The Margin of Conservatism (MoC) in the IRB approach: defining and measuring the general estimation error

di[1] Franco Varetto (Politecnico di Torino) e Silvio Cuneo (Intesa San Paolo)

Abstract (English)

Since the Basel 2 Accord, the Regulators have raised the issue of errors arising in the internal model estimation process, calling for margins of conservatism to cover possible underestimation in capital. However, this issue has been dealt with just a few general statements into to the primary regulations, until the EBA has devoted a material share of its 2017 Guidelines on model estimation to build up a framework to define, classify and quantify the Margin of Conservatism (MoC).

In this article, we surveyed the regulatory requirements and their methodological foundations, coming to propose to a method for MoC quantification based on a simple k-sigma formula.

Both the components of the method (i.e. quantifying the standard deviation of the parameter estimator and setting the k factor) are analysed from the methodological and the operational point of view:

- For sigma, some alternative approaches are put forward, as well as an approach to models designed as the aggregation of multiple components;
- For k, whose calibration is crucial to avoid (or at least to limit) double counting of risks and capital requirements, a closed-formula and a Monte Carlo approach are discussed, and the latter is elaborated in depth with a final calibration proposal.

The obtained results are used to perform an impact analysis on RWAs according to different calibration scenarios for k and sigma. The main outcome is that MoC, if not properly calibrated, can be very burdensome for small portfolios, and especially for low default portfolios.

Finally, some operational issues are explored concerning the implementation of the MoCs on “real life” portfolios, where different segmentations across risk components may involve technical problems, and the extension of MoCs to defaulted asset RWAs computation. Also, the lifecycle of MoCs should follow the one of the related models, thus implying a synchronization in the lifecycle of the models for all risk components for each portfolio segment.

Abstract (Italiano)

Fin dall’Accordo di Basilea 2, i Regulators hanno posto il problema degli errori derivanti dal processo di stima dei modelli interni, richiedendo margini di prudenzialità per coprire possibili sottostime del capitale. Tuttavia, il tema è stato affrontato soltanto in termini generali nella normativa primaria, finché l’EBA ha dedicato una significativa porzione delle sue Linee Guida per la stima dei modelli del 2017, che stabiliscono un quadro per definire, classificare e quantificare il Margin of Conservatism (MoC).

In questo articolo ripercorriamo i requisiti regolamentari e i loro fondamenti metodologici, pervenendo ad una proposta per la quantificazione del MoC basata sulla semplice formula k-sigma.

Entrambe le componenti del metodo (ovvero la quantificazione della deviazione standard dello stimatore e la calibrazione del fattore k) sono analizzate dal punto di vista metodologico e operativo:

- Relativamente a sigma, vengono proposti alcuni approcci ed esaminato un metodo specifico per i modelli basati sull’aggregazioni di componenti;
- Per quanto riguarda il fattore k, la cui calibrazione è cruciale per evitare (o quantomeno limitare) il “double counting” di rischi e requisiti, vengono discussi un approccio a formula chiusa e uno basato sul metodo Monte Carlo; quest’ultimo è elaborato in dettaglio con una proposta finale per la calibrazione di k.

I risultati ottenuti sono stati utilizzati per effettuare un’analisi di impatto sui RWA secondo differenti scenari di calibrazione per k e sigma. L’evidenza principale è che il MoC, qualora non calibrato in modo appropriato, può risultare estremamente penalizzante per portafogli piccoli, e in particolare per i “low default portfolios”.

Da ultimo, sono stati analizzati alcuni problemi operativi riguardanti l’implementazione del MoC su portafogli reali, dove segmentazioni diverse tra i fattori di rischio possono comportare problemi tecnici, e l’estensione dei MoC al calcolo dei RWA per le esposizioni deteriorate. Infine, il ciclo di vita dei MoC dovrebbe seguire quello dei relativi modelli, implicando la necessità di una sincronizzazione tra tutti i modelli per i diversi fattori di rischio di ciascun segmento del portafoglio.

1 Foreword

Since the Basel 2 Accord, the Regulators have raised the issue of errors arising in the internal model estimation process, calling for margins of conservatism to cover possible underestimation in capital.

However, this issue has been dealt with just a few general statements into to the primary regulations, namely the 2013 Capital Requirement Directive and the EU Regulation 575/2013 for European Union.

More recently, the EBA has devoted a material share of its 2017 Guidelines on model estimation to build up a framework to define, classify and quantify the Margin of Conservatism (MoC). Furthermore, regulation has been brought out on Model Risk, which appears to have close relations with MoCs, though still not thoroughly elaborated.

To discuss the issue and put forward methodological proposals for the application of regulatory requirements on Margin of Conservatism in banks models, AIFIRM set up a dedicated Commission, coordinated by Silvio Cuneo (Intesa Sanpaolo) and Franco Varetto (Politecnico of Torino).

---

[1] L’articolo è stato inoltre sottoposto a doppio referaggio anonimo, pervenuto in data 31/01/2011 e accettato il 19/03/2019.

This article summarizes the main findings of the Commission works, which can be read in detail into the AIFIRM Position Paper n. 13. (cfr. http://www.aifirm.it/position-paper-2/)

2 The Margin of Conservatism in the Regulation

The story begins with the Basel 2 Accord, which at par. 451 states:
"In general, estimates of PDs, LGDs, and EADs are likely to involve unpredictable errors. In order to avoid over-optimism, a bank must add to its estimates a margin of conservatism that is related to the likely range of errors. Where methods and data are less satisfactory and the likely range of errors is larger, the margin of conservatism must be larger".
The risk of introducing undesired and excessive capital burdens to the bank is quite concrete. This seemed to be already apparent to Basel Regulators since the beginning of the Accord: in the 2005 BCBS Paper "An Explanatory Note on the Basel II IRB Risk Weight Functions" (see [29]), one can read at par. 5.1: "The high confidence level was also chosen to protect against estimation errors, that might inevitably occur from banks' internal PD, LGD and EAD estimation, as well as other model uncertainties".

However, the issue has been dealt with just a few general statements into to the primary regulations, namely the 2013 Capital Requirement Directive and the EU Regulation 575/2013 for European Union.

These high-level principles have been interpreted by the banking industry and national supervisors in various ways, leading to a wide variety of practices. Hence, in the context of its project to harmonize industry practices and reduce RWA variability, the EBA has devoted a material share of its 2018 *"Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures"* to build up a framework to define (EBA GLs), classify and quantify the Margin of Conservatism (MoC).
The EBA GLs introduces the following taxonomy (par. 42):

- ✓ Category A: MoC related to data and methodological deficiencies …;
- ✓ Category B: MoC related to relevant changes to underwriting standards, risk appetite, collection and recovery policies …;
- ✓ Category C: the general estimation error.

Here we are focusing only on the last MoC category[2], whose requirements can be summarized as follows:

- ✓ It should reflect the dispersion of the distribution of the statistical estimator (par. 43)
- ✓ It should be quantified at least for every calibration segment (par. 43)
- ✓ It must always be greater than zero (par. 47)
- ✓ (together with the other MoCs) It should not distort with excessive adjustments the estimates of the risk parameters and the resulting own funds requirements (par. 48).

It is also convenient mentioning what the GLs say about the aggregation of MoCs at Chapter 2 Background and Rationale:
"…for the purpose of harmonisation and in order not to impose over sophistication, it has been decided to require the aggregation of MoC between the categories based on a simple sum. However, different aggregation techniques may be used within each of the categories".
Though this is related to the different MoC types, we notice that the same issue must be coped with for the aggregation of the MoCs on the different risk components. Specifically, general estimation errors can be reasonably assumed independent on one another while separately adjusting each risk parameter estimator for its own MoC involves an implicit assumption of correlation between errors. MoC aggregation should hence be designed in a way not produce "excessive adjustments", as stated before. In this regard, it worth noticing that the EBA GLs themselves clearly states in par. 43. to "[…] quantify MoC for the identified deficiencies referred to in paragraphs 36 and 37 [i.e. categories A and B], <u>to the extent not covered by the general estimation error</u>, […]", thus putting attention if the generalized estimation error, in the way in which is quantified, might already factorize the additional buffer of conservatism stemming from a category A or B deficiency.

A further regulatory reference, still published as a consultative document in September 2018, is the *"ECB guide to internal models, Risk type specific chapters"*. The MoC paragraph, specifically concerning MoC of type C, extends to CCFs the content of the EBA GLs on PD and LGD estimation. Furthermore, for PDs it specifies that the MoC should "account for statistical uncertainty/sampling error affecting the LRA estimate at grade level stemming from the variability of each year's default rate and from the period considered…".
This raises some concerns:

- Estimating the MoC at the grade level alters the shape of the discriminant function. An inversion of the curve (leading to a reduction of predicting power) can even happen if the MoC is inversely proportional to PD in some segment of the curve, as it might easily be the case if the MoC is proportional to the sample size. Moreover, the estimation becomes rather arbitrarily dependent on the number of grade levels of the model. For that reasons, we believe that this provision should be interpreted in the sense that the MoC must be calculated at the model level and then allocated to each grade in a way that preserves the predicting power of the model;
- Considering separately the default rate of each year can confuse the variability of the default rate time series with that of its estimator: only the latter should contribute to defining the type C MoC. We believe that this statement should

[2] In what follows we will refer shortly to MoC as the type C MoC.

be interpreted as a recommendation to consider the whole time series in order to evaluate the variability of the estimation error.

Both these issues are further explored in the methodological analysis and in the proposals illustrated forward below.

Finally, Regulation has been brought out on Model Risk, which appears to have close relations with MoCs, though still not thoroughly elaborated. A specific issue relates to the definition of a capital buffer to cover model risk: if such a capital buffer is deemed necessary, a mapping of prudential elements applied to face specific regulatory request should be done in order to avoid any double counting on economic capital measures with respect to MoCs (and generally speaking add-ons) already included in Pillar I measures. Probably, most part of these prudential measures are considered also in the economic capital ones, which justifies the fact of not applying a model risk capital buffer to Pillar I risks that are also applied in Pillar II (ICAAP). Banks should calculate an economic capital to face model risk only if no "general estimation error" MoC is included in Pillar II.

3 The methodological framework

In order to define a method to calculate the MoC in compliance with the regulatory requirements, we started surveying the regulatory requirements and the methodological foundations of the MoC related to the general estimation error.

After recalling the fundamental difference between the variability of the observed variable and that of its estimator, several approaches have been examined, ranging from parametric and non-parametric inferential statistics to Bayesian approach.

We refer to the position paper for the full methodological review.

It comes out that there is no simple way to address the issue consistently with methodological soundness and meaningfulness of the results.

The difficulties arise already at the level of a single risk parameter (PD, LGD and CCF), but become overwhelming when we tackle the issue of the aggregation of MoCs on different risk parameters.

We thus formulate a set of general principles that should be followed to comply with regulatory requirement while avoiding excessive and unjustified burdens:

1. In order to define the size of the adjustment, it should be kept in mind that the confidence interval of Basel formulas is fixed at 99.9% for PD parameter in order to "protect against estimation errors": category C MoC determines a further implicit increase of this confidence level, which is already quite high.
2. Category C MoC is to be computed at calibration segment level rather than rating grade level (once computed at calibration segment level, an application at grade level is always possible). Two main reasons are motivating the choice: MoC concerns capital quantification, thus it should affect calibration which is the final layer in the parameters determination; furthermore, if applied at rating grade level, it could introduce distortions in the measures: for example, it can break the monotonicity of a PD scale, thus reducing the predictive power of the model (the same is true for the other risk parameters).
3. Consequently, the granularity should follow the calibration, generally coinciding with (regulatory) model level. The rationale is to avoid an excessive granularity, potentially leading to distortions.
4. In the case of models built through a "component" approach, the potential overestimate of the error due to the aggregation of independent measures must be limited.
5. A similar consideration holds for the aggregation in the Expected Loss measure of MoCs calculated separately for PDs, LGDs and CCFs: given the correlations among the three MoC measures, a cap on the total level of the add-on should be imposed and then single parameters levels should be consequently adjusted.
6. The estimate error measurement must be (inversely) proportional to total sample dimension, while time dimension *per se* should not be relevant.

We propose then a simple method to calculate the MoC as an adjustment to each risk parameter (PD, LGD and CCF) resulting from the product of a factor (k) times the standard deviation (sigma) of the parameter estimator:

MoC = K*sigma

In what follows, we discuss how to define and compute the sigma parameter and how to calibrate the k parameter in a way that is consistent with the principles set above.

4 The k-sigma approach: estimation of sigma

In order to define sigma, it must be recalled that we have to measure the "dispersion of the distribution of the statistical estimator". Hence, sigma should measure the volatility of the statistical estimator. Even if it may appear trivial, it is important to never confuse the variability of the population with the one of the estimator. This confusion comes even easier when talking about the LGD and CCF.

We have identified two alternative methods, one based on the calibration sample volatility and the second one based on the "within" variance of the estimates. The table below summarizes the main features of the proposed methods.

	Method 1	Method 2
PD	$\sigma = \sqrt{\frac{p \cdot (1-p)}{n}}$	$\sigma = \frac{\sqrt{s_W{}^2}}{\sqrt{N}} = \frac{\sqrt{\sum_{j=1}^{J} N_j \cdot \frac{N_j}{(N_j-1)} \cdot (PD_j - DR_j)^2}}{N}$
LGD/EAD (CCF)	Bootstrapping	$\sigma = \frac{\sqrt{s_w^2}}{\sqrt{N}} = \frac{\sqrt{\sum_{j=1}^{J} \frac{N_j}{N} \frac{1}{(N_j-1)} \sum \left(x_{ij} - \hat{x}_j\right)^2}}{\sqrt{N}}$
Legend of symbols	σ = standard deviation p = probability of default n = sample size	σ = standard deviation s_W = within standard deviation PD_j = probability assigned to cluster *j* DR_j = observed default rate of cluster *j* x_{ij} = *i-th* observation of cluster *j* $\hat{x}_j$ = estimator for cluster *j* N_j = size of cluster *j* n = sample size
Rationale	Variance of the binomial random variable for PD, numerical approximation for LGD and EAD	Decomposition between-within of estimation error variance and utilization of the within component
Pros and cons	• Direct measurement of the unavoidable estimation error, inversely proportional to the sample size • Extremely penalizing for low default portfolios	• Sigma inversely proportional to both the sample size and discriminant power ➡ risk sensitive approach which provides correct incentives to the estimation process • The MoC can theoretically (in case of "perfect" model) result equal to zero

It must be stressed again that sigma must be computed at the level of calibration sample and not at each rating grade level, in order to avoid several undesired consequences: arbitrariness in defining the master scale (especially the number of rating grades), disproportionate impacts on low populated rating grades and finally the risk of altering the rank order of the rating and diminishing its predictive power. However, it will be always possible to allocate the resulting MoC to each rating grade, for example via a Bayesian approach, rather than simply adding it to each grade of the scale, depending on the model design that the bank has adopted.
In summary, the two methods can be compared on several angles:

- ✓ On the methodological standpoint, both methods appear to be compliant to the requirement to assess the "dispersion of the distribution of the statistical estimator". This is a feature of the sample size and distribution and cannot be eliminated by the estimation model. The main difference is that method 2 admits theoretically a zero result for sigma, but only in the unrealistic case of a "perfect" model.
- ✓ From a practical point of view, method 2 has a number of positive features: it is risk sensitive, in the sense that a better model will have a lower sigma and consequently a lower MoC, which seems to be sound; it is computed at the rating grade/grid cell level, though recombined at the calibration level (thus resulting suitable in a context of calibration by grade or pool whereas Method 1 results more immediately applicable in presence of direct calibration at segment level); being sensitive to the discriminant power (in addition to sample numerosity), it can mitigate disproportionate results in low default portfolios.

Finally, the paper also explores an alternative approach, based on bootstrapping and the computation of semi-variance, that is the computation of variance only on "positive" estimation errors.

<u>The case of model components</u>

A further issue that should be treated is the calculation of sigma for models whose model design envisages the aggregation of model components. Instead of proposing a general approach, we explored a simple but common case: a LGD model based on two components, a Loss Given Loss and a Cure Rate:

$$LGD = d \cdot LGL + (1-d) \cdot LGC$$

where *LGL* is the loss incurred in the workout procedure, *LGC* is the loss incurred in case of cure (returning to performing status) and *1-d* is the cure rate, or *d* is the probability to enter the workout procedure ("danger rate").
If we assume for simplicity that *LGC≡0* and neither loss nor recoveries are incurred in the pre-workout phase, our model reduces to:

$$LGD = d \cdot LGL$$

Then, by assuming that the *d* and *LGL* estimators are independent variables, it can be shown that sigma can be computed according to the formula:

$$\sigma_{LGD} = \sqrt{\sigma_d^2 \cdot \sigma_{LGL}^2 + d^2 \cdot \sigma_{LGL}^2 + LGL^2 \cdot \sigma_d^2}$$

Where both σ_d σ_{LGL} can be obtained with either method 1 or method 2 described above.

To obtain this result, we made a couple of assumptions that deserve to be discussed.
Starting with independence between the estimators of model component, we note that this assumption is sound, since it is not related to the two variables, which may actually be correlated, but to their estimators: there is no reason why the errors made in the estimation of cure rate should be correlated with those made in the estimation of loss given loss.
On the other hand, it can be shown that a simple aggregation rule, consisting of computing separately the MoCs on cure rate and *LGL* and calculating the LGD with MoC as the product of the two post MoC components, embeds the implicit assumption of perfect correlation between cure rate and *LGL* estimation errors and leads to excess of conservatism.

The second one is more a simplification than an assumption, consisting in reducing the model to only two components. Actually, the extension to more model components is feasible though a closed formula for the variance of the non-linear aggregation of several components is not generally available and then the complexity can explode. We recommend striking the balance between theoretical perfection and simplicity, by limiting the computation of sigma only to the components that judgementally can be deemed as more relevant in the whole model.

Empirical results

The proposed methods have been tested on sample portfolios of a large bank.
Table 1 shows that the two methods yield very similar results for LGD and CCF, while Method 1 is far more conservative for PD and especially burdensome for low default portfolios. Obviously, these results could change over the time, likely in the downward direction, because of larger sample size (as long as time series expand) and more discriminant models.

Table 1: sigma (expressed in share of the average value of estimate)

	Sigma PD		Sigma LGD*		Sigma CCF	
	Method 1	Method 2	Method 1	Method 2	Method 1	Method 2
LDPs	4-11%	≈2%	4-10%	3-10%	N/A	N/A
Corporate	≈0.5%	≈0.2%	≈0.5%	≈0.5%	≈6%	≈5%
Retail	0.2-0.4%	0.2-0.4%	0.1-0.9%	0.1-0.9%	0.6-1.2%	0.6-1.2%

* computed only on workout exposures, can be regarded as a loss given loss
NB: floors to 1 bp have been applied on sigma absolute values

Table 2 finally reports the results obtained for model component LGD, where for sake of simplicity we only used Method 1 for the danger rate sigma (note also that sigma *LGL* are the values taken from LGD in table 1).

Table 2: sigma for model component LGD (expressed in share of the average value of estimate)

	Sigma *LGL*		Sigma *d*	Sigma LGD	
	Method 1	Method 2	Method 1	Method 1/1	Method 2/1
LDPs	4-10%	3-10%	1,5-2,5%	4-10%	4-10%
Corporate	≈0.5%	≈0.5%	≈1.50%	≈1.50%	≈1.50%
Retail	0.1-0.9%	0.1-0.9%	0.2-0.5%	0.2-1%	0.2-1%

5 The k-sigma approach: calibration of k

The calibration of K coefficient, which identifies the upper bound of the confidence interval adopted for the general estimation error, is a crucial issue of the entire procedure for the MoC quantification. It is possible to evaluate various methodological alternatives but, in this paper, a scheme seemingly rationale and in general straightforward is privileged.
The starting point, as underlined above, derives from the choices adopted by the Basel Committee about the regulatory formula for the incorporation of prudence within internal rating systems parameters:

a) The PD confidence interval set at 99.9% "was also chosen to protect against estimation errors, that might inevitably occur from banks' internal PD, LGD and EAD estimation, as well as other model uncertainties. The confidence level is included into the Basel risk weight formulas and … used to provide the appropriately conservative value of the single risk factor" (see [29]);
b) The prudence perspective must involve also the LGD and EAD evaluation. "A bank must estimate an LGD for each facility that aims to reflect economic downturn conditions where necessary to capture the relevant risk. … banks may make reference to the average of loss severities observed during periods of high credit losses, forecasts based on appropriately conservative assumption, or other similar methods" (see ref. 235 of [30]). "… for exposures for which EAD estimates are volatile over the economic cycle, the bank must use EAD estimates that are appropriate for economic downturn, if these are more conservative than the long-run average" (see ref. 42 of [30]).

While for PD parameter a specific regulatory formula exists, in which the prudence is explicitly incorporated through a confidence interval, for LGD parameter, as well as for EAD, without specific regulatory formulas, the prudence is defined in terms of downturn adjustments with respect to the long run weighted averages of the same variables.

In this context, it appears evident that the MoC inclusion within PD, LGD and EAD estimates inevitably determines an increase in the overall confidence interval for the computation of the capital requirement. Nevertheless, as cited above, the 99.9% threshold for PD and the references to downturn scenarios for LGD and EAD appear to have been chosen by the Committee to indeed include uncertainties and errors within the models and the relative estimates of parameters. Apparently, it turns out to be a duplication, at least partial, of corrections for prudence. The underlying idea developed by the AIFIRM working group considers the MoC inclusion as making explicit a share of uncertainties and protections against errors already implicitly incorporated within regulatory formulas.

As specified in the general principles defined above, the adoption of MoC for general estimation error should not be a surreptitious tool to ask banks for an increase of their regulatory capital requirements.
Based on this principle, the approach adopted in the paper for the k calibration has the goal to reduce as much as possible the double counting of regulatory cautions, nonetheless recognizing that, with the aim of not altering the regulatory formulas, this double counting cannot be entirely eliminated. In what follows we outline the general features of the approach; for further details, we refer to the Position Paper.

By referring, for sake of simplicity, to Corporate segment, the regulatory formula can be expressed as:

$$\text{Capital requirement} = EAD_d \cdot \left\{ \Phi \left[\frac{\Phi^{-1}(PD) + \sqrt{\rho}\Phi^{-1}(0.999)}{\sqrt{1-\rho}} \right] \cdot LGD_d - PD \cdot LGD_d \right\} \cdot M$$

where EAD_d = EADdownturn, Φ = normal probability distribution, ρ = asset correlation, LGD_d = LGDdownturn, M = maturity adjustment.

By ignoring the maturity adjustment and working on the Total Loss (TL, loss at the quantile of the desired confidence interval corresponding to the sum of the expected loss and the requirement for the unexpected one), defined as the product of EAD, LGD and PD adjusted for prudence, it is possible to obtain the stressed total loss before the deduction for the expected loss:

$$\text{TL(R)} = EAD_d \cdot \left\{ \Phi \left[\frac{\Phi^{-1}(PD) + \sqrt{\rho} \cdot \Phi^{-1}(0.999)}{\sqrt{1-\rho}} \right] \cdot LGD_d \right\}$$

where TL(R) represents the Total Loss of Basel regulation.

If we want to avoid the double counting of regulatory cautions, the TL inclusive of MoC should be written as:

$$\text{TL(MoC)} = (\text{EAD} + k_E \sigma_{EAD}) \cdot \left\{ \Phi \left[\frac{\Phi^{-1}(PD + k_P \sigma_{PD}) + \sqrt{\rho_{|k_P \sigma_{PD}}} \cdot \Phi^{-1}(F)}{\sqrt{1-\rho_{|k_P \sigma_{PD}}}} \right] \cdot (LGD + k_L \sigma_{LGD}) \right\}$$

where EAD and LGD variables are computed before downturn adjustment and the macro factor (F) is left free to assume values lower than the threshold of the 99.9% regulatory confidence interval; the asset correlation, finally, takes into account the MoC impact on PD. Both F and the three k parameters have standard normal distribution. To avoid double counting, the value assumed by TL (MoC) in the 99.9% probability scenario [TL(MoC)|0.999] should not be higher than TL(R).

Therefore, the comparison introduced for the k calibration is between TL(R) and TL(MoC)|0.999.
To this aim two different paths can be followed: the first alternative concerns a solution in closed form (or partially closed), while the second one is based on a numerical solution through the Monte Carlo simulation.
The latter allows for more degrees of freedom compared to the solution in closed form, whose quality entirely depends on the goodness of the fit (in the tail) of the theoretical distribution applied to the effective one of TL(MoC). Thus, we summarize here only the Monte Carlo simulation method, referring to the Position Paper for the full discussion of both approaches.

We organized the Monte Carlo simulation as follows:

1) Two million random values with an underlying Uniform distribution are generated with a repetition for n run (n must be sufficiently high to guarantee stable results, in our case we use 2.000). The inverse of the normal standard distribution at 99.9 percentile is used as a cap for each simulated macroeconomic factor (the regulatory stressed macro scenario), thus $\Phi^{-1}(0.999) = 3.09023$
2) With the same approach, three i.i.d. normal standard random variables are generated independent from the macro factor. They are marked as k_E, k_L and k_P
3) Since the Regulation foresees that MoC should always be positive, the extraction of three normal standard random variables can be constrained to assume only positive values of k. This goal can be achieved either by adopting the absolute value or by generating a normal standard distribution truncated on 0, which is the preferred option in the process here described

4) Starting from the simulated scenarios of the four random variables, the TL(MoC) distribution is computed, from which it is possible to derive the 99.9th percentile. Since this percentile is very much inside the tail of the distribution, it is essential that the number of simulated scenarios is remarkable in order to fill the tail with a sufficient number of values and to stabilise the percentile value
5) After having obtained the 99.9th percentile of TL(MoC), the process continues by finding the single k value that, correspondent to the macro factor (F) above obtained, substituted to k_E, k_L and k_P and applied to the original PD, LGD and CCF values, maintains unchanged the value obtained of TL(MoC)|0.999.

In the following table we report, for each regulatory segment, the average *k* for the 2.000 runs of the process above described:

Table 3: average of *k*

Segment	Average
Low Default Portfolios	0.71-0.91*
Corporate	0.81
Retail	0.81-0.82*

*Depending on the analysed portfolio

Now, since the regulatory formula should not be altered, the k value above obtained must be inserted within the computation of TL(R) corrected for MoC:

$$\mathrm{TL(R+MoC)} = (\mathrm{EAD}_d + k\sigma_{EAD}) \cdot \left\{ \Phi \left[\frac{\Phi^{-1}(PD + k\sigma_{PD}) + \sqrt{\rho_{|k\sigma_{PD}}} \cdot \Phi^{-1}(0.999)}{\sqrt{1-\rho_{|k\sigma_{PD}}}} \right] \cdot (LGD_d + k\sigma_{LGD}) \right\}$$

The result, including the regulatory Macro factor and the Downturn adjusted LGD and EAD, is indeed higher than the TL(R) previously computed; the magnitude of the difference is quantified by generating a Macro factor distribution of 2 Million scenarios whose 99.9th percentile corresponds exactly to the regulatory one and consequently to the regulatory TL. By mapping the value of TL(R + MoC) on this distribution it is possible to derive the equivalent percentile reached by the MoC introduction combined with the Macro factor and the Downturn adjustments. Results are reported in detail in table 4, always computed as the average of the different runs and differentiated by regulatory segment: the differences with 0.9990 represent the impact of the double counting introduced by MoC.

Table 4: equivalent percentile of TL(R + MoC) on TL(R)

LDPs	Corporate	Retail
99.91-99.93	99.91	99.90-99.91

6 Summary of results and final remarks

In the previous paragraphs, we reported both methodologies and results obtained on realistic sample portfolios.
Given these ingredients, the calculation of *MoC=k*sigma* is obviously trivial. In what follows we report the impact analysis on performing exposures RWAs of the results obtained, using the sigma computed according to method 2 and assuming *k*=0.8 as per our main finding. For the purpose of sensitivity analysis, we also show the impacts with *k*=1, to be interpreted as an upper bound of the parameter in case of extremely conservative choice.
Table 5 shows that, as expected, the impacts are very burdensome for low default portfolios and manageable for retail. Corporate lies in the middle and, considering that it is for most bank the largest portfolio, it is likely to be representative of the total impact.
Taking also into account that we have chosen the less penalizing method for sigma, we conclude that a calibration of 0.8 for *k* is already quite punishing and lower values could be chosen, at least for low default portfolios. We recall once again that *k* has been calibrated assuming that estimation errors are always “on the wrong side”, and that is actually altering the 99.9% level of confidence stated by the Regulation.

Table 5: RWA impacts of MoC-C on an Italian large bank performing portfolio

	Delta RWA%	
	k=0.8	k=1
LDPs	3-8%	4-10%
Corporate	≈1.5%	≈2%
Retail	0.3-0.9%	0,3-1.1%

Before going to the conclusions, we want to address some operational issues.
First, we have worked under the implicit assumption that the level of application of MoC is the model segment and this coincides for all the risk components.
Nevertheless, it is easy to see that it is not the general situation. Take for example the case of a bank estimating PD for individuals at counterpart level, while LGD models are differentiated among Mortgage and Other Individuals (and EAD models possibly following a further segmentation): following the approaches described above, we should calibrate at least two different *k*, one for Mortgage and one for Other Individuals (ignoring EAD for sake of simplicity); but when we must apply the MoC to PD, we need to have a single *k* value.
Since we do not have a methodological solution for this issue, we can define a practical rule of aggregation, such as taking the average of obtained *k*, or conservatively the maximum of the two.

Another issue relates to application of MoC framework to the Defaulted Asset LGD and defaulted exposures RWAs. Regulation prescribes that RWAs on defaulted exposures must stem from the estimate of the increase of loss rate caused by possible additional unexpected losses during the recovery period, which should be broken down into three fundamental components: the downturn conditions component calibrated on the downturn adjustment to the long-run average LGD, any component covering for potential additional unexpected losses during the recovery period and the MoC component where the latter encompasses the three categories A, B and, above all, C.
The regulatory requests on defaulted assets clearly increase the importance of the MoC quantification since all the three categories defined by the *EBA* must be part only of the LGD in-default and not of the ELBE. Therefore, together with the Downturn, they characterize the unexpected loss component which, multiplied by 12.5 times the EAD, determine the RWA figures.
The Method 2 approach ("within variance") presented for the sigma definition on LGD (the unique risk parameter significant for defaulted assets valuation) can be applicable as well to defaulted exposures where a further issue can be represented by the granularity introduced by the reference date LGD observation (i.e. the differentiation of the risk parameter by the time spent in default). In fact, the regulatory requirements for defaulted assets LGD estimation underline the need to analyse the LGD by time in-default and recoveries realised so far, determining at least a further axis of analysis with respect to Performing LGD. The greater granularity together with the increase of the number of observations (i.e. the same facility is repeated by the years spent in default) cause a different result in terms of sigma from the Performing LGD estimation. More problems concern the application of Method 1 (bootstrapping), unless a constraint is imposed for the facility extraction on the combination of risk drivers (the time spent in default causes significant differences in LGD outcomes).
For the default statuses before the litigation process, the same methodological issues underlined for model components approach on Performing LGD can be extended. On the other side, given the framework introduced for the *k* calibration and the impossibility to replicate it on the formula for defaulted exposures, the proposal is to adopt the same *k* obtained for performing exposures of the same calibration segment under evaluation.
Again, the risk of imposing excessive and undue burdens to banks must be carefully considered in the calibration of parameters; the request to apply MoC only to the LGD in-default is an additional source of prudence with respect to the Downturn component which originally represented the add-on for unexpected loss.

A further operational issue to tackle is the lifecycle of our type C MoC. Clearly the sigma depends on the models and their development samples, thus they should be recomputed every time the model is updated. The *k* parameter, on the other hand, is a more structural parameter, depending on the methodology chosen for defining the MoC, rather than the features of the model.
Therefore, a convenient choice can be fixing the values of *k* once and for all at the moment of the definition of the MoC policy of the bank. In this sense, the MoC will be part of the periodic updating of the time series, not requiring any validation activity apart from the classification of the change itself.
The *k* could possibly be redefined only in case of structural re-design of the models, involving an update of the whole MoC policy. This would normally happen in the event of changes classified as model changes, thus requiring a full validation activity.

Fortunately, the results obtained for *k* calibration seem to simplify our issues: as we saw in paragraph 4, *k* is quite stable across the spectrum of sigma input values we used in the simulations, around a value of 0.8. Furthermore, this result suggests that *k* is stable whether type C MoC is the only one included or if it is added with other type A and B MoC.
In conclusion, therefore, a value of *k* in that range can serve the purpose of defining the type C MoC, encompassing the objectives of both conservativeness and methodological soundness.

Silvio Cuneo e Franco Varetto

AIFIRM

RISK MANAGEMENT MAGAZINE

Anno 14, numero 1 **Gennaio – Aprile 2019**

Poste Italiane - Spedizione in abbonamento postale – 70% aut. DCB / Genova nr. 569 anno 2005

TESTATA INDIPENDENTE CHE NON PERCEPISCE CONTRIBUTI PUBBLICI (legge 250/1990)

IN QUESTO NUMERO

ARTICOLI A CARATTERE DIVULGATIVO

ARTICOLI A CARATTERE SCIENTIFICO
(sottoposti a referaggio)

Risk Management Magazine
Anno 14, n° 1 Gennaio - Aprile 2019

Direttore Responsabile:
Maurizio Vallino

Condirettore
Corrado Meglio

Proprietà, Redazione e Segreteria:
Associazione Italiana Financial Industry Risk Managers (AIFIRM), Via Sile 18, 20139 Milano

Registrazione del Tribunale di Milano n° 629 del 10/9/2004

ISSN 2283-7329

E-mail: amministrazione@aifirm.it; Tel. 389 6946315

Stampa: Algraphy S.n.c. - Passo Ponte Carrega 62-62r 16141 Genova

Le opinioni espresse negli articoli impegnano unicamente la responsabilità dei rispettivi autori

SPEDIZIONE IN ABBONAMENTO POSTALE AI SOCI AIFIRM RESIDENTI IN ITALIA, IN REGOLA CON L'ISCRIZIONE

Rivista in stampa il 2 Aprile 2019

ACCADEMIA ITALIANA DI ECONOMIA AZIENDALE

Rivista accreditata AIDEA