Risk Integrated

3 July 2013

# Enterprise Risk Management and CRE Lending

## Introduction

Five years after the worst of the financial crisis, companies are moving from the highly re-active patching of their risk management infrastructure to considering the bigger picture of enterprise risk management. Enterprise risk management (ERM) is the set of policies, controls and risk measurement tools put in place so that an institution can coherently and consistently manage its risks. This paper discusses how ERM can be achieved while still giving the commercial real estate lending teams useful risk tool that meet the special requirements of commercial real estate (CRE) lending.

The implementation of ERM is difficult because it means bringing together risk management both vertically up through the different risk measurement functions within one asset-class and horizontally across all asset-classes. Within one asset-class, vertical integration requires consistency for all the risk measurement functions applied to that asset e.g., from structuring, pricing and approvals through to portfolio reporting, capital management and stress-testing. However, in most institutions the data and models are highly fragmented, especially for CRE. For example, there may be five different risk models in use: one each for underwriting, pricing, grading, economic capital and stress-testing. One consequence is that the different models produce different risk estimates for the same asset. Another consequence is that the data entered into the models is scattered across spreadsheets and disparate systems, requiring multiple rounds of data entry and causing inconsistent reporting.

Horizontally, the challenge is to measure risk consistently across the different asset-classes. It is common for there to be a patchwork of models for grading different assets and typically the models are not consistent. An example of inconsistency within CRE is that a loan to a property with 49% office income and 51% retail income may be rated very differently than a loan to a property with 51% office income and 49% retail income because the retail and office models are separate.

A more subtle set of horizontal problems is embedded in the details of the underlying risk methodology for each model e.g., whether the rating philosophy is Point-in-Time vs. Through-the-Cycle vs. Through-the-Life, or multi-year vs. single-year[1]. These subtleties can have a

[1] Models can be inconsistent as to whether they are attempting to measure the risk at a point in time (PIT), through the life (TTL) or through the cycle (TTC). They can also be inconsistent in the way they

significant effect on the risk numbers reported and therefore a significant effect on pricing, provisions and capital.

One of the greatest horizontal challenges is estimating the loss correlations across asset-classes, e.g., the correlation between loss on a CRE deal and loss in retail lending. These correlations are important for estimating the required economic capital, for stress-testing, and, in some cases, for pricing. Often the estimation of correlations is done in yet another separate model, but as discussed later it is also possible to integrate loss correlation with the rest of the risk analysis.

## The Special Challenges Posed for ERM by CRE

Before we go-on to look at solutions to these challenges, let us first consider why it is especially difficult to implement ERM for CRE lending. Unlike lending to retail or plain corporate customers, CRE lending is long-term and there is a high degree of structure in the underlying asset. It is also greatly affected by the overall market conditions and there is a strong linkage between the cause of default and the degree of subsequent loss.

From an ERM perspective this has two important consequences. One is that the long-term risk profile is very unlike that of the shorter term, less structured asset-classes: a typical CRE default profile, as shown below, is highly structured and has spikes of risk corresponding to events in the deal such as lease expirations, changes in the interest rate from fixed to floating and refinancing risk. Each of these spikes has a different correlation with the market and with the other assets.

The other consequence of CRE's long-term structured nature is that the scorecard approaches frequently used to assess the risk for retail and C&I loans do not work well for the multi-year complexity of CRE deals and therefore an approach such as cashflow simulation is required to capture the interdependencies. The challenge then is to use cashflow simulation in such a way that it is consistent with the models used for other classes and does not violate the integrity of the overall ERM framework.

look at multi-year risks when bringing them back to equivalent one-year risks, and they can simply be inconsistent in the way they map the grade to the probability of default (PD), loss given default (LGD) and/or expected loss (EL).

*Example of the Structured Risk Profile for a CRE Loan*

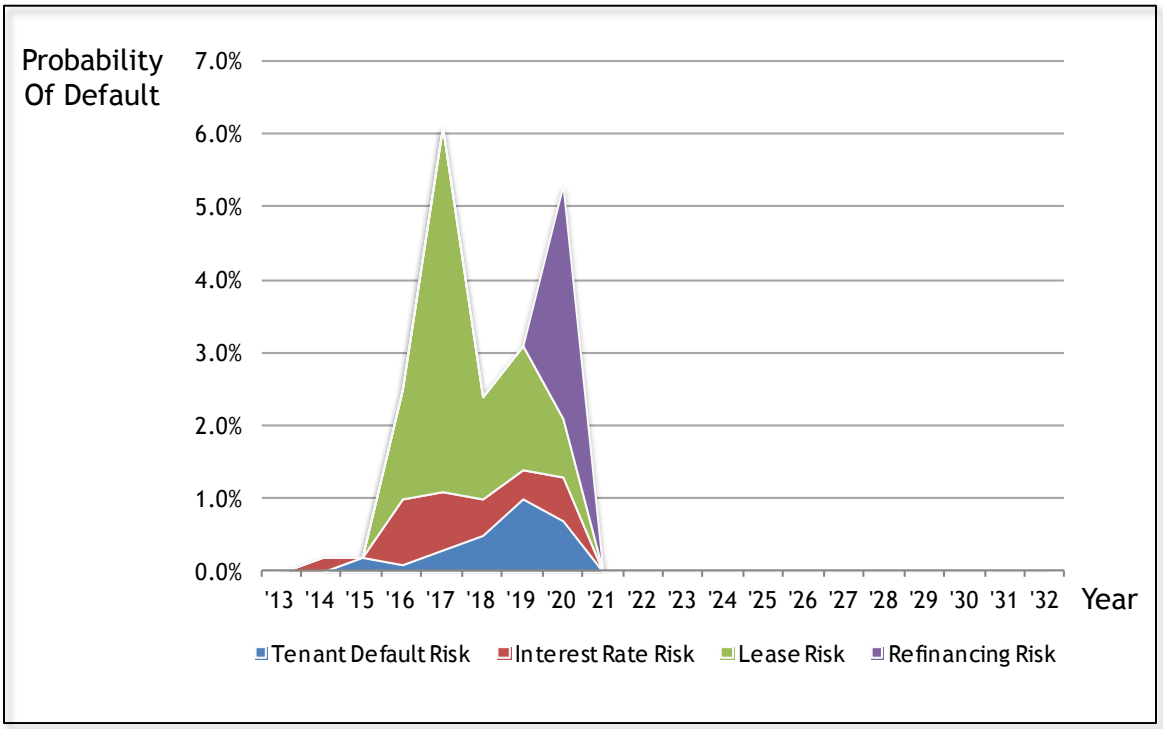


## Achieving Vertical Integration

For vertical integration the ideal is that for a given asset, the risk results should be the same whether it is being assessed for underwriting, pricing, or as part of stress-testing[2]. For vertical integration one of the main problems is that the models typically used are not sufficiently flexible for all the required tasks, so that instead of there being one consistent model for all tasks, there is one model per task. This is especially the case for CRE and a typical situation in CRE lending is described in Appendix 1.

Risk Integrated's solution for vertical integration is to ensure that there is a single set of models and central data-base for all stages of analysis. There are three key elements to the approach:

1. Use a model which is comprehensive and has sufficient flexibility and detail to provide all of the following:
    - the cashflows desired by underwriters
    - the risk grades for PD and LGD

[2] the base case used for grading and the base case used for stress reporting should be the same but of course the results conditional on a stress should be different from the base-case grading results.

- the regulatory capital
- the loan-loss provisions
- the risk-theoretical price
- the multi-year loss profile required for stress-testing
- the loss correlations used for economic capital

This model is built to allow users to enter different levels of deal information depending on the data available and the purpose. Fully detailed information of is entered for deal underwriting, when the lending officer has full information and there is the power to accept, reject, restructure or re-price the deal. Limited data is entered when assessing new portfolios or for initial deal structuring. This flexibility allows a single model to be used for all functions and allows data to be added incrementally as it becomes available.

Although the use of the model is relatively straightforward for any single deal, behind the scenes the underlying model needs to be relatively sophisticated so that it is sufficiently flexible to cater for all deal types and to cater for all contingencies such as tenant defaults. In order to give the bank faith in the sophisticated underlying model, it is fully exposed to their analysts to understand and change if desired.

2. The second key is to have a centralized set of data, accessed with web-technology, and with maximum flexibility for adding to the data structure as required[3]. The sources must include data typed into web-forms, lease information uploaded from spreadsheets, and loan information updated automatically from the central banking systems. Importantly, because the model is also used for generating the cashflows during underwriting, the lenders have a strong incentive to type-in accurate detailed data because the results will be shown as part of the credit review. The consequence is that there is a central set of detailed, accurate, and up-to-date information for generating portfolio reports for senior managers, regulators and investors.
3. The third key piece of functionality needed for vertical integration is the ability for a portfolio manager to run a batch of deals (e.g., all live deals) and re-grade them using the latest models, forecasts and data without requiring the lending teams to manually re-grade their assets. This functionality means that the portfolio manager has use of the full granularity that is available for assessing individual deals and same rich data-set.

[3] For this we use an XML data structure in an open SQL database with blending to collate multiple data-sources. XML also allows new information to be captured quickly, e.g., for new markets or custom data fields, without the need to refactor a database.

These principles underpin the dataflow embedded in Risk Integrated's Specialized Finance System, as illustrated below.

***Dataflow, Model Management and Reporting in the Specialized Finance System***

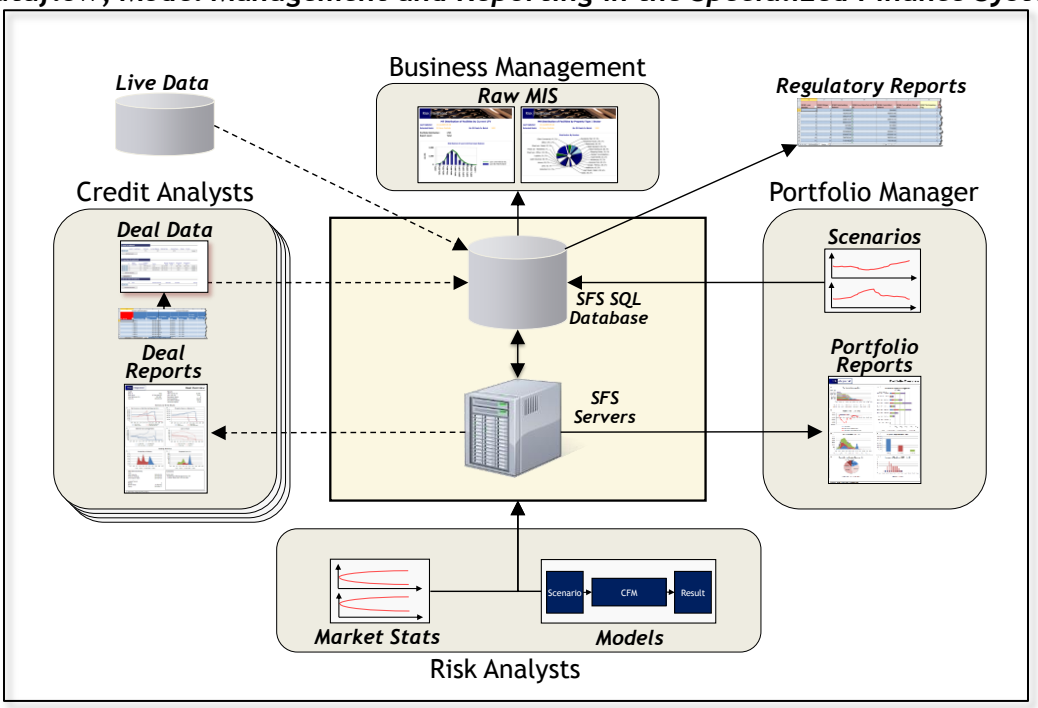


Within the framework of the SFS there are four main user groups: the risk analysts, credit analysts, portfolio managers and business managers:

- The risk analysts control the models, the risk parameters and the forecasts. Once they approve and load a model into the system it becomes available to the other users. The other users can be granted different levels of access to the system. Typically the portfolio analyst has access to all models and settings but the credit analysts can use only one model as specified by the risk analysts.
- The credit analysts use the system to assess individual deals. They type deal-data into web screens, upload it from other systems and/or upload it from templates, e.g., spreadsheets. This data is stored in the SQL database and then used when the credit analyst wants to generate the cashflows and risk statistics. Once the credit analyst has tried different financing combinations and is satisfied with the risk/return profile, he or she can print out a credit application form containing all the financial and risk projections for the deal. This same information is also stored in the database for use by the portfolio manager and business managers.

- The portfolio manager can select sets of deals to be run in a batch, e.g., all live deals, or all deals flagged as being in the pipeline, or all deals with a particular tenant. That batch can be run against any model and set of forecasts, including stressed forecasts. This means that the portfolio manager can easily and quickly run stress-tests and what-if analyses without burdening the credit teams.
- The business managers can of course get the same loan-level and portfolio-level reports being generated by the credit teams and portfolio managers, but they can also get detailed management information (MI) reports using the raw data that was captured during the underwriting process. For example they can see the total balance lent to each sector and geography, they can see the LTV of the deals originated by each team, and they can see the distribution of exit yields for the loans maturing in one, two or three years. This data can be presented with minimal statistical analysis[4], but because it is detailed and reliable, it can strongly inform the intuition of the businesses managers about the level of risk in the portfolio.

The principles described above can be used to achieve vertical integration for any asset-class, but they are especially important for CRE where there is so much information that has a meaningful impact on the understanding of the risk profile.

## Achieving Horizontal Integration

For horizontal integration across the different asset-classes we need to address the following issues:

- No jump in grade when an asset is categorized as one type vs. another
- Consistency in rating philosophy: e.g., PIT vs. TTC
- Consistency in mapping multi-year risks
- Meaningful calculation of correlations
- Pricing

### *Avoiding Discontinuities in Grades*

Jumps in grades can be avoided by designing the models to be continuous and avoiding strict categorization so far as possible. For example rather than classifying the asset as being in a particular sector and geography, recognize the degree to which it is related to each sector and geography and the extent to which it is, for example, under-construction, being rented-up or stabilized. This continuity is natural in a full cashflow simulation model because each property

[4] For example graphs of LTV require simply dividing the current balance by the current value, whereas estimating exit yields requires more complex cashflow projection, and risk estimates require full statistical analysis.

and lease responds to its own market influences, e.g., all the retail units respond to changes in the retail rental market and office units respond to the office rental market. Similarly, large leases naturally have a larger impact than small leases so it is not necessary to have a category of "anchor tenant", instead we just use the amount of rent for each tenant and larger tenants naturally have a larger impact on the results.

***Ensuring Consistent Rating Philosophies***

Ensuring consistency in rating philosophy between different models can be difficult, especially since there is often inconsistency within any individual philosophy. Cashflow simulation models naturally behave as multi-year Point-in-Time models. The results from cashflow simulation can easily be converted to Through-the-Cycle grades, but because cashflow simulation naturally takes into account the whether the market is currently high or low, the results are much more stable than typical "Through-the-Cycle" scorecards[5].

***Mapping Multi-year Risk Profiles to a Single Grade***

The mapping of multi-year risks becomes an issue when looking at CRE deals because CRE deals have an irregular risk profile over their lives, whereas C&I loans are assumed to have a smooth increase in risk over time. For conventional C&I loans, the risk in year-one completely defines the assumed risk for all subsequent years, therefore it has become customary for purposes like Basel capital to expect that a single number will be used to represent "the grade". To create a single grade for CRE there must be a way to average the fundamental multi-year risk profile into a single number. There are multiple ways of doing this averaging, and it primarily depends on the purpose of the final grade, but the final single number grade should be directly comparable with the grade for vanilla C&I loans[6]. However, if the models and data are vertically integrated as discussed earlier, the single grade becomes less important because multi-year risks such as stress losses can be estimated directly rather than being projected from the single "one-year" grade numbers for the probability of default and loss given default.

***Estimating Correlations***

The estimation of loss correlations typically uses two approaches: top-down and bottom-up. Top-down calculation looks at the history of losses on overall portfolios and estimates how the losses moved together. This has the great advantage of giving direct empirical evidence, but

---

[5] For example most "Through the Cycle" (TTC) scorecards keep the models' weights fixed over time, but the outputs vary widely because the inputs vary widely. Consider a CRE scorecard that assigned PD based simply on the Loan-to-Value (LTV). At the top of the market the scorecard would assign a relatively low PD to a loan with a 65% LTV. At the bottom of the market the same loan might now have a 95% LTV and now be assigned a very high PD. The result is that typical "TTC" models produce very unstable results and assign low risk, capital and reserves at the top of the market and high capital after the fall.

[6] For example a CRE loan graded as a "5" might be defined as having the same loss rate over its life as a C&I loan graded "5".

has two major disadvantages: loss data is rarely available for long time periods, and past correlations may not reflect future correlations because the asset mix may have changed[7] or the loss driver may change[8]. These limitations mean that top-down estimates of correlation are almost always supplemented by bottom-up estimates based on the estimated volatility of asset values and the Merton model[9]. For CRE the same process can be accomplished with more granularity and subtlety by co-simulating defaults driven by market and economic conditions.

The plainest example is to simulate two CRE deals together across thousands of alternative market environments and calculate the correlation between them by looking at how often they default together. A slightly more subtle example is to consider the simulation of a commercial real estate deal next to a C&I deal where the risk of the C&I deal includes the state of the economy[10]. In this case the driving economic factors are correlated and the correlation between losses on the CRE deal and the C&I deal is estimated from the number of scenarios in which they default jointly and separately.

Given this mechanism for estimating correlations, there are two alternative paths to implementation: one is to put all the assets into the simulation engine and directly generate the loss distribution. The other alternative is to take the correlation coefficients estimated by the simulation engine and put those coefficients into some other calculation engine for economic capital. Both approaches give results that allow the CRE results from cashflow simulation to be used with results for the other asset-classes and the details are discussed in Appendix 2.

***Integrated Loan Pricing***

Pricing is where all the aspects of risk measurement come together and impact the business directly. There are four primary approaches to pricing using quantitative risk measurement:

- Grade grid
- RARORC (Risk Adjusted Return on Regulatory Capital)
- RAROEC (Risk Adjusted Return on Economic Capital)
- CAPM (Capital Asset Pricing Model)

---

[7] For example if the current portfolio has higher LTVs than the historical portfolio.

[8] For example in the previous credit-cycle, losses on two different portfolios may have moved closely together because in that crisis both portfolios were driven by a fall in the economy, whereas the next losses may be driven by a rise in interest rates and if one portfolio is primarily floating and the other is long term fixed, there will be little correlation.

[9] See Chapter 20 of *The Fundamentals of Risk Measurement*, C. Marrison, McGraw Hill, 2001.

[10] Either because the original model was built accounting for the state of the economy in its regressors, or because sensitivity to the economy was "retro-fitted"( *See Linking Stress Tests to the Real Economy*, http://www.riskintegrated.com/DigitalAsset.aspx?id=92)

The advantages and disadvantages of each approach are discussed in Appendix 3. Each approach can be used for commercial real estate, and cashflow simulation can support each, but in our opinion an approach like CAPM is the most direct.

Overall, cashflow simulation improves horizontal integration by providing risk metrics that are consistent with the metrics being provided for the other asset classes and by giving a framework to evaluate the correlation across assets.

## Conclusions

To summarize, Enterprise Risk Management needs to be supported by integrated risk measurement, and that measurement needs to be integrated both vertically across functions and horizontally across asset-classes. It is particularly challenging for commercial real estate because of the nature of the assets and the nature of the risk models that are best suited to assessing those assets. However, the cashflow simulation framework gives the flexibility to have a single model for all of the vertical functions within CRE and output results that are horizontally consistent with the modeling frameworks used for other asset-classes.

Dr. Chris Marrison, Founder & CEO, Tel: +1 (845) 598-1620, Chris.Marrison@RiskIntegrated.com

 

**Appendix 1: The Typical Profusion of Models used for Managing CRE Risks**

Many banks have found it to be difficult to achieve vertical integration for risk measurement of the CRE assets and instead have developed different models for different functions. This situation arises because it is assumed that detailed data cannot be made available at the portfolio level and also because it is assumed that no single model can fulfill all the requirements from deal structuring to economic capital. The typical situation is described below:

- For structuring and underwriting a new loan, the lending team makes a spreadsheet with a cashflow model for net income and debt servicing over time. The spreadsheet is copied and modified from one deal to the next because it does not capture the "unique" features of each new deal. The result is that some of the richest data is stranded in spreadsheets on desktops. Also, managers can never quite be sure of the consistency of the model being used from one deal to the next, and of course there is a strong possibility of spreadsheet errors.
- For grading the asset there is typically some scorecard which requires a new set of information to be typed in to a system. The inputs are usually manipulated a little until the grade that comes out matches expectations. That grade is then added to a credit application form, along with the results of the spreadsheet to form the basis of the approval and pricing decision. A few institutions have the ability to re-run the grading analysis on batches of loans as loan balances, forecasts or model parameters change. But in most cases, this level of integration and automation does not yet exist so the resulting PD, LGD or grade is manually typed into another system. In this case, if re-grading is needed, it is done manually by the lending team.
- For economic capital, the PDs are dumped from the grading system or are retyped into a system that will feed the capital calculator. The capital calculator then takes these PDs and estimates some correlation with the other assets to come up with the overall capital. By this stage all the detail of the cashflow and risk structure over time has been lost so it is not possible to tell if the risk of an individual deal is correlated to interest rates, rental rates, etc., or which years are the most risky. This loss of information means that the economic capital results are indicative and may be directionally correct, but are not really actionable.
- Similarly for portfolio stress-testing, the information on the underlying property, lease and financing structure is typically lost and the risk is only represented by the base-case PD and LGD. With some assumptions those risk numbers can be stressed with changes in the economy and possibly with assumed changes in the CRE market, but again the results are so general that it is not possible to draw detailed conclusions to direct how new assets should be structured.

**Appendix 2: Approaches to Estimating the Correlated Loss Distributions**

There are two primary approaches to using the Merton model with cashflow simulation to estimate correlated loss distributions. One is to assess the losses together within the simulation. The other is to use the simulation to estimate the correlation and export that correlation to a different loss calculator:

- In the first case, in which we generate the loss distribution directly in the simulation, we simply record the sum of the losses from all assets in each scenario as representing the portfolio loss for that scenario. We then generate thousands of scenarios and thousands of scenario losses to generate the loss distribution and thereby read-off the required economic capital. In this approach we also record the loss contribution of each asset. In this approach models can be used for each individual asset, or the assets may be lumped together: for example instead of assessing each individual credit card holder we might use a model for a sub-portfolio of credit cards within a given FICO band.
- In the second case where there is an economic capital calculator which is external to the simulation engine there are broadly two options: we can give the standalone loss distribution for the CRE portfolio, plus the correlation between CRE and each of the other asset-classes, or at the other extreme we can give the PD and LGD for each CRE loan, plus the correlation between that loan and the rest of the portfolio.

**Appendix 3: Risk-based Pricing for CRE Loans**

There are four primary approaches to pricing CRE loans using quantitative risk models:

- Grade grid
- RARORC (Risk Adjusted Return on Regulatory Capital)
- RAROEC (Risk Adjusted Return on Economic Capital)
- CAPM (Capital Asset Pricing Model)

The **grade grid** is the most common approach. In this approach the risk models assign a grade, possibly with some modification from the lending team, and then a "grid" is published by management giving the target price depending on the grade and possibly the size and asset-class. The flexibility in defining the prices in the grid provides management with a pragmatic tool for controlling the lending whilst taking into consideration the prevailing competitive conditions and the bank's appetite for different asset-classes. This type of pricing is fairly easy to integrate across assets so long as the grades for each asset-class reflect the same amount of risk, e.g., a corporate loan grade "5" has the same risk as a CRE loan graded "5".

Some banks adopt **RARORC** whereby they price according to the expected loss and the cost of the *regulatory* capital to be held. This has the advantage that there is no need to calculate inter-asset correlations and that there is an external benchmark as to the amount of capital required, but it has the disadvantage that regulatory capital is only expected to be reasonably accurate at the portfolio level and it may be a very poor estimate of the risk for individual assets. To use RARORC the bank needs to estimate the expected loss over the life of the loan, and it needs to have a calculator embedded in the system reflecting the latest capital rules.

**RAROEC** requires an estimation of the marginal contribution of the deal to the bank's *economic* capital. The first problem in assessing RAROEC is in estimating the correlation with other assets. This is especially difficult for CRE where the correlation for one deal may be very different from another. Even within a deal the correlation may change from year to year: in one year the driver of loss might be changes in rental rates and in another year it might be changes in interest rates. The full estimation of economic capital for a CRE deal therefore requires a multi-year simulation of the asset alongside a model for the shape of the bank's future portfolio. This is quite possible once a form has been decided for the future portfolio. However there are a couple of problems with RAROEC that are exposed by CRE. One is that as other business units change their asset mix over the life of the CRE deal, its correlation to that portfolio will change and so its measured capital will also change. The other problem is that RAROEC does not take into account the fact that when a new asset is added to the portfolio it not only changes the amount of required economic capital, but also changes the cost of the capital because it changes the correlation between the bank's capital and the rest of the market. In extreme cases it can cause a significant distortion in pricing and behaviors, and on

frequent occasions this effect leads to arguments between business groups who thought they had locked in their cost of economic capital, but then found it was changed because of the actions of another group.

In RAROEC the cost of capital for asset i is defined as follows:

$$\text{Capital Cost}_i = hm\overline{\rho_i}UL_i$$

Where

$h$ is the hurdle rate for the bank

$m$ is the capital multiplier

$\overline{\rho_i}$ is the average correlation between loss on asset i and loss on the portfolio

$UL_i$ is the standard deviation of loss for asset i

In theory, $h$ is set for the bank as a whole and depends on the expected return on equity for the bank. The expected return depends on the volatility of credit losses, the market's price for risk, and the correlation between the equity value and the general market.

The RAROEC approach has several well recognized operational issues such as determining how, or whether, the cost of capital the measured profitability should change if the composition of the portfolio changes and therefore causes a change in $\overline{\rho}$. This particular issue leads business units to be in conflict if the actions of one unit affect the correlation and pricing of another unit. For multi-year assets, another challenge is to determine what the future portfolio composition is likely to be, especially if there are alternative strategic plans with different implications for $\overline{\rho}$.

These challenges are well recognized, but a more subtle issue is that the addition of each new asset changes the hurdle rate. To see this, consider the following example.

> Consider a bank with a portfolio which has zero correlation to the general market level. The equity of this bank would be a good diversifying asset for any portfolio manager and would be priced according to the expected cashflows discounted at the risk-free rate, i.e., the hurdle rate would be the risk-free rate.
>
> Now consider a business unit who wants to buy a large block of the S&P 500. It would be uncorrelated with the bank's portfolio and therefore $\overline{\rho_i}$ would be zero and there would be no capital cost. However, by adding this asset, the bank's whole portfolio is now correlated a little with the market and the hurdle rate will rise for all the business units.

In the RAROEC framework this problem might be tackled by a policy rule, or by complicated math to relate $\overline{\rho_i}$ and $h_i$. A more direct approach is to directly focus on the correlation between asset i and the market. This eventually collapses into a "direct-market" approach such as CAPM. Economic capital is an important metric to manage the soundness of the institution as a whole, but for pricing individual transactions, it is easier and more correct to use CAPM.

In **CAPM** the price of the asset still depends on its EL and UL, but $hm\overline{\rho_i}$ is replaced with beta and the market's price for risk. The direct-market approaches have several advantages compared with RAROEC:

- The pricing is consistent with asset managers and traders (i.e., a similar bond and loan to the same company will have the same price)
- There is no interference between one business unit and another
- There are no conditions in which assets are underpriced relative to their full market value
- They are much easier to implement than an $hm\overline{\rho_i}$ approach

CAPM has the advantage of not relying primarily on estimates of the correlation between the bank's assets because the correlation is assessed relative to the broad market. The disadvantage of CAPM is that it does not, on its own, dictate higher prices for assets that concentrate the portfolio and lead to a higher probability of the bank's default. The typical solution to the correlation problem is that the portfolio manager orders limits to the origination of any asset-class which is causing excessive concentration. However, the theoretically pure answer is that a small cost should be added for the asset's contribution to the probability of the bank defaulting, multiplied by the cost of a bank default. This small correction does require the correlation to be known with the rest of the portfolio, but if this term is neglected, it is possible to assess the price of the asset on a standalone basis.