UNIVERSITY OF TURKU
Turku School of Economics

Tail Credit Losses in Debt Portfolios: Evidence and Measurement in the Vasicek Framework

Evidence from the U.S. Corporate Bond Market

Master's Thesis
in Accounting and Finance

Author:
Tuomas Kerkkonen

Supervisor:
Professor Luis Alvarez Esteban

21/04/2026
Turku

Master's thesis

Subject: Accounting and Finance
Author: Tuomas Kerkkonen
Title: Tail Credit Losses in Debt Portfolios: Evidence and Measurement in the Vasicek Framework
Supervisor: Professor Luis Alvarez Esteban
Number of pages: 69+2
Date: 21/04/2026

Tail credit losses in corporate debt portfolios are highly relevant for risk management, capital planning, and stress testing, because losses in systemic credit downturns can rise far above average loss levels. Despite the practical importance of this issue, there is limited transparent empirical evidence on how issuer-level credit risk inputs translate into portfolio tail losses in large corporate debt universes.

This thesis examines tail credit losses in a large debt-only corporate issuer universe in the United States corporate bond market using the one-factor Vasicek and asymptotic single risk factor (ASRF) framework. It focuses on how issuer-level default probabilities, exposure measures, and loss-given-default assumptions map into portfolio tail risk within a transparent and replicable modelling framework.

The empirical analysis combines issuer-level one-year physical default probability data with instrument-level bond and note data to construct a debt-only portfolio and a seniority-based loss framework. The model is then used to estimate expected and tail loss measures under alternative dependence and severity assumptions, while the asymptotic approximation is also compared with Monte Carlo simulation to assess finite-portfolio effects.

The findings show a clear gap between average and tail outcomes. Tail losses are driven primarily by dependence assumptions, while LGD mainly affects loss severity. The results also suggest that finite-portfolio effects can be economically meaningful in concentrated portfolios, implying that asymptotic approximations may understate risk when concentration is material. Overall, the thesis shows that the framework remains a useful and interpretable tool for portfolio tail credit risk analysis.

Keywords: Credit risk, Tail risk, Vasicek model, Portfolio credit risk

Pro gradu -tutkielma

Oppiaine: Laskentatoimi ja rahoitus
Tekijä: Tuomas Kerkkonen
Otsikko: Tail Credit Losses in Debt Portfolios: Evidence and Measurement in the Vasicek Framework
Ohjaaja: Professori Luis Alvarez Esteban
Sivumäärä: 69+2
Päivämäärä: 21/04/2026

Yrityslainaportfolioiden häntäpään luottotappiot ovat keskeisiä riskienhallinnan, pääomasuunnittelun ja stressitestauksen kannalta, koska systeemisissä luottosuhdanteen heikkenemisissä tappiot voivat nousta selvästi keskimääräisiä tappiotasoja suuremmiksi. Aiheen käytännön merkityksestä huolimatta läpinäkyvää empiiristä näyttöä siitä, miten liikkeeseenlaskijatasoiset luottoriskisyötteet muuntuvat portfoliotason häntätappioiksi laajoissa yritysvelkaportfolioissa, on edelleen rajallisesti.

Tässä tutkielmassa tarkastellaan häntäpään luottotappioita laajassa debt-only-muotoisessa yritysliikkeeseenlaskijaportfoliossa Yhdysvaltain yrityslainamarkkinalla Vasicekin luottoriskimallin ja asymptoottisen yksifaktorisen kehikon avulla. Tutkielma keskittyy siihen, miten liikkeeseenlaskijatasoiset maksukyvyttömyystodennäköisyydet, vastuumitat ja tappioosuusoletukset välittyvät portfolion häntäriskiin läpinäkyvässä ja toistettavassa mallikehikossa.

Empiirinen analyysi yhdistää liikkeeseenlaskijatasoisen yhden vuoden fyysisiä maksukyvyttömyystodennäköisyyksiä sisältävän aineiston sekä instrumenttitason bond- ja note-aineiston. Mallia sovelletaan odotettujen tappioiden ja häntätappiomittareiden estimoimiseen vaihtoehtoisilla riippuvuus- ja severity-oletuksilla, ja asymptoottista approksimaatiota verrataan lisäksi Monte Carlo -simulointiin äärellisen portfoliokoon vaikutusten arvioimiseksi.

Tulokset osoittavat selvän eron keskimääräisten ja häntäpään tappioiden välillä. Erityisesti riippuvuusoletuksilla on keskeinen vaikutus häntätuloksiin, kun taas LGD vaikuttaa ennen kaikkea tappioiden mittakaavaan. Tulokset viittaavat myös siihen, että äärellisen portfoliokoon vaikutukset voivat olla taloudellisesti merkityksellisiä keskittyneissä portfolioissa, jolloin asymptoottiset approksimaatiot voivat aliarvioida riskiä. Kokonaisuutena tutkielma osoittaa, että kehikko on hyödyllinen ja tulkittava väline portfoliotason häntäluottoriskin analysointiin.

Avainsanat: Luottoriski, Häntäriski, Vasicekin malli, Portfolion luottoriski

CONTENTS

FIGURES

TABLES

NOTATIONS

The main notations and abbreviations used in this thesis are summarised below. The list is not necessarily exhaustive; additional symbols are defined locally when first introduced.

T	Risk horizon
i	Issuer / obligor index
j	Instrument index (bond, note, etc.)
N	Number of issuers / obligors in the sample
$\mathcal{J}(i)$	Set of retained debt instruments of issuer i
EAD_{ij}	Issuance-based exposure proxy for instrument j of issuer i (EUR)
EAD_i	Issuer-level exposure, $\mathrm{EAD}_i = \sum_{j\in\mathcal{J}(i)} \mathrm{EAD}_{ij}$
EAD	Total portfolio exposure, $\mathrm{EAD} = \sum_{i=1}^{N} \mathrm{EAD}_i$
w_i	Exposure weight of issuer i, $w_i = \mathrm{EAD}_i/\mathrm{EAD}$
$\mathcal{T}_{10}$	Set of the ten largest issuers by exposure
Top10	Sum of the exposure weights of the ten largest issuers, $\mathrm{Top10} = \sum_{i\in\mathcal{T}_{10}} w_i$
HHI	Herfindahl–Hirschman concentration index, $\mathrm{HHI} = \sum_{i=1}^{N} w_i^2$
δ	Granularity measure; in this thesis, $\delta = \mathrm{HHI} = \sum_{i=1}^{N} w_i^2$
N_{eff}	Effective number of names, $N_{\mathrm{eff}} = 1/\delta = 1/\mathrm{HHI}$
τ_i	Default time of issuer i
$I_i(T)$	Default indicator over horizon T, $I_i(T) = \mathbf{1}\{\tau_i \le T\}$
$p_i(T)$	Physical default probability of issuer i over horizon T
p_i	One-year physical probability of default (PD) for issuer i, $p_i = p_i(1)$
p	Common one-year PD in the homogeneous Vasicek benchmark
p_w	Exposure-weighted mean PD, $p_w = \sum_{i=1}^{N} w_i p_i$
$\mathrm{PD}_{\mathrm{wtd}}$	Exposure-weighted mean PD; equal to p_w
LGD_{ij}	LGD assigned using the seniority-based mapping
LGD_i	EAD-weighted average of LGD
LGD	Common deterministic LGD in the homogeneous Vasicek benchmark
$\bar{\mathrm{LGD}}_w$	Exposure-weighted mean LGD, $\bar{\mathrm{LGD}}_w = \sum_{i=1}^{N} w_i \,\mathrm{LGD}_i$
$\mathrm{LGD}_{\mathrm{wtd}}$	Exposure-weighted mean LGD; equal to $\bar{\mathrm{LGD}}_w$
R_i	Recovery rate of issuer i, $R_i = 1 - \mathrm{LGD}_i$
ρ	Common asset correlation in the homogeneous Vasicek benchmark

ρ_i	Asset correlation of issuer i (IRB function or constant)
$\rho(p)$	Basel IRB-style corporate asset-correlation function as a function of PD
ρ_{wtd}	Exposure-weighted mean correlation, $\rho_{\text{wtd}} = \sum_{i=1}^{N} w_i \rho_i$
Y	Systematic factor, $Y \sim \mathcal{N}(0,1)$
ε_i	Idiosyncratic factor, $\varepsilon_i \sim \mathcal{N}(0,1)$
X_i	Latent asset variable, $X_i = \sqrt{\rho_i}\, Y + \sqrt{1-\rho_i}\, \varepsilon_i$
$\Phi(\cdot)$	Standard normal cumulative distribution function
$\Phi^{-1}(\cdot)$	Inverse standard normal cumulative distribution function
y_α	Systematic stress-state realisation corresponding to confidence level α, $y_\alpha = \Phi^{-1}(1-\alpha)$
$p(y)$	Conditional PD in the homogeneous Vasicek benchmark given $Y = y$
$p_i(y)$	Conditional PD of issuer i given $Y = y$
$p_i(y_\alpha)$	Conditional PD of issuer i in the α-tail stress state
$D(y)$	Conditional default rate given $Y = y$
$L(y)$	Conditional portfolio loss rate given $Y = y$
N_D	Number of defaults over the horizon
D	Exposure-weighted default rate
D_α	α-quantile of the default-rate distribution
$\mathcal{L}(T)$	Portfolio loss amount over horizon T
$\mathcal{L}$	Portfolio loss amount when the horizon argument is omitted
$L(T)$	Portfolio loss rate over horizon T
L	Portfolio loss rate when the horizon argument is omitted
α	Confidence level (baseline $\alpha = 0.999$)
EL	Expected portfolio loss rate
VaR_α	Value-at-Risk at level α (loss rate)
$\text{VaR}_\alpha^{\text{ASRF}}$	ASRF / asymptotic Value-at-Risk at level α
$\text{VaR}_\alpha^{\text{MC}}$	Monte Carlo Value-at-Risk at level α
ES_α	Expected shortfall at level α (loss rate)
UL_α	Unexpected loss, $\text{UL}_\alpha = \text{VaR}_\alpha - \text{EL}$

ABBREVIATIONS

ASRF	Asymptotic Single Risk Factor
IRB	Internal Ratings-Based
BCBS	Basel Committee on Banking Supervision
CDS	Credit Default Swap
PD	Probability of Default
LGD	Loss Given Default
EAD	Exposure at Default
VaR	Value-at-Risk
UL	Unexpected Loss
EL	Expected Loss
HHI	Herfindahl–Hirschman Index
MC	Monte Carlo
TRBC	Thomson Reuters Business Classification

1 Introduction

1.1 Background

Portfolio credit risk is inherently a tail-risk problem. In normal states, credit losses in diversified corporate debt portfolios are often modest relative to portfolio size. In adverse systematic states, however, defaults can become strongly clustered and loss severity can deteriorate at the same time. For this reason, the key risk-management question is not only how large average credit losses are, but how large portfolio losses can become when many obligors are exposed to the same adverse shock.

A central response to this problem has been the development of portfolio credit risk models that map obligor-level probability of default, exposure at default, loss given default, and a dependence structure into a portfolio loss distribution over a fixed horizon. In regulatory and applied credit risk measurement, the most influential benchmark is the Vasicek–ASRF framework. In banking regulation, the internal ratings-based approach operationalises this logic by converting PD, EAD, and LGD inputs into capital requirements calibrated to a high confidence level (Basel Committee on Banking Supervision, 2006, 2005a). In the academic literature, the appeal of the Vasicek framework lies in its analytical tractability: it provides a transparent mapping from issuer-level default risk and dependence assumptions to portfolio tail-loss estimates (Vasicek, 2002; Gordy, 2003).

This tractability is also the framework's central limitation. The asymptotic single-risk-factor benchmark is most accurate when portfolios are highly granular, dependence can be represented by a Gaussian one-factor structure, and input parameters are measured without major distortion. Real portfolios, however, need not satisfy these conditions. They may exhibit meaningful name concentration, sector heterogeneity, incomplete exposure aggregation, and recovery risk that worsens in downturn conditions. Moreover, empirical implementation is non-trivial because issuer-level default probabilities, instrument-level exposures, and priority-sensitive loss severity assumptions are often drawn from different data sources that do not align one-to-one. These frictions imply that credible tail-loss measurement requires attention not only to the benchmark model, but also to portfolio construction, concentration, parameter sensitivity, and model risk (Altman and Kishore, 1996; Schuermann, 2004; Board of Governors of the Federal Reserve System, 2011).

Against this background, the thesis treats the Vasicek–ASRF framework as a transparent benchmark for applied portfolio credit risk measurement and studies how issuer-level data construction, dependence assumptions, concentration, and loss severity jointly shape portfolio tail-loss estimates in a corporate debt setting. What remains relatively limited in the literature, however, is transparent issuer-level empirical work that documents portfolio construction from non-aligned data sources and evaluates tail-loss sensitivity to granularity, dependence, and severity assumptions within one coherent implementation.

1.2 Motivation

The motivation of this thesis is both empirical and methodological. Empirically, applied tail-loss estimation depends on how issuer-level inputs are constructed from imperfectly overlapping PD, debt-instrument, and seniority datasets. The resulting portfolio is therefore not given ex ante, and its coverage, representativeness, and concentration affect how tail-loss estimates should be interpreted. Methodologically, the Gaussian one-factor benchmark is attractive because it is transparent and tractable, but prior literature also shows that its performance depends on granularity, diversification, dependence modelling, and recovery assumptions (Gordy, 2003; Gordy and Lütkebohmert, 2007; Frey and McNeil, 2003; Schuermann, 2004). The practical issue is therefore not whether the benchmark is elegant in theory, but how informative it remains in an issuer-level corporate debt setting once finite granularity, richer dependence structures, and state-dependent loss severity are taken seriously.

1.3 Research questions, objectives, and contribution

The thesis takes a portfolio risk measurement perspective and studies how issuer- and instrument-level information translates into portfolio tail losses within a credit portfolio model. The empirical setting is a matched issuer-level debt-only portfolio drawn from the U.S. corporate bond market, and the baseline modelling approach is deliberately parsimonious in order to keep the mapping from inputs to tail outcomes transparent.

The first research question concerns empirical portfolio construction. How can an issuer-level credit portfolio be built by combining one-year physical default probabilities with an internally consistent exposure proxy derived from debt-instrument data and with seniority-based loss severity assumptions that reflect priority structure? Because the relevant data sources are not aligned one-to-one at issuer level, the portfolio is defined by the intersection of issuers for which the required inputs are simultaneously available. The construction is therefore accompanied by explicit coverage, representativeness, and concentration diagnostics.

The second research question concerns baseline tail risk and parameter sensitivity. What level of expected and tail loss does a baseline Vasicek implementation produce for the constructed debt-only issuer portfolio at high confidence levels, and how sensitive are these tail measures to dependence and recovery assumptions? The objective is to identify which modelling inputs primarily shape the far tail and which mainly scale average loss.

The third research question concerns model risk beyond the asymptotic Gaussian benchmark. When the portfolio is not well approximated by the fine-grained ideal, and when dependence or recovery assumptions are relaxed beyond the standard one-factor specification, how much do tail-loss estimates change? This question links the theoretical ASRF benchmark to practical concentration risk, dependence-model risk, and severity-model risk by quantifying the effects of finite-default granularity, sector structure, heavy-tailed dependence, and state-dependent LGD on extreme-loss outcomes.

The contribution of the thesis is fourfold. First, it constructs a transparent issuer-level debt portfolio from non-aligned PD and debt-instrument datasets and makes the resulting coverage and representativeness limitations explicit. Second, it quantifies baseline expected and tail losses in a standard Vasicek–ASRF framework. Third, it evaluates how strongly these tail estimates change once finite-portfolio effects, concentration, and alternative dependence and severity assumptions are introduced. Fourth, it clarifies how these modelling choices affect the interpretation of issuer-level portfolio tail-loss estimates in an applied credit-risk setting.

1.4 Methods, limitations, and thesis structure

The empirical implementation proceeds in three steps. First, issuer default probabilities are obtained from an LSEG vendor export and interpreted as one-year physical probabilities of default. Second, instrument-level debt data are aggregated to issuer level to form an exposure proxy and portfolio weights, while seniority classifications are mapped into a transparent loss-given-default grid. Third, the resulting issuer-level inputs are fed into a baseline Vasicek framework and selected Monte Carlo extensions in order to compute expected loss and high-quantile tail measures under alternative dependence and severity assumptions. The full data-cleaning and modelling pipeline is implemented in R with reproducible scripts and exported intermediate tables.

The approach is subject to several limitations that guide interpretation. Default probabilities are physical rather than risk-neutral, so the results should be interpreted as risk estimates rather than pricing-implied quantities. The exposure proxy is constructed from issuance-based instrument information rather than contractual exposure at default, which limits the interpretation of absolute euro magnitudes in settings where utilisation, holdings, or risk mitigation determine actual exposure. Recoveries are not observed directly and must therefore be imposed through seniority-based assumptions, which makes LGD sensitivity central to the analysis. In addition, the empirical dataset is a point-in-time cross-sectional snapshot rather than a historical panel of portfolio vintages and realised losses. The thesis therefore evaluates tail-loss estimates, robustness, and model sensitivity, but it does not conduct a full out-of-sample backtesting exercise of realised default or realised portfolio loss outcomes. Finally, the baseline Vasicek framework is intentionally stylised and abstracts from dynamic rating migration, endogenous spread-default feedback, and full market-consistent pricing dynamics.

The remainder of the thesis is structured as follows. Section 2 reviews the theoretical background on claim priority, loss severity, and portfolio tail-loss modelling. Section 3 discusses the role of portfolio credit risk models in practice and motivates the empirical design. Section 4 describes the data, matching procedure, issuer-level input construction, and baseline implementation choices. Section 5 presents the empirical results, beginning with coverage, representativeness, and baseline portfolio characteristics, and then extending the analysis to sensitivities, finite-portfolio effects, alternative dependence structures, and state-dependent severity assumptions. Section 6 discusses the results in relation to prior literature and applied credit risk measurement. Section 7 concludes the thesis.

2 Theoretical Background

This chapter establishes the conceptual and modelling foundations required for the empirical portfolio credit risk analysis. The objective is to connect instrument-level credit inputs—one-year default probabilities, exposure measures, and loss severities—to decision-relevant portfolio tail metrics over a fixed horizon. The discussion proceeds in three steps. First, it motivates why claim priority and recovery mechanisms matter for loss severity and why recoveries may be state dependent. Second, it introduces the portfolio loss framework and the tail measures used in practice. Third, it develops the Vasicek–ASRF framework and discusses finite-portfolio frictions such as granularity and concentration, calibration and estimation uncertainty, and robustness considerations that guide the empirical scenario design. Throughout, the focus is on a transparent horizon-loss perspective. The chapter does not aim to develop a full dynamic pricing model of credit spreads, but rather provides the minimum theoretical structure needed to interpret the empirical tail-loss results and sensitivity analyses.

2.1 Primitives and scope

This thesis treats portfolio credit risk as a fixed-horizon loss measurement problem. Let T denote the risk horizon, and let τ_i denote the default time of obligor i. Default over the horizon occurs when $\tau_i \leq T$, and the corresponding physical default probability is

$$p_i(T) = \mathbb{P}(\tau_i \leq T).$$

In the empirical implementation of this thesis, the horizon is set to one year, so the relevant input is the one-year physical probability of default,

$$p_i = \mathbb{P}(\tau_i \leq 1).$$

Let $I_i(T) := \mathbf{1}\{\tau_i \leq T\}$ denote the default indicator over horizon T. Portfolio losses are measured through the horizon loss amount

$$\mathcal{L}(T) = \sum_{i=1}^{N} \mathrm{EAD}_i \, \mathrm{LGD}_i \, I_i(T).$$

The corresponding loss rate is

$$L(T) = \frac{\mathcal{L}(T)}{\mathrm{EAD}}.$$

With the horizon fixed, write $L := L(T)$ and $\mathcal{L} := \mathcal{L}(T)$ when no ambiguity arises. In this framework, EAD_i captures the amount at risk at the time of default, while LGD_i captures loss severity conditional on default and embeds recovery assumptions, workout costs, and discounting conventions as required by the intended application. For tractability in the baseline specification, EAD_i and LGD_i are treated as fixed inputs over the chosen horizon. Dependence between obligors is introduced through a systematic component that drives default clustering and therefore tail outcomes; the baseline implementation uses a one-factor structure in the spirit of the ASRF framework, while ASRF-style closed-form results additionally rely on asymptotic granularity (Gordy, 2003; Basel Committee on Banking Supervision, 2005a).

The scope of the theoretical framework is deliberately aligned with the empirical design. The objective is not to model the full time path of credit spreads or instrument prices, but to obtain a transparent mapping from empirically constructed issuer-level PD, LGD, and EAD inputs, together with dependence assumptions, into portfolio tail-loss measures over a fixed horizon. Accordingly, the discussion emphasises realised-loss measurement and economic or regulatory capital interpretation under physical probabilities. Risk-neutral probabilities and dynamic intensity models are introduced only to clarify the distinction between risk measurement and valuation, and to avoid conflating market-implied default likelihoods with physical default probabilities (Hull, Predescu and White, 2005). Similarly, rating migration dynamics and multi-period specifications are discussed as extensions and robustness considerations, rather than as the baseline engine used for the empirical tail estimates.

2.2 Claim priority, recovery, and loss severity

Recoveries are not only heterogeneous across instruments but also plausibly state dependent. Seniority, collateralisation, and the security package affect expected loss severity across claims, while adverse macro-financial conditions may simultaneously increase default incidence and depress recoveries. For portfolio tail measurement, this means that constant unconditional LGD assumptions can understate losses precisely in the stress states that drive capital-style risk metrics. This provides a direct motivation for the LGD sensitivity scenarios used in the empirical analysis.

2.2.1 Capital structure and subordination as a loss-allocation mechanism

Firms finance assets with a mix of equity and debt, and the capital structure determines how cash flows and control rights are allocated across claimholders. In the context of credit losses, the key dimension is not only the debt–equity split but the priority structure within debt: which claims are paid first and which absorb losses first in distress (Brealey, Myers and Allen, 2019; Baird, 2017).

In formal insolvency and reorganisation settings, priority is often benchmarked by the absolute priority rule (APR), although its implementation can vary by jurisdiction and procedure, for example through bargaining in U.S. Chapter 11 and through relative-priority features in EU preventive restructuring frameworks (Baird, 2017; European Parliament and Council of the European Union, 2019). In this thesis, such institutional variation is abstracted into a stylised priority ordering that is operationalised empirically through seniority-based LGD assumptions.

From a valuation perspective, priority can be illustrated through a stylised payoff waterfall consistent with the contingent-claim view of corporate liabilities (Merton, 1974; Geske, 1977). Let enterprise value at resolution be W, and let the face values of senior and junior debt be D_{sen} and D_{jun}. Abstracting from interim coupons, costs, and deviations from strict priority, and assuming $W \geq 0$, payoffs follow:

$$\begin{aligned}
\text{Payoff}_{\text{sen}}(W) &= \min\{D_{\text{sen}},\, W\}, \\
\text{Payoff}_{\text{jun}}(W) &= \min\{D_{\text{jun}},\, \max(W - D_{\text{sen}}, 0)\}, \\
\text{Payoff}_{\text{eq}}(W) &= \max(W - D_{\text{sen}} - D_{\text{jun}}, 0).
\end{aligned}$$

For debt claims $k \in \{\text{sen}, \text{jun}\}$, with face value D_k, define instrument-level recovery as payoff divided by face value,

$$R_k(W) = \frac{\text{Payoff}_k(W)}{D_k},$$

and the realised loss rate as

$$\ell_k(W) = 1 - R_k(W).$$

Priority in law is therefore a useful benchmark, but realised recoveries are ultimately economic objects shaped by the resolution process and contractual frictions. In particular, realised recoveries depend on collateral coverage, enforcement and bankruptcy costs, and the timing of cash recoveries, and they can deviate from strict priority due to bargaining, court discretion, or special-priority claims in reorganisation (Baird, 2017; Altman, Hotchkiss and Wang, 2019). In practice, even when APR provides the conceptual ordering, recovery outcomes can differ materially across otherwise identical seniority labels because of differences in security packages, collateral valuation haircuts, and the composition of claims ahead of or pari passu with a given instrument. This motivates treating the seniority-to-LGD mapping in this thesis as an operational recovery grid that captures the dominant cross-sectional effect of subordination while abstracting from jurisdiction- and process-specific deviations that are difficult to observe consistently in large instrument datasets.

In the IRB framework, loss given default is measured as economic loss given default expressed as a percentage of exposure at default, and economic loss incorporates discount effects as well as material direct and indirect workout and collection costs. Because recoveries are often realised with delay, economic LGD is naturally interpreted in present-value terms at default rather than as an undiscounted ultimate recovery (Basel Committee on Banking Supervision, 2006).

In applications, the LGD parameter is typically interpreted as the expected value of this loss rate under the relevant conditioning set, such as downturn conditions and workout costs, rather than as a single-state realised outcome (Schuermann, 2004; Basel Committee on Banking Supervision, 2005b). Under this notation, subordination mechanically reallocates losses across instruments. Junior claims absorb losses earlier, while senior claims are protected until junior is exhausted. Empirically, recoveries vary systematically by seniority, for example across secured, unsecured, and subordinated claims (Altman and Kishore, 1996).

A further tail-risk-relevant friction is that default likelihood and recovery severity need not be independent. Empirical and institutional considerations suggest that recoveries tend to be lower precisely in states when defaults are more clustered. Market-wide distress reduces buyer capacity, tightens financing conditions for distressed transactions, and depresses collateral and going-concern values, producing so-called downturn recoveries (Basel Committee on Banking Supervision, 2005b). In portfolio terms, this introduces a form of wrong-way dependence between PD and LGD, where systematic stress simultaneously increases default incidence and loss severity. Since the Vasicek ASRF mapping is highly sensitive to systematic stress in the tail, allowing LGD to rise in adverse states is economically important and provides a concrete rationale for the conservative LGD calibration and LGD-sensitivity scenarios employed in the empirical analysis.

2.2.2 Structural intuition

Structural credit risk models link default to the evolution of firm value relative to promised payments. In the canonical contingent-claims formulation, equity can be viewed as a residual claim on firm assets and risky debt as a defaultable claim whose payoff depends on whether asset value remains above the debt obligation at the relevant horizon. In the baseline Merton setup, this horizon is maturity (Merton, 1974). Default is thus an endogenous event generated by the firm's balance-sheet leverage and the volatility of its asset value: higher leverage reduces the distance to the default boundary, and higher asset risk increases the probability that the boundary is crossed over a given horizon. Extensions with an explicit default barrier, such as safety-covenant or first-passage formulations, sharpen this intuition by allowing default to occur when firm value hits a threshold before maturity and by emphasising how contractual design, for example through covenant-type triggers, can effectively shift the relevant boundary (Black and Cox, 1976).

While the empirical analysis in this thesis does not estimate structural default probabilities from equity prices or balance-sheet dynamics, structural logic remains useful as an organising interpretation for how capital structure and financing arrangements feed into portfolio loss measurement. First, leverage and the tightness of promised payments shape the likelihood of distress and default, so issuer-level PD inputs can be interpreted as summarising the firm's effective distance to default under the prevailing information set. Second, contractual priority and security packages determine how enterprise value at resolution is allocated across claimholders, connecting the same underlying firm-value realisation to heterogeneous recoveries across instruments. In structural terms, senior secured debt is protected by both priority and collateral, whereas junior or subordinated claims absorb losses earlier, implying higher expected loss severity conditional on default. The seniority-to-LGD mapping employed in this thesis can therefore be viewed as a reduced-form representation of the structural payoff waterfall under limited observability of collateral values, enforcement costs, and bargaining outcomes. Taken together, structural intuition provides a coherent link between financing choices and the two primitives that drive horizon loss distributions in portfolio models: default likelihood and loss severity. Variations in leverage, asset risk, and contractual design are reflected in the empirical inputs through issuer-level PDs and instrument-level LGD assumptions, which are then combined with a dependence structure to quantify portfolio tail losses over the one-year horizon considered in the empirical analysis (Hull, 2021).

2.3 Portfolio tail risk measurement

2.3.1 Portfolio loss, expected loss, and tail measures

Portfolio credit risk concerns the distribution of losses over a given horizon due to defaults across many obligors. The core inputs are exposure at default (EAD), probabilities of default (PDs), loss given default (LGD), and a dependence structure governing default co-movement (McNeil, Frey and Embrechts, 2015). A standard one-period portfolio loss rate can be written as

$$L(T) = \sum_{i=1}^{N} w_i \, \mathrm{LGD}_i \, I_i(T), \qquad w_i = \frac{\mathrm{EAD}_i}{\mathrm{EAD}}, \qquad \sum_{i=1}^{N} w_i = 1.$$

where $I_i(T) := \mathbf{1}\{\tau_i \leq T\}$ is the default indicator over the horizon T, and w_i is the exposure weight of obligor i in the portfolio.

Two summary objects are central for the empirical analysis. The expected loss is

$$\mathrm{EL} = \mathbb{E}[L(T)].$$

Tail risk is measured by a high-quantile loss VaR_α and the corresponding unexpected-loss component

$$\underset{\alpha}{\mathrm{UL}} = \underset{\alpha}{\mathrm{VaR}} - \mathrm{EL},$$

with a one-year horizon and a high confidence level motivated by regulatory and risk-management practice (Basel Committee on Banking Supervision, 2006).

For completeness, expected shortfall at level α is defined as the conditional mean loss beyond the VaR threshold,

$$\mathrm{ES}_\alpha = \mathbb{E}\Big[L(T) \mid L(T) \geq \underset{\alpha}{\mathrm{VaR}}\Big].$$

Expected shortfall is not a central reported risk measure in the empirical analysis. The main empirical focus remains on expected loss, VaR, and unexpected loss, while an expected-shortfall-style proxy is reported only for completeness in the baseline results (Yamai and Yoshiba, 2002).

2.3.2 One-factor tail mapping

To translate obligor-level inputs into portfolio tail loss, this thesis adopts a one-factor framework aligned with the IRB and ASRF logic, in which systematic stress drives default clustering and tail losses (Gordy, 2003). The basic idea is that, over a fixed horizon, obligor defaults are conditionally independent given a single common risk factor that summarises the macro-financial credit state. A sufficiently adverse realisation of this factor raises many conditional default probabilities at the same time, so portfolio tail losses arise from default clustering rather than from isolated single-name events. In this way, the one-factor framework provides a tractable link from individual credit quality to portfolio loss quantiles.

Formally, the ASRF representation can be viewed as a latent-variable threshold model. Each obligor i is associated with an unobserved asset-return variable that depends on a common systematic factor and an idiosyncratic component, and default occurs when the latent variable falls below a threshold chosen to match the obligor's unconditional one-year PD (Vasicek, 2002). Dependence enters through the factor loading, or equivalently through the asset-correlation parameter ρ_i (or ρ in the homogeneous benchmark), which governs how strongly conditional default probabilities respond to systematic stress. More generally, asset correlation governs default clustering in adverse states and therefore the thickness of the portfolio loss tail (Gordy, 2003).

Within this mapping, LGD determines how clustered default events are translated into clustered loss realisations. Conditional on a given systematic state, higher LGD raises loss severity name by name, while higher dependence increases the mass of joint-default scenarios. The quantile mapping therefore combines three ingredients: obligor PDs, which anchor default thresholds; a dependence assumption, which governs the stressed conditional default environment; and LGD, which converts default events into monetary or rate losses. The resulting portfolio loss quantile can be interpreted as the loss corresponding to a sufficiently adverse realisation of the common credit factor and is therefore naturally suited to capital-style tail measurement.

The one-factor mapping used in this thesis is a disciplined approximation rather than a complete description of credit crises. It abstracts from multi-factor sector structure, rating-migration dynamics, and endogenous feedback effects, and it relies on a granularity condition under which idiosyncratic risk is diversified away. These assumptions motivate the robustness checks reported in the empirical section, including finite-portfolio comparisons, sector-structure extensions, and alternative dependence specifications. The specific Vasicek–ASRF quantile mapping and its empirical implementation are presented formally in Section 2.5.

2.4 Finite-portfolio effects and concentration

2.4.1 Granularity, concentration, and model risk

The ASRF representation relies on a granularity condition under which idiosyncratic risk diversifies away in large, fine-grained portfolios (Gordy, 2003). Real portfolios can violate this assumption because of name concentration and sector tilts. A convenient concentration diagnostic is the Herfindahl index of exposure weights,

$$\delta = \sum_{i=1}^{N} w_i^2,$$

with the associated effective number of names

$$N_{\text{eff}} = \frac{1}{\delta}.$$

Under equal weights, $\delta = 1/N$, so larger values of δ and smaller values of N_{eff} indicate greater concentration. When concentration is meaningful, idiosyncratic default risk remains relevant in the tail and can cause asymptotic tail approximations to understate extreme-loss risk (Gordy, 2003; Gordy and Lütkebohmert, 2013).

In the empirical analysis, finite-portfolio model risk is therefore assessed by comparing asymptotic tail estimates with a Monte Carlo or discrete-default implementation under the same one-factor structure and by weight-structure experiments that vary portfolio granularity, for example equal-weight portfolios versus top-k portfolios. The purpose of this comparison is to isolate the effect of relaxing the asymptotic approximation within the same dependence specification, rather than to compare VaR levels across different issuer universes.

2.5 The Vasicek ASRF framework and one-factor extensions

2.5.1 Model setup and the asymptotic loss distribution

The Vasicek one-factor model, pioneered by Oldřich Vašíček, adopts an obligor-level latent-variable view of default that yields a tractable asymptotic default-rate distribution. In the asymptotic limit of large, well-diversified portfolios, idiosyncratic shocks diversify away and the model provides a closed-form mapping from obligor-level default inputs to portfolio default-rate quantiles (Vasicek, 2002). Its basic intuition is parsimonious: a single systematic factor drives default co-movement, while idiosyncratic shocks wash out in sufficiently granular portfolios.

Following Vasicek (2002), consider a large granular portfolio of initially homogeneous obligors with unconditional default probability p and asset correlation ρ. Here, p denotes the common one-year default probability in the homogeneous Vasicek benchmark and should not be confused with the empirical exposure-weighted mean PD, denoted by p_w in the empirical sections. Defaults are represented through the latent variable

$$X_i = \sqrt{\rho}\, Y + \sqrt{1-\rho}\, \varepsilon_i, \qquad \text{default if } X_i < \Phi^{-1}(p),$$

where $Y, \varepsilon_i \sim \mathcal{N}(0,1)$ are independent. Here, Y is the common systematic factor and ε_i is the obligor-specific idiosyncratic shock.

Let $I_i(T) := \mathbf{1}\{\tau_i \leq T\}$ denote the default indicator over horizon T, and define the default count and exposure-weighted default rate as

$$N_D := \sum_{i=1}^{N} I_i(T), \qquad D := \sum_{i=1}^{N} w_i\, I_i(T), \qquad \sum_{i=1}^{N} w_i = 1.$$

Under equal weights $w_i = 1/N$, D reduces to the default fraction $N^{-1} \sum_{i=1}^{N} I_i(T)$.

As portfolio size grows and the granularity condition $\sum_{i=1}^{N} w_i^2 \to 0$ holds, the default-rate distribution converges to the Vasicek limit with cumulative distribution function

$$\Pr[D \le x] = \Phi\left(\frac{\sqrt{1-\rho}\,\Phi^{-1}(x) - \Phi^{-1}(p)}{\sqrt{\rho}}\right), \quad 0 \le x \le 1. \tag{1}$$

The closed-form CDF in (1), and the corresponding closed-form quantile in (2), pertain to the canonical homogeneous-portfolio Vasicek limit with common (p, ρ) and asymptotically small exposure shares. In heterogeneous portfolios, ASRF implementations proceed through conditional expected loss given the systematic factor rather than by applying (1) directly.

The mean of the limiting default rate is $\mathbb{E}[D] = p$. Conditional on a realised systematic state $Y = y$, the obligor default probability becomes

$$p(y) = \Phi\left(\frac{\Phi^{-1}(p) - \sqrt{\rho}\,y}{\sqrt{1-\rho}}\right),$$

so for a granular portfolio the default rate satisfies $D(y) \approx p(y)$. This representation is especially useful for tail-risk interpretation: sufficiently low realisations of y correspond to systematic stress and therefore to clustered defaults.

The same construction can also be interpreted as a one-factor Gaussian copula model for default events over a fixed horizon. Writing

$$U_i = \Phi(X_i), \qquad i = 1, \dots, N,$$

implies that each $U_i \sim \mathrm{Unif}(0, 1)$ marginally, while their joint dependence is governed by the Gaussian copula induced by the one-factor normal structure. In this sense, the Vasicek ASRF framework combines one-year default marginals with a Gaussian dependence structure. This copula interpretation is useful because it makes the dependence assumption explicit and therefore provides a natural benchmark for later robustness checks under heavier-tailed alternatives such as t-copulas or heavy-tailed systematic factors (Li, 2000; McNeil et al., 2015).

The horizon and dynamics of the static one-year ASRF framework differ from those of dynamic credit models. ASRF is a fixed-horizon model: inputs are calibrated to the event $\{\tau_i \le T\}$, typically with $T = 1$, and portfolio risk is assessed through the distribution of the horizon default rate or horizon loss. It does not model the timing of defaults within the horizon, nor the joint evolution of spreads and prices. When the objective is instead to model loss evolution across maturities or to price credit-sensitive claims consistently across a term structure, one typically moves to dynamic specifications such as copula-based time-to-default models or reduced-form intensity models (Hull and White, 2008;

Duffie and Singleton, 1999).

Several limiting properties are useful for tail-risk interpretation. As $\rho \to 0$, idiosyncratic diversification becomes complete and the default-rate distribution collapses to the degenerate limit $D(y) = p$. As $\rho \to 1$, defaults become perfectly clustered and the portfolio behaves in an all-or-nothing fashion. More generally, upper-tail portfolio outcomes correspond to sufficiently adverse systematic states. Finally, Gaussian dependence implies zero tail dependence in the copula sense. Heavier-tailed dependence structures, such as t-copulas or heavy-tailed systematic factors, generate positive tail dependence and therefore larger probabilities of joint extreme default realisations, which motivates robustness checks beyond the Gaussian ASRF baseline (McNeil et al., 2015).

2.5.2 Capital mapping: loss quantiles and the IRB ASRF link

For this reason, the framework underpins modern portfolio credit capital calculations, including the IRB ASRF approach (Gordy, 2003). The ASRF assumption, meaning Asymptotic Single Risk Factor, combines (i) an asymptotically fine-grained portfolio and (ii) a single systematic factor, implying that portfolio tail losses are driven by systematic stress rather than idiosyncratic noise.

Write D_α for the α-quantile of D. Inverting (1) gives

$$D_\alpha = \Phi\left(\frac{\Phi^{-1}(p) + \sqrt{\rho}\,\Phi^{-1}(\alpha)}{\sqrt{1-\rho}}\right). \tag{2}$$

In the Basel II IRB framework, capital is calibrated at a high confidence level, commonly implemented by setting the systematic stress via $\Phi^{-1}(0.999)$ and interpreting PD as a one-year PD (Basel Committee on Banking Supervision, 2005a). Since, for a granular portfolio, $D(y) \approx p(y)$, the quantile can be interpreted as a stressed default rate. With the sign convention in (1), a convenient choice is

$$y_\alpha = \Phi^{-1}(1-\alpha),$$

so that $y_{0.999} < 0$, and hence

$$D_\alpha \approx p(y_\alpha).$$

In the homogeneous benchmark, let $\mathrm{LGD}_i \equiv \mathrm{LGD}$ for all i. With deterministic LGD and homogeneous, asymptotically small exposure shares, the portfolio loss-rate quantile satisfies

$$\operatorname*{VaR}_\alpha(L) \approx \mathrm{LGD}\, D_\alpha,$$

since in the homogeneous-LGD case $L = \mathrm{LGD} \cdot D$. More generally, with deterministic but heterogeneous severities, the ASRF loss mapping for the portfolio loss rate can be written as

$$L(y) \approx \sum_{i=1}^{N} w_i \, \mathrm{LGD}_i \, p_i(y).$$

As a special case, if obligors remain homogeneous on the default side so that they share the same conditional default probability $p(y)$, then

$$L(y) \approx p(y) \, \bar{\mathrm{LGD}}_w, \qquad \bar{\mathrm{LGD}}_w := \sum_{i=1}^{N} w_i \, \mathrm{LGD}_i,$$

implying

$$\underset{\alpha}{\mathrm{VaR}}(L) \approx \bar{\mathrm{LGD}}_w \, D_\alpha, \qquad \mathrm{EL} \approx \bar{\mathrm{LGD}}_w \, p.$$

In amount terms, the corresponding portfolio loss amount is

$$\mathcal{L} = \Big(\sum_{i=1}^{N} \mathrm{EAD}_i \Big) L,$$

so that

$$\underset{\alpha}{\mathrm{VaR}}(\mathcal{L}) \approx \Big(\sum_{i=1}^{N} \mathrm{EAD}_i \Big) \underset{\alpha}{\mathrm{VaR}}(L).$$

Accordingly, under homogeneous LGD,

$$\underset{\alpha}{\mathrm{VaR}}(\mathcal{L}) \approx \Big(\sum_{i=1}^{N} \mathrm{EAD}_i \Big) \mathrm{LGD} \, D_\alpha,$$

while in the special case above with deterministic heterogeneous severities,

$$\underset{\alpha}{\mathrm{VaR}}(\mathcal{L}) \approx \Big(\sum_{i=1}^{N} \mathrm{EAD}_i \Big) \bar{\mathrm{LGD}}_w \, D_\alpha.$$

A common capital proxy is unexpected loss at level α,

$$\underset{\alpha}{\mathrm{UL}} = \underset{\alpha}{\mathrm{VaR}}(L) - \mathrm{EL} \approx \bar{\mathrm{LGD}}_w (D_\alpha - p),$$

which reduces to

$$\underset{\alpha}{\mathrm{UL}} \approx \mathrm{LGD}(D_\alpha - p)$$

in the homogeneous-LGD case. In amount terms, the corresponding unexpected loss amount is

$$\Big(\sum_{i=1}^{N} \mathrm{EAD}_i \Big) \bar{\mathrm{LGD}}_w (D_\alpha - p).$$

2.5.3 Inputs, baseline calibration and parameter estimation

The parameters of the Vasicek ASRF framework have clear economic roles. The unconditional default probability p is an input that can be calibrated to reflect a long-run average one-year PD in capital applications. In Basel II IRB applications, p is interpreted as a one-year PD estimated from long-run experience. The asset-correlation parameter ρ governs dependence through the common factor and should not be interpreted as a pairwise default correlation; default correlation is instead a nonlinear function of both p and ρ. Loss-given-default translates default rates into losses measured as economic loss. Regulatory practice specifies ρ through supervisory correlation prescriptions by exposure class, often as a function of PD, and uses downturn LGD rather than average LGD to ensure conservatism (Basel Committee on Banking Supervision, 2006, 2005a).

In practice, PD inputs may be based on ratings and historical default experience, market-implied measures under identifying assumptions, or model-based signals. LGD can be calibrated from historical recoveries with a downturn overlay, while dependence parameters may be inferred from default co-movement, factor or probit specifications, asset-return proxies, or supervisory prescriptions (McNeil et al., 2015; Hull, 2021). Scenario analysis and stress testing enter naturally by shifting the common factor to reflect adverse macro-financial conditions (Vasicek, 2002). Among model-based PD approaches, the Moody's KMV framework provides a structural route from equity data to empirical default probabilities by inferring asset value and asset volatility from equity-market information, although such estimation is not implemented in this thesis (Crosbie and Bohn, 2003; Bharath and Shumway, 2008).

A practical implementation of the Vasicek ASRF framework hinges not only on baseline calibration but also on the uncertainty of the inputs (p, ρ, LGD). While the mapping from inputs to tail loss is analytically transparent, these parameters are estimated with non-trivial uncertainty, and such parameter risk can be material for high-quantile measures such as $\mathrm{VaR}_{0.999}$ and the resulting capital estimates. In particular, default dependence and the underlying asset-correlation parameter are difficult to estimate reliably because defaults are rare events and available panel lengths are often short relative to the frequency of systemic stress episodes. Finite-sample estimators can exhibit substantial dispersion and may be biased, which may translate into an understatement of portfolio tail risk when point estimates are used mechanically (Gordy and Heitfield, 2002).

For PD, a simple binomial benchmark illustrates the sampling problem. If one-year default events were independent across obligors and $N_D \sim \mathrm{Bin}(N, p)$, then the sample default-rate estimator

$$\hat{p} = \frac{N_D}{N}$$

would satisfy

$$\mathrm{Var}(\hat{p}) = \frac{p(1-p)}{N}, \qquad \mathrm{se}(\hat{p}) = \sqrt{\frac{p(1-p)}{N}}.$$

In practice, the standard error may be estimated by

$$\sqrt{\frac{\hat{p}(1-\hat{p})}{N}}.$$

This independence benchmark is useful only as a lower-complexity reference, because in credit portfolios defaults are typically clustered rather than independent.

The effective information content decreases further when defaults are clustered, that is, when conditional independence given Y holds but the unconditional default indicators remain dependent. For ρ, identification is especially weak in short samples. Dependence is inferred indirectly from default co-movement, factor or probit specifications, or asset-return proxies, and uncertainty in ρ is amplified in tail measures because the upper tail of the loss distribution is highly sensitive to dependence assumptions (Gordy and Heitfield, 2002). LGD estimates, in turn, are inferred from recovery data and are subject to measurement noise, selection effects, and cyclicality, which motivates conservative or downturn-consistent overlays in capital applications.

Taken together, uncertainty in PD, correlation, and LGD provides a strong rationale for robustness analysis around baseline ASRF point estimates. In the empirical part of this thesis, this motivation is reflected in correlation sensitivities, LGD stress scenarios, finite-portfolio comparisons, and alternative dependence specifications reported alongside the baseline results (Gordy and Heitfield, 2002).

2.5.4 Finite granularity, concentration, and risk contributions

If a few names are large, idiosyncratic risk does not wash out. Let

$$\delta = \sum_{i=1}^{N} w_i^2$$

denote the Herfindahl index of exposure shares, and define the effective number of names as

$$N_{\mathrm{eff}} = \frac{1}{\delta},$$

so that under equal weights $\delta = 1/N$. As δ increases, that is, as the portfolio becomes more concentrated, the law-of-large-numbers argument that collapses conditional losses to $p(y)$ weakens, and idiosyncratic risk contributes non-negligibly to tail outcomes. Granularity adjustments therefore add finite-N corrections to asymptotic tail measures; formal treatments and practical approximations are provided by Gordy and Lütkebohmert (2013).

Concentration also matters for attribution. In a concentrated portfolio, a small set of large exposures can account for a disproportionate share of tail risk. A general theoretical benchmark for such attribution is Euler allocation of positively homogeneous risk measures. Let $x_i \geq 0$ denote the scalable position size of exposure i, for example euro EAD, and define the portfolio loss amount as

$$\mathcal{L}(x) = \sum_{i=1}^{N} x_i \,\mathrm{LGD}_i\, \mathbf{1}\{\tau_i \leq T\}.$$

Under standard regularity conditions, a positively homogeneous risk measure ϱ admits the decomposition

$$\varrho(\mathcal{L}(x)) = \sum_{i=1}^{N} x_i \,\frac{\partial \varrho(\mathcal{L}(x))}{\partial x_i}.$$

For VaR_α, this motivates defining risk contributions as

$$\mathrm{RC}_i := x_i \,\frac{\partial \mathrm{VaR}_\alpha(\mathcal{L}(x))}{\partial x_i},$$

so that portfolio tail risk can be decomposed into name-level components (Tasche, 1999). Intuitively, RC_i measures how strongly portfolio tail loss responds to a marginal scaling of exposure i, so large names naturally emerge as dominant contributors when exposures are uneven.

In the empirical section, this logic motivates reporting both concentration metrics and tail-risk attribution measures. In particular, Top-10 exposure shares, Herfindahl-based concentration measures, and issuer- and sector-level tail-state contributions are reported jointly in order to show how exposure concentration translates into tail-risk concentration. The key point is that finite granularity is not only a question of approximation error relative to the ASRF limit, but also a question of which names and sectors account for a large share of the portfolio tail.

2.5.5 Extensions beyond the one-factor, one-period baseline

The Vasicek ASRF setup provides a parsimonious baseline in which a single systematic factor governs default co-movement over a fixed horizon. Two natural extensions are particularly relevant in the context of this thesis: allowing for richer dependence structures and relaxing the Gaussian one-factor benchmark.

In practice, credit portfolios are exposed to several macroeconomic and sectoral sources of systematic risk, suggesting a dependence structure richer than a single common factor.

A generic multi-factor generalisation replaces the one-factor loading by a vector of factor loadings $\beta_i \in \mathbb{R}^m$ satisfying $\|\beta_i\| \leq 1$,

$$X_i = \beta_i^\top Y + \sqrt{1 - \|\beta_i\|^2}\, \varepsilon_i, \qquad Y \in \mathbb{R}^m, \;\; Y \sim \mathcal{N}(0, I_m),$$

so that obligor dependence is driven by a low-dimensional systematic vector Y, for example global, regional, and sector factors. While exact tail calculations become more involved, such structures provide a natural way to capture heterogeneous co-movement beyond the single-factor ASRF benchmark (Pykhtin, 2004). This extension is directly relevant for the empirical robustness analysis in this thesis, where sector-based multi-factor dependence is used to assess how strongly the results rely on the one-factor approximation.

A second limitation of the baseline ASRF model is its fixed-horizon nature. It targets the distribution of $L(T)$ at a given horizon, typically $T = 1$ year, but does not describe how credit risk evolves across time. Multi-period extensions of the Vasicek framework can be used to study loss distributions across dates or maturities while retaining part of the analytic tractability of the one-period benchmark (García Céspedes and Moreno, 2017). Such extensions are useful when the research question concerns maturity transformation, refinancing risk, or horizon-to-horizon dynamics, but they are outside the scope of this thesis, which remains focused on one-year tail loss measurement.

2.5.6 Model risk and robustness checks

The Vasicek ASRF implementation adopted in this thesis is intentionally transparent, but it is also a stylised approximation. Robustness checks are therefore not merely an appendix to the analysis, but a necessary component in interpreting tail-loss estimates as decision-relevant quantities. Model risk in the far tail arises both from misspecification of the loss-generation mechanism, most importantly of dependence and severity, and from uncertainty in key inputs and calibration choices such as ρ and LGD, as well as from measurement noise in empirical PDs and constructed exposure weights. These uncertainties are amplified at high confidence levels because tail quantiles are driven by adverse systematic states in which default incidence and loss severity interact non-linearly.

A central modelling convenience in the ASRF baseline is Gaussian one-factor dependence with a fixed asset correlation ρ. Conditional on the systematic factor, defaults are independent, and the Gaussian specification implies asymptotic tail independence, so joint extremes do not become mechanically more likely solely because the systematic driver is extreme (McNeil et al., 2015). In credit portfolios, however, default clustering can strengthen in adverse states as common funding conditions tighten, balance-sheet constraints bind, and market liquidity deteriorates. A closely related concern is misspecification of the systematic driver itself. Heavier-tailed systematic factors, or equivalent copula specifications, can increase the probability of joint extreme default realisations relative to the Gaussian benchmark even when marginal PDs are held fixed. In the empirical analysis, this issue is addressed through heavier-tailed dependence robustness checks designed to assess whether the main conclusions remain stable beyond the Gaussian baseline.

Tail risk is shaped by loss severity as well as by default incidence. The decomposition

$$L(T) = \sum_{i=1}^{N} w_i \, \mathrm{LGD}_i \, \mathbf{1}\{\tau_i \leq T\}$$

makes explicit that tail loss depends on loss severity through LGD in addition to default clustering through (p, ρ). Treating loss given default as constant across time and states is analytically convenient, yet it is a strong empirical restriction. Recoveries are plausibly state dependent: adverse macro and market conditions associated with elevated default rates tend to coincide with depressed collateral and going-concern values and higher workout frictions, producing downturn recoveries and higher loss given default (Schuermann, 2004). This creates a natural wrong-way channel in which systematic stress increases both the incidence of default and the severity of loss, so a single unconditional LGD assumption can understate tail losses precisely in the states that matter for economic and regulatory capital (Basel Committee on Banking Supervision, 2005b).

A further limitation concerns the asymptotic granularity condition underpinning the ASRF representation. The approximation assumes a large, fine-grained portfolio in which idiosyncratic risk diversifies away, leaving the tail primarily driven by systematic risk (Gordy, 2003). Real portfolios can violate this condition due to name concentration, sector tilts, and dispersed exposure weights. In such settings, idiosyncratic default risk can remain relevant in the tail and widen the loss distribution even under an otherwise correct one-factor structure, implying that asymptotic tail approximations may understate extreme-loss risk in concentrated portfolios.

These limitations motivate a robustness design that focuses on a small set of economically interpretable stress channels rather than on a proliferation of alternative models. In this thesis, the main robustness dimensions are correlation sensitivity, sector multi-factor dependence, heavier-tailed dependence alternatives, LGD stress and state dependence, and checks beyond the asymptotic granularity assumption. Concretely, correlation sensitivity runs with higher ρ assess how strongly tail losses depend on default co-movement assumptions; LGD stress scenarios capture cyclical severity and potential co-movement between default incidence and recovery outcomes; heavier-tailed dependence alternatives assess the role of tail dependence beyond the Gaussian benchmark; and finite-N checks compare asymptotic tail estimates with Monte Carlo or discrete-default simulations under the same basic structure, supplemented by granularity and concentration experiments such as equal-weight and top-k portfolios (Basel Committee on Banking Supervision, 2005b; McNeil et al., 2015; Gordy and Lütkebohmert, 2013). Collectively, these exercises do not validate the baseline model in an absolute sense; rather, they quantify how sensitive the main conclusions are to empirically plausible deviations from the stylised ASRF assumptions and thereby provide a robustness envelope for tail-loss estimates.

2.5.7 Physical and risk-neutral probabilities

Throughout this thesis, default probabilities used for risk measurement are physical, real-world probabilities over the fixed one-year horizon defined in Section 2.1, rather than risk-neutral probabilities inferred from market prices. Loss severity is represented through LGD, and in sensitivity analysis the calibration may be shifted toward conservative or downturn levels to reflect state-dependent recoveries. This physical-probability perspective is natural for realised-loss measurement and economic capital, because it links portfolio loss outcomes to actual default frequencies and recovery outcomes rather than to market prices.

By contrast, the valuation of traded credit claims is conducted under the risk-neutral measure. For a one-year default event, one may write

$$p_i^{\mathbb{Q}} = \mathbb{Q}(\tau_i \leq 1), \qquad p_i^{\mathbb{P}} = \mathbb{P}(\tau_i \leq 1).$$

Alternatively, in reduced-form terms, default may be modelled through an intensity process. Under such specifications, risk-neutral default likelihoods are calibrated to market prices, such as CDS spreads and bond prices, conditional on recovery conventions, rather than obtained as a direct transformation of physical default probabilities (Duffie and Singleton, 2003; Lando, 1998). In particular, market-implied default probabilities are not identified without assumptions on recovery and risk premia. CDS calibration typically fixes a recovery rate, or equivalently an LGD convention, so changes in assumed recovery mechanically affect the implied risk-neutral default intensity $\lambda^{\mathbb{Q}}(t)$ and the corresponding $p_i^{\mathbb{Q}}$ even when observed spreads are unchanged (Duffie and Singleton, 2003; Hull, 2021).

The distinction between physical and risk-neutral default probabilities is economically important. Physical probabilities are appropriate for realised-loss measurement and capital-style tail-risk analysis, whereas risk-neutral probabilities are appropriate for pricing and hedging traded credit claims across maturities. In addition, observed CDS and bond spreads may reflect liquidity-related components as well as compensation for default risk, so market-implied default intensities cannot be interpreted as objective default probabilities without additional assumptions (Duffie and Singleton, 2003; Longstaff, Mithal and

Neis, 2005).

This separation aligns with the modelling distinction adopted in this thesis. The ASRF Vasicek framework is a one-period horizon loss model designed to map physical one-year PD inputs, dependence assumptions, and LGD severity into tail-loss measures relevant for capital and risk management. Reduced-form term-structure models, by contrast, are dynamic default-time models designed for pricing and hedging across maturities. The former is disciplined by real-world default likelihoods, whereas the latter is disciplined by market prices under $\mathbb{Q}$. Accordingly, risk-neutral quantities are introduced here only to clarify the conceptual distinction between risk measurement and market valuation and to avoid conflating market-implied default likelihoods with the physical PD inputs used in the empirical framework.

2.6 Synthesis and implications for the empirical design

The theoretical discussion in this chapter implies a concrete empirical design. Portfolio tail risk is measured over a one-year horizon using physical issuer-level default probabilities, instrument-level loss-severity assumptions, and an explicit dependence structure. The baseline engine is the Vasicek ASRF mapping, which provides a transparent link from $(p_i, \rho_i, \mathrm{LGD}_i, \mathrm{EAD}_i)$ inputs to expected loss and high-quantile tail loss.

These theoretical foundations motivate five central design choices in the empirical analysis. First, the analysis focuses on one-year tail loss rather than on mark-to-market repricing, which makes physical PDs the relevant probability inputs. Second, dependence assumptions are treated as a first-order driver of tail outcomes, which motivates the correlation and dependence robustness checks. Third, loss severity is modelled through seniority-based LGD assumptions and supplemented by conservative and stress-style LGD scenarios in order to reflect the possibility of lower recoveries in adverse states. Fourth, the asymptotic baseline is complemented by finite-portfolio diagnostics and Monte Carlo comparisons in order to assess the economic relevance of granularity and concentration effects. Fifth, because tail estimates depend materially on data construction and calibration choices, the empirical results are interpreted together with robustness checks that make model risk explicit.

More generally, the empirical implementation is interpreted as a disciplined approximation rather than as a complete model of credit crises. The purpose is not to claim that a single reduced-form framework can fully describe all dimensions of portfolio credit risk, but to show transparently how tail-loss estimates respond to plausible changes in dependence, severity, concentration, and input construction. In this sense, the empirical analysis is designed not only to produce a baseline estimate, but also to quantify the robustness envelope around that estimate.

3 Portfolio credit risk models in practice

Chapter 2 established the theoretical logic of fixed-horizon portfolio credit loss measurement and the one-factor tail mapping used in the empirical analysis. The purpose of the present chapter is narrower and more practical. Rather than extending the theory, it explains how portfolio credit risk models function as operational tools in institutional settings and why issues such as input construction, calibration philosophy, validation, and governance matter for interpreting the empirical results of this thesis.

The chapter therefore provides an institutional bridge between the theoretical framework and the empirical implementation. Its purpose is not to survey the entire credit risk modelling literature, but to clarify the practical setting in which tail-loss measures are used for capital assessment, internal limits, stress testing, and model-based risk management. In particular, it highlights why the empirical outputs reported later in the thesis should be interpreted as the results of a disciplined but stylised modelling pipeline rather than as frictionless mechanical estimates.

3.1 Model use in practice and the position of this thesis

At a general level, portfolio credit risk models translate granular exposure and credit information into portfolio-level loss metrics over a chosen horizon. This translation is valuable in practice because it supports the identification of concentrations, the assessment of diversification, and the use of loss distributions in capital planning, stress testing, and internal risk control (Gordy, 2003; Board of Governors of the Federal Reserve System, 2011).

A useful practical distinction is between default-mode models and migration-based or mark-to-market models (Gordy, 1998). Default-mode models focus on losses arising from default events over a fixed horizon and typically operate on inputs such as one-year PD, EAD, and LGD. They therefore align naturally with banking-book credit loss measurement and capital-style applications. Migration- or mark-to-market models, by contrast, translate rating changes and spread movements into changes in instrument values over the horizon and are more naturally suited to traded-credit portfolios in which repricing effects matter even without realised default.

The empirical framework of this thesis is intentionally positioned in the default-mode tradition. Its objective is to quantify realised tail loss over a one-year horizon rather than mark-to-market repricing. For this reason, the core inputs are physical one-year default probabilities, exposure proxies, and loss-severity assumptions, which are then combined with an explicit dependence structure to generate tail-loss measures. In this sense, the modelling goal of the thesis is closer to capital-oriented portfolio loss measurement than to valuation-oriented migration modelling.

3.2 Regulatory capital and the IRB benchmark

A central institutional benchmark for default-mode portfolio credit risk modelling is the internal ratings-based framework under Basel II and related supervisory rules. In this setting, the key modelling inputs are probability of default, exposure at default, and loss given default, and these are translated into regulatory capital requirements through supervisory risk-weight functions calibrated to a high confidence level (Basel Committee on Banking Supervision, 2005a, 2006). The analytical benchmark underlying this logic is closely related to the asymptotic single risk factor framework, while Gordy (2003) shows that portfolio invariance under value-at-risk essentially implies an asymptotic single-factor structure.

This benchmark is important for the interpretation of the present thesis even though the thesis does not estimate regulatory capital mechanically. The IRB framework clarifies why one-year PDs, dependence assumptions, and conservative LGD treatment are central when tail-loss measures are read through a capital-oriented lens. It also clarifies why asymptotic one-factor models remain influential in practice despite their stylised nature: they provide a transparent and operational mapping from obligor-level risk inputs to portfolio-level tail outcomes.

At the same time, the IRB benchmark is embedded in a broader regulatory architecture that includes supervisory review, disclosure, and stress testing. Where real portfolios deviate materially from the asymptotic fine-grained ideal, supervisory assessment is expected to consider concentration risk, stress results, and other institution-specific features not fully captured by the benchmark formula (Basel Committee on Banking Supervision, 2005a, 2006). This is directly relevant for the empirical design of this thesis, which supplements the asymptotic benchmark with finite-portfolio comparisons, concentration diagnostics, and robustness checks.

3.3 Implementation frictions, validation, and governance

In practice, the credibility of tail-loss metrics depends not only on model form but also on implementation quality. Even when the modelling core is theoretically coherent, empirical frictions arise because defaults are rare, recoveries are resolved with lags, exposures are observed through heterogeneous systems, and dependence must be inferred from limited data. These frictions can affect capital and risk decisions in economically meaningful ways and therefore belong to the interpretation of model outputs rather than to a purely technical appendix (Board of Governors of the Federal Reserve System, 2011).

A first practical friction is definitional and operational. Tail-loss models require a consistent default definition and a coherent portfolio object before PDs can be estimated, exposures aggregated, and outcomes compared with model predictions. In large-scale institutional settings, this is inseparable from risk data aggregation. Without reliable aggregation across business lines and legal entities, model inputs and realised outcomes may not refer to the same underlying portfolio, which weakens both measurement and validation (European Banking Authority, 2016; Basel Committee on Banking Supervision, 2013).

A second friction concerns dependence. In one-factor portfolio models, the tail is dominated by systematic states, so high-quantile loss measures are highly sensitive to the assumed strength and form of dependence. Yet dependence is difficult to estimate reliably from default data because defaults are rare and systemic stress episodes are infrequent, which makes parameter uncertainty material precisely where the tail is measured (Gordy and Heitfield, 2002). For this reason, dependence assumptions should be treated as a primary robustness dimension rather than as a fixed background parameter.

A third friction concerns loss severity. Recovery outcomes are heterogeneous across instruments and are plausibly cyclical, so LGD tends to be higher in states where defaults are also more clustered. This creates a practical wrong-way channel between default incidence and loss severity and implies that models using fixed unconditional LGD assumptions may understate downturn tail risk (Altman, Brady, Resti and Sironi, 2005; Frye, 2000; European Banking Authority, 2017).

A fourth friction concerns concentration. Asymptotic benchmarks abstract from undiversified idiosyncratic risk, yet real portfolios may contain large-name concentrations that remain economically important in the tail. This is why practical tail-risk analysis often supplements asymptotic benchmarks with concentration diagnostics, finite-portfolio comparisons, or granularity adjustments (Gordy and Lütkebohmert, 2013; Wilde, 2001; Emmer and Tasche, 2005).

These implementation frictions also shape validation practice. Because realised credit tail events are infrequent and recoveries are observed with delay, classical backtesting has limited power for high-quantile credit loss measures. Validation therefore tends to rely on a broader combination of benchmarking, sensitivity analysis, outcomes analysis, and stress testing, embedded in a governance framework that ensures documentation, effective challenge, and appropriate use (Board of Governors of the Federal Reserve System, 2011). In this sense, model governance is not external to tail-risk measurement; it is part of the conditions under which the output becomes decision-relevant.

3.4 PD, EAD, and LGD as practical modelling inputs

Probability of default, exposure at default, and loss given default form the empirical bridge between the conceptual loss framework and the concrete implementation choices made in this thesis. Although the terminology is standard, these inputs are not purely mechanical objects in practice. Their meaning depends on how default is defined, how exposures are aggregated, what horizon is used, and whether parameter estimates are intended to reflect average or downturn conditions.

For PD, the key practical issue is the construction of a coherent obligor-level default likelihood over the chosen horizon. In this thesis, the relevant object is the one-year physical probability of default, consistent with fixed-horizon realised-loss measurement rather than market-implied pricing probabilities. In practice, this requires not only a numerical PD estimate but also a consistent obligor definition and a default identification framework that allows comparisons across names and across time (European Banking Authority, 2016, 2017).

For EAD, the main practical issue is that observed exposure is often only a proxy for the amount truly at risk at default. While this problem is especially important for revolving facilities and commitments, it also arises whenever empirical work must approximate exposure using available instrument information rather than contractual default-time balances. In this thesis, exposure is represented by an issuance-based proxy derived from instrument-level bond and note data. The resulting tail-loss measures should therefore be interpreted conditional on that exposure proxy rather than as estimates of full contractual exposure dynamics.

For LGD, the practical issue is that loss severity depends both on cross-sectional protection features and on the economic state. Seniority and collateralisation are among the most important observable determinants of recovery, but realised recovery also depends on enforcement, workout costs, resolution timing, and cyclical conditions. In the present thesis, LGD is operationalised through a transparent seniority-based mapping, with robustness checks designed to assess how strongly the results depend on conservative or downturn-style loss-severity assumptions.

Taken together, these considerations explain why the empirical implementation later in the thesis places particular emphasis on data construction, representativeness, concentration, and robustness. The PD, EAD, and LGD inputs used in the model are not treated as frictionless primitives, but as practical modelling choices that must be made explicit in order for the resulting tail-loss estimates to be interpreted responsibly.

3.5 Implications for the empirical chapters

The practical discussion in this chapter has a direct implication for the way the empirical results should be read. The tail-loss estimates reported later are not simply outputs of a theoretical mapping; they are outputs of a modelling pipeline that combines a one-year default-mode framework, an asymptotic benchmark, empirical input construction, and explicit robustness analysis. This is why the later empirical chapters do not report only a single baseline number, but also concentration diagnostics, dependence sensitivities, LGD stress scenarios, and finite-portfolio comparisons.

Accordingly, the role of the empirical analysis is not to claim that one stylised model fully captures all dimensions of portfolio credit risk. Its role is to show transparently how tail-loss estimates change when key practical modelling choices are varied within a coherent framework. In that sense, the usefulness of the model lies not only in the baseline estimate it produces, but also in the robustness envelope it provides around that estimate.

4 Research Data and Implementation Methods

4.1 Data sources and sample construction

The baseline portfolio construction combines two LSEG exports retrieved in February 2026 and interpreted as a point-in-time snapshot of vendor coverage and classifications at the download date. The first file is an issuer-level default probability dataset for U.S.-headquartered corporate issuers. It reports one-year physical default probabilities through the Credit Combined PD (%) field and includes TRBC-based sector classifications. The second file is an instrument-level bond dataset for U.S.-incorporated corporate issuers. It contains instrument descriptors such as instrument type, seniority, and bond grade, together with an issuance-based size field reported as Amount Issued in euros.

The baseline empirical universe consists only of debt instruments. Concretely, the bond file is restricted to instruments classified as bond or note, and non-target instrument classes such as preferred-type instruments are excluded from the baseline universe. This yields a baseline universe with a clear interpretation as straight corporate debt exposure. In the present dataset, this debt-only filter retains approximately 96.1% of the total issuance-based euro amount in the raw bond file.

Because the portfolio model is defined at the obligor level, the instrument-level bond data are aggregated to the issuer level before the portfolio loss calculations are performed. This step is necessary for model coherence. Multiple instruments issued by the same firm do not represent separate default events in an obligor-default framework. Treating them as separate names would therefore mechanically overstate portfolio granularity and introduce spurious within-issuer diversification.

Issuer matching across the two files is performed using a standardised name key. Company and issuer names are converted to uppercase, common legal suffixes are removed, non-alphanumeric characters are stripped, and repeated spaces are collapsed. The resulting key is then used to merge issuer-level PD information with aggregated issuer-level bond exposures. Since name-based matching can affect both sample size and representativeness, the resulting coverage and selection patterns are analysed explicitly in Section 5.

An additional source of sample selectivity arises from the fact that the two vendor extracts are not defined over identical underlying universes: the PD file covers U.S.-headquartered corporate issuers, whereas the bond file covers U.S.-incorporated corporate issuers. The matched sample is therefore shaped not only by name-key matching and missing values, but also by this difference in source-universe definitions.

The final working portfolio is the intersection of the processed PD universe and the debt-only issuer universe under the adopted matching protocol and the underlying source-universe definitions. After filtering and merging, the baseline sample contains $N = 765$ issuers with non-missing one-year PDs and positive issuance-based exposure. This matched issuer sample is the portfolio used in the baseline Vasicek implementation and in the subsequent sensitivity, robustness, and model-extension exercises.

In addition to the matched issuer portfolio constructed from the two LSEG exports, the model-extension analysis also uses an auxiliary external dependence proxy panel. This panel consists of weekly adjusted return series for 11 U.S. sector ETFs and is used only to diagnose whether a low-dimensional multi-factor dependence structure is empirically plausible in market-based sector data. The proxy panel does not alter the issuer-level portfolio construction, PD inputs, exposure aggregation, or LGD mapping. Instead, it provides an external diagnostic for assessing whether the one-factor restriction is a reasonable first-order approximation or whether sector-level co-movement appears to contain additional common components. Because the full 11-sector overlap is constrained by the shorter histories of some ETFs, the common weekly sample used in the PCA diagnostic runs from 6 July 2018 to 2 April 2026.

4.2 Issuer-level model inputs and baseline implementation

This subsection defines how the theoretical and practical framework developed in Chapters 2 and 3 is parameterised in the empirical implementation. For each issuer $i = 1, \dots, N$, the baseline model requires a one-year probability of default p_i, an exposure proxy EAD_i, a loss-given-default parameter LGD_i, and the corresponding portfolio weight

$$w_i = \frac{\mathrm{EAD}_i}{\mathrm{EAD}}.$$

These issuer-level inputs are then mapped into portfolio loss measures under the one-factor Vasicek framework.

Default probabilities are taken from the vendor's one-year Credit Combined PD (%) field and converted from percent units into fractions $p_i \in (0, 1)$. When multiple PD observations exist for the same standardised issuer key, the baseline retains the maximum value. This choice is conservative in a tail-risk setting, because understating issuer-level default intensity can bias extreme portfolio loss estimates downward.

Exposures are proxied using the bond-file field Amount Issued in euros. Let EAD_{ij} denote the issuance-based exposure proxy for instrument j of issuer i. For each issuer i, the issuer-level exposure proxy is then constructed by aggregating issuance amounts across the set of debt instruments $j \in \mathcal{J}(i)$:

$$\mathrm{EAD}_i = \sum_{j \in \mathcal{J}(i)} \mathrm{EAD}_{ij}.$$

This construction yields an internally consistent issuer-level size measure and, therefore, a transparent set of portfolio weights. At the same time, it should be interpreted as an issuance-based proxy rather than as contractual exposure at default, current holdings, or a regulatory EAD measure.

Loss given default is not observed directly in the vendor files. Instead, instrument-level seniority labels are mapped into a transparent LGD grid and then aggregated to issuer level using issuance-based weights. For issuer i,

$$\mathrm{LGD}_i = \frac{\sum_{j \in \mathcal{J}(i)} \mathrm{LGD}_{ij}\, \mathrm{EAD}_{ij}}{\sum_{j \in \mathcal{J}(i)} \mathrm{EAD}_{ij}}.$$

If seniority text is missing or cannot be mapped reliably, the fallback instrument-level assumption is $\mathrm{LGD}_{ij} = 60\%$, corresponding to a conventional senior-unsecured benchmark. The seniority-based LGD mapping used in the baseline implementation is reported in Table 1.

Table 1: LGD mapping from seniority text used in the baseline implementation.

Seniority text (examples)	**LGD used**
First lien / first-priority secured	35%
Senior secured or second lien	45%
Senior unsecured	60%
Subordinated, hybrid Tier 2 debt	75%
Junior, junior subordinated	80%
Missing or ambiguous	60%

Systematic default dependence is captured through the standard one-factor Vasicek framework. In the baseline specification, issuer asset correlations are assigned by applying the Basel IRB-style corporate correlation curve to the issuer-level PDs,

$$\rho(p) = 0.12 \cdot \frac{1 - e^{-50p}}{1 - e^{-50}} + 0.24 \times \left(1 - \frac{1 - e^{-50p}}{1 - e^{-50}}\right),$$

which maps lower PDs into higher asset correlations and yields values in the interval $[0.12, 0.24]$. This provides a transparent benchmark without requiring additional market-implied dependence inputs. Alternative dependence assumptions, including constant-correlation cases, are treated explicitly as sensitivity analysis rather than folded into the baseline specification.

To complement the economically specified dependence extensions, an auxiliary principal components analysis (PCA) is applied to the standardised weekly return series of the 11-sector proxy panel. The objective is not to replace the baseline issuer-level loss model with a separate PCA-calibrated benchmark, nor to re-estimate issuer-level PDs, exposures, or LGDs from market data. Rather, the PCA is used as an external dependence diagnostic: it indicates whether the historical sector proxy panel is largely one-dimensional or whether a small number of additional common components contributes materially to cross-sector co-movement. In this sense, the PCA results are used to motivate and interpret the sector multi-factor simulations in Section 5, rather than to define a separate primary portfolio model.

Portfolio losses are reported primarily as loss rates, that is, as fractions of total portfolio exposure. The main tail-risk measure is $\mathrm{VaR}_{0.999}$, computed at confidence level $\alpha = 0.999$. EL and $\mathrm{UL}_{0.999}$ are reported alongside this benchmark, and selected extensions also consider simulation-based tail measures and higher-tail summary quantities. Reporting results in loss-rate units is especially appropriate here because the exposure measure is issuance-based rather than institution-specific regulatory EAD.

The empirical baseline is therefore defined by five choices: a debt-only issuer universe constructed from bonds and notes, issuer-level maximum one-year PDs, issuance-based issuer exposures, seniority-mapped issuer LGDs, and the IRB-style correlation function. Deviations from this setup, such as constant-correlation scenarios, downturn-LGD stress, finite-portfolio simulation, multi-factor extensions, t-copula dependence, and state-dependent LGD, are treated explicitly as sensitivity, robustness, or model-extension exercises in Section 5. In addition, an auxiliary PCA-based sector proxy diagnostic is used to assess whether a low-dimensional multi-factor dependence structure is empirically supported by market-based sector data. Table 2 summarises the key baseline inputs, modelling choices, and auxiliary dependence diagnostics used in the empirical implementation.

All cleaning rules and transformations—name-key normalisation, EU-number parsing, PD scaling, debt-only filtering, exposure aggregation, and LGD mapping—are implemented in R using reproducible scripts. The resulting issuer-level inputs are exported into intermediate analysis tables before the portfolio loss calculations are performed. The empirical design is therefore intentionally transparent: a matched issuer portfolio is constructed from vendor data, transformed into issuer-level PD, EAD, LGD, and weight inputs, and then analysed under a clearly documented one-factor credit-risk specification.

The empirical design also imposes clear limitations. The PDs are one-year physical probabilities rather than risk-neutral measures. The exposure proxy is based on issuance amounts rather than outstanding balances or institution-specific exposures. The LGD

Table 2: Summary of key inputs and baseline settings in the empirical implementation.

Item	**Baseline setting**
Empirical unit	Issuer
Default probability horizon	One-year physical PD from vendor export
PD aggregation	Maximum PD per issuer key if duplicates exist
Exposure definition	Issuance-based amount issued in euros aggregated to issuer
Portfolio weights	$w_i = \text{EAD}_i/\text{EAD}$
LGD	Seniority-based mapping; fallback 60% if missing or ambiguous
Baseline dependence	Basel IRB corporate curve $\rho(p)$
Correlation sensitivities	Constant $\rho = 0.12$ and $\rho = 0.24$
Main tail measure	$\text{VaR}_{0.999}$
Baseline universe	Debt-only (bonds and notes)
External dependence proxy	Weekly adjusted returns for 11 U.S. sector ETFs
PCA diagnostic	Standardised weekly sector returns, 6 Jul 2018–2 Apr 2026 ($T = 405$)

mapping is stylised and depends on the adopted recovery grid. Finally, the matched sample is selective because it is defined by the overlap of the two vendor datasets under the chosen matching protocol. These limitations do not invalidate the analysis, but they define how the results should be interpreted: as internally consistent tail-loss estimates for a constructed issuer portfolio rather than as direct forecasts for a specific investor or bank balance sheet. In addition, the PCA-based dependence diagnostic relies on market-based sector ETF proxies rather than on a historical panel of issuer-level default or spread data, so it should be interpreted as supportive evidence on low-dimensional co-movement rather than as a direct estimate of default dependence.

5 Results

5.1 Coverage and representativeness

The baseline modelling portfolio is economically meaningful but clearly selective relative to the underlying source universes. Because the empirical analysis combines an issuer-level PD file with an instrument-level debt dataset, the baseline sample is defined by the intersection of the two sources after issuer aggregation and name-key matching. Table 3 summarises both the coverage of the matched debt-only portfolio and the main selection patterns relative to unmatched observations.

The matched debt-only portfolio contains 765 issuers and retains 58.0% of the issuance-based exposure in the debt-only bond universe. The sample is therefore large enough to support meaningful portfolio tail-loss analysis, but it should not be interpreted as a full representation of the U.S. corporate debt market. Instead, all subsequent results describe a matched issuer portfolio constructed from the overlap between the PD dataset and the debt-only bond universe.

Table 3 also shows that the matched sample differs systematically from the unmatched observations. Relative to unmatched issuers in the PD dataset, matched issuers exhibit lower default risk, with both mean and median PD materially below those of the unmatched group. Relative to unmatched issuers in the debt-only bond universe, matched issuers are larger in exposure terms, as indicated by the substantially higher median issuer EAD. In addition, the matched debt sample is materially less secured than the unmatched bond universe, while subordinated exposure remains economically small in both groups.

These patterns matter for interpretation. The reported tail-loss estimates describe a portfolio tilted towards larger, lower-PD, and more unsecured issuers. The main empirical value of the analysis therefore lies in internally consistent comparisons across model specifications within this matched sample, rather than in claiming that the baseline portfolio is fully representative of the broader debt market.

5.2 Baseline portfolio characteristics

The matched issuer portfolio is moderately concentrated, exhibits a strongly right-skewed distribution of issuer default risk, and has a portfolio-level LGD profile close to conventional senior-unsecured debt. These baseline characteristics are important because they shape how the one-factor framework maps issuer-level inputs into aggregate tail-loss outcomes.

Table 3: Coverage and representativeness of the matched debt-only portfolio.

Panel	Metric	Value
A. Join coverage		
	PD issuers (unique)	4,461
	Debt-only bond issuers (unique)	1,674
	Joined issuers (unique)	765
	Issuer match rate vs PD universe	17.15%
	Issuer match rate vs debt-only bond universe	45.7%
	Total EAD in debt-only bond universe	EUR 4.24×10^{12}
	Total EAD retained after join	EUR 2.46×10^{12}
	EAD retention rate	58.0%
	EAD lost due to join	EUR 1.78×10^{12}
B. Representativeness diagnostics		
	Matched vs unmatched PD mean	0.514% vs 1.563%
	Matched vs unmatched PD median	0.152% vs 0.371%
	Matched vs unmatched median issuer EAD	EUR 1,604.9m vs EUR 867.5m
	Matched vs unmatched secured EAD share	5.83% vs 26.61%
	Matched vs unmatched subordinated EAD share	0.12% vs 0.57%

Table 4 summarises the main descriptive features of the matched debt-only issuer portfolio. The portfolio contains $N = 765$ issuers and total issuance-based exposure of approximately EUR 2.46×10^{12}. Exposure is not distributed uniformly across issuers, yet the portfolio is not dominated by only a handful of names. The ten largest issuers account for 14.49% of total exposure, and the Herfindahl–Hirschman index equals 0.00479, corresponding to an effective number of names of roughly $N_{\text{eff}} = 1/\text{HHI} \approx 209$. Concentration is therefore economically relevant for tail-risk measurement, but not yet extreme in the baseline sample.

Default risk is strongly right-skewed across issuers. In unweighted terms, the one-year PD distribution has a median of 0.152%, a mean of 0.514%, and a 99th percentile of 7.80%. Because the portfolio model aggregates losses using exposure weights, the exposure-weighted PD summaries are more informative for tail-loss measurement. The EAD-weighted mean PD is substantially lower, at 0.294722%, and the corresponding EAD-weighted median is 0.119%. This indicates that the largest exposures in the matched sample are, on average, safer than the typical issuer in the portfolio.

Loss severity is comparatively stable at portfolio level. The issuer-level mapped LGD distribution has an unweighted mean of 58.5%, an unweighted median of 60.0%, and an EAD-weighted mean of 58.85%. This concentration around the 60% benchmark is consistent with the composition of the underlying debt-only instrument universe, which is

Table 4: Baseline characteristics of the matched debt-only issuer portfolio.

Panel	**Metric**	**Value**
A. Portfolio size and concentration		
	Number of issuers	765
	Total EAD	EUR 2.46×10^{12}
	Mean issuer EAD	EUR 3,217m
	Median issuer EAD	EUR 1,605m
	Top10	14.49%
	HHI of issuer weights	0.00479
	Effective number of names N_{eff}	209
B. PD distribution (1-year, %)		
	Unweighted mean	0.514%
	Unweighted median	0.152%
	Unweighted 99th percentile	7.80%
	EAD-weighted mean	0.294722%
	EAD-weighted median	0.119%
	EAD-weighted 99th percentile	3.71%
C. LGD distribution (mapped, %)		
	Unweighted mean	58.5%
	Unweighted median	60.0%
	EAD-weighted mean	58.85%
	EAD-weighted median	60.0%
D. Debt-only instrument mix (pre-join, EAD *share)*		
	Unsecured / other	85.1%
	Secured	14.6%
	Subordinated / junior	0.306%

economically dominated by unsecured instruments. Secured debt forms a meaningful but still clearly smaller share, while subordinated and junior instruments account for only a negligible proportion of issuance-based exposure.

Taken together, the baseline portfolio combines moderate name concentration, a skewed cross-section of issuer default risk, and a relatively stable LGD structure centred near conventional senior-unsecured debt. These characteristics provide the economic backdrop for the next subsection, which examines how the baseline portfolio inputs translate into expected and tail loss under the Vasicek framework.

5.3 Baseline tail loss estimates

The baseline Vasicek specification produces a pronounced gap between average and extreme loss, indicating that portfolio credit risk in the matched sample is primarily a tail phenomenon rather than an expected-loss problem. Baseline results are computed using the matched debt-only issuer portfolio together with the IRB-style correlation function and the mapped issuer-level LGD specification.

Under this baseline setup, the portfolio expected loss equals 0.165% of total exposure, while $\mathrm{VaR}_{0.999}$ equals 3.02%. The corresponding unexpected-loss component, defined as $\mathrm{UL}_{0.999} = \mathrm{VaR}_{0.999} - \mathrm{EL}$, is 2.85%. For completeness, the higher-quantile expected-shortfall proxy reported by the pipeline equals 4.04%. All figures are reported primarily as loss rates relative to total portfolio exposure, because the exposure measure is based on issuance amounts rather than on regulatory EAD measured for capital purposes.

The economic implication is straightforward. Average credit loss in the constructed portfolio remains small, but losses become materially larger in extreme systematic states. In other words, the baseline one-factor framework implies that the main risk-management challenge is not the unconditional mean loss rate, but the magnitude of clustered losses in the far tail. This baseline result provides the natural reference point for the sensitivity, finite-portfolio, and model-risk extensions analysed in the following subsections.

5.4 Baseline sensitivities: dependence and loss severity

Within the baseline one-factor framework, tail-loss estimates are more sensitive to the dependence assumption than to simple level shifts in unconditional LGD, although stressed LGD still increases tail losses materially. The purpose of this subsection is therefore to separate the systematic-clustering channel from the loss-severity channel by varying correlation and LGD assumptions around the baseline specification.

Table 5 summarises the baseline specification together with the main dependence and severity sensitivity scenarios. A clear pattern emerges. Changes in correlation have the largest effect on the tail-loss quantile, whereas changes in unconditional LGD mainly rescale loss severity. Downturn-style LGD stress scenarios provide an intermediate case: they leave the dependence structure unchanged, but still widen the tail materially because loss severity increases precisely in adverse states.

The correlation sensitivity results show that tail loss is highly responsive to the assumed degree of common default dependence. Holding the mapped LGD specification fixed, lowering correlation to a constant $\rho = 12\%$ reduces $\mathrm{VaR}_{0.999}$ from 3.02% in the baseline to

Table 5: Baseline sensitivity analysis: dependence and loss severity. All loss measures are reported as percentages of total portfolio exposure.

Scenario group	**Scenario**	$\mathrm{LGD}_{\mathrm{wtd}}$	ρ_{wtd}	EL	$\mathrm{VaR}_{\mathbf{0.999}}$	$\mathrm{UL}_{\mathbf{0.999}}$
Baseline	IRB-style $\rho(p_i)$ + mapped LGD	58.9%	22.7%	0.165%	3.02%	2.85%
Correlation sensitivity	Constant $\rho = 12\%$	58.9%	12.0%	0.165%	1.68%	1.52%
Correlation sensitivity	Constant $\rho = 24\%$	58.9%	24.0%	0.165%	3.45%	3.28%
LGD sensitivity	Constant LGD = 45%	45.0%	22.7%	0.133%	2.36%	2.23%
LGD sensitivity	Constant LGD = 60%	60.0%	22.7%	0.177%	3.15%	2.97%
Downturn-LGD stress	Mapped LGD ×1.15	67.7%	22.7%	0.190%	3.47%	3.28%
Downturn-LGD stress	Mapped LGD ×1.25	73.6%	22.7%	0.207%	3.77%	3.56%

1.68%, while increasing correlation to $\rho = 24\%$ raises it to 3.45%. Expected loss remains unchanged at 0.165% across these scenarios because the PD and LGD inputs are held fixed. The main effect of correlation is therefore not on the average loss level, but on the thickness of the loss tail generated by systematic default clustering.

Changing LGD while holding the IRB-style correlation structure fixed has a different interpretation. Relative to the mapped-LGD baseline with $\mathrm{LGD}_{\mathrm{wtd}} = 58.9\%$, imposing a constant LGD of 45% lowers expected loss from 0.165% to 0.133% and reduces $\mathrm{VaR}_{0.999}$ from 3.02% to 2.36%. By contrast, a constant LGD of 60% increases expected loss to 0.177% and raises $\mathrm{VaR}_{0.999}$ to 3.15%. These results indicate that unconditional LGD changes mainly act as a severity scaler: they shift both expected and tail losses, but do not alter the dependence structure that governs default clustering.

A more conservative severity view emerges in the downturn-LGD stress scenarios. When the mapped LGD is multiplied by 1.15, the portfolio-weighted LGD rises from 58.9% to 67.7% and $\mathrm{VaR}_{0.999}$ increases to 3.47%. Under the 1.25 stress, the weighted LGD rises further to 73.6% and $\mathrm{VaR}_{0.999}$ reaches 3.77%. The corresponding unexpected-loss measure rises from 2.85% in the baseline to 3.28% and 3.56%, respectively. These stress scenarios therefore show that even without changing the dependence assumption, worsening recovery conditions can widen the loss tail in economically meaningful ways.

Taken together, the sensitivity analysis suggests that correlation is the dominant driver of tail amplification within the standard one-factor setup, while LGD primarily scales the level of loss once defaults occur. At the same time, the downturn-LGD results show that recovery assumptions remain materially important for conservative tail-risk measurement, especially when loss severity is allowed to deteriorate in stress conditions rather than remain fixed at its unconditional level.

Table 6: Finite-portfolio effects under alternative granularity and concentration structures. All VaR measures are reported as percentages of total portfolio exposure.

Portfolio specification	N	**Top-10 weight**	**HHI**	$\mathrm{VaR}^{\mathbf{ASRF}}_{\mathbf{0.999}}$	$\mathrm{VaR}^{\mathbf{MC}}_{\mathbf{0.999}}$	**Gap**
Baseline (original weights)	765	14.49%	0.00479	3.02%	3.20%	0.185 pp
Equal-weight (more granular)	765	1.31%	0.00131	3.89%	3.86%	-0.030 pp
Top-50 only (more concentrated)	50	38.6%	0.0267	2.30%	4.42%	2.12 pp
Top-10 only (very concentrated)	10	100.0%	0.112	2.06%	9.49%	7.43 pp

5.5 Finite-portfolio effects and concentration

Finite-portfolio effects are modest in the baseline portfolio but become economically large when exposure concentration increases. This subsection compares the analytical ASRF approximation with Monte Carlo simulation under the same one-factor structure in order to isolate the effect of discrete defaults and imperfect granularity.

Table 6 summarises the comparison across the baseline portfolio and three alternative granularity or concentration specifications. A clear pattern emerges. When the portfolio is sufficiently granular, the ASRF approximation is very close to the finite Monte Carlo benchmark. When exposure is concentrated in fewer names, however, the gap between the two measures widens sharply.

In the baseline portfolio, the ASRF approximation yields $\mathrm{VaR}_{0.999} = 3.02\%$, whereas the corresponding one-factor Monte Carlo estimate is 3.20%. The implied finite-portfolio correction is therefore 0.185 percentage points. Expected loss is essentially unchanged, with the Monte Carlo mean loss equal to 0.164% versus 0.165% in the analytical baseline. The baseline evidence thus suggests that the asymptotic approximation remains informative in a moderately concentrated portfolio, but still understates the tail once discrete issuer defaults are explicitly simulated.

The role of granularity becomes clearer when the portfolio weight structure is varied. Under an equal-weight specification, concentration falls sharply, with the top-10 share declining to 1.31% and the HHI to 0.00131. In that case, the ASRF and Monte Carlo tail measures are almost identical: $\mathrm{VaR}_{0.999}$ equals 3.89% under ASRF and 3.86% under Monte Carlo. This indicates that once the portfolio becomes sufficiently fine-grained, the asymptotic approximation tracks the finite-default benchmark very closely.

By contrast, when the portfolio is restricted to the 50 largest issuers and the weights are renormalised, concentration rises materially, with a top-10 share of 38.6% and an HHI of 0.0267. The corresponding ASRF $\mathrm{VaR}_{0.999}$ is 2.30%, while Monte Carlo $\mathrm{VaR}_{0.999}$ rises to 4.42%, implying a gap of 2.12 percentage points. Under the extreme top-10 specification, where the portfolio is fully concentrated in ten names and the HHI rises to 0.112, the difference becomes dramatic: ASRF $\mathrm{VaR}_{0.999}$ is 2.06%, whereas Monte Carlo $\mathrm{VaR}_{0.999}$ reaches 9.49%, implying a gap of 7.43 percentage points.

The interpretation is that the ASRF framework measures tail loss through the conditional mean portfolio loss given the systematic factor, which is appropriate when idiosyncratic risk is well diversified. In finite and concentrated portfolios, however, large-name default realisations remain discrete and economically important even in the tail, so the asymptotic approximation can materially understate extreme loss. The results therefore suggest that the ASRF benchmark is most reliable in sufficiently granular portfolios, whereas concentrated portfolios require explicit finite-default treatment, concentration stress testing, or other granularity adjustments.

Taken together, the finite-portfolio results reinforce the broader message of the chapter. In the baseline matched portfolio, asymptotic tail measurement provides a useful first-order approximation, but its accuracy depends materially on the degree of concentration. Once granularity weakens, the distinction between conditional mean loss and the actual discrete loss distribution becomes economically important, and the analytical ASRF benchmark can cease to be a reliable proxy for extreme portfolio risk.

5.6 Beyond the Gaussian one-factor benchmark

Moving beyond the Gaussian one-factor benchmark materially changes the estimated loss tail. The baseline Vasicek specification remains useful as a transparent reference point, but the extended simulations show that tail-loss estimates are highly sensitive to how systematic dependence and loss severity are modelled. This subsection therefore examines three departures from the benchmark: a Gaussian sector multi-factor structure, heavy-tailed dependence through a t-copula, and state-dependent LGD.

Before turning to the simulated loss outcomes, it is useful to ask whether a low-dimensional multi-factor dependence structure is also supported by external market-based proxy data rather than only by economic intuition. To address this question, a historical weekly sector proxy panel is constructed from 11 U.S. sector ETFs, and principal components analysis (PCA) is applied to the standardised sector return series over the common sample from 6 July 2018 to 2 April 2026, yielding $T = 405$ weekly observations. The first principal component explains 67.8% of total standardised return variation, while the first two and three components explain 76.4% and 84.1%, respectively. The evidence therefore suggests that dependence is dominated by a broad common market component, but is not fully summarised by a single factor. The PCA exercise is used here as an external dependence diagnostic rather than as a separate benchmark model in the loss simulations.

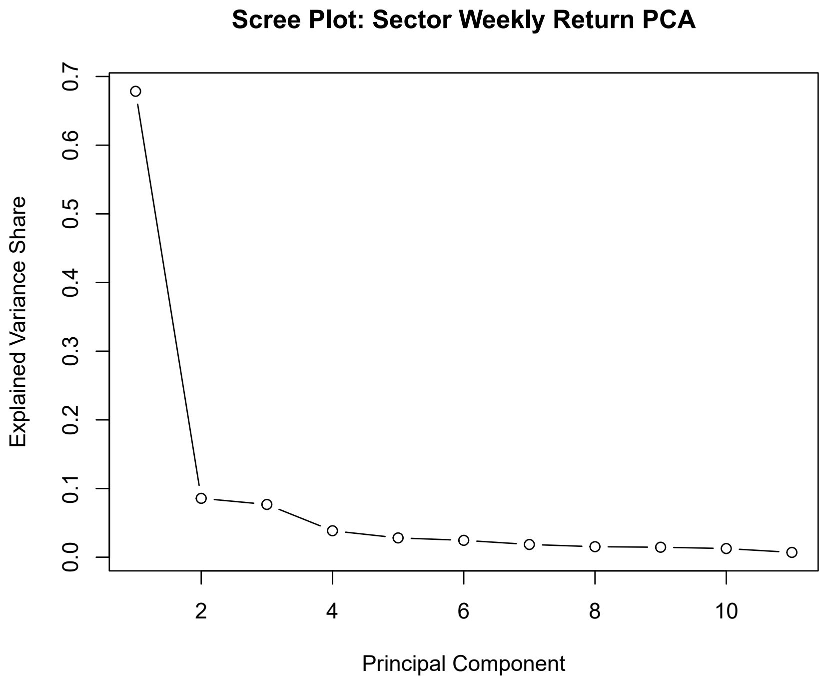

Figure 1: Scree plot for the PCA of the 11-sector weekly proxy return panel. The first principal component explains the largest share of total standardised return variation, while the second and third components remain non-trivial. This pattern is consistent with one dominant common component, but not with a purely one-dimensional dependence structure.

It nevertheless provides empirical support for relaxing the single-factor restriction and for interpreting the sector multi-factor simulations reported below as a low-dimensional dependence extension rather than as a purely ad hoc specification.

As Figure 1 shows, the first principal component clearly dominates the sector proxy panel, but the drop after PC1 is not large enough to treat sector-level dependence as purely one-dimensional. This visual evidence reinforces the interpretation that the Gaussian one-factor benchmark remains a useful first-order approximation, while low-dimensional multi-factor extensions are also empirically plausible.

Table 7 summarises the main extensions. The results show that the Gaussian one-factor benchmark is neither uniformly conservative nor uniformly non-conservative. Rather, both the direction and the magnitude of the change depend on which modelling restriction is relaxed. Sector segmentation lowers tail loss relative to a pure one-factor Gaussian benchmark, whereas heavy-tailed dependence and state-dependent LGD both increase it materially. Because the panels are based on separate simulation blocks, the Gaussian benchmark values differ slightly across panels; these small differences should be interpreted as Monte Carlo noise rather than as economically meaningful model differences.

Table 7: Extensions beyond the Gaussian one-factor benchmark. All tail measures are reported as percentages of total portfolio exposure.

Panel	**Scenario**	**Key parameter(s)**	$\text{VaR}_{\mathbf{0.999}}$
A. Gaussian sector multi-factor structure			
	One-factor MC benchmark	$global\_share = 1.00$	3.26%
	Sector MF MC	$global\_share = 0.75$	2.65%
	Sector MF MC	$global\_share = 0.50$	2.24%
	Sector MF MC	$global\_share = 0.25$	1.89%
	Two-factor MC (PC1-calibrated)	$global\_share = 0.678$	2.48%
	Three-factor MC (per-issuer PCA)	PCA-based ($PC1, PC2, PC3$) loadings	2.26%
B. Heavy-tailed dependence (t-copula)			
	Gaussian one-factor MC benchmark	Gaussian	3.20%
	t-copula one-factor	$\nu = 10$	7.47%
	t-copula one-factor	$\nu = 5$	11.3%
	t-copula one-factor	$\nu = 3$	14.6%
	Gaussian sector MF benchmark	$global\_share = 0.50$	2.24%
	t-copula sector MF	$global\_share = 0.50,\ \nu = 5$	9.55%
C. State-dependent LGD in sector multi-factor MC			
	Constant-LGD benchmark	$global\_share = 0.50$	2.24%
	State-LGD, mild	$\lambda_g = 0.25,\ \lambda_s = 0.10$	3.53%
	State-LGD, medium	$\lambda_g = 0.50,\ \lambda_s = 0.20$	3.67%
	State-LGD, strong	$\lambda_g = 0.75,\ \lambda_s = 0.30$	3.69%

The Gaussian sector multi-factor simulations indicate that tail losses can fall relative to the pure one-factor benchmark when systematic dependence is distributed across sector-specific factors rather than concentrated entirely in a single common driver. In Panel A, the one-factor Monte Carlo benchmark yields $\text{VaR}_{0.999} = 3.26\%$. Under the sector multi-factor specification, $\text{VaR}_{0.999}$ declines to 2.65% when the common factor remains dominant ($global\_share = 0.75$), to 2.24% in the balanced benchmark case ($global\_share = 0.50$), and to 1.89% when sector segmentation is strongest ($global\_share = 0.25$). The implication is that the Gaussian one-factor model can overstate portfolio-wide tail loss when it forces all systematic dependence to load on a single common factor, whereas a sector structure allows part of the systematic variation to remain segmented rather than perfectly market-wide. This interpretation is consistent with the PCA evidence above: one common component is clearly dominant, but the dependence structure is not purely one-dimensional.

To assess whether the ad hoc `global_share` values reported in Panel A are empirically defensible, two PCA-calibrated specifications are also considered. In the first, the two-factor structure is retained, but `global_share` is set equal to the PC1 variance share of

0.678, that is, the dominant common component identified in the sector proxy panel. In the second, each sector inherits a unit-norm direction vector in $(PC1, PC2, PC3)$ from the PCA loadings, and each issuer's factor-loading vector is obtained by scaling this direction by $\sqrt{\rho_i}$, so that the IRB asset correlation is preserved. The PC1-calibrated two-factor specification produces $\text{VaR}_{0.999} = 2.48\%$, which lies between the ad hoc `global_share` $=$ 0.75 and `global_share` $= 0.50$ benchmarks, close to where linear interpolation would place a 0.678 value. The three-factor specification yields $\text{VaR}_{0.999} = 2.26\%$, essentially coinciding with the balanced `global_share` $= 0.50$ case. Taken together, these results suggest that the earlier ad hoc multi-factor range is broadly consistent with the external PCA evidence.

The dependence results change sharply once heavy-tailed joint extremes are allowed. In Panel B, the Gaussian one-factor Monte Carlo benchmark equals 3.20%, but the t-copula one-factor specification raises $\text{VaR}_{0.999}$ to 7.47% for $\nu = 10$, 11.3% for $\nu = 5$, and 14.6% for $\nu = 3$. A similar pattern appears in the sector multi-factor setting: the Gaussian sector model with $global\_share = 0.50$ produces $\text{VaR}_{0.999} = 2.24\%$, whereas the corresponding t-copula specification raises it to 9.55%. These results show that the dependence assumption is a first-order model-risk channel. In particular, the Gaussian benchmark appears far less conservative once tail dependence is introduced, because joint extreme credit outcomes become substantially more likely.

A related amplification mechanism appears when LGD is allowed to worsen endogenously in adverse systematic states. In the constant-LGD sector multi-factor benchmark, $\text{VaR}_{0.999}$ equals 2.24%. Under state-dependent LGD, $\text{VaR}_{0.999}$ rises to 3.53% in the mild scenario, 3.67% in the medium scenario, and 3.69% in the strong scenario. In the medium scenario, the average LGD on defaulted exposure rises to 73.4%, while tail LGD on default approaches 94.7%. This implies that tail loss is shaped not only by the probability of clustered defaults, but also by the fact that loss severity itself may become materially worse in stressed states.

Taken together, the extended simulations substantially widen the range of economically plausible tail outcomes around the Gaussian one-factor benchmark. The Gaussian one-factor model therefore remains valuable as a disciplined and interpretable baseline, but not as a uniquely robust estimate of extreme-loss risk. In the present sample, the modelling of dependence and state-contingent severity matters at least as much as, and in some cases more than, the baseline parameter choices studied earlier.

Table 8: Top-10 issuers by baseline tail-state loss contribution.

Rank	**Issuer**	**Sector**	w_i	p_i	$p_i(y_{0.999})$	**VaR share**
1	Verizon Communications Inc.	Technology	2.29%	0.154%	4.64%	2.12%
2	Fiserv Inc.	Industrials	0.47%	1.18%	15.1%	1.42%
3	Boeing Co.	Industrials	0.99%	0.297%	7.16%	1.42%
4	AT&T Inc.	Technology	2.06%	0.100%	3.42%	1.40%
5	Centene Corp.	Healthcare	0.76%	0.368%	8.16%	1.23%
6	Oracle Corp.	Technology	0.91%	0.245%	6.33%	1.14%
7	AbbVie Inc.	Healthcare	1.56%	0.102%	3.48%	1.08%
8	Broadcom Inc.	Technology	2.03%	0.0648%	2.49%	1.00%
9	American Tower Corp.	Real Estate	0.74%	0.264%	6.63%	0.98%
10	CoreWeave Inc.	Technology	0.26%	1.88%	18.5%	0.98%

5.7 What drives baseline tail loss?

Baseline tail loss is driven by a relatively small set of large issuers and by sectoral exposure concentration rather than by unconditional PD levels alone. To make the aggregate VaR result more interpretable, this subsection examines the composition of the baseline ASRF tail state by issuer and by sector.

Table 8 reports the ten largest issuer-level contributors to baseline tail loss. The ranking is not determined by unconditional PD alone. Instead, tail-state contribution reflects the joint effect of portfolio weight, mapped LGD, and the issuer's conditional default probability in the 99.9% systematic stress state. The top contributor is Verizon Communications, followed by Fiserv, Boeing, and AT&T. The top-10 issuers together account for approximately 12.8% of the baseline tail-state loss.

An important implication of Table 8 is that tail-loss contribution does not simply mirror point-in-time default risk. Some issuers rank highly because they carry very large portfolio weights despite low unconditional PDs, whereas others enter the top group because more moderate exposure is combined with sharply elevated conditional default probability in the 99.9% stress state. Tail-risk contribution should therefore be interpreted as the joint outcome of exposure size and stress-state default intensity, rather than as a simple ranking by standalone credit quality.

The sector decomposition reinforces the same conclusion at a more aggregated level. Table 9 shows that Technology is the largest sector driver of baseline tail loss, accounting for 24.2% of $\mathrm{VaR}_{0.999}$, followed by Consumer Cyclicals at 17.0%, Healthcare at 14.4%, and Industrials at 14.0%. Energy contributes a further 10.6%, while Consumer Non-Cyclicals and Real Estate each account for roughly 7–8% of the baseline tail. Overall, the baseline

Table 9: Sector-level decomposition of baseline tail-state loss.

Sector	**Issuers**	**VaR share**	**UL share**
Technology	171	24.2%	24.3%
Consumer Cyclicals	142	17.0%	16.8%
Healthcare	88	14.4%	14.6%
Industrials	134	14.0%	14.1%
Energy	70	10.6%	10.6%
Consumer Non-Cyclicals	59	7.36%	7.44%
Real Estate	46	7.26%	7.15%
Basic Materials	52	5.10%	5.01%
Academic & Educational Services	3	0.086%	0.087%

tail is not evenly distributed across the portfolio, but instead concentrated in sectors that combine large exposure shares with non-trivial stress sensitivity.

5.8 Robustness and estimation uncertainty

The main conclusions are not materially driven by minor data-construction choices, although model-based tail estimates still retain non-trivial sampling uncertainty. This subsection therefore distinguishes between light implementation-level robustness checks and broader uncertainty around the reported point estimates.

Table 10 shows that the light robustness checks leave the baseline tail-loss picture essentially unchanged. First, aggregating issuer-level PDs by the maximum rather than the mean has a negligible effect on portfolio tail loss. Under max-PD aggregation, $\mathrm{VaR}_{0.999}$ equals 3.02%, whereas under mean-PD aggregation it is 3.01%, a difference of only -0.01 percentage points. Second, restricting the portfolio to the set of issuers that cumulatively accounts for the top 95% of total exposure reduces the issuer count from 765 to 526 and lowers $\mathrm{VaR}_{0.999}$ only modestly, from 3.02% to 2.94%, that is, by -0.08 percentage points. These checks suggest that the baseline tail-loss estimates are not mechanically driven by the choice of issuer-level PD aggregation rule or by the extreme lower tail of very small exposures.

Taken together, these light robustness checks support the interpretation that the baseline findings are stable to modest implementation-level changes. Small adjustments in PD aggregation or in the inclusion of very small exposures do not materially alter the estimated tail-risk profile. At the same time, this stability should not be interpreted as eliminating model uncertainty more broadly; rather, it indicates that the main conclusions do not hinge on a narrow set of mechanical data-construction choices.

Table 10: Light robustness checks for the baseline tail-loss estimate. All loss measures are reported as percentages of total portfolio exposure.

Scenario group	**Scenario**	N	$\mathrm{PD}_{\mathrm{wtd}}$	$\mathrm{VaR}_{\mathbf{0.999}}$	$\Delta\mathrm{VaR}_{\mathbf{0.999}}$
PD aggregation	Max PD per issuer (baseline)	765	0.294722%	3.02%	0.00 pp
PD aggregation	Mean PD per issuer	765	0.294%	3.01%	−0.01 pp
Exposure filter	All issuers (baseline)	765	0.294722%	3.02%	0.00 pp
Exposure filter	Top 95% EAD issuers	526	0.274%	2.94%	−0.08 pp

Table 11: Bootstrap uncertainty for selected $\mathrm{VaR}_{0.999}$ specifications. Intervals report the 5% and 95% bootstrap quantiles.

Model	**Point estimate**	**P5**	**P95**
ASRF analytical	3.02%	2.86%	3.21%
One-factor MC	3.26%	2.97%	3.79%
Sector MF MC ($global\_share = 0.50$)	2.24%	2.02%	2.50%
State-LGD Sector MF ($global\_share = 0.50$, medium)	3.67%	3.30%	4.14%

Bootstrap results provide a complementary view of uncertainty around the headline model outputs. Table 11 reports the point estimate together with the 5%–95% bootstrap interval for the main $\mathrm{VaR}_{0.999}$ specifications discussed in the results section. The analytical ASRF baseline $\mathrm{VaR}_{0.999}$ of 3.02% lies within a bootstrap interval of 2.86% to 3.21%. The one-factor Monte Carlo estimate of 3.26% lies within a corresponding interval of 2.97% to 3.79%. The Gaussian sector multi-factor benchmark with $global\_share = 0.50$ yields a point estimate of 2.24% within a bootstrap interval of 2.02% to 2.50%. The state-dependent LGD sector multi-factor specification remains clearly higher, with a point estimate of 3.67% and a bootstrap interval of 3.30% to 4.14%. In all four cases, the reported point estimate lies comfortably within the corresponding bootstrap interval. Thus, while the bootstrap confirms that tail-loss estimates are subject to non-trivial uncertainty, it does not overturn the ranking of the main model specifications.

Taken together, the robustness checks and bootstrap results support two main conclusions. First, small implementation-level choices at the data-construction margin do not materially alter the baseline tail-risk picture, whereas broader modelling choices continue to matter strongly for the level of estimated extreme loss. Second, the bootstrap evidence indicates that the headline $\mathrm{VaR}_{0.999}$ estimates should be interpreted as ranges rather than exact point values. This is fully consistent with the model-risk perspective adopted throughout the thesis: the empirical results are most informative when read as a disciplined baseline together with an explicit uncertainty envelope around it.

6 Discussion of results

6.1 Main findings in relation to prior literature

The empirical findings are broadly consistent with the portfolio credit risk literature in which tail outcomes are driven primarily by systematic dependence rather than by average default rates alone. In the baseline specification, expected loss remains small relative to total portfolio exposure, whereas the far-tail loss measure is materially larger. This pattern is fully in line with the Vasicek–ASRF logic, where the systematic factor governs the clustering of defaults in adverse states and thereby steepens the loss distribution in the tail (Vasicek, 2002; Gordy, 2003). In this sense, the baseline results support the central theoretical intuition developed by Vasicek and by the ASRF literature: portfolio credit risk is fundamentally a tail problem because diversification that is effective in average states becomes much less effective when obligors are simultaneously exposed to the same adverse systematic shock (Vasicek, 2002; Gordy, 2003).

The finite-portfolio results are likewise consistent with the granularity literature. In the baseline portfolio, the gap between analytical ASRF VaR and one-factor Monte Carlo VaR is present but still moderate, suggesting that the asymptotic benchmark remains informative in a portfolio that is not extremely concentrated. However, once the exposure structure is made more concentrated, the gap widens sharply. This is precisely the pattern emphasised in the work on granularity and concentration risk: asymptotic approximations perform well when portfolios are sufficiently fine-grained, but they can materially understate extreme loss once name concentration becomes economically meaningful (Gordy, 2003; Gordy and Lütkebohmert, 2013; Wilde, 2001). The present evidence therefore reinforces the view that concentration is not a minor technical detail but a first-order model-risk channel in applied portfolio credit measurement (Gordy and Lütkebohmert, 2013).

The extended simulations also clarify where the Gaussian one-factor benchmark is informative and where it becomes too restrictive. The Gaussian sector multi-factor results suggest that a pure one-factor structure may overstate portfolio-wide tail loss when systematic risk is not fully market-wide but partly segmented across sectors (Pykhtin, 2004). At the same time, the t-copula results show that the Gaussian benchmark can also be too optimistic once heavier-tailed joint extremes are allowed (Li, 2000; McNeil et al., 2015). Similarly, the state-dependent LGD simulations show that tail loss is shaped not only by clustered default incidence but also by the possibility that loss severity itself deteriorates in the same adverse states, which is consistent with the literature on downturn recoveries and cyclical LGD (Schuermann, 2004). Taken together, these findings are consistent with the broader credit risk literature that treats dependence misspecification, tail dependence, and downturn recoveries as central model-risk concerns rather than as secondary refinements (Li, 2000; McNeil et al., 2015; Schuermann, 2004).

A useful additional perspective comes from the PCA-based sector proxy diagnostic. The historical weekly sector panel is dominated by a broad common component, which is consistent with the continued relevance of the one-factor benchmark as a first-order approximation. At the same time, the second and third principal components retain non-trivial explanatory power, indicating that sector-level co-movement is not fully one-dimensional. This finding supports the broader literature in which multi-factor dependence structures are treated as empirically plausible extensions of the single-factor benchmark rather than as purely ad hoc modifications. In the context of this thesis, the PCA evidence therefore helps explain why the sector multi-factor simulations are economically meaningful even though the one-factor model remains a useful baseline.

Against this background, the contribution of the thesis is not to reject the Vasicek framework, but to locate more precisely where its benchmark value remains informative in an applied issuer-level setting. In particular, the results show that portfolio construction, concentration, dependence specification, and severity assumptions all matter materially for the interpretation of tail-loss estimates. The thesis therefore contributes to the applied literature by combining transparent issuer-level portfolio construction with evidence on representativeness, finite-default effects, and model-risk channels within one coherent empirical framework.

6.2 Interpreting the benchmark in the present dataset

The empirical interpretation of the tail-loss estimates must be conditioned on the way the portfolio is constructed. The final modelling sample is not a full market portfolio, but a matched issuer subset obtained from the intersection of the PD dataset and the debt-only bond universe. As the results section showed, this matched sample is tilted toward larger, lower-PD, and more unsecured issuers than the unmatched observations. The reported loss estimates should therefore not be read as direct statements about the entire corporate debt market, but rather as measurements for a selective, economically meaningful issuer portfolio built from the available overlap between the two source universes.

A second interpretive boundary concerns the exposure measure. Portfolio size and issuer weights are based on issuance amounts rather than on institution-specific exposure-at-default, current holdings, or contractual utilisation profiles. This makes the exposure proxy useful for internal scaling, concentration diagnostics, and scenario comparison within the constructed sample, but it limits the interpretation of absolute euro magnitudes. The euro-denominated loss figures are therefore best understood as indicative portfolio-scale translations of the loss rates rather than as institution-specific loss forecasts. In contrast, the loss-rate results themselves remain informative because they describe how the same internally consistent portfolio reacts to alternative assumptions on dependence, severity, and granularity.

Within these limits, the benchmark remains highly useful. Its value in the present thesis is not that it reproduces a bank's exact regulatory or economic capital model, but that it provides a transparent mapping from issuer-level PD, exposure, and LGD inputs into tail-loss outcomes. This makes it possible to identify which conclusions are robust within the matched sample, which depend strongly on modelling choices, and where the Gaussian one-factor approximation becomes too restrictive. In that sense, the benchmark is best interpreted as a disciplined empirical reference point rather than as a fully general measure of credit tail risk.

6.3 What the model extensions imply

The model extensions show that the estimated tail loss depends materially on whether one abstracts from finite-default effects, richer dependence structures, and deterioration in LGD during stressed states. The baseline model remains a useful reference point, but the extensions reveal how strongly tail estimates can move once these additional channels are taken seriously.

A first qualification comes from the finite-portfolio evidence. In the baseline portfolio, the gap between analytical ASRF VaR and one-factor Monte Carlo VaR is present but still moderate, which suggests that the asymptotic approximation remains informative under the observed degree of concentration. However, the concentration experiments show that this conclusion is not structurally stable. Once the portfolio becomes dominated by a smaller number of large names, the Monte Carlo tail rises far above the asymptotic benchmark. This implies that the one-factor benchmark is most credible when interpreted together with a granularity condition: it is informative for moderately diversified issuer portfolios, but can become materially optimistic once concentration risk is economically meaningful (Gordy, 2003; Gordy and Lütkebohmert, 2013; Wilde, 2001).

A second qualification concerns the modelling of dependence itself. The Gaussian sector multi-factor simulations suggest that the pure one-factor benchmark may overstate tail loss when systematic risk is partly segmented across sectors rather than fully market-wide (Pykhtin, 2004). At the same time, the t-copula results point in the opposite direction and indicate that the Gaussian benchmark may understate tail loss when joint extremes are heavier-tailed than the normal benchmark permits (Li, 2000; McNeil et al., 2015).

These two findings are jointly important. They imply that the main model-risk issue is not simply whether the benchmark is conservative or non-conservative in one direction, but that the tail estimate depends strongly on how dependence is structured. The baseline one-factor model is therefore best viewed as a transparent simplification of dependence rather than as a neutral representation of it.

The PCA evidence sharpens this interpretation. The sector proxy panel suggests that dependence in market-based sector data is strongly but not exclusively driven by one common component. This means that the Gaussian one-factor benchmark is not contradicted by the data, but neither is it fully sufficient to describe all systematic co-movement. In that sense, the PCA diagnostic supports a middle-ground interpretation: the one-factor model remains a useful organising benchmark, while low-dimensional multi-factor extensions are empirically defensible when one wants to relax the assumption that all systematic risk is perfectly market-wide.

A third qualification arises from the state-dependent LGD results. Once loss severity is allowed to worsen in adverse common states, tail losses increase materially relative to the constant-LGD benchmark. This means that the benchmark can miss an important amplification mechanism even if the default side of the model is otherwise left unchanged. The key implication is that tail credit loss is jointly determined by dependence, finite granularity, and severity cyclicality (Schuermann, 2004). Taken together, the extensions show that the baseline Gaussian one-factor estimate remains useful as an interpretable anchor, but that any serious tail-risk assessment should recognise a wider range of plausible outcomes around that anchor.

6.4 Remaining limitations and unresolved issues

Even after the empirical extensions, the analysis remains a stylised fixed-horizon credit risk measurement exercise rather than a fully validated forecasting model or an institution-specific capital framework. The thesis therefore identifies how tail-loss estimates behave under a transparent set of modelling assumptions, but not how any one specification would perform as a complete realised-loss forecasting model in practice.

A first limitation concerns the empirical inputs. The portfolio is built from a selective matched issuer sample, and the exposure measure is based on issuance amounts rather than on actual holdings, contractual exposure-at-default, or utilisation profiles. Likewise, LGD is imposed through seniority-based mapping and stress scenarios rather than estimated from realised recovery data. Most importantly, the analysis does not include a historical backtesting exercise against realised defaults or realised portfolio loss rates. The results should therefore be interpreted as disciplined risk estimates for a constructed issuer portfolio, not as validated out-of-sample forecasts of realised credit losses.

A second limitation concerns model structure. Although the analysis moves beyond the Gaussian one-factor benchmark through finite-default simulation, sector multi-factor specifications, t-copula dependence, and state-dependent LGD, the framework remains static and one-period. It does not model rating migration, spread dynamics, refinancing effects, liquidity stress, or feedback between macro conditions and default timing. Nor are the alternative factor structures estimated from a full structural or market-implied dependence model. Accordingly, the extensions widen the range of plausible tail outcomes, but they do not eliminate model uncertainty.

These limitations define the scope of the thesis rather than invalidate its contribution. The results remain informative for understanding how issuer-level PD, exposure, LGD, dependence assumptions, and concentration interact in the formation of portfolio tail loss. However, applications aimed at institution-specific capital planning, validation, or realised loss forecasting would require richer exposure data, observed recovery information, and a design that permits explicit backtesting over time.

7 Conclusions

This thesis examined portfolio tail credit losses in a matched debt-only corporate issuer sample using a transparent issuer-level implementation of the Vasicek–ASRF framework. The main objective was not to propose a new portfolio credit risk model, but to show how one-year physical default probabilities, an issuance-based exposure proxy, mapped loss-given-default assumptions, and dependence structure translate into portfolio tail-loss estimates, and how sensitive those estimates are to concentration, finite-default effects, and model extensions beyond the Gaussian one-factor benchmark.

The first main conclusion is that portfolio credit risk in the present sample is clearly a tail phenomenon rather than an expected-loss problem. In the baseline specification, EL = 0.165% of total exposure, whereas $\mathrm{VaR}_{0.999} = 3.02\%$, implying $\mathrm{UL}_{0.999} = 2.85\%$. Thus, the far-tail loss rate is many times larger than the average loss rate. This is fully consistent with the underlying Vasicek–ASRF logic: when a severe adverse systematic state occurs, conditional default probabilities rise jointly across issuers and portfolio losses increase in a strongly non-linear way. The central empirical implication is therefore that even portfolios with modest average PDs can still exhibit substantial extreme-loss exposure once default clustering is taken seriously.

The second main conclusion is that the economic interpretation of these tail-loss estimates must be conditioned on the way the portfolio is constructed. The final modelling sample is a selective matched issuer subset obtained from the overlap of the PD dataset and the debt-only bond universe. It contains 765 issuers and retains 58.0% of the issuance-based exposure in the debt-only bond universe. Moreover, the matched sample is tilted toward larger, lower-PD, and more unsecured issuers than the unmatched observations. The results should therefore not be interpreted as direct statements about the entire corporate debt market. Rather, they describe tail-loss behaviour in a constructed but economically meaningful issuer portfolio. The strongest empirical content of the thesis accordingly lies in internally consistent comparisons across model specifications within this matched sample.

A third conclusion concerns the main drivers of tail loss within the baseline one-factor structure. The results show that tail estimates are more sensitive to dependence assumptions than to simple unconditional LGD changes. Holding the mapped LGD structure fixed, changing correlation from a lower constant level to a higher one produces large differences in $\mathrm{VaR}_{0.999}$, whereas constant-LGD changes primarily rescale loss severity. At the same time, downturn-style LGD stresses increase tail loss materially, which indicates that severity assumptions remain economically important even if they are not the primary source of default clustering. The overall conclusion is that correlation governs the thickness of the tail, while LGD governs how costly clustered defaults become once they occur.

A fourth conclusion is that the asymptotic ASRF approximation remains informative in the baseline portfolio, but it can materially understate tail risk once concentration becomes pronounced. In the baseline sample, the gap between analytical $\mathrm{VaR}_{0.999}^{\mathrm{ASRF}}$ and one-factor Monte Carlo $\mathrm{VaR}_{0.999}^{\mathrm{MC}}$ is present but moderate. However, when the portfolio is made more concentrated, the gap widens sharply. Conversely, in a much more granular equal-weighted portfolio, the gap becomes negligible. This implies that the usefulness of the ASRF benchmark depends crucially on portfolio granularity. In moderately diversified issuer portfolios it remains a useful approximation and an interpretable benchmark, but in concentrated portfolios finite-default effects become first-order and explicit Monte Carlo treatment or concentration adjustments are needed.

The fifth conclusion is that moving beyond the Gaussian one-factor benchmark substantially widens the range of plausible tail outcomes. A PCA-based diagnostic applied to an external weekly sector proxy panel further suggests that dependence is dominated by a broad common component, but not fully summarised by a single factor. This supports a middle-ground interpretation of the benchmark. The Gaussian one-factor model remains a useful first-order approximation, but low-dimensional multi-factor extensions are also empirically defensible when systematic co-movement is not fully market-wide. Consistent with this interpretation, the Gaussian sector multi-factor simulations show that tail loss may fall relative to the pure one-factor benchmark when systematic dependence is partly segmented across sectors. By contrast, the t-copula specifications imply much larger tail losses, indicating that the Gaussian benchmark can be materially optimistic once heavier-tailed joint extremes are allowed. Likewise, the state-dependent LGD simulations show that tail loss increases substantially when recovery severity deteriorates endogenously in stressed states. Taken together, these extensions show that the Gaussian one-factor model is best interpreted as a transparent reference point rather than as a uniquely robust estimate of extreme-loss risk.

A sixth conclusion is that baseline tail loss is not evenly spread across issuers or sectors. Tail-risk contributions are concentrated in a limited set of economically large names and in sectors that combine high exposure shares with non-trivial stress sensitivity. In the baseline decomposition, the top-10 issuers account for a meaningful share of total tail-state loss, and Technology is the single largest sector contributor. This finding is important from an applied perspective because it shows that portfolio tail risk is shaped by the joint distribution of exposure size, conditional tail-state default intensity, and loss severity, rather than by unconditional PDs alone. Tail-risk measures are therefore most informative when complemented by contribution analysis.

The robustness checks support the stability of the main empirical messages. Small implementation-level changes, such as using mean rather than maximum issuer-level PDs or trimming the smallest exposures, do not materially alter the baseline tail-risk picture. Bootstrap results also indicate that the ranking of the main model specifications is stable even though point estimates are subject to non-trivial estimation uncertainty. The main conclusions of the thesis are therefore not driven by a single fragile aggregation rule or by small-exposure noise. Instead, they appear structurally linked to dependence assumptions, concentration, and severity modelling.

These conclusions should nevertheless be interpreted within clear limits. The analysis is based on one-year physical PDs, an issuance-based exposure proxy, and a stylised seniority-based LGD mapping. The empirical design is cross-sectional and point-in-time, not a historical panel of realised portfolio losses. The thesis therefore measures and interprets tail-loss risk under transparent assumptions, but it does not provide a full backtesting exercise of realised one-year default or loss outcomes. Nor does it constitute an institution-specific capital model. Applications aimed at bank portfolio management, validation, or capital planning would require richer exposure data, observed recovery information, and a design that permits explicit out-of-sample performance assessment.

Overall, the thesis supports a clear final conclusion. The Vasicek–ASRF framework remains a useful and interpretable benchmark for portfolio tail credit risk measurement, especially when the objective is to understand how issuer-level PDs, exposures, LGDs, and dependence assumptions combine to produce extreme losses. At the same time, the results show that tail-loss estimates are highly sensitive to concentration, dependence structure, and severity cyclicality. The practical implication is that Gaussian one-factor tail estimates should be interpreted as disciplined benchmarks, not as mechanically sufficient answers. Reliable tail-risk assessment therefore requires explicit attention to sample construction, concentration, model risk, and the possibility that the dependence structure and loss severity become materially more adverse in the same stress states. In applied portfolio risk management, this means that transparent benchmark models remain useful, but they should be complemented by concentration diagnostics, dependence stress tests, and conservative severity scenarios rather than interpreted in isolation. A natural next step would be to combine the present cross-sectional framework with a historical panel design that permits explicit backtesting of realised default clustering and loss outcomes over time.

References

Altman, E. I., Brady, B., Resti, A. – Sironi, A. (2005) The link between default and recovery rates: Theory, empirical evidence, and implications. *The Journal of Business*, vol. 78 (6), 2203–2228.

Altman, E. I., Hotchkiss, E. – Wang, W. (2019) *Corporate Financial Distress, Restructuring, and Bankruptcy*. John Wiley & Sons, Hoboken, NJ, 4 edn.

Altman, E. I. – Kishore, V. M. (1996) Almost everything you wanted to know about recoveries on defaulted bonds. *Financial Analysts Journal*, vol. 52 (6), 57–64.

Baird, D. G. (2017) Absolute priority, relative priority, and the costs of bankruptcy. *University of Pennsylvania Law Review*, vol. 166, 741–785.

Basel Committee on Banking Supervision (2005a) An explanatory note on the Basel II IRB risk weight functions. Tech. rep., Bank for International Settlements, URL: `https://www.bis.org/bcbs/irbriskweight.pdf`, explanatory note on the ASRF/IRB risk weight functions (Basel II).

Basel Committee on Banking Supervision (2005b) Guidance on paragraph 468 of the framework document. Tech. Rep. BCBS 115, Bank for International Settlements, URL: `https://www.bis.org/publ/bcbs115.pdf`.

Basel Committee on Banking Supervision (2006) International convergence of capital measurement and capital standards: A revised framework (comprehensive version). Tech. rep., Bank for International Settlements, URL: `https://www.bis.org/publ/bcbs128.pdf`.

Basel Committee on Banking Supervision (2013) Principles for effective risk data aggregation and risk reporting. Tech. Rep. BCBS 239, Bank for International Settlements, URL: `https://www.bis.org/publ/bcbs239.pdf`.

Bharath, S. T. – Shumway, T. (2008) Forecasting default with the Merton distance to default model. *Review of Financial Studies*, vol. 21 (3), 1339–1369.

Black, F. – Cox, J. C. (1976) Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance*, vol. 31 (2), 351–367.

Board of Governors of the Federal Reserve System (2011) SR 11-7: Guidance on model risk management. Tech. rep., Board of Governors of the Federal Reserve System, URL:

https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm, supervision and Regulation Letter SR 11-7; issued jointly with the OCC.

Brealey, R. A., Myers, S. C. – Allen, F. (2019) *Principles of Corporate Finance*. McGraw-Hill Education, New York, NY, 13 edn.

Crosbie, P. J. – Bohn, J. R. (2003) Modeling default risk. Moody's KMV technical document, Moody's KMV methodology note.

Duffie, D. – Singleton, K. J. (1999) Modeling term structures of defaultable bonds. *Review of Financial Studies*, vol. 12 (4), 687–720.

Duffie, D. – Singleton, K. J. (2003) *Credit Risk: Pricing, Measurement, and Management*. Princeton University Press, Princeton, NJ.

Emmer, S. – Tasche, D. (2005) Calculating credit risk capital charges with the one-factor model. *The Journal of Risk*, vol. 7 (2), 85–101, published version of arXiv:cond-mat/0302402.

European Banking Authority (2016) Guidelines on the application of the definition of default under Article 178 of Regulation (EU) No 575/2013. Tech. Rep. EBA/GL/2016/07, European Banking Authority.

European Banking Authority (2017) Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures. Tech. Rep. EBA/GL/2017/16, European Banking Authority.

European Parliament – Council of the European Union (2019) Directive (EU) 2019/1023 of the European Parliament and of the council of 20 June 2019 on preventive restructuring frameworks, on discharge of debt and disqualifications, and on measures to increase the efficiency of procedures concerning restructuring, insolvency and discharge of debt. EUR-Lex, URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019L1023, CELEX: 32019L1023.

Frey, R. – McNeil, A. J. (2003) Dependent defaults in models of portfolio credit risk. *Journal of Risk*, vol. 6 (1), 59–92.

Frye, J. (2000) Depressing recoveries. Tech. Rep. S&R-2000-8, Federal Reserve Bank of Chicago.

García Céspedes, J. C. – Moreno, M. (2017) An approximate multi-period Vasicek credit risk model. *Journal of Banking & Finance*, vol. 81, 105–113.

Geske, R. (1977) The valuation of corporate liabilities as compound options. *Journal of Financial and Quantitative Analysis*, vol. 12 (4), 541–552.

Gordy, M. B. (1998) A comparative anatomy of credit risk models. Tech. Rep. FEDS 1998-47, Board of Governors of the Federal Reserve System, URL: `https://www.federalreserve.gov/pubs/feds/1998/199847/199847pap.pdf`.

Gordy, M. B. (2003) A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation*, vol. 12 (3), 199–232.

Gordy, M. B. – Heitfield, E. (2002) Estimating default correlations from short panels of credit rating performance data. Tech. rep., Board of Governors of the Federal Reserve System, working paper dated 29 January 2002.

Gordy, M. B. – Lütkebohmert, E. (2007) Granularity adjustment for Basel II. Tech. Rep. 01/2007, Deutsche Bundesbank.

Gordy, M. B. – Lütkebohmert, E. (2013) Granularity adjustment for regulatory capital assessment. *International Journal of Central Banking*, vol. 9 (3), 33–71.

Hull, J., Predescu, M. – White, A. (2005) Bond prices, default probabilities and risk premiums. *Journal of Credit Risk*, vol. 1 (2), 53–60.

Hull, J. – White, A. (2008) Dynamic models of portfolio credit risk: A simplified approach. *The Journal of Derivatives*, vol. 15 (4), 9–28.

Hull, J. C. (2021) *Options, Futures, and Other Derivatives*. Pearson, Hoboken, NJ, 11 edn.

Lando, D. (1998) On Cox processes and credit risky securities. *Review of Derivatives Research*, vol. 2 (2–3), 99–120.

Li, D. X. (2000) On default correlation: A copula function approach. Tech. Rep. 99-07, RiskMetrics Group.

Longstaff, F. A., Mithal, S. – Neis, E. (2005) Corporate yield spreads: Default risk or liquidity? New evidence from the credit-default swap market. *The Journal of Finance*, vol. 60 (5), 2213–2253.

McNeil, A. J., Frey, R. – Embrechts, P. (2015) *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, Princeton, NJ, 2 edn.

Merton, R. C. (1974) On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, vol. 29 (2), 449–470.

Pykhtin, M. (2004) Multi-factor adjustment. *Risk*, 85–90.

Schuermann, T. (2004) What do we know about Loss Given Default? In *Credit Risk Models and Management*, ed. D. Shimko, 249–274, Risk Books, London, chapter in an edited volume.

Tasche, D. (1999) Risk contributions and performance measurement. Tech. rep., Technische Universität München, working paper.

Vasicek, O. (2002) Loan portfolio value. *Risk*, vol. 15 (12), 160–162.

Wilde, T. (2001) Probing granularity. *Risk*, vol. 14 (8), 103–106.

Yamai, Y. – Yoshiba, T. (2002) Comparative analyses of Expected Shortfall and Value-at-Risk: Their estimation error, decomposition, and optimization. *Monetary and Economic Studies*, vol. 20 (1), 87–121, institute for Monetary and Economic Studies, Bank of Japan.

Disclosing the use of AI:

I used AI tools such as ChatGPT to improve my language and choices of words. Also, the abstract was translated to Finnish using AI. Some of the R scripts were created with the help of AI.

APPENDIX 1: Data management plan

Table 12: Research data used in the thesis.

Type of material	Contains personal data	Produced by me	Produced by others	Other notes
PD data export (LSEG / Refinitiv)	No	No	Yes	Issuer-level one-year physical PD data used as baseline model input.
Bond data export (LSEG / Refinitiv)	No	No	Yes	Instrument-level bond and note data used for issuer exposure aggregation and seniority-based LGD mapping.
Sector ETF proxy panel	No	No	Yes	Weekly adjusted return series for 11 U.S. sector ETFs used only for the PCA-based external dependence diagnostic in the model-extension analysis.
Derived issuer-level dataset	No	Yes	No	Constructed by cleaning, matching, and aggregating the source datasets into issuer-level PD, EAD, LGD, and portfolio-weight inputs.
Derived sector proxy return panel and PCA outputs	No	Yes	No	Constructed from the external sector ETF proxy data; includes standardised weekly return panel, PCA explained-variance summaries, and associated diagnostic output files.
R code and output files	No	Yes	No	Used for cleaning, aggregation, baseline Vasicek implementation, Monte Carlo simulations, PCA diagnostics, robustness checks, and result tables.

Handling of personal data

This thesis does not use personal data. The empirical material consists of firm-level financial and credit-risk data obtained from external commercial databases, market-based sector ETF proxy data, and derived datasets constructed for the empirical analysis.

Storage of material during research

The material used in the thesis was stored in the author's secured personal storage during the research process. Raw data exports, external proxy data, derived datasets, and R scripts were stored in separate folders to maintain version control and reproducibility. Intermediate analysis tables and model outputs were retained alongside the scripts used to generate them.

Documentation and metadata of the material

The empirical material consists of two main vendor source files: issuer-level one-year probability of default data and instrument-level bond and note data. These source files were cleaned, matched, and aggregated into an issuer-level modelling dataset. In addition, the model-extension analysis uses an auxiliary external sector proxy panel consisting of weekly adjusted return series for 11 U.S. sector ETFs. This auxiliary panel is used only for the PCA-based dependence diagnostic and does not alter the baseline issuer-level portfolio construction. Variable definitions, transformations, and implementation choices are documented in the thesis and in the accompanying R code used for the empirical analysis.

Material after completion of the study

The thesis text is publicly available through the university repository. The original vendor source data are subject to third-party licensing restrictions and therefore cannot be made publicly available. The external sector ETF proxy data are market-based source data obtained from third-party services and should be treated in accordance with the relevant source terms. The author's own code and derived non-licensed output files may be retained for a limited period for verification and potential later reference.