THE CYCLICAL BEHAVIOR OF DEFAULT AND RECOVERY RATES

Kyriakos Chourdakis*

FitchSolutions and CCFEA, Essex University

Abstract

Using a regime switching framework we investigate the determinants of default clustering. We find that a common credit cycle, modeled as a two-state Markov chain, can account for a large portion of default correlations, with the residual clustering being captured by a factor structure. During credit crunches default rates increase, and so does the conditional residual correlation. Using data for the period 2000-2006 we find that this common cycle is robust in explaining a large part of the default behavior of individual industries and recovery rates. We also find that structured products that depend on the distribution tail are the first to respond to a deterioration of the credit environment.

1 *Introduction*

Loss distributions of portfolios that are subject to credit events are sensitive to a number of factors and their interplay. The patterns of defaults are cyclical, suggesting that a dynamic model of default intensities will be better suited to explain the variability of credit losses. Furthermore, it is natural to assume that the propensity to default across different names is driven by common factors: as the credit cycle deteriorates more names default, a feature that is clearly illustrated by the common movements of credit spreads.

Default probabilities evolve through time, not only due to a global "credit cycle", but also due to events that are country- or industry-specific. Default correlations can be generated by such correlated hazard rate paths, even if the events themselves are independent when conditioned on each path. Nevertheless, recent empirical studies suggest that this is insufficient: default events appear to be correlated to a degree that extends beyond the one that can be explained by common spread moves (Das, Duffie, Kapadia, and Saita, 2007).

Here we develop, implement and test a simple dynamic model of hazard rates which also allows for conditional default correlations. Although there is a

*The author wishes to thank Gareth Stoyle, Matthias Neugebauer, Damiano Brigo, Fares Triki and Dirk Tasche for useful discussions. Email: kyriakos.chourdakis@fitchsolutions.com

multitude of models for assets that are subject to default, most focus in interpolating across various (sometimes quite illiquid) assets to recover no-arbitrage prices for nonstandard bespoke contracts. Academic research on the nature of default correlations and probabilities is scarce. As Darrell Duffie points out in a 2004 FT article, "[credit risk] is one area of finance where our ability to structure financial products may be running ahead of our understanding of the implications" (Duffie, 2004).

The seminal work of Merton (1974) recognizes that default is essentially an option on the assets of the firm, held by the equity holders. In this framework, the firm will choose to default if the value of its assets declines below a threshold which is determined by the debt burden. Setting the dynamics of the firms asset value, together with the dependence structure between asset values of different firms gives rise to models of correlated defaults. The simplest case of such a *structural form* model, corresponds to the popular Gaussian copula model (see Li, 2000). In this specification a Gaussian distribution dictates the firm value, and default can happen on maturity only. Extensions in this framework have either taken the form of specifying different distributional forms,[1] or relaxing the assumption of default on the terminal date.[2] We show that a large class of these models can be though of as conditional Gaussian copula models, if we allow for a randomization of the default probability. We then build on a recent model by Wang et al. (2006) and consider a Gaussian copula model that exhibits correlations and default probabilities that depend on an underlying random state. Such a model allows for Large Homogeneous Portfolio (LHP) formulas that are very simple to implement.

A second family of models focuses on hazard rates, rendering them stochastic. In particular, hazard rates are modeled as a Cox process, that is to say default is identified as the first jump of a Poisson process with stochastic intensity. Correlation of defaults is induced by the common behavior of hazard processes, as firms will tend to default in clusters with increasing hazard rates.[3] As we have already built a copula and default probability structure that is dependent on an underlying state, we render the state dynamic using a regime switching specification. In this way we bridge the benefits and intuition of a static Gaussian copula framework with a dynamic process that attempts to mimic the credit cycula. The regime switching structure ensures that in doing so we don't sacrifice any of the tractability of the Gaussian framework.

We estimate the parameters of a two-state regime switching process on market implied default probabilities (MIDP) provided by Fitch-QFR.[4] The data

[1] For examples see Hull and White (2004), Kalemanova, Schmid, and Werner (2005), and Moosbrucker (2005). Wang, Rachev, and Fabozzi (2006) give a detailed overview of the recent literature.

[2] This introduces a barrier option to default rather than a plain vanilla one, as in Black and Cox (1976).

[3] Early examples of this *reduced form* approach include Madan and Unal (1998), Jarrow and Turnbull (1995) and Lando (1997, 1998).

[4] Although the default probabilities are market-implied, they are not regarded under the risk neutral measure. See Liu, Kocagil, and Gupton (2007) for details on the construction methodology of market implied quantities.

consist of one- and five-year implied default rates that span 16 years (1990-2006) and are collected monthly. The default probabilities are viewed as losses of a large well diversified portfolio with zero recovery, and the likelihood function is based on the conditional LHP density. This allows us to estimate not only the two distinct default probabilities that prevail over different parts of the credit cycle, but also different correlation structures. The estimation procedure identifies two distinct and very persistent regimes: the "bad times" during the market turmoil of 2000-3, and the "good times" before and after that period. We also find that during credit crunches default correlation increases substantially, causing the loss distribution to flatten further and shifting probability mass to extreme adverse events.

A two regime specification is pursued in Das and Geng (2004) who find that a regime structure is an important specification of defaults. Unlike Das and Geng (2004) we do not split the data in two regimes *ad hoc*, but allow for the correct structure to be endogenously determined. An important byproduct of our methodology is the conditional probability of being in the "bad" or "good" regime at each point in time, which serves as the weighting factor of the LHP densities. A natural interpretation of this probability is as the position of the market in the credit cycle. Bruche and González-Aguado (2007) also take an approach that utilizes regime switching, but defaults in their model are conditionally independent resembling a discretized hazard rate specification. For that reason their approach does not allow one to make inference on the correlation structure of defaults within each regime.[5]

Having established the path probabilities of credit cycle movements, we drill further into investigating the conditional behavior of individual industries. We extend the factor structure of the copula to allow for inter-industry correlation uplifts. The estimation is then carried out conditional on the regime switching path that was filtered out of the aggregate default frequencies. We find that overall, industry data support the presence of a global factor: allowing the copula parameters to vary across the cycle has a dramatic effect on the likelihood value, when compared to a single copula specification.

Recovery rates are another important ingredient of the loss distribution. There has been a growing body of literature not only on the distribution of recovery rates across seniorities and industries, but also on the relationship between recovery rates and the credit cycle.[6] We confirm the recent empirical literature which suggests that recovery rates are negatively correlated with default rates, by setting the recovery distribution to depend explicitly on the underlying state variable.

The paper is organized in the following way: the next section outlines the copulas used, and provides motivation for the use of a two-regime framework.

[5] In the Bruche and González-Aguado (2007) framework the conditional LHP portfolio will be distributed as a Dirac mass on the default probability point, if recovery rates are zero. Some variability can be introduced by rendering recoveries stochastic.

[6] Recent work on recovery models and the determinants of recovery can be found in (Frye, 2000a,b), Hu and Perraudin (2002), Altman, Resti, and Sironi (2004), Altman, Brady, Resti, and Sironi (2005) and Gupton and Stein (2005).and the references therein.

Section 2 introduces the idea of a conditional Gaussian copula and illustrates the connections with other popular families of copulas. In section 3 we explain the regime switching framework, and give a detailed overview of the estimation procedure. We investigate the temporal structure of recovery rates in section 4, while section 5 presents a small Monte-Carlo experiment that investigates the impact of credit cycle movements on different tranches. The last section concludes.

2 The copula structure

In this section we introduce the state dependent copula that will be subsequently cast in a dynamic setting. We also show how this copula can be seen as a generalization of families of copulas that are commonly used.

2.1 Gaussian copulas and extensions

In the Gaussian copula framework we define a *default driver* Z that serves as a proxy of the asset value. In particular we set up a structure which depends on a number of independent Gaussian factors X_j. For the n-th asset we write

$$Z_n = \sum_j \sqrt{\alpha_{\ell,j}} X_j$$

We normalize the variance to be equal to one by requiring the factor loadings to satisfy $\sum_j \alpha_j = 1$. The name will default if the realization $Z_n < C_n$, with C_n a predetermined threshold. This approach, motivated by Merton's model for corporate defaults, was popularized by Li (2000) and is currently the benchmark for the analysis of credit-risky securities portfolios.

The default drivers of the n-th and m-th assets will be correlated, as

$$\mathrm{E}\{Z_n Z_m\} = \sum_j \sqrt{\alpha_{n,j}\alpha_{m,j}}$$

Typically, the first factor is "global" and its realization will affect all assets in the portfolio, that is to say $\alpha_{n,1} \neq 0$ for all n. One can then easily set uplifts that are domicile- or industry-specific. For example, one can put $\alpha_{n,k} \neq 0$ for all names that belong to the same industry, and $\alpha_{n,k} = 0$ otherwise. Then the k-th factor will be associated with this particular industry.

A popular special case is where there is a single common factor with constant loading, where the drivers are given by

$$Z_n = \sqrt{\alpha} X + \sqrt{1-\alpha} X_n$$

Then, the correlation of the drivers is equal to α for all name pairs. In this special case one can recover the loss distribution for a large homogeneous portfolio (LHP). In particular, given the value of X all names default independently, and therefore a large portfolio will realize a loss which is exactly

$$L|X = \Phi\left(\frac{C - \sqrt{\alpha}X}{\sqrt{1-\alpha}}\right)$$

Therefore, the loss distribution will be equal to

$$\mathrm{P}\{L \leqslant \ell\} = \mathrm{E}_X\{\mathrm{P}\{L \leqslant \ell | X\}\}$$

As the loss given X is deterministic, the probability $\mathrm{P}\{L \leqslant \ell|X\}$ will be either one, if $L|X \leqslant \ell$, or zero otherwise. Thus, the loss distribution can be rewritten as

$$\mathrm{P}\{L \leqslant \ell\} = \mathrm{P}_X\left\{\Phi\left(\frac{C - \sqrt{\alpha}X}{\sqrt{1-\alpha}}\right) \leqslant \ell\right\} = \Phi\left(\frac{\Phi^{-1}(\ell)\sqrt{1-\alpha} - C}{\sqrt{\alpha}}\right)$$

The failure of the Gaussian copula to generate loss distributions that match the historical experience has led to the introduction of models where the drivers exhibit fat tails. Popular extensions include drivers with Student-t or Lévy distributions, such as the Variance Gamma (VG) or the Normal Inverse Gaussian (NIG). We will see in the next subsection that such models can be viewed as *conditional* Gaussian copula models.

2.2 Conditional Gaussian copulas

For a large family of tractable factor models the default drivers can be cast in the form

$$Z_n = \varpi\sqrt{Y_n}Z_n^\star + \vartheta Y_n + \varphi$$

where $Z_n^\star$ is a standardized Gaussian random variable. A name would then default if the driver takes a value below a threshold C_n that is determined by the probability of default, or

$$\text{name defaults} \Leftrightarrow Z_n \leqslant C_n$$

If we consider symmetric distributions with fat tails, then $\vartheta = 0$, while setting $\vartheta \neq 0$ defines skewed drivers. In particular, the parameters φ and ϖ are chosen in a way that ensures the normalization $\mathrm{E}\{Z_n\} = 0$ and $\mathrm{E}\{Z_n^2\} = 1$. This amounts to

$$\begin{aligned} \varphi &= -\vartheta\mu_Y \\ \varpi &= \sqrt{\frac{1 + \vartheta^2\mu_Y^2 - \vartheta\sigma_Y^2}{\mu_Y}} \end{aligned}$$

For example, in the t-copula case, ν/Y^2 is distributed as a χ^2_ν distribution, and $\vartheta = 0$. In this case $\mu_Y = \sqrt{\nu/(\nu-2)}$, and therefore $\varpi = \sqrt{(\nu-2)/\nu}$. Since every Lévy process is essentially a time-changed Brownian motion, the NIG and VG copulas are defined with Y following the Gamma or the Inverse Gaussian distribution, respectively.

One then introduces a standard Gaussian factor structure on $Z_n^\star$

$$Z_n^\star = \sum_j \sqrt{\alpha_{n,j}} X_j$$

Then, the default criterion can be written as

$$\text{name defaults} \Leftrightarrow \sum_j \sqrt{\alpha_j} X_j \leqslant \frac{C - \varphi - \vartheta Y}{\varpi \sqrt{Y}} = C(Y)$$

This representation allows us to view such models as Gaussian copula specifications, where the default probability is randomized by the variable Y. In fact, the LHP forms of those models will reflect that, as they are a weighted average of Gaussian copula LHPs, with weights defined by the randomized default threshold. In particular, for the one factor case we can write the distribution for the loss L as

$$\begin{aligned} \mathrm{P}\{L \leqslant \ell\} &= \mathrm{E}_Y \left\{ \Phi\left(\frac{\Phi^{-1}(\ell)\sqrt{1-\alpha} - C(Y)}{\sqrt{\alpha}} \right) \right\} \\ &= \int_{\mathbf{R}+} \Phi\left(\frac{\Phi^{-1}(\ell)\sqrt{1-\alpha} - C(y)}{\sqrt{\alpha}} \right) \mathrm{d}F_Y(y) \end{aligned}$$

There are two elements that determine the loss distribution in these models: the asset correlation and the distribution of the default probabilities. Observe that in the limit, where the names are uncorrelated, the loss distribution is a simple transformation of the distribution for Y

$$\mathrm{P}\{L \leqslant \ell\} = \mathrm{P}\left\{C(Y) \leqslant \Phi^{-1}(\ell)\right\}$$

This indicates that if we depart from the Gaussian copula, then we can match any loss distribution with a set of uncorrelated (but certainly dependent) drivers.

Using a specific copula (like the t, VG or NIG) sets a specific structure on the distribution of Y, which overall will not resemble the historical default experience. It is therefore appealing to consider a simple discrete state distribution for Y, which will be flexible enough to be calibrated to a given loss distribution. In addition, in a discrete state setting we can also render the correlations dependent on Y. This can capture scenarios where higher default probabilities are accompanied by higher correlations of the underlying assets.

For example, if we postulate $\theta = 0$ and Y taking values y_k with probabilities p_k, we can write the loss distribution as

$$\mathrm{P}\{L \leqslant \ell\} = \sum_k \Phi\left(\frac{\Phi^{-1}(\ell) \cdot \sqrt{1-\alpha_k} - C_k}{\sqrt{\alpha_k}} \right) p_k$$

Such an expression can replicate the features of the Lévy copula models discussed above, using a model that is very flexible and fast to calibrate. Historical default rates do not vary wildly across time, exhibiting instead well defined cyclical patterns. The discrete space model can be simulated in a way that can maintain these dynamical patterns, for example by setting a Markov chain that controls the transitions across states.

	correlation	threshold	fit	other parameters
Gaussian	5.64%	-2.413	167.86	
NIG	5.66%	-2.414	167.62	$\alpha = 9.22$
	10.99%	-2.728	123.72	$\alpha = 19.64$, $\beta =$ -18.59
Gaussian mix	3.30%	-2.211	47.66	$p = 44.37\%$
	0.39%	-2.633		

Table 1: Calibrated parameters for default frequencies. The objective function that was minimized is the sum of squared differences between the actual and the theoretical cumulative densities. The column labeled 'fit' gives the minimum value of this objective function.

3 *Data and estimation methodology*

We use the market implied default probabilities (MIDP) supplied by Fitch-QFR. They are recorded monthly from 1990 to the end of 2005, and give the default frequency embedded in equity prices recovered using the QFR methodology (see Liu et al., 2007, for details). The data set includes aggregate one- and five-year default probabilities, and also default probabilities across different industries. Table 2 gives a list of the data series. For comparison purposes, figure 1 presents the time series of the MIDPs, together with the actual default rates of investment grade bonds collected by Standard and Poor's. For the analysis on recovery rates we use the 30-day post default prices of all rated corporate bonds that defaulted in the period 2000-2006.

Initially we take a "static" look at the data, and fit the Gaussian copula LHP loss distribution, together with the NIG and the Gaussian-mixture copula. After having established the superiority of the mixture approach, we cast the dynamics of the underlying state in the regime switching framework.

3.1 Unconditional estimation

The initial calibration is based on the historical distribution of the MIR series. In particular, we are looking for the parameter set in a one factor setting that will replicate the distribution of default rates. As we want to focus on the distribution tails, we minimize the distance between the empirical and theoretical cumulative distribution functions.

At this stage we look solely at the unconditional distribution of default rates, and ignore possible dynamics. For that reason we coin the calibration methodology "unconditional". In the next subsection we will take a closer look at the features of default frequencies through time. Also, in this subsection we focus on the distribution of aggregate default rates across all sectors.

The results of this calibration exercise are given in table 1. The Gaussian copula calibrates with a correlation of 5.64%, and a threshold that is roughly at 80bps. This is in line with the historical average default rate, as given by the MIDP data set.

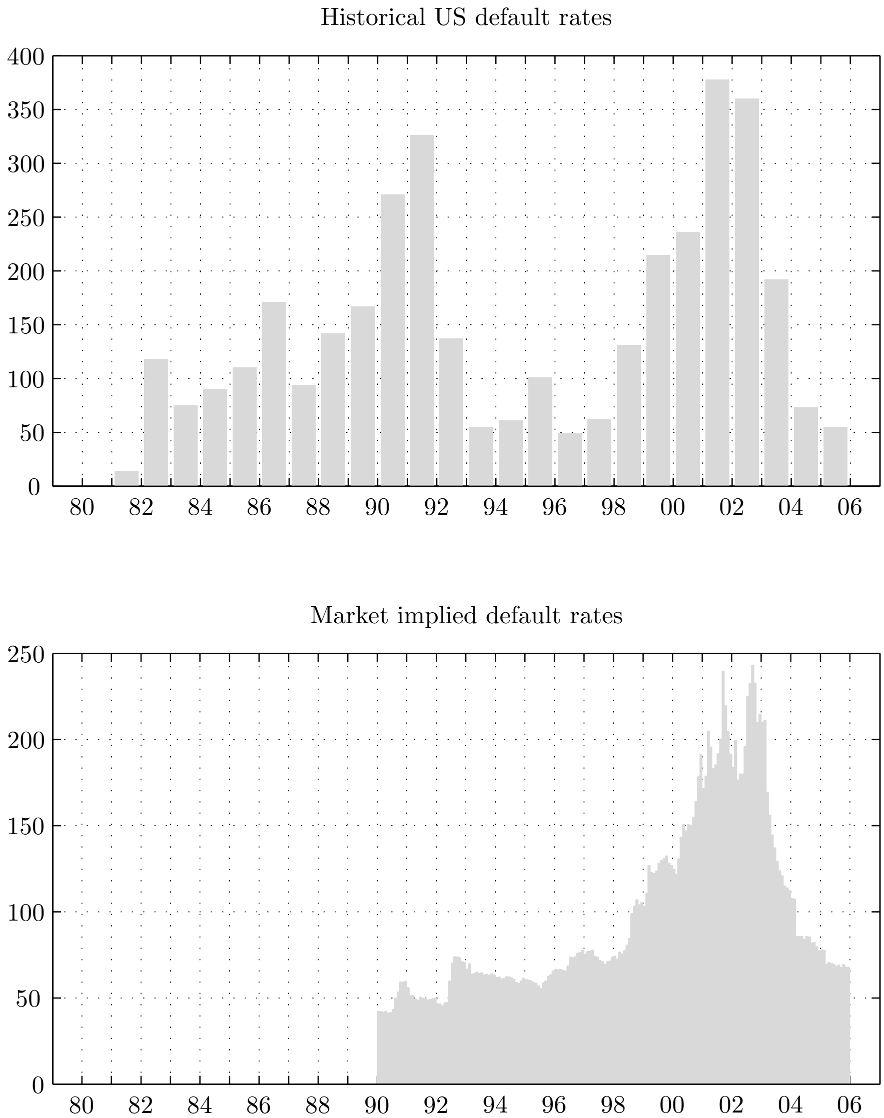

Figure 1: Default rates: historical and market implied.

Calibrating the NIG copula introduces extra parameters. Parameter α controls the kurtosis of the distribution, while β can be used to set the skewness. We calibrate the symmetric version, and the resulting distribution is virtually indistinguishable from the Gaussian. When we allow for skewness, the fit is somewhat improved, with a driver distribution that is heavily skewed ($SK = -0.65$) and slightly leptokurtic ($KU = 3.37$).

Both copulas fail to capture the salient features of the default rate distribution. These families of copulas implicitly assume that defaults occur in a manner that is stable through time. A casual inspection of the historical default rates reveals that this is not the case: defaults rates exhibit a cyclical pattern, with well defined regimes. In the MIDP sample, default rates increased dramatically after 1999 from about 50bps to over 150bps. They remained at their high levels until 2004, subsequently reverting back to their previous levels. The copula models fitted above cannot account for this sudden switches, and spuriously attempt to fit the resulting distribution with higher correlation values.

The Gaussian-mix process is by construction designed to simultaneously calibrate to the "good" times before 1999, and the "bad" times afterwards. In the good times, the calibrated default probability is low (47bps) and so is the underlying asset correlation (0.39%); in the bad times, higher default probabilities (135bps) are accompanied by higher asset correlation (3.30%).

3.2 A regime switching framework

In the previous subsection we demonstrated the superiority of the Gaussian-mix model, against other fat-tailed alternatives, in capturing the salient features of historical default frequencies. The cyclical variation of historical default rates motivates us to view the Gaussian-mix as the unconditional density of a regime switching model that alternates between "good" and "bad" times of the credit cycle. It is therefore natural to attempt to calibrate the parameters of the Gaussian-mix model under this framework, using maximum likelihood estimation.

We adopt a two-regime structure with a base global correlation, and an uplift that is industry specific. Say that the state of the world during period t is s_t. The default driver of an asset in the n-th industry will be given by

$$Z_n = \sqrt{\alpha_G(s_t)}X_G + \sqrt{\alpha_n(s_t)}X_n + \sqrt{1 - \alpha_G(s_t) - \alpha_n(s_t)}\epsilon$$

and the asset will default if $Z_n \leqslant C_n(s_t)$. In the above relationship all α_G, α_n are non-negative. Therefore, assets that belong to different industries will exhibit correlation $\alpha_G(s_t)$, while assets in the n-th industry will have a higher correlation $\alpha_G(s_t) + \alpha_n(s_t)$.

3.3 Global parameters and the regime structure

To calibrate the above model we use a two-step procedure. First, we use the aggregate default frequencies and compute the global correlations α_G, together with the probabilities of remaining in the two regimes, p_{11} and p_{22}. We use

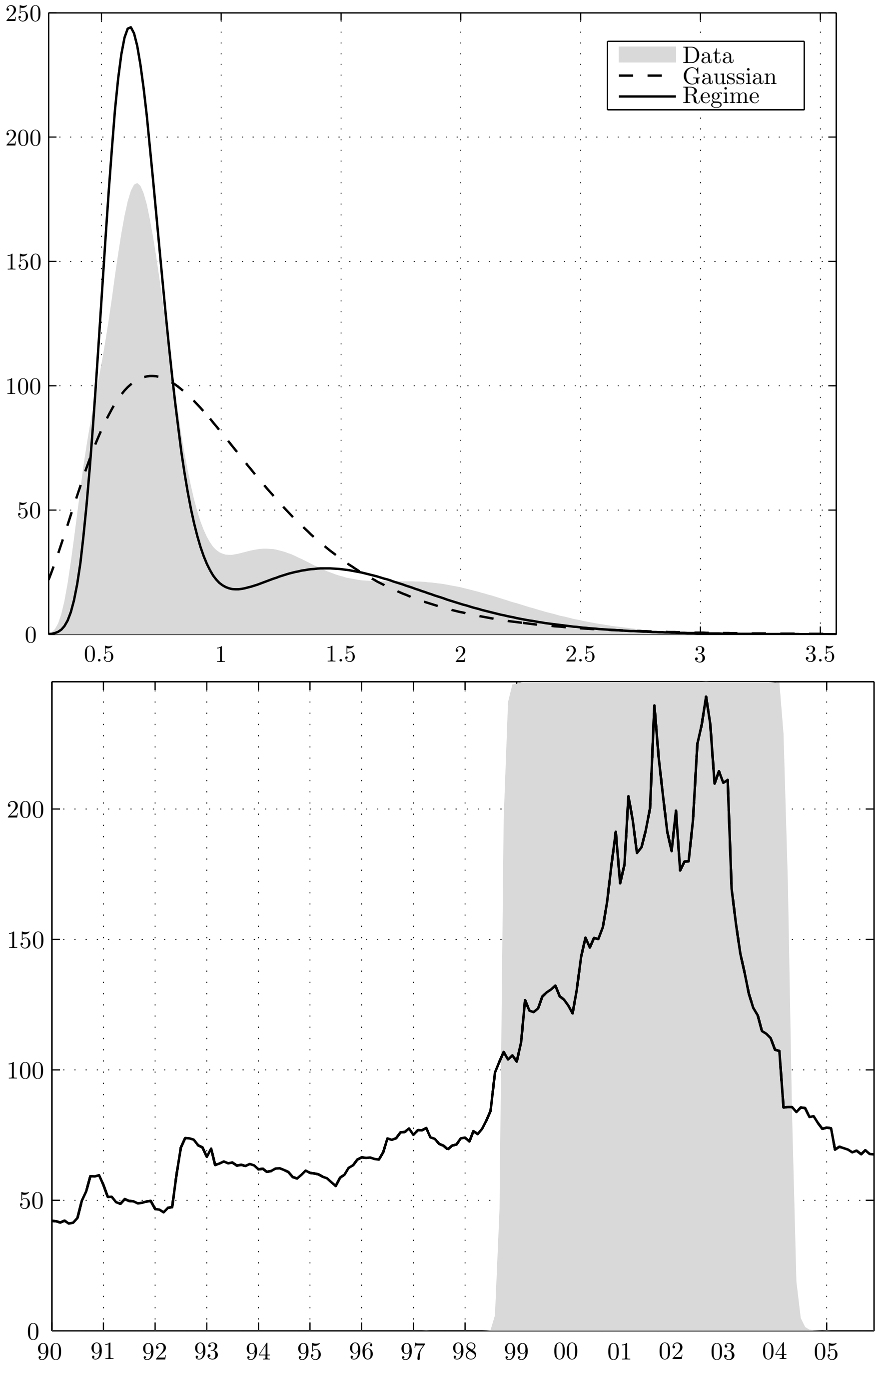

Figure 2: Estimation of the parameters for the regime switching loss distribution model, based on the aggregate ("Global") one year default probability. The shaded area in the top figure gives a non-parametric estimator for the loss distribution. The dashed line gives the loss distribution of a LHP where the factors are Gaussian. The solid line gives the unconditional loss distribution for the two-state regime switching loss model. The raw data are given in the bottom figure, together with the shaded area that represents the probability of being in the regime with high default probabilities.

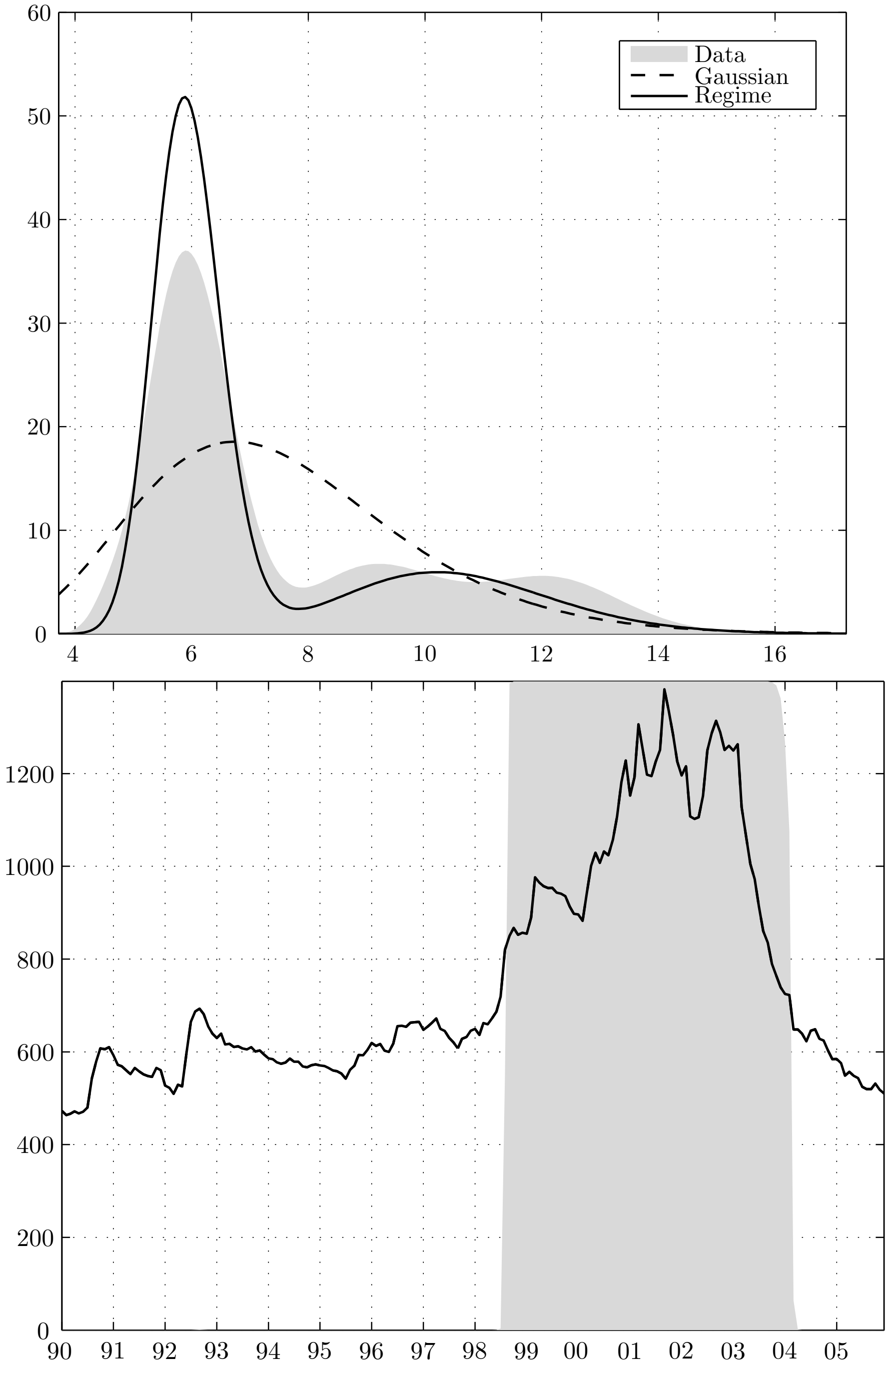

Figure 3: Estimation of the parameters for the regime switching loss distribution model, based on the aggregate ("Global") five year default probability. The details are as in figure 2.

Aggregate series	
Five year Global index	GBL5
One year Global index	GLB1

Industry-specific series	
Consumer Discretionary	COD
Consumer Staples	COS
Energy	ENE
Health care	HEA
Industrial	IND
Information Technology	INF
Materials	MAT
Telecommunications	TEL
Utilities	UTI

Table 2: The aggregate and industry-specific industries used in this study.

Hamilton's filter to carry out the maximum likelihood estimation. An important byproduct of this approach is the filtered probability of actually being in each regime at time t (given the information up to time t), denoted with

$$\boldsymbol{\xi}_{t|t} = \mathrm{P}\{s_t = j|\mathscr{F}_t\}$$

Essentially we make the assumption that the global portfolio of corporate bonds is diversified enough across industries to render industry specific effects negligible. Then, for elements of the global portfolio distance to default will take a single-factor form, namely

$$Z = \sqrt{\alpha_G(s_t)}X_G + \sqrt{1-\alpha_G(s_t)}\epsilon$$

with default happening if $Z \leqslant C(s_t)$. Therefore, conditional on the regime, the loss distribution will satisfy

$$\mathrm{P}\left\{L \leqslant \ell|s_t\right\} = \Phi\left(\frac{\Phi^{-1}(\ell)\cdot\sqrt{1-\alpha(s_t)} - C(s_t)}{\sqrt{\alpha(s_t)}}\right)$$

Differentiating yields the probability density of the credit loss, which is the fundamental block of the maximum likelihood estimation. In particular

$$f(\ell|s_t) = \frac{\sqrt{1-\alpha(s_t)}}{\sqrt{\alpha(s_t)}\varphi(\Phi^{-1}(\ell))}\varphi\left(\frac{\Phi^{-1}(\ell)\cdot\sqrt{1-\alpha(s_t)} - C(s_t)}{\sqrt{\alpha(s_t)}}\right)$$

where φ denotes the probability density of the standard normal distribution.

The regime switching filter and likelihood is based on a simple Bayesian update of $\boldsymbol{\xi}_{t|t}$, based on the lagged filtered probabilities $\boldsymbol{\xi}_{t-1|t-1}$, the observed loss ℓ_t and the transition probability matrix $\boldsymbol{P}$. A detailed overview can be found in Hamilton (1994), but we give here a brief overview of the algorithm:

Gaussian specification

	GLB5	GLB1	COD	COS	ENE	HEA	IND	INF	MAT	TEL	UTI
α(pc)	2.50	3.04	2.08	1.48	1.32	2.66	2.21	7.57	4.81	12.12	0.69
	(0.25)	(0.30)	(0.21)	(0.15)	(0.13)	(0.27)	(0.22)	(0.72)	(0.47)	(1.09)	(0.07)
C	-1.43	-2.34	-2.37	-2.60	-2.64	-2.15	-2.47	-2.31	-2.42	-1.97	-2.67
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.01)
log-lik	461.8	781.1	827.1	966.4	995.1	710.8	870.3	703.7	786.3	531.9	1068.3

Regime switching specification

	GLB5	GLB1	COD	COS	ENE	HEA	IND	INF	MAT	TEL	UTI
α_1(pc)	0.76	0.96	0.42	0.16	0.17	0.00	0.38	1.58	0.00	7.73	0.00
	(0.16)	(0.19)	(0.33)	(0.24)	(0.21)		(0.34)	(0.45)		(1.43)	
C_1	-1.25	-2.14	-2.23	-2.48	-2.53	-2.05	-2.30	-2.03	-2.22	-1.64	-2.57
	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.07)	(0.02)	(0.02)	(0.26)	(0.04)	(0.01)
α_2(pc)	0.13	0.33	0.26	0.12	0.00	2.34	0.00	0.40	1.58	3.33	0.00
	(0.03)	(0.06)	(0.09)	(0.07)		(0.35)		(0.13)	(1.23)	(0.49)	
C_2	-1.56	-2.48	-2.46	-2.67	-2.71	-2.22	-2.58	-2.58	-2.60	-2.26	-2.72
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.06)	(0.02)	(0.01)
p_{11}(pc)	98.0	98.1									
	(1.6)	(1.6)									
p_{22}(pc)	99.3	99.3									
	(0.6)	(0.7)									
log-lik	639.9	931.5	902.0	1031.9	1074.1	770.9	1015.9	875.3	904.2	621.4	1133.3

Table 3: Maximum likelihood estimation for the Gaussian-mix and the regime switching specifications. **Gaussian-mix:** Correlation and default thresholds are reported. **Regime switching:** We report the probabilities of remaining in the current regime. Industry-specific estimates are conditional on the global regime structure. Also, excess correlation values are estimated. For example, the correlation across Industrials in the first regime will be $0.96\% + 0.38\% = 1.34\%$. Standard errors in parentheses.

- Market participants observe loss realizations ℓ_t over time, and therefore the information $\mathscr{F}_t$ is the one generated by these realizations
$$\mathscr{F}_t = \{\ell_t, \ell_{t-1}, \ldots, \ell_1\}$$
For this reason we can write $\mathrm{P}\{s_t = j|\mathscr{F}_t\} = \mathrm{P}\{s_t = j|\ell_t = \ell, \mathscr{F}_{t-1}\}$.

- We invoke Bayes' formula, and change the argument of the probability
$$\mathrm{P}\{s_t = j|\ell_t = \ell, \mathscr{F}_{t-1}\} = \frac{\mathrm{P}\{\ell_t = \ell|s_t = j, \mathscr{F}_{t-1}\}\mathrm{P}\{s_t = j|\mathscr{F}_{t-1}\}}{\mathrm{P}\{\ell_t = \ell|\mathscr{F}_{t-1}\}}$$

- The denominator of the above expression is actually the conditional likelihood, which we see to sum over all times and maximize. We can condition on the current regime, $s_t = k$, and write
$$\mathrm{P}\{s_t = j|\ell_t = \ell, \mathscr{F}_{t-1}\} = \frac{\mathrm{P}\{\ell_t = \ell|s_t = j, \mathscr{F}_{t-1}\}\mathrm{P}\{s_t = j|\mathscr{F}_{t-1}\}}{\sum_k \mathrm{P}\{\ell_t = \ell|s_t = k, \mathscr{F}_{t-1}\}\mathrm{P}\{s_t = k|\mathscr{F}_{t-1}\}}$$

- The quantities $\mathrm{P}\{\ell_t|s_t, \mathscr{F}_{t-1}\}$ are the conditional densities $f(\ell_t|s_t)$ which are known in closed form. The quantities $\mathrm{P}\{s_t|\mathscr{F}_{t-1}\}$ are the state forecasts given the information at time $t-1$, which we denote $\boldsymbol{\xi}_{t|t-1}$. Therefore they will be given as the elements of the product $\boldsymbol{\xi}_{t|t-1} = \boldsymbol{P} \cdot \boldsymbol{\xi}_{t-1|t-1}$.

Given a set of parameter values we can loop through the steps above to compute the filtered probabilities $\boldsymbol{\xi}_{t|t}$. During each time step we store the likelihood value
$$L(\ell_t|\mathscr{F}_{t-1}; \boldsymbol{\vartheta}) = \sum_k f(\ell_t|s_t)\mathrm{P}\{s_t = k|\mathscr{F}_{t-1}\}$$
Our objective is to find the parameter values $\boldsymbol{\vartheta} = \{\alpha_1, \alpha_2, C_1, C_2, p_{11}, p_{22}\}$ that maximize the sample log-likelihood
$$L(\boldsymbol{\ell}; \boldsymbol{\vartheta}) = \sum_t \log L(\ell_t|\mathscr{F}_{t-1}; \boldsymbol{\vartheta})$$
In order to start the recursion one needs to input the initial regime probabilities $\boldsymbol{\xi}_{0|0}$. Typically, one uses the unconditional distributions, which for the two-state Markov chain are given by $\mathrm{P}\{s_t = j\} = \frac{p_{jj}}{p_{11}+p_{22}}$.

The estimated density and the corresponding filtered probabilities are given in figure 2. One can readily observe that the probability of being in a "bad" period increases from virtually zero to one at around mid-1999, and then drop back to zero during 2003. The parameter estimates are given in table 3. We also calibrate the model to the five year default probabilities, and the resulting densities are given in figure 3. The parameter estimates between the two horizons are indeed consistent, with the implied Gaussian correlation for the five year maturity being lower.

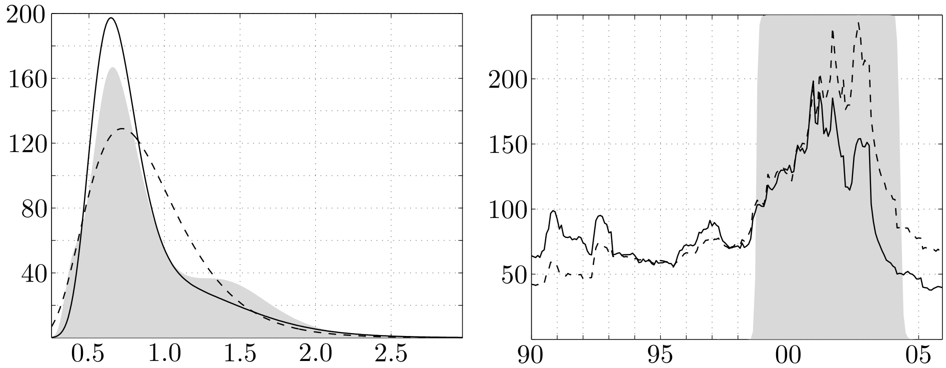

(a) Consumer Discretionary

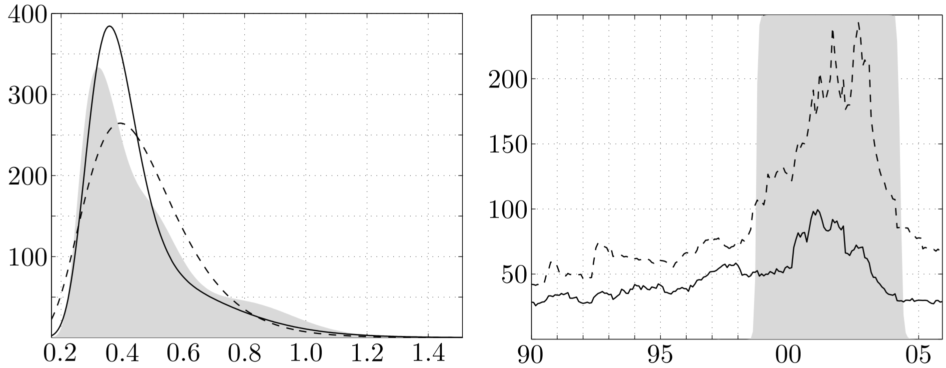

(b) Consumer Staples

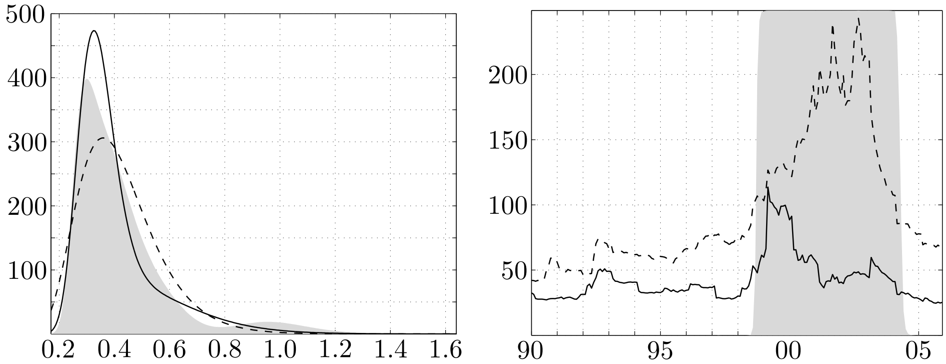

(c) Energy

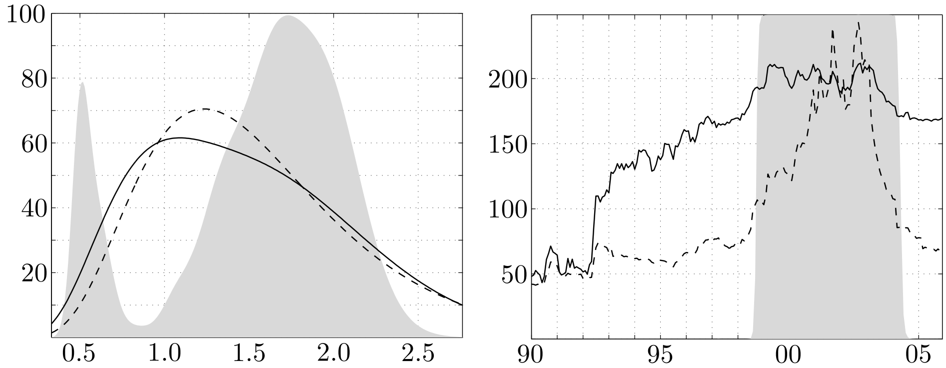

(d) Health Care

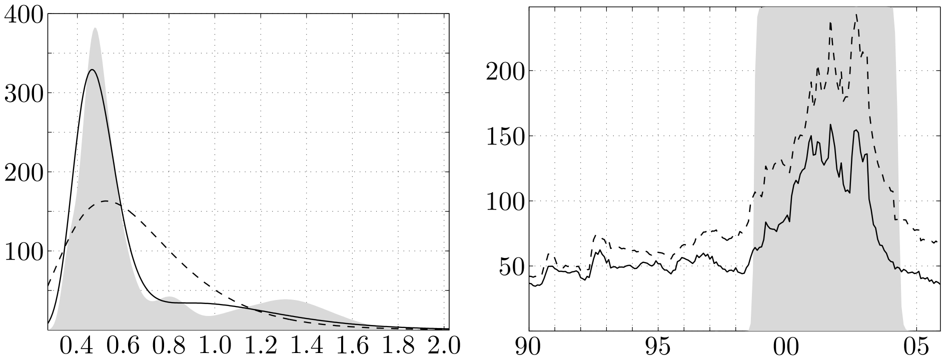

(e) Industrials

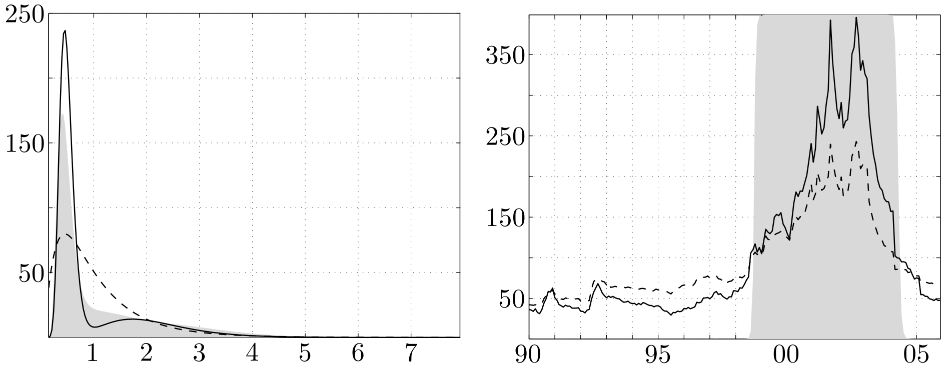

(f) Information Technology

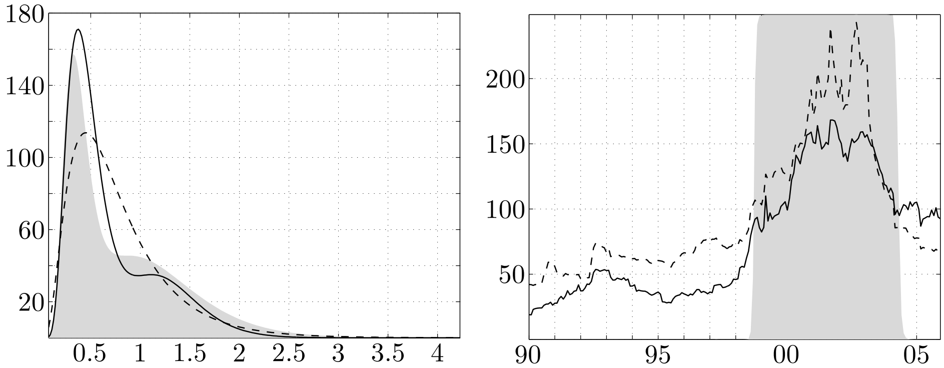

(g) Materials

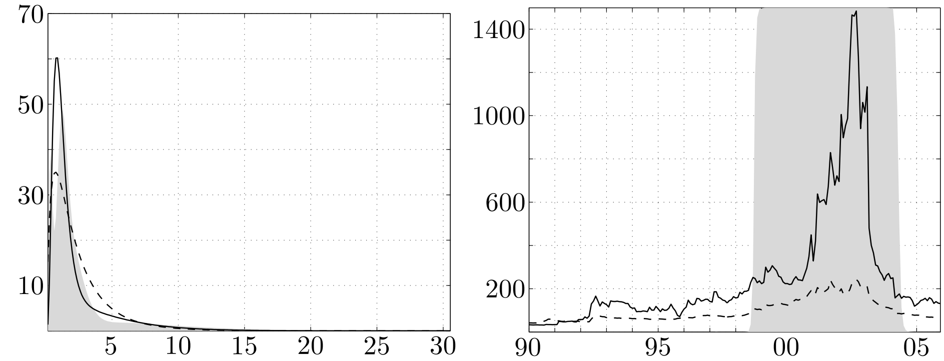

(h) Telecommunications

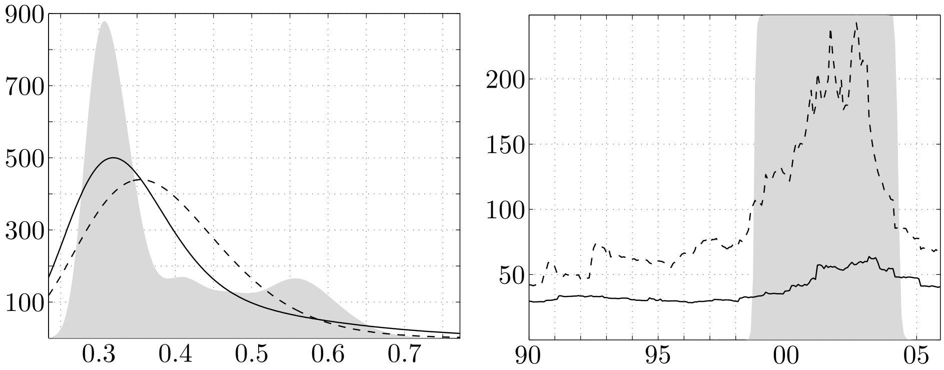

(i) Utilities

Figure 4: Estimation of the parameters for the regime switching loss distribution model across different sectors. See figure 2 for a description. The line in the bottom figure gives the default probabilities for the corresponding sector. The dashed line is the aggregate probability of default, as in figure 2.

3.4 Industry-specific parameters

Given the regime structure $\boldsymbol{\xi}_{t|t}$ and the global correlations α_G, we now use the industry-specific historical default rates to estimate the uplifts α_n.

Within each industry the global and industry-specific factor affects all securities, and the distance to default can be written in the form

$$Z_n = \sqrt{\alpha_n^\star(s_t)} X_n^\star + \sqrt{1 - \alpha_n^\star(s_t)} \epsilon$$

where $\alpha_n^\star(s_t) = \alpha_G(s_t) + \alpha_n(s_t)$, and $X_n^\star$ follows a standardized normal. This representation is valid, since the distribution of

$$\sqrt{\alpha_G(s_t)} X_G + \sqrt{\alpha_n(s_t)} X_n \sim \mathrm{N}(0, \alpha_G(s_t) + \alpha_n(s_t))$$

and, as we pointed out, within an industry there is no distinction between the two factors, as they affect all securities in the same fashion.

We make the assumption that the credit cycle is a global factor. For that reason we estimate the parameters conditional on the process $\mathrm{P}\{s_t|\mathscr{F}_{t-1}\}$, which has been filtered out based on the global default experience. This also means that we do not need to re-estimate the transition matrix, and the only free parameters that we do estimate for each industry are the uplifts and the default thresholds $\boldsymbol{\vartheta}_n = \{\alpha_{n,1}, \alpha_{n,2}, C_{n,1}, C_{n,2}\}$. The log-likelihood that we maximize is therefore

$$L(\boldsymbol{\ell}; \boldsymbol{\vartheta}_n) = \sum_t \log \left\{ \sum_k f(\ell_t | s_t; \boldsymbol{\vartheta}_n) \mathrm{P}\{s_t = k | \mathscr{F}_{t-1}\} \right\}$$

As correlation within an industry must be at least as high as the global, we restrict the uplifts to ne non-negative, $\alpha_n \geqslant 0$. The data and the calibrated densities are given in figures 4(a–i), while the parameter estimates are presented in table 3.

Overall, the fit is substantially improved when a regime switching structure is used, even though the timing of the switches is determined solely by the aggregate default frequencies. This is in line with the high correlation that default rates exhibit across different industries. The only exceptions is the Health Care sector, where the correction of 1993 appears to have a permanent effect that distorts the results.

4 Recovery rates across regimes

We now turn our attention to the analysis of recovery rates. Recent evidence suggests that the unconditional distribution is bimodal (see for example Schuermann, 2004). This distributional feature can be linked to a regime switching model, where the conditional recovery densities are unimodal but unconditionally they can produce a bimodal mixture. Various links of this state variable to the credit or business cycle have been offered in the literature, as in Frye (2000b), Hu and Perraudin (2002), and Altman et al. (2005), among others.

The recovery database spans seven years, from 2000 to the end of 2006. Defaults are given across industries, seniorities and time of default. Overall the database consists of 1382 defaulted high yield bonds, issued by 561 entities. The average time to default of the bonds is about four years and five months. For the purpose of this paper we pool together bonds of the same issuer, and compute a proxy recovery rate which is weighted by the bond notional. The large majority (94%) of issuers have issued bonds at the same seniority level. For the remaining 6% we also compute a "weighted seniority" value, which we then map to the closest seniority class. Thus, we end up with recovery rates and the corresponding seniorities of 561 issuers.

4.1 Estimation of regime-dependent recovery rates

The data set of recoveries is substantially shorter than the one used for the default process, but fortunately is spans both regimes. In particular, the first four years are identified by the regime switching filter as "bad times", with the switch to "good times" around the end of 2003. Figure 5 attempts to illustrate the variability of recovery rates across the two regimes, and the negative correlation that they exhibit with respect to the default frequencies. Low recoveries are clearly coupled with the high default rates of the first three years up to 2003. As the market moves towards the low default rate regime recovery rates increase, and during the "good" times of the period after 2004 recoveries are overall higher.

Motivated by this correspondence of the credit cycle with historical recovery rates, we estimate the recovery rate distribution that is conditional on the filtered regime probabilities over the available period, $\boldsymbol{R} = \{R_1, \ldots, R_T\}$. In

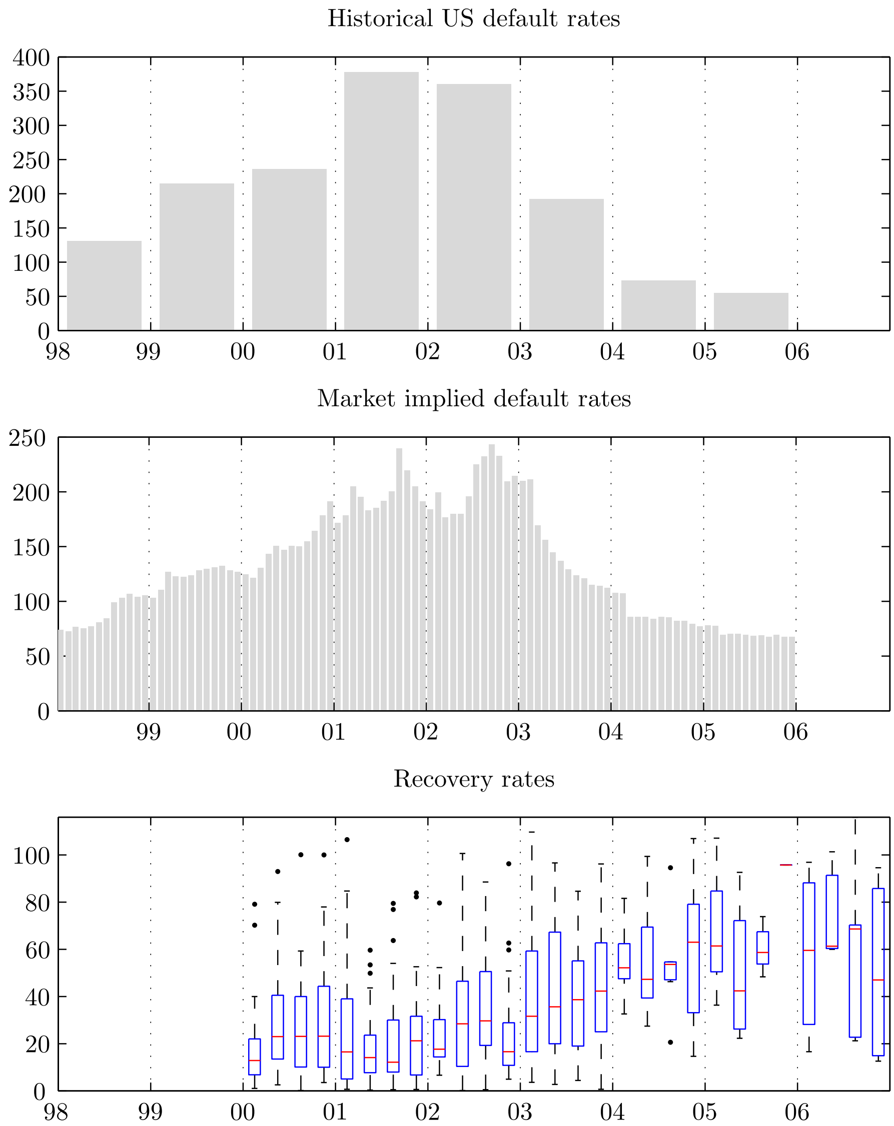

Figure 5: Recovery rates and historical default experience. The top figure gives the historical default rates of high yield securities, collected by Standard and Poor's. The middle graph gives the market-impled probabilities of default that are used in this study. In the bottom graph we present the distribution of recovery rates, pooled quarterly. There is an inverse relationship between default and recovery rates over this period of time.

	before 2004		after 2004	
	double-bounded	Beta	double-bounded	Beta
ρ_1	0.873	0.864	1.767	1.887
ρ_2	2.155	2.206	1.500	1.509

Table 4: Estimated recovery rate parameters for the double-bounded and the Beta distribution. See figure 6 for plots of the corresponding probability densities.

particular, we maximize the likelihood function

$$L(\boldsymbol{R}; \boldsymbol{\rho}) = \sum_t \log \left\{ \sum_k g(R|s_t; \boldsymbol{\rho}) \mathrm{P}\{s_t = k|\mathscr{F}_{t-1}\} \right\}$$

In the above expression $g(R_t|S_t; \boldsymbol{\rho})$ is the conditional density of recovery rates, and $\boldsymbol{\rho}$ is the corresponding parameter vector.

The Beta distribution has been the natural choice to model the magnitude of recoveries, as its support is naturally bounded over the unit interval. Here we also implement a very simple alternative, the double-bounded distribution. This is also bounded and resembles the Beta, but is a lot simpler to simulate in practical applications as it does not need to numerically invert the incomplete Beta function. The probability and cumulative densities for both distributions are given by

Beta distribution:

$$g(R; \boldsymbol{\rho}) = \frac{1}{B(\rho_1, \rho_2)} R^{\rho_1 - 1}(1 - R)^{\rho_2 - 1}$$
$$G(R; \boldsymbol{\rho}) = \frac{1}{B(\rho_1, \rho_2)} \int_0^R r^{\rho_1 - 1}(1 - r)^{\rho_2 - 1} \mathrm{d}r$$

Double-bounded distribution:

$$g(R; \boldsymbol{\rho}) = \rho_1 \rho_2 R^{\rho_1 - 1} \left(1 - R^{\rho_1}\right)^{\rho_2 - 1}$$
$$G(R; \boldsymbol{\rho}) = 1 - \left(1 - R^{\rho_1}\right)^{\rho_2}$$

The maximum likelihood estimates confirm that the credit cycle switch also has a negative impact on the recovery rates. Table 4 gives the estimated parameter values, and the probability densities are presented in figure 6. The parameters and in particularly the fitted densities are indistinguishable, but we will use the double-bounded in our simulation experiments, as it is a lot easier to simulate from.

During "bad times" recovery rates tend to be very low, with over two thirds below the 50% mark. As default rates drop in "good times" recovery rates

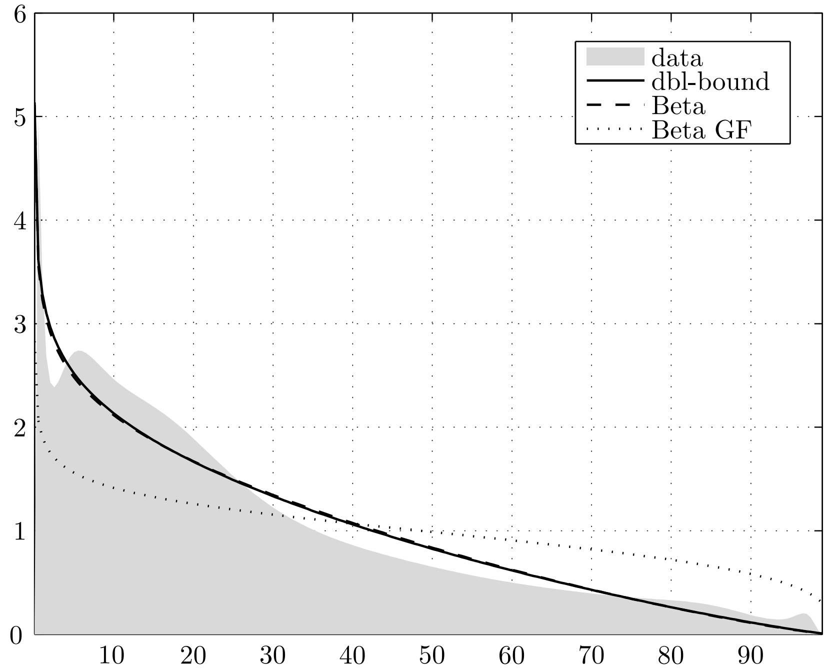

(a) “bad” times, before 2004

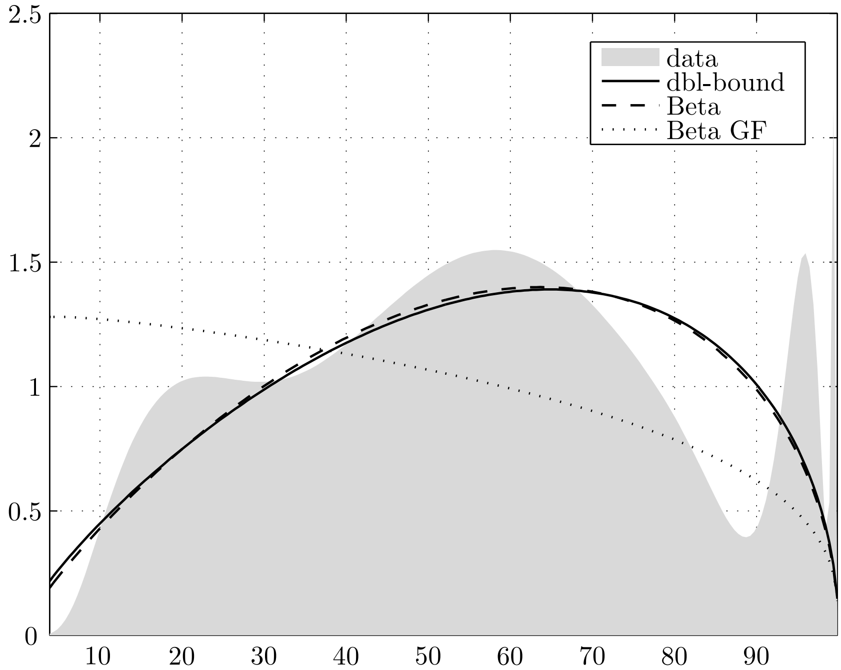

(b) “good” times, after 2004

Figure 6: Recovery rates in “bad” and “good times”. The empirical distribution is presented, along with the fitted Beta and double-bounded probability density functions. The estimated parameters used are given in table 4. The Beta distribution estimated in Gifford Fong Associates (2004) is also presented.

become distributed in a more uniform manner. Note, that only 72 out of the 561 defaults (13%) happened during "good times". A two-sample Kolmogorov-Smirnov (KS) test to assess if the differences of the two distributions are statistically insignificant returns a p-value of zero, suggesting that the behavior of recovery rates is markedly different across credit regimes. Gifford Fong Associates (2004) arrive to a similar conclusion using a more extensive data set, but conditioning on the official business cycles, rather than endogenously determining the state. The Beta densities of Gifford Fong Associates (2004) are also plotted in figure 6.

4.2 Seniority and industry dependence

Data are available for the three most senior classes, namely senior subordinated, senior unsecured and senior secured debt. The nonparametric density reveals small differences between senior subordinated and senior unsecured bonds. We test the hypothesis that the two samples come from the same distribution using the two-sample KS test, which returns a p-value of 3.5%. On the other hand, senior secured instruments appear to be distributed differently from the other two, and the KS test confirms that at all confidence levels, returning a p-value of zero. The empirical distributions of the recovery rates across different seniority classes are given in figure 7.

We investigate further by looking at recovery rates across seniorities, conditional on the two regimes. Perhaps surprisingly, there is a monotonic relationship between default rates and seniority. 9% (14 out of 153) senior subordinated bonds defaulted during the "good times", as opposed to 13% (46/343) of senior unsecured and 19% (12/58) of senior secured bonds. The recovery rate distribution is consistent across all seniority classes during "bad times", but the small sample does not allow for a conclusive analysis of the period after 2004.[7]

The recovery data are also split into 27 industries (identified by Fitch), and we now attempt to investigate any such dependence of recovery rates. Due to the small sample sizes we pool the recovery rates across industries to form nine sectors, which correspond roughly to the sectors used in the analysis of default rates.[8] The distributions of recovery rates are presented in figure 8.

We also perform all pairwise KS tests to assess if the recovery rate distributions are significantly different. In total, 45 pairs are examined. We find that only the pairs that include the Utilities sector are significantly different from all other pairs, at the standard 5% confidence level. Also significant are three pairs that include the Technology sector (with Industrials, Materials and Services), and the Services-Consumer pair. The results confirm the findings of Altman and Kishore (1996) and the more recent study of Acharya, Bharath, and Srinivasan (2006), who find strong evidence that industries with tangible assets (such as utilities) exhibit higher recovery rates.

[7] More detailed results are available upon request.
[8] With the exception of Financials which is not included in the analysis of default rates.

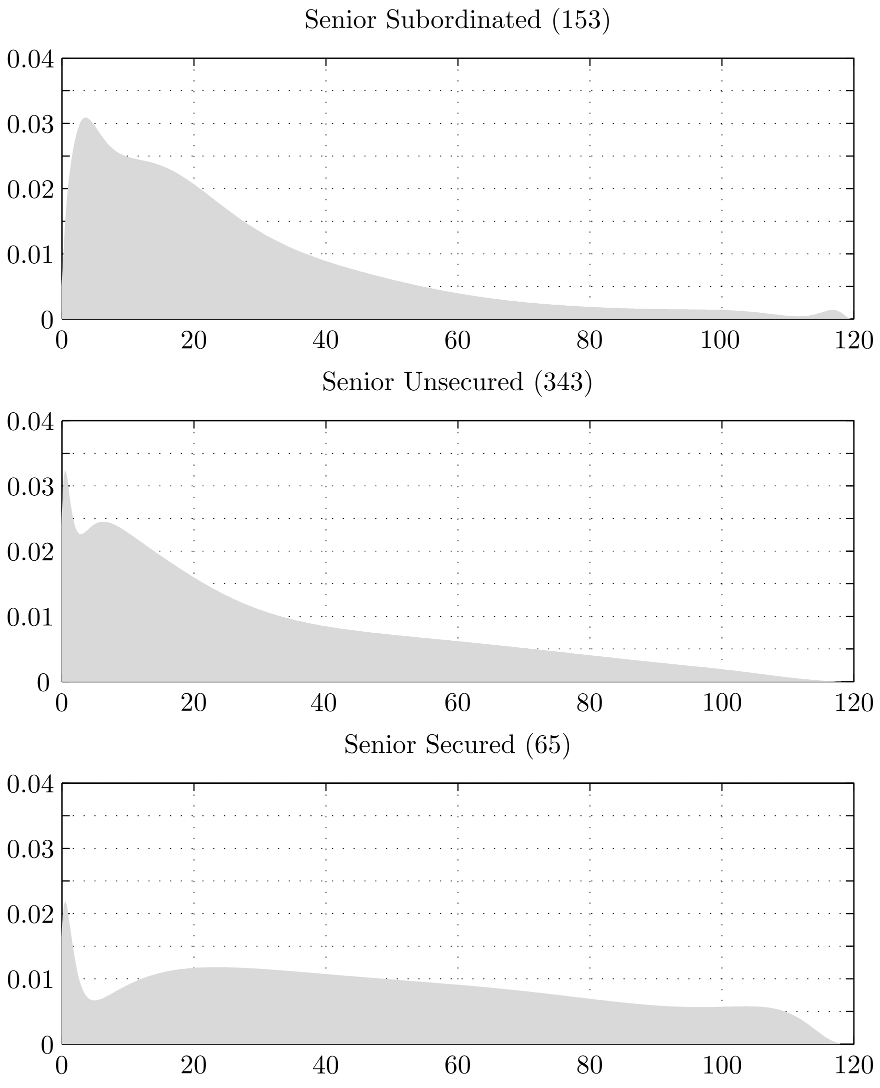

Figure 7: Empirical distributions of recovery rates across seniorities. The number of available observations within each class is given in parentheses.

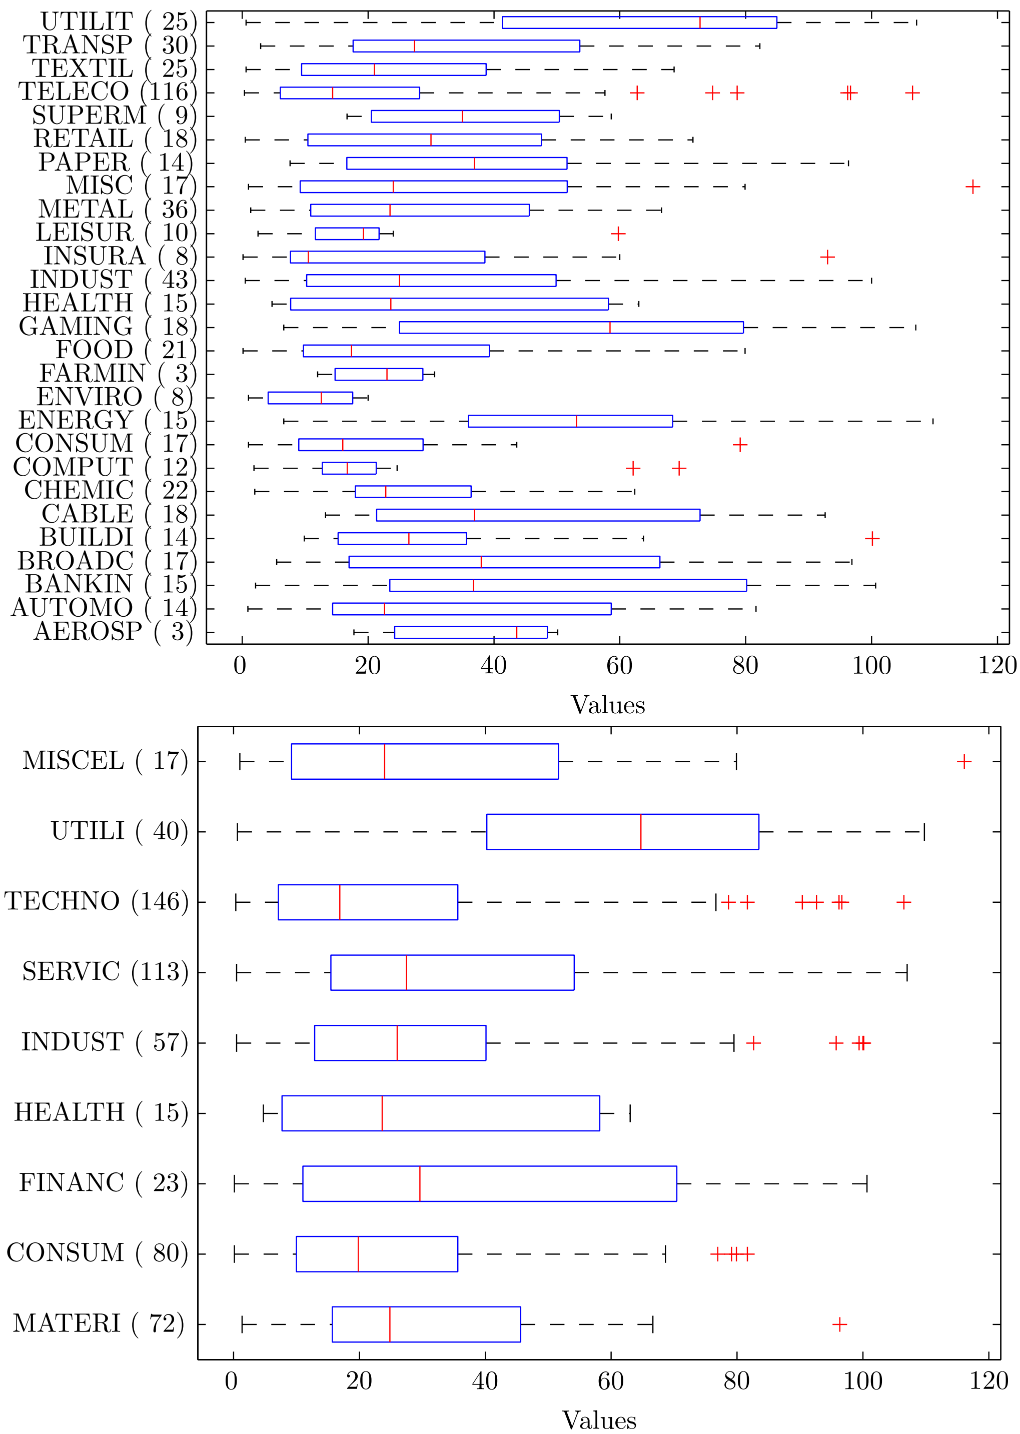

Figure 8: Empirical distributions of recovery rates across industries and sectors. The number of available observations within each class is given in parentheses.

5 A simulation exercise

In this section we use the results of our analysis to investigate the loss distribution across different regime probabilities. As pricing is carried out under an equivalent probability measure, we also stress the various parameters to investigate the effect of this measure change. We focus on determining the attachment point of different senior tranches, which is in principle the same as solving for the Value-at-risk point for different confidence levels.

To construct the loss distributions we simulate 100,000 scenarios. We consider 8 industries and select 20 names from each industry, constructing a portfolio of 160 names. This resembles a typical CDO structure. The global correlations across the two regimes are set close to their maximum likelihood values, namely $\alpha_{G,1} = 1.00\%$ and $\alpha_{G,2} = 0.35\%$. We set an industry uplift of $\alpha_{n,1} = \alpha_{n,2} = 0.30\%$, which is constant across all industries and across the two regimes. This is a choice that resembles the average estimated uplifts. The default probabilities vary and are given in the following table

n	$C_{n,1}$		$C_{n,2}$	
1	-2.00	(230bps)	-2.30	(110bps)
2	-2.00	(230bps)	-2.30	(110bps)
3	-2.00	(230bps)	-2.30	(110bps)
4	-2.20	(140bps)	-2.60	(45bps)
5	-2.20	(140bps)	-2.60	(45bps)
6	-2.20	(140bps)	-2.60	(45bps)
7	-2.50	(60bps)	-2.70	(35bps)
8	-2.50	(60bps)	-2.70	(35bps)

These values capture a wide range of default probabilities across both regimes, broadly capturing the market implied default rates used in the estimation part of this paper. The recovery rates across the two regimes follow the double-bounded distribution with parameters $\boldsymbol{\rho}_1 = (0.90, 1.80)$ and $\boldsymbol{\rho}_2 = (2.20, 1.50)$. We do not differentiate across industries or seniorities in this exercise.

We also stress the correlation parameters, as we are interested in a more conservative setting that would be plausible under the risk-adjusted equivalent probability measure. In distributional terms these adjustments will produce a loss distribution that exhibits a fatter right tail. Stressing the default thresholds would pretty much shift the loss probability to the right, but would not substantially alter its shape. For that reason we focus on correlation stresses.

The impact of different stresses will vary with respect to the initial regime. In the regime switching framework the current market regime is considered unobserved, and investors only take a probabilistic view on the position over the cycle. When it comes to the loss distribution, it will be the outcome of a mixture of the two conditional loss distributions, weighted with the respective probability values. In that respect it makes sense to talk about a market which is 75% in recession, meaning an appropriate weighting of the loss distribution. Note that the unconditional probabilities would just reflect a non-informative view of the current market conditions.

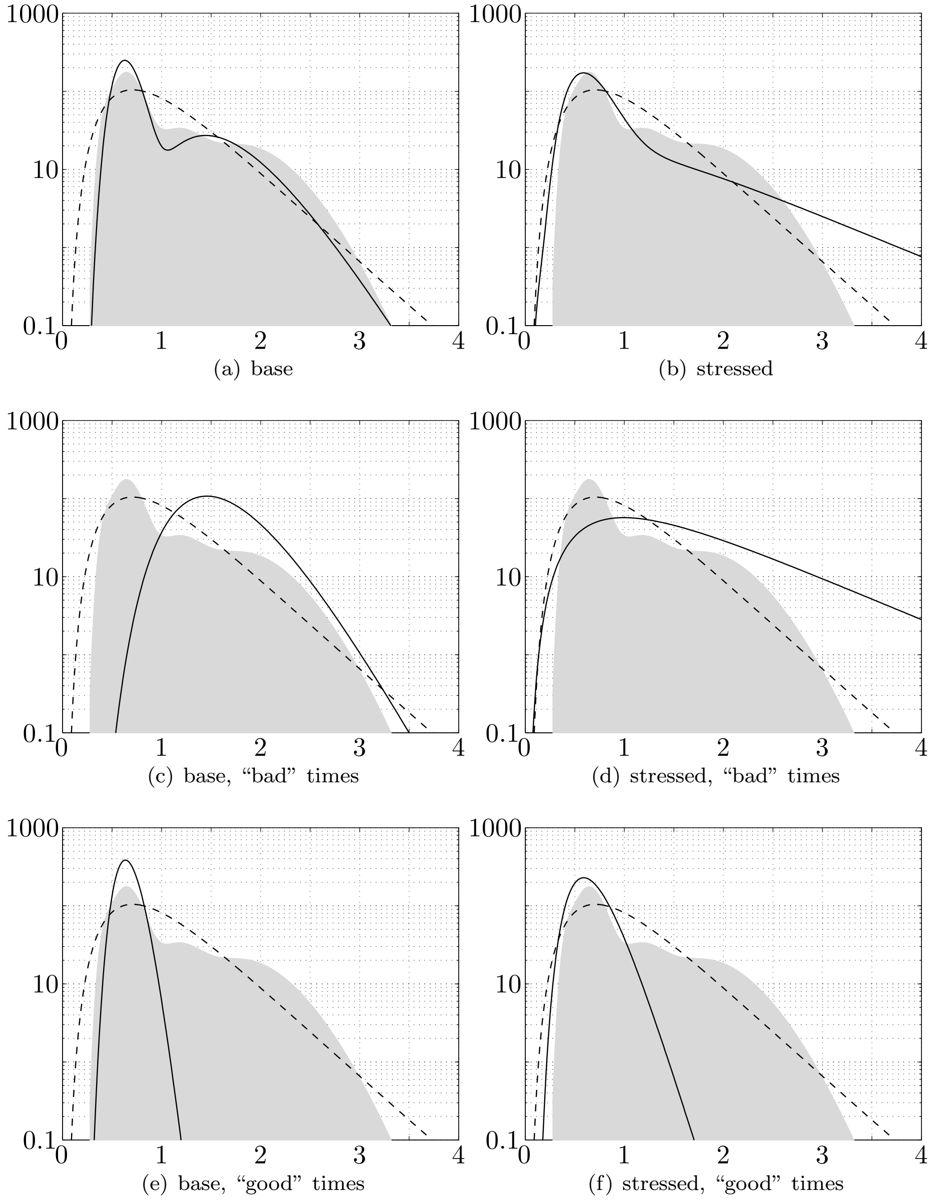

Figure 9: Conditional and stressed probability log-densities. The “base” cases correspond to the maximum likelihood stimates $\alpha_1 = 0.35\%$ and $\alpha_2 = 1\%$. The “stressed” parameters are $\alpha_1 = 1\%$ and $\alpha_2 = 5\%$. The dashed line corresponds to the unstressed Gaussian loss distribution that was fitted with maximum likelihood.

In this exercise we increase the correlation in the "good" regime from 0.35% to 1%, and the correlation of the "bad" regime from 1% to 5%. In 9 we present logarithmic plots of the conditional densities, together with the unconditional density. We can see that the impact on the right tail of the distribution is substantial. Also, we can verify that the Gaussian copula produces an unconditional loss distribution that attaches spuriously higher than needed correlation, as it attempts to produce the humped shape of the loss distribution.

In figure 10 we give the loss distributions that are conditional on the probability of currently being in the "bad" regime, for both base and stressed parameter configuration. We also illustrate the Value-at-Risk points, where a senior tranche would be attached to in order to sustain losses with a predefined probability. These probabilities range from 0.50% to 20%.

We start with investigating a worsening credit cycle, moving down from figure 10(a) to 10(g). In this case the probability of the current regime being the "bad" one increases towards one, and as expected all attachment points shift to the right. There are two observations that one can make with the respect to the nature of these shifts: Firstly, even a relatively small increase in the probability of a recessions results in a substantial shift of the attachment points. This is illustrated by the way the tightly packed attachment points at $p = 0$ spread when the probability rises to $p = 25\%$. Subsequent increases to 75% or 100% do not have such an impact. Secondly, these shifts are not uniform. The range between the highes and lowest attachment points rises initially from 1.4% to 3.6%, but then drops down to 3.1%. Higher quality senior tranches are the first to react to the worsening cycle, as they are most responsive to the fattening of the extreme tail. Further deterioration does not have such an impact. On the other hand, the attachment points of lower quality tranches do not shift that erratically, and increase in a steadier fashion.

If we turn to the impact of correlation stresses, we observe that attachment points are more responsive in the downside of the credit cycle. Also, it is easy to verify that higher quality assets are more sensitive to such stresses than lower quality ones. The attachment points of class F assets hardly move as a result of the stresses, across all market conditions. On the other hand, the attachment points of class A assets increase by as much as 2%.

6 Concluding remarks

In this paper we have attempted to decompose the loss distribution in two dimensions: On a global scale, we model the impact of the credit cycle, using a regime switching specification. This is responsible for the multi-modality that is routinely observed in historical default rates. On a more granular scale, given the position of the market in the cycle, we use a Gaussian copula to capture the residual correlation that is not captured by common default rate moves. We implement a regime switching model to estimate the parameters. This has the advantage of endogenously determining the peaks and troughs of the cycle, giving a probabilistic weighting at each point in the sample. We find indeed

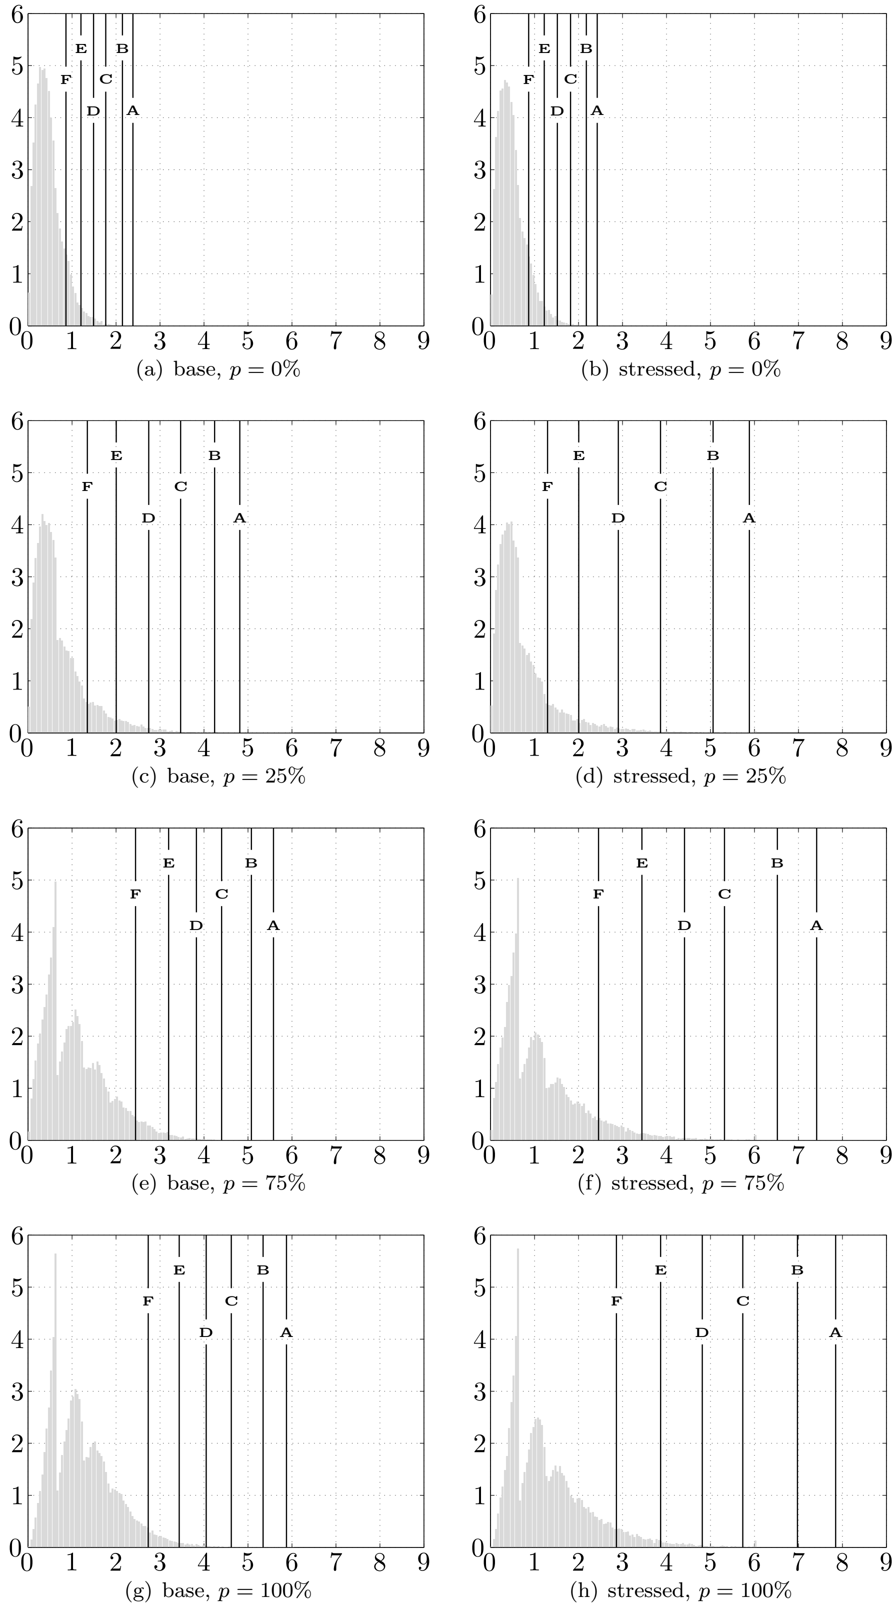

(a) base, $p = 0\%$

(b) stressed, $p = 0\%$

(c) base, $p = 25\%$

(d) stressed, $p = 25\%$

(e) base, $p = 75\%$

(f) stressed, $p = 75\%$

(g) base, $p = 100\%$

(h) stressed, $p = 100\%$

Figure 10: Loss distributions for different regime weights. Also given are attachment points for senior trances that can sustain losses with probabilities $(A, B, C, D, E, F) \mapsto (0.5, 1.0, 2.5, 5.0, 10.0, 20.0)$ percentage points.

that a relatively low correlation is needed to fit the historical experience. The high correlation values that are required in single state Gaussian copula models are more likely to be due to model mis-specification.

In the next step we assess the impact of the global credit cycle on individual industries. We find that using a common factor does a very good job in explaining the variabilities across most industries. An investigation of recovery rates indicates that they too depend heavily on the underlying regime.

A small simulation experiment is conducted to investigate the impact of the credit cycle and parameter stresses on loss distributions of hypothetical portfolios.

Overall, we conclude that

- The credit market exhibits two distinct regimes that we coin "good times" and "bad times".
- "Good times" have lower default probabilities and the correlation across assets is also low. Defaults are rare and when they happen they are due to idiosyncratic factors, rather than industry specific. Recoveries in such epochs are distributed uniformly.
- In "bad times" default probabilities increase substantially, and so do the correlations. Now defaults more dependent on the systematic and the industry factors. During "bad times" recoveries are substantially lower, and also markedly skewed.
- Different asset classes are more responsive to changes in the credit environment than others. A small probability of a downturn is sufficient for a large impact on the behavior of high quality assets. In contrast, the effect on assets with lower quality is more gradual.

References

Acharya, V. V., S. T. Bharath, and A. Srinivasan (2006). Does industry-wide distress assect defaulted firms? Evidence from creditor recoveries. Technical report, Unpublished Manuscript.

Altman, E. I., B. Brady, A. Resti, and A. Sironi (2005). The link between default and recovery rates: Theory, empirical evidence and implications. *Journal of Business 78*(6), 2203–2227.

Altman, E. I. and V. M. Kishore (1996). Almost everything you wanted to know about recoveries on defaulted bonds. *Financial Analysts Journal*, 57–64.

Altman, E. I., A. Resti, and A. Sironi (2004). Default recovery rates in credit risk modeling: A review of the literature and empirical evidence. *Economic Notes 33*(2), 183–208.

Black, F. and J. C. Cox (1976). Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance 31*(2), 351–367.

Bruche, M. and C. Conzález-Aguado (2007). Recovery rates, default probabilities and the credit cycle. Unpublished Manuscript, CEMFI.

Das, S. R., D. Duffie, N. Kapadia, and L. Saita (2007). Common failings: How corporate defaults are correlated. *Journal of Finance 62*.

Das, S. R. and G. Geng (2004). Correlated default processes: A criterion-based copula approach. *Journal of Investment Management 2*.

Duffie, D. (2004). Irresistable reasons for better models of credit risk (Comment). Financial Times, April 16, 2004.

Frye, J. (2000a, April). Collateral damage. *Risk Magazine*, 91–94.

Frye, J. (2000b, November). Depressing recoveries. *Risk Magazine*, 108–111.

Gifford Fong Associates (2004). Recovery rate modeling. Unpublished Manuscript.

Gupton, G. M. and R. M. Stein (2005). LossCalc V2: Dynamic prediction of LGD. Technical report, Technical Report, Moody's-KMV.

Hamilton, J. D. (1994). *Time Series Analysis.* Princeton, NJ: Princeton University Press.

Hu, Y.-T. and W. Perraudin (2002). The dependence of recovery rates and defaults. Working paper, Birkbeck College, University of London.

Hull, J. C. and A. White (2004). Valuation of a CDO and an n^{th} to default CDS without Monte Carlo simulation. *Journal of Derivatives 12*(2), 8–23.

Jarrow, R. A. and S. M. Turnbull (1995). Pricing options on financial securities subject to default risk. *Journal of Finance 50*, 53–86.

Kalemanova, A., B. Schmid, and R. Werner (2005). The normal inverse Gaussian distribution for synthetic CDO pricing. Unpublished manuscript.

Lando, D. (1997). Modeling bonds and derivatives with default risk. In M. Dempster and S. Pliska (Eds.), *Mathematics of Derivatives Securities*, pp. 369–393. Cambridge University Press.

Lando, D. (1998). On Cox processes and credit-risky securities. *Review of Derivatives Research 2*, 99–120.

Li, D. X. (2000). On default correlation: A copula function approach. *Journal of Fixed Income*, 43–54.

Liu, B., A. E. Kocagil, and G. M. Gupton (2007). Fitch equity implied rating and probability of default model. Technical paper, QFR, Fitch Ratings.

Madan, D. and H. Unal (1998). Pricing the risks of default. *Review of Derivatives Research 2*(2-3), 121–160.

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance 2*, 449–471.

Moosbrucker, T. (2005). Pricing CDOs with correlated Variance Gamma distributions. unpublished manuscript.

Schuermann, T. (2004). What do we know about loss-given-default? In D. Shimko (Ed.), *Credit Risk Models and Management* (2nd ed.). London, UK: Risk Books.

Wang, D., S. T. Rachev, and F. J. Fabozzi (2006). Pricing tranches of a CDO and a CDS index: Recent advances and future research. Working paper.