

Regime-Switching Asset Correlation in Hybrid Structural–Intensity Credit Risk Models: A Monte Carlo Study of Portfolio Tail Risk

Peter Mutua Malonza¹

¹Independent Researcher, Kitengela, Kenya

September 2026

Abstract

Standard credit portfolio risk models typically assume a constant asset-return correlation across the economic cycle. This assumption is increasingly difficult to reconcile with empirical evidence that default correlation rises sharply during periods of macroeconomic stress. This paper develops a hybrid structural–intensity credit risk model in which a firm’s default intensity is driven endogenously by its simulated distance-to-default, and in which the systematic correlation between firms’ asset returns follows a two-state Markov regime-switching process (a “normal” regime and a “stress” regime). Using a Monte Carlo simulation of a 100-firm equally-weighted portfolio over a five-year horizon, we compare the resulting loss distribution to a benchmark model with constant correlation calibrated to the same long-run average. We find that while the two specifications produce nearly identical expected losses (6.39% vs. 6.47% of notional), the regime-switching specification produces materially fatter tails: value-at-risk at the 99.9% confidence level is approximately 23% higher, and expected shortfall at the 97.5% level is approximately 15% higher, than under the constant-correlation benchmark. These results are consistent with the procyclicality concerns raised in the Basel III and IFRS 9 literature, and suggest that point-in-time capital models calibrated only to average correlation may materially understate tail capital requirements. We discuss implications for regulatory capital adequacy and expected credit loss (ECL) provisioning, and outline extensions including empirical calibration and factor-based regime identification.

Keywords: credit risk, structural models, intensity models, default correlation, regime-switching, Monte Carlo simulation, portfolio credit risk, Basel III, IFRS 9

This is a working paper / preprint. It has not been submitted to a peer-reviewed journal. All empirical results reported are generated from a simulation study using synthetic data; no proprietary or confidential market data is used.

1 Introduction

The 2008 global financial crisis and, subsequently, the COVID-19 shock and the 2022–2023 interest-rate tightening cycle each demonstrated a common empirical pattern in credit markets:

corporate defaults cluster in time far more than a model of independent, or even constantly-correlated, default events would predict (Das et al., 2007; Duffie et al., 2009). This clustering has direct consequences for the tail of a credit portfolio’s loss distribution, and therefore for the capital a bank or investor must hold against unexpected losses.

Two broad families of single-name credit risk models dominate the literature. *Structural models*, originating with Merton (1974) and extended by Black and Cox (1976), treat default as endogenous: a firm defaults when the market value of its assets falls below a liability-related barrier. *Reduced-form (intensity) models*, developed by Jarrow and Turnbull (1995) and Duffie and Singleton (1999), instead treat default as an exogenous jump process governed by a stochastic hazard rate, calibrated to observed credit spreads. Each approach has complementary strengths: structural models offer an economically intuitive default trigger tied to firm fundamentals, while intensity models are more tractable for pricing and can flexibly accommodate time-varying and correlated hazard rates.

A smaller strand of the literature combines the two approaches into *hybrid* models, in which the default intensity is itself a function of a latent or observable structural state variable (e.g. distance-to-default), rather than being specified purely exogenously. This paper adopts such a hybrid specification, in the spirit of the broader literature on frailty and contagion in default correlation (Giesecke, 2004; Duffie et al., 2009), but focuses specifically on a dimension that has received comparatively less attention in the applied practitioner literature: *regime-dependent* systematic correlation.

The workhorse model for regulatory credit portfolio capital, the asymptotic single risk factor (ASRF) model underlying the Basel II/III internal-ratings-based (IRB) capital formulas, is built on the single-factor Gaussian copula framework of Vasicek (2002), in which a single correlation parameter governs the co-movement of all obligors’ asset returns. In practice, that correlation parameter is typically calibrated to a long-run average. If true correlation instead varies with the state of the economy – rising in downturns and falling in expansions – then a model calibrated only to the average will be well-specified for the body of the loss distribution but misspecified in exactly the tail region that determines regulatory capital and economic capital buffers.

This paper makes three contributions. First, it specifies a tractable hybrid structural–intensity model with a two-state Markov regime-switching correlation structure that nests the constant-correlation benchmark as a special case. Second, it implements a full Monte Carlo simulation of the model for a stylised but realistic 100-firm portfolio and quantifies the difference in loss-distribution tail risk between the regime-switching specification and a constant-correlation benchmark calibrated to match the same long-run average correlation and expected loss. Third, it discusses the implications of these findings for Basel III capital adequacy and IFRS 9 expected credit loss (ECL) provisioning, drawing on the author’s applied background in prudential risk and compliance.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 specifies the model. Section 4 describes the simulation design. Section 5 presents results. Section 6 discusses regulatory implications. Section 7 concludes and outlines extensions.

2 Literature Review

2.1 Structural models

Merton (1974) models a firm’s equity as a call option on its assets, with default occurring if the terminal asset value falls below the face value of debt at maturity. Black and Cox (1976) extend this to a first-passage-time formulation in which default can occur at any time the asset value process breaches a (possibly time-varying) barrier, which is the formulation adopted in the present paper’s structural component.

2.2 Reduced-form / intensity models

Jarrow and Turnbull (1995) model default as the first jump of a Poisson process with constant intensity, calibrated from observed credit spreads. Duffie and Singleton (1999) generalise this to a doubly stochastic (Cox) process in which the intensity itself follows a stochastic process, enabling richer term-structure dynamics and, importantly for portfolio applications, correlated intensities across obligors.

2.3 Default correlation and clustering

Vasicek (2002) derives the loss distribution of a large homogeneous portfolio under a single-factor Gaussian copula, which remains the analytical basis for the Basel IRB capital formula. Das et al. (2007) provide empirical evidence that observed default clustering in U.S. corporate data exceeds what a doubly stochastic model with constant intensity correlation can explain, motivating frailty and contagion channels. Giesecke (2004) and Duffie et al. (2009) develop models in which unobserved common (“frailty”) factors or direct contagion mechanisms generate additional default clustering beyond that implied by observable covariates.

2.4 Regime-switching and procyclicality

The prudential regulation literature has separately raised concerns about the procyclicality of both IRB capital requirements and, more recently, IFRS 9 expected-credit-loss provisioning, insofar as both are sensitive to point-in-time versus through-the-cycle parameter estimation. This paper connects the asset-correlation regime-switching literature directly to that procyclicality discussion by quantifying, within a single simulation framework, how much of the potential understatement of tail risk is attributable specifically to the assumption of constant correlation, holding the average correlation and average default rate fixed.

3 Model Specification

3.1 Portfolio and structural process

Consider a portfolio of N obligors. For obligor $i \in \{1, \dots, N\}$, define a standardised distance-to-default process $D_i(t)$, evolving in discrete time steps of length Δt as

$$D_i(t + \Delta t) = D_i(t) + \left(\mu - \frac{1}{2}\sigma^2\right) \frac{\Delta t}{\sigma} + \varepsilon_i(t)\sqrt{\Delta t}, \quad (1)$$

where μ and σ are the drift and volatility of the (log) asset-return process and $\varepsilon_i(t)$ is a standardised shock described below. Structural default occurs the first time $D_i(t)$ crosses a barrier $\bar{D} = 0$ (a Black–Cox-style first-passage default), i.e. default at

$$\tau_i^{\text{struct}} = \inf\{t : D_i(t) \leq \bar{D}\}. \quad (2)$$

3.2 Systematic factor and regime-switching correlation

The standardised shock is generated by a one-factor model,

$$\varepsilon_i(t) = \sqrt{\rho(t)} Z(t) + \sqrt{1 - \rho(t)} \eta_i(t), \quad (3)$$

where $Z(t)$ is a common systematic factor, $\eta_i(t)$ is obligor-specific idiosyncratic noise, both i.i.d. standard normal and mutually independent, and $\rho(t) \in (0, 1)$ is the time-varying systematic asset correlation. The correlation follows a two-state Markov chain $S(t) \in \{\text{Normal}, \text{Stress}\}$,

$$\rho(t) = \rho_N \mathbf{1}\{S(t) = \text{Normal}\} + \rho_S \mathbf{1}\{S(t) = \text{Stress}\}, \quad \rho_S > \rho_N, \quad (4)$$

with transition probabilities per period $p_{N \rightarrow S}$ and $p_{S \rightarrow N}$ governing, respectively, the expected time between stress episodes ($1/p_{N \rightarrow S}$) and the expected duration of a stress episode ($1/p_{S \rightarrow N}$). The constant-correlation benchmark model is nested as the special case $\rho_N = \rho_S$.

3.3 Hybrid default intensity

In addition to the structural default trigger, each surviving obligor is exposed to a hazard-driven default arrival with intensity coupled to its current distance-to-default,

$$\lambda_i(t) = \lambda_0 + \kappa \cdot \max(0, D^* - D_i(t)), \quad (5)$$

where λ_0 is a baseline hazard rate applying even to firms far from default, $\kappa > 0$ governs the sensitivity of the hazard to deteriorating fundamentals, and D^* is the distance-to-default level below which the hazard begins to rise. This specification nests a pure reduced-form model in the limit $\kappa \rightarrow 0$ (constant intensity) and approaches a pure structural model as $\kappa \rightarrow \infty$ (default triggered essentially exactly at the structural barrier). In discrete time, conditional on survival to t , obligor i defaults via the hazard channel in $(t, t + \Delta t]$ with probability

$$p_i(t) = 1 - \exp(-\lambda_i(t) \Delta t). \quad (6)$$

The obligor's overall default time is $\tau_i = \min(\tau_i^{\text{struct}}, \tau_i^{\text{hazard}})$.

3.4 Portfolio loss

Each obligor carries equal exposure $1/N$ and a fixed loss-given-default $\ell \in (0, 1)$. Portfolio loss over horizon T is

$$L(T) = \frac{\ell}{N} \sum_{i=1}^N \mathbf{1}\{\tau_i \leq T\}. \quad (7)$$

4 Simulation Design

The model is implemented via Monte Carlo simulation with $N = 100$ obligors, a five-year horizon ($T = 5$) discretised into monthly steps ($\Delta t = 1/12$), and 30,000 independent portfolio-level scenarios. Table 1 reports the parameterisation. Correlation and hazard parameters are chosen to be broadly representative of an investment-grade-tilted corporate portfolio rather than calibrated to any specific proprietary dataset; as such, the results below should be read as illustrating the *qualitative and relative* effect of regime-switching correlation on tail risk, not as point estimates of any particular institution’s capital requirement.

Table 1: Model parameters used in the simulation

Parameter	Symbol	Value
Number of obligors	N	100
Horizon	T	5 years
Time step	Δt	1/12 year
Monte Carlo scenarios	–	30,000
Asset return drift	μ	2.0% p.a.
Asset return volatility	σ	25.0% p.a.
Initial distance-to-default	$D_i(0)$	4.0
Hazard-trigger distance	D^*	2.0
Baseline hazard	λ_0	0.4% p.a.
Hazard sensitivity	κ	0.35
Normal-regime correlation	ρ_N	0.15
Stress-regime correlation	ρ_S	0.55
Expected time between stress spells	$1/p_{N \rightarrow S}$	5 years
Expected stress-spell duration	$1/p_{S \rightarrow N}$	1 year
Loss given default	ℓ	45%

Two specifications are simulated under identical random seeds for the idiosyncratic and hazard-innovation draws, differing only in the correlation process: (i) a **constant-correlation benchmark** with $\rho(t) \equiv \rho_N = 0.15$ for all t ; and (ii) the **regime-switching specification** described in Section 3.2. Both specifications share the same long-run average correlation, since the regime-switching chain spends the large majority of time in the Normal state given the calibrated transition probabilities.

5 Results

5.1 Loss distribution and tail risk

Table 2 reports summary statistics of the simulated five-year portfolio loss distribution under both specifications.

Despite near-identical expected losses (a difference of 0.08 percentage points, i.e. within Monte Carlo noise for this sample size), the two specifications diverge materially in the tail. Loss volatility is approximately 15.6% higher under regime-switching correlation. The 99.9% VaR – the metric most closely aligned with the Basel IRB solvency standard – is 23.1% higher under regime-switching correlation than under the constant-correlation benchmark (28.80% vs.

Table 2: Simulated portfolio loss distribution: constant vs. regime-switching correlation

Statistic	Constant correlation	Regime-switching correlation
Expected loss	6.465%	6.388%
Loss volatility (std. dev.)	3.771%	4.358%
Value-at-risk, 99%	18.450%	20.700%
Value-at-risk, 99.9%	23.400%	28.800%
Expected shortfall, 97.5%	18.202%	20.974%

23.40% of notional). Expected shortfall at the 97.5% level, a metric increasingly emphasised in stress-testing and economic capital frameworks, is 15.2% higher (20.97% vs. 18.20%).

Figure 1 plots the simulated loss distributions and the VaR profile across confidence levels. The two histograms overlap substantially in the body of the distribution, consistent with the similar expected losses; the divergence is concentrated in the upper tail, where the regime-switching specification’s occasional excursions into the high-correlation stress state generate loss outcomes materially beyond what the constant-correlation model can produce with comparable frequency.




Figure 1: Left: histogram of simulated five-year portfolio losses under the two specifications. Right: value-at-risk as a function of confidence level. The two specifications are calibrated to the same long-run average correlation and produce similar loss distributions in the body, but the regime-switching specification produces materially higher tail risk.

5.2 Default counts

Table 3 reports selected percentiles of the number of defaulting obligors (out of 100) over the five-year horizon. The median and 90th-percentile default counts are similar across specifications, but the 99.9th percentile is markedly higher under regime-switching correlation (64 vs. 52 defaults), reflecting the fatter right tail driven by scenarios in which the portfolio spends an extended period in the high-correlation stress regime.

Table 3: Percentiles of simulated default counts (out of 100 obligors, five-year horizon)

Specification	p50	p90	p99	p99.9
Constant correlation	13	26	41	52
Regime-switching correlation	12	27	46	64

6 Discussion: Implications for Capital Adequacy and ECL Provisioning

These results speak directly to two areas of prudential practice. First, under the Basel III IRB framework, regulatory capital for credit risk is set with reference to a 99.9% VaR-type solvency standard, using an asset correlation parameter that is, for corporate exposures, fixed by regulatory formula rather than estimated period-by-period. If true correlation is regime-dependent and a bank’s internal models (or the regulatory formula itself) reflect only a long-run average, the results here suggest tail capital could be understated by a magnitude on the order of 20% in a portfolio broadly similar to the one simulated – a gap that would only be revealed empirically during, or immediately preceding, a stress episode, which is precisely when a capital shortfall is costliest to discover.

Second, under IFRS 9, lifetime expected credit losses must be recognised once an exposure experiences a significant increase in credit risk, and provisioning is explicitly forward-looking. A hybrid model of the type proposed here, in which default intensity is directly linked to an observable or estimable distance-to-default measure, offers a natural mechanism for identifying such staging transitions endogenously, rather than relying solely on ratings migration triggers. Incorporating a regime indicator (for example, derived from credit spread indices, equity volatility indices, or macro-financial conditions indices) into the correlation and intensity specification could allow provisioning models to become more explicitly forward-looking with respect to systemic, rather than purely idiosyncratic, credit deterioration.

More broadly, the exercise illustrates a general point of relevance to model risk management: a model can be well calibrated on average (as measured by expected loss) while still being materially misspecified in the region of the distribution that matters most for solvency and provisioning decisions. Validation exercises that focus primarily on average calibration accuracy may therefore miss exactly the model risk that matters most for capital adequacy.

7 Conclusion and Extensions

This paper develops a hybrid structural–intensity credit risk model with regime-switching systematic correlation and quantifies, via Monte Carlo simulation, the extent to which ignoring correlation regime-switching understates portfolio tail risk relative to a constant-correlation benchmark calibrated to the same average correlation and expected loss. For the stylised 100-obligor portfolio considered here, 99.9% VaR is understated by approximately 23% and 97.5% expected shortfall by approximately 15% under the constant-correlation assumption.

Several extensions would strengthen and generalise these findings. First, the regime process here is exogenous and unobserved; an empirically estimable version would link $\rho(t)$ to observable

macro-financial indicators (e.g. credit spread indices, implied equity volatility) via a filtering approach (e.g. a hidden Markov model estimated by maximum likelihood or a particle filter), enabling out-of-sample regime inference rather than assumed regime dynamics. Second, calibrating μ , σ , λ_0 , κ , and the regime transition probabilities to actual portfolio or index-level default and spread data (rather than the illustrative parameters used here) would allow the framework to produce institution-specific capital estimates rather than qualitative comparisons. Third, the current specification uses a single systematic factor; a multi-factor or sector-correlated extension would allow the regime-switching mechanism to operate heterogeneously across industries or geographies, which is likely to be relevant for internationally diversified portfolios. Finally, extending the loss-given-default assumption from a fixed constant to a stochastic, correlated LGD process (consistent with the empirically documented negative correlation between default rates and recovery rates) would likely further amplify the tail-risk gap identified here.

Disclosure: This is a working paper prepared as an independent research exercise. All quantitative results are generated from a Monte Carlo simulation using illustrative, non-proprietary parameters, and are intended to demonstrate the qualitative effect of regime-switching correlation on portfolio tail risk rather than to represent a calibrated estimate for any specific institution or portfolio.

References

- Merton, R.C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance*, 29(2), 449–470.
- Black, F., & Cox, J.C. (1976). Valuing Corporate Securities: Some Effects of Bond Indenture Provisions. *Journal of Finance*, 31(2), 351–367.
- Jarrow, R.A., & Turnbull, S.M. (1995). Pricing Derivatives on Financial Securities Subject to Credit Risk. *Journal of Finance*, 50(1), 53–85.
- Duffie, D., & Singleton, K.J. (1999). Modeling Term Structures of Defaultable Bonds. *Review of Financial Studies*, 12(4), 687–720.
- Vasicek, O. (2002). Loan Portfolio Value. *Risk*, 15(12), 160–162.
- Das, S.R., Duffie, D., Kapadia, N., & Saita, L. (2007). Common Failings: How Corporate Defaults Are Correlated. *Journal of Finance*, 62(1), 93–117.
- Giesecke, K. (2004). Correlated Default with Incomplete Information. *Journal of Banking & Finance*, 28(7), 1521–1545.
- Duffie, D., Eckner, A., Horel, G., & Saita, L. (2009). Frailty Correlated Default. *Journal of Finance*, 64(5), 2089–2123.
- Basel Committee on Banking Supervision (2011). *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Bank for International Settlements.

International Accounting Standards Board (2014). *IFRS 9 Financial Instruments*. IFRS Foundation.