Regression Sensitivities for Initial Margin Calculations

Claudio Albanese[1, 2], Simone Caenazzo[1, 2] & Oliver Frankel[1, 3]

This version: June 14, 2016

Abstract

Implementations of the Standard Initial Margin Model (SIMM) [5] and the Sensitivity Based Approach (SBA) in the Fundamental Review of the Trading Book (FRTB), both call for the calculation of sensitivities with respect to a standardised set of risk factors. Since standard factors are generally collinear and pricing functions are possibly rough, finding sensitivities qualifies as a mathematically ill-posed problem for which analytical derivatives do not provide a robust solution. Numerical instabilities are particularly problematic since they hamper reconciliation and make collateral optimisation strategies inefficient.

In this article, we introduce a method for calculating sensitivities based on ridge regressions to keep sensitivities small and stable. We find that a drift term and FX cross-gammas significantly improves the accuracy of the P&L explain achieved in the SIMM methodology. The method implies rigorous upper bounds on errors in P&L explain, on which basis we adjust Initial Margin conservatively in order to pass back-testing benchmarks.

Satisfying the IM requirements hinges upon the ability of reconciling and agreeing the amounts being called. While the BCBS261 gave lip service to IM reconciliation, the regulators produced rules without those requirements. However, they fully understand that without counterparts agreeing the IM amount, with only the undisputed amount being posted, a counterpart could be under-protected. This is particularly important in a time of stress, not just because that is when the IM is more likely to have to be used, but also because that's when counterparts' portfolio sensitivities, on which SIMM (and all practical IM models for non-cleared derivatives) depend, are most likely to materially differ.

The reason for this is embedded not so much in the counterparts' different pricing models but in their market models, that is, the assumptions they make on how market levels are interrelated. A clear example of this is the assumption a firm makes on how market volatility will change with its market level. While not a large contributor to sensitivity differences in times of low and stable volatility, when volatility jumps to stressed levels, large differences in sensitivity calculations are bound to emerge. And correspondingly, just at such times, using the practice to post just the undisputed amount will systematically lead to bilateral under-margining with respect to the BCBS261 requirements.

The most straightforward model for computing sensitivities is to calculate them at trade level, possibly using either closed form solutions, or Adjoint Algorithmic Differentia-

[1] *IMEX*, `http://www.imex-global.net`
[2] *Global Valuation*, `http://www.global-valuation.com`
[3] *Bilateral Risk Management, New York*

tion (AAD) [3] or small bumps. Sensitivities are then mapped into the SIMM risk factors and aggregated linearly. See the SIMM technical model for a discussion [5].

Mapping to SIMM factors involves degrees of freedom which potentially jeopardise reconciliation among market participants. As an example, if two banks agree on the payoff of a 5-into-10 year swaption and are both using SABR, there is no guarantee that they will agree on the sensitivity to the 7 year rate since interest rate deltas depend on volatility.

Another problem is that analytical sensitivities are sometimes singular and numerically very large, as for instance happens when the spot is near the strike of an option and the maturity is very short. This is not particularly difficult when considering trades individually. However, when considering large portfolios, these exceptions occur with large probability and are able to spoil the final results if not properly regularised, thus adding further noise to the reconciliation process. Sensitivities computed by bumping are more robust in this respect, as long as the bumps are sufficiently large. But what is the right scale for bumps? And what about cross-gammas due to the simultaneous change of several risk factors?

A third possible issue is that the accuracy of analytical or finite difference sensitivities is difficult to assess, especially when they are used to find the impact of fairly extreme shocks that populate the 99% quantile of the P&L distribution. It would be useful to estimate errors, but how can one do so if power series expansions do not converge in general?

Last but not least, gamma sensitivities to SIMM factors are not straightforward to find analytically. The SIMM document calls for using only diagonal gammas as the cross-gammas are generally considered to be too difficult to find. One solution is to find the diagonal gammas by finite bumping. Another solution outlined in the SIMM document is to imply gamma sensitivities from vegas.

Inaccuracies in the gamma can be problematic in the case of hedged portfolios which are delta-neutral by construction. In this case, higher order non-linearities and cross-gammas as a rule dominate over the diagonal gammas in delta-hedged portfolios.

To calculate sensitivities robustly and with verifiable accuracy, we propose a methodology based on "regression sensitivities". The idea is to optimise directly the quality of P&L explain for a given set of SIMM standard sensitivities and assess errors in order to infer upper bounds. In its most primitive form, the method is based on the solution of a regression equation of the form:

$$(P_s - P_0) = \sum_i \Delta_i(\mathrm{RF}^i_s - \mathrm{RF}^i_0) + \frac{1}{2}\Gamma_i(\mathrm{RF}^i_s - \mathrm{RF}^i_0)^2 + \epsilon_s \tag{1}$$

Here, s is an index for a scenario obtained in a risk neutral simulation, P_0 is the portfolio spot valuation, P_s is the portfolio valuation in two weeks time under scenario s, RF^i_0 is the spot valuation of the i-th SIMM factor and RF^i_s is the valuation of SIMM factors in two weeks time under scenario s. We interpret this as a linear system of equations, one for each scenario s, whereby the sensitivities Δ_i and Γ_i are treated as the unknowns. The system is solved in the least square sense by seeking to minimise the sum of squares of residuals, i.e.

$$\text{minimize} \sum_s \epsilon_s^2 + \lambda \sum_i (\Delta_i\sigma_i)^2 + \left(\Gamma_i\sigma_i^2\right)^2 \tag{2}$$

In this equation, σ_i is the volatility of the $i-$th factor and λ is the Tikhonov regularisation parameters, see [6]. By choosing λ slightly positive, we ensure that sensitivities are as small as they can be notwithstanding collinearities between risk factors and without spoiling much the quality of P&L explain. This regularisation is also called method of "ridge sensitivities" in statistics, see [4].

Robust and small sensitivities are particularly useful for optimisation purposes.

Based on our experience, we recommend using at least 100,000 primary risk neutral scenarios to carry out a calculation of regression sensitivities, followed by an additional 5,000 secondary scenarios branching off at the 2-weeks time horizon from each primary scenario to find future valuations of exotic derivatives.

The final objective of this calculation is to evaluate 99% VaR by applying historical shocks to the SIMM risk factors and using sensitivities. An alternative way to go about doing this calculation would be to map the shocks of the SIMM factors into shocks of the calibration inputs for all pricing models that are used and then carry out a full revaluation of the portfolio. However, the mapping of SIMM factor shocks is a difficult and error prone procedure. For our calculations, we use the same risk system for XVA analytics also used to write Risk articles such as [1] and [2]. In this setup, the model recalibration step would become a numerical bottleneck leading to an unacceptable level of performance degradation. Instead, risk neutral simulations under a global calibration are much faster to compute, even in case the number of risk neutral scenarios is of the order of a billion. Using analytical or discrete sensitivities would be faster to implement but have the problems we discussed earlier. Regression sensitivities allow one to bypass the problem of specifying a scheme for SIMM factor mapping and arrive to the best possible degree of P&L explain for a given regression model. As a second step, one can then apply historical shocks while controlling approximation errors, insofar as one uses pricing models of sufficiently high quality to have calibrated parameters that are fairly stable across time.

An approximate upper bound U on errors for initial margin requested can be estimated based on the distribution of the residuals ϵ_s. We use the following definition:

$$U_+ = \min\{\xi : \mathrm{Prob}(\epsilon < \xi \mid P - P_0 > \mathrm{VaR}_+(95\%)) \geq 0.99\} \tag{3}$$

and

$$\mathrm{VaR}_+(95\%) = \min\{\xi : \mathrm{Prob}(P - P_0 < \xi) \geq 0.95\} \tag{4}$$

where ϵ and P are random variables distributed as the residuals ϵ_s and the valuations P_s, respectively. This definition of upper bound probes the upper 95% quantile of the return distribution. Similarly, for initial margin received on can define the upper bounds U_- and VaR_- by flipping the sign of portfolio returns.

Upper bounds on residuals are useful as they allow one to arrive to conservative estimates for IM, even when the calculation is carried out using historical shocks as required by SIMM in order to be consistent with back-testing benchmarks. Figures 1-4 show four instances of how upper bounds to IM deviate from the rigorous value for IM when both are computed under the risk-neutral measure.

Figure 1 compares the exact IM with an upper bound estimate obtained using the regression model in equation 1 for a sample of about 2,000 fixed income derivative portfolios with a total of about 100,000 trades. The difference between the two graphs is that the left hand one was obtained using the original portfolio while the right hand graph was obtained by adding hedge trades to make the delta neutral.

Figure 3 uses the same regression model and portfolio as in Figure 1 except that the portfolios are assumed to be delta-neutral. This is an important case as delta hedging is a risk-reducing strategy to be expected and regulators place particular emphasises on the validation of IM models under delta-neutrality conditions. Hedged delta-neutral portfolios are far harder to approximate with a sensitivity based model since they are dominated by higher order non-linearities. This is particularly problematic when gamma terms are obtained from vegas. Using regression sensitivities instead, real gammas yield better fits.

Once we depart from analytical sensitivities, we can easily build regression models that go beyond power series expansions. A more elaborate but also far more accurate regression model is the following:

$$(P_s - P_0) = \alpha + \sum_i \Delta_i (X_s^i \mathrm{RF}_s^i - X_0^i \mathrm{RF}_0^i) + \frac{1}{2}\Gamma_i (X_s^i \mathrm{RF}_s^i - X_0^i \mathrm{RF}_0^i)^2 + \epsilon_s \tag{5}$$

Here, X_s^i is the exchange rate between the reference currency and the currency of the $i-$th SIMM risk factor in two weeks time in the scenario s. The difference between the regression model in 1 and 5 consists in the drift term α which is also optimized and in the FX-cross gammas which are accounted for by sake of inserting the exchange rate scenarios X_s^i.

As Figures 2 and 4 show, the introduction of α and the FX-cross gammas lead to remarkable improvements in the quality of fit that conservative upper bounds provide. The SIMM technical document [5] already hints at the possibility to have a drift term, although it does not suggest a method to compute it analytically. Furthermore, the same document suggests that diagonal gammas could be derived from vegas and that cross-gammas should be neglected. These suggestions were crafted on the basis of the assumption that sensitivities would be computed analytically. Interestingly enough, the method of regression sensitivities we propose in this article easily allows us to go further and compute optimal drifts, rigorous diagonal gammas and FX cross-gammas.

References

[1] C. Albanese, L. Andersen, and S. Iabichino. FVA: Accounting and Risk Management. *Risk*, February 2015.

[2] C. Albanese, S. Caenazzo, and S. Crepey. Capital and Funding. *Risk*, May, 2016.

[3] M. Giles and P. Glasserman. Smoking Adjoints: Fast Monte Carlo Greeks. *Risk*, January, 2006.

[4] A.E. Hoerl and R.W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67, 1970.

[5] ISDA. SIMM Methodology. *Available at the ISDA website*, 2016.

[6] A.N. Tikhonov, A.S. Leonov, and Yagola A.G. *Nonlinear Ill-Posed Problems*. Chapman and Hall, 1998.

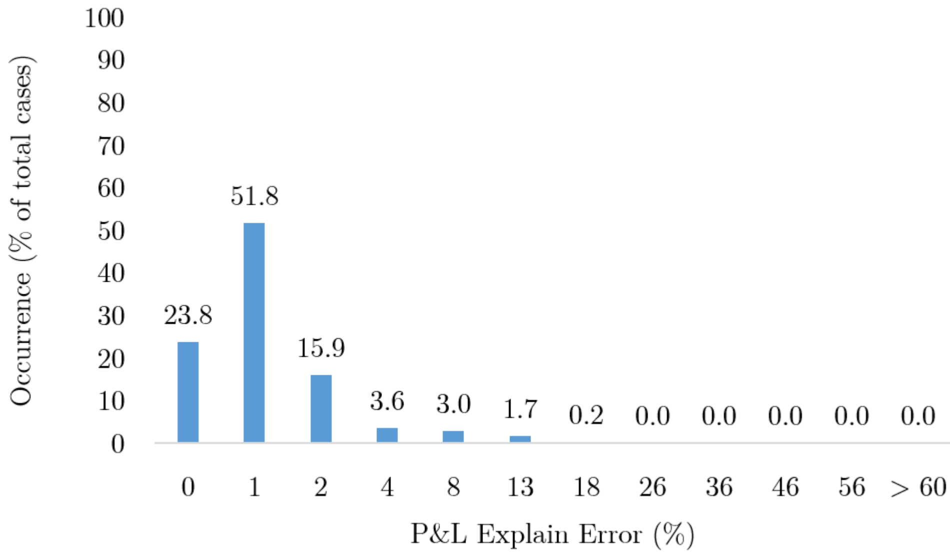

Figure 1: This graph refers to the original portfolios (which are typically delta-dominated), and shows the distribution of relative errors for IM in a risk neutral simulation. We estimate IM conservatively as upper bounds based on the regression model in Equation (1) which is the exact same as in the standard SIMM model.

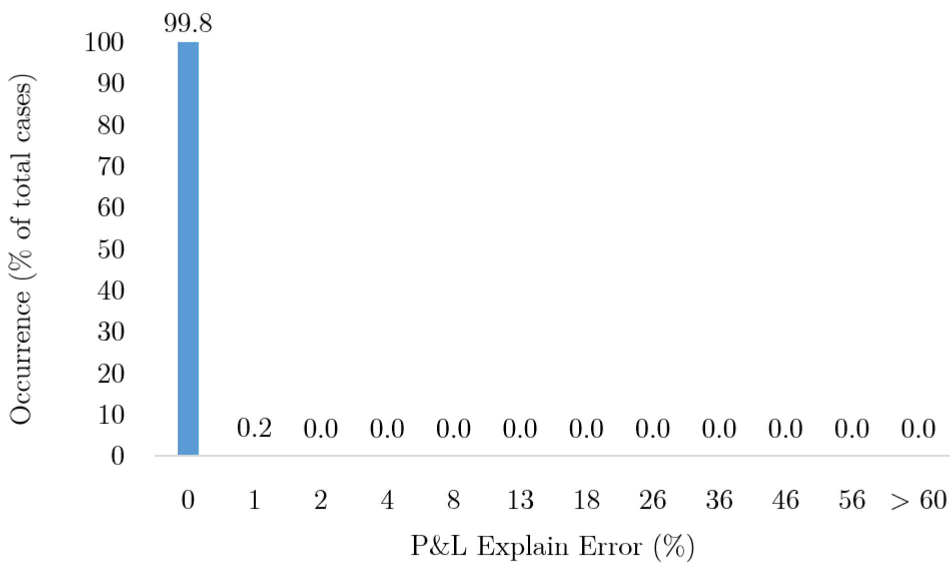

Figure 2: Same as Figure 1 except that the regression model accounts for a drift α and FX-cross gammas as in Equation (5).

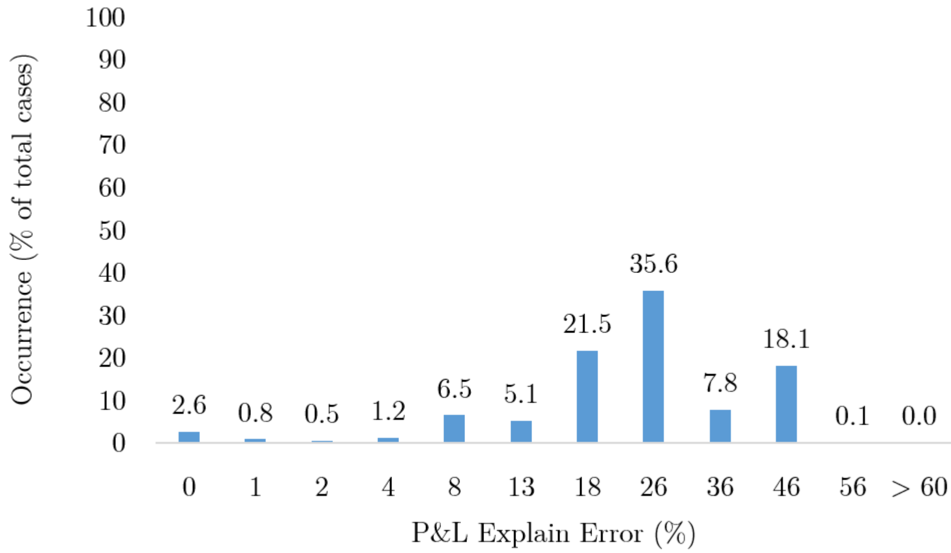

Figure 3: Same as Figure 1 except that our test portfolios are assumed to be delta-hedged.

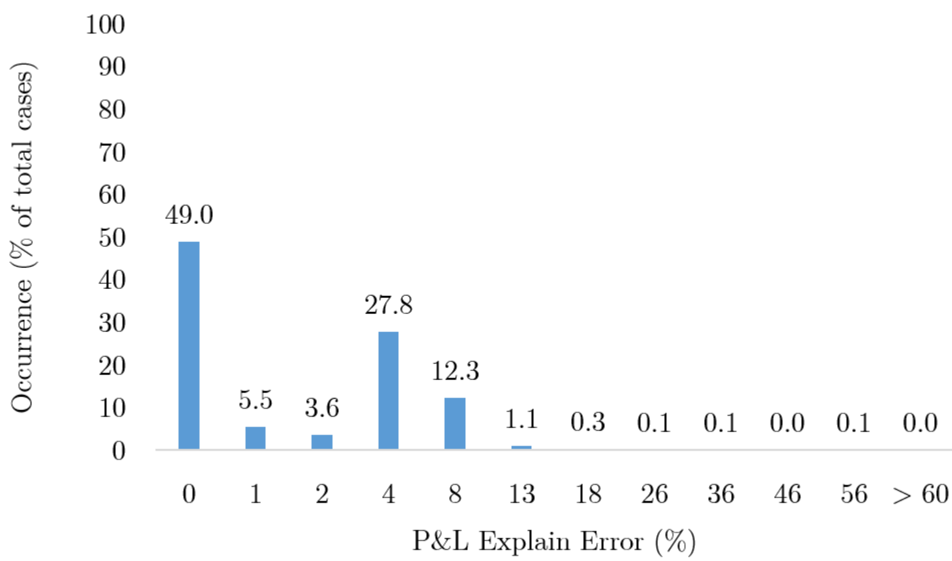

Figure 4: Same as Figure 2 except that our test portfolios are assumed to be delta-hedged.