Article from:

Risks and Rewards Newsletter

February 2002 – Issue No. 38

Exploring C1 Risk

by Thomas Merfeld

Editor's Note: This is part two of a two-part article. The previous part ran in the July 2001 issue of Risks and Rewards.

The first part of this article provided an overview of an excess return-based C1 model. This second part provides richer discussion of C1 risks in five areas: 1) the economics of C1 risk, 2) portfolio optimization, 3) policy, 4) risk interaction and 5) institutional aspects of C1 risk. It has no more than a loose organization and you don't need to read it in sequence.

Economics of C1 Risk

Difficult markets

Asset classes become more correlated to one another during difficult markets. As a result, some of the portfolio's diversification slips away and C1 risk increases. It is a good idea to have several sets of correlations at hand; these indicate the amount of C1 risk in different economic environments.

Derivative contracts

Imagine that all derivatives are one of two types. The first is interest rate-based contracts, such as futures, options, caps, floors, and swaps with LIBOR or treasury indexes. These contracts all have different payoff profiles depending on the course of the underlying interest rates. That is, they may have substantial C3 variability and may net with other asset and liability variability to produce C3 risk. But they do not have material C1 variability.

The second type of derivative is asset-based contracts, such as futures, options and swaps with S&P 500, Lehman corporate and Merrill Lynch high yield bond indexes (along with many others). These contracts have payoff profiles that may depend on the course of risk-free rates and always depend on the risk premium for the asset class. They can provide or reduce exposure to an asset class without using cash. In effect, derivatives use debt to provide the exposure—they are naturally leveraged instruments. Since debt imposes a non-contingent claim on the assets of the company, it effectively reduces the proportion of the only non-contingent asset on the company's balance sheet. That is, debt effectively reduces cash. Table 3 increases the Table 2 portfolio by .36% by adding 20% to the bond class by means of a derivative contract. Equation 1 measures portfolio variance without regard to the signs of asset class weights.

Table 3

Class	Original Weight	Derivative	New Weight	Standard Deviation
Treasury	10%	(20%)	(10%)	0%
Bond	55%	20%	75%	2.41%
Mortgage	25%		25%	3.80%
Stock	10%		10%	14.05%
Total	100%	0%		
Original Portfolio				2.29%
New Portfolio				2.65%

Normal return distributions

Some early studies of stock market total returns, such as the original Mandelbrot article, demonstrated excess kurtosis. That is, they have more probability mass in their tails than a normal distribution. Time-varying levels of volatility can explain this. Usually the stock total return series distribution is normal. But in times of financial crisis, prices of the riskiest portion of the corporate capital structure can fall by much more than a normal distribution would suggest.

It isn't as clear, however, that non-stock total return series are leptokurtotic. And, as it relates to C1 risk, excess return series consistently appear to have less kurtosis than their analogous nominal return series.

I think it is appropriate to assume normality for baseline C1 risk assessment and optimization. And it is prudent to be comfortable with the loss associated with the downside tail coming to bear in sensitive classes.

Investment classes

Ultimately, identifying a set of asset classes to use has an empirical element and a judgment element. If two excess return series are not highly correlated or if they have different liquidity or different RBC characteristics, then they may represent separate investment classes. So the empirical element involves close review of time series data, after adjusting for C3 variability. It also involves comparing bid-offer relationships and RBC algorithms to determine whether two instruments have similar levels of liquidity and receive similar RBC treatment.

Professional judgment decides how much correlation is enough and how similar liquidity must be to warrant a distinct asset class. You want enough classes to assure diversification. But you don't want so many that the distinctions between classes lose their meaning. In most well-diversified general accounts, I think 8 to 12 classes is appropriate.

New asset classes

Sometimes the capital markets develop new asset classes. You can't perform time series analysis on them because they have no return history.

Usually, however, apparently new classes are merely combinations of existing asset classes for which you have data. An example is the commercial mortgage-backed security. In theory, high quality CMBS are a combination of a risk-free instrument and mortgage whole loans. As you move down in quality, the risk-free portion becomes less and the whole loan portion becomes greater. At some point the risk-free element turns to debt and from that point on, the security becomes more and more leveraged. You can model this class as a combination of mortgage whole loan and treasury, with the treasury piece being either positive or negative.

Truly new asset classes, although rare, force you to make judgments. Even eliminating these from the portfolio is a tacit articulation of their risk parameters.

Investment time horizon

Parameters for risk and correlation can be different depending on the length of time over which you estimate them. One extreme is to measure the parameters over a very long time, since the company has a perpetual life and it rarely needs to liquidate assets. Another extreme is to measure them over the period of time that it takes a product to reprice.

I suggest that two guidelines are useful. First, estimate the parameters over the entire return cycle of the investment class. Real estate cycles tend to be long; high quality corporate cycles tend to be shorter; municipal bond cycles tend to be variable. But at least you allow the data to speak. Second, estimate the parameters over the length of time that you're likely to be holding the funds. For funds associated with a two-year GIC, two years is an appropriate time horizon. For an annuity, it may be until the end of the surrender charge schedule.

As a practical matter, I think five years is about right for most companies. You need to use a consistent time horizon to correlate two time series. Five years is long enough for most series to have gone through a cycle. And most companies have products with different time horizons as well as perpetual capital. So five years is a useful rule of thumb while recognizing the conceptual nuances.

Other analysts have suggested exponential smoothing, which weights more recent data more heavily, or a formal GARCH approach.

Scenario testing

In the general case, you shouldn't need to run special portfolio scenarios, since excess return distributions are assumed to be normal and there are no path dependencies. So closed-form statistical procedures can tell you everything you need to know about portfolio excess return dispersion.

Simulated methods may be useful. For example, you may wonder about a particular path for a macroeconomic variable such as industrial production. After estimating relationships, you could simulate C1 exposure to this environment. Money center banks and dealers have successfully used a Cholesky decomposition process for this purpose.

Extreme events

Unusual market conditions occur every couple of years, for one reason or another. The associated flight to quality quickly lowers the prices on the riskiest classes. That is, markets become more volatile and more correlated at the same time. So portfolio values move by much more than you would estimate under baseline assumptions.

Some professionals express dismay that a "multiple sigma" event would occur within their working lives. But that is a strikingly naive position to take when the data on time-varying return variances in some markets and increased correlation of all risky markets during periods of stress are readily available.

A reasonable approach to address the phenomenon is in Table 4. Note that risks are doubled and correlations set to unity. Portfolio risk is almost triple the base portfolio risk of 2.29%. So this would be a three standard deviation event that you can assume will happen regularly. Management and the board need to know whether the company's operations remain viable under these circumstances.

(continued on page 8)

Exploring C1 Risk

continued from page 7

Table 4

Class	Weights	Risk	Correlations Treasuries	Bonds	Mortgages	Stocks
Treasuries	10%	0%	1			
Bonds	55%	4.82	0	1		
Mortgages	25%	7.60	0	1	1	
Stock	10%	28.10	0	1	1	1
Portfolio	100%	7.36				

Portfolio Optimization

Expected excess returns

Excess return, for a period, is what remains after removing the risk-free component of an index total return. You do this after identifying the total return on a treasury instrument that has a duration equal to the index. In general, the excess return will be positive if the index's risk premium remains the same or falls. It may be negative if the premium rises. The average of the residuals from this process is the expected excess return. It makes sense because the company's liabilities are generally priced as a function of the risk-free rate. That is why investment year method models have been important to insurance companies. Endowments and foundations often use real returns and estimate other parameters from their real return series.

Sometimes this process yields an absurd result. One series' risk and return may be dominated by another's. In these cases, a Bayesian approach can be useful, the prior conviction of which is that the asset class' variance is more reliable then its excess return. You can then assign an expected excess return to the class that is consistent with its risk and the risk/return pairs of other classes.

This describes excess returns expected to prevail over longer periods of time. They are appropriate for relatively permanent allocations and may be called strategic excess returns. By contrast, most classes may be quoted at current levels. It is not clear what these quotes mean since the excess return premium and realized excess return can change during the holding period. They are commonly called current returns.

Constraints

Consider constraining the optimization in four ways to make the portfolio more robust. First, the portfolio needs a minimum level of liquidity. A reasonable approach is to assign degrees of liquidity by relative bid/ask spreads. A better approach is to measure relative asset class salability during periods of financial crisis. You need to keep your institution viable during the most stressful periods; markets often don't function well during these times. So treasury instruments would have a liquidity index of 100, venture capital 0 and everything else in between. The actual constraint depends on the certainty of the company's funding and how much transaction cost it is willing to incur to accommodate disintermediation.

Second, the portfolio should not encumber more than a certain amount of risk-based capital. Finance functions often complain about the portfolio using too much capital. A simple means of addressing the issue is to assign capital loads by asset class. Then constrain the optimization to a maximum capital encumberance. In this way, allocation decisions play within the rules of the NAIC and private rating agencies, but are made on the basis of sound economics.

Third, since some constituents care about the distinction between debt and equity, the allocation can be so constrained. Fourth, managers may have reasons—such as unrealized taxable gains they are unwilling to trigger—to slow the pace of reallocation.

Optimization of excess returns

A simple constrained non-linear optimization routine maximizes expected excess return at a target risk.

Be prepared to work with the allocation through trial and error. Two common problems arise. First is the robustness problem of a barbelled portfolio. The second problem arises because the partial derivatives in the maximization function can be almost equal. The routine may load up on one class even though there is only immaterial benefit to doing so.

Table 5 on the next page shows some optimization runs. Your position on portfolio parameters and constraints can dramatically affect what your optimal asset allocation is. Indeed, your allocations almost tip off what you consider to be important.

Table 5

Class	Return	Risk	Liquidity	RBC	Base	Allocation Percentages 1	2	3	4	5
Treasury	0.0%	0.0%	1.0	0.0%	10.0%	23.0%	19.0%	28.0%	9.0%	0.0%
Bonds	0.73	2.41	0.85	1.0	55.0	38.0	40.0	37.0	67.0	83.0
Mortgages	0.93	3.80	0.4	4.0	25 .0	27.0	28.0	22.0	15.0	12.0
Stocks	6.30	14.05	0.9	30.0	10 .0	12.0	13.0	13.0	9.0	5.0
Return					1.26%	1.29%	1.36%	1.28%	1.20%	1.03%
Risk					2.29%	2.29%	2.41%	2.29%	2.29%	2.29%
Liquidity					0.76	0.77	0.76	0.80	0.80	0.80
RBC					4.50%	5.00%	5.30%	5.10%	4.00%	2.80%

(1) Optimize at base risk with no constraints.
(2) Optimize at bond-quality portfolio with no constraints.
(3) Optimize at base risk with minimum liquidity of .8.
(4) Optimize at base risk with minimum liquidity of .8 and maximum RBC of 4%.
(5) Optimize at base risk with minimum liquidity of .8 and maximum RBC of 4% and maximum equity of 5%.

Strategic, actual and tactical allocations

A strategic allocation represents relatively permanent and maximized strategic excess returns. You can use it for many benchmarking and return attribution purposes.

The actual allocation will naturally wobble around the strategic allocation due to operational cash flows, uneven asset class maturities and market value changes. To a degree, these fluctuations are trivial and you shouldn't try to manage them. Beyond this degree, they represent a fundamentally different allocation and you might want to take steps to nudge the actual toward the strategic allocation. Fluctuations of riskier classes upset economically equivalent portfolios more than fluctuations of less risky classes. Bands placed around strategic allocations—such that all allocations within the bands are stipulated to have economically equivalent levels of strategic excess return and C1 risk—should reflect this uneven impact.

A tactical allocation is one in which you adjust a class weight above or below the strategic bands in the belief that it will perform better at the new weight than at the strategic weight. This performance ought to reveal itself within a short time, perhaps a year. You will want to measure how your tactical decisions perform along with how much the move affected the portfolio's C1 risk. Relative valuation models are reasonable bases for these moves.

Policy

Insurance regulation

The commissioner's office ought to be responsible for imposing a maximum level of C1 risk by company. It should depend on the reliability of the company's funding sources, its other risks and its capital. Maximum C1 risk ought to represent consistent potential for ruin from company to company. Given the clear concept of C1 and the understanding of the company's entire portfolio of risks, consistent statistical measurements of ruin likelihood should be feasible.

Beyond these steps, however, the company is responsible for establishing its C1 target within this limit.

Investment policy

The company's board of directors has a grave responsibility in establishing the investment policy. It needs a basic understanding of the company's funding sources and capital position. In discharging its responsibility, the board will establish a risk target. It may also establish liquidity limits and may go as far as placing upper limits on individual classes.

Beyond these steps, however, management is responsible for deploying its C1 risk budget.

(continued on page 10)

Exploring C1 Risk
continued from page 9

Risk interaction

C3 risk
It's useful to develop cohorts of risk and correlation estimates under rising and falling risk-free rates. To the extent that they are different, portfolio risk may be similarly different. It's good to know how sensitive the portfolio risk is.

It's probably not useful to estimate excess return cohorts or to optimize excess return under each rate environment. Doing so would assume that you know what environment will prevail over the investment horizon.

Exchange rate risk
In general, currency risk reflects the possibility that the market value of surplus would fall as a result of changes in foreign currency values relative to the U.S. dollar. It is complicated by at least two factors.

First, a company may have insurance operations—premium and reserves—denominated in a foreign currency. The operations may take place under a branch or subsidiary arrangement and may be wholly or partially owned. Although the financial reporting may differ under any of the arrangements, there is no significant exchange rate risk to foreign operations. This is because, whatever happens to the dollar-denominated value of the assets also happens in rough proportion to the dollar-denominated value of the liabilities; surplus is almost unaffected.

For foreign currency-denominated assets unmatched by reserves, there is a second complication. In contrast to the single-event nature of C3 risk, foreign currency risk is driven by how multiple specific currencies combine to affect surplus. So foreign currencies also have a portfolio effect. And to measure company exchange rate risk, you need to measure the variability of each currency, relative to the dollar, to which the company is exposed. You then need to measure how these currencies interact with one another.

It may often make sense to accept some exposure to foreign exchange risk. Think of securities in each currency being fairly priced for their level of credit risk. Now if you can offset some of the foreign exchange risk, then you have produced value.

It's more intriguing to reflect on decoupling foreign currency variability from the pure C1 risk in foreign stocks and bonds. Foreign exchange markets allow companies to hedge the currency risk in these markets without affecting the actual C1 exposures themselves.

Liquidity and crises
Sharply curtailed liquidity, although different in nature from C1 risk, often accompanies the same economic downturn that heightened asset price volatility does.

Insurance risks
Most life and property risks have nothing in common with financial markets. So insurance exposures don't reflect back to raise or lower C1 risk.

But they do influence the company's C1 risk target. Here's the argument: the company's overall financial risk limit ought to be expressed as a tolerable percentage of company surplus. Assume the same amount of surplus in both of the following cases. If relatively few insurance risks are present, then the company may be able to bear greater C1 risk. Additional units of C1 risk, however, will translate directly into proportionally greater overall financial risk and the company may have little room to adjust its C1 exposure. If relatively greater insurance risks are present, then the company may be able to bear less C1 risk. Additional units of C1 risk, however, may have little incremental influence on the company's overall portfolio of financial risk.

Financial reinsurance
Co- and Mod-co contracts remind us that not all C1 exposures are present on the company's balance sheet, and that not all balance sheet items bear C1 risk. Financial reinsurance contracts can leverage or deleverage the economic balance sheet. Proper C1 assessment reflects the reality of the contracts rather than the fantasy of published statements.

Institutional aspects of C1 risk

Allocating investment income
Imagine that C1 optimizations take place at two levels. The more important global process allocates all consolidated general account invested assets to specific classes as I've described. The subsequent process drives the global allocation down to specific lines of business. The output of the second optimization is the investment spread that attaches to each line of business. It is one of two important components of equitable investment income allocation. Here are some details.

Each line should receive the risk-free rate prevailing at the time reserves came into the company or were reinvested. This is a standard investment year method with two clarifications. First, in this component of investment income, I'm referring only to the risk-free rate, not the total return. Second, it should receive the rate that is consistent with the duration of the reserves. That is, if a line has a liability C3 variability (as measured by duration) of 4.5, then the risk-free rate should be that of whatever treasury has a duration of 4.5. As treasury rates migrate to different levels, the line will develop a weighted average risk-free rate. This rate will respond quickly to changing treasury rates in short duration lines and more slowly in long duration lines.

Each line should then receive a C1 spread on top of its risk-free rate. This is merely the weighted average spread of the asset classes allocated to it.

The product level C1 optimization begins with the company's actual asset allocation. It then swaps assets in each class for each product line in and out of surplus. You can assign unique risk targets and liquidity and RBC constraints to each line. The objective function of the optimization is to maximize the aggregate product line spreads subject to their risk targets and boundary conditions. That is, the process will conclude when the partial derivatives of return per unit of risk are equal across product lines.

You can establish product line C1 parameters for risk, liquidity and RBC on the same basis as you establish them for the company as a whole. How long is it until the cash is needed? Are there important insurance contingencies? Are earnings strong? Is capital plentiful?

This process can yield some intriguing results. Consider the following in a Life/P&C group of companies. It's easy for me to imagine allocating some municipal bond excess returns to a life insurance line even though the life company wouldn't hold the actual bonds. But since the life reserves allowed the group to hold a greater dollar amount of municipal bonds (at a given allocation), it can share synthetically in the excess returns.

Product line managers may want to receive units of excess return based on RBC rather than true C1 risk. They may further want to receive current weighted average spreads rather than long term spreads. Both of these make sense to me, provided there is a consensus among managers.

Product segments and surplus investments

One approach is to isolate the two portfolios according to the source of funds—surplus and reserves—and then optimize them independently. A nice feature to this approach is your ability to specify objective functions and constraints in each. But there are theoretical and practical problems with the approach. Optimization theory tells us that the aggregate of two efficient portfolios may not be efficient. So it's better to have one portfolio with well-crafted parameters. Furthermore, from a practical perspective, the two portfolios bleed between one another. So when a product line has earnings, the same dollar moves from reserves to surplus. And, in times of stress, it moves back to reserve. The notion of separate portfolios have a long history in industry convention, but is not real.

I think it's best to have just one portfolio. Then carefully target the portfolio's C1 risk, reflecting the company's surplus level and other factors, and establish an efficient allocation. This will give you good economic characteristics. Subsequent to all of this, direct investment excess returns to various products and to surplus.

Multiple companies in a group

It is almost always more efficient to construct just one portfolio rather than aggregate several. I like to look through group entity structures in assessing and optimizing C1 risk. This supports the reality that management provides capital to subsidiaries under stress or sells them and bears the loss.

Two qualifications are useful here. First, it remains true that, subsequent to the group's overall allocation, real securities need to be placed in real insurance companies. Each entity demands a reasonable, albeit not efficient, allocation. Secondly, there may be contractually understood circumstances to pass specific returns to specific claimants, even outside of the variable product context. Examples include certain participating policies and CBO structures in which the company owns the entire equity portion and are consolidated under GAAP.

But in general it's best to think of the group having just one portfolio. Then make reasonable allocations.

Mutual and stock charters

We've known for a long time that, if a corporation—including a stock insurance company—doesn't see good opportunities for its capital, then its board should dividend the capital out for shareholders to deploy in another equity venture. By analogy, if a mutual insurance company doesn't see good opportunities for its capital, then it should increase its exposure to equities. That is, as a proxy for shareholder stock investment, the mutual can make the same investment itself, in theory for the benefit of policyholders. In practice, many stock company boards do not declare this type of signaling dividend, acting more like a mutual company. In any case—stock or mutual, dividend or not—greater amounts of free capital are appropriately allocated to equity investments.

More generally, greater levels of capitalization are associated with greater risk bearing capacity within the insurance company. So the target risk, around which to optimize the portfolio, can be greater. This is consistent with the dividend irrelevance theorem of Modigliani and Miller. And the type of legal charter doesn't change the economics.

Life and property companies

Investment professionals, who are lay asset-liability managers, are fond of telling company management that they have structured the portfolio to be "consistent with the nature of the liabilities." Almost without meaning in itself, investment managers use the assertion as a way to avoid the question of why they have mortgages and corporates in the life company and municipals and stocks in the property and casualty company.

But the question remains. And it becomes more poignant once you get past the C3 variability issue, which truly may be different for life versus property and casualty. On what basis do you structure the C1 exposure so differently?

I look to the drivers of C1 capacity: available capital, other insurance risks, reliability of funding sources and board temperament. These factors are not fundamentally different between life and property companies. Indeed, some groups contain both types of charters. With capital transferability, consolidated risks and common boards, I submit that the C1 risk targets ought not to be fundamentally different. Most differences in portfolio structure ought to relate to taxation or potential rating agency scrutiny.

Differences in the way life and property companies ought to consider, assess and target C1 risk are trivial.

(continued on page 12)

Exploring C1 Risk

continued from page 11

Income and capital appreciation

Income returns are deemed to have an endogeneity that capital gains do not have. So income returns are capitalized more heavily in company valuations. Visible capital gains give the impression of earnings volatility. And so gain streams are heavily discounted in company valuations. Management often believes that it is adding value to shareholders by substituting relatively modest income in place of relatively robust capital gains.

In truth, the distinction between capital gains and investment income is economically meaningless. Long-established financial reporting conventions distinguished between monies clearly owed the investor—declared dividends and timely coupons—and the balance of the investment corpus. The distinction begins to crack with high premiums and deep discount bonds because monies clearly owed are adjusted by amortization and accretion to derive interest income. And it crumbles when applied to interest-only securities in which there is no ultimate principal balance. Financial engineers exploit the convention by ascribing amortized cost and NAIC 1 treatment to a 30 year note that is roughly five parts S&P 500 and one part U.S. treasury zero coupon bond that will ultimately pay its par value on all six parts.

True C1 risk reflects the variability of the asset class' total return series, net of its C3 variability, without regard to the character of its return components. Management focused on producing investment income may be responding to shareholders' current desires, but it is probably not building economic value.

Investment performance measurement

Actual returns demand a context, so people compare to benchmarks. In the context of C1, it is clear to me that the first comparison is between actual total return and excess return. This will tell you whether and by how much your particular subset of an asset class outperformed the class as a whole.

It's useful, though, to ask the next question. In an asset allocation context, the investor holds a particular asset class to play a certain role in the portfolio. At times, the manager of an asset class generates a positive excess return by investing in securities that are not members of the asset class. Doing so may reflect sound portfolio management, but only if the actual securities portfolio has statistical properties similar to the index. Otherwise, the efficiency of the portfolio has been compromised. Indeed, a good measure is the amount of excess return a portfolio has earned per unit of tracking error. The pension investment literature calls this an information ratio. So, think of the average excess return divided by the standard deviation of the excess return series. It is analogous to a Sharpe ratio, which indicates excess return per unit of risk. A high information ratio acknowledges that excess return is valuable. It also acknowledges that the asset class has a role to play in the overall portfolio and that a high excess return that compromises this role is less useful.

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

1998-99 Redington Prize Awarded

To promote investment research, the Investment Section sponsors a biennial prize of $2000. The prize is named after F. M. Redington, the eminent British actuary who coined the term "immunization" in a 1952 paper in the Journal of the Institute of Actuaries. This is the sixth award since the prize was first established.

The Council would like to thank all those who took the time to send in nominations. The Prize Committee received a total of 11 nominations. Many worthy papers were submitted, and therefore, the Committee's decision was not an easy one. For the 1997–98 publication period, the Prize has been awarded to:

"Term Structure Models: A Perspective from the Long Rate" by Yong Yao (A.S.A.) in *NAAJ,* Vol. 3, No. 3, (July 1999).

The paper responds to an important need in the actuarial profession, which is the valuation of long term insurance or pension plan cash flows where the yield curve, as measured by prices of traded securities, does not exist.

The paper partially answers two questions: in frictionless markets having no arbitrage, what should the behavior be; and, in known term structure models, what can the behavior be. In frictionless markets having no arbitrage, yields of all maturities should be positive and uniformly bounded from above. The yield curve should level out as term to maturity increases. Slopes with large absolute values occur only in the early maturities. The paper goes on to show how the long rate behaves in well known term structure models. Practitioners using these models at these longer durations should be aware of their behavior.

On behalf of the Investment Section, the Council would like to congratulate and thank Mr. Yao for the exceptional work he has accomplished. The Council also expresses its gratitude to the members of the Prize Committee. These are Nino Boezio, Luke Girard, Jeremy Gold, David Li, John Manistre, Robert Reitano, Michael Sherris, Elias Shiu, Ken Seng Tan and Richard Wendt.

The next Redington Prize will be awarded in 2003 for papers published in 2000–01.