A New Approach to Credit Ratings*

Giorgi Pertaia
Industrial and Systems Engineering
University of Florida

Artem Prokhorov
University of Sydney Business School
& St.Petersburg State University

Stan Uryasev
Applied Mathematics and Statistics
Stony Brook University

May 2020
Submitted for publication to the Journal of Banking and Finance

Abstract

Credit ratings are fundamental in assessing the credit risk of a security or debtor. The failure of the Collateralized Debt Obligation (CDO) ratings during the financial crisis of 2007-2008 and the massive undervaluation of corporate risk leading up to the crisis resulted in review of rating approaches. Yet the fundamental metric that guides the construction of credit ratings has not changed. We study the inadequacies of the old metric in simple models of investment and in structured finance portfolio optimization tasks, and we propose a new methodology based on a buffered probability of exceedance. The new approach offers a conservative risk assessment, with substantial conceptual and computational benefits. We illustrate the new approach using several examples and report the results of a structuring step-up CDO case study, with details available in an online Supplement.

Keywords: Credit Rating, Probability of Exceedance, Buffered Probability of Exceedance, Expected Shortfall, Conditional Value at Risk, CVaR, Value at Risk, VaR, Loss Given Default, Collateralized Debt Obligation, CDO, Tranche Structuring, Portfolio Optimization, Credit Default Swap, CDS

*Helpful comments from Buhui Qiu and the participants of INFORMS 2019 and SEQURA 2019 are gratefully acknowledged.

1 Introduction

At the height of the financial crisis of 2008, American International Group, Inc. (AIG), once the largest insurance company in the US, was rescued from bankruptcy by a US government bailout worth $85 bn [see, e.g., Sjostrum, 2009]. This was part of the Troubled Asset Relief Program (TARP) that cost the US taxpayer in excess of $245 bn. What caused the companies that enjoyed stable AAA credit ratings to fail abruptly and what role did credit ratings play in the failure?

Early post-crisis literature focused on issues of risk mispricing caused by using dependence models that fail to accommodate realistic tail behavior of joint defaults and on issues around structured finance where loan securitization obscured the true riskiness of the collateral. For example, Coval et al. [2009] and Zimmer [2012] provide two different perspectives at how securitized risky debt was repackaged as virtually risk-free. What is common in the two papers is that they show how rating agencies were simply unfamiliar with assessing creditworthiness of financial instruments that cannot be ascribed to a single company and instead involve a pool of loans, bonds and mortgages from various sources. Thus, the subsequent issues of claims – known as synthetic instruments – against those assets, were not supported by a robust methodology for pricing their riskiness.

Coval et al. [2009] make the point that the new developments in structured finance amplified errors in risk assessment, while Zimmer [2012] shows that the commonly used dependence assumption known as the Gaussian copula was inappropriate. As a result, relatively minor imprecisions in credit risk estimation could have led to variations in default risk of the synthetic securities that were large enough to cause an AAA-rated security to default with a high probability.

Moreover, Ashcraft et al. [2011] looked at a large number of mortgage-backed securities (MBS), collaterized debt obligations (CDO) and other structured finance securities and found empirical evidence that higher credit ratings were closely associated with higher MBS prices after controlling for a large set of security fundamentals. They report that, in terms of value, 80 to 90 percent of sub-prime MBS initially received AAA ratings but were in effect 6-10 rating notches lower[1]. This offers support for the widely held belief that more conservative credit ratings would have muted the crisis by making credit more expensive and providing a more reliable information about synthetic instruments to less informed investor. Bolton et al. [2012] describe the various conflicts of interest that may have added to the inability or unwillingness of credit rating agencies to do that.

Moody's, Standard and Poor's and Fitch Group – the three major credit rating agencies known as

[1]Rating agencies commonly use 21-22 notch scales, from AAA to C or AAA to D.

the Big Three – have evolved since then. They are now more mindful of joint tail risk and synthetic instruments are hardly new any more. More recent papers on this topic focus on how credit rating inflation is affected by competition between agencies, by regulation of the industry and by the business cycle [see, e.g., Dilly and Mahlmann, 2016, Alp, 2013, Altman and Rijken, 2004, Amato and Furfine, 2004, Baghai et al., 2014, Opp et al., 2013, He et al., 2012, Bar-Isaac and Shapiro, 2013, Rabanal and Rud, 2018]. For example, Dilly and Mahlmann [2016] find evidence that credit ratings are inflated during the boom periods and Bar-Isaac and Shapiro [2013] present a model where ratings quality is counter-cyclical. It is noteworthy that the "boom bias" in these papers does not result from changes in legislation or competitive pressures surrounding rating agencies. Rather it comes from the rating agencies' incentive conflicts.

Credit ratings continue to form the basis of credit assessment. They serve as inputs into numerous risk assessment tools such as CreditMetrics of JP Morgan, and they are widely used to determine optimal debt ratio and other aspects of firm's investment decisions [see, e.g., Kisgen, 2006, 2019]. For example, Standard and Poor's now rates over $10 tr in bonds and other assets in more than 50 countries.

In this paper we argue that the fundamental properties of credit ratings have not been given sufficient attention. Incentive conflicts aside, the current credit ratings are prone to massive underestimation of risk. The reason is that they are still primarily guided by the probability of exceedance (PoE), a risk measure which, in addition to suffering from a number of computational issues, estimates the chance of a default-level loss, not the loss given default. We argue that extreme risk exposure can still be concealed behind a high credit rating, which has far reaching implications for financial modeling and operation.

We offer an alternative, more conservative, approach based on a buffered version of PoE. This new measure – referred to as *Buffered Probability of Exceedance* (bPoE) – is tied to a loss threshold, akin to PoE. Unlike PoE, bPoE takes into account the magnitude of tail outcomes exceeding the threshold. It is possible to stretch the tail of the loss distribution and increase the exposure without increasing PoE, but not without increasing bPoE.

More formally, bPoE is the probability of a tail event with the mean tail value equal to a specified threshold. Therefore, by definition, bPoE controls both the average magnitude of the tail and its probability, adding a "buffer" to PoE. The probability measure bPoE is an inverse function to the Expected Shortfall (ES) coherent risk measure, which is also called Conditional Value-at-Risk (CVaR), Superquantile, Average VaR and Tail VaR. In this paper we will use interchangeably the

ES and CVaR terms (the ES term is included in the financial regulations and the CVaR term is used in risk management and optimization applications, which we are referring to in this paper). In the engineering literature, the concept of bPoE has been introduced by Mafusalov and Uryasev [2018] as an extension of the buffered failure probability proposed by Rockafellar [2009] and explored by Rockafellar and Royset [2010].

From the computational perspective, bPoE has considerable advantages compared to PoE. First, bPoE has an analytic representation through a minimization formula [see Mafusalov and Uryasev, 2018], similar to CVaR [see Rockafellar and Uryasev, 2002]. Moreover, bPoE is quasi-convex [see, e.g., Mafusalov and Uryasev, 2018], similar to CVaR which is convex [see, e.g., Rockafellar and Uryasev, 2002]. This means that there are efficient algorithms for solving optimization problems involving these measures. Second, bPoE is a monotonic function of the underlying random variable and a strictly decreasing function of the threshold on the interval between the mean value and the essential supremum. This avoids discontinuity of PoE for discrete distributions.

The link between bPoE and ES is not surprising but has been overlooked. In response to the 2007-2009 crisis, the Basel Committee on Banking Supervision, among other measures, moved from using an unconditional Value-at-Risk (VaR) to ES in order to provide an additional buffer to capital reserve requirements of financial institutions. Yet, no equivalent move has been implemented in the way credit ratings are constructed. Similar to capital reserve requirements, the difference between bPoE and PoE is most pronounced for extremely heavy tailed distributions of losses, so PoE-based ratings fail when they are needed most – at times of distress. Regarding numerical implementation of risk constraints, there is an equivalence between risk constraints on CVaR and bPoE [see Mafusalov and Uryasev, 2018], similar to the equivalence of risk constraints on VaR and PoE. Therefore, bPoE risk constraints can be replaced by CVaR constraints, as described by Rockafellar and Uryasev [2002].

The paper is organized as follows. Section 2 discusses how credit rating construction is guided by the probability of exceedance. Section 3 provides additional motivation for using bPoE. Section 4 studies the disparity between the two measures under the most popular statistical distributions used in structured finance and discusses how we can estimate bPoE. In Section 5 we analyze the adjustments to traditional credit ratings needed to reflect the use of the new measure. Sections 6 and 7 offer several special cases where the distinction bPoE vs PoE matters. In particular, we show (a) what happens to creditworthiness of an insurance company as it accumulates exposure in the way AIG did in the early 2000s, (b) how to solve the problem of optimal CDO structure under credit rating constraints when the use of standard ratings is suboptimal. Additionally, Section 7 contains some detail of a numerical case study which is posted online in its entirety, including codes, data and results.[3] Section 8 concludes.

2 Credit Ratings and Probability of Exceedance

As a risk measure, bPoE has gained initial popularity in areas where tail events can be catastrophic. For instance, in engineering it has been used to assess tropical storm damages [see, e.g., Davis and Uryasev, 2016] and to optimize network infrastructure [see, e.g., Norton and Uryasev, 2018b]. Now the popularity is extending to other areas of risk analytics. For example, in machine learning, it has been used to improve on data mining algorithms [see, e.g., Norton and Uryasev, 2018a, Norton et al., 2017]. However, it has not been introduced to finance, except, perhaps, in asset and liability management [Shang et al., 2016].

Traditionally, credit ratings are driven by historical default rates. These rates are used to estimate the likelihood of a financial loss exceeding the default threshold for a given security or debtor [see, e.g., Trueck and Rachev, 2009, Ch. 2]. Of course, credit ratings are assigned to different entities in different ways. For example, for large issuers, agencies initiate the construction of a rating; for others, a debtor approaches an agency. Rating of some securities and debtors involves a large amount of non-quantitative information collected by credit analysts; for others, only quantitative information is used.

For example, for assigning a rating grade to a company, credit agency analysts usually request financial information about it, consisting of several years of audited annual financial statements, operating and financing plans, management policies, and other credit factors affecting the risk profile of the entity. Some agencies claim to incorporate the extend of potential loss and recovery rates into the risk profile, however, the way it is done is not disclosed and, at best, this information affects ratings indirectly as an element of the risk profile.

All this information goes into constructing a credit score, reflecting the likelihood of default, obtained using a rating algorithm such as a logit model, discriminant analysis and, more recently, machine learning classification techniques such as support vector machines and artificial neural networks. Usually, securities or debtors with a similar risk profile will be assigned to the same rating grade. Sometimes, expert judgments override a rating assignment produced by the algorithm.

Based on the credit scores, probabilities of default are assigned. Using the historical data available to a rating agency and the risk profile of the security or debtor, an agency assigns a rating if

Rating \Time	1	2	3	4	5	6	7	8	9	10
AAA	0	0.03	0.13	0.24	0.35	0.46	0.52	0.61	0.66	0.72
AA	0.02	0.06	0.13	0.23	0.34	0.45	0.55	0.63	0.71	0.79
A	0.06	0.15	0.26	0.4	0.55	0.72	0.92	1.1	1.28	1.48
BBB	0.19	0.53	0.91	1.37	1.84	2.3	2.71	3.11	3.5	3.89
BB	0.73	2.25	4.07	5.86	7.51	9.03	10.34	11.49	12.53	13.45
B	3.77	8.56	12.66	15.82	18.27	20.26	21.89	23.19	24.32	25.37
CCC/C	26.36	35.54	40.83	44.05	46.43	47.28	48.24	49.05	49.95	50.6
Investment grade	0.1	0.28	0.48	0.73	0.98	1.24	1.49	1.72	1.94	2.17
Speculative grade	3.8	7.44	10.6	13.15	15.24	16.94	18.38	19.58	20.65	21.61
All rated	1.49	2.94	4.21	5.27	6.17	6.92	7.57	8.12	8.62	9.09

Table 1: S&P global corporate average cumulative default rates (1981-2015) (%).

probability of default, that is PoE of a given default threshold, is inside a range of default probability characterizing that specific rating grade. Agencies publish tables of default probabilities for each rating grade over a given time horizon. Table 1 contains Standard and Poor's global corporate average cumulative default rates. The data for Table 1 is taken from 2016 Corporate Default S&P Study [Standard and Poor, 2016]. For example, the BBB rating is assigned to an entity with one-year PD in the range $0.08\% < PD \leq 0.23\%$.

As an illustration of how PoE guides the construction of credit ratings, it is useful to think of a synthetic instrument within a simple Merton [1974] model. This is a security for which rating agencies usually use complicated, but exclusively quantitative, models reflecting the various assumptions and approaches involved in constructing the instruments.

Suppose that a firm finances its operation by issuing a single zero-coupon bond with face value B_T payable at time T. Assume that at every time $t \in [0, T]$ the company has total assets A_t. It is standard in the Merton model to assume that A_t follows a Geometric Brownian motion and that default of the company occurs when the firm has no capital (equity) to pay back the debt holders. Because the zero-coupon pays only at time T, default can occur only at T.

The probability of default at time T equals $\mathbb{P}(\text{default}) = \mathbb{P}(A_T < B_T)$. This formula can be rewritten in terms of PoE by changing the sign of assets and liabilities,

$$\mathbb{P}(\text{default}) = \mathbb{P}(A_T < B_T) = \mathbb{P}(-A_T > -B_T).$$

Thus, PD is a PoE of the random variable $-A_T$ exceeding the threshold $-B_T$ and the probabilities in Table 1 can be used to convert the PD into a rating and vice versa.

Figure 1 illustrates PoE as the shaded area $1 - \alpha$. If we define Value-at-Risk (VaR) as the loss

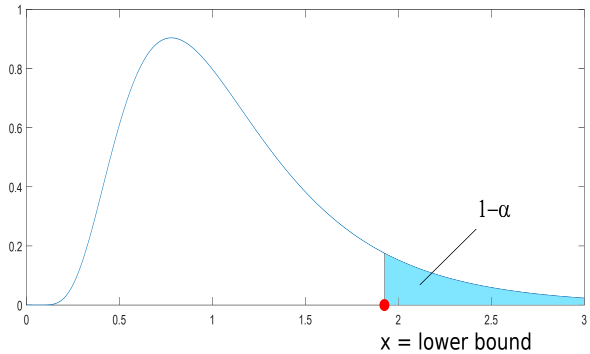

Figure 1: Relationship between PoE and VaR.

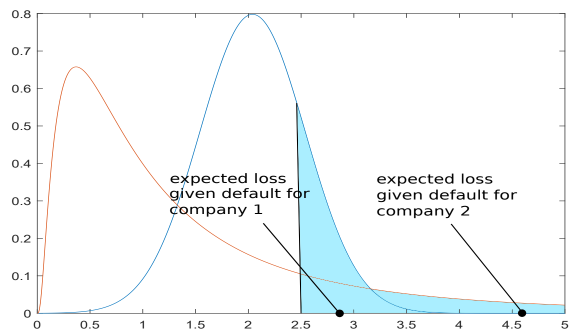

Figure 2: Loss distributions for two companies with equal PoE.

that is exceeded no more than a given (small) proportion of time $1 - \alpha$, then it is clear from the figure that PoE is simply one minus the inverse of VaR. Consequently, PoE-based constraints are equivalent to VaR-based constraints. Hence they are equivalent in terms of rating-based constraints employed by firms in capital structure and investment decisions.

3 Motivation for bPoE-based ratings

The PoE-VaR pair has been criticized on a number of conceptual and computational grounds. First, VaR is not a coherent risk measure because it fails the sub-additivity condition, which implies in essence that a diversified portfolio may have a higher, rather than lower, VaR [see, e.g., Ibragimov, 2009, Ibragimov and Prokhorov, 2016]. Second, VaR is discontinuous, non-differential and non-convex for empirical distributions – a major numerical difficulty for optimization algorithms [see, e.g., Uryasev, 2000]. In particular, when it is necessary to minimize PoE or to impose a PoE constraint, the resulting optimization models are often intractable. Most importantly, the PoE-VaR pair does not account for the magnitude of the loss given default (LGD). Losses of vastly different expected value can have the same VaR and thus rating-based constraints may obscure massive risk exposure.

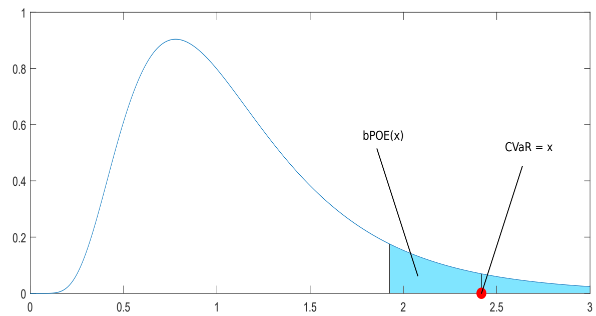

Figure 3: Relationship between bPoE and CVaR.

The PoE-VaR pair offers an overly optimistic measure of risk due to insensitivity to the tail properties of the distribution of losses. For two loss distributions, one with a heavier tail than the other, PoE-based ratings can be identical (in some cases, the instrument with heavier-tailed losses might have a higher rating). Figure 2 illustrates this situation using the normal and log-normal distributions.

It is not difficult to see that CVaR is related to bPoE in the same way as VaR is related to PoE: bPoE is simply one minus the inverse of CVaR. Figure 3 illustrates this relationship with bPoE represented by the shaded area. It is clear from the figure that bPoE measures the probability of a tail event with expected loss equal to CVaR, which captures LGD. This recognizes the shortcomings of the PoE-VaR pair that have led the Basel Committee to adopt CVaR for capital reserve calculations.

In the setting of credit ratings, the 'buffer' interpretation of bPoE is obtained by setting the threshold CVaR at the value of VaR and recovering the difference between bPoE and PoE. It is no surprise that bPoE is no less than PoE for any non-degenerate distribution and the difference can be viewed as a cushion for the LGD implicit in the rating, similar to how CVaR provides a cushion in capital reserves. Since VaR has been supplemented by CVaR in financial industry, it follows naturally that PoE needs the same upgrade.

As mentioned in the introduction, credit ratings have direct influence on prices of financial assets and firm's capital structure and investment decisions and thus debtors might have incentives to inflate ratings. Often, credit ratings work on the 'issuer pays' basis, where the issuer of the security pays to get a rating from the Big Three, which implies an incentive conflict. Debtors may have incentive to exploit weaknesses of existing credit rating models.

Comprehensive risk profiling by a rating agency may be able to spot an excessive LGD and this will translate into a hopefully higher PoE of a default loss. However, even if rating agencies are incentivized to do that, the credit scores they produce are based on historical default rates. Thus they reflect the likelihood of default, not the buffered likelihood of default. In order words, by not accounting for LGD explicitly, traditional ratings do not distinguish clearly between securities and debtors that have a major impact on the investor in case of default, from those that do not.

In certain cases, credit ratings are assigned using only quantitative information and securitization conceals LGD. In particular, for the synthetic instruments such as CDOs, it is possible to clearly define an event of default of its tranches – see Section 7. This event is usually expressed as the loss exceeding some threshold value, known as an attachment point of the tranche. Having a clear definition of default of a tranche and well specified distribution of losses for the asset pool underlying the CDO, we can compute the probability of default. Based on this a rating grade of a tranche is assigned. However, the joint loss distribution for a large and diverse pool of assets, e.g., bonds and mortgages, may be unavailable to the agency. In this case, credit rating of securitized debt may raise both conceptual and computational difficulties.

Contrary to PoE, bPoE has exceptional mathematical properties. It is a quasi-convex function of the loss which makes it a desirable function for optimization models. In particular, minimization models with bPoE inequality constraints can lead to convex or even linear programming problems, which can be solved very efficiently, in contrast to discontinuous and non-convex problems associated with PoE constraints. This offers a potential for efficient solutions to firm's investment and capital structure problems with rating-based constraints [see, e.g., Hung et al., 2017, Wojewodzki et al., 2018].

4 bPoE Definition and Estimation

We now turn to the mathematical and statistical properties of the bPoE-CVaR pair and compare them to the PoE-VaR counterparts. We start by re-emphasizing that non-sub-additivity makes VaR a non-convex function. Non-convexity means that optimization problems with VaR constraints or with a VaR objective function are, in general, intractable for large dimensions. At the same time, optimization problems involving CVaR constraints or with CVaR as the objective function, are usually solvable in polynomial time, using convex or even linear programming methods.

Despite the wide adoption of CVaR in financial industry, there has not been an analogous substitution of the PoE-based methodologies with bPoE-based methodology, even though bPoE inherits similar desirable mathematical properties from CVaR. For example, it is possible to solve a large dimensional portfolio optimization problem with a constraint on bPoE of the portfolio loss at a speed many times faster than the equivalent problem with PoE-based constraints.

We now define the relevant risk measures formally. Let X denote a random loss, F_X its cumulative distribution function and $\alpha \in [0, 1)$ some confidence level, then VaR (or a quantile of X) can be defined as

$$VaR_\alpha(X) = \inf \left\{ v \in R \mid F_X(v) \geq \alpha \right\}$$

and CVaR can be defined as

$$CVaR_\alpha(X) = \inf_t \left\{ t + \frac{1}{1-\alpha} \mathbb{E}[X - t]^+ \right\} , \tag{1}$$

where $[x]^+ = \max\{x, 0\}$ (see, e.g., Rockafellar and Uryasev [2002] for more detail on these definitions). If $\alpha = 0$ then CVaR(X) $= \mathbb{E}(X)$ and as $\alpha \to 1$, CVaR(X) goes to infinity if X is unbounded. It can be shown that CVaR is a coherent risk measure as it satisfies the translation invariance, sub-additivity, positive homogeneity and monotonicity properties, whenever it is defined. The convexity of CVaR follows from the sub-additivity and positive homogeneity.

The relationship between bPoE and PoE is similar to that between CVaR and VaR. The value of bPoE with threshold v of a random variable X equals to the probability mass in the right tail of the distribution of X such that the average value of X in this tail is equal to the threshold v. It is convenient to define bPoE formally as follows, by using the minimization representation [see, e.g., Mafusalov and Uryasev, 2018]

$$bPoE_v(X) = \inf_{a \geq 0} \mathbb{E}[\, a(X - v) + 1]^+ \, . \tag{2}$$

Because bPoE is equal to one minus the inverse function of CVaR, where CVaR gives the average value in the tail having probability $1 - \alpha$, bPoE equals to PoE for the right tail with CVaR equal to v.

The asymptotic results by Mafusalov et al. [2018] suggest a simple estimator of bPoE. Let $\{x_i\}_{i=1}^n$ denote an iid sample of realizations of X. Under fairly general conditions, any quantile of the distribution of X can be consistently estimated by its empirical counterpart. The corresponding CVaR is just the sample mean over the observations exceeding the relevant empirical quantile. Given these estimates, it is natural to estimate bPoE by the sample equivalent of the population problem in (2) as follows

$$\widehat{bPoE}_v = \min_{a \geq 0} \left\{ \frac{1}{n} \sum_{i=1}^{n} [a(x_i - v) + 1]^+ \, I\{v < \max(x_1, \ldots, x_n)\} \right\},$$

where $I\{x\}$ is the indicator function and v can take any estimated CVaR value.

The resulting estimator converges to bPoE uniformly in v at the $\sqrt{n}$-rate. If quantiles are unique, then the solution is

$$a^* = \frac{1}{v - q_\alpha(X)},$$

where $q_\alpha(X)$ is the $(1-\alpha)\%$ quantile of X. In this case,

$$\widehat{bPoE}_v = \frac{1}{n} \sum_{i=1}^{n} [a^*(x_i - v) + 1]^+$$

and Mafusalov et al. [2018] show that

$$\sqrt{n}(\widehat{bPoE}_v - bPoE_v(X)) \to N(0, \sigma_v^2), \text{for any } v,$$

where $\sigma_v^2 = Var\left([a^*(X - v) + 1]^+\right)$.

For any consistent estimator $\hat{a}$ of a^*, a consistent estimator of σ_v^2 can be obtained as follows

$$\hat{\sigma}_v^2 = \frac{1}{n-1} \sum_{i=1}^{n} \left([\hat{a}(x_i - v) + 1]^+ - \widehat{bPoE}_v \right)^2.$$

This gives grounds to statistical inference about the buffer in terms of economic and statistical significance.

An important consequence of these asymptotic results is that standard models of dynamic quantiles of financial returns, including popular GARCH specifications and quantile regressions, can be effectively used in evaluating bPoE. To rating agencies, they permit credit scoring to be based on models of buffered likelihood of default and the resulting credit ratings to include an explicit buffer for LGD.

A suite of standard statistical results also follow from this asymptotic distribution. For example, the $(1-\beta)100\%$ asymptotic confidence bands for bPoE at a given quantile v can be written as

$[\bar{q}_L^\beta(v), \bar{q}_U^\beta(v)]$, where

$$\bar{q}_L^\beta(v) = \max\left(0, \widehat{bPoE}_v(X) - \Phi^{-1}(\beta)\frac{\hat{\sigma}_v}{\sqrt{n}}\right) \tag{3}$$

$$\bar{q}_U^\beta(v) = \min\left(1, \widehat{bPoE}_v(X) + \Phi^{-1}(\beta)\frac{\hat{\sigma}_v}{\sqrt{n}}\right) \tag{4}$$

and $\Phi^{-1}(\cdot)$ is the inverse of standard normal cdf. Using formulas (3) and (4), it is easy to calculate the sample size needed in order to achieve a given precision in estimating bPoE with a desired confidence.

5 bPoE Ratings

The idea of the proposed methodology is to use bPoE to guide the construction of credit ratings and the use of rating-based constraints. This means, in order to assign a rating grade, we propose estimating bPoE for the same loss threshold as before and assigning credit grades using a revised conversion table.

The most obvious revision is to scale the probabilities of default in Table 1 in order to align them with bPoE, not PoE. This will be scaling up since bPoE is no less than PoE evaluated at the same threshold. For example, if losses are distributed according to the standard normal distribution, bPoE for this distribution is roughly 2.4 times higher than PoE calculated at the commonly used thresholds; if losses are log-normally distributed with parameters $\mu = 0$ and $\sigma = 1$ then bPoE is roughly 3.2 time higher. As an illustration, Figure 4 plots the ratio bPoE/PoE for standard normal distribution as a function of PoE (left panel) and as a function of quantile threshold v (right panel). The question, however, is what loss distribution to use.

In principle, each security or debtor has its own loss distribution and, in general, credit rating agencies do not have access to this information even if it exists. For example, risk profiles traditionally constructed by the agencies in order to assign a debtor to a rating grade do not include historical distribution of losses of the debtor. At best, they have access to historical recovery rates of similar-profile debtors. However, in the case of synthetic instruments, the loss distributions can usually be evaluated by simulation under the assumptions that govern the construction of such instruments. Once we obtain an estimate of bPoE for the rated entity, we can assign a grade to it based on the revised conversion table.

As a benchmark adjustment to the conversion table we propose scaling the default probabilities by the factor $e = 2.72$. This adjustment factor will not seem ad hoc if we notice that this is the bPoE/PoE ratio for the exponential distribution. Therefore, this is the buffer required to account for the loss given default when losses have exponentially decaying tails of the distribution. There are two reasons why the exponential distribution is a good candidate for a benchmark scaling factor. First, the exponential distribution can serve as a 'demarcation line' between heavy- and light-tailed distributions, where a distribution is called heavy-tailed if

$$\lim_{v \to \infty} e^{\lambda v} \mathbb{P}(X \geq v) = \infty \quad \forall \lambda > 0,$$

that is, if the distribution has heavier tails than exponential with arbitrary parameter λ. A security or debtor with a heavier-tailed loss distribution than exponential will have a higher bPoE and thus will receive a lower rating. Second, the $bPoE(v)/PoE(v)$ ratio for the exponential distribution with arbitrary parameter $\lambda > 0$ and arbitrary threshold value $v > \mathbb{E}X$ is constant. It is easy to show that bPoE for the exponential distribution equals $e^{1-\lambda v}$ [see, e.g., Mafusalov et al., 2018], while PoE is $e^{-\lambda v}$. Thus no adjustment is needed to the various legislated VaR thresholds.

When done across all rating grades, such a scaling will simply replace PoE-based definitions of rating grades with bPoE-based definitions. Table 2 implements this rescaling using the probabilities in Table 1. In practice, when no information is available about the loss distribution of a security or debtor, traditional credit scoring algorithms can be used prior to scaling and the point of the transformed conversion table is to produce more conservative credit ratings. However, when the agency has a way of assessing the loss distribution then bPoE can be estimated and the rating grade will reflect the security-specific loss distribution tail index.

As an example, consider assignment of a rating grade to two synthetic instruments, both are CDOs but structured differently:

Case I: the loss distributions for the underlying asset pool is such that the pooled loss has an exponential distribution with parameter $\lambda > 0$. Then, the bPoE credit rating for this CDO will be exactly the same as its PoE rating. The scaled conversion table will place that CDO into the same grade as before the conversion and this will be regardless of the value of λ.

Case II: the the loss distributions of the assets in the pool is such that the pool has a Pareto distributed loss with parameters $\alpha > 0$ and $x_m > 0$, that is,

$$F(x) = 1 - \left(\frac{x_m}{x}\right)^{\alpha}, \quad x \geq x_m.$$

Similarly to the exponential distribution, the $bPoE_v/PoE_v$ ratio for the Pareto distribution does not depend on threshold v, however it depends on parameter α and is equal to

$$bPoE_v/PoE_v = \left(\frac{\alpha}{\alpha-1}\right)^{\alpha}, \quad \alpha > 1. \tag{5}$$

Assume independence of asset losses in the pool and note that the right hand side of (5) goes to ∞ as $\alpha \to 1$ and it goes to e as $\alpha \to \infty$. Therefore, the corresponding CDO will have a higher bPoE and hence a lower rating. In particular, its bPoE will be $\left(\frac{\alpha}{\alpha-1}\right)^{\alpha}/e$ times higher compared to the exponential distribution (irrespective of threshold v). For example, if $\alpha = 1.1$, which is not unusual especially for financial returns on emerging markets, the bPoE will be more than 5 times higher. Using the one-year ahead values from Tables 1 and 2, if this CDO used to be AA, it is now BBB!

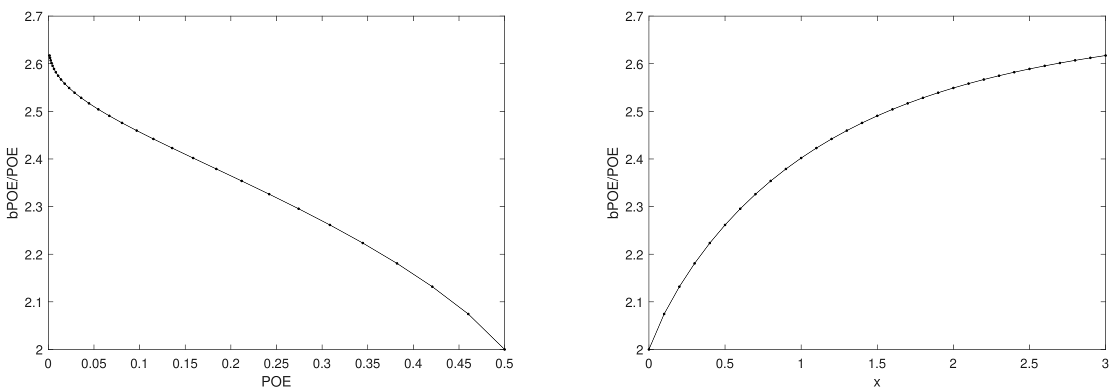

Figure 4: Left graph shows the relationship between $PoE_v(X)$ and $bPoE_v(X)/PoE_v(X)$ for the standard normal distribution. Right graph shows relationship between quantile v and $bPoE_v(X)/PoE_v(X)$ for the standard normal distribution.

Rating \Time	**1**	**2**	**3**	**4**	**5**	**. . .**	**10**
AAA	0.00	0.08	0.35	0.65	0.95	. . .	1.96
AA	0.05	0.16	0.35	0.63	0.92	. . .	2.15
A	0.16	0.41	0.71	1.09	1.50	. . .	4.02
BBB	0.52	1.44	2.47	3.72	5.00	. . .	10.57
BB	1.98	6.12	11.06	15.93	20.41	. . .	36.56
B	10.25	23.27	34.41	43.00	49.66	. . .	68.96
CCC/C	71.65	96.61	100.00	100.00	100.00	. . .	100.00
Investment grade	0.27	0.76	1.30	1.98	2.66	. . .	5.90
Speculative grade	10.33	20.22	28.81	35.75	41.43	. . .	58.74

Table 2: Revised ratings for buffered probability of default.

In order to further illustrate the effect of switching from PoE to bPoE consider the case of two assets, one with normally distributed losses with parameters $\mu = 10$ and $\sigma = 3$ (Asset 1) and another with log-normal losses with parameters $m = 0.02852$ and variance $v = 1$ (Asset 2). Suppose that each asset defaults if the loss is greater than 18.7 (default threshold). Then, probability of default of each asset is 0.018% and therefore, each will have an AA-rating. However, the expected loss for Asset 1 in case of default is 19.6, while the expected loss for Asset 2 in case of default is 26.4. Clearly Asset 2 is a more risky investment but the ratings do not differential between them. If we now turn to bPoE, Asset 1 has bPoE=0.049%, corresponding to an AA-rating (unchanged), while Asset 2 has bPoE = 0.059%, giving it an A-rating. This is, of course, the reflection of the fact that the loss distribution of Asset 2 has a heavier tail.

6 Uncovered Call Options Investment Strategy: Exposure Incentive

We illustrate how bPoE-based ratings prevent overuse of uncovered call options investment strategies similar to those employed in the industry around the time of the AIG debacle. The idea is that the conventional credit rating incentivizes the strategies that load the book with upper tranches of CDOs without appropriate hedging. This leads to accumulation of uncovered exposure which is not reflected in the credit rating. Effectively, the tail of the loss distribution is made arbitrary heavy and the credit rating fails to reflect this.

Consider the following simple model. Suppose that a portfolio manager sells a number of uncovered call options with the same underlying asset and strike price K. Without loss of generality, assume zero interest rates and let $P(K)$ be the price of the option with strike price K. We assume that the portfolio has no capital, except proceeds from selling the call options. The portfolio manager sells $n_{_K} = \frac{1}{P(K)}$ options so that the proceeds from the sale are equal to \$1. Given an underlying asset price S_T at maturity time T, the call option payoff at time T equals $f_K = \max\{S_T - K, 0\}$.

The portfolio will have a negative balance, when $n_{_K} f_K - 1 > 0$, which is defined to be default. Thus, the default probability $\mathbb{P}(n_{_K} f_K - 1 > 0)$ is PoE of the random variable $n_{_K} f_K - 1$ with the threshold at 0.

Assume that S_t evolves over time according to the geometric Brownian motion

$$S_t = S_0 e^{(\mu - \sigma^2/2)t + \sigma W_t} \,,$$

where W_t is the Wiener process. Thus, price S_t is log-normally distributed with cumulative distribution function

$$F_t(x) = \frac{1}{2} + \frac{1}{2}\text{erf}\left(\frac{\ln(x) - m_t}{\sqrt{2}s_t}\right)$$

with parameters $m_t = \ln S_0 + (\mu - 0.5\sigma^2)t$ and $s_t = \sigma\sqrt{t}$, for each time instance $t \in (0, T]$ (see, e.g., Chapter 14 in Hull [2009]).

Then, it is easy to show that the probability of default equals

$$\mathbb{P}(n_{_K} f_K - 1 > 0) = \mathbb{P}\left(\max\{S_T - K, 0\} > \frac{1}{n_{_K}}\right) = \mathbb{P}\left(\max\{S_T - K - \frac{1}{n_{_K}}, -\frac{1}{n_{_K}}\} > 0\right)$$

$$= \mathbb{P}\left(S_T > K + \frac{1}{n_{_K}}\right) = 1 - F_T\left(K + \frac{1}{n_{_K}}\right) .$$

Furthermore,

$$PoE_0(n_{_K} f_K - 1) = \mathbb{P}(n_{_K} f_K - 1 > 0) \to 0 \quad \text{as} \quad K \to \infty . \tag{6}$$

In other words, we can reduce PoE by simply increasing the strike price K. It is no surprise in this setting, that top-notch ratings can be obtained using PoE-based credit rating by increasing the exposure of the portfolio through a sufficiently high strike price.

Now consider the bPoE for $n_{_K} f_K - 1$ at the threshold value of 0. The probability density function of random variable $X_K = n_{_K} f_K - 1$ has a single atom located at point $X_K = -1$ with probability $\mathbb{P}(n_{_K} f_K - 1 = -1) = \mathbb{P}(S_T \leq K)$. The cdf of $X_K \geq x$ where $x \in (-1, \infty)$ is

$$\mathbb{P}(X_K \geq x) = \mathbb{P}(n_{_K} f_K - 1 \geq x) = \mathbb{P}\left(f_K \geq \frac{x+1}{n_{_K}}\right) = \mathbb{P}\left(S_T \geq \frac{x+1}{n_{_K}} + K\right).$$

Thus, for values grater than -1, the distribution of X_K is the same as log-normal distribution corresponding to S_T, however it is shifted left by $K + 1/n_{_K}$ and is scaled by $n_{_K}$.

In order to calculate $bPoE_0(X_K)$ we need to consider two possibilities. In the first case, when $\mathbb{E}\left[X_K I\{X_K > -1\}\right] \leq 0$, the value $bPoE_0(X_K)$ can be calculated using only the log-normal part of the distribution of X_K. In the second case, when $\mathbb{E}\left[X_K I\{X_K > -1\}\right] > 0$, it is necessary to take some fraction of the probability of the atom, in order to bring the conditional expectation down to 0. Thus, in the second case, $bPoE_0(X_K)$ is calculated as

$$bPoE_0(X_K) = \mathbb{P}(X_K > -1) + p_{-1},$$

where p_{-1} is the fraction of the probability of the atom, such that

$$0 = \mathbb{E}\left[X_K I\{X_K > -1\}\right] + (-1)p_{-1}. \tag{7}$$

Because $n_{_K}$ is always positive,

$$bPoE_0(X_K) = bPoE_0(n_{_K} f_K - 1) = \mathbb{P}(f_K > 0) + p_{-1} = \mathbb{P}(S_T > K) + p_{-1}\ .$$

Note that $I\{X_K > -1\} = I\{n_{_K} f_K - 1 > -1\} = I\{S_T > K\}$. Then, from (7), we have

$$\begin{aligned}
bPoE_0(X_K) &= \mathbb{P}(S_T > K) + \mathbb{E}\left[X_K I\{X_K > -1\}\right] \\
&= \mathbb{P}(S_T > K) + \mathbb{E}\left[(n_{_K} f_K - 1) I\{S_T > K\}\right] \\
&= \mathbb{P}(S_T > K) + n_{_K}\mathbb{E}\left[f_K I\{S_T > K\}\right] - \mathbb{P}(S_T > K) \\
&= n_{_K}\mathbb{E}\left[f_K\right]. \qquad (8)
\end{aligned}$$

From the fundamental theorem of asset pricing (see Chapter 14 in Hull [2009]) and our assumption that interest rates (including risk free) are 0, we have that

$$n_{_K} = \frac{1}{P(K)} = \frac{1}{\mathbb{E}\left[f_K\right]}\ . \tag{9}$$

Substituting the right-hand side of (9) in (8) we get

$$bPoE_0(X_K) = n_{_K}\mathbb{E}\left[f_K\right] = \frac{1}{\mathbb{E}\left[f_K\right]}\mathbb{E}\left[f_K\right] = 1, \tag{10}$$

so bPoE-based rankings would be informative of the extremely high riskiness of the strategy.

7 Application to Optimal Step-Up CDO Structuring

A significant advantage of bPoE-based ratings is the possibility they offer to solve to optimality complicated portfolio optimization and structuring problems. This section discusses the problem of CDO structuring in order to demonstrate this advantage with a practical example.

A CDO consists of a pool of assets generating a cash flow. This asset pool is repackaged in a number of tranches with ordered priority on the collateral in the event of default. Each of these

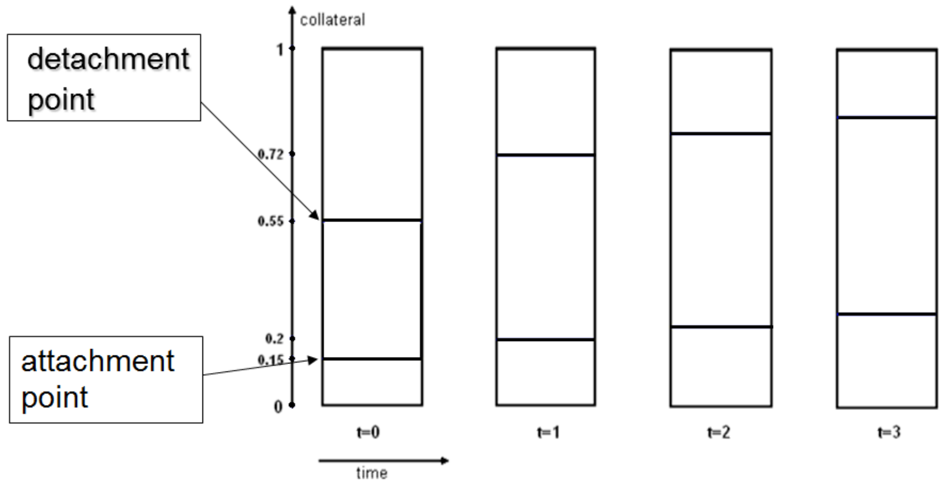

Figure 5: CDO attachment and detachment points

tranches comes with a separate rating assigned to it. Top-quality tranches, often called *senior tranches*, have the first priority on collateral payment in the event of default. They have a higher rating compared to other tranches, often called *mezzanine and equity tranches*.

Each tranche has an *attachment and detachment point* that controls the amount of loss absorbed by the tranche in the event of default. The detachment point of a given tranche is the attachment point of the following upper tranche. A CDO tranche defaults when the cumulative loss reaches its attachment point. At each time t, there is a set of attachment/detachment points that determine the width of a tranche as illustrated in Figure 5. Traditionally, the tranche rating is calculated based on the PoE of the loss using the attachment point of the tranche as a threshold. Tranches are sold to investors as separate assets and the payoff (spread) of the tranche depends on the assigned rating.

A CDO consisting of a pool of credit default swaps (CDSs) is called *synthetic*. CDS buyers make payments (*CDS spreads*) every time period to the CDO originator. The CDO originator "repackages" these spreads and makes payments (*tranche spreads*) to buyers of CDO tranches. A tranche spread depends on an attachment point and is driven by the tranche rating.

We now show how the bPoE-based ratings can be used for structuring of synthetic CDOs. Given a fixed pool of assets, a common objective pursued in structured finance is to select positions in the pool of assets and optimal attachment points.

7.1 Optimal CDO Structuring with PoE-Based Ratings

We consider the optimization problem faced by an originator of a synthetic CDO: to find such positions in a pool of CDSs and such CDO attachment points that result in maximum profit. That is, we minimize the expected sum of discounted spread payments subject to constraints on tranche ratings (to ensure CDO tranche spreads) and a constraint on the cost of the purchased pool [see, e.g., Veremyev et al., 2012, Problem B]. We solve this optimization problem for various costs of purchased pool. The most profitable CDO has the largest difference between the received CDS spreads and paid tranche spreads.

We start with describing the problem of calculating optimal attachment points assuming a fixed pool of CDSs. Then, we extend the problem and include a CDS portfolio optimization component.

Let M denote the number of tranches, T the number of time periods, L_t the loss at each time $t = 1, \ldots, T$ and s_m the spread payment for each tranche $m = 1, \ldots, M$. The total payment for a tranche with the given attachment/detachment points (x^t_m, x^t_{m+1}) at time t can be written as follows

$$\sum_{m=1}^{M} (x^t_{m+1} - \max(x^t_m, L_t))^+ s_m \ . \tag{11}$$

Given a discount rate r, the goal is to minimize the expected present value of the total future payments with respect to tranche attachment points $\{x^t_2, \ldots, x^t_m\}$

$$\sum_{t=1}^{T} \frac{1}{(1+r)^t} \sum_{m=1}^{M} \mathbb{E}(x^t_{m+1} - \max(x^t_m, L_t))^+ s_m \ . \tag{12}$$

The lowest attachment point is assumed to be fixed at zero, $x^t_0 = 0$, $t = 1, \ldots, T$.

Veremyev et al. [2012] proved that the objective function (12) has the following equivalent representation

$$\sum_{t=1}^{T} \frac{1}{(1+r)^t} \sum_{m=1}^{M} \Delta s_m \mathbb{E}[(x^t_{m+1} - L_t)^+] \ , \tag{13}$$

where $\Delta s_m = s_m - s_{m+1}$. Being a sum of expectations of convex functions in attachment points x^t_m, this representation is more attractive as it is convex in x^t_m, which is a desirable property in optimization.

By construction, each tranche in the CDO has a predefined rating. Let $\hat{p}_m$ denote the probability of default at any time point up to T, corresponding to a given tranche rating. Then, the rating constraints on the tranche attachment points are written as follows

$$1 - \mathbb{P}(L_1 \le x_m^1, \ldots, L_T \le x_m^T) \ \le \ \hat{p}_m \quad m = 2, \ldots, M\ . \tag{14}$$

These constraints bound default probabilities of tranches. Note that in (14) index m starts from 2 because the attachment point of the lowest tranche is fixed at zero. Additionally, the attachment points should satisfy the monotonicity constraints

$$x_m^t \ge x_{m-1}^t \quad m = 3, \ldots, M;\ t = 1, \ldots, T. \tag{15}$$

Let us denote by $\boldsymbol{x} = \{x_m^t\}_{m=2,\ldots,M}^{t=1,\ldots,T}$ the vector of attachment points. By combining (13)-(15), we write the optimization problem as follows

$$\min_{\boldsymbol{x}} \sum_{t=1}^{T} \frac{1}{(1+r)^t} \sum_{m=1}^{M} \Delta s_m \mathbb{E}[(x_{m+1}^t - L_t)^+] \tag{16}$$

subject to the constraints

$$\begin{aligned}
&1 - \mathbb{P}(L_1 \le x_m^1, \ldots, L_T \le x_m^T) \le \hat{p}_m && m = 2, \ldots, M && (17)\\
&x_m^t \ge x_{m-1}^t && m = 3, \ldots, M; \quad t = 1, \ldots, T && (18)\\
&0 \le x_m^t \le 1 && m = 2, \ldots, M;\ \ t = 1, \ldots, T && (19)
\end{aligned}$$

Further we include optimization of the portfolio of CDSs into problem (16)-(19). Let K denote the number of available CDSs. Let $y_k,\ k = 1, \ldots, K$, denote the weight of the k-th asset in the asset pool and c_k the annual income spread payment of the CDS. Assume that the CDS portfolio should obtain an annual spread of at least ζ, where ζ is a parameter. Let θ_k^t denote the cumulative loss of asset k at time t. Then, the total loss of the CDS pool at time t is $L(\boldsymbol{\theta}^t, \boldsymbol{y}) = \sum_{k=1}^{K} \theta_k^t y_k$, where $\boldsymbol{\theta}^t = (\theta_1^t, \ldots, \theta_K^t)$ and $\boldsymbol{y} = (y_1, \ldots, y_K)$.

The following optimization problem finds an optimal portfolio allocation as well as optimal attachment points:

$$\min_{\boldsymbol{x},\boldsymbol{y}} \sum_{t=1}^{T} \frac{1}{(1+r)^t} \sum_{m=1}^{M} \Delta s_m\, \mathbb{E}\left[\, \left(x_{m+1}^t - L(\boldsymbol{\theta}^t, \boldsymbol{y})\right)^+ \right] \tag{20}$$

subject to the constraints

$$1 - \mathbb{P}(\, L(\boldsymbol{\theta}^1, \boldsymbol{y}) \le x_m^1 \,, \ldots ,\; L(\boldsymbol{\theta}^T, \boldsymbol{y}) \le x_m^T \,) \;\le\; \hat{p}_m \qquad m = 2, \ldots, M \tag{21}$$

$$\sum_{k=1}^{K} c_k y_k \;\ge\; \zeta \tag{22}$$

$$\sum_{k=1}^{K} y_k \;=\; 1 \tag{23}$$

$$y_k \;\ge\; 0 \qquad k = 1, \ldots, K \tag{24}$$

$$x_m^t \;\ge\; x_{m-1}^t \qquad m = 3, \ldots, M; \quad t = 1, \ldots, T \tag{25}$$

$$0 \;\le\; x_m^t \;\le\; 1 \qquad m = 2, \ldots, M; \quad t = 1, \ldots, T \tag{26}$$

The stated optimization problem (20)-(26) finds an optimal CDO structure for a fixed annual income spread payment ζ. However, the objective is to find a CDO with a minimal difference between total discounted income spread payment of the CDS portfolio and total expected spread payments of tranches. To accomplish this task, we can solve problem (20)-(26) for a grid of parameter ζ and take the solution with the highest expected profit

$$\sum_{t=1}^{T} \frac{1}{(1+r)^t}\,\zeta \;-\; \sum_{t=1}^{T} \frac{1}{(1+r)^t} \sum_{m=1}^{M} \Delta s_m \,\mathbb{E}\left[\, \left(x_{m+1}^t - L(\boldsymbol{\theta}^t, \boldsymbol{y}) \right)^+ \right] . \tag{27}$$

Probability constraints (17) and (21) are non-convex. Problems with such constraints are difficult to solve to optimality. Veremyev et al. [2012] used Portfolio Safeguard (PSG) software [2] for solving problems with probability constraints. PSG has pre-coded probability functions and specially designed heuristic algorithms for optimization with probability constraints, similar to the heuristic described by Larsen et al. [2002] for Value-at-Risk optimization. PSG provides reasonable solutions for these non-convex problems, but does not guarantee optimality.

Data, codes and solutions for six simplified variants of problem (16)-(19) described in detail in Veremyev et al. [2012] are posted online [3]; see Problems 1-6 and the description of the case study posted on the website.

[2] Portfolio Safeguard (PSG) is a product of American Optimal Decisions: `http://aorda.com`

[3] Online Supplement "Case study: structuring step-up CDO" is available at `http://uryasev.ams.stonybrook.edu/index.php/research/testproblems/financial_engineering/structuring-step-up-cdo/`

7.2 Optimal CDO Structuring with bPoE-Based Ratings

In this section we demonstrate how the non-convex risk management problems with PoE constraints can be converted into convex problems with bPoE constraints. Specifically we consider extensions of optimization problem (20)-(26) which find, in one shot, the optimal attachment points of a CDO and the optimal portfolio component allocation. We emphasize that this problem has not, until now, been solved to optimality.

Constraint (21) is equivalent to

$$PoE_0\left(\max\left\{L(\boldsymbol{\theta}^1,\boldsymbol{y})-x_m^1\,,\ldots,\;L(\boldsymbol{\theta}^T,\boldsymbol{y})-x_m^T\right\}\right)\;\leq\;\hat{p}_m\;, \tag{28}$$

which defines a non-convex region for variables $\boldsymbol{x},\boldsymbol{y}$. However, a similar constraint with bPoE

$$bPoE_0\left(\max\left\{L(\boldsymbol{\theta}^1,\boldsymbol{y})-x_m^1\,,\ldots,\;L(\boldsymbol{\theta}^T,\boldsymbol{y})-x_m^T\right\}\right)\;\leq\;\bar{p}_m \tag{29}$$

defines a convex region, where $\bar{p}_m$ is some scaled bound on probability. The convexity of the feasible region of constraint (29) follows from the quasi-convexity of bPoE in random variable and convexity of the function $\max\left\{L(\boldsymbol{\theta}^1,\boldsymbol{y})-x_m^1\,,\ldots,\;L(\boldsymbol{\theta}^T,\boldsymbol{y})-x_m^T\right\}$ in $(\boldsymbol{x},\boldsymbol{y})$ [see Mafusalov and Uryasev, 2018, Proposition 4.8]. More importantly, bPoE constraint (29) is more conservative compared to PoE constraint (28) and the effect of this replacement is greater for loss distributions with fatter tails. Under the assumption of an exponential tail, the following bPoE constraint is equivalent to (28)

$$bPoE_0\left(\max\left\{L(\boldsymbol{\theta}^1,\boldsymbol{y})-x_m^1\,,\ldots,\;L(\boldsymbol{\theta}^T,\boldsymbol{y})-x_m^T\right\}\right)\leq e\,\hat{p}_m\;, \tag{30}$$

where $e = 2.7182$. However, generally the two constraints are not equivalent.

An additional benefit is that bPoE and CVaR constraints are equivalent [see Mafusalov and Uryasev, 2018]. Constraint (30) is equivalent to the following CVaR constraint

$$CVaR_{1-e\,\hat{p}_m}\left(\max\left\{L(\boldsymbol{\theta}^1,\boldsymbol{y})-x_m^1\,,\ldots,\;L(\boldsymbol{\theta}^T,\boldsymbol{y})-x_m^T\right\}\right)\leq 0\;. \tag{31}$$

bPoE is one minus the inverse function of CVaR, therefore, an appropriate confidence level in CVaR constraint (31) is $1 - e\,\hat{p}_m$. CVaR-based objective functions and constraints are coded in many software packages. For instance, MATLAB has a financial toolbox that included CVaR functions.

In addition, it is worth mentioning that the probability function in the left hand side of equations

(21) and (28) is the high-dimensional CDF of random vector $\boldsymbol{L(\theta, y)} = \Big(L(\boldsymbol{\theta}^1, \boldsymbol{y}), \dots, L(\boldsymbol{\theta}^T, \boldsymbol{y})\Big)$ at point $\boldsymbol{x}$, i.e.,

$$
\begin{aligned}
CDF_{\boldsymbol{L(\theta,y)}}(\boldsymbol{x}) &= \mathbb{P}(\, L(\boldsymbol{\theta}^1, \boldsymbol{y}) \le x_m^1 \,, \dots, \, L(\boldsymbol{\theta}^T, \boldsymbol{y}) \le x_m^T \,) \\
&= 1 - PoE_0 \left(\max \left\{ L(\boldsymbol{\theta}^1, \boldsymbol{y}) - x_m^1 \,, \dots, \, L(\boldsymbol{\theta}^T, \boldsymbol{y}) - x_m^T \right\} \right) .
\end{aligned}
$$

We can define *buffered* CDF of the random vector $\boldsymbol{L(\theta, y)}$ at point $\boldsymbol{x}$ as follows

$$
bCDF_{\boldsymbol{L(\theta,y)}}(\boldsymbol{x}) = 1 - bPoE_0 \left(\max \left\{ L(\boldsymbol{\theta}^1, \boldsymbol{y}) - x_m^1 \,, \dots, \, L(\boldsymbol{\theta}^T, \boldsymbol{y}) - x_m^T \right\} \right) .
$$

Therefore, convex constraint (29) is a constraint on buffered CDF $bCDF_{\boldsymbol{L(\theta,y)}}(\boldsymbol{x})$ and it is equivalent to constraint (31) on CVaR.

Next we demonstrate the effect of replacing PoE with bPoE and CVaR in the CDO structuring problem described in the previous section. We numerically solve problem (20)-(26) with three types of risk constraints for tranches:

Problem PoE: problem (20)-(26) with PoE constraint (21);

Problem bPoE: problem (20)-(26) with PoE constraint (21) replaced by bPoE constraint (30);

Problem CVaR: problem (20)-(26) with PoE constraint (21) replaced by CVaR constraint (31).

Codes, data, and solutions for these three problems are posted online [3]; see Problems 7-9.

We consider a CDO with 5 tranches ($M = 5$) using the data from Veremyev et al. [2012]. The most "Senior" tranche has the highest credit rating (AAA), followed by the "Mezzanine 1" (AA), "Mezzanine 2" (A), "Mezzanine 3" (BBB) and finally "Equity" tranche (no rating). The planning horizon is 5 years ($T = 5$). The interest rate is assumed to be $r = 7\%$ and, for simplicity, discounting is done in the middle of the year. With Standard and Poor's credit ratings and corresponding default probabilities from Table 1, maximum default probabilities for tranches are presented in Table 3 in column "PoE DP". Column "bPoE DP" of Table 3 contains default probabilities based on bPoE, see Table 2; column "Rating" shows ratings of the tranches.

Tranche	PoE DP	bPoE DP	Rating
Senior	0.35%	0.95%	AAA
Mezzanine 1	0.34%	0.92%	AA
Mezzanine 2	0.55%	1.50%	A
Mezzanine 3	1.84%	5.00%	BBB

Table 3: "Tranche" = tranche name; "PoE DP" and "bPoE DP" = maximum default probability for a CDO tranche based on PoE and bPoE of attachment point; "Rating" = rating of tranche. Equity tranche does not have a rating.

In Table 3, Senior tranche has higher default probability than Mezzanine 1 tranche (both PoE and bPoE), which is unusual but follows from the actual S&P default rates. A CDS pool is the underlying asset for the considered CDO. Simulations of default scenarios for the CDSs in the pool were done using Standard and Poor's CDO Evaluator. The dataset contains a list of defaults and recovery rates for CDSs in the pool for 3×10^5 scenarios. This number of scenarios is considered adequate for low probability events, such as a default of the AAA Senior tranche.

For optimization we used PSG[2] version 2.3, running on a Windows 10 PC, with Intel Core i7-8550U CPU. We used the VANGRB PSG solver which minimizes a sequence of quadratic programming problems by calling the GUROBI solver. We note that although *problems with bPoE* and *CVaR* constraints are convex in decision variables, they are quite challenging from the numerical perspective. A reduction to linear programming with additional variables results in dimension exceeding 1.2×10^7 variables. PSG is capable of solving such problems to optimality because it has pre-coded bPoE and CVaR functions with algorithms designed to solve such problems. PSG also has a pre-coded PoE function and optimization problems with PoE functions can be solved to optimality for problems with a small number of scenarios. However, for a large number of scenarios, PSG uses a heuristic suggested by Larsen et al. [2002], which does not guarantee optimality.

Risk constraint	Solving time (sec)	Objective value	Duality Gap
PoE	3223	544.54	N/A
bPoE	286	545.78	10^{-5}
CVaR	285	545.78	10^{-5}

Table 4: Numerical results for CDO structuring problem with three types of risk constraints.

Table 4 reports solving time, minimal objective value and solution precision (Duality Gap) for the problems with PoE, bPoE and CVaR risk constraints. First, we note that *Problem bPoE* and *Problem CVaR* provide identical objective values and duality gaps. This is not surprising as these two problems are equivalent. Also, we observe that changing the constraint type from PoE to bPoE or CVaR leads to a significant improvement in solution time with a similar objective value. The solving-time improvement is particularly remarkable as *Problem bPoE* and *Problem CVaR* were solved to optimality (Duality Gap = 10^{-5}) while there is no optimality guaranty for *Problem PoE* (Duality Gap=N/A (not applicable)).

To illustrate the difference between PoE- and bPoE-based ratings, the dataset from Veremyev et al. [2012] was modified for 500 scenarios with a 50% default rate of each CDS in the CDO. These stressed losses happen in the 5-th year of CDS payments and ensure that the tail of the loss distribution is heavy. However, the probability of high losses for stressed scenarios is small: $500/(3 \times 10^{-5}) = 1.6 \times 10^{-3} = 0.166\%$. This probability is about twice lower than the AAA default probability of 0.35%, see Table 3.

As discussed earlier, this implies a significant difference between PoE- and bPoE-based ratings and illustrates how bPoE ratings accurately account for heavy tails. The calculation results for *Problem PoE* and *Problem bPoE*, based on the data including the modified stressed scenarios, are provided in Tables 5-7.

Risk constraint	**Solving time (sec.)**	**Objective value**	**Duality Gap**
PoE	3976	550.24	N/A
bPoE	182	589.73	0

Table 5: Numerical results for *Problem PoE* and *Problem bPoE* with stressed scenarios.

Tranches	**PoE sol**	**bPoE sol**	**PoE rating**	**bPoE rating**
Senior	0.34%	1.93%	AAA	BBB
Mezzanine 1	0.34%	1.92%	AA	BBB
Mezzanine 2	0.55%	2.65%	A	BBB
Mezzanine 3	1.84%	6.52%	BBB	BB

Table 6: Solution of *Problem PoE* with stressed scenarios. “PoE sol” and “bPoE sol” = PoE and bPoE for tranches at optimal point of *Problem PoE*; “PoE rating” and “bPoE rating” = PoE and bPoE rating according to Table 3.

Tranches	PoE sol	bPoE sol	PoE rating	bPoE rating
Senior	0.18%	0.92%	AAA	AAA
Mezzanine 1	0.18%	0.92%	AAA	AA
Mezzanine 2	0.25%	1.50%	AAA	A
Mezzanine 3	1.33%	5.00%	A	BBB

Table 7: Solution of *Problem bPoE* with stressed scenarios. “PoE sol“ and “bPoE sol” = PoE and bPoE for tranches at optimal point of *Problem bPoE*; “PoE rating” and “bPoE rating” = PoE and bPoE rating according to Table 3.

Comparison of Table 4 and Table 5 shows that the objective value of *Problem PoE* with stressed scenarios has increased only by about 1% from 544.54 to 550.24. The PoE constraint has a small impact on profitability of CDO because it is not sensitive to the stressed scenarios with a small probability. In contrast, the objective value of *Problem bPoE* has increased by 8% from 545.78 to 589.73. The bPoE constraint is sensitive to the small-probability events and decreased profitability of CDO. Also, Table 5 shows that *Problem bPoE* is solved 22 times faster than *Problem PoE*.

Table 6 demonstrates that PoE-based ratings do not reflect the increased riskiness coming from the stressed scenarios. In this Table, “PoE sol” and “bPoE sol” stand for PoE and bPoE at the solution of *Problem PoE*, respectively; “PoE rating” and “bPoE rating” stand for tranche ratings calculated using PoE- and bPoE-based ratings, respectively, see Table 3. The PoE tranche ratings in Table 6 coincide with the ratings in Table 3 of the original *Problem PoE* without the stressed scenarios. However, the corresponding bPoE ratings in Table 6 are drastically lower than the PoE ratings, reflecting the actual risk of the additional 500 stressed scenarios.

Finally, Table 7 shows the solution of *Problem bPoE*, similar to Table 6 showing the solution of *Problem PoE*. Here, the CDO is calibrated using the bPoE rating constraints. Again, the PoE-based ratings are severely inflated. In this case, the high PoE-based ratings can be interpreted as the ratings that would have been required to correctly reflect the riskiness of the additional stressed scenarios.

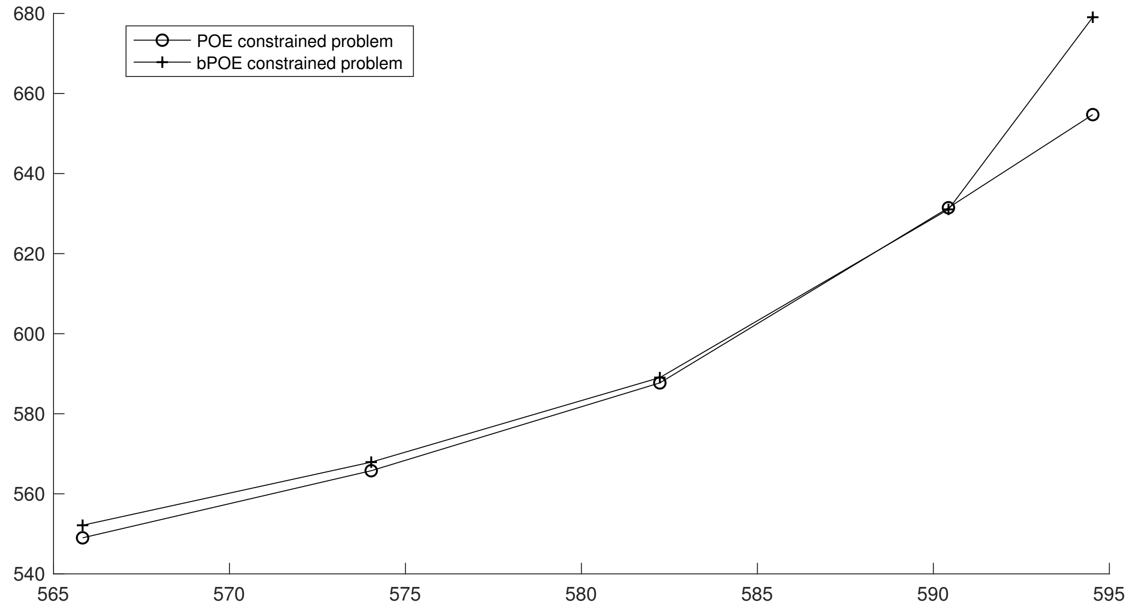

Figure 6: The horizontal axis is the total discounted income, over 5 years, payed by the CDS pool of the CDO (in basis points); the vertical axis is the total discounted spread payments of the CDO (in basis points).

Now we return to the comparison of the solutions of *Problem PoE* and *Problem bPoE* for the original dataset without stressed scenarios. The optimized objective values of these two problems are very close, see Table 4. But there is still the question of whether using one or the other constraint type generates significantly different cash flows. Figure 6 shows total discounted spread payments of all tranches of the CDO (vertical axis in basis points) versus total discounted income generated by the CDS pool underlying the CDO (horizontal axis in basis points). The calculations are made for the two problems with $\zeta = \{138, 140, 142, 144, 145\}$ in budget constraint (22). The solutions of the two problems are quite close, except at the highest value of income generated by the CDS pool where the bPoE-based solution gives a higher spread. We note that CDO profitability is measured by the difference between the horizontal and vertical scale in the graph. We observe that the highest profitability is achieved for the smallest income values. This means that for the considered dataset, the highest profitability is achieved by a portfolio of CDSs with a low spread (and low default probability) for both the PoE- or bPoE-constrained portfolios.

8 Conclusion

This paper presents a new approach to credit ratings based on the bPoE risk function. bPoE-based ratings have a number of attractive features compared to PoE-based ratings. They explicitly account for the magnitude of loss given default via their dependence on the tail of the loss distribution. bPoE is a quasi-convex function in the random variable, which implies that risk optimization problems are much easier to solve than PoE optimization problems.

We show that bPoE-based ratings address crucial inadequacies characterizing traditional credit ratings. These include incentivizing excess risk exposure and encouraging credit risk mispricing. We demonstrate bPoE's advantages using several examples, including an uncovered call options investment strategy with the incentive to open exceedingly large positions with low default probabilities.

With CDO structuring problems, we argue that the new approach shows exceptional promise from the computational perspective as it makes use of the quasi-convexity of bPoE-based constraints and of the reduction to convex and linear programming. PoE-based ratings do not capture high-value low-probability risks. bPoE-based ratings overcome this deficiency for loss distributions with long tails.

References

A. Alp. Structural shifts in credit rating standards. *The Journal of Finance*, 68(6):2435–2470, 2013.

E. I. Altman and H. A. Rijken. How rating agencies achieve rating stability. *Journal of Banking and Finance*, 28(11):2679 – 2714, 2004.

J. D. Amato and C. H. Furfine. Are credit ratings procyclical? *Journal of Banking and Finance*, 28 (11):2641 – 2677, 2004.

A. Ashcraft, P. Goldsmith-Pinkham, P. Hull, and J. Vickery. Credit ratings and security prices in the subprime mbs market. *The American Economic Review*, 101(3):115–119, 2011.

R. P. Baghai, H. Servaes, and A. Tamayo. Have rating agencies become more conservative? implications for capital structure and debt pricing. *The Journal of Finance*, 69(5):1961–2005, 2014.

H. Bar-Isaac and J. Shapiro. Ratings quality over the business cycle. *Journal of Financial Economics*, 108(1):62 – 78, 2013.

P. Bolton, X. Freixas, and J. Shapiro. The credit ratings game. *The Journal of Finance*, 67(1): 85–111, 2012.

J. Coval, J. Jurek, and E. Stafford. The economics of structured finance. *Journal of Economic Perspectives*, 23:3–25, 2009.

J. Davis and S. Uryasev. Analysis of tropical storm damage using buffered probability of exceedance. *Natural Hazards*, 2016.

M. Dilly and T. Mahlmann. Is there a boom bias in agency ratings? *Review of Finance*, 20(3): 979–1011, 2016.

J. J. He, J. Q. Qian, and P. E. Strahan. Are all ratings created equal? the impact of issuer size on the pricing of mortgage-backed securities. *The Journal of Finance*, 67(6):2097–2137, 2012.

J. Hull. *Options, Futures, and other Derivatives (7 ed.)*. Pearson, 2009.

C.-H. Hung, A. Banerjee, and Q. Meng. Corporate financing and anticipated credit rating changes. *Review of Quantitative Finance and Accounting*, 48(4):893–915, 2017.

R. Ibragimov. Portfolio diversification and value at risk under thick-tailedness. *Quantitative Finance*, 9(5):565–580, 2009.

R. Ibragimov and A. Prokhorov. Heavy tails and copulas: Limits of diversification revisited. *Economics Letters*, 149:102 – 107, 2016.

D. Kisgen. Credit ratings and capital structure. *Journal of Finance*, 61(3):1035–1072, 2006.

D. Kisgen. The impact of credit ratings on corporate behavior: Evidence from moody's adjustments. *Journal of Corporate Finance*, 58:567–582, 2019.

N. Larsen, H. Mausser, and S. Uryasev. *Algorithms for Optimization of Value-At-Risk*, pages 19–46. 2002.

A. Mafusalov and S. Uryasev. Buffered probability of exceedance: Mathematical properties and optimization. *SIAM Journal on Optimization*, 28(2):1077–1103, 2018.

A. Mafusalov, A. Shapiro, and S. Uryasev. Estimation and asymptotics for buffered probability of exceedance. *European Journal of Operational Research*, (270):826–836, 2018.

R. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29:449–470, 1974.

M. Norton and S. Uryasev. Maximization of auc and buffered auc in binary classification. *Mathematical Programming*, 174:575–612, 2018a.

M. Norton, A. Mafusalov, and S. Uryasev. Soft margin support vector classification as buffered probability minimization. *Journal of Machine Learning Research*, (18):1–43, 2017.

M. A. Norton, M. and S. Uryasev. Cardinality of upper average and its application to network optimization. *SIAM Journal on Optimization*, 28(2):1726–1750, 2018b.

C. C. Opp, M. M. Opp, and M. Harris. Rating agencies in the face of regulation. *Journal of Financial Economics*, 108(1):46 – 61, 2013.

J. P. Rabanal and O. A. Rud. Does competition affect truth telling? an experiment with rating agencies. *Review of Finance*, 22(4):1581–1604, 2018.

R. Rockafellar. Safeguarding Strategies in Risky Optimization. In *Presentation at the International Workshop on Engineering Risk Control and Optimization, Gainesville, FL*, 2009.

R. Rockafellar and J. Royset. On buffered failure probability in design and optimization of structures. *Reliability Engineering and System Safety*, 95(5):499–510, 2010.

R. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26:1443–1471, 2002.

D. Shang, V. Kuzmenko, and S. Uryasev. Cash flow matching with risks controlled by buffered probability of exceedance and conditional value-at-risk. *Annals of Operations Research*, pages 1–14, 2016.

W. Sjostrum. The aig bailout. *Washington and Lee Law Review*, 66:941–991, 2009.

Standard and Poor. Standard and poor's global ratings. 2015 annual global corporate default study and rating transitions. 2016.

S. Trueck and S. Rachev. *Rating Based Modeling of Credit Risk*. 2009.

S. Uryasev. *Introduction to the Theory of Probabilistic Functions and Percentiles (Value-at-Risk)*, pages 1–25. Kluwer Academic Publishers, Dordrecht, 2000.

A. Veremyev, P. Tsyurmasto, and S. Uryasev. Optimal structuring of cdo contracts: Optimization approach. *Journal of Credit Risk*, 8, 2012.

M. Wojewodzki, W. Poon, and J. Shen. The role of credit ratings on capital structure and its speed of adjustment: an international study. *European Journal of Finance*, 24(9):735–760, 2018.

D. Zimmer. The role of copulas in the housing crisis. *The Review of Economics and Statistics*, 94: 607–620, 2012.