Credit risk assessment of fixed income portfolios: an analytical approach

Bernardo K. Pagnoncelli

Business School. Universidad Adolfo Ibañez, Santiago, Chile.
bernardo.pagnoncelli@uai.cl

Arturo Cifuentes

CREM, Financial Regulation and Stability Center.
Faculty of Economics and Business.
Universidad de Chile, Santiago, Chile.
arturo.cifuentes@fen.uchile.cl

Abstract

We propose a model to assess the credit risk features of fixed income portfolios assuming that such portfolios can be characterized by two parameters: their default probability and their default correlation. We rely on explicit expressions to assess their credit risk and demonstrate the benefits of our approach in a complex leveraged structure example. We show that using the expected loss as a proxy for credit risk is both misleading and dangerous since it does not capture the dispersion effects introduced by correlation. Moreover, we also challenge the prevailing wisdom in relation to the so-called "fat tail" effects in portfolio loss assessments. The implications of these findings are relevant for improving current risk management practices, as well as developing a sound regulatory framework for the fixed income market.
Keywords: Credit risk; expected loss; Risk management; portfolio losses; correlation;

1 Introduction

Fixed income portfolios are normally subject to credit risk, that is, the possibility that some of their assets could default, which, in turn, could result in monetary losses. For example, a key aspect of portfolio management consists of monitoring the credit risk profile of its assets with the aim of triggering corrective actions whenever the risk exceeds some specific level. Additionally, regulators need to understand and quantify the credit risk exposure of the entities they supervise (for instance, banks, insurance companies, pension funds, swap counterparties, to name a few) to make assessments regarding their solvency. Finally, and more broadly, we could speculate that the recent financial (subprime) crisis was caused by an inadequate understanding of the distribution of credit risk within the global financial system.

The naïve approach to deal with portfolios subject to credit risk is to ignore the correlation altogether and treat the defaults as independent random variables. This leads to the conventional binomial distribution. Unfortunately, most real-life situations involve portfolios that do exhibit some degree of correlation and thus this approach is quite misleading. Recognition of the importance of correlation lead to Standard & Poor's (S&P), a rating agency, to deal with correlation—at least initially—by artificially increasing the value of the portfolio default probability, p, according to some ad hoc rules [18]. In essence, S&P treated the pool of assets as being independent with the hope that the effect of correlation could be, somehow, accounted for by modeling the behavior of the pool using a higher (conservative?) value of p. In recent years S&P has abandoned this approach.

Another rating agency, Moody's, suggested a method that recognized the importance of correlation and addressed it via the so-called Diversity Score (DS) concept in combination with the Binomial Expansion Technique (BET), proposed by [2]. The idea behind this method was simple: it relied on the assumption that one could mimic the default behavior of the original pool of assets by using an "ideal" pool made up of uncorrelated bonds, but having the same aggregate notional amount of the actual portfolio. This method was the *de facto* and preferred approach by many analysts for a number of years. Moreover, Moody's relied extensively on the BET method for structured products ratings [6], [17].

More recently, after [12] introduced the Gaussian copula concept to the financial community as a way of generating correlated random vectors (based on a given correlation matrix), Monte Carlo simulations flourished. Notwithstanding the popularity and increasing level of sophistication of these Monte Carlo methods, an unpleasant fact remains: they rely on controlling the so-called asset correlation (a variable related to, but different than, the default correlation, which is the variable we should care about). A full discussion of this issue is beyond the scope of this paper but it has been treated in detail by [1].

In this article we introduce a far simpler model, which relies only on two parameters: the assets default probability (p) and their default correlation (ρ). We do not rely on asset correlations (presumably estimated from returns, another variable that—even though it might be easy to observe—is not obvious that it captures default correlations). In the last few years, several authors such as [8], [9], [13] and [15] have proposed more elaborated models that incorporate the possibility that an asset might default at any time between 0 and some far enough horizon T, but introduce dependence in some ad hoc manner.

It might seem that our model is overly simplistic. However, despite this simplicity it proves to be far more insightful (and accurate) than both, the BET method and Monte Carlo simulations based on the Gaussian copula. In fact, as we will see with a simple example, our model captures a salient feature of the portfolio loss distribution that so far has gone unnoticed: two humps, one at each end of the distribution, which cannot be captured by re-scaling or adjusting the so-called "fat-tail" distributions [11], [16] and [19]. Monte Carlo simulations based on the Gaussian copula—for all their complex correlation matrices and different default probabilities for each asset—cannot account for this salient feature either. Which, we may add, it is critical to assess the likelihood of experiencing extreme events. Moreover, our example will also show that relying of the expected loss as a proxy to assess credit risk is quite misleading since this metric misses the dispersion behavior of the loss distribution. Again, giving an inaccurate appraisal of the portfolio credit risk.

With that as context, we claim that having reliable tools to assess the credit risk characteristics of fixed income portfolios is a problem of significant practical importance. We hope that our paper will contribute to this end.

2 An explicit expression for portfolio defaults and losses

The default behavior of a single asset (or bond, as these two terms will be used interchangeably in this presentation) is a binary event—default/ no default—that can be fully described with one parameter: its default probability p. The problem at hand is more interesting because we consider a portfolio of assets. Hence, we assume that we have n homogeneous bonds, all having the same notional amount, and the same default probability p. But this requires to incorporate an additional piece of information, namely, the correlation. To be precise, the correlation we are referring to in this context is the correlation between the default behavior of any two assets within the portfolio, which we take it to be the same for all pair of assets. Thus, if Y_i and Y_j are index variables (1=default; 0=no default) characterizing the default behavior of assets i and j respectively, we assume that $\mathrm{Corr}(Y_i, Y_j) = \rho > 0 \;\forall\; i, j = 1, \ldots, n$ $(i \neq j)$. Finally, the random variable $X = Y_1 + Y_2 + \ldots + Y_n$ is the number of defaults that the portfolio could experience (clearly, $0 \leq X \leq n$). It can be shown that the density function for $X = Y_1 + \ldots + Y_n$ is as follows:

$$\mathbb{P}\left(X = x | n, p, \rho\right) = \binom{n}{x} p^x (1-p)^{n-x}(1-\rho) + \begin{cases} (1-p)\rho, & \text{if } x = 0, \\ p\rho, & \text{if } x = n, \\ 0, & \text{otherwise,} \end{cases} \tag{1}$$

for $x \in \{0, 1, \ldots, n\}$. For any other value of x the value of the probability in (1) is zero. The mean and variance of X are given by

$$\mathbb{E}[X] = np, \tag{2}$$

$$\mathbb{V}[X] = p(1-p)(n + \rho n(n-1)). \tag{3}$$

Such model appeared in several works that investigate the phenomena of overdispersion in Poisson regression data (see [4], [14] for details). To the best of our knowledge the earliest appearance of the correlated binomial model is [10], in which the authors analyze toxicological experiments. A full derivation of the density of X can be found in [5].

If W refers to the portfolio notional (or par) amount, the loss L experienced by the pool is defined as $L = X(W/n)$, assuming no recoveries. [1] From (2) and (3) we have

$$\mathbb{E}[L] = Wp, \tag{4}$$

$$\mathbb{V}[L] = (W^2/n)p(1-p)(1+\rho(n-1)). \tag{5}$$

It might seem overly simplistic, at first sight, to rely on a single parameter, ρ, to characterize the entire correlation structure of a pool of n assets. Thus, some observations are in order. First, as a practical matter, estimating the entries of a more specialized correlation matrix, namely, a different $\rho_{i,j}$ for each pair (i, j) is impossible due to insufficient data. Notice that defaults (unlike equity prices which can be observed every trading day) do not occur that often. And second, the strong variability exhibited by correlation values as a function of time (high at times of financial stress, low at "normal" times), coupled with the consistency across same asset classes, do not seem to warrant the use of a specialized correlation matrix. This issue is discussed in more detail in [20].

3 A few insightful observations

Let us illustrate our model with a numerical example. Consider a pool of $n = 50$ identical assets, having an aggregate notional amount $W = \$100$, and each a default probability $p = .27$ and correlation $\rho = .1836$. From (4) and (5), we have that $\mathbb{E}[L] = \$27$ and $\mathbb{V}[L] = \$19.85$.

Using the explicit expression (1), we show the density function of the losses L in Figure 1. The x-axis shows the losses in monetary units and the y-axis corresponds to the exact probability of observing those losses. Its features—not only the fact that is tri-modal but also the fact that has two local peaks at both ends of the graph, indicated by arrows in Figure 1—should act as a warning. It indicates that the suggestion of using fat-tailed distributions to improve modeling accuracy, at least in the context of fixed income portfolios, is misguided. Instead of fat-tailed distributions, our model shows that defaults follow thin-tail distributions with protuberances at both ends of the spectrum. In short, we are witnessing a quite different phenomenon.

In order to take into account the correlation between the defaults, some dispersion measure should be used in addition to (or in conjunction with) the expected loss. The natural candidate is a combination of the expected loss with

[1] Without loss of generality we are assuming that each default inflicts a loss of W/n on the pool. In the examples we will allow for a recovery rate greater than 0%.

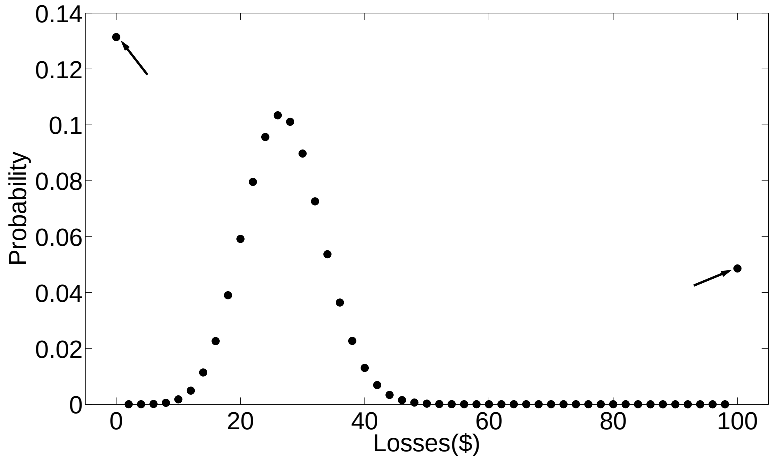

Figure 1: The Correlated binomial density for $n = 50, p = .27, \rho = .1836$ and $W = \$100$. The arrows indicate the probabilities of having 0 and 100 defaults.

the standard deviation, which we will denote by ELSD:

$$\begin{aligned} \mathrm{ELSD}[L] &= \mathbb{E}[L] + k\sigma[L] \\ &= Wp + k\sqrt{(W^2/n)p(1-p)(1+\rho(n-1))}, \; k \in [0, \infty). \end{aligned} \tag{6}$$

The coefficient k represents the level of risk aversion. To illustrate (6), we go back to our initial example with $k = 2$ and calculate ELSD for five values of ρ. The results are on Table 1. There are several other possible choices for measures that incorporate dispersion. The choice is a compromise between which aspects of dispersion are the most relevant (e.g. extreme events, deviation from a threshold) and mathematical tractability. Table 1 highlights the sensitivity of ELSD with respect to correlation, a property that the expected loss does not posses.

ρ	ELSD
0	39.55
.2	68.26
.5	90.41
.8	106.61
1	115.79

Table 1: The ELSD with $k = 2$ for a pool of assets with $n = 50, p = .27$ and notional amount of \$100.

To conclude our observations, we depart slightly from the initial example and allow the number of assets in the pool n to go to infinity. We have

$$\lim_{n\to\infty} \text{Variance}[L] = W^2 \rho p(1-p). \tag{7}$$

For a large n, we observe that the variance of the portfolio exhibits the same behavior described for a fixed value of n: it grows linearly with the correlation and quadratically with the default probability. This feature is relevant for it shows that increasing the number of assets in the pool to achieve more diversification (and reduce risk) has only a limited effect after a certain point. As n grows, a saturation effect takes place and the variance of the portfolio asymptotically approaches $W^2\rho p(1-p)$, as shown in (7), and cannot be reduced any further.

3.1 Comparison with the BET method

Theorem 1 *Consider two homogeneous portfolios with the same notional amount W and made up by assets having identical default probabilities p. The first (actual) portfolio has n assets and the correlation is ρ; the second (ideal) portfolio consists of DS uncorrelated assets. Furthermore, let X be the random variable that denotes the number of defaults associated with the first portfolio and L the corresponding losses. And let $\hat{X}$ and $\hat{L}$ be the corresponding variables in relation to the second portfolio. We claim that if*

$$DS = \frac{n}{n\rho + 1 - \rho},$$

Then

$$\mathbb{E}[L] = \mathbb{E}(\hat{L}) \quad \textit{and} \quad \textit{Variance}[L] = \textit{Variance}[\hat{L}].$$

Proof: Let us notice that $L = (W/n)X$ and $\hat{L} = (W/\text{DS})\hat{X}$ are the random variables representing the losses in the actual and ideal portfolios respectively. Using that in the BET model $\hat{X}$ follows a binomial distribution with parameters DS and p, we have

$$\mathbb{E}[\hat{L}] = \frac{W}{\text{DS}}\, p\text{DS} = Wp,$$

which shows the first part of the result. For the variance, we have

$$\text{Variance}[\hat{L}] = \frac{W^2}{\text{DS}^2}\text{DS}p(1-p) = \frac{W^2(n\rho+1-\rho)^2}{n^2}\frac{n}{n\rho+1-\rho}p(1-p) = \frac{W^2(n\rho+1-\rho)}{n}p(1-p) = \frac{W^2}{n}p(1-p)(1+\rho(n-1)),$$

which completes the proof.

Using as a reference the same portfolio described at the beginning of Section 3, we can perform an interesting analysis using the BET method. In this case the diversity score DS is 5. This value results from evaluating

$$\text{DS} = \frac{n}{n\rho+1-\rho} = \frac{50}{50 \times .1836 + 1 - .1836} = 5.$$

Number of defaults	Probability	Losses (in $)
0	.2073	0
1	.3833	20
2	.2835	40
3	.1048	60
4	.0193	80
5	.0014	100

Table 2: BET computations.

Table 2 shows, for each of the six default scenarios, its probability of occurrence and its corresponding portfolio loss.

The portfolio expected loss and its corresponding standard deviation, as calculated from the data displayed on Table 2, are, respectively, $27 and $19.85. The fact that these values coincide with those obtained previously based on the correlated binomial distribution is not surprising since the DS is derived, as Theorem 1 shows, to enforce the condition that the first and second moment of the loss distribution of both portfolios (the real and the ideal) should match. That said, it would be naïve to expect much comfort beyond this agreement.

More to the point, the fact that the actual distribution (see Figure 1) is tri-modal is in stark contrast with the features of the binomial distribution, on which the BET relies. Accordingly, one should expect significant errors for any computation in which tail-events play an important role. And this is indeed what happens. For example, according to Table 2, the probability of having no defaults is 20.7%. However, using formula (1) we have that the actual value is 13.4% Thus, the BET overestimates the probability of this outcome giving a false sense of security.

Furthermore, suppose an investor wishes to estimate the likelihood of a catastrophic event, say, losing 80% (40 defaults) or more of the value of the portfolio. The BET approximation produces an estimate of roughly 2% for such outcome (the probability of having 4 or more defaults in Table 2). On the other hand, the actual probability of suffering losses exceeding this amount (equivalent to having 40 defaults or more) is, again using expression (1), almost 5%. One more time, the BET grossly underestimates the probability of this event.

In summary, although the BET does a good (indeed, perfect) job at estimating the expected loss and standard deviation of portfolio losses, it is inadequate to provide reliable estimates of extreme events, such as the probability of having a very high (or zero) number of defaults.

4 Structures with leverage (CBOs)

In this section we explore the suitability of the expected loss as a proxy for estimating credit risk in the context of structures with leverage. Figure 2 shows a leveraged structure known as collateralized bond obligation (CBO). The assets consist of a diversified pool of bonds, characterized in terms of their credit-risk properties. The liabilities, or capital structure, consist of three so-called tranches (senior, mezzanine, and equity) which exhibit each a different risk profile. As defaults occur in the reference portfolio (assets), losses are assigned in reverse order of seniority (from the equity to the senior position). It follows then that the senior tranche is the most "conservative" while the equity tranche (also referred to as the first-loss position) is the most speculative.

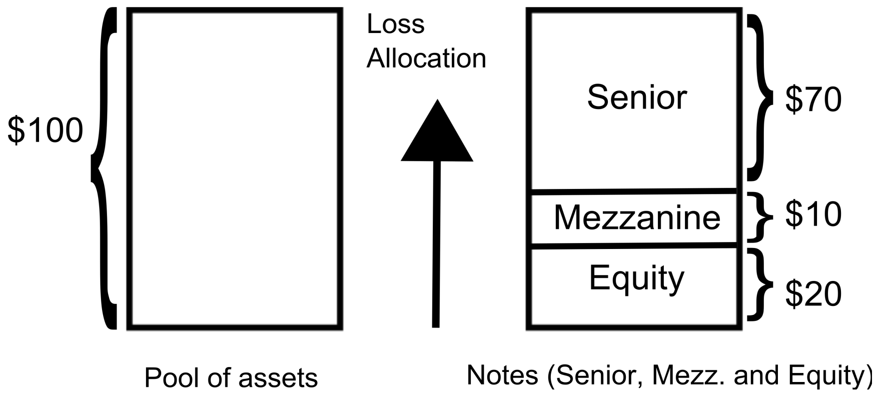

Figure 2: A Simple CBO Structure.

CBO-A refers to a CBO structure like the one described in Figure 2, with a notional amount of \$100 and 3 tranches, senior, mezzanine and equity, with values \$70, \$10 and \$20 respectively. The pool of assets is specified by $p = 12\%, \rho = 0.1$ and $n = 40$. Unlike the previous example, we assume a recovery rate equals to 40%, which means that each default will cause in the pool a loss equal to (\$100/40).60 = \$1.5. It is also assumed that the assets have a 4-year maturity.

Consider now CBO-B, which has the same capital structure as CBO-A, but it is based on a portfolio with different characteristics, namely: $p = 43\%, \rho = 0$ and $n = 45$, with the same recovery rate of 40%. Hence, each default amounts to a (\$100/45).60 = \$1.33 loss. Table 3 shows the expected loss and standard deviation for the senior tranches of CBO-A and CBO-B.

	CBO-A	CBO-B
Expected Loss (%)	.51	.59
Standard deviation (%)	4.66	1.76

Table 3: Expected loss and standard deviation for the senior tranche of CBO-A and CBO-B.

A naïve investor, say, an investor who bases his or her decisions on the value of the expected loss might feel tempted to select the senior tranche of CBO-A instead of CBO-B as a better choice. However, considering the value of the standard deviation of the loss, or more precisely, the coefficients of variation

(9.13 and 2.98 for CBO-A and CBO-B, respectively) one might arrive at a different conclusion. Moreover, the relevance of taking the dispersion into account becomes more apparent if we consider that the threshold between a Baa-rating (investment grade) and a Ba-rating (speculative grade) is, for 4-year assets, $(.66\% + 3.74\%)/2 = 2.2\%$ (Table 4).

In summary, although both senior tranches have a Baa-rating, the senior tranche of CBO-A appears more likely to be downgraded from Baa to Ba. This is not a trivial consideration from a portfolio management viewpoint, as such downgrades normally have unpleasant effects: from being forced to sell the asset to being unable to count it as reserves. A similar phenomena was observed in [7] for CPDO's, a product that has similarities with CBOs. In the recent financial crisis instruments that received the highest possible rating defaulted or lost their value significantly within a two-year horizon because the models used by the two main rating agencies assumed an insignificant probability for events that turned out to occur quite frequently in the market. As our study shows, the errors were not due to parameter uncertainty but rather to conceptual mistakes in the structure of the model themselves.

Therefore, the previous example highlights the need to discard the expected loss as a reliable metric. Ideally, a good metric should capture both the dispersion behavior of the loss as well as the risk aversion profile of the decision maker. The ELSD metric defined in (6) represents an attempt at incorporating both factors into a single metric. For this particular example, any investor with a risk aversion parameter $k > .027$ would prefer the senior tranche of CBO-B rather than that of CBO-A.

Rating	Expected Loss target (%)
Aaa	.00099
Aa	.02585
A	.18975
Baa	.66000
Ba	3.7400
B	9.9715
Caa	24.130

Table 4: Moody's expected loss targets for assets with a 4-year maturity.

	Expected Loss (%)		
ρ	Senior	Mezzanine	Equity
1	5.14	12.00	12.00
0.9	4.62	10.80	14.39
0.5	2.57	6.00	24.00
0.1	.51	1.20	33.59
0	$< 10^{-8}$	.002	36.00

Table 5: Expected loss for the 3 tranches of CDO-A with different values of ρ.

Finally, Table 5 shows that, using as a reference CBO-A, increasing values of correlation tend to reduce the percentage of expected loss of the mezzanine and senior tranches at the expense of the equity position. In essence, changing the correlation redistribute the expected portfolio losses ($\$7.2 = \$100 \times .12 \times .6$) among the 3 tranches. However, the aggregate of the tranche losses remains constant. To illustrate this, consider $\rho = .5$ on Table 5. Calculating per tranche, the expected loss is equal to $.0257 \times \$70 + .06 \times \$10 + .24 \times \$20 = \$7.2$, as stated. This phenomenon has also been observed in [3] and is consistent with the market view that equity investors are "long correlation".

5 Conclusions

In recent years, Monte Carlo methods combined with copula-based correlation characterizations have become the most popular choice to analyze fixed income portfolios—and more complex structures such as CBOs—from a credit risk viewpoint. However, the correlated binomial approach introduced in this paper offers significant advantages compared to such simulation techniques. The first is obvious: its explicit nature. The second advantage is less obvious but perhaps more relevant: it permits to account for the correlation between defaults directly. We must emphasize that Monte Carlo methods handle the default correlation indirectly, using the so-called asset correlation as a proxy for default correlation. The errors caused by this simplification, which are seldom acknowledged, are yet to be fully quantified and understood.

Leaving aside the previous consideration, it might be argued that Monte Carlo methods are better, at least in theory, because they offer the possibility of using more specialized correlation matrices (different values of ρ for each (i, j) pair). In practice, however, this advantage is rendered irrelevant by the fact that correlation coefficients are very difficult to estimate due to insufficient data. The best that an analyst can hope for is to have one overall reliable correlation estimate.

Our analysis has also cast doubts on the appropriateness of using fat-tailed distributions to improve the accuracy of models based on the normal distribution when it comes to defaults. More precisely, this study shows that in general the distribution of defaults is tri-modal, as opposed to unimodal densities such as the binomial. Neglecting the first peak ($X = 0$) might overestimate the probability of having no defaults, while the third peak at $X = n$ implies that the probability of occurrence for default scenarios between the second and third peaks are for all practical purposes zero. This situation, again, contradicts another conventional view: namely, that in a CBO, increasing the subordination level (the "thickness" of the equity tranche) affords more protection to senior or mezzanine investors.

It might seem naïve to address the default behavior of a diversified pool with only one parameter (p) when, at least in theory, each asset might have a different default probability. Let us recognize, however, that the ability to estimate for each asset in the pool its own default probability might be a daunting task. Given that, the thought of capturing, with one "representative" value the overall default behavior of the pool, does not seem far-fetched. Moreover, what is far more relevant is that our approach, for all its simplicity, it does capture the most salient feature of the pool loss distribution: the humps at both ends. A Gaussian copula-based Monte Carlo simulation, no matter how detailed the characterization of the assets default probabilities, will simply miss this critical phenomenon.

We argue that far more prudent than relying on the above-mentioned Monte Carlo is to perform, using our simple formula, a number of sensitivity analysis with different values for p and ρ. For example suppose an investor wishes to evaluate the risk of his/her position in a scenario of crisis. In such a situation defaults become highly correlated. Therefore, it is much more relevant for a potential investor to explore how the expected loss (and its standard deviation) evolve as a function of ρ rather than having (accurate?) estimates of these quantities for only one specific scenario. The simple formula presented in this study is ideal to carry out such analysis.

Finally, our study should force us to rethink the financial regulatory framework in relation to credit ratings. Credit ratings are used, among other things, to determine suitability requirements (for example, the type of assets insurance companies and pension funds can buy); to specify capital reserve levels (in the case of banks); and to trigger liquidation events for entities posting collateral. It is therefore important that the metric employed to define credit ratings should capture well the key dimensions of credit risk. The expected loss, as we have seen here, does a poor job at this task since it fails to account for dispersion. Moreover, the recent performance of credit ratings during, and after, the subprime crisis (that is, massive downgrades of allegedly very safe products) provides additional empirical evidence to the unsuitability of the expected loss as a reliable metric. Our suggestion is to explore metrics that combine the expected loss with some dispersion metric, as indicated previously. In this work we used the standard deviation as the dispersion metric, but there are certainly other options such as the Value-at-Risk, the Mean Absolute Deviation or the Conditional Value-at-Risk, to name a few. One avenue of future research is to compare those metrics and to understand the advantages and disadvantages of using them for different financial products.

References

[1] A. Cifuentes and G. Katsaros. The one-factor gaussian copula applied to CDOs. *The Journal of Structured Finance*, 13(3):60–71, 2007.

[2] A. Cifuentes and G. O'Connor. The binomial expansion method applied to CBO/CLO analysis. *Moodys Investor Services, December*, 13, 1996.

[3] J. Coval, J. Jurek, and E. Stafford. The economics of structured finance. *The Journal of Economic Perspectives*, pages 3–26, 2009.

[4] C.B. Dean. Testing for overdispersion in poisson and binomial regression models. *Journal of the American Statistical Association*, 87(418):451–457, 1992.

[5] C.A.R. Diniz, M.H. Tutia, and J.G. Leite. Bayesian analysis of a correlated binomial model. *Brazilian Journal of Probability and Statistics*, 24(1):68–77, 2010.

[6] I. Fender and J. Kiff. CDO rating methodology: Some thoughts on model risk and its implications. BIS Working Papers 163, Bank for International Settlements, 2004.

[7] M.B. Gordy and S. Willemann. Constant proportion debt obligations: A postmortem analysis of rating models. *Management Science*, 58(3):476–492, 2012.

[8] J.C. Hull and A.D. White. Valuing credit derivatives using an implied copula approach. *The Journal of Derivatives*, 14(2):8–28, 2006.

[9] A. Kalemanova, B. Schmid, and R. Werner. The normal inverse gaussian distribution for synthetic cdo pricing. *The Journal of Derivatives*, 14(3):80–94, 2007.

[10] LL Kupper and J.K. Haseman. The use of a correlated binomial model for the analysis of certain toxicological experiments. *Biometrics*, pages 69–76, 1978.

[11] J.A.C. Lau. *Fat Tails and their (Un) happy Endings: Correlation Bias and its Implications for Systemic Risk and Prudential Regulation (EPub)*. International Monetary Fund, 2011.

[12] D. Li. On default correlation: A copula function approach. *The Journal of Fixed Income*, 9(4):43–54, 2000.

[13] F.A. Longstaff and A. Rajan. An empirical analysis of the pricing of collateralized debt obligations. *The Journal of Finance*, 63(2):529–563, 2008.

[14] A. Luceno. A family of partially correlated poisson models for overdispersion. *Computational statistics & data analysis*, 20(5):511–520, 1995.

[15] D. O'Kane. *Modelling single-name and multi-name credit derivatives*, volume 545. Wiley, 2008.

[16] PIMCO. Understanding tail risk. Technical report, education series, Allianz, 2012.

[17] P.J. Schönbucher. *Credit derivatives pricing models: models, pricing and implementation*. Wiley, 2003.

[18] Standard & Poor's. Global CBO/CLO criteria. Technical report, Standard & Poor's, 1999.

[19] S.V. Stoyanov, S.T. Rachev, B. Racheva-Yotova, and F.J. Fabozzi. Fat-tailed models for risk estimation. *The Journal of Portfolio Management*, 37(2):107–117, 2011.

[20] S. Williams, D. Fenn, M. McDonald, D. Bloom, and P. Mackel. Risk on risk off: the full story. Currency: Quant special, HSBC, 2010.