

ERM for Strategic Management – Status Report

Gary G. Venter, FCAS, MAAA

Abstract

Much of the push for ERM has come from regulators and rating agencies, but it is being applied in internal company decision-making as well. This paper reviews the progress made and the needs still outstanding in two key areas of application: optimal capital level for an insurer and risk-adjusted profitability of business units. The basic conclusion is that progress has been made in these areas, but more is needed. Building models is not emphasized – it will be assumed that a state-of-the art model is available. The emphasis is on using such models in decision-making.

Keywords. Enterprise risk management; capital allocation; risk measures, value of risk

“Ev’rythin’s up to date in Kansas City. They’ve gone about as fur as they c’n go!”

Oscar Hammerstein II, 1943

“You came a long way from St. Louis, but baby, you still got a long way to go.”

Bob Russell, 1948

“I’m from Missouri, you’ve got to show me.”

Willard Duncan Vandiver, 1899

1 INTRODUCTION

It would be naïve to suppose that ERM has gone about as far as it can go, but it indeed has come a long way. Some skepticism about its conclusions is still prudent, however. In particular at this point ERM is only partially able to characterize the capital need and profit adequacy of insurers. This paper reviews some of the concrete progress that ERM has made and where more is needed. Two key areas of application are reviewed: capital needs and risk-adjusted profitability. A good deal of the history and development of quantitative risk models has been spurred by the regulatory environment, but the emphasis here is on use of the models for decision making in insurance firms. Thus the shareholders’ perspective is used here, rather than the policyholders’.

2 CAPITAL NEEDS OF AN INSURER

2.1 Historical Capital Requirements

Historically the capital adequacy of an insurer was measured by its premium-to-surplus ratio. Three was thought to be a good number, but this gradually declined over time. As reserves began to grow, the reserve-to-surplus ratio was also reported, and eventually the liabil-

ities-to-surplus ratio was used, where liabilities included reserves and unearned premium. Surplus around $\frac{1}{4}$ th of liabilities might be considered sufficient, but more was better.

One weakness with this is that different types of premiums and reserves impose different degrees of risk. Inadequate premiums and reserves of course pose more risk, but different lines of business, territories, etc. do as well. Risk-based capital (RBC) was a response. Under that scheme, various classes of assets and liabilities all get their own risk factors. Sometimes an adjustment for the postulated independence of different risks is also included.

While an improvement over the liabilities-to-surplus ratio, RBC does not directly measure risk either. Inadequate premiums and reserves can still generate lower capital charges than adequate levels of the same accounts, for instance. And there is still an unmeasured difference in the risk of different companies that write the same lines of business, due to differences in understanding of the business, different pricing and reserving practices, etc.

2.2 Economic Capital

RBC is then a natural setup for quantitative risk modeling. If you can build a model that quantifies the risk of loss to an insurance company from all sources, perhaps that model can be used to give a probabilistic answer to the question of how much capital an insurer needs to operate prudently. The first answer in that direction is economic capital. This is usually thought of as the amount of capital needed to get the one-year probability of ruin below some target threshold. Finally, here is a capital need that is quantified probabilistically.

There are problems with economic capital, however, starting with what the target threshold should be. Usually the goal is formulated as keeping the ruin probability below α , where $1/\alpha$ is a large round number. For instance, retrospective studies have found that some highly rated corporate bonds have a default probability of $3/10,000$, which suggests taking $1/\alpha = 3333$. Values between 2000 and 3333 are typical. In practice it is not considered prudent for economic capital to be more than actual capital, so α is chosen to be as large a round number as possible that will keep actual capital above economic capital.

While sometimes criticized for being just a single quantile on the probability distribution, for regulatory purposes this can be thought of as an advantage of economic capital. There is often an assumption that regulatory information becomes public information, so a regulatory framework should provide useful information for judging risk, but not be too revealing to

competitors. While information paucity can be useful in a regulatory framework, it is less so for strategic management. Alternative ways of quantifying capital adequacy may thus be more valuable to company management.

Another drawback of economic capital is that it is beyond the capacity of current ERM models to quantify. Just the catastrophe component of loss risk for exposed insurers is not well known at levels like 1 in 2000. The different catastrophe models, all with strong science behind them, diverge greatly even at lower probability levels. Loss reserve inadequacy, failure to understand and control risk, fraud, and other risk issues are prominent in the list of causes of insurance company failure, and these can be even more difficult to quantify.

When companies publish loss levels at remote probabilities, they are usually modeling the standard deviation of results fairly carefully, then assuming a distributional form, like lognormal, to read off the probabilities. The distribution assumed can make a large difference in the projected loss levels, however.

Due to the modeling difficulties at remote probabilities, more sophisticated companies are beginning to express capital as a multiple of lower loss levels. For instance in a recent financial report, Swiss Re said that its capital was about 3.5 times its 1-in-100 probability point and 3 times the 1-in-200 level. That means that if the 1-in-200 year were to occur, Swiss Re would lose 1/3rd of its capital. Other companies have tried related measures. For instance, a Munich Re representative once stated that if it had two 1-in-100 years in a row, it would lose 2/3rd of its capital. Such quantification is more realistic than the economic-capital approach, but it is more difficult to interpret. Companies can track such ratios over time and compare with statements of other companies, so benchmarking is one possibility. Another use of these ratios could be to determine at what probability the company would have to stop writing new business, or to non-renew existing business.

2.3 Capital Needed Practically

Relating actual capital to several loss probability levels reflects another reality: the probability level does not determine the capital need, but is calculated as a check after the capital has been established. Determining the capital an insurer needs goes beyond selecting a probability level. In most general terms, the ideal capital can be related to the market value of the insurer. Maximizing the franchise value, i.e., market minus book, could be a good target.

The capital level that would meet such criteria would not be a simple risk measurement, but would have to take into account numerous market responses. The frictional costs of holding capital, such as taxation of investment income, would limit the ideal capital from going too high, and other frictional costs, like the costs of raising new capital and financial distress costs (imposed when capital is perceived by regulators and rating agencies to be too low) push up the ideal capital level. Also customer perception and risk attitudes become important here. Studies have found that customers tend to be more risk averse in their insurance purchases than market principles might suggest, probably because frictional costs tend to make them non-diversified in their insurance purchases. For an insurer, some way of gauging the insurance market reactions to various capital levels, such as studying reasons that quotes get accepted and rejected, can help shed light on this impact.

There is a fair amount of ongoing research on the optimum capital level of an insurer from the point of view of maximizing franchise value. Some of this is reviewed in Appendix A. The main conclusion is that progress is being made, but this is still an unsolved problem.

The ideal capital level for an insurer is thus not a purely risk-based measurement that can be done with an ERM model. The models can help quantify the risk inherent in the business in comparison to the capital decisions that have been made, however. They can also permit what-if analysis along the lines of finding how the capital need might change if the same relationship of risk and capital is maintained, but other strategic changes are made, such as changing reinsurance purchased, writing different business, etc.

In summary, being able to say that capital is 4 times the 1-in-100-year loss demonstrates a far greater understanding of the company's risk than saying that capital is 1/4th of liabilities. Thus in understand capital and the risk to it, ERM has come a long way. Yet ERM as it exists today is not able to definitively answer the question of how much capital an insurer should hold. The optimal capital needed to maximize franchise value may be the way to answer that question. What ERM can do now is to provide various measures of the risks of the firm in comparison to the capital held, and test those risks against strategic alternatives.

3 RISK AND PROFITABILITY OF BUSINESS UNITS

3.1 Historical Risk Measures

So-called silo risk measures were used historically for different lines and categories of ex-

posure. These are risk measures that can be used for some risks, but not all, or that became standard in some lines. Premium volume is a crude risk measure in any line. Exposures in force are used when available, as are limits in force and number of policies. In property lines probable maximal loss, or PML is a widely used risk measure. That sounds perhaps probabilistic, but has always been hard to pin down to a clear probabilistic definition. Investment risk is quantified relative to the performance of similar market portfolios, as well as with the greeks, which are sensitivity measures usually abbreviated by greek letters.

3.2 ERM Risk Measures

All the silo measures are still in use, but it is increasingly clear that they are hard to compare across business units and risk categories. Statistical risk measures like standard deviation have been widely used in quantifying financial risk, but due to the skewness of losses, this can be a somewhat limiting measure for insurers. ERM has introduced measures like value at risk (VaR) and tail value at risk (TVaR), also called conditional tail expectation. These are defined relative to probability levels. At level α , VaR is just the α^{th} quantile of the distribution of the account being quantified. Thus ERM has two new names for quantile: VaR and economic capital. TVaR is the conditional mean of values greater than the VaR point.

VaR and TVaR have their critics. VaR is not subadditive. That is, the VaR at a fixed probability level for a number of risks combined can be greater than the sum of the VaRs of the individual risks. This can be a problem when trying to quantify the benefits of diversification. However for most ERM applications subadditivity is not necessary. Another problem with VaR is that it is just a single point on a distribution, so it does not give very much information. It is also difficult to allocate reasonably to business units, which is discussed in more detail later.

TVaR is sometimes criticized for including losses excess of insolvency, which presumably do not affect shareholders. This is not really a problem, however, as these losses can hurt policyholders, whose attitudes affect growth and profits, and hence do affect shareholders. A more serious issue with TVaR is that it is linear for losses in the excess region. This is contrary to usual risk attitudes, where a loss twice as big is regarded as more than twice as bad.

Alternative risk measures can address that problem. For instance, weighted-TVaR, or WTVaR, uses an adjusted probability distribution to calculate the conditional expectation.

The probabilities of larger losses are increased, which gives them more weight in the average. They are still less likely than smaller losses, but relatively more likely than with the actual probabilities. A similar adjustment is risk-adjusted TVaR, or RTVaR. This is defined as the conditional tail mean plus some fraction of the conditional tail standard deviation. It can be thought of as a pricing of the tail risk, using a standard deviation loading. It was introduced by Furman and Landsman [1], who call it the tail standard deviation premium.

There is a wide variety of probabilistic risk measures. Some are discussed in Appendix B. The use of risk measures in evaluating the profitability of business units is addressed next.

3.3 Allocation of Risk

3.31 Basic Methods

Units with higher risk should earn higher profits to justify taking that risk, but getting that general rule down to specifics can be difficult. A standard ERM approach is to allocate economic capital to business unit based on risk, and then divide the unit profit by the allocated capital to get a risk-adjusted return. The allocation might be done by the proportional method: VaR at a high level is computed for each unit, these are added up, and each unit's share of the total is taken to be its share of the economic capital. This is in part developed from the idea that firm capital is determined by a high VaR level, so that is the risk measure to use to allocate capital to each unit.

A problem with this method is that it is not marginal, and so the risk-adjusted profit calculation does not maintain the financial principle of comparing marginal profits with marginal costs. A marginal allocation would charge each unit with the additional capital it requires. Or allocation could be incremental marginal: each small bit of business of the unit would be charged with the additional capital needed to keep the overall company risk measure the same with or without it.

Of course a problem with incremental marginal allocation is that it might not add up to the total risk. This has led to allocation in proportion to the marginal impacts, instead of using the marginal impacts themselves. Another problem is that it treats each increment as if it were the last one in. An alternative that addresses this is the Shapley method from game theory, in which each line is allowed to form hypothetical companies by forming coalitions with other lines. The allocation is then the average impact of the line in all such coalitions. This

gets cumbersome, however, and still does not give the exact marginal impacts. Fortunately, there is an alternative that in some cases will result in incremental marginal impacts that add up to the entire risk measure for the company, so no proportional allocation is needed.

3.32 Incremental Marginal Decomposition

This alternative is what Patrik et al. [2] call the Euler method, and Venter et al. [3] call marginal decomposition. It works only for risk measures that are homogeneous¹, i.e., risk measures $\rho(Y)$ such that $\rho(aY) = a\rho(Y)$ for any constant $a > 0$. This holds for standard deviation but not for variance². It turns out that VaR and TVaR defined for probability level α are homogeneous, but they are not if they are defined excess of a fixed monetary amount. For a homogeneous risk measure ρ on a random variable Y that is the sum of component random variables X_1, \dots, X_n , the marginal decomposition of $\rho(Y)$ to the X 's is defined for X_k by:

$$r(X_k) = \lim_{\varepsilon \rightarrow 0} \frac{\rho(Y) - \rho(Y - \varepsilon X_k)}{\varepsilon}$$

This looks at the reduction in the risk measure for Y from reducing Y by reducing X_k by a small proportion ε , say by using a small quota-share reinsurance treaty. This is the incremental impact on $\rho(Y)$ from the last little bit of X_k . Dividing by ε scales up this incremental change to all of X_k . A theorem of Euler shows that these marginal impacts add up across the components to the whole risk measure $\rho(Y)$. This is a marginal decomposition of the risk measure to its component units. The risk measure is not allocated to units. Rather the impact of the business unit on the overall company risk measure is calculated for each unit.

Marginal decomposition is closely linked to a more general method of allocation from game theory called the Aumann-Shapley method. This is used in cost accounting to allocate common costs, like those from a production line that is used in the manufacturing of several products. See Billera et al. [4] for further details. For homogeneous risk measures, the Aumann-Shapley method reduces to marginal decomposition. For other measures it differs in ways that are probably not appropriate for insurance application, as it averages costs over zero to full production, which is not usually the range of strategic choices for insurers.

¹ Technically, positively homogeneous of degree 1

² Which is homogeneous of degree 2.

3.33 Co-Measures

Venter et al. [3] link marginal decomposition to the method of co-measures. If a risk measure can be defined in a specified way as a conditional expected value, it has a purely additive decomposition into co-measures defined from that expected value. In particular, when $\rho(Y)$ is expressed as a conditional expected value:

$$\rho(Y) = E[\sum_i h_i(Y) L_i(Y) \mid \text{condition on } Y],$$

where the h_i 's are additive functions, i.e., $h(V+W) = h(V)+h(W)$, and the L_i 's are any functions for which this conditional expected value exists, then the co-measure for component k is the same formula but with Y replaced by X_k in the argument of the h functions. That is, the co-measure r is defined by:

$$r(X_k) = E[\sum_i h_i(X_k) L_i(Y) \mid \text{condition on } Y]$$

By the additivity of the h 's this satisfies $\rho(Y) = \sum_k r(X_k)$. As an example, excess tail value at risk (XTVaR) excess of level b can be defined as:

$$\rho(Y) = E[(Y - EY) \mid Y > b]$$

Now $h(X)$ is $X - EX$, $L(Y) = 1$, the condition is $Y > b$, so $r(X_k) = E[(X_k - EX_k) \mid Y > b]$. This is an additive decomposition, but this ρ is not homogeneous, so this is not a marginal decomposition. However for XTVaR defined excess of probability level α :

$$\rho(Y) = E[(Y - EY) \mid F(Y) > \alpha]$$

The co-measure $r(X_k) = E[(X_k - EX_k) \mid F(Y) > \alpha]$ is marginal and ρ is homogeneous. The marginal co-measure for TVaR is $r(X_k) = E[X_k \mid F(Y) > \alpha]$. The marginal co-measure for VaR is $r(X_k) = E[X_k \mid F(Y) = \alpha]$, but this can be difficult to estimate. WTVaR has the same marginal decomposition as TVaR, just using the adjusted probabilities in the expectations.

Since it is homogeneous, standard deviation has a marginal co-measure. In Venter et al. [3] it is shown that this co-standard deviation is $r(X_k) = \text{Cov}(X_k, Y) / \text{Std}(Y)$. This can also be used to decompose RTVaR. For instance, taking a loading of $1/2$ of a standard deviation, $\rho(Y) = E[Y \mid F(Y) > \alpha] + 1/2 \text{Std}(Y \mid F(Y) > \alpha)$. Then the co-measure is $r(X_k) = \text{co-TVaR}_\alpha + 1/2 \text{Cov}(X_k, Y \mid F(Y) > \alpha) / \text{Std}[Y \mid F(Y) > \alpha]$.

Another advantage of marginal decomposition is that it meets the requirement of Tasche [5] for suitable allocation. Tasche points out that if capital is allocated by a risk measure in order to calculate risk-adjusted return by line, you would like to be able to conclude that growing a line with a higher-than-average return will increase the return for the whole company. Venter et al. [3] show that marginal decomposition meets this requirement. This is only guaranteed to hold for proportional growth, like taking a higher percentage of business already written, but it usually also holds for adding business units similar to those already held.

3.34 Risk and Profit

Probably the most reasonable risk measures discussed above for use in calculating risk-adjusted profitability are WTVaR and RTVaR. They avoid the problems with TVaR being linear and VaR being just a single point. Another advantage is that they conceivably relate to the value of risk transfer. Under marginal decomposition, the risk attributed to a business unit is the increase in firm risk due to that business unit. But if profit is to be compared to the risk measure to determine if it is high enough to meet the risk taken, the measure should in some way reflect the value of the risk.

RTVaR is based on a simple historical risk-pricing mechanism, the standard deviation loading. While not really meeting modern financial standards as a risk-pricing methodology, standard deviation loading is used by insurers and reinsurers as a convenient benchmark. It is somewhat weak in the insurance setting as it is a purely quadratic measure, so does not capture the heavy tail of most insurance business. This is sometimes compensated for in practice by layering the business and using higher percentages of the standard deviation for extreme tail covers, perhaps 50% instead of 30%. WTVaR, as discussed further in Appendix B, can capture higher-order effects, but it is harder to calculate and it is not clear how to apply it consistently in a multi-line company. Further experimentation with WTVaR in practical settings could help clarify these issues.

A related issue is what probability level to use. When capital is thought of as there to support the extreme loss scenarios, allocation is often done on the same basis. But we have seen that capital does not come from risk measures and it can be related to any risk measure, not just extreme tail measures. For instance, in the worst year in 7 a company might not want to lose more than 5% of capital. Then capital is 20 times the 1-in-7 loss. A company can lose value when losses occur that are well below the extreme thresholds, and the risk of such

losses is charged for. It appears, in fact, that failure to meet plan is punished in the stock market, sometimes far more than the financial shortfall. Thus a reasonable probability level might be the probability of not meeting plan, and the shortfall from plan could be the random variable measured. This is a much lower probability level than those often used.

Under TVaR such a low probability level can tend to attribute too little risk to the extreme loss scenarios. This is part of the problem of the linearity of TVaR. RTVaR and WTVaR should be more appropriate at the low probability thresholds.

3.4 Problems with and Alternatives to Allocation

Capital allocation is in a sense artificial, in that capital allocated is not actually assigned to the business units. The units can have losses higher than their allocation and use more capital than they have been allocated. Allocation also is somewhat arbitrary, as it involves a selection of risk measure and other parameters, such as probability level. Also the idea that each unit will have the same target return on allocated capital is suspect, since this implies that the allocation has really captured the value of risk. Really this is making capital allocation into a risk-pricing methodology. To evaluate this, an excursion into pricing theory is required.

A popular alternative to capital allocation is what Mango [6] calls capital consumption. This follows the basic setup of Merton and Perold [7] who suggest a value-added approach to measuring risk-adjusted profitability. Instead of allocating capital and computing a risk-adjusted return, they advocate calculating the cost to the firm of bearing the risk of each business unit, and then subtracting this from unit profits to get the value added. The cost of bearing the risk of a unit is the value of the unit's right to access the capital of the firm. The unit can use any or all of the capital, depending on how big its losses are. The right to access capital is thus a contingent claim, and so options-pricing methodology can be applied.

In Merton and Perold's examples, typical options-pricing formulas are used. But the insurance risk case is more complicated. There is no fixed date when capital will be attached: whenever the unit runs out of the premiums it charged, and the investment income they generated, it can start using the firm's capital. Then it does not stop until all its losses are paid, whenever that is. The loss distributions are usually heavy tailed as well. Contingent-claims pricing theory says that the price for the option should be the mean of a martingale transform of the loss process. Møller [8] discusses martingale transforms for the compound

Poisson process, and some of this is reviewed in Appendix B.

Another way to view this option is that the firm implicitly provides each business unit each year with stop-loss reinsurance with a retention at break-even. Then the cost of carrying the unit is the value of this stop loss. Reinsurance pricing methodology can then be used to get a handle on the value of that cover. This would usually be less than an open-market quote for such a cover, as the market price could have substantial provisions for asymmetric information and behavior disincentives that an internal cover would not need.

One issue with the value-added framework is what measure of profit to use. Year-to-year fluctuations can distort the measure, so some expected value would be more appropriate. Actually the profit the firm gets is in the form of a contingent claim: the firm gets all the profit if it is positive and none otherwise. Thus the profit measure from which the capital cost should be subtracted is an options price itself. But the combination of the option of the firm and the option of the unit comes down to the firm paying all the losses and getting all the profit. The net position is thus not a contingent claim. Thus the value of the unit comes down to regular risk-capital asset pricing, not options pricing.

Capital consumption avoids the problems of artificial and arbitrary allocation but it cannot dodge the issue of value of risk transfer. A martingale transform methodology might still apply, and there may be useful information in separating out the values of the firm and unit options, but the covariance with the market, frictional costs of capital, and other typical risk-pricing issues cannot be avoided by using the capital consumption methodology.

4. CONCLUSIONS

ERM is able to quantify the risk of an insurance business through multiple risk measures (Appendix B) and compare these with capital. It is not yet able to say what the target capital level should be, but the work on optimal capital in Appendix A is promising. Measuring risk-adjusted profit is really risk pricing, and so the pricing approaches discussed in Appendix C provide the long-term direction for this. Risk measures that are pricing related, like RTVaR and WTVaR are most likely to approximate true risk pricing. The capital consumption framework avoids some of the artificiality of capital allocation in this regard.

References

- [1] Furman, Edward and Zinoviy Landsman 2006 “Tail Variance Premium with Applications for Elliptical Portfolio of Risks,” *ASTIN Bulletin*, 36:2, 433-462.
- [2] Patrik, Gary, Stefan Bernegger and Marcel Beat Rüegg 1999: “The Use Of Risk Adjusted Capital To Support Business Decision-Making,” *CAS Forum*, Spring
- [3] Venter, Gary, John Major and Rodney Kreps 2006 “Marginal Decomposition of Risk Measures”: *ASTIN Bulletin*, 36:2, 375-413.
- [4] Billera, Louis J, David C. Heath and Robert E. Verrecchia 1981 “A Unique Procedure for Allocating Common Costs from a Production Process,” *Journal of Accounting Research*, 19:1, 185-196.
- [5] Tasche, Dirk 2000 “Risk contributions and performance measurement,” Report of the Lehrstuhl für Mathematische Statistik, T.U. München.
- [6] Mango, Donald 2005. “Insurance Capital as a Shared Asset,” *ASTIN Bulletin*, 35:2, 471-486.
- [7] Merton, Robert and Andre Perold 1993, “Theory of Risk Capital in Financial Firms,” *Journal of Applied Corporate Finance*, 6:3, Fall, 16-32.
- [8] Møller, Thomas 2004 “Stochastic Orders in Dynamic Reinsurance Markets,” *Finance and Stochastics*, 8, November, 479-499", also see 2003 *ASTIN Colloquium Papers*.
- [9] De Finetti, Bruno 1957 “Su un’ Impostazione Alternativa della Teoria Collettiva del Rischio.” *Transactions of the XVth International Congress of Actuaries* 2, 433-443.
- [10] Gerber, Hans U. and Elias S. W. Shiu 2006 “On Optimal Dividend Strategies in the Compound Poisson Model,” *North American Actuarial Journal* 10:2, 76-93.
- [11] Bather, J. A. 1969 “Diffusion Models in Stochastic Control Theory.” *J. Royal Statistical Society A.*, 132, 335-352.
- [12] Asmussen, S., B. Højgaard, and M. Taksar 2000 “Optimal Risk Control and Dividend Distribution Policies: Example of Excess-of Loss Reinsurance for an Insurance Corporation.” *Finance and Stochastics*, 4:3, 299-324.
- [13] Froot, Kenneth, D. Scharfstein, and J. Stein 1993 “Risk Management: Coordinating Corporate Investment and Financing Policies,” *Journal of Finance*, 48, 1629-1658.
- [14] Peura, S. “Essays on Corporate Hedging.” Academic dissertation, University of Helsinki (2003).
- [15] Froot, Kenneth 2003 “Risk Management, Capital Budgeting and Capital Structure Policy for Insurers and Reinsurers.” *NBER Working Papers* 10184, National Bureau of Economic Research, Inc.
- [16] Kahneman, Daniel, and Amos Tversky 1979, “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica*, Vol. 47, No. 2, March, 263-291.
- [17] Major, John 2007 “Measuring the Market Value of Risk Management: An Introduction to FLAVORED Models” in *Enterprise Risk Management for Property & Liability Insurance Companies*, Paul J. Brehm, Geoffrey R. Perry, Gary G. Venter, and Susan E. Witcraft, editors, Guy Carpenter & Company, LLC, New York.
- [18] Myers, Stewart C., and Nicholas S. Majluf, 1984, “Corporate Financing and Investment Decisions when Firms Have Information That Investors Do Not Have,” *J. of Financial Economics* 13, 187-221.
- [19] Venter, Gary G. 1991, “Premium Implications of Reinsurance without Arbitrage,” *ASTIN Bulletin* 21:2, 223-230.
- [20] Rubinstein, M., 1973 “A Comparative Statics Analysis of Risk Premiums,” *Journal of Business* 46, 605-615.
- [21] Chung, Y. Peter, Herb Johnson and Michael J. Schill 2006 “Asset Pricing when Returns Are Nonnormal,” *Journal of Business* 79, March, 923-940.
- [22] Hung, Chi-Hsiou 2007 “Momentum, Size and Value Factors versus Systematic Co-moments in Stock Returns,” WP-105, Apr. 07 ISSN: 1749-3641 (Online).
- [23] Buhlmann, Hans 1984 “The General Economic Premium Principle,” *ASTIN Bulletin* 14, 13-21.
- [24] Myers, Stewart C. and James A. Read Jr. 2001 “Capital Allocation for Insurance Companies,” *Journal of Risk and Insurance* 68:4, 545–580.
- [25] Phillips, Richard D., J. David Cummins and Franklin Allen 1998 “Financial Pricing of Insurance in the Multiple-Line Insurance Company,” *The Journal of Risk and Insurance*, 65:4. 597-636.
- [26] Gründl, Helmut and Hato Schmeiser 2007 “Capital Allocation for Insurance Companies—What Good Is It?” *The Journal of Risk and Insurance*, , 74:2, 301-317
- [27] Doherty, Neil A., and James R. Garven 1986 “Price Regulation in Property-Liability Insurance: A Contingent-Claims Approach,” *Journal of Finance*, 41, 1031-1050.

Appendix A – Optimal Capital

Since De Finetti [9], optimal capital for an insurer has been phrased as finding the capital strategy that over time would maximize the expected present value of cash flows to shareholders. Although they can take various forms, it is customary in this literature to refer to these cash flows as “dividends.” The expected present value is viewed as the value of the firm, which is thus the statistic to be optimized. Optimizing value in this context is part of the subject of stochastic control, which seeks to find strategies to optimize some function of an ongoing process through manipulation of available controls. The early papers on optimal capital used dividend strategy as the control parameter, so would seek a strategy that would optimize the expected present value of the dividends. Paying too much too soon would likely lead to slow growth or even insolvency, but delaying payment would reduce the present value. A recent paper in this tradition is Gerber and Shiu [10], who solve this problem for a compound Poisson process with severity a mixture of exponential distributions.

Another potential control variable for an insurer is reinsurance ceded. Several authors have extended the stochastic control problem for firm value to include both dividend policy and reinsurance. Initially this just looked at ceding proportional reinsurance when capital levels were low enough, e.g., as in Bather [11]. Ceding proportional reinsurance is like writing less business, so this provides a form of a premium-writing control variable. Otherwise this literature takes premiums as a given. Later papers, like Asmussen et al. [12] consider using excess-of-loss reinsurance, which in effect changes the severity distribution. The pricing rule for the excess-of-loss cover becomes one of the conditions of the optimization, and solving under more realistic pricing rules is one possible direction of this literature.

This actuarial approach takes the capitalization of the firm as a starting point, and usually finds that the firm should let profits build up until an ideal level of capital is reached, then dividend out any excess capital. It is closely related to earlier actuarial exercises in minimizing ruin probability, since too high a probability of early ruin reduces the expected present value of the dividend stream. However this literature does not consider the possibility that the firm can raise additional capital if its capital level gets too low.

De Finetti was writing in the late 1950’s. The financial literature of that time was dominated by the ideas of Modigliani and Miller, who assumed that firms can raise new capital without imposing costs on existing shareholders. One of their conclusions was that risk

management, such as reinsurance, and capital structure are irrelevant to firm value. So while actuaries were essentially assuming that the cost of refinancing was infinite, financial economists were assuming it was zero. The truth, of course, is somewhere in-between.

Froot et al. [13] developed the concept of costly external finance, and found that risk management and capital policy then can indeed make a difference. Peura [14] brings the idea of external finance at a finite cost into the actuarial approach. Froot [15] looks at capital issues specific to insurance companies, particularly the fact that policyholders tend to be non-diversified in their insurance purchases and hence tend to be more risk adverse towards insurance company failure than capital-market theory might suggest. This increases the capital need and the value of risk management for insurance companies. Kahneman and Tversky [16] give some support to this idea through studies of decision-making under uncertainty.

Major [17] explores ways to bring these ideas into the stochastic-control actuarial framework. Unfortunately, there is only a fairly sparse literature on policyholder risk aversion, so it is not clear how best to represent it. Major does this by assuming that if capital drops below a target level, the profit rate reduces substantially. Also, while the general concept of costly external finance is well-supported, the specifics of the actual costs and how they relate to current capital levels are still under development. Myers and Majluf [18] provide a potentially useful framework for modeling the cost to existing shareholders of new capital infusions. This may help quantify the costs of external finance.

The recent optimal capital papers use an approach to optimization known as Bellman's method. Further advances using it may be possible, particularly in the mix of forms of reinsurance, the pricing rule for reinsurance, how to represent policyholder risk aversion, and how to formulate the costs of external finance, especially in relationship to capital impairment. However it may be easier to incorporate other complex modeling of insurer finances, such as variable growth of premiums by time and line, cyclical profitability, complex reinsurance deals, etc. if the optimization framework is abandoned in favor of measuring relative changes to the franchise value from different strategies.

Modeling franchise value is a promising tool in insurer capital management. Many advances have been made, but a comprehensive model of the value of capital is yet to come.

Appendix B – Risk Measures

B.1 Background

Using risk measures can provide a consistent way of comparing the various types of risk a company may have, but different measures have different implications for comparative risk.

B.1.1 Which Risk Measure?

A recent focus has been risk measures for tail losses. In part, this arises because capital is needed to cover extreme events or combinations of events. However tail risk is not the only financially relevant risk: a company's capital can be considered insufficient if it is too low to take advantage of profitable business opportunities.

For a company whose capital is sufficient, there is a chance it could become insufficient. Any risks that contribute to this are of concern. Essentially, a risk for which a price is charged is relevant. For investments, this includes any risk for which a return greater than risk-free is needed. These issues suggest that tail-risk measures at high probabilities miss some risk that is important for management purposes.

B.2 Classification of Risk Measures

Broadly speaking, risk measures can be classified as moment based, tail based, or transformed distribution measures. These categories, however, can overlap.

B.2.1 Moment based measures

Using moments to measure risk is common. Standard deviation quantifies distance from the mean, but does not distinguish favorable from unfavorable deviations. Skewness shows whether favorable or unfavorable deviations are more likely, and which are likely to be larger. The semi-standard deviation uses unfavorable deviations only, so has some of the advantages of both standard deviation and skewness.

Moment-based measures have the advantage that they incorporate the whole distribution of results, or at least adverse results. Negative and fractional moments are possible. Generalized moments are expected values of any function of the variable, like $E[Ye^{cY/EY}]$. This uses the whole distribution, but emphasizes the tail, which can give a useful perspective.

B.2.2 Tail based measures

A number of popular risk measures look only at the tail of the distribution. These include:

- Probability of default

- Value at risk (VaR)
- Tail value at risk (TVaR or CTE = conditional tail expectation)
- Excess tail value at risk (XTVaR = TVaR – EX)
- Risk-adjusted tail value at risk (RTVaR)
- Expected policyholder deficit (EPD)

EPD at probability level α is $(1-\alpha)[TVaR_\alpha - VaR_\alpha]$. The idea is that if capital is carried at VaR, then TVaR – VaR is the average uncovered loss in default. Multiplying by the probability of default makes it an expected loss. However α does not have to be the probability of default to make this work. EPD can be calculated at any probability level. However EPD at the probability of default is significant to policyholders. Empirical work suggests that premium discounts of up to 20 times this expected deficit are needed to induce purchase of the insurance, absent an effective default fund. Thus insurers are motivated to keep EPD low.

The distinction between moment and tail measures is not always a bright line. In a sense, semi-standard deviation could be considered a tail-based measure, because it uses only results worse than the mean. If p is the probability that results are better than the mean, then TVaR at p is the average of results worse than the mean, so is closely related to semi-standard deviation. Often high probabilities are selected for tail measures, but this is not necessarily appropriate. TVaR at the probability of a financial loss is a risk measure that makes sense as the average of the losses when there is a loss.

B.2.3 Transformed distribution measures

Transformed distribution measures change the probabilities, giving more weight to adverse results, then take the mean or use some other risk measure with the transformed probabilities. They use the entire distribution of loss events, but put more weight in tails by increasing the probabilities of large losses. Some well known transforms are the Esscher transform, the Wang transform, and the proportional hazards (PH) transform.

There are two general classes of probability transforms: transforms of the probabilities of financial results, and transforms of the probabilities of the underlying events that lead to the results. The two classes turn out to have different properties.

B.2.3.1 Transforming Probabilities of Financial Results

One type of risk measure based on transforming result probabilities is spectral measures. These are of the form $\rho(Y) = E[Y \cdot \eta(F(Y))]$ for some nonnegative scalar function η . TVaR is actually a spectral measure. It is defined by $TVaR_q = E[Y | F(Y) > q] = \int_{y > F^{-1}(q)} y f(y) / (1 - q)$

dy, so η is given by:

$$\eta(p) = \begin{cases} 0, & p \leq q \\ 1/(1-q), & p > q \end{cases}$$

If η is an arbitrary step function, you get linear combinations of TVaR. For example,

$$\eta(p) = \begin{cases} 0, & p \leq q \\ 1/(r-q), & q < p \leq r \\ 0, & r < p \end{cases}$$

results in a measure $\rho = [(1-q)/(r-q)]\text{TVaR}_q - [(1-r)/(r-q)]\text{TVaR}_r$. This might be referred to as the risk of a “managed layer.” This measure is not coherent because it ignores the risk in the upper r tail, and therefore fails subadditivity. It can be viewed as a blurred VaR at $(q+r)/2$, as it approaches that VaR as $q \rightarrow r$. Usually allocation of VaR in practice allocates the mean in a symmetric interval around VaR, so is actually an allocation of this measure.

Another spectral measure is given by:

$$\eta(p) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{p-q}{\sigma}\right)^2\right).$$

This is a blurred VaR using a normal distribution spread. It, like VaR, is not coherent, failing subadditivity. In the limit, as $\sigma \rightarrow 0$, the η function becomes a Dirac delta-function at $p=q$ and the risk measure ρ becomes the VaR.

Closely related to spectral measures are distortion risk measures. These start with a non-decreasing distortion function g to and from the unit interval with $g(0) = 0$ and $g(1) = 1$. Such a function is a cdf on the unit interval. Usually $g(p) \geq p$ makes sense in applications, although it is not a formal requirement. If $S(x) = 1 - F(x)$ is the survival function, then the distortion measure using g is defined as $\rho(X) = \int_0^\infty g[S(x)]dx$.

As an example consider $g(p) = p^a$. In general $g(p) = p^a$, $0 < a < 1$, is the proportional hazards, or PH, transform, so called because it changes $\log S(x)$ by a factor. Any distribution function on the unit interval can be used for g , for instance the generalized beta. Only some parameters give $g(p) \geq p$, however.

To understand distortion measures, first note that $g[S(x)]$ is itself a survival function, so denote it as $g[S(x)] = S^*(x) = 1 - F^*(x)$. Then recall that the full integral of the survival function is the mean. This is usually proved by looking at the limited expected value at y which is $\int_0^y xf(x)dx + yS(y)$. Using integration by parts, this is $xF(x)|_0^y - \int_0^y F(x)dx + yS(y)$, which simplifies to $yF(y) - \int_0^y [1 - S(x)]dx + yS(y)$ which in fact is just $\int_0^y S(x)dx$.

Thus $\rho(X)$ is the expected value of X using the distorted distribution F^* . That can be expressed as $\rho(X) = E^*(X) = \int_0^\infty xf^*(x)dx$. Here $f^*(x) = -\partial g[S(x)]/\partial x = g'[S(x)]f(x)$. This shows that $\rho(X)$ is a spectral measure, with $\eta(p) = g'(1-p)$.

The point of having $g(p) \geq p$ is to get $\rho(X) = \int_0^\infty g[S(x)]dx \geq \int_0^\infty S(x)dx = E(X)$, so if ρ is used to price the risk, the risk load would never be negative.

Another well-known distortion measure is the Wang transform with parameters a, b :

$$g(p) = 1 - T_a[\Phi^{-1}(1-p) - b]$$

Here T_a is the t-distribution function with a degrees of freedom, a not necessarily an integer, and Φ is the standard normal distribution. In empirical tests Wang found that b around 0.45 and a around 5.5 closely matched market pricing for corporate bonds and cat bonds.

Under the idea that risk worth pricing is risk worth measuring, the pricing principles are promising risk measures. They use the entire distribution of results even though they put more emphasis on tail risk. Transforming the probabilities of end results is a direct way to do this, but it has one disadvantage – it is not arbitrage-free.

B.2.3.2 Transforming Probabilities of Underlying Events

Transformed probabilities are often thought to produce arbitrage-free pricing, but this is only so for transforms of the probabilities of underlying events, not for transforms of the probabilities of results of deals. In the insurance case the underlying events are ground-up claim counts and costs. There is an additional technical requirement on these transforms: the zero-probability events have to be the same pre and post transform.

For a compound Poisson process, the required transformed frequency and severity probabilities can be calculated in a coordinated fashion. With Poisson frequency λ and density $g(y)$ for the loss size variable Y , this method uses a function $\phi(y)$, with the only restriction

that $\phi(y) > -1$ for all positive losses y . The transformed frequency parameter is then $\lambda^* = \lambda[1+E\phi(Y)]$ and the severity density gets transformed to $g^*(y) = g(y)[1+\phi(y)]/[1+E\phi(Y)]$.

This theory works in a single-period framework and also generalizes to continuous processes. The transformed process is required to be a martingale, which means it has no expected growth. This can be accomplished by making the transformed mean equal to the mean plus profit loading from the original process.

Two interesting choices for the $\phi(y)$ function come from the minimum martingale transform and the minimum entropy martingale transform of the theory of pricing in incomplete markets. In a complete market there is often a perfect hedging strategy available which is associated with a single specific martingale transform. In an incomplete market where this is not possible, it might be desirable to find the hedging strategy which will minimize the variance of the payoff risk. The martingale that produces this strategy is the minimum martingale transform. Thus it is minimal in the sense of quadratic risk.

The relative entropy between two measures P and Q is $E_P[dQ/dP \log(dQ/dP)]$. This is a distance of a sort, as it is zero if $P=Q$ and is otherwise positive. However it is not symmetric in P and Q . Minimizing the relative entropy is popular and is related to optimizing a fit given the information available, according to principles of information theory. In the insurance pricing case, P is the real-world measure and a martingale Q is sought that will minimize the relative entropy – of course under the constraint that the transformed mean loss is the loaded expected loss in the premium. Q is then the martingale closest to P in the sense of relative entropy.

The $\phi(y)$ functions that give the minimum martingale measure (MMM) and minimum entropy measure (MEM) for the surplus process (loaded premium less compound Poisson claims) are known. For the MMM the transform is expressed with a positive constant $s < 1$:

$$\phi_M(y) = (sy/EY)/(1-s)$$

$$1+E\phi_M(Y) = 1/(1-s)$$

$$\lambda_M = \lambda/[1-s]$$

$$g_M(y) = g(y)[1-s + sy/EY]$$

The severity probability can be seen to increase for losses above the mean and decreases for losses below the mean.

For MEM, the frequency mean and severity probabilities are also multiplied by factors:

$$\phi_E(y) = e^{y/c} - 1, \text{ so}$$

$$1 + E\phi_E(Y) = Ee^{Y/c}$$

$$\lambda_E = \lambda Ee^{Y/c}$$

$$g_E(y) = g(y)e^{y/c}/Ee^{Y/c}.$$

It follows that $\lambda_E E_Y = \lambda E[Ye^{Y/c}]$, so the risk loading factor is $1 + \theta = E[Ye^{Y/c}]/EY$. This can be used to define c if θ is known. For the MMM,

$$\theta = \frac{s}{1-s} \frac{EY^2}{[EY]^2}$$

The MEM can be shown to be an Esscher transform for the compound Poisson. An Esscher transform in general is defined implicitly by $E^*[Z] = E[Ze^{Z/c}]/Ee^{Z/c}$. It is clear that the MEM severity transform is an Esscher transform of severity. There is a general theorem that the MEM transform is always an Esscher transform.

The combined frequency-severity transforms can be more complex to apply than the aggregate transforms, since this has to be done before calculating the aggregate distribution. Like the transforms of aggregate results, they use the whole distribution of risk, but put more emphasis on the tail.

B.3 Which Risk Measures to Use?

Some moment measures, like semi-standard deviation, capture features of the risk that are financially important. Tail measures have become popular because of their direct connection to solvency needs. However they ignore some portion of risk that must be of some concern to companies because they are not willing to take it for free. RTVaR and WTVaR at low probability levels get around much of this problem. Transformed probability measures use the entire distribution but emphasize the tail, which seems appropriate. They are also linked to pricing of risk, which is a good idea if profit is going to be judged in comparison to the risk measure.

Appendix C – Risk Pricing Issues

An academic point of view might be that target profitability is mainly about getting the risk pricing right. But insurance risks tend to be diverse, and rating plans do not capture all the diversity. Risk selection then becomes an important skill, and prices are modified based on subjective assessments of risk characteristics. Thus evaluating the profitability of a book of insurance business involves comparing the profit to the aggregated risk profile. This exercise can be thought of as a retrospective pricing and risk evaluation of the entire book of business. Risk pricing concepts become relevant at this level. Risk-adjusted profitability can be looked at as risk pricing of the book of business. Typically this will involve more detailed work than capital allocation via allocation of a risk measure.

Risk pricing is often viewed as having two chief components: return for taking on risk; and compensation for the frictional costs of holding capital. These are addressed separately.

C.1 Return for Risk Taking

The financial literature has two primary pricing paradigms: the capital asset pricing model, CAPM, and no-arbitrage pricing. No-arbitrage pricing is often used for contingent claims (options) since in complete markets unique prices can be derived. However in incomplete markets like insurance this is not the case. No-arbitrage pricing can determine relationships among prices for various instruments in incomplete markets, but in itself cannot even provide a finite range for the profit provisions.

CAPM gives prices that provide for adequate returns on risk for diversified investors. Options pricing in complete markets is able to ignore the diversification because unique prices are determined based only on individual risk characteristics. This does not hold in incomplete markets, so it is not clear that contingent claims in such markets can be priced without reference to the market portfolio.

The main tool for no-arbitrage pricing even in incomplete markets is to price as the mean of a transformed probability distribution. Venter [19] shows that CAPM can often be viewed as a transformed mean. The idea is to transform the probability distribution $f(y)$ of the risk Y to $g(y)$, which can be facilitated through an auxiliary function $h(y)$ so $g(y) = f(y)h(y)$. This implies that $Eh(Y) = 1$. To get to CAPM this way, take $h(y) = 1 + b(E[M|y] - EM)$, where M is the market portfolio and b is any constant small enough that h is not negative. It is easy

to see that $Eh(Y) = 1$. The transformed mean is $E[Yh(Y)] = EY + b\{E[YE(M|Y)] - [EM][EY]\} = EY + b\{E[YM] - [EM][EY]\} = EY + b\text{Cov}[Y,M]$. This is equivalent to the CAPM pricing formula as long as the b needed for CAPM is small enough so that $h > 0$.

The usual formulation of CAPM is in terms of returns: $ER_Y = R_F + \gamma\text{Cov}[R_Y, R_M]$, where R_F is the risk-free rate. In a single-period framework, $R_Y = Y_1/Y_0 - 1$ and similarly for M . This gives $EY_1 = (1+R_F)Y_0 + (\gamma/M_0)\text{Cov}(Y_1, M_1)$. If Y is assets, its expected growth is the risk-free rate plus the covariance loading. Since $b = \gamma/M_0$, b is probably small, so $h > 0$.

CAPM dates from the 1960's and is based on a normal distribution of returns. By 1973 Rubenstein [20] had shown that for non-normal returns, maximizing investor utility requires loading for higher co-moments as well. These are defined similarly to other co-measures. The n^{th} co-moment $s_n(Y, M) = E\{[Y - EY][M - EM]^{n-1}\}$. Note that this is not commutative, that is it is not necessarily the same as $s_n(M, Y)$ for $n \neq 2$. Investors tend to prefer positive odd moments and low even moments. But returns often have negative odd moments, so require negative loadings for these. Working with the negative of returns is often convenient as then all moments would require positive loadings.

Chung et al. [21] find that observed risk factors such as those identified by Fama and French become insignificant if enough higher co-moments are included in the pricing formula. Hung [22] shows that including the 3rd and 4th co-moments gives better fits than does CAPM with those additional risk factors.

Higher co-moments can also be represented as probability transforms. For instance, expanding the co-3rd moment shows that it is $E[YM^2] - EM^2EY - 2E[YM]EM + 2EY(EM)^2$. Then similar arguments as for covariance show that setting $h(y) = 1 + b\{E[M^2|y] - EM^2 + 2[EM]^2 - 2E[M|y]EM\}$ gives the loaded co-3rd moment as a transformed mean.

Thus no-arbitrage pricing has the potential to represent the risk needs of a diversified investor. However to do this requires certain types of probability transforms.

A related issue is the pricing of discontinuous jumps in the process. Authors tend to assume that jumps are priced, or that they are not. A typical argument that jumps are not priced is that their effect on the moments is already included in moments pricing. In a single-period model this might be true. However a lot of no-arbitrage pricing is done in continuous models. Utility theory is typically done in a single-period framework, but there might be an

extension of it to a continuous process. In that case, an investor would always be interested in current wealth. You would think that a risk of bigger changes would be worse than smaller changes, so jump risk would be disfavored. Also adding jump risk to a complete market makes it incomplete, and some jump risk cannot be completely hedged. Investors probably would not like that either. This reasoning suggests that jump risk should be priced. But it is probably the jump risk in the portfolio, and thus in the market, that would be the issue. This leads to the idea of pricing co-jumps with the market for individual securities.

Defining co-jump risk is most readily done specific to a particular discontinuous model of pricing movements. One possibility would be a compound-Poisson model, or other such compound frequency-severity model. Another possibility would be a Levy process. In general these processes allow infinitely many small jumps in a finite period, but only finitely many of these are greater than ε no matter how small ε is. Thus it should not be too troublesome to ignore the small jumps and consider a model that is a continuous process plus a compound-Poisson process. The jumps would be at the Poisson events, and the co-jump risk could simply be a co-measure of the security jumps with the market jumps. For something like earthquake insurance, the insurance losses could correspond greatly with the market losses, so a reasonable first approximation to the co-jump would be the risk's own jump. Thus just pricing jump risk as idiosyncratic risk may be reasonable as a starting point.

The minimum-entropy martingale transform discussed in Appendix B is a popular transform in incomplete markets. For a compound-Poisson process this applies to the jump risk component by simply increasing the expected number of jumps. A possible way to incorporate market risk for the diversified investor is provided by the technique of exponential tilting defined by Bühlmann [23]. The minimum-entropy martingale is an Esscher transform, which is closely related to the moment-generating function. If many co-moments with the market are needed to reflect investor risk attitudes, perhaps a co-moment generating function could work, at least for the higher moments. This might be especially applicable if the negative of profit were the risk variable, as then all moment preferences would have the same sign.

The general exponential tilting for a compound Poisson process could be defined as: $\lambda^* = \lambda E e^{M/c}$; $g^*(y) = g(y) E[e^{M/c} | y] / E e^{M/c}$. This is similar to the minimum entropy martingale but with the market risk replacing the loss risk at some points.

3.2 Frictional Costs of Holding Capital

Costs that accrue from holding capital whether or not risk is taken, such as taxes on investment income, are considered frictional costs of holding capital. Agency risk, that is the risk that management will use capital according to its own priorities rather than those of the shareholders, is thought to reduce the value of a company, and so would also be a frictional cost. A related risk is that of putting your investments under someone else's control. This risk may be a reason that closed-end mutual funds often trade below book value.

Allocating capital seems to make sense as a way to allocate the frictional costs of holding capital. Myers and Read [24] suggest an allocation method that seems to be designed for this purpose. They constrain capital so that the value of the default put option as a percentage of loss liabilities is constant. The reduction in this capital from a slight percentage decrease in the writings of a line is then the allocation of capital to that portion of the line. This is a marginal allocation that adds up. Myers and Read calculate this using standard option pricing assumptions, and include earnings from investment income as well as underwriting in their calculation. Venter et al. [3] apply this same principle more abstractly with separate random variables for underwriting and total profit, and no assumptions on distributional form.

The Myers-Read allocation makes sense for frictional costs as keeping the target ratio of default put to losses seems to be a reasonable interpretation of the function of capital. However there are other reasonable interpretations of the function of capital, such as keeping the probability of default low, or keeping the expected losses given default below a target level, or having enough to renew all existing business after a 100 year loss. Each of these would lead to a different marginal allocation of capital that adds up to total capital.

The default put option is favored by academics, but remember that the market value of this option is the added value to the shareholders of having the option to default. While this can be quite a valuable option in some cases (Lloyds not having had it comes to mind) usually it adds only a small amount to the shareholders' portfolio value. For the unfortunate policyholders who get left uncovered, who are usually not diversified in their insurance purchases, this option has a much greater value. Phillips et al. [25] find in fact that policyholders demand a much greater discount in premium than the market value of this option. While this does not totally negate the significance of the value of the default put, it does reduce any advantage it may have had over the other interpretations of the value of capital above.

This again leaves capital allocation as an arbitrary choice, even in the realm of allocation of frictional costs of holding capital. Gründl and Schmeiser [26] provide a take-no-prisoners critique of capital allocation, and argue that choosing among such arbitrary alternatives is bound to lead to erroneous conclusions:

“According to the common cost literature, informational limitations leave us with no nonarbitrary common cost allocation for purposes of performance measurement and pricing. Instead, the generally accepted response is to develop a set of desired properties for the allocation process itself and proceed with the method that best satisfies these properties. It is inherent in such a process, however, that whatever allocation method used will result in distortions and the question future research ought to investigate is the extent to which those distortions exist under various allocation methods.”

They do provide an alternative, however. They suggest that each frictional cost be treated individually and included in a contingent claims approach to pricing. For instance for corporate taxation they suggest, following Doherty and Garven [27], that the government holds an option on the profits of the firm. This in actuality is a fairly complex option, depending on the current and historical profitability of the firm. But given the right martingale transform, the value of this option should be calculable without allocating capital, and then it can be added to the target profitability of each business unit. Essentially this approach converts frictional costs to another risk charge, but for a risk different than insurance losses.

3.3 Conclusions

Calculating risk-adjusted return is risk pricing applied to the entire business unit. Risk pricing theories for insurance need to take into account the heavy-tailed distributions of returns and the presence of jumps. Financial risk-pricing theory is advancing in this direction and has the potential to capture these effects. Both the cost of bearing risk and frictional costs could be covered by this approach. Selecting a risk measure and allocating it is not likely to match the results of such a calculation, however.

Biography of the Author

Gary Venter is managing director at Guy Carpenter, LLC. He has an undergraduate degree in philosophy and mathematics from the University of California and an MS in mathematics from Stanford University. He has previously worked at Fireman’s Fund, Prudential Reinsurance, NCCI, Workers Compensation Reinsurance Bureau and Sedgwick Re, some of which still exist in one form or another. At Guy Carpenter, Gary develops risk management and risk modeling methodology for application to insurance companies. He also teaches a graduate course in loss modeling at Columbia University.

917.937.3277

gary.g.venter@guycarp.com