Lecture notes on risk management, public policy, and the financial system

Credit risk models

Allan M. Malz

Columbia University

Overview of credit risk analytics

Credit risk metrics and models
Intensity models and default time analytics

Single-obligor credit risk models

Key metrics of credit risk

Probability of default π_t defined over a time horizon t, e.g. one year

Exposure at default: amount the lender can lose in default

- For a loan or bond, par value plus accrued interest
- For *OTC derivatives*, also driven by market value
  - Net present value (NPV) $\lesseqqgtr 0$ ($\rightarrow$counterparty risk)
  - But exposure at default ≥ 0

Recovery: creditor generally loses fraction of exposure $R < 100$ percent

Loss given default (LGD) equals exposure minus recovery (a fraction $1 - R$)

Expected loss (EL) equals default probability $\times$ LGD or fraction $\pi_t \times (1 - R)$

- Credit risk management focuses on **unexpected loss**

Credit Value-at-Risk related to a quantile of the credit return distribution

- Differs from market risk in *excluding* EL
- Credit VaR at confidence level of α defined as:

$$\alpha\text{-quantile of credit loss distribution} - \text{EL}$$

Estimating default probabilities

Risk-neutral default probabilities based on market prices, esp. credit spreads

- Data sources include credit-risky securities and CDS
- Risk-neutral default probabilities may incorporate risk premiums
- Used primarily for market-consistent pricing

Physical default probabilities based on fundamental analysis

- Based on historical default frequencies, scenario analysis, or credit model
- Associated with credit ratings
- Used primarily for risk measurement

Types of credit models

Differ on inputs, on what is to be derived, and on assumptions:

Structural models or **fundamental models** model default, derive measures of credit risk from fundamental data

- Firm's balance sheet: volumes of assets and debt
- Standard is the **Merton default model**

Reduced-form models or **intensity models** take estimates of default probability or LGD as inputs

- Often used to simulate default times as one step in portfolio credit risk modeling
- Often risk-neutral
- Common example: **copula models**

Factor models: company, industry, economy-wide fundamentals, but highly schematized, lends itself to portfolio risk modeling.

Some models fall into several of these categories

What risks are we modeling?

Credit risk: models are said to operate in

- **Migration mode** taking into account credit migration a well as default, or
- **Default mode** taking into account default only

Spread risk: credit-risk related market risk

Rating migration rates, 1920–2016

From/To:	Aaa	Aa	A	Baa	Ba	B	Caa	Ca–C	WR	Default
Aaa	86.7	7.8	0.8	0.2	0.0	0.0	0.0	0.0	4.4	0.0
Aa	1.1	84.2	7.6	0.7	0.2	0.0	0.0	0.0	6.1	0.1
A	0.1	2.7	85.0	5.6	0.6	0.1	0.0	0.0	5.7	0.1
Baa	0.0	0.2	4.3	82.7	4.6	0.7	0.1	0.0	7.0	0.3
Ba	0.0	0.1	0.5	6.1	73.9	6.9	0.7	0.1	10.6	1.2
B	0.0	0.0	0.2	0.6	5.6	71.7	6.2	0.5	11.9	3.3
Caa	0.0	0.0	0.0	0.1	0.6	6.9	67.3	2.9	13.7	8.4
Ca–C	0.0	0.0	0.1	0.0	0.6	3.0	8.0	48.4	18.7	21.1

Average one-year letter rating migration rates, 1920-2016, percent. Each row shows the probability of starting the year with the rating in row heading and ending with the rating in the column heading. “WR” denotes withdrawn rating. *Source*: Moody’s Investor Service.

Modeling default time

- Occurrence of default event for single company over discrete time horizon t can be modeled as Bernoulli distribution
  - Default *occurrence* the random variable
- Alternatively: **default intensity models**, model the time—specific instant τ—at which default occurs
  - Default *time* the random variable
  - Look out over horizon from now (time 0) to time t
  - Default probability over $[0, t)$: $\mathbf{P}\left[0 \le \theta < t\right]$, firm defaults by time t
- Example of **jump** or **Poisson process** with exactly one jump possible
- Default probability increases as horizon grows longer
- Every firm defaults eventually: $\lim_{t \to \infty} \mathbf{P}\left[0 \le \tau < t\right] = 1$
- **Survival probability** is $1 - \mathbf{P}\left[0 \le \tau < t\right]$:
  - Obligor remains solvent until at least time t

Default time distributions

- Default probabilities can be expressed through **cumulative default time distribution function**
- Simple form: $\mathbf{P}[0 \leq \tau < t] = 1 - e^{-\lambda t}$
    - Survival probability is $e^{-\lambda t}$
    - If t expressed in years, 1-year default probability is $1 - e^{-\lambda}$
- Corresponding p.d.f. is derivative w.r.t. $t : \lambda e^{-\lambda t}$
    - For tiny time interval dt

$$\mathbf{P}[t \leq \tau < t + dt] \cong \lambda e^{-\lambda t}$$

Hazard rates

- Default can only occur once $\Rightarrow$

$$\mathbf{P}\left[t \leq \tau < t + dt\right] = \mathbf{P}\left[t \leq \tau < t + dt \cup t < \tau\right]$$

- Conditional probability of default, given it has not occurred before t:

$$\begin{aligned} \mathbf{P}\left[t \leq \tau < t + dt \mid t < \tau\right] &= \frac{\mathbf{P}\left[t \leq \tau < t + dt \cup t < \tau\right]}{\mathbf{P}\left[t < \tau\right]} \\ &= \frac{\lambda e^{-\lambda t}}{e^{-\lambda t}} \\ &= \lambda \end{aligned}$$

- λ called **hazard rate** or **default intensity**
  - Viewed from time 0, probability of default over dt is λdt
- λ can be modeled as a constant or as changing over time
- In insurance, **force of mortality,** probability of death of a population member over next short time interval

Conditional default probability

- **Conditional default probability**: probability of default over a future time horizon, *given* no default before then
- With constant hazard rate:
  - Unconditional one-year default probability lower for more remote years
  - But time to default *memoryless*: *if* no default occurs next year, probability of default over subsequent year is same as next year
- λ: *instantaneous* conditional default probability
  - Probability of default over next instant, given no prior default

Default probability analytics: example

Hazard rate	λ	0.15
1-yr. default probability	$1 - e^{-\lambda}$	0.1393
2-yr. default probability	$1 - e^{-2\lambda}$	0.2592
1-yr. survival probability	$e^{-\lambda}$	0.8607
1-yr. conditional default probability	$1 - e^{-\lambda}$	0.1393

Default time distribution

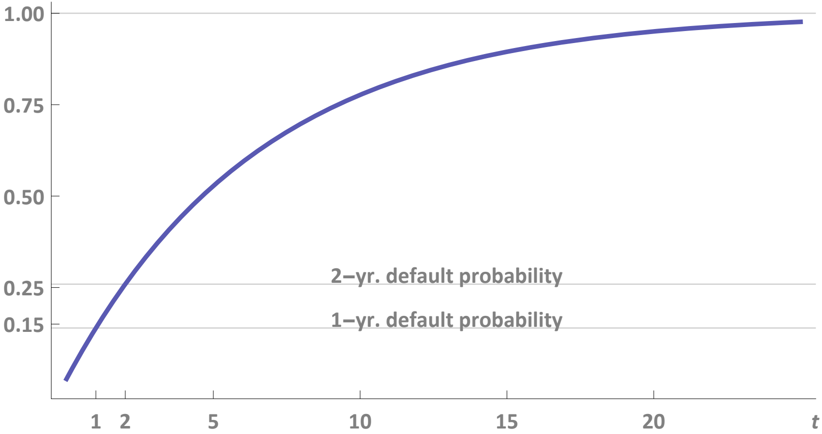

Cumulative default time distribution function π_t, constant hazard rate $\lambda = 0.15$, t measured in years, π_t and λ at an annual rate.

Merton model: overview

- Widely-used structural model based on fluctuations in debt-issuing firm's asset value
- Default occurs when asset value falls below **default threshold**, at which
  - Equity value extremely low or zero
  - Asset value close to par value of debt (plus accrued interest)
- Simplest version:
  - Default occurs when equity value hits zero
  - Default threshold equals par value of debt (plus accrued interest)

Equity and debt as options

- Assets assumed to display return volatility, so can apply option-pricing theory
- Equity can be viewed as a long call on the firm's assets, with a strike price equal to the par value of the debt
- Debt can be viewed as a portfolio:
  - A riskless bond with the same par value as the debt
  - Plus an implicit short put on the firm's assets, with a strike price equal to the par value of the debt
- If the lender bought back the short put, it would immunize itself against credit risk
  - $\Rightarrow$The value of the implicit short a measure of credit risk

Merton default model

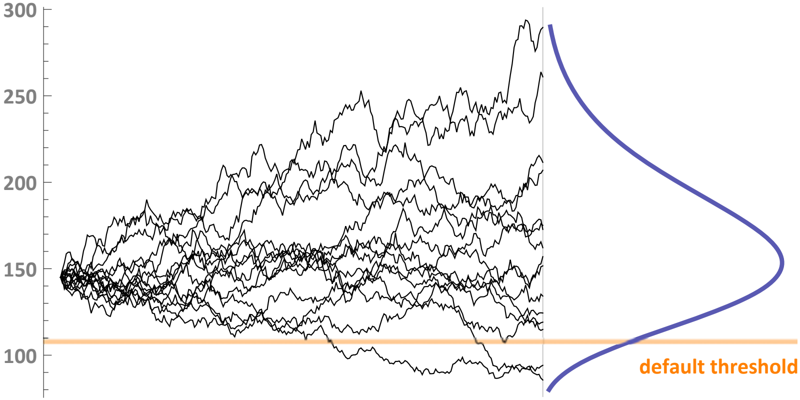

Left: 15 daily-frequency sample paths of the geometric Brownian motion process of the firm's assets with a drift of 15 percent and an annual volatility of 25 percent, starting from a current value of 145. Right: probability density of the firm's asset value on the maturity date, one year hence, of the debt. The grid line represents the debt's par value (100) plus accrued interest at 8 percent.

Applying the Merton default model

- Immediate consequence: higher volatility (risk) benefits equity at expense of debtholders
- Model can be used to compute credit spread, expected recovery rate
- Two ways to frame model, depending on how mean of underlying return process interpreted

  Risk-neutral default probability: expected value equal to firm's dividend rate

  Physical default probability: expected value equal to asset rate of return
- Model timing of default, compute default probability
- Moody's KMV (and other practitioner applications):
  - Equity vol plus leverage→asset vol
  - Plus book value of liabilities→default threshold
  - Historical data+secret sauce to map into default frequency

Structure of single-factor model

- Basic similarity to Merton model
  - Default occurs when asset value falls below default threshold
- Asset returns depend on two random variables:

  Market risk factor m affects all firms, but not in equal measure
  - Expresses influence of general business conditions, state of economy on default risk
  - Latent factor: not directly observed, but influences results indirectly via model parameters

  Idiosyncratic risk factor ϵ affects just one firm
  - Expresses influence of individual firm's situation on default risk
- Fixed time horizon, e.g. one year
- Returns and shocks are measured as deviations from expectations or from a "neutral" state
- Most often used to model portfolio credit risk rather than single obligor

Parameters of single-factor model

Default probability π or, equivalently, default threshold k

- Combination of adverse market and idiosyncratic shocks sufficient to push borrower into default

Correlation β of asset return to market risk factor m

- High correlation implies strong influence of general business conditions on firm's default risk
- Correlations of individual firms' asset returns key driver of extent to which defaults of different firms coincide
- $\rightarrow$Portfolio credit models and **default correlation**

Single-factor model: asset return behavior

- Merton model framework: default threshold is hit when firm's asset return r large and negative
- Asset return a function of market and idiosyncratic risk factors m, ϵ:

$$r = \beta m + \sqrt{1 - \beta^2}\,\epsilon$$

- β: correlation between firm's asset return and market factor m
- m and ϵ uncorrelated standard normal variates:

$$\begin{aligned} m &\sim \mathcal{N}(0,1) \\ \epsilon &\sim \mathcal{N}(0,1) \\ \mathrm{Cov}[m,\epsilon] &= 0 \end{aligned}$$

- Therefore r is a standard normal variate, expressed in "volatility units": $r \sim \mathcal{N}(0,1)$, with
  - r and m expressed as deviations from neutral state of business cycle

$$\begin{aligned} \mathbf{E}\,[r] &= 0 \\ \mathrm{Var}[r] &= \beta^2 + 1 - \beta^2 = 1 \end{aligned}$$

Asset and market returns in the single-factor model

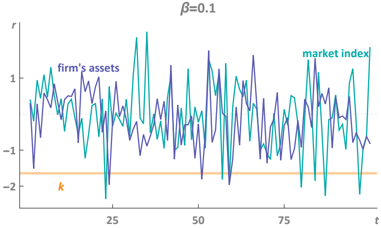

Each panel shows a sequence of 100 simulations from the single-factor model. Cyan plot: returns on the market index m. Purple plot: associated returns $r = \beta m + \sqrt{1-\beta^2}\epsilon$ on firm's assets with the specified β to the market. Plots are generated by simulating m and ϵ as a pair of uncorrelated $\mathcal{N}(0,1)$ variates, using the same random seed for both panels.

Single-factor model: default probability

- Default probability an assigned parameter
  - Rather than an *output*, as in the Merton model, the default probability is an *input* in the single-factor model
- Expressed via default threshold k or initial **distance-to-default**
  - Default threshold a negative number, distance-to-default initially equal to $-k$
  - Default if r negative and large enough to wipe out equity:

$$\beta m + \sqrt{1-\beta^2}\,\epsilon \leq k$$

- Or, equivalently, **distance-to-default** $-k = |k|$
- Finding the initial default threshold: set k to match stipulated default probability π via

$$\pi = \mathbf{P}\left[r \leq k\right] \Leftrightarrow k = \Phi^{-1}(\pi),$$

  where $\Phi(\cdot)$ is the standard normal CDF
- **Example**:

	$\pi = 0.01$	$\pi = 0.10$
Distance-to-default ($-k$)	2.33	1.28

Single-factor model: default probability

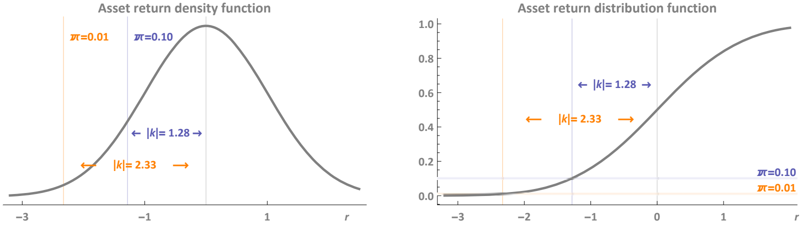

Vertical grid lines mark the default threshold corresponding to default probabilities of 0.01 and 0.10.

Single-factor model and CAPM

- Single-factor model vs. CAPM beta
  - Since $\text{Var}[r] = 1$, β analogous to the *correlation* of market and firm, rather than CAPM beta
  - Relationship of asset rather than equity values to market factor
- Systematic and idiosyncratic risk: fraction of asset return variance explained by variances of
  - Market risk factor: β^2
  - Idiosyncratic risk factors: $1 - \beta^2$
  - Example:

	$\beta = 0.40$	$\beta = 0.90$
Market factor β^2	0.16	0.81
Idiosyncratic factor $1 - \beta^2$	0.84	0.19

Market factor and conditional independence

- Suppose we know the "state of the economy,", i.e. the particular realization $\bar{m}$ of m
- Obligor i asset return r_i now has only one random driver: idiosyncratic factor ϵ_i

$$r_i = \beta_i \bar{m} + \sqrt{1 - \beta_i^2}\epsilon_i, \qquad i = 1, 2, \ldots$$

- Distance-to-default—the default-triggering return—becomes $-k_i + \beta_i \bar{m}$
  - Default threshold k_i itself hasn't changed
- ϵ_i independent⇒conditional returns of two different obligors

$$\sqrt{1 - \beta_i^2}\epsilon_i, \sqrt{1 - \beta_j^2}\epsilon_j, i \neq j$$

  are independent
- ⇒**Conditional independence**: defaults of two firms independent
  - Conditioning is on realization of market risk factor
  - $\not\leftrightarrow$ discussion of hazard rates, in which "conditional" default probability refers to non-default in prior period

Conditional default distribution of a single obligor

- Conditional on $m = \bar{m}$:
  - **Mean** of the return distribution **changes**: $0 \rightarrow \beta_i \bar{m}$
  - **Variance** of the return distribution **reduced**: $1 \rightarrow 1 - \beta_i^2$
    - Because we have eliminated market factor as source of variation
    - And **standard deviation** from $1 \rightarrow \sqrt{1 - \beta_i^2}$
  - **Distance-to-default changes**: $-k_i \rightarrow -(k_i - \beta_i \bar{m})$
    - In standard units: $-k_i \rightarrow -\frac{k_i - \beta_i \bar{m}}{\sqrt{1-\beta_i^2}}$
  - **Default probability changes**: $\pi_i = \Phi(k_i) \rightarrow \Phi\left(\frac{k_i - \beta_i \bar{m}}{\sqrt{1-\beta_i^2}}\right)$
    - With $\Phi(x)$ the CDF of a standard normal variate x
- $\Rightarrow$**Conditional default probability distribution function**:

$$p_i(m) = \mathbf{P}\left[r_i \leq k_i | m\right] = \Phi\left(\frac{k_i - \beta_i m}{\sqrt{1 - \beta_i^2}}\right), \qquad i = 1, 2, \ldots$$

Conditional default probability: given market shock

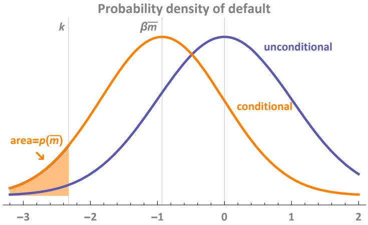

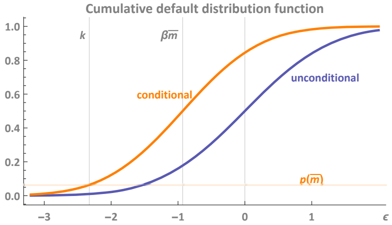

Density and cumulative probability as a function of idiosyncratic shock. Graph assumes $\beta_i = 0.4, k_i = -2.33 \quad (\Leftrightarrow \pi_i = 0.01)$, and $\bar{m} = -2.33$. The unconditional default distribution is a standard normal, while the conditional default distribution is $\mathcal{N}(\beta_i \bar{m}, \sqrt{1 - \beta_i^2}) = \mathcal{N}(-0.9305, 0.9165)$. The orange area in the density plot and horizontal grid line in the cumulative distribution plot identify $p(\bar{m})$, as in the example.

Conditional default distributions: example

- Firm: $\beta_i = 0.4$, $k_i = -2.33$ (so $\pi_i = 0.01$)
- Market shock: $\bar{m} = -2.33$ (sharp downturn)

	Unconditional	Conditional	Change
Mean return	0	−0.9305	−0.9305
Return variance	1	0.8400	−0.1600
Return std. deviation	1	0.9165	−0.0835
Distance-to-default	2.33	1.3958	−0.9305
(standardized)	2.33	1.5230	−0.8034
Default probability	0.01	0.0639	0.0539

Properties of the conditional distribution

- Once the market factor is realized, the default distributions of individual loans/obligors are independent
- But the market factor continues to be a random variable—together with idiosyncratic risk—driving default
- Both parameters β_i and k_i continue to influence the shape of the distribution function

Conditional default distributions

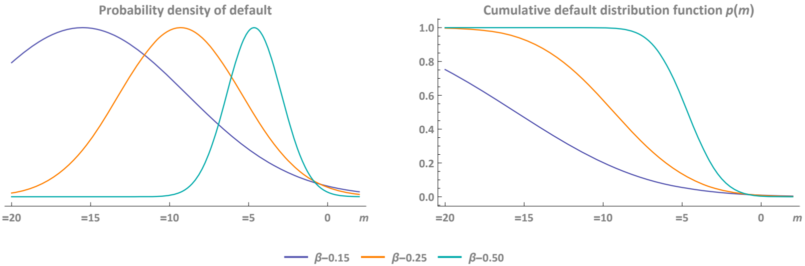

Probability of default of a single obligor, conditional on the realization of m (x axis). Default probability 1 percent ($k = -2.33$). conditional cumulative distribution function of default $p(m)$. Values of the distribution function run from 1 to 0 because it is plotted against m rather than $\frac{k-\beta m}{\sqrt{1-\beta^2}}$.