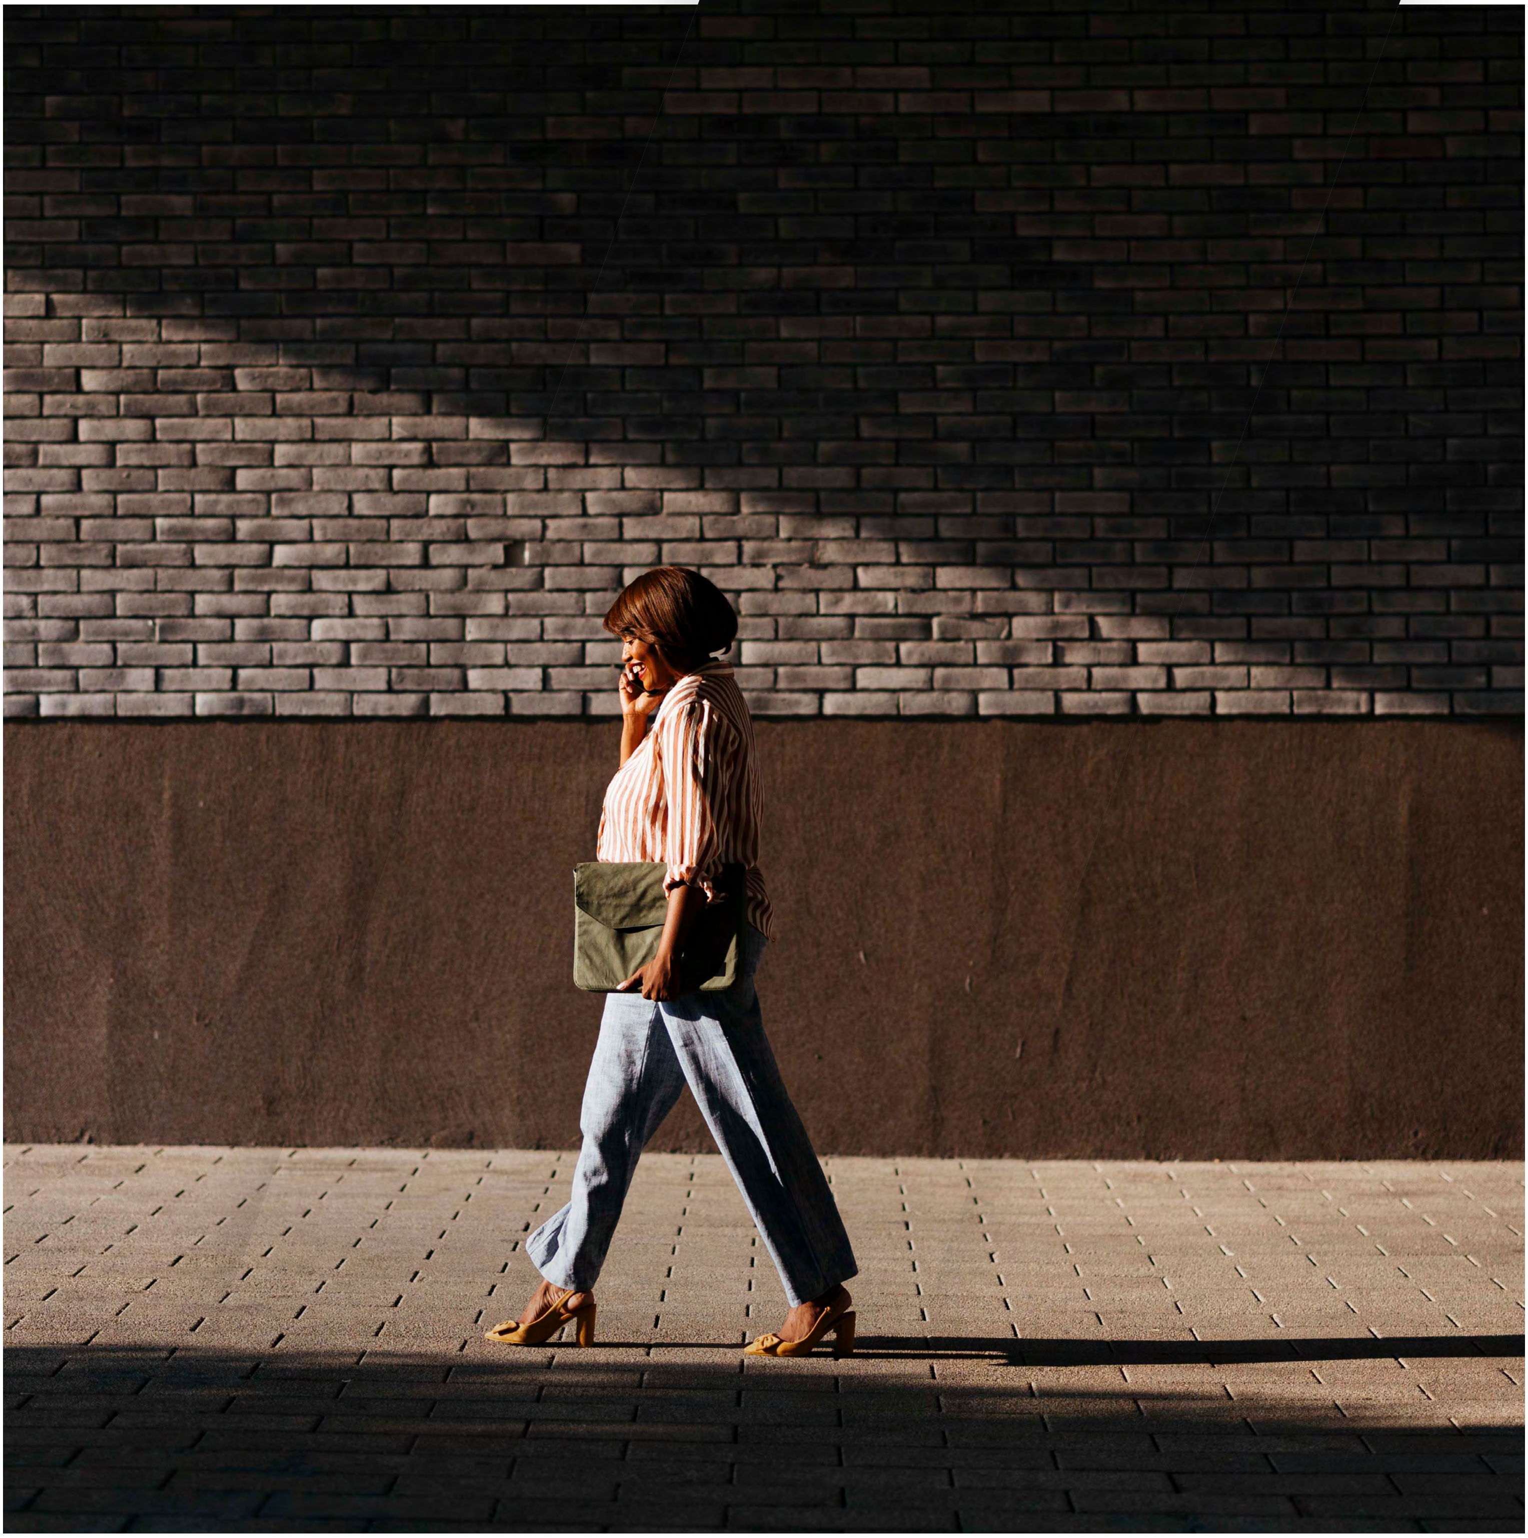

Quantitative take on double default modelling

Contents

1 Introduction

Credit risk modelling is fundamentally concerned with the distinction between probability and outcome. A probability of default describes the likelihood of an event, but it does not determine whether that event will occur. When an exposure is guaranteed, the problem becomes more complex: the relevant risk depends not only on the creditworthiness of the obligor and guarantor individually, but also on the likelihood that both fail within the relevant horizon.

Different players have different expectations. Regulators want models that lean cautious, ensuring banks can weather financial storms while investors and accountants look for models that reflect real economic risk, not just worst-case scenarios.

These two objectives, prudence vs realism, do not always align, though both frameworks recognize that some degree of conservatism may be warranted in the face uncertainty. One area where this difference is especially clear is in the treatment of credit guarantees.

For example, under regulatory capital frameworks, the effect of an eligible guarantee on the protected portion is generally recognised by reference to the protection provider's risk characteristics, subject to the applicable approach, eligibility requirements and regulatory floors. This provides a relatively simple and prudentially constrained treatment, but it does not explicitly model the joint relationship between obligor default and protection-provider performance.

Under IFRS 9, one is not limited to simplified treatments such as Probability of Default (PD) substitution. Instead, the standard states that:

"An entity shall measure expected credit losses of a financial instrument in a way that reflects an unbiased and probability-weighted amount".

It is important to note that "unbiased" in this context means avoiding deliberate over- or understatement of expected losses, not ignoring parameter uncertainty or model risk. Indeed, IFRS 9 requires consideration of "reasonable and supportable information" including multiple scenarios where appropriate. This opens the door to more refined approaches that better reflect the economic substance of credit enhancements such as the double default approach for guaranteed exposures. However, a joint-default probability should not be interpreted as a reduction in the obligor's standalone PD; rather, it is a modelling input intended to capture the combined occurrence of obligor default and guarantor non-performance on the protected portion.

This paper focuses on the quantitative modelling of double default. In practice, any proposed application should also be assessed through an accounting lens to determine whether, and to what extent, the relevant guarantee may be reflected in the measurement of Expected Credit Losses (ECL) under the applicable requirements.[1]

This paper is not an academic exercise. It is intended for risk practitioners, model developers & validators, and anyone interested in understanding the principles behind double default modelling and how it can be applied in practice. Our goal is to provide a balanced, technically grounded, and actionable overview, including both the benefits and limitations of this approach, rather than a full theoretical or empirical treatment.

[1] For example, IFRS 9 paragraph B5.5.55 addresses the circumstances in which cash flows expected from collateral and other credit enhancements are reflected in expected cash shortfalls.

2 Regulatory boundaries

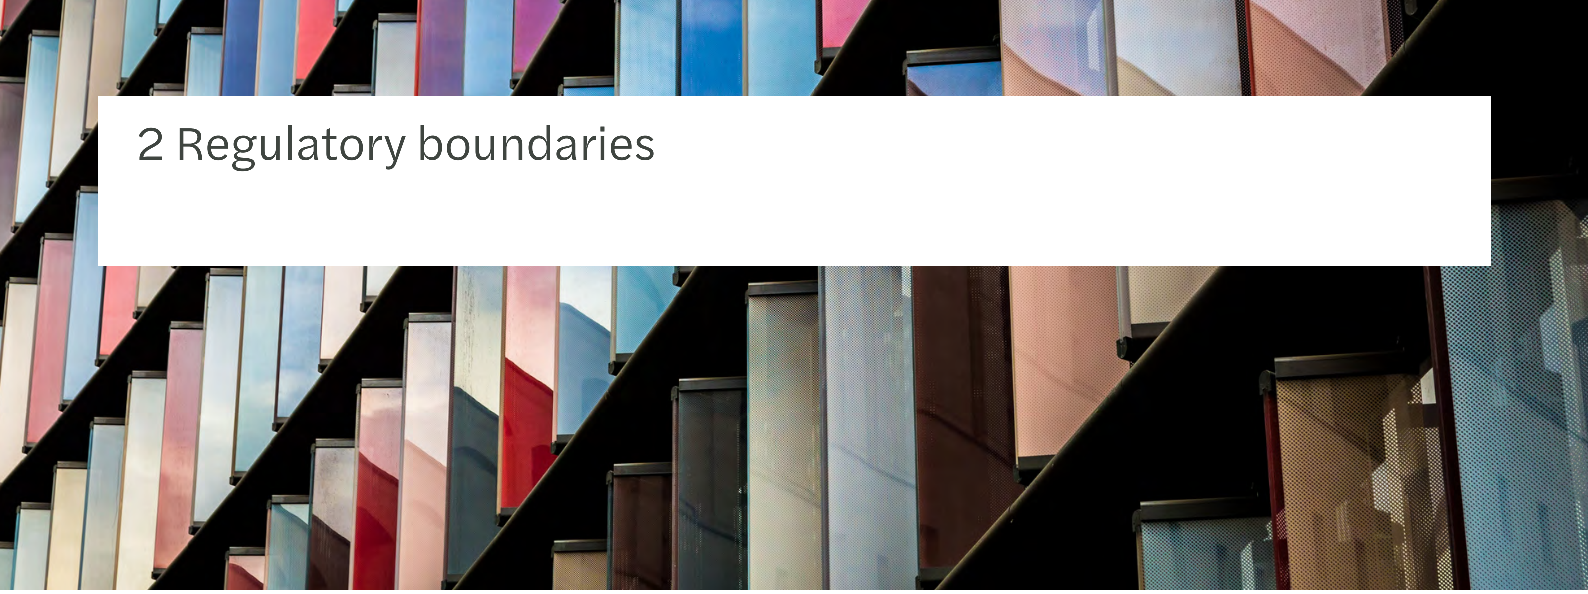

Before exploring the mechanics of double default modelling, it is important to distinguish its use for expected credit loss modelling from its treatment under regulatory capital frameworks.

The regulatory treatment has evolved over time. During the development of Basel II, the QIS3 FAQ (Section E)[2] required institutions to disregard double-default effects when calculating capital requirements, noting that the magnitude of the effect depends significantly on the correlation between the obligor and the guarantor. A dedicated double-default treatment was subsequently introduced under the Basel II framework, and reflected in the earlier EU Capital Requirements Regulation, for certain guaranteed exposures and credit derivatives meeting prescribed eligibility and operational requirements.

The final Basel III framework has moved away from that dedicated treatment[3]. Under the current Basel framework, the recognition of guarantees and credit derivatives must not reflect the effect of double default. The resulting risk weight for the protected exposure cannot be lower than the risk weight of a comparable direct exposure to the protection provider.

The EBA's 2024[4] report on credit insurance provides a current European perspective on this development. The report explains that CRR3 removes the previous dedicated double-default treatment and reinforces the role of the risk-weight floor for unfunded credit protection. As a result, institutions cannot recognise, for regulatory capital purposes, the potentially risk-reducing effect arising from having recourse to both the obligor and the protection provider. The EBA also notes that this is not specific to credit insurance: double default and dual recourse are features of unfunded credit protection more generally, including bank guarantees.

The regulatory position should not, however, be interpreted as suggesting that dependence between an obligor and a guarantor is economically irrelevant. Rather, the prudential framework deliberately limits the extent to which this effect may reduce capital requirements reflecting the objective of simplicity and conservatism. The quantitative relationship remains relevant when considering the economic behaviour of a guarantee and provides the basis for the modelling framework explored in this paper.

[2] QIS3 FAQ: E. credit risk mitigation

[3] CRE32 - IRB approach: risk components

[4] Report on credit insurance.pdf

3 Guarantees: substitution vs joint default

Guarantees serve as a form of credit enhancement by introducing a second source of payment if the obligor fails to meet its contractual obligations. A conventional guarantee triggered by obligor non-payment does not necessarily change the obligor's standalone probability of default; rather, it may reduce the lender's expected cash shortfall through recovery from the guarantor.

- **Collateral** affects the Loss Given Default (LGD) by providing a security that can be recovered after a default, reducing the severity of loss.
- **Guarantees** can impact both PD and LGD depending on their structure. The impact on LGD depends on factors such as guarantee coverage ratio, payment timing, and the guarantor's financial condition at the time of claim.

In this section, we compare the substitution and the joint default approaches for guarantees and how they affect the measurement of credit risk, and PD in particular.

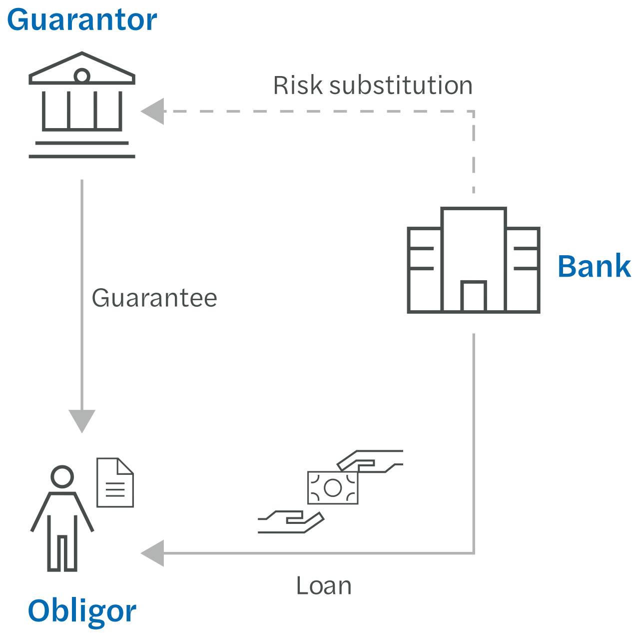

3.1 The substitution approach

The substitution approach recognises credit protection by applying the protection provider's risk characteristics to the covered portion of an eligible exposure, subject to the rules of the relevant framework. The uncovered portion continues to be treated by reference to the underlying obligor. Its precise implementation may involve more than replacing one PD with another, including the application of the relevant risk-weight function, LGD treatment and coverage adjustments.

The approach is comparatively straightforward and avoids the need to estimate an explicit obligor-guarantor dependence parameter. However, it does not separately model the joint occurrence of obligor default and guarantor non-performance: it assumes the guarantor's performance is unconditional and immediate; an oversimplification in many real-world situations. Under a joint default view, a loss materializes only if the obligor defaults and the guarantor fails to perform, so the substitution approach represents a conservative simplification of the true workings of guaranteed exposures.

3.2 The joint default approach

The joint default method offers a more refined perspective. In the stylised case of a full and effective guarantee, a credit loss is modelled to arise when the obligor defaults and the guarantor is unable to perform its contractual obligation. This captures the conditional nature of the protection: default of the obligor triggers recourse to the guarantor, while loss on the protected portion depends on the guarantor's ability to perform.

3 Guarantees: substitution vs joint default

This representation is intentionally simplified to isolate the double default effect. In practice, the amount and timing of any loss may also be affected by the percentage of the exposure covered, payment delays, contractual exclusions, recovery costs and other legal or operational factors. These features would normally be reflected elsewhere in the ECL calculation, including through the estimated recovery cash flows or LGD.

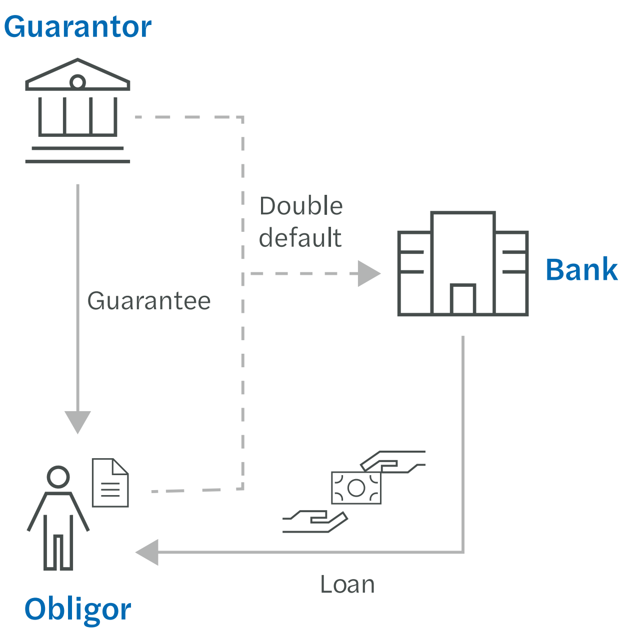

In this case, we need to model both PDs separately and account for how closely linked the two parties are. This is where correlation comes in. Techniques such as copulas are used to combine distributions into a joint distribution, reflecting the economic linkage between the two entities. The result is a model that attempts to capture the economic sense of guarantees, especially where the guarantor's creditworthiness is not independent of the obligor's.

Moreover, this approach can be economically advantageous for institutions. Since the joint probability of obligor default and guarantor non-performance cannot exceed either marginal probability, a joint default approach may produce a lower modelled probability of loss than a substitution approach, depending on the dependence (correlation) assumption and the other features of the guarantee. It should be noted, of course, that this reduction in ECL is only appropriate if:

- The correlation parameter is accurately estimated;
- The independence assumption between defaults (conditional on the modelled factors) holds;
- Legal and operational risks are minimal;
- The model is properly validated against realized losses.

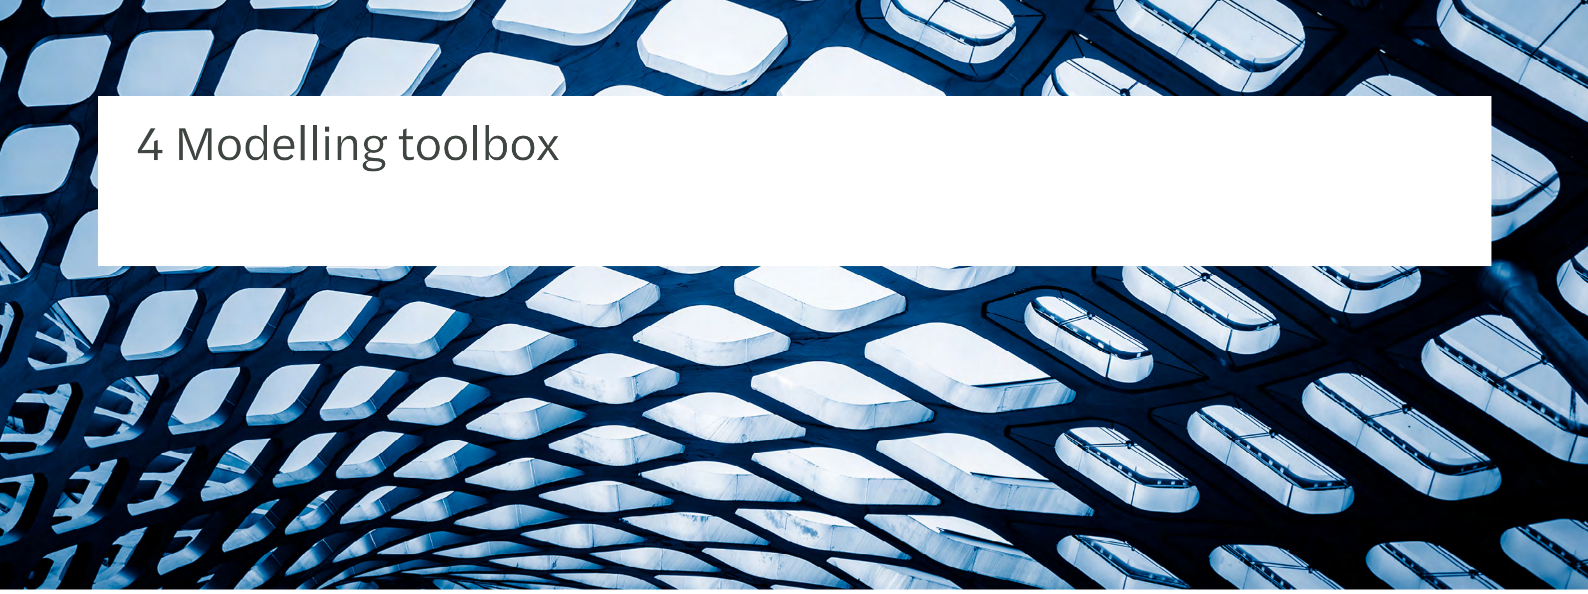

4 Modelling toolbox

Having established the general framework for joint default modelling, we can now turn to the quantitative component that makes it work.

4.1 Copulas

A key component of double default modelling is how to address the question of dependence. How do we combine two marginal default probabilities into a single, joint default probability that reflects the economic relationship between obligor and guarantor?

Copulas provide the mathematical framework to do exactly this. A copula is a function that links marginal distributions to form a joint distribution while preserving their individual properties.

The Gaussian Copula

The Gaussian copula[5] is a commonly applied tool in credit risk. For two entities with marginal default probabilities PD_A and PD_B, and correlation ρ, the joint probability of default can be expressed as:

$$Joint\ PD = \Phi_2(\Phi^{-1}(PD_A), \Phi^{-1}(PD_B), \rho)$$

Where Φ^{-1} is the inverse of the standard normal distribution and Φ_2 is the cumulative distribution of the bivariate normal with correlation ρ.

The figure below illustrates this relationship by fixing the borrower's PD at 5% and the guarantor's PD at 0.5%, and varying the correlation parameter ρ. As the correlation increases from 0 toward 1, the joint probability rises smoothly, moving from the independence case (approximately $PD_A \times PD_B$) towards the lower of the two marginal PD.

This simple experiment highlights why the dependence structure cannot be ignored when modelling guarantees. Even with unchanged marginal PDs, the assumed correlation can substantially alter the joint default probability. As the correlation increases (e.g. when a parent company provides a guarantee for its subsidiary), the two defaults become increasingly linked.

In the extreme case of perfect correlation, the joint probability of default converges to the lower of the two marginal PDs, which is, in practice, the guarantor's PD. This shows that the substitution approach is not fundamentally different, but rather the limiting case of the joint default framework under perfect dependence.

4.2 Default correlation

As we have seen in the previous section, in a double-default setting, loss occurs only if the obligor and the guarantor are both in distress. The degree of correlation (or dependence) between them therefore has a direct impact on the joint PD.

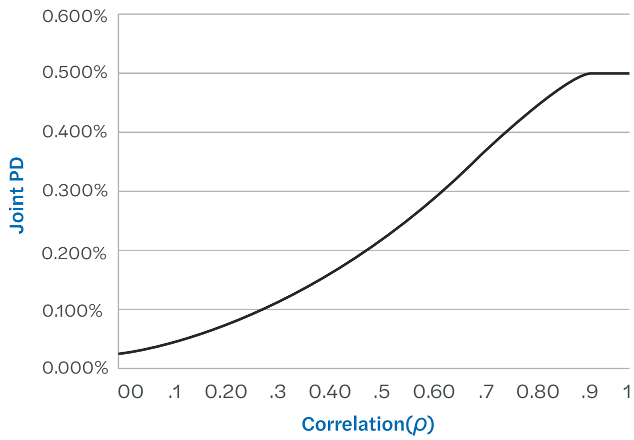

[5]This copula formulation was popularized in the CreditMetrics framework (JP Morgan, 1997) and later used in David X. Li's famous 2000 paper "On Default Correlation: A Copula Function Approach"

Estimating this correlation is not straightforward. Unlike marginal PDs, which can be derived from ratings, market spreads, or internal models, there is rarely direct data on joint defaults. Instead, practitioners must rely on their Quant toolbox, using proxies and specific calibration techniques to estimate dependence. Some of the common approaches include:

- **Market-implied measures**, extracted from instruments sensitive to joint credit risk (e.g. basket CDS), where data is available.
- **Historical proxies**, typically based on market observables such as equity returns or credit spreads. Equity prices in particular are often used as a proxy for the "health" or survival strength of a company, making their co-movements a natural indicator of potential joint stress.
- **Structural or macro-based estimates**, where correlation is linked to shared exposure to systematic risk factors such as sector, region, or parent–subsidiary relationships.

Even within these methods, the choice and interpretation of the statistical measure matter. Pearson correlation may be used in a Gaussian copula when the dependence is based on approximately linear co-movement in suitably transformed variables. But market data does not necessarily behave in a perfectly Gaussian way. Rank-based measures such as Spearman correlation or Kendall's Tau are often more robust, since they capture the strength of monotonic dependence without being distorted by outliers or distributional shape. On their own, they do not measure tail risk, but they provide a solid basis for calibrating copulas that can incorporate tail dependence, such as the t-copula. This is particularly important when considering the convex behaviour of ECL models, as detailed in our article: [Addressing tail risk in expected credit loss models.]

Practitioners should be particularly aware that correlation parameters may be unstable over time and tend to increase during stress periods (e.g. considering the increasing default correlations during the 2007-2008 GFC). This phenomenon requires careful consideration in model implementation and validation.

Default dependence deserves a dedicated paper in its own right. Proxy selection, data windows, statistical transformations, stress behaviour, wrong-way risk and the choice of copula family can each materially affect the estimated joint PD. A future paper will explore these calibration questions in greater depth. For present purposes, the key point is that correlation is not merely a sensitivity parameter: it is a model input requiring a clear economic interpretation, empirical support and appropriate validation.

4.3 IFRS 9 integration

To illustrate the practical implications of joint default modelling in an ECL context, we conducted an impact analysis on a corporate loan portfolio. The sample contains 120 loans with a total exposure of €2.3bn, of which 99 are in Stage 1 and 21 in Stage 2. Loan sizes vary between €0.5m and €150m. The portfolio carries a relatively low coverage ratio of 0.5%, so these results should not be used for benchmarking purposes without further context. They do, however, provide useful insights into the mechanics of substitution versus joint default.

The figure below compares ECL under substitution with those under joint default, using average correlation assumptions of 0%, 25%, 50%, and 100%. For the purposes of this illustrative analysis, the substitution comparator is represented by the lower of the obligor and guarantor PDs. This is a simplified study convention designed to isolate the effect of the PD methodology; it should not be interpreted as a prescribed regulatory or IFRS 9 substitution method. In the small number of cases where the obligor PD was lower than the guarantor PD, the obligor PD was therefore retained. EAD and LGD were held constant across all scenarios so that the results isolate the effect of the selected PD treatment.

4 Modelling toolbox

The first clear observation is the material difference between the simplified substitution proxy and the joint default scenarios considered. In this illustrative portfolio, the substitution proxy produces materially higher ECL because it applies the lower marginal PD directly, whereas the joint default approach estimates the probability of obligor default and guarantor non-performance occurring within the relevant modelling horizon. Within the joint default framework, the dependence assumption is also a material driver: lower correlation assumptions produce lower joint PDs and, consequently, lower ECL.

The results are presented separately for the Stage 1 and Stage 2 sub-portfolios. In both cases, the simplified substitution proxy produces higher ECL than the joint default scenarios considered, with the proportional difference being more pronounced for Stage 1.

It is noted, however, that the two stages contain different exposures. The observed difference may therefore also reflect their respective PD distributions, exposure profiles, guarantor characteristics and obligor-guarantor dependence assumptions. It should therefore not be attributed solely to the distinction between 12-month and lifetime ECL.

The mechanics can be illustrated using the example in Section 4.1. If the obligor PD is 5% and the guarantor PD is 0.5%, the joint PD under independence (ρ=0%) is approximately 0.025%. By contrast, the simplified substitution proxy applies the lower marginal PD of 0.5%, which is twenty times the joint PD under independence.

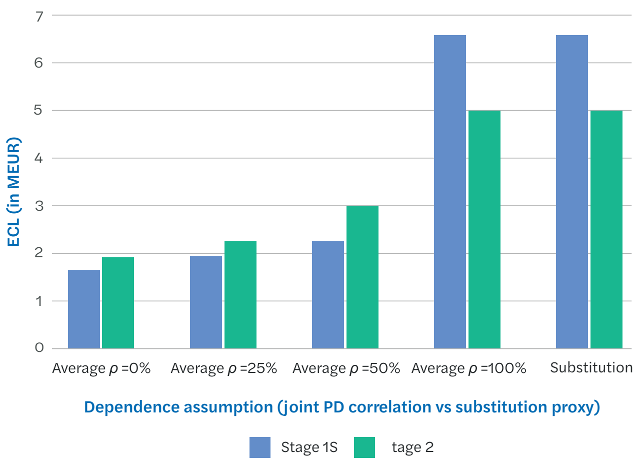

For the Stage 2 sub-portfolio, the proportional difference is smaller but remains material. This is a portfolio-specific observation and should not be interpreted as evidence that the effect of joint default modelling is inherently smaller for Stage 2 exposures.

For a more granular ECL calculation on longer-dated Stage 2 exposures, a single lifetime cumulative joint PD would generally be insufficient because it would not identify when the joint credit event is expected to occur. A more granular implementation would construct obligor and guarantor marginal or incremental PD term structures over the remaining life of the exposure and apply a consistent dependence framework at each time interval. The resulting time-specific probabilities of joint failure could then be combined with the corresponding EAD, expected recoveries and discount factors. A cumulative PD may still describe the probability of default by a specified horizon, but using it as the sole input to a lifetime ECL calculation can obscure the timing of loss and the evolution of dependence over the life of the exposure.

5 Conclusion

The purpose of this paper has been to provide a balanced assessment of double default modelling, with a particular focus on its relevance under IFRS 9. Guarantees are a central feature of corporate lending, yet their treatment in risk models often falls back on substitution. This method is simple but represents a simplification of the true mechanics of loss. While regulatory capital frameworks deliberately exclude double default effects, IFRS 9 sets a different standard: it requires that expected credit losses be measured in an unbiased and probability-weighted way. This creates scope for institutions to consider whether a more risk-sensitive methodology would provide a better-supported representation of expected cash shortfalls arising from guaranteed exposures.

We have walked through the core building blocks of the methodology: copulas as the mathematical link between PDs, and default correlation as the critical (and most challenging) parameter to estimate. Both areas warrant deep study, particularly correlation, which can fundamentally alter joint default outcomes and is itself a subject of ongoing research.

Our limited impact analysis on a small corporate loan portfolio demonstrated the potential magnitude of these effects. The simplified substitution proxy produced higher ECL than the joint default scenarios considered in both sub-portfolios, with the proportional difference being more pronounced for Stage 1. However, because the Stage 1 and Stage 2 populations contain different exposures, this result should not be attributed solely to the distinction between 12-month and lifetime ECL. Across both sub-portfolios, lower dependence assumptions produced lower joint PDs and correspondingly lower ECL under the joint default approach.

In short, while double-default benefits cannot be recognised for regulatory capital where prohibited by the applicable prudential framework, a joint-default framework may provide one possible means of modelling expected cash shortfalls associated with a qualifying guarantee, provided that the approach faithfully represents the contractual coverage, timing and probability of guarantee recoveries and satisfies the broader IFRS 9 requirements for unbiased, probability-weighted ECL.

For banks, this is not just a theoretical exercise but a tangible opportunity: better models, better risk capture, and potentially better financial outcomes.

Contacts

Xavier Larrieu
Partner
Co-head of Quantitative Solutions UK
T: +44 7881 283791
E: xavier.larrieu@mazars.co.uk

Andreas Mavrocordatos
Director
Quantitative Solutions
T: +44 7483 662604
E: andreas.mavrocordatos@mazars.co.uk