

3 July 2019

Discussion Paper

PRUDENTIAL STANDARD SPS 515

STRATEGIC PLANNING AND MEMBER OUTCOMES

By David Bell [REDACTED]

David Bell is currently completing his PhD at UNSW. Previous career experiences include industry (Mine Super (CIO), St Davids Rd Advisory (Founder), and Colonial First State), and education (Macquarie University (lecturer), and Cuffelinks (co-Founder)). He was a Lead Author of the Member's Default Utility Function and made contributions to many industry reviews over the past decade.

Executive Summary

The review of SPS 515 represents an opportunity for APRA to frame how outcomes are assessed on a pre and post-basis in the superannuation industry.

To make the most of the opportunity I suggest that APRA could be a little more prescriptive, primarily at a high level, detailing some work items that they expect will be undertaken.

I spend a large part of this Discussion Paper outlining how lifecycle outcomes can be modelled and assessed. This can then be applied to address a range of important issues in the industry (e.g. default solution design, and the *"net return fallacy"*).

Additionally, I make some recommendations to improve the insight and assessment of investment objectives, and have some suggestions regarding business plans.

My review of the public submissions on SPS 515 confirms the shortcomings of the industry in developing universal benchmarking and assessment techniques. Working out ways to develop these solutions is critical, otherwise the industry will stagnate.

Finally, some of the points made in public submissions regarding concern around consistency of SPS 515 with present and future regulation / legislation (for instance AIST, ISA, Mercer, and ASFA) appear well made; I do not address these issues.

Approach: Prescriptive versus Self-Determination

Much of the feedback from public submissions focussed on the challenges associated with self-determination with respect to different areas of outcomes assessment. If viewed on the scale of (zero) complete prescriptiveness (prescribing a set of exact tests, templates etc.) through to (ten) fully self-determination approach, I consider APRA's reviewed SPS 515 to be close to ten.

There are benefits of prescriptive and self-determination approaches.

A prescriptive approach:

- Generates information which can be directly compared across funds
- Represents framing of how regulators and policymakers are viewing superannuation
- Creates a baseline standard for the industry to adhere to in its outcomes assessment

A self-determination approach:

- Allows for discovery of the degree of consideration being undertaken by each super fund
- Aids the discovery of new thinking and best practice across a range of complex topic areas
- Provides APRA with new learnings around assessment in some of these complex topic areas

I acknowledge that APRA may not have the authority to be overly prescriptive. However, there may be an opportunity to be a little more prescriptive, in the form of required pieces of work, rather than specific templates etc.

I will return to this issue towards the end of this document.

Outcomes Modelling and an Outcomes-Based Objective

I recommend that a baseline level of outcome modelling be required by all funds. This would create greater accountability of funds regarding the range of outcomes that their default solutions may produce. It has been my long-held view that an outcomes models (or an outcomes engine) should be at the heart of every super fund (see <https://cuffelinks.com.au/outcome-engines-heart-business/>).

Super funds determine an overarching default strategy for members which consists of a fund strategy (lifecycle or default), and an investment strategy, which consists of return, risk and fee targets. Ultimately a fund spends member's capital to implement an investment strategy and takes investment risk on their behalf. The risk / return (hence fee) target may vary through time. Theoretically this could occur for two reasons: (1) accounting for characteristics of the individual such as age (a 'lifecycle' strategy) may generate varying asset allocations; and (2) changing expectations around return for a given level of risk.

What appears missing, across the industry, and in documents such as SPS 515, SPS 530 and SPG 530, is an overarching strategic objective for members. Effectively it represents a quantifiable interpretation of the Objective of Superannuation (acknowledging that this is a Bill which has not yet been passed). This objective would relate to the outcome for the member; the outcome itself needs to be defined, the two most realistic possibilities being:

- Account value at retirement – related to retirement outcomes¹;
- Income in retirement – more consistent with stated objective of superannuation, but much more difficult to model, and generates cohort issues (balance cohorts and the interaction with Age Pension), which isn't bad, just complex!

¹ For instance, if retirement solutions are likely to be life annuities then account value at retirement is not the optimal objective; the appropriate objective is the amount of lifetime income that can be purchased at retirement. This introduces what is known as annuity purchase price risk, and to manage this risk it may be appropriate to transition a greater share of assets into long-dated fixed income securities as retirement approaches.

Once the outcome is determined the trade-off between achieving a higher expected outcome (linked to returns) and the range of possible outcomes (linked to risk) needs to be determined. Establishing such a trade-off is difficult, but best practice is to have clear objectives. Fortunately, academic researchers have been addressing problems of this nature for over 50 years. They use techniques such as utility functions, which formalise the trade-off between expected outcomes and the range of possible outcomes. Typically, such functions are designed to penalise poor outcomes more heavily than the rewards attributed to outsized positive outcomes (i.e. they acknowledge that people are risk averse), which feels an appropriate assumption for superannuation.

Consider the range of outcomes presented in Figure 1 below: has the Trustee of a fund confirmed that they are comfortable with the range of outcomes presented? How could they compare this range of outcomes against another one with a different shape (perhaps a lower expected outcome but a lower standard deviation and a higher minimum outcome)? A utility function can objectively calculate a score based on the range of outcomes, thereby enabling the comparison of alternatives which produce different outcome profiles. In this respect it is surprising that the submission made by Optimum Pensions suggests everything but a utility function as a candidate retirement metric! Balanced approaches of multiple measurements always face the difficulty of how to weigh the various measurements. By comparison a utility function creates discipline by making one establish those trade-offs (between different objectives) up-front (via the choice of parameter values); from there a utility function provides complete objectivity.

The Member’s Default Utility Function represents a somewhat complex example (because it focuses on the more complicated problem of retirement income) – see [http://www.aist.asn.au/policy/member%E2%80%99s-default-utility-function-\(mduf\).aspx](http://www.aist.asn.au/policy/member%E2%80%99s-default-utility-function-(mduf).aspx). Associate Professor Geoff Warren from ANU has a paper soon to be published in the Financial Analysts Journal which provides an excellent introduction into utility functions and how to use them (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3207865).

Without an objective based on member outcomes that can be quantified, we are left with a difficult situation:

	Super Funds	Policymakers, Review Groups, Regulators
Default Option Design	No objective basis for determining appropriate default design. Left with heuristic approaches and marketing input.	Restricted ability to compare default strategies. E.g. Productivity Commission found it difficult to compare balanced and lifecycle approaches.
Level of Risk	No objective approach for determining the appropriate level of risk which drives the range of outcomes and influences the expected outcome.	Difficult to compare MySuper defaults which target different levels of risk.
Cohorts	The consideration of cohorts is limited in the absence of a quantifiable objective.	This is a present issue as funds are compared more directly to each other yet may be acting appropriately for their cohort of members.

The Future: Personalisation based on age, balance, contribution rate, and gender i.e. the concept of a “cohort of one” ² .	No objective basis for determining benefits, leaving funds hesitant to spend on innovation if they are unsure of the benefits.	Difficult to place a value on future potential innovations, which may result in policies which discourage innovation.
---	--	---

Table 1: Challenges in the absence of a clear, quantifiable, outcomes objective.

From here I provide a short note on lifecycle outcomes modelling and then provide some case studies illustrating how such a framework can be used to address important problems that industry, policymakers, review groups and regulators are struggling to address.

Introduction to a Basic Lifecycle Outcomes Model

Much of this Discussion Paper is supported by a basic lifecycle outcomes model. This model was developed in Excel and is shared with APRA.

One of the important unanswered questions is whether funds should be assessed on a “to or through retirement” basis. Modelling outcomes on a “through” basis is much more complex. At a minimum it requires mortality modelling, Age Pension interaction modelling, and potentially consideration of more complex retirement solutions. The MDUF Working Group provides a free open-source lifecycle model with many of these features

([http://www.aist.asn.au/media/1102613/static_model_calculator - mduf v1 - 20171212.xlsx](http://www.aist.asn.au/media/1102613/static_model_calculator_-_mduf_v1_-_20171212.xlsx)).

For the purposes of this Discussion Paper, modelling on a “to” basis is sufficient to provide case studies on key issues.

Our basic lifecycle model is detailed in Appendix 1. An example of the output used to address some case studies is included in Figure 1.




Figure 1: Summary analysis from basic lifecycle model (accumulation balance at retirement).

The table contains some self-explanatory summary statistics, while the chart provides a distribution chart of outcomes.

Appendix 2 contains an initial basic attempt of a standardised industry objective.

² A “cohort of one” concept ascribes a unique asset allocation to each member of a MySuper default based on MySuper permitted characteristics such as age, balance, contribution rate, and gender.

Case Study 1: It is not just about net returns...

One of the themes in the industry, and indeed the submissions, is that it is all about net returns. The key word missing here is risk. Funds spend member’s savings to implement an investment strategy and they take investment risk on their behalf. There exists a reasonable dispersion amongst the risk levels targeted by different MySuper offerings.

Consider the following simple case, viewed through the eyes of a risk-agnostic party whose focus is on net returns only. There is a general acceptance that higher risk is rewarded with higher returns over the long-term, so a risk-agnostic investor would rationally target higher risk. If leveraged exposure (through non-liability vehicles such as geared share trusts) can be taken then this is also rationally consistent with a risk-agnostic, net return focused investor.

Clearly this is not the case, but it does highlight the need to account for risk, especially given the varying levels of risk targeted amongst MySuper options.

This case study illustrates that it is not just about net returns. I manufacture a specific case:

Fund 1: 10% volatility, targeting net returns of CPI + 4.5%

Fund 2: 7% volatility, targeting net returns of CPI + 4%

The summary results for each are presented below:




Figure 2: Summary analysis from basic lifecycle model (accumulation balance at retirement), Fund 1 assuming 10% volatility, targeting net returns of CPI + 4.5%.

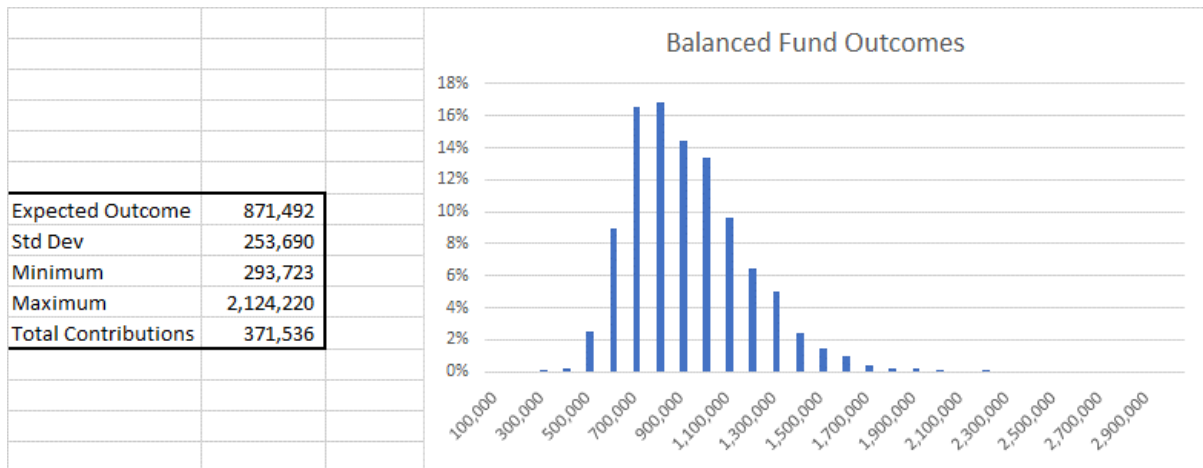


Figure 3: Summary analysis from basic lifecycle model (accumulation balance at retirement), Fund 2 assuming 7% volatility, targeting net returns of CPI + 4%.

Comparing Figure 2 against Figure 3 reveals some interesting observations:

- Fund 1 has a higher expected balance at retirement (\$972,302) compared with Fund 2 (\$871,492).
- This appears driven by a long tail of very high outcomes. The maximum simulated outcome for Fund 1 was \$4,804,655, compared with \$2,124,220 for Fund 2.
- However, Fund 2 delivers a narrower range of outcomes, as evident from the diagrams, and statistically by the lower standard deviation (\$253,690 for Fund 2 versus \$436,615 Fund 1).
- The worst-case outcome for Fund 2 (\$293,723) is also higher than for Fund 1 (\$248,377).

How does a Trustee choose between such strategies, as they are obliged to? How do policymakers, reviewers and regulators compare two strategies (which could readily represent MySuper defaults)?

A basic utility function considers all possible outcomes and considers. In this case Fund 2 is considered to provide a better balance of outcomes for members (the worse cases in Fund 1 are penalised more heavily than the very high outcomes).

This example demonstrates that it is not all about net returns. Where trustees have the responsibility to determine an appropriate risk target, it is all about acknowledging the range of possible returns associated with a strategy. Only once a risk target is determined (whether legislated (e.g. all MySuper defaults must target 7% standard deviation) or determined by trustees), can the focus switch solely to net returns.

From this example observe that:

- This represents a case study where, arguably, Fund 2, delivering lower net returns, is a better strategy
- Without accounting for risk, there exists the potential for strategies targeting high risk, underperforming in terms of converting risk into returns, may still delivering high net returns, to be viewed favourably by a net-return focused industry.

Case Study 2: Reviewing appropriateness of default strategy

There is mixed evidence on the ability to forecast market returns³. Frontier Investment Consulting recently observed that 25% of super funds have reduced their investment objectives over the last three years. Has this resulted in a change in the default strategy design? In nearly all cases no, yet the question for trustees is an important one: *“Is it appropriate to take so much risk on behalf of my members if I now expect this to be less rewarded?”*.

Let’s explore this case a little further. I analyse the outcomes of a fund before and after reviewing its investment objectives.

Fund 1 (pre-review): 8% volatility, targeting net returns of CPI + 4%

Fund 1 (post-review): 8% volatility, targeting net returns of CPI + 3.5%

The summary results for each are presented below:




Figure 4: Summary analysis from basic lifecycle model (accumulation balance at retirement), Fund 1 (pre-review) assuming 8% volatility, targeting net returns of CPI + 4%.




Figure 5: Summary analysis from basic lifecycle model (accumulation balance at retirement), Fund 1 (post-review) assuming 8% volatility, targeting net returns of CPI + 3.5%.

Comparing the outcomes in Figure 4 against those in Figure 5, it becomes obvious that a 0.5% reduction in investment objectives has a large impact on the range of possible outcomes. It would be

³ I personally agree with the research of Nobel Laurette Robert Shiller whose research suggests that valuation provides an insight into longer term return expectations.

appropriate for a trustee to review their default strategy: it would be appropriate to at least consider a lower return / lower risk default design in an environment where they don't expect risk to be as well rewarded. Consideration of the optimal strategy is a more complex case study not explored in this Discussion Paper.

Case Study 3: Comparing lifecycle against balanced strategies

There remains an ongoing debate in the industry around whether lifecycle strategies or balanced strategies represent more appropriate default strategies. This uncertainty has carried across to reviews such as the recent work undertaken by the Productivity Commission. Unfortunately, subjective arguments and anecdotes appear to dominate objective analysis based on the range of outcomes members may experience.

A very brief introduction on why lifecycle strategies have merit:

- Basic finance theory teaches us that we manage risk best when we diversify our risk. A related concept is time diversification: if we are exposed to the same amount of investment risk each period of time, we maximise time diversification.
- If we have a single amount to invest for the long-term then a balanced strategy broadly delivers time diversification.
- However, if we acknowledge a lifetime of contributions to super, then our dollar exposure to investment risk through time is no longer constant – we are investing a much larger amount of contributions closer to retirement.
- A lifecycle strategy partly addresses this by taking more investment risk when contributed balances are small and less risk when contributed balances are large, somewhat evening out the dollar investment risk exposure through time.

Reasons to be cautious:

- Future contributions are not certain (though I would suggest this highlights the motivation for personalised defaults – enabling a fund to account for whether a member is an active contributor).
- The analysis becomes more complex when we consider “*through retirement*” strategies and interaction with the Aged Pension.

To aid industry / policymaker resolution around this issue, I believe a lifecycle outcomes model complemented with an outcomes-based objective helps provide clarity. I now provide an illustration comparing a lifecycle and a balanced default strategy.

Lifecycle Strategy:

- First 20 years: 10% volatility, targeting net returns of CPI + 5%
- Next 10 years: 8% volatility, targeting net returns of CPI + 4%
- Final 10 years: 6% volatility, targeting net returns of CPI + 3%

Balanced Strategy: 8% volatility, targeting net returns of CPI + 4%

The summary results for each are presented below:




Figure 6: Summary analysis from basic lifecycle model (accumulation balance at retirement), Lifecycle Strategy.




Figure 7: Summary analysis from basic lifecycle model (accumulation balance at retirement), Balanced Strategy⁴.

The analysis seems to suggest these two strategies are expected to produce a similar range of outcomes for their respective members. The Lifecycle strategy delivers a narrower range of outcomes and the worst-case outcomes are improved, but the Balanced strategy has a higher average outcome. Weighting the range of outcomes using our utility function, the utility scores are nearly identical! A fence-sitters case study(!) but one which illustrates that the debate around which strategy is more appropriate for default design is complex but can be assessed objectively.

Drilling into Investment Objectives

There is the opportunity to create greater accountability in the setting of investment objectives and assessment of investment performance.

⁴ Note that these results should be identical to those in Figure 4, but they are not. This highlights the variation in results produced by simulations. In practice a greater number of simulations would produce more similar results.

Investment objectives are impacted by the following: level of risk, strategic asset allocation (SAA), implementation, and tactical asset allocation. Forward-looking and post-review accountabilities can be created at each of these points. A breakdown (with example numbers) is provided below:

Component	Expected Return	Justification	Review
Level of Risk - Taking investment risk is expected to deliver returns	CPI + 3% pa (based off 8% risk target)	Reference portfolio. Risk target determined by Trustee based on outcomes modelling.	Was our forecasting approach appropriate?
Strategic Asset Allocation - Performance of strategic mix of assets against the basic mix of assets in a reference portfolio.	0.7% pa	Internal modelling based on expected returns from strategic asset allocation compared to expected returns from reference portfolio. Total tracking error (to reference portfolio): 1% pa.	Performance assessment of SAA versus reference portfolio performance. What worked, and what did not? Actions?
Implementation - Performance of implemented portfolio against SAA.	0.2% pa	E.g. 1: Implementation through active fund managers – we expect to achieve extra returns net of fees. Total tracking error (to SAA): 0.5% pa. E.g. 2: Implementation through internal implementation - we expect to achieve extra returns net of fees. Total tracking error (to SAA): 0.5% pa.	Performance assessment of portfolio versus SAA performance. What worked, and what did not? Actions?
Tactical Asset Allocation - Performance of active asset allocation tilts relative to SAA.	0.1% pa	E.g. 1: Implemented internally – we expect to achieve extra net returns. Total tracking error (to SAA): 0.3% pa. E.g. 2: Outsourced via a TAA overlay - we expect to achieve extra returns net of fees. Total tracking error (to SAA): 0.3% pa.	Performance assessment of portfolio versus SAA performance. What worked, and what did not? Actions?
	Total: CPI + 4%		

Table 2: Framework for generating investment objectives and improving accountability and review.

The table above refers to the concept of reference portfolios. These are reasonably commonplace now; they represent a simple, implementable portfolio of mainstream asset classes (typically, listed stock indices and bond indices). They have strong application for performance review – an excellent case study is NZ Super (<https://nzsuperfund.nz/performance-investment/returns-vs-ref-portfolio>). It would be ideal for industry level reference portfolios to exist, as identified by the Productivity Commission.

The framework outlined in Table 2 could readily be incorporated into SPS 515 or SPG 530 – there is some potential crossover between the two. The forced accountability on super funds (in my view

many funds are weak on performance assessment, accountability, and plans to address), along with insights gleaned for APRA represent significant benefits.

Business Plans

I offer some suggestions for further enhancing the Business Plan section:

1. Funds should detail the strategic initiatives they are not pursuing (i.e. considered but excluded), with reasoning provided, and detail their self-assessment of how this leaves them with a competitive offering. Examples are highly varied and could include internalisation of investment capabilities, ESG, retirement solutions, and financial advice.
2. Funds should identify the areas where they consider they take significant operational peer group risk. Examples include internal administration, fee levels, product offering (e.g. QSuper, Christian Super), investment model (e.g. retail funds are liquid asset only), financial advice services, and other relevant areas (e.g. major system development).
3. It should be made clear that APRA will cross-check the Business Plan against likely regulatory changes, to see that they have been accounted and planned for (at least the contingency) in the present Business Plan.
4. The Business Plan should include a financial stress-test scenario which considers a fall in asset markets, and how the fund would respond.

Clarity on Cohorts

Some submissions (e.g. Super Consumer Australia, ASFA) raise the issue of how to define cohorts. Cohorting is an interesting topic, as it goes much against some of the themes represented by the Productivity Commission (basic defaults, non-lifecycle, enabling more direct performance comparison).

Cohorts are best considered objectively through the lens of a lifecycle outcomes framework. In this sense having a primary cohort (or test case) is most important. Subsequent cohorts could differ in many different dimensions (and combinations of), including income, retirement age, contribution breaks, additional contributions, and if we consider retirement outcomes then household structure, home ownership, and other assets.

The governance required around cohort design would make the development of a collection of cohorts a project with no near-term solution – discussed further below. However, this is an important area to consider as the value of present (multi-factor lifecycle solutions) and potential future innovations (such as “*cohort of one*” solutions) can only be ascertained by reflecting a fund membership as a collection of different cohorts, through a lifecycle outcomes model.

Weighting of Components of Outcomes Assessment

Some submissions (notably AIST and ISA) note that a self-determined weighting process creates ambiguity, is open to manipulation, and is potentially inconsistent with (what they identify as) the primary focus of net returns to members (noting previous section on the need to consider default design and level of risk).

The issues of default design, risk, and net returns could be reconciled via a lifecycle outcomes model and accompanying utility function. This would create an objective outcomes assessment test based on measurable components.

How to best incorporate the more subjective components? One solution is a balanced scorecard approach, with self-determined weightings. The approach suggested is to make the financial outcomes assessment front-and-centre and then to have a table which details other benefits. The onus could be placed on the super fund to provide an objective-based argument for the benefit of other services and features to retirement outcomes.

Industry Benchmarking

Some submissions note the need for the provision of standardised information to facilitate more direct comparison. The industry collectively has failed to provide solutions, and self-determination is not working well in some areas, impacting comparability and consumer outcomes. It seems unclear whether industry or regulatory groups should produce this information⁵. Related to outcomes assessment, I believe the following areas are most in need of standardisation:

- Growth / defensive asset categorisation – currently funds are not being compared on a like-for-like basis
- Reference portfolios to cover the entire risk spectrum (low to high risk)
- Reference portfolio performance – to enhance accountability
- An open-architecture basic lifecycle outcomes assessment tool
- A standardised utility function to facilitate objective performance assessment
- A standard cohort for assessing lifecycle outcomes
- Alternative cohorts, to enable consideration of the unique member mix of individual super funds

Suggestions and Conclusion

The review of SPS 515 represents an opportunity for APRA to frame how outcomes are assessed on a pre and post-basis in the superannuation industry.

To make the most of the opportunity I suggest that APRA be a little more prescriptive, but only at a high level, detailing some work items that they expect will be undertaken. To this end, I would recommend:

1. Each fund be required to project the range of lifecycle outcomes that they expect their default strategy (style of default, risk and return targets) to produce.

⁵ It should be noted that nearly all major achievements in the superannuation industry have been driven by policy and regulation: the track record of participant-based solutions is modest.

- a. Given this is new to much of the industry it could be based on accumulation balances at retirement and could be as simple as the model detailed in Appendix 1. It could also be flagged that more complex modelling may be required at a future point in time (for instance retirement outcomes).
2. Make an assertion that the Trustee has considered the range of outcomes produced by their default strategy and that they consider it provides an appropriate range of outcomes relative to other strategies.
3. Require funds to provide a detailed breakdown of how they formed their investment objectives, based on the framework outlined in Table 2.
4. Require funds to attribute and self-assess their investment performance against the framework outlined in Table 2.
5. Additional business planning enhancements:
 - a. Funds should detail the strategic initiatives they are not pursuing
 - b. Funds should identify the areas where they consider they take significant operational peer group risk
 - c. It should be made clear that APRA will cross-check the Business Plan against likely regulatory changes, to see that they have been accounted and planned for (at least the contingency) in the present Business Plan
 - d. The Business Plan should include a financial stress-test scenario which considers a fall in asset markets.
6. Consider whether it is possible to design and disclose an initial cohort, clarifying that this will be reviewed and enhanced through time.
7. Make the financial outcomes assessment test front-and-centre and then have a table which details other benefits. The onus could be on super funds to provide an objective-based argument for the benefit of other services and features to outcomes.
8. Consider ways in which development of a range of industry benchmarking activities can be accelerated.

I do not consider that any of the above recommendations would be viewed as too prescriptive. In each case a fund would have the right to complement the analysis with analysis based on their own assumptions and insights. Each recommendation would raise industry standards in terms of the work they undertake.

Appendix 1

Basic Lifecycle Model

The model used is quite basic and has the following characteristics:

- An individual is assumed to earn an amount similar to AWOTE for their entire career
- Earnings growth slightly exceeds inflation
- 9.5% Superannuation Guarantee
- Taxes ignored (but should be included)
- Investment returns defined on a case-by-case basis
- Investment returns are assumed to be normally distributed
- Model can account for lifecycle strategies
- 1,000 simulations to generate the range of possible outcomes

Appendix 2

Straw Man 'Standard' Objective Function

A basic utility function which stays quite close to the academic literature. It has the following characteristics:

- Focuses on accumulation balance at retirement (not income in retirement)
- Assumes constant relative risk aversion (known as CRRA), specifically people take risk in proportion to their assets, and not a dollar amount of risk. This is an assumption the whole default industry is based on
- The formula appears as follows:

$$U(W_T) = \frac{W_T^{1-A_R}}{1-A_R}$$

Where:

- W_T is accumulation balance (wealth) at the time of retirement (T)
- A_R is the coefficient of relative risk aversion
- I assume A_R is 5, consistent with values seen in the academic literature