Journal of Financial Economics 125 (2017) 454–474

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Debt correlations in the wake of the financial crisis: What are appropriate default correlations for structured products?[☆]

Jordan Nickerson[a,*], John M. Griffin[b]

[a] Carroll School of Management, Boston College140 Commonwealth Avenue, Chestnut Hill, MA 02467, USA
[b] McCombs School of Business, The University of Texas at Austin, 1 University Station, B6600, Austin, TX 78712, USA

ARTICLE INFO

Article history:
Received 27 August 2015
Revised 19 July 2016
Accepted 25 July 2016
Available online 3 July 2017

JEL classification:
G14
G24
G28
G32

Keywords:
Credit ratings
Financial crises
Structured finance
Default correlations

ABSTRACT

This paper proposes several frameworks to estimate the appropriate default correlations for structured products, each of which jointly considers the role of co-movements in modeled risk characteristics and unmodeled systematic risk, or 'frailty.' We contrast our estimates with credit rating agencies' default correlation assumptions, which were only 0.01 for Collateralized Loan Obligations (CLOs) pre-crisis and have increased to 0.03 post-crisis. In contrast, the joint consideration of observable risk factors and frailty leads to substantially higher estimates of 0.12. We show that this translates into CLOs with credit risk understated by 26%, suggesting caution for the post-crisis structured finance market.

1. Introduction

During the financial crisis, over 13,250 AAA-rated tranches with an issuance value of $1.26 trillion consequently defaulted on their claims.[1] A commonly perceived force of this activity is that actors did not understand the highly correlated nature of Mortgage Backed Securities (MBS), Collateralized Debt Obligations (CDO), and other structured finance collateral's default risk. A *Financial Times* article concisely summarizes: "Simply stated, what was supposed to be correlated in a certain way turned out to be correlated in a completely different fashion."[2] Federal Reserve Chairman Ben Bernanke told the Financial Crisis Inquiry Committee, "They did not take into account the appropriate correlation between [and] across the categories of mortgages."[3] Despite the level of attention paid to default correlations, the discussion remains qualitative in nature. No work has quantified what exact default

☆ We thank Pierluigi Balduzzi, Jonathan Cohn, Jess Cornaggia, Jennifer Dlugosz (Carnegie Mellon Credit Rating discussant), Darrell Duffie, Cesare Fracassi, Jay Hartzell, Oğuzhan Karakaş, Darren Kisgen, Tim Landvoigt, Andrew Lo, Jeff Pontiff, Uday Rajan, Jon Reuter, Alessio Saretto, Mark Seasholes, Gustavo Schwenkler (Cavalcade discussant), Tyler Shumway, Phil Strahan, Dragon Tang, Jérôme Taillard, Sheridan Titman, James Vickery (WFA discussant), and seminar participants at Boston College, University of Cambridge, the Cavalcade 2015 meetings, The University of Texas at Austin, and Western Finance Association 2015 meetings for valuable discussions. We also thank the McCombs Research Excellence Fund for financial support. Griffin consults on issues regarding credit ratings, MBS, and CDOs and is also a party to litigation against major banks regarding MBS fraud. The analyses and conclusions of the paper are unaffected. Additional results are available in an Online Appendix at http://www.JGriffin.info and http://www.JordanNickerson.com.

* Corresponding author.
E-mail addresses: jordan.nickerson@bc.edu (J. Nickerson), john.griffin@mail.utexas.edu (J.M. Griffin).

[1] Our calculation is based on data pulled from Bloomberg on the universe of 2,350 structured products issued between January 2000 and December 2007 that defaulted between January 2008 and May 2014.

[2] Pablo Triana, July 26, 2010, *Financial Times*.

[3] Financial Crisis Inquiry Commission (2010, p.149)

http://dx.doi.org/10.1016/j.jfineco.2017.06.011
0304-405X/

correlations were assumed by rating agencies prior to 2007 and to what extent these estimates have been updated in the wake of the financial crisis—an important question given the revival of the structured finance market with over $3.10 trillion in securities issuance between January 2010 and June 2015.[4]

Our first objective is to obtain a sense of what pre-crisis correlations were. To examine this question, we back out default correlation estimates from pre-crisis rating agency data. We build several intuitive methodologies to derive default correlation estimates across a set of modeling frameworks. We then compare these estimates to those used by rating agencies for what industry sometimes refers to as '2.0' (or post-crisis) structured products.

The infrequent nature of default events makes it difficult to model their underlying correlated nature. Additionally, an important unmodeled or omitted variable can emerge and cause a wave of defaults at a particular point in time. Taleb (2007) popularized one narrative of this concept known as the 'black swan event' where previously unforeseen events cause chaos on bank risk models and lead to a crisis. More importantly, Duffie, Eckner, Horel and Saita (2009) propose a method to capture the tail loss risk associated with an unmodeled systematic risk factor. Even after controlling for a broad spectrum of firm-specific and macro-explanatory variables, they find that an unobservable time-varying factor referred to as 'frailty' can significantly help in explaining default clustering. They discuss the potential importance of frailty for portfolios of assets, such as those found in a CDO, but they do not show how it can be incorporated into default correlations, nor compare this to estimates currently used in practice. Our paper is the first to detail an approach to incorporate the effects of both frailty as well as the co-movement of observable risk characteristics into an estimate of default correlations. We document the default correlations assumed by the rating agencies, contrast their assumption with our estimates, and quantify the effects of our frailty-incorporated default correlation estimates on the appropriate size of actual senior AAA-rated CDO tranches.

The traditional theoretical literature focuses on the correlation in default intensities of assets.[5] In contrast, our study examines the correlation in the realization of defaults between assets. While the modeling of correlations among default intensities has clean mathematical properties, it is conceptually difficult to map such default intensity correlations to actual defaults. Ultimately, the creditworthiness of a structured finance product is dictated by the realized defaults of its underlying collateral pool. For this reason, practitioners primarily focus on the correlation of realized defaults. Credit rating agencies specifically mention a concern for achieving the appropriate correlation of realized defaults (Moody's, 2010) and base their final metrics of both collateral correlation and collateral risk on the distribution of realized asset defaults (Standard & Poor's, 2013). Thus, by estimating default correlations from realized defaults, we are able to directly compare the operating assumptions of rating agencies to estimates of joint collateral risk under our framework.

As a benchmark for common practice, we begin by asking what correlation levels were assumed by rating agencies for structured finance products leading up to the financial crisis. We back out default correlations from rating agency data and find that Standard and Poor's (S&P) and Moody's assumed an average default correlation from 1997 to 2007 of 0.01. To provide some economic context for the relevance of default correlations for CDOs, we show that a change in default correlation from 0.005 to 0.035 leads to approximately a 10% increase in the proportion of subordinated tranches needed to protect the claim of a senior AAA tranche.

Given the importance of default correlations, we use multiple distinct methodologies, each of which is based on systematic changes in both observable and unobservable risk factors, to estimate their appropriate level. The first class of models we consider is based on clustering in credit rating upgrades and downgrades using different characterizations of a state-dependent rating transition matrix. Ashcraft, Goldsmith-Pinkham and Vickery (2010) show that there is variation in performance beyond initial credit ratings based on other observable risk characteristics. In a similar manner, we also consider a second class of models which evaluates the importance of a panel of macroeconomic variables in explaining default risk. For each model, we then incorporate unobservable systematic changes in default risk, or 'frailty,' utilizing the framework of Duffie, Eckner, Horel and Saita (2009). With these tools, we are able to estimate default correlations for CDOs backed by corporate debt. In addition, by considering multiple models we are able to evaluate the sensitivity of our default correlation estimates to the choice of modeling assumptions.

For corporate bonds before the financial crisis (1986 to 2006), our estimated pairwise default correlation is only 0.002 when using only the state-dependent rating transition matrix. However, when allowing only for model frailty, the average pairwise bond default correlation jumps to 0.086. These default correlations are more than eight times those used by rating agencies for CLOs prior to the crisis. Furthermore, the inclusion of both rating changes and model frailty increases the average default correlation to 0.10. This estimate increases by roughly 25% to 0.125 when incorporating information contained in the financial crisis and estimating the models using a sample ending in December 2012. Overall, our findings show that the joint consideration of co-movement in observable risk factors and frailty can add considerable thickness to the right tail of the default distribution.

We now turn our attention to the extent to which rating agencies incorporated information gained from the financial crisis by examining a set of post-financial crisis CLOs. Using a small sample of 136 CLOs rated by S&P, we find that the average default correlation assumed by rating agencies has increased to 0.033 (as compared to 0.01 pre-crisis). Unfortunately, this number is considerably below

[4] Issuances are from the Securities Industry and Financial Markets Association reports (SIFMA) from 2010 through the second quarter of 2015. These totals are estimates and may be missing smaller categories.

[5] For example, see Azizpour, Giesecke and Kim (2011), Das, Duffie, Kapadia and Saita (2007), Giesecke (2004), Giesecke and Weber (2004), Koopman, Lucas and Schwaab (2012), Lando and Nielsen (2010), and Li (1999), among others.

our estimates across the varying frameworks. We find that the measure of AAA collateral risk currently assumed by S&P increases 20% under the models we consider, suggesting that the AAA tranche sizes of recently certified CLOs are too optimistic. In partial support of this, we find that yield spreads of AAA CDO tranches have increased substantially following the crisis. To avoid over-parameterizing, we have purposefully taken a straightforward and intuitive approach to incorporate rating changes and frailty, one that leaves room for future extensions.

For structured products, the most important determinant of credit risk is appropriately gauging the level of default risk and the correlation among the underlying assets. We focus on the latter, while Cornaggia, Cornaggia and Hund (2017) focus on the former by comparing the default probabilities across all rated asset classes. Griffin and Tang (2012) argue that there was a significant subjective component of CDO ratings beyond credit risk modeling that increased from 2002 to 2007. Griffin, Nickerson and Tang (2013) find that this subjective component was positively correlated with increased competition due to a competitor's more favorable rating assumptions, which includes default correlation. These upward adjustments were commonly in the 4–8% range. To maintain the same level of AAA notes, the incorporation of our default correlation estimates would require an additional upward adjustment of 17%, indicating that the correlation effects we document are quite sizeable. Given the findings of earlier research that rating agencies cater to their clients and keep ratings high (Becker and Milbourn, 2011; Griffin, Nickerson and Tang, 2013; and Cornaggia, Cornaggia and Hund, 2017), it is possible that rating agencies are not using higher correlation assumptions that incorporate frailty for fear of losing business to their competitors. Our evidence of rating agencies using low correlation assumptions post-crisis raises concerns of continued agency problems in credit rating agencies, despite much attention from regulatory bodies.

In terms of assessing other potential problems in structured finance products, Coval, Jurek and Stafford (2009) demonstrate that CDO valuation models hinge on a high degree of confidence in the parameter inputs. Our analysis shows that such confidence was unjustified, as correlation assumptions can vary widely. We hope increased research and transparency of underlying correlation assumptions will facilitate a better understanding of a re-emerging structured finance market.

2. Default correlation modeling overview

This section provides a brief discussion of default correlations and examines the default correlations assumed by credit rating agencies leading up to the financial crisis.

2.1. Default correlations background

The evaluation of the underlying collateral supporting the asset balance sheet of a CDO involves the measurement of two characteristics of the pool—the collateral's quality and default correlation. While the quality of the underlying collateral determines the mean of the pool's default distribution, the collateral's default correlation determines the joint likelihood of default across multiple underlying assets, and thus the thickness of the default distribution's upper tail. This correlation has the largest effect on senior CDO note holders who receive credit protection on their claims provided by the subordinated debt.[6]

The extant literature largely examines the joint credit risk of multiple assets by modeling the correlation of default intensities or asset lives.[7] An advantage of these approaches is their ability to abstract away from a specific time horizon in which assets' defaults are realized. While such techniques lend themselves to the study of joint credit risk in a broad sense, a limitation of this approach is the difficulty in applying the resulting estimates to the joint default risk of a specific set of assets. Specifically, the joint default risk for specific assets is also a function of the assets' unconditional default probabilities, as discussed in more detail below.

The default intensity approach also commonly relies on the doubly stochastic assumption that, conditional on the observed paths of risk factors, realized defaults are independent. Duffie, Eckner, Horel and Saita (2009) note that the full set of risk factors is not observed by the econometrician, and hence from the econometrician's vantage point, defaults are not doubly stochastic. This leads to their use of a time-varying frailty component. We follow their lead in incorporating additional systematic risk resulting from a failure in the doubly stochastic assumption.[8]

In contrast to the use of default intensities, we take an alternative approach which measures default correlation based on the realized defaults of assets. The use of realized defaults has two main benefits. First, the approach is advantageous in that it mirrors concerns of an investor who is concerned with the risk of default inherent in a specific set of assets over the life of a CDO. This motivation is highlighted by rating agency methodologies, which are based on the simulation of realized defaults (Standard & Poor's, 2013; Moody's, 2014) and the explicit statement of their concern for achieving an appropriate correlation of realized defaults (Moody's, 2010). Consistent with this reasoning, realized default correlations are strongly correlated with the measure of credit risk used by S&P, the Scenario Default Rate (SDR). In contrast, the correlation of default intensities can be unrelated to SDR. Internet Appendix Fig. IA.1 demonstrates this with an example in which the default correlation and SDR of assets are related, while the correlation of the assets' default intensities is held fixed at unity. Hence, the correlation of default intensities can be

[6] While the evaluation of structured finance products involves modeling both asset credit risk and cash-flow waterfalls, both S&P and Moody's methodologies treat the two components independently of each other. Thus, by focusing on the collateral risk of the underlying asset pool, we can cleanly evaluate the default correlation assumptions used by credit rating agencies relative to our estimates separately from deal-specific cash-flow protections.

[7] For example, Das, Duffie, Kapadia and Saita (2007) study the extent to which default intensities modeled as a log-linear function of firm and macro observables can explain clustering in firm defaults. Feldhütter and Nielsen (2012) use Credit Default Swap (CDS) and CDS index (CDX) spreads to estimate a firm's default intensity as the linear function of a common and idiosyncratic component.

[8] Nonetheless, we must parameterize the frailty process to do so, which if misspecified would only partially address this failure.

disconnected from the tail-risk of a collateral pool while the correlation of realized defaults is highly related.

Focusing on realized default correlation leads to a second main benefit, namely, that we can directly compare our estimates to the credit rating agencies. While the correlation in default intensities is not easily compared to the methods used by rating agencies, we are able to derive closed-form solutions of realized default correlations from the proprietary correlation metrics used by both S&P and Moody's at the deal level.

A potential drawback of using correlations of realized defaults is that the estimate varies with the maturity of the assets. More specifically, as the probability of default for the underlying assets approaches zero, the correlation in realized defaults attenuates towards zero. This asymptotic result can occur by reducing the time-horizon to zero over which defaults are measured. Thus, when using the correlation in realized defaults to estimate the joint risk of default for two assets over a very short period, such as a week, the estimated default correlation will be close to zero for mechanical reasons. Alternatively, this result arises when estimating the default correlation between assets with less credit risk (and thus a smaller probability of default). Aware of this concern, Moody's (2010) states that they "increase the asset correlation assumptions in the investment-grade rating categories...in order to generate default correlations in line with current observations." We explore the sensitivity of our estimates with respect to this conditional dependence in Section 5.4, and also use the rating agencies' deal-specific weighted-average maturity in order to effectively compare our estimated agency assumptions in a matched sample. Overall, we believe that the choice of estimating correlations from realized defaults rather than from default intensities is beneficial given the allowance for frailty, its relation to the tail-risk of CDO collateral pools, and direct comparability with rating agency assumptions.

2.2. Pre-crisis correlation assumptions for Moody's and S&P

Before introducing our approach to estimating default correlations, it is useful to examine the assumptions used by credit rating agencies (CRAs) prior to the financial crisis. Examining rating agency default assumptions provides us with a baseline to which we can compare the estimated default correlations stemming from our models prior to the crisis. If rating agency assumptions and our estimates are similar in magnitude pre-crisis, then it may suggest that the recent structured finance crisis more closely resembled a black swan. If our estimates are higher than the rating agencies, then it may suggest a risk factor present in the pre-crisis data that is unaccounted for by pre-crisis models used by CRAs.

We now turn to a sample of CDOs backed by bond and loan collateral.[9] Although default correlation assumptions of rating agencies have been a topic of speculation, they have escaped systematic examination. This may be related to the fact that rating agencies do not directly report the default correlation that is assumed for a CDO, but instead report proprietary measures such as 'diversity score' and 'correlation measure.' Thus, we first back out default correlation estimates from other model outputs provided by the rating agencies.[10]

Table 1 reports summary information for our sample. The average (median) collateral pool for a deal in our sample consists of roughly 140 (131) obligors. The potential diversification effect of such a large asset pool highlights the key role that correlation plays in determining risk to senior tranche holders. Table 1 also shows that for the average corporate debt-backed deal in the sample, S&P assumes a pairwise default correlation in the underlying collateral pool of 0.0099. Moody's average default correlation is only slightly higher at 0.011. To illustrate the economic meaning of these assumptions, consider the following example. Suppose there are two corporate bonds with identical probabilities of default p. Each bond's realization of default is stochastically determined by either a macro factor or a firm-specific idiosyncratic factor.[11] A default correlation of 0.01 implies that the probability that both bonds' default realization is determined by the common macro factor is 1.00%. In this case, they will either survive or default together, while they will otherwise perform independently of each other.

Panel B of Table 1 indicates that there is very little variation in default correlations over time, with similarly small default correlations across different years of origination. The sole exception is the slight increase in the default correlation assumed by S&P (0.0187) in the 2006–2007 period. In addition, there does not seem to be substantial disagreement between the two rating agencies' assumptions.

Finally, Fig. 1 reports the distribution of correlation assumptions made by S&P. The figure illustrates that the majority of Collateralized Bond Obligations (CBOs) and CLOs had an average pairwise default correlation of less than 0.0075. Note, there is a small portion of deals with a default correlation greater than 0.03. S&P's estimate of default correlation is based on correlation assumptions estimated within and across industries. Thus, the underlying collateral in such deals likely has a higher degree of industry concentration relative to the other deals in the sample.

2.3. Impact of default correlations on SDR

While the previous section documents the default correlations assumed by rating agencies, it is not clear how much these assumptions influenced the estimated collateral risk of the underlying pool of assets.

Fig. 2 illustrates how collateral risk is affected by changes in the default correlation and collateral quality

[9] Further details concerning the data can be found in Internet Appendix A.

[10] These derivations are shown in Appendix A. While S&P publishes methodological documents reporting inter- and intra-industry asset correlation assumptions, these reported asset correlations are not the same as default correlations which capture correlation in the realization of default. For a representative deal an asset correlation of 0.10 (0.20) is roughly equivalent to a default correlation of 0.05 (0.105).

[11] For simplicity, assume that the components share a common likelihood of default of p.

Table 1
Summary statistics.
This table reports summary statistics collected from S&P Presale and New Issue reports for 838 CDOs over the time period between January 2000 and December 2007. Panel A reports deal-level characteristics. Summary statistics for the average pairwise default correlation assumed are through time reported in Panel B.

Panel A: summary statistics

	N	Mean	Median	Std. dev.	p10	p90
S&P default correlation	838	0.0099	0.0065	0.0079	0.0042	0.0237
Moody's default correlation	831	0.0110	0.0095	0.0125	0.0076	0.0141
S&P collateral default probability	838	0.2367	0.2291	0.0745	0.1582	0.3241
Moody's collateral default probability	831	0.1908	0.1852	0.0554	0.1485	0.2480
No. of obligors	838	140.29	131.27	64.86	67	223
Deal size ($M)	834	354.75	353.50	272.71	0.46	588.00
Weighted-average maturity	838	5.91	5.63	1.33	4.83	7.56

Panel B: correlation by origination year

	S&P			Moody's		
	N	Mean	Median	*N*	Mean	Median
Pre-2004	266	0.0060	0.0052	264	0.0129	0.0099
2004	69	0.0065	0.0059	69	0.0098	0.0094
2005	102	0.0070	0.0060	102	0.0100	0.0091
2006	193	0.0077	0.0068	192	0.0098	0.0089
Post-2006	166	0.0187	0.0221	162	0.0098	0.0092

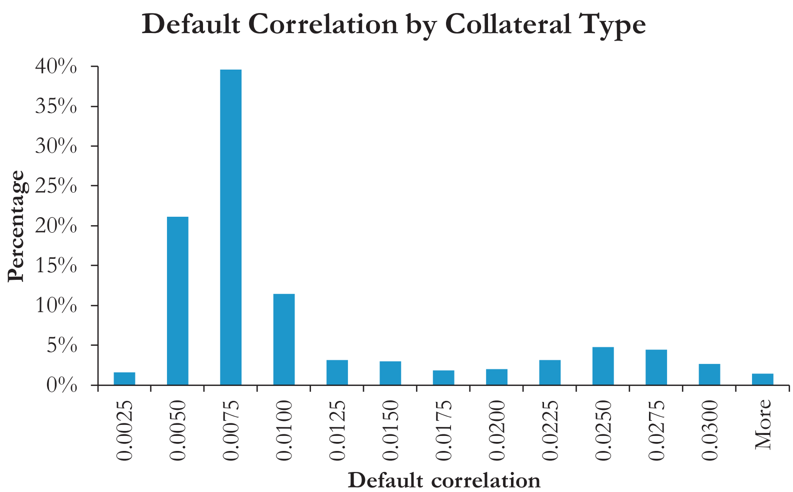

Fig. 1. Average default correlation assumptions. This figure reports the histogram of the average pairwise default correlation assumed in the S&P model for Collateralized Bond Obligations and Collateralized Loan Obligations. The sample consists of 838 CLOs and CBOs issued over the time period between January 2000 and December 2007. Details on the derivation of default correlations can be found in Appendix A.

of the underlying pool. S&P represents the riskiness of the collateral pool with the Scenario Default Rate (SDR), which is equivalent to the value-at-risk (VaR) of the collateral pool's default distribution.[12] Specifically, the SDR for a given rating and Weighted Average Maturity (WAM) (for instance, the AAA SDR for a CDO with a 7-year WAM) is the VaR with a confidence interval equal to the expected default probability of an asset with an equal credit rating and maturity (in this case a 7-year AAA bond). We provide an example of the SDR computation in detail in Appendix B.

Fig. 2 indicates a concavity where the SDR is particularly sensitive to changes in the correlation assumption for lower correlation levels. This makes the choice of default correlation particularly crucial for CDO modeling. Changing the default correlation from 0.005 to 0.015 increases the SDR of the pool by roughly 0.06, or 25% of its prior value. A CDO comprised of collateral of median quality with a correlation near zero has an SDR of about 32%, whereas a CDO with a default correlation of 0.035 has approximately 13% more SDR (45%). CDOs with a default correlation of 0.10 have a further 0.10 increase in their SDRs to slightly more than 55%. Many junior AAA tranches consisted of only 10% of the capital structure. Thus, even such a small change in correlation, could lead to a large economic effect on the amount of capital that can obtain a AAA rating or, alternatively, the rating itself.

[12] When the timing of defaults and cash-flow protection for a deal is ignored, the required credit support necessary to obtain a AAA rating is linearly increasing in the SDR.

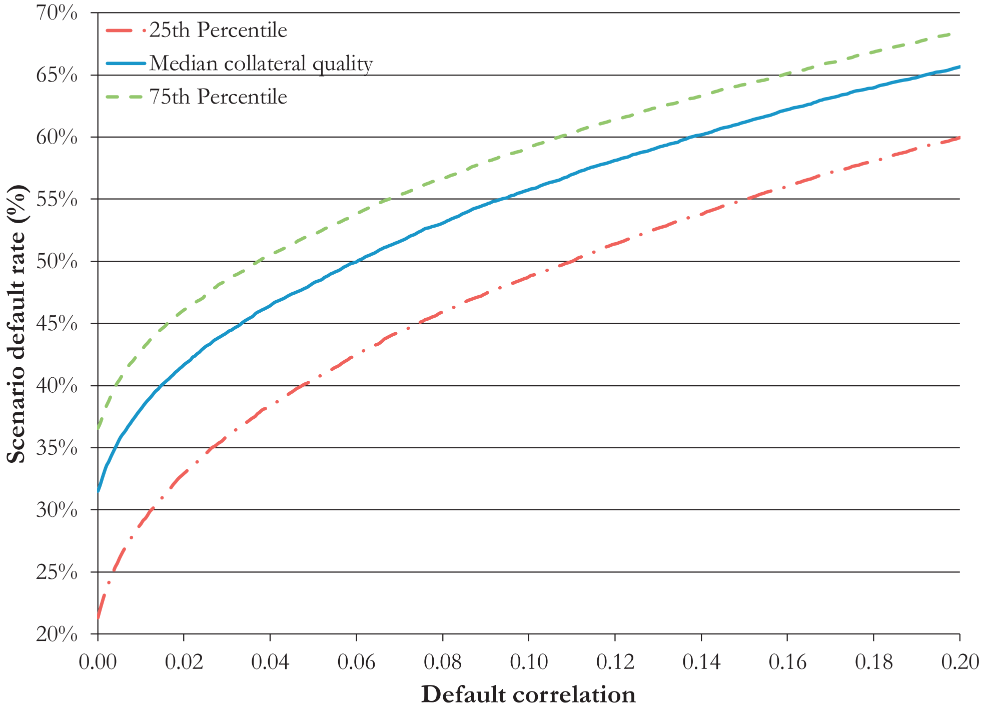

Fig. 2. Sensitivity of scenario default rate (SDR) to default correlation. This figure reports the surface plot of the change in SDR with respect to default correlation. The figure plots the SDR for a given range of default correlations for a CDO estimated with collateral quality set to the 25th (low quality), 50th (median quality), and 75th (high quality) percentiles of our sample. The number of equal-sized assets in the collateral pool is set to 131, equal to the median number of obligors in the CDOs of our sample.

Our preliminary examination has shown that the pre-crisis default correlations assumed by credit rating agencies for corporate debt were typically between 0.005 and 0.025, and that upper tail-risk is particularly sensitive to even small changes in the assumed default correlation at such levels. While these assumed levels seem low, we need a methodology to accurately assess if they are consistent with the empirical evidence.

3. Modeling joint default risk of CDO collateral

Estimating the default correlation of a CDO's asset pool requires assessing the joint default risk of the underlying assets. Our approach estimates this common risk component by incorporating the effects of observable risk factors and unobserved systematic risk in 100,000 simulated cross-sections of asset defaults for a given CDO's collateral pool, from which we are able to estimate a default correlation.[13]

Additionally, our approach also allows for an unobservable systematic component of default risk not captured by observables, or model *frailty*.[14] We begin by discussing the generic modeling of frailty before discussing the specific set of observable factors included.

3.1. Unobserved frailty

We model the arrival time of default for an asset as a Poisson process with intensity:

$$\lambda_{it} = e^{\alpha+\beta X_{it}+Y_t} \tag{1}$$

where X_{it} is a set of asset-specific covariates, and Y_t is a time-varying unobservable macro factor that affects the likelihood of default for all assets, or model frailty. Eq. (1) represents an exponential survival model, with the exception of the unobservable factor.

While the set of asset-specific covariates represents risk factors observable to the econometrician, default correlations should also account for the unmodeled risk whose systematic nature has implications for the severity of tail losses. We now propose an adaption of the frailty model of Duffie, Eckner, Horel and Saita (2009), which identifies the unobservable frailty path using time-varying deviations

[13] Specifically, we first estimate the time-series properties of explanatory variables that predict asset default. Using these estimates, for each CDO we draw 100,000 paths of all explanatory variables, such as rating transitions, with length equal to the structured product's weighted-average maturity. We then compute each asset's probability of default under each path drawn from the fitted values of a hazard model. From these default probabilities we simulate realizations of default within the pool.

[14] To be technical, 'frailty' typically refers to the case of what is left over when including an inexhaustible set of explanatory variables. For our first set of models which only include ratings, a more technical term for the second source of risk is 'unmodeled systematic risk.' However, since 'frailty' and 'unmodeled systematic risk' are conceptually and empirically similar, we use the term 'frailty' for brevity.

of realized firm defaults relative to expectations given observable risk factors.[15]

Given the unobservable nature of the model frailty, Y, for tractability structure must be imposed on its time-series dynamics. Following Duffie, Eckner, Horel and Saita (2009) we assume that the unobservable macro factor follows an Ornstein-Uhlenbeck (O-U) process:

$$dY_t = -\kappa Y_t dt + dB_t \tag{2}$$

where κ captures the speed of mean-reversion for the process while the series of innovations, dB, follows a standard Brownian motion. However, by assuming an O-U process we are fixing the variance of innovations. Ultimately, it is the volatility of innovations and the speed of mean-reversion that drive the correlation in asset defaults. Thus, we multiply Y by the scaling parameter η, such that the default intensity from (1) now becomes:

$$\lambda_{it} = e^{\alpha+\beta X_{it}+\eta Y_t} \tag{3}$$

where η is a scaling parameter which determines the role that the macro factor plays in a firm's default intensity. Intuitively, a larger value of η suggests that the unobservable macro factor plays a larger role in determining the likelihood of default for all firms in the economy and increases the correlation of defaults within a collateral pool.

To estimate the parameter set (η, κ), we follow the process outlined in Duffie, Eckner, Horel and Saita (2009) which uses an expectations maximization algorithm and Gibbs sampling to estimate the conditional distribution of Y. The full procedure is outlined in Internet Appendix B.

3.2. Explanatory covariates of collateral default

Rating agencies infer collateral risk for a CDO from the credit ratings of the underlying assets at the deal's origination date. Our first class of models relies on this same covariate when predicting asset default. However, in contrast to the rating agencies, we allow ratings to be updated contemporaneously when predicting asset default. Rating agency methodologies infer collateral risk by mapping credit ratings to historical default probabilities. We mirror this approach and use the default intensity implied by a firm's credit rating in Eq. (3).[16]

The second class of models we consider augments the ability of credit ratings to predict default with the additional information contained in macroeconomic explanatory variables. Specifically, in addition to a firm's credit rating we include the following covariates: the trailing one-year market return, 3-month Treasury rate, AAA credit spread over the 10-year Treasury yield, and the seasonally adjusted civilian unemployment rate from Federal Reserve Economic Data (FRED).[17]

3.3. Ratings and macro covariates as predictors of default

To estimate the models, we obtain S&P's long-term corporate credit ratings updated at a monthly level for roughly 2,000 firms from January 1, 1986 to December 31, 2013 from Compustat. Additionally, primary default information is obtained from the deletion date and reason fields provided by the Center for Research in Securities Prices (CRSP)/Compustat.

We use Maximum Likelihood Estimation (MLE) to estimate (3) where covariates are updated on a monthly basis. Thus, the model is estimated over a given horizon using the pooled sample of firm-month observations within this time period. Note that we use the universe of firms with credit ratings provided by Compustat when estimating the model parameters.[18] The MLE parameter estimates that we obtain from the full sample are reported in Table 2 for both the rating-based models (specifications 1 and 2) as well as the model incorporating macroeconomic covariates (specifications 3 and 4). We report standard errors in parentheses. Only the credit rating-implied default intensity is included in the first specification. A coefficient of one would translate into differences in realized defaults across bonds sorted into rating buckets that perfectly matched the difference in default intensities across the rating categories. At the other extreme, a coefficient of zero would indicate a constant probability of default regardless of the credit rating bucket, suggesting the ratings contained no information about credit risk. Thus, the coefficient of 0.872 indicates that there is less spread in realized defaults than suggested by rating-implied default intensities. The inclusion of a frailty component in the second specification leaves this coefficient virtually unchanged.

Interestingly, the estimated standard deviation of the frailty path which is captured by the scaling factor, η, is quite large with a value of 0.203 when using only credit ratings to predict default risk. If mean-reversion is ignored, this translates to a yearly standard deviation of 0.703. An increase of this magnitude in the frailty factor raises each firm's default intensity by a factor of roughly $e^{0.703} \approx 2.02$. This effectively doubles each firm's likelihood of default. To give economic content to this estimate, suppose that a bond currently holds a long-run credit rating of 'Baa1' which maps to an unconditional 5-year probabil-

[15] Azizpour, Giesecke and Schwenkler (2015) find that frailty is still a significant source of default clustering when also considering contagion amongst firms, while Koopman, Lucas and Schwaab (2011) show that frailty represents the combined effects of multiple one-off events not quantified by macro variables and He and Xiong (2012) relate frailty to correlated bond liquidity shocks.

[16] For each credit rating S&P issues, we solve for the value of X which when inserted into (1) as the sole explanatory variable would result in an expected likelihood of default that matches S&P's one-year implied probability of default. Specifically, for a credit rating with an implied one-year probability of default of p, the (monthly) value of the explanatory variable necessary to match implied default probabilities is $X = \ln(-\ln(1-p)/12)$.

[17] We select each covariate from a set of candidates based on their marginal improvement in explanatory power. Other covariates considered and rejected include Gross Domestic Product (GDP) growth, 10-year Treasury yield, Baa-Aaa spread, West Texas Intermediate (WTI) Crude spot price, University of Michigan's consumer sentiment measure, and FRED's U.S. Recession Probability. Shumway (2001) and Campbell, Hilscher and Szilagyi (2008) document important firm-specific default covariates.

[18] While one could restrict the sample to only pieces of collateral used in CLOs, this would lead to less precisely estimated parameters given the already small number of defaults in the full sample. Credit rating agencies (Moody's, 2004) also use the universe of all corporate bonds when developing their methodologies.

Table 2
Corporate debt parameter estimates.
This table reports the results of an exponential hazard model of corporate defaults reported in Compustat from January 1986 to December 2012. *Rating implied intensity* is the implied default intensity from S&P's one-year rating default probabilities. *AAA spread* is difference in AAA corporate debt yields and the 10-year Treasury rate and 3-*Month yield* is the Treasury yield, both reported by FRED. *Unemployment* is the seasonally adjusted U.S. civilian unemployment rate. 12-*Month market return* is the lagged annual CRSP value-weighted return. Frailty volatility is the scaling factor, η, from Eq. (3). Frailty mean-reversion is the speed of mean-reversion, κ, from Eq. (2). Data are at the monthly level; all covariates are measured at month end and used to predict default during the following month. Standard errors are reported in parentheses. *Log-likelihood* is the average loglikelihood across all frailty paths drawn from the Gibbs sampler.

	(1)	(2)	(3)	(4)
Rating implied intensity	0.872	0.860	0.848	0.849
	(0.0174)	(0.0175)	(0.0175)	(0.0177)
AAA spread			0.658	0.526
			(0.1376)	(0.2130)
3-Month yield			0.048	0.078
			(0.0338)	(0.0311)
Unemployment			−0.215	−0.171
			(0.0393)	(0.0472)
12-Month market return			−1.954	−1.944
			(0.2913)	(0.2897)
Frailty volatility, η		0.203		0.142
		(0.0105)		(0.0128)
Frailty mean-reversion, κ		0.020		0.029
		(0.0118)		(0.0242)
Intercept	−2.119	−2.092	0.002	−0.151
	(0.0708)	(0.0717)	(0.6109)	(0.6081)
No obs.	670,465	670,465	670,465	670,465
Log-likelihood	−3453.11	−3348.34	−3330.95	−3263.44

ity of default of roughly 1.10% (Araya, Mahdavi and Ouzidane, 2006). The estimate of η indicates that a one (yearly) standard deviation increase in the frailty component will increase the 5-year default probability to 2.21% which is equivalent to roughly a 1.5 notch decrease in the bond's credit rating.

Furthermore, the frailty process exhibits a relatively slow rate of mean-reversion (κ of 0.020). To illustrate this, consider a first-order auto-regressive process, the discrete analog of the O-U process. Here, if the frailty component of a firm's default intensity receives a shock, 78.7% of the shock will still be present in the default intensity of the firm 12 months later.[19] This indicates that credit ratings either understate or overstate credit risk for extended periods of time. We discuss possible explanations for such persistence in Section 5.2.

The estimation results incorporating only credit ratings are pertinent to examining CBO and CLO credit ratings since rating agencies only use credit ratings to evaluate the riskiness of CBO or CLO collateral pools. However, when only relying on a single predictive factor, a credit rating, one would expect the unexplained systematic risk to play a greater role. Therefore, the third and fourth specifications include a set of macro-explanatory variables of firm default. Their inclusion reduces the role played by unmodeled risk; the scaling factor, η, decreases from 0.203 (in Specification 2) to 0.142 (in Specification 4). The mean-reversion speed, κ, also increases from 0.020 to 0.029 indicating a decrease in the persistence of a shock to the frailty component.

4. Methodology for co-movement in default risk

The first class of models we consider focuses solely on how asset-specific credit ratings co-move together and the implications this has for co-movements in credit risk. When rating a CDO, collateral risk must be inferred from credit ratings at the deal's origination date. Any common change in these ratings in the future constitutes a systematic risk that should be incorporated into the default correlation. Part of Moody's methodology is based on the co-movement of credit ratings, referred to as the directional rating transition matrix (DRTM) approach (Moody's, 2004). This technique estimates asset correlations based on the joint-likelihood of two ratings being upgraded or downgraded together over a one-year horizon. However, it assumes no autocorrelation in the joint-likelihood of directional rating changes across years. Empirically, it is well known that an abnormally large percentage of credit rating downgrades in a given year is often followed by increased rating downgrades in the subsequent year, which would not be captured by the DRTM approach.[20]

Our first set of models addresses this concern, where the simplest model characterizes credit ratings as following one of two transition matrices, corresponding to a *good* and *bad* state. In the *good* state, firms are more likely to have their credit ratings upgraded relative to the unconditional likelihood of an upgrade, and likewise downgraded in the *bad* state. We model the transition between the two regimes using a two-state hidden Markov model.

We begin by ranking all credit ratings in ascending probability of default (*AAA*, *AA*, etc.), and indexing these ratings with ordinal values $1, 2, \ldots, n$, respectively. A rating transition matrix, Π, is a matrix with dimension $n \times n$ where each element π_{ij} is the probability that an asset that has an initial credit rating indexed by i at time t transitions to a rating indexed by j at time $t+1$. The likelihood of a set of rating changes, Z, being generated from a transition matrix Π in a given period is proportional to:

$$\mathcal{L}(\Pi|Z) = \prod_{i=1}^{n}\prod_{j=1}^{n}\pi_{ij}^{N_{ij}} \tag{4}$$

where N_{ij} is the number of credit ratings that transition from state i to state j.

However, our interest is in identifying periods when the likelihood of rating changes is systematically biased upwards or downwards due to changes in the macroeconomy. To do this, the first model we propose combines two rating transition matrices, $\Pi^{(g)}$ and $\Pi^{(b)}$, with a hidden Markov model (HMM) governing the switching between *good* (g) and *bad* (b) states. In our implementation, for each credit rating i we constrain the expected ordinal credit rating

[19] In a discrete setting, the speed of mean-reversion becomes $1 - e^{-0.02}$, thus the AR(1) parameter is $e^{-0.02} = 0.980$ when measured at a one-month interval. Thus, the effect after 12 months is $0.980^{12} = 0.787$.

[20] Nickell, Perraudin and Varotto (2000) relate these time-varying transition probabilities to the business cycle.

in the following period to be lower (corresponding to a higher expected credit quality) in the *good* state than in the *bad* state:

$$\sum_{j=1}^{N}\left[j\cdot\pi_{ij}^{(g)}\right] < \sum_{j=1}^{N}\left[j\cdot\pi_{ij}^{(b)}\right]\ \forall\ i. \tag{5}$$

Following this restriction, the full-information likelihood of a set of rating changes, Z, coinciding with transition matrices $\Pi^{(g)}$ and $\Pi^{(b)}$ and probability p_s of being in state s is proportional to an average of the state-dependent likelihoods from (4), weighted by the probability of the economy being in that state:

$$\mathcal{L}\left(\Pi^{(g)},\Pi^{(b)}|Z\right) = \sum_{s\in\{g,b\}} p_s\cdot\mathcal{L}\left(\Pi^{(s)}|Z\right). \tag{6}$$

Thus, the full estimation procedure requires estimating two state-dependent transition matrices, $\hat{\Pi}^{(g)}$ and $\hat{\Pi}^{(b)}$, probabilities of being in the *good* and *bad* state for each period, and the accompanying 2× 2 transition matrix characterizing the probabilities of switching between the hidden states. Following common practice, we estimate the HMM using the Baum-Welch algorithm as outlined in Appendix C with monthly data on credit rating changes.

While implementing a two-state HMM allows us to model rating changes in *good* and *bad* states, confining the economy to only two states may be overly restrictive if periods frequently occur in which ratings do not exhibit a large degree of upgrades or downgrades. Thus, our second model allows for three state-dependent rating transition matrices, representing *good, moderate*, and *bad* states.[21] Fortunately, the only modification to the two-state framework needed is to extend the constraint in (5) to accommodate the *moderate* state:

$$\sum_{j=1}^{N}\left[j\cdot\pi_{ij}^{(g)}\right] < \sum_{j=1}^{N}\left[j\cdot\pi_{ij}^{(m)}\right] < \sum_{j=1}^{N}\left[j\cdot\pi_{ij}^{(b)}\right]\ \forall\ i. \tag{7}$$

Simply modeling correlated changes in ratings will fail to capture systematic default risk not contained in the ratings, such as business cycle conditions (Nickell, Perraudin and Varotto, 2000). Therefore, we extend these rating-based models to also include macroeconomic covariates, modeling their time-series dynamics using a first-order vector auto-regression AR(1), following Duffie, Saita and Wang (2007). MLE results are reported in Appendix D.

5. Estimation results for corporate bonds

5.1. Rating-based estimation

Table 3 reports the estimation results for the class of HMMs characterizing rating transitions. Panel A reports the difference in the expected monthly change in the credit rating of a firm between the *bad* and *good* states, conditional on the firm's current rating. The table indicates that the expected rating for a firm whose initial rating is 'B,' the most common underlying collateral rating in our sample, is 0.014 greater in the *bad* state of the economy than in the *good* state. Intuitively, this is equivalent to an increased likelihood of being downgraded by one notch to 'CCC' of 1.4% per month in the *bad* state compared to the *good* state. This difference is generally smaller in investment-grade debt with less default risk.

Panel B reports the estimated probabilities of transitions between the hidden states for the two-state HMM. The estimates suggest that the state of the economy, which governs the likelihood of rating upgrades and downgrades, is relatively persistent. Additionally, there is asymmetry in the switching probabilities. Given that the economy is currently in the *good* state in a given month, the probability of switching to the *bad* state in the following month is only 4.3%. However, conditional on currently being in the *bad* state, the likelihood of the economy transitioning to the *good* state in the next month is 14.6%.

However, restricting the economy to two states may exacerbate these switching probabilities if credit rating transitions frequently resemble some intermediate state between the *good* and *bad* states. To examine this possibility, Panel C reports the estimated switching probabilities for the three-state HMM. Interestingly, the switching probabilities conditional on being in the *good* or *bad* state are more symmetric relative to their two-state counterparts reported in Panel B. Additionally, the probability of remaining in the same state for two consecutive periods is relatively constant at approximately 85% across the *good, moderate*, and *bad* states.[22]

5.2. Model frailty and good vs. bad states

We now examine the segmentation of rating changes into the *good* and *bad* states and model frailty over the sample period. Fig. 3 illustrates the mean of the frailty processes drawn from the Gibbs sampler for the ratings-only model, the model using the full set of macro covariates, and the posterior marginal probabilities of the two-state HMM. Note that for each model, we initialize the frailty path to a value of zero in January 1986.

The mean frailty path when incorporating default information from only credit ratings indicates that corporate bond credit ratings tend to overstate the likelihood of default over the time interval from 1989 to 1998 before understating the probability of firm default from 1999 to 2003 in relative terms. Bond ratings are relatively conservative in the period from 2003 to the first half of 2007, in contrast to their performance in the second half of 2007 to the beginning of 2009. Recall that credit ratings are updated monthly in the estimation process. Overall, this indicates that there is considerable variation in an unobservable systematic component of firm default through time that contemporaneous credit ratings do not account for. This variation is not confined to the recent financial crisis, with the frailty path exhibiting substantial variation prior to 2007. However, it is perilous to draw absolute inferences

[21] For robustness, we also propose an alternative model based on the continuous analog to the finite-spaced HMMs used here. The complete model details and estimation procedure are reported in Internet Appendix B, with results reported in the Internet Appendix tables.

[22] The difference in expected monthly rating change between the *bad* and *good* states is quantitatively similar to the two-state HMM reported in Panel A.

Table 3
Rating transition parameter estimates.

This table reports the results of the estimates for the two-state HMM model of rating transitions (Panels A and B) and three-state HMM model (Panel C). Included are all rating changes reported in Compustat from January 1986 to December 2012. Data and estimates are at the monthly level. Bootstrapped standard errors are reported in parentheses.

Panel A: two-state HMM: rating transitions

	AAA	AA	A	BBB	BB	B	CCC	CC
E(Rating$_{t+1}$ | Bad) - E(Rating$_{t+1}$ | Good)	0.007	0.010	0.006	0.005	0.009	0.014	0.021	0.003
	(0.0044)	(0.0026)	(0.0011)	(0.0009)	(0.0013)	(0.0015)	(0.0040)	(0.0188)

Panel B: two-state HMM: state transitions

Current state:	Good	Bad
Prob(Good) at $t+1$	0.957	0.146
	(0.0291)	(0.0291)
Prob(Bad) at $t+1$	0.043	0.854
	(0.0523)	(0.0523)

Panel C: three-state HMM: state transitions

Current state:	Good	Moderate	Bad
Prob(Good) at $t+1$	0.838	0.103	0.018
	(0.0967)	(0.0648)	(0.0496)
Prob(Moderate) at $t+1$	0.122	0.857	0.143
	(0.0835)	(0.0818)	(0.0618)
Prob(Bad) at $t+1$	0.039	0.039	0.839
	(0.0460)	(0.0350)	(0.0609)

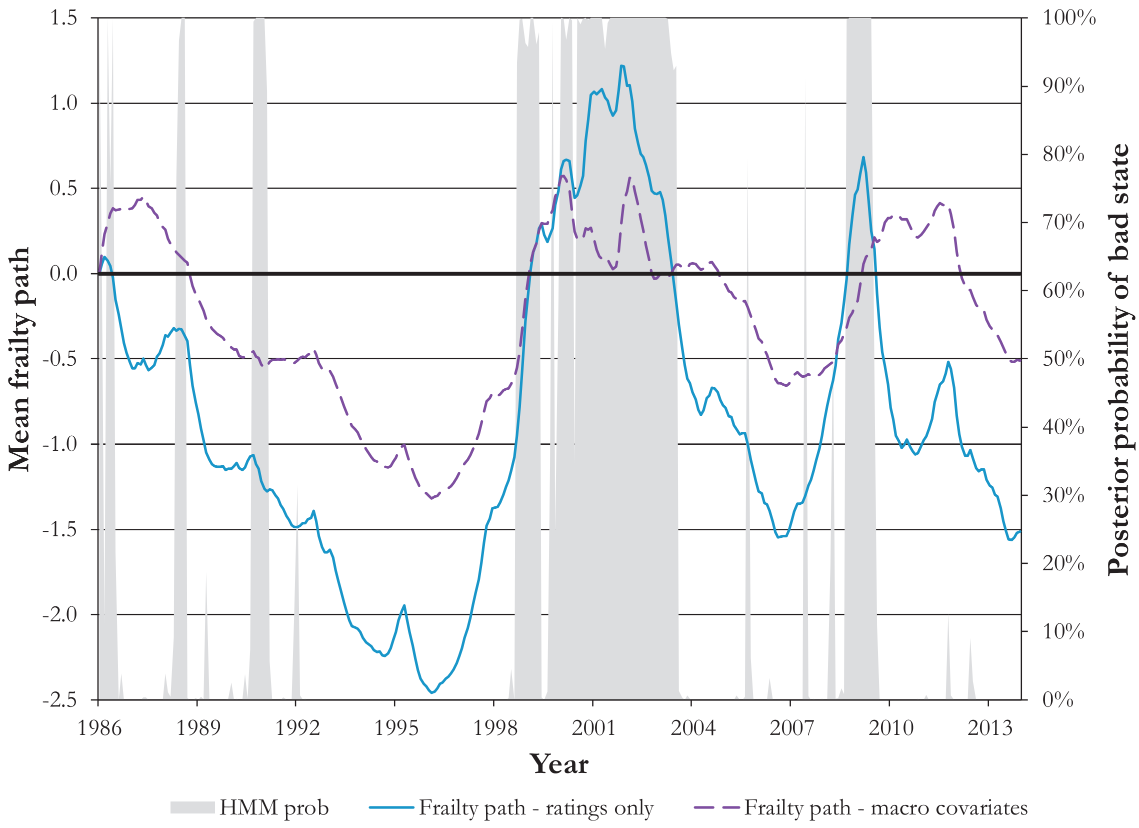

Fig. 3. Estimated frailty and rating transition paths for corporate bonds. This figure illustrates the conditional mean of the frailty path from a hazard model fitted using only the firm's credit rating lagged by one month (solid line) and also including the set of macro covariates reported in Table 2 (dashed line). 4,800 paths were drawn from a Gibbs sampler using the estimated coefficients from the fitted frailty models reported in Table 2. Each path has been scaled by its respective scaling parameter, η. '*HMM prob*' represents the posterior probability of the *bad* state (shaded region) from the two-state HMM. Posterior probabilities were calculated using the Forward–Backward algorithm.

about credit rating standards from this figure since such a comparison also takes into consideration the coefficient on rating-implied default intensity and intercept reported in Table 2.

Fig. 3 also plots the frailty path of the second class of models that incorporate macro-explanatory variables, yielding an interesting set of findings. The first thing of note is that the macro-model's frailty path exhibits less time-variation, confirming the estimated coefficients from Table 2. Furthermore, while the inclusion of macro covariates helps explain some of the predictive shortcomings of ratings, the model frailty present when using macro covariates is correlated with the unmodeled risk not accounted for in credit ratings.

This reduction in model frailty when including macro covariates suggests that credit ratings do not fully capture the credit risk of firms. One possible driver of such persistence in frailty when incorporating only the information in credit ratings is the tendency for ratings to become 'stale' over time (Löffler, 2005; Hand, Holthausen and Leftwich, 1992). Another explanation consistent with a credit rating that understates risk during periods of economic downturns is the tendency of agencies to rate "through-the-cycle" (Standard & Poor's, 2003). When contrasting the estimated frailty path against the business cycle, the figure suggests that ratings understate risk in economic downturns and overstate risk in expansionary periods, consistent with agencies rating through-the-cycle. Thus, the true default risk of a CDO's collateral pool can be systematically understated or overstated depending on the current economic conditions, resulting in correlated default risk of the underlying assets.

We now turn to an examination of systematic rating changes. To the extent that credit ratings of the collateral pool change in a systematic fashion, this adds another potential source of tail risk to the tranche holders. This risk is confirmed when examining the posterior marginal probabilities of the *bad* state of the economy (the shaded area in Fig. 3). Fig. 3 indicates that there are extended periods of time when credit ratings are more likely to be downgraded or upgraded in a systematic fashion. Interestingly, the positive correlation between the mean frailty path and the probability of being in the *bad* state indicates that in times when a credit rating is more likely to be downgraded, the harsher rating is also more likely to understate the firm's default risk.

Overall, the results in this section suggest both model frailty and systematic changes in observable default risk factors possess considerable time-series dynamics which lead to periods where risk is understated or overstated.

5.3. Default correlation in corporate bonds

We now use the results from the previous subsection to estimate the default correlation present in a CBO's pool of corporate bonds. Intuitively, we use the estimated time-series dynamics of the observable covariates and frailty to repeatedly simulate possible cross-sections of hazard rates and thus defaults of a CBO's collateral pool.

For each model and CBO, we begin with the initial credit ratings of the structured product's collateral pool. We then draw 100,000 paths of the frailty process with length equal to the structured product's weighted-average maturity. For each model, we also simulate hidden states of the economy, their associated rating transition matrices, and time-series of macro covariates when considering appropriate models.[23] Using these transition matrices, we next simulate credit rating changes for each asset over the CDO's life. To incorporate the effects of frailty and observable covariates, we first take the resulting draws of credit ratings, frailty, and macro covariate time-series and compute each asset's probability of default using Eq. (3) and the MLE parameter estimates obtained. From these default probabilities we simulate realizations of default within the pool. The result is 100,000 cross-sections of asset defaults for a given CBO's collateral pool. It should be noted that up to this point, our approach only differs from that of S&P in one dimension. We simulate cross-sections of default from a time-varying default intensity that incorporates frailty while the rating agency generates them from a Gaussian Copula. From these default realizations, we calculate the average pairwise default correlation of the collateral pool. Specifically, when computing the average we value-weight each asset-pair's default correlation, ρ_{ij}, by the product of the two assets' sizes, w_i and w_j, respectively.

Table 4 reports the average default correlation across our sample of CBOs issued prior to 2007 for each model. We begin by restricting ourselves to the data available to rating agencies prior to the financial crisis. Thus, Panel A reports the results when estimating model parameters using data observable from January 1986 to December 2006, and then using these parameter estimates to simulate realizations of defaults and default correlations. To begin, we focus on the default correlation attributable to rating transitions alone in the two-state HMM. The average pairwise default correlation for our sample is 0.0017 when considering only a *good* and *bad* state transition matrix and increases marginally to 0.0019 under the three-state HMM. These results suggest that the risk of common default across multiple firms appears to be slight when examining only co-movement in credit ratings.

The previous specifications assume that credit ratings perfectly reflect default risk at all times. We now turn our attention to frailty, or the default risk not captured by credit ratings. The second column of Table 4 reports the pairwise default correlations for our sample of CBOs when subject to the common frailty factor. The results are striking. The default correlation due to frailty (*Unmodeled risk*) is 0.0856 for our sample. This value is close to a full order of magnitude larger than either the default correlation due to *Modeled risk factors* or the average default correlation used by S&P prior to the financial crisis. Note, the default correlation due to *Unmodeled risk* is identical under the first three models because each model uses the same set of covariates in predicting firm default.

Finally, we consider the joint effect when accounting for both co-movement in ratings and frailty in the *Both* column. The average default correlation for our sample

[23] Note, we do not model the correlation across determinants of credit risk, instead drawing frailty paths, hidden states, and macro covariates independently of each other.

Table 4
Corporate bond default correlations.
This table reports the average pairwise default correlation for a sample of corporate-backed CDOs issued from 2000 to 2007. Reported are the results obtained when estimating each model using data from January 1986 to December 2006 (Panel A) and from January 1986 to December 2012 (Panel B). Modeled risk factors denotes the default correlation resulting from co-movement in observable risk factors considered in each model specification. Unmodeled risk denotes the default correlation resulting from systematic, unmodeled risk. Both denotes the default correlation when considering both co-movement in modeled risk factors and systematic unmodeled risk. The default correlation for each deal is calculated as the value-weighted pairwise default correlation between the underlying assets, where each pairwise weight is equal to the product of the two assets' sizes. Bootstrapped standard errors are reported in parentheses.

Methodology	Modeled risk factors	Unmodeled risk	Both
Panel A: pre-2007 sample correlation estimates			
Ratings only:			
2-State HMM	0.0017	0.0856	0.1019
	(0.0002)	(0.0076)	(0.0075)
3-State HMM	0.0019	0.0856	0.1016
	(0.0002)	(0.0076)	(0.0073)
Ratings & macro covariates:			
2-State HMM	0.0397	0.0335	0.0778
	(0.0069)	(0.0091)	(0.0117)
3-State HMM	0.0412	0.0335	0.0791
	(0.0069)	(0.0091)	(0.0117)
Panel B: full sample correlation estimates			
Ratings only:			
2-State HMM	0.0013	0.1117	0.1218
	(0.0001)	(0.0074)	(0.0071)
3-State HMM	0.0016	0.1117	0.1225
	(0.0002)	(0.0074)	(0.0071)
Ratings & macro covariates:			
2-State HMM	0.0314	0.0421	0.0749
	(0.0060)	(0.0071)	(0.0122)
3-State HMM	0.0315	0.0421	0.0749
	(0.0060)	(0.0071)	(0.0122)

increases relative to either individual component of default risk and is 0.1019 using the two-state HMM and 0.1016 with the three-state model.

The second class of models incorporates a set of macro-explanatory variables able to explain firm default risk. The inclusion of these factors increases the default correlation due to observables to between 0.0397 for the two-state HMM and 0.0412 under the three-state HMM. As expected, the addition of more observable covariates decreases the role played by unmodeled risk, resulting in a default correlation due to frailty alone of 0.0335. Interestingly, when considering the combined effect due to changes in observables and frailty, the correlated default risk falls between 0.0778 (two-state HMM) and 0.0791 (three-state HMM), estimates which fall short of the models only incorporating information in rating change co-movements and frailty.

However, we view these default correlation estimates as a lower bound on the true values. In our estimation procedure, we independently draw states of the world, realizations of the macro covariates, and frailty paths for tractability. A positive correlation in these factors driving default would lead to increased loadings on the included variables if a subset of the variables was omitted (e.g., omitted-variable bias). This bias in coefficients would partially account for the positive correlation between the included and omitted variables. However, the inclusion of all the variables combined with a failure to model the positive correlation between the factors in the simulation would result in a downward bias in the estimated default correlation.[24]

To gauge the impact of information revealed through the financial crisis, we re-estimate our model using the full sample. Panel B reports default correlations when using parameter estimates from the sample ending in December 2012 to simulate default. The average pairwise default correlation for the ratings-based models with frailty is roughly 0.125, or 25% larger relative to their pre-crisis sample counterparts. Such a change in estimated default correlations may appear small relative to the severity of the recent financial crisis. However, recall that the frailty paths illustrated in Fig. 3 exhibited a considerable amount of variation prior to 2007, first decreasing from 1988 to 1995, then dramatically increasing from 1997 to 2002 before again decreasing through 2006. Thus, it is plausible that default correlations which incorporate frailty estimates based on pre-crisis data do not differ substantially from those estimated over the full sample.

Overall, these findings suggest that default correlations due to co-movement in observable risk factors and model frailty can add considerable thickness to the tails of CBO and CLO collateral pool losses.

5.4. Default correlations across asset maturities

As mentioned in Section 2.1, correlations based on realizations of default are also a function of the underlying assets' default probabilities. This implies that these default correlations are also influenced by the maturity of the underlying assets when keeping collateral quality constant.

To illustrate this, consider the following scenario. Suppose that two assets are exposed to a risk factor, x_t, which is the sole determinant of their default intensity $\lambda_t = e^{x_t}$. Additionally, assume that x_t follows a random walk with a standard deviation of monthly innovations equal to our full-sample parameter estimates, $\sigma = 0.203$. Using this simple framework, we estimate the correlation of realized defaults between two assets for 100,000 draws of the process x for maturity values that fall within the 5th and 95th percentiles of WAMs used by credit rating agencies for our sample of corporate-backed CDOs.[25] Fig. 4 plots how the resulting estimate of default correlation varies across the range of weighted-average maturities. The figure illustrates that the estimated default correlation increases with the

[24] Confirming this intuition, we find a positive correlation among the three components, consistent with decreased default correlation estimates when the states of the world and macro covariates are not jointly drawn. Thus, modeling the factors independently should understate the correlated default risk, producing a lower bound.

[25] Note, we calibrate the model such that the 5-year expected default probability matches the average default probability of collateral in our sample.

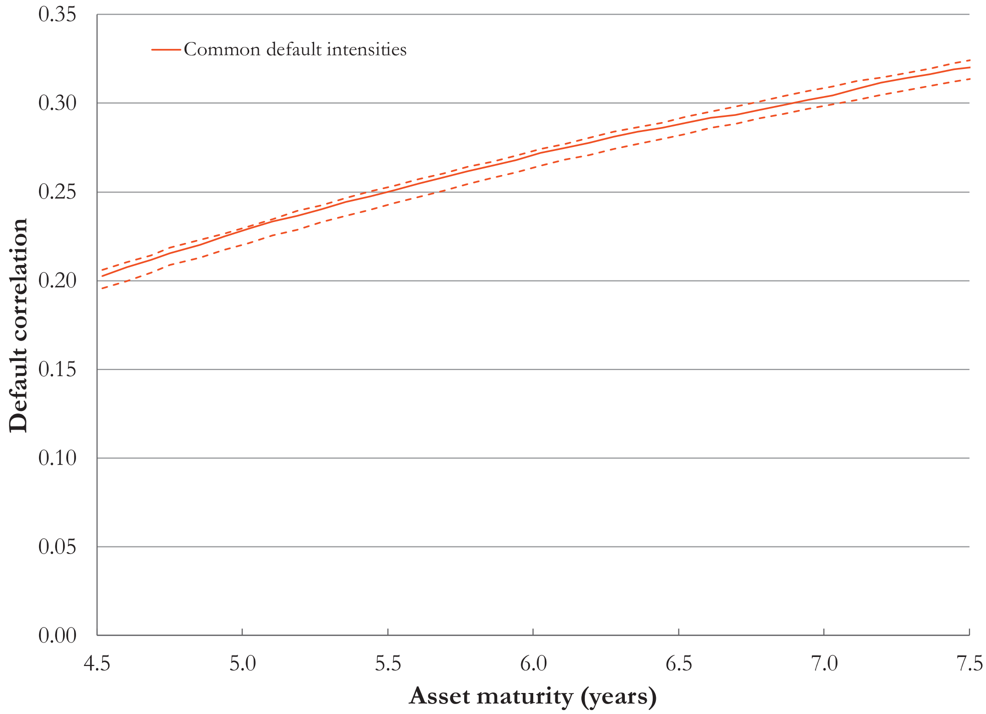

Fig. 4. WAM dependent default correlations. This figure plots the default correlation between two assets when exposed to a common log-normal default intensity which is time varying with a standard deviation of monthly innovations of 0.203. The model is calibrated such that the 5-year expected default probability matches the average default probability of collateral in our sample. Reported is the estimated default correlation for asset maturities ranging from 4.5 to 7.5 years, approximately equal to the 5th and 95th percentiles of the deal WAMs in our sample, respectively. Dashed lines indicate 95% confidence bands, generated from 100 bootstrapped simulations.

maturity of the assets over the range of WAMs common in our sample. This confirms that correlations attenuate over shorter horizons. This conditional dependence raises the concern that our methodology may have low power (high Type II error) to detect a difference relative to rating agency assumptions for CLOs with a short WAM. As a diagnostic test of our methodology's power, we repeat the estimation procedure 100 times and plot the 95% confidence interval for these estimates (dashed lines). If our test lacked precision, with little power to distinguish between differing default correlation values, the resulting estimates would demonstrate a wide range of values. In contrast, the estimated default correlation exhibits relatively tight confidence bands, suggesting our test possesses sufficient power to detect differences in default correlation estimates.

In addition, the need to estimate model parameters which are measured with error raises the concern of a Type I error, in which we falsely reject the null that the difference between our estimated default correlation and rating agency assumptions is zero. To alleviate this concern, throughout the paper we report bootstrapped standard errors for each parameter estimate and default correlation estimate to ensure the differences between our estimates and those of the rating agencies are not being driven by misestimated parameters.

Nonetheless, given the effect of asset maturity on default correlation estimates, we now re-examine the sensitivity of our previous results to the WAM of the underlying collateral pool. Recall that we estimate the average pairwise default correlation for each CDO collateral pool using the deal's observed WAM. This same WAM is used by rating agencies when evaluating the pool's collateral risk, allowing us to effectively compare our estimate with that of the CRAs within a matched sample. Thus, we would expect that both the rating agencies' assumed default correlations as well as our estimates increase with WAM.

For this reason, Table 5 partitions our sample of CDOs into terciles based on the collateral pool WAMs and reports the average default correlation estimate for each subset.[26] The table indicates that the default correlation assumed by S&P is monotonically increasing in the collateral's WAM, increasing from 0.0071 to 0.0128. In contrast, Moody's assumed default correlations only exhibit an increase in the subset of deals with long maturities.[27] Furthermore, the estimated default correlation from each of the modeling frameworks exhibits a monotonically increasing relation. The estimated values are generally larger than S&P's assumed default correlations by a factor of five. This set of results confirms the conditional nature of default correlations with respect to the underlying assets' default probabilities and maturities while highlighting the robustness of

[26] The average WAMs for the terciles are 4.9 years, 5.6 years, and 7.27 years. The specific cutoffs are reported in the table header.

[27] This is partly explained by the methodology used by Moody's to capture correlation in collateral defaults, which is based on the pool's industry concentration but not the maturity of the assets or simulated realizations of default.

Table 5
Corporate bond default correlations by WAM.
This table reports the average pairwise default correlation for a sample of corporate-backed CDOs issued from 2000 to 2007. Reported are the results obtained when estimating the model which considered both co-movement in modeled risk factors and systematic unmodeled risk using data from January 1986 to December 2006. The primary difference between this table and Panel A of Table 4 is that the sample has been split into terciles based on the underlying collateral's weighted-average maturity to compare estimates within maturity bins. The WAM for CDOs in the Short bin range from 3.5 to 5.38 years, in the Medium bin range from 5.39 to 5.93 years, and in the Long bin from 5.94 to 12.4 years. These maturities are obtained from the stated maturities issued by S&P. Bootstrapped standard errors are reported in parentheses.

	Deal WAM		
Methodology	Short	Medium	Long
CRA assumptions:			
S&P default correlation	0.0071	0.0092	0.0128
Moody's default correlation	0.0104	0.0095	0.0122
Ratings only:			
2-State HMM	0.0822	0.1034	0.1201
	(0.0099)	(0.0109)	(0.0111)
3-State HMM	0.0822	0.1031	0.1197
	(0.0091)	(0.0105)	(0.0111)
Ratings & macro covariates:			
2-State HMM	0.0644	0.0798	0.0891
	(0.0155)	(0.0157)	(0.0157)
3-State HMM	0.0656	0.0811	0.0906
	(0.0156)	(0.0157)	(0.0156)

our findings. These results also highlight the importance of our matched-sample approach which uses the same deal-specific weighted-average maturities that the rating agencies assume.

6. Implications for current ratings

Default correlations are pivotal in the rating of structured finance products that make up a large part of credit rating business, and rating agencies have considerable data and resources at their disposal. In the aftermath of the financial crisis and the substantial negative publicity received, rating agencies made a considerable number of press releases regarding methodological adjustments to make their modeling more robust. The comparison of rating agency assumptions after the crisis and our default correlation estimates over the full sample gives a sense of any additional risk possibly absent from the evaluation of post-crisis structured finance products.

Using surveillance reports, we obtain data for 136 CLOs from S&P issued from June 2011 to June 2014. The AAA tranches make up 61.9% of the deal on average, while tranches AA and higher are 72.9%. The average collateral pool in the sample is made up of 179 obligors, suggesting there is potential for considerable diversification. Panel A of Fig. 5 reports the total par amount of corporate bonds and loans by industry which make up the collateral pools. The panel indicates that there is significantly more corporate debt from the 'Healthcare' sector ($16.8B) than that of the second largest industry, 'Electronics' ($10.7B). Given this over-concentration in the healthcare industry, we next examine the industry concentration within each CLO. Specifically, for each deal we compute the Herfindahl index (HHI) of the percent of the collateral pool represented by each industry classification. Panel B of Fig. 5 reports the histogram of HHIs for the CLOs in our sample. The majority of the CLOs appear to be relatively diverse with an HHI below 0.07. In contrast, the HHI for the total par amount by industry from Panel A is 0.05, suggesting the overall sample is fairly representative of each CLOs industry concentration.[28] Finally, we examine the homogeneity in credit quality for each collateral pool. For each deal, we compute the HHI of the percent of the collateral pool made up of each broad credit rating (*BBB*, *BB*, ...). Panel C reports the histogram of credit rating HHIs for the CLOs in our sample. The panel suggests that the collateral in each pool is relatively homogeneous; the majority of the CLOs have a rating HHI greater than 0.50.[29] The most common ratings are B, BB, and CCC with 75.2%, 22.1%, and 1.62% of collateral, respectively.

For this sample S&P assumes an average default correlation of 0.033 with a range of 0.012 to 0.060. Thus, S&P has increased their correlation estimates in the aftermath of the crisis. Panel A of Table 6 reports the average default correlation estimated with our approach outlined in Section 5.3 above. The estimated default correlations generated across our methodologies are three to four times greater, ranging from 0.103 when modeling default risk using a combination of macroeconomic covariates and a two-state HMM to 0.124 when modeling ratings with the three-state HMM but excluding macro covariates. The lower assumptions being used by credit rating agencies indicate that tranches given AAA ratings by rating agencies are likely exposed to more risk than their current credit rating merits.

However, this in itself does not give a sense of the economic magnitude embedded in the differential between our estimates of default correlation and those used by rating agencies. S&P uses the scenario default rate (SDR), which is equivalent to a value-at-risk, to measure the losses in the tail of the distribution.

Fig. 6 reports the scatter-plot of S&P's AAA SDR levels relative to SDR estimates under our modeling framework. Reported for convenience is the 45-degree line, which segments the sample into deals in which our methodology yields a higher level of collateral risk (above the line), and deals in which S&P estimates the collateral as being riskier (below the line). The figure reports the estimated SDR when considering the risk of rating co-movements, modeled with a two-state HMM, and model frailty. While there is a relationship between the AAA SDR reported by S&P and our estimates (correlation of 0.356), every observation lies above the 45-degree line. Confidence intervals are constructed from bootstrapped parameter estimates in a similar fashion to the standard errors of Table 6. This indicates that for every CDO considered, evaluating the collateral under our methodology leads to a higher level of extreme collateral default risk relative to S&P.

[28] Additionally, the HHI of a perfectly diversified collateral pool of the 34 industries is 0.03.

[29] The average HHI of credit ratings is 0.49 when using the full credit rating scale (*BBB+*, *BBB*, ...).

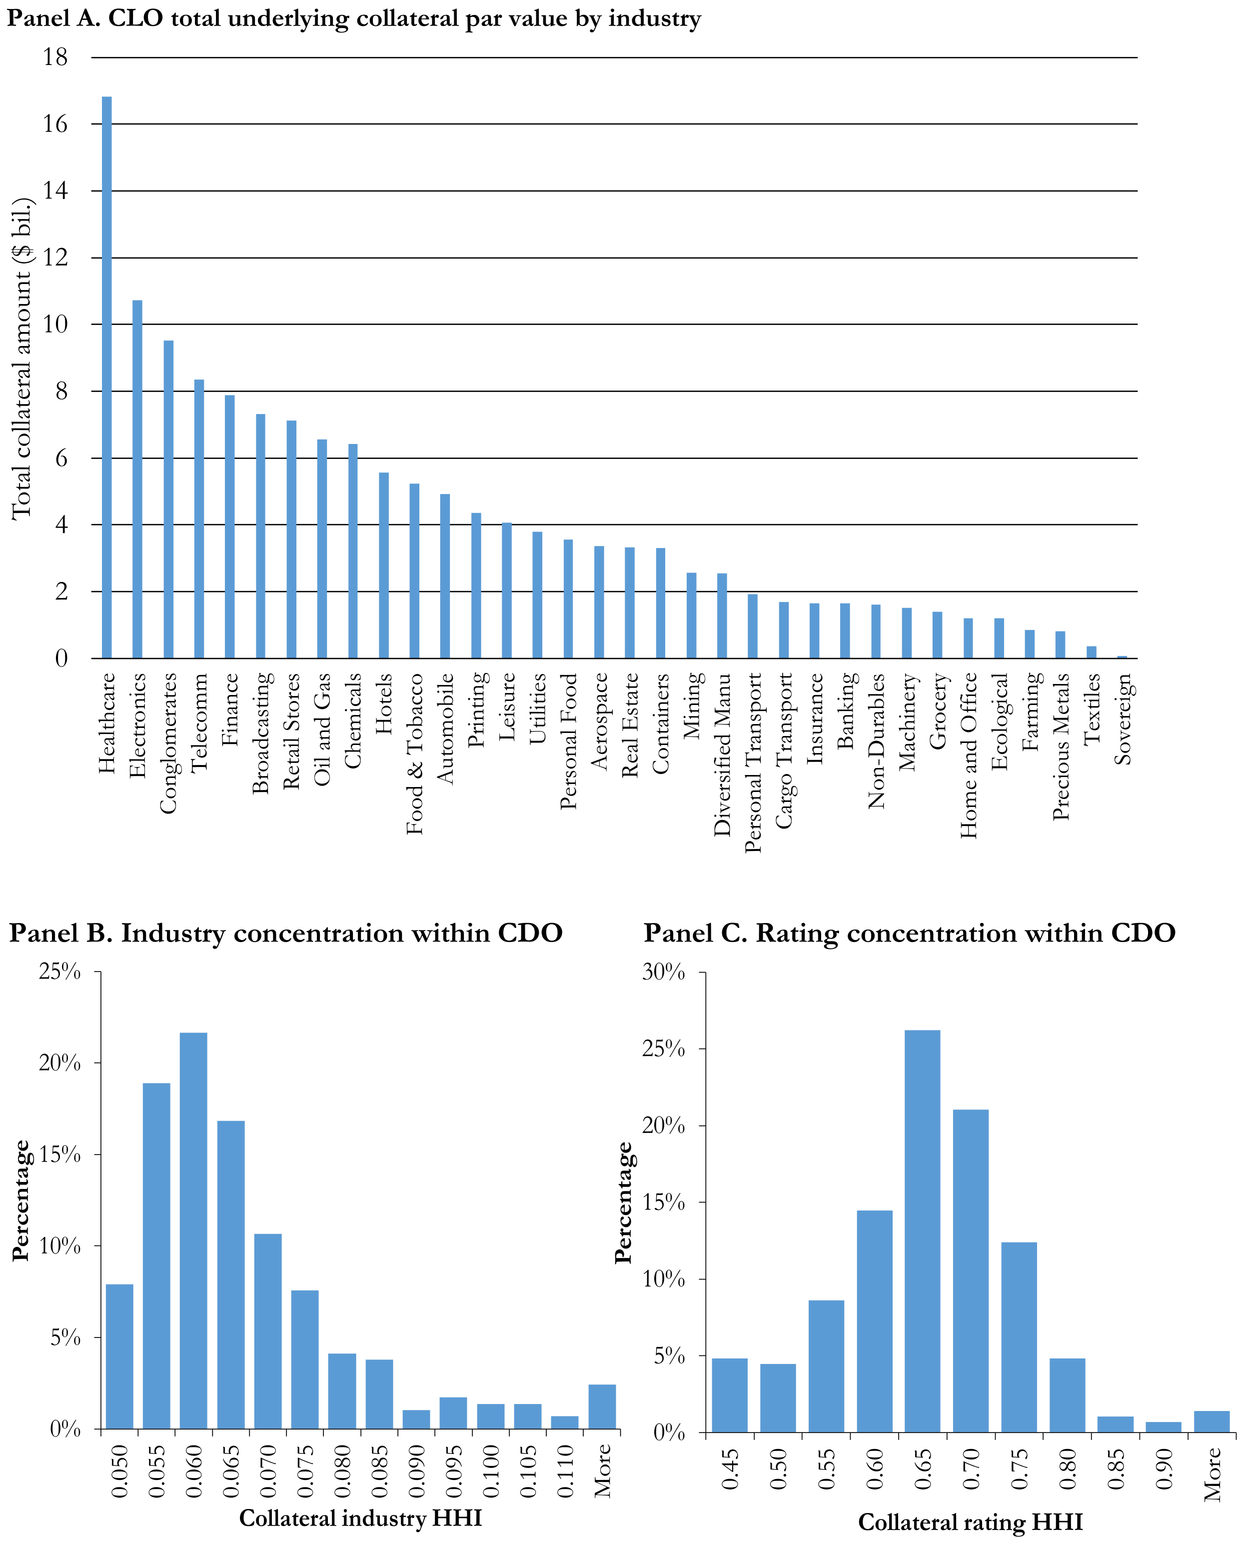

Fig. 5. CBO underlying collateral composition by industry and rating. This figure reports the total par amount of underlying collateral by industry (Panel A), the industry concentration by deal (Panel B), and rating concentration by deal (Panel C) for a sample of 136 CDOs recently issued. Panel A plots the total par amount by industry, as classified in the surveillance reports. Panel B plots the histogram of Herfindahl-Hirschman Indices (HHI) with respect to the industry composition making up the underlying collateral pool for each CDO in the sample. Panel C plots the histogram of HHI with respect to the rating composition making up the underlying collateral pool for each CDO in the sample. The sample consists of collateralized loan obligations (CLOs) and collateralized bond obligations (CBOs) issued from June 2011 to December 2013.

Table 6
Post-crisis portfolio default risk estimates.

This table reports the average pairwise default correlation (Panel A) and scenario default rate (Panel B) for a sample of 136 corporate-backed CDOs issued from 2011 to 2014. *Modeled risk factors* denotes the default correlation resulting from co-movement in observable risk factors considered in each model specification. *Unmodeled risk* denotes the default correlation resulting from systematic, unmodeled risk. Both denotes the default correlation when considering both co-movement in modeled risk factors and systematic unmodeled risk. The default correlation for each deal is calculated as the value-weighted pairwise default correlation between the underlying assets. Panel B reports summary statistics for AAA SDRs estimated by S&P, under our methodology when modeling rating changes and frailty (*Ratings only*) and under our methodology when modeling rating changes, macroeconomic covariates, and model frailty (*Ratings & macro covariates*). Bootstrapped standard errors are reported in parentheses.

Panel A: pairwise default correlations

Methodology	Modeled risk factors	Unmodeled risk	Both
Ratings only:			
2-State HMM	0.0014	0.1126	0.1236
	(0.0002)	(0.0109)	(0.0102)
3-State HMM	0.0017	0.1126	0.1240
	(0.0003)	(0.0109)	(0.0102)
Ratings & macro covariates:			
2-State HMM	0.0317	0.0420	0.1030
	(0.0083)	(0.0107)	(0.0153)
3-State HMM	0.0319	0.0420	0.1027
	(0.0083)	(0.0107)	(0.0155)

Panel B: portfolio SDRs

Methodology	Mean	Median	Std. dev.
CRA assumptions:			
S&P's SDRs	0.661	0.660	0.028
Ratings only:			
2-State HMM	0.835	0.841	0.030
	(0.022)	(0.022)	
3-State HMM	0.833	0.836	0.030
	(0.022)	(0.023)	
Ratings & macro covariates:			
2-State HMM	0.698	0.703	0.032
	(0.029)	(0.030)	
3-State HMM	0.701	0.704	0.030
	(0.030)	(0.030)	

Panel B of Table 6 summarizes the results of the estimated SDRs under each model we consider along with those assumed by S&P. The average SDR reported by S&P for the sample is 66.1% (median of 66.0%). Recall that the model which includes macro covariates is not able to account for a positive correlation between the default intensity implied by macro covariates and frailty, whereas the correlation between these two risk factors is 0.151 in our sample. Thus, we think that the model specification with rating co-movement and frailty is the most appropriate. The average SDR under models which consider only rating co-movement and frailty range from 83.3 to 83.5%; amounting to at least a 17.2% movement in SDR or a 26% (83.3/66.1 - 1) relative increase in the credit risk.

7. Economic importance and relevance

A final point of comparison should be made. While credit rating agencies like S&P separate the modeling of risk into credit and cash-flow components, the correlation assumptions we are examining only affect the credit risk of the CLOs. Griffin and Tang (2012) compare the difference between the amount of AAA issued and the amount implied from S&P's credit risk model (1-SDR) as a measure of the aggressiveness of internal cash-flow modeling.

Griffin, Nickerson and Tang (2013) document that both Moody's and S&P respond to competitive pressure resulting from the other's favorable credit risk modeling by issuing positive cash-flow adjustments. It is possible that the use of higher default correlation assumptions could be offset by larger cash-flow adjustments. However, the increase in SDR of 17% that we document here due to a more conservative correlation assumption is much larger than the adjustments in response to competitive pressure of 4–8% documented in Griffin, Nickerson and Tang (2013). Thus, a rating agency may find it difficult to relax its cash-flow modeling criteria by enough to offset the incorporation of frailty in its correlation assumptions. Hence, it is possible that rating agencies have not adopted more aggressive correlation assumptions due to concerns regarding market share. Such concerns are consistent with prior findings regarding business considerations modeling choices. For example, the Statement of Facts released in the Department of Justice (DOJ) settlement with S&P states that S&P began testing a default matrix proposed by the heads of groups, "in part based upon business decisions, and considerations."[30] Additionally, Moody's (2004) considers two methods for calculating default correlations but ultimately chooses the method yielding the lower estimates.

Another possibility is that cash-flow modeling has become more conservative post-crisis to the point of offsetting the increase in SDR under larger default correlations. This does not seem to be the case. For our sample from 2011 to 2014, the average CLO has 61.9% of the deal rated AAA, whereas S&P's credit risk model alone generates only 33.9%.[31] Thus, the AAA amount issued is 28% above the amount implied from S&P's credit risk model. This is considerably higher than the 16% for CLOs found by Griffin and Tang (2012) from 1997 to 2007. Unless there has been a substantial change in the waterfall structure of CLOs, this could indicate that S&P is currently using even more aggressive cash-flow modeling practices and provides even more caution regarding structured finance CLOs.

While our findings suggest that rating agencies' models may be understating senior tranche credit risk, it is interesting to observe whether market prices reflect higher risk or are pricing AAA CDOs similar to corporate AAA debt. Fig. 7 plots the average quarterly spread of CBO AAA tranche yields above the London Interbank Offered Rate (LIBOR). The figure indicates that spreads narrowed considerably from 2004 to 2007 at the same time that the estimated frailty path was decreasing and there were few defaults. Spreads hovered around ten basis points in 2006, while the average yield spread from 2011 through the first quarter of 2014 is between 100 and 150 basis points. These yields suggest that the marginal post-crisis investors gener-

[30] http://www.justice.gov/file/338706/download

[31] The average CLO has 72.6% AAA in the sample of Griffin and Tang (2012).

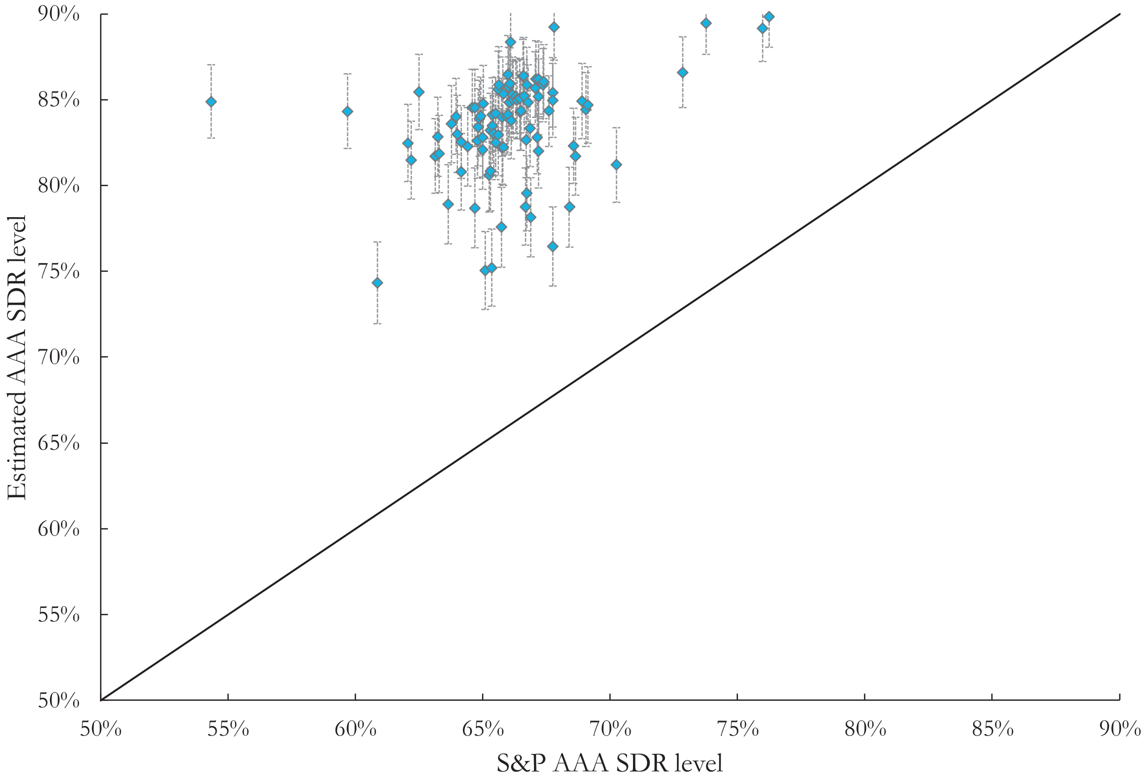

Fig. 6. Comparison of recent AAA scenario default rates. This figure reports the scatter-plot of AAA scenario default rates (SDRs) used by S&P (*x*-axis) and calculated from our methodology (*y*-axis). The figure illustrates the estimated SDR under our methodology when modeling only rating changes under a two-state HMM and including model frailty. Dashed line segments denote 95% confidence intervals estimated from bootstrapped parameter estimates. The sample consists of collateralized loan obligations (CLOs) and collateralized bond obligations (CBOs) issued from June 2011 to December 2013.

ally believe AAA CLOs' risk to be considerably greater than their rating. This lends additional credence to our conclusion that '2.0' ratings are still too aggressive.

One might be tempted to conclude that inaccurate ratings are unimportant as long as the marginal price incorporates additional credit risk. Yet, there are some problems with this line of reasoning. First, the marginal investor could still be an investor who does not understand the risk and misprices the security.[32] Second, accurate ratings are still important to regulators and third-party investors who, for example, might invest in a AAA money market fund and rely on the credit ratings. Third, inaccurate ratings can facilitate the extent to which institutional investors "reach for yield" (Becker and Ivashina, 2015) or cause excessive risk-taking (Becker and Opp, 2013).

8. Conclusion

A commonly assumed lesson from the financial crisis is that default correlations were not well understood. Despite this period of massive default, almost no academic work has been done to understand how we should assess default correlations for structured products. Our findings point to the importance of incorporating default risk due to both systematic changes in observable risk factors and frailty. Even when estimating our model using pre-crisis data, the correlations used by rating agencies for CLOs were considerably lower than those we obtain. Thus, the low default correlations used prior to the crisis are not solely attributable to a lack of crises in the data, but are also a result of methodological choices. Recent CLO issuances reinforce this conclusion. Post-crisis CLO correlation estimates do not appear to incorporate frailty and appear to be structured too aggressively. If frailty were incorporated, our methodology suggests that credit risk on AAA tranches may be understated by 26%.

Our default correlation estimates that include the financial crisis indicate substantial diversification benefits across structured finance instruments, hence indicating a role for structured finance in the future. Nevertheless, we must also temper this enthusiasm by noting that credit cycles can create waves of similar financial instruments in the precise time periods where frailty correlations could be increased through amplification mechanisms such as those noted by Brunnermeier (2009). Thus, while our findings indicate that the waves of newly issued structured products may not be as creditworthy as their ratings suggest, there are potentially other reasons for caution as well. Future work should focus on further understanding default correlations as well as other aspects of structured finance modeling, such as incorporating parameter uncertainty. Without such research and more transparency in both detailed

[32] More extensive modeling with more detailed data on particular deals would need to be performed to determine if the 100–150% basis points is sufficient for the additional frailty risk.

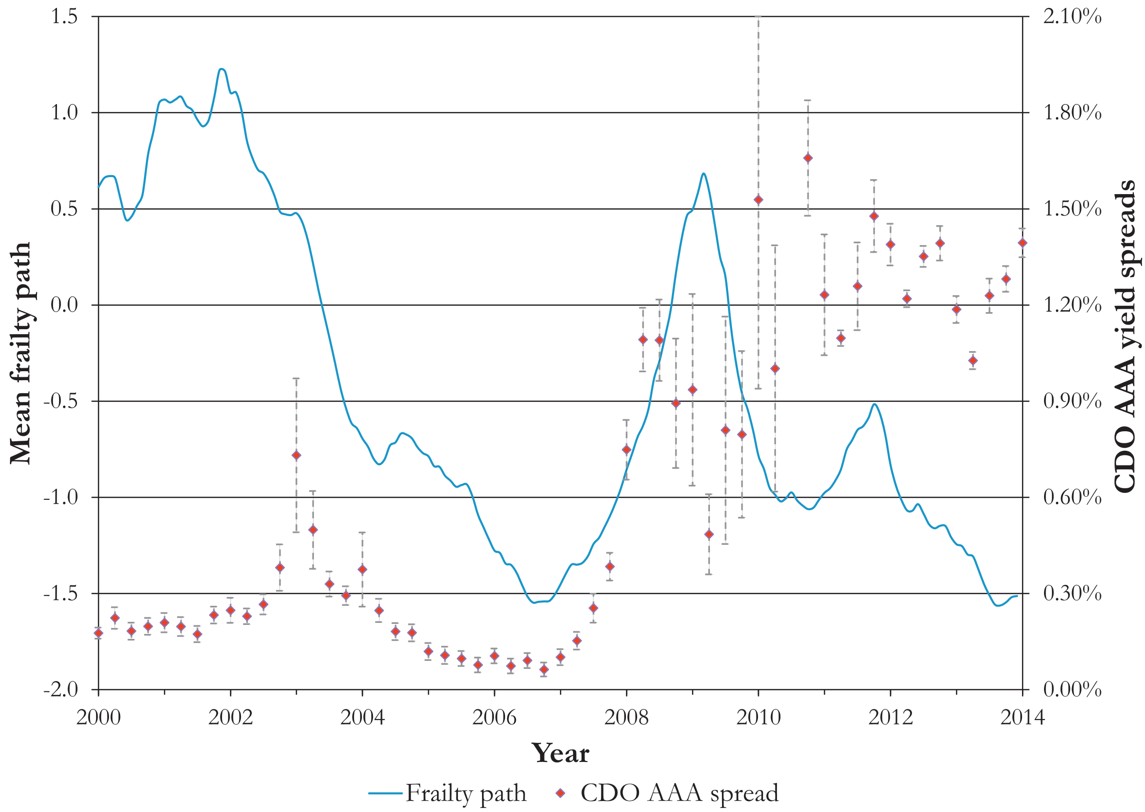

Fig. 7. Estimated frailty path and CDO yield spreads. This figure illustrates the conditional mean of the frailty path from a hazard model fitted using the firm's credit rating lagged by one month (solid line) and the mean yield spread by quarter (diamonds). The frailty path has been scaled by the appropriate scaling parameter, η. Reported is the average yield spread for the junior-most tranche rated AAA at issuance across all CLOs and CBOs by quarter. A one standard deviation range is denoted by the dashed line segments.

data and assumptions by credit rating agencies, structured finance ratings should be viewed cautiously.

Appendix A

This appendix details the closed-form solution which maps S&P and Moody's proprietary correlation metric to default correlation. A previous version of these derivations can be found in Griffin, Nickerson and Tang (2013). Both derivations are based on the variance of the underlying collateral pool's realized default distribution, σ_p^2. This variance can be decomposed into asset default variances and asset-pair covariances as follows:

$$\sigma_p^2 = \sum_{i=1}^{N} w_i^2\sigma_i^2 + 2\times\sum_{i=1}^{N}\sum_{j=i+1}^{N} w_i w_j \sigma_{ij} = \sum_{i=1}^{N} w_i^2\sigma_i^2 + 2\times\sum_{i=1}^{N}\sum_{j=i+1}^{N} w_i w_j\sqrt{\sigma_i^2\sigma_j^2}\rho_{ij} \tag{8}$$

where σ_i^2 is the variance of the realization of default for asset i and w_i is the percentage of asset i's size relative to the total asset pool. Assuming assets are of equal size w, are of the same collateral quality and thus have a common variance of default σ^2, and share a common pairwise correlation ρ, the equation simplifies to:

$$\sigma_p^2 = \frac{1}{N}\sigma^2 + \frac{N-1}{N}\sigma^2\rho = \left[\frac{1}{N} + \frac{N-1}{N}\rho\right]\sigma^2. \tag{9}$$

Moody's proprietary metric, the Diversity Score, is defined as the number of independent assets the collateral pool can be represented by. Thus, for a pool with a diversity score of DS:

$$\sigma_p^2 = \frac{1}{DS}\sigma^2. \tag{10}$$

Therefore, to match the second moment the average pairwise correlation must satisfy the following:

$$\begin{aligned}\frac{1}{DS}\sigma^2 &= \left[\frac{1}{N} + \frac{N-1}{N}\rho\right]\sigma^2\\ \rho &= \frac{N-DS}{DS(N-1)}.\end{aligned} \tag{11}$$

In contrast, S&P's correlation metric, CM, is defined as the ratio of the standard deviation of the portfolio default distribution with correlation divided by the standard deviation of the distribution when ignoring correlation between asset defaults:

$$\begin{aligned}CM^2 &= \frac{\frac{1}{N}\sigma^2 + \frac{N-1}{N}\sigma^2\rho}{\frac{1}{N}\sigma^2}\\ \rho &= \frac{CM^2-1}{N-1}.\end{aligned} \tag{12}$$

Appendix B

This appendix provides a detailed example of S&P's computation of the Scenario Default Rate (SDR) for a collateral pool, which is equivalent to the value-at-risk of the collateral pool's default distribution. Suppose the collateral pool of a CLO is made up of 100 assets, with identical default probabilities $p_i = 0.15\ \forall\, i$ and relative sizes $s_i = \frac{1}{100}$. For simplicity, assume the default risk of the assets is uncorrelated. Thus, the probability that x percent of the assets in the pool default follows a binomial distribution:

$$Pr\left(\sum_{i=1}^{100} d_i s_i = x\right) = \binom{100}{x} p^x (1-p)^{100-x} \tag{13}$$

where d_i is an indicator which equals one if asset i defaults, and zero otherwise. In contrast, when default risk of the underlying assets is correlated S&P estimates the probability distribution of (14) using a Gaussian Copula Monte Carlo simulation.

In a generic sense, the value-at-risk for a probability α, VaR_α, is the default percentage threshold such that $Pr(\sum_{i=1}^{100} d_i s_i \geq VaR_\alpha) = \alpha$. Thus, to compute the SDR what remains is the choice of α. For this, S&P references a table of expected default probabilities. For instance, to estimate the AAA SDR for a CDO with a weighted-average maturity of five years, α is set to S&P's 5-year AAA expected default rate (0.061%). Therefore, in this example the AAA SDR for a CDO with a 5-year WAM is $SDR_{AAA} \approx 34\%$.[33]

Appendix C

This appendix outlines the estimation of the two-state hidden Markov model (HMM) of credit rating changes using the Baum-Welch algorithm. We begin by defining the sequence of observed rating changes as $\{O_1, \ldots, O_T\}$, where O_t is the observed set of rating changes at time t. We also denote the sequence of hidden states as $\{S_1, \ldots, S_T\}$ where $S_t \in \{1\text{: } good,\ 2\text{: } bad\}$ is the state at time t, and the *good* state of the world is index by 1. Finally, credit rating changes at time t follow the state-dependent transition matrix $\Pi(S_t) \in \{\Pi(1)\text{: } good,\ \Pi(2)\text{: } bad\}$.

We assume that the hidden state evolves over time according to the two-state transition matrix, A, and denote the distribution of the initial state by the vector Q:

$$\begin{aligned} a_{ij} &= P(S_t = j | S_{t-1} = i) \\ q_i &= P(S_1 = i). \end{aligned} \tag{14}$$

The probability of observing the set of rating changes O_t given the economy is in state k is proportional to:

$$b_k(O_t) = P(O_t | S_t = k) = P(O_t | \Pi(k)) = \prod_{i=1}^{N} \prod_{j=1}^{N} \pi(k)_{ij}^{n_{ij}(t)} \tag{15}$$

where N is the number of possible credit ratings, $n_{ij}(t)$ is the number of firms transitioning from rating i to rating j at time t, and $\pi(k)_{ij}$ is the probability of transitioning from rating i to rating j according to the rating transition matrix $\Pi(k)$. To estimate the full parameter set $\Theta = \{A, Q, \Pi(\cdot)\}$ we implement the Baum-Welch algorithm. For this, we define the forward probability with the following recursive relationship:

$$\begin{aligned} \alpha_i(t) &= P(O_1, \ldots, O_t, S_t = i | \Theta) \\ \alpha_i(1) &= q_i \cdot b_i(O_1) \\ \alpha_j(t+1) &= \sum_{i=1}^{2} \left(\alpha_i(t) \cdot a_{ij}\right) \cdot b_j(O_{t+1}). \end{aligned} \tag{16}$$

Additionally, we define the backward probability with the following recursive relationship:

$$\begin{aligned} \beta_i(t) &= P(O_{t+1}, \ldots, O_T | S_t = i, \Theta) \\ \beta_i(T) &= 1 \\ \beta_i(t) &= \sum_{j=1}^{2} \left(a_{ij} \cdot b_j(O_{t+1}) \cdot \beta_j(t+1)\right). \end{aligned} \tag{17}$$

The probability of being in state i at time t given $\{O_1, \ldots, O_T\}$ becomes:

$$\gamma_i(t) = \frac{\alpha_i(t) \cdot \beta_i(t)}{\sum_{j=1}^{2} \alpha_j(t) \cdot \beta_j(t)}. \tag{18}$$

Thus, $\gamma_1(t)$ represents the probability of being in the *good* state of the world at time t given the full information of all rating changes over the entire sample period. Finally, the probability of transitioning from state i at time t to state j at time $t+1$ is:

$$\xi_{ij}(t) = \frac{\gamma_i(t) \cdot a_{ij} \cdot b_j(O_{t+1}) \cdot \beta_j(t+1)}{\beta_i(t)}. \tag{19}$$

Given these definitions, we now proceed to the procedure used to estimate $\hat{\Theta}$, such that:

$$\hat{\Theta} = \arg\max_{\Theta} P(O_1, \ldots, O_T | \Theta). \tag{20}$$

We begin by initializing the parameters with an initial guess, $\Theta^{(0)}$. We set the transition matrix, A, and distribution of the initial state, Q, to the following:

- $A^{(0)} : a_{ij}^{(0)} = 0.5\ \forall\, i, j$
- $Q^{(0)} : q_i^{(0)} = 0.5\ \forall\, i$

In addition, we must also begin with an initial guess for the two state-dependent transition matrices. To do this, we first compute the average credit rating across all firms for each period. Next, we assign each period whose average credit rating is greater than the median credit rating across all periods to the *good* state and assign the remainder of the periods to the *bad* state. Finally, we use this classification to assign weights to each period's observed rating changes from which we estimate the state-dependent transition matrices. Specifically, let r be the ordinal credit rating for a firm where AAA corresponds to one. Thus, we initialize the transition matrices to the following:

[33] Given the discrete nature of the default distribution generated from the Monte Carlo simulations, S&P uses linear interpolation to compute the SDR.

$$\hat{\pi}(1)_{ij}^{(0)} = \frac{\sum_{t=1}^{T}[1(\bar{r}_t < r_{median})(0.9) + 1(\bar{r}_t \geq r_{median})(0.1)] \cdot n_{ij}(t)}{\sum_{t=1}^{T}\sum_{k=1}^{N}[1(\bar{r}_t < r_{median})(0.9) + 1(\bar{r}_t \geq r_{median})(0.1)] \cdot n_{ik}(t)}$$
$$\hat{\pi}(2)_{ij}^{(0)} = \frac{\sum_{t=1}^{T}[1(\bar{r}_t < r_{median})(0.1) + 1(\bar{r}_t \geq r_{median})(0.9)] \cdot n_{ij}(t)}{\sum_{t=1}^{T}\sum_{k=1}^{N}[1(\bar{r}_t < r_{median})(0.1) + 1(\bar{r}_t \geq r_{median})(0.9)] \cdot n_{ik}(t)} \quad (21)$$

where $n_{ij}(t)$ is the number of firms that transitioned from rating i to rating j at time t.

Given this initialization of $\Theta^{(0)}$, we calculate $\alpha_i^{(0)}(\cdot)$, $\beta_i^{(0)}(\cdot)$, $\gamma_i^{(0)}(\cdot)$, and $\xi_{ij}^{(0)}(\cdot)$. The parameter estimates are then updated in the following iterative fashion:

- $\hat{A}^{(x+1)} : \hat{a}_{ij}^{(x+1)} = \dfrac{\sum_{t=1}^{T-1}\xi_{ij}^{(x)}(t)}{\sum_{t=1}^{T-1}\gamma_i^{(x)}(t)}$
- $\hat{Q}^{(x+1)} : \hat{q}_i^{(x+1)} = \gamma_i^{(x)}(1)$
- $\hat{\Pi}(k)^{x+1} : \hat{\pi}(k)_{ij}^{x+1} = \dfrac{\sum_{t=1}^{T-1}\gamma_k^{(x)}(t) \cdot n_{ij}(t)}{\sum_{t=1}^{T-1}\sum_{l=1}^{N}\gamma_k^{(x)}(t) \cdot n_{il}(t)}$

This procedure is repeated until the values converge, yielding $\hat{\Theta}$.

Appendix D

This appendix outlines the time-series estimation of macroeconomic covariates used to predict the default risk of firms. We opt to model each macro variable using a first-order auto-regressive structure.

Specifically, we model the civilian unemployment rate U_t and trailing one-year market return S_t at time t as independent AR(1) processes:

$$\begin{pmatrix} U_{t_1} \\ S_{t_1} \end{pmatrix} = \begin{pmatrix} \alpha_U \\ \alpha_S \end{pmatrix} + \begin{pmatrix} \rho_U & 0 \\ 0 & \rho_S \end{pmatrix}\begin{pmatrix} U_t \\ S_t \end{pmatrix} + \begin{pmatrix} \sigma_U & 0 \\ 0 & \sigma_S \end{pmatrix}\varepsilon_{t+1} \quad (22)$$

where ε is a two-dimensional vector of independent standard random normal variables. Over the full sample, we obtain the following parameter estimate:

$$\begin{pmatrix} \hat{\alpha}_U \\ \hat{\alpha}_S \end{pmatrix} = \begin{pmatrix} 0.0303 \\ 0.0970 \end{pmatrix} \quad \begin{pmatrix} \hat{\rho}_U \\ \hat{\rho}_S \end{pmatrix} = \begin{pmatrix} 0.9942 \\ 0.9130 \end{pmatrix}$$
$$\begin{pmatrix} \hat{\sigma}_U \\ \hat{\sigma}_S \end{pmatrix} = \begin{pmatrix} 0.1572 \\ 0.0696 \end{pmatrix}. \quad (23)$$

For the 3-month interest rate and AAA corporate spread over the 10-year rate, we jointly model the three interest rates using a first-order vector auto-regression. We begin by including all lagged coefficients and sequentially exclude statistically insignificantly lagged values until no insignificant values remain. The final model specification is as follows:

$$r_{t+1} = \begin{pmatrix} r_{t+1}^{3m} \\ r_{t+1}^{10y} \\ r_{t+1}^{corp} \end{pmatrix} = \begin{pmatrix} \alpha_{3m} \\ \alpha_{10y} \\ \alpha_{corp} \end{pmatrix} + \begin{pmatrix} \varphi_{11} & \varphi_{12} & 0 \\ \varphi_{21} & \varphi_{21} & 0 \\ 0 & \varphi_{32} & \varphi_{33} \end{pmatrix} r_t + L\,\eta_{t+1}$$
$$LL^T = \Sigma \quad (24)$$

where η_{t+1} is a three-dimensional independent standard random normal variable and L is a 3×3 lower triangular matrix so that LL^T is the covariance matrix Σ of innovations across the system of time-series. The full sample yields the following parameter estimates:

$$\begin{pmatrix} \hat{\alpha}_{3m} \\ \hat{\alpha}_{10y} \\ \hat{\alpha}_{corp} \end{pmatrix} = \begin{pmatrix} -0.0728 \\ 0.0243 \\ 0.1795 \end{pmatrix}$$
$$\hat{\varphi} = \begin{pmatrix} 0.9785 & 0.0237 & 0 \\ -0.0015 & 0.9928 & 0 \\ 0 & 0.0585 & 0.9240 \end{pmatrix}$$
$$\hat{\Sigma} = \begin{pmatrix} 0.0359 & 0 & 0 \\ 0.0404 & 0.0594 & 0 \\ 0.0081 & 0.0200 & 0.0376 \end{pmatrix}. \quad (25)$$

Note: All interest rates and the civilian unemployment rate are expressed in percentage terms before performing the estimation procedure.

References

Araya, R., Mahdavi, Y., Ouzidane, R., 2006. Moody's approach to rating US REIT CDOs. Doc ID no. SF72848. Moody's Investors Service.
Ashcraft, A.B., Goldsmith-Pinkham, P., Vickery, J.I., 2010. MBS ratings and the mortgage credit boom. Staff report no. 449. Federal Reserve Bank of New York.
Azizpour, S., Giesecke, K., Kim, B., 2011. Premia for correlated default risk. Journal of Economic Dynamics and Control 35 (8), 1340–1357.
Azizpour, S., Giesecke, K., Schwenkler, G., 2015. Exploring the sources of default clustering. Unpublished working paper. Stanford University.
Becker, B., Ivashina, V., 2015. Reaching for yield in the bond market. Journal of Finance 70 (5), 1863–1902.
Becker, B., Milbourn, T., 2011. How did increased competition affect credit ratings? Journal of Financial Economics 101 (3), 493–514.
Becker, B., Opp, M., 2013. Regulatory reform and risk-taking: replacing ratings.NBER Working Paper No. 19257. NBER.
Brunnermeier, M.K., 2009. Deciphering the liquidity and credit crunch 2007–08. Journal of Economic Perspectives 23, 77–100.
Campbell, J.Y., Hilscher, J., Szilagyi, J., 2008. In search of distress risk. Journal of Finance 63 (6), 2899–2939.
Cornaggia, J., Cornaggia, K.R., Hund, J., 2017. Credit ratings across asset classes: a long-term perspective. Review of Finance 21 (4), 465–509.
Coval, J., Jurek, J., Stafford, E., 2009. The economics of structured finance. Journal of Economic Perspectives 23 (1), 3–26.
Das, S.R., Duffie, D., Kapadia, N., Saita, L., 2007. Common failings: how corporate defaults are correlated. Journal of Finance 62 (1), 93–117.
Duffie, D., Eckner, A., Horel, G., Saita, L., 2009. Frailty correlated default. Journal of Finance 64 (5), 2089–2123.
Duffie, D., Saita, L., Wang, K., 2007. Multi-period corporate default prediction with stochastic covariates. Journal of Financial Economics 83 (3), 635–665.
Feldhütter, P., Nielsen, M.S., 2012. Systematic and idiosyncratic default risk in synthetic credit markets. Journal of Financial Econometrics 10 (2), 292–324.
Financial Crisis Inquiry Commission, FCIC, 2010. The financial crisis inquiry report: final report of the national commission on the causes of the financial and economic crisis in the United States.FCIC, Washington, DC.
Giesecke, K., 2004. Correlated default with incomplete information. Journal of Banking & Finance 28 (7), 1521–1545.
Giesecke, K., Weber, S., 2004. Cyclical correlations, credit contagion, and portfolio losses. Journal of Banking & Finance 28 (12), 3009–3036.
Griffin, J.M., Nickerson, J., Tang, D.Y., 2013. Rating shopping or catering? An examination of the response to competitive pressure for CDO credit ratings. Review of Financial Studies 26 (9), 2270–2310.

Griffin, J.M., Tang, D.Y., 2012. Did subjectivity play a role in CDO credit ratings? Journal of Finance 67 (4), 1293–1328.
Hand, J.R., Holthausen, R.W., Leftwich, R.W., 1992. The effect of bond rating agency announcements on bond and stock prices. Journal of Finance 47 (2), 733–752.
He, Z., Xiong, W., 2012. Rollover risk and credit risk. Journal of Finance 67 (2), 391–430.
Koopman, S.J., Lucas, A., Schwaab, B., 2011. Modeling frailty-correlated defaults using many macroeconomic covariates. Journal of Econometrics 162 (2), 312–325.
Koopman, S.J., Lucas, A., Schwaab, B., 2012. Dynamic factor models with macro, frailty, and industry effects for US default counts: the credit crisis of 2008. Journal of Business & Economic Statistics 30 (4), 521–532.
Lando, D., Nielsen, M.S., 2010. Correlation in corporate defaults: contagion or conditional independence? Journal of Financial Intermediation 19 (3), 355–372.
Li, D. X., 1999. On default correlation: a copula function approach. Working paper. AIG Asset Management.
Löffler, G., 2005. Avoiding the rating bounce: why rating agencies are slow to react to new information. Journal of Economic Behavior & Organization 56 (3), 365–381.
Moody's, 2004. Moody's revisits its assumptions regarding corporate default (and asset) correlations for CDOs.Available at http://www.alacrastore.com/moodys-credit-research/Moody-s-Revisits-Its-Assumptions-Regarding-Corporate-Default-and-Asset-Correlations-for-CDOs-PBS_SF47403.
Moody's, 2010. Moody's updated modeling parameters for rating corporate synthetic CDOs and cash flow CLOs.Available at http://post.nyssa.org/nyssa-news/2010/03/moodys-updated-modeling-parameters-for-rating-corporate-synthetic-cdos-and-cash-flow-clos.html.
Moody's, 2014. Moody's CDOROM v2.12 user guide.Available at http://www.moodysanalytics.com/~/media/Brochures/Structured-Analytics-Valuation/CDOROM/CDOROM-User-Guide.pdf.
Nickell, P., Perraudin, W., Varotto, S., 2000. Stability of rating transitions. Journal of Banking & Finance 24 (1), 203–227.
Shumway, T., 2001. Forecasting bankruptcy more accurately: a simple hazard model*. The Journal of Business 74 (1), 101–124.
Standard & Poor's, 2003. Corporate ratings criteria.Available at goo.gl/mZRI83.
Standard & Poor's, 2013. CDO evaluator system version 5.1 user guide.Available at http://wenku.baidu.com/view/9c5ed1395727a5e9856a611c.html.
Taleb, N.N., 2007. The Black Swan: The Impact of the Highly Improbable. Random House.