Additional comments on Adrija Djurovic's presentation of Concentration Risk
Alan Forrest 12/07/25

Concentration Metrics This link gives Andrija's outline of the basic metrics for measuring Concentration among PD risk grades, and the European regulatory requirements for reporting them. Andrija's presentation gives helpful commentary and links to useful code – R and Python – to calculate these metrics. This should be read and explored by all students of quantitative risk management and by model managers, as an example of best practice and regulatory requirement for Credit Risk IRB.
Concentration Risk This linked folder contains Andrija's presentation summary of the IMF's Partial Portfolio Approach to Concentration Risk, with illustrating data and code in R and Python. The IMF's hybrid model of portfolio credit risk has a structure that maintains good performance even in the presence of high name-exposure concentrations, and therefore is an important model option and a mitigation of Concentration Risk. This should be read by model developers and model risk managers as an example where model structure has mitigated a potential model error, so helping manage the specification risk of the model.

---

1 Further thoughts

A good presentation and style grabs the interest of the reader and sets their thoughts running independently. I have found this the case with Andrija Djurovic's two sets of slides on Concentration Risk, linked above, and I use this space to add my own thoughts. These are aimed at different audiences but I hope are interesting to all:

- More about concentration measures - for the student
- Statistical tests of Concentration - for the model manager
- Concentration as a Model Risk - for the model risk manager

Both Andrija and I refer to the following documents, linked in Andrija's presentations and in the text below. I am grateful for Andrija's feedback on this.

1.1 References

[ECB] European Central Bank, Instructions for reporting the validation results of internal models, Feb 2019
[EBA] Committee of European Banking Supervisors, Guidelines on the management of concentration risk under the supervisory review process (GL31) , September 2010.
[IMF] Pierpaolo Grippa and Lucyna Gornicka, Measuring Concentration Risk A Partial Portfolio Approach, IMF Working Paper WP/16/158, August 2016

2 Measuring Concentration

In his slides, Andrija describes the importance of measuring portfolio concentration and the risks of operating an unbalanced grading. He also details the calculations and tests required for European IRB Banks by ECB guidance [ECB , Section 2.5.5.3] to assess concentration, and he puts the calculations and simulations helpfully into code. Here I expand on the approximate relationships between these concentration metrics and their statistical properties, and consider other measures and tests as well.

Start with a portfolio of N obligors, distributed into K risk grades (or other categorisation), with a proportion R_i in grade i. This proportion can be a proportion of number of obligors, or EAD-weighted, or of other suitable weighting. In any case $R_i \geq 0$ and $\sum_i R_i = 1$.

The Coefficient of Variation CV for this distribution is defined in the ECB guidance, and can be rewritten:

$$CV^2 = K\sum_i \left(R_i - \frac{1}{K}\right)^2 = K\left(\sum_i R_i^2 - 1\right)$$

The ECB paper also details the Herfindahl Index (also called Herfindahl-Hirschmann Index)

$$HI = 1 + \log\left(\frac{CV^2+1}{K}\right)/\log K = 1 + \frac{\log(\sum R_i^2)}{\log K}$$

(all logarithms are natural) which simplifies then approximates as

$$HI = \frac{\log(CV^2+1)}{\log K} \simeq \frac{CV^2}{\log K}$$

the approximation being strongest for small values of CV. Thus HI is approximately CV^2 scaled to take account of the number of grades.

The Coefficient of Variation has strong foundations in classical statistics. If R_i are count proportions, then $N.CV^2$ is the classic Chi-squared statistic that compares the observed portfolio grade distribution with a uniform distribution (same number of obligors in each grade). Under the hypothesis that R is uniform, $N.CV^2$ is distributed as Chi-squared with $K-1$ degrees of freedom.

CV can also be approximated by other classic metrics, for example Population Stability Index (PSI).

$$CV^2 = \sum_i \left(R_i - \frac{1}{K}\right)(R_iK-1) \simeq \sum_i \left(R_i - \frac{1}{K}\right)\log(R_iK) = PSI(R, \frac{1}{K})$$

where the (natural) logarithm approximation is strongest when R_iK are close to 1.

2.1 Other kinds of concentration

The set up above considers a portfolio distributed among risk grades. The same construction can also explore concentration in geographies and in other natural segmentations of the portfolio – eg industry sectors in a corporate loan portfolio. Note that assessment of such concentrations is required by regulation, for instance CEBS Guidelines [EBA, guidelines 7, 12 etc.]

The CV metric can be repurposed also to assess single name concentration – where a large fraction of the portfolio's exposure is concentrated in a small number of obligors. Single name concentration is also a concern of Capital regulation.

To do this we push the construction to its extreme and put one obligor in each grade (so $K = N$) and weight by exposure. For this discussion I'll take $R_i = EAD_i/TEAD$, where $TEAD$ is the portfolio total EAD. The CV and HI can be computed using the same formulae as before, so

$$CV^2 = \left(\frac{N.\sum EAD_i^2}{TEAD^2}\right) - 1$$

Large values of CV indicate that much of the portfolio's EAD is concentrated on only a few obligors.

Another metric of EAD concentration is the Gini index, computed as:

$$\frac{1}{N}\left(\frac{1}{2} + \frac{1}{TEAD}\sum_i \sum_{j:\ EAD_j > EAD_i} EAD_j\right)$$

This has no close relation with the CV and HI metrics, but is also used in assessing concentration and in Granularity Adjustment - see later.

3 Using concentration metrics for monitoring and testing

To test for changes of grade distribution over time, the ECB guidance [ECB , Section 2.5.5.3] cites the Feltz-Miller test. This tests the null hypothesis $CV = CV^o$, where CV^o is the coefficient of variation of an initial portfolio. Deviations of the current CV from this initial value are measured and detected by the Feltz-Miller statistic whose distribution is normal under this hypothesis. Too great a deviation (or too small a p-value) suggests the current distribution has changed from the initial distribution. The ECB requires this test's p-value to be reported, though it does not give guidance what values to expect and act on.

Given an initial distribution $R^o = (R_i^o)$, we can test the hypothesis $R = R^o$ by the statistic:

$$N\sum_i (R_i - R_i^o)^2/R_i^o$$

which under this hypothesis is distributed as Chi-squared with $K - 1$ degrees of freedom. If this statistic is significantly large (at an agreed level) then we declare that the current portfolio distribution has changed from the initial distribution.

How does this Chi-squared test compare with the Feltz-Miller test?

There is no doubt that the Chi-squared test is more powerful – it will detect true differences more reliably, without increasing the chance of false alarms. You can see this by comparing the information used by each test: the F-M test compares aggregate CVs, the Chi-squared test compares the distributions behind these aggregates, grade by grade: Chi-squared beats F-M in the same way that a paired test beats an unpaired test.

But the Chi-squared test depends on a pairing, not just the summary CVs, and this imposes requirements on the data and context: we need the same grading scale initially and currently, and we need to know the distributions R^o and R.

The ECB may have had this point in mind when it preferred the less powerful F-M test for its requirements. In practice, grade information may be lost and grade scales may change over time, so that CVs might be the only comparable statistics and the F-M test the only option.

However, if paired grade-level information is available it is important to exploit its additional power. After meeting the ECB's regulatory obligations, I recommend following up with the Chi-squared test described above to detect changes in grade distributions more powerfully.

One final note of caution - when using any statistical test, a significant result should never be the only driver of a business decision. For instance in the Chi-squared test above, the factor N in the formula can drive large values of Chi-squared in the presence of even tiny changes in distribution, especially for a large portfolio, say $N > 10000$. This is reasonable because the statistical test asks quite a weak question: is there evidence that the distribution among grades has changed in any way? This is not the same as the business question which is alert to distribution changes that affect business outcomes, of a certain shape or materiality. Therefore the test above may be too sensitive or unspecific for business use, for which it then needs adjustment or interpretation. Note that this mismatch between business significance and statistical significance also applies to the F-M test, but its relative lack of power tends to dampen its sensitivity in practice.

4 Concentration as a Model Risk

The modeling difficulties caused by concentration are presented clearly and readably in the IMF paper [IMF], which also proposes a solution – the Partial Portfolio Approach. Andrija's slides summarise this approach clearly and give code to implement it. Here I consider this development from the point of view of model risk: Concentration as a source of model risk and the Partial Portfolio Approach as an example of effective model specification risk mitigation.

The IMF paper examines the credit risk of a portfolio by two possible model approaches:

1. simulate each exposure and exposure interactions in detail and extract aggregate risk information from formulae or simulations;

2. use a portfolio-level model, with approximating or simplifying assumptions about the portfolio.

The first approach above is accurate but expensive in time and resources, especially if it is applied to a large portfolio. The second approach is usually adopted, with the Asymptotic Single Risk Factor (ASRF) model as the industry and regulatory standard.

ASRF builds on the strong foundations of the Merton-Vasicek model of default, and is sophisticated enough to take account of exogenous influences, idiosyncratic behaviours and a network of correlations between exposures. But like all models the ASRF also rests on assumptions, idealisations and approximations of reality: in particular the ASRF assumes the portfolio has infinite granularity - the mathematical limit of a large number of obligors and close to zero concentration.

Infinite granularity clearly does not happen in real portfolios, and this broken assumption is a source of model risk and possible model error. When concentrations are large the assumption of infinite granularity becomes untenable and ASRF needs correction. To do this, various kinds of Granularity Adjustment (GA) can be applied to ASRF, as described and referenced in the IMF paper. In practice, banks create tables mapping HI or Gini values to the appropriate GA, which then adjusts Capital through Basel Pillar 2. However, these GAs are also a source of model risk, as they opt among various methods using various principles, all with approximation and uncertain stability.

Thus under high concentration, both the ASRF model and its adjustments are beset with model risk that should be quantified and managed, a significant task in addition to the implementation and operation of these models. Must we now bring the first approach above back into play?

The IMF's Partial Portfolio Approach answers this question: "yes, we should use the first approach but only as far as we need to."

High concentration, measured by HI or Gini, is typically caused by a few high EAD obligors. If those obligors are excluded then the remainder of the portfolio (the Granular part) is likely to have low concentration, and the ASRF can be applied to it. On the single names with concentrated EAD (the Non-granular part) a full simulated model will work efficiently because the obligors are few. Each model is set up on its part of the portfolio and they are joined under a common simulation of exogenous effects. Sufficient simulations are run to achieve good estimates of the 99.9%-ile, as required by regulation. Andrija's code can control the number of simulations and is illustrated with $100,000$ which is sufficient reliably to achieve about 3bp accuracy.

The tactic of applying both models partially is intelligent and reduces model risk significantly and structurally. It applies each model exactly where it is most accurate and where its assumptions work best, while coordinating them under a common exogenous environment. Therefore model risk is mitigated well by the structure of this method: we still have to keep alert of course, but now Specification Risk is no longer an urgent issue.

But even good ideas come at some cost. To the model risk manager, concerns remain about the modelers' choice where to draw the line between granular and non-granular parts. In the absence of general principles – the IMF paper gives no guidance here – sensitivity analysis might be the best option. Such analysis compares the performance of several challenger models built with various choices of granular cut-off. Depending on the results and our comfort with their spread, we might set up model corrections or other model risk actions as appropriate.

Alan Forrest, Edinburgh, July 2025