A Hidden Markov Model of Default Interaction

Martin Crowder, Mark Davis and Giacomo Giampieri

Department of Mathematics, Imperial College, London SW7 2AZ

Abstract. The occurrence of defaults within a bond portfolio is modeled as a simple hidden Markov process. The hidden variable represents the risk state, which is assumed to be common to all bonds within one particular sector and region. After describing the model and recalling the basic properties of hidden Markov chains, we show how to apply the model to a simulated sequence of default events. Finally, we consider a real scenario, with default events taken from a large database provided by Standard & Poor's. We are able to obtain estimates for the model parameters, and also to reconstruct the most likely sequence of the risk state.

1 Introduction

Interaction effects are a key component of portfolio credit risk, but how to quantify these effects in a credible way has been the subject of some controversy. For large portfolios of, say, $n = 50$ bonds, it is generally infeasible to model the default risk of each individual issuer and the 'correlation' (however defined) with other issuers, since this leads to a high-dimensional model with an enormous number of parameters, which cannot be reliably estimated. Instead, one is looking for a simple description of the interaction process, justifiable on economic and/or empirical grounds, that is characterized by a small number of parameters. A number of such models have been proposed and some of them are in widespread industrial use. A good example is Moody's 'binomial expansion techniques' (BET) (see [4] or §11.3 of Duffie and Singleton [3]). In the BET, the original portfolio of size n is replaced by a smaller portfolio of size $n' < n$, the members of which are supposed to default independently, leading to a binomial distribution for the number of defaults over a fixed time horizon. The number n' is determined by a 'diversity score' analysis in which issuers in different industry sectors are deemed independent while those in the same sector are coupled in a quantified way. The bottom line is that portfolios with low diversity have greater tail risk. A similar effect was obtained by Davis and Lo [1] in an infection model which assumes that a defaulting bond may trigger off defaults in other bonds. The model only has two parameters, an individual default probability p and an infection parameter q. As the latter is increased, default distributions very similar to the Moody's model are produced.

These models are *static* in that they only concern the total number of defaults in a specified period. For applications such as CDOs (collateralized debt obligations) the *timing* of defaults is as important as the total number, and one needs a dynamic – i.e. stochastic process – model. In [2], Davis and Lo define the so-called *enhanced risk* model as a dynamic version of infectious defaults. The portfolio is assumed to be in one of two states: normal risk and enhanced risk. It starts in normal risk, but as soon as a default occurs it moves to enhanced risk, where the hazard rates for all remaining issuers are multiplied by an enhancement factor $\kappa > 1$. The portfolio stays in the enhanced risk state for an exponentialy-distributed random time before dropping back to normal risk. The two states can be thought of as a general 'good times/bad times' economic variable. This interpretation is the one examined below in this paper. A somewhat similar approach, in which default of 'primary issuers' affects the hazard rate of 'secondary issuers', has been taken by Jarrow and Yu [5].

A very common idea, which one sees in, for example, the Credit Metrics model [8] or the CDO model in §11.3 of [3], is to suppose that each issuer is exposed to three default risk factors: a business cycle factor affecting all issuers, an industry-specific factor affecting only firms in the same industry sector, and an idiosyncratic factor specific to the issuer itself. One of the purposes of this research is to determine whether these factors are supported by the data.

Turning to the work described below, we consider a simplified enhanced risk model along the lines of Davis and Lo [2] with two, not directly observed, states corresponding to normal and enhanced risk. There is no 'infection' effect: we suppose that the hidden variable is a two-state Markov process in discrete time (time is quantised into intervals of 90 days), not depending on the default events. Within each time period defaults are supposed to be binomially distributed, with higher mean in the enhanced risk state. This is a 'hidden Markov model', for which theory and estimation algorithms are available [6],[7]. We estimate model parameters and most likely paths for the hidden state using a large database of default histories made available to us by Standard & Poor's. Section 2 below describes the model, while Section 3 describes the estimation techniques and gives parameter estimates and most likely hidden Markov process sample paths for for data drawn from four different industry sectors. Error estimates for the parameters can be obtained by a bootstrapping technique described in Section 4. As a further diagnostic test, we investigate in Section 5 whether the model prediction of binomial default distribution within each risk category is supported by the data. Section 6 is a preliminary examination of the influence of global versus industry-specific economic factors. We now aggregate the data from all the sectors previously considered and re-estimate to find the most likely path for a global economic factor. It seems that the global factor probably accounts for most of the interaction between issuers, but that one could possibly distinguish secondary effects related to industry sectors. Concluding remarks and suggestions for further work are given in the final section, Section 7.

2 Description of the model

A discrete state, discrete time hidden Markov model (HMM) consists of a set of n measurements (nodes), each of which is associated with a set of m possible observations. The parameters of the model include an initial state π which describes the distribution over the initial node, a transition matrix a_{ij} for the transition probability from node i to node j conditional on node i, and an observation matrix $b_i(m)$ for the probability of observing m conditional on node i.

More specifically, in our case the hidden state is associated to the risk state, which can take two values: 0 (normal risk), and 1 (enhanced risk). For simplicity, we assume the initial state has equal probability of being in one of the two possible states. In the normal risk state, the number m of observed defaults in each time step is distributed binomially, with parameter λ:

$$p_0(m) = \binom{N_s}{m} \lambda^m (1-\lambda)^{(N_s-m)} \tag{1}$$

where N_s is the number of surviving bonds. In the enhanced risk state, the p.d.f. $p_1(m)$ is still given by eq. (1), after multiplying λ by a factor $\kappa \geq 1$. In the limiting case $\kappa = 1$, the two hidden states become equivalent, i.e., $p_0(m) = p_1(m)$. The transition matrix a_{ij} is assumed to be constant, i.e.,

$$a_{ij} = \begin{pmatrix} q & 1-q \\ 1-p & p \end{pmatrix} \tag{2}$$

where q is the probability of remaining in the normal risk state, and p is the probability of remaining in the enhanced risk state. Thus, our model is fully described by four parameters: λ, κ, q, p. Note that, although a_{ij} is time independent, the model turns out to be time dependent, since the number of defaults in each time step depends on the number of surviving bonds at the beginning of the period (see eq. (1)).

3 Estimation of parameters

Given the model and the observation sequence, the model parameters can be estimated with standard Maximum Likelihood techniques. In particular, an algorithm developed by Baum and Welch for signal processing applications (see Rabiner [7]) can be applied. Implementation of the Baum-Welch algorithm involves computation of two different probability terms. First, the forward path probability $\alpha_t(i) = P(m_1 m_2 \ldots m_t, i)$ is defined as the joint probability of having generated a partial observation sequence in the forward direction (i.e., from the start of the sequence) and having arrived at a certain state i at time t. Next, the backward path probability $\beta_t(i) = P(m_{t+1} \ldots m_T | i)$ denotes the probability of generating a partial observation sequence in the reverse direction (from the final time T), given that the state sequence starts from a certain state i at time t. In addition, the probability $\gamma_t(i) = P(i|\mathcal{M})$ of being in a given state i at time t, given the whole observation sequence $\mathcal{M} = m_1, \ldots, m_T$, can be expressed in terms of the forward-backward variables as

$$\gamma_t(i) = \frac{\alpha_t(i)\beta_t(i)}{\sum_{j=0}^{1} \alpha_t(j)\beta_t(j)} \tag{3}$$

The parameters of the HMMs can be estimated by using these probabilities. Full details of the Baum-Welch procedure for parameter estimation, as well as the various implementation issues, are described in [7].

3.1 Simulation

First, we apply the Baum-Welch algorithm for estimating the parameters of the hidden Markov model, using a simulated data set. We have started with N=1000 bonds, and simulated a sequence of defaults events over a time period of 20 years, using a time step of 90 days. The parameters chosen in this simulation are $\lambda = 0.004, \kappa = 5, q = p = 0.9$, meaning that in each of the 90-day periods, 0.4% of surviving bonds will, on average, default, and five times more if we are in the enhanced risk state. The probability of jumping from one risk state to the other is 10% during each time step. We stress that the Baum-Welch algorithm only finds local maxima of the probability function, so that the starting model has to be chosen with care. In our case, we start from the initial guess $\lambda = 0.001, \kappa = 2, q = p = 0.5$, and after only five iterations we find $\lambda = 0.0038, \kappa = 4.6727, q = 0.9075, p = 0.9043$, in excellent agreement with the true values. The simulated time sequence, along with the true and estimated hidden sequence, is shown in Fig. 1.

As one may expect, the number of iterations increases, and the final accuracy decreases, as κ gets smaller. For example, decreasing κ from 5 to 3 (keeping all other parameters the same) produces the following output: $\lambda = 0.0046, \kappa = 2.5733, q = 0.8353, p = 0.8740$. We have also verified that in the degenerate case $\kappa = 1$, the transition probabilities q and p become meaningless.

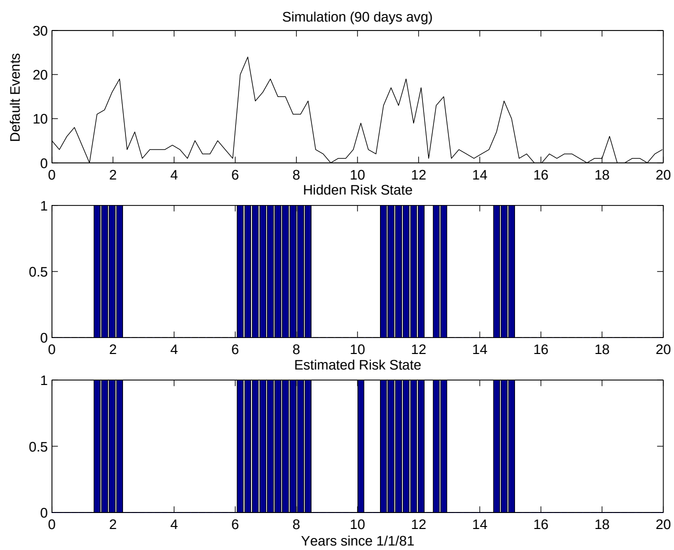

Figure 1: Simulated data. The top graph gives the number of defaults as function of time. The solid bars in the middle plot indicate the enhanced risk periods. The bottom plot shows the reconstruction of the risk state, which agrees very well with the true distribution.

3.2 Results from S&P's database

We now consider the actual time sequence of default events, as provided by Standard & Poor's. In particular, we have extracted the default times of bonds in the following US sectors: automotive, consumer, energy, and media. The S&P database provides also the time when each bond was first rated. These times define a deterministic birth process, which we have included in our HMM model. We have grouped the default events in 90-day periods, and applied the estimation algorithm to the resulting sequence.

Table 1 shows, for each sector, the total number of bonds issued over the whole period, the number of defaulted bonds, and the parameters obtained from the forward-backward procedure.

Figures 2-5 show the time sequence of default events, along with the estimated risk-state sequence, for each of the four sectors considered. Note how the risk state is correlated among different sectors, especially in the few final years, which raises the possibility of investigating a cross-sectors infection effect. This issue will be briefly addressed in the final paragraph.

Table 1: Results from S&P's database

Sector	N_{tot}	N_{def}	λ	κ	q	p
automotive	820	167	0.00121	2.77610	0.91367	0.96549
consumer	1041	251	0.00180	7.38824	0.89472	0.76040
energy	420	71	0.00123	8.25494	0.92074	0.84073
media	650	133	0.00269	6.75460	0.95939	0.82885

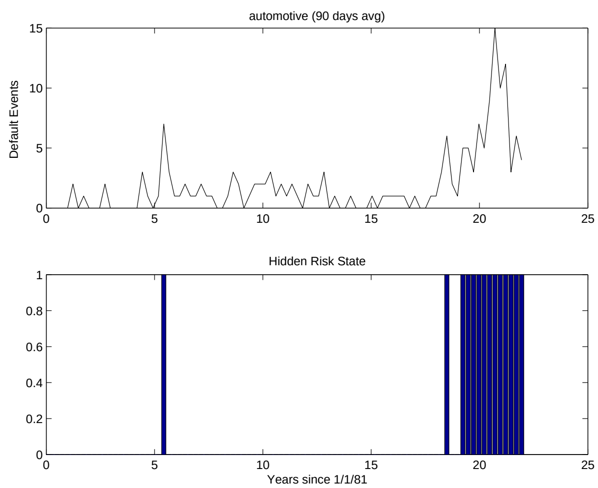

Figure 2: automotive

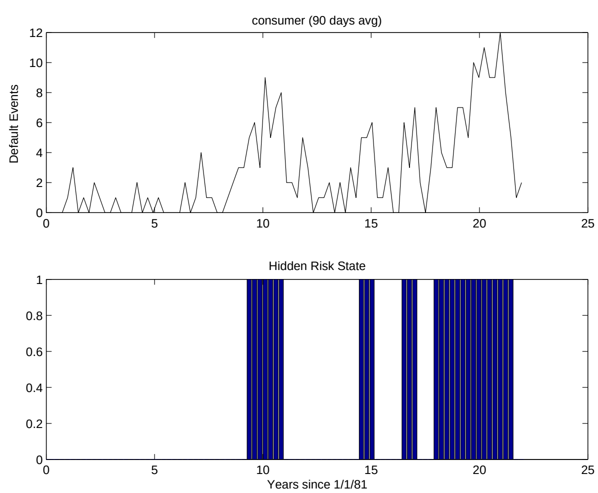

Figure 3: consumer

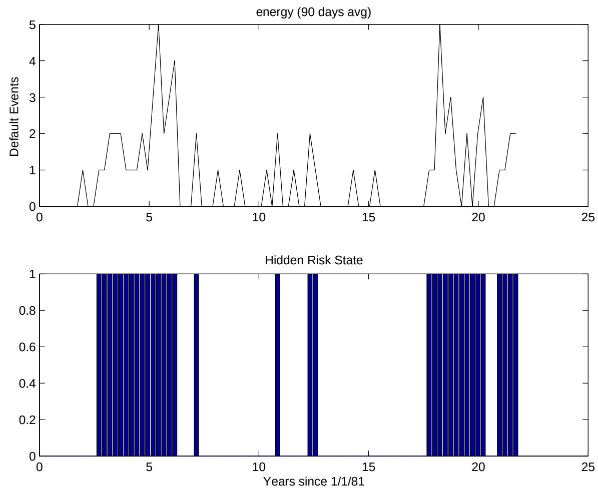

Figure 4: energy

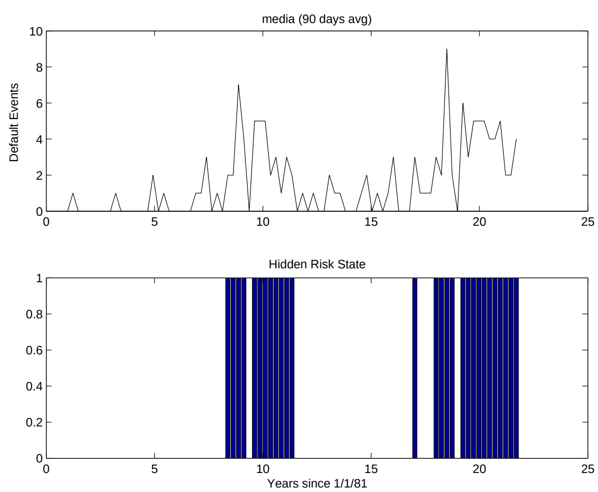

Figure 5: media

4 Parametric Bootstrap

We now show how to estimate the covariance matrix for the four parameters found with the Maximum Likelihood estimator. This can be done via a parametric bootstrap technique. For simplicity, we focus our attention to the one of the four cases considered in Table 1, namely to the US consumer sector. The bootstrap technique consists in simulating a large (e.g., $N = 50$) number of realization of the fitted HMM, that is, the HMM with the parameters shows in the second line of Table 1. For each generated realization, we estimate the four parameters with the same Baum-Welch algorithm used above. Each result is stored in a vector θ_i, and the covariance matrix estimator is then given by

$$\mathbf{C} = \frac{1}{N-1} \sum_{i=1}^{N} \left(\theta_i - \hat{\theta}\right)' \cdot \left(\theta_i - \hat{\theta}\right) \tag{4}$$

where

$$\hat{\theta} \equiv \frac{1}{N} \sum_{i=1}^{N} \theta_i \tag{5}$$

For the US consumer case we find:

$$\mathbf{C} = \begin{pmatrix} 0.000000102 & \text{-}0.000677963 & 0.000004177 & 0.000013423 \\ \text{-}0.000677963 & 8.169546678 & \text{-}0.008282575 & \text{-}0.196507296 \\ 0.000004177 & \text{-}0.008282575 & 0.001610009 & 0.001228519 \\ 0.000013423 & \text{-}0.196507296 & 0.001228519 & 0.010093449 \end{pmatrix} \tag{6}$$

In addition, we can determine the distribution of the parameters θ_i around the mean $\hat{\theta}$. Fig. 6 shows the four histograms corresponding to each of the four estimated parameters. We can see that the distribution of λ, q and p is reasonably close to being normal (however, note that q and p, being bounded between 0 and 1, cannot be, strictly speaking, normally distributed). The fourth parameter, κ, has a flatter distribution, skewed towards higher values.

5 Test of the Binomial distribution

As explained in §2, our model assumes that, at each time step, the number m of defaults is distributed according to the binomial p.d.f. $p_0(m)$ or $p_1(m)$, depending on the risk state. We can now check whether this is verified in practice. In doing so, we must remember that the parameters of the binomial distribution (for each of the two risk states) are not constant, because of the fact that N_s varies with time. Thus, when we look at the distribution of the number of defaults, conditioned on being in risk state 0 and 1 respectively, we do not expect to find a pure binomial curve. We have considered the enhanced risk periods for the same four US sectors considered in §3.2 (solid bars in figs. 2-5). Since N_s is far from being constant, one can expect a deviation from a purely binomial curve. The distributions found using the actual data (fig. 7) are in reasonably good agreement with the expected ones, indicating that the binomial distribution is indeed a viable assumption.

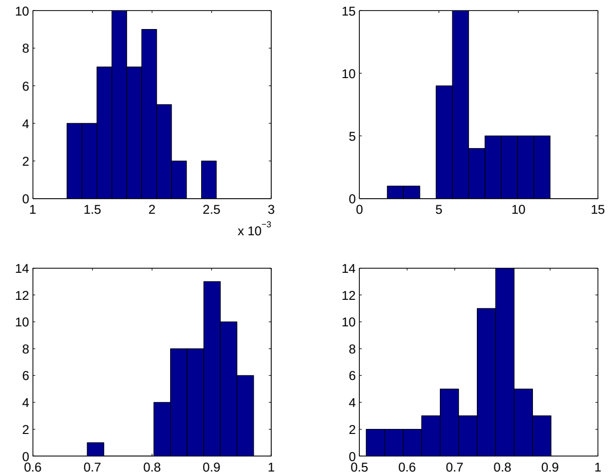

Figure 6: Frequency distribution of estimated parameters produced by the parametric bootstrap technique.

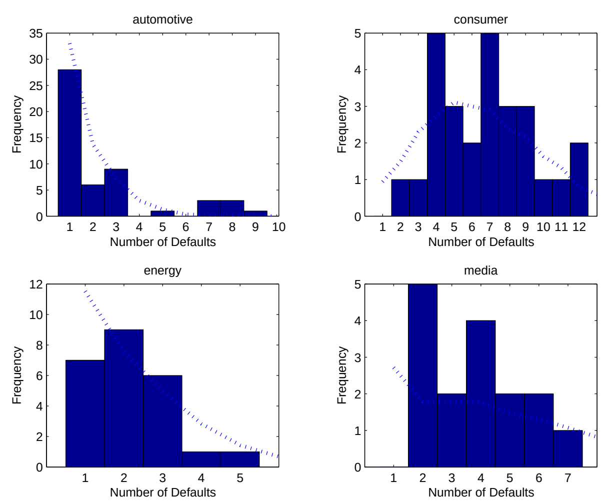

Figure 7: Frequency distribution of defaults during the enhanced risk periods in the four sectors: automotive, consumer, energy, and media. The dotted line in each plot is the expected curve for a HMM, with the estimated parameters obtained from Table 1.

6 Global effects

The main motivation behind our work is to model the risk factor affecting all individual firms within a specific sector. However, in some cases, the enhanced risk could be related to global economic factors, affecting all sectors and all regions at the same time, as noted at the end of §3.2. We have thus applied our model to the whole database, without distinction for industry type, country, etc. The total number of bonds considered is 9928, of which 1238 defaulted during the considered period. The Baum-Welch procedure, applied to the whole set of data, produced the following parameters: $\lambda = 0.0016, \kappa = 6.9012, q = 0.9262, p = 0.5951$. The defaults sequence, along with the estimated hidden state, is shown in fig. 8. Note that the presence of global enhanced risk periods, especially the one at the end of the observing time, could have easily been guessed by overlapping the analogous plots for the individual sectors (as shown, for example, in figs. 2-5). However, most of the enhanced risk periods shown in figs. 2-5 do not have a counterpart in fig. 8. Thus, in this way one may be able to disentangle, a posteriori, the default risks associated with the global economy from those specific to a particular sector.

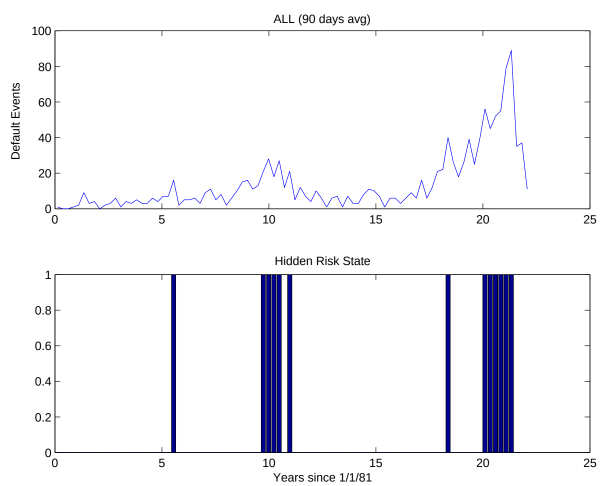

Figure 8: Same as figs. 2-5, with all available data.

7 Concluding Remarks

The model we have introduced is certainly very simple, but our empirical analysis shows that it has good explanatory power. We have already mentioned several extensions that could be pursued, for example explicit inclusion of separate industry-specific and general economic variables. One could also consider a hidden process with more than two states, though an interpretation of this might be problematic.

One area that we have so far ignored entirely is changes of rating. We have only used information about realized default events, but in fact the database contains far more. A more sophisticated model might posit interaction effects in rating changes as well, something more in the spirit of the Credit Metrics approach [8].

Acknowledgements.

We are very grateful to Standard & Poor's, and particularly to our colleagues there Arnaud de Servigny and Olivier Renault, for allowing us access to S&P's database of default history. We also thank Jonathan Staples for helpful discussions. The work of Giacomo Giampieri was supported by the UK EPSRC under Grant GR/R80131/01.

References

[1] M. Davis and V. Lo, *Infectious defaults*, Quantitative Finance **1** (2001) 382-386

[2] M. Davis and V. Lo, *Modelling default correlation in bond portfolios*, in *Mastering Risk Volume 2: Applications*, ed. Carol Alexander, Financial Times Prentice Hall 2001, pp 141-151

[3] D. Duffie and K. J. Singleton, *Credit Risk: Pricing, Measurement and Management*, Princeton University Press 2003

[4] Moody's KMV Company, *The Binomial Expansion Technique*, Technical document 1997 `http://www.moodyskmv.com`

[5] R. Jarrow and F. Yu, *Counterparty risk and the pricing of defaultable securities*, J. Finance **56** (2001) 1765-1799

[6] I.L. MacDonald and W. Zucchini, *Hidden Markov and Other Models for Discrete-valued Time Series*, Chapman and Hall/CRC, Boca Raton 1997

[7] L.R. Rabiner, *A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition*, Proc. IEEE, **77**, 257-286.

[8] Risk Metrics Group, *Credit Metrics Technical Document*, `http://www.riskmetrics.com/cmtdovv.html`