



## **Response to CBI paper on Risk Appetite**

This response is submitted in an individual capacity. I currently act as an INED for three life assurance companies and the views expressed herein are entirely independent.

**Brian Woods BSc FIA FSAI**

**July 2014**

## Introduction

I welcome this opportunity to participate in the discussion on risk appetite prompted by the CBI's paper. Although my contribution is in an individual capacity it will be from my perspective as an actuary in the life insurance industry and in particular from my acquaintance with VA business. Also, my experience has invariably been with subsidiaries of much larger companies. I restrict my comments to the capital aspects of risk appetite.

The CBI paper is addressed to the broader financial services community and therefore it is not surprising that some of the concepts are less relevant from my narrower perspective. For example, risk diversification and individualised risk limits will be highly relevant for, say, compiling a loan book, building a commercial lines insurance portfolio or running a derivatives trading platform. By contrast, at the retail level of life assurance the focus is on marketing, distribution and customer service with the risk aspects being the preserve of the technical experts and, of course, the Board, and that in my opinion is how it should be.

This variety of perspectives underpins the view, expressed in the paper, that a generic type approach by the Supervisor to risk appetite would not be helpful.

## Word games

The paper notes that the noun "risk" is used to qualify four other nouns, sometimes interchangeably, viz. "appetite", "capacity", "tolerance" and "limits". The dictionary definition of "appetite" places a strong emphasis on the concept of "desire". It is therefore somewhat unfortunate that risk **appetite** has become the overarching theme. Gamblers might "desire" risk. Rational players in pursuit of profitable opportunities accept that they must "tolerate" a level of risk. To my mind it would have been preferable if risk **tolerance** was the dominant motif. However, we are where we are and even I will find myself using the word "appetite" when I really mean "tolerance". That is because "appetite" does also convey a somewhat different meaning additional to "desire". It also refers to "capacity" and that is a much more relevant concept. The FSB definition describes the situation succinctly.

<p><b>Risk Appetite:</b> The aggregate level and types of risk an organisation is willing to assume within its risk capacity to achieve its strategic objectives and business plan.</p>
---

My only slight quibble is that I would have preferred the word "prepared" to "willing" as the latter could be misconstrued as indicating an element of desire. For a VA business the central driver is capacity. Subsidiaries of much larger companies usually find that in practical terms no capacity constraint will be applied to writing profitable new business – the parent will be only too willing to supply any capital for this purpose. Where the rationing of capacity occurs for a VA writer is in deciding what market risks to hedge and what insurance risks to reinsure. We thus see that the above FSB definition accurately describes that situation.

One might usefully distinguish between those risks which have been transferred from the policyholder and those risks which are incremental. The latter include operational and counterparty risk. Longevity and mortality risks are clear examples of the former. Market risk is less obviously of the transfer variety as the policyholder's market risk will have been created by taking out the policy in the first place! However, upon reflection, what is happening is that the policyholder is persuaded that an element of market risk is necessary to enhance prospective returns and the VA writer then puts a floor on the policyholder's risk exposure. Basis risk stands out as a "bad" risk in this analysis. The policyholder garners no value from basis risk whilst a VA writer that chooses to hedge will find that basis risk consumes capital. Basis risk should be avoided if at all possible. The use of volatility targeting in product design is a more complex example. To be sure, it is attempting to reduce the VA hedger's "realised volatility" risk which is somewhat akin to basis risk in that it consumes capital with no obvious benefit for the policyholder. However, it comes at the expense of a complication of the customer proposition.

Having designed the product and having decided the hedging and reinsurance policy, all approved at Board level, the broader organisation is then tasked with performing their various roles in marketing, servicing and indeed hedging. In my opinion, they can perform these roles confident that the risk aspects are being managed at the appropriate level. This is even true of the hedging process. Whilst some limited discretion will be delegated on the timings of re-balancing, the hedging personnel will generally not be making active risk allocation decisions – these will already have been made at the strategic level.

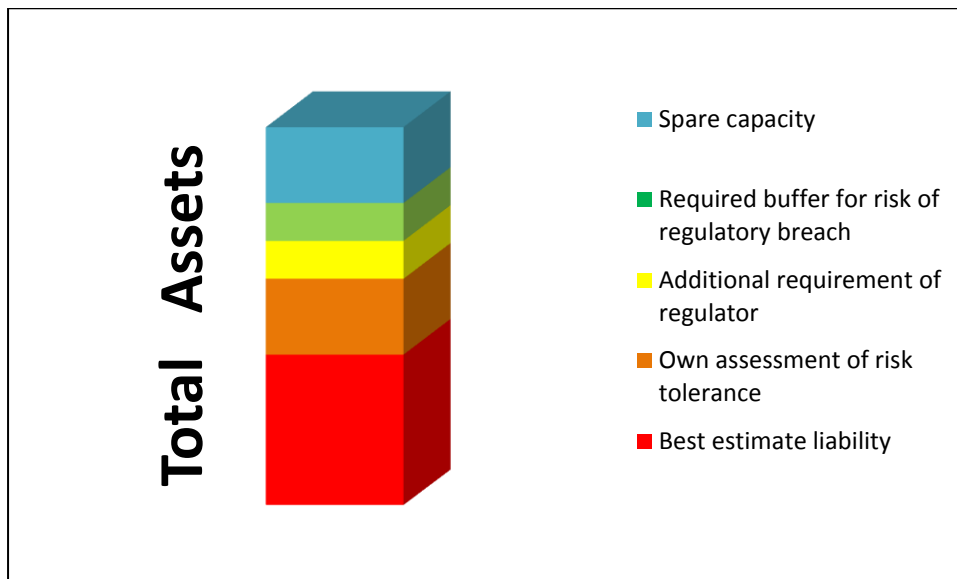
Other than the day to day management of operational risk the concept of a risk culture imbuing the whole organisation has not in my experience been the normal condition of a retail life assurance company and nor do I think it would be particularly helpful if it were. The big "disasters" in the UK were more of the miss-selling variety. Arguably a more robust ethics culture would have been more effective in preventing these failures than an enhanced risk culture. The failings on the banking side are more obviously laid at the door of risk management.

### Remaining within risk capacity

The Board has approved the products the company is to market and the risks that it will accept and those that it will hedge or reinsure. It will now need to monitor that the business remains within its **risk capacity**. That immediately poses the questions as to how to **measure** risk and how to measure capacity.

Under Solvency II these are more or less prescribed by regulation. We measure capacity as the excess of assets over technical provisions, the so called AFR, and we measure risk as the fall in this capacity that would occur on ½% of occasions over the following 12 months, the so called SCR. The AFR must exceed the SCR, it is as simple as that. In this determination we are required to consider all the risks that the company is exposed to and not just those prescribed by the Standard Formula. For VA writers this invariably involves an internal model.

The following figure illustrates the situation.



The bottom red layer of assets are required to meet best estimate liabilities. Under Solvency II these are separately identified but not so under Solvency I. Left to its own devices a Board would then require a margin for prudence<sup>1</sup>, the orange layer, to ensure that it will meet policyholder liabilities in all reasonably foreseeable circumstances. Under Solvency I the combination of the red and orange layers is identified as the technical provisions but there is no such separate identification under Solvency II. The yellow layer is an additional margin set by the regulator. Under Solvency I this is the usual EU minimum margin formula as enhanced by the CBI. Under Solvency II the yellow and orange layers combined are the sum of the Risk Margin and the SCR i.e. the split between yellow and orange is not explicit. Under Solvency II it is not really necessary for the Board to evaluate the orange layer in the above figure.

And now to the main point of the figure. The Board identifies a completely separate risk from that of failing to meet policyholder liabilities. That is the risk of breaching the regulatory requirement itself. It is sometimes stated that the company has zero tolerance or appetite for being in breach of the regulatory requirement. That is a somewhat fatuous statement as this zero tolerance is dictated by regulation. However, it can only be a statement about the present balance sheet at time 0. It cannot be interpreted as a zero tolerance for a possible breach of the regulatory requirement over some future period. That would require infinite capital.

This leads to the concept of a buffer over the regulatory requirement. The situation is really quite similar under either regulatory regime and a similar buffer would seem to be consistent. It is to be emphasised that the buffer is not there so that policyholder liabilities will be even more certainly met than is required by regulation - the regulatory requirement is quite adequate protection in that regard. When establishing a buffer the Board is addressing a completely new risk – the risk that at some time in the future the company might breach the regulatory requirement.

---

<sup>1</sup> Strictly speaking, under Solvency I this margin for prudence is determined by the Appointed Actuary.

Assuming the regulator is not going to prescribe the appropriate level of buffer, the company and its Board are now genuinely on their own in determining their tolerance for the risk of a future breach of the regulatory requirement. The issues which they would consider when assessing what that tolerance might be are, amongst others:

*On the downside*

- i) The reputational risk (with the regulator) in incurring a breach.
- ii) Perception in the marketplace and peer pressure in general.
- iii) And simply the sheer hassle factor in having to address a breach.

Note that the threat to meeting policyholder liabilities is not listed above. That is because the “yellow” layer in the above figure, representing the additional regulatory prudence over any self assessment, should ensure the matter is addressed long before the Board perceives that an unacceptable risk has developed on this front.

*In mitigation (of the risk of future regulatory breach):*

- i. The ease of access to fresh capital. A Board which had an undertaking from a financially sound parent to supply capital upon request could increase tolerance (reduce the buffer) in this regard.
- ii. Issued but unpaid capital could have a similar effect. Note that whilst this would be inadmissible for meeting the regulatory requirement, the Board is free to “admit” it when considering the buffer.
- iii. Any contingency measures in place to rapidly facilitate drawing down additional capital.

In actually quantifying the buffer (the green layer in the above figure) the Board may take some approach such as having less than an X% chance of a breach of the regulatory requirement in the next 12 months.

For ease of communication the buffer might be expressed as a percentage of the SCR (under Solvency II). The following table makes the simplifying assumption that the company’s aggregate risk profile is Normally distributed and that the SCR is the 99.5<sup>th</sup> centile of that distribution. It shows the probability of a breach of SCR in the next 12 months given a particular percentage of SCR buffer.

***Probability of a breach of SCR in next twelve months***

<b>Buffer % of SCR</b>	<b>Probability of breach</b>
10%	40%
20%	30%
30%	22%
40%	15%
50%	10%
60%	6%
70%	4%
80%	2%
90%	1%

Unlike the SCR, which in theory gets continually reset, the buffer would be reset annually for if it is continually reset we get into a silly exponential spiral of needing a buffer for the buffer. Though a traffic light system could also be applied whereby if we cross some alert levels mid year the Board may consider mitigating actions, of which the most obvious would be tapping into any available sources of additional capital.

## Risk Limits

As mentioned above the concept of risk limits delegated down to the individual agents of the company is not very prevalent in life assurance<sup>2</sup>. However, for VA companies in particular, it is highly desirable that the Board monitors the main segmental risk exposures, possibly setting alert levels, or limits if you will, for each segment.

This segmental analysis of risk exposures serves several purposes. Firstly it can give early warning signs of the possibility that overall capacity might become exhausted. A second reason is to spot concentrations building up as these may undermine any assumed diversification benefits in the determination of capacity. Another purpose is to highlight shifts in risk exposure. This can happen especially with VA business as the risks can change quite significantly relative to each other according to the extent to which the guarantees are in or out of the money. This in turn can cause rethinks on whether or not to hedge the various market risks.

The limits are set to best provide these various signals. One approach is to allocate the overall risk capacity between the various risk segments in proportion to what the exposures would be expected to be under the business plan. This in turn will depend on such matters as the chosen hedging strategy. But for clarity, in my opinion, these segmental alert levels or limits are just signals; they are not indications of breaches of risk appetite.

## Risk and Return

I think Figure 2 in the CBI paper makes a bit more sense (to me) if the axes are reversed and the Risk Capacity becomes a red horizontal line.

It seems to derive from the CAPM model in which the appetite (this time meaning “desire”) for an asset varies directly with the expected return on the asset and inversely with the risk attaching to the asset.

This would seem to overcomplicate the RAS. A simpler approach is to consider the minimum return required on some measure of risk itself. We could reasonably set this minimum return to be a constant irrespective of the source of the risk. And what better measure of risk than the capital which it consumes and what better measure of capital than the SCR (plus Risk Margin) set by the regulator (under Solvency II).

Different players will have different requirements for their return on this “½% risk” capital though Solvency II posits that the average requirement in the marketplace would be 6%. It is moot whether companies should also require a return (in excess of risk free) on their buffer capital. It should be noted that whilst the buffer may easily diminish the buffer capital is itself, by definition, at less than ½% risk and so a requirement to have an excess return on this buffer capital could possibly be ignored.

---

<sup>2</sup> An exception would be underwriting limits such as the maximum sum assured that would be accepted on an individual life or without a medical. These are set at a strategic level and applied more or less mechanically.

In short, the interaction between return and risk is best catered for within the RAS by having a minimum uniform requirement for the internal rate of return on the capital consumed by new business.

***Brian Woods, July 2014***

*Bio:*

Former Appointed Actuary and Finance Director at Eagle Star (to 1990) and Ark Life (1990 – 2006)

Currently INED for AXA Life Europe, AXA Life Invest Reinsurance and SMI

Chairman of the SAI working party established to respond to the 2009 VA consultation conducted by the CBI

Presented a paper on VA hedging effectiveness to the SAI in April this year