WHITEPAPER

VSTOXX 101: Understanding Europe's volatility benchmark

Hamish Seegopaul, Global Head of Index Product Innovation, STOXX Thomas Shuttlewood, Product Research and Development, Associate Vice President, STOXX

October 2024



Contents

1.	Introduction	3
2.	Volatility theory	3
	 2.1 Interpreting volatility 2.2 Calculating realized volatility 2.3 Calculating implied volatility 2.4 Capturing implied volatility – Challenges with single options 2.5 Capturing implied volatility using variance swaps 	3 5 6 8 9
3.	The VSTOXX – A (re) introduction	12
	 3.1 Intuition 3.2 The EURO STOXX 50 3.3 A (very) brief background to the VSTOXX 3.4 A brief overview of the VSTOXX 	12 12 12 13
4.	The VSTOXX – Core methodology	15
	 4.1 Option price screening 4.2 Determining the ATM strike price 4.3 Variance and subindex calculation 4.4 Option weights 4.5 Settlement values 4.6 Tick verification 	16 17 19 21 22 22
5.	A deeper dive – A historical analysis	23
	 5.1 VSTOXX history – An overview 5.2 VSTOXX versus realized volatility 5.3 The VSTOXX and the EURO STOXX 50 – An overview 5.4 Exploring the correlation 5.5 Inherent behavior of the VSTOXX 	23 24 26 26 28
6.	Use cases	31
	6.1 Trading strategies6.2 Additional STOXX volatility indices	31 32
7.	Conclusion	33
8.	References	34
9.	Appendix	35
10.	. Offices and contacts	36

1. Introduction

Uncertainty is a core feature of investment markets: A range of different performance outcomes is possible at any given minute, day, month or year. *Volatility* helps us quantify the probabilities of these potential outcomes, which, in turn, can help us anticipate the future.

Volatility's unique behavioral characteristics – its tendencies to revert to the mean, to be negatively correlated with its underlying asset and to cluster in regimes – have also led to its acceptance as a separate asset class. Today, it can be an important part of an investor's toolkit.

The VSTOXX[®] (EURO STOXX Volatility) index was launched in 2005 and represented a quantum leap forward in the measurement of European volatility at the time. In 2009 and 2010, Eurex launched futures and options on the VSTOXX, providing investors with new tools to access the asset class. The index is still as relevant as ever today, at a time when the era of globalization is ending and there is greater potential differentiation by region.

This paper reintroduces the VSTOXX. It is divided into four sections: the index's theoretical underpinnings, its calculation methodology, resulting behavior and potential use cases.

2. Volatility theory

Before discussing the approach taken by the VSTOXX, we will look at the theory behind the concept of volatility, and how the latter can be measured. An entire book could be written about this, but we will scratch the surface of the issues just enough to understand how the VSTOXX methodology has been derived. Please note, however, that this section contains some math and formulae. A summary of the key points is given at the start of section 3 for anyone who wants to skip this.

2.1. Interpreting volatility

There are a number of different types of volatility that investors may come across:

- Realized volatility i.e., past volatility calculated using historical returns
- Predicted volatility i.e., future volatility estimated using historical returns and a model
- Implied volatility i.e., market-implied volatility modeled using option prices

Underpinning all these calculations is the assumption that the logarithmic returns¹ of asset prices are normally distributed. This means that, when plotted as a histogram, the returns should be shaped like a bell ("bell curve"). The distribution of returns for the EURO STOXX 50 index (SX5E) is shown in Figure 1.

¹ Asset prices are generally bounded by 0 and infinity, and are therefore said to be lognormally distributed.



Figure 1: Distribution of daily EURO STOXX 50 (price return in EUR) log returns since 1999.

Keen observers will note that the distribution is not perfectly symmetric (this phenomenon is known as "skewness"), and that it also contains some "fat tails," i.e., some observations are further away from the mean than expected (a phenomenon known as "kurtosis"). This is to be expected for any asset, and is why volatility must always be considered along with the assumptions made in connection with it.

Volatility is often quoted as a one standard deviation move per year. This leads to a straightforward interpretation: For example, given a volatility level of 20%, using the normal distribution – the "bell curve" previously seen, where the shape is determined by the standard deviation – implies that:

- There is a ~16% chance that the return over one year will be 20% higher than the current level
- There is a ~16% chance that the return over one year will be 20% lower than the current level
- There is a ~2% chance that the return over one year will be 40% higher than the current level
- ...and so on

These likelihoods are why volatility is said to be a measure of uncertainty – the higher the volatility, the wider the range of likely potential outcomes. What is more, this effect can be compounded over time. Illustrating this, the charts below show randomly simulated returns for 10,000 portfolios with different volatility levels over a 10-year period. Normally distributed annual log returns are assumed.

Source: STOXX, Jan 1999 - May 2024.



Figure 2: Simulated stock returns under different volatility assumptions.

Source: STOXX.

2.2. Calculating realized volatility

The formula used to calculate volatility for a single asset is the square root of the variance of an asset's price from its mean, i.e., its standard deviation:

Equation (1) Annualized realized volatility =
$$\sqrt{\frac{1}{n-1} \times \sum_{i=1}^{n} \log returns^2} \times \sqrt{reasurement periods per year}$$

As we can see, if price movements are small, the resulting standard deviation will be small as well. The reverse applies if prices are much further away from the average.

Realized index volatility can be measured using the same formula; however, it is also important to note how the relationships between the underlying assets will impact the volatility. The method used to calculate the standard deviation of a portfolio is shown below, with weight (w), standard deviation (σ) and correlation (ρ) – which measures the degree to which two assets move in relation to each other – being used as inputs:

Equation (2) Portfolio volatility (two-asset portfolio) =
$$\sqrt{w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2w_1 w_2 \sigma_1 \sigma_2 \rho_{12}}$$

Equation (3) Portfolio volatility (three-asset portfolio) =

 $\sqrt{w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + w_3^2\sigma_3^2 + 2w_1w_2\sigma_1\sigma_2\rho_{12} + 2w_1w_3\sigma_1\sigma_3\rho_{13} + 2w_2w_3\sigma_2\sigma_3\rho_{23}}$

Here, the role that correlation plays in relation to index volatility is evident – the latter rises in line with the former. This is an important feature of the volatility asset class.

Illustrating this, the chart below shows a portfolio volatility heat map for two theoretical assets, both with volatilities of 15%. The y-axis varies their correlation between -1 and 1, while the x-axis varies the weight of Asset 1 (the weight of Asset 2 is 100% – Asset 1). As can be seen, the higher the correlation (or concentration), the higher the volatility. The portfolio volatility for each square was simulated using corresponding correlation levels and weights.





Source: STOXX.

2.3 Calculating implied volatility

Options confer the right (but not the obligation) to buy or sell an asset at a specified price and within a specified time window. They are defined by unique underlyings, which may be calls (rights to buy), puts (rights to sell), strike prices (exercise prices) or expiry dates. Call options are said to be "in the money" (ITM) when the stock price exceeds the strike price (strike) and "out of the money" (OTM) when the strike price exceeds the stock price, while the opposite applies to put options. Options that can only be exercised at expiry are referred to as "European-style options," whereas options that can be exercised at any time are called "American-style options" (other variations also exist).

Traded option prices are rooted in market dynamics, i.e., in supply/demand. Prices can also be modeled based on the theoretical cost to hedge an option position – this breakthrough by Black and Scholes in 1976 resulted in the now-famous formula shown below. For simplicity's sake, we have focused on European-style call options for stocks that do not pay dividends.

Equation (4) Call option price = $S_0N(d_1) - Ke^{-rT}N(d_2)$, with $d_1 = \frac{\ln \frac{S_0}{K} + \left(r + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T}$,

where S_0 = the current stock price, K = the strike price of the option, r = the risk-free interest rate, T = the time to expiration (in years), N = the cumulative distribution function of the standard normal distribution and σ = the standard deviation of the stock's returns.

Although the formula may appear daunting to the uninitiated, if one understands that the intrinsic value of a call option is the greater of (stock price – strike price) and 0, then the Black-Scholes formula can be thought of as a probability-adjusted version of this, i.e.,

Equation (5) Call option price = stock price (adjusted for probability of being above strike) – strike price (present value adjusted for probability of being exercised)

The probabilities used are derived from concepts already presented in this paper, namely the normal distribution and volatility/standard deviations. Note that this also means that the model outputs must be taken with an equally large grain of salt.

Using an observed option price and a model such as Black-Scholes, one can reverse engineer the *implied volatility*, i.e., the level of volatility needed to make the modeled option price match the actual price. Put-call parity dictates that the implied volatility of puts and calls for given strikes and expiration dates will be the same. Given this direct relationship, many volatility traders will think about, and quote, options in terms of implied volatility.

Figure 4: Pricing inputs.

Option price inputs	Implied volatility inputs
Stock price	Stock price
Strike price	Strike price
Risk-free rate	Risk-free rate
Time to expiry	Time to expiry
Stock volatility	-
-	Option price

Source: STOXX.

While it is tempting to simply use the resulting implied volatility figure as a volatility forecast, it is worth stressing again that this model of volatility is based on prices, which in turn are based on supply and demand. There is also lasting evidence of a volatility risk premium, which means that implied volatility typically exceeds realized volatility. This can be thought of as the price of uncertainty about future volatility.

2.4. Capturing implied volatility – Challenges with single options

Unfortunately, ascertaining whether implied volatility is going up or down is not as simple as buying or selling an option. There are two main challenges:

2.4.1 Exposure to "the Greeks"

Option prices are influenced by several factors, which are commonly measured in terms of sensitivities referred to as "the Greeks." The sensitivity of an option price to each of the Greeks can also be extrapolated from the Black-Scholes model.

- Delta: Sensitivity to changes in the underlying
- Gamma: Sensitivity of the delta to changes in the underlying
- Theta: Sensitivity to the passage of time
- Vega: Sensitivity to changes in volatility

In practical terms, this means that the implied volatility of an option can change simply as a result of a change in (for example) the price of the underlying stock. Continuous delta hedging (buying or selling stock to offset the delta exposure) can mitigate this sensitivity, but requires active management and introduces path dependency.

2.4.2 The shape of volatility

Since implied volatility is based on option prices, which are set by supply/demand, it is not static across strike prices or expiration dates.

As we saw in section 2.1, one of the main assumptions underlying volatility is that returns are normally distributed. This implies that the chance of (for example) a three standard deviation move occurring is only 0.4%. However, in the 1,000 or so trading days in 2020, 2021, 2022 and 2023, a three standard deviation daily move occurred 12 times on the EURO STOXX 50 (roughly 1.2% of the total). As such "fat tails" are not uncommon, traders are unwilling to sell far OTM options without an additional premium over and above what the model implies. This can be seen from the "volatility smile" in the chart below, in which the ends (i.e., those areas further away from the at-the-money (ATM) strike) show an elevated level of implied volatility. The chart also shows that this relationship changes depending on the tenor of the option, a concept known as the "term structure."



Figure 5: Implied volatility surface of the EURO STOXX 50 index.

Source: STOXX, Refinitiv. As of July 2024.

2.5 Capturing implied volatility using variance swaps

If using options is a challenging way of providing pure exposure to volatility, what would be the alternative? This is where variance swaps come into the picture. These are instruments that pay the difference between the realized variance and the variance level agreed at the start of the swap. More important, however, is the fact that the payoff can be replicated using a prespecified portfolio of options. This was highlighted in a seminal paper on the topic from Goldman Sachs (Demeterfi, 1999). The VSTOXX itself was originally developed by Goldman Sachs and Deutsche Börse.

The theory behind variance swap pricing is more challenging to intuitively explain than Black-Scholes, but is based on similar underpinnings. As stated in the previous section, the underlying asset price can influence the price of an option. The sensitivity to changes in variance is greatest when the underlying asset is closest to the strike price, and also increases as the underlying price rises.

The figures on page 10 are taken from the Goldman paper and show the exposure to variance (y-axis) versus the stock price (x-axis). Each curve represents the exposure of a single option with a different strike, and illustrates the higher sensitivities at the ATM level and for higher prices.



Figure 6: Exposure to variance by different option strikes.

Source: Goldman Sachs (Demeterfi, 1999).

This sensitivity to changes in the underlying asset price can be neutralized by creating a portfolio of OTM puts and calls, weighted by the inverse of the squared strike (1/strike²). As can be seen in the series of graphics below, which are again taken from Demeterfi (1999), the inverse method yields the least sensitivity (until you run out of options).

Figure 7: Exposure to variance, aggregated by methods in Figure 6.



Source: Goldman Sachs (Demeterfi, 1999).

Deriving a formula to provide this exposure starts with solving the level of variance the market expects, i.e., the strike of a variance swap at which the expected value of the variance, minus the realized variance, is equal to 0. Much of the math and intuition will be skipped in this section, but can be found in the papers given in the "References" section.

To start with, the variance is the average of all the variances (measured using log returns) during the lifetime of the swap (t = 0 to t = T).

Equation (6) Variance $_{T} = \frac{1}{T} \int_{0}^{T} \sigma_{t}^{2} dt$,

where σ is the standard deviation of the asset.

The variance swap rate (K_{var}), which is the expected value of the future variance, can be derived by taking the expected value of a continuously rebalanced stock position (which captures intraperiod returns) for a contract that pays off the negative log return of the stock over the full time period (t = 0 to t = T).

Equation (7)
$$K_{var} = E \left[Variance_T \right] = \frac{2}{T} E \left[\int_0^T \frac{dS_t}{S_t} - \ln \frac{S_t}{S_0} \right],$$

where S_0 is the asset price at the start of the swap period, S_t is the price at each time interval and S_T is the price at the end.

The first term in the bracket can be reduced to simply the risk-free rate. For the second term, a contract that pays off the log return does not readily exist but can be manufactured using a combination of stock positions and the portfolio of OTM put and call options. OTM options are favored because they are generally more liquid. If one also assumes a forward price (F_T) that is the future expected stock price (using the risk-free rate), the formula is:

Equation (8)
$$K_{var} = \frac{2}{T} e^{rT} \left\{ \int_{0}^{F_{T}} \frac{1}{K^{2}} P(K) dK + \int_{F_{T}}^{\infty} \frac{1}{K^{2}} C(K) dK \right\},\$$

where r is the risk-free rate, P(K) is the put price for strike K and C(K) is the call price for strike K.

Note that the formula is continuous, i.e., it works perfectly when there is an infinite pool of options with reliable prices to choose from, and does not exhibit price jumps. To make this practical for use with the listed options market, the formula must be translated into a version that allows for a discrete set of options to be used (see below). Note the additional term at the end, which accounts for switching from forward prices in Equation 8 to discrete strike prices in Equation 9.

Equation (9)
$$K_{var}^* = \frac{2}{T} e^{rT} \left(\sum_{K_j \le K_0} \frac{\Delta K_j}{K_j^2} P(K_j) + \sum_{K_j > K_0} \frac{\Delta K_j}{K_j^2} C(K_j) \right) - \frac{1}{T} \left(\frac{F}{K_0} - 1 \right)^2$$
,

where K_0 is the strike price for the ATM option, ΔK_j is the increment between strike prices and F is the forward price of the asset at time T.

Equation 9 is used as the basis for calculating the VSTOXX. However, this method has additional limitations:

- This paper has focused on volatility so far, but here we are switching to variance. This is because replicating a variance swap using options is easier than replicating a volatility swap, making the former a more popular instrument. As variance is volatility squared, this also means that variance is convex with respect to volatility (for example, where the volatility versus the original variance swap strike is higher, the variance is also much higher).
- Since the portfolio is weighted by 1/K², more weight is given to lower strikes, i.e., far OTM puts (which tend to be less liquid/more expensive).

In the case of the VSTOXX, the convexity issue is solved by recalculating the variance at every index tick. This is a trade-off: It makes the index impractical to replicate (since each print would require a slightly different options portfolio) but ensures consistent measurement of volatility. To solve the liquidity issue, options are screened for trade data (see the following sections for more details).

3. The VSTOXX – A (re)introduction

Although the VSTOXX was launched in 2005, we are taking the opportunity here to provide a high-level overview of its objective and importance. Then, in the section after this, we take a deeper dive into the calculation methodology.

3.1 Intuition

For those readers who skipped the "Volatility theory" section, the intuition can be summarized as follows:

- Volatility measures the uncertainty of a stock price's returns.
- Option pricing theory allows us to derive the market's expectation of a stock's volatility, and therefore to rely on the "wisdom of crowds."
- An index's volatility is based both on the volatility of its components and on these components' relationships with one another.
- Financial engineering allows us to construct a portfolio of options which can track the market's expectation of volatility, and only this. This is what underlies the VSTOXX's construction.
- The higher the uncertainty surrounding a stock price is, the higher the market expectation of volatility, the higher the option price, the higher the cost of the option portfolio and the higher the VSTOXX.

3.2 The EURO STOXX 50

While this paper clearly focuses on the VSTOXX, it is important to consider the index, and the associated derivatives, that it is based on.

The EURO STOXX 50 was launched in 1998 with the express aim of becoming the benchmark for the European market. It captures the largest companies domiciled in the eurozone, while ensuring diversification by selecting the leading blue-chip entities from each of the 20 Industry Classification Benchmark (ICB) supersectors. This ensures it covers all parts of the European market, from technology to real estate, with additional capping reducing the chances of oversaturation and concentration. The EURO STOXX 50 has gone on to become a barometer for the eurozone economy and one of the most significant benchmarks in the world. It currently underlies more than EUR 28 billion in ETF assets, and its index derivatives are the most actively traded euro-denominated equity index derivatives on Eurex (data as of December 2023).² The EURO STOXX 50 also serves as the basis for a number of derived strategy indices, which are used across a broad suite of structured products. This level of market participation, and especially the deep derivatives market, facilitate calculation of the VSTOXX.

3.3 A (very) brief background to the VSTOXX

The launch of the VSTOXX in 2005 was a significant moment for the European market. For many years previously, realized volatility had been used as a metric to inform investment decisions; a publicly available measure of implied volatility was less ubiquitous due to the complexity involved in its calculation. The difficulties are obvious: Real-time measurement requires accurate and reliable option prices across a range of strikes, and these must also be delivered in a timely manner. However, the advantages of such an index are equally clear: Using implied volatility, one can successfully design a product or trading strategy to reflect market expectations of volatility over a given time frame.

² STOXX.com/index/SX5E/

Not only did the VSTOXX become the gauge of market volatility, but it also helped broaden access to volatility as an investable asset class in Europe. Prior to the launch of listed derivatives on the VSTOXX, volatility trading strategies – in the form of volatility swap- or variance swap-like payoffs – were out of reach for the average investor. The VSTOXX, and the deep suite of derivatives based on it, opened up this investment space, with high liquidity and interest showcasing the market's appetite. This can be seen from Figure 8, which shows the ten most liquid Eurex-traded derivatives based on indices calculated by STOXX.

Underlying index	Contract type	Traded contracts	Open interest (in EUR)
EURO STOXX 50	Futures	25,155,416	228,395,716,661
EURO STOXX 50	Options	24,298,936	1,294,723,660,123
EURO STOXX Banks	Futures	7,796,017	12,166,820,644
EURO STOXX Banks	Options	5,297,599	55,459,143,700
STOXX Europe 600	Futures	3,302,424	12,163,658,925
DAX	Futures	2,340,851	26,221,007,890
VSTOXX (EURO STOXX 50 Volatility)	Futures	1,531,003	375,448,805
DAX	Options	1,248,512	89,847,338,450
VSTOXX (EURO STOXX 50 Volatility)	Options	1,045,562	1,803,783,050
EURO STOXX 50 DVP	Futures	558,948	9,394,437,020

Figure 8: Eurex derivatives data for indices calculated by STOXX (as of June 2024).

Source: Eurex.

3.4 A brief overview of the VSTOXX

As previously mentioned, the VSTOXX is designed to reflect market expectations of volatility, and captures implied volatility by taking the square root of the implied variance (as seen in the "Volatility theory" section). The VSTOXX covers all listed EURO STOXX 50 options (Eurex product ID: OESX) with a given time to expiration, assuming that these meet certain pricing criteria.

Although so far we have largely talked about the "VSTOXX," which measures 30-day volatility, we should perhaps have referred to the "VSTOXX main indices." A total of twelve such indices cover a range of time horizons, since they have been calculated for rolling 30-, 60-, 90-, 120-, 150-, 180-, 210-, 240-, 270-, 300-, 330- and 360-day expiries. These indices effectively have constant times to expiry and do not expire, eliminating the effects of volatility fluctuations occurring close to option contract expiries. The VSTOXX main indices are calculated using a linear combination of two out of a total of eight subindices covering EURO STOXX 50 option expiries, ranging from one month to two years. A list of these main indices and subindices is given in Figure 9.

Figure 9: VSTOXX main indices and subindices.

Index name	Index type	STOXX index symbol
VSTOXX	Main index	V2TX
VSTOXX 60 Days	Main index	VSTX60
VSTOXX 90 Days	Main index	VSTX90
VSTOXX 120 Days	Main index	VSTX120
VSTOXX 150 Days	Main index	VSTX150
VSTOXX 180 Days	Main index	VSTX180
VSTOXX 210 Days	Main index	VSTX210
VSTOXX 240 Days	Main index	VSTX240
VSTOXX 270 Days	Main index	VSTX270
VSTOXX 300 Days	Main index	VSTX300
VSTOXX 330 Days	Main index	VSTX330
VSTOXX 360 Days	Main index	VSTX360
VSTOXX 1M	Subindex	V6I1
VSTOXX 2M	Subindex	V6I2
VSTOXX 3M	Subindex	V6I3
VSTOXX 6M	Subindex	V6I4
VSTOXX 9M	Subindex	V6I5
VSTOXX 12M	Subindex	V6I6
VSTOXX 18M	Subindex	V6I7
VSTOXX 24M	Subindex	V6I8

Source: STOXX.

The VSTOXX model allows volatility to be (technically) replicated using an option portfolio that does not react to price fluctuations but only to changes in volatility. As explained in the "Volatility theory" section, this is achieved by directly replicating variance rather than volatility itself, with options covering a range of strikes (with individual weightings) meeting this goal.

4. The VSTOXX – Core methodology

This section takes a deeper dive into the VSTOXX methodology. We will start by looking at the main index formula. The indices are calculated as a time-weighted average of two subindices as follows:

Equation (10)	Main index _{tm} = 100 × $$	$\left[\frac{T_{st}}{T_{365}}X ight]$	(Subindex _{st}) 100	$\Big)^2 \times \frac{T_{lt} - T_{tm}}{T_{lt} - T_{st}} -$	$+\frac{T_{lt}}{T_{365}}\times ($	(Subindex _{lt}) 100	$\Big)^{2} \times \frac{T_{tm} - T_{st}}{T_{lt} - T_{st}} \Big] \times$	T ₃₆₅ T _{tm}

where:

- tm = the fixed time to expiry, in days, targeted by the index (i.e., 30 days for the VSTOXX and 60 days for the VSTOXX 60 Days).
- Subindex $_{st}$ = the subindex with the shorter expiry used in the calculation
- Subindex_{It} = the subindex with the longer expiry used in the calculation

T _{st}	= the number of seconds to expiry of Subindex _{st}
T _{It}	= the number of seconds to expiry of Subindex _{It}
T _{tm}	= the number of seconds in tm
T ₃₆₅	= the number of seconds in a standard year of 365 days (31,536,000 seconds)

At first glance, this may not seem to be an overly complex formula, but it is subtler than it appears. The main indices are calculated using a linear interpolation of the subindices whose times to expiry better represent the targeted fixed time to expiry of the main index. If two subindices exist whose times to expiry bracket the time to expiry targeted by the main index, the main index is calculated as an interpolation of the two subindices. When the expiry of the two subindices used to calculate a main index approaches, the times to expiry may not bracket the fixed time to expiry of the main index: In this case, the algorithm extrapolates between the two subindices. However, as soon as an interpolation between two other subindices becomes possible, the algorithm switches to the new subindex pair.

The examples below illustrate this with reference to the 30-day VSTOXX (V2TX).

Figure 10: Case 1 – VSTOXX targeted expiry and associated subindices.

Index calculation date	V2TX targeted expiry	Subindex _{st} expiry	Subindex _{It} expiry	
August 16, 2024	September 15, 2024	September 20, 2024	October 18, 2024	

Source: STOXX.

Figure 10 shows the case in which the targeted expiry is not bracketed by the subindices; in this case extrapolation is performed.

Figure 11: Case 2 – VSTOXX targeted expiry and associated subindices.

Index calculation date	V2TX targeted expiry	Subindex _{st} expiry	Subindex _{lt} expiry	
August 21, 2024	September 20, 2024	September 20, 2024	October 18, 2024	

Source: STOXX.

Figure 11 shows the case in which the targeted expiry is equal to the expiry of a subindex. In this case, the other subindex term will go to zero (as can be seen in Equation 10), meaning that only the subindex level with the shorter term will be used in the calculation.

Figure 12: Case 3 – VSTOXX targeted expiry and associated subindices.

Index calculation date	V2TX targeted expiry	Subindex _{st} expiry	Subindex _{It} expiry
August 26, 2024	September 25, 2024	September 20, 2024	October 18, 2024
August 27, 2024	September 26, 2024	September 20, 2024	October 18, 2024
August 28, 2024	September 27, 2024	September 20, 2024	October 18, 2024

Source: STOXX.

Figure 12 shows the simplest case, in which the targeted expiry of the V2TX is bracketed by shorter- and longer-term subindices, meaning that interpolation takes place.

Calculation of these subindex levels is intricate and complex. We must therefore dive even deeper, starting this time at the beginning with the task of cleansing raw options data.

4.1 Option price screening

Ensuring that the VSTOXX represents the desired economic reality depends on selecting accurate and reliable option prices. To do this, the trade price, the mid-quote (i.e., the average of the bid and ask prices) and the daily settlement price (plus the corresponding time stamps) are identified for all EURO STOXX 50 index options within the given expiry.

This is a multistep process: First, a filter is applied that ignores any trade price, mid-quote or daily settlement price below 0.5 points, so as to exclude all negligible values. Mid-quotes are then calculated if the following criteria are met:

- Both bid and ask quotes are available
- Both the bid and ask quotes are equal to or greater than 0.1 points (the mid-quote must be 0.5 or greater)
- The bid-ask spread does not exceed the following thresholds:
- In a normal market: 8% of the bid quote, with a minimum value of 1.2 points and a maximum of 18 points
- In a stressed market: 16% of the bid quote, with a minimum value of 2.4 points and a maximum of 36 points

The price with the most recent time stamp is selected for each option used in calculating a subindex. If prices with identical time stamps exist, the price used is selected on the basis of the following hierarchy:

- 1. Trade price
- 2. Mid-quote
- 3. Daily settlement price

An example of this process is given below, with the highlighted prices being the ones selected at the point in time concerned:

Strike	Settlement	Bid (time)	Ask (time)	Mid (time)	Last-traded (time)	Price
4,050	76.70	-	-	-		76.70
4,100	53.71	-	-	-	54.01 (09:05)	54.01
4,150	37,51	33.70 (09:04)	34.40 (09:05)	34.05 (09:05)		34.05
4,200	22,54	17.29 (09:04)	19.53 (09:05)	18.41 (09:05)	20.21 (09:01)	18.41

Figure 13: Example of option price selection.

Source: STOXX.

4.2 Determining the ATM strike price

Since the methodology requires OTM options, the next step in calculating the subindex is to identify the ATM strike price, which is then used as the basis for calculating the variance. The ATM strike price is defined as the highest strike not exceeding the forward ATM price for the option expiry given. This forward ATM price is derived by finding the strike with the smallest absolute difference between the call and put prices. Once identified, the forward ATM price is calculated using the following formula:

Equation (11) $F_i = K_{\min |C-P|} + R_i \times (C-P)$

where:

F_i = the forward ATM price for option expiry date i

 $K_{min|C-P|}$ = the strike for which the absolute difference between call and put prices is the smallest

R_i = the refinancing factor for option expiry i

with C and P referring to the relevant call and put prices respectively.

The refinancing factor above is defined by:

Equation (12) $R_i = e^{r_i \times (T_i / T_{365})}$

where:

 T_i = the time (in seconds) to the option expiry I, T_{365} has been previously defined and r_i is the interpolated risk-free rate valid for option expiry i. This interpolation is performed by considering the time to option expiry and incorporating an upper and lower interest rate with periods bounding the time to expiry. Further information on calculating the interpolated risk-free rate can be found in the Appendix.

Once the forward ATM price has been calculated, the highest strike not exceeding this value is selected and defined as the ATM strike.

4.2.1 ATM strike price selection – A sample calculation

It may be useful at this point to go through the steps involved using a sample calculation.

If we start by defining the variables T_i and r_i as follows:

T_i = 1,908,000 seconds

r_i = 1.41296%

then we can calculate the refinancing factor:

 $R_{i} = e^{1.41296\% \times (1,908,000 / 31,536,000)} = e^{1.41296\% \times 0.0605022831} = 1.0008552403$

The smallest absolute difference between the call and put prices, measured across all available strike prices, is then identified. This process is shown in Figure 14. It should be noted that the larger number of OTM puts is a signal of the skewness in which puts (i.e., downside risk) are often of greater significance to market participants.

K (strike price) Call price		Put price	Call price – put price
2,350	472.00	0.60	471.40
2,400	422.30	1.00	421.30
2,450	372.80	1.50	371.30
2,500	322.40	2.30	320.10
2,550	273.50	3.30	270.20
2,600	225.15	4.60	220.55
2,650	177.85	6.70	171.15
2,700	132.40	12.00	120.40
2,750	90.90	21.00	69.90
2,800	57.90	35.40	22.50
2,850	29.50	58.25	28.75
2,900	13.10	92.00	78.90
2,950	5.00	134.10	129.10
3,000	1.50	180.90	179.40
3,050	0.70	229.55	228.85
3,100	0.60	230.00	229.40

Figure 14: Option data for example portfolio.

Source: STOXX.

It can be seen that $K_{min|C-P|}$ = 2,800. We can then calculate the implied forward ATM price:

F_i = 2,800 + 1.0008552403 × 22.50 = 2,822.51924290675

The highest strike price not exceeding this value, i.e., the ATM strike, is K = 2,800. OTM option positions are taken into account using put prices for strikes below this value and call prices for strikes above this value. The price of the ATM strike is then the simple arithmetic mean of the put and call prices.

4.3 Variance and subindex calculation

Once the ATM strike has been calculated and all prerequisite inputs are in place, Equation 9 can be revisited in order to calculate implied variance for a given option expiry. Translating this into STOXX nomenclature and considering a portfolio of j options, the implied variance, σ_i^2 for the ith option expiry is expressed as:

Equation (13)
$$\sigma_i^2 = \frac{2}{T_i / T_{365}} \times \left(\sum_j \frac{\Delta K_{i,j}}{K_{i,j}^2} \times R_i \times M_{K_{i,j}} \right) - \frac{1}{T_i / T_{365}} \times \left(\frac{F_i}{K_{i,0}} - 1 \right)^2$$

where:

K_{i.0} = the ATM option strike price previously discussed

 $K_{i,i}$ = the strike price of the jth OTM option, with options being sorted by strike in ascending order

- $\Delta K_{i,j} = \text{the average price difference of the two options struck immediately above and below K_{i,j}.$ In the case of the highest and lowest strike prices, this is simply taken to be the price difference between the preceding or succeeding strike prices respectively.
- M_{Ki,0} = the inclusion price of the ATM strike option. This is simply calculated as the arithmetic mean of the call and put prices.
- M_{Ki,j} = the inclusion price of the OTM option j. As previously discussed, this is defined as the put price for strikes below K_{i,0} and the call price for strikes above K_{i,0}.

The variables T_i , T_{365} , F_i and R_i were defined earlier.

Once the implied variance for the option expiry has been calculated, we can revert to volatility by taking the square root. This leads us to the final subindex formula for the given option expiry date of:

Equation (14) Subindex_i = $100 \times \sqrt{\sigma_i^2}$

4.3.1 A sample variance calculation

At this point it may again be helpful to showcase a sample calculation using the inputs and portfolio of options from the previous section.

Expanding Figure 14, we now have:

K _{i, j} (strike price)	∆K _{i,j}	Call price	Put price	Call price – put price	M _{Ki,j}	$\frac{\Delta K_{i,j}}{K_{i,j}^2} \times R_i \times M_{K_{i,j}}$
2,350	50	472.00	0.60	471.40	0.60	0.0000054370
2,400	50	422.30	1.00	421.30	1.00	0.0000086880
2,450	50	372.80	1.50	371.30	1.50	0.0000125055
2,500	50	322.40	2.30	320.10	2.30	0.0000184157
2,550	50	273.50	3.30	270.20	3.30	0.0000253966
2,600	50	225.15	4.60	220.55	4.60	0.0000340528
2,650	50	177.85	6.70	171.15	6.70	0.0000477446
2,700	50	132.40	12.00	120.40	12.00	0.0000823749
2,750	50	90.90	21.00	69.90	21.00	0.0001389617
2,800	50	57.90	35.40	22.50	46.65	0.0002977672
2,850	50	29.50	58.25	28.75	29.50	0.0001817497
2,900	50	13.10	92.00	78.90	13.10	0.0000779501
2,950	50	5.00	134.10	129.10	5.00	0.0000287520
3,000	50	1.50	180.90	179.40	1.50	0.0000083405
3,050	50	0.70	229.55	228.85	0.70	0.0000037656
3,100	50	0.60	230.00	229.40	0.60	0.0000031244

Figure 15: Option data and calculations for example portfolio.

Source: STOXX.

Adding up the values in the column on the far right of the table above, we find that:

$$\sum_{j} \frac{\Delta K_{i,j}}{K_{i,j}^{2}} \times R_{i} \times M_{K_{i,j}} = 0.0009750263$$

We can then use the previously calculated parameters:

K_{i,0} = 2,800

F_i = 2,822.51924290675

$$T_i / T_{365} = 0.0605022831$$

to calculate the implied variance of the portfolio:

$$\sigma_i^2 = \frac{2}{0.0605022831} \times 0.0009750263 - \frac{1}{0.0605022831} \times \left(\frac{2,822.51924290675}{2,800} - 1\right)^2$$

resulting in: σ_i^2 = 0.0311619545863044

This then leads to a subindex value of: Subindex_i = $100 \times \sqrt{0.0311619545863044} = 17.65274896$

The subindex levels calculated in this way can then be used as inputs for the main index formula, resulting in the final VSTOXX value.

4.4 Option weights

The VSTOXX is designed around the principle of using an options portfolio to measure market expectation of volatility. While a discrete set of options makes this theoretically feasible, tick-by-tick reweighting of the portfolio means it is impractical. Nonetheless, determining option weightings in the VSTOXX main indices is both insightful and useful. The weight of the individual options can be obtained by simply expanding the main index formula and rearranging its terms. If we ignore the ATM adjustment term in the subindex formula, variance can be approximated as:

Equation (15)
$$\sigma_i^2 \approx \frac{2}{T_i / T_{365}} \times \sum_j \frac{\Delta K_{i,j}}{K_{i,j}^2} \times R_i \times M_{K_{i,j}} = \frac{2}{T_i / T_{365}} \times S_{K_{i,j}}$$

Inserting this expression in the main index formula results in the following:

Equation (16) Main index_{tm} = 100 ×
$$\sqrt{\left(\frac{T_{st}}{T_{365}} \times \frac{2}{T_{st}/T_{365}} \times S_{K_{i,j}} \times \frac{T_{lt} - T_{tm}}{T_{lt} - T_{st}} + \frac{T_{lt}}{T_{365}} \times \frac{2}{T_{lt}/T_{365}} \times S_{K_{lt,j}} \times \frac{T_{tm} - T_{st}}{T_{lt} - T_{st}}\right) \times \frac{T_{365}}{T_{tm}}}{T_{tm}}$$

This can be simplified to produce:

Equation (17) Main index_{tm}
$$\approx 100 \times \sqrt{2 \times \left(S_{K_{st,j}} \times \frac{T_{lt} - T_{tm}}{T_{lt} - T_{st}} + S_{K_{lt,j}} \times \frac{T_{tm} - T_{st}}{T_{lt} - T_{st}}\right) \times \frac{T_{365}}{T_{tm}}}$$

After this, we use the following definitions:

Equation (18)
$$WS_{st} = \frac{T_{st}}{T_{tm}} \times \left(\frac{T_{lt} - T_{tm}}{T_{lt} - T_{st}}\right), \qquad WS_{lt} = \frac{T_{lt}}{T_{tm}} \times \left(\frac{T_{tm} - T_{ts}}{T_{lt} - T_{st}}\right)$$

where WS_{st} and WS_{lt} are defined as the weights of the shorter- and longer-term subindices respectively. We can then define individual option values, RS, as:

$$\textbf{Equation (19)} \ \ \textbf{RS}_{j \in st} = 2 \times \textbf{R}_{st} \times \left(\frac{\Delta K_{st,j}}{K_j^2} \times \textbf{M}_{K_{i,j}} \right) \times \textbf{WS}_{st} \times \frac{\textbf{T}_{365}}{\textbf{T}_{st}}, \ \ \textbf{RS}_{j \in It} = 2 \times \textbf{R}_{It} \times \left(\frac{\Delta K_{It,j}}{K_j^2} \times \textbf{M}_{K_{i,j}} \right) \times \textbf{WS}_{It} \times \frac{\textbf{T}_{365}}{\textbf{T}_{It}},$$

and we can restate the main index approximation as:

Equation (20) Main index_{tm} \approx 100 $\times \sqrt{RS_{st} + RS_{lt}}$

with the terms RS_{st} and RS_{lt} representing the values of the option portfolios in the shorter- and longerterm subindices respectively.

The contribution of each jth option to the total market value of the portfolio is simply RS_j over the sum of RS_{st} and RS_{lt}. This leads us to the following formula for the individual option weight:

Equation (21) $W_j = \frac{RS_j}{RS_{st} + RS_{lt}}$

Daily files containing these option weights can be found on the relevant index pages on the STOXX website.

4.5 Settlement values

Determining the final index settlement values is critical given the wide-ranging and complex ecosystem of financial products based on the VSTOXX.

The final settlement day for each main index is defined as the 30th calendar day preceding the expiry of the EURO STOXX 50 options. The settlement level of each main index is calculated on the settlement day as the average of all valid ticks that index during the settlement time window. Historically, this window has been between 11:30:00 and 12:00:00 CET. However, as of the September 2024 settlement, this window has been changed to between 11:00:00 and 12:00:00 CET. The change was made to facilitate the unwinding of larger positions at expiry, hence improving liquidity and making executing volatility strategies easier.

4.6 Tick verification

Given the significance of the VSTOXX indices, each index tick is verified before publication. The maximum deviation allowed is $\pm 20\%$ for subindices and $\pm 8\%$ for main indices.

5. A deeper dive – A historical analysis

This section focuses on the VSTOXX's historical performance. An analytical focus on the data, combined with the underlying theory previously discussed, allows us to draw conclusions about the inherent behavior of the index. All analyses were performed on the 30-day expiry index (STOXX index symbol: V2TX).

5.1 VSTOXX history – An overview

We first observed the index level history as a simple, but natural starting point. Figure 16 shows the index level track of the VSTOXX from inception to the data cut-off date (July 29, 2024). Whilst it may seem a simple graph, there are a number of takeaways, one of which relates to the nature of market volatility itself. Increases in volatility can be very sharp and sudden as investors react to market shocks. These spikes are then often followed by a gradual decline in volatility over the following months as markets adapt and return to equilibrium. This return to the status quo displays the mean-reverting nature of volatility: High or low levels tend to return to the average state.

The average level across the lifetime of the index was calculated to be 23.48 index points. As discussed in the initial "Volatility theory" section, this implies (with a number of assumptions) that the EURO STOXX 50 would change by less than $\pm 23.48\%$ per annum, 68% of time.



Figure 16: Historical VSTOXX index levels.

Source: STOXX.

The graph also allows us to pinpoint milestone events in the European market. Spikes in the early period can be attributed to the September 11, 2001, attacks, and to the 2002 stock market downturn triggered by the dot-com crash and major accounting scandals. The two most prominent volatility spikes resulted from the periods of greatest market distress in recent history: the global financial crisis (GFC) of 2007–08 (which peaked in October 2008) and the COVID-19 pandemic. The most recent notable peak can be seen in early 2022, coinciding with the Russian invasion of Ukraine. Figure 17 provides details on periods that saw extremely high V2TX levels (over 45 index points) and their causes.

Figure 17: Notable events causing spikes in the VSTOXX index.

Period	Cause	Index peak value
September 2001	September 11 attacks in New York City	57.74
July 2002 – April 2003	General stock market downturn caused by the bursting of the dot.com bubble and major accounting scandals, among other things	62.73
October 2008	Global financial crisis	87.51
May 2010	US flash crash ³	49.87
August – October 2011	European sovereign debt crises	53.55
March – April 2020	COVID-19 pandemic	85.62
March 2022	Russian invasion of Ukraine	49.64

In addition to turbulence, the graph also clearly shows periods of calm, and the bullish market conditions accompanying them in most cases. Examples of this include (roughly speaking) the four years preceding the GFC and the period between 2012 and the COVID-19 pandemic.

To sum up, one can conclude that the VSTOXX market not only acts as a barometer for near-term volatility but also highlights the patterns associated with, and tells the story of, major world events and the European economy as a whole.

5.2 VSTOXX versus realized volatility

It may be useful at this point to highlight the differences between implied and realized volatility that were already discussed earlier. Comparisons can easily be made by looking at the VSTOXX and the 30-day realized volatility for the EURO STOXX 50, but let us first refresh our memories as to what causes these differences.

Realized volatility simply measures historical deviations from an average price over a given period. By contrast, implied volatility is based on market expectations. Option prices are used to measure expected volatility, and these intrinsically incorporate a risk aversion component. Since option prices tend to increase with market nervousness, participants are willing to pay a premium for protection.

³ <u>https://corporatefinanceinstitute.com/resources/equities/2010-flash-crash/</u>

The options' forward-looking nature also prices in upcoming events, from earnings reports to elections, which may dramatically alter market volatility (Fassas et al, 2021). This inherent uncertainty often leads to overestimations when compared to realized volatility.

Figure 18 shows the comparison between the 30-day realized volatility (forward-looking as of the calculation date) of the EURO STOXX 50 and the VSTOXX; the subsequent overestimation of implied volatility in the VSTOXX is clearly visible.





Source: STOXX (annualization factor of 252 used).

The average realized volatility of the EURO STOXX 50 over this period was 19.75%, with the average VSTOXX level being 23.48 index points. If we instead compare the median, and hence negate the skew caused by the most extreme values, we see values of 16.92% and 21.53 index points for the EURO STOXX 50 and VSTOXX respectively. Comparing these results, we can surmise that implied volatility was on average around 4% higher than realized volatility across the time period.

5.3 The VSTOXX and the EURO STOXX 50 – An overview

This section investigates the relationship between the VSTOXX and the EURO STOXX 50, beginning with a simple index-level comparison.





Source: STOXX.

Superimposing the data sets as in Figure 19 enables us to easily see the intrinsic nature of volatility once again. Comparing the volatility of daily returns between the VSTOXX and the EURO STOXX 50 produces values of 103.77% and 22.35% respectively. It is also clear that the indices have an inverse relationship, or, in other words, that they are negatively correlated. Significant decreases in the EURO STOXX 50 during times of market distress are met with clear spikes in VSTOXX volatility. The milestone events that have already been discussed (such as the GFC and COVID-19) are examples of this. Conversely, periods of relative calm in the EURO STOXX 50, and generally bullish market runs, coincide with markedly low levels on the part of the VSTOXX.

5.4 Exploring the correlation

Figure 20 visualizes the relationship between the returns generated by the EURO STOXX 50 and the VSTOXX indices more clearly. The daily returns for the VSTOXX and EURO STOXX 50 showed a strong overall negative correlation of –0.72.



Figure 20: Negative correlation of the VSTOXX and EURO STOXX 50 (price return in EUR) daily returns.

Source: STOXX.

Figure 21 shows the rolling 30-day correlation between the VSTOXX and EURO STOXX 50 indices. During this rolling time period, the average correlation is –0.78, with similar values being seen over 90-day and one-year periods. We can therefore conclude that, even though the overall –0.72 correlation is strong, outliers may have dampened it, with smoothing resulting in even greater values.



Figure 21: Rolling 30-day correlation between the VSTOXX and EURO STOXX 50 (price return in EUR).

Source: STOXX.

We can dive deeper into this data and observe the correlation during periods that were significant for the market. For instance, the 30-day correlation reached –0.95 at the beginning of the COVID-19 pandemic, with the same level also being seen during the summer of the GFC. On the other hand, we also saw these levels in the summer of 2021. Spurred on by (possibly premature) post-COVID-19 optimism, a strong equity rally coincided with lower volatility, highlighting the fact that the correlation relationship is not just dominated by periods of market distress.

There are potentially two explanations for this negative correlation, with the first element being known as volatility feedback. If volatility is persistent and is priced, then an unexpected increase in it raises the expected future volatility and hence the required return on stocks. The result is an immediate negative impact on the current stock price (Kim et al, 2007). The second potential explanation is based on the "leverage effect". A decline in the price of a company's shares (and hence in its overall equity value) increases the company's balance sheet leverage. This makes its equity riskier and consequently increases the volatility of its share price (Schwert, 1989). These relationships between volatility and stock returns mean volatility products are now a commonly used way of hedging equity investments.

5.5 Inherent behavior of the VSTOXX

It is generally accepted that volatility has certain behavioral characteristics. This section explains these in greater detail and examines whether they can be seen when looking at the VSTOXX.

5.5.1 Volatility and stock returns

The following section separates the VSTOXX into volatility regimes. By dividing the data into terciles, we can assign VSTOXX levels to low, mid- and high volatility regimes. Figure 22 shows this breakdown.

Regime	Percentile	VSTOXX levels
Low	< 33.34	< 18.5
Mid	33.34 - 66.66	18.5 – 25.0
High	> 66.67	> 25.0

Figure 22: VSTOXX regimes.

Source: STOXX.

These distinctions then allow us to compare behavior at differing expectation levels for market volatility.

Figure 23 illustrates the price changes in the EURO STOXX 50 over a 30-day period for different VSTOXX regimes. The VSTOXX regime was established and the change in the price of the EURO STOXX 50 for the following 30-day period was calculated for each day (rounded to the nearest 2% interval, from –20% to 20%). Viewed from this perspective, VSTOXX levels are indeed a good estimate of market volatility and subsequent price changes (regardless of direction) in the EURO STOXX 50. Low regime VSTOXX levels corresponded to the smallest price change in the underlying index, while the fat tails of the high regime VSTOXX showed increased divergence. A slight upside bias can also be seen from the positive skew of all distributions.





Source: STOXX.

5.5.2 Mean reversion

Volatility is typically thought to be mean-reverting, and we examined whether the VSTOXX index levels support this claim. Mean reversion is a statistical characteristic in which a series of values is more likely to move towards its longer-term mean than away from it. Volatility is often observed moving between temporary or short-term highs and lows before reverting to a mean value, especially when viewed over longer time periods (Fouque et al, 2000).

Figure 24 shows the evolution of low regime VSTOXX levels over various time periods. We see a decreasing proportion of low regime volatility, and subsequent increases in mid- and high regime volatility.





Source: STOXX.

The inverse can be seen in Figure 25, with high regime levels reverting to low and mid-regime levels over time. This behavior aligns with mean reversion phenomena and supports previous theoretical thinking.



Figure 25: Evolution of VSTOXX levels from a high volatility regime.

Source: STOXX.

6. Use cases

As previously mentioned, volatility has become accepted as a separate asset class, and the VSTOXX index methodology enables it to be used in volatility-focused investment strategies. Our final section offers examples of how the VSTOXX is used in trading strategies and within indices.

6.1 Trading strategies

Earlier in this paper, we saw that the construction of the VSTOXX index does not allow for practical continuous replication, since a broad array of options are selected and reweighted throughout the trading day. Because of this, there is no concept of a "cash" or "spot" equivalent of a VSTOXX replicator. However, there is a rich ecosystem of tradable products in the derivatives market in the form of futures and options that are listed on Eurex.

In terms of mechanics, these derivatives have a slightly different expiry schedule than conventional index options. Since the VSTOXX index measures volatility over a 30-day window, VSTOXX futures expiry is set to 30 days from the standard third Friday option expiry. This means that the final hedge is a portfolio of options with a single expiry date. The ability to more easily hedge the futures enhances the liquidity of the derivatives market.

The figure below shows sample use cases for these instruments, which provide exposure to future volatility expectations.

Market gauge	Volatility indices are often referred to as "fear gauges", with higher levels pointing to greater uncertainty (which some observers interpret as fear). Note, however, that a high level on the VSTOXX does not forecast performance, but rather simply points to a broader range of possible outcomes.
Directional views	Using a single option to express a view on volatility exposes the position to many potentially unwanted risks. An instrument that tracks the VSTOXX allows for purer exposures to movements in volatility.
Volatility selling	Being short volatility can result in exposure to the volatility risk premium (i.e., the tendency for implied volatility to exceed realized volatility). As with any carry trade, this is not a risk-free strategy, since the short seller is exposed to volatility shocks.
Dispersion trading	The variance of a portfolio depends on the variance of its components and their correlation to one another. Trading the volatility of an index versus the volatility of its components allows the trader to express a view on correlation.
Spread trading	Relationships between regions, asset classes and future expectations evolve over time, and can lead to trading opportunities. In addition, event risk can arise; this may be most relevant to the EURO STOXX 50 (e.g., European elections). The resulting opportunities can take numerous forms, including: • Regions, e.g., US volatility versus European volatility • Term structure, e.g., long-dated versus short-dated volatility • Asset class, e.g., equities versus credit volatility
Hedging	As mentioned earlier, volatility is typically negatively correlated with underlying prices, and so can provide a valuable diversifier. This is not a free lunch, however: there is a cost to being long volatility because you are paying the volatility risk premium.
Trading trigger	Volatility tends to cluster in certain regimes. Since the performance of trading strategies can vary based on the market regime, the VSTOXX can be used to specify the regime/ environment within systematic strategies.
Risk control	Indices which target specific volatility levels are popular underlyings for structured products, as they can give more certainty to issuers who write options on these indices. Using the VSTOXX as an input to the risk control mechanism can incorporate the market's view of expected volatility for the EURO STOXX 50.

Figure 26: Use cases.

6.2 Additional STOXX volatility indices

There are a number of STOXX volatility indices beyond the VSTOXX, which are used to support some of these investment strategies:

Figure 27: Live STOXX volatility indices.

The EURO STOXX 50 Volatility (VSTOXX) Short-Term Futures replicates the performance of a long position in constant-maturity one-month forward, one-month implied volatilities on the underlying EURO STOXX 50.
The VSTOXX Short-Term Futures Investable is similar to the Short-Term Futures, but takes into account the bid-ask spread in the roll procedure.
The VSTOXX Short-Term Futures Inverse Investable is similar to the VSTOXX Short-Term Futures Investable, but replicates a short position.
The EURO STOXX 50 Volatility Mid-Term Futures replicates a constant five-month forward, one-month implied volatility exposure.
The EURO STOXX 50 Investable Volatility is designed as a rolling index targeting a constant three-month (90-day) forward, three-month maturity volatility exposure.
The EURO STOXX 50 Volatility-Balanced couples a base investment in the EURO STOXX 50 Net Return index with a dynamic allocation to index volatility via the VSTOXX Short-Term Futures.

Source: STOXX.

More information is provided in the STOXX Strategy Index Guide.

7. Conclusion

The VSTOXX measures market expectation of the volatility of the EURO STOXX 50 index over the next 30 days. In this paper, we examined the VSTOXX from start (theory) to finish (use cases), explaining the methodology and behavior in between.

From a theoretical perspective, volatility is a measure of uncertainty. The higher the volatility, the less certainty there is about the range of performance outcomes. Volatility can be measured in several ways, and the VSTOXX measures implied volatility, which is derived from option prices. However, options are exposed to a range of factors that influence their price. By creating a portfolio of options that is weighted by the inverse of their strike price squared, one can capture implied volatility without the sensitivities that individual options are subject to.

Launched in 2005, the VSTOXX uses this intuition to measure the volatility of the EURO STOXX 50 index, which serves as a barometer for Europe. The VSTOXX is fully rules-based, and at each calculation interval recreates a portfolio of OTM options that represents the implied variance of the EURO STOXX 50. The square root of this is the implied volatility. The VSTOXX also uses available pricing data to screen for option prices at each calculation interval.

A look at the historical performance of the VSTOXX allows us to validate some theories about the performance of volatility itself. These are that it tends to cluster in regimes, based on the market environment, and that it tends to mean-revert, meaning that extremely high or low levels tend to be followed by more moderate levels. Volatility also exhibits a strong negative correlation to index prices: In other words, large drops in the index tend to correspond to periods of high volatility, while steadily rising markets tend to correspond to periods of low volatility.

Finally, volatility became established as a separate asset class decades ago. The VSTOXX index and the ecosystem of derivatives based on it enable market participants to express views on European volatility, establish carry trades, hedge positions and determine which market regimes might offer optimal positioning. In this way, the VSTOXX enriches the investor toolkit – one of STOXX's core missions.

8. References

Allen, P., Einchcomb, S. and Granger, N. (2006) "Variance swaps," *European Equity Derivatives Research,* J.P. Morgan Securities Ltd., November 17, 2006.

Bae, J., Kim, C.J., Nelson, C. (2007) "Why are stock returns and volatility negatively correlated?", *Journal of Empirical Finance 2007*, pp. 41–58.

Black, F. and Scholes, M. (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, pp. 637–659.

Demeterfi, K., Derman, E., Kamal, M. and Zou, J. (1999) "A Guide to Volatility and Variance Swaps," *Journal of Derivatives*, pp. 6, 9–32; "More Than You Ever Wanted to Know About Volatility Swaps," *Goldman Sachs Quantitative Strategies Research Notes*, March 8, 1999.

Fassas, A., Siriopoulous, C. (2021) "Implied Volatility Indices – A Review," *The Quarterly Review of Economics and Finance*, 2021.

Fouque, J.P., Papanicolaou, G. (2000) "Mean-Reverting Stochastic Volatility," *International Journal of Theoretical and Applied Finance*, Vol. 03, No. 01, pp. 101–142, 2000.

Nielsen, L. (1993) "Understanding N (d1) and N (d2): Risk-Adjusted Probabilities in the Black-Scholes Model," *Finance*, 14, pp. 95–106.

Rouah, F. "Variance Swap," <u>https://www.frouah.com/finance%20notes/Variance%20Swap.pdf</u>.

Schwert, G.W. (1989) "Why Does Stock Market Volatility Change over Time?", *Journal of Finance*, 44, pp. 1115–1153.

STOXX. EURO STOXX 50 Index, https://stoxx.com/index/SX5E/.

STOXX. (2024) "STOXX Strategy Index Guide," <u>https://www.stoxx.com/document/Indices/Common/</u> Indexguide/stoxx_strategy_guide.pdf, June 2024.

9. Appendix

Interpolated risk-free rate calculation

The interpolated risk-free rate calculation uses the following rates:

Figure 28: Rates used for interpolation.

Rate	Period	ISIN
€STR	1 day	EU000A2X2A25
EURIBOR 1 month	1 month	EU0009659937
EURIBOR 3 months	3 months	EU0009652783
EURIBOR 6 months	6 months	EU0009652791
EURIBOR 12 months	12 months	EU0009652809
REX 2-years (price index)	2 years	DE0008469149

Source: STOXX.

The interpolated risk-free rate for the ith option expiry is given by:

$$r_{i} = \frac{T_{lt} - T_{tm}}{T_{lt} - T_{st}} \times r_{st} + \frac{T_{tm} - T_{st}}{T_{lt} - T_{st}} \times r_{lt}$$

where:

 T_{tm} = the number of seconds to the expiry of the i^{th} option.

 T_{st} = the number of seconds in the period immediately lower than T_{tm} .

 T_{lt} = the number of seconds in the period immediately greater than T_{tm} .

 r_{st} = the rate used for interpolation in which the period is immediately lower than T_{tm} .

 r_{lt} = the rate used for interpolation in which the period is immediately greater than T_{tm} .

10. Offices and contacts

Learn more about STOXX & DAX Indices on STOXX.com

Zug

Theilerstrasse 1A 6300 Zug Switzerland P|41 43 430 71 60

London

8 Old Jewry 4th Floor London EC2R 8DN United Kingdom P|+44 20 7862 7680

Frankfurt

Taunus Tower Mergenthalerallee 73–75 65760 Eschborn Germany P|+49 69 2 11 0

Paris

5 Rue du Renard 75004 Paris France P|+33155 27 38 38

Prague

Futurama Business Park Building E Sokolovska 662/136e 186 00 Prague 8 Czech Republic P|+420 228 889 234 New York

1177 Avenue of the Americas 14th Floor New York, NY 10036 USA P|+1 646 680 6350

Hong Kong

Unit 901, 9/F 100 Queen's Road Central Hong Kong P|+852 3107 8030

Singapore

80 Robinson Road, #02-00 Singapore 068898 Singapore P|+852 3107 8030

Sydney

68 Pitt Street Level 17 Sydney, NSW 2000 Australia P|+61 2 8074 3104

Call a STOXX representative

Customer support customersupport@stoxx.com P|+41 43 430 72 72



Part of





DEUTSCHE BÖRSE GROUP

STOXX Ltd. (STOXX) and ISS STOXX Index GmbH (together "STOXX") research reports are for informational purposes only and do not constitute investment advice or an offer to sell or the solicitation of an offer to buy any security of any entity in any jurisdiction. Although the information herein is believed to be reliable and has been obtained from sources believed to be reliable, we make no representation or warranty, expressed or implied, with respect to the fairness, correctness, accuracy, reasonableness or completeness of such information. No guarantee is made that the information in this report is accurate or complete, and no warranties are made with regard to the results to be obtained from its use. STOXX will not be liable for any loss or damage resulting from information obtained from this report. Furthermore, past performance is not necessarily indicative of future results. Exposure to an asset class, a sector, a geography or a strategy represented by an index can be achieved either through a replication of the list of constituents and their respective weightings or through investable instruments based on that index. STOXX does not sponsor, endorse, sell, promote or manage any investment product that seeks to provide an investment return based on the performance of any index. STOXX makes no assurance that investment products based on any STOXX® or DAX® index will accurately track the performance of the index itself or return positive performance. The views and opinions expressed in this research report are those of the author and do not necessarily represent the views of STOXX. This report may not be reproduced or transmitted in whole or in part by any means – electronic, mechanical, photocopying or otherwise – without STOXX's prior written approval.