DISCUSSION OF PAPER PUBLISHED IN VOLUME LXXVII

RISK LOADS FOR INSURERS

SHOLOM FELDBLUM

DISCUSSION BY STEPHEN PHILBRICK

VOLUME LXXVIII

AUTHOR'S REPLY TO DISCUSSION

VOLUME LXXX

DISCUSSION BY TODD BAULT[1]

1. INTRODUCTION

I have been following with great interest this discussion "thread" in the *Proceedings* [1, 2, 3], along with the recent papers of Rodney Kreps [4] (with Daniel Gogol's reply [5]) and Glenn Meyers [6] (with Ira Robbin's reply [7] and Meyers's response [8]). Obviously, this is an important topic for the CAS, as evidenced by the amount of discussion it has generated; and it is of particular interest to me, given my current specialization in rate of return, surplus need, and related areas of financial actuarial practice. The focus of the Feldblum/Philbrick discussion has been five methods of setting risk loads and the relative merits and deficiencies of each. The other papers by Kreps and by Meyers deal with related approaches and issues. I wish to add two observations to the discussion:

[1]The author would like to thank Mr. Randall Holmberg for the training, insight, and encouragement he has provided me over the years, and for the many stimulating discussions we have had, one of which led to this paper.

- All of the methods are more *similar* than different, including the methods discussed by Kreps and Meyers, and if care is taken to use a common set of assumptions, the methods are nearly *equivalent*;

- *None* of these methods, *including* CAPM, resolves several fundamental problems; and any risk loads derived from these methods must still contain a great deal of subjectivity, more than is implied by Feldblum's discussion of CAPM.

I wrote this discussion to question the level of certainty conveyed in Feldblum's initial paper and to keep the topic open. Although Philbrick's comments help in this regard, he does not go far enough. I am concerned that inexperienced actuaries will see betas published in the *Proceedings* and will feel justified in rushing off to use these in setting profit loads, despite Feldblum's warning that his calculations were for illustration only. There are still many unresolved issues regarding the measurement of risk and its application to profit provisions. The research to date is encouraging and highly connected, as we shall see, but there is still much left to do.

2. THE FIVE RISK LOADS—ARE THEY EQUIVALENT?

I start this analysis with the work of Rodney Kreps—his paper already describes most of the connections I want to demonstrate, but they have not been fully integrated. We can use his equations to show that variance, standard deviation, ruin theory, and CAPM describe similar (and nearly equivalent) concepts. Although I could not incorporate utility theory and the reinsurance method with sufficient mathematical rigor, there is reasonable evidence that these latter approaches are also strongly related to the others.

We begin with Kreps's equation for surplus supporting insurance variability of a given portfolio. Kreps assumes this portfolio represents a company's book of business, and Feldblum agrees that this assumption is appropriate for adapting CAPM to insurance. However, I will assume this portfolio is the *industry* portfolio (in the next section, I will support this position—one could use a company portfolio in this analysis and reach the same conclusions). The equation is

$$V = zS - R, \tag{2.1}$$

where V is the surplus, S is the standard deviation of the portfolio, R is the return in dollars, and z is the standard normal percentile value associated with a given probability of ruin (i.e., exceeding needed surplus). Kreps does not show explicitly that a *ruin theory* equation produces this formula; it is

$$\Pr(\mathbf{L} + E > \mathrm{E}(\mathbf{L}) + E + R + V) < e, \tag{2.2}$$

where $\mathbf{L}$ is the random variable for loss (boldface will always be used for random variables), E is expense, and e is the threshold probability of ruin corresponding to the value z. Note that the standard deviation of $\mathbf{L}$ is S. Standardizing $\mathbf{L}$ to $(\mathbf{L} - \mathrm{E}(\mathbf{L}))/S$ produces the solution for z, from which (2.1) follows:

$$z = (R + V)/S. \tag{2.3}$$

Kreps then produces the equation for the marginal surplus required for a new risk $\mathbf{x}$. We shall assume $\mathbf{x}$ is very small in magnitude compared to $\mathbf{L}$, for both means and standard deviations. The equation for the marginal surplus is

$$V' - V = z(S' - S) - r, \tag{2.4}$$

where V' and S' are the surplus and standard deviation, respectively, for the portfolio with $\mathbf{x}$ added, and r is the return for $\mathbf{x}$. Further, Kreps solves for $S' - S$ as

$$S' - S = \sigma(2SC + \sigma)/(S' + S), \tag{2.5}$$

where σ is the standard deviation of $\mathbf{x}$, and C is the correlation coefficient of $\mathbf{x}$ and $\mathbf{L}$.

Gogol, in his discussion of Kreps's paper, noted that this approach is highly "order dependent." That is, Equation 2.5 shows the increase in standard deviation if x is the *last* risk added to the portfolio. If we assume that L is a portfolio of risks identical to x, then x's contribution would be σ if it were the *first* risk in the portfolio, a smaller number than σ (or at least not larger) for the second risk, and so on, despite the fact that the risks are identical. Measuring each risk's surplus requirement based only upon its marginal risk contribution will underestimate the total surplus need of the portfolio. Gogol developed a formula to allocate the total surplus need to all individual risks based upon an average of the risk's standard deviation on a "first-in" and "last-in" basis. (Please see Gogol's paper for more details.) This is an important adjustment for practical implementations of this method, but please note that Gogol's technique is not the only way to do this—another approach will be discussed at the end of this section.

Up to this point, we could call Kreps's method a ruin theory approach, because ruin theory is the basis of his equations. But is this method related to any other approaches? Suppose that x and L are *independent* (rare, but the usual assumption), so that $C = 0$. Then

$$S' - S = \sigma^2/(S' + S) \approx \sigma^2/(2S), \tag{2.6}$$

because σ is small compared to S. Thus, the marginal surplus is a function of the *variance* of the new risk. (Feldblum cites this formula in Footnote 1 of his reply to Philbrick, but does not mention explicitly the independence requirement.) Conversely, if x and L are completely dependent (also rare, but illustrative), then $C = 1$ and

$$S' - S = \sigma(2S + \sigma)/(S' + S) \approx \sigma. \tag{2.7}$$

Again, this is true because σ is small compared to S. Now the marginal surplus is a function of the *standard deviation* of the new risk. So in the most common situation, where x is slightly correlated with L, the marginal surplus will be a linear combination of the variance and standard deviation related to the *covariance*. In my opinion, this makes the whole "variance vs. standard deviation" debate much less interesting, because both are simply special cases of a unifying covariance framework. Actuaries may continue to choose one or the other method on the basis of tractability concerns (and measuring covariance *is* very difficult), but they should be aware of what these decisions imply and whether or not their assumptions are appropriate.

We have two important results so far:

1. The distinction between variance and standard deviation methods is somewhat artificial. Which method to use is a function of the correlation between the new and existing risks, and in most cases, the "correct" answer is a marginal risk approach that incorporates covariance.
2. Marginal risk methods (including variance and standard deviation methods as special cases) are closely related to a ruin theory approach.

Let us examine $S' - S$ further. Define P as the premium associated with the industry portfolio, and let p be the premium associated with the new risk. From Equation 2.2, it should be clear that S/P is the standard deviation for the industry return on premium. Further, $(S' - S)/p$ is the *marginal contribution* to the standard deviation of the return on premium from the new risk. Using Equation 2.5 and some algebraic manipulation, note that

$$\begin{aligned}
[(S' - S)/p]/[S/P] &= [\sigma(2SC + \sigma)/((S' + S) \times p)]/[S/P] \\
&= [(2SC\sigma + \sigma^2)/(p \times 2P)]/[(S/P) \times (S' + S)/2P] \\
&\approx [C \times (\sigma/p) \times (S/P)]/[(S/P) \times (S/P)] \\
&\qquad \text{(because } \sigma^2/(pP) \text{ is small)} \\
&\approx \mathrm{cov}(\boldsymbol{x}/p, \boldsymbol{L}/P)/\mathrm{var}(\boldsymbol{L}/P). \qquad (2.8)
\end{aligned}$$

The last part of Equation 2.8 looks remarkably like a CAPM beta. In fact, this *is* the formula for beta proposed by Feldblum, so let us "set" β equal to $[(S' - S)/p]/[S/P]$. This is *not* how beta would be derived in practice, but it serves as a link in the chain of reasoning of this analysis. However, Feldblum might disagree with this characterization, as I have just said, in effect, that the variance in profit equals the variance in loss, and Feldblum produced at least three examples to demonstrate that this is not true. Before proceeding, then, I should justify this simplification.

Two of Feldblum's examples are rather naive descriptions of how to measure the variance of losses, and it is fairly easy to remedy the problems he describes. His first example, retrospective rating, could be fixed by measuring the variance of the insurer's *effective* loss distribution, which is zero in his idealized example. His third example, heterogeneous mix of risks, is fixed by using homogeneous groups, or by adjusting for the heterogeneity in a reasonable fashion.

But Feldblum does indicate some important sources of risk that are not derived from the loss distribution, including inflation, investment, default, and parameter risk. All of these need to be measured, but the method for doing so does *not* require measuring the variance of profits directly. In fact, given the problems with calendar year measures of profitability in insurance (which Feldblum used, although he did acknowledge that problems existed), it may be preferable to measure the variance of profitability in other ways, such as starting with the variance of accident year losses and modeling additional sources of risk as required. (This will be discussed further in the next section.) Thus, as long as x and L reflect these additional sources of risk, it is appropriate to use β as I have defined it. Therefore, I shall assume that x and L are so stated.

Feldblum will probably have one final point of disagreement: Covariance and ruin theory approaches usually do *not* reflect these additional sources of risk, so to claim that x and L consider these factors alters his initial assumptions to represent a situation much more like CAPM, making this an unfair comparison. Further, adjusting loss distributions to reflect these risk sources is non-trivial, and probably non-objective as well. These are both reasonable points, but for most lines of business (after adjusting for special features like retrospective rating), why is the loss distribution *not* a reasonable first approximation for measuring the variance of profits? With the exception of parameter risk, most of the other components have very low variance compared to losses. Parameter variance must be included, but it is not clear that measuring calendar year variance of profits directly is the best way to do this.

Returning to the analysis, consider the following return on equity equation, which is a form of the equation used by Ferrari [9]:

$$R_m = R_f + (P/(zS\beta))R_p, \tag{2.9}$$

where R_m is the target return on equity for risk x (and, in fact, for all risks), R_f is the risk-free rate of return obtained from the supporting surplus, R_p is the return on premium, and $P/(zS\beta)$ is the premium-to-surplus (or *leverage*) ratio appropriate for x. But wait—shouldn't this last statement be a question? *Is* $P/(zS\beta)$ an appropriate leverage ratio for x? The answer is yes, *if* $P/(zS)$ is the appropriate leverage ratio for the industry portfolio. To see this, consider the standard CAPM equation:

$$R_e = R_f + \beta(R_m - R_f), \tag{2.10}$$

where R_m is the return on equity for the industry portfolio. To solve for the return on premium, R_p, needed to produce R_e, according to Feldblum, we subtract the risk-free rate and divide by the portfolio leverage ratio to obtain:

$$R_p = (R_e - R_f)/(P/(zS)) = (R_m - R_f)/(P/(zS\beta)). \tag{2.11}$$

But Equation 2.11 is also equivalent to Equation 2.9 *if* R_m in Equations 2.9 and 2.10 means the same thing. Note that Equations 2.9 and 2.10 are two different approaches to the same question: How do we determine the needed profit load? Under Equation 2.10, the CAPM approach, each risk requires a different rate of return, but a common leverage ratio is used for all lines. Under Equation 2.9, which I will call the *leverage ratio approach*, we target a common rate of return but vary the leverage requirements; in effect, all lines are scaled to the market return.

This last result might seem a little odd—aren't the measures of actual ROE by line of business different under the two approaches? The answer is yes, but this is because the approaches use different leverage ratios. Under CAPM, the industry leverage ratio $P/(zS)$ is used for *all* lines of business, whereas under the leverage ratio approach, the leverage varies by line: $P/(zS\beta)$. If you accept a single leverage ratio, then you must demand differing rates of return on equity based upon the line's beta. But if you adjust the surplus requirements *by* beta, then you can accept an *equal* rate of return on equity across all lines, and this will equal the industry rate. In practice, the distinction is largely academic—regardless of which formula we use, the resulting profit load is the same. Thus, the meanings of R_m in Equations 2.9 and 2.10 are equivalent, and we have our third important result:

3. CAPM and the leverage approach, which are based upon the covariance method, are equivalent for computing insurance profit loads.

We also obtain an additional result, a counterargument to those who suggest that surplus shouldn't be allocated to line of business for pricing purposes. If we allocate surplus in proportion to a line's beta, we obtain a profit load rule that is equivalent to that produced by CAPM. We also normalize the by-line ROEs towards the industry average, allowing the convenience of targeting a single ROE for all lines instead of varying the ROE target by line. You can obtain the same results by *not* allocating surplus, using the industry leverage for all lines and varying the target ROE according to CAPM. I agree with Feldblum that this allocation of surplus has nothing to do with *solvency* considerations, but that it produces a *pricing* rule that makes economic sense. This allocation of surplus differs from Gogol's method in that it is simply a "grossing up" of the marginal surplus requirements so that they balance back to the total.

But now we have a problem. In creating this nice link to CAPM, we seem to have lost our way back to the original ruin theory equation. According to CAPM and the leverage approach, the appropriate surplus for the industry is zS, but Equation 2.1 says this value should be $zS - R$. R is *not* small compared to zS, and there is no adjustment that brings these approaches into line. However, let us recall Philbrick's concern for "the overly simplistic binary division of the world into solvent and insolvent companies. Gradations of insolvency are important...." We could reflect this by adopting a more aggressive ruin constraint: for example, that the sum of loss and expense, *minus profit*, may not exceed premium plus available surplus (i.e., just surviving is not good enough). In this case, the needed surplus would now be zS, as per CAPM and the leverage approach. This is no longer strictly ruin theory, but it is certainly related, and our analysis provides evidence that this is a more "financially sound" approach than pure ruin theory.

That leaves us with the two final approaches, utility theory and the reinsurance method. As Feldblum noted in his paper, neither of these approaches has straightforward equations with which to work, so this part of the analysis will be less rigorous, and more brief!

The reinsurance approach is not really an independent method, but is, as Philbrick pointed out, "a powerful reality check." Presumably, reinsurers are subject to the same market forces as primary insurers, and assuming that marginal risk methods are correct for primary insurers, they should work for reinsurers also. In fact, this was the context in which Kreps presented his findings. Further, the "reality checking" feature should help both primary insurers and reinsurers calibrate their estimates from other approaches and verify that they make sense in the context of the larger market. But this is the extent of this "method's" usefulness—it cannot determine risk loads from scratch, and it is not the only way one can check the market viability of other methods.

Utility theory is a more complicated issue. Feldblum notes correctly that there is no good method for determining exactly what utility function should be used for determining investor preferences (or insurer risk loads). However, CAPM requires *as an assumption* that investors have utility functions of a certain form—specifically, risk-averse functions with known first and second moments [10]. Suppose we consider all of the assumptions required by CAPM *except* for the utility requirement, and furthermore, suppose we *assume* that investors will value risk as per the CAPM formula. The question is: what does this say about investor utility? Clearly, it still implies that investors are risk-averse, because they demand higher returns for taking on more risk. Do we really need to know anything else? CAPM can price the risk loads, so why do we need a corresponding utility method to do the same?

One might argue that we could better price the risk if we knew more specifics about the market utility function, but this seems equivalent to knowing how the market rewards different levels of risk, at least for fairly "well behaved" utility functions. One could certainly conceive of investor utility functions so complex that CAPM no longer applies, but such functions could probably be shown to fit into the framework of something like the arbitrage pricing model (APM), which is a generalization of CAPM [10]. This is not a trivial step—APM is considerably more complex than CAPM, in that it allows investors to use information other than mean and variance statistics to price risk. This strikes me as an important insight—utility theory at least holds out the potential of using more information than just the first two moments of a portfolio's probability distribution to determine investor preferences. The following chart describes this potential shift in approach:

Current Methods		*Future Methods*
moment-based	→	moments plus other data
CAPM	→	APM
simple utility functions	→	complex utility function
simple ruin theory	→	complex ruin theory

As used here, "simple" basically means "tractable and understandable." Most people involved in the field of financial research know that CAPM and related approaches are approximations (hopefully good ones) of a more complex reality. But tractability becomes less of an issue every year as computing power increases and research progresses. Understandability is a more serious issue and may slow progress more than tractability.

I don't want to pursue this direction any further in this paper—the subject would fill a book. As for a verdict on utility theory: 1) for a fairly large class of tractable utility functions, there is consistency with CAPM and related methods, so it seems unimportant that we don't actually have a method to determine what utility functions to use; and 2) even if more complex utility functions might model market preferences more accurately, there are probably other equivalent methods, like APM, that would be used in practice.

In summary, it seems clear that the five approaches have more in common than Feldblum, or even Philbrick, would admit. The key is to carefully state the initial assumptions and eliminate the various shortcuts and approximations that are so often used with these approaches. If actuaries continue to ignore covariance considerations when setting risk loads, for example, then the approaches will not agree, and many of Feldblum's and Philbrick's criticisms will be completely justified.

3. REMAINING PROBLEMS

This analysis seemingly produces a good result, in that we now have a single approach for setting profit loads that is objective, agrees with financial theory, and could be used in practice. However, the conclusion I have reached is that *none* of the five approaches deals with some very important and practical considerations, and without a resolution to these problems, we end up with risk loads that are driven largely by subjective considerations. This conclusion does *not* imply that these methods are unusable—I use a form of the leverage ratio approach in practice—but Feldblum's article might leave one with the impression that CAPM solves more subjectivity problems than it actually does. There are many problems that require further attention, but the following are examples that loom large in my mind.

What is the Industry Leverage Ratio?

This is a very important question that CAPM does not answer and that Feldblum appears to have overlooked. In fact, Feldblum seems to imply in his paper that once you have computed your return-on-premium betas, you need only use the "Kenney Rule" (2-to-1 premium-to-surplus ratio) to convert CAPM return on equity targets to profit loads! The exact leverage value is not important—the point is that Feldblum seems to be saying that $P/(zS)$ (using the above notation) is *known* for the industry, when most certainly it is not. zS is definitely not statutory surplus, nor even GAAP equity, because these accounting measures don't use components that are stated economically (e.g., reserves aren't discounted), and we cannot rely upon any given year-end snapshot of equity to be free of distortions and random fluctuations. Even if we came up with a way to measure S properly (does risk-based capital do this?—I have my doubts), what is the correct value for z? The answer must be something like "whatever the market *says* z should be," but this doesn't help us to *compute* a value for z.

No, there is only one answer to this question at the present time: $P/(zS)$ must be *selected*, giving due consideration to the amount of risk the market and company senior management are willing to bear (as correctly discerned by Kreps). Once this key leverage ratio is selected, the other calculations become possible, and it *is* key because it impacts the profit load for every individual line of business. So the most CAPM can accomplish is to compute profit load *relativities*, which is no better than ISO's approach, old or new. Perhaps this is what Glenn Meyers means when he says that CAPM requires an allocation of surplus. Strictly speaking, CAPM does not require one to know leverage ratios by line, as that is what it computes, but CAPM most certainly *does* require that one know the overall leverage ratio or, equivalently, the leverage for one line of business. It would certainly be worthwhile to try to develop ways to evaluate the choice of overall leverage ratio and its accompanying return on equity (apart from obvious *ad hoc* methods like comparisons to industry figures, or other industries with similar risk characteristics, etc.), but that is a subject worthy of a paper of its own.

Why Industry Leverage Over Company Leverage?

In my analysis, I specifically assumed that the existing large portfolio was the *industry* portfolio, rather than an individual *company* portfolio as specified by Feldblum. This difference in assumption does not affect the conclusions of my analysis *per se*, but it could produce different risk loads. Indeed, Feldblum notes that a "small- or moderate-size insurer needs a slightly larger risk load than that indicated by the industry-wide experience," in order to pick up some of the specific risk. I question this: why would an insured be willing to pay this additional charge? One could argue that a small insurer may be less "solid" than a large insurer because the small insurer is more affected by random fluctuations in experience. The risk of insolvency is higher and thus the small insurer offers a "lower quality product" and thus demands a *lower* premium. This is a simplistic argument with problems of its own, of course, but I have heard it made. Although I agree that an *insurer* may possess additional risks versus other companies, I don't see why an *insured* would pay for this difference.

Why shouldn't "equivalent" lines of business demand equal risk premiums in a competitive market? This question is almost tautological. The answer that "every insurer is different" might be a hard sell to insureds, particularly less-sophisticated insureds (e.g., as in personal lines, where the products are relatively simple risk-transfer mechanisms and are largely interchangeable between companies). A riskier insurer needs to *do something*, probably via reinsurance or a portfolio change, to "steer" its portfolio towards the market optimal portfolio that is less risky. Using a "market equilibrium"-type argument, shouldn't insureds pay only the competitive equilibrium risk charge for all interacting companies, and doesn't this mean that beta should therefore be measured against the market return as opposed to a company's overall return?

Also, do not assume that "market" means the *insurance* market—in view of overall concerns for asset/liability management, why shouldn't we measure risk against the *entire* market? Actually, this is perhaps too big a stretch—in his paper, Feldblum points out some valid reasons why insurance contracts differ from financial instruments. But surely the risk inherent in the investment portfolios varies among insurers. The extent to which the investment portfolio does not interact efficiently with the insurer's underwriting book is another risk for which insureds may not be willing to pay. This line of reasoning starts to touch on areas outside of underwriting risk. For example, insurers are exposed to asset risks that are not directly related to their underwriting risk, such as the risks associated with stocks, real estate, or venture capital. From a stockholder's perspective, these asset risks are important components of an insurer's beta; but, arguably, these forms of investment should pay their own way and should not be charged back to the policyholder in the form of a higher risk load. It would seem that only those risks that arise from the interaction of investment and underwriting that cannot otherwise be diversified away should be included in insurance risk loads. Realistic examples of this seem hard to come by: one could envision a deal to pay an insured a guaranteed rate of interest on the funds held for a large deductible account, and the rate might be higher than current Treasury rates. The insurer would certainly have a right to charge for this, but I suspect that such an arrangement would more likely be struck, with little consideration for an adequate rate, simply to get the account.

There is another reason why the distinction between industry and company risk is important. If one measures risk against a company portfolio only, it is possible that individual transactions could unduly influence the risk calculation. An example would be large assumed reinsurance contracts. Although such considerations are important to the insurer, there is still the question of how much of this cost to pass down to the insured. In his discussion of Kreps's paper, Gogol correctly identified this issue as a problem of "order dependence" (i.e., that the risk load changes depending upon when the risk is written) and developed a formula to correct for this. Similarly, using a larger market base forces the risk measurement of individual contracts closer to the margin, which equalizes risk charges and better satisfies CAPM assumptions.

There are no definitive answers to these questions. The practical effect of these concerns would be to shift an insurer's total risk load up or down, equivalent to changing the overall industry leverage ratio, and in practice this value is selected as noted above. The point is that the CAPM methodology proposed by Feldblum has not resolved these issues, although it is a very good framework within which to further discuss the problems.

How to Compute Covariance?

The fact that CAPM is a theory that applies specifically to financial securities means that assumptions will be needed to adapt the approach to measuring insurer risk. For example, Feldblum states in a footnote that CAPM "has obviated the need for quantifying covariance." This may be true for stocks, but not for insurance profit loads. As stated before, empirical profit information is not the best starting point for this calculation. Feldblum mentioned at least two problems that require attention:

- adjusting for reserve deficiencies and redundancies (i.e., getting to an accident year basis); and
- using discounted cash flow to allocate investment income.

There is a third problem, and it's a big one. To estimate the by-line betas, we need a series of historical operating ratios by line and in total in order to perform the required regression. However, what we want is an estimate of the *current* beta for a line. Doesn't that require our data to be at "current level"? Moreover, "current level" comprises a lot more here than just rates and trends—changes in mix of limits, legal climate, social conditions, and the like are much more important in an analysis of risk than in an analysis of expected cost. Add to this the numerous other calendar year distortions faced by insurance companies, and calendar year data becomes very messy indeed. It isn't clear that the most fruitful approach is to start with calendar year data and to expend a lot of effort cleaning it up. Actuaries simply have more troublesome data problems than do stock analysts in this instance!

The reviewers of this paper brought up a good point. Bringing data to current level has the effect of reducing the variance of historical loss ratios that resulted from shifting conditions, but these shifts reflect legitimate risks to the company and should be included in the cost of capital. I agree with this *to a point*. By bringing data to current level, my hope is to obtain a good measure of *process* risk. However, this procedure does eliminate valid sources of *parameter* risk that somehow must be measured and included. I prefer to separate the two measures and try to obtain a clean estimate of each. Further, some of this perceived risk *cannot* be passed on to insureds—for example, the risk due to a period of deliberate underpricing to gain market share. This is something that a company inflicts upon itself, and we cannot expect future policyholders to accept risk loads computed using past "price volatility."

Rather than a straight CAPM approach with calendar year data, begin instead with a model of current accident year losses, adjusted for all current conditions and including a measure of parameter risk. The advantage of this approach is that models of this form probably already exist for pricing and/or reserving purposes. One problem, of course, is that we must also include a measure of this distribution's covariance to the market. In practice, the only source for such information is the same kind of calendar year industry data used by Feldblum, but such data are very difficult to work with even for this more narrow purpose, and don't produce very "intuitive" results (such as the low beta for surety computed by Feldblum and noted by Philbrick). In most cases, it is necessary to *ignore* the covariance terms and to use instead a simplification that is more practical (e.g., one based upon standard deviation). It is preferable not to do this, but we must realize that this *is* an approximation to the correct answer. Further, we should continue to explore ways to better measure and incorporate covariance.

It boils down to a choice between simplifying assumptions: use CAPM with calendar year data adjusted "top-down" as best you can, or start "bottom-up" with an accident year model and reflect as many sources of risk as possible. I prefer the latter, and I presume Feldblum would advocate the former; but *both* are approximations and need more research.

4. CONCLUSION

Feldblum and the many other contributors to this subject should be congratulated and encouraged to continue the discussion. Most of what they have said has significantly advanced the state of the art in measuring risk loads. My message is directed primarily to those less familiar with these issues, and that message is 1) the show has just begun, and 2) the show to date has largely consisted of variations on a common theme. My concerns are only a sampling of the issues needing resolution—this topic should be fertile ground for inquiry for some time to come.

REFERENCES

[1] Feldblum, Sholom, "Risk Loads for Insurers," *PCAS* LXXVII, 1990, pp. 160–195.

[2] Philbrick, Stephen W., Discussion of Feldblum: "Risk Loads for Insurers," *PCAS* LXXVIII, 1991, pp. 56–63.

[3] Feldblum, Sholom, Author's Reply to Discussion of Feldblum: "Risk Loads for Insurers," *PCAS* LXXX, 1993, pp. 366–379.

[4] Kreps, Rodney, "Reinsurer Risk Loads from Marginal Surplus Requirements," *PCAS* LXXVII, 1990, pp. 196–203.

[5] Gogol, Daniel F., Discussion of Kreps: "Reinsurer Risk Loads from Marginal Surplus Requirements," *PCAS* LXXIX, 1992, pp. 362–366.

[6] Meyers, Glenn G., "The Competitive Market Equilibrium Risk Load Formula for Increased Limits Ratemaking," *PCAS* LXXVIII, 1991, pp. 163–200.

[7] Robbin, Ira, Discussion of Meyers: "The Competitive Market Equilibrium Risk Load Formula for Increased Limits Ratemaking," *PCAS* LXXIX, 1992, pp. 367–384.

[8] Meyers, Glenn G., Author's Reply to Discussion of Meyers: "The Competitive Market Equilibrium Risk Load Formula for Increased Limits Ratemaking," *PCAS* LXXX, 1993, pp. 396–415.

[9] Ferrari, J. Robert, "The Relationship of Underwriting, Investment, Leverage, and Exposure to Total Return on Owners' Equity," *PCAS* LV, 1968, pp. 295–302.

[10] D'Arcy, Stephen P. and Neil A. Doherty, *The Financial Theory of Pricing Property-Liability Insurance Contracts*, Monograph No. 15, Homewood, Illinois: Richard D. Irwin, 1988.