//DWS

Optimizing Life-Cycle Strategies for Savings Plans

A Practitioner's Perspective

Dr. Sebastian Kring

Senior Structuring
Manager, DWS

IN A NUTSHELL

— Life-Cycle (LC) strategies for savings plans are well-supported by the personal finance literature. Their use is widespread in defined contribution (DC) plans, as well as other private savings solutions.
— However, optimizing these LC strategies according to an individual investor's goals and objectives can be challenging, due to the multi-period, and long-term investment context.
— Practitioners, including asset managers and individual financial advisors (IFAs), often employ a select and benchmark (SB) approach to develop LC strategies customized for individual investors.
— Arguably a better approach may be through optimization techniques, such as dynamic programming (DP), which seeks to enable the direct optimization of LC strategies to align with the specific goals and objectives of an individual investor within a multi-period investment framework.
— Individualized Asset Management (IAM) platforms may provide individual investors with efficient and scalable access to such optimized LC strategies tailored to their goals and objectives.

Introduction

Thoughtful, long-term savings plan solutions are crucial for individual investors because they may help address the pension gap created by state pension systems in many countries. In the developed world, state pension systems face structural challenges due to low birth rates over several decades, and the resulting demographic changes. To close the pension gap for individual investors, policymakers globally are promoting savings plan solutions through occupational and private pension schemes.

In the academic personal finance literature, LC strategies are considered one of the most favored investment strategies for these savings plans. Many practitioners, including asset managers and independent financial advisors (IFAs), also advocate for LC strategies in these solutions.

Typically, LC strategies share a common principle: young individual investors with a long investment horizon until their retirement date, and a relatively small amount of financial assets in their savings plan, compared to their future lifetime savings, tend to have a more "return-seeking" allocation in their savings plan accounts. Middle aged investors tend to have a more "balanced" allocation in their accounts. Finally, older investors, with a short investment horizon until their retirement date, and a substantial amount of financial assets, typically have a more "conservative" allocation in their accounts. The key question today, therefore, is how to develop a prudent LC strategy that reduces the risk-return profile of the allocation in savings plan solutions as the investor ages. This reduction aims to achieve the investors' goals and objectives, which, in most cases, means closing the pension gap.

Practitioners, such as asset managers, banks, and individual financial advisers, typically derive LC strategies via several steps. They begin with a single-period portfolio optimization approach, such as mean-variance optimization, to create several optimized portfolios. The risk-return profiles these portfolios range from a risk-reduced to a return-seeking positioning.

Next, they construct different LC strategies using an age-based risk-reduction strategy - referred to as a "Glide-Path" in this paper. A Glide-Path determines the transition from a return-seeking to a risk-reduced portfolio over time by allocating the investor's savings plan assets into these optimized portfolios over time. According to their multi-period nature these LC strategies are then evaluated using Monte Carlo simulations by projecting the development of various representative savings plans into the future. Finally, the simulation results are used to determine the overall risk-return profiles of the LC strategies and to benchmark them against the objectives, goals, and risk-return profiles of the individual investors.

From discussions with other practitioners this approach has shown effective and has served us and other practitioners well in the past. We refer to this method as the "select and benchmark" (SB) approach because it involves selecting optimized portfolios and Glide-Paths, as well as benchmarking the resulting LC strategies against the objectives, goals, and risk-return profiles of the individual investors. We will call LC strategies derived using the SB approach "SB-LC-strategies" in this report.

Today however - driven by technological progress, digitalization, and a substantial increase in computational power - more systematic approaches such as dynamic programming (DP) can be used to derive optimal investment strategies in a multi-period investment context – approaches that still align directly with the goals and objectives of their individual investors. In this DP approach, the goals and objectives of the individual investor become the objective function. Then, through an optimization process performed in a multi-period investment environment, an optimized investment strategy is derived that maximizes the objective function of the individual investor in question.

As we will see later, for well-established objective functions - such as a Constant Relative Risk Aversion (CRRA) utility - these optimized investment strategies also exhibit an age-based de-risking pattern in the allocation of the savings plan. Compared to the SB approach, this method has the advantage of directly deriving optimized strategies by solving the multi-period investment problem, rather than selecting LC strategies and benchmarking their behavior against the investor's objectives. We refer to LC strategies derived using this approach as "DP-LC-strategies" in this document.

Savings plan solutions with LC strategies are implemented in target date funds, also known as life-cycle funds (Viceira, 2008), or in individualized account-based solutions. Target date funds allow for very efficient implementation of a savings plan solution: investors select a target date fund whose target date aligns with the estimated end of their savings period as closely as possible, and then simply purchase shares over this period. In an individual account-based solution, each investor has their own account, which needs to be administered, but tailoring the LC strategy to the objectives of the account owner, i.e., the individual investor, is also possible.

The target date funds solution has the advantage of an efficient implementation as a one-size-fits-all solution. However, it has the disadvantage of not being able to account for the individual objectives and risk-return preferences of the investor. Conversely, the account-based solutions offer the advantage of customization to individual needs but require more administrative effort.

Historically, the trade-off between these two implementation approaches has often fallen in favor of target date funds due to the ease of their implementation. However, due to the technological progress mentioned above - which has led to efficient account platforms, web-based front-end solutions, investment apps, and investment engines providing rule-based and customized investment strategies – our view is that the trade-off will increasingly favor individualized account-based solutions for new savings plan solutions. This is why banks and asset managers are offering individualized asset management (IAM) platforms to provide more bespoke investment solutions.

This paper is organized as follows. In Section One, we will present the SB and DP approaches. In Section Two, we compare these two approaches by applying them to the savings plan of an individual investor who intends to maximize the expected CRRA-utility derived from the final value of the savings plan at the plan's end date. The last section concludes and provides an outlook.

1 / Concepts

In this chapter, we take a closer look into the SB and DP approaches for constructing LC strategies. We first note that financial regulators in many jurisdictions have established frameworks to classify individual investors based on their financial knowledge, capacity, and willingness to bear financial risks, resulting in an overall risk profile for each investor. And practitioners, such as banks, asset managers, and financial advisors, must develop savings plan solutions that adhere to the key principles of these regulatory frameworks.

Typically, during the development process, product providers use a version of the SB approach. This involves heuristic selection of the investment universe, investment constraints, and glide paths, followed by benchmarking the solution against the overall risk profile in one of the final steps.

Conversely, the DP approach directly translates the overall risk profiles and objectives into an objective function. Once the objective functions for an investor are determined, the optimized investment strategy for the savings plan solutions can be derived by solving the optimization problem.

We will now detail these two approaches.

Select and Benchmark Approach

In the SB approach, LC strategies for savings plan solutions are typically derived using the following six step procedure.

1. **Selects an asset class-specific investment universe**: This includes capital market assumptions for the investment universe, such as expected returns, and the covariance matrix, as well as investment constraints.

2. **Calculates the efficient frontier**: Based on the assumptions from step one, uses a portfolio optimization technique, such as mean-variance optimization, to calculate the efficient frontier.

3. **Selects model portfolios**: Typically, at least three model portfolios with distinct risk-return profiles are chosen from the efficient frontier. These are commonly labeled as "defensive", "balanced", and "dynamic", representing low, medium, and high risk-return profiles, respectively.

4. **Selects Glide Paths**: Choose several Glide Paths that allocate investments into these model portfolios. Note that the combination of model portfolios selected in step three and the Glide Paths determine the LC strategy.

5. **Calculates the risk-return profile**: Simulate the performance of different savings plans into the future using a Monte Carlo Simulation to calculate the risk-return profile of the LC strategies.

6. **Benchmarks the LC strategies**: Based on the calculations from step five, benchmark the LC strategies against the risk-return profiles and objectives of the individual investor(s) in scope. Choose the LC strategy that performs best in the benchmark test. If none of the simulation results are convincing, return to step three or four and proceed accordingly.

Let us take a closer look into this six-step procedure.

Table One shows the capital market assumptions used in the first step. In this example, we have three main asset classes: money market (MM), fixed income (FI), and equities (EQ).

	Exp. Returns	Volatility
MM	1.4%	0.5%
FI	3.5%	4.1%
EQ	6.5%	14.2%

Correlation	MM	FI	EQ
MM	100%	11%	-18%
FI	11%	100%	9%
EQ	-18%	9%	100%

Source: DWS - *For illustrative purposes only*

Table One: Capital market assumptions – expected returns, volatilities and correlations for the asset classes money market (MM), fixed income (FI) and equities (EQ)

In the second step, based on these capital market assumptions, the efficient frontier is derived using mean-variance optimization, as shown in Figure One (see below). In the third step, three model portfolios (Defensive, Balanced, and Dynamic) are selected on the efficient frontier. These portfolios have distinct risk-return profiles: low, medium, and high, respectively. The expected return per annum (p.a.) and the volatility p.a. for these portfolios are described in Table Two.

Based on these optimized model portfolios, different Glide Paths are selected in the fourth step. Figure Two shows three different Glide Paths. In Glide Paths One and Two, the de-risking process starts 15 years before the target date, typically the envisaged retirement of the individual investor. This means that the investor is 100% invested in the Dynamic model portfolio 15 years or more before the target date. Then, the allocation is shifted, on a pro-rata basis, from the Dynamic model portfolio to the Balanced portfolio.

In Glide Path One, 10 years before the target date, 100% is invested in the Balanced model portfolio, while in Glide Path Two, this 100% level is reached 7.5 years before the target date. In Glide Path One, from 5-10 years before the target date, the allocation is shifted from 100% Balanced to 100% Defensive on a pro-rata basis, and, in Glide Path Two, this is done from 0-7.5 years before the target date. Glide Path Two has a higher risk-return profile compared to Glide Path One, because, over the last 15 years before the target date, the allocation in Glide Path Two always has a higher risk-return profile than Glide Path One.

Glide Path Three has an even higher risk-return profile compared to Glide Path Two. This is because Glide Path Three starts the de-risking period 10 years before the target date by shifting the allocation from 100% Dynamic (10+ years before the target date) to 100% Balanced (5 years before the target date) on a pro-rata basis, and then from 100% Balanced to 100% Defensive over the last 5 years before the target date.

The fifth and sixth step will be discussed in the next section, where we perform the corresponding Monte-Carlo simulation.

Efficient Frontier
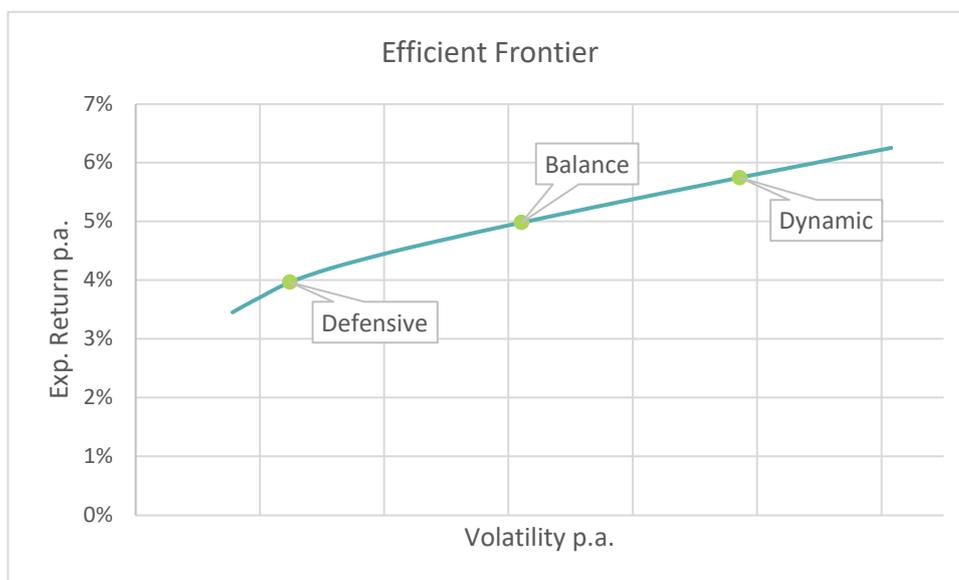
Source: DWS - For illustrative purposes only

Figure One shows the efficient frontier based on the capital market assumptions. Three model portfolios Defensive, Balanced and Dynamic with distinct risk-return profiles (low, medium, and high) are selected on the efficient frontier.

When deriving the LC strategy with this procedure, the model portfolios and the Glide Paths are the two components that determine the LC strategies. There is no algorithm or rule involved in the initial selection of these components. The selection in steps One, Three, and Four is based on preferences, experience, and expert judgment of the parties involved in the design process of the savings plan solution. This refers to the selection (S) part in this SB-approach. Typical selection criteria are:

- The model portfolio should not have an equity allocation much above 80%, while the defensive portfolio should not have an equity allocation much below 20%.

- The de-risking period should not be substantially below 10 years; otherwise, the LC strategies become too risky and vulnerable to historic stress scenarios.

Step Six refers to the benchmarking (B) in this approach. Typically, there are no strict rules on how to perform the benchmarking. This assessment is also performed based on experience and expert judgment of the parties involved in the design process of the solution.

	Exp. Returns	Volatility
Defensive	4.0%	4.5%
Balanced	5.0%	8.2%
Dynamic	5.7%	11.7%

Source: DWS - For illustrative purposes only

Table Two: Expected return p.a. and volatility p.a. of the selected model portfolios Defensive, Balanced and Dynamic.
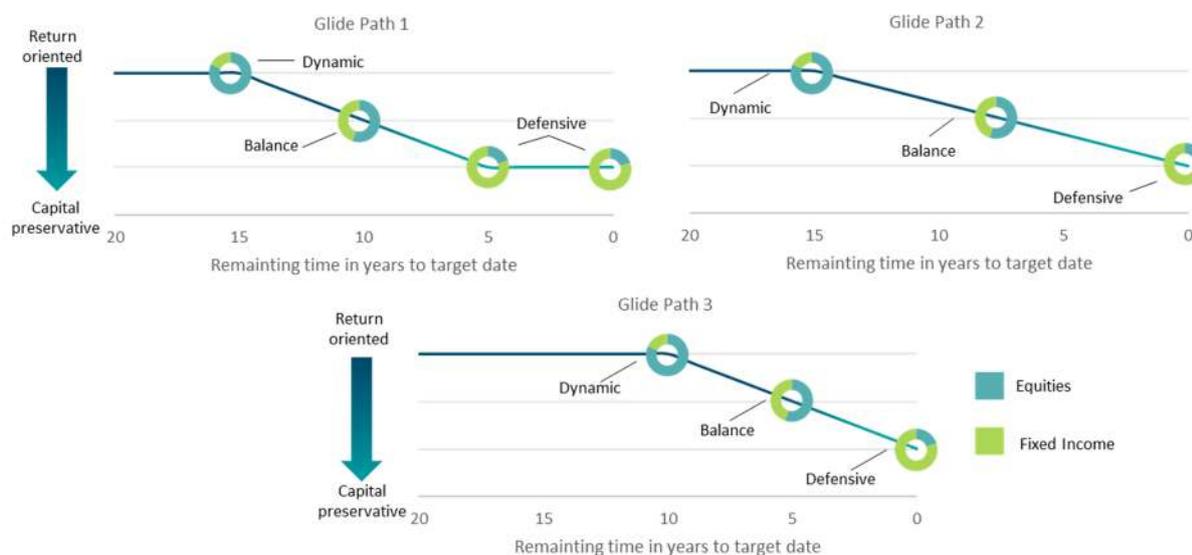
Source: DWS - For illustrative purposes only

Figure Two shows three different Glide Paths. The risk-return-profile of the SB-LC-strategies implied by the Glide Paths increases from Glide Path One to Three.

Dynamic Programming Approach

We switch now from the SB approach to the DP approach (an Introduction to Dynamic Programming and its application in solving multi-period savings plan problems is provided in the Appendix). Essentially, DP LC strategies can be derived using the following six step procedure:

1. **Selects an asset class-specific investment universe**: Defines the investment universe, including capital market assumptions such as expected returns, the covariance matrix, and investment constraints.

2. **Calculates the efficient frontier**: Uses a portfolio optimization technique, such as mean-variance optimization, to calculate the efficient frontier based on the assumptions from step One.

3. **Selects a grid of model portfolios**: Chooses a range of model portfolios along the efficient frontier, from low to high-risk profiles.

4. **Formulates goals and objectives**: Defines the individual investor's goals and objectives, and formulate them into an objective function.

5. **Derives the DP LC strategy**: By solving the multi-period optimization problem using dynamic programming, it determines the DP LC strategy.

6. **Reviews with Monte-Carlo simulation (optional)**: Optionally, reviews the DP LC strategy using Monte-Carlo simulations to understand the risk and return characteristics.

Let us take a closer look at this procedure. The first and second steps are the same as in the SB approach. In the third step, we select a discrete version of the efficient frontier, which essentially determines how many portfolios we want to evaluate in the fifth step. In the fourth step, the goals and objectives need to be mathematically formulated into an objective function representing the individual investor's goals and objectives. This step is crucial as it significantly influences the DP LC strategy. There is extensive literature on objective functions for investors – for a good introduction, see (Campbell & Viceira, 2002).

In this paper, we will use the CRRA-utility function in the next section. The fifth step is algorithmic, where the multi-period savings plan problem is solved using dynamic programming as outlined previously. The sixth step is optional because, by construction, the DP LC strategy is optimal for the objective function. However, particularly in the product concept phase, it is advisable to perform Monte-Carlo simulations for different savings plan configurations to gain an understanding of the strategy's risk-return profile. We will conduct this analysis in the next section as well.

Compared to the SB approach, the DP approach involves much less expert judgment. The most significant expert judgment is required when determining the objective function and mapping the individual investors' goals and objectives to this function. In the next section, we will apply these two approaches.

2 / Application

The savings plan problem with CRRA-utility

In this section, we apply the SB and DP approaches to a concrete savings plan problem. We assume that an individual investor will start contributing to a savings plan over the next 20 years by making quarterly contributions of EUR 1000. This results in 80 contributions over the 20-year period. The investor's goal is to maximize their expected Constant Relative Risk Aversion CRRA-utility $E[u_\gamma(v_T)|\pi^*]$[1] derived from the savings plan value at the target date, given an optimized investment strategy π^*.

$$E[u_\gamma(v_T)|\pi^*] = \max_{\pi \epsilon \Pi} E[u_\gamma(v_T)|\pi]$$

Portfolio value at target date (v_T)

Regular Payments

...

Remaining time to target date   target date
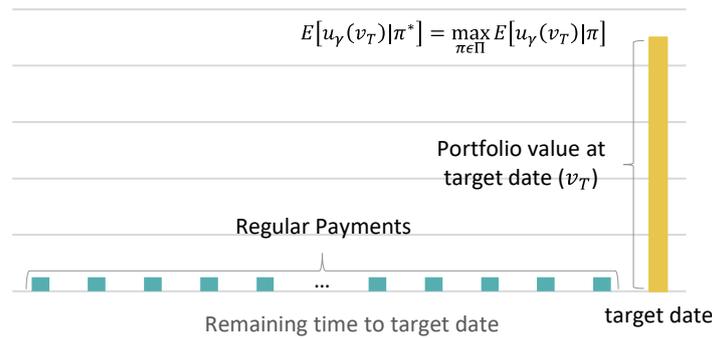
Source: DWS - For illustrative purposes only

Figure Three illustrates the savings plan problem: An optimized investment strategy π^* needs to be derived which maximizes the expected CRRA utility for all investment strategies $\pi \epsilon \Pi$ where Π is the set of all possible investment strategies for this problem.

We derive the optimized investment strategy π^* using dynamic programming, as outlined in the previous section, for different risk aversion parameters $\gamma = 2, 5, 8, 12, 20, 30$. We then compare these optimized investment strategies, i.e., the DP LC strategies, with different SB LC strategies described earlier. Please note that the risk aversion of investors increases with the risk aversion parameter γ.

Figure Four

Figure visualizes the DP LC strategy for an individual investor with a risk aversion parameter $\gamma = 2$. For the DP LC strategy, the portfolio allocation depends on two input parameters: the portfolio value, and the remaining time to the target date. Depending on these two parameters, the (optimized) DP LC strategies are depicted in the three asset classes: money market (MM), fixed income (Bonds), and equities (Equities). It starts with an initial savings plan value of EUR 500, 20 years before the retirement date (time (in years) = -20 in Four), with an allocation of 100% Equities.

---

[1] u_γ is the CRRA utility function with risk aversion parameter γ which is defined is $u_\gamma(v) = \frac{1}{1-\gamma}(v^{1-\gamma} - 1), \gamma > 0, \gamma \neq 1$ and $u_\gamma(v) = \log v, \gamma = 1$. An investor is risk averse if $\gamma > 0$. The risk aversion increases with increasing γ.
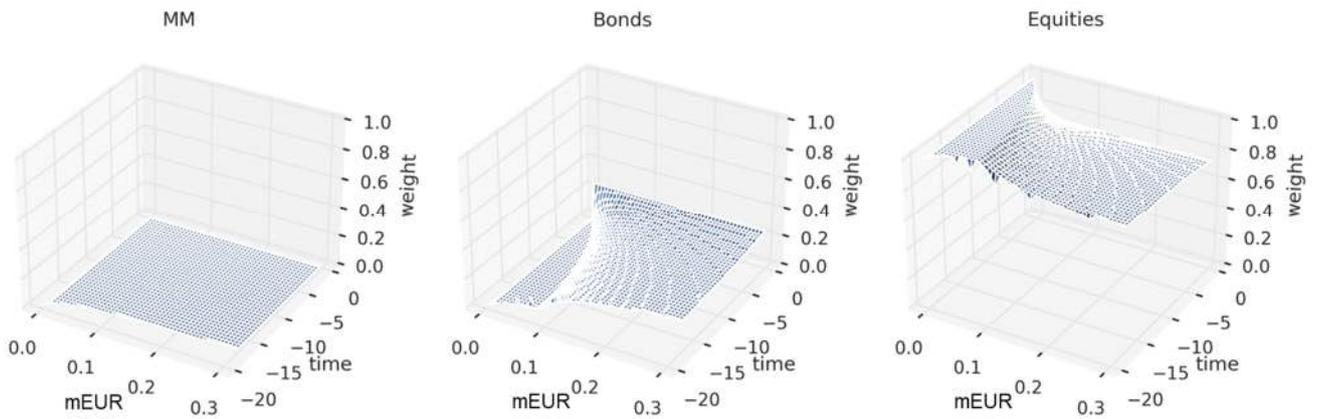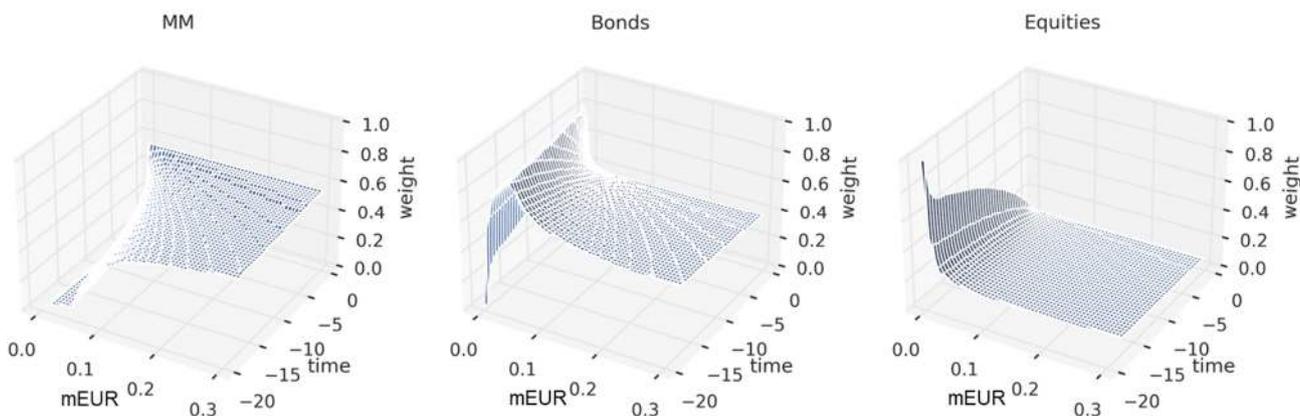
Source: DWS – For illustrative purposes only

Figure Four visualizes the DP LC strategies for risk aversion parameter $\gamma = 2$.

From Figure Four we can derive the following three observations:

1. For a fixed portfolio value, the portfolio allocation becomes more conservative over time.

2. For a fixed remaining time to the target date, the allocation becomes more conservative as the portfolio value increases.

3. At the last allocation decision time ($time = -0.25$), the allocation does not change with varying portfolio values.

The economic interpretation is as follows: at any allocation decision time t with portfolio value v, the DPLC strategy determines the optimal allocation by considering the present value of the future contribution stream, and the current portfolio value v. The present value of future contributions can be considered a risk-free asset ("Contribution Asset") because the future contribution stream is deterministic under the modelling assumptions.

This implies that the total portfolio, consisting of the savings plan portfolio plus the Contribution Asset, has a very conservative allocation at the beginning of the savings plan because we have only EUR 1000 in the savings plan portfolio and the present value of EUR 79,000 ($= 79 \cdot 1000$) in the Contribution Asset (treated as risk-free). In contrast, at the last allocation decision date, after all contributions are made, the total portfolio is entirely invested in the savings plan portfolio, since the Contribution Asset is zero.
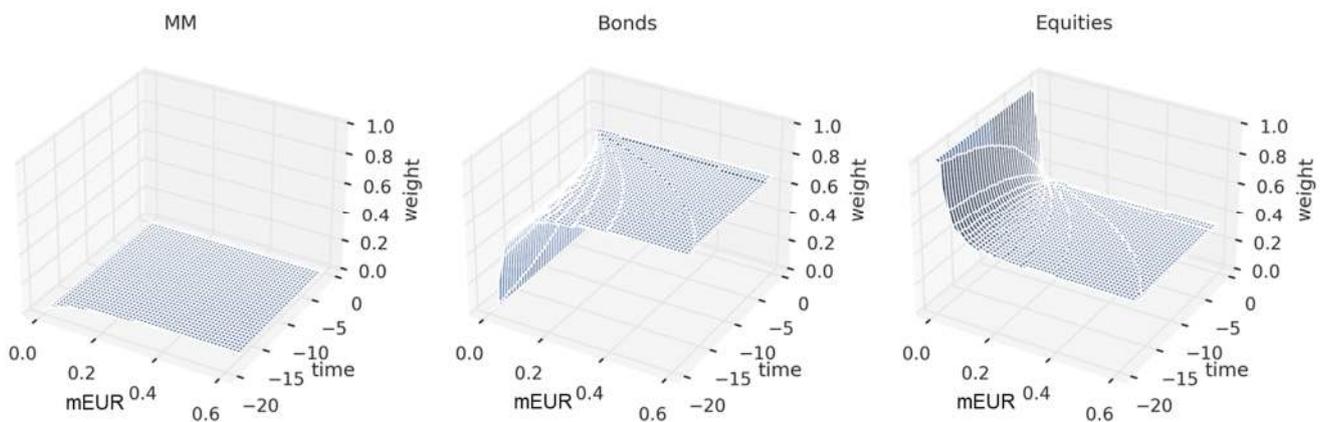
Source: DWS – For illustrative purposes only

Figure Five visualizes the DP LC strategies for risk aversion parameter $\gamma = 5$.

It is well-known in the financial literature that the final optimal allocation ("FOA") for a CRRA-investor without future contributions does not depend on the portfolio value, which we observe according to the third item above – see (Samuelson, 1969). At each allocation decision time with future contributions, the DP LC strategy tries to bring the risk-return profile of the total portfolio allocation as close as possible to the FOA by having a higher risk-return profile of the DP LC strategy compared to the FOA.

In conclusion, if $v \ll Contribution\ asset$, which is in the first phase of the savings plan, the DP LC strategy has a high risk-return profile, and if $v \gg Contribution\ Asset$, which is in the last phase of the savings plan, the DP LC strategy is close to the FOA. Please note, we would get the same DP LC strategy π^* for any contribution amount, so instead of contributing EUR 1000, we could also contribute EUR c with $c > 0$.
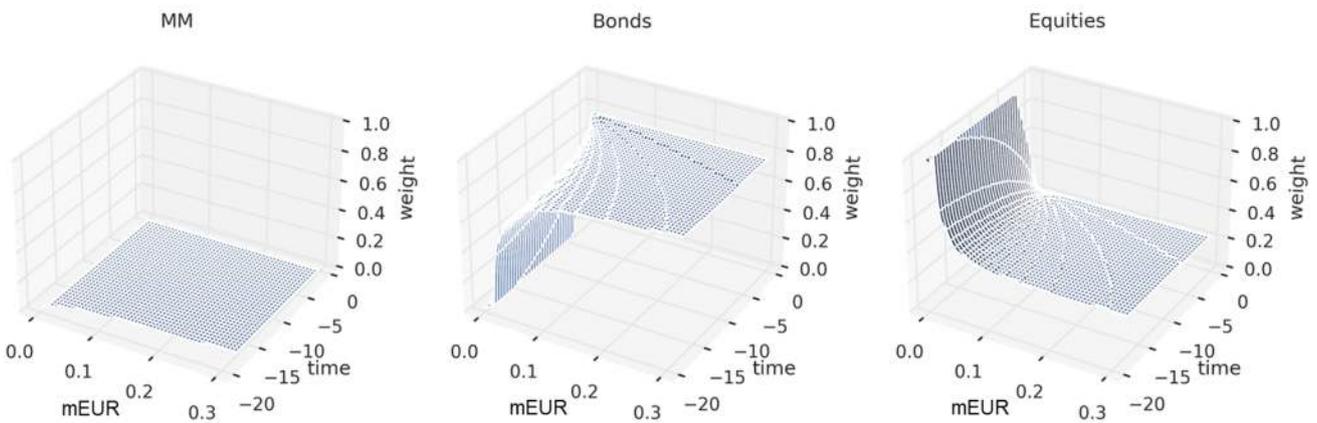
Source: DWS – For illustrative purposes only

Figure Six visualizes the DP LC strategies for risk aversion parameter $\gamma = 8$.

Figures Five to Nine visualize the DP LC strategies for risk aversion parameters $\gamma = 5, 8, 12, 20, 30$. These figures show the same properties as Figure Four, only the allocations have in general a lower risk return profile and the same economic interpretation for these properties holds. Additionally, please note for risk aversion parameters $\gamma = 5, 8, 12$ the DP LC strategy does not invest into MM.
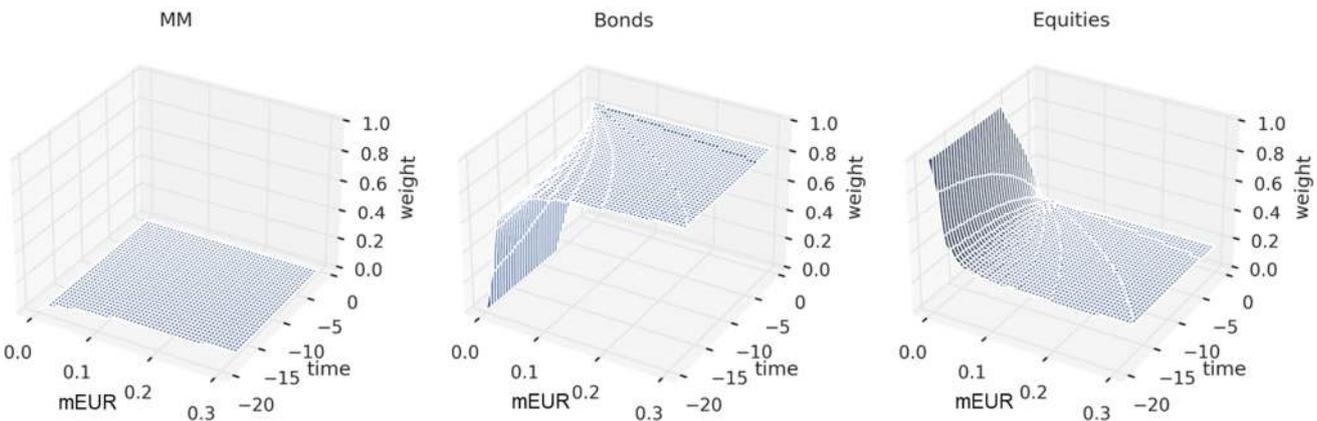
Source: DWS – For illustrative purposes only

Figure Seven visualizes the DP-LC-strategies for risk aversion parameter $\gamma = 12$.
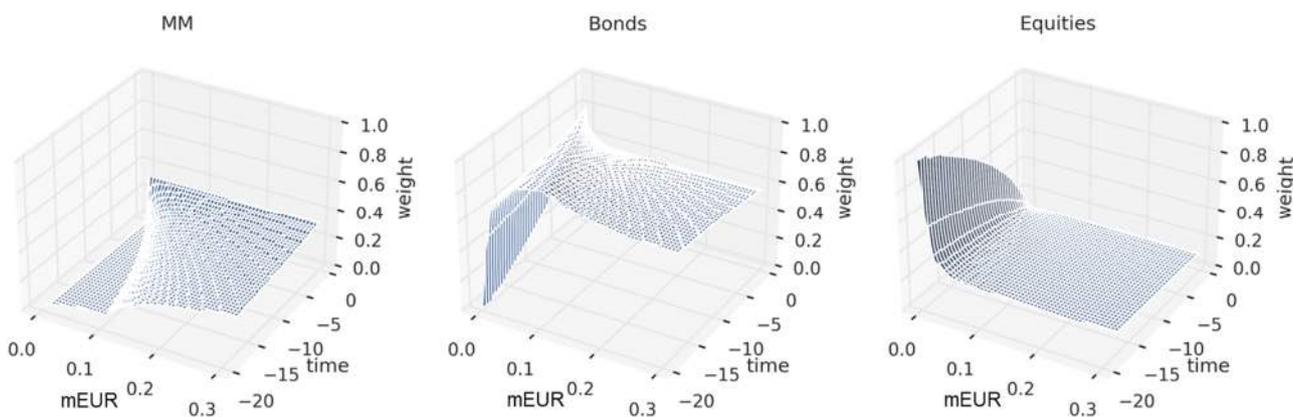
Source: DWS – For illustrative purposes only

Figure Eight visualizes the DP LC strategies for risk aversion parameter $\gamma = 20$.

Source: DWS – For illustrative purposes only

Figure Nine visualizes the DP LC strategies for risk aversion parameter $\gamma = 30$.

Table Three depicts the different FOAs. With increasing CRRA risk aversion parameters, the risk-return profile of the FOA decreases. In contrast, at the beginning of the savings plan, after the first contribution is made, the initial optimized allocation is 100% equities (EQ) for $\gamma = 2, 4, 8, 12, 20$, only for $\gamma = 30$ do we have an initial allocation of 4% fixed income (FI) and 96% equities.

FOA	CRRA risk aversion parameter (γ)					
	2	5	8	12	20	30
MM	0.0%	0.0%	0.0%	0.0%	32.7%	55.4%
FI	25.6%	67.1%	76.7%	83.1%	56.4%	37.2%
EQ	74.4%	32.9%	23.3%	16.9%	10.8%	7.4%

Source: DWS – For illustrative purposes only

Table Three shows the final optimized allocation (FOA) for different CRRA risk aversion parameters.

Comparison of DP and SB approaches with Monte-Carlo Simulations

In this section, we compare the DP and SB approaches using Monte-Carlo Simulations. For these simulations, we generated 10,000 capital market scenarios over the next 20 years, applying the capital market assumptions described in One. Each path describes the evolution of the three asset classes – Money Market, Fixed Income, and Equity. We used a multivariate geometric Brownian motion, parametrized according to One, to generate these paths.

The individual investor contributes a fixed quarterly amount, unchanged over the next 20 years. At each contribution date, after the contribution has been invested, the allocation is rebalanced to the target allocation of the LC strategy π in question. We consider the DP LC strategies π_γ^* for the CRRA-risk aversion parameters $\gamma = 2, 5, 8, 12, 20, 30$, and the three SB LC strategies $\pi_i^{LC}, i = 1, \ldots, 3$ as depicted in Figure Two.

At the target date, and for each path, we calculate the CRRA-utility $u_\gamma(v_{t_n}^i), i = 1, \ldots 10k$, based on the final savings plan value and determine the estimated expected CRRA utility $\hat{E}[u_\gamma(v_T)|\pi] = {}^1/_{10k} \sum_{1=1}^{10k} u_\gamma(v_{t_n}^i(\pi))$, which serves as an estimator for $E[u_\gamma(v_T)|\pi]$. To ensure numerical stability, especially for $\pi_\gamma^*, \gamma = 20, 30$, we consider the quarterly contributions to be 0.001 EUR, which does not affect the optimized DP LC strategies since scaling does not impact the optimization process[2].

Table Four shows the estimated expected CRRA utility derived from the Monte-Carlo Simulations. The simulations suggest that for a given risk aversion level γ, the estimated utility of DP LC strategies is higher than that of the SB LC strategies. However, these numbers are difficult to interpret directly. The standard approach is to calculate the certainty equivalent of the estimated expected utilities.

---

[2] In theory we could use also EUR 1000 as quarterly contributions, or any other contributions amounts. In particular large final savings plan values at target dates lead to very small utility amounts, when working with $\gamma = 20, 30$.

Estimated Expected CRRA Utility		DP-LC-strategies	Select and benchmark LC-strategies based on glide path		
			1	2	3
Risk aversion parameter (γ)	2	-6.36	-6.87	-6.69	-6.58
	5	-1126.55	-1162.45	-1139.97	-1174.89
	8	-483381	-510448	-550754	-650628
	12	-2.5455E+09	-3.2214E+09	-4.3530E+09	-6.7006E+09
	20	-1.6546E+17	-4.3473E+17	-1.0662E+18	-3.0490E+18
	30	-1.8198E+27	-2.1012E+28	-1.1392E+29	-7.1887E+29

Source: DWS - For illustrative purposes only

Table Four shows the estimated expected CRRA utility derived from the Monte-Carlo Simulation.

Table Five describes the estimated certainty equivalents[3]. In the context of the savings plan problem, the economic interpretation of the certainty equivalent is as follows: an individual investor with utility function u is indifferent to receiving the portfolio value v_T^{ce} with certainty or an uncertain portfolio value with expected utility $E[u(v_T)]$, when investing into the plan with strategy π.

Table Five again shows that for a given CRRA risk aversion parameter γ, the DP LC strategies tend to have a higher certainty equivalent than the SB LC strategies. For $\gamma = 2, 5, 8$, the difference between the certainty equivalents of the DP LC strategies and the SB LC strategies is small.

Estimated Certainty Equivalents		DP-LC-strategies	Select and benchmark LC-strategies based on glide path		
			1	2	3
Risk aversion parameter (γ)	2	0.1359	0.1271	0.1301	0.1319
	5	0.1220	0.1211	0.1217	0.1208
	8	0.1167	0.1158	0.1146	0.1119
	12	0.1123	0.1099	0.1069	0.1028
	20	0.1063	0.1010	0.0964	0.0912
	30	0.1022	0.0940	0.0886	0.0832

Source: DWS - For illustrative purposes only

Table Five describes the estimated certainty equivalents.

To better understand why the certainty equivalents are close, we analyze these strategies further. Figure Ten depicts the mean allocation of DP LC strategies over time for CRRA risk aversion parameters $\gamma = 5, 8$ and the allocation for the SB LC strategies with glide path Two and One. As observed, the two strategy pairs ($\gamma = 5$, glide path Two) and ($\gamma = 8$, glide path One) have different shapes, but result in quite similar estimated certainty equivalents $(0.1220, 0.1217)$ and $(0.1167, 0.1158)$, respectively. Table Six describes the mean and volatility of the final savings plan values, and the internal rate of returns (IRRs) derived from the 10,000 paths of the Monte-Carlo Simulations.
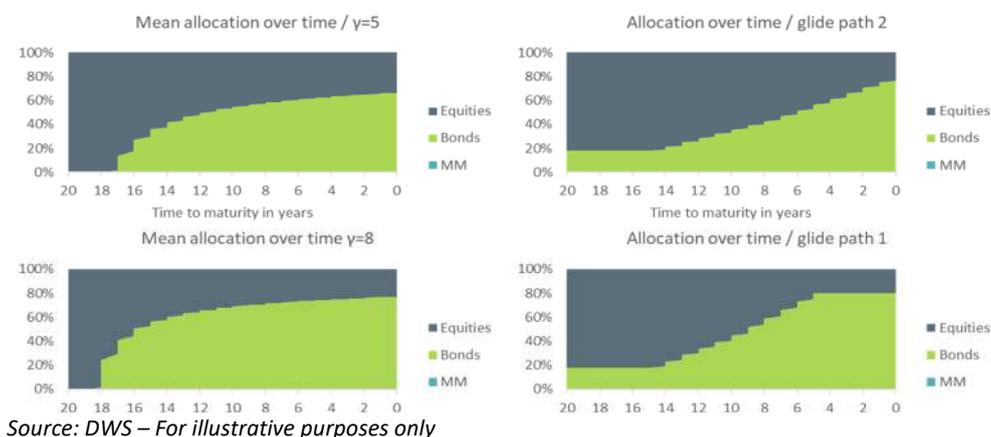
Source: DWS – For illustrative purposes only

Figure Ten depicts the mean allocation for the DP-LC-strategies with CRRA risk aversion parameter $\gamma = 5, 9$ and the allocation of the select and benchmark LC-strategies with glide path 2 and 1.

When comparing the two strategy pairs ($\gamma = 5$, glide path Two) and ($\gamma = 8$, glide path One), the mean and volatilities implied by these strategy pairs are quite different, although the certainty equivalents are very similar. This suggests that there is an investment strategy that optimizes the expected CRRA utility of the individual investor, namely the DP LC strategies, but that there are many investment strategies with different risk-return profiles that have quite similar certainty equivalents.

The reason for this is that expected CRRA-utility increases with a higher expected final savings plan value, all else being equal, and decreases with high volatility of the final savings plan values. This trade-off allows for many good solutions in terms of the certainty equivalent, which are all close to the optimal solution. In other words, one can deviate from the optimal solution in the risk-return space without a significant loss in CRRA utility.

If we assume that many individual investors are at least implicitly CRRA-utility maximizers, then there is a variety of investment strategies with different risk-return profiles that are suitable for these investors. This might explain why we see many investment strategies with different risk-return profiles from different product manufacturers targeting similar investor profiles. The CRRA-utility function offers a wide field of "good" solutions.

Finally, this application of dynamic programming demonstrates that dynamic programming can be applied to solve multi-period optimization problems. Furthermore, it shows that LC strategies derived under the SB approach still yield satisfactory results under the CRRA-framework.

Distribution ana- lyses		DP-LC-strategies						Select and benchmark LC- strategies based on glide path		
		2	5	8	12	20	30	1	2	3
v_T	Mean	0.1548	0.1331	0.1273	0.1240	0.1176	0.1105	0.1316	0.1366	0.1407
	Vol.	0.0588	0.0254	0.0191	0.0163	0.0120	0.0081	0.0252	0.0311	0.0375
IRR	Mean	5.51%	4.60%	4.26%	4.05%	3.60%	3.05%	4.49%	4.76%	4.94%
	Vol.	3.09%	1.69%	1.35%	1.19%	0.93%	0.68%	1.68%	1.97%	2.26%

Source: DWS - For illustrative purposes only

Table Six depicts the mean and volatility of the final savings plan values and IRRs derived from the 10k paths of the Monte-Carlo Simulations.

3 / Conclusion and Outlook

Both the SB and the DP approach can be used to derive LC strategies for savings plans. In the SB approach, a significant amount of expert judgment is applied when designing the model portfolio and the different glide paths. Conversely, in the DP approach, once the objective function reflects the goals and objectives of the individual investor, the entire DP LC strategy is derived through the dynamic programming optimization process.

With advances in digitalization, increased access to individualization platforms, and greater computational power, we anticipate that goal-based approaches for savings and decumulation plans, solved by dynamic programming or similar methods, will become an additional option for product providers in the future. The CRRA-utility is certainly one potential objective function to be used. However, as discussed in the previous sections, the CRRA-utility can be challenging for individual investors and practitioners to interpret within the risk-return space, where risk is measured as volatility and return is interpreted as the expected return.

Market participants tend to think in terms of expected target returns, volatilities, and quantiles of the final wealth distribution of a savings plan. The CRRA-utility function does not directly align with these notions. Looking ahead, one potential route is to apply dynamic programming to solve multi-period savings and decumulation plan problems but using objective functions that more directly translate into the individual investor's goals and objectives, as described above.

Bibliography

Campbell, J. Y., & Viceira, L. M. (2002). *Strategic Asset Allocation for Long-Term Investors.* Oxford University Press.

Das, S., Ostrov, D., Radhakrishnan, A., & al., e. (2019). Dynamic portfolio allocation in goals-based wealth management.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics, 51(3)*, 239-246.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction.* MIT Press.

Viceira, L. (2008). *Life-Cycle Funds. In Overcoming the Saving Slump: How to Increase the Effectiveness of Financial Education and Saving Programs.* University of Chicago Press.

Glossary

CRRA Utility (Constant Relative Risk Aversion) – A mathematical model that describes how much risk an investor is willing to take – used to tailor investment strategies to individual preferences.

Capital Market Assumptions – Estimates about future returns, risks, and relationships between asset classes – used to guide investment strategy design.

Certainty Equivalent – The amount of money an investor would prefer to receive for sure instead of taking a risky investment with uncertain outcomes. It helps compare strategies by showing how much risk an investor is willing to avoid for a similar level of satisfaction.

Contribution Asset – The total value of all future contributions an investor plans to make into their savings plan – treated as a safe and predictable part of the portfolio because it's not affected by market fluctuations.

Covariance Matrix – A table that shows how different investments move in relation to each other – used to manage and diversify risk in a portfolio.

Defined Contribution (DC) Plan – A retirement savings plan where the investor regularly contributes a fixed amount – the final value depends on how the investments perform over time, not on a guaranteed payout.

Dynamic Programming (DP) – A mathematical technique used to optimize investment decisions over time – it calculates the best strategy by considering future contributions, market changes, and investor goals at each stage.

Efficient Frontier – A curve that shows the best possible return for each level of risk – helping investors choose portfolios that offer the most value for their risk tolerance.

Final Optimized Allocation (FOA) – The ideal mix of investments at the end of a savings plan – based on the investor's risk tolerance and financial goals.

Glide Path – A schedule that gradually shifts investments from higher-risk to lower-risk portfolios – designed to reduce risk as the investor approaches retirement.

Individualized Asset Management (IAM) – Customized investment strategies tailored to each investor's financial goals, risk tolerance, and retirement timeline.

Internal Rate of Return (IRR) – The annualized return earned on an investment – accounting for all contributions and changes in value over time.

Life-Cycle (LC) Strategy – An investment approach that adjusts the risk level of a portfolio based on the investor's age – typically more aggressive when the investor is younger and has a longer time until retirement, and more conservative as they get older and approach retirement.

Mean-Variance Optimization – A method for building investment portfolios that aim to maximize expected return for a given level of risk – commonly used to identify efficient investment combinations.

Model Portfolios (Defensive, Balanced, Dynamic) – Predefined investment mixes with different levels of risk and return – Defensive portfolios are low-risk, Balanced are moderate-risk, and Dynamic are high-risk and return-seeking.

Monte Carlo Simulation – The random sampling of thousands of possible paths from a specified statistical distribution used to recreate, and garner insights from, multiple future states of the world.

Objective function – The end goal of a process that is trying to be solved for or optimized.

Select and Benchmark (SB) Approach – A method used to build investment strategies by selecting portfolios and comparing them to investor goals – it relies on expert judgment and simulations to find a suitable match for the investor's needs.

Utility function – A mathematical attempt to quantify a person's preferences formulaically.

Appendix – Intro to Dynamic Programming

Dynamic programming, developed by Richard Bellman in the 1950s, is a method used to solve problems involving sequential decision-making in fields such as operational research, economics, and finance. For an introduction to DP, refer to (Sutton & Barto, 2018).

In finance, the multi-period savings plan problem can be framed as a sequential decision-making problem under uncertainty. The savings phase of an individual investor, spanning a target period (e.g., 20 years), is divided into single investment periods (e.g., 3 months). At the beginning of each investment period, a new allocation is determined based on the current value of the savings plan, the envisaged future contributions, and the remaining time to the target date. This process aims to achieve an optimal outcome at the target date, aligned with the investor's goals and objectives, which are formulated as an objective function depending on the value of the savings plan at the target date.

These problems are typically challenging to solve for several reasons:

1. An allocation decision at the beginning of one period impacts all subsequent periods.

2. Each allocation decision is influenced by all preceding allocation decisions.

3. The stochastic development of the savings plan value precludes an isolated path-wise assessment of an allocation decision. Instead, decisions must be evaluated based on all potential future savings plan values and allocation decisions.

Dynamic programming addresses these challenges by determining the optimal investment strategy for the last investment period (a single-period investment) and then recursively solving the investment problem for all preceding periods by working backward in time. In the following paragraph, we explain this DP approach in more detail for the savings plan problem.

Dynamic programming applied to the savings plan problem

In the following, we describe how the optimized DP LC strategy is derived by applying dynamic programming. This approach is detailed in (Das, Ostrov, Radhakrishnan, & al., 2019).

First, the problem is discretized with respect to all future allocation decision times $t_i, i = 0, \dots, n-1$, plus the target date t_n, which is not an allocation decision date, and all savings plan values. We refer to $(t_i)_{i=1,\dots,n}$ as the time grid. All future savings plan values are represented by a sufficiently large, discretized savings plan value grid $(v(t_i, j))_{i=0,\dots,n,j=0,\dots m_i}$. We have a list of $m_i + 1$ savings plan values for grid time t_i, meaning the number of grid values depends on t_i. Typically, $m_i \le m_j$ for $i < j$ because we start with a single value at t_0 and the potential range of savings plan values increases over time, necessitating a sufficient grid density.

The allocation decision space is reduced and discretized as follows: instead of considering all allocations in the asset class universe under predetermined investment constraints, we first reduce this portfolio space to the portfolios on the efficient frontier based on the capital market assumptions of the universe and the constraints. We then discretize the efficient frontier by selecting a grid of portfolio allocations $(a_k)_{k=1,\dots K}$ on this efficient frontier, consisting of, for example, $K = 10$ to 160 efficient portfolios ranging from a low to a high risk-return profile.

This savings plan problem is then solved backwards in time: the grid values on the target date $(v(t_n, j))_{j=0,\dots m_n}$ are evaluated with the objective function u, leading to $(u(v(t_n, j)))_{j=0,\dots m_n}$. At the beginning of the last single investment period, a portfolio is selected from the efficient portfolio allocation grid, and the transition probabilities $p(v(t_n, j)|v(t_{n-1}, l), a_k)$, from each wealth grid point in t_{n-1} to each wealth grid point in t_n calculated considering the inflows from t_{n-1} to t_n[4]. For each wealth grid point $v(t_{n-1}, l), l = 0, \dots, m_{n-1}$ the so-called Q-function $Q(v(t_{n-1}, l), a_k)$ is calculated according to

$$Q(v(t_{n-1}, l), a_k) = \sum_{j=0}^{m_n} u(v(t_n, j)) \cdot p(v(t_n, j)|v(t_{n-1}, l), a_k).$$

Equation One is the Bellman equation applied to the savings plan problem.

Equation One is called the Bellman equation. Based on the Q-function the optimal portfolio allocation and the optimal value function V^* is calculated in t_{n-1}

$$a^*(t_{n-1}, v(t_{n-1}, l)) = \underset{k=1,\ldots,K}{\operatorname{argmax}} Q(v(t_{n-1}, l), a_k) \text{ and}$$

$$V^*(t_{n-1}, v(t_{n-1}, l)) = Q\left(v(t_{n-1}, l), a^*(t_{n-1}, v(t_{n-1}, l))\right).$$

The same steps are performed to go from t_i to t_{i-1} and to derive $a^*(t_{i-1}, v(t_{i-1}, l))$ for all $i = 0, \ldots, n-1$ and $l = 0, \ldots, m_i$. Please note that in Equation One, instead of $u((t_n, j))$ $V^*(t_i, v(t_i, l))$ is used. Based on this procedure the optimized portfolio allocations $(a^*(t_i, v(t_i, l)))_{i=0,\ldots,n, j=0,\ldots m_i}$ are defined on the wealth grid $(v(t_i, j))_{i=0,\ldots,n, j=0,\ldots m_i}$. For an arbitrary savings plan value v on an allocation decision date t_i DP LC strategy is then defined as portfolio allocation a^* in t_i whose grid value is close to v.