An Investor's Perspective on Volatility as an Asset Class: Evidence from the European Stock Market

Reinhold Hafner*, Martin Wallmeier†

Paper submitted to the 10th Symposium on Finance, Banking, and Insurance.

The paper would be presented by Martin Wallmeier.

Abstract

Volatility movements are known to be negatively correlated with stock index returns. Hence, investing in volatility appears to be attractive for investors seeking risk diversification. The most common instruments for investing in pure volatility are variance swaps, which now enjoy an active over-the-counter market. This paper investigates the risk-return tradeoff of variance swaps on the "Deutscher Aktienindex" (DAX) and EuroStoxx50 index (ESX) over the time period of 1995 to 2004. We synthetically derive variance swap prices from the smile in option prices, which we estimate using transaction data. Our objective is to analyze the relationship between index and variance swap returns and to draw conclusions for investors.

Empirically, the profile of log swap returns against log index returns on average resembles the payoff of a long put position. This highlights that variance swaps provide crash protection to investors. However, the market price of crash protection is suprisingly high. In line with previous U.S. evidence, we find a strongly negative volatility risk premium at the German as well as the European stock market. Its magnitude is not compatible with standard equilibrium pricing models. Thus, selling realized volatility seems to be a profitable strategy. Our backtests result in significant portfolio weights of the short volatility position during the sample period. Thus, our findings contradict recommendations of major investment banks and investment consultants to integrate long volatility positions into equity portfolios.

JEL classification: G10; G12; G13
Keywords: Implied Volatility; Smile; Variance Swap; Volatility Risk Premium

*Dr. Reinhold Hafner, risklab germany GmbH, Nymphenburgerstr. 112-116, D-80636 Munich, reinhold.hafner@risklab.de. This article represents the personal opinion of the author and does not necessarily reflect the opinion of risklab germany GmbH.

†Prof. Dr. Martin Wallmeier, Department of Finance and Accounting, University of Fribourg, Avenue de Beauregard 11, CH-1700 Fribourg, martin.wallmeier@unifr.ch. Corresponding author.

An Investor's Perspective on Volatility as an Asset Class: Evidence from the European Stock Market

First Version: May 2005.
This Version: May 2005.

Abstract

Volatility movements are known to be negatively correlated with stock index returns. Hence, investing in volatility appears to be attractive for investors seeking risk diversification. The most common instruments for investing in pure volatility are variance swaps, which now enjoy an active over-the-counter market. This paper investigates the risk-return tradeoff of variance swaps on the "Deutscher Aktienindex" (DAX) and EuroStoxx50 index (ESX) over the time period of 1995 to 2004. We synthetically derive variance swap prices from the smile in option prices, which we estimate using transaction data. Our objective is to analyze the relationship between index and variance swap returns and to draw conclusions for investors.

Empirically, the profile of log swap returns against log index returns on average resembles the payoff of a long put position. This highlights that variance swaps provide crash protection to investors. However, the market price of crash protection is suprisingly high. In line with previous U.S. evidence, we find a strongly negative volatility risk premium at the German as well as the European stock market. Its magnitude is not compatible with standard equilibrium pricing models. Thus, selling realized volatility seems to be a profitable strategy. Our backtests result in significant portfolio weights of the short volatility position during the sample period. Thus, our findings contradict recommendations of major investment banks and investment consultants to integrate long volatility positions into equity portfolios.

JEL classification: G10; G12; G13
Keywords: Implied Volatility; Smile; Variance Swap; Volatility Risk Premium

1 Introduction

In recent times, institutional and private investors have shown an increased interest in volatility as an investment vehicle. At first sight, investing in volatility appears attractive since volatility movements are known to be negatively correlated with stock index returns. Thus, adding volatility exposure to a portfolio of common stocks promises to improve risk diversification. In addition, past experience indicates that negative correlation is particularly pronounced in stock market downturns, offering protection against stock market losses when it is needed most.

If markets are efficient, the favorable characteristics of volatility are naturally reflected in higher prices of volatility instruments. Therefore, from an investor's point of view, the crucial question is which price surplus is charged to get access to this protection against stock market crashes.

Researchers distinguish between different notions of volatility, such as local volatility, forward volatility, realized volatility and implied volatility. Much research has been directed towards modeling the dynamics of these different variables and towards understanding how they are related. Theoretically, several kinds of volatility might qualify as trading objects in volatility claims. In practice, however, volatility trading is actually concentrated on *realized* variance or volatility.

A volatility trade is defined as a position which provides pure exposure to volatility alone without being affected by directional movements of the underlying asset. Classic methods for trading volatility, such as buying at-the-money (ATM) straddles, do not meet the demand of pure volatility exposure. They require frequent rebalancing to keep the options portfolio delta-neutral, which imposes high transaction costs.[1] Therefore, a more convenient solution should provide a payoff directly tied to measures of realized variance or realized volatility. The most common claim of this type is a variance swap.[2] At expiration, the buyer of this forward contract 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 forward stock holding and a static options portfolio including long positions in out-of-the-money (OTM) options for all strikes from zero to infinity.[3] 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, a robust replication strategy still exists, generating only small approximation errors in realistic settings.[4] It is important to note that these results are valid for arbitrary stochastic processes of volatility. Thus, using this approach, there is no need to assume a specific model of stochastic volatility in order to synthetically create a variance payoff. Certainly, this does not mean that the stochastic process of volatility does not influence the fair value of variance swaps. Rather, its influence is captured

[1] See Ederington/Guan (2002).

[2] Demeterfi et al. (1999) give an overview over the properties of variance swaps.

[3] See Neuberger (1994).

[4] See Carr/Wu (2004).

indirectly through the market prices of options. Given these prices, we do not need to know the stochastic behavior of volatility in order to infer the variance swap rate.[5]

Due to their payoff structure and the existence of a robust replication strategy, variance swaps constitute an ideal instrument for investing in volatility and for studying the market pricing of volatility risk. The fair value of the instrument, which is equal to the initial cost of the replication portfolio, can be inferred from the range of (plain vanilla) option prices over all strikes. The mean difference between the realized variance and the fair value of the variance swap rate quantifies an estimate of the variance risk premium. This approach was developed and first applied by Carr/Wu (2004). The authors synthesize variance swap rates for the most common US stock indices and 35 individual US stocks. Carr and Wu report strongly negative variance risk premia. Only a small part of the estimated premia can be explained by the negative correlation between index returns and volatility movements. The magnitude of the negative premium appears to be too large to be compatible with either the CAPM or the Fama/French (1993) three-factor-model.

A negative variance risk premium provides a possible explanation for the well-known finding of Jackwerth/Rubinstein (1996) that ATM implied volatilities of index options are typically greater than realized volatilities. Other studies, using a variety of methods, confirm the observation of a strongly negative volatility risk premium at the US stock market.[6] There is some evidence that its magnitude depends on the level of implied volatility and the time to maturity[7], the buying pressure for index puts[8] and the uncertainty of forward volatility[9]. The estimated premium shows significant temporal dependencies.[10] Apart from these observations, it is still an open question which economic factors and settings cause selling variance to be perceived as risky enough to charge a high premium. An alternative interpretation is that a high premium does not have an economic meaning but reflects systematic overpricing of out-of-the money calls and puts.

Since options over the whole range of possible strikes enter the replicating portfolio, variance swaps provide a direct link between implied and realized volatilities. The variance swap rate can be regarded as a representation of the implied volatility structure ("smile"). Thus, the large body of literature on the shape and dynamics of implied volatility surfaces is relevant for understanding the determinants of variance swap rates. An overview over the main findings of this line of research can be found in Hafner (2004).

To the best of our knowledge, this is the first empirical study on synthetically constructed variance swaps on European stock indices. We focus on the "Deutscher Aktienindex" (DAX) and the European stock index EuroStoxx50 (ESX) as the indices with the most liquid options

[5] For a derivation of the fair delivery price for variance swaps assuming a specific stochastic volatility model see, e.g., Howison et al. (2005).

[6] See Chernov/Ghysels (2000), Coval/Shumway (2001), Pan (2002), Bakshi/Kapadia (2003), Eraker et al. (2003), Driessen/Maenhout (2003), Doran/Ronn (2004b), Doran/Ronn (2004a), Bondarenko (2004), Moise (2004), Santa-Clara/Yan (2004).

[7] See Bliss/Panigirtzoglou (2004).

[8] See Bollen/Whaley (2004).

[9] See Carr/Wu (2004).

[10] See Bollerslev et al. (2005) and Santa-Clara/Yan (2004).

and futures trading. Our database consists of tick data for DAX options (introduced at the Eurex in 1994) over the time period 1995-2004 and for ESX options (introduced in 1999) over the years from 2000 to 2004.

Apart from providing results for Europe which can be compared to the U.S. evidence, our study contributes to the literature in the following ways: We estimate the volatility smile from transaction data instead of settlement prices. The latter are often unreliable for OTM options, and it is usually impossible to achieve a perfect matching to the stock index level. In addition, using tick data allows us to account for specific features of the indices (dividend and tax effects) which are important in Europe. We decompose variance swap prices into four components related to smile characteristics such as slope and curvature. In measuring the performance of variance swaps, we account for the non-normality of returns by using adjusted betas. Finally, our overall objective is to analyze the relationship between index and variance swap returns and to draw conclusions for investors.

The paper is organized as follows: In Section 2, we derive estimates of variance swap prices from the implied volatility structure of index options. We also decompose these prices to illustrate the influence of different smile characteristics. In Section 3, we analyze the distribution of variance swap returns and test, whether the sample mean returns are compatible with standard equilibrium models. Section 4 deals with the implications of our findings for investors. We determine "optimal" portfolio weights of variance swaps in a mean-variance framework and under power utility. The paper concludes with a brief summary.

2 Inferring Variance Swap Rates from the Smile in Option Prices

2.1 Valuation Formula

Given a stock price process S sampled on N equidistant points in time $0 = t_0 < t_1 < t_2 < \ldots < t_N = T$, where Δt denotes the length of the sampling interval, a *variance swap* is a forward contract on the realized variance of stock returns. Its payoff at expiry is:

$$VARS_T = \left(\hat{\upsilon}_T^2(N) - K_{VARS}\right) \cdot \mathcal{N}, \tag{1}$$

where $\hat{\upsilon}_T^2(N)$ is the realized variance (quoted in annual terms) over the life of the contract $[0, T]$, K_{VARS} is the *delivery price for variance* and $\mathcal{N}$ is the notional amount of the swap in euros per annualized volatility point squared. At expiry, the holder of a variance swap receives $\mathcal{N}$ euros for every point by which the realized stock variance $\hat{\upsilon}_T^2(N)$ exceeds the delivery price for variance K_{VARS}. The procedure for computing the realized variance has to be specified in the contract. In a typical contract, the stock price is sampled each trading day at the official close, i.e. $\Delta t = 1/252$, and the mean of daily stock returns is assumed to be zero. Formally, $\hat{\upsilon}_T^2(N)$ is usually defined as:

$$\hat{\upsilon}_T^2(N) = \frac{1}{\Delta t\,(N-1)} \sum_{i=1}^{N} \left(R_{t_i}\right)^2, \tag{2}$$

where $R_{t_i} = \ln(S_{t_i}) - \ln(S_{t_{i-1}})$ for $i = 1, \ldots, N$.[11]

Suppose the stock price process $S = \{S_t : t \in [0,T]\}$ is a pure diffusion with drift process $\mu = \{\mu_t : t \in [0,T]\}$ and volatility process $\upsilon = \{\upsilon_t : t \in [0,T]\}$:[12]

$$dS_t = S_t \mu_t dt + S_t \upsilon_t dW_t, \qquad \forall t \in [0,T], \tag{3}$$

where $W = \{W_t : t \in [0,T]\}$ is a one-dimensional standard Brownian motion. Volatility may be constant, deterministic or stochastic. For a given price history, the *realized continuously sampled variance* w_T over the interval $[0,T]$ is defined by:

$$w_T = \frac{1}{T}\int_0^T \upsilon_t^2 dt, \tag{4}$$

where the integral $\int_0^T \upsilon_t^2 dt$ is known as the (realized) *total variance* over the interval $[0,T]$. The continuously sampled variance w_T is a good approximation to the realized (discretely sampled) variance $\hat{\upsilon}_T^2(N)$ of daily returns used in the contract specifications of most variance swaps, i.e. $w_T \approx \hat{\upsilon}_T^2(N)$.[13]

In principle, valuing a variance swap is no different from valuing any other contingent claim. According to the risk-neutral valuation formula, the arbitrage free value of a variance swap at time $t \in [0,T]$ is the discounted expected value of the future payoff under the risk-neutral measure $\mathbb{Q}$. Since the *fair value of variance* is the delivery price K_{VARS} that makes the swap value zero today ($VARS_0 = 0$), it follows that:[14]

$$K_{VARS} = \mathbb{E}_{\mathbb{Q}}\left[w_T|\mathcal{F}_t\right] = \frac{1}{T}\mathbb{E}_{\mathbb{Q}}\left[\int_0^T \upsilon_t^2 dt\right]. \tag{5}$$

Generally, given a deterministic or stochastic volatility process, the expectation on the right-hand side of equation (5) can either be computed analytically or it must be numerically approximated.

Alternatively, the valuation can be based on a model-independent trading strategy that exactly replicates variance.[15] In fact, it can be shown that if the stock price process is a diffusion and interest rates are constant, then the fair value of variance is given by the value of an infinite strip of European options:

$$K_{VARS} = \frac{2}{T}e^{rT}\left(\int_0^{F_0(T)} \frac{1}{K^2}P_0(K,T)dK + \int_{F_0(T)}^{\infty} \frac{1}{K^2}C_0(K,T)dK\right), \tag{6}$$

where $C_0(K,T)$ and $P_0(K,T)$, respectively, denote the current market price of a put and a call option of strike K and maturity T, and $F_0(T) = S_0 e^{rT}$ is the stock's T-maturity forward price.

[11] Sometimes the maximum likelihood estimator of variance $\frac{1}{\Delta t N}\sum_{i=1}^{N}(R_{t_i})^2$ is used instead of the sample variance $\frac{1}{\Delta t(N-1)}\sum_{i=1}^{N}(R_{t_i})^2$. The difference is, however, usually very small.

[12] The drift process μ and the volatility process υ are required to be progressively measurable.

[13] Exact equality is given in the limit: $\lim_{N\to\infty} \hat{\upsilon}_T^2(N) = w_T$.

[14] Note that $\mathbb{E}_{\mathbb{Q}}[X|\mathcal{F}_0] = \mathbb{E}_{\mathbb{Q}}[X]$.

[15] See Carr/Madan (1997) and Demeterfi et al. (1999).

Defining $\sigma_0(K,T)$ as the Black-Scholes implied volatility at time $t=0$ of an option with strike K and maturity T, we can express the fair swap delivery price as:

$$\begin{aligned} K_{VARS} &= \frac{2}{T}e^{rT}\int_0^{F_0(T)} \frac{1}{K^2}P_{BS}\left(K,T,\sigma_0(K,T)\right)dK \\ &+\frac{2}{T}e^{rT}\int_{F_0(T)}^{\infty} \frac{1}{K^2}C_{BS}(K,T,\sigma_0(K,T))dK, \end{aligned} \tag{7}$$

where the suffix BS indicates that prices come from the Black-Scholes model.

2.2 Data

Our database comes from the joint German and Swiss options and futures exchange, Eurex.[16] It contains all reported transactions of options and futures

- on the German stock index DAX from January 1995 to December 2004 (option "ODAX", future "FDAX"), and
- on the European stock index Dow Jones Euro Stoxx 50 (ESX) from January 2000 to December 2004 (option "OESX", future "FESX").

The options are European style. At any point in time during the sample period, at least eight option maturities were available. However, trading is heavily concentrated on the nearby maturities. The contract values amount to 5 euros (ODAX) and 10 euros (OESX), respectively. Trading hours changed several time during our sample period, but both products were traded at least from 9:30 a.m. to 4:00 p.m.

The ODAX and OESX options have been more heavily traded than any other index option in Europe. The *average number of trades per day* of all DAX options went up from about 2'000 in 1995 to 3'000 in 2004, with a peak of 3'700 trades in the year 2000. The *average size* of ODAX trades varied between a minimum of 33 contracts in 2000 and a maximum of 54 contracts in 2004. This contrasts with the significantly higher mean trade size of about 115 (in 2000) to 240 (in 2004) OESX contracts. The *number of trades* in OESX options strongly increased from an average of 260 trades per day in 2000 to over 1'000 trades per day in 2004. Since 2003, the total number of traded OESX contracts has exceeded the number of ODAX contracts (63.8 million versus 40.5 million in 2004).

To calculate the implied volatility for each transaction, it is crucial to accurately match the corresponding underlying index level. We derive the stock index price $S_{t,n}$ on day t at minute n from the current price $F_{t,n}$ of the futures contract most actively traded on that day. The maturity of this contract, which is typically the nearest available, is denoted by T_F. The value $F_{t,n}(T_F)$ corresponds to the average transaction price observed in the T_F-futures contract in minute n on day t. To obtain the corresponding index level, we solve the theoretical futures pricing model

$$F_{t,n}(T_F) = S_{t,n}e^{r(T_F-t)}$$

[16] We are very grateful to the Eurex for providing the data.

for $S_{t,n}$. This procedure ensures that the time stamps of options and underlying prices diverge by not more than one minute.

The DAX is a performance index in which dividends are reinvested. However, as Hafner/Wallmeier (2001) point out, the DAX-calculation does not accurately consider corporate and personal taxes on dividends. Therefore, the dividends assumed to be reinvested might be different from the dividends the stockholders actually receive. This discrepancy can be interpreted as a (positive or negative) "extra" dividend not accounted for in the index. If we neglect this tax effect, the implied volatilities of calls and puts with the same strike diverge, falsely indicating violations of put-call parity. To avoid this problem, we increase or diminish all stock prices $S_{t,n}$ on trading day t for options with maturity T_O by an amount A_{t,T_O} such that ATM puts and calls have the same implied volatility. Thus, the adjusted underlying price used to calculate implied volatilities is defined by:

$$\check{S}_{t,n} = F_{t,n}(T_F)e^{-r(T_F-t)} + A_{t,T_O},$$

where A_{t,T_O} is a *put-call parity consistent correction amount.* This adjustment is separately done for each combination of option (ODAX or OESX), trading day and maturity.

We calculate risk-free interest rates from money market rates for 1, 3, 6, and 12 months. All interest rates are converted to continuously compounded rates and expressed in the daycount convention Actual/Actual. For an arbitrary time period τ, the τ-period risk-free interest rate r is obtained by linear interpolation between the available rates enclosing τ.

2.3 Estimating ODAX and OESX Volatility Smiles

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:

$$M(t, n, S, T, K) = \frac{\ln\left(\frac{K}{\check{S}_{t,n}e^{r(T-t)}}\right)}{\sqrt{T-t}}.$$

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

$$\sigma = \beta_0 + \beta_1 M + \beta_2 M^2 + \beta_3 D \cdot M^3 + \varepsilon, \tag{8}$$

where σ is the implied volatility, $\beta_i,\ i = 0, 1, 2, 3$ are regression coefficients, ε is a random error, and D is a dummy variable defined as:

$$D = \begin{cases} 0 & , \quad M \leq 0 \\ 1 & , \quad M > 0 \end{cases}.$$

The dummy variable accounts for an asymmetry of the pattern of implied volatilities around the (forward) ATM strike ($M = 0$). Typically, the "smile" is better characterized by a "sneer", with the negative relation between implied volatility and moneyness extending clearly beyond $M = 0$ (see, as an example, Figure 1). 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.

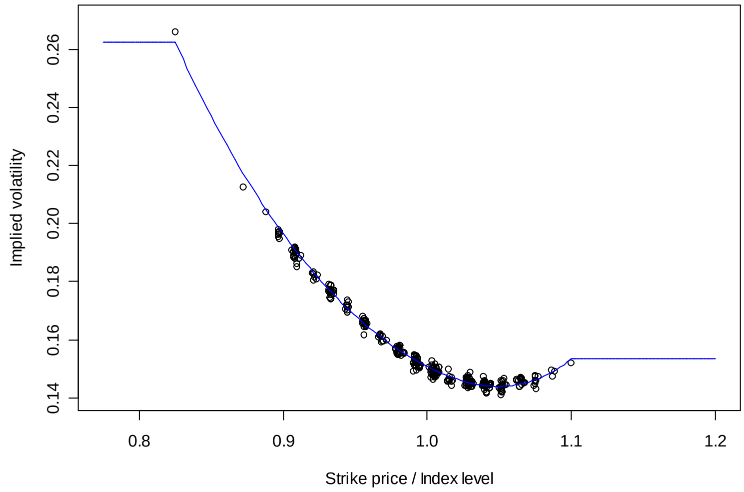

Figure 1: Scatterplot of implied volatility across moneyness for all trades of DAX options with a time to maturity of 42 days on December 10, 2004. The line shows the estimated regression function with constant extension outside the range of moneyness observations.

A quadratic or cubic regression without differentiating between $M \leq 0$ and $M > 0$ does not capture this increase. The regression function (8) is twice differentiable, which ensures that the corresponding risk-neutral density is continuous.

The implied volatility of deep in-the-money calls and puts is very sensitive to minor non-synchroneity of options and futures prices. Thus, the variance of the regression error term is supposed to increase as options go deeper in-the-money. To account for this heteroskedasticity of the disturbances we apply a weighted least squares estimation assuming that the disturbance variance is proportional to the positive ratio of the option's delta and vega.[17] This ratio indicates how an increase in the index level by one (marginal) point affects the implied volatility of an option, given the option's market price.

In view of the large number of intraday transactions it is clear that some extreme deviations occur representing "off-market" implied volatilities. They can, for example, be the result of a faulty and unintentional input by a market participant. In this case, the trade can be annulled if certain conditions are fulfilled. To exclude such unusual events we discard all observations corresponding to large errors of more than four standard deviations of the regression residuals where the standard deviation is computed as the square root of the weighted average squared residuals. We then repeat the estimation on the basis of the reduced sample until no further observations are discarded (trimmed regression). We examined the impact of this exclusion of outliers and found it to be negligible in all but very few cases.

[17] The delta and vega are computed using the implied volatility of the corresponding option. The delta of puts is multiplied by -1 to obtain a positive ratio.

The smile estimation according to equation (8) 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 τ calendar days, we linearly interpolate between the implied variances $\sigma_t^2(M, T_1)$ and $\sigma_t^2(M, T_2)$ of the two neighbouring maturities which are available (see, e.g., Wilmott (1998), p. 290). Formally:

$$\sigma_t^2(M, T = t + \tau \text{ days}) = \frac{T_2 - T}{T_2 - T_1}\sigma_t^2(M, T_1) + \frac{T - T_1}{T_2 - T_1}\sigma_t^2(M, T_2) \tag{9}$$

where

T : assumed (fictitious) expiration date τ calendar days in the future;
T_1 : latest available expiration date before T or equal to T;
T_2 : earliest available expiration date after T.

In this study, we choose $\tau = 45$ calendar days, because ODAX and OESX option series with lifetimes between 30 and 60 days are the most liquid contracts. This ensures an accurate estimation of the smile.

Average estimates of the smile regression coefficients are given in Table 1. The average is taken over the trading days of one year, where each trading day is represented by the option with a lifetime nearest to 45 calendar days. The yearly average of the ATM implied volatility β_0 varies between 12% and 35%. The other coefficients have the expected signs: $\bar{\beta}_1$ is always negative, whereas $\bar{\beta}_2$ and $\bar{\beta}_3$ are typically positive. The mean number of observations ($\bar{N}$) available for estimating the smile is markedly higher for ODAX compared to OESX. With the exception of 1995, the mean R^2-coefficient is always larger than 95%.

	ODAX						**OESX**					
	$\bar{\beta}_0$	$\bar{\beta}_1$	$\bar{\beta}_2$	$\bar{\beta}_3$	$\bar{N}$	$\bar{R}^2_{adj}$	$\bar{\beta}_0$	$\bar{\beta}_1$	$\bar{\beta}_2$	$\bar{\beta}_3$	$\bar{N}$	$\bar{R}^2_{adj}$
1995	0.144	-0.092	0.101	0.578	508	0.921						
1996	0.123	-0.158	0.075	1.964	503	0.967						
1997	0.233	-0.147	0.017	0.911	622	0.955						
1998	0.305	-0.190	-0.005	0.456	674	0.978						
1999	0.257	-0.184	0.011	0.615	862	0.985						
2000	0.234	-0.114	0.051	0.258	837	0.970	0.242	-0.137	0.067	0.813	65	0.962
2001	0.249	-0.136	0.059	0.313	806	0.965	0.258	-0.142	0.063	0.377	141	0.960
2002	0.346	-0.149	0.046	0.188	755	0.973	0.337	-0.162	0.048	0.344	177	0.969
2003	0.318	-0.148	0.041	0.292	719	0.978	0.291	-0.163	0.051	0.492	254	0.971
2004	0.187	-0.155	0.081	0.760	734	0.982	0.171	-0.180	0.081	1.547	236	0.982

Table 1: Yearly averages of estimated smile regression coefficients. Each trading day is represented by the option with the time to maturity nearest to 45 days.

We assume the smile function of equation (8) to be valid in a moneyness-range between the lowest and highest moneyness of all observations considered in the last step of the trimmed regression, i.e. excluding outliers according to the 4-sigma rule. Outside this range, we assume implied volatilities to be constant on the volatility level of the relevant moneyness boundary (see Figure 1). This corresponds to a conservative estimate of the fair values of options far in

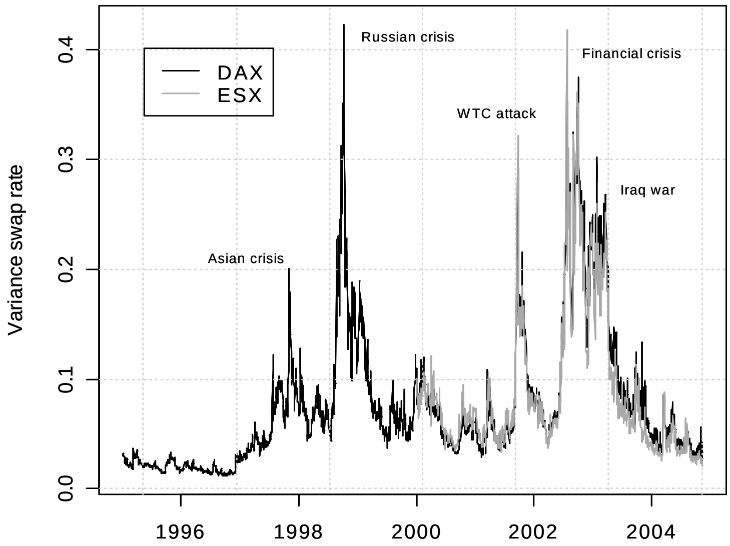

Figure 2: DAX and ESX variance swap rates over the time period January 1995 through November 2004.

or out-of-the-money. Other extrapolation techniques would provide higher variance swap rates and (even) lower variance returns.

2.4 Variance Swap Rates

Descriptive Analysis

Figure 2 displays daily swap rates for variance swaps on the DAX index (black line) and ESX index (gray line) with a time to maturity of 45 calendar days over the sample period January 1995 to November 2004 (DAX) and January 2000 to November 2004 (ESX). Table 2 presents some additional information on the distribution of these series.

Both series exhibit a strong mean-reverting behaviour with relatively well-defined lower and upper bounds. As can be seen from the graph, variance swap rates are exceptionally high in situations of crisis. For example, the swap rate of a 45 calendar days variance swap on the DAX initiated during the Russian crisis in 1998 amounted to 0.4232 (or 65.06% in volatility space) compared with a mean swap rate of 0.0770 (or 27.75%). Although the maximum DAX variance swap rate drops in the period from 1995 to 1999, swap rates were generally higher in the period from 2000 to 2004. Actually, the mean swap rate in the second period is roughly 55% higher than the mean swap rate in the first period. The high standard deviations of both series indicate substantial fluctuations over time. The estimated skewness and kurtosis suggest a highly non-normal distribution of DAX and ESX variance swap rates. Since the distributions are skewed to the right, we observe volatility shocks more often to the upside than to the downside.

	Period	Mean	Minimum	Maximum	Std.dev.	Skewness	Kurtosis
DAX	95-04	0.0770	0.0115	0.4232	0.0643	1.84	6.34
		(0.2775)	(0.1076)	(0.6506)			
	95-99	0.0605	0.0115	0.4232	0.0535	2.18	9.37
		(0.2450)	(0.1076)	(0.6506)			
	00-04	0.0941	0.0241	0.3746	0.0696	1.61	4.80
		(0.3067)	(0.1553)	(0.6121)			
ESX	00-04	0.0887	0.0205	0.4187	0.0658	1.78	5.81
		(0.2978)	(0.1435)	(0.6471)			

Table 2: Summary statistics for daily swap rates of DAX and ESX variance swaps with a time to maturity of 45 calendar days. Summary statistics in volatility space are given in parentheses.

Comparing DAX and ESX swap rates, the DAX swap rate generally exceeds the ESX swap rate. On average, the difference amounts to roughly 1 volatility point. The lower ESX swap rates may be explained by better diversification of the ESX index.

Influence of Smile Characteristics on Variance Swap Prices

Figure 3 shows the *spread* of the swap rate K_{VARS} over the forward ATM implied variance $\sigma_t^2(K = F_t(T), T)$ for DAX and ESX variance swaps. To control for level effects, we normalize the spread by dividing through the ATM implied variance.[18] As theory suggests, this *relative spread* is always positive. It ranges from about zero to roughly 100%. As is apparent from the figure, the relative spread fluctuates substantially. This results from a changing smile pattern over time. When the smile becomes more pronounced, the spread increases, and vice versa. Spread variations appear to have increased over time. When comparing the evolution of the variance swap rates in Figure 2 to the evolution of the spreads in Figure 3, we find little or no relationship. In fact, the correlation is 0.069 in the case of the DAX and -0.085 in the case of the ESX.

In a further analysis, we decompose the difference between the variance swap rate and the ATM implied variance (called *absolute spread*), into three components: a slope, a curvature, and an asymmetry component. The slope component is defined as the difference between the variance swap rate computed on the basis of the linear smile function $\sigma_t(M, \tau) = \beta_0 + \beta_1 M$ and the variance swap rate computed on the basis of the flat smile function $\sigma_t(M, \tau) = \beta_0$ (i.e. the forward ATM implied volatility). The difference between the two swap rates can hence be attributed to the linear smile term $\beta_1 M$. The curvature component $\beta_2 M^2$ and the asymmetry component $\beta_3 D M^3$ are defined analogously. Figure 4 shows the *average values* of the three components. Clearly, the slope component dominates. Almost 60% to 70% of the average absolute spread can be traced back to this component. An additional 20 to 30% of the average spread may be attributed to the curvature component. The remainder of roughly 10% of the average spread goes back to the asymmetry component.

[18] Note that the forward ATM implied variance equals the fair price of a variance swap for a flat smile.

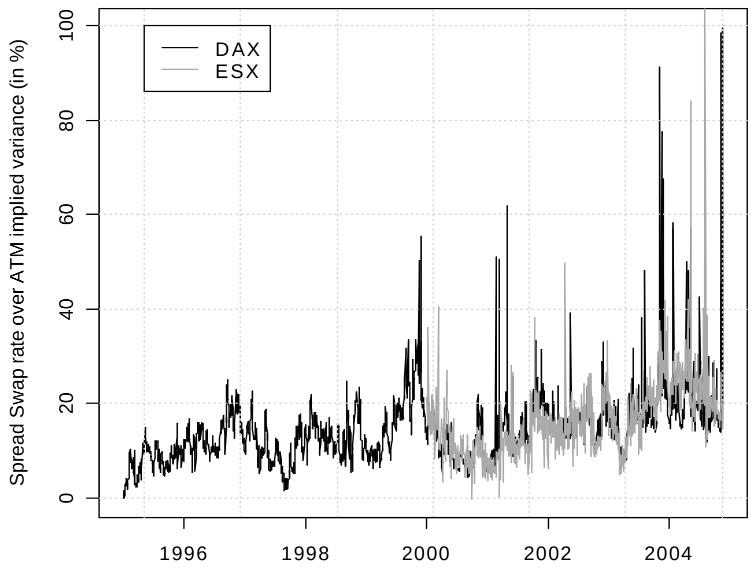

Figure 3: Spread of the variance swap rate over the ATM implied variance in % of the ATM implied variance. Sample period: 1995-2004 (DAX) and 2000-2004 (ESX).

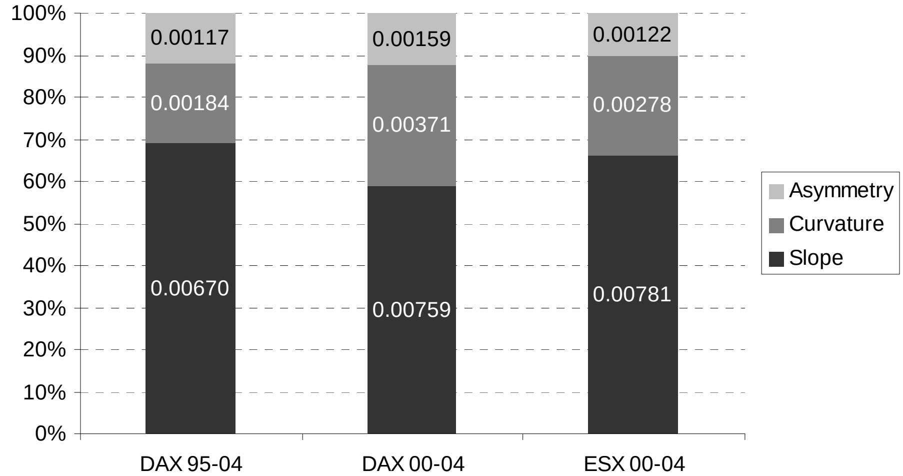

Figure 4: Decomposition of the *average absolute spread* of the variance swap rate over the ATM implied variance in a slope, a curvature, and an asymmetry component.

3 Empirical Evidence on Variance Swap Returns

3.1 Distribution of Variance Swap Returns

In this section, we investigate the distributional properties of variance swap returns. Using daily data over the period 1995-2004 (DAX) and 2000-2004 (ESX), respectively, we compute on each trading day $t \in \{1, ..., N_i\}, i \in \{DAX, ESX\}, N_{DAX} = 2472$, $N_{ESX} = 1209$, the following return measures for DAX and ESX variance swaps with a fixed time to maturity of 45 calendar days (i.e. $T = t + 45$ calendar days):

- *Payoff* $VARS_T$: the euro amount of money the holder of a long position in a variance swap with notional amount $\mathcal{N}$ =100 EUR receives at expiry T :

$$VARS_T = \left(\hat{\upsilon}_T^2(N) - K_{VARS}\right) \cdot 100 \text{ EUR}. \tag{10}$$

- *Discrete return* R_{VARS}: the simple net return of a long position in a variance swap contract:

$$R_{VARS} = \frac{\hat{\upsilon}_T^2(N)}{e^{-rT} K_{VARS}} - 1. \tag{11}$$

  This definition is motivated by the fact that if an investor creates the fixed part of the variance swap payoff by purchasing at the contract's outset $t = 0$ the proper replicating portfolio, the initial cost is $e^{-rT} K_{VARS}$, and the terminal payoff at expiry T is the realized variance over the contract's lifetime $\hat{\upsilon}_T^2(N)$.[19]

- *Log return* r_{VARS}: the continuously compounded return of a long position in a variance swap contract:

$$r_{VARS} = \ln(1 + R_{VARS}) = \ln\left(\frac{\hat{\upsilon}_T^2(N)}{e^{-rT} K_{VARS}}\right) = \ln\left(\hat{\upsilon}_T^2(N)\right) - \ln\left(e^{-rT} K_{VARS}\right) \tag{12}$$

By subtracting the (continuously compounded) risk-free rate from the discrete or log variance swap return, we obtain the *variance* or *log variance risk premium.*[20]

Table 3 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 payoffs, discrete returns as well as for log returns. To test their statistical significance, we construct robust t-statistics, using the serial-dependence adjusted Newey/West (1987) estimator for the standard deviation with a lag of 33.[21] The large t-statistics suggest that the mean returns are significantly different from zero.[22] The implication of this is that variance swap levels over-estimate subsequently realized

[19] Such a return is sometimes called "unleveraged", because the economic exposure underlying the swap contract is fully collateralized by the purchase of risk-free money market instruments. This terminology is, e.g., standard in the field of commodity futures indices.

[20] See, e.g., Carr/Wu (2004) and Bondarenko (2004). Note that Carr/Wu (2004), p. 28 label the variance swap payoff $VARS_T$ as the *variance risk premia.*

[21] Our choice of a fixed time to maturity of 45 calendar days corresponds to approximately 33 trading days.

[22] 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
$VARS_T$	Mean	-1.291	-0.988	-1.597	-1.894
	t-statistic Mean	-3.348	-2.220	-2.547	-3.210
	Median	-1.324	-0.691	-2.053	-2.133
	Minimum	-27.716	-27.716	-17.576	-23.899
	Maximum	20.016	15.246	20.016	18.266
	Std. dev.	4.342	3.521	5.019	4.771
	Skewness	0.979	-0.392	1.503	0.788
	Kurtosis	8.502	10.022	7.048	6.662
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 3: Summary Statistics for payoffs, 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.

variance. It is well-known that ATM implied volatility tends to be higher than subsequent realized volatility.[23] Due to the option's smile, variance swap levels are even higher than the ATM variance (see Figure 3), and so this 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 sellers of variance and are providing insurance to the market, require a significantly positive risk premium.

[23] See, e.g., Jackwerth/Rubinstein (1996).

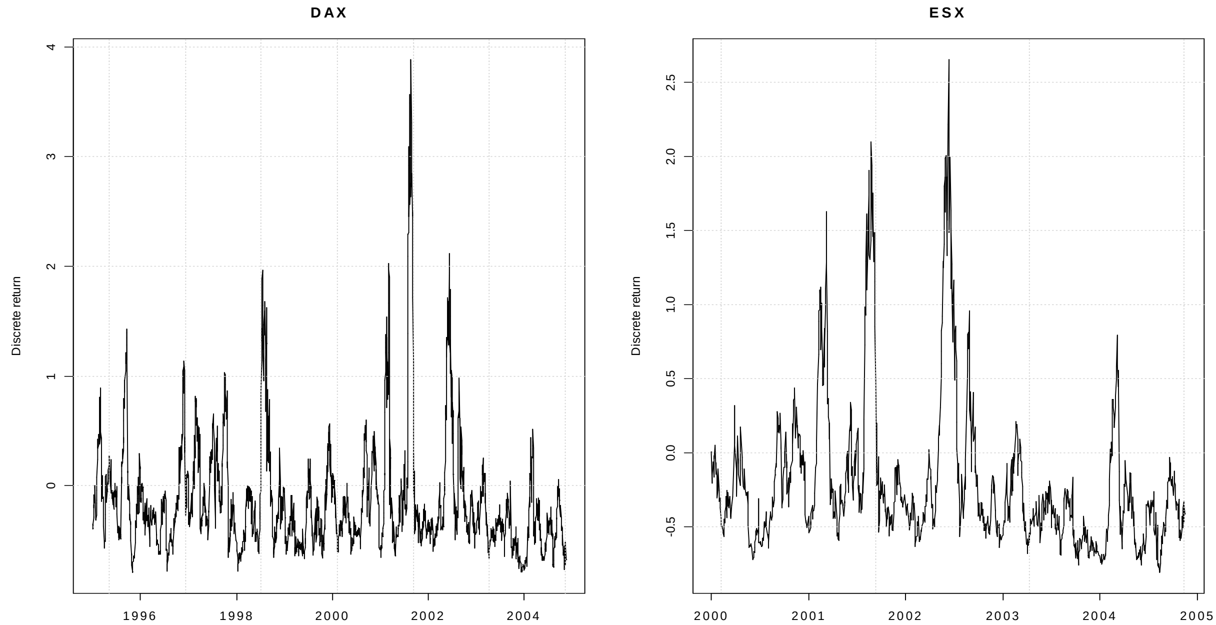

Figure 5: Discrete variance swap returns over over time.

The fact that selling variance swaps is profitable on average, does not mean that all sales of DAX or ESX variance have been profitable. In fact, the high standard deviations indicate considerable variation in the variance swap returns over time. As Figure 5 illustrates, there have been a number of high realized variance periods that were not anticipated in variance swap levels (most obviously the months in the forefront of major crises: July/August 1998, February 2001, August/September 2001, and June 2002). Selling variance swaps during these periods caused substantial losses. Comparing Figure 5 with Figure 2 and Figure 3, it appears that selling variance swaps has been more profitable in higher variance environments and periods with steeper smiles.

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 6 illustrates these results. 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. Comparing the distributions of DAX and ESX log variance swap returns, they are in most cases close to each other (see Figure 7). 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).

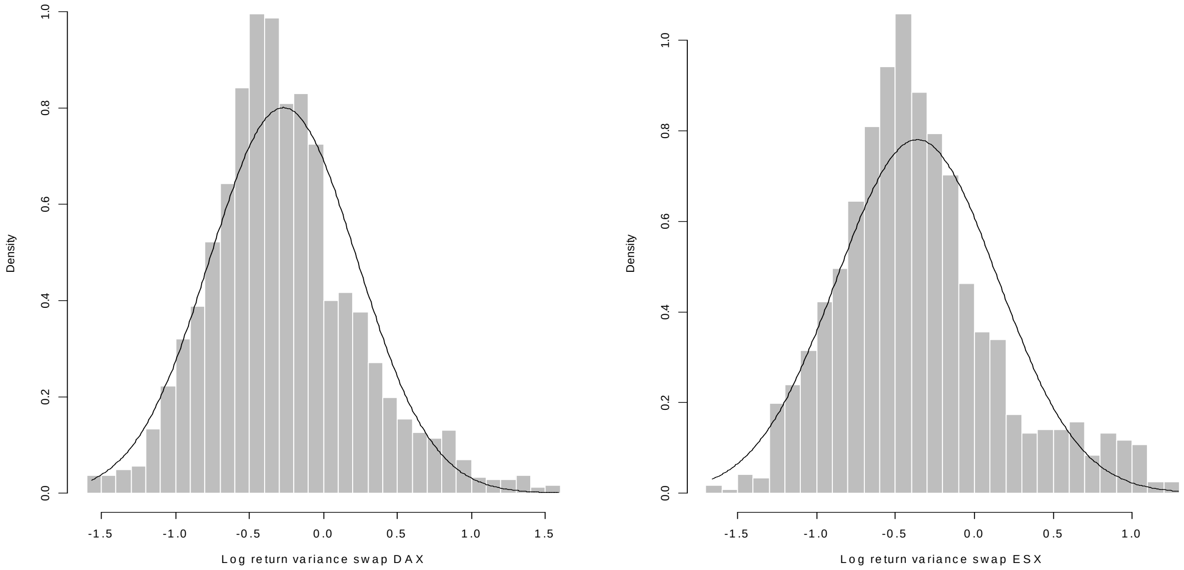

Figure 6: Histogram of 45 calendar day log returns of DAX variance swaps (left graph) and ESX variance swaps (right graph) over the sample periods 1995-2004 (DAX) and 2000-2004 (ESX).

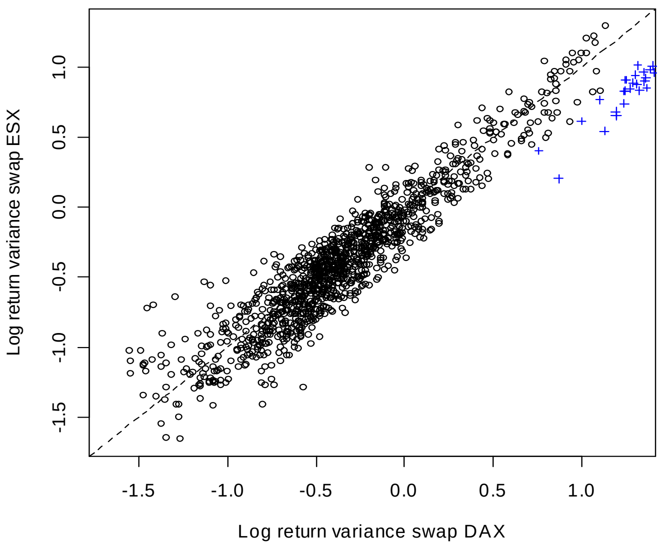

Figure 7: Log return variance swap DAX versus log return variance swap ESX within the period 2000-2004. Blue crosses mark all dates where September 11 is within the swap's time to maturity.

3.2 Relationship with Index Returns

As first discussed by Black (1976), the volatility of stock index returns tends to increase when stock prices drop. This negative correlation between past stock returns and future volatility is often referred to as "asymmetric volatility".[24] Early studies hypothesized that the higher volatility might be explained by the increase in debt-to-equity ratios induced by a stock market downturn.[25] However, this "leverage effect" explanation turned out to be incompatible with empirical observations. Analyzing the individual stocks in the S&P100 index and the index itself, Figlewski/Wang (2000) conclude that changes in leverage cannot account for the observed magnitude of the volatility-return correlation. There is considerable evidence that the impact of negative stock returns on volatility is much stronger than the response of volatility to positive returns.[26] Thus, the observed relationship might be better characterized as a "down market effect".[27]

The stylized fact that negative returns tend to be followed by high return fluctuations does not necessarily mean that the effect is return-driven. The causality could also run in the opposite direction from volatility to returns. If volatility is associated with systematic risk, an anticipated increase in volatility will immediately lower the fundamental stock price. According to this "volatility feedback" hypothesis, the stock price decline is caused by a reassessment of future volatility.

Under both of these explanations, we expect to find a negative correlation between the relative profit or loss from a long position in a variance swap and the contemporaneous stock return. If, for example, a decline in stock prices shortly after initiation of the variance swap leads to a higher subsequent volatility, this *negative* stock return is *good news* for the variance swap holder. The magnitude of the negative correlation depends on the temporal structure of the volatility-return relationship.

Figures 8 and 9 are scatterplots of variance swap returns versus index returns, conditioned on falling or rising stock markets (left and right graphs, respectively). In each scatterplot, the solid straight line represents a linear OLS regression of the form

$$r_{VARS,t} = a + br_{S,t} + \epsilon_t,$$

where a, b are regression coefficients, $r_{VARS,t}$ is the log-return of a long position in a variance swap starting at time t with a time to maturity of 45 calender days, $r_{S,t}$ is the concurrent stock index return, and ϵ_t is a random error with zero expected value. A new variance swap is bought each day, so that t corresponds to the trading days in the sample period.

The regression equation is separately estimated for all observations where $t \in \{t \mid r_{S,t} < 0\}$ ("down markets") and for $t \in \{t \mid r_{S,t} \geq 0\}$ ("up markets"). The estimated regression coefficients and t-statistics based on Newey/West (1987) adjusted standard errors are given in Table 4. In

[24] See, e.g., Wu (2001) with futher citations.

[25] See Black (1976) and Christie (1982).

[26] See, e.g., Haugen et al. (1991) and Bekaert/Wu (2000).

[27] See Figlewski/Wang (2000).

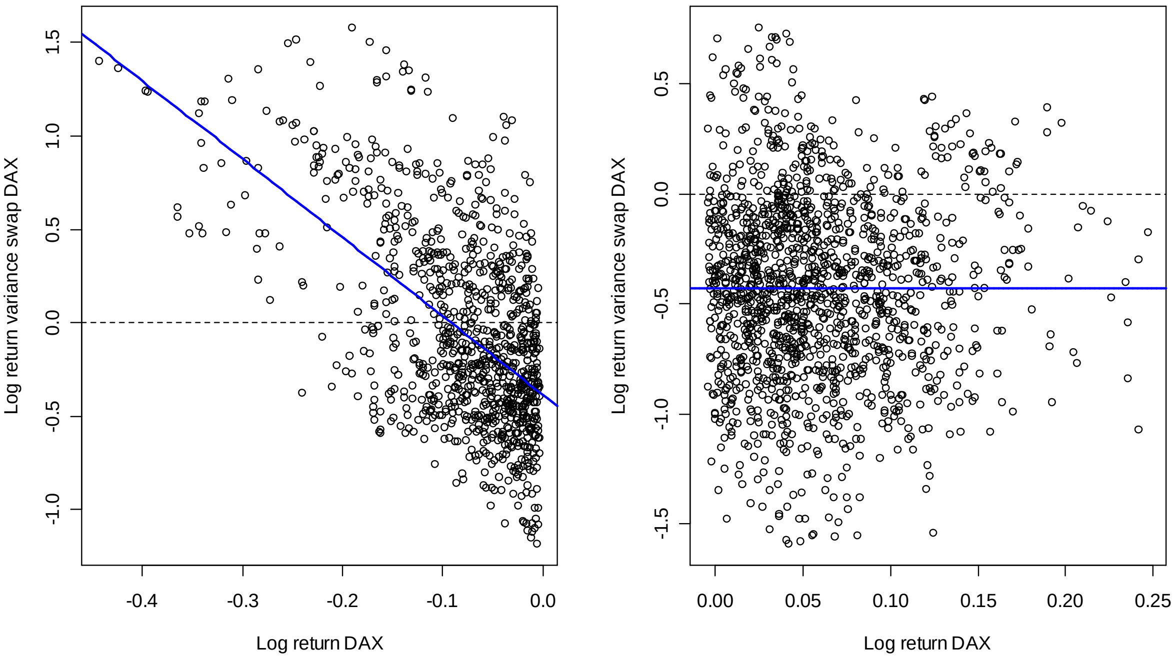

Figure 8: Scatterplot of log return variance swap versus log return DAX over all intervals of 45 days in the sample period from 1995 to 2004. Left (right): observations with negative (positive) DAX returns.

down markets, we find a significantly negative estimate of b for DAX as well as ESX variance swaps, ranging from -4.2 to about -5.0. The index return explains 30 to 38% of the variation of swap returns. In up markets, the estimated regression line is almost flat. In agreement with the descriptive return statistics in Section 3, all estimates of a are significantly negative. Thus, the holder of a DAX or ESX variance swap on average incurred a constant relative loss when stock returns were zero or positive. This loss can be interpreted as a premium he has to pay in order to receive rising gains in case of a stock price decline, especially a crash. The crash-intervals provide the highest returns to the variance swap holder. Yet, these returns are well captured by the linear regression line. Adding a quadratic regression term yields insignificant estimates for the slope coefficient of the squared index return. In this sense, the crash returns do not appear extraordinary.

In all, as Figures 8 and 9 show, the return profile in the sample period resembles the payout structure of an ATM long put. In the next section, we will analyze whether the magnitude of the "option premium" is compatible with standard equilibrium pricing models.

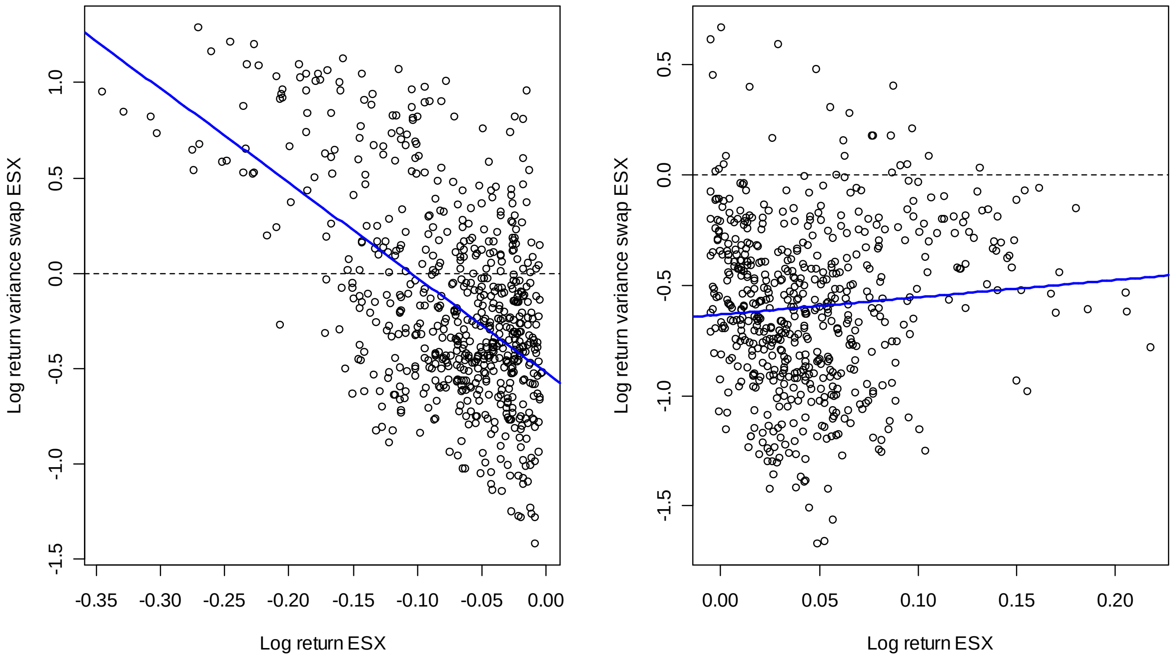

Figure 9: Scatterplot of log return variance swap versus log return ESX over all intervals of 45 days in the sample period from 2000 to 2004. Left (right): observations with negative (positive) ESX returns.

	$r_S < 0$			$r_S \geq 0$		
	a	b	R^2	a	b	R^2
ODAX 95-04	−0.3842 (−6.13)	−4.20 (−6.60)	0.3217	−0.4254 (−6.77)	−0.02 (−0.02)	0.0000
ODAX 95-99	−0.2724 (−2.38)	−4.58 (−6.31)	0.3806	−0.3032 (−3.34)	−0.73 (−0.58)	0.0057
ODAX 00-04	−0.4659 (−6.76)	−4.29 (−5.41)	0.3361	−0.6084 (−9.51)	0.95 (1.39)	0.0140
OESX 00-04	−0.5250 (−6.56)	−4.98 (−5.25)	0.3027	−0.6325 (−9.47)	0.78 (0.98)	0.0066

Table 4: Regression of variance swap returns on index returns in up- and down-markets. The return interval is 45 days.

3.3 Equilibrium Analysis

We apply three partial equilibrium models of expected asset returns in order to evaluate the appropriateness of the negative variance risk premia. The first model is the single period CAPM which concludes that the systematic risk of any asset should be equal to its marginal contribution to the overall portfolio risk.[28] Deviations from the security market line can be measured by Jensen's alpha

$$\alpha_i = \mathbb{E}(R_i^e) - \mathbb{E}(R_M^e)\beta_i,$$

where R_i^e denotes the excess return of asset i, R_M^e is the excess return of a market proxy, and β_i is the covariance between R_i and R_M, divided by the variance of R_M. We obtain an estimate $\hat{\alpha}_i$ of the outperformance of asset i by running the regression:[29]

$$R_{i,t}^e = \alpha_i + \beta_i R_{M,t}^e + \epsilon_{i,t}.$$

The CAPM rests on the assumption that either asset returns follow a multivariate normal distribution or that investors have mean-variance preferences. It is common knowledge that none of these assumptions is satisfactory, since discrete stock returns are bounded at -100% and the first two moments do not fully capture the relevant characteristics of return distributions. Since investors typically prefer positively skewed returns, CAPM-based performance measures are particularly suspect if a portfolio includes options or other assets with highly skewed return distributions.

Based on earlier research by Rubinstein (1976), Breeden/Litzenberger (1978), Brennan (1979) and He/Leland (1993), Leland (1999) presents a modification of the CAPM which does not require symmetrical distributions of asset returns. It is suitable for analyzing the performance of portfolios including options, which is important here due to the option-like return profile of variance swaps. The model assumes frictionless financial markets, continuous trading and iid returns of the market portfolio at each moment in time. These assumptions imply that market portfolio returns - but not asset returns - are log-normal, and the representative investor can be characterized by a power utility function. Equilibrium expected returns must then satisfy

$$\mathbb{E}(R_i^e) = \mathbb{E}(R_M^e)\beta_{L,i},$$

where

$$\beta_{L,i} = \frac{\mathrm{cov}\left[R_i, -(1+R_M)^{-\theta}\right]}{\mathrm{cov}\left[R_M, -(1+R_M)^{-\theta}\right]}. \tag{13}$$

The parameter θ measures the risk aversion of the representative investor. It is equal to the market price of risk defined as

$$\theta = \frac{\ln\left[\mathbb{E}(1+R_M)\right] - \ln(1+R_f)}{\mathrm{var}\left[\ln(1+R_M)\right]}.$$

Analogous to Jensen's alpha, the modified performance measure is given by

$$\alpha_{L,i} = \mathbb{E}(R_i^e) - \mathbb{E}(R_M^e)\beta_{L,i}.$$

[28] See Sharpe (1964), Lintner (1965), and Mossin (1966).

[29] See Jensen (1968).

We first obtain an estimate of $\beta_{L,i}$ from the return observations in the sample period (see equation 13). In the second step, we run a regression of $R^e_{i,t} - \widehat{\beta}_{L,i} R^e_{M,t}$ on a vector of ones according to

$$R^e_{i,t} - \widehat{\beta}_{L,i} R^e_{M,t} = \widehat{\alpha}_{L,i} + \epsilon_{i,t}.$$

This yields an estimate of the modified alpha $\widehat{\alpha}_{L,i}$.

Our third model is the continuous-time version of the CAPM developed by Merton (1973). It assumes that all asset returns are lognormally distributed.[30] In addition, either the investment opportunity set must be constant[31], or preferences must be restricted such that the investors are not interested in hedging against changes in the opportunity set. The latter condition is fulfilled if all investors have logarithmic utility.[32] Under these and further technical assumptions[33], the CAPM holds for continuously compounded rates of return:

$$\mathbb{E}(r_i) = r_f + [\mathbb{E}(r_M) - r_f]\, \beta_{c,i},$$

where beta is defined as

$$\beta_{c,i} = \frac{\operatorname{cov}[r_i, r_M]}{\operatorname{var}[r_M]}.$$

The risk-return tradeoff in this continuous-time version of the CAPM is the same as in the standard CAPM except that continuously compounded returns have replaced rates of return over discrete intervals of time. Thus, the estimation of the corresponding performance measure - the continuous-time equivalent of Jensen's alpha - is straightforward.

We apply the three models to measure the performance of the DAX variance swap and the variance swap on the European stock index ESX. As proxies for the market portfolio, we consider the index underlying the variance swap (Table 5) and the MSCI World Index (Table 6). The latter is a free float-adjusted market capitalization index that is designed to measure performance of global developed equity markets. It includes reinvested dividends.

The t-statistics given in parentheses in Tables 5 and 6 are based on robust standard errors according to Newey/West (1987) with a lag of 33. In the modified CAPM, the risk aversion coefficient θ is rather arbitrarily set to 2, but the results are insensitive to this choice.[34] The results for the DAX variance swap are shown for the full period (DAX 95-04) as well as for two subperiods of equal length.

A comparison of Tables 5 and 6 shows almost identical results for the two market proxies. The only difference worth mentioning is that the R^2-coefficients are typically lower when using the MSCI World Index.

[30] As a consequence, market returns are not lognormal. This is an important difference between the continuous-time CAPM by Merton (1973) and the modified CAPM proposed by Leland (1999).

[31] This assumption is questionable since it produces internal inconsistencies in the continuous-time CAPM as put forward by Rosenberg/Ohlson (1976).

[32] Logarithmic utility corresponds to a market price of risk θ equal to one.

[33] See Ingersoll (1987), Chapter 13.

[34] We computed the results for values between 1 and 10 to check robustness.

In all cases, beta is significantly negative on the 1% level. Its absolute value is highest in the Leland model and lowest in the standard CAPM. However, the beta adjustment does not fully explain the on average negative variance swap returns, as is apparent from the uniformly negative alphas in all three models. All alphas are significantly different from zero with the exception of $\widehat{\alpha}$ and $\widehat{\alpha}_L$ in the first subperiod. Whereas the standard CAPM and the Leland-modification provide similar alpha estimates, the underperformance of variance swaps appears larger in magnitude when measured against the continuous-time CAPM. The t-statistics (absolute value) for $\widehat{\alpha}_c$ are always highest among the three models. In sum, the return premia of DAX and ESX variance swaps realized in the sample period cannot be fully explained by the negative correlation to the market within standard equilibrium models.

	Single period CAPM			Leland-modification		Continuous-time CAPM		
	$\widehat{\alpha}$	$\widehat{\beta}$	R^2	$\widehat{\alpha}_L$	$\widehat{\beta}_L$	$\widehat{\alpha}_c$	$\widehat{\beta}_c$	R^2
DAX 95-04	−0.1055 (−2.32)	−3.19 (−4.29)	0.253	−0.1016 (−2.39)	−3.69	−0.2658 (−6.75)	−2.69 (−5.85)	0.229
DAX 95-99	−0.0352 (−0.64)	−3.04 (−3.79)	0.258	−0.0240 (−0.51)	−3.45	−0.1713 (−3.33)	−2.83 (−4.52)	0.206
DAX 00-04	−0.1859 (−3.11)	−3.63 (−3.44)	0.283	−0.1934 (−2.87)	−4.24	−0.3632 (−6.80)	−3.01 (−5.05)	0.306
ESX 00-04	−0.2465 (−4.98)	−3.74 (−3.85)	0.299	−0.2528 (−4.58)	−4.22	−0.4219 (−7.83)	−3.51 (−5.02)	0.289

Table 5: Equilibrium analysis of variance swap returns. Underlying index serves as market proxy.

	Single period CAPM			Leland-modification		Continuous time CAPM		
	$\widehat{\alpha}$	$\widehat{\beta}$	R^2	$\widehat{\alpha}_L$	$\widehat{\beta}_L$	$\widehat{\alpha}_c$	$\widehat{\beta}_c$	R^2
DAX 95-04	−0.1180 (−2.54)	−3.29 (−4.25)	0.183	−0.1166 (−2.60)	−3.68	−0.2728 (−6.64)	−2.93 (−5.75)	0.177
DAX 95-99	−0.0340 (−0.58)	−3.26 (−4.03)	0.229	−0.0264 (−0.55)	−3.56	−0.1619 (−2.97)	−3.32 (−5.31)	0.207
DAX 00-04	−0.2192 (−3.54)	−4.12 (−3.43)	0.203	−0.2282 (−3.14)	−4.60	−0.3880 (−6.54)	−3.51 (−4.71)	0.230
ESX 00-04	−0.2641 (−5.05)	−3.56 (−3.48)	0.225	−0.2721 (−4.59)	−3.98	−0.4369 (−7.45)	−3.36 (−4.53)	0.222

Table 6: Equilibrium analysis of variance swap returns. MSCI world index serves as market proxy.

4 Implications for Investors

4.1 Mean-Variance Efficient Portfolios

Assuming that the underperformance of variance swaps in the sample period does not reflect rationally priced risk factors other than beta, how many variance swap contracts should investors sell to profit from the apparent mispricing? 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.

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 implicitly "deleveraged" 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 10 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 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

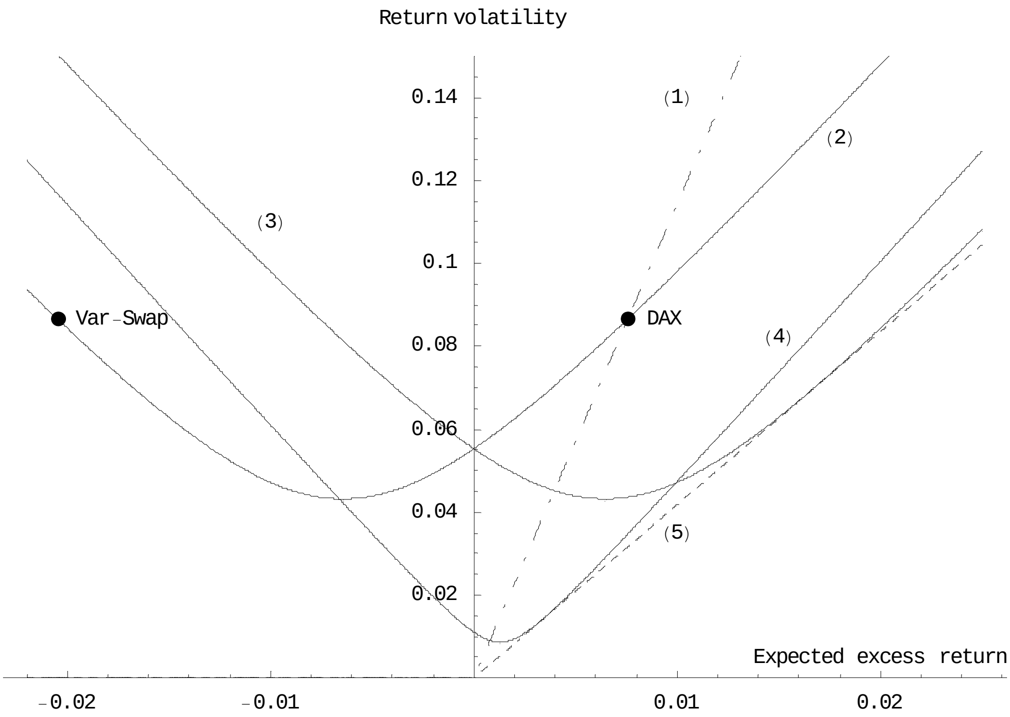

Figure 10: Mean variance analysis of DAX and variance swap investments. The expected returns and variances are estimated from the sample of return observations in all intervals of 45 days in the period from 1995 to 2004.

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 7 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 7. 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. 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

	SR_0	SR	x_{VARS}	x_S	x_{rf}	$LL: \frac{x_{VARS}}{x_S}$	$HL: \frac{x_{VARS}}{x_S}$
		Case 1: Base case					
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 7: Characteristics of mean-variance efficient portfolios.

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 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.

4.2 Backtesting Under Power Utility

Table 8 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								
	Base case			+2 STD			-2 STD		
α	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								
	Base case			+2 STD			-2 STD		
α	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								
	Base case			+2 STD			-2 STD		
α	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								
	Base case			+2 STD			-2 STD		
α	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 8: Optimal portfolio weights under power utility.

5 Conclusion

The idea of treating volatility as a separate asset class is attractive because of the highly negative correlation of volatility to stock market indices. In fact, several major investment banks and investment consultants have been advocating for some time to integrate long volatility positions in the form of variance swaps or forward-starting straddles into equity portfolios.[35] However, using daily series of variance swap prices, our results do not support this recommendation. In line with previous research for the U.S., we find a strongly negative volatility risk premium - defined as the difference between the log realized variance and the log variance swap rate - at the German and European stock market. Its magnitude is not compatible with standard equilibrium pricing models. Thus, selling realized volatility turns out to be a more profitable strategy. The performance characteristics for passive variance swap selling strategies compare favorably with many other asset classes. This may also explain the popularity of such strategies among hedge funds.[36] Our backtests reveal that the short volatility position received a significant weight in the optimal portfolios during the sample period.

The persistence of the negative premium over a period of several years indicates that it is driven by persistent economic factors. One important feature of variance swaps is the protection against stock market crashes. Our analysis shows that the profile of log swap returns against log index returns resembles the payoff of a long put position. This is due to a strong negative correlation in down markets and an almost zero correlation in up markets. It is still an open question, why the premium for buying this put protection is so high. This issue is closely related to the question why the smile in option prices is so steep. Although we did not directly aim at this puzzle, we show which characteristics of the smile are most important for the magnitude of variance swap rates. On average, about 60-70% of the absolute spread of the variance swap rate over the ATM implied variance can be attributed to the slope of the smile, another 10-20% to its curvature. The remainder is due to the asymmetry of the smile.

Our research design makes use of a comprehensive high-quality data set of tick-by-tick option prices. Following Carr/Wu (2004), we first estimate the smile and then derive smile-consistent prices of variance swaps. This approach offers promising extensions for further research. For instance, it can be used to determine the importance of the extrapolation of the smile function to OTM options, to test different hedging strategies, and to investigate the impact of trading costs. The approach also allows us to analyze the term structure of variance swap rates.

[35] See, e.g., Bowler et al. (2003).

[36] Bondarenko (2004) shows that the variance risk factor accounts for a considerable portion of hedge fund historical returns.

References

Bakshi, G./Kapadia, N. (2003), Volatility risk premiums embedded in individual equity options: Some new insights, *Journal of Derivatives* **11**(1), pp. 45–55.

Bekaert, G./Wu, G. (2000), Asymmetric volatility and risk in equity markets, *Review of Financial Studies* **13**, pp. 1–42.

Black, F. (1976), The pricing of commodity contracts, *JFE* **3**, pp. 167–179.

Bliss, R./Panigirtzoglou, N. (2004), Option-implied risk aversion estimates, *Journal of Finance* **59**, pp. 407–446.

Bollen, N./Whaley, R. (2004), Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance* **59**, pp. 711–754.

Bollerslev, T./Gibson, M./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.

Bowler, B./Ebens, H./Davi, J./Kolanovic, M. (2003), *Volatility - The Perfect Asset?*, Merill Lynch Global Securities Research and Economics Group, December.

Breeden, D./Litzenberger, R. (1978), Prices of state-contingent claims implicit in option prices, *JB* **51**, pp. 621–651.

Brennan, M. (1979), The pricing of contingent claims in discrete time models, *JF* **24**, pp. 53–68.

Carr, P./Madan, D. (1997), *Towards a Theory of Volatility Trading*, Working Paper.

Carr, P./Wu, L. (2004), Variance risk premia, *Working paper*, Courant Institute.

Chernov, M./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 Economis* **56**, pp. 407–458.

Christie, A. (1982), The stochastic behavior of common stock variances, *JFE* **10**, pp. 407–432.

Coval, J./Shumway, T. (2001), Expected option returns, *Journal of Finance* **56**, pp. 983–1009.

Demeterfi, K./Derman, E./Kamal, M./Zou, J. (1999), A guide to volatility and variance swaps, *Journal of Derivatives* **6**(4), pp. 9–35.

Doran, J./Ronn, E. (2004a), *The Bias in Black-Scholes/Black Implied Volatility*, Working Paper Florida State University.

Doran, J./Ronn, E. (2004b), *On the Market Price of Volatility Risk*, Working Paper Florida State University.

Driessen, J./Maenhout, P. (2003), *The World Price of Jump and Volatility Risk*, Working Paper University of Amsterdam.

Ederington, L./Guan, W. (2002), Why are those options smiling?, *Journal of Derivatives* **10**, pp. 9–35.

Eraker, B./Johannes, M./Polson, N. (2003), The impact of jumps in volatility and returns, *Journal of Finance* **58**, pp. 1269–1300.

Fama, E./French, K. (1993), Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **Vol. 33**, pp. 3–56.

Figlewski, S./Wang, X. (2000), *Is the 'Leverage Effect' a Leverage Effect?*, Working Paper Stern School of Business, New York.

Hafner, R. (2004), *Stochastic Implied Volatility*, Lecture Notes in Economics and Mathematical Systems 545, Springer -Verlag Berlin Heidelberg.

Hafner, R./Wallmeier, M. (2001), The dynamics of Dax implied volatilities, *International Quarterly Journal of Finance* **1**, pp. 1–27.

Haugen, R./Talmor, E./Torous, W. (1991), The effect of volatility changes on the level of stock prices and subsequent expected returns, *Journal of Finance* **46**, pp. 985–1007.

He, H./Leland, H. (1993), On equilibrium asset price processes, *Review of Financial Studies* **6**, pp. 593–617.

Howison, S./Rafailidis, A./Rasmussen, O. (2005), On the pricing and hedging of volatility derivatives, *Applied Mathematical Finance* **11**, pp. 317–346.

Ingersoll, J. (1987), *Theory of Financial Decision Making*, Rowman and Littlefield Publishers.

Jackwerth, J./Rubinstein, M. (1996), Recovering probability distributions from option prices, *JF* **51**, pp. 1611–1631.

Jensen, M. (1968), The performance of mutual funds in the period 1945 - 1964, *Journal of Finance* **23**, pp. 389–416.

Leland, H. (1999), Beyond mean-variance: Performance measurement in a nonsymmetrical world, *Financial Analysts Journal* **55**, pp. 27–36.

Lintner, J. (1965), The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* **47**, pp. 13–37.

Merton, R. (1973), Theory of rational option pricing, *Bell Journal of Economics and Management Science* **4**(Spring), pp. 141–183.

Moise, C. (2004), *Is Market Volatility Priced?*, Working Paper University of Chicago.

Mossin, J. (1966), Equilibrium in a capital asset market, *Econometrica* **34**, pp. 768–783.

Neuberger, A. (1994), The log contract, *Journal of Portfolio Management* (Winter), pp. 74–80.

Newey, W./West, K. (1987), A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**, pp. 703–708.

Pan, J. (2002), The jump-risk premia implicit in option prices: Evidence from an integrated time-series study, *Journal of Financial Economics* **63**, pp. 3–50.

Rosenberg, B./Ohlson, J. (1976), The stationary distribution of returns and portfolio separation in capital markets: A fundamental contradiction, *Journal of Financial and Quantitative Analysis* **11**, pp. 393–401.

Rubinstein, M. (1976), The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics and Management Science* **7**, pp. 407–425.

Santa-Clara, P./Yan, S. (2004), *Jump and Volatility Risk and Risk Premia: A New Model and Lessons from S+P 500 Options*, NBER Working Paper 10912.

Sharpe, W. (1964), Capital asset prices. a theory of market equilibrium under conditions of risk, *Journal of Finance* **19**, pp. 425–442.

Wilmott, P. (1998), *Derivatives: The Theory and Practice of Financial Engineering*, New York et al.

Wu, G. (2001), The determinants of asymmetric volatility, *Review of Financial Studies* **14**, pp. 837–860.