ullint6b.tex am Wed 7.3.2018

VI. CREDIT RISK

1. Introduction

As general introductions to the area of credit risk, we refer to [BM, VII, Ch. 21 - 23], [BPT], [BMP], [B]. Also useful:
[L] David LANDO, *Credit risk modelling*, Princeton University Press, 2004.

As we all know, and as the events of the Crash and after daily remind us, obligations are not always fulfilled: there is the possibility of *default.* We turn now to introducing this possibility into our models.

We begin here, as we began the course, with the simplest and most basic interest-rate product – the *zero-coupon bond (ZCP).* This pays 1 at time T (maturity); its price at time $t \in [0, T]$ is $P(t, T)$. Now suppose that the issuer of the bond (a company, or even a government) may *default.* In this case the payoff of the ZCB at time T is:

1 with no default; 0 with default.

We write the bond price at time t now as $\overline{P}(t, T)$. Clearly,

$$0 \leq \overline{P}(t.T) \leq P(t, T).$$

When considering default, we have a random time τ at which the bond issuer defaults ($\tau = \infty$ if there is no default – the waiting time for something that never happens is infinite).

The value of the bond issued by the company and promising payment of 1 at time T is, as usual (MATL480 and Ch. I) *the risk-neutral expectation of the discounted payoff.* We now have *two* relevant filtrations (information flows):

$\{\mathcal{F}_t\}$, the default-free filtration we have used till now (still relevant to default-free market variables, such as the risk-free short rate r_t);

$\{\mathcal{G}_t\}$, the $\mathcal{F}$-filtration augmented by information on whether or not default has occurred by time t:

$$\mathcal{G}_t := \mathcal{F}_t \vee \sigma(\{\tau \leq u\}, 0 \leq u \leq t),$$

where, used between two σ-fields, $\vee$ means 'the σ-field generated by (both of these)'.

With $D(t, T)$ the stochastic discount factor between two dates, as before, we now have

$$P(t, T) = E[D(t, T)\ 1\ |\mathcal{F}_t],$$

$$I\{\tau > t\}\overline{P}(t, T) = E[D(t, T)I\{\tau > T\}|\mathcal{G}_t].$$

Here, for an event A, its *indicator function* is

$$I(A) := 1 \quad \text{if } A \text{ occurs}, \qquad 0 \quad \text{otherwise}.$$

So $I\{\tau > T\}$ is the payoff for a defaultable bond: the contract pays 1 if the issuer (usually a company) has not defaulted, 0 if it has.

In particular, as $\tau > 0$ (so $I(\tau > 0) = 1$),

$$\overline{P}(0, T) = E[D(0, T)I\{\tau > T\}]. \qquad (*)$$

Note. 1. This ignores issues involving *recovery.* Usually, a company will only default if it is going bankrupt. Then – as its liabilities exceed its assets – the company dies, and its assets are distributed to the creditors by the receivers or liquidators. We will not go into this in detail, as bankruptcy laws vary with country and over time.
2. To be specific: under US law, what a company usually does is to file under Chapter 11 for protection from creditors. Homework: explore this on the Internet.
3. We have met enlargement of filtrations before, in MATL480, in the context of *insider trading*, and its detection. The analogy between financial crime there, and default here, is interesting. But there, the focus is on *detection*; here, it is on *prediction* or *prevention* if possible – and coping with it if not.

Taking recovery into account: the discounted payoff becomes

$$D(t, T)I\{\tau > T\} + REC\ D(t, \tau)I\{\tau \leq T\}$$

if any such recovery is made at the time τ of default, and

$$[D(t, T)I\{\tau > T\} + REC\ D(t, T)I\{\tau \leq T\}$$

if it is made instead at the time T of maturity. Taking $E[.|\mathcal{G}_t]$ in each case gives the price of the bond.

Financial forensics

When a major default occurs, it is necessary to find out *why*. This is just as necessary for the health of the financial world as finding out who did it and why after a major crime – and for the same reasons. The leading exponent of financial forensics was Steve (Professor Stephen A.) Ross (1944-2017) (of the Cox-Ingersoll-Ross model and the Cox-Ross-Rubinstein binomial tree), who also pioneered the arbitrage pricing technique (APT) in the 1970s. He was an excellent speaker, and writer, on this important area.

2. Credit default swaps (CDS)

Credit default swaps (CDSs) are basic protection contracts that became quite liquid on a large number of products after their introduction in 1991-94. CDSs reached a notional value of $ 3.7t (trillion) in 2003, $ 62.2t in 2007, $ 38.6t in 2008, $ 25t in 2012 (ISDA figures – International Swaps and Derivatives Association). They are now actively traded, and are a basic product in the area of credit derivatives, analogous to IRSs and FRAs (II.2) being basic products in the world of interest-rate derivatives.

We do not need a model to value CDSs; rather, we need a model that can be *calibrated* to CDSs – that is, to take CDSs as model *inputs* (rather than outputs), in order to price more complex derivatives.

Regarding options: single-name CDS options have never been liquid. What is liquid is CDS *index* options – options on an index (Footsie, S&P, Dax etc.) – giving protection against an index (which reflects the state of a national economy) rather than an individual company. As above, we may expect to have to incorporate CDS index options into our model, rather than price them, as with CDSs themselves.

A CDS contract gives protection against default. Two companies, A (protection *buyer*) and B (protection *seller*) agree on the following. If a third company C (reference credit) defaults at time τ, with

$$T_a < \tau < T_b,$$

B pays A a certain deterministic cash amount LGD, the *loss given default*. In return, A pays to B a *rate* R at times $T_{a+1}, \cdots, T_b$ or until default – at time τ_C (infinite if C does not default). Set $T_0 = 0$,

$$\alpha_i := T_i - T_{i-1}.$$

Typically, LGD is the notional, 1, or notional-minus-recovery, $1 - REC$. This just says that money is either *lost* or *recovered* – true to a first approximation (but neglects such things as the costs of the administrators and liquidators who must take over the company's assets and see that the creditors are repaid as much as possible).

A typical stylised case occurs when A has bought a corporate bond issued by C, and is waiting for the coupons (if any) and final notional payment from C. If C defaults before the corporate bond maturity, A does not receive such payments. To guard against this, A goes to B and buys some protection against this risk, asking B for a payment (in case of default) that roughly amounts to the loss on the bond (e.g. notional minus deterministic recovery) that A would face in case C defaults. Again, A may hold a portfolio of several instruments with a large exposure to counterparty C. To partly hedge such exposure, A enters into a CDS, by buying protection from a bank B against default by C.

What counts as a credit event triggering τ_C? Possibilities include:
(a) bankruptcy of C;
(b) failure of C to pay (or pay on time)[1];
(c) obligation acceleration, when C is required to pay ahead of schedule because of C's failure to meet the terms of the loan;
(d) restructuring, when C undergoes reorganisation to consolidate its debt. There are several types of restructuring; definitions and legislation vary, e.g. between Europe and the USA.

What happens in a CDS contract at default of C?
(a) *Cash settlement.*

The protection seller pays to the buyer the loss value of the referenced instruments (e.g. bonds issued by C), following the credit event. The bonds or loans themselves are not transferred. When more instruments can be referenced the cheapest-to-deliver price variation is used (see below).
(b) *Physical settlement.*

The protection buyer receives a cash payment, typically the insured face value, from the seller, and the seller takes possession of the defaulted loan instrument or bonds for an equivalent notional amount.

[1] Sharp practice exists in the business world regarding paying on time. Of course one should do this, morally and to fulfil one's legal obligations, and for the sake of one's business reputation, which is a valuable asset and hard to replace. But, big companies have leverage over small companies, because of the asymmetry of the relationship: an order from a big company is harder for the small company to replace. This can be exploited by an unscrupulous company. For example, in the construction industry, contractors employ sub-contractors; the practice of deliberately paying them late is called "subby-bashing".

Here most CDSs allow the protection buyer to choose deliverables from a pool of defaulted bonds with equal seniority. The cheapest-to-deliver bond is typically chosen (different value in a reorganisation, higher accrued interest, ...).

Physical settlement: Auction.

If there are not enough bonds to match the insured face value, a credit-event *auction* occurs, and the payment received is usually substantially less than the face value of the loan.

Recovery rate, REC.

The recovery rate REC is implicitly defined by these procedures and by market-value decline after a credit event, and is very hard to estimate before the event.

Before the Crash in 2007, REC = 40 % was a typical figure, and REC = 50 % for financials. But, Lehman REC in immediate auction was 8.625 %! Lehman asset liquidations are still ongoing.[2] Recovery has led to legal battles. The final recovery might exceed 40 %.

ISDA in 2009 recommends REC 20 % or 40 %.

Analysis is mostly possible in aggregate on large pools of bonds or loans with similar ratings (one can estimate more precisely for a large sample than for an individual case).

There are only a few studies available. In aggregate, there is an inverse relationship between *recovery rates REC* and *credit risk/spread or default rates* – as one would expect. One can postulate such an inverse relationship between spreads and recoveries, but there is no consensus on how to make this precise.

We may write the running CDS discounted payoff to B at time $t < T_a$ as

$$\begin{aligned}\Pi_{RCDS}(t) \;:=\;& RD(t,\tau)(\tau - T_{\beta(\tau)-1})I(T_a < \tau < T_b)\\ &+R\sum_{i=a+1}^{b} D(t,T_i)\alpha_i\; I(\tau > T_i)\\ &-LGDI(T_a < \tau \leq T_b)D(t,\tau),\end{aligned}$$

where $T_{\beta(\tau)}$ is the first of the T_is following τ. Here, the three terms in the payout on the right correspond to:

[2] Over 350 firms participated in the auction following the collapse of Lehman.

1. Discounted accrued rate at default. This is supposed to compensate the protection seller for the protection he provided from the last T_i before default to default at τ.
2. CDS rate premium payments if there is no default. This is the premium received by the protection seller for the protection being provided.
3. Payment of protection at default if this happens before the final T_b.
These are random discounted cash flows, not yet the CDS price.

To find the price, we take the risk-neutral expectation, as usual:

$$CDS_{a,b}(t, R, LGD) := E[\Pi_{RCDS}(t)].$$

This pricing formula depends on the assumptions on interest-rate dynamics (as in previous chapters), and assumptions on the default time τ. These are of various kinds: reduced-form models, structural models, etc. (see below).

As usual, we will not use the resulting formulas to price CDSs already quoted in the market. Rather, we will invert these formulas for the corresponding CDS market quotes to calibrate our models to the CDS quotes themselves. We will see examples later.

Model-independent formulas

We will assume that the default times τ are *independent* of the stochastic discount factors $D(s, t)$. This is reasonable:
(a) the default is at *company* level (micro-economic), while the discount factors are at national or international level (macro-economic);
(b) we cannot proceed without such an assumption.

Nevertheless, full independence may well not hold. Major events at world level (the Crash; wars; major political events – Brexit, Trump, ...; major terrorist attacks; oil crises; natural disasters (tsunamis, earthquakes, hurricanes, ...) affect both.

The price of the premium leg of the CDS is

$$\begin{aligned} PremiumLeg_{a,b}(R) &= E[D(0,\tau)(\tau - T_{\beta(\tau)-1})\ R\ I(T_a < \tau < T_b) \\ &\quad + \sum_{i=a+1}^{b} D(0,T_i)\alpha_i\ R\ I(\tau > T_i)] \end{aligned}$$

$$= E[\int_0^\infty D(0,t)(\tau - T_{\beta(\tau)-1})\ R\, I(\tau \in [t, t+dt]) + R \sum_{i=a+1}^{b} D(0,T_i)\alpha_i E[I(\tau \geq T_i)]$$

$$= E[\int_{T_a}^{T_b} D(0,t)(\tau - T_{\beta(\tau)-1})\; R\, I(\tau \in [t, t+dt]] + R \sum_{i=a+1}^{b} P(0,T_i)\alpha_i \mathbb{Q}(\tau \geq T_i)$$

(by independence: $E[XY] = E[X].E[Y]$ with X, Y independent)

$$= R \int_{T_a}^{T_b} E[D(0,t)(\tau - T_{\beta(\tau)-1})]E[I(\tau \in [t, t+dt])] + R \sum_{i=a+1}^{b} P(0,T_i)\alpha_i \mathbb{Q}(\tau \geq T_i)$$

$$= R \int_{T_a}^{T_b} P(0,t)(\tau - T_{\beta(\tau)-1})\mathbb{Q}(\tau \in [t, t+dt]) + R \sum_{i=a+1}^{b} P(0,T_i)\alpha_i \mathbb{Q}(\tau \geq T_i),$$

using independence again.

We can reduce to the $R = 1$ case:

$$PremiumLeg_{a,b}(R; P(0,.), \mathbb{Q}(\tau > .) = R\, PremiumLeg1_{a,b}(P(0,.), \mathbb{Q}(\tau > .)),$$

where

$$\begin{aligned} PremiumLeg1_{a,b}(P(0,.), \mathbb{Q}(\tau > .)) \;\; &:= \;\; \int_{T_a}^{T_b} P(0,t)(\tau - T_{\beta(\tau)-1})d_t\mathbb{Q}(\tau \leq t) \\ &\quad + \sum_{i=a+1}^{b} P(0,T_i)\alpha_i \mathbb{Q}(\tau \geq T_i). \end{aligned}$$

This model-independent formula uses the initial market ZCB curve (bonds) at time 0 (i.e. $P(0,.)$) and the survival probabilities $\mathbb{Q}(\tau \geq .)$ at time 0.

A similar formula holds for the protection leg, again under independence between default τ and interest rates:

$$ProtectLeg1_{a,b}(P(0,.), \mathbb{Q}(\tau > .)) = \int_{T_a}^{T_b} P(0,t)d_t\mathbb{Q}(\tau \leq t).$$

This too is model-independent.

So one obtains CDS prices:

$$\begin{aligned} CDS_{a,b}(t, R, LGD; \mathbb{Q}(\tau \leq .)) \;\; &= \;\; -LGD[\int_{T_a}^{T_b} P(0,t)d_t\mathbb{Q}(\tau \leq t)] \\ &\quad + R[\int_{T_a}^{T_b} P(0,t)(\tau - T_{\beta(\tau)-1})\mathbb{Q}(\tau \in [t, t+dt]) \\ &\quad + \sum_{i=a+1}^{b} P(0,T_i)\alpha_i \mathbb{Q}(\tau \geq T_i)]. \end{aligned}$$

The integrals here in the survival probabilities are Stieltjes integrals, and can be approximated numerically by Riemann-Stieltjes sums, by using a small enough discretisation time-step.

CDS stripping

The market quotes, at time 0, the fair $R = R_{0,b}^{mktMID}(0)$ coming from bid and ask quotes for this fair R (MID: average the bid and ask quotes). This fair R equates the two legs for a set of CDSs with initial protection time $T_a = 0$ and final protection time

$$T_b \in \{1y, 2y, 3y, 4y, 5y, 6y, 7y, 8y, 9y, 10y\},$$

although often only a subset of these maturities ($\{1, 3, 5, 7, 10\}$ say) is available.

So: solve

$$CDS_{0,b}(t, R_{0,b}^{mktMID}(0), LGD; \mathbb{Q}(\tau > .)) = 0$$

for the relevant cases of $\mathbb{Q}(\tau > .)$:
(i) starting from $T_b = 1y$, and finding the market-implied survival

$$\{\mathbb{Q}(\tau \geq t), t \leq 1y\};$$

(ii) plugging this into the $T_b = 2y$ CDS legs formula, and then solving the same equation with $T_b = 2y$, we find the market-implied survival

$$\{\mathbb{Q}(\tau \geq t), t \in (1y, 2y]\},$$

and so on up to $T_b = 10y$.

This method is called *CDS stripping*. This is a way to strip survival (or equivalently, default) probabilities from CDS quotes in a model-independent way. There is no need to assume an intensity or a structural model for default here.

However, the market in doing the above stripping typically resorts to intensities (also called hazard rates), assuming existence of intensities associated with the default time. We turn to intensity models next.

3. Intensity models; stochastic intensity; Cox processes; Lando's formula

Recall the work on *hazard rates* in MATL480 Ch. VII, in the context of *survival analysis* in life insurance.

In intensity models, the random default time τ is assumed to be exponentially distributed.

Stochastic intensity

The deterministic intensity or hazard rates above account for credit-spread structure. But they do not account for volatility. To do this, we need to move to stochastic intensity – a *Cox process.* The deterministic

$$t \mapsto \gamma(t), \qquad \Gamma(t) = \int_0^t \gamma(u)du$$

now becomes the stochastic

$$t \mapsto \lambda(t) = \lambda_t, \qquad \Lambda(t) = \int_0^t \lambda(u)du.$$

A strictly positive stochastic process

$$t \mapsto \lambda_t,$$

called the *default intensity* or *hazard rate*, is given for the bond issuer or the CDS reference name.

The *cumulative intensity* or *hazard function* is the integrated process

$$\Lambda : \quad t \mapsto \Lambda_t := \int_0^t \lambda_s ds.$$

The *default time* τ can then be defined as the inverse of the process Λ applied to an exponentially distributed random variable ξ with mean 1 and independent of λ: thus

$$\xi \sim E(1): \qquad \mathbb{Q}(\xi > u) = e^{-u}, \quad \mathbb{Q}(\xi < u) = 1 - e^{-u}, \quad E[\xi] = 1,$$

$$\tau = \Lambda^{-1}(\xi), \qquad \xi = \Lambda(\tau) \sim E(1), \quad \text{independent of } \lambda.$$

Now the probability of surviving for time t is

$$\begin{aligned}
\mathbb{Q}(\tau > t) &= \mathbb{Q}(\Lambda^{-1}(\xi) > t) \\
&= \mathbb{Q}(\xi > \Lambda(t)) \\
&= E[I(\xi > \Lambda(t))] \\
&= E[E[I(\xi > \Lambda(t))|\mathcal{F}_t]] \qquad \text{(Conditional Mean Formula)} \\
&= E[e^{-\Lambda(t)}] \qquad (\xi \sim E(1)) \\
&= E[\exp\{-\int_0^t \lambda_s ds\}].
\end{aligned}$$

This looks exactly like a bond price if we replace r by λ!

We can see this coming!:

The reason we obtain the bond-price formula is because *(compound) interest is exponential*, by its very nature (see e.g. MATL480, I.1).

The reason we obtain the above survival-probability formula is because *defaults are exponentially distributed.*

We are now ready to price a defaultable ZCB, $\overline{P}(0,T)$ (with zero recovery, for simplicity). Recall (II.1) that in the non-defaultable case,

$$P(t,T) = E_t[\frac{B_t}{B_T}1] = E_t[\exp(-\int_t^T r_s ds)] = E_t[D(t,T)]. \qquad (P)$$

The result below is due to Lando in 1994 (thesis); see also

[Lan1] D. LANDO, Modelling bonds and derivatives with default risk, *Mathematics of derivative securities* (ed. M. A. H. Dempster & S. R. Pliska) 369-393, CUP, 1997,

[Lan2] D. LANDO, On Cox processes and credit-risky securities. *Review of Derivates Research* **2** (1998), 99-120.

Theorem (Lando's formula). The price of a defaultable bond with default intensity λ is just the price of a default-free bond, *where the risk-free short-rate r is replaced by $r + \lambda$.*

Proof.

$$\begin{aligned}
\overline{P}(0,T) &= E[D(0,T)I(\tau > T)] \qquad \text{(by } (*)) \\
&= E[\exp\{-\int_0^T r_s ds\}I(\Lambda^{-1}(\xi) > T)] \qquad \text{(def. of } D;\ \tau = \Lambda^{-1}(\xi))
\end{aligned}$$

$$
\begin{aligned}
&= E[\exp\{-\int_0^T r_s ds\} I(\xi > \Lambda(T))] \qquad (\text{apply } \Lambda \text{ inside } I(.)) \\
&= E[E[\exp\{-\int_0^T r_s ds\} I(\xi > \Lambda(T)) | \Lambda, r]] \qquad (\text{Tower property}) \\
&= E[\exp\{-\int_0^T r_s ds\}] E[I(\xi > \Lambda(T)) | \Lambda] \qquad (\text{independence}) \\
&= E[\exp\{-\int_0^T r_s ds\}] \ \mathbb{Q}(\xi > \Lambda(T)) | \Lambda) \qquad (E[I_A] = Prob(A)) \\
&= E[\exp\{-\int_0^T r_s ds\} \exp\{-\Lambda(T)\}] \qquad (\xi \sim E(1)) \\
&= E[\exp\{-\int_0^T r_s ds\} \exp\{-\int_0^T \lambda_s ds\}] \qquad (\text{defn. of } \Lambda) \\
&= E[\exp\{-\int_0^T (r_s + \lambda_s) ds\}]. \qquad //
\end{aligned}
$$

Note.
1. The *independence* assumption in Lando's formula is its weak point, and is open to question. When the economy is in trouble, *both* interest rates (including the spot rate r) *and* default intensities λ are affected. They will thus not be independent, as they both respond to the same macro-economic situation. For a more exact result, we would need the *joint distribution* of $(e^{-r_t}, e^{-\lambda_t})$. This would be difficult to model, and Lando's formula works well in practice, so it is widely used.
2. Cox processes (D. R. Cox (1924-) in 1955) were first introduced by Cox in statistical studies of fibre strength in the textile industry. They were used by Lando (1994, 1998) above, for the pricing of defaultable bonds. They are also widely used in the modelling of geophysical events such as volcano eruptions, earthquakes, tsunamis etc. See e.g.
[DVJ] D. J. DALEY and D. VERE-JONES, *An introduction to the theory of point processes. Vol. I, Elementary theory and methods; Vol. II, General theory and structure*. Wiley, 2003 (2nd ed.; 1st ed., one volume, 1998).
3. The use of martingale methods for point processes (Poisson processes, Cox processes etc.) is Poisson-based. Although it is easier than the martingale theory of processes based on Brownian motion (MATL480), it came later. For background, see
[Bre] Pierre BRÉMAUD, *Point processes and queues: Martingale dynamics.*

Springer, 1981.
Just as Brownian-based stochastic integration and martingales rest on (we give both notations)

$$(dB_t)^2 = dt : \qquad (dW(t))^2 = dt$$

(Lévy's theorem on quadratic variation of Brownian motion), its Poissonian analogue is based on

$$(dN_t)^2 = dN_t.$$

Here N_t ('N for number') counts e.g. renewals up to time t, or insurance claims, etc. (MATL480, VII). This just says that dN_t is 0 or 1: the count goes up by 1 when an event occurs.
4. Both the Wiener and the Poisson processes belong to the wider class of *Lévy processes* – processes with *stationary independent increments.* A stochastic calculus exists for them too, including the Wiener and Poisson cases above. See
D. Applebaum, *Lévy processes and stochastic calculus*, 2nd ed., CUP, 2009 (1st ed. 2004).
Stochastic calculus was developed for the much more general *semimartingales* (local martingale + finite variation) in the 1960s and 70s by P.-A. Meyer, but this is harder.