Superfunds: Some Brief Notes on Risk, Certainty, Capital and Value

The Pensions Regulator's recent guidance[1] provides some immediate clarity on what it requires of superfunds as they go about their business of consolidating UK Defined Benefit (DB) pension funds. The guidance has been widely welcomed, and it is easy to see why: the UK DB pension fund sector is a problem that is in need of solutions.

According to the Pension Protection Fund's 2019 Purple Book[2], the 5,436 schemes in the PPF's eligible universe had total assets at 31st March 2019 of £1,615bn and total liabilities (on a buy-out basis) of £2,091bn, implying an aggregate deficit of some £476bn. Market movements since that valuation date have moved against DB pension funds' aggregate buy-out solvency positions. Falls in equity markets, increases in credit spreads and falls in interest rates (both real and nominal) are all movements in the wrong direction for the aggregate buy-out solvency level of UK DB pension funds. The aggregate buy-out deficit of the PPF eligible universe today is likely to be materially in excess of half a trillion pounds.

Moreover, the financial robustness of the employer sponsors who are expected to, eventually, plug this funding gap will, in many cases, have been weakened by the economic consequences of the Covid-19 crisis. It increasingly stretches credulity to suppose that half a trillion pounds of future UK plc profits will find its way into predominantly closed DB pension funds. Of course, it is quite possible that future market movements will eradicate the buy-out deficit without any further sponsor contributions. Interest rates could increase by a few hundred basis points; equity markets might double. Optimists are usually popular but, to quote the cliché, hope is not a strategy.

In this context, innovative thinking and fresh capital are likely to receive a fair hearing. And that is what has happened. The Pensions Regulator has clearly set out its requirements and expectations with regards to the appropriate and sound management of superfunds, particularly in their transactions with DB pension funds and their sponsors.

Assessing the adequacy of the superfund capital buffer

When a superfund acquires a DB pension fund's assets and liabilities, the pension fund's employer sponsor is relieved of all future obligations in respect of the pension fund. Of course, the sponsor is likely to make a final contribution at the point of transferring the obligations to the superfund in exchange for this relief (and this contribution may go directly into the pension fund or into the superfund's capital buffer, or a combination of the two). And the superfund will put some of its own capital into the capital buffer.

From the perspective of measuring the security of members' benefits that have been transferred to a superfund, the key question is not how much of the capital buffer is funded by the sponsor and how much by the superfund, but whether the buffer is big enough to provide a high probability that the members' benefits can be paid in full. This is naturally one of TPR's main areas of focus. Their recent guidance states that members whose pensions have been transferred to a superfund should receive their full benefits with '*a high degree of certainty*'.

The guidance sets out some specific quantitative measures by which it will consider the adequacy of the capital funding of a superfund. In particular, the TPR will assess the adequacy of the capital

[1] https://www.thepensionsregulator.gov.uk/en/document-library/regulatory-guidance/db-superfunds

[2] https://www.ppf.co.uk/sites/default/files/2020-01/Purple%20Book%202019.pdf

buffer by whether the total assets (scheme assets plus capital buffer) are sufficient to imply that there is a 99% probability that the market value of the total assets will exceed the scheme's Technical Provisions after five years. This capital requirement metric is similar in form to the Value-at-Risk capital definition used in Solvency II. Solvency II's Solvency Capital Requirement is based on a 1-year projection horizon and a 99.5% probability that the market value of total assets exceed liability values. TPR's choice of a 5-year horizon instead of 1-year is essentially a matter of taste, though it may be aligned with the expected 'interim period' before superfund legislation becomes operational and supersedes this guidance. And, as we will see below, there are other aspects of the guidance that also use a 5-year horizon.

It is also important to note that TPR is adopting a principle-based approach to the capital assessment, similar in spirit to that used in Solvency II Internal Models. That is, the regulator does not prescribe the model that should be used in the capital assessment. Instead, the onus is put on the superfund to develop an appropriate capital assessment and then (privately) disclose and justify its methods to the regulator.

Risk and a High Degree of Certainty

In my recent paper on actuarial methodology[3], I argued that probabilistic quantitative capital metrics of the kind prescribed in the TPR guidance are profoundly difficult to calculate with reliable accuracy. This is especially the case where economic and financial market risks are a significant contributor to the capital requirement. It is likely that investment risk-taking will be fundamental to superfund economics and that economic and financial market risks will therefore be the main drivers of the TPR capital requirement.

Why do I assert that these probabilistic capital measures are very difficult to reliably implement? After all, haven't they been used globally for decades in banking and insurance? Isn't there today a vast array of sophisticated statistical and computational techniques available for the implementation of advanced financial modelling methods? Undoubtedly, the answers to both these questions are 'yes'. But the historical performance track record of these capital models is fairly patchy. And they all rely on limited historical data and structural assumptions about how the future is related to the past. There is a 'sceptical' school of thought that argues that the future of the economy and the future behaviour of financial markets (and other forms of social phenomena) are the result of processes and conditions that are inherently too complex, unobservable and non-stationary to be amenable to reliable probabilistic projections.

When a student nuclear physicist learns that today the probability of an alpha particle tunnelling through the wall of the nucleus of a U-238 atom is 10^{-38}, they have good reason to believe that this will also be true tomorrow. At the very least, they will have an awareness of the conditions necessary and sufficient for this probability to apply tomorrow. In contrast, we simply don't have a firm logical basis for asserting that the 10-year interest rate 5 years from today will be greater or less than X% with some measurable probability; or what specific conditions are necessary or sufficient for X to be of a given size.

We can, of course, specify a model and calibrate it to historical financial and economic data. The model may provide a very satisfactory description of historical behaviour, and the quality of the model's fit to data may be assessed using rigorous statistical methods. But the relevance of this analysis is contingent on the model being the 'right' one for the economic dynamics of the future.

---

[3] Chapter 5 of https://craigturnbullfia.com/wp-content/uploads/2020/04/On-The-Methodology-of-Actuarial-Science.pdf

And our understanding of the economy and financial markets is not sufficient for us to have any means of determining the 'true' stochastic process for these phenomena, or how it changes as the myriad influences on the economy unpredictably change[4].

To be clear, this is hardly a new perspective. Many readers will be familiar with Frank Knight's famous distinction between risk and uncertainty. And, over the past one hundred years or so, economists as varied and as notable as Max Weber, John Maynard Keynes, Karl Popper and Friedrich Hayek have expressed similar perspectives on the severe epistemic limits that probabilistic modelling of the economy must face. Frank Redington, who is often regarded as the most important actuary of the twentieth century[5], also expressed views of this kind. To quote from his essay *Prescience and Nescience*:

"The continued instabilities we observe in our experience are not the regular instabilities of probability theory…the conditions never repeat themselves".

The thinkers mentioned above expressed these views many years ago. It might be thought that I am describing an antiquated perspective that has been made redundant by modern advances in statistics and computers. But the intellectual fashion amongst financial thinkers in recent years seems to be moving towards rather than away from this philosophical outlook. Since the Global Financial Crisis, an array of thought-leaders in financial policymaking circles have voiced strong concerns about the limitations of probabilistic models when used in assessing the capital requirements of financial institutions[6].

To illustrate this point in an actuarial context, we can consider the past performance of actuarial models in making long-term forecasts of the probability distributions of the things that drive the future solvency of a DB pension fund. How accurate a job have major actuarial models done of estimating the probabilities associated with the behaviour of key actuarial phenomena? This is clearly a very broad question that could form the basis of a very extensive empirical investigation. It is also a bit tricky, as we are interested here in multi-year modelling performance, and so it inevitably takes a while to establish the empirical performance of a model. And as we are particularly interested in the tails, we would like to observe many independent multi-year experiences in order to rigorously measure the performance of the models. These models have only been around for a few decades and the models and calibrations are not always publicly available. So, we do not have ready access to the huge pool of independent empirical results that would allow us to make a comprehensive analysis. Here, we will merely try to illustrate the point anecdotally with an example from a major DB pension actuarial model.

Suppose we are in the mid-1990s and we are interested in the multi-year projection of the long-term index-linked gilt yield. The 5-year change in this yield is likely to be a material driver of the capital

---

[4] Social phenomena are exposed to an interesting source of change that is generally absent of natural phenomena: a change in the human understanding of the behaviour of the phenomena may influence and change its behaviour. Philosophers have discussed this idea since at least the late 19th century, and it has been a significant theme in macroeconomics since Lucas' Nobel prize-winning work of the 1970s.

[5] This is a view expressed in Tan Suee Chieh's recent IFoA Presidential Address. The Address is also more generally aligned to the perspective of this discussion on uncertainty and the epistemic limits of modelling social phenomena, see: https://www.actuaries.org.uk/system/files/field/document/200616-IFoA-Presidential-Address-Strategy-Culture-and-Imagination.PDF

[6] For example, Mervyn King, the former Governor of the Bank of England (his 2016 book *The End of Alchemy*); the economist John Kay (the 2020 book, *Radical Uncertainty*, co-written with King); and Andrew Haldane, former Chief Economist of the Bank of England (his 2012 paper, *The Dog and the Frisbee*).

metric prescribed by the TPR guidance for many DB pension funds (and superfunds). An updated version of the Wilkie model was published in the British Actuarial Journal in 1995 with a calibration for this yield. At the time, this model was the most popular stochastic asset model for use in UK DB pension actuarial work.

The chart below shows how the actual outcome for the 10-year index-linked yield compares with the projected probability distribution produced by the model published in 1995[7].

Exhibit 1: Percentiles of the projected index-linked gilt yield distribution from Wilkie (1995) model and the actual 10-year index-linked yield path

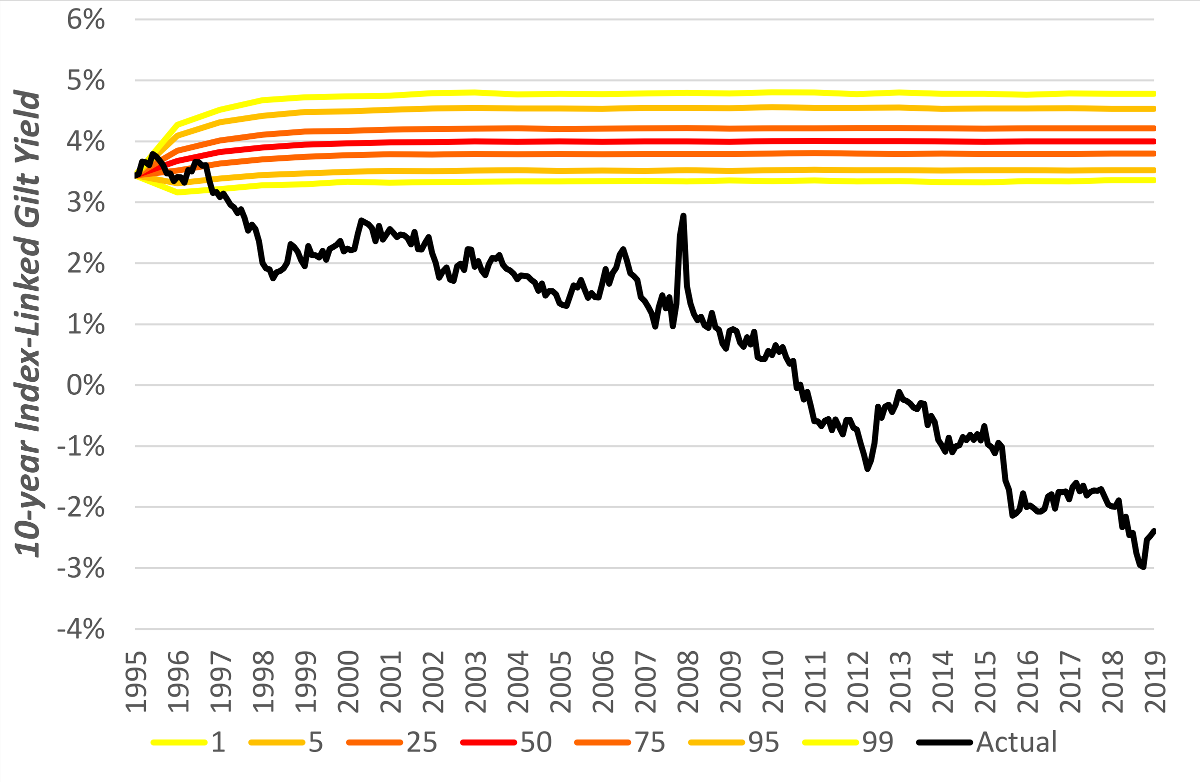

The chart shows that the model attached no meaningful probability to a real interest rate outcome as low as the one actually experienced. Indeed, the model structure assumed there was zero probability that real interest rates could ever go below zero. But this is not intended as a criticism of the Wilkie model. It is not the specific modelling and calibration methods of this particular model that are of interest to us in this discussion.

This anecdotal example is used to make a much more general point: making reliable probabilistic forecasts for economic phenomena such as the long index-linked gilt yield is really beyond us. Professor Wilkie made some difficult expert judgments when constructing this model. He could have used different modelling assumptions and different calibration methods. Some of these alternative modelling approaches may have been more sophisticated than those used in his model. However, it is contended here that no conventional method of modelling and calibrating a probability model for the long-term projection of the index-linked gilt yield in 1995 would have attached a meaningful probability to the yield being as low as -2% by the mid-2010s. *And if the model did suggest such a*

[7] The historical data for the 10-year index-linked bond yield has been sourced from the Bank of England website: https://www.bankofengland.co.uk/statistics/yield-curves

scenario had a meaningful probability, it would almost certainly have been rejected with incredulity by experts at the time.

When this model was calibrated in 1995, the index-linked gilt had only been in issuance for 14 years. Now it has been around for almost 40 years. With so much more data, perhaps we now have a much better grasp of how it may behave in the future. Moreover, it might be argued that statisticians, actuaries, data scientists and others have made great strides in the statistical modelling of economic phenomena since 1995, and we can therefore now produce a much more reliable estimate of the 99th percentile 5-year change in the long index-linked yield. I would treat such optimism with scepticism. As Redington put it, "If we cannot foresee the future then we cannot! No proliferation of polysyllables or multivariate analysis will succeed (other than in deceiving ourselves)."

At the time of writing, the UK 10-year index-linked bond yield is around -3.0%. There is no way of reliably estimating the 99th percentile of today's probability distribution for the 2025 index-linked yield. The same goes for nominal interest rates, the level of the FTSE 100, investment-grade credit spreads and most of the other exposures that will drive the TPR superfund capital assessment.

So, perhaps I have convinced you that we have no idea what the 99th percentile of the 5-year change in the index-linked yield or the 99th percentile 5-year equity return may be. Perhaps you didn't need convincing in the first place. So what? I argue in the methodology paper that these models are nonetheless of significant value, despite their inherent fallibility. But, I argue, their value is derived from the insights they can deliver to an expert who understands the model's limitations and uses it as an aid to their professional work. The value is not derived by using the model directly as an answer-generating machine.

A superfund might reasonably say, however, that if TPR compels us to calculate the number, we have to do it and do it as well as we possibly can, no matter how convincing my methodological reservations about the epistemic limits of the calculation may be. That is, of course, an entirely appropriate position. To the superfunds who have to calculate the numbers, and TPR who must assess the adequacy of their modelling, this sceptical perspective can still be used constructively. Points such as these naturally follow from this perspective:

- Do not be overly impressed by 5 significant figure maximum likelihood estimates for model parameter fits to historical data. Such statistics are only meaningful if we know the specified model is the actual process that generates the modelled phenomena. In reality, for economic and financial phenomena, we can be very sure it is not.
- Use forward-looking and independently developed stress scenarios to 'sanity-check' the probability distributions. But, again, be aware that expert views of the future may ultimately suffer from the same forms of bias as the probability model calibration - an implicit reluctance to recognise that the future could be profoundly and almost unimaginably different to the past. If we asked a group of experts in 1995 how low index-linked yields might be in 2020, very few would have said -3%. And those that did would probably have been dismissed as cranks.
- Current option prices can be a useful source of information about forward-looking risk. Option prices on equity indices and (nominal) interest rates are easily found (though that is not the case for index-linked yields). For example, the price of a 5-year option on a 10-year (nominal) interest rate swap is currently consistent with an interest rate volatility of around 0.6% pa. There are many good reasons why we shouldn't interpret this number as a 'best estimate' for future interest rate volatility: the model that derives the volatility from the option price is a simplification of the real world, and there are many factors that may be

empirically important for option pricing that are not explicitly captured in the pricing formula. So, we have good reason to expect the 'option-implied' volatility to be an upwardly biased indicator of future volatility. But if the VaR model assumes a volatility that is, say, half of that implied by option prices, it should at least give some serious pause for thought.

- And remember the chart above when interpreting what a VaR number can really tell us.

Comparing superfund risk and employer covenant risk

The above perspective may lead to the conclusion that meeting TPR's 'high degree of certainty' will require a lot of capital. Economic and financial market risk inevitably comes with a lot of uncertainty, and a high degree of certainty must therefore entail a high degree of capital. However, the above discussion perhaps also suggests that a high degree of certainty is a rather aspirational target for a DB pension fund sector with an aggregate deficit of half a trillion pounds or so. Perhaps, for many DB pension fund members, the days of a high degree of certainty are unfortunately in the past (at least in terms of receiving full promised benefits).

It might be argued from here that the important question is not whether the superfund arrangement provides a high degree of certainty that members' benefits will be paid in full, but whether it provides a *higher* degree of certainty than the status quo. This implies a more complicated relativism: if two pension funds are currently 70% funded on TPR's Technical Provisions basis, and one has a AA-rated sponsor and the other a B-rated sponsor, a lower superfund capital buffer is necessary to deliver an improvement in member security for the members of the latter pension fund than the former. The trustees of the pension fund with the AA-rated sponsor may be quite relaxed about the security of their members' benefits and may have no intention of considering the superfund route. For the pension fund with the B-rated sponsor, is it reasonable to suppose any commercially viable route can deliver a high degree of certainty for the full payment of member benefits? In the absence of financial alchemy, it sounds a bit implausible. But perhaps the superfund route isn't intended for pension funds in that situation either.

So, in what type of pension fund situations does TPR expect superfunds to participate? The TPR guidance makes clear that the superfund route is not for schemes that are capable of funding a buy-out now or who are on course to do so within the foreseeable future (which the guidance suggests could be considered as the next five years). So, the guidance requires superfunds to deliver a 'high degree of certainty' to pension fund members whose pension funds do not have a realistic prospect of being able to afford a buy-out in the next 5 years. Superfund arrangements may be of limited interest to DB pension funds with very secure sponsors. But the sponsor will need to be in a financial position to contribute to the superfund arrangement to a degree sufficient to make the attainment of a high degree of certainty for the payment of full members' benefits a commercially viable proposition for the superfund. This might work well for some segment of the pension fund universe, but it seems unlikely to be a formula, in its current form, that can really move the dial for the aggregate UK DB pension fund position. To understand the sort of scenarios where the superfund solution is potentially workable for all interested parties (sponsor, superfund, TPR and members), we need to dig a bit deeper into the economics of superfunds.

Superfund economics

As noted above, the profitability of the superfund business model is not of direct concern to the regulator. The regulator's concern is with the overall security of the members' benefits, and this is determined by, amongst other things, the total size of the capital buffer - how much of the capital buffer is funded by superfund capital and how much by sponsor capital is largely a commercial matter between the superfund and the sponsor. Nonetheless, we might be interested in better understanding the commercial viability of superfunds…after all, superfunds are not charities, and it is unlikely superfunds will play a useful role in the sector without an economically viable business model.

From an economic perspective, a key concept underlying the superfund business is that the superfund's equity takes the form of a call option on the total assets of the pension fund arrangement (scheme assets plus capital buffer), with a strike equal to the cost of the scheme liabilities. If the assets prove inadequate to meet the cost of the liabilities, the superfund receives nothing and loses the capital it has invested in the buffer (which can be thought of as the option premium). If the total assets prove more than sufficient to meet the cost of the liabilities, all the upside belongs to the superfund.

Let us take an extremely simplified example to conceptually illustrate this point. Suppose the scheme consists of a single, certain liability cashflow that will be payable in 10 years. Let's suppose this liability cashflow is 1.22, the risk-free rate is 2.0% (continuously-compounded) and the liability therefore has a present value of 1 today. The market value of the scheme assets today is 0.75. The superfund implements the TPR VaR calculation and determines that, given the investment strategy for the assets, the capital buffer today must have a value of 0.45 to meet the 5-year 99% VaR requirement (so, bringing the total assets available to support the scheme to 0.75 + 0.45 = 1.2). Further suppose that the outcome of the negotiation between the superfund and the scheme sponsor is that the superfund will fund 0.25 of the required capital buffer of 0.45 and the sponsor will fund the remaining 0.2. For simplicity, we assume both the scheme assets and the capital buffer are invested in the same static asset strategy, and that this strategy produces an asset volatility of 9% pa.

The economic value of the superfund's equity is a 10-year call option with an underlying asset value of 1.2, a strike of 1.22, an asset volatility of 9% and a risk-free rate of 2.0%. If we plug these numbers into the Black-Scholes formula, we obtain a call option value of 0.25. So, in this case, the superfund is economically 'breaking even': it has invested 0.25 into a structure and its stake in this structure has a current value of 0.25.

It might be argued that this option valuation relies on probabilities for future asset prices of the form that we have argued above are incalculable. But financial markets do trade uncertainty, and at least some forms of financial options trade with observable prices. The process of valuing an illiquid asset consistently with observable market prices is certainly not without its challenges and will generally involve some subjective judgment. But I would argue that valuation is an intrinsically more objective act of extrapolation (from observable prices) than the estimation of the probabilities that may underlie those prices. For the avoidance of doubt, more sophisticated option pricing techniques could be employed than the one used in this extremely simplified conceptual example.

Clearly, the reality of superfund economics is much more complex in real life. We need to consider buy out costs, the path dependency created by the TPR's funding triggers, dynamic and differential asset strategies, longevity risk and so on. But the essence of the economics of superfunds is determined by the optionality of their exposure to the pension fund's assets and liabilities. A few fundamental insights immediately follow from the example, despite its artificiality.

The superfund perspective: From the superfund's perspective, *the economic profitability of the transaction critically depends on the level of the TPR capital requirement*. In the above example, if the required capital buffer is assessed at 0.4 instead of 0.45, and this is reflected solely in the amount of capital invested by the sponsor, the economic profitability to the superfund becomes positive. The superfund would make an initial outlay of 0.2, and the equity valuation would be 0.211 (i.e. the above call option valuation with an underlying asset value of 1.15 instead of 1.2).

Alternatively, if a capital buffer of, say, 0.5 was required instead of 0.45, and this increase was again entirely funded by the superfund and not the sponsor, the transaction would become economically unprofitable. The superfund would be required to make an initial capital outlay of 0.30, and their equity valuation in this case would be only 0.289 (i.e. the above call option valuation re-calculated with an underlying asset value changed from 1.2 to 1.25). This result implies that, in this case, the superfund may be able to find alternative cheaper ways of obtaining the financial market exposures that would be created by the transaction.

We can also note that the optionality of the superfund's exposure means the superfund is economically incentivised to increase investment strategy risk (to the extent it does not result in a value-offsetting increase in the size of the capital buffer required by TPR).

The sponsor perspective: In the above example, the scheme sponsor makes a contribution of 0.2. Is this an attractive proposition for the sponsor? The status quo leaves the sponsor with an obligation to fund the 0.25 deficit (i.e. scheme assets of 0.75 less liability value of 1). But the sponsor is not obliged to fund all of that 0.25 deficit today. From an economic theory perspective, the 0.2 immediate contribution makes economic sense to the sponsor if that amount is less than the present value of their future deficit contributions after allowing for the sponsor default risk associated with those contributions. The value of the sponsor 'option to default' will depend on the timing of the deficit contributions under the status quo and the sponsor's credit quality[8]. The lower the sponsor's credit quality and the longer they can otherwise defer future deficit contributions, the less attractive the superfund route will be to the sponsor.

The TPR perspective: As noted above, the TPR guidance does not prescribe how to assess the 5-year VaR capital metric. The earlier discussion of uncertainty highlights that there is a very wide range of plausible assumptions that could be used in making this capital assessment: is the 99th percentile of the 2025 10-year index-linked yield -3% or -5%, for example. At this stage, TPR has not provided any more explicit guidance on how it expects the numbers to be produced, beyond the modelling being 'appropriate'. *Much therefore hinges on exactly how the TPR 'enforces' the assessment of the capital requirement, and what minimum standards emerge either from further public guidance or in private bilateral feedback to superfunds*.

The use of a principle-based rather than prescribed modelling approach has the advantage that it allows firms to tailor their modelling to the specific risks that their strategy exposes them to. But it can lead to unnecessary inconsistencies: if all pension funds are exposed to falls in the 10-year index linked yield, it makes little sense for different superfunds to use different estimates of how much the index-linked yield might fall over the next 5 years; the process would be a lot simpler if the regulator simply told everyone what 99th percentile 5-year fall in the 10-year index-linked yield is acceptable for use in the capital assessment. (I use the yield example for continuity, but the same point applies to any economic or financial market risk generated by the superfund investment strategy.)

Finally, from the TPR perspective, the uncertainty and judgment inherent in their defined capital metric means that they may find alternative indicators of the economic viability of the transaction,

---

[8] I explored the idea of valuation of the credit-risky stream of sponsor's deficit contributions in a paper published in the British Actuarial Journal in 2014, '*Market-Consistent of a Defined Benefit Pension Fund's Employer Covenant and its Use in Risk-Based Capital Assessment*'. (And I humbly note in passing that it won the IFoA's 2013 Peter Clark Prize!)

such as the above option analysis, useful in setting (formally or informally) the capital measure's minimum standards.

In Brief

The recent guidance from TPR clarifies what it expects from superfunds as they go about their business of consolidating UK DB pension funds. An important component of this regulatory guidance is the requirement for superfunds to assess the capital required to ensure there is at least a 99% probability that the market value of the total assets supporting the pension scheme (pension scheme and capital buffer) will be greater than the pension scheme's Technical Provisions after five years.

This note has explored a few conceptual points that are relevant in this context. There is inevitably very considerable uncertainty associated with an estimate of this type of quantitative probabilistic capital metric; the commercial viability of the superfund model is highly dependent on the level of capital it must make available to support the pension funds it acquires; and much therefore hinges on what minimum standards the TPR sets (publicly or privately) for superfunds' capital calculations.