POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Credit Value Adjustment in Reinsurance: A Quantitative Framework and Case Study

Tesi di Laurea Magistrale in
Mathematical Engineering - Quantitative Finance

Author: **Marco Enrico Baggio**

Student ID: 226481
Advisor: Prof. Alessandro Calvia
Academic Year: 2024-25

Abstract

Reinsurance allows insurance companies to transfer and diversify underwriting risk, but it also introduces exposure to counterparty credit risk arising from the possible default of the reinsurer. In recent years, this risk has gained increasing relevance, particularly within regulatory frameworks such as Solvency II, which require its explicit quantification. Despite this, quantitative models for reinsurance counterparty credit risk remain less developed than qualitative ones. Building on a model developed by Ceci, Colaneri, Frey and Kock, this thesis aims to calculate the Credit Value Adjustment (CVA) for specific, economically and financially relevant cases. Aggregate losses are modeled through a Cox process with stochastic intensity driven by an Ornstein–Uhlenbeck factor, capturing mean-reverting dynamics of the latent factor. Reinsurer default risk is described by a reduced-form intensity that depends on a two-state continuous-time Markov chain representing different market regimes. Contagion effects are introduced through a jump in the claims arrival intensity at the default time, reflecting increased exposure following counterparty failure. Within this setting, the CVA is characterized as the solution of a system of backward partial integro-differential equations. A regime-averaged CVA is also considered under the stationary distribution of the Markov chain. The theoretical analysis is complemented by a numerical implementation based on Monte Carlo simulations and PIDE methods. The results highlight the impact of regime switching and stochastic claim intensities on CVA, confirming the relevance of dynamic and regime-dependent modeling for reinsurance counterparty credit risk.

Keywords: Reinsurance Counterparty Credit Risk, Credit Value Adjustment,
Regime Switching, Ornstein–Uhlenbeck Process.

Abstract in lingua italiana

La riassicurazione consente alle compagnie di assicurazione di trasferire e diversificare il rischio tecnico, ma introduce un'esposizione al rischio di credito di controparte derivante dalla possibile insolvenza del riassicuratore. Negli ultimi anni questo rischio ha assunto crescente rilevanza, in particolare nel contesto regolamentare di Solvency II. Nonostante ciò, i modelli quantitativi per la valutazione del rischio di credito di controparte in ambito riassicurativo risultano meno sviluppati rispetto a quelli utilizzati nel settore bancario. Partendo da un modello proposto da Ceci, Colaneri, Frey e Kock, questa tesi si pone l'obiettivo di calcolare il Credit Value Adjustment (CVA) per casi specifici di interesse economico e finanziario, adottando un approccio quantitativo. Le perdite aggregate dell'assicurazione sono modellate tramite un processo di Cox con intensità stocastica, guidata da un fattore di tipo Ornstein–Uhlenbeck, in grado di rappresentare dinamiche di sinistro caratterizzate da un fattore latente con comportamento di ritorno alla media. Il rischio di default del riassicuratore è descritto attraverso un modello a forma ridotta, in cui l'intensità di default dipende da una catena di Markov a tempo continuo a due stati, rappresentativa di differenti regimi di mercato. Il modello tiene inoltre conto di effetti di contagio al momento del default della controparte, introducendo un salto nell'intensità dei sinistri che riflette l'aumento dell'esposizione dell'assicuratore a seguito della perdita della copertura riassicurativa. In questo contesto, il CVA viene caratterizzato come soluzione di un sistema di equazioni integro-differenziali alle derivate parziali di tipo backward. Viene inoltre considerata una formulazione media di regime del CVA, ottenuta mediante la distribuzione stazionaria della catena di Markov. L'analisi teorica è affiancata da un'implementazione numerica basata sia su simulazioni Monte Carlo sia sulla risoluzione numerica delle PIDE. I risultati mostrano come la presenza di intensità stocastiche e di dinamiche di regime nel rischio di default influisca in modo significativo sul valore del CVA, confermando la rilevanza di modelli dinamici e regime-dipendenti per la valutazione del rischio di credito di controparte nella riassicurazione.

Parole chiave: Riassicurazione, Rischio di Credito di Controparte, Credit Value Adjustment.

Contents

Introduction

Reinsurance plays a central role in the global insurance industry. By transferring part of their underwriting exposure to specialized counterparties, primary insurers reduce loss volatility, stabilize earnings, optimize capital allocation and enhance underwriting capacity. In economic terms, reinsurance contributes directly to solvency management and risk diversification, allowing insurers to underwrite larger and more complex portfolios than would otherwise be feasible. The size of the reinsurance market reflects this systemic importance. According to recent industry reports (e.g. Swiss Re Sigma Reports; Munich Re annual reports), global reinsurance premiums amount to several hundreds of billions of dollars annually, with the life and non-life segments both playing significant roles in global risk transfer. In particular, the increasing frequency and severity of catastrophic events, together with financial market volatility, have reinforced the strategic relevance of reinsurance as a capital and risk management tool. However, reinsurance is not a risk-free transfer mechanism. When an insurer cedes part of its liabilities to a reinsurer, it becomes exposed to the creditworthiness of that counterparty. If the reinsurer defaults before fulfilling its contractual obligations, the primary insurer may suffer significant financial losses. This exposure is known as Reinsurance Counterparty Credit Risk (RCCR). The relevance of RCCR has grown markedly over the last two decades. Episodes of systemic financial distress — most notably the 2008 global financial crisis — have demonstrated that even highly rated financial institutions may become vulnerable under severe macroeconomic shocks. A paradigmatic example is the near-collapse of American International Group (AIG), whose substantial exposure to credit default swaps linked to U.S. subprime mortgages led to massive losses and an unprecedented government intervention in 2008. Prior to the crisis, AIG was considered a highly rated and financially sound institution, yet its deterioration revealed how interconnected financial exposures and liquidity pressures can rapidly undermine even large and diversified groups. The episode significantly altered market perceptions of counterparty risk in the insurance and reinsurance sectors, highlighting that creditworthiness cannot be treated as static and that extreme market conditions may simultaneously impair asset values, capital positions and contractual performance. In this context, the assessment of reinsurance counterparty credit risk gained renewed importance, both from a regulatory and from a risk management perspective. Empirical evidence following the 2008–2009 financial crisis shows that systemic stress significantly affected insurers' capital positions, asset valuations and risk perceptions (Cummins & Weiss, 2014; The Geneva Association, 2010). At the same time, reinsurance has been shown to play a stabilizing role in periods of financial distress by reducing earnings volatility and supporting capital adequacy (Cummins et al., 2008; Cole & McCullough, 2006). Regulatory developments have reinforced this attention. Under the European Solvency II framework, insurers are required to explicitly quantify counterparty default risk within the Solvency Capital Requirement (SCR). The standard formula and internal model approaches both require the evaluation of expected losses and unexpected losses arising from counterparty exposures, including reinsurance recoverables. In particular, the calculation involves the exposure at default (EAD), the probability of default (PD) and the loss given default (LGD), thereby embedding RCCR within a capital adequacy framework that directly affects insurers' strategic decisions. Despite its importance, the quantitative modeling of RCCR remains less developed than counterparty risk modeling in banking. In the banking sector, Credit Value Adjustment (CVA) has been extensively studied, both theoretically and numerically, and is now a standard component of derivative pricing. In contrast, the insurance literature has historically focused more on underwriting risk and reserving, with counterparty risk often treated in a simplified or qualitative manner. Only more recent contributions have begun to address RCCR within a structurally consistent valuation framework that jointly models insurance losses and counterparty default risk. In particular, the approach developed by Ceci, Colaneri, Frey and Köck (2020) provides a reduced-form framework in which reinsurance valuation and counterparty risk are treated in a unified stochastic setting. Their model serves as the theoretical foundation for the extension developed in this thesis. From a modeling perspective, RCCR presents a distinctive challenge: it requires the joint treatment of two fundamentally different sources of uncertainty. On one side, insurance losses are driven by stochastic claim arrivals and claim severities. On the other, default risk depends on the financial condition of the reinsurer and on broader market factors. These risks are typically correlated: defaults tend to occur in stressed economic environments, which may simultaneously influence claim frequencies, replacement costs and discount factors. Capturing this interaction in a coherent and tractable framework is therefore a central issue in the valuation of reinsurance contracts under counterparty risk. This thesis addresses the problem within a reduced-form modeling framework, building on the approach introduced by Ceci, Colaneri, Frey and Köck (2020). In reduced-form models, default is represented through a stochastic intensity process rather than through an explicit modeling of the firm's asset dynamics. More precisely, the default time is modeled as the first jump of a Cox process with stochastic intensity λ_t^R. The intensity λ_t^R can be interpreted as the instantaneous conditional probability of default, given survival up to time t. This approach allows for flexible modeling of time-varying credit risk and is particularly suitable for valuation applications. Within this setting, counterparty risk is incorporated through a Credit Value Adjustment (CVA), defined as the discounted expected loss due to the possible default of the reinsurer. In simplified form, the CVA at time t can be written as

$$\mathrm{CVA}_t = (1 - R)\,\mathbb{E}\left[\int_t^T D(t,s)\,\lambda_s^R\,\mathrm{EE}_s\,ds\right],$$

where R denotes the recovery rate, $D(t,s)$ is the discount factor and EE_s represents the exposure at time s. In the reinsurance context, the exposure corresponds to the value of future recoverables from the reinsurer, which depends on the stochastic evolution of claims. The main contribution of this thesis is the development and numerical analysis of an extension of the model of Ceci et al., designed to jointly capture the stochastic claim intensity dynamics, the regime-dependent default risk and the contagion effects between default events and insurance exposure, within a unified valuation framework that remains analytically and numerically manageable. More specifically, the intensity of claim arrivals is driven by a latent stochastic factor following an Ornstein–Uhlenbeck (OU) process. This specification captures mean-reverting dynamics frequently observed in claim frequencies: periods of elevated loss activity — for instance following natural catastrophes or systemic events — are typically followed by gradual normalization. Empirical evidence on insurance loss cycles and catastrophe clustering indicates that claim frequencies exhibit significant temporal dynamics rather than being constant over time. Classic studies document cyclical behavior in property-liability markets (Cummins, Harrington & Klein, 1995; Cummins & Phillips, 2005) and non-stationary patterns in catastrophe frequencies (Barnett & Zehnwirth, 2011; Embrechts & Puccetti, 2006). These findings support the use of mean-reverting processes, such as Ornstein–Uhlenbeck dynamics, to represent aggregate claim intensity. The OU process offers a parsimonious representation of this behavior while preserving the Markov property required for dynamic valuation. Counterparty default risk is modeled through a reduced-form intensity that depends on a two-state continuous-time Markov chain (CTMC). The two regimes represent normal and stressed market conditions, each associated with a different level of default intensity. This regime-switching structure captures macro-financial fluctuations in credit quality and allows for endogenous transitions between low and high-risk environments. A further element is the introduction of contagion effects: at the default time of the reinsurer, the claim arrival intensity is allowed to jump. Economically, this reflects the sudden loss of reinsurance protection and the potential increase in replacement costs or retained exposure faced by the insurer. The model therefore captures both pre-default dependence (through regime switching) and at-default discontinuities (through jump effects). From a theoretical standpoint, the valuation problem is characterized through a system of backward Partial Integro-Differential Equations (PIDEs). In the OU–CTMC framework, the CVA satisfies a coupled system of regime-dependent PIDEs. The thesis derives the corresponding system, establishes its well-posedness under suitable assumptions, and analyzes its structural properties. From a numerical perspective, two complementary methodologies are implemented and compared: on one hand a Monte Carlo simulation, based on the discretization of the OU process and the simulation of the Markov chain and default time, on the other hand a finite-difference scheme for the solution of the coupled PIDEs. The comparison is performed in terms of pricing accuracy, computational time and numerical stability across different parameter configurations, with a careful analysis of sensitivity to changes in model parameters. The results indicate that, under the stochastic OU specification, the discrepancy between Monte Carlo and PIDE valuations becomes more pronounced relative to the constant-intensity benchmark, reflecting the stronger sensitivity of the PIDE system to the additional volatility factor driving exposure dynamics. From a computational standpoint, the introduction of stochastic claim intensity substantially amplifies the simulation burden: as shown in numerical analysis, specifically in Table 4.2, while the deterministic case requires only a few seconds under Monte Carlo, the OU specification increases runtime to several minutes due to the need to simulate an additional continuous factor and its interaction with default risk. Although the PIDE method also experiences an increase in computation time due to the higher-dimensional structure of the valuation problem, it remains significantly more efficient in terms of computational time than the Monte Carlo approach across all specifications considered. Finally, the inclusion of contagion effects generates discontinuities in the exposure profile at default, capturing a jump behavior that is absent in standard frameworks where exposure dynamics are purely diffusive and therefore continuous. These findings have direct implications for insurers' internal model development under Solvency II, as they highlight the importance of incorporating dynamic dependence structures in RCCR assessment. The remainder of the thesis is organized as follows. Chapter 1 reviews the literature on reinsurance counterparty risk and presents the reduced-form framework of Ceci et al.(2020). Chapter 2 develops the proposed model extension with OU-driven claim intensity and regime-switching default risk. Chapter 3 describes the numerical implementation of both Monte Carlo and PIDE methods. Chapter 4 presents the quantitative results and sensitivity analyses. Finally, Chapter 5 concludes and outlines possible directions for future research.

1 | Reinsurance Counterparty Credit Risk: Background

The diversification of risk has always been a widely adopted practice to mitigate potential unexpected losses, both in the banking and insurance sectors. A first-level solution for insurance companies is to enter into a reinsurance contract, which simultaneously achieves a high degree of diversification and contributes to the stabilization of results. As defined in [21], a reinsurance contract is an agreement through which an insurance company (the ceding insurer or primary insurer) transfers part of the risks it has assumed under its insurance policies to another insurance company, referred to as the reinsurer, in exchange for the payment of a reinsurance premium. In simple terms, reinsurance can be viewed as insurance for insurance companies, allowing insurers to protect themselves against large or unexpected losses. The primary purpose of reinsurance is risk management. By transferring a portion of potential losses, the ceding insurer reduces the volatility of its technical results, improves its financial stability, and optimizes capital allocation under regulatory constraints. From an economic perspective, reinsurance can be interpreted as a mechanism of risk sharing among insurance and reinsurance companies, particularly suited to low-frequency, high-severity events such as natural catastrophes (see, e.g., [8]). It is important to emphasize that a reinsurance contract does not involve the final policyholder; in fact, the contractual relationship between the insured and the primary insurer remains unchanged, while reinsurance operates exclusively at the inter-company level. Among the most commonly used forms of reinsurance, two non-proportional structures are particularly relevant: excess-of-loss (XL) and stop-loss reinsurance. In an excess-of-loss contract, coverage applies to each individual claim exceeding a predefined retention level M, up to an overall upper limit $\overline{K}$. In contrast, stop-loss reinsurance provides coverage on the aggregate amount of claims once total losses exceed a lower threshold $\underline{K}$, up to an upper limit $\overline{K}$. Excess-of-loss reinsurance is typically preferred in scenarios characterized by low-frequency, high-severity events, whereas stop-loss reinsurance is more suitable when losses arise from a larger number of events with lower individual severity (see, e.g., [4]). Section 2.1 provides a more detailed quantitative analysis of this distinction, as formulated in Equations (1.1) and (1.2). However, as highlighted in [5], this transfer does not eliminate risk but rather changes its nature: technical risk is partly replaced by counterparty credit risk. This risk transfer introduces a new source of vulnerability: the Reinsurance Counterparty Credit Risk (RCCR), namely the possibility that the reinsurer may fail to meet its contractual obligations. Indeed, credit risk has long been recognized as a structural component of insurance risk management. Insolvency events or credit downgrades involving major reinsurers can undermine the financial stability not only of the ceding companies but also of the entire insurance market.

Publications such as [19, 20] show that the credit quality of reinsurance companies has a first-order impact on the demand for reinsurance coverage by primary insurers. In particular, lower-rated reinsurers are associated with a reduction in the volume of reinsurance purchased, a higher degree of risk retention by ceding insurers, and a reallocation of demand toward reinsurers with stronger balance sheets.

From a quantitative perspective, empirical evidence suggests that counterparty credit risk is priced through both quantities and contractual terms. Primary insurers tend to require higher reinsurance capacity, stricter collateralization clauses, or lower attachment points when dealing with reinsurers of weaker credit quality. Conversely, when reinsurers are highly rated, insurers are willing to cede larger portions of risk and accept lower retention levels, reflecting the lower perceived default risk.

Empirical evidence shows that a deterioration in a reinsurer's credit rating leads to a statistically significant decline in reinsurance demand, even after controlling for price effects ([20]).

These findings highlight that reinsurance contracts embed an implicit credit risk premium, reflecting the expected loss associated with reinsurer default. As a result, counterparty credit risk affects not only pricing but also the optimal structure of reinsurance programs, including retention levels, coverage limits, and contract type. This interaction between underwriting risk and credit risk provides a strong motivation for explicitly modeling reinsurer default risk within a reduced-form credit risk framework. Moreover, downgrades of large international reinsurers can generate contagion and loss amplification effects, especially during periods of financial stress, thereby highlighting the systemic nature of RCCR.

This awareness became definitive after the 2008 financial crisis, when the default or fragility of entities previously considered sound, including globally significant reinsurers, demonstrated that capital adequacy alone does not ensure absolute security under turbulent market conditions. This episode prompted regulators and legislators to strengthen the prudential framework of the insurance industry, with an increasing focus on risks arising from exposures to external counterparties.

In Europe, the Solvency II Directive explicitly established the importance of counterparty risk assessment, incorporating it as one of the core modules of the Solvency Capital Requirement (SCR). Solvency II is a European regulatory framework that sets risk-based capital requirements for insurers and reinsurers, aiming to ensure their financial stability and the protection of policyholders. Supervisory authorities such as IVASS (*Istituto per la Vigilanza sulle Assicurazioni*) in Italy and EIOPA (*European Insurance and Occupational Pensions Authority*) at the European level have issued guidelines and methodologies for the accurate estimation and management of exposures arising from reinsurance treaties. These guidelines recommend that insurers quantify the expected recoverables from each reinsurance contract while taking into account the default probability of the reinsurer and potential correlations with the underlying insurance losses. Formally, the expected reinsurance recoverable R can be expressed as:

$$R = \mathbb{E}\big[\text{Recoverable} \times \mathbf{1}_{\{\text{reinsurer survives}\}}\big] = \text{Recoverable} \times (1 - PD)$$

where PD is the probability of default of the reinsurer over the relevant period. Such calculations are central to the SCR under Solvency II, ensuring that capital is held to cover both underwriting losses and potential counterparty defaults. However, as noted in [11], qualitative approaches still dominate industrial practice, which often prevent a precise and dynamic evaluation of actual risk and fail to quantify the economic impact of default or credit deterioration events.

At the same time, academic research has begun to explore more advanced quantitative approaches, inspired by financial modeling developed in the banking context. In particular, studies such as [4] have introduced models for the valuation of reinsurance counterparty risk based on Credit Value Adjustment (CVA) metrics, and have also studied dynamic hedging strategies. Similarly, [7] proposed a partial internal model compliant with Solvency II, which explicitly accounts for the reinsurer's default risk in the solvency capital calculation. These contributions demonstrate the evolution of research towards a more rigorous and coherent quantification of counterparty risk, while also revealing the persistent gap between theory and practice: quantitative models remain scarcely adopted in the industry due to their implementation complexity and calibration challenges with real-world data.

This scenario highlights a clear methodological gap between the insurance and banking sectors. In banking, counterparty credit risk has long been the subject of extensive theoretical study and detailed regulation. In particular, Basel III, a global regulatory framework issued by the Basel Committee on Banking Supervision, provides a comprehensive package of prudential rules aimed at strengthening the resilience of the banking system. Within this framework, counterparty credit risk arising from derivatives, securities financing transactions, and repurchase agreements is explicitly addressed through dedicated capital requirements.

A central element is the introduction of the Credit Value Adjustment (CVA) capital charge, which requires banks to hold capital against the potential mark-to-market losses caused by the deterioration of a counterparty's creditworthiness, even in the absence of default. Basel III allows institutions to compute this charge either through a standardized approach or through internal models, subject to supervisory approval, reflecting differences in modeling sophistication and risk management practices.

From a quantitative standpoint, banking regulation has fostered the development of advanced models for exposure-at-default, default probabilities, and loss given default, typically within reduced-form or structural credit risk frameworks. This regulatory and methodological maturity in the banking sector provides a natural benchmark for the analysis of counterparty credit risk in insurance and reinsurance markets, where similar exposures arise but have only more recently received comparable regulatory attention. This discrepancy underscores a key research objective: to develop state-of-the-art quantitative models similar to those used in banking, yet consistent with Solvency II requirements and market realities, without becoming impractical from a computational or implementation standpoint.

In light of these considerations, the present work aims to contribute to the study and modeling of Reinsurance Counterparty Credit Risk, with particular focus on CVA-based quantification methodologies, following the approach introduced in [4]. The goal is to investigate a modeling approach that, while remaining consistent with market dynamics and the principles of quantitative finance, is sufficiently tractable to allow a detailed analysis of a specific case study derived from Ceci et al., together with a corresponding numerical implementation of the selected case, illustrating the practical implications of the theoretical framework.

1.1. Model of Ceci, Colaneri, Frey, Kock (2020)

The aim of this section is to discuss the model introduced in the paper by Ceci et al., using it as the theoretical foundation for the study and developments presented in this thesis. The main objective is to construct a stochastic, market-consistent framework for the determination of the Credit Value Adjustment (CVA).

Compared to traditional static models, this reduced-form model allows for capturing dynamic dependencies between the loss risk and the counterparty's default risk, as well as modeling contagion effects that arise at the time of the reinsurer's default. The model proposed by Ceci et al. thus belongs to the research line that aims to adapt quantitative techniques developed in the banking sector to the reinsurance context, while maintaining a rigorous probabilistic structure and a clear economic interpretation.

Let us introduce a filtered probability space $(\Omega, \mathcal{G}, \mathbb{G}, \mathbb{Q})$, where:

- $\mathbb{Q}$ is the risk-neutral measure, consistent with market valuation, and
- $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$ is a complete, right-continuous filtration.

Two counterparties are considered, denoted respectively by I (insurance company) and R (reinsurance company), which enter into a reinsurance contract with time horizon T.

The process describing the cumulative claim amount underlying the reinsurance contract is denoted by $L = (L_t)_{t\in[0,T]}$. It represents the aggregate sum of all claim sizes occurred up to time t.

To model it, a counting process $N = (N_t)_{t\in[0,T]}$ is introduced, where N_t represents the total number of loss events observed up to time t. Each event i is associated with a positive and measurable random variable Z_i, describing the severity of the claim. The process L_t can thus be expressed as:

$$L_t = \sum_{i=1}^{N_t} Z_i, \quad t \in [0, T].$$

Hence, L_t represents the cumulative claim amount process, which constitutes the basis for the valuation of exposures and liabilities arising from the reinsurance contract.

In what follows, N_t is assumed to be a Cox process describing the number of loss events over time. The variables $(Z_i)_{i\geq 1}$ represent the claim severities and are assumed to be independent and identically distributed (i.i.d.) with law ν, and independent of the counting process N_t. This assumption allows for the separation of frequency and severity components in the modeling of aggregate losses.

Next, two processes are introduced to describe the stochastic dynamics of loss and default events. Specifically, the model features two distinct intensity processes: the loss intensity λ^L and the default intensity λ^R.

The loss intensity λ_t^L represents the instantaneous frequency at which loss events occur in the insurance portfolio, conditional on the information available up to time t. It is defined as the $\mathbb{G}$-predictable process such that

$$\mathbb{E}\left[\int_0^t C_s \, dN_s\right] = \mathbb{E}\left[\int_0^t C_s \lambda_s^L \, ds\right]$$

for all bounded $\mathbb{G}$-predictable process C. We can interpret λ^L as

$$\lambda_t^L = \lim_{\Delta t \downarrow 0} \frac{\mathbb{P}(N_{t+\Delta t} - N_t = 1 \mid \mathcal{G}_t)}{\Delta t}.$$

It is possible to show that the compensated process

$$M_t := N_t - \int_0^t \lambda_s^L \, ds$$

is a $\mathbb{G}$-martingale. From an interpretative perspective, λ_t^L captures the pure underwriting or technical risk associated with claim occurrences in the insurance portfolio and does not entail any counterparty default event. In probabilistic terms, λ_t^L can be interpreted as the conditional probability of observing a loss event in an infinitesimal interval following t. The default intensity λ_t^R is associated with the insolvency risk of the reinsurance company and quantifies the instantaneous probability of default of the counterparty within an infinitesimal time interval. Let τ_R denote the default time of the reinsurance company and define the default indicator process $\{H_t^R\}_{t\geq 0}$ as $H_t^R := \mathbf{1}_{\{\tau_R \leq t\}}$. The default intensity λ^R is defined as the $\mathbb{G}$-predictable process satisfying

$$\mathbb{E}\left[\int_0^t C_s \, dH_s\right] = \mathbb{E}\left[\int_0^t C_s \lambda_s^R (1 - H_{s^-}) \, ds\right]$$

for all bounded $\mathbb{G}$-predictable process C. We can interpret λ^R as

$$\lambda_t^R = \lim_{\Delta t \downarrow 0} \frac{\mathbb{P}(\tau_R \in (t, t + \Delta t] \mid \mathcal{G}_t)}{\Delta t}.$$

It is possible to show that the compensated process

$$M_t := H_t^R - \int_0^t \lambda_s^R \, (1 - H_s^R) \, ds$$

is a $\mathbb{G}$-martingale. The process λ_t^R therefore represents the hazard rate of the reinsurer within a reduced-form credit risk framework. In [4], intensities λ^R and λ^L are modeled as deterministic functions of two underlying factor processes X and Y, which govern their stochastic dynamics. More precisely, $\lambda_t^L = \lambda^L(X_t)$ and $\lambda_t^R = \lambda^R(Y_t)$, where $\lambda^L : \mathbb{R} \to \mathbb{R}^+$ and $\lambda^R : \mathbb{R} \to \mathbb{R}^+$. $X = (X_t)_{t\in[0,T]}$ and $Y = (Y_t)_{t\in[0,T]}$ satisfy the following system of SDEs:

$$
\begin{aligned}
dX_t &= \gamma^X(X_{t-})\, dH_t^R + b^X(X_t)\, dt + \sigma^X(X_t)\, dW_t^1, \qquad X_0 = x_0 \in \mathbb{R},\\
dY_t &= b^Y(Y_t)\, dt + \sigma^Y(Y_t)(\rho\, dW_t^1 + \sqrt{1-\rho^2}\, dW_t^2), \qquad Y_0 = y_0 \in \mathbb{R},
\end{aligned}
$$

For some $\rho \in [0,1]$ and measurable functions $b^X, b^Y : \mathbb{R} \to \mathbb{R}$, $\sigma^X, \sigma^Y : \mathbb{R} \to \mathbb{R}^+$ and for continuous and increasing function $\gamma^X : \mathbb{R} \to \mathbb{R}^+$. With this modeling choice, both the loss counting process $\{N_t\}_{t\geq 0}$ and the default indicator process $\{H_t^R\}_{t\geq 0}$ fall within the class of Cox processes, also referred to as doubly stochastic Poisson processes.

This modeling choice introduces a dynamic dependency structure between the evolution of risks and the economic or market conditions and it allows for the incorporation of external factors such as claim frequencies, weather conditions, or sectoral shocks (affecting λ^L) and credit spreads, market volatility, or liquidity and capital indicators (affecting λ^R) into the stochastic calculation of the intensities. From a theoretical standpoint, the introduction of the stochastic drivers X and Y enables the model to overcome the temporal independence assumption typical of constant-intensity models and to capture the correlation between insurance and counterparty risks through a correlation factor ρ in the system of stochastic differential equations (SDEs).

Now it is the turn of the substantive part of the contract, namely how the amount of risk transferred to the reinsurer is determined and under which terms. The paper considers two types of reinsurance contracts: Stop-Loss Reinsurance and Excess-of-Loss Reinsurance.

In a Stop-Loss Reinsurance contract, coverage applies to the aggregate amount of losses accumulated over a given period. The reinsurer intervenes only when the total cumulative loss L_T exceeds a pre-specified attachment point, denoted by $\underline{K}$. The contract specifies a maximum coverage limit $\overline{K}$, beyond which losses revert to the insurer:

$$
\phi(l) = \min\{\overline{K}, [l - \underline{K}]^+\} \tag{1.1}
$$

This type of contract is particularly suitable for protecting the insurer from high-frequency, high-severity aggregate events, such as years characterized by an exceptionally large number of claims. The insurer thus benefits from limiting the overall risk exposure, while the reinsurer assumes the risk of rare but potentially very costly aggregated events.

In contrast, in an Excess-of-Loss Reinsurance contract, coverage applies to each individual loss event rather than to the aggregate. Here, the reinsurer intervenes for each claim when its severity Z_i exceeds the retention threshold M, possibly with a maximum coverage limit per event $\overline{K}$. Thus, the total amount borne by the reinsurer at time T is given by:

$$\phi(l) = \min\left\{\overline{K}, \sum_{i=1}^{N_T}[Z_i - M]^+\right\} \tag{1.2}$$

With all these elements, it is then possible to compute the market value of the reinsurance contract, given by the conditional expected value of the discounted future reinsurance payments:

$$V_t^{\phi} = \mathbb{E}^{\mathbb{Q}}\left[e^{-r(T-t)}\phi(L_T)\,\middle|\,\mathcal{G}_t\right], \qquad 0 \leq t \leq T \tag{1.3}$$

1.2. CVA Calculation According to Ceci et al. (2020)

The contract value derived in the previous section represents a default-free valuation at time t, obtained under the implicit assumption that the reinsurer will honor its contractual obligations with certainty. In other words, this value corresponds to the price of the reinsurance agreement in a hypothetical setting where counterparty credit risk is absent. In reality, the reinsurer is subject to default risk, and there exists a non-zero probability that it may fail to meet its payment obligations, either partially or entirely, over the lifetime of the contract. As a result, the default-free price does not account for the risk of reinsurer default, and therefore overestimates the true economic value of the reinsurance coverage for the ceding insurer. To account for this discrepancy, the contract value must be adjusted to reflect the expected losses arising from a potential default of the counterparty. This adjustment is performed through the introduction of the Credit Value Adjustment (CVA), which represents the market-consistent cost of counterparty credit risk. Formally, the CVA is defined as the difference between the default-free value of the contract and its credit-risk-adjusted value, and can be interpreted as the discounted expected loss due to the possible default of the reinsurer over the relevant time horizon. By incorporating the CVA, the resulting price reflects both the underwriting risk embedded in the reinsurance contract and the credit quality of the reinsurer, thereby providing a more accurate and economically meaningful valuation.

In the event that the reinsurer R defaults before the maturity of the contract, i.e., at some time $\tau_R \leq T$, the ceding insurer would need to replace the original reinsurance coverage by entering into a new reinsurance contract with value $V_{\tau_R}^{\phi}$, representing the post-default value of the contract.

It is assumed that, at the time of default, the insurer receives a recovery payment corresponding to a fraction of the post-default contract value, given by $(1 - \delta^R)\, V_{\tau_R}^{\phi}$, where

$\delta^R \in [0,1]$ denotes the Loss Given Default (LGD) of the reinsurer. The LGD represents the proportion of the contract value that is lost due to the reinsurer's default and is a standard measure in credit risk modeling. A value of $\delta^R = 0$ corresponds to full recovery, i.e., no loss, whereas $\delta^R = 1$ corresponds to total loss, with the ceding insurer receiving nothing. In economic terms, δ^R quantifies the severity of the counterparty default and directly affects the expected loss component in the valuation of reinsurance contracts under counterparty credit risk considerations.

In the insurance context, the CVA is generally computed through a standard formulation under the risk-neutral measure $\mathbb{Q}$ and LGD δ^R:

$$\mathrm{CVA}_t = \mathbb{E}^{\mathbb{Q}}\left[\int_t^T e^{-r(s-t)}\delta^R\,\mathrm{EAD}_s\,\mathbf{1}_{\{\tau_R \leq s\}}\,\middle|\,\mathcal{G}_t\right], \tag{1.4}$$

Here, EAD_s (Exposure at Default) represents the economic exposure of the ceding insurer to the reinsurer at time s, i.e., the value of the reinsurance contract that would be lost if the reinsurer were to default at that time. The EAD quantifies the size of the potential credit loss, and its evolution over time is essential for the accurate computation of the CVA.

In [4], using the stochastic model described earlier, the CVA is computed as the market-consistent value of future credit losses, where EAD corresponds to the post-default value of the reinsurance contract, so that

$$\mathrm{CVA}_t = \mathbb{E}^{\mathbb{Q}}\left[\int_t^T \delta^R\,\mathrm{V}_s^{\phi}\,e^{-r(s-t)}\mathrm{dH}_s^R\,\middle|\,\mathcal{G}_t\right], 0 \leq t \leq T \tag{1.5}$$

Another way to interpret the CVA is as a monetary reserve that must be held at time t to account for the counterparty risk with the reinsurance company.

Within this framework, an alternative characterization method for the CVA is developed, showing that both the contract value and the CVA can be computed as solutions to a PIDE (Partial Integro-Differential Equation).

PIDEs represent a fundamental tool for modeling the value of financial or insurance contracts in the presence of jumps or discontinuities in the underlying processes. In the context of credit risk, such discontinuities may reflect sudden events such as counterparty default or sharp changes in credit intensity.

In this setting, the valuation of instruments sensitive to credit risk (such as credit default swaps, insurance-linked bonds, or reinsurance contracts with loss clauses) naturally leads to the formulation of a PIDE. The general form of the equation is:

$$\frac{\partial V}{\partial t} + \mathcal{L}V - rV = 0, \tag{1.6}$$

where:

- $\mathcal{L}$ is the infinitesimal generator of the state processes in the model;
- r is the risk-free rate;
- V is the pricing function of the contract.

The PIDE formulation thus allows one to consistently account for jump events, providing a mathematical framework that links credit risk dynamics with the economic valuation of contracts.

In particular, in a credit reinsurance context, a PIDE can be used to compute the expected value of the contract conditional on the survival of the reinsured counterparty, or to compute the expected loss in the presence of jumps in the default intensity or in the recovery rate.

The use of PIDEs in credit risk and reinsurance models thus enables more realistic and robust valuations, capable of capturing the discrete and sudden nature of financial and insurance shocks—overcoming the limitations of purely diffusive models.

1.2.1. Contagion-Free Market Framework

Before introducing the elements required for the PIDE formulation, it is convenient to define an auxiliary market setting in which contagion effects induced by the default of the reinsurer are removed. This construction, commonly referred to as a contagion-free market, plays a central role in the valuation of counterparty credit risk and, in particular, in the computation of the Credit Value Adjustment.

In the full model, the default of the reinsurer at time τ_R may induce abrupt changes in the dynamics of the underlying risk factors, reflecting contagion effects between the reinsurer's creditworthiness and the insurance portfolio. In particular, the process X, which drives the loss intensity of the insurance portfolio, may experience a jump at τ_R, thereby coupling underwriting risk and counterparty credit risk.

To isolate and control for these contagion effects, a set of auxiliary processes $(\widetilde{X}, \widetilde{N}, \widetilde{L})$ is introduced. These processes are defined on the same probability space and filtration and describe the evolution of the system in the absence of contagion. More precisely, the process $\widetilde{X}$ is specified to have the same drift and diffusion as X, except that it does not experience any jump at the reinsurer's default time τ_R. Hence, $\widetilde{X}$ evolves continuously across τ_R and represents the dynamics of the loss-driving factor in a market where the default of the reinsurer has no direct impact on the insurance risk dynamics.

Correspondingly, $\widetilde{N}$ and $\widetilde{L}$ denote the counting process of loss events and the associated aggregate loss process driven by $\widetilde{X}$. They are defined via a stochastic time-change construction as follows.

Let $\eta = (\eta_t)_{t\geq 0}$ be a standard Poisson process with unit intensity, and define the random time-change $\widetilde{\theta}_t := \int_0^t \lambda^L(\widetilde{X}_s)\, ds$. Then the loss counting process is given by $\widetilde{N}_t := \eta_{\widetilde{\theta}_t}$ with $t \geq 0$. Since η has unit intensity, its compensator is given by t. Under the time-change $\widetilde{\theta}_t$, the compensator becomes $\widetilde{\theta}_t$. Hence, $\widetilde{N}$ admits compensator $\int_0^t \lambda^L(\widetilde{X}_s)\, ds$ and therefore has stochastic intensity $\lambda^L(\widetilde{X}_t)$. Moreover, let $(Z_n)_{n\geq 1}$ be a sequence of i.i.d. non-negative random variables with law ν representing individual loss severities, independent of η and $\widetilde{X}$. Define the compound Poisson process $M_t := \sum_{n=1}^{\eta_t} Z_n$, $t \geq 0$ and set $\widetilde{L}_t := M_{\widetilde{\theta}_t}$, $t \geq 0$. Consequently, $\widetilde{N}$ is a Cox process with stochastic intensity $\lambda^L(\widetilde{X}_t)$, and $\widetilde{L}$ is a compound Cox process representing the cumulative claim amount up to time t. Finally, let us define the filtration $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$ by $\mathcal{G}_t = \mathcal{F}_t \vee \mathcal{H}_t, t \geq 0$ where $\mathcal{H}_t$ is the filtration generated by the counterparty default indicator H^R and $\mathcal{F}_t = \mathcal{F}_t^W \vee \mathcal{F}_t^{\widetilde{L}}$ with $\mathcal{F}_t^W$ is the filtration generated by the Brownian motion W driving $\widetilde{X}$ and $\mathcal{F}_t^{\widetilde{L}}$ is the filtration generated by the cumulative loss process $\widetilde{L}$. These processes are constructed to describe the evolution of insurance losses independently of the default event of the reinsurer.

A key property of this construction is that, up to the default time τ_R, the original and contagion-free models are indistinguishable. That is, for all $t < \tau_R$, $(X_t, N_t, L_t) = (\widetilde{X}_t, \widetilde{N}_t, \widetilde{L}_t)$, a.s. As a consequence, before default occurs, the two sets of processes can be used interchangeably. This property allows conditional expectations and value adjustments to be computed within the contagion-free market, while preserving full consistency with the original model dynamics prior to default.

The introduction of contagion-free processes is particularly useful in the derivation of the PIDE satisfied by the contract value and its CVA adjustment, as it enables a clear separation between intrinsic insurance risk and the additional effects induced by reinsurer default and contagion.

Returning to the development of the PIDE, as a first step the generator of the two-dimensional Markov process $(\widetilde{L}, \widetilde{X})$ is defined:

$$\begin{aligned}\mathcal{L}^{(\widetilde{L},\widetilde{X})} f(t,l,x) =& \frac{\partial f}{\partial x}(t,l,x)\, b^X(x) + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}(t,l,x)\,(\sigma^X(x))^2 \\ &+ \int_{\mathbb{R}^+} \Big(f(t,l+z,x) - f(t,l,x)\Big)\,\lambda^L(x)\,\nu(dz).\end{aligned} \tag{1.7}$$

Then, through the proof of the following proposition, it is shown that the value of the contract V at the time of counterparty default (i.e., at τ) can be described as the solution to a backward PIDE.

Assumption A.1

Define the instantaneous covariance matrix of (X, Y) as

$$\Sigma(x,y) = \begin{pmatrix} (\sigma^X(x))^2 & \rho\,\sigma^X(x)\sigma^Y(y) \\ \rho\,\sigma^X(x)\sigma^Y(y) & (\sigma^Y(y))^2 \end{pmatrix}, \qquad (x,y) \in \mathbb{R}^2.$$

(A1) The functions b^X, b^Y, σ^X, and σ^Y are Lipschitz.

(A2) There exists a constant $\beta > 0$ such that, for all $(x, y) \in \mathbb{R}^2$ and all $w \in \mathbb{R}^2$,

$$w^\top \Sigma(x,y) w \geq \beta \|w\|^2.$$

(A3) The functions λ^L and λ^R are Lipschitz continuous and bounded.

(A4) The claim–size distribution ν has finite second moment.

Proposition 1.1. *Under assumption A.1, there exists a unique bounded classical solution v^ϕ (i.e., continuous, $\mathcal{C}^1$ in t and $\mathcal{C}^2$ in x) to the backward PIDE*

$$\frac{\partial v^\phi}{\partial t}(t,l,x) + \mathcal{L}^{(\widetilde{L},\widetilde{X})} v^\phi(t,l,x) = r\, v^\phi(t,l,x), \qquad (t,l,x) \in [0,T) \times \mathbb{R}^+ \times \mathbb{R}, \tag{1.8}$$

with terminal condition

$$v^\phi(T,l,x) = \phi(l).$$

Moreover, for $\tau_R \leq T$ it holds that

$$V^\phi_{\tau_R} = v^\phi\Big(\tau_R,\ \widetilde{L}_{\tau_R}, \widetilde{X}_{\tau_R} + \gamma^X(\widetilde{X}_{\tau_R})\Big)\,.$$

Proof. The proof is available in [4]; it is reproduced here in detail for completeness. The process $(\widetilde{L}, \widetilde{X})$ is a two-dimensional Markov process, where $\widetilde{L}$ is a pure-jump process and generator $\mathcal{L}^{(\widetilde{L},\widetilde{X})}$ given in equation (1.7)

Existence: by definition, we have

$$v^{\phi}(t,l,x) = \mathbb{E}^{\mathbb{Q}}\Big[e^{-r(T-t)}\phi(\widetilde{L}_T) \mid \widetilde{L}_t = l,\ \widetilde{X}_t = x\Big].$$

define $M_s := e^{-r(s-t)}\, v^{\phi}(s, \widetilde{L}_s, \widetilde{X}_s)$ with $t \le s \le T$

$$M_s = e^{-r(s-t)}\, \mathbb{E}^{\mathbb{Q}}\Big[e^{-r(T-s)}\phi(\widetilde{L}_T) \mid \widetilde{L}_s = l,\ \widetilde{X}_s = x\Big] = \mathbb{E}^{\mathbb{Q}}\Big[e^{-r(T-t)}\phi(\widetilde{L}_T) \mid \widetilde{L}_s = l,\ \widetilde{X}_s = x\Big]$$

By assumption, ϕ is bounded; therefore the random variable $e^{-r(T-t)}\phi(\widetilde{L}_T)$ is integrable, and consequently M_s, being the conditional expectation with respect to $\widetilde{L}_s$ and $\widetilde{X}_s$, is a martingale.

Now, by applying Itô's formula to the martingale M_s:

$$d(e^{-r(s-t)}\, v^{\phi}(s, \widetilde{L}_s, \widetilde{X}_s)) = e^{-r(s-t)}\left(\tfrac{\partial v^{\phi}}{\partial s} + \mathcal{L}^{(\widetilde{L},\widetilde{X})}v^{\phi} - r\, v^{\phi}\right)(s, \widetilde{L}_s, \widetilde{X}_s)ds + e^{-r(s-t)}dF_s$$

where dF_s denotes the local martingale term containing the stochastic components: it includes the diffusion term from $\widetilde{X}$ and, if applicable, the compensated jump term from $\widetilde{L}$. Knowing that M_s is a martingale, its finite variation part must vanish along the paths of $(\widetilde{L}_s, \widetilde{X}_s)$. To deduce the PIDE pointwise, we assume that the support of $(\widetilde{L}_s, \widetilde{X}_s)$ is $\mathbb{R}^+ \times \mathbb{R}$. Under this assumption, we can conclude that

$$\frac{\partial v^{\phi}}{\partial s}(s,l,x) + \mathcal{L}^{(\widetilde{L},\widetilde{X})}v^{\phi}(s,l,x) - r\, v^{\phi}(s,l,x) = 0, \qquad \forall (l,x) \in \mathbb{R}^+ \times \mathbb{R}.$$

Letting $s \to t$ then yields the PIDE for $v^{\phi}(t,l,x)$.

Uniqueness: Assume that u is a classical solution to the PIDE, with $u(T,l,x) = \phi(l)$. Defining $N_s = e^{-r(s-t)}\, u(s, \widetilde{L}_s, \widetilde{X}_s)$, with $t \le s \le T$ as before, by applying Itô's formula to N gives

$$dN_s = e^{-r(s-t)}\left(\frac{\partial u}{\partial s} + \mathcal{L}^{(\widetilde{L},\widetilde{X})}u - ru\right)(s, \widetilde{L}_s, \widetilde{X}_s)\, ds + dM_s,$$

where dM_s collects all the stochastic terms (diffusion and compensated jump parts). Since u satisfies the PIDE, the drift term in parentheses vanishes, and therefore $dN_s = dM_s$. Moreover, since u is bounded and $e^{-r(s-t)}$ is bounded, N is uniformly integrable, and consequently is a true martingale. So, in particular, it is true that

$u(t,l,x) = N_t = \mathbb{E}^{\mathbb{Q}}\Big[N_T \mid \widetilde{L}_t = l,\ \widetilde{X}_t = x\Big]$ but $N_T = e^{-r(T-t)}\phi(\widetilde{L}_T)$

and so finally $u(t,l,x) = \mathbb{E}^{\mathbb{Q}}\Big[e^{-r(T-t)}\phi(\widetilde{L}_T) \mid \widetilde{L}_t = l,\ \widetilde{X}_t = x\Big] = v^{\phi}(t,l,x)$

By the strong Markov property, on the event $\{\tau_R \le T\}$ we have

$$V^{\phi}_{\tau_R} = v^{\phi}(\tau_R, L_{\tau_R}, X_{\tau_R}).$$

From the construction of the model, before the default we have the identities

$$L_{\tau_R} = \widetilde{L}_{\tau_R}, \qquad X_{\tau_R} = \widetilde{X}_{\tau_R} + \gamma_X(\widetilde{X}_{\tau_R}),$$

with $\widetilde{X}_{\tau_R^-} = \widetilde{X}_{\tau_R}$. Inserting these expressions yields

$$V^{\phi}_{\tau_R} = v^{\phi}\Big(\tau_R,\ \widetilde{L}_{\tau_R},\ \widetilde{X}_{\tau_R} + \gamma_X(\widetilde{X}_{\tau_R})\Big),$$

which concludes the proof. □

Finally, the equation defining the computation of the CVA in this framework (via the solution of a PIDE) is obtained through the following proposition.

Proposition 1.2. *Under Assumption A.1, the value of the CVA is given by*

$$CVA_t = \delta^R(1 - H^R_t) f^{CVA}(t, L_t, X_t, Y_t) \tag{1.9}$$

where $f^{CVA} : [0,T] \times \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^+$ is a classical solution (i.e., continuous, $\mathcal{C}^1$ in t and $\mathcal{C}^2$ in (x,y)) of the backward PIDE:

$$\frac{\partial f^{CVA}}{\partial t} + \mathcal{L}^{(\widetilde{L},\widetilde{X},Y)} f^{CVA} + \lambda^R(y) v^{\phi}(t,l,x+\gamma^X(x)) = (\lambda^R(y) + r)\, f^{CVA} \tag{1.10}$$

where f^{CVA} is evalueted in (t,l,x,y) with terminal condition $f^{CVA}(T,l,x,y) = 0$. The generator $\mathcal{L}^{(\widetilde{L},\widetilde{X},Y)}$ is given by:

$$\begin{aligned}\mathcal{L}^{(\widetilde{L},\widetilde{X},Y)} f = {} & \frac{\partial f}{\partial x}\, b^X(x) + \frac{\partial f}{\partial y}\, b^Y(y) + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}\, (\sigma^X(x))^2 + \frac{1}{2}\frac{\partial^2 f}{\partial y^2}\, (\sigma^Y(y))^2 \\ & + \frac{\partial^2 f}{\partial x \partial y}\, \rho \sigma^X(x) \sigma^Y(y) + \int_{\mathbb{R}^+} \Big(f(t,l+z,x,y) - f(t,l,x,y)\Big)\, \lambda^L(x)\, \nu(dz).\end{aligned} \tag{1.11}$$

where f is always evalueted in (t,l,x,y).

Proof. The proposition proved before allows to express CVA at τ_R in terms of contagion-free quantities:

$$\mathrm{CVA}_t = \mathbb{E}^{\mathbb{Q}}\left[\int_t^T \delta^R v^\phi(s, \widetilde{L}_s, \widetilde{X}_s + \gamma^X(\widetilde{X}_s))\, e^{-r(s-t)} \mathrm{dH}_s^R \,\middle|\, \mathcal{G}_t\right] \tag{1.12}$$

In this expression it is possible to replace $\mathcal{G}_t$ with its definition $\mathcal{F}_t^W \vee \mathcal{F}_t^{\widetilde{L}} \vee \mathcal{H}_t$, since these sigma fields coincide up to time τ_R. By theorem 10.19 of [17] is known that τ_R is doubly stochastic with respect to the background filtration and also $\mathbf{Q}(\tau_R > t \mid \mathcal{F}_\infty^W \vee \mathcal{F}_\infty^{\widetilde{L}}) = e^{-\int_0^t \lambda^R(Y_s)\, ds}$. So one gets:

$$\mathrm{CVA}_t = \delta^R(1 - H_t^R)\mathbb{E}^{\mathbb{Q}}\left[\int_t^T v^\phi(s, \widetilde{L}_s, \widetilde{X}_s + \gamma^X(\widetilde{X}_s))\lambda^R(Y_s)\, e^{-\int_t^s (r+\lambda^R(Y_u))\, du} ds \,\middle|\, \mathcal{F}_t\right] \tag{1.13}$$

The process $(\widetilde{L}, \widetilde{X}, Y)$ is Markovian with respect to the filtration $\mathbb{F}$ with generator as in (1.11). It follows that, similar to Proposition 1.1, there exists a function f^{CVA} : $[0, T] \times \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^+$ which is a classical solution of the PIDE (1.10).
In order to have the expression of the proposition, note that on the event $\{\tau_R > t\}$, $1 - H_t^R = 1$ and also $\widetilde{L}_t = L_t$ and $\widetilde{X}_t = X_t$, which concludes the proof. □

2 | Extensions and Case Studies of the Model

This chapter is devoted to the analysis and discussion of case studies of the reduced-form framework proposed in [4], with the objective of providing an in-depth analysis of two special cases of the main state processes with respect to the more general formulations considered before, while preserving the same financial interpretation and mathematical consistency of the original framework. In particular, the chapter focuses on alternative dynamics for the default-related process and the loss intensity, which are jointly applied to a revised version of Example 3.4 from [4].

The first change concerns the modeling of the default environment. To this end, the default-related factor process Y is specified as a two-state continuous-time Markov chain (CTMC) (we refer the reader to Appendix A for the theoretical framework). The market is assumed to have two possible states, denoted as state 0 and state 1, representing normal and stressed market conditions, respectively. Each state is associated with a specific level of default intensity λ^R, reflecting the different probabilities of default consistent with the prevailing market regime. The stochastic evolution of Y between these states is governed by constant transition intensities, which capture both the frequency and the persistence of regime switches.

The use of a regime-switching framework for default intensities is well established in the credit risk literature. From an economic perspective, such models reflect the empirical evidence that default risk varies with macroeconomic and financial conditions and tends to increase during periods of systemic stress. A finite-state Markov chain provides a simple and effective way to represent these regime-dependent dynamics. This approach has been widely adopted both in general regime-switching models, as discussed in [9, 12], and in reduced-form credit risk settings (see, e.g., [2, 15]).

From a quantitative perspective, the CTMC specification offers clear analytical advantages. Since the regime process takes values in a finite state space, the Markov property ensures that the valuation problem can be formulated in a finite-dimensional state space, leading to a system of coupled partial integro-differential equations (PIDEs).

Regarding the claims arrival process, it is modeled as a Cox process, in line with [4]; however, the factor process X is here assumed to follow Ornstein–Uhlenbeck (OU) dynamics (we refer the reader to Appendix B for the theoretical framework). This specification represents a particular case within the more general class of processes considered in the original framework, while remaining economically meaningful and financially consistent.

The economic motivation for modeling claims arrival intensity through a mean-reverting process is supported by empirical evidence (see, e.g. [16]). Loss frequencies tend to exhibit cyclical behavior: periods of elevated loss activity are often followed by phases of normalization due to underwriting adjustments, changes in risk appetite, and portfolio rebalancing, while unusually low loss levels are typically not persistent in the long run. Mean-reverting intensity models have therefore been widely adopted in the actuarial and insurance mathematics literature (see, e.g. [6]).

From an analytical and numerical perspective, the Ornstein–Uhlenbeck process offers a convenient balance between flexibility and simplicity. Its Gaussian and Markovian structure simplifies the derivation of valuation equations and the implementation of numerical schemes. Moreover, the OU process naturally leads to linear drift terms in the associated generators, resulting in partial integro-differential equations that are well suited for finite-difference or Fourier-based methods. These features make the OU specification particularly appealing for a systematic numerical analysis, as also highlighted in [3] in the context of interest rate models.

It is worth emphasizing that, within the proposed framework, the roles of the two stochastic processes remain clearly distinct. The CTMC process Y governs the evolution of the general economic environment and captures regime-dependent counterparty risk, while the Ornstein–Uhlenbeck process X drives the intensity of loss occurrences. These processes are, however, implicitly linked: at the default of Y, the claim arrival intensity X experiences a proportional jump, maintaining consistency with the base case considered in [4]. This separation allows different sources of risk to be modeled in a coherent and interpretable way, while remaining fully consistent with the reduced-form approach adopted throughout the thesis.

2.1. Case Study: Y_t CTMC and X_t OU

In this section, the post-default value and the CVA are computed within the previously described framework, namely with X as an OU process and Y as a two-state CTMC process.Before proceeding, it is necessary to prove a version of Proposition 4.2 of [4] under the following assumptions

Assumption A.2

(A1) The function λ^L is Lipschitz continuous and bounded.

(A2) The claim–size distribution ν has finite second moment.

Proposition 2.1 (CVA with OU and two–state CTMC). *Under Assumptions A.2, suppose that the stochastic factor X follows an Ornstein–Uhlenbeck diffusion:*

$$dX_t = \kappa(\theta - X_t)\,dt + \sigma\,dW_t,$$

where $\kappa > 0$ is the speed of mean reversion, $\theta > 0$ is the long–run mean level, $\sigma > 0$ is the volatility parameter, and W is a Brownian motion under the risk–neutral measure.

Moreover, let $Y = (Y_t)_{t\geq 0}$ be a two–state continuous–time Markov chain with state space $\{0, 1\}$ and generator matrix

$$Q = \begin{pmatrix} -q_{01} & q_{01} \\ q_{10} & -q_{10} \end{pmatrix},$$

where $q_{01} > 0$ and $q_{10} > 0$ denote the transition intensities from state 0 to state 1 and from state 1 to state 0, respectively.

Finally, assume that the counterparty default intensity satisfies

$$\lambda^R(Y_t) = \begin{cases} \lambda_0^R, & \text{if } Y_t = 0, \\ \lambda_1^R, & \text{if } Y_t = 1. \end{cases}$$

Then, the Credit Value Adjustment is given by

$$\mathrm{CVA}_t = \delta^R(1 - H_t^R)\, f^{\mathrm{CVA}}(t, L_t, X_t, Y_t), \tag{2.1}$$

where

$$f^{\mathrm{CVA}} : [0, T] \times \mathbb{R}^+ \times \mathbb{R} \times \{0, 1\} \longrightarrow \mathbb{R}^+$$

is a classical solution of the following backward system of integro–differential equations. For each state y_i with $i \in \{0,1\}$ and $f_i(t,l,x) := f(t,l,x,y_i)$ it holds

$$
\begin{aligned}
&\frac{\partial f_i}{\partial t}(t,l,x) + \kappa(\theta - x)\frac{\partial f_i}{\partial x}(t,l,x) + \frac{1}{2}\sigma_X^2 \frac{\partial^2 f_i}{\partial x^2}(t,l,x) \\
&\qquad + \int_{\mathbb{R}^+} \Big(f_i(t,l+z,x) - f_i(t,l,x)\Big)\lambda^L(x)\,\nu(dz) \\
&\qquad + q_{i(1-i)}\big(f_{(1-i)}(t,l,x) - f_i(t,l,x)\big) + \lambda_i^R\, v^\phi\big(t,l,x+\gamma_X(x)\big) \\
&\qquad = (r + \lambda_i^R)\, f_i(t,l,x),
\end{aligned}
\tag{2.2}
$$

with terminal condition $f_i(T,l,x) = 0, \qquad i,j \in 0,1$

Proof. The CVA is a payment–at–default claim. As in the general framework considered in the paper, Proposition 4.1 of [4] allows one to express the payoff at the default time τ_R exclusively in terms of contagion–free quantities. Therefore, for any $t \in [0,T]$,

$$
\mathrm{CVA}_t = \mathbb{E}^{\mathbb{Q}}\left[\int_t^T \delta^R\, v^\phi\big(s,\widetilde{L}_s,\widetilde{X}_s + \gamma_X(\widetilde{X}_s)\big)e^{-r(s-t)}\,dH_s^R \,\middle|\, \mathcal{G}_t\right]. \tag{2.3}
$$

Since the filtration $\mathcal{G}_t = \mathcal{F}_t \vee \mathcal{H}_t$ coincides with $\mathcal{F}_t$ up to the default time τ_R, $\mathcal{G}_t$ can be replaced by $\mathcal{F}_t$.

Consider the three-dimensional Markov process $(\widetilde{L},\widetilde{X},Y)$. Its infinitesimal generator $\mathcal{L}$ acts on a sufficiently smooth function $f_i(t,l,x) = f(t,l,x,y_i), \quad i \in [0,1]$ as the sum of the generators of the jump-diffusion $(\widetilde{L},\widetilde{X})$ and the discrete CTMC Y:

$$
\begin{aligned}
\mathcal{L}f_i(t,l,x) = &\underbrace{\kappa(\theta - x)\frac{\partial f_i}{\partial x}(t,l,x) + \frac{1}{2}\sigma_X^2\frac{\partial^2 f_i}{\partial x^2}(t,l,x)}_{\text{OU dynamics}} \\
&+ \underbrace{\int_{\mathbb{R}^+}\big(f_i(t,l+z,x) - f_i(t,l,x)\big)\lambda^L(x)\nu(dz)}_{\text{Claim jumps}} \\
&+ \underbrace{q_{i(1-i)}\big(f_{(1-i)}(t,l,x) - f_i(t,l,x)\big)}_{\text{CTMC switching}}.
\end{aligned}
$$

Here, the first two terms correspond to the continuous OU dynamics of $\widetilde{X}_t$ and the jump dynamics of $\widetilde{L}_t$, while the last term arises from the infinitesimal generator of the two-state Markov chain Y_t. In particular, for a small Δt one has

$$
\mathbb{E}[f(t+\Delta t,l,x,Y_{t+\Delta t}) \mid Y_t = y_i] = f_i(t,l,x) + q_{i(1-i)}\big(f_{(1-i)}(t,l,x) - f_i(t,l,x)\big)\Delta t + o(\Delta t),
$$

which is the standard definition of the generator of a CTMC.

Using the intensity-based representation of default, the CVA expectation can be written as

$$\mathrm{CVA}_t = \delta^R (1 - H_t^R)\, \mathbb{E}^{\mathbb{Q}} \left[\int_t^T e^{-\int_t^s (r + \lambda^R(Y_u))\, du}\, \lambda^R(Y_s)\, v^\phi(s, \widetilde{L}_s, \widetilde{X}_s + \gamma_X(\widetilde{X}_s))\, ds \;\middle|\; \mathcal{F}_t \right].$$

Applying the Feynman–Kac theorem for jump-diffusions with regime switching (see, e.g.,[22]), this expectation is represented as a classical solution f^{CVA} of the backward PIDE:

$$\frac{\partial f_i}{\partial t}(t,l,x) + \mathcal{L} f_i(t,l,x) + \lambda_i^R\, v^\phi(t,l,x + \gamma_X(x)) = (r + \lambda_i^R)\, f_i(t,l,x),$$

with $f_i(T,l,x) = 0, \quad i = 0,1$. Finally, on the set $\{\tau_R > t\}$, the contagion-free processes coincide with the original ones: $\widetilde{L}_t = L_t$, $\widetilde{X}_t = X_t$, and $1 - H_t^R = 1$. Therefore, one obtains

$$\mathrm{CVA}_t = \delta^R (1 - H_t^R)\, f^{\mathrm{CVA}}(t, L_t, X_t, Y_t),$$

providing the desired representation in terms of the observable processes and concluding the proof. □

Proposition 2.1 shows that, also in presence of a regime–switching factor modeled by a continuous–time Markov chain Y, the CVA admits a state–dependent representation. In particular, the value of the CVA at time t depends on the current realization of the regime Y through the function $f^{\mathrm{CVA}}(t, L_t, X_t, Y_t)$.

From a probabilistic point of view, this representation corresponds to a valuation conditional on the event $\{Y_t = y_i\}$, with f_i^{CVA} denoting the CVA function associated with regime y_i. As a consequence, the model yields different CVA levels depending on the prevailing regime, reflecting heterogeneous default intensities and exposure dynamics across regimes.

While this conditional representation is appropriate when the regime is observable or can be reliably inferred, it may be less suitable in situations where the initial state of the Markov chain is unknown or when the regime represents a latent economic or credit condition. In such cases, it is natural to consider a valuation that averages over the possible regimes, leading to a regime–averaged notion of CVA.

Corollary 2.1 (Regime–averaged CVA under stationary distribution)**.** *Assume that the Markov chain Y is ergodic and that its initial distribution coincides with the stationary distribution $\pi = (\pi_0, \pi_1)$. Moreover, assume that Y is independent of the processes L and X. Then, the CVA averaged over the regimes is given by*

$$\overline{\mathrm{CVA}}_t := \mathbb{E}\Big[\mathrm{CVA}_t \;\Big|\; \mathcal{F}_t^{L,X}\Big] = \delta^R(1 - H_t^R)(\pi_0\, f_0^{\mathrm{CVA}}(t, L_t, X_t) + \pi_1\, f_1^{\mathrm{CVA}}(t, L_t, X_t)), \quad (2.4)$$

where f_i^{CVA} with $i \in \{0, 1\}$ denotes the solution of the backward integro–differential equation associated with regime y_i as defined in Proposition 2.1.

Sketch of proof. By definition, the CVA process can be written in regime–dependent form as

$$\mathrm{CVA}_t = \delta^R(1 - H_t^R)\, f_{Y_t}^{\mathrm{CVA}}(t, L_t, X_t),$$

Since H_t^R depends only on Y and Y is assumed independent of L and X, it follows that H_t^R is independent of $\mathcal{F}_t^{L,X}$. Therefore,

$$\mathbb{E}\Big[\mathrm{CVA}_t \;\Big|\; \mathcal{F}_t^{L,X}\Big] = \delta^R(1 - H_t^R)\mathbb{E}\Big[f_{Y_t}^{\mathrm{CVA}}(t, L_t, X_t) \;\Big|\; \mathcal{F}_t^{L,X}\Big].$$

Since the Markov chain Y is assumed to be independent of (L, X) and its initial distribution coincides with the stationary distribution π, ergodicity implies that Y_t is distributed according to π for all $t \geq 0$. Consequently,

$$\mathbb{P}(Y_t = i \mid \mathcal{F}_t^{L,X}) = \pi_i, \qquad i \in \{0, 1\}.$$

It follows that

$$\mathbb{E}\Big[f_{Y_t}^{\mathrm{CVA}}(t, L_t, X_t) \;\Big|\; \mathcal{F}_t^{L,X}\Big] = \pi_0\, f_0^{\mathrm{CVA}}(t, L_t, X_t) + \pi_1\, f_1^{\mathrm{CVA}}(t, L_t, X_t),$$

which concludes the sketch of the proof. □

The regime–averaged CVA does not, in general, satisfy a closed backward equation, since the averaging is performed at the level of the solution rather than at the level of the generator. By assuming that the Markov chain has reached its stationary distribution, the valuation is performed ex ante, averaging over all possible regimes according to their long–run probabilities. The resulting CVA can therefore be interpreted as a long–term or unconditional credit adjustment, which abstracts from short–term regime fluctuations while still capturing their impact through the regime–specific solutions f_i^{CVA}.

From a practical perspective, this regime–averaged CVA is well suited for pricing, risk management and regulatory applications, where robustness with respect to the initial regime is desirable and where regime switching is intended to capture structural rather than transitory features of credit risk.

The regime–averaged CVA introduced in Corollary 2.1 provides a valuation framework that is particularly meaningful when the regime–switching factor represents a latent or only partially observable economic condition. This approach becomes especially relevant in models with multiple regimes, where identifying or estimating the initial state of the Markov chain may be difficult or unreliable. Although the present thesis focuses on a two–state Markov chain, where the identification of the initial regime is comparatively simpler, the regime–averaged CVA is also adopted here for consistency and practical relevance. From a numerical perspective, this choice leads to a single CVA value, which facilitates comparison across different model specifications and provides a common benchmark for subsequent sensitivity analyses and stress–testing exercises.
Having proved the proposition also in this setup, one can proceed with the calculation in a manner analogous to the Example 3.4 in [4], varying only the components of the process Y. Let the loss intensity function λ^L be defined as a non-negative function of the stochastic factor X, i.e.,

$$\lambda_t^L = \lambda^L(X_t), \quad \lambda^L : \mathbb{R} \to \mathbb{R}^+.$$

The stochastic factor X is defined as

$$X_t = \left(\theta + (x_0 - \theta)e^{-\kappa t} + \sigma \int_0^t e^{-\kappa(t-s)}\,\mathrm{d}W_s\right)(1+\gamma)^{H_t^R},$$

which is the solution of the equation

$$\mathrm{d}X_t = \kappa(\theta - X_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t + \gamma X_{t\_}\,\mathrm{d}H_t^R$$

for constant $\kappa > 0$,$\sigma > 0$,$\gamma > 0$ and $\theta > 0$, where the additional component with respect to the classical OU process models the percentage change in the loss intensity at τ_R. By exploiting the previous propositions, we can compute the CVA and the post-default value via a PIDE with generator in the contagion-free regime; therefore, the generator $\mathcal{L}$ is:

$$\begin{aligned}\mathcal{L}^{(\widetilde{L},\widetilde{X})} f(t,l,x) =& \kappa(\theta - x)\frac{\partial f}{\partial x}(t,l,x) + \frac{1}{2}\sigma^2\frac{\partial^2 f}{\partial x^2}(t,l,x) \\ &+ \int_{\mathbb{R}^+} (f(t,l+z,x) - f(t,l,x))\lambda^L(x)\,\nu(dz)\end{aligned} \tag{2.5}$$

Define the function $(t,l,x) \to v^\phi(t,l,x)$ as the solution of the backward integral equation (1.6), where V is substituted with $v^\phi(t,l,x)$ with $(t,l,x) \in [0,T) \times \mathbb{R}^+ \times \mathbb{R}$ and with terminal condition $v^\phi(T,l,x) = \phi(l)$. Then thanks to the previous propositions the postdefault value of the reinsurance contract is given by

$$V^\phi_{\tau_R} = v^\phi(\tau_R, \widetilde{L}_{\tau_R}, \widetilde{X}_{\tau_R}).$$

It is now possible to calculate the CVA, and bearing in mind the previously assumed hypothesis that the chain will reach stationarity, we obtain the equation (2.4) in Proposition 2.1, where the function f_i^{CVA} is calculated in $(t, \widetilde{L}_t, Y_t)$ and it is the solution of the backward equation:

$$\begin{aligned} \frac{\partial f_i^{CVA}}{\partial t}(t,l,x) + \mathcal{L}^{(\widetilde{L},\widetilde{X},Y)} f_i^{CVA}(t,l,x) + \lambda_i^R v^\phi(t,l,x(1+\gamma)) \\ = (\lambda_i^R + r)\, f_i^{CVA}(t,l,x) \end{aligned} \tag{2.6}$$

Denoting by λ_i^R the default intensity in regime i, for every $(t,l,x) \in [0,T) \times \mathbb{R}^+ \times \mathbb{R}$ with terminal condition $f_i^{CVA}(T,l,x) = 0$, the infinitesimal generator $\mathcal{L}^{(\widetilde{L},\widetilde{X},Y)}$ acts componentwise on the vector-valued function $f = (f_0, f_1)$, that is, for each $i \in \{0,1\}$ it operates on f_i separately as follows:

$$\begin{aligned} \mathcal{L}^{(\widetilde{L},\widetilde{X},Y)} f_i(t,l,x) = & \int_{\mathbb{R}^+} \big(f_i(t,l+z,x) - f_i(t,l,x)\big)\lambda^L(x)\,\nu(dz) \\ & + q_{i(1-i)}\big(f_{(1-i)}(t,l,x) - f_i(t,l,x)\big) \\ & + \kappa(\theta - x)\frac{\partial f_i}{\partial x}(t,l,x) + \frac{1}{2}\sigma^2 \frac{\partial^2 f_i}{\partial x^2}(t,l,x), \qquad i \in \{0,1\}. \end{aligned} \tag{2.7}$$

3 | Numerical Implementation

3.1. Numerical Modeling Framework

This chapter presents the numerical implementation and analysis of the credit risk model introduced in the theoretical part of the thesis. The aim is twofold: on the one hand, to illustrate how the model can be effectively implemented using different numerical techniques; on the other hand, to assess the impact of key modeling assumptions on the resulting Credit Value Adjustment (CVA).

The numerical analysis is structured around a common modeling framework, which is then explored under different assumptions on the dynamics of the underlying processes and solved using both Monte Carlo and Partial Integro-Differential Equation (PIDE) methods. This section introduces the core components of the model that are used across all numerical experiments.

The process X represents the stochastic factor driving the intensity of claim arrivals affecting the loss process of the reinsurance contract. Two alternative specifications are considered. In the first case, the intensity is assumed to be constant over time and serves as a benchmark scenario. In the second case, the intensity evolves stochastically according to an Ornstein–Uhlenbeck (OU) dynamics, allowing for time-varying claim activity.

In both specifications, the intensity process is subject to a multiplicative jump of relative size $\gamma \in \mathbb{R}^+$ at the default time of the reinsurer. This jump captures the increase in replacement cost following a default event, reflecting the fact that a newly negotiated contract is typically associated with higher risk and, therefore, higher expected losses.

In the constant-intensity specification, the claim arrival intensity is given, as already considered in [4], by

$$\lambda^L(X_t) = X_t = x_0\big(1 + H_t^R \gamma\big),$$

where $H_t^R = \mathbf{1}_{(\tau_R \leq t)}$ denotes the default indicator of the reinsurer. Under this assumption, the simulation of the loss process is simplified, as the claim arrival rate depends only on the occurrence of default. In this setting, the only source of randomness affecting default probabilities is the regime-switching process Y. This case provides a natural benchmark for the numerical analysis and is particularly useful for validating the Monte Carlo implementation and for comparing stochastic simulations with deterministic results obtained through the PIDE approach.

While the constant-intensity model offers a tractable starting point and captures a basic form of wrong-way risk through the default-induced jump, it does not account for the temporal variability in claim intensity typically observed in reinsurance portfolios. To address this limitation, the model is extended by allowing X to evolve stochastically over time.

In the absence of counterparty default, the stochastic factor X follows an Ornstein–Uhlenbeck dynamics:

$$\mathrm{d}X_t = k(\theta - X_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t,$$

where $k > 0$ is the speed of mean reversion, $\theta > 0$ is the long-term mean level, and $\sigma > 0$ is the volatility parameter. At the default time of the reinsurer τ_R, the factor undergoes a multiplicative jump

$$X_{\tau_R} = (1+\gamma)\,X_{\tau_R^-}, \qquad \gamma \geq 0,$$

where τ_R^- denotes the time just prior to the reinsurance default.

In this scenario, since the probability of the process X attaining negative values is strictly positive, the loss intensity is defined as a lower-bounded transformation of the stochastic factor,

$$\lambda^L(X_t) = \max(\varepsilon, X_t),$$

where $\varepsilon > 0$ is introduced to guarantee strict positivity of the claim arrival intensity.

The choice of the lower bound should ensure positivity without materially altering the intrinsic dynamics of the Ornstein–Uhlenbeck factor. To formalize this requirement, ε is determined through a uniform probabilistic criterion: the lower bound is required to be activated with probability at most p, uniformly in time, where p is a small tolerance level. More precisely, p is such that

$$\sup_{t\geq 0} \mathbb{P}(X_t \leq \varepsilon) \leq p, \qquad 0 < p \ll 1.$$

This condition ensures that the truncation affects the dynamics only in rare events whose probability can be explicitly controlled.

The Ornstein–Uhlenbeck process is Gaussian at every time t, with distribution

$$X_t \sim \mathcal{N}\big(m(t), v(t)\big),$$

where

$$m(t) = \theta + (X_0 - \theta)e^{-kt}, \qquad v(t) = \frac{\sigma^2}{2k}\left(1 - e^{-2kt}\right).$$

The variance $v(t)$ is increasing in time and converges to its stationary value $\sigma^2/(2k)$, while the mean converges to θ. Consequently, for any fixed ε,

$$\mathbb{P}(X_t \leq \varepsilon) = \Phi\left(\frac{\varepsilon - m(t)}{\sqrt{v(t)}}\right),$$

where Φ denotes the standard normal distribution function. Since the variance attains its maximum asymptotically and the mean converges to θ, the supremum over time is achieved in the stationary limit. Hence,

$$\sup_{t\geq 0} \mathbb{P}(X_t \leq \varepsilon) = \Phi\left(\frac{\varepsilon - \theta}{\sigma/\sqrt{2k}}\right).$$

Imposing the probabilistic constraint yields

$$\Phi\left(\frac{\varepsilon - \theta}{\sigma/\sqrt{2k}}\right) \leq p.$$

Defining

$$z := \frac{\theta - \varepsilon}{\sigma/\sqrt{2k}},$$

this condition becomes $\Phi(-z) \leq p$, which is equivalent to $\Phi(z) \geq 1 - p$. Therefore,

$$z = \Phi^{-1}(1 - p),$$

and the corresponding lower bound is explicitly given by

$$\varepsilon(p) = \theta - \Phi^{-1}(1 - p)\frac{\sigma}{\sqrt{2k}}.$$

However, since the claim arrival intensity must remain non–negative by construction, the effective lower bound is defined as

$$\varepsilon := \max\{0, \varepsilon(p)\}.$$

In addition to satisfying the probabilistic constraint, the model parameters must ensure that the truncation level is strictly positive whenever possible. In particular, the condition

$$\theta > \Phi^{-1}(1-p)\frac{\sigma}{\sqrt{2k}} \tag{3.1}$$

guarantees that $\varepsilon(p) > 0$, so that the probabilistic criterion alone determines the lower bound.

If this compatibility condition is violated, then $\varepsilon(p) \leq 0$ and the effective lower bound reduces to $\varepsilon = 0$. In this case, the intensity is simply reflected at zero and the probabilistic constraint cannot be simultaneously enforced with strict positivity for the given parameter configuration.

The tolerance level p will be specified in the numerical implementation section. With this specification, the intensity process remains non–negative by construction, while the probability that the lower bound alters the natural dynamics of X is uniformly controlled whenever the model parameters are compatible with the probabilistic requirement.

A useful limiting case is obtained by setting $\sigma = 0$ and $k = 0$. When the initial condition satisfies $X_0 = \theta$, the intensity remains constant over time and the stochastic-intensity model reduces to the constant-intensity benchmark.

Counterparty default risk is modeled independently from the claim arrival process through a regime-switching default intensity. Specifically, the default intensity is governed by a two-state CTMC $Y = (Y_t)_{t\geq 0}$ with $Y_t \in \{0, 1\}$, with transition intensities q_{01} and q_{10}.

Each state of the chain corresponds to a different level of default intensity:

$$\lambda^R(Y_t) = \begin{cases} \lambda_0^R, & \text{if } Y_t = 0, \\ \lambda_1^R, & \text{if } Y_t = 1. \end{cases}$$

This formulation allows the default intensity to switch between different regimes, capturing sudden changes in the counterparty's credit quality. The CTMC framework is particularly convenient for both Monte Carlo simulation and PIDE-based valuation, as it leads to a system of coupled equations, consisting of two integro-differential equations, one for each regime of Y.

3.1.1. Construction of the Default Time

The default time τ is defined using the following standard construction. Let $\Theta \sim \mathrm{Exp}(1)$ be an exponential random variable independent of Y. The default time is given by

$$\tau := \inf \left\{ t \geq 0 : \int_0^t \lambda(Y_s)\, ds \geq \Theta \right\}. \tag{3.2}$$

For numerical purposes, especially in the Monte Carlo framework, τ is simulated using an iterative competing-risks algorithm, that is, an estimation method for event probabilities in the presence of multiple causes of failure that compete to determine the primary outcome (we refer the reader to Appendix C.1 for the algorithm). On each time interval during which Y remains in a fixed state, two exponential clocks are sampled: one corresponding to the next state transition of Y, and one corresponding to default under the current intensity level. In particular, let $\overline{\tau}$ denote the default time constructed iteratively and consider the following iterative simulation algorithm:

- start at time $t = 0$ from the initial state $Y_0 = i$;
- on each time interval during which Y remains in state i, generate two independent exponential random variables

  $$T_d \sim \mathrm{Exp}(\lambda_i), \qquad T_y \sim \mathrm{Exp}(q_i),$$

  representing, respectively, the default time and the next state–transition time;
- if $T_d < T_y$, set $\overline{\tau} = t + T_d$;
- if $T_y < T_d$, update $t \leftarrow t + T_y$, change state and repeat;
- if a fixed time horizon T is exceeded, set $\overline{\tau} = T$.

The equivalence between τ and $\overline{\tau}$ is formalized in the following proposition.

Proposition 3.1 (Equivalence between integral and iterative constructions of the default time). *Let $Y = (Y_t)_{t \geq 0}$ be a two–state continuous–time Markov chain (CTMC) with state space $\{0, 1\}$ and transition intensities $q_0 := q_{01}, q_1 := q_{10}$. Let $\lambda : \{0, 1\} \to \mathbb{R}^+$ be a default intensity function and let $\Theta \sim \mathrm{Exp}(1)$ be a standard exponential random variable, independent of the process Y.*

Define the default times

$$\tau := \inf \left\{ t \geq 0 : \int_0^t \lambda(Y_s)\, ds \geq \Theta \right\},$$

and $\overline{\tau}$ costructed as above. Then the random variables τ and $\overline{\tau}$ have the same law.

Proof. The proof proceeds by induction over the successive inter–jump intervals of the process Y.

First interval. Assume that $Y_0 = i$ and denote by

$$S := \inf\{t > 0 : Y_t \neq i\}$$

the time of the first state transition. Since Y is a two–state CTMC,

$$S \sim \mathrm{Exp}(q_i),$$

and S is independent of Θ.

Conditionally on $S = s$, the default intensity is constant and equal to λ_i on the interval $[0, s)$. By definition of τ, for any $t < s$,

$$\{\tau \leq t\} = \left\{ \Theta < \int_0^t \lambda_i \, du \right\} = \{\Theta < \lambda_i t\}.$$

Therefore, the conditional distribution function of τ given $S = s$ is

$$\mathbb{P}(\tau \leq t \mid S = s) = \mathbb{P}(\Theta \leq \lambda_i t) = 1 - e^{-\lambda_i t}, \qquad 0 \leq t < s.$$

Differentiating with respect to t yields the conditional density of τ on $[0, s)$,

$$f_{\tau|S=s}(t) = \frac{d}{dt}\mathbb{P}(\tau \leq t \mid S = s) = \lambda_i e^{-\lambda_i t}, \qquad 0 \leq t < s,$$

that is, the density of an $\mathrm{Exp}(\lambda_i)$ random variable truncated at s. In particular,

$$\mathbb{P}(\tau < S \mid S = s) = 1 - e^{-\lambda_i s} = \mathbb{P}(T_d < s).$$

Therefore, on the first interval, comparing T_d with S correctly reproduces both the event of default before the state transition and the conditional distribution of the default time.

State transition and threshold update. Assume now that $T_d \geq S$, i.e., no default occurs on the interval $[0, S]$. Then

$$\Theta > \int_0^S \lambda_i \, du = \lambda_i S.$$

Define the residual threshold

$$\Theta' := \Theta - \lambda_i S.$$

Since $\Theta \sim \text{Exp}(1)$, by the memoryless property (Theorem 2.3.1 of [18]) of the exponential distribution,

$$\Theta' \mid \{\Theta > \lambda_i S\} \sim \text{Exp}(1),$$

and Θ' is independent of the past.

Consequently, at the state transition time the problem regenerates with the same probabilistic structure as at time zero: the residual default time is given by

$$\inf \left\{ u \geq 0 : \int_0^u \lambda(Y_{S+v}) \, dv > \Theta' \right\},$$

with $\Theta' \sim \text{Exp}(1)$ independent of the future evolution of Y.

Inductive step. Repeating the same argument on each inter–jump interval of the process Y, it follows that the iterative algorithm generates a default time $\overline{\tau}$ with the same distribution as τ obtained from the integral construction. $\square$

This result provides a theoretical justification for simulating default times by comparing independent exponential random variables. The construction is standard in reduced–form credit risk models and is widely used in Monte Carlo simulations for credit and counterparty risk. At the same time, the CTMC structure of Y allows the default component to be incorporated naturally into the PIDE framework through a system of coupled equations, one for each regime.

In the remainder of this chapter, this common modeling framework is employed to compare Monte Carlo and PIDE methods under different assumptions on the stochastic factor X, starting from the constant-intensity benchmark and progressively extending to the stochastic OU case.

3.2. Monte Carlo Method

This section describes the Monte Carlo methodology adopted to estimate the CVA within the modeling framework introduced in the previous section. The Monte Carlo approach provides a flexible and intuitive benchmark, as it allows for a direct simulation of all the stochastic components of the model, including the claim arrival intensity, the regime-switching default intensity, and the default time.

The CVA estimator is based on the pathwise simulation of the three stochastic processes driving the model, namely the stochastic factor X, the regime process Y and the default time τ_R. For each Monte Carlo path, the trajectory of Y and the default time τ_R are simulated jointly using the competing-risks algorithm described in Section 3.1.1. Conditional on the realization of τ_R, the loss process is evolved up to default and the exposure at default is computed accordingly.

Let N_{paths} denote the total number of Monte Carlo simulations. The CVA is approximated by the empirical average

$$\widehat{\mathrm{CVA}}^{MC} = \frac{1}{N_{paths}} \sum_{m=1}^{N_{paths}} \delta^R \, e^{-r\tau_R^{(m)}} \, \Phi\big(L^{(m)}_{\tau_R^{(m)}}\big) \, \mathbf{1}_{\{\tau_R^{(m)} \le T\}},$$

where Φ denotes the exposure at default and δ^R is the loss-given-default parameter. An estimate of the statistical error is obtained from the sample standard deviation of the Monte Carlo estimator, allowing for the construction of confidence intervals and for a quantitative assessment of numerical accuracy.

In the baseline specification, the stochastic factor X is assumed to be constant, taking value x_{pre} before default and x_{post} after, where $x_{\text{post}} = (1+\gamma)x_{\text{pre}}$, in accordance with the jump-to-default hypothesis of the model. For each simulated path, the default time of the reinsurer is generated as a first step; if default occurs before maturity (i.e. $(\tau_R < T)$), the cumulative loss process is simulated up to the default time τ_R using the pre-default intensity x_{pre} and once the default event has occurred, a nested Monte Carlo procedure is employed to estimate the continuation value $v^{\phi}(\tau_R, L_{\tau_R})$, by simulating multiple post-default trajectories of the loss process over the interval $[\tau_R, T]$ under the post-default intensity x_{post}. The CVA is then obtained as the discounted average of losses across all paths in which default occurs.

The model is subsequently extended by allowing the claim arrival intensity, which becomes $\lambda^L(X_t), t \ge 0$, to evolve stochastically over time. In this case, X follows an Ornstein–Uhlenbeck dynamic, introducing mean reversion and capturing time-varying claim activity. The main additional numerical ingredient is the discretization of the stochastic factor, which is performed using an Euler–Maruyama scheme over a fixed time grid. Conditional on default, the post-default evolution of both $\lambda^L(X_t)$ and the cumulative loss process L is simulated using a nested Monte Carlo approach. To mitigate the computational burden induced by the additional source of randomness, the inner simulations are vectorized, allowing multiple post-default scenarios to be generated simultaneously.

Overall, the Monte Carlo method provides a flexible and robust valuation tool, capable of handling complex dynamics and path-dependent features. Its main limitation lies in the relatively high computational cost required to achieve a satisfactory level of statistical accuracy, especially in the presence of stochastic intensities and nested simulations. For this reason, Monte Carlo estimates are primarily used as a reference benchmark in the numerical analysis, while more efficient deterministic methods are considered.

3.3. PIDE Approach

In this section, the valuation of the CVA is addressed through a deterministic Partial Integro-Differential Equation (PIDE) approach. By exploiting the Markovian structure of the model, the CVA can be characterized as the solution of a backward equation, providing a computationally efficient alternative to Monte Carlo simulation in specific settings. This approach shifts the problem from the probabilistic domain of pathwise simulations to the deterministic framework of numerical analysis. In particular, the PIDE describes the evolution of the contract value by combining a differential operator associated with continuous dynamics and an integral operator accounting for the jump structure of the loss process and the default event.

The cumulative loss process L is defined on a truncated spatial domain $[0, L_{\max}]$, where $L_{\max}$ is chosen sufficiently large so that truncation effects are negligible. A uniform spatial grid is introduced as

$$l_i = i\Delta l, \qquad i = 0, \dots, N_l - 1,$$

with spacing $\Delta l = L_{\max}/(N_l - 1)$. Losses evolve through upward jumps whose sizes follow a Gamma distribution with density $g(z)$. The integral part of the infinitesimal generator acting on a test function f is given by

$$\mathcal{L}f(l) = \int_0^\infty \big[f(l+z) - f(l)\big] g(z)\, dz.$$

The integration domain is truncated at $z_{\max}$ and discretized into J_z intervals. A midpoint quadrature rule is applied, yielding

$$\int_0^{z_{\max}} g(z)\,dz \approx \sum_{j=1}^{J_z} g(z_j)\,\Delta z.$$

Any residual probability mass beyond $z_{\max}$ is assigned to the last quadrature node in order to preserve normalization. Since the shifted location $l_i + z_j$ generally does not coincide with a grid point, linear interpolation between the two nearest nodes is employed. This construction leads to a sparse matrix representation of the integral operator,

$$\mathcal{L}f \;\approx\; \mathcal{M}f,$$

where $\mathcal{M} \in \mathbb{R}^{N_l \times N_l}$ is sparse. As a consistency check, the matrix approximately satisfies

$$\mathcal{M}\mathbf{1} \approx 0,$$

ensuring mass conservation.

Having defined the spatial discretization, consider the value function $v^{\phi}(t,l)$ associated with a terminal payoff $\phi(l)$. In the constant claim arrival intensity case, the pricing PIDE reads

$$\frac{\partial v^{\phi}}{\partial t}(t,l) + X\mathcal{M}v^{\phi}(t,l) - rv^{\phi}(t,l) = 0, \qquad v^{\phi}(T,l) = \phi(l).$$

Time discretization is performed using the Crank–Nicolson scheme, which is second-order accurate and unconditionally stable. The resulting semi-discrete formulation is

$$\left(I - \tfrac{1}{2}\Delta t A\right) v^{n} = \left(I + \tfrac{1}{2}\Delta t A\right) v^{n+1}, \qquad A = X\mathcal{M} - rI.$$

The associated linear system is solved backward in time. Since the matrix on the left-hand side does not depend on time, its LU factorization is computed once and reused across all time steps, resulting in an efficient numerical procedure. The full value surface $v^{\phi}(t_n, l_i)$ is obtained as output.

Let $f_0^{\mathrm{CVA}}(t,l)$ and $f_1^{\mathrm{CVA}}(t,l)$ denote the CVA value functions corresponding to the two default regimes of the continuous-time Markov chain Y. These functions satisfy the coupled PIDE system

$$\begin{aligned}
\frac{\partial f_0^{\mathrm{CVA}}}{\partial t} + X\mathcal{M}f_0^{\mathrm{CVA}} - (r + \lambda_0 + q_{01})f_0^{\mathrm{CVA}} + q_{01}f_1^{\mathrm{CVA}} + \lambda_0 v^{\phi} = 0,\\
\frac{\partial f_1^{\mathrm{CVA}}}{\partial t} + X\mathcal{M}f_1^{\mathrm{CVA}} - (r + \lambda_1 + q_{10})f_1^{\mathrm{CVA}} + q_{10}f_0^{\mathrm{CVA}} + \lambda_1 v^{\phi} = 0,
\end{aligned}$$

with terminal conditions

$$f_0^{\mathrm{CVA}}(T,l) = f_1^{\mathrm{CVA}}(T,l) = 0.$$

The system is discretized in space using the same matrix $\mathcal{M}$ and in time using the Crank–Nicolson scheme. The two equations are assembled into a block matrix system and solved backward in time, with the source terms involving v^{ϕ} integrated consistently via trapezoidal time discretization.

The CVA is finally extracted from the solution by evaluating the functions at the initial loss level $l_0 = 0$ and time $t = 0$,

$$\mathrm{CVA}_i = \delta^R f_i^{\mathrm{CVA}}(0, l_0), \qquad i \in \{0, 1\}.$$

Assuming that the default regime starts from its stationary distribution $\pi = (\pi_0, \pi_1)$, the overall CVA is given by

$$\mathrm{CVA} = \delta^R\big(\pi_0 f_0^{\mathrm{CVA}}(0, l_0) + \pi_1 f_1^{\mathrm{CVA}}(0, l_0)\big).$$

The framework can be naturally extended to the case of stochastic claim arrival intensity by allowing X to follow an Ornstein–Uhlenbeck dynamics. The infinitesimal generator associated with the OU process acts on a test function V is given by

$$\mathcal{L}_X V = \kappa(\theta - x)\,\frac{\partial V}{\partial x} + \frac{1}{2}\sigma^2\,\frac{\partial^2 V}{\partial x^2},$$

which is added to the operators already present in the model. The intensity variable is discretized on a uniform grid $\{x_i\}_{i=0}^{N_x-1}$, where N_x is chosen sufficiently large to ensure numerical stability, with spacing Δx. The drift term is approximated using an upwind finite-difference scheme, ensuring numerical stability, while the diffusion term is discretized using second-order centered differences. For an interior grid point x_i, with $i = 1, \dots, N_x - 2$, the second derivative is approximated as

$$\frac{1}{2}\sigma^2\,\frac{\partial^2 V}{\partial x^2}(x_i) \;\longrightarrow\; \frac{\sigma^2}{2\Delta x^2}\left(V_{i+1} - 2V_i + V_{i-1}\right).$$

The resulting operator is represented by a sparse tridiagonal matrix, preserving the local structure of the OU dynamics and enabling efficient computation.

The multiplicative jump in intensity at counterparty default introduces a non-local effect in the intensity dimension, therefore, at default, the value function is evaluated at the shifted intensity level $x' = (1+\gamma)x$. Since this point does not generally coincide with a grid node, linear interpolation is employed. At the discrete level, this mechanism is captured by a jump matrix J_x, which maps the value function from the pre-default to the post-default intensity. The matrix is sparse, with at most two non-zero entries per row, ensuring both interpolation accuracy and computational efficiency.

Overall, the extension to a stochastic intensity integrates directly into the PIDE framework developed for the constant-intensity case. The OU generator replaces the null operator associated with X, the jump matrix accounts for the default-induced intensity shift, the overall structure of the PIDE and its numerical resolution remains unchanged. From a numerical standpoint, the same routines are employed to compute both the contract value v^{ϕ} and the CVA in the constant and stochastic intensity settings, with the only modification concerning the dynamics of X, as illustrated by the pseudocodes reported in Appendix C.2 and C.3. This formulation allows for a direct and coherent comparison between the constant and stochastic intensity models, isolating the impact of intensity dynamics on the CVA.

4 | Results and Analysis

This chapter presents the numerical results of the model and the corresponding sensitivity analysis of the Credit Value Adjustment (CVA). First, the set of financial, credit-related and numerical parameters adopted in the computations is introduced and motivated. Then, the consistency between the Monte Carlo and the PIDE numerical approaches is assessed in a benchmark configuration. Finally, a detailed sensitivity analysis is carried out, focusing separately on the parameters governing the loss process X, the credit regime process Y, and the remaining contractual and credit-related parameters, with particular attention devoted to those exhibiting a significant economic or financial impact on the CVA.

4.1. Parameter specification

In the following analysis, the model setup is specified as follows. The financial and contractual parameters, namely the yearly risk-free interest rate $r = 0.03$, the contract maturity $T = 10$y, the upper and lower limits of the stop-loss contract $\overline{K} = 3$ and $\underline{K} = 0.9$, and the loss given default $\delta^R = 0.6$, are taken from [4] or are consistent with values commonly adopted in the related literature. In particular, the stop-loss thresholds $\overline{K}$ and $\underline{K}$ are chosen to be consistent with the scale of individual claim amounts, so that the contract limits are meaningful relative to the distribution of losses. Similarly, the parameters governing the dynamics of the stochastic factor X, namely the volatility $\sigma = 0.1$, the long-term mean $\theta = 1$, the mean-reversion speed $k = 0.5$, and the loss amplification parameter $\gamma = 0.2$, are calibrated according to the same reference. In order to ensure strict positivity of the claim arrival intensity while preserving the intrinsic dynamics of the Ornstein–Uhlenbeck factor, the tolerance level introduced in Section 3.1 is fixed at $p = 10^{-5}$. This choice implies that the probability of the stochastic factor X being truncated by the lower bound remains uniformly below 10^{-5} over time, provided that the compatibility condition on the parameters is satisfied. With the baseline volatility $\sigma = 0.1$, the resulting lower bound $\varepsilon(p)$ is strictly positive, so that the truncation is activated only in extremely rare events and does not materially affect the natural dynamics of the factor.

The claim size distribution is assumed to follow a Gamma distribution with parameters $(1, 1)$, a common specification in collective risk modelling due to its analytical tractability and flexibility (see, e.g., [14]). This choice ensures that the stop-loss limits $\overline{K}$ and $\underline{K}$ are appropriately scaled relative to the typical claim size, maintaining coherence between the claim distribution and the contract thresholds. The numerical parameters, including the number of discretization points for the grids of the state variables and time, namely $N_x = 101$, $N_L = 251$ and $N_t = 1000$, as well as the upper boundaries of the computational domain, set to $L_{\max} = 20$ and $X_{\max} = 10$, are selected on the basis of a robustness analysis. This analysis aims at ensuring negligible approximation errors while maintaining a reasonable balance between numerical accuracy and computational efficiency. This aspect is particularly relevant for the Monte Carlo approach, where computational costs grow significantly with the number of simulated paths, which is set to $N_{\text{paths}} = 50{,}000$ throughout the analysis. The parameters driving the regime-switching process Y admit a clear macroeconomic interpretation, as the process is a continuous-time Markov chain describing the evolution of the economic environment. In particular, following [13], the default intensity in the normal regime, $\lambda_0 = 0.01$, is set at levels consistent with historical estimates for investment-grade financial institutions, while a significantly higher intensity $\lambda_1 = 0.05$ is assumed in the stressed regime, reflecting the increase in credit risk typically observed during systemic crises. Concerning the transition intensities, in line with [1] and with empirical evidence on financial and business cycles duration reported in [10], the calm state is assumed to have an average duration of approximately eight years, corresponding to a transition rate $q_{01} = 0.125$. Conversely, the stressed state is characterized by a shorter average duration of about one and a half years, leading to a transition rate $q_{10} = 0.66$.

All parameters introduced in this section will be varied in the subsequent sensitivity analysis in order to assess their impact on the CVA and to provide an economic interpretation of the model behaviour.

Table 4.1: Parameters

r	T	$\overline{K}$	$\underline{K}$	δ^R	σ	k	θ	p	γ	λ_0	λ_1	q_{01}	q_{10}
0.03	10	3	0.9	0.6	0.1	0.5	1	10^{-5}	0.2	0.01	0.05	0.125	0.66

Table 4.2: Comparison MC and PIDE approach

	MC CVA	CI	Accuracy	CT MC	PIDE CVA	CT PIDE
X const	0.1969	[0.1928 , 0.2010]	2.09%	5s	0.1953	1s
X OU	0.1969	[0.1928 , 0.2010]	2.08%	10m	0.1939	2m

4.2. Monte Carlo vs PIDE: base case comparison

Table 4.2 reports the CVA estimates obtained under both constant and stochastic claim intensity specifications. For each configuration, the CVA is computed using both the Monte Carlo and the PIDE approaches in order to assess the consistency between the two numerical frameworks.

In both cases considered, the CVA obtained through the PIDE approach lies well within the Monte Carlo confidence interval. This result confirms the consistency of the two numerical methods and provides validation for the PIDE implementation, which will be employed in the subsequent sensitivity analysis due to its numerical stability and significantly lower computational cost.

The role of the initial credit regime is further illustrated in Figure 4.1, which shows the time evolution of the CVA obtained with the PIDE approach in the stochastic intensity setting, conditional on the initial state of the regime-switching process. Specifically, the CVA is computed assuming that the system starts at time $t = 0$ in the normal regime ($Y_0 = 0$), in the stressed regime ($Y_0 = 1$), and under the stationary distribution of the Markov chain. Under the selected calibration of the transition intensities q_{01} and q_{10}, the stationary distribution of the two-state regime-switching Markov chain is given by

$$\pi_0 = \frac{q_{10}}{q_{01} + q_{10}} = 0.8407, \qquad \pi_1 = \frac{q_{01}}{q_{01} + q_{10}} = 0.1593.$$

As expected, the CVA is significantly higher when the initial state corresponds to the

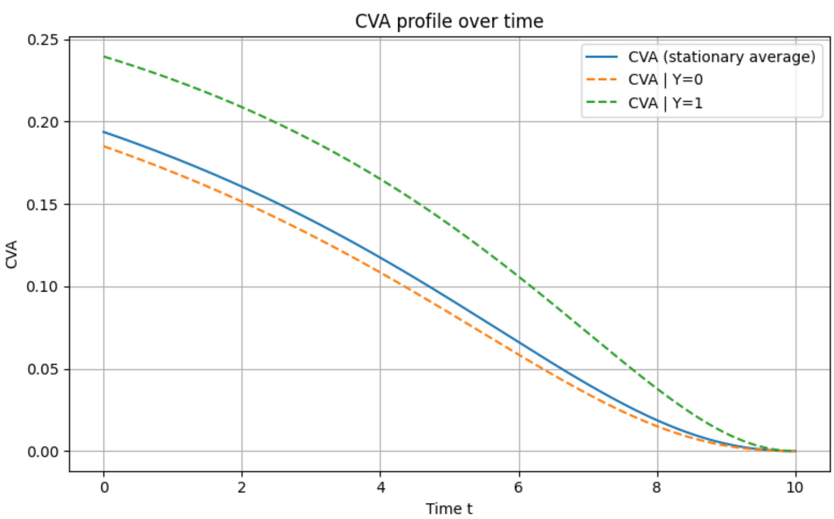

Figure 4.1: CVA vs time on different initial state of Y

stressed regime, reflecting the higher default intensity. Over time, the difference between the conditional CVA profiles gradually decreases as the effect of the initial condition vanishes and the distribution of the regime converges to its stationary measure. For this reason, all subsequent numerical analyses are conducted assuming that the initial regime is distributed according to the stationary distribution.

Based on the strong agreement between the Monte Carlo and PIDE results and on the analysis of the initial regime effect, the PIDE framework with process X following an Ornstein–Uhlenbeck dynamics is adopted in the remainder of the section to perform a detailed sensitivity analysis.

4.3. Sensitivity analysis

The sensitivity analysis focuses on the impact of the model parameters on the CVA. Particular attention is devoted to those parameters that produce economically and financially meaningful variations, while parameters with marginal effects are discussed only briefly.

Parameters of the process X. Figure 4.2 reports the sensitivity of the CVA with respect to the volatility σ, the long-term mean θ and the mean-reversion speed k of the Ornstein–Uhlenbeck process. The parameter σ captures the variability of the claim intensity over time, that is, how much and how often the claim intensity fluctuates around its long-term level. An increase in σ does not change the average level of the intensity, but makes its dynamics more irregular, increasing the likelihood of observing both high-intensity and low-intensity phases. At first sight, these two effects might appear to offset each other. However, the numerical results show that the CVA decreases as volatility increases. This behavior can be explained by the way the CVA is constructed and by the structure of the stop–loss contract. The CVA depends on the value of the contract at the time of counterparty default, which in turn is determined by the losses accumulated up to that moment. These losses build up over time and are capped by the stop–loss payoff. When volatility is low, periods of high claim intensity tend to be persistent. This allows losses to accumulate steadily, increasing the likelihood that the contract has a high value when default occurs. In this case, the exposure to counterparty risk is larger, leading to a higher CVA. When volatility is high, the claim intensity becomes much less stable. Although high intensity levels may occur, they tend to be short-lived and are often followed by rapid reversions to lower levels. As a result, it is less likely that a phase of high intensity lasts long enough to generate substantial loss accumulation before default. At the same time, periods of low intensity significantly reduce the contract value, while periods of high intensity have a more limited impact because of the payoff cap. Overall,

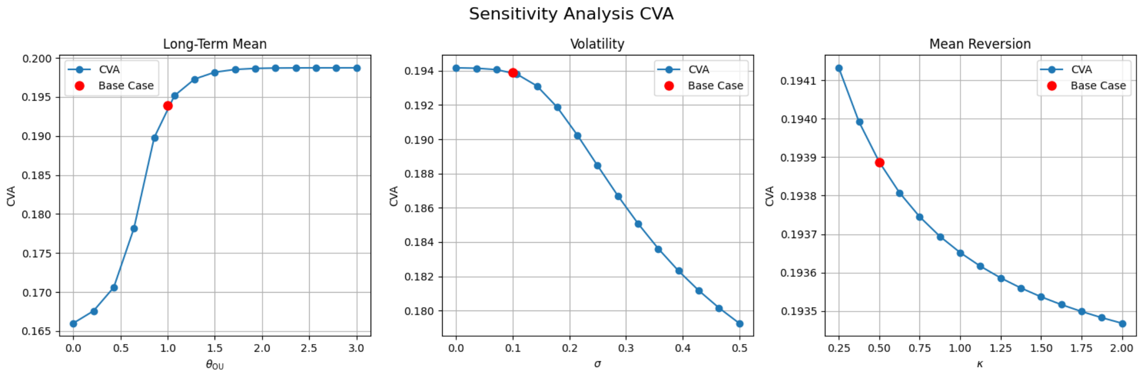

Figure 4.2: sensitivity of the OU parameters

the negative effects of low intensity phases are not fully offset by the positive effects of high intensity phases. Higher volatility therefore reduces the persistence of insurance risk over time and lowers the expected exposure at default, resulting in a lower CVA. It is worth noting that the lower-bound construction discussed in Section 3.1 interacts with the volatility parameter. Given the baseline values $\theta = 1$, $k = 0.5$ and $p = 10^{-5}$, the compatibility condition (3.1) is satisfied only for $\sigma < \sigma^* \approx 0.234$. For volatility levels exceeding this threshold, the quantity $\varepsilon(p)$ becomes non–positive and the effective lower bound reduces to $\varepsilon = 0$. In this region, the intensity process remains non–negative by construction, but it is no longer guaranteed that the truncation occurs with probability smaller than 10^{-5}. Therefore, for $\sigma > 0.234$, the probabilistic interpretation of the lower bound as a rare-event correction is no longer strictly valid, and the truncation mechanism effectively acts as a reflection at zero.

The effect of θ on the CVA is strongly non-linear: while low levels of the structural intensity lead to a significant reduction in counterparty risk, for higher values the CVA exhibits a saturation behavior due to the capped nature of the stop–loss payoff. This confirms that the CVA is particularly sensitive to the long-term level of credit risk embedded in the intensity process.

The CVA exhibits only a mild decreasing trend as k increases, indicating that the speed of mean reversion plays a secondary role compared to the level and volatility of the intensity.

Parameters of the regime-switching process Y. Figure 4.3 illustrates the sensitivity of the CVA with respect to the default intensity λ_1 associated with the stressed regime and the effect of the transition intensities q_{01} and q_{10} by jointly scaling both parameters by a common multiplicative factor. The CVA increases sharply as λ_1 rises, highlighting the dominant role of stressed credit conditions in the determination of counterparty credit risk. Even moderate increases in λ_1 lead to substantial changes in the CVA, emphasizing

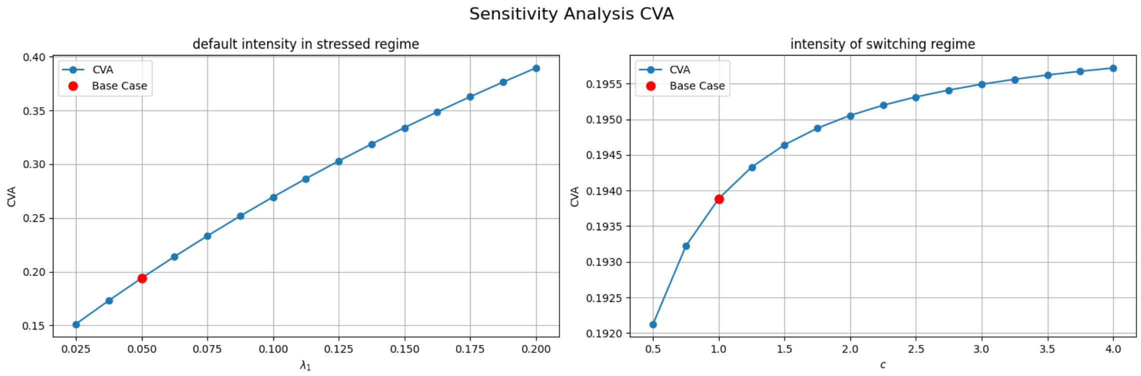

Figure 4.3: sensitivity of the default intensity parameters

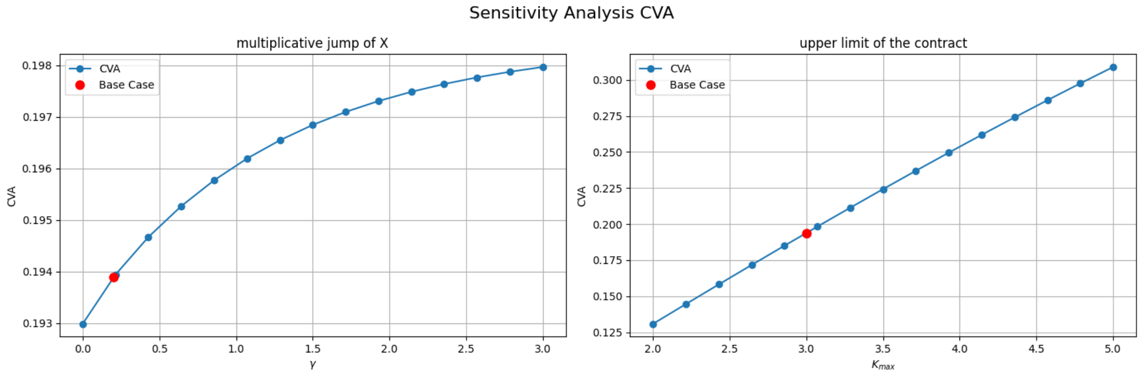

Figure 4.4: sensitivity of γ and $K_{\max}$

the importance of accurately modelling extreme credit scenarios.

An increase in the transition rates results in a moderate increase in the CVA, reflecting a higher likelihood of entering the stressed regime over the contract horizon. However, the impact remains limited compared to that of λ_1, suggesting that the level of credit risk within each regime is more relevant than the frequency of regime switches.

Other credit and contractual parameters. Figure 4.4 reports the sensitivity of the CVA with respect to the parameter γ and the impact of the upper limit $K_{\max}$ of the stop-loss contract.

The CVA increases in a nonlinear fashion as γ rises, with the slope gradually decreasing and the curve flattening for higher values of the parameter. This concave behaviour indicates that, while larger values of γ amplify credit losses, their marginal impact on the CVA becomes progressively weaker. This result is consistent with the interpretation of γ as a loss amplification parameter, capturing the effect of adverse market conditions on default-related losses. Moreover, this analysis emphasizes the importance of explicitly accounting for γ in CVA computations, as it provides a mechanism to model the increase in market-wide risk and its transmission to counterparty credit exposure, particularly during periods of financial stress. The CVA increases significantly with $K_{\max}$, as higher loss caps imply larger exposure in default scenarios. This result highlights the strong interaction between contractual features and counterparty credit risk.

Overall, the sensitivity analysis shows that the CVA is mainly driven by parameters affecting the level of credit risk and the contractual exposure, while parameters governing the dynamics of the intensity and the regime-switching mechanism play a secondary, though non-negligible, role.

4.4. Monte Carlo consistency check

To further validate the results obtained with the PIDE approach, a consistency check is performed using Monte Carlo simulations as a benchmark. Starting from a baseline calibration, we consider alternative specifications in which one parameter at a time is shifted to a stressed value, while all remaining parameters are kept fixed at their benchmark levels. This approach allows for a direct comparison between the two numerical methods without repeating the entire sensitivity analysis. For each stressed configuration, the CVA is computed using both Monte Carlo simulation and the PIDE approach, assuming that the process X follows an Ornstein–Uhlenbeck dynamics. Table 4.3 reports, for each single-parameter stressed, the Monte Carlo CVA estimate together with its 95% confidence interval, and the corresponding PIDE value. In each row, only the parameter indicated in the first column is modified with respect to the baseline calibration; all other parameters remain unchanged. For all stressed configurations considered, the PIDE estimate lies within the Monte Carlo confidence interval. Moreover, both the direction and the magnitude of the CVA variations induced by the parameter shifts are consistent across the two numerical methods. This confirms that the sensitivity patterns identified in the PIDE-based analysis are not numerical artefacts, but reflect genuine model dynamics captured coherently by both Monte Carlo simulation and the PIDE framework.

Table 4.3: Comparison MC and PIDE on different parameter values

	MC CVA	Confidence Interval	PIDE CVA
$\gamma = 2$	0.1972	[0.1931 , 0.2014]	0.1973
$\sigma = 0.3$	0.1902	[0.1862 , 0.1943]	0.1897
$\lambda_1 = 0.2$	0.3878	[0.3826 , 0.3930]	0.3892
$K_{\max} = 4$	0.2588	[0.2534 , 0.2642]	0.2535

5 | Conclusions and future developments

This thesis has investigated the problem of Reinsurance Counterparty Credit Risk (RCCR) within a reduced–form, market–consistent framework, with particular focus on the computation of the Credit Value Adjustment (CVA) for reinsurance contracts. Starting from the model proposed by Ceci, Colaneri, Frey and Köck [4], the analysis combined theoretical modeling and numerical implementation in order to assess the impact of reinsurer default risk on the valuation of reinsurance agreements. The work has progressively moved from the general formulation of the framework to a specific and tractable case study, concluding with a numerical comparison between Monte Carlo simulation and PIDE–based methods.

The modeling framework adopted throughout the thesis is based on an intensity–based approach, in which both insurance losses and counterparty default are described by Cox processes with stochastic intensities. Aggregate losses are generated by a counting process whose intensity is driven by a latent factor process, while the default time of the reinsurer is modeled through a reduced–form hazard rate. This structure allows for a flexible and economically meaningful representation of insurance risk and credit risk within a unified probabilistic setting. A key feature of the framework is the explicit treatment of contagion effects at the time of default, whereby the claims arrival intensity may experience a jump when the reinsurer defaults, capturing the empirical interaction between counterparty distress and insurance risk dynamics. The introduction of a contagion–free auxiliary market further enables a clean and rigorous derivation of valuation equations, leading to backward Partial Integro–Differential Equations (PIDEs) for both the contract value and the CVA.

Within this general setting, the main contribution of the thesis consists in a specific extension of the model of Ceci et al. to the case where the claims arrival intensity factor X follows an Ornstein–Uhlenbeck (OU) process, while the default–related factor Y is modeled as a two–state continuous–time Markov chain (CTMC). This specification introduces a regime–switching default environment, representing normal and stressed market conditions, each associated with a distinct default intensity. At the same time, the OU dynamics for X capture mean–reverting behavior in claim arrivals, which is well supported by empirical evidence in insurance and reinsurance markets. The resulting model preserves the structural consistency of the original framework while significantly enhancing its economic interpretability and numerical tractability.

From a theoretical perspective, the thesis shows that, under the OU–CTMC specification, the CVA admits a regime–dependent representation and satisfies a system of coupled backward integro–differential equations. This result extends the original propositions of Ceci et al. to a hybrid continuous–discrete setting, where diffusion, jump and regime–switching components coexist. The introduction of a regime–averaged CVA, computed under the stationary distribution of the Markov chain, provides an economically meaningful valuation when market regimes are latent or only partially observable. This formulation is particularly relevant for practical applications, as it delivers a single, robust CVA measure that incorporates long–run regime effects without relying on precise identification of the current state.

On the numerical side, the thesis implements both Monte Carlo and PIDE–based methods to compute the CVA under the proposed model. The comparison between the two approaches highlights their respective strengths and limitations. Monte Carlo simulation offers flexibility and intuitive implementation, especially in higher–dimensional settings, but at the cost of higher computational effort and variance. The PIDE approach, on the other hand, exploits the Markovian structure of the model and provides faster and more stable results in low–dimensional cases, making it particularly suitable for sensitivity analysis and stress testing. The numerical results confirm the internal consistency of the framework and illustrate the quantitative impact of regime switching and mean reversion on CVA levels.

Several directions for future research naturally arise from this work. A first extension concerns the default–environment process Y. While the thesis focuses on a two–state CTMC, the framework can be generalized by allowing the state space of Y to be countable, or more generally by modeling Y as a pure jump process. This generalization would enable a more flexible representation of the credit environment, where transitions are not restricted to a finite set of regimes but may reflect a potentially large (or even infinite) collection of market conditions. From a modeling perspective, a countable-state or pure jump specification allows the intensity dynamics to incorporate rare but economically meaningful events, clustering effects, and heterogeneous stress levels without imposing an a priori coarse regime partition. Such an extension would shift the mathematical structure from a finite system of coupled PIDEs to an infinite-dimensional system (or to an integro-differential formulation driven by the generator of the jump process), thereby increasing analytical and numerical complexity but yielding a richer and more structurally consistent description of regime-driven credit risk.

A second and particularly relevant development concerns the introduction of dependence between the stochastic factor of the claims arrival intensity process X and the regime process Y. In the current setup, Y affects the model only through the default intensity of the reinsurer, while the dynamics of X are independent of the prevailing regime, except for the jump at default. A natural extension would be to allow the parameters of the OU process—such as the mean reversion level, speed or volatility—to depend explicitly on the state of Y. This would reflect the empirical observation that claim frequencies and market conditions are often jointly driven by macroeconomic and systemic factors, leading to a tighter interaction between underwriting risk and credit risk.

Further developments could also include more general claim size distributions, stochastic recovery rates and short rate, or the calibration of the model to real market data within a Solvency II internal–model framework. Overall, the results of this thesis confirm that CVA–based methodologies, originally developed in the banking sector, can be effectively adapted to reinsurance counterparty risk, providing a coherent and quantitative tool for pricing, risk management and regulatory assessment.

Bibliography

[1] A. Ang, A. Timmermann, and C. for Economic Policy Research (Great Britain). *Regime Changes and Financial Markets.* Discussion paper. Centre for Economic Policy Research, 2011.

[2] T. Bielecki and M. Rutkowski. *Credit Risk: Modeling, Valuation and Hedging.* Springer, 2002.

[3] D. Brigo and F. Mercurio. *Interest Rate Models — Theory and Practice: With Smile, Inflation and Credit.* Springer, 2006.

[4] C. Ceci, K. Colaneri, R. Frey, and V. Köck. Value adjustments and dynamic hedging of reinsurance counterparty risk. *SIAM JOURNAL ON FINANCIAL MATHEMATICS.*, pages 788–814, 2020.

[5] G. Cesari, J. Aquilina, N. Charpillon, Z. Filipovic, G. Lee, and I. Manda. *Modelling, Pricing, and Hedging Counterparty Credit Exposure: A Technical Guide.* Springer, 2010.

[6] C. W. Chang, J. S. Chang, and W. Lu. Pricing catastrophe options with stochastic claim arrival intensity in claim time. *Journal of Banking Finance*, 34(1):24–32, 2010.

[7] M. Crisafulli. A solvency ii partial internal model considering reinsurance and counterparty default risk. *Journal of Risk and Financial Management*, 17(4), 2024.

[8] J. D. Cummins and P. Trainar. Securitization, insurance, and reinsurance. *Journal of Risk and Insurance*, 76(3):463–492, 2009.

[9] R. Elliott, L. Aggoun, and J. Moore. *Hidden Markov Models: Estimation and Control.* Applications of mathematics. Springer, 1995.

[10] M. Gallegati. Financial and business cycles in the us: A non-parametric time–frequency investigation. *The North American Journal of Economics and Finance*, 81:102547, 2026.

[11] M. Gatumel and S. Lemoyne de Forges. Understanding and monitoring reinsurance counterparty risk. *Bulletin Français d'Actuariat*, 13(26):121–138, 2013.

[12] J. D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometria*, 57:357–384, 1989.

[13] J. C. Hull, M. Predescu, and A. White. Bond prices, default probabilities and risk premiums. *Journal of Credit Risk*, 1(1):53–60, 2005.

[14] S. Klugman, H. Panjer, and G. Willmot. *Loss Models: From Data to Decisions*. Wiley Series in Probability and Statistics. Wiley, 2019.

[15] D. Lando. On cox processes and credit risky securities. *Review of Derivatives Research*, 2:99–120, 1998.

[16] C. Levi and C. Partrat. Statistical analysis of natural events in the united states. *ASTIN Bulletin*, 21(2):253–276, 1991.

[17] A. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools - Revised Edition*. Princeton Series in Finance. Princeton University Press, 2015.

[18] J. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.

[19] S. C. Park and X. Xie. Reinsurance and systemic risk: The impact of reinsurer downgrading on property–casualty insurers. *Journal of Risk Insurance*, 81(3):587–621, 2014.

[20] S. C. Park, X. Xie, and P. Rui. The sensitivity of reinsurance demand to counterparty risk: Evidence from the u.s. property–liability insurance industry. *Journal of Risk Insurance*, 86(4):915–946, 2019.

[21] SwissRe. The essential guide to reinsurance, 2015.

[22] C. Zhu, G. Yin, and N. Baran. Feynman-kac formulas for regime-switching jump diffusions and their applications. *Stochastics An International Journal of Probability and Stochastic Processes*, 87, 2014.

A | Continuous-Time Markov Chains Theory

This appendix outlines the fundamental mathematical definitions and theorems regarding Continuous-Time Markov Chains (CTMC), based on the treatment by J.R. Norris in *Markov Chains* (Cambridge University Press, 1997).

Before defining the process, we establish the properties of the state space and the generator matrix.

Definition A.1 (Distribution, p. 1, Norris, 1997). *Let I be a countable set. Each $i \in I$ is called a state and I is called the state-space. We say that $\lambda = (\lambda_i : i \in I)$ is a measure on I if $0 \leq \lambda_i < \infty$ for all $i \in I$. If in addition the total mass $\sum_{i \in I} \lambda_i$ equals 1, then we call λ a distribution.*

The evolution of a CTMC is governed by a transition rate matrix, known as a Q-matrix.

Definition A.2 (Q-matrix, pp. 60-61, Norris, 1997). *A Q-matrix on I is a matrix $Q = (q_{ij} : i, j \in I)$ satisfying the following conditions:*

(i) $0 \leq -q_{ii} < \infty$ *for all* i;

(ii) $q_{ij} \geq 0$ *for all* $i \neq j$;

(iii) $\sum_{j \in I} q_{ij} = 0$ *for all* i.

Note: We often denote $q_i = -q_{ii}$ as the rate of leaving state i.

To construct the process formally, we derive the Jump Matrix from the Q-matrix.

Definition A.3 (Jump Matrix, p. 87, Norris, 1997)**.** *The jump matrix* $\Pi = (\pi_{ij} : i, j \in I)$ *of* Q *is defined by*

$$\pi_{ij} = \begin{cases} q_{ij}/q_i & \textit{if } j \neq i \textit{ and } q_i \neq 0 \\ 0 & \textit{if } j \neq i \textit{ and } q_i = 0 \end{cases}$$

$$\pi_{ii} = \begin{cases} 0 & \textit{if } q_i \neq 0 \\ 1 & \textit{if } q_i = 0 \end{cases}$$

The CTMC is defined constructively via its embedded discrete chain and holding times.

Definition A.4 (Markov Chain, p. 88, Norris, 1997)**.** *A minimal right-continuous process* $(X_t)_{t\geq 0}$ *on* I *is a Markov chain with initial distribution* λ *and generator matrix* Q *if its jump chain* $(Y_n)_{n\geq 0}$ *is discrete-time Markov(*λ, Π*) and if for each* $n \geq 1$*, conditional on* $Y_0, \dots, Y_{n-1}$*, the holding times* $S_1, \dots, S_n$ *are independent exponential random variables of parameters* $q(Y_0), \dots, q(Y_{n-1})$ *respectively. We say* $(X_t)_{t\geq 0}$ *is Markov(*λ, Q*) for short.*

The transition probabilities link the Q-matrix to the time-dependent behavior of the chain via the Backward Equation.

Definition A.5 (Transition Probabilities, p. 109, Norris, 1997)**.** *Associated to any Q-matrix is a semigroup* $(P(t) : t \geq 0)$ *of sub-stochastic matrices* $P(t) = (p_{ij}(t) : i, j \in I)$*. As the name implies we have*

$$P(s)P(t) = P(s+t) \quad \textit{for all } s, t \geq 0.$$

The entries $p_{ij}(t)$ *are the transition probabilities.*

The relationship between $P(t)$ and Q is formally defined as follows:

Theorem A.1 (Theorem 2.8.3, p. 96, Norris, 1997)**.** *Let* Q *be a Q-matrix. Then the backward equation*

$$P'(t) = QP(t), \quad P(0) = I$$

has a minimal non-negative solution $(P(t) : t \geq 0)$*. [...] We call* $(P(t) : t \geq 0)$ *the minimal non-negative semigroup associated to* Q*.*

We analyze the long-term behavior of the chain through invariant measures and convergence theorems.

Definition A.6 (Invariant Measure, p. 117, Norris, 1997)**.** *We say that a measure* λ *is invariant if* $\lambda Q = 0$*.*

The existence of a stationary distribution is linked to the positive recurrence of the states.

Theorem A.2 (Theorem 3.5.3, p. 118, Norris, 1997). *Let Q be an irreducible Q-matrix. Then the following are equivalent:*

(i) every state is positive recurrent;

(ii) some state i is positive recurrent;

(iii) Q is non-explosive and has an invariant distribution λ.

Moreover, when (iii) holds we have $m_i = 1/(\lambda_i q_i)$ for all i [where m_i is the expected return time].

Theorem A.3 (Theorem 3.6.2, p. 122, Norris, 1997). *Let Q be an irreducible non-explosive Q-matrix with semigroup $P(t)$, and having an invariant distribution λ. Then for all states i, j we have*

$$p_{ij}(t) \to \lambda_j \quad \text{as } t \to \infty.$$

Crucially, if a chain begins with the stationary distribution, it maintains that distribution for all future times.

Theorem A.4 (Theorem 3.5.6, p. 121, Norris, 1997). *Let Q be an irreducible non-explosive Q-matrix having an invariant distribution λ. If $(X_t)_{t\geq 0}$ is Markov(λ, Q) then so is $(X_{s+t})_{t\geq 0}$ for any $s \geq 0$.*

This theorem implies that if the process starts in equilibrium, the probability of being in state i remains $\mathbb{P}(X_t = i) = \lambda_i$ for all $t \geq 0$.

B | Gaussian Ornstein-Uhlenbeck Processes

This appendix presents the mathematical foundation of the Gaussian Ornstein-Uhlenbeck (OU) process, a continuous-time stochastic process characterized by its mean-reverting property. The following definitions and theorems are based on Maller et al. (2009).

The Gaussian OU process is defined as the solution to a linear stochastic differential equation driven by Brownian motion.

Definition B.1 (Gaussian Ornstein-Uhlenbeck Processe, p. 2, Maller et al., 2009). *The (one-dimensional) Gaussian OU process $X = (X_t)_{t\geq 0}$ can be defined as the solution to the stochastic differential equation (SDE):*

$$dX_t = \gamma(m - X_t)dt + \sigma dB_t, \quad t > 0, \tag{B.1}$$

where γ, m, and $\sigma \geq 0$ are real constants, and B_t is a standard Brownian Motion (SBM) on $\mathbb{R}$.

The OU process admits a closed-form solution obtained through the method of integrating factors.

Definition B.2 (Stochastic Integral, p. 2, Maller et al., 2009). *Alternatively, we could define X in terms of a stochastic integral:*

$$X_t = m(1 - e^{-\gamma t}) + \sigma e^{-\gamma t} \int_0^t e^{\gamma s} dB_s + X_0 e^{-\gamma t}, \quad t \geq 0.$$

It is easily verified that X satisfies B.1 for any γ, m, σ, and choice of X_0; it is the unique, strong Markov solution to B.1.

To conclude, the principal moments of the process are given by:

Definition B.3 (Expectation and Covariance Functions, p. 2, Maller et al., 2009)**.** *Conditional on X_0, and assuming X_0 has finite variance, X_t is Gaussian with expectation and covariance functions given by*

$$\mathbb{E}X_t = m(1 - e^{-\gamma t}) + e^{-\gamma t}\mathbb{E}X_0, \quad t \geq 0,$$

and

$$Cov(X_u, X_t) = \frac{\sigma^2}{2\gamma} e^{-\gamma u}(e^{\gamma t} - e^{-\gamma t}) + e^{-\gamma(u+t)}\, Var X_0, \quad u \geq t \geq 0.$$

C | Algorithms of the model

Algorithm C.1 Simulation of the Markov Chain Y and Default Time

Set $t \leftarrow 0$

Initialize the state:

$$Y_0 = \begin{cases} 0, & \text{with probability } \pi_0, \\ 1, & \text{with probability } 1 - \pi_0 \end{cases}$$

while $t < T$ **do**

  Set switching intensity:

$$q \leftarrow \begin{cases} q_{01}, & \text{if } Y_t = 0, \\ q_{10}, & \text{if } Y_t = 1 \end{cases}$$

  Sample switching time $\tau_Y \sim \mathrm{Exp}(q)$

  Set default intensity:

$$\lambda \leftarrow \begin{cases} \lambda_0, & \text{if } Y_t = 0, \\ \lambda_1, & \text{if } Y_t = 1 \end{cases}$$

  Sample default time $\tau_D \sim \mathrm{Exp}(\lambda)$

  if $\tau_D \leq \tau_Y$ **and** $t + \tau_D \leq T$ **then**

    return default time $t + \tau_D$ and state Y_t

  end if

  if $t + \tau_Y \geq T$ **then**

    return $+\infty$ (no default before maturity)

  end if

  Update time: $t \leftarrow t + \tau_Y$

  Switch state: $Y_t \leftarrow 1 - Y_t$

end while

return $+\infty$ (no default before maturity)

Algorithm C.2 Backward Solution of the Value Function v^ϕ via Crank–Nicolson Scheme

Input: model type X, generator matrix M, grid parameters
if $X =$ constant **then**
  Set operator $L \leftarrow x \cdot M$
  Set effective dimension dim $\leftarrow N_\ell$, $N_x^{\text{eff}} \leftarrow 1$
else if $X =$ OU **then**
  Construct identity I_ℓ and diagonal matrix $\text{diag}(x)$
  Set operator:

$$L \leftarrow \text{diag}(x) \otimes M + D_x \otimes I_\ell$$

  Set effective dimension dim $\leftarrow N_x N_\ell$, $N_x^{\text{eff}} \leftarrow |\mathcal{X}|$
else
  error: invalid model specification
end if
Construct identity matrix I of size dim
Define operator $A_v \leftarrow L - rI$
Construct Crank–Nicolson matrices:

$$L_{\text{CN}} \leftarrow I - \tfrac{1}{2}\Delta t\, A_v,$$
$$R_{\text{CN}} \leftarrow I + \tfrac{1}{2}\Delta t\, A_v$$

Precompute LU factorization of L_{CN}
Terminal condition
if $X =$ constant **then**
  Set $v_T^\phi(\ell) \leftarrow \Phi(\ell)$
  Initialize v_T^ϕ identically in both regimes
else
  Replicate $v_T^\phi(\ell)$ over the x-grid
  Initialize $v_T^\phi(x, \ell)$ identically in both regimes
end if
Store terminal value at time $t = T$

Algorithm C.2 Backward Computation of Value Function ϕ (continued)

Backward iteration
for $n = N_t, \ldots, 1$ **do**
    for regime $i = 0, 1$ **do**
        Solve linear system:

$$v_i^{\phi^{n-1}} = L_{\mathrm{CN}}^{-1} R_{\mathrm{CN}} v_i^{\phi^n}$$

    end for
    Store solution at time step $n-1$
end for
return value function $v^\phi(t, x, \ell)$ on the full grid

Algorithm C.3 Backward Computation of CVA via Coupled Crank–Nicolson Scheme

Input: model type X, contract value $v^\phi(t)$, generator M, jump size γ
if $X =$ constant **then**
    Set operator $L \leftarrow x \cdot M$
    Set effective dimension $\dim \leftarrow N_\ell$, $N_x^{\mathrm{eff}} \leftarrow 1$
else if $X =$ OU **then**
    Construct diagonal matrix $\mathrm{diag}(x)$ and identity I_ℓ
    Set operator:

$$L \leftarrow \mathrm{diag}(x) \otimes M + D_x \otimes I_\ell$$

    Set effective dimension $\dim \leftarrow N_x N_\ell$, $N_x^{\mathrm{eff}} \leftarrow |\mathcal{X}|$
else
    error: invalid model specification
end if
Construct identity matrix I of size dim
Define block operators:

$$\begin{aligned} A_{11} &= L - q_{01} I - (r + \lambda_0) I, \\ A_{12} &= q_{01} I, \\ A_{21} &= q_{10} I, \\ A_{22} &= L - q_{10} I - (r + \lambda_1) I \end{aligned}$$

Assemble block generator:

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

Construct Crank–Nicolson matrices:

$$\begin{aligned} L_{\mathrm{CN}} &\leftarrow I_{2\dim} - \tfrac{1}{2} \Delta t\, A, \\ R_{\mathrm{CN}} &\leftarrow I_{2\dim} + \tfrac{1}{2} \Delta t\, A \end{aligned}$$

Algorithm C.3 Backward Computation of CVA (continued)

Precompute LU factorization of L_{CN}
Initialize CVA value $f_T = 0$ in both regimes
Allocate storage for $f(t)$
Jump operator for OU model
if $X = \text{OU}$ **and** $\gamma \neq 0$ **then**
  Construct jump matrix J_x corresponding to $x \mapsto x(1+\gamma)$
  Extend jump operator: $J_x^{\text{big}} \leftarrow J_x \otimes I_\ell$
end if
Backward iteration
for $n = N_t, \ldots, 1$ **do**
  Extract contract values v^{ϕ^n} and $v^{\phi^{n-1}}$ in both regimes
  if $X = \text{OU}$ **and** jump operator is defined **then**
    Apply jump: $v^\phi \leftarrow J_x^{\text{big}} v^\phi$
  end if
  Construct source terms:

$$s^n = (\lambda_0 v_0^{\phi^n}, \lambda_1 v_1^{\phi^n}),$$
$$s^{n-1} = (\lambda_0 v_0^{\phi^{n-1}}, \lambda_1 v_1^{\phi^{n-1}})$$

  Compute right-hand side:

$$\text{rhs} = R_{\text{CN}} f^n + \tfrac{1}{2}\Delta t (s^n + s^{n-1})$$

  Solve linear system:

$$f^{n-1} = L_{\text{CN}}^{-1} \text{rhs}$$

  Store CVA values at time t_{n-1}
end for
return CVA surface $f(t, x, \ell)$

List of Figures

List of Tables