BARNETT WADDINGHAM
beyond the expected

Briefing

Lloyd's post-diversified capital stability

Within the Lloyd's market, capital is allocated to individual risk types using the co-SpreadVaR method. The default Lloyd's co-SpreadVaR window does not take full advantage of the stability benefits of running a model on a high simulation count. In fact, the higher the simulation count, the smaller the default Lloyd's window[1] is. In this briefing, we discuss how using a window size that is a fixed percentage (e.g. 0.3%) of the simulation count gives a more stable allocation, particularly for models with high simulation counts.

Lloyd's allows managing agents to use an alternative window to calculate post-diversified risks with justification, particularly when submitting their results to Lloyd's. However, few managing agents have yet chosen to take this option.

Improvements in modelling software and a desire for more stable results have led many agents to run their internal models on a higher simulation count. Therefore, it is now more important than before to consider alternative window sizes to calculate post-diversified capital.

Introduction

Post-diversified risk values are calculated to understand how different sub-risks, or risk categories, contribute to the total risk capital. Under Solvency II, the total risk capital that needs allocating to sub-risks is the 99.5th value-at-risk (VaR) and is equal to the Solvency Capital Requirement (SCR).

The method prescribed by Lloyd's to calculate post-diversified figures is the co-SpreadVaR method[2]. Lloyd's prescribes this method for calculating risk category level contributions to the total SCR and also for calculating class of business level contributions to Insurance risk.

The focus of this paper is at the total SCR level. The high-level risk types contributing to the SCR are insurance risk, market risk, credit risk and operational risk; and the sum of the post-diversified figures for these risks equals the value of the SCR.

Post-diversified figures are receiving increased regulatory scrutiny, leading to capital loadings for syndicates with low submitted post-diversified risk values. This is because post-diversified figures are generally considered to be unstable and therefore difficult to rationalise. A lack of rationale may then cast doubt over the validity of the low submitted post-diversified risk values.

In this paper we recommend action that managing agents could take to improve stability in the post-diversified figures.

1 In this paper, "window" refers to a band of simulations centred on the 99.5th percentile of the total risk. The Co-SpreadVaR method allocates capital based on the averages over these simulations.

2 An introduction to the co-SpreadVaR method and the Lloyd's default window size is provided in the Appendix.

Analysis and conclusions

Overview of analysis

In the following analysis we construct a simple example model and calculate co-SpreadVaR post-diversified allocations for various different window sizes and simulation counts. This analysis is repeated for 1,000 different seeds, to allow us to assess the stability of the post-diversified allocations.

Our example model

We have used a simple illustrative model for the analysis, comprising of four shifted-lognormal[3] distributions, representing insurance, market, credit and operational risk. A Gaussian copula defines the dependencies between these risk types. The parameters of each distribution, along with the correlation matrix used for the copula, were selected to be 'typical' based on a range of models we have observed in the market.

We have performed this analysis on a wide range of marginal distributions and correlations, which all gave the same conclusion. Therefore, we consider the example we present here to be broadly illustrative of the general case.

Producing co-SpreadVaR results

To investigate the stability of co-SpreadVaR at the 99.5th percentile, we ran the example model on 1,000 different seeds and for 20 different sim counts (between 10,000 and 500,000). For each of the 20,000 model runs we calculated the co-SpreadVaR contribution of each risk type to the 99.5th percentile using different window sizes; 20 different window sizes are used within our analysis, varying between 0.018%[4] and 0.5%. A window size of 0.2%, for example, would correspond to calculating co-SpreadVaR over the window between the 99.4th and the 99.6th percentiles of the sum of all four risks.

We express these results as percentage contributions to total VaR, in order to focus on the contribution of each risk rather than variability in total capital. This is consistent with the Lloyd's method of scaling co-SpreadVaR values such that they sum up to the total VaR.

3 If Y has a lognormal distribution with parameters μ and σ then X=Y+s has a shifted-lognormal distribution with shift s.

4 0.018% is the smallest window size used in the analysis because it is just below the Lloyd's default window size for a simulation count of 10,000.

Standard deviation calculated across 1,000 runs with different seeds, on various sim counts

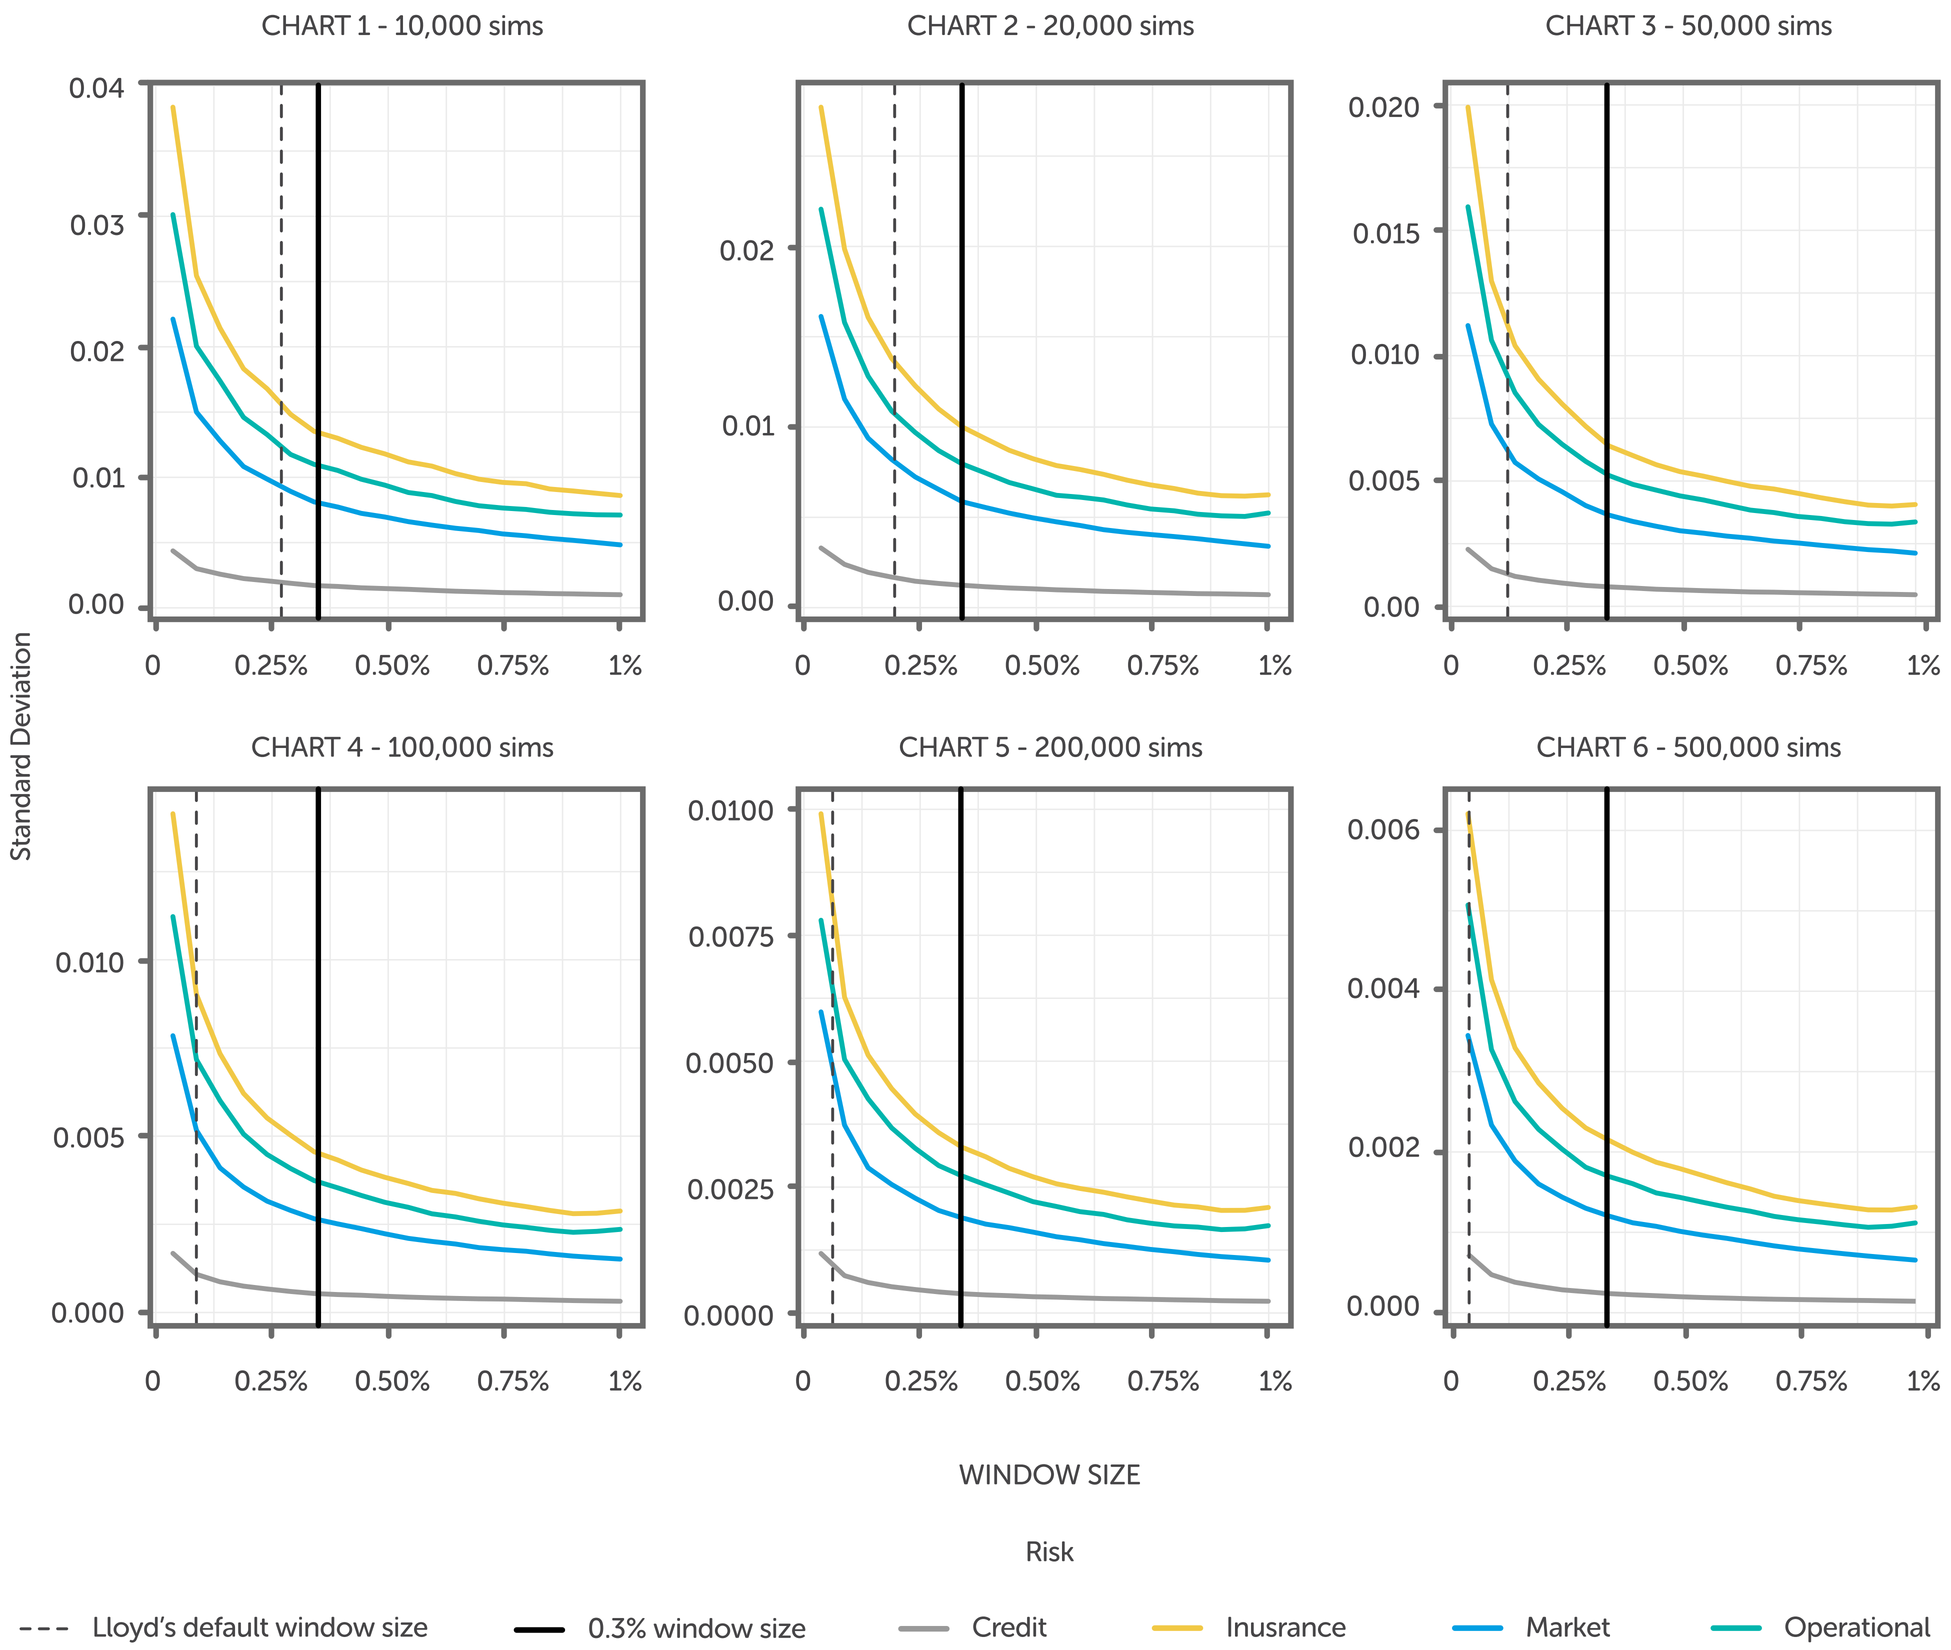

Results

To measure the instability of co-SpreadVaR, we took the standard deviation across the 1,000 seeds[4] for each sim count and window size. Where co-SpreadVaR is less stable, there will be a larger spread of results from using different seeds, and hence a higher standard deviation. The results are shown in the charts 1-6 above, presented for 20 different window sizes and six different simulation counts. In these charts the dotted line represents where the Lloyd's default window size sits, which varies by sim count.

From the charts opposite we can see that the contributions are more stable when a larger window is used, in line with expectations. However, using a larger window has the downside that you may include less relevant simulations. For example, the largest symmetrical window size possible runs from the 99th percentile to the maximum modelled value. The composition of the maximum and 99th percentile losses are less likely to be representative of the true composition of risk at the 99.5th percentile.

The instability curve shapes are similar for each risk type and sim count, initially showing a sharp decrease in standard deviation when the window size is increased by a small amount, followed by diminishing returns. From the analysis to the right, a full window size of 0.3% appears to give a reasonable balance between stability and relevance across the different scenarios. Other window sizes can also be justifiable depending on the number of simulations used.

Increasing the sim count increases the stability of the contributions (the scale of the graph axes changes). However, the Lloyd's window size becomes smaller relative to the sim count as the sim count increases. This partly offsets the stability benefits of using a larger sim count. **As we increase the sim count, the Lloyd's default window size becomes relatively less stable compared to what is possible for a given sim count.**

4 We performed the same analysis using range (the maximum contribution less the minimum contribution across different seeds) instead and it gave the same conclusions, therefore we only present the results using standard deviation

In Charts 7-10 we explore this further. We present the standard deviation of the allocations for 20 different simulation counts, relative to that of 10,000 simulations, on three bases:

1. The post-diversified capital using a fixed (relative to sim count) 0.3% full-window size;
2. The post-diversified capital using the Lloyd's default window and;
3. The pre-diversified, i.e. standalone VaR.

We see that the Lloyd's default window approach limits the stability. This is shown by the divergence of the orange and blue lines on the graph: the blue line shows greater improvements in stability with increasing sim count while the orange line quickly levels off.

Using a window size fixed as a percentage of sim count (in this example 0.3%) allows the stability of the post-diversified result to scale in line with sim count.

This is shown by the grey line closely tracking the blue line; i.e. the stability of the post-diversified result sees a similar benefit to the pre-diversified result when sim count increases.

The Lloyd's default window size works best for insurance risk, where the graph shows less divergence between the orange and blue lines. This

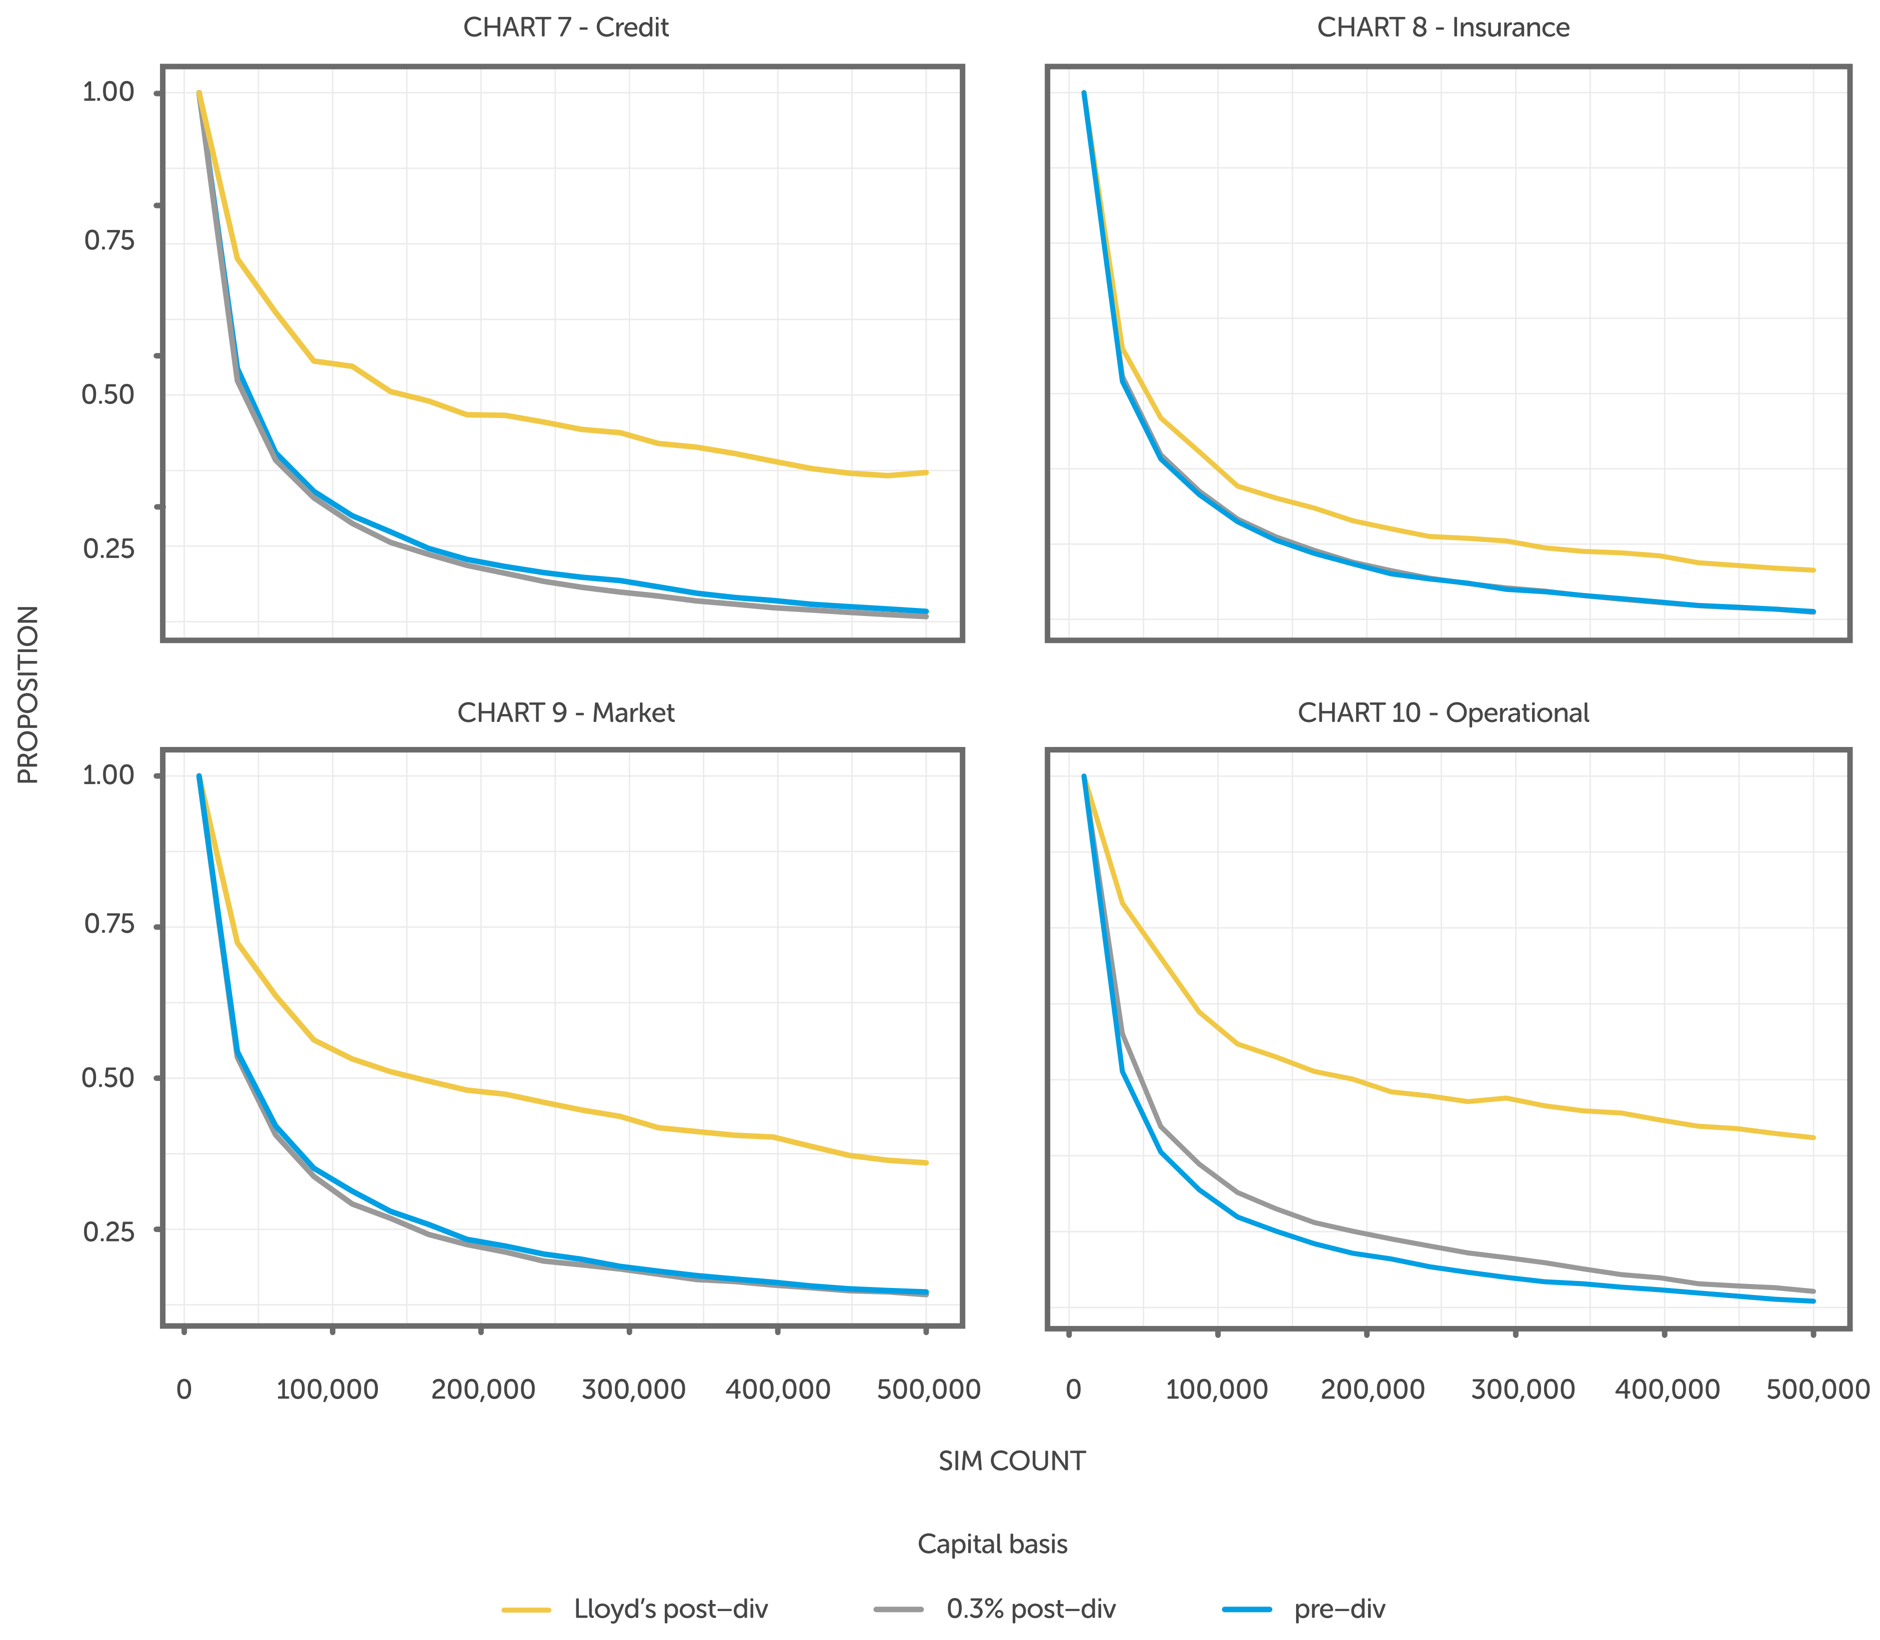

is expected, as insurance risk is the main driver of total capital. Therefore, as the 95% confidence interval is designed to give stability in the total risk measure, the approach works well for the co-measure on insurance risk too.

The other risk categories do not generally drive the total capital and hence the confidence interval approach based around total risk works less well. This is consistent with Lloyd's feedback we have observed in the market, which often highlights issues with the post-diversified figures for these smaller risk types.

Conclusion

In conclusion, for models using a high simulation count, we recommend using a wider co-SpreadVaR window than the Lloyd's default window. This is because the default Lloyd's window size definition becomes less optimal when a model is run on a high simulation count. We propose that the window size is defined as a fixed percentage of simulation count to ensure that post-diversified risk stability scales (with sim count) in line with total SCR stability. A wider window sacrifices relevance for stability and there is a balance to be struck. Based on our analysis, a fixed 0.3% window size gave a reasonable balance.

Appendix

Why co-SpreadVaR?

The purpose of calculating post-diversified figures for Lloyd's SCR reporting is to understand which risks are key drivers of the total aggregated risk and, implicitly, the dependencies between them. The current Lloyd's method aims to break down the capital setting scenario directly into the risks responsible for that scenario. We considered analysing alternative methods of allocating capital, such as Bodoff percentile layer or co-TVaR. However, these were considered less appropriate for determining the drivers of total VaR than the co-SpreadVaR method because their methodologies were less well aligned with the specific purpose we consider here. This sub-section explains the motivation for using the co-Spread VaR method for the purpose of Lloyd's SCR reporting.

For a particular model run, the actual contribution of each risk to the 99.5th percentile is determined by the single simulation at the 99.5th percentile of the total distribution, the co-VaR method, as per the following example.

Example: 1,000 simulations, with each row representing a simulation. We can sort the rows by the total risk column and look at the row corresponding to the 99.5th percentile of the total.

Percentile of sampled total	Insurance risk	Market risk	Credit risk	Operational risk	Total risk
99.5th	150	30	20	15	215

The total 99.5th VaR is 215 and the post-diversified values under the co-VaR method are 150, 30, 20 and 15 for each of the risks respectively.

However, this is prone to significant simulation error and rerunning the model on an alternative seed could lead to wildly different post-diversified contributions

To determine the 'true' contribution[5] of each risk to the SCR, ideally the model would be run many times and the post-diversified figures would be calculated as the average of all of these runs based on the relevant single simulation of each run. But this is impractical as the post-diversified risk values need to be calculated for each model run, of which there are many in the run up to a typical capital submission[6].

5 The 'true' contribution of each risk being the contribution with no simulation error.

6 The co-SpreadVaR method is generally defined as the method of taking averages over simulations in the window. However, when we refer to co-SpreadVaR throughout this paper, we include the additional step of scaling the averages such that the post-diversified figures sum up to the total VaR. This is consistent with the Lloyd's application of the SpreadVaR method in cases where the total SCR is calculated based on the 99.5th VaR instead of SpreadVaR.

Hence, in an attempt to estimate the 'true' contributions based on a single model run, the average of a range of simulations close to the 99.5th VaR is taken – this is the co-SpreadVaR method .

Example (continued): We choose a window for co-SpreadVaR between the 99.4th-99.6th percentiles as an example.

Percentile of sampled total	Insurance risk	Market risk	Credit risk	Operational risk	Total risk
99.4th	170	10	15	17	212
99.5th	150	30	20	15	215
99.6th	172	5	25	16	218
99.4thto 99.6th co-Spread VaR allocation	164	15	20	16	215

The total 99.5th VaR is 215 and, in this example, the average total risk over the window is also 215. We calculate the post-diversified values by averaging each risk category over the window. If the total of the average was different to the value of the SCR, we would scale the post-diversified figures accordingly such that they sum up to the SCR.

Default Lloyd's "window"

Lloyd's prescribes the use of the co-SpreadVaR method for calculating post-diversified figures and also provides a default method of calculating the window of simulations (around the 99.5th percentile of the total SCR distribution).

The default method sets the co-SpreadVaR window size using a confidence interval approach. This window corresponds to the 95% confidence interval of the 'true' 99.5th percentile loss. As a consequence, the width of this interval relative to the simulation count decreases with the square root of the simulation count. This is because a larger 'sample size' gives lower simulation error.

That is, when a model is run on a higher simulation count, we can be more confident that the true 99.5th of the total distribution sits within a smaller range around the observed 99.5th.

However, this logic does not directly translate to the calculation of post-diversified figures. If a model is run on a higher simulation count, it does not translate that the true allocation is the average over the smaller simulation count defined by this confidence interval. This is because this confidence interval relates only to the 99.5th percentile of the total and is not a confidence interval for the co-SpreadVaR values determining the contributions to the total VaR. Often, individual risk categories may demonstrate unstable behaviours at the 99.5th (e.g. when dealing with highly skewed distributions) which can invalidate this premise. premise; e.g. when dealing with highly skewed distributions.

As a leading independent UK professional services consultancy across insurance, risk, pensions and investment, we're free to do the right thing by all our clients at all times. With over 1,400 dedicated people, we act for a wide range of clients, both large and small, including over 15% of FTSE 350 companies.

Our insurance experts provide market-leading advice and support on actuarial, investment, longevity and risk management solutions. We employ a personal and partner-led approach, with a growing client base covering a broad range of life and general insurance firms.

BARNETT WADDINGHAM
beyond the expected

Please contact your Barnett Waddingham consultant if you would like to discuss any of the above topics in more detail. Alternatively get in touch via the following:

wan.heah@barnett-waddingham.co.uk

020 7776 3819

cherry.chan@barnett-waddingham.co.uk

020 7776 3882

Contributors to this briefing note

Alpesh Patel, Matthew Bett and John Townhill

www.barnett-waddingham.co.uk