The Problem with Risk: Second Order Risk and How Standard Simplification Methods Fail

Sean Sonni
email Sean.Sonni@PLATSON.io

Richard Platt
email Richard.Platt@PLATSON.io

May 28, 2018

Abstract

Vanilla interest rate derivatives are normally given risk via Taylor Series expansions of the valuation function. This proves to be simple to estimate via first difference methods to first order, but computationally intensive for second order risk (gamma, vanna, volga). Standard methods to simplify this problem cause significant P&L noise, and impair risk management of portfolios. While the simplifications are well understood, the extent to which they impair risk estimation are not. Via a simplification of the yield curve specification, dynamics and volatility space, the severity of this error in benign market moves will be demonstrated and it is concluded that the most tractable solution is to "bite the bullet" and generate risk without simplifications. While this comes at a significant computational cost, it allows for order of magnitude improvement in hedging, tractability of traditionally "difficult" products, and unique insights into the risk of books.

1 Introduction

1.1 Risk Generation and Computation Time

Interest rate derivatives risk is generally computed via first difference approach, which brings significant overhead as it requires rebuilding the yield curve object and volatility surface object for every risk point calculated. A trade-off has to be made between computational time and achieving granularity and accuracy of the higher order risk.

If a portfolio, P, is valued on a yield curve with N inputs and a volatility surface with V inputs, first order risk (delta and vega) will require N curve rebuilds and $N+V$ vol surface rebuilds (assuming a skew model in which the implied volatility is a function of the forward rate). Assume that the time required to rebuild market data and revalue the portfolio for a bump of either a volatility surface point or a yield curve point (or any combination) is roughly constant, T.

The computation time, C_1, for first order risk is calculated as

$$C_1 = T(N+V) \tag{1}$$

For complete second order risk (gamma, vanna, volga), the computation time becomes

$$C_2 = T(N+V)^2 \tag{2}$$

and in general for i$^{\text{th}}$ order risk

$$C_i = T(N+V)^i \tag{3}$$

Inputting some reasonable estimates for N, V, and T shows why this is so expensive even at the second order. It is obviously not feasible to consider any-

N	V	T	C_1	C_2	C_3
130	1000	1 second	19 minutes	14 days	45 years

thing beyond the second order, and even the second order is so computationally intensive as to either require certain simplifications or significant hardware investment. The specifications for curves continues to compound in complexity (single curve, tenor basis curves, discounting curves, clearing house curves, CSA curves etc.) Volatility surfaces also continue to increase parameters to offer better pricing and hedging flexibility. While modern computer systems could handle full second order risk given market data from 10 years ago, it has become a figurative arms race to keep up with modern changes to models. (For fully specified CSA risk,the low-ball estimate of 130 curve points above would be significantly under-specified).

1.2 Common Simplifications: A Mathematical Description

The most common approach to solving the computation time issue for higher order risk is to ignore volga and vanna and simplify gamma by assuming the yield curve moves only in parallel. Approximations for volga and vanna can be given via scenario analysis. However this leads to the question of which bumped volatility surface and yield curve most accurately reflect the move on any given day.

A curve builder, B is a function that takes a series of yields (deposits, futures, swaps, forward swaps), Y_i, and builds a yield curve. From this yield curve a function is applied to return discount factors for any tenor:

$$\delta(t) = B(Y_i) = B\left(\vec{Y}\right) \tag{4}$$

Any swap rate,S, can be extracted from the function $\delta(t)$

$$S = f\left(\vec{Y}\right) \tag{5}$$

Notation for derivatives of this function is defined:

$$\frac{\partial S}{\partial Y_i} = S_i \tag{6}$$

And to higher order

$$\frac{\partial^2 S}{\partial Y_i \partial Y_j} = S_{ij} \tag{7}$$

The Taylor series expansion to second order for a derivative, $V = f(Y_i)$, noting that the change in yields to the curve are given by dr^i

$$dV \approx V_i dr^i + \frac{1}{2} V_{ij} dr^i dr^j \tag{8}$$

Calculation of the points, V_{ij}, is computationally intensive so the problem is simplified by defining a dynamic to the curve,

$$\forall \quad i, j \quad dr^i = dr^j \tag{9}$$

which leads to

$$V_{ii} = \sum_j V_{ij} \tag{10}$$

$$dV \approx V_i dr^i + \frac{1}{2} V_{ij} dr^i dr^j = V_i dr^i + \frac{1}{2} V_{ii} dr^i dr^i \tag{11}$$

A parallel shift assumption may suffice in markets where the term structure is primarily determined by a single driver but significant errors in risk bucketing are introduced when this assumption breaks down.

The terms, V_{ii}, are easily computed by parallel shifting all yield curves by 1 basis point and then generating a new delta ladder via first difference methods. For a curve with N points, this simplification means the computation time for gamma is reduced from TN^2 to TN, an ever more significant gain as curves have become more complex. Vanna and volga are usually excluded, reducing total computation time for second order risk by several orders of magnitude. This is then supplemented by generating several gamma ladders, to try and capture the non-linearity of gamma via a grid, hoping that there is good offset between the loss of accuracy locally (relevant on most days) and a slight gain globally (on the rare high volatility days). Unfortunately this approach still fundamentally relies on the assumption that the high volatility days are also parallel in nature, but given computational limitations it seems to be the best currently available solution.

2 The Breakdown of Current Second Order Risk

2.1 Simplified Example

For products or portfolios of significant complexity, the above assumptions can prove to be inadequate. A portfolio hedged using the parallel assumptions can be seemingly well balanced but produce unexpected risks and valuation changes when the curve moves in a non-parallel fashion. For example, a simple midcurve, hedged with spot starting European swaptions carries second order risk that looks very similar (locally) to a spread option. The spot starting European swaptions hedge the parallel shift gamma quite well (especially when valued on a relatively high term correlation), but fail to capture any of the curve gamma embedded in a midcurve option.

New terminology is now introduced - Traditional Gamma(Γ_p) is defined as:

$$\Gamma_p = \frac{1}{2} V_{ii} dr^i dr^i \tag{12}$$

And complete Gamma (Γ) is

$$\Gamma_f = \frac{1}{2} V_{ij} dr^i dr^j \tag{13}$$

To highlight the implications of parallel shift gamma and how it can fail, a simple example is demonstrated. A market consisting of only 3 assets, A, B, and C[1] where:

$$\begin{aligned}
di &= \sigma_i dW_i \quad i = A, B && (14)\\
dW_A dW_B &= \rho && (15)\\
VAR(A) &= \sigma_A^2 && (16)\\
VAR(B) &= \sigma_B^2 && (17)\\
C &= 2B - A && (18)\\
VAR(C) &= \sigma_C^2 = \sigma_A^2 + 4\sigma_B^2 - 4\rho\sigma_A\sigma_B && (19)
\end{aligned}$$

This model market is further simplified by assuming the implied volatility of A and B are fixed. A 1 month option on asset C is examined to determine the percentage of the non-linear P&L of the option on C that is unexplained if gamma is expressed purely in terms of the primary assets A and B (i.e. using the assumption of a parallel shift). 1 month optionality is used as gamma risk stability is inversely proportional to expiry time, hence traditional gamma ladders should be able to explain a large proportion of the P&L. 2000 daily simulations were run and the total unexplained P&L under Traditional Gamma (Γ_p) and the percent of this error explained by removing the parallel shift assumption were calculated. The majority of the P&L error was found to be remediated by eliminating the parallel shift assumption, and it is thus concluded that extending Traditional Gamma with ladders is not a valuable metric to estimate P&L.

Changing the hedging parameters and values of ρ, σ_A, and σ_B were further investigated (the simulated data is generated using the same implied volatility

[1] The structure specified can be best thought of as replicating a 1y1y swaption, where A and B are the 1 year and 2 year rate, respectively and C is the 1y1y forward rate

as pricing so the ratio $\frac{\sigma_B}{\sigma_A}$ is the main driver). Data on predicted first order risk changes due to second order risk[2] was deemed superfluous, as if the P&L prediction of the book was lacking this is deemed due to inadequate predictions for changes in first order risk.

2.2 Test Structure and Results

Two calculations are performed.

$$E_p = \frac{\sum_i ABS\left(dV^i - \Gamma_p^i\right)}{\sum_i ABS\left(dV^i\right)} \tag{20}$$

$$E_f = \frac{\sum_i ABS\left(dV^i - \Gamma_f^i\right)}{\sum_i ABS\left(dV^i\right)} \tag{21}$$

$$W = 1 - \frac{E_f}{E_p} \tag{22}$$

where i is an index over each simulation. This metric gives more weight to days with significant absolute non-linear P&L.
Simulations are run across various volatility ratios and correlations for A and B, focusing on the level of failure for traditional methods, E_p, and how much of this error is due specifically to the parallel shift assumption, W. To correct for issues regarding the usefulness of a full expansion, the simulated summary data for only the most extreme 2% of days is also returned, extreme being defined as the days with the largest magnitude non-linear P&L[3].

Traditional Gamma Error as a percent of nonlinear PnL								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	88%	6%	7%	14%	20%	28%	32%
	90%	59%	10%	11%	17%	22%	29%	33%
	85%	53%	12%	13%	19%	24%	30%	34%
	80%	47%	14%	15%	20%	25%	31%	35%
	75%	52%	17%	16%	21%	25%	31%	35%
	70%	51%	19%	18%	22%	26%	32%	35%
	65%	57%	20%	19%	23%	26%	32%	36%
	60%	61%	23%	21%	24%	27%	32%	36%
	55%	62%	25%	22%	25%	28%	33%	36%
	50%	66%	29%	23%	25%	28%	32%	36%

Figure 1: The error in P&L prediction using Traditional Parallel Shift Gamma

[2] i.e. Delta change due to gamma or vanna
[3] The extreme 2% of cases shows the value of a full expansion in explaining risk when markets are most volatile, which are the most important days to have correct risk and P&L.

Percent of Traditional Gamma Error Explained							
	Vol Ratio, σ^2/σ^1						
Rho	0.6	0.8	1	1.25	1.5	2	2.5
95%	100%	94%	95%	98%	98%	99%	99%
90%	100%	96%	97%	98%	98%	99%	99%
85%	100%	97%	97%	98%	99%	99%	99%
80%	100%	98%	98%	98%	99%	99%	99%
75%	100%	98%	98%	99%	99%	99%	99%
70%	100%	99%	98%	99%	99%	99%	99%
65%	100%	99%	99%	99%	99%	99%	99%
60%	100%	99%	99%	99%	99%	99%	99%
55%	100%	99%	99%	99%	99%	99%	99%
50%	100%	99%	99%	99%	99%	99%	99%

Figure 2: The percent of P&L prediction error explained by using the full Gamma

Traditional Gamma 2% Error Rate							
	Vol Ratio, σ^2/σ^1						
Rho	0.6	0.8	1	1.25	1.5	2	2.5
95%	36%	1%	5%	14%	20%	27%	32%
90%	11%	1%	7%	16%	22%	29%	33%
85%	9%	2%	9%	17%	23%	30%	34%
80%	20%	2%	10%	18%	23%	30%	34%
75%	31%	4%	12%	19%	24%	30%	34%
70%	37%	5%	12%	19%	24%	31%	35%
65%	34%	14%	13%	20%	24%	31%	35%
60%	38%	25%	13%	21%	25%	31%	35%
55%	40%	34%	17%	20%	25%	31%	35%
50%	39%	38%	26%	21%	26%	31%	35%

Figure 3: The error in Profit prediction using Traditional Parallel Shift Gamma

There are several important points to take away from this analysis. Firstly, the volatility surface shape is a more significant driver of error than the correlation between points (see Figure 1). Extremely flat ($sigma_i \approx sigma_j$) volatility surfaces (if the implied volatility is a good measure of the realized volatility) will generate dynamics that approximate a parallel shift. In fact, a reduction in correlation from 99% to 54% is roughly equal to changing the volatility ratio from 1 to 2. Secondly, almost all of the unexplained P&L is captured by eliminating the parallel shift assumption (Figure 2) and this also holds true when considering only the extreme days(Figure 4).

The implication is that for the majority of market structures, traditional gamma

Percent of Traditional Gamma Error Explained, Most Extreme 2%								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	98%	-15%	86%	95%	97%	98%	98%
	90%	96%	-26%	91%	96%	97%	98%	98%
	85%	95%	67%	93%	96%	97%	98%	98%
	80%	98%	77%	94%	96%	97%	98%	98%
	75%	99%	87%	95%	97%	97%	98%	98%
	70%	99%	92%	95%	97%	97%	98%	98%
	65%	99%	97%	96%	97%	98%	98%	98%
	60%	99%	99%	97%	98%	98%	98%	98%
	55%	99%	99%	98%	98%	98%	98%	98%
	50%	99%	99%	98%	98%	98%	98%	98%

Figure 4: The error in Profit prediction using Traditional Parallel Shift Gamma

ladders are found to be wholly inadequate in minimising the largest P&L prediction errors. Traditional gamma ladders may show significant value when dealing with options expiring over very short horizons (¡1 week) and for options whose risk is driven by a single yield curve point (ie. not midcurves, spread options, or European swaptions across a curve built from futures), but are severely limited in quantifying the dynamics of most risk in a trading book.

3 Results for a Hedged Portfolio

Previously, a single trade was analyzed, specifically with the goal of clarifying the severity of the problem and the only appropriate solution. The next level of complexity is then addressed via a modestly hedged portfolio. Specifically, a portfolio is built whereby,

$$N_i = -\frac{dC}{di} \tag{23}$$

$$i = A, B, C \tag{24}$$

with the purpose of analyzing the performance of the parallel shift assumption when reasonable market hedges are put in place. The hedges are only completely natural when $\rho \approx 100\%$, but these can be understood as the set of hedges that generate a dispersion portfolio. More importantly, for well correlated assets, this is a natural first order hedge when the volatility surface is relatively flat. The key results highlight:

- Figure 5 highlights how severe the error regularly is, almost never less than 100% of the actual PnL.
- Figures 6 and 8 almost all the error could be captured by quantifying the full gamma, and that gamma ladders are of little value in risk attribution.

- In steep volatility curves, the error becomes so egregious both in general, and at the 2% level, that traditional gamma metrics become almost meaningless at quantifying portfolio performance.
- Effective risk management for non-linear portfolios should either include full revaluations or eliminate the parallel shift assumption.

Traditional Gamma Error as a percent of nonlinear PnL								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	225%	54%	108%	357%	697%	1544%	2412%
	90%	117%	49%	98%	250%	440%	945%	1510%
	85%	79%	46%	93%	204%	351%	695%	1073%
	80%	67%	46%	84%	168%	285%	555%	854%
	75%	61%	48%	78%	153%	242%	471%	727%
	70%	56%	42%	70%	135%	226%	409%	617%
	65%	56%	44%	68%	122%	186%	370%	538%
	60%	52%	44%	64%	110%	170%	318%	501%
	55%	53%	43%	64%	100%	161%	296%	444%
	50%	52%	45%	60%	101%	148%	257%	390%

Figure 5: The error in P&L prediction using Traditional Parallel Shift Gamma

Percent of Traditional Gamma Error Explained								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	100%	99%	99%	100%	100%	100%	100%
	90%	99%	98%	99%	100%	100%	100%	100%
	85%	99%	98%	99%	100%	100%	100%	100%
	80%	99%	98%	99%	99%	100%	100%	100%
	75%	99%	98%	99%	99%	100%	100%	100%
	70%	99%	98%	99%	99%	100%	100%	100%
	65%	99%	98%	99%	99%	100%	100%	100%
	60%	99%	98%	99%	99%	99%	100%	100%
	55%	99%	98%	99%	99%	99%	100%	100%
	50%	99%	98%	99%	99%	99%	100%	100%

Figure 6: The percent of P&L prediction error explained by using the full Gamma

Traditional Gamma 2% Error Rate								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	150%	86%	119%	163%	210%	317%	395%
	90%	44%	11%	70%	164%	211%	329%	387%
	85%	19%	4%	49%	155%	277%	421%	649%
	80%	12%	5%	39%	125%	217%	453%	746%
	75%	10%	6%	35%	103%	185%	378%	612%
	70%	7%	5%	29%	86%	153%	316%	504%
	65%	6%	5%	28%	70%	132%	268%	418%
	60%	5%	5%	23%	65%	116%	232%	365%
	55%	5%	5%	22%	58%	103%	208%	326%
	50%	25%	4%	20%	51%	91%	180%	281%

Figure 7: The error in Profit prediction using Traditional Parallel Shift Gamma

Percent of Traditional Gamma Error Explained, Most Extreme 2%								
		Vol Ratio, σ^2/σ^1						
		0.6	0.8	1	1.25	1.5	2	2.5
Rho	95%	99%	100%	100%	100%	100%	100%	100%
	90%	97%	88%	98%	100%	100%	100%	100%
	85%	92%	61%	97%	99%	100%	100%	100%
	80%	88%	68%	96%	99%	99%	100%	100%
	75%	85%	71%	95%	98%	99%	100%	100%
	70%	80%	70%	94%	98%	99%	99%	100%
	65%	78%	70%	94%	98%	99%	99%	100%
	60%	72%	70%	93%	97%	99%	99%	100%
	55%	74%	68%	93%	97%	98%	99%	99%
	50%	96%	66%	93%	97%	98%	99%	99%

Figure 8: The error in Profit prediction using Traditional Parallel Shift Gamma

4 Conclusion

A simplified set of examples has shown the pitfalls of assuming a parallel shift when calculating second order risk. This is especially marked if the volatility surface is steep, or if used as an estimate of extreme outcomes (ie. VaR or Expected Shortfall) and finally, if the portfolio has reasonable market hedges applied. The performance in an absolute sense worsens when correlations break down, as long as the vol ratio is above a critical value, which is determined by the relationship between A,B, and C. However, the sensitivity to correlation is not universal, and in some well hedged portfolios, poor correlation leads to a lower relative error. Most importantly when calculating Value-at-Risk (and similarly derived Expected Shortfall calculations), in a portfolio that is modestly well hedged, it is crucial to remove the parallel shift assumption, as, this assumption, combined with the hedged nature of the book leads to severe errors in P&L prediction.