

IMPERIAL

IMPERIAL COLLEGE LONDON

DEPARTMENT OF MATHEMATICS

On the Hedgeable–Unhedgeable Boundary: A Comparative Study of Capital Valuation Adjustment in Derivative Pricing

Author: Chengdong Song (CID: 02030540)

A thesis submitted for the degree of
MSc in Mathematics and Finance, 2024-2025

Declaration

The work contained in this thesis is my own work unless otherwise stated.

Acknowledgements

As I finish this work, I'm amazed at how quickly time has passed. Four years in the Department of Mathematics at Imperial College London have gone by so fast. What began in autumn 2021 now ends in summer 2025. These years have gently closed one chapter of my life while softly opening the next. Though it weighs at my heart to leave behind the familiar skin of a student, what remains is a deep and humble gratitude.

Imperial's Mathematics Department has shaped me in ways I could never have imagined. At times it felt almost like a Stockholm syndrome with mathematics—intense, demanding, occasionally overwhelming, yet impossible to let go, and ultimately becoming a defining part of who I am. For that, I am grateful beyond words.

I've been incredibly fortunate to have Professor Damiano Brigo as my supervisor. A leading legend in quantitative finance, it was his course "SDEs in Financial Modelling" in my third year that first kindled my fascination with being a quant—setting alight a path I am now determined to walk. Since then, he has been my lecturer, personal tutor, and thesis supervisor, offering not only knowledge but also care and encouragement. I will always be thankful for his presence along this journey.

As a student in the MSc in Mathematics and Finance in the 2024-2025 cohort, I want to thank all the faculty who made this program so exceptional, especially Course Directors Professor Jack Jacquier and Dr. Eyal Neuman, whose dedication defined much of our journey. Among them, Dr. Yufei Zhang deserves particular gratitude. He offered not only knowledge but also genuine mentorship. Our conversations often wandered beyond the syllabus, into futures and intellectual curiosities, for which I am deeply thankful.

To my friends, thank you for being part of my journey. You've shared both the struggles and the celebrations with me. The late-night conversations, the laughter, the glasses of alcohol that sometimes turned into philosophy, and the friendship have meant everything. Life often feels like being gently pushed along by its current, and how lucky I was to drift alongside each of you. Now as we each find our own currents, may your paths be kind, your skies clear, and your hearts always light.

Most importantly, I thank my family, especially my parents. For twenty-two years, they have given me their love, patience, and unwavering support. Every achievement I've reached has been possible because of their foundation. Their dedication is woven into every achievement I call mine, and their quiet strength lives in me. No words can capture what they mean to me. I can only wish for them years filled with health and peace.

And finally, a word to myself. This thesis, written during nights and weekends while working separately as an Off-Cycle Intern in Market Quantitative Research at J.P. Morgan, is more than an academic endeavor. Though the twin demands of work and study often seemed overwhelming, it is precisely their tension that has made the completion of this thesis feel so significant. It's a testament to twenty-two years of trying, growing, and becoming. Life rarely shows us what the next five years will bring. But that is what makes the journey meaningful—constantly building while holding on to the nostalgic weight of the past, walking into the unknown with both hope and doubt. Time passes quickly, and all I can do is face each challenge with courage, and cherish every moment with gratitude.

Keep walking, and let each step teach me who I am.

Abstract

The Capital Valuation Adjustment (KVA) represents the cost of regulatory capital in derivative pricing and raises the question of whether capital should be modeled as hedgeable or unhedgeable. This thesis compares two alternative approaches. In the semi-replication framework, KVA is incorporated within an extended risk-neutral setup, treated in parallel with other valuation adjustments. By contrast, indifference pricing views KVA as an unhedgeable cost, embedding it through the requirement that the bank remain economically indifferent between entering or not entering the trade so that the chosen performance or utility measure is preserved.

A consistent simulation environment is developed to implement and compare both methodologies on representative derivatives and portfolio settings. Special attention is given to the impact of netting effects, which can materially affect capital requirements and hence the size of KVA, as well as to wrong-way risk (WWR), where exposure and counterparty credit quality are correlated, affecting both capital charges and valuation adjustments.

The analysis shows how the hedgeability assumption influences valuation outcomes and portfolio effects, with semi-replication producing results more aligned with traditional XVA practice, while indifference pricing ties valuation directly to capital requirements and shareholder objectives. The study highlights how the treatment of KVA, alongside considerations such as netting and WWR, affects both deal pricing and portfolio management, and suggests that banks may combine the two perspectives in practice.

Contents

1	Introduction	6
2	Literature Review	9
2.1	Basel III Capital Framework	9
2.2	Key Literature on KVA	10
3	Theoretical Framework	13
3.1	Semi-Replication Method	13
3.2	Indifference Pricing Method	19
3.2.1	RAROC-based Indifference Method	20
3.2.2	Optimization-based Indifference Method	21
3.2.3	Alternative Perspective	23
3.2.4	KVA for Indifference Method	24
3.3	Exposure Profiles and Capital Requirements	26
3.4	Netting, Wrong-Way Risk and Double Counting	27
3.4.1	Netting of Exposures	27
3.4.2	Wrong-Way Risk Mathematical Framework	27
3.4.3	Double Counting	28
3.5	Example Instruments and Models	29
3.5.1	European Call Option under Black-Scholes Model	29
3.5.2	Interest Rate Swap under the Vasicek Model	30
4	Numerical Implementation Strategies	33
4.1	Methodological Alignment and Unified Parameter	33
4.2	Hedgeability Cases for XVA	35
4.3	Detailed Implementation Strategies	37
4.3.1	Simulation Framework for Exposure Calculation	37
4.3.2	Parameter Calibration under Indifference Pricing	39
4.3.3	Netting Implementation Strategy	39
4.3.4	Wrong Way Risk Implementation Strategy	41
5	Numerical Experiments Results	44
5.1	XVA Baseline Calculation	44
5.1.1	European Call Option Case	46
5.1.2	Interest Rate Swap Case	47
5.2	Netting Result Analysis	48
5.3	Wrong Way Risk Analysis	50
6	Integrating KVA into a Comprehensive XVA Framework	53
6.1	Unified XVA Architecture	53
6.2	KVA as the Cost of Retained Capital	56
6.3	KVA in the Trading Desk	56

7 Conclusion	58
Bibliography	61

List of Figures

3.1	Interest Rate Swap Illustration	31
5.1	Expected Exposure (EE) profiles for Call Option and Interest Rate Swap. .	44
5.2	Simulated Vasicek Short Rate Paths Illustration	47
5.3	Bank and Counterparty Default Intensity Sample Paths	50
5.4	Joint Correlation Between Stock Prices and Default Intensities	51
6.1	XVA Framework Illustration	54

List of Tables

3.1	Comprehensive List of Parameters Used Across XVA Models	14
5.1	Model Parameters for Call Option and IRS Experiments	45
5.2	Call Option BS price, capital requirement, and XVA adjustments	46
5.3	Interest Rate Swap BS price, capital requirement, and XVA adjustments . .	48
5.4	Netting Results	49
5.5	Semi-Replication XVA components: pre/post netting and reductions	49
5.6	Wrong-way and right-way risk analysis results	51

Chapter 1

Introduction

The 2008 global financial crisis highlighted serious shortcomings in the way banks were regulated and capitalized, particularly in relation to the risks associated with over-the-counter (OTC) derivatives. In response, regulators introduced a series of reforms—most notably Basel III and subsequent Basel IV—aimed at improving financial system stability through stricter capital, liquidity, and counterparty risk requirements.

These changes have had a lasting impact on the pricing of derivatives. Banks now apply a range of valuation adjustments (XVAs) to reflect credit risk, funding costs, and, more recently, capital costs. Among these, Capital Valuation Adjustment (KVA) has gained growing attention as it represents the opportunity cost of regulatory capital that must be held against derivative exposures. KVA is essentially an upfront charge to ensure new trades meet internal return targets, compensating shareholders for the capital tied up by the trade.

KVA differs from other valuation adjustments in both nature and purpose. Credit Valuation Adjustment (CVA) accounts for expected losses due to counterparty default, reducing the derivative's value to reflect credit risk. Debit Valuation Adjustment (DVA), in contrast, reflects the benefit a bank might receive if it were to default on its own obligations¹. Funding Valuation Adjustment (FVA) captures the cost of funding uncollateralized exposures, i.e., the incremental interest expense or loss of yield from posting collateral or funding the position over time. KVA, by contrast, measures the economic cost of capital usage itself, typically benchmarked against the internal return hurdle rate. Unlike CVA or FVA, which relate to specific cash flows or spreads that can often be hedged (e.g., CVA via credit default swaps, FVA via funding instruments), KVA introduces a long-term profitability constraint not directly linked to market-tradable assets.

This distinction makes KVA's hedgeability highly controversial. In principle, a bank can hedge CVA and FVA, whereas capital costs cannot be fully hedged or traded away—any hedge to reduce capital usage tends to itself require capital and create additional KVA. As a result, whether one treats KVA as a hedgeable adjustment or an unhedgeable cost leads to fundamentally different pricing approaches.

No single agreed-upon methodology for KVA has been universally adopted. In practice, institutions largely agree on the need for KVA, but not on its precise definition or implementation [35]. Broadly, two schools of thought have emerged. The semi-replication Approach extends the arbitrage pricing by including capital as an added cost in the replicating portfolio. This method, pioneered by Burgard and Kjaer [22] for CVA and FVA and later applied to KVA by Green, Kenyon and others [26], assumes that the incremental cost of capital can be priced as deterministic cash flows within a risk-neutral valuation framework. On the other hand, indifference pricing approach treats KVA as an unhedgeable

¹Though recognizing DVA as a gain is controversial, since it is only realized in the event of the bank's failure.

risk that must be priced via the bank’s risk appetite. Introduced by Brigo, Francischello and Pallavicini [13], this approach does not assume a perfect replicating hedge for capital costs. Instead, the trade’s price is set such that the bank is economically indifferent to entering the trade. For example, the desk’s risk-adjusted return on capital (RAROC) remains at target levels, or the bank’s maximized expected payoff is unchanged.

This dichotomy has practical significance. If KVA (and other XVAs) were fully hedgeable, one could in theory extend the no-arbitrage pricing machinery to include them all, using the semi-replication approach. However, if some components are assumed to be not hedgeable, then traditional replication breaks down and one must resort to the indifference approach. As we will explore, assuming XVAs are hedgeable or unhedgeable yields different results.

This thesis conducts a comparative study of KVA pricing under different hedgeability assumptions. We develop two main methodologies—the semi-replication and indifference pricing frameworks, and examine their theoretical foundations and quantitative outcomes. In particular, we consider three scenarios:

1. All XVA components unhedgeable: Only the indifference approach applies, as no replication is possible.
2. All XVA components hedgeable: Only the semi-replication approach is used.
3. CVA/DVA/FVA hedgeable, but KVA unhedgeable: A hybrid approach where we hedge credit and funding risks, but use an indifference charge for the residual capital cost.

A unified simulation and parameter alignment strategy is designed to ensure comparability across methods, harmonizing key inputs such as hurdle rates and cost-of-capital specifications. Practical portfolio-level considerations—including counterparty-level netting, wrong-way risk (WWR), and the avoidance of double counting across overlapping adjustments—are integrated into the framework to preserve consistency of the full XVA stack.

The empirical analysis is performed on two representative instruments—a European call option under the Black-Scholes model and an interest rate swap under the Vasicek model. Beyond baseline results, we assess the impact of netting and WWR on KVA, highlighting how portfolio structure and exposure-default correlation reshape valuation outcomes across hedgeability regimes.

Finally, we embed KVA within a comprehensive XVA architecture, showing how capital costs can be integrated alongside CVA, DVA, and FVA into a coherent framework for trading desks. This perspective clarifies how different methodological choices—semi-replication versus indifference—translate into pricing, transfer-pricing, and governance implications.

The remainder of this dissertation is structured as follows: Chapter 2 provides a review of the relevant regulatory background and literature on capital costs and KVA. Chapter 3 develops the theoretical framework. It first presents the semi-replication approach, then a family of indifference pricing formulations (RAROC-based, optimization-based, shareholder-perspective, and median-based). The chapter also introduces two representative instruments—a European call option under the Black-Scholes model and an interest-rate swap under the Vasicek model—which serve as testbeds for later numerical analysis. Practical issues such as netting, wrong-way risk, and avoiding double counting are formalized within the framework. Chapter 4 describes the numerical implementation strategy. Beyond parameter alignment across methods, it sets out the simulation scheme for exposure profiles, provides pseudo-code for the indifference optimization problems, and details the implementation of netting sets and WWR correlations. Chapter 5 reports the results of the numerical experiments. We analyze baseline XVA adjustments for

the call and swap, quantify the effect of netting, and study the sensitivity to wrong-way risk, comparing outcomes across hedgeability regimes. Chapter 6 discusses how KVA can be integrated into a complete XVA framework alongside CVA, DVA and FVA, and discusses implications for transfer pricing and trading-desk governance. Chapter 7 concludes that the two approaches reflect complementary perspectives—market-consistent replication and shareholder-oriented indifference—and that practice often requires combining them. It also points to future research on unified frameworks that integrate these views into a consistent approach to KVA.

Chapter 2

Literature Review

2.1 Basel III Capital Framework

The introduction of Basel III by the Basel Committee on Banking Supervision marked a fundamental shift in the regulatory landscape for financial institutions. Motivated by the lessons of the 2008 financial crisis, Basel III introduced a set of reforms that significantly tightened capital adequacy standards. These reforms introduced more stringent definitions of eligible capital and higher minimum capital ratios. Additionally, Basel III mandated new buffers, including the leverage ratio and capital conservation buffer [6]. The overarching goal was to ensure banks hold sufficient high-quality capital to absorb unexpected losses and reduce systemic risk.

Crucially for derivatives, Basel III introduced a specific CVA capital charge: banks must hold regulatory capital against potential mark-to-market losses from counterparty defaults. In addition, the framework revised how counterparty exposure is measured via new methodologies like the Standardized Approach for Counterparty Credit Risk (SA-CCR), which provides more risk-sensitive measures of potential future exposure compared to the previous Current Exposure Method, replacing simpler measures. Basel III also included the leverage ratio which, being non-risk-sensitive, does not fully recognize the benefits of netting or collateral for derivatives exposures. The net effect was that derivatives, especially uncleared bilateral contracts, became much more capital-intensive. Trades that were economically immaterial or hedgeable under Basel II [5] could still consume significant capital under Basel III. In effect, capital costs became a material consideration in trading OTC derivatives, not just an afterthought.

This had direct implications for valuation adjustments. If a new trade increases the bank's risk-weighted assets or CVA capital, it imposes a capital opportunity cost until that exposure runs off. Banks began asking how to charge this cost to the counterparty, to avoid diluting their return on equity or similar metric. This gave rise to KVA: an ex-ante fee intended to offset the drag on ROE caused by the trade's capital usage. The stronger the capital requirements (e.g., higher CVA risk weights, additional buffers), the larger the KVA for a given trade. For example, global reforms that encourage central clearing of derivatives effectively make uncleared trades more expensive, in part due to higher capital charges—a phenomenon documented by regulators. In a 2018 study, the Financial Stability Board [25] noted that the cost difference between uncleared and cleared trades is explained largely by capital (KVA) charges, with uncleared trades bearing a significantly higher KVA while cleared trades have negligible KVA. This creates an incentive to clear standardized trades, aligning with policy goals.

In summary, Basel III's capital framework has made capital costs explicit and often dominant in XVA. It set the stage for KVA methodologies by quantifying how much capital each trade ties up and for how long. The industry's response in pricing models and internal

transfer pricing has been to incorporate these capital costs into deal pricing—however, as the next sections show, there is debate on exactly how to do so, leading to the different approaches for KVA.

2.2 Key Literature on KVA

Academic and practitioner interest in incorporating capital costs into derivative pricing began in the early 2010s, in response to these regulatory changes. This development occurred in parallel with debates on FVA and the broader suite of XVAs. Several key works have shaped the modern treatment of KVA; we outline a few main contributions below.

Green et al. [26] extended the semi-replication framework originally developed by Burgard-Kjaer [21, 22] for CVA and FVA to account for capital costs. In their setup, the dealer sets up a replicating portfolio for a derivative that now includes equity capital as one of the components. Building on a Black-Scholes-type PDE, they define KVA as the discounted expectation of future capital requirements incurred over the lifetime of the trade, calculated under the risk-neutral measure. This formulation allows KVA to be integrated into the pricing PDE similarly to CVA and FVA. However, unlike those adjustments, one cannot dynamically hedge away the capital term by trading, because it depends on regulatory conditions and losses in stressed scenarios. They also acknowledge that capital cannot be dynamically reduced as the trade evolves due to regulatory constraints - one must hold capital until maturity or run-off. They also discuss that any hedging of KVA is problematic: hedging transactions would themselves require capital, thus one cannot eliminate KVA without incurring more. Despite this limitation, the semi-replication literature provides a tractable, formula-driven approach to computing KVA consistent with risk-neutral pricing.

This replication method motivates computational work. Building upon the semi-replication framework, Trevisani et al. [35] address the computational challenges inherent in KVA calculations by developing specialized numerical methods for the resulting high-dimensional PDEs. Their work acknowledges that while Green et al.’s approach provides theoretical consistency, practical implementation faces significant numerical difficulties due to the non-hedgeable characteristics of KVA. The authors systematically review existing methodologies, identifying a critical gap between theoretical sophistication and computational tractability. They emphasize that standard numerical techniques often suffer from stability and convergence issues when applied to KVA problems, necessitating the development of robust computational methods specifically tailored to handle the mathematical complexity while maintaining efficiency for practical applications.

In contrast to replication, Brigo et al. [13] introduce an economically motivated alternative through an indifference pricing approach. At its core, the indifference approach requires that a trader be economically indifferent between entering and not entering a new trade—that is, the relevant performance measure, such as expected utility, RAROC, or expected payoff, must remain unchanged before and after the trade, after accounting for the capital impact. In this framework, KVA emerges as the compensation necessary to preserve the trader’s economic position when taking on additional trades under capital constraints. The idea is that if a trade imposes additional capital constraints—forcing the trader to tie up more resources or absorb more risk—they will demand an adjustment (KVA) such that their ex-ante performance metric is preserved.

The indifference pricing framework naturally comprises two main implementations. The first is based on Risk-Adjusted Return on Capital (RAROC), where the trading desk aims to maintain its target RAROC level before and after the inclusion of a new trade. In this setting, KVA is defined as the minimal additional charge required to preserve the

desk’s RAROC, typically relative to a predefined hurdle rate. This approach is intuitive and closely aligned with how banks assess internal profitability and performance. The second is a more general formulation where one sets up an optimization problem for the bank’s expected equity value, subject to constraints like regulatory capital and limited liability. This formulation is explored from both the whole-bank and shareholder perspectives. When viewed from the shareholder’s perspective, the problem becomes nonlinear due to limited liability, and the resulting KVA reflects the cost required to preserve shareholder value while accounting for both risk and capital requirements. Importantly, Brigo et al. [13] also note that assigning a utility function to the bank introduces considerable model risk, as the relationship between changes in the utility function and changes in pricing outcomes is not always straightforward or interpretable.

Running in parallel to both approaches, Albanese and Crépey [1, 2] took a more holistic, balance-sheet-centered approach. In a series of papers, they formulated XVA, including KVA, in terms of backward stochastic differential equations (BSDEs). This approach is rooted in an accounting view: they write the bank’s balance sheet and P&L evolution including the effects of defaults, funding, margin, and capital. KVA in their formulation is defined via a fixed-point: it is the solution to a BSDE reflecting the cost of carrying capital and its feedback on the portfolio. The result is a fully nonlinear valuation framework—essentially a control problem where KVA is the “value” of the capital constraint. While theoretically rigorous and consistent with a broader XVA architecture, the high computational burden of solving BSDEs makes these models difficult to apply to large-scale portfolios in practice.

Morini and Prampolini [34] interpret KVA from a regulatory perspective, asking how regulatory requirements, like leverage ratio and RWA calculations, effectively transform some of the bank’s future profits into a valuation adjustment today. They treat KVA as arising from “frictions”—constraints that force the bank to allocate profit to capital buffers instead of paying it out. In their analysis, KVA is not derived from replication, but from how regulatory capital rules (like the leverage ratio not recognizing collateral fully, or RWA formulas) cause a drag on performance. They demonstrate how KVA interacts with other XVAs in various scenarios and emphasize that KVA (in their view) is fundamentally about transforming part of the trade’s future profit into an upfront cost due to regulatory conservatism.

Linking the balance sheet to who benefits, Mats Kjaer in a pair of papers [29, 30] further examines the foundations of KVA from a dealer’s balance sheet, deriving the capital cost for a new trade from first principles. Kjaer treats KVA not as a hedgeable fee but as a consequence of regulatory frictions: part of the derivative’s profit is locked away and cannot be distributed to shareholders, creating a wealth transfer. Matthias Arnsdorf [3, 4] makes this explicit. In *KVA as a Transfer of Wealth* [3], the additional capital requirement reduces default risk and benefits creditors at shareholders’ expense; KVA is the charge required to offset this transfer. In *KVA is Incomplete* [4], the market for capital risk is incomplete; KVA is a risk premium for bearing unhedgeable capital consumption risk. These perspectives bridge theory and practice: KVA compensates shareholders for a constraint that cannot be hedged away.

Notably, Hull and White [27] questioned whether capital costs and funding costs should be included in derivative pricing. They argued that these are internal costs: a well-functioning market would price derivatives at a rate that does not include a specific bank’s cost of capital or funding. This argument aligns with accounting standards and some investors’ views, but banks have generally moved toward passing as many of these costs to clients as competition allows. In contrast, all previous literature has argued that capital usage should be priced into derivatives and passed on to clients, albeit through differing frameworks and justifications. Indeed, the industry in recent years started re-

porting FVA in its financial statements, and KVA is used internally to evaluate trades. A 2018 survey by IACPM [28] found that while CVA/FVA frameworks were well-established, implementation of KVA and MVA (margin VA) varied widely, reflecting different views and operational challenges. Nonetheless, an overarching trend is that banks have begun to demand an additional premium from clients for capital-intensive trades, even if the exact terminology of “KVA” is not always used externally.

In parallel to the development of KVA methodologies, significant research has addressed the impact of netting agreements and wrong-way risk (WWR) on counterparty credit risk metrics. This legal view of nettings follows from ISDA agreements and is further confirmed by legal opinions covering major jurisdictions for the United States, United Kingdom, France, and Germany, as documented in [23, 33, 31, 32]. Early foundational work by [16] established general formulas for CVA under netting arrangements, demonstrating how aggregating exposures can reduce overall credit risk. Building on this foundation, [19] developed models for wrong-way risk in interest rate swap exposures, showing that positive correlation between exposure and counterparty default probability can substantially increase CVA.

Subsequent studies have extended these fundamental insights across multiple dimensions. [12] examined the impact of spread volatility and default correlation on counterparty risk for credit default swaps, while [9] focused on accurate counterparty risk valuation for energy-commodity swaps. [18] incorporated collateralization effects and studied credit calibration with structural models under counterparty risk. [11] presented an important unified framework for arbitrage-free bilateral counterparty risk valuation that incorporates collateralization effects.

More recent contributions have focused on advanced methodological developments. [20] proposed efficient numerical approaches for pricing credit valuation adjustments via change of measures to address WWR. Contemporary work by [14] examined nonlinear valuation under credit, funding, and margins, while [10] developed nonlinear valuation approaches with XVAs using two converging methods. Most recently, [15] provided mild to classical solutions for XVA equations under stochastic volatility.

These works collectively demonstrate that netting agreements can mitigate exposure (and hence CVA/KVA), whereas WWR can amplify credit exposure and capital requirements. Our analysis leverages these insights by incorporating both netting and WWR effects in the simulation framework, ensuring that the KVA analysis reflects realistic portfolio risk dynamics.

In summary, capital is expensive, and ignoring it can lead to suboptimal or even loss-making trades. The next chapter builds the theoretical foundations of the two main approaches we will compare: the semi-replication approach and the indifference pricing approach, before we move on to implementation and numerical comparisons.

Chapter 3

Theoretical Framework

In this chapter, we develop the theoretical foundations of the main KVA methodologies. We present the derivations for the semi-replication approach and for two forms of the indifference pricing approach, one based on maintaining RAROC and one based on an optimization criterion. For each, we highlight assumptions and simplifications used. We then discuss how these methods differ and how they might be combined for hybrid scenarios. Finally, we address important practical considerations—considering netting agreements, wrong-way risk, and avoiding double counting of overlapping risk factors in KVA and other XVAs. We also introduce two representative instruments—a European call option under the Black-Scholes model and an interest rate swap under the Vasicek model—which serve as testbeds for the subsequent numerical analysis.

To begin, let us define the notation and basic setup common to all approaches. We denote by $\hat{V}(t)$ the economic or fair value of the derivative or portfolio of derivatives at time t , including all valuation adjustments. We can decompose as:

$$\hat{V}(t) = V(t) + U(t),$$

where $V(t)$ denotes the risk-free value of the derivative at time t , calculated under the classical no-arbitrage framework, assuming perfect collateralization and no default or funding frictions, and $U(t)$ captures the aggregate valuation adjustments.

At time $t = 0$, from the bank's point of view, CVA is charged to the counterparty and the bank receives an upfront fee to compensate for counterparty credit risk, whereas DVA is effectively paid to the counterparty and the bank concedes a lower price due to its own default risk. Thus, at inception CVA appears as a positive adjustment to the bank's price and DVA as a negative adjustment. However, after inception, in mark-to-market valuations, the bank will carry CVA as a reduction to its asset value (a liability for potential default losses) and DVA as an asset (reflecting the gain the bank would realize if it defaulted). This sign convention is analogous to the contract's value V itself: a positive V means the trade is an asset to the bank (owed by the counterparty), while a negative V means a liability of the bank.

Table 3.1 summarizes key parameters and notation that will be used across the models for reference.

With these in mind, we proceed with each approach.

3.1 Semi-Replication Method

The semi-replication approach attempts to extend the classic no-arbitrage replication paradigm to include capital costs. The starting point is the insight from Burgard and Kjaer that one can write a self-financing portfolio consisting of the derivative itself and various hedging instruments. If this portfolio is properly constructed, its value process

Parameter	Description
X	Collateral
K	Capital requirement
Π	Replicating portfolio
S	Underlying stock
μ_S	Stock drift
σ_S	Stock volatility
P_C	Counterparty bond (zero recovery)
P_i	Issuer bond with recovery R_i , $i = \{1, 2\}$ with $R_1 \neq R_2$
$d\bar{\beta}_S$	Growth in the cash account associated with stock
$d\bar{\beta}_C$	Growth in the cash account associated with counterparty bond
$d\bar{\beta}_X$	Growth in the cash account associated with collateral
$d\bar{\beta}_K$	Growth in the cash account associated with capital
r	Risk-free rate
r_C	Yield on counterparty bond
r_i	Yield on issuer bonds
r_X	Yield on the collateral position
α_C	Holding of counterparty bonds
α_i	Holding of issuer bond i
δ	The stock position
γ_S	Stock dividend yield
q_S	Stock repo rate
q_C	Counterparty bond repo rate
J_C	Default indicator for counterparty
J_B	Default indicator for issuer
g_B	Value of derivative portfolio after issuer default
g_C	Value of derivative portfolio after counterparty default
R_i	Recovery on issuer bond i
R_C	Recovery on counterparty derivative portfolio
λ_C	Effective financing rate of counterparty bond $\lambda_C = r_C - r$
λ_B	Spread of a zero-recovery issuer bond, $(1 - R_i)\lambda_B = r_i - r$
s_F	Funding spread in one-bond case $s_F = r_F - r$
s_X	Spread on collateral
$\gamma_K(t)$	Cost of capital, or hurdle rate
$\Delta\hat{V}_B$	Change in value of derivative on issuer default
$\Delta\hat{V}_C$	Change in value of derivative on counterparty default
ϵ_h	Hedging error on default of issuer, often decomposed as $\epsilon_h = \epsilon_{h0} + \epsilon_{hK}$
P	$P = \alpha_1 P_1 + \alpha_2 P_2$, value of the bond portfolio prior to default
P_D	$P_D = \alpha_1 R_1 P_1 + \alpha_2 R_2 P_2$, value of the bond portfolio after default
ϕ	Fraction of capital available for derivative funding
q	Trade notional size
C_0	Initial investable capital (free cash)
h	hurdle rate
m	Quantile parameter for VaR (e.g., $m = 2.33$ for 99% confidence)
RAROC ₀	Target RAROC of existing portfolio
θ	Trading strategy (portfolio weights)
Y	Payoff of new OTC trade
$P(q)$	Price function for trade with size q
$\chi(q)$	Lagrange multiplier from optimization-based indifference method
μ	Drift of asset returns (real-world measure)
A	Covariance matrix of asset returns
b	Covariance vector between trade payoff and asset returns
σ_Y^2	Variance of Y

Table 3.1: Comprehensive List of Parameters Used Across XVA Models

should evolve without arbitrage at the risk-free rate. Any shortfall from achieving a perfect hedge manifests as a valuation adjustment to the derivative's price.

We consider a setting with a single defaultable counterparty. The clean value V of the derivative, which excludes any costs related to credit or funding, satisfies a standard Black-Scholes partial differential equation.

Building on the approach of Burgard and Kjaer [22], incorporating counterparty credit risk introduces a correction CVA term, while the inclusion of funding costs leads to an additional FVA term.

To capture capital effects, we assume that the bank is required to hold a capital amount $K(t)$ at each time t , associated with this transaction. A parameter $\phi \in [0, 1]$ is introduced, which indicates the proportion of this capital that can be used to finance the trade. When $\phi = 0$, no capital can be used for financing (full opportunity cost), when $\phi = 1$, all capital can be utilized (no opportunity cost beyond funding spread).

Let $(\Omega, \mathcal{G}, \mathbb{P})$ denote a probability space, where $\{\mathcal{F}_t\}_{t \geq 0}$ represents the market filtration generated by a Brownian motion $\{W_t\}$, and $\{\mathcal{J}_t\}_{t \geq 0}$ represents the filtration associated with default events. The full filtration is taken as $\mathcal{G}_t = \mathcal{F}_t \vee \mathcal{J}_t$.

Remark 3.1.1. For simplicity of notation, we omit the explicit time subscript t . It is important to emphasize that the time dependence is still present; we simply refrain from writing the subscript at every occurrence. The measure and filtration assumptions remain unchanged, and the omission of subscripts does not affect measurability or adaptiveness. When we need to stress a specific point in time, we will write X_0 , X_T , etc. explicitly.

We model a derivative agreement between a bank B and a counterparty C . Following a standard sign convention, positive cash values are from the perspective of the bank. The underlying dynamics of the relevant financial instruments are specified as follows:

$$\begin{aligned} dS &= \mu S dt + \sigma S dW, \\ dP_C &= r_C P_C dt - P_C dJ_C, \\ dP_i &= r_i P_i dt - (1 - R_i) P_i dJ_B, \quad i \in \{1, 2\}, \end{aligned}$$

where P_C is the price of a zero-recovery zero-coupon bond issued by the counterparty, while P_1 and P_2 are the bank's own bonds with different recovery rates R_1 and R_2 , respectively. Default events J_B and J_C are modeled by independent jump processes that switch from 0 to 1 upon default.

We assume zero basis between the bank's bonds, implying

$$r_i - r = (1 - R_i) \lambda_B,$$

where r is the risk-free interest rate and λ_B denotes the hazard rate of the bank's zero-recovery bond.

Let the economic value of the derivative at time t be represented by $\hat{V}(t, S, J_B, J_C)$. Upon default by either party, the derivative value transitions to:

$$\begin{aligned} \hat{V}(t, S, 1, 0) &= g_B(M_B, X), \\ \hat{V}(t, S, 0, 1) &= g_C(M_C, X), \end{aligned}$$

with M_B and M_C being close-out values, and X representing collateral posted (assumed static in time). A typical close-out convention gives:

$$\begin{aligned} g_B &= (V - X)^+ + R_B(V - X)^- + X, \\ g_C &= R_C(V - X)^+ + (V - X)^- + X, \end{aligned}$$

where for any real x , write $x^+ = \max(x, 0)$ and $x^- = \max(-x, 0)$.

The funding equation ensuring the trade is funded becomes:

$$\hat{V} - X + \alpha_1 P_1 + \alpha_2 P_2 - \phi K = 0. \quad (3.1.1)$$

Assume cash accounts evolve according to:

$$\begin{aligned} d\beta_S &= \delta(\gamma_S - q_S)S dt, \\ d\beta_C &= -\alpha_C q_C P_C dt, \\ d\beta_X &= -r_X X dt, \\ d\beta_K &= -\gamma_K K dt. \end{aligned}$$

Applying Itô's lemma to the economic value $\hat{V}$, we obtain:

$$d\hat{V} = \frac{\partial \hat{V}}{\partial t} dt + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \hat{V}}{\partial S^2} dt + \frac{\partial \hat{V}}{\partial S} dS + \Delta \hat{V}_B dJ_B + \Delta \hat{V}_C dJ_C.$$

Definition 3.1.2 (Trading Strategy). Assume we hold Θ and δ units of S and B respectively at time t , then we have a portfolio whose time- t value is $\Theta S_t + \delta B_t$. The assumptions of Black-Scholes are that we have a frictionless market, meaning that S and B can be traded in arbitrary amounts with no transaction costs, and short positions are allowed. In particular this means we can invest in, or borrow from, the riskless account at the same rate r of interest. A trading strategy is then a triple $(\Theta_t, \delta_t, x_0)$, where x_0 is the initial endowment and (Θ_t, δ_t) is a pair of adapted processes satisfying

$$\int_0^T \Theta_t^2 dt < \infty \quad \text{a.s.}, \quad \int_0^T |\delta_t| dt < \infty.$$

Definition 3.1.3 (Self-financing). We say a portfolio is *self-financing* if

$$\Theta_t S_t + \delta_t B_t - \Theta_s S_s - \delta_s B_s = \int_s^t \Theta_u dS_u + \int_s^t \delta_u dB_u.$$

The increase (decrease) in portfolio value is entirely due to gains (losses) from trade.

For a self-financing replication portfolio Π , the change in value is given by:

$$\begin{aligned} d\Pi &= \delta dS + \delta(\gamma_S - q_S)S dt + \alpha_1 dP_1 + \alpha_2 dP_2 + \alpha_C dP_C \\ &\quad - \alpha_C q_C P_C dt - r_X X dt - \gamma_K K dt. \end{aligned}$$

Adding the dynamics of the portfolio and the derivative gives:

$$\begin{aligned} d\hat{V} + d\Pi &= \left[\frac{\partial \hat{V}}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \hat{V}}{\partial S^2} + \delta(\gamma_S - q_S)S + \alpha_1 r_1 P_1 + \alpha_2 r_2 P_2 \right. \\ &\quad \left. + \alpha_C r_C P_C - \alpha_C q_C P_C - r_X X - \gamma_K K \right] dt \\ &\quad + \epsilon_h dJ_B + \left[\delta + \frac{\partial \hat{V}}{\partial S} \right] dS + \left[g_C - \hat{V} - \alpha_C P_C \right] dJ_C, \end{aligned}$$

where the hedging error at bank default is

$$\begin{aligned} \epsilon_h &= \Delta \hat{V}_B - \alpha_1(1 - R_1)P_1 - \alpha_2(1 - R_2)P_2 \\ &= \Delta \hat{V}_B - (P - P_D) \\ &= \hat{V}(t, S, J_B = 1, J_C) - \hat{V}(t, S, J_B = 0, J_C) - P + P_D \\ &= g_B - X + P_D - \phi K \\ &= \epsilon_{h_0} + \epsilon_{h_K}, \end{aligned}$$

where in the final line ϵ_h is split into a term which does not depend on capital, ϵ_{h_0} and a term which does depend on capital ϵ_{h_k} . Similarly, the term before dJ_C is calculated by

$$\begin{aligned}\Delta \hat{V}_c - \alpha_C P_C &= \hat{V}(t, S, J_B, J_C = 1) - \hat{V}(t, S, J_B, J_C = 0) - \alpha_C P_C \\ &= g_C - \hat{V} - \alpha_C P_C\end{aligned}$$

Assuming the replication holds perfectly outside default, we set

$$d\hat{V} + d\Pi = 0.$$

By eliminating residual risk sources, we impose:

$$\begin{aligned}\delta &= -\frac{\partial \hat{V}}{\partial S}, \\ \alpha_C P_C &= g_C - \hat{V}.\end{aligned}$$

This yields the following PDE for $\hat{V}$:

$$\begin{aligned}0 &= \frac{\partial \hat{V}}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 \hat{V}}{\partial S^2} - (\gamma_S - q_S)S \frac{\partial \hat{V}}{\partial S} - (r + \lambda_B + \lambda_C)\hat{V} \\ &\quad + g_C \lambda_C + g_B \lambda_B - \epsilon_h \lambda_B - s_X X - \gamma_K K + r\phi K, \\ \hat{V}(T, S) &= H(S),\end{aligned}$$

where we used the fact that $s_x = r - r_x$, $\lambda_c = r_c - q_c$ and

$$\begin{aligned}\alpha_1 r_1 P_1 + \alpha_2 r_2 P_2 &= \alpha_1 r P_1 + \alpha_2 r P_2 + \alpha_1 P_1(1 - R_1)\lambda_B + \alpha_2 P_2(1 - R_2)\lambda_B \\ &= (r + \lambda_B)P - \lambda_B P_D \\ &= (r + \lambda_B)(X - \hat{V} + \phi K) - \lambda_B(\epsilon_h - g_B + X + \phi K) \\ &= rX - (r + \lambda_B)\hat{V} - \lambda_B(\epsilon_h - g_B) + r\phi K.\end{aligned}$$

Decomposing the economic value as $\hat{V} = V + U$, and letting V solve the standard Black-Scholes PDE:

$$\begin{aligned}\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} - (\gamma_S - q_S)S \frac{\partial V}{\partial S} - rV &= 0, \\ V(T, S) &= H(S),\end{aligned}$$

We derive the PDE for the adjustment term U :

$$\begin{aligned}\frac{\partial U}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 U}{\partial S^2} - (\gamma_S - q_S)S \frac{\partial U}{\partial S} - (r + \lambda_B + \lambda_C)U &= \\ \lambda_C(V - g_C) + \lambda_B(V - g_B) + \lambda_B \epsilon_h + s_X X + \gamma_K K - r\phi K & \\ U(T, S) &= 0.\end{aligned}$$

Theorem 3.1.4 (Feynman-Kac Theorem). *Let X_t be the solution to the SDE*

$$dX_t = m(X_t) dt + g(X_t) dW_t,$$

where m, g are Lipschitz continuous functions with linear growth:

$$|m(x) - m(y)| \leq K|x - y|, \quad |m(x)| \leq K(1 + |x|),$$

and similarly for g . Define the infinitesimal generator $\mathcal{A}$ as

$$\mathcal{A}f(x) = m(x)\frac{\partial f}{\partial x} + \frac{1}{2}g^2(x)\frac{\partial^2 f}{\partial x^2}.$$

Suppose $v(t, x) \in C^{1,2}([0, T] \times \mathbb{R})$ solves the PDE

$$\begin{aligned} \frac{\partial v}{\partial t} + \mathcal{A}v(t, x) - r(t, x)v(t, x) &= 0, \quad 0 \leq t < T, \\ v(T, x) &= v_T(x), \end{aligned}$$

where $r(t, x)$ and $v_T(x)$ are given functions.

Then the solution $v(t, x)$ admits the probabilistic representation:

$$v(t, x) = \mathbb{E}_t \left[\exp \left(- \int_t^T r(s, X_s) ds \right) v_T(X_T) \right].$$

Applying the Feynman-Kac representation leads to the decomposition:

$$U = \text{CVA} + \text{DVA} + \text{FCA} + \text{COLVA} + \text{KVA},$$

with

$$\text{CVA} = - \int_t^T \lambda_C(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [V(u) - g_C(V(u), X(u))] du \quad (3.1.2)$$

$$\text{DVA} = - \int_t^T \lambda_B(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [V(u) - g_B(V(u), X(u))] du \quad (3.1.3)$$

$$\text{FCA} = - \int_t^T \lambda_B(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [\epsilon_{h_0}(u)] du \quad (3.1.4)$$

$$\text{COLVA} = - \int_t^T s_X(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [X(u)] du \quad (3.1.5)$$

$$\text{KVA} = - \int_t^T e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [(\gamma_K(u) - r(u)\phi) K(u) + \lambda_B \epsilon_{h_K}(u)] du. \quad (3.1.6)$$

This shows that KVA can be understood as the discounted expected stream of shareholder hurdle costs, adjusted for the bank's own intensity.

Remark 3.1.5.

1. In the semi-replication framework, only the funding cost adjustment (FCA) is retained, while the funding benefit adjustment (FBA) is omitted. This is because:
 - FBA would effectively duplicate the debit valuation adjustment (DVA), as both capture gains arising from the dealer's own credit spread; including both would result in double counting;
 - In practice, surplus cash cannot be invested at the dealer's unsecured funding rate, since liquidity is centralized and remunerated at collateral or OIS rates. Consequently, the replication argument leads naturally to a one-sided FCA term linked to positive funding needs only.
2. By contrast, collateral valuation adjustment (COLVA) remains because collateral accounts generate observable, contractually specified cash flows at the collateral rate r_X , which are symmetric and replicable; COLVA therefore captures the impact of the $(r - r_X)$ basis on valuation.

To derive more tractable and implementable formulas, we make the assumptions that the collateral is ignored or assumed to be zero ($X = 0$). Using the same setting as Burgard and Kjaer's work [22], we assume the first issuer bond to have zero recovery and use this bond to invest or fund the difference between $\hat{V}$ and V . The P_2 bond position has recovery $R_2 = R_B$ and is given by the funding constraint in Equation (3.1.1). Hence we have

$$\begin{aligned}\alpha_1 P_1 &= -U \\ \alpha_2 P_2 &= -(V - \phi K).\end{aligned}$$

Thus

$$P_D = -R_B(V - \phi K)$$

and hence

$$\epsilon_h = (1 - R_B)[V^+ - \phi K]$$

with

$$\epsilon_{h_0} = (1 - R_B)V^+ \quad \text{and} \quad \epsilon_{h_K} = -(1 - R_B)\phi K.$$

Under these assumptions, the formulas (3.1.2), (3.1.3), (3.1.4) and (3.1.6) simplify to the explicit expressions below:

$$\text{CVA} = -(1 - R_C) \int_t^T \lambda_C(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [(V(u))^+] du, \quad (3.1.7)$$

$$\text{DVA} = -(1 - R_B) \int_t^T \lambda_B(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [(V(u))^-] du, \quad (3.1.8)$$

$$\text{FCA} = -(1 - R_B) \int_t^T \lambda_B(u) e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [(V(u))^+] du, \quad (3.1.9)$$

$$\text{KVA} = - \int_t^T e^{-\int_t^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}_t [(\gamma_K(u) - r_B(u)\phi)K(u)] du, \quad (3.1.10)$$

where we use the fact that $r(u) + (1 - R_B)\lambda_B = r_B(u)$. In the following numerical experiments, we will compute the XVA adjustments based on these simplified expressions under representative transaction scenarios.

In this framework, we are effectively treating CVA, DVA and FVA as hedgeable or at least as directly priceable adjustments. For KVA, although we derived a formula for it, one must recognize that we assumed the cost of capital is just an additive expected cost and there is no stochastic hedge for it. Semi-replication thus assumes that it is valid to charge the expected cost $hK(t)$ without any additional risk premium. In other words, the semi-replication approach effectively treats capital costs as if they were hedgeable fees, analogous to how one might treat a known funding spread. This is a strong assumption, given that capital cost is not linked to a traded asset. We highlight that Green et al. explicitly noted the inability to truly hedge KVA. Nevertheless their method prices KVA under the risk-neutral measure as if it were a predictable cost.

3.2 Indifference Pricing Method

The indifference pricing approach takes a different perspective: rather than expanding the replicating portfolio, it asks what price must be charged such that taking on the trade does not worsen the bank's position. The notion of indifference can be defined in terms of utility, maximum of expected value, or an internal performance metric. The key is that we do not assume we can hedge away the risky components of the cost. Instead, we require compensation for bearing those risks. This approach is naturally suited to KVA, since

capital cost is an untradeable risk. But it can also be applied to other unhedgeable costs if we assume some cost is unhedgeable. In contrast to the semi-replication, the indifference approach derives the cost of capital from first principles, ensuring that the trade's price makes the shareholders indifferent adding the new trade to the bank's portfolio.

This section introduces the theoretical formulation of the indifference method. We discuss two concrete implementations: one based on RAROC preservation, and another based on expected equity optimization, considering both whole-bank and shareholder perspectives.

We consider a one-period model in a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$. At time 0, the bank chooses a trading strategy $\theta \in \mathbb{R}^d$ in a frictionless market with asset prices S_0 and random terminal values $S_1 \in \mathbb{R}^d$. The bank has initial capital C_0 , trades at borrowing rate $r + \lambda$, and considers entering a client-proposed OTC trade of q units with random payoff Y .

The mathematical expression for assets and liabilities of our bank at time 1, for our chosen strategy θ are:

$$\begin{aligned}\text{Assets} &= (\theta^\top S_1)^+ + (qY)^+ + (C_0 - \theta^\top S_0 + P(q))^+(1 + r + \lambda) \\ \text{liabilities} &= (\theta^\top S_1)^- + (qY)^- + (C_0 - \theta^\top S_0 + P(q))^- (1 + r + \lambda).\end{aligned}$$

Then the equity at time 1 is:

$$\begin{aligned}X(\theta, q) &= \text{Assets} - \text{liabilities} \\ &= qY + \theta^\top (S_1 - S_0(1 + r + \lambda)) + (C_0 + P(q))(1 + r + \lambda).\end{aligned}$$

3.2.1 RAROC-based Indifference Method

A pragmatic method for incorporating capital costs in deal pricing uses the concept of Risk-Adjusted Return on Capital (RAROC). RAROC is a performance metric defined as the ratio of a deal's expected profit to the allocated economic capital for that deal. Banks typically set a target RAROC, or hurdle rate h , for each business line or deal type to ensure adequate risk-adjusted returns. If a proposed trade has an expected RAROC that falls below this hurdle rate, it would dilute the bank's overall return on capital and may be deemed uneconomical. The RAROC-based KVA approach addresses this by charging an additional fee on the trade sufficient to raise its expected RAROC to meet or exceed the target hurdle rate.

If we make the assumption that the Value at Risk of our portfolio is proportional to the standard deviation of the portfolio, the RAROC metric is then defined as:

$$\text{RAROC}(\theta, q) = \frac{\mathbb{E}[X(\theta, q)] - C_0}{m\sqrt{\text{Var}[X(\theta, q)]}}.$$

The price $P(q)$ is set such that the RAROC is preserved and equal to the hurdle rate:

$$h = \text{RAROC}(\theta, 0) = \text{RAROC}(\theta', q),$$

where θ' is the new optimal strategy when including the trade. θ' is somehow different from θ and will possibly include a hedge of the product and will need to take into account capital constraints. To simplify, we assume the new product is contains only unhedgeable risks, independent from the other marketed product, and that we can treat $\theta' = \theta$.

Then we have

$$h = \frac{\mathbb{E}[X(\theta, q)] - C_0}{m\sqrt{\text{Var}[X(\theta, q)]}}.$$

Under the assumption of independence between Y and S_1 we have

$$h = \frac{\mathbb{E}[X(\theta, 0)] + q\mathbb{E}[Y] + P(q)(1 + r + \lambda) - C_0}{m\sqrt{\text{Var}[X(\theta, 0)] + q^2\text{Var}[Y]}},$$

which leads to

$$P(q) = \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] - \mathbb{E}[X(\theta, 0)] + C_0 + hm\sqrt{\text{Var}[X(\theta, 0)] + q^2\text{Var}[Y]} \right).$$

If we expand $h = \text{RAROC}(\theta, 0)$, we get

$$(\mathbb{E}[X(\theta, 0)] - C_0) \left(\sqrt{1 + \frac{q^2\text{Var}[Y]}{\text{Var}[X(\theta, 0)]}} \right) - \mathbb{E}[X(\theta, 0)] + C_0 = q\mathbb{E}[Y] + P(q)(1 + r + \lambda).$$

For small q , expanding using Taylor Expansion around $q = 0$ gives:

$$(\mathbb{E}[X(\theta, 0)] - C_0) \left(1 + \frac{q^2\text{Var}[Y]}{2\text{Var}[X(\theta, 0)]} \right) - \mathbb{E}[X(\theta, 0)] + C_0 \approx q\mathbb{E}[Y] + P(q)(1 + r + \lambda).$$

Hence by rearranging the term we have

$$\begin{aligned} P(q) &\approx \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] + \frac{\mathbb{E}[X(\theta, 0)] - C_0}{m\sqrt{\text{Var}[X(\theta, 0)]}} \frac{mq^2\text{Var}[Y]}{2\sqrt{\text{Var}[X(\theta, 0)]}} \right) \\ &= \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] + h \frac{mq^2\text{Var}[Y]}{2\sqrt{\text{Var}[X(\theta, 0)]}} \right). \end{aligned}$$

So the acceptable price can be computed as the discounted expectation of the cash flows due to our new product plus an adjustment proportional to the variance of the new product and the hurdle rate h . This RAROC rule is straightforward and often used as a quick heuristic in banks.

3.2.2 Optimization-based Indifference Method

An alternative performance metric for indifference method is an optimization problem for expected payoff of equity. We suppose that the bank would choose θ to maximize its value as the whole bank's strategy. More precisely for a certain set $\Theta(q) \subseteq \mathbb{R}^d$ of admissible strategies, we define the optimal strategy as computed via the constrained maximization

$$\tilde{v}(q) = \sup_{\theta \in \Theta(q)} \mathbb{E}[X(\theta, q)]. \quad (3.2.1)$$

To model regulatory capital constraints, we impose a risk constraint approximated by:

$$\Theta(q) = \{\theta \in \mathbb{R}^2 | C_0 = \nu \sqrt{\theta^\top A \theta + 2q\theta^\top b + \sigma_Y^2 q^2}\}, \quad (3.2.2)$$

where A is the covariance of S_1 , b is the covariance vector between S_1 and Y , and σ_Y^2 is the variance of Y .

Similar with the RAROC-based indifference method, we look for a price $P(q)$ such that the bank is indifference with respect to enter the contract. So we look for $P(q)$ such that

$$\tilde{v}(0) = \tilde{v}(q).$$

The solution to our problem is given by the following proposition.

Proposition 3.2.1. Consider, for a fixed q , the problem of finding P such that

$$\tilde{v}(q) = \tilde{v}(0), \quad (3.2.3)$$

where $\tilde{v}(q)$ and $\Theta(q)$ are defined in (3.2.1) and (3.2.2) respectively.

Define $\hat{\mu} = (S_1 - S_0(1 + r + \lambda))$, and $\mu = \mathbb{E}[(S_1 - S_0(1 + r + \lambda))]$, its expected value under the real world measure. Then equation (3.2.3) has a solution and is given by

$$P(q) = \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] + \frac{1}{2} \left(\frac{1}{\chi(0)} - \frac{1}{\chi(q)} \right) \mu^\top A^{-1} \mu + qb^\top A^{-1} \mu \right),$$

where $\chi(q)$ is a Lagrange multiplier satisfying:

$$\chi(q) = \frac{\sqrt{\mu^\top A^{-1} \mu}}{2\sqrt{\frac{C_0^2}{\nu^2} - q^2(\sigma_Y^2 - b^\top A^{-1} b)}}.$$

Moreover, for small deal such that q is small, we have the following:

$$P(q) \approx \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] + \frac{1}{2} \left(\frac{(\sigma_Y^2 - b^\top A^{-1} b)q^2\nu}{C_0} \right) \mu^\top A^{-1} \mu + qb^\top A^{-1} \mu \right).$$

The proof can be found in [13]. This proposition tells us that the price can be decomposed into three terms. The first term in the indifference price expression corresponds to the discounted expected payoff under the real-world measure, and can be interpreted as the baseline economic value of the trade in the absence of capital constraints or valuation adjustments. The remaining terms then constitute the valuation adjustment, compensating for the cost of regulatory capital and the inability to perfectly hedge the associated risks.

Note that, unlike classical risk-neutral pricing, the indifference framework operates under the real-world measure, as it seeks to preserve economic performance rather than enforce no-arbitrage replication.

To gain further insight into the economic interpretation of the pricing formula, we now consider the expected P&L associated with the optimal strategy under the capital constraint. Let $\tilde{v}(0)$ denote the expected value of equity under optimal allocation without the new trade. From the previous derivation, we have:

$$\tilde{v}(0) - C_0 = \left(\frac{\mu^\top A^{-1} \mu}{2\chi(0)} \right) + C_0(r + \lambda) = \frac{\sqrt{\mu^\top A^{-1} \mu}}{\nu} C_0 + C_0(r + \lambda).$$

This expression motivates a natural definition of RAROC, defined as the ratio of expected excess return to capital employed. Using the same risk metric (VaR at 99%) for both the capital constraint and RAROC denominator, we define:

$$\text{RAROC}(\theta, 0) = \frac{\sqrt{\mu^\top A^{-1} \mu}}{\nu} + r + \lambda =: h, \quad (3.2.4)$$

where h is the effective hurdle rate for capital usage, incorporating both the intrinsic return of the optimal strategy and the external funding cost $r + \lambda$.

Now substituting this result into the indifference pricing formula, we obtain a simplified and more interpretable expression for $P(q)$ under the small- q approximation:

$$P(q) \approx \frac{1}{1 + r + \lambda} \left(-q\mathbb{E}[Y] + \frac{1}{2} \left(\frac{(\sigma_Y^2 - b^\top A^{-1} b)q^2\nu^2}{C_0} \right) (h - r - \lambda) + qb^\top A^{-1} \mu \right). \quad (3.2.5)$$

The term $(h - r - \lambda)$ captures the excess return required to compensate for the cost of capital beyond basic funding, and hence plays a central role in quantifying the capital constraint's pricing impact.

3.2.3 Alternative Perspective

Beyond the RAROC and optimization-based formulations, other variants of the indifference approach have been proposed. One such perspective is from shareholders under limited liability, where the optimization is applied to the positive part of equity, reflecting that downside losses are absorbed by creditors. Another is to replace expectation with the median as the performance criterion, aligning the interests of shareholders and bondholders.

While these approaches provide additional economic interpretations of capital costs, they are only briefly introduced here, as our numerical analysis focuses on the RAROC and optimization-based methods, which are more tractable and commonly used in practice.

Indifference Pricing from Shareholder's Perspective

From the shareholder's viewpoint, limited liability leads to the optimization:

$$v(q) = \sup_{\theta \in \Theta(q)} \mathbb{E} [(X(\theta, q))^+], \quad (3.2.6)$$

where $\Theta(q)$ is defined in (3.2.2).

We seek a price $P(q)$ such that the shareholder's value remains unchanged:

$$v(q) = v(0).$$

To obtain a tractable solution, we expand $v(q)$ near $q = 0$ as a second-order Taylor approximation:

$$v(q) = v(0) + q \left. \frac{dv}{dq} \right|_{q=0} + \frac{q^2}{2} \left. \frac{d^2v}{dq^2} \right|_{q=0} + o(q^2).$$

Imposing the indifference condition $v(q) = v(0)$ implies:

$$q \left. \frac{dv}{dq} \right|_{q=0} + \frac{q^2}{2} \left. \frac{d^2v}{dq^2} \right|_{q=0} = 0.$$

The following proposition gives a solution to the limited liability problem in the small deal approximation.

Proposition 3.2.2. *Consider, for fixed q , the problem of finding P such that*

$$v(q) = v(0), \quad (3.2.7)$$

where $v(q)$ and $\Theta(q)$ are defined in (3.2.6) and (3.2.2) respectively. The second-order marginal price $P(q)$ that satisfies equation (3.2.7) is given by

$$P(q) \approx \frac{\mathbb{E}[-\mathbf{1}_{D^c}Y] + \chi(0)2\theta^\top b}{\mathbb{P}[D^c](1+r+\lambda)}q - \frac{1}{2} \cdot \frac{\psi(0) - \chi(0)2\sigma_Y^2 - 2q^2\chi'(0)(2\theta^\top b) + \mathcal{C}(0)}{(1+r+\lambda)}q^2,$$

where:

- D^c denotes the survival event: $D^c = \{\omega \in \Omega \mid X(\theta, q) > 0\}$.
- $\chi(0)$ is the Lagrange multiplier at $q = 0$.

- $\psi(0)$ captures convexity under the default boundary:

$$\psi(0) = \lim_{h \rightarrow 0} \mathbb{E} \left[(Y + P(0)(1 + r + \lambda)) \cdot \frac{\mathbf{1}_{D^c(h)} - \mathbf{1}_{D^c}}{h} \right],$$

- $\mathcal{C}(0)$ is the second-order term due to the sensitivity of the optimal strategy $\theta^*(q)$ to q :

$$\mathcal{C}(0) = \sum_{i=1}^d \sum_{j=1}^d \left(\chi(0) \frac{\partial^2 g}{\partial \theta_i \partial \theta_j} \Big|_{q=0} - \frac{\partial^2 f}{\partial \theta_i \partial \theta_j} \Big|_{q=0} \right) \frac{\partial \theta_j^*}{\partial q} \Big|_{q=0} \frac{\partial \theta_i^*}{\partial q} \Big|_{q=0}.$$

The proof can be found in [13]. This pricing formula consists of several interpretable components. The term $\mathbb{E}[-\mathbf{1}_{D^c} Y]$ captures the expected loss conditional on survival. The covariance terms $-q\chi(0)2\theta^\top b$ and $-2q^2\chi'(0)(2\theta^\top b)$ reflect the interaction between the trade and the existing portfolio. The quantities $\psi(0)$ and $\mathcal{C}(0)$ represent convexity effects and sensitivity of the optimal strategy. Lastly, the capital penalty term $-\chi(0)2\sigma_Y^2$ accounts for the cost of capital due to unhedgeable risk.

Median-based Indifference Method

As an alternative to expected value, we consider the median as the objective in the indifference valuation problem under limited liability. Let $\mathcal{M}[X]$ denote the median of a real-valued random variable X , defined as:

$$\mathcal{M}[X] := \inf \left\{ m \mid \mathbb{P}[X \leq m] > \frac{1}{2} \right\}.$$

We assume the existence of an admissible strategy $\bar{\theta} \in \Theta(q)$ such that the bank remains solvent on average. The solution of median approach can be formulated as below.

Proposition 3.2.3. *Consider, for a fixed q , the problem of finding P such that*

$$w(q) = w(0), \tag{3.2.8}$$

where

$$w(q) = \sup_{\theta \in \Theta(q)} \mathcal{M} \left[qY + \theta^\top (S_1 - S_0(1 + r + \lambda)) + (C_0 + P(q))(1 + r + \lambda) \right]^+.$$

and Θ is defined in equation (3.2.2). Equation (3.2.8) has the same solution of Equation (3.2.3) and is hence given by

$$P(q) = \frac{1}{1 + r + \lambda} \left(q\mathbb{E}[-Y] + \frac{1}{2} \left(\frac{1}{\chi(0)} - \frac{1}{\chi(q)} \right) \mu^\top A^{-1} \mu + 2qb^\top A^{-1} \mu \right).$$

The proof can be found in [13]. Thus, the median-based indifference price coincides with the linear case. In practice, using the median aligns shareholder and bondholder interests, as both are protected from downside risk, reinforcing the economic interpretation of capital costs in the valuation process.

3.2.4 KVA for Indifference Method

We define the unhedgeable valuation adjustment (UVA) as

$$\text{UVA}(q) := P(q) - V. \tag{3.2.9}$$

By construction, $\text{UVA}(q)$ captures the entire premium the bank requires on top of the replication (risk-neutral) price V in order to remain economically indifferent once non-replicable frictions are taken into account: regulatory capital, shareholders' hurdle rate, funding asymmetries not spanned by traded instruments, and the interaction with the legacy portfolio under capital constraints. Hence, in the indifference framework, $\text{UVA}(q)$ plays the role of “all residual XVAs that cannot be priced by pure no-arbitrage replication”. In particular, the capital component KVA is embedded in $\text{UVA}(q)$.

A central theoretical subtlety of equation (3.2.9) lies in the fact that its two constituent legs are evaluated under different probability measures. The indifference price $P(q)$ is obtained under the real-world measure $\mathbb{P}$, whereas the clean, replication-based price V is computed under the risk-neutral measure $\mathbb{Q}$. This misalignment of measures is not a mere technicality; it has far-reaching implications for both interpretation and implementation of the overall valuation adjustment $\text{UVA}(q) = P(q) - V$.

The difference between the two expectations, $\mathbb{E}^{\mathbb{P}}[\cdot] - \mathbb{E}^{\mathbb{Q}}[\cdot]$, embeds the entire market price of risk. Hence, even in a stylised world with no capital constraints, $P(q)$ would in general deviate from V solely because of the measure change. In other words, $\text{UVA}(q)$ always contains a pure risk-premium component in addition to any capital-cost effect. It follows that equating the whole of $\text{UVA}(q)$ to all the unhedgeable valuation adjustment is conceptually imprecise, since part of the difference would persist even if the regulatory capital requirement were zero.

Additionally, the result in the indifference framework is non-linear. Because $P(q)$ is the solution to a constrained optimization problem—featuring both limited liability and capital constraints—the mapping $q \mapsto P(q)$ is generally non-linear. Consequently,

$$\text{UVA}(q_1 + q_2) \neq \text{UVA}(q_1) + \text{UVA}(q_2),$$

undermining the additivity property that underpins most desk-level XVA allocation and back-testing methodologies.

Definition 3.2.4 (Radon–Nikodym Derivative). Let $\mathbb{P}$ and $\mathbb{Q}$ be two equivalent probability measures defined on the same filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]})$. The Radon–Nikodym derivative of $\mathbb{Q}$ with respect to $\mathbb{P}$, denoted by

$$\xi_T := \left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_T},$$

is a positive $\mathcal{F}_T$ -measurable random variable that enables us to translate expectations under $\mathbb{Q}$ into expectations under $\mathbb{P}$. Specifically, for any $\mathcal{F}_T$ -measurable random variable X , we have

$$\mathbb{E}^{\mathbb{Q}}[X] = \mathbb{E}^{\mathbb{P}}[\xi_T X].$$

Theorem 3.2.5 (Girsanov's Theorem). Let $\mathbb{P}$ and $\mathbb{Q}$ be two equivalent probability measures on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]})$. Suppose X_t follows a diffusion process under $\mathbb{P}$:

$$dX_t = f^{\mathbb{P}}(X_t) dt + \sigma(X_t) dW_t^{\mathbb{P}},$$

and we define a new measure $\mathbb{Q}$ via the density process

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} := \exp \left\{ -\frac{1}{2} \int_0^t \left(\frac{f^{\mathbb{Q}}(X_s) - f^{\mathbb{P}}(X_s)}{\sigma(X_s)} \right)^2 ds + \int_0^t \frac{f^{\mathbb{Q}}(X_s) - f^{\mathbb{P}}(X_s)}{\sigma(X_s)} dW_s^{\mathbb{P}} \right\},$$

then under $\mathbb{Q}$, X_t evolves as

$$dX_t = f^{\mathbb{Q}}(X_t) dt + \sigma(X_t) dW_t^{\mathbb{Q}},$$

where $W_t^{\mathbb{Q}}$ is a Brownian motion under $\mathbb{Q}$. A sufficient condition ensuring the validity of this measure change is the Novikov condition:

$$\mathbb{E}^{\mathbb{P}} \left[\exp \left\{ \frac{1}{2} \int_0^T \left(\frac{f^{\mathbb{Q}}(X_s) - f^{\mathbb{P}}(X_s)}{\sigma(X_s)} \right)^2 ds \right\} \right] < \infty.$$

From a mathematical perspective, by Definition 3.2.4 and Theorem 3.2.5, let $\xi_T := \frac{d\mathbb{P}}{d\mathbb{Q}}|_{\mathcal{F}_T}$ be the Radon–Nikodym derivative. A purely algebraic decomposition highlights the misalignment term:

$$\begin{aligned} P(q) &= \frac{1}{1+r+\lambda} \mathbb{E}^{\mathbb{P}}[qY] + (\text{capital \& constraint terms}) \\ &= \frac{1}{1+r+\lambda} \mathbb{E}^{\mathbb{Q}}[\xi_T qY] + (\text{capital \& constraint terms}), \\ P(q) - V &= \underbrace{\left[\frac{1}{1+r+\lambda} \mathbb{E}^{\mathbb{Q}}(\xi_T qY) - \mathbb{E}^{\mathbb{Q}}\left(e^{-\int_0^T r_u du} qY\right) \right]}_{\text{measure/discount misalignment}} \\ &\quad + \underbrace{(\text{capital, hurdle-rate and portfolio-interaction terms})}_{\text{KVA-like}}. \end{aligned}$$

The first bracketed term is a measure-alignment adjustment and disappears only if (i) ξ_T is proportional to the risk-neutral discount factor and (ii) funding/discount conventions coincide. Otherwise, UVA(q) cannot be interpreted as a “pure KVA”.

3.3 Exposure Profiles and Capital Requirements

To evaluate risk profiles and derivative prices numerically, we first define standard exposure processes in continuous time and then introduce their Monte Carlo estimators. Let $V(t)$ denote the net mark-to-market of the portfolio from the bank’s perspective. The usual exposure measures are:

$$\begin{aligned} \text{EE}(t) &= \mathbb{E}[V(t)^+], \\ \text{ENE}(t) &= \mathbb{E}[V(t)^-] = \mathbb{E}[(-V(t))^+], \\ \text{PFE}_{\alpha}(t) &= \inf\{x \geq 0 : \mathbb{P}(V(t)^+ \leq x) \geq \alpha\}, \quad \alpha \in (0, 1), \\ \text{EPE} &= \frac{1}{T} \int_0^T \text{EE}(t) dt. \end{aligned}$$

In regulatory frameworks such as Basel III/IV, the capital requirement K_t is typically defined as a multiple of these exposure measures, for example through potential future exposure (PFE) or expected positive exposure (EPE). Thus, the specification of K_t is closely tied to the choice of exposure metric.

In this thesis, for tractability and comparability across methodologies, we adopt a simplified representation:

$$K_t = \alpha \cdot \text{EE}(t), \tag{3.3.1}$$

where $\alpha > 0$ is a regulatory multiplier. This simplification is deliberate. The actual calculation of Risk-Weighted Assets (RWA) under Basel frameworks is not particularly precise—it assumes risks are additive when in reality they are not, and accounts for this imprecision through various “magic multipliers”, such as the 1.4 multiplier for EPE. These multipliers are regulatory compromises rather than economically derived values. Given these inherent limitations in the regulatory framework, and since our focus is on the

fundamental XVA methodology rather than regulatory minutiae, we do not delve deeply into the complexities of RWA calculations. The simplified form (3.3.1) is sufficient to illustrate the key concepts of XVA pricing and the double counting challenges, while avoiding unnecessary complications from regulatory artifacts that lack strong economic foundations.

3.4 Netting, Wrong-Way Risk and Double Counting

Having developed the core methodologies, we now extend the discussion to address three practical complexities that any comprehensive KVA framework must consider: netting of exposures, wrong-way risk (WWR) between exposure and counterparty default, and potential double counting of risks across overlapping valuation adjustments. These factors are important for ensuring that our pricing adjustments are not only theoretically sound but also practically applicable in a portfolio context.

3.4.1 Netting of Exposures

In reality, derivatives are traded under netting agreements where all trades with a counterparty are netted for the purposes of default close-out and often for collateral. Netting can dramatically reduce the exposure profile of a portfolio compared to the sum of exposures of individual trades. From a capital perspective, regulatory capital is calculated on a netting set basis—meaning a counterparty’s portfolio is evaluated as one combined exposure.

This means the incremental capital usage of a new trade depends on what else is in the portfolio with that counterparty. If a bank adds a trade that is negatively correlated with existing exposures (hedging them), it might add very little incremental capital or even reduce total capital. If one priced each trade separately ignoring netting, one could overcharge KVA. For example, two trades that largely offset each other in exposure would each, on a standalone basis, seem to require capital, but together they might require almost none. Charging KVA on both could double count capital that in reality the bank does not need to hold because of the risks offset.

The semi-replication approach can handle netting by simply using portfolio-level net exposure

$$E_t^{\text{net}} = \mathbb{E}[(V_A(t) + V_B(t))^+]$$

in the CVA/FVA integrals, where this represents the expected exposure after netting within each netting set.

The indifference approach naturally incorporates netting effects through the portfolio’s joint risk distribution. When computing RAROC or utility, diversification benefits are captured in the covariance terms: if Y_A and Y_B are negatively correlated, i.e., $\text{Cov}[Y_A, Y_B] < 0$, the total portfolio variance

$$\text{Var}[Y_A + Y_B] = \text{Var}[Y_A] + \text{Var}[Y_B] + 2\text{Cov}[Y_A, Y_B]$$

will be reduced, leading to lower capital requirements and improved risk-adjusted returns. Our implementation considers the combined distribution of multiple trades when optimizing the overall portfolio utility.

3.4.2 Wrong-Way Risk Mathematical Framework

Wrong-way risk (WWR) and right-way risk (RWR) refer to the dependence between counterparty exposure and default probability. The classification depends on whether this

dependence increases or decreases the credit valuation adjustment:

$$\begin{aligned}\text{Wrong-way risk: } \text{Corr}[V^+(t), \lambda_C(t)] > 0 &\Rightarrow |\text{CVA}| \text{ increases} \\ \text{Right-way risk: } \text{Corr}[V^+(t), \lambda_C(t)] < 0 &\Rightarrow |\text{CVA}| \text{ decreases}\end{aligned}$$

where $V^+(t)$ represents the positive exposure at time t and $\lambda_C(t)$ is the counterparty's default intensity.

Under the independence assumption, CVA is calculated as:

$$\text{CVA}_{\text{indep}} = (1 - R_C) \int_0^T \lambda_C(u) e^{-\int_0^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} \mathbb{E}[V(u)^+] du$$

When incorporating dependence between exposure and default risk, CVA becomes:

$$\text{CVA}_{\text{dep}} = (1 - R_C) \int_0^T \mathbb{E}[V(u)^+ \lambda_C(u)] e^{-\int_0^u (r(s) + \lambda_B(s) + \lambda_C(s)) ds} du$$

The dependence effect is captured through the covariance decomposition:

$$\mathbb{E}[V(u)^+ \lambda_C(u)] = \mathbb{E}[V(u)^+] \mathbb{E}[\lambda_C(u)] + \text{Cov}[V(u)^+, \lambda_C(u)]$$

The impact on CVA depends on the sign of the covariance:

- When $\text{Cov}[V(u)^+, \lambda_C(u)] > 0$: Wrong-way risk increases $|\text{CVA}|$
- When $\text{Cov}[V(u)^+, \lambda_C(u)] < 0$: Right-way risk decreases $|\text{CVA}|$

In our setup, regulatory capital is defined in Equation (3.3.1). WWR affects $K(t)$ only through its impact on the exposure profile. $\text{EE}_{\text{WWR}}(t_k) = \mathbb{E}[(V(t_k))^+]$ is computed under the joint dynamics with exposure–credit correlation used to model WWR.

3.4.3 Double Counting

Double counting represents one of the most significant practical challenges in XVA implementation. It occurs when the same economic risk or cost is inadvertently charged through multiple adjustment channels, leading to an overstatement of the total valuation adjustment.

Consider a simple example: A bank's exposure to a counterparty drives both funding costs (through FVA) and regulatory capital requirements (through KVA). If both adjustments independently charge for the same exposure-related costs, the bank would be double-charging for essentially the same economic risk. This not only distorts pricing but can make the bank uncompetitive in the market.

Rather than attempting a precise mathematical definition, which proves elusive for many risks in mathematical finance, we define double counting conceptually as the risk of pricing the same economic risk in two separate adjustments. This informal approach acknowledges that double counting is fundamentally about economic substance rather than mathematical correlation.

The challenge of defining double counting precisely parallels other areas in quantitative finance where universal mathematical definitions remain contentious. Just as liquidity risk has multiple competing definitions without consensus, double counting resists clean mathematical characterization while remaining an important practical concern.

Remark 3.4.1. For

$$U_0 = \text{CVA}_0 + \text{DVA}_0 + \text{FCA}_0 + \text{COLVA}_0 + \text{KVA}_0,$$

if we write each adjustment in stochastic-integral form:

$$\text{XVA}_i = \int_0^T e^{-\int_0^u r(s)ds} \mathbb{E}_0^Q[\Xi_i(u)] du,$$

where the process $\Xi_i(u)$ is the driver of XVA_i , we should notice that a purely statistical “no-overlap” rule, such as requiring zero covariance between adjustment, i.e.,

$$\text{Cov}^Q[\Xi_i(t), \Xi_j(t)] = 0 \quad \forall t \in [0, T],$$

is too strong. Because several XVAs respond to related drivers, some statistical dependence between them is inevitable and does not, by itself, imply double counting. For example, CVA and DVA naturally reference the same exposure with opposite sign and will typically be correlated; that correlation is not double counting.

Historically, genuine double counting was noticed when FVA, and in particular FBA under a treasury “net borrower” policy, effectively collapsed to DVA. In such situations, charging both would amount to paying for the same risk twice.

For KVA, common example occurs with FCA. Both adjustments are fundamentally driven by the bank’s exposure to counterparties: FCA captures the cost of funding positive exposure to counterparties, and KVA captures the cost of holding regulatory capital, which is often calculated based on the same positive exposure.

Since regulatory capital requirements are frequently calculated using Expected Exposure (EE), both FCA and KVA end up charging for costs related to the same underlying exposure. This creates a situation where the bank is essentially pricing the same risk twice—once through the direct funding cost and again through the capital cost.

The most straightforward approach to address double counting is to ensure that each economic risk is captured in only one adjustment. In practice, this typically means adjusting the capital calculation to exclude components already accounted for in other XVAs.

For example, when calculating KVA, banks can modify their capital requirements to exclude the exposure-related components that are already captured in FCA. This involves separating capital requirements into:

- Exposure-related capital (already captured in funding adjustments)
- Non-exposure-related capital (captured in KVA)
- Other regulatory capital components unrelated to counterparty exposure

This separation ensures that the economic cost of counterparty exposure is charged once through the funding adjustment, while other capital costs are appropriately captured through KVA.

3.5 Example Instruments and Models

In this section, we introduce the two financial instruments that will be used in the numerical experiments, namely a European call option and an interest rate swap. We also present the market models employed for the simulations: the Black–Scholes model for the equity derivative and the Vasicek model for the interest rate product.

3.5.1 European Call Option under Black-Scholes Model

A European call option is a contract that gives the holder the right, but not the obligation, to buy an underlying asset at a predetermined price K , the strike price at a specified future date T , the maturity. The payoff of a call option at maturity is:

$$\text{Payoff} = \max(S_T - K, 0) = (S_T - K)^+$$

where S_T is the underlying asset price at maturity. The seller of the call receives an upfront premium, which is the option price and is obligated to pay the payoff at T if exercised.

Under the Black-Scholes [7] framework, we assume the underlying asset follows a geometric Brownian motion under the risk-neutral measure $\mathbb{Q}$:

$$dS_t = rS_t dt + \sigma S_t dW_t^{\mathbb{Q}}$$

where r is the risk-free interest rate, σ is the volatility, and $W_t^{\mathbb{Q}}$ is a Brownian motion under $\mathbb{Q}$.

The Black-Scholes equation for a derivative price $V(t, S_t)$ is given by:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 V}{\partial S_t^2} + rS_t \frac{\partial V}{\partial S_t} - rV = 0$$

with terminal condition $V(T, S_T) = (S_T - K)^+$.

Using Theorem 3.1.4 yields the Black-Scholes formula for the call price. At inception $t = 0$, the call's fair price is given in closed form by:

$$V_{BS}(t, S_t) = S_t \Phi(d_1) - Ke^{-r(T-t)} \Phi(d_2)$$

where:

$$d_1 = \frac{\ln(S_t/K) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}$$

$$d_2 = d_1 - \sigma\sqrt{T - t}$$

and $\Phi(\cdot)$ denotes the cumulative standard normal distribution function.

The European call option is chosen as a simple yet insightful test case for several reasons. First, it is one of the most fundamental derivative contracts, with well-understood properties and an analytic pricing solution in the Black-Scholes model. This allows us to validate our numerical methods against a known benchmark. Second, the call option's exposure profile is simple, which is a one-factor model driven by the underlying stock, making it easier to isolate the impact of various XVA adjustments. The call's payoff depends only on the terminal stock price and has no intermediate cashflows or early exercise features, which simplifies counterparty exposure modeling.

3.5.2 Interest Rate Swap under the Vasicek Model

An Interest Rate Swap (IRS) is a fundamental over-the-counter derivative instrument that facilitates the exchange of cash flows between two counterparties based on different interest rate indices [17]. The contract structure involves periodic payments over a predetermined schedule, typically with one leg paying a fixed rate and the other paying a floating rate tied to a reference index such as LIBOR or SOFR.

As shown in Figure 3.1, which is adapted from [8], consider an IRS contract initiated at time t with maturity T_β and payment dates $\{T_{\alpha+1}, T_{\alpha+2}, \dots, T_\beta\}$, where $T_\alpha < t \leq T_{\alpha+1}$. Let $\tau_i = T_i - T_{i-1}$ denote the day count fraction for the i -th period, and K represent the fixed rate. The cash flows at each payment date T_i are structured such that the fixed leg payment equals $\tau_i K N$ and the floating leg payment equals $\tau_i L(T_{i-1}, T_i) N$, where N is the notional principal and $L(T_{i-1}, T_i)$ represents the floating rate for the period $[T_{i-1}, T_i]$.

The present value of an IRS can be expressed as the difference between the present values of the fixed and floating legs. For a receiver IRS, which involves receiving fixed rate payments and paying floating rate payments, the value at time t is given by:

$$\text{IRS}(t) = \sum_{i=\alpha+1}^{\beta} P(t, T_i) \tau_i (K - F(t; T_{i-1}, T_i)) N$$




Figure 3.1: Interest Rate Swap Illustration

where $P(t, T_i)$ denotes the price of a zero-coupon bond maturing at time T_i , and $F(t; T_{i-1}, T_i)$ represents the forward rate for the period $[T_{i-1}, T_i]$ as observed at time t .

The spot linear rate $L(T_{i-1}, T_i)$ realized at time T_{i-1} for the period $[T_{i-1}, T_i]$ can be expressed in terms of bond prices as:

$$L(T_{i-1}, T_i) = \frac{P(T_{i-1}, T_{i-1}) - P(T_{i-1}, T_i)}{\tau_i P(T_{i-1}, T_i)} = \frac{1 - P(T_{i-1}, T_i)}{\tau_i P(T_{i-1}, T_i)}$$

The relationship between the spot linear rate and forward rate is established through the risk-neutral expectation:

$$\mathbb{E}_t^Q[D(t, T_i)L(T_{i-1}, T_i)] = P(t, T_i)F(t; T_{i-1}, T_i)$$

where $\mathbb{E}_t^Q[\cdot]$ denotes the risk-neutral expectation conditional on information at time t , and $D(t, T_i)$ is the discount factor from time t to time T_i . This fundamental relationship ensures that the forward rate $F(t; T_{i-1}, T_i)$ provides an unbiased predictor of the future spot linear rate under the risk-neutral measure, which is essential for the no-arbitrage valuation of interest rate derivatives.

The Vasicek model, introduced by Oldřich Vašíček in 1977 [36], represents one of the earliest and most influential short-rate models in mathematical finance. This single-factor model describes the evolution of the instantaneous short rate through a mean-reverting stochastic differential equation.

The Vasicek model postulates that the short rate $r(t)$ follows the stochastic differential equation:

$$dr(t) = \kappa(\theta - r(t))dt + \sigma dW(t)$$

where:

- $\kappa > 0$ is the mean reversion parameter, governing the speed of adjustment toward the long-term mean
- θ represents the long-term mean level of the short rate
- $\sigma > 0$ is the volatility parameter
- $W(t)$ is a standard Brownian motion under the risk-neutral measure

The Vasicek model exhibits several desirable mathematical properties that facilitate analytical tractability. The drift term $\kappa(\theta - r(t))$ ensures that the short rate exhibits mean-reverting behavior, with stronger reversion when $r(t)$ deviates significantly from θ . Also, the distribution of short rate $r(t)$ and zero coupon bond price is very clear.

Theorem 3.5.1. *The short rate $r(t)$ follows a Gaussian distribution at any future time, with conditional mean and variance given by:*

$$E[r(T)|r(t)] = \theta + (r(t) - \theta)e^{-\kappa(T-t)}$$

$$\text{Var}[r(T)|r(t)] = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa(T-t)})$$

Additionally, under the Vasicek framework, the price of a zero-coupon bond with maturity T at time t admits a closed-form solution:

$$P(t, T) = A(t, T) \exp(-B(t, T)r(t))$$

where the functions $A(t, T)$ and $B(t, T)$ are given by:

$$B(t, T) = \frac{1 - e^{-\kappa(T-t)}}{\kappa}$$

$$A(t, T) = \exp \left\{ \left(\theta - \frac{\sigma^2}{2\kappa^2} \right) (B(t, T) - (T - t)) - \frac{\sigma^2 B(t, T)^2}{4\kappa} \right\}$$

The proof of theorem can be found in [17]. Although Vasicek model is limited by its allowance for negative rates, constant parameters, and a single-factor structure that may not fully capture real-world yield curve dynamics, the Vasicek model remains a cornerstone in interest rate modeling due to its mathematical elegance and analytical tractability, making it an ideal benchmark for more sophisticated models and a valuable tool for understanding fundamental interest rate dynamics.

Chapter 4

Numerical Implementation Strategies

Having established the theoretical foundations, we now turn to their numerical implementation. The purpose of this chapter is to construct a coherent computational framework that allows for meaningful comparison across different methodologies. A central difficulty arises from the fact that the semi-replication and indifference pricing approaches rely on distinct sets of inputs and assumptions. To overcome this, we define a unified parameterization and tried to impose consistent modelling conventions as much as we can. We also introduce hedgeability conditions, under which XVA components are treated differently. In addition, we outline the numerical strategies employed—covering discretization choices, simulation methods, and calibration procedures—so that the subsequent results can be interpreted within a common and transparent framework.

4.1 Methodological Alignment and Unified Parameter

The fundamental challenge in comparing KVA methodologies lies in their different parameterizations and underlying assumptions. Each approach was developed within distinct theoretical frameworks, leading to incompatible input requirements. This section establishes a mapping protocol to enable meaningful numerical comparison.

To enable direct comparison between semi-replication and the indifference pricing approaches, we establish a common parameter framework that reconciles the different theoretical foundations while preserving the economic intuition of each method.

Capital Financing Parameter

In the semi-replication framework, the parameter $\phi \in [0, 1]$ determines the proportion of regulatory capital that can be used for trade financing. For numerical consistency, we set:

$$\phi = 0.$$

This assumption reflects the most conservative regulatory interpretation where regulatory capital cannot be used for financing trades, ensuring that the full opportunity cost of capital is captured in the KVA calculation.

Cost of Capital Alignment

A critical alignment requirement is establishing consistency between the hurdle rate in the indifference approach and the capital cost parameter in semi-replication:

$$h = \gamma_K.$$

This ensures that both methodologies use the same benchmark for capital profitability. The hurdle rate h represents the minimum risk-adjusted return on capital that shareholders demand, while γ_K captures the opportunity cost of capital in the semi-replication framework.

Funding Spread

Throughout the numerical experiments we set the bank funding spread using a first-order credit-risk approximation. For a bank with risk-neutral default intensity λ_B and loss-given-default $1 - R_B$, the credit-risk adjusted funding spread is:

$$\lambda_{\text{cred}} = (1 - R_B) \lambda_B.$$

In a reduced-form Duffie-Singleton [24] framework, this represents the fair spread above the risk-free rate that compensates for expected credit losses. Specifically, for a defaultable zero-coupon bond with recovery rate R_B , the credit spread over the risk-free rate equals the product of default intensity and loss severity. For comparison, a zero-recovery bond ($R_B = 0$) would have a fair spread equal to the full default intensity λ_B , while our formulation accounts for partial recovery.

In practice, real unsecured funding spreads typically embed additional risk premia beyond pure credit risk, including:

- **Liquidity premium:** compensation for funding illiquidity
- **Regulatory premium:** costs associated with capital and liquidity buffers
- **Term structure effects:** maturity-dependent risk adjustments

When absolute accuracy is required, the total funding spread can be expressed as:

$$\lambda = (1 - R_B) \lambda_B + \lambda_{\text{liq}} + \lambda_{\text{reg}},$$

where λ_{liq} captures liquidity-related costs and λ_{reg} represents regulatory adjustments calibrated to the bank's funding curve.

Because the present work focuses on a like-for-like comparison of CVA methodologies, we set $\lambda_{\text{liq}} = \lambda_{\text{reg}} = 0$, ensuring that numerical differences between approaches are driven solely by the underlying modeling philosophy rather than by different spread assumptions. This provides a clean basis for evaluating the relative performance of different CVA calculation methods.

Risk Measure Standardization

Indifference approaches require risk quantification. We standardize the risk measure using:

$$\text{VaR}_{99\%} = m\sigma,$$

where $m = 2.33$ corresponds to the 99% confidence level under normal distribution assumptions. This parameter m appears in the RAROC denominator and ensures consistent risk-adjustment across methodologies.

It should be noted that this standardization applies only to the indifference methodologies and does not play any role in the semi-replication approach. Conversely, parameters such as R_c (recovery rate) and λ_C (counterparty intensity) are specific to the semi-replication framework and have no counterpart in the indifference setting. As a result, these elements cannot be fully aligned across the two approaches.

4.2 Hedgeability Cases for XVA

We now consider three distinct scenarios for the hedgeability of XVA risks, ranging from a fully unhedged world to a fully hedged one. These cases correspond to different assumptions about which XVA components (credit, debt, funding, capital adjustments) a bank can hedge or replicate using market instruments, and which must be treated as residual risks requiring capital and risk premia.

Case I: All XVAs Unhedgeable

In the most conservative scenario, none of the XVA components can be hedged. The bank cannot trade CDSs to offset counterparty risk, cannot arrange funding to neutralize funding spreads, and cannot transfer regulatory capital costs.

Here, all XVAs are treated as genuinely unhedgeable risks. The valuation framework must therefore be fully indifference-based (see Section 3.2), with the entire XVA vector forming an integrated risk complex.

This case is rarely encountered in practice, since at least partial hedging of credit and funding risks is typically feasible. For this reason, our numerical analysis will mainly focus on the other two scenarios (partial and full hedgeability), which are more relevant to practical applications.

Case II: Partially Hedgeable

This scenario reflects the most typical real-world situation. Credit and funding risks are hedgeable through market instruments, whereas capital risk (KVA) remains unhedgeable because there is no liquid market to transfer regulatory capital requirements. The hedgeability of DVA is somewhat controversial; however, whether DVA is treated as hedgeable or not does not affect our main argument.

We assume that CDS hedging is implemented by purchasing only the protection leg cashflow, and that the premium leg is either negligible or charged outside the hedge, e.g., via a separate XVA charge. Formally:

$$\textbf{Assumption A. } \Pi_t^{\text{prem}}(\text{CDS}) \approx 0 \text{ in PV accounting of the hedge.} \quad (4.2.1)$$

Consider a CDS on C with deterministic, $\mathcal{F}_t$ -measurable time-varying notional

$$N(u) = \mathbb{E}_t[V(u)^+], \quad t \leq u \leq T.$$

Under intensity pricing, the (discounted) protection leg PV at time t is

$$\Pi_t^{\text{prot}}(\text{CDS}) = \int_t^T \tilde{D}(t, u) (1 - R_C) \lambda_C(u) N(u) du = (1 - R_C) \int_t^T \lambda_C(u) \tilde{D}(t, u) \mathbb{E}_t[V(u)^+] du, \quad (4.2.2)$$

where

$$\tilde{D}(t, u) \equiv \exp\left(-\int_t^u (r(s) + \lambda_C(s) + \lambda_B(s)) ds\right), \quad t \leq u \leq T,$$

Comparing (4.2.2) with (3.1.7) yields the identity

$$\boxed{\Pi_t^{\text{prot}}(\text{CDS}) = -\text{CVA}_t.} \quad (4.2.3)$$

Therefore, under Assumption (4.2.1), the protection-only CDS hedge exactly offsets the CVA_t of the client trade in risk-neutral present value.

For FVA, let the treasury borrow unsecured cash in the amount of the expected positive exposure, i.e.

$$C_u^- = \mathbb{E}_t[V(u)^+], \quad t \leq u \leq T,$$

at an instantaneous funding spread $s_B(u)$ consistent with the bank's own intensity,

$$s_B(u) = LGD_B \lambda_B(u).$$

The discounted PV of the funding cashflows is then

$$\Pi_t^{\text{fund}} = - \int_t^T \tilde{D}(t, u) s_B(u) C_u^- du = - LGD_B \int_t^T \lambda_B(u) \tilde{D}(t, u) \mathbb{E}_t[V(u)^+] du. \quad (4.2.4)$$

Comparing (4.2.4) with (3.1.9) shows

$$\boxed{\Pi_t^{\text{fund}} = -\text{FCA}_t.} \quad (4.2.5)$$

Hence the funding desk exactly offsets the FCA in risk-neutral PV.

For DVA, it represents the benefit from the bank's own default risk. To hedge this, the bank should sell protection on itself. This can be done by buying bonds issued earlier. In reality this is very difficult. Most of the times instead of selling protection on oneself, one sells protection on a number of names that one thinks are highly correlated to oneself.

Let $P_t^{\text{rf}} = \mathbb{E}_t[\tilde{D}(t, T)Y_T]$ denote the clean (risk-free) price of the client payoff Y_T . We quote to the client

$$P_t = P_t^{\text{rf}} - (\text{CVA}_t + \text{FCA}_t + \text{DVA}_t + \text{KVA}_t). \quad (4.2.6)$$

Augmenting the position by the protection-only CDS hedge and the treasury funding hedge, and invoking (4.2.3)–(4.2.5) together with Assumption (4.2.1), we obtain

$$P_t + \Pi_t^{\text{prot}}(\text{CDS}) + \Pi_t^{\text{fund}} = P_t^{\text{rf}} + \text{DVA}_t + \text{KVA}_t. \quad (4.2.7)$$

If DVA is also treated as hedgeable (e.g., via a proxy), its hedge PV cancels DVA_t on the right-hand side as well.

Thus, within our protection-only convention, the hedgeable valuation adjustments (CVA and FCA) are fully neutralized by traded hedges in risk-neutral PV; only the non-hedgeable components (potentially DVA, KVA) remain.¹

Remark 4.2.1.

1. Equations (4.2.3) and (4.2.5) are equalities, not approximations. All approximation is confined to Assumption (4.2.1), which ignores the CDS premium leg in the hedge book (its recognition can be shifted to a separate XVA charge or deemed negligible in PV).
2. A pathwise (scenario-by-scenario) neutralization of the counterparty default jump would require a contingent CDS (CCDS) with notional $N(u) = V(u)^+$; the present analysis targets exact PV cancellation under intensities.

In an idealized frictionless market, the time-0 cost of establishing each hedge exactly equals the corresponding XVA received or paid; hence the hedged position delivers the clean payoff at maturity. Denoting by Y' the payoff of the hedged position (including XVAs and hedge cash flows), we obtain

$$\text{In the perfect-hedge limit, } Y' = Y,$$

so that the replication restores the clean trade.

Operationally, hedgeable components (CVA, FCA, and optionally DVA) are neutralized by traded hedges in PV, while the non-hedgeable component (KVA) is priced via indifference pricing. This separation is the key methodological consequence of the partial-hedge assumption.

¹If DVA and/or KVA are priced but not hedged, they persist as residual valuation items. If they are excluded from pricing (e.g., accounting purposes), (4.2.7) reduces to $P_t + \Pi_t^{\text{prot}} + \Pi_t^{\text{fund}} = P_t^{\text{rf}}$.

Case III: All XVAs Hedgeable

In this idealized scenario, we assume all XVA risks are hedgeable or fully transferable, including capital costs. This scenario provides a useful benchmark. Hedgeability of KVA may be interpreted as:

1. Residual risk is minimal due to sophisticated hedging strategies
2. Regulatory capital requirements are predictable and can be funded at known rates
3. Insurance-like contracts to cover regulatory capital charges.

This is the semi-replication (Section 3.1) case: all adjustments can be neutralized, and the clean payoff is perfectly replicated.

4.3 Detailed Implementation Strategies

Having defined the theoretical framework, products, models, and hedgeability assumptions, we now describe the numerical implementation strategy for computing XVAs and KVA in our examples. This includes the simulation of exposure profiles, calibration of key parameters for the indifference pricing method, and the calculation of valuation adjustments under both the semi-replication and indifference approaches. We maintain a consistent methodological framework across the cases to allow a clear comparison of results.

4.3.1 Simulation Framework for Exposure Calculation

To evaluate risk profiles and derivative prices numerically, we adopt a Monte Carlo simulation framework applicable to a wide class of stochastic models, including geometric Brownian motion for equity underlyings and Ornstein-Uhlenbeck processes (such as the Vasicek model for interest rates).

We simulate N independent scenarios for the underlying state variables X_t over $[0, T]$, using the Euler-Maruyama discretization for the chosen Itô process. For example, for the Vasicek model:

$$r_{t_{m+1}} = r_{t_m} + \kappa(\theta - r_{t_m}) \Delta t + \sigma \sqrt{\Delta t} Z_{m+1}, \quad Z_{m+1} \sim \mathcal{N}(0, 1).$$

At each time t_m , we compute the mark-to-market value $V^{(i)}(t_m)$ for scenario i . For path-dependent products, this involves a pricing subroutine; for European-style products, a closed-form conditional pricing formula may be used.

To evaluate risk profiles and derivative prices numerically, we form exposure statistics based on Section 3.3:

$$\begin{aligned} \text{EE}(t_m) &= \frac{1}{N} \sum_{i=1}^N (V^{(i)}(t_m))^+, & \text{ENE}(t_m) &= \frac{1}{N} \sum_{i=1}^N (-V^{(i)}(t_m))^+, \\ \text{PFE}_\alpha(t_m) &= \text{Quantile}_\alpha\left(\{(V^{(i)}(t_m))^+\}_{i=1}^N\right), & \text{EPE} &= \sum_{m=1}^M \frac{\Delta t_m}{\sum_k \Delta t_k} \text{EE}(t_m). \end{aligned}$$

These will feed into the computation of valuation adjustments and associated hedging strategies. The procedure is summarized in Algorithm 1.

The time step Δt is chosen to balance computational efficiency and numerical accuracy; convergence is checked by increasing N and refining the grid. This simulation framework provides a unified approach to modelling derivative cashflows under stochastic dynamics, forming the quantitative backbone for semi-replication and indifference pricing.

Algorithm 1: Monte Carlo for Exposure (EE/ENE/EPE/PFE)

Input:

- Time grid $0 = t_0 < \dots < t_M = T$ with steps Δt_m .
- Number of paths N .
- Model type $\in \{\text{GBM}, \text{Vasicek}, \text{Generic It}\hat{o}\}$ and parameters.
- Pricing routine for the portfolio (analytic or sub-algorithm).
- Quantile level α for PFE.

Output: $V^{(i)}(t_m)$, and profiles $\text{EE}(t_m)$, $\text{ENE}(t_m)$, $\text{PFE}_\alpha(t_m)$, EPE .

for $i \leftarrow 1$ **to** N **do**

```
Initialize state (e.g.,  $S_{t_0}^{(i)}$  or  $r_{t_0}^{(i)}$ );  
for  $m \leftarrow 0$  to  $M - 1$  do  
    Draw  $Z_{m+1} \sim \mathcal{N}(0, 1)$  // use vector  $Z$  + Cholesky for multi-factor  
    if  $model = \text{GBM}$  then  
        //  $dS_t = S_t((r - q)dt + \sigma dW_t)$   
         $S_{t_{m+1}}^{(i)} \leftarrow S_{t_m}^{(i)} + S_{t_m}^{(i)}((r - q)\Delta t_m + \sigma\sqrt{\Delta t_m} Z_{m+1});$   
    else  
        if  $model = \text{Vasicek}$  then  
            //  $dr_t = \kappa(\theta - r_t)dt + \sigma dW_t$   
             $r_{t_{m+1}}^{(i)} \leftarrow r_{t_m}^{(i)} + \kappa(\theta - r_{t_m}^{(i)})\Delta t_m + \sigma\sqrt{\Delta t_m} Z_{m+1};$   
        else  
            // Generic Itô:  $dX_t = \mu(X_t, t)dt + \sigma(X_t, t)dW_t$   
             $X_{t_{m+1}}^{(i)} \leftarrow X_{t_m}^{(i)} + \mu(X_{t_m}^{(i)}, t_m)\Delta t_m + \sigma(X_{t_m}^{(i)}, t_m)\sqrt{\Delta t_m} Z_{m+1};$   
  
    for  $m \leftarrow 1$  to  $M$  do  
        if path-dependent payoff then  
            Compute  $V^{(i)}(t_m)$  via pricing sub-algorithm (e.g., backward  
            induction/LSM);  
        else  
            Compute  $V^{(i)}(t_m)$  via analytic formula conditional on the simulated  
            state at  $t_m$ ;
```

for $m \leftarrow 1$ **to** M **do**

```
EE( $t_m$ )  $\leftarrow \frac{1}{N} \sum_{i=1}^N (V^{(i)}(t_m))^+;$   
ENE( $t_m$ )  $\leftarrow \frac{1}{N} \sum_{i=1}^N (-V^{(i)}(t_m))^+;$   
PFE $_\alpha$ ( $t_m$ )  $\leftarrow \text{Quantile}_\alpha(\{ (V^{(i)}(t_m))^+ \}_{i=1}^N);$ 
```

$\text{EPE} \leftarrow \sum_{m=1}^M \frac{\Delta t_m}{\sum_k \Delta t_k} \text{EE}(t_m);$

return $V^{(i)}(t_m)$, $\text{EE}(t_m)$, $\text{ENE}(t_m)$, $\text{PFE}_\alpha(t_m)$, EPE ;

4.3.2 Parameter Calibration under Indifference Pricing

To ensure consistency with the bank's target hurdle rate h under the indifference-pricing framework, we use a two-stage calibration of the terminal covariance matrix $\mathbf{A}$ and the risk-aversion parameter θ . Stage 1 rescales a baseline semi positive definite covariance $\mathbf{A}_0$ to enforce the hurdle-rate constraint in (3.2.4), namely $\mu^\top \mathbf{A}^{-1} \mu = (h - r - \lambda_{\text{total}})^2 \nu^2$, which aligns the marginal compensation for risk with the target return; concretely, set $\mathbf{A} := c^* \mathbf{A}_0$ with $c^* = (\mu^\top \mathbf{A}_0^{-1} \mu) / [(h - r - \lambda_{\text{total}})^2 \nu^2]$. Stage 2 then tunes θ by solving the one-dimensional root condition $\text{RAROC}(\theta) = h$ implied by the indifference price, so that pricing outcomes are consistent with the bank's capital-allocation objectives.

The full procedure is summarized in Algorithm 2, which calls the subroutine Algorithm makeSPD 3 to enforce positive definiteness of the covariance matrix.

Constraint (3.2.4) links excess returns, covariance, and risk capacity so that the marginal compensation for risk is consistent with the hurdle rate. Rescaling a baseline $\mathbf{A}_0$ by a scalar c^* preserves positive definiteness and the correlation structure while enforcing the constraint exactly.

Feasibility requires $\mu \neq 0$ and $g := h - r - \lambda_{\text{total}} \neq 0$. if $g = 0$ the constraint degenerates and Stage 1 becomes immaterial. The mapping $\theta \mapsto \text{RAROC}(\theta)$ is in general nonlinear and not guaranteed to be monotone (numerator affine, denominator $\sqrt{\text{quadratic}}$). Therefore we solve $\text{RAROC}(\theta) = h$ by a bracketed one-dimensional root finder (bisection/secant/Newton with safeguards); if multiple roots arise, we select the smallest θ achieving h (minimal-risk solution).

Algorithm 3: MakeSPD: minimal SPD repair (inline)

Input: $\mathbf{B}$ (symmetric or near-symmetric)

Output: $\tilde{\mathbf{B}}$ (SPD, well-conditioned)

1. Symmetrize: $\mathbf{B} \leftarrow (\mathbf{B} + \mathbf{B}^\top)/2$.
 2. Eigendecompose: $\mathbf{B} = \mathbf{Q} \text{diag}(\lambda_i) \mathbf{Q}^\top$.
 3. Clip eigenvalues: $\lambda'_i \leftarrow \max(\lambda_i, \delta)$ with tiny $\delta > 0$.
 4. Optional shrinkage: $\mathbf{B}' \leftarrow (1 - \rho) \mathbf{Q} \text{diag}(\lambda'_i) \mathbf{Q}^\top + \rho \tau \mathbf{I}$
// $\rho \in [0, 1)$, $\tau = \frac{1}{d} \text{tr}(\mathbf{B})$; **stabilizes small-sample noise**.
 5. Condition cap: if $\kappa(\mathbf{B}') > \kappa_{\text{max}}$, increase ρ until $\kappa \leq \kappa_{\text{max}}$.
 6. Return $\tilde{\mathbf{B}} \leftarrow \mathbf{B}'$.
-

4.3.3 Netting Implementation Strategy

The numerical implementation of netting effects necessitates a comprehensive aggregation framework operating at the portfolio level to capture diversification benefits inherent in legally enforceable netting agreements.

For a portfolio of N options, the pre-netting Expected Exposure (EE) is computed as the sum of individual option EEs, considering only long positions:

$$\text{EE}_{\text{pre}}(t) = \sum_{i=1}^N \mathbf{1}_{\{\text{position}_i > 0\}} \cdot \text{notional}_i \cdot \mathbb{E}^{\mathbb{Q}} [\max(V_i(t), 0)]$$

where $V_i(t)$ represents the mark-to-market value of option i at time t under the risk-neutral measure $\mathbb{Q}$.

Algorithm 2: Two-Stage Calibration under Indifference Pricing

Input:

- Target hurdle rate h , risk-free rate r , total funding spread λ_{total} ;
set $g := h - r - \lambda_{\text{total}}$.
- Excess expected return vector $\boldsymbol{\mu} = \mathbb{E}[\mathbf{S}_T] - \mathbf{S}_0(1 + r + \lambda_{\text{total}})$.
- Baseline (near-)SPD covariance $\mathbf{A}_0$ for $\mathbf{S}_T$ (empirical or model-implied).
- Risk-capacity parameter $\nu > 0$; tolerance $\varepsilon > 0$.

Output: Covariance matrix $\mathbf{A}$; risk-aversion parameter θ^* .

Stage 1: Covariance rescaling (constraint enforcement)

1. $k \leftarrow 0$, set $\mathbf{A}^{(0)} \leftarrow \mathbf{A}_0$.
2. **while** $|\boldsymbol{\mu}^\top (\mathbf{A}^{(k)})^{-1} \boldsymbol{\mu} - (g\nu)^2| > \varepsilon$ **do**

- **SPD repair/shrinkage:** $\mathbf{A}^{(k)} \leftarrow \text{MakeSPD}(\mathbf{A}^{(k)})$;
// Empirical $\mathbf{A}_0$ may be only psd / ill-conditioned due to sampling error, missing data, collinearity, or rounding; inversion and Cholesky require SPD.
 - Compute $q^{(k)} := \boldsymbol{\mu}^\top (\mathbf{A}^{(k)})^{-1} \boldsymbol{\mu}$.
 - Set $c^{(k+1)} := \frac{q^{(k)}}{(g\nu)^2}$ and update
$$\mathbf{A}^{(k+1)} \leftarrow c^{(k+1)} \mathbf{A}_0.$$
 - // Scaling enforces $\boldsymbol{\mu}^\top \mathbf{A}^{-1} \boldsymbol{\mu} = (g\nu)^2$ while preserving correlations.
 - $k \leftarrow k + 1$.
3. **Output** $\mathbf{A} \leftarrow \mathbf{A}^{(k)}$.

Stage 2: Risk-aversion calibration (RAROC matching)

1. Initialize $\theta^{(0)}$; set $k \leftarrow 0$.
 2. **while** $|\text{RAROC}(\theta^{(k)}) - h| > \varepsilon$ **do**

- Compute indifference price at $\theta^{(k)}$ and evaluate $\text{RAROC}(\theta^{(k)})$.
 - Update θ via a 1D root-finder for $f(\theta) = \text{RAROC}(\theta) - h$ (bisection/secant/Newton with safeguards).
 - $k \leftarrow k + 1$.
 3. **Output** $\theta^* := \theta^{(k)}$ once $|\text{RAROC}(\theta^{(k)}) - h| \leq \varepsilon$.
-

The post-netting exposure accounts for portfolio-level offsetting effects. For each simulation path ω and time t , we first compute the net portfolio value:

$$V_{\text{net}}(\omega, t) = \sum_{i=1}^N \text{position}_i \cdot \text{notional}_i \cdot V_i(\omega, t)$$

The post-netting EE is then:

$$\text{EE}_{\text{post}}(t) = \mathbb{E}^{\mathbb{Q}} [\max(V_{\text{net}}(\omega, t), 0)] = \frac{1}{M} \sum_{j=1}^M \max \left(\sum_{i=1}^N \text{position}_i \cdot \text{notional}_i \cdot V_i(\omega_j, t), 0 \right)$$

where M denotes the number of Monte Carlo paths.

4.3.4 Wrong Way Risk Implementation Strategy

We incorporate wrong way risk (WWR) by coupling the underlying asset with bilateral default intensities. The stock follows GBM and the (nonnegative) default intensities follow correlated CIR processes:

$$\begin{aligned} dS_t &= \mu S_t dt + \sigma S_t dW_t^S, \\ d\lambda_t^B &= \kappa_B(\theta_B - \lambda_t^B) dt + \sigma_B \sqrt{\lambda_t^B} dW_t^B, \\ d\lambda_t^C &= \kappa_C(\theta_C - \lambda_t^C) dt + \sigma_C \sqrt{\lambda_t^C} dW_t^C, \end{aligned}$$

where B denotes the bank and C the counterparty.

We simulate wrong-way risk by correlating the exposure and credit processes, a positive correlation $\rho > 0$ means high exposure scenarios coincide with high default intensity, thus exacerbating CVA. So the instantaneous correlation matrix of (W^S, W^B, W^C) would be

$$\Sigma = \begin{pmatrix} 1 & \rho_{SB} & \rho_{SC} \\ \rho_{SB} & 1 & \rho_{BC} \\ \rho_{SC} & \rho_{BC} & 1 \end{pmatrix}.$$

To ensure numerical stability we repair Σ to be positive semi definite by eigenvalue clipping:

$$\Sigma = \mathbf{Q} \Lambda \mathbf{Q}^\top, \quad \tilde{\Lambda}_{ii} = \max(\Lambda_{ii}, \epsilon), \quad \epsilon = 10^{-8}, \quad \tilde{\Sigma} = \mathbf{Q} \tilde{\Lambda} \mathbf{Q}^\top,$$

then take a Cholesky factor $\mathbf{L} = \text{chol}(\tilde{\Sigma})$ and set $(dW_t^S, dW_t^B, dW_t^C)^\top = \mathbf{L} (dZ_t^1, dZ_t^2, dZ_t^3)^\top$, with dZ_t^i independent standard Brownian increments.

We use Euler–Maruyama. For CIR we employ full truncation to preserve nonnegativity:

$$\begin{aligned} S_{t+\Delta} &= S_t \exp \left(\left(\mu - \frac{1}{2} \sigma^2 \right) \Delta + \sigma \sqrt{\Delta} Z^S \right), \\ \lambda_{t+\Delta}^B &= \max(0, \lambda_t^B + \kappa_B(\theta_B - \max\{\lambda_t^B, 0\})\Delta + \sigma_B \sqrt{\max\{\lambda_t^B, \epsilon\}} \sqrt{\Delta} Z^B), \\ \lambda_{t+\Delta}^C &= \max(0, \lambda_t^C + \kappa_C(\theta_C - \max\{\lambda_t^C, 0\})\Delta + \sigma_C \sqrt{\max\{\lambda_t^C, \epsilon\}} \sqrt{\Delta} Z^C). \end{aligned}$$

Define the survival–discount kernel along a path ω by

$$\tilde{D}(0, t; \omega) = \exp \left(- \int_0^t \{r(u) + \lambda_B(u, \omega) + \lambda_C(u, \omega)\} du \right).$$

Let $V(t, \omega)$ be the (clean) netting-set mark-to-market on the path. The semi-replication Feynman-Kac representations with WWR are

$$\begin{aligned} \text{CVA}_{\text{WWR}} &= -(1 - R_C) \mathbb{E} \left[\int_0^T \lambda_C(t) \tilde{D}(0, t) (V(t))^+ dt \right], \\ \text{DVA}_{\text{WWR}} &= -(1 - R_B) \mathbb{E} \left[\int_0^T \lambda_B(t) \tilde{D}(0, t) (-V(t))^+ dt \right], \\ \text{FCA}_{\text{WWR}} &= -\mathbb{E} \left[\int_0^T s_B(t) \tilde{D}(0, t) (V(t))^+ dt \right], \end{aligned}$$

where $s_B(t) = (1 - R_B)\lambda_B(t)$ is the banks unsecured funding spread.

On a grid $\{t_k\}_{k=1}^n$ with steps Δt , define

$$\text{EE}_{\text{WWR}}(t_k) = \frac{1}{M} \sum_{m=1}^M (V^{(m)}(t_k))^+.$$

Using a proportional capital proxy $K^{\text{WWR}}(t_k) = \alpha \text{EE}_{\text{WWR}}(t_k)$, we approximate

$$\text{KVA}_{\text{WWR}} \approx \sum_{k=1}^n \bar{\tilde{D}}(0, t_k) (h - r(t_k)) \alpha \text{EE}_{\text{WWR}}(t_k) \Delta t, \quad \bar{\tilde{D}}(0, t_k) = \frac{1}{M} \sum_{m=1}^M \tilde{D}^{(m)}(0, t_k).$$

We also have discrete estimators for CVA, DVA and FVA as:

$$\begin{aligned} \text{CVA}_{\text{WWR}} &\approx -(1 - R_C) \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^n \lambda_C^{(m)}(t_k) \tilde{D}^{(m)}(0, t_k) (V^{(m)}(t_k))^+ \Delta t, \\ \text{DVA}_{\text{WWR}} &\approx -(1 - R_B) \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^n \lambda_B^{(m)}(t_k) \tilde{D}^{(m)}(0, t_k) (-V^{(m)}(t_k))^+ \Delta t, \\ \text{FCA}_{\text{WWR}} &\approx -(1 - R_B) \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^n \lambda_B^{(m)}(t_k) \tilde{D}^{(m)}(0, t_k) (V^{(m)}(t_k))^+ \Delta t. \end{aligned}$$

The full procedure is summarized in Algorithm 4

Remark 4.3.1.

1. Under semi-replication, XVA admits a Feynman-Kac form with intensity weights and the survival-discount kernel $\tilde{D}(0, t) = \exp(-\int_0^t (r + \lambda_B + \lambda_C) du)$. Hence Monte Carlo estimators use $\lambda(t_k) \tilde{D}(0, t_k) \Delta t$ rather than sampling default times. In the $\Delta t \rightarrow 0$ limit this is equivalent to indicator-based estimators, e.g.,

$$\mathbb{E}[D(0, t_k) \mathbf{1}_{\{\tau_C \in (t_{k-1}, t_k]\}} V^+(t_k)] \approx \mathbb{E}[\lambda_C(t_k) \tilde{D}(0, t_k) V^+(t_k) \Delta t],$$

while avoiding rare-event variance, step-size bias, and book-keeping of first jumps.

2. Many regulatory/IMM implementations use EEPE (Effective EPE): first form the non-decreasing $\text{EE}^{\text{eff}}(t_k) = \max(\text{EE}^{\text{eff}}(t_{k-1}), \text{EE}(t_k))$ and then time-average it (often over a capped horizon, with supervisory multipliers), so capital uses $K(t_k) = \alpha \text{EE}^{\text{eff}}(t_k)$. In this thesis we focus on the main idea and adopt the simpler EPE-based proxy $K(t_k) = \alpha \text{EE}(t_k)$ in KVA. If a stricter EEPE treatment is desired, one can replace $\text{EE}(t_k)$ by $\text{EE}^{\text{eff}}(t_k)$ in the KVA sum without changing the rest of the algorithm.

Algorithm 4: WWR XVA Algorithm

Input: Correlation matrix Σ , number of paths M , steps n , step Δt , GBM/CIR parameters, R_B, R_C , rate $r(t_k)$, capital multiplier α , hurdle h .

Output: CVA_{WWR} , DVA_{WWR} , FCA_{WWR} , KVA_{WWR} , and EPE_{WWR} .

// SPD repair is necessary: empirical Σ can be only psd / ill-conditioned; Cholesky and the kernel $\tilde{D}$ require SPD for stability.

Compute SPD-repaired $\tilde{\Sigma}$ by eigenvalue clipping; set $L = \text{chol}(\tilde{\Sigma})$;

Initialize accumulators: $CVA \leftarrow 0$, $DVA \leftarrow 0$, $FCA \leftarrow 0$;

Initialize arrays: $\text{avgD}[1..n] \leftarrow 0$, $\text{avgVp}[1..n] \leftarrow 0$;

for $m = 1$ **to** M **do**

Initialize $S \leftarrow S_0$, $\lambda_B \leftarrow \lambda_0^B$, $\lambda_C \leftarrow \lambda_0^C$, and $\tilde{D} \leftarrow 1$;

for $k = 1$ **to** n **do**

Draw $Z = (Z_1, Z_2, Z_3)^\top \sim \mathcal{N}(0, I_3)$; set $(dW^S, dW^B, dW^C)^\top = LZ \sqrt{\Delta t}$;

// Euler / full truncation updates

$S \leftarrow S \cdot \exp((\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma dW^S)$;

$\lambda_B \leftarrow \max\{0, \lambda_B + \kappa_B(\theta_B - \max\{\lambda_B, 0\})\Delta t + \sigma_B \sqrt{\max\{\lambda_B, \epsilon\}} dW^B\}$;

$\lambda_C \leftarrow \max\{0, \lambda_C + \kappa_C(\theta_C - \max\{\lambda_C, 0\})\Delta t + \sigma_C \sqrt{\max\{\lambda_C, \epsilon\}} dW^C\}$;

// update survival--discount kernel

$\tilde{D} \leftarrow \tilde{D} \cdot \exp(-(r(t_k) + \lambda_B + \lambda_C)\Delta t)$;

// pathwise mark-to-market and positive/negative parts

$V \leftarrow \text{PricePortfolio}(S, t_k)$; $V^+ \leftarrow \max(V, 0)$;

$V^- \leftarrow \max(-V, 0)$;

// semi-replication quadrature for CVA/DVA/FCA

$CVA += (1 - R_C) \lambda_C \tilde{D} V^+ \Delta t$;

$DVA += (1 - R_B) \lambda_B \tilde{D} V^- \Delta t$;

$FCA += (1 - R_B) \lambda_B \tilde{D} V^+ \Delta t$;

// store for EPE/KVA

$\text{avgD}[k] += \tilde{D}$; $\text{avgVp}[k] += V^+$;

end

end

// Aggregate across paths

for $k = 1$ **to** n **do**

$\tilde{D}(0, t_k) \leftarrow \text{avgD}[k]/M$; $EE_{WWR}(t_k) \leftarrow \text{avgVp}[k]/M$;

end

Compute $EPE_{WWR} = \sum_{k=1}^n \frac{\Delta t}{\sum_{j=1}^n \Delta t} EE_{WWR}(t_k)$;

Compute $KVA_{WWR} = -\sum_{k=1}^n \tilde{D}(0, t_k) (h - r(t_k)) \alpha EE_{WWR}(t_k) \Delta t$;

return $CVA_{WWR} = -CVA/M$, $DVA_{WWR} = -DVA/M$, $FCA_{WWR} = -FCA/M$, KVA_{WWR} , EPE_{WWR} ;

Chapter 5

Numerical Experiments Results

In this chapter, we report numerical results for the European call option and the interest rate swap. While both products are used to benchmark the KVA methodologies, specific features such as netting effects and wrong-way risk are illustrated through the call option example. All computations follow the implementation framework outlined in Chapter 4.

5.1 XVA Baseline Calculation

We begin with baseline pricing results for the standalone European call option and the interest rate swap, under each methodology. Key model parameters are summarized in Table 5.1. The Black-Scholes price serves as a reference ‘clean’ valuation without credit, funding and capital effects. We then compute valuations with XVA adjustments under the two methodologies:

- **Semi-replication approach:** treating XVA costs (including capital) as hedgeable that can be priced via replication in a risk-neutral framework.
- **Indifference pricing approach:** treating capital as unhedgeable risks that require an extra charge to satisfy the bank’s internal return targets. We consider both the RAROC-based and the optimization-based formulations of indifference pricing, though in practice these yielded very similar numerical results for our cases (reflecting that both formulations preserve the bank’s hurdle rate in pricing).




Figure 5.1: Expected Exposure (EE) profiles for Call Option and Interest Rate Swap.

Before turning to the valuation adjustments themselves, we first examine the expected exposure (EE) profiles of the two test products.

Table 5.1: Model Parameters for Call Option and IRS Experiments

Parameter	Value	Which Model
Call Option Market Parameters		
Initial Stock Price (S_0)	100	All
Strike Price (K)	100	All
Risk-free Rate (r)	0.05	All
Volatility (σ)	0.5	All
$\mathbb{P}$ drift μ_P	0.06	Indifference
Maturity (T)	1.0	All
IRS Parameters		
Initial short rate (r_0)	0.040	All
Mean reversion speed under $\mathbb{P}$ (κ_P)	0.50	Indifference
Long-term mean under $\mathbb{P}$ (θ_P)	0.060	Indifference
Mean reversion speed under $\mathbb{Q}$ (κ_Q)	0.50	All
Long-term mean under $\mathbb{Q}$ (θ_Q)	0.056	All
Volatility of short rate (σ)	0.020	All
Fixed rate of IRS (K)	0.050	All
Notional (N)	1000	All
Payment interval / Tenor (τ)	0.25	All
Swap maturity (T)	5.0	All
Default and Recovery Parameters		
Bank Default Intensity (λ_b)	0.04	All (indiff through λ)
Counterparty Default Intensity (λ_c)	0.08	Semi-replication
Bank Recovery Rate (R_b)	0.4	All (indiff through λ)
Counterparty Recovery Rate (R_c)	0.4	Semi-replication
Borrowing Spread (λ)	0.016	Indifference (via λ_b, R_b)
Capital and Regulatory Parameters		
Proportion of Capital Usable (ϕ)	0	Semi-replication (set 0)
Hurdle Rate / Target RAROC (h/γ_k)	0.1	All
Initial Capital (C_0)	1000	Indifference
Proportion of VaR and std (m)	2.33	Indifference
Regulatory Capital Constant (α)	1.4	Semi-replication
Constant in Constraint of θ (ν)	2.33	Optimization-based Indifference
Indifference Model Parameters		
Initial Stock $S_{\text{indiff}0}$	[100, 100, 100]	Indifference
$\mathbb{P}$ drift μ_{indiff}	[0.06, 0.07, 0.065]	Indifference

Figure 5.1 highlights two distinct dynamics. For the call option, convexity of the payoff dominates time decay: as maturity approaches, the distribution of S_t widens, so $\text{EE}(t)$ rises monotonically. This leads to persistent exposures and sustained capital charges throughout the contract.

By contrast, the swap profile is hump-shaped. Early on, rate uncertainty grows and exposures increase; around mid-life, the sensitivity of the remaining cashflows to interest rates is largest, so exposures peak. As payments are gradually exchanged, exposures decline toward zero. Hence, swaps concentrate capital requirements in the middle of their life, whereas options impose long-term, steadily growing consumption.

5.1.1 European Call Option Case

Standing as a bank that long a call option, the results are summarized in Table Call option. The Black-Scholes fair value of the option is about 21.79, which is the market price the

Table 5.2: Call Option BS price, capital requirement, and XVA adjustments

Metric		Semi-replication	Indiff (RAROC)	Indiff (Opt.)
Black-Scholes call price		-21.7926	-21.7926	-21.7926
Capital requirement		-20.2275	-20.2275	-20.2275
XVAs	Total XVA	-2.6163	-0.3813	-0.3900
	CVA -0.4866	FVA -0.2163		
	DVA 0.0000	KVA -1.9134		
Adjusted Price		-19.1763	-22.1739	-22.1826

bank would pay to enter the trade absent any XVA. The semi-replication method produces an adjusted valuation of roughly 19.18, which is about 12% lower than the Black-Scholes price. In other words, from the bank’s perspective the trade is worth less once credit, funding, and capital costs are considered. This implies a net negative XVA adjustment to the clean price. The largest contributor to this downward adjustment is the capital valuation adjustment (KVA), on the order of 1.9 in absolute value, with smaller deductions from credit valuation adjustment (CVA) and funding valuation adjustment (FVA). There is no debit valuation adjustment (DVA) in this case since the position is an asset to the bank, the bank’s own default would not relieve any obligation.

The economic interpretation is clear: under replication, capital and other XVA costs act as a drag on the trade’s value, so the bank would only be willing to do this trade at a discounted price. In effect, the bank demands compensation upfront for the anticipated costs—it would pay about 19.18 instead of 21.79 to enter the option. This lower price can be viewed as a form of capital relief: by paying less initially, the bank sets aside the saved amount to cover future capital and funding costs. The semi-replication framework thus treats the required regulatory capital as an opportunity cost that reduces the option’s value. The bank is only willing to pay the full market price if the counterparty covers those XVA costs separately. If forced to pay the full 21.79 without adjustment, the bank’s shareholders would effectively be subsidizing the trade’s credit and capital requirements—hence the trade would not meet the bank’s acceptance criteria in a hedgeable-cost worldview.

By contrast, the indifference pricing approach yields a valuation for the same call option that is slightly above the Black-Scholes price. In our results, the RAROC-based indifference price is about 22.17 and the optimization-based indifference price about 22.18. This corresponds to a small positive total XVA adjustment added on top of the BS price.

This upward adjustment should not be read as a “capital surcharge” in the same sense as semi-replication. Instead, it reflects that from a portfolio perspective the additional trade is actually advantageous. The higher indifference valuation means the bank is effectively willing to pay more to enter the trade, since doing so improves the overall RAROC or the portfolio hurdle rate. In other words, the indifference approach embeds the economic logic that a trade which enhances portfolio-level performance has a positive value beyond its risk-neutral price.

However, practice diverges from theory in how dealers implement this adjustment. In many dealer-corporate relationships, banks are asymmetric: they will add a KVA charge when a transaction worsens their RAROC, but they will not pass on a premium to the counterparty when a transaction improves their RAROC. In such cases, the indifference price serves more as an internal decision tool than as a symmetric quotation mechanism.

By contrast, in interbank markets where pricing symmetry is expected, the adjustment is applied in both directions, so a favorable RAROC effect would indeed be reflected in the price offered.

This asymmetry highlights a broader tension: while indifference pricing provides a consistent theoretical framework grounded in shareholder return and hurdle preservation, its practical application depends on the counterparty type and competitive dynamics. For the long call option, this means the indifference framework conceptually justifies a price premium relative to Black-Scholes, but whether that premium is actually “paid” to the counterparty depends on market convention.

5.1.2 Interest Rate Swap Case

We next examine a 5-year payer interest rate swap (IRS), where the bank pays floating to the client and receives fixing. We chose the fixed rate such that the swap has a near-zero fair value at inception.

Figure 5.2 illustrates simulated Vasicek short-rate paths used in the pricing/exposure engine. The paths display clear mean reversion toward the long-run level θ , while cross-sectional dispersion grows with horizon. Several scenarios spend time near the zero/low-rate region, whereas others drift materially higher; the mean path stays close to the long-run mean. Because we receive fixed and pay floating, low-rate scenarios make the fixed leg relatively more valuable (positive exposure to the bank), while high-rate scenarios generate negative exposure. This two-sided, long-horizon exposure with a noticeably heavy positive tail (driven by low-rate scenarios) magnifies CVA and the associated capital metrics relative to the short-dated option.

In our calibration it’s actually slightly positive, about 0.889 to the bank—meaning if we ignore XVA, the swap is almost a breakeven trade. Despite this tiny baseline value, the XVA impacts for the swap are proportionally much larger than for the short-term option, because the swap’s longer maturity and ongoing exposure profile introduce significant credit and capital costs. Table 5.3 summarizes the results.




Figure 5.2: Simulated Vasicek Short Rate Paths Illustration

Under the semi-replication approach, the total XVA adjustment for the swap is about -4.39 . In other words, entering this swap is expected to impose a considerable cost on the bank. Summing up the components, the swap suffers a substantial CVA charge (around -1.41 , reflecting the counterparty’s default risk over 5 years) and an FVA charge

Table 5.3: Interest Rate Swap BS price, capital requirement, and XVA adjustments

Metric		Semi-replication	Indiff (RAROC)	Indiff (Opt.)
Black-Scholes IRS price		0.8892	0.8892	0.8892
Capital requirement		7.7682	7.7682	7.7682
XVAs	Total XVA	-4.3937	-12.2781	-12.5534
	CVA -1.4098	FVA -0.5708		
	DVA 1.0520	KVA -3.4651		
Adjusted Price		5.2830	13.1674	12.5141

(around -0.57 for funding the uncollateralized exposure). Unlike the call option, here there is a meaningful DVA benefit ($+1.05$) because the swap can be a liability to the bank in scenarios where interest rates fall—if the bank were to default, it might be released from making some of those costly fixed payments, which is a windfall to shareholders in an extreme event. However, this DVA gain is far outweighed by a very large KVA (approximately -3.47) reflecting the capital the bank must hold against a long-maturity OTC contract. The KVA is the dominant term, making the overall XVA strongly negative. To make the trade acceptable, the bank would require compensation for this—effectively charging an additional amount equal to the XVA cost. Indeed, the adjusted semi-replication price comes out to about 5.28. This means the bank would ask the counterparty for roughly 5.28 on a trade that otherwise would be worth only 0.889 in a risk-neutral world.

The indifference pricing approach for the swap yields an even more dramatic price increase. To satisfy a 10% hurdle rate on a low-margin, risk-heavy trade, the indifference price comes out to roughly 12.5 or 13.2 depending on the method. This is an order of magnitude above the clean value. Put differently, the indifference framework implies a total XVA adjustment of $+12.3$ or more on the swap—a huge premium. This reflects the fact that from a risk-adjusted return perspective, the swap is very capital-inefficient: it ties up a lot of capital while offering almost no baseline profit. To maintain the bank’s ROE or RAROC, the price must be hiked significantly. The swap results underscore that when capital is treated as an unhedgeable constraint, the required charge can be extremely high for trades with long tenors and heavy exposures. Compared with semi-replication approach, the indifference approach is more conservative: it demands a higher compensation for the same trade. This wide gap highlights the practical impact of the hedgeability assumption: a fully hedgable-cost view (semi replication) versus an unhedgeable-capital view (indifference) can lead to very different pricing of the same risk, especially for trades that heavily engage regulatory capital.

5.2 Netting Result Analysis

In practice, trades are often part of a portfolio with netting agreements, rather than stand-alone transactions. Netting allows the bank and counterparty to offset exposures of multiple trades: positive exposure on one trade can be reduced or canceled by negative exposure on another in the event of default. This can significantly alter the XVA profile because credit exposure and capital requirements are computed on a net portfolio basis under such agreements. We therefore investigated how netting affects valuations under each methodology by looking at some simple portfolio scenarios. Rather than exhaustively simulating many portfolios, we chose three representative combinations of calls for our portfolio:

1. a “long” portfolio where all positions are long calls with strikes randomly near K

($\pm 10\%$).

2. a “mixed” portfolio with alternating long/short calls and strikes randomized a bit wider ($\pm 15\%$).
3. a “hedged” portfolio with one long lower-strike calls ($0.9K$) and one short higher-strike calls ($1.1K$).

For each portfolio, we computed the total XVA-adjusted prices in two ways: pre-netting (summing the XVA results of each trade computed independently, as if no netting agreement) and post-netting (computing exposure and XVA on the combined portfolio as one netting set). Table 5.4 and 5.5 summarizes the results for each method.

Table 5.4: Netting Results

Portfolio	BS Sum	Semi-Replication			Indifference		
		Pre	Post	Benefit	Pre	Post	Benefit
long	−67.2409	−58.7466	−58.7091	0.0375	−67.3355	−68.9235	−1.5880
mixed	−5.2114	−0.6006	−2.2241	−1.6235	−6.2045	−5.7918	0.4127
hedged	10.1155	14.0706	12.7512	−1.3194	8.6002	10.4161	1.8159

Table 5.5: Semi-Replication XVA components: pre/post netting and reductions

Portfolio	Pre-netting					Post-netting					Reduction (Pre – Post)			
	CVA	DVA	FVA	KVA	Total	CVA	DVA	FVA	KVA	Total	CVA	CVA %	Total	Total %
long	−1.5819	0.0	−0.7031	−6.2094	−8.4943	−1.5888	0.0	−0.7062	−6.2367	−8.5317	0.0070	−0.44	0.0374	−0.44
mixed	−0.8582	0.0	−0.3814	−3.3712	−4.6108	−0.5548	0.0	−0.2466	−2.1858	−2.9873	−0.3034	35.35	−1.6235	35.21
hedged	−0.7384	0.0	−0.3282	−2.8885	−3.9551	−0.4924	0.0	−0.2189	−1.9244	−2.6357	−0.2459	33.31	−1.3194	33.36

Note: Percentage columns use $(\text{Post} - \text{Pre})/|\text{Pre}| \times 100\%$

Tables 5.4 and 5.5 show how legal netting alters valuations under two distinct logics, semi-replication, which prices hedgeable adjustments from exposures, and indifference pricing, which sets a portfolio-level premium so that the book meets the hurdle h . While both recognise the same netting set, they aggregate risk through different channels and therefore respond differently to portfolio composition.

Under semi-replication each component is a pathwise integral of the positive or negative part of the portfolio mark-to-market, weighed by an intensity-discount kernel. Hence the netting effect is governed by how much netting reduces the positive part of the combined V at each date.

This mechanism matches the patterns in Table 5.5. When the portfolio consists entirely of long call positions, all exposures move in the same direction. In such a setting there is virtually no scope for offsetting or diversification, and therefore the combined exposure profile is essentially the sum of the individual trades. The values are almost identical, the difference is within Monte Carlo and time-discretisation noise.

By contrast, the mixed and hedged portfolios generate material internal offsets, shrinking V^+ at the netting-set level and with it the adjustments. These cases, by design, produces strong internal offsetting, and the net exposure after netting is small. The semi-replication framework captures this in a straightforward way: CVA, FVA, and KVA each decline materially once netting is applied. This outcome is fully consistent with expectations, since netting directly lowers expected exposure to the counterparty and reduces the capital required to support the trade.

Indifference prices are fixed at the portfolio level: the premium is whatever is required so that the combined book (existing positions plus the candidate trades) attains the target

hurdle rate h . Capital enters through a tail-risk measure of aggregate P&L, not through additive tradewise exposures, so aggregation is intrinsically nonlinear.

This portfolio logic explains the entries in Table 5.4. In the all-long case, pooling same-direction trades concentrates losses, raises economic capital, and yields a combined price that is worse than the sum of stand-alone quotes. This is not a failure of netting but a reflection that co-movement increases portfolio tail risk and therefore the capital charge required to keep the hurdle. In the mixed and hedged cases, moving offsetting legs into the same netting set removes redundancy in capital compensation that would exist if each leg were priced in isolation: required capital contracts at the portfolio level and the indifference value improves. The magnitude of the improvement is larger than the semi-replication relief because the indifference premium re-optimizes the capital charge for the package, not just the exposure integral.

Semi-replication answers: “what is the PV of hedgeable adjustments given the portfolio exposure profile?”; the response to netting is therefore near-linear and exposure-proportional. Indifference pricing answers: “should the bank accept this package at the hurdle?”; it embeds a portfolio capital constraint and is therefore more sensitive to diversification and the elimination of double-counted capital charges. In aligned portfolios (all-long) the two logics coincide in predicting that netting does little; in offsetting portfolios (mixed/hedged) both signal sizeable relief, with indifference amplifying the effect because the capital charge is recomputed for the combined book.

5.3 Wrong Way Risk Analysis

Wrong-way risk (WWR) arises when the counterparty’s likelihood of default is positively linked with the bank’s exposure to that counterparty. In such cases, adverse scenarios for exposure coincide with deteriorating credit quality, amplifying expected losses and capital charges. Figures 5.3 provide simulated sample paths for the stochastic intensities of bank and counterparty defaults, while Figure 5.4 depicts the joint correlation structure between stock prices and default intensities at maturity. Together, these plots illustrate how market dynamics and credit drivers may interact in shaping WWR.

As shown in Figure 5.3, the bank intensity paths exhibit moderate fluctuations around a mean of 2 – 3%, with upward shifts in some scenarios reflecting rising bank-specific default risk. Counterparty intensity paths display a wider dispersion, with peaks as high as 6 – 7%, consistent with higher intrinsic credit volatility of counterparties compared to large banks. The 3D scatter in Figure 5.4 further confirms that at maturity, states of lower stock prices are associated with higher counterparty default intensities (λ_C), while the bank’s own intensity (λ_B) shows weaker dependence. This joint distribution underpins the WWR adjustments observed in Table 5.6.




Figure 5.3: Bank and Counterparty Default Intensity Sample Paths




Figure 5.4: Joint Correlation Between Stock Prices and Default Intensities

Table 5.6: Wrong-way and right-way risk analysis results

Scenario	Correlations			XVA components					Total XVA	Adjusted price	Regulatory capital
	ρ_{SB}	ρ_{SC}	ρ_{BC}	BS price	CVA	DVA	FVA	KVA			
No Correlation	0.0	0.0	0.0	-14.2313	-0.3272	0.0	-0.1394	-1.6368	-2.1034	-12.1279	35.6260
Mild WWR	0.1	0.2	0.1	-14.2313	-0.3647	0.0	-0.1403	-1.6451	-2.1501	-12.0812	35.7924
Moderate WWR	0.3	0.4	0.2	-14.2313	-0.3926	0.0	-0.1400	-1.6412	-2.1738	-12.0575	36.0163
Strong WWR	0.4	0.6	0.3	-14.2313	-0.4286	0.0	-0.1414	-1.6501	-2.2201	-12.0112	35.9317
Extreme WWR	0.5	0.7	0.4	-14.2313	-0.4386	0.0	-0.1398	-1.6399	-2.2183	-12.0130	35.5890
Mild RWR	-0.1	-0.2	0.1	-14.2313	-0.3034	0.0	-0.1377	-1.6299	-2.0709	-12.1604	35.7924
Moderate RWR	-0.3	-0.4	0.2	-14.2313	-0.2807	0.0	-0.1410	-1.6513	-2.0730	-12.1583	36.0163
Strong RWR	-0.4	-0.6	0.3	-14.2313	-0.2520	0.0	-0.1389	-1.6296	-2.0205	-12.2108	35.9317
Extreme RWR	-0.5	-0.7	0.4	-14.2313	-0.2437	0.0	-0.1420	-1.6562	-2.0419	-12.1894	35.5890

Table 5.6 summarizes the semi-replication XVA results under different correlation scenarios, revealing several important patterns in valuation adjustments.

Remark 5.3.1. In Table 5.6, all prices are reported from the bank’s perspective. Hence, a long call position appears as a negative base price (initial outflow), and valuation adjustments (CVA, FVA, KVA) also enter with negative signs as they represent additional costs. In the discussion below, we refer to their absolute magnitudes for clarity.

First, the Black-Scholes base price of the option is 14.23, providing a neutral benchmark. This is the price the bank would need to pay to the counterparty to long one call option. Under the uncorrelated benchmark, the total XVA is 2.10, dominated by KVA (1.64) with CVA (0.33) and FVA (0.14) playing smaller roles. This corresponds to an adjusted price of 12.13.

The most pronounced variation across scenarios appears in the CVA. In the uncorrelated benchmark, CVA is 0.327, but as the stock-counterparty correlation becomes more positive (our WWR cases), CVA rises monotonically, reaching 0.439 under extreme WWR. This is fully consistent with the payoff structure of a long call option: when the stock price is high, the option exposure $V^+(t)$ is large, and if default risk also increases in such states, losses at default become more severe. The covariance term $\text{Cov}(\lambda_C, V^+)$ is therefore positive, inflating CVA. By contrast, under right-way risk (negative stock-counterparty correlation), defaults are more likely when the option is out-of-the-money and exposures

are small, creating a natural hedge. Accordingly, CVA falls to 0.244 in the extreme RWR case.

The funding adjustment (FVA) remains comparatively stable across all correlation specifications, fluctuating only between 0.138 and 0.143. This reflects the fact that FVA is primarily driven by expected positive exposure discounted at the bank's funding spread, which in our setup is linked to the bank's own default intensity but only weakly affected by stock-counterparty correlation. Correlation perturbs the shape of the exposure distribution, but because the call option's exposure is negligible when the stock is low, the time-averaged funding requirement does not change significantly.

KVA is consistently the dominant component, ranging between 1.63 and 1.66 across all cases. This stability arises because KVA depends on expected (or effective expected) positive exposure profiles aggregated over the full horizon. While correlation affects the joint distribution of exposure and counterparty default, the impact is concentrated in states where the option is deeply out-of-the-money. These states contribute little to the capital calculation, which averages exposures across time and scenarios. As a result, KVA shows only minor variation under either WWR or RWR, confirming that capital costs remain the principal driver of the total XVA.

Overall, the decomposition reveals that correlation primarily influences CVA: stronger WWR amplifies it, while RWR dampens it. FVA and KVA remain broadly unchanged, with the latter continuing to dominate the total adjustment. Consequently, total XVA increases under WWR, reducing the adjusted option price, whereas RWR reduces the total XVA and raises the adjusted price. This highlights that the direction of wrong-way effects is product-dependent: for monotone increasing payoffs like calls, stronger WWR aggravates CVA, while for products with opposite monotonicity the effect would reverse.

Chapter 6

Integrating KVA into a Comprehensive XVA Framework

In this chapter, we synthesize the previous theoretical and numerical findings into a unified framework that incorporates the Capital Valuation Adjustment alongside the more traditional valuation adjustments. We begin by interpreting a financial architecture diagram that illustrates how these adjustments arise within a banking institution, clarifying the roles of the counterparty, trading desk, treasury, external funders, and capital/regulatory layers. We then develop an integration of KVA with the other XVAs. This includes a discussion of how each XVA term enters the pricing equation and how KVA fits into the broader picture. Next, we explain how regulatory capital requirements and the cost of capital are introduced into derivative pricing—contrasting the semi-replication approach and the indifference pricing approach outlined in Chapter 3—and ensure consistency in notation and assumptions. Finally, we discuss the practical interpretation of KVA as the cost of retained capital and how an XVA desk might incorporate KVA into pricing decisions. Throughout, we maintain a narrative focus on conceptual clarity and the interactions among the adjustments, rather than detailed mathematics (which were covered in earlier chapters).

6.1 Unified XVA Architecture

Figure 6.1, which is based on [8], provides a schematic overview of how the main valuation adjustments arise inside a bank. The diagram highlights the principal actors in a derivatives trade: the counterparty, the trading desk (front office), the treasury in charge of liquidity and funding, external funders such as bondholders or money markets, and finally the shareholders' capital and regulatory layer. Arrows indicate either cash movements or risk transfers; each dashed arrow corresponds to one of the XVAs. In what follows we read the figure component by component, clarifying how the adjustments are generated and how they interact.

Counterparty and Trading Desk

The counterparty initiates a derivatives trade with the bank's trading desk. This trade creates two immediate elements: market exposure and counterparty credit risk. The latter gives rise to the Credit Valuation Adjustment (CVA), defined as the expected discounted loss due to counterparty default, as in Equation (3.1.7). Symmetrically, the bank's own default generates a Debit Valuation Adjustment (DVA), as in Equation (3.1.8). However, DVA's recognition is controversial in practice—while it appears as a gain to the bank when its own credit quality worsens, it is only realized in the event of the bank's default,




Figure 6.1: XVA Framework Illustration

at which point shareholders may not actually benefit. Nonetheless, for a bilateral trade, the pricing framework includes both CVA and DVA to adjust for counterparty and own default risk symmetrically.

Many derivative trades require posting collateral to mitigate credit risk. Collateralization reduces counterparty exposure but introduces a spread cost. As shown in Equation (3.1.5), ColVA is negative when collateral earns less than the bank's effective funding rate (collateral posting is costly). If $s_X < 0$, it may even represent a benefit. ColVA therefore links the credit risk mitigation effect of collateral with its funding consequences.

Trading Desk and Treasury

The trading desk does not keep large cash balances; it turns to the treasury whenever cash is needed to fund losses or to post collateral. A request F^+ denotes borrowing from treasury, while a return F^- denotes repayment when the desk is flush with cash. Treasury aggregates these flows and raises external debt (bonds B , loans L) at a spread over the risk-free rate.

The resulting Funding Valuation Adjustment (FVA) measures the cost of unsecured funding. Decomposed into a Funding Cost Adjustment (FCA) and a Funding Benefit Adjustment (FBA), a representative FCA expression is in Equation (3.1.9).

FVA is then obtained as FCA minus FBA. Intuitively, when the desk must borrow at a spread above r , the cost enters as FCA; conversely, when it can lend surplus cash below r , the benefit appears as FBA. Both are summarized by

$$\text{FVA} = \text{FCA} - \text{FBA}.$$

Because the spread reflects the bank's own default risk, careful definitions are needed to avoid double counting between FVA and DVA.

Trading Desk and Shareholders/Capital

A new derivative trade also triggers a regulatory capital requirement. Under Basel III/IV, the bank must hold a capital profile $K(u)$ against future losses. This capital comes from shareholders' equity, who in turn demand a hurdle return h .

The cost of this retained capital is captured by the Capital Valuation Adjustment (KVA). A closed-form expression under semi-replication method is Equation (3.1.10).

Intuitively, each unit of $K(u)$ ties up scarce equity on which shareholders expect h . If the trade does not generate enough profit to cover this charge, the shortfall appears as KVA. Unlike CVA or FVA, KVA is not a hedgeable cash flow but an internal cost: it represents a transfer from the trading desk to shareholders to satisfy their required return.

This distinction is crucial. CVA, DVA, and FVA arise from tradable market frictions (credit spreads, collateral, funding rates). KVA instead reflects a prudential constraint: it ensures that pricing covers the economic cost of capital allocation, thereby aligning derivative valuation with regulatory capital rules and shareholder expectations.

This distinction also explains why two approaches coexist in the literature. In a semi-replication framework, KVA is treated as an additive discounted cost stream, on the same footing as CVA or FVA. In contrast, the indifference pricing method interprets KVA as the premium required to restore the desk's utility once a capital constraint is imposed. Whether one regards KVA as hedgeable or not leads to different pricing adjustments, and we saw that in Chapter 5 the two methods produces different results. This motivates our unified treatment of both perspectives.

Treasury, Shareholders, and Regulator

Beyond the main cash flows, the treasury connects funding and capital. It not only channels FVA charges to the desk but may also levy an internal charge for capital usage. Many banks enforce this via a RAROC framework: each trade must clear the shareholder hurdle rate once both funding and capital costs are applied. In the diagram, this is shown by the feedback arrow between Treasury and Capital, reflecting a transfer-pricing mechanism.

The regulator, in turn, defines the capital rules (credit RWA, CVA capital, leverage ratio, buffers) that determine the profile $K(u)$. These constraints limit trading capacity: a bank cannot scale exposures beyond its available capital headroom, making KVA a binding component in business decisions.

Summary

Figure 6.1 thus illustrates a unified XVA architecture:

- **CVA/DVA** capture bilateral default risk,
- **ColVA** reflects the cost/benefit of collateral carry,
- **FVA** prices unsecured funding spreads,
- **KVA** charges the cost of equity capital required by regulation.

Each arises from a concrete interaction in the diagram: counterparty flows, treasury funding, shareholder capital, and regulatory rules. Valuation adjustments can therefore be read as “value transfers” inside the bank: CVA reserves against counterparty loss, FVA compensates treasury for funding, ColVA accounts for collateral spreads, and KVA ensures shareholders are remunerated for trapped equity.

Methodologically, it is worth emphasizing that CVA, DVA, and FVA are naturally framed under semi-replication, since they correspond to hedgeable or market-spread cash flows. KVA is different: capital is not a traded asset. This is why both semi-replication (treating KVA as a discounted cost stream) and indifference pricing (treating it as a utility premium under a capital constraint) are used in the literature. Whether one regards KVA as hedgeable or not leads to different adjustments—and potentially different deal prices.

Our unified framework therefore keeps both perspectives, to ensure consistency while acknowledging the unique nature of capital costs.

Remark 6.1.1.

1. DVA is symmetric to CVA in pricing but rarely benefits shareholders, so many banks exclude it from performance metrics.
2. A clean split of FVA into FCA/FBA avoids overlap with DVA. We treat FVA net of funding benefits, and keep DVA explicitly for own default risk.

6.2 KVA as the Cost of Retained Capital

The Capital Valuation Adjustment (KVA) represents the cost of retaining shareholder equity to support risky derivative positions. Unlike CVA or FVA, which correspond to hedgeable cash-flow based risks, KVA reflects an unhedgeable burden: capital trapped by regulation that could otherwise be deployed to generate returns elsewhere. From a theoretical perspective, KVA can therefore be interpreted as the opportunity cost of equity, directly linked to the bank’s target return on equity (ROE) or risk-adjusted return on capital (RAROC).

This interpretation highlights a fundamental methodological divide. Under semi-replication, KVA is often embedded into valuation as if it were hedgeable, producing linear, additive adjustments in line with other XVAs. Indifference pricing instead treats KVA as genuinely unhedgeable, pricing it via the bank’s risk preferences and hurdle rates. The former frames KVA as a deduction from value, akin to funding costs; the latter frames it as a required compensation to shareholders, effectively increasing the quoted deal price. Thus, the theoretical status of KVA—as a cost of retained capital—directly conditions whether it is treated as an external charge or an internal profitability requirement.

6.3 KVA in the Trading Desk

From the perspective of a trading desk, KVA is not just a theoretical construct but a binding component of deal economics. It enters pricing and risk management in at least two ways. First, KVA may be charged directly in quotes through internal transfer pricing, appearing alongside CVA and FVA. Second, it operates as a profitability filter: new trades are only approved if their post-KVA returns exceed the bank’s hurdle rate. Our comparative analysis provides several insights into how these mechanisms function in practice.

The portfolio netting experiments show that under semi-replication, XVA charges are additive and scale proportionally with exposure, leading to only modest capital relief when trades are combined. By contrast, under indifference pricing, diversification benefits are amplified: concentrated risks are penalized while offsetting positions receive significant capital relief. For a trading desk, this means that KVA is not determined at the trade level alone but depends critically on portfolio context. Similarly, the wrong-way risk study reveals that while CVA is highly sensitive to correlation assumptions, KVA remains relatively stable, underscoring its role as a persistent and “sticky” driver of total adjustments.

These findings underline KVA’s strategic role on the desk. A replication-style treatment may keep client quotes competitive but risks systematically undercharging capital-intensive trades. An indifference-style treatment protects shareholder returns but may reduce deal flow if competitors quote more aggressively. In practice, many institutions adopt hybrid approaches: client-facing prices closer to semi-replication values, combined with internal RAROC filters for deal approval. Our results provide a theoretical foundation for this practice and quantify the consequences of choosing one lens over the other.

Ultimately, KVA should be understood by the trading desk not as a passive surcharge, but as an active determinant of portfolio composition, competitiveness, and shareholder value.

Chapter 7

Conclusion

This thesis investigated how assumptions about the hedgeability of capital costs affect the pricing of the Capital Valuation Adjustment (KVA). We compared two contrasting methodologies: semi-replication, which treats KVA as if hedgeable and embeds it in a risk-neutral valuation, and indifference pricing, which treats KVA as unhedgeable and links it directly to the bank's risk appetite and return targets. Using both theoretical analysis and numerical experiments on options, swaps, netting sets, and wrong-way risk, we conducted a systematic comparison of the two approaches.

Our findings show that hedgeability assumptions lead to materially different pricing outcomes. Semi-replication produces additive and largely linear adjustments: XVA charges scale with exposure and diversify proportionally under netting. Indifference pricing, in contrast, is non-linear: it penalizes concentrated risks and rewards diversification more strongly, reflecting capital relief at the portfolio level. Overall, semi-replication provides a market-consistent benchmark, while indifference pricing ensures compensation for shareholders by enforcing a hurdle rate.

In practice, these differences imply a trade-off. Pricing purely by semi-replication can make a bank competitive on capital-intensive trades, but at the risk of systematically eroding returns. Indifference pricing protects profitability but risks losing flow business to competitors. Many institutions therefore adopt hybrid approaches: semi-replication prices as outward quotes, supplemented by internal RAROC or KVA charges for deal approval. Our results provide a theoretical rationale for this practice by showing how each methodology captures different dimensions of risk.

Looking forward, a promising direction is to develop unified frameworks that blend replication-based valuation with explicit risk premia for unhedgeable costs. Such models could help reconcile accounting fair value with economic profitability, producing consistent pricing that satisfies both external and internal stakeholders. Ultimately, our study underscores that KVA is not just a technical adjustment but a strategic variable: whether treated as hedgeable or not fundamentally shapes both valuation and capital allocation. We hope these insights contribute to a clearer dialogue between trading desks, risk management, and finance, and stimulate further research into the valuation of unhedgeable risks in derivatives portfolios.

Bibliography

- [1] Claudio Albanese, Simone Caenazzo, and Stéphane Crépey. Capital valuation adjustment and funding valuation adjustment. *arXiv preprint arXiv:1603.03012*, 2016.
- [2] Claudio Albanese, Simone Caenazzo, and Stéphane Crépey. Credit, funding, margin, and capital valuation adjustments for bilateral portfolios. *Probability, Uncertainty and Quantitative Risk*, 2:1–26, 2017.
- [3] Matthias Arnsdorf. Kva as a transfer of wealth. *Risk, December*, 2020.
- [4] Matthias Arnsdorf. Kva is incomplete. 2020.
- [5] Basel Committee on Banking Supervision. International convergence of capital measurement and capital standards: A revised framework. Technical report, Bank for International Settlements, Basel, Switzerland, June 2004. Also known as Basel II.
- [6] Basel Committee on Banking Supervision. Basel iii: A global regulatory framework for more resilient banks and banking systems. Technical report, Bank for International Settlements, December 2011.
- [7] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [8] Damiano Brigo. Msc in mathematics and finance math70111 interest rate models with credit risk, collateral, funding and liquidity risk lecture notes. Imperial College London, internal teaching material, 2024. Unpublished lecture notes.
- [9] Damiano Brigo and Imane Bakkar. Accurate counterparty risk valuation for energy-commodities swaps. *Energy Risk*, March 2009. Trade magazine article.
- [10] Damiano Brigo, Cristin Buescu, Marco Francischello, Andrea Pallavicini, and Marek Rutkowski. Nonlinear valuation with xvas: Two converging approaches. *Mathematics*, 10(5):791, 2022.
- [11] Damiano Brigo, Agostino Capponi, and Andrea Pallavicini. Arbitrage-free bilateral counterparty risk valuation under collateralization and application to credit default swaps. *Mathematical Finance*, 24(1):125–146, January 2014.
- [12] Damiano Brigo and Kyriakos Chourdakis. Counterparty risk for credit default swaps: Impact of spread volatility and default correlation. *International Journal of Theoretical and Applied Finance*, 12(07):1007–1026, 2009.
- [13] Damiano Brigo, Marco Francischello, and Andrea Pallavicini. An indifference approach to the cost of capital constraints: Kva and beyond. *arXiv preprint arXiv:1708.05319*, 2017.

- [14] Damiano Brigo, Marco Francischello, and Andrea Pallavicini. Nonlinear valuation under credit, funding, and margins: Existence, uniqueness, invariance, and disentanglement. *European Journal of Operational Research*, 274(2):788–805, 2019.
- [15] Damiano Brigo, Federico Graceffa, and Alexander Kalinin. Mild to classical solutions for xva equations under stochastic volatility. *SIAM Journal on Financial Mathematics*, 15(1):215–254, 2024.
- [16] Damiano Brigo and Massimo Masetti. Risk neutral pricing of counterparty risk. In Michael Pykhtin, editor, *Counterparty Credit Risk Modelling: Risk Management, Pricing and Regulation*, chapter 10. Risk Books, London, 2006.
- [17] Damiano Brigo, Fabio Mercurio, et al. *Interest rate models: theory and practice*, volume 2. Springer, 2001.
- [18] Damiano Brigo, Massimo Morini, and Marco Tarengi. Credit calibration with structural models and equity return swap valuation under counterparty risk. In Tomasz R. Bielecki, Damiano Brigo, and Frédéric Patras, editors, *Credit Risk Frontiers: Subprime Crisis, Pricing and Hedging, CVA, MBS, Ratings, and Liquidity*, pages 457–484. Wiley / Bloomberg Press, Hoboken, NJ, 2011.
- [19] Damiano Brigo and Andrea Pallavicini. Counterparty risk pricing under correlation between default and interest rates. In John Miller, David Edelman, and John Appleby, editors, *Numerical Methods for Finance*. Chapman & Hall/CRC, Boca Raton, 2007.
- [20] Damiano Brigo and Frédéric Vrins. Disentangling wrong-way risk: pricing credit valuation adjustment via change of measures. *European Journal of Operational Research*, 269(3):1154–1164, 2018.
- [21] Christoph Burgard and Mats Kjaer. In the balance. *C. Burgard, M. Kjaer. In the balance, Risk*, November, pages 72–75, 2011.
- [22] Christoph Burgard and Mats Kjaer. Funding strategies, funding costs. *Risk*, 26(12):82, 2013.
- [23] Chapman and Cutler. Use of excess posted credit support, July 2022.
- [24] Darrell Duffie and Kenneth J. Singleton. Modeling term structures of defaultable bonds. *The Review of Financial Studies*, 12(4):687–720, 1999.
- [25] Financial Stability Board, Basel Committee on Banking Supervision, Committee on Payments and Market Infrastructures, and International Organization of Securities Commissions. Incentives to centrally clear over-the-counter (otc) derivatives. Joint report, Financial Stability Board, Basel, Switzerland, November 2018.
- [26] Andrew Green and Chris Kenyon. Kva: Capital valuation adjustment. *arXiv preprint arXiv:1405.0515*, 2014.
- [27] John Hull and Alan White. Valuing derivatives: Funding value adjustments and fair value. *Financial Analysts Journal*, 70(3):46–56, 2014.
- [28] International Association of Credit Portfolio Managers and Fintegral. Making the most of xva: Practitioner perspectives. Technical report, International Association of Credit Portfolio Managers, New York, NY, 2018. In conjunction with Fintegral.
- [29] Mats Kjaer. Kva from the beginning. *SSRN Electronic Journal*, July 2017.

- [30] Mats Kjaer. Kva unmasked. *SSRN Electronic Journal*, February 2018.
- [31] Linklaters LLP. Cross-collateralisation under the french master agreement, September 2022.
- [32] Linklaters LLP. Cross-collateralisation under the german master agreement (drv), September 2022.
- [33] Linklaters LLP. Cross-collateralisation under the isda master agreement, September 2022.
- [34] Andrea Prampolini and Massimo Morini. Derivatives hedging, capital and leverage. *Capital and Leverage (July 8, 2018)*, 2018.
- [35] Davide Trevisani, José Germán López-Salas, Carlos Vázquez, and José Antonio García-Rodríguez. Mathematical models and numerical methods for a capital valuation adjustment (kva) problem. *Applied Mathematics and Computation*, 488:129105, 2025.
- [36] Oldrich Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.