

Credit Risk Analytics

with Python

A practical handbook for students, instructors,
model developers, validators, and banking practitioners

From exposure-level expected loss to correlated portfolio simulation, credit VaR, Basel IRB benchmarks, IFRS 9 lifetime ECL, stress testing, monitoring, model governance, ratings migration, CDS valuation, structural models, counterparty CVA, CreditRisk+, CreditMetrics, and spread risk.

Built from the Credit Risk Lab executable package
Real-Data Expanded Edition, Version 2.1 – October 2026

How to use this handbook

This book connects mathematical definitions, executable Python, interpretation, and institutional controls. Each technical chapter follows the same pattern:

1. define the risk quantity and its estimand;
2. derive the principal equations and assumptions;
3. implement the calculation in transparent Python;
4. show results generated by the accompanying package;
5. explain validation, governance, and common failure modes;
6. finish with exercises or questions for review.

The material follows the University of Economics in Katowice course *Measurement of Credit Risk*: loss distributions, expected and unexpected loss, PD, LGD, EAD, implied PD, dependence, credit VaR, and stress testing. It extends that foundation with IFRS 9 mechanics, Basel IRB benchmarks, model validation, data lineage, ratings migration, traded credit, structural models, counterparty credit risk, portfolio migration models, spread risk, and evidence packs used in contemporary institutions.

All reported numerical examples were executed with the companion `credit-risk-lab` package. Version 2.1 bundles dated observed-data snapshots: 1,000 UCI credit applicants, 792 joint FRED corporate-spread dates, 9,716 FRED credit-condition dates, 360 World Bank country-indicator-years, and 7,169 ECB daily exchange-rate observations. Every CSV has a metadata sidecar with provider, source URL, retrieval date, row count, and `synthetic=false`.

The separate demonstration portfolio contains 30 fictional exposures with total EAD of EUR 45.144 million. It remains only because public customer-level bank portfolios, derivative netting sets, CSAs, and matched CDS histories are not openly redistributable. The package labels every such input and never presents it as an observation.

Observed-data coverage

Figure 1 is drawn directly from the bundled FRED, World Bank, and ECB snapshots. The UCI records feed the PD and CreditRisk+ laboratories. The download script can refresh every snapshot while preserving its provenance sidecar.




Figure 1: Observed public data bundled with the executable package

Control warning. This is an educational and prototyping handbook. It is not a turnkey lending, provisioning, pricing, stress testing, or regulatory capital system. Definitions, data, legal requirements, accounting policy, validation, approval, access control, and change governance remain institution-specific.

Suggested routes

Student. Read Chapters 1–12 in order, then use Chapters 17–24 for markets, counterparty, and portfolio-model specialization.

Instructor. Use the worked figures in class, then alter one assumption at a time and ask learners to predict direction and magnitude.

Practitioner. Start with Chapters 2, 13, 14, and 24; then use the component chapters as transparent benchmarks and challenger specifications.

Validator. Focus on estimands, leakage, calibration, uncertainty, sensitivity, implementation verification, and limitations in every chapter.

Contents

How to use this handbook	i
I Foundations and institutional workflows	1
1 Credit risk as a measurable system	2
1.1 From event to economic consequence	2
1.2 Different questions require different loss measures	3
1.3 The sample portfolio	3
2 Data, time, and definitions	5
2.1 Write the data contract before the model	5
2.2 Time axes that must not be confused	5
2.3 Snapshot and lineage	5
3 Loss distributions	7
3.1 Why credit losses are not normally distributed	7
3.2 Shape diagnostics	8
4 Expected and unexpected loss	9
4.1 Expected loss	9
4.2 Unexpected loss	9
5 Probability of default	10
5.1 The estimand	10
5.2 Logistic scorecard	10
5.3 Discrimination and calibration	11
5.4 Calibration to a target central tendency	11
5.5 Validation design	12
6 PD implied by prices and bank rates	13
6.1 Reduced-form approximation	13
6.2 Why it is not a physical PD without adjustment	13
7 Loss given default	15
7.1 Workout LGD	15
7.2 Long-run and downturn LGD	15
7.3 Modelling choices	15
8 Exposure at default and credit conversion factors	17
8.1 Operational complications	18

9	Dependence and portfolio simulation	19
9.1	Gaussian one-factor construction	19
9.2	Correlated severity	19
9.3	Simulation controls	20
10	Credit VaR, expected shortfall, and capital	21
10.1	Quantiles and tail means	21
10.2	Basel IRB analytical benchmark	21
11	Stress testing and reverse stress	23
11.1	Scenario transformation	23
11.2	Reverse stress	24
11.3	Governance questions	24
12	IFRS 9 expected credit losses	25
12.1	Probability-weighted cash shortfalls	25
12.2	Stages and significant increase in credit risk	26
13	Validation, monitoring, and drift	27
13.1	Population stability index	27
13.2	Outcome tests	27
13.3	Threshold governance	27
14	Production workflow and evidence	29
14.1	A controlled analytical chain	29
14.2	Reproducibility manifest	29
14.3	Model inventory and change control	30
15	Integrated case study	31
15.1	Portfolio results	31
15.2	Committee interpretation	31
15.3	Reproduction	31
16	Teaching plan and assignments	33
16.1	Thirty-hour sequence	33
16.2	Capstone specification	33
16.3	Advanced extensions	34
II	Advanced credit markets and portfolio models	35
17	Ratings, migration, and default intensity	36
17.1	External ratings, internal grades, and what they estimate	36
17.2	Transition matrices	36
17.3	Cumulative, marginal, and conditional default probability	37
18	Credit derivatives, spreads, and basis	39
18.1	Single-name CDS cash flows	39
18.2	Indices and institutional uses	40
18.3	Bond spreads, asset swaps, and CDS-bond basis	40
19	Structural credit models	41
19.1	Equity as an option on firm assets	41
19.2	Implementation and validation	42

20 Counterparty exposure, netting, and collateral	43
20.1 Exposure is a future distribution	43
20.2 Netting sets and collateral agreements	43
20.3 Simulation design and controls	44
21 CVA, DVA, incremental risk, and wrong-way risk	45
21.1 Discrete unilateral CVA	45
21.2 Incremental CVA and allocation	45
21.3 Wrong-way and right-way risk	45
21.4 CVA market risk and governance	46
22 Portfolio credit engines: Vasicek, CreditRisk+, and CreditMetrics	47
22.1 Three questions, three model families	47
22.2 Vasicek conditional default rate	47
22.3 CreditRisk+: Poisson counts and stochastic intensity	47
22.4 CreditMetrics: migration and revaluation	48
23 Credit-spread risk and constant-risk portfolios	50
23.1 Spread duration and convexity	50
23.2 Jump to default and nonlinear instruments	51
23.3 Constant level of risk versus buy and hold	51
24 Advanced integration, governance, and teaching cases	52
24.1 A single measure cannot serve every decision	52
24.2 Expanded executed results	53
24.3 Institutional implementation blueprint	53
24.4 Advanced laboratory sequence	53
24.5 Advanced capstone	54
A Formula reference	55
B Python cookbook	56
B.1 Load and validate	56
B.2 Fetch public macro data	56
B.3 Simulate and calculate tail risk	56
B.4 Stress and retain both states	57
B.5 Monitoring	57
B.6 Ratings, CDS, and structural credit	57
B.7 Counterparty and portfolio models	57
C Glossary	58
D References and official resources	59
Closing note	61

List of Figures

1	Observed public data bundled with the executable package	ii
1.1	Demonstration portfolio profile. Calculated from the supplied offline dataset. . .	3
1.2	Name concentration in the demonstration portfolio.	4
3.1	Correlated one-year loss simulation. The horizontal axis is truncated beyond the 99.98th percentile for readability.	7
5.1	PD holdout diagnostics on the UCI Statlog German Credit dataset. The random split is pedagogical because the source has no observation date.	11
6.1	Sensitivity of five-year implied PD to spread and recovery.	13
8.1	Observed CCF examples. Small undrawn denominators amplify ratio instability.	17
9.1	Executed sensitivity across asset correlations. Tail measures change much more than the mean.	20
10.1	Illustrative corporate IRB benchmark across PD and LGD.	22
11.1	Base and adverse expected loss by segment under the illustrative scenario.	23
12.1	Illustrative forward-looking PD paths and scenario-specific lifetime ECL.	25
17.1	One-year rating transitions and the five-year Markov projection	37
17.2	Observed rating OAS and constant-hazard survival proxies	38
18.1	CDS proxy from observed BBB OAS and leg accumulation	40
19.1	Merton risk-neutral PD sensitivity to debt and equity volatility	42
20.1	Expected exposure, collateralized exposure, and PFE	44
22.1	CreditRisk+ tails anchored to the observed UCI default count	48
22.2	CreditMetrics migration-loss distribution and upper tail	49
23.1	Observed FRED OAS shocks applied to hypothetical current positions	51

List of Tables

1.1	Common credit-risk measures and their principal use	3
12.1	Simplified IFRS 9 staging view	26
13.1	A balanced monitoring dashboard	27
15.1	Executed case-study results	31
16.1	Suggested course delivery plan	33
22.1	Selecting a portfolio model by purpose	49
24.1	Measure-purpose map	52
24.2	Version 2.1 advanced-workflow results and provenance	53
24.3	Additional advanced laboratories	53
A.1	Core formula sheet	55

Part I

Foundations and institutional workflows

Chapter 1

Credit risk as a measurable system

1.1 From event to economic consequence

Credit risk is the possibility that a counterparty's failure to meet a financial obligation causes an economic loss. Measurement is not simply classification. It combines four questions:

1. What event counts as default, and over what horizon?
2. How much exposure exists when the event occurs?
3. What fraction of that exposure is lost after recoveries, costs, and time value?
4. How do individual losses combine under dependence and concentration?

For exposure i , define the one-period random loss

$$L_i = D_i \times \text{LGD}_i \times \text{EAD}_i, \quad D_i \in \{0, 1\}. \quad (1.1)$$

Here D_i is a default indicator, LGD_i is loss severity conditional on default, and EAD_i is exposure at the default date. Portfolio loss is

$$L = \sum_{i=1}^n L_i. \quad (1.2)$$

Equation (1.1) is simple; its components are not. PD may be point-in-time or through-the-cycle. LGD may be long-run, downturn, accounting, or economic. EAD may include future drawings, accrued interest, or netting. Default definitions and reference dates must align across all three.

1.2 Different questions require different loss measures

Table 1.1: Common credit-risk measures and their principal use

Measure	Question	Typical use
Expected loss	What loss is expected on average?	Pricing, budgeting, provisioning inputs
Unexpected loss	How variable is loss around the mean?	Economic capital and risk appetite
VaR	What loss threshold is exceeded with probability $1 - \alpha$?	Tail limit and benchmark
Expected short-fall	What is average loss beyond VaR?	Tail severity and sensitivity
IFRS 9 ECL	What probability-weighted discounted cash shortfall is expected?	Financial reporting
IRB capital	What regulatory benchmark follows the applicable framework?	Prudential capital
Stress loss	What happens under a specified severe but plausible path?	Planning and resilience

Practice point. Never use the word “risk” without naming the horizon, population, event, loss definition, conditioning information, and intended decision.

1.3 The sample portfolio

Figure 1.1 shows the demonstration portfolio used throughout the book. Corporate names dominate EAD, while consumer and SME assets carry higher expected-loss rates. The concentration curve in Figure 1.2 shows why exposure totals alone do not describe diversification.




Figure 1.1: Demonstration portfolio profile. Calculated from the supplied offline dataset.




Figure 1.2: Name concentration in the demonstration portfolio.

Chapter 2

Data, time, and definitions

2.1 Write the data contract before the model

A production-shaped input contract should state required columns, types, units, allowed ranges, uniqueness, reference time, and relationships. At minimum, the one-year portfolio examples require unique exposure identifiers and finite PD, $\text{LGD} \in [0, 1]$, with $\text{EAD} \geq 0$.

Listing 2.1: A strict but readable portfolio contract

```
1 def validate_portfolio(frame):
2     require_columns(frame, ["exposure_id", "pd", "lgd", "ead"])
3     if frame["exposure_id"].duplicated().any():
4         raise DataContractError("exposure_id must be unique")
5     frame["pd"] = require_probability(frame["pd"], name="pd")
6     frame["lgd"] = require_probability(frame["lgd"], name="lgd")
7     frame["ead"] = require_nonnegative(frame["ead"], name="ead")
8     return frame.copy()
```

Range checks catch obvious errors but not semantic errors. A PD in $[0, 1]$ can still be wrong if it is lifetime rather than one-year, includes a cure event, uses a different default definition, or belongs to a later observation date.

2.2 Time axes that must not be confused

Credit datasets commonly contain at least four clocks:

- observation date: information available to the model;
- outcome window: period over which default is measured;
- default date: trigger for recovery and EAD measurement;
- reporting/as-of date: portfolio scope for aggregation.

For development, every feature must be available at the observation date. Borrowers and connected groups should not leak between training and test data. Random row splitting can overstate performance when repeated customers, vintages, or macro regimes appear on both sides.

2.3 Snapshot and lineage

An evidence-grade snapshot should preserve provider, request URL, retrieval time, as-of or vintage, units, schema, row count, raw hash, transformations, exclusions, missing-value policy,

owner, and permitted use. Revisions matter: a macro model validated on final revised values may not represent a model run on information available in real time.

Listing 2.2: Saving a public-data snapshot with metadata

```
1 macro = fred_series(["UNRATE", "DRALACBN", "BAA10Y"], start="2000-01-01")
2 save_snapshot(
3     macro,
4     "data/fred_credit_cycle_2026-10-04.csv",
5     metadata={
6         "provider": "Federal Reserve Bank of St. Louis",
7         "retrieved_at_utc": "2026-10-04T14:00:00Z",
8         "series": ["UNRATE", "DRALACBN", "BAA10Y"],
9         "purpose": "credit-cycle teaching example",
10    },
11 )
```

Chapter 3

Loss distributions

3.1 Why credit losses are not normally distributed

Many obligors do not default in a one-year horizon, so the distribution has a large mass at zero or at low loss. A smaller set of joint defaults and severe recoveries creates a long right tail. In the handbook run, 48.0% of 150,000 portfolio scenarios had zero loss, while the 99.9% VaR was EUR 5.69 million.




Figure 3.1: Correlated one-year loss simulation. The horizontal axis is truncated beyond the 99.98th percentile for readability.

For a Bernoulli default with fixed LGD and EAD,

$$\mathbb{E}[L_i] = p_i \text{LGD}_i \text{EAD}_i, \quad (3.1)$$

$$\mathbb{V}[L_i] = p_i(1 - p_i)(\text{LGD}_i \text{EAD}_i)^2. \quad (3.2)$$

If losses were independent, portfolio variance would be the sum of exposure variances. With dependence,

$$\mathbb{V}[L] = \sum_i \mathbb{V}[L_i] + 2 \sum_{i < j} \text{Cov}(L_i, L_j). \quad (3.3)$$

The covariance terms explain why tail risk can rise sharply without a similar increase in average PD.

3.2 Shape diagnostics

Useful diagnostics include mass at zero, mean, standard deviation, skewness, excess kurtosis, quantiles, tail conditional means, and concentration. A chart should always display the population, horizon, currency, scenario count, seed or run identifier, and whether losses are gross or net of recoveries.

Listing 3.1: Generating and inspecting a loss distribution

```
1 losses = simulate_credit_losses(  
2     portfolio,  
3     n_scenarios=150_000,  
4     asset_correlation=0.15,  
5     lgd_systematic_correlation=0.25,  
6     seed=2026,  
7 )  
8 print(losses.mean(), losses.std(ddof=1))  
9 print(value_at_risk(losses, 0.999))  
10 print(expected_shortfall(losses, 0.999))
```

Review questions

1. Why can two portfolios with the same EL have different tail risk?
2. What changes when exposure size is concentrated in one borrower?
3. Why is fitting a normal distribution to the center dangerous for capital?

Chapter 4

Expected and unexpected loss

4.1 Expected loss

When PD, LGD, and EAD are deterministic parameters over the same horizon,

$$EL_i = PD_i \times LGD_i \times EAD_i, \quad EL_{\text{portfolio}} = \sum_i EL_i. \quad (4.1)$$

For example, a facility with PD 2%, LGD 40%, and EAD EUR 1,000,000 has

$$EL = 0.02 \times 0.40 \times 1,000,000 = \text{EUR } 8,000.$$

The demonstration portfolio analytical EL is EUR 275,678.63, or 0.611% of EAD. The correlated simulation mean is EUR 306,480.16 because the simulation includes positive systematic default-LGD dependence. This is not numerical error: bad systematic states simultaneously raise default incidence and severity.

4.2 Unexpected loss

A basic volatility definition is

$$UL_\sigma = \sqrt{\mathbb{V}[L]}. \quad (4.2)$$

Economic capital is often expressed as a high loss quantile less expected loss,

$$EC_\alpha = \text{VaR}_\alpha(L) - \mathbb{E}[L]. \quad (4.3)$$

The exact institutional definition may differ. The confidence level and horizon must be explicit, and the measure must not be confused with accounting allowance or regulatory RWA.

Listing 4.1: Reconciliation-ready portfolio summary

```
1 summary = portfolio_summary(portfolio, group_by="segment")
2 summary[["segment", "exposures", "ead", "expected_loss",
3         "el_rate", "ead_weighted_pd", "hhi"]]
```

The Herfindahl-Hirschman concentration index used here is

$$HHI = \sum_i w_i^2, \quad w_i = \frac{EAD_i}{\sum_j EAD_j}. \quad (4.4)$$

It is informative but incomplete: sector, geography, collateral, guarantor, and common-driver concentration also matter.

Chapter 5

Probability of default

5.1 The estimand

A one-year PD is a conditional probability

$$\text{PD}_i(t, t+1) = \Pr(D_i = 1 \mid \mathcal{I}_t), \quad (5.1)$$

where $\mathcal{I}_t$ is information available at time t . Point-in-time PDs respond to current conditions. Through-the-cycle estimates abstract from some cyclical variation. A model should state which concept it targets.

5.2 Logistic scorecard

For feature vector x_i , logistic regression specifies

$$\eta_i = \beta_0 + x_i^\top \beta, \quad (5.2)$$

$$p_i = \frac{1}{1 + \exp(-\eta_i)}, \quad (5.3)$$

$$\log \frac{p_i}{1 - p_i} = \beta_0 + x_i^\top \beta. \quad (5.4)$$

The model is interpretable on log odds, but interpretation requires stable feature definitions, sensible transformations, and awareness of correlation and selection.

Listing 5.1: A reproducible mixed-feature PD pipeline

```
1 numeric_pipe = Pipeline([
2     ("impute", SimpleImputer(strategy="median")),
3     ("scale", StandardScaler()),
4 ])
5 categorical_pipe = Pipeline([
6     ("impute", SimpleImputer(strategy="most_frequent")),
7     ("onehot", OneHotEncoder(handle_unknown="ignore", min_frequency=0.01)),
8 ])
9 features = ColumnTransformer([
10     ("num", numeric_pipe, numeric_features),
11     ("cat", categorical_pipe, categorical_features),
12 ])
13 model = Pipeline([
14     ("features", features),
15     ("model", LogisticRegression(C=0.5, max_iter=2000,
16                                     random_state=42)),
17 ])
```

Control warning. Class weighting can improve a classification objective while destroying raw probability calibration. If weights or over-sampling are used, calibrate on an appropriate untouched sample before interpreting outputs as PDs.

5.3 Discrimination and calibration

Discrimination ranks defaults above non-defaults. With area under the ROC curve AUC, the Gini coefficient is

$$\text{Gini} = 2 \text{AUC} - 1. \quad (5.5)$$

The Kolmogorov-Smirnov statistic is

$$\text{KS} = \max_c |F_0(c) - F_1(c)|. \quad (5.6)$$

The Brier score evaluates squared probability error,

$$\text{Brier} = \frac{1}{n} \sum_{i=1}^n (y_i - p_i)^2. \quad (5.7)$$

On the held-out UCI data, the executed logistic pipeline produced AUC 0.803, Gini 0.606, KS 0.498, and Brier score 0.156. Mean predicted PD was 30.24% against an observed default rate of 30.00%. Figure 5.1 shows that good ranking does not guarantee perfect calibration in every decile.




Figure 5.1: PD holdout diagnostics on the UCI Statlog German Credit dataset. The random split is pedagogical because the source has no observation date.

5.4 Calibration to a target central tendency

When development and target default priors differ, a transparent intercept shift is

$$\text{logit}(p_i^*) = \text{logit}(p_i) + \text{logit}(\pi^*) - \text{logit}(\pi_{\text{sample}}). \quad (5.8)$$

This assumes slope stability and prior-probability shift; it does not repair covariate shift, concept drift, label bias, or an invalid population mapping.

For small or zero-default grades, a beta-binomial estimate is

$$\hat{p} = \frac{d + \alpha}{n + \alpha + \beta}, \quad (5.9)$$

where d defaults occur among n observations. A Jeffreys prior uses $\alpha = \beta = 1/2$. Pooling and conservatism still require policy and expert review.

5.5 Validation design

At minimum, assess:

- population coverage, exclusions, missingness, and feature availability;
- time- and borrower-separated holdout performance;
- discrimination, calibration, grade migration, and concentration;
- benchmark/challenger performance and uncertainty;
- stability through vintages, sectors, geographies, and adverse periods;
- fairness, prohibited variables, overrides, and decision outcomes;
- implementation parity between development and production.

Chapter 6

PD implied by prices and bank rates

6.1 Reduced-form approximation

If a credit spread s compensates only for expected default loss, recovery is R , and a constant hazard rate λ applies, then

$$\lambda \approx \frac{s}{1-R}, \quad \text{PD}(0, T) = 1 - \exp(-\lambda T). \quad (6.1)$$

At a 200 basis-point spread, 40% recovery, and five-year horizon,

$$\lambda = \frac{0.02}{0.60} = 3.33\%, \quad \text{PD}(0, 5) = 1 - e^{-0.0333 \times 5} = 15.35\%.$$




Figure 6.1: Sensitivity of five-year implied PD to spread and recovery.

Listing 6.1: Market-implied PD

```
1 def market_implied_pd(spread, recovery_rate=0.40, horizon_years=5):  
2     hazard = spread / (1 - recovery_rate)  
3     return 1 - np.exp(-hazard * horizon_years)
```

6.2 Why it is not a physical PD without adjustment

Observed spreads also reflect liquidity, risk premia, taxes, capital, market technical factors, embedded options, and measurement conventions. Therefore Equation (6.1) is a sensitivity and benchmark, not a proof that the result equals a real-world default frequency.

A simplified bank-rate break-even approximation is

$$\text{PD}_{\text{BE}} \approx \frac{r_{\text{loan}} - r_{\text{funding}} - m_{\text{operating}}}{1 - R}. \quad (6.2)$$

A real pricing system additionally handles capital charges, liquidity, optionality, fees, taxes, prepayment, utilization, acquisition costs, and target return.

Chapter 7

Loss given default

7.1 Workout LGD

For EAD at default E_0 , recovery cash flows R_k , direct and allocated costs C_k , payment times t_k , and effective interest rate r ,

$$\text{LGD}_{\text{workout}} = 1 - \frac{\sum_{k=1}^K \frac{R_k - C_k}{(1+r)^{t_k}}}{E_0}. \quad (7.1)$$

The worked example uses EUR 1,000,000 EAD, recoveries of EUR 100,000, EUR 250,000, and EUR 300,000 at 0.5, 1.5, and 3 years, costs of EUR 20,000, EUR 35,000, and EUR 45,000, and a 6% effective rate. The resulting workout LGD is 51.12%.

Listing 7.1: Discounted workout LGD

```
1 lgd = workout_lgd(  
2   ead_at_default=1_000_000,  
3   recoveries=[100_000, 250_000, 300_000],  
4   times_years=[0.5, 1.5, 3.0],  
5   costs=[20_000, 35_000, 45_000],  
6   effective_interest_rate=0.06,  
7 )
```

7.2 Long-run and downturn LGD

A simple transparent add-on is

$$\text{LGD}_{\text{DT}} = \min(1, \text{LGD}_{\text{LR}} + \max(\text{LGD}_{\text{obs,DT}} - \text{LGD}_{\text{LR}}, m)), \quad (7.2)$$

where m is a minimum margin. A robust institutional method must define downturn periods, representativeness, incomplete recoveries, secured/unsecured treatment, cure and re-default, collateral valuation, cost allocation, and uncertainty.

7.3 Modelling choices

LGD is bounded and often has mass near zero and one. Common approaches include:

- segment averages with credibility weighting;
- fractional logit or beta regression for interior values;
- two-stage models for cure/full recovery versus severity;

- survival or cash-flow models for time to recovery;
- tree ensembles used with monotonic, stability, and explainability controls.

Open recoveries create censoring. Treating every open workout as if the final recovery were already known biases estimates toward faster resolutions. Vintage analysis and explicit completion methods are essential.

Chapter 8

Exposure at default and credit conversion factors

For a revolving facility with current balance B , undrawn commitment U , and credit conversion factor c ,

$$\text{EAD} = B + cU. \quad (8.1)$$

The observed CCF at default is

$$\hat{c} = \frac{B_{\text{default}} - B_{\text{reference}}}{L_{\text{reference}} - B_{\text{reference}}}, \quad (8.2)$$

with limit L . Implementations often floor or cap the ratio, but such treatment must be documented rather than silently applied.




Figure 8.1: Observed CCF examples. Small undrawn denominators amplify ratio instability.

Listing 8.1: CCF and EAD calculation

```
1 ccf = np.clip(  
2     (balance_at_default - current_balance) / undrawn_commitment,  
3     0, 1,  
4 )  
5 ead = current_balance + ccf * undrawn_commitment
```

8.1 Operational complications

The reference date and default horizon determine which drawings count. Limit reductions, temporary excesses, blocked accounts, interest and fees, currency, netting, multiple facilities, and connected borrower behavior can all distort CCF. Validate in both ratio and currency space: a small ratio error on a large facility may be more material than a large ratio error on a nearly fully drawn account.

Practice point. Prevent post-default leakage. Collections actions, limit blocks, and default-date balances must not appear as predictors if they were unavailable at the reference date.

Chapter 9

Dependence and portfolio simulation

9.1 Gaussian one-factor construction

Let Y be a common systematic factor and ε_i an idiosyncratic factor, both standard normal and mutually independent. The obligor latent variable is

$$Z_i = \sqrt{\rho_i} Y + \sqrt{1 - \rho_i} \varepsilon_i. \quad (9.1)$$

Default occurs when

$$Z_i < \Phi^{-1}(p_i). \quad (9.2)$$

The threshold matches the unconditional PD, while the common factor generates joint defaults. Conditional on $Y = y$,

$$p_i(y) = \Phi \left(\frac{\Phi^{-1}(p_i) - \sqrt{\rho_i} y}{\sqrt{1 - \rho_i}} \right). \quad (9.3)$$

9.2 Correlated severity

The package optionally drives beta-distributed LGD through the same systematic factor with reversed sign, so weak systematic states produce both more defaults and higher severity. This wrong-way risk increases both mean and tail loss unless parameters are calibrated to preserve the desired unconditional LGD definition.

Listing 9.1: Vectorized correlated loss simulation

```
1 factor = rng.standard_normal((scenario_count, 1))
2 latent = (np.sqrt(rho) * factor
3           + np.sqrt(1 - rho) * rng.standard_normal((scenario_count, n)))
4 defaults = latent < norm.ppf(pd)
5
6 lgd_uniform = norm.cdf(
7     np.sqrt(lgd_corr) * (-factor)
8     + np.sqrt(1 - lgd_corr) * rng.standard_normal((scenario_count, n))
9 )
10 severity = beta.ppf(lgd_uniform, alpha, beta_parameter)
11 loss = (defaults * severity * ead).sum(axis=1)
```




Figure 9.1: Executed sensitivity across asset correlations. Tail measures change much more than the mean.

9.3 Simulation controls

1. Check PD and marginal LGD recovery from large simulations.
2. Reconcile mean loss to the intended analytical quantity.
3. Increase scenarios until VaR and ES are numerically stable.
4. Test seeds, chunk sizes, and vectorized versus reference implementations.
5. Repair or reject non-positive-semidefinite correlation matrices explicitly.
6. Report Monte Carlo uncertainty and preserve the random generator/seed.
7. Challenge Gaussian tail dependence using alternative factors or scenarios.

Chapter 10

Credit VaR, expected shortfall, and capital

10.1 Quantiles and tail means

For loss distribution function F_L , credit VaR is

$$\text{VaR}_\alpha(L) = \inf\{\ell : F_L(\ell) \geq \alpha\}. \quad (10.1)$$

Expected shortfall is

$$\text{ES}_\alpha(L) = \mathbb{E}[L \mid L \geq \text{VaR}_\alpha(L)] \quad (10.2)$$

for a continuous distribution; discrete implementations require a careful tail convention. The handbook run gave 99.9% VaR of EUR 5.69 million and ES of EUR 6.84 million.

VaR is a threshold, not the worst case. ES describes average severity beyond the threshold but needs more tail observations and therefore has greater estimation uncertainty. Bootstrap intervals describe Monte Carlo sampling uncertainty, not model risk from PD, LGD, dependence, or data definitions.

10.2 Basel IRB analytical benchmark

For a corporate exposure, the Basel-style asset correlation function is

$$R(p) = 0.12 \frac{1 - e^{-50p}}{1 - e^{-50}} + 0.24 \left(1 - \frac{1 - e^{-50p}}{1 - e^{-50}} \right). \quad (10.3)$$

At confidence $q = 99.9\%$, conditional default probability is

$$p_q = \Phi \left(\frac{\Phi^{-1}(p)}{\sqrt{1-R}} + \sqrt{\frac{R}{1-R}} \Phi^{-1}(q) \right). \quad (10.4)$$

With maturity adjustment

$$b(p) = (0.11852 - 0.05478 \ln p)^2, \quad (10.5)$$

$$\text{MA}(p, M) = \frac{1 + (M - 2.5)b(p)}{1 - 1.5b(p)}, \quad (10.6)$$

the unexpected-loss capital rate benchmark is

$$K = [\text{LGD } p_q - p \text{LGD}] \text{MA}(p, M), \quad \text{RWA} = 12.5 K \text{ EAD}. \quad (10.7)$$




Figure 10.1: Illustrative corporate IRB benchmark across PD and LGD.

Control warning. The implemented curve is a transparent benchmark, not a complete jurisdiction-specific capital engine. Floors, exposure classes, supervisory parameters, supporting factors, output floors, defaulted assets, provisions, and current local rules must be implemented and independently verified.

Chapter 11

Stress testing and reverse stress

11.1 Scenario transformation

Stress scenarios should start with coherent macro paths, then map those paths to risk parameters through documented satellite relationships. A stable way to shock PD is on log odds:

$$p_i^{\text{stress}} = \text{logit}^{-1}(\text{logit}(p_i) + \Delta\eta_i). \quad (11.1)$$

A macro satellite may define

$$\Delta\eta_t = \gamma_{\text{GDP}}\Delta\text{GDP}_t + \gamma_u\Delta u_t + \gamma_r\Delta r_t + \gamma_h\Delta h_t. \quad (11.2)$$

Coefficients require estimation and stability testing; the expression alone does not establish causality.

The illustrative adverse scenario applies a PD log-odds shift of 0.55, an additional PD multiplier of 1.10, LGD multiplier 1.08 plus a four percentage-point add-on, and EAD multiplier 1.05. Portfolio EL rises from EUR 275,679 to EUR 630,441.




Figure 11.1: Base and adverse expected loss by segment under the illustrative scenario.

Listing 11.1: Traceable scenario application

```
1 stressed = apply_scenario(portfolio, {  
2     "pd_log_odds_shift": 0.55,  
3     "pd_multiplier": 1.10,  
4     "lgd_multiplier": 1.08,  
5     "lgd_addon": 0.04,
```

```
6     "ead_multiplier": 1.05,  
7 })  
8 assert {"base_el", "stressed_el", "el_change"}.issubset(stressed)
```

11.2 Reverse stress

Reverse stress fixes a loss capacity C and solves for the smallest shock s^* that breaches it:

$$s^* = \inf \left\{ s : \sum_i \text{logit}^{-1}(\text{logit}(p_i) + s) \text{LGD}_i \text{EAD}_i \geq C \right\}. \quad (11.3)$$

For a capacity equal to twice base EL in the sample, the solved common log-odds shift is 0.7264. Reverse stress is a vulnerability diagnostic, not a probability forecast.

11.3 Governance questions

- Are paths internally coherent and severe but plausible for the purpose?
- Does a PIT PD already embed the same macro deterioration?
- Are downturn LGD, collateral stress, and overlays double counted?
- Do revolving exposures draw as borrower quality deteriorates?
- Are second-round, concentration, and management actions credible?
- Are scenarios, mappings, overlays, and approvals separately versioned?

Chapter 12

IFRS 9 expected credit losses

12.1 Probability-weighted cash shortfalls

IFRS 9 ECL is a probability-weighted estimate of discounted cash shortfalls. For scenario s and period t , a parameterized representation is

$$\text{ECL} = \sum_{s=1}^S w_s \sum_{t=1}^T \text{mPD}_{s,t} \text{LGD}_{s,t} \text{EAD}_{s,t} D F_t, \quad \sum_s w_s = 1. \quad (12.1)$$

If Q_t is cumulative PD through t , then unconditional marginal PD is

$$\text{mPD}_t = Q_t - Q_{t-1}, \quad Q_0 = 0. \quad (12.2)$$

Care is required because some systems use conditional hazards instead. Mixing a conditional hazard with an unconditional marginal-PD formula double counts survival.




Figure 12.1: Illustrative forward-looking PD paths and scenario-specific lifetime ECL.

With scenario weights 20% upside, 60% baseline, and 20% downside, the worked five-year example produces probability-weighted lifetime ECL of 5.238 per EUR 100 initial EAD.

Listing 12.1: Probability-weighted lifetime ECL

```
1 cumulative_pd = np.array([upside_curve, baseline_curve, downside_curve])
2 marginal_pd = marginal_pd_from_survival(cumulative_pd)
3 ecl = lifetime_ecl(
4     marginal_pd,
5     lgd=0.45,
6     ead=[100, 85, 70, 55, 40],
```

```

7     discount_factors=1 / (1.05 ** np.arange(1, 6)),
8     scenario_weights=[0.20, 0.60, 0.20],
9 )

```

12.2 Stages and significant increase in credit risk

Table 12.1: Simplified IFRS 9 staging view

Stage	Credit condition	Allowance concept
1	Performing without significant increase since origination	12-month ECL
2	Significant increase in credit risk, not credit-impaired	Lifetime ECL
3	Credit-impaired	Lifetime ECL; interest treatment changes

An illustrative policy may combine relative lifetime-PD movement, an absolute increase, 30-day past-due backstop, 90-day default backstop, watchlist status, forbearance, and qualitative indicators. Thresholds are accounting policy choices supported by analysis; they should not be tuned merely to achieve a desired allowance.

Control warning. A full implementation also needs contractual cash-flow schedules, prepayment, revolving-life policy, modifications, POCI assets, guarantees, write-offs, effective-interest discounting, scenario governance, overlays, reconciliations, controls, and disclosures.

Chapter 13

Validation, monitoring, and drift

13.1 Population stability index

For development-bin shares e_j and current shares a_j ,

$$\text{PSI} = \sum_{j=1}^J (a_j - e_j) \ln \left(\frac{a_j}{e_j} \right). \quad (13.1)$$

Small-cell protection is required when a share is zero. PSI is descriptive, not a universal statistical test. A high PSI can be benign if business strategy changed as intended; a low PSI can coexist with calibration failure if outcomes changed.

13.2 Outcome tests

Monitoring should cover multiple layers:

Table 13.1: A balanced monitoring dashboard

Layer	Examples
Data	Timeliness, completeness, ranges, missingness, category and join drift
Population	Volume, mix, scores, grades, exposures, concentration, PSI
PD outcome	Defaults, calibration by grade/time, AUC/Gini/KS, Brier, migration
LGD outcome	Recovery timing, costs, realized LGD, open-workout bias, downturn split
EAD outcome	CCF and currency error, utilization, limit changes, horizon alignment
Use	Overrides, policy exceptions, manual adjustments, approval rates, fairness
Operations	Failures, latency, fallbacks, version parity, access and incidents

For d defaults among n exposures under grade PD p , an exact binomial test is

$$H_0 : D \sim \text{Binomial}(n, p). \quad (13.2)$$

Dependence, overlapping horizons, grade migration, and low counts may invalidate the simple test or reduce its power. Confidence intervals and multi-period analysis should accompany traffic lights.

13.3 Threshold governance

Green/amber/red levels are policy decisions, not constants found in statistics textbooks. Each metric needs population, frequency, minimum count, threshold, owner, action, escalation deadline,

override authority, and evidence of closure. Thresholds should reflect model materiality and the cost of false alarms versus missed deterioration.

Chapter 14

Production workflow and evidence

14.1 A controlled analytical chain

1. Freeze immutable source snapshots and metadata.
2. Validate schemas and reconcile exposure to a trusted control total.
3. Build features with point-in-time availability and leakage controls.
4. Apply versioned model components and preserve raw outputs.
5. Apply calibration, scenario adjustments, margins, and overlays in separate columns.
6. Aggregate with traceable keys and independent reconciliation.
7. Execute unit, integration, regression, and outcome checks.
8. Write exposure results, summaries, metrics, exceptions, and a run manifest.
9. Route findings, approvals, and remediation through accountable governance.

14.2 Reproducibility manifest

Listing 14.1: Evidence manifest for an analytical run

```
1 manifest = run_manifest(  
2     inputs=["data/portfolio_2026-09-30.csv"],  
3     parameters={  
4         "model_id": "CR-PD-RETAIL-4.2",  
5         "scenario_set": "2026Q3_APPROVED",  
6         "seed": 2026,  
7         "n_scenarios": 500_000,  
8     },  
9     model_id="CR-PORTFOLIO-1.0",  
10 )  
11 write_evidence_pack(  
12     "evidence/2026Q3",  
13     portfolio_results=results,  
14     summary=summary,  
15     metrics=metrics,  
16     manifest=manifest,  
17 )
```

The manifest records input SHA-256 hashes, parameters, timestamp, Python and platform versions, and Git revision when available. It supports replication but does not replace access controls, release approval, environment locking, data retention, or human accountability.

14.3 Model inventory and change control

An inventory record should include model ID/version, owner, developer, validator, approver, purpose, prohibited uses, materiality, jurisdictions, data, estimand, methodology, dependencies, limitations, implementation location, monitoring, approval conditions, and retirement plan.

Changes should be classified before release. A code refactor with identical output, a recalibration, a new variable, a new population, and a definition change have different validation needs. Every change needs impact analysis, test evidence, approval, versioning, deployment plan, and rollback.

Chapter 15

Integrated case study

15.1 Portfolio results

Table 15.1 brings together the executed results used in the book.

Table 15.1: Executed case-study results

Metric	Result
Exposures	30
Total EAD	EUR 45,144,000
Analytical EL	EUR 275,678.63
Correlated simulated mean loss	EUR 306,480.16
99.9% VaR	EUR 5,693,113.20
99.9% expected shortfall	EUR 6,835,267.62
Zero-loss simulation frequency	48.024%
Adverse expected loss	EUR 630,441.31
Reverse-stress shift for 2× base EL	0.7264 log-odds
PD holdout AUC	0.8032
PD holdout Brier score	0.1562
IFRS 9 weighted lifetime ECL	5.2385 per EUR 100

15.2 Committee interpretation

The corporate segment dominates exposure and absolute EL, while SMEs and consumer loans show higher loss rates. The simulated distribution is highly right-skewed. Tail loss is sensitive to dependence, and positive default-LGD association lifts the simulated mean above simple component EL. Under the teaching adverse scenario, EL more than doubles. These results justify attention to concentration, scenario coherence, and component calibration rather than relying on portfolio average PD.

The analysis is not decision-ready because the portfolio is fictional, the dependence model is a transparent benchmark, scenarios are illustrative, and the UCI model lacks a time field. A real committee submission would need portfolio reconciliation, approved definitions, current institution data, independent validation, limitations, scenario ownership, uncertainty, and management actions.

15.3 Reproduction

Listing 15.1: End-to-end reproduction command sequence

```
1 python -m venv .venv
2 .venv\Scripts\Activate.ps1
3 python -m pip install -e "[dev,jupyter]"
4 pytest
5 python examples/00_end_to_end.py
6 python examples/03_pd_from_real_data.py
7 python examples/07_correlated_losses.py
8 python examples/09_stress_and_ifrs9.py
```

Chapter 16

Teaching plan and assignments

16.1 Thirty-hour sequence

Table 16.1: Suggested course delivery plan

Lab	Theme	Hours	Assessed output
1	Loss shape	2	Distribution note with mass at zero, skew, tail and concentration
2	EL and UL	3	Exposure-to-portfolio reconciliation and interpretation
3	Observed PD	4	Scorecard, holdout validation, calibration and model-card excerpt
4	Implied PD	2	Spread/recovery/horizon sensitivity with bias discussion
5	Workout LGD	3	Recovery cash-flow schedule and downturn adjustment
6	Revolver EAD	3	CCF reconstruction and vintage/denominator analysis
7	Dependence	4	Independent versus common-factor simulation
8	VaR and capital	4	Tail-risk memo, uncertainty and IRB comparison
9	Stress and IFRS	5	Scenario pack, staging example and reverse stress

9

16.2 Capstone specification

Teams receive a dated, de-identified portfolio and data dictionary. They must:

1. define PD, LGD, EAD, horizons, and target populations;
2. reconcile inputs and identify point-in-time leakage risks;
3. build transparent component benchmarks and one challenger;
4. validate out of time where the data permit;
5. simulate loss under at least two dependence assumptions;
6. produce VaR, ES, stress, reverse stress, and sensitivity results;
7. document limitations, model use, monitoring, and governance;
8. deliver code, tests, evidence manifest, and committee memo.

Suggested grading: 15% scope and estimand, 15% data/lineage, 20% estimation, 20% validation, 15% stress/tail interpretation, 10% reproducibility, and 5% communication. Target leakage, irreproducibility, hidden adjustments, or material exposure non-reconciliation should be pass gates rather than minor deductions.

16.3 Advanced extensions

- discrete-time survival PD with censoring and competing risks;
- multi-state migration, cure, and re-default processes;
- incomplete-recovery LGD with vintage completion;
- t-copula, sector factors, and granularity adjustments;
- extreme-value tail estimation and importance sampling;
- reject inference and selection-bias sensitivity;
- explainability stability and fairness outcome review;
- scenario-conditioned lifetime ECL and overlay backtesting.

Part II

Advanced credit markets and portfolio models

Chapter 17

Ratings, migration, and default intensity

17.1 External ratings, internal grades, and what they estimate

A rating is an ordered credit state, not a probability by itself. External agency ratings support capital-market comparison; internal grades support underwriting, limits, pricing, monitoring, and capital. A useful rating system therefore needs an explicit obligor or facility level, horizon, default definition, calibration philosophy, review frequency, and treatment of overrides.

Point-in-time (PIT) grades respond strongly to current conditions. Through-the-cycle (TTC) grades suppress some temporary variation. Neither label is a complete specification: two TTC systems can differ materially in cyclicalities and default rate. Institutions should measure grade migration, override rates, realized default rates, and stability through time, not rely on the label alone.

Accounting ratios can provide transparent screening variables. For example, a classic Altman-style manufacturing score takes the form

$$Z = 1.2 \frac{\text{working capital}}{\text{total assets}} + 1.4 \frac{\text{retained earnings}}{\text{total assets}} + 3.3 \frac{\text{EBIT}}{\text{total assets}} + 0.6 \frac{\text{market equity}}{\text{book liabilities}} + 1.0 \frac{\text{sales}}{\text{total assets}}. \quad (17.1)$$

The coefficients and interpretation belong to the population and era in which the model was developed. The score is a useful feature or challenger, not a universal cut-off. A bank must re-estimate or validate it for its legal entities, industries, accounting regimes, and observation dates.

17.2 Transition matrices

Let P be a one-year matrix whose element p_{ij} is the probability of moving from grade i to grade j . Every row must be nonnegative and sum to one. Default is normally absorbing, so its row has probability one in the default column. Under a time-homogeneous first-order Markov assumption,

$$P(n) = P^n. \quad (17.2)$$

For a fractional horizon t , a continuous embedding is sometimes approximated by $P(t) = \exp(t \log P)$. Not every empirical matrix has a valid generator; complex or materially negative entries are a warning, not values to clip silently.

Figure 17.1 shows an intentionally compact teaching matrix and its five-year power. The default column accumulates probability, while migration mass spreads away from the original grade.



Figure 17.1: One-year rating transitions and the five-year Markov projection

The executed example produces five-year default probabilities of 1.75% for AAA, 5.68% for A, 14.76% for BBB, and 37.51% for BB. These are properties of the stylized matrix, not agency estimates. Their purpose is to make matrix mechanics auditable.

Listing 17.1: Rating-matrix horizon conversion

```

1 import pandas as pd
2 from creditrisk_lab.ratings import transition_matrix_horizon
3
4 annual = pd.read_csv(
5     "data/stylized_transition_matrix.csv", index_col=0
6 )
7 five_year = transition_matrix_horizon(annual.to_numpy(), 5)
8 default_probability = dict(zip(annual.index, five_year[:, -1]))
9 print(default_probability)

```

The Markov assumption excludes rating momentum: after a downgrade, the chance of another downgrade may be higher than the grade alone implies. A practical test adds previous direction, time in grade, watch-list status, or outlook to the transition model and compares out-of-time likelihood and calibration.

17.3 Cumulative, marginal, and conditional default probability

Let $F(t) = \Pr(\tau \leq t)$ and $S(t) = 1 - F(t) = \Pr(\tau > t)$. The marginal default probability for interval $(t_{k-1}, t_k]$ is

$$q_k = S(t_{k-1}) - S(t_k), \quad (17.3)$$

and the conditional interval probability is

$$\Pr(t_{k-1} < \tau \leq t_k \mid \tau > t_{k-1}) = \frac{q_k}{S(t_{k-1})}. \quad (17.4)$$

The hazard rate is the instantaneous conditional default intensity,

$$\lambda(t) = \lim_{\Delta t \downarrow 0} \frac{\Pr(t < \tau \leq t + \Delta t \mid \tau > t)}{\Delta t}, \quad S(t) = \exp \left[- \int_0^t \lambda(u) du \right]. \quad (17.5)$$

For a rough constant-intensity estimate from a credit spread s and recovery R , $\lambda \approx s/(1 - R)$. The package also builds piecewise-constant forward hazards from maturity-specific spreads.

Figure 17.2 instead uses the latest complete observed FRED ICE BofA rating OAS cross-section in the bundled snapshot, dated 2 October 2026: AAA 41 bp, AA 60 bp, A 72 bp, BBB 104 bp, and BB 191 bp. These are cash-bond index OAS observations, not CDS quotes.




Figure 17.2: Observed rating OAS and constant-hazard survival proxies

At 40% recovery, the BBB OAS gives a 1.733% constant hazard proxy and 84.09% ten-year survival. The rating curves are comparisons across credit quality, not a maturity term structure. A production bootstrap instead uses maturity-matched tradable quotes and solves premium and protection legs jointly with the discount curve, schedule, accrual on default, recovery, and quotation conventions.

Practice point. Keep historical or physical PD separate from market-implied, risk-neutral PD. They answer different questions. Pricing includes risk premia, liquidity, and technical effects; forecasting targets realized frequencies.

Chapter 18

Credit derivatives, spreads, and basis

18.1 Single-name CDS cash flows

A protection buyer pays a periodic coupon on notional N until maturity or credit event. If default occurs, the protection payment is approximately

$$\text{protection payment} = N(1 - R), \quad (18.1)$$

where R is the contract recovery or auction recovery used in settlement. Cash settlement dominates contemporary standardized markets; deliverable obligations and cheapest-to-deliver optionality still matter to documentation and basis risk.

With payment dates t_k , accrual fractions α_k , discount factors D_k , and survival S_k , a discrete premium annuity is

$$A = \sum_k \alpha_k D_k S_k + \frac{1}{2} \sum_k \alpha_k D_k (S_{k-1} - S_k), \quad (18.2)$$

where the second term approximates accrued premium on default. The protection leg per unit notional is

$$B = (1 - R) \sum_k D_k (S_{k-1} - S_k), \quad (18.3)$$

and the par spread is $s^* = B/A$. For contractual coupon c , the buyer's clean upfront amount under this sign convention is

$$U = N(B - cA). \quad (18.4)$$

Listing 18.1: CDS legs, par spread, and upfront amount

```
1 import numpy as np
2 import pandas as pd
3 from creditrisk_lab.cds import cds_legs, cds_upfront
4
5 oas = pd.read_csv("data/real/fred_corporate_oas_daily.csv")
6 rates = pd.read_csv("data/real/fred_credit_conditions.csv")
7 spread = oas["BAMLCOA4CBBB"].dropna().iloc[-1] / 100
8 risk_free = rates["DGS5"].dropna().iloc[-1] / 100
9 t = np.arange(0.25, 5.01, 0.25)
10 discount = np.exp(-risk_free * t)
11 survival = np.exp(-(spread / 0.60) * t)
12 legs = cds_legs(t, discount, survival, recovery_rate=0.40)
13 upfront = cds_upfront(
14     10_000_000, legs["par_spread"], 0.01, legs["premium_annuity"]
15 )
16 print(legs["par_spread"] * 10_000, upfront)
```

Using the observed 104 bp BBB OAS on 2 October 2026 as a clearly labelled hazard proxy and the observed 5.06% five-year Treasury rate from FRED, the calculation gives a five-year par spread of 104.66 bp. A buyer paying a standardized 100 bp coupon pays about EUR 19,624 upfront on EUR 10 million notional. With 35% recovery, a default settlement on the same notional would be EUR 6.5 million. Figure 18.1 shows the CDS proxy and its accumulating legs. It is not represented as an observed CDS market quote.




Figure 18.1: CDS proxy from observed BBB OAS and leg accumulation

18.2 Indices and institutional uses

Credit indices provide liquid exposure to a defined basket. Their fixed coupons, roll dates, constituent changes, and defaulted-name treatment mean that quoted index spread, price, and intrinsic single-name value are not interchangeable. Institutions use CDS and indices for hedging, relative-value analysis, warehouse risk, macro overlays, and counterparty-risk mitigation. Hedge effectiveness depends on reference entity, seniority, maturity, currency, restructuring clause, documentation, and liquidity.

18.3 Bond spreads, asset swaps, and CDS-bond basis

A corporate bond spread includes expected default loss, systematic risk premium, liquidity, funding, tax, optionality, benchmark choice, and technical supply and demand. Therefore it is not a pure PD quote. The practical spread depends on the selected risk-free proxy: government, overnight-index swap, or another curve.

An asset swap packages a bond with an interest-rate swap to turn fixed coupons into a floating reference plus spread. The CDS-bond basis is often reported as

$$\text{basis} = s_{\text{CDS}} - s_{\text{bond}}, \quad (18.5)$$

provided both are matched for maturity, seniority, currency, and curve convention. A positive basis is not automatically an arbitrage: funding, repo, counterparty exposure, deliverability, liquidity, capital, and transaction costs can prevent convergence. Version 2.1 deliberately reports no numerical basis: an honest basis requires licensed, matched CDS and cash-bond quotes. Fabricating one would defeat the data-provenance objective.

Control warning. Do not convert a quoted spread into an approved PD by one formula without documenting recovery, curve, risk premium, liquidity, instrument features, and calibration purpose. The simple hazard approximation is a diagnostic benchmark, not a full market calibration.

Chapter 19

Structural credit models

19.1 Equity as an option on firm assets

In the Merton framework, firm asset value V_A follows a diffusion and debt with face value D is due at horizon T . Equity is a call option on assets:

$$E = V_A \Phi(d_1) - De^{-rT} \Phi(d_2), \quad (19.1)$$

$$d_1 = \frac{\ln(V_A/D) + (r + \frac{1}{2}\sigma_A^2)T}{\sigma_A \sqrt{T}}, \quad d_2 = d_1 - \sigma_A \sqrt{T}. \quad (19.2)$$

Observed equity volatility supplies a second equation,

$$\sigma_E E = \Phi(d_1) \sigma_A V_A. \quad (19.3)$$

The two nonlinear equations are solved for unobserved asset value V_A and asset volatility σ_A .

The risk-neutral default probability is

$$\text{PD}^Q = \Phi(-d_2). \quad (19.4)$$

For forecasting under an assumed physical asset drift μ , distance to default and physical PD can be written

$$\text{DD} = \frac{\ln(V_A/D) + (\mu - \frac{1}{2}\sigma_A^2)T}{\sigma_A \sqrt{T}}, \quad \text{PD}^P = \Phi(-\text{DD}). \quad (19.5)$$

Using r in place of μ changes the measure and the interpretation.

Listing 19.1: Solving the Merton equations

```
1 from creditrisk_lab.structural import merton_model
2
3 result = merton_model(
4     equity_value=3.0,
5     equity_volatility=0.80,
6     debt_face_value=10.0,
7     risk_free_rate=0.05,
8     maturity_years=1.0,
9     asset_drift=0.09,
10 )
11 print(result)
```

The executed result is $V_A = 12.3954$, $\sigma_A = 21.23\%$, distance to default 1.1408, risk-neutral PD 12.70%, physical PD 9.19%, and implied debt market value 9.3954 in the same currency units. Figure 19.1 shows the expected increase in PD as leverage and observed equity volatility rise.




Figure 19.1: Merton risk-neutral PD sensitivity to debt and equity volatility

19.2 Implementation and validation

A workplace implementation should use a defensible default point rather than blindly treating all liabilities as one-year debt; synchronize equity value, volatility, debt, currency, and date; handle thinly traded equity; record solver diagnostics; and compare structural PD with observed defaults, market spreads, and internal grades. Newton or generic root solvers can converge to economically invalid roots, so positivity, residual size, monotonicity, and sensitivity tests are mandatory.

The model is most natural for listed firms with observable equity. Its basic form assumes lognormal assets, constant volatility, a single liability date, frictionless markets, and default only at the horizon. Extensions add first passage, jumps, stochastic rates, multiple debt classes, and accounting information, but each extension introduces new identification and data questions.

Practice point. Structural, accounting-score, and reduced-form models are complements. Structural models connect market leverage to default; scorecards rank using observed borrower attributes; intensity models price event timing. Model choice starts with the decision and measure, not the most elaborate method.

Chapter 20

Counterparty exposure, netting, and collateral

20.1 Exposure is a future distribution

For a derivative portfolio with simulated future mark-to-market V_t , unsecured positive and negative exposures are

$$E_t^+ = \max(V_t - C_t, 0), \quad E_t^- = \max(C_t - V_t, 0), \quad (20.1)$$

where C_t is collateral held, under the implemented sign convention. Summary profiles include

$$EE(t) = \mathbb{E}[E_t^+], \quad ENE(t) = \mathbb{E}[E_t^-], \quad (20.2)$$

$$PFE_q(t) = \text{quantile}_q(E_t^+), \quad DEE(t) = D(0, t)EE(t). \quad (20.3)$$

Expected positive exposure over a horizon is commonly summarized as expected positive exposure (EPE), but regulatory effective EPE and internal measures can apply additional rules. Names that sound similar must not be substituted.

20.2 Netting sets and collateral agreements

Exposure is aggregated inside a legally enforceable netting set, not automatically by customer or trade. For trade values $V_{j,t}$ in a set,

$$V_t = \sum_j V_{j,t}, \quad E_t^+ = \max(V_t - C_t, 0). \quad (20.4)$$

Computing $\sum_j \max(V_{j,t}, 0)$ loses the diversification provided by valid close-out netting. Conversely, aggregating across contracts without enforceable netting understates risk.

Collateral simulation should reflect threshold, minimum transfer amount, independent amount or initial margin, eligible collateral, haircuts, settlement lag, margin frequency, disputes, cure period, and rehypothecation. A simple lagged variation-margin rule is

$$C_t = \max(V_{t-\ell} - H, 0) \quad (20.5)$$

after applying minimum-transfer logic, where H is threshold and ℓ is the operational lag. Even daily margining leaves gap risk during the margin period of risk.

Listing 20.1: Simulated exposure with lagged collateral

```
1 from creditrisk_lab.counterparty import (  
2     collateral_from_lagged_mtm, exposure_profile
```

```

3 )
4
5 collateral = collateral_from_lagged_mtm(
6     mtm_paths,
7     lag_steps=1,
8     threshold=100_000,
9     minimum_transfer_amount=25_000,
10 )
11 profile = exposure_profile(
12     mtm_paths,
13     collateral_paths=collateral,
14     discount_factors=discount,
15     pfe_quantile=0.975,
16 )

```

Figure 20.1 is generated from 80,000 market paths for a decaying five-year position. Collateral reduces expected exposure substantially, yet the one-period lag and threshold preserve a nonzero tail. The maximum 97.5% PFE is about EUR 927,017.




Figure 20.1: Expected exposure, collateralized exposure, and PFE

20.3 Simulation design and controls

Exposure engines usually simulate market factors under a risk-neutral measure for valuation, then revalue trades or use validated approximations at future dates. Controls include trade-to-ledger reconciliation, curve and market-data lineage, time-grid adequacy, collateral-state testing, legal-netting evidence, benchmark revaluation, convergence, sensitivity, and reproduction of limit reports. Downgrade triggers may demand collateral precisely when a counterparty's liquidity deteriorates; they should be scenario variables rather than assumed benefits.

Chapter 21

CVA, DVA, incremental risk, and wrong-way risk

21.1 Discrete unilateral CVA

Let $q_k^Q = S^Q(t_{k-1}) - S^Q(t_k)$ be the counterparty's marginal risk-neutral default probability, D_k a discount factor, and EE_k expected positive exposure. With deterministic recovery,

$$\text{CVA} = (1 - R_C) \sum_k D_k q_{C,k}^Q EE_k. \quad (21.1)$$

This factorization assumes exposure and default are independent after the modeled conditioning. If they are dependent, CVA is a scenario expectation of exposure times default loss and discounting, not the product of separate averages.

Own-default valuation adjustment uses negative exposure and own marginal default probability,

$$\text{DVA} = (1 - R_O) \sum_k D_k q_{O,k}^Q ENE_k. \quad (21.2)$$

A simplified bilateral value is $V = V_{\text{risk-free}} - \text{CVA} + \text{DVA}$. Close-out, funding, collateral, and accounting policies make production valuation more complex; the formula is an attribution scaffold.

The executed collateralized example gives CVA EUR 5,674.73 and DVA EUR 3,445.24. The amounts are deliberately small and are not a market quote. They verify that the exposure, survival, recovery, and discounting interfaces reconcile.

21.2 Incremental CVA and allocation

The incremental CVA of a new trade is

$$\Delta\text{CVA} = \text{CVA}(\text{netting set} + \text{trade}) - \text{CVA}(\text{netting set}). \quad (21.3)$$

It need not equal standalone CVA because netting and collateral are nonlinear. A trade with positive standalone exposure can reduce portfolio CVA if it offsets existing exposure in the relevant states. Desk pricing, limit usage, accounting valuation, and regulatory capital may allocate different quantities; policies should name the one being charged.

21.3 Wrong-way and right-way risk

Wrong-way risk (WWR) occurs when exposure is high in states where counterparty credit quality is weak. Right-way risk is the opposite. General WWR comes from broad macro dependence;

specific WWR arises from a direct link, such as accepting a counterparty's own securities as collateral or writing protection on a closely related entity.

In a scenario engine, default likelihood can depend on the same market factor that drives exposure. The package's illustration increases interval default weights in adverse market states and then renormalizes the scenario-by-time probabilities. CVA rises from EUR 5,674.73 to EUR 7,513.33, a 32.4% increase.

Listing 21.1: Independent and wrong-way CVA comparison

```
1 from creditrisk_lab.counterparty import cva, scenario_cva
2
3 independent = cva(marginal_pd, expected_exposure, discount)
4 wrong_way = scenario_cva(
5     positive_exposure_paths,
6     conditional_default_probability_paths,
7     discount,
8     recovery_rate=0.40,
9 )
```

Regulatory multipliers sometimes represent residual dependence or model uncertainty. A scalar multiplier is not a substitute for identifying material specific WWR, validating joint dynamics, or setting limits.

21.4 CVA market risk and governance

CVA changes with credit spreads, interest rates, FX, volatility, correlation, collateral terms, and the portfolio itself. A CVA desk therefore manages both the level and sensitivities. Finite-difference Greeks should use common random numbers, stable bump sizes, and numerical error analysis. Hedge P&L should be attributed to spread, market, new-trade, time-decay, model, and residual effects. Model risk includes curve construction, proxy spreads, sparse counterparties, recovery, WWR, simulation, and legal terms.

Control warning. DVA is not a reserve that can necessarily be monetized. CVA, DVA, funding valuation adjustments, accounting fair value, and regulatory CVA capital must be reconciled by framework and owner. Similar names do not imply identical cash flows, measures, or governance.

Chapter 22

Portfolio credit engines: Vasicek, CreditRisk+, and CreditMetrics

22.1 Three questions, three model families

Portfolio models differ in the state variable they simulate. The asymptotic single-risk-factor Vasicek model conditions Bernoulli default on one systematic factor and is useful for capital benchmarks. CreditRisk+ models default counts with Poisson mixtures and is computationally efficient for default-mode loss. CreditMetrics models rating migrations and revaluation, capturing spread changes as well as default. These are alternative views, not successive accuracy levels.

22.2 Vasicek conditional default rate

For unconditional PD p , asset correlation ρ , and systematic realization $Y = y$, conditional PD is

$$p(y) = \Phi \left(\frac{\Phi^{-1}(p) - \sqrt{\rho} y}{\sqrt{1 - \rho}} \right). \quad (22.1)$$

At confidence α , the adverse conditional default rate is commonly written

$$\text{WCDR}_\alpha = \Phi \left(\frac{\Phi^{-1}(p) + \sqrt{\rho} \Phi^{-1}(\alpha)}{\sqrt{1 - \rho}} \right). \quad (22.2)$$

The asymptotic result assumes a fine-grained homogeneous portfolio. Name and sector concentration require simulation or a granularity adjustment.

22.3 CreditRisk+: Poisson counts and stochastic intensity

For n homogeneous obligors with rare default probability p , default count is approximated as $N \sim \text{Poisson}(\mu)$ with $\mu = np$. If the intensity Λ is gamma distributed with mean μ and variance σ_Λ^2 , mixing the Poisson distribution produces a negative-binomial count. One useful parameterization is

$$r = \frac{\mu^2}{\sigma_\Lambda^2}, \quad p_{\text{NB}} = \frac{r}{r + \mu}, \quad (22.3)$$

where this p_{NB} is a distribution parameter, not obligor PD. As intensity volatility increases, the upper tail becomes much heavier.

Listing 22.1: CreditRisk+ default-count tail

```
1 from creditrisk_lab.creditrisk_plus import (
```

```

2     default_count_distribution, default_count_quantile
3 )
4
5 pmf = default_count_distribution(
6     mean_defaults=300.0, # observed UCI defaults
7     std_defaults=75.0,   # sensitivity: 25% of the mean
8     max_defaults=650,
9 )
10 q999 = default_count_quantile(300.0, 75.0, 0.999)

```

The UCI snapshot contains 1,000 observed applicants and 300 defaults. Using 300 as the mean count, the 99.9% quantile is 355 under fixed intensity, 470 when intensity standard deviation is 15% of the mean, and 593 at 25%. The mean is observed; the intensity volatility is explicitly a sensitivity because the UCI cross-section has no time axis. The duration-band severity illustration produces 99.9% loss VaR of about EUR 866,706 and ES of EUR 903,105.




Figure 22.1: CreditRisk+ tails anchored to the observed UCI default count

Production CreditRisk+ implementations may use sector factors, exposure bands, probability generating functions, and recursion. Controls must address the Poisson approximation, obligor concentrations, sector mapping, stochastic intensity calibration, severity discretization, and tail truncation.

22.4 CreditMetrics: migration and revaluation

CreditMetrics begins with a transition matrix. For each starting grade, cumulative transition probabilities map to standard-normal thresholds. A correlated latent variable selects the terminal state; the instrument is then revalued using the state-specific value. With initial value V_0 and terminal value V_s ,

$$L = V_0 - V_s. \quad (22.4)$$

Migration gains can therefore be negative losses. Default recovery and non-default spread curves must be consistent with horizon and instrument cash flows.

Listing 22.2: Correlated CreditMetrics migration loss

```

1 from creditrisk_lab.creditmetrics import (
2     simulate_migration_losses, migration_risk_metrics
3 )
4
5 losses = simulate_migration_losses(

```

```

6     initial_states,
7     transition_matrix,
8     current_values,
9     state_value_factors,
10    asset_correlation=0.20,
11    n_scenarios=100_000,
12    seed=41,
13 )
14 metrics = migration_risk_metrics(losses, confidence=0.999)

```

The executed portfolio has mean migration loss EUR 181,619, 99.9% VaR EUR 2.74 million, and ES EUR 3.183 million. Figure 22.2 separates the central distribution from the extreme tail.




Figure 22.2: CreditMetrics migration-loss distribution and upper tail

Table 22.1: Selecting a portfolio model by purpose

Model	Main state	Strength	Main limitation
Vasicek	Default indicator	Transparent systematic-factor benchmark	Asymptotic and usually one-period
CreditRisk+	Default counts	Fast, tractable default-mode tail	Migration and mark-to-market effects omitted
CreditMetrics	Rating state and value	Captures migration and default revaluation	Transition and valuation inputs are demanding
Full simulation	Market and credit states	Flexible netting, WWR, and path dependence	Expensive and model-risk intensive

Chapter 23

Credit-spread risk and constant-risk portfolios

23.1 Spread duration and convexity

For small parallel spread change Δs , price loss relative to current price can be approximated by

$$\frac{\Delta L}{P} \approx \text{SD} \Delta s - \frac{1}{2} \text{SC}(\Delta s)^2, \quad (23.1)$$

where SD is spread duration and SC is spread convexity under the chosen sign convention. The approximation must use decimal spread changes: 100 bp is 0.01. Key-rate or bucketed spread durations are preferable when curve twists matter.

Historical simulation applies joint historical changes to current sensitivities. It preserves empirical cross-name and cross-tenor dependence in the selected window but assumes that historical scenarios are relevant to today's portfolio. The package implementation accepts a matrix of joint spread changes and returns scenario P&L, VaR, and ES.

Listing 23.1: Joint historical spread VaR and ES

```
1 from creditrisk_lab.spread_risk import historical_spread_risk
2
3 result = historical_spread_risk(
4     market_values,
5     spread_durations,
6     historical_changes_bps, # rows are dates; changes are basis points
7     confidence=0.99,
8     spread_convexities=spread_convexities,
9 )
10 print(result["var"], result["expected_shortfall"])
```

The executed example uses 775 observed joint daily changes in FRED AAA, BBB, and BB option-adjusted spreads. Applied to explicitly hypothetical current positions, it gives 99% spread VaR EUR 15,230 and ES EUR 24,564. Thus the shocks are real; only the position vector and durations are teaching assumptions.




Figure 23.1: Observed FRED OAS shocks applied to hypothetical current positions

23.2 Jump to default and nonlinear instruments

Spread history alone can miss the discontinuous loss at default, particularly when defaulted issuers disappear from the clean time series. A complete credit trading risk framework therefore combines spread migration, jump-to-default, recovery uncertainty, basis risk, curve risk, and option nonlinearity. CDS options, callable bonds, and tranche products require full revaluation or validated nonlinear approximations.

23.3 Constant level of risk versus buy and hold

There are two distinct backtesting questions. A buy-and-hold simulation follows the original positions as ratings, maturities, and market values evolve. A constant-risk simulation periodically rebalances to restore today-like rating, duration, and concentration. The second design is often closer to a risk measure for an actively managed desk, but it embeds a trading strategy and transaction assumption. Both should disclose reinvestment, turnover, defaults, exits, new issuance, and liquidity. Comparing the two reveals how much reported tail risk is driven by portfolio evolution rather than the chosen statistical model.

Control warning. Never splice spread histories without survivor-bias controls. Retain defaulted, downgraded, merged, and delisted names; define how stale or missing quotes enter scenarios; and document the current-to-history mapping.

Chapter 24

Advanced integration, governance, and teaching cases

24.1 A single measure cannot serve every decision

Table 24.1 is a practical map of measures, data, and uses. The critical separation is between physical measures for forecasting and provisioning, risk-neutral measures for valuation, and stress or regulatory transforms for capital and limits.

Table 24.1: Measure-purpose map

Decision	Measure	Primary evidence	Key control
Underwriting/limits	Physical PIT/TTC PD, LGD, EAD	Borrower and performance history	Date integrity and target definition
IFRS 9	Probability-weighted discounted ECL	Historical performance plus forecasts	Staging, scenario governance, overlays
CDS/fair value	Risk-neutral hazard and recovery	Tradable curves and contracts	Quote, curve, and convention controls
Counterparty pricing	CVA/DVA and sensitivities	Exposure simulation and credit curves	Netting, collateral, WWR, reserves
Economic capital	Loss distribution, VaR/ES	Portfolio, dependence, severity	Concentration and tail uncertainty
Trading limits	Spread, jump, basis, migration risk	Positions and joint market history	Full revaluation and survivor bias

24.2 Expanded executed results

Table 24.2: Version 2.1 advanced-workflow results and provenance

Metric	Executed result
Observed FRED rating OAS date	2 October 2026
Observed AAA / BBB / BB OAS	41 / 104 / 191 bp
BBB hazard proxy / ten-year survival	1.733% / 84.086%
Five-year CDS proxy / upfront at 100 bp	104.658 bp / EUR 19,623.60
Observed UCI records / defaults	1,000 / 300
Merton risk-neutral / physical PD	12.697% / 9.189%
Collateralized CVA / DVA	EUR 5,674.73 / EUR 3,445.24
Wrong-way CVA	EUR 7,513.33
CreditRisk+ 99.9% count, $\sigma_{\Lambda} = 25\%\mu$	593 defaults
CreditMetrics 99.9% VaR / ES	EUR 2.740m / EUR 3.183m
Spread 99% VaR / ES, 775 observed shocks	EUR 15,230 / EUR 24,564

The OAS, UCI, FRED Treasury, and historical spread-shock rows use observed snapshots. The Merton firm, derivative exposure paths, CVA/DVA netting set, and CreditMetrics position book remain explicitly simulated because no equivalent open, redistributable production portfolios exist. Their purpose is algorithm and control verification, not empirical benchmarking.

24.3 Institutional implementation blueprint

An end-to-end implementation can be governed as eight linked layers:

1. contract, customer, facility, collateral, and market-data ingestion;
2. dated definitions, populations, horizons, and legal netting sets;
3. physical PD/LGD/EAD estimation and market hazard calibration;
4. exposure, migration, default, recovery, and spread simulation;
5. valuation, ECL, capital, CVA, limits, and scenario aggregation;
6. independent validation and implementation verification;
7. monitoring, P&L or loss attribution, thresholds, and escalation;
8. controlled release, evidence retention, and accountable use.

Each layer needs an owner and reconciliation. A sophisticated simulator cannot repair a stale legal agreement, wrong sign, mismatched curve date, leaked target, or incomplete default history.

24.4 Advanced laboratory sequence

Table 24.3: Additional advanced laboratories

Lab	Theme	Hours	Practitioner deliverable
10	Ratings and hazard curves	3	Matrix validation, five-year migration, hazard and survival memo

Lab	Theme	Hours	Practitioner deliverable
11	CDS and basis	3	Leg reconciliation, upfront quote, recovery and curve sensitivity
12	Structural model	3	Solver evidence, distance-to-default and market comparison
13	Counterparty exposure	4	EE/PFE profiles, netting and collateral agreement tests
14	CVA and WWR	4	CVA/DVA attribution, incremental trade and WWR scenario
15	CreditRisk+	3	Count-tail calibration and sector-severity sensitivity
16	CreditMetrics	4	Migration thresholds, correlated loss, VaR and ES
17	Spread risk	3	Joint historical VaR, jump-to-default and concentration note
18	Real-data audit	2	Snapshot inventory, provenance checks and refresh evidence

24.5 Advanced capstone

Students or practitioner teams should combine a bond/CDS book and derivative netting set. Required outputs are: (i) physical and risk-neutral PD term structures; (ii) a reconciled CDS quote and basis analysis; (iii) a structural-model challenger; (iv) collateralized EE and PFE; (v) CVA, DVA, incremental CVA, and one WWR scenario; (vi) CreditRisk+ and CreditMetrics tail comparisons; (vii) joint spread VaR with jump-to-default overlay; and (viii) a model-use and limitations memorandum.

The assessment should reward reconciliation and interpretation as much as model complexity. Every headline number must link to source data, parameter version, code version, execution timestamp, and an owner-approved definition.

Appendix A

Formula reference

Table A.1: Core formula sheet

Quantity	Formula
Exposure loss	$L_i = D_i \text{LGD}_i \text{EAD}_i$
Portfolio loss	$L = \sum_i L_i$
Expected loss	$\text{EL}_i = \text{PD}_i \text{LGD}_i \text{EAD}_i$
Independent Bernoulli variance	$\mathbb{V}[L_i] = p_i(1 - p_i)(\text{LGD}_i \text{EAD}_i)^2$
Economic-capital form	$\text{EC}_\alpha = \text{VaR}_\alpha - \mathbb{E}[L]$
Logistic PD	$p = (1 + e^{-x^\top \beta})^{-1}$
Prior shift	$\text{logit}(p^*) = \text{logit}(p) + \text{logit}(\pi^*) - \text{logit}(\pi)$
Hazard approximation	$\lambda \approx s/(1 - R)$
Cumulative implied PD	$1 - e^{-\lambda T}$
Workout LGD	$1 - \sum_k (R_k - C_k)(1 + r)^{-t_k} / E_0$
CCF	$(B_D - B_0)/(L_0 - B_0)$
EAD	$B + cU$
One-factor latent variable	$Z_i = \sqrt{\rho_i} Y + \sqrt{1 - \rho_i} \varepsilon_i$
VaR	$\inf\{\ell : F_L(\ell) \geq \alpha\}$
Expected shortfall	$\mathbb{E}[L \mid L \geq \text{VaR}_\alpha]$
IFRS 9 ECL	$\sum_s w_s \sum_t \text{mPD}_{s,t} \text{LGD}_{s,t} \text{EAD}_{s,t} D F_t$
PSI	$\sum_j (a_j - e_j) \ln(a_j / e_j)$
HHI	$\sum_i (\text{EAD}_i / \sum_j \text{EAD}_j)^2$
Transition horizon	$P(n) = P^n$
Survival from hazard	$S(t) = \exp[-\int_0^t \lambda(u) du]$
CDS par spread	$s^* = (1 - R) \sum_k D_k (S_{k-1} - S_k) / A$
CDS upfront	$U = N(B - cA)$
Merton equity	$E = V_A \Phi(d_1) - D e^{-rT} \Phi(d_2)$
Distance to default	$[\ln(V_A/D) + (\mu - \sigma_A^2/2)T] / (\sigma_A \sqrt{T})$
Expected exposure	$\text{EE}(t) = \mathbb{E}[\max(V_t - C_t, 0)]$
CVA	$(1 - R_C) \sum_k D_k q_{C,k}^Q \text{EE}_k$
DVA	$(1 - R_O) \sum_k D_k q_{O,k}^Q \text{ENE}_k$
Vasicek WCDR	$\Phi\{[\Phi^{-1}(p) + \sqrt{\rho} \Phi^{-1}(\alpha)] / \sqrt{1 - \rho}\}$
Gamma-Poisson shape	$r = \mu^2 / \sigma_\Lambda^2$
Spread loss	$\Delta L / P \approx \text{SD} \Delta s - \frac{1}{2} \text{SC}(\Delta s)^2$

Appendix B

Python cookbook

B.1 Load and validate

```
1 import pandas as pd
2 from creditrisk_lab.schemas import validate_portfolio
3 from creditrisk_lab.portfolio import portfolio_summary
4
5 portfolio = validate_portfolio(pd.read_csv("data/sample_portfolio.csv"))
6 print(portfolio_summary(portfolio, group_by="segment"))
```

B.2 Fetch public macro data

```
1 from creditrisk_lab.data import fred_series, world_bank_indicator
2
3 fred = fred_series(["UNRATE", "DRALACBN", "BAA10Y"], start="2000-01-01")
4 gdp = world_bank_indicator(
5     "NY.GDP.MKTP.KD.ZG",
6     countries=["PL", "DE", "CZ"],
7     start_year=2000,
8     end_year=2025,
9 )
```

B.3 Simulate and calculate tail risk

```
1 from creditrisk_lab.simulation import simulate_credit_losses
2 from creditrisk_lab.tail import value_at_risk, expected_shortfall
3
4 losses = simulate_credit_losses(
5     portfolio,
6     n_scenarios=500_000,
7     asset_correlation=0.15,
8     lgd_systematic_correlation=0.25,
9     seed=2026,
10 )
11 print(value_at_risk(losses, 0.999))
12 print(expected_shortfall(losses, 0.999))
```

B.4 Stress and retain both states

```

1 from creditrisk_lab.stress import apply_scenario
2
3 result = apply_scenario(portfolio, scenario)
4 result[["exposure_id", "base_el", "stressed_el", "el_change"]]

```

B.5 Monitoring

```

1 from creditrisk_lab.validation import (
2     binary_model_metrics,
3     calibration_table,
4     population_stability_index,
5 )
6
7 metrics = binary_model_metrics(default_flag, predicted_pd)
8 calibration = calibration_table(default_flag, predicted_pd, bins=10)
9 psi = population_stability_index(development_score, current_score)

```

B.6 Ratings, CDS, and structural credit

```

1 from creditrisk_lab.ratings import transition_matrix_horizon
2 from creditrisk_lab.hazard import bootstrap_piecewise_hazard
3 from creditrisk_lab.cds import cds_legs
4 from creditrisk_lab.structural import merton_model
5
6 five_year = transition_matrix_horizon(annual_matrix, 5)
7 hazard = bootstrap_piecewise_hazard(maturities, spreads, recovery_rate=0.40)
8 legs = cds_legs(times, discount, survival, recovery_rate=0.40)
9 merton = merton_model(3.0, 0.80, 10.0, 0.05, 1.0)

```

B.7 Counterparty and portfolio models

```

1 from creditrisk_lab.counterparty import exposure_profile, cva, dva
2 from creditrisk_lab.creditrisk_plus import default_count_distribution
3 from creditrisk_lab.creditmetrics import simulate_migration_losses
4 from creditrisk_lab.spread_risk import historical_spread_risk
5
6 profile = exposure_profile(
7     mtm_paths,
8     collateral_paths=collateral_paths,
9     discount_factors=discount,
10 )
11 cva_value = cva(counterparty_marginal_pd, profile["expected_exposure"], discount)
12 dva_value = dva(own_marginal_pd, profile["expected_negative_exposure"], discount)
13 count_pmf = default_count_distribution(3.0, 2.0, 80)
14 migration_losses = simulate_migration_losses(
15     initial_states, P, current_values, value_factors
16 )
17 spread = historical_spread_risk(values, durations, joint_spread_changes)

```

Appendix C

Glossary

- Calibration.** Agreement between predicted probabilities and observed frequencies.
- CCF.** Credit conversion factor applied to an undrawn commitment.
- CDS.** Credit default swap; a contract transferring defined credit-event risk.
- CreditMetrics.** Migration-and-revaluation portfolio credit framework.
- CreditRisk+.** Default-mode portfolio framework based on Poisson mixtures.
- CVA.** Valuation adjustment for counterparty default risk.
- Default.** A precisely defined credit event; definition varies by purpose and framework.
- Discrimination.** Ability to rank higher-risk observations above lower-risk observations.
- EAD.** Exposure expected at the moment of default.
- ECL.** Expected credit loss, commonly referring to IFRS 9 measurement.
- EE.** Expected positive future exposure at a specified date.
- EL.** Long-run or conditional average loss under stated parameters.
- ES.** Average loss in the tail at or beyond a VaR threshold.
- Hazard rate.** Instantaneous conditional default intensity.
- LGD.** Economic loss as a share of EAD, conditional on default.
- PIT.** Point-in-time; sensitive to current and forecast conditions.
- PFE.** A selected quantile of positive future exposure.
- PD.** Probability of default over a specified horizon and definition.
- PSI.** Population stability index, a descriptive distribution-shift measure.
- SICR.** Significant increase in credit risk under IFRS 9.
- TTC.** Through-the-cycle; intended to abstract from some cyclical variation.
- WWR.** Wrong-way risk; adverse dependence between exposure and credit quality.
- UL.** Unexpected loss, defined relative to the mean or a selected tail measure.
- VaR.** Loss quantile at a specified confidence level and horizon.

Appendix D

References and official resources

Bibliography

- [1] Basel Committee on Banking Supervision. *The Basel Framework*. Bank for International Settlements. <https://www.bis.org/baselframework/BaselFramework.pdf>
- [2] European Banking Authority. *Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures*, EBA/GL/2017/16. <https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/model-validation/guidelines-pd-estimation-lgd>
- [3] IFRS Foundation. *IFRS 9 Financial Instruments*. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>
- [4] University of Economics in Katowice. *Measurement of Credit Risk: Course Description Card*, ARIMA 2021/2022. https://www.ue.katowice.pl/fileadmin/user_upload/WFiU/dziekanat/syllabusy_studia_po_angielsku/ARIMA_2021/sem2/M4.2_ARIMA_Measurement_of_Credit_Risk.pdf
- [5] Federal Reserve Bank of St. Louis. *FRED API Documentation*. <https://fred.stlouisfed.org/docs/api/fred/overview.html>
- [6] European Central Bank. *ECB Data Portal API Documentation*. <https://data.ecb.europa.eu/help/api/data>
- [7] World Bank. *Indicators API Documentation*. <https://datahelpdesk.worldbank.org/knowledgebase/articles/889392>
- [8] UCI Machine Learning Repository. *Statlog German Credit Data*. <https://archive.ics.uci.edu/dataset/144/statlog+german+credit+data>
- [9] Hull, J. C. *Risk Management and Financial Institutions*, 5th edition. Wiley, 2018.
- [10] Crouhy, M., Galai, D., and Mark, R. *The Essentials of Risk Management*. McGraw-Hill.

Closing note

Good credit-risk work is not defined by the sophistication of a model in isolation. It is the disciplined alignment of definitions, point-in-time data, estimation, calibration, uncertainty, validation, use, controls, and accountable decision-making. Transparent benchmarks remain valuable even when an institution uses complex models: they make assumptions visible and create a basis for challenge.