Journal Of Investment Management, Vol. 14, No. 2, (2016), pp. 49–70

www.joim.com

PORTFOLIO DIVERSIFICATION IN CONCENTRATED BOND AND LOAN PORTFOLIOS

Paul Kupiec[a,1]

I develop an algorithm to approximate the loss rate distribution for fixed income portfolios with obligor concentrations. The approximation requires no advanced mathematics or statistics, only the summation of large exposures and the evaluation of binomial probabilities. The approximation is model-independent and can be used after removing default dependence using any risk modeling approach. It is especially useful for capital calculations given its inherent accuracy in the upper tail of the cumulative portfolio loss rate distribution. The approximation provides a simple way to calculate the capital benefits of risk mitigation or the capital needed when a marginal credit is added to a concentrated portfolio.

1 Introduction

Compound interest and risk diversification, if not among the most powerful forces in nature, are still perhaps the two most important forces in finance.[1] The modern theory of portfolio diversification began when Markowitz (1952) emphasized the importance of efficient mean-variance portfolios for investment management. Markowitz's insights lead to the development of Sharpe's (1964) Capital Asset Pricing Model, the first equilibrium model that links risk and expected return.

Despite the fundamental importance of diversification, it took almost 50 years after Markowitz's original insights before formal diversification techniques were adopted to manage high-quality loan and bond portfolios. For example, according to Altman and Saunders (1998), " [*While*] one might expect that these very same [*Markowitz*] techniques would (and could) be applied to the fixed income area . . . there has been, however, very little published work in the bond area and a recent survey of practices by commercial banks found fragmented and untested efforts." (p. 1728)

There are many reasons why fixed income managers were slow to adopt formal portfolio

[a]Resident Scholar, the American Enterprise Institute, Washington, D,C., USA. Phone: 202-862-7167. E-mail: paul.kupiec@aei.org

[1]The views in this paper are those of the authors alone. They do not represent the official views of the American Enterprise Institute.

diversification models. One is that it is not so obvious how diversification works for fixed income investments, given the abbreviated nature of their positive return tails. Moreover, fixed income investments often are not actively traded and most lack the return histories necessary to construct Markowitz efficient portfolios. Finally, fixed income investments tend to be discrete, meaning that they come in prepackaged sizes that may be large and not easily disaggregated and traded. This discrete, illiquid nature makes it inherently difficult and expensive to diversify a portfolio of fixed income claims and consequently many credit portfolios contain obligor concentrations—large unbalanced exposures to a borrower or multiple borrowers. Obligor concentrations can significantly reduce portfolio diversification.

In this paper, I develop a simple algorithm to approximate the loss rate distribution of a fixed income portfolio with obligor credit concentrations. The intuition that underlies the approximation is easy to understand and the approximation calculations require no advanced mathematics or statistics—only the summation of a portfolio's largest loss exposures and an evaluation of binomial probabilities. Unlike the so-called "granularity adjustment" approach for measuring concentration risk, this approximation is not model-dependent.[2] It can be used after removing obligor default dependence using any risk modeling approach.

In closely related work on CDO and CDS pricing, Hull and White (2004) develop a "probability bucketing" algorithm to approximate a credit portfolio loss rate distribution. The Hull and White approach is also model independent and accounts for obligor concentration risk, but it involves substantially more computation than the approach proposed in this paper. My approximation also provides a simpler method for calculating the "value-at-risk" capital benefit of a risk mitigation strategy or the capital increment that is required when a new obligor is added to a concentrated credit portfolio.

The approximation results show that, to a very close approximation, the value-at-risk capital required to fund a portfolio is given by the sum of the q-largest portfolio exposures, where the portfolio exposure of an individual obligor is defined as the product of the credit's loss given default and exposure at default, normalized by the total portfolio exposure at default. The number of large exposures, q, that must be added up to calculate the portfolio capitalization rate is determined by the desired capital coverage rate (e.g. 95 or 99 percent), the credits' correlation and unconditional probabilities of default, and the binomial cumulative probability distribution. The approximation is very accurate for typical value-at-risk capital coverage rates.

This paper is organized as follows. Section 2 provides an abbreviated overview of the development of formal portfolio diversification models for high-quality fixed income portfolios including the granularity adjustment approach for measuring obligor concentration risk and the Hull and White approach. Section 3 reviews the structure of the Vasicek (1987, 1991) model for measuring default risk diversification including the so-called asymptotic single-factor model. Section 4 discusses concentration risks that arise in finite portfolios of obligors with uniform risk and exposure characteristics. Section 5 derives the portfolio loss rate distribution when there is obligor concentration risk generated by varying obligor exposure or loss rate characteristics. For portfolios of even moderate size, the calculation of the exact portfolio loss rate distribution may be impractical because of its demands on computing capacity. Section 6 introduces the approximation algorithm which is easily computed even for a very large number of obligors. Section 7 uses the approximation to construct a value-at-risk style capital requirement for the marginal credit in a portfolio with obligor concentrations. Section 8 summarizes the paper's findings.

2 Background

The intuition behind portfolio diversification with stock returns is simple. By distributing invested funds among a broad set of individual stocks with less than perfectly correlated returns, unexpectedly large positive returns on some stocks will tend to offset unexpectedly large negative returns on others. Consequently, the overall return variation on a well-diversified portfolio will be smaller than the variation of the return on a portfolio with fewer more concentrated holdings. But when it comes to understanding the mechanics of diversification in a portfolio of high-quality loans or bonds, the intuition underlying stock return diversification falls short.

Unlike stocks, upside payoffs on bonds and loans are capped. Performing credits do not provide an outsized gain to offset the large losses generated by defaulting credits. Moreover, the covariance terms needed to construct Markowitz mean-variance efficient portfolios are not easily estimated for loans and bonds. Unlike stocks, most credit claims do not actively trade, and when they are traded, day-to-day return variation must be parsed among multiple causes including a changing term structure of default free interest rates, variation in the market-wide default risk premium and changing expectations for the performance of individual credits. Modeling diversification for a fixed income portfolio requires a framework that can recognize the unique features of returns on bond and loan investments.

Practical approaches for measuring diversification in credit portfolios began with Vasicek (1987, 1991). Vasicek (1991) formulates a single common factor approach for modeling default correlations and shows that this structure becomes especially parsimonious for a so-called asymptotic portfolio, a portfolio with an infinite number of obligors with identical risk and exposure characteristics. Vasicek's asymptotic single-factor model was embraced by bank regulators (Gordy, 2003) and, in modified form, was eventually adopted in 2006 as an international standard for setting minimum regulatory capital requirements for internationally active banks. This so-called Basel II approach sets minimum regulatory capital requirements for banks using the Vasicek portfolio loss distribution under a specialized and restrictive set of assumptions. Bank portfolios are assumed to be comprised of infinitely many loans of identical size, with identical default probabilities, default correlations, and loss rates in default. Default correlations are assumed to be driven by a single latent common factor.

The highly restrictive assumptions of the asymptotic single common factor model greatly simplifies the computation of the portfolio loss rate distribution. However, the assumptions rule out credit risk concentration in any form. There are no outsized exposures to any single borrower and idiosyncratic default risks are assumed to be completely diversified away. The only factor driving portfolio performance is the single latent common factor that in part determines individual bond or loan defaults.

The single-factor asymptotic model's failure to recognize credit risk concentrations is a serious shortcoming. Indeed, even the official Basel II documentation states, "Risk concentrations are arguably the single most important cause of major problems in banks." Despite this ominous warning, Basel capital regulations include no formal models for analyzing credit risk concentration but instead identify concentration risk as an issue to be addressed by national supervisors on an ad hoc basis.

Various authors, including Vasicek (1991), recognized the need to measure obligor concentration risk in credit portfolios. A common approach is to treat concentration risk as a perturbation from the asymptotic portfolio's loss rate distribution function.[3] In this approach, for a given realized value of the common factor, concentration risk causes the conditional portfolio loss rate to deviate from the conditional asymptotic portfolio loss rate. The true portfolio loss rate distribution is the sum of the asymptotic portfolio loss rate distribution and a mean zero idiosyncratic loss rate distribution. The conditional composite loss rate distribution including concentration risk is approximated using a second-order Taylor series expansion around the conditional asymptotic portfolio loss rate distribution. The Taylor series approximation will differ according to the statistical properties of the specific modeling approach that is used to model default correlation.[4] Different models require different granularity adjustments. The granularity adjustment is the difference between conditional portfolio loss rate calculated using the Taylor series approximation and the conditional asymptotic portfolio loss rate.

The granularity adjustment does not appear to have been widely adopted in practice. The original adjustment never made it into the formal Basel Capital Accord because it was considered too complicated to impose as a regulation.[5] And the subsequent academic literature developing the granularity adjustment is even more complex. It requires familiarity with advanced probability theory before one can become comfortable with the intuition behind the granularity adjustment and the required calculations. The granularity adjustment is also model-dependent, and also dependent on the quantile of the loss distribution that is being evaluated. So different modeling approaches for capturing portfolio default dependence require bespoke granularity adjustment factors, and even these bespoke factors vary depending on the loss quantile of interest.

Hull and White (2004) discuss an alternative approach for approximating a credit portfolio loss rate distribution in the context of pricing CDOs and *n*th to default CDS contracts. Rather than focus on the Vasicek (Gaussian) copula model, they adopt a generalized copula model to generate default correlation. After conditioning on specific value(s) for the common copula model factor(s), they use a "probability bucketing" algorithm to approximate a conditional loss rate distribution. Similar to the approach I develop in the paper, the Hull and White bucketing algorithm can be used with any default correlation model by operating on conditional default rates. The user defines a series of loss rate buckets that span the loss space. The algorithm iterates through all the individual credits in the portfolio. At each iteration, it adds the credit's conditional default probability to an appropriate loss bucket and adjusts the probabilities in the remaining user-defined loss buckets.

The Hull and White approximation is flexible and can be used in the presence of obligor concentration risk. However, the approach requires that every portfolio credit be evaluated and assigned to a loss bucket. In other words, the entire loss rate distribution must be approximated, which involves substantially more computation than the approached proposed in this paper. Moreover, the extra work required by the Hull and White approximation is completely unnecessary for capital requirement calculations as my approximation is guaranteed to be very accurate for the upper tails of the cumulative portfolio loss rate distribution.

3 The Vasicek single-factor model of portfolio credit risk

The Vasicek single common factor model of credit diversification assumes that credits in a portfolio have an identical size, probability of default, loss given default, and default correlation. An individual credit's default is determined by the realized value of a random variable, $\tilde{V}_i$, with the following properties:

$$\begin{aligned} \tilde{V}_i &= \sqrt{\rho_V}\tilde{e}_M + \sqrt{1-\rho_V}\tilde{e}_i \\ \tilde{e}_M &\sim \phi(e_M) \\ e_{id} &\sim \phi(e_i), \\ E(\tilde{e}_i\tilde{e}_j) &= E(\tilde{e}_M\tilde{e}_j) = 0, \quad \forall\, i,\, j, \end{aligned} \tag{1}$$

where $\phi(z)$ represents the value of the standard normal density function evaluated at z. $\tilde{V}_i$ has a standard normal distribution[6] and is often interpreted as a proxy for the market value of the creditor firm. The common factor in expression (1), $\tilde{e}_M$, induces correlation among credit defaults, $\rho_V = \frac{Cov(\tilde{V}_i, \tilde{V}_j)}{\sigma(\tilde{V}_i)\sigma(\tilde{V}_j)}$. Defaults are less than perfectly correlated because each credit also has a latent idiosyncratic risk factor, $\tilde{e}_i$ that also governs the default process.

Individual obligor default

Credit i is assumed to default when $\tilde{V}_i < D_i$. The unconditional probability that credit i defaults is $PD = \Phi(D_i)$, where $\Phi(z)$ represents the value of the cumulative standard normal density function evaluated at z. All credits in a portfolio are assumed to have the same unconditional probability of default, $D_i = D, \forall\, i$. Time is not an independent factor in this model, but is implicitly recognized through the calibration of input values for PD.

An indicator function can be used to record the default status of individual credits,

$$\tilde{I}_i = \begin{cases} 1 & \text{if } \tilde{V}_i < D \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

$\tilde{I}_i$ has a binomial distribution with an expected value of $\Phi(D)$. By construction, the indicator functions of individual credits are correlated through the common factor $\tilde{e}_M$. The default indicator function for credit i conditional on a specific realized value of e_M is,

$$[\tilde{I}_i|\tilde{e}_M = e_M] = \begin{cases} 1 & \text{if } \dfrac{D - \sqrt{\rho}e_M}{\sqrt{1-\rho}} < \tilde{e}_i \\ 0 & \text{otherwise} \end{cases}. \tag{3}$$

Expression (3) shows that the conditional default threshold for credit i changes as the realized value of the common factor, $\tilde{e}_M$, changes. A positive value of e_M lowers the credit's default threshold, thereby decreasing the probability that the credit will default. A negative value of e_M increases the credit's default threshold, increasing the probability that the credit will default. The expected value of the conditional default indicator, conditioned on a specific realization of the common factor, is given by,

$$E[\tilde{I}_i|\tilde{e}_M = e_M] = \Phi\left(\frac{D - \sqrt{\rho}e_M}{\sqrt{1-\rho}}\right). \tag{4}$$

Portfolio default rate distribution

Let $\tilde{X}_N$ be the portfolio default rate on a portfolio comprised of N individual credits; that is, the proportion of credits in the portfolio that default, $\tilde{X}_N = \frac{\sum_{i=1}^N \tilde{I}_i}{N}$. $\tilde{X}_N$ is the average value of the indicator functions of credits included in a portfolio.

Portfolio loss rate distribution

Let EAD_i represent the exposure at default created by credit i (the loan balance or maturity value), and LGD_i represent the loss rate experienced should credit i default. The loss rate at default is measured relative to EAD_i.

In the asymptotic portfolio model, $EAD_i = EAD\ \forall\, i$, and $LGD_i = LGD$, $\forall i$. Under the assumptions that EAD and LGD are respectively uniform across all credits, the loss rate on a portfolio of n credits is

$$\frac{\sum_{i=1}^{n} \tilde{I}_i \, EAD \, LGD}{nEAD} = LGD \, \tilde{X}_n. \qquad (5)$$

If EAD is measured as the maturity value of a credit, then $LGD \, \tilde{X}_n$ represents the portfolio loss rate from the contract maturity value caused by portfolio defaults.[7]

3.1 The asymptotic single-factor model

Let $[\tilde{X}_N | \tilde{e}_M = e_M]$ be the proportion of n credits that default in the portfolio conditional on the realization of a specific value of the single common factor, $[\tilde{X}_N | \tilde{e}_M = e_M] = \frac{\sum_{i=1}^{N} [\tilde{I}_i | \tilde{e}_M = e_M]}{N}$. Individual credit's conditional indicator functions are uncorrelated random variables since their randomness is determined solely by idiosyncratic risk, $\tilde{e}_i$.

In an asymptotic portfolio, the number of individual credits is assumed to increase without bound, $N \rightarrow \infty$. The law of large numbers ensures that the sample average of a random sample of independently identically distributed observations converges almost surely to the expected value of the underlying distribution as the sample sizes increases without bound. Thus, in an asymptotic portfolio,

$$\tilde{X}_\infty = \lim_{N \to \infty} [\tilde{X}_N | \tilde{e}_M = e_M] \xrightarrow[a.s.]{} \Phi\left(\frac{D - \sqrt{\rho} e_M}{\sqrt{1 - \rho}}\right). \qquad (6)$$

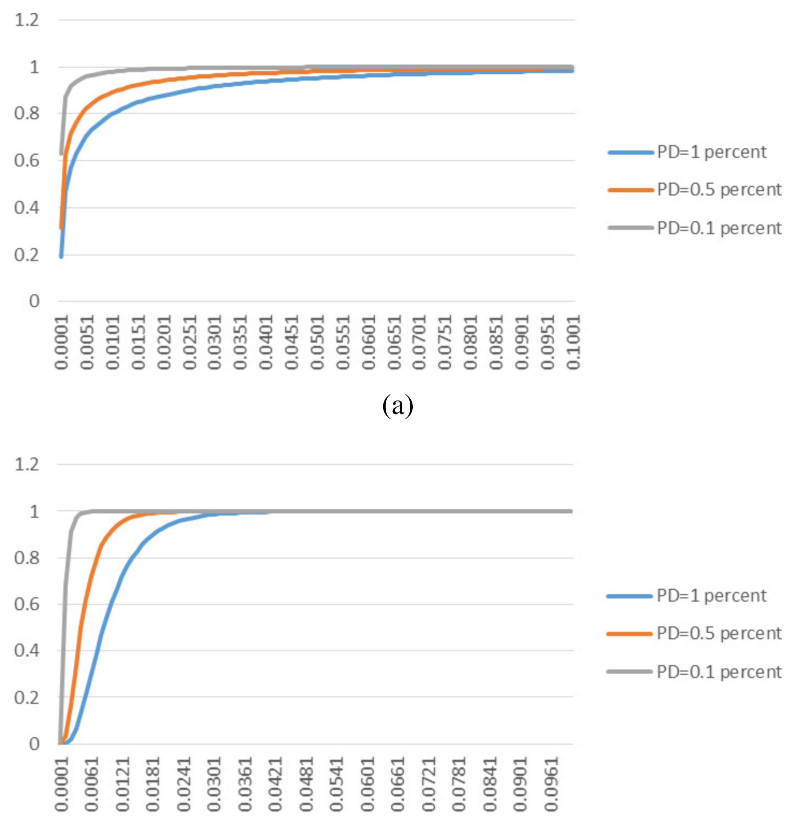

(b)

Figure 1 (a) Cumulative asymptotic default rate distribution when $\rho = 0.4$. (b) Cumulative asymptotic default rate disribution when $\rho = 0.05$.

Expression (6) shows that, in the limit, idiosyncratic default risk is completely diversified away in an asymptotic single common factor portfolio and default rate uncertainty is driven by the common market factor alone.

The unconditional distribution function of the asymptotic portfolio's default rate is given by,

$$\Pr[\tilde{X}_\infty \leq x] = \Phi\left(\frac{\sqrt{1-\rho}\,\Phi^{-1}(x) - \Phi^{-1}(PD)}{\sqrt{\rho}}\right), \quad x \in [0,1] \tag{7}$$

where Expression (7) makes use of the identity $D = \Phi^{-1}(PD)$. Expression (7) implies that an asymptotic portfolio's default rate is a random variable with a probability distribution that is determined by two parameters, the credits' unconditional default rate, *PD*, and the default correlation parameter, ρ. Figures 1(a) and 1(b) illustrate the shape of the cumulative probability distribution for the default rate of an asymptotic portfolio for selected default correlation values (ρ) and unconditional probability of default (*PD*) characteristics.

4 Idiosyncratic default risk in a finite portfolio with uniform obligor exposures

Concentration risk arises when any of the assumptions underlying the asymptotic single common factor portfolio model are violated. As a first step, I consider the implications of relaxing the assumption that the portfolio contains an infinite number of individual credits. I maintain the uniformity assumptions for *EAD*, *LGD*, *PD*, and ρ, but assume that the portfolio contains only a finite number of independent credits. The assumption of an infinite number of independent identically distributed credits assumes that the portfolio achieves the maximum possible risk reduction from diversification. In reality, all portfolios include only a finite number of independent credits and so all portfolios will have some remaining idiosyncratic risk.

Conditional on a specific realized value of the common market factor, $\tilde{e}_M = e_M$ the probability that a credit defaults is $\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)$. Also, conditioned on a specific realized value for the single common factor, individual credit defaults are uncorrelated. The independence of conditional defaults implies that, in a portfolio of N individual credits, the probability of realizing exactly n defaults is

$$\binom{N}{n}\left(\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^n \times \left(1-\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^{N-n}. \tag{8}$$

The probability of experiencing a default rate less than or equal to $\frac{n}{N}$ is equal to the probability of experiencing n or fewer defaults in N independent trials, or,

$$\Pr[\tilde{X}_N|\tilde{e}_M = e_M] \leq \frac{n}{N} = \sum_{i=0}^{n}\binom{N}{n} \times \left(\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^n \times \left(1-\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^{N-n}, \quad n \in [0,1,2,\ldots,N]. \tag{9}$$

The conditional portfolio loss rate distribution is constructed by multiplying the conditional default rate distribution [Expression (9)] by the uniform loss given default rate, *LGD*.

The unconditional portfolio loss rate density is a discrete function in three-dimensions: a realization of the (continuous) common factor, e_M, a realization of the (discrete) portfolio loss rate, $LGD \times X_N$ and the probability that the specific values of e_M and $LGD \times X_N$ jointly occur. This joint probability, $prob\left(e_{M,}LGD \times \frac{n}{N}\right)$, is

$$prob\left(e_{M,}LGD \times \frac{n}{N}\right) = \phi(e_M)\binom{N}{n}\left(\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^n \times\left(1-\Phi\left(\frac{\Phi^{-1}(PD)-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right)^{N-n} \quad \text{for } e_M \in (-\infty,\infty), \quad n \in \{0,1,2,3,\ldots,N\}. \tag{10}$$

The unconditional portfolio loss rate probability densities for two different examples of finite portfolios with uniform obligor exposures are shown in Figure 2. The top panel of Figure 2 represents the portfolio loss rate density for a portfolio with 30 uniform credits, each with a $PD = 0.01$, $LGD = 0.5$, and a default correlation parameter of $\rho = 0.20$. While distributions in Figure 2 are discrete, the graphs include a "mesh" that interpolates between discrete event probabilities to improve the visualization of these densities. The bottom panel of Figure 2 represents a portfolio of 100 credits with the same individual characteristics as in the top panel.

A comparison of the top and bottom panels of Figure 2 illustrates the impact of diversification in idiosyncratic default risk. For every possible realization of the common factor, e_M while the mean of the expected portfolio loss rates are identical in the two panels, the range of possible portfolio loss rates is much larger for the portfolio with 30 obligors. The variance of the portfolio loss rate is an inverse function of N, the number of credits in the portfolio.[8] The mean of both portfolios equals $PD \times LGD$, and the variance equals, $\frac{LGD^2}{N} \times PD(1-PD)$, and so the variance of the loss rate on the 100 obligor portfolio is only 30 percent of the loss rate variance on the 30 obligor portfolio.

Concentration risk, even in this simplest form — a portfolio containing less than an infinite number of portfolio credits with uniform exposure characteristics—has a dramatic effect on the portfolio default rate distribution. For a second example, consider the effect of diversification on the

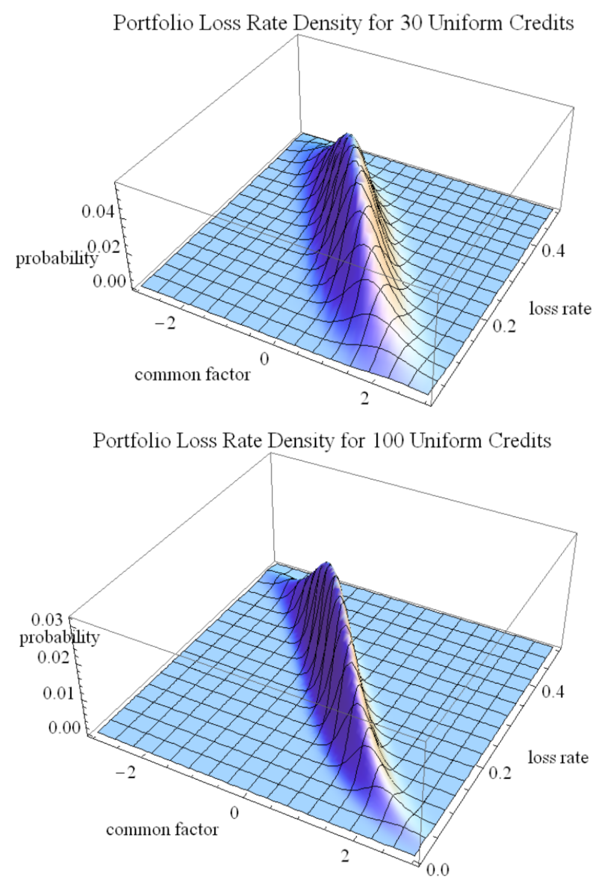

Figure 2 Unconditional loss rate density for two finite portfolios with uniform exposures.

Credits have uniform *EAD*, uniform $LGD = 0.5$, and uniform $\rho = 0.20$. The portfolio pictured in the top panel has 30 individual credits; the portfolio in the bottom panel has 100 credits.

cumulative portfolio default loss rate that encompasses 99 percent of all potential portfolio losses. Estimates of extreme critical loss distribution values like the 99 percent quantile of the portfolio loss rate distribution are frequently used to set value-at-risk style minimum capital requirements for regulated banks and other financial intermediaries.

Table 1 reports the loss rates associated with the 99 percent cumulative probability loss rate thresholds for portfolios with a different number of identical credits. Each credit is assumed to be of identical size, with identical values for ρ, *LGD*, and *PD*. To further simplify, I assume $LGD = 1$, so Table 1 is equivalent to the default rate distribution. The common market factor is set equal to its 1 percent quantile value, $e_M = \Phi^{-1}(.01) = -2.32635$, which generates high conditional default rates for the portfolio credits. The rows in Table 1 report the 99 percent cumulative default rate thresholds for finite portfolios with different numbers of obligors. The columns differ according to the assumed unconditional default rate (*PD*) for the individual portfolio credits. The last line in Table 1 reports the 99 percent cumulative default rate threshold for an asymptotic portfolio.

The elements in Table 1 show that the default rate thresholds for finite portfolios are multiple times larger than the default rate thresholds for an asymptotic portfolio of similar credits. From the values reported, it is possible to construct a concentration risk multiplier—the ratio of the exact critical default rate for a finite portfolio divided by the critical value for an otherwise similar asymptotic portfolio. When these multipliers are applied to the asymptotic portfolio critical default rate values, they reproduce the true critical values for portfolios with a finite number of credits.

Table 2 reports the value of these credit risk multipliers. These concentration risk multipliers are larger in magnitude the smaller the number of credits in the portfolio, and the smaller is a credit's unconditional probability of default. For example, for a portfolio of 50 high-quality loans with an unconditional probability of default equal to 0.1 percent, the concentration risk multiplier is nearly 5.5, implying that minimum economic capital needed to achieve 99 percent coverage for the concentrated portfolio is almost 5.5 times

Table 1 Portfolio default rate that provides at least 99 percent loss coverage when credits have uniform size and loss given default.

	PD = 1 percent		PD = 0.5 percent		PD = 0.25 percent		PD = 0.10 percent	
Number of portfolio credits	Number of defaults	Default rate	Number of defaults	Default rate	Number of defaults	Default rate	Number of defaults	Default rate
50	9	18.000	6	12.000	4	8.000	3	6.000
100	14	14.000	10	10.000	7	7.000	4	4.000
500	52	10.400	33	6.600	21	4.200	12	2.400
1,000	95	9.500	59	5.900	36	3.600	19	1.900
5,000	420	8.400	249	4.980	147	2.940	73	1.460
10,000	814	8.140	478	4.780	278	2.780	134	1.340
Asymptotic		7.525		4.301		2.412		1.096

Table 2 Selected concentration risk multipliers when credits have uniform size and loss given default

Number of portfolio credits	Unconditional probability of default			
	1 Percent	0.5 Percent	0.25 Percent	0.1 Percent
50	2.39	2.79	3.32	5.48
100	1.86	2.33	2.90	3.65
500	1.38	1.53	1.74	2.19
1,000	1.26	1.37	1.49	1.73
5,000	1.12	1.16	1.22	1.33
10,000	1.08	1.11	1.15	1.22
Asymptotic default rate	7.525	4.301	2.412	1.096

larger than the capital suggested by the asymptotic model.

5 Idiosyncratic risk and obligor concentrations in a finite portfolio

When the individual credits in a portfolio differ in size (*EAD*) or loss given default (*LGD*), then the portfolio loss rate distribution depends not only on the portfolio default rate, but on the exposure characteristics of the individual credits. Under these conditions some credits will create larger potential losses for the portfolio should they default. Exposure differences complicate the calculation of the portfolio loss rate distribution. The logic behind the construction of the loss distribution is transparent, but the calculations, while simple, are voluminous and can quickly exhaust desktop computer memory.

The calculation of the portfolio loss distribution requires the enumeration and ranking of each possible loss outcome and its attached probability. After loss outcome possibilities are enumerated, losses must be ranked from smallest to largest. The probabilities associated with each ranked loss are then accumulated to generate the cumulative portfolio loss distribution.

Table 3 Concentration risk example.

Credit ID	*EAD*	*LGD*	Loss in default	Portfolio loss rate
1	30	0.4	12	0.100
2	20	0.4	8	0.067
3	70	0.4	28	0.233
Total	120		48	0.4

Consider a simple example of this process using only three credits. Table 3 lists the obligors' characteristics. Each credit is assigned a unique identifier (credit ID). The portfolio loss rate is calculated as the potential loss associated with each credit measured as a proportion of total portfolio exposure, $\frac{LGD_i \times EAD_i}{\sum EAD_i}$.

For purposes of this example, I assume each credit has a probability of default of 5 percent and individual obligor defaults are independent.[9] The first three columns of Table 4 enumerate the entire sample space of potential outcomes—all the possible default combinations that could occur. The first column enumerates the possible default events and the second column reports the portfolio loss rates that are generated by each specific default event. The third column reports the

Table 4 Event space and cumulative loss rate distribution.

Event space			Cumulative loss rate distribution		
Credit defaults	Portfolio loss rate	Probability of event	Loss events rank ordered	Portfolio loss rate	Cumulative probability
0	0.0000	0.85738	0	0.0000	0.85738
1	0.1000	0.04513	2	0.0667	0.90250
2	0.0667	0.04513	1	0.1000	0.94763
3	0.2333	0.04513	1,2	0.1667	0.95000
1,2	0.1667	0.00238	3	0.2333	0.99513
1,3	0.3333	0.00238	2,3	0.3000	0.99750
2,3	0.3000	0.00238	1,3	0.3333	0.99988
1,2,3	0.4000	0.00013	1,2,3	0.4000	1.00000

Each credit is assumed to have a probability of default of 5 percent and individual credit defaults are assumed to be independent.

probability that the specific event occurs. To take a particular example, the probabilities associated with experiencing one default are given in rows 2, 3, and 4 of the third column. When there are three obligors, each with an independent probability of default of 5 percent, the probability of experiencing a single default is 0.135375.[10] There are three unique ways that the portfolio could experience a single default.[11] Since each of these possibilities is equally likely, the probability of any one of these outcomes is $\frac{0.135375}{3} = 0.045125$. The remaining entries in column 3 of Table 4 represent the outcomes of similar calculations for the probabilities associated with specific defaults events that involve 0, 2, and 3 credits.

The final three columns of Table 4 represent the cumulative portfolio loss rate distribution for the portfolio. The portfolio loss rate distribution is calculated from the event space by ranking the possible credit loss events (column 2) from smallest to largest beginning with the 0 default event. The resulting loss event ranking appears in columns 4 and 5. The probabilities associated with each specific event are accumulated. For example, the probability of 0 losses is .85738. The next smallest possible loss rate is 0.0667, and the probability of experience a loss rate that is at most 0.0667 is 0.90250, and so on.

Obligor concentration changes some important features of the cumulative portfolio loss rate distribution. When portfolio credits have identical *EAD*s and *LGD*s, the number of unique outcomes in the sample space is reduced. For example, with three credits of identical size and *LGD*, there are only four possible loss rate outcomes: those associated with 0, 1, 2, or 3 defaults. Whereas, when the credits differ in size, *LGD,* or in both dimensions, there are eight possible loss rate outcomes. Uniformity reduces the size of the sample space because when one credit defaults, the loss is the same no matter which of the individual credits defaults—and there is no need to keep track of individual obligor performance in order to calculate the associated portfolio loss rate.

Figure 3 compares the portfolio loss distribution example in Table 4 with a portfolio of equivalent size and total exposure, but with uniform credits. The blue points in Figure 3 represent the portfolio loss distribution for the portfolio with obligor concentrations—the three-credit example

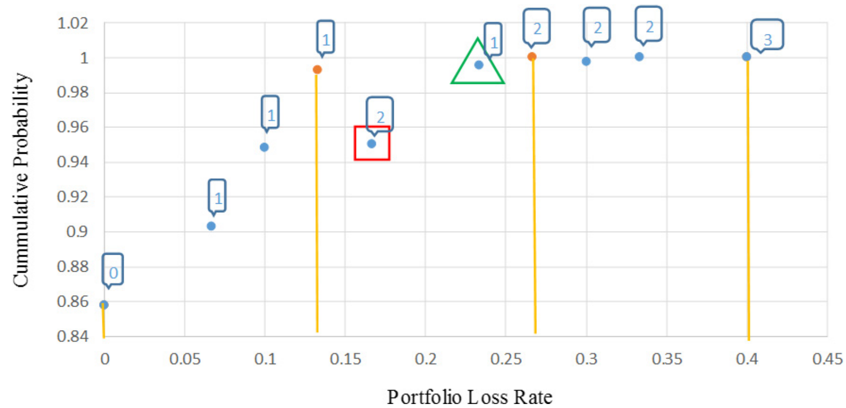

Figure 3 Portfolio loss rate distribution with concentration risk.

from Table 4. This portfolio has total *EAD* of 120 and a possible worst-case loss rate of 40 percent. The orange columns in Figure 3 represent the portfolio loss distribution associated with a portfolio comprising three independent uniform credits, each with $PD = 5$ percent, $EAD = 40$, and $LGD = 40$ percent.[12] This uniform portfolio has the same size and total loss potential as the portfolio with obligor concentration. The numbers in the call-out boxes in Figure 3 represent the number of defaulted credits associated with each point in the respective distributions.

Figure 3 shows the reduced unique number of outcomes in the sample space associated with the uniform credit portfolio relative to the portfolio with obligor concentrations. While the two portfolio loss distributions have two points in common {(0, 0.86), (0.4, 1)}, the uniform credit distribution has only two possible intermediate loss rates, while the portfolio with obligor concentration has six possible intermediate loss rates.

A general feature associated with obligor concentration risk is that individual credit exposure differences can cause events with fewer defaults to have portfolio loss severities that exceed events with a larger number of defaults. An example of this phenomenon appears in Figure 4 where the joint default of credits 1 and 2 (red square) produces a smaller portfolio loss rate than the default of the single credit with the largest exposure (green triangle). This feature implies that the cumulative probability associated with the largest single default exposure in the concentrated portfolio will always be at least as large as the cumulative probability of a single default in a comparable uniform obligor portfolio. In other words, in terms of Figure 4, the cumulative probability (height) of the point inside the green triangle will always be as large, or larger, than the cumulative probability of the orange column associated with one default. While single default events are equally likely regardless of their severity, when there are obligor concentrations, two-default events can rank below the largest single default event.

When portfolios include obligor concentration risk, the calculation of the loss distribution follows the logic outlined in Table 4. However, even for portfolios with a modest number of obligors, the number of calculations required to construct the exact loss distribution can quickly become unmanageable. For example, in a portfolio with 25 obligors with different exposure characteristics, there are 3,268,760 unique ways the portfolio can experience 10 obligor defaults. When there are 100 obligors, the number of unique combinations of 10 obligor defaults exceeds, 1.73×10^{13}, and the total number of unique default combinations in the entire event space is approximately 1.27×10^{30}. As these examples illustrate, the full enumeration of the possible set of loss outcomes becomes impractical except when the portfolio has only a modest number of obligors. This motivates the need for the approximation outlined in the following section.

6 Obligor concentration risk and an approximate portfolio loss rate distribution

When portfolios with obligor concentrations include more than a modest number of credits, it becomes infeasible to enumerate the entire sample space and construct the exact portfolio loss rate distribution. However, it is possible to approximate the portfolio loss rate distribution. The approximation is computationally simple, yet it produces reasonably accurate quantile estimates for the true cumulative portfolio loss rate distribution, especially for quantiles in the upper tails of the distribution—the quantiles that are typically used to set capital allocations or regulatory capital requirements. I will explain the intuition behind the portfolio loss rate approximation by referencing the three-obligor example in the previous section.

When credits are uniform, after conditioning on the common market factor realization, the uncertainty in portfolio loss rate distribution is entirely determined by the binomial distribution that specifies the probability associated with experiencing each possible integer number of defaults. When credits are non-uniform in size or *LGD* (or both), the portfolio loss associated with "n" defaults depends on which specific "n" credits default. Moreover, it is also possible that the loss generated by "$n+1$" defaults can be smaller than the loss generated by "n" defaults (or even "$n-1$", "$n-2$", or "$n-j$" defaults) depending on which particular credits default.

The algorithm to approximate the portfolio loss rate distribution in the presence of obligor concentrations uses a specific isomorphic portfolio in the approximation. The isomorphic portfolio has the same number of credits, identical *PD* and correlation parameters, the same total portfolio exposure, and same maximum loss rate, but the individual credits have uniform exposure characteristics.

Suppose there are N independent obligors in the credit portfolio. Let TE represent the total portfolio exposure at default; TL the maximum possible portfolio loss; LR the maximum portfolio loss rate, and plr_i the portfolio loss rate associated with the default of credit i,

$$TE = \sum_{i=1}^{N} EAD_i;$$

$$TL = \sum_{i=1}^{N} EAD_i \times LGD_i;$$

$$LR = \frac{TL}{TE}; \quad \text{and}$$

$$plr_i = \frac{EAD_i \times LGD_i}{TE}.$$

Let Plr represent the rank-ordered vector of individual credit portfolio loss rates,

$$Plr = \{plr_1, plr_2, plr_3, \ldots, plr_N\}, \quad \text{where,}$$

$$plr_1 \leq plr_2 \leq plr_3 \leq \cdots \leq plr_N.$$

Without loss of generality, I will assume the portfolio loss rate attached to each credit to be unique.[13]

Now consider the isomorphic credit portfolio. For this portfolio, $LGD_i = \frac{LR}{N}$, and $EAD_i = \frac{TE}{N}$. This portfolio has the same underlying binomial probability structure and the same maximum portfolio default rate as the credit portfolio with obligor concentration risk.[14]

Both portfolio loss rate distributions are discrete. The quantile values of the isomorphic uniform exposure distribution are defined as follows. For $q \in [0, 1]$, and a set of integers, $K = \{0, 1, 2, \ldots, N\}$, the quantile q of the portfolio loss rate distribution is the smallest portfolio loss rate that has a cumulative probability at least as large as q, or the loss rate $k_q \times \frac{LR}{N}$, such that,

$$k_q = \inf\left\{K : \sum_{i=0}^{k} \binom{N}{i} PD^i (1 - PD)^{N-i} \geq q\right\}. \quad (11)$$

In Expression (11), k_q represents the number of defaults that are needed to generate a cumulative probability at least as large as q. Depending on the value of q selected, k_q can be any integer value between 0 and N. Finally, let $F^i\left(k_q \times \frac{LR}{N}\right) = \sum_{i=0}^{k_q} \binom{N}{i} PD^i (1 - PD)^{N-i}$ represent the cumulative probability distribution function associated with k_q defaults under the isomorphic portfolio loss rate distribution.

The quantiles of the portfolio loss rate distribution with obligor concentrations can be approximated as follows. Let $F(lr), lr \in [0, LR]$ represent the cumulative probability distribution for the loss rate on the portfolio with obligor concentrations. Select the desired quantile q, and use Expression (11) to solve for k_q. Construct, $\widehat{LR}_q$,

$$\widehat{LR}_q = \sum_{i=1}^{k_q} plr_{N-i+1} \quad \text{for } plr_i \in Plr. \quad (12)$$

$\widehat{LR}_q$ is a "conservative" estimate[15] for the portfolio loss rate that generates a cumulative probability of at least q under the true cumulative probability distribution $F(lr)$,

$$F(\widehat{LR}_q) \geq F^i\left(k_q \times \frac{LR}{N}\right). \quad (13)$$

A formal proof of this inequality is given in the Appendix.

In plain English, the algorithm says the following: (1) use the isomorphic distribution and solve for the number of defaults that are required to reach the desired quantile of the cumulative portfolio loss rate distribution (k_q defaults); (2) calculate the sum of the k_q largest individual credit portfolio loss rates in the portfolio with obligor concentrations; (3) the sum of the k_q largest individual credit portfolio loss rates will have a true cumulative probability that is at least as large a q.

For small values of q (relatively small portfolio loss rates) the approximation for the cumulative probability associated with $\widehat{LR}_q$ will likely understate the true value, $F(\widehat{LR}_q)$. However, the approximation becomes very good as q gets large, and it becomes exact as q approaches 1,

$$\lim_{q \to 1}\left[F(\widehat{LR}_q) - F^i\left(k_q \times \frac{LR}{N}\right)\right] = 0. \quad (14)$$

Expression (14) implies that for high quantile values [for example, $q = 0.95$, or $q = 0.99$], there is very little error involved is using $F^i\left(k_q \times \frac{LR}{N}\right)$ as an approximation for $F\left(\widehat{LR}_q\right)$. Formal justification for this claim is provided in the Appendix.

Consider a step-by-step example of the approximation algorithm for a 10-credit portfolio with obligor concentration risk. Table 5 provides the details on the individual obligor exposure characteristics. Each credit is assumed to have $PD = 0.01$, $\rho = 0.20$, $LGD = 0.40$, and $\rho = 0.2$. Total portfolio exposure at default is 1680, and the worst-case default losses are 672, which implies a maximum portfolio loss rate of 40 percent.

Table 6 illustrates the approximation. The isomorphic portfolio is identical to the concentrated portfolio in all characteristics except its credits have a uniform $EAD = 168$. To remove default correlation, I condition on the 1 percentile of the common market factor distribution (i.e. $\tilde{e}_M = -2.32635$). After conditioning, defaults will be independent

Table 5 Ten-credit portfolio with obligor concentrations.

Credit ID	*EAD*	Loss in default	Portfolio loss rate
1	50	20	0.0119
2	100	40	0.0238
3	110	44	0.0262
4	125	50	0.0298
5	150	60	0.0357
6	170	68	0.0405
7	200	80	0.0476
8	225	90	0.0536
9	250	100	0.0595
10	300	120	0.0714
Total	1680	672	0.4

and each credit will have a conditional probability of default equal to 0.0752508.

Using the relevant parameters, I construct the portfolio loss rates and the cumulative probabilities associated with $n = \{1, 2, 3, \ldots, 10\}$ defaults under the isomorphic portfolio loss rate distribution. I calculate the loss rates associated with the sum of the k largest individual credit portfolio loss rates for $k = \{1, 2, 3, \ldots, 10\}$. I then construct the full event space for the concentrated portfolio, rank-order the possible outcomes, and calculate its true portfolio loss rate distribution. The true probabilities associated with the largest "k" exposures are reported in the Table 6 column labeled "Concentrated Distribution Cumulative Probability".

Table 6 Approximating the portfolio loss rate distribution when there is obligor concentration risk.

Numbers defaults	Uniform portfolio loss rate	Uniform loss distribution cumulative probability	Portfolio loss rate from largest *k* exposures	Concentrated distribution cumulative probability	Δ Loss rate	Percentage change in portfolio loss rate	Δ Cumulative probability
0	0.00	0.4573	0.0000	0.4573	0.0000	0.00	0.0000
1	0.04	0.8295	0.0714	0.8850	0.0314	78.57	0.0555
2	0.08	0.9658	0.1310	0.9864	0.0510	63.69	0.0207
3	0.12	0.9953	0.1845	0.9989	0.0645	53.77	0.0035
4	0.16	0.9996	0.2321	0.9999	0.0721	45.09	0.0004
5	0.20	1.0000	0.2726	1.0000	0.0726	36.31	0.0000
6	0.24	1.0000	0.3083	1.0000	0.0683	28.47	0.0000
7	0.28	1.0000	0.3381	1.0000	0.0581	20.75	0.0000
8	0.32	1.0000	0.3643	1.0000	0.0443	13.84	0.0000
9	0.36	1.0000	0.3881	1.0000	0.0281	7.80	0.0000
10	0.40	1.0000	0.4000	1.0000	0.0000	0.00	0.0000

Portfolio loss rate distribution approximation calculations. The true portfolio exposures are listed in Table 5. For each obligor, $PD = 0.01$, $\rho = 0.2$, and $LGD = 0.4$. All probabilities are conditional probabilities calculated with the common factor equal to its 1 percentile value (-2.33). The conditional probability of default for each credit is 0.0753. The isomorphic portfolio has 10 credits, each with $EAD = 168$, $LGD = 0.4$, $PD = 0.01$, $\rho = 0.2$ and conditional probabilities of default = 0.075. The column Δ Loss rate reports the portfolio loss rate associated with the largest n exposures from the concentrated portfolio less the portfolio loss rate associated with n defaults for the isomorphic portfolio. Δ Cumulative probability reports the true probability associated with the largest n exposures and the approximated probability for n defaults using the isomorphic loss rate distribution. Δ Loss rate and Δ Cumulative probability measure the accuracy of the approximation.

The final 3 columns in Table 6 provide information on the accuracy of the portfolio loss rate approximation. The column labeled Δ Loss rate compares the loss rates associated k defaults under the uniform isomorphic distribution to the true loss rate for the k largest portfolio default exposures. This column represents the loss rate underestimation that occurs when the concentrated portfolio is modeled as a portfolio with uniform exposures. The adjacent column to the right expresses this loss rate estimation error as a percentage using the uniform isomorphic loss rate as the base. The final column in Table 6 represents the difference between the true cumulative probabilities associated with the portfolio loss rates generated by the largest k exposures and the cumulative probability assigned by the approximation algorithm. Notice that as the number of defaults increases and the quantile of the cumulative probability increases, the cumulative probability approximation error (Δ Cumulative probability) monotonically declines to the point that there is no measureable error in the cumulative probability beyond $n = 5$ defaults. This particular approximation is illustrated in Figure 3.

The blue points in Figure 4 represent the actual loss rate distribution for the concentrated portfolio. The red columns represent the loss rate distribution of the isomorphic portfolio with uniform exposures. The gray columns represent the actual points on the concentrated credit portfolio's loss distribution that correspond to the sum of the largest k default exposures, for $k = 1, 2, 3, \ldots, 10\}$. The call-out box with arrows represent the probability approximations that apply to each of these reference loss rates of the true concentrated portfolio loss distribution.

Figure 5 plots the approximate probability density for a portfolio of 100 credits with obligor concentration risk. The portfolio credits in this example have $PD = 0.01, LGD = 0.4$, and $\rho = 0.2$. The exposure sizes associated with each credit are given by the sequence, $EAD_i = \{105, 110, 115, \ldots, 600\}$. The isomorphic portfolio with uniform EAD has $EAD = 352.50$.

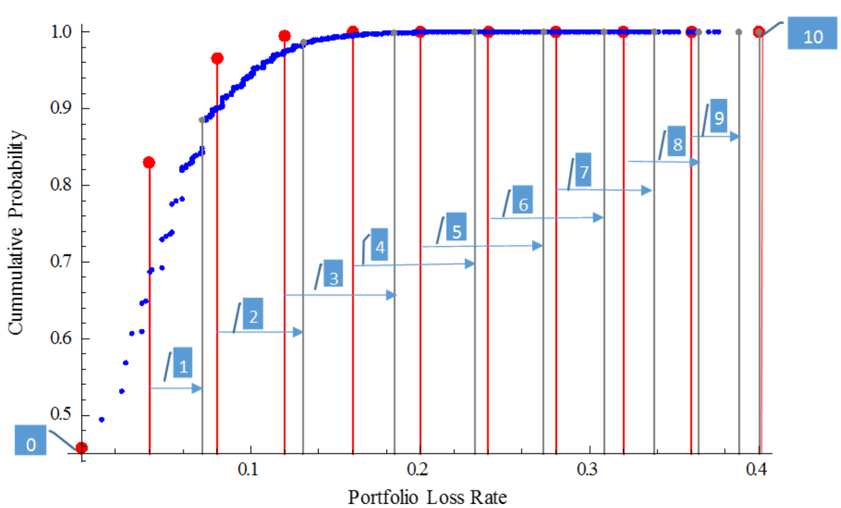

Figure 4 Approximating the portfolio loss rate distribution in the presence of obligor concentrations.

Graphic illustration of the approximation in Table 6. The blue points represent the true cumulative portfolio loss rate distribution for the concentrated portfolio. The red columns represent the cumulative loss rate distribution for the isomorphic portfolio. The gray colums represent the points on the true cumulative probability distribution associates with the defaults of the largest "k" exposures. $k = \{1, 2, \ldots, 10\}$. The call out boxes with the arrows link the approximating point on the isomorphic loss rate distribution with the true point on the concentrtated portfolio loss rate distribution.

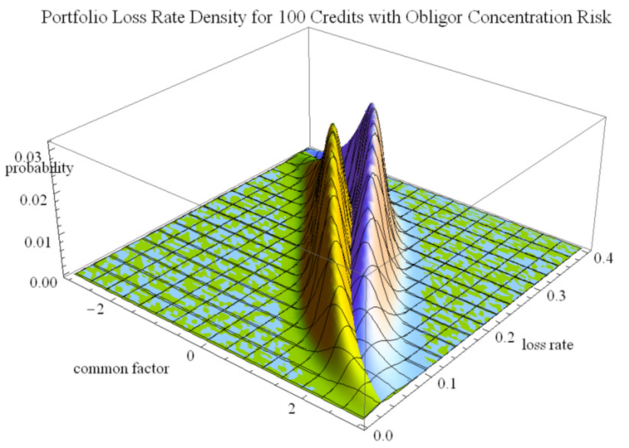

Figure 5 Approximate loss rate density for a portfolio with obligor concentrations.

The blue paraboloid is the approximate probability density for the loss rate on a portfolio with 100 credits, each credit having, $PD = 0.01$, $LGD = 0.4$, and $\rho = 0.2$. The portfolio *EAD*s differ by 5 and range from 105 to 600. The yellow paraboloid is the probability density of the isomorphic portfolio with 100 credits which is identical to the concentrated portfolio in all respects except that its credits have a uniform $EAD = 352.50$.

The yellow paraboloid in Figure 5 is the density for the isomorphic portfolio loss rate distribution. The blue paraboloid is the portfolio loss rate density for the portfolio with obligor concentrations. Figure 5 provides a clear illustration that obligor concentration risk increases the portfolio loss rates for any common factor realization. The imbalances in the portfolio's obligor concentrations reduce the diversification of idiosyncratic risk.

Unlike the granularity assumption, this approximation for the portfolio loss rate distribution does not depend on the default correlation modeling assumptions that are used to generate correlation among defaults. Whatever mechanism used to model common factors that drive defaults [e.g. Vasicek model, CreditRisk$^+$, etc], once the common factors are controlled, individual defaults can be modeled as independent Bernoulli events and the algorithm can be applied to approximate the portfolio loss rate distribution in the presence of concentration risk.

7 Value-at-risk capital for a marginal credit

Economic capital allocation decisions and regulatory capital requirements often use value-at-risk to set required investment capitalization rates.[16] In general, capitalization requirements are often set so that the equity used to fund the portfolio will not be exhausted by portfolio losses except in exceptionally rare circumstances. Such a rule is often operationalized by setting the share of equity used to fund the portfolio equal to a high quantile of the portfolio's loss rate distribution. Typical coverage rates used in capital allocation models range between 95 and 99 percent, although the Basel Committee on Banking Supervision sets the coverage rate at 99.9 percent.

Within a value-at-risk capital framework, it is of interest to know the additional capital that will be needed should a new credit be added to an existing portfolio. Under the assumptions of the Vasicek asymptotic single common factor model, the capitalization rate that applies to any new credit is independent of the composition on the portfolio and equal to the capitalization rate for all the credits already in the portfolio. This invariance arises because idiosyncratic risk is fully diversified and there is no additional diversification benefit from adding an additional credit. However, in most cases, portfolios are not asymptotic and the capitalization rate of the marginal credit will depend on the composition of the existing portfolio.

In the simplest setting, where portfolio credits are uniform in size and default risk characteristics, the capitalization rate on a new marginal credit with exposure and risk characteristics identical to the credits already in the portfolio tends to decline as the number of credits in the existing portfolio increases. The tendency for declining capitalization rates is upset periodically as N increases as a consequence of the discrete nature of the default rate distribution. Discrete jumps in the critical value of the default rate used to set capital create a capitalization rate that declines with N, but with an irregular saw tooth style pattern.

Using the notation defined earlier, the q-quantile of a portfolio loss rate distribution is associated with k_q defaults. If equity capital is set to cover q percent of all possible portfolio losses, the capitalization rate on the portfolio (and each of its credits) will be $LDG \times \frac{k_q}{N}$.

Now consider adding an additional credit to the portfolio. For finite distributions, the loss rate distribution is discrete, and k_q may not change for $N+1$ credits. In such case, the capitalization rate on the marginal credit, that is the change in the total required capital for the portfolio of $N+1$ credits divided by the EAD of the new credit is $LGD \times \left(\frac{N-1}{N}\right)\left(\frac{k_q}{N+1}\right)$. The marginal capital is smaller than $\left(\frac{k_q}{N+1}\right)$ because idiosyncratic risk is better diversified in the new portfolio generating capital savings on the original N credits. This is accounted for by the factor $\frac{N-1}{N} < 1$. When N is small, the extra diversification benefit can be large, but as N increases, the benefit of additional idiosyncratic diversification diminishes.

As N increases, and more credits with identical characteristics are added to the portfolio, so will the value of k_q. Binomial probabilities are associated with the number of discrete default events, and once a sufficient number of additional credits are added to the portfolio, k_q will increase. The increase in k_q as N increases creates a declining sawtooth pattern in required capitalization rates. A specific example of this sawtooth capitalization rate pattern is illustrated in Figure 6.

When there are obligor concentrations in the portfolio, an additional factor enters into the capitalization rate calculations. With concentration risk, the value-at-risk capitalization rate is equal to the exposure generated by the specific k_q credits that individually generate the largest portfolio loss rates, $\widehat{LR}_q = \sum_{i=1}^{k_q} plr_{N-i+1}$

When a credit is added to the portfolio, new portfolio loss rates must be computed for the individual credits that are contained in $\widehat{LR}_q$,

$$plr'_j = \frac{EAD_j \times LGD_j}{\sum_{i=1}^{N+1} EAD_i}, \quad \text{for } plr_j \in \widehat{LR}_q. \tag{15}$$

The portfolio loss rate of the newly added credit must be calculated and compared with the portfolio loss rates plr'_j of the credits in $\widehat{LR}_q$. If the new credit's portfolio loss rate is less than any plr'_j that is included in $\widehat{LR}_q$, the new portfolios' capitalization rate is $\frac{\sum_{i=1}^{N} EAD_i}{\sum_{i=1}^{N+1} EAD_i} \times \widehat{LR}_q$. In this instance, the benefit of additional diversification is measured by $\frac{\sum_{i=1}^{N} EAD_i}{\sum_{i=1}^{N+1} EAD_i} < 1$.

Should the new credit generate a portfolio loss rate that exceeds the smallest portfolio loss rate plr'_j in $\widehat{LR}_q$, the new credit will replace the smallest loss rate, and the new larger loss rate associated with the quantile must be calculated, $\widehat{LR}'_q > \widehat{LR}_q$. Consequently, in cases when k_q is unaffected by the addition of a new credit, the change in the portfolio's capitalization rate required by the addition of a new credit is

$$\widehat{LR}'_q - \left(\frac{\sum_{i=1}^{N} EAD_i}{\sum_{i=1}^{N+1} EAD_i}\right)\widehat{LR}_q. \tag{16}$$

Of course, in some instances the addition of a new credit will require a unit increase in the value of k_q. In these cases $\widehat{LR}'_q$ will also increase because an additional large portfolio loss rate will be added to the sum that determines $\widehat{LR}'_q$. Consequently, the capitalization rates will exhibit a declining sawtooth pattern, but unlike Figure 6, the jump increments will be irregular, with a size that

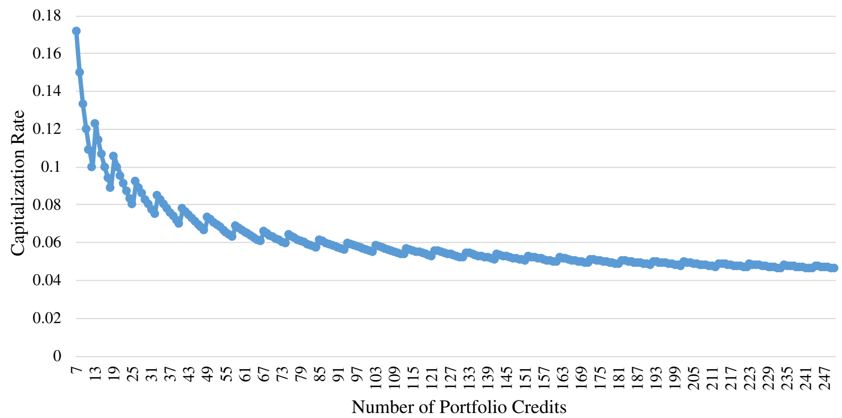

Figure 6 Idiosyncratic risk diversification for uniform credit exposures.

The capitalization rate needed to achieve 99 percent coverage of the loss rate distribution in a portfolio with N credits with equal EAD, $PD = 0.01$, $\rho = 0.2$, and $LGD = 0.4$.

depends on the loss exposure concentrations in the portfolio.

8 Conclusion

The benefit of portfolio diversification is one of the only true "free lunches" available to investors in equilibrium. Risk can be reduced merely by the judicious structuring of portfolio investments. Given the practical benefits to be gained, the academic literature offers surprisingly few practical approaches for assessing the impact of obligor concentration risk on the diversification of credit risk portfolios. In this paper, I analyze obligor concentration risk and present a new algorithm that can be used to approximate the loss rate distribution for a fixed income portfolio with credit risk concentrations. The intuition behind the approximation is easily understood using simple set theory without the need for advanced mathematics or statistics. The approximation is independent of the modeling structure assumed to generate default correlation and is highly accurate in the upper quantiles of a portfolio's loss rate distribution. Its accuracy makes it especially useful for estimating economic capital allocations or setting regulatory capital requirements for credit risk portfolios with obligor concentration risk.

Appendix

Consider a portfolio of N credits. Assume each credit has a probability of default of PD, and default events are independent. The individual credit $EADs$ and $LGDs$ can be arbitrary admissible values

$[EAD_i > 0, 0 \leq LGD_i \leq 1, \forall i]$. Consistent with the text, define: $TE = \sum_{i=1}^{N} EAD_i$; $TL = \sum_{i=1}^{N} EAD_i \times LGD_i$; $LR = \frac{TL}{TE}$; and $plr_i = \frac{EAD_i \times LGD_i}{TE}$. Let $F(LR)$, $LR \in [0, 1]$ be the cumulative probability distribution for the loss rate on this portfolio.

Consider a hypothetical isomorphic portfolio of N credits with uniform obligor exposure characteristics. Each credit has a probability of default of PD. Default events are independent. Each credit has $EAD_i = \frac{TE}{N}$ and $LGD_i = \frac{LR}{N}$. Let $F^i(LR)$, $LR \in [0, 1]$ represent the cumulative probability distribution for the loss rate on this isomorphic portfolio.

For either portfolio, the probability of experiencing exactly k defaults is $\binom{N}{k} PD^k (1 - PD)^{N-k}$. There are $\binom{N}{k}$ unique combinations in which individual credits in either portfolio can experience k defaults. Each unique combination of k defaults has a probability of $p_k = PD^k(1-PD)^{N-k}$. For the isomorphic portfolio, the portfolio loss rate is identical for each of the $\binom{N}{k}$ combinations of k defaults.

In the isomorphic uniform obligor portfolio, the portfolio loss rate increases monotonically with the number of portfolio defaults. In the portfolio with arbitrary credit exposure characteristics, the portfolio loss rate need not be a monotonic function of the number of defaults.

Let R_k be the set of $M(k) = \binom{N}{k}$ portfolio loss rates generated by loss events with exactly k defaults, rank-ordered (from smallest to largest) according to the event's total portfolio loss rate,

$$R_k = \{elr_1^k, elr_2^k, elr_3^k, \ldots, elr_{M(k)}^k\}.$$

elr_1^k is the minimum portfolio loss rate in R_k; it is the smallest portfolio loss rate generated by k defaults. Similarly, $elr_{M(k)}^k$ is the maximum portfolio loss rate in R_k; the largest portfolio loss rate that can be generated by k defaults. Each individual event in R_k has a probability of p_k.

R_1 has N elements corresponding to the number of unique ways to generate one default from the N individual credits in the portfolio. Let R_1^+ be the set of events in $\{R_2, R_3, \ldots, R_k, \ldots R_N\}$ where $elr_i^k \leq elr_N^1$ for all i, and $k = \{2, 3, 4, \ldots, N\}$. Let $R_k^{1C} = R_k \backslash R_1^+$ where the notation $R_k \backslash R_1^+$ indicates elements in set R_k that are not in R_1^+. Let $prob(R_k \cap R_1^+)$ represent the cumulative probability of events that are in the intersection of sets R_k and R_1^+.

The probability of observing a loss at least as large as elr_N^1 is given by $prob(R_1 \cup R_1^+) = p_0 + p_1 + \sum_{i=2}^N prob(R_i \cap R_1^+)$. Since $F^i\left(\frac{LR}{N}\right) = p_0 + p_1$, and $\sum_{i=2}^N prob(R_i \cap R_1^+) \geq 0$, it follows that $F(elr_N^1) \geq F^i\left(\frac{LR}{N}\right)$. That is, the probability of experiencing a portfolio loss rate that is less than or equal to the portfolio loss rate caused by the default of the largest single loss exposure in the concentrated portfolio is always greater than or equal to the probability of experiencing one or fewer defaults in the isomorphic uniform credit portfolio.

In order to demonstrate, $F(elr_{M(2)}^2) \geq F^i(2 \times \frac{LR}{N})$, let R_2^+ be the set of events in $\{R_3, R_4, \ldots, R_k, \ldots R_N\}$ where $elr_i^k \leq elr_{M(2)}^2$ for all i, and $k = \{3, 4, 5, \ldots, N\}$. Let $R_k^{2C} = R_k \backslash R_2^+$ for $k = \{3, 4, 5, \ldots, N\}$.

Using the fact that, $prob(R_1 + (R_1^+ \cap R_2) \cup R_2^{1C}) = prob(R_1 + R_2)$, the probability of observing a portfolio loss rate at least as large as $elr_{M(2)}$ is given by $prob(R_1 \cup R_2 \cup R_2^+) = p_0 + p_1 + p_2 + \sum_{i=3}^N prob(R_i \cap R_2^+)$. Since $F^i\left(2 \times \frac{LR}{N}\right) = p_0 + p_1 + p_2$, and $\sum_{i=3}^N prob(R_i \cap R_3^+) \geq 0$, it follows that $F(elr_{M(2)}^2) \geq F^i\left(2 \times \frac{LR}{N}\right)$. In plain language, the last inequality says that the probability of a portfolio loss rate that is less than or equal to the portfolio loss rate caused by the default of the largest two loss exposures in the concentrated portfolio will always be greater than or equal to the probability of experiencing two or fewer defaults in the isomorphic uniform credit portfolio.

The remaining inequalities, $F(elr_{M(3)}^3) \geq F^i\left(3 \times \frac{LR}{N}\right)$, $F(elr_{M(4)}^4) \geq F^i\left(4 \times \frac{LR}{N}\right), \ldots, F(elr_{M(N)}^N) \geq F^i(LR)$, are established by induction.

While this proves that the approximation is always conservative, it does not provide any evidence in the accuracy of the approximation. The two cumulative probability distributions are actually equal by construction for $NF(elr_{M(N)}^N) = F^i(LR) = 1$. The distributions also must agree for $N-1$ defaults because the $prob(R_{N-1}^+) = 0$; that is, the single event that corresponds to N defaults must be larger than the largest loss rate generated by $N-1$ defaults $elr_{M(N-1)}^{N-1}$.

Moving back from $q = 1$, in the direction of $q = 0$, for $N-2$ defaults, the term $prob(R_{N-2}^+)$

can be larger than 0. For example, consider a portfolio of 10 credits each with $EAD = 100$. If the loss given default on these credits are the integer values from 1 to 10, and $\{plr_i\}$ represents the vector of individual credit portfolio loss rates ranked in ascending order, the sum of the largest eight credit portfolio loss rates $\sum_{i=3}^{10} plr_i = .52 > \sum_{i=1}^{9} plr_i = .45$, and so $prob(R_{N-2}^{+}) > 0$, and $F(elr_{M(N-2)}^{N-2}) > F^i\left(\frac{N-2}{N} LR\right)$.

In the far right tail of the distribution, the probabilities associated R_{N-2}^{+} are generally very small for the probabilities of default in the ranges normally encountered in fixed income portfolios. Defaults are distributed binomially and so most of the probability mass for the default distribution is located within two standard deviations $[\pm 2 \times \frac{LGD}{\sqrt{N}} \times \sqrt{PD(1-PD)}]$ of the expected value of the distribution [a portfolio loss rate of $LGD \times PD$]. Progressing back towards $q = 0$, values for R_{N-j}^{+} that are associated with loss rates far above the mean portfolio loss rate will also be very small and so the approximation will be highly accurate in the upper tail region of the distribution.

Notes

1 Legend has it that Albert Einstein once called compound interest "the most powerful force in the universe" or "the greatest invention in human history." However, there is no official record or transcript that supports this claim.

2 A granularity adjustment for obligor concentration risk was first introduced by Vasicek and refined by Wilde (2001), Martin and Wilde (2002), Gordy (2003), Gordy and Lütkebohmert (2013), Gordy and Marrone (2012) among others.

3 The so-called "granularity adjustment" for concentration risk was first proposed by Vasicek in 1991. The Basel Committee on Banking Supervision (2001) proposed a granularity adjustment based on estimates from Monte Carlo simulations using commercial risk measurement software [CreditRisk$^+$]. Subsequently, Wilde (2001), Martin and Wilde (2002), Gordy (2003), Gordy and Lütkebohmert (2013), Gordy and Marrone (2012) provided generalized statistical theory to support the granularity adjustment in a number of model settings.

4 For example, the Vasicek and CreditRisk$^+$ portfolio models have different probability structures that drive defaults, so they have different granularity adjustment factors. The volume edited by Gundlach and Lehrbass (2004) includes a discussion of the CreditRisk$^+$ model and various generalizations.

5 Wilde (2001).

6 $E(\tilde{V}_i) = 0$, and $\sigma^2(\tilde{V}_i) = E(\tilde{V}_i^2) - E(\tilde{V}_i)^2 = 1$.

7 Alternatively, EAD can be measured as the initial loan balance. Here LGD would exclude the loss of accrued interest and Expression (7) is the loss rate of the initial portfolio that owes to defaults.

8 The portfolio loss rate is given by $\left(\frac{LGD}{N}\right)\tilde{n}$, where n is the number of defaults in a portfolio of N credits. $\tilde{n}$ is distributed binomially with a mean, $E(\tilde{n}) = N \times PD$ and variance of $Var(\tilde{n}) = N \times PD(1-PD)$.

9 The default events will be independent after conditioning on a specific value of the common factor. To keep the discussion as simple as possible, I assume the conditioning step has already been done.

10 The probability of experiencing one success in three independent Bernoulli trials, where the probability of a successes on each Bernoulli trial is 5 percent, $\binom{N}{k} PD^k (1-PD)^{N-k} = \binom{5}{1} .05(.95)^2$ where $\binom{N}{k} = \frac{N!}{k!(N-k)!}$

11 The number of unique combinations of k obligor defaults in a portfolio of N obligors is $\binom{N}{k}$.

12 The credits' individual portfolio loss rates equal $16/120 = 13.33$ percent.

13 There is no conceptually difficultly incorporating credits with identical portfolio default rates. However, it would needlessly complicate the discussion.

14 The orange bars in Figure 2 represent the isomorphic loss rate distribution that corresponds to the obligor concentration portfolio loss rate distribution (blue points).

15 By conservative, I mean, by way of example: if $q = 0.95$ and $\widehat{LR}_q$ sets the required capitalization rate for the portfolio, the true probability of default associated with a capitalization rate of $\widehat{LR}_q$ will always be 5 percent or less.

16 The equity used to fund the portfolio is set equal to a value-at-risk estimate for the portfolio return or loss rate distribution. For addition discussion, see for example, Kupiec (2004, 2007).

References

Altman, E. and A. Saunders (1998). "Credit Risk Measurement: Developments over the last 20 years," *Journal of Banking and Finance* **21**, 1721–1742.

Basel Committee on Banking Supervision (2001). "Basel II: The New Basel Capital Accord, Second Consultative Paper," The Bank for International Settlements.

Gordy, M. (2003). "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules," *Journal of Financial Intermediation* **12**(3), 199–232.

Gordy, M. and Marrone, J. (2012). "Granularity Adjustment for Mark-to-Market Credit Risk Models," *Journal of Banking and Finance* **36**, 1896–1910.

Gordy, M. and Lütkebohmert, E. (2013). "Granularity Adjustment for Regulatory Capital Assessment," *International Journal of Central Banking* **9**(3), 33–71.

Gundlach, V. M. and Lehrbass, F. B. (eds.) (2004). *CreditRisk$^+$ in the Banking Industry*. Heidelberg: Springer-Verlag.

Hull, J. and White, A. (2004). "Valuation of a CDO and an *n*-th to Default CDS Without Monte Carlo Simulation," *Journal of Derivatives* **12**(2), 8–23.

Kupiec, P. (2004). "Estimating Economic Capital Allocation for Market and Credit Risks," *The Journal of Risk* **6**(4), 11–29.

Kupiec, P. (2007). "Capital Allocation for Portfolio Credit Risk," *The Journal of Financial Services Research* **32**(1–2), 103–122.

Kupiec, P. (2008). "A Generalized Single Common Factor Model of Portfolio Credit Risk," *The Journal of Derivatives* **15**(3), 25–40.

Markowitz, H. (1952). "Portfolio Selection," *The Journal of Finance* **7**(1), 77–91.

Martin, R. and Wilde, T. (2002). "Unsystematic Credit Risk," *Risk* **15**(11), 123–128.

Sharpe, W. F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," *The Journal of Finance* **19**(3), 425–442.

Vasicek, O. (1987). *Probability of Loss on Loan Portfolio*. KMV Corporation Working Paper.

Vasicek, O. A. (1991). "Limiting Loan Loss Probability Distribution," KMV Corporation Working Paper. Subsequently published as, Vasicek, O. (2002). "Loan Portfolio Value," *Risk* **15**(12), 160–162.

Wilde, T. (2001). "Probing Granularity," *Risk* **14**(8), 103–106.

Keywords: Portfolio diversification; idiosyncratic default risk; obligor concentration; Vasicek single common factor model of credit risk; credit value-at-risk; Basel bank capital requirements