

MIMS – Measuring and Exploiting Volatility Risk Premium

Markets and Alternatives Research Team

Report – December 2022

As hinted by the Macro Research Team in the report concerning the effectiveness of sanctions against Russia, equity implied volatility increased (yet for just a very short time horizon) dramatically following the Russian invasion of Ukraine.

With the aim of understanding more in-depth the topic of volatility investing, this report delves into a set of possible strategies to exploit volatility in the market. In order to provide the reader with a precise view of the topic under discussion, we first introduce the notion of volatility risk premium (VRP, hereafter). This concept refers to the phenomenon that option-implied volatility tends to exceed realized volatility of the same underlying asset, thereby generating a profit opportunity for volatility sellers. Then, we show how this relationship has held almost systematically over the past 20 years and introduce a set of strategies that can be put in place to exploit this phenomenon. Finally, we analyze the result we would have obtained by implementing the simplest possible strategy to harvest VRP, i.e., put writing (that is, shorting a put option), focusing both on the US and European equity markets.

INTRODUCTION

Markowitz and the birth of portfolio theory

In recent years, volatility has evolved from an academic idea into a risk management tool. Indeed, it has nowadays become something investors can trade, just like a stock or a bond.

However, the notion of volatility was firstly introduced as a proxy for the dispersion of portfolio returns (Markowitz, 1952). In 1952, Markowitz published a groundbreaking paper titled "Portfolio Selection". In that circumstance, he argued that fund performance should be judged compared to the amount of risk it bears (and not just on the performance itself).

According to Markowitz's theory, the ideal combination of weights is that which allows the investor to maximize the level of expected excess return, given a certain level of volatility; or, equivalently, to minimize the level of volatility, given a target expected excess return. As a result, according to this view, volatility is a handy proxy for portfolio risk, rather than a fully-fledged asset class.



Head of Research Area

Emanuele Licari: +39 346 956 1880

Head of Markets & Alternatives Research

Antonio Messina: +39 348 382 3037

Markets and Alternatives Research Analysts

Matteo Negrini: +39 334 795 6252

Bartu Sen: +39 331 278 1412

Edoardo Calleri di Sala: +39 345 995 7753

Angelo Perri: +39 389 101 6478


Deniz Aktas: +39 347 784 4647

Black-Scholes model

For insurance and speculation purposes, options trading was first introduced in 1973, with a model to efficiently calculate the price of options, partly based on their volatility. This became popularly known as the Black - Scholes Model (B&S hereafter).

Nonetheless, this was not quite enough to truly turn volatility into an asset class like equities or bonds. Indeed, as a consequence of the hypotheses underlying the B&S model, there should be a direct relationship between the level of implied volatility and the corresponding price of the option. More precisely, according to the B&S formula, all the options on the same underlying should be priced with the same implied volatility, even though they have different strikes. Hence, if the B&S model was perfectly correct, we should observe a constant level of implied volatility across strikes, thus obtaining a flat volatility surface.

However, empirical evidence on the S&P 500 shows that implied volatility is, in fact, not constant across strike prices, as shown by the graph below:



Source: *Implied volatility and skew as an indicator of market direction*, Agarwal, 2021

This graph shows that the distribution of stock returns is slightly skewed relative to that suggested by the B&S model. This argument clarifies what assumption causes the model not to perfectly hold in practice, providing at the same time a great intuition behind option pricing models. In addition, further contributions (not pivotal for the sake of this report) address the B&S model's main flaw, namely that volatility is (wrongly) supposed to be constant. On the contrary, it suffers from mean reversion, as reflected e.g., by the Heston model (Heston, 1993).

How do we measure volatility?

Before delving into the details of VRP, the concepts of realized and implied volatilities must be introduced.

On the one hand, realized volatility (RV) is defined as the actual movement that occurs in a given underlying over a defined past period.

On the other hand, implied volatility (IV) is the level of volatility that, inserted in the B&S formula, provides the analyst with the observed market price of the option.

As for RV, we need to bear in mind that volatility cannot be observed. Therefore, there will always be a measurement error. In order to measure volatility in such a way as to minimize the measurement error, the realized volatility estimator has been developed.

Let $r(\tau, h) = \log(S_\tau) - \log(S_{\tau-h})$ be the return on the time interval $\tau, \tau + h$ of stock S_τ . The realized volatility estimator for the volatility of an asset price S_τ on a time interval $(t - x, t)$ at resolution h is given by:

$$RV\left(t, x; \frac{1}{h}\right) = \sqrt{\sum_{j=1}^{\lfloor x/h \rfloor} r(t-x+jh; h)^2}$$

RV is computed as the square root of the sum of the square returns over the sub-periods of length h .

Using high-frequency data, it is possible to take h to 0. For $h \rightarrow 0$, and under the assumption that price paths are continuous, the realized variance estimator converges to the integrated variance:

$$RV^2\left(t, x; \frac{1}{h}\right) \rightarrow \text{Int}V(t, x) \text{ as } h \rightarrow 0$$

These two measures are going to be of pivotal importance when we are going to discuss the difference between realized and expected volatility in a systematic manner.

In order to properly introduce the IV, we first present the B&S formula for pricing a European call option on a stock (whose current price is S_t) that does not pay dividends. Let π be the quoted price of such an option and $c^{B\&S}$ the B&S formula for the option price (with strike K and time to maturity $T - t$):

$$c^{B\&S} = S_t N(d_1) - Ke^{-r(T-t)} N(d_2)$$

where d_1 and d_2 :

$$d_1 = \frac{\ln\left(\frac{S_t}{x}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}$$

$$d_2 = \frac{\ln\left(\frac{S_t}{x}\right) + \left(r - \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}$$


The previous formulation can be summarized via the following mathematical identity:

$$c^{B\&S}(S_t, x, T-t, r, \sigma) = \pi$$

Now, by taking the current price of the option as given, we can solve the reverse problem with the aim of computing the level of volatility implied in the actual market price. In other words, IV is going to be the level of volatility that, plugged into the B&S formula, makes the price estimated adopting B&S formula equal to the observed market price.

As already discussed, there is a misalignment with respect to the B&S environment as far as the relationship between implied volatility and strike price is concerned. Further, if we account for time-to-maturity, we see how differences are even more pronounced. By looking at the implied volatility surface (plotted in the next page), it is apparent that to different levels of strike price and time to maturity correspond different levels of implied volatility. Recalling the fact that both call and put prices are increasing in volatility, if we believe the B&S model to hold a trivial arbitrage opportunities would be:

- Selling Out of the Money (OTM) puts as they are overpriced for very low strikes;
- Buying In the Money (ITM) puts as they are underpriced for very high strikes.



Source: *Dynamics of Implied Volatility Surfaces*, Cont & DeFonseca, 2002

Therefore, the price of OTM puts is generally higher than that of ATM or ITM calls, since they incorporate an intrinsic “insurance value”. Notice that according to the convention in the graph above, moneyness is proxied by the ratio $\left(\frac{K}{S_t}\right)$. With volatility not yet considered as an asset class, in 1993, the idea of an option-based volatility index turned into reality with the CBOE Volatility Index, popularly called VIX. This index is derived from a calculation designed to produce a measure of the constant, 30-day expected volatility of the US stock market, as implied by option prices. As of today, it is one of the most recognized measures of volatility, widely reported as a daily market indicator.

To sum up, we can therefore say that the realized variance (or quadratic variation) cannot be measured until the uncertainty has unfolded (i.e., it can only be measured ex-post). The ex-ante expectation of this quantity over the next 30 days is instead captured by the VIX.

Introduction to Volatility Risk Premium

Once the components of the topic have been introduced, it is possible to define the notion of volatility risk premium. The VRP reflects the compensation investors earn for providing insurance against market losses. In other words, it is the reward for bearing an asset’s downside risk.

More technically, VRP refers to the spread between an option’s expected volatility (usually the VIX) and the underlying asset’s subsequent RV. On average, the expected value is well above the realized one, therefore implying a positive VRP (except for very special occurrences). In fact, if the market were on average correct, the estimate would be unbiased (i.e., the average of the error should be 0; or, equivalently, the expected value of the estimator should be the true parameter).

When constructing the VIX, it was mentioned that σ may depend on any information prior to T , on top of the final level of the underlying index. More specifically, assume to operate in a market where realized (local) volatility estimation works perfectly. In this setting, the level of volatility is a deterministic function of the underlying index, and it is possible to hedge any put and call option by constructing the so-called “local volatility tree”. Hence, after delta-hedging the underlying risk, there are no other sources of risk left unhedged, which is equivalent to saying that the position is riskless.

On the contrary, we can observe a bias in the error made by the market, which is systematically overestimating the level of expected volatility. In real markets, despite being protected from the risk associated with the underlying security (i.e., of being delta-hedged), investors charge an extra premium equal to the difference between the two curves below (which refer to the US equity market in 2022). In doing so, they seek protection against the so-called “stochastic volatility risk” (that is, the component of the risk-driving volatility fluctuations that is not accounted for by the conventional local volatility setting). In short, the level of volatility estimated ex-ante through the VIX is systematically higher than the RV, and such a difference is due to the additional premium asked by risk-averse individuals for selling protection.

To conclude, the whole comparison between the VIX and RV is inconsistent within the B&S framework, according to which volatility is constant over time, there is no volatility risk, and (as a result) there is no reason to ask for a volatility risk premium.



Source: *Historical Norm, VIX Index vs Future Realized Volatility (21 Trading Day, Annualized)*

Beware that a positive VRP is observed not only in stock markets but across different asset classes. For instance, investors can gain exposure to this premium to generate income from selling volatility also in metal or agricultural commodities markets.

Moreover, the VRP is present across volatility regimes. Even in a low-volatility environment, IV has been higher than RV, meaning that an option seller in such environments has been profitable on average. In conclusion, we notice that the rationale for the existence of the VRP, i.e., providing insurance against large market moves, prevails regardless of the current level of volatility.

On a minor note, behavioral insights may give an answer to VRP existence, as IV is higher than RV because investors tend to overestimate the probability of extreme market events.

According to a survey conducted by Yale University many retail and institutional investors believe that the probability of a “catastrophic stock market crash” in the next six months is greater than 10%. However, data show that their estimation is wrong: the likelihood that this kind of event happens is just 1%.


People overestimate unlikely risk events in financial markets because market movements are controlled by external forces. In fact, Investors do not control, at least not directly, government policies, central banks, wars, and natural disasters such as pandemics. People feel less risk when they control something, even though the risk could be higher.

Furthermore, investors are more motivated to avoid a loss rather than to gain a win. Consequently, they tend to hedge risk through insurance instruments, like options.

EMPIRICAL EVIDENCE

Volatility Risk Premium

The following graph provides an illustration of the evolution of the volatility risk premium for the S&P 500 over the last 20 years.



Sources: Bloomberg, MIMS estimate

The graph is the result of the difference between the VIX (i.e., expected 1-month daily volatility) and the subsequent 1-month daily realized volatility. RV has been calculated using the standard deviation of the closing prices, via the following formula (which is a simplified version of the one reported in the previous section of the report):


$$RV_t = \sqrt{\frac{252}{N} \sum_{i=0}^{N-1} \ln\left(\frac{P_{t-i}}{P_{t-i-1}}\right)^2} \times 100$$

where:

- RV_t : level of realized volatility on day t
- P_t : closing price on day t
- N : number of trading days in the lookback period

The VRP has been mostly positive, except for some exceptional cases (2008 and 2020). However, despite other drops below zero in the 2010s, the 6-month daily moving average has only turned negative in 2008 and 2020.


A similar pattern is observable in the case of the Euro Stoxx 50, analyzed in the following graph, where the same calculation method has been adopted.



Sources: Bloomberg, MIMS estimate

There are, however, some differences between Euro Stoxx 50 and S&P 500. Among these, as for the Euro Stoxx 50, there is a drop in VRP in 2010 (because of the effects of the 2008 financial crises in Europe) and a higher standard deviation of the VRP. Finally, we observe that the 6-month moving average did not drop below zero in 2008, contrarily to the S&P 500.

Relationship between expected and realized volatility



Source: Bloomberg

Volatility has been consistently overestimated, as previously demonstrated. However, by looking at the difference between expected and realized volatility for the last 12 months, we can get a grasp of the relationship between them. When the realized volatility exceeded expectations (i.e., in March and May), the market expectations for future volatility increased, even if the subsequent realized volatility did not resemble those expectations. As a result, the VRP, after turning negative, spiked again rapidly. We will also delve into this phenomenon later when analyzing VRP across asset classes.

Reasons behind sharp drops in VRP

As previously stated, the VRP consists of the difference between implied and realized volatility. A well-known empirical regularity is that volatility tends to be negatively correlated with current and past asset returns. In other words, volatility tends to be much higher when asset prices drop than when markets rally (Lombardi & Schimpf, 2014). During the 2008 and 2020 financial crises, market prices plummeted and realized volatility sharply increased, overcoming the previous estimates of the market, thus leading to a sudden drop in VRP.

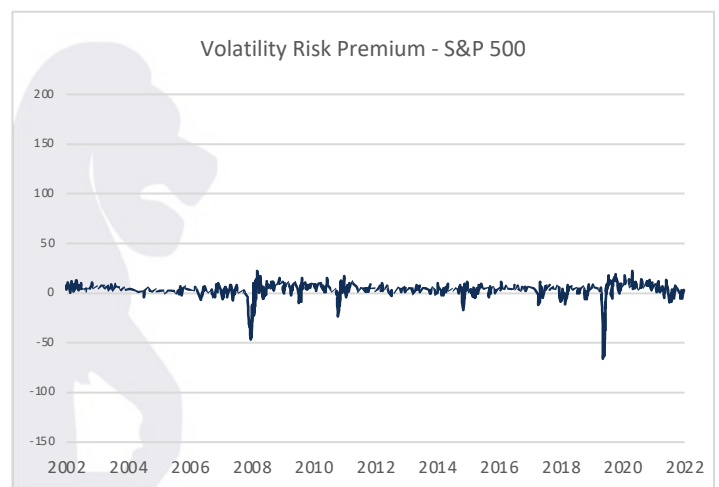
VRP across sectors

Empirical evidence shows that volatility is not the same across sectors, and this can be confirmed by observing the VRP trends across representative examples taken from different industries. We consider some examples of growth stocks (e.g., Apple and Amazon), for which the VRP is typically higher, and compare them with reflationary sectors such as financials (e.g., Goldman Sachs), for which volatility tends to be lower. Data on implied volatility have been extracted from the following indices: VXAPL, VXAZN, VXGS.


Data provide evidence in favor of the previous relationships, apart from 2020, as the financial industry was severely hit by the pandemic. In fact, the VRP for Goldman Sachs has been significantly lower over time, when compared with Apple and Amazon. Finally, we observe similar behaviors for Apple and Amazon, the latter showing a higher standard deviation of the VRP.

VRP across asset classes


The following graphs allow us to analyze the VRP across equity (i.e., S&P 500) and commodities (i.e., gold and crude oil). The data on implied volatility are extracted respectively from the following indices: VIX, GVZ and OVX; whereas RV are calculated using the previously described formula. The results confirm that the VRP for equity has been significantly higher than for gold, which tends to be less volatile over time, yet lower than crude oil, which shows overall the second highest VRP, after that of the Euro Stoxx 50. The reasons underpinning such values are related to the uncertainty about the future of crude oil markets and imbalances between supply and demand (which are key drivers for oil prices). We can again observe sudden spikes after a negative downturn in VRP.



Sources: Bloomberg, MIMS estimate



Sources: Bloomberg, MIMS estimate



Sources: Bloomberg, MIMS estimate

Finally, the table provides a comparison of VRP, implied volatility, and realized volatility among different assets. In order to avoid the negative impact of outliers on our analysis, we consider the median of those variables. An interesting aspect is the difference between the S&P 500 and the Euro Stoxx 50, the latter showing the highest median VRP, therefore providing the best environment for harvesting it.

	Median VRP	Median Implied volatility	Median Realized Volatility
S&P 500	4.025	15.52	11.852
Euro Stoxx 50	6.118	23.10	16.595
Apple Inc.	4.283	28.39	24.680
Amazon	4.317	30.55	25.458
Goldman Sachs	3.989	26.73	22.815
Crude Oil	4.684	34.36	27.999
Gold	2.553	15.94	13.290

Sources: Bloomberg, MIMS estimate

Put-Call Parity and VRP: the role of implied dividend yield

Put-Call Parity (PCP hereafter) defines the relationship that should hold between the value of a position that is long call and short put and the value of a position that is long the stock and borrows an amount of cash equal to the strike price. The general PCP rule does not apply to American options as they can be exercised before expiration. As such, PCP is the most prominent example of no-arbitrage (NA) condition.

In this part of the report, we deal with implied dividend yield and its role in detecting potential arbitrage opportunities arising from systematic mispricing of (typically) put options.

First, we provide the basic formulation of PCP:

$$c_T - p_T = S_T - K$$

where:

- c_T is the price of the call option;
- p_T is the price of the put option;
- S_T is the current price of the stock underlying such options;
- K is the strike price of such options.

Furthermore, the law of one price suggests that if PCP holds at expiration date T , it should hold at any t prior to expiration. If this was not the case, we could buy the cheaper option and sell the more expensive one. Therefore, we can rewrite PCP as follows:

$$c_t - p_t = S_t - Ke^{-r(T-t)}$$

Then, we can transcribe PCP in terms of future prices, simply by performing the following substitution:

$$F_t^T = S_t e^{r(T-t)}$$

Then, PCP can be rewritten as:

$$p_t - c_t = e^{-r(T-t)}(K - F_t^T)$$

All three contracts involved should be purchased on the same given date; otherwise, we cannot guarantee that prices are in equilibrium.

As we accept that PCP holds in generic time t , we also accept that we can be dealing with future price F_t^T as well as spot price S_t , since we can be interested in exploiting arbitrage in the future. Holding a forward (or an option) contract, nevertheless, does not entitle the owner of dividends, which, therefore, should be deducted from the future price as follows:

$$F_t^T = S_t e^{(r-q)(T-t)}$$

However, an investor cannot know the dividend q with certainty as q depends on multiple company-specific, structural, and macroeconomic reasons. This creates so-called dividend risk. At this point, PCP rule assists investors in estimating possible dividend yields in the future as:

$$p_t - c_t = e^{-r(T-t)}(K - F_t^T)$$

holds true. Notice that, in this circumstance, F_t^T incorporates the impact of dividends.

Since, dividend yield q , future price, and $e^{r(T-t)}$ do not depend on the strike price, it is perfectly feasible to find implied dividend yield curve from quoted option bid/ask price and quoted future price. Therefore, implied dividend yield $q(K)$ at strike price K is:

$$q(K) = -\frac{1}{T-t} \log \frac{(Ke^{-r(T-t)} - (p_t - c_t))}{S_t}$$

Clearly, the lower this value, the higher is the willingness of the option buyer to hedge against very unlikely market crashes (e.g., in the case of deep OTM put options). No-arbitrage conditions occur when $q(K)$ is independent of K . This is reasonable as the dividend is related to the company's earnings rather than the strike price of the option under assessment.

Lastly, we explain what would have happened if arbitrage opportunities existed in the market. If PCP holds at the expiry date, T , but not in any generic date t , then we would encounter:

$$p_t - c_t \neq e^{-r(T-t)}(K - F_t^T)$$

It is, then, very feasible and convenient for a trader to simply buy the cheaper options strategy and sell the more expensive one.

Empirical test on the implied dividend yield

After giving a brief introduction on PCP and NA and their importance for computing implied dividend, we empirically test whether we can rely on PCP as a tool to assess NA in the market. Firstly, we use the S&P 500 Options Index for options expiring in 3 and 5 months and for the Euro Stoxx 50 Options Index expiring in 3 and 5 months. We assume the price of a put is the simple average of the bid and ask price of the underlying asset class. The same simplification is adopted for call options.

$T - t$ is 0.25 for 3-month options and 0.35 for 5-month options as we assume that t is as of the day data have been extracted. To clarify, as our time frame comprises daily data extracted (t) until the expiry date (T), it can be said that $T > t$. Lastly, 3-month S&P 500 options are multiplied by 10 to be consistent with its 5-month options and for better visualization purposes.

Both tables include implied dividend yields. As for the Euro Stoxx 50, the current future price as of the day data are extracted was between 3450-3475 and for S&P 500 that level was between 3830-3835.

S&P 500 and Euro Stoxx 50 tables show a very interesting but expected result about the implied dividend yield: around ATM level it starts stabilizing. As regards S&P 500, 3-month options, that level is around -0.12%, while for 5-month options it is around -1.5%.

Moreover, as for Euro Stoxx 50, we notice a stabilization of around -0.35% for 3-month options and -0.39% for 5-month options. A similar feature is also observed for extremely high strike prices (deep ITM puts) relative to current future prices. Indeed, both indices' implied dividend yields fall into a certain range. It is easier to comprehend this from graphs. For example, notice that 3-month S&P 500 options' implied dividend yields are restricted between -1.8% and -2%. A similar pattern is observable in the case of the Euro Stoxx 50, as dividend yields are constrained between -0.38% and -0.40%.

Similarly, we also expected to encounter this situation, as we already established that dividend yields and strike prices are independent of each other: the strike price is an intrinsic characteristic of an option contract, specifying its class, while dividend yields should be more connected to profitability of companies listed in the indices and, more specifically, what percentage of the net income the board of directors approves to pay to shareholders. Overall, stabilized implied dividend yields in both scenarios prove that strike prices and dividend yields are independent of each other.

By looking, again, at the Euro Stoxx 50 options, we realize that extremely low strike prices relative to current future prices have considerably lower negative implied dividend yield percentages than those of higher strike prices.

The latter is a simple explanation of the so-called "OTM Put Puzzle", which occurs when OTM Puts are priced by investors way higher than their intrinsic value, computed coherently with Put-Call Parity.

Historically, the consequences of such evidence became apparent in the aftermath of circumstances such as the 1987 market crash. Prior to that event, deep OTM put options were priced according to the Black-Scholes formula, which (given the assumption that returns are distributed according to the gaussian distribution) did not incorporate the insurance premium previously mentioned. As a result, as the market crashed, buyers of OTM put options exercised them, causing severe losses to option sellers.

Going back to the analysis of implied dividend yield, we can argue that those who invested in option contracts in Euro Stoxx 50 still believe that an imminent extreme market crash is unlikely. Hence, deep OTM puts are still overpriced. The first implication of such a result is the close relationship between the option price and the possibility of it being executed in the future. Intuitively, a lower likelihood of being executed yields a cheaper price of the option. Now, let us briefly comment on the results.

On the one hand, as for S&P500, 5-month options have the lowest implied dividend yield among others. This entails that buyers of 5-month S&P 500 options are very concerned with the current economic situation within a 5-month time horizon relative to that within a 3-month time horizon. They are, therefore, willing to pay much more for such options, that is, they are willing to accept a much more negative implied dividend yield in order to be insured against the risk of a market crash.

As regards 3-month options, we observe an implied dividend yield ranging from approximately -0.13% to -0.20%; whereas, for 5-month options, we observe implied dividend yields ranging from approximately -1.57% to -1.61%.

On the other hand, as for Euro Stoxx 50 options, 3-month options are slightly cheaper than 5-month options across the whole set of strikes. Looking at the Euro Stoxx 50, for 3-month options, the implied dividend yield is -0.40% for the lowest strike price and -0.32% for the highest strike price. As for the 5-month options, we observe an implied dividend yield ranging from approximately -0.38% to -0.41%.

Finally, since both indices' options have negative implied dividend yields, option sellers require an additional risk premium to sell options, in addition to the price they should ask according to PCP. To be even more precise, the rationale behind negative dividend yield and option being over-priced is that since option holders are not entitled to receive dividends, we need to deduct them. As dividends decrease with negative dividend yields, the intrinsic value of stocks decreases as well; however, this, in our case, is not reflected in options prices because negative dividend yields indeed lead to higher (over-priced) option prices, which should be the other way around.

Implied Dividend Yields for S&P 500 Options		
Strike Price	3 Months S&P 500	5 Months S&P 500
3785	-0.1290%	-1.5706%
3790	-0.1238%	-1.5655%
3795	-0.1186%	-1.5677%
3800	-0.1343%	-1.5737%
3805	-0.1239%	-1.5797%
3810	-0.1291%	-1.5745%
3815	-0.1292%	-1.5731%
3820	-0.1448%	-1.5828%
3825	-0.1292%	-1.5776%
3830	-0.1240%	-1.5910%
3835	-0.1241%	-1.5896%
3840	-0.1866%	-1.5844%
3845	-0.2022%	-1.5904%
3850	-0.1918%	-1.5889%
3855	-0.1918%	-1.5949%
3860	-0.1971%	-1.5897%
3865	-0.1919%	-1.5920%
3870	-0.1867%	-1.5943%
3875	-0.1919%	-1.6040%
3880	-0.1972%	-1.6062%

Sources: Bloomberg, MIMS estimate

Implied Dividend Yields for Euro Stoxx 50 Options		
Strike Price	3 Months Euro Stoxx 50	5 Months Euro Stoxx 50
3225	-0.4023%	-0.4129%
3250	-0.3937%	-0.4026%
3275	-0.3735%	-0.4048%
3300	-0.3649%	-0.4069%
3325	-0.3736%	-0.3842%
3350	-0.3650%	-0.3864%
3375	-0.3679%	-0.3802%
3400	-0.3708%	-0.3906%
3425	-0.3564%	-0.3886%
3450	-0.3651%	-0.3989%
3475	-0.3565%	-0.3928%
3500	-0.3536%	-0.3990%
3525	-0.3450%	-0.4011%
3550	-0.3422%	-0.3868%
3575	-0.3278%	-0.3930%
3600	-0.3249%	-0.3951%
3625	-0.3336%	-0.4013%
3650	-0.3250%	-0.4034%
3675	-0.3106%	-0.3973%
3700	-0.3251%	-0.3871%

Sources: Bloomberg, MIMS estimate

To exploit opportunities arising from the studied phenomenon, market operators construct different strategies by using derivatives to short volatility. In the following sections of the report, we are going to present the three main strategies adopted in the field to harvest VRP. First, we are only going to introduce and make reference to the main advantages and drawbacks of two pivotal strategies in the context of volatility trading: Variance swaps and Short straddle. The calibration of such strategies is, however, left to further research. Finally, we are going to systematically test the actual performance and diversification role of one of them (i.e., Put writing).

Shorting variance swaps strategy

A variance swap is a forward contract on future realized price variance. At expiry, the receiver of the “floating leg” pays (or owes) the difference between the realized variance (or volatility) and the agreed-upon strike price.

At inception, the strike K is chosen so that the fair value of the swap is zero (i.e., it is based on expected volatility). As long as realized variance is lower than the expected one, the floating leg receiver would be able to profit from this strategy.

While the main advantage of trading volatility swaps is the one-to-one exposure to the VRP, the other side of the coin is given by its exposure to volatility spikes and the low transparency, as they are traded OTC.

Short straddle strategy

A short straddle is a widely used strategy to harvest VRP. The strategy consists of periodically selling an ATM call option and an ATM put option with the same exercise date, allowing the investor to get exposure to the volatility of markets. Unlike buying a straddle, selling it lets the investors position themselves according to the likely possibility of low realized volatility compared to implied volatility.

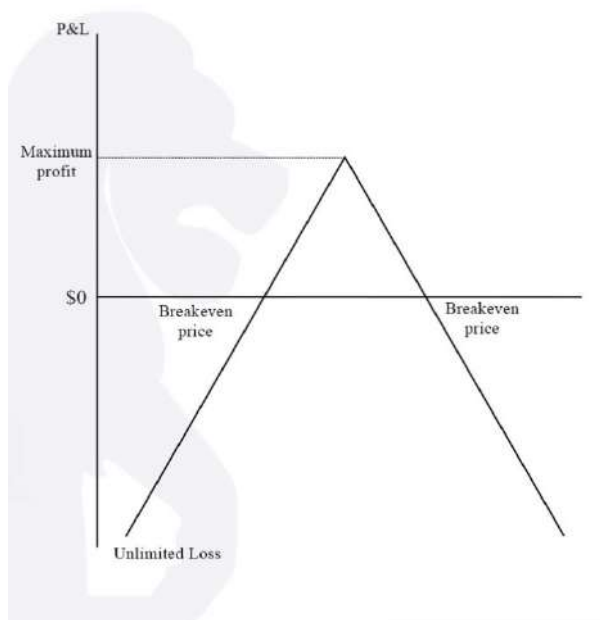
This is why, short straddles are listed under the strategies used to "sell volatility" (Eraker, 2009). Selling options grants the investor a premium, which the option buyer pays. Investing periodically, the straddle seller gains a premium and profits if the underlying asset's price traces a path within a narrow channel determined by breakeven prices.

The potential profit of short straddles is limited to the premium received. However, it is essential to note that if the underlying asset price increases above the breakeven price, the loss might be unlimited, as a stock price may theoretically increase infinitely.

If the price of the stock drops, the loss is substantial yet limited, as the underlying asset cannot drop below zero (as shown in the graph below).

Another risk of selling straddles is an early exercise, which applies to American-style options. As these options can be exercised any business day before expiration, the straddle seller might be obliged to buy the underlying asset at a higher price unexpectedly.

Payoff Diagram of a Short Straddle



It is now worth commenting on the greeks of such exotic strategy.

- Delta Exposure (i.e., the impact on the market value of the option of a price change in the underlying): short straddle strategy has a near-zero delta value, though its sign is unclear and depends on gamma. This implies that the position is not sensitive to minor price changes in the underlying asset.
- Gamma Exposure (i.e., delta sensitivity): as the price drifts from the strike price, a short straddle's delta shifts towards -1, which means the gamma is negative. This infers that a change that exceeds one of the breakeven prices will cause the position to lose money.
- Theta Exposure (i.e., time decay): the option loses value as it gets closer to the exercise date. Straddle sellers benefit from a positive theta value and are exposed to double the theta value of a naked option since they sold two options.

- Vega Exposure (i.e., sensitivity to volatility): the most paramount greek for short straddles used in VRP harvesting is vega, which is negative for the strategy. An increase in the implied volatility contradicts the straddle seller's prediction that the market volatility will not increase. This causes the option prices to go up, affecting the strategy negatively. A straddle seller is exposed to volatility twice as they sell two options.

Short straddles can be used to diversify an options portfolio when an investor recognizes the risk aversion of the market participants. This is usually done in between important news, including FED meetings and company announcements. They can be utilized to manage portfolio greeks such as a decreasing agent for the time decay effect of the overall portfolio or to harvest Volatility Risk Premium without affecting the delta of the portfolio.

Put writing strategy

The put writing strategy requires the market operator to sell OTM put options, based on the hypothesis that the realized volatility will not move the actual price of the underlying asset above the strike price, and the counterparty is willing to pay a premium for the perceived volatility, typically greater than the ex-post one. While profits are limited to the premium received, losses can be substantial in case the hypothesis is proved to be invalid a posteriori.


Now, the comparison is conducted as follows:

First, we present the CBOE S&P 500 PutWrite Index, which tracks the value of a hypothetical portfolio of securities (i.e., the PUT portfolio) that yields a buffered exposure to S&P 500 stock returns. More precisely, the PUT portfolio is composed of 1- and 3- T-bills (in order for that to be delta-hedged) and of a short position in put options on the S&P 500 index (SPX puts). Then, we compare the performance of such an index with that of the S&P500 itself.

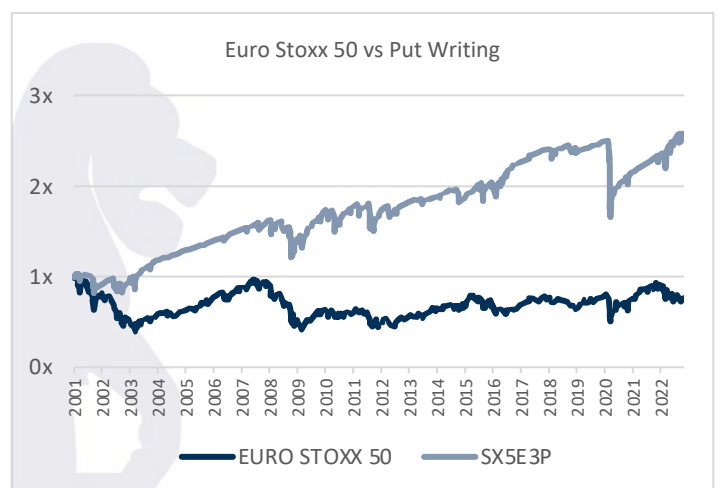
Second, we introduce the Euro Stoxx 50 PutWrite Index (SX5E3P), which aims to replicate the performance of a collateralized put option strategy.

This last index is based on a quarterly scheme with monthly put option tranches where: i) the investment notional is invested into the 3-month Euribor market; ii) Monthly put options are written in three tranches; iii) Intra-quarter put options are cash settled by borrowing in the one-month Euribor market (if necessary).

The following graphs study the evolution of the previous strategies, assessing by what multiplicative factor the price index of both strategies has evolved over the past 22 years relative to the benchmark. The systematic put strategy has consistently overperformed the Euro Stoxx 50 over the past 20 years, whereas the comparison with the US leads to a more conservative result, yet in line with the performance of the benchmark, even after discounting for crises of 2008 and 2020.



Sources: Bloomberg, MIMS estimate



Sources: Bloomberg, MIMS estimate

By analyzing the graphs, it is apparent that, at least in the long run, the put writing strategy gains from periods when equity markets are stable. Such phases occur with a higher probability relative to that assumed by put sellers when pricing OTM put options. This is what rationalizes the overall positive performance of these two indices relative to their benchmarks.


Then, we compare the put writing strategy with other long-short portfolios on the basis of the suggestion of the Carhart four-factor model (Carhart, 1997).

The Carhart model is an expansion of the Fama-French three-factor model (Fama and French, 1993), that, on top of the index proxying the market portfolio, incorporates three additional factors depicting three spread positions:

- SMB (Small minus big market capitalization): small companies tend to outperform big companies a posteriori;
- HML (High book-to-market minus low): stocks with high book-to-market (i.e., firms that are currently underpriced, also called value stocks) tend to outperform firms with low book-to-market (i.e., growth stocks);
- UMD (Up minus down, also called “Momentum”): stocks that have performed well over the past 6 months tend to outperform stocks that have performed badly over the past 6 months.

As the US Equity market is the most liquid, we decide to choose the MSCI US Indices tackling Growth, Momentum, Small, and Value stocks. Such long-short portfolios are constructed following the rationale described above when presenting the Carhart model.

The main result is that, over the past 22 years, the put writing strategy provided an efficient diversification tool when the monthly return of the S&P500 turned out to be negative (108 times), relative to the factor-mimicking portfolios. Indeed, as can be seen from the chart below, conditioning on S&P 500 return being negative, the frequency of positive returns registered by the put writing strategies is much higher than that of the factors under assessment. We have seen how the put writing strategy tends to suffer in case of severe market crashes. This notwithstanding, this last result seems to suggest that, other things being equal, if the market performance is (only slightly) negative, the put writing strategy proved to be an appropriate instrument to diversify out the risk of a portfolio.



Source: Bloomberg, MIMS estimate

In this report, we have analyzed the reasons behind the existence of the so-called volatility risk premium and the empirical evidence justifying the willingness of investors to systematically short OTM put options.

After having introduced the difference between realized and expected volatility, we have shown how the latter was (almost) systematically higher than the former over the past 22 years, both in the US and in Europe. Then, we have provided empirical evidence in favor of this phenomenon and discussed the impact of different asset classes and sectors. After that, we have assessed the argument connecting option pricing and dividend yields, introducing the notion of OTM put puzzle, which, however, proves to be not crucial under current market conditions.

In the last section of the report, we have presented a set of potential strategies to harvest volatility risk premium and studied deeply the strategy entailing a systematic short position on put options (i.e., Put Writing). In that section, we emphasized the differences between the implementation of such a strategy in the EU and US equity markets. More precisely, one euro invested in 2001 according to such a strategy would be worth EUR 2.60 as of November 2022, compared to EUR 0.78 received by investing in the Euro Stoxx 50 itself. The difference is not as extreme in the case of the US equity market: one dollar invested in 2001 would be worth USD 3.25 by investing in the equivalent strategy on the US equity market as of November 2022, compared to USD 2.94 received by investing in the S&P 500 itself. Finally, we have provided evidence in favor of the diversification benefits offered by this strategy. This last feature might reveal to be crucial under current market conditions.

- Markowitz, H. "Portfolio Selection", 1952
- Black, F. & Scholes, M. "The pricing of options and corporate liabilities", 1973
- Agrawal, S. "Implied volatility and skew as an indicator of market direction", 2021
- Heston, S.L. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", 1993
- Cont, R. & De Fonseca, J. "Dynamics of implied volatility surfaces", 2002
- Goetzmann, W., Kim D. & Schiller R. "Crash beliefs from investor surveys", 2016
- Lombardi, M. & Schrimpf, A. "Volatility concepts and the risk premium", 2014
- Eraker, B. "The volatility premium", 2009
- Carhart, M. "On persistence in mutual fund performance", 1997
- Fama, E. & French, K., "Common risk factors in the returns on stock and bonds", 1993
- Deutsche Börse Group. "Stoxx® Strategy Index Guide", 2022
- <https://www.cboe.com/us/indices/dashboard/put/>

Disclaimer

This research material has been prepared by Minerva Invest. Minerva Invest specifically prohibits the redistribution of this material in whole or in part without the written permission of Minerva Invest. The research officer(s) primarily responsible for the content of this research material, in whole or in part, certifies that their views are accurately expressed, and they will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this research material. Whilst we have taken all reasonable care to ensure that the information contained in this publication is not untrue or misleading at the time of publication, we cannot guarantee its accuracy or completeness, and you should not act on it without first independently verifying its contents. Any opinion or estimate contained in this report is subject to change without notice. We have not given any consideration to and we have not made any investigation of the investment objectives, financial situation or particular needs of the recipient or any class of persons, and accordingly, no warranty whatsoever is given and no liability whatsoever is accepted for any loss arising whether directly or indirectly as a result of the recipient or any class of persons acting on such information or opinion or estimate. You may wish to seek advice from a financial adviser regarding the suitability of the securities mentioned herein, taking into consideration your investment objectives, financial situation, or particular needs, before making a commitment to invest in the securities. This report is published solely for information purposes, it does not constitute an advertisement and is not to be construed as a solicitation or an offer to buy or sell any securities or related financial instruments. No representation or warranty, either expressed or implied, is provided in relation to the accuracy, completeness or reliability of the information contained herein. The research material should not be regarded by recipients as a substitute for the exercise of their own judgement. Any opinions expressed in this research material are subject to change without notice.

© 2022 Minerva Investment Society