Fast and accurate simulation of differently seasoned loan defaults in a Merton-style framework in discrete time

Varsanyi, Zoltan

August 2008

Fast and accurate simulation of differently seasoned loan defaults in a Merton-style framework in discrete time

Zoltan Varsanyi*

August, 2008

In this paper I present a method for the simulation of the default of such loans that have two important properties: they are seasoned – maybe even being at different points of the seasoning curve – and they evolve in an asset-value based framework. This latter model allows us to introduce correlation between the loan defaults. Although these two features are widely considered in modelling, linking them into one single (simulation) framework might not be that common. However, the most important merit of this paper is showing a fast and accurate simulation algorithm for the asset values.

The Merton-framework is probably the most widespread way to credit risk modelling. Its importance is shown, for example, by its application in risk rating (e.g. KMV) and risk regulation (the Basel II framework).

The model's basic idea is to let the asset value of a firm evolve according to a stochastic process (random walk) and as soon as this asset value reaches a lower trigger value (the "default trigger") consider that firm being in default status. There are two important variations: in the first the firm defaults whenever its asset value reaches the trigger; alternatively, the firm defaults if *at the end of a given time horizon* its asset value is below the trigger. The Basel II regulation is based on the second approach; in another paper (Varsanyi [2006]) I examine whether the two approaches (the second as applied under Basel II) are really equivalent.

When one tries to simulate from this model one faces a problem. In the above first approach (when defaults are triggered whenever the default trigger is hit anytime during the period of the analysis) a timescale that is very densely broken up into small time intervals has to be applied. For example assuming that the length of the full analytical period equals 1, in order to get accurate results sub-periods of a length of, say, one-hundredth or even smaller are advisable. Of course, the longer the full period the higher number of sub-periods we should have. Since then the asset value is simulated over a high number of sub-periods it can take a very considerable amount of time to carry out the analysis. If, in turn, we apply a lower number (and hence longer) sub-periods, we will get inaccurate estimates – to be more specific, we will underestimate the default probability (since defaults are only checked at the end of each sub-period we lose all the default-trigger hits which occurred between two such end-points but given that the asset value came back above the trigger level by the end of the sub-period).

---

* varsanyiz77@yahoo.com

The second approach doesn't require the full period to be broken up. On the one hand it makes simulation much quicker; however, it doesn't allow the analysis of the timely evolution of the asset value during the period: we only look at the asset value at the end of the full horizon So, for example, looking at a period of one year, we can say nothing about the evolution of defaults during the year – we can only tell whether the asset defaulted at the end of the year or not.

Fortunately, there is a very simple solution to the problem. A solution whereby we only have to simulate a relatively small number of periods (in the vast majority of real life problems it should be enough, for example, to simulate a loan default on a monthly basis) and still we get fully accurate results. The idea is as follows. Let's take a period, say, one year. Our purpose is to simulate defaults over the period using a monthly time-scale. First, let's take the seasoning curve of the asset (over one year – this could be substituted by any length):

— cumulative defaults over time

The curve shows, for example, that up until (and including) month 4 the probability that the asset defaults is about 6 percent.

Now, a very easy way of modelling the default of an asset over consecutive months with the above default curve is to generate a standard normal random variable for each month and check whether the value of these variables is below the default trigger. The default triggers are in turn given by applying the inverse of the standard normal distribution over the monthly increments of the default curve. For example, if the probability of being in default increases from month 4 to month 5 from 6.1 percent to 6.8 percent, then in month 5 we would compare the value of that month's random variable with the inverse of 0.7 percent (6.8-6.1), which is -2.46: if the random variable is below -2.46 then we say that the asset defaulted in that month (given that it hasn't defaulted earlier). It is easy to see why we get the correct results: the default outcomes for the different sub-periods are independent, thus their expected values can be summed to arrive at the full-period default

rate; while these sub-period default rates have expected values implied by the seasoning curve.

The approach can be extended to simulate correlated defaults of assets having different seasoning curves. Applying the asset-value framework we can write the asset value (R_t) depending on a factor (X_t) and an idiosyncratic shock (e_t) as:

$$R_t = wX_t + \sqrt{1-w^2}\, e_t.$$

Assuming both X and e are standard normal and they are independent R will also be standard normal. In the above equation w can be called the asset correlation which links the asset value to the factor.

If several assets depend on the same factor it will create correlation among their asset values. In this case the simulation using the seasoning curves would go as follows. We simulate two standard normal variables for each asset and for each month. The first variable is for the factor the other is for the idiosyncratic shock. We calculate the asset value using the above equation. We know that this asset value is standard normal – so we can compare it to the default trigger calculated as described above, for each asset. We start by month one, then if the trigger is not hit then we move on to month 2, and so on. Of course, in this case we can't speak about "asset value" as before, since R is not a cumulative process. Nevertheless, by this procedure we can correctly simulate correlated defaults which also obey the seasoning curve.

As a demonstration I take two exposures. The two seasoning curves are shown in the figure below:

I examine a high asset correlation case and a low one. In the previous one w equals 0.95 for both loans, in the latter one it is 0.2. I will run each scenario 3000 times and check 1.

the average default probability at the end of the 12 sub-periods (should be 10% and 3% for loan 1 and loan 2, respectively) and 2. the timely arrival of defaults.

In the high asset correlation case, as regards the average default probability I get 10.43% for loan 1 and 3.13% for loan 2 in one set of 3000 runs. Regarding the seasoning curves, the theoretical curves and their empirical counterparts are shown below:

In case of both the average default rate and the approximation of the theoretical curves we cannot expect 100% accuracy due to sampling error. A simple measure of dependence between the two loans can be the number of periods when the two loans defaulted simultaneously. In this high correlation case it is 65, while the total number of loan 1 and loan 2 defaults is 313 and 94, respectively.

In the low asset correlation case I got 10.2% and 2.5% for the respective averages, and also the empirical seasoning curves gave a good approximation of the theoretical ones. According to the simple measure of the dependence between loan 1 and loan 2 there were 11 cases (out of 306 of loan 1 and 75 of loan 2 defaults) when the two loans defaulted together.

Reference:

Varsanyi, Z. (2006): "**The Basel II IRB approach revisited: do we use the correct model?**", MPRA Paper 1244, University Library of Munich, Germany.