

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Kupiec, Paul H.

Working Paper Capital for concentrated credit portfolios

AEI Economics Working Paper, No. 2015-06

Provided in Cooperation with: American Enterprise Institute (AEI), Washington, DC

Suggested Citation: Kupiec, Paul H. (2015) : Capital for concentrated credit portfolios, AEI Economics Working Paper, No. 2015-06, American Enterprise Institute (AEI), Washington, DC

This Version is available at: https://hdl.handle.net/10419/280526

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Capital for Concentrated Credit Portfolios

Paul H. Kupiec American Enterprise Institute

AEI Economic Policy Working Paper 2015-06

May 1, 2015

Capital for Concentrated Credit Portfolios Paul Kupiec¹ May 1, 2015

ABSTRACT

Most credit portfolios contain obligor concentration risk and yet international bank regulatory capital rules and many industry models assume perfect diversification. Multiple methods are available to calculate the approximate capital needs of a concentrated credit portfolio, but many of these involve advanced mathematical arguments, substantial computation time, and fail to clearly identify the most important credits causing concentration risk. In this article, I illustrate three approaches for calculating loss distributions and value-at-risk capital requirements. Of these the large exposure approach proposed by Kupiec (2015) is especially easy to implement. It produces accurate estimates of the economic capital required for a concentrated portfolio and immediately identifies the obligors most responsible for generating concentration risk.

Key Words: Portfolio diversification, idiosyncratic default risk, obligor concentration, Vasicek single common factor model of credit risk, credit value at risk, Basel bank capital requirements

¹ Resident Scholar, the American Enterprise Institute. The views in this paper are those of the authors alone. They do not represent the official views of the American Enterprise Institute. Email: <u>paul.kupiec@aei.org</u> phone: 202-862-7167

Capital for a Concentrated Credit Portfolio

I. Introduction

Diversification is perhaps the only "free lunch" offered in well-functioning financial markets. By merely selecting investments in the correct proportions, a financial institution can reduce unnecessary risk exposures and reduce the economic risk capital it needs to fund operations. Yet, despite the unquestioned benefits of holding a diversified portfolio, the literature on measuring credit risk diversification, to me at least, seems underdeveloped and mathematically daunting.

International bank regulators have focused on credit risk and regulatory capital for the last 15 years, and so one might think there would be useful guidance in the Basel capital regulations. However hard you look, you will be disappointed. For example, on concentration risk, Basel II only offers the following wisdom: "Risk concentrations are arguably the single most important cause of major problems in banks."

Maybe it is unfair to be so hard on the Basel Committee on this point, but then again, maybe not. If concentration is indeed the biggest risk facing banks, and the benefits to reducing concentration risk are available at minimal cost, why haven't international bank regulators devoted more time and energy on developing practical concentration risk methodologies that will improve all banks' safety and soundness? Perhaps one reason is that credit risk concentration is complex and can stem from many sources.

One source that gives rise to concentration risk is omitted common factors that drive defaults. If the true default correlation is driven by 3 factors, and a default correlation model assumes there is only one factor, any measurements produced by the model will have unmeasured concentration risk against the two missing factors. Another source of concentration risk could arise from the choice of an incorrect factor probability distribution (or so-called copula model) for default correlation. Use of the wrong distributional assumption (e.g. Gaussian instead of Student-t) will generate an unmeasured concentration of portfolio defaults in the tail of the portfolio's loss distribution. Concentration risk can also arise because of unmeasured correlation between credits' conditional probability of default and their loss given default. Unless it is accounted for, such a correlation could produce unmeasured concentrations in the tail of the loss distribution.² The final source of concentration risk I will mention is obligor exposure concentrations—a portfolio with large unhedged exposures to individual borrowers.

All of these sources of concentration risk are potentially important, and there are various ways a risk manager might deal with each of them. In this article I will focus on methods to measure the most straight-forward source of concentration risk in a credit portfolio—obligor concentration risk. Obligor

² See for example, the credit default model developed in Kupiec (2008).

concentration risk arises because individual borrowers have different loan sizes and loss rates. Obligor concentration risk is pervasive in practice and yet there is no widely accepted technique for measuring its impact on credit portfolio loss distributions.

A search of the literature will find various "granularity adjustments" that can be used in conjunction with the alternative credit default models including Basel II and other frameworks. These adjustments involve complex mathematical approximations and offer no intuitive appeal—at least not to me.

An alternative approach is to approximate the entire conditional loss distribution using a generalization of the loss-bucketing approach proposed in CDO and n-th to default CDS pricing models of Andersen, Sidenius and Basu (2003), and Hull and White (2004). This approaches make good intuitive sense, and it can be very accurate, but it requires a significant amount of numerical computation.

A third alternative is to use a new "largest exposure approach" for approximating the critical values of the portfolio loss distribution developed in Kupiec (2015). This approach is very simple and accurate and the answers are very intuitive. It also directly identifies the obligors that are the biggest source of concentration risk. Regardless of which method you prefer, the "bad news" is that properly accounting for obligor concentration risk has the potential to substantially increase the economic capital requirements for many portfolios.

An outline of the article follows. The influences of common factors that drive default correlations must first be removed before calculating the impact of obligor concentrations on the loss distribution and economic capital for a credit portfolio. There are various ways to model default correlation, but I will focus on the approach pioneered by Vasicek (1987, 1991) and adopted by the Basel Committee on Banking Supervision for use in Basel II. Section II discusses default correlation and the calculation of conditional default probabilities. The Basel II assumes that idiosyncratic default risk is fully diversified. Section III discusses alternative ways of measuring the true idiosyncratic risk that remains in credit portfolios because of obligor concentration risk and the associated methods for estimating value-at-risk economic capital. Section IV concludes.

II. The Vasicek model of default correlation

The Vasicek (1987, 1991) model of default correlation assumes that an individual credit's default is determined by the realized value of a random variable, \tilde{V}_i , with the following properties:

$$\widetilde{V}_{i} = \sqrt{\rho} \, \widetilde{e}_{M} + \sqrt{1 - \rho} \, \widetilde{e}_{i}
E(\widetilde{e}_{i} \, \widetilde{e}_{j}) = E(\widetilde{e}_{M} \, \widetilde{e}_{j}) = 0, \, \forall i, j,$$
(1)

If \tilde{e}_M and \tilde{e}_i have standard normal distributions, \tilde{V}_i also has a standard normal distribution.³ \tilde{V}_i is often interpreted as a proxy for the total market value of the obligor *i*'s firm. The common factor, \tilde{e}_M , induces correlation among the end-of-period total market values of different obligor's firms, $\rho = \frac{Cov(\tilde{V}_i, \tilde{V}_j)}{\sigma(\tilde{V}_i)\sigma(\tilde{V}_j)}$,

which in turn induces correlation among obligor defaults. Still, defaults are less than perfectly correlated because each firm has a latent idiosyncratic risk factor, \tilde{e}_i , that governs the default process.

Obligor *i* is assumed to default when $\tilde{V}_i < D$. The unconditional probability that credit *i* defaults is $PD = \Phi(D)$, where $\Phi(z)$ represents the value of the cumulative standard normal probability evaluated at *z*. All credits in a portfolio are assumed to have the same unconditional probability of default.

The probability that an individual credit defaults depends on the realized value of the common factor, \tilde{e}_M . For a given realized value, $\tilde{e}_M = e_M$, obligor *i* will default if the total market value of from *i* is less than the boundary *D*,

$$\left(\tilde{V}_{i}|\tilde{e}_{M}=e_{M}\right)=\sqrt{\rho}\;e_{M}+\sqrt{1-\rho}\;\tilde{e}_{i}< D,$$
(2)

$$\frac{D - \sqrt{\rho} e_M}{\sqrt{1 - \rho}} < \tilde{e}_i. \tag{3}$$

Since \tilde{e}_i has a standard normal distribution, the conditional probability of default is $\Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)$.

If all obligors in a credit portfolio have the same unconditional probability of default and the same default correlation parameter, then conditional on $\tilde{e}_M = e_M$, portfolio obligor defaults will be independent Bernoulli random variables, each with a default probability of $\Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)$ and a survival probability of $\left[1 - \Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right]^4$.

If there are *N* obligors in a credit portfolio then, conditional on the realized value of the common market factor, $\tilde{e}_M = e_M$, the number of portfolio defaults, *n*, has a binomial distribution,

$$Prob(defaults = n) = {\binom{N}{n}} \left[\Phi\left(\frac{D - \sqrt{\rho}e_M}{\sqrt{1 - \rho}}\right) \right]^n \left[1 - \Phi\left(\frac{D - \sqrt{\rho}e_M}{\sqrt{1 - \rho}}\right) \right]^{N-n}.$$
(4)

The components of this binomial probability are: (1) probability of any specific set of n credits defaulting,

or,

⁴ Conditional obligor defaults will be independent because default is determined by the realization of \tilde{e}_i , i =

³ $E(\widetilde{V}_i) = 0$, and $\sigma^2(\widetilde{V}_i) = E(\widetilde{V}_i^2) - E(\widetilde{V}_i)^2 = 1$.

 $^{\{1,2,3,\}ldots,N\}$ where N is the number of obligors in the portfolio and \tilde{e}_i is independent of \tilde{e}_j , for $i \neq j$.

$$\left[\Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right]^n \left[1-\Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)\right]^{N-n},\tag{5}$$

And, (2) the number of unique combinations of n defaults in a portfolio with N obligors,

$$\binom{N}{n} = \frac{N!}{n!(N-n)!}.$$
(6)

When loss exposures differ among credits, expressions (5) and (6) are used to generate the conditional portfolio loss distribution.

III. The Conditional Portfolio Loss Distribution

The credit risk model that underlies Basel II assumes that all credit portfolios are infinitely granular, meaning that there are an enormous number of obligors generating potential loss exposures that are approximately the same size. There is no allowance in the Basel II model for portfolios with large exposures to a relatively few borrowers. Basel II assumes that all idiosyncratic default risk has been diversified away. The consequence of this assumption is that the only source of default risk in the Basel II model is the common market factor. Each realization of the common market factor is associated with a single specific portfolio loss. In other words, the conditional loss rate distribution is not a distribution at all, but a single point. Portfolio loss outcomes are not so tidy in the real world.

To generate the true underlying conditional portfolio loss distribution, all potential conditional loss outcomes for a portfolio must be evaluated and ranked. In the most extreme example of portfolio risk exposure diversity, each credit generates a uniquely valued portfolio loss when it defaults. In this case, each possible combination of n obligor defaults will also generate a unique portfolio loss, and the total number of unique possible portfolio loss outcomes is,

$$\sum_{i=0}^{N} \frac{N!}{i!(N-i)!}.$$
(7)

As the number of obligors in the portfolio, *N*, increases, the total number of unique portfolio loss outcomes grows exponentially. For example, when there are 3 obligors in the portfolio, there are only 8 possible loss outcomes. For 10 obligors, 1024 possible loss outcomes must be evaluated. For a portfolio with 100 borrowers, there are 1.267×10^{30} possible portfolio loss outcomes. As *N* increases, it quickly becomes cumbersome if not impossible to evaluate and rank all these loss outcomes simultaneously.

Calculating the True Conditional Portfolio Loss Distribution

In the remainder of this section, I will consider a simple specific example of the steps that must be taken to construct the exact conditional portfolio loss distribution. Subsequently, I will use this example to illustrate two alternative approaches for approximating this distribution.

Consider constructing a portfolio loss distribution for a portfolio comprised of only 4 individual credits.⁵ Assume I have already conditioned on a common factor realization (e.g., the first percentile value, $e_M = -2.326$) and after conditioning, each credit has a 5 percent conditional probability of default.⁶ Table 2 lists the probabilities attached to each possible outcome: there could be no defaults; a single default, two defaults, three defaults or all credits in the portfolio could default.

There are multiple ways that the portfolio could experience 2 or 3 defaults. Table 3 lists the possibilities. The corresponding probabilities attached to each unique combination of multiple defaults are reported in Table 2. If all credits default, the loss will be 500; the probability of this unlikely outcome is 0.00000625.

Table 1: Potential Credit Losses		Table 2: F	Table 2: Probability of Alternative			Table 3: Potential Losses From Multiple Defaults				
			Outcomes		combinations of		combinations of			
	loss in		probability	number of	2 defaults	loss	3 defaults	loss		
credit ID	default	number of	of single	possible	1,2	160	1,2,3	360		
1	100	defaults	outcome	outcomes	1,3	300	1,2,4	300		
2	60	0	0.81450625	1	1,4	240	1,3,4	440		
3	200	1	0.04286875	4	2,3	260	2,3,4	400		
4	140	2	0.00225625	6	2,4	200				
		3	0.00011875	4	3,4	340				
		4	0.00000625	1						

The exact portfolio loss distribution is constructed by listing all the possible default outcomes and their associated probabilities of occurrence, and then ranking the list of outcomes from smallest to largest. The ranked losses and their associated probabilities comprise the portfolio's conditional loss distribution; the associated cumulative probabilities define the portfolio's cumulative loss distribution. Figure 1 is the portfolio's conditional loss probability distribution; Figure 2 is the portfolio's cumulative loss distribution.

⁶ The conditional probability of default in the Vasicek model (and Basel II) is, $\Phi\left(\frac{D-\sqrt{\rho}e_M}{\sqrt{1-\rho}}\right)$. It is also possible to construct the loss distribution when conditional loss probabilities differ among portfolio obligors. The process is the same, but there is added complexity. Single defaults may have unique conditional probabilities and so the unique combination of 2 and 3 defaults may also have unique probabilities.

⁵ Here I am constructing the portfolio loss distribution. To construct the loss rate distribution, I would use exposure at default for each credit to calculate portfolio loss rates and calculate the loss rate probability distribution in terms of possible portfolio loss rates.



The value-at-risk economic capital required by this portfolio is the capital needed to cover, at a minimum, the loss associated with a user-selected quantile of the portfolio loss distribution, typically 95- or 99- percent. Basel II requires 99.9-percent loss coverage. For this portfolio, a loss of 340 has a cumulative probability of 99.964-percent, so 340 is the required risk capital to meet the Basel II requirement of 99.9-percent loss coverage.⁷

The problem with calculating economic capital requirements using the true portfolio loss distribution in the presence of obligor concentration risk is that, if you add an additional credit to the portfolio, you must recalculate the entire loss distribution to find out how much capital the new portfolio requires. You cannot evaluate the new credit in isolation and evaluate how much extra capital will be required should you fund the new borrower.

The Loss-Bucketing Approach

The loss-bucketing approach generates an approximation for the portfolio loss distribution. The benefit of the loss-bucketing approach is that it limits the number of outcomes that must be considered simultaneously by building up an approximation for the loss distribution using an iterative algorithm.

To begin the loss-bucketing algorithm, one must first decide on a grid of loss buckets for the approximation. The finer the grid, the closer the approximation will be to the true distribution. The loss buckets need to cover the range of possible losses—in this case they must range between 0 and 500. The range of loss spanned by each loss bucket need not be equal; the buckets can vary in width. For this example, I will choose a 0 loss bucket, and then sequence of buckets of width 50.

The algorithm for constructing the approximation to the loss distribution has two parts. The first part calculates the probabilities that should be attached to each loss bucket. The second part calculates an estimate of the average loss that is associated with each loss bucket.

⁷ The next small potential portfolio loss outcome, 300, covers only 99.738-percent of possible loss outcomes.

The algorithm iterates through all possible default outcomes and assigns outcome probabilities to the appropriate loss bucket. At the end, probabilities are summed for each loss bucket. The process is illustrated in Table 4.

Consider first, the possibility of 0 defaults. This outcome generates 0 losses, and so the probability associated 0 defaults is entered into the 0 loss bucket. Next, the algorithm evaluates events associated with 1 default. There are 4 ways that a single obligor could default. Each event has a probability of 0.042869. Each of these events generates a unique loss, and the probability of the event is entered into the corresponding loss bucket. Since, in this example, the default of obligor 1 or obligor 2 generates a loss between 51 and 100, the sum of these event probabilities is entered into the corresponding loss bucket. The process continues as probabilities are assigned to loss buckets for all 2-default events, all 3-default events, and finally the default of all 4 obligors. The sum of the probabilities (across the number of defaults columns) yields the probability associated with each bucket.

						total	
loss	number of defaults					bucket	cummulative
bucket	0	1	2	3	4	probability	probability
0	0.814506					0.814506	0.814506
1 to 50						0.000000	0.814506
51 to 100		0.085738				0.085738	0.90024375
101 to 150		0.042869				0.042869	0.9431125
151 to 200		0.042869	0.004513			0.047381	0.99049375
201 to 250			0.002256			0.002256	0.99275
251 to 300			0.004513	0.000119		0.004631	0.99738125
301 to 350			0.002256			0.002256	0.9996375
351 to 400				0.000238		0.000238	0.999875
401 to 450				0.000119		0.000119	0.99999375
451 to 500					0.000006	0.000006	1

Table 4: Loss Bucketing Algorithm for Approximating Loss Distribution Probabilities

The second step of the process yields an estimate of the average loss associated with each bucket. The average loss can be taken as the mid-point of each bucket, but the approximation is improved by calculating the actual average loss associated with events in each loss bucket. In this stage, the average losses associated with each probability entry in Table 4 are calculated by adding up the losses associated with each event that contributes probability to a bucket, and then dividing the sum by the number of events that contributed to the bucket.⁸ For example, if two double-default events and a single three-

⁸ In a full-on computer algorithm, one would keep a third table to count the number of events in each bucket to calculate the bucket averages. To safe space, I have omitted that table.

default event contribute probability to a loss bucket, this counts as three events. These calculations are illustrated in Table 5.

						U
loss	total losses in bucket					loss per
bucket	0	1	2	3	4	event
0	0					0.0
1 to 50						0.0
51 to 100		160				80.0
101 to 150		140				140.0
151 to 200		200	360			186.7
201 to 250			240			240.0
251 to 300			560	300		286.7
301 to 350			340			340.0
351 to 400				760		380.0
401 to 450				440		440.0
451 to 500					500	500.0

Table 5: Average Loss Per Bucket Calculations

average

A comparison of the true conditional loss distribution to the loss distribution approximation generated by the loss-bucketing algorithm appears in Figure 3. The comparison shows that, provided that the loss bucketing grid is appropriately granular, the loss bucking algorithm will produce a close approximation to the true underlying loss distribution. The fewer the number of loss buckets, the greater the averaging that will take place and the less the approximation will look like the true underlying distribution. Regardless, the loss bucketing algorithm requires a lot of computation, especially when there are a large number of obligors in the portfolio.

If one uses this loss-bucketing algorithm to evaluate the 99.9-percent coverage economic capital required by this portfolio, the capital estimate is 340, and the estimate coverage rate is 99.964. Similar to the full loss distribution approach, when a borrower is added to the portfolio, the entire loss distribution approximation must be recalculated to estimate out how much additional capital the new portfolio requires. The marginal capital requirement cannot be evaluated the new credit in isolation form the rest of the portfolio.



Figure 3: True vs Approximate Conditional Cumulative Loss Distribution

The Largest Exposure Approximation

There is a much simpler approach for approximating value-at-risk capital for a portfolio loss distribution with concentration risk. The technique is fully developed in Kupiec (2015). To a very close approximation, the value-at-risk capital required to fund a portfolio with obligor concentrations is given by the sum of the q-largest portfolio loss exposures. The number of large exposures, q, that must be added up to calculate the portfolio capital 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.

In "simple" English, the algorithm is: (1) assume all exposures are identical and evaluate their cumulative binomial probability distribution using the conditional probability of default and solve for the number of defaults that are required to reach the desired quantile, q, of the cumulative portfolio loss rate distribution (k_q defaults); (2) calculate the sum of the k_q largest individual obligor losses in the true portfolio; (3) the sum of the k_q largest individual credit portfolio losses will have a true cumulative probability that is at least as large a q. This approximation is very accurate for typical value-at-risk capital coverage rates.

If all obligors generate the same identical loss in default, there are only 5 possible loss outcomes for a portfolio with 4 obligors, 0,1,2,3 or 4 defaults. The probabilities associated with each of these outcomes is given by expression 4. In the case of multiple defaults, no matter which credits default, the loss is the same, and the portfolio loss grows in proportion to the number of defaults. Table 6 reports the cumulative

probabilities for a binomial distribution with a probability of default equal to 5-percent.⁹ From these cumulative probabilities, it is readily apparent that more than 99.9-percent of all potential loss outcomes are smaller than the loss generated by an outcome with two obligor defaults when all obligor losses are identical.

But in the portfolio with obligor concentrations, all obligor losses are not equal. It turns out that by selecting the individual obligors that generate the largest two individual losses we can be sure that the sum of these losses will cover at least 99.9 percent of all losses in the portfolio with obligor concentration risk. The formal mathematical proof of this can be found in Kupiec (2015). So, for the 4-obligor example, the large exposure approximation indicates that the 99.9-percent economic capital requirement is 340, the sum of the losses generated by borrowers 3 and 4.

Table 6: Cumulative Binomial Probabilities

cumulative
probability
0.81451
0.98598
0.99952
0.99999
1.00000

Should a new credit be added to this portfolio, the economic capital can change for two reasons: (1) q, the number of the largest losses that must be summed could increase; or, (2) the new borrower could generate a larger loss than any of the large exposures currently determining the portfolio's economic capital in which case the new credit would replace the smallest loss in the summation. For example, Table 7 shows that when the portfolio grows from 4 to 5 obligors, the 99.9-percent loss threshold increases from 3 to 4 defaults. In this case the addition of a new borrower will increase the economic risk capital for the portfolio, and the magnitude of the increase would be equal to the larger of the new credit's potential loss or the loss generated by the next largest exposure [credit 1=100] that not currently determining the portfolio's economic capital.

For large portfolios, q, the number of large exposures that must be accumulated to estimate the economic capital is "sticky" meaning that the cumulative probability associated with a given value q does not change very much for high quantiles of the loss distribution as an additional credit is added to the portfolio. But when enough credits are added, q will increase, and there will be a threshold jump in economic capital required to cover the required percentage of potential portfolio losses.

⁹ Again, this is the conditional probability of default after controlling for the common factor(s) driving default.

number of cumulative					
defaults	probability				
0	0.773781				
1	0.977408				
2	0.998842				
3	0.999970				
4	1.000000				
5	1.000000				

_

Table 7: Cumulative Binomial Probabilities for a Uniform 5-Obligor Portfolio

The largest exposure approximation provides a very convenient and accurate method for calculating the capital required for a portfolio with obligor concentrations. It also identifies the exact obligors who are most responsible for generating the concentration risk in the portfolio. These credits are perhaps the highest priority candidates for hedging or other risk mitigation actions. The marginal economic capital required to fund a marginal credit or the capital released by new loss mitigation measures are easily calculated without the need to re-approximate the entire conditional portfolio loss rate distribution.

IV. Conclusion

Concentration risk is particularly important for many fixed income portfolios in part because the methods that have been available to evaluate concentration risk are non-intuitive and computationally cumbersome. In this brief overview, I have reviewed some alternative approaches for measuring concentration risk and calculating economic capital. Among these, the largest exposures approach proposed by Kupiec (2015) is by far the easiest to implement.

References

Andersen, Leif, Jakob Sidenius and Susanta Basu, (2003). "All your hedges in one basket," *Risk*, November 2003, pp. 67-72.

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

Gordy, Michael, (2003). "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules," *Journal of Financial Intermediation*, Vol. 12, No. 3, pp. 199-232.

Gordy, Michael, and James Marrone, (2012). "Granularity adjustment for mark-to-market credit risk models," *Journal of Banking and Finance*, Vol. 36, pp. 1896-1910.

Gordy, Michael, and Eva Lütkebohmert, (2013). "Granularity Adjustment for Regulatory Capital Assessment," Vol. 9, No. 3, pp. 33-71.

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

Hull, John, and Alan White, (2004). "Valuation of a CDO and an n-th to Default CDS Without Monte Carlo Simulation," *Journal of Derivatives* Vol. 12, No. 2, pp. 8-23.

Kupiec, Paul, (2008). "A Generalized Single Common Factor Model of Portfolio Credit Risk," *The Journal of Derivatives*, Vol. 15, No. 3, pp. 25-40, 2008.

Kupiec, Paul, (2015). "Portfolio Diversification in Concentrated Bond and Loan Portfolios," forthcoming, Journal of Investment Management, http://www.aei.org/publication/portfolio-diversification-inconcentrated-bond-and-loan-portfolios/

Martin, Richard, and Tom Wilde, (2002). "Unsystematic Credit Risk," Risk, Vol. 15, No.11, pp. 123-128.

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*, Vol. 15, No. 12, pp. 160-162.

Wilde, Tom, (2001). "Probing Granularity," Risk, Vol.14, No. 8, pp. 103-106.