Measuring cross-gamma risk

Harold Chicot

"Theory and practice sometimes clash. And when that happens, theory loses. Every single time."

Linus Torvalds

1 Introduction

During the last decade many financial institutions have seen a growth in the complexity of their products from a pricing point of view. There are multiple reasons for this, among which:

1. - the need to price within a new regulation framework (for instance interest rate instruments now require using different curves for rates diffusion and cash flows discounting, taking into account the credit risk of each counterparty, the nature of collateral, margin payments, ...)
2. - the need to take into account model dynamics, for instance the stochastic aspect of parameters such as volatility or correlation in order to manage mismatches between pricing and actual hedging during the life of the product...
3. - the rise in computing power which now allows running in real time multi-factor monte-carlo processes that couldn't be handled fast enough a few years back.

As a result, in terms of pricing, many simple market products have turned into complex multi-factor exotics. A bond option for instance was originally priced using a plain Black-Scholes closed formula, assuming a log-normal process for the bond price. In a more modern framework the same option can be seen as a payoff depending on a (stochastic) rate curve and a credit curve (with stochastic default process), ie a credit/rates hybrid. Going from single to multi-factor pricing implies not only more complex modeling but also an exponential growth of risk parameters, in particular cross risks now have to be taken into account.

In this paper, we look at different types of unhedged delta risks that appear in trading books due to various constraints, including gamma and cross-gamma. We then provide some risk metric tools that help us decide whether the extra $P\&L$ variance generated by these risks is acceptable or not, an important step in the validation of an exotic product and in the sizing of a risk framework for that product.

2 Some theoretical background on expected P&L and its variance

2.1 P&L of a delta-hedged product

We follow [1] and consider a short position in an exotic option with price $P(t, S)$, where S is a tradable stock. At time t we hedge this option with a quantity Δ of stock. Over a short timelapse δt the $P\&L$ of the hedged position is given by:

$$P\&L = -[P(t+\delta t, S+\delta S) - P(t,S)] + rP(t,S)\delta t + \Delta(\delta S - rS\delta t + qS\delta t)$$

where r is the (deterministic) interest rate and q the (deterministic) repo rate including dividend yield.

Choosing $\Delta = \frac{\partial P}{\partial S}$ to cancel the first-order term in δS and expanding the $P\&L$ to second order powers of δS and δt, we get the following standard expression:

$$P\&L = -\left(\frac{\partial P}{\partial t} - rP + (r-q)S\frac{\partial P}{\partial S}\right)\delta t - \frac{1}{2}S^2\frac{\partial^2 P}{\partial S^2}\left(\frac{\delta S}{S}\right)^2$$

Defining the lognormal historical volatility $\hat{\sigma}$ by $<(\frac{\delta S}{S})^2>= \hat{\sigma}^2\delta t$, where $<>$ is the "average" operator, and using the risk-management argument stating that the portfolio should not lose or make money on average the previous formula can be rewritten [1] as:

$$P\&L = -\frac{S^2}{2}\frac{\partial^2 P}{\partial S^2}\left(\frac{\delta S^2}{S^2} - \hat{\sigma}^2\delta t\right) \tag{1}$$

Note that, when we work in the Black-Scholes model framework (S follows a lognormal process with constant volatility $\sigma = \hat{\sigma}$), the sum of the above $P\&Ls$ vanishes as $\delta t \to 0$. This is the classic "delta replication" principle by which an option price can be replicated by continuous delta-hedging. In reality actual prices do not follow lognormal processes and delta-hedging is discrete, typically done on a daily basis, so δt is small enough to use the power expansion above, but we cannot use the replication argument: the sum of $P\&Ls$ will not vanish.

For multi-asset products the price is of the form $P(t, S_1, S_2, ..., S_n)$ and assuming all the S_i are delta-hedged we get in a similar way:

$$P\&L = -\frac{1}{2}\sum_{i,j}\Phi_{i,j}\left(\frac{\delta S_i}{S_i}\frac{\delta S_j}{S_j} - C_{ij}\delta t\right) \tag{2}$$

where $\Phi_{i,j} = S_i S_j \frac{\partial^2 P}{\partial S_i \partial S_j}$ (known as **dollar gamma** for $i = j$ and **dollar cross-gamma** for $i \neq j$), and $(C)_{i,j}$ is a positive matrix which can be interpreted as a covariance matrix.

Equations (1) and (2) are special cases of the more general formula:

$$P\&L = -\frac{1}{2}\sum_{i,j}\Phi_{i,j}\left(\frac{\delta x_i}{x_i}\frac{\delta x_j}{x_j} - C_{ij}\delta t\right) \tag{3}$$

where the x_i can be linearly hedged (locally) by market instruments (such as plain linear products ("delta-one") but also volatility or gamma for example).As before $\Phi_{i,j} = S_i S_j \frac{\partial^2 P}{\partial x_i \partial x_j}$, and $(C)_{i,j}$ is a positive matrix known as **Implied Covariance** (or Model Covariance). Another way to express (3) is therefore:

$$P\&L = -\frac{1}{2}\sum_{i,j}\Phi_{i,j}\left(RealizedCovar(\delta x_i, \delta x_j) - ImpliedCovar(\delta x_i, \delta x_j)\right) \tag{4}$$

In what follows we will be looking at a 2-dimensional version of (2), which can be written after separating gamma and cross-gamma as:

$$P\&L = \underbrace{-\frac{1}{2}\sum_{i}\Phi_{i,i}\left(\left(\frac{\delta S_i}{S_i}\right)^2 - \hat{\sigma_i}^2\delta t\right)}_{\text{Gamma P\&L}}\underbrace{-\frac{1}{2}\sum_{i,j,i\neq j}\Phi_{i,j}\left(\frac{\delta S_i}{S_i}\frac{\delta S_j}{S_j} - C_{ij}\delta t\right)}_{\text{Cross-Gamma P\&L}} \tag{5}$$

[1] See [1] p4 for details

The gamma part of the right hand-side can be hedged locally by trading options on S_1 and S_2. In general there are no liquid market instruments to hedge the cross-gamma part, so we will look at ways to estimate the impact of this part on the $P\&L$ variance in order to decide if the product has a reasonable risk-reward and if we should decide to trade it or not. Note that we want to concentrate on the *structural* $P\&L$ variance due to the payoff, so for simplicity we will put aside the impact of realized vs implied covariance by assuming no drift on the final $P\&L$, ie

$$\mathbb{E}\left[\sum_{n=0}^{N}\left(-\frac{1}{2}\sum_{i,j}\Phi_{i,j}\left(\frac{\delta S_i}{S_i}\frac{\delta S_j}{S_j}-C_{ij}\delta t\right)\right)_{|t=t_n}\right]=0$$

(after choosing a time discretization $[0=t_0,t_1,...,t_N=T]$ where T is the option maturity). In brief we are interested in the following:

$$StDev\left(\sum_{n=0}^{N}\left(-\frac{1}{2}\sum_{i,j,i\neq j}\Phi_{i,j}\left(\frac{\delta S_i}{S_i}\frac{\delta S_j}{S_j}-C_{ij}\delta t\right)\right)_{|t=t_n}\right) \tag{6}$$

3 Real world fluctuations of delta risk

Before dwelling on cross-gamma per se we will describe other sources of delta variations as they sometimes appear on trading desks. This section is independent of the rest of the article.

3.1 Deltas for the practitioner

Consider a product with price $P(S)$, where S is a tradable asset, and its theoretical delta: $\Delta=\frac{dP}{dS}$. In the real world delta is neither instantaneous nor unique (we will not discuss the fact that delta may not even be defined or meaningful, for example when an option is close to maturity - anyone who had to hedge a decent option size close to the strike on expiry day will understand). In fact you can choose to define it in a number of ways:

1. - Choosing the shift type: additive or multiplicative

Additive Delta	Multiplicative Delta
$\frac{P(S+\epsilon)-P(S)}{\epsilon}$	$\frac{P(S\times(1+\epsilon))-P(S)}{\epsilon S}$

This choice is usually dictated by the underlying itself. On the equity side, with thousands of underlyings and prices in a wide range of values, using an additive shift would mean choosing a different one for each underlying, which is impractical to say the least. Also, equity diffusion processes have been traditionally modelled as lognormal, and in that respect a multiplicative shift is more natural . On the rates side , with current rates level being close to 0 and possibly positive or negative, we have to use an additive shift or else delta would depend too much on the rates level, and on top of that using a multiplicative shift would mean using a "down delta" (as described below) on negative rates versus an "up delta" on positive rates. [2]

[2] Nowadays choosing a normal-like process for rates, or at least one that allows for negative rates, goes without saying, but it wasn't the case when rates were high, in fact not so long ago swaptions were still priced using a black-scholes lognormal process for the underlying swap (in which case a multiplicative delta could make sense). It's only in the late 90s when JPY short-term rates started to get close to zero that traders started switching from "% vols" to "bp vols".

2. Choosing the discretized formula: up, down or average are common choices. Here is a summary of delta formulas depending on shift type and discretization type:

	Up	Down	Average
Additive Delta	$\frac{P(S+\epsilon)-P(S)}{\epsilon}$	$\frac{P(S)-P(S-\epsilon)}{\epsilon}$	$\frac{P(S+\epsilon)-P(S-\epsilon)}{2\epsilon}$
Multiplicative Delta	$\frac{P(S(1+\epsilon))-P(S)}{\epsilon S}$	$\frac{P(S)-P(S(1-\epsilon))}{\epsilon S}$	$\frac{P(S(1+\epsilon))-P(S(1-\epsilon))}{2\epsilon S}$

Depending on this choice, the resulting delta can be quite different for products with large gamma and will therefore imply different $P\&L$s and $P\&L$ variances.

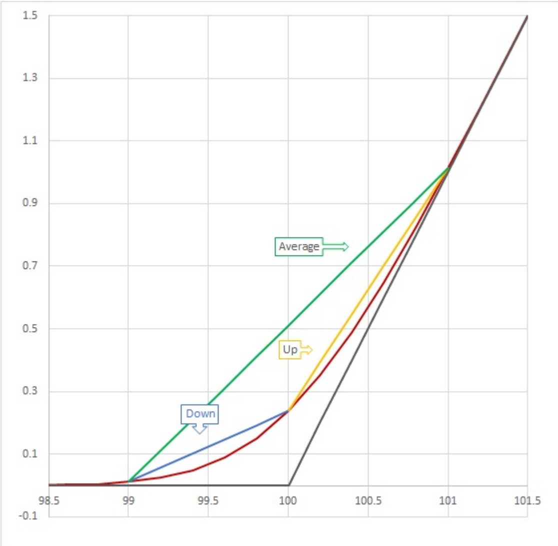

Figure 1: Different deltas for a long atm call position, depending on the choice of discretized formula. The Average Delta is the closest to theoretical value, the Up Delta will induce overhedging, the Down Delta will induce underhedging.

3. Choosing the shift value. If it's small you are closer to theoretical value but if it is too small the difference in price might be close to the approximation error of the price and you would be pricing noise (obviously noise is very limited when using an analytical formula but here we have in mind a general exotic payoff). If the shift is too big you will miss out on local behaviour and ultimately "mishedge" (in the sense that delta hedging could be very different from the one implied by a theoretical "no hedge cost and continuous hedging" price). Also note that it is reasonable to ask that the shift be larger than the bid-ask spread of the asset price in the market, this may sound obvious but should be taken into account for less liquid markets. [3].

Example:

Consider the following plain vanilla equity call on an asset S: expiry $T = 3$ months , forward $F = 100$, strike $K = 110$, volatility $\sigma = 20\%$. It has the following (multiplicative) deltas

	Analytic	Up 1%	Down 1%	Average +/-1%
Delta	18.3%	19.7%	17.0%	18.4%
Delta Spread	0%	1.4%	-1.3%	0.1%
Implied Volatility Spread	0%	-1.0%	1.0%	-0.1%

[3]For instance there were times when the daily moving range of TWD swap rates (mid) was smaller than the bid-ask spread on the swap ! This would obviously compromise the economic rationale of delta hedging.

On the last line of the table we have converted the delta spread into a volatility spread (for example in the Up1% case we have $AnalyticDelta(marketvolatility - 1\%) = Up1\%Delta(marketvolatility)$).

We need to be careful when interpreting a delta spread in terms of volatility spread since obviously the relation between the two is not linear, which is another way of saying that vanna ($\frac{\partial^2}{\partial\sigma\partial S}$) is not constant. However, in the above example, using an Up 1% delta, when the forward F is in the range $[95; 105]$ vanna is in the range $[0.9; 1.3]$ and implied vol spreads in $[-0.7\%; -2.0\%]$ for the Up 1% Delta case. So it is fair to consider that, in that forward range, the implied volatility is at least 0.7% lower than market volatility. If the forward level is above the strike, say in the range $[115; 125]$ there is a similar implied volatility spread but in the opposite direction: the volatility spread is positive. Altogether using the Up 1% Delta from your system is equivalent to changing the volatility smile (here by at least 0.7% on the 10% OTM strikes).

The above example is very simple but the idea stays the same for more complex payoffs: depending on your choice of definition for delta, your realized delta-hedging strategy is different from the one implied by the theoretical price, thereby changing your expected $P\&L$ and its variance.

3.2 Sources of variance in delta-hedging

In the previous paragraph we saw how a *choice* of delta calculation parameters would impact the $P\&L$ variance. We now look at reasons for delta slippage that are external to the trading desk.

3.2.1 Exotic desk vs rest-of-the-world

Exotic desks (hybrid desks in particular) are often using different booking systems at the same time, for instance one for the exotic product and one (or more) for their vanilla options hedges (delta hedges are booked in either or both systems...). More often than not there are difference in definitions or parameters between the systems.

For example consider an exotic option that is locally gamma-hedged with vanilla options. Suppose that it is priced in Monte-Carlo using an Up1% delta to avoid pricing too much noise [4]. If the vanilla system is using an analytic formula for deltas (ie the "exact" derivative), you will have a residual delta from the difference in delta definitions as seen before. One way to avoid this mismatch would be to use an Up1% delta on the vanilla side but vanilla systems are usually set up for vanilla traders, and they certainly wouldn't like to hedge using a delta that doesn't match their analytic delta (because they would suffer from extra PL variance, as seen before). As a result, the exotic desk will effectively be forced to hedge using deltas that do not match its model deltas, with a direct impact on the expected $P\&L$ and its variance.

This is a simple illustration of the following global issue: choosing between full flexibility of trading parameters for each trading book and global pricing coherence within a financial institution. Let's take a look at this dilemma from different angles:

1. Traceability : obviously a unique set of parameters is easier for reporting.
2. Risk and $P\&L$ variance control : a unique set of parameters for all trading books is close to being the promise of a running disaster for exotic desks. Indeed, by definition, exotic products depend on a number of parameters that are not market parameters, and the ability to keep the sales margin during the life of the product depends heavily on the choice of the hedging strategy. Further, controling pricing parameters for each desk means being able to reduce $P\&L$ variance for each book and therefore the global $P\&L$ variance of the company.

[4] Even though the Average Delta looks like a better proxy, systems could be using an Up Delta for historical reasons, internal conflicts, lack of resources to change the setup, computation time (calculating Price and Average Delta means calculating 3 prices instead of 2 for Up Delta), etc...

Electronic copy available at: https://ssrn.com/abstract=3502251

3. $P\&L$ friction between desks : a unique set of parameters certainly helps to check that an internal deal is booked correctly (as it would be neutral in terms of global $P\&L$). Unfortunately, it could imply $P\&L$ transfer and risk transfer from one book to another, as described in the following example.

Example: An exotic desk has sold a product with digital risk on, say, the Hang Seng Index (HSI), and is hedging some of this risk by selling a binary put strike 95% atm. The hedge is done internally with the vanilla desk. As the strike is below spot and HSI can be quite volatile, the vanilla desk decides to book the digital option conservatively as a vanilla put spread $(95\% - \alpha\%, 95\%)$, where $\alpha > 0$ and can be adapted depending on the risk appetite for hedging large gamma positions. Now if the exotic desk is forced to manage this hedge using the vanilla booking, it will be in the opposite position, ie an inflated $P\&L$ *and* an aggressive booking which will be more difficult to hedge ! Further, since the vanilla desk is likely to change the value of α during the life of the hedge, the exotic desk will be torn apart as it suffers from the resulting forced delta hedges and unwanted changes in its trading strategy.

Note that this situation also occurs for some "observable parameters" such as correlation. Given the many definitions of correlation (short term vs long term, window span, sampling period, instantaneous vs terminal, historical vs implied, etc...) and the very limited source of market prices that can be used for calibration, correlation is usually one of the most overlooked and debatable booking parameters. Given the variety of payoff sensitivities to correlation, reducing $P\&L$ variance would imply using different correlation parameters *for each product type*, which is often not the choice made by institutions because of the resulting extra checking processes.

3.2.2 Cross Risks are (nearly) everywhere

Now let's assume we have reasonable control of our pricing parameters. We are left with one of the main goals of exotic desks: understanding cross-risks ie how different parameters or risk metrics influence each other. For a single underlying we might look at the way gamma changes when spot changes (vanna), the way vega changes when spot volatility changes (volga), the way volatility changes when spot changes (smile), etc... For multiple underlyings new correspondances (correlations) come into play and in general we need to consider all the cross-risks: the way the delta of S_1 changes when the spot S_2 moves (cross-gamma), the way the gamma of S_1 changes when the spot S_2 moves, the way the vega of S_1 changes when the correlation between S_1 and S_2 changes, etc.. Recall from the introduction that nowadays products are priced using more complex sets of parameters, which implies dealing with more cross-risks. Dealing with these risks can be addressed in many different ways: static/dynamic hedging, global/local hedging, payoff shifts, parameters shifts, choice of model dynamics, and so on... A whole book could be written on this section but one of the key points to remember is that most of these cross risks *cannot* be hedged with market instruments so it is vital to examine their behavior and measure the associated risks before trading them.

In the following section we will take a detailed look at concrete examples involving cross-gamma, and we shall see that some "natural" hedging strategies are sometimes ineffective, costly, and may even result in an increase in $P\&L$ variance.

4 Cross-gamma and $P\&L$ variance

The effect of the combined moves of two underlyings, or cross-gamma, is the most basic cross-risk in the sense that it is based on the first level of observability: the spot prices of the underlyings [5]. Regardless of the system level of volatility or correlation, in general cross-gamma cannot be hedged

[5] As explained earlier nowadays the pricing of most products requires at least one stochastic underlying, one discounting rates curve and one credit spread curve, which implies cross-gamma risk between these 3 underlyings, even if the last 2 are deterministic they will be updated from time to time and induce delta rehedging on the stochastic underlying.

Electronic copy available at: https://ssrn.com/abstract=3502251

using other vanilla instruments, even locally: cross gammas must be accounted for and their risk is born through the life of the product. [6].

The effects of cross-gamma are three-fold:

1. on the expected $P\&L$ through the difference between realized and implied covariance as seen in equation (4) - as explained previously we will not consider this effect here.
2. on the expected $P\&L$ through hedging costs (crossing bid-offers).
3. on the $P\&L$ variance since the lack of local cross-gamma hedge implies potentially significant daily delta rebalancing until maturity.

We now focus on the following exotic product: a **Worst-Of Binary Call** (also known as a Double Binary Call), a product which is implicitly traded as a component of the more complex and actively traded Worst-Of Autocall for example. We consider a *portfolio made of a Worst-Of Binary Call and its hedges* (delta and gamma hedges as specified below) and examine the behaviour of its $P\&L$ for different strategies.

Setup:

Underlyings: S_1 and S_2, spot values $S_1(0) = S_2(0) = 100$.

Payoff: In 3 months the buyer of the option receives 100% if $Min(S_1, S_2) > 105$.

Model: independant lognormal processes with constant volatility $\sigma_1 = \sigma_2 = 20\%$, no drift, priced using a flat rate curve at 0%.

Delta Hedging with the underlyings (with a typical transaction cost of a liquid asset: 0.01% notional per delta trade).

Gamma Hedging with binary calls strike 105 expiring on the same date as the exotic (with a typical transaction cost of 0.5% shift on the strike).

Numerical results can be found below for each hedging scenario, assuming the exotic desk is short the exotic product. In the case of multiple gamma rehedging, in order to simulate a more realistic hedging frequency (thereby reducing hedging costs), we do not hedge gamma on a daily basis but only when the exotic gamma and hedge gamma differ by a relative 20% or more. The calculations are done by Monte-Carlo simulation of *the exotic book* (= exotic product + delta hedges+ gamma hedges), in particular, although the exotic payoff is European (ie depends only on the final values of S_1 and S_2), the book $P\&L$ is very much path-dependent. For comparison we also give the results for a single underlying binary call.

Legend:

- Initial Delta: a boolean which is TRUE if we delta hedge the product at inception.
- Initial Gamma: a boolean which is TRUE if we gamma hedge the product at inception.
- Initial Gamma Ratio: the ratio of hedge notional vs exotic notional for the inital gamma hedge.
- Rehedge Delta: a boolean which is TRUE if we delta rehedge the product daily until maturity.

[6] Note that this is not necessarily the case for other types of parameters: take an emerging market correlation for example, say 5y10y KRW IRS (forward swap rate) vs 5y2y KRW IRS. The fact that it is hard or impossible to find an implicit level for that correlation means that its risk management is smoother to some extent: it is unlikely to be updated frequently (for a good reason: it is hard to come by) and can be booked conservatively depending on the targeted volatility of the $P\&L$ and the bank's appetite for risk

- Rehedge Gamma: a boolean which is TRUE if we gamma rehedge the product until maturity.
- Initial Gamma: a boolean which is TRUE if we gamma hedge the product at inception.
- Book PL : the expected $P\&L$ assuming no transaction costs.
- Book PLStDev : the $P\&L$ variance assuming no transaction costs.
- Deltahedgecost : the transaction cost for all delta hedges.
- Deltahedgecost : the transaction cost for all gamma hedges.
- Total seller $P\&L$: expected $P\&L$ + all transaction costs.

Scenario#	1	2	3
Price	29.43%	29.43%	29.43%
Initial delta	TRUE	TRUE	FALSE
Initial gamma	TRUE	FALSE	FALSE
Initial gammaratio1	100.00%		
Initial gammaratio2	0.00%		
Rehedge delta	TRUE	TRUE	FALSE
Rehedge gamma	FALSE	FALSE	FALSE
Book PL wo hedge cost	0.00%	-0.06%	0.10%
Book PLSTDEV	0.00%	13.66%	46.74%
Deltahedgecost	-0.00%	-0.65%	0.00%
Gammahedgecost	-1.71%	0.00%	0.00%
Total seller PL	-1.71%	-0.71%	0.10%
PL/EXOPRICE	-5.79%	-2.41%	0.34%
STDEV/EXOPRICE	0.00%	46.42%	158.82%

Table 1: Binary Call

Is is worth stressing the fact that the $P\&L$ variance of the single underlying Binary Call which is only delta-hedged is huge. In a market that has poor liquidity in vanilla calls or puts, this would be the kind of variance we have to bear once liquidity has dried out (on top of the large transaction costs associated to poor liquidity).

Scenario#	1	2	3	4	5	6	7
Price	10.2338%	10.2338%	10.2338%	10.2338%	10.2338%	10.2338%	10.2338%
Initial delta	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE
Initial gamma	TRUE	TRUE	TRUE	TRUE	FALSE	FALSE	FALSE
Initial gammaratio1	32%	50%	10%	32%			
Initial gammaratio2	32%	50%	10%	32%			
Rehedge delta	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	FALSE
Rehedge gamma	TRUE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
Book PL wo hedge cost [7]	-0.01%	-0.01%	0.04%	-0.04%	0.03%	-0.01%	-0.05%
Book PLSTDEV wo hedge cost	5.19%	8.65%	8.80%	8.19%	10.09%	25.90%	30.37%
Deltahedgecost	-0.32%	-0.52%	-0.46%	-0.46%	-0.48%	-0.03%	0.00%
Gammahedgecost	-2.83%	-1.48%	-0.30%	-0.95%	0.00%	0.00%	0.00%
Total seller PL	-3.16%	-2.01%	-0.72%	-1.45%	-0.45%	-0.04%	-0.05%
PL/EXOPRICE	-30.89%	-19.65%	-7.04%	-14.17%	-4.40%	-0.39%	-0.49%
STDEV/EXOPRICE	50.67%	84.53%	85.99%	80.05%	98.55%	253.07%	296.77%
Totalhedgecost/price	-30.73%	-19.53%	-7.34%	-13.69%	-4.65%	-0.31%	0.00%

Table 2: Binary Worst-Of Call

In the above table we have simulated different strategies (with or without delta or gamma hedges, with different inital ratios, etc...) to try and find a compromise between variance reduction and transaction costs.

Key results:

1. Even with frequent gamma rehedging (Table 2 scenario 1) the stdev/price ratio of the Worst-Of Binary Call hedged portfolio is extremely high: at 50% it is comparable to that of a single digit which is only delta hedged (Table 1 scenario 2)!

2. The transaction costs of re-hedging the double digit kill any realistic margin on the product.

The above results tell us not to trade that product as it is, and we might ask ourselves which part of its risk profile has the most impact on variance. Let's look at the price surface as a function of S_1 and S_2 in the outright case and in the gamma hedged case (without delta hedge):

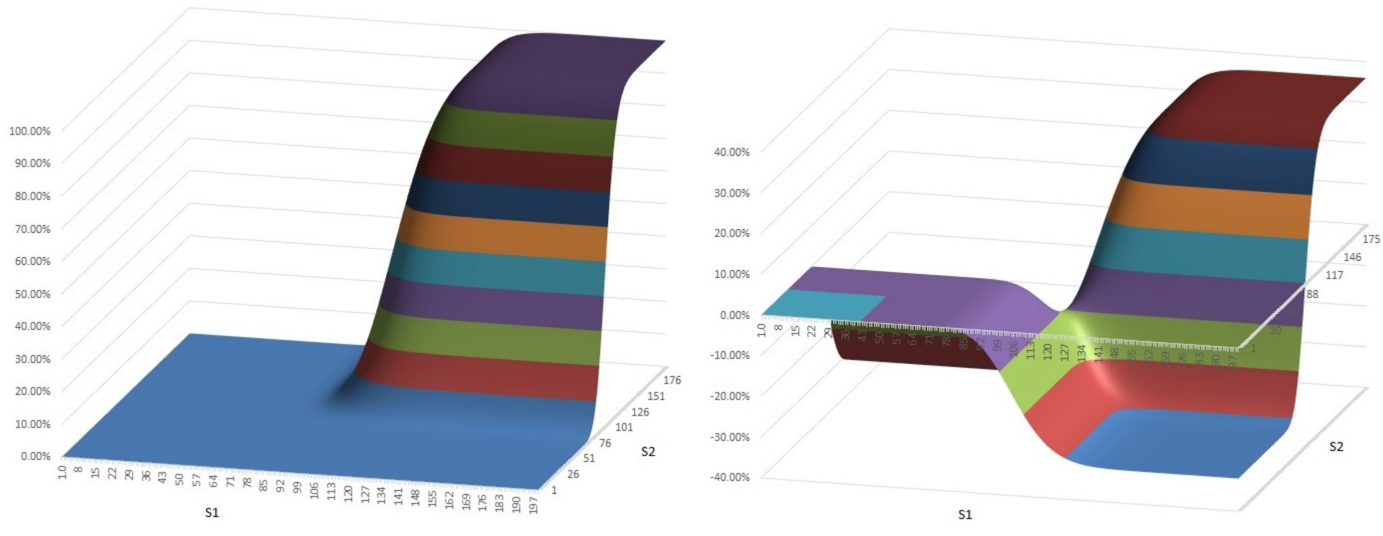

Figure 2: Price surface of a Binary Worst-Of Call (left) and the same with initial gamma hedges (right), no delta hedges, 3 months to maturity.

The above figure shows that even if gamma hedging neutralizes gamma *along the underlying axes* there is still a lot of gamma in other directions, especially when the underlyings are close to the strike. From a geometrical point of view the price surface still shows a lot of *curvature*.

Let's take a look at another product with cross-gamma: a **Spread Option**.

Setup:

Underlyings: S_1 and S_2, spot values $S_1(0) = S_2(0) = 100$.

Payoff: In 3 months the buyer of the option receives $Max(S_1 - S_2, 0)$.

Model: independant lognormal processes with constant volatility $\sigma_1 = \sigma_2 = 20\%$, no drift, priced using a flat rate curve at 0%.

Delta Hedging with the underlyings (with a typical transaction cost of a liquid asset: 0.01% notional per delta trade).

Gamma Hedging with a vanilla call for the first hedge and then with vanilla call spreads: for S_1 we trade a call strike "the spot value of S_2 at hedging time" MINUS the existing call hedge, both expiring on the same date as the exotic (with a typical transaction cost of 0.5 vega).

The numerical results are listed below:

Electronic copy available at: https://ssrn.com/abstract=3502251

Scenario	1	2
price	4.77	4.77
Initial delta	TRUE	TRUE
Initial gamma	FALSE	FALSE
Initial gammaratio1		
Initial gammaratio2		
Rehedge delta	TRUE	FALSE
Rehedge gamma	FALSE	FALSE
Book PL wo hedge cost	-0.02	-0.04
Book PLSTDEV wo hedge cost	0.46	3.54
Deltahedgecost	-0.12	-0.02
Gammahedgecost	0.00	0.00
Total seller PL	-0.10	0.03
PL/EXOPRICE	-2%	1%
STDEV/EXOPRICE	10%	74%
Totalhedgecost/price	-2.45%	-0.35%

Table 3: Vanilla Call

Scenario	1	2	3	4	5	6	7
Price	6.74	6.74	6.74	6.74	6.74	6.74	6.74
Initial delta	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE
Initial gamma	TRUE	TRUE	TRUE	TRUE	FALSE	FALSE	FALSE
Initial gammaratio1	71%	71%	50%	25%			
Initial gammaratio2	71%	71%	50%	25%			
Rehedge delta	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	FALSE
Rehedge gamma	TRUE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
Book PL wo hedge cost	0.12	-0.02	0.02	-0.02	-0.05	0.39	-0.08
Book PLSTDEV wo hedge cost	4.04	0.54	0.54	0.55	0.62	4.81	9.96
Deltahedgecost	-0.37	-0.09	-0.09	-0.10	-0.12	-0.02	0.00
Gammahedgecost	-1.51	-0.08	-0.06	-0.03	0.00	0.00	0.00
Total seller PL	-2.00	-0.16	-0.17	-0.11	-0.07	0.37	-0.08
PL/EXOPRICE	-30%	-2%	-3%	-2%	-1%	5%	-1%
STDEV/EXOPRICE	60%	8%	8%	8%	9%	71%	148%
Totalhedgecost/price	-27.91%	-2.57%	-2.23%	-1.93%	-1.72%	-0.25%	0.00%

Table 4: Spread Option

Key results:

1. The $P\&L$ variance of the strategies with initial gamma-hedge, no subsequent gamma rehedge and daily delta rehedge (scenario 2,3,4,5 Table 4) is of the same order of magnitude as a for a plain vanilla option with daily delta rehedge (scenario 1 Table3).

2. Gamma rehedging actually results in a larger $P\&L$ variance (and obviously a higher transaction hedge cost) than daily delta hedging alone.

3. The transaction costs of gamma re-hedging kill any possible margin on the product.

The above Spread Option is therefore a reasonable product to trade as long as we hedge daily delta only (and possibly the initial gamma). The graphs below show that the product gamma and cross-gamma is indeed limited in all directions:

Electronic copy available at: https://ssrn.com/abstract=3502251

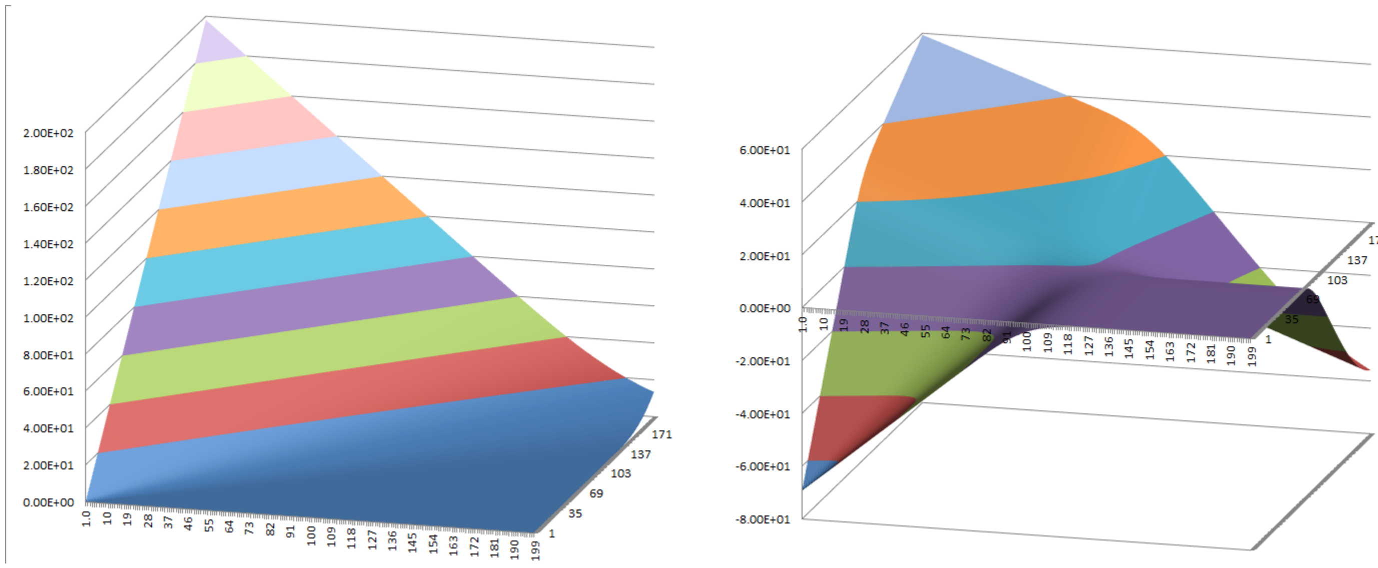

Figure 3: Price surface of a Spread Option as a function of the two underlyings (left), and the same with initial gamma hedges, no delta hedges, 3 months to maturity (right).

In the above examples we applied a generic method that can be used to decide if a product is worth trading or not: we specify the hedging strategy (instruments + hedging rules) and run forward Monte-Carlo simulations to compute the $P\&L$ expected value and standard deviation for the *portfolio* made of the exotic product and its hedging instruments. It is effective but requires some significant computer resources (it cannot be performed on a "live" basis). In the next section we give an alternative tool to measure cross-gamma risk and decide if it is sustainable or not for a given product.

5 Some metrics for cross-gamma risk

In the previous section we looked at the price surfaces of different products and saw that the impact of cross-gamma on $P\&L$ variance is important at a point $M(S_1, S_2, t)$ where there is a large "amount of curvature". For a random path from inception to expiry, we are therefore trying to evaluate the "total amount of curvature for all the points on the trajectory". The average of the above will give us a proxy for cross-gamma risk. Now, following [1] let us consider for a moment a one-dimensional product and a time period with daily subdivision $[t_{n-1}, t_n] = [t_{n,0}, ..., t_{n,M}]$ (M days) where dollar gamma ($S^2 \frac{\partial^2 P}{\partial S^2}$) is considered to be constant, and express the $P\&L$ over the period by:

$$P\&L_n = -\frac{S_0^2}{2}\frac{\partial^2 P}{\partial S_0^2}(t_0, S_0)\sum_m(\sigma_m^2 z_m^2 - \hat{\sigma}^2)\delta t$$

where σ_m is realized volatility for day m and z_m is centered with unit variance. Note that the z_m are independent but the σ_m may be correlated. The Standard Deviation is then:

$$StDev(P\&L)_n = \left|\frac{S_0^2}{2}\frac{\partial^2 P}{\partial S_0^2}(t_0, S_0)\right| StDev(\sum_m(\sigma_n^2 z_n^2 - \hat{\sigma}^2)\delta t)$$

so it is equal to the absolute value of dollar gamma times an expression which is independent of the product payoff. Now back to the 2 dimensional case, the expression is similar and given by:

$$StDev(P\&L)_n = \sqrt{\frac{1}{4}\sum_{i,j,k,l}|\Phi_{i,j}\Phi_{k,l}|\sum_m \mathbb{E}\left((\sigma_i^2 z_i^2 \sigma_j^2 z_j^2 - C_{ij}\delta t)(\sigma_k^2 z_k^2 \sigma_l^2 z_l^2 - C_{kl}\delta t)\right)_{|t=t_{n,m}}} \quad (7)$$

Next we use the inequality (at time t):

$$\left|\frac{\Phi_{i,j}(a,b)}{S_i S_j}\right| < \max(MaxCurvature(a,b), |MinCurvature(a,b)|) \tag{8}$$

We call the last expression **Maximum Absolute Curvature (MAC)** at (a,b,t) (it gives the maximum absolute value for gamma along all vectors of the tangent space at a point of the price surface). Recall from differential geometry on a parametric surface in $\mathbb{R}^3$ that $MaxCurvature$ and $MinCurvature$ (also called *principal* curvatures), which we rename K_1 and K_2 for short, are found using:

$$K_1 = H + \sqrt{H^2 - K}$$

and

$$K_2 = H - \sqrt{H^2 - K}$$

where K (the *Gaussian* curvature) and H (the *mean* curvature) are given by

$$K = \frac{P_{xx}P_{yy} - P_{xy}^2}{(1 + P_x^2 + P_y^2)^2}$$

$$H = \frac{(1 + P_x^2)P_{yy} + (1 + P_y^2)P_{xx} - 2P_x P_y P_{xy}}{2(1 + P_x^2 + P_y^2)^{3/2}}$$

where

$$P_x = \frac{\partial P}{\partial S_1}, P_y = \frac{\partial P}{\partial S_2}, P_{xx} = \frac{\partial^2 P}{\partial S_1^2}, P_{yy} = \frac{\partial^2 P}{\partial S_2^2}, P_{xy} = \frac{\partial^2 P}{\partial S_1 \partial S_2}.$$

Injecting (8) in (7) we get (for the time period $[t_{n-1}, t_n]$) :

$$StDev(P\&L)_n \leq MAC(t_n) \times \sqrt{\frac{1}{4}\sum_{i,j,k,l} S_i S_j \sum_m \mathbb{E}\left((\sigma_i^2 z_i^2 \sigma_j^2 z_j^2 - C_{ij}\delta t)(\sigma_k^2 z_k^2 \sigma_l^2 z_l^2 - C_{kl}\delta t)\right)_{|t=t_{n,m}}} \tag{9}$$

and the term under the square root is independent of the product payoff.

We now proceed to characterize $P\&L$ variance in terms of this curvature. We choose a time discretization such that over each period the dollar gamma is constant and we can apply the above formula. In order to measure the amount of curvature met on average for all random paths $(S_1(t), S_2(t))_{t\in[0,T]}$ from inception to expiry we calculate a weighted sum of curvatures, call it the **Cumulated Curvature**:

$$CC = \int_{[0,T]\times\mathbb{R}_+{}^2} MAC(a,b,t)dP_{a,b,t}$$

with

$$dP_{a,b,t} = \mathbb{P}(a < S_1(t) < a + da, b < S_2(t) < b + db)_{[t,t+dt]}$$

It can be thought of as the expected amount of curvature through the life of the product. We then divide $[0,T]$ into N time periods $[0 = t_0, t_1], ..., [t_{N-1}, t_N = T]$, such that curvature is supposed to be constant on each time period:

$$CC = \sum_{n=1}^{N} \int_{[t_{n-1},t_n]\times\mathbb{R}_+{}^2} MAC(a,b,t_n)dP_{a,b,t}$$

Using a finite grid $[a_0, .., a_{Na}] \times [b_0, .., b_{Nb}]$ for the discretization of $\mathbb{R}_+^2$, we get the approximation:

$$CC = \sum_{n=1}^{N} \sum_{i=1}^{N_a} \sum_{j=1}^{N_b} MAC(a_i, b_j, t_n) \int_{[t_{n-1}, t_n]} dP_{a_i, b_j, t_n}$$

or

$$CC = \sum_{n=1}^{N} \sum_{i=1}^{N_a} \sum_{j=1}^{N_b} MAC(a_i, b_j, t_n) \int_{[t_{n-1}, t_n]} \mathbb{1}_{[a_{i-1}, a_i] \times [b_{j-1}, b_j]}(S_1(t), S_2(t)) d(S_1(t), S_2(t))$$

The integral expression above can be seen as the value of a double-range double-underlying range accrual (up to the discount factor), call it $RA(A_i, B_j, t_n)$, with $A_i = [a_{i-1}, a_i]$ and $B_j = [b_{j-1}, b_j]$. We get the final expression:

$$CC = \sum_{n=1}^{N} \sum_{i=1}^{N_a} \sum_{j=1}^{N_b} MAC(a_i, b_j, t_n) RA(A_i, B_j, t_n)$$

The $RA(A_i, B_j, t_n)$ term can be approximated by

$$RA(A_i, B_j, T_n) = \frac{1}{2}[DBO(A_i, B_j, t_{n-1}) + DBO(A_i, B_j, t_n)] \times \Delta T_n$$

where $\Delta T_n = t_n - t_{n-1}$ and $DBO(A_i, B_j, t)$ is the price of the *double range double binary option* that pays 1 if S_1 and S_2 are in A_i and B_j respectively at maturity t, 0 otherwise.

Suppose for a moment that the S_i follow independent processes. The double binary option can then be expressed as a product of binary options and the latter are approximated by a linear combination of vanilla calls/puts, which means that if we have a full call/put surface (as a function of strike and maturity) from market option prices, the calculation of the $RA(A_i, B_j, t_n)$ is instantaneous. Also,if $MAC(x, y, t)$ is "slowly monotonous" as a function of t, we might be able to factor out a single price surface ($MAC(x, y, 0)$ or $MAC(x, y, T)$ or their average for example) to compute a rough approximation of CC.

Now, in the general case, suppose we want to decide if a product is acceptable for trading in terms of cross-gamma risk, we first decide on a reference product (for example one that has a known (cross-)gamma risk) and calculate the cumulated curvature CC_{ref} for that product.Then we decide on a maximum acceptability ratio M such that the new product should have a cumulated curvature CC_{ref} satisfying:

$$\frac{CC}{CC_{ref}} \leq M$$

The above method has notable advantages:

- It's cheap: the calculation of the curvature involves only algebraic functions of first and second order derivatives of the price (deltas and gammas) so their value might already exist in the system (as part of daily risk analysis, VAR calculation or stress tests) and if not they should be easy to get (no major development required).
- It's fast: in the case of a european payoff the calculation of the cumulated curvature approximation relies on prices of range accruals or binary options (in the best case these can be approximated by algebraic combinations of option spreads), so the calculation requires only vanilla-style computation time. It follows from the previous points that the calculation time of CC is roughly equivalent to that of model VAR.

Going back to the case of the 3 month Worst-Of Binary Option described in Section 4, we can compare $P\&L$ variance and CC value through monthly period. Results (shown as percentages of

Electronic copy available at: https://ssrn.com/abstract=3502251

total value over 3 months) are given in the following table and show that the above metric is coherent (though CC is steeper as a function of time).

Period	0-1m	1m-2m	2m-3m
PLVariance(relative)	7%	39%	54%
CC(relative)	18%	29%	53%

Table 5:

The following diagram illustrates the evolution the MAC through time for the Worst-Of Binary Option:

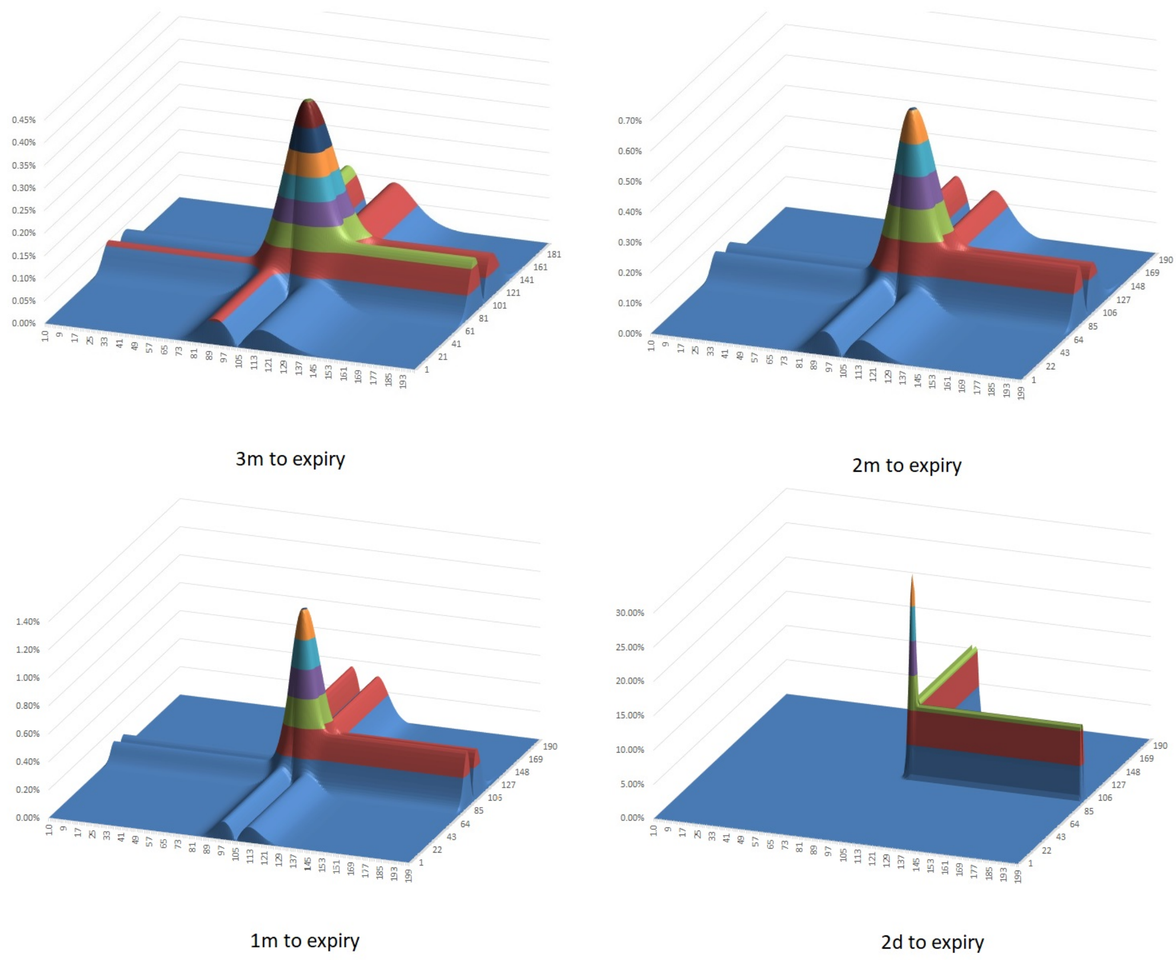

Figure 4: Evolution of the Maximum Absolute Curvature surface over time for a gamma hedged WO Binary option (note that the vertical axes each have different scales)

The above figure gives a clear view of the relative importance of cross-gamma compared to gamma along the axes. Now let's look at the Maximum Absolute Curvature weighted by the RA factors:

Electronic copy available at: https://ssrn.com/abstract=3502251

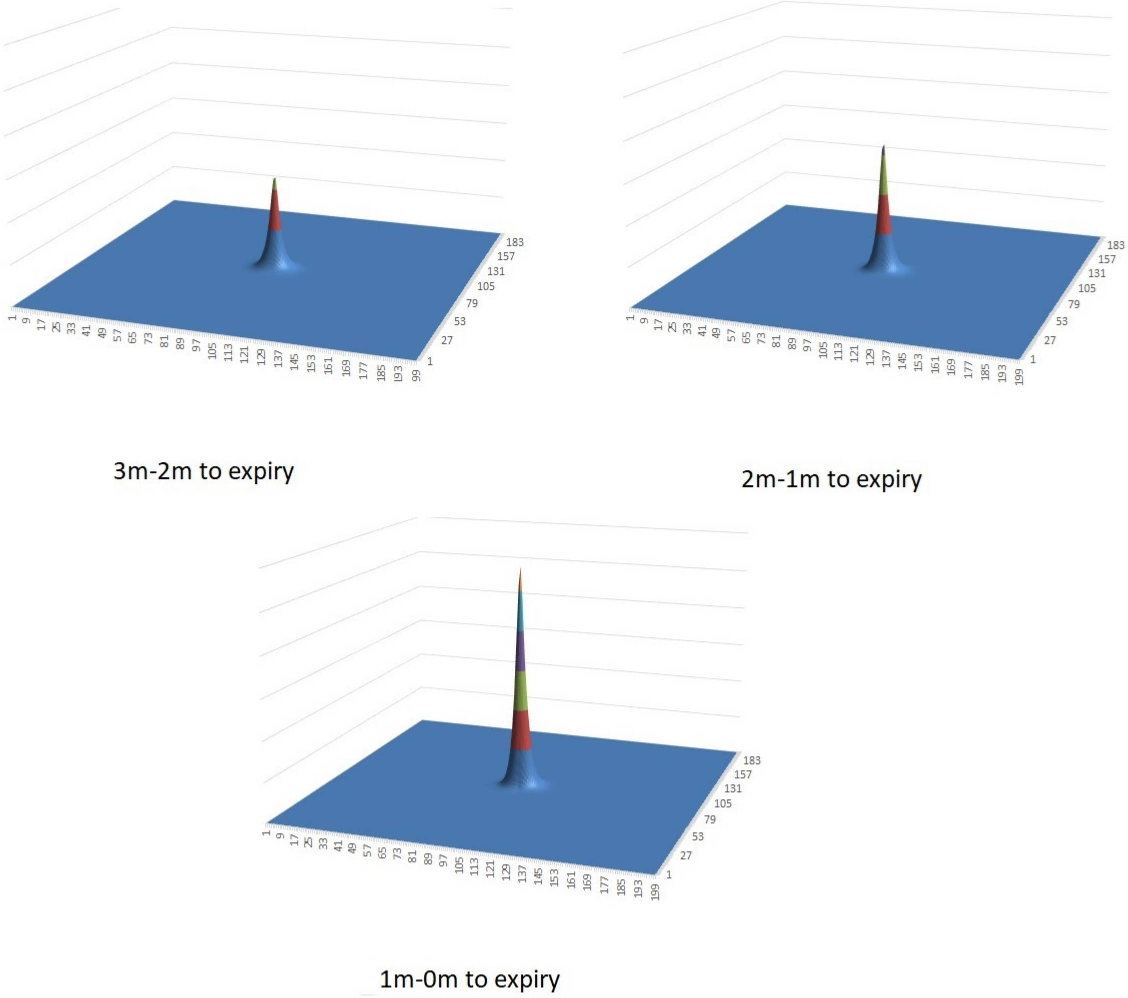

Figure 5: Evolution of the Weighted Maximum Absolute Curvature surface over time for a gamma hedged WO Binary option (same scale for vertical axes)

As expected, the Weighted Maximum Absolute Curvature is higher as we reach expiry, and reflects the higher $P\&L$ variance on that period.

Notes and further remarks:

1. The simple hypotheses on spot, drift, discount, correlation, diffusion model were chosen to simplify the exposure and concentrate on the variance due to the geometry of the payoff. The method is of course easily adaptable to more realistic sets of parameters (and models).

2. Note that we are not giving a formula for an approximation of the $P\&L$ variance itself but only defining an indicator with similar behaviour.

3. The method can be applied to other types of cross-risks besides cross-gamma. For example if we are interested in checking if the vanna risk of a given product is acceptable we can use S and σ as parameters (so the price is a parametric function $P(S, \sigma)$) and compare the cumulated curvature to that of a reference product. Note that σ could be stochastic or not.

4. Though the method is most effective for european payoffs, it could be adapted to payoffs with cancellable features although this would require a monte-carlo diffusion instead of closed formulas

Electronic copy available at: https://ssrn.com/abstract=3502251

to compute the "range accrual" weights so it might not be much faster than running the full simulation for the hedged portfolio to calculate the real standard deviation.

5. If we consider a strategy where delta and gamma are both hedged daily then $P_x = P_y = P_{xx} = P_{yy} = 0$ so the MAC is exactly $|\Phi_{1,2}(a, b)|$.

6. It is important to note that the above study is of geometrical nature AND parameter dependant: the same product might be a poor trade with one set of parameters but a reasonable one with another. For example if the correlation between the underlyings is set at 99% the WO Binary Call behaves more or less like a single Binary Call and the $P\&L$ variance can be reduced significantly.

7. The above might be adapted to dimensions higher than 2 since by definition the price is a parametric surface and we can define its principal curvatures.

8. Depending on computational resources and how "nice" the price surface looks, we could decide to use other metrics: for example instead of calculating the *sum* over time we could calculate the *max* over time, if the curvature is a "slowly monotonous" function of time this would simplify the calculation (at the expense of being more conservative on the risk metric). Another possible improvement would be the use of an adaptive mesh around points of high curvature.

6 Conclusion

There are many sources of delta variation for an exotic payoff, and as we have seen previously, as soon as cross-risks are involved, some products will show a lot of residual delta risk that cannot be hedged using liquid market products. The above metric gives a cheap and fast indicator of $P\&L$ variance due to cross-gamma risk from the point of view of the payoff. This leaves us with a lot of room for development. In particular we might be interested in applying the same method to non-european payoffs, to other types of cross-risks besides cross-gamma, and in higher dimensions to accommodate for multiple underlyings (an obvious but critical requirement being the ability to provide meaningful correlation parameters). There will be limitations in terms of computation time but then again the above method is not meant to be processed daily. Nonetheless it looks like a reasonably flexible tool to quantify the proportion of hedgeable vs non-hedgeable risk in an exotic product.

References

[1] Lorenzo Bergomi, *Stochastic Volatility Modelling*. Chapman & Hall CRC, 2016.

Electronic copy available at: https://ssrn.com/abstract=3502251