

The Bayesian Approach to Default Risk: A Guide

Michael Jacobs Jr.¹

Nicholas M. Kiefer²

Office of the Comptroller of the Currency and Cornell University³

March 3, 2010

¹US Department of the Treasury, Office of the Comptroller of the Currency, Credit Risk Analysis Division.

²Cornell University, Departments of Economics and Statistical Science, 490 Uris Hall, Ithaca, NY 14853-7601, US. email:nicholas.kiefer@cornell.edu; US Department of the Treasury, Office of the Comptroller of the Currency, and CREATES, funded by the Danish Science Foundation, University of Aarhus, Denmark.

³*Disclaimer: The statements made and views expressed herein are solely those of the author, and do not necessarily represent official policies, statements or views of the Office of the Comptroller of the Currency or its staff.*

Abstract

A Bayesian approach to default rate estimation is proposed and illustrated using a prior distribution assessed from an experienced industry expert. The principle advantage of the Bayesian approach is the potential for coherent incorporation of expert information - crucial when data are scarce or unreliable. A secondary advantage is access to efficient computational methods such as Markov Chain Monte Carlo. After a preliminary discussion of elicitation of expert beliefs, all steps in a thorough Bayesian analysis of a default rate are illustrated. Using annual default rate data from Moody's (1999-2009) and a prior elicited from an industry expert, we estimate three structural credit models in the asymptotic single risk factor (ASRF) class underlying the Basel II framework (Generalized Linear and Generalized Linear Mixed Models), using a Markov Chain Monte Carlo technique.

Keywords: Basel II, risk management, prior elicitation, maximum entropy, MCMC

1 Introduction

All competent statistical analyses involve a subjective or judgmental component. Sometimes the importance of this input is minimized in a quest for objectivity. Nevertheless, it is clear that specification of a model, definition of parameters or quantities of interest, specification of the parameter space, identification of relevant data, all require judgment and are subject to criticism and require justification. Indeed this justification is an important part of the validation procedure expected of financial institutions OCC (2000). However, estimation of parameters after these judgments are made typically proceeds without regard for potential nondata information about the parameters, again in an attempt to appear completely objective. But subject-matter experts typically have information about parameter values, as well as about model specification, etc. For example, a default rate should lie between zero and one (definition of the parameter space), but if we are considering a default rate for a particular portfolio bucket, we in fact have a better idea of the location of the rate. The Bayesian approach allows formal incorporation of this information, formal combination of the data and nondata information using the rules of probability. A simple example in the case of estimating default rates is sketched in Kiefer (2007).

The Bayesian approach is most powerful and useful when used to combine data and nondata information. There is also an advantage in that powerful computational techniques such as Markov Chain Monte Carlo (MCMC) and related techniques are available. These are widely discussed in the economics literature and have been applied in the default estimation setting. These applications invariably specify a "prior" which is convenient and adds minimal information - there is no such thing as an uninformative prior - allowing computationally efficient data analy-

sis. This approach, while valuable, misses the true power of the Bayesian approach, the coherent incorporation of expert information.

The difficulty in Bayesian analysis is the elicitation and representation of expert information in the form of a probability distribution. This requires thought and effort, rather than mere computational power, and is therefore not commonly done. Further, in "large" samples data information will typically overwhelm nondogmatic prior information, so the prior is irrelevant asymptotically, and economists often justify ignoring prior information on this basis. However there are many settings in which expert information is extremely valuable. In particular, cases in which data may be scarce, costly, or of questionable reliability. These issues come up in default estimation, where data may not be available in quantity for low-default assets or for new products, or where structural economic changes may raise doubts about the relevance of historical data.

We go through the steps in a Bayesian analysis of a default rate. Estimation of long-run default rates for groups of homogeneous assets is essential for determining adequate capital. The Basel II (B2) framework (Basel Committee on Banking Supervision (2006)) for calculating minimum regulatory capital requirements provides for banks to use models to assess credit (and other) risks. In response to the credit crisis, the Basel Committee has stressed in a document for comment the continuing importance of quantitative risk management, see Basel Committee on Banking Supervision (2009a). Our emphasis is on the incorporation of nondata information, so we focus on elicitation and representation of expert information and then on the Bayesian approach to inference in the context of a simple model of defaults. Uncertainty about the default probability should be modeled the same way as uncertainty about defaults – represented in a probability distribution. A future default either occurs or doesn't, given the definition. Since we do not know

in advance whether default occurs or not, we model this uncertain event with a probability distribution. Similarly, the default probability is unknown. But there is information available about the default rate in addition to the data information. The simple fact that loans are made shows that some risk assessment is occurring. This information should be organized and incorporated in the analysis in a sensible way, specifically represented in a probability distribution. We discuss elicitation generally in section 2. Then we run through the steps of a formal Bayesian analysis in a particular example. Definition of a parameter of interest requires a model, so we turn to specification of a sequence of simple models (each generating a likelihood function) in section 3. The first two models are consistent with the asymptotic single-factor model underlying B2. The third adds temporal correlation in asset values, generalizing B2 and perhaps in line with the validation expectations of Basel Committee on Banking Supervision (2009b), in that accounting for autocorrelation in the systematic factor speaks to how a model should be forward looking. Section 4 goes through the actual elicitation and representation for a particular expert. Section 5 sketches the MCMC approach to calculating the posterior distribution (the distribution resulting from combining data and expert information coherently using the rules of probability) and illustrates all of the steps using Moody's data on corporate defaults. Section 6 is the conclusion.

To summarize: The steps in a Bayesian analysis are 1) Specify an economic model precisely defining the quantity of interest and generating a likelihood function for well-defined data (section 3). 2) Identify a qualified expert and elicit information about the quantity of interest (section 4). 3) Quantify this information in a probability distribution (section 4). 4) Use the rules of probability to combine coherently the likelihood and the prior distribution, generating the posterior distribution of the quantity of interest (section 5). 5) Analyze the posterior distribution

using MCMC (section 5).

2 Elicitation of Expert Information

A general definition of elicitation that we may offer in this context is a structured process or algorithm for transforming an expert's beliefs regarding an uncertain phenomenon into a probability distribution. In deploying Bayesian statistical technology in the domain of credit risk, elicitation arises as a method for specifying a prior distribution for one or more unknown parameters governing a model of credit risk (i.e., a probability of default - PD), where the expert may be an experienced statistician or a non-quantitatively oriented risk specialist (e.g., a loan officer). In this setting, the prior distribution will be combined with a data likelihood through Bayes' Theorem, to derive the posterior distribution of the measure of risk (i.e., the distribution of the default rate). While our focus is on formulating a probability distribution for uncertain quantities for the purpose of inference about a parameter, especially when there are no or very limited data, we note here that this is not the only context in which elicitation is important. This situation also arises in decision making where uncertainty about "states of nature" needs to be expressed as a probability distribution in order to derive and maximize expected utility. Similarly, this methodology arises in the application of mechanistic models built in almost all areas of science and technology to describe, understand, and predict the behavior of complex physical processes. In that application a model developer will typically propose sensible model parameters in order to obtain outputs in cases where in general there is uncertainty about the inputs' true values. As in our application, this highlights the importance of having a coherent approach to represent that uncertainty.

A useful way to frame the elicitation is to identify the model developer or econo-

metrician as a facilitator, who helps the expert transform the “soft data” (i.e., experience or opinion) into a form amenable to statistical inference, which is the process of crafting expert’s knowledge into probabilistic form. Elicitation is a complex process that, if it done well, calls for a facilitator to be skilled and multi-faceted, as the role of the facilitator is central to the process of elicitation. Therefore, not only is the facilitator an econometrician, but should also have knowledge of the business of making loans and issues in managing credit risk, as well as be a strong communicator.

We may be able to set criteria for the quality of an elicitation. In doing so, we might believe that a meaningful distinction exists between the quality of an expert’s knowledge on the one hand, and the accuracy with which that knowledge is translated into probabilistic form on the other. Therefore, we say that an elicitation is done well under the condition that the distribution so derived is an accurate representation of the expert’s knowledge, no matter the quality of that knowledge. A good facilitator asking probing questions may also be able to determine whether the expert really does have valuable information. We do not pursue this line here.

We may conclude that accurate elicitation of expert knowledge is by no means a straightforward task. This remains the case even if all we wish to elicit is expert’s beliefs regarding only a single of event or hypothesis, an example in credit risk being the proportion of defaulted obligors in a particular rating class (or segment) over a given horizon. Here we seek an assessment of probabilities, but it is entirely possible that either the expert may be either unfamiliar with the meaning of probabilities, or if he can think intuitively in these terms then it still may be hard for him to articulate precise statements about probabilities. Even in the case where the expert is comfortable with probabilities and their meaning, still it is challenging to accurately assess numerical probabilities with respect to a relatively rare event such

as that of default for a set of obligors, especially if they are highly rated and there is a lack of historical performance data on a portfolio of similar credits.

Let us now consider the task of eliciting a distribution for a continuous parameter θ , the proportion of customers in a given rating class defaulting. How may we proceed with this? One direct approach, impractical, involves implicitly eliciting an infinite collection of probabilities regarding this uncertain proportion (itself a probability), which we may write in terms of the distribution function for all of the possible values of θ . Note the symmetry here, as we characterize the uncertainty regarding the unknown probability governing the distribution of the default rate PD itself in terms of probabilities. However, we realize rather early on in this process that it is clearly impossible to do this, as in practice an expert can make only a finite number (and usually a rather limited number) of statements of belief about θ . It is likely that the best that we could hope for is that such statements might take the form of a small set of either individual probabilities, or a few quantiles of the distribution of θ ; or possibly this might involve other summaries of the distribution, such as modes. In the case of a joint distribution for a collection of random quantities, for example default rates in conjunction with loss severities, the elicitation task is much more complex.

Given the apparent formidable difficulties involved in the elicitation process, a reasonable observer may question if it is worth the effort to even attempt this. The answer to why this is a worthy endeavor lies in the use of elicitation as part of the business decision making. We often find that a sensible objective for elicitation is to measure salient features of the expert's opinion, so that exact details may not be of the highest relevance with respect to the decision to be reached. For example, in forming a prior for the distribution of the default rate, a general sense of where it is centered (5 bps, 1% or 10%?), and degree to which the tail is elongated,

may be enough to inform the data at hand. Note the similarity to the issue of specification of the likelihood function, where typically an infinity (for continuous data) of probabilities are specified as a function of a small number of parameters. The point is strongly made in the normal case, when the whole set of probabilities are specified as functions of a mean and variance. This can hardly be credible as an exact description of a real data distribution, but nevertheless its usefulness has been proven in countless applications. Similarly, we specify a prior distribution on the basis of a small number of elicited properties. Even for cases in which the decision is somewhat sensitive to the exact shape of the elicited distribution, it may not be the decision but another metric that is of paramount importance, for example the regulatory capital impact or the expected utility of the supervisor, which in many cases may be quite often robust to details of the expert's opinion. Another reason that supports the importance of elicitation is its use in statistical inference, in particular in the estimation of posterior distributions or predictive densities. This is a case in which elicitation promotes a careful consideration on the part of both the expert and the facilitator regarding the meaning of the parameters being elicited. This process results in two beneficial effects. First, it results in an analysis that is closer to the application, through requiring an attention to the subject of the modeling exercise; in our application, this insures that the focus is upon a set of plausible observed default rates, over a set horizon, with respect to obligors of a particular credit quality. Second, this discipline is useful in giving rise to a posterior distribution that when finally calculated, is a meaningful object. By way of illustration, this process not only produces a PD estimate that can be used in a compliance exercise, but a complete predictive distribution of the default rate that is potentially useful in other risk management contexts, such as credit decisioning, account management or portfolio stress testing.

A natural interpretation of elicitation is to conceive of it as part of the process of statistical modeling. When statisticians write down a likelihood function for an applied problem, this is nothing more than an informed (one hopes) opinion regarding a data generation process, which is conditional on a parameter set. In hierarchical frameworks, examples being random-effects models or models incorporating latent variables, we have distributions on a sub-set of parameters that are conditional on another set of parameters. Therefore, what we term “elicitation” in this context can be interpreted as nothing more than the final step in such a hierarchy, the statement of the form of the probability distribution of the highest-level parameters. This highlights that we should not lose sight of the fact that all of the usual principles of statistical modeling apply to elicitation as well.

A stylized representation of the elicitation process consists of four separate stages. First, in the setup stage, we prepare for the elicitation by identifying the expert, training the expert, and identifying what aspects of the problem to elicit. The second step, clearly the heart of the process, is to elicit specific summaries of the experts’ distributions for those aspects. This is followed by the fitting of a probability distribution to those summaries elicited in the second step. Note that in practice, there may be overlap between this and the previous phase of the process, in the sense that the choice of what data to elicit often follows from the choice of distributional form that the facilitator prefers to fit. For example, if we prefer a simple parametric distribution such as a beta to describe the prior of the PD, then a few quantiles may suffice; whereas the more data intensive choice of a non-parametric kernel density may require other pieces of information. Finally, we note that elicitation is in almost all cases an iterative process, so that the final stage is an assessment of the adequacy of the elicitation, which leaves open the possibility of an iterative return to earlier stages in order to gather more summaries from the

expert. For example, the fitted prior distribution of the PD parameter may be presented to the expert, and if the expert is not comfortable with the shape for whatever reason, we may try to gather more quantiles, re-fit and return later to make further assessments.

Thus far, we have framed the process of conducting an elicitation as that of formulating in probabilistic terms the beliefs regarding uncertainty from an expert, which we argue is the appropriate way to think about credit risk. However, in this context, who is the expert? There are two aspects here: the qualification of the expert and the basis of his information. For the first, we look at education and experience, particularly at experience in related risk-management situations. For the second, we evaluate the quality of the arguments - would they be convincing to other experts? Are they based on reasoning from particular similar portfolios or configurations of economic conditions? In practice, the choice of expert or experts must be justified. In our context the experts are not that difficult to identify: they are the individuals who are making risk-management decisions for the relevant portfolio in a successful financial institution.

In summary, we outline suggested criteria for the conduct of elicitations, in the context of formulating beliefs regarding parameters governing default risk. While some of these aspects may be ignored in an informal elicitation, they become considerations of utmost importance wherever substantive decisions or inferences may depend on the expert's knowledge, such as deriving a distribution of the default rate for either risk management or regulatory capital purposes. First, we should keep in mind that the objective is elicitation of a PD rate distribution, that represents the credit expert's current knowledge on the risk inherent in a portfolio, and it is very useful in this regard to have a summary of what that knowledge is based on (e.g., state of the credit cycle, industry condition of the obligors or average features

of the portfolio that are drivers of default risk). Second, we must be wary of any financial or personal interest that the credit expert may possess, and that any inferences or decisions that will depend on the expert's distribution so elicited should be declared upfront (e.g., if the credit executive's bonus is a function of the regulatory capital charge on his portfolio). Next, it is also paramount that training should be offered in order to familiarize the expert with the interpretation of probability, as well as whatever other concepts and properties of probability will be required in the elicitation. It may be helpful in this regard to perform a "dry-run" through an elicitation exercise with a view toward providing practice in the protocol that the facilitator proposes to use. Finally, the elicitation should be well documented. Ideally, this should set out all the questions asked by the facilitator, the expert's responses to those, and the process by which a probability distribution was fitted to those responses (e.g., details of any moment matching or smoothing performed, such as well-commented computer code). These documentation requirements for expert elicitation fit well in the the supervisory expectations for the documentation of developmental evidence.

3 Statistical Models for Defaults

Before elicitation can proceed the quantities of interest need to be defined precisely. This requires a model. The simplest probability model for defaults of assets in a homogeneous segment of a portfolio is the Binomial, in which the defaults are assumed independent across assets and over time, and occur with common probability $\theta \in [0, 1]$. The Basel requirements demand an annual default probability, estimated over a sample long enough to cover a full cycle of economic conditions. Thus the probability should be marginal with respect to external conditions. Perhaps this

marginalization can be achieved within the binomial specification by averaging over the sample period, thus many discussions of the inference issue have focussed on the binomial model and the associated frequency estimator. Suppose the value of the i th asset in time t is

$$v_{it} = \epsilon_{it}$$

where ϵ_{it} is the time and asset specific shock (idiosyncratic risk) and default occurs if $v_{it} < T^*$, a default threshold value. A mean of zero is attainable through translation without loss of generality. We assume the shock is standard normal with distribution function $\Phi(\cdot)$. Let d_i indicate whether the i th observation was a default ($d_i = 1$) or not ($d_i = 0$). The distribution of d_i is $p(d_i|\theta) = \theta^{d_i}(1 - \theta)^{1-d_i}$, where $\theta = \Phi(T^*)$, our Binomial parameter. Let $D = \{d_i, i = 1, \dots, n\}$, $n \in I^+$ denote the whole data set and $r = r(D) = \sum_i d_i$ the count of defaults. Then the joint distribution of the data is

$$\begin{aligned} p(D|\theta) &= \prod_{i=1}^n \theta^{d_i}(1 - \theta)^{1-d_i} \\ &= \theta^r(1 - \theta)^{n-r} \end{aligned} \tag{1}$$

Since this distribution depends on the data D only through r (n is regarded as fixed), the sufficiency principle implies that we can concentrate attention on the distribution of r

$$p(r|\theta) = \binom{n}{r} \theta^r (1 - \theta)^{n-r} \tag{2}$$

a Binomial(n, θ) distribution. This is Model I. Model I underlies what the ratings agencies assume in their PD estimation.

The Basel II guidance suggests there may be heterogeneity due to systematic temporal changes in asset characteristics or to changing macroeconomic conditions.

There is some evidence from other markets that default probabilities vary over the cycle. See Nickell, Perraudin, and Varotto (2000) and Das, Duffie, Kapadia, and Saita (2007). The B2 capital requirements are based on a one-factor model due to Gordy (2003) that accommodates systematic temporal variation in asset values and hence in default probabilities. This model can be used as the basis of a model that allows temporal variation in the default probabilities, and hence correlated defaults within years. The value of the i th asset in time t is modeled as

$$v_{it} = \rho^{1/2}x_t + (1 - \rho)^{1/2}\epsilon_{it} \quad (3)$$

where ϵ_{it} is the time and asset specific shock (as above) and x_t is a common time shock, inducing correlation $\rho \in [0, 1]$ across asset values within a period. The random variables x_t are assumed to be standard normal and independent of each other and of the ϵ_{it} . The overall or marginal default rate we are interested in is $\theta = \Phi(T^*)$. However, in each period the default rate θ_t depends on the systematic factor x_t . The model implies a distribution for θ_t . Specifically, the distribution of v_{it} conditional on x_t is $N(\rho^{1/2}x_t, 1 - \rho)$. Hence the period t default probability (also referred to as the conditional default probability) is

$$\theta_t = \Phi[(T^* - \rho^{1/2}x_t)/(1 - \rho)^{1/2}] \quad (4)$$

Thus for $\rho \neq 0$ there is random variation in the default probability over time. The distribution function for $A \in [0, 1]$ is given by

$$\begin{aligned} \Pr(\theta_t \leq A) &= \Pr(\Phi[(T^* - \rho^{1/2}x_t)/(1 - \rho)^{1/2}] \leq A) \\ &= \Phi[((1 - \rho)^{1/2}\Phi^{-1}[A] - \Phi^{-1}[\theta])/ \rho^{1/2}] \end{aligned} \quad (5)$$

using the standard normal distribution of x_t and $\theta = \Phi(T^*)$. Differentiating gives the density $p(\theta_t|\theta, \rho)$. This is the Vasicek distribution, see e.g. Bluhm, Overbeck, and Wagner (2003) Section 2.5, for details. The parameters are θ , the marginal or mean default probability and the asset correlation ρ . The conditional distribution of the number of defaults in each period is (from (2))

$$p(r_t|\theta_t) = \binom{n_t}{r_t} \theta_t^{r_t} (1 - \theta_t)^{n_t - r_t} \quad (6)$$

from which we obtain the distribution conditional on the underlying parameters

$$p(r_t|\theta, \rho) = \int p(r_t|\theta_t) p(\theta_t|\theta, \rho) d\theta_t$$

Since different time periods are independent, the distribution for $R = (r_1, \dots, r_T)$ is

$$p(R|\theta, \rho) = \prod_{t=1}^T p(r_t|\theta, \rho) \quad (7)$$

where we condition on (n_1, \dots, n_T) , i.e. they are considered to be known. Regarded as a function of (θ, ρ) for fixed R , (7) is the likelihood function. This is Model II.

Model II allows clumping of defaults within time periods, but not correlation across time periods. This is the next natural extension. Specifically, let the systematic risk factor x_t follow an $AR(1)$ process

$$x_t = \tau x_{t-1} + \eta_t$$

with η_t iid standard normal and $\tau \in [-1, 1]$. Now the formula for θ_t (4) still holds but the likelihood calculation is different and cannot be broken up into the period-

by-period calculation, cf. (7). Write using (6)

$$p(R|\theta_1, \dots, \theta_T) = \prod_{t=1}^T p(r_t|\theta_t(x_t, \theta, \rho))$$

emphasizing the functional dependence of θ_t on x_t as well as θ and ρ . Now we can calculate the desired unconditional distribution

$$p(R|\theta, \rho, \tau) = \int \cdots \int \prod_{t=1}^T p(r_t|\theta_t(x_t, \theta, \rho)) p(x_1, \dots, x_T|\tau) dx_1 \cdots dx_T \quad (8)$$

where $p(x_1, \dots, x_T|\tau)$ is the density of a zero-mean random variable following an $AR(1)$ process with parameter τ . Regarded as a function of (θ, ρ, τ) for fixed R , (8) is the likelihood function. This is Model III.

Model I is a very simple example of a Generalized Linear Model (GLM) (McCullagh and Nelder (1989)), see also Chapter 1, Section 2.1, in this volume. Models II and III are in the form of the General Linear Mixed Model (GLMM), a parametric mixture generalization of the popular GLM class. These models were analyzed using MCMC in the default application by McNeil and Wendin (2007) using convenience priors and focussing on default rate estimation, and by Kiefer (2009) using an elicited prior and focussing on predictability of default rates.

4 Elicitation: Example

We have asked an expert to consider a portfolio bucket consisting of loans that might be in the middle of a bank's portfolio. These are typically commercial loans to unrated companies. If rated, these might be about Moody's Ba-Baa or S&P BB-BBB. The elicitation method included a specification of the problem and some specific questions over e-mail followed by a discussion. Elicitation of prior distributions is

an area that has attracted attention. General discussions of the elicitation of prior distributions are given in Section 2 of this volume and also by Garthwaite, Kadane, and O'Hagan (2005), O'Hagan, Buck, Daneshkhah, Eiser, Garthwaite, Jenkinson, Oakley, and Rakow (2006) and Kadane and Wolfson (1998). Our expert is an experienced industry (banking) professional with responsibilities in risk management and other aspects of business analytics. He has seen many portfolios of this type in different institutions. The elicitation took place in 2006. The expert found it easier to think in terms of the probabilities directly than in terms of defaults in a hypothetical sample. This is not uncommon in this technical area, as practitioners are accustomed to working with probabilities. The mean value was 0.01. The minimum value for the default probability was 0.0001 (one basis point). The expert reported that a value above 0.035 would occur with probability less than 10%, and an absolute upper bound was 0.3. The upper bound was discussed: the expert thought probabilities in the upper tail of his distribution were extremely unlikely, but he did not want to rule out the possibility that the rates were much higher than anticipated (prudence?). Quartiles were assessed by asking the expert to consider the value at which larger or smaller values would be equiprobable given the value was less than the median, then given the value was more than the median. The median value was 0.01. The former, the .25 quartile, was 0.0075. The latter, the .75 quartile, was assessed at .0125. The expert, who has long experience with this category of assets, seemed to be thinking of a distribution with a long and thin upper tail but otherwise symmetric. After reviewing the implications, the expert added a .99 quantile at 0.02, splitting up the long upper tail.

At this point a choice must be made on the representation of the elicited information. Of course, without further assumptions, we do not have enough information to specify a probability distribution. In principle that would require an infinity

of elicitation. However, choosing a parametric form for a statistical distribution allows determination of the parameters on the basis of the assessed information (assuming standard identification properties - one cannot assess a median alone and uniquely determine a $k > 1$ - parameter distribution). This is the most common approach in practice and parallels the usual practice in specifying the data distribution - a parametric form based (one hopes) on an economic model, allowing an infinity (or large number in the discrete case) of probabilities to be determined by finitely many parameters. This approach is illustrated in Kiefer (2010a), where the elicited information was used to fit a truncated Beta distribution. The disadvantage of this approach is that there is rarely good guidance beyond convenience on the choice of functional form. Thus, this choice can insert information not elicited from the expert nor really intended by the analyst. Based on experience, we prefer a nonparametric approach (really, less parametric), the maximum entropy (ME) approach (Kiefer (2010b)).

The ME approach provides a method to specify the distribution that meets the expert specifications and imposes as little additional information as possible. Thus, we maximize the entropy (minimize the information) in the distribution subject to the constraints indexed by k given by the assessments. Entropy is

$$H(p) = - \int \log(p(x))dP$$

Entropy is a widely used measure of the information in an observation (or an experiment). Further discussion from the information theory viewpoint can be found in Cover and Thomas (1991). The general framework is to solve for the distribution p

$$\max_p \left\{ - \int p \ln(p(x))dx \right\} \tag{9}$$

$$\begin{aligned}
& s.t. \int p(x)c_k(x)dx = 0 \text{ for } k = 1, \dots, K \\
& \text{and } \int p(x)dx = 1
\end{aligned}$$

In our application the assessed information consists of quantiles. The constraints are written in terms of indicator functions for the α_k quantiles q_k ; for example the median constraint corresponds to $c(x) = I(x < \text{median}) - 0.5$. To solve this maximization problem, form the Lagrangian with multipliers λ_k and μ and differentiate with respect to $p(x)$ for each x . Solving the resulting first-order conditions gives

$$p_{ME}(\theta) = \kappa \exp\left\{\sum_k \lambda_k (I(\theta < q_k) - \alpha_k)\right\} \quad (10)$$

The multipliers are chosen so that the constraints are satisfied. For details see Cover and Thomas (1991) or for an approach not using the Lagrangian Csiszar (1975).

This gives a piecewise uniform distribution for θ . It can be argued that the discontinuities in $p_{ME}(\theta)$ are unlikely to reflect characteristics of expert information and indeed this was the view of the expert. Smoothing was accomplished using the Epanechnikov kernel with several bandwidths h chosen to offer the expert choices on smoothing level (including no smoothing). Specifically, with $p_S(\theta)$ the smoothed distribution with bandwidth h we have

$$p_S(\theta) = \int_{-1}^1 K(u)p_{ME}(\theta + u/h)du \quad (11)$$

with $K(u) = 3(1 - u^2)/4$ for $-1 < u < 1$. Since the density $p_{ME}(\theta)$ is defined on bounded support there is an endpoint or boundary "problem" in calculating the

kernel-smoothed density estimator. Specifically, $p_S(\theta)$ as defined in (11) has larger support than $p_{ME}(\theta)$, moving both endpoints out by a distance $1/h$. We adjust for this using reflection, $p_{SM}(\theta) = p_S(\theta) + p_S(a - \theta)$ for $a \leq \theta < a + 1/h$, $p_{SM}(\theta) = p_S(\theta)$ for $a + 1/h \leq \theta < b - 1/h$, and $p_{SM}(\theta) = p_S(\theta) + p_S(2b - \theta)$ for $b - 1/h \leq \theta \leq b$. The resulting smoothed density has support on $[a, b]$ and integrates to 1. See Schuster (1985). The prior distribution for θ is shown in Figure 1.

Model 2 requires a prior on the asset correlation ρ . Here B2 provides guidance. For this portfolio bucket, B2 recommends a value of approximately 0.20. We did not assess further details from an expert on this parameter. There appears to be little experience with correlation, relative to expert information available on default rates. There is agreement that the correlation is positive (as it has to be asymptotically if there are many assets). Consequently, we choose a Beta prior with mean equal to 0.20 for ρ . Since the B2 procedure is to fix ρ at that value, any weakening of this constraint is a generalization of the model. We choose a Beta(12.6, 50.4) distribution, with a standard deviation of 0.05. This prior is illustrated in Figure 2. Thus, the prior specifications on the parameters for which we have no expert information beyond that given in the B2 guidelines reflect the guidelines as means and little else. The joint prior for θ and ρ is obtained as the product, which is the maximum-entropy combination of the given marginals. Here, it does not seem to make sense to impose correlation structure in the absence of expert information.

As to τ , here we have little guidance. We take the prior to be uniform on $[-1, 1]$. It might be argued that τ is more likely to be positive than negative, and this could certainly be done. Further, some guidance might be obtained from the literature on asset prices, though this usually considers less homogeneous portfolios. Here we choose a specification that has the standard B2 model at its mean value, so that allowing for nonzero τ is a strict generalization of existing practice.




Figure 1: Prior on the long-run default probability θ




Figure 2: Prior on the asset correlation ρ

5 Inference

Writing the likelihood function generically as $p(R|\phi)$ with $\phi \in \{\theta, (\theta, \rho), (\theta, \rho, \tau)\}$ depending on whether we are referring to the likelihood function (2), (7), or (8), and the corresponding prior $p(\phi)$, inference is a straightforward application of Bayes rule. The joint distribution of the data R and the parameter ϕ is

$$p(R, \phi) = p(R|\phi)p(\phi)$$

from which we obtain the marginal (predictive) distribution of R ,

$$p(R) = \int p(R, \phi) d\phi \tag{12}$$

and divide to obtain the conditional (posterior) distribution of the parameter ϕ :

$$p(\phi|R) = p(R|\phi)p(\phi)/p(R) \tag{13}$$

Given the distribution $p(\phi|R)$, we might ask for a summary statistic, a suitable estimator for plugging into the required capital formulas as envisioned by Basel Committee on Banking Supervision (2006). A natural value to use is the posterior expectation, $\bar{\phi} = E(\phi|R)$. The expectation is an optimal estimator under quadratic loss and is asymptotically an optimal estimator under bowl-shaped loss functions.

In many applications the distribution $p(\phi|R)$ can be difficult to calculate due to the potential difficulty of calculating $p(R)$ which requires an integration over a possibly high dimensional parameter. Here, the dimensions in models 1, 2, and 3 are 1,2, and 3. The first model can be reliably integrated by direct numerical integration, as can model 2 (requiring rather more time). Model 3 becomes very

difficult and simulation methods are more efficient. Since many applications will require simulation and efficient simulation methods are available, and since these methods can replace direct numerical integration in the simpler models as well, we describe the simulation approach. Here we describe the Markov Chain Monte Carlo concept briefly and give details specific to our application. For a thorough and wide-ranging description see chapter 2, especially section 4.2. and Robert and Casella (2004).

Markov Chain Monte Carlo methods are a wide class of procedures for calculating posterior distributions, or more generally sampling from a distribution when the normalizing constant is unknown. We consider here a simple case, the Metropolis method. The idea is to construct a sampling method generating a sample of draws $\phi^0, \phi^1, \dots, \phi^N$ from $p(\phi|R)$, when $p(\phi|R)$ is only known up to a constant. The key insight is to note that it is easy to construct a Markov Chain whose equilibrium (invariant, stationary) distribution is $p(\phi|R)$. Begin with a proposal distribution $q(\phi'|\phi)$ giving a new value of ϕ depending stochastically on the current value. Assume (for simplicity - this assumption is easily dropped) that $q(\phi'|\phi) = q(\phi|\phi')$. This distribution should be easy to sample from and in fact is often taken to be normal: $\phi' = \phi + \epsilon$ where ϵ is normally distributed with mean zero and covariance matrix diagonal with elements chosen shrewdly to make the algorithm work. Then, construct a sample in which ϕ^{n+1} is calculated from ϕ^n by first drawing ϕ' from $q(\phi'|\phi^n)$ then defining $\alpha(\phi', \phi^n) = p(R, \phi')/p(R, \phi^n) \wedge 1$ and defining $\phi^{n+1} = \phi'$ with probability $\alpha(\phi', \phi^n)$ or ϕ^n with probability $(1 - \alpha(\phi', \phi^n))$. Note that $p(R, \phi)$ is easy to calculate (the product of the likelihood and prior). Further, the ratio $p(R, \phi')/p(R, \phi^n) = p(\phi'|R)/p(\phi^n|R)$ since the normalizing constant $p(R)$ cancels. The resulting sample $\phi^0, \phi^1, \dots, \phi^N$ is a sample from a Markov Chain with equilibrium distribution $p(\phi|R)$. Eventually (in N) the chain will settle down and the

sequence will approximate a sequence of draws from $p(\phi|R)$. Thus the posterior distribution can be plotted, moments calculated and expectations of functions of ϕ can be easily calculated by sample means. Calculation of standard errors should take into account that the data are not independent draws. Software to do these calculations with a user-supplied $p(R, \phi)$ exists. We use the the `mcmc` package (Geyer (2009)) used in R (R Development Core Team (2009)). Some experimentation with these methods is useful to gain understanding. Valuable guidance and associated warnings are available on the website noted in the package documentation. Generally, an acceptance ratio of about 25% is good (see Roberts, Gelman, and Gilks (1997)). The acceptance rate is tuned by adjusting the variances of ϵ . Long runs are better than short. There is essentially no way to prove that convergence has occurred, though nonconvergence is often obvious from time-series plots. For our illustrative application M samples from the joint posterior distribution were taken after a 5000-sample burnin. Scaling of the proposal distribution allowed an acceptance rate between 22 and 25 percent. This procedure was used for Model II ($M=10000$) and for Model III ($M=40000$). Calculation of posterior distributions of the parameters and the functions of parameters considered below are based on these samples.

We construct a segment of upper tier high-yield corporate bonds, from firms rated Ba by Moody's Investors Service, in the Moody's Default Risk ServiceTM (DRSTM) database (release date 1-8-2010). These are restricted to U.S. domiciled, non-financial and non-sovereign entities. Default rates were computed for annual cohorts of firms starting in January 1999 and running through January 2009. In total there are 2642 firm/years of data and 24 defaults, for an overall empirical rate of 0.00908. The data are shown in Figure 3.




Figure 3: Default Rates




Figure 4: Model I, $p(\theta|R)$

The analysis of the binomial model is straightforward using direct calculations involving numerical integration to calculate the predictive distribution and various moments (recall we are not in a conjugate-updating framework due to the flexible form of the prior representation).

The posterior distribution for the binomial model is shown in Figure 4.

This density has $E(\theta|R = r = 24) = 0.0098$ and $\sigma_\theta = 0.00174$. Note that this is higher than the empirical default rate of 0.0091. The right skewness of the distribution is evident, which has flowed through from the prior distribution. The 95% credible interval for θ is (0.00662, 0.0134), which corresponds to a relative uncertainty of about 68% for the estimated PD.

Model II has asset value correlation within periods, allowing for heterogeneity in the default rate over time (but not correlated over time) and clumping of defaults. The marginal posterior distributions are shown in Figures 5 and 6.

We observe that the estimate of the probability of default in this model is




Figure 5: Model II, $p(\theta|R)$




Figure 6: Model II, $p(\rho|R)$

higher than in the 1-parameter model, this density having $E(\theta|R) = 0.0105$ and $\sigma_\theta = 0.00175$. The 95% credible interval for θ is (0.0073, 0.0140). This density has $E(\rho|R) = 0.0770$ and $\sigma_\rho = 0.0194$, so that there is a higher degree of variability relative to the mean in the estimated distribution of the asset value correlation, as compared to the probability of default parameter. The 95% credible interval for ρ is (0.0435, 0.119). Note that the prior mean (0.2) is well outside the posterior 95% confidence interval for ρ . Analysis of the Vasicek distribution shows that the data information on ρ comes through the year-to-year variation in the default rates. At $\theta = 0.01$ and $\rho = 0.2$ the Vasicek distribution implies an intertemporal standard deviation in default rates of 0.015. With $\rho = 0.077$, the posterior mean, the implied standard deviation is 0.008. In our sample, the sample standard deviation is 0.0063. This is the aspect of the data which is moving the posterior to the left of the prior.

The marginal posterior distributions for Model III are shown in Figures 7-9.

We observe that the estimate of the probability of default in this model is slightly higher than in the 1-parameter model, this density having $E(\theta|R) = 0.0100$ and $\sigma_\theta = 0.00176$. This density has $E(\rho|R) = 0.0812$ and $\sigma_\rho = 0.0185$ with a 95% credible interval of (0.043, 0.132). The density of the autocorrelation parameter in the latent systematic factor has $E(\tau|R) = 0.162$ and $\sigma_\tau = 0.0732$. The 95% credible interval is (-0.006, 0.293).

In summary, the picture on the default probability is pretty clear: it is around 0.01 in all models. The asset value correlation is around 0.08, estimated to be somewhat but only slightly higher in model III than in model II. This is substantially less than the value specified in B2. The temporal correlation in the systematic factor is only present in model III. The evidence is sparse here (recall there are only 11 years of data and the prior information was as uninformative as possible) but it appears to be slightly positive.

Posterior Density: Probability of Default (3-Parameter Model)




Figure 7: Model III $p(\theta|R)$

Markov Chain Monte Carlo Posterior Density: Asset Value Correlation (3-Parameter Model)




Figure 8: Model III, $p(\rho|R)$

Markov Chain Monte Carlo Posterior Density: Autocorrelation in Systematic Factor (3-Parameter Model)




Figure 9: Model III, $p(\tau|R)$

6 Conclusion

In this and related applications the econometrician faces the dual chore of modeling the data distribution with a specification of a statistical distribution and modeling expert information with a statistical distribution. Adding the latter task substantially increases the range of applicability of econometric methods. This is clearly an area for further research. Our application has gone through the steps of a formal Bayesian analysis, focussing on the default probability, a key parameter which is required to be estimated under B2 by a large number of institutions worldwide. We concluded our analysis by generating the posterior distributions for the parameters of a nested sequence of models and calculating summary statistics. The mean default probability would be a natural estimator to use for calculating minimum regulatory capital requirements using the formulas provided by B2. In practice, these distributions have many uses, and the analysis would be ongoing. For example, institutions might want to use the entire distribution of the default probability in pricing credit and in setting in-house capital levels. Having an entire distribution for the PD could also be useful in stressing of IRB models - i.e., plugging in a high quantile of PD into the capital formula to model an adverse scenario. The more general models provide insight into the extent to which default rates over time are predictable, and to the extent to which risk calculations should look ahead over a number of years. An analysis of Loss Given Default (LGD) using Bayesian methods would be useful - here there is substantial experience and a joint analysis of LGD and the default probability is likely to be extremely interesting. These and many other possible analyses build on the methods illustrated here.

References

- BASEL COMMITTEE ON BANKING SUPERVISION (2006): “International Convergence of Capital Measurement and Capital Standards: A Revised Framework, Comprehensive Version,” Bank for International Settlements.
- (2009a): “Proposed enhancements to the Basel II framework: Consultative Document,” Discussion paper, Bank for International Settlements.
- (2009b): “Range of practices and issues in economic capital frameworks,” Discussion paper, Bank for International Settlements.
- BLUHM, C., L. OVERBECK, AND C. K. J. WAGNER (2003): *An Introduction to Credit Risk Modeling*. Chapman & Hall-CRC, Financial Mathematics Series, London.
- COVER, T. M., AND J. A. THOMAS (1991): *Elements of Information Theory*. John Wiley & Sons.
- CSISZAR, I. (1975): “I-divergence Geometry of Probability Distributions and Minimization Problems,” *The Annals of Probability*, 3, 146–158.
- DAS, S. R., D. DUFFIE, N. KAPADIA, AND L. SAITA (2007): “Common Failings: How Corporate Defaults are Correlated,” *Journal of Finance*, 62, 93–117.
- GARTHWAITE, P. H., J. B. KADANE, AND A. O’HAGAN (2005): “Statistical Methods for Eliciting Probability Distributions,” *Journal of the American Statistical Association*, 100, 780–700.
- GEYER, C. J. (2009): *mcmc: Markov Chain Monte Carlo*. R package version 0.6.

- GORDY, M. B. (2003): “A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules,” *Journal of Financial Intermediation*, 12, 199–232.
- KADANE, J. B., AND L. J. WOLFSON (1998): “Experiences in Elicitation,” *The Statistician*, 47(1), 3–19.
- KIEFER, N. M. (2007): “The Probability Approach to Default Probabilities,” *Risk*, pp. 146–150.
- (2009): “Correlated Defaults, Temporal Correlation, Expert Information and Predictability of Default Rates,” Discussion paper, Cornell University.
- (2010a): “Default Estimation and Expert Information,” *Journal of Business and Economic Statistics*, 28(2), 320–328.
- (2010b): “Default Estimation, Correlated Defaults and Expert Information,” *Journal of Applied Econometrics*, forthcoming.
- MCCULLAGH, P., AND J. NELDER (1989): *Generalized Linear Models*, 2nd ed. Chapman & Hall, London.
- MCNEIL, A. J., AND J. P. WENDIN (2007): “Bayesian inference for generalized linear mixed models of portfolio credit risk,” *Journal of Empirical Finance*, 14, 131149.
- NICKELL, P., W. PERRAUDIN, AND S. VAROTTO (2000): “Stability of Rating Transitions,” *Journal of Banking and Finance*, 24, 203–227.
- OCC (2000): “OCC Bulletin 2000-16, Subject: Risk Modeling, Description: Model Validation,” Discussion paper, Office of the Comptroller of the Currency.

- O'HAGAN, A., C. E. BUCK, A. DANESHKHAH, J. R. EISER, P. GARTHWAITE, D. J. JENKINSON, J. E. OAKLEY, AND T. RAKOW (2006): *Uncertain Judgments: Eliciting Experts' Probabilities*. Chichester: John Wiley & Sons.
- R DEVELOPMENT CORE TEAM (2009): *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- ROBERT, C., AND G. CASELLA (2004): *Monte Carlo Statistical Methods (2nd edition)*. New York: Springer-Verlag.
- ROBERTS, G. O., A. GELMAN, AND W. R. GILKS (1997): "Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms," *The Annals of Applied Probability*, 7(1), 110–120.
- SCHUSTER, E. F. (1985): "Incorporating support constraints into nonparametric estimators of densities," *Communications in Statistical Theory and Methods*, 14, 1123–1136.