Society of Actuaries

Article from:

Risk Management

August 2014 – Issue 30

Down But Not Out: A Cost of Capital Approach to Fair Value Risk Margins

By B. John Manistre

Editor's Note: *This is a conceptual introduction to a much longer technical paper with the same title to be presented at the 2014 ERM Symposium.*

INTRODUCTION

There is a well-known quote, due to George E.P. Box, which goes, "All models are wrong but some are useful."[1] All of the methods outlined in this article take this concept to heart in the sense that the model structures themselves recognize that the models are wrong and will require adjustment as new information becomes available. The models are therefore intended to be applied in the context of a principles based, fair valuation system where continuous model improvement is an integral part of the process. One possible application would be to an internal economic capital model or an Own Risk and Self-Assessment (ORSA) process.

The cost of capital concept itself has been part of actuarial culture for many decades and this paper assumes the reader already has some familiarity with the idea. At a high level, the idea is that if a contract requires the enterprise to hold economic capital in the amount EC then we need to build an annual expense πEC into the value of the contract to price in the risk. The quantity π here is the cost of capital rate and it can vary from application to application. For non-hedgeable life insurance risk a typical cost of capital rate is $\pi = .06$.

THREE THEMES

There are three themes or common denominators that run through all of the methods presented in the complete research paper. These are (1) Down but not out, (2) Linearity and (3) The basic risk modeling process.

1 - Down but not Out

The idea is that if a 1 in N year event wipes out the economic capital of a risk enterprise there should still be enough risk margin on the balance sheet that the company can either attract a new investor to replace the lost capital or, equivalently, pay a similar healthy enterprise to take on its obligations. The chart at the top of column 2 illustrates the idea graphically.

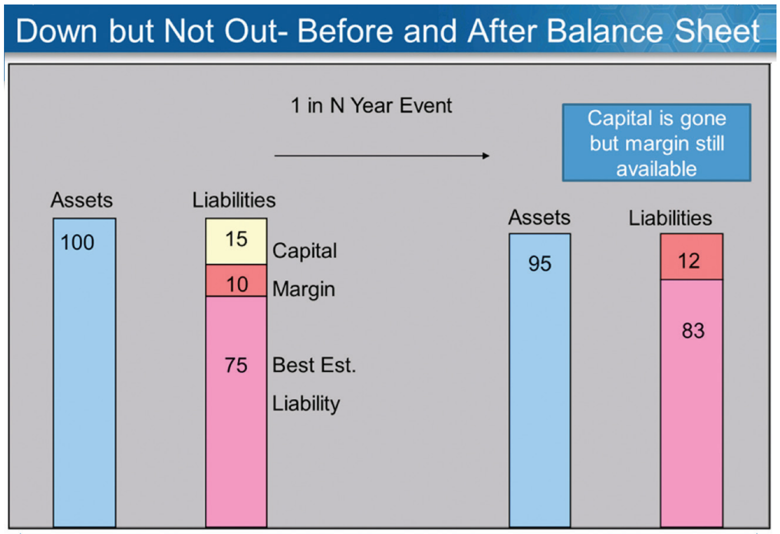

On the left side of the chart we see the risk enterprise's economic balance sheet at the beginning of the year. The right side of the chart shows the fair value balance sheet after a bad year. As a result of both poor experience in the current year and adverse assumption revisions all of the economic capital is gone. The risk enterprise is down. However, the economic balance sheet is still strong enough that it can either attract a new investor to replace the lost capital or pay another enterprise to take on its obligations i.e., the risk enterprise is not out because appropriate risk margins are still available.

This is clearly a desirable theoretical property for a model to have. In order to actually work in practice the revised balance sheet on the right must have enough credibility with the outside world that a knowledgeable investor would actually put up the funds necessary to continue. One way to get the needed credibility is for the actuarial profession to devel-

B. John Manistre, FSA, CERA, FCIA, MAAA, is a research actuary at GGY AXIS in Baltimore, Md. He can be reached at *john.manistre@ggyaxis.com*.

CONTINUED ON PAGE 28

op standards of practice that are rigorous enough for the shocked balance sheet to be credible.

2 - Linearity

All of the methods considered here can be formulated as systems of linear stochastic equations. This has two very general consequences.

1. As is well known, a linear problem usually has a dual version. If you can solve the primal problem you can also solve the dual to get the same answer. In this case the primal version of the problem looks like an "actuarial" calculation where we project capital requirements into the future and then compute margins as the present value of the cost of capital.

   As formulated here, the dual version of the problem looks more like a "financial engineering" calculation. The process above is reversed by starting with a concept of risk neutral or risk loaded mortality, lapse etc. and then determining the corresponding implied economic capital by seeing how the margins unwind over time.

   Put another way, if the present value of margins M and the economic capital EC are related by an equation of the form

$$\frac{dM}{dt} = (r + \mu)M - \pi EC,$$

then the primal version of the method starts by projecting EC and then uses the above relation to calculate margins by discounting. The dual approach calculates M first and then uses a version of the relation above to estimate an implied economic capital EC.

2. A second useful consequence of using linear models is that they allow us to avoid the "stochastic on stochastic" issue that bedevils many other approaches to the margin issue. Linear models can be calculated scenario by economic scenario. Any errors we make by ignoring the "stochastic on stochastic" nature of the problem average out to zero when we sum over a large set of risk neutral scenarios. [2]

With this result we can develop the cost of capital ideas in a simple deterministic economic model, and be confident that the results developed will continue to apply when we go to a fully stochastic economic model.

Looking at the dual approach gives us both new theoretical insight and an alternative way to compute any given model. In particular, the dual approach adds transparency in the sense that it tells us what the implied "risk neutral" assumptions for mortality, lapse etc. are.

For any particular application, the primal and dual approaches are equivalent but can differ in practice for a variety of reasons. One of the paper's general conclusions is that solving the primal problem works well for simple applications but the dual approach can be preferable as the complexity of the application increases. The main problem with the dual approach is the effort required to understand why the theory works. The actual implementation is not that difficult.

We take the view that both the primal and dual versions of a model should make theoretical sense and this leads to a critique of some approaches. For example, the primal version of the prospective model used in Europe usually looks simple and reasonable but the dual version may not. This is illustrated in the main paper by looking at the example of a lapse supported insurance product. It is possible for the dual problem to exhibit negative risk loaded lapse rates. We offer a modification to the method, as well as several other approaches, that can resolve this issue.

3 - The basic risk modeling process

This article assumes a three step process for putting a value on non-hedgeable risk. In a bit more detail, the steps are:

1. Develop a best estimate model that is appropriate to the circumstances of the application. Detailed discussion of this step is outside the scope of even the main paper although we do provide a number of examples from life insurance. The key assumption we make is that our best estimate models are not perfect and are subject to revision.

"The risk enterprise is not out because appropriate risk margins are still available."

2. Hold capital and risk margins for a contagion event i.e., the risk that current experience may differ substantially from our best estimate.

Imagine, for the sake of clarity, that our best estimate model is a traditional actuarial mortality table. Even if our table is right on average, we could still have bad experience in any given year. The classic example of a contagion event would be a repeat of the 1918 flu epidemic—hence the name contagion risk.

More recent examples of contagion risk events would be the North American commercial mortgage meltdown in the early 1990s[3] and the well-known problems with the U.S. residential mortgage market that led to the financial crisis of 2008.

A risk enterprise should have sufficient capital and margins that it can withstand a plausible contagion event and still be able to continue as a going concern without regulatory intervention. We show that traditional, static, risk loadings in our parameters can usually deal with this issue.

3. Hold capital and margins for parameter risk: new information might arrive in the course of a year that causes the risk enterprise to revise one or more models. To the extent these model revisions cause the fair value of liabilities to increase, we need economic capital to absorb the loss. Again we need a margin model that allows the risk enterprise to withstand the loss and carry on without regulatory intervention. To deal with this issue, we introduce the concept of a dynamic margin which arises naturally out of the dual approach.

Static and dynamic loadings differ in the way margin gets released into income over time. If best estimate assumptions are realized, then any static margin emerges as an experience gain in the current reporting period. The risk loading is engineered so that the resulting gain is equal to the cost of holding capital for contagion risk. This is what most actuaries would expect.

By contrast, a dynamic margin is a time dependent loading to a parameter which is equal to zero at the valuation date and then grades to an ultimate value we discuss later. There is very little experience gain in the current reporting period. The risk margin gets released into income by pushing out the grading process as time evolves i.e., when we come to do a new valuation, we establish a new dynamic margin which restarts from zero at the new valuation date. If we get the math right, this process releases the correct amount of margin to pay for the cost of holding economic capital for parameter risk, while still leaving sufficient margin on the balance sheet for the future.

Chart 1 below shows a simple example of the risk loading ideas introduced above.

In this example we have a model parameter whose best estimate value is $\theta_0 = 100\%$ and a static contagion loading of 5% has been added. At the valuation date

Risk Loading Example

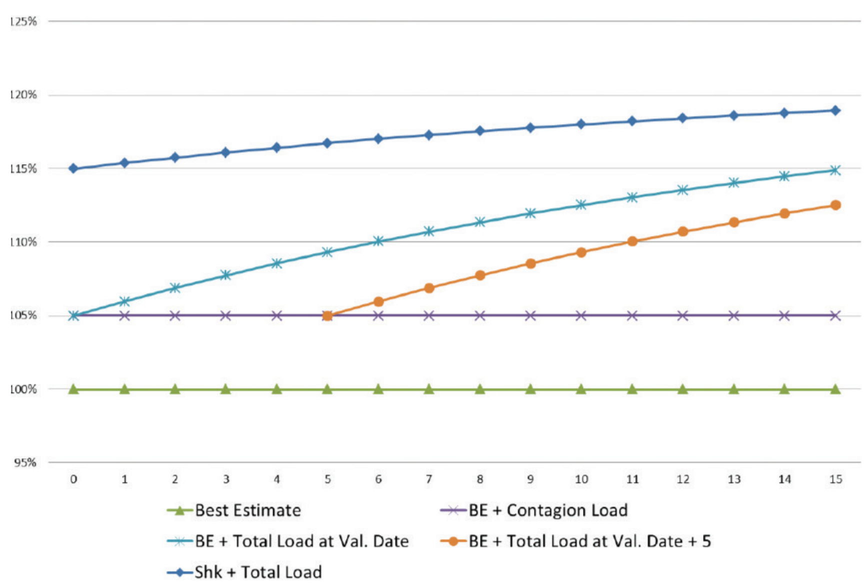

$(t = 0)$, we have added a dynamic load that takes the parameter up to the value of 115% over a 15 year period. This is the parameter path used to compute a fair value. A shocked fair value is calculated assuming a shocked path that starts at 115% (base + 10%) and then grades to about 119%. Economic capital, for parameter

CONTINUED ON PAGE 30

risk, is the difference between the shocked and base fair values.

When we come to do a new valuation five years later, the contagion loading has not changed but the dynamic loading for parameter risk has been recalculated to start at zero again. The risk margin released into income, if the assumptions do not change, is engineered to provide a target return on the risk capital.

SUMMARY

A high level summary of the paper's theory is that the cost of capital method for calculating risk margins is, for most practical purposes, equivalent to using an appropriate combination of static and dynamic risk loadings.

The process described above is much easier to implement than it looks. The full paper discusses a number of reasonable simplifying assumptions that allow the risk loaded parameters to be calculated fairly easily. None of the methods discussed require any computationally expensive "stochastic on stochastic" or "projection within projection" algorithms.

Two additional versions of this paper are available for further reading: 1) a condensed version which summarizes the theory of this approach and provides several practical examples, and 2) a full detailed theoretical development. The condensed version is available on the Risk Management website at: *http://www.soa.org/Professional-Interests/Joint-Risk-Management/Joint-Risk-Management-Section.aspx.* The full version of the paper is forthcoming, and will be introduced at the SOA ERM Symposium in September 2014. ■

ENDNOTES

1 George E.P. Box (FRS) in 1987.

2 This is a standard result in stochastic calculus which is outlined in the main technical paper.

3 This was caused by the overbuilding of office space during the 1980's in many North American cities. When the oversupply became apparent, office rents plummeted. This dragged down property values and triggered defaults on many of the mortgages used to finance the office towers.