Chapter 16

Credit

In these notes we price the instruments that trade on default. Chapter 2 built the default time and observed that its survival probability is formally a discount factor; here we collect on that observation, deriving the credit default swap, the relation between a spread and a hazard rate — which turns out to be exact rather than approximate — and the bootstrap that turns a handful of quotes into a curve. We end with the one thing a single-name model cannot do, which is the subject of chapter 18.

16.1 What We Already Have

Chapter 2 did most of the modelling. Default is the first jump of a Cox process with intensity λ_t, the intensity exists because default arrives at a totally inaccessible time, and the survival probability is

$$Q(t) = \mathbb{P}(\tau > t) = \mathbb{E}\left[e^{-\int_0^t \lambda_s \, ds}\right], \tag{16.1}$$

which is chapter 7's bond price with the intensity where the short rate was.

What is missing is everything on the instrument side: what trades, how it is priced, and how λ is recovered from what trades. That is this chapter, and the reader who has been through chapters 7 and 8 will find the shape familiar — an instrument that pays a stream against a contingent payment, a par rate at which the two balance, and a bootstrap. It is the swap market again with a different underlying.

Throughout, R is the recovery rate: the fraction of notional a claim is worth after default. The loss given default is $1 - R$. We take λ deterministic, which makes (16.1) an ordinary integral, and note where stochastic intensity would change things.

16.2 The Credit Default Swap

Definition 16.1 (Credit default swap)**.** A credit default swap is an agreement in which the buyer of protection pays a periodic spread s on a notional until the earlier of maturity and default, and in exchange receives $1 - R$ at default if it happens before maturity.

It is an insurance contract written as a swap, and like a swap it has two legs to value.

Calculation 16.2 (The protection leg)**.** The buyer receives $1 - R$ at τ if $\tau \le T$. Default happens in $[t, t + dt]$ with probability $Q(t)\lambda_t\, dt$ — survive to t, then default immediately — so

$$\text{Prot} = (1 - R) \int_0^T P(0, t)\, \lambda_t\, Q(t)\, dt, \tag{16.2}$$

with $P(0, t)$ the ordinary discount factor of chapter 7. Note that $-dQ/dt = \lambda_t Q(t)$, so this is the loss integrated against the density of the default time.

Calculation 16.3 (The premium leg)**.** The seller receives s per period while the name survives. On payment dates t_i with accruals δ_i, the expected discounted receipts are

$$s \sum_i \delta_i P(0, t_i) Q(t_i),$$

plus one correction. If default falls partway through a period, the protection buyer still owes the spread accrued up to that moment; conventionally this is taken as half a period, contributing

$$s \sum_i \tfrac{1}{2} \delta_i P(0, t_i) \big(Q(t_{i-1}) - Q(t_i)\big).$$

The sum of the two, per unit of spread, is the *risky annuity* RiskyAnn — chapter 7's annuity with survival probabilities inserted alongside the discount factors.

Definition 16.4 (Par spread)**.** The par spread is the s making the contract worth zero:

$$s = \frac{\text{Prot}}{\text{RiskyAnn}}. \tag{16.3}$$

This is chapter 7's par swap rate again — a contingent leg over an annuity — and everything structural follows as it did there. The par spread is what is quoted. The risky annuity is the sensitivity of the contract's value to the spread, so it is the desk's risk measure, quoted in the same units and for the same reason.

16.3 The Credit Triangle

There is a relation between spread, hazard and recovery that every credit desk uses to convert between them in their head. It is usually presented as an approximation. It is not.

Theorem 16.5 (The credit triangle)**.** *With a flat hazard rate λ and premium paid continuously, the par spread is*

$$s = \lambda\, (1 - R) \tag{16.4}$$

exactly, for any deterministic interest rate curve.

Proof. With λ constant, $Q(t) = e^{-\lambda t}$. Writing the premium leg as a continuous integral, the two legs are

$$\begin{aligned} \text{Prot} &= (1-R)\int_0^T P(0,t)\,\lambda\, e^{-\lambda t}\,dt, \\ \text{RiskyAnn} &= \int_0^T P(0,t)\, e^{-\lambda t}\,dt. \end{aligned}$$

The integrands differ by the constant factor $(1-R)\lambda$ and by nothing else. So the ratio (16.3) is that constant, whatever $P(0,\cdot)$ happens to be. □

The interest rate curve cancelling entirely deserves a pause. A credit default swap is a swap between two streams that stop at the same moment and are discounted by the same factors, so the discounting divides out. That is why a credit spread can be quoted and traded without reference to the rate curve, and why the credit market and the rates market can move independently without either repricing the other.

Example 16.1 (Converting in your head)**.** A name quoted at 150 basis points, at the conventional forty percent recovery, is defaulting at

$$\lambda = \frac{0.0150}{0.6} = 0.025,$$

two and a half percent a year. Over five years that is a survival probability of $e^{-0.125} = 88\%$, so the market is pricing about a one in eight chance of this name failing within five years. Running it the other way, a desk that believes a name defaults with probability one in twenty over five years believes $\lambda \approx 0.01$ and should be paid about 60 basis points to insure it.

Doing this conversion quickly is most of what the triangle is for. It turns a quoted price into a statement about the world.

Calculation 16.6 (What the payment convention costs)**.** Real contracts pay quarterly, not continuously, so (16.4) is not exactly the quoted spread.

A premium paid at the end of each period is on average half a period later than one paid continuously, and that delay is discounted by both the interest rate and the survival probability. But the accrued-interest convention — the half-period term in the premium leg — was put there precisely to compensate the survival part of the delay. What it does not compensate is the interest rate part. So the error should be about half a period of interest,

$$\frac{s_{\text{quarterly}} - \lambda(1-R)}{\lambda(1-R)} \approx \frac{r\,\delta}{2}, \tag{16.5}$$

with δ the accrual period. At a three percent rate and quarterly payment this predicts 0.375%, and the model in `quant/src/credit.rs` gives 0.376%. At annual payment it predicts 1.5% and gives 1.517%. At a zero interest rate it predicts nothing at all, and the computed error is zero to eight decimal places at every spread tried.

Notice what does *not* appear in (16.5): the spread. The error is 0.38% at thirty basis points and 0.36% at twelve hundred. This is the opposite of the usual guess — that the triangle "breaks down for distressed names" — and the reason is Theorem 16.5: the hazard rate cancels between the legs whatever it is, so nothing that survives the cancellation can depend on it. What survives is the interest rate, which the theorem removed only because continuous payment let it cancel too.

Remark (Recovery cannot be fitted)**.** Equation (16.4) pins down the product $\lambda(1-R)$ and not the two factors. A quoted spread of 125 basis points is equally consistent with a two percent hazard and forty percent recovery, or a three percent hazard and sixty percent recovery, and no amount of staring at that quote will separate them.

This is why recovery is fixed by convention — forty percent for senior unsecured, conventionally — rather than calibrated. It is not laziness. There is nothing in the price to calibrate it against, and pretending otherwise would produce a number that looked estimated and was in fact assumed. Separating the two requires an instrument sensitive to recovery on its own, which is what a recovery swap or a pair of bonds at different seniority provides.

Compare this with the identification problems of chapters 10 to 12. It is the same shape — a product of parameters is observable, the factors individually are not — but with an unusually honest resolution, since the market agreed on a convention rather than each desk picking its own.

16.4 Bootstrapping the Hazard Curve

One quote gives one hazard rate. A curve of quotes — one, three, five, seven and ten years are the liquid points — gives a term structure, and the procedure is chapter 7's.

Take the hazard rate to be piecewise constant between quoted maturities. Then the one year quote involves only the first segment and fixes it. The three year quote involves the first segment, now known, and the second, which it therefore fixes. And so on outward.

Remark. The parallel with chapter 7 is exact enough:

Rates	**Credit**
short rate r_t	hazard rate λ_t
discount factor $P(0,T)$	survival probability $Q(T)$
par swap rate	par CDS spread
annuity	risky annuity
bootstrap from swaps	bootstrap from CDS
forward rate	forward hazard

And the correspondence is not an analogy but an identity of formulas, for the reason chapter 2 gave: (16.1) and the bond price are the same expectation.

Example 16.2 (The curve behind the figure)**.** The five quotes drawn opposite, at forty percent recovery and a three percent rate, bootstrap to[1]

Tenor	**Quote**	**Triangle**	**Forward hazard**
1y	80 bp	133 bp	133 bp
3y	105 bp	175 bp	197 bp
5y	125 bp	208 bp	264 bp
7y	138 bp	230 bp	295 bp
10y	150 bp	250 bp	311 bp

[1] `the whole worked example, checked against the bootstrap`.

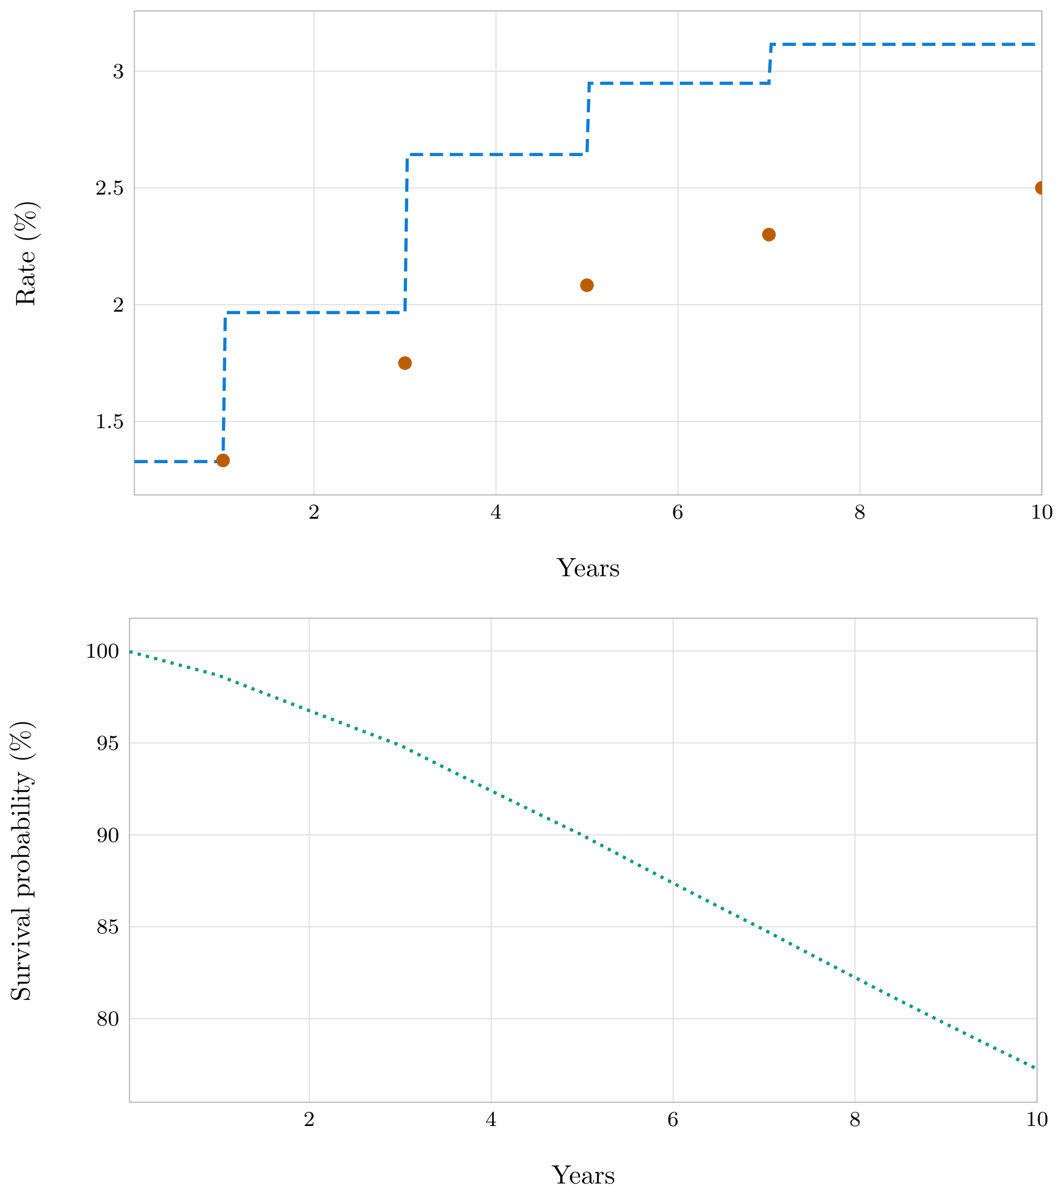

Figure 16.1: A hazard curve bootstrapped from five quoted spreads, with the survival curve it implies. The circles are the triangle applied to each quote individually, which is what a desk computes in its head; the line is the forward hazard the bootstrap actually implies. They agree at the front, where there is nothing before to average, and separate as the curve extends — the forward hazard reaching over three hundred basis points where the ten year triangle reads two hundred and fifty. This is chapter 7's lesson repeating: a quote is an average over its whole life, forwards move further than the averages that contain them, and it is the forwards that anything between the quoted tenors prices off.

Drawn by `bootstrap_hazards` and `par_spread`.

Read the last two columns against each other. At one year they agree, because there is nothing before the first quote for it to be an average of. By ten years the triangle reads 250 basis points and the forward hazard is 311, a quarter higher — the ten year quote is an average over the whole decade, and the market is saying the last three years of it are considerably worse than the first.

A desk pricing a seven-to-ten year forward CDS off the ten year triangle would be using 250 where the curve says 311, and would sell that protection a fifth too cheap. This is chapter 7's warning again in a new market: a quote is an average, a product prices off a forward, and the two are not interchangeable.

16.5 Marking a Position

The par spread prices a new contract. An existing one was struck at some spread K that the market has since left behind.

Calculation 16.7 (The value of an off-market contract)**.** A protection buyer paying K when the market charges s is paying $K - s$ too much, on a stream that survives exactly as the annuity says. So the position is worth

$$V = (s - K) \times \text{RiskyAnn}, \tag{16.6}$$

to the buyer of protection, which is chapter 7's swap valuation with the risky annuity in place of the ordinary one.

Concretely, using the curve above: a five year contract struck at 100 basis points, marked against a market at 125, with a risky annuity of 4.43, is worth[2]

$$0.0025 \times 4.43 = 1.11\%$$

of notional to the buyer of protection. Twenty five basis points of spread became a hundred and eleven basis points of value, because the position runs for about four and a half years' worth of annuity.

Remark (Why contracts trade with an upfront)**.** Equation (16.6) explains a market convention that otherwise looks arbitrary. Standard contracts do not trade at their par spread at all; they pay a fixed coupon, 100 basis points for investment grade names and 500 for high yield, and the difference between that coupon and the par spread is settled as a cash payment at the start.

That payment is exactly V. Fixing the coupon means every contract on a name has identical cash flows regardless of when it was struck, so two positions can be netted against each other and torn up, instead of accumulating as an ever-growing book of offsetting trades with different coupons. The convention exists to make positions fungible, and (16.6) is the formula that converts between the running spread a trader thinks in and the upfront the contract actually settles.

Remark (The annuity is shorter than it looks)**.** The riskless five year quarterly annuity at a three percent rate is 4.63.[3] The risky one above is 4.43, four percent shorter. The missing four percent is the premium that will not be paid because the name may not survive to pay it — which is to say the risky annuity is doing real work, and substituting the ordinary annuity into (16.6) would overstate the position by that much. On a distressed name the gap is far larger.

[2] `the risky annuity and the mark to market, checked`.

[3] `the riskless annuity, checked`.

16.6 The Two Risk Numbers

A credit position carries two risks that cannot be netted, and chapter 2 explained why: they are the diffusive part and the jump part of the same position, and they require different instruments.

Definition 16.8 (CS01)**.** The change in value for a one basis point widening of the spread curve. By (16.3) it is the risky annuity times a basis point, exactly as a swap's DV01 is its annuity.

Definition 16.9 (Jump to default)**.** The change in value if the name defaults immediately. For a protection buyer it is a gain of $(1 - R)$ less the value already marked; for a bond holder a loss of $(1 - R)$ times notional.

CS01 is the risk of the spread drifting. Jump to default is the risk of the event. A position can be flat on one and enormously exposed on the other — selling protection on a name trading at fifty basis points has a small CS01 and a jump to default of sixty percent of notional — which is precisely why they sit on separate lines of a risk report rather than being combined into one number.

16.7 Making the Intensity Random

Everything so far took λ deterministic.

A deterministic intensity makes the whole future spread curve known today. So an option on a spread is worth its intrinsic value, a position in spread volatility does not exist, and the only uncertainty left in the model is *when* the default arrives rather than *what the market will think* of the name in a year. That is the same complaint chapter 9 made about a deterministic volatility, arriving in a different market, and it has the same remedy.

Let λ be a process. The survival probability is then

$$Q(0,T) \;=\; \mathbb{E}\Big[\exp\Big(-\int_0^T \lambda_s\,ds\Big)\Big], \tag{16.7}$$

and the reader should recognise that expectation, because it is the one chapter 8 spent a chapter on with r in place of λ. A credit curve is a discount curve on a different rate, so every model in this book's term structure chapters transfers to it unchanged — and so does every one of their limitations.

Remark (Which model, and the one constraint credit adds)**.** The choice is not quite free, because an intensity is an arrival rate and a negative arrival rate is meaningless. Hull-White, which allows r to go negative and is none the worse for it in rates, is not admissible here.

The usual answer is a square-root process, $d\lambda = \kappa(\theta - \lambda)dt + \eta\sqrt{\lambda}\,dZ$, whose diffusion switches off as the intensity approaches zero. That is chapter 10's backbone question with a hard constraint attached: the exponent is not being fitted to data here but forced to a half by the requirement of positivity.

The reward is that the model is affine, so (16.7) is exponential-affine in λ_0 with coefficients solving a Riccati pair, by the theorem of chapter 14 — and for this process the pair has a closed form. The credit curve is then computed by evaluating a formula rather than by simulating, which is what makes calibration to a strip of CDS quotes practical.[4]

[4] `credit::CirIntensity`, checked against a simulation sharing none of its algebra.

Calculation 16.10 (What randomness alone does to the spread)**.** Randomness in the intensity changes the curve even when it changes nothing about the intensity's average, and the direction is fixed rather than a matter of parameters.

Take $\lambda_0 = \theta = 2\%$, so the intensity's mean is flat at two percent for every horizon, and $\kappa = 0.5$, $R = 40\%$. With $\eta = 0$ the credit triangle gives a flat five year spread of 120 basis points. Turning on the vol of intensity lowers it monotonically, and at $\eta = 0.20$ it has fallen by 4.2 basis points.[5]

The reason is Jensen: (16.7) is a convex function of the integrated intensity, so $\mathbb{E}[e^{-X}] > e^{-\mathbb{E}[X]}$ and a random intensity survives better than its own average. A desk bootstrapping a deterministic curve from quotes is therefore not recovering the mean intensity; it is recovering something below it, by an amount that is a pure convexity and grows with the volatility of the spread.

16.8 Options on Credit

With the intensity random there are credit derivatives to price, and the machinery for them is already in this chapter.

Definition 16.11 (Payer and receiver swaptions on a CDS)**.** A payer gives the right at T to buy protection to T_b at a strike spread K; a receiver, the right to sell it.

The valuation follows chapter 8's swaption argument with one substitution. Equation (16.6) says a position struck at K is worth $(s-K)\text{RiskyAnn}$, so the payer's payoff at T is $\text{RiskyAnn}(T)\,(s(T)-K)^+$ — an option on the forward spread, with the *risky* annuity multiplying it. The risky annuity is a positive traded portfolio, so it is an admissible numeraire; the forward spread is a traded value divided by it, so it is a martingale in the corresponding measure; and the option prices by Black in that measure, exactly as a swaption does.

So a credit desk quotes spread volatility, runs a smile, and inherits every question of chapter 10 — with lognormal the usual convention, since a spread cannot go negative and is far from zero.

Remark (What is different: the numeraire can vanish)**.** The parallel with swaptions is exact but for one thing. The risky annuity is not merely positive: it can jump to zero, because the name can default before the option expires. A numeraire that can vanish is not a numeraire on that event, and the argument above says nothing about what happens there.

The market resolves this by contract rather than by mathematics, and in two different ways. A single name option that *knocks out* on default is worthless if the name goes, and for it the Black argument holds on the surviving paths, which are the only paths that matter. An index option does not knock out — it carries *front-end protection*, paying the losses incurred between today and expiry — and its value is then the Black term plus the value of that protection, which is not an option at all but a forward-starting exposure to the jump risk of §16.6.

The two are quoted in the same units and are not the same instrument. It is also the cleanest illustration in the book of a numeraire argument being contract-dependent: what makes the pricing work is a clause, not a theorem.

Constant maturity CDS. A contract paying a periodically reset spread of fixed tenor, rather than a fixed one. It is chapter 15's problem in credit: the rate is observed under one measure and

[5] `measured`.

paid under another, so it carries a convexity adjustment, and the same replication over the spread smile prices it under the same kind of mapping assumption.

Recovery locks and digital default swaps. This chapter's own identification problem was that λ and R enter the par spread only through their product, so no CDS quote separates them. A recovery lock trades R directly, and a digital swap pays a fixed amount on default rather than $1-R$. Where they trade, the identification failure stops being an assumption and becomes a market price — which is the honest resolution of a degeneracy, and rarer than one would like.

16.9 What One Name Cannot Tell You

Everything above concerns a single name, and for a single name the model is in good order: the intensity is identified up to the recovery convention, the curve bootstraps, and the two risks are separately measurable and separately hedgeable.

Portfolios are a different matter, and the reason is not a defect in the machinery so far but a question it never asked. A collateralised loan obligation, an index tranche, or the counterparty exposure of chapter 24 all depend on *how many* names default together, and nothing in a collection of single-name curves says anything about that. Each curve fixes its own marginal default distribution and the joint distribution is left entirely open — which is chapter 9's Gyongi problem in another costume: the marginals are pinned by the market and everything about the dependence is unconstrained.

Assuming independence is not neutral, it is a strong and badly wrong assumption. Defaults cluster: they share exposure to the economy, to funding conditions, and to each other through direct claims. A portfolio model that assumes independence will price the senior tranche of anything at nearly zero, because it assigns essentially no probability to many names failing at once, and that specific error has a history.

How dependence is modelled, what the industry standard gets wrong, and what to use instead is the subject of chapter 18.

References

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

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

- O'Kane, D. (2008). *Modelling Single-name and Multi-name Credit Derivatives.* Wiley.