

# Recovery Rates, Default Probabilities and the Credit Cycle\*

Max Bruche and Carlos González-Aguado

CEMFI<sup>†</sup>

First version: November 28, 2006

This version: January 23, 2008

## Abstract

In recessions, the number of defaulting firms rises. On top of this, the average amount recovered on the bonds of defaulting firms tends to decrease. This paper proposes an econometric model in which this joint time-variation in default rates and recovery rate distributions is driven by an unobserved Markov chain, which we interpret as the “credit cycle”. This model is shown to fit better than models in which this joint time-variation is driven by observed macroeconomic variables. We use the model to quantitatively assess the importance of allowing for systematic time-variation in recovery rates, which is often ignored in risk management and pricing models.

**JEL classifications:** G21, G28, G33

**Keywords:** credit, recovery rate, default probability, business cycle, capital requirements, Markov chain


---

\*We would like to thank Ed Altman for giving us access to the Altman-NYU Salomon Center Default Database. We thank Manuel Arellano, Stéphane Bonhomme, Abel Elizalde, René Garcia, Juri Marcucci, Javier Mencía, Simon Potter, Rafael Repullo, Olivier Scaillet, Enrique Sentana and Javier Suárez, as well as seminar participants at the New York Fed, Brandeis, the Econometric Society European Meeting 2007, and the CEMAF/ ISCTE - Nova Credit Risk Conference for helpful comments.

<sup>†</sup>Centro de Estudios Monetarios y Financieros, C/ Casado del Alisal 5, 28014 Madrid.  
Email: bruche@cemfi.es and cgonzalez@cemfi.es.

# 1 Introduction

It has been noted that default probabilities and default rates (i.e. the fraction of defaulting firms in the economy) and average recovery rates are negatively correlated (see e.g. Altman et al. (2005); Acharya et al. (2007)). Both variables also seem to be driven by the same common factor that is persistent over time and clearly related to the business cycle: in recessions or industry downturns, default rates are high and recovery rates are low (Figure 1).



**Figure 1. Historical default rates and average recovery rates**

Historical default rates (fraction of defaulting firms) against time (from Standard and Poor's) and average recovery rates from the Altman data (see Section 3).

This paper attempts to empirically characterize the time-series behaviour of default probabilities and recovery rates distributions using an econometric model in which both depend on an unobserved two-state Markov chain.<sup>1</sup> The estimated states correspond to bad times, in which default probabilities are high and recovery rates are low, and good times, in which the reverse is true, so that we can interpret this Markov chain as

---

<sup>1</sup>Since the Markov chain is unobserved it creates (unconditional) dependence between defaults across firms, and between defaults and recovery rates, and therefore also plays a similar role to a dynamic frailty variable. Dynamic frailty variables have been used in the context of defaults e.g. by Duffie et al. (2006).

a “credit cycle”.

We can use the estimated model to characterize systematic credit risk. For example, we can explore the effect of allowing for time-variation in recovery rate distributions on estimates of credit risk. Whereas time-variation in default probabilities is almost always taken into account when calculating loss distributions or pricing credit-risk sensitive instruments, it is often assumed that recovery rates are either constant, or that recovery rates are independent of default probabilities. Given the negative relationship between default probabilities and recovery rates, this seems likely to be a bad idea. From the point of view of a holder of a diversified portfolio of corporate bonds the fact that recovery rates are low precisely in situations in which many companies default is important because the negative relationship between recoveries and default probabilities amplifies the systematic risk of the portfolio. The question is: By how much do you underestimate risk (overprice) if you ignore this negative relationship?

Other issues that can be addressed are in what way time-variations in default probabilities and recovery rate distributions relate to the business cycle, and whether recoveries on bonds of different seniorities vary in the same way over the cycle.

Related questions have been addressed in the literature. The relationship between default probabilities and the business cycle, for example, has already been documented by e.g. Nickell et al. (2000), who estimate different rating transition matrices for periods of high, medium and low GDP growth and find that default probabilities especially seem to be affected, or Bangia et al. (2002), who estimate separate rating transition matrices for NBER recessions and expansions, and find that economic capital of banks should be about 30% higher in recessions. They do not consider recovery rates, however.

Altman et al. (2005) regress average recovery rates on default rates and macroeconomic variables, and find that recovery rates and default rates are closely linked, and that macroeconomic variables become insignificant once default rates are included as explanatory variables. They suggest that this might be due to inelastic demand for defaulted securities; their hypothesis is that typical investors in defaulted securities (vulture funds) have limited capacity, and as a result, the price of defaulted securities falls by a large amount when many defaults occur. In related work Altman et al. (2001) also try to calculate by how much the 99% VaR increases when taking into account the negative relation between recovery rates and default rates. On the basis of a static model which is unrelated to their regressions, into which they substitute guessed

parameter values, they claim that the 99% VaR for a representative portfolio might increase from a percentage loss of about 3.8% to 4.9% when moving from a model in which recovery rates are constant to one in which there is negative dependence between recovery rates and default probabilities.

Acharya et al. (2007) take an argument by Shleifer and Vishny (1992) as a starting point: Suppose an industry is in distress and firms in this industry default. If the assets of the defaulting firms consist of industry-specific assets, the firms best able to put these assets to good use might also be experiencing problems, and hence might be unable to buy the assets. This could lead to low prices of these assets, and hence low recovery values. In a sense, this argument is similar to that proposed by Altman et al. (2005), except that here, the resale price of *real* assets is depressed whereas in the case of Altman et al. (2005), the resale price of financial assets (the defaulted bonds) is depressed. Acharya et al. (2007) find that especially in industries in which assets are very industry-specific, industry distress is related to low recoveries, providing evidence for the hypothesis of Shleifer and Vishny (1992). They argue that their results should have implications for pricing; which will of course only be the case if industry distress is associated with aggregate fluctuations such that industry credit risk has a systematic (and hence priced) component. Unfortunately, they do not investigate implications of their results for systematic risk.

Chava et al. (2006) propose a combined model of default probability and recovery rates to model expected losses. In their specification, default arrivals are governed by a continuous hazard that depends on observable firm-specific and economy-wide covariates and an unobserved industry-specific and constant frailty variable. Probit-transformed recovery rates are modelled as a linear function of possibly covariates and an error. In this approach, the dynamics of recovery rate distributions and their relation to default rates are driven by time-variation in observed covariates, whereas in our approach, this time-variation is driven by the unobserved Markov chain, i.e. we essentially exploit the persistence in the factor that seems to be driving both recovery rate distributions and default rates.

Using our approach, we find that a latent credit cycle describes default rates and recovery rates relatively well. For instance, on the basis of a model that only uses bond seniority and our latent cycle as explanatory variables, we calculate the out-of-sample rolling root mean square error (RMSE) for predicted recovery rates as Chava et al.

(2006) do. We obtain a value of 22.86%, which compares favourably with the lowest RMSE of 24.96% that they achieve. This seems to indicate that an approach that utilizes the persistence over time in aggregate recovery rates and default rates might have some advantages, although note that the datasets are not directly comparable.

We examine to which extent the latent “credit cycle” simply picks up macroeconomic fluctuations by estimating versions of the model in which default probabilities and recovery rate distributions also depend on macroeconomic and other economy-wide variables. We find that macroeconomic variables in general are significant determinants of default probabilities but not so for recovery rate distributions. However, the unobserved credit cycle variable is highly significant, whether or not macroeconomic variables are included, and the fit of the model seems to depend mainly on whether or not the unobserved credit cycle is included, and not on whether or not macroeconomic variables are included. This indicates that recovery rates and default rates are more tightly related to each other than to macroeconomic variables. Also, in our model the beginning of the state with high default probabilities and low recoveries that we interpret as a credit downturn precedes the start of a recession, and it can continue until after the end of a recession. The average duration of a credit downturn is in the range of 3.5 years. In contrast, a typical estimate of the average duration of a recession would be in the range of 4.1 to 4.7 *quarters* (Hamilton (1989)). It appears that the cyclicalities in credit variables is related to, but somewhat distinct from the macroeconomic cycle.

Given our estimated models, we can show that credit risk is much higher in a dynamic model in which both default probabilities and recovery rates are allowed to vary, than in a static model. For a well-diversified representative portfolio, the 99% VaR is a percentage loss of about 2.9% in the dynamic case, as opposed to 1.9% in the static case. This increase is due mainly to time-variation in default probabilities, however: In a model in which only default probabilities vary over time, the 99% VaR is already about 2.6% (in contrast, in a model in which only recovery rate distributions are allowed to vary over time, the 99% VaR is about 1.9%, i.e. equal to the VaR in the static case). This suggests that although variation in recovery rate distributions over time does have an impact on systematic risk, this impact is small relative to the importance of the time variation in default probabilities. Also, in relative terms, it is much smaller than the impact suggested by the calculations of Altman et al. (2001).

Using the calculated loss distributions to price bonds in the context of a consumption-

based asset pricing model also suggests that time-variation in recovery rate distribution is likely to have an impact on spreads - mainly that it amplifies the large swings in spreads over the cycle caused by time-variation in default probabilities. On average over the cycle, however, the effect on the level of spreads are likely to be very small even if risk aversion is very large.

The rest of this paper is structured as follows: The model is presented in section 2. In section 3 we describe the data set used. Section 4 discusses the various different versions of the models we estimated and the estimation results, and sections 5 and 6 explore the implications for credit risk management and pricing respectively. Finally, section 7 concludes.

## 2 The model

In terms of the data generating process, the model can be thought of as follows: Time is discrete. In each period, the state of the credit cycle is determined by the evolution of a two-state Markov chain. This implies that the dynamics of the state of the credit cycle are parameterized by two probabilities, i.e the economy can either be in either state 1 or state 0, and if the economy is in state 1 it will remain in this state with a probability  $p$  and move into state 0 in the next period with probability  $1 - p$ . If it is in state 0, it will remain in this state with probability  $q$  and move into state 1 with probability  $1 - q$ . Given the state of the credit cycle, the number of defaulting firms is then drawn using a state-dependent default probability. For each defaulting firm, recovery rates are drawn from a state-dependent recovery rate distribution.<sup>2</sup>

We assume that conditional on observable information and the unobserved state of the credit cycle, defaults are independent, recoveries in different default events are independent and the number of defaulting firms and recoveries are also independent. As a consequence, unconditional dependence between default events (i.e. dependence when not conditioning on the unobserved credit cycle) is driven entirely by the unobserved state of the credit cycle. Intuitively, when the state of the credit cycle is not known, observing the default of company A would make it more likely that the credit cycle is

---

<sup>2</sup>We do not impose the negative relationship between default rates and average recovery rates that is apparent in the data as part of the model. This negative relationship does, however, emerge once the model is estimated, which will allow us to label the state with high default rates and low recoveries, or low default rates and high recoveries as a credit ‘downturn’ and ‘upturn’ respectively.

in a downturn, especially if we were to observe a low recovery. This would increase the probability that we would attribute to seeing company B default with a low recovery. The events of company A and company B defaulting with low recovery are therefore not independent if we do not know the state of the credit cycle. The credit cycle here plays a role very similar to a dynamic frailty variable; in the context of defaults, dynamic frailty variables have been used previously e.g. by Duffie et al. (2006).

In our case, we will be able to estimate the probability of being in either state of the credit cycle by using either data on default rates, or data on recovery rates, or both, since both default probabilities as well as recovery rate distributions vary with the cycle.

We also need to specify the relationship between recoveries in the same default event. Although we assume that recoveries across different default events are independent conditional on the state, we clearly cannot assume that the recovery on a senior bond and the recovery on a junior bond issued by the same firm, which are observed during the same default event, are independent. In this case, we let parameters of the recovery rate distribution of any junior bond depend on the recoveries of all bonds senior to this bond. More precisely, we let the parameters depend on the average deviation of the senior recoveries from the mean recovery of their respective seniority class in the whole sample.

We also make the following assumptions with respect to specific functional forms of default probabilities and recovery rate distributions. Firstly, conditional on the state of the credit cycle (and possibly other explanatory variables), the arrival of defaults are described by discrete hazards of the form

$$\lambda_t = (1 + \exp \{ \gamma_0 + \gamma_1 c_t + \gamma_2 X_t \})^{-1}, \quad (1)$$

where  $c_t$  is the state of the cycle, a binary variable that is 0 or 1, and  $X_t$  represent other variables that might or might not be included in various versions of the model - for example, we might let hazards depend on log GDP growth.<sup>3,4</sup> This functional form of the discrete hazard has been used by Shumway (2001).

---

<sup>3</sup>Since we do not want to specify laws of motion for macroeconomic variables, we interpret the likelihood estimates as partial likelihood estimates whenever these variables are included in the different versions of the model.

<sup>4</sup>Note that we do not include firm-specific determinants of default probabilities, although it would be possible to do so in this framework. This is due to limitations of our data, which we discuss in the next section.

Secondly, we assume that the recovery rate  $y_{tis}$  for a default at time  $t$  of firm  $i$  on debt with seniority class  $s$  is drawn from a beta distribution. This distribution is well suited to modelling recoveries as it has support  $[0, 1]$ , is relatively flexible and requires only two parameters (which we call  $\alpha$  and  $\beta$ ). It is in fact often used by rating agencies for this purpose (see e.g. Gupton and Stein (2002)). We let the parameters of this beta distribution depend on the state of the cycle  $c_t$  and macroeconomic variables ( $X_t$ ) as in the case of hazard rates, as well as the industry sector of the firm ( $Z_i$ ) and the seniority of the instrument on which recovery is observed ( $S$ ). In order to ensure positivity of  $\alpha, \beta$ , we choose an exponential specification:

$$\alpha_{tis} = \exp \{ \delta_0 + \delta_1 c_t + \delta_2' Z_i + \delta_3' Z_i c_t + \delta_4' S + \delta_5' S c_t + \delta_6 X_t \} \quad (2)$$

$$\beta_{tis} = \exp \{ \zeta_0 + \zeta_1 c_t + \zeta_2' Z_i + \zeta_3' Z_i c_t + \zeta_4' S + \zeta_5' S c_t + \zeta_6 X_t \} \quad (3)$$

We obtain a likelihood function of the model as described in Appendix A by employing a slightly modified version of the method proposed by Hamilton (1989).

### 3 Data

The recovery data and the information on the defaulting firms is extracted from the Altman-NYU Salomon Center Corporate Bond Default Master Database. This data set consists of more than 2,000 defaulted bonds of US firms from 1974 to 2005. Each entry in the database lists the name of the issuer of the bond (this means that we can determine its industry as described by its SIC code), the date of default, the price of this bond per 100 dollars of face value one month after the default event, and a ‘bond category’.

The bond category is one of either Senior Secured, Senior Unsecured, Senior Subordinated, Subordinated or Discount. Altman and Kishore (1996) and Altman et al. (2005) label this bond category as ‘seniority’, and we will follow this convention here at times. It should be pointed out, though, that it also appears to contain information on the presence or absence of collateral (Senior Secured versus the other categories), and the type of coupon (Discount versus the other categories), which we will take into account when interpreting results.

Typically, for a given firm and date, the database contains the prices of many bonds for a given seniority class, so we need to aggregate. We do this by taking averages



(weighted by issue size).<sup>5</sup>

We also aggregate data across time into periods corresponding to calendar years.<sup>6</sup> We assume that a default of the same firm within twelve months of an initial default event (e.g. in December and then March of the following year) represent a single default event.

We calculate the recovery rate as the post-default price divided by the face value. Some of the recovery rates calculated in this way are larger than 1. This probably reflects the value of coupons. We scale the recovery rates by a factor of .9 to ensure that our observations lie in the support of the beta distribution. We could achieve the same effect by defining our distribution to be a beta distribution over the support that equals the range of our actual data.

The data set does not contain the total size of the population of firms from which the defaults are drawn, or equivalently, the number of non-defaulting firms. This variable is necessary in order to calculate the likelihood of observing a given number of defaults in a given period (cf. equation [5]). To determine the population size we use default frequencies as reported in Standard & Poor's Quarterly Default Update from May 2006 (taken from the CreditPro Database). Dividing the number of defaulting firms in each year in the Altman data by Standard & Poor's default rate, we can obtain a number for the total population of firms under the assumption that both data sets track the same set of firms, or that both data sets exhibit the same default rate. The available default rate data ranges from 1981 to 2005, so we lose observations from periods 1974-1980. However, these years cover only 9 observations in the Altman data.

After all adjustments, we obtain a combined data set which is structured as follows: We have 25 years of data. In each year, we know the size of the population of firms, and know the subset of the firms which defaults. For each default event, we observe one or more recoveries. In total, the final data set contains 1,124 recovery rate observations related to 898 default events. Descriptive statistics for these observations are presented in Tables V and VI in Appendix D. Table V reports the yearly statistics of our adjusted

---

<sup>5</sup>This is something commonly used in the literature. See for example Varma and Cantor (2005).

<sup>6</sup>We have to aggregate to calendar years rather than quarters because our default rate data is annual. Aggregating to quarters might allow obtaining more accurate transition probabilities. Nonetheless, the use of the coarser annual periods does not suppose a big loss in terms of our estimations. Also, in a lot of applications, e.g. the Basel II regulatory framework, the risk measures of interest are annual and not quarterly.

data set. Table VI reports the number of recovery rate observations and the mean and standard deviation of recoveries classified by seniority and industry respectively. These are in line with those reported in other papers, although on average recoveries are slightly lower here.

In some of our specifications we let default rates and recovery rates distributions depend on macroeconomic and economy-wide variables. We use GDP, investment growth and unemployment; the first two variables are extracted from NIPA tables, provided by the Bureau of Economic Analysis of the United States, while unemployment data comes from the Bureau of Labor Statistics of the U.S. Department of Labor. We also use annual returns on Standard & Poor’s 500 index, annual realized volatility of the Standard & Poor’s 500 index, and the VIX volatility index (all obtained from Yahoo Finance), the slope of the term structure (the 10-year yield minus 2-year yield, obtained from the Federal Reserve website) and NBER recessions as published by the NBER.

An important caveat of our analysis is related to the fact that we do not have firm-specific information on the non-defaulting firms, which implies that we cannot estimate default probabilities that depend on firm-specific information. This means that we can only make statements about aggregate default risk. Since we have firm- and security-specific information for defaulting firms, we can make some statements about specific recovery risk. The lack of firm-specific information on non-defaulting firms, does, however, limit the extent to which we can decompose the combined aggregate risk into its constituents.

## 4 Estimation results

In this section, we compare various estimated versions of the model across several dimensions. In the following, we present some short summary tables, for full tables see Appendix D.

**Macroeconomic variables versus the cycle** We address the question of whether our “credit cycle” is more or less useful than macroeconomic variables in explaining default rates and recovery rates by comparing a completely static version of the model (Model 1) in which  $\lambda_t$ ,  $\alpha_{tis}$  and  $\beta_{tis}$  do not vary (cf. equations [1],[2],[3]) with versions in which these parameters are either functions of the unobserved credit cycle or

macroeconomic variables, or both (Models 2-4). In order to check the fit of the different versions, we report Akaike information criteria and the Bayesian information criteria in Table I (for full parameter estimates see Appendix D).

**Table I**  
**Summary of selected models**

Akaike Information Criterion, Bayesian Information criterion for four different versions of the model.

	Model 1	Model 2	Model 3	Model 4
Explanatory variables for $\lambda$ (default probability)				
	constant	constant cycle	constant log GDP growth	constant cycle log GDP growth
Explanatory variables for $\alpha, \beta$ (recovery rates)				
	constant	constant cycle	constant log GDP growth	constant cycle log GDP growth
	seniority <sup>a</sup>	seniority <sup>a</sup>	seniority <sup>a</sup>	seniority <sup>a</sup>
Log Likelihood	84.023	295.87	135.58	307.88
AIC	-37.73	-405.31	-128.96	-412.79
BIC	0.0727	-0.1549	0.0188	-0.1290

<sup>a</sup>This includes seniority dummies, dummies that indicate whether a more senior recovery is observed for a given default event, and various interaction terms. See the paragraph on ‘Seniority’ for a full description.

It can be seen that introducing an unobserved credit cycle (Model 2) leads to a great improvement in fit over a completely static model (Model 1) as measured in terms of the information criteria. The coefficients on the cycle are highly significant. We also check the recovery rate densities implied by the models using a method proposed by Diebold et al. (1998), described in Appendix C. The basic idea is that applying the probability integral transform based on the predicted (filtered) distributions to the actual observations (the recovery rates in our case) should yield an i.i.d.-uniform series (the PIT series) under the null hypothesis that the density forecasts are correct. The correlograms and QQ-plots of PIT series presented in Appendix E indicate that the model with the unobserved credit cycle does a much better job at matching the time-variation in recovery rate distributions.

Including a macroeconomic variable instead of the cycle, in this case log GDP growth

(Model 3), also produces an improvement over the fit of the static model (Model 1), but it can be seen that this improvement is much more modest. Although log GDP growth is a significant determinant of default probabilities, it is not a significant determinant of recovery rates. In terms of the BIC, which strongly favours more parsimonious models, the model with constant and cycle (Model 2) fits even better than the model with constant, cycle and log GDP growth (Model 4).

We also ask whether the cycle is necessary at all, and compare a model in which we have only log GDP growth as an explanatory variable (Model 3) to one in which we have log GDP growth and the cycle as explanatory variables (Model 4). Like Duffie et al. (2006) we can calculate the Bayes factor, which is twice the difference in the log likelihood of the model with the latent variable (Model 4 in this case) and the log likelihood of the model without the latent variable (Model 3 in this case). This statistic is on the same scale as the likelihood ratio test statistic, but has a Bayesian interpretation. Duffie et al. (2006) cite literature that suggests that a Bayes Factor between 2 and 6 provides positive evidence for the model with the latent variable, and a Bayes factor between 6 and 10 suggest strong evidence. Our Bayes Factor is 344.60, which indicates very strong evidence in favor of the model that includes the latent cycle.

We also compare a version of Model 3 against a version of Model 4 with a frequentist test. Note that under the null that the cycle is not necessary, some coefficients, and the transition probabilities are not identified, which means that standard asymptotic results do not apply. We therefore use the test proposed by Hansen (1992) (cf. also Hansen, 1996), where in both models we drop some interactions terms for reasons of computational feasibility.<sup>7</sup> The p-value of the test statistic is 0.0019, which given the conservative nature of this test is very strong evidence against the null that the cycle does not matter.

The picture looks very similar for the other macroeconomic and economy-wide variables that we considered (unemployment, investment, the slope of the nondefaultable term-structure, the S&P 500 return, the realized volatility of the S&P 500 return, the

---

<sup>7</sup>The test requires restricted estimation at each point of a grid whose dimension depends on the number of unidentified parameters. Through dropping interactions of the cycle with other variables in Model 4, we reduce the dimensionality of the grid from 29 to 5. Following Hansen, for each transition probability we consider 10 values, and for each coefficient 20 values. We also drop interactions of log GDP growth with other variables in Model 3 and Model 4 to speed up estimation for each point on the grid.

VIX<sup>8</sup> and an NBER indicator).

We interpret this as saying that the default rate and recovery rates are much more tightly related to each other than to macroeconomic variables; or that there is more information in recovery rates about default rates, and vice versa, than there is information in e.g. log GDP growth about default rates. These results are consistent with those of Altman et al. (2005).

We also calculate the out-of-sample RMSE of our recovery rate predictions in a manner similar to the calculation performed by Chava et al. (2006), using the expected recovery rate of Model 2. We estimate the model on the data up until 1996, predict the state of the credit cycle for 1997 given the information up until 1996, and calculate the expected recovery rate under the model for each recovery rate observation in 1997, and calculate the corresponding prediction error. We then roll the estimation forward, estimate using data up until 1997, and predict recovery rates for 1998, and so on. The RMSE for these out-of-sample predictions is 22.86%, which compares favourably with the lowest out-of-sample RMSE obtained by Chava et al. (2006), which is 24.96%, even though these numbers are probably not directly comparable since they are based on different datasets. Chava et al. (2006) try models based on many interesting observed firm-specific, security-specific and economy-wide covariates, but do not consider lagged aggregate default rates or recovery rates.


To make sure that our assumptions about dependence between recoveries on the same default event do not drive the results, we also estimated all models dropping all default events for which we observe more than one recovery (this is roughly a quarter of our data). Parameters estimates, filtered and smoothed probabilities of the cycle, as well as the relative sizes of the incremental improvement in fit produced by including either the cycle or macroeconomic variables are virtually unaffected.

To qualify how the dynamics of macroeconomic variables differ from the dynamics of credit variables, we plot the (smoothed) probabilities of being in the state of the Markov chain which we label as a credit downturn over time and compare these to NBER recession dates in Figure 2 (this state is associated with a higher default probability and a lower mean recovery). As can be seen, we do not pick up the recession of 1981, probably because although average recovery rates are very low for 1981, default rates

---

<sup>8</sup>Since we only have data for the VIX starting in 1990, we had to restrict ourselves to a subsample of the data in this case.

are also very low in our data, and we only have a single recovery observation for this year. It can be seen that the credit cycle as measured by our Markov chain had two major downturns, around the recession of the early 90s and around 2001, but in each case, the credit downturn started before the recession, and in the latter case ended after it.



**Figure 2. Probabilities of being in the credit downturn versus NBER recessions (annual)**

Annual smoothed probability of being in a credit downturn as estimated on the basis of the version of the model that only uses the unobserved credit cycle as an explanatory variable. This is contrasted with NBER recessions (grey areas).

To further investigate the difference in timings of credit and business cycle downturns, we check whether one of the credit variables leads the business cycle by running bivariate Granger causality test of default rates and log GDP growth, as well as average recovery rates and log GDP growth. While there appears to be no Granger causality in either direction between default rates and log GDP growth for our data, we find that average recovery rates at lags of 1 year, or in terms of quarterly data at lags of 2 and 4 quarters strongly Granger-cause log GDP growth.

In the following, we focus on models that contain only the credit cycle as a time-varying variable, since the unobserved “credit cycle” seems to be a better way of de-

scribing the data than using macroeconomic explanatory variables.<sup>9</sup>

**Default Rates** In the static model (Model 1), the model-implied default probability is constant at 1.55%. In one version of a model including a cycle (Model 2), the default probability in downturns is 2.77% and 0.87% in upturns (these are similar in all models that include a cycle).

To check the goodness of fit of the cycle-implied default rates, we regress observed default rates on a constant and estimated default probabilities of Model 2. These estimated default probabilities are obtained by mixing, for each period, the default probabilities of both states where the weightings are the (smoothed) probabilities of being in one state or the other. Given the coefficients of this regression we can test the default rate component of both the static and dynamic model. The null hypothesis for the dynamic model (that Model 2 correctly captures time-variation in default rates) is that the constant and the coefficient of the estimated default probability are equal to zero and one respectively. The p-value associated with this restriction is 0.9641, indicating that this null is not rejected.<sup>10</sup> We also explicitly consider restricting the coefficient of the estimated default probability to zero, and letting the constant be unrestricted. This restriction would correspond to the static model (Model 1). The p-value associated with this restriction is less than  $10^{-5}$ , indicating that the static model (Model 1) does not describe the data well.

**Seniority (bond category)** We can examine the behaviour of recovery rate distributions for different seniorities over the cycle on the basis of Model 2, where here we use the rather broad definition of seniority used by Altman and Kishore (1996) and Altman et al. (2005), i.e. their ‘bond category’ which also includes information on the presence of collateral and coupons (see Section 3 for details).

We can distinguish two types of default events in our data, one for which we observe only a single recovery associated with a given default event, and one for which we observe multiple recoveries of different seniority classes associated with the default event. For

---

<sup>9</sup>Note that this also has the effect of simplifying forecasting. In order to forecast from models that use macroeconomic variables as covariates we would need to posit an appropriate law of motion for these variables, an issue which we avoid here.

<sup>10</sup>The p-values are those of standard t-tests. In order to verify that these are appropriate, we bootstrapped the t-statistics of the regression to confirm that they do indeed follow a t-distribution.

the second kind of event, we create a dummy which is one for the most senior bond, and zero otherwise, as well as a dummy that is one for the junior bonds and zero otherwise. We interact these dummies with the seniority dummies. The idea here is that seniority can mean different things depending on the debt structure of the firm. For example, the recovery on a Senior Unsecured bond might be higher in situations where we also observe recovery on a junior bond, since in this case the first losses are taken by the holders of the junior bond; the junior bond performs the role of a “junior debt cushion”.

**Table II**  
**Implied Mean Recoveries**

Implied mean recoveries in percent for Model 2, by state of the credit cycle (upturn or downturn), bond category (ranging from Senior Secured to Discount), and by number of recovery observations per default event. “Single rec.” means events for which recovery is observed only on a single bond, “Multiple rec., senior bond” means recoveries on the senior bond in default events for which several recoveries are observed, and “Multiple rec., junior bonds” means recoveries on the junior bonds in default events for which several recoveries are observed. Note that here, we follow the convention of Altman and Kishore (1996) in labeling the bond category as ‘seniority’.

		Upturn	Downturn
Single rec.	Sen. Sec.	54.7	32.0
	Sen. Unsec.	47.5	31.2
	Sen. Sub.	41.6	32.0
	Sub.	41.8	31.7
	Discount	33.0	19.1
Multiple rec., senior bond	Sen. Sec.	73.8	35.7
	Sen. Unsec.	48.9	32.8
	Sen. Sub.	41.2	27.4
	Sub.*	-	-
Multiple rec., junior bonds	Sen. Unsec.	30.8	37.2
	Sen. Sub.	21.5	26.8
	Sub.	29.7	20.7
	Discount	26.1	17.5

\* We have no observations for which we observe several recoveries, and the recovery on a Subordinated bond is the recovery on the most ‘senior’ bond.

We calculate the average recoveries implied by our estimates for different seniority classes for all types of events in Table II. We can see for example that in default events where recovery is observed only on a single bond, senior unsecured bonds have a higher recovery in upturns on average (47.5%), they have very similar (low) recoveries to bonds



of all other seniorities in downturns (31.2%). Similarly to e.g. Acharya et al. (2007), we also find that this effect is more pronounced for senior securities. Acharya et al. (2007) also provide evidence that in their data, the effect is not present if the debt is secured. Comparing our estimated recoveries on Senior Secured bonds and Senior Unsecured bonds in both states of the cycle, we find that in our data, senior debt is more strongly affected by a credit downturn if it is secured. The difference in results might be due to differences in the definition of secured debt, or due to differences between industry distress as defined by Acharya et al. (2007) and the credit cycle as defined here.

Lastly, one would expect that a “junior debt cushion” in a given default event matters especially for Senior Unsecured bonds, since in this case, all bonds share the same collateral and any losses are first borne by the junior bond(s). Conversely, for a Senior Secured bond the existence of a “junior debt cushion” should matter less, since in this case, the Senior Secured bond has bond-specific collateral. Surprisingly, the existence of a “junior debt cushion” seems to matter (in upturns) for secured debt, but not for unsecured debt; we can see that Senior Secured bonds recover more in upturns in situations in which a junior bond exists (73.8% versus 54.7% in upturns), whereas the recoveries on Senior Unsecured bonds are not really affected by whether or not a junior bond exists. To check that this is a feature of the data and not an artifact of any modeling assumption, we also regress recoveries of Senior Secured bonds on a constant and a dummy that is one if a cushion exists (a recovery is observed on another bond). For our data the estimated coefficients are 39.1% and 11.6% on the constant and on the dummy respectively, implying that recoveries on Senior Secured bonds are 39.1% on average if no cushion exists, and 50.7% on average when some type of cushion does exist. The coefficient on the dummy is also significant at 1%. This does appear to be a feature of the data.

**Industry** Lastly, we also consider the industry of the issuer as an additional explanatory variable for  $\alpha_{tis}$  and  $\beta_{tis}$ . Since for some industries, we have a very low number of recovery and default observations in upturns, we aggregate industries into three groups: Group A: Financials, Leisure, Transportation, Utilities; Group B: Consumer, Energy, Manufacturing, Others; and Group C: Construction, Mining, Services, Telecoms. For our data, these groupings roughly correspond to industries with high, medium and low

mean recovery rates respectively.<sup>11</sup> All industry groups show recovery rate distributions that vary over the cycle. Detailed results for this model (Model 5) are in Appendix D.

## 5 Implications for Risk Management

Altman et al. (2005) cite previous work (Altman et al., 2001) that states that allowing for dependence between default rates and recovery rates produces a 99% VaR of a percentage loss of about 4.9% for a representative portfolio, whereas the VaR which is calculated on the basis of a model that assumes no dependence between default rates and recovery rates is around 3.8%. In relative terms, this is an increase of around 29% in the VaR. These number are based on a very simple static model into which they substitute guessed parameter values. We are in a position to compare the VaR calculated when assuming recovery rate distributions are either static or time varying, but on the basis of an estimated model.

We calculate (by simulation of 10,000 paths) the one-year loss distribution of a hypothetical portfolio of 500 senior unsecured bonds, issued by 500 separate issuers that are representative of the firms in our dataset. We compare the 99% VaR assuming that the bonds in the portfolio are either well described by a static model (Model 1) or by dynamic model which includes the credit cycle as an explanatory variable, where we look at a version of the model in which both default probabilities and recovery rates are functions of the credit cycle (Model 2), and a version in which only default probabilities are a function of the credit cycle and recovery rate distributions are static (Model 2a), and only recovery rate distributions are a function of the credit cycle and default probabilities are static (Model 2b). This allows us to examine the impact on risk calculations of ignoring the time-variation either in recovery rate distributions or in default probabilities.

Note that when estimating the models, the identification of the state can come either through recovery rates, or through the default rate, or through both. If we assume that recovery rate distributions vary, but default probabilities are constant over the cycle, the state is identified only through the time-variation in recovery rate distributions. Intuitively, this is the case because the number of defaults (or the default rate) in this

---

<sup>11</sup>The groupings were chosen to minimize the LR statistic of a test that restricts the coefficients on industry dummies to be the same within these groupings, i.e. to choose the most likely grouping given the data. The p-value of the likelihood ratio test for the given groupings is 10%.


case does not contain information about the state of the cycle (for a mathematical argument, see Appendix B). Conversely, if we assume that default probabilities vary with the cycle, but recoveries do not, then the state is identified only via the variation in the number of defaults (the default rate). When we estimate Models 2, 2a and 2b, the smoothed probabilities of being in either state as well as the transition probabilities vary very little across the models, even though identification is coming through different subsets of the variables. This reflects the close link between recovery rates and default rates evident e.g. in Figure 1.

Once the models are estimated, we can simulate from them. For this we have to choose the probability that we attach to being in a credit downturn today. We examine cases with this probability being equal to one (we know that we are in a downturn today), zero (we know that we are in an upturn today) and 33.5%, which corresponds to the unconditional probability of being in a downturn, given our estimated transition probabilities, i.e. under the assumption that we have no information on the current state of the cycle. The 99% VaRs calculated in this fashion are in Table III.

A representative loss density, the one implied by the dynamic model (Model 2) based on the unconditional probability of being in a downturn is compared to the loss density implied by the static model (Model 1) in Figure 3 (for more plots of loss densities, see also Figure 6 in Appendix E). The dynamic model loss density is bimodal, reflecting the possibility of ending up either in an upturn with low default probabilities and high recoveries, or ending up in a downturn, with high default probabilities and low recoveries. As can be seen, the tail of the loss distribution implied by the dynamic model is much larger.

Table III shows that the 99% VaR ranges from a loss of 2.7% to a loss of 3.1% of the portfolio assuming the world is dynamic as described in Model 2, versus a 1.9% loss on the portfolio if we had assumed that the world is static as in Model 1. Even supposing that we are in an upturn today, in which cases losses over one year are likely to be smaller, the dynamic model (Model 2) still produces a 99% VaR of 2.7%, which is larger than the VaR based on the static model (Model 1). This is because in the dynamic model, even though we are in an upturn today, we might go into a credit downturn tomorrow, with the associated higher default rates and lower recoveries.

How much of this extra risk in the dynamic model is due to the variation in recovery rates, though? Assuming that recovery rates are drawn independently from distribu-



### Figure 3. Loss Densities

One-year Loss densities implied by a static model (Model 1) and a dynamic model (Model 2) for a portfolio of 500 senior unsecured bonds issued by 500 different issuers. The initial probability of being in an upturn is assumed to be equal to the unconditional probability. See Figure 6 in Appendix E for more plots of loss densities.

tions that do not vary over time, or that recovery rates are static parameters is very common in pricing and risk management applications. We therefore also calculate the loss distribution on the basis of Model 2a, which assumes that recovery rate distributions are constant, but that default probabilities vary with the credit cycle. The 99% VaR produced by this Model 2a ranges from 2.4% to 2.8%, which is only slightly lower than the VaR produced by the fully dynamic model (Model 2). Suppose we focus on the case in which we assume that the probability of being in an upturn initially is equal to the unconditional probability of being in an upturn. Going from a model in which both default probabilities and recovery rate distributions are static to one in which default probabilities and recovery rate distributions vary with the cycle increases the VaR by a factor of 1.53 or by 53%. We can decompose this increase by comparing VaRs of Model 1, 2 and 2a. We can see that of these 53%, 37% are due to letting the default probabilities vary, and only about 16% are due to the additional amplification effect

**Table III**  
**Value-at-Risk**

One-year 99% VaRs for hypothetical portfolios of 500 senior unsecured bonds (issued by 500 different issuers) based on a static and two dynamic models. In the dynamic models, default probabilities depend on the cycle, and recovery rate distributions may or may not depend on the cycle. For these models, we calculate loss distributions assuming that the probability of being in a downturn initially are 0 (“upturn”), equal to the unconditional probability (“uncond.”), or equal to 1 (“downturn”).

MODEL SPECIFICATIONS				
	Model 1	Model 2	Model 2a	Model 2b
Explanatory variables for $\lambda$ (default probability)				
	constant	constant cycle	constant cycle	constant
Explanatory variables for $\alpha, \beta$ (recovery rates)				
	constant	constant cycle	constant	constant cycle
	seniority	seniority	seniority	seniority
Init. state				
	ANNUAL 99%-VAR			
upturn	1.9%	2.7%	2.4%	1.8%
uncond.	1.9%	2.9%	2.6%	1.9%
downturn	1.9%	3.1%	2.8%	2.1%

of letting recovery rates vary over the cycle as well. Although certainly not negligible, this is smaller than the 29% suggested by the calculations of Altman et al. (2001).

Conversely, we can estimate a model in which default probabilities are constant, but recovery rate distributions vary over the cycle (Model 2b). The VaR that this model produces is essentially the same VaR of the static model, indicating again that the variation in default probabilities over the cycle has a larger impact on losses than the variation in recovery rate distributions.

For different seniorities, or looking at the model that includes industry group dummies the picture looks very similar, so we do not report the corresponding VaRs here. Essentially, it appears that although the time-variation in recovery rate distributions does amplify losses, quantitatively speaking, having about 4-6 times the number of defaulting firms in downturns is more important than losing an extra 15% or so per default in downturns.

We can also phrase this point in the context of the expected loss on a portfolio (for this, we do not even need our estimated models). Altman et al. (2001) suggest the following numerical example: Suppose that the default probability in upturns is  $r_U = 2\%$  and the default probability in downturns is  $r_D = 10\%$ , and the loss given default (1 minus the recovery rate) is a constant  $L_U = 30\%$  in upturns and a constant  $L_D = 70\%$  in downturns. Assume the probability of being in an upturn is  $1/2$ .

The expected loss for a large portfolio, calculated using the assumption that there is no relationship between loss given default and default probability would be

$$\text{expected loss} = E[L] \times E[r] = 6\% \times 50\% = 3\%.$$

Of course, expected loss and default probability are related here through the state of the cycle. Taking this into account, we have

$$\begin{aligned} \text{expected loss} &= E[L \times r] = L_U \times r_U \times \frac{1}{2} + L_D \times r_D \times \frac{1}{2} \\ &= \frac{1}{2} \times 30\% \times 0.02 + \frac{1}{2} \times 70\% \times 0.1 = 3.8\% \end{aligned}$$

The difference is  $E[L \times r] - E[L] \times E[r] = 80\text{bp}$ . We could calculate this on the basis of our estimated models, but note that this difference is just the covariance of  $L$  and  $r$ ,  $E[Lr] - E[L]E[r] = \text{Cov}(L, r)$ . If the annual default rate is a reasonable estimator of the annual default probability, and the annual average recovery rate is a reasonable estimator of 1 minus the annual loss given default, then the sample covariance between these two quantities in our data can be interpreted as an estimator of this difference in expected losses. This covariance is 5bp (rather than 80bp, as suggested by the above calculation) in our data. This indicates that the amplification effect of recovery rates on expected loss is relatively small.

## 6 Implications for Pricing

The time variation in recovery rate distributions should also have an effect on bond prices and spreads; since risks are larger in a world where recovery rate distributions vary over time than in one where they are static, spreads should be higher. We explore this in the context of a consumption-based asset pricing model. It is well known that the price  $P_t$  of any asset with a single payoff  $X_{t+1}$  can be expressed in terms of a stochastic

discount factor  $M_{t+1}$  as

$$\begin{aligned}
 P_t &= E_t [M_{t+1} X_{t+1}] \\
 P_t &= E_t [M_{t+1}] E_t [X_{t+1}] + \text{Cov}_t [M_{t+1}, X_{t+1}] \\
 &= \frac{1}{1 + r_f} E_t [X_{t+1}] + \text{Cov}_t [M_{t+1}, X_{t+1}]
 \end{aligned} \tag{4}$$

where  $r_f$  is the risk-free return (see e.g. Cochrane, 2001), i.e. the price can be decomposed into an expected discounted payoff component, and a risk premium. To obtain a large negative impact of recovery rates on prices, we would need for recovery rates to have a large negative effect on the expected payoffs, and / or have a large negative effect on payoffs especially in situations in which the stochastic discount factor is high, such that the risk premium is high.

In a standard power utility formulation, we have that  $M_{t+1} = \beta (C_{t+1}/C_t)^{-\gamma}$ , where  $\beta$  is the subjective discount rate,  $\gamma$  is the coefficient of relative risk aversion, and  $C_t$  is aggregate consumption at  $t$ . Through our econometric model, we have already specified the dynamics of default probabilities and recovery rate distributions. In order to use [4] to price, we need to make additional assumptions about a filtration (the information available to agents), and the dynamics of consumption, and how they relate to the dynamics of default probabilities and recovery rate distributions. An easy way to relate the dynamics is to assume that all variables are driven by the same latent Markov chain, the “credit cycle”. When matching a given unconditional variance of consumption, this is likely to overstate the magnitude of the (negative) correlation between consumption and bond payoffs, and therefore the risk premium.<sup>12</sup> With respect to the filtration, we assume that agents know the current state of the credit cycle, to simplify the pricing exercise.<sup>13</sup>

We assume that the value of log consumption growth is 1% in downturns and 4% in upturns; this is higher than the average values in years which we identify as credit cycle downturns and upturns respectively, but produces an unconditional variance of log consumption growth which roughly matches the one in our sample. This also produces a (positive) correlation between average recovery rates and log consumption growth that is much larger (by a factor of about 2.5) than the one observed in the data,

---

<sup>12</sup>Bond payoffs and the stochastic discount factor are not perfectly dependent here because bond payoffs are not known conditional on knowing the state.

<sup>13</sup>This seems reasonable given that our filtered probabilities are either zero or 1 most of the time.

and a (negative) correlation between default rates and log consumption growth whose magnitude is about 3 times the magnitude of the one we observe in the data. We initially set the coefficients of relative risk aversion to  $\gamma = 2$ , and set the subjective discount factor to  $\beta = 0.98$ . Collectively, these assumptions imply a risk-free interest rate of about 9.7% in upturns and 5.7% in downturns. We then also try a coefficient of relative risk aversion of  $\gamma = 35$ , which produces unrealistically high risk-free rates, but maybe more realistic risk premia.

Combining these assumptions with our previous estimates for transition probabilities, default probabilities and recovery rate distributions, we can calculate the resulting spreads on a senior unsecured bond, for a world in which recovery rate distributions are static (Model 2a), and for one in which recovery rate distributions vary (Model 2). These are shown in Table IV. We also calculate the spreads that would result if agents were risk neutral, to gauge what proportion of the spread is a risk premium, and what proportion compensates for expected losses.

**Table IV**  
**Calibrated Spreads**

Spreads of zero coupon bonds with a maturity of 1 year (based on [4]) for a coefficient of relative risk aversion of  $\gamma = 2$  and  $\gamma = 35$ , as well as spreads under risk-neutrality, based on prices calculated as expected discounted payoffs, conditional on the initial state of the cycle, which is either an upturn or downturn.

PANEL A: No variation in recovery rates (Model 2a)			
Init. state	$\gamma = 2$	$\gamma = 35$	risk-neutral
upturn	70	91	69
downturn	147	166	146
PANEL B: Variation in recovery rates (Model 2)			
Init. state	$\gamma = 2$	$\gamma = 35$	risk-neutral
upturn	64	86	63
downturn	149	171	147

We can see that for the given consumption-based asset pricing model, the risk premium component is negligible unless the coefficient of absolute risk aversion is very high, in spite of the fact that we have overstated the magnitude of the negative covariance between consumption growth and payoffs. Also, spreads are much larger in downturns than in upturns in both models and regardless of the assumptions about risk aversion,



indicating that the fluctuations in the spread over the cycle here are caused by differences in expected losses in downturns and upturns, which in turn are driven mainly by variation in default probabilities. Letting recovery rate distributions vary increases the amplitude of the swings in spreads over the cycle slightly, since time-variation in recovery rate distributions amplify the effect of time variation in default probabilities. We can see, for example, that if we suppose that  $\gamma = 2$ , the difference in spreads between upturns and downturns in a model in which recovery rates vary (Model 2) is  $149-64=85$ bp. In a model in which recovery rates do not vary (Model 2a), it is slightly lower,  $147-70=77$ bp. However, the average spread (either the simple average over the two states, or weighting by unconditional probabilities) changes at most by a few basis points in moving from a model in which recovery rate distributions are static (Model 2a) to a model in which recovery rate distributions vary with the cycle (Model 2). The conclusion therefore is that allowing for time variation of recovery rate distributions in a pricing model could allow a slightly better matching of spreads over the cycle, but is unlikely to alter the ability of the model to match the level of spreads.<sup>14</sup>

Although the approach here could be extended by considering different types of stochastic discount factors, note that Chen et al. (2006) (using the discount factor suggested by Campbell and Cochrane (1999) which can produce higher risk premia) also find that accounting for variation in recovery rates only alters average spreads by at most a couple of basis points.

## 7 Conclusions

This paper proposes and estimates an econometric model of the systematic time-variation in recovery rate distributions and default probabilities. On the basis of our estimated model, we can state that the time-variation in recovery rate distributions does amplify risk, but that this effect is much smaller than the contribution of the time variation in default probabilities to systematic risk. On the basis of a consumption-based asset pricing model, we show that it is likely that the impact on average spreads is in the range of a couple of basis points, while potentially having a relatively larger impact on the variability of spreads across the cycle. We also present evidence that

---

<sup>14</sup>Using a recursive version of [4] to price longer maturity bonds does not produce qualitatively different results.

indicates that default rates and recovery rates are more tightly related to each other than to macroeconomic variables, and that credit downturns seem to be only imperfectly aligned with recessions; they start before recessions and last longer. The different phases of the business and the credit cycle seem to be particularly evident in recovery rates - average recovery rates actually Granger-cause log GDP growth in our data set.

The results here also suggest some interesting avenues for future research. For example, a drawback of our dataset is that it contains no firm-specific information on non-defaulting firms, which limits the extent to which we can decompose the time-variation in default probabilities. The methodology proposed here, however, could easily be extended to such datasets. This would allow examining not only the systematic time-variation in credit risk, but also any possible relationship between recovery risk and firm-specific default probabilities over time, which could yield important insights. Also, the different phases of the credit cycle and the business cycle which are particularly apparent in the lead-lag relationship between recovery rates and GDP growth warrant further theoretical and empirical attention.

# Appendix

## A The likelihood function

Let  $c_t$  be the unobservable state of the cycle (our Markov chain). Also, let  $p$  be the probability of remaining in state  $c_t = 1$ , and  $q$  the probability of remaining in state  $c_t = 0$ .<sup>15</sup>

Let  $d_t$  be the number of defaulted firms observed in period  $t$ . Firms are indexed by  $i$ . We denote the recovery observed for a default event at  $t$  of firm  $i$  on debt of seniority class  $c$  as  $y_{tis}$ . A firm can recover on more than a single seniority class. We collect the recoveries for a given default event in  $t$  for a given firm  $i$  into a vector which we call  $Y_{it}$ . All recoveries observed in a given period (across all firms) are denoted as  $Y_t$ .

It is easy to see that the number of defaulting firms  $d_t$  conditional on the state  $c_t$  of the cycle in any period (and the variables that influence hazards,  $X_t$ ) will be binomially distributed, given that we have assumed that firms default independently conditional on this information. The conditional default probability here is the discrete hazard described above. We can write

$$\Pr(d_t|c_t) = \binom{N_t}{d_t} \lambda_t^{d_t} (1 - \lambda_t)^{N_t - d_t}. \quad (5)$$

where

$$\lambda_t = (1 + \exp \{ \gamma_0 + \gamma_1 c_t + \gamma_2 X_t \})^{-1}, \quad (6)$$

This probability is also to be interpreted as being conditional on  $X_t$  and  $N_t$ , which is not explicit in the above notation.

For each default at  $t$  associated with firm  $i$  on debt of seniority class  $c$ , we observe a  $y_{tis}$ . The density of  $y_{tis}$  conditional on the state is

$$g(y_{tis}|c_t) = \frac{1}{B(\alpha_{tis}, \beta_{tis})} (y_{tis})^{\alpha_{tis}-1} (1 - y_{tis})^{\beta_{tis}-1}, \quad (7)$$

where

$$\alpha_{tis} = \exp \{ \delta_0 + \delta_1 c_t + \delta_2' Z_i + \delta_3' Z_i c_t + \delta_4' S + \delta_5' S c_t + \delta_6 X_t \} \quad (8)$$

$$\beta_{tis} = \exp \{ \zeta_0 + \zeta_1 c_t + \zeta_2' Z_i + \zeta_3' Z_i c_t + \zeta_4' S + \zeta_5' S c_t + \zeta_6 X_t \} \quad (9)$$

---

<sup>15</sup>In our estimated models,  $c_t = 0$  will correspond to a state with high default probabilities and low recoveries which we label as a ‘credit downturn’ and  $c_t = 1$  to a state with low default probabilities and high recoveries which we label a ‘credit upturn’, but we do not impose this in the structure of the model.

This density is the *marginal* density with respect to recoveries on other seniority classes (but associated with the same default event), but *conditional* on the state.

If we have more than one recovery event per default event, we allow for dependence between these recoveries in the following way. For a vector  $Y_{ti} = (y_{ti1}, y_{ti2}, y_{ti3}, y_{ti4})'$ , we can write its density as

$$\begin{aligned} h(Y_{ti}|c_t) &= g(y_{ti1}, y_{ti2}, y_{ti3}, y_{ti4}|c_t) \\ &= g(y_{ti1}|c_t)g(y_{ti2}|y_{ti1}, c_t)g(y_{ti3}|y_{ti2}, y_{ti1}, c_t)g(y_{ti4}|y_{ti3}, y_{ti2}, y_{ti1}, c_t) \end{aligned}$$

The density of the most senior bond  $y_{ti1}$  can then be treated as above. We allow for the fact that the conditional densities of  $y_{ti2}$ ,  $y_{ti3}$  and  $y_{ti4}$  in this expression can each be a function of the more senior recoveries by calculating variables  $\bar{m}_{ti2}$ ,  $\bar{m}_{ti3}$ ,  $\bar{m}_{ti4}$  which are functions of the more senior recoveries, and then letting  $\alpha_{ti2}$ ,  $\beta_{ti2}$ ,  $\alpha_{ti3}$ ,  $\beta_{ti3}$  and  $\alpha_{ti4}$ ,  $\beta_{ti4}$  depend on these  $\bar{m}$  respectively. The  $\bar{m}_{tis}$  are defined as the average deviation of the senior recoveries from the sample mean for that respective seniority class, that is

$$\bar{m}_{tiS} = \frac{1}{S-1} \sum_k^{S-1} m_{tik} \quad (10)$$

where

$$m_{tik} = y_{tik} - \frac{1}{T} \sum_t \frac{1}{d_t} \sum_i y_{tik}. \quad (11)$$

As a special case, this implies that when we do not observe recoveries across several seniority classes for a default event (this is the case for the majority of our observations), we simply have

$$h(Y_{ti}|c_t) = g(y_{tis}|c_t). \quad (12)$$

Putting recoveries and the number of defaults together, it is then easy to see that conditional on the state, the density associated with observing a given number of defaulting firms  $d_t$  and associated recoveries  $Y_t$  in time period  $t$  is given by

$$f(Y_t, d_t|c_t) = \binom{N_t}{d_t} \lambda_t^{N_t} (1 - \lambda_t)^{N_t - d_t} \prod_{i=1}^{d_t} h(Y_{it}|c_t), \quad (13)$$

where the understanding is that the last product is 1 if  $d_t = 0$ , and  $N_t$  is the number of firms in the population at time  $t$ .

The density of the data observed in period  $t$ , not conditional on the state, but conditioned on past information  $\Omega_{t-1}$  can be written as

$$f(Y_t, d_t | \Omega_{t-1}) = f(Y_t, d_t | c_t = 1) \Pr(c_t = 1 | \Omega_{t-1}) + f(Y_t, d_t | c_t = 0) \Pr(c_t = 0 | \Omega_{t-1}). \quad (14)$$

Note that given our assumptions, the past information  $\Omega_{t-1}$  is not relevant for the conditional density of  $Y_t, d_t$ . Since the state is not observed, we need a way of deriving the filtered (and smoothed) probabilities of being in either state. Obviously,

$$\Pr(c_t = 1 | \Omega_{t-1}) = p \cdot \Pr(c_{t-1} = 1 | \Omega_{t-1}) + (1 - q) \cdot \Pr(c_{t-1} = 0 | \Omega_{t-1}), \quad (15)$$

where we used the notation defined previously,  $p = \Pr(c_t = 1 | c_{t-1} = 1, \Omega_{t-1})$ , the probability of remaining in state 1, and  $(1 - q) = \Pr(c_t = 1 | c_{t-1} = 0, \Omega_{t-1})$ , the probability of moving from state 0 to state 1. We can rearrange [15] to obtain

$$\Pr(c_t = 1 | \Omega_{t-1}) = (1 - q) + (p + q - 1) \Pr(c_{t-1} = 1 | \Omega_{t-1}) \quad (16)$$

The probability  $\Pr(c_{t-1} = 1 | \Omega_{t-1})$  can be obtained via a recursive application of Bayes' rule:

$$\Pr(c_{t-1} = 1 | \Omega_{t-1}) = \frac{f(Y_{t-1}, d_{t-1} | c_{t-1} = 1) \Pr(c_{t-1} = 1 | \Omega_{t-2})}{f(Y_{t-1}, d_{t-1} | \Omega_{t-2})}. \quad (17)$$

Note that there is information in recoveries  $Y_t$  and  $d_t$ , about the state, i.e. identification of the state is obtained through both the number of defaults and recoveries.

We now have all elements to maximize our log likelihood function

$$L = \sum_{t=1}^T \log f(Y_t, d_t | \Omega_{t-1}), \quad (18)$$

which is a partial likelihood since we omit specifying the law of motion of  $X_t$  and  $N_t$ .

For a given parameter vector, the likelihood can therefore be calculated recursively, given some suitable initial conditions, e.g. the unconditional probabilities of being in each state implied by the transition matrix.

## B Identification of the state of the cycle

The state of the cycle is identified only via variables that depend on it, in the following sense:

1. When default probabilities are independent of the cycle and recovery rates are not, the filtered and smoothed probabilities of being in either state only depend directly on the number of defaults.
2. When recoveries are independent of the cycle and default probabilities are not, the filtered and smoothed probabilities of being in either state only depend directly on the recoveries.

To see that this is the case, suppose first that all values of recoveries  $Y_{it}$  do not depend on the cycle, i.e.  $h(Y_{it}|c_t) = h(Y_{it})$ , but that the number of defaults  $d_t$  does. We need to show that in the recursive calculation of the probabilities of the state of the cycle,  $Y_{it}$  does not enter. The prediction step, i.e. calculating  $\Pr(c_t = 1|\Omega_{t-1})$  from  $\Pr(c_{t-1} = 1|\Omega_{t-1})$  works as before and does not involve any data (see [16]). The filtering step, i.e. the step involving the calculation of  $\Pr(c_{t-1} = 1|\Omega_{t-1})$  is now slightly simpler:

As before, we define  $Y_t = \{Y_{it}\}$ . First note that

$$f(Y_{t-1}, d_{t-1}|c_{t-1}) = f(Y_{t-1}|d_{t-1}, c_{t-1})f(d_{t-1}|c_{t-1}) = f(Y_{t-1}|c_{t-1})f(d_{t-1}|c_{t-1})$$

because the value of recoveries and the number of defaults are independent conditional on the cycle as before. Additionally, if the  $Y_{it}$  now do not depend on the cycle, we know that  $f(Y_t|c_t) = f(Y_t)$ , and we can write

$$f(Y_{t-1}, d_{t-1}|c_{t-1}) = f(Y_{t-1})f(d_{t-1}|c_{t-1}). \quad (19)$$

and

$$f(Y_{t-1}, d_{t-1}|c_{t-1} = 1) = f(Y_{t-1})f(d_{t-1}|c_{t-1} = 1). \quad (20)$$

Combining the previous equations with [14] indicates that for the given case, we can write  $f(Y_{t-1}, d_{t-1}|\Omega_{t-2})$  as

$$\begin{aligned} & f(Y_{t-1}, d_{t-1}|\Omega_{t-2}) \\ &= f(Y_{t-1}) \{f(d_{t-1}|c_{t-1} = 1) \Pr(c_{t-1} = 1|\Omega_{t-2}) \\ & \quad + f(d_{t-1}|c_{t-1} = 0) \Pr(c_{t-1} = 0|\Omega_{t-2})\} \\ &= f(Y_{t-1})f(d_{t-1}|\Omega_{t-2}). \end{aligned} \quad (21)$$

Inserting both [20] and [21] into our original equation for the filtering step, i.e. [17], we obtain

$$\begin{aligned}
\Pr(c_{t-1} = 1|\Omega_{t-1}) &= \\
&= \frac{f(Y_{t-1}, d_{t-1}|c_{t-1} = 1) \Pr(c_{t-1} = 1|\Omega_{t-2})}{f(Y_{t-1}, d_{t-1}|\Omega_{t-2})} \\
&= \frac{f(Y_{t-1})f(d_{t-1}|c_{t-1} = 1) \Pr(c_{t-1} = 1|\Omega_{t-2})}{f(Y_{t-1})f(d_{t-1}|\Omega_{t-2})} \tag{22} \\
&= \frac{f(d_{t-1}|c_{t-1} = 1) \Pr(c_{t-1} = 1|\Omega_{t-2})}{f(d_{t-1}|\Omega_{t-2})}.
\end{aligned}$$

Neither the calculation of  $\Pr(c_{t-1}|\Omega_{t-1})$ , nor the calculation of  $\Pr(c_t|\Omega_{t-1})$  involves recovery rates directly.

In the converse case where default rates are not a function of the credit cycle, but recovery rates are, an equivalent argument will show that the filtering algorithm does not use the default rates.

Intuitively, if either recovery rates or default rates are independent of the cycle, they do not contain information about it - and they are therefore not useful in deducing the probability of being in either state of the credit cycle.

## C The Diebold-Gunther-Tay transform

Diebold et al. (1998) propose a way to evaluating density forecasts. The basic idea is that since under the null hypothesis the forecasts are equal to the true densities (conditioned on past information), applying the cumulative distribution function (the probability integral transform) to the series of observations should yield a series of iid uniform-[0,1] variables. Whether the transformed variables are iid uniform can be checked in various ways. Diebold et al. (1998) suggest plotting histograms and autocorrelation functions to visualize the quality of the density forecasts.

In order to apply the DGT transform to our predicted recovery rate densities, we create a vector  $y^\dagger$  in which we stack all recovery rate observations. For each element in the vector, we can now create a density forecast from our estimated model which conditions on all previous elements in this vector, and uses the filtered state probabilities. Due to our independence assumptions, this is mostly straightforward (most previous elements do not enter the conditional distributions).

Applying the cumulative distribution function associated with these density forecasts to the vector  $y^\dagger$  yields a vector of transformed variables. Under the null hypothesis that the density forecasts are correct, the elements of the vector of transformed variables should be an iid uniform series. Serial correlation of the series would indicate that we have not correctly conditioned on the relevant information. A departure from uniformity would indicate that the marginal distributions are inappropriate.



## D Supplementary tables

**Table V**  
**Recovery Rates Statistics by Year**

This table reports some annual statistics of the sample used in the paper. First column figures are default frequencies extracted from Standard & Poor's (2006). The other three columns are the number of observations and the mean and standard deviation for recovery rates.

Year	Default frequency	Number of observations	Mean Recovery	Standard Deviation
1981	0.14%	1	12.00	-
1982	1.18%	12	39.51	14.90
1983	0.75%	5	48.93	23.53
1984	0.90%	11	48.81	17.38
1985	1.10%	16	45.41	21.87
1986	1.71%	24	36.09	18.82
1987	0.94%	20	53.36	26.94
1988	1.42%	30	36.57	17.97
1989	1.67%	41	43.46	28.78
1990	2.71%	76	25.24	22.28
1991	3.26%	95	40.05	26.09
1992	1.37%	35	54.45	23.38
1993	0.55%	21	37.54	20.11
1994	0.61%	14	45.54	20.46
1995	1.01%	25	42.90	25.25
1996	0.49%	19	41.90	24.68
1997	0.62%	25	53.46	25.53
1998	1.31%	34	41.10	24.56
1999	2.15%	102	28.99	20.40
2000	2.36%	120	27.51	23.36
2001	3.78%	157	23.34	17.87
2002	3.60%	112	30.03	17.18
2003	1.92%	57	37.33	23.98
2004	0.73%	39	47.81	24.10
2005	0.55%	33	58.63	23.46

**Table VI**  
**Recovery Rates by Seniority and Industry**

Panel A: Number of observations and the mean and standard deviation of recovery rates in our sample classified by seniority, for the whole sample (all default events), for default events for which we only observe recovery on a single instrument (with only one seniority), and for default events for which we observe recoveries on at least two different seniorities.

Panel B: Number of observations and the mean and standard deviation of recovery rates in our sample classified by industry.

PANEL A: RECOVERY RATES BY SENIORITY			
Seniority	Number of observations	Mean Recovery	Standard Deviation
All default events			
Senior Secured	203	42.08	25.48
Senior Unsecured	366	36.88	23.29
Senior Subordinated	326	32.90	23.77
Subordinated	154	34.51	23.05
Discount	75	21.29	18.48
Default events with single recovery			
Senior Secured	145	39.29	23.35
Senior Unsecured	239	36.45	22.45
Senior Subordinated	209	34.52	23.30
Subordinated	87	37.86	20.22
Discount	29	21.72	19.67
Default events with multiple recoveries			
Senior Secured	58	49.04	29.25
Senior Unsecured	127	37.68	24.87
Senior Subordinated	117	30.00	24.42
Subordinated	67	30.16	25.79
Discount	46	21.03	17.91
PANEL B: RECOVERY RATES BY INDUSTRY			
Industry			
Building	10	33.56	36.24
Consumer	149	35.66	22.21
Energy	47	36.47	16.66
Financial	95	35.60	25.54
Leisure	69	41.43	29.40
Manufacturing	395	35.08	23.83
Mining	14	35.52	17.50
Services	65	34.16	28.09
Telecom	169	29.43	20.90
Transportation	66	38.07	23.79
Utility	23	51.34	27.97
Others	22	37.94	19.30

**Table VII**  
**Parameter estimates**

Parameter estimates and measures of fit for various models that differ in the combination of variables that influence the hazard rate  $\lambda$  and the two parameters of the beta distribution  $\alpha$  and  $\beta$ . Note that with the specification of Shumway (2001), a positive coefficient on an explanatory variable for  $\lambda$  implies that  $\lambda$  falls when the variable rises. sen2, sen3, sen4, sen5 are seniority dummies for Senior Unsecured, Senior Subordinated, Subordinated and Discount respectively, multsen is a dummy that is one for observations corresponding to default events for which we observe multiple recoveries, for the most senior recovery, and multjun is a dummy that is one for observations corresponding to the same default events, but for the junior recoveries. meanrec is the variable  $\bar{m}$  described in Appendix A, cycle is the unobserved credit cycle, and macro is log GDP growth (versions with other macroeconomic variables are omitted). indB and indC are dummies corresponding to industry groups B and C. Coefficients individually significant at 5% are in bold.

	M1		M2		M3		M4		M5	
DEFAULT RATES	$\lambda$		$\lambda$		$\lambda$		$\lambda$		$\lambda$	
constant	<b>4.16</b>		<b>3.56</b>		<b>3.69</b>		<b>3.37</b>		<b>3.56</b>	
cycle			<b>1.17</b>				<b>1.09</b>		<b>1.17</b>	
macro					<b>15.54</b>		10.44			
RECOVERY RATES	$\alpha$	$\beta$	$\alpha$	$\beta$	$\alpha$	$\beta$	$\alpha$	$\beta$	$\alpha$	$\beta$
constant	<b>0.44</b>	<b>1.03</b>	<b>0.40</b>	<b>1.31</b>	<b>0.36</b>	<b>1.14</b>	<b>0.40</b>	<b>1.32</b>	<b>0.28</b>	<b>1.14</b>
sen2	-0.00	<b>0.11</b>	-0.01	0.03	<b>-0.05</b>	0.04	-0.05	-0.05	-0.03	-0.04
sen3	<b>-0.22</b>	-0.04	<b>-0.16</b>	<b>-0.16</b>	<b>-0.33</b>	<b>-0.15</b>	<b>-0.19</b>	<b>-0.21</b>	<b>-0.26</b>	<b>-0.29</b>
sen4	<b>0.27</b>	<b>0.31</b>	-0.01	0.00	<b>0.15</b>	<b>0.22</b>	0.11	0.01	0.04	0.10
sen5	<b>-0.28</b>	<b>0.51</b>	<b>-0.22</b>	<b>0.45</b>	<b>-0.27</b>	<b>0.53</b>	<b>-0.21</b>	<b>0.43</b>	<b>-0.26</b>	<b>0.37</b>
multsen	<b>-0.28</b>	<b>-0.60</b>	<b>-0.19</b>	<b>-0.35</b>	<b>-0.48</b>	<b>-0.74</b>	<b>-0.38</b>	<b>-0.62</b>	<b>-0.19</b>	<b>-0.32</b>
multsen×sen2	0.08	<b>0.35</b>	0.07	0.17	<b>0.23</b>	<b>0.47</b>	0.08	0.19	0.11	0.17
multsen×sen3	0.14	<b>0.51</b>	0.05	<b>0.42</b>	<b>0.30</b>	<b>0.60</b>	-0.02	<b>0.40</b>	0.00	0.32
multsen×sen4*	-	-	-	-	-	-	-	-	-	-
multjun	<b>-0.27</b>	<b>-0.30</b>	<b>-0.34</b>	<b>-0.59</b>	<b>-0.41</b>	<b>-0.56</b>	<b>-0.49</b>	<b>-0.90</b>	<b>-0.31</b>	<b>-0.59</b>
multjun×sen3	0.06	<b>0.50</b>	-0.05	<b>0.44</b>	<b>0.19</b>	<b>0.77</b>	0.10	<b>0.82</b>	-0.07	<b>0.43</b>
multjun×sen4	-0.26	<b>0.30</b>	-0.13	<b>0.68</b>	-0.08	<b>0.56</b>	-0.16	0.69	-0.22	0.46
multjun×sen5	-0.03	0.18	0.03	0.38	0.12	0.36	0.00	0.52	0.01	0.38
meanrec	0.41	<b>-2.92</b>	-0.04	<b>-2.42</b>	-0.78	<b>-4.18</b>	0.11	<b>-2.86</b>	0.34	-2.19
cycle			<b>0.60</b>	<b>-0.28</b>			<b>0.49</b>	<b>-0.41</b>	<b>0.51</b>	<b>-0.39</b>
cycle×sen2			-0.16	0.07			-0.11	0.14	-0.15	0.09
cycle×sen3			<b>-0.51</b>	-0.03			<b>-0.50</b>	-0.01	<b>-0.34</b>	0.17
cycle×sen4			0.08	<b>0.54</b>			0.01	0.52	-0.07	0.31
cycle×sen5			-0.44	-0.28			-0.57	-0.30	-0.39	-0.19
cycle×multsen			<b>0.47</b>	-0.08			0.54	0.05	<b>0.48</b>	-0.11
cycle×multsen×sen2			-0.64	-0.07			-0.07	0.46	-0.54	0.07
cycle×multsen×sen3			-0.49	-0.13			0.09	0.29	-0.36	-0.02
cycle×multjun			-0.29	0.63			-0.06	1.12	-0.11	<b>0.81</b>
cycle×multjun×sen3			<b>1.02</b>	0.77			0.55	-0.03	0.76	0.58
cycle×multjun×sen4			0.40	-0.57			<b>1.38</b>	0.41	0.52	-0.25
cycle×multjun×sen5			0.63	-0.07			0.64	-0.38	0.50	-0.23
cycle×meanrec			1.47	-1.36			1.01	-1.17	1.32	-1.21
macro					5.61	-1.05	2.75	2.61		
macro×multsen					3.54	1.83	2.92	2.80		
macro×multjun					2.58	2.96	2.66	2.98		
indB									<b>0.27</b>	<b>0.37</b>
indC									<b>0.09</b>	<b>0.25</b>
TRANSITION PROB.										
p			0.8787				0.8146		0.8808	
q			0.7109				0.7271		0.7191	
MEASURES OF FIT										
Log Likelihood	84.023		295.87		135.58		307.88		301.38	
AIC	-114.05		-477.74		-203.16		-487.76		-480.76	
BIC	0.019		-0.17		-0.0288		-0.1479		-0.1550	

\* We have no observations for the case where multsen = 1 and sen4 = 1, and can therefore not estimate the coefficient on this interaction effect.


**Table VIII**  
**Parameter estimates (2)**

Parameter estimates and measures of fit for various models that differ in the combination of variables that influence the hazard rate  $\lambda$  and the two parameters of the beta distribution  $\alpha$  and  $\beta$ . Note that with the specification of Shumway (2001), a positive coefficient on an explanatory variable for  $\lambda$  implies that  $\lambda$  falls when the variable rises. sen2, sen3, sen4, sen5 are seniority dummies for Senior Unsecured, Senior Subordinated, Subordinated and Discount respectively, multsen is a dummy that is one for observations corresponding to default events for which we observe multiple recoveries, for the most senior recovery, and multjun is a dummy that is one for observations corresponding to the same default events, but for the junior recoveries. We have no observations for the case where multsen = 1 and sen4 = 1, and can therefore not estimate this interaction effect. meanrec is the variable  $\bar{m}$  described in Appendix A, cycle is the unobserved credit cycle, and macro is log GDP growth (versions with other macroeconomic variables are omitted). indB and indC are dummies corresponding to industry groups B and C.

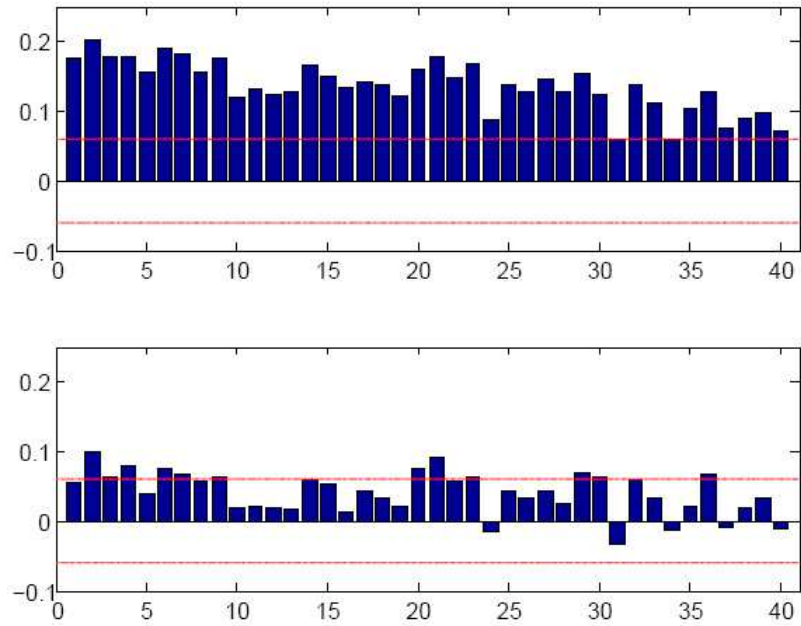
	M2a		M2b	
DEFAULT RATES	$\lambda$		$\lambda$	
constant	<b>3.59</b>		<b>4.16</b>	
cycle	<b>1.18</b>			
RECOVERY RATES	$\alpha$	$\beta$	$\alpha$	$\beta$
constant	<b>0.44</b>	<b>1.04</b>	<b>0.51</b>	<b>1.46</b>
sen2	-0.01	<b>0.09</b>	0.00	0.12
sen3	<b>-0.22</b>	<b>-0.05</b>	<b>-0.24</b>	-0.19
sen4	<b>0.27</b>	<b>0.31</b>	<b>-0.56</b>	<b>-0.71</b>
sen5	<b>-0.30</b>	<b>0.47</b>	-0.25	0.29
multsen	<b>-0.25</b>	<b>-0.61</b>	<b>-0.35</b>	<b>-0.45</b>
multsen×sen2	0.08	<b>0.38</b>	0.2	0.41
multsen×sen3	0.08	<b>0.52</b>	0.12	0.36
multsen×sen4*	-	-	-	-
multjun	<b>-0.27</b>	<b>-0.30</b>	-0.43	<b>-0.84</b>
multjun×sen3	0.03	<b>0.48</b>	0.06	0.87
multjun×sen4	<b>-0.23</b>	<b>0.32</b>	0.92	<b>2.69</b>
multjun×sen5	-0.07	0.13	0.05	0.68
meanrec	0.58	<b>-2.82</b>	-0.54	-2.16
cycle			<b>0.29</b>	<b>-0.54</b>
cycle×sen2			-0.05	<b>0.02</b>
cycle×sen3			<b>-0.25</b>	<b>0.07</b>
cycle×sen4			<b>0.69</b>	<b>1.32</b>
cycle×sen5			<b>-0.52</b>	0.07
cycle×multsen			0.56	<b>0.07</b>
cycle×multsen×sen2			<b>-0.56</b>	-0.29
cycle×multsen×sen3			-0.29	<b>0.11</b>
cycle×multjun			<b>-0.19</b>	0.77
cycle×multjun×sen3			<b>0.60</b>	<b>-0.36</b>
cycle×multjun×sen4			-0.89	-2.86
cycle×multjun×sen5			<b>1.01</b>	<b>-0.35</b>
cycle×meanrec			1.77	-0.61
TRANSITION PROB.				
p	0.8721		0.9055	
q	0.7271		0.6106	
MEASURES OF FIT				
Log Likelihood	221.4		171.89	
AIC	-380.81		-231.78	
BIC	-0.2002		0.0441	

\* We have no observations for the case where multsen = 1 and sen4 = 1, and can therefore not estimate the coefficient on this interaction effect.

## E Supplementary Figures

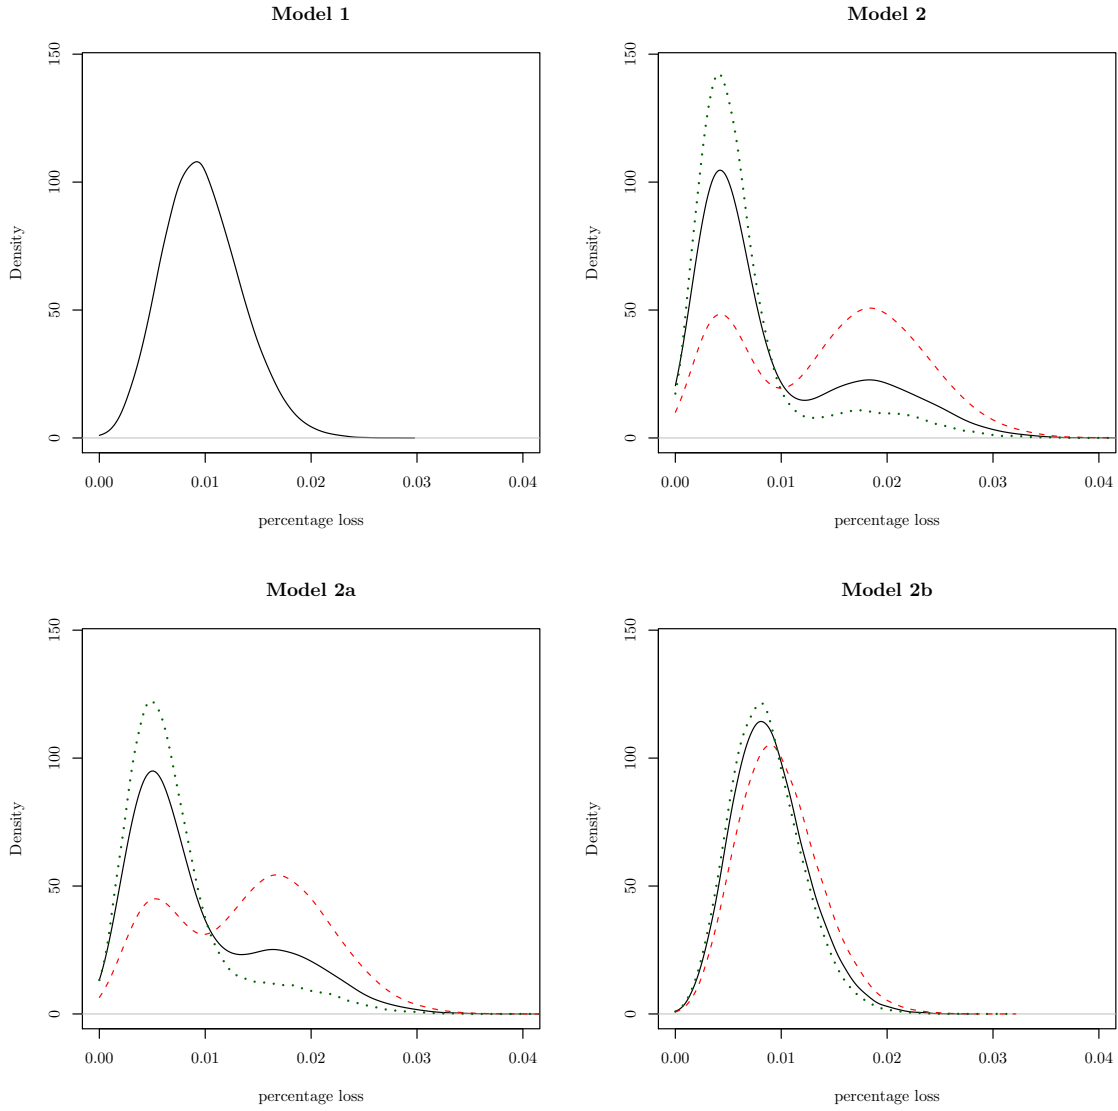


**Figure 4. Q-Q plots of the PIT series of the static and cycle models.** Dashed line is the static model (M1), solid line is a dynamic model (M2). The upper panel is a Q-Q plot for periods which are identified as upturn by the dynamic model, the lower panel is a Q-Q plot for periods which are identified as downturns by the dynamic model.



**Figure 5. Correlograms of the PIT series of the static and cycle model.**

Upper panel is the correlogram of the static model (M1), the lower panel is the correlogram of a dynamic model (M2). Horizontal lines are 5% two-sided confidence intervals for a single autocorrelation coefficient.



**Figure 6. Loss densities for Models 1, 2, 2a, and 2b**  
 Loss densities for Models 1, 2, 2a and 2b. The solid black line is the unconditional loss density (i.e. assuming the probability of being in an upturn is equal to the unconditional probability), dotted green line is the upturn loss density (probability of being in upturn of 1), the red dashed line is the downturn loss density (probability of being in upturn of 0).

## References

- Acharya, V. V., S. T. Bharath, and A. Srinivasan, 2007. Does Industry-wide Distress Affect Defaulted Firms? - Evidence from Creditor Recoveries. *Journal of Financial Economics* 85(3):787–821.
- 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* pages 57–64.
- Altman, E. I., A. Resti, and A. Sironi, 2001. Analyzing and Explaining Default Recovery Rates. Report submitted to ISDA, London.
- Bangia, A., F. X. Diebold, A. Kronimus, C. Schagen, and T. Schuermann, 2002. Ratings Migration and the Business Cycle, with Application to Credit Portfolio Stress Testing. *Journal of Banking & Finance* 26:445–474.
- Campbell, J. Y. and J. H. Cochrane, 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107(2):205–251.
- Chava, S., C. Stefanescu, and S. M. Turnbull, 2006. Modeling Expected Loss. Working Paper.
- Chen, L., P. Collin-Dufresne, and R. S. Goldstein, 2006. On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle. Working Paper.
- Cochrane, J. H., 2001. *Asset Pricing* (Princeton University Press, Princeton, N.J.).
- Diebold, F. X., T. A. Gunther, and A. S. Tay, 1998. Evaluating Density Forecasts with Applications to Financial Risk Management. *International Economic Review* 39(4):863–883.
- Duffie, D., A. Eckner, G. Horel, and L. Saita, 2006. Frailty Correlated Default. Working Paper.



- Gupton, G. M. and R. M. Stein, 2002. LossCalc<sup>TM</sup>: Model for Predicting Loss Given Default (LGD). Moody's KMV, New York.
- Hamilton, J. D., 1989. Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2):357–384.
- Hansen, B. E., 1992. The Likelihood Ratio Test under Nonstandard Conditions: Testing the Markov Switching Model of GNP. *Journal of Applied Econometrics* 7:S61–S82.
- , 1996. Erratum: The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP. *Journal of Applied Econometrics* 11(2):195–198.
- Nickell, P., W. Perraudin, and S. Varotto, 2000. Stability of Rating Transitions. *Journal of Banking & Finance* 24:203–227.
- Shleifer, A. and R. W. Vishny, 1992. Liquidation Values and Debt Capacity: A Market Equilibrium Approach. *Journal of Finance* 47(4):1343–1366.
- Shumway, T., 2001. Forecasting Bankruptcy More Accurately: A Simple Hazard Model. *Journal of Business* 74(1):101–124.
- Standard & Poor's, 2006. Quarterly Default Update and Rating Transitions. New York.
- Varma, P. and R. Cantor, 2005. Determinants of Recovery Rates on Defaulted Bonds and Loans for North American Corporate Issuers: 1983-2003. *Journal of Fixed Income* 14(4):29–44.