เทคนิคในการวิเคราะห์เชิงปริมาณ สำหรับพอร์ตโฟลิโอของตราสารเครดิต

สิระ สุจินตะบัณฑิต*

บทคัดย่อ

กุญแจสำคัญในการวิเคราะห์เชิงปริมาณสำหรับพอร์ตโฟลิโอที่ประกอบไปด้วยตราสารเครดิต คือการบ่งชี้การกระจายและค่าคาดหมายของความสูญเสียของพอร์ตโฟลิโอที่เกิดจากการผิดชำระ บทความนี้ทบทวนทฤษฎีการสร้างแบบจำลองทางคณิตศาสตร์สำหรับพอร์ตโฟลิโอตราสารเครดิต และวิเคราะห์ระเบียบวิธีที่ไม่ต้องพึ่งการจำลองสถานการณ์ที่ใช้กันอย่างแพร่หลายในการการคำนวณการกระจายและค่าคาดหมายของความสูญเสีย แต่เดิมวรรณกรรมเกี่ยวกับระเบียบวิธีเหล่านี้มักมุ่งเน้นไปที่การคำนวณการกระจายของความสูญเสียเป็นหลัก โดยละการคำนวณค่าคาดหมายของความสูญเสียไว้ ซึ่งสามารถคำนวณได้โดยปริยายจากการกระจายของความสูญเสีย ในทางตรงกันข้าม บทความนี้นำเสนอรูปแบบปรับปรุงของระเบียบวิธีเหล่านั้นโดยเน้นที่การคำนวณค่าคาดหมายของความสูญเสียเป็นหลัก และอธิบายถึงวิธีการคำนวณย้อนกลับไปหาการกระจายของความสูญเสียได้ถ้าต้องการ เราเริ่มด้วยการทบทวนว่าระเบียบวิธีที่ใช้กันทั่วไปนั้นมีขั้นตอนการปฏิบัติเพื่อการคำนวณการแจกแจงของความเสียหายอย่างไร พร้อมทั้งชี้ให้เห็นถึงข้อบกพร่องของขั้นตอนการปฏิบัติดังกล่าว จากนั้นเราได้เสนอการปรับปรุงระเบียบวิธีให้ดีขึ้น และชี้แจงถึงประโยชน์ที่ตามมาที่จะสามารถช่วยบรรเทาข้อบกพร่องนั้นได้ และได้นำเสนอตัวอย่างเชิงเลขที่แสดงให้เห็นถึงข้อดีของวิธีการที่แนะนำ

* อาจารย์ภาควิชาการธนาคารและการเงิน จุฬาลงกรณ์มหาวิทยาลัย

Techniques in the Quantitative Analysis of Portfolios of Credit Instruments

Sira Suchintabandid*

Abstract

The key to the quantitative analysis of a portfolio of credit instruments is the determination of the distribution and the expected value of the portfolio's default loss. This article reviews the theory behind the mathematical modeling of credit portfolios, and analyzes the widely-used, non-simulation-based methods for computing the loss distribution and the expected losses. Existing literature on these methods typically focuses on the computation of loss distribution as the primary quantity of interest, leaving the expected loss to be implied from the loss distribution. This article, on the other hand, proposes an improved version of these methods that focuses on computing the expected loss as the primary quantity of interest, and explains how to retrieve the loss distribution if desired. We first provide a concise review of how the methods are commonly implemented to compute the loss distribution, and point out the drawbacks of such an implementation. We then propose an improved variation of the methods and discuss its benefit in alleviating those drawbacks. Numerical examples that demonstrate the advantage of the suggested methods are also provided.

* Lecturer, Department of Banking and Finance, Chulalongkorn Business School

1. Introduction

Credit instruments generally refer to financial products that are subject to default loss, such as defaultable bonds and loans. Thanks to the developments in the derivative market, credit instruments have expanded to include the so-called credit derivatives — a class of derivative contracts whose payoff are determined by the credit events of certain reference entities. A Credit Default Swap (CDS), a contract which pays off based on the reduction in market value of the underlying bond when it defaults, is an example of credit derivatives and can be considered a type of credit instrument as well.

Quantifying the credit risk is necessary for both the risk management and the pricing of credit instruments. Credit risk management, an important process that impacts many areas such as reserve requirement and collateral provision, typically requires assessing the probability of default and the size of the loss upon default. Likewise, the pricing of credit instruments, such as the determination of bond spread or the valuation of CDS contracts, also requires one to quantify the credit risk and assess the loss that can occur from default.

In many situations, quantifying the credit risk requires one to adopt the so-called *portfolio view*. Consider, for example, the risk management of a bank's portfolio of loans. In addition to assessing the risk-return characteristic of each loan individually, the bank must also consider the effect of the correlation between the loans, and assess the risk-return characteristic of the entire portfolio as a whole.

Another situation in which the portfolio view of credit risk must be adopted is the analysis of the so-called basket credit derivatives. Such derivative instruments have payoff that is based, not on one single credit instruments, but on a *portfolio* of credit instruments. Collateralized Debt Obligations (CDOs), for example, is a form of asset-backed securities that results from securitizing a *portfolio* of debt instruments. A CDO conceptually behaves like a fixed-income instrument, paying periodic coupon based on a notional principal that declines as obligors in the portfolio default. CDOs are often structured into tranches, where each tranche absorbs different portion of the portfolio's default loss. For example, a CDO structure may contains five tranches: 0% – 3%, 3% – 6%, 6% – 9%, 9% – 12%, 12% – 22%. This means that the first tranche (known as the equity tranche) absorbs the first 3% of the portfolio's loss, the second tranche absorbs the loss in excess of the first 3% up to 6%, and so on. (The levels 3%, 6%, ... are called attachment points.) It is known that the risk-return characteristic of CDO tranches is very sensitive to the default correlation among obligors. Thus, the pricing and risk management of these tranches require one to adopt the portfolio view.

The key to the portfolio view of credit risk is the assessment of the portfolio's total loss that occurs from aggregating the default losses from its constituents. Depending on which obligors default and how much loss is incurred, different scenarios, each with its own likelihood, will result in different values of total loss. Thus, the aggregate loss, when viewed as a random variable, follows a *probability distribution* that depends on the probability of default of the constituents, their loss-given-default, and the default correlation. Computing this probability distribution is key to the risk management of credit portfolios. For example, the *Credit Value-at-Risk* (VaR), a widely-used measure of credit risk, is defined in terms of the quantile of the probability distribution.

Closely related to the probability distribution, the expected loss is another key quantity of interest. In particular, we often want to compute not only the expected loss of the entire portfolio, but also the expected loss of certain *portions* of the portfolio. For example, pricing a CDO tranche requires the expectation of the portion of the loss that occurs in that tranche. As another example, computing the *expected shortfall*, a risk measure usually reported along with VaR, amounts to computing the expectation of the portion of the loss that exceeds a certain threshold.

In practice, computing the loss distribution or the expected loss is quite challenging. This is because, in general, there exists no closed-form formula for computing such quantities. (Closed-form formulas exist only in some special cases, such as the case of large homogeneous portfolio; see Vasicek (1991).) Monte Carlo simulation, while generally applicable, is not recommended because it produces estimates with large variance. (This is due partly to the extreme-but-rare nature of default events.) This gives rise to the need for alternative, semi-analytical methods that are more computationally effcient and suitable for large non-homogeneous portfolios.

In addition to providing an overview of the theoretical framework that underlies the modeling of credit portfolios, this article aims at reviewing and analyzing two widely-used numerical methods for computing the loss distribution and the expected loss, providing further recommendation of how to effciently implement the methods. Both methods are based on the *Conditionally Independent Default* (CID) framework (which will be described in section 2.2 below). The first method, referred to as the *recursive* method, is developed by Andersen et al. (2003) and Hull et al. (2004). The second method, referred to as the *transform* method, makes use of the classical theory of Fourier transform and is applied to the context of portfolio credit risk by researchs such as Burtschell et al. (2009), Gregory et al. (2003), Gregory et al. (2004), Glasserman et al. (2012).

This article seeks to present the different ways in which each of these two methods can be implemented. While existing literature typically focuses on applying the methods to compute probability distribution as the primary quantity of interest, this paper will propose a variation of the methods customized for computing expected loss. We will also explain in detail why the version that we propose is beneficial in implementation.

Owing to the theoretical nature of this article, certain parts of the article will be mathematically oriented. For example, the article occasionally deals with concepts in probability and complex analysis. However, we will review these mathematical concepts in a manner that is self-contained, requires minimal prior background, and is tailored specifically to the context of credit portfolios. Much care has been taken to make sure that notations are kept simple and intuitive.

The rest of the article is organized as follows. In Section 2, we describe the mathematical modeling of credit portfolios, review the literature on CID models, and establish the connection between our two quantities of interests: the probability distribution and the expectation of the loss. In Section 3, we review the widely-used recursive method for computing the loss distribution. We point out the technical issues inherent in the method, and then suggest a modification that resolves those issues. Section 4 overviews the Fourier transform approach, and discusses the challenges that exist in practice. Then we suggest a version of the Fourier transform approach that alleviate those challenges. Section 5 gives the conclusion.

2. Mathematical model of portfolio loss

2.1 Mathematical Representation

The modeling of a credit portfolio starts with a mathematical representation of the portfolio's loss. Consider a portfolio of N credit instruments. For each instrument j, let c_j denote its loss-given-default (assumed known). The portfolio loss, which we will denote by L, is the sum of the default loss of its N constituents, namely:

$$L = \sum_{j=1}^{N} c_j Y_j$$

where the random variable Y_j, called the default indicator of obligor j, is equal to 1 if the jth obligor default and 0 if it does not. Therefore, when $Y_j = 0$, obligor j does not contribute any loss to the portfolio loss L. But when $Y_j = 1$, obligor j contributes a loss of c_j to the portfolio. Let $p_j = P(Y_j = 1)$ denote the default probability of obligor j. We assume that c_j and p are known.

2.2 Conditionally Independent Default

To capture default correlations, one is required to specify the correlation structure among Y_j's. When the number of obligors N is large, the correlation structure among Y_j's can be very complex. In this article, however, we shall restrict our attention to a widely-adopted class of default correlation models known as *Conditionally Independent Default (CID)* models, which can greatly simplify the correlation structure among defaults even when the number of obligors is large.

To understand the CID model, let us first consider, as an example, a sub-class of CID models known as the *one-factor copula models*. Such models include the well-known *one-factor Gaussian copula model* proposed by Li (2000), which has since become the industry standard, as well as various other variations such as *t*-copula, Marshall-Olkin, Clayton, etc. (For reference, see, for example, Laurent et al. (2003), Gregory et al. (2004)). In a one-factor copula model, correlation is assumed to be governed by a single random variable, say, Z, typically referred to as the "market factor". The type of copula specifies the range of values that Z can take and its probabilistic distribution. (In the Gaussian copula case, for example, the market factor can take any values between $-\infty$ and ∞ and has Gaussian distribution.) But when Z is fixed at a particular value, obligors Y_j's are assumed to be independent; the conditional default probability, denoted by $P(Y_j = 1|Z)$, is assumed to be a known function of Z. Because the default probabilities of all obligors are affected by Z, the random variable Z in effect produces the correlation among defaults.

Extensions of the CID models that include many factors are referred to as multi-factor models, studied in works such as Glasserman et al. (2007) and Glaserman et al. (2012). A more advanced class of CID models assume that correlation is governed by a *process* rather than a single random variable. In such case, obligors become independent conditioned on the *path* of that process. Examples include Schoesser et al. (2009), Brunlid (2006), as well as the CID *intensity-based models* discussed in Duffie et al. (2001) and Mortensen (2006). The analysis in this article, while focused on the one-factor copula case, can be extended to these more advanced settings as well.

2.3 Quantities of interest

In this section, we highlight some important quantities that we want to compute. The first quantity, which is fundamental in the analysis of credit portfolios, is the probability of the form:

$$P(L \leq x) \;\;,\;\; x \in R \tag{2.1}$$

This is the probability that the loss does not exceed a threshold x. When considered as a function of x, (2.1) is typically referred to as the *cumulative distribution function* (cdf) of L. (The cdf of a random variable is commonly used to describe the distribution of the random variable.) Another related probability is the "tail probability" of the form:

$$P(L \geq x) \;\;,\;\; x \in R \tag{2.1}$$

This is the probability that the portfolio loss is x or above. The relationship between (2.1) and (2.2) is given by $P(L \geq a) = 1 - \lim_{x \uparrow a} P(L \leq x)$. The tail probability (2.2) is important in risk management. For example, computing the *value-at-risk* (VaR), which is defined as a quantile of the distribution of L, inevitably requires computing the tail probability (2.2).

Many applications require one to compute not only the probabilistic distribution of the portfolio loss but also the expected amount of loss. More specifically, one is often required to compute expectation of the form:

$$E(L \;\mathrm{L}\; x) \;\;,\;\; x \in R \tag{2.3}$$

where $L \;\Lambda\; x = \min\{L, x\}$; that is to say, $L \;\Lambda\; x$ takes the value of either L or x, whichever is smaller. The quantity $L \;\Lambda\; x$ represents the amount of portfolio loss that does not exceed x. Therefore, the quantity (2.3) above represents not the expectation of the entire portfolio loss L but rather the expectation of *only the portion of L that does not exceed x*. In the context of CDOs, this is the expected loss of the (equity) tranche which absorbs the first x dollar of default losses. As for a CDO tranche with lower attachment point a and upper attachment point b, it can be shown that the expected loss is $E(L \;\Lambda\; b) - E(L \;\Lambda\; a)$. Pricing and risk management of CDOs, therefore, require computing expectation of the form (2.3).

A closely related expression is:

$$E(L - x)^+ \ , \ x \in R \tag{2.4}$$

This is the expected loss in excess of x (in other words, the expected loss absorbed by the tranche). It can be shown that $E(L \ \Lambda \ x) + E(L - x)^+ = E[L] = \Sigma_{j=1}^{N} p_j c_j$. Because of this relationship, one can easily retreive (2.3) from (2.4), and vice versa. Apart from its relevance in CDO pricing, the quantity (2.4) is also closely linked to the concept of *expected shortfall*, which is defined as the expected loss in excess of a certain threshold and is used along with the value-at-risk for the purpose of risk management.

As an emphasis of this article, we will now argue that the probabilities (2.1-2.2) and the expectations (2.3-2.4) are, in fact, dual quantities, in the sense that the knowledge of one implies the knowledge of the other. To see this, note the fact[2] that

$$E(L - x)^+ = \int_x^\infty P(L \geq x)\, dx \ \text{, or } \ P(L \geq x) = -\frac{d}{dx} E(L - x)^+ \tag{2.5}$$

where the second equality holds for every x at which $E(L - x)^+$ is differentiable. Identities (2.5) complete the linkage among (2.1), (2.2), (2.3), and (2.4); the knowledge of any one of them (for all x) implies the knowledge of the others.

In practice, computing the loss distribution (2.1-2.2) and the expected loss (2.3-2.4) is quite challenging, because there exists no closed-form formula in general. For the special case of homogeneous portfolio, in which all obligors have the same probability of default and loss-given-default, Vasicek (1991) has derived a closed form formula for (2.2) for a granula (infinitely large) portfolio under the 1-factor Gaussian copula correlation structure. But for the general, more realistic case of non-homogeneous portfolio, computing (2.1-2.4) is more challenging. The rest of the papers detail the various methods for computing such quantities, along with our recommendation of how to implement them.

3. The Recursive Method

The first method that we will discuss is the recursive method, as appeared Andersen et al. (2003) and Hull et al. (2004). We review the method in Section 3.1 below, and point out the challenges behind the implementation. Then, in Section 3.2, we present a variation of the method that addresses those challenges.

[2] The outline of the proof is as follows. First, note the identity $(L = x)^+ = \int_x^\infty 1\{L \geq u\}\, du$, where $1\{\cdot\}$denote the indicator function. Taking the expectation on both sides yield the first part of (2.5). which in turn implies the second part.

3.1 Probability Bucketing

The idea behind the method presented in Hull et al. (2004) is as follows.

Let us first consider the case of independent defaults. Let $0 = x_0 < x_1 < x_2 < \cdots$ be a sequence of increasing real numbers. Consider a portfolio of n obligors. Let $w_k := P(x_{k-1} < L \leq x_k)$ denote the probability that the default loss of the portfolio consisting of obligors takes value in the k th "bucket" (x_{k-1}, x_k) . We assume further that the probability w_k is concentrated at a point $a_k \in (x_{k-1}, x_k)$. Therefore, the loss distribution can be represented by the set of pairs $\{(a_k, w_k) : k = 1, 2, \ldots\}$, which can be visualized as a bar chart in which vertical bars of height w_k's are placed at points a_k's on the horizon axis. (This plot is similar in concept to the *probability mass function* of L.) In this setting, the cdf (2.1) is given by $P(L \leq x) = \Sigma_{k : ak \leq x} W_k$.

The method attempts to determine what will happen to the loss distribution if one more obligor is added into the portfolio. Let $\tilde{p}$ and $\tilde{c}$ denote the default probability and the loss-given-default of the new obligor. The loss distribution, after the new obligor is added, will be a combination of the following two cases. With probability $1 - \tilde{p}$, the new obligor does not default and the portfolio incurs no additional loss. So the loss distribution for this case is the same as before, but with the w_k's adjusted by the chance that the new obligor does not default, i.e., $\{(a_1, (1 - \tilde{p})w_1), (a_2, (1 - p)w_2), \ldots\}$. But if the new obligor defaults, the portfolio loss will increase by $\tilde{c}$ and the loss distribution, after adjusted by the default probability $\tilde{p}$, will shift to the right by $\tilde{c}$, i.e., $\{(a_1, + \tilde{c}, \tilde{p}w_1), (a_2, + \tilde{c}, \tilde{p}w_2), \ldots\}$. Combining these two cases, the new loss distribution, denoted by $\{(\tilde{a}_1, + \tilde{w}_1), (\tilde{a}_2, \tilde{w}_2), \ldots\}$, is obtained as follows:

$$\tilde{w}_k = (1 - \tilde{p})w_1 + \tilde{p}w_j \tag{3.6}$$

$$\tilde{a}_k = (1 - \tilde{p})w_k a_k / \tilde{w}_k + \tilde{p}w_j (a_j + \tilde{c}) / \tilde{w}_k \tag{3.7}$$

where j is such that $a_j + \tilde{c} \in (x_{k-1}, x_k)$ The above is the result of combining $(a_k, (1 - p)w_k)$ from bucket k of the non-default case, and $(a_j + \tilde{c}, \tilde{p}w_j)$ that gets "shifted" from bucket j of the default case. For $\tilde{w}_k$, we combine the two cases simply by adding the two probabilities $(1 - \tilde{p})w_k$ and $\tilde{p}w_j$. Although these two probabilities are not concentrated at the same point (one at a_k and the other at $a_j + \tilde{c}$), the method makes a rough approximation by "lumping" the probabilities together at a single point $\tilde{a}_k$, which is a weighted average of a_k and $a_j + \tilde{c}$.

Equations (3.6-3.7) enable one to compute the loss distribution of a portfolio of any size. By starting with an empty portfolio and successively adding one obligor at a time, one can iterate (3.6-3.7) to arrive, ultimately, at the loss distribution of the full portfolio.

We can extend this method to the case of correlated defaults under the CID model. Under the CID model with one market factor Z, for example, the tail probability can be computed by

$$P(L \geq x) \;=\; \int_{-\infty}^{\infty} P(L \geq x|Z = z)\, f\,(z)\; dz$$

Because of the CID property, the conditional probability in the integrand can be computed using the method described above.

To complete the discussion of the method, let us note a few issues that require some care. First, the above method is valid when all buckets are of equal size. This ensures index j in (3.6-3.7) is uniquely defined. In practice, one may want to use finer partition (smaller buckets) for parts of the distribution that we are particularly interested in (such as the tail of the distribution), and use coarser partition (bigger buckets) for the other parts. In such a case of non-uniform bucket size, the indexing in (3.6-3.7) needs to be modified, and the resulting formulas will become slightly more complex.[3]

A more serious drawback of this method is attributed to the *discretization error*. As mentioned, the probability mass $\tilde{w}_k$ in bucket k can, in actuality, be scattered at many points in the interval $(x_{k-1}, x_k]$. But because the method "lumps up" the whole probability $\tilde{w}_k$ at a single point $\tilde{a}_k$, it loses some information about the true distribution and, as a result, produces an approximation error. This discretization error gets more pronounced when bucket size is large and when more and more obligors are added to the portfolio.

These problems prompt us to seek an alternative method in the next section.

3.2 Suggested Alternative Method

To address the problems inherent in the method of the previous section, in this section we recommend an alternative procedure in computing the portfolio loss. The proposed method, and its benefits over the traditional method, will be discussed presently.

---

[3] More specifically, if we let B_k denote the set of all j's such that $a_j + \tilde{c} \in (x_{k-1}, x_k]$, then the update formulae should be modified to :

$$\tilde{w}_k \;= (1 - \tilde{p})w_k + \tilde{p} \sum_{J \in B_k} w_j$$

$$\tilde{a}_k \;= (1 - \tilde{p})w_k a_k \tilde{w}_k + \tilde{p} \sum_{J \in B_k} w_j (a_j + \tilde{c})\tilde{w}_k$$

Similar to the previous section, our method constructs the portfolio loss by recursively adding obligors one by one. But instead of recursively computing the *probabilities* (w_k's), we recursively compute the *expected loss*. To describe the method in more detail, first consider the case of independent defaults. Let $h(x) := E(L \wedge x)$, where is a portfolio loss containing n obligors, and $x \in R$. Assume that $h(x)$ is known for $x = x_1, x_2, \ldots$. Now, suppose another obligor, with default probability $\tilde{p}$ and loss-given-default $\tilde{c}$, is added to the portfolio. Let $\tilde{L}$ denote the loss of the new portfolio (which contains $n + 1$ obligors), so that $\tilde{L} = L + \tilde{c}$ if the new obligor defaults (with probability $\tilde{p}$), and $\tilde{L} = L$ otherwise. Therefore, if we let $\tilde{h}(x) = E(\tilde{L} \wedge x)$, then

$$\tilde{h}(x_j) = E(\tilde{L} \wedge x_j) = (1 - \tilde{p})E(L \wedge x_j) + \tilde{p}E[(L + \tilde{c}) \wedge x_j]$$

$$(1 - \tilde{p})h(x_j) + \tilde{p}\left[\tilde{c} + E[L \wedge (x_j - \tilde{c})]\right] \quad (3.7)$$

where the last equality follows from the fact that $(L + \tilde{c}) \wedge x_j = \tilde{c} + [L \wedge (x_j - \tilde{c})]$ Since $h(x) = E(L \wedge x)$ is a continuous function of x, we can approximate $E[L \wedge (x_j - \tilde{c})]$ by linearly interpolating between $h(x_{\ell-1})$ and $h(x_\ell)$, where ℓ is such that $x_{\ell-1} < x_k - \tilde{c} \leq x_\ell$. Therefore, (3.10) becomes

$$\tilde{h}(x_k) = (1 - \tilde{p})h(x_k) + \tilde{p}\left[\tilde{c} + \frac{x_\ell - x_k + \tilde{c}}{x_\ell - x_{\ell-1}} h(x_{\ell-1}) + \frac{x_k - \tilde{c} - x_{\ell-1}}{x_\ell - x_{\ell-1}} h(x_\ell)\right] \quad (3.11)$$

Similarly to the previous section, if one starts with an empty portfolio (for which the expected loss is zero) and successively adds one obligor at a time, then iterative applications of (3.11) will enable one to obtain the value of $E(L \wedge y)$ where L is a portfolio of any size.

If defaults are correlated under the CID structure, then we use the fact that obligors becomes independent upon conditioning on a value of the market factor $Z = z$, and use the above recursion to compute the conditional expected loss $E(L \wedge x|Z = z)$. The unconditional expected loss can then be obtained by

$$E(L \wedge x) = \int_{-\infty}^{\infty} E(L \wedge x|Z = z)\, f(z)\, dz$$

where the integral is approximated by numerical integration.

Since the recursion (3.11) computes the expected loss without having to compute the probability distribution first, it is more appropriate for pricing basket credit derivatives (or for computing expected shortfall). And thanks to the relationship (2.5), the expected loss from recursion (3.11) can be used to imply the loss distribution, if so desired.

Let us now discuss other important advantages of this method. Recall that (3.6-3.7) are valid only for the case where x_1, ..., x_k are equally-spaced (equal bucket size). The formula (3.11), on the other hand, is valid regardless of whether or not the buckets are equally-spaced. More importantly, in contrast to the method in the previous section, recursion (3.11) uses linear interpolation to alleviate the discretization error, taking advantage of the fact that $E(L \Lambda x)$ is continuous in x. In fact, it can be shown that $E(L \Lambda x)$ is piece-wise linear in x, hence there are instances in which the linear interpolation produces no approximation error at all. This is the main reason why it is preferable to use the recursive method to build up the portfolio's expected loss, rather than the loss distribution.

3.3 Numerical Demonstration

To demonstrate the accuracy of the suggested method, consider the following numerical example. Consider a credit portfolio with 125 obligors. (This is the typical number of obligors in index portfolios such as the iTraxx.) Assume that the default probabilities of the obligors are uniformly distributed between 2% and 5%, and the losses-given-default (i.e., exposure less recoverable amounts) are between 0 to 5 (million). The correlation structure among obligors is assumed to follow a one-factor Gaussian copula with correlation 0.5. For this sample portfolio, the risk measures obtained from using the above methods are shown below. Here, we set the bucket size x_k - x_{k-1} to be equal to 1.

		Bucketing Method (Section 3.1)	Suggested Method (Section 3.2)	Actual Value
90%	- Value-at-Risk	30.4	31.1	31.2
	- Expected shortfall	67.6	68.1	68.1
95%	- Value-at-Risk	54.0	55.0	54.7
	- Expected shortfall	94.4	95.1	95.1
99%	- Value-at-Risk	121.1	122.0	121.6
	- Expected shortfall	157.2	160.3	160.2

Table 1 Shows the values-at-risk and the expected shortfall at three probability levels. The column 'Actual Values' displays the true value calculated as the limit as the bucket size goes to zero.

As can be seen, the suggested method, in addition to being easier to implement (as explained in Section 3.2), is more accurate overall. While both methods give roughly the same VaRs, it should be noted that the suggested method gives noticeably better result for the expected shortfall. This is because the suggested method is designed to take advantage of the continuity property of the expected loss function (2.3-2.4). Therefore, even with large bucket size of 1, the suggested method gives remarkably accurate results. To confirm this point, the following plot shows the rate of convergence of the two methods as the bucket size decreases.

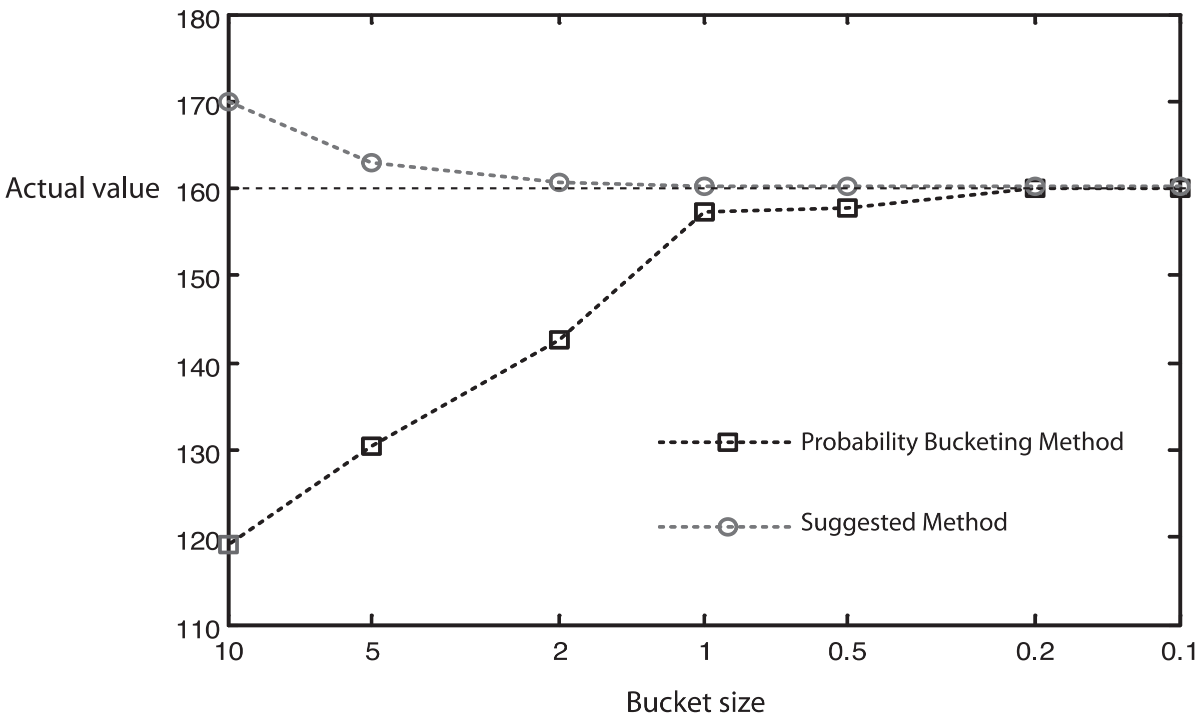

Figure 1 Compare the rate of convergence of the two methods as the bucket size decreases. The vertical axis shows the expected shortfall at 99%. The horizontal axis shows the bucket size on the logarithmic scale.

As seen, the suggested method produces better approximation for the actual value of the expected shortfall.

4. Fourier Transform Methods

This section discusses the transform method for analyzing portfolio of credit instruments. A concise review of the general method of Fourier transforms is given in Section 4.1. In Section 4.2, we use the Fourier framework to derive a formula for computing the loss distribution, and then point out some problematic issues of the approach. Section 4.3 then suggests a different approach that resolves those issues.

4.1 Mathematical Background on Fourier Transforms and Inversions

A comprehensive background on the transform method can be found in Abate et al. (1992) and Abate et al. (1995). Here, we provide a concise review of the theory relevant to our analysis.

Let $g(x)$ be a complex-valued function. Assume that g is integrable, i.e., $\int_{-\infty}^{\infty} |g(x)| dx$ is is finite. For any real number ω, define $\psi(\omega)$ as:

$$\psi(\omega) := \int_{-\infty}^{\infty} e^{i\omega x} g(x)\, dx$$

where $i = \sqrt{-1}$ is the imaginary unit. The complex-valued function ψ is called the Fourier transform of g. The value of g can be retrieve from ψ using the so-called *Fourier inversion formula*:

$$g(x) := \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-i\omega x} \psi(\omega)\, d\omega \qquad (4.12)$$

In many applications, the quantity of interest is a function $g(x)$ that is difficult to compute directly, but its Fourier transform turns out to be relatively easy to compute. In such situations, equation (4.12) allows one to use the Fourier transform ψ to retrieve the quantity of interest $g(x)$. For the special case where g is real and defined only on $x \geq 0$, the inverse formula can be shown to simplify to:

$$g(x) = \frac{2}{\pi} \int_{0}^{\infty} \cos \omega x \,\mathrm{Re}\, \psi(\omega)\, d\omega = \frac{2}{\pi} \int_{0}^{\infty} \sin \omega x \,\mathrm{Im}\, \psi(\omega)\, d\omega \qquad (4.13)$$

Since the portfolio loss is a real, nonnegative random variable, we shall base the following analysis on this simplified inversion formula.

4.2 Computing Loss Distribution using Fourier Inversion Method

Fourier transform method has been widely used to analyze the loss portfolio (see Burtschell et al. (2009), Laurent et al. (2003), Gregory et al. (2004), Glasserman et al. (2012), etc.) The most direct application of this techniques is to compute the tail probability $P(L \geq x)$ that defines the distribution of the portfolio. The technique makes use of the observation that, while $P(L \geq x)$ usually cannot be computed directly by an analytical formula, its Fourier transform can be computed quite easily. It can be shown that the Fourier transform of $P(L \geq x)$ is

$$\int_0^\infty e^{-i\omega x} P(L \geq x)\, dx = \frac{\phi(\omega) - 1}{i(\omega)} \tag{4.14}$$

where $\phi(\omega) := \mathrm{E}[e^{i\omega L}]$ is the so-called *characteristic function* of L, which can be computed easily as follows. If defaults are independent, $\phi(\omega)$ is given by:

$$\phi(\omega) = \mathrm{E}[e^{i\omega L}] = \prod_{j=1}^{N} \mathrm{E}[e^{ic_j Y_j}] = \prod_{j=1}^{N} (1 - p_j + p_j e^{i\omega c_j}) \tag{4.15}$$

The second-to-last equality follows from the independent assumption, and the last equality makes use of the observation that $\mathrm{E}[e^{ic_j Y_j}] = 1 - p_j + p_j e^{i\omega c_j}$. If defaults are not independent but has a CID structure with one market factor Z, one can compute $\phi(\omega)$ by

$$\phi(\omega) = \int_{-\infty}^{\infty} \mathrm{E}[e^{i\omega L} | Z = z]\, f(z)\, dz$$

Since obligors become independent conditioned on $Z = z$, the expectation in the integrand above can be computed by (4.15) with p_j replaced by the conditional default probability.

Once the Fourier transform (4.14) is obtained, we can use the inversion formula to compute $P(L \geq x)$. The knowledge of $P(L \geq x)$ implies the distribution of the portfolio loss, which one can subsequently use to compute the expected loss.

4.3 Suggested Procedure for Computing Expected Loss

We now present a more direct procedure for computing the expected loss. Rather than evaluating the inversion integral to obtain the *probability distribution* of the portfolio loss, we derive an inversion integral for the *expectation* of the portfolio loss instead. We now describe the method.

Our method begins by deriving the Fourier transform of : $E(L - x)^+$:

$$\int_0^\infty e^{i\omega x} E(L - x)^+ \, dx = \frac{E[L]}{i\omega} + \frac{1}{i\omega}\int_0^\infty e^{i\omega x} P(L \geq x)\, dx = \frac{1 - \phi(\omega) + i\omega E[L]}{\omega^2} \tag{4.16}$$

Here, the first equality results from integrating by parts and then using (2.5). The second equality follows by substituting (4.14). By taking the real part of (4.16) and applying inverse formula (4.13), one obtains the following formula for computing the expected loss $E(L - x)^+$:

$$E(L - x)^+ = \frac{2}{\pi}\int_0^\infty \cos \omega x \, \frac{1 - \operatorname{Re} \phi(\omega)}{\omega^2}\, d\omega$$

$$= E[L] - \frac{2}{\pi}\int_0^\infty (1 - \cos \omega x)\, \frac{1 - \operatorname{Re} \phi(\omega)}{\omega^2}\, d\omega \tag{4.17}$$

where the second inequality follows from the fact that $\frac{2}{\pi}\int_0^\infty \frac{1 - \operatorname{Re} \phi(\omega)}{\omega^2}\, d\omega = E[L]$ (this fact can be easily proved by letting $x = 0$ in the above equation). The last integral in (4.17) is preferable to the integral preceding it because the last integrand is more stable at $\omega = 0$. (This is because, unlike $\cos \omega x/\omega^2$, the value of $(1 - \cos \omega x)/\omega^2$ is bounded around the neighborhood of $\omega = 0$.) In the last integral (4.17), therefore, the error (if any) that one incurs when approximating the characteristic function will not be adversely affected by the factor $1/\omega^2$ in the neighborhood of $\omega = 0$.

Similarly, one can use the second part of (4.13) to show that

$$E(L - x)^+ = \frac{2}{\pi}\int_0^\infty \sin \omega x \, \frac{\omega E[L] - \operatorname{Im} \phi(\omega)}{\omega^2}\, d\omega$$

$$= E[L] - \frac{2}{\pi}\int_0^\infty \sin \omega x \, \frac{\operatorname{Im} \phi(\omega)}{\omega^2}\, d\omega \tag{4.18}$$

(Here, the second equality follows from the fact that $\int_0^\infty \frac{\sin \omega x}{\omega^2}\, d\omega = \frac{\pi}{2}$ for all x.)

With the characteristic function $\phi(\omega)$ computed in the manner described in the previous section, equations (4.17) and (4.18) provide a direct way to compute the expected loss without the need to compute the probability distribution of L. This is beneficial in situations that require only the values of expected loss such as computing expected shortfall, or pricing basket credit derivatives. (In applications that require the probability distribution of L, one can make use of (2.5).)

Although (4.17) and (4.18) are theoretically equivalent, in practice we suggest using (4.18) rather than (4.17). In computing the integrals (4.17-4.18) numerically, one unavoidably incurs the truncation error that comes from approximating the indefinite integral with a definite integral. Because the integrand in (4.17) is nonnegative (as opposed to the integrand in (4.18) which alternates signs as ω varies), truncating the integral (4.17) will result in over-estimating the true value of $E(L - x)^+$. On the other hand, in practice we find that (4.18) converges quickly to the true value of $E(L - x)^+$.

4.4 Numerical Demonstration

We now demonstrate the use of (4.17-4.18) in computing the expected shortfall of a credit portfolio. Consider the sample portfolio from Section 3.3. The expected shortfall at various probability level, computed from (4.17-4.18), are shown below. Here the integrals (4.17-4.18) are truncated at $\omega = 10$.

	Using (4.17)	Using (4.18)	Actual Value
Expected shortfall @ 90%	68.4	68.1	68.1
Expected shortfall @ 95%	95.7	95.1	95.1
Expected shortfall @ 99%	163.6	160.3	160.2

Table 2 Compare the expected shortfall at three probability levels, as computed using the inversion integral (4.17) versus using (4.18).

As seen from Table 2, while both integrals (4.17-4.18) approximate the actual value of the expected shortfall quite well, the suggested method (4.18) yields much more accurate results. The expected shortfall computed from (4.17) over-approximates the actual value because, as explained previously, the non-negativity of the integrand in (4.17) accentuates the truncation error. On the other hand, the integral (4.18) converges quite quickly to the actual value. The following plot compares the rate of convergence of both integrals.

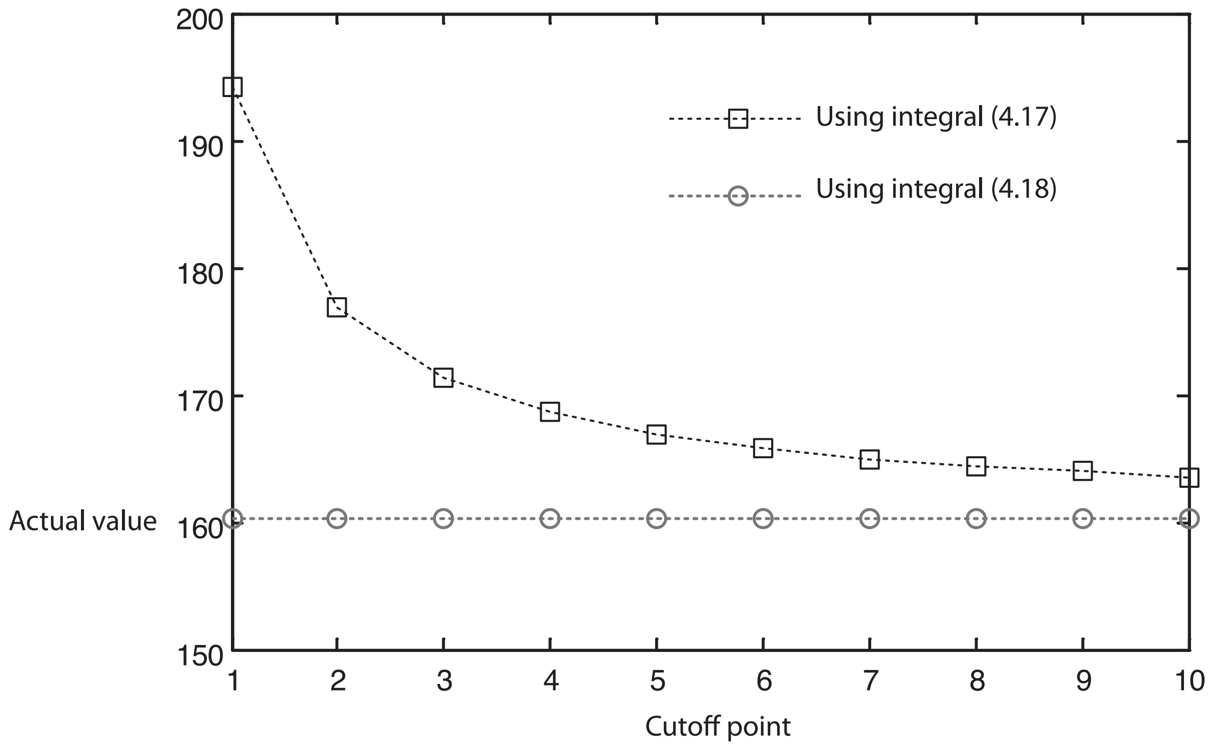

Figure 2 Expected shortfall at probability 99%, computed using (4.17-4.18). The cutoff point is the point at which the indefinite integrals (4.17-4.18) are truncated.

As seen, the suggested method (4.18) produces accurate result even for low cutoff points, while the integral (4.17) takes longer to converge to the actual value. This makes it more preferable to use (4.18) in practice, as it provides better approximation with lower computational effort.

5. Summary

The management of credit risk and the pricing of credit derivatives often require one to adopt the portfolio view. The default risk of credit instruments in a portfolio must be assessed not only on the individual basis but also in the aggregate manner as well. This article discusses the theoretical framework and the numerical techniques involved in the quantitative analysis of a portfolio of credit instruments.

The framework under discussion is the conditionally-independent default (CID) framework. In such a framework, correlation among defaults is driven by a set of market factors, conditioned upon which obligors become independent. This framework facilitates the analysis of large portfolios and has been widely adopted as the industry standard.

In the quantitative analysis of a credit portfolio, the key quantities of interest are the distribution and the expectation of the amount of total loss that results from the defaults of the portfolio's constituents. This article emphasizes the close relationship between the two quantities: The expected loss of any portion of the portfolio can be implied if one knows the distribution of the loss amount. Conversely, the loss distribution can be implied if one knows the expected loss of any portions of the portfolio.

Computing the loss distribution and the expected loss can be challenging because, in general, there exists no closed-form formula with which one can compute those quantities. This article analyzes two widely-used numerical methods for computing the portfolio loss. The recursive method reconstructs the portfolio loss by successively adding one constituent at a time to the portfolio. The transform method computes the portfolio loss indirectly through the inversion of its Fourier transform.

In addition to reviewing the methods, we point out the challenges in their implementation and provide alternatives that address those challenges. The literature traditionally focuses on computing the probability distribution of the portfolio loss, leaving the expected loss to be implied from the loss distribution. The suggested methods in this article, in contrast, are tailored to compute the expected loss in a given portion of the portfolio. Because the suggested methods compute the expected loss without having to compute the loss distribution first, they are particularly attractive for the purpose of pricing. The loss distribution, if desired, is shown to be easily retrievable from the knowledge of the expected loss of different portions of the portfolio. More importantly, the suggested methods take advantage of certain features of the problem to alleviate numerical issues related to estimation and discretization. We provide a detailed description of the suggested methods and discuss the benefits that such methods bring to implementation.

References

[1] Abate, J., & Whitt, W. (1992). The Fourier-Series Method for Inverting Transforms of Probability Distributions. Queueing Systems: Theory and Applications, 10(1–2), 5-88.

[2] Abate, J., & Whitt W. (1995). Numerical Inversion of Laplace Transforms of Probability Distributions. ORSA Journal of Computing, 7, 36-43.

[3] Andersen, L., Basu, S., & Sidenius J. (2003). All Your Hedges in One Basket. Risk, November, 67-72.

[4] Brunlid, H. (2006). A comparative analysis of hyperbolic copulas induced by a one factor Levy model. Working paper.

[5] Burtschell, X., Gregory, J., & Laurent, J.-P. (2009). A comparative analysis of CDO pricing models under the factor copula framwork. Journal of Derivatives, 16, 9-37.

[6] Duffie, D., & Garleanu, N. (2001). Risk and valuation of collateralized debt obligations. Financial Analysts Journal, 57, 41-59.

[7] Glasserman, P., & Suchintabandid, S. (2007). Correlation Expansion for CDO pricing. Journal of Banking & Finance, 31, 1375-1398.

[8] Glasserman, P., & Suchintabandid, S. (2012). Quadratic Transform Approximations. SIAM Journal of Financial Mathematics, 3(1), 137-162.

[9] Gregory, J., & Laurent, J.-P. (2003). I Will Survive. RISK, June, 103-107.

[10] Gregory, J., & Laurent, J.-P. (2004). In the Core of Correlation. RISK, October, 87-91.

[11] Hull, J. & White, A. (2004). Valuation of a CDO and an nth to default CDS without Monte Carlo simulation. The Journal of Derivatives, 12(2), 8-23.

[12] Laurent, J.-P., & Gregory, J. (2003). Basket default swaps, CDOs and factor copulas. Journal of Risk, 7(4), 103-122.

[13] Li, D. (2000). On default correlation: A copula function approach. Journal of Fixed Income, March, 115–118.

[14] Mortensen, A. (2006). Semi-analytical valuation of basket credit derivatives in intensity-based models. Journal of Derivatives, 13, 8-26.

[15] Schoesser, A., & Zagst, R. (2009). Crash-NIG copula model: Regime switching credit portfolio modeling through the crisis. Working paper.

[16] Vasicek, O. (1991). Limiting Loan Loss Distribution. KMV Corporation.