Capital Need, Allocation and Risk Pricing Example

Gary G. Venter, FCAS, CERA

Abstract

Wacek (2014) discusses a simplified example of capital need, allocation and pricing. The current paper uses that example to illustrate some of the more recent applications of actuarial finance to such calculations. Most of this has been puslished somewhere sometime, so there is nothing absolutely new here except possibly in combining the ingredients, perhaps also illustrating the adage "the old ones are the great ones."

Keywords. Capital need, risk pricing, allocation, distortion, transform.

1. CAPITAL NEED

A classical heist movie, in the team recruitment stage, featured the following dialogue:

I don't know, I don't feel right about ripping those people off.

Oh, don't worry about them – they have insurance for all that.

But then, aren't we just ripping off the insurance company?

No, they need to pay a few like this, to keep 'em buying.

Perhaps the project planner was on to something here. Before being shut down by regulators, for example, Alien Abduction Insurance did actually pay a few claims. But to keep 'em buying, a company also needs to convince the customers that it is good for the potential losses. That's what the capital need derives from. The fundamental role of insurance companies is to reduce the anxiety of their customers, which means that capital has to be large enough to inspire confidence.

Prospect theory has found that the amount people are willing to pay to avoid extreme risks is greater than utility theory might suggest, perhaps due to a tendency to exaggerate small probabilities. If a family has a 0.1% chance of their home burning down and buys insurance that has a 0.1% chance of failing to pay, there will be a tendency to assess these rare events as being likely to occur simultaneously.

Phillips, Cummins, and Allen (1998) indeed found that homeowners were willing to pay substantially more than the expected loss differential for insurance with a reduced probability of default, except when state guarantee funds were in place. Thus it could be worthwhile for an insurer to maintain a substantial level of capital to attract certain types of customers and to charge them higher prices. Could be but not necessarily – not all insureds have the same risk preferences. There are others willing to pay less for less security. How much capital an insurer needs is not a simple optimization. It depends on what market they have access to and the risk preferences of that market.

Wacek uses as an example a capital requirement at the 1:250 VaR level. But something like that is not a capital requirement – it is a capital choice. It is implicitly targeting a population that will pay enough for the insurer to justify holding that level of capital. Well-capitalized insurers tend to play up such capital choices as part of their marketing efforts. One for instance announced a target of maintaining 2.5 times 1:100 VaR. A less model-intensive target could be three times regulatory minimum capital. Companies compete on such grounds. However these examples are of companies looking for high-margin business from highly risk-adverse clients. There is money to be made in other market segments that work on significantly less capital and margin, perhaps substandard auto.

Bottom line: the capital requirement is whatever is needed to keep 'em buying.

2. THE EXAMPLE

Wacek builds on an example that originated in Mango et al. (2013). A sidecar reinsurer has been set up for a year to write catastrophe reinsurance. It is subject to five possible loss events, each with binomial probability of 1% (but apparently which are mutually exclusive), with losses of 100, 200, 300, 400, and 500. The reinsurer is organized into five of what I will call business units, each of which will pay either zero or 100 of the loss. These amounts are arranged in layers, so the first unit will pay 100 in the event of any loss, the second unit will pay 100 if the loss exceeds 100, etc. They then each have loss severity of 100 but loss probabilities of 5%, 4%, 3%, 2%, and 1% and so expected loss of 5, 4, 3, 2, and 1. The company has expected loss of 15.

Wacek calculates the capital and profit goals for the company based on assumptions of 1:250 VaR and 15% post-tax return in the US. The biggest possible loss has 1% probability, so the VaR requirement means that there must be sufficient funds to pay all of the losses. This is arguably what is also needed to keep 'em buying. If an opportunistic hedge fund wants to set up a sidecar for one year to take a quite limited range of losses, it better be able to pay all of them. The 15% is presumably the funders' required expected return needed to get the money flowing. It is not unusual for a target like this to be exogenous to the actuarial calculations.

He specifies additional conditions that the premium is paid at the beginning of the year, the losses are paid at the end, and funds earn 3% risk-free. Then the constraints are all satisfied with a premium of 91.39 and capital of 394.07, totaling 485.46, which by the end of the year grows to 500.

Now the fun part: allocating the premium and capital to business unit.

3. FINANCIAL PRICING THEORY

Actuaries were exploring financial pricing prior to Venter (1991), but probably without such ire from traditional mean-variance pricers, judging by the comments of the reviewers. This discussion did inspire quite good research by other actuaries, however, including especially by Shaun Wang in a series of increasingly sophisticated papers, leading to the Wang transform. See the references.

The kernel of my argument was that even though insurance markets are not complete, insurance pricing still has to be arbitrage-free, or else competition will erode the arbitrage opportunities, even if the instruments to end up in an actual arbitraged position do not exist. Finance theory tells us that arbitrage-free pricing requires that risk prices have to be expected values from transformed probability distributions, sometimes referred to as Q-measures. Usually these transformed means will be higher than the actual means, so a risk charge is included. Also the transforms have to be made on the probabilities of events, not on the probability distributions of returns of deals, so all the deals get the same transform. The transform is uniquely determined in complete markets, but it is not in incomplete markets, so insurance pricing by transformed distributions still includes some need for understanding market impacts.

3.1 So Which Transform?

Reesor and McLeish (2003) systematically lay out a number of probability transforms. (An earlier version of their paper is on the web as a working paper.) Two we will look at are the exponential transform and what they call the normal transform, which is the original Wang transform. They are not too complicated. Here we express the transformed cumulative probability q in terms of the actual probability p. Each transform has a parameter, m for the Wang transform and b for the exponential. The Wang transform uses the standard normal cdf Φ:

$$q = \Phi(\Phi^{-1}(p) - m)$$

The percentile at p is shifted downward by m, which produces a lower probability. That means that more of the probability goes above any particular percentile.

The exponential transform starts with the exponential distribution with mean $1/b$ limited to the unit interval, with cdf F. The transform is then

$$q = 1 - F(1 - p), \text{ or}$$

$$q = [\exp(pb) - 1] \,/\, [\exp(b) - 1]$$

This also moves the probabilities away from the smaller losses. It turns out that the Wang transform moves more of the probability to the extreme upper tail, as we will see below. But for some applications even more of a shift to the extreme tail is needed. For this John Major suggested using a mixed t – normal transform. Wang (2004) applies this idea. In practice we tend to use the form below, which I call the normal – t. Here T_v is the t cdf with v dof.

$$q = \Phi(\Phi^{-1}(p) - m) \text{ if } p < \Phi(m)$$

$$q = T_v(\Phi^{-1}(p) - m) \text{ otherwise}$$

The t and normal distributions both have value 0.5 at zero, so this is continuous. Note that v does not have to be an integer. The transformed beta distribution interpolates the t-distribution like the gamma function interpolates factorial. Using Excel as notation, the cdf is:

$$T_v(w) = \tfrac{1}{2} + \tfrac{1}{2}\,\text{sign}(w)\text{betadist}\mathbf{[}w^2/(v+w^2), \tfrac{1}{2}, v/2\mathbf{]}$$

This also makes it convenient to estimate v by methods that work only for continuous functions, like Newton's method.

Niederau and Zweifel (2009) use the exponential transform in some actuarial pricing applications. They report that this transform is the minimum entropy transform That basically means it is the transform that is closest to the real-world distribution in information distance. People who know what that means seem to believe it is a good thing. Actually people used to think that the Esscher transform was the minimum entropy transform but it turns out upon closer examination that transform is a function of the losses while the exponential transform is the Esscher transform applied to the probabilities, and that is what is needed to minimize entropy = disorder. We will look at all three of these transforms, but exponential does seem to be the most reasonable in the end.

3.2 Trying It Out

The total risk premium required in this example has been specified by investor requirements at 91.39. This is over six times the expected loss of 15, but with payments up to 500 being quite possible, the expected loss does not mean much in this context. High cost/low probability risks often get such margins. For instance Chen et al. (2014) find that AAA/AA rated bonds have credit spreads about seven times their default probability, while for BB bonds it is about twice. The credit spread also compensates for liquidity risk, however. Wang (2004) was actually a parallel modeling exercise for bond and cat bond credit spreads, which ended up implying similar transforms.

The problem now is to come up with prices for the five business units (layers) that incorporate their relative riskiness and add to 91.39. The way to do this with transforms is to find parameters for the transforms such that the transformed expected loss for the whole business is 91.39, then use those probabilities to price each business at its transformed mean.

Table 1 – Transformed Probabilities

Loss	Probability		Exponential	Wang	Normal-t
0	0.95	0.95	0.709	0.740	0.769
100	0.96	0.01	0.051	0.033	0.021
200	0.97	0.01	0.054	0.037	0.023
300	0.98	0.01	0.058	0.043	0.026
400	0.99	0.01	0.062	0.054	0.033
500	1	0.01	0.067	0.092	0.127

Table 1 shows the resulting transformed probability distributions for each transform, with a rather extreme value of $v = 2$ for the normal-t. Wang (2004) gets v around 5. The way these probabilities were computed for each transform was:

1. Guess a value for the transformation parameter m or b.
2. Compute the transformed cumulative probabilities by the formulas for each transform.
3. Difference to get the transformed incremental probabilities.
4. Multiply each incremental probability by the corresponding loss and total these.
5. Subtract 91.39 from the total.
6. Increment on the parameter, e.g., by using Goal Seek, to make the difference zero.

This process would be basically the same starting from a simulated distribution of outcomes.

The resulting parameters were, in the order above, 6.8781 1.0003 0.7419. Each business unit has a loss of zero or 100 so has a price = transformed expected loss of 100 times its probability of attaching. Table 2 shows the probabilities for each layer for each transform and the ratio of price to mean loss. The layer probabilities are upward sums of the loss-size probabilities.

Table 2 – Prices by Layer for Each Transform

Layer	Probability	Exponential	Wang	Normal-t	Ratio of Price to Expected Loss		
1	0.05	0.291	0.260	0.231	5.8	5.2	4.6
2	0.04	0.241	0.227	0.210	6.0	5.7	5.2
3	0.03	0.187	0.189	0.186	6.2	6.3	6.2
4	0.02	0.129	0.146	0.160	6.4	7.3	8.0
5	0.01	0.067	0.092	0.127	6.7	9.2	12.7

The relative probabilities and prices across the transforms reverse as you go to higher layers. The Wang and normal-t transforms put more probability into the extreme right tail. A very large international insurer allocates profit targets to its business units using this method. They use the normal-t transform, but some practical actuary has capped the extreme probabilities. "I don't believe these and I can't sell them to the business units." His intuition seems in line with using one of the other transforms. The theoretical and practical leanings are thus towards the exponential transform.

Seeing these transforms as implied for the entire distribution function provides further insight.

Figure 1: Transformed probabilities at each actual probability.

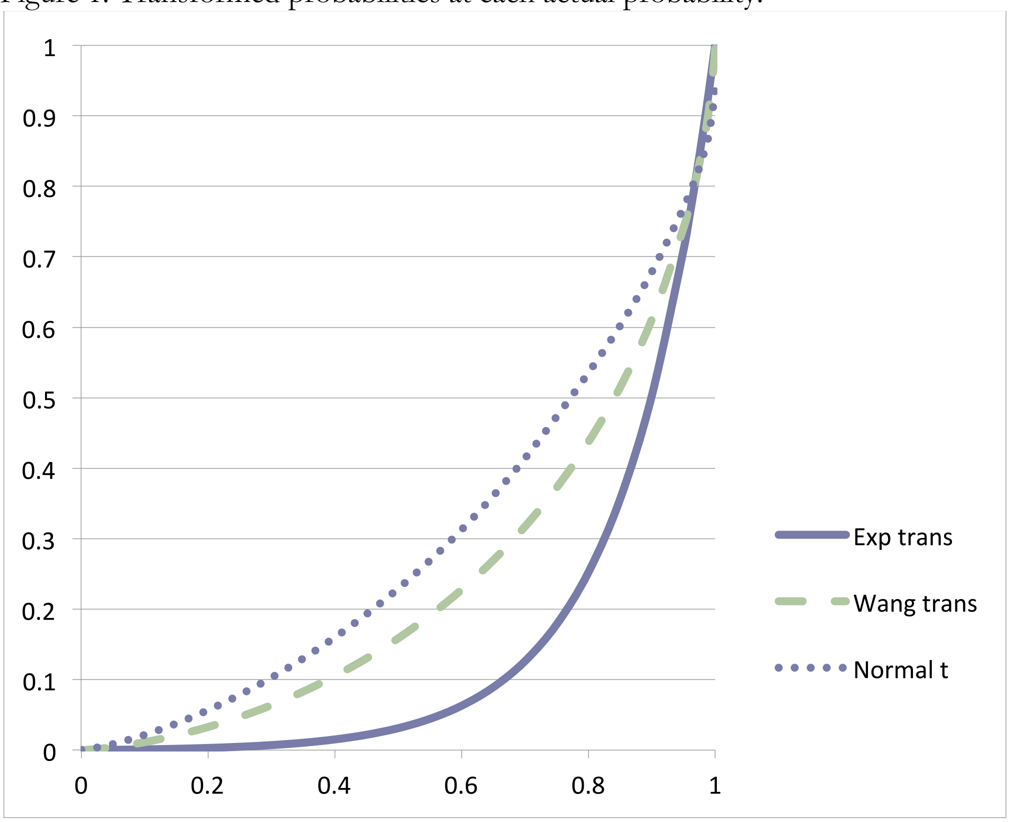

The actual cumulative probabilities would be a straight line corner to corner. Transformed probabilities are always lower than the actual since they show the probability below each point. The transforms cross near the end, with the normal-t having a fair amount of probability remaining at high percentiles. Blowing up the probabilities and graphing as a function of 1 – p on a log scale shows this in more detail.

Figure 2: Transformed probabilities at the extreme tail

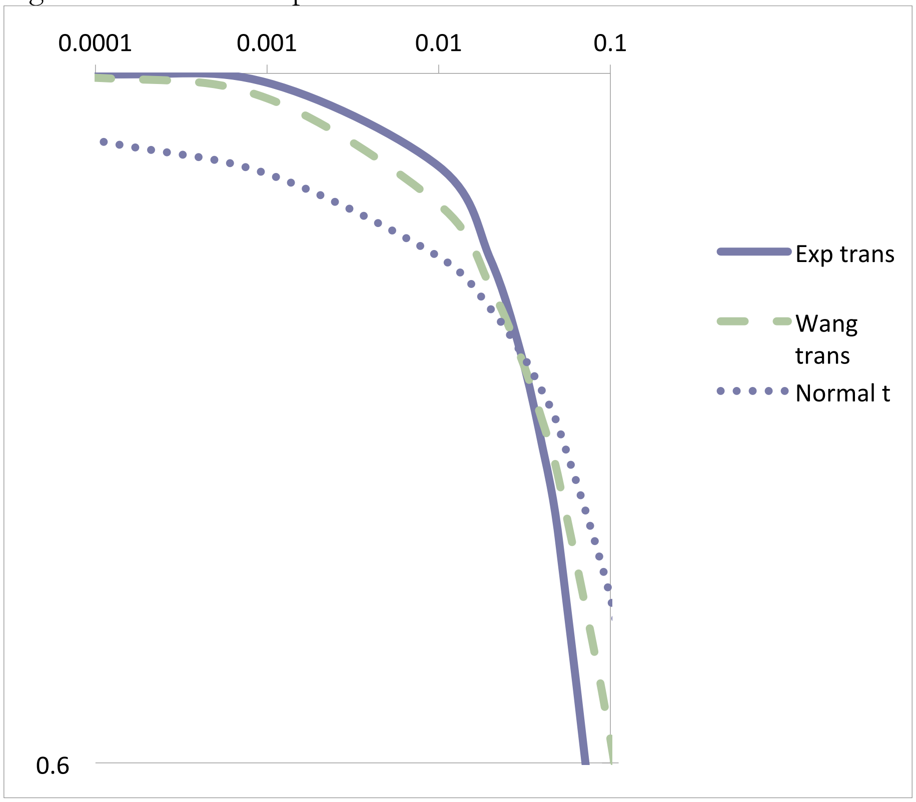

The normal-t packs in a lot of probability in the extreme tail.

4. CAPITAL ALLOCATION

Once pricing has been done the need for capital allocation becomes less pressing. Still it provides a handy way to compare returns. If comparison is the purpose, the most direct method of allocation would be in proportion to the target profit. Then business units' returns can be compared directly. If some other allocation is done, the comparison shifts to difference from target. Somehow it sounds better to say "return of 1%" rather than "8 points below target." At least it is a positive number and also obviates the skepticism that is sometimes associated with the targets. Sherris (2006), a prize winning paper, discusses such allocation in a probability transform pricing context. Wacek uses standard deviation pricing so also does not need capital allocation for pricing, but ends up allocating capital by an allocation of VaR. That seems mostly harmless but a little awkward as target returns vary by unit. We discuss allocation of VaR below.

4.1 Capital as a Shared Resource

Capital allocation may give the impression that each unit ends up with so much capital and that's all it has, but actually any business unit has access to the full capital of the firm. Merton and Perold (1993) start from this premise and postulate that the capital cost of a business unit to the firm is the value of the option the unit has to put any and all losses to the firm. Thus to add value, the unit must plan to earn more than that cost.

Ken Froot, in a private conversation, noted that this is not the full story. The firm also has an option to take all the profits of the business unit. His point was that in combination, these two options mean that the firm takes all the profits and pays all the losses, so there is no inherent optionality. Nonetheless, businesses might find it useful to compare the value of the call option they hold on the profits to the value of the put option the unit has on the liabilities.

Let's price these options using a probability transform with expected value operator E*[X]. Consider a unit with premium P and losses L, where q is the transformed probability that L > P. Then the values of the put and call to the firm are:

- Put = E*[P – L | L > P]q
- Call = E*[P – L | P > L](1–q)

Here L is a random variable and the value of P is sought. The value of the unit to the firm is the sum of these option values, which is E*[P – L]. The lowest acceptable P would be when that is zero, or P = E*(L). A probability transform is needed to apply to the whole firm, so the transformed mean of L has to be the actual mean plus the firm wide target profit. This is the same result as provided by financial pricing, but now split into profit and loss components.

4.2 Capital Allocation Other Than for Pricing

A non-pricing use for capital allocation arises when capital requirements are dictated by a regulatory rule, such as VaR at 1:200 for solvency two. A company might want to know how its business units contribute to that requirement even if they are not going to price that way. One way to approach that is to see how much a small change in the volume of the business unit would proportionally change the capital requirement for the firm.

In mathematical terms, this is the derivative of the firm VaR with respect to the volume of the business unit. It could be quantified by looking at the proportional capital effect of ceding a small percentage of the unit through a purely proportional reinsurance treaty. A theorem of Euler implies that the sum of these derivatives across the business units will be the firm VaR[1]. This fact makes it reasonable to consider that derivative as the contribution of the unit to the VaR of the firm.

Patrik et al. (1999) introduced the Euler method. Venter et al. (2006) call this method risk measure decomposition, and the derivatives co-measures, which they calculate for a few risk measures, including standard deviation and VaR. They show that co-VaR is the conditional expected loss of the business unit given that the firm loss is at the VaR level. These clearly add up across business units to the firm VaR and represent a plausible decomposition of VaR. The difficulty comes when the firm VaR has been estimated by simulation of all the business unit losses. VaR then is a single scenario, and how the losses in that one scenario split by business unit is pretty arbitrary. In practice an interval around that scenario is used to estimate both VaR and co-VaR. A straight or weighted average based on distance can be used. One cute example is to use a normal pdf centered at the VaR scenario to determine the weights. This has nothing in particular to recommend it, however.

This method of determining VaR and co-VaR is sometimes called fuzzy VaR allocation. In this context it can be used to estimate the contribution of each business unit to the capital requirement but has no justification as a risk-pricing methodology. Nonetheless, knowing what businesses are driving the regulatory capital requirement could be useful in capital management.

5. CONCLUSIONS

- Capital need is driven by the market segment an insurer is targeting, so for the same probability distribution of financial results could be different for different insurers.
- Firmwide target return typically comes from the requirements of the capital providers.
- Transformed distribution pricing is financially supported and not difficult to implement.
- The latest and greatest transform is the exponential, which has theoretical backing and appears practically reasonable.
- Capital allocation is not relevant for risk pricing, but may be useful for comparison across units and understanding regulatory requirements.

---

[1] Technically this is because VaR is a so-called homogenous function of premium volume. That basically means that it is denominated in dollar terms, as opposed to variance, which is quantified in dollars squared.

References

[1] Chen, Hui, Rui Cui, Zhiguo He, and Konstantin Milbradt 2014. "Quantifying Liquidity and Default Risks of Corporate Bonds over the Business Cycle," http://www.mit.edu/~huichen/bondliquidity.pdf

[2] Mango, Don, John Major, Avi Adler, and Claude Bunick. 2013. "Capital Tranching: A RAROC Approach to Assessing Reinsurance Cost Effectiveness," *Variance* 7(1):82-91.

[3] Merton, Robert and Andre Perold, 1993. 'Theory of Risk Capital In Financial Firms," *Journal of Applied Corporate Finance* 6: 16-32.

[4] Niederau, Harry and Zweifel, Peter. 2009. "Quasi Risk-Neutral Pricing in Insurance," *ASTIN Bulletin* 39(1): 317-337.

[5] Patrik, Gary, Stefan Bernegger and Marcel Beat Rüegg 1999. "The Use of Risk Adjusted Capital to Support Business Decision-Making," *CAS Forum* Spring

[6] Phillips, Richard D., J. David Cummins, and Franklin Allen, 1998. "Financial Pricing of Insurance in the Multiple Line Insurance Company," *Journal of Risk and Insurance* 65: 597-636.

[7] Reesor, R. Mark and Don L. McLeish. 2003. "Risk, Entropy, and the Transformation of Distributions," *NAAJ* 2003 7(2): 128-144

[8] Sherris, Michael, 2006, "Solvency, Capital Allocation and Fair Rate of Return in Insurance," *Journal of Risk and Insurance*, 73(1) :71- 96.

[9] Venter, Gary G. 1991. "Premium Calculation Implications of Reinsurance without Arbitrage," *ASTIN Bulletin* 21:223–230.

[10] Venter, Gary G., Rodney E. Kreps and John A Major. 2006. "Marginal Decomposition of Risk Measures," *ASTIN Bulletin* 36(2): 375–413

[11] Wacek, Michael G. 2014. "Discussion of 'Capital Tranching: A RAROC Approach to Assessing Reinsurance Cost Effectiveness Variance 7.1.'" *Variance* 7(2).

[12] Wang, Shaun S. 1995. "Insurance Pricing and Increased Limits Ratemaking by Proportional Hazards Transforms," *Insurance: Mathematics and Economics* 17: 43-54.

[13] Wang, Shaun S. 1996a. "Ordering of Risks Under PH-Transforms," *Insurance: Mathematics and Economics* 18: 109-114.

[14] Wang, Shaun S. 1996b. "Premium Calculation by Transforming the Layer Premium Density," *ASTIN Bulletin* 26(1): 71-92.

[15] Wang, Shaun S. 2000. "A Class of Distortion Operators for Pricing Financial and Insurance Risks," *Journal of Risk and Insurance* 67(1): 15-36.

[16] Wang, Shaun S. 2004. "Cat Bond Pricing Using Probability Transforms." Geneva Papers: Etudes et Dossiers, special issue on Insurance and the State of the Art in Cat Bond Pricing, 278(January): 19–29.

Biography of the Author

Gary Venter is Head of PC Model Validation at AIG. 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, Guy Carpenter and Sedgwick Re, some of which still exist in one form or another. He also teaches a graduate course in quantitative risk management at Columbia University.

212.551.5471 **gary.venter@aig.com**