Reserving for annuities: a complicated business!

The current UK implementation of the Solvency II Matching Adjustment (MA) framework has resulted in a complicated system for reserving and assessing capital requirements for annuities. Reserving for annuities involves MA asset and liability eligibility rules, Fundamental Spread (FS) specifications (with their own array of LTAS, PD, CoD elements), numerous matching tests, some involving MA portfolio components defined with reference to the above elements of the FS, and a Risk Margin calculation that is set to become more elaborate under proposed reforms. On top of that, solvency capital assessment typically then involves an internal model that includes each firm's unique approach to modelling the same vanilla corporate bonds (and other more illiquid assets) together with a model of how the calibration methodology for Fundamental Spreads might be changed in stress scenarios.

Complicated regulatory rules tend to beget complicated solutions. And the MA framework provides an interesting case study in this respect in the form of the tens of billions of pounds of internal securitisations – complicated structures that are engineered solely to meet the MA eligibility rules.

The UK's implementation of the Matching Adjustment (MA) framework is currently undergoing reform as part of the broader process of implementing a UK-specific version of the Solvency II framework (Solvency UK). The indications provided by HM Treasury (HMT)'s SII Consultation Response of November 2022 and its draft legislation published in June 2023 suggest that this reform opportunity will not result in a simpler MA framework. On the contrary, changes in asset eligibility requirements together with the introduction of additional PRA supervisory tools are likely to result in an already overly complicated framework becoming even more so.

Why is the MA framework so complicated?

It might be argued that the complexity of the UK's MA implementation is evidence of an over-zealous UK regulator. In my personal view, the answer is a bit deeper and more interesting than that. This complexity is a near inevitable consequence of the way in which the fundamentals of the MA are developed in the SII legislation. The root source of the convoluted character of the UK's MA implementation is the ambiguity around what the MA is technically supposed to represent in the SII legislative framework. This may sound like a rather abstract line of argument. But when probabilistic modelling methods are used to determine overall levels of required capital, a lack of clarity or coherence in the model's fundamental metrics can create the need for *ad hoc* bells and whistles to try to ensure the model's overall outcomes behave reasonably. The result is a regulatory framework that is more complicated and costly and less transparent than it would need to be if the starting point of the capital metric was clearer and more coherent.

More specifically, the fundamentals of the MA framework represent a confused hybrid of two very different approaches to defining capital requirements (one based on short-term valuation solvency and the other based on long-term cashflow funding). At the inception of SII this hybrid approach was arguably necessary because of the need to fit the MA into the confines of the wider SII 1-year Value-at-Risk (VaR) capital framework. The Solvency UK MA reforms could remove the confusion and complexity of the hybrid approach by deciding to use either one of the two fundamentally different approaches to capital that the MA currently conflates. This would result in a simpler, more transparent and less costly regulatory solvency system without necessarily having material implications for the levels of capital held by MA firms. These points are developed further below.

The MA discount rate is not the illiquid risk-free rate

In the Solvency II Value-at-Risk capital framework, a best estimate liability valuation is calculated by discounting the liability's projected cashflows using a risk-free yield curve. This is a basic step in the determination of a hypothetical liability transfer value, which plays a fundamental role in SII's 1-year VaR capital framework.

The MA fits into this framework by allowing MA-eligible liability cashflows to be discounted using an adjusted yield. The most obvious rationale for this adjustment would be that MA-eligible liabilities are highly illiquid (the policyholder is not entitled to surrender their annuity back to the insurance firm for a cash settlement), and illiquid assets / liabilities should be valued at a discount to otherwise equivalent liquid assets / liabilities, if this is what is implied by market prices. So, following this line of thinking, the MA adjustment could simply represent the difference between how liquid and illiquid risk-free assets are currently valued in the market.

This 'illiquidity adjustment' would be a perfectly logical extension to the Solvency II system – SII uses a 1-year Value-at-Risk (VaR) capital metric, and, if illiquidity impacts on market prices and hence on illiquid liability transfer values, this effect should be captured in the framework's treatment of illiquid liabilities (in base and in stress). The illiquidity adjustment isn't an easy quantity to observe in market prices (illiquid assets generally don't have very transparent market prices), but there are plausible ways of estimating this 'illiquidity adjustment' and hence deriving an applicable illiquid risk-free discount rate.

But the MA framework does not attempt to estimate the market's illiquid risk-free rate. We can see this by noting that the market's illiquid risk-free rate is not a function of the specific assets that insurance firms choose to use to back illiquid liabilities (just like the liquid risk-free rate does not depend on the specific assets that back liquid liabilities). The market's illiquid risk-free rate is also not related to the quality of cashflow matching that the MA asset portfolio has achieved, and nor is it dependent on the nature of the asset portfolio's cashflows. The calculation of the MA discount rate is a complicated process that is a function of several things that are clearly irrelevant to an objective estimate of the market's illiquid risk-free rate and to the estimation of a current transfer value for MA liabilities.

VaR and transfer values

The illiquid risk-free rate would most naturally be estimated by observing the rates at which very low-risk illiquid lending is being originated in the market (and perhaps extrapolating these rates to an estimated zero-risk case). I have written about possible approaches to implementing this previously[1], but before we go too far down this rabbit hole, let's take a step back. We are interested in the illiquid risk-free rate because we would like to use it as input to the assessment of the current transfer values of long-term illiquid liabilities with guaranteed cashflows. This points towards a set of market prices that are more relevant than illiquid asset origination rates: the prices at which these illiquid liabilities are currently being transferred from pension funds to insurance firms.

What better guide to the current transfer values of annuity liabilities can there be than the market prices at which these liabilities are currently being transferred? We won't find these prices on a Bloomberg screen, but the UK pension risk transfer market is now a very well-established competitive market with frequent transactions. Annuity firms, pension funds and regulators all have a good understanding of prevailing market pricing levels.

[1] The Mortgage-Implied Illiquidity Premium | Craig Turnbull FIA

If market buy-out pricing levels were used as the basis for annuity technical provisions, the need for a whole swathe of complicated reserving calculations would be removed at a stroke: all the MA calculations and Risk Margin calculations that are constructed to produce a hypothetical transfer value would be made redundant by the direct use of observed transfer values to set the basis for technical provisions. And there would be less need for special matching tests or asset eligibility requirements to regulate the amount of MA benefit that is being generated, as the MA benefit would no longer be a function of firm-specific asset choices.

At current market pricing levels, the use of a buy-out basis to value technical provisions would typically result in a small increase in technical provisions relative to what is produced by the current MA framework. This is to be expected. As is discussed further below, the MA discount rate incorporates (at least a significant part of) the expected return on the credit-risky MA asset portfolio into the liability discount rate. If this discount rate was used in pricing, it would imply the annuity firm's capital provider does not require compensation for bearing the credit risk that they are exposed to when backing the business with an MA asset portfolio. (And MA portfolios do tend to take material credit risk. This is why MA firms hold material credit risk solvency capital requirements for the credit risk in their MA portfolios.) Annuity firms naturally require compensation for bearing this risk, and this is reflected in the price they charge to take on the liabilities. So, it is economically intuitive that the market buy-out basis tends to be stronger than the SII MA basis, given how the MA's Fundamental Spread is determined.

Back-of-the-envelope calculations (for example, using published industry estimates of current buy-out pricing and the MA disclosures in Solvency and Financial Condition Reports) suggest the increases in technical provisions implied by this approach would consume only a fraction of the very significant capital buffers that annuity firms tend to currently hold. The implications of this approach for the Solvency Capital Requirement element are less clear-cut – it would depend on how buy-out transfer values are assumed to behave in stress, and how this compares to the way internal models currently assume Fundamental Spreads behave in stress. The SCR implications of this change in approach to defining the technical provisions would probably be second order.

The above points imply the current capitalisation levels of the industry would be quite adequate under this more direct transfer value VaR approach. And this simpler, stronger approach to setting reserves and capital requirements could naturally be accompanied by a reduction in the size of capital buffers that firms typically target (for example, a 130% SCR buffer under a transfer value VaR might be regarded as comparable to a 150% buffer under the current SII MA system). So, the net result of the above approach need not be a change in the amount of capital that annuity firms hold. Rather, its main impact would be a substantial simplification of the regulatory framework and a reduction in its associated costs that would benefit all stakeholders.

What is the current MA discount rate?

If the MA discount rate isn't an estimate of the illiquid risk-free rate, then what is it? As currently implemented, the MA discount rate is a measure of the current expected return of the MA asset portfolio, on the assumption its credit quality is re-balanced back to its starting position at the end of every year (and subject to the LTAS cap). MA asset portfolios tend to bear some investment-grade quality credit risk, and those assets are priced to offer an expected return for bearing credit risk - a credit risk premium (CRP). A credit-risky asset's CRP is not directly observable, and its estimation is subject to judgment. But it is difficult to plausibly argue it is zero.

I argued above that the MA is not an illiquidity valuation adjustment, and it is notable that the MA refers to matching and not to illiquidity. Cashflow matching plays a central role in the MA

framework. This is because the run-off cashflows of a well-matched portfolio, managed on a hold-to-maturity basis, have no exposure to changes in market prices that are solely the result of movements in credit spreads (i.e. where the spread change is not accompanied by any change in the asset's expected future default losses, and hence the spread change simply represents a change in the expected return of the asset). This is the essence of the rationale for the MA: matched assets and liabilities, held on a buy-and-hold basis, are not exposed to credit spread changes, and therefore should not hold capital for risk related purely to variations in market credit spreads.

The difficulty with this argument is that, in a VaR framework, risk capital *should* be required for 'pure' credit spread variations, even when assets and liabilities are perfectly cashflow matched. Credit spread variations will impact on the market value of credit-risky assets in a way that will not impact on the transfer cost of illiquid risk-free liabilities - these transfer values should be a function of the risk-free illiquidity premia but not asset risk premia. Changes in illiquid asset values that are driven by changes in illiquidity premia may be offset by corresponding changes in transfer costs for illiquid liabilities. But a rise in asset risk premia will negatively impact on asset values in a way that may not be wholly transmitted to the transfer values of guaranteed liabilities. Whilst the providers of the risk capital that backs transferred liabilities will expect a higher return for a given asset strategy following a rise in asset risk premia, they will also rationally *require* a higher return as participation in the higher risk premia that they can access in alternative ways in the market. So, an increase in asset risk premia does not automatically equate to lower transfer values in the way that it does automatically equate to lower asset market values for credit-risky assets.

An empirical analysis of the relative volatilities of risk premia and illiquidity premia is no mean feat, but there is strong empirical evidence that liquid assets (such as equities) exhibit material variation in risk premia over time, and there is no obvious reason to suppose this doesn't similarly applying to credit-risky assets including those that back annuity business.

So, the no-capital-required- for-credit-spread-volatility-when-cashflow-matched argument doesn't naturally apply in a VaR capital measure, even for illiquid liabilities. And the expected return on the MA asset portfolio isn't a relevant variable in a VaR capital measure. But both play a major role in the MA framework. The no-capital-for-credit-spread-vol argument and the expected return metric are not relevant to a VaR capital measure, but they *do* feature prominently in the most obvious alternative approach to VaR as a probabilistic solvency capital methodology: a 'run-off' capital measure, where a total asset requirement is determined by assessing the amount of assets that is required to fund all liability cashflows as they fall due over the run-off of the business, at some defined probabilistic threshold.

This suggests the SII MA framework is really an attempt to implement a run-off capital measure within the confines of the overall SII VaR framework. There is nothing wrong with that objective. The run-off approach to defining a probabilistic capital requirement is a perfectly valid and respectable alternative to VaR and is one that has been used extensively in other actuarial settings. The two approaches have interesting relative merits and this has resulted in each one moving in and out of fashion over time through the decades since actuaries started using probabilistic capital measures for long-term liabilities.

Like the MA, a run-off approach is sensitive to some challenging assumptions about the long-term expected returns of assets. But a direct implementation of a run-off capital measure would be a lot simpler and more transparent than the current MA framework. For example, the risk benefits of cashflow matching would be 'automatically' captured by the run-off modelling implementation.

Matching tests and asset eligibility rules would therefore be redundant (and hence so too would be the costly asset re-structuring exercises that are a response to those rules).

The current approach of dressing a run-off calculation in VaR clothing has created substantial cost, complexity and confusion. *If everyone agrees that the purpose of the MA is simply to move the reserving and capital requirements of MA business from a VaR capital basis to a run-off basis, it seems a pity not to find a way to simplify our world accordingly!*

It could reasonably be argued that the constraints of Solvency II meant that squeezing the intended MA run-off capital calculation into SII's VaR framework was unavoidable. The Solvency UK reform process opened the door to resolving this issue. The Solvency UK framework could clarify what capital measure is intended for MA business – transfer value VaR or run-off – and then implement the simplified framework that would be implied by either choice. *Either choice is quite reasonable and both choices would result in a simpler, more transparent and effective solvency framework for illiquid liabilities than the current MA.*

So what?

If there is one thing that HMT, the PRA and the ABI all seem to agree on it is that there is little appetite for the sort of fundamental re-wiring of the MA framework that is alluded to above. There will likely be lots of MA reforms, but they will ultimately be incremental adjustments to the current framework that will tend to add to the complexity of the regulatory solvency framework rather than reduce it.

This is probably not an unusual feature of the general dynamics of regulatory rulemaking (in any industry), and the important question that follows is: does it really matter? Annuity firms hold significant capital buffers in the current system, and the above discussion has argued that the current overall capitalisation levels of MA balance sheets are very likely adequate. The MA framework probably keeps actuaries and securities lawyers a bit busier than is strictly necessary, but these are long-term promises written on a huge scale, and perhaps it is just the inevitable nature of the real world that there are more complicated regulatory rules and processes than appear to be theoretically optimal.

I think that's a fair and reasonable counterpoint to the above discussion. But the Solvency UK reform process offers an opportunity to address the fundamental difficulties at the very core of the MA. If this nettle is not grasped, the MA seems destined to continue to be a source of unnecessary regulatory complexity and cost for decades to come. After seven difficult years with the current MA, this reform opportunity appears fortuitously timed. But without bolder thinking, it may come to be recognised in future years as an opportunity missed.

Nonetheless, the above analysis may provide actuaries with some insight into ways to explain and validate the strength of the capital positions produced by the MA framework – for example, by comparing and contrasting their MA reserve and capital results with the results produced by either a run-off approach or a transfer value VaR approach. Such analysis can provide independent validations of the MA capital position that may support work in formal attestations or in developing the firm's own economic understanding of their capital position.

Craig Turnbull FIA

June 2023.