AFIN8003 Week 7 - Credit Risk II: Loan Portfolio and Concentration Risk

Banking and Financial Intermediation

Dr. Mingze Gao

2026-04-21

Table of contents

1 Credit Risk II: Loan Portfolio and Concentration Risk

1.1 Why this week matters

! A cautionary tale

In March 2023, **Silicon Valley Bank** collapsed in 36 hours. Every one of its loans was performing. What killed it? **Concentration.** Its depositors were almost all Bay-Area tech startups, and its asset book was loaded with long-duration Treasuries. When one sector caught cold, the whole bank got pneumonia.

- In the last couple of years, three of the four largest bank failures in U.S. history (SVB, Signature, and First Republic) were concentration stories, not one-loan stories.
- Individual due diligence is not enough. **Portfolio risk sum of loan risks.**
- This week: how to measure concentration, diversify it away, or sell it off with derivatives.

1.2 Roadmap

- In Week 6, we measured default risk on *individual* loans.
- FIs hold *portfolios* of loans — so credit risk must be measured and managed in a portfolio context.
- This week we cover:
  1. Simple models of concentration risk (migration analysis, concentration limits)
  2. Modern Portfolio Theory applied to loans (Moody's RiskFrontier)
  3. Regulatory approaches to concentration
  4. Credit derivatives (forwards, options, CDS) — how FIs separate credit risk from the lending relationship

2 Simple models of loan concentration risk

2.1 Simple models

Caution

A large credit risk exposure to a single borrower — or a group of borrowers exposed to the **same risk factor** — poses a potential threat to a bank's safety and soundness.

- Regulations limit such exposure, so individual loans rarely cause bank failures outright.
- The real danger: pools of loans that **look different but share the same risk driver** behave identically in a downturn.
- Think of the 2008 GFC — mortgages, CDOs, structured credit, MBS, and derivatives were marketed as *different products* but were all bets on U.S. house prices. When the one underlying factor turned, everything turned.

The hidden-correlation trap

Products with different names, sold by different business units, can still share the same underlying risk. A bank's "diversified" book may be one bet in disguise.

Two simple models widely used to measure concentration risk:

1. Migration analysis
2. Concentration limits

2.2 Simple model: migration analysis

- Track credit ratings of certain types of loans or certain sectors, either externally from credit rating agencies or internally.
- If actual rating deteriorates faster than historical experience, limit lending to that loan class or sector.
- Historical credit migration measured through **loan migration matrix** (or transition matrix).

Table 1: Example loan migration matrix

	AAA-A	BBB-B	CCC-C	Default
AAA-A	0.85	0.10	0.04	0.01
BBB-B	0.12	0.83	0.03	0.02
CCC-C	0.03	0.13	0.80	0.04

Table 1, for example, shows the *transition probabilities* of loans that began the year with a certain credit rating being upgraded/downgraded to a certain rating, or default.

- The probability of AAA-rated loan at the start of a year being downgraded to BBB to B by the year's end is 0.10.
- The probability of AAA-rated loan at the start of a year being downgraded to CCC to C by the year's end is 0.04.
- The probability of AAA-rated loan at the start of a year defaults by the year's end is 0.01.

2.3 Simple model: migration analysis (cont'd)

In practice, FIs use migration matrices with many more rating classes (S&P uses 20+).

Migration analysis is also applied to credit card and consumer loan portfolios.

⚠ Where migration analysis can mislead you

- **Historical data only** — "last war" risk. Matrices built on 2010s data missed COVID-era dislocations.
- **Rating agencies lag** — downgrades typically arrive *after* trouble is visible. By the time Moody's cuts, the market has already repriced.
- **Point-in-time vs through-the-cycle** — the same rating can mean different things across agencies and time.

ℹ Discussion

Lehman Brothers was rated A by S&P five days before its September 2008 bankruptcy. What does that tell you about relying on migration analysis alone?

2.4 Simple model: concentration limits

- Caps the maximum loan size to an individual borrower, sector, or geographic area.
- Used to reduce exposure to some industries and increase it in others.
- Aggregate limits applied to industries whose performance is highly correlated.
- **Regulatory floors:**
  - **US (OCC):** loans to a single borrower capped at 15% of bank's capital (25% if fully secured).
  - **Australia (APRA APS 221):** large exposures to a single counterparty (or group) generally capped at 25% of Tier 1 capital; tighter 15% limit applies between domestic G-SIBs.

$$\text{Concentration limit} = \text{Maximum loss as a percentage of capital} \times \frac{1}{\text{Loss rate}}$$

Example

If an FI's manager is unwilling to permit losses exceeding 15% of the FI's capital, with an estimated loss rate in a particular industry of 40 per cent, then the manager should set a concentration limit on the exposure to that industry as $15\% \times \frac{1}{0.4} = 37.5\%$.

2.5 Simple model: concentration limits (cont'd)

Let's look at a real Australian bank. BOQ's FY2023 credit exposure — notice anything?

Table 2: Proportionate credit exposures of lending activities of BOQ FY23

Sector	$m	% of Total Exposure
Residential mortgages	62,738	77.8
Property and construction	6,887	8.5
Healthcare	2,763	3.4
Professional services	2,431	3.0
Agriculture	1,232	1.5
Transportation	606	0.8
Manufacturing and mining	682	0.8
Hospitality and accommodation	841	1.0
Other	2,453	3.0
Total	80,633	100.0

Discussion

77.8% in residential mortgages. Is BOQ "diversified" across 9 sectors — or is it essentially one big bet on Australian house prices? Compare this to the Big Four. This pattern is typical of Australian banks and is one reason APRA stress-tests them on a property-market downturn.

3 Loan portfolio diversification and Modern Portfolio Theory (MPT)

3.1 Loan portfolio diversification and MPT

MPT can be used to measure and control an FI's aggregate credit risk exposure.

Any model that seeks to estimate an efficient frontier for loans needs to determine and measure three things:

- The expected return on individual loans
- The risk of individual loans
- The correlation of default risks between loans

Expected return R_p of a portfolio of N assets:

$$R_p = \sum_{i=1}^{N} X_i R_i$$

where

- R_i is the expected return on the ith asset
- X_i is the proportion of the asset portfolio invested in the ith asset (the desired concentration amount)

3.2 Loan portfolio diversification and MPT (cont'd)

Variance of returns (or risk) of the portfolio σ_p^2 can be calculated as

$$\begin{aligned}\sigma_i^2 &= \sum_{i=1}^{N} X_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{j=1}^{N} X_i X_j \sigma_{ij} \\ &= \sum_{i=1}^{N} X_i^2 \sigma_i^2 + \sum_{i=1}^{N} \sum_{j=1}^{N} X_i X_j \rho_{ij} \sigma_i \sigma_j\end{aligned}$$

where

- ρ_{ij} is the correlation between the returns on the ith asset jth asset
- σ_i^2 is the variance of returns on the ith asset
- σ_{ij}^2 is the covariance of returns between the ith asset and jth asset

3.3 Loan portfolio diversification and MPT (cont'd)

The fundamental lesson of MPT is that by taking advantage of its size, an FI can diversify considerable amounts of credit risk as long as the returns on different assets are imperfectly correlated with respect to their default risk adjusted returns.

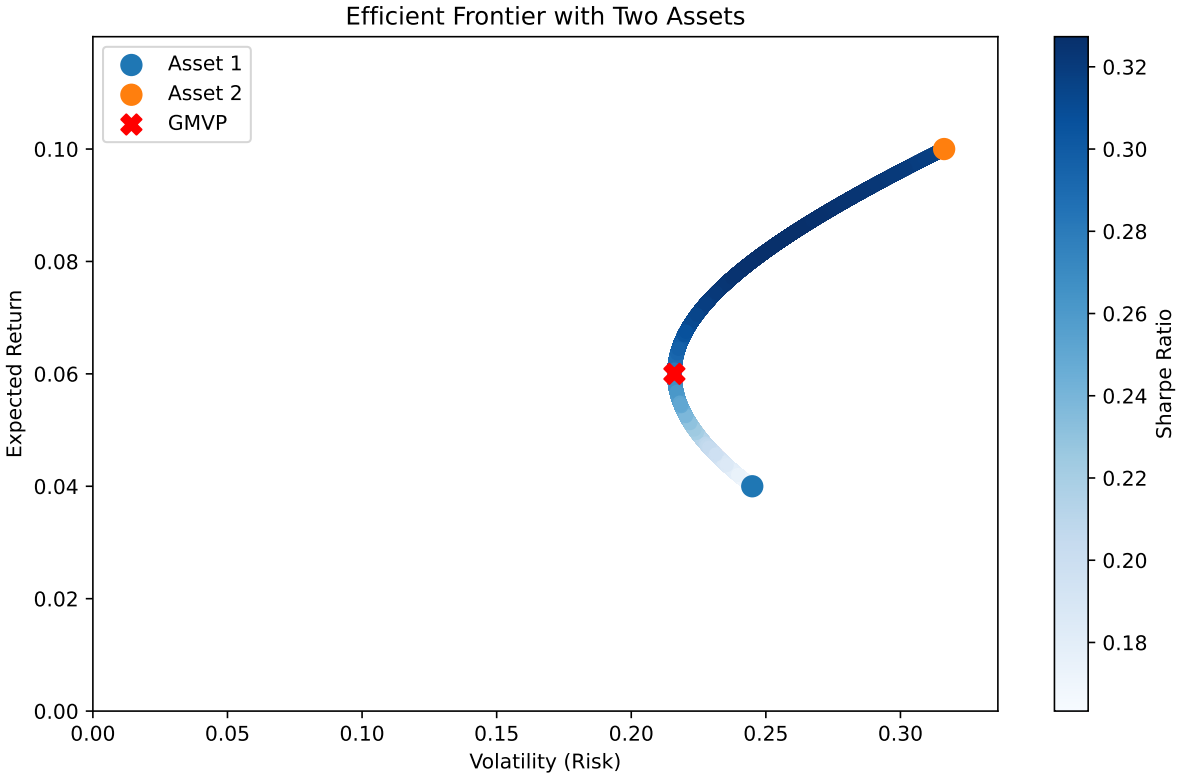

Figure 1: MPT and efficient frontier

Minimum risk portfolio (Global Minimum Variance Portfolio, GMVP)

- Combination of assets that reduces the variance of portfolio returns on the lowest feasible level.

3.4 Moody's RiskFrontier: the big picture

The problem: MPT needs **three inputs** — but for a loan portfolio, none of them are directly observable.

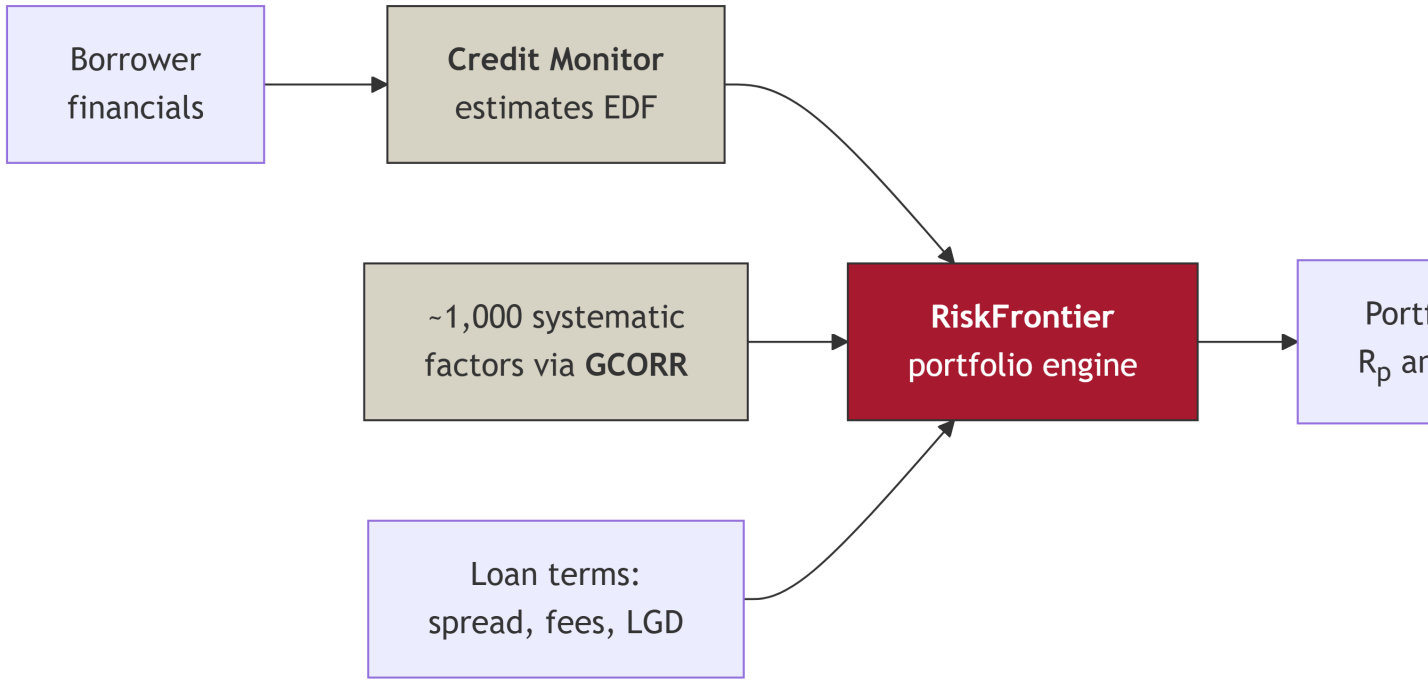

Two Moody's models in sequence:

1. **Credit Monitor** → produces EDF (Expected Default Frequency) for each borrower
2. **RiskFrontier** → plugs EDFs, LGDs, and correlations into MPT to get portfolio R and

> Why not just use historical correlations?
>
> Most loans never trade. There is no daily price series to correlate. RiskFrontier works around this by computing correlations from **shared exposure to ~1,000 systematic factors** (country, industry, region) rather than from return history.

3.5 RiskFrontier: the three MPT inputs

The whole model is just **three numbers per loan**, fed into standard MPT formulas.

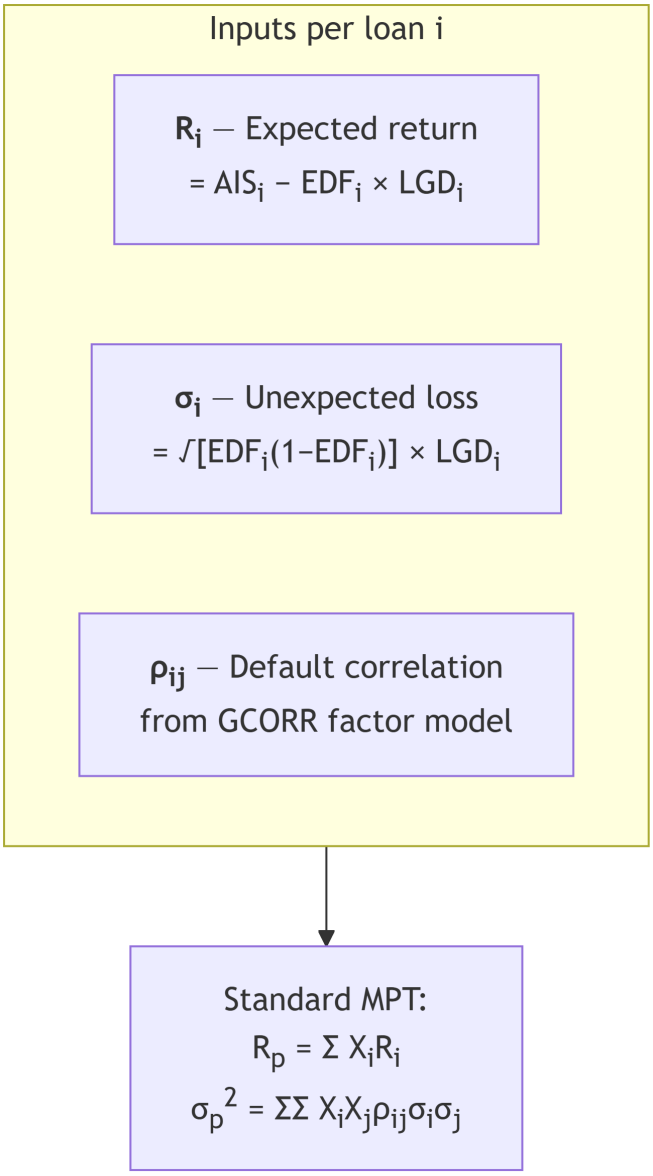

Symbol	Meaning	Source
AIS	All-in-drawn spread (loan rate − cost of funds + fees)	Loan contract
EDF	Prob. of default in the next year	Credit Monitor
LGD	Fraction lost if default occurs	Basel floors or bank estimate[1]
	Default correlation	GCORR factor model

3.6 RiskFrontier: expected return and risk

Expected return — earn the spread, lose the expected loss:[2]

[1] Basel Foundation IRB (CRE32): senior unsecured = **45%**, secured by other physical collateral = **40%**, secured by receivables or CRE/RRE = **35%**, subordinated = **75%**.

[2] In Saunders et al. (2023), this equation has a typo: $E(L_i)$ is written as $R(L_i)$.

$$\underbrace{R_i}_{\text{net return}} = \underbrace{AIS_i}_{\text{spread + fees}} - \underbrace{EDF_i \times LGD_i}_{E(L_i),\ \text{expected loss}}$$

Unexpected loss — default is binomial, so has a closed form:[3]

$$\sigma_i = UL_i = \underbrace{\sqrt{EDF_i(1-EDF_i)}}_{\sigma \text{ of a } 0/1 \text{ default event}} \times \underbrace{LGD_i}_{\text{loss if default}}$$

Intuition

- **Expected loss** (EDF × LGD) is already *priced into* the spread — it's the cost of doing business.
- **Unexpected loss** ( ) is what you hold capital against — the surprise.

3.7 RiskFrontier: where correlations come from (GCORR)

Default correlations between two loans cannot be directly observed. GCORR computes them via a **factor model**: two borrowers are correlated to the extent they share exposure to the *same underlying risk factors*.

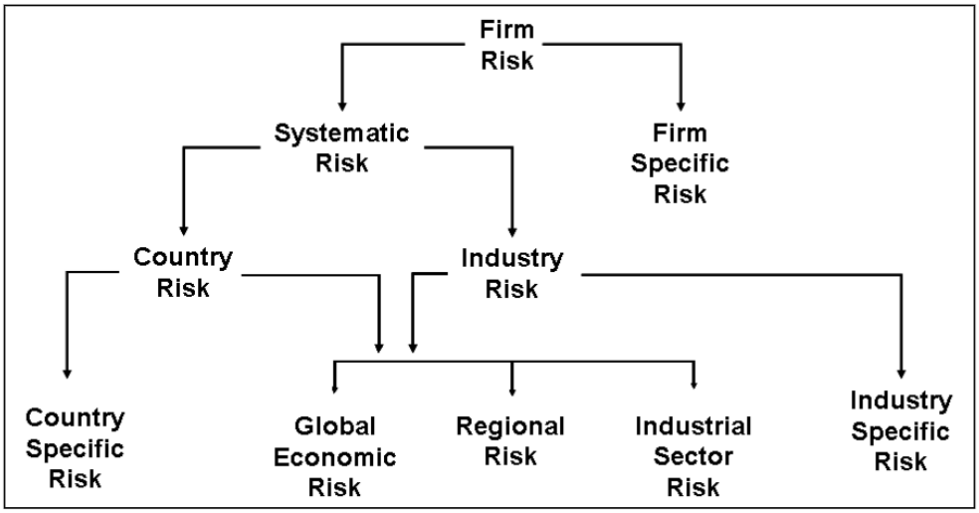

Figure 2: Moody's GCorr Corporate factor structure

Read Figure 2 this way

Each borrower's asset return = **global economy** + **region/country** + **industry** + **firm-specific noise**. Two borrowers are correlated only through the **shared** branches of the tree. A Sydney miner and a Perth miner share the "Australia + Materials" branches — high . A Sydney miner and a Berlin software firm share almost nothing — near zero.

[3] In Saunders et al. (2023), this equation is incorrectly written.

3.8 Moody's Analytics RiskFrontier Model (example)

Suppose that an FI holds two loans with the following characteristics. Assume that the correlation $\rho_{12} = -0.25$, what are the return and risk of the portfolio?

Loan i	X_i	Spread between loan rate and FI's cost of funds	Fees	LGD	EDF
1	0.6	5%	2%	25%	3%
2	0.4	4.5%	1.5%	20%	2%

The return and risk on loan 1 are:

$$R_1 = (0.05 + 0.02) - (0.03 \times 0.25) = 0.0625$$
$$\sigma_1 = \sqrt{0.03 \times 0.97} \times 0.25 = 0.04265$$

The return and risk on loan 2 are:

$$R_2 = (0.045 + 0.015) - (0.02 \times 0.2) = 0.056$$
$$\sigma_2 = \sqrt{0.02 \times 0.98} \times 0.2 = 0.028$$

The return and risk of the portfolio are then:

$$R_p = 0.6 \times 0.0625 + 0.4 \times 0.056 = 0.0599 \text{ or } 5.99\%$$
$$\sigma_p^2 = (0.6)^2(0.04265)^2 + (0.4)^2(0.028)^2 + 2(0.6)(0.4)(-0.25)(0.04265)(0.028) = 0.0006369$$
$$\sigma_p = \sqrt{0.0006369} = 0.0252 = 2.52\%$$

3.9 Regulatory models

- **Fed's 1994 Ruling on Credit Concentration Risk**
  - Subjective approach based on examiner discretion.
  - Rejected technical models — data and methods were too undeveloped at the time.
- **2006 regulatory tightening**
  - **BIS:** 10 principles on credit risk assessment and supervisory evaluation.
  - **OCC:** guidance on sound risk management for commercial real estate (CRE) lending.
- **OCC/Fed Joint Study (April 2013) — the numbers are striking:**

Bank profile during GFC	Failure rate
Construction loans > 100% of capital	**13%**
Exceeded BOTH construction AND total CRE criteria	**23%**
Did not exceed either criterion	0.5%

> **Takeaway**
>
> Banks that ignored the 2006 CRE concentration guidance were ~**45× more likely to fail** during the GFC than compliant banks. Concentration guidance is advisory — but the statistical case for it is overwhelming.

4 Use of derivatives to manage credit risk

4.1 Credit derivatives

- Diversification is the first line of defence. Derivatives are the second.
- The key insight: **credit derivatives separate the credit risk from the lending relationship**. A bank can keep the client, service the loan, and still offload the default risk to someone else.
- Three main instruments: credit **forwards**, credit **options**, credit **swaps** (including CDS).

> Why this innovation matters
>
> Before credit derivatives (pre-1990s), the only way a bank could reduce credit exposure to a big client was to **refuse the loan** or **sell it** — both of which damage the relationship. Credit derivatives let the bank say "yes" and still cap its downside.

4.2 Credit forward contracts and credit risk hedging

- A credit forward is a forward agreement that hedges against an increase in default risk of a firm (borrower).[4]
- Specifies a credit spread on a benchmark bond issued by the borrower.
- Used to **hedge against credit deterioration** (spread widening)
- Spread is measured vs. a risk-free Treasury
- Example: BBB bond trades at 2% spread over Treasury

4.3 Credit forward contracts and credit risk hedging (cont'd)

Market outcome	Long position (hedger)	Short position
Spread widens → credit quality worsens	Gains (receives payment)	Loses (makes payment)
Spread tightens → credit quality improves	Loses (makes payment)	Gains (receives payment)

- If $\phi_T > \phi_F$: Long receives $(\phi_T - \phi_F) \times MD \times A$
- If $\phi_T < \phi_F$: Long pays $(\phi_F - \phi_T) \times MD \times A$

where

- ϕ_F is the credit spread on which the credit forward contract is written
- ϕ_T is the actual credit spread on the bond when the credit forward matures
- MD is the modified duration on the benchmark bond
- A is the principal amount of the forward agreement
- **Long position** protects against borrower credit quality getting worse
- **Short position** benefits if borrower credit improves
- Acts like a **put-style hedge** for lenders

[4] See A Note on Credit Spread Forwards for some discussion on the potential confusion due to different definitions.

4.4 Credit options

Use of options to hedge credit risk is a relatively new phenomenon. In September 2021, commercial bank holdings of credit options totaled $233 billion, which represented 6.05% of all credit derivatives outstanding.

- A **credit spread call option** is a call option whose payoff increases as the (default) risk premium or yield spread on a specified benchmark bond of the borrower increases above some exercise spread.
- A **digital default option** is an option that pays a stated amount in the event of a loan default (the extreme case of increased credit risk).

4.5 Credit (default) swaps (CDS)

The most important — and most controversial — credit derivative.

Explosive growth, then regulatory pushback:

Date	US bank credit-derivative notional	Note
2000	$426 billion	Nascent market
March 2008	$16.44 trillion	Pre-GFC peak
2009	$13.44 trillion	Post-Lehman retrenchment
September 2011	$15.66 trillion ($15.31tn CDS)	CDS dominance
September 2021	$3.9 trillion ($3.4tn CDS)	Post-Dodd-Frank, central clearing

! What caused the collapse?

Dodd-Frank (2010) pushed most standardized CDS onto **central clearinghouses** (CCPs), with initial-margin and variation-margin requirements that made speculative CDS much more expensive. Bilateral inter-dealer books shrank dramatically.

Why CDS exist:

1. Credit risk is still the #1 cause of FI failure — more than interest-rate or FX risk.
2. CDS let FIs keep **long-term customer relationships** while offloading default risk.

4.6 Basics of CDS

- **CDS Payments**: The buyer makes periodic payments to the seller (quarterly, semi-annually, or annually) until the end of the swap or a credit event (e.g., default) occurs.
- **Settlement**: Upon default, settlement involves either physical delivery of bonds (loans) or a cash payment.
  - The protection buyer receives a payment upon the occurrence of a credit event trigger, but the swap "expires worthless" if no trigger occurs.
- No requirement that the CDS buyer actually owns the underlying reference securities.
- The CDS buyer hedges its exposure to default risk, but there is still counterparty credit risk in the event that the seller fails to perform their obligations under the terms of the contract.[5]

We examine two types of credit swaps:

1. total return swap
2. pure credit swap

[5] In September 2008, **AIG** had written roughly **$440 billion of CDS protection** — much of it on subprime MBS — with insufficient collateral posted. When spreads blew out, AIG faced margin calls it could not meet. The US Treasury/Fed extended $182 billion in rescue financing. The lesson: CDS protection is only as good as the counterparty. This episode is why **central clearing** is now mandatory for standardized CDS.

4.7 Credit swaps: total return swaps

A **total return swap** involves *swapping* an obligation to pay interest at a specified fixed or floating rate *for* payments representing the total return on a loan or a bond (interest and principal value changes) of a specified amount.

The figure below illustrates a total return swap.[6]

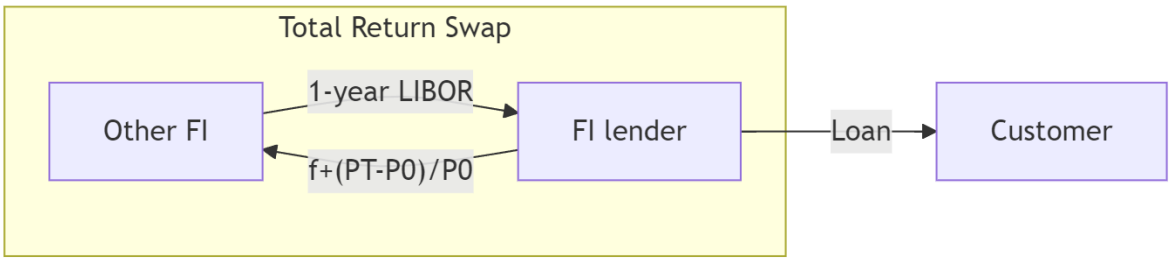

The FI lender pays a fixed annual rate f plus changes in the market value of the loan and receives a variable rate payment (historically **1-year LIBOR**; post-2023 **SOFR** in USD, **AONIA/BBSW** in AUD since LIBOR's cessation in June 2023).

- Interest rate risk remains — the TRS transfers credit *and* market risk of the asset, but the fixed vs floating mismatch is still an interest-rate exposure.

4.8 Credit swaps: total return swaps (cont'd)

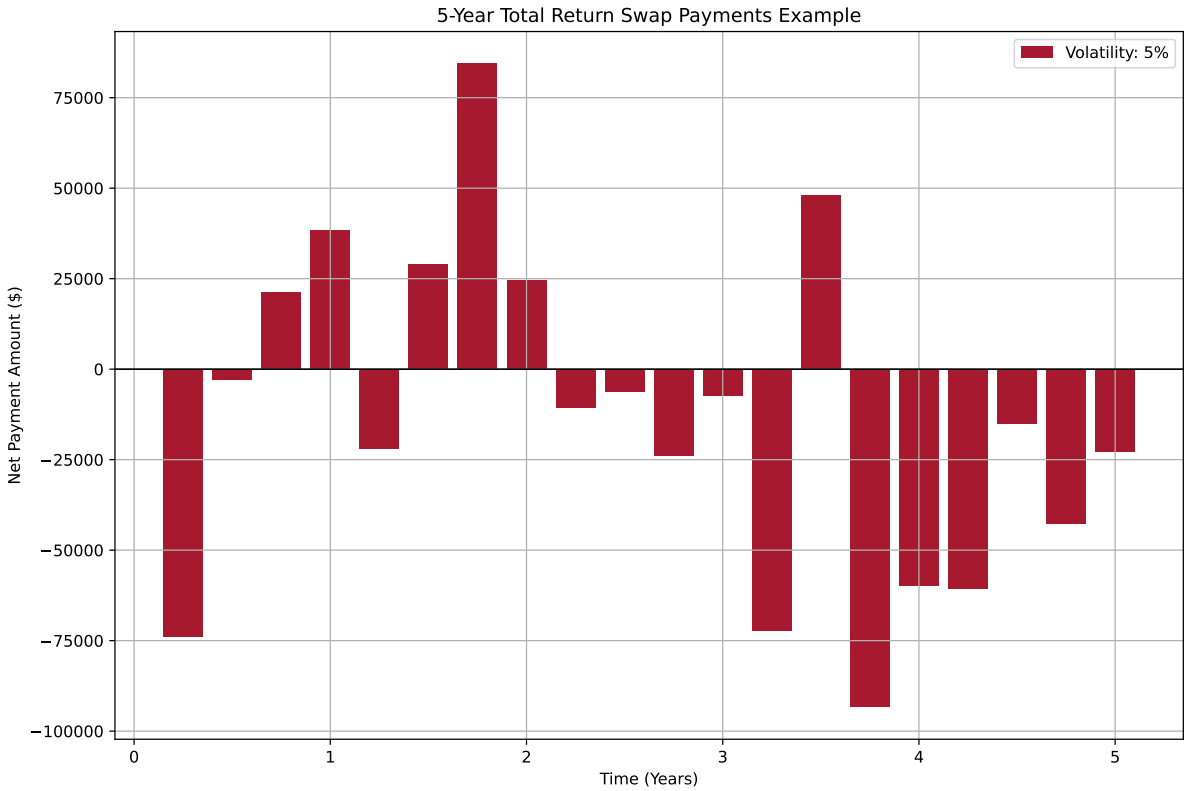

Figure 3: Total Return Swap Payments Example

[6] In the textbook Saunders et al. (2023), Figure 25-4 incorrectly has $f - \frac{P_t - P_0}{P_0}$.

4.9 Credit swaps: pure credit swaps

Interest-rate sensitive element stripped out leaving only the credit risk.

Similar to buying an insurance:

- If default, the counterparty makes the default payment
- If not default, the FI receives nothing from the counterparty.

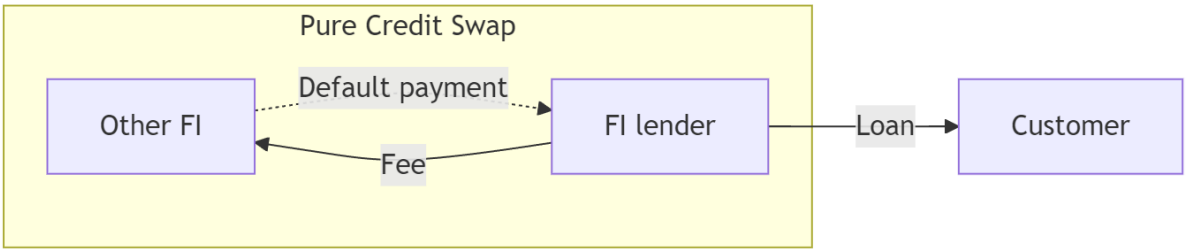

4.10 Credit swaps: pure credit swaps (cont'd)

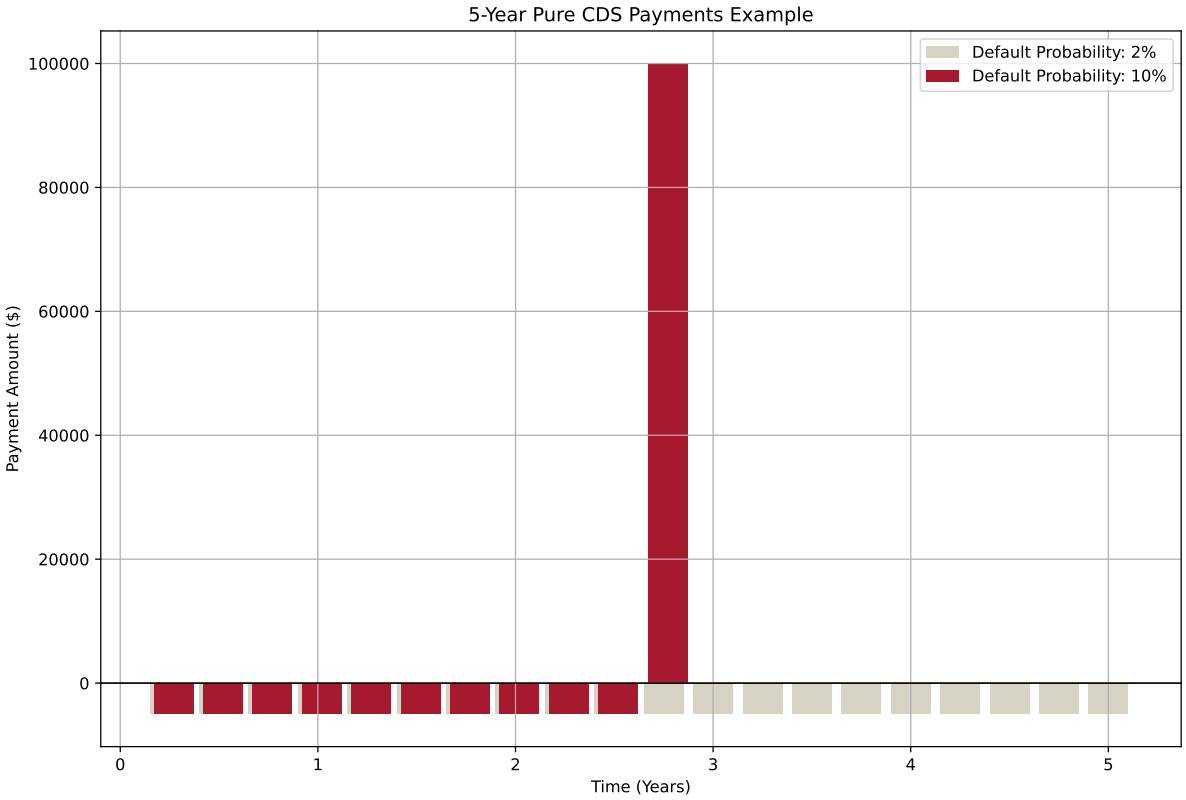

Figure 4: Pure CDS Payments Example

5 Finally...

5.1 Key takeaways

1. **Concentration, not individual default, is the usual killer.** Loans that "look different" often share one risk factor.

2. **Migration analysis** is useful but backward-looking — rating agencies lag reality.
3. **MPT applies to loans**, with adaptations (RiskFrontier): expected loss = EDF × LGD, risk = √(EDF(1-EDF)) × LGD, correlations come from factor models (GCORR), not history.
4. **Derivatives separate credit risk from the lending relationship.** TRS transfers total return; pure CDS transfers only default risk.
5. **Counterparty risk is real** — AIG 2008 is the cautionary tale. Central clearing is the post-GFC response.

5.2 Suggested readings

- Moody's Analytics — An Overview of Modelling Credit Portfolios
- BIS CRE32 — IRB approach: risk components (for the LGD floors)
- OCC Quarterly Report on Bank Trading and Derivatives Activities (for current credit-derivative notionals)

References

Saunders, Anthony, Marcia Millon Cornett, and Otgo Erhemjamts. 2023. *Financial Institutions Management ISE*. 11th ed. McGraw Hill.