Evaluating Design Choices in Economic Capital Modeling: A Loss Function Approach[1]

by

Nicholas M. Kiefer

and

C. Erik Larson

Abstract

This paper considers issues relating to the segmentation or grouping of credit exposures and the potential impact upon economic capital allocation and attribution. When discussing capital allocation, we refer to the assessment of total capital at the portfolio level, while our discussion of capital attribution focuses on getting capital assigned appropriately at the bucket level.

We emphasize that a loss or value function must be specified so as to quantify the gains and losses from choosing a more or less granular asset segmentation scheme. Our chosen loss function considers the trade-off between the decrease in sampling variance obtained by combining data to increase sample size and the bias resulting from characterizing unlike assets with the same default probability.

The implications are illustrated with several numerical examples that consider accuracy in the estimation of both portfolio-level and asset-level capital requirements. The suggested technique can be used to quantify whether a loss in accuracy from grouping or segmentation is outweighed by the decrease in variance of estimated capital. That is, the "loss" from grouping is small when the evaluation criterion is the accuracy of estimation of the required total capital; grouping is of more concern when we are interested in getting capital attributed correctly at the bucket level.

[1] *The statements made and views expressed herein are solely those of the authors, and do not represent official policies, statements, or views of the Office of the Comptroller of the Currency, the U.S. Department of the Treasury, or its staff*

Evaluating Design Choices in Economic Capital Modeling: A Loss Function Approach

by

Nicholas M. Kiefer

and

C. Erik Larson

I. Introduction

The concept of economic capital (also referred to as "risk capital" or "risk-based capital") is increasingly being adopted by banks and other financial institutions as a standard by which to determine the amount of capital needed to protect against financial distress in the event of unexpectedly large losses.

When calculating portfolio economic capital requirements, most models estimate critical values corresponding to extreme tail percentiles of a portfolio or whole-bank loss distribution. Economic capital requirements are then set to cover a measure of "unexpected loss," defined as the difference between the estimated mean of the loss distribution and the estimated loss level corresponding to the chosen critical tail percentile.

By their design, economic capital models are complex, and usually take as input the output of several other modeling exercises, including but not limited to the estimation of asset-level default probabilities (PD), loss given default (LGD) rates, and cross-asset correlations of these same parameters. Because these models are inherently complex, financial institutions must assess the risks of using them, as well as their associated driver-models. Model risk is defined for this purpose as the potential for loss from incorrect predictions or incorrect decisions resulting from the misuse of models. Such misuse usually occurs when a model is misapplied or its results are misinterpreted. Model risk is assessed in the context of the intended use of models and best-known practices used to build models. Credit risk decision models are evaluated with respect to sample design, modeling techniques, validation procedures, and re-validation procedures. This paper considers issues relating to the segmentation or grouping of credit exposures and the potential impact upon economic capital allocation and attribution. When discussing capital allocation, we refer to assessing total capital at the portfolio level, while our discussion of capital attribution refers to assigning capital appropriately at the bucket level. We discuss whether a model's logical structure fits its application. As referenced in OCC Bulletin 2000-16, "Risk Modeling — Model Validation," this assessment is essential to the first stage of model validation.

In most quantitative approaches to assessing expected loss and reserves, or the appropriate amount of economic capital to support a portfolio of assets, the risk ratings of assets and their associated estimates of PD and LGD are key inputs. PD and LGD can be estimated using a variety of techniques including simple descriptive statistical analysis, statistical and econometric regression models, and structural finance models. Whatever the approach, these metrics are almost impossible to estimate uniquely for each asset – there is simply not enough available information. Assets are therefore grouped, or segmented, into categories – buckets – and PDs are estimated by bucket. This results in PD estimates that are actually average PDs for assets within categories.

Since models that yield estimates of economic capital requirements are typically nonlinear in PD, how assets are grouped or bucketed has implications for economic capital. That is, estimation usually poses the following trade-off: As the size of each group increases, PD estimates of group averages, although more precise, are less relevant because more heterogeneous assets are grouped together. And as the size of each group decreases, PD estimates become less accurate.

This paper analyzes exactly this trade-off in the context of economic capital allocation and attribution. We employ the Basel II specification in our analysis since it is built upon a very simplified economic capital model, the Asymptotic Single Risk Factor (ASRF) model, which allows for marginal portfolio capital charges to be computed based upon exposure-level characteristics. (See Vasicek (1997) and Gordy (2000) for a detailed discussion of the ASRF.) The ASRF model enables a bank to calculate its minimum regulatory capital requirement for total portfolio credit risk as the sum of exposure-level capital charges, which in turn are strictly functions of PD, LGD, and a single portfolio-level asset correlation coefficient. However, this simplicity does not come without cost, since one can justify computing portfolio capital charges in this way only if (1) there is a single systematic risk factor driving correlations across obligors and (2) no exposure in a portfolio accounts for more than an arbitrarily small share of total exposure.

The Basel II implementation process is devoting considerable resources to defining standards and procedures by which to judge the readiness and ability of financial institutions to estimate loan characteristics including PD and LGD. Supervisory authorities are developing detailed specifications of the validation standards for these drivers. We therefore do not focus on issues relating to the validation of models used to estimate the drivers of, or inputs to, economic capital models. Our focus is instead on the application of the economic capital model, and we emphasize that a loss or value function must be specified so as to quantify the gains and losses from choosing a more or less granular scheme of asset segmentation. The numbers and types of alternate loss functions that could be specified are great, and they vary with the ultimate business uses of the capital estimates. Nevertheless, a natural starting point is to consider the mean-square error implications (MSE) of alternate segmentations or groupings of assets for economic capital. We illustrate the implications with several numerical examples.

II. Parameter Estimation

Consider first the case of two types of assets, with the second being the riskier (higher PD) asset. The question is whether to combine assets 1 and 2 into the same risk bucket for purposes of estimating PD and capital. Suppose there is a sample of experience on loans of each type, n_1 observations on loans of type 1 and n_2 on loans of type 2. Presumably (but not necessarily) $n_1>n_2$, so that there are fewer of the riskier type of asset. Let x_1 and x_2 be the observed average default rates of assets 1 and 2. Now, suppose that x_1 and x_2 are normally distributed with mean vector θ and variance matrix Σ. This makes sense if n_1 and n_2 are fairly large, or if x_1 and x_2 are suitable transformations of the default rates, for example, logits. We proceed with the actual rates, so that the situation is one of estimation of two binomial probabilities, noting that the results easily apply more generally. In this case the variance has a simple structure, with

$$\Sigma_{11} = \theta_1(1-\theta_1)/n_1,\ \Sigma_{22} = \theta_2(1-\theta_2)/n_2.$$

To simplify matters, we will assume here that $\Sigma_{12} = \Sigma_{21} = 0$.

The single "restricted" estimator, x_r, that results from combining type 1 and type 2 assets into one group is given by

$$x_r = (n_1x_1+n_2x_2)/n, \text{ where } n = n_1+n_2.$$

Its expectation is

$$E[x_r] = (n_1\theta_1 + n_2\theta_2)/n$$

The biases of x_r as an estimator of θ_1 and θ_2 are

$$E(x_r-\theta_1) = n_2(\theta_2 - \theta_1)/n, \quad E(x_r-\theta_2) = -n_1(\theta_2 - \theta_1)/n.$$

These are sensible: the higher risk asset has an underestimated PD and the lower risk an overestimated PD, and the position of the average between these two PDs depends on the relative sample sizes. The gain from allowing this bias is a variance reduction relative to the unrestricted estimator. The variance of x_r is

$$V(x_r) = (n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}$$
$$= (n_1\theta_1(1-\theta_1) + n_2\theta_2(1-\theta_2))/n^2$$
$$= (n_1\theta_1(1-\theta_1) + n_2\,\theta_2(1-\theta_2))/n^2$$

III. Estimating Capital Requirements

Rather than simply considering the variability or bias in estimation of PD, we want to focus on the variability in estimation of risk capital. As mentioned earlier, we will consider capital to be determined by the risk weight formula for corporate, sovereign, and bank (CSB) exposures, which is specified in the proposed revisions to the Basel accord (BIS, 2004). Actually, we will use a somewhat simplified version of the Basel II function, considering the case where asset maturity is fixed at one year and LGD=100%, .

Let $W(\theta)$: $[0,1]\rightarrow[0,1]$ denote the curve giving the capital risk weight (in fractions of loss given default, LGD) as a function of the probability of default. We have that

$$W(\theta) = N[G(\theta)\,(1-R)^{-0.5} + G(0.999)\,(R/(1-R))^{-0.5}] - \theta$$

where $R = 0.12(1+EXP(-50\theta))$, $N(x)$ denotes the cumulative distribution function for a standard normal random variable, and $G(z)$ denotes the inverse cumulative normal distribution. We have made a further simplification by approximating the term EXP(-50) appearing in the published formula by zero. The actual value is less than 10^{-20}. Note that this risk weight curve is generally a concave function in PD, as illustrated in figure 1.

Figure 1: Capital Risk Weight Curve

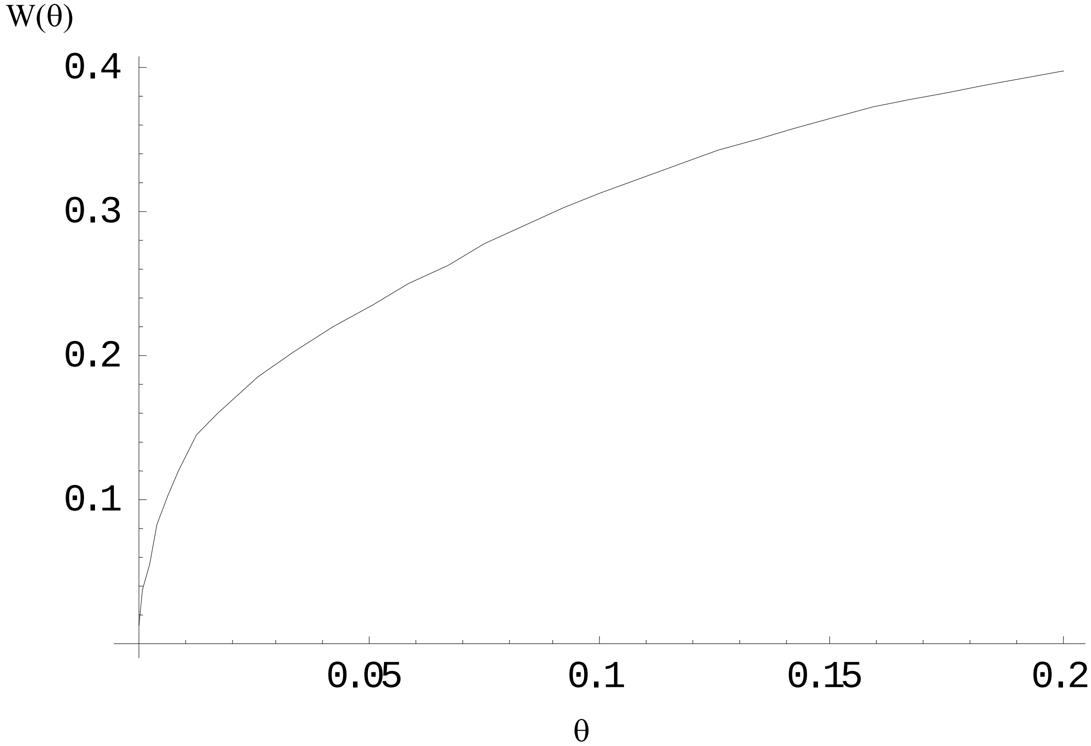

It is interesting to note at the outset that, while in general the unrestricted observed default rate is an unbiased estimator for the true group default probability θ, *the risk weight corresponding to the observed default rate, W(x), does not yield an unbiased estimate of the risk weight corresponding to the true default probability, W(θ).* Indeed, since the risk curve is concave, we have E[W(x)] < W(E[x]) = W(θ), by Jensen's inequality. Thus, "plugging in" an unbiased estimator for PD and evaluating the risk curve there leads to a higher-than-appropriate capital estimate.[2]

---

[2] To see this, take a Taylor series expansion of $W(x_1)$ around the true default rate θ_1. This yields

$$W(x_1) = W(\theta_1) + \partial W(\theta_1)/\partial x_1 * (x_1-\theta_1) + (1/2)\partial^2 W(\theta_1)/\partial x_1^2 * (x_1-\theta_1)^2 + r$$

where r is small (with a maximum order of $(x_1-\theta_1)^3$). Taking expectations gives

$$E[W(x_1)] = W(\theta_1) + (1/2)\, \partial^2 W(\theta_1)/\partial \theta_1^2 * V(x_1)$$

IV. Loss Functions

We are now at a point where we can discuss the alternative loss functions that could be considered when assessing the consequences of bucketing decisions on economic capital estimates. We must distinguish between capital allocation and capital attribution. When discussing capital allocation, the corresponding loss function will focus on the variation in the average risk weight across buckets. In contrast, when considering capital attribution, a loss function for assessing attributed capital will be driven by a weighted average of variations in bucket-level capital risk weights.

Capital Allocation

The average capital risk weight for our portfolio containing n_1 assets of type 1 and n_2 assets of type 2 is given by:

$$(n_1 W(\theta_1) + n_2 W(\theta_2))/n$$

It will also be useful to use a quadratic approximation to the concave W() function:

$$W(x) = ax - bx^2 + k$$

Using this approximation, the average capital risk weight, when evaluated using the unrestricted estimates x_1 and x_2, has expected value

$$(n_1/n)(aE[x_1] - bE[x_1]^2 - b\Sigma_{11} + k) + (n_2/n)(aE[x_2] - bE[x_2]^2 - b\Sigma_{22} + k)$$

$$= (n_1/n)(a\theta_1 - b\theta_1^2 - b\Sigma_{11}) + (n_2/n)(a\theta_2 - b\theta_2^2 - b\Sigma_{22}) + k$$

If the default probabilities θ were known, the average capital risk weight would be correctly calculated as

$$(n_1/n)(a\theta_1 - b\theta_1^2) + (n_2/n)\,(a\theta_2 - b\theta_2^2) + k$$

Hence the bias in the average capital risk weight is negative and equal to

$$-b(n_1\Sigma_{11}+n_2\Sigma_{22})/n.$$

Similarly, we calculate the variance of the average portfolio risk weight to be

$$E[n_1(a(x_1-\theta_1) - b(x_1^2-\theta_1^2))+ n_2(a(x_2-\theta_2) - b(x_2^2-\theta_2^2))]^2/n^2$$

---

since $E(x_1-\theta_1) = 0$. Since W is concave, the second term is negative and the random variable $W(x_1)$ has expectation smaller than $W(E[x_1]) = W(\theta_1)$. Kiefer and Larson (2003) investigate this bias in detail and propose corrections.

Using the assumed independence of x_1 and x_2, and noting that the normal third moments are zero and fourth moments are $3\Sigma_{jj}^2$, this simplifies to

$$(n_1^2 (a^2\Sigma_{11} + b^2(3\Sigma_{11}^2)) + n_2^2 (a^2\Sigma_{22} + b^2(3\Sigma_{22}^2)))/n^2$$

We now have enough information to show that when risk weights are calculated by plugging the unbiased, unrestricted estimators into the W function (as envisioned by Basel II), the mean square error in the average risk weight is given by

$$\mathbf{MSE_u} = (b^2(n_1\Sigma_{11}+n_2\Sigma_{22})^2 + n_1^2 (a^2\Sigma_{11} + 3b^2\Sigma_{11}^2) + n_2^2 (a^2\Sigma_{22} + 3b^2\Sigma_{22}^2))/n^2.$$

We now turn to the average capital risk weight that is obtained using the restricted estimator that combines assets into a single bucket. We consider the calculation of W at x_r. Again using our quadratic approximation to W, taking expectations yields

$$E[W(x_r)] = (aE[x_r] - b(E[x_r]^2 - bV(x_r)+k).$$

Recall that $E[x_r]= (n_1\theta_1 + n_2\theta_2)/n$ and that $V(x_r) = (n_1^2/n^2) \Sigma_{11} + (n_2^2/n^2) \Sigma_{22}$.

Thus, the bias in using the restricted estimator $W(x_r)$ for $W(\theta)$ is given by

$$a((n_1\theta_1 + n_2\theta_2)/n) - b((n_1\theta_1 + n_2\theta_2)^2/n^2) - b((n_1^2/n^2) \Sigma_{11} + (n_2^2/n^2) \Sigma_{22})$$
$$- (n_1/n)(a\theta_1 - b\theta_1^2) - (n_2/n)(a\theta_2 - b\theta_2^2).$$

The variance of $W(x_r)$ is given by

$$V(W(x_r)) = a^2V(x_r) + b^2(3V(x_r)^2)$$
$$= a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2) \Sigma_{11} + (n_2^2/n^2) \Sigma_{22})^2)$$

and the $\mathbf{MSE_r}$ is of course the variance plus the squared bias.

The best quadratic approximation to $W(\theta)$ around $\theta = 0.05$ is given by

$$W(\theta) = 0.130922 + 2.33006\,\theta - 5.17491\,\theta^2$$

Note that this quadratic approximation is quite accurate, with a maximum absolute relative error of less than 0.2 percent for $0.015<\theta< 0.1$ (i.e., the maximum error as a fraction of the actual $W(\theta)$ is less than 0.002)

Figures 2 and 3 illustrate the impact of choosing to combine or segment asset classes for the purposes of allocating economic capital. The figures graph the difference between the two mean square error measures of the average risk weight as functions of θ_1 and θ_2, the true rates of default for the two asset classes in the portfolio.

The surfaces have been shaded to illustrate the regions where, for the indicated portfolio sizes, the difference between θ_1 and θ_2 results in either positive or negative differences in restricted less unrestricted MSE. When the difference is positive, a granular bucketing system is to be preferred to one which pools asset types for the purposes of minimizing MSE in total capital allocation. When the difference is negative, a pooling of asset types results in lower MSE.

Comparing figures 2 and 3 illustrates the impact of larger sample sizes. We see that, as expected, the restriction is better when the range of PD values for each bucket is small. Larger sample sizes lead to restrictions being less desirable.

Capital Attribution

Capital attribution is concerned with bucket-level or segment-level accuracy in estimation. We therefore want to formulate a loss function that is sensitive to variation in bucket-specific estimates of risk weights.

When attributing capital to each of our two assets, using the unrestricted estimators, the expected value of the bucket-specific risk weights are given by

$$E[W(x_1)] = (aE[x_1] - bE[x_1]^2 - b\Sigma_{11} + k) = (a\theta_1 - b\theta_1^2 - b\Sigma_{11} + k)$$

$$E[W(x_2)] = (aE[x_2] - bE[x_2]^2 - b\Sigma_{22} + k) = (a\theta_2 - b\theta_2^2 - b\Sigma_{22} + k)$$

If the true segment-specific default rates were known, then the risk weights would be computed as

$$W(\theta_1) = a\theta_1 - b\theta_1^2 + k$$

$$W(\theta_2) = a\theta_2 - b\theta_2^2 + k$$

This allows us to compute the unrestricted estimate bucket-level risk weight biases as

$$E[W(x_1)-W(\theta_1)] = -b\Sigma_{11}$$

$$E[W(x_2)-W(\theta_2)] = -b\Sigma_{22}.$$

The variances of the bucket-level risk weight estimates are given by

$$V[W(x_1)] = E[(a(x_1-\theta_1) - b(x_1^2-\theta_1^2))^2]$$

$$V[W(x_2)] = E[(a(x_2-\theta_2) - b(x_2^2-\theta_2^2))^2]$$

Figure 2: Restricted MSE Minus Unrestricted MSE of Allocated Capital

(n_1 = 500, n_2 = 100)

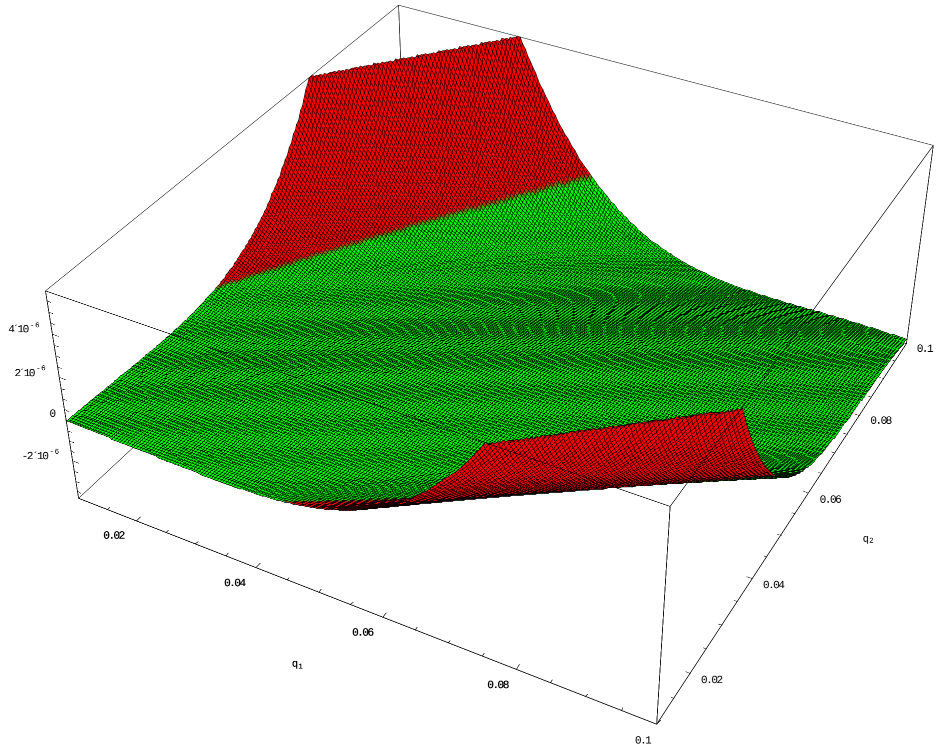

Figure 3: Restricted MSE Minus Unrestricted MSE of Allocated Capital

(n_1 = 1000, n_2 = 200)

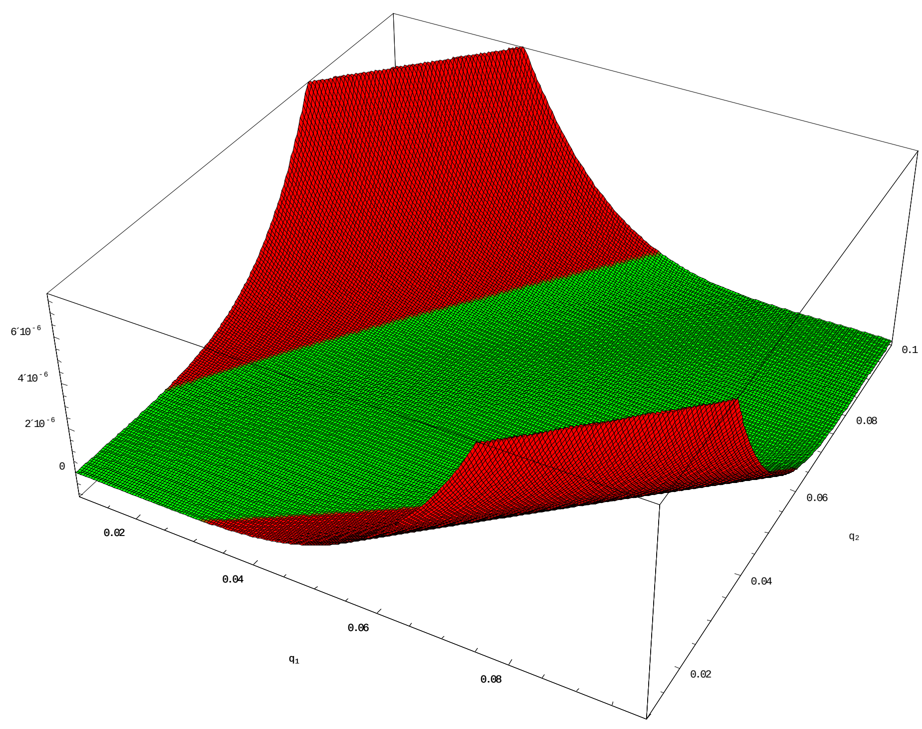

Since the normal third moments are zero and fourth moments are $3\,\Sigma_{jj}^2$, these simplify to

$$V[W(x_1)] = (a^2\Sigma_{11} + b^2(3\Sigma_{11}^2))$$

$$V[W(x_2)] = (a^2\Sigma_{22} + b^2(3\Sigma_{22}^2))$$

By adding the squared bias, the bucket-level unrestricted mean square errors are given by

$$MSE_{u1}= a^2\Sigma_{11} + 4b^2\Sigma_{11}^2$$

$$MSE_{u2}= a^2\Sigma_{22} + 4b^2\Sigma_{22}^2$$

Which allows us to compute the weighted-average unrestricted MSE across buckets as

$\mathbf{MSE_u} = (n_1/n)(a^2\Sigma_{11} + 4b^2\Sigma_{11}^2) + (n_2/n)(a^2\Sigma_{22} + 4b^2\Sigma_{22}^2)$

Turning to the restricted estimator, we have from our previous work that

$$E[W(x_r)] = a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})+k.$$

and

$$V(W(x_r)) = a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})^2)$$

We compute the restricted estimate bucket-level risk weight biases as

$E[W(x_r)-W(\theta_1)] =$

$$a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_1 - b\theta_1^2)$$

$E[W(x_r)-W(\theta_2)] =$

$$a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_2 - b\theta_2^2)$$

Again, by adding the variance and squared bias, the bucket-level MSEs from using the restricted estimator are given by

$$MSE_{r1}= a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})^2) +$$
$$(a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_1 - b\theta_1^2))^2$$

$$MSE_{r2}= a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})^2) +$$
$$(a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_2 - b\theta_2^2))^2$$

which allows for the weighted-average restricted MSE across buckets to be computed as

$$\mathbf{MSE_r} = (n_1/n)(a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})^2) +$$
$$(a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_1 - b\theta_1^2))^2) +$$
$$(n_2/n)(a^2((n_1^2\Sigma_{11} + n_2^2\Sigma_{22})/n^2)+ b^2(3((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22})^2) +$$
$$(a((n_1\theta_1 + n_2\theta_2)/n) - b(((n_1\theta_1 + n_2\theta_2)/n)^2 - b((n_1^2/n^2)\,\Sigma_{11} + (n_2^2/n^2)\,\Sigma_{22}) - (a\theta_2 - b\theta_2^2))^2)$$

Figures 4 and 5 illustrate the differences in the MSE of risk weights that arise for various combinations of θ_1 and θ_2. Again, we see that the restrictions are desirable only when the range of PDs is small. Here, in contrast to the case of total capital, it is not only the difference between the PDs that matters. When PDs are small, restrictions are less desirable for a given distance between them. Concern for bucket-level accuracy will lead to less combining of estimators.

IV. Conclusions

This paper illustrates an approach to capital model assessment by considering the following trade-off: A bank can decrease sampling variance by combining data to increase sample size, but as the bank increases sampling size, its estimates become less accurate because increasingly unlike assets are assigned the same default probability. We considered accuracy in the estimation of both portfolio-level and asset-level capital requirements using a specification from the proposed revisions to the Basel accord.

Our technique can be used to quantify whether the decrease in variance of estimated capital outweighs the loss of accuracy that results from making segments more heterogeneous. Although these numbers are specific to the example, it is likely that the relative ranking of the criteria holds more generally. That is, the "loss" from grouping is small when the evaluation criterion is the accuracy of estimation of the required total capital; grouping is of more concern when we are interested in getting capital attributed correctly at the bucket level.

Note that we have not here suggested practical methods for deciding the granularity of a bucketing procedure. We have simply considered the effects of using different criteria to judge the effects of pooling buckets. A classical approach is to "pretest," perhaps with a t-test for differences in means, and then decide whether to pool on the outcome of such a test (Mosteller, 1948) Classically, the pretest is done on the difference between parameter estimates. The pretest, if desired, might be better done on the estimated capital requirements directly.

Figure 4: Restricted MSE Minus Unrestricted MSE of Attributed Capital

(n_1 = 500, n_2 = 100)

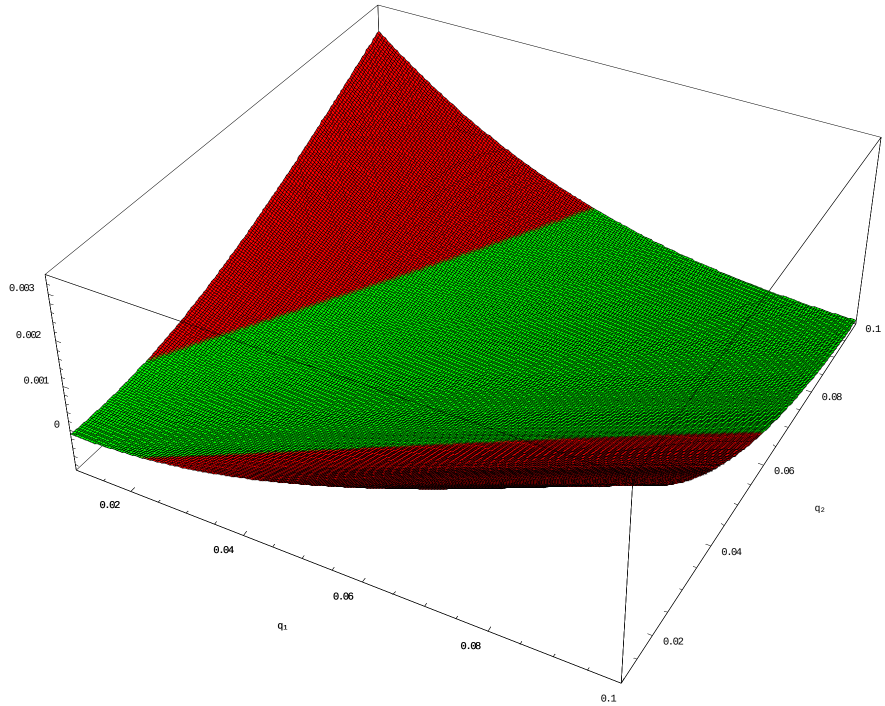

Figure 5: Restricted MSE Minus Unrestricted MSE of Attributed Capital

(n_1 = 1000, n_2 = 200)

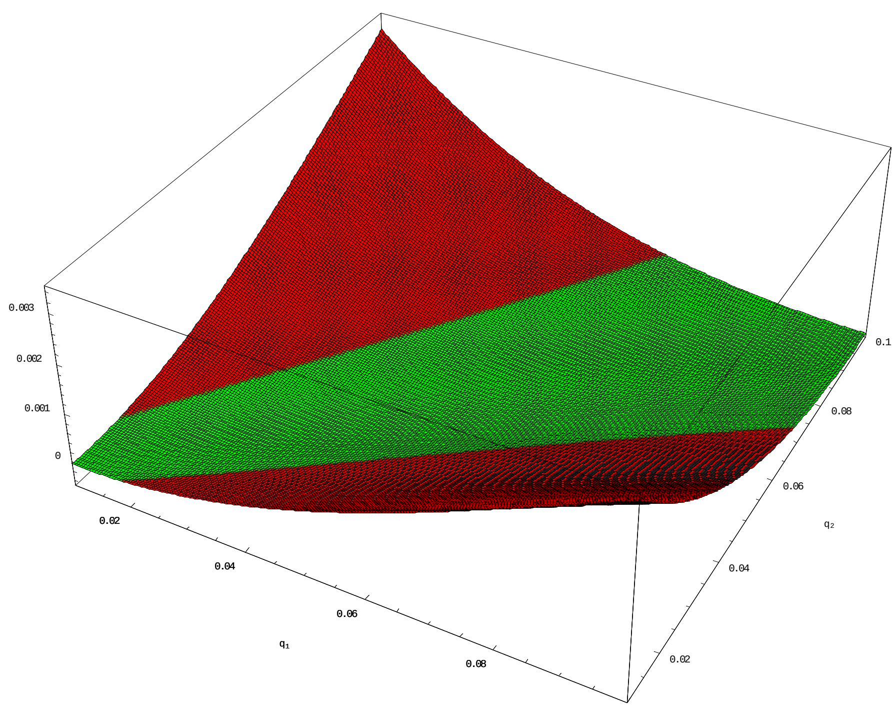

References:

Basel Committee on Banking Supervision (2004). "International Convergence of Capital Measurement and Capital Standards: A Revised Framework," Basel: Bank for International Settlements.

Gordy, Michael B. (2000). "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules," Working Paper, Board of Governors of the Federal Reserve System.

Kiefer, Nicholas M., and C. Erik Larson (2003). "Biases in Default Estimation and Capital Allocations Under Basel II," Working Paper, Cornell University.

Mosteller, F. (1948). “On Pooling Data,” *Journal of the American Statistical Association*, 43, 231-242.

Vasicek, Oldrich A. (1997). “The Loan Loss Distribution,” Technical Report, KMV Corporation.