SOCIETY OF ACTUARIES

Article from:

Risks and Rewards Newsletter

July 2001 – Issue No. 37

RISKS and REWARDS

The Newsletter of the Investment Section of the Society of Actuaries

NUMBER 37 JULY 2001

Exploring C1 Risk

by Thomas Merfeld

Editor's Note: This is part one of a two-part article. The second part will run in the next issue of Risks and Rewards.

Most of us consider insurance companies to be expert risk managers. One of these risks reflects the possibility that their investments perform poorly. We call this C1 risk.

I've spent years wondering how to articulate the possibility that investments perform poorly. Is an investment that you mark-to-market on the statutory filing riskier than if you could hold it at historical cost? Are private placements riskier because they don't enjoy a ready market? Are derivatives risky? Are stocks riskier in the short run than over long investment horizons? How do you isolate C1 risk from C3 risk? What is a sufficient asset reserve? Should product managers care if returns fall short of pricing assumptions? Are bond defaults worse than other causes of bond value declines? Should a P&C company own commercial mortgages? How much risk is enough? Does the character of return—income versus capital appreciation—matter? Should stocks back reserves? How bad can things get?

(continued on page 4)

CIA Task Force on Segregated Fund Investment Guarantees

excerpt from the Canadian Institute of Actuaries

Editor's Note: The CIA Task Force on Segregated Fund Investment Guarantees was founded in 1999 and charged with developing recommended approaches for the use of stochastic techniques to measure the obligations created by segregated fund investment guarantees (i.e., where an underlying level of investment performance is guaranteed by an insurer). The Task Force issued a 64-page report in August 2000 and recommended that Canadian actuaries use stochastic techniques to establish liabilities for these guarantees. The following passage on investment return models is excerpted from Section 2 of the report, and should be of particular interest to readers of this newsletter. The full report is available at the CIA Web site as accession number 20020. Also see the announcement on page 34 of this issue for the 2001 Symposium on this subject.

Policy liabilities for segregated funds, as for other policy liabilities, should be based on a prospective analysis of asset and liability cash flows. Because of the uncertainty of the underlying investment returns on which the liability costs and revenues are based, a stochastic approach is required to estimate these values.

(continued on page 9)

In This Issue

Exploring C1 Risk

continued from page 1

I think I understand how to answer these questions in a generalized asset risk model. I'll draw heavily from investment theory and pension fund management, applying these to an insurance general account. This article has two parts. The first part, contained in this issue of Risk and Reward, describes the basic C1 model itself. The second part discusses various implications of the basic C1 model, especially related to managing the company's investment portfolio.

Stylized Facts

Here are some assertions that I take to be true:

Portfolio variance is a good measure of risk. Portfolios are comprised of components we call asset classes; classes are comprised of individual issues. Issues within an asset class respond in substantially similar ways to economic stimuli. So asset classes are more interesting than individual issues for C1.

The investment literature characterizes asset classes by the moments of their return series. Think of returns as the sum of periodic interest or dividend payments and the change in market value. This is usually called "total return" in the investment literature. So a series' mean represents its expected value and its variance represents its risk.

We characterize portfolio risk by component asset class risks and their interactions. If classes comprising most of the portfolio respond in similar ways to economic stimuli, then the portfolio risk is almost equal to the weighted average of component variance. If component classes respond dissimilarly, then some component variability nets with other component variability and portfolio variance is less than the asset class weighted average. Write:

portfolio variance is less than the asset class weighted average. Write:

$$\sigma_p^2 = \omega'_p * \Sigma * \omega_p \qquad [1]$$

where ω' denotes the transpose of the vector ω,

N = number of asset classes in the universe,

ω = vector of portfolio weights of the N classes and

Σ = covariance matrix of the N classes.

C3 risk measures the harm associated with a change in the basic cost of money in the economy. C3 variability measures the degree to which the fair values of assets and liabilities change in response to interest rate changes. Most invested assets—especially bonds and mortgages—have identifiable C3 variability. Many insurance liabilities also have it. The right amount of asset variability neutralizes liability variability. If you have the right amount, a change in the cost of money in the economy causes no residual harm. If you do not have the right amount, then you have C3 risk.

The NAIC and private rating agencies provide measures of capital adequacy that include a C1 risk component. The algorithms they use map their impressions of the potential for investment loss into asset class capital loads. An "adequately capitalized" company has more assets left over after reflecting these loads than an inadequately capitalized company.

These impressions of risk have deep roots in industry lore but a more tenuous grip on reality. Nevertheless, they represent an important boundary condition and must be reckoned with.

The Basic Model

I begin by removing C3 variability from the scope of my C1 concern. This is fundamental. Asset class returns vary around their means because the basic cost of money in the economy changes from period to period; this is C3 variability. And whether or not I am matching C3 asset variability with C3 liability variability, in either case it is independent of C1 and outside my C1 concern.

Asset class returns also vary around their means because the market risk premium—the market spread—changes from period to period; this is C1 variability for an asset class. High quality asset classes have almost no residual variability when you've removed the C3 variability. Low quality and equity-like classes retain almost all of their variability.

Asset class returns are the means of the risk premia themselves. Let's call this "excess return." These returns have distributions that are nearly normal. Furthermore, individual asset class excess return series are correlated in various degrees to each other.

So now I have all the elements—C1 excess return series with two moments and correlations between them—that a Markowitz portfolio has. And then my articulation of C1 risk is almost trivial.

Under normal circumstances, C1 risk is a measure of the variability of portfolio returns due solely to changes in market risk premia from period to period.

A Simple Example

Imagine that a company's investment portfolio consisted of diversified portfolios within four asset classes: U.S. Treasuries, investment grade corporate bonds, commercial mortgage loans and common stock. Consider their total returns over a recent 23-year period.[1]

Chart 1 - Total Returns

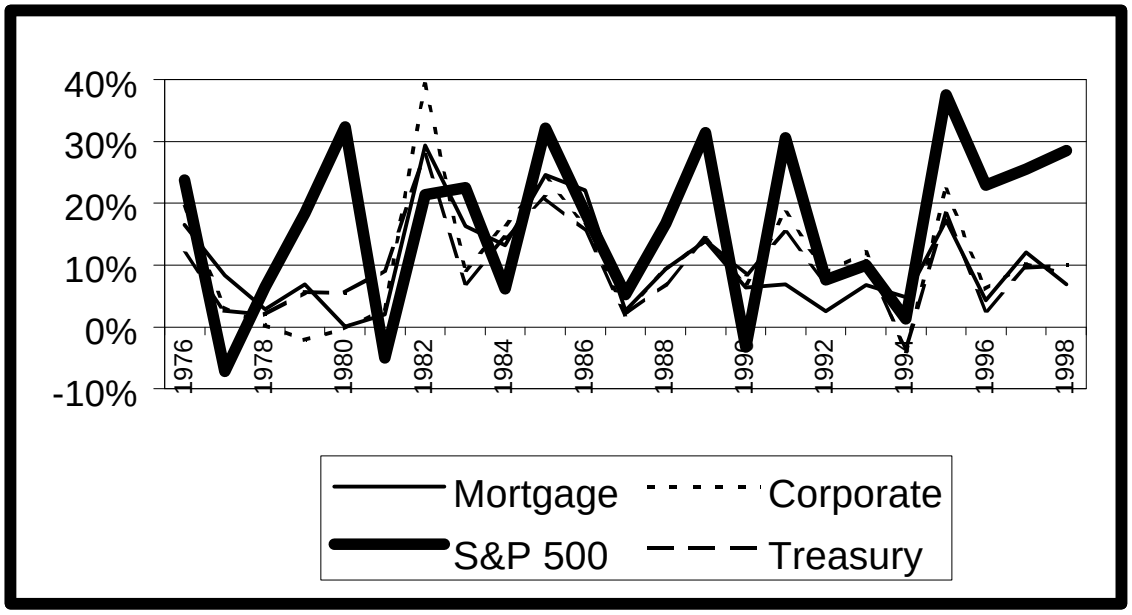

Now consider the same series after adjusting for C3 variability.

Chart 2 - Excess Returns

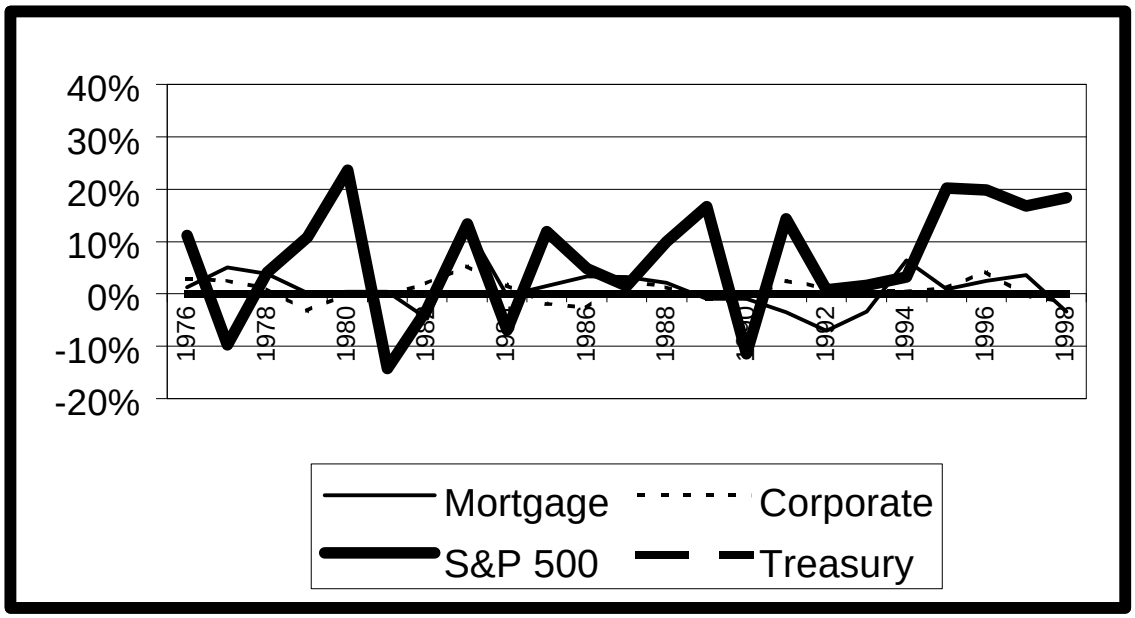

A few features of the return series are apparent:

- Means of the excess return series are lower than the total return series as a result of the C3 variability adjustment. They fall by the total return of a duration matched treasury instrument. That is why the mean of the treasury series falls to zero. Means of other series represent their risk premia.

- All excess return series are less volatile than their total return series. That is, the basic cost of money in the economy—and its change from period to period—explains some of the total return of every asset class. Indeed, by definition, it explains the entire treasury return. That is why the variance of the treasury series falls to zero. Other series retain some variability in response to spread dynamics.

(continued on page 6)

Exploring C1 Risk

continued from page 5

- All excess return series are less correlated than their total return series. If you remove one factor that affected each of them, then they will be less correlated with one another. Indeed, treasury excess returns have no relationship to the excess returns of any other asset class.

Table 1 contains summary statistics of the excess return series.

Table 1

	Expected Returns	Standard Deviation	Correlations			
			Treasuries	Bonds	Mortgages	Stocks
Treasuries	0%	0%	1			
Bonds	0.73	2.41	0	1		
Mortgages	0.93	3.80	0	0.03	1	
Stock	6.30	14.05	0	0.17	-0.03	1

Table 2, using Table 1 figures, assumes some portfolio weights and provides asset class weighted average risk. It also estimates actual portfolio risks by equation 1. The portfolio effect of mixing less correlated asset classes removes about 38% of the weighted average risk. The portfolio will earn an expected premium of 126 basis points per year plus or minus 229 basis points at one standard deviation. Alternatively, if you assume a treasury expected total return of six percent, then this portfolio would have an expected total return of 7.26% and the C1 component of overall risk at one standard deviation would explain total returns of between 9.55 and 4.97%.

Table 2

Class	Portfolio Weight	Expected Returns	Standard Deviation
Treasury	10%	0%	0%
Bond	55%	0.73%	2.41%
Mortgage	25%	0.93%	3.80%
Stock	10%	6.30%	14.05%
	100%		
Weighted Average		1.26%	3.68%
Portfolio		1.26%	2.29%

What Doesn't Matter

Insurance companies balance the needs of several constituents. Flowing from the basic model, here are five issues that, while important to these constituents, have no bearing on C1 risk.

Financial reporting. Statutory reporting assumes risk discontinuities where none exist in reality. It requires that common stocks and NAIC 6 bonds be carried at their market values. This reveals the C1 risk of these assets and makes statutory balance sheet information useful to readers. That is good as far as it goes. But reporting conventions allow performing fixed income classes to be carried at

amortized cost. This presumes that these classes are free of risk, a designation appropriate only for treasuries. The result is that smaller amounts of riskier assets reveal their C1 risk in the balance sheet, but much larger amounts of less risky assets hide their C1 risk.

In fact, C1 risk is present on a continuum across asset classes, but you won't find this in any financial reporting convention recognized by the industry. The way in which assets are reported indicates little about the C1 risk actually borne by a company.

In managing their companies' risk capital, management should try to obtain the highest excess return for the risk it is able and willing to bear. That is, it should diversify its portfolio in an efficient way. Company portfolios are often inefficient because management deploys too much of its risk in classes that appear less risky than they are in reality.

The best argument for attaching importance to financial reporting goes something like this. Regulators, rating agencies and stock analysts prefer smooth earnings and stable reported surplus. Since these constituents significantly affect the company's prospects, the financial reporting convention effectively represents the economic reality of the company.

Ultimately, however, the risk of the portfolio will play itself out, regardless of whether you have distorted risk in your financial disclosures. It will usually show up in the company having either greater than reported risk or an inefficient mix. I am reminded of the famous cave scene in Plato's *Republic*; looking for C1 risk in financial reports is like watching shadows rather than the reality itself.

C3 variability. We often consider C3 variability and duration to be the same thing. But the price variability for an asset can be described by a rich polynomial function, the primary independent term of which is a parallel shift in the treasury yield curve. Other terms include: 1) squared and cubed parallel shifts, 2) curve segment shifts and their squared shifts, and 3) combinations of curve segment shifts. This function is much more general than simple duration and convexity or key rate duration. And it is an important concept for the asset/liability management aspects of spread management.

But none of it matters for C1. Indeed, the residual from this function—the part of asset total return variability not explained by the generalized C3 function—is precisely C1 risk.

Bond defaults and mortgage delinquencies. Since the Hickman studies, groups—including the Society of Actuaries—have estimated historical losses on various fixed income classes. The approach usually nets default and recovery rates. More sophisticated studies use a transition matrix to reflect a time dimension. Some analysts point to municipal bonds and agency-backed residential mortgage securities as low risk instruments because they have low historical loss rates.

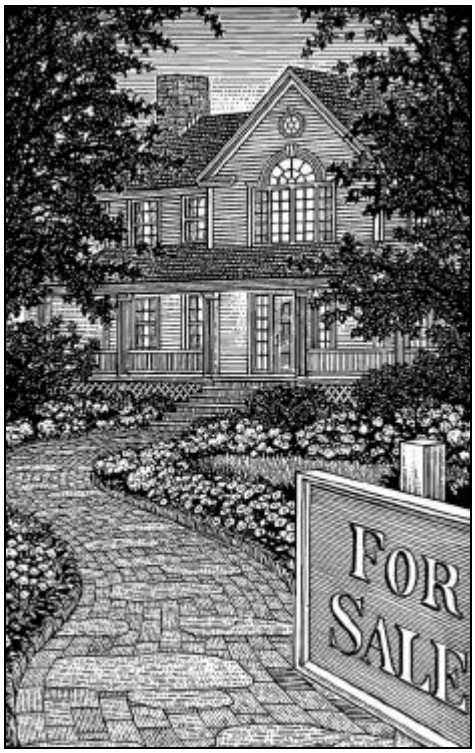

For me, focusing on the discrete event of a default loss ignores the continuous information provided by market prices. For example, when the legislature threatens to change the tax law on municipal bond income, municipal bond prices fall. That is risk, but no bond has defaulted. For another example, as a bond transitions down through ratings on its way to default, the market price falls. That, again, is risk, even though no bond has defaulted.

Market prices generally reflect readily available relative value information. Total returns, which reflect periodic changes in market prices, provide superior risk information than simple defaults and foreclosures. Indeed, you could run a bond portfolio without ever experiencing a default—simply sell the bond when its price falls by more than other bonds' prices. Risk would be manifest, but you would have no defaults.

Liquidity risk. For me, C1 risk is a measure of asset class return variability. Liquidity risk is a measure of asset class salability. At the most pure level, the two are not related.

Some asset classes have a great deal of C1 risk, but very little liquidity risk (i.e. large capitalization growth stocks). Other asset classes have very little C1 risk but a great deal of liquidity risk (i.e. highly covenanted private placements). Still other classes have a proportional amount of each (i.e. high yield bonds). But C1 risk and liquidity risk are not related.

Institutional constraints. The insurance industry is laden with constraints from regulatory and rating agencies. Many are established in the name of controlling investment risks. If you bump up against them, you may hear that your portfolio is too risky. Often these constraints bear only sketchy resemblance to true C1 risk. Companies need to live within their constraints; but the constraints themselves are not C1 risk.

I find it most useful to treat these constituent requirements as boundary conditions. That is, I prefer to mix the assets in an economically efficient way, subject to the constraints imposed on management.

(continued on page 8)

Exploring C1 Risk

continued from page 7

What Matters

Here are elements that are important to C1 risk. I give short treatment to some, with greater discussion later.

Variance of excess returns. On a class by class basis, excess return variance is the best estimate of risk. You will find that variance measures are different depending on what time period you measure them. The investment literature is divided over whether variance actually falls with longer measurement periods. You will want to measure the returns over a consistent period for each class. I like to think of measuring them over a period that approximates a class' return cycle and the investor's holding period.

Correlation of returns. You need to complete a correlation matrix to describe how asset classes interact with one another. Correlations change over time, often following an autoregressive process. So if you correlate excess returns across overlapping 60 month periods, you will often detect a pattern that demonstrates some momentum. Filling in the correlation matrix is important and there is more than a little judgment involved.

Crises. During financial crises asset classes become more volatile and they become more correlated to one another. This is why C1 risk increases so much at these times.

Measurement, deployment, barbells, policy and implementation. For many purposes, it is sufficient merely to measure the C1 risk that an investment portfolio has. You can do that by equation 1, after completing the difficult work of estimating class excess return variances and inter-class correlations.

For other purposes, it is useful to ask the next question. That is, given this level of C1 risk, what is the maximum expected excess return the portfolio can achieve? This is now a simple non-linear optimization problem.

> *"You will want to be careful about barbells. Optimization math will lead you to extreme portfolios. These portfolios place too much pressure on a few cells of the correlation matrix. Barbelled portfolios are not adequately diversified, even if the model claims they are."*

You will want to be careful about barbells. Optimization math will lead you to extreme portfolios. These portfolios place too much pressure on a few cells of the correlation matrix. Barbelled portfolios are inadequately diversified, even if the model claims they are. A measure of professional judgment avoids spectacular mistakes.

The next question asks what the right target C1 risk for the company ought to be. It becomes the central investment policy question for the board of directors. I think the board's policy ought to reflect the term of the company's funding sources, its capital position, the broader insurance risk portfolio the company is bearing and the board's attitude toward risk.

The smallest level of question, then, is implementation: capital gains taxes, new money flows, tactical allocations, asset class management expertise. It tends to be where we spend most of our time.

This completes the basic articulation of the C1 model. Part two of this article will expand on various management and governance issues relative to C1 risk.

Tom Merfeld is vice president at Century Investment Management in Madison, WI. He can be reached at thomas.merfeld@CUNAMutual.com.

Endnote

(1) Returns are for the period 1976 – 1998. Treasury returns are the medium term treasury index. Bond returns are the Lehman intermediate term corporate index. Mortgage returns are the Giliberto-Levy index. And equity returns are the S&P 500. I calculated excess returns by comparing monthly total returns on the risky classes with the monthly return of a duration-matched treasury. Although the empirical duration of the S&P 500 is considerably less, the modified duration of the index is approximated by 1/i. Consequently, in order to better reflect the timing of the cash flows, I related the return on the S&P 500 to the return on a 10 year treasury bond.