Target Benefit Retirement Schemes

Securing sufficient funds to ensure a comfortable old age income is becoming an increasingly troublesome undertaking globally. Defined benefit schemes are frequently underfunded and under more and more pressure to reduce pension outcomes whilst defined contribution funds deliver below par results. Although this is mainly due to demographic factors, increasing longevity and disappointing returns it is clear that in addition especially for defined contribution systems a crucial ingredient is lacking, which is the notion of an explicit target outcome. The lack of such a target implies the absence of a relevant measure of risk as well. Consequently asset allocations in DC systems are usually rudderless in their essential objectives, or are at best driven by the reliance on time diversification, which implies a decreasingly volatile asset allocation over time. In this paper we make use of techniques common in the defined benefit world to bridge the gap between DC and DB through a "target benefit" approach that leads to much more stable and appropriate pension outcomes.

Multi-Asset Solutions Research Papers - Issue 4

For institutional and adviser use only

First Sentier Investors

Introductions

This paper presents a flexible and innovative target benefit approach for a collective defined contribution pension system or fund. Defined contribution schemes are becoming increasingly prevalent in retirement systems globally. However, questions remain about the effectiveness of many of the designs in achieving the ultimate goal: assisting participants to save for retirement. Flexibility, scalability and innovation are essential to build a program that can address the defined contribution requirements to deliver appropriate retirement income streams. An innovative concept is required to address the investment problem before retirement (the accumulation phase) and the provision of income after retirement (the decumulation phase) in a consistent manner. Our approach provides a flexible solution that borrows heavily from concepts originally developed in defined benefit schemes, and which thus far have found little application elsewhere. By explicitly targeting a level of income after retirement, we can design an investment strategy which, over time, minimizes the probability of not providing this income stream or its annuity equivalent. This is where we bridge the gap between defined benefit and defined contribution schemes – hence the moniker "targeted benefit." Essential ingredients are:

Liability driven	asset allocation based on a notional pension liability at retirement age
Focused on participants	actuarial and human capital factors taken into account
Global	use of a wide and diverse range of asset
Dynamic	regular rebalancing to adjust risk/return profile appropriately
Long-term	continued integrated management after retirement

"Inflation Plus" investment targets suit many types of investor, and are usually derived from an underlying liability structure or spending pattern.

The importance of the right Risk/Return trade-off

Choosing the appropriate risk/return trade-off is vitally important in its implications for the design of a target benefit scheme or in fact, any investment problem. Crucial in this trade-off is the necessity to define and measure the outcomes of the scheme in relevant quantities. Portfolio return and volatility are important, but not necessarily relevant to participants. Instead participants want to know how much income they can reasonably generate from their accumulated capital. In order to do this the value of the accumulated capital is expressed as the capital equivalent of an annuity as percentage of (latest) income. The concept of risk then becomes clear: it is not volatility around an expected return that matters, but the probability of not accumulating sufficient capital to meet the funding requirement for such an annuity target. This translates into a notional liability at retirement age. That is, we view the capital required to purchase an adequate income stream as a future liability that must be funded. The ability and likelihood of funding this liability depends to a great extent on the investment strategy during the accumulation phase. In addition, after retirement the risk of the capital running out during the lifetime of the retired participant is obviously relevant. This is where insurance aspects may come into play in using part of the accumulated capital to acquire actual (lifelong) insured annuities.

The ability to express the risk and return in terms of annuity (equivalents) at retirement age

also sets the targeted benefit apart from traditional lifecycle and target date schemes. These typically only target accumulated wealth in isolation without regard to the true purpose of the wealth accumulation, which is retirement income. This leads to well-attested suboptimality for traditional lifecycle[1] funds as documented by Booth and Yakoubov 2000 using both historical as well as simulated data, and also more exhaustively in terms of available investment strategies by Blake, Cairns and Dowd 2001. The latter looked at the riskiness of various defined contribution schemes relative to a defined benefit retirement target, finding that the asset allocation of schemes is the most important factor determining the riskiness and that portfolios with high static equity weights would have done better than dynamic asset allocation strategies over time. This point is also made by Shiller 2006 in whose results a 100% equity portfolio produced higher wealth than the lifecycle proxies. This is in part due to the so-called "size effect" as explored by Basu and Drew 2009 which reflects the fact that a return made later in life affects a larger accumulated corpus due to earlier capital gains and contributions. As a result portfolios with high constant equity weights can be expected to accumulate greater wealth but without a proper notion of risk. A rebuttal of the size effect can be found in Pfau 2011, but as even Shiller 2006, Basu and Drew 2009, and Basu, Byrne and Drew 2011 argue this context-free unanchored approach to wealth accumulation points to the underlying problem of the surveyed lifecycle funds.

Traditional lifecycle and target date funds operate without an explicit and relevant retirement target, making them suboptimal for most investors.

However, the approach presented here addresses exactly this shortcoming of traditional lifecycle and target date funds. The moment a target is established the move to a less aggressive portfolio with increasing age and capital becomes the superior strategy. Here we optimize for annuity-based retirement targets throughout the lifetime of the participants allowing them (with the scheme's guidance) to achieve their objectives with much less overall relevant risk, where the relevant risk is the risk of not having enough capital to buy their desired annuity. The customization inherent to our approach affords the participants a risk-minimizing path to accumulating the required capital that is specific to their circumstances in terms of their human capital, built-up financial capital and time remaining until retirement. A very different approach is explored in Kyrychenko 2008, where city-specific human capital factors, housing and business assets are treated as asset categories in an unconstrained mean-variance approach. This leads to heavily leveraged portfolios in some cases and, while academically interesting, the results are not likely to find implementation in practice. Moreover the optimization does not include the contribution cash flows, which do play a role in the risk/return trade-off that needs to be made. However the paper does show that even participants living in different cities and working in different industries will have different retirement objectives and will not be best served by a traditional lifecycle fund. The implied, intuitively obvious, result that "one size does not fit all" is made explicitly by Bodie and Treussard 2007 in a broader context.

[1] The terms "lifecycle" and "lifestyle" are used interchangeably here and refer to funds that reduce the equity weight in the portfolio as the participants near retirement. In that sense our target benefit funds are a special case of the lifecycle funds, the key differentiator being the explicit target and the resultant target-relative risk criteria. Most existing lifecycle funds wuse a much less sophisticated concept of risk and return, sometimes even just reducing equity weight by one percentage point per year.

The Participant Grid concept offers enormous flexibility in design, both before and after retirement, while never losing sight of the target benefit

Designing the Scheme

Focusing specifically on the retirement needs of the participants allows us to create a system in which participants are actively guided to their retirement goal. To determine optimal asset allocations for a target benefit scheme, Asset Liability Management tools are required. Contributions, actuarial discount factors and career patterns as well as the economic environment also play an important role in the design of the scheme. Purely analytical optimization approaches are unequal to the task.

Analysis of annuity risk is complex and requires simulations of all ingredients. By generating numerous scenarios through random drawings of macro-economic variables and returns[2], an annuity risk analysis is possible by interpreting fractions of outcomes as probabilities. By assessing these probabilities we can optimize the allocations for individual participants. These optimal allocations are not static, but change over time. They form a guided path that takes actually achieved investment returns, the level of accumulated capital and the remaining investment horizon into account.

Basics of the Design

The basis for the design is the so-called Participant Grid (see Figure 1) in which generic participants are defined in terms of their age and their number of participation years in the retirement savings system.

Figure 1: Participant Grid

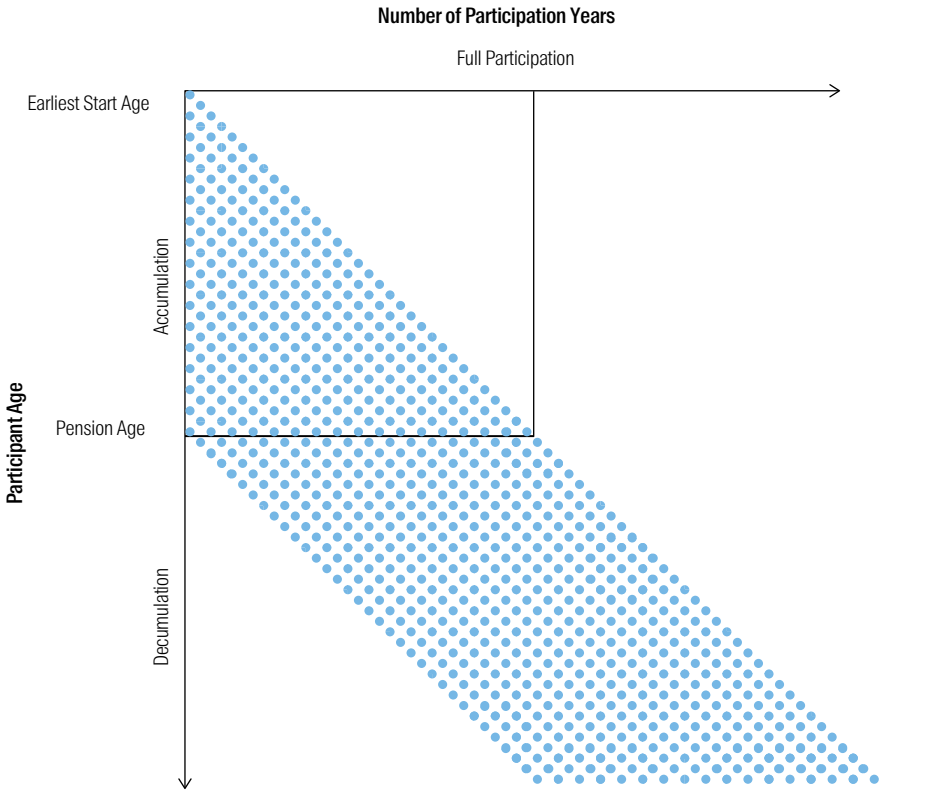

Source: First Sentier Investors

[2] We use a VAR(1) model.

The Participant Grid captures any participant in two dimensions and serves as the basis for the design, adding additional dimensions for each point in the grid.

The Participation Grid reflects all participants of a scheme in terms of age and years of participation. Capital shortfalls can be translated to an equivalent in terms of number of missed participation years. In essence the horizontal axis is a function of accumulated capital, but conceptually it is beneficial to think in terms of participation years to increase the general applicability of the Participant Grid. The upper diagonal of this Participation Grid represents the participants that are "at target[3]," as defined in the construction of the scheme.

Along the vertical axis (i.e. with zero participation years) the Participation Grid represents new entrants into the system without any previous capital brought in. The design consists of defining tailored annuity or capital objectives and risk tolerances for every point in this entire grid. An optimization process then leads to customized asset allocations for all participants. That is, each point in the Participant Grid has an optimal asset allocation. For an entrant at the top left of the grid the scheme minimizes the probability that this participant will fall short of the capital required as represented by the bottom right corner.

Annuity Targets versus Expectation Values

There is an important distinction to be made between the targeted annuity and the expected annuity. We illustrate this with the simple case where an investor requires a 5% return. Investing in a portfolio with an expected return of 5% the probability of meeting or exceeding this target would be roughly 50%[4], which is obviously not a very satisfactory outcome. Hence the investor requires a portfolio with a higher expected return in order to increase the probability of making 5%. The same concept applies here as well albeit in annuity space. In order to minimize the probability of not achieving the scheme target annuity percentage, the expected annuity will have to be substantially higher.

As this pertains specifically to the accumulation phase we look at only the pre-retirement part of the grid to illustrate the consequences of this point. For a participant starting at the earliest possible entry age (at the top left of the grid), the annuity target is represented by the point at the bottom right of the grid. In order to minimize the probability of not achieving this, the scheme needs an expectation value for the annuity that is much higher, i.e. much further to the right. This is what we show in Figure 2.

[3] We will use the "at target" shorthand throughout this paper to refer to participants who have just enough capital at any given point in time to buy an annuity equal to the scheme target at that point. Participants can also fall behind, which leads to their moving to the interior of the grid as their equivalent number of participation years will be lower.

[4] Exactly 50% assuming a normal distribution, but typically close to 50% with lognormal distributions due to skewness.

Figure 2: Sample Participant Accumulation Path

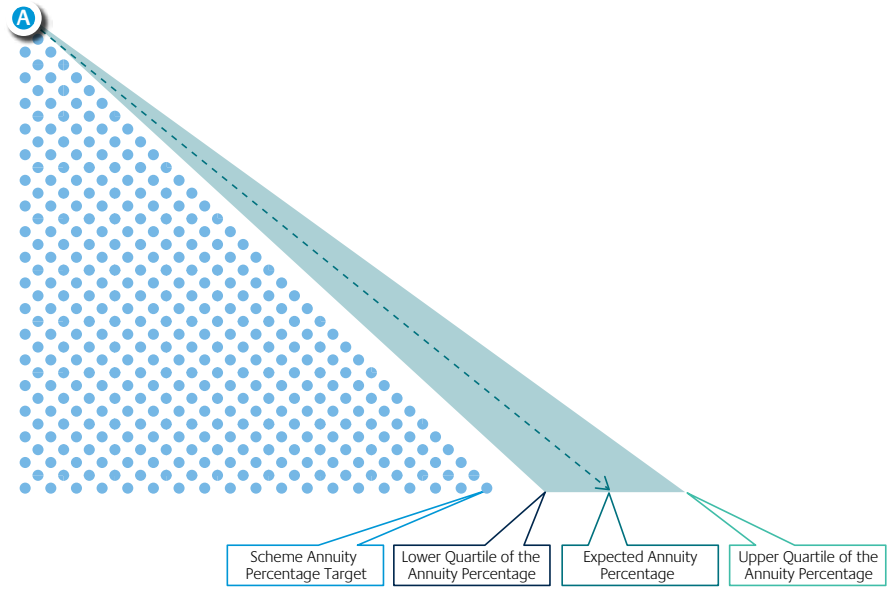

Source: First Sentier Investors

Participant A's capital accumulation is actually outside and to the right of the grid. The probability distribution of the annuity outcomes is shown as the blue shaded area, which here represents the interquartile range of outcomes[5].

Taking two participants, B and C, we further illuminate the concept in Figure 3. They both enter the system at a later age, but B brings in no capital at all, while C is "at target," having just enough capital to buy an annuity equal to the scheme target for his age. For both B and C we can plot an expected annuity path as well as the interquartile range.

Figure 3: Accumulation Paths for Model Participants

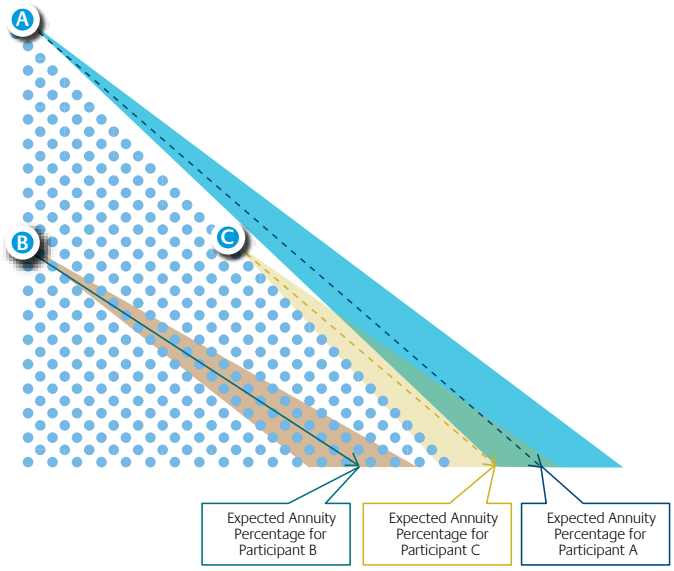

Source: First Sentier Investors

[5] This conceptual diagram is not to scale. We show actual outcomes later in this paper.

Participant B entered too late with too little capital to make the overall scheme target at retirement age[6], and hence the entire interquartile range of outcomes also falls within the grid. For Participant C the situation is different in that the expected annuity outcome is higher than the scheme target at retirement age, but there is also a significant probability of falling short of the target.

It is also instructive to note that Participant C differs from A fundamentally in that C is not just an older version of A, which is due to the difference in built-up capital. An older version of A would fall on the blue dashed line outside of the grid. There may also be differences in terms of starting salary and other parameters.

Assumptions and Design Parameters

In the remainder of this paper we will explain the workings of the scheme by going through an Australia-focused sample design. In this sample design we assume that the earliest age at which participants[7] can enter the scheme is 20 and the assumed retirement age is 65. This means that the full accumulation period is 45 years. In the accumulation phase the contributions are 12% of salary[8] and the annual withdrawals in the decumulation phase are assumed to be 3% of available capital.

In the sample scheme we target a capital at retirement equivalent to an annuity of 30% of salary for participants that are at target or less than 10 years behind[9]. If participants have a participation shortfall of 10 years or more the target annuity level at retirement decreases to 1% of salary at age 64, with a remaining investment horizon 1 year. The above design parameters are depicted in Figures 4 (a), 4 (b) and 4 (c).

Figure 4 (a): Start Age/ Retirement Age

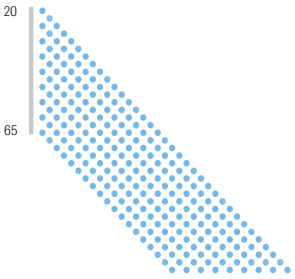

Figure 4 (b): Contributions/ Withdrawals

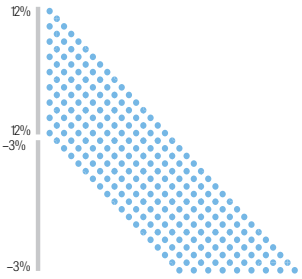

Figure 4 (c): Annuity Targets at Retirement

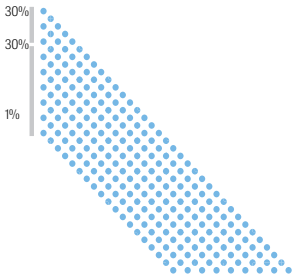

Source: First Sentier Investors

[6] With a reasonable probability. The real probability of making the scheme target is not actually zero, but will be very close to it.

[7] We assume all participants to be male in this example.

[8] Benefits are indexed and include old age pension and a 60% insured spouse pension in case of death of participant.

[9] All actuarial calculations are based on Australian Life Tables 2008-2010. Source: Australian Bureau of Statistics.

For each path parallel to the main diagonal[10] in the Participation Grid (and therefore for any participant) one can determine the annually required actuarial return to arrive at the target annuity level at retirement. For the at-target participant, to arrive at the 30% goal, the annually required actuarial return is 4.36%. The required return as a function of missed participation years increases to an annual rate of 6% at 44 years of missed participation. Using the required actuarial return one can calculate the annual target annuity levels over the course of the participation. Figure 5 (a) shows the target annuity levels at retirement depending on the missed participation years as described above. In addition, in Figure 5 (b) the target annuity path is shown for the at-target participant.

Figure 5 (a): Target Annuity Level

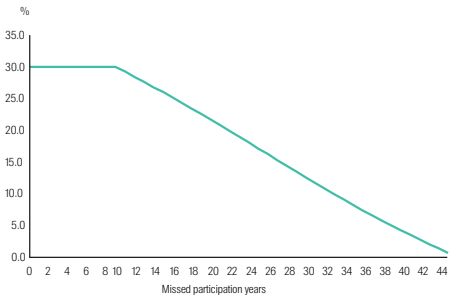

Figure 5 (b): Capital Annuity Equivalent

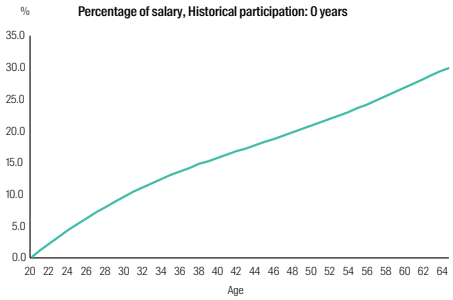

Source: First Sentier Investors

After retirement participants can use part of the accumulated capital to buy an indexed annuity. For the remainder of the capital (the part not used to buy the indexed annuity) we target an annual return of 5.5% in the decumulation phase, which is equal to the assumed withdrawal of 3% plus 2.5% expected Australian inflation. This way we strive for a perpetuity after retirement through aiming to protect the real capital amount.

With this we have defined the level of contribution (or withdrawal) for every point in the Participant Grid as well as the annuity target at retirement. To each of these points we then assign a confidence level with which these targets need to be attained given full participation in the scheme from that point onwards. The confidence level to achieve the 30% annuity target for a new entrant at age 20 is 68%. This required confidence level increases as this participant invests in the scheme over time, rising to 80% at retirement age, and then increasing even further to 90-95% in the decumulation phase, where the confidence level then is applied to the required withdrawals. The general scheme is shown in Figure 6.

The potential target annuity depends on many exogenous scheme characteristics, such as contribution levels and retirement age.

[10] This being the upper diagonal going from the top left of the grid down to the bottom right.

Figure 6: Required Confidence Level To Attain Target

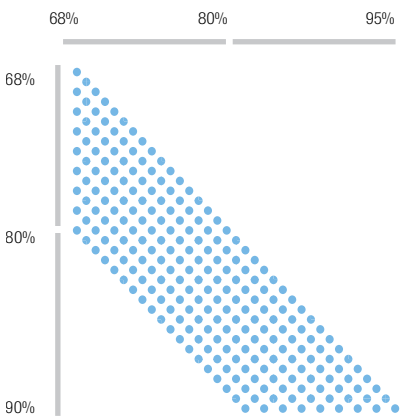

Source: First Sentier Investors

Having set all of the above scheme parameters we still need to model participants in terms of their salary development. Like many aspects of the design this is greatly dependent on the actual population at hand and should be looked at in each specific case.

In our sample design we assume age-dependent start salaries for all participants in the scheme, with, for example, a 20 year old participant assumed to have a start salary of 35,000 AUD and a 40 year old participant to have a start salary of 65,000 AUD. In Figure 7 (a) we show the assumed start salaries for each participant by age. On top of this, the career development of a participant (expressed in terms of annual salary percentage increases) greatly influences the ratio of annuity and salary over time. In Figure 7 (b) we show the assumed age dependent career development for each participant. As one can see the career related salary increases decrease with age, becoming zero at age 60[11].

Figure 7 (a): Start Salary

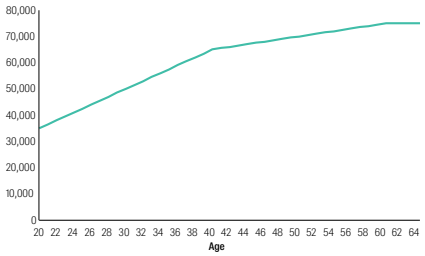

Figure 7 (b): Age Dependent Career Pattern

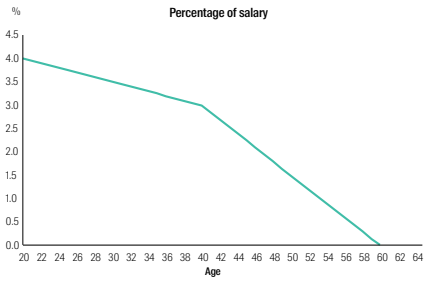

Source: First Sentier Investors

[11] These are obviously examples only and should be tailored to a specific population in an actual implementation of the scheme.

From Asset Universe to Glide Surface

The next step in designing a target benefit scheme is to define a universe of relevant and acceptable asset categories. In our sample case we choose Australian nominal and index-linked fixed income, Australian equities, international equities, and emerging markets equities, with all international exposure unhedged. Table 1 shows the risk and return assumptions of the asset categories included in the analysis. These assumptions are based on the output of our stochastic Long-Term Asset Return Model (LTARM)[12] applied within a macro economic climate with a long term inflation expectation for Australia of 2.5%.

Table 1: Expected Asset Class Characteristics

			Correlations					
Asset Classes	Expected Return (%)	Expected Volatility (%)	AUD Cash	Australian Bonds	Australian ILBs	Australian Equities	World Equities	Emerging Markets Equities
AUD Cash	3.5	0.3	1.00	0.12	0.07	–0.10	0.01	–0.14
Australian Bonds	4.0	4.0	0.12	1.00	0.63	0.06	0.00	–0.06
Australian ILBs	4.0	5.9	0.07	0.63	1.00	0.08	0.07	–0.01
Australian Equities	8.5	14.4	–0.10	0.06	0.08	1.00	0.52	0.63
World Equities	8.5	13.3	0.01	0.00	0.07	0.52	1.00	0.59
Emerging Markets Equities	10.5	21.2	–0.14	–0.06	–0.01	0.63	0.59	1.00

Source: First Sentier Investors

To determine the optimal asset allocation for each position in the Participant Grid we first have to restrict ourselves to a manageable range of candidate strategies. These candidate strategies are derived by conducting a Conditional Value-at-Risk portfolio optimization using the assumptions in Table 1[13]. Figure 8 shows the result of this optimization. The varying weights of the asset categories along this CVaR-efficient frontier are depicted as vertical slices in the chart, the strategies as a whole ranging from the most conservative strategy on the left to the most aggressive on the right. The horizontal axis shows the nominal expected return of each strategy.

Figure 8: Mean-Expected Conditional Value at Risk Efficient Frontier Weights

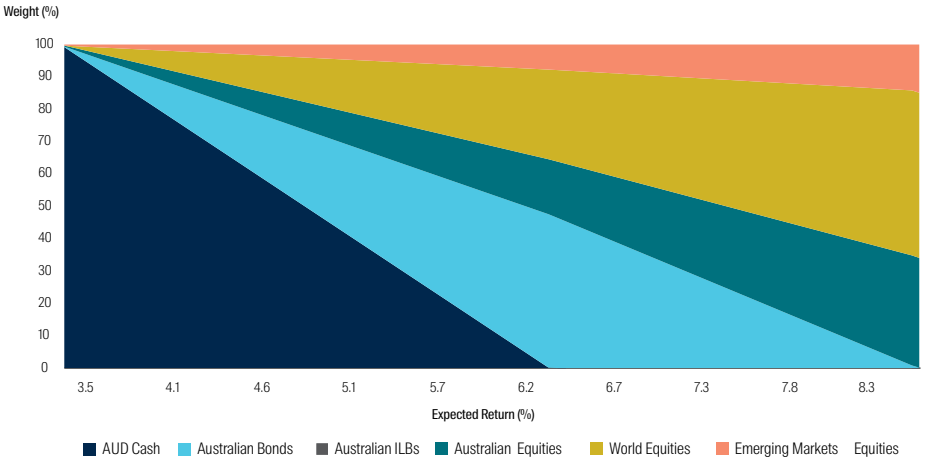

Source: First Sentier Investors

[12] Further information on our Long-Term Asset Return Model is available upon request.

[13] Other ways of shrinking the set of possible candidate strategies are possible too without loss of generality for the scheme design concept presented here. For instance, a risk-weighted or factor-weighted approach can equally shrink the set to a manageable number, which then can be used instead of the optimization we employed in this example.
Any forecasts in the table above represent hypothetical numbers and are purely for illustrative purposes only.
The numbers do not represent actual or future performance.

The next step is to determine which of the above candidate strategies is optimal given the tailored target annuity criteria and confidence levels for each point in the Participation Grid[14]. In the accumulation phase, the objective for the participants obviously is to maximize the accumulated capital. The optimal allocation can be found by restricting ourselves to only those candidate strategies that satisfy a pre-defined acceptable Value at Risk (VaR) relative to the target annuity level, or equivalently to the corresponding actuarially required return. Out of the range of acceptable strategies the one with highest expected return outcome is considered optimal. We show the actual allowed shortfall in terms of VaR and the corresponding confidence levels in Table 2.

Table 2: Sample Participant Grid Parameters

	Accumulation Phase				Decumulation Phase			
Participants (Missed Participation Years)	Required Return (%)	Allowed Shortfall (%)	Confidence level at start age (%)	Confidence level at retirement (%)	Required Return (%)	Allowed Shortfall (%)	Confidence level at retirement age (%)	Confidence level at scheme end (%)
At Target	4.4	2.5	68.0	80.0	5.5	3.5	80.0	95.0
10 MPY	5.4	4.0	70.7	80.0	5.5	4.0	80.0	93.8
15 MPY	5.5	4.3	72.0	80.0	5.5	4.2	80.0	93.2
20 MPY	5.6	4.6	73.3	80.0	5.5	4.4	80.0	92.7
25 MPY	5.7	4.9	74.7	80.0	5.5	4.6	80.0	92.1
30 MPY	5.8	5.2	76.0	80.0	5.5	4.9	80.0	91.6
35 MPY	5.8	5.5	77.3	80.0	5.5	5.1	80.0	91.0
40 MPY	5.9	5.8	78.7	80.0	5.5	5.3	80.0	90.4
44 MPY	6.0	6.0	79.7	80.0	5.5	5.5	80.0	90.0

The acceptable VaR level relative to the required return is set at 2.5% for participants who are at target. Essentially, in the worst outcome as measured by VaR we accept a loss relative to our actuarially required return (4.4%) that is equal to expected inflation (i.e. 2.5%)[15]. However as the "at target" participant nears retirement the confidence level required in the VaR does increase, going from 68% at age 20 to 80% at age 64. In the decumulation phase we use a similar increase in the confidence level for the VaR, increasing from 80% at age 65 to 95% at age 106. This VaR is calculated relative to the allowed shortfall we defined vis-à-vis the 5.5% annual nominal withdrawals.

In Table 2 we also show the treatment of participants who are not at target. Every row in the table represents the evolution of the parameters along a line parallel to the upper diagonal in the participation grid. As the missed participation years increase, so does the required annual return as well as the allowed shortfall relative to that required return. This reflects that the trade-off between making up lost ground (through higher investment returns by allowing more risk) and protecting accumulated capital is different for the different participants. For both the accumulation and decumulation phases all relative VaRs are calculated at the specified confidence levels from the grid in Figure 6 and with the other parameters as shown in Figure 4. Using this set of confidence levels, allowed shortfalls and required returns we can find the highest-returning portfolio for each point in the participation grid that still meets the criteria.

Where conventional schemes have a equity weight "glide path," ours produces a glide surface due to the higher dimensionality of the analysis.

[14] Strictly speaking each point represents a one-year age cohort in one dimension, and a one missed-participation year cohort in the other dimension. For the sake of brevity though we will also refer to participants at these points in the grid, rather than the cohorts. So a "Model Participant" could equally be read to mean "Model Cohort."

[15] In this specific instance. In general this need not be equal to the inflation rate but it does provide for conceptual consistency in that we demand at least a positive nominal outcome in the worst case.

Any forecasts in the table above represent hypothetical numbers and are purely for illustrative purposes only. The numbers do not represent actual or future performance.

This process leads to a glide surface with explicit asset allocations for all points in the Participation Grid. To show this we present in Figure 9 the equity weights of these optimized asset allocations. Color changes denote the change in equity weight within the Participation Grid, where aggressive portfolios are red and conservative portfolios are green. For example a participant at retirement age with full participation will have an equity weight of 30%, whereas the allocation for a participant of the same age just entering the system would be 65%.

Figure 9: Glide Surface Equity Weights

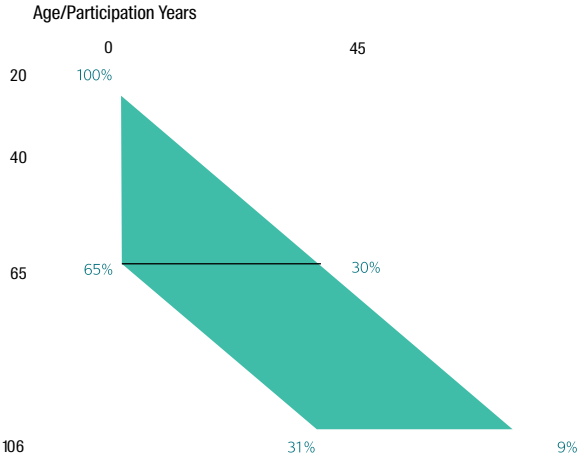

Source: First Sentier Investors

The equity allocation for participants whose capital is above target and therefore fall outside of the grid is set equal to the equity weight of the "at target" participant of the same age. Additional lock- in mechanisms for heavily overfunded participants are possible within the scheme too, for instance by locking in any excess capital by investing that portion in a risk-free asset.

Projecting outcomes forward to retirement age provides a glimpse into the underlying ranges of numbers.

Pension Results – Accumulation Phase

In the full program design we repeat the preceding analysis for every point (i.e. every participant) in the Participation Grid. In Table 3 we show some basic results for a cross-section of model participants. Consider the 35 year old with zero missed years, who therefore can be considered "at target." The start allocation to equity is 71%. Since this participant is "at target" the equity allocation at retirement age is 30% as it is for all other "at target" participants. Since this participant has no missed participation years the target annuity is 30%. Obviously the capital at retirement is not known with certainty but based on our stochastic output we calculate that the average capital balance is $2,454,000 and that there is a 5% probability that the accumulated capital will be less than $1,744,000[16]. This then implies that this participant, if the outcome is right at the expectation value, will be able to fund an annuity that delivers an income stream equivalent to 40.5% of projected salary 30 years from now. In the worst 5% of cases the participant will be able to fund an annuity at most equivalent to 27.2% of salary. Note also that for some of the sample participants shown in Table 3 the annuity target has changed to a lower value than the 30% overall scheme target. This is due to the large number of missing participation years which make attaining the original target infeasible.

[16] Note that these dollar numbers are capital balances in 30 years' time, incorporating both contributions and capital returns throughout this period.

Table 3: Sample Participant Outcomes

		Equity Weight		Capital at Retirement			Annuity Percentage at Retirement				
Age	Missed Years	Start (%)	At Retirement (%)	5% Worst Case	Average	5% Best Case	5% Worst Case (%)	Average (%)	5% Best Case (%)	Target (%)	Prob Achieving Target (%)
20	0	100.0	30.0	1744	2454	3508	29.9	44.5	65.7	30.0	94.9
35	0	71.0	30.0	1048	1480	2042	27.2	40.5	57.5	30.0	89.0
45	0	50.0	30.0	673	917	1207	26.0	36.5	49.4	30.0	82.2
35	15	93.0	41.0	750	1085	1520	19.6	29.7	43.1	25.1	71.5
45	10	62.0	38.0	574	822	1138	22.2	32.7	46.7	29.2	64.6
45	25	80.0	49.0	328	459	630	12.7	18.3	25.4	16.2	66.5

Capital shown in thousands of dollars

As one can see the average outcomes overshoot the targets significantly[17] and in all cases there is at least a two-thirds probability of achieving the targets, with full participation in the system delivering the highest probability of achieving the target. The fact that on average the outcomes overshoot the target reflects the risk criterion in the asset allocation. That is, we use VaR, a downside risk measure, as opposed to volatility around the annuity outcome.

To get a better feel for the implications of the design of the Participation Grid for actual pension outcomes we will show more in depth simulated results for two sample model participants. The first example, Participant A, is the at-target participant who joins the program at age 20, also presented in the first row of Table 3. The second example, Participant B, is a 45 year old participant who enters the program with no accumulated capital whatsoever, meaning that this participant missed out on 25 participation years as shown in the bottom row in Table 3[18].

In Figures 10 (a) and 10 (b) we show the dynamic asset allocation for Participant A in the accumulation phase. As one can see the equity allocation becomes more conservative over time[19]. In the first few years the equity allocation is 100% and in the last year before retirement the equity allocation is 30%.

Figure 10: Dynamic Allocations for Participant A

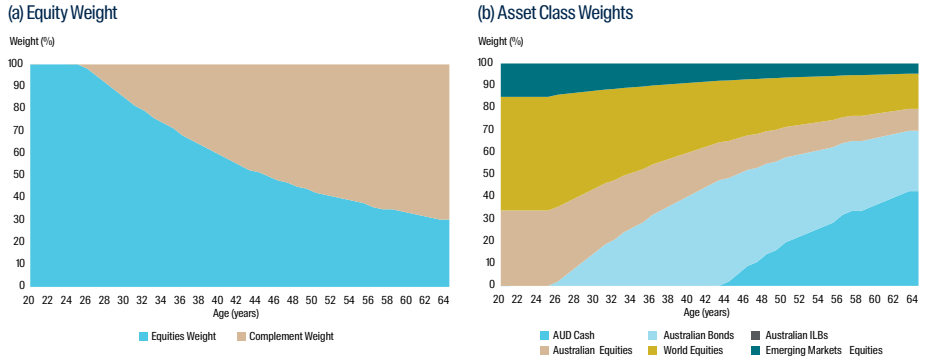

Source: First Sentier Investors

[17] This is also what we showed conceptually in Figures 2 and 3.

[18] Participants A and B also appear in the conceptual design grid in Figure 3.

[19] While this may appear to be in line superficially with traditional lifecycle funds, the fact that these portfolios have all been individually optimized with respect to an explicit annuity target with built-in dynamic guidance, the overall pension scheme avoids the documented suboptimality of existing lifecycle and target date funds.

Any forecasts in the table above represent hypothetical numbers and are purely for illustrative purposes only. The numbers do not represent actual or future performance

In Figures 11 (a) and 11 (b) we show the dynamic asset allocation for Participant B in the accumulation phase. In the first year the equity allocation is around 80% and in the last year before retirement the equity allocation is 49%, which is a significantly different pattern than the one for Participant A. This is due to the late entry into system without any accumulated capital. That is, Participant B has a higher required return and higher risk tolerance than Participant A. This reflects, in part, the greater amount of capital that Participant A has accumulated.

Figure 11: Dynamic Allocations for Participant B

Source: First Sentier Investors

Using the optimized allocation paths for the chosen model participants we can show the specific bandwidths of outcomes in terms of capital and annuity percentages.

By implementing the dynamic asset allocation for all participants we can estimate the probability distributions of built-up capital and built-up annuities as percentage of salary over time.

These results are shown in Figures 12 (a) and 12 (b) (Participant A) and 13 (a) and 13 (b) (Participant B). Figures 12 (a) and 12 (b) show the annual bandwidth of outcomes, with the big blue dot in the center being the median outcome. We also show the upper and lower quartiles (these being the purple and green dots) as well as the 5th and 95th percentiles of outcomes at the extremes of the bandwidth, represented by small lilac and gray dots. We also show the "at target" path in the gray shaded area. As Figures 12 (a) and 12 (b) show, Participant A's median accumulated capital and median annuity equivalent end well above the target. The worst case annuity outcomes at retirement age still satisfy the target annuity objective. This is as it should be, for this was the optimization criterion in the first place. In other words, the probability of not achieving the annuity target of 30% at retirement is approximately only 5% for Participant A. This is also consistent with Figure 2 with the grid represented by the gray area, the dashed blue line by the median and the interquartile range by the distance between the purple and green dots. Figure 2 does not contain the 5th and 95th percentile outcomes though.

Figure 12 (a): Percentiles of Participant A's Capital

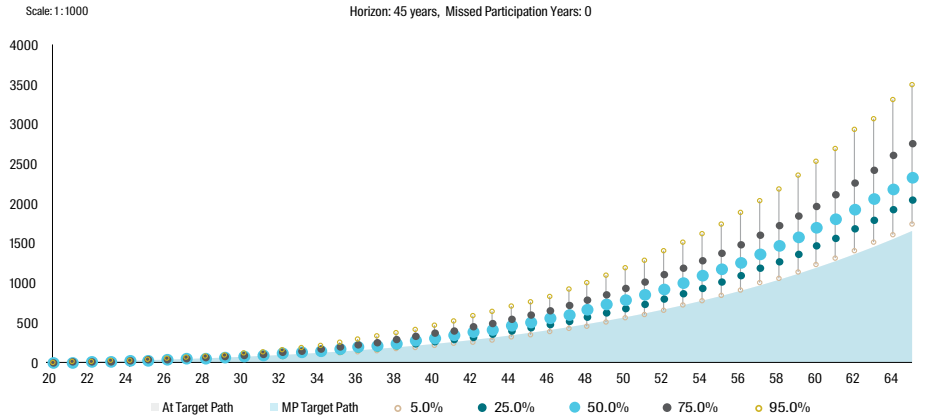

Source: First Sentier Investors

Figure 12 (b): Percentiles of Participant A's Annuity Percentage

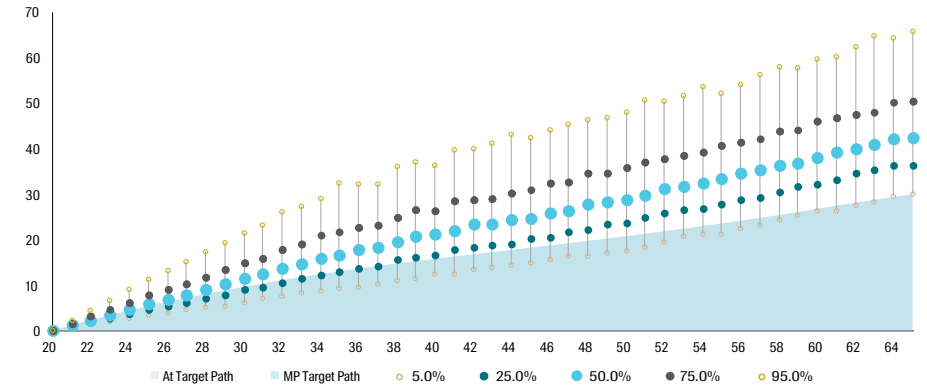

Source: First Sentier Investors

The decreasing level of risk tolerance from young to old age is clearly visible in the annuity chart in Figure 12 (b). Here we see that the 5th and 25th percentiles initially fall within the gray target area, implying a lower level of certainty of attaining the target at that point. By increasing the certainty of outcomes over time we can achieve the desired outcome in terms of annuity without foregoing higher return potential over the full duration of participation in the scheme. This facet is neatly accounted for in the design and set up of the retirement savings system.

For Participant B the estimated probability distributions of capital and corresponding annuity levels are shown in Figures 13 (a) and 13 (b). As this participant has missed 25 years of capital build up in the scheme, the charts are very different from Participant A's. Since this participant has zero capital start with and only 20 years to go to retirement, the original scheme target of 30% of annuity is no longer feasible. Instead the modified target for Participant B is an annuity equal to 16% of salary. We show both the overall scheme target path of 30% annuity and the modified target path of 16% annuity in Figure 13 (a) and 13 (b), represented respectively by the light and dark gray shaded areas.

Figure 13 (a): Percentiles of Participant B's Capital

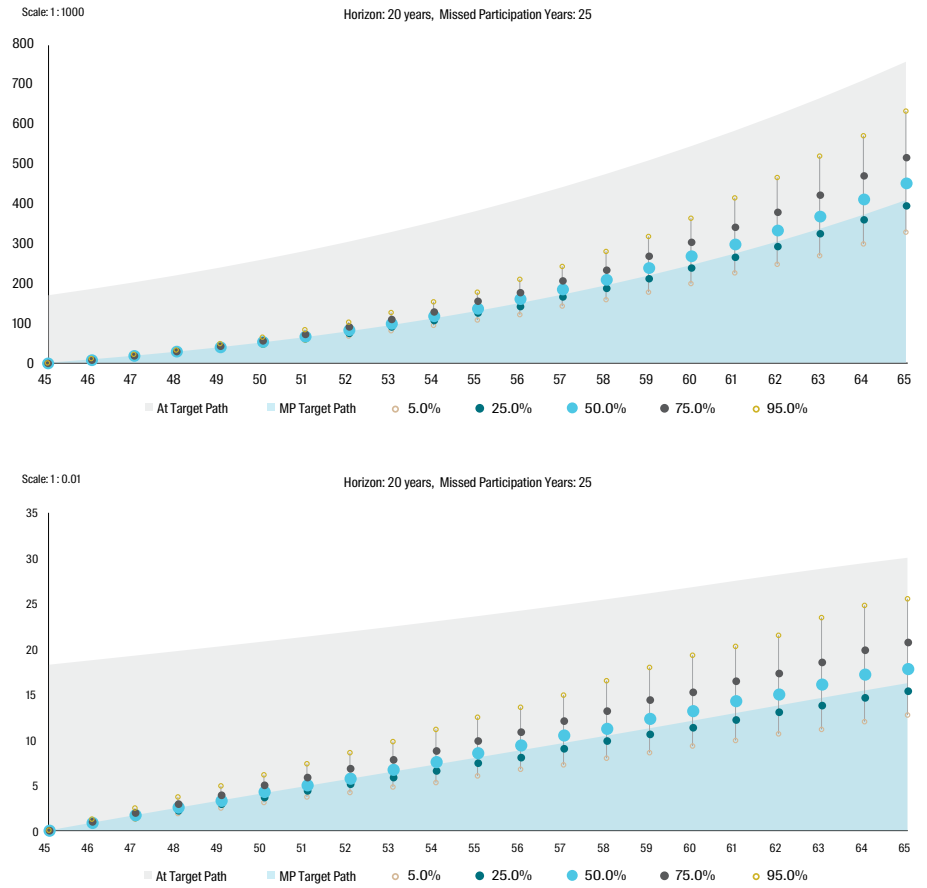

Source: First Sentier Investors

Pension Results – Decumulation Phase

Moving on to the decumulation phase, we now highlight the optimization results for the period after retirement. In Figures 14 and 15 the dynamic asset allocations for both participants are shown for the decumulation phase, with these charts essentially being diagonal slices from Figure 9. Again we note the difference in the results between the two sample participants.

The lower equity weight for Participant A reflects the protection aspect of the allocation as this participant has a large amount of capital built up. For instance, the equity allocation in the first year after retirement is 27% for Participant A and 20 years later it has decreased to 15%. On the other hand for Participant B the fraction of the portfolio invested in equities is 47% in the first year after retirement and after 20 years later the equity allocation has dwindled to 31%.

These declining equity weights are the results of the optimization relative to the parameters for retirement we outlined in an earlier section.

Figure 14: Dynamic Allocations for Participant A

Source: First Sentier Investors

Figure 15: Dynamic Allocations for Participant B

Source: First Sentier Investors

Again we can look at the impact of actually investing the participants' portfolios according to these dynamic asset allocations over time, and estimate the probability distributions of capital, capital translated into annuity levels and total income in the decumulation phase. The results after retirement depend on the amount of accumulated capital that is used to actually acquire indexed annuities. We will exemplify this in the next few Figures.

As we previously mentioned the system allows for flexibility in retirement in buying annuities. The participants can decide to use all or part of their built-up capital to buy an insured lifelong annuity. This flexibility is important not only because it allows the participants to choose how much certainty they are willing to trade for additional upside potential, but also because it allows this scheme to be used more widely with a guided or designed annuity purchase scheme.

There is no unambiguous optimal allocation to annuities, or mixes of annuity of various lengths as the trade-off will be an individual one. For instance the optimal age to buy annuities depends on the participants' bequest utility as well as the asset allocation of the portfolio during retirement as described in Blake, Cairns and Dowd 2003. Investigating the aggregate welfare of annuitization options for the UK Einav, Finkelstein and Schrimpf 2010 conclude that forcing participants to choose the longest possible guarantee period maximizes aggregate welfare, however also concludes that actually mandating this is unlikely to be practicable. Trying to find a more practical approach to the issue of how to invest during the decumulation phase some propose a benchmark of laddered TIPS and an indexed annuity, as for instance Sexauer, Peskin and Cassidy 2012. This latter approach can be easily accommodated within the framework we propose here.

In order to illustrate the effect of various choices of participants we show three cases here:

1. The participant does not buy any annuities but continues to invest the accumulated capital

2. The participant uses 50% of the accumulated capital to buy an annuity, and leaves the other 50% invested

3. The participant uses all capital to buy an annuity and does not invest any capital after retirement

The scheme can accommodate flexibility for participants to buy annuities at any time with any fraction of their capital, even before retirement.

We also assume that the participants only buy annuities at one specific point in time. In practice the scheme allows for the purchase of annuities at any point. In the following section we will present sets of three charts for both participants for all three options enumerated above. The three charts per set show respectively the bandwidths of accumulated capital, the annuity percentage this capital translates into, as well as the actual benefits percentage. The latter represents the actual income extracted from the capital plus the income from any annuities purchased. The shaded gray area after retirement represents the targeted withdrawal.

In Figure 16 the results are shown for the case where Participant A chooses not to use any capital to buy annuities but to keep the portfolio fully invested. In that case, after retirement a 3% annual withdrawal from the available capital is assumed as actual income and this is reflected in the bending downwards of the bandwidths at age 65 in the capital percentiles chart 16 (a).

Since the expected return is higher than this 3% the remaining capital will generally grow, leading to an increase in absolute and even real value of these withdrawals over time. However, this does not really reflect an attractive pay-off, since the income shortly after retirement is relatively low and there is a lot of "unused" capital left at very old ages as is visible in Figure 16 (a)[20]. Figure 16 (b) shows the value of the annuity that the participant could buy with this accumulated capital. Figure 16 (c) represents the income from the capital; in this case this is just the 3% annual withdrawal as the capital is left fully invested and Participant A buys no annuities.

Figure 16: Percentiles of Outcomes for Participant A

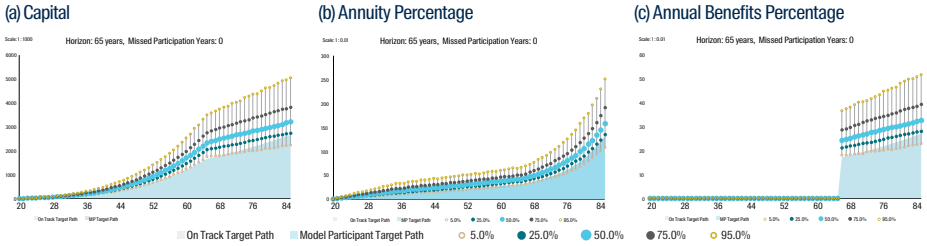

Source: First Sentier Investors

In Figure 17 the equivalent results are shown for Participant B. The results are analogous to the ones for Participant A with the exception that there is a greater likelihood of a capital shortfall being experienced by Participant B and consequently there is a higher risk that the desired retirement income will not be generated.

Figure 17: Percentiles of Outcomes for Participant B

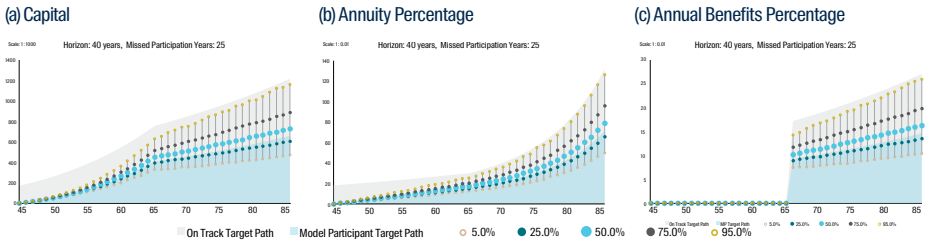

Source: First Sentier Investors

[20] We have cut off the graphs at 85 years of age for the sake of clarity in the charts but the actual scheme design runs until the end of the life tables at age 107.

Figure 18 shows the results in case Participant A chooses to use half of the accumulated capital at retirement to buy an indexed annuity. Figure 18 (a) shows the depletion of the capital as half of it used to purchase the annuity at age 65, while the other half remains invested and then continues to grow. In addition there is still a 3% withdrawal from this remaining half of the capital. Figure 18 (b) is again the translation of the remaining capital into its annuity equivalent. Figure 18 (c) now is the sum of this 3% withdrawal from capital and the income from the annuity that the participant bought. Compared to the situation shown in Figure 16 (c) we now see more stability in the income as might be expected from the annuity component. That is, the 5% of worst outcome sits wholly below the target path in Figure 16 (c), whereas in Figure 18 (c), the 5% of worst outcomes sits roughly at the target path. This is however balanced by having less upside as there is only half as much capital invested.

Figure 18: Percentiles of Outcomes for Participant A

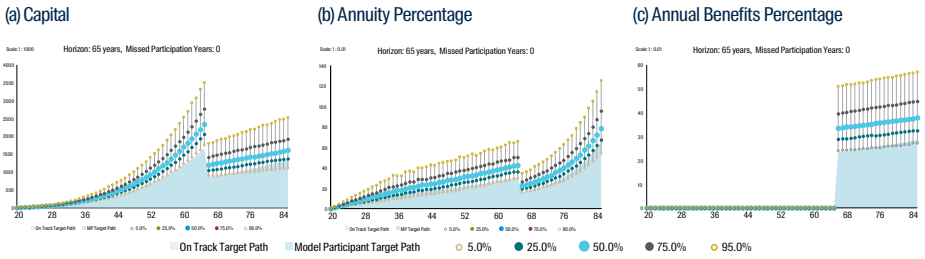

Source: First Sentier Investors

Figure 19 shows the results for Participant B with broadly similar changes as for Participant A going from buying no annuities to buying annuities with 50% of capital at age 65.

Figure 19: Percentiles of Outcomes for Participant B

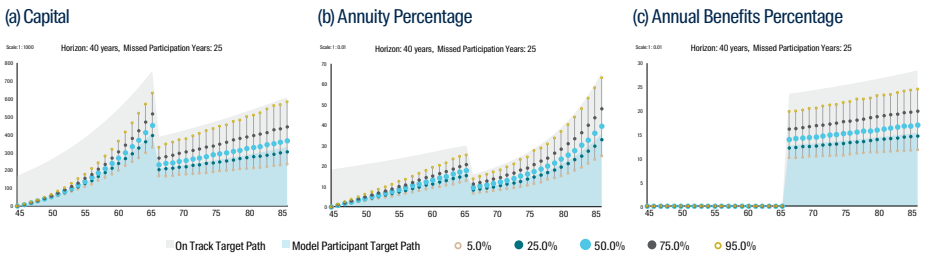

Source: First Sentier Investors

Finally we look at the case where both participants choose to cash in their entire capital at age 65 and use it to buy lifelong indexed annuities, which we show in Figures 20 and 21 respectively. Figures 20 (a) and 21 (a) show the complete depletion of capital and the bandwidths obviously also collapse to zero. Figures 20 (b) and 21 (b), which are translations of the capitals into its annuity equivalent, also correspondingly collapse to zero. The income post retirement, as shown in Figures 20 (c) and 21 (c), now becomes certain as income is sourced solely from the purchased annuity. Essentially the distribution of capital at retirement is set in stone the moment you buy the annuity. The bandwidths do persist though as there is a point risk at the time of retirement which directly influences the annuity that can be purchased.

In this case there is no "unused" capital, so there is no upside potential in the annuity level either. Also, by spending all accumulated capital in acquiring an annuity at retirement there is no discretion at all anymore for the participant with regards to the use of that capital. For example it is no longer possible to acquire additional tranches of annuities later based on favorable returns or making extra withdrawals if needed, which makes it a very safe, but the least flexible, solution.

Figure 20: Percentiles of Outcomes for Participant A

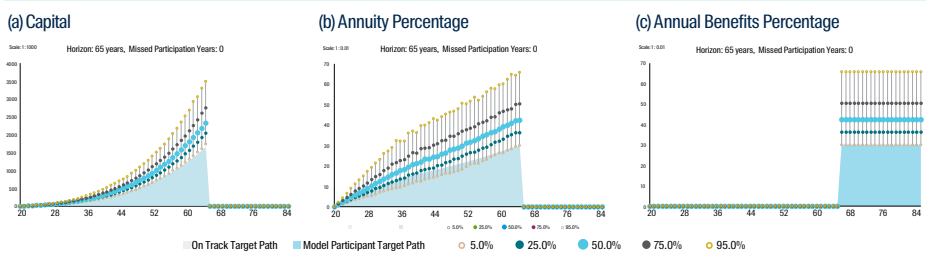

Source: First Sentier Investors

Figure 21: Percentiles of Outcomes for Participant B

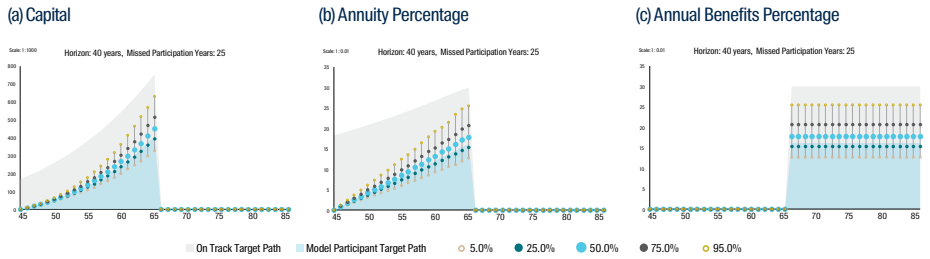

Source: First Sentier Investors

Active Management: Adding Tracking Error and Alpha

Having established optimal dynamic strategies for each participant in the Participation Grid, the next step is one of implementation. We have used only pure β exposure up until this point, which in most cases can be replicated passively. In many instances some form of active implementation does make sense though, and it is crucially important to assess the viable leeway for active management in the same risk/return space as the β – only analysis. Setting a risk budget is a natural extension of setting a strategy, and a comprehensive methodology is described in Baars, Kocourek and van der Lende 2012.

Taking the dynamic asset allocations for a participant as the starting point, we perturb the allocations in three ways:

1. By adding tracking error with zero alpha to the portfolio
2. By adding uncorrelated alpha with zero tracking error to the portfolio
3. By adding a combination of uncorrelated alpha and tracking error with a fixed information ratio to the portfolio

For the at-target model Participant A, we calculate for each perturbation the expected annuity percentage and the 5% worst case annuity percentage at retirement age. The results are shown in Figure 22.

Active management of the portfolios can add valuable additional alpha at reasonable information ratios.

Figure 22: Model Participant Risk Analysis

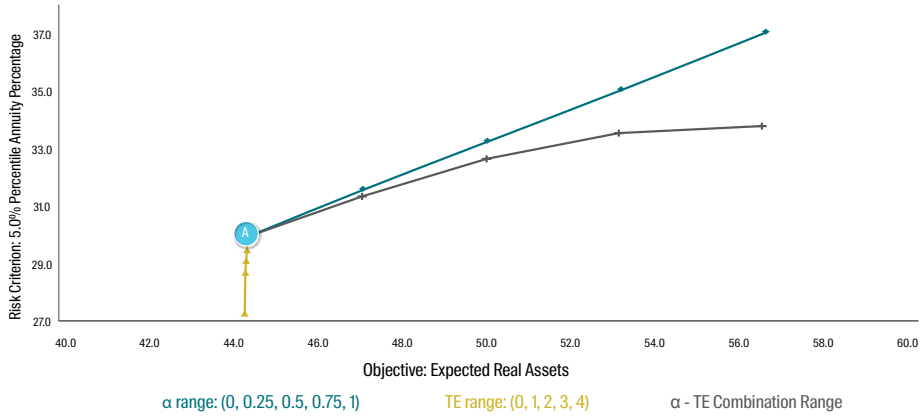

Source: First Sentier Investors

This chart shows that adding tracking error with no α (represented by the green line) results in deterioration in the worst case outcome, but has no effect on the expected outcome. This is also what one would expect as it is merely adding volatility without changing the mean of a distribution. Adding uncorrelated α improves both the expected outcome as well as the worst-case outcome (represented by the purple line). After all, we are adding "pure performance" to the portfolios without any penalty. Looking at the combination of adding both α and tracking error with a fixed information ratio of 0.25 we actually observe a flattening of the gray line.

This means that the worst-case outcome ceases to improve once the tracking error exceeds about 3%, and the incremental benefit of adding tracking error becomes marginal. The expected outcome still shows improvement though.

Alternative Implementation Options

In the previous section we explained how one can construct a comprehensive solution for a target benefit approach. We can actually apply this methodology to produce a spectrum of granularity in the implementation. In Figure 23 we show how the Participant Grid allows for a number of implementation options in the accumulation phase.

The first option is to use a range of target benefit funds for the full participation case only. In this option the asset allocation is done within the funds. This option is operationally the least complex and only differs from industry standard target date funds by having explicit targets and risk profiles with the main disadvantage being that it only caters to the modeled, or "typical" participants. The second option is to use target benefit funds for certain model participants in the Participant Grid. This model participant approach uses a tailored number of strategically chosen participant cohorts and is operationally of medium complexity, which allows the range of funds to cover a much wider range of potential real participants, rather than just the "typical" ones of the first option. The third option is the most comprehensive solution with annual cohorts by age and accumulated capital as described in the previous sections in this paper and allows for the best matching of actual participants with the modeled ones. This option has the highest operational complexity.

Operational complexity can be a stumbling block for implementation, but there are ways of ameliorating this.

Figure 23

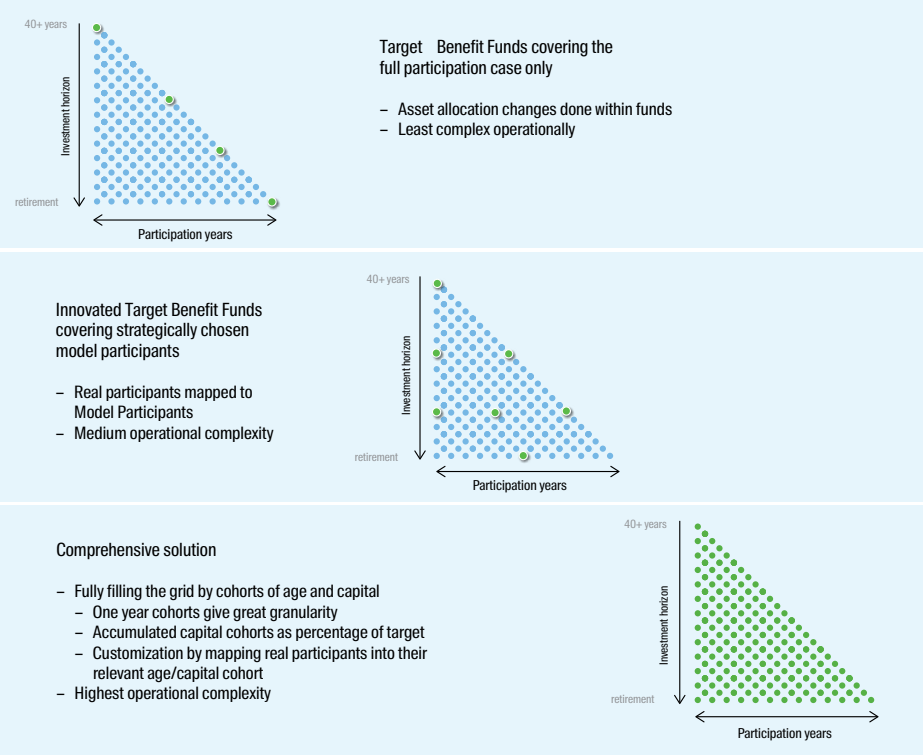

Source: First Sentier Investors

Example of Strategic Target Benefit Funds

In Figure 24 we have depicted one practical application of the general target benefit concept. In this example the implementation of the concept uses six strategically chosen target benefit funds for the accumulation phase to cover a cross section of the participant population.

- Participants younger than 30 years are all assumed to be at target, represented by TBF 57F
- Participants between 30-40 are all assumed to be at target, represented by TBF 47F
- Participants between 40-50 represented by TBF 37F for those who are at target and TBF 37U for those who are underfunded.
- Participants between 50-60 represented by TBF 27F for those who are at target and TBF 27U for those that are underfunded.

As explained in the previous sections, in the decumulation phase a participant can choose to continue investing (part of) his accumulated capital or to acquire annuities from it. In case a participant chooses to continue investing, one of the options is to buy into a CPI-plus type product, e.g. CPI-plus 3%. This type of product targets an optional perpetual withdrawal of 3% annually whilst preserving real capital. This is a somewhat simpler solution than the one we have described in the sample design in this paper, since we do not vary the confidence levels in this case. In our paper Baars, Kocourek and van der Lende 2012b, we describe how a CPI-plus type product can be designed.

Figure 24

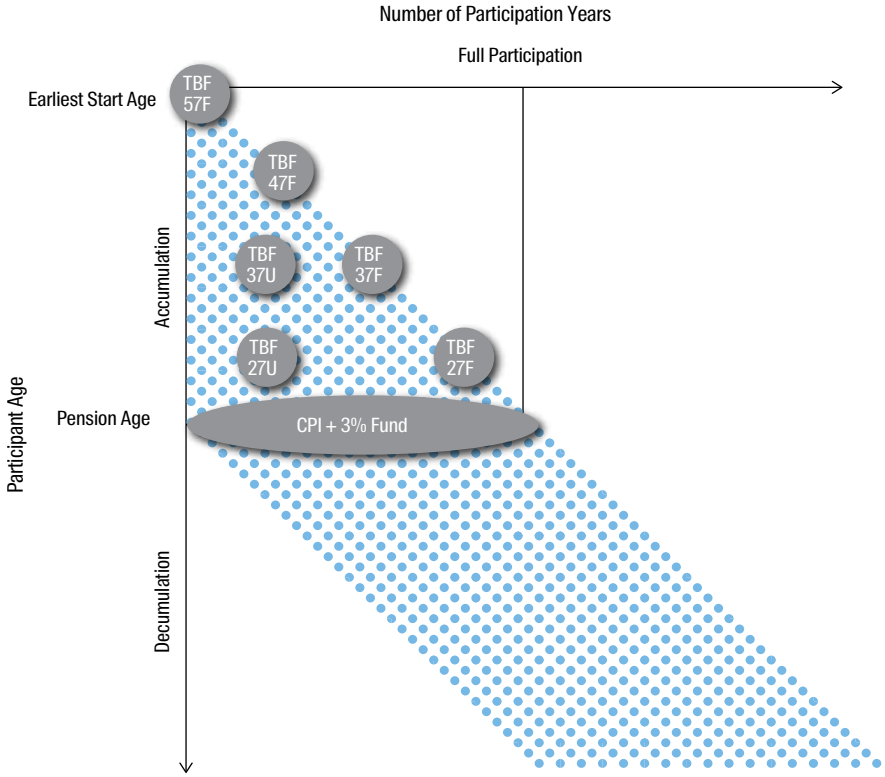

Global Applications

This paper used the Australian superannuation system as an example to show the mechanics and outcomes of designing a target benefit scheme. As we have noted previously though, the applicability of this general concept is global and can be used wherever there are defined contribution pension systems. One obviously example would be the 401(k) market in the United States, where target date funds have seen large inflows since the passage of the Pension Protection Act of 2006 with its "safe harbor" protection for auto-enrollment. Using the target benefit approach as presented here it is possible to create appropriate target date funds that avoid the pitfalls of the ones currently in existence.

The multifondos approach used in Latin American pensions can also provide a framework for target benefit funds. For instance in the case of Chile the legislation provides for five different risk categories of funds ranging from very aggressive to very conservative with limits on equity exposure for each category. Since participants are allowed to choose any combination of two funds this target benefit scheme could be applied as well by carefully choosing the model participants in the Grid. The pension system in Chile also allows the option to buy annuities at retirement or to make scheduled withdrawal in the decumulation phase. The multifondos in Peru and Colombia offer less flexibility as only three categories of riskiness exist.

In Singapore the Central Provident Fund (CPF) is offering four options for annuitization known as CPF Life with the possibility of the government mandating full annuitization upon retirement to prevent cash-outs and rapid depletion of the cash. This is another example that shows the necessity of providing an integrated solution covering both the accumulation and decumulation phase and the various attempts to address this globally.

Summary

We presented our approach on designing a target benefit scheme that addresses the shortcomings of tradition lifecycle and target date funds. In the design we make essential use of our Asset Liability Management toolkit which is required to derive optimal asset allocations for participants of differing backgrounds in terms of capital and age. By explicitly targeting a pre-defined annuity level at retirement we are able to derive for each participant in the target benefit scheme an appropriate investment strategy that minimizes the probability of not providing this annuity. In this sample design we have also included a consistent way to optimize investment strategies for participants after retirement. The scheme presented here also provides flexibility in actual implementation, allowing it to be tailored to practically feasible representations of the underlying concept.

References

Anson, M., "Strategic versus Tactical Asset Allocation", Journal of Portfolio Management, Winter 2004.

Baars, Jan, Petr Kocourek, and Epco van der Lende. "Integrated Approach to ALM, Risk Budgeting and Alpha." First Sentier Investors Multi-Asset Research Papers, no. 2 (June 2012a).

Baars, Jan, Petr Kocourek, and Epco van der Lende. "Inflation-Linked Investment Objectives with Liquid Asset Portfolios." First Sentier Investors Multi-Asset Solutions Research Papers, no. 3 (June 2012b).

Basu, Anup K., Alistair Byrne, and Michael E. Drew. "Dynamic Lifecycle Strategies for Target Date Retirement Funds." The Journal of Portfolio Management 37, no. 2 (2011): 83-96.

Basu, Anup K., and Michael E. Drew. "Portfolio Size Effect in Retirement Accounts: What Does It Imply for Lifecycle Asset Allocation Funds?" The Journal of Portfolio Management 35, no. 3 (2009): 61-72.

Blake, David, Andrew J.G. Cairns, and Kevin Dowd. "Pensionmetrics 2: Stochastic pension plan design during the distribution phase." Mimeo (Pensions Institute, Birkbeck College), 2003.

Blake, David, Andrew J.G. Cairns, and Kevin Dowd. "Pensionmetrics: stochastic pension plan design and value-at-risk during the accumulation phase." Insurance: Mathematics and Economics 29 (2001): 187-215.

Bodie, Zvi, and Jonathan Treussard. "Making Investment Choices as Simple as Possible, but Not Simpler." Financial Analysts Journal 63, no. 3 (May-June 2007): 42-47.

Booth, Philip, and Yakoub Yakoubov. "Investment Policy for Defined-Contribution Pension Scheme Members Close to Retirement: An Analysis of the "Lifestyle" Concept." North American Actuarial Journal 4, no. 2 (2000): 1-19.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the UK Annuity Market." Econometrica 78, no. 3 (May 2010): 1031-1092.

Kyrychenko, Vladyslav. "Optimal Asset Allocation in the Presence of Nonfinancial Assets." Financial Services Review 17 (2008): 69-86.

Pfau, Wade D. "The Portfolio Size Effect and Lifecycle Asset Allocation Funds: A Different Perspective." The Journal of Portfolio Management 37, no. 3 (2011): 44-53.

Sexauer, Stephen C., Michael W. Peskin, and Daniel Cassidy. "Making Retirement Income Last a Lifetime." Financial Analysts Journal 68, no. 1 (2012): 74-84.

Shiller, Robert J. "The Life-Cycle Personal Accounts Proposal for Social Security: An Evaluation of President Bush's Proposal." Journal of Policy Modelling 28 (2006): 427-444.