Moody's
ANALYTICS

MODELING METHODOLOGY

Quantifying Risk Appetite in Limit Setting

Authors
Andrew Kaplin
Amnon Levy
Qiang Meng
Libor Pospisil

Contact Us

Americas
+1.212.553.1653
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia-Pacific (Excluding Japan)
+85 2 3551 3077
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Abstract

An institution's Risk Appetite Statement (RAS) specifies the aggregate level and types of risks the firm is willing to take (or avoid) in order to achieve its business goals. The translation of the RAS into limits allows an organization to achieve its strategic objectives and business plan while adhering to its risk capacity.

In this paper, we explore leveraging an organization's economic capital framework to quantify the RAS via risk- and macro scenario-based limits. Risk-based limits create a level playing field, using risk-based metrics to align with the organization's risk appetite. Macro scenario-based limits control exposures to adverse macroeconomic scenarios and are frequently viewed as more intuitive and more tangible than risk-based limits.

We also describe a number of approaches for setting risk- and macro scenario-based limits: top-of-the-house (TOTH) risk limits, standalone sub-portfolio (SASP) risk limits, portfolio referent sub-portfolio (PRSP) risk limits, as well as Stressed Expect Loss (SEL) Limits, and macro risk-based limits. To demonstrate the approaches for limit setting, we use a sample portfolio that consists of C&I, CRE, and retail exposures. We illustrate the relationship between varying sub-portfolio characteristics and the various approaches to setting limits.

Table of Contents

1. Introduction

An organization's risk appetite statement (RAS) is the aggregate risk level and types of risk it is willing to accept (or avoid) in order to achieve its business objectives. The RAS includes qualitative statements as well as quantitative measures, expressed relative to earnings, capital, risk measures, liquidity, and other relevant measures.[1] The RAS translation feeds into quantitative risk-based limits and allows an organization to achieve its strategic objectives and business plan, while adhering to its risk capacity. Risk-based limits allocate the financial institution's aggregate risk appetite statement (e.g. measure of loss or negative events) to business lines, legal entities as relevant, specific risk categories, concentrations, and other levels.

There is no question that articulating stakeholder preference and appetite for risk is challenging. The challenges are intimately linked with the implications RAS has on metrics related to its monitoring and enforcement. Risk limits are a key mechanism of enforcement and part of the overall risk appetite framework.

The limit setting process is multifaceted, with risk appetite playing a central role. And while the quantitative measures in the risk appetite statement provide limit setting guidance, the qualitative portions must be incorporated as well.The challenge is in quantifying an RAS that may not have obvious quantitative translations. A specific statement such as "maintain an Aa rating or better" (perhaps with a certain likelihood), was part of the Royal Bank of Canada's (RBC) RAS, profiled in the 2011 Institute of International Finance (IIF) study on risk appetite. [2] This statement can be linked to quantitative performance measures by quantifying the requirements needed to maintain the Aa rating. Some judgment is required, however, given agency ratings are not entirely based on quantitative metrics.

An example of a more open-ended statement, "maintain low exposure to 'stress events,'" is another item in RBC's RAS. The process of quantifying this statement requires defining "stress events" as well as "low exposure." In the U.S., natural interpretations of stressed events are the CCAR Adverse or Severely Adverse scenarios. A low exposure can be related to the sensitivity of expected loss to the associated macro-economic variables.

In some cases, one can relate multiple elements in the RAS to allow for a more holistic quantitative view of risk appetite. For example, the two statements above can describe RBC's aversion to losing their Aa rating under the Adverse or Severely Adverse scenarios. This can be made more precise by interpreting the risk appetite to target a portfolio composition and capital buffer so that, say, the likelihood of losses exceeding those that would result in a downgrade are no higher than one-in-25 years, conditional on the Adverse scenario.

Quantifying risk appetite should rely on sound economic principles, along with strategic overlay. It is natural to leverage an economic capital framework in this context. For example, an economic capital framework can help us understand the likelihood of a portfolio incurring losses that drive the institution to a downgrade. An EC framework can also allow further drill-down by providing estimates for the likelihood of, for example, losses related to a sub-portfolio resulting in the downgrade. Alternatively, the framework can be used to help understand portfolio loss dynamics, conditional on macroeconomic scenarios, to define limits that will help avoid losses associated with stress events.

A subtle point worth mentioning is the role that limits play and how their role relates to other tools used to align incentives across the organization. Limits are generally broad-brush mechanisms that ensure risk exposure conforms to the level and likelihood of loss. Other tools, such as RORAC and EVA-style measures used in deal pricing and incentive compensation, allow an organization to optimize their risk-return profile. The various tools complement each other, as the complex nature of the organization's desire to align various incentives and achieve its goal of maximizing return, while adhering to a desired risk profile, is too complex to achieve using a single mechanism. We analyze a number of quantitative approaches in translating an RAS to limits, and we use case studies to better understand the ensuing dynamics.

The remainder of this paper is organized as follows:

- » Section 2 considers risk-based limits.

- » Section 3 considers macro scenario-based limits.

- » Section 4 considers practical issues when applying an economic capital framework when setting limits.

- » Section 5 concludes.

---

[1] Financial Stability Board, "Principles for an Effective Risk Appetite Framework." November 18, 2013.

[2] Institute of International Finance, "Implementing Robust Risk Appetite Frameworks to Strengthen Financial Institutions." June 2011.
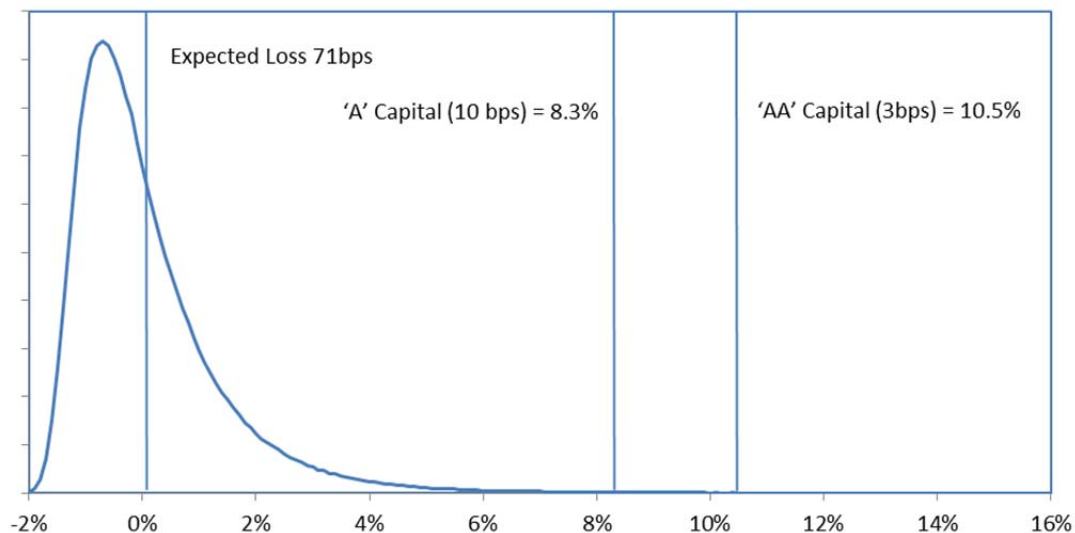
2. Risk-Based Limits

As defined by the IIF, risk limits represent quantitative measures based on forward-looking assumptions that allocate the financial institution's aggregate risk appetite statement (e.g., measure of loss or negative events) to business lines, legal entities as relevant, specific risk categories, concentrations, and, as appropriate, other levels. This approach to limit setting is a departure from traditional, notional-based limits. Risk-based limits create a level playing field, utlizing risk-based metrics that align with the organization's risk appetite. In this section, we use a quantification case study to demonstrate how to design risk-based limits that address a statement such as, maintain an Aa rating (referenced in the Introduction). There are three sub-sections: 2.1 describes how top-of-the house (TOTH) risk limits can be defined, 2.2 describes how stand-alone sub-portfolio (SASP) risk limits can be defined, and 2.3 describes how portfolio-referent sub-portfolio (PRSP) risk limits can be defined.

2.1 Top-of-the House Risk Limits

Maintaining a high credit rating is not an easy task. There were 32 banks rated Aa or higher in 2014 (2004 – 2006 there were 58 such banks). When quantifying downgrade risk, one must understand the likelihood of various loss levels under the portfolio's existing composition and the required buffer that allows retention of an entity's Aa rating. At the top-of-the house, a rating transition matrix can provide a rough sense of the likelihood of a downgrade. An economic capital framework can refine this process by more accurately mapping out the portfolio loss distribution, specifying the existing buffer, and quantifying the likelihood of losses exceeding the threshold associated with a downgrade. In other words, the currency loss would be the difference between capital thresholds associated with the current capital buffer and an A rating, as in Figure 1. The figure depicts the loss distribution for the U.S. corporate sub-portfolio of the IACPM portfolio. Table 2 describes details of the portfolio. Under the assumption[3] that an Aa rating is associated with a 3bps – 10bps range of target probabilities, the portfolio maintains its Aa rating as long as the existing capital falls into that range. The notional TOTH limit can be calculated using the ratio that reflects how much growth the current capital can support without additional contributions, assuming that the portfolio composition scales up proportionally. For example, if the current capital is at 10.5%, then the portfolio notional can grow by (10.5% – 8.3%) / 8.3% = 26.5% and still maintain that rating without injecting additional capital. This can be expanded by assuming a certain rate of growth for capital that leads to higher limits, keeping everything else constant, as well as using a more conservative lower limit (higher than 8.3%), which results in lower notional limits.

Two aspects to this approach are worth recognizing. First, the portfolio composition does not change as it grows. And, the buffer is driven entirely by loss/gain on assets (i.e., liabilities are not changing). Section 4 further discusses these assumptions and also considers practical applications.

Figure 1   Loss Distribution and Capital Thresholds Associated with Aa and A Ratings

Expected Loss 71bps

'A' Capital (10 bps) = 8.3%

'AA' Capital (3bps) = 10.5%

---

[3] We assume that the Aa rating will be maintained up to the 10bp threshold. Alternative thresholds are excluded for exposition.

2.2 Sub-portfolio Risk Limits

There are a number of ways to translate RAS into risk-based sub-portfolio limits. In this sub-section, we explore two approaches that define sub-portfolio limits: stand-alone sub-portfolio and portfolio-referent sub-portfolio limits. We conduct the analysis within the context of the case study above, where there is an appetite to maintain the Aa rating to construct industry-based risk limits.

Proceeding with the case study, we analyze the portfolio in question for 14 industry segments. Figure 3 provides the notional amount as well as the Risk Contribution. Notice the difference in the rank-ordering of the notional and the risk as it relates to limit setting.

A practical consideration in setting limits is the wide cross-sectional variation in commitment amount that can proxy for the organization's strategic investment opportunity set. When setting limits, the idea of constructing a level playing field in the notional or risk space makes sense from a theoretical perspective, but it does not consider the practicalities of the organization's investment strategies. The annual operating plan (AOP) must be overlaid in a way that allows the limits to provide realistic guidelines that can relate back the RAS, which provides guidelines on the risk the financial institution is willing to accept (or avoid), in order to achieve its business objectives. Section 4 discusses these considerations in greater detail.

SASP RISK LIMITS

We can construct stand-alone industry risk limits using each industry's loss distribution and mapping the likelihood that loss exceeds a threshold associated with a downgrade for the overall organization. Conceptually, the limit is not portfolio-referent, in that, it does not account for correlation across industries, and the likely losses associated with correlated industries, when the industry in question deteriorates. This property has drawbacks and benefits. By not accounting for portfolio concentration in correlated industries, the approach does not recognize that more highly-correlated industries are riskier from a portfolio-referent perspective. This said, that very characteristic can make it easier to manage, as various stakeholders can manage proximity to their relevant limits without concern over how investments are managed elsewhere in the organization. This point is discussed further in Section 2.3, which describes portfolio-referent sub-portfolio limits.

Table 1 provides some basic characteristics of the portfolio, segmented based on industry. Included is the commitment, stand-alone Unexpected Loss (UL), risk contribution, and other relevant parameters of each sub-portfolio.

Table 1

Characteristics: U.S. Corporates Sub-portfolio

INDUSTRY	NOTIONAL	NOTIONAL WEIGHTED PD	NOTIONAL WEIGHTED LGD	NOTIONAL WEIGHTED RSQ	NOTIONAL WEIGHTED RC	STANDALONE UL
Agriculture	103,000,000	0.61%	38.66%	17.45%	0.71%	1,480,307
Banks and S&Ls	1,743,000,000	0.72%	45.80%	46.27%	1.72%	49,526,730
Business Products	176,000,000	3.49%	38.56%	11.26%	1.33%	5,612,098
Consumer Products	5,922,000,000	1.12%	39.27%	34.86%	0.92%	67,191,014
Equipment	1,147,000,000	3.36%	43.56%	32.23%	4.13%	81,834,392
Financials (Other)	2,525,000,000	0.33%	40.39%	51.54%	0.50%	19,801,516
High Tech	1,703,000,000	0.11%	44.38%	45.94%	0.38%	13,298,554
Materials/Extraction	9,057,000,000	1.81%	40.53%	33.02%	1.62%	169,944,117
Medical	212,000,000	1.04%	38.23%	11.45%	0.59%	2,610,645
Real Estate	2,166,000,000	0.82%	37.98%	34.64%	1.05%	53,584,208
Services	1,455,000,000	0.80%	41.59%	20.34%	0.62%	16,106,039
Tel/Cable/Printing/Publishing	6,604,000,000	1.74%	37.93%	35.26%	1.24%	114,211,073
Transportation	2,207,000,000	1.14%	40.43%	31.06%	1.21%	47,408,142
Utilities	1,212,000,000	3.07%	41.11%	38.50%	2.30%	36,953,750

SASP risk limits can be quantified using a risk space (e.g., UL or a tail-based) as the level playing field, whereby risk at the industry segment is considered. An approach to defining sector limits would be to require that even a catastrophic case in that sector would not downgrade the overall portfolio. For this process, we calculate capital for each sector at a very small target probability and then calculate what sub-portfolio increase would lead to a currency-based loss leading to a downgrade

In the TOTH risk limit calculation in Section 2.1, the threshold for a currency-based loss is 2.2% (difference between 10.5% and 8.3%) of portfolio notional. Suppose Banks and S&L sector capital is 69.6%, and we know that it represents approximately 4.8% of the overall portfolio (by commitment). Then we can increase the holdings in Banks and S&Ls (assuming no injection of new capital and current capital is at 10.5%) by approximately 2.2%/69.6% or 3.2%[4] of the overall portfolio. So the limit is (3.2%+4.8%)*Current Overall Portfolio Commitment.

2.3 Portfolio-referent Sub-portfolio Risk Limit

Portfolio-referent sub-portfolio risk limits consider the risk of a segment in the context of the overall portfolio. The limits recognize segments likely to incur substantial loss when the rest of the portfolio performs poorly as more risky. In this sense, segments that contribute to a larger portion of loss when the overall portfolio deteriorates to the point of a downgrade would face a tighter notional limit.

The choice of risk allocation measure, tail-based or Risk Contribution (RC)-based, is an important one, and a choice specific to an institution's risk preferences. For the purpose of this discussion, we focus on RC-based risk allocation, which measures an instrument or portfolio segment's contribution to the portfolio UL. RC is formally defined as the change in portfolio standard deviation (in currency unit) resulting from one additional currency unit of this particular instrument or portfolio segment. It can be shown that RC equals UL of the instrument or segment multiplied by the correlation between the value of the instrument or

segment and value of the entire portfolio. Mathematically represented as follows: $RC_i \equiv \dfrac{\partial UL_P}{\partial w_i} = \rho_{i,P} \cdot UL_i$

Visually, Figure 2 presents a traditional decomposition of an instrument or segment's UL in the context of an overall portfolio. The market. UL represents the stand-alone standard deviation, and it can be decomposed as a diversified portion, a diversifiable (but not diversified) portion, and a systematic or non-diversifiable portion. The Risk Contribution is the sum of the systematic portion and the diversifiable portion.

Figure 2   Decomposition of UL
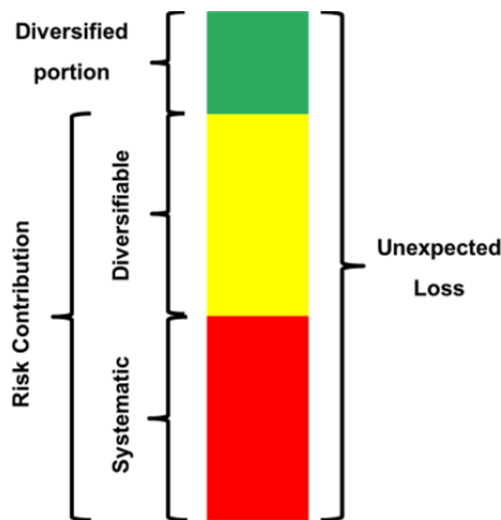
From a diversification perspective, the worst situation is that the segment is perfectly correlated with the entire portfolio. This issue occurs if the entire portfolio consists of only one segment, or if all segments are perfectly correlated. For such portfolios, none of the segment's stand-alone risk can be diversified away, and the segment RC is equal to segment UL. On the other hand, in the ideal case where all diversifiable portions in segment UL are diversified away, RC will consist of only the systematic portion of UL. In the real world, it is usually not possible to diversify away all diversifiable portions in segment UL, and the lower bound for RC

---

[4] Capital rate associated with a catastrophic case — a 1bps event in this example — for the Banks and S&L sector.

is typically higher than the systematic portion in segment UL. This lower bound is portfolio-specific and we refer to it as portfolio-referent systematic portion in this paper.

Continuing with the case of downgrade avoidance, portfolio-referent sub-portfolio risk limits can be constructed by measuring the proportion of risk associated with each segment when a downgrade occurs. As a starting point, Figure 3 presents Commitment and normalized Risk Contribution for each industry segment as a proportion of the overall portfolio Commitment and Unexpected Loss, respectively. While many industries have a similar relationship between the proportion of risk and commitment, there are a number of striking exceptions, including Financials and High Tech, whose portfolio-referent risk statistics (measured as proportion of the portfolio's UL) are lower than half the sector's commitment as a proportion of portfolio's commitment. On the flip side, a proportion of the portfolio's UL is more than three times higher than the sector's commitment, as a proportion of the portfolio's commitment to the Equipment segment.

Figure 3   Commitment and Risk Contribution (as a proportion of the portfolio Commitment and Unexpected Loss, respectively)
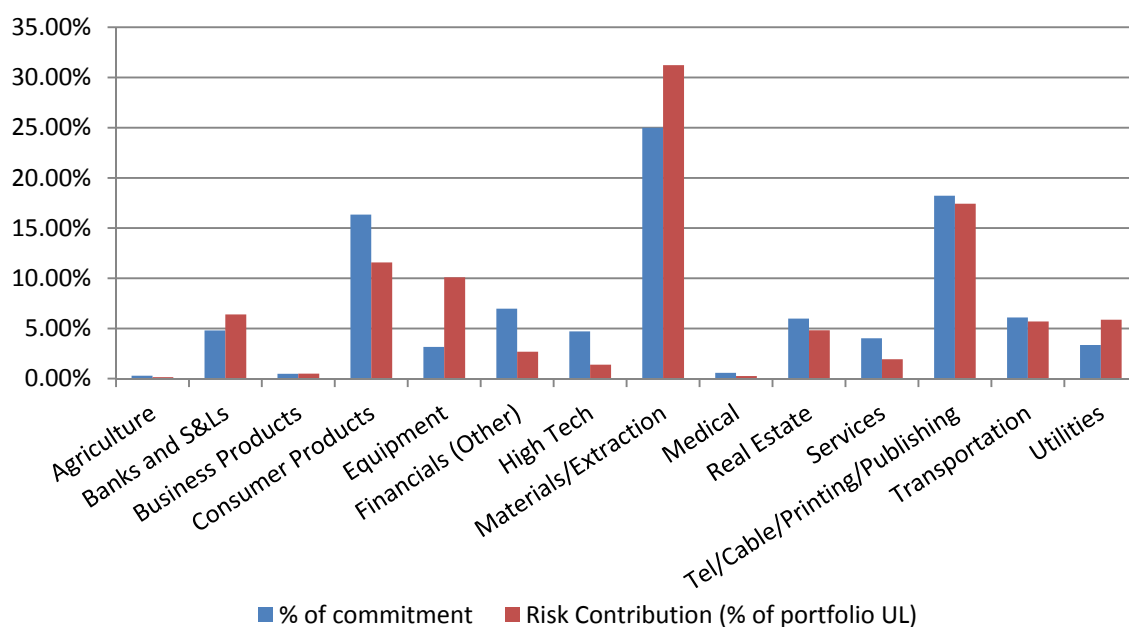
When defining risk based limits, it is necessary to understand the relationship between the risk and amount invested in a segment. Figure 4 illustrates how the RC of a segment changes when the weight of that segment increases.and presents the relationships for the 14 sectors. Take Equipment as an example. The current RC is 4.13% and comprises 10% of the portfolio UL. As discussed earlier, if the Equipment were 100% of the portfolio, the RC would represent the UL of Equipment. On the other hand, the normalized RC when Equipment is set at 0% represents a measure of portfolio-referent systematic risk related to Equipment.

We can now apply the analysis more broadly and define Risk Contribution-to-Portfolio (RCP) as the segment RC measured as a proportion of the entire portfolio UL: the share of portfolio UL attributed to a particular segment. The RCP is a function of segment weight and increases from 0, when the weight is 0, to 100%, when the weight is 100%. Now we are ready to translate the RC-based limits to notional-based limits.

Suppose no more than 35% of portfolio risk should be allocated to any segment in order for a bank to maintain an Aa rating.[5] For RC-based capital allocation, this mandate is equivalent to requiring that the RCP should not exceed 35% for any segment. Since RCP is an increasing function of segment weight and takes a value between 0 and 100%, there is a unique value for segment weight, such that, RCP equals 35%. Figure 4 illustrates the translation from RC-based limits to notional-based limits.

[5] In practice, sub-portfolio limits should not be determined purely based on the requirement of maintaining a high credit rating. For instance, an organization can maintain an Aa rating by investing in only high credit quality assets. But this strategy does not allow the organization to recognize high fees. In this paper, we assume that the upper bound for RCP not only satisfies the requirement of a high rating but also allow business development in areas where the organization has strategic advantage.

Figure 4   Relationship Between RC and the Amount Invested in a Segment
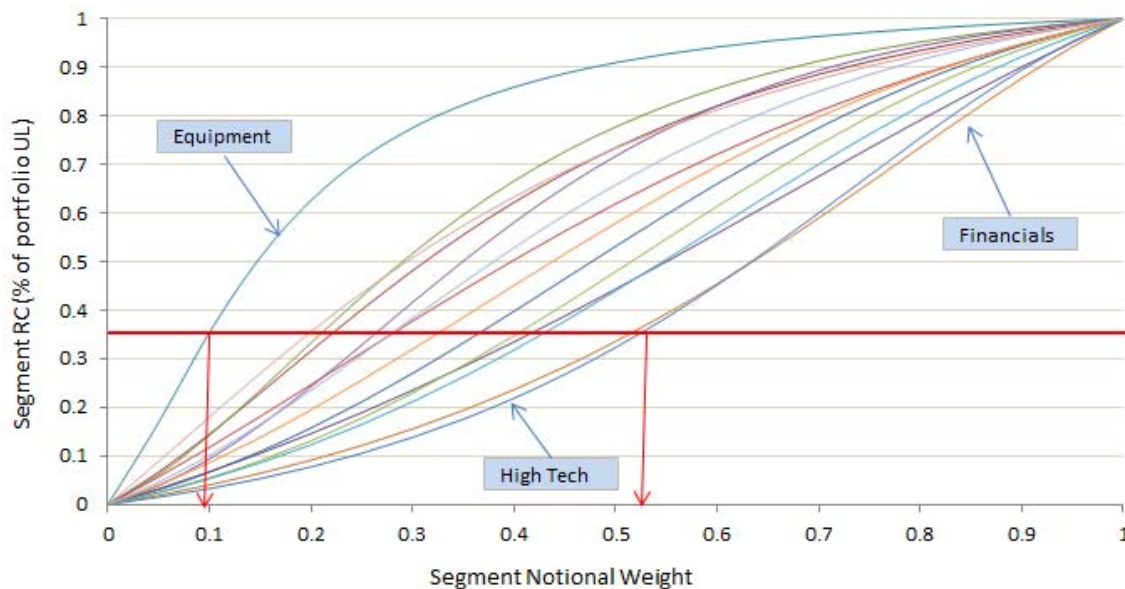
It is interesting to notice that the notional-based limits for the segments take a wide range of values, from around 9.9% (for Equipment sector) to 52.2% (for High Tech segment), even though the RC-based limits are identical for all segments. This trait is due to the different characteristics of each segment. Generally speaking, for a fixed RC-based limit, the notional-based limit is higher if the segment has lower stand-alone credit risk (as measured by PDs) and lower concentration risk (reflected by low correlation with other segments in the portfolio). Figures 5 and 6 demonstrate the notional-based limits as a function of sector's average PD and R-squared values, respectively. It is evident that the notional-based limit is lower for segments with higher PD values (higher credit risk). This finding is intuitive and corresponds to the traditional rating-based limits. But the rating-based limits miss the other, equally important side of the story, namely the concentration risk. As Figure 6 shows, the RC-based approach to limits allows us to capture these dynamics naturally.

Generally, the more concentration risk added by a sector, the more limiting the result will be. Thus, RC-based limits take into account both credit and concentration risks and actually allow risk managers to reflect the effects of both in a single limits measure. For example, just looking at the credit risk, it might seem counterintuitive that the limit for the Business Products sector is twice as large as the limit for the Equipment sector (20.9% vs. 9.9%) even though Business Products is riskier than Equipment (average PD values are 3.49% and 3.36%, respectively). But the concentration risk graph explains why this is the case. We see the Equipment sector sub-portfolio introduces considerably more concentration when compared to Business Products (average R-squareds are 32.23% and 11.26%, respectively). Similar dynamics can be observed when the limits are higher, even if the concentration risk brought by the sector's sub-portfolio is higher (for example, compare High Tech and Equipment limits, average PDs and R-squareds).

Figure 5   Notional Limit as a Function of Sector's Average PD
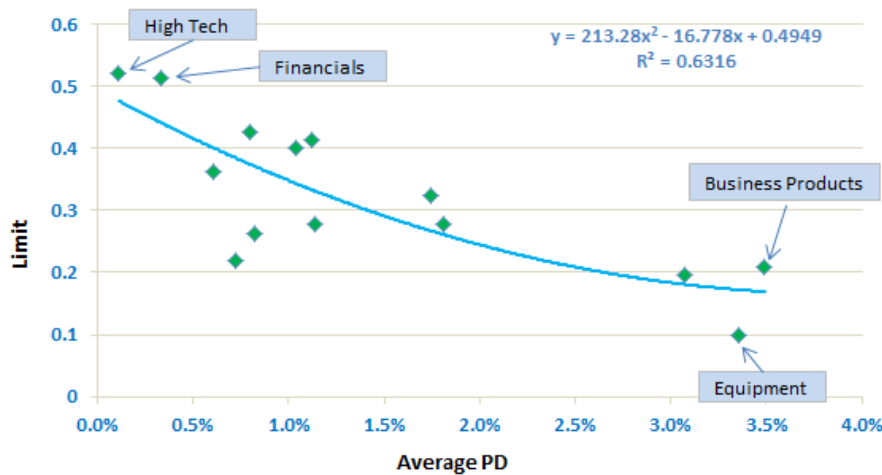
Figure 6   Notional Limit as a Function of Sector's Average RSQ[6]
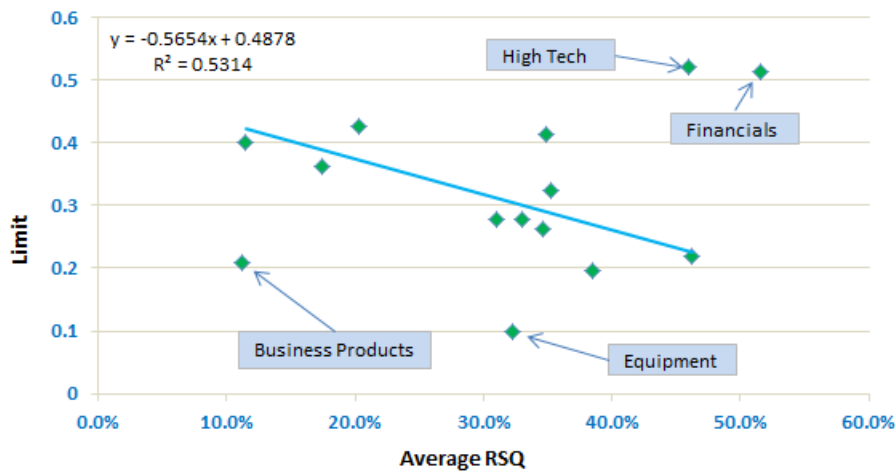
For strategic planning, it is also helpful to compare the notional-based limits with the weights of each segment in the current portfolio, illustrated in Figure 7. All segments in the graph fall below the 45-degree line, meaning that the current weights are below the notional-based limits. For the segments close to the 45-degree line, such as Materials/Extraction, the current weight is just slightly below the notional-based limit. The implication is that even a relatively small increase in the investment in this industry could jeopardize the banks Aa rating. On the contrary, the High Tech industry lies far below the 45-degree line, implying that the bank can significantly increase the weight of the High Tech industry in the entire portfolio and still maintain the Aa rating.

---

[6] The trendline and the corresponding regression equation are calculated after omitting the four outliers.

Figure 7   Limits and Current Weights
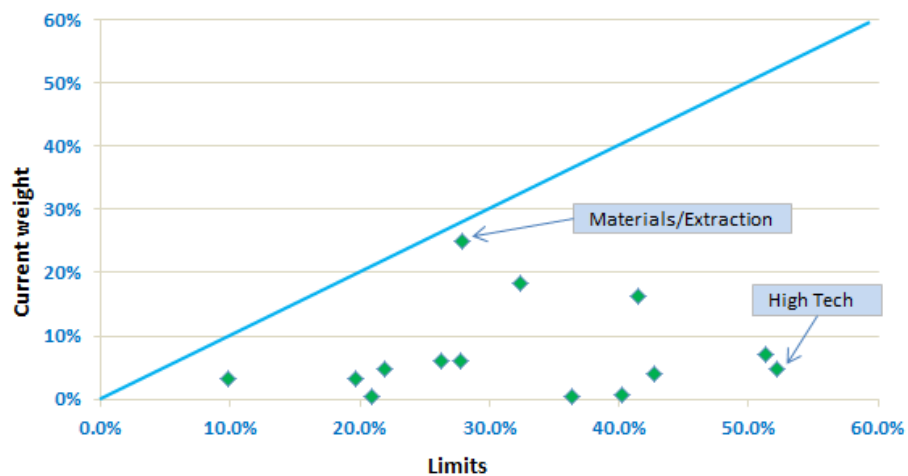
3. Macro Scenario (MS) Limits

MS limits control exposure to adverse macroeconomic scenarios. Macro scenario limits have some appealing qualities. In the U.S. and other jurisdictions, regulators require organizations to maintain capital levels that can withstand adverse economic environments. Defining limits that ensure adherence to these requirements has appeal from a regulatory compliance perspective. Separately, groups outside of risk functions can find risk-based limits abstract and difficult to intuit at times. Meanwhile, macro-based analysis is frequently viewed as more tangible and thus easier to relate with.

There are at least two broad classes of macro-based limits: Stressed Expect Loss (SEL) Limits and Macro-Risk-based limits. Similar to notional or EL-based limits, SEL grows linearly with notional and does not capture concentration or diversification effects. It speaks to traditional loss stress testing and is typically viewed as transparent. SEL is additive, and it does not require as complex a modeling infrastructure. Meanwhile, Macro-Risk-based limits consider portfolio effects. A limit ensuring that an organization maintains an IG rating under an adverse economic scenario with a certain probability is an example of a Macro-risk-based limit. Such an analysis requires expansion of an economic capital framework linking portfolio loss with macroeconomic scenarios.[7]

To demonstrate the risk appetite quantification process in this context, we create a sample portfolio containing exposures to large U.S. and global corporate entities, U.S. retail and U.S. CRE. The sub-portfolios have characteristics typical of those asset classes (e.g., corporate entities have lower PDs and higher RSQs than retail). The portfolio is designed to illustrate variation in sensitivity to macroeconomic variables across asset classes.

Table 2 provides summary statistics for each of the sub-portfolios and the combined portfolio.

---

[7] Libor Pospisil, Andrew Kaplin, Amnon Levy, and Nihil Patel, "Applications of GCorr™ Macro: Risk Integration, Stress Testing, and Reverse Stress Testing." Moody's Analytics, September 2014.

Table 2

Summary Statistics for Sample Sub-portfolios and Combined Portfolios

	U.S. CORPORATES	GLOBAL CORPORATES	U.S. CRE	U.S. RETAIL	AGGREGATE PORTFOLIO
Total Commitment	36 billion USD	26 billion USD	13 billion USD	14 billion USD	89 billion USD
# exposures	2,266	3,734	130	498 (homogenous pools)	6,130 individual exposures plus 498 homogeneous pools
# counterparties	1,133	1,867	130	498(homogeneous pools)	3,130 counterparties plus 498 homogeneous pools
Weighted Avg. PD	1.40%	1.58%	1.44%	2.41%	1.62%
Weighted Avg. LGD	40.25%	40.73%	25.00%	30.00%	36.55%
Weighted Avg. RSQ	35.61%	39.19%	35.05%	9.38%	32.46%
Locations and Types of Exposures	Diversified across U.S. industries (Agriculture, Banks and S&Ls, Business Products, High Tech, …).	Exposures from Japan, Europe, Australia. Diversified across industries, similarly to the U.S. portfolio.	U.S. Commercial	Diversified across U.S. industries (Agriculture, Banks and S&Ls, Business Products, High Tech, …).	Exposures from Japan, Europe, Australia. Diversified across industries, similarly to the U.S. portfolio.
Concentration	The 10 largest exposures account for 16% of the total commitment.	The 10 largest exposures account for 8% of the total commitment.	The 10 largest exposures account for 66% of the total commitment.	The 10 largest pools account for 23% of the total commitment	The 10 largest pools account for 10% of the total commitment
EL	0.71%	0.86%	0.49%	0.64%	0.71%
Capital wrt EL, 10bps	8.35%	7.84%	13.99%	5.44%	6.97%

3.1 Top-of-the House Macro Scenario Limits

MACRO SEL LIMITS

As referenced above, SEL-based limits are appealing for a number of reasons. In jurisdictions that require financial institutions to adhere to regulatory stress tests, SEL-based limits are a natural mechanism to help ensure compliance. In the context of the sample portfolio, one can define a limit requiring SEL to fall below the regulatory thresholds, perhaps with an additional buffer. The process of setting macro scenario limits requires estimating SELs, which can be involved and is sub-portfolio specific. For the purposes of this document, we provide only a rough set of the steps needed, but for those interested, please see Pospisil, et. al., "Using GCorr™ Macro for Multi-Period Stress Testing of Credit Portfolios."

To begin, each sub-portfolio differs in its sensitivity to macro-economic variables.[8] Thus, modeling SEL frequently entails relating sub-portfolios with a different set of macroeconomic variables. For example, HPI for retail portfolios and the DJ Total Stock Market Index for U.S. corporate exposures. Notice, one can mechanically aggregate the SEL across the sub-portfolios to arrive at the TOTH SEL, despite conditioning on a different set of variables for each sub-portfolio. This will be a more complex issue when considering macro scenario risk, where conditioning is conducted at the portfolio level, and we require a consistent set of macro variables across sub-portfolios.

In analyzing the sample portfolio, Figure 10 below provides SEL for each sub-portfolio, assuming a scenario based on Unemplyment Rate, BBB Corporate Spread, Dow Jones Total Stock Market Index, and the VIX, and their changes from 2014Q3 – 2015Q3 under the CCAR 2015 Severely Adverse scenario.[9] One can imagine defining a limit based on SEL limit that recognizes the organization's capital structure, along with the organization's appetite to the sensitivity of loss to macroeconomic variables. For

---

[8] A variable selection procedure leads to the following sets of macroeconomic variables that best describe each sub-portfolio from both a statistical perspective (a parsimonious model containing only significant variables, high adjusted R-squared value when regressing the portfolio losses on the macroeconomic variables), as well as an economic perspective:

> » U.S. Corporates: Unemployment Rate, BBB Corporate Spread, Dow Jones Total Stock Market Index, and VIX; with adjusted R-squared of 49%.
> » Global Corporates: Japanese Equity Market Index, UK Equity Market Index, UK Unemployment Rate, Eurozone BBB Corporate Spread; with adjusted R-squared of 44%.
> » U.S. CRE: Real GDP Growth, Dow Jones Total Stock Market Index, CRE Index; with adjusted R-squared of 39%.
> » U.S. Retail: Unemployment Rate, Dow Jones Total Stock Market Index, House Price Index (HPI); with adjusted R-squared of 60%.

The variable selection employed here is described in Pospisil, et al., "Applications of GCorr™ Macro: Risk Integration, Stress Testing, and Reverse Stress Testing."

[9] Note, these are the four macroeconomic variables constituting a model best describing the U.S. Corporates portfolio. Thus, when we use these variables for stress testing another sub-portfolio, the explanatory power of this model is lower than the explanatory power of the best model selected for that sub-portfolio. As an example, these variables lead to an adjusted R-squared of 49% for the U.S. Retail portfolio, while the best model, which contains HPI, explains 60% of variation in losses.

example, CCAR requires at least a 5% capital buffer remaining under the Severely Adverse scenario. In the interest of exposition, we assume that: (1) portfolio composition does not change under stress, (2) the capital buffer before stress is 11.5%, (3) EL is 0.7% and SEL is 5.0%, and (4) the buffer absorbs the SEL. Under these assumptions, the amount of capital buffer consumed to absorb SEL is 4.3% (i.e., 5.0%-0.7%) of notional. The remaining capital buffer is 11.5%-4.3% = 7.2%. The TOTH SEL limit can now be set as 144% of current notional amount, since the available capital buffer 7.2% is 144% of the CCAR required capital buffer 5%. If we further assume that the capital buffer under the Severely Adverse scenario should be higher than the CCAR required capital rate to account for variation in modeling across the institution and the Fed, the TOTH SEL limit will be lower than 144%. For instance, if the desired capital rate under the Severely Adverse scenario is 7%, the limit should be set as 103% (i.e., 7.2%/7%) of current notional amount.

MACRO SCENARIO RISK LIMITS

The preceding section focused on SEL as a portfolio statistic under a scenario. Despite its value for understanding portfolio behavior under a scenario, SEL itself does not measure risk at the portfolio level, in the sense that it does not depend on portfolio loss distribution, nor does it account for interactions among instruments in the portfolio (instrument level SELs are agnostic to the composition of the portfolio).

We therefore consider other analyses that take into account the entire distribution of a portfolio under a macroeconomic scenario. Figure 8 shows the unconditional and the conditional distribution of the U.S. Corporate portfolio from Table 2. We base the scenario utilized for that conditional simulation on shocks to Unemplyment Rate, BBB Corporate Spread, Dow Jones Total Stock Market Index, and the VIX and their changes from 2014Q3 to – 2015Q3 under CCAR 2015 Severely Adverse scenario.

Figure 8   Conditional and Unconditional Dstributions of U.S. Corporates Sub-portfolio
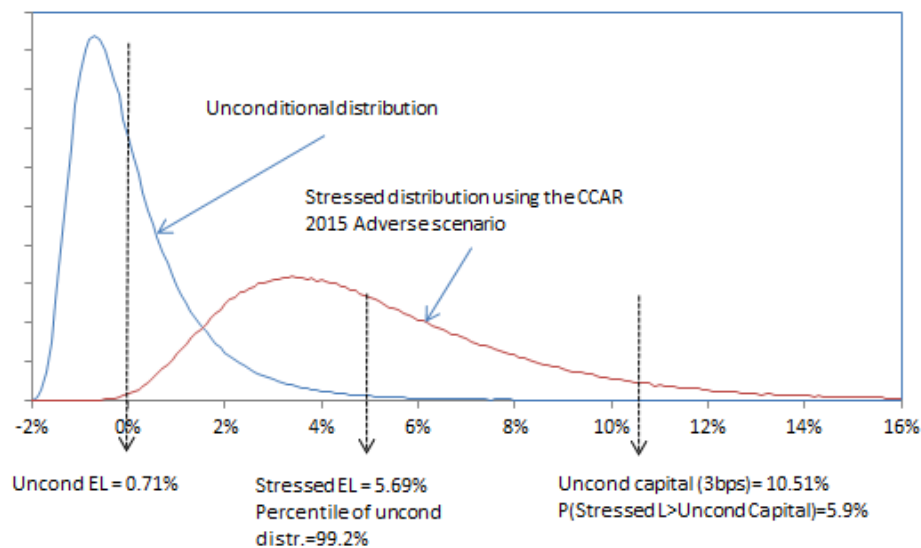
As Figure 8 shows, the conditional loss distribution centers around higher losses, because it assumes the Adverse scenario. But another important point: there is still a dispersion in the conditional distribution, due to the fact that the risk in the portfolio is unrelated to the macroeconomic environment. Such a risk can arise from name concentration (idiosyncratis risks of dominating borrowers affect loss distribution even after conditioning on macroeconomic variables) and country- or industry-specific risks uncorrelated with the macroeconomic variables considered in the scenario (scenario macroeconomic variables are typically broad economic indicators that cannot capture shocks to individual industries not neccesarily associated with economy-wide downturns). The U.S. Corporate portfolio is sufficiently diversified across names and contains only U.S. names. Therefore, the dispersion in its conditional distribution comes primarily from industry-specific effects.

Questions of interest related to Figure 8 can be of two types:

» What are the quantiles of the conditional loss distributions?

» What is the probability that the loss will exceed a certain threshold under a macroeconomic scenario?

Both of these types of metrics can be used to define limits: if a portfolio's metric exceeds the limit, an action is triggered.

Let us give an example of the second metric — the threshold could be set to the level consistent with an institution retaining a certain rating (such as investment grade rating). The metric of interest is the probability that the portfolio loss will be high enough to breach this threshold under a macroeconomic scenario. If the probability exceeds a limit, then an action must be taken — such as moving exposures to industries less sensitive to the scenario. We explore this example in more detail in Section 3.2, within the context of multiple portfolios, including a discussion of how to determine a loss threshold associated with a given rating.

Figure 9 plots two series of rating-implied-EDF measures (for ratings Aa3 and Baa 3). As Figure 9 shows, the meaning of rating, in terms of default probability, substantially varies over time. For example, an entity rated Aa3 had a one-year default probability of 2bps in 2008Q1, but this probability reached 9bps in 2009Q1, more than a four-fold increase.

Figure 9   EDF Value Change Over Time
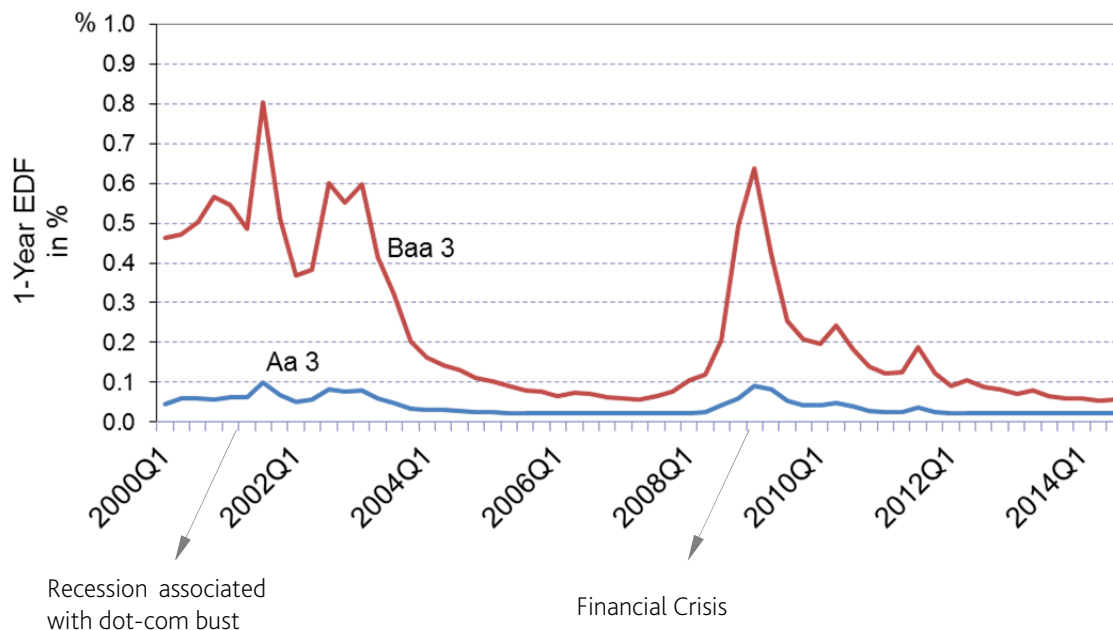
3.2 Sub-portfolio Macro Scenario Limits

MACRO SEL LIMITS

Sub-portfolio SEL limits are similar to the sub-portfolio risk-based limits, except that SEL is being used as the level playing field. Figure 10 illustrates one approach of setting sub-portfolio SEL limits. In the example, the SEL threshold for each sub-portfolio is set at 2.5 billion, which is very different when translated to notional, when looking across the various portfolios. The relationship between SEL and notional is linear, and the slope represents the SEL as a percentage of the notional amount.

It is worth highlighting some of the drives that result in different limits. Table 2 provides summary statistics that demonstrates, while CRE has the lowest EL, global corporates has the highest EL. This said, CRE has the tightest limit when translated to notional, and global corporates has the highest. This finding is driven by the senativity of loss to stressed macro factors. In this case, global corporates are much less sensative to the scenarios than CRE.

Figure 10   Stressed EL-based Limits
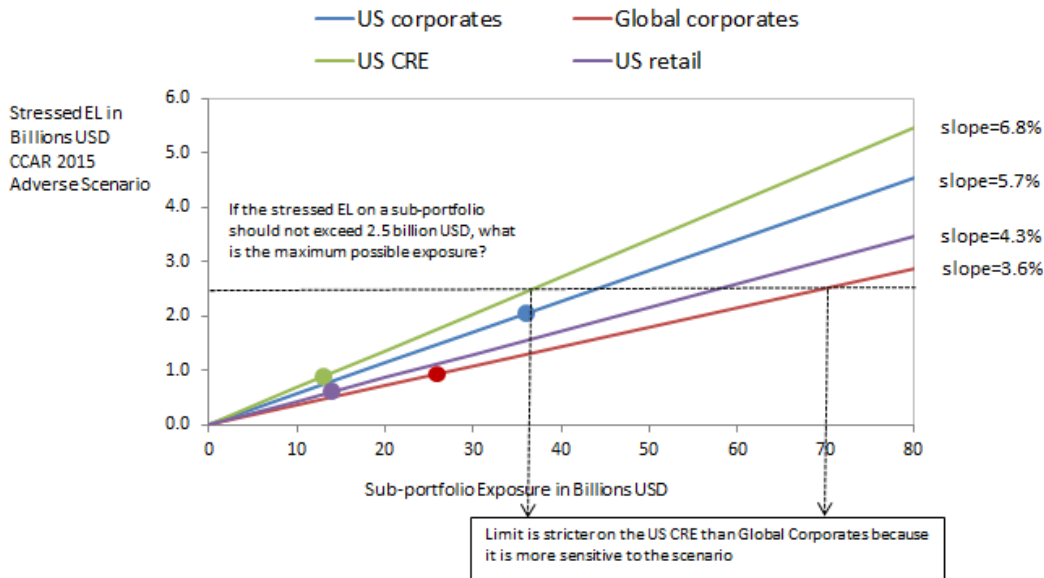
MACRO SCENARIO RISK LIMITS

Similar to sub-portfolio risk-based limits, one can define limits based on the extent to which a sub-portfolio might contribute to loses resulting in an organization falling below investment grade. Following the discussion above on risk-based limits, macro-scenario risk limits can be stand-alone or portfolio-referent. For exposition, we focus the discussion on the stand-alone-macro-scenario risk limits . Figure 11 depicts how each of the four sub-portfolios behaves under the stressed scenario . The shape of the distributions reflect how the sensitivity of each portfolio is to the stressed macrovariables. One approach to defining the limit is to consider the likelihood that each sub-portfolio will incur losses associated with the organization falling below investment grade under the stressed scenario.

Figure 11   Conditional and Unconditional Distributions of Sample Sub-portfolios and the Aggregate Portfolio
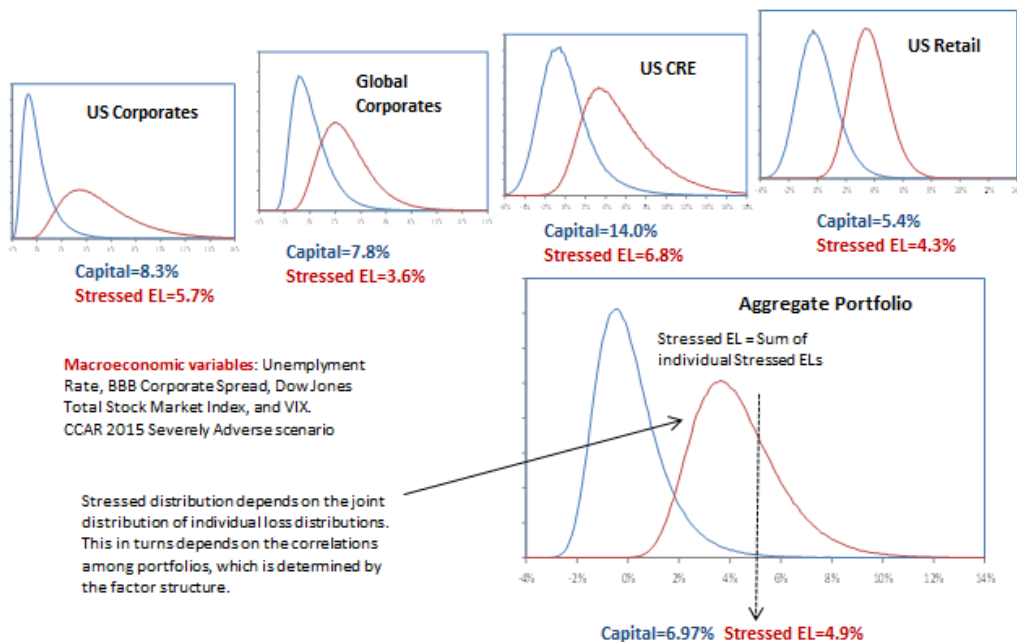
In order to incorporate ratings into the analysis, we must translate rating downgrade/upgrade events in terms of loss levels experienced by a portfolio. This process can be achieved by mapping ratings to default probabilities, which, however, must reflect the macroeconomic environment, because the limits are set in the context of a macroeconomic scenario. To this end, we leverage rating-implied EDF (Expected Default Frequency) values, point-in-time default probability measures estimated for each rating class.

An institution often holds such an amount of equity capital against a portfolio that the probability of that capital being wiped out by losses over a given period is low enough to ensure that the instituion retains a certain minimum rating. Let us consider an example, where the desired rating is Aa3 in 2014Q3 (analysis date for CCAR 2015 scenario), which means that the institution must hold equity capital that could be exceeded by portfolio losses with a probability of, at most, 2bps. This is the default probability associated with Aa3 rating in that quarter. Note, this is exactly the same default probability for Aa3 as in 2008Q1. The amount of equity capital now can be determined using the (unconditional) portfolio loss distribution. Considering the aggregate portfolio from Table 2, the equity capital should be 8.85%.[10]

Now let us assume that the institution is interested in the case when it can potentially lose its investment grade rating if the CCAR 2015 Severly Adverse scenario occurs between 2014Q3 and 2015Q3. Losing the investment grade rating means that the institution's equity capital over that one-year period will deplete to such an extent that at the end of the period, the institution's default probability will be at least 64bps. This is the default probability associated with the Baa3 rating — the poorest investment grade rating — as of 2009Q1, and we use it as proxy for Baa3 default probability as of 2015Q3. Having a higher default probability, therefore, leads to a downgrade to a speculative grade.

In our example, we assume that this downgrade to a speculative grade occurs if losses on the aggregate portfolio over the one-year period exceed threshold $Th^L_{Spec} = 5.39\%$.[11]

Let us shift the discussion to setting limits on individual sub-portfolios under the macroeconomic scenario. Within our example, we define the limit on a sub-portfolio in terms of the probability (under the scenario) that this sub-portfolio will drive the losses on the aggregate portfolio beyond the threshold Th^L_{Spec} and, thus, trigger a downgrade to a speculative rating. To determine this probability, we must make an assumptiona about the the losses on the other sub-portfolios under the scenario. We assume they are equal to stressed expected losses in excess of the uncondtional losses. We can, therefore, write the probability in the following way:

$$P^{Scenario}\left(L_{1,Sub,i} + \sum_{j \neq i} EL^{Scenario}_{1,Sub,j} > Th^L_{Spec} \cdot V_0 \right)$$

Where $L_{1,Sub,i}$ is the loss on a sub-portfolio i over the one year period, $EL^{Scenario}_{1,Sub,j}$ are the stressed expected losses on the remaining sub-portfolios, and V_0 is the value of the aggregate portfolio.

We present these probabilities in Table 3.

---

[10] The value 8.85% is the 99.98[th] percentile of the distribution of aggregate portfolio losses with respect to expected loss. Expected loss should not be included in the capital, as it is typically accounted for in allowance for losses on loans and leases (ALLL).

[11] This calculation can be carried out with certain assumptions, which are institution-specific. Let us provide a highly stylized description of this calculation. Denoting equity capital at the beginning of the period EqC_0 and at the end of the one year period EqC_1, suppose that the EqC_0 is set so that $P(L_1 > EqC_0) = 2bps$, where L_1 is the portfolio loss with respect to unconditional EL_1. EqC_1 depends on several factors — the losses experienced over the period (L_1), new volume, amortization, and maturities of exposures, and importantly on the equity capital raised over the period. Many of these factors depend on investment strategy and capital planning by the institution. However, for our purposes, the most important is dependence of EqC_1 on L_1: $EqC_1(L_1)$. Downgrade to a speculative grade occurs if $EqC_1(L_1)$ drops below a certain threshold Th^{EqC}_{Spec} , defined by equation $P(L_2 > Th^{EqC}_{Spec}) > 64bps$, where $64bps$ is the default probability associated with Baa3 rating as of 2015Q3, and L_2 is the loss on the institution's portfolio over the period 2015Q3 – 2016Q3. This means that the threshold Th^{EqC}_{Spec} depends on the characteristics of the portfolio as of 2015Q3 — including its composition and credit risk parameters, such as PDs, LGDs, and correlations. The last step is to translate the threshold on capital to threshold on losses over the first period — that is, condition $EqC_1(L_1) < Th^{EqC}_{Spec}$ implies a condition $L_1 > Th^L_{Spec}$, which, in our example, derives as $Th^L_{Spec} = 5.39\%$.

Table 3

Probability of Causing Downgrade for Sample Sub-portfolios

RANKING OF STRATEGY RETURNS	U.S CORPORATES	GLOBAL CORPORATES	U.S. CRE	U.S. RETAIL
Total Commitment	36 billion USD	26 billion USD	13 billion USD	14 billion USD
Probability of causing downgrade to a speculative grade	19.1%	6.1%	5.5%	<1bps

It is worth noting that patterns in Table 3 depend on several sub-porttfolio characteristics — its overall weight in the aggregate portfolio, as well as its sensitivity to the scenario and residual volatility after conditioning on a scenario. As a result, the U.S. Corporates portfolio, which constitutes the largest sub-portfolio of the aggregate portfolio, also has the highest probability of driving the aggregate portfolio to the downgrade. However, this probabability for the U.S. CRE portfolio is not that lower than for the Global Corporates portfolio, even though the U.S. CRE portfolio has a substantially lower commitment, which can be attributed to high stressed EL and high conditional volatility on the U.S. CRE portfolio. On the other hand, the dispersion of the U.S. Retail portoflio losses is not high enough to cause a downgrade on the aggregate portoflio.

A limit can, for example, prescribe that the probability of causing the downgrade should not exceed 25%. If that limit were breached, a possible action could be to shift the portfolio balance to exposures less sensitive to the macroeconomic variables considered in the scenario.

4. Practical Applications

There are a number of practical considerations to implementing the methodologies outlined above. Most notable is ensuring the limit setting process aligns with an organization's operating plan and core competency. As an example, a regional bank focused on domestic lending in the U.S. will find that its country risk limit outside the States will be beyond anything the organization would consider. Meanwhile, the limit for the U.S. will be too conservative, given that the entire portfolio is domestic.

To address this issue, the limit setting process must consider the geographic composition of the organization's investment market. In this example, appropriate geographic-based limits would likely be regional classifications within the United States and account for the organization's market concentration. As we formalize below, a more practical risk limit can be defined as a risk-based growth rate above each of the existing portfolio segments. Where segmentation can be defined based on business process, to more easily manage incentives and, with roughly similar level of overall risk, to more easily classify the limit hierarchy.

The remainder of this sections formalizes a few approaches that address cases with practical challenges associated with using quantitative methods when setting limits. The approaches can be used in broader contexts. And while challenges are very specific to an institution's specific business process, the spirit of how to deal with these challenges typically has some commonality.

4.1 Business Strategy and Operating Plan

In the introduction to this section, we presented an example of a bank entirely focused on domestic investing. We provided intuition for why setting country limits using the methodology outlined in the Introduction makes sense. As is obvious, a naive application of the risk limits at the country-level does not serve a useful business purpose.

A more useful approach defines limits based on geographic segmentation, limits that align with the organization's structure. In the example, one can imagine lending business segmented by U.S. region, perhaps Northwest, Southwest, Central, Northeast, and Southeast. This structure allows the respective business owners to have clear ownership over their respective portfolio risks (not to say that limits should only be defined by business lines). Even within this structure, one can imagine that the organization has an unusually high investment concentration in, for example, the Northeast, and an application of the risk limits methodology to the proposed geographic segmentation still not properly aligned with business practice; limits would likely be too tight for the Northeast and too lax for the other geographic classifications.

To define risk-based limits that are in-line with business operations, one can begin with the existing portfolio, confirm sub-portfolio risk-profiles are in-line with the organization's risk appetite, and then define risk limits that allow for each sub-portfolio to grow. The degree of growth can be defined in the risk space (e.g., each sub-portfolio can increase risk allocation by no more than 10%). The proportion of growth must be defined so that the portfolio conforms with the organization's risk appetite when a limit is hit.

4.2 Portfolio Dynamics

In some instances portfolios are unusually dynamic, with sub-portfolios exhibiting extreme growth rates. For example, when an organization enters new markets or existing markets. Using the current portfolio as a starting point to define sub-portfolio growth in risk allocation may be distortive, as risk concentrations change drastically, resulting in portfolio-referent risk characteristics of other parts of the portfolio being affected. Instead, an organization can consider using scenario-based reference portfolios as the starting point for defining the growth limits. The portfolio for each scenario can be defined either through expert judgment as the organization considers its desired market share for a business line. Alternatively, portfolios can be defined using quantitative methods with dynamics related to various macroeconomic scenarios.[12] What matters in the end is that relevant reference portfolios are used in defining portfolio referent risk measures, as well as the degree to which risk growth is considered appropriate in the limit setting process.

4.3 Sub-portfolio Limits and TOTH Limits

In some cases, it is important to define limits so that at least one sub-portfolio limit is breached when the TOTH limit is breached. This process can be done by defining the sub-portfolio limit dynamically as a function of, for example, how far the TOTH risk measure is to its limit. We use the following algorithm as an example:

Define k_j to be the relative distance of the TOTH risk measure is to its limit, and sub-portfolio j's risk measure is to its limit at time of the limit definition, $t = 0$:

$$\mathrm{sub-portfolio}_j \mathrm{risk\ measure}_{t=0} - \mathrm{sub-portfolio}_j \mathrm{limit}_{t=0} = k_j(\mathrm{TOTHrisk\ measure}_{t=0} - \mathrm{TOTHlimit})$$

We can now define a dynamic sublimit for j that will be violated whenever the TOTH limit is violated:

$$\mathrm{sub-portfolio}_j \mathrm{risk\ measure}_{t=\tau} - \mathrm{sub-portfolio}_j \mathrm{limit}_{t=\tau} = k_j(\mathrm{TOTHrisk\ measure}_{t=\tau} - \mathrm{TOTHlimit})$$

Alternative mechanisms can produce the same outcome, this is simply an example of one.

5. Conclusion

Quantifying risk appetite is central to the limit setting process. While an organization's RAS can be qualitative and difficult to translate into quantitative metrics, this paper defines a number of intuitive approaches. We define both risk and macro-scenario-based limits that can be directly linked to common components of an organization's RAS. Aside from advances related to the quantification of limits, the methodologies this paper outlines align with an organization's economic capital framework, which provides a foundation for a granular description of credit portfolio risk.

We illustrate this alignment using a number of case studies that demonstrate the relationships between varying sub-portfolio characteristics and the various approaches to setting limits. In practice, of course, applying these approaches requires not just quantitative methods, but practical considerations for strategic business needs.

---

[12] Levy, Amnon, "Quantifying PPNR Modeling," Moody's Analytics White Paper, 2014.

References

Levy, Amnon, "Quantifying PPNR Modeling." Moody's Analytics White Paper, 2014.

Financial Stability Board, *Principles for an Effective Risk Appetite Framework*. November 18, 2013.

Institute of International Finance*, Implementing Robust Risk Appetite Frameworks to Strengthen Financial Institutions*. June 2011.

Pospisil, Libor, Andrew Kaplin, Amnon Levy, and Nihil Patel, "Applications of GCorr™ Macro: Risk Integration, Stress Testing, and Reverse Stress Testing." Moody's Analytics White Paper, 2014.

Pospisil, Libor, Jimmy Huang, Mariano Lanfranconi, Albert Lee, Amnon Levy, Marc Mitrovic, Olcay Ozkanoglu, Nihil Patel, and Kevin Yang, "Using GCorr™ Macro for Multi-Period Stress Testing of Credit Portfolios." Moody's Analytics White Paper, 2015.