Recovery Rates, Default Probabilities and the Credit Cycle

Max Bruche and Carlos González-Aguado*

CEMFI†

November 17, 2006

Abstract

Recovery rates are negatively related to default probabilities (Altman et al., 2005). This paper proposes and estimates a model in which this dependence is the result of an unobserved credit cycle: When times are bad, the default probability is high and recovery rates are low; when times are good, the default probability is low and recovery rates are high. The proposed dynamic model is shown to produce a better fit to the data than a standard static approach. It indicates that ignoring the dynamic nature of credit risk could lead to a severe underestimation of credit risk (e.g. by a factor of up to 1.7 in terms of the 95% VaR). Also, the model indicates that the credit cycle is related to but distinct from the business cycle as e.g. determined by the NBER, which might explain why previous studies have found the power of macroeconomic variables in explaining default probabilities and recoveries to be low.

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 Stéphane Bonhomme, Abel Elizalde, Rafael Repullo, Enrique Sentana and Javier Suárez for helpful comments.

†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 suggested that default probabilities and / or rating transitions might be related to the business cycle. Bangia et al. (2002), for example, estimate separate rating transition matrices for NBER recessions and expansions,[1] and Nickell et al. (2000) estimate different rating transition matrices for periods of high, medium and low GDP growth and find that rating transition matrices are significantly different across these phases. In both cases, a downturn in the business cycle seems to go hand in hand with an increase in default probabilities.

It has also been noted, however, that macroeconomic variables seem to be unable to fully explain default probabilities. Das et al. (2007) produce evidence to suggest that apart from the macroeconomic and other covariates that they use to model default intensities, there are unobserved covariates or factors that drive default probabilities. This is corroborated by for example Koopman et al. (2006), who find that a latent or unobserved state variable is important in explaining default probabilities or rating transitions, even in the presence of macro variables.

What this suggests is that a credit cycle that drives default probabilities exists, but that it cannot be identified by using only macroeconomic data.

Another strand of literature examines the relationship of default probabilities or default rates, and finds that default rates and recovery rates are negatively correlated (see e.g. Altman et al., 2005).

The contribution of this paper is to combine these strands of literature and examine to what extent it is possible to identify a credit cycle using only the data on recovery rates and default rates, i.e by letting the credit data speak. Intuitively, we formalize the idea that in "good times", default rates are low and recoveries are high, and in "bad times", default rates are high and recoveries are low. So far, no attempt has been made to identify the credit cycle on the basis of both recoveries and default events[2].

Credit risk can conceptually be thought of as consisting of both default and recovery risk. If default rates and recovery rates are negatively correlated through an unobserved

[1] Bangia et al. (2002) also estimate the NBER business cycle transition matrix, and use this to simulate the unconditional rating transition matrix by this dynamic model, which they compare to the static rating transition matrix to find that the static rating transition matrix understates default risk.

[2] Frye (2000) proposed a model based on Vasicek (1987) that is closest in spirit to what we propose here. His model, however, is phrased in a static context, and cannot therefore be used to identify a dynamic credit cycle in a consistent manner.

credit cycle, this could change the view of the credit risk of portfolios quite dramatically, which has implications for risk management and possibly also for pricing.

In many commercial credit risk models, the default rate and recovery rates are assumed to be independent,[3] or constant.[4] If realizations of recoveries are low exactly at times when many firms default, assuming that recoveries are independent of default rates or constant will mean an underestimation of credit risk.

In many pricing models, the recovery rate is also assumed to be a constant, or independent of other factors in the model, either in the model itself or in its empirical implementation. This will lead to models that understate risk, and hence produce prices that are too high or spreads that are too low. It has been argued by Huang and Huang (2003), for instance, that many structural models of credit produce spreads that are too low once calibrated to match average default probabilities and average recovery rates. They suggest that this might be due to assuming a risk premium that is constant, when in reality it is time-varying. The approach presented here does not understate risk and hence is likely to produce lower prices and higher yields, even without reference to time-varying risk premia.

Our credit cycle will be an unobserved two-state Markov chain. Making some distributional assumptions about default rates and recoveries will then allow us to estimate a model akin to the one presented by Hamilton (1989). We show that the resulting model can be interpreted as a simple version of a reduced form (or intensity) model of credit (Jarrow and Turnbull, 1995).

Our estimated model allows comparing our credit downturns to recession periods as determined by the NBER. The beginning of our credit downturns typically precedes the start of a recession, and continue until after the end of a recession (see e.g. Figure 1). The average duration of a credit downturn is in the range of 4.9 years, as opposed to an estimated average duration of 4.1 to 4.7 *quarters* (Hamilton, 1989). Our estimated credit cycle indicator exhibits some, but not perfect correlation with various macroeconomic variables (e.g. a correlation of 36% with the S&P 500 return, or a correlation of 34% with GDP growth, see Table 14) This indicates that the credit cycle is related to, but distinct from the macroeconomic cycle. We argue that this explains why previous studies have found that macroeconomic variables explain only a small proportion in the

[3]E.g. J.P. Morgan's CreditMetrics^{TM} model

[4]E.g. CSFB's CreditRisk+^{TM} model

variation of recovery rates (Altman et al., 2005).

We show how the model can be used to calculate loss distributions for portfolios. Allowing for dependence between default probabilities and recovery rates via the state of the credit cycle can increase e.g. the 95% VaR of a given portfolio by a factor of up to 1.7, and even in credit upturns, the VaR can increase by a factor of up to 1.3. Altman et al. (2005) obtain a very similar difference in VaRs on the basis of a hypothetical simulation exercise; we can confirm on the basis of our estimated model that their numbers are very, very plausible. Also, Hu and Perraudin (2002) estimate the tails of actual loss distributions on the basis of extreme value theory, and compare that with the tail estimate based on losses calculated by historical simulation assuming independence between recoveries and default probabilities, and produce a qualitatively similar result, i.e. they find that for a given confidence level, quantiles of actual losses are much large than quantiles of hypothetical losses assuming independence. They do not identify the credit cycle, however.

The difference in VaRs is economically significant, and should be taken into account e.g. by banks attempting to implement the Advanced Internal Ratings-Based Approach of Basel II, under which they can calculate their own default probabilities and estimates of loss given default.

Several hypotheses explaining why recoveries might be low in times when many firms default have been suggested. Altman et al. (2005) argue that the markets for defaulted securities have limited capacity (i.e. demand for these securities is not perfectly elastic as standard asset pricing theory would suggest), and when many of these appear at the same time, this depresses the price of defaulted securities. If recovery is measured as the price of a defaulted security as a fraction of par (as is standard practice), then of course this would depress recoveries in times of large default rates. Regressing recovery rates on e.g. the aggregate default rate as an indicator of the aggregate supply of defaulted bonds, they find a negative relationship. When adding macroeconomic variables such as GDP growth, they also find that these do not contribute much to explaining recovery rates.

Two competing hypotheses were examined by Acharya et al. (forthcoming), who argue that recoveries might be low in times when the industry of the defaulting firm is in distress because this implies both a lower economic worth of the defaulting firm's assets, as well as a lower resale value of the firms assets because of the firms of the same industry, i.e. those firms that could put the assets to the most productive use, are financially constrained. This second hypothesis was first developed by Shleifer and Vishny (1992). Acharya et al. (forthcoming) produce evidence to suggest that the second "fire-sales" hypothesis is a likely explanation, and note that industry distress dummies explain a large part of the apparent effect of the aggregate supply of defaulted bonds. The data seems to indicate that industries are in distress at the same time; this leaves the question as to why they would be in distress simultaneously.

Kiyotaki and Moore (1997) analyze the effects of productivity shocks on the business cycle in the presence of interactions between borrowing constraints and asset prices. Similarly to Shleifer and Vishny (1992), they argue that given a negative productivity shock, the net worth of firms diminishes and the demand for assets (in the presence of credit constrains) decreases. This lowers the price of assets and thus the net worth of these constrained firms. But this might reduce the future net worth of these firms. So in the presence of credit limits, effects of a productivity shock can be transmitted into the future. This creates a business and credit cycle out of productivity shocks and credit constraints. Although they do not explicitly discuss defaults or recoveries, it is possible to imagine a version of the model in which more firms default and less is recovered in situation when prices of assets are depressed.

Suárez and Sussman (2007) present a model that is similar in spirit in which the presence of financial constraints together with the effects of liquidations on asset prices can produce endogenous fluctuations (i.e. without productivity shocks).

In conclusion, there is empirical evidence as well as theoretical considerations that make it plausible that recovery rates and default rates would be related to a credit cycle, which in turn could be related to the business cycle.

The rest of this paper is structured as follows: The model is presented in section 2. Relationships between the model and standard reduced-form models of credit are explored in section 3. In section 4 we describe the data set used. Section 5 discusses the estimation and various tests, and section 6 explores the implications for credit risk management and pricing. Finally, section 7 concludes.

2 The model

The basic idea of the model is to let default probabilities and recovery rates be related by letting them depend only on the state of the credit cycle. Default correlation between firms, dependence between recoveries of different firms, as well as the relationship between recovery rates and default rates are driven entirely by the credit cycle. While this might seem restrictive, this is of course a lot less restrictive than assuming that variables are independent, and relationships are static.

The credit cycle is described by a two-state Markov chain. Let s_t be the unobservable state of the cycle, with $s_t = 0$ corresponding to a credit downturn and $s_t = 1$ corresponding to a credit upturn. Also, let p be the probability of remaining in an upturn and q the probability of remaining in a downturn. Then, following Hamilton (1989), we can describe the dynamics of the state of the cycle as

$$s_t = (1-q) + (p+q-1)s_{t-1} + v_t \tag{1}$$

where v_t is a martingale difference sequence that can take a finite number of values. Given the specification in (1) and taking into account that s_t is a binary variable, we can easily obtain the probabilities of being in one state or another conditioned on the state in the previous period.

Each firm in the population has a default probability which only depends on the state of the cycle. We call the default probabilities in downturns r_0, and the default probability in upturns r_1.

We assume that conditional on the cycle, defaults of firms are independent, i.e. that any (unconditional) dependency is driven entirely by the cycle. The default indicator $\mathbf{1}_{it}$ of a firm i is zero if a firm has not defaulted in the period and is equal to 1 in the period in which the firm defaults. In terms of this variable, our assumption is that conditional on the state, $\mathbf{1}_{it} \perp\!\!\!\perp \mathbf{1}_{jt}$ for firms $i \neq j$. As a consequence of this assumption, the number of defaults d_t in the population of size N_t will be binomially distributed with parameter r_s, conditional on knowing the state.

The recovery rate for a default 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 α and β). 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 s, the seniority class of the debt on which default occurs c and the industry of the issuer k, such that the recovery rate of firm i at time t on debt of class c is drawn from a density

$$f(y_{tic}) = \frac{1}{B(\alpha_{sck}, \beta_{sck})} y_{tic}^{\alpha_{sck}-1}(1 - y_{tic})^{\beta_{sck}-1}. \tag{2}$$

We assume that conditional on the state, $y_{tic_1} \perp\!\!\!\perp y_{tjc_2}$ for firms $i \neq j$ and $\forall c_1, c_2$, i.e. once we condition on the unobserved state, recovery rates of one defaulting firm are assumed independent of recovery rates of all other defaulting firms. This means that any dependence in the unconditional recoveries across firms will be driven entirely by the unobserved state.

Also, we assume that recoveries are independent of the default events, i.e. $\mathbf{1}_{it} \perp\!\!\!\perp y_{tjc}$ for firms $i \neq j$, such that (unconditional) dependence between recoveries and the aggregate number of defaults is entirely driven by the state of the credit cycle. This assumption is of course crucial, and its validity will be examined below (see section 5).

Often, a single firm defaults on more than one class of debt. Recoveries associated with one firm but across several classes of debt are unlikely to be independent. It is likely that a high firm-level recovery implies higher instrument-level recoveries, and that therefore recoveries on instruments issued by the same firm but of different seniority classes exhibit positive dependence. Also, it would be natural to expect that higher seniority classes observe a higher recovery. As it turns out, in the data, it is relatively frequently the case that higher seniority classes recover less than the lower seniority classes, however.[5]

In the absence of a model that relates firm-level recoveries to instrument-level recoveries (see e.g. Carey and Gordy, 2004, for a discussions of some of the issues), we propose the relatively simple assumption that the dependence structure of recovery rates across different seniority classes for the same firm is given by a Gaussian copula, such that we can specify a correlation matrix $\mathbf{\Gamma}$ of recoveries across seniority classes. This means that we do not impose a dependence structure that implies that higher seniorities will always recover more, since this is not what we see in the data. When a

[5]This is likely to reflect the fact that the specific terms of debt contracts of a given seniority class vary across firms, and that the labels of "Senior Unsecured" etc. do not always carry the same meaning across firms. On average, higher seniority classes recover more, however, so these labels do contain some information on potential recovery rates. It is also possible that this reflects aggregation issues and the structure of our data, see section 4 for details.

firm defaults on only one class of debt, we interpret [2] as the density of the recovery rate, and in the case where a firm defaults on more than one class of debt, we interpret [2] as the *marginal* density of recoveries.

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 Interpretations of the model

As a consequence of having probabilities of being in an upturn or downturn of the credit cycle, the particular density describing possible recoveries at any particular point in time will be a mixture of the two different beta densities for the two different states, with the mixing probabilities being given by the probabilities of being in either state. This gives considerably more flexibility in matching observed unconditional distributions of recovery rates, and is in some sense similar to the nonparametric kernel density approach of Renault and Scaillet (2004).[6]

We can also view the model as a discrete version of a reduced-form or intensity model. According to the model, conditional on the state of the credit cycle, the firms default independently, either with probability r_0 (in downturns), or with probability r_1 in upturns. This implies that the number of defaults in a given period, d_t, is binomially distributed. Taking limits of the binomial distribution as the size of the population $N_t \to \infty$ and the success probability $r_s \to 0$, such that $N_t r_s = \lambda_{ts}$, where λ_{ts} is constant (for a given state in a given time period), the distribution of d_t tends to a Poisson distribution. In fact, since our N_t is always above 1000, and our per-period default probabilities do not exceed 3%, the Poisson distribution is a very good approximation of our implicitly assumed binomial distribution. We can view the counting process describing the aggregate number of defaulted firms d_t as the sum of N_t individual counting processes that describe the default of individual firms. For this to work, the intensities of the individual Poisson processes need to sum up to λ_{ts}. This could be achieved by e.g. assuming the individual default intensity of firms to be r_s. These intensities depend on an unobserved state; such that we can view the default process of an individual firm as a Cox process, where the relevant conditioning information is

[6] It is similar in the sense that we are also using a mixture of beta distributions. The difference is that obviously, a kernel density estimator will mix over a much larger number of densities (equal to the number of observations), with mixing probabilities constant and equal across the kernel densities.

the (unobserved) state of the credit cycle, i.e. our two-state Markov chain. Recoveries are also assumed to depend on this unobserved Markov chain.

If the default counting process conditional on knowing the states $s = \{s_0, s_1, \ldots, s_T\}$ has a Poisson distribution, the default time τ of an individual firm will be exponentially distributed. The survival probability of an individual firm from time $t = 0$ until time T is given by

$$\Pr(\tau > T|s) = \exp\left\{\sum_{t=0}^{T} r_{s_t}\right\}, \tag{3}$$

and the unconditional probability is given by

$$\Pr(\tau > T) = E\left[\exp\left\{\sum_{t=0}^{T} r_{s_t}\right\}\right], \tag{4}$$

where the expectation is given over the probabilities of the states.

In fact, this is simply a discrete time version of the continuous-time models e.g. used by Duffie and Singleton (1999) or Duffee (1999), apart from the fact that here the unobserved process driving the default intensity is not a continuous-time Feller square-root process, but a discrete time two-state Markov chain. Also, the recovery assumptions are different; here, conditional on the state recovery is a beta-distributed fraction of par.

4 Data

The data 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 seniority of the bond, the date of default and the price of this bond per 100 dollars of face value one month after the default event.

Typically, the database contains the prices of many bonds for a given firm and seniority class on a given date, so we need to aggregate. We do this by taking weighted averages (weighted by issue size).[7]

[7] This is something commonly used in the literature. See for example Varma and Cantor (2005).

We also aggregate data across time into periods corresponding to calendar years. 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 [8].

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. Equivalently, we could view our beta distribution as being defined over a 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. This variable is necessary in order to calculate the probability of observing d_t defaults in [10]. 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 frequency, we can obtain a number for the total population of firms under the assumption that both data sets track the same set of firms. The available default frequency 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 these adjustments, the final data set contains 1,078 observations. Descriptive statistics for these observations are presented in Tables 1 to 3.

Table 1 reports the yearly statistics of our adjusted data set. The first column shows the annual default frequencies as reported by Standard & Poor's, while the other columns refer to the main data set, containing the number of observations per year and their means and standard deviation. It can already be seen that typical recession years (as published by the NBER: 1990-91 and 2001) show higher default frequencies and lower recoveries than years of economic expansion. This will of course play an important role below.

[8] It is possible that the fact that we sometimes observe higher recoveries for bonds of lower seniorities is related to this aggregation. In situations where default occurs on a junior security before they occur on the senior security, it could between the earlier date and the later date bond prices fall as more negative information about the firm becomes available, and that the recovery on junior securities could therefore be higher.

Tables 2 and 3 report the number of 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 Table 2, it is shown that mean recovery increases with increasing seniority, while standard deviation remains more or less constant across seniorities. In Table 3, we see that the mean recovery is highest for utilities and lowest for telecoms.

5 Estimation and tests

The work of Altman and Kishore (1996) among others suggests that other important determinants of recovery rates are seniority and industry. Ultimately, we are therefore interested in constructing a model that takes into account the effects of the credit cycle, but also industry and seniority. Initially, however, we examine how well the model does using only the assumption about state dependence, ignoring the effects of industry and seniority (i.e. we look at a version of the model where marginal recovery rate distributions do not vary according to industry and seniority). Also, to sidestep the issue of dependence of recoveries on issues of the same firm, but with different seniorities, we initially drop all default events for which we observe recoveries on more than one seniority (the remaining observations are roughly representative of the whole sample). We estimate a *basic static model*, in which recovery rate distributions and default probabilities do not depend on the state, and contrast this with a *basic dynamic model* in which both recovery rate distributions and default probabilities are state dependent.

We then proceed towards a model that takes into account industry and seniority. Since we do not have a sufficient number of default events for every possible combination of industry, seniority and state of the credit cycle, however, we will need to aggregate industries into broad groups.

In order to check that the aggregation is reasonable, we first allow marginal distributions to vary across industries only, and check which industries can be grouped. Next, we allow marginal distributions to vary across seniorities only. We can then look at the dependence of recoveries on issues of the same firm but with different seniorities. For this we need to add the observations dropped at the earlier stage back in, and look at the full data set.

Finally, we combine information on industry and seniority to construct the *dynamic industry/ seniority model*. This will be contrasted with a *static industry/ seniority model*.

At the various stages, different tests will be performed.

5.1 Basic model

5.1.1 The basic static model

In the *static model* recovery rate distributions and default probabilities do not depend on the state. We also ignore industry and seniority. We estimate on the sub-sample of observations for which we do not observe recoveries across different seniorities for the same firm.

Estimates of this model are provided in Table 4. With the estimated distribution parameters we obtain an implied mean recovery of 37%.[9] The default probability in this static case is found by taking the average of all the default frequencies in our data set and it is equal to 1.47%.

5.1.2 The basic dynamic model

In the *basic dynamic model* recovery rate distributions and default probabilities depend on the state, but we ignore industry and seniority. We estimate on the sub-sample of observations for which we do not observe recoveries across different seniorities for the same firm.

The estimated parameters are reported in Table 4. Average recoveries are much lower in downturns (31% versus 47%), and default probabilities are higher (2.69% versus 0.86%).

We test whether the estimated parameters values for downturns are significantly different from the estimated parameter values for upturns via likelihood ratio tests: We first test whether the recovery rate distribution parameters parameters are different, and then test whether the default probabilities are different. The p-value are less than 0.01% for both tests, so that the null hypothesis of no difference across states is rejected both for recovery rate distributions and default probabilities. The credit cycle matters.

[9]The mean of a beta distribution with parameters α and β is $\frac{\alpha}{\alpha+\beta}$. Note that we also have to divide the resulting number by 0.9 in order to undo the scaling adjustment.

We check the recovery rate densities implied by the models using a method proposed by Diebold et al. (1998), described appendix B. 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. Departures from uniformity are easily visible when plotting histograms, and departures from i.i.d. are visible when plotting autocorrelation functions of the PIT series. We also test for departures from uniformity using the Kolmogorov-Smirnov (K-S) test,[10] and departures from i.i.d-ness using the Ljung-Box (L-B) test.

The p-value of the K-S statistic for the basic static model is 2.32%, versus 6.38% for the basic dynamic model, indicating that the dynamic model does a better job at matching recovery rate distributions. The p-value of the L-B test is virtually zero for the basic static model, indicating that the static model clearly does not describe the dynamics of recovery rates adequately. For the basic dynamic model, the p-value of the L-B test is 20.56%, indicating that it is much better at describing the dynamics of recovery rates. This can also be seen by examining the histograms and correlograms of the transformed series (see figures 2 and 3).

We also check the goodness of fit of our estimated default rates. For this purpose, we regress observed default rates on a constant and estimated default rates of the dynamic model (for each of its variants). These estimated default rates are obtained by mixing, for each period, the default rates 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 both the static and dynamic models. The static model would be correct if the coefficient of the variables is not significantly different from zero, so that the only explanatory variable was the constant (the static default rate). The null hypothesis for the dynamic model is that the constant and the coefficient of the variable are equal to zero and one respectively. Table 12 reports the p-values for all these hypothesis. We see how the static model is clearly rejected, while we obtain very high p-values (and then accept) the dynamic model. [11]

[10] The distribution of this test is derived under the null that the distribution to which the empirical distribution is compared is known, and not estimated, which is not the case in our application. Some care therefore has to be taken in interpreting the p-values.

[11] 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.

We also plot the (smoothed) probabilities of being in the credit downturn over time and compare these to NBER recession dates in figure 1. As can be seen, the credit cycle had two major downturns, around the recession of the early 90s and around 2001, but in each case, the credit downturn started well before the recession and ended after it.[12] This difference between the credit cycle and the business cycle might explain the low explanatory variable of macroeconomic variables documented by Altman et al. (2005).

Raw correlations between the smoothed upturn indicator and various macroeconomic variables are reported in Table 14. We can see for instance that the upturn indicator is positively correlated with consumption growth (a correlation of 34%).

We also examine our assumption that recoveries and default probabilities are independent conditional on the state of the credit cycle (see equation 7). We regress annual mean recoveries on default rates for the whole sample. The coefficient is significantly different from zero at 5% (the correlation between the two variables is -0.43). We also run two regressions separately in the two subsamples given by the periods for which the smoothed probabilities of being in a credit upturn or downturn are unambiguous (i.e. either 1 or 0). The coefficients of the regressions in the two sub-samples are not significantly different from zero at 5% (the correlations are 0.47 in upturns and -0.12 in downturns). This seems to support our assumption.

5.2 Adding information on industry

We now let the parameters of the recovery rate density depend on industry, as well as on the unobserved credit cycle. We have 12 industries, which implies that we need to estimated 48 recovery rate distribution parameters (2 states, 12 industries, 2 parameters per beta distribution), 2 transition probabilities and 2 default probabilities. The estimated parameters of the recovery rate distributions and the implied mean recoveries are provided in Table 5.

As discussed, in order to construct a combined industry / seniority model, due to data limitations, we will have to aggregate. We aggregate industry into three groups:

1. Group A: Financials, Leisure, Transportation, Utilities
2. Group B: Consumer, Energy, Manufacturing, Others

[12] The "semi-downturn" of 1986 consists mostly of defaulting oil companies. It is probable that this is related to the drop in oil prices at the time.

3. Group C: Building, Mining, Services, Telecoms

For the given sample, these groupings roughly correspond to industries with high, medium and low mean recovery rates respectively. We test whether the parameters of the recovery rate distribution are significantly different within these groups (i.e. whether it the groupings are reasonable). The p-value of the likelihood ratio test is 10%, which indicates that the null of the same parameters within the groups cannot be rejected at conventional levels of significance.

5.3 Adding information on seniority

We now estimate a model in which the parameters of the recovery rate densities depend on the state of the credit cycle and seniority (and ignore industry for the time being). We observe 4 different categories of seniority (Senior Secured, Senior Unsecured, Senior Subordinated, Subordinated). The estimated recovery rate parameters and the mean recoveries implied by these parameters are shown in Table 6. We observe that although senior secured bonds have a higher recovery in upturns on average (52%), they have very similar (low) recoveries to bonds of all other seniorities in downturns (29%). As has been pointed out by Frye (2000), this might have important consequences for risk management, as instruments that are though of as "safe" turn out to be safe only in good times. Acharya et al. (forthcoming) also find that in times of distress, senior securities recover significantly less, while recoveries on bank debt and junior securities are not significantly affected. They hypothesize that it is especially holders of senior debt that lose their bargaining power in times of distress, and hence recover less. Acharya et al. (forthcoming) also argue that that if the senior debt is collateralized, this effect is mitigated. Based on our estimation results, we cannot corroborate this result. For our data, senior debt is strongly affected by a credit downturn, regardless of whether or not it is secured.

5.3.1 Dependence across seniorities

So far we have dealt only with observations for which we do not observe several recoveries across different seniorities for the same firm, which we now add back in. It is unreasonable to assume that these observations of recoveries represent independent draws. As described above, letting our marginal recovery rate densities be beta densities as before, we model this dependence with a Gaussian copula.

In order to estimate this dependence, we add the observations of default events for which we observe recovery rates across instruments of different seniorities back in. For these observations, seniority seems to mean something quite different, as can be seen from Table 2. One explanation for this would be that recovery on a bond that is labelled as "senior unsecured" (for example) depends not so much on this label, but much more on whether or not there exists a cushion of junior debt.

Since categories of seniority seem to have a very different meaning when we observe recoveries across instruments of different seniorities, we allow the parameters of our marginal recovery rate densities for a given seniority to be distinct for the cases where either recovery only on a single seniority is observed, or whether recovery is observed on more than one seniority.

The estimated parameters of the marginal beta densities are reported in Table 7, and the estimated correlation matrix is presented in Table 8.

Implied mean recoveries of this model are presented in Table 7. We can see that junior debt recovers less and senior debt recovers more for default events for which we observe recoveries for more than one seniority. For example, in an upturn, Senior Unsecured debt would recover (on average) 46% and Subordinated debt would recover 34% if recoveries are observed on more than one seniority, whereas they would recover (on average) 42% and 37% respectively if recovery was only observed on a single seniority.

At 5% significance, only the correlations between Senior Unsecured, Senior Subordinated and Subordinated seniorities are significantly different from zero. These correlations are positive, indicating that a high recovery on e.g. Senior Unsecured debt for a defaulting firm indicates a likely high recovery on its Senior Subordinated Debt. The recovery on Senior Secured debt seems to be less strongly related to recoveries on different seniorities. This could reflect that the existence of a junior debt cushion matters less, and the quality of the collateral matters more for Senior Secured debt.

5.4 The industry / seniority dynamic model

Having decided on groupings of industries such that we now have a sufficient amount of data for each possible combination of seniority group, industry group, and state of the business cycle, we are now in a position to estimate a combined industry / seniority model. The recovery rate density parameters are in tables 9 and 11.

For comparison purposes, we also estimate a static industry / seniority dynamic model, which is identical with the dynamic industry / seniority model except for the fact that the static model ignores the credit cycle. We calculate the PIT series for both models, and plot histograms and correlograms in figures 2 and 3. We can see that the unconditional distributions seem to be closer to uniformity, suggesting that taking into account industry and seniority helps in describing distributions of recovery rates. Looking at the correlograms, it is also apparent, however, that the static version of the industry / seniority model is again unable to explain the dynamics of recovery rates. We conclude that allowing the parameters of the density of recoveries to depend on the credit cycle allows for a better match of the empirically observed recovery rates.[13]

6 Implications

Altman et al. (2005) suggest that a correlation between default rates and recovery rates can imply much higher VaR numbers than those implied by independence between these two variables. In a completely static model, they compare VaRs calculated in the case of independence, and in the case of perfect rank correlation between default rates and recovery rates.

We are in a position to calculate the loss distributions and statistics of loss distributions (such as the VaR) implied by our estimated model; i.e. taking into account the estimated degree of dependence between default probabilities and recovery rates, which in our case is driven entirely by the credit cycle.

We calculate (by simulation of 10,000 paths) the one-year loss distribution of a hypothetical portfolio of 500 bonds. For this calculation, we have to chose 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 (we have no information on the current state of the cycle).

[13] We also calculate a Kolmogorov-Smirnov test on the PIT series, and find a p-value of 9.91% for the static model, and a p-value of 14.37% for the dynamic model. Note that the caveat mentioned above applies.

Comparing the *basic static model* and the *basic dynamic model* (i.e. ignoring industry and seniority for the time being), we can see from Table 13 that the 95% VaR is a 2.39% loss on the portfolio assuming the world is dynamic as described in our model, versus a 1.58% loss on the portfolio if we had assumed that the world is static. This is a very sizeable difference in risk. It arises because the dynamic model allows for credit downturns, in which not only the default rate is very high, but also, the recovery rate is very low. This amplifies losses vis-a-vis the static case.

Even supposing that we are in an upturn today, the dynamic model still produces a 95% VaR of 1.96%, which is larger than the static VaR because even though we are in an upturn today, we might go into a credit downturn tomorrow, with higher default rates and lower recoveries.

The loss density implied by the dynamic model based on the unconditional probability of being in a downturn is compared to the loss density implied by the static model in figure 4. 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.

We also show this comparison for the case of assuming that we are in an upturn today in figure 4. It can be seen that even being in an upturn today, the possibility of going into a downturn tomorrow produces a bimodality in the loss distribution (albeit smaller) implied by the dynamic model, and a larger tail than that produced by the static model.

Looking at our industry / seniority model, we can see that the underestimation of the VaR of the static model seems to be most pronounced for Group C, our low recovery industries (including e.g. Telecoms), where the VaR based on the dynamic model can be up to 1.7 times as large.

It is possible that these results have implications for capital requirements. Altman et al. (2005) argue that credit risk is cyclical, which would lead to capital requirements calculated under the Basel II Internal Ratings-based Approach to be cyclical. Concretely, capital requirements should increase in credit downturns, as the estimated loss given default and the default probability rise simultaneously. They note that regulation should encourage banks to use "long-term average recovery rates" (and presumably default probabilities) in calculating capital requirements to dampen the procyclicality of capital requirements. In terms of the model presented here, one could argue that the appropriate "long-term" recovery rates and default probabilities to use would be the unconditional ones. This would be tantamount to discarding information on the state of the cycle when calculating capital requirements, but at least it would avoid calculating them on the of a wrong (static) model.

Of course, this assumes that the data on bond defaults and recoveries is indicative of bank loans (which might not necessarily be the case). The fact that Acharya et al. (forthcoming) report that recoveries on bank debt seem to be largely unaffected by their distress variable (while reporting other results that appear to be largely consistent with our results) casts doubt on this idea, although it is likely that default rates of bank loans and bonds are more similar. To answer the question conclusively, our model would have to be estimated on bank loan data.

Also, the increase in risk associated with negatively correlated default frequencies and recovery rates is likely to have an effect on pricing, and might (at least in part) explain the "corporate spread puzzle" (see e.g. Chen et al., 2006), i.e. the fact that structural models of corporate debt that assume that the recovery rate is a fixed parameter seem to be unable to match observed corporate bond yields with reasonable risk-aversion parameters.

7 Conclusions

This paper formalizes the idea that default probabilities and recovery rates are related through the credit cycle: it proposes a model that is estimated and found to fit the data reasonably well in various ways. We demonstrate that taking into account the dynamic nature of credit risk in this way implies dramatically higher risk, as evident in our VaR calculations, since default probabilities and recovery rates are negatively related. The estimated credit downturns seem to start much earlier and end later than the recessions as reported by the NBER, indicating that the credit cycle is to some extent distinct from the business cycle, which might explain why Altman et al. (2005) find that macroeconomic variables in general have low explanatory power for recovery rates.

References

Acharya, Viral V., Sreedhar T. Bharath, and Anand Srinivasan, "Does Industry-wide Distress Affect Defaulted Firms? - Evidence from Creditor Recoveries," *Journal of Financial Economics*, forthcoming.

Altman, Edward, Brooks Brady, Andrea Resti, and Andrea Sironi, "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications," *Journal of Business*, 2005, *78* (6), 2203–2227.

Altman, Edward I. and Vellore M. Kishore, "Almost everything you wanted to know about recoveries on defaulted bonds," *Financial Analysts Journal*, November/December 1996, pp. 57–64.

Bangia, Anil, Francis X. Diebold, André Kronimus, Christian Schagen, and Til Schuermann, "Ratings Migration and the Business Cycle, with Application to Credit Portfolio Stress Testing," *Journal of Banking & Finance*, 2002, *26*, 445–474.

Carey, Mark and Michael Gordy, "Measuring Systemic Risk in Recoveries on Defaulted Debt I: Firm-Level Ultimate LGDs," 2004. Working Paper.

Chen, Long, Pierre Collin-Dufresne, and Robert S. Goldstein, "On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle," 2006. Working Paper.

Das, Sanjiv R., Darrell Duffie, Nikunj Kapadia, and Leandro Saita, "Common Failings: How Corporate Deafults are Correlated," *Journal of Finance*, February 2007, *62* (1).

Diebold, Francis X., Todd A. Gunther, and Anthony S. Tay, "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, November 1998, *39* (4), 863–883.

Duffee, G., "Estimating the Price of Default Risk," *Review of Financial Studies*, 1999, *12*, 197–226.

Duffie, Darrell and Kenneth J. Singleton, "Modelling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 1999, *12* (4), 687–720.

Frye, Jon, “Depressing Recoveries,” 2000. Risk Magazine.

Gupton, Greg M. and Roger M. Stein, “LossCalc$^{\text{TM}}$: Model for Predicting Loss Given Default (LGD),” 2002. Moody’s KMV, New York.

Hamilton, James D., “Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 1989, *57* (2), 357–384.

Hu, Yen-Ting and William Perraudin, “Recovery Rates and Defaults,” 2002. Birbeck College Working Paper.

Huang, J. and M. Huang, “How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk?,” 2003. 14th Annual Conference on Financial Economics and Accounting (FEA); Texas Finance Festival.

Jarrow, R. and S. Turnbull, “Pricing derivatives on financial securities subject to credit risk,” *Journal of Finance*, 1995, *51* (1), 53–85.

Kiyotaki, Nobuhiro and John Moore, “Credit Cycles,” *Journal of Political Economy*, 1997, *105* (2), 211–248.

Koopman, Siem Jan, Roman Kräussl, André Lucas, and André Monteiro, “Credit Cylces and Macro Fundamentals,” 2006. Tinbergen Institute Discussion Paper.

Nickell, Pamela, William Perraudin, and Simone Varotto, “Stability of Rating Transitions,” *Journal of Banking & Finance*, 2000, *24*, 203–227.

Renault, Olivier and Olivier Scaillet, “On the way to recovery: A nonparametric bias free estimation of recovery rate densities,” *Journal of Banking & Finance*, 2004, *28*, 2915–2931.

Shleifer, Andrei and Robert W. Vishny, “Liquidation Values and Debt Capacity: A Market Equilibrium Approach,” *Journal of Finance*, September 1992, *47* (4), 1343–1366.

Standard & Poor’s, “Quarterly Default Update and Rating Transitions,” May 2006. New York.

Suárez, Javier and Oren Sussman, “”Financial distress, bankruptcy law and the business cycle,” *Annals of Finance*, January 2007, *3* (1).

Varma, Praveen and Richard Cantor, “Determinants of Recovery Rates on Defaulted Bonds and Loans for North American Corporate Issuers: 1983-2003,” *Journal of Fixed Income*, March 2005, *14* (4), 29–44.

Vasicek, Oldrich, “Probability of loss on loan portfolio,” 1987. KMV Working Paper.

A The likelihood function

Let d_t be the number of defaulted firms observed in t. Firms are be indexed by i. In terms of the previously defined default indicator, $d_t = \sum_i \mathbf{1}_{it}$. Different seniority classes are indexed by c (c ranges from 1 to C).

Let Y_t be a matrix of dimension $d_t \times C$ containing observations on the d_t defaulted firms, one row per firm, for all different seniority classes c. Denote the typical element as y_{tic}, and denote each row in this matrix as $y_{ti.}$.

Then our objective is to maximize the following likelihood function

$$L = \sum_{t=1}^{T} \log f(Y_t, d_t | \Omega_{t-1}), \tag{5}$$

where Ω_{t-1} is all information available at $t-1$, which includes observed defaults up to that point. Note that we have made explicit that there is information in both the matrix of recoveries Y_t and the number of its rows, d_t, about the state, i.e. identification of the state is obtained through both the number of defaults and recoveries.

Letting s denote state as above, we can write

$$\begin{aligned} f(Y_t, d_t|\Omega_{t-1}) = f(Y_t, d_t|s_t = 1, \Omega_{t-1}) \Pr(s_t = 1|\Omega_{t-1}) \\ + f(Y_t, d_t|s_t = 0, \Omega_{t-1}) \Pr(s_t = 0|\Omega_{t-1}). \end{aligned} \tag{6}$$

We assumed independence between recovery rates and defaults. This implies that the values of the recovery rates in the rows of Y_t, $y_{ti.}$, are independent of d_t conditional on the state. Once the state is known, the number of defaults does not contain information about the likely values of recoveries, and the values of recoveries do not contain information about the likely number of defaults. This allows us to write

$$\begin{aligned} f(Y_t, d_t|\Omega_{t-1}) = f(Y_t|s_t = 1, \Omega_{t-1}) \Pr(d_t|s_t = 1, \Omega_{t-1}) \Pr(s_t = 1|\Omega_{t-1}) \\ + f(Y_t|s_t = 0, \Omega_{t-1}) \Pr(d_t|s_t = 0, \Omega_{t-1}) \Pr(s_t = 0|\Omega_{t-1}). \end{aligned} \tag{7}$$

The contribution to the likelihood function of a given period is given by the sum of two products, one for each state. The components of these products are the state-conditional density function of the observed recovery rates, the state-conditional probability of the number of defaults and the probability of being in that state. These components are now described separately.

A.1 State-conditional recoveries

We have assumed that conditional on the state, recoveries of different firms are independent. This allows us to factorize, such that the conditional density for a given state s_t can be written as

$$f(Y_t|s_t, \Omega_{t-1}) = \prod_{i=1}^{d_t} f(y_{ti.}|s_t, \Omega_{t-1}). \tag{8}$$

The parameters of our marginal distributions of recovery rates vary according to state s, seniority class c and industry k. The marginal density was assumed to be given by

$$f(y_{tic}|s_t, \Omega_{t-1}) = \frac{1}{B(\alpha_{sck}, \beta_{sck})} y_{tic}^{\alpha_{sck}-1}(1 - y_{tic})^{\beta_{sck}-1}. \tag{9}$$

We assumed that the copula of the elements of $y_{ti.}$ is Gaussian with correlation matrix $\mathbf{\Gamma}$. Together with the assumption of the marginal distribution in [9] this determines the densities $f(y_{ti.}|s_t, \Omega_{t-1})$ in [8].

A.2 State-conditional number of defaults

Since we assumed that conditional on the state, defaults are independent with probability r_s, the probability of observing d_t defaults is binomial conditional on the state. Let N_t be the number of (defaulted and non-defaulted) firms in the population in period t, then we can calculate the probability of observing d_t defaults in a given state s_t as

$$\Pr(d_t|s_t, \Omega_{t-1}) = \binom{N_t}{d_t} r_s^{N_t}(1 - r_s)^{N_t - d_t} \tag{10}$$

A.3 Probabilities of the states

Finally, the probability of being in one state can be derived from the total probability theorem. Obviously,

$$\Pr(s_t = 1|\Omega_{t-1}) = p \cdot \Pr(s_{t-1} = 1|\Omega_{t-1}) + (1 - q) \cdot \Pr(s_{t-1} = 0|\Omega_{t-1}), \tag{11}$$

where we used the notation defined previously, $p = \Pr(s_t = 1|s_{t-1} = 1, \Omega_{t-1})$, the probability of remaining in state 1, and $(1 - q) = \Pr(s_t = 0|s_{t-1} = 0, \Omega_{t-1})$, the probability of moving from state 0 to state 1.

We can rearrange [11] to obtain

$$\Pr(s_t = 1|\Omega_{t-1}) = (1-q) + (p+q-1)\Pr(s_{t-1} = 1|\Omega_{t-1}) \tag{12}$$

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

$$\Pr(s_{t-1} = 1|\Omega_{t-1}) = \frac{f(Y_{t-1}|s_{t-1} = 1, \Omega_{t-2})\Pr(d_{t-1}|s_{t-1} = 1, \Omega_{t-2})\Pr(s_{t-1} = 1|\Omega_{t-2})}{f(Y_{t-1}|\Omega_{t-2})} \tag{13}$$

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 The Diebold-Gunther-Tay test

Diebold et al. (1998) propose a test (DGT test) for 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 tested in various ways, e.g. via a Kolmogorov-Smirnov test. Departures from uniformity or iid-ness are also easily visible when looking at histograms and autocorrelation functions of the transformed series.

In order to apply the DGT test to our predicted recovery rate densities, we create a vector $y^{\dagger}$ with typical element $y_t^{\dagger} = \text{vec}\left(Y_t^{\intercal}\right)$. 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). A slight complication arises due to the assumption of a Gaussian copula between recoveries of the same firm but across different seniorities. For these observations, conditional densities can easily be calculated via the conditional Gaussian distribution, though.

Applying the cumulative distribution function associated with these density forecasts to the vector $y^{\dagger}$ yields a vector of transformed variables which we call z. Under the null hypothesis that the density forecasts are correct, the elements of z 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.

C Tables and Figures

C.1 Tables

Table 1
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	39.64	14.27
1983	0.75%	5	48.24	20.35
1984	0.90%	11	48.88	16.59
1985	1.10%	14	48.17	21.28
1986	1.71%	26	35.19	18.16
1987	0.94%	19	52.89	27.05
1988	1.42%	35	37.19	20.33
1989	1.67%	41	43.55	28.29
1990	2.71%	81	25.49	21.80
1991	3.26%	94	40.37	26.27
1992	1.37%	37	51.50	24.02
1993	0.55%	21	37.58	19.61
1994	0.61%	16	43.77	24.88
1995	1.01%	24	43.76	24.69
1996	0.49%	18	43.59	23.79
1997	0.62%	23	54.95	23.76
1998	1.31%	32	46.55	24.52
1999	2.15%	94	30.29	19.92
2000	2.36%	112	28.03	23.81
2001	3.78%	137	24.71	18.05
2002	3.60%	100	29.78	16.69
2003	1.92%	58	39.24	23.48
2004	0.73%	35	50.59	24.13
2005	0.55%	32	58.71	23.41

Table 2
Recovery Rates by Seniority

This table reports the 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.

Seniority	Number of observations	Mean Recovery	Standard Deviation
All default events			
Senior Secured	210	42.26	25.76
Senior Unsecured	376	36.86	23.54
Senior Subordinated	334	32.73	23.66
Subordinated	158	34.17	23.00
Default events with single seniority only			
Senior Secured	158	40.11	23.90
Senior Unsecured	276	35.62	22.09
Senior Subordinated	226	33.69	23.29
Subordinated	90	37.91	20.21
Default events with multiple seniorities only			
Senior Secured	52	48.81	29.77
Senior Unsecured	100	40.30	26.85
Senior Subordinated	108	30.70	24.28
Subordinated	68	29.24	25.42

Table 3
Recovery Rates by Industry

This table reports the number of observations and the mean and standard deviation of recovery rates in our sample classified by industry.

Industry	Number of observations	Mean Recovery	Standard Deviation
Building	15	32.19	30.32
Consumer	152	35.56	22.89
Energy	50	37.65	17.09
Financial	104	36.91	25.99
Leisure	53	46.03	27.77
Manufacturing	368	35.78	22.89
Mining	15	35.05	18.67
Services	76	33.41	25.83
Telecom	123	31.53	21.14
Transportation	66	38.99	24.01
Utility	23	46.93	28.14
Others	33	38.01	23.44

Table 4
Basic Model - Parameter Estimates

This table reports the parameter estimates for the the basic model (no industry and seniority), in the static version (i.e. default probability and recovery rate distribution are constant, there is no credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle). The first column indicates the model being estimated. The second column refers to the state of the cycle (when applicable). The next three columns show the estimated parameters α and β and the mean of the estimated recovery rate distribution. The last two columns are the default and transition probability of each state respectively.

	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Basic Static		1.4474	2.9288	0.3675	0.0147	-
Basic Dynamic	Upturn	1.9860	2.7241	0.4685	0.0086	0.8707
	Downturn	1.4181	3.5990	0.3141	0.0269	0.7408

Table 5
Industry Model - Parameter Estimates

This table reports the parameter estimates for the industry dynamic model. The first column indicates the industry, while the second column refers to the state of the cycle. The next three columns show the estimated parameters α and β and the mean of the estimated recovery rate distribution. The last two columns are the default and transition probability of each state respectively.

Industry	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Building	Upturn	2.9565	4.2768	0.4541	0.0086	0.8745
	Downturn	0.8783	6.0405	0.1410	0.0269	0.7488
Consumer	Upturn	1.7818	2.8330	0.4290	0.0086	0.8745
	Downturn	1.6379	3.4976	0.3544	0.0269	0.7488
Energy	Upturn	3.6112	5.7415	0.4290	0.0086	0.8745
	Downturn	4.8878	8.4901	0.4060	0.0269	0.7488
Financial	Upturn	1.7054	2.1217	0.4951	0.0086	0.8745
	Downturn	1.2534	1.7697	0.4607	0.0269	0.7488
Leisure	Upturn	1.4490	1.8521	0.4877	0.0086	0.8745
	Downturn	1.7519	2.9508	0.4139	0.0269	0.7488
Manufacturing	Upturn	2.2361	2.7325	0.5001	0.0086	0.8745
	Downturn	1.4409	4.0063	0.2939	0.0269	0.7488
Mining	Upturn	2.9327	3.6264	0.4968	0.0086	0.8745
	Downturn	2.7132	7.1598	0.3053	0.0269	0.7488
Services	Upturn	2.6556	4.1652	0.4326	0.0086	0.8745
	Downturn	0.9950	3.0480	0.2734	0.0269	0.7488
Telecom	Upturn	2.7992	5.0656	0.3955	0.0086	0.8745
	Downturn	1.6347	5.2445	0.2640	0.0269	0.7488
Transportation	Upturn	1.3672	2.4168	0.4015	0.0086	0.8745
	Downturn	2.1564	5.3268	0.3202	0.0269	0.7488
Utility	Upturn	3.3697	2.7382	0.6130	0.0086	0.8745
	Downturn	1.0017	2.2649	0.3407	0.0269	0.7488
Others	Upturn	3.5737	4.5783	0.4871	0.0086	0.8745
	Downturn	2.1649	5.6149	0.3092	0.0269	0.7488

Table 6
Seniority Model - Parameter Estimates

This table reports the parameter estimates for the dynamic seniority model, estimated only on default events for which we do not observe recoveries across seniority classes, so that we can ignore dependence across seniorities. The first column indicates the seniority, while the second column refers to the state of the cycle. The next three columns show the estimated parameters α and β and the mean of the estimated recovery rate distribution. The last two columns are the default and transition probability of each state respectively.

Seniority	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Senior Secured	Upturn	2.8179	2.6014	0.5777	0.0084	0.8249
	Downturn	1.5080	3.7242	0.3202	0.0267	0.6868
Senior Unsecured	Upturn	2.1664	2.9563	0.4699	0.0084	0.8249
	Downturn	1.5872	4.1010	0.3100	0.0267	0.6868
Senior Subordinated	Upturn	1.4265	2.2403	0.4323	0.0084	0.8249
	Downturn	1.2165	3.1425	0.3101	0.0267	0.6868
Subordinated	Upturn	2.8277	4.7524	0.4145	0.0084	0.8249
	Downturn	1.6246	3.4969	0.3525	0.0267	0.6868

Table 7
Dependence across Seniorities Model - Parameter Estimates (Marginals)

This table reports the parameter estimates of the dynamic seniority model, estimated on all default events (including those with recoveries across multiple seniority classes). Here, only the parameters for marginal distributions are reported. Please see Table 8 for the estimated correlations between seniorities. Estimates are grouped into *Single seniority only* and *Multiple seniority only*. The second column indicates the seniority, while the third column refers to the state of the cycle. The next three columns show the estimated parameters α and β and the mean of the estimated recovery rate distribution. The last two columns are the default and transition probability of each state respectively.

Type	Seniority	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Single seniority only	Senior Secured	Upturn	2.8207	2.6349	0.5745	0.0086	0.8154
		Downturn	1.4953	3.6863	0.3206	0.0271	0.6261
	Senior Unsecured	Upturn	2.1788	2.9797	0.4693	0.0086	0.8154
		Downturn	1.5809	4.0832	0.3101	0.0271	0.6261
	Senior Subordinated	Upturn	1.4170	2.2870	0.4251	0.0086	0.8154
		Downturn	1.2117	3.1484	0.3088	0.0271	0.6261
	Subordinated	Upturn	2.8452	4.9285	0.4067	0.0086	0.8154
		Downturn	1.4951	3.1791	0.3554	0.0271	0.6261
Multiple seniority only	Senior Secured	Upturn	5.8617	2.2875	0.7992	0.0086	0.8154
		Downturn	1.3292	2.5514	0.3816	0.0271	0.6261
	Senior Unsecured	Upturn	1.4327	1.6664	0.5137	0.0086	0.8154
		Downturn	1.0724	2.4544	0.3379	0.0271	0.6261
	Senior Subordinated	Upturn	1.1786	2.1021	0.3992	0.0086	0.8154
		Downturn	0.8875	2.4601	0.2946	0.0271	0.6261
	Subordinated	Upturn	1.5457	3.0064	0.3773	0.0086	0.8154
		Downturn	0.6590	2.8774	0.2071	0.0271	0.6261

Table 8
Dependence across Seniorities Model - Parameter Estimates (Correlations)

These are the correlations implied by the Dependence across Seniorities Model. None of the correlations involving Senior Secured are different from zero at a 5% level of significance.

	S.Sec.	S.Uns.	S.Sub.	Sub.
S.Sec.	1.0000			
S.Uns.	0.2942	1.0000		
S.Sub.	0.1935	0.5449	1.0000	
Sub.	-0.2427	0.4976	0.8008	1.0000

Table 9
Industry / Seniority Model - Parameter Estimates
(Non-Multiple Marginals)

This table reports the parameter estimates for the Industry / Seniority Model, for default events for which default is observed only on a single seniority (non-multiple default events). Here, only parameters of the marginal distributions are reported. For the parameter estimates for default events for which default is observed across multiple senorities, please see Table 10. Please see Table 11 for the estimated correlations across these seniorities. The first two columns indicate the industry and seniority, while the third column refers to the state of the cycle. The next three columns show the estimated parameters α and β and the mean of the estimated distribution. The last two columns are the default and transition probability of each state respectively.

Industry	Seniority	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Group A	Senior Secured	Upturn	2.9140	2.5464	0.5930	0.0085	0.8752
		Downturn	1.3481	2.8525	0.3566	0.0267	0.7487
	Senior Unsecured	Upturn	1.5688	2.9588	0.3850	0.0085	0.8752
		Downturn	1.5668	2.5826	0.4196	0.0267	0.7487
	Senior Subordinated	Upturn	0.9822	1.2307	0.4932	0.0085	0.8752
		Downturn	1.7561	3.7186	0.3564	0.0267	0.7487
	Subordinated	Upturn	2.6826	4.5252	0.4135	0.0085	0.8752
		Downturn	1.0515	1.4742	0.4626	0.0267	0.7487
Group B	Senior Secured	Upturn	2.9346	2.8156	0.5671	0.0085	0.8752
		Downturn	2.0018	5.5988	0.2926	0.0267	0.7487
	Senior Unsecured	Upturn	2.4316	2.8778	0.5089	0.0085	0.8752
		Downturn	1.6957	4.2500	0.3169	0.0267	0.7487
	Senior Subordinated	Upturn	1.6778	3.0337	0.3957	0.0085	0.8752
		Downturn	1.2955	3.1757	0.3219	0.0267	0.7487
	Subordinated	Upturn	2.5997	4.1187	0.4299	0.0085	0.8752
		Downturn	2.1275	5.4169	0.3133	0.0267	0.7487
Group C	Senior Secured	Upturn	1.5760	2.1701	0.4674	0.0085	0.8752
		Downturn	1.3022	3.3210	0.3130	0.0267	0.7487
	Senior Unsecured	Upturn	3.4871	5.4421	0.4339	0.0085	0.8752
		Downturn	1.9317	7.0903	0.2379	0.0267	0.7487
	Senior Subordinated	Upturn	1.8059	2.7522	0.4402	0.0085	0.8752
		Downturn	0.9438	3.0932	0.2598	0.0267	0.7487
	Subordinated	Upturn	2.5959	4.3902	0.4129	0.0085	0.8752
		Downturn	1.3569	4.1573	0.2734	0.0267	0.7487

Table 10
Industry / Seniority Model - Parameter Estimates (Multiple Marginals)

This table reports the parameter estimates for the Industry / Seniority Model, only for default events for which default is observed on more than one seniority (non-multiple default events). Here, only parameters of the marginal distributions are reported. For the parameter estimates for default events for which default is observed only one seniorities, please see Table 9. Please see Table 11 for the estimated correlations across the seniorities. The first two columns indicate the industry and seniority, while the third column refers to the state of the cycle. The next three columns show the estimated parameters α and β and the mean of the estimated distribution. The last two columns are the default and transition probability of each state respectively.

Industry	Seniority	Cycle	α	β	Implied Mean Recovery	Default Probability	Probability of staying in the same state
Group A	Senior Secured	Upturn	4.2663	1.9468	0.7630	0.0085	0.8752
		Downturn	2.7358	2.8585	0.5434	0.0267	0.7487
	Senior Unsecured	Upturn	1.2148	2.6416	0.3500	0.0085	0.8752
		Downturn	0.9257	1.6984	0.3920	0.0267	0.7487
	Senior Subordinated	Upturn	0.8468	1.3719	0.4241	0.0085	0.8752
		Downturn	0.8376	2.4929	0.2794	0.0267	0.7487
	Subordinated	Upturn	1.5556	3.4942	0.3423	0.0085	0.8752
		Downturn	0.6408	2.5614	0.2223	0.0267	0.7487
Group B	Senior Secured	Upturn	2.8220	1.7865	0.6804	0.0085	0.8752
		Downturn	1.4238	4.5619	0.2643	0.0267	0.7487
	Senior Unsecured	Upturn	2.2624	2.5336	0.5241	0.0085	0.8752
		Downturn	1.1351	1.6423	0.4541	0.0267	0.7487
	Senior Subordinated	Upturn	2.2301	4.7745	0.3538	0.0085	0.8752
		Downturn	0.9744	2.7496	0.2907	0.0267	0.7487
	Subordinated	Upturn	2.1047	5.7251	0.2987	0.0085	0.8752
		Downturn	0.6795	3.1701	0.1961	0.0267	0.7487
Group C	Senior Secured	Upturn	2.6236	1.2105	0.7603	0.0085	0.8752
		Downturn	1.5581	3.4703	0.3443	0.0267	0.7487
	Senior Unsecured	Upturn	4.7929	3.0807	0.6764	0.0085	0.8752
		Downturn	0.9945	3.0158	0.2755	0.0267	0.7487
	Senior Subordinated	Upturn	2.2040	2.5248	0.5179	0.0085	0.8752
		Downturn	0.7572	1.3666	0.3961	0.0267	0.7487
	Subordinated	Upturn	2.8524	3.6745	0.4856	0.0085	0.8752
		Downturn	0.5384	1.2774	0.3295	0.0267	0.7487

Table 11
Industry / Seniority Model - Parameter Estimates (Correlations)

These are the correlations implied by the Full Model. None of the correlations involving Senior Secured are different from zero at a 5% level of significance.

	S.Sec.	S.Uns.	S.Sub.	Sub.
S.Sec.	1.0000			
S.Uns.	0.4216	1.0000		
S.Sub.	0.1606	0.5281	1.0000	
Sub.	-0.2429	0.6911	0.7595	1.0000

Table 12
Diagnostic Tests

This table presents diagnostic tests of the recoveries and default frequencies. Recovery rates tests are in Panel A. The first column is the p-value for the Kolmogorov-Smirnov (K-S) test where the null hypothesis is that the transformed series is drawn from a uniform(0,1) distribution. The second column reports the p-values of the Ljung-Box (L-B) test, where the null is that the transformed series is white noise. Panel B reports the test for the estimated default frequencies. The p-value reported is under the null hypothesis that the model is correctly specified.

Panel A: Recovery Rates		
Model	K-S	L-B
Basic Static	0.0232	0.0000
Basic Dynamic	0.0638	0.2056
Full Static	0.0991	0.0000
Full Dynamic	0.1437	0.4892

Panel B: Default rates	
Model	p-value
Basic Static	0.0000
Basic Dynamic	0.9641
Full Static	0.0000
Full Dynamic	0.9190

Table 13
95% VaR for Selected Models.

This table presents the VaRs of simulated one-year loss distribution for a hypothetical portfolio of 500 bonds. The VaRs implied by the basic model is contrasted with the VaRs implied by the full model, both for the dynamic and static specifications. For the dynamic specifications, the initial probability of being in an upturn matters. We contrast the cases where this probability is equal to 1 ("Upturn") with the situation where this is equal to zero ("Downturn"), and the situation where this probability is equal to the unconditional probability of being in an upturn (i.e. we have no knowledge of the initial state of the credit cycle).

		Static specification	Dynamic specification		
			Upturn	Unconditional	Downturn
Basic model					
		0.0158	0.0196	0.0239	0.0263
Full model					
Group A	Senior Secured	0.0148	0.0194	0.0235	0.0260
	Senior Unsecured	0.0155	0.0176	0.0214	0.0240
	Senior Subordinated	0.0153	0.0189	0.0233	0.0258
	Subordinated	0.0153	0.0166	0.0203	0.0227
Group B	Senior Secured	0.0160	0.0209	0.0252	0.0280
	Senior Unsecured	0.0159	0.0200	0.0245	0.0273
	Senior Subordinated	0.0171	0.0196	0.0246	0.0272
	Subordinated	0.0159	0.0202	0.0246	0.0276
Group C	Senior Secured	0.0172	0.0203	0.0246	0.0275
	Senior Unsecured	0.0184	0.0221	0.0268	0.0299
	Senior Subordinated	0.0179	0.0215	0.0269	0.0294
	Subordinated	0.0169	0.0216	0.0265	0.0288

Table 14

Correlations of Macro Variables with the Credit Cycle.

This table presents the correlation matrix for various macroeconomic variables and the smoothed credit cycle. The variables are gdp: GDP growth, unemp: Unemployment, invest: Investment, stock: S&P 500 return; r2y: 2 year treasury yield, r10y: 10 year treasury yield, r_slope: difference between 10 year and 2 year yield, upturn: the smoothed probability of being in a credit upturn.

gdp	unemp	invest	stock	r2y	r10y	r_slope	upturn	
1.0000	−0.2250	0.8714	0.4277	0.1958	0.1160	−0.3242	0.3443	gdp
	1.0000	−0.0909	0.0029	0.5335	0.6799	0.3316	0.3985	unemp
		1.0000	0.3084	0.2099	0.1625	−0.2204	0.4178	invest
			1.0000	0.1915	0.0845	−0.4142	0.3591	stock
				1.0000	0.9624	−0.4205	0.4288	r2y
					1.0000	−0.1583	0.3679	r10y
						1.0000	−0.3301	r_slope
							1.0000	upturn

C.2 Figures

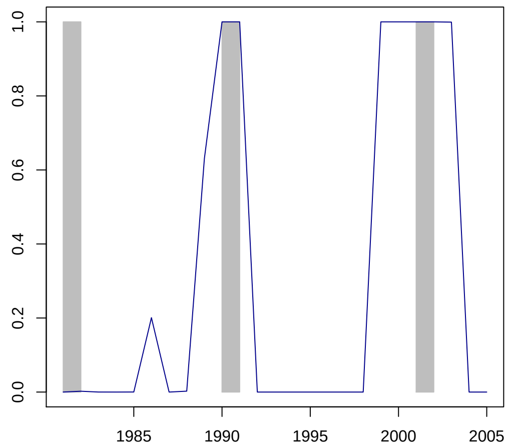

Figure 1. Probabilities of being in the credit downturn versus NBER recessions (annual)
This figure plots the annual smoothed probability of being in a credit downturn as estimated on the basis of the basic dynamic model (no industry and seniority), assuming that the last quarter of 1981 was a credit downturn with a probability corresponding to the unconditional probability of being in a downturn. This is contrasted with NBER recessions (grey areas).

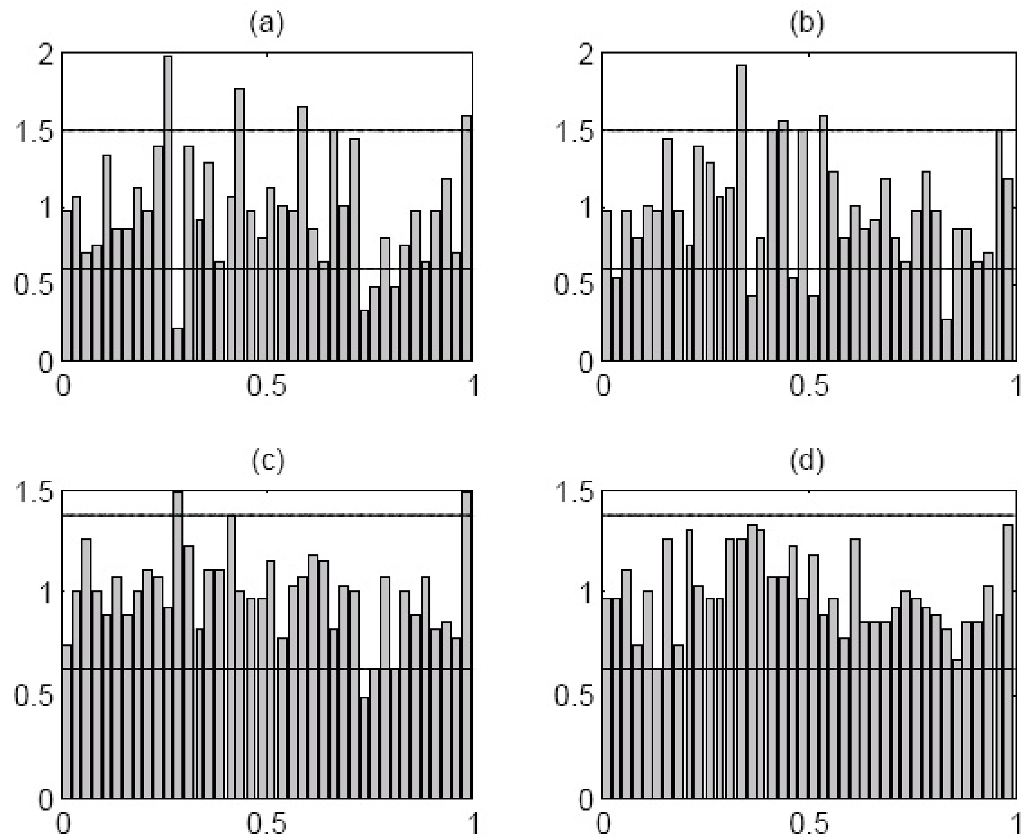

Figure 2. Histograms of the DGT-transformed recoveries.
This figure presents the histograms of the DGT-transformed recoveries for (a) the basic static model, (b) the basic dynamic model, (c) the full static model, and (d) the full dynamic model.

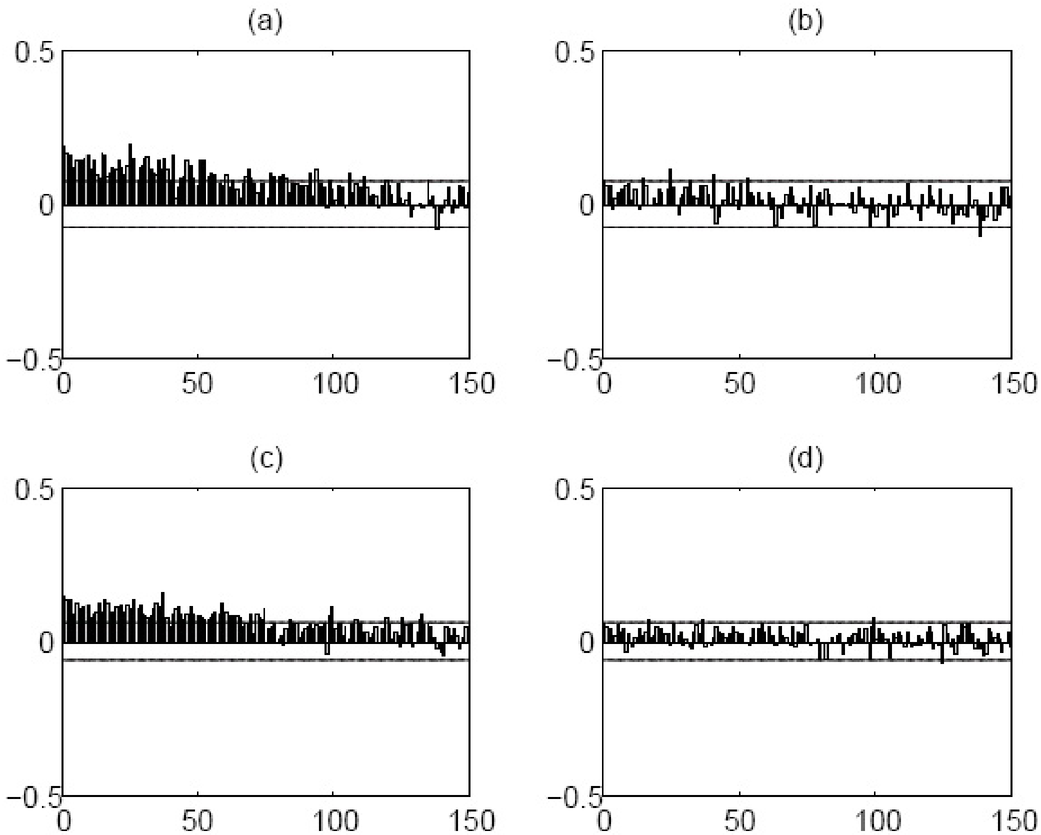

Figure 3. Correlograms of the DGT-transformed recoveries
This figure presents the correlograms of the DGT-transformed recoveries for (a) the basic static model, (b) the basic dynamic model, (c) the full static model, and (d) the full dynamic model

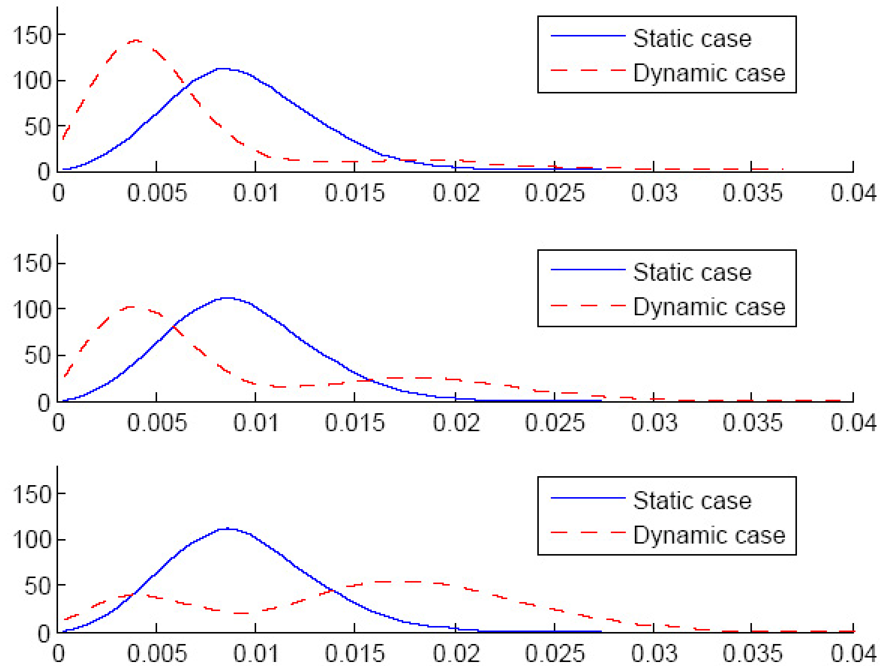

Figure 4. Simulated loss density
This figure contrasts the simulated loss density (pdf), of the basic model (no industry and seniority), for the static version (i.e. default probability and recovery rate distribution do not vary with the credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle). For the dynamic model, the probability of being in a credit downturn is assumed to be equal to zero in the upper panel, the unconditional probability in the middle one and one in the lower panel.