

Working Paper

Optimal Investments in Financial
Instruments with Heavily Skewed
Return Distributions: The Case of
Variance Swaps

Reinhold Hafner
Martin Wallmeier

Juni 2007

risklab germany GmbH
Nymphenburger Straße 112-116
80636 München

Optimal Investments in Financial Instruments with Heavily Skewed Return Distributions: The Case of Variance Swaps

Reinhold Hafner*, Martin Wallmeier†

June 2007

Abstract

Mean returns of DAX and ESX variance swaps over the time period from 1995 to 2004 were strongly negative, and only part of the negative premium can be explained by the negative correlation of variance swap returns with stock market indices. We analyze the implications of this observation for optimal portfolio compositions. We are particularly interested in the tradeoff between Sharpe ratio, skewness and kurtosis, because this tradeoff is important for many hedge fund strategies that produce heavily skewed return distributions.

Mean-variance efficient portfolios are characterized by sizable short positions in variance swaps. Typically, the stock index is also sold short to make use of the negative return correlation and achieve a better portfolio diversification. The results are similar if we take the perspective of an investor who maximizes expected power utility. However, just one risk aversion parameter cannot capture heterogeneous preferences for higher moments. To be able to explicitly control such preferences, we use a variant of the Polynomial Goal Programming method, which attempts to find an optimal combination of the degrees of achievement of different objectives. We assume that investors strive for a high Sharpe ratio, high skewness and low kurtosis. These objectives are highly conflicting; in particular, the short positions required to profit from negative variance swap returns inevitably lead to an undesired negative skewness. Our analysis shows that a balanced tradeoff between Sharpe ratio and skewness often does not exist. Investors tend to the extreme portfolios (Sharpe ratio driven, skewness driven or kurtosis driven) and avoid to be stuck in the middle. This ‘all-or-nothing’ characteristic is reflected in jumps of asset weights when certain thresholds of preference parameters are crossed. The findings make clear why many investors are so reluctant to implement option-based short-selling strategies and why a sophisticated risk management for such strategies is indispensable.

JEL classification: G10; G12; G13

Keywords: Variance Swap; Volatility Risk Premium; Portfolio Analysis; Higher Moments; Polynomial Goal Programming

*risklab germany GmbH, Nymphenburgerstr. 112-116, D-80636 Munich, reinhold.hafner@risklab.de

†Chair of Finance and Accounting, University of Fribourg, Bd. de Pérolles 90, 1700 Fribourg / Switzerland, martin.wallmeier@unifr.ch (corresponding author).

1 Introduction

Apart from its prominent role as financial risk measure, volatility has now become established as an asset class of its own. Initially, the main motive for trading volatility was to manage the risk in option positions and to control the vega exposure independently of the position's delta and gamma. With the growth of the volatility trading segment, other market participants became aware of volatility as an investment vehicle. They were attracted by the strong negative correlation between volatility movements and stock index returns. This negative correlation seems to be particularly pronounced in stock market downturns, offering protection against stock market losses when it is needed most. In practice, the 90/10 rule of investing 10% of a stock portfolio into volatility is often suggested. Empirical studies, however, indicate that this kind of crash protection might be expensive. The volatility risk premium is found to be strongly negative at the US as well as the European stock market (see, e.g., Jackwerth und Rubinstein (1996), Chernov und Ghysels (2000), Coval und Shumway (2001), Pan (2002), Bakshi und Kapadia (2003), Eraker et al. (2003), Driessen und Maenhout (2003), Doran und Ronn (2004b), Doran und Ronn (2004a), Bondarenko (2004), Moise (2004), Santa-Clara und Yan (2004), Carr und Wu (2005), Hafner und Wallmeier (2008)). There is some evidence that its magnitude depends on the level of implied volatility and the time to maturity (see Bliss und Panigirtzoglou (2004)), the buying pressure for index puts (see Bollen und Whaley (2004)) and the uncertainty of forward volatility (see Carr und Wu (2005)). The estimated premium shows significant temporal dependencies (see Bollerslev et al. (2005) and Santa-Clara und Yan (2004)). In all, the negative premium appears to be too large to be explained by the negative covariance with market returns within standard equilibrium models (see Carr und Wu (2005), Hafner und Wallmeier (2008)). Therefore, selling volatility on a regular basis might be a profitable strategy. This is compatible with the finding of Bondarenko (2004) that the variance risk factor accounts for a considerable portion of hedge fund historical returns.

Although many studies on the variance risk premium have been published, the implications for investors are rarely addressed in any detail. To draw conclusions from the negative risk premium from an investor's point of view, it is necessary to specifically analyze optimal portfolio compositions in realistic settings. This is the focus of our paper. We are particularly interested in the tradeoff between Sharpe ratio, skewness and kurtosis, because this tradeoff is important for many hedge fund strategies that produce heavily skewed return distributions. Such strategies are characterized by small profits in 'normal' situations and potentially large losses when some rare event occurs.

In the next section, we introduce variance swaps as an investment instrument. We synthetically derive values of variance swaps over the time period from 1995 to 2004 using tick-by-tick data for options on the German stock index DAX and the European index Euro STOXX 50 (ESX). The historical variance swap returns serve as a basis for our portfolio analysis (Section 3). The mean-variance framework (Subsection 3.1) neglects the return properties of negative skewness and excess kurtosis. These are taken into account implicitly in expected utility maximization (Subsection 3.2) and explicitly in the *Polynomial Goal Programming (PGP)* optimization (Subsection 3.3). The last section addresses implementation aspects, especially the concern that a selling strategy might involve a high conditional value-at-risk because large losses will ac-

cumulate if a large unexpected increase in variance occurs. The paper concludes with a brief summary.

2 Variance Swap Returns

2.1 Variance Swaps and Valuation

Variance swaps (also called variance contracts) offer pure exposure to realized future variance. At expiration, the swap buyer receives a payoff equal to the difference between the annualized variance of log stock returns and the swap rate fixed at the outset of the contract. The swap rate is chosen such that the contract has zero present value. Thus, it can be interpreted as the risk neutral expectation of unconditional future variance. Variance swaps on the most common stock indices now enjoy an active over-the-counter market. This was made possible by theoretical work designing a robust replication strategy. It consists of a continuously adjusted stock holding and a static options portfolio including long positions in out-of-the-money (OTM) options for all strikes from zero to infinity (see Neuberger (1994)). Based on this replication, we can express the fair swap delivery price K_{VARS} as:

$$K_{VARS} = \frac{2}{T} e^{rT} \int_0^{F_0(T)} \frac{1}{K^2} P_0(K, T) dK + \int_{F_0(T)}^{\infty} \frac{1}{K^2} C_0(K, T) dK, \quad (1)$$

where $C_0(K, T)$ and $P_0(K, T)$ denote the current market price of a put and a call option of strike K and maturity T , r is the risk-free rate and $F_0(T)$ is the stock's T -maturity forward price.

In a perfect market, the replication is exact if options with arbitrary strikes are available and the stock price process is continuous. If stock price jumps occur and the number of strike prices is limited, the formula still offers a good approximation in realistic settings (see Carr and Wu (2005)). Our comparison with OTC quotes from two major investment banks in November and December 2004 confirms that the theoretical values from (1) are close to market prices. With the exception of two days, the theoretical values always fall into the spread between bid and ask quotes. This spread seems sufficiently narrow (about 1 volatility point) to make the comparison meaningful.

To apply formula (1) in the empirical part, we first estimate the strike price structure of implied volatilities ('smile'). The strike-dependent implied volatilities are then put into the Black-Scholes formula (or equivalently the Black (1976) model, if forward prices for the underlying asset are used) to obtain the strike price structure of option prices $C_0(K, T)$ and $P_0(K, T)$.

2.2 Data and Estimation

Our database contains all reported transactions of the two stock index options and futures with the highest trading volume in Europe. Option and future data for the German stock index DAX are available from January 1995 to December 2004, those on the Euro STOXX 50 index (ESX)

from January 2000 to December 2004. Both European style options ('ODAX' and 'OESX') are traded at the joint German and Swiss options and futures exchange, Eurex.¹

To calculate an implied volatility for each transaction, it is crucial to accurately match the corresponding forward price. As we use time-stamped tick-by-tick data, matching of option and future prices is straightforward. We apply the method of Hafner und Wallmeier (2001) to account for dividend effects and to ensure put-call-parity consistent estimates of implied volatilities.

For each trading day and each time to maturity available on that day, we estimate a smooth curve of implied volatilities across strike prices. Let K denote the strike price of an option with time to maturity $T - t$. Each trade is assigned a moneyness according to:

$$M(t, T, F_t(T), K) = \frac{\ln \frac{K}{F_t(T)}}{\sqrt{T - t}},$$

where $F_t(T)$ is the (intraday) forward price at the time of this trade.

Suppressing the arguments of moneyness, we chose the cubic regression function:

$$\sigma = \beta_0 + \beta_1 M + \beta_2 M^2 + \beta_3 D \cdot M^3 + \varepsilon, \quad (2)$$

where σ is the implied volatility, β_i , $i = 0, 1, 2, 3$ are regression coefficients, ε is a random error, and D is a dummy variable which assumes the value one for positive moneyness and zero otherwise. The dummy variable accounts for an asymmetry of the pattern of implied volatilities around the at-the-money strike ($M = 0$). Typically, the 'smile' is better characterized by a 'sneer', with the negative relation between implied volatility and moneyness extending clearly beyond $M = 0$. Only when the call (put) is deep out-of-the-money (in-the-money) the implied volatility function forms a minimum and eventually rises slightly. A quadratic or cubic regression without differentiating between $M \leq 0$ and $M > 0$ does not capture this increase. The regression function (2) is twice differentiable, which ensures that the corresponding risk-neutral density is continuous.

The smile estimation according to Eq. (2) is based on all trades of one day in options with the same time to maturity. In order to obtain an estimate of the smile for a given, pre-specified time to maturity of x calendar days, we linearly interpolate between the implied variances of the two neighbouring maturities which are available (see, e.g., Wilmott (1998), p. 290). In this study, we choose $\tau = 45$, because ODAX and OESX option series with lifetimes between 30 and 60 days are the most liquid contracts, which ensures an accurate estimation of the smile. The average R^2 -coefficient over all trading days of one year is 92% in 1995 and larger than 95% in all later years.

We assume the smile function of Eq. (2) to be valid in a moneyness-range between the lowest and highest moneyness of all observations. Outside this range, following Carr und Wu (2005) and Jiang und Tian (2005), we assume implied volatilities to be constant on the volatility level of the relevant moneyness boundary. This corresponds to a conservative estimate of the fair values of options far in and out-of-the-money. Other extrapolation techniques would provide higher variance swap rates and (even) lower variance returns.

¹ We are very grateful to the Eurex for providing the data.

Having estimated the smile structure, we convert implied volatilities from the regression function (2) into call and put option prices (functions $C_0(K, T)$ and $P_0(K, T)$). It is then straightforward to numerically calculate the integrals in (1). The calculations result in an estimate of the variance swap value for each trading day in the sample period.

2.3 Return Distributions

Institutional investors will typically hold a variance swap position over a certain horizon. We therefore assume a buy-and-hold strategy over the lifetime of each contract (45 days). In order to obtain a complete picture of possible returns over this holding period, we assume that one contract is bought on each trading day during the period under study. The resulting return series is strongly autocorrelated because of overlapping return periods (day 1 to 45, day 2 to 46, etc.). We account for this characteristic by using serial-dependence adjusted Newey/West (1987)-standard errors (with a lag of 33 trading days, corresponding to 45 calendar days). As a robustness check, we repeated all tests for non-overlapping periods and obtained similar standard error estimates.

Table 1 presents summary statistics on the distribution of DAX and ESX variance swap returns. The mean returns are negative in all time periods. This holds for discrete returns R_{VARS} as well as log returns r_{VARS} . The large t -statistics suggest that the mean returns are significantly different from zero.² The implication of this is that variance swap levels over-estimate subsequent realized variance. It is well-known that ATM implied volatility tends to be higher than subsequent realized volatility. Due to the option's smile, variance swap levels are even higher than the ATM variance, which increases the spread further. As the table shows, it has been more profitable to initiate a short position in DAX variance swaps in the second period from 2000 to 2004 than in the first period from 1995 to 1999. In the second period, however, it would have been even more advantageous to shorten ESX variance swaps. The mean discrete return for 45-days ESX variance swaps in the period from 2000 to 2004 is minus 19.3%, while it is only minus 13.6% for DAX variance swaps. The above results show that on average, investors are willing to accept a heavily negative risk premium for being long in realized variance. Equivalently, investors who are providing insurance to the market, i.e. who are sellers of variance, require a significantly positive risk premium.

The payoff and discrete return distributions of DAX and ESX variance swaps are clearly non-normal: they show positive skewness and excess kurtosis. The log transformation of discrete returns to continuously compounded returns, however, reduces the skewness estimate for DAX (ESX) variance swap returns in the period from 1995-2004 (2000-2004) from 2.689 (2.282) to 0.560 (0.676) and the kurtosis estimate from 13.246 (8.793) to 3.690 (3.545). The log return distributions appear to be close to normal distributions, though standard tests (Jarque-Bera test, Kolmogorov-Smirnov goodness-of-fit test, etc.) reject normality. Figure 1 illustrates these findings. It shows the empirical density functions for the log returns of DAX and ESX variance swaps along with normal distributions having the same means and the same variances as those estimated from the samples. The distributions of DAX and ESX log variance swap returns

² There is only one mean return (DAX, period 2000-2004) that is not statistically significant at the 5% level.

		DAX			ESX
		95-04	95-99	00-04	00-04
R_{VARS}	Mean	-0.125	-0.114	-0.136	-0.193
	t -statistic Mean	-2.468	-1.988	-1.626	-2.754
	Median	-0.277	-0.228	-0.321	-0.342
	Minimum	-0.792	-0.792	-0.787	-0.809
	Maximum	3.886	1.966	3.886	2.653
	Std. dev.	0.548	0.451	0.632	0.520
	Skewness	2.689	1.435	3.037	2.282
	Kurtosis	13.246	5.308	13.921	8.793
r_{VARS}	Mean	-0.271	-0.235	-0.307	-0.360
	t -statistic Mean	-5.877	-3.927	-4.412	-5.270
	Median	-0.324	-0.259	-0.387	-0.419
	Minimum	-1.569	-1.569	-1.547	-1.658
	Maximum	1.586	1.087	1.586	1.295
	Std. dev.	0.498	0.472	0.521	0.510
	Skewness	0.560	0.176	0.890	0.676
	Kurtosis	3.690	2.781	1.475	3.545

Table 1: Summary Statistics for discrete returns and log returns of variance swaps on the DAX and ESX index with a time to maturity 45 calendar days. Robust t -statistics are calculated using the Newey-West estimator for the standard deviation with a lag of 33.

are in most cases close to each other (see Figure 2). In some cases, however, they strongly deviate. For example, during the period July 30th to September 10th 2001 the log return of DAX variance swaps is significantly higher than the log return of ESX variance swaps (points marked as crosses in the graph). Investors who bought variance swaps in this period realized substantial profits, since September 11 is within the swap's time to maturity, so that realized variance was extremely high.




Figure 1: Histogram of 45 calendar day log returns of DAX variance swaps (left graph) and ESX variance swaps (right graph) over the sample periods 1995-2004 (DAX) and 2000-2004 (ESX).




Figure 2: Log return variance swap DAX versus log return variance swap ESX within the period 2000-2004. The crosses mark all dates where September 11 is within the swap's time to maturity.

3 Portfolio Analysis

3.1 Mean-Variance Efficient Portfolios

Given the historical return distributions, how should investors structure their portfolios to profit in the best possible way from the attractive returns of selling positions of variance swaps or the diversification benefit that buying positions provide? To tackle this question, we first compare mean-variance efficient portfolios with and without considering variance swaps and secondly analyze optimal portfolios under power utility. In the third subsection, we will explicitly incorporate skewness and kurtosis preferences into the analysis using the *Polynomial Goal Programming* method. Our asset universe always consists of variance swaps, its underlying stock index, and the risk-free asset.

So far, we have defined variance swap returns as the quotient of realized variance and the present value of the delivery price K_{VARS} . We thereby assume that the swap buyer makes an up-front payment of $e^{-rT}K_{VARS}$ in order to receive a payment of 1 euro times realized variance at delivery. In reality, though, except for the margin requirements, it costs nothing to enter into the variance swap contract. Since there is no initial investment, we can neither characterize the profit or loss in relative terms nor determine the weight of variance swaps in the investor's portfolio. To overcome this problem, we 'deleverage' the contract by introducing an up-front payment in the form of a risk-free investment. The proceeds of the risk-free asset enhance the net payoff at expiry. An investment of $e^{-rT}K_{VARS}$ seems to be the natural choice since this amount corresponds to the present value of a net payoff equal to the realized variance. However, this still implies a high degree of leverage compared to stocks, as can be seen from higher return fluctuations. Therefore, as an alternative, we assume a risk-free investment of $f \cdot e^{-rT}K_{VARS}$, where the factor f is chosen such that the volatility of variance swap returns is equal to the index return volatility in our sample period. This makes it easier to interpret the portfolio weights of variance swaps and to compare them with the weights of stocks. It is evident, that the set of mean variance efficient portfolios will be the same for any choice of f as long as the risk-free asset is part of the asset universe. We refer to the first return definition ($f = 1$) as *HL* (High Leverage) and to the second definition (same volatility as stock index) as *LL* (Low Leverage).

Figure 3 illustrates the mean-variance analysis for the *LL*-case using estimates from our sample period. The asset universe consists of the DAX index (weight x_S), the DAX variance swap (x_{VARS}) and the risk-free asset (x_{rf}). The sample average and the sample standard deviation of DAX returns over all intervals of 45 days in the period from 1995 to 2004 were 0.76% and 8.66%, respectively. Over the same set of intervals, an *LL*-mean return of -2.05% was observed for DAX variance swaps. Line (1) is the efficient frontier without considering variance swaps ($x_S + x_{rf} = 1$), whereas line (2) represents all combinations of DAX and variance swaps without the risk-free asset ($x_{VARS} + x_S = 1$). If we allow all three assets to enter into the portfolio, we obtain the new efficient line (5). All portfolios on this line are characterized by the same ratio x_{VARS}/x_S but different weights of the risk-free asset. Pre-specifying the weight x_{rf} and maximizing the Sharpe ratio, we obtain one point on the efficient line (5). For instance, line (4) with tangency portfolio T_1 represents all portfolios with $x_{rf} = 1.1$, line (3) with tangency portfolio T_2 represents all portfolios with $x_{rf} = 2$. As we move on the efficient line towards




Figure 3: Mean variance analysis of DAX and variance swap investments. The expected returns and variances are estimated from the sample of return observations in all intervals of 45 days in the period from 1995 to 2004.

combinations of higher risk and return, the short sales of the risky part of the portfolio increase, meaning that the sum $x_{VARS} + x_S$ becomes more negative. To the same extent, the weight of the risk-free asset increases. This increase results in a riskier portfolio since the risk-free investment is financed by short selling risky assets. In the *HL*-case of our return definition, we obtain the same efficient line (5). The efficient portfolios are merely characterized by a different ratio of variance swap and stock index weights.

Table 2 summarizes characteristics of mean-variance efficient portfolios. The first part of the table ('Base case') is based on sample estimates of mean returns, standard deviations and the correlation coefficient. In order to examine the sensitivity of the results to errors in the estimated variance risk premium, we then increase ('Case 2') or decrease ('Case 3') the mean return of variance swaps by twice the Newey-West standard error of the mean estimate. Leaving all other input parameters as they are, we obtain the results shown in the second and third part of Table 2. SR_0 denotes the Sharpe ratio without variance swaps, and SR is the Sharpe ratio of efficient portfolios including variance swaps.

The weight of variance swaps is always negative with the one exception of the first subperiod with raised DAX swap returns (Case 2). Typically, the stock index also enters into the efficient portfolios with a negative weight. This short-selling fits to the short position in variance swaps, because in this way, investors make use of the negative return correlation to achieve better diversification. In Case 2, however, short-selling the index is only suitable in the second subperiod.

		Case 1: Base case					
SR_0		SR	x_{VARS}	x_S	x_{rf}	$LL : \frac{x_{VARS}}{x_S}$	$HL : \frac{x_{VARS}}{x_S}$
DAX 95-04	0.0876	0.2394	< 0	< 0	> 1	6.1140	0.9655
DAX 95-99	0.3632	0.3743	< 0	> 0	< 1	-0.3394	-0.0566
DAX 00-04	-0.1336	0.3728	< 0	< 0	> 1	1.1667	0.1709
ESX 00-04	-0.1740	0.5924	< 0	< 0	> 1	1.2438	0.1816
		Case 2: Higher return of variance swap (+2 STD)					
SR_0		SR	x_{VARS}	x_S	x_{rf}	$LL : \frac{x_{VARS}}{x_S}$	$HL : \frac{x_{VARS}}{x_S}$
DAX 95-04	0.0876	0.0885	< 0	> 0	< 1	-0.1888	-0.0298
DAX 95-99	0.3632	0.4143	> 0	> 0	< 1	0.4814	0.0803
DAX 00-04	-0.1336	0.1399	< 0	< 0	> 1	0.3055	0.0448
ESX 00-04	-0.1740	0.3059	< 0	< 0	> 1	0.8882	0.1297
		Case 3: Lower return of variance swap (-2 STD)					
SR_0		SR	x_{VARS}	x_S	x_{rf}	$LL : \frac{x_{VARS}}{x_S}$	$HL : \frac{x_{VARS}}{x_S}$
DAX 95-04	0.0876	0.4413	< 0	< 0	> 1	3.0466	0.4811
DAX 95-99	0.3632	0.5260	< 0	> 0	< 1	-3.1768	-0.5301
DAX 00-04	-0.1336	0.6683	< 0	< 0	> 1	1.4188	0.2079
ESX 00-04	-0.1740	0.8980	< 0	< 0	> 1	1.4043	0.2050

Table 2: Characteristics of mean-variance efficient portfolios.

The Sharpe ratio SR for DAX amounts to 0.37 in both subperiods, but only 0.24 in the full sample. This is due to the different portfolio structures in the two subperiods. We need a long index position in the first period and a short position in the second to reach the higher Sharpe ratio of 0.37. If we only compose one portfolio for the full sample, this will be suboptimal in both subperiods.

The optimal ratio of x_{VARS} and x_S strongly differs along time period and underlying index. In the Cases 1 and 3, the variance swap weight typically exceeds the stock index weight. It is interesting to note that the short position in variance swaps is not necessarily extended when assuming more strongly negative variance swap returns. For instance, over the full period of DAX returns, efficient portfolios are characterized by a ratio x_{VARS}/x_S of 6.1, compared to 3.0 in Case 2. Thus, variance swaps are less aggressively sold compared to the underlying index, although the absolute variance risk premium has been raised. The reason for this counter-intuitive observation is that the risk reduction resulting from less divergent weights x_{VARS} and x_S is larger than the loss in expected return.

3.2 Backtesting under Power Utility

Table 3 shows optimal portfolio weights for an investor who maximizes his expected utility based on the power utility function with risk aversion parameter α . As in the previous section, we differentiate between three cases. In the base case, the bivariate distribution of excess returns of variance swaps and the underlying stock index is set equal to the observed distribution in the sample period. The columns ‘+2 STD’ result from shifting all variance swap returns by twice the Newey-West adjusted standard error of the volatility risk premium. In the case ‘-2 STD’, the adjustment goes in the opposite direction, so that the negative risk premium gets even larger. The table is based on variance swaps that are levered such that their sample return volatility equals the volatility of the stock index return (LL -definition of previous section). The portfolio weights are restricted to lower and upper bounds of -3.0 and 3.0 , respectively.

The results can be summarized as follows: In the base case, the weights x_{VARS} are all negative. The size of the short position goes down with a higher degree of risk aversion. The variance swap weight is always lowest in the case ‘-2 STD’ and highest in the case ‘+2 STD’. In the second subperiod, the investor also takes a short position in stocks, but its weight is smaller than x_{VARS} . Since the optimal portfolio typically contains short positions in the index and in variance swaps, the risk-free asset often has a heavy weight. The period from 1995 to 1999 provides substantially different results. The weights of the short position in variance swaps are rather small, and the risky part of the portfolio is strongly concentrated on long index holdings. This is certainly due to high stock returns during that period of ‘irrational exuberance’.

Overall, the results are similar to the preceding mean variance analysis.

DAX 1995 - 2004									
α	Base case			+2 STD			-2 STD		
	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}
1	-1.56	-0.31	2.87	-0.10	0.89	0.21	-1.73	-0.27	3.00
1.5	-1.22	-0.25	2.47	-0.08	0.60	0.48	-1.58	-0.42	3.00
2	-0.98	-0.20	2.18	-0.06	0.45	0.61	-1.44	-0.56	3.00
5	-0.44	-0.09	1.53	-0.03	0.18	0.85	-0.73	-0.33	2.06
10	-0.22	-0.04	1.26	-0.01	0.09	0.92	-0.39	-0.17	1.56
DAX 1995 - 1999									
α	Base case			+2 STD			-2 STD		
	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}
1	-0.21	2.87	-1.66	0.80	3.20	-3.00	-2.65	0.65	3.00
1.5	-0.23	2.17	-0.94	1.13	2.87	-3.00	-2.33	0.41	2.92
2	-0.21	1.71	-0.50	1.34	2.66	-3.00	-1.88	0.37	2.51
5	-0.12	0.73	0.39	0.73	1.30	-1.03	-0.84	0.19	1.65
10	-0.06	0.37	0.69	0.37	0.66	-0.03	-0.43	0.10	1.33
DAX 2000 - 2004									
α	Base case			+2 STD			-2 STD		
	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}
1	-1.19	-0.81	3.00	-0.40	-1.60	3.00	-1.64	-0.36	3.00
1.5	-1.10	-0.90	3.00	-0.30	-1.10	2.40	-1.43	-0.57	3.00
2	-1.05	-0.95	3.00	-0.23	-0.83	2.06	-1.30	-0.70	3.00
5	-0.67	-0.69	2.36	-0.10	-0.34	1.44	-1.01	-0.98	2.99
10	-0.34	-0.34	1.68	-0.05	-0.17	1.22	-0.53	-0.50	2.03
ESX 2000 - 2004									
α	Base case			+2 STD			-2 STD		
	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}	x_{VARS}	x_S	x_{rf}
1	-1.64	-0.36	3.00	-0.81	-1.19	3.00	-2.18	0.18	3.00
1.5	-1.42	-0.58	3.00	-0.87	-1.13	3.00	-1.85	-0.15	3.00
2	-1.31	-0.69	3.00	-0.89	-1.11	3.00	-1.64	-0.36	3.00
5	-1.08	-0.92	3.00	-0.73	-0.85	2.58	-1.20	-0.80	3.00
10	-0.70	-0.65	2.35	-0.37	-0.43	1.80	-0.99	-0.88	2.87

Table 3: Optimal portfolio weights under power utility.

3.3 Polynomial Goal Programming

Maximization of expected utility is based on the return distribution as a whole and therefore implicitly takes into account higher moments such as skewness and kurtosis. However, it is generally not evident how to adjust the risk aversion parameter in the utility function to capture heterogeneous investor preferences with respect to different moments of the return distribution. To be able to explicitly control the influence of higher moments, we use *Polynomial Goal Programming (PGP)* which was introduced by Lai (1991) and later applied by Chunnachinda et al. (1997), Prakash et al. (2003), Sun und Yuxing (2003), Canela und Colla (2004), and Davies et al. (2004). This method allows to incorporate different objectives whose relative importance depends on individual preference parameters. We assume that investors aim at a high Sharpe ratio SR , a high (positive) skewness SK and a low kurtosis KT , where

$$SR = \frac{E[R]}{\sigma}; \quad SK = \frac{E[(R - E[R])^3]}{\sigma^3}; \quad KT = \frac{E[(R - E[R])^4]}{\sigma^4}.$$

R denotes excess returns and σ is the standard deviation of R . If the investor attempts to find an optimal compromise between these objectives, she has to weigh up a better achievement of one objective against a loss in another. We assume that her overall assessment is based on the discrepancy between the degree of actual achievement and the maximal degree of achievement (SR^{\max} , SK^{\max} , KT^{\min}) that is possible if one objective alone is considered. Since Sharpe ratio, skewness and kurtosis are of different magnitudes, we have to rescale these measures to make them comparable. Therefore, we express the discrepancies between the maximal degree and the actual degree of achievement in *relative* terms. Our overall objective function to be minimized is defined as:

$$Z(x_{VARS}, x_S) = (1 + d_{SR}(x_{VARS}, x_S))^\alpha + (1 + d_{SK}(x_{VARS}, x_S))^\beta + (1 + d_{KT}(x_{VARS}, x_S))^\gamma, \quad (3)$$

where:

$$d_{SR}(x_{VARS}, x_S) = \frac{SR^{\max} - SR(x_{VARS}, x_S)}{SR^{\max}}, \quad (4)$$

$$d_{SK}(x_{VARS}, x_S) = \frac{SK^{\max} - SK(x_{VARS}, x_S)}{SK^{\max}}, \quad (5)$$

$$d_{KT}(x_{VARS}, x_S) = \frac{KT(x_{VARS}, x_S) - KT^{\min}}{KT^{\min}}. \quad (6)$$

The parameters $\alpha \geq 0$, $\beta \geq 0$, and $\gamma \geq 0$ represent the investor's preference on Sharpe ratio, skewness and kurtosis. The weight of the risk-free asset is given by $x_{rf} = 1 - x_{VARS} - x_S$.

The ratios d_{SR} , d_{SK} , and d_{KT} according to Eq. (4) to (6) are invariant with respect to multiplying x_{VARS} and x_S by the same constant $c > 0$. Such a transformation changes portfolio variance by the factor c^2 . This means that the solution of the PGP optimization problem is compatible with any degree of portfolio variance. The desired level of variance is simply achieved by adjusting the weight of the risk-free asset without changing the structure of the risky sub-portfolio. Thus, we can pre-specify variance without restricting the space of solutions. For convenience, we choose $\sigma^2 = 1$. As before, the expected values needed as input parameters are estimated by the respective mean values in the historical return sample. In addition to

the parameters entering in the mean-variance optimization, the parameter set also includes the skewness and kurtosis of variance swap and underlying index returns as well as the co-movements $E[(R_{VARS} - E[R_{VARS}])^x, (R_S - E[R_S])^y]$ with $(x, y) \in \{(1, 2), (2, 1), (1, 3), (3, 1), (2, 2)\}$.

Figure 4 illustrates some important elements of the optimization results for DAX variance swaps over the total period from 1995 to 2004. Since short selling is allowed, the asset weights for portfolios with the same variance are given by an ellipse that is concentric about the point of the minimum variance portfolio characterized by $x_{VARS} = x_S = 0$ (and therefore $x_{rf} = 1$). The ellipse shown in Figure 4a) represents a portfolio variance of 1. Corresponding to this high volatility of 100%, the asset weights are rather extreme. But as already mentioned, it is only the *ratio* of asset weights that matters in our analysis. Their magnitude can be easily scaled down to be compatible with a lower variance.

On the line in Figure 4a), we highlight the portfolio compositions resulting from separately optimizing one of the objectives Sharpe ratio, skewness and kurtosis.³ As is already known from the mean-variance analysis in Section 3.1, the portfolio with the highest Sharpe ratio is short in variance swaps as well as the underlying index, with a far more heavy (negative) weight on variance swaps. Maximizing portfolio skewness results in a completely different portfolio at the opposite side of the isovariance ellipse. The KT^{\min} -portfolios are again very different from both the SR^{\max} - and SK^{\max} -portfolio. This illustrates to what degree the three objectives are conflicting. The conflict of skewness and sharpe ratio objectives is also apparent from the different alignment of the isovariance lines of $1 - d_{SR} = SR/SR^{\max}$ and $1 - d_{SK} = SK/SK^{\max}$ against x_{VARS} in Figure 4b). The best portfolio with respect to skewness delivers almost the worst results with respect to Sharpe ratio, and vice versa.

The investor's optimal portfolio depends on her preference structure represented by the triple (α, β, γ) . Figure 4c) shows how the asset weights x_{VARS} and x_S depend on β when the other parameters are fixed at $\alpha = 1$ and $\gamma = 0$. In the case of $\beta = 0$, the Sharpe ratio is the only relevant objective, so that the asset weights correspond to the SR^{\max} -portfolio. With increasing β , skewness becomes more and more important relative to the Sharpe ratio. At first, the extent of short selling stocks rises without significant changes in the variance swap position. The optimal portfolio thus moves from the point SR^{\max} in Figure 4a) in the direction of the lower KT^{\min} -portfolio. This smooth adjustment stops when β exceeds a level of about 1.32. At this threshold, the portfolio weights jump to the SK^{\max} -portfolio composition. The reason for this abrupt rearrangement of the optimal portfolio is that the objective function $Z(x_{VARS}, x_S)$ has two local minima, one at a strongly positive variance swap weight, and the other at a strongly negative weight (see Figure 4d). Portfolios with x_{VARS} in-between these values are apparently not attractive. For β lower than 1.32, the investor chooses the 'left' local minimum which represents a structure close to the SR^{\max} -portfolio, and for β higher than 1.32, the investor chooses the 'right' local minimum which is almost identical to the SK^{\max} -portfolio. Thus, the tradeoff between Sharpe ratio and skewness seems to be such that investors cannot always find a reasonable compromise. The core of the matter is that they primarily face a decision between two real *alternatives* so that the optimization problem resembles an 'all or nothing' decision.

³ There are two portfolios with minimal kurtosis, since $KT(x_{VARS}, x_S) = KT(-x_{VARS}, -x_S)$.

To further investigate this observation, Figure 4e) plots the relative degree of achievement of the Sharpe ratio objective (measured by $1 - d_{SR} = SR/SR^{\max}$) against the degree of achievement of the skewness objective (measured by $1 - d_{SK} = SK/SK^{\max}$) for all portfolios with unit variance. Given an arbitrary level of skewness achievement, investors will always prefer the portfolio with a higher Sharpe ratio. Therefore, only the upper line is relevant and might be called ‘efficient’. On the extreme left, the Sharpe ratio is maximal, on the extreme right, skewness is maximal. The efficient line is concave in its left and convex in its right part. Starting on the left, skewness can at first be improved without much loss in Sharpe ratio. But the higher the slope coefficient (in absolute terms) of the efficient line, the more ‘expensive’ in terms of Sharpe ratio is a further increase in portfolio skewness. If an investor is willing to pay this price, it makes sense to even go a step further and choose a portfolio near SK^{\max} . The reason is that the slope eventually diminishes again, rendering a further increase in skewness less costly. With increasing β in the setting ($\alpha = 1, \beta, \gamma = 0$), the optimal portfolio moves from SR^{\max} to T ($\beta = 1.32$) and then jumps to a point near SK^{\max} . The middle part of the efficient line is no-man’s-land: Those who go beyond T will not stop before reaching SK^{\max} .

If we include kurtosis preference into the analysis, the portfolio structure once again does not slowly shift towards the kurtosis minimal portfolio, but it exhibits a pronounced jump at a certain γ -threshold (see Figure 4e). As is already known, in the case of $\alpha = 1$ and $\beta = 1.33$, the optimal portfolio without considering kurtosis ($\gamma = 0$) is close to the SK^{\max} -portfolio. The asset weights are stable as long as γ is below 0.21. But as soon as γ exceeds this threshold, the optimal solution shifts to one of the two KT^{\min} -portfolios. In the case of DAX variance swaps, the KT^{\min} -portfolio with negative asset weights is chosen since it represents a more attractive combination of skewness and Sharpe ratio.

The corresponding results for portfolios composed of the European index ESX and the ESX variance swaps are shown in Figure 5. These results are based on the shorter time period from 2000 to 2004. The main observations are very similar to the DAX analysis. In particular, there is once again a sharp distinction between skewness driven optimal portfolios and Sharpe ratio driven investments. This separation is possible even if the investor population represents a continuum of α - and β -preferences (see the jump of asset weights in Figure 5c). If we incorporate kurtosis preference into the optimization, investors with a strong preference for low kurtosis will select the better of the two KT^{\min} -portfolios. But once again, the transition from the SK^{\max} - to the KT^{\min} -portfolio is very abrupt which means that the mixture of both portfolios is not attractive.

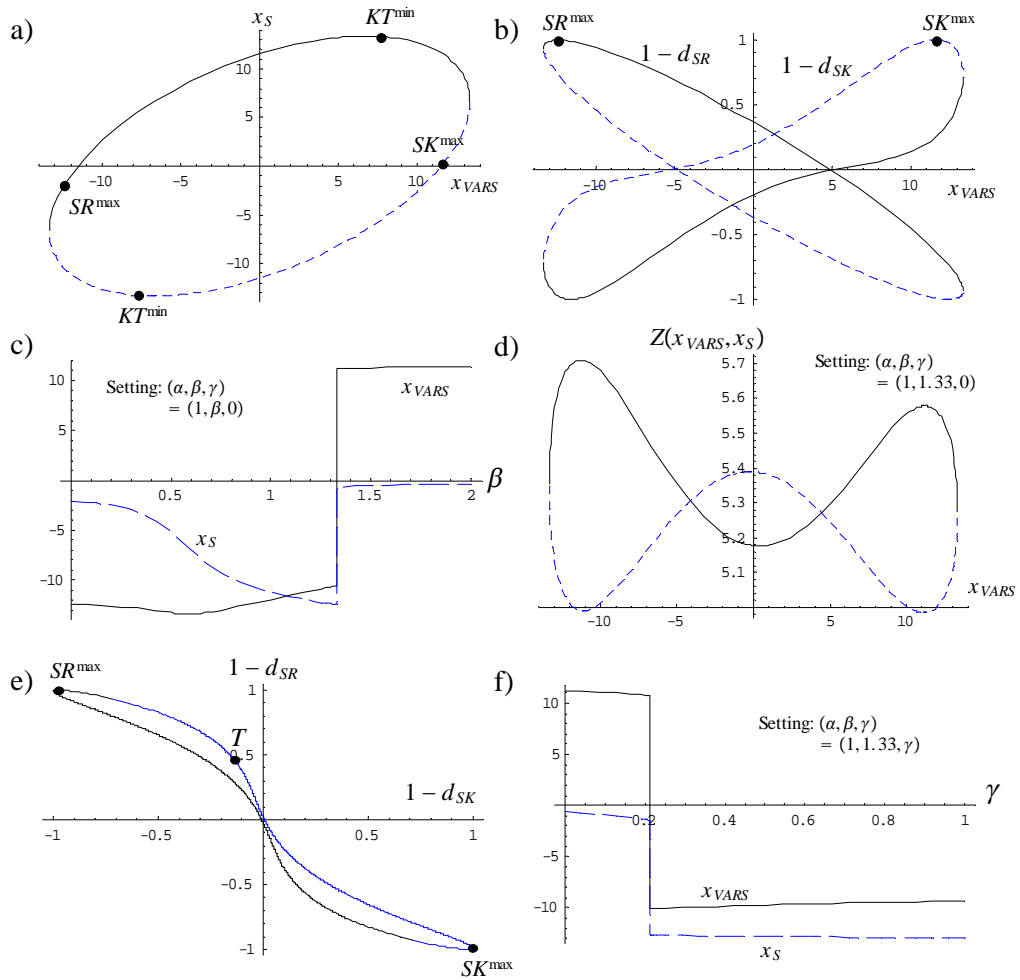


Figure 4: PGP optimization for portfolios of DAX, DAX variance swaps and the risk-free asset based on historical returns over the time period from 1995 to 2004. The lines show the coordinates of all unit variance portfolios. In figures a), b) and d), the solid lines always represent the same set of unit variance portfolios, whereas the coordinates of the remaining portfolios are represented by broken lines.

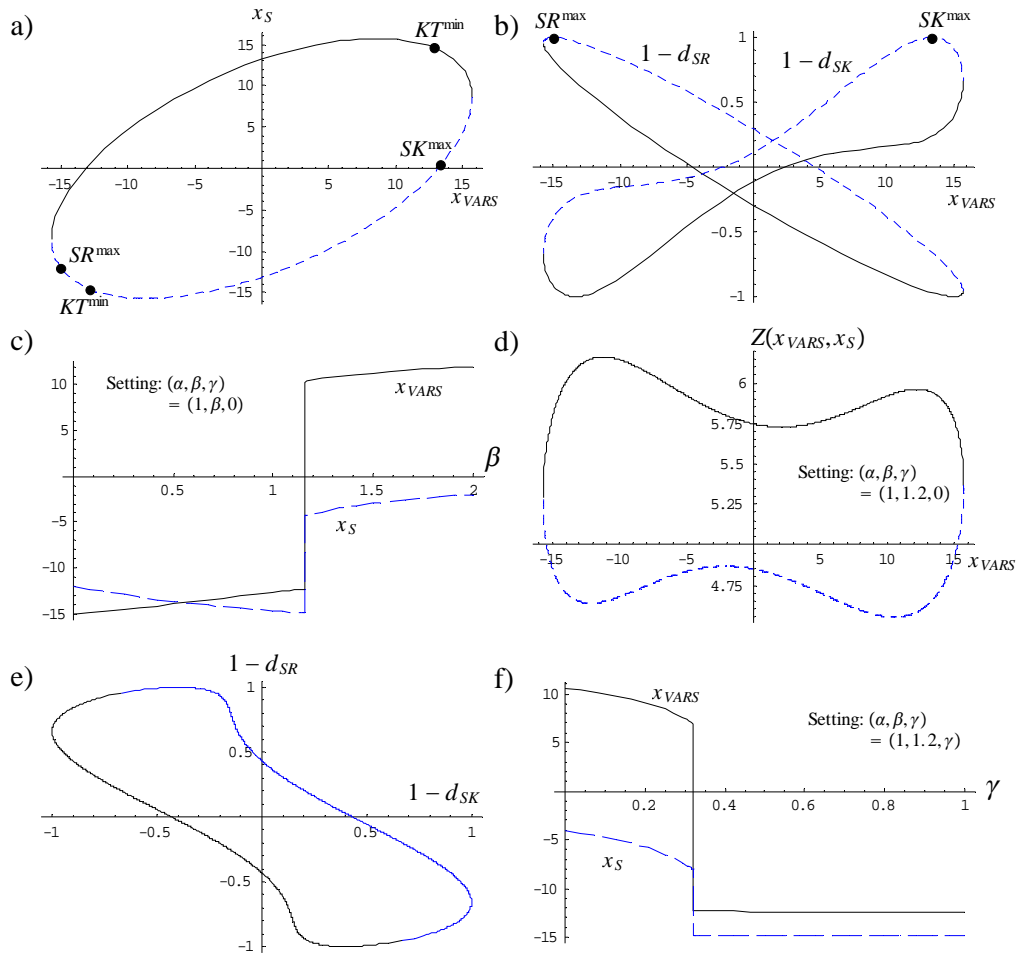


Figure 5: PGP optimization for portfolios of ESX, ESX variance swaps and the risk-free asset based on historical returns over the time period from 1995 to 2004. The lines show the coordinates of all unit variance portfolios. In figures a), b) and d), the solid lines always represent the same set of unit variance portfolios, whereas the coordinates of the remaining portfolios are represented by broken lines.

4 Implementation Aspects

4.1 Risk Management Issues

As was shown before, the risk-return profile of the (short) variance trading strategy is different from that of traditional investment vehicles, such as stocks and bonds. Most notably, the negatively skewed return distribution of selling variance reflects the risk that the strategy will cause substantial losses if volatility suddenly rises. Historically this has happened in situations of (financial) crisis. A prominent example is the collapse of the *Long Term Capital Management (LTCM)* hedge fund. LTCM built up large short positions in equity volatility in the forefront of the Asian and Russian crisis. When volatility increased dramatically, LTCM was confronted with huge losses. The losses from short volatility trades were estimated at 1.3 bn USD – roughly 30% of LTCM's total loss (see Lowenstein (2000)).

The analysis of LTCM's volatility strategy reveals at least two important aspects for risk managing a short variance strategy. First, LTCM measured risk in terms of value at risk (VaR). VaR summarizes the worst loss of a position or portfolio over a given period of time with a given level of confidence. However, VaR does not take into account the severity of an incurred damage event. A measure that is able to quantify losses beyond VaR is expected shortfall or conditional value at risk (CVaR).⁴ Such a measure appears to be more appropriate for risk managing a short variance strategy. In addition, sophisticated stress testing techniques should be applied. The second issue in running the short-selling strategy is liquidity. Although the strategy requires no initial investment (except of a small initial margin), marking positions to market may require additional margin payments when the strategy moves against the investor, i.e. when volatilities rise. Therefore, the investor must hold enough liquidity to compensate for margin calls, otherwise positions will be liquidated. This is particularly important here, because losses are potentially unlimited.⁵ One way to alleviate the effects of an adverse market movement is to hedge part of the tail risk of the strategy, e.g. by buying OTM variance call options. These options are now regularly traded in the OTC market.

4.2 Persistence of Results

It is an important question for investors whether the negative risk premium for realized variance is a temporary phenomenon or results from a structural imbalance which is likely to persist in the future. On the one hand, there are at least two reasons why the phenomenon might be temporary. First, as knowledge about the effect becomes widespread, more investors are willing to exploit it by selling variance. This should lower the prices of variance swaps and diminish any arbitrage profits over time. Second, if the negative risk premium is due to the illiquidity of variance, the premium is again likely to decrease, because liquidity in variance and volatility instruments has substantially improved over the last years.

⁴ CVaR has the additional advantage to be *sub-additive*. This means that the risk of a portfolio is always smaller than or equal to the sum of the risks of its components.

⁵ In practice, variance contracts often exhibit a cap of 250% in volatility terms.

On the other hand, there are a number of reasons why variance swaps could stay overpriced compared to the expected variance during the swap life (see Bennett (2005)). First and most importantly, large institutional investors such as pension funds tend to buy rather than sell large amounts of (OTM) put options to protect themselves, their sponsors and their beneficiaries against big losses. Therefore, the option's price is lifted to accommodate for the costs and risks of delta hedging. This increase is most pronounced for short-dated options, because an option's gamma increases when time to maturity decreases. Hence, being short a near-dated option implies a high gap risk for the trader which has to be compensated. Second, as the (absolute) premium earned from selling deep OTM options is very low, traders are typically reluctant to sell such options. The profits earned are not sufficient to make a big difference in their overall profit and loss, whereas the potential loss in the event of a large move in the underlying asset may be substantial. Third, many investors seem to strongly weigh the risk of crashes, so that they tend to use low strike puts for hedging, thereby bidding their prices up. As low strike puts have a large weight in the variance swap formula, the impact on variance swap prices is substantial. This makes selling variance an attractive strategy for investors who are less risk averse.

5 Conclusion

Mean returns of DAX and ESX variance swaps over the time period from 1995 to 2004 were strongly negative, which is compatible with the observation of a negative volatility risk premium reported in previous studies for the US and Europe. Part of the negative premium can be attributed to the negative correlation of variance swap returns with stock market indices, but according to previous studies, beta alone does not suffice to explain its magnitude (see Carr und Wu (2005), Hafner und Wallmeier (2008)). As a consequence, mean-variance efficient portfolios are characterized by sizable short positions in variance swaps. Typically, the stock index is also sold short to make use of the negative return correlation with variance swaps and in that way achieve a better portfolio diversification. These results are robust with respect to reasonable estimation errors for expected returns. Results are also similar if we take the perspective of an investor who maximizes expected power utility.

Since expected utility maximization is based on the return distribution as a whole, it implicitly takes into account higher moments such as skewness and kurtosis. However, just one risk aversion parameter cannot capture heterogeneous preferences for the moments of the portfolio return distribution. To be able to explicitly control such preferences, we use a variant of the Polynomial Goal Programming method (PGP) proposed by Lai (1991). This method attempts to find an investor-specific optimal combination of the degrees of achievement of different objectives. We assume that investors strive for a high Sharpe ratio, high skewness and low kurtosis. Unfortunately, these objectives are highly conflicting. In particular, the short positions required to profit from negative variance swap returns inevitably lead to an undesired negative skewness. Our analysis shows that a balanced tradeoff between Sharpe ratio and skewness often does not exist. Investors tend to the extreme portfolios (Sharpe ratio driven, skewness driven or kurtosis driven) and avoid to be stuck in the middle. This ‘all-or-nothing’ characteristic is reflected in jumps of asset weights when certain thresholds of preference parameters are crossed.

In this study, we did not investigate whether the variance risk premium can be regarded as an adequate compensation for negative skewness and excess kurtosis. To answer this question, an equilibrium model incorporating higher moments of return distributions is needed. But the analysis clearly shows to what degree a short-selling strategy implies a negatively skewed distribution. It is therefore important to implement effective risk management procedures. We shortly touch upon this requirement in the last section.

An important question for practitioners is whether the negative risk premium for realized variance is a temporary phenomenon or results from a structural imbalance that is likely to persist in the future. On the one hand, more and more hedge funds are willing to exploit the risk premium and increasingly act as variance sellers. On the other hand, the vast majority of market participants, among them most large institutional investors, are still reluctant to sell variance because it implies selling large amounts of out-of-the-money put options. This view is supported by the stability of the smile structure during recent years. Thus, there is some reason to believe that our analysis of historical returns is indicative of future return distributions and relevant for current investor decisions.

References

- Bakshi, G. und Kapadia, N. (2003). Volatility risk premiums embedded in individual equity options: Some new insights, *Journal of Derivatives* **11**(1): 45–55.
- Bennett, C. (2005). *Selling Volatility*, Research Report DrKW, November.
- Black, F. (1976). The pricing of commodity contracts, *Journal of Financial Economics* **3**: 167–179.
- Bliss, R. und Panigirtzoglou, N. (2004). Option-implied risk aversion estimates, *Journal of Finance* **59**: 407–446.
- Bollen, N. und Whaley, R. (2004). Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance* **59**: 711–754.
- Bollerslev, T., Gibson, M. und Zhou, H. (2005). *Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities*, Working Paper.
- Bondarenko, O. (2004). Market price of variance risk and performance of hedge funds, *Working paper*, University of Illinois.
- Canela, M. und Colla, E. (2004). *Portfolio Selection with Skewness in Emerging Markets*, Technical report IESE Business School Barcelona.
- Carr, P. und Wu, L. (2005). Variance risk premia, *Working paper*, Courant Institute and Baruch College, First draft July 2003.
- Chernov, M. und Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation, *Journal of Financial Economics* **56**: 407–458.
- Chunhachinda, P., Dandapani, K., Hamid, S. und Prakash, A. (1997). Portfolio selection and skewness: Evidence from international stock markets, *Journal of Banking and Finance* **21**(2): 143–167.
- Coval, J. und Shumway, T. (2001). Expected option returns, *Journal of Finance* **56**: 983–1009.
- Davies, R., Kat, H. und Lu, S. (2004). *Fund of Hedge Funds Portfolio Selection: A Multiple-Objective Approach*, Working paper ISMA Centre, University of Reading.
- Doran, J. und Ronn, E. (2004a). *The Bias in Black-Scholes/Black Implied Volatility*, Working Paper Florida State University.
- Doran, J. und Ronn, E. (2004b). *On the Market Price of Volatility Risk*, Working Paper Florida State University.
- Driessen, J. und Maenhout, P. (2003). *The World Price of Jump and Volatility Risk*, Working Paper University of Amsterdam.
- Eraker, B., Johannes, M. und Polson, N. (2003). The impact of jumps in volatility and returns, *Journal of Finance* **58**: 1269–1300.
- Hafner, R. und Wallmeier, M. (2001). The dynamics of Dax implied volatilities, *International Quarterly Journal of Finance* **1**: 1–27.

- Hafner, R. und Wallmeier, M. (2008). Volatility as an asset class: European evidence, *forthcoming: European Journal of Finance* .
- Jackwerth, J. und Rubinstein, M. (1996). Recovering probability distributions from option prices, *JF* **51**: 1611–1631.
- Jiang, G. und Tian, Y. (2005). The model-free implied volatility and its information content, *Review of Financial Studies* **18**: 1305–1342.
- Lai, T.-Y. (1991). Portfolio selection with skewness: A multiple-objective approach, *Review of Quantitative Finance and Accounting* **1**(3): 293–305.
- Lowenstein, R. (2000). *When Genius Failed*, Random House, New York.
- Moise, C. (2004). *Is Market Volatility Priced?*, Working Paper University of Chicago.
- Neuberger, A. (1994). The log contract, *Journal of Portfolio Management* (Winter): 74–80.
- Pan, J. (2002). The jump-risk premia implicit in option prices: Evidence from an integrated time-series study, *Journal of Financial Economics* **63**: 3–50.
- Prakash, A., Chang, C.-H. und Pactwa, T. (2003). Selecting a portfolio with skewness: Recent evidence from US, european, and latin american equity markets, *Journal of Banking and Finance* **27**(7): 1375–1390.
- Santa-Clara, P. und Yan, S. (2004). *Jump and Volatility Risk and Risk Premia: A New Model and Lessons from S+P 500 Options*, NBER Working Paper 10912.
- Sun, Q. und Yuxing, Y. (2003). Skewness persistence with optimal portfolio selection, *Journal of Banking and Finance* **27**(6): 1111–1121.
- Wilmott, P. (1998). *Derivatives: The Theory and Practice of Financial Engineering*, New York et al.