

Optimal investments in volatility

Reinhold Hafner · Martin Wallmeier

Published online: 31 March 2008
© Swiss Society for Financial Market Research 2008

Abstract Volatility has evolved as an attractive new asset class of its own. The most common instruments for trading volatility are variance swaps. Mean returns of DAX and ESX variance swaps over the time period of 1995 to 2004 are 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 composition. Mean-variance efficient portfolios are characterized by sizable short positions in variance swaps. Typically, the stock index is also sold short to achieve a better portfolio diversification. To capture heterogeneous preferences for higher moments, we use a variant of the polynomial goal programming method. We assume that investors strive for a high Sharpe ratio, high skewness, and low kurtosis. Our analysis reveals that it is often not possible to achieve a balanced tradeoff between Sharpe ratio and skewness. Investors are advised to hold the extreme portfolios (Sharpe ratio driven, skewness driven, or kurtosis driven) and avoid the middle ground. This “all-or-nothing” characteristic is reflected in jumps of asset weights when certain thresholds of preference parameters are crossed. These empirical findings can explain why many investors are so reluctant to implement option-based short-selling strategies.

Keywords Variance swap · Volatility risk premium · Portfolio analysis · Higher moments · Polynomial goal programming · Hedge funds

JEL Classification G10 · G12 · G13

R. Hafner
risklab germany GmbH, Nymphenburgerstr. 112-116, 80636 Munich, Germany
e-mail: reinhold.hafner@risklab.de

M. Wallmeier (✉)
University of Fribourg, Bd. de Pérolles 90, 1700 Fribourg, Switzerland
e-mail: martin.wallmeier@unifr.ch

1 Introduction

In addition to its prominent role as a financial risk measure, volatility has evolved to become an established asset class of its own. Volatility can be traded by creating contracts whose payoff depends on the volatility realized during the time to maturity. The buyer of the contract is “long volatility” and receives a payment from the seller when realized volatility is high. The seller, who is “short volatility”, benefits from a low realized volatility. There are various volatility instruments, each of which differ in the details of the payoff function.

Initially, the main motive for trading volatility was to manage risk exposure in option positions and to control the vega exposure independently of the position’s delta and gamma. The growth of the volatility trading segment attracted the interest of other market participants, who were particularly interested in the strong negative correlation between volatility movements and stock index returns. One common explanation for this phenomenon is the so-called leverage effect, which occurs where the risk of the equity position increases when stock prices fall and the market value of debt stays constant. The opposite relationship holds when stock prices rise, thus decreasing the leverage ratio in market value terms.¹ According to this explanation, the negative correlation is stock return driven. In contrast, the competing feedback hypothesis assumes that volatility changes are the reason for this effect: higher volatility induces a higher risk premium and, therefore, higher discount rates, thus lowering fundamental stock values. Another explanation why volatility and stock prices should be negatively correlated is to prevent arbitrage opportunities in options markets in the presence of a volatility skew (see Hafner 2004, p. 46). From an investor’s point of view, it seems attractive that the negative correlation between volatility and stock index returns is particularly pronounced in stock market downturns, thus offering protection against stock market losses when such is needed most. This is one reason why financial advisors often recommend the 90/10 rule of investing, i.e., 10% of a stock portfolio should be invested in volatility.

Empirical studies, however, indicate that this kind of downside or crash protection might be expensive. The volatility risk premium is found to be strongly negative in the US as well as in the European stock markets (see, e.g., Jackwerth and Rubinstein 1996; Chernov and Ghysels 2000; Coval and Shumway 2001; Pan 2002; Bakshi and Kapadia 2003; Eraker et al. 2003; Driessen and Maenhout 2003; Doran and Ronn 2004a; Doran and Ronn 2004b; Bondarenko 2004; Moise 2004; Santa-Clara and Yan 2004; Carr and Wu 2005; Hafner and Wallmeier 2007). This means that long volatility positions provide expected returns below the risk-free interest rate. There is some evidence that the magnitude of the premium depends on the level of implied volatility and the time to maturity (see Bliss and Panigirtzoglou 2004), the buying pressure for index puts (see Bollen and Whaley

¹This relationship between stock returns and leverage corresponds to an argument by Welsh (2004) who proposes that changes in leverage are only due to changes in the market value of equity. See also Bessler et al. (2008) for European evidence. However, according to empirical findings of Figlewski and Wang (2000), the leverage changes cannot fully explain the extent of the negative correlation of stock return and volatility changes.

2004) and the uncertainty of forward volatility (see Carr and Wu 2005). The estimated premium has significant temporal dependencies (see Bollerslev et al. 2005; Santa-Clara and Yan 2004). All in all, the negative premium appears too large to be explained by the negative covariance with market returns within standard (partial) equilibrium models (see Carr and Wu 2005; Hafner and Wallmeier 2007). Consequently, 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 historical hedge fund returns.²

It is interesting that a strategy of systematically selling volatility earns, on average, a positive risk premium. The economic rationale behind this empirical observation is that investors receive a premium for providing other market participants with insurance against potential market losses. However, the return distribution of short volatility is negatively skewed (see Carr and Wu 2005; Hafner and Wallmeier 2007), which means 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 had built up large short positions in equity volatility previous to the Asian and Russian crisis in 1997/1998. When volatility dramatically increased, LTCM suffered huge losses. The losses from short volatility trades were estimated to be about US\$1.3 billion, i.e., roughly 30% of LTCM's total loss (see Lowenstein 2000).

The negative skewness in the return distribution of a (short) volatility trading strategy makes the volatility asset class different from traditional investment vehicles, such as stocks and bonds, which has several important implications for asset allocation and risk management. First, higher moments, such as skewness and kurtosis, need to be considered in the asset allocation decision process in addition to the first two moments of the portfolio return distributions. Second, in managing the risk of a short variance strategy, risk measures able to quantify the financial impact of potentially damaging events should be used, one of which is the conditional value at risk (CVaR). Third, losses from a short variance strategy are potentially unlimited.³ 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 out-of-the-money variance call options. These options are now regularly traded in the OTC market.

Although there are many studies on the variance risk premium, the implications for investors have rarely been addressed in any detail to date. To draw any conclusions from the negative risk premium from an investor's point of view, it is necessary to specifically analyze optimal portfolio composition in a realistic setting and this is exactly the focus of our research. We are particularly interested in the tradeoff between Sharpe ratio, skewness, and kurtosis. This tradeoff is important not only for the short volatility strategy presented here, but also for many other hedge fund strategies that

²See also Le Moigne and Savaria (2006), who identify volatility as an important attribute in explaining cross-sectional variations in hedge fund returns, and Füss and Kaiser (2007), who analyze comovements between hedge fund strategies and different asset classes.

³In practice, variance contracts often exhibit a cap of 2.5 times the current swap rate.

produce heavily skewed return distributions.⁴ In general, such strategies are characterized by small profits in “normal” situations and potentially large losses when some rare event occurs.

The results of our analysis suggest that it is often not possible to achieve a balanced tradeoff between Sharpe ratio, skewness, and kurtosis. Investors are advised to hold the extreme portfolios—Sharpe ratio driven, skewness driven or kurtosis driven—and avoid the middle ground. This “all-or-nothing” characteristic is reflected in jumps of asset weights when certain thresholds of preference parameters are crossed. These empirical findings can explain why many investors are so reluctant to implement option-based short-selling strategies.

The remainder of the paper is organized as follows. 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 (Sect. 3). The mean-variance framework (Sect. 3.1) ignores the return properties of negative skewness and excess kurtosis. These are taken into account implicitly in expected utility maximization (Sect. 3.2) and explicitly in the *polynomial goal programming* (PGP) optimization (Sect. 3.3). The paper concludes with a brief summary of our findings.

2 Variance swap returns

2.1 Variance swaps and valuation

Variance swaps, also known as variance contracts, offer pure exposure to the realized future variance. The payment at expiration is equal to the difference between the annualized variance of log stock returns and the swap rate fixed at the outset of the contract. If this difference is positive, the swap buyer receives payment from the swap seller. If it is negative, the swap seller receives payment from the buyer. The swap rate is chosen such that the contract has zero present value. Thus, the swap rate can be interpreted as the risk-neutral expectation of the unconditional future variance. Variance swaps on the most common stock indices now have an active over-the-counter market. This market growth occurred as the result of theoretical research deriving a robust replication strategy. This strategy 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:

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

⁴In analyzing hedge fund returns, Bondarenko (2004) finds a significant relation to variance returns. This indicates that short variance strategies are in part responsible for skewness and kurtosis in hedge fund returns.

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. We assume that there is no credit risk involved. According to (1), the replication portfolio includes out-of-the-money put and call options with weights inversely proportional to the square of their strike. This weighting scheme is necessary to obtain a constant portfolio exposure with respect to realized variance independent of stock price moves (see Demeterfi et al. 1999 for a detailed derivation and economic motivation).

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 2 days, the theoretical values always fall into the spread between bid and ask quotes. This spread seems sufficiently narrow (about one volatility point) to make the comparison meaningful.

To employ (1) in the empirical analysis, we first estimate the strike price structure of implied volatilities ("smile"). The strike-dependent implied volatilities are then inserted 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 on the joint German and Swiss options and futures exchange Eurex.⁵

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 employ the method of Hafner and 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:

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

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

⁵We are very grateful to the Eurex for providing the data.

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 that assumes the value of 1 for positive moneyness and 0 otherwise. The dummy variable accounts for the asymmetry pattern of implied volatilities around the at-the-money strike ($M = 0$). Typically, the “smile” is better characterized by a “skew,” with the negative relation between implied volatility and moneyness extending clearly beyond $M = 0$. Only when the call (put) is deeply out-of-the-money (in-the-money), does the implied volatility function dip to a minimum before eventually rising slightly. A quadratic or cubic regression without differentiating between $M \leq 0$ and $M > 0$ will not capture this increase. The regression function of (2) is twice differentiable, which ensures that the corresponding risk-neutral density is continuous.

The smile estimation derived from (2) is based on all trades on one day in options with the same time to maturity. To obtain an estimate of the smile for a given, prespecified time to maturity of x calendar days, we linearly interpolate between the implied variances of the two neighboring maturities that 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 that the smile function of (2) is valid in a moneyness range between the lowest and highest moneyness of all observations. Outside this range, following Carr and Wu (2005) and Jiang and Tian (2005), we assume that implied volatilities are constant at the volatility level of the relevant moneyness boundary. This corresponds to a conservative estimate of the fair values of options deeply in and out-of-the-money. Other extrapolation techniques would provide higher variance swap rates and (even) lower variance returns.

Having estimated the smile structure, we convert implied volatilities from the regression function of (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 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). 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 problem 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 nonoverlapping periods and obtained similar standard error estimates.

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

		DAX			ESX
		95–04	95–99	00–04	00–04
r_S	Mean	0.008	0.022	-0.013	0.012
	Median	0.022	0.029	<-0.001	-0.006
	Minimum	-0.437	-0.337	-0.437	-0.340
	Maximum	0.249	0.245	0.249	0.221
	Std. dev.	0.089	0.075	0.102	0.078
	Skewness	-0.902	-0.848	-0.670	-0.554
	Kurtosis	1.898	2.219	0.992	0.964
r_{VARS}	Mean	-0.271	-0.235	-0.307	-0.360
	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
<i>Corr</i>		-0.476	-0.421	-0.574	-0.534

Table 1 presents summary statistics of the distribution of DAX and ESX index log returns (r_S) and variance swap log returns (r_{VARS}). The mean variance swap returns are significantly negative in all time periods,⁶ implying that the variance swap rates overestimate the subsequently realized variance. Indeed, it is well known that ATM-implied volatility tends to be higher than subsequently realized volatility. Due to the option's smile, variance swap rates are even higher than the ATM variance, which increases the spread further. As shown in Table 1, it was more profitable to initiate a short position in DAX variance swaps in the second period (2000 to 2004) than in the first period (1995 to 1999). In the second period, however, it would have been even more advantageous to shorten ESX variance swaps. The mean log return for 45-day ESX variance swaps in the period from 2000 to 2004 is minus 36.0%, whereas it is only minus 30.7% for DAX variance swaps. These results indicate that, on average, investors are willing to accept a heavily negative risk premium in exchange for being long in realized variance. Equivalently, investors who are providing insurance to the market, i.e., sellers of variance, require a significantly positive risk premium.

The sample Pearson correlation coefficient of r_S and r_{VARS} is given as *Corr* in Table 1. As expected, it is significantly negative in all cases.

The payoff and discrete return distributions of DAX and ESX variance swaps are clearly nonnormal: they show positive skewness and excess kurtosis. The log trans-

⁶There is only one mean return (DAX, period 2000–2004) that is not statistically significant at the 5% level.




Fig. 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)

formation 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 close to normal, although standard tests (Jarque–Bera test, Kolmogorov–Smirnov goodness-of-fit test, etc.) reject normality. Figure 1 illustrates these findings by showing 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 are in most cases close to each other (see Fig. 2). In some cases, however, they deviate greatly. For example, during the period July 30 to September 10, 2001, the log return of DAX variance swaps was significantly higher than the log return of ESX variance swaps (points marked as crosses in the figure). 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.⁷

3 Portfolio analysis

3.1 Mean-variance efficient portfolios

Given the historical return distributions, the interesting question from an investor perspective is how to structure a portfolio in order to profit from the attractive return opportunities that arise from either selling variance swaps or from the diversification benefit that buying positions provide. To tackle this question, we first compare

⁷See Hafner and Wallmeier (2007) for a description of variance swap rates during other crises.




Fig. 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

mean-variance efficient portfolios with and without variance swaps, and second, analyze optimal portfolios under power utility.⁸ In Sect. 3.3, we explicitly include the investor’s skewness and kurtosis preferences by using the *polynomial goal programming* method. Our asset universe always consists of variance swaps, the underlying stock index, and the risk-free asset.⁹

In Sect. 2, we defined variance swap returns as the ratio of realized variance and the present value of the delivery price K_{VARS} . We thus assume that the buyer of the swap 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, there is no cost involved in entering into the variance swap contract. Since there is no initial investment, we cannot characterize the profit or loss in relative terms; nor can we 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 increase the net payoff at maturity. An investment of $e^{-rT} K_{VARS}$ appears to be the natural choice because 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 investing in stocks, as is evident from the higher return fluctuations. Therefore, as an alternative, we assume a risk-free investment of $f \cdot e^{-rT} K_{VARS}$, where 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 port-

⁸These analyses are based on discrete returns as required by the usual calculations in portfolio theory.

⁹We follow an absolute-performance approach. See Illmer and Marty (2007) for an analysis of the return decomposition of absolute-performance multi-asset class portfolios.

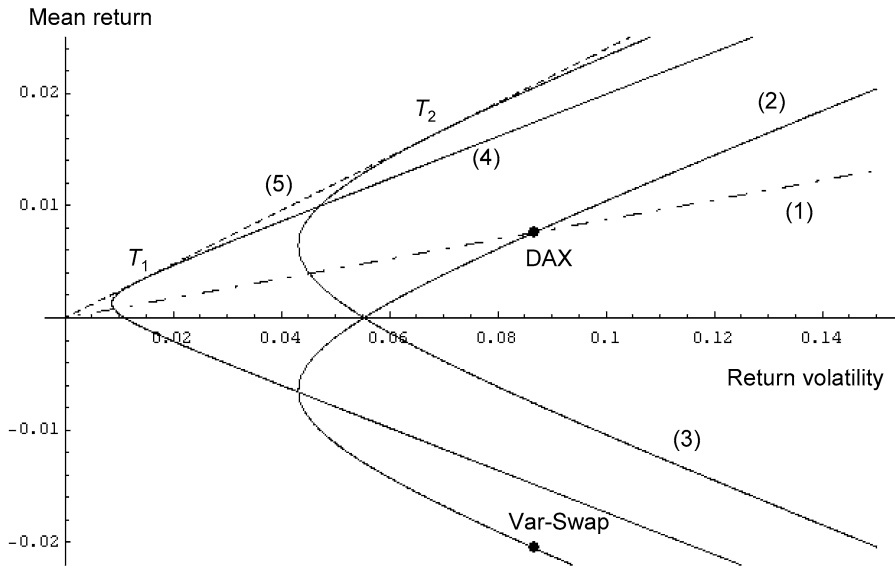


Fig. 3 Mean variance analysis of DAX and variance swap investments. The expected (excess) returns and variances are estimated from the sample of return observations in all intervals of 45 days in the period from 1995 to 2004

folio weights of variance swaps and to compare them to the weights of stocks. It is evident, that the set of mean variance efficient portfolios will be the same for any value 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 45-day intervals 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 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 with different weights of the risk-free asset. Prespecifying 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 along the efficient line toward 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. Similarly, the weight of the risk-free asset increases. This increase results in a riskier portfolio because the risk-free investment is financed by short selling risky assets. In the *HL* case of our return definition, we

Table 2 Characteristics of mean-variance efficient portfolios. Case 1 is based on sample estimates of mean returns, standard deviations, and the correlation coefficient. Cases 2 and 3 result from increasing or decreasing the mean return of variance swaps by twice the standard error of the mean estimate. SR_0 denotes the Sharpe ratio without variance swaps; SR is the Sharpe ratio including variance swaps

	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

obtain the same efficient Line (5). The efficient portfolios vary merely by a different ratio of variance swap and stock index weights.

Table 2 summarizes characteristics of mean-variance efficient portfolios, where we restrict portfolio weights to a range between -3.0 and 3.0.¹⁰ The first part of the table (Case 1) is based on sample estimates of mean returns, standard deviations, and the correlation coefficient. To examine the sensitivity of the results to errors in the estimated variance risk premium, we then either 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 constant, we obtain the

¹⁰It is reasonable to restrict the magnitude of individual asset positions and the extent to which investors are allowed to sell assets short. Most institutional investors have to meet such restrictions. We assume that the individual long and short positions cannot be larger than three times as much as net investment. However, our results are robust with respect to the values of the upper and lower boundaries.

results shown in the second and third part of Table 2. SR_0 denotes the Sharpe ratio without variance swaps; SR is the Sharpe ratio of efficient portfolios including variance swaps. It is important to note that our allocation analysis does not take general equilibrium conditions into account. Because variance swaps are, on aggregate, in zero net supply, large positive or negative weights of these instruments are generally not compatible with general equilibrium.

The weight of variance swaps is negative, with the exception of the first subperiod for DAX swap returns in Case 2. Typically, the stock index also enters the efficient portfolios with a negative weight. This short selling of the stock index matches the short position in variance swaps because in this way, investors benefit from the negative correlation by achieving better diversification. In Case 2, however, selling the index short is appropriate only in the second subperiod. The Sharpe ratio (SR) for DAX is 0.37 in both subperiods, but only 0.24 in the full sample, which is due to the different portfolio structures in the two subperiods. In fact, a long index position is required in the first period and a short position in the second period to reach the higher Sharpe ratio of 0.37. If we construct only one portfolio for the entire sample period, it would be suboptimal in both subperiods.

The optimal ratio of x_{VARS} and x_S varies widely along time period and underlying index. In Cases 1 and 3, the variance swap weight typically exceeds the stock index weight. Interestingly, 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 sold less aggressively compared to the underlying index, although the absolute variance risk premium has increased. The explanation for this counterintuitive 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

In Table 3, we present optimal portfolio weights for an investor who maximizes 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 of “−2 STD,” the adjustment shifts in the opposite direction so that the negative risk premium becomes even larger. The table is based on variance swaps that are levered so that their sample return volatility equals the volatility of the stock index return (LL , see definition in previous section). As before, 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 decreases with a higher degree of risk aversion. The weight of the variance swap is always lowest in the case of “−2 STD” and highest in the case of “+2 STD.” During the second subperiod, the investor also

Table 3 Optimal portfolio weights under power utility for different risk aversion coefficients α . The “Base case” is based on sample estimates of mean returns, standard deviations, and the correlation coefficient. “+2 STD” and “-2 STD” result from increasing or decreasing the mean return of variance swaps by twice the standard error of the mean estimate.

α	DAX 1995–2004								
	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
α	DAX 1995–1999								
	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
α	DAX 2000–2004								
	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
α	ESX 2000–2004								
	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

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 large weight. Our analysis of the period from 1995 to 1999 reveals 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 in long index holdings. This phenomenon is undoubtedly due to the high stock returns during that period of “irrational exuberance.” Overall, the results are similar to the preceding mean variance analysis.

3.3 Polynomial goal programming

Maximization of expected utility is based on the return distribution as a whole, and, therefore, implicitly takes into account the higher moments, such as skewness and kurtosis. However, it is generally not obvious how to adjust the risk-aversion parameter in the utility function to capture heterogeneous investor preferences with respect to the different moments of the return distribution. To explicitly control the influence of higher moments, we use *polynomial goal programming (PGP)*, which was introduced by Lai (1991) and later applied by Chunhachinda et al. (1997), Prakash et al. (2003), Sun and Yuxing (2003), Canela and Colla (2004), and Davies et al. (2004). This method allows us 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 . To discover an optimal tradeoff between these objectives, the investor needs to balance a gain in one objective against a loss in another. We assume that the investor’s 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 need 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:

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

where

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

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

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

The parameters $\alpha \geq 0$, $\beta \geq 0$, and $\gamma \geq 0$ express the investor's preference with respect to the Sharpe ratio, skewness, and kurtosis.¹¹ 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 (4)–(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 , meaning that the solution of the PGP optimization problem is compatible with any degree of portfolio variance. The desired level of variance is achieved simply by adjusting the weight of the risk-free asset without changing the structure of the risky subportfolio. Thus, we can pre-specify the variance without restricting the solution space. For convenience, we choose $\sigma^2 = 1$. As before, the expected values needed as input parameters are estimated by the respective mean values of 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 comovements $E[(R_{VARS} - E[R_{VARS}])^x, (R_S - E[R_S])^y]$ 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 entire 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 Fig. 4(a) represents a portfolio variance of 1. At this high volatility of 100%, the asset weights are rather extreme. However, as previously mentioned, it is only the *ratio* of asset weights that is of any consequence to our analysis and their magnitude can easily be scaled down to be compatible with a lower variance.

On the line in Fig. 4(a), we highlight the portfolio compositions resulting from separately optimizing one of the objectives, i.e., either Sharpe ratio, skewness, or kurtosis.¹² As is already known from the mean-variance analysis in Sect. 3.1, the portfolio with the highest Sharpe ratio is short in variance swaps as well as in the underlying index, with a much larger (negative) weight on variance swaps. Maximizing portfolio skewness results in a completely different portfolio, which is at the opposite side of the isovariance ellipse. The KT^{\min} portfolios are again very different from both the SR^{\max} and SK^{\max} portfolios, graphically illustrating the extent to which the three objectives conflict. The conflict between the 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 Fig. 4(b). The best portfolio with respect to skewness provides almost the worst results with respect to the Sharpe ratio, and vice versa.

The investor's optimal portfolio depends on his or her preference structure, as represented by the triple (α, β, γ) . Figure 4(c) 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

¹¹Skewness preference is an implication of decreasing absolute risk aversion. The experimental and empirical evidence of preference for positively skewed returns is summarized in Harvey et al. (2004). For a discussion of kurtosis preference, see Dittmar (2002).

¹²There are two portfolios with minimal kurtosis, since $KT(x_{VARS}, x_S) = KT(-x_{VARS}, -x_S)$.

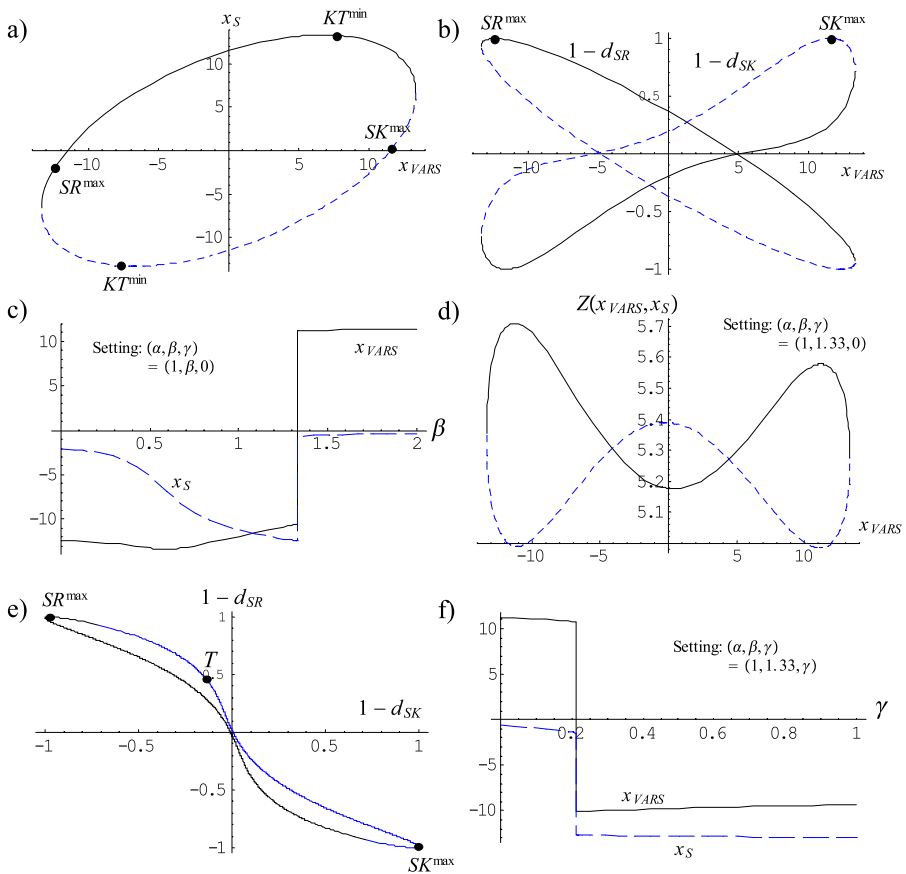


Fig. 4 PGP optimization for portfolios of DAX, DAX variance swaps and the risk-free asset based on historical returns over the time period from 1995 to 2004. The lines show the coordinates of all unit variance portfolios. In figs. (a), (b) and (d), the solid lines always represent the same set of unit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

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 increasingly important relative to the Sharpe ratio. At first, the extent of short selling stocks rises without significant change in the variance swap position. The optimal portfolio thus moves from the point SR^{\max} in Fig. 4(a) 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 where the positive variance swap weight is large, and the other where the negative weight is large (see Fig. 4(d)). Portfolios with x_{VARS} in-between these values do not appear to be 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 between the two. The crux of the matter is that investors are faced with a decision between two real *alternatives* so that the optimization problem resembles an “all or nothing” decision.

To further investigate this observation, Fig. 4(e) 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 a one 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 has a maximum; on the extreme right, skewness has a maximum. The efficient line is concave to its left and convex to its right. Starting on the left, skewness can at first be improved without much loss in Sharpe ratio. However, the higher the slope coefficient (in absolute terms) of the efficient line, the more “expensive,” in terms of Sharpe ratio, a further increase in portfolio skewness becomes. If an investor is willing to pay this price, it makes sense to go even a step further and choose a portfolio near SK^{\max} because 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 SK^{\max} to T ($\beta = 1.32$) and then jumps to a point near SK^{\max} . The middle part of the efficient line is a no-man’s land: those who venture beyond T will not stop until reaching SK^{\max} .

If we include kurtosis preference in the analysis, the portfolio structure once again does not shift continuously toward the kurtosis minimal portfolio, but instead exhibits a pronounced jump at a certain γ threshold (see Fig. 4(e)). 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 Fig. 5. These results are based on the shorter time period of 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, a distinction that persists even if the investor population represents a continuum of α and β preferences (see the jump of asset weights in Fig. 5(c)). 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. However, 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 unattractive.

4 Conclusion

Mean returns of DAX and ESX variance swaps for the time period from 1995 to 2004 were strongly negative. This is in accordance with the observation of

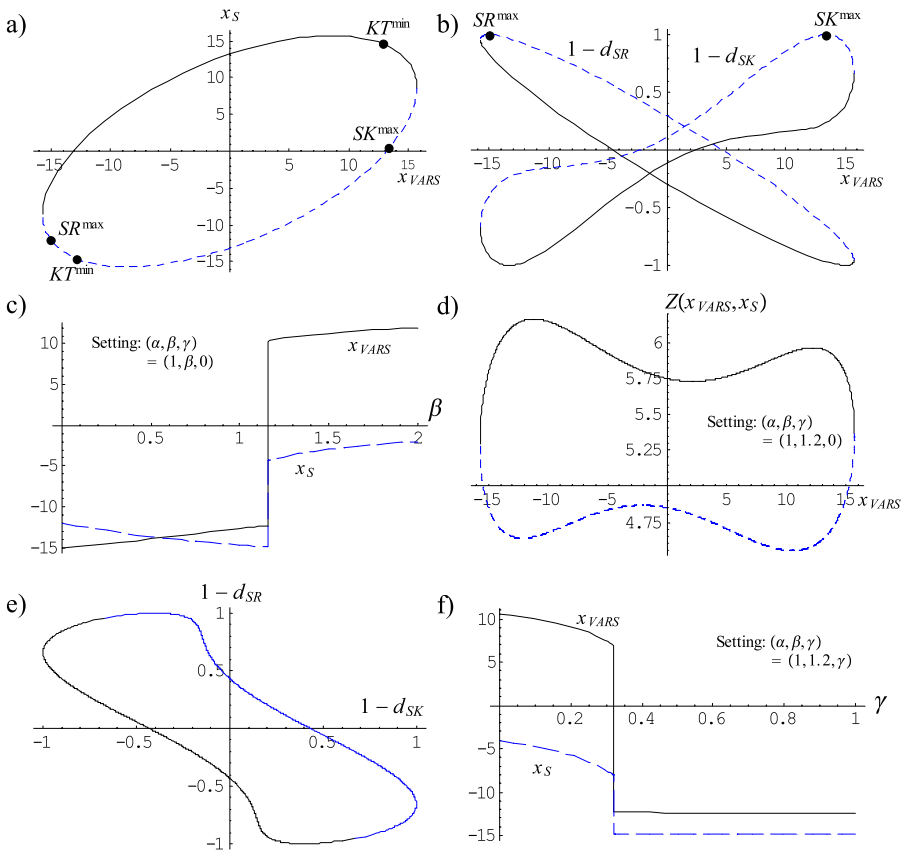


Fig. 5 PGP optimization for portfolios of ESX, ESX variance swaps, and the risk-free asset based on historical returns over the time period from 1995 to 2004. The lines show the coordinates of all unit variance portfolios. In figs. (a), (b), and (d), the *solid lines* always represent the same set of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

a negative volatility risk premium reported in previous studies for the United States and Europe. Part of the negative premium can be attributed to the negative correlation of variance swap returns with stock market index returns. However, beta alone does not suffice to explain its magnitude (see Carr and Wu 2005; Hafner and Wallmeier 2007). 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 take advantage of the negative return correlation with variance swaps and in that way achieve better portfolio diversification. These results are robust with respect to reasonable estimation errors for expected returns. Results are similar if we take the perspective of an investor who maximizes expected power utility.

Because expected utility maximization is based on the entire return distribution, it implicitly takes higher moments such as skewness and kurtosis into account. It is not surprising, however, that merely one risk-aversion parameter cannot capture

investors' heterogeneous preferences with respect to the higher moments of the portfolio return distribution. 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 for achieving different objectives. In our analysis, we assume that investors strive for a high Sharpe ratio, high skewness, and low kurtosis, which is in agreement with standard assumptions on the economic behavior of investors. Unfortunately, the individual objectives are in serious conflict. In particular, the short positions required to profit from negative variance swap returns inevitably lead to an undesired negative skewness. Our analysis shows that it is often not possible to achieve a balanced tradeoff between Sharpe ratio and skewness. Investors tend to the extreme portfolios (Sharpe ratio driven, skewness driven, or kurtosis driven) and avoid the middle ground. 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 viewed as adequate compensation for negative skewness and excess kurtosis. An equilibrium model incorporating higher moments of return distributions is needed to answer this question. However, our 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.

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 into the future. On the one hand, an increasing number of hedge funds are willing to exploit the risk premium and act as variance sellers. On the other hand, the vast majority of market participants, including 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.

Acknowledgement We would like to thank Wolfgang Bessler (special issue editor) and two anonymous referees for their helpful comments and suggestions.

References

- Bakshi, G., Kapadia, N.: Volatility risk premiums embedded in individual equity options: Some new insights. *J. Derivativ.* **11**(1), 45–55 (2003)
- Bessler, W., Drobetz, W., Pensa, P.: Do managers adjust the capital structure to market value changes? Evidence from Europe. *Zeitschrift für Betriebswirtschaft* **78** (2008, forthcoming)
- Black, F.: The pricing of commodity contracts. *J. Financ. Econ.* **3**, 167–179 (1976)
- Bliss, R., Panigirtzoglou, N.: Option-implied risk aversion estimates. *J. Finance* **59**, 407–446 (2004)
- Bollen, N., Whaley, R.: Does net buying pressure affect the shape of implied volatility functions? *J. Finance* **59**, 711–754 (2004)
- Bollerslev, T., Gibson, M., Zhou, H.: Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. Working Paper (2005)
- Bondarenko, O.: Market price of variance risk and performance of hedge funds. Working paper, University of Illinois (2004)

- Canela, M., Colla, E.: Portfolio selection with skewness in emerging markets. Technical report IESE Business School Barcelona (2004)
- Carr, P., Wu, L.: Variance risk premia. Working paper, Courant Institute and Baruch College, July 2003 (2005)
- Chernov, M., Ghysels, E.: A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *J. Financ. Econ.* **56**, 407–458 (2000)
- Chunhachinda, P., Dandapani, K., Hamid, S., Prakash, A.: Portfolio selection and skewness: Evidence from international stock markets. *J. Bank. Finance* **21**(2), 143–167 (1997)
- Coval, J., Shumway, T.: Expected option returns. *J. Finance* **56**, 983–1009 (2001)
- Davies, R., Kat, H., Lu, S.: Fund of Hedge funds portfolio selection: a multiple-objective approach. Working paper ISMA Centre, University of Reading (2004)
- Demeterfi, K., Derman, E., Kamal, M., Zou, J.: A guide to volatility and variance swaps. *J. Deriv.* **6**(4), 9–35 (1999)
- Dittmar, R.: Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *J. Finance* **57**, 369–403 (2002)
- Doran, J., Ronn, E.: The bias in black-scholes/black implied volatility. Working Paper Florida State University (2004a)
- Doran, J., Ronn, E.: On the market price of volatility risk. Working paper Florida State University (2004b)
- Driessen, J., Maenhout, P.: The world price of jump and volatility risk. Working Paper University of Amsterdam (2003)
- Eraker, B., Johannes, M., Polson, N.: The impact of jumps in volatility and returns. *J. Finance* **58**, 1269–1300 (2003)
- Figlewski, S., Wang, X.: Is the ‘Leverage Effect’ a leverage effect? Working Paper Stern School of Business, New York (2000)
- Füss, R., Kaiser, D.: The tactical and strategic value of hedge fund strategies: A cointegration approach. *Financ. Mark. Portfolio Manag.* **21**(4), 425–444 (2007)
- Hafner, R.: Stochastic Implied Volatility. Lecture Notes in Economics and Mathematical Systems, vol. 545, Springer, Berlin (2004)
- Hafner, R., Wallmeier, M.: The dynamics of Dax implied volatilities. *Int. Q. J. Finance* **1**, 1–27 (2001)
- Hafner, R., Wallmeier, M.: Volatility as an asset class: European evidence. *Eur. J. Finance* **13**(7), 621–644 (2007)
- Harvey, C., Liechty, J., Liechty, M., Müller, P.: Portfolio selection with higher moments. Working Paper Duke University (2004)
- Illmer, S., Marty, W.: Return decomposition of absolute-performance multi-asset class portfolios. *Financ. Mark. Portfolio Manag.* **21**(1), 121–134 (2007)
- Jackwerth, J., Rubinstein, M.: Recovering probability distributions from option prices. *J. Finance* **51**, 1611–1631 (1996)
- Jiang, G., Tian, Y.: The model-free implied volatility and its information content. *Rev. Financ. Studies* **18**, 1305–1342 (2005)
- Lai, T.-Y.: Portfolio selection with skewness: A multiple-objective approach. *Rev. Q. Finance Account.* **1**(3), 293–305 (1991)
- Le Moigne, C., Savaria, P.: Relative importance of hedge fund characteristics. *Financ. Mark. Portfolio Manag.* **20**(4), 419–441 (2006)
- Lowenstein, R.: *When Genius Failed*. Random House, New York (2000)
- Moise, C.: Is market volatility priced? Working Paper University of Chicago (2004)
- Neuberger, A.: The log contract. *J. Portfolio Manag.* **20**(Winter), 74–80 (1994)
- Pan, J.: The jump-risk premia implicit in option prices: Evidence from an integrated time-series study. *J. Financ. Econ.* **63**, 3–50 (2002)
- Prakash, A., Chang, C.-H., Pactwa, T.: Selecting a portfolio with skewness: Recent evidence from US, european, and Latin American equity markets. *J. Bank. Finance* **27**(7), 1375–1390 (2003)
- Santa-Clara, P., Yan, S.: Jump and volatility risk and risk premia: a new model and lessons from $S + P$ 500 options. NBER Working Paper 10912 (2004)
- Sun, Q., Yuxing, Y.: Skewness persistence with optimal portfolio selection. *J. Bank. Finance* **27**(6), 1111–1121 (2003)
- Welsh, I.: Capital structure and stock returns. *J. Political Econ.* **112**, 106–132 (2004)
- Wilmott, P.: *Derivatives: the theory and practice of financial engineering*. New York et al. (1998)



Reinhold Hafner is Managing Director of risklab. He began his professional career as a bank apprentice. After that he founded and headed the company “Financial Software Systems.” In 1997, he joined RiskLab GmbH in Munich as a financial engineer, where he worked as Head of Market Analytics since 2001. In 2004, he received his Ph.D. degree on factor-based stochastic implied volatility models from the Economics Faculty at the University of Augsburg. Reinhold has published numerous articles in finance journals and has received several awards, including the Best Paper Award from the Swiss Society for Financial Market Research in 2003.



Martin Wallmeier graduated in Economics from the University of Münster. He obtained his Ph.D. degree in 1997 and his habilitation in 2002, both from the University of Augsburg. His dissertation and habilitation monographs deal with stock market anomalies and information contained in option prices. Since 2002, he is Professor of Finance and Accounting at the University of Fribourg. His primary research interests are in asset management, financial engineering, and firm valuation with a focus on empirical finance. His articles have been published, among others, in the *Journal of Computational Finance*, the *European Financial Management* journal, and the *European Journal of Finance*.