$$

Note that for the variance metric, hedging effectiveness relates to variance reduction whereas for the utility based measures it relates to an increase in utility.

3.   UTILITY

In this section we provide further detail on the two characterisation of investor utility that we employ.

3.1   Quadratic Utility

The quadratic utility function has had broad application in finance and economics in areas including portfolio theory, asset pricing and hedging. In the hedging literature it underlies much of the literature on optimal hedging which is based on the expected utility maximisation paradigm (Baron, 1977, Ederington, 1979).

It is defined as follows:

$$U(W) = W - aW^2, \qquad a > 0 \qquad (8a)$$

Define U ( ) as the utility function and W as wealth, a is a positive scalar parameter measuring risk aversion. The first and second derivatives of this are given by:

$$U^{'}(W) = 1 - 2aW \qquad (8b)$$

$$U^{''}(W) = -2aW \qquad (8c)$$

To be consistent with non-satiation where utility is an increasing function of wealth implying more is preferable to less, the following restriction is placed on W:

$$U^{'}(W) = 1 - 2aW > 0$$

The relative risk aversion measure is:

$$R(W) = \frac{2aW}{1 - 2aW} \qquad (8d)$$

9

The quadratic utility function is consistent with a hedger who decreases the percentage invested in risky assets as wealth increases. This implies increasing relative risk aversion[12]. This characterization of risk aversion is consistent with an economic rationale whereby wealthy investors have less need for higher return investments and can pursue a conservative investment strategy by investing in lower risk and lower return assets as their wealth increases.

3.2   Exponential Utility

Exponential utility has been broadly applied as a useful representation of investor behaviour (see for example, Ho, 1984). It has been applied in a hedging context by a number of papers including, Lien (2003), and Brooks, Cerny and Miffre (2012). It is defined as follows:

$$U(W) = -e^{-aW}, a > 0 \tag{9a}$$

$$U'(W) = ae^{-a*W} \tag{9b}$$

$$U''(W) = -a^2 e^{-a*W} \tag{9c}$$

$$R(W) = \frac{-W(-a^2 e^{-aW})}{ae^{-aW}} = -Wa \tag{9d}$$

Investors with exponential utility invest decreasing proportional amounts in risky assets as their wealth increases. Therefore this utility function is also consistent with increasing relative risk aversion.

4.   DATA

For the empirical analysis, we use spot and futures prices for two asset classes; energy and equities. The energy contract used is West Texas Intermediate (WTI) Light Sweet Crude. For equities we use the S&P500 Composite Equity Index. Both of these assets are traded on CMEGROUP and they represent highly liquid spot and futures markets with a long and robust returns history[13]. They were chosen as they represent two of the most important

---

[12] For evidence and a discussion of increasing relative risk aversion see Eisenhauer and Ventura (2003).
[13] Contract details are available from http://www.cmegroup.com/trading/energy/crude-oil/light-sweet

pricing benchmarks for energy and equities respectively. Our full sample runs from 1st January 1990 until the 5th of September 2011 and includes data at daily (1-day), weekly (5-day) and monthly (20-day) frequencies. This large timeframe was chosen as it contains both tranquil and volatile periods, is of sufficient length to examine daily, weekly and monthly time horizons, and because it incorporates a number of key events for both energy and equity prices. This allows us to compare hedging scenarios that reflect the different holding periods of hedgers and includes periods of intense volatility which is useful for an examination of hedging performance. All data were obtained from Datastream and returns were calculated as the differenced logarithmic prices. A continuous series was formed with the contract being rolled over by largest volume.

Figure 1 provides a time series plot and Table 1 provides descriptive statistics of the data for the full sample period. The characteristics of the data are in line with those commonly observed for energy and equity assets, namely, the presence of both skewness and kurtosis. Jarque-Bera (J-B) statistics indicate non-normality for each series at all frequencies[14]. For both assets we see a positive mean which is indicative of the strong price rises over the period. This is particularly marked for oil which shows monthly average returns of almost half a per cent. The strong price increases for oil in the period up to 2008 has been linked to uncertainty about the supply of oil based products due to large increases in demand attributed to the BRIC[15] countries. Both series also exhibit periods of marked volatility, most notably associated with the financial crisis of 2008. The large changes in price and the rise in volatility for our sample period is useful as it allows us to examine whether hedging would be effective in circumstances where investors seek to reduce their exposures.

[FIGURE 1 HERE]

[TABLE I HERE]

---

crude_contract_specifications.html.
http://www.cmegroup.com/trading/equity-index/us-index/sandp-500_contract_specifications.html

[14] There are some differences in the distributional characteristics of the data at the different frequencies, This is of relevance to utility based hedgers in that where data is closer to normal (e.g. for the monthly frequency), hedging strategies may be more similar across different utility functions, whereas significant skewness and excess kurtosis will affect the hedging strategies and cause them to diverge for different utility functions such as the quadratic and exponential.
[15] Brazil, Russia, India, China

4.1 Estimation Procedure

To estimate the utility maximizing hedges for use in the period $t-n$ to t, from equations (2) and (4) we require estimates of $E(r_{ft}), \sigma^2_{ft}, \sigma_{sft}, \tau$, and κ, as well as the estimated risk aversion parameter λ. The first hedge uses data from January 1990 until January 2003. To allow the hedge ratio to vary over time, we have adopted a rolling window approach with a window length of 13 years[16] whereby the newest observation is added and the oldest observation is dropped therefore keeping the window length unchanged. In this way we generate 1560 in-sample hedges at the daily frequency, 312 at the weekly frequency and 170 hedges at the monthly frequency for the five year period from January 2003 to December 2008. The remaining three years of observations are used to generate out-of-sample hedges. We use two different values of the risk aversion coefficient, 5 and 10 to allow for different risk preferences. These are kept constant to allow the underlying changes in the variance covariance matrix to be the key driver of the OHR's[17] We also calculate a time-varying MVHR using the same rolling window methodology but employing (3). This doesn't incorporate the risk aversion parameter and is based on the variance covariance matrix alone.

The second set of OHR's we estimate are 1-step ahead forecast hedges for use in period t+1. To do this, we reserved a sub-period of three years of data to allow us to generate forecasted OHR's in a consistent manner. The time-varying hedges were forecast using a random walk process whereby we used the estimates from the t-period hedges to generate hedges for use in period t+1[18]. Using this methodology, we obtained 700 hedges for use in period t+1 at the daily frequency, 140 at the weekly frequency and 35 hedges at the monthly frequency. These covered the period from January 2009 to September 2011. The advantage of our approach is that it allows us to generate sufficiently large numbers of hedges for analysis using relatively low frequency data.

---

[16] This allows for sufficient observations at the weekly and monthly frequency for robust estimation

[17] In related work, we estimated and applied a time varying risk aversion coefficient to match the time varying characteristic of the variance covariance matrix. See Cotter and Hanly (2010) for more detail.

[18] Although not reported, we carried out Stationarity tests on the OHR's using the Adjusted Dickey Fuller Test. Results indicated that each of our model based hedges were non stationary for both Equities and Oil. This finding motivated our use of a random walk approach to forecast one step ahead hedge ratios for use out-of-sample. See Poon and Granger (2003) for further discussion.

5.   EMPIRICAL RESULTS

5.1   Optimal Hedges

We examine the optimal hedges for each of the two assets, S&P500 and Oil. Figure 2 plots a comparison of the OHR's for each of the different utility functions, quadratic and exponential together with the MVHR. Weekly frequency is chosen for illustration. Table 2 presents summary statistics for each of the hedges for the in-sample period.

[FIGURE 2 HERE]

[TABLE 2 HERE]

We can see from Figure 2 that there are large differences between the MVHR and both of the utility based hedges for both S&P500 and Oil. More specifically we can see that the utility based hedges are larger than the MVHR. For example, from Table 2 reading across, we can see that mean values for the OHR for S&P500 using the quadratic and exponential utility functions with a CRRA of 5 are 1.183 and 1.442 respectively as compared with 0.934 for the MVHR. Similar differences are observed for each frequency. This finding shows that for non-normal data, there will be substantial differences between not just the MVHR and the quadratic utility hedge but also between the quadratic utility and exponential utility based hedges. This relates to the presence of skewness and kurtosis in both equity and oil returns data and supports Lien (2007) but differs from Brooks Cerny and Miffre (2012) who find only small differences between MVHR's vs. utility based hedges. The implication is that hedge strategies that focus only on risk reduction may not provide the best outcomes in terms of expected utility.

We can also see that for the utility based hedges the risk aversion parameter has a clear impact with the spread between the hedges narrowing when we use the more risk averse CRRA value of 10 as compared with the moderate value of 5. For example, reading across Table 2 for the daily frequency using Oil, the mean OHR for the quadratic utility function with a CRRA of 5 is 0.990. This is significantly different than the equivalent value of 0.959 which is estimated using a CRRA of 10. Note that for the S&P500 in all cases the utility based OHR's are larger than one. For Oil, both of the utility hedges are larger than one for the weekly and monthly but not the daily frequency. These findings indicate that allowing

13

the expected return to play a part in determining the OHR via a utility function results in most cases in hedgers holding futures in excess of their initial spot position in the underlying asset. They also show the importance of using a risk aversion coefficient that reflects the hedgers individual risk preferences. These findings support earlier work by deVille deGoyet, Dhaene and Sercu (2008).

Table 3 further illustrates this by a comparison of the absolute differences between the mean of each of the OHR's. In all cases there are significant differences between the MVHR and the utility based hedge ratios, with differences smaller but still significant for the larger CRRA of 10. Using the daily frequency for the S&P500 with CRRA 5 for example, the difference between the exponential and quadratic OHR's is 0.258 as compared with 0.130 for a similar comparison using CRRA 10. This reflects the fact that as risk aversion increases, the speculative component plays a smaller part in the utility based hedges and they then tend to converge towards the MVHR.

[TABLE 3 HERE]

We next compare the OHR's for hedgers with different hedging horizons. Again, from Table 2 we can see that there are significant differences between the hedge ratios for different frequencies in every single case across all three hedge ratios, both risk aversion values and for both assets. For example, reading down, mean HR values for S&P500 using the quadratic utility function with a CRRA of 5 ranges from an average of 1.183 for daily data to 1.23 and 1.28 for weekly and monthly data respectively. The comparable values for Oil are 0.990, 1.044 and 1.084. This shows that using a hedge estimated at one hedging horizon, for example daily, which may be useful for an Oil commodities trader is inappropriate for a hedger such as an airline for whom a monthly hedging time horizon is more appropriate.

Looking next at the two different assets, from the previous example we can see that the OHR values are generally lower for Oil as compared with the S&P500. The differences between the two assets may indicate that for equity hedgers the speculative component plays a larger part in determining the optimal hedge than for Oil hedgers. A possible reason is that there may be a larger proportion of Oil market participants such as long oil hedgers who are consumers and who are in the market for reasons other than speculation.

We now turn to hedging performance which we examine using three performance metrics, (HE$_1$) variance, (HE$_2$) quadratic utility and (HE$_3$) exponential utility which were chosen to match the underlying optimisation goals of the model based hedge ratios. Table 4 presents underlying values for these metrics which are calculated for each OHR; MVHR, quadratic utility hedge and exponential utility hedge together with No hedge which is used as a benchmark for comparison purposes. Values are presented for both CRRA of 5 and a CRRA of 10.

[TABLE 4 HERE]

From Table 4 we can observe the general effectiveness of hedging in that each model yields a large improvement in the underlying optimization criteria as compared with an unhedged position. We also note that the best hedging strategy for a given performance metric is the hedging model that is designed to optimize that metric; be it variance reduction or utility maximisation. For example, for both assets and across frequency and risk aversion values, we can see that in all cases the MVHR returns the minimum variance while the quadratic HR returns the best quadratic utility, as does the exponential hedge for the exponential utility. These results show that it is important to match the performance criterion with the process for estimating the hedge ratio. It also indicates that applying a MVHR which has been commonly done in the literature as an approximation for all hedgers irrespective of their utility functions or attitudes towards risk may not be appropriate. [19]

Turning next to the relative performance of the various hedging models, Table 5 presents performance measures (HE$_1$), (HE$_2$) and (HE$_3$) which show the percentage change for a given metric as compared with a No Hedge scenario. Taking the S&P500 first, and looking at daily hedges for a CRRA of 5, reading across Table 5 we can see that hedging using the MVHR reduces (HE$_1$), the variance criterion by 93.69% as compared with 86.25% for the quadratic HR but only 63.97% for the exponential HR. This shows that there is a significant performance differential between the model based hedges depending on the metric used to measure performance. If variance reduction is the optimisation criterion chosen, both of the utility based hedges will return relatively poor results for an equity hedger. For the

---

[19] A large number of the papers in the literature on optimal hedging use OLS or GARCH based methodologies to optimize the hedge ratio based on the minimum variance criterion. See for example Lien (2002) for a review.

higher CRRA 10, the performance differential persists albeit at a lower level. This relates to the tendency for quadratic and utility based hedges to move closer to the MVHR for higher levels of risk aversion. This result is broadly similar at all frequencies so it will apply to hedgers irrespective of hedging horizon.

[TABLE 5 HERE]

If we use (HE$_2$) and (HE$_3$) the utility based performance metrics, we can see that in many cases, hedging more than doubles the utility of a hedger as compared with a no hedge strategy. For example, using CRRA 5, the exponential hedge increases the exponential utility by 107% for daily, 111% for weekly and 115% for monthly frequencies. In terms of a comparison there are some differences in the relative performance of the model based hedges however they are not as large as the differentials we find using the variance criterion. Using (HE$_2$), the quadratic criterion, the best performer is the quadratic HR which tends to outperform both the MVHR and exponential HR by around 4% to 6% depending on frequency. For (HE$_3$) the exponential criterion, the exponential hedge outperforms the quadratic hedge by 3% to 4% but outperforms the MVHR hedge by an average of 15% across daily weekly and monthly frequencies.

Taken together, these results show that for equity hedgers there are large performance differences between the minimum variance and utility based approaches. The hedging performance differentials mimic the differences in the hedge ratios themselves with the MVHR and exponential HR showing the largest differences. This means that the quadratic HR which lies between the two may be a good option for hedgers who are looking for an approach that offers good performance across both risk minimising and utility maximising spectrums.

For oil hedgers the results are somewhat different. Broadly speaking, what we find is that the relative performance differentials between the model based hedges are quite small irrespective of which performance metric is used. For example, if we look at daily hedges with CRRA of 5, the MVHR outperforms the quadratic and exponential hedges in terms of variance reduction by just 0.4% and 1.5% respectively in absolute terms and just 0.1% and 0.3% respectively using a CRRA of 10. For weekly and monthly frequencies relative

performance differentials are similarly small. What this means in practice is that for oil hedgers, although the different hedge strategies produce different hedge ratios, the performance differentials are not economically significant and therefore there is little point in adding additional complexity to the hedging decision by using the utility maximising as distinct from the variance minimising approach.

We also find that hedging is not quite as effective for oil as for the S&P500. For example, variance reduction for oil hedgers using the MVHR at the daily frequency is 76.29% as compared with 93.69% for the S&P500 for the equivalent hedge. This is confirmed as we find average effectiveness across all performance measures and hedging frequencies of about 96% for the S&P and 94% for oil. These figures also show the efficacy of hedging across the different hedging models using a broad spectrum of performance criteria.

Turning next to the out-of-sample performance of the different hedging strategies, Table 6 presents mean, variance, quadratic utility and exponential utility values while Table 7 presents hedging performance metrics. The first thing we note is that the out-of-sample hedging performance is good across each hedging model irrespective of the performance criteria used. A quick glance at Table 7 confirms this with improvements in performance in excess of 95% for both assets and in some cases of hedgers more than doubling their utility as compared with a No Hedge position. Looking at the relative performance of the difference HR's, from Table 7, when (HE_1) - the variance is used as the performance benchmark, we observe only very small differences. For example, for S&P500 using weekly data and a CRRA of 5 the difference between the best performing quadratic HR and the worst performing exponential HR is just 0.29%. The results for oil are broadly similar with only small differentials in performance between models irrespective of frequency or risk aversion value.

[TABLE 6 AND 7 HERE]

17

For the utility based performance criteria (HE_2), and (HE_3) some differences emerge when we compare results for the two different risk aversion values and across hedging frequency. Looking first at (HE_2), the quadratic utility for the S&P500, we can see that for the lower risk aversion value CRRA 5, there are some significant performance differences. For example, using daily data, the performance for the MVHR, quadratic and exponential HR's is 95%, 98% and 101% respectively. An interesting result is that for weekly hedgers the performance differences between the model based hedges are quite small irrespective of the performance criteria used, in all cases being less than 1%. At the monthly frequency only the exponential HR differs significantly from both the MVHR and quadratic HR, returning the best performance with an increase in utility of over about 110%.For (HE_3) the exponential utility the results are quite similar as for (HE_2), with large performance differentials opening up between the exponential HR and both the MVHR and quadratic utility HR for daily hedges and especially for monthly hedges. If we make similar comparisons for hedgers with higher risk aversion value CRRA 10 we find that the performance differentials narrow considerably, particularly at weekly and monthly frequencies.

Moving on to oil, the results are quite similar to the equity hedges although performance differences between hedging models are generally smaller than for the S&P500. For daily hedgers with CRRA5, differences are of the order of 1.0 % to 3.0% using (HE_2) and 2.0% to 6.0% using (HE_3), while for monthly and weekly hedges the differences are smaller again. Finally for the higher risk aversion of CRRA10 we again see a convergence in hedging performance across models and for each performance metric. An interesting observation is that the model based hedges don't uniformly perform best out-of-sample in terms of their given optimisation criterion. Using (HE_1) and daily data for oil as an example, the best performer at both risk aversion levels is the quadratic HR, and if we use (HE_2) and (HE_3), for the same data, the best performer is the exponential HR.

In order to illustrate these results in economic terms we also report Value at Risk (VaR) in Table 8 for each of the out-of-sample hedges based on the 95% confidence interval.[20] VaR is

---

[20] VaR has increasingly been used as a performance measure given its wide applicability in risk management and relative ease of understanding as compared with statistical measures.

estimated using the commonly applied Historical Simulation method. A quick glance confirms the overall effectiveness of hedging with each hedging model showing large reductions in excess of 80% in the VaR across different hedging horizons and for both assets. For the S&P500, we can see that VaR figures are broadly similar across each of the hedging models at the daily and weekly frequencies. For example, daily VaR ranges from $4,029 for the best performing MVHR to $4,111 for the worst performer which is the exponential HR for CRRA 5, a difference of just $82 which can be termed insignificant. For the monthly frequency however, there are some significant differences particularly between the MVHR ($5,201) and the quadratic ($7,296) and exponential ($13,786) utility based hedges for the lower CRRA 5. For oil the differences in the VaR from the different hedging models are quite small even at lower frequencies with the quadratic HR for CRRA 5 yielding the best performing monthly VaR of $17,748 which beats the worst performing exponential HR with a VaR of $20,759, a difference of just $3,011.

In terms of which model provides the best performance, the results depend on frequency. For S&P500 the MVHR returns the best VaR at daily and monthly frequencies with quadratic HR performing best for weekly hedgers. For oil the position is reversed with the quadratic HR performing best for daily and weekly hedgers and MVHR the best at the weekly interval. The economic performance of the exponential HR is poorest but is only notably worse for monthly hedgers for the S&P500. Taken together, these results seem to indicate that for S&P500, using performance metrics based on the optimisation criteria will yield some significant performance differentials especially for investors with moderate risk aversion, and for daily and monthly frequencies, while for oil the differences are very slight. Finally when estimating the loss associated with the different hedges using VaR, large performance differences are to be found only for monthly hedgers and especially for the S&P500.

6. CONCLUSION

We estimate and compare optimal utility based hedges using performance metrics based on the underlying optimization criteria. We apply this approach for energy and equity market participants across daily weekly and monthly hedging frequencies. By matching the performance appraisal metrics to the underlying methods used to generate the optimal hedges we can make a comprehensive comparison of the relative performance of utility based hedges and contrast these approaches with the more traditional variance minimising hedge. By focusing on the role of both utility and risk aversion in the hedging decision, we are reflecting real world investor concerns which encompass both risk and expected return.

Our findings indicate that there are significant differences between the quadratic and exponential hedge ratios and also between each of the utility based hedges and the variance minimising approach. The OHR's also differ significantly for each of the different hedging horizons, daily weekly and monthly, which implies that optimal hedge ratio estimation should be tailored to the specific hedging frequency of the individual hedger. We also find that using performance methods based on the optimization criteria will yield significant performance differentials for equity hedger's but not for oil hedgers in a comparison between different utility based hedging models. The differences are more marked for investors with moderate risk aversion and at the monthly hedging frequency.

Taken together these results suggest that a utility based approach may be worthwhile for equity market hedgers whereas for oil market participants hedging performance is not significantly different enough to warrant using a more complicated utility based approach as compared with the simpler minimum variance based hedge.

Bibliography

Arrow, K. (1971). Essays in the theory of risk bearing, Markham, New York.

Alexander, C., 2008. Quantitative Methods in Finance, Wiley, England.

Baron, D. (1977). On the utility theoretic foundations of mean-variance analysis. Journal of Finance, 32, 1683 – 1697.

Brooks, C., O. Henry, and G. Persand, (2002). The Effects of Asymmetries on Optimal Hedge Ratios, Journal of Business 75, 333 – 352.

Brooks, C., Cerny, A., & Miffre, J. (2012). Optimal hedging with higher moments. Journal of Futures Markets, 32, 909-944.

Chen, S., Lee, C., Shrestha, K., (2001). On a mean-generalized semi-variance approach to determining the hedge ratio. Journal of Futures Markets, 21, 581-598.

Cotter, J., Hanly, J. (2006). Re-examining hedging performance. Journal of Futures Markets, 26, 657-676.

Cotter, J. & Hanly, J. (2010). A utility based approach to energy hedging. Energy Economics, 32, 432 - 441.

Demirer, R., & Lien, D. (2003). Downside risk for short and long hedgers. International Review of Economics and Finance, 12, 25 – 44.

deVille deGoyet, C., Dhaene, G., & Sercu. P. (2008). Testing the martingale hypothesis for futures prices: implications for hedgers. Journal of Futures Markets, 28, 1040 – 1065.

Ederington, L. (1979). The hedging performance of the new futures markets. Journal of Finance, 34, 157 – 170.

Eisenhauer, J & Ventura, L. (2003). Survey measures of risk aversion and Prudence. Applied Economics, 35, 1477–1484

Fung, H.G, Jarrett, J., & Leung, W.K. (1990). Optimal hedging with higher moments. Managerial Finance, 16, 19-23.

Ghysels, E. Santa-Clara, P. & Valkanov, R. (2005). There is a risk return trade-off after all. Journal of Financial Economics, 76, 509 - 548.

Guo, H., & Whitelaw, R. (2006). Uncovering the risk–return relation in the stock market. The Journal of Finance, 61, 1433 – 1463.

Harris, R. D. F., & Shen, J. (2006). Hedging and value at risk. Journal of Futures Markets, 26(4), 369–390.

Ho, T. (1984). Intertemporal Commodity Futures Hedging and the Production Decision. Journal of Finance, 39, 351 – 376.

Juhl, T. Kawaller, G. and Koch, P. (2012). The effect of the hedge horizon on optimal hedge size and effectiveness when prices are cointegrated. Journal of Futures Markets. 32, 837 – 876.

Lien, D. (2003). The effect of liquidity constraints on futures hedging. Journal of Futures Markets, 23, 603-613.

Lien, D. (2007). Optimal futures hedging: quadratic versus exponential utility functions. Journal of Futures Markets 28, 208 - 211.

Lien, D. (2012). A note on utility-based futures hedging performance measure. Journal of Futures Markets, 32(1), 92-98.

Mehra, R. & Prescott, E. (1985). The equity premium: A puzzle. Journal of Monetary Economics, 15, 145-161.

Miffre, J. (2004). Conditional OLS minimum variance hedge ratios. Journal of Futures Markets, 24(10), 945–964.

Poon, S. & Granger, C. (2003). Forecasting Volatility in Financial Markets: A Review. Journal of Economic Literature, 41, 478–539.

Sadorsky, P. (2002). Time varying risk premiums in petroleum futures prices. Energy Economics, 24, 539–556.