Ratings-based credit risk modelling: an empirical analysis

Pamela Nickell — Moody's-KMV
William Perraudin — Imperial College, UK
Simone Varotto* — ISMA Centre, UK

May 6, 2005

Abstract

Banks have recently developed new techniques for gauging the credit risk associated with portfolios of illiquid, defaultable instruments. These techniques could revolutionise banks' management of credit risk and could in the longer term serve as a more risk-sensitive basis for calculating regulatory capital on banks' loan books than in Basel 2, the new regulatory capital framework. In this paper we implement a popular credit risk model that exploits the information in credit ratings to determine a portfolio's value-at-risk. Using price data on large eurobond portfolios, we assess, on an out-of-sample basis, how well the model tracks the risks it is supposed to measure.

Keywords: Credit risk models, Value-at-risk

JEL Classification: G11, G21, L51

*Corresponding author. Address: ISMA Centre, University of Reading, Whiteknights Park, PO Box 242, Reading RG6 6BA, UK. Tel.: +44(0)118 378 6655, fax: +44(0)118 931 4741. Email: s.varotto@ismacentre.rdg.ac.uk. We thank Angus Guyatt for outstanding research support. We thank Reza Bahar, Patricia Jackson and other Bank of England colleagues for valuable comments and Angus Guyatt and Steve Grice for research assistance. The authors are very grateful to Peter Bryant of Reuters for facilitating the collection of the data used in this study.

1 Introduction

The systematic application of Value-at-Risk (VaR) models[1] by large international banks has significantly enhanced their ability to measure and hedge their trading book risks. A valuable side-effect of the new emphasis on VaR modelling is that regulators have been able to reduce the distortionary impact of prudential capital requirements for banks' trading portfolios by basing such requirements on VaRs generated by banks' internal risk management models.[2]

Recently, regulators have begun to consider the use of internal models for measuring credit risk to increase the effectiveness of the capital requirement regime. As with VaR applied to trading books, the possibility has arisen because banks themselves are exploring the use of credit risk models for measuring the riskiness of their portfolios. A crucial question preoccupying both firms and financial policy makers has been whether internal models could be used to assess required capital for banking books.[3]

The fundamental difficulty in assessing credit risk is that most credit exposures have no easily observable market price. The lack of price information obliges one to base credit risk estimates on other kinds of data.[4] The two feasible approaches in current use are (i) ratings-based methods (exemplified by JP Morgan's Creditmetrics), and (ii) equity price-based techniques (advocated by, for example, Moody's KMV).

In this paper we focus on the former type of model. Ratings-based techniques attribute a rating to each defaultable investment in a portfolio. Then, they estimate the probability of upward or downward moves in ratings using historical data on ratings transitions for different traded bond issues. The probabilities are collectively termed the ratings transition matrix. By simulating rating scenarios that are consistent with the transition probabilities one can derive the empirical distribution of the value of the portfolio and calculate the portfolio's value-at-risk.

The credit exposures we examine are dollar-denominated eurobonds. The unusually rich dataset we employ includes 1,294 bond price histories observed from 1988 to 1998. All the bonds are straight bonds with no call or put features. To implement the model,

[1]Such models estimate the VaR on a portfolio, i.e. the loss that will be exceeded on some given fraction of occasions if the portfolio in question is held for a particular period.

[2]See Basel Committee on Banking Supervision (1996).

[3]The conclusion of the Basel Committee has been that this would be premature, although it accepts that model-based capital calculation may be feasible in the longer term.

[4]If prices were observed in a reasonably liquid market, one might attempt to assess risk through a relatively simple VaR calculation.

we also collected the ratings histories for all the bonds in our sample.

We implement the model month by month, calculating in each period a credit risk VaR for the following year. We are careful only to employ lagged data which would have been available to an analyst implementing the model in the given period. To assess the model's performance, we then compare the estimated VaRs with the actual outturn for the portfolio in question one year later.[5] If the model supplies unbiased VaR estimates, the fraction of occasions on which losses exceed the VaRs will roughly equal the VaR confidence level.

Our major conclusion is that the model as commonly implemented may significantly under-estimate the risks involved in holding our eurobond portfolios. The implication of our analysis is not that credit risk models have no value. Rather it suggests that models should be employed cautiously with conservative parameterisations.

Although credit risk models are a very recent development, they are the subject of a rapidly growing literature. Surveys of the techniques employed may be found in Basel Committee on Banking Supervision (1999a) and Crouhy, Galai and Mark (2000). Discussion of issues related to the regulatory use of credit risk models is provided by Mingo (2000), Jones (2000), and Jackson and Perraudin (2000).

Empirical investigations of these models has so far been limited. Lopez and Saidenberg (2000) suggest techniques for assessing models through cross-sectional evaluation of their risk measures but do not implement their suggestions on actual data. Gordy (2000) and Kiesel, Perraudin and Taylor (2003) implement ratings-based models on stylised portfolios, studying how the risk measures vary across different types of portfolio. Gordy (2000) and Crouhy et al (2000) also compare the VaRs implied for the similar portfolios at one point in time by different models. Nickell et al (2000) discuss the stability of ratings transition matrices, the central component of ratings-based credit risk models. To our knowledge, the current study is the only out-of-sample assessment of a credit risk model using time series data.

The structure of our paper is as follows. In Sections 2 we describe the ratings-based credit risk modelling approach. In Section 3 we detail the substantial datasets that we created in order to perform our analysis. Section 4 gives the results of our investigation, and Section 5 concludes, drawing out the lessons learnt and discussing further work that is needed to investigate aspects of credit risk modelling methodologies.

[5] It is possible to do such comparisons only because we apply the credit risk model to liquid, defaultable instruments like bonds.

2 A ratings-based credit risk model

The ratings-based approach, of which the canonical example is the Creditmetrics model,[6] provides a framework for estimating the distribution of the future value of a portfolio of credit-sensitive exposures. The assumptions on which Creditmetrics is based are as follows. First, it is supposed that an obligor's credit standing is completely summed up by its rating. Second, it is assumed that the spreads at which the market will in future discount payoffs by obligors with particular ratings are known. These two assumptions immediately imply that if one knows, for all i and j, the probabilities that an i-rated obligor will be j-rated in the future, then one may deduce the distribution of the value of the exposure at the future date in question.

To estimate the distribution of a portfolio of credit exposures, however, requires that one knows more than the univariate distributions of the individual components of the portfolio. In particular, one must know the correlation structure (or more generally the dependence) of the different exposures. The third assumption of Creditmetrics is that the dependence between ratings changes stems from the fact that individual ratings changes are driven by latent variables that are themselves correlated. The correlations are for simplicity taken to equal those of weighted sums of equity indices where the weights are chosen to reflect the degree to which the obligor in question is engaged in different industries or markets.

The Creditmetrics manual by Gupton, Finger and Bhatia (GFB) (1997), describes a wide range of techniques of varying technical sophistication and practical use. We base our analysis on what appears to us to represent the core set of techniques. These involve the derivation of the empirical distribution of bond portfolios from data on (i) rating transition matrices, (ii) bond spreads, (iii) equity index correlations, and (iv) the level of idiosyncratic risk in the return of each obligor's equity.

To understand how Creditmetrics works, suppose that the probability that a firm will default is fully described by its current rating, say j. Consider a discount function $D_j(t, t')$ corresponding to the market price at date t of a promise to deliver 1 dollar at date t' in the event that a given j-rated bond issue does not default before t'. Let $D_0(t, t')$ denote the price of a default-free Treasury strip. We can implicitly define the credit spread for a

[6] Another ratings-based approach (in which correlations are modelled differently) has been implemented by McKinsey Inc. (see Wilson (1997)). CSFP (1997) have proposed a model similar to ratings-based approaches in which obligors occupy just two 'ratings categories', default or non-default. Among ratings-based approaches, we focus in this study on the Creditmetrics model.

j-rated obligor from t to t', denoted $S_j(t,t')$ as:

$$D_j(t,t') \equiv D_0(t,t') \exp[-S_j(t,t')(t'-t)] \tag{1}$$

If there are $J+1$ rating categories of which the $J+1$st represents default, the price at t of a j-rated, defaultable, coupon bond may be expressed as:

$$B_t^{(j)} = \begin{cases} \sum_{i=1}^{N} D_0(t,t_i) \exp[-S_j(t,t_i)(t_i-t)]\, c_i & j=1,2,\ldots,J \\ \sum_{i=1}^{N} \xi\, D_0(t,t_i)\, c_i & j=J+1 \end{cases} \tag{2}$$

where $c_1, c_2, \ldots, c_N$ are the promised cash flows (including the repayment of principal) subsequent to t at respective payment dates $t_1, t_2, \ldots, t_N$, and the recovery rate, ξ, is assumed fixed.

Let $\pi_{ij}(t,T)$ denote the probability that a bond issue rated i at t will be rated j at T. The matrix of such probabilities, $[\pi_{ij}(t,T)]$, is termed the ratings transition matrix. Estimates of transition matrices covering different sample periods and for different types of obligor may be obtained from ratings agency publications (see, for example, Lucas and Lonski 1992 and Carty 1997) or from academic articles (see Altman and Kao 1992 and Nickell et al 2000).

Suppose that at date t, the ratings-contingent discount functions which will pertain at $T > t$, namely $D_j(T,s)$, are known. From equation (2), one may deduce the value that the bond will have at T conditional on knowing its rating at T. Using transition probabilities, $\pi_{ij}(t,T)$, it is then straightforward to deduce moments of the bond price conditional on information at t:

$$\text{Mean}_{it} = \sum_{j=1}^{J+1} \pi_{ij} B_T^{(j)} \equiv \mu^{(i)} \tag{3}$$

$$\text{Variance}_{it} = \sum_{j=1}^{J+1} \pi_{ij} (B_T^{(j)} - \mu^{(i)})^2 \tag{4}$$

The assumption made in Creditmetrics that the future ratings-contingent discount factors are known is a strong assumption. Kiesel et al (2003) show how it may be relaxed and argue that allowing for randomly varying spreads is important if one is modelling the risk of portfolios comprising high credit quality exposures.

2.1 Modelling correlations

The way in which Creditmetrics allows for correlated ratings transitions consists of assuming that each obligor's ratings transitions are driven by a normally distributed latent variable. More formally, suppose that for an i-rated bond issue, there exists a normally distributed variable, R_i, such that, for a set of cut-off points, $(Z_{i,0}, Z_{i,1}, \ldots, Z_{i,J})$, if R_{ik} falls in the interval $(Z_{i,j-1}, Z_{i,j})$, the bond is rated j at date T.

Given an estimate of a ratings transition matrix, $[\pi_{i,j}(t,T)]$, the cut-off points $Z_{i,j}$ may be deduced directly using the recursive equations:

$$\begin{cases} \pi_{i,J}(t,T) &= 1 - \Phi(Z_{i,J-1}) \\ \pi_{i,j}(t,T) &= \Phi(Z_{i,j}) - \Phi(Z_{i,j-1}) \qquad j = 2, \ldots, J \\ \pi_{i,1}(t,T) &= \Phi(Z_{i,1}) \end{cases} \tag{5}$$

The approach of allowing multinomial transitions to be driven by an underlying latent variable with a continuous distribution is widely applied in the discrete choice econometrics literature. When the latent variable is normally distributed it corresponds to an ordered probit approach. The major benefit of this approach for credit risk modelling is that it permits one to model the correlation between different ratings transitions in a straightforward fashion.

Creditmetrics makes the simple assumption that correlations between the latent variables driving transitions for different bond issues, say R and R', equal those of the firms' respective log equity values. Without loss of generality, one may standardise the variances and means of the latent variables to unity and zero, respectively. The joint probability of ratings transitions for a pair of obligors with initial rating i and l and end-of-period rating of j and m is then:

$$\begin{aligned} \pi_{(i,j),(l,m)} &= \text{Prob}\left\{Z_{i,j-1} < R < Z_{i,j} \; , \; Z_{l,m-1} < R' < Z_{l,m}\right\} \\ &= \int_{Z(i,j-1)}^{Z(i,j)} \int_{Z(l,m-1)}^{Z(l,m)} \phi(s, s'|\sigma, \rho) ds ds' \end{aligned} \tag{6}$$

Here, ϕ is a standard bivariate normal density with a correlation coefficient, ρ. Rather than estimate ρ using time series data on equity values, the Creditmetrics manual suggests that one construct proxies consisting of weighted averages of industry and country equity indices, assuming that the given equity return also contains some additional amount of idiosyncratic risk.

This approach is particularly simple when there is a single index for each obligor since one may then express the standardised (*i.e.* unit variance) log return of the nth firm's equity value as:

$$r_n = \omega_{1n} r_M + \omega_{2n} \hat{r}_n \tag{7}$$

where r_M and $\hat{r}_n$ denote the standardised return on the index and the idiosyncratic component of the firm's equity return. If ω_{1n} equals α, then the standardisation implies that

$\omega_{2n} = \sqrt{1-\alpha^2}$.

Thus, given the correlation matrix for the standardised indices, $[\zeta_{mn}]$, and an assumption about the fraction of volatility that is idiosyncratic for each obligor (*i.e.* a choice of α for each obligor), one may deduce the correlation of the obligors' latent variables as:

$$\rho_{n,m} = \alpha_n \alpha_m \zeta_{nm} \tag{8}$$

Note that what we have just described is simpler than the approach set out in Creditmetrics since we have supposed that there exists a single index for each obligor rather than several national and industry indices with known weights, α_{mh}.

2.2 Portfolio distribution

By knowing the correlation structure of the latent variables one can analytically derive the portfolio variance. The ingredients that are necessary for this calculation are (1) the correlations obtained in the previous sections ($\alpha_n \alpha_m \zeta_{nm}$), (2) the assumption that the multivariate distribution of the latent variables is normal, (3) the cut-off points $(Z_{i,0}, Z_{i,1}, \ldots, Z_{i,J})$ for each rating i and, (4) the value that each bond takes when a migration from the current rating to any rating j occurs (i.e. $B_t^{(j)}$, see equation 2). The role of correlation is to determine the shape of the multivariate normal and hence the probability associated with any pair of obligors in each of the possible future rating combinations which they may migrate to. The covariance between the values of two bond issues with initial ratings i and l will then be,

$$\sum_{j,m} \left(B_T^{(j)} - \mu^{(i)}\right) \left(B_T^{(m)} - \mu^{(l)}\right) \pi_{(i,j),(l,m)}$$

Portfolio volatility can be easily calculated by combining covariances with the individual bond variances computed as shown in (4). At this point, if one is prepared to assume normality of portfolio values, the Value-at-Risk of the portfolio is easily derived (see Section 4).

An alternative and more precise method for the calculation of portfolio VaRs is suggested in GFB. Their approach is based on Monte Carlo simulations. The idea is to generate values for the latent variables driving future rating scenarios, where the latent variables are assumed to be jointly normally distributed with correlation matrix $\left[\rho_{n,m}\right]$. Given the cut-off points $(Z_{i,0}, Z_{i,1}, \ldots, Z_{i,J})$, it is possible to associate each simulated value of our latent variables with a specific future rating j and bond value $B_T^{(j)}$. By adding the simulation-generated values for all the bonds in the sample, we obtain a "realisation" of the portfolio distribution,

$$\sum_{\delta}^{S} B_{T,\delta}^{(j_\delta)}$$

where j_δ is the simulated end-of-period rating of bond δ in a particular simulation scenario. We use 5,000 scenarios to build the empirical distribution and compute 99% VaRs.

3 Data

The data requirements of our study are considerable since we calculate monthly risk measures generated by the model and then compare them with outturns for large portfolios of bonds. Let us start by describing the substantial bond price dataset we created.

This comprised 1,294 US dollar-denominated bonds[7] selected from the much larger number of bonds listed on the Reuters 3000 price service. Our criteria in selecting the bonds were (i) that they were neither callable nor convertible, (ii) that a rating history was available, (iii) that the coupons were constant with a fixed frequency, (iv) that repayment was at par, and (v) that the bond did not possess a sinking fund. To arrive at the 1,294, we further eliminated bonds for which the price and rating histories did not overlap for more than a year, and very illiquid bonds with price histories which contained at least one gap of more than 100 days.

The prices we used were Reuters composite bids. The Reuters composite is the best bid reported at close of trading by a market-maker from which Reuters has a data feed. Our data included comprehensive information about the cash flows, ratings and price histories of the bonds, and the name, domicile and industry code of the obligor.

To conduct our various analyses, we created a series of different portfolios. We use the term 'total portfolio' to denote a portfolio comprising one unit of every bond available on a given day. (The composition of the portfolio therefore evolves over time.) We subdivided the total portfolio into bonds issued by US and non-US obligors, and into bonds issued by banks and financials or by other obligors. We also examined 'quartile samples' made up of four randomly selected sub-samples of the total sample, each containing a quarter of the bonds available. Lastly, we studied two concentrated portfolios made of fifty highest and lowest rated securities.

[7] Of these, 90% were eurobonds, the remainder being national bonds from several countries.

The composition of the total portfolio is shown in Table 1. 48% of bonds had issuers domiciled in the United States. A further 33% had issuers domiciled in Japan, the Netherlands, Germany, France or the United Kingdom. A large fraction of bonds, namely 66%, was issued by firms in the financial services or banking industries. 57% of bonds were unsecured.

Additionally, to implement Creditmetrics, we also needed: (i) transition matrices, (ii) default spreads and default-free yield curves over time, (iii) equity index data, (iv) a set of weights linking individual obligors to the equity indices, and (v) an assumption about the fraction of equity volatility that is idiosyncratic for each obligor. The transition matrix we employed was the unconditional Standard and Poor's transition matrix provided by the Creditmetrics software Credit Manager. Default-free interest rates and spreads for different ratings categories were taken from Bloomberg.[8]

To obtain equity indices for the obligors, we created a time series dataset going back to 1983, comprising 243 country and industry-specific MSCI indices. For each obligor, we then chose one of these indices as the source of non-idiosyncratic risk.

4 Value-at-Risk and pricing errors

To assess the performance of different credit risk models, we compare VaR measures for a one-year holding period with the actual outturns of different portfolios. These comparisons are complicated, however, by the fact that the model described above abstracts from interest volatility in calculating risk measures. To see how well the model measures credit risk, one must, therefore, remove from the portfolio value realisation that part of the value change that is attributable to changes in the default-free term structure.

To explain the adjustments we made, we adopt the following notation. Let P_t denote the value at time t of a bond portfolio and let $P_{t,T}$ denote the expectation conditional on information at t of the portfolio value at T. Furthermore, let Q_t represent the price at t of a portfolio of default-free bonds having the same contractual payments as the defaultable bonds in our portfolio. Let $Q_{t,T}$ equal the forward price of the default-free portfolio at t.

To remove the effects of default-free interest rate changes on our portfolio return, we

[8]We used spreads for United States industrials since these had the longest series and the fewest missing observations.

work with the following adjusted return:

$$\frac{P_T - P_{t,T} - (Q_T - Q_{t,T})}{P_{t,T}} \tag{9}$$

This quantity represents the net return one would have obtained by investing a unit amount in the portfolio of defaultable bonds if value changes due to changes in general interest rates had been hedged. Suppose that the empirical distribution generated by the credit risk model indicates that for some confidence level, c, $\mathrm{Prob}_t(P_T < \gamma) = c$ for a cut-off point or 'VaR quantile', γ, then we can compare the return in equation (9) with the quantity:

$$\frac{\gamma - P_{t,T}}{P_{t,T}} \tag{10}$$

If (9) falls below (10), then the loss on the position has exceeded the VaR.

γ can be directly deduced from the portfolio value empirical distribution or by assuming normality. In the latter case, the 99% VaR, may be calculated by inverting the probability statement:

$$\mathrm{Prob}\left\{P_T < \gamma\right\} = \Phi\left(\frac{\gamma - \mu_p}{\sigma_p}\right) = 0.01 \tag{11}$$

to obtain:

$$\gamma = \mu_p + \Phi^{-1}(0.01)\ \sigma_p \tag{12}$$

where μ_p and σ_p are the portfolio's analytical mean and volatility.

A second problem that we face in gauging the accuracy of the model's risk predictions is that the model is effectively being used to *price* portfolios as well as to measure their risk. Let $P_t^{(m)}$ denote the value that the model attributes to the portfolio at time t. To correct for pricing errors in the expected price, $P_{t,T}$, we add the pricing discrepancy at the initial date, t. In (9) and (10), we therefore replace $P_{t,T}$ with

$$\tilde{P}_{t,T}^{(m)} \equiv P_{t,T}^{(m)} + P_t - P_t^{(m)} \tag{13}$$

5 Results

Annual VaR estimates and annual portfolio returns calculated month-by-month for the period January 1989 to February 1998 for various portfolios are reported in Figures 2 to 9. The Figures indicate actual portfolio returns adjusted for interest rate risk as in equation (9) and for pricing errors as shown in equation (13). The VaRs are based on a one-year holding period and a 99% confidence level. Returns and VaRs are in unit of percent (multiplied by 100).

In the terminology of the Basel Committee on Banking Supervision (1996), an exception occurs when the outturn loss on a portfolio exceeds the VaR measure supplied by a VaR model. In Figure 2, such an exception takes place when the solid line representing year-on-year returns[9] falls below one of the VaR levels, which appear in the Figure as dashed and dashed-and-dotted lines. If the credit risk models were correctly measuring risk, and we had non-overlapping observations, the returns would cross the VaR quantiles approximately once every 100 years. (With overlapping observations, the crossing will be more frequent.)

In the Figures we plot the return of the total portfolio and VaRs derived from a Monte Carlo generated portfolio distribution for various values of α and for a 50% recovery rate in case of default (which is consistent with the dominant seniority of the bonds in our sample[10]), unless otherwise stated. The VaR level and the number of exceptions, are highly sensitive to the choice of α. For all the obligors in our sample for which we have equity prices (367 of them) we estimate α as the correlation between the firm's equity and its national industry index. This follows from GFB if one assumes that obligors are solely operating in one industry sector. Only 2% of our obligors are classified by Reuters as "multi-industry". In these few cases obligors are allocated to their dominant sector.

Since, for several companies in the sample, stock prices are only available in later years in the sample period, we compute αs whenever the equity and index series overlap. Results are reported in Figure 1. The average estimated α is 0.263 and its standard deviation 0.174. Less than 5% of the observations exceed 0.6. In our analysis, for the various portfolios we consider, we shall use a common α for all obligors represented in the portfolio. We do so, as the paucity of our equity data does not allow us to estimate individual αs for all companies at each point in time in the sample period. In Figure 2, we look at the total return of the whole portfolio and compare it with Monte Carlo VaRs obtained with αs equal to 0.25 (that represents the average value), 0.6 and 0.95. A value of 0.95 may appear an unreasonable choice. However, it may be appropriate to consider

[9] As a check on our returns data, we compared our return series with changes in Bloomberg spreads and found they were broadly consistent. We also noted that the two most marked losses narrowly preceded the quarters in which defaults as measured by Moody's peaked within our sample period.

[10] More than 70% of our bonds are unsecured, (57.34% of unsecured proper and 14.14% of senior unsecured) which, according to Carty and Lieberman (1996) have an average recovery rate of 51.13%. Bonds with lower seniority (that is, subordinated), and hence with a lower recovery rate, account for only 1.55% of the total sample.

high αs, in that, if regulators allowed banks to use credit risk models for setting capital requirements they would likely ask banks to take a conservative approach when choosing the models' parameters. Also, Figure 2 clearly highlights that Creditmetrics can be quite a poor predictor of portfolio risk if ones uses the average α, as indicated by the large number of exceptions. For α equal to 0.6 we have three exceptions over the 10 years of the sample period, again well above the one exception in one hundred years one should expect when using 99% annual VaRs. However, with α equal to 0.95 no exception occurs.

The conclusion we can draw from this Figure is that a standard parameterisation of CreditMetrics may lead to underestimation of risk. An explanation may be that, although sound on paper, the CreditMetrics model may need to be calibrated with current market data, as is the case, for example, with other credit risk models now popular in the industry (see, for example, the KMV approach in Crouhy et al 2000). Another option would be to use backtesting and adjust the model's parameters according to its past performance. This approach is popular with regulators and forms the basis of bank capital rules for market risk. Lack of long time series data should not hamper the feasibility of backtesting since large cross-sections of data may be used instead, as suggested by Lopez and Saidenberg (2000).

Figure 3 compares the average Monte Carlo VaR with the VaR derived by assuming that portfolio values are normally distributed. The large difference between the two risk measures is due, as one should expect, to the pronounced negative skewness and leptokurtosis of loan portfolio distributions, caused by the potentially large downfalls that follow default events.

In Figures 4 and 6 we consider four sub-sample portfolios, banks and financials, other than banks and financials, US obligors, other than US obligors. Their VaR is estimated with an α of 0.25 and 0.95. In all cases the average α is clearly too low. With $\alpha = 0.95$ in Figure 4, the portfolio of non-US obligors exhibits two exceptions in consecutive months, February and March 1995. Figure 5a and 5b show the Monte Carlo 99% VaR and the 99% quantiles of the VaR estimates for the portfolio of non-US obligors. Any portfolio return below the dashed line delimiting the 99% quantile, is a statistically significant exception at the 1% confidence level. In Figure 5a, VaRs are estimated by assuming a 50% recovery rate. In Figure 5b we assume a 0% recovery. With a 50% recovery rate the p-value of the February exception is less than 0.001, hence highly statistically significant, while the March exception (with a p-value of 0.27) is not significant. With 0% recovery the two exceptions disappear. This suggests that a conservative parameterisation may cure the risk underestimation bias that appears to affect the model.

Figure 6 shows returns and VaRs for four quartiles of the whole portfolio. Quartile number 4 exhibits two exceptions in early 1990 and early 1991, which are both not significant as shown in Figure 7a and 7b, where the 50% and 0% recovery assumptions are used. Finally, we analyse the performance of the model in concentrated high an low quality portfolios. Figure 8a and 9a show returns, VaRs and VaR 99% quantiles (α=0.95) for the fifty lowest rated bonds and the fifty highest rated bonds, with a 50% recovery rate. In the lowest rated bond portfolio we obtain six exceptions, one each month from March to June 1995 and again in November and December of the same year, all statistically significant (p-values are always less than 0.001). For the highest rated bond portfolios we observe three non-significant exceptions (September and October 1992 and December 1997) and two highly significant exceptions in August and November 1997 (with p-values less than 0.001). When using a 0% recovery rate, as in Figure 8b and 9b, the number of exceptions that are statistically significant decreases. However, in the high quality portfolio (Figure 9b), we still have two significant exceptions in August and November 1997. This is surprising given the conservative parameterisation adopted (α=0.95 and 0% recovery). The result may reflect the fact that spreads of high quality bonds are strongly influenced by factors other than credit risk namely, market risk and liquidity risk (see, for example Elton et al 2001 and Perraudin and Taylor 2003).[11]

6 Conclusion

In this paper, we conduct the first out-of-sample evaluation of a new type of credit risk models. The model we study derives risk estimates for a portfolio of credit exposures by exploiting the information embedded in the exposures' credit rating. The most important feature of the new methodology is to provide a way to model the correlation of rating migrations.

Our approach consists of implementing the model over a ten year period on large portfolios of eurobonds. Month-by-month, we calculate the risk measures implied by the model and compare them with the actual outcomes as credit spreads move around. We are careful in each period only to employ lagged data so that the evaluation is genuinely out-of-sample.

We conclude that a standard implementation of the model may lead to severe un-

[11] The authors also estimate tax effects. Most of the bonds in our analysis, however, are eurobonds which are bearer securities. Interest on eurobonds during the sample period was paid gross to the issuer, without any deduction of withholding tax.

derestimation of portfolio risk. We find this problem can be eased by using conservative parametrisations. Our results suggest that it would be prudent to build in safety margins into capital allocation decisions and regulatory capital calculations if at a future date they were based on output from the current generation of rating-based credit risk models.

References

[1] Altman, E I and Kao, D L (1992), 'Rating drift of high yield bonds', Journal of Fixed Income, pages 15-20.

[2] Basel Committee on Banking Supervision (1996), 'Overview of the amendment to the capital accord to incorporate market risk', unpublished mimeo, Bank for International Settlements, Basel.

[3] Basel Committee on Banking Supervision (2004), 'International Convergence of Capital Measurement and Capital Standards: a Revised Framework', unpublished mimeo, Bank for International Settlements, Basel.

[4] Carty, L V (1997), 'Moody's rating migration and credit quality correlation, 1920-1996', Special comment, Moody's Investors Services, New York.

[5] Carty, L V and Lieberman, D (1996) 'Corporate bond default and default rates 1938-1995', Moody's Investors Service, Global Credit Research, January.

[6] Credit Suisse Financial Products (1997), 'Credit risk+ : technical manual', Discussion paper, CSFP.

[7] Crouhy, M, Galai, D and Mark, R (2000), 'A comparative analysis of current credit risk models', Journal of Banking and Finance, 24(1-2), pages 59-117.

[8] Duffee, G R (1998), 'The relation between treasury yields and corporate bond yield spreads', Journal of Finance, 53(6), pages 2,225-41.

[9] Elton, E, Gruber, M, Agrawal, D and Mann, C (2001), 'Explaining the rate spread of corporate bonds', Journal of Finance, Vol. LVI, No. 1, February, pages 247–77.

[10] Gordy, M (2000), 'A comparative anatomy of credit risk models', Journal of Banking and Finance, 24(1-2), pages 119-49.

[11] Hendricks, D (1996), 'Evaluations of value-at-risk models using historical data', Federal Reserve Bank of New York Economic Policy Review, Vol.2(1), pages 39-69.

[12] Jackson, P D and Perraudin, W R (2000), 'Regulatory implications of credit risk modelling', Journal of Banking and Finance, 24(1-2), pages 1-14.

[13] Jones, D (2000), 'Emerging problems with the accord: regulatory capital arbitrage and related issues', Journal of Banking and Finance, 24(1-2), pages 35-58.

[14] Jones, E P, Mason, S P and Rosenfeld, E (1984), 'Contingent claims analysis of corporate capital structures: an empirical investigation', Journal of Finance, Vol.39(3), pages 611-27.

[15] Gupton, G M, Finger, C C and Bhatia, M (1997), Creditmetrics-Technical Document, JP Morgan, New York.

[16] Kiesel, R, Perraudin, W R and Taylor, A (2003) 'The structure of credit risk: spread volatility and ratings transitions', The Journal of Risk, Vol. 6 No. 1, Fall.

[17] Lopez, J A and Saidenberg, M R (2000), 'Evaluating credit risk models', Journal of Banking and Finance, 24(1-2), pages 151-65.

[18] Lucas, D J and Lonski, J G (1992), 'Changes in corporate credit quality 1970-1990', Journal of Fixed Income, pages 7-14.

[19] Mingo, J (2000), 'Policy implications of the federal reserve study of credit risk models at major US banking institutions', Journal of Banking and Finance, 24(1-2), pages 15-33.

[20] Morris, C, Neal, R and Rolph, D (1998), 'Credit spreads and interest rates', mimeo, Indiana University, Indianapolis.

[21] Nickell, P, Perraudin, W R and Varotto, S (2000), 'The stability of ratings transitions', Journal of Banking and Finance, 24(1-2), pages 203-25.

[22] Perradin, W R and Taylor A (2003) 'Liquidit and Bond Market Spreads', Manchester Business School working paper.

[23] Pritsker, M (1996), 'Evaluating value at risk methodologies: accuracy versus computational time', Working paper, Federal Reserve Board, Washington, DC.

[24] Wilson, T (1997), 'Credit risk modelling: a new approach', unpublished mimeo, McKinsey Inc., New York.

Table 1
Total Portfolio Characteristics

Domicile	No.	%	Sector	No.	%	Seniority	No.	%
US	618	47.76	Financial Services	567	43.82	Unsecured	742	57.34
Japan	124	9.58	Banking	292	22.57	Guaranteed	262	20.25
Netherlands	102	7.88	Utilities, Elect.+Gas	65	5.02	Senior Unsecured	183	14.14
Germany	74	5.72	Energy Sources	43	3.32	Government Guaranteed	50	3.86
UK	63	4.87	Telecommunications	40	3.09	Secured	22	1.70
France	61	4.71	Beverage+Tobacco	31	2.40	Subordinated	20	1.55
Neth. Antille	40	3.09	Health+Personal Care	27	2.09	Mortgaged	11	0.85
Cayman Is.	24	1.85	Bus.+Public Services	27	2.09	Collateralised	4	0.31
Mexico	24	1.85	Merchandizing	27	2.09			
Australia	20	1.55	Food+Hshld Product	20	1.55			
Other	144	11.13	Multi-Industry	18	1.39			
			Other	137	10.59			

Note: The whole sample consists of all, liquid dollar-denominated Eurobonds available in each period.

Figure 1: Distribution of Alphas

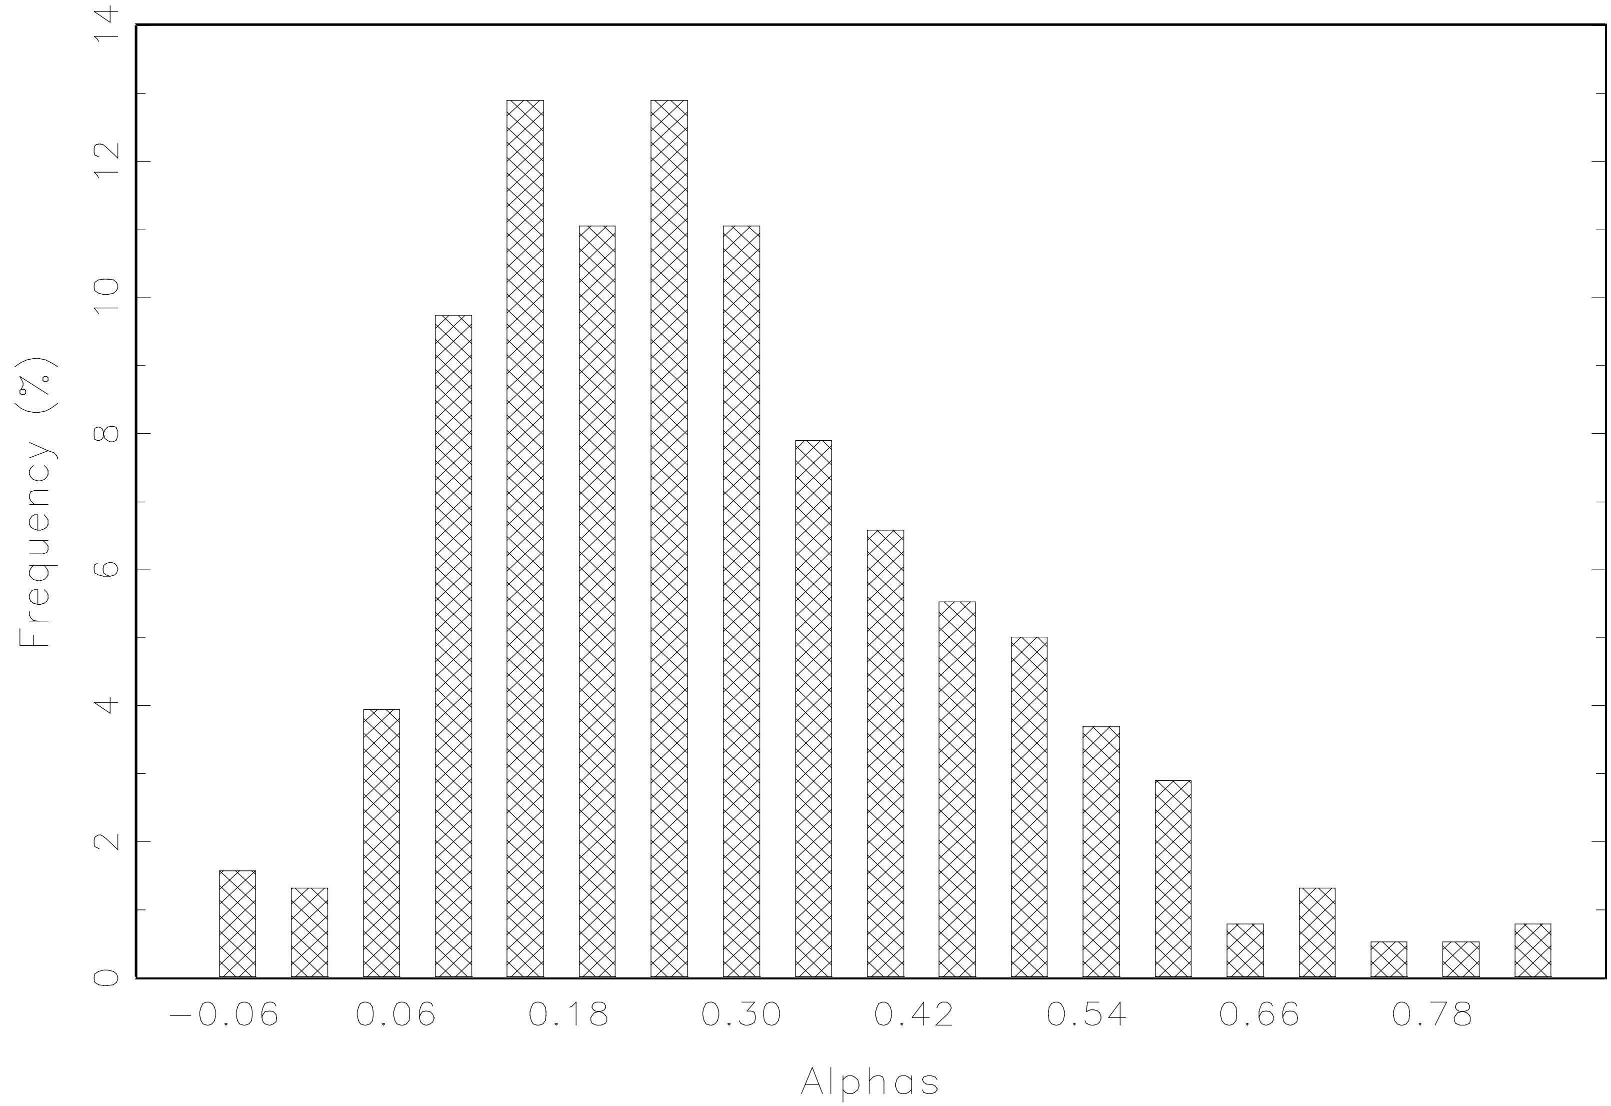

Figure 2: Total Portfolio Returns and Monte Carlo VaRs

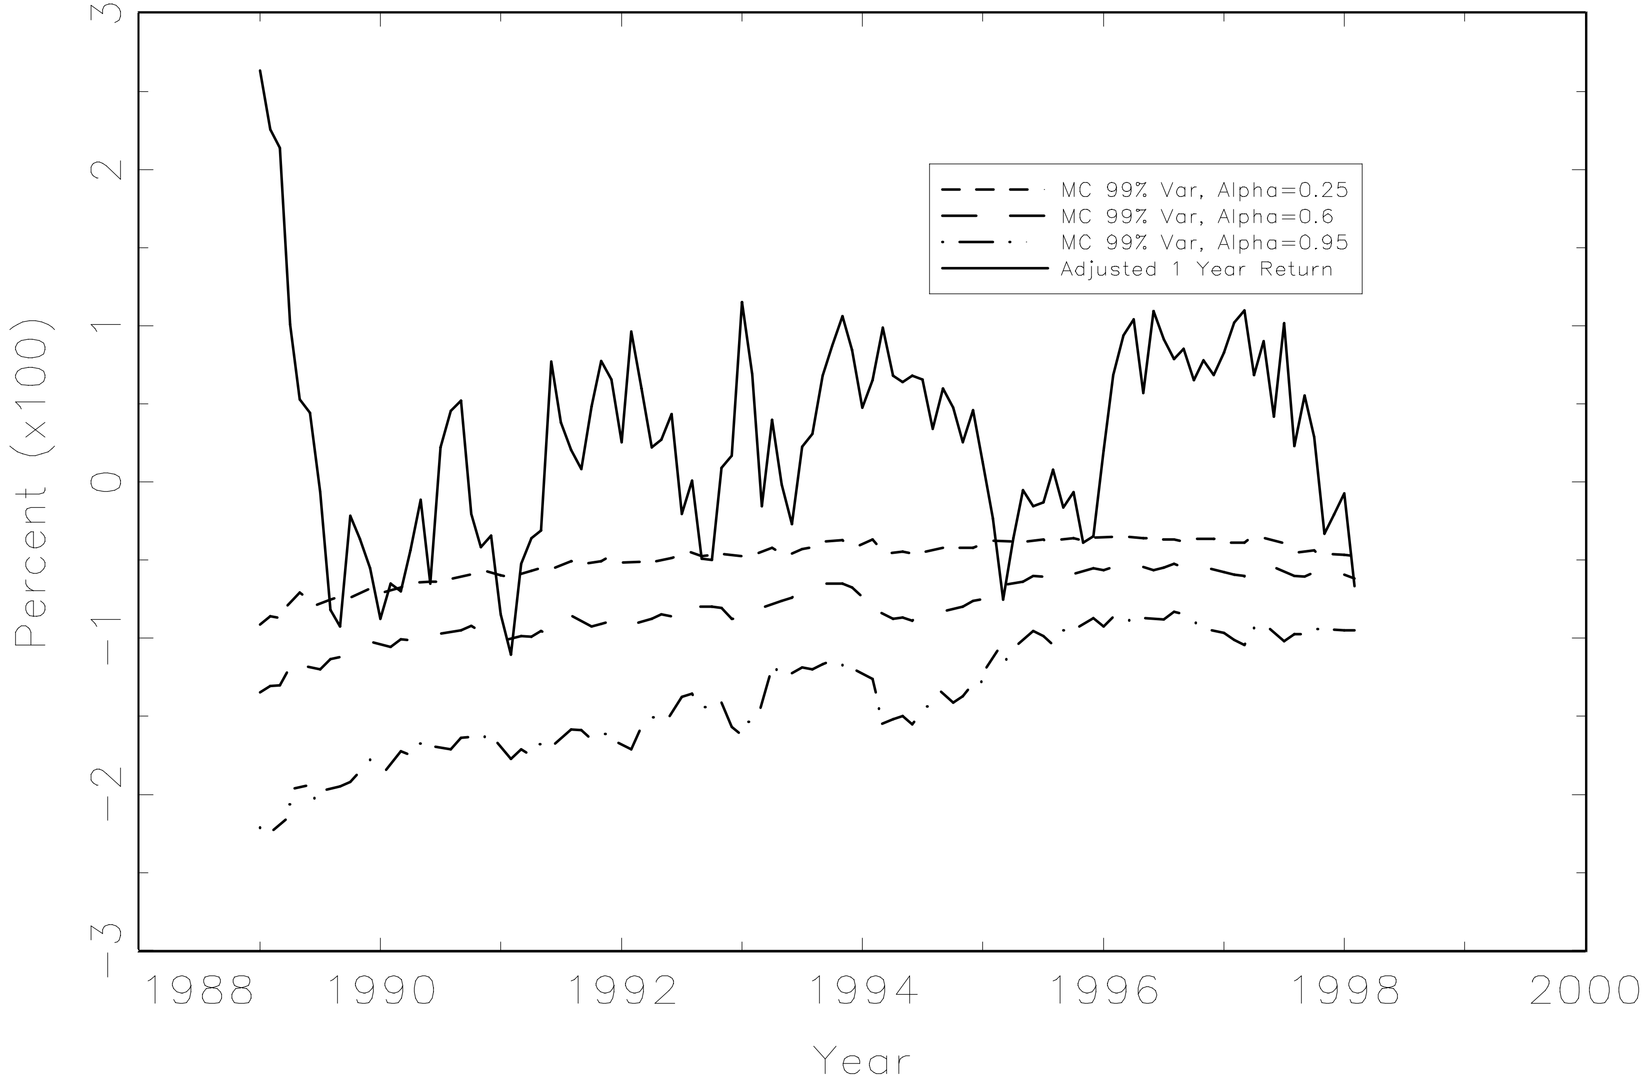

Figure 3: Normal versus Monte Carlo Based VaRs

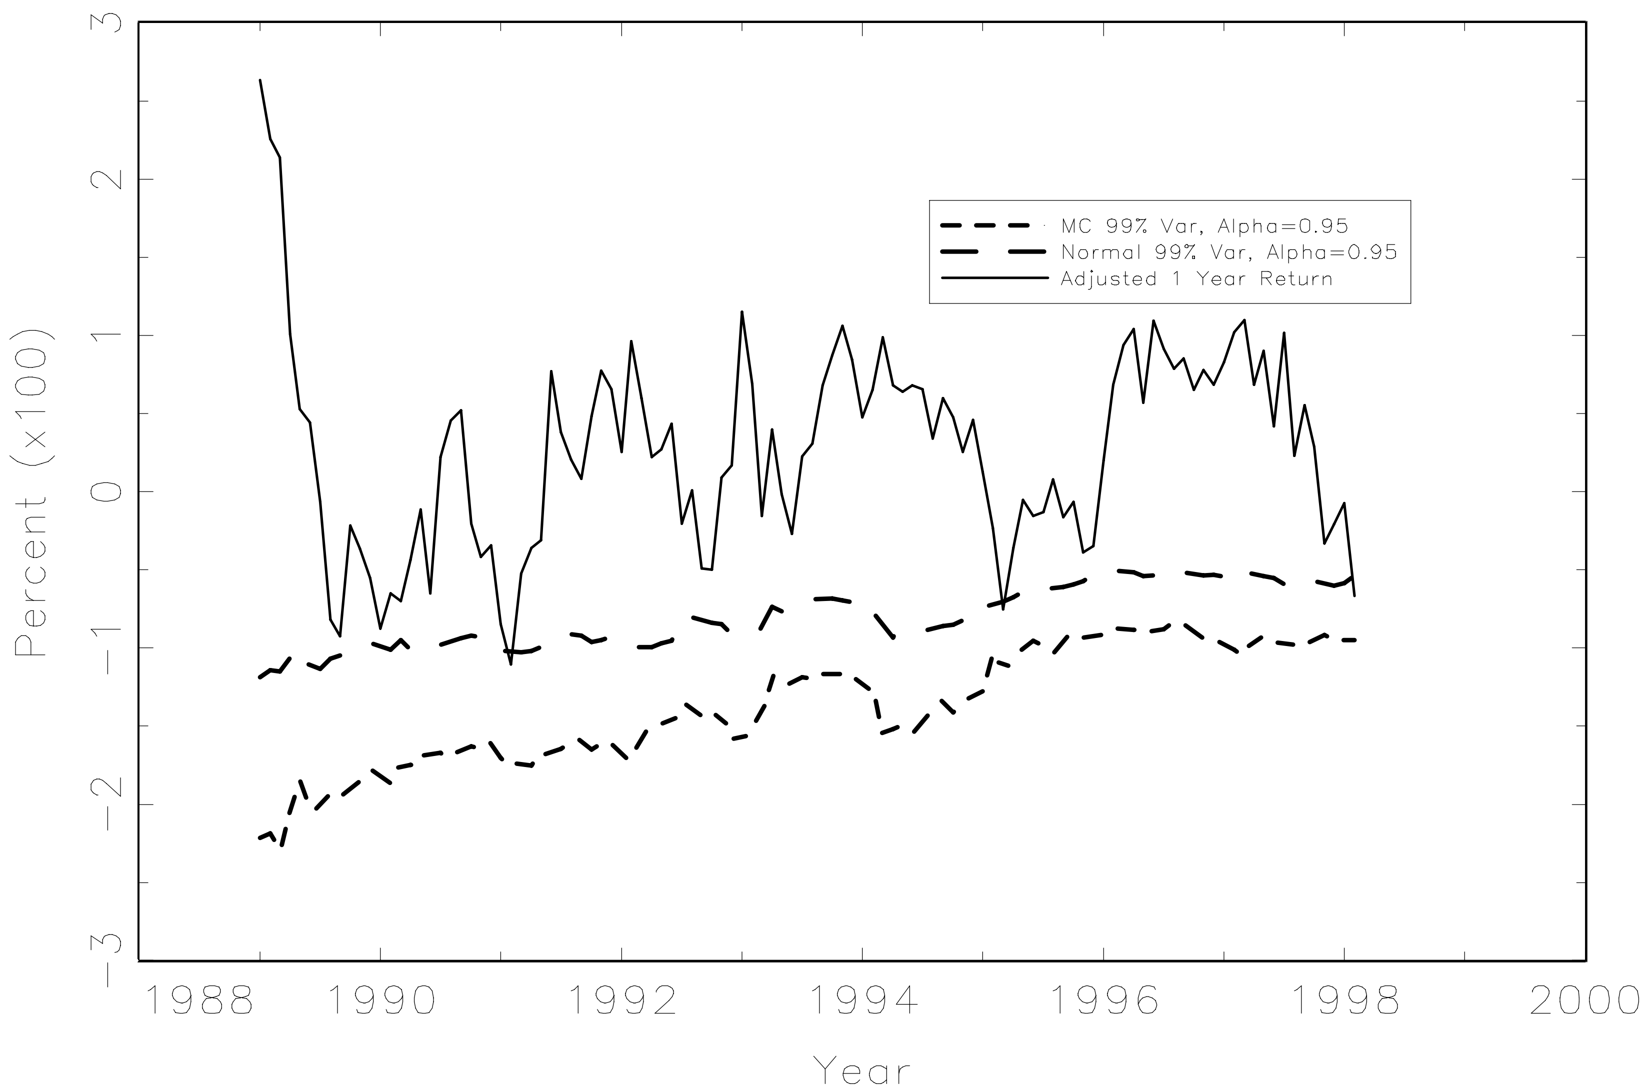

Figure 4: Industry and Domicile: Monte Carlo VaRs

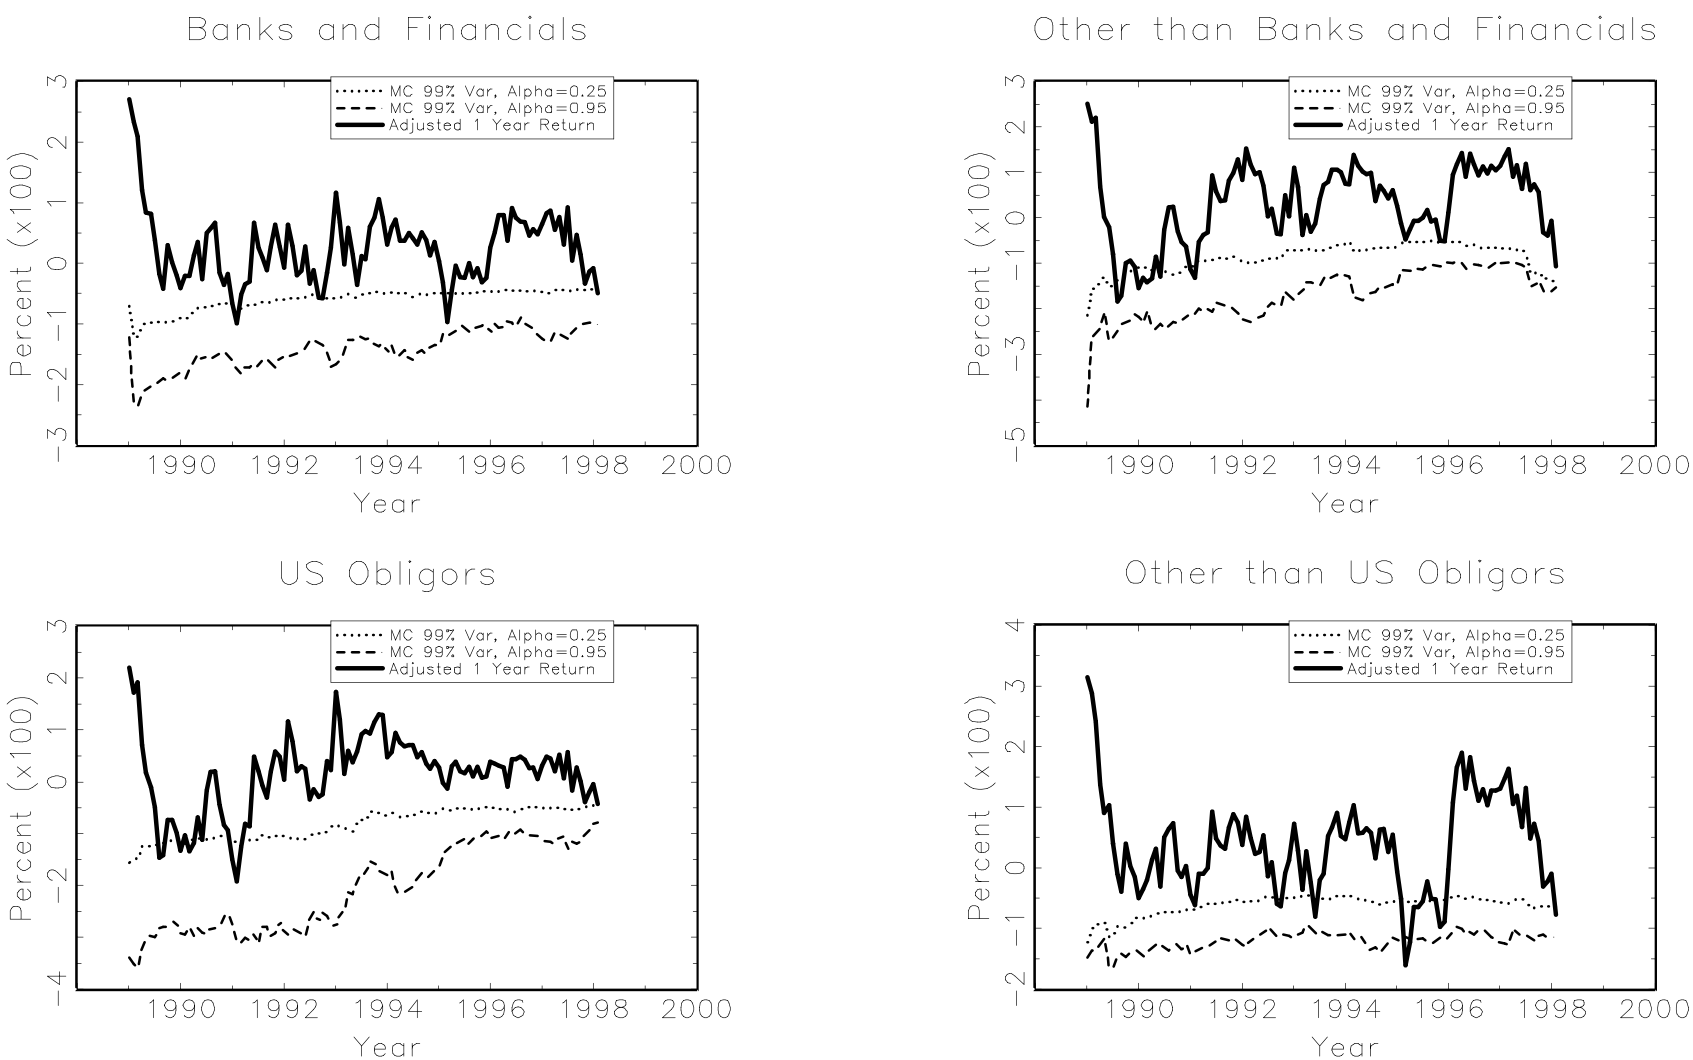

Other than US Obligors

Figure 5a: 50% Recovery rate

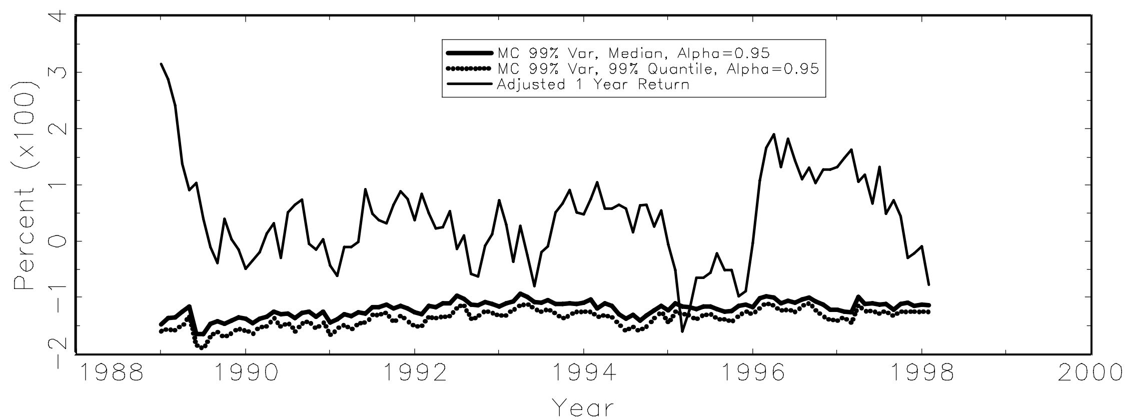

Figure 5b: 0% Recovery rate

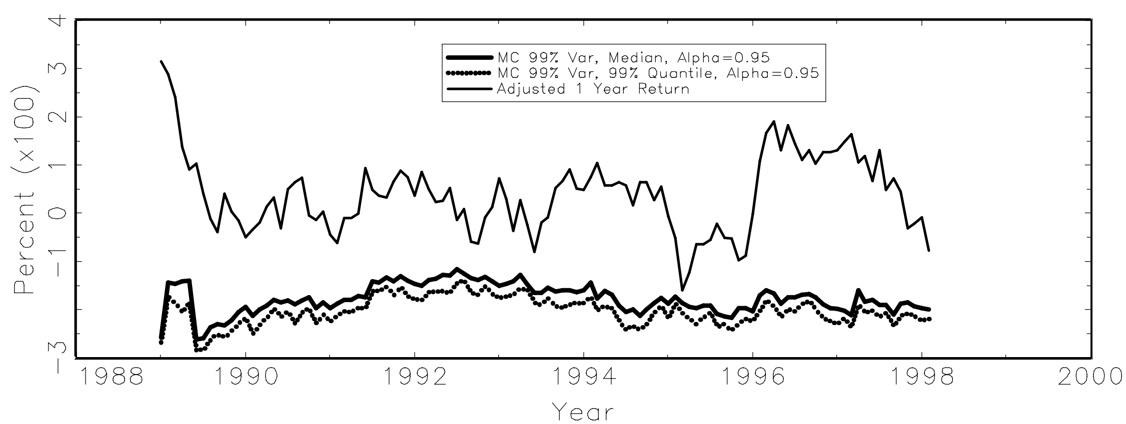

Figure 6: Quartile Portfolios: Monte Carlo VaRs

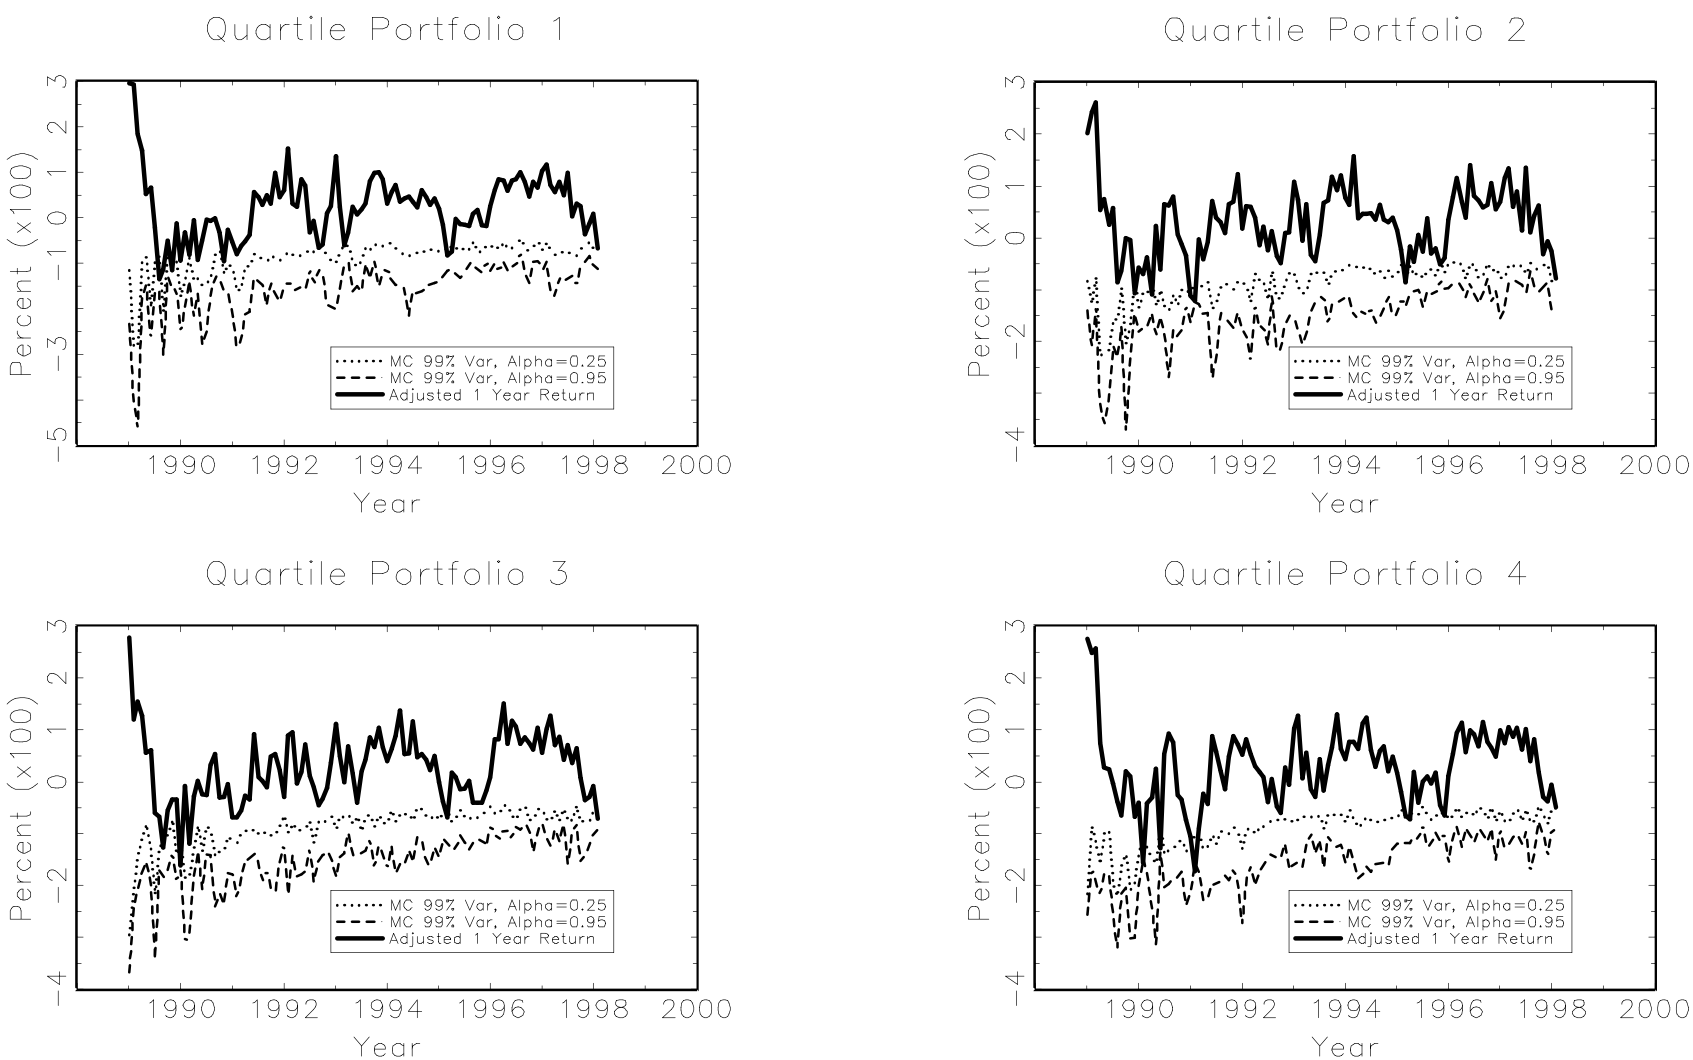

Quartile Portfolio 4

Figure 7a: 50% Recovery rate

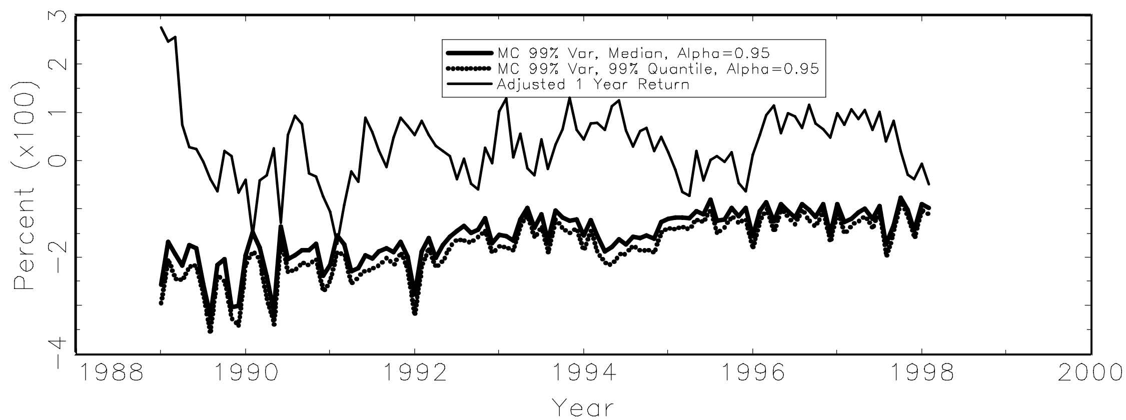

Figure 7b: 0% Recovery rate

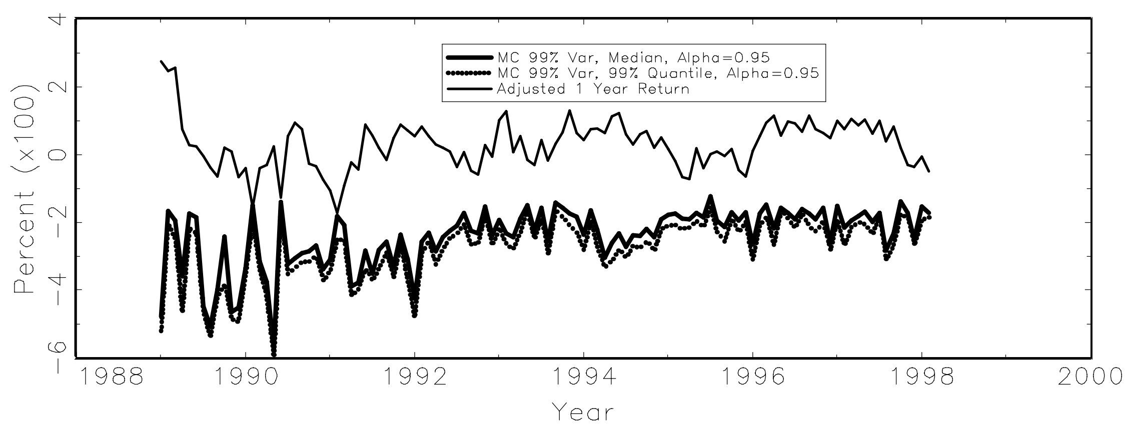

Fifty Lowest Rated Bonds

Figure 8a: 50% Recovery rate

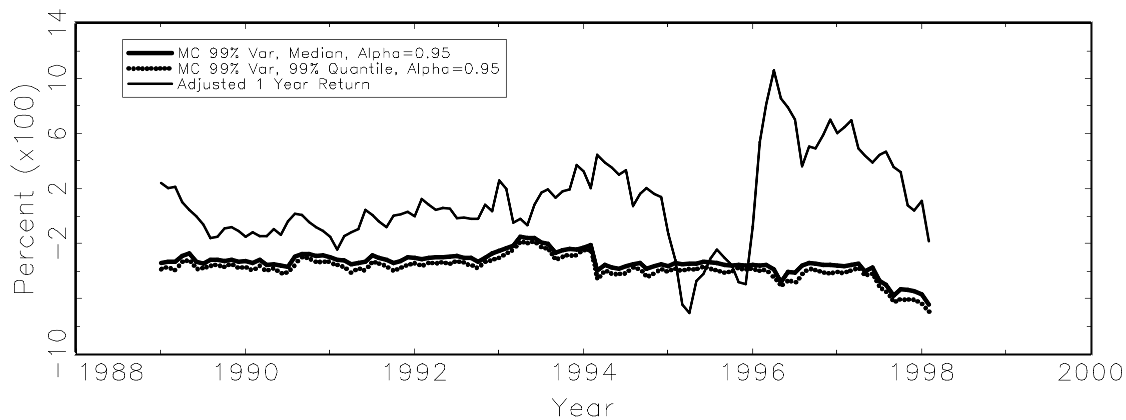

Figure 8b: 0% Recovery rate

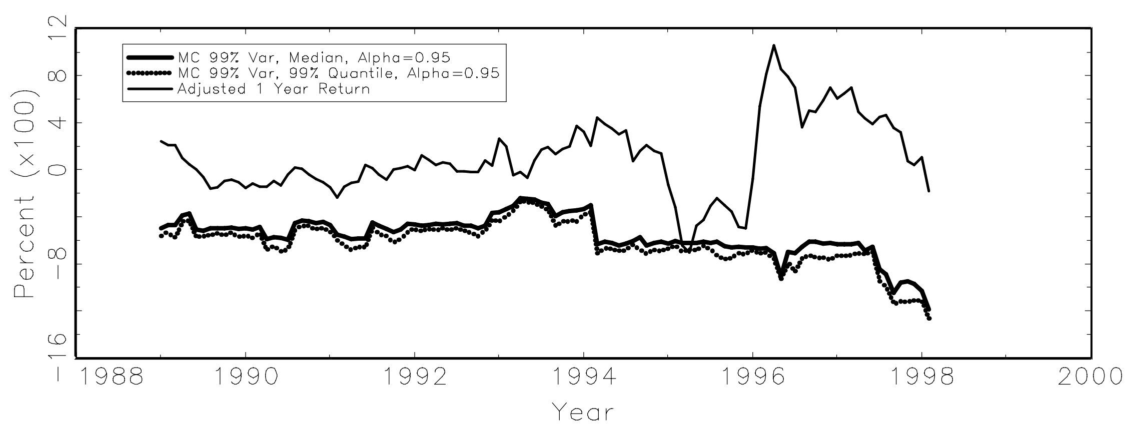

Fifty Highest Rated Bonds

Figure 9a: 50% Recovery rate

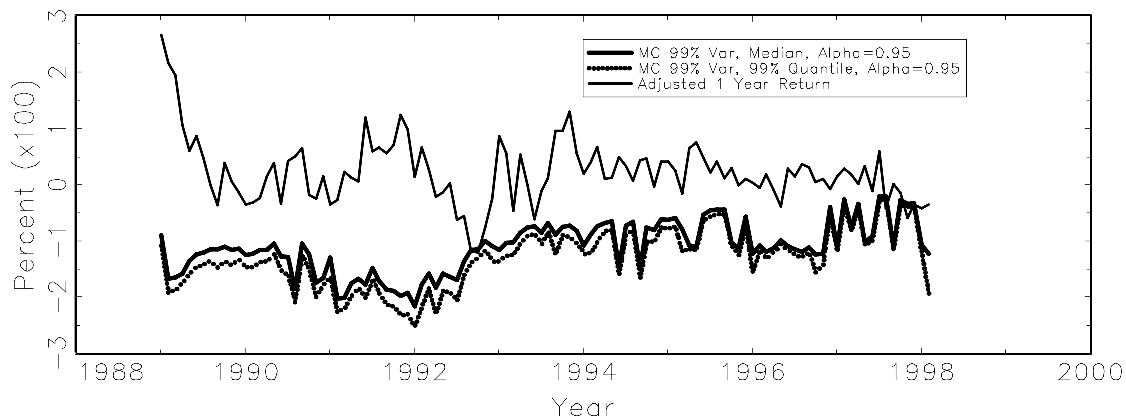

Figure 9b: 0% Recovery rate