THE USE OF RISK ADJUSTED CAPITAL

TO SUPPORT BUSINESS DECISION-MAKING

By Gary Patrik

Stefan Bernegger

Marcel Beat Rüegg

Swiss Reinsurance Company

Casualty Actuarial Society and Casualty Actuaries in Reinsurance

1999 Call for Papers

CARe meeting, Baltimore, Md.

June 6-8, 1999

OUTLINE:

0. Abstract

1. Pressures on Capital (Surplus)

2. How Much Capital Do You Really Have?

   Risk Bearing Capital (RBC)

3. How Much Capital Do You Really Need?

   Risk Adjusted Capital (RAC)

4. How To Estimate RAC: Underwriting Model

5. Modeling Man-made Major Catastrophes (Threats)

6. How To Estimate RAC: Investment Model

7. How To Estimate RAC: Credit Model

8. How To Estimate RAC: Putting It Together

9. How to Allocate RAC to Line, Product, etc.

10. Managing RAC To Optimize Risk and Return

11. Conclusion

12. Bibliography

13. Appendices

0. Abstract

This is partially a conceptual paper about the reasons why an insurance company should address risk and capital issues in a methodical manner and about the problems encountered doing so. But it also offers some mathematical methods for dealing with some of the problems. We do not offer the reader the final answer, since we certainly don't have it. But we do offer some ideas and some procedures for obtaining useful measurements. Without reasonably accurate parameter estimation, the most sophisticated dynamic financial analysis model is simply a black box mapping information according to the "garbage in, gospel out"–syndrome (let us all bow down to our computers and worship their unarguable output!).

Modelers of insurance risk may find value in the discussion of modeling man-made major catastrophes via the construction of threat scenarios. The section on modeling investment risk discusses possible ways of using the prevailing Value at Risk model and some problems in doing so. The section on credit risk outlines the modeling problems encountered here. The reader may find the discussion of capital allocation to be particularly enlightening. In the section on managing risk adjusted capital (RAC), we attempt to show, as simply as possible, how the concept of RAC can be used by management to steer the course of business decision-making. The Bibliography lists some references which the reader can use to learn more about the ideas presented in this paper. And the Appendices contain more mathematics about some of the models and their estimation.

At this stage we also want to mention that the overall capital estimation and allocation

methodology described in this paper is intended for a company's internal risk management. It cannot be used in the same way by external parties such as regulators, rating agencies, etc. These external parties need a standard model for the whole industry, and they must rely only upon publicly available information.

1. Pressures on Capital (Surplus)

We use the terms capital and surplus interchangeably throughout this paper.

The pressures on insurance industry capital are intense and conflicting. Company shareholders, policyholders, insurance regulators and rating agencies are all pushing and pulling in different directions. The shareholders want their capital to perform, that is, earn a higher return. But there are many obstacles. Industry returns-on-equity (RoE) have been weak historically; most critics see them as being less than commensurate with the risk level. In addition, long-latent claims are still a drag upon the results of many companies. But yet, many people believe there is excess capacity currently in the insurance industry. Rates are decreasing, thus driving down profit margins. It is a situation of too much capital chasing too little business. To satisfy shareholders hoping to obtain a higher return, there is an intense competitive push to assume more risk in order to use capital more efficiently.

Meanwhile, policyholders, insurance regulators and rating agencies are all pulling in the direction of higher capitalization. They are concerned about insurance company solvency in light of the recent greater recognition of the industry's extreme exposure to natural catastrophes, the emergence of claims stemming from many long-latent man-made exposures and the threats of future claims from many similar exposures. The recent savings and loan crisis in the US has made insurance regulators and rating agencies aware that such a crisis could also possibly occur in the insurance industry if a claims shock is accompanied by a financial shock. In order to pull insurance companies to a higher, more conservative capital base, the NAIC has formulated the concept of risk-based capital to define relatively high capital thresholds for companies operating in the US [ref. 1.1].

With these intense, conflicting pressures, insurers need a better concept of capital in order to measure capital adequacy and to help steer decision-making throughout their companies.

We will discuss various kinds of capital. The three main types we distinguish are:

(1.1) Types of capital

- Publicly-perceived capital: This has more than one value. These are the various values of capital calculated by the statutory or GAAP financial statements, the NAIC, A.M. Best, Standard & Poors, etc. These are external views.

- Risk bearing capital (RBC): In section 2, we will define a simple calculation of the capital that the company has available to support its business. Note that the RBC gives an internal view, and is distinctly different from the NAIC risk based capital concept, which we would classify as one of the publicly-perceived types of capital.

- Risk adjusted capital (RAC): In section 3, we will define a simple calculation of the capital that the company needs to support its business. Again, this will be an internal view.

2. How Much Capital Do You Really Have?

Risk Bearing Capital (RBC)

The simplest answer to the question of how much capital you really have is financial statement capital, either statutory or GAAP. This of course equals financial statement assets minus financial statement liabilities. It has the advantage of being very simple. It also has the advantage of being audited; it is independently verified and signed-off by professionals who are potentially liable for negligence if, for example, the future run-off of loss reserves turns out significantly different from that stated. It is also public information, printed in black and white, for review and comment by any critics or other interested parties.

A problem with financial statement capital is that it doesn't give a complete picture of the value of an insurance company. The time value of money is generally not recognized for property and casualty companies. It does not recognize various "hidden" values such as goodwill. But worst of all, it is a snapshot picture. It is not a dynamic view of an ongoing, active company. It is not forward looking. It looks backward only to previous exposure.

A better view of how much capital a company really has to support the risk generated by its business is given by the concept of risk bearing capital, or RBC. A very simple, operational view of RBC can be obtained as follows.

(2.1) RBC = financial statement capital

plus any unrealized capital gains not included above

plus the discount inherent in the loss reserves

plus other "hidden" values

minus "latent taxes"

The latent taxes are those that would occur if the three plusses listed above flowed through income.

One can argue ad nauseam about which financial statement to start with: statutory or GAAP. Clearly, whole chapters of lengthy financial accounting books can be written about the exact treatment of the other items. And actuaries can go on for days about how to discount loss reserves, at what interest rates, etc. We do not wish to prescribe too much here. Since we intend this to be an internal view, we believe it is up to the individual companies and their technical staffs to decide how exact a measurement they want. The main thing is to do something along these lines. Don't worry too much about dotting i's and crossing t's. Since there are so many fuzzy issues and difficult-to-measure variables in any endeavor like this, trying to be overly exact is wasted effort.

The main point is to devise for your company some measure of risk bearing capital in order to give management a reasonably accurate picture of the amount of capital they have available to support current and possible future business.

3. How Much Capital Do You Really Need?

Risk Adjusted Capital (RAC)

Perhaps the question that should be asked first is: Why must we have any capital at all? If we were dealing with a situation where run-off of current liabilities and the earnings on current assets were completely predictable and where the company was not about to assume additional liabilities and assets in the coming year, there would be no need for capital. The concept of capital makes no sense without the concept of risk. And risk has entirely to do with the unpredictability of future events. The claims run-off will never be the same as predicted; the future earnings on and the future values of current assets will be, except in unusual circumstances, also unpredictable. And of course future assumed risk-transfer business by definition is unpredictable.

Capital is necessary for the future. It is not a static concept for either claims run-off or for ongoing business.

This tells us why NAIC risk-based capital or most rating agency models do not give us good answers to the question of how much capital do you really need. These are relatively simple models; they are quick 'n' dirty, geared to adequacy, not optimization. They usually have "faulty thermostats", employing simple ratios of premium or loss reserves which indicate less required capital when rates or reserves are inadequate, which is of course exactly when you need more capital. They also don't recognize management or shareholder risk-level preferences. Like financial statement capital, they are retrospective, not prospective. Finally, they are not easily translatable to lines of business, types of contracts, profit centers, etc., in a manner that is useful for supporting business decision-making.

We may classify the risks we should consider when defining risk adjusted capital for an insurance company.

(3.1) Insurance company "technical" risks

Underwriting risk

- Claims
- Rating system biases: parameters, formulas, etc. Some actuaries would further split this according to the concepts of parameter risk versus process risk, or split even further with the concept of model specification risk.
- Underwriting cycles

Investment risk

- Market risk: stock market, interest rate, foreign exchange
- Default

- Liquidity
- Etc.

Credit risk

- Reinsurance
- Accounting balances due
- Letters of credit
- Etc.

Note that we consider only “technical” risks, ignoring softer concepts which might be gathered under the rubric “management risk”. These other kinds of risks are better handled outside of a technical model.

What criteria do we want our definition of risk adjusted capital to satisfy? We want RAC to be the level of capital an insurance company needs to write its business. Among many possible criteria, we select the following.

(3.2) Criteria for RAC

- It must meet specified management risk and survival criteria.
- It should quantify the risk/return trade-off for all risk exposures.
- It must be useful for making appropriate risk-based business decisions.

Management risk criteria have to do with the public statements of the company’s results. These criteria may have to do with year by year fluctuation of results or with downside potential. Although fluctuations are annoying, the real fear of management is downside potential. They fear that an event or a series of events may occur that might cause the company's publicly-perceived capital to fall far enough to interfere with the company's ability to continue normal business and also to raise serious questions about the company's continuation.

Since it is absolutely critical that the model reflects management risk tolerance, it is absolutely critical for management to understand enough of the model so that they can understand what is being asked of them, so that their opinions are translated accurately into the model structure and parameterization. Thus we propose a very simple, non-black box model for RAC.

There are three steps in constructing our very simple RAC model.

(3.3) A very simple RAC model

- Management specifies a simple risk tolerance rule.
- The probability distribution of the company's result is estimated for the time period specified by the management risk tolerance rule.
- RAC is set to be the minimum RBC necessary at the beginning of the time period so that the RBC at the end of the time period satisfies the management risk tolerance rule.

As an example, in this paper we will use a very simple management risk tolerance rule.

(3.4) A very simple management risk tolerance rule

- Define the management Risk Tolerance Level (RTL) to be that value of RBC necessary to maintain a given external rating, e.g., A.M. Best A rating, S&P BBB rating, etc.

- With a specified probability, e.g., 90%, 95% or 99%

- At the end of a time period of one, two or n year(s).

You can see that this is indeed a very simple management risk tolerance rule. To simplify our discussion, let us assume that the probability is 99% and the time period is one year. A time period of one year is not long enough to completely model the effects of potential shocks to capital caused by the manifestation of long-latent claims. And of course a single probability level is also very simplistic: why not use 99.5%? 98%? And of course management has more on their mind than the company's A.M. Best rating. However, let us walk with management before we make them (and ourselves) run.

This definition of RTL is the same as the definition of u_C, the early warning limit of capital in Daykin, Pentakäinen and Pesonen's *Practical Risk Theory for Actuaries* [ref. 3.1, p.365]. Note that the RTL is defined to be a minimal level of RBC, but yet the first criterion refers to a company's selected external rating that utilizes one of the publicly-perceived types of capital. Thus we must construct a mapping of the RTL level of RBC onto the selected external rating's capital level. Rather than getting bogged down in this construction, we leave this detail to the reader (remember not to worry too much about i's and t's.).

By the company’s result during the time period (in our case, one year), we mean simply the change in RBC from beginning to end.

(3.5) Company’s Result

CR = { ending RBC} minus { beginning RBC}

Note that CR is a calendar year concept.

CR may be thought of, and modeled, as the sum of three components.

(3.6) CR = company’s calendar year underwriting result
plus company’s calendar year investment result
plus company’s calendar year credit result

A simple graph of the probability density function of CR is illustrated in (3.7).

(3.7) CR probability density function

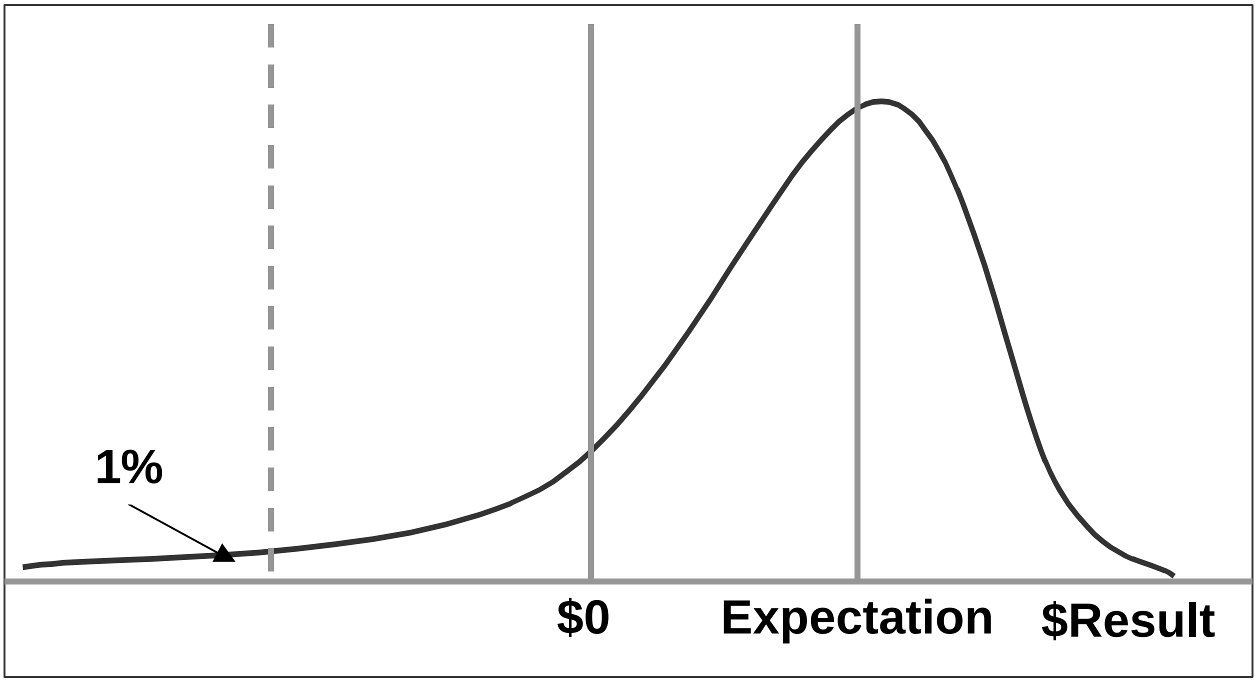

Note that the distribution, whose random variable is essentially calendar year premium plus investment result minus expenses minus incurred claims, is decidedly non-normal. It is skewed to the left since the distribution of incurred claims is skewed to the right [ref. 3.1-3.3] and the distribution of the investment result is usually modeled to be essentially normal [ref. 3.4]. What else can we say about the distribution of CR? The expectation is the expected (planned?) result for the calendar year. The distribution obviously depends upon the degree of diversification in the company's underwriting, investment and credit risk portfolios: the more diversified (lower correlation), the narrower the spread.

We can now construct a simple picture of how RAC is defined.

(3.8) Definition of RAC

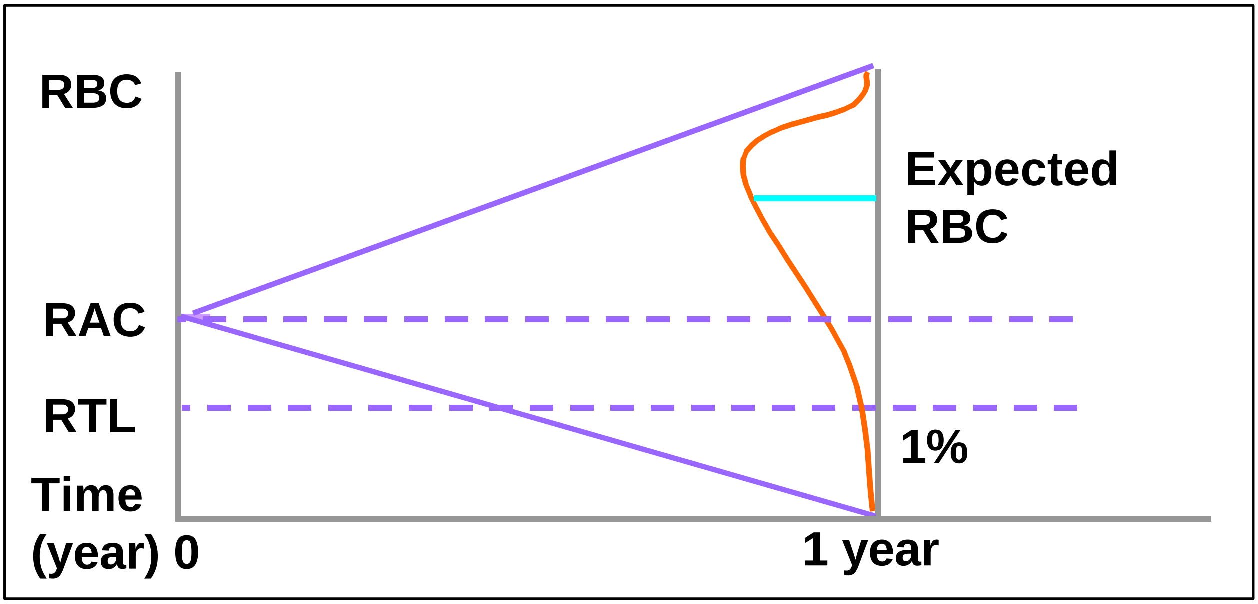

Given this simple model and definition of RAC, we have the following relationships.

(3.9) RTL = RAC + $CR_{1\%}$

or

RAC = RTL - $CR_{1\%}$

where $CR_{1\%}$ is the first percentile of the distribution of CR.

In order to test our understanding of RAC (may we say RACology?), we can ponder the following results.

(3.10) Some simple RAC results

- As diversification increases (lower correlation among risk portfolio segments), RAC decreases. Thus the amount of RAC relative to "risk volume", e.g., premium, loss reserves, assets, etc., provides one measure of diversification.

- As RTL decreases, RAC decreases.
- As the probability (99%) decreases, RAC decreases.
- As the time period decreases, RAC decreases.
- As RAC decreases, the company's RBC can be reduced or additional business can be written. As a consequence, the company's needed target pricing margins will decrease.

4. How To Estimate RAC: Underwriting Model

Our simple RAC depends upon the distribution of CR, the company's calendar year result, and CR has three components: underwriting, investment and credit results. The first step is to estimate the distribution of the company's calendar year underwriting result. We define the underwriting result, CR_U, as follows:

(4.1) CR_U = calendar year net earned premiums
minus calendar year net expenses
minus calendar year net incurred losses

all discounted at selected risk-free investment returns

The net result should however not be modeled directly, but via a model for the gross result and a separate risk transfer model for ceded reinsurance. This approach then allows us to address separately in the credit result model the impact of the credit risk arising from the ceded reinsurance. In the following we will describe the modeling of the gross result.

Why complicate the underwriting model by reflecting risk-free investment income? Why not simply leave all the investment income in the investment result? The reason is that once you have a model for RAC, inevitable questions arise about the risk/return contribution of the various components of the model. Management will quite rightly want to know which business segments are contributing their fair share and which are not. To push the RAC model of returns down to business segment level requires a corporate-wide business evaluation system reasonably consistent with the RAC methodology. It is easy to see that if such a business evaluation model is to properly evaluate underwriting results, it must reflect some kind of investment income to be able to accurately evaluate the relative values of the various short and longer-tail business segments. Many actuaries considering this issue have opted to reflect some kind of risk-free investment income only. The thought is that adding in total investment income would unduly distort the distribution of underwriting results because of the inclusion of too much investment risk. It is thought that it is better to account for this additional risk elsewhere, in the investment model.

But how should investment income be reflected in calendar year results? In order to have a good underwriting evaluation model, calendar year results must be modeled by first modeling accident year or policy year results, so that premiums, expenses and losses can be tied together for different exposure periods. Risk-free investment returns are used to discount all cash flows arising from each accident or policy year to a single evaluation date. These risk-free returns may differ for each year. For the calendar year result, the calendar year discounted cash flows can reflect any changes in each accident or policy year's results during the calendar year that were not expected at the beginning of the calendar year. For the simplicity that arises from having non-overlapping exposure periods, you might choose to model accident years instead of policy years, as we do.

So now we are in the situation of modeling the distribution of the discounted change in result during the next calendar year for each past accident year, and also for the next accident year. A stochastic model for premium and loss development is most helpful here. The premium model is simpler than the loss model except for the inclusion of retrospective or other later premium adjustments and payments. But since the really significant premium changes rely upon changes in loss evaluation, and can thus be modeled as a function of the losses, let us concentrate upon the modeling of the losses.

We want a stochastic model for accident year losses and the calendar year incurred loss development thereof. There are many such models in the actuarial literature to choose from [ref. 4.1 - 4.3]. Whichever model is selected, we need a reasonably good parameterization in order to produce reasonably good answers – remember the garbage-in-gospel-out syndrome? A reasonably good parameterization for the accident year loss distribution can be obtained from historical information suitably adjusted to future level by the following steps.

(4.2) Loss distribution modeling steps

1. Define future potential major catastrophes.

2. Filter the major cats out of the historical loss data.

3. Model the filtered, non-major cat losses, adjusted to future level.

4. Model the future major cat scenarios.

5. Glue them back together.

The major cats are those large loss events whose presence or absence from the historical loss data distorts the estimation of future loss occurrence or loss run-off potential. These major cats can be either natural or man-made catastrophes. Examples of realized major cats distorting recent loss data are Hurricane Andrew, the Kobe earthquake, asbestos and pollution clean-up. Depending upon the company's insurance portfolio, there may be others. This is an important point: what constitutes a major cat event for a particular company depends upon the company's particular insurance portfolio.

An example of a major cat whose absence distorts the historical loss data, and thus the extrapolation to the estimate of future loss potential, is the non-occurrence of an earthquake centered on the New Madrid fault. The future occurrence of this event is believed to be very possible and of very large loss potential, but yet there has not been an occurrence since 1812.

Major cats are very rare events that may be difficult to identify and are certainly difficult to quantify. Yet these events are the key to understanding and measuring insurance company underwriting risk.

Let us separate the major cat scenario modeling into two types: natural and man-made. Natural major cats are easier to model than man-made. Much data exists for smaller and medium-sized natural catastrophes, and data exists for some larger events like Hurricane Andrew and the Kobe earthquake. Reinsurers especially have devoted much time and effort analyzing insurance exposure to natural catastrophes, and there are some reasonably good, commercially available models to quantify the exposure of any insurance portfolio. The models seem to be doing a reasonably good job estimating loss severity. The main problem is the estimation of loss frequency, the inverse of which is sometimes referred to by underwriters as “return period”.

The more difficult problem is the modeling of man-made major cats. These are loss events arising from the unknown risks of technological, economic, legal and social development. This development may include the expansion of insurance coverage for what had been considered to be “business risks”. If we take asbestos and pollution clean-up as canonical examples, we can say that man-made major cats have the following characteristics.

(4.3) Characteristics of man-made major cats

- Arising from technological, economic, legal and social development
- No single, well-defined event causing all the claims
- Long claims discovery periods
- Unknown number of claimants
- No geographical limitation

Modeling man-made major cats is very difficult. There is great uncertainty regarding appropriate model structures and realistic parameter values. Historically, we have had so far only two major events which might be considered to be man-made major cats: asbestos and pollution clean-up, and the ultimate insurance losses from these two events are still very much unknown [ref. 4.4]. But since these man-made major cats are so important to the evaluation of future loss occurrences and loss run-off, and thus to the evaluation of risk adjusted capital, we must do something.

The modeling of these man-made major cats is so complex and yet so important that it deserves a section to itself. So we will put off this modeling discussion until the next section. Let us temporarily assume that we have successfully modeled these man-made major cats, and thus continue the discussion of the underwriting model.

The filtered loss distributions may be modeled using standard actuarial methods. We will want to create the distributions of the next calendar year incurred losses arising from each previous or current accident year and also the next accident year. The standard methodologies tell us to segment the company's insurance portfolio into its major components. Some of these components may be single contracts large enough and significant enough to be analyzed and modeled separately. In addition to line of business, a well-diversified company may also need to consider geographic area. Suitable historical exposure and loss data must be analyzed, model structures determined and parameters estimated. The models can then be extrapolated to the next calendar year. Note that this extrapolation itself increases uncertainty because of future inflation and market risks. Most actuaries would say that this increases parameter risk.

It is important to also model and estimate correlation among the business segments, both loss event correlation and pricing or rate-level correlation. The loss modeling obviously depends upon the exposure estimates by year. And let us not forget that the modeling of premiums and expenses depends upon the rating/underwriting cycle and the degree of pricing or rate-level correlation among the lines.

The model for the filtered losses is combined with the models for natural and man-made major cats to yield the distribution of gross incurred loss for the next calendar year. Combine this with the model for premiums and expenses and risk-free investment income on the cash flow to obtain the distribution of the underwriting result for the next calendar year.

Now let us return to the modeling of man-made major cats.

5. Modeling Man-made Major Catastrophes

Before describing an approach to the modeling of man-made major cats, we want to point out the goal once more. Both the natural and man-made major cats have a significant impact on the company's result, affecting the tail of the loss distribution, and thus the calculation of RAC. But since they occur with such low frequency, they cannot be modeled appropriately on the basis of historical loss records only. This is especially true for the man-made cats. So instead, we model them using additional information. Think of it this way. Instead of using only statistical and actuarial methods to model the tail of the loss distribution, we want to take into account as much information as possible. This information comprises hard data such as past loss experience, portfolio information and market information, but also soft factors such as expert-knowledge and gut feeling. In addition, quite a bit of pure guesswork is necessary.

Some important issues related to the modeling of man-made major cats are illustrated with following example. Suppose we ask ten experts to estimate the frequency of an oil tanker pollution event with an impact comparable or worse to the Exxon Valdez accident. We will probably get more than ten answers! This is not only because of different opinions among experts, but also because of a different understanding of the question: what do we mean by comparable impact? Is it the impact on the environment, on the society, on the economy or on the insurance industry? Are we asking for the expected frequency in the next year, in the next 10 years or in the next 100 years? Are we asking for the frequency of oil tanker accidents only, or do we want to include oil platforms as well? Are we asking for the frequency just for Alaska or are we interested in the worldwide number of similar events?

Because of this ambiguity, we must accept the fact that there cannot be a single correct answer. We will always end up with a set of possible solutions. Obviously, there is not only uncertainty related to the stochastic nature of the modeled events. The distributions used to describe this stochasticity are also uncertain. This is usually described as model and/or parameter risk. Whenever possible, this uncertainty should also be included in the overall RAC model.

The tail of the loss distribution to be modeled should represent all events which might have a major financial impact on the company. In order to derive a reasonably accurate distribution, we want to quantify the impact (severity) and the occurrence probability of all relevant events. This is impossible. Therefore we must restrict the evaluation to a limited number of representative scenarios. In doing so, we can exactly define which events we want to consider and which assumptions we want to make. For each such specific event, it is then possible to quantify the financial impact on the company (this is a sort of stress testing). Much more difficult is the estimation of the occurrence frequency of the events represented by each representative scenario. We could of course also start with the frequency of the events and then try to quantify their severity in a second step (what is the impact of an event expected to occur with a frequency of 0.01?).

It should be clear that there is not a unique methodology for modeling man-made major cats. The modeling method depends on the nature of the risks, the line of business, the available information, etc. A model for man-made major cats will never be complete, but will have to be adapted and modified over time in an ongoing process.

The modeling of future man-made major cats involves underwriters, actuaries, scientists, engineers, claims people and financial analysts. The steps involved are as follows.

(5.1) Modeling man-made major cats

- Define the characteristics of a major cat event
- Identify potential major cat exposures
- Estimate the frequency probability of each major cat event
- Estimate the severity distribution of each major cat event
- Translate all this into insurance coverage

We have already discussed some defining characteristics of man-made major cats. But let us restate them somewhat differently.

(5.2) Characteristics of man-made major cats

- No single, triggering event
- Unknown temporal duration

  (development of loss exposure over many years)
- Unknown geographical boundaries

  (geographic development from local to global)

- Limited knowledge of event frequency and severity probabilities

- Limited knowledge of insured values because of unknown:

  - lines of business exposed
  - types of claims (BI, PD, financial)
  - number of claims, plaintiffs, insureds
  - amount of compensation per claim

- Limited knowledge of the reaction of society and insurance markets

  (If losses are to be paid only in the future and if the industry can collect enough additional premiums and if the losses need not be reserved too quickly, there may be less calendar year impact.)

The first step in identifying the major cat exposures is to hold brainstorming sessions of underwriters, actuaries, claims people, scientists, engineers and financial analysts to list scenarios which might generate major cats. Let us call these scenarios "threats". An example of such a list might look something like this.

(5.3) Potential threat scenarios (from brainstorming sessions)

Additives (food and other)	Aluminum
Animal feed	Architects/Engineers
Asbestos	Banks
Bio-gen technology	BSE (mad-cow disease)
Building materials	Chemicals
Collapse of bridge or tunnel	Directors and Officers
EMF	Food
Implants (medical devices)	Lawyers/Accountants
Lead	Pharmaceuticals
Pollution	RSI
Surveyors	Terrorism
Tobacco	Transport of hazardous material

The next step is to eliminate threat scenarios which are judged unlikely to fulfill the characteristics of a major cat event or whose likelihood of occurrence is sufficiently remote. The retained "identified" threats are those which have had some insurance industry claims activity. A pruned down list might look something like this.

(5.4) Identified threats

- Asbestos
- Building materials

- EMF (electro-magnetic force)
- Implants
- Pharmaceutical
- Pollution
- RSI (repetitive stress injuries)
- Y2K

Our next step is to classify these identified threats according to industry activities from which they might arise.

(5.5) Threats and industry activities

Threat	Industry Activity
Asbestos	Asbestos manufacturing, construction, building
Building materials	Construction, chemicals, building, owner/management
EMF	Electronics, power, railroads, communications
Implants	Medical devices
Pharmaceutical	Pharmaceuticals
Pollution	All industries; especially chemical, paper, machinery
RSI	All industries; especially construction, equipment
Y2K	Software and chip manufacturers, all industries D&O

The next step is to establish a list of triggering events from a purely scientific/technical point of view, so that we would have some early warning of possible occurrence. For each identified threat, we estimate its "degree of maturity".

The next step is to determine which insurance lines of business the identified threats might impact.

(5.6) Threats and insurance lines of business

Threat	General Liability	Products Liability	Workers Comp.
Asbestos	yes	yes	yes
Building materials	yes	yes	yes
EMF	yes	yes	yes
Implants	no	yes	no
Pharmaceuticals	no	yes	no
Pollution	yes	yes	no
RSI	no	yes	yes
Y2k	yes	yes	no

The next step is to estimate the frequency of occurrence for each threat. Since none of the identified threats will occur exactly as modeled, we must estimate the frequency for a larger set of similar scenarios. As the probabilities might change over time, it might also be necessary to specify the considered time period. As we have seen and as we may guess, past experience is most likely insufficient and is not necessarily representative of future probabilities. Because of the long development period involved, the loss probabilities are subject to technological, economic, financial, social, legislative and jurisdictional changes.

The estimation of insurance loss severity is even more difficult for all of the above reasons. In addition, we must analyze current and possible future insurance coverage standard wording regarding possible claim triggers, geographical scope of coverages, claim series clauses, cost in addition to limits and aggregation of limits. The particular insurance portfolio, past and future, must be analyzed to identify exposed insureds and coverages. We want to end up with portfolio-wide loss frequency and severity probability distributions for each identified threat. For these events, we must remember that there is a strong correlation among the losses over many accident years. This is a bottom-up threat exposure analysis for a particular insurance portfolio.

An alternate estimate for a particular insurance portfolio is a top-down market share assessment once a total or industry-wide estimate has been made of an identified threat's severity, For a very immature threat, this may be the best that can be done.

Note that the estimation procedures outlined above are the same as those used over the last 15 years to estimate asbestos and pollution losses.

6. How To Estimate RAC: Investment Model

We now turn to the modeling of the distribution of the company's calendar year investment result. We define the investment result, CR_I, as follows.

(6.1) CR_I = calendar year change in assets accounted for in RBC

minus items already accounted for in CR_U

(such as premiums, expenses, losses, risk-free investment income calculated for underwriting)

minus items that will be accounted for in credit risk, CR_C

Since we are modeling the change in a company's assets over a calendar year, the risks that must be taken into account are those arising from the nature of the company's assets.

(6.2) Asset Risks

- Interest rate risk
- Default risk
- Stock market/equity risk
- Real estate risk
- Foreign exchange risk

Constraints upon investment policy stemming from insurance regulation in many jurisdictions are a complication for insurance companies.

The modeling of the distribution of a company's calendar year investment result is a very complex and difficult problem. We will only be able to discuss some concepts and some well-recognized tools for attacking the problem. The interested reader can find most of the details in the references, but will have to decide for himself how exact the modeling should be. This modeling is closely related to the modeling of market risks in the banking world. However, there is one major difference: the time horizon, as will be discussed below. Probably the best known tool for measuring market risks is RiskMetrics™ by J.P. Morgan [ref. 6.1], also known as Value at Risk (VaR). There is a great amount of information about VaR on the Internet; a comprehensive list of published and working papers can be found in reference [6.2]. The description of VaR methodology in Appendix 13.1 is based on reference [6.3].

The general idea of VaR is to combine a stochastic model for the basic risk factors together with a deterministic model that links the value of the financial instruments in a portfolio to the random changes of the risk factors. The distribution of the change in value of the portfolio is obtained by fitting a parametric distribution (very often a normal distribution) to the first two moments.

VaR is defined as a percentile (usually the first percentile) of the result distribution of the assets. However, we are interested in the whole distribution, not in single percentiles. But since the distribution is defined by the set of all of its percentiles, it is possible to use VaR-methodology to model the whole distribution of investment risk.

However, there are several shortcomings.

(6.3) Problems with the Use of VaR To Model the CR_I

- The asset portfolio is actively managed
- VaR is a *point-in-time* concept
- VaR assumes that market risks can be modeled by means of a Markov process
- VaR assumes complete liquidity of all investment instruments
- VaR is usually based on normal distributions

In contrast to the fairly stable insurance portfolio, an investment portfolio is more actively managed over the one year time period we are considering. Thus, management interaction can influence the risk. The second point is that a fundamental assumption of VaR is that the time horizon is short, one day if not infinitesimal, because some asset portfolios may not be constant in their composition and may have time dependent characteristics. Since we need the distribution of CR_I over a period of one year, the question is how to scale the distribution to this longer time horizon. A problem is that the various methods of scaling VaR to different time horizons do not consider the influence of active management. Thirdly, the assumption that market risks can be modeled by means of a Markov process, wherein the future probabilities depend only upon the current state, not upon previous history, is somewhat questionable. Fourthly, it is clear that there is a great variation in the liquidity of investment instruments. Some are very liquid (e.g., stocks of big public companies or some standardized derivatives), while others take an intermediate position (e.g., corporate bonds) and some are very illiquid (e.g. mortgage backed securities and especially real estate). Liquidity in general depends upon economic conditions. Fifthly, the assumption of linearity between changes of the risk factors and changes of the values of the financial instruments is an approximation only. And finally, VaR is usually based on normal distributions. Since we are interested in the tail of the overall result distribution, it might be necessary to replace the normal distributions by heavier-tailed distributions.

To overcome these limitations, it is necessary to incorporate the dynamics of the portfolio management into the model. We are not in a position to offer a satisfactory solution for this problem. Instead, we will only outline the standard VaR concept and discuss a few approaches for extending this simple model. More details on VaR are in Appendix 13.1.

If a diffusion process without memory effects (i.e. no serial correlation) is used for modeling the changes of the risk factors underlying investment risk, then all variances and covariances of the components of the investment portfolio will increase linearly with time. Because all variances and covariances are scaled with time, the portfolio variance is also scaled by the same factor. If the distribution is also scaled with the same factor, it follows immediately that VaR based on one day can be scaled to k times one day by multiplying it with the square root of k.

(6.4) VaR for k days

$$VaR_{k\ days} = VaR_{1\ day} \times \sqrt{k}$$

The time horizon to which VaR can be scaled should be related to the period within which no management action is taken, or, in short, the holding period of the asset mix. However, our time horizon is one year, thus contradicting the assumption of a fixed portfolio.

So, how can we estimate the portfolio component covariance matrix for a one year time period? One path is to construct VaR from historical time series. There are many different ways of doing so. The sample mean of monthly data could be taken and be scaled to one year by taking k = 12. Using daily data would result in k=250, this being the number of trading days of a year (one could also take k=365 in arguing that the risk factors continue to change over weekends, etc.). We could also estimate the covariance matrix by using an exponential weighting scheme. However, the resulting covariances would be much more responsive to the near past and would be more volatile over time. Thus they might not be a good basis for determining the portfolio risk over a one year time horizon. Choosing a more sophisticated methodology for estimating the behavior of risk factors over one year does not relax the assumption of having no portfolio management action taking place, but may yield a better description of the risk factors.

Another modeling and estimation problem is liquidity. The vast universe of securities is very heterogeneous with respect to liquidity. Liquidity may change due to introduction of new products superseding others. Markets can get sticky because of general downturns and imitative behavior. Liquidity measured as a discount to ordinary prices can evaporate due to the size of the particular lot to be liquidated.

Big lots often trigger widening spreads and thus incur so-called “market impact costs”.

Because of this, it would seem that the liquidation period is the maximum time period that the VaR model could be extended for a single investment instrument. But, as different instruments have different liquidation periods, some means of finding a proxy is needed. One obvious choice is a weighted average liquidation period. Another choice is to incorporate the instrument specific liquidation period into the exposure. This leads to some sort of exposure weighted time horizon. Even if such an approach were adopted, giving a better picture of the today’s loss potential, the problem of expanding the horizon to one year is still not solved.

The most difficult modeling and estimation problem is that of active management. Over a one year time period, portfolio managers will revise their portfolio mix in order to realize opportunities where they foresee them, conditional on some restrictions, be they legal or other, and with an eye on loss potential. Their action is aimed at pushing the portfolio value up, or in extreme cases to stop further losses.

A strategy is a set of rules, which are applied when some instance of the state variables materialize. We can think of this as being a multi-sequential two player game where a passive strategy of the environment is confronted by the active strategy of the portfolio manager. The environment makes a move into a new state of reality. The manager then makes a move by changing his forecast and revising his portfolio. Doing nothing is also an action. Within this framework, the distribution of the final value of the portfolio can only be assessed by simulation.

The approach implemented should be a reliable model of reality, but also be pragmatic and intelligible by a broad audience. The basic VaR model is somewhat understood by many senior managers. Therefore, the VaR model should serve as a starting point. Ideally, the one-year extended VaR model percentile would be a function of the basic VaR model percentile. The simplest way of doing so is a multiplication rule of the following form.

(6.5) One-year extended VaR model percentile

$$VaR_{1year} = g(\theta) \cdot VaR_{1day} \cdot \sqrt{250}$$

where g(.) is a function of the strategy θ and the other considerations discussed above.

Several ways of estimating g(.) are pragmatically possible. We could study history to find out what range g covered in the past. Another method would be to interview the portfolio managers regarding their strategy, and then simulate the outcomes. Also, we could look at the guidelines and predefined strategies in the case of an emergency.

7. How To Estimate RAC: Credit Model

We turn now to the modeling of the distribution of the company's calendar year credit result. For convenience, we repeat and expand the list of components of credit risk noted in (3.1).

(7.1) Credit risk

- Reinsurance ceded
- Accounting balances due
- Letters of credit
- Credit surety business assumed
- Financial guarantee business assumed
- Credit derivatives
- Etc.

Note that we are including here as credit risk part of what may be thought of as part of the underwriting risk: credit surety and financial guarantee business. The reason is that this insurance coverage is much more dependant upon general economic conditions and upon the financial condition of the insured party than is the coverage for the more typical underwriting insureds.

Modeling the ceded reinsurance component is just like modeling the underwriting portfolio, except now we must account for the additional risk of future non-collectibility of claims payments by the reinsurers. This ceded reinsurance result is obviously highly correlated with the gross underwriting result.

In a very simple approach, the distribution of the credit result could be modeled by combining historical results and threat scenarios in a manner similar to the modeling of the underwriting result as described in sections 4 and 5. However, a potential credit loss not only depends upon the credit exposure and the size of some future random events, but also upon the financial strength of the counterparties (debtors) involved, i.e., upon their default probabilities. Since the default of several counterparties can be triggered by a common event (e.g., a large earthquake, a recession, etc.), there is potential strong correlation within the credit portfolio and between the credit portfolio and the underwriting and investment portfolios.

Whenever possible, the credit model should combine all relevant risk factors, i.e., credit risk should be modeled bottom up. This means that we should first model the risk of each component separately, and then, by modeling the correlation of the components, we can evaluate the credit portfolio's overall risk. In principle we will find that the ideas driving the modeling of the credit result are similar to those discussed for the underwriting and investment models.

Some credit models have so far been established in practice – see references [7.1]-[7.6]. These all follow different approaches and rely also upon different assumptions. The selection of an appropriate model depends upon the purpose and the goals to be achieved. We cannot in this paper list all the advantages and disadvantages of the referenced models. Instead, we will simply present some of their common ideas that drive the estimation of credit risk. Further details can be found in the references.

Since we believe that the credit risk should be modeled bottom up, we will discuss the modeling of a single counterparty's credit risk. We believe that this will give you some of the flavor of the modeling issues. Later we will very briefly discuss the modeling of the whole credit risk portfolio. But the full discussion of this very much more difficult problem must be left to the above noted references.

The credit risk for a standalone transaction is primarily dependant upon the counterpart credit quality (rating) and upon the contract type (ceded reinsurance, financial guarantee, zero coupon bond, floating rate, coupon bond, credit and surety, default swap, default digital, default option, etc.). A simple model has three components.

(7.2) Components of a simple model for a single transaction (contract)

- default frequency
- default exposure
- default loss severity

Usually these components can be modeled independently and glued together at the end. This separation is motivated by the fact that for simple contracts, these three components are almost independent. The default frequency usually depends just upon the counterpart's credit quality, the default exposure upon the type and limit of the contract, and finally, the default loss severity depends mainly upon the seniority of the contract. However, these three components are not necessarily independent for the case of ceded reinsurance, since the default probability of an assuming reinsurer and the default loss severity might both depend upon the size of a specific catastrophic event (e.g., an earthquake in California). And yet if we are clever enough in the construction of our model, we may still be able to use this simple model for the default of a reinsurance cession.

We will now discuss the modeling of these three components, with special attention paid to the default probability because it is the most sensitive part of the integrated model.

(7.3) Two methods for estimating the default frequency

- Rating agency statistics

- Black-Scholes-Merton model

The default frequency corresponds to the probability that the company will default during the specified time period (in our case, one year).

Rating agencies analyze the credit worthiness of each company by looking at their cash flow, profitability, financial flexibility, industry sector, market position, competitors, management, controls, financial reporting and legal structure, etc.. This process is known either as fundamental analysis or as the rating process. Its outcome is a valuation of the company's credit worthiness. Both Moody's and S&P can estimate the default probability of each rating class over various time horizons based upon historical data. The following table shows default probability percentages as a function of rating class at the beginning of the time period and the length of the time period.

(7.4) Default frequency (probability) by S&P rating class [ref. 7.7]

Term	**1**	**2**	**3**	**4**	**5**	**7**	**10**	**15**
AAA	0.00	0.00	0.07	0.15	0.24	0.66	1.40	1.40
AA	0.00	0.02	0.12	0.25	0.43	0.89	1.29	1.48
A	0.06	0.16	0.27	0.44	0.67	1.12	2.17	3.00
BBB	0.18	0.44	0.72	1.27	1.78	2.99	4.34	4.70
BB	1.06	3.48	6.12	8.68	10.97	14.46	17.73	19.91
B	5.20	11.00	15.95	19.40	21.88	25.14	29.02	30.65
C	19.79	26.92	31.63	35.97	40.15	42.64	45.10	45.10

Note that these default frequencies are based upon historical data. However, past default frequencies may not accurately predict project future frequencies because of the presence of economic cycles. Thus our estimates of future default frequencies must take into account the current state of the economy.

(7.5) Black-Scholes-Merton Model

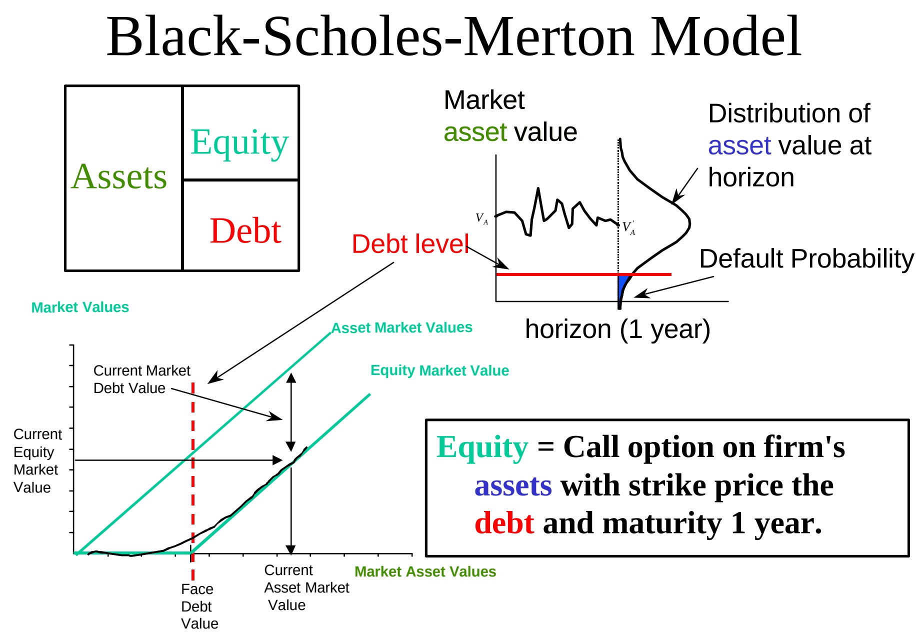

From the above picture, we see that the equity market value of a company may be thought of as the price of a call option on the market value of its assets, with strike price equal to its debt (liability) value. To evaluate the call option price, we can assume that the probability distribution of asset values at the end of the year satisfy certain Black-Scholes-Merton properties, where the call option price can be calculated by the BSM-formula.

(7.6) Black-Scholes-Merton formula

Equity = BSM[underlying spot price (asset value); volatility (asset volatility);

time to maturity (1 year); risk-free rate; strike price (*debt*)]

In case we do not know the volatility, we look for the market price of a call option and then invert the BSM-formula. Similarly, since we may not be able to directly observe neither asset volatility nor asset prices, we can calculate them from market equity prices, inverting the BSM-formula and using the delta hedging formula. If we can assume that asset returns are normally distributed (mean = actual asset value, volatility = BSM implicit volatility), then we can derive the probability that the asset value falls below the threshold of debt-liability value. This probability is defined to be the expected default frequency (EDF). Note that in this model, EDF is derived from the equity market value, which is based on investors' future expectations about the company.

The second component of the single transaction model is the default exposure. The default exposure depends upon the particular contract and not much upon a company's credit quality. For the case of ceded reinsurance, the credit risk default exposure is given by the amount of loss ceded to the reinsurance contract and the shares underwritten by the various reinsurers.

Default exposure patterns are usually assumed to be flat (= average default exposure) over the time period, except for contracts that have a long duration. We must remark here that this strong assumption does not hold for derivative products. Their default exposure can increase or decrease dependent upon market movements. For these dynamic default exposures, we can calculate a default exposure envelope. This is an estimate of the default exposure probability distribution during the time period. We must also remark that collateral or netting agreements can substantially reduce default exposure. The modeling of these is straightforward, but adds considerable complication.

The third component of the single transaction model is the default loss severity. The default loss severity is the percent of the total owed which is not paid following a default. Often it is more convenient to look at the complement,
1 – {default loss severity}, which is the recovery percentage. Default loss severity mainly depends upon the contract structure (seniority), and usually less upon the rating. A standard model calculates the probability distribution of the recovery via a beta distribution fit to historical data with the key input variables being the contract seniority and rating. We finally observe that the default loss severity usually does not depend on the duration (time period).

Let us assume that we have estimated the three model components for a single transaction. We can now put them together. Since these three components are all independent, we simply convolute their distributions. The standalone capital needed for a single transaction then corresponds to a percentile (for example, 99%) minus the expected default loss, which is calculated by multiplying the expected values of the components.

(7.7) Expected default loss

Expected default loss =

EDF $\times$ E[default loss severity] $\times$ E[default exposure]

The more difficult problem is to step up from the modeling of single transactions to the modeling of a whole portfolio of transactions. To do so, we must estimate the correlation between single transactions. For the case of ceded reinsurance, some correlation between various ceded reinsurance contracts can be accounted for via threat scenario modeling as in section 5. For each scenario, we can estimate the percentage of loss which will not be recovered (the default loss severity) from each reinsurer.

Even though the modeling of the risk of the whole credit portfolio is the really difficult and more interesting modeling problem, addressing it would require a substantially longer paper. Thus we must leave this to the interested reader, who can find many good ideas in the references by Tom Wilson, [ref. 7.4–7.6]. An analytical procedure for calculating the aggregate loss distribution is too complicated (especially for a portfolio containing hundreds of contracts). The only possibility for estimating the final loss distribution is the use of Monte-Carlo simulation.

8. How To Estimate RAC: Putting It Together

The last step in the estimation of RAC is the construction of the distribution of CR, the company's calender year result. Once we have an estimate for the distribution of CR, the management risk tolerance rule defining RTL will produce a value for RAC via (3.9).

Remember that CR is simply the sum of the underwriting, investment and credit results (3.6). After doing the hard work of constructing the underwriting, investment and credit result models in sections 4-7, putting them together is almost anticlimactic. The problem here is to account for correlation and other higher moment dependencies among the components. As already discussed, the ceded reinsurance part of the credit result is highly correlated with the underwriting result. We have no magic solutions to offer to this highly complex problem. Again we will suggest some simple methods.

The problem is how to construct the distribution function for CR from the models for the underwriting, investment and credit results together with information about their interdependencies. We assume that our modeling of the underwriting, investment and credit results also produces information about the correlation matrix for these results, but that we have no information about higher moment dependencies – it is difficult enough to estimate correlation, much less higher moments. So our information is not sufficient to determine the common overall result distribution, since the information regarding higher moments is incomplete.

In principle, with reasonable assumptions, the distribution of CR can be obtained analytically. In practice however, this is often too complicated, so simulation is used to generate the distribution of CR from the component distributions together with their correlation. Thus we have modeling alternative 1.

(8.1) CR Modeling Alternative 1: Simulation using the component distributions and their correlation

In the following, we will present some alternative approaches to the construction of the distribution function for CR. These methodologies permit analytical constructions. The first is an old actuarial standby.

(8.2) CR Modeling Alternative 2: Fitting a selected parametric distribution via the Method of Moments

The simplest modeling approach is to determine the first two moments of CR from the moments and correlations of the components, and then fit an appropriate parametric distribution via the Method of Moments. This is the methodology used for modeling asset risks via the VaR methodology as described in section 6 and Appendix 13.1.

The missing information with respect to higher moments is filled in via the selection of the distribution type for CR. Since the ratios (percentile minus expectation)/standard deviation strongly depend on the shape of the distribution, the selection of an appropriate type of distribution is most important. The decision of which distribution type to use must be based on the shape and relative importance of the distributions of the components. For the case of a reinsurance company, the underwriting result distribution is almost always heavily skewed. This will most likely cause the distribution of CR to be also skewed as we pictured in (3.7).

(8.3) CR Modeling Alternative 3: Building a hierarchical stochastic compound model based upon common underlying risk factors

A slightly more complicated modeling approach is to build a model based upon common underlying risk factors. This kind of hierarchical stochastic model is called a compound model. In a compound model, deterministic event frequencies are replaced by a stochastic variable.

For example, suppose in the underwriting model that the distribution of the number of losses (or events) N_i within each model component i depends upon a loss frequency parameter λ_i, and the distribution of the each parameter λ_i depends upon a common random variable χ in the following way.

(8.4) $\lambda_i = \lambda_{0,i} \times \chi$ with $E[\lambda_i] = \lambda_{0,i}$ $\forall i$

where χ is a random variable with $E[\chi] = 1$

Then all the component models will be correlated via χ. This approach can of course be extended by introducing a set of common risk factors as outlined in sections 6 and 7 for the case of the investment risk model and the credit risk model, respectively (see also Appendix 13.1). The loss frequency of each component model is then given by a linear combination of the fundamental risk factors.

Clearly, this type of modeling can be extended to all three components of CR simultaneously in order to build a compound model for the distribution of CR.

(8.6) CR Modeling Alternative 4: Constructing Loss Frequency Curves

Our last approach is especially useful for modeling correlations for major cat events, whether they cause underwriting, investment or credit losses or any combination of these types of loss. Since RAC is determined by major cat loss events of whatever type, we can use this modeling approach to determine the left hand tails of the results distributions via the right hand tails of the potential loss event distributions. In this modeling approach, the different model components can be interpreted and related to observable quantities. However, this model is not restricted to catastrophe-like processes; it can be applied in many cases where the stochastic process is described by a compound model [ref. 8.1].

The problem with the use of this model will be the difficulty of transforming the annual aggregate models for investment and credit risk into the framework of compound models with event frequencies and severities. If this is possible, then this model can be used to construct the distribution of CR.

In a compound model, the number of loss events and the individual event loss severity are modeled separately. For the loss frequency curve model, instead of representing the loss frequency λ_0 and the loss severity F(x) independently, we can use the loss frequency curve (LFC) for a joint representation.

(8.7) Loss frequency curve (LFC)

Frequency: $E[N] = \lambda_0$ and $Var[N] = Q \times \lambda_0$

Severity: cdf $F(X)$

LFC: $\lambda(x) = \lambda_0 \times (1 - F(x))$

Inverse LFC: $\lambda^{-1}(\lambda_0 \times y) = F^{-1}(1 - y)$ for $0 < y \leq 1$

Since $F(x)$ is the probability that the severity of an loss X is less than or equal to x, then $\lambda(x)$ is the expected number of losses above the threshold x. Conversely, given an LFC $\lambda(x)$, then the corresponding loss frequency $\lambda_0 = \lambda(0)$ and the loss severity distribution $F(x) = 1 - \lambda(x)/\lambda(0)$ are easily constructed.

We introduce three fundamental operations on LFC's.

(8.8) Operations on LFC's: 1) Aggregation in frequency direction

The joint LFC of two independent subportfolios is given by the sum of the LFC's. In this case, the expected number of losses above a certain threshold x is simply the sum of the frequencies:

$\lambda(x) = \lambda_1(x) + \lambda_2(x)$

In the case of Poisson loss frequency, the aggregation in the frequency direction is equivalent to convolution. Instead of using using Panjer recursion to calculate the distributions of the component distributions, we can simply add the LFC's and "Panjerize" the aggregated LFC. In the non-Poisson case, the situation is more complicated. If the two LFC's may be calculated via Panjer recursion, but $Q_i \neq 1$, then the joint distribution will generally not belong to the Panjer class. However, it is still possible to find a representative from the Panjer class fitting the first two moments.

(8.9) Operations on LFC's: 2) Aggregation in loss direction

The strongest correlation between two subportfolios is obtained in the case of comonotonicity. Then, the losses are deterministic functions (given by the inverse LFC's) of a single random variable. In this case, the LFC's are aggregated in loss direction:

$$\lambda^{-1}(y) = \lambda_1^{-1}(y) + \lambda_2^{-1}(y) \qquad \text{for } 0 < y \leq \lambda_{1,0} = \lambda_{2,0}$$

Comonotonicity means that two portfolios are affected by the same events, i.e., $Q_1 = Q_2$ by definition. If the expected numbers of losses are different in the two subportfolios, we can always add an appropriate number of "zero"-losses to the portfolio with the smaller frequency in order to achieve $\lambda_{1,0} = \lambda_{2,0}$.

(8.10) Operations on LFC's: 3) Frequency split

This is the opposite operation of aggregation in frequency direction. If the losses observed in a subportfolio can be separated into two independent classes, then we can construct LFC's for each class in such a way that the sum of these LFC's equals the total LFC.

We have now seen how to aggregate independent and fully dependent (comonotonic) distributions. With the help of the three fundamental operations, we can construct any correlation between two subportfolios as follows.

(8.11) Combining two correlated subportfolios

- Split the LFC of each sub-portfolio according to (8.10) into a “local” noncorrelated component and a “global” correlated component.
- Aggregate the two subportfolio “global” components in loss direction according to (8.9) to obtain an overall “global” component.
- Aggregate the overall “global” component and the two “local” components in frequency direction according to (8.8) to obtain the combined LFC.
- Calculate the overall loss frequency and loss severity from the overall LFC, and calculate the overall loss frequency variance. The overall aggregate loss distribution is then obtained via Panjer recursion.

The model thus distinguishes between events affecting one of the subportfolios only (“local” events) and events affecting both subportfolios at the same time (“global” events).

For the case of two subportfolios, it is sufficient to introduce one group of “global” events. For modeling the correlation structure between more than two subportfolios, this simple procedure can be generalized by the introduction of several “global” event groups.

(8.12) Example 1: Excess layers

The concept of comonotonicity is best known to reinsurers writing shares of several layers of a nonproportional reinsurance program. The loss amounts ceded to different layers are deterministic functions of each ground up loss. The overall loss to be paid by the reinsurer is the sum of all these ceded layer losses. The reinsurer's LFC is obtained by aggregation in loss direction. Losses affecting only the first layer are "zero" losses for the second layer, while losses affecting the second layer are total losses for the first layer.

(8.13) Example 2: Natural catastrophes

In the case of a natural catastrophe, each insurance company operating in the affected area will suffer a loss which is very nearly determined by its market share times the total loss. In this case, a good approximation for the market LFC is obtained by the aggregation of the company specific LFC's in loss direction (comonotonicity).

(8.14) Example 3

Windstorms in Europe usually affect not just a single country, but several countries at the same time. The LFC for a reinsurer's continent-wide windstorm exposure thus strongly depends on its exposure in each country (or region). The LFC's per country (region) can be obtained by scaling the corresponding normalized market LFC with the reinsurer's exposure in the country. The overall LFC can then be constructed as follows:

- Determine the specific LFC's for all relevant countries: e.g., for France, the United Kingdom, Benelux, Germany, Denmark etc.

- Define a set of event-groups: e.g., “Europe” (all countries); “North” (countries in northern Europe); “South” (countries in southern Europe); “Continent” (Europe without United Kingdom).
- The LFC’s per country are split into LFC’s for “local” events and LFC’s for each event group. The number of events per group can be derived from the analyses of the wind fields of historical storms.
- All LFC’s belonging to the same event group are aggregated in loss direction (comonotonicity) in order to obtain the overall LFC per event group.
- The overall LFC’s per event group and all “local” LFC’s are aggregated in frequency direction to obtain the overall LFC.

The model introduced above is best suited for modeling natural catastrophes where the assumption of comonotonicity is best fulfilled. However, clearly it can also be applied to the modeling of man-made threats.

However, as we discussed above, in order to use this model for calculating the distribution of CR, the problem will be the modeling of the investment and credit results via compound models.

9. How To Allocate RAC To Line, Product, etc.

Once a company has estimated its overall RAC, it is useful to allocate it down to variously defined risk subportfolios (line, product, profit center, etc.) of the overall risk portfolio for the following reasons.

(9.1) Reasons to allocate RAC to risk subportfolio

- Measure performance on a risk-adjusted and consistent basis
- Determine risk-adjusted profit margins for insurance product pricing
- Evaluate alternate business strategies on a risk-adjusted basis

These three items are not independent of each other. To encourage consistent decision-making, it is desirable to have a consistent business measurement and evaluation structure. One way of accomplishing this is to allocate RAC (or a similar risk measure) down to individual business units: profit centers, lines of business, products and even individual contracts. In this section, we discuss methods for doing so.

There is no general answer to the question of how RAC should be allocated. We will discuss various criteria which might be desirable for an allocation method, and then evaluate various suggested allocation methods according to these criteria. In fact, allocation of RAC is not necessary to accomplish (9.1). But we will see that a RAC allocation defines a coherent structure which can make it easier to see what is going on and encourage consistent decision-making.

The RAC allocation is a virtual allocation, not a physical or legal allocation, just as RAC itself is a virtual internal “required” capital to support risk. But since RAC is defined in relation to the determination of available risk bearing capital, RBC, different corporate structures may give rise to different overall RAC, and may lead to different allocations.

One benefit of a RAC allocation is that it puts us in a position to allocate an overall corporate return on equity (RoE) target to risk subportfolios in a risk-adjusted and consistent manner. However, since the overall RAC may be very different from the equity used in the RoE denominator (usually one of the publicly-perceived capitals in (1.1)), the RoE target must be transformed into a RoRAC (Return on RAC) target.

If we can properly allocate down the RAC and RoRAC target to risk subportfolios, it is clear we have the first part of (9.1); we can measure subportfolio performance on a risk-adjusted and consistent basis. A tool that compares actual results with target results is a powerful instrument in the hands of controllers and management. But management action is then required to affect the future business mix and future results of the company. Therefore, management must decide how to steer the company with the help of this tool, as we shall see in section 10. Only after the overall business goals (what should be optimized) and the steering process are clearly defined, can we decide how to allocate the overall targets to risk subportfolios in order to best support the steering process. Management must also decide how to translate the allocation of RoRAC into risk-adjusted profit margins for insurance product pricing in order to meet corporate objectives.

You might ask: Why not accomplish the RAC allocation by calculating the standalone RAC for each subportfolio? The problem with this method is that it ignores any diversification benefits created by combining the various risk elements. Diversification here means that the RAC for the overall risk portfolio is less than the sum of the RACs of the individual subportfolios. This is true for our simple RAC in section 3 based upon a percentile (unless the subportfolios are completely stochastically dependent) because the results distribution of the portfolio is more compressed in a relative sense than are the results distributions of the various subportfolios. Theoretically, the percentiles of the subportfolios are not necessarily sub-additive, but in practice they usually are. The diversification effect is greater the more the subportfolios are stochastically independent. The various subportfolios differ in the degree to which they increase the diversification of the overall portfolio. Thus it is important to see how each subportfolio fits into the total; its contribution to the overall RAC may be less (or more) than is indicated by its standalone RAC.

Another question arises: If we have an overall RAC model, can we use a different model for the RAC allocation? Our simple RAC calculation in section 3 was based upon a percentile of the company's calendar year result. If we use a different method for allocating RAC, we must check whether or not the overall goals of the company can still be met.

We turn now from talking about the particular, simple RAC calculation method described in section 3 to the general problem of allocating any overall RAC down to subportfolio.

Before deciding which specific method should be used for RAC allocation, we should want to define various properties that might be fulfilled by any allocation method.

Some desired properties can be derived from mathematical principles; others depend on economic, organizational or practical considerations. The final selection of properties that are most important is possible only in the context of management's overall goals and their steering process.

The following list of properties of allocation methods does not represent a set of consistent axioms. Instead, the list describes various criteria which we might want the RAC allocation to satisfy. The list is not complete or mutually exclusive: some properties exclude each other and some are already contained in others.

(9.2) Properties of allocation methods

- Risk adjusted: The RAC allocated to a subportfolio accounts for the riskiness of the subportfolio, as seen from the company's overall perspective.
- Partiality: If two subportfolios share exactly the same risk elements (contracts, equities, etc.), X% for the first, Y% for the second, then the RAC allocated to the first and second is exactly in the proportion X to Y.
- Linearity: The RAC allocation is additive; this also means that the RAC allocated to a particular subportfolio doesn't depend upon which larger subportfolio it is contained in. If this criterion were not fulfilled, there would be arbitrage opportunities within the company.
- Account for diversification and dependency: The RAC allocation accounts for dependencies among risk subportfolios as well as independence.
- Organisational independence: The RAC allocated to a particular subportfolio does not depend upon its particular organizational structure nor upon the organisational structure of the rest of the company (this is closely related to linearity).
- Measure dependent: The RAC allocation is the inverse of the aggregation process used to determine the overall RAC. This criterion is fulfilled if the RAC allocated to each subportfolio is based upon the same underlying risk

model as is used for the overall portfolio, and if it also takes into account dependencies and diversification in an appropriate way.

- Based upon reliable information: The RAC allocation is based on available and reliable information as far as possible.
- Practicality: It is possible to implement numerically in such a way that the RAC allocation is stable and reliable, and can be obtained within a reasonable time and with justifiable effort.
- Reflect internal risk perception: The RAC allocation is based on internal risk perception only. This may be in contrast to the overall RAC, which might be composed of internal (percentile) and external (RTL) risk perception, as in our simple case.
- Ex post additivity: The sum of the allocated subportfolio RAC (on the basis of ex ante information) is equal to the overall ex post overall RAC for the entire portfolio.
- Consistency with statutory requirements: The RAC allocation takes statutory requirements into account.

This is quite a list of possible properties. No single RAC allocation method can satisfy all of them.

When discussing allocation methods, we would first like to distinguish between three different classes. The first class consists of those allocation methods that are based on "local" risk measures. The overall RAC is allocated in proportion to the individual riskiness of each risk subportfolio as measured by the local risk measure without taking into account dependencies among the subportfolios and diversification.

(9.3) Examples of “Local” risk measures

- Standard deviation
- Variance
- Percentiles
- Shortfall risk (conditional expected excess loss)

Using local risk measures without considering the actual dependencies between subportfolios does not mean that such dependencies are irrelevant for allocation. For example, the standard deviation principle assumes full correlation between all subportfolios; the variance principle assumes independence; the percentile and the shortfall principles assume comonotonicity.

The second class consists of those allocation methods that are based on some kind of volume measure, such as solvency ratios. These methods play a special role since the local volume measure (as used by the first class of measures) equals the contribution to the overall volume (property of the methods belonging to the third class). This is due to the fact that expected values are additive and cannot be diversified.

The third class consists of those allocation methods which are based upon the contribution of each subportfolio risk to the overall “global” risk. The overall RAC is allocated in proportion to the contribution of each subportfolio to the overall risk. Thus, these methods take into account dependencies as well as diversification (as far as modeled).

(9.4) Examples of “Global” risk measures

- Marginal principle (variation, or “with and without” as defined in Appendix 13.2)
- “Euler”-principle (special case of marginal principle – to be described below and in Appendix 13.2)
- Covariance (“Euler” for standard deviation)
- Higher co-moments (“Euler” for higher moments)
- Linear combinations of above

A RAC allocation based upon the marginal principle (as defined in Appendix 13.2 and not to be confused with the similar “marginal change allocation” as defined below) or on the Euler principle depends on the definition of an overall risk measure, since these methods quantify the contribution of each subportfolio to the overall risk. So, let us define risk measure.

(9.5) Risk Measure: A risk measure ρ is a mapping from the space of volume measures of subportfolios of a risk portfolio into the positive real numbers.

We assume that there exists a volume measure V_j associated with each subportfolio R_j. For any subportfolio R_j, $\rho(V_j)$ is the RAC allocated to R_j.

The Euler principle allocates RAC by reversing the aggregation process. This is achieved by answering the following question: What is the contribution of each subportfolio to the overall RAC? The answer is of course trivial for those components in the RAC formula which are strictly additive - as is the RTL in section 3, since it is based on solvency ratios. But what is the contribution of those RAC components which depend on diversification, such as percentiles? Fortunately, there exists a very general and conceptually easy solution to this problem. Since the basic mathematical principle behind it was found by the famous Leonhard Euler more than two centuries ago, we call it the "Euler" principle.

(9.6) Euler Principle: $\rho(V_j) = \{(\partial/\partial V_j)\rho(V)\} \times V_j$

where ρ is a homogeneous risk measure differentiable on the space of volume measures of subportfolios R_i of the portfolio R,

$\{R_1, R_2, \ldots R_n\}$ is any partition of R, and

V is the volume measure of R.

There is a more thorough development of this concept in Appendix 13.2.

The marginal principle is an approximation for "Euler". The covariance principle is a special case of the "Euler" principle (obtained for the case where the overall risk is defined by the standard deviation, as shown in Appendix 13.2).

The following special capital allocation method, also belongs to class 3, and is similar but not identical to the marginal principle "with and without".

(9.7) “Marginal Change Allocation”: The RAC allocated to a subportfolio is given by the marginal change of the overall RAC as a consequence of adding the subportfolio to the pre-existing portfolio excluding this subportfolio.

This allocation method has the feature that the overall RAC always equals the sum of the RAC’s allocated to the subportfolios in the portfolio (ex post additivity). However, it has a very serious shortcoming: the allocated RAC depends upon the ordering of the risks; the diversification between a new subportfolio and the existing portfolio is fully attributed to the new subportfolio. Thus there is no clear rule on how to treat existing risks in a mature portfolio. Also, the implementation of this method requires a very powerful and fully integrated online risk modeling tool.

In the following table, we comment upon the properties that one might want fulfilled by an allocation method. We omit the property “consistency with statutory requirements”, since this must be determined locally.

(9.8) Properties of allocation methods

	“Local” risk measures (class 1)	**“Volume” measures** (class 2)	**“Global” risk measures** (class 3)
Risk adjusted Consider contribution to overall risk	**not fulfilled** Considers only full correlation between subportfolios or no correlation at all	**not fulfilled** Volume alone does not contain information about riskiness	**fulfilled** Subportfolio risk is measured as part of the overall risk
Partiality Allocate the same RAC for the same risk	**partly fulfilled** OK for most methods, but, e.g., not for variance!	**fulfilled**	**fulfilled** In some cases (e.g., marginal principle) not exactly
Linearity RAC allocated is the sum of the individual RAC’s	**not fulfilled**	**fulfilled**	**fulfilled** In some cases (e.g., marginal principle) not exactly
Dependency Dependencies between risk are considered	**not fulfilled**	**fulfilled** Volume cannot be diversified	**fulfilled**
Organisational independence Allocated RAC does not depend on the structure of the company	**not fulfilled**	**fulfilled**	**fulfilled** In some cases (e.g., marginal principle) not exactly

	“Local” risk measures (class 1)	**“Volume” measures** (class 2)	**“Global” risk measures** (class 3)
Measure dependent Reflects the overall risk measure	**not fulfilled**	**fulfilled** If overall risk is defined by volume	**fulfilled**
Information Based on all reliable information	**partly fulfilled** Information regarding dependencies is not considered	**fulfilled** If overall risk is defined by volume	**partly fulfilled** Higher co-moments (for which only little information is available) may be important
Practicality	**fulfilled**	**fulfilled**	**partly fulfilled** OK for covariance, difficult for others
Reflect internal risk perception Be derived from overall distribution and not from volume	**fulfilled**	**not fulfilled**	**partly fulfilled** OK for higher co-moments, but not necessarily for “Euler” and “marginal”
Ex post additivity The sum of the RAC allocated to business units always equals the overall RAC	**not fulfilled**	**fulfilled** The expected value cannot be diversified	**partly fulfilled** Exactly fulfilled by one special method, almost fulfilled by other methods

10. Managing RAC To Optimize Risk and Return

As you may recall from section 3, we defined RAC to meet certain criteria, which we repeat here for convenience.

(10.1) Criteria for RAC

- It must meet specified management risk and survival criteria.
- It should quantify the risk/return trade-off for all risk exposures.
- It must be useful for making appropriate risk-based business decisions.

The first criterion is part of the definition of RAC in section 3. In sections 4-9, we discussed the quantification of the risk/return trade-off via the estimation of RAC and RoRAC for each segment of an insurance company's risk portfolio and the calculation of the company's underwriting, investment and credit results. In this section, we will deal with the use of RAC and RoRAC for making business decisions about the company's underwriting, investment and credit portfolios. Among the many goals one might have for the use of RAC and RoRAC to optimize risk and return are the following.

(10.2) Some goals

- Achieve capital efficiency
- Maximize shareholder value
- Achieve a target RoE

Two particular strategic actions we will explore are the following.

(10.3) Two possible strategic actions

- Move RAC or RBC toward equality
- Maximize RoE with respect to management's risk tolerance rule

Let us consider the first strategic action.

(10.4) If RAC $>$ RBC, then move toward RAC = RBC by one or more of the following tactical actions:

- Decrease underwriting risk
- Cede more risk to (secure) reinsurers
- Decrease investment risk
- Decrease credit risk
  - buy more secure reinsurance (buy diversification)
- Divest non-core and RAC-intensive operations
- Raise new capital

This is the case where the available capital, RBC, is not enough to properly support the company's risk portfolio and also meet management's risk tolerance rule. Thus the company's capital is exposed to more risk than desired. The first five tactical actions decrease RAC; the last increases RBC.

Note that buying reinsurance decreases underwriting and also investment risk (if it is not placed on a funds-withheld basis). As discussed earlier, reinsurance transforms an underwriting risk, and usually also an investment risk, into a credit risk. The credit risk is the possible non-payment or decreased payment of ceded claims caused by a financial default by the reinsurer. The cost of this credit risk must be thought of as part of the cost of reinsurance. If the reinsurance is secure, the company's overall risk level decreases.

The sum of the company's and reinsurer's RAC may be less after the placement of reinsurance. This can occur if the reinsurer is better diversified than the ceding company, or if the assumed reinsurance is less correlated with the remainder of the reinsurer's portfolio than it is with the ceding company's. In these cases, the risk transfer causes the ceding company's RAC to decrease more than the reinsurer's RAC increases, if their management risk tolerance rules are the same.

An international reinsurer may be better diversified than a ceding company in the following ways.

(10.5) Well-diversified reinsurers typically have better diversification by

- Number of insureds (smaller shares of many)
- Types of risks (a broader underwriting portfolio)
- Concentration of risk, e.g., geographically

Let us return to the first strategic action of moving RAC or RBC toward equality. The alternate case is as follows.

(10.6) If RAC $<$ RBC, then move toward RAC = RBC by one or more of the following tactical actions:

- Increase underwriting risk

  – expand into new business

  – write higher risk business

  – develop new products

  – cede less reinsurance

- Increase investment risk

- Increase credit risk

- Buy another company

- Return capital to shareholders

This is the case where the available capital, RBC, is more than necessary to support the company's risk portfolio and also meet management's risk tolerance rule. Thus the company's capital is not being fully utilized to maximize the RoE with respect to risk level. The first four tactical actions increase RAC; the last two decrease RBC.

Now let us explore the second strategic action of maximizing RoE with respect to management's risk tolerance rule. The conditions to do so are clear.

(10.7) Maximize RoE with respect to management's risk tolerance rule

Simultaneously achieve the following:

- Achieve RAC = RBC
- Maximize RoRAC

RoRAC, as we have seen, is simply the company's result divided by the company's RAC.

(10.8) RoRAC = CR / RAC

If RAC = RBC, and RoRAC is maximized, then it is clear that the RoE is maximized in a manner satisfying management's risk tolerance rule.

There may be, and probably is, more than one way to achieve (10.7). What we will describe here is an iterative process. Our starting base case is the assumption that the company's current underwriting, investment and credit risk portfolios will continue through next year.

(10.9) Base Case: Current risk portfolio

Risk portfolio components indexed via j = 1, 2, . . . , n.

$CR_0[j]$ = rv for next year's result for component j

$CR_0 = \Sigma_j CR_0[j]$ = overall result for base portfolio

RBC_0 = RBC for base portfolio calculated via section 2

RTL_0 = risk tolerance level for base portfolio

RAC_0 = RAC for base portfolio calculated via (3.9)

$E[RoRAC_0]$ = $E[CR_0] / RAC_0$ via (10.8)

$RAC_0[j]$ = standalone RAC for component j

$ARAC_0[j]$ = overall RAC_0 allocated to component j via section 9

We need the distinction between standalone RAC and allocated RAC.

Now compare the expected return $E[CR_0[j]]/ ARAC_0[j]$ for each component to $E[RoRAC_0]$. Let us first consider the case where overall $RAC_0 < RBC_0$. Let us assume that management has no wish to decrease RBC, but instead wishes to increase RAC. The first possibility for component j is that its expected return divided by allocated RAC exceeds the overall expected RoRAC. Component j is thus the kind of assumed risk we would like to have more of.

(10.10) Case 1: $RAC_0 < RBC_0$

and $E[CR_0[j]] / ARAC_0[j] > E[RoRAC_0]$

Then increase RAC[j] by increasing the risk level of the jth component.

The risk level can be increased by one of the tactical actions listed in (10.6). This will also change E[CR[j]] and ARAC[j], and we should not assume that the changes will be linearly related to the change in RAC[j]. Note that a change in the risk portfolio will also change RBC.

For example, let us assume that component j is an underwriting component. Thus the base case defines a certain level of pure premium $PP_0[j]$ (premium minus all external (commissions, brokerage fees, etc.) and internal expenses). The expected result $E[CR_0[j]]$ is simply the expected underwriting return (including risk-free investment income). This relates to the expected underwriting economic margin $EUEM_0[j]$ as a percent of pure premium as follows.

(10.11) Expected underwriting economic margin

$$EUEM_0[j] = E[CR_0[j]] / PP_0[j] \qquad \text{(definition)}$$

Likewise we can define a target pricing margin, or target underwriting economic margin, $TUEM_0[j]$, and then relate it to the target result, $TCR_0[j]$, for component j.

Thus it is clear that, given any assumed risk portfolio, any designation of an overall target RoE and any allocation of RAC can be translated into a target company return and thus into unique target economic underwriting margins for each risk component j. And vice versa, any set of target economic underwriting margins for each risk component j for an assumed risk portfolio translates back into an overall RoE and exactly one RAC allocation.

It should also be clear that for an underwriting component, as RAC[j] increases by increasing the volume and type of business being written, E[CR[j]] cannot increase indefinitely due to market constraints. At some point, E[CR[j]] / ARAC[j] reaches a maximum and begins to decrease. Thinking about this for a minute should convince you, because the fact that volume is increasing in a competitive market means that underwriting standards are decreasing, so that lower profit business is being assumed. Please note that the pure premium PP[j] at which E[CR[j]] / ARAC[j] reaches a maximum for the jth subportfolio may not be the pure premium at which the overall portfolio return is maximized.

Let us deal with the second possibility for component j, where its expected return divided by allocated RAC is less than the overall expected RoRAC, also in case 1 where overall $RAC_0 < RBC_0$. This is the kind of risk assumption we would like to have less of.

(10.12) Case 1: $RAC_0 < RBC_0$

and $E[CR_0[j]] / ARAC_0[j] < E[RoRAC_0]$

Then decrease RAC[j] by decreasing the risk level of the jth component.

The risk level can be decreased by one of the tactical actions listed in (10.4). Again, as noted above, this will also change E[CR[j]], ARAC[j] and RBC. We want to feed some of the ARAC for component j into components where (10.10) holds, thereby increasing the overall expected RoRAC.

As noted above, it should also be clear that for an underwriting component, as RAC[j] decreases by decreasing the volume and type of business being written, E[CR[j]] / ARAC[j] will tend to increase, because the fact that volume is decreasing in a competitive market usually means that underwriting standards are increasing. Lower profit business is being discarded, so that for the remaining portfolio, the expected underwriting economic margin in (10.11), EUEM[j] = E[CR[j]] / PP[j], increases. Again, this cannot increase indefinitely. As volume shrinks, eventually the ratio of underwriting expenses to pure premium will grow enough to decrease EUEM[j]. Also, as volume shrinks, at some point the company begins to lose underwriting expertise because it is losing experienced underwriters. This will also tend to decrease EUEM[j].

We may think of this problem as one of constrained optimization. Define the function f as follows.

(10.13) $f(X[1], \ldots, X[n]) = \Sigma E[R[j] \mid X[j]]$

where $X[j] = RAC[j]$

and $R[j] = CR[j]$

The problem can then be stated as follows.

(10.14) Maximize f with respect to $\{X[1], \ldots, X[n]\}$

such that $\Sigma A[j] = RBC \mid \{X[1], \ldots, X[n]\}$

where $A[j] = ARAC[j] \mid \{X[1], \ldots, X[n]\}$

Standard techniques can be used to solve this problem. But please remember that CR[j] is not a linear function of RAC[j]. In our age of fast desktop PCs, we can bash this out in a reasonable length of time.

The optimization for Case 2, where overall $RAC_0 > RBC_0$, is similar. Let us assume that management has no wish to increase RBC, but instead wishes to decrease RAC. This can be solved as was Case 1. Clearly, the solution will certainly entail decreasing RAC[j] for components where $E[CR_0[j]] / ARAC_0[j] < E[RoRAC_0]$.

11. Conclusion

We have discussed the reasons why an insurance company should address risk and capital issues in a methodical manner and we have discussed some of the problems encountered doing so. We hope that the reader has discovered some good ideas and some procedures for obtaining useful measurements. We hope that the reader found the discussion of threat scenarios and their estimation particularly useful, since we believe that, whatever RAC measure is used, a more accurate measurement of future underwriting risk is crucial to proper RAC determination. We hope to see less black box modeling of this problem, and more thought put into the estimation of underwriting parameters especially. We hope that someone devises better ways of modeling and estimating investment risk over extended time periods. We hope that the modeling and estimation of credit risk is improved. We hope that the discussion of the various possible properties of capital allocation methods will help improve future discussions of this topic. We also hope that the reader also found the discussion of managing RAC to be interesting and informative.

12. Bibliography

[0.1] Swiss Re, *From Risk To Capital: An Insurance Perspective*, Swiss Re publication, 1999.

[1.1] Feldblum, S., “NAIC Property/Casualty Insurance Company Risk-Based Capital Requirements”, *PCAS* LXXXIII, 1996.

[3.1] Daykin, C.D., Pentikäinen, T. and Pesonen, M., *Practical Risk Theory for Actuaries*, Chapman and Hall, 1994.

[3.2] Klugman, S., Panjer, H. and Wilmot, G., *Loss Models: From Data to Decisions*, John Wiley & Sons, 1998.

[3.3] Embrechts, p., Kluppelberg, C. and Mikosch, T., *Modelling Extremal Events for Insurance and Finance*, Springer-Verlag, 1997.

[3.4] Maginn, J. and Tuttle, D., Eds., *Managin Investment Portfolios: A Dynamic Process*, Warren, Gorham & Lamont, 1990, p. 2-17.

[4.1] Klugman, S., Panjer, H. and Wilmot, G., op. cit.

[4.2] Embrechts, p., Kluppelberg, C. and Mikosch, T., op. cit.

[4.3] Taylor, G., *Claim Reserving in Non-Life Insurance*, North-Holland, 1991.

[4.4] A.M. Best, “Property/Casualty A&E Losses Plunge, But Concerns Remain for Individual Companies”, *BestWeek Special Report: Property/Casualty Edition*, September 21, 1998.

[6.1] J.P. Morgan/Reuters, *RiskMetrics*™ – Technical Document, New York, December 1996, free download from http://www.riskmetrics.com/.

[6.2] A comprehensive alphabetical listing of published value-at-risk papers can be found at http://www.gloriamundi.org/var/pub.html and a list with working papers at http://www.gloriamundi.org/var/wps.html.

[6.3] Franzetti, C., “Value at Risk for Longer Time Horizons”, Swiss Re internal document, 1998.

[7.1] JP Morgan, *CreditMetrics*^{TM} – Technical Document, New York, April 1997, free download from www.JPMorgan.com.

[7.2] Credit Suisse Financial Products, *CreditRisk*$^{+}$*, A Credit Risk Management Framework*, 1997, free download from www.csfp.ch.

[7.3] KMV, *Credit Monitor*, information can be downloaded from www.KMV.com.

[7.4] Wilson, T., *Measuring and Managing Credit Portfolio Risk*, technical document by McKinsey & Co..

[7.5] Wilson, T., “Portfolio Credit Risk” (1), *Risk 10*, 1997, pp.111-117.

[7.6] Wilson, T., “Portfolio Credit Risk” (2), *Risk 10*, 1997, pp. 56-61.

[7.7] S&P *Credit Week*, April 15, 1996.

[8.1] Krieter, F., Bloch, A., Meyer, P., Schmid, E. and Bernegger, S., “Catastrophe Portfolio and Capital Needs of a Reinsurance Company”, *Transactions of the 25th International Congress of Actuaries*, Brussels, 1995.

[9.1] Artzner, P., Delbaen, F., Eber, J-M. and Heath, D., “Thinking Coherently”, *Risk 10*, 1997.

[9.2] Artzner, P., “Application of Coherent Risk Measures to Capital Requirements in Insurance”, *North American Actuarial Journal*, vol. 3, 1999.

[9.3] Christoffersen, P., Diebold, F and Schuermann, T., “Horizon Problems and Extreme Events in Financial Risk Management”, *Economic Policy Review*, Federal Reserve Bank of New York, 1998.

[9.4] Embrechts, P., Resnick, S. and Samorodnitsky, G., “Extreme Value Theory as a Risk Management Tool, *North American Actuarial Journal*, vol. 3, 1999.

13. Appendices

Appendix 13.1. Value at Risk

(13.1.1) Goal

- Model and/or estimate a specific percentile of the distribution of the change in value of an investment portfolio over a given time period.

(13.1.2) Assumptions

- The randomness of the portfolio value stems from the randomness of the values of the component financial instruments, which is induced by random risk factors.

- Most often it is assumed that the changes of risk factors are normally distributed. More precisely, the logarithms of the risk factors are modeled as independent identically multivariate normally distributed increments. This assumption is not questioned here.

- In the simplest form of VaR, the functional dependencies of the instrument prices upon the risk factors is linearized for the non-linear instruments. These are all instruments with some optionality features.

(13.1.3) Definitions

$$
\begin{aligned}
&Q_j(t): Risk\ \ factor,\ \ j = 1,2,...,J \\
&Y_j(t) := log(Q_j(t)) \\
&X_j(t) = \frac{Q_j(t) - Q_j(t-1)}{Q_j(t)} \approx Y_j(t) - Y_j(t-1) = log(Q_j(t)/Q_j(t-1)) \\
&P_k(Q,t): Instrument\ \ price\,(value),\ \ k = 1,2,...,K \\
&V(Q,t) = \textstyle\sum_{k=1}^{K} P_k\ :Total\ \ portfolio\ value \\
&\acute{o}_{ij}(t) = Cov(X_i(t), X_j(t))
\end{aligned}
$$

We assume a portfolio of K instruments which are exposed to J risk factors. Risk factors may be interest rates, stock indices, foreign exchange rates and commodity prices. Some definitions above have been introduced for the sake of simplification.

The linearization yields the following.

(13.1.4) Linearization

$$
\Delta P_k(Q,t) = \sum_{j=1}^{J} \frac{\partial P_k}{\partial Q_j} \Delta Q_j
$$

$$
\Delta V = \sum_{k=1}^{K} \sum_{j=1}^{J} \frac{\partial P_k}{\partial Q_j} \Delta Q_j = \sum_{j=1}^{J} \left(\sum_{k=1}^{K} \frac{\partial P_k}{\partial Q_j} \cdot Q_j \right) \frac{\Delta Q_j}{Q_j} := \sum_{j=1}^{J} \left(w_j \right) \frac{\Delta Q_j}{Q_j} = \mathbf{w}^T \cdot \mathbf{X}
$$

with

$$
\begin{aligned}
&\mathbf{X} = (X_1, X_2, ..., X_J)^T, \\
&\mathbf{w} = (\textstyle\sum_{k=1}^{K} \frac{\partial P_k}{\partial Q_1} Q_1, \sum_{k=1}^{K} \frac{\partial P_k}{\partial Q_2} Q_2, \ldots, \sum_{k=1}^{K} \frac{\partial P_k}{\partial Q_J} Q_J)^T .
\end{aligned}
$$

The **X** vector describes the logarithms of the risk factor changes and **w** is the vector of exposures or sensitivities.

The variance and the standard deviation of the portfolio value change can be calculated as follows.

(13.1.5) Variance and standard deviation

$$Var(\Delta V) = Var(\mathbf{w}^T \mathbf{X}) = \mathbf{w}^T \, Var(\mathbf{X})\mathbf{w} = \mathbf{w}^T \mathbf{R}\mathbf{w}$$
$$\sigma_{\Delta V} = \sqrt{\mathbf{w}^T \mathbf{R}\mathbf{w}}$$

where the matrix **R** is the so-called covariance matrix

$$\mathbf{R}(t) = \begin{bmatrix} \sigma_{11} & \sigma_{11} & . & \sigma_{1J} \\ \sigma_{21} & \sigma_{22} & . & . \\ . & . & . & . \\ \sigma_{J1} & . & . & \sigma_{JJ} \end{bmatrix}.$$

The Value at Risk is defined as a percentile, most often the 95% or 99% percentile, of the assumed normal distribution.

(13.1.6) Value at Risk, VaR, with respect to percentile γ

$$Prob(\Delta V > VaR) = 1 - \gamma .$$

(13.1.7) VaR for the normal distribution

$$\mathrm{Va}R = \alpha_{\gamma} \cdot \sigma_{\Delta V}$$
$$\alpha_{99\%} = 2.326$$

Note: We have omitted the derivative with respect to time. This is defensible on the grounds that its effect do not show up in the VaR calculation due to the deterministic nature of time and because the resulting so-called drift term may be added afterwards.

So far we have not specified the time interval, but it is implicit in the risk factor change.

(13.1.8) Risk factor change

$$X_j(t) = log(Q_j(t)/Q_j(t-1)).$$

Therefore ΔV is based on the same interval as the change in risk factors. Let us assume it to be one day. Under the so-called OLS conditions, VaR based upon one day can be scaled to k times one day by multiplying with the square root of k. First, let us define OLS conditions.

(13.1.9) OLS conditions:

- Equal expectations: $E(Y_j(t)) = \text{ì}$ for all t
- Equal variances: $Var(Y_j(t)) = \text{ó}$ for all t, and
- No serial correlation: $Cov(Y_j(t), Y_j(s)) = 0$, $s \neq t$

Under these conditions, the following holds.

(13.1.10) OLS conditions imply:

$$Var(Y_j(t+k)-Y_j(t))=k\cdot Var(Y_j(t+1)-Y_j(t))$$

and

$$Cov(\{Y_j(t+k)-Y_j(t)\},\{Y_i(t+k)-Y_i(t)\})=k\cdot Var(\{Y_j(t+1)-Y_j(t)\},\{Y_i(t+1)-Y_i(t)\}).$$

Because all variances and covariances are scaled with k, the portfolio variance is also scaled by the same factor. It follows immediately that a Value at Risk based on one day can be scaled to k times one day by multiplying with the square root of k.

(13.1.11) VaR for k days

$$VaR_{k\ Days}=VaR_{1Day}\cdot\sqrt{k}\ .$$

For determining the capital adequacy of the company, for planning and allocation purposes and for making risks comparable with other risk sources, e.g., credit and insurance or underwriting, we are interested in a risk measure that has the following form.

(13.1.12) Risk Measure Form

$$Prob\,(\text{Value in one year - Expected value in one year} > \text{LVaR}) = \text{p}$$

LVaR is a percentile to be found. The main difficulty consists in estimating realistically the distribution of the value of the portfolio in one year given both exogenous uncertainties, e.g., risk factor changes, and endogenous actions, especially the strategies of the portfolio managers. Therefore, more precisely we are looking for a probability distribution which is contingent upon the strategy, call it θ.

(13.1.13) Risk Measure contingent upon strategy θ:

$$P\text{rob}\,(\text{Value in one year - Expected value in one year} > \text{LVaR}\,|\,\theta) = \text{p}$$

Appendix 13.2 The Marginal and the "Euler" Allocation Principles

The determination of the overall RAC C is based upon a set of volume measures V_i (premiums, liabilities, number of shares, bonds, etc.) which represent the portfolio and upon models for the various risk factors (natural catastrophes, man-made catastrophes, interest rates, etc.). For a given risk model, C can be represented as a function ρ, called a risk measure, of the volume measures V_i.

(13.2.1) RAC as a function of volume

$$\rho\,(V_1\,, V_2\,, \ldots, V_N\,) = \rho\,(\overline{V}) = C$$

If the volume V_i of unit i is slightly modified, i.e., if V_i is replaced by $V_i + \Delta V_i$, then the outcome of the RAC formula will also differ slightly from C.

(13.2.2) Change in RAC as volume changes

$$\rho\,(V_1\,, V_2\,, \ldots, V_{i-1}\,, V_i + \Delta V_i\,, V_{i+1}\,, \ldots, V_N\,) = C + \Delta C_i$$

The ratio $\Delta C_i / \Delta V_i$ measures the sensitivity of the overall RAC, C, with respect to the risks belonging to unit i. The higher this ratio, the stronger the overall RAC depends upon the risks of unit i. Of course, the contribution also depends upon the actual volume V_i of the risks. If the RAC allocated to unit I, C_i, should depend upon the contribution of the unit to the overall risk, then it is very natural to allocate it in proportion to the sensitivity and the volume.

(13.2.3) Marginal Allocation: RAC allocation in proportion to sensitivity and volume

$$C_i \propto \frac{\Delta C_i}{\Delta V_i} V_i \qquad i.e. \quad C_i = C \frac{\frac{\Delta C_i}{\Delta V_i} V_i}{\sum_{j=1}^{N} \frac{\Delta C_j}{\Delta V_j} V_j}$$

The first term and the denominator are required for the purpose of normalization, i.e., in order to make sure that the sum of the allocated RAC's, C_i, equals the total RAC, C. The marginal principle is additive for the case of infinitesimal volume changes (see Euler principle below). For the case of non infinitesimal changes (e.g., for "with and without" as defined below) it is only almost additive as demonstrated for following example.

(13.2.4) Example

$$\begin{aligned}
&V^* = V_1 + V_2 \\
&\Delta V^* = a \cdot V^* = a \cdot V_1 + a \cdot V_2 = \Delta V_1 + \Delta V_2 \\
&\Delta C^* \approx \frac{\partial C}{\partial V_1} \Delta V_1 + \frac{\partial C}{\partial V_2} \Delta V_2 = a \cdot \left(\frac{\partial C}{\partial V_1} V_1 + \frac{\partial C}{\partial V_2} V_2 \right) \\
&\Rightarrow \frac{\Delta C^*}{\Delta V^*} V^* = \frac{\Delta C^*}{a} \approx \frac{\partial C}{\partial V_1} V_1 + \frac{\partial C}{\partial V_2} V_2 \approx \frac{\Delta C_1}{\Delta V_1} V_1 + \frac{\Delta C_2}{\Delta V_2} V_2
\end{aligned}$$

The risk measure ρ is homogeneous when the following condition holds: If all volumes are multiplied by the same factor λ, then the overall RAC will also change by the factor λ.

(13.2.5) Homogeneity of ρ

$$\rho(\lambda \cdot V_1, \lambda \cdot V_2, \ldots, \lambda \cdot V_N) = \rho(\lambda \cdot \overline{V}) = \lambda \cdot C$$

If ρ is homogeneous, then the following relation is obtained by taking the derivative on both sides.

(13.2.6)

$$\frac{\partial}{\partial \lambda} \rho\left(\lambda \cdot \overline{V}\right) = \sum_{i=1}^{N} \left. \frac{\partial \rho\left(\overline{U}\right)}{\partial U_i} \right|_{U_i = \lambda V_i} \cdot V_i \qquad \forall \lambda$$

$$\frac{\partial}{\partial \lambda}\left(\lambda \cdot C\right) = C = \rho\left(\overline{V}\right)$$

$$\Rightarrow \quad \sum_{i=1}^{N} \frac{\partial \rho\left(\overline{V}\right)}{\partial V_i} V_i = \rho\left(\overline{V}\right) = C$$

The “Euler” principle splits the overall RAC into contributions from the individual components. It is therefore natural to allocate the RAC in the same way.

(13.2.7) The Euler principle allocation

$$C_i = \frac{\partial \rho\left(\overline{V}\right)}{\partial V_i} V_i$$

The derivatives represent the above sensitivity (for the case of infinitesimal variations) and thus the Euler principle is a special case of the marginal principle. It can therefore be easily interpreted and it has the nice property that it is already normalized.

(13.2.8) Example

Let’s assume that the risk measure ρ is defined as a multiple of the overall standard deviation. In this case, the Euler-principle is equivalent to the covariance principle:

$$\rho = k\sigma V = k\sqrt{\sum_i (\sigma_i V_i)^2 + \sum_{i,j:j\neq i} \rho_{ij}\left(\sigma_i V_i \sigma_j V_j\right)}$$

$$\frac{\partial \rho}{\partial V_i} = k\,\frac{2\sigma_i^2 V_i + 2\sigma_i \sum_{j:j\neq i} \rho_{ij}\sigma_j V_j}{2\sqrt{\sum_i (\sigma_i V_i)^2 + \sum_{i,j:j\neq i} \rho_{ij}\left(\sigma_i V_i \sigma_j V_j\right)}}$$

$$\sum_i \frac{\partial \rho}{\partial V_i} V_i = k\,\frac{\sum_i \left(\sigma_i^2 V_i + \sigma_i \sum_{j:j\neq i} \rho_{ij}\sigma_j V_j\right) V_i}{\sqrt{\sum_i (\sigma_i V_i)^2 + \sum_{i,j:j\neq i} \rho_{ij}\left(\sigma_i V_i \sigma_j V_j\right)}} = k\sigma V\,\frac{\sum_i \operatorname{cov}\left[X_i, X\right]}{\operatorname{var}\left[X\right]} = \rho$$

Homogeneity is one of the most important properties to be fulfilled by coherent risk measures as defined by Artzner, Delbean, Eber and Heath [ref. 9.2]. It simply says that if one blows up (or down) all volume measures by the same factor, the overall RAC will also increase (decrease) by the same factor. Apart from capital measures expressed in absolute and not relative terms, such as minimum statutory capital requirements for example, all relevant risk measures used in practice (percentages of volume, standard deviation, percentile, shortfall, etc.) are homogeneous. From this perspective, the Euler principle would be the natural candidate for RAC allocation. However, if certain RAC components like the RTL are defined on the basis of external risk perception, the Euler principle will also determine the contribution of a unit to this component. That is, RAC allocation would then be at least partly based on external risk perception, which is not necessarily in line with the objectives of RAC based business decision-making. Of course, the Euler principle cannot only be applied to the overall RAC, but to any homogeneous risk measure. It will therefore have to be considered as an allocation candidate even in the case where different risk measures are used for RAC determination and RAC allocation.

(13.2.9) Examples of allocation via the Euler principle

- If the overall RAC (or one of its components as, e.g., the RTL) is given as the sum of weighted volumes, then there is no diversification and the allocated RAC of unit i equals the standalone RAC of unit i.
- If the overall RAC is given as a function of the standard deviation of the overall result distribution and the if the homogeneity principle is fulfilled, then the RAC can only be a multiple of the standard deviation! In this case, the Euler principle corresponds to the covariance principle as shown above.
- If the overall RAC is derived from percentiles, shortfall risk or similar risk measures, it can be represented as a function of all or several moments of the overall distribution. According to the Euler principle, the contribution of the individual components can be represented as a linear combination of co-moments.

Another special case of the marginal principle is the “with and without” principle. Here, the overall RAC is evaluated with and without the risk of unit i, and the RAC is then allocated in proportion to the RAC reductions. This can be achieved by setting the volume change ΔV_i equal to $-V_i$ in the above formula for the marginal RAC allocation.