

Hedging index exchange traded funds

C. Alexander^{a,*}, A. Barbosa^{a,b}

^a *The University of Reading, ICMA Centre, Business School, P.O. Box 242, Reading RG6 6BA, United Kingdom*

^b *Brazilian Securities Commission, Rua Sete de Setembro 111, Rio de Janeiro, 20050-901, Brazil*

Received 21 April 2006; accepted 23 March 2007

Available online 6 July 2007

Abstract

This paper presents an empirical comparison of the out of sample hedging performance from naïve and minimum variance hedge ratios for the four largest US index exchange traded funds (ETFs). Efficient hedging is important to offset long and short positions on market maker's accounts, particularly imbalances in net creation or redemption demands around the time of dividend payments. Our evaluation of out of sample hedging performance includes aversion to negative skewness and excess kurtosis. The results should be of interest to hedge funds employing tax arbitrage or leveraged long–short equity strategies as well as to ETF market makers.

© 2007 Elsevier B.V. All rights reserved.

JEL classification: C32; G10; G15

Keywords: Exchange traded fund; Minimum variance hedge ratio; Exponential utility; Adjusted Sharpe ratio

1. Introduction

An exchange traded fund (ETF) is an instrument for investment in a basket of securities. It is similar to an open-ended fund, but it can be transacted at market price any time during the trading day. ETF market makers publicly quote and transact firm bid and offer prices, making money on the spread, and buy or sell on their own account to counteract temporary imbalances in supply and demand and hence stabilize prices. A basic regulatory requirement for ETFs is that shares can only be created and redeemed at the fund's net asset value (NAV) at the end of the trading day. Of course, the bid-ask spread itself is the normal compensation to any market maker for bearing order imbalances and inventory risk. But increasingly efficient electronic trading platforms have induced a very competitive environment and a downward pressure on bid-ask spreads. As a result market maker's are seeking the best ways of reducing the uncertainty of exposures arising from

order imbalances: if the market risk of inventories can be effectively hedged then the market maker can reduce bid-ask spreads and thereby become more competitive.

Hedging is a particularly important tool in the ETF markets where market makers seek the cheapest ways to reduce the uncertainty of their exposures. Being a basket product, the creation/redemption of ETF shares requires trading on several component stocks, many of them relatively illiquid with high transaction costs. The cost of reducing the market risk of inventory through creations/redemptions is therefore much higher than the cost of hedging with futures. Besides, daily net creation or redemption demands can be huge, especially around the time of dividend payments when the tax treatment of dividends on ETF investments significantly increases trading volume. These demands may be too great for a market maker to close the position out at the end of the day by buying or selling the index component stocks, especially for small cap ETFs. In that case the market maker can either attempt to borrow from or lend to another market maker or he can take large long or short positions onto his own account, until they are offset by an opposite demand or supply of the ETF from investors or until he can close

* Corresponding author. Tel.: +44 1183786431.

E-mail address: c.alexander@icmacentredg.ac.uk (C. Alexander).

the positions in the market. If he takes large positions on his own account the incentive for hedging is increased, and since the demand and supply of different ETFs from investors is heterogeneous he may first consider to what extent a long position on one ETF offsets a short position on another correlated ETF.

The hedging behaviour of market makers is discussed by Froot and Stein (1998), who show how firm-level risk management considerations should be factored into the pricing of those risks that cannot be easily hedged. They claim that hedging tradable risks is the aim of all financial intermediaries and that dealers' behaviour is driven by the unhedgeable risks of inventories. In a similar vein, Naik and Yadav (2003) argue that the decentralized structure of market-making causes trading and pricing decisions to be based on the 'ordinary' inventory of a dealer and emphasize that their empirical findings do not preclude 'macro' hedging at the level of the firm, i.e., the cross-hedging of opposite positions in correlated securities. The fact that correlated exposures may influence dealer's behaviour in a competitive market was also noted many years ago, by Ho and Stoll (1983). Hence, although individual dealers decide their bid-ask spread and inventory levels on an individual basis, cross-hedging of correlated exposures and hedging baskets with single index futures are important risk management functions within the market making firm. Therefore as well as hedging with individual index futures this paper will access the performance of cross-hedging with other ETFs and hedging all ETFs with single index futures.

The ETF hedging problem is not only of interest to market makers. The ability to short, coupled with low transaction costs has fuelled a significant increase in demand for ETFs from hedge funds. Common strategies include tax arbitrage, ETF pairs trading and market timing of a futures hedge on a long ETF position. Hedge funds may apply leverage to hedged ETF portfolios, but if the portfolio returns have a high negative skewness and a significant excess kurtosis, and if these higher moments are ignored when considering a leveraged strategy, the hedge fund runs a small risk of a very large loss. For this reason, in addition to a standard variance reduction criterion, our out of sample hedging results will be analysed using an adjusted information ratio and a utility-based performance criterion, both of which include aversion to negative skewness and positive excess kurtosis.

We shall compare the hedging performance of 'naïve' (1:1) hedging with minimum variance hedging. The returns on US index ETFs are extremely highly correlated with the index futures returns and we therefore expect that it will be difficult to improve upon the 1:1 hedge ratio even over very short term horizons. This would agree with the findings of Alexander and Barbosa (2007b) who demonstrate that no minimum variance futures hedge can improve on the 1:1 hedge in major stock index markets. However econometric minimum variance hedge ratios should be useful for computing the short position on one ETF that provides the best

hedge for a long position on another ETF, because the correlation between different ETFs is generally much lower. We employ a variety of econometric models that have become standard in the hedging literature. In particular we extend the work of many authors that have employed econometric methods to investigate the efficiency of short-term minimum variance futures hedging of stock indices, but not of ETFs.¹ Many previous studies have not used contemporaneous data on the spot index and the futures, thus increasing the apparent effectiveness of minimum variance hedging. They have also measured performance by variance reduction criteria only, thus ignoring the potential for leveraged hedged positions to incur large losses when returns have negative skewness and high positive excess kurtosis. Neither have previous studies employed statistical tests for significant differences in results. None of these concerns arise with the present study.

A considerable body of academic research on index ETFs has examined: arbitrage opportunities between the ETF, index and futures (Switzer et al., 2000; Ackert and Tian, 2000, 2001; Chu and Hsieh, 2002; Kurov and Lasser, 2002); their price characteristics and the reasons for their underperformance relative to the index and index funds (Elton et al., 2002; Kostovetsky, 2003; Gastineau, 2004; Engle and Sarkar, 2006); their role in the price discovery process (Chu et al., 1999); the tax and cost advantages relative to index funds (Poterba and Shoven, 2002; Gastineau, 2002, Chapter 4; Bergstresser and Poterba, 2002); the impact of the NYSE entry into the trading of AMEX-listed ETFs (Boehmer and Boehmer, 2003); and the effect of index ETF trading on the liquidity of the underlying stocks (Hegde and McDermott, 2004). Hence this paper is the first major study of hedging ETFs.²

The important new contributions of this paper are: to examine the mispricing and basis risk of the four major US index ETFs, especially around the time of dividend payments; to analyse the performance of different minimum variance hedge ratios for ETFs; to investigate the extent to which the risk of a long position on one index ETF can be hedged by a short position on another correlated index ETF; and to assess the effectiveness of hedging all ETFs with a single index futures. The remainder of this paper is structured as follows. Section 2 describes the characteristics of the four largest US index ETFs. Section 3 analyses the empirical properties of mispricing and basis

¹ See Hill and Schneeweis (1984), Figlewski (1984, 1985), Junkus and Lee (1985), Peters (1986), Graham and Jennings (1987), Merrick (1988), Lindahl (1991, 1992), Bera et al. (1993), Stoll and Whaley (1993), Benet and Luft (1995), Park and Switzer (1995), Geppert (1995), Lien et al. (2002), Brooks et al. (2002), Miffre (2004), Alexander and Barbosa (2007b) and many others.

² The most closely related work is by Alexander and Barbosa (2005) who show that minimum variance hedging of the Spider with the SP500 index futures performs no better than the 1:1 naïve hedge and that hedging the Spider with SP500 futures is more efficient than hedging the SP500 index with its futures.

risk in ETFs. Section 4 describes the methodology for computing minimum variance hedge ratios and evaluating their performance. Section 5 present our empirical results on the hedging ETFs and Section 6 summarizes and concludes.

2. Index ETFs in the US

We shall examine the risks of trading and market making in four funds that by the end of 2005 together accounted for one-third of the assets invested in US passive ETFs. These are³:

- The ‘Spider’, i.e., the S&P 500 SPDR that was listed on the AMEX in 1993: ticker symbol SPY. It remains by far the largest passive ETF with 56.12 bn\$ under management by 20 September 2006. The Spider share price corresponds to 1/10th of the S&P 500 index value.
- The ‘Cubes’, i.e., the Nasdaq-100 ETF: ticker symbol QQQQ. This is the second largest ETF in the US, launched in March 1999 and by 20 September 2006 having 16.89bn\$ under management. The Cubes share price is approximately 1/40th of the Nasdaq-100 index value.
- The ‘Diamond’, i.e., the ETF tracking the Dow Jones Industrial Average (DJIA) index: ticker symbol DIA. It began trading in January 1998 and by 20 September 2006 had 5.89 bnUS\$ under management. The Diamond share price is approximately 1/100th of the DJIA index value.
- The Russell 2000 iShare: ticker symbol IWM. It was launched in May 2000 and had 10.57 bnUS\$ under management by 20 September 2006. The Russell iShare price corresponds to 1/5th of the Russell iShare index value until 9 June 2005 when it had a 2:1 stock split and now it corresponds to 1/10th of the index.

Market orders for creation and redemption may be placed until 4:00 p.m. EST and at this time the market maker needs to decide whether to create or redeem shares, to lend or borrow shares from other market makers, to keep an open position on their own account, or to hedge their position. If they choose to hedge with the index futures then the hedge would be affected at 4:15 p.m., or just before.

Our study is based on Bloomberg daily data on these four ETFs from May 2000 to September 2006. Fig. 1 shows how the total market value of each fund evolved over the period. The growth in the Spider’s market value represents a significant increase in shares outstanding. The other funds’ market values grew little and the Cubes in particular did not increase in market value at all between January

2001 and September 2006.⁴ Very large daily net creations or redemptions of around 5% of the NAV of the fund are quite normal and it is not uncommon for redemption or creation demand to be over 10% of NAV especially around the time of dividend payments.⁵ The holder of the ETF on the ex-dividend date is entitled to receive the dividends, no matter how long the share has been held. But if the share is sold during the ex-dividend period the registered investor loses the dividends and any tax advantage or disadvantage related to it. Moreover ETFs traded on the secondary market do not necessarily include the dividend or cash components. Hence the very large creation and redemption demands observed arise because there is considerable scope for tax arbitrage around the time of dividend payments.

Table 1 compares the daily average of net creations and redemptions of the total sample with the daily average around the ex-dividend dates. The positive mean in each case is a result of the huge net creation of ETF shares over the period. The standard deviation measures the extent of creation/redemption activity. The middle section of Table 1 shows a marked increase in creation/redemption activity around dividend dates for the Spider and the Cubes and the lower part of Table 1 displays very low correlations between the creation/redemption series of the four funds. We conclude that market makers are likely to face a quite heterogeneous demand for long and short positions in different ETFs, especially around the time of dividend payments in the Spider and the Diamond.

Since quantity correlations are low the market maker is likely to face a high demand for creation (or redemption) of shares in one fund when there is no offsetting redemption (or creation) demand in a correlated fund. At such times the prior cross-hedging of correlated ETFs before future hedging is not relevant and market makers will hedge each fund with its own index futures. Nevertheless Table 5 below will show that the daily returns correlations are high. So when a market making firm does have offsetting positions on correlated ETFs the cross-hedging of these positions prior to affecting a futures hedge is a costless way of reducing uncertainty.

Following Ackert and Tian (2000) we have adjusted each fund’s price by deducting the value of the cash component. The Spider, Cubes and Russell iShare pay quarterly dividends that coincide with the date of the expiration of the futures. Hence there is no dividend uncertainty

⁴ The number of Cubes shares outstanding has increased, but the Nasdaq-100 index fell by 37.5% between May 2000 and January 2001, by 26.5% between January 2001 and December 2004 and had a slight 1.45% increase between December 2004 and September 2006.

⁵ The regular quarterly ex-dividend date for the Spider and the Cubes is the third Friday in each of March, June, September and December. However, from inception until the end of 2004 the Cubes paid dividends only twice, in December 2003 and December 2004. In the years 2005 and 2006, the Cubes paid dividends in 5 occasions. The Diamond has monthly dividend payments and the dividend stream of the Russell iShare, although quarterly, does not coincide with that of the Spider.

³ All four trusts issue and redeem shares in creation units of 50000. The portfolios studied in this paper were based on a block size of ETFs corresponding to one unit of the underlying index in order to match the futures contract trading unit based on the spot value of the index. That is, for each trade unit we hedge 10, 100, 40 and 5 (10 after 09/06/2005) shares of Spider, Diamond, Cubes and iShares, respectively.




Fig. 1. Total market value of ETF trusts. Closing price of the four ETF shares multiplied by the number of outstanding shares in each trust fund at the end of trading day. Measured in Billion USD.

Table 1
Net daily creations and redemptions

	SPY	DIA	QQQQ	IWM
<i>Net daily creations and redemptions (total sample)</i>				
Mean	0.09%	0.08%	0.08%	0.31%
St. dev.	1.48%	2.28%	1.31%	4.21%
<i>Net daily creations and redemptions (around dividend dates only)</i>				
Mean	0.26%	0.05%	0.16%	0.26%
St. dev.	2.30%	2.10%	1.52%	4.22%
<i>Correlations of net daily creations and redemptions</i>				
SPY	1	-0.02072	0.04900	0.12429
DIA		1	0.05884	0.04218
QQQQ			1	-0.03649
IWM				1

Mean, standard deviation and correlations between the net number of shares created or redeemed each day for the four funds. Mean and standard deviations are expressed as a percentage of the number of shares outstanding on the previous day.

included in the arbitrage relation between the fund and the index futures as all dividends, expected and paid, are isolated in the cash account. This is not true for the Diamond as it pays monthly dividends. For this reason, besides the cash component adjustment made to all four funds, we also adjusted the Diamond theoretical futures price for dividends paid before the expiration of the futures contract.

In Table 2 we present the first four sample moments of the ETFs returns for the entire sample and over three separate and quite distinct two-year periods of the US equity market: the bear market from January 2001 until December 2002, the recovery phase from January 2003 until December 2004 and finally the sideways market from January 2005 to September 2006. All four funds performed badly over the first period and this period is by far the most

Table 2
Descriptive statistics of the ETF returns

	SPY	DIA	QQQQ	IWM
<i>2001–2006</i>				
Average annual return	0.11%	1.42%	-6.71%	7.46%
Volatility (annualized)	17.76%	17.37%	32.31%	21.09%
Skewness	0.0621	-0.0293	0.2583	-0.0733
XS kurtosis	2.5512	4.5582	5.0856	0.3947
<i>2001–2002</i>				
Average annual return	-20.68%	-12.87%	-44.84%	-12.07%
Volatility (annualized)	24.67%	24.27%	48.64%	25.79%
Skewness	0.19	0.05	0.35	0.06
XS kurtosis	0.81	2.12	1.46	0.09
<i>2003–2004</i>				
Average annual return	16.09%	12.96%	24.58%	27.20%
Volatility (annualized)	14.21%	13.71%	21.40%	18.69%
Skewness	-0.04	0.05	-0.06	-0.22
XS kurtosis	0.98	1.39	0.74	-0.40
<i>2005–2006</i>				
Average annual return	5.16%	4.26%	0.28%	6.77%
Volatility (annualized)	10.66%	10.41%	14.83%	17.44%
Skewness	-0.10	-0.17	-0.01	-0.08
XS kurtosis	0.37	0.57	0.42	0.26
<i>Around dividends</i>				
Average annual return	-42.25%	-5.75%	22.62%	-0.40%
Volatility (annualized)	19.22%	17.80%	13.55%	20.54%
Skewness	-0.37	-0.3	0.28	0.1
XS kurtosis	2.19	5.45	0.75	0.39

Mean, volatility, skewness and excess kurtosis of the return series of four ETFs from January 2001 to September 2006.

volatile, as it covers the aftermath of the technology bubble and the terrorist attack on the US. The period 2003–2004 was much less volatile, as markets began to recover the losses made between 2000 and 2002. Volatility was even

lower in the period 2005–2006 but returns, whilst positive, decreased significantly. In the three periods the Diamond, being based on Blue Chip stocks, was the least volatile and the Cubes and the Russell iShare the most volatile. This reflects continued uncertainty surrounding performance of technology and small cap stocks. Apart from this, the higher moments indicate the heavy-tailed and slightly skewed nature of the fund's returns distributions.⁶

We have explained above why tax arbitrage activities in ETFs are expected around the time of dividend payments. From the 1555 observations in the total sample, there are about 273 days in the weeks before and after dividend payments for the Spider, 260 days for the Russell iShare, 793 for the Diamond but only 69 days for the Cubes. Tax arbitrage activities are indeed evident from the last section of Table 2, which shows that returns on the ETFs are very significantly lower and more variable around the time of dividends. The exception is the Cubes, where the mean return is positive and it is not significantly different from the overall mean. It is clear that hedging the Spider will be particularly important around dividend dates. Based on the 273 days around the dividends the mean annualized return was –42%, with a volatility of 19% and a very large excess kurtosis, compared with a mean return of 0.11% over the entire sample with a volatility of 18%. Hence the dividend effect on the Spider is very significant. The Diamond and Russell iShare have similar but less significant dividend effects.⁷

Clearly risk exposure in ETF markets is greatest at the time of a dividend payment, so the short-term hedging of index ETFs is of particular interest to market makers and tax arbitrage investors at this time. Thus we shall single out the periods before and after dividend payments on the ETFs and examine whether the efficiency of futures hedge ratios are different at these times, and if so, how they should be adjusted.

3. Mispricing and basis risk

The fair value of a futures contract at any time t prior to the expiry time T is given by:

$$F_{T,t}^* = \exp((r - q)(T - t))S_t, \quad (1)$$

⁶ The standard error is approximately $\sqrt{(6/T)}$ for the skewness and $\sqrt{(24/T)}$ for the excess kurtosis where T is the sample size. In our case, with T approximately equal to 500 in each sub-sample, the approximate standard error for the skewness coefficient is 0.11 and for the excess kurtosis it is approximately 0.22. Note that the excess kurtosis is significantly different from zero (except for the IWM) but that it was at a relatively low level compared with the 1990s. For instance, from 1993 until December 2004 the sample excess kurtosis of the S&P 500 index daily returns was 3.69, having achieved a maximum of 10.91 during September 1998, although over the entire period 2001–2006 the excess kurtosis of the index was only 2.59. Sub-period results are available from the authors.

⁷ The mean annualized return was –5.75% when averaged over all 793 days in the dividend periods for the Diamond but only 1.42% overall. However the overall volatility was about 17.37% and so –5.75% is well within one standard deviation. Similarly for the Russell iShare: the mean annualized return was –0.40% during dividend periods and 7.5% overall, with a volatility of 21%.

where S_t is the spot price at time t , r is the risk-free interest rate and q is the dividend yield of maturity $T - t$. The market price of the futures contract may be expressed in the form:

$$F_{T,t} = F_{T,t}^* + x_{T,t}S_t, \quad (2)$$

where the ‘mispricing’

$$x_{T,t} = \frac{F_{T,t} - F_{T,t}^*}{S_t}, \quad (3)$$

is the difference between the market and fair futures prices, as a proportion of the spot price.⁸ Thus in the basis, $S_t - F_{T,t}$ there are two distinct sources of variability: changes in the ‘fair basis’ (the difference between the spot price and the fair price of the futures contract) and the mispricing variation. Since it depends only on time, discount rates and dividend yields there is relatively little uncertainty about the fair value of the basis. Hence it is the mispricing variation that dominates the basis risk and it is this that needs to be hedged in an optimal manner.

We now consider the mispricing series (3) where S_t is the spot ETF and F_t is the index futures. We use the index futures because futures contracts on ETFs are very recent products (e.g., Spider futures started trading only in June 2005)⁹ and the trading volume on ETF futures remains very much lower than that on index futures contracts. Note that the effect of using an ETF in place of the spot index for index futures arbitrage is to reduce the no-arbitrage range for the market price of index futures about the fair price, and thus reduce the mispricing. When the futures is sold and the spot index is bought, and even more so when hedge portfolio is long the futures and short the index, the trading costs are high.¹⁰ These present a barrier to arbitrage, and the no-arbitrage range will be relatively wide. However if the market has an ETF and this is used for spot-futures

⁸ Note that although many authors refer to (3) as the ‘mispricing’ of the market price of the futures compared with its fair value, it is really the spot rather than the futures that is mispriced because it is the futures contract that serves the dominant price discovery role, even when the market has an actively traded ETF. See Chu et al. (1999).

⁹ Futures contracts on the (Switzer et al., 2000) iShare were launched by Chicago Mercantile Exchange in June 2005. While there were 1355 futures traded on the Cubes, trading of the Nasdaq-100 futures reached 1 772 666 contracts from January to September 2006. For the Spider there were 1705 futures traded on the ETF during the same period, compared with 11 234 324 futures on the S&P 500 and the Russell 2000 iShare futures traded only 667 contracts while the futures on the index traded 472 797 during the same period. The futures contract on the Diamond started trading in November 2002 at OneChicago. From January to September 2006 only 6389 futures contracts were traded while the futures on the Dow Jones traded at Chicago Board of Trade had a volume of 1 427 375 contracts. Source: Chicago Mercantile Exchange for futures contracts on Spider, Cubes and Russell 2000 iShare and corresponding indices, Chicago Board of Trade for futures on Dow Jones Industrial Average and Bloomberg for futures on Diamonds.

¹⁰ Trading in the spot index requires a large amount of capital and trading in some relatively illiquid stocks, but the trading costs for index replication are not reflected in the index point value.

arbitrage in place of an index replicating portfolio, costs are significantly reduced. Furthermore, like futures, ETFs are not held to the up-tick rule so short arbitrage is easier. Consequently the no-arbitrage range for the index futures should be smaller in the presence of an ETF as an arbitrage vehicle. In particular the incidence of negative mispricing, where the market price of the futures is much less than the fair price, will be reduced.

Table 3 reports the sample statistics of each funds' mispricing relative to the index futures. The results are presented for the entire period since there was little variation over the three sub-samples. The extent of the basis risk in each market is captured by the volatility and higher moments of the mispricing series. Note that these statistics essentially capture the results of the naïve hedging strategy, i.e., long S (or its close substitute, F^*) and short F . In particular the volatility of mispricing is linked to the variance of the naïve hedged portfolio, and hence to the Ederington effectiveness measure in Table 6 below. The volatility of mispricing is very much lower than the volatility of the ETF returns, indicating that naïve hedge will be highly effective and may be difficult to improve upon using a minimum variance hedging strategy. But also note that the large excess kurtosis of the mispricing series on the Diamond, the Cubes and the Russell iShare indicates that hedges of these ETFs could fail significantly on some days. We shall consider this issue in more detail in the next section.

The average mispricing of the ETF relative to the index futures contract depends on the handling of dividends and the transactions costs. The Spider has the largest negative mispricing because its dividends are relatively high and its costs are relatively low: in particular the Spider has the lowest turnover, the lowest expense ratio and the smallest transactions costs. Indeed Table 4 shows that the two funds with the most negative mispricing (the Spider and

Table 3

Descriptive statistics of mispricing: Mean, volatility, skewness, excess kurtosis, maximum and minimum of the mispricing series of four ETFs

2001–2006	SPY	DIA	QQQQ	IWM
Mean (daily)	−0.48%	−0.04%	−0.02%	0.12%
Volatility (annualized)	2.68%	3.21%	4.67%	4.03%
Skewness	0.2731	0.0408	−0.4365	0.3095
XS kurtosis	0.4100	2.2479	4.4353	3.9673
Maximum	0.36%	1.17%	1.42%	1.75%
Minimum	−1.17%	−0.96%	−1.59%	−1.49%

The mispricing series is the market futures price minus the theoretical futures price, as percentage of the spot price of the ETF.

Table 4

Dividend yield, expense ratio and turnover: Average dividend yield for the benchmark index between 2000 and 2006, expense ratio and portfolio turnover during 2006

	SPY	DIA	QQQQ	IWM
Benchmark index dividend yield	1.54%	1.98%	0.31%	1.30%
ETF expense ratio	0.10%	0.18%	0.20%	0.20%
ETF portfolio turnover	2.23%	3.88%	6.60%	20.00%

Table 5

Daily returns and mispricing correlations: Correlation between the return series (bold) and between the mispricing series (normal) for each ETF pair

2001–2006	SPY	DIA	QQQQ	IWM
SPY	1	0.958	0.831	0.864
DIA	−0.460	1	0.747	0.809
QQQQ	0.093	0.321	1	0.779
IWM	0.281	0.013	0.209	1

Diamond) have higher dividend yields, lower expense ratios, and lower turnover than the other two funds. So even after the cash account adjustment, these funds are being priced at a premium because the benefits of holding such ETFs outweigh the costs. By contrast the Russell iShare is normally priced at a discount to its index because it has higher trading costs. The trading costs for the Cubes decreased as the dividend stream became more significant at the end of 2004, since which time the mispricing for the Cubes has been near zero.

Table 5 displays for the four ETFs the daily returns correlations above the diagonal and the mispricing correlations below the diagonal. Daily returns were very highly correlated, with the highest correlation between the Diamond and the Spider (as expected since they share many common stocks) and the lowest correlation between the Diamond and the Cubes. Significant mispricing correlations are found between the Diamond and Cubes, the Diamond and iShares and between the Spider and iShares.¹¹ More detailed analysis shows that returns correlations and mispricing correlations were similar in each of the two-year sub-samples, indicating that the relative performance of different hedging strategies may be independent of the market regime. We shall discuss this possibility in Section 5.

4. Hedging methodologies and performance criteria

When hedging with futures it normally makes little difference whether we estimate regression based minimum variance hedge ratios using the spot return or the futures return as the dependent variable. One hedge ratio is simply the other hedge ratio multiplied by the relative variance and, since spot and futures have a relative volatility near to unity, the two estimated hedge ratios are very similar. But if two ETFs have different volatilities then the choice of dependent variable is important. It is straightforward to show that one should take the fund having lower returns volatility as the dependent variable to obtain the hedged portfolio with the smaller variance. Thus the minimum variance hedge ratio for a hedge of duration τ is:

$$\beta_t^*(\tau) = \frac{\sigma_{12,t}(\tau)}{\sigma_{2,t}^2(\tau)}, \quad (4)$$

¹¹ Significance is based on a t -ratio of $r(n-2)^{1/2}(1-r^2)^{-1/2}$ where r is the sample correlation and n is the number of observations. For instance a correlation of 0.2 based on 484 observations has a t -ratio of 4.48, i.e., highly significant.

where $\sigma_{12,t}(\tau)$ denotes the covariance of the τ -period returns and $\sigma_{2,t}^2(\tau)$ denotes the variance of the τ -period returns on the more volatile ETF.

Time-varying minimum variance hedge ratios are obtained using three different econometric models: ordinary least squares (OLS) with a rolling in sample estimation periods of six months,¹² exponentially weighted moving average (EWMA) with a smoothing constant of 0.95, and a multivariate generalised autoregressive conditionally heteroscedastic model (VAR-GARCH).¹³ There is weak evidence of cointegration between the Diamond and the Spider, with Johansen and Juselius (1990) trace and maximal eigenvalue tests being significant at 10%, but no evidence of cointegration was apparent in other fund pairs over the sample period. Hence we do not include an error correction term in the following conditional mean equation for the GARCH hedge ratios on long-short ETF positions:

$$y_t = \mu + \sum_{i=1}^n \Gamma_i y_{t-i} + \varepsilon_t,$$

where y_t is the vector of τ -period log returns on the two funds, ε_t is the vector of unexpected returns and μ and Γ are constants, with

$$y_t = \begin{pmatrix} r_{1,t}(\tau) \\ r_{2,t}(\tau) \end{pmatrix}, \quad \mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \quad \Gamma_i = \begin{pmatrix} \Gamma_{i,11} & \Gamma_{i,12} \\ \Gamma_{i,21} & \Gamma_{i,22} \end{pmatrix}$$

and $\varepsilon_t = \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix}$.

When hedging with futures, however, the cointegration implies that the conditional mean equation should include the carry cost C (as in Ghosh, 1993), with $\pi = (\pi_1 \pi_2)'$ as a vector of constants and the equation becomes an error correction model¹⁴:

$$y_t = \mu + \sum_{i=1}^n \Gamma_i y_{t-i} + \pi C_{t-n} + \varepsilon_t.$$

¹² We also used one-year estimation period for the OLS hedge ratio and the results were very similar to the 6-months OLS hedge ratios, the latter performing slightly better on specific occasions. The difference however is not statistically significant. We report only the 6-month results so as to cover the longest period in our results. Results for the 1-year OLS hedged portfolio are available from the authors on request.

¹³ The EWMA methodology is frequently considered amongst alternative minimum variance hedge ratios, due to its simplicity. Moreover Laws and Thompson (2005) find that EWMA presents a better overall performance than more sophisticated econometric hedge ratios and Brooks et al. (2002) also find that EWMA provides a useful and simple alternative to GARCH hedge ratios.

¹⁴ To see why, take logarithms of (1) giving: $\ln F_t^* - \ln S_t = (r - q)(T - t)$. Hence if the carry cost, $C_t = (r - q)(T - t)$ is stationary the logarithm of the spot price and the logarithm of the fair value of the futures price should be cointegrated with cointegrating vector (1, -1). However the carry cost need not be the most stationary linear combination of the log of the market price of the futures and the log of the spot price. Nevertheless since the mispricing of the futures relative to its fair value is so small it is reasonable to assume the error correction term in the error correction model is equal to the carry cost.

We model time-varying minimum variance hedge ratios based on a conditional bivariate GARCH of the unexpected returns, exploring a variety of GARCH(1, 1) parameterisations such as the scalar and diagonal BEKK specifications (Engle and Kroner, 1995), the t-BEKK, and the DCC (Engle, 2002). However results for all these models were very similar and for reasons of space we report only the diagonal BEKK results.

Each day we estimate the minimum variance hedge ratio to determine the position to be taken at the end of the day until the following day. The sample is then rolled one day, the hedge ratios re-estimated, and the hedge re-balanced and held until the end of the next day. We thus form an ‘out of sample’ hedge portfolio returns series. Since the minimum variance criterion is applied in sample and the hedging performance is tested out of sample there is no guarantee that minimum variance hedging will produce more effective hedges than a simple 1:1 hedge.

Regarding the hedging performance, Ederington (1979) applies the proportional reduction in variance as a measure of hedging effectiveness (E). This is widely used even though it is known to favour the OLS hedge (see Lien, 2005). Also it takes no account of the effect of variance reduction on skewness and kurtosis. Minimum variance hedged portfolios are designed to have very low returns volatility but a high kurtosis indicates that the hedge can be spectacularly wrong on just a few days and a negative skewness indicates that it would be losing rather than making money. Therefore, following Scott and Horvath (1980), Cremers et al. (2004), Harvey et al. (2004), Patton (2004) and others, our second measure of hedge effectiveness accounts for skewness and kurtosis in out of sample performance of hedged portfolios. We shall compute the certainty equivalent (CE) derived from an exponential utility for the hedger, based on the out of sample portfolio returns.¹⁵

The exponential utility function is:

$$U(x) = -\lambda \exp(-x/\lambda), \tag{5}$$

where x is wealth and λ is the coefficient of risk tolerance, which defines the curvature of the utility function and which is measured in the same units as wealth. The CE is that level of wealth such that $U(x) = E[U(x)]$ where $E[U(x)]$ is the expected utility associated with a profit and loss distribution. Applying the expectation operator to a Taylor expansion of $U(x)$ about $U(\mu)$, the utility associated with the mean return provides a simple approximation for the CE associated with any utility function:

$$E[U(x)] = U(\mu) + U'(x)|_{x=\mu} E[x - \mu] + \frac{1}{2} U''(x)|_{x=\mu} E[(x - \mu)^2] + \frac{1}{3!} U'''(x)|_{x=\mu} E[(x - \mu)^3] + \dots,$$

¹⁵ It is natural to base utility on an investor’s level of wealth although it is more intuitive empirically to use the moments of portfolio returns in the CE, as for instance in Harvey et al. (2004).

with the exponential utility, setting $x = \text{CE}$ the above gives an approximation:

$$\exp(-\text{CE}/\lambda) \approx \exp(-\mu/\lambda) \left(1 + \frac{E[(x-\mu)^2]}{2\lambda^2} - \frac{E[(x-\mu)^3]}{6\lambda^3} + \frac{E[(x-\mu)^4]}{24\lambda^4} \right).$$

Thus the certainty equivalent associated with the exponential utility function is approximated as:

$$\text{CE} \approx \mu - \frac{\sigma^2}{2\lambda} + \frac{\varphi}{6\lambda^2} - \frac{\kappa}{24\lambda^3}, \quad (6)$$

where μ and σ are the mean and the standard deviation of x , $\varphi = E[(x-\mu)^3]$ and $\kappa = E[(x-\mu)^4]$. The formulation (6) shows that when the parameter $\lambda > 0$ there is an aversion to risk associated with increasing variance, negative skewness and increasing kurtosis.

In order to capture higher moment effects we have chosen to calculate CE based on the sample moments of the out of sample daily returns using and $\lambda = 10\%$.¹⁶ This represents an average level of risk aversion and our results are qualitatively robust to small changes in this value. Note that if λ is much greater than 10% then the aversion to variance and higher moments would be inconsequential, since (6) would be dominated by the mean. On the other hand a much lower value of λ would emphasise the skewness and kurtosis of the distribution, rather than the mean and variance.

We shall also consider the position of investors seeking to leverage returns and hence also report the adjusted information ratio:

$$\text{AIR} = \text{IR} + \left(\frac{\hat{\tau}}{6}\right)\text{IR}^2 + \left(\frac{\hat{\kappa}}{24}\right)\text{IR}^3. \quad (7)$$

Here IR denotes the ordinary information ratio (the ratio of the annualized mean return to the volatility of the return) and $\hat{\tau}$ and $\hat{\kappa}$ denote the sample skewness and excess kurtosis respectively.

Finally, we test for significant differences between the returns generated by different hedging strategies using the Kolmogorov–Smirnov distance metric (see Siegel, 1956), i.e.,

$$\text{KS} = \sup_x |F_1(x) - F_2(x)|, \quad (8)$$

where $F_1(x)$ and $F_2(x)$ are the empirical distribution functions of the returns on the two hedged portfolios.

5. Empirical analysis of hedging performance

In this section we use the minimum variance hedge ratios described above, compare their performance accord-

ing to the three different criteria with that of the naïve hedge. In the tables below the best performing hedge according to each criterion will be highlighted in bold type.

Table 6 presents the out of sample performance measures for each ETF hedged with its own index futures. Note that in all cases the Ederington effectiveness criterion is marginally higher for the naïve 1:1 hedge than for any minimum variance hedge. However except for the Russell iShare there is a pronounced negative skewness and a very highly significant positive excess kurtosis in the hedged portfolio returns, which is particularly pronounced when the naïve hedge is applied. For this reason the CE and AIR criteria indicate a preference for minimum variance hedging. However, the difference between the hedged portfolio returns distributions is very small and the Kolmogorov–Smirnov tests indicate that none of the returns distributions differ significantly. We have also analysed the performance of futures hedging over each sub-sample, finding that the hedging performance of the different strategies varies little across different market regimes.¹⁷

Isolating the weeks before and after the time that a dividend is paid on an ETF allows us to investigate whether the characteristics of basis risk and hedging efficiency are different at these times. Not surprisingly we find that all hedges become noticeably less efficient at reducing variance around the time of dividend payments. Again it is the minimum variance hedge ratios that provide returns with smaller negative skewness and excess kurtosis, thus being more suitable than a naïve hedge for leverage tax arbitrage strategies. However it is not possible to discern which econometric model provides the best minimum variance hedge, because Kolmogorov–Smirnov tests indicate no significant difference between the different hedge portfolio returns, provided they are obtained from minimum variance hedging.¹⁸

Recall that Table 5 revealed several significant and positive mispricing correlations. This raises the possibility of macro hedging at the level of the firm, a possibility that is considered in Froot and Stein (1998), Naik and Yadav (2003). We therefore considered whether the basis risk from a long position on one ETF be effectively offset by a short position on another ETF. Our results show that a minimum variance hedged positions in two ETFs provides a greater variance reduction than a simple matched long-short position in the two ETFs. Moreover it reduces the skewness and excess kurtosis more effectively than any index futures hedges. The CE criterion also favours a minimum variance hedge ratio, but the AIR is higher when matching a long position on one ETF with a short position on another ETF because the mean return is often significantly greater than zero.¹⁹

¹⁶ Since utilities are only unique up to positive affine transformations it is admissible to apply a linear transformation to the result provided the transformation is the same for all series that are being compared, and we have done this merely to present the CE figures on an intuitive scale.

¹⁷ Results are shown in Alexander and Barbosa (2007a).

¹⁸ Results for the hedging performance around dividends are shown in Alexander and Barbosa (2007a).

¹⁹ Results for cross-hedging ETFs with other ETFs are shown in Alexander and Barbosa (2007a).

Table 6
Out of sample futures hedging performance

		Mean return (%)	Volatility (%)	Skewness	Kurtosis	E (%)	CE	AIR
SPY	Un-hedged	0.11	17.76	0.0621	2.5512		−2567.87	0.0059
	Naïve	0.82	2.08	−1.6962	15.4399	98.63	56.14	0.3871
	OLS	0.87	2.09	−1.6464	15.6387	98.61	61.53	0.4162
	EWMA	0.84	2.11	−1.4638	15.2497	98.58	58.35	0.4001
	GARCH	0.85	2.12	−1.5743	15.2492	98.58	58.77	0.4000
DIA	Un-hedged	1.42	17.37	−0.0293	4.5582		−3121.19	0.0817
	Naïve	0.46	1.59	−1.6485	26.0777	99.17	31.16	0.2900
	OLS	0.49	1.60	−1.7116	27.1943	99.15	33.86	0.3090
	EWMA	0.51	1.62	−1.6643	26.6009	99.13	35.98	0.3223
	GARCH	0.48	1.60	−1.8069	27.5421	99.15	33.13	0.3034
QQQQ	Un-hedged	−6.71	32.31	0.2583	5.0856		−27537.85	−0.2078
	Naïve	−0.10	3.28	−0.6806	9.1805	98.97	−72.84	−0.0319
	OLS	−0.01	3.30	−0.6798	9.1622	98.96	−64.34	−0.0043
	EWMA	−0.11	3.33	−0.6325	8.8470	98.94	−74.62	−0.0324
	GARCH	−0.03	3.31	−0.6614	9.0439	98.95	−66.12	−0.0085
IWM	Un-hedged	7.46	21.09	−0.0733	0.3947		−1916.28	0.3531
	Naïve	0.23	4.18	−0.2131	7.0331	96.07	−76.15	0.0545
	OLS	0.18	4.19	−0.3366	6.5391	96.05	−81.97	0.0439
	EWMA	−0.37	4.23	−0.3883	6.2644	95.97	−139.93	−0.0883
	GARCH	0.14	4.21	−0.2920	6.6247	96.02	−86.68	0.0335

Ederington (E), certainty equivalent (CE) and adjusted information ratio (AIR) performance measures, annual average returns, volatility, skewness and kurtosis of the four ETFs: un-hedged; with a naïve hedge (1:1) composed of one long position on the ETF and one short position on the futures; and when hedged using the estimated OLS, EWMA and GARCH hedge ratios, the last being based on the ECM-BEKK model. The Ederington performance measure is the percentage variance reduction achieved by the hedge. The certainty equivalent is based on the return series and an exponential utility function, to account for investor's aversion to negative skewness and positive excess kurtosis. The adjusted information ratio is the usual information ratio adjusted to account for skewness and excess kurtosis of the return series of the hedged portfolios.

Another possibility for macro hedging is to hedge correlated positions taken by the firm using only a few of the individual futures. To investigate the extent to which this

is possible with ETFs, Table 7 shows the performance of hedging the Diamond, Cubes and Russell iShare with the S&P 500 futures only. As expected the best performance

Table 7
Out of sample hedging performance when hedging with S&P 500 futures only

	Mean return (%)	Volatility (%)	Skewness	XS kurtosis	E (%)	CE	AIR
<i>Naïve</i>							
DIA-SPX	2.13	5.16	−0.0194	6.5807	91.19	60.11	0.43
QQQQ-SPX	−6.00	20.20	0.3555	8.4160	60.91	−7991.47	−0.30
IWM-SPX	8.17	10.73	−0.1965	0.5196	74.11	172.71	0.75
<i>OLS</i>							
DIA-SPX	1.53	5.08	−0.4323	7.7019	91.46	−6.31	0.30
QQQQ-SPX	−0.06	17.07	0.0760	6.9672	72.10	−3861.32	−0.0033
IWM-SPX	6.68	10.13	−0.1977	1.0330	76.90	75.19	0.66
<i>EWMA</i>							
DIA-SPX	1.07	5.00	−0.1985	5.9289	91.71	−37.48	0.22
QQQQ-SPX	1.29	16.88	0.1105	6.0404	72.70	−3251.40	0.08
IWM-SPX	5.08	10.05	−0.1909	1.2050	77.27	−80.53	0.50
<i>GARCH</i>							
DIA-SPX	1.64	5.36	0.0144	5.1553	90.47	2.60	0.31
QQQQ-SPX	0.99	17.04	0.1012	6.4219	72.18	−3527.62	0.06
IWM-SPX	6.18	10.24	−0.2247	1.2850	76.44	−5.03	0.60

Performance measures of the out of sample returns when each ETF is hedged with the S&P 500 futures. Each ETFs is hedged with one short position on the S&P 500 futures and using the estimated OLS, EWMA and GARCH hedge ratios, the last being based on the ECM-BEKK model. The Ederington performance measure is the percentage variance reduction achieved by the hedge. The certainty equivalent is based on the return series and an exponential utility function, to account for investor's aversion to negative skewness and positive excess kurtosis. The adjusted information ratio is the usual information ratio adjusted to account for skewness and excess kurtosis of the return series of the hedged portfolios.

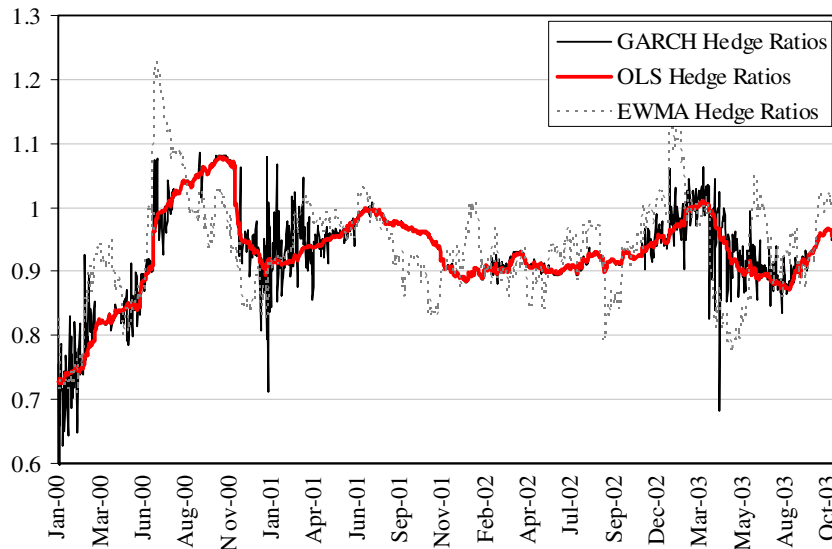


Fig. 2. Time series of minimum variance hedge ratios estimated daily on a rolling sample for the Diamond when hedged with the S&P 500 futures contract.

is achieved on hedging the Diamonds. For all three ETFs the EWMA minimum variance hedge performs best according to the Ederington measure, and Fig. 2 provides some intuition behind this result. The EWMA hedge ratio reflects short term risks better than the OLS hedge ratio, and it is more stable than the GARCH hedge ratio which is prone to periods of instability. Notice however that when higher moments are taken into account the naïve hedge beat the minimum variance hedge for both the Diamond and the Russell iShare.

6. Summary and conclusions

This paper has investigated the effectiveness of minimum variance hedging for the four largest US index ETFs using three different performance criteria, including aversion to negative skewness and excess kurtosis as well as effective reduction in variance, over a very long out of sample period. The ETF hedging problem is of considerable practical interest to market makers, since hedging is the most cost effective way of reducing the market risk of inventories, thus hedging enables market makers to reduce bid-ask spreads in a competitive environment. It is also an interesting problem for hedge funds employing tax arbitrage, or leveraged long-short equity strategies that may ignore the higher moment properties of hedge portfolio returns.

First we examined the characteristics of their mispricing relative to the index futures, showing that the average mispricing could be related to the handling of dividends and transactions costs. The Spider and the Diamond are priced at a premium to the index futures because their dividends are relatively high and their costs are relatively low, compared with the Russell iShare which is priced at a discount to its index futures because it has higher trading costs. The uncertainty in mispricing represents the basis risk that may

be hedged with index futures, the ETF futures being too new and illiquid for us to consider. We first compared the performance of a naïve 1:1 index futures hedge with OLS, EWMA and GARCH minimum variance hedge ratios. There was no evidence that minimum variance hedging could improve on the naïve hedge. However when the large negative skewness and positive excess kurtosis of the hedged portfolio returns were accounted for in the performance criterion there was a clear preference for minimum variance hedging. We also found that any type of hedging with futures is less efficient around the time of dividend payments and that at these times there is a significant amount of tax arbitrage activity, particularly in the Spider.

Further results concerned the extent to which opposite positions on correlated ETFs can reduce risks prior to hedging, and the effectiveness of the S&P 500 futures for hedging the Diamond, Cubes and Russell iShare. In the cross-hedging study we found much less skewness and excess kurtosis in the hedge portfolio returns than when hedging with futures. Also, minimum variance cross-hedging significantly outperformed the naïve matching of long and short positions on different ETFs. Minimum variance hedging of any other ETF with the S&P 500 futures was also generally preferred to a naïve hedge. However we found no single econometric model for minimum variance hedging that performed best according to all of our criteria.

Our results considered daily hedging over a long out of sample period from January 2001 until September 2006, during which three distinct regimes were evident in the US equity markets. Since there was very little difference in the performance of each hedging strategy over the three sub-periods corresponding to each regime our conclusions appear to be independent of the market regime.

Acknowledgement

We would like to thank one of the two anonymous referees for many useful comments which considerably improved the presentation of this paper.

References

- Ackert, L.F., Tian, Y.S., 2000. Arbitrage and valuation in the market for Standard and Poor's depository receipts. *Financial Management* 29, 71–88.
- Ackert, L.F., Tian, Y.S., 2001. Efficiency in index options markets and trading in stock baskets. *Journal of Banking and Finance* 25, 1607–1634.
- Alexander, C., Barbosa, A., 2005. The Spider in the hedge. *Review of Futures Markets* 11 (1), 93–116.
- Alexander, C., Barbosa, A., 2007a. Hedging and cross hedging ETFs. ICMA Centre Discussion Papers in Finance DP2007-01. Available from: <www.icmacentre.rdg.ac.uk>.
- Alexander, C., Barbosa, A., 2007b. Effectiveness of minimum variance hedging. *Journal of Portfolio Management* 33 (2), 46–59.
- Benet, B.A., Luft, C.F., 1995. Hedge performance of SPX index options and S&P 500 futures. *Journal of Futures Markets* 15, 691–717.
- Bera, A., Bunnys, E.L., Park, H.Y., 1993. ARCH effects and efficient estimation of hedge ratios for stock index futures. In: Chance, D.M., R Trippi, R. (Eds.). In: *Advances in Futures and Options Research*, vol.6. JAI Press Inc., pp. 313–328.
- Bergstresser, D., Poterba, J.M., 2002. Do after-tax returns affect mutual fund inflows? *Journal of Financial Economics* 63, 381–414.
- Boehmer, B., Boehmer, E., 2003. Trading your neighbor's ETFs: Competition or fragmentation? *Journal of Banking and Finance* 27, 1667–1703.
- Brooks, C., Henry, O.T., Persaud, G., 2002. The effect of asymmetries on optimal hedge ratios. *Journal of Business* 75, 333–352.
- Chu, Q.C., Hsieh, W.G., 2002. Pricing efficiency of the S&P 500 index market: Evidence from the Standard & Poor's depository receipts. *Journal of Futures Markets* 22, 877–900.
- Chu, Q.C., Hsieh, W.G., Tse, Y., 1999. Price discovery on the S&P 500 index markets: An analysis of spot index, index futures, and SPDRS. *International Review of Financial Analysis* 8, 21–34.
- Cremers, J.-H., Kritzman, M., Page, S., 2004. Portfolio Formation with Higher Moments and Plausible Utility. *Revere Street Working Paper Series, Financial Economics* 272-12.
- Ederington, L.H., 1979. The hedging performance of the new futures markets. *Journal of Finance* 34, 157–170.
- Elton, E., Gruber, M.J., Comer, G., Li, K., 2002. Spiders – where are the bugs? *Journal of Business* 75, 453–472.
- Engle, R., 2002. Dynamic conditional correlation – A simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* 20, 339–350.
- Engle, R.F., Kroner, K., 1995. Multivariate simultaneous generalised ARCH. *Econometric Theory* 11, 122–150.
- Engle, R., Sarkar, D., 2006. Premiums–discounts and exchange traded funds. *Journal of Derivatives* 13 (4), 27–45.
- Figlewski, S., 1984. Hedging performance and basis risk in stock index futures. *The Journal of Finance* 39, 657–669.
- Figlewski, S., 1985. Hedging with stock index futures: Theory and application in a new market. *Journal of Futures Markets* 5, 183–199.
- Froot, K., Stein, J., 1998. Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics* 47, 55–82.
- Gastineau, 2002. *The Exchange Traded Funds Manual*. John Wiley and Sons, New York.
- Gastineau, G., 2004. The benchmark index ETF performance problem. *Journal of Portfolio Management* 30, 96–103.
- Geppert, J.M., 1995. A statistical model for the relationship between futures contract hedging effectiveness and investment horizon length. *Journal of Futures Markets* 15, 507–536.
- Ghosh, A., 1993. Cointegration and error correction models: Intertemporal causality between index and futures prices. *Journal of Futures Markets* 13, 193–198.
- Graham, D., Jennings, R., 1987. Systematic risk, dividend yield and the hedging performance of stock index futures. *Journal of Futures Markets* 7, 1–13.
- Harvey, C., Liechty, J., Liechty, M. Muller, P., 2004. Portfolio selection with higher moments, Working Paper 70, Fuqua School of Business, Duke University.
- Hegde, S.P., McDermott, J.B., 2004. The market liquidity of Diamond, Qubes and their underlying stocks. *Journal of Banking and Finance* 28, 1043–1067.
- Hill, J.M., Schneeweis, T., 1984. Reducing volatility with financial futures. *Financial Analysts Journal* 40, 34–40.
- Ho, T., Stoll, H., 1983. The dynamics of dealer markets under competition. *Journal of Finance* 38, 1053–1074.
- Johansen, S., Juselius, K., 1990. Maximum likelihood estimation and inference on cointegration – with applications to the demand for money. *Oxford Bulletin of Economics and Statistics* 52, 169–210.
- Junkus, J.C., Lee, C.F., 1985. Use of three stock index futures in hedging decisions. *Journal of Futures Markets* 5, 201–222.
- Kostovetsky, L., 2003. Index mutual funds and exchange traded funds: A comparison of two methods of passive investment. *Journal of Portfolio Management* 29, 80–92.
- Kurov, A., Lasser, D., 2002. The effect of the introduction of Cubes on the Nasdaq-100 index spot-futures pricing relationship. *The Journal of Futures Markets* 22, 197–218.
- Laws, J., Thompson, J., 2005. Hedging effectiveness of stock index futures. *European Journal of Operational Research* 163, 177–191.
- Lien, D.D., 2005. The use and abuse of the hedging effectiveness measure. *International Review of Financial Analysis* 14, 277–282.
- Lien, D.D., Tse, Y.K., Tsui, A.K.C., 2002. Evaluating the hedging performance of the constant correlation GARCH model. *Applied Financial Economics* 12, 791–798.
- Lindahl, M., 1991. Risk-return hedging effectiveness measures for stock index futures. *Journal of Futures Markets* 11, 399–409.
- Lindahl, M., 1992. Minimum variance hedge ratios for stock index futures: Duration and expiration effects. *Journal of Futures Markets* 12, 33–53.
- Merrick, J.J., 1988. Hedging with mispriced futures. *Journal of Financial and Quantitative Analysis* 23, 451–464.
- Miffre, J., 2004. Conditional OLS minimum variance hedge ratios. *Journal of Futures Markets* 24, 945–964.
- Naik, N., Yadav, P., 2003. Do dealer firms manage inventory on a stock-by-stock or a portfolio basis? *Journal of Financial Economics* 69, 325–353.
- Park, T.H., 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.
- Patton, A., 2004. On the out of sample importance of skewness and asymmetric dependence for asset allocation. *Journal of Financial Econometrics* 2, 130–168.
- Peters, E.E., 1986. Hedged equity portfolios: Components of risk and return. In: Fazio, F.J. (Ed.). In: *Advances in Futures and Options Research A Research Annual Futures*. vch, pp. 75–91.
- Poterba, J., Shoven, J.B., 2002. Exchange traded funds: A new investment option for taxable investors. *American Economic Review* 92, 422–427.
- Scott, R., Horvath, P., 1980. On the direction of preference for moments of higher order than the variance. *The Journal of Finance* 35, 915–919.

- Siegel, S., 1956. *Nonparametric Statistics for the Behavioural Sciences*. McGraw-Hill, New York.
- Stoll, H.R., Whaley, R.E., 1993. *Futures and Options: Theory and Applications*. South-Western Publishing Co.
- Switzer, L.N., Varson, P.L., Zghidi, S., 2000. Standard & Poor's depository receipts and the performance of the S&P 500 index futures market. *Journal of Futures Markets* 20, 705–716.