STOCKHOLM SCHOOL OF ECONOMICS
Master's Thesis in Finance
Spring 2008

Variance Risk Premiums on the S&P 500 and OMXS30

Magnus Petersson $^{\mathbb{P}}$
Joso Šarić $^{\mathbb{Q}}$

Abstract

In this thesis we analyze and quantify the variance risk premium, defined as the average difference between realized variance and its risk-neutral expectation. The risk-neutral expectation is obtained by using the notion of a variance swap, a contract that pays the difference between realized variance and a predetermined variance swap rate. Synthetically constructing variance swap rates from observed option prices on the S&P 500 and OMXS30 equity indexes we are able to calculate their variance risk premiums. The calculated variance swap rates and thereby our results are found to be robust to computational assumptions and correspond well to actual over-the-counter (OTC) variance swap quotes obtained from two investment banks. The average variance risk premiums are found to be strongly negative and statistically significant for the S&P 500 and on the short-term horizon for the OMXS30 index. In addition, we show that investors have historically been able to earn superior risk-adjusted returns by short-selling variance risk. Since the variance risk premium on the OMXS30 has previously not been studied, the results on this smaller Swedish equity index can be readily compared to the variance risk premiums on other indexes.

Tutor: Professor Tomas Björk
Presentation date: May 23^{rd}, 2008 at 15:15
Discussants: Stefan Björklund and Björn Johansson
Keywords: variance, risk premium, variance swap, option, OTC

Acknowledgements: Firstly we would like to thank our tutor Professor Tomas Björk for his help during the writing of this thesis. We are also grateful to Associate Professor Per-Olov Edlund for providing us with his valuable econometric insights. In addition we owe a debt of gratitude to Aaron Brask PhD, Global Head of Equity Derivatives Research at Barclays Capital and Stephen Hawinkels at Goldman Sachs, at whose discretion we obtained the data on OTC variance swap quotes. We would like to thank Lukas Larsson at OMX for providing us with Swedish option data. Finally we are grateful to our friends, Maxim Berglund, Fredrik Berglund Sonesson and Alexander Temml for valuable suggestions and comments. The authors alone assume responsibility for any errors or views contained in this thesis.

$^{\mathbb{P}}$mape82@gmail.com
$^{\mathbb{Q}}$joso.saric@gmail.com

Contents

1 Introduction

Investors face risk in all investment decisions. Volatility or variance is commonly used as the simplest proxy measure of "risk" by investors. However, volatility itself is uncertain. It is a well-known fact that stock markets tend to exhibit large volatility shocks. Events such as the Long-Term Capital Management (LTCM) debacle, 9/11 and the most recent financial crisis of 2007-2008 are examples of extremely high volatility states. Furthermore, these volatility shocks seem to be negatively related to index returns. Investors thus face risk not only from returns but also from the variance of the returns. It is therefore important to understand how investors deal with variance risk, both for derivatives markets and risk management purposes. In addition, it is an important factor for asset pricing and asset allocation decisions.

Financial contracts, primarily on variance, first appeared in the mid-1990s. It was not until the late 1990s that pricing and hedging methodologies were developed that enabled investors to trade these contracts on a larger scale. Today the over-the-counter (OTC) market of the most common volatility contract, the variance swap, has increased dramatically. A long position in this contract provides a payoff equal to the difference between future realized variance over a given period of time and a predetermined rate. It therefore enables market participants to take a position on pure variance compared to alternative methods such as straddles.

In this thesis we use the notion of a variance swap to provide a model-free approach to extracting and quantifying the price of variance risk. Similar to any ordinary forward contract, the variance swap rate that sets the value of the contract equal to zero, represents the risk-neutral expectation of future variance for a given period of time. Consequently the variance risk premium is defined as the difference between the actual realized variance and its risk-neutral expectation for any given period.

The attempts of examining and more notably, quantifying, the variance risk premium, can only be found relatively recently in the literature. Even though the number of academic contributions to and the interest of practitioners in this topic is growing, there are still areas to be explored. In this thesis, the analysis of the variance risk premium is restricted to two equity indexes. More specifically:

To what extent does the market price variance risk on the S&P 500 and OMXS30 indexes?

Since the market of variance swaps is OTC, transaction data is not publicly available to researchers. We therefore synthetically construct variance swap rates from publicly traded option prices on the S&P 500 and OMXS30. Given a measure of ex post realized variance, these synthetic rates are then used in the calculation of the variance risk premium.

A large share of the studies on the market price of variance risk found in the literature focuses on the short-term S&P 500 index options only. Our methodology in combination with our extensive dataset enable us to also analyze the longer-term maturities. In addition, it is our hope to contribute to the literature by analyzing the Swedish OMXS30 index as well. To the best of our knowledge, the market price of variance risk on the Swedish equity option market has never been investigated before using a model-free approach.

Our study also distinguishes itself by the extensiveness of our datasets. We have been able to retrieve more than eight years of daily options data on both the S&P 500 and OMXS30. More notably, we are able to compare the results we generate from the option data to a unique proprietary dataset of actually traded variance swap rates on the S&P 500. This leads us to conclude that the results from our method for determining the variance risk premiums produces results that are in line with the variance risk premiums extracted from actually quoted variance swap rates on the S&P 500. In addition, our results are robust against changes in computational assumptions giving further validity to our results.

In conclusion, we find that there exists negative variance risk premiums on both indexes, but only for short maturities on the OMXS30. In addition, it seems as if the market price of variance risk is overpriced in relation to the risks it entails. The consequence of this conclusion is that investors can possibly take advantage of this opportunity. Hedge funds, in particular, have been exploiting this empirical phenomenon in recent years. However, as we show in this thesis, taking advantage of the variance risk premium has not been – and should not be considered an entirely safe business in the future either.

2 Theoretical Framework

In this section, we commence the investigation of the variance risk premium by defining it more formally. We then proceed to the definition of realized variance and its risk-neutral expectation.

2.1 Defining the variance risk premium

The variance risk premium for any time horizon between t and T is defined as the difference between realized variance under the physical probability measure $\mathbb{P}$ and the expectation of future variance under the risk-neutral measure $\mathbb{Q}$, implied by the market. Also, as more formally derived in Appendix A.2, the variance risk premium, $RP_{t,T}$, for the period is defined as:

$$RP_{t,T} = \mathbb{E}_t^{\mathbb{P}} \left[\text{Realized variance}_{t,T} \right] - \mathbb{E}_t^{\mathbb{Q}} \left[\text{Realized variance}_{t,T} \right] \tag{1}$$

The time-series conditional mean of the realized return variance can only be measured historically and is by definition unknown ex ante. Therefore the variance risk premium can only be measured on an ex post basis. It is reasonable to assume that errors in expectations by market participants, if any, over the sample period average out to zero over time.

In order to estimate the ex ante risk-neutral expectation of realized variance, we use the notion of a variance swap. This approach was initially used by Carr and Wu (2007) for reasons outlined in Section 3.2.6 and has since then been used extensively in the literature to quantify the variance risk premium.[3]

The variance swap is a financial contract that pays the difference between the annualized realized variance (RV) during any time t and T (the "floating rate") as defined in the contract and a fixed swap rate (K_{VAR}) set at the outset. Formalized, the payoff at the predetermined maturity date is:

$$\left[RV_{t,T} - K_{\text{VAR}t,T}\right] \cdot \aleph \tag{2}$$

where $\aleph$ is the notional of the contract. Valuing a variance swap is no different to valuing any other derivative security. The arbitrage-free value of a variance swap at time t with maturity T equals the present value of the expected payoff under a given martingale measure. We employ the most commonly used approach in derivatives pricing of risk-neutral valuation, assuming a deterministic and constant short rate.[4] The discounted expected value of a variance swap under the $\mathbb{Q}$ measure equals:

$$\Pi_t = e^{-r(T-t)}\mathbb{E}_t^{\mathbb{Q}}\left[RV_{t,T} - K_{\text{VAR}t,T}\right] \tag{3}$$

where r is the continuously compounded risk-free interest rate at time t between t and T. The subscripts are henceforth left out for notational clarity. Even though it is not necessary, the fixed rate K_{VAR} is generally set as for the value of the contract to equal zero. As such, the swap rate de facto becomes the forward variance price. It follows that when the fair value of the contract is equal to zero, the fixed rate equals the risk-neutral expected value of future realized variance:

$$K_{\text{VAR}t,T} = \mathbb{E}_t^{\mathbb{Q}}[RV_{t,T}] \tag{4}$$

The rigorous theoretical framework and the practical implementation procedure needed to find the variance swap rate, and consequently the variance risk premium, is outlined in Section 2.3 and 5.2.

[3]The first working paper by Carr and Wu on this topic is from 2003. Other papers include amongst others, Bondarenko (2004, 2007), Egloff et al. (2007), Hafner and Wallmeier (2007) and Panigirtzoglou and Lynch (2003).

[4]See Appendix A.1 for further treatment of this topic.

2.2 Defining variance

In order to measure the risk premium as defined in equation (1) and also to be able to calculate the variance swap payoff of equation (2), we have to define a measure of variance. Assuming that the underlying S follows a diffusion process with drift μ_t and volatility σ_t:

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_t \tag{5}$$

where W is the standard Brownian motion, it then follows that the *continuously sampled realized variance* ($\nu_{t,T}$) over the time interval $[t,\ T]$ is defined as:

$$\nu_{t,T} = \frac{1}{(T-t)} \int_t^T \sigma_t^2 \ dt \tag{6}$$

where σ_t is strictly positive. *Continuously sampled realized variance* is commonly used in theory for increased analytical tractability.

2.3 Theoretically constructing the variance swap rate

Defining variance as in equation (6) implies that the risk-neutral expectation of realized variance at time t over the period t to T equals:

$$K_{\mathrm{VAR}t,T} = \mathbb{E}_t^{\mathbb{Q}} \left[\frac{1}{(T-t)} \int_t^T \sigma_t^2 \ dt \right] \tag{7}$$

Finding this risk-neutral expectation, either analytically or by numerical approximation, gives the variance swap rate. In order to provide the payoff equal to future realized variance, Dupire (1994) and Neuberger (1994) first showed the possibility of forming a replicating portfolio by using the *log contract* given a continuous process of the underlying. This instrument, that only provides payoff equal to the logarithm of the ratio between the underlying price at time T and initial time t, is however an abstraction used by financial mathematicians for increased analytic tractability and is not traded in reality.

Based on these insights, Carr and Madan (1998) and Demeterfi et al. (1999) construct a replicating portfolio of a continuum of European options that perfectly hedges the *log contract's* payoff. The fair value of future continuously sampled realized variance is given by the value of the replicating portfolio:

$$\begin{aligned} K_{\mathrm{VAR}t,T} = \frac{2}{T} \Bigg(& r(T-t) - (\frac{S_0}{S_*} e^{r(T-t)} - 1) - \log\left(\frac{S_*}{S_0}\right) \\ & + e^{r(T-t)} \int_0^{S_*} \frac{1}{K^2} P(K)\, dK + \int_{S_*}^{\infty} \frac{1}{K^2} C(K)\, dK \Bigg) \end{aligned} \tag{8}$$

where S_0 is the spot price, $P(K)$ and $C(K)$ are European call and put prices at strike K and expiry date T, r is the deterministic risk-free interest rate and S_* is an arbitrary stock price. If S_* is chosen to be the forward price F at time T, the above expression turns into:[5]

$$K_{\mathrm{VAR}t,T} = \frac{2}{T}e^{r(T-t)}\left(\int_0^F \frac{1}{K^2}P\left(K\right)dK + \int_F^\infty \frac{1}{K^2}C\left(K\right)dK\right) \tag{9}$$

That is, the future realized variance can be perfectly hedged with a static position in a linear combination of put and call options, weighted by the inverse of their strike squared. The intuitive explanation of this weighting scheme is that we attempt to create an options portfolio which has a constant sensitivity to variance. Because sensitivity to variance will increase for options with higher strikes as the underlying price increases, we eliminate the dependence on the underlying level S by underweighting options with higher strikes relative to options with lower strikes.

The only assumptions we make in arriving at the equation above are that: (*i*) the markets are arbitrage-free and frictionless, (*ii*) there exists a continuum of options with maturity T over an infinite range of strikes and (*iii*) the evolution of the underlying is diffusive, that is, any continuous price process is acceptable.

3 Overview of the Literature

In this section we provide an overview of the literature related to the variance risk premium. Firstly, the arguments for the existence of a risk premium are given. Following this, we briefly present the previous approaches of gaining the exposure to variance risk or quantifying its price. In light of these approaches, the advantages of our theoretical framework are explained. Finally the previous empirical findings are presented.

3.1 Theoretical reasons for the existence of a variance risk premium

If there is a significant difference between the price of variance in the market and the subsequently realized variance, the explanations for the existence of such a phenomenon are still ambiguous. A consistent spread could be attributed to either a systematic risk premium being paid to investors holding this type of risk by investors willing to insure themselves from it – or as inefficiencies in the market. In the case of the latter, the interpretation of an existence of a risk premium would be economically meaningless, stemming only from the overpricing of puts and calls. Since we do not find it probable that a potential risk premium can in its entirety be attributed to investor's systematic irrationality, we turn to alternative sources of explanations.

[5]The full derivation of equations (8, 9) can be found in Appendix B for the interested reader.

Low and Zhang (2005) claim it would be reasonable to assume that risk-averse investors dislike volatile states of the world. In addition, Bakshi and Madan (2006) state that finance theory suggests that risk-averse investors pay more attention to large losses. Investors who are exposed to short variance risk can bear unlimited losses as variance spikes, whereas they have a limited upside as variance cannot be negative. If variance risk introduces an additional non-diversifiable risk to investors' portfolios, affecting the overall portfolio risk unfavorably, investors ought to require compensation for this risk.

Further on as argued by Bakshi and Kapadia (2003a,b), out-of-the-money (OTM) put options provide hedges against large market declines. Amongst others, Bollen and Whaley (2004) and Bondarenko (2003), find that the buyers of these OTM puts pay an abnormal premium for this characteristic. Since these options are also a part of the replicating portfolio that hedges variance risk, this premium drives up the expectation of future variance, consequently giving rise to a negative variance risk premium.

A similar explanation for the existence of a variance risk premium comes from the findings of Bakshi et al. (2000) as well as Buraschi and Jackwerth (2001). The authors suggest that equity index options are not redundant securities, meaning they cannot be perfectly replicated by holding the underlying and a risk-free asset as often suggested in the literature. Since the option redundancy property is empirically violated, it implies option prices in reality incorporate other risk factors than the underlying price level. This allows for the interpretation of variance risk being priced into options. However, such a conclusion cannot determine the sign nor the magnitude of a possible variance risk premium.

Another explanation for a potential risk premium is that unexpected increases in volatility and equity returns are negatively correlated, argued by for example French et al. (1987), Bekaert and Wu (2000) and Wu (2001). Assuming the equity markets to be a significant part of the aggregate investor portfolio, it would be reasonable to expect investors willing to pay a premium for the characteristics of volatility protecting them against negative equity returns. In the example of a variance swap, investors should be more willing to receive the floating leg of the swap, as variance will tend to increase when markets decline. This demand for the floating leg should reasonably result in higher fixed rates and a negative risk premium between the two.

3.2 Alternatives to estimating the variance risk premium

3.2.1 Straddles

Straddles are one of the most common strategies, used in studies such as that of Coval and Shumway (2001) or Low and Zhang (2005), of capturing the variance of a given underlying. By purchasing a put and a call of the same strike and maturity, investors can profit from increases in volatility. However, since the sensitivity to variance of this strategy is not a constant function of the underlying, this method is susceptible to increasing errors as the underlying moves away from the strike price of the options. Therefore, it is impossible to exactly quantify the variance risk premium, making this method imprecise for our purposes.

3.2.2 Delta-hedged option portfolios

The approach of Bakshi and Kapadia (2003a,b) examines the profit and loss of a long position in a call option and a short position in its dynamically delta-hedged replicating portfolio. The combination of the two is neutral to directional movements, but sensitive to changes in volatility. However, this method heavily relies on a general stochastic volatility framework for modeling the delta of the hedged position. This method is therefore susceptible to misspecification of the framework, termed as "model risk" by Mougeot (2005). Additionally, the method of dynamic delta-hedging often incurs large transaction costs which make it prohibitively expensive in practice.

3.2.3 Parametric or structural approaches

Several studies including Guo (1998), Doran and Ronn (2008) as well as Bollerslev et al. (2007), use some form of parametric approach to extract the volatility risk premium from option prices in the market. By specifying the return and variance processes in advance, the parameters of the specified risk-neutral and physical processes, calibrated to real data, can be compared. The difference in estimated parameters can then be attributed to a variance risk premium. However, the parametric approaches suffer from potential misspecification errors in the processes or the option pricing model. As a consequence, rather strong assumptions are imposed on the estimation of the risk premium or its functional form by using these approaches.

3.2.4 Black-Scholes implied volatilities

The Black-Scholes at-the-money (ATM) implied volatilities have been shown by, for example Christensen and Prabhala (1998), to be efficient, though biased, estimates of future realized variance in the literature. However, Duffie (2001) amongst others, points out the flaws in the assumptions of the Black-Scholes model which empirically manifest themselves through the implied volatility smiles or smirks. The strongest objection towards using the Black-Scholes model is however that the model itself relies on constant future volatility. It is obvious that this is not the case in reality. Moreover, the existence of a volatility risk premium would be inconsistent with the model.

3.2.5 The VIX indexes

The VIX index is commonly used as the market expectation of future short-term volatility. The original VIX was launched in 1993 by the Chicago Board of Options Exchange (CBOE) and is calculated as the average of eight near-ATM Black-Scholes implied volatilies for options on the S&P 100 as to create a 22-trading day implied volatility.[6] The approximate conversion to trading days makes it impossible to compare the index to annualized realized volatility for the index and has according to Carr and Wu (2006) been criticized by both academics and practitioners. In addition, it suffers from the aforementioned flaw by being dependant on the Black-Scholes model.

In 2003 the way of calculating the index was changed. The new VIX index is based on the S&P 500 and only takes the market prices of options into account. In essence, it is a rather rough approximation to the model-free measure we use in our thesis. The main difference is that the new VIX only includes options with strikes actually traded in the market and as such lacks the necessary accuracy according to Andersen and Bondarenko (2007).

This inaccuracy is due to what Jiang and Tian (2007) term truncation and discretization errors to the true risk-neutral expectation. These two sources of errors will be explained in detail in Section 8.2.2. Due to these errors, the VIX index tends to systematically underestimate the true volatility when the overall volatility is relatively high and vice versa. As Jiang and Tian (2007) suggest, these errors misestimate the risk premium to an economically significant degree. Therefore, it is more theoretically appealing to use the model-free measure of future variance, rather than squaring VIX for our purposes.

3.2.6 The model-free approach to estimating the risk-neutral expectation

The model-free approach used in this thesis for extracting and directly *quantifying* the risk-neural expectation of future realized variance through variance swaps has several advantages over the other methods found in the literature. Variance swaps provide pure exposure to future realized variance, independent of the underlying price or other factors at any time during the contract's lifetime. This is a major benefit compared to classic straddle strategies. On the contrary to the commonly used Black-Scholes implied volatilities and the parametric approaches, no parametric structure, price or variance dynamics or option-pricing models are imposed. This does not imply that stochastic volatility is not taken into account in variance swap prices, but rather that these features are already incorporated into market prices of the replicating portfolio. In other words, the option prices in the market are the only sources for the expectation of future variance.

[6]The old VIX has now changed its name to the VXO index, but it is still calculated from the implied volatilities of the two puts and calls with strike prices around the ATM index level that are interpolated between the two nearest available maturities.

In order to form a replicating portfolio of the variance swap rate, a static position in an options portfolio is needed. As the replicating portfolio is static, there is no rebalancing, reducing the transactions costs over time compared to dynamic delta-hedging alternatives.

3.3 Previous empirical findings

In this section we provide an overview, to the best of our knowledge, of the literature dealing with the volatility or variance risk premiums up to this date. To make this information more transparent to the reader we present it in a table format.

Most authors estimate a negative volatility or variance risk premium. It is also apparent that a large share of the studies are performed on the S&P 500 for varying time periods and with daily data. This is not surprising as S&P 500 is considered as one of the most important market indexes and one of the most, if not the most, liquid options market in the world.

Author(s)	**RP estimation technique**	**Measure of realized vol. or var.**	**Sign, magnitude of RP**	**Underlying**	**Findings**
Bakshi and Kapedia (2003a)	Long option and short dynamically delta-hedged position	Standard deviation of log returns	-0.05% of the index level, over all strikes and maturities. -$0.43 per ATM option	S&P 500, daily (1988-1995)	RP is the dominant explanatory factor of delta-hedged differences.
Bakshi and Kapedia (2003b)	Long option and short dynamically delta-hedged position	Standard deviation of returns (sample mean equals zero)	-0.03% RP of the underlying stock price, over all strikes and maturities.	S&P 500 and 25 US stocks, daily (1991-1995)	Less negative RP for options on individual stocks than indexes. Idiosyncratic volatility not priced.
Bakshi and Madan (2006)	Model-free	Annualized 28-day st. dev. of daily returns	-2.31% average for 28-day RP	S&P 100, daily (1984-1999)	RP linked to higher-moments of physical distribution
Bollerslev, Gibson and Zhou (2007)	VIX	Sum of 5-minutes squared returns	-7.40% average for 30-day RP	S&P 500, tick-data and monthly (1990-2004)	Evidence of temporal dependencies of the RP.
Bondarenko (2004)	Model-free	Sum of squared daily returns	-26.66% discrete excess return on monthly variance contract	S&P 500 futures, daily (1988-2000)	Little correlation of RP to market returns.
Bondarenko (2007)	Model-free	Sum of squared daily returns	-31.24% discrete excess return on monthly variance contract	S&P 500 futures, daily (1990-2000)	Superior risk-adj. characteristics of short RP.
Carr and Wu (2006)	VIX squared	Annualized 30-day st. dev. of daily log returns (mean=0)	-1.59%, annualized	S&P 500, daily, (1990-2005)	CAPM cannot fully explain, time-varying RP.

Author(s)	RP estimation technique	Measure of realized vol. or var.	Sign, magnitude of RP	Underlying	Findings
Carr and Wu (2007)	Model-free	30-day sum of squared daily returns	-66% log excess return on 30-day variance contract (S&P 500) *(equivalent to)* \$2.7 loss per \$100 notional	Five US stock indexes and 35 stocks, daily (1996-2003)	Variance RP generated by independent risk factor. Superior risk-adj. returns from short RP.
Coval and Shumway (2001)	Zero-beta ATM straddles	n/a	-3% per week for S&P 500 -0.5% per day for S&P 100	S&P 500, weekly (1990-1995) S&P 100, daily (1986-1995)	Other risk factors besides market risk priced into option returns; volatility RP.
Doran and Ronn (2008)	Black-Scholes implied volatility	Sample st. dev. of the daily return on the futures over remaining life of option	Negative and significant estimate of the volatility RP parameter in the risk-neutral process across markets.	<1yr futures on energy commodities (1994-2004)	Given the parametric specification, RP generates the difference between implied and realized volatility.
Egloff, Leippold and Wu (2007)	Model-free	Sum of x day squared daily log returns, annualized (x = variance swap maturity)	-55% log excess return on 60-day variance contract *(equivalent to)* \$1.7 loss per \$100 notional	Proprietary weekly data from BofA; OTC quotes on variance swaps with maturities of 2, 3, 6, 12 and 24 months on the S&P 500 (1996-2007)	Term structure of variance swaps decomposed into instantaneous variance and a central tendency factor. RP more heavily priced for the former.
Guo (1998)	Parametric approach (Heston model)	60-day st. dev. of the log FX rate	Non-zero estimates of the volatility RP parameter.	USD/DEM exchange rate, daily (1987-1992)	RP is time-varying.
Hafner and Wallmeier (2007)	Model-free	45-day sum of squared daily log returns, annualized	-27% (DAX); -36% (ESX) log excess return on 45-day variance contract	DAX, EURO STOXX 50 futures, tick-data (1995-2004)	Short RP generates positive Jensen's alpha in CAPM regression.
Low and Zhang (2005)	Delta-neutral straddles	n/a	Negative returns, between -0.5% (CHF) and -2.4% (GBP) for 30-day.	GBP, EUR, JPY and CHF currency pairs, daily (1996-2002)	RP has a decreasing term structure.
Lynch and Panigirtzoglou (2003)	Model-free	Sum of squared high-intensity intra-day futures returns	Average annualized variance RP between -0.005 for interest rates and -0.01 for equities.	S&P 500, FTSE 100, eurodollar and short sterling futures	Magnitude of equity RP greater than short-term interest rate RP.

Table I: **Previous empirical findings regarding the variance risk premium in the literature**

4 Hypotheses

Based on our initial research question and the previous discussion in Section 3, a set of null hypotheses naturally follows. The hypotheses are simultaneously examined on both the S&P 500 and OMXS30 indexes.

Hypothesis 1: The variance risk premium is zero.

In order to be able to investigate the variance risk premium we must first establish its existence or lack thereof. If the difference on average is negative between the realized variance, as defined in equation (6), and the risk-neutral expectation as defined in equation (9), then the variance risk premium embedded in equity index options is negative, and vice versa. Under the null hypothesis, the sample average of equation (1) equals zero and investors are not compensated for holding variance risk.

However, we expect to reject this hypothesis in favor of evidence for negative risk premiums. Given the existence of a variance risk premium, we find both theoretical support from Section 3.1 and a general consensus in the empirical literature for its negativity as seen in Section 3.3. Regardless of the method, all studies, in particular of the equity markets, empirically find a negative variance or volatility risk premium. Therefore, the null hypothesis is tested against an alternative of a negative variance risk premium.

Hypothesis 2: Investors exposed to variance risk do not earn abnormal risk-adjusted returns.

The second hypothesis follows as a logical consequence of the rejection of the first hypothesis. If there are significant, negative variance risk premiums over time, it is of interest to examine whether the size of the risk premiums is fair relative to the risks carried. By construction, investors paying the floating leg of the variance swap and receiving the fixed leg (that is short-selling variance swaps) will have a positive payoff *on average* if the risk premium is negative. Even though the strategy of short-selling variance swaps may generate profits on average, this does not mean that *all* transactions generate profits. As a matter of fact when variance spikes unpredictably, such a strategy can incur very large losses – infinite in theory. On the other hand, variance cannot be negative, so the upside for investors is clearly limited. The question is then whether the price for the premium is justifiable relative to the risks investors are exposed to or not.

We find reasons to expect a rejection of the null hypothesis. As can be seen in Table I, previous studies suggest that investors are rewarded in excess of the risks. The market seemingly prices variance risk (too) heavily. Given the negative correlation between equity returns and volatility, it is of interest to determine whether the risk premiums are priced beyond their correlation to equity returns.

Hypothesis 3: The variance risk premium is constant over time.

Under the null hypothesis of a constant variance risk premium over time, the absolute size of the risk premium is hypothesized to remain unchanged over the sample period. Disregarding the case of a zero risk premium, it is of interest to examine whether the risk premium exhibits any temporal dependencies or not. The consequence of a time-varying risk premium is that investors are not only faced with having to estimate future variance but also with estimating the size of the risk premium charged for holding this risk over any given period.

Theory and previous empirical studies from Table I have found that the variance risk premium is time-varying. It would be reasonable to expect that the compensation for variance risk charged by investors changes depending on new information affecting the expectations of future variance. More specifically, the absolute magnitude of the risk premium should be positively correlated to the increases in variance levels to reflect the higher degrees of uncertainty in general. We expect our third hypothesis to be rejected as well.

5 Methodology

In this section we adapt the theoretical framework of Section 2 to more realistic circumstances by developing the practical implementation procedure needed to find the variance risk premium. We then proceed to defining three key empirical measures of the variance risk premium. Finally, we provide the more formal tests needed to answer the hypotheses in Section 4.

5.1 Measuring realized variance

Since asset prices are in most cases frequently, but not continuously observable, the notion of *continuously sampled realized variance* as defined in Section 2.2 does not exist in reality. It is de facto impossible to observe this idealized measure of variance. According to Hafner and Wallmeier (2007), the continuously sampled realized variance is well-approximated by various measures of *discretely sampled realized variance*. In fact, in almost all variance swap transactions, realized variance is defined as:

$$\nu_{t,T} \approx \widehat{\nu}_{t,T}(N) = \widehat{RV}_{t,T} = \frac{1}{\triangle t \cdot N} \sum_{i=1}^{N} (R_{t_i})^2 \tag{10}$$

where the returns on underlying are computed as log returns on the interval between i and $i-1$.

$$R_{t_i} = \ln S_{t_i} - \ln S_{t_{i-1}} \tag{11}$$

Equation (10) converges to the integrated variance of equation (6) as Δt approaches 0. Using $\Delta t = \frac{1}{252}$ where 252 is the assumed number of trading days in a year, equation (10) then becomes the annualized realized variance of the log returns over N daily observations.

This is standard practice for most variance swap contracts and is also the measure used in computing the payoff in the variance swap.

5.2 Constructing the variance swap rate in reality

The assumption of a continuum of options made in Section 2.2 does not hold in reality. In fact, there is only a finite number of options traded with discretely spaced strike prices. In order to calculate the fair value of future variance in practice equation (9) must be discretized and truncated:

$$K_{\mathrm{VAR}t,T} \approx \widehat{K}_{\mathrm{VAR},t,T} = \frac{2}{T} e^{r(T-t)} \left(\sum_{K_L \leq K < F} \frac{1}{K^2} P\left(K\right) \triangle K + \sum_{F \leq K \leq K_H} \frac{1}{K^2} C\left(K\right) \triangle K \right) \tag{12}$$

where K_L is the lower bound close to zero and K_H is a large upper bound chosen to minimize approximation errors. On this interval, the option strikes K are assumed to be equally spaced with a distance of $\triangle K$. This expression converges to the theoretical $K_{\mathrm{VAR}t,T}$ as $\triangle K$ and K_L tend to zero as well as K_H approaches infinity.

Traded options are only available on the interval $[K_{-}, K_{+}]$, where $K_L < K_{-}$ and $K_H > K_{+}$. Even on this shorter interval, options tend to be unevenly spaced with relatively large distances across strikes. We therefore must determine the option prices at each increment of $\triangle K$ on the interval $[K_{-}, K_{+}]$.

In order to solve this problem, we employ an interpolation technique between the traded options for each day t. Since it is standard practice to quote options in terms of Black-Scholes implied volatilities, we calculate these from available option prices. These volatilities are then used in the interpolation of the implied volatility function across strikes with a mesh fineness of $\triangle K$. These interpolated implied volatilities are then converted back into prices using the same Black-Scholes formula.[7] This procedure does not imply that we rely on the Black-Scholes model in any way, it only serves as a function that maps interpolated market implied volatilities into option prices.

The interpolation of the implied volatilities is made by using natural cubic splines, which is in line with Jiang and Tian (2007). This interpolation technique is chosen to ensure (i) an exact fit to observed implied volatilities and (ii) to ensure a smooth implied volatility function. The second point is important in order to ensure monotonically decreasing prices of calls over strike prices (the opposite holds true for puts), which is consistent with no arbitrage.

Another issue arises as there are no observed market implied volatilities on the intervals $[K_{-},\ K_L)$ and $(K_{+},\ K_H]$. We therefore hold the implied volatilities of the traded options with

[7]The Black-Scholes formula we employ is based on the forward prices and as such it takes dividends into account as well.

the lowest and highest strikes constant on these intervals. This is in line with the method used by Carr and Wu (2007) and as we show in Section 8.2.2, it is robust to alternative methods.

Following Carr and Wu (2007) we choose the interval $[K_L, K_H]$ to encompass eight standard deviations around the ATM forward level at maturity T. We use the ATM Black-Scholes implied volatilities as a proxy for the standard deviations. The number of discrete steps $\triangle K$ over the strike interval $[K_L, K_H]$, is set at 5,000 to generate a fine grid of Black-Scholes implied volatilities.

The exact maturity T of the variance swap is seldom available *per se* as there is only a limited number of option maturities traded at any time t. Unless there are options across several strikes available that have precisely the same option maturity as the variance swap, we must resort to an alternative approach. We therefore linearly interpolate between the computed variance swap rate of the closest maturity before T and the rate for the closest maturity after the expiry date of the variance swap to obtain the desired maturity.

5.3 Measuring the variance risk premium

We define three return measures used for variance swaps which have also been used most extensively in the literature. These are also concurrently empirical measures of the variance risk premium.

5.3.1 Monetary payoff

The payoff in monetary units that investors, obliged to pay the fixed rate K_{VAR} at the final maturity date T, receive or pay in net, for a notional of $\aleph = 100$ monetary units, equals:

$$\widehat{RP}_{payoff} = (\widehat{RV}_{t,T} - \widehat{K}_{\text{VAR},t,T}) \times \aleph \tag{13}$$

This measures the constant-notional exposure to a long position in variance swaps. The monetary payoff is also an estimate of the absolute magnitude of the variance risk premium in percentage units.

5.3.2 Discrete return

The discrete return between t and T for a long position in a variance swap is defined as:

$$\widehat{RP}_{discrete} = \frac{\widehat{RV}_{t,T}}{e^{-r(T-t)}\widehat{K}_{\text{VAR}t,T}} - 1 \tag{14}$$

where r is the continuously compounded risk-free rate between t and T at time t. In this case the initial capital "invested" is the discounted forward cost of future variance, K_{VAR}. This is a more realistic measure as the notional is seldom exchanged or held by the two counterparties in real variance swap transactions. The discounting takes place in order to take into account the interest generated on the known payment of K_{VAR} between the time from entering the swap t and final maturity date T. Discounting also takes the time value of the risk premium into account, as the variance swap payoff does not take place until the final maturity time T, as well as transforming the returns into excess returns. The main benefit of the *discrete variance risk premium* over the *monetary payoff* is that it provides the *relative* size of the risk premium.

5.3.3 Log return

Similar to the discrete return measure, the continuously compounded return of a long position in a variance swap is defined as:

$$\widehat{RP}_{log} = \ln \widehat{RV}_{t,T} - \ln \widehat{K}_{\text{VAR},t,T} - r\,(T-t) = \ln\left(\frac{\widehat{RV}_{t,T}}{e^{-r(T-t)}\widehat{K}_{\text{VAR},t,T}}\right) = \ln(1+\widehat{RP}_{discrete}) \quad (15)$$

In line with Bondarenko (2007), Carr and Wu (2007) as well as Hafner and Wallmeier (2007), we also define this as the *log variance risk premium.*

5.4 Hypotheses testing

5.4.1 Testing Hypothesis 1 of a significant variance risk premium

In order to test for the significance of a possible variance risk premium, we perform a Student's t-test using a robust t-statistic. Since we calculate the risk premium daily, overlapping periods of realized variance could entail the existence of autocorrelation. In order to test for the existence of autocorrelation in our data we use the Breusch-Pagan test. Further on we suspect heteroscedasticity in our data due to the intrinsic nature of variance itself. To test for this we employ the modified Breusch-Pagan Lagrange test developed by Koenker and Basset which is robust in the presence of autocorrelation and non-normality as suggested by Greene (2002). If the existence of these properties are found, we calculate the Newey and West (1987) autocorrelation and heteroscedasticity consistent standard deviations, where the number of lags corresponds to the number of calendar days during the maturity of the swap.

5.4.2 Testing Hypothesis 2 of abnormal, risk-adjusted variance risk premiums

In order to investigate the risks of a position in a short variance swap, it is important to analyze the return distributions and the risk-adjusted performance of this strategy. Only in this way can we possibly justify the size of the risk premium relative to the risks. We assume the sample distributions to be unbiased estimates of their true counterparts. It is important to bear this assumption in mind as extreme events are most often the cause of large losses. If these events do not occur in the sample period, it may appear that the risks of short-selling variance are lower than what they actually are.

This has to do with what Bondarenko (2003) terms the *Peso problem*, which refers to "a rare but influential event [that] could have reasonably happened but did not happen in the sample". Investors may actually correctly price the probability of crash occurrences. However, because of sample size limitations, the ex post sample realizations differ from the ex ante investor beliefs, causing the riskiness of the variance risk premium to appear understated. The *Peso problem* is however alleviated as the sample size increases.

As an attempt to account for the risks associated with holding variance risk, we define three risk-adjusted return measures. We simply scale the mean discrete and log risk premiums by some form of risk measure. The classic Sharpe ratio uses the standard deviation as a measure of riskiness. In order to take into account the tail risks of the distributions, we also compute value-at-risk (VaR) and conditional value-at-risk (CVaR) risk measures, based on the empirical distributions for varying confidence levels $\theta = \{0.1\%,\ 0.5\%,\ 1\%,\ 5\%,\ 10\%\}$.[8]

In line with Carr and Wu (2007) and also Hafner and Wallmeier (2007), we also employ an equilibrium model, the Capital Asset Pricing Model (CAPM), to be able to reconcile the observed variance risk premiums with theory and to justify their size. As previously mentioned in Section 3.1, studies have found strong evidence of the negative correlation between equity returns and equity volatility. Therefore, the CAPM should at least to some extent explain the variance risk premiums. Consequently, we estimate the following regression for the S&P 500 and OMXS30 variance risk premiums:

$$-\widehat{RP}_{t,T} = \alpha_i + \beta_i R^M_{e,T} + \varepsilon_{i,T} \tag{16}$$

where RP is the variance risk premium, corresponding to the excess returns on the variance swaps as previously defined, computed either in discrete or log terms. R^M_e represents the excess return on the market portfolio that is computed either discretely or with continuously compounding so as to remain consistent with the compounding regime of the regressand. The MSCI World index is used as a proxy for the market portfolio.

The estimate of the alpha is the excess return that investors would have obtained from neutralizing the market risk of the variance risk premium, also termed Jensen's alpha. In the ordinary least squares (OLS) procedure, we compute the standard errors according to Newey and West (1987) with the appropriate number of lags after having tested for the existence of autocorrelation and heteroscedasticity as described in Section 5.4.1.

[8] A more extensive treatment of these risk measures can be found in Appendix C.

5.4.3 Testing Hypothesis 3 of a time-varying variance risk premium

In order to understand the variance risk premium over time, we reside to the same test used by Carr and Wu (2007). We form the two following *expectation hypothesis regressions*, with the second regression in log terms:

$$\widehat{RV}_{t,T} = a + b\widehat{K}_{\mathrm{VAR}t,T} + \omega_{t,T} \tag{17}$$

$$\ln \widehat{RV}_{t,T} = c + d\ln \widehat{K}_{\mathrm{VAR}t,T} + \varphi_{t,T} \tag{18}$$

These two equations test if there are time-varying risk premiums. An estimated slope coefficient other than one indicates that the risk premium is time-varying and correlated with the variance swap rate. That is, a *constant* monetary or log risk premium must produce a perfect correlation between the risk-neutral expectation and the realized variance.

Under Hypothesis 3 of constant variance risk premiums, we test if the slope equals zero against the alternative of a slope coefficient that is greater than or less than zero. The significance of our estimated coefficients is tested by computing t-statistics with Newey-West standard errors, again after having tested for the presence of autocorrelation and heteroscedasticity.

6 Dataset

6.1 Dataset description

We collect daily option data for our two indexes from different sources. For the S&P 500 index, the dataset consists of all European options on the S&P 500 index traded on the CBOE at the closing of each trading day. The data is obtained from Optionmetrics, which is also the data source of Carr and Wu (2007). The sample period stretches from January 4^{th}, 1996 through June 29^{th}, 2007 which is Optionmetrics complete dataset. In total, the dataset comprises of 2879 days and 1,697,634 unique option observations. From an average of approximately 400 unique options per trading day during the sample period up until 2003, this figure increases to the double for the latter part of the sample period. This is due to the increased trading in options during recent years. Applying our cleaning procedure, as described in the next section, reduces the dataset by approximately 30%. The dataset also contains information on the bid-ask prices, implied volatilities based on the mid price calculated by Optionmetrics, traded volume and the respective strike and maturity of the option for each trading day. The spot level of the S&P 500 was obtained directly from the CBOE.

Since we were not able to obtain any data on futures on the S&P 500 that we could reconcile to our option sample, we calculate the forward levels for any maturity T ourselves. This is done with the help of the expected dividend yield provided by Optionmetrics and risk-free interest rates, together with the spot level. For each trading day in our sample, we extract the continuously compounded yield curve from the British Bankers' Association (BBA) London Interbank Offered Rate (LIBOR) rates for maturities up to one year, accounting for appropriate day count conventions. We use cubic interpolation of the yield curve to obtain the interest rate for maturity T.

Given the computed forward levels for the maturity T, we compute our own Black-Scholes implied volatilities from the average of the bid-ask prices for all the options in our sample. Plotting the implied volatilities over the strike levels produces for most days the characteristic smile or smirk shape.

The data on the Swedish equity index OMXS30, composed of the 30 most traded stocks on the Stockholm Stock Exchange, was provided directly by the exchange operator OMX. It comprises of daily closing prices of options and futures on the OMXS30 during the sample period January 2^{nd}, 1999 to December 28^{th}, 2007. This amounts to 2096 days with 240,284 unique observations in total which is the most extensive dataset we could obtain. Cleaning for inconsistent data reduces this number by approximately 20%.

For the OMXS30 index we do not need to calculate the forwards ourselves as futures data is available. The expected dividend yield is already incorporated into the futures prices. However, the futures price only equals the forward price at the settlement time. As we do not have settlement data, the closing prices may be different from the settlement prices, giving rise to a small discrepancy, which is assumed to equal zero. This equivalency of futures and forwards itself relies again on the implicit assumption of deterministic interest rates as proposed by Björk (2004), amongst others. In practice, the difference is very small and this minor assumption is frequently made by both practitioners and academics.

Given the forward price at maturity T that we obtain by interpolation from the futures curve, we can extract implied volatilities from the mid option prices. For any discounting purposes, we use the Stockholm Interbank Offered Rate (STIBOR) rates up to one year obtained from Datastream. These are cubically interpolated in order to match the correct maturity T, as in the case of the S&P 500. The index levels of OMXS30 are taken from Datastream, as well as the levels of MSCI World index that are used as the market index in the analysis.

In addition to our option datasets, we have obtained proprietary data on variance swap quotes on the S&P 500 from two of the largest investment banks and dealer-brokers in the OTC variance swap market. The dataset provided by Goldman Sachs contains daily variance swap rates on the 90, 180 and 365-day maturities between April 29^{th} 1997 and March 1^{st}, 2008. At the discretion of Barclay's Capital, we have also obtained daily variance swap quotes on the S&P 500 for the same maturities, but from March 4^{th}, 2002 through February 29^{th}, 2008. The convention is to quote the variance swap rates in units of volatility squared, even though the payoff is in variance units. To be able to compare these quotes to our results, we square these in order to obtain results in variance space. This transformation occurs at no loss.

6.2 Data processing

Such a large dataset as we posses is susceptible to errors. Some observations have been erroneously recorded as they are clearly unreasonable or they violate basic arbitrage restrictions. There are several reasons for this, but these are clearly beyond the scope of this study. In addition to the common cleaning procedures, as suggested for example by Bondarenko (2004, 2007), we remove all options with the same maturity if there are less than four options in total for the given maturity. It is also required that there is at least one option that has a strike price higher and at least one lower than the ATM forward level. This is done in order to increase the robustness of the implied volatility interpolation. For example, given only a small number of deep-OTM options, the interpolation technique may produce unreasonable results for the implied volatility function across all strikes. Appendix D provides a more thorough discussion of the complete list of data cleaning measure.

7 Empirical Results and Analysis

We calculate the three measures of the risk premium across varying time horizons for the S&P 500 and the OMXS30 for every available trading day with a valid K_{VAR}. We chose to compute the risk premiums for maturities of 30, 60, 90, 180 and 365 calendar days. On the contrary to a large share of previous studies that only focus on the short-term maturities (30-60 days), we also examine the long-term variance risk premiums.

We do not however compute the risk premium for longer maturities than 365 days. Firstly, longer-term maturities require longer time-series of the underlying index in order to compute realized variance, which effectively shortens our sample period. Secondly, after cleaning the data there are not enough valid options with maturities greater than 365 days. Therefore there is a trade-off between the sample period length, and consequently the significance of our results, and the duration of the longest maturity examined.

Using the Breusch-Godfrey test, with a lag length equal to the number of calendar days during the swap period, we reject the hypothesis of no autocorrelation in the time series of the risk premiums and in the regressions. In addition we find evidence of heteroscedasticity using the modified Breusch-Pagan Lagrange test developed by Koenker and Basset. Because of the strong presence of these properties, we modify the ordinary t-test statistica according to Newey and West (1987).[9]

[9]Because of their high significance, the results of these tests are not presented in the summary tables in Appendix E.

7.1 Are there significantly negative variance risk premiums?

7.1.1 S&P 500

It is clear from Table II that all measures of the risk premium are negative for the S&P 500 across all maturities. The average monetary payoff for a notional of 100 USD is -2.1 USD for the 30-day horizon and increases slowly over all investigated maturities to -1.6 USD for the 365-day. Similarly, the mean $RP_{discrete}$ progressively increases from -41% for the 30-day to -29% for the 365-day. This increasing pattern is also displayed by the mean of RP_{log}. The log measures are however consistently below the means of $RP_{discrete}$.

Even though both the discrete and log measures provide an estimate of the relative size of the premiums, there are several important considerations in comparing the two. Carr and Wu (2007) claim that the main benefit of the logarithmic transformation is to obtain a distribution closer to normality, which is true for leveraged derivatives such as the variance swap. However the drawback according to Bondarenko (2007) is that it also introduces a downward bias of the variance risk premium.

The reason for RP_{log} being lower than $RP_{discrete}$ is that the average log return is lower than the average discrete return, by construction. As Bondarenko (2007) claims, even if there is no risk premium the average log return will not be zero, but instead negative. Fortunately we are not faced with having to judge whether the risk premiums are truly negative, as all our results point to their unquestionable negativity.

Assuming simply that $\mathbb{E}_t\left(K_{\mathrm{VOL}}\right) \approx \mathbb{E}_t(\sqrt{K_{\mathrm{VAR}}})$, without adjusting for the convexity bias, we can also interpret our results in volatility space. Due to Jensen's inequality, the size of the "volatility" risk premium is slightly overestimated. Nevertheless, we obtain the risk premium to be on average approximately between -5 and -4 percentage units during the sample period. That is, investors are willing to pay a hefty premium of approximately up to a fourth of the annual realized volatility, which is around 20% on the S&P 500, to hedge variance risk.

Even though the variance risk premium is negative for 89% of all overlapping 30-day periods and 82% for our 365-days periods, there are a number of very large positive observations. As can be seen in Figure 1, large positive short-term discrete risk premiums of more than 200% tend to be clustered around the Asian financial crisis of 1997, the Russian crisis of 1998 and the fall of LTCM, as well as 9/11. These are events which caused spikes in realized variance and are impossible to predict. For long-term risk premiums, the stock market downturn of 2002-2003 provides a source of prolonged, unforeseen period of high variance. The outliers are not as large as for the short-term maturities, but remain in the range of 75-100% in discrete terms.

What is even more interesting is that the largest *negative* risk premiums tend to follow these extreme events. The reason for this may be that fear and uncertainty sentiments in the market following unexpected events drive up the market price of variance risk. Following these events, variance does not remain at the same high levels as initially expected; leading to the large negative RP. In times of distress, the market participants may irrationally overprice a certain source of risk. The effect seems to be even more pronounced for the short-term maturities. A formal test of a time-varying risk premiums will follow in Section 7.3.

Turning to the significance of our results, we see from the p-values in Table II that all our estimated variance risk premiums are significant on all reasonable levels. Because of the significant evidence of autocorrelation and heteroscedasticity, the Newey-West consistent standard deviations are used in the calculation of the robust t-statistica. We can therefore safely reject the Hypothesis 1 of a zero variance risk premium on the S&P 500.

7.1.2 OMXS30

The results on the OMXS30 are in line with those on the S&P 500 index. Although not as exemplary, our results on the OMXS30 indicate that there is evidence of negative variance risk premiums on the Swedish equity options market. From Table III, it is clear that the monetary payoff on a notional of 100 SEK is on average approximately -2.0 SEK on the 30-day horizon, increasing to -1.0 SEK for the 365-day maturity. The discrete risk premiums are also more negative for the short-term maturities, with -17% in discrete terms for the 30-day horizon. The increasing trend for the OMXS30 is however not as clear as on the S&P 500, but the size of the risk premiums visibly diminish for longer maturities.

The discrete variance risk premiums on the OMXS30 are less than half of their counterparts on the S&P 500. Only the 30-day monetary payoff, providing an absolute size of the risk premium, is approximately equal to that of the S&P 500. It would be reasonable to expect that the relative sizes of the risk premiums are equal as well, but since the absolute levels of variance are higher on the OMXS30, the discrete and log measures are depressed to even lower levels than on the S&P 500.

The only positive mean risk premium estimates we obtain on the OMXS30 are for the discrete measures on the 180-day and 365-day horizons. The reason for this is that the discrete measure is scaled by a smaller denominator K_{VAR} and is disproportionately affected by large positive realizations of the risk premium. None of them are however significantly negative *or* positive on any reasonable level. However the median risk premiums, are still greatly negative with -23% for the 180-day and -21% for 365-day horizons. The skewness estimates in Table III can possibly explain this observation.

Examining the time-series of the risk premiums more closely from Figure 2, it is clear that there are several extremely positive observations during the sample period which can all be attributed to the extremely volatile period during May 2006. After a long period of relatively low volatility, global macroeconomic concerns caused the volatility in the Swedish equity market to spike. With discrete excess returns in the range of 300-1000% on the short horizons and 150-250% on the long-term horizons, these observations clearly constitute extreme outliers. Out of curiosity, attempting to remove just the 10 largest observations associated with May 2006, the 30-day discrete risk premium is depressed to -21% and the 180-day and 365-day discrete risk premiums are virtually zero. The reason why the Swedish market prices variance risk less heavily than the S&P 500 can possibly be found in more than simply positive extreme outliers that are specific to OMXS30. One fundamental difference between the two indexes is that the Swedish index only has 30 constituents and as such is more exposed to idiosyncratic volatility than the well-diversified S&P 500.

Bakshi and Kapadia (2003b) find that idiosyncratic volatility risk is not priced. The "variance beta" of Carr and Wu (2007) also supports this hypothesis to a great extent. In other words, because the OMXS30 contains relatively more idiosyncratic volatility than the S&P 500, the size of the systematic component of the variance risk premium relative to the total realized variance is lower. A second explanation as to the difference between the S&P 500 and the OMXS30 may be the so-called "hedging pressure" hypothesis by Bollen and Whaley (2004). According to this hypothesis, in line with the findings of Bondarenko (2003), investors are prepared to pay a premium for being able to hedge stock market downturns through OTM puts on a major index such as the S&P 500, which is assumed to be the best representative of the aggregate investor portfolio. Evidence of this is that the characteristic Black-Scholes implied volatility smile practically did not exist before the stock market crash of 1987. Therefore, this specific hedging premium on the S&P 500 causes the variance risk premium to become more negative compared to other indexes such as the OMXS30.

Because of the evidence of autocorrelation and heteroscedasticity in the time series of the variance risk premium, we compute the robust Newey-West t-statistic. Unlike the S&P 500, the long-term maturities are not significant. We can therefore not conclude that there exist variance risk premiums on these horizons. The only consistently negative and significant risk premium across all three measures is for the 30-day maturity on the OMXS30. We therefore find evidence that it is possible to safely reject Hypothesis 1 of a zero risk premium for this maturity on the OMXS30. Henceforth, we only focus on this maturity for the OMXS30 in our further analysis.

7.2 Could investors have made abnormal risk-adjusted returns from short-selling variance risk?

The significant negative variance risk premium we find evidence of implies that a short position in a variance swap generates positive returns *on average*. This is because short-sellers of variance swaps provide insurance to the market by engaging in these transactions, taking on the risks associated with variance. In order to more intuitively analyze these risks, we restrict our focus in this section to the short-sales of variance risk, that is short positions in variance swaps.

7.2.1 Return distributions

In Figure 3,4 and 5 the return distributions from short-selling variance swaps on the S&P 500 and OMXS30 are presented. By visual inspection and from the descriptive statistics in Table II and III, it is obvious that none of the distributions conform to normality. The return distributions are skewed to the right, with a large amount of small, positive realizations. However, the distributions display a relatively small amount of extremely large negative realizations and "fat tails".

More formally, we test for normality using the Jarque-Bera and Kolmogorov-Smirnov tests. As can be seen in Table II and III we find that none of the distributions for either the S&P 500 nor the OMXS30 are normally distributed at any reasonable significance level. The log variance risk premiums conform better to a normal distribution, but are nonetheless rejected by these tests. The non-normality of the distributions is important to consider in calculating the risk-adjusted return measures, especially the Sharpe ratio.

7.2.2 The risk-adjusted variance risk premiums

In this section we discuss the risk-adjusted return measures in the same order as defined in Appendix C. In order to gauge the profitability in short-selling variance risk, we also compare these results to the risk-adjusted return measures for an alternative zero-cost strategy.

From Table IV, the annualized Sharpe ratios calculated for the discrete premium measures on the S&P 500 index range from 3.48 on the 30-day horizon and decrease steadily to 0.80 for the 365-day horizon. The ratios calculated on the log return measures are somewhat larger for reasons already mentioned in Section 7.1.1. The Sharpe ratio for the 30-day OMXS30 is only 0.79, which indicates that it would have been more profitable to short sell variance on the S&P 500.

From Table V the annualized Sharpe ratio for the excess returns on the MSCI World index is 0.14. At least using this classic risk-adjusted measure, it seems as it would have been far more lucrative to short-sell variance risk than borrowing to invest in the equity portfolio.

However, Sharpe ratios calculated for derivative payoffs should be treated with caution. Goetzmann et al. (2002) suggest that Sharpe ratios may be biased measures of the attractiveness of an investment strategy. In particular, this is true for return distributions that display a high degree of asymmetry – which indisputably is the case of short-selling variance swaps. The Sharpe ratio implicitly assumes returns to be linearly related to their volatility, which does not necessarily hold for asymmetric return distributions. The authors state that investment strategies generating high overall Sharpe ratios, with regular but modest profits, are commonly plagued by occasional crashes or the *Peso problem*. Asymmetric distributions and ex ante expectations of rare, but fatal losses, may explain the rather excessive Sharpe ratios from selling variance risk.

A similar explanation of the excessive Sharpe ratios from shorting variance risk is found in Taleb (2004). The author labels asymmetric payoffs from derivatives such as the variance swaps as “bleed” or “blowup” strategies. A long position in a variance swap “bleeds” by incurring small, but frequent losses – but largely gains from rare events. Investment strategies such as our short variance swap generate steady returns, but face “blowup” risks, thereof Taleb’s name for this type of strategy.

Turning to our VaR and CVaR risk-adjusted measures, our conclusions do not change. At least for our purposes, VaR and in particular CVaR, are superior measures of abnormal returns to the classic Sharpe ratios. The VaR and CVaR ratios obtained on the excess returns on the MSCI World index indicate that short-selling variance risk on the S&P 500 provides approximately two to three times less “tail risk” for any given excess return. This is true for all examined time horizons and confidence levels of VaR and CVaR. The figures are somewhat smaller for OMXS30, but short-selling 30-day variance still outperforms the MSCI World index by up to 30% for both the VaR and CVaR. That is, even when the most critically negative realizations are taken into account, short-selling variance still is as lucrative a strategy on the S&P 500 and the 30-day horizon on the OMXS30.

7.2.3 Equilibrium analysis

As can be seen in Table VI, we obtain positive and highly significant intercept and slope coefficients for the regression of equation (16) on all maturities for the S&P 500 and the 30-day horizon on the OMXS30. These results are consistent with the results obtained by Carr and Wu (2007, 2006) as well as Hafner and Wallmeier (2007).

The significant, positive alpha estimates imply that short-selling variance generates abnormal positive returns. They also suggest that the beta coefficients cannot fully account for a linear relationship between the variance risk premiums and the market risk premium. The size of the alpha coefficient is almost as large as the variance risk premium itself for all maturities on the S&P 500 and for the 30-day maturity of the OMXS30. This implies that the CAPM can only explain a very small part of the excess returns for the variance swaps.

The absolute magnitude of the beta coefficients, as well as the R^2 of the regressions, increases for longer maturities on the S&P 500. This indicates that even though the explanatory power of the regressions are generally low, the variance risk premiums for longer-term maturities seem relatively more compatible in an equilibrium setting.

Assuming that the CAPM holds, an estimated alpha value that is significantly different from zero implies investors do not exploit systematic mispricing opportunities in the market.

We do not believe this to be the case and the only remaining interpretation is that the variance risk premium constitutes a separate risk source independent of the market premium. This is consistent with the findings of Bondarenko (2007) explaining how the variance risk premium is to a large extent priced beyond its correlation with the market return.

The most probable reason for this finding is that the CAPM relies on investors having mean-variance preferences or on the normality for asset returns. Neither of these assumptions hold in our case. Clearly, the variance risk premiums are not nearly normally distributed and cannot be described by only the first and second moments.

7.2.4 How could shrewd investors have exploited the abnormal risk-adjusted risk premiums?

In order to demonstrate how investors can exploit the apparent overpricing of variance risk, we devise a relatively simple trading strategy by rolling the short 30-day variance swaps forward. That is, we enter into new short variance swaps contract at each expiry day of the preceding variance swap. If there is no available K_{VAR} for that day, we take the following date with available quotes.

Because this strategy comes at no cost, just as an ordinary forward contract, we devise a comparable zero-cost strategy for an equity portfolio. We lend a certain amount which is invested into the MSCI World index, simultaneously accounting for the costs of borrowing. Essentially, we compute the excess payoffs for a fixed investment on the MSCI World index.

In Figure 6 and 7 we show the cumulative payoffs of these two strategies over time and in Table VII we present the mean payoffs and standard deviations. The notional of the variance strategy is set so as for the volatility of the two strategies to equal. From the figures it is clear that short-sales of variance provides superior payoffs over equities. Since we have non-overlapping periods, we avoid the inherent problems of autocorrelation found in the time-series of the risk premium. Compared to our previous risk-adjusted results for overlapping periods, this zero-cost strategy is more robust to this error source.

Similar strategies to the one present have also been popular by the investment community, in particular by hedge funds. Bondarenko (2004) finds variance risk to account for a large share of many hedge fund returns. It is therefore appealing to attribute short-selling of variance risk to many famous hedge fund demises.

In theory, any single-period payoff on a short variance swap can be infinitely negative. This is often avoided in reality as there are often caps to the maximum payoff from variance swaps. Even in this case, investors can potentially lose more than the size of the notional. Since this has not happened during our sample, the actual risks may appear to be understated. As with all investments, past performance is not necessarily indicative of future performance.

The section above exemplifies the superiority of the risk-adjusted returns of short-selling variance risk that we have found evidence of in the previous sections. Both classical and more sophisticated risk measures indicate that the historical risk-return tradeoff has been more favorable than a comparable zero-cost investment in the equity portfolio. In addition, shorting variance risk is not consistent with an equilibrium framework, providing strongly positive Jensen's alpha and thereby generating abnormal returns in a CAPM framework. We therefore reject Hypothesis 2 of no abnormal risk-adjusted returns from short-selling variance risk for the S&P 500 and the OMXS30.

7.3 Are variance risk premiums time-varying?

In Figure 1 and 2 we plot the variance risk premiums over time for both short-term (30-day) on the S&P 500 and OMXS30 and long-term (365-day) maturities for the S&P 500 only. Even though it is difficult to see any trends over time, it is clear that the risk premiums are not constant over time.

The more formal results from the expectation hypothesis regressions (17,18) are presented in Table VIII. The slope coefficients are significantly less than one in line with Carr and Wu (2007), but positive and decreasing across maturities. As the slope coefficients are different from one, we can also draw the conclusion that the risk premiums are time-varying. As expected, the size of the risk premium seems to be positively related to the size of the variance swap rate. The decreasing slope coefficients over increasing maturities can be explained by the inability of the market to accurately price variance risk for longer time horizons.

Because realized variance and the variance swap rate show evidence of asymmetry, we also run the regression in log variables. The slope coefficients are much closer to one than before, indicating that even though the log variance risk premium may be negative, the logarithm of realized variance is closer to being uncorrelated to the log of the variance swap rate. In other words, the log risk premium is more of a constant time-series than the risk premium defined as in the monetary payoff measure.

Both regressions provide evidence of a less than perfect correlation between the variance swap rate and realized variance itself. As such, we conclude that the difference between the two, that is the variance risk premium, is not constant over time. In fact, we find evidence of it to be positively related to the variance swap rate. We therefore reject Hypothesis 3 of a constant variance risk premium over time for all maturities on the S&P 500 and the 30-day horizon on the OMXS30 index. We do not investigate the time-varying properties of the longer-term maturities of the OMXS30 as we cannot be certain that there exist variance risk premiums on these horizons.

8 Robustness Analysis

In this section we scrutinize the most critical aspects in arriving at the estimates of the variance risk premium. First, we attempt to redefine *discretely realized variance* in a set of ways to see if variance measurement errors could give rise to unwanted discrepancies. After having eliminated all reasonable doubts associated with the first component of the variance risk premium, we turn to the second component – the risk-neutral expectation. In order to verify the robustness of the calculated fair future realized variance, $\widehat{K}_{\text{VAR}}$, we examine some of the assumptions, parameters and techniques used in order to obtain this estimate.

Lastly, we examine the unique dataset on the OTC quotes of variance swap on the S&P 500. Comparing this data to our results, it is clear our method provides estimates close to the actual market rates. Our methodology is thus able to robustly quantify the variance risk premium given only option data.

8.1 Robustness of realized variance

In order to verify the robustness of the first component of the variance risk premium, the realized variance, we attempt several ways of computing the discretely sampled variance of equation (10). Firstly, instead of using the spot index levels, we compute daily log returns of the forward prices whose maturity T is fixed and equals the maturity of the variance swap.

$$R^F_{t_i} = \ln F_{t_i;T} \; - \ln F_{t_{i-1};T} \tag{19}$$

The variance of these log returns is as previously defined:

$$RV^F_{t,T} = \frac{1}{\triangle t \cdot N} \sum_{i=1}^{N} \left(R^F_{t_i}\right)^2 \tag{20}$$

As the forward prices incorporate information on the future dividends, we not only capture the variance of the underlying but also of its expected dividends and the interest rate. This should increase our estimate of realized variance and thus decrease the magnitude of the variance risk premiums. However, by using the forward prices, we avoid the additional variance caused by the actual payments of the dividends. Therefore, the difference in variance between the two measures mainly arises if there are differences between the variance of the expectations and the actual dividends. We find that the difference to our original realized variance measure is small.

Secondly, even though this is not the convention in the variance swap market according to Bossu et al. (2005), we compute realized variance with subtracting the mean return on the index during the period of the swap $[t,\ T]$. In doing so, we lose one degree of freedom, which explains why we subtract one from the number of observations (N) in the equation below:

$$RV^{mean}_{t,T} = \frac{1}{\triangle t(N-1)} \sum_{i=1}^{N} \left(R_{t_i} - \overline{R_{t,T}}\right)^2 \tag{21}$$

Subtracting the mean obviously lowers the realized variance, and hence the magnitude of the risk premium will increase due to this. Our original method thus provides a more conservative estimate of the variance risk premium. As can be seen in Table IX, the difference is however minor and this change does not significantly affect our results. Similarly, changing the annualization convention from trading days to calendar days does not significantly affect our results. The same holds true for calculating realized variance for the discrete returns rather than the log returns.

8.2 How fair is the fair price of future variance?

8.2.1 Cleaning the data

It follows that an accurate estimate of K_{VAR} relies on the accuracy of the option price data entering the replicating portfolio. Inaccurate data leads to erroneously estimated variance swap rates. It is therefore essential to first examine the data we have at our disposal.

After having cleaned options observations that violate no-arbitrage restrictions or are clearly unreasonable, we turn to the relatively more arbitrary cleaning choices from Section 6.2. In the original case we impose the requirement that there must be at least four or more valid option prices across the entire strike interval for any given day. This is done to avoid calculating K_{VAR} for days with only one option available for the replicating portfolio, which would most likely yield strange results. We also test by lowering the threshold value to two. In this case we obtain a valid K_{VAR} for 25% more days on the OMX30 and 1% for the S&P 500. As can be seen in Table IX this does not change the size of the risk premiums on any maturity for any of the two indexes.

Additionally, we attempt to change the number of options with strikes that must be greater or lower than the ATM forward level at time T. This is done to avoid clusters of OTM options that may bias the entire implied volatility shape. From the original value of one, we increase the threshold to two. This gives us 13% less days to analyze on the OMX30 and 2% less on the S&P 500. The risk premiums do not change greatly. This indicates our original cleaning procedure to be robust with regards to the magnitude of the risk premiums compared to more restrictive cleaning assumptions.

8.2.2 Possible errors in the practical implementation procedure

There are several approximation errors that arise due to difference between the theoretical equation (9) and the discretized equation (12) implemented in practice. In this section we attempt to outline the sources of errors that may arise.

Firstly, due to the limited availability of options with a continuum of strike prices, we must reside to a finite strike range on the interval $[K_L, K_H]$. This leads to so-called *truncation* errors:

$$\varepsilon_{truncation} = -\frac{2}{T}e^{r(T-t)}\left(\int_0^{K_L}\frac{1}{K^2}P\left(K\right)dK + \int_{K_H}^{\infty}\frac{1}{K^2}C\left(K\right)dK\right) \tag{22}$$

As can be seen from the equation above, the expected sign of the error is negative. This implies that our method underestimates K_{VAR} and thus provides a more conservative measure of the variance risk premiums. It can also be seen that the size of the error decreases for larger intervals $[K_L, K_H]$. Originally, we define this interval to encompass ± 8 standard deviations around the ATM forward level. Jiang and Tian (2007) claim truncation errors are typically negligible for values above three.

To gauge the sensitivity of our results to changes in this parameter, we first lower the interval to only encompass ± 2 standard deviations around the ATM forward level. As expected, this consistently decreases the average risk premiums by a small amount. On both the S&P 500 and the index OMXS30, the differences are however negligible. On the other hand, increasing this parameter to sixteen standard deviations, the new variance risk premiums in Table IX remain identical to the original results. This shows our choice of eight standard deviations to be reasonable and that our method is robust to truncation errors.

The second source of error arise due to the approximation of the integral of equation (9) on the interval $[K_L, K_H]$ by summation, we term this the *discretization* error. The sign of the discretization error is difficult to predict, but the size of the error will depend on the size of ΔK. More intuitively, it depends on the fineness of the spacing between options with varying strikes. The discretization error will disappear as ΔK approaches zero. Jiang and Tian (2007) explain that discretization errors are typically negligible for values of ΔK that are less than 0.5% of the underlying level.

The fineness of the mesh is a tradeoff between computational intensity and error minimization. We originally create a fine grid of 5,000 points on the interval $[K_L, K_H]$, which can be compared to the 2,000 points of Carr and Wu (2007). Even though the authors use a lower number than ours, they still find that their discretization error is virtually zero.

Similarly we do not find that increasing the number of points in the mesh to 50,000 changes our estimates of the variance risk premiums on either index. The results are identical up to at least eight decimal places. On the other hand lowering the number to only 50 points, does affect our results; but only by approximately 0.1% in the OMXS30 and 1% in the S&P 500. That is, even if we were to use a very coarse grid, it would not greatly affect our results. Our method is therefore robust to discretization errors.

The third source of error is attributable to the interpolation and extrapolation schemes of the Black-Scholes implied volatility function. Instead of natural cubic splines we attempt to interpolate linearly between implied volatility observations, in line with Carr and Wu (2007). The estimates of the risk premiums with linear interpolation from Table IX, lower the estimates of the monetary risk premiums by approximately 1% compared to the original results. Therefore the original method provides a more conservative estimate of the variance risk premium. The drawback of linear interpolation is that it creates kinks in the implied volatility functions. This is not consistent with no arbitrage as it may give rise to negative risk-neutral densities. Therefore our original approach of natural cubic splines is preferable.

The extrapolation scheme by which we and Carr and Wu (2007) hold the implied volatilities constant at the end-points is different to the approach Jiang and Tian (2007) advocate. Their alternative instead linearly extrapolates on the interval with unobserved option prices by setting the slopes of the interpolation and extrapolation schemes equal at the end points. Most frequently, this will produce small, positive slopes at the end points which consequently drive up the Black-Scholes implied volatilities at the deep-OTM option ends. This in turn causes higher option prices and therefore the replicating portfolio becomes more expensive, increasing the magnitude of the variance risk premiums. Because of this our original method produces more conservative estimates of the variance risk premiums, which can also be verified in Table IX. The differences are not very large, and can be explained by the extremely small values of these extrapolated deep-OTM options relative to the value of the entire replicating portfolio.

The fourth source of error originates from the linear interpolation scheme employed to find the variance swap at the correct maturity T. Only for 3% of the days with a variance swap rate, the option maturity actually corresponds to the maturity of the swap. For all other days, we interpolate between the maturities surrounding the maturity T. This gives rise to an approximation error if the variance swap rate is a non-linear function of maturity.

The sign of the approximation error will depend on the non-linear shape of the term structure of the variance swap rate. There are however no conclusive results on the possible shapes of the latter in the limited literature available. Investigating the term structure of the variance swap rate is beyond the scope of our study and would command a separate thesis in itself. We can therefore not quantify this most likely, small, source of error – if any.

8.2.3 Subsample analysis

In order to test if there are any differences between different periods in our sample, we extract the variance risk premiums for a set of subsamples. We shorten both the S&P 500 and the OMXS30 respectively so as to match the sample periods between 1999 and 2007 that are common to both. This is done in order to be able to fully compare the results. However, the differences between the two indexes are not due to the time period investigated. The OMXS30 long-term maturities are still not significant and the magnitude of the variance risk premium for the S&P 500 in fact increases slightly, especially for the long-term maturities when using this sample period.

Secondly, we divide the sample on the S&P 500 and OMXS30 on two periods that approximately correspond to two cycles of the economy. The period up to 2002 contains both the unprecedented bull market of the late 1990s as well as the following downturn. The second period from 2003 up to 2007 is characterized by a strengthening economic environment. The two periods also approximately correspond to two different variance states, in which the variance in the earlier period is generally higher and characterized by occurrences with variance spikes, which can also be seen in Figure 1.

By dividing our sample into these subsamples, it is also possible to see if the increasing trading and liquidity of the options market influences the variance risk premiums in any way. The absolute magnitude of the variance risk premium could have decreased as a result of more investors, especially hedge funds, short-selling variance risk.

Surprisingly, the summary results in Table IX show that the variance risk premiums on the S&P 500 become consistently even more negative for the latter subsample period between 2003 and 2007. This cannot be attributed to less efficient markets or decreasing liquidity and suggests that investors have demanded higher premiums in recent years. Such a conclusion itself may have many interesting explanations.

The subsamples for the OMXS30 displays similar features as the S&P 500. In the latter subsample between 2003 and 2007, the variance risk premium measures are more negative across all maturities, apart from the 30-day measures. More interestingly, the discrete variance risk premium on the 30-day horizon is no longer significant for the latter subsample. In part, this may be explained by the lower number of valid observations. Now however, all measures of the 60-day variance risk premium are negative and significant instead.[10]

8.3 Using OTC variance swap market quotes

In order to validate the robustness of our method in a more realistic setting, we compare our calculated variance swap rates to the OTC market quotes provided by two large broker-dealers of variance swaps on the S&P 500. We restrict our analysis to the 90, 180 and 365 day variance swap rates as these maturities are common to both of the datasets.

Since there is no public market for variance swaps, there is no unique market price at any given time. The variance swap from Goldman Sachs and Barclays Capital differ slightly, but the discrepancies are extremely small. On average, the two datasets differ by approximately -0.8% on the 90-day and 2% on the 365-day maturity.

It is apparent that the variance swap rates from the two broker-dealers generally lie in a close proximity to each other. Despite this, we synthesize the two datasets into one combined dataset by taking the lowest of the variance swap rate for the days which are common in both datasets. On the contrary to Egloff et al. (2007) that only use variance swap rates from one broker-dealer, we avoid potential observations that would negatively bias the size of the

[10]These results are however not presented in Table IX because of space limitations, but can be obtained upon request.

risk premium. That is, we obtain more conservative estimates of the size of the variance risk premiums.

Comparing the OTC variance swap rates with our estimate of K_{VAR}, we find that the average difference is very small. Our original estimates of K_{VAR} are on average lower by approximately -0.1% on the 365-day to -2% on the 90-day maturities. This indicates that our results are not only close to the real market quotes, but that they also provide a slightly more conservative estimate of the variance risk premiums.

The major benefit of using OTC variance swap rates over our synthetically constructed rates is that the former also reasonably should price empirical anomalies such as jumps. This is something that our method cannot take into account. After accounting for a maximum bid-ask spread of 100 basis point that is typically encountered in the variance swap market according to Hafner and Wallmeier (2007), we see that approximately 93% of our calculated variance swap rates fall within this range. This shows that our method performs well in more realistic settings.

We calculate the return and risk-adjusted measures for the new OTC dataset from April 29^{th} 1997 to March 1^{st} 2008. From Table IX, it can be seen that the magnitude of the risk premiums generally decreases compared to our original results. The risk premiums are however still strongly negative and highly significant. Similarly the risk-adjusted measures are slightly smaller, but all of them are still far superior to the risk-adjusted excess returns on the MSCI World index.

These results show that our methodology is robust and that the results we obtain are in line with the risk-return characteristics of real-world short variance swap investments. More importantly, this robustness verification shows that even in reality, short-sellers of variance swaps can earn abnormal returns on a risk-adjusted basis.

Because there is no variance swap market for the OMXS30, or there is very little trading, we do not have any actual data with which we can compare to our results. However we implement exactly the same methodology on the OMXS30 as on the S&P 500 index. Since we can we infer that our method is robust in reality for the S&P 500, it would also be reasonable to expect that our methodology is robust on the OMXS30 index as well.

8.4 Liquidity

In order to gauge the effect of transaction costs on the variance risk premium, we use the bid prices of the options on the S&P 500 and OMXS30 indexes. In other words, instead of using the *mid* Black-Scholes implied volatilities based on the average of the bid and ask prices, we use the lower of the two. Because of this the price of the replicating portfolio is lower. As a consequence, this will provide us with a lower boundary of magnitude of the variance risk premium that investors can "lock-in" in order to be able to short-sell variance.

Using bid prices of options decreases the monetary variance risk premiums on the S&P 500 by approximately 20% over all maturities. Interestingly, all variance risk premiums are still strongly negative and significant on all reasonable significance levels. This finding is in line with Carr and Wu (2007). The results for the OMXS30 index are however not quite robust enough to the use of bid prices. The only significant risk premium on the 30-day horizon is now insignificant. The reason for this is that the liquidity on the Swedish option market is lower than on the S&P 500, which is not surprising. However the original method still provides fair values of future variance if the bid-ask spread is symmetric around the mid prices, which is commonly assumed to be the case.

For our combined OTC variance swap rate dataset on the S&P 500 we also introduce a bid-ask spread that takes liquidity concerns into account. The variance swap market is very liquid and according to market participants the bid-ask spread is typically within 30 basis points in volatility space. In order to account for transaction costs in the OTC market, we subtract the *entire* spread from the square root of the OTC variance swap rates. We then square this to obtain K_{VAR} in units of variance. The variance risk premiums that we obtain after adjusting for this spread are greatly negative and highly significant. This suggests that market participants in the OTC variance swap market are able to capture the variance risk premium even after accounting for transaction costs.

9 Conclusion

In this thesis we find that the market prices variance risk on both the S&P 500 and OMXS30 indexes. Using a model-free approach, the average variance risk premiums extracted from the S&P 500 are strongly negative for all maturities. For the OMXS30, the variance risk premium is only significantly negative for the 30-day maturity and its magnitude is smaller than for the S&P 500. For the long-term maturities on the OMXS30, the variance risk premiums are not statistically significantly different from zero. Their median values are however consistently negative, which can be explained by a small number of extreme positive realizations.

Given that the average variance risk premium is negative, there is an opportunity to profit from this phenomenon. Through the short sales of variance swaps, a derivative security on variance, we explain how investors are able to exploit negative variance risk premiums. Over a relatively extensive sample period, we show that the risk premium charged by investors for holding variance risk is not in proportion to its risks. This suggests that investors have previously been able to exploit the superior risk-return characteristics of short-selling variance swaps on both the S&P 500 and OMXS30 index options compared to other alternative investments.

We also find evidence of the time-varying properties of the variance risk premium. This suggests that investors continuously reappraise the riskiness of variance itself. Moreover, this riskiness seems to be positively related to the absolute levels of variance.

The existence of a negative variance risk premium on the S&P 500 has been known at least since 2000-2003 when the first studies of this empirical phenomenon emerged in the literature. Also investors, in particular hedge funds, have in recent years taken advantage of the negative variance risk premiums on the S&P 500 and other major indexes such as the EURO STOXX 50 or the DAX. On the contrary to the OMXS30, these indexes have liquid OTC variance swap markets enabling institutional investors to efficiently hedge or take views on future variance.

Our approach of synthetically constructing variance swap rates in a model-free way, with the purpose of quantifying the variance risk premium on the OMXS30, is therefore the first time such a study has been implemented on the options of this Swedish equity index. Our findings suggest that negative variance risk premiums are not specific to a major global index and that there are opportunities for investors to earn superior risk-adjusted returns on smaller markets as well.

We also find our method for estimating the variance risk premiums is robust under a range of computational assumptions and for more realistic circumstances. We see that our method performs well in comparison to a unique dataset on traded OTC variance swaps for the S&P 500 index. Moreover, investors are equally able to take advantage of the superior risk-adjusted returns of the short variance swaps through the OTC market, even after accounting for transaction costs.

This leads us to the question: how can we rationalize these abnormal risk-adjusted returns from shorting variance swaps? Taleb (2007) coins the concept of a *black swan*, similar to the *Peso* concept we previously introduced. An example of a black swan is 9/11. It is precisely these unexpected, but rare events that do not repeat themselves, that cause short-sellers of variance risk to incur very large losses. Investors would thus be irrational not to take these risks into account. The complication arises if these events do not happen during a given time period, even if this period is long, it may appear as if these risks do not exist! But even *if* the black swan event does occur, it will not repeat itself – by definition. This explains why these events are difficult to consider in any risk-adjusted performance measure. One must have historical data of several centuries in order to be able to somewhat accurately calculate the probability of an adverse event that happens every 100 years. Certain events occur even less frequently.

In light of this, despite the relatively short period of time that our study covers, we do provide a point of entry to these issues. Short-selling variance risk will only be truly profitable if investors systematically undervalue the true (risk-neutral) probabilities of large negative realizations of these instruments. However, a consistent mispricing is an obvious violation of efficient markets. The variance risk premium therefore provides an insight into the nature of the elusive concept of risk and investors' views on the future.

References

Acerbi, C. and Tasche, D. (2002). On the Coherence of Expected Shortfall, *Journal of Banking and Finance* **26**(7): 1487–1503.

Andersen, T. G. and Bondarenko, O. (2007). *Construction and Interpretation of Model-Free Implied Volatility*, Risk Books, London, pp. 141–181.

Artnzer, P., Delbaen, F., Eber, J.-M. and Heath, D. (1999). Coherent Measures of Risk, *Mathematical Finance* **9**(3): 203–228.

Bakshi, G., Cao, C. and Chen, Z. (2000). Do Call Prices and the Underlying Stock Always Move in the Same Direction?, *The Review of Financial Studies* **13**(3): 549–584.

Bakshi, G. and Kapadia, N. (2003a). Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *The Review of Financial Studies* **16**(2): 527–566.

Bakshi, G. and Kapadia, N. (2003b). Volatility Risk Premiums Embedded in Individual Equity Options: Some New Insights, *The Journal of Derivatives* **11**(1): 45–54.

Bakshi, G. and Madan, D. (2006). A Theory of Volatility Spreads, *Management Science* **52**(12): 1945–1956.

Bekaert, G. and Wu, G. (2000). Asymmetric Volatility and Risk in Equity Markets, *The Review of Financial Studies* **13**(1): 1–42.

Björk, T. (2004). *Arbitrage Theory in Continuous Time*, second edn, Oxford University Press, Oxford.

Bollen, N. P. B. and Whaley, R. E. (2004). Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, *The Journal of Finance* **59**(2): 711–753.

Bollerslev, T., Gibson, M. and Zhou, H. (2007). Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities. FEDS Working Paper No. 2004-56.

Bondarenko, O. (2003). Why are Put Options So Expensive. University of Illinois at Chicago Working Paper.

Bondarenko, O. (2004). Market Price of Variance Risk and Performance of Hedge Funds. University of Illinois at Chicago Working Paper.

Bondarenko, O. (2007). Variance Trading and Market Price of Variance Risk. University of Illinois at Chicago Working Paper.

Bossu, S., Strasser, E. and Guichard, R. (2005). *What You Need to Know About Variance Swaps*, second edn, JPMorgan Quantitative Research, London.

British Bankers' Association (2008-03-15). `http://www.bba.org.uk/`.

Buraschi, A. and Jackwerth, J. (2001). The Price of a Smile: Hedging and Spanning in Option Markets, *The Review of Financial Studies* **14**(2): 495–527.

Carr, P. and Madan, D. (1998). *Towards a Theory of Volatility Trading*, Risk Publications, London, pp. 417–427.

Carr, P. and Wu, L. (2007). Variance Risk Premia. forthcoming Review of Financial Studies.

Carr, W. and Wu, L. (2006). A Tale of Two Indices, *The Journal of Derivatives* **13**(3): 13–29.

Chichago Board of Exchange (2008-03-15). `http://www.cboe.com/`.

Christensen, B. J. and Prabhala, N. R. (1998). The relation between implied and realized volatility, *The Journal of Financial Economics* **50**(2): 125–150.

Coval, J. D. and Shumway, T. (2001). Expected Option Returns, *The Journal of Finance* **56**(3): 983–1009.

Demeterfi, K., Derman, E., Kamal, M. and Zou, J. (1999). A Guide to Volatility and Variance Swaps, *The Journal of Derivatives* **6**(4): 9–32.

Doran, J. S. and Ronn, E. I. (2008). The Market Price of Volatility Risk in the Energy Commodity Markets. Forthcoming Journal of Banking and Finance.

Duffie, D. (2001). *Dynamic Asset Pricing Theory*, third edn, Princeton University Press, Princeton, New Jersey.

Dupire, B. (1994). Pricing with a Smile, *Risk* **7**(1): 18–20.

Egloff, D., Leippold, M. and Wu, L. (2007). Variance Risk Dynamics, Variance Risk Premia, and Optimal Variance Swap Investments. EFA 2006 Zurich Meetings Working Paper.

French, K. R., Schwert, W. G. and Stambaugh, R. F. (1987). Expected Stock Returns and Volatility, *The Journal of Financial Economics* **19**(1): 3–29.

Gagliardini, P. and Trojani, F. (2007). Asset Pricing Course Notes. University of Sankt Gallen.

Goetzmann, W., Ingersoll, J., Spiegel, M. and Welch, I. (2002). Sharpening Sharpe Ratios. NBER Workning Paper No. 9116.

Greene, W. H. (2002). *Econometric Analysis*, fifth edn, Prentice Hall, Upper Saddle River, New Jersey.

Guo, D. (1998). The Risk Premium of Volatility Implicit in Currency Options, *Journal of Business & Economic Statistics* **16**(4): 498–507.

Hafner, R. and Wallmeier, M. (2007). Volatility as an Asset Class: European Evidence, *European Journal of Finance* **13**(7): 621–644.

Jiang, G. J. and Tian, Y. S. (2007). Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index, *The Journal of Derivatives* **14**(3): 35–60.

Low, B. S. and Zhang, S. (2005). The Volatility Risk Premium Embedded in Currency Options, *Journal of Financial and Quantitative Analysis* **40**(4): 803–832.

Mougeot, N. (2005). *Volatility Investing Handbook*, BNP Paribas Equities & Derivatives Research, Paris.

Neuberger, A. (1994). The Log Contract: A New Instrument to Hedge Volatility, *Journal of Portfolio Management* **20**(2): 74–80.

Newey, W. K. and West, K. D. (1987). A Simple Positive Semi-definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* **55**(3): 703–708.

Optionmetrics (2008-02-23). `http://www.optionmetrics.com/`.

Panigirtzoglou, N. and Lynch, D. (2003). Option Implied and Realised Measures of Variance. Bank of England Working Paper.

Sharpe, W. F. (1966). Mutual Fund Performance, *The Journal of Business* **39**(1): 119–138.

Taleb, N. N. (2004). The Practioner's Perspective: Bleed or Blowup? Why Do We Prefer Asymmetric Payoffs?, *The Journal of Behavioral Finance* **5**(1): 2–7.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*, Random House, New York.

Thomson Datastream (2008-03-15). `http://www.datastream.com/`.

Wu, G. (2001). The Determinants of Asymmetric Volatility, *The Review of Financial Studies* **14**(3): 837–859.

A Formally Deriving the Variance Risk Premium

A.1 Assuming a deterministic and constant short rate

In our approach to quantifying the variance risk premium, we assume deterministic and constant risk-free interest rates. The purpose of this thesis is not to study the short-rate r_s. In other words, we assume that the payoff and the short rate r_s are independent under the $\mathbb{Q}$ measure. This is the only way of safely eliminating the covariance term that may arise between the two. The value of the variance swap:

$$\Pi_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r_s ds} \left(RV_{t,T} - K_{\mathrm{VAR}t,T} \right) \right] \tag{23}$$

therefore splits into:

$$\Pi_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r_s ds} \right] \mathbb{E}_t^{\mathbb{Q}} \left[RV_{t,T} - K_{\mathrm{VAR}t,T} \right] \tag{24}$$

where $\mathbb{Q}$ is the martingale measure that has the money-market account as a numéraire. Under the assumption of a deterministic and constant short rate:

$$\mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r_s ds} \right] = e^{-r(T-t)} \tag{25}$$

where r is the continuously compounded risk-free interest rate at time t between t and T as previously defined. This gives equation (3).

A.2 The asset pricing approach to the variance risk premium

Carr and Wu (2006, 2007) show that by using the notion of stochastic discount factors (SDF) found in asset pricing theory the variance risk premium can be derived more formally. The SDF process Λ is defined as:

$$\Lambda\left(t\right) = e^{-\int_0^t r_s ds} L(t) \tag{26}$$

where $L(t)$ is defined as the likelihood ratio, or also known as the Radon-Nikodym derivative of $\mathbb{Q}$ with respect to $\mathbb{P}$. That is:

$$L\left(t\right) = \frac{d\mathbb{Q}}{d\mathbb{P}}, \qquad \text{on } \mathcal{F}_t \tag{27}$$

where $\mathcal{F}_t$ can be considered to be the information generated on the interval $[0,\ t]$. Under no arbitrage it is shown amongst others, by Duffie (2001) that there exists at least one SDF for all assets. The arbitrage free price, $\Pi_t(X)$, of any payoff X under the $\mathbb{P}$ probability measure can be shown to equal:

$$\Pi_t(X) = \mathbb{E}_t^{\mathbb{P}} \left[\frac{\Lambda(T)}{\Lambda(t)} X | \mathcal{F}_t \right] \tag{28}$$

Under no arbitrage this value must equal the value under the risk-neutral measure $\mathbb{Q}$.

$$\Pi_t(X) = \mathbb{E}_t^{\mathbb{P}} \left[\frac{\Lambda(T)}{\Lambda(t)} X | \mathcal{F}_t \right] = E_t^{\mathbb{Q}} \left[e^{-\int_t^T r_s ds} X | \mathcal{F}_t \right] \tag{29}$$

Using that the payoff of a variance swap with a notional of $\aleph = 1$ equals:

$$X = RV_{t,T} - K_{\mathrm{VAR}t,T} \tag{30}$$

and assuming deterministic interest rates as outlined in the section above, the value of the variance swap at time t under the risk-neutral measure $\mathbb{Q}$ then equals:

$$\begin{aligned} \Pi_t(X) = \Pi_t\left(RV_{t,T} - K_{\mathrm{VAR}t,T}\right) &= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_s ds}(RV_{t,T} - K_{\mathrm{VAR}t,T})|\mathcal{F}_t\right] \\ &= e^{-r(T-t)}\mathbb{E}_t^{\mathbb{Q}}\left[RV_{t,T}|\mathcal{F}_t\right] - e^{-r(T-t)}K_{\mathrm{VAR}t,T} \end{aligned} \tag{31}$$

Here we use the fact that $K_{\mathrm{VAR}t,T}$ is known at time t and can therefore be lifted out of the expectation. Since we would like to compute the fair value of future realized variance, the value of the variance swap is set to zero. The discount factors cancel out and this implies that:

$$0 = \Pi_{\mathrm{t}}\left(RV_{t,T} - K_{\mathrm{VAR}t,T}\right) \Rightarrow K_{\mathrm{VAR}t,T} = \mathbb{E}_t^{\mathbb{Q}}\left[RV_{t,T}|\mathcal{F}_t\right] \tag{32}$$

From equation (28), it is clear that:

$$K_{\mathrm{VAR}t,T} = \mathbb{E}_t^{\mathbb{Q}}\left[RV_{t,T}|\mathcal{F}_t\right] = \mathbb{E}_t^{\mathbb{P}}\left[D_{t,T}RV_{t,T}|\mathcal{F}_t\right] \tag{33}$$

where $D_{t,T} = \frac{\Lambda(T)}{\Lambda(t)}$. The basic properties of covariance stipulate that for any variable X and Y:

$$Cov[X,Y] = E\left[X \cdot Y\right] - E\left[X\right] \cdot E\left[Y\right] \tag{34}$$

Given that $\mathbb{E}_t^{\mathbb{P}}\left[\frac{\Lambda(T)}{\Lambda(t)}|\mathcal{F}_t\right] = 1$, equation (33) above can be decomposed into:

$$K_{\mathrm{VAR}t,T} \equiv \mathbb{E}_t^{\mathbb{Q}}\left[RV_{t,T}|\mathcal{F}_t\right] = \mathbb{E}_t^{\mathbb{P}}\left[D_{t,T}RV_{t,T}|\mathcal{F}_t\right] = \mathbb{E}_t^{\mathbb{P}}\left[RV_{t,T}|\mathcal{F}_t\right] + Cov_t^{\mathbb{P}}\left[D_{t,T}, RV_{t,T}|\mathcal{F}_t\right] \tag{35}$$

In asset pricing theory, asset payoffs that covary negatively with the SDF command a positive risk premium and vice versa. Asset payoffs uncorrelated with the SDF have zero risk premiums. Therefore $-Cov_t^{\mathbb{P}}\left[D_{t,T}, RV_{t,T}\right]$ defines the variance risk premium. The first term in equation (35) above is the conditional mean of the realized variance. Removing the conditioning on the $\mathcal{F}_t$-measureable for simplicity, we then arrive at the definition of the variance risk premium of equation (1) in Section 2.1:

$$\mathbb{E}_t^{\mathbb{Q}}\left[\text{Realized variance}_{t,T}\right] = \mathbb{E}_t^{\mathbb{P}}\left[\text{Realized variance}_{t,T}\right] - RP_{t,T} \tag{36}$$

A.3 Estimation of the variance risk premium

Defining the true measure of realized variance as $RV_{t,T}$ with its estimate defined as $\widehat{RV}_{t,T} \equiv RV_{t,T} + \varepsilon_{t,T}$ where $\varepsilon_{t,T}$ is a random source of error such that $\mathbb{E}_t[\varepsilon_{t,T}] = 0$. The variance risk premium is then defined as:

$$\begin{aligned} \mathbb{E}_t^{\mathbb{P}}\left[\widehat{RP}_{t,T}\right] \equiv \mathbb{E}_t^{\mathbb{P}}\left[\widehat{RV}_{t,T}\right] - \mathbb{E}_t^{\mathbb{P}}\left[D_{t,T}\widehat{RV}_{t,T}\right] &= -Cov_t^{\mathbb{P}}\left[D_{t,T}, \widehat{RV}_{t,T}\right] \\ = \mathbb{E}_t^{\mathbb{P}}[RP_{t,T}] - Cov_t^{\mathbb{P}}\left[D_{t,T}, \varepsilon_{t,T}\right] & \end{aligned} \tag{37}$$

Gagliardini and Trojani (2007) show that the estimated risk premium only equals the true risk premium when the covariance term with the error equals zero. That is, assuming that the sample variance is an unbiased estimate of the true realized variance only holds when the error does not represent a systematic source of risk. We do not find this plausible and therefore disregard this error source.

B Deriving the Fair Value of Future Variance

Assuming that the stock price process S_t follows:

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_t \tag{38}$$

where the drift and volatility processes μ_t and σ_t are arbitrary functions of time and other parameters. By using Itô on $\log(S_t)$, we obtain:

$$d(\log(S_t)\ ) = \left(\mu_t - \frac{1}{2}\sigma_t^2\right) dt + \sigma_t^2 dZ_t \tag{39}$$

Subtracting (39) from (38) after dividing (38) by S_t on both sides gives:

$$\frac{dS_t}{S_t} - \ d(\log(S_t)\ ) = \frac{1}{2}\sigma_t^2 dt \tag{40}$$

This removes all dependence on the drift μ_t and leaves only exposure to volatility. The variance is now given by:

$$\nu_{t,T} \equiv \int_t^T \sigma_t^2 \ dt = \frac{2}{T-t}\left[\int_t^T \frac{dS_t}{S_t} \ - \ \log\left(\frac{S_T}{S_t}\right)\ \right] \tag{41}$$

Taking the risk-neutral expectation now gives the fair variance swap rate:

$$K_{\mathrm{VAR}t,T} = \frac{2}{T-t}\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T \frac{dS_t}{S_t} \ - \ \log\left(\frac{S_T}{S_t}\right)\ \right] \tag{42}$$

Since $\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T \frac{dS_t}{S_t}\ \right] = r(T-t)$ by construction:

$$K_{\mathrm{VAR}t,T} = \frac{2}{T-t}(r(T-t) - \mathbb{E}_t^{\mathbb{Q}}\left[\log\left(\frac{S_T}{S_t}\right)\ \right] \tag{43}$$

Now given the fact that:

$$\log\left(\frac{S_T}{S_t}\right) \ = \log\left(\frac{S_T}{S_*}\right) \ + \log\left(\frac{S_*}{S_t}\right) \tag{44}$$

for any S_* we notice that only the first term is dependent on S_T and in need of replication (the second term is constant). Further on we have that:

$$-\log\left(\frac{S_T}{S_*}\right) \ = -\frac{S_T - S_*}{S_*} + \int_0^{S_*} \frac{1}{K^2}\max(K - S_T, 0)dK + \int_{S_*}^{\infty} \frac{1}{K^2}\max(S_T - K, 0)dK \tag{45}$$

Inserting equation (44) and (45) into (43) and taking the expectation now gives:

$$K_{\mathrm{VAR}t,T} = \frac{2}{T}\Bigg(r(T-t) - (\frac{S_t}{S_*}e^{r(T-t)} - 1) - \log\left(\frac{S_*}{S_0}\right) + e^{r(T-t)}\int_0^{S_*}\frac{1}{K^2}P(K)\,dK + \int_{S_*}^{\infty}\frac{1}{K^2}C(K)\,dK\Bigg) \tag{46}$$

Setting $S_* = F = e^{r(T-t)}S_t$, that is the forward price, yields:

$$K_{\mathrm{VAR}t,T} = \frac{2}{T}e^{r(T-t)}\left(\int_0^{F}\frac{1}{K^2}P(K)\,dK + \int_{F}^{\infty}\frac{1}{K^2}C(K)\,dK\right) \tag{47}$$

which can be discretized and calculated numerically.

C Defining the Risk-Adjusted Performance Measures

C.1 The Sharpe ratio

As first developed by Sharpe (1966), the Sharpe ratio assumes that the riskiness of the investment can be quantified by the return standard deviation.

$$\widehat{SR} = \frac{\mathbb{E}(r_i)}{\sqrt{var(r_i)}} \approx \frac{\overline{RP}_{discrete,log}}{\sigma_{RP}} \tag{48}$$

where r_i are the excess returns and RP is estimated by either $RP_{discrete}$ or RP_{log} and σ_{RP} is the standard deviation of these estimates of the risk premium. There is no need to subtract the risk-free rate to obtain the excess return, as it is already accounted for in the definitions of the RP.

In order to be able to compare across varying maturities and with comparable investment opportunities, we annualize the Sharpe ratios by multiplying with $\sqrt{365/(T-t)}$, where T – t is the length of the variance swap maturity in calendar days.

C.2 Value-at-Risk (VaR)

This measure of quantifying risk was fist invented by JPMorgan in the early 1990s. The measure is commonly defined as the maximum loss that is not exceeded given a certain $\theta \in (0, 1)$ confidence level during a given time period. In probabilistic terms, it is the θ quantile of the profit-loss distribution. For our purposes, it more convenient to calculate the VaR in return space and the basic concept does not change.

The main benefit of VaR is that it takes a distributional approach to measuring risk. It also focuses on left-tail realizations as the source of these risks, which are of utmost interest to us as well. It is precisely these extreme realizations that are driving the risk-return tradeoff for variance swaps.

In calculating the VaR we implement the empirical distribution approach. On the contrary to other approaches, it does not rely on any given distribution of the returns, but rather taking the historical realizations into account. It does rely on the assumption that the sample distribution is an unbiased estimate of the true population distribution. For our purposes, this assumption is alleviated by the long sample period we use.

The commonly used alternative Monte-Carlo approach is not used for several reasons. Primarily, we do not feel comfortable with specifying a certain variance dynamics as *(i)* it would contradict our primary motivation of the *model-free* approach of extracting the variance swaps, which is that we do not rely on any specification of the underlying or variance dynamics *(ii)* there is to our knowledge no model specification of the variance risk premium dynamics that has proven to be empirically viable.

In order to be able to consistently compare our results, we scale the average return over a certain period by the absolute value of its historical VaR for a given confidence level θ which is also computed in return space.

$$\frac{\overline{RP}}{|\mathrm{VaR}_\theta|} \tag{49}$$

We arbitrarily set θ to 0.1%, 0.5%, 1%, 5% and 10%. It should also be noted that VaR is usually a forward-looking measure of risk, whereas we compute it on an ex post basis.

C.3 Conditional Value-at-Risk (CVaR)

CVaR was first presented by Artnzer et al. (1999) as a superior measure of risk compared to VaR. Apart from its axiomatic properties, the main advantage of CVaR over VaR is that the *entire* left tail is taken into account. CVaR or *expected shortfall* is defined as the expected value of losses given that a loss occurs at or below the θ quantile for a given time horizon. More formally defined,

$$\mathrm{CVaR}_\theta = \mathbb{E}\left(z \mid z \leq \mathrm{VaR}_\theta\right) \quad \text{where} \quad Prob\left(z \leq \mathrm{VaR}_\theta\right) = \theta \tag{50}$$

where z represents the loss and θ is the confidence level. For any integrable and continuous probability distribution of the profit and losses, $F(z)$, CVaR is defined as:

$$\mathrm{CVaR}_\theta = \frac{1}{1-\theta}\int_\theta^1 q_u\left[F\left(z\right)\right]du \tag{51}$$

where $q_u\left[F\left(z\right)\right]$ is the quantile function of $F(z)$, that is the VaR value of the same distribution as a function of any confidence level u. We omit the more complicated formula for discontinuous distributions developed by Acerbi and Tasche (2002) and instead we outline the formula that we implemented in practice.

$$\mathrm{CVaR}_\theta = \frac{\sum_{t=1}^{[n\theta]} R_{i;n}}{[n\theta]} \tag{52}$$

where $\{R_1, \ \ldots, \ R_n\}$ are returns from our empirical distributions ordered by increasing size and $[n\theta]$ denotes the largest integer not exceeding $n\theta$. As for VaR, we scale the average return over a period by the absolute value of its historical CVaR which is computed in return space.

D Data Cleaning

We perform the following cleaning procedures to our dataset:

1. Removing options for which the bid or ask price is not greater than zero. Obviously, negative or zero option prices are not reasonable and should be removed.

2. Removing options whose bid price is greater than the ask price. That is, we eliminate opportunities for the investor simultaneously to sell at higher prices than at which they buy options. The so-called crossed market is clearly unnatural under no-arbitrage.

3. Removing few options that violate basic no-arbitrage option price bounds. Only options that satisfy:

   - $Ke^{-r(T-T)} > [\textit{put bid or ask price}] > \max(K - S_T, 0)$
   - $S_T > [\textit{call bid or ask price}] > \max(S_T - K, 0)$

   are kept in the sample. S_T is the forward price.

4. Removing puts and calls that have the same strike, but whose Black-Scholes implied volatilities differ by more than 1%. Ideally according to put-call parity, the options should have the same implied volatilities but we allow for a small error margin. Puts and calls that fall outside this margin are assumed to be erroneously recorded and thus removed.

5. Removing call with prices that are not monotonically decreasing over increasing strike prices and puts that are not monotonically increasing over increasing strikes. Violating this monotonicity requirement gives rise to negative risk-neutral densities which are inconsistent with no arbitrage. We conduct this cleaning procedure by visual inspection for the days that clearly violate this restriction.

6. Removing options have the same bid and ask prices as the preceding trading day, given that the trading volume the same day is zero. Such observations are assumed to be the effect of stale prices, as market participants have not updated their quotes for the option in question. This is most frequently the case for deep-OTM options whose value is very small.

7. Removing options whose calculated Black-Scholes implied volatility is greater than 100%. We assume that such observations are the effect of erroneously recorded prices. This is frequently the case of deep-OTM option price, as they are quoted in small discrete steps of 0.01. If their prices are close to zero, even the smallest error has a large impact on implied volatilities.

8. Removing options where bid and ask differed by more than 50%. By averaging the bid and ask prices, the obtained mid price may not be correct if the spread is very large. This cleaning procedure is implemented as a precaution in order to remove very illiquid options whose prices are greatly distorted by transaction costs.

9. Removing options with a maturity that is less than 7 calendar days. This is done in order to avoid possible microstructure effects in the last week of trading and is in line with other studies found in the literature, such as Carr and Wu (2007).

10. Finally, after performing these cleaning measures, we only keep the days where there are at least four option prices across all strikes, of which at least one option must have a strike which is lower and one that is higher than at the forward ATM index level. This is done in order to avoid implied volatility smiles where all the options are concentrated at one OTM end of the smile. Basing the extrapolation scheme on such smile would most likely yield erroneous results if applied to the entire strike interval.

E Tables

	30-day RP			60-day RP			90-day RP		
	Payoff	Discrete	Log	Payoff	Discrete	Log	Payoff	Discrete	Log
Average RP	-2.10	-41.4%	-71.5%	-1.77	-35.5%	-60.1%	-1.72	-33.8%	-56.9%
Average RP (in volatility)	-5.79			-4.93			-4.75		
Median	-1.88	-52.3%	-74.0%	-1.65	-48.8%	-67.0%	-1.58	-48.2%	-65.8%
Maximum	11.8	250.6%	125.4%	8.76	226.8%	118.4%	7.25	209.9%	113.1%
Minimum	-35.1	-97.0%	-351.8%	-19.8	-94.0%	-281.2%	-17.1	-94.1%	-283.1%
Upper quartile	-0.86	-30.6%	-36.5%	-0.70	-24.3%	-27.9%	-0.66	-20.1%	-22.4%
Lower quartile	-3.06	-67.2%	-111.5%	-2.86	-62.7%	-98.6%	-2.87	-61.6%	-95.7%
Standard deviation	2.84	0.41	0.59	2.59	0.42	0.55	2.51	0.43	0.54
Skewness	-1.05	2.63	0.24	-0.32	2.08	0.49	-0.30	2.07	0.52
Kurtosis	17.02	12.85	3.60	8.24	8.49	3.06	7.06	8.50	3.02
Number of observations	2335			2802			2849		
Number of observations < 0	2083			2372			2369		
t-test, statistica, N-W adjusted	-12.00	-13.77	-16.65	-7.24	-8.11	-10.68	-5.92	-6.34	-8.45
t-test, p-value, N-W adjusted	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Jarque-Bera test, p-value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Kolmogorov-Smirnov test, p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

	180-day RP			365-day RP		
	Payoff	Discrete	Log	Payoff	Discrete	Log
Average RP	-1.71	-32.0%	-52.7%	-1.62	-28.9%	-46.2%
Average RP (in volatility)	-4.61			-4.34		
Median	-1.40	-45.4%	-60.6%	-1.34	-34.7%	-42.6%
Maximum	6.34	182.2%	103.7%	3.63	109.1%	73.8%
Minimum	-34.8	-90.2%	-231.9%	-13.27	-83.6%	-181.0%
Upper quartile	-0.59	-17.1%	-18.7%	-0.54	-9.2%	-9.7%
Lower quartile	-2.81	-60.0%	-91.5%	-2.73	-57.6%	-85.8%
Standard deviation	2.39	0.40	0.52	1.96	0.36	0.49
Skewness	-1.30	1.88	0.38	-0.66	1.09	0.08
Kurtosis	19.05	7.62	2.83	5.19	4.13	2.17
Number of observations	2857			2748		
Number of observations < 0	2398			2347		
t-test, statistica, N-W adjusted	-4.78	-4.61	-5.89	-3.70	-3.28	-3.47
t-test, p-value, N-W adjusted	0.0000	0.0000	0.0000	0.0001	0.0005	0.0003
Jarque-Bera test, p-value	0.001	0.001	0.001	0.001	0.001	0.001
Kolmogorov-Smirnov test, p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table II: **Statistics of the variance risk premium for S&P 500 during 1996-2007.** **The monetary payoff measure is in USD. The risk premium (RP) in volatility space is computed as the square root of the realized variance and $K_{\mathbf{VAR}}$ measures at each available day. The robust t-statistica is defined as: $t - \mathbf{stat} = \frac{\overline{RP} - \phi}{\sigma_{RP}^{N-W}}$ where ϕ equals zero, as under the null Hypothesis 1 and σ_{RP}^{N-W} denotes the Newey-West (1987) standard deviations adjusted for autocorrelation and heteroscedasticity, with the number of lags corresponding to the number of calendar days during the maturity of the variance swap.**

	30-day RP			60-day RP			90-day RP		
	payoff	discrete	log	payoff	discrete	log	payoff	discrete	log
Average RP	-2.03	-17.3%	-37.1%	-0.81	-8.5%	-22.4%	-0.43	-3.0%	-17.5%
Average RP (in volatility)	-4.09			-2.12			-1.46		
Median	-1.73	-34.7%	-42.7%	-1.03	-25.1%	-28.8%	-0.86	-23.3%	-26.5%
Maximum	18.7	1009.4%	240.6%	15.5	479.3%	175.7%	12.6	355.4%	151.6%
Minimum	-21.2	-91.9%	-251.7%	-14.6	-80.1%	-161.6%	-29.2	-81.7%	-169.6%
Upper quartile	-0.26	-7.1%	-7.4%	0.36	7.3%	7.0%	0.78	13.5%	12.6%
Lower quartile	-4.13	-51.9%	-73.2%	-2.85	-43.1%	-56.4%	-2.52	-41.6%	-53.8%
Standard deviation	4.69	0.77	0.54	3.97	0.58	0.49	3.79	0.63	0.51
Skewness	0.46	7.09	0.82	0.92	3.12	0.63	0.57	2.40	0.75
Kurtosis	6.28	74.90	5.25	5.88	17.72	3.89	6.54	10.05	3.44
Number of observations	1883			1821			1503		
Number of observations < 0	1483			1303			1031		
t-test, statistica, N-W adjusted	-4.67	-2.64	-7.98	-1.61	-1.27	-3.58	-0.72	-0.32	-2.11
t-test, p-value, N-W adjusted	0.0000	0.0041	0.0000	0.0537	0.1026	0.0002	0.2371	0.3745	0.0177
Jarque-Bera test, p-value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Kolmogorov-Smirnov test, p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

	180-day RP			365-day RP		
	payoff	discrete	log	payoff	discrete	log
Average RP	-0.39	5.9%	-13.8%	-1.00	3.8%	-13.4%
Average RP (in volatility)	-1.21			-1.86		
Median	-0.73	-23.8%	-27.2%	-0.70	-21.2%	-23.8%
Maximum	11.5	259.4%	127.9%	6.77	157.1%	94.4%
Minimum	-29.8	-81.6%	-169.1%	-13.5	-78.3%	-155.8%
Upper quartile	1.73	27.4%	24.2%	1.47	58.0%	45.7%
Lower quartile	-2.63	-43.6%	-57.2%	-2.28	-47.3%	-64.1%
Standard deviation	4.10	0.73	0.60	2.99	0.58	0.60
Skewness	-0.29	1.43	0.49	-0.55	0.49	-0.05
Kurtosis	8.71	4.05	2.51	3.72	1.81	1.57
Number of observations	864			465		
Number of observations < 0	570			267		
t-test, statistica, N-W adjusted	-0.51	0.34	-0.96	-1.31	0.23	-0.82
t-test, p-value, N-W adjusted	0.3060	0.3661	0.1699	0.0946	0.4095	0.2057
Jarque-Bera test, p-value	0.001	0.001	0.001	0.001	0.001	0.041
Kolmogorov-Smirnov test, p-value	0.000	0.000	0.000	0.000	0.000	0.000

Table III: **Statistics of the variance risk premium for OMXS30 during 1999-2007.**
The monetary payoff measure is in SEK. The risk premium (RP) in volatility space is computed as the square root of the realized variance and $K_{\mathbf{VAR}}$ measures at each available day. The robust t-statistica is defined as: $t-\mathbf{stat} = \frac{\overline{RP}-\phi}{\sigma_{RP}^{N-W}}$ where ϕ equals zero, as under the null Hypothesis 1 and σ_{RP}^{N-W} denotes the Newey-West (1987) standard deviations adjusted for autocorrelation and heteroscedasticity, with the number of lags corresponding to the number of calendar days during the maturity of the variance swap.

	30-day (OMXS30) RP		30-day (S&P 500) RP		60-day (S&P 500) RP	
	discrete	log	discrete	log	discrete	log
Annualized Sharpe Ratio	0.79	2.38	3.48	4.26	2.06	2.70
$\lvert RP/\mathrm{VaR}_\theta \rvert$						
$\theta = 0.1\%$	0.21	0.22	0.44	0.27	0.40	0.29
$\theta = 0.5\%$	0.22	0.24	0.46	0.31	0.43	0.34
$\theta = 1\%$	0.23	0.26	0.48	0.37	0.44	0.37
$\theta = 5\%$	0.25	0.32	0.52	0.45	0.47	0.43
$\theta = 10\%$	0.28	0.37	0.54	0.50	0.50	0.48
$\lvert RP/\mathrm{CVaR}_\theta \rvert$						
$\theta = 0.1\%$	0.20	0.18	0.43	0.22	0.39	0.25
$\theta = 0.5\%$	0.21	0.22	0.45	0.27	0.41	0.30
$\theta = 1\%$	0.22	0.23	0.46	0.30	0.42	0.32
$\theta = 5\%$	0.24	0.28	0.49	0.39	0.45	0.39
$\theta = 10\%$	0.25	0.31	0.51	0.43	0.47	0.42

	90-day (S&P 500) RP		180-day (S&P 500) RP		365-day (S&P 500) RP	
	discrete	log	discrete	log	discrete	log
Annualized Sharpe Ratio	1.60	2.13	1.13	1.45	0.80	0.94
$\lvert RP/\mathrm{VaR}_\theta \rvert$						
$\theta = 0.1\%$	0.39	0.27	0.37	0.26	0.37	0.31
$\theta = 0.5\%$	0.42	0.34	0.40	0.33	0.39	0.33
$\theta = 1\%$	0.43	0.36	0.41	0.35	0.39	0.34
$\theta = 5\%$	0.46	0.43	0.44	0.41	0.41	0.38
$\theta = 10\%$	0.49	0.48	0.47	0.46	0.43	0.41
$\lvert RP/\mathrm{CVaR}_\theta \rvert$						
$\theta = 0.1\%$	0.37	0.23	0.36	0.24	0.36	0.29
$\theta = 0.5\%$	0.40	0.30	0.39	0.29	0.38	0.32
$\theta = 1\%$	0.41	0.33	0.40	0.32	0.38	0.33
$\theta = 5\%$	0.44	0.38	0.42	0.37	0.40	0.35
$\theta = 10\%$	0.46	0.42	0.44	0.40	0.41	0.37

Table IV: **Risk-adjusted return measures on the S&P 500 and the OMXS30.**
Lists the return of shorting variance against the risk measures volatility, Value-at-Risk (VaR) and Conditional-Value-at-Risk (CVaR) computed for varying confidence levels. The table includes 30, 60, 90, 180 and 365 calendar day horizons for the S&P 500 and the 30-day horizon on the OMXS30.

	Arithmetic returns	Log returns
Average yearly excess return	2.37%	4.03%
Annualized Sharpe ratio	0.14	0.24
$\lvert RP/\mathrm{VaR}_\theta \rvert$		
$\theta = 0.1\%$	0.16	0.18
$\theta = 0.5\%$	0.17	0.19
$\theta = 1\%$	0.17	0.20
$\theta = 5\%$	0.21	0.25
$\theta = 10\%$	0.27	0.32
$\lvert RP/\mathrm{CVaR}_\theta \rvert$		
$\theta = 0.1\%$	0.16	0.18
$\theta = 0.5\%$	0.16	0.18
$\theta = 1\%$	0.17	0.18
$\theta = 5\%$	0.19	0.22
$\theta = 10\%$	0.21	0.25

Table V: **Risk-adjusted excess returns for the MSCI World index.**
The table lists risk measures for MSCI World using volatility, Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR). The measures are calculated on the period 1996-2007.

	30-day (OMXS30) RP		**30-day (S&P 500) RP**		**60-day (S&P 500) RP**	
	discrete	log	discrete	log	discrete	log
α coefficient	0.16	0.37	0.40	0.70	0.33	0.58
t-stat, N-W adjusted	2.484	8.409	14.16	18.19	8.85	12.94
p-values, N-W adjusted	0.007	0.000	0.000	0.000	0.000	0.000
β coefficient	3.46	2.69	3.36	4.67	3.60	4.68
t-stat, N-W adjusted	2.50	3.61	5.09	5.97	5.38	7.51
p-values, N-W adjusted	0.006	0.000	0.000	0.000	0.000	0.000
R^2	9.1%	11.1%	12.8%	12.5%	26.3%	26.7%

	90-day (S&P 500) RP		**180-day (S&P 500) RP**		**365-day (S&P 500) RP**	
	discrete	log	discrete	log	discrete	log
α coefficient	0.30	0.54	0.27	0.49	0.25	0.43
t-stat, N-W adjusted	6.85	10.52	4.70	7.04	3.16	3.78
p-values, N-W adjusted	0.000	0.000	0.000	0.000	0.001	0.000
β coefficient	3.34	4.33	2.29	3.00	1.00	1.43
t-stat, N-W adjusted	5.00	7.26	3.84	6.11	2.18	2.43
p-values, N-W adjusted	0.000	0.000	0.000	0.000	0.015	0.008
R^2	32.5%	34.1%	33.9%	35.1%	21.6%	22.7%

Table VI: **CAPM regressions results for the S&P 500 and OMXS30.**
The estimated regression is given by:
$-\widehat{RP}_{t,T} = \alpha_i + \beta_i R^M_{e,T} + \varepsilon_{i,T}$
where $\widehat{RP}_{t,T}$ is the risk premium from t to T and $R^M_{e,T}$ is the market excess returns for the same period. The t-statistica are adjusted using Newey-West (1987). The Jarque-Bera, Breusch-Godfrey and Koenker-Basset tests for normality, autocorrelation (lag = calendar days during variance swap maturity) and heteroscedasticity respectively, have been employed, rejecting their respective null hypotheses.

S&P 500 (1996-2007)	Payoff of short 30-day variance swap	MSCI World (excess returns)
Average monthly payoff (USD)	2.988	0.375
Monthly volatility	4.308%	4.308%
OMXS30 (1999-2007)	Payoff of short 30-day variance swap	MSCI World (excess returns)
Average monthly payoff (SEK)	3.073	0.392
Monthly volatility	6.289%	6.298%

Table VII: **Average payoffs and standard deviations of the two zero-cost variance investment strategies for the S&P 500 and OMXS30.**

Displays the average monthly payoff and volatility of a rolling strategy entering a short variance swap contract with 30-day maturity at each expiry day of the preceding one for the period given. For S&P 500 the period is given by 1996-2007 and for OMXS30 1999-2007. The notional for the variance swap is set such that the volatility of the short variance strategy corresponds to a portfolio borrowing in the money market at the risk-free rate in order to buy the MSCI World index. The figures relate to the payoff on a zero-cost investment in the MSCI World index for the same period is given as reference.

	30-day (OMXS30) RP		**30-day (S&P 500) RP**		**60-day (S&P 500) RP**	
	discrete	log	discrete	log	discrete	log
Intercept coefficient $(a,\ c)$	0.01	-0.69	0.00	-0.92	0.01	-0.80
t-stat, N-W adjusted	1.85	-4.34	1.02	-4.06	1.88	-3.10
p-values, N-W adjusted	0.032	0.000	0.154	0.000	0.030	0.001
Slope coefficient $(b,\ d)$	0.65	0.89	0.52	0.94	0.50	0.94
t-stat, N-W adjusted	-5.18	-2.16	-6.28	-0.92	-6.27	-0.79
p-values, N-W adjusted	0.000	0.016	0.000	0.178	0.000	0.214
R^2	47.2%	64.0%	36.6%	50.0%	31.8%	50.5%
	90-day (S&P 500) RP		**180-day (S&P 500) RP**		**365-day (S&P 500) RP**	
	discrete	log	discrete	log	discrete	log
Intercept coefficient $(a,\ c)$	0.01	-0.81	0.01	-0.86	0.01	-0.79
t-stat, N-W adjusted	-6.1	-0.83	-6.96	-0.88	-4.86	-0.52
p-values, N-W adjusted	0.000	0.203	0.000	0.19	0.000	0.301
Slope coefficient $(b,\ d)$	0.46	0.93	0.38	0.90	0.42	0.91
t-stat, N-W adjusted	5.27	10.57	4.35	8.17	3.46	5.38
p-values, N-W adjusted	0.000	0.000	0.000	0.000	0.000	0.000
R^2	28.2%	49.0%	23.0%	45.6%	27.6%	44.1%

Table VIII: **Expectation hypothesis regression results on the S&P 500 and OMXS30.**

The estimated regression is given by:

$\widehat{RV}_{t,T} = a + b\widehat{K}_{\mathbf{VAR},t,T} + \omega_{t,T}$

and

$\ln \widehat{RV}_{t,T} = c + d \ln \widehat{K}_{\mathbf{VAR}t,T} + \varphi_{t,T}$

The t-statistica are adjusted using Newey-West (1987). The Jarque-Bera, Breusch-Godfrey and Koenker-Basset tests for normality, autocorrelation (lag = calendar days during variance swap maturity) and heteroscedasticity respectively, have been employed, rejecting their respective null hypotheses.

	30-day (OMXS30) RP			30-day (S&P 500) RP			60-day (S&P 500) RP		
	payoff	discrete	log	payoff	discrete	log	payoff	discrete	log
Real. forward var.	-1.99***	-16.3%*	-36.7%	-2.09	-41.2%	-70.3%	-1.73	-34.7%	-59.5%
Real. var. (-mean)	-2.05	-17.5%**	-37.1%	-2.14	-41.5%	-71.7%	-1.77	-35.5%	-60.1%
Number of op. $>$ 2	-2.01	-17.1%**	-37.0%	-2.07	-40.2%	-70.2%	-1.73	-34.6%	-59.1%
$>$ 1 op. around ATM	-2.03	-17.3%**	-37.1%	-2.10	-41.3%	-71.3%	-1.77	-35.4%	-60.0%
Trunc. ($\pm$2 st dev.)	-2.03	-17.3%**	-37.1%	-2.10	-41.4%	-71.5%	-1.77	-35.5%	-60.1%
Trunc. ($\pm$16 st dev.)	-2.03	-17.3%**	-37.1%	-2.09	-41.4%	-71.5%	-1.77	-35.5%	-60.1%
Discr. (50,000)	-2.03	-17.3%**	-37.1%	-2.10	-41.4%	-71.5%	-1.77	-35.5%	-60.1%
Discr. (50)	-2.03	-17.3%**	-37.1%	-2.16	-38.8%	-69.7%	-1.79	-35.2%	-60.0%
Linear interp.	-2.06	-17.6%**	-37.4%	-2.36	-43.5%	-76.0%	-1.91	-36.8%	-62.7%
Linear extrap.	-2.06	-17.4%**	-37.2%	-2.14	-41.5%	-71.8%	-1.77	-35.5%	-60.1%
Subsample (-'03)	-2.91	-22.5%***	-38.9%	-2.37	-34.8%	-61.1%	-1.93	-27.4%	-48.6%
Subsample ('03-'07)	-1.19**	-12.3%^{ns}	-35.3%	-1.78	-50.5%	-85.8%	-1.62	-48.2%	-77.6%
Liquidity (bid)	0.16^{ns}	4.7%^{ns}	-12.4%**	-1.46	-30.2%	-52.9%	-1.37	-28.3%	-48.7%

	90-day (S&P 500) RP			180-day (S&P 500) RP			365-day (S&P 500) RP		
	payoff	discrete	log	payoff	discrete	log	payoff	discrete	log
Real. forward var.	-1.68	-32.9%	-55.4%	-1.72	-31.1%	-50.1%	-1.70**	-27.9%**	-43.3%**
Real. var. (-mean)	-1.72	-33.8%	-56.8%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Number of op. $>$ 2	-1.70	-33.4%	-56.5%	-1.71	-31.9%	-52.7%	-1.62***	-28.8%*	-46.0%**
$>$ 1 op. around ATM	-1.72	-33.7%	-56.8%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Trunc. ($\pm$2 st dev.)	-1.72	-33.8%	-56.8%	-1.71	-31.9%	-52.6%	-1.61***	-28.9%***	-46.0%***
Trunc. ($\pm$16 st dev.)	-1.72	-33.8%	-56.9%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Discr. (50,000)	-1.72	-33.8%	-56.9%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Discr. (50)	-1.73	-33.8%	-57.0%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Linear interp.	-1.80	-34.7%	-58.5%	-1.71	-32.3%	-53.0%	-1.63***	-29.0%***	-46.4%***
Linear extrap.	-1.72	-33.8%	-56.9%	-1.71	-32.0%	-52.7%	-1.62***	-28.9%***	-46.2%***
Subsample (-'03)	-1.84	-24.9%***	-44.4%	-1.78***	-22.1%**	-37.5%	-1.58**	-17.4%*	-26.7%**
Subsample ('03-'07)	-1.62	-48.0%	-76.5%	-1.70	-48.5%	-77.6%	-1.83***	-49.0%	-80.4%
Liquidity (bid)	-1.39	-28.0%	-47.9%	-1.47	-27.7%	-46.3%	-1.44***	-25.8%**	-41.8%***
OTC (GS & BarCap)	-1.46	-29.3%	-50.0%	-1.58***	-29.7%***	-50.5%	-1.86***	-32.2%***	-52.4%***
OTC (-30 bp bid-ask)	-1.34	-27.0%	-46.7%	-1.46***	-27.4%**	-47.4%	-1.73***	-30.1%**	-49.3%***

Table IX: **Robustness analysis.**

Estimates of the average variance risk premiums and significance levels. Note that the earlier subsample on the S&P is from 1996-2003, whereas for OMXS30 it is only during 1999-2007. The significance of the average risk premiums is indicated by the following superscripts: ^{ns} = not significant; $^{*} = 0.01 < p \leq 0.05$; $^{**} = 0.001 < p \leq 0.01$; $^{***} = 0.0001 < p \leq 0.001$; no star indicates that the p-value is lower than 0.0001.

F Figures

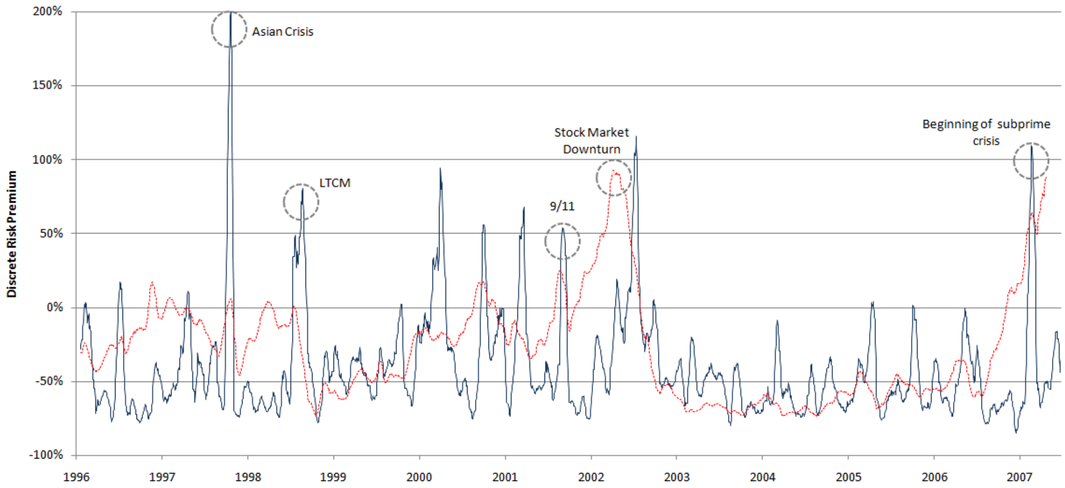

Figure 1: **15-day moving average of $RP_{discrete}$ between January 1996 and June 2007 for S%P 500, 30-day (blue line) and 365-day (red dotted line).**

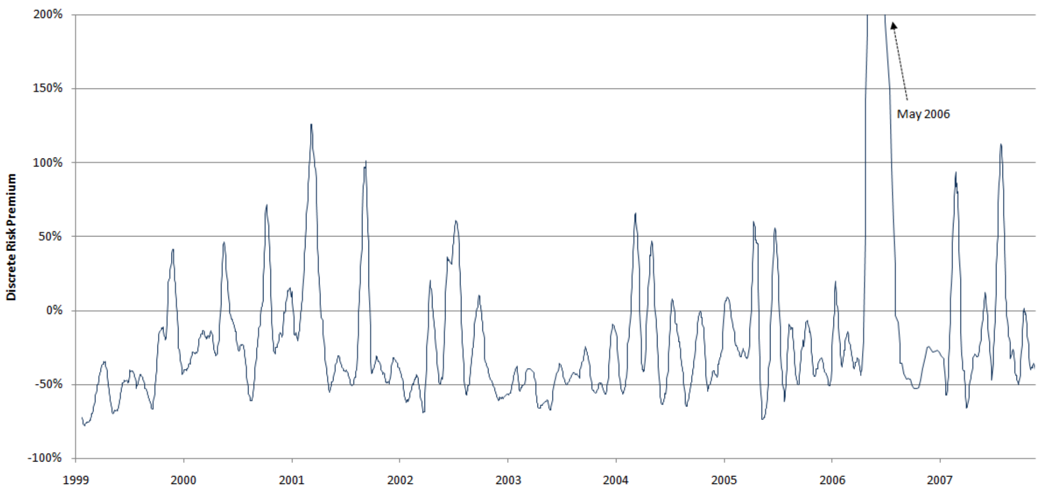

Figure 2: **15-day moving average of 30-day $RP_{discrete}$ between January 1999 and December 2007 for OMXS30.**

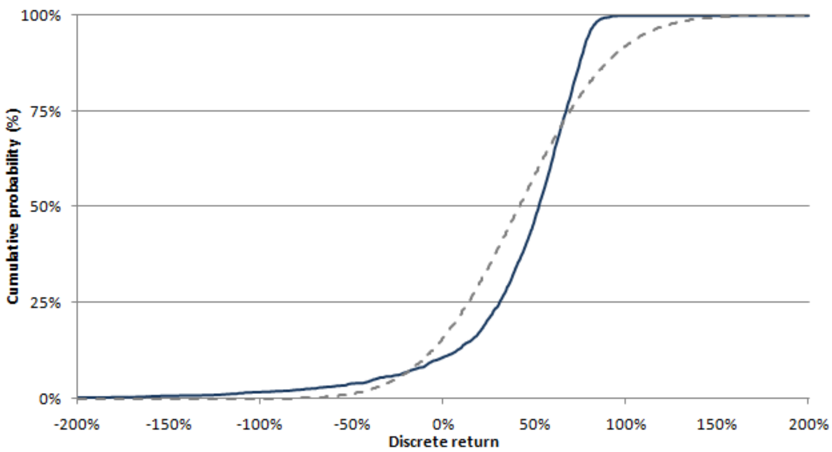

Figure 3: **Historical cumulative distribution of short $RP_{discrete}$ for 30-day S&P 500. The dotted line is the superimposed cumulative normal distribution.**

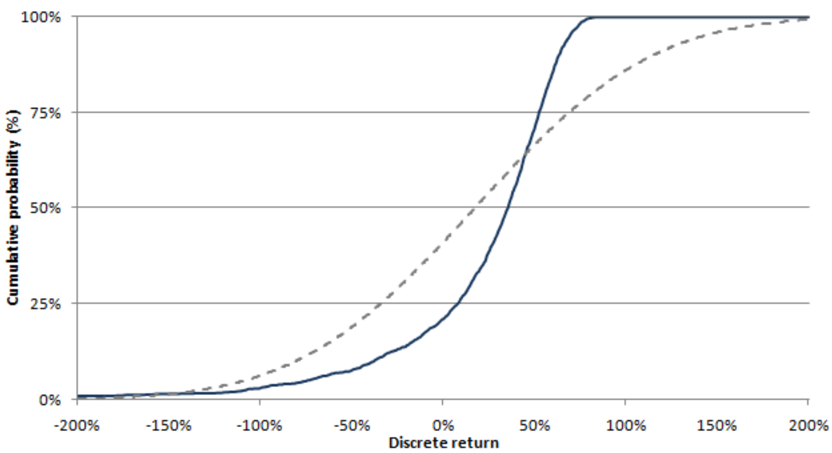

Figure 4: **Historical cumulative distribution of short $RP_{discrete}$ for 30-day OMXS30. The dotted line is the superimposed cumulative normal distribution.**

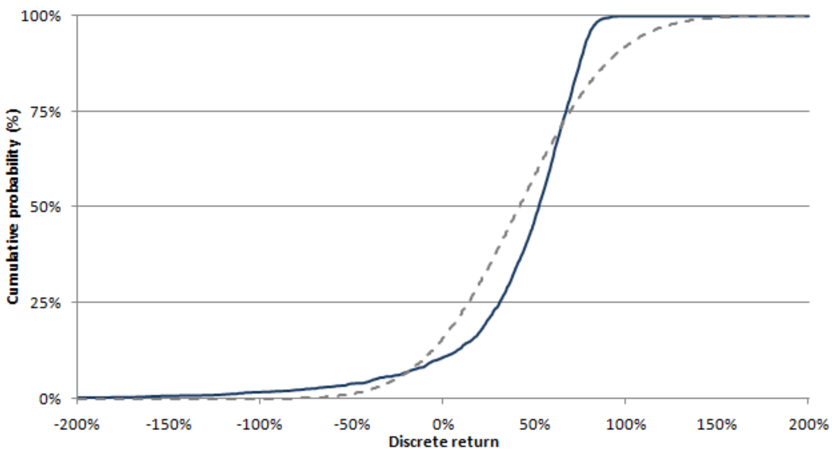

Figure 5: **Historical cumulative distribution of short $RP_{discrete}$ for 365-day S&P 500. The dotted line is the superimposed cumulative normal distribution.**

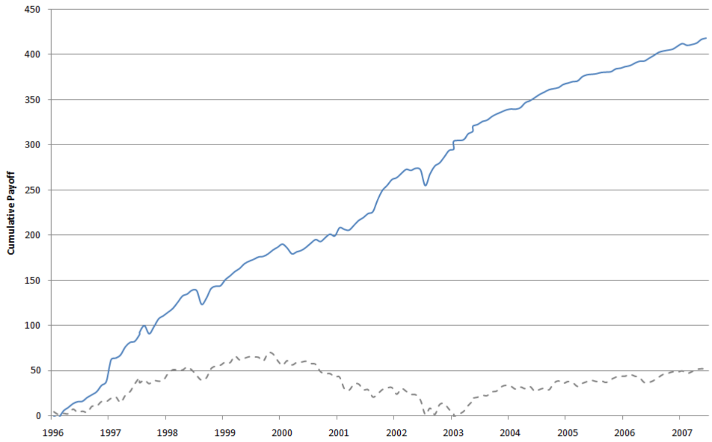

Figure 6: **Cumulative payoffs of zero-cost strategy of shorting 30-day variance swaps on the S&P 500 versus a zero-cost investment in the MSCI World index (dotted line) during 1996-2007.**

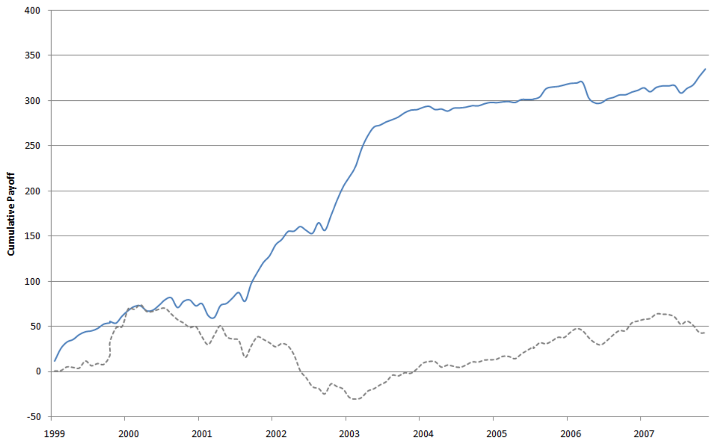

Figure 7: **Cumulative payoffs of zero-cost strategy of shorting 30-day variance swaps on the OMXS30 versus a zero-cost investment in the MSCI World index (dotted line) during 1999-2007.**