A Story of Volatility Investing and Trading Part III - The Need for Tailored Solutions : the Third Generation of Volatility

Study carried out by the Quantitative Practice
Special thanks to Pierre-Edouard THIERY

canopee
empowering *ecosystem*

www.canopee-group.com

Summary

Introduction 1

1. Corridor Variance Swaps: Constraints on the Underlying Price 1

1.1. Up-Variance or Down-Variance Corridor Swaps 1

1.2. Replication of Corridor Variance Swaps 2

2. Conditional Variance Swaps Are Close But Different To Corridors 2

3. Gamma Swaps 2

4. Beyond the World of Variance Swaps: Options on Realized Variance 4

Conclusion 4

References 4

canopee
empowering *ecosystem*

Introduction

This work is the third paper from a series of three entitled "A Story of Volatility Investing and Trading". The first paper was dedicated to the options-based approaches, and the second one introduced variance swaps, and how those contracts can be traded when it comes to volatility investing.

Over the past decades, the demand for such products has increased, and investors have grown ever more accustomed to them: variance swaps are now widely used within the industry. However the appetite has also grown for new products, enabling investors to take ever more precise views on volatility. For instance, investors want to be able to bet on volatility while taking into account the context: is the market going up or down? To do so, a new wave of volatility instruments, known as the third generation of volatility products, has emerged, with the rise of variations on variance swaps, like corridor variance swaps, conditional variance swaps or gamma swaps.

Those products can then be used for many purposes by investors: to take specific views on the skew of implied volatility, to get better prices than those offered by traditional variance swaps, or to allow easier hedging strategies.

In this third and final paper we provide the reader with an extensive presentation of the third generation of volatility products.

1 Corridor Variance Swaps: Constraints on the Underlying Price

Corridor variance swaps, as the name implies, are a kind a variance swap: it pays the realized variance against a predetermined strike. However the computations of the realized variance are tweaked: they take into account the fact the price of the underlying must be in a given interval called "the corridor".

Definition 1 (Corridor Realized Variance)
If we denote $[L, U]$ the corridor, T the maturity of the contract, $t = 0$ the inception time, and if we assume that the time period $[0, T]$ is made of $N + 1$ days: $t_0 = 0 < t_1 < \cdots < t_N = T$, then the corridor realized variance is:

$$\sigma^2_{corridor} = \frac{252}{N} \sum_{i=1}^{N} 1_{L < S_{t_{i-1}} \leq U} \ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

So the payoff of a corridor variance swap is simply:

$$\left(\sigma^2_{corridor} - K_v^{corridor}\right) \mathcal{N}$$

where $\mathcal{N}$ is the notional amount and $K_v^{corridor}$ the variance strike of the contract.

It is straightforward that, if $L = 0$ and $U = \infty$, a corridor variance swap is equivalent to a variance swap. The addition of the corridor is a way of gaining exposure to the volatility of the underlying only when its price is comprised within some boundaries.

When it comes to corridor variance swaps, one should pay attention to one important detail regarding the corridor realized variance: even though the sum contains less than N elements, because of the activation function $1_{L < S_{t_{i-1}} \leq U}$, the result is divided exactly by N.

In practical terms, it means that the corridor realized volatility tends to be smaller than the realized volatility over the same period of time, insofar as we do not take into account all the returns of the period. This explains why we expect the variance strikes of corridor variance swaps to be inferior to the strikes of variance swaps.

Corridor variance swaps are often offered in one of the two typical forms: up-variance or down-variance.

1.1 Up-Variance or Down-Variance Corridor Swaps

An up-variance corridor swap is a corridor swap where the lower bound L is a fixed number, and the upper bound U is infinite:

$$\sigma^2_{up-variance} = \frac{252}{N} \sum_{i=1}^{N} 1_{L < S_{t_{i-1}}} \ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

An down-variance corridor swap is a corridor swap where the lower bound L is zero, and the upper bound U is equal to a fixed number.

$$\sigma^2_{down-variance} = \frac{252}{N} \sum_{i=1}^{N} 1_{S_{t_{i-1}} \leq U} \ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

The interest of those two types of corridor variance swaps lies in the fact that they allow investors to easily take a view on the skew of implied volatility.

Let us assume that an investor expects the implied volatility skew to steepen, meaning that the 90% implied volatility increases while the ATM implied volatility remains constant or decreases. One solution consists in buying a down-variance corridor swap, and selling an up-variance corridor swap. Here is a heuristic which explains why such a strategy may be profitable.

If the 90% implied volatility increases, it means that the price of the corresponding European put option also increases. This may be explained by a flare-up of risk-aversion when market participants anticipate a sharp decrease in the price of the underlying. So the investor believes that one consequence of his or her bet is that the price will decline below a given level L, and that realized volatility will be high since volatility jumps when the price crashes.

By being long a down-variance corridor swap, the investor receives the corridor realized variance: if the price actually goes below L, many returns will be taken into account when computing the corridor realized volatility, and the investor gets this volatility at maturity. On the contrary, being short an up-variance corridor swap is a way of getting the corridor variance strike at maturity, without having to pay anything if, for instance, the price of the underlying has stayed below the level U during the lifetime of the contract.

This quick example shows how corridor variance swaps may be used to take asymmetric bets on volatility: the investor bet on the volatility in a context of decreasing or increasing prices.

1.2 Replication of Corridor Variance Swaps

The replication of such swaps has been studied extensively by Carr and Lewis [1]. We only set forth here the main results regarding this aspect of corridor variance swaps.

Intuitively, it is possible to have insight in the replication strategy once we notice that, under a few conditions:

$$\frac{252}{N}\sum_{i=1}^{N}\ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

$$= \frac{252}{N}\sum_{i=1}^{N}1_{X<S_{t_{i-1}}}\ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2 + \frac{252}{N}\sum_{i=1}^{N}1_{S_{t_{i-1}}\leq X}\ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

meaning that a variance swap can be seen as the sum of an up-variance and a down-variance corridor swaps.

So the replication of a up/down-variance corridor variance swaps is similar in weighting to the one of a variance swap. However it requires less options, especially the ones which are the most expensive or which do not even exist. For instance, when considering an up-variance corridor swap, its replicating strategy does not involve put options with very low strikes, which are among the most problematic options for the replication of a variance swap.

2 Conditional Variance Swaps Are Close But Different To Corridors

Conditional variance swaps are very similar to corridor variance swap. There is only one major, albeit small, difference in the definition: when computing the realized variance, we divide the sum by the exact number of elements in it. The realized variance is based on a true average of the log returns which are taken into account.

So, if we use the same notations as above, the payoff of a conditional variance swap is the following:

$$\left(\sigma^2_{cond} - K_v^{cond}\right)\mathcal{N}$$

where

$$\sigma^2_{cond} = \frac{252}{\sum_{i=1}^{N}1_{L<S_{t_{i-1}}\leq U}}\sum_{i=1}^{N}1_{L<S_{t_{i-1}}\leq U}\ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

This slight difference may have a huge impact in practical terms. Indeed, if an investor is long a corridor variance swap, he or she receives the corridor realized variance. However, each return which is not taken into account because the underlying price is outside the corridor has a negative impact on the amount the investor gets at maturity: the sum does not increase because of the activation function, but the denominator will still be the total number of days N. With a conditional variance swap, such returns are not considered, and so have no impact at all on the investor long the conditional variance swap.

Since the investor is purely exposed to the volatility between the chosen range $[L, U]$, the risk no longer resides in the fact that the underlying trades outside the range, but inside it. To put it in simple words: a corridor variance swap is a story of the underlying price between inside or outside the range, whereas a conditional variance swap is a story of how the underlying moves within the range.

Let us consider the following simple example: a conditional variance swap with maturity equal to six months, and a given range $[L, U]$. We assume that the underlying price spends only two days within the range between the inception of the contract and maturity. A corridor realized variance would have been extremely low, since we divide a sum of two values by the total number of trading days during six months, i.e. 126. But, when it comes to the conditional variance, this number is no longer small: it is only the average of the two returns. So, how the underlying moves during those two days is of massive importance when being long a conditional variance swap.

Once those differences are well understood, conditional variance swaps can be used for the same purposes as corridor variance swaps: for instance, down-variance conditional variance swaps allow investors to take a view on the levels of volatility when market is declining.

When it comes to replication, the replicating strategy of a conditional variance swap is fairly simple, insofar as it involves an initial static hedge, and then only some delta-hedging. For further details, the reader can have take a look at [2].

3 Gamma Swaps

Gamma swaps have been devised in order to take into account a characteristic of variance swap: the cash gamma of such a contract is constant, meaning that the gamma exposure of a variance swap does not depend on the underlying's price.

Indeed, as shown by [3], the gamma of a variance swap is equal to:

$$\Gamma = \frac{2}{T}\frac{1}{S_t^2}$$

So, if we compute the cash gamma:

$$\Gamma_{cash} = \Gamma\frac{S_t^2}{100} = \frac{2}{100T}$$

The cash gamma is not a function of either the price of the underlying, or time.

Based on this, gamma swaps are meant to be variance swaps where the notional is proportional to the level of the underlying, and thus to have a constant gamma. Such products are helpful for investors who think in terms of units of the underlying, instead of the cash value of such units.

Definition 2 (Gamma Swap)
If we denote $\mathcal{N}$ the notional, T the maturity date, K_{gamma} the gamma strike level, $t = 0$ the inception time, and if we assume that the time period $[0, T]$ is made of $N + 1$ days: $t_0 = 0 < t_1 < \cdots < t_N = T$, the payoff of a gamma swap is

$$\left(\gamma - K_{gamma}\right) \mathcal{N}$$

where

$$\gamma = \frac{252}{N} \sum_{i=1}^{N} \frac{S_{t_i}}{S_{t_0}} \ln\left(\frac{S_{t_i}}{S_{t_{i-1}}}\right)^2$$

A gamma swap contract displays some advantage over a variance swap contract in case of sharp downward move of the underlying price. Let us assume that the underlying defaults: at some date t_i before the end of the contract, S_{t_i} falls to 0. In the case of a variance swap, the realized variance would spike up to the infinity:

$$\ln\left(\frac{S_{t_i} = 0}{S_{t_{i-1}}}\right)^2 = \infty$$

So the investor which is short the variance swap should pay an infinite realized variance. A solution to circumvent this risk is to cap the realized variance when writing the contract. This is no longer necessary with a gamma swap. Indeed, if the underlying stock defaults, the factor $\frac{S_{t_i}}{S_{t_{i-1}}}$ in the sum ensures that the corresponding term is equal to zero instead of the infinity:

$$\frac{S_{t_i} = 0}{S_{t_0}} \ln\left(\frac{S_{t_i} = 0}{S_{t_{i-1}}}\right)^2 = 0$$

This means that there is no need to impose a cap on the gamma realized variance in a gamma swap. Subsequently, a gamma swap has a lower fair strike than a variance swap, since its payoff is usually lower than the one of a variance swap.

The replication of gamma swaps is very similar to the one of variance swaps. It is indeed possible to apply the same approach, based on the Carr formula, in order to replicate a gamma swap. By definition, the continuous gamma realized variance is:

$$\Gamma_{0,T} = \frac{1}{T} \int_0^T \sigma_t^2 \frac{S_t}{S_0} dt$$

Since this is the unknown quantity at the inception of the contract $t = 0$ which has to be paid at maturity, we want to find a means to replicate it. Similarly to what we did for variance swaps [4], we apply the Ito formula to the function:

$$f(F_t, t) = e^{rt}\left(F_t \ln\left(\frac{F_t}{F_0}\right) - F_t + F_0\right)$$

where F_t is the futures price of the underlying, having the same instantaneous volatility than the underlying price.

The Ito formula gives us:

$$\frac{1}{2}\int_0^T S_t\sigma_t^2 dt = e^{rT}\left[F_T \ln\left(\frac{F_T}{F_0}\right) - F_T + F_0\right]$$

$$-r\int_0^T e^{rt}\left[(1 + F_t)\ln\left(\frac{F_t}{F_0}\right) - F_t + F_0\right] dF_t$$

Using the Carr formula: for any payoff $g(F_T)$, and any κ

$$g(F_T) = g(\kappa) + g'(\kappa)\left[(F_T - \kappa)^+ - (\kappa - F_T)^+\right]$$

$$+ \int_0^{\kappa} g''(K)(K - F_T)^+ \, dK + \int_{\kappa}^{\infty} g''(\kappa)(F_T - K)^+ \, dK$$

In our situation, we use $\kappa = F_0$ and the function $g(x) = x\ln\left(\frac{x}{F_0}\right) - x + F_0$:

$$F_T \ln\left(\frac{F_T}{F_0}\right) - F_T + F_0 = \int_0^{F_0} \frac{1}{K}(K - F_T)^+ \, dK + \int_{F_0}^{\infty} \frac{1}{K}(F_T - K)^+ \, d$$

Therefore, to replicate the gamma realized variance at maturity, a trader must buy a continuum of put options with strikes ranging from 0 to F_0, and call options with strikes ranging from F_0 to ∞. The weight for each option is proportional to the inverse of the strike. Then the trader needs to roll a futures position, holding at date t:

$$-\frac{2r}{TS_0}\left[(1 + F_t)\ln\left(\frac{F_t}{F_0}\right) - F_t + F_0\right]$$

The initial cost of the replicating strategy is given by the part involving options:

$$\frac{2e^{rT}}{TS_0}\left[\int_0^{F_0} \frac{1}{K}P(K)dK + \int_{F_0}^{\infty} \frac{1}{K}C(K)dK\right]$$

where P and C refer to the prices of put and call options. Therefore, once we take into account the discount factor, we get the fair value of the gamma strike:

$$\frac{2e^{2rT}}{TS_0}\left[\int_0^{F_0} \frac{1}{K}P(K)dK + \int_{F_0}^{\infty} \frac{1}{K}C(K)dK\right]$$

A gamma swap can be used to take position on the implied volatility skew. Let us assume that an investor expects the skew to steepen. he or she can make such a bet by going short a gamma swap and long a variance swap. Such a combination allows the investor to be long downside variance and short upside variance.

4 Beyond the World of Variance Swaps: Options on Realized Variance

In parallel to the development of variance swaps and all the closely-related products, options written on realized variance have emerged.

Definition 2 (Options on Realized Variance)
If we denote T the maturity date, K_v a variance strike level, $t = 0$ the inception time, and if we assume that the time period $[0, T]$ is made of $N + 1$ days: $t_0 = 0 < t_1 < \cdots < t_N = T$, the payoff of a call on realized variance is

$$\left(\sigma^2 - K_v\right)^+$$

where

$$\sigma^2 = \frac{252}{N} \sum_{i=1}^{N} \ln \left(\frac{S_{t_i}}{S_{t_{i-1}}} \right)^2$$

Similarly, the payoff of a put option on the realized variance is

$$\left(K_v - \sigma^2\right)^+$$

Those products suffer from the same danger as variance swaps: the realized variance may spike over the lifetime of the contract, thus exposing those investing in such products to high levels of risk. To avoid such situations, those options are often capped or floored.

When it comes to replication, options on realized variance, contrary to variance swaps and their variations, are not replicable in a model-free manner. Pricing those options requires a model. Once such assumptions are made, it is possible to find practical formulas: we do not delve into such details in this paper, so the reader interested in a full presentation of such mathematical questions can take a look at [5] and [6].

Conclusion

The journey across the world of volatility investing and trading has come to an end. Those three papers were the opportunity for a full presentation of the most salient aspects of volatility investing, be it from a historical and academic point of view or from a more practical one.

In this third paper, we introduced some of the most complex volatility products. Most of them are derived from the variance swaps. Corridor and conditional variance swaps have been devised in order to enable investors to bet on volatility while taking into account the context. Thanks to those contracts, it is now possible to take a position on the implied volatility skew without having to rely on options. Gamma swaps are a refinery of variance swaps. Tweaking the definition of the realized variance is a way of limiting the potential loss an investor may face; therefore it reduces the risk of such a contract, and provides a cheaper alternative to variance swaps. Beyond variance swaps and all their variations, options on realized variance have also emerged. Those options however, even though their definition may sound fairly simple, as not easy to handle, especially when it comes to valuation.

References

[1] P. Carr and K. Lewis. Corridor variance swaps. *Risk*, 2004.

[2] P. Allen, S. Einchcomb, and N. Granger. Conditional variance swaps. *Product Note. JP Morgan*, 2006.

[3] K. Demeterfi. A guide to volatility and variance swap. *Journal of Derivatives*, 1999.

[4] PE. Thiery. A story of volatility trading, part i: Historical approaches through options. *Awalee Papers*, 2021.

[5] P. Carr and R. Lee. Realized volatility and variance: options via swaps. *Risk*, 2007.

[6] A. Sepp. Pricing options on realized variance in the heston model with jumps in returns and volatility. *Journal of Computational Finance*, 2008.

Nous sommes Canopee, un cabinet de conseil indépendant, spécialiste en Finance, DATA et transformation digitale.

Nous sommes l'*empowering ecosystem* ! Un écosystème en perpétuel mouvement.
Un écosystème qui donne le pouvoir à nos collaborateurs, nos projets et nos clients de se nourrir de l'émulation collective pour lui donner de la force.

Depuis 2009, nous intervenons dans les secteurs de la BFI de l'Asset Management, la banque privée ou de détail, les services financiers et l'assurance et cela auprès de clients tels que SG, HSBC,BNP Paribas ou encore ALLIANZ GI et AXA IM. Des écosystèmes réglementaires et spécifiques qui ont été nos premiers terrains de jeu.
C'est ici que nous avons su développer notre agilité, nos compétences, notre sens de l'engagement. Aujourd'hui, nous grandissons et continuons à nous investir et nous engager sur des écosystèmes spécifiques, techniques, mais diversifiés tels que l'industrie, le retail ou encore la pharma.
Toujours autour de nos 3 expertises que sont la finance, la data et la transformation digitale.

canopee
empowering *ecosystem*

www.canopee-group.com