Supplementary notes and comments, Section 11

1   Credit spread

Suppose a risk free bond with a given coupon and maturity has a total yield Y_{Tr} (Tr for Treasury). A bond with the same maturity and coupons may have a higher total yield Y_R (R for risky). (We compute the effective total yield from the market price in the same way for both bonds: what constant discount rate Y applied to all the payments, coupon and principal, makes the value of the bond equal to its quoted price.) The *credit spread* is $s = Y_R - Y_{Tr}$.

The credit spread for a particular bond reflects the market's view of the riskiness of that bond. There are so many bonds with so many features from so many companies and other entities (cities, universities, etc.) that it is hard for an individual investor to form an educated opinion of the credit quality of a given bond. Most bonds (certainly by number, if not by total value) trade so rarely that there is no meaningful market price for that individual bond. For this reason, *rating agencies* offer the service of offering an opinion about the default probability of individual bonds, placing them into broad categories ranging from AAA (the best) to "junk".

Bonds with the same rating and duration ($\neq$ maturity) often trade (when they do trade) at a similar credit spread. But markets may and do form different implicit views of the credit quality of a particular bond than rating agencies. The market has still other views of credit quality of individual bonds, including rates charged for Credit Default Swaps (CDS).

2   Continuous time default

I present the material on default in continuous rather than discrete time. The formulas are possibly slightly less realistic (though not much), but they are simpler. Consider a bond that is supposed to pay a continuous coupon $C\,dt$ in time dt. The maturity time is T_M and the principle is one unit of currency. If the bond defaults, the default time is T_D. This is a random time. To make discussions easier, there may be a probability distribution for T_D that makes any positive value of T_D possible. We say that the bond *defaults* if $T_D \leq T_M$. If a bond does default, it makes a payment R (the *recovery rate*) at time T_D. The recovery is between 0 and 1. A typical value might be about $.4 = 40\%$, but it could go as high as, maybe, 70% and as low as zero. One could consider R to be random, though that makes no difference in what we do here.

First, consider a bond with no possibility of default. Suppose the risk free rate, r, is known and constant. Suppose that $T_M = \infty$, in order to simplify calculations. We compute the total present value of all the coupon payments.

A payment $C\,dt$ at time t has present value $e^{-rt}C\,dt$. The integrated value of all payments is

$$\int_0^\infty Ce^{-rt}\,dt \;=\; \frac{C}{r}\;. \tag{1}$$

The bond sells at par $= 1$ if its value is one, which is to say that $C = r$ (DUH). Conclusion: a risk free bond sells at par if pays the risk free rate (again DUH, this is just to understand the integral.).

We now put in the simplest model of bond default, the constant intensity model. In this model, bonds default according to an *intensity* parameter, λ. The model is that bond default is a completely unpredictable event. If a bond has not defaulted before time t, the probability that it will default just after is

$$\Pr\left(\,t \le T_D \le t + dt \,|\, T_D \ge t\,\right) \;=\; \lambda\,dt\;. \tag{2}$$

This says that a bond that has not defaulted in ten years is just as likely to default in year eleven (the first year after year ten) as it was to default in year one. Mathematically, this makes single name bond default a Markov process. Knowing that age of a bond does not help predict its future default behavior.

The formula (2) leads to a formula for the probability density of T_D, which we call $f(t)$. In fact, if $Q(t) = \int_t^\infty f(t)\,dt$ is the *survival probability*, then $f(t) = -\partial_t Q(t)$, and (think this through)

$$\Pr\left(\,t \le T_D \le t + dt \text{ and } T_D \ge t\,\right) \;=\; \Pr\left(\,t \le T_D \le t + dt\,\right) \;=\; f(t)\,dt \;=\; -Q'(t)\,dt\;.$$

Therefore

$$\begin{aligned}
\lambda\,dt \;&=\; \Pr\left(\,t \le T_D \le t + dt \,|\, T_D \ge t\,\right) \\
&=\; \frac{\Pr\left(\,t \le T_D \le t + dt \text{ and } T_D \ge t\,\right)}{\Pr\left(T_D \ge t\right)} \\
&=\; \frac{-Q'(t)\,dt}{Q(t)}\;.
\end{aligned}$$

This gives the differential equation $Q'(t) = -\lambda Q(t)$. Together with the initial condition $Q(0) = \Pr(T_D \ge 0) = 1$, it leads to $Q(t) = e^{-\lambda t}$, and $f(t) = \lambda e^{-\lambda t}$. That is:

$$\Pr\left(\,t \le T_D \le t + dt\,\right) \;=\; e^{-\lambda t}\,\lambda\,dt\;. \tag{3}$$

The right side has a simple interpretation. The first factor, $e^{-\lambda t}$ is the survival probability. The second term is the probability that a surviving bond defaults soon after t. We multiply the two probabilities because the default is a Markov process.

We pause for a moment to check the units of these formulas. Any time you take the exponential or log or sine or cosine of something, that something must be dimensionless - a "pure number" without units. The default intensity, λ, measures number of defaults per unit time. Its units are 1/Time. Thus, the exponent, λt, is dimensionless, as it should be. The probability density is supposed to have units 1/Time, because you multiply it by dt to get an actual

probability. A probability is a number between zero and one, which makes sense only if it is dimensionless. Our formula $f(t) = \lambda e^{-\lambda t}$ gives f the correct units.

We calculate the expected total discounted payout of an infinite maturity bond in this simple default model. The cash flow from the time interval $(t, t+dt)$ is $C \, dt$. Its present discounted value, as before, is $e^{-rt} C \, dt$. The probability of receiving that cash is $e^{-\lambda t}$. Altogether, the present expected value of the potential cash flow from future interval $(t, t + dt)$ is $e^{-\lambda t} e^{-rt} C \, dt$. For the moment, neglect the recovery payout R that would happen at time T_D. The present value of all these possible cash flows is

$$ C \int_0^\infty e^{-(r+\lambda)t} \, dt \;=\; \frac{C}{r + \lambda} \,. $$

We set this to one by taking

$$ C \;=\; r + \lambda \,. \tag{4} $$

In this simplest of simple models, the coupon, C, is also the rate of return on the bond. Since r is the yield on the comparable risk free asset, (4) shows that the credit yield spread is exactly equal to the default intensity, in this simplest of simple models.

We can include a non-zero recovery into the value computations above. The present value of payment R at time T_D is Re^{-rT_D}. The probability that this payment is made in the time interval $(t, t + dt)$ is, by (3), $\lambda e^{-\lambda t} \, dt$. Therefore, the present expected value of this possible recovery payment is

$$ R \int_0^\infty e^{-rt} \lambda e^{-\lambda t} \, dt \;=\; \frac{\lambda R}{r + \lambda} \,. $$

With recovery, the total (discounted expected) value is the sum of the value of the coupons and the recovery, which is

$$ \frac{C + \lambda R}{r + \lambda} \,. $$

Setting this equal to 1 gives the coupon for which the bond is valued at par:

$$ C \;=\; r + \lambda(1 - R) \,. $$

This gives the default intensity in terms of the yield spread, s, as

$$ \lambda \;=\; \frac{s}{1 - R} \,. $$

For the same yield spread, the default intensity increases as the recovery increases, for a given yield spread.

Computations of the same kind can be done in discrete time. You can find them in the two chapters of Hull on credit, and also in the lecture notes of Steve Allen, Section 11. The results are more complicated, but the qualitative behavior is similar.

If you had the yield spreads for a collection of bonds from the same firm with different maturities, you could use a procedure like bootstrapping to get a term structure for the default intensity, $\lambda(t)$. Analogous to the derivation of (3), we have

$$\partial_t Q(t) \;=\; -\lambda(t)Q(t)\,, \quad Q(0) = 1\,, \quad \Pr\left(\,t \leq T_D \,\leq\, t + dt\,\right) \;=\; \lambda(t)Q(t)\,dt\,.$$

3 Risk neutral probabilities

Historical default intensities – the number of bonds of a given rating that default per year – are far below default intensities implicit in market yield spreads. (This is even more true today than a year ago. Yield spreads are many times their historical levels.) Both Allen and Hull discuss the common view that this represents risk aversion on the part of people buying bonds.

The default intensity parameters, λ, estimated from yield spreads (more properly, probabilities such as (3)) are risk neutral probabilities. Recall the definition: risk neutral probabilities are those probabilities that give market prices as expected discounted cash flows. Therefore, *by definition*, parameters, λ, estimated from market yield spreads are risk neutral default intensities. The recovery fraction, R, also is the risk neutral expected recovery rate upon default. What makes this not a giant tautology is that it is useful to model λ or $\lambda(t)$ to get an idea what yield spreads might be for bonds without explicit market quotes. For example, a firm might be trying to mark to market thinly traded bonds or planning a new issue.

One also could imagine a more sophisticated model of the bond default process. For example, Risk Metrics treats changes in ratings of a bond as a Markov chain process and estimates upgrade and downgrade probabilities. These are also in the form of intensities: $\lambda_{AA \to A}$ would be the probability per unit time that a bond rated AA would be downgraded to A. In this model, a bond initially rated highly would have to drift down the ratings before it because a significant default risk. It's default probability per unit time would increase with time appropriately. Of course, these would be "historical" (i.e. actual) transition probabilities rather than risk neutral probabilities. Presumably, one could estimate risk neutral upgrade and downgrade intensities from credit spread or other market price data.

My final remark concerns bond default correlation. This is a major reason the market for defaultable bonds is so risk averse. If bond defaults for different firms were independent events, diversification would produce a return based on historical rather than risk neutral default rates. If defaults were independent and one owned bonds from many companies, the law of large numbers would imply that the actual number of defaults (divided by the number of names) would be approximately given by historical default probabilities. These, in turn, could be estimated from historical default data.

This estimation/diversification idea fails in the real world because of default correlation. When times are good, nobody defaults. When times are bad,

everyone does. For this reason, the realized return on a diversified basket of bonds is still a random variable with very large uncertainty. Moreover, and this point goes beyond Hull or Allen, it also makes it difficult to estimate historical default probabilities.

In the extreme case, suppose there is one single time, T_D, at which every single bond defaults. Suppose this time is an exponential random variable specified by a constant default intensity, λ. If this default intensity were equal to a 2% yield spread, the expected time between massive bond defaults would be fifty years. There would be a reasonable chance not observe such an event between the crash of 1929 and now. Estimating λ from only a hundred years of historical data would be almost impossible.

Given recent events (written November 20, 2008), default intensities of 2%/year above historical rates do not seem so unreasonable. In fact, the actual "historical" bond default intensity for a typical bond probably is much higher than one would estimate from data that excludes the past few months.