A Plan-Based Glide Path

Tom Breaden, CFA

Introduction

SUMMARY

- Target date retirement fund analysis can be enhanced by comparing participants' return requirements to a glide path's performance potential.
- Doing so requires computation and data resources but can help reassure fiduciaries that their funds reasonably match their Plan's demographics to follow regulatory guidance.

Since their establishment as Qualified Default Investment Alternatives (QDIAs) under the Pension Protection Act of 2006, target date retirement funds (TDFs)[1] are now featured in many defined contribution plans such as 401(k) and governmental 457 plans. In response to this increased popularity, the Department of Labor (DoL) issued a bulletin in February 2013, "Target Date Retirement Funds – Tips for ERISA Plan Fiduciaries," specifying areas for plan fiduciaries to consider in selecting TDFs. One section of this bulletin ("Establish a process for comparing and selecting TDFs") suggests that fiduciaries "consider how well the TDF's characteristics align with eligible employees' ages and likely retirement dates" and other characteristics such as pension plan participation, salary levels, cash flow patterns, and others. The DoL did not mandate or prohibit any particular approach; qualitative, quantitative, or otherwise, to vet these considerations; leaving plan fiduciaries with both latitude to perform their analyses and uncertainty as to their appropriateness. The investment industry also has yet to converge on a set of generally accepted glide path selection procedures reflecting the DoL's guidance. To bridge the gap between regulatory guidance and industry practice, this paper proposes a quantitative approach that, while technical, makes efficient use of plan and investment performance data to provide fiduciaries with a process to identify TDFs that may be appropriate for their plan.

Step One: Finding a Target Rate of Return

This paper draws on the simple, fundamental belief that participants rely on investment returns to achieve their retirement savings goals. To that effect, there is a target return that, if achieved, can produce an adequate retirement income for each participant. Fiduciaries who are evaluating TDFs may wish to limit their evaluations to TDFs with reasonably high probabilities of achieving these target returns and with minimal shortfall risk. The "Internal Rate of Return" (IRR) metric, which equates expected savings with expected expenditures, is a way to calculate such a target return. The graphic on this page provides a simple IRR representation, its general functionality being that as

Saving for Retirement | **Retirement Spending**

$$\frac{\$\ Saved_n}{(1+IRR)^n} + \frac{\$\ Saved_{n+1}}{(1+IRR)^{n+1}} = \frac{\$\ Spent_{n+2}}{(1+IRR)^{n+2}} + \frac{\$\ Spent_{n+3}}{(1+IRR)^{n+3}}$$

[1] TDFs are portfolios that are primarily composed of fixed income and equity investments and whose name denotes a "target" retirement year (for example a Target Date 2050 fund would be designed for an investor who expects to retire in 2050). The portfolio gradually assumes a more conservative allocation as that target year approaches to remain appropriate for the investor's time horizon. The rate at which the asset allocation becomes increasingly conservative is commonly referred to as a "glide path".

expected savings (or spending) increase, the IRR will decrease (or increase).

In this paper, calculating a per-participant IRR requires fairly basic plan demographic data. The "Plan Demographics Data" table shows each piece of participant data that is required from the recordkeeper to calculate an IRR and how the IRR is impacted as a given data point increases, all else held equal. In addition to Plan Demographic Data, the user will need to make reasonable assumptions about the future. The assumptions and their IRR impact are stated in the "Plan-Based Assumptions" table. These assumptions should reasonably reflect plan-specific features. For example, for a governmental public safety retirement plan, the user might assume an earlier retirement age and a greater percentage of retirement income coming from a pension. Another example is a growing private sector company with a younger workforce whose retirement income will primarily come from a safe harbor 401(k) plan and Social Security. These different plans may have unique return needs and correspondingly suitable TDFs.

The full IRR calculation methodology is provided in the Appendix. To be conservative and efficient, this paper assumes no additional outside wealth or liability. Generally, the extent that outside wealth is understated will place upward bias on the IRR and glide path selection. The next section of this paper will suggest, however, that the estimated IRRs have been obtainable over financial markets' history. In summary, the IRR calculation steps outlined above may provide plan sponsors with a way to efficiently incorporate the data points suggested by the DoL in determining a suitable glide path.

Age	Expected Retirement Date	Plan Participation	Salary	Cash Flow Patterns
✓	✓	✓	✓	✓

Plan Demographics Data

Variable	IRR Impact if Var. ↑
Account Balance	↓
Age	↑
Gender	↑/↓ F/M
Net Cash Flows	↓
Salary	↑
Date of Hire	↑

Plan-Based Assumptions

Variable	IRR Impact if Var. ↑
Salary Growth Rate	↓
Retirement Age	↓
Income % Replacement Ratio	↑
Pension + Social Sec. Income	↓
Inflation in Retirement	↑
Life Expectancy	↑

Once each participant's IRR is calculated, the IRRs can be aggregated across various demographic lines. This paper's approach is to divide participants into four groups: those with less than ten years to retirement, between ten and twenty, between twenty and thirty, and between thirty and forty years to retirement and; then calculate the IRR for each group[2].

[2] Participant IRRS can be divided into age groups in order to perform some amount of analytical drill-down, which may uncover demographic anomalies that may warrant consideration.

A Sample Exercise

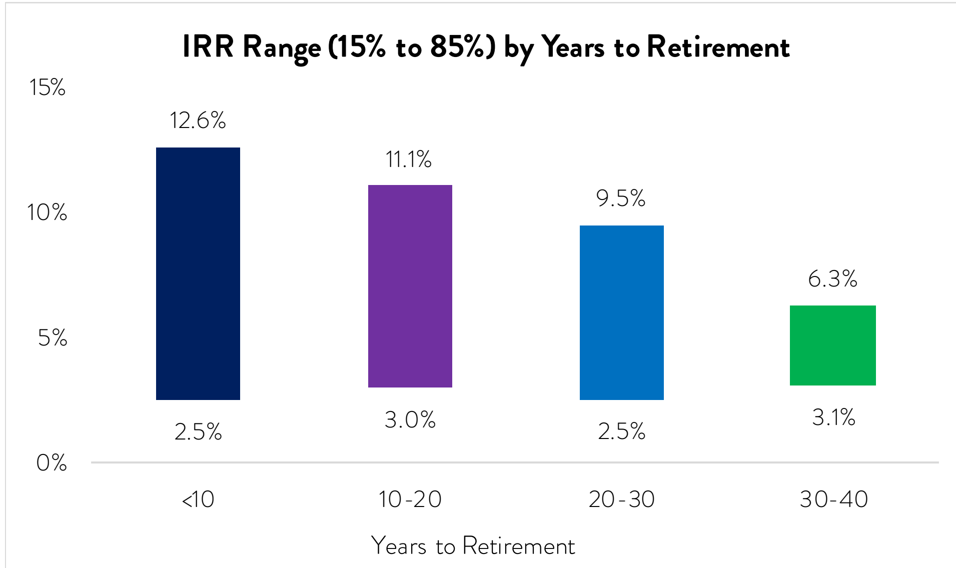

The "IRR Range by Years to Retirement" chart shows the 15th to 85th percentile range of IRRs of individual participants who are assumed to retire at age 65, live until the ages of 86 and 91 for men and women respectively, and seek to replace 80% of their inflation-adjusted final working income from a combination of Social Security income and 401(k) balances. The IRR calculations also incorporate adjustments for assumed prior working and savings history. In this example, the median IRR per age bracket declines from 8.1% for those with less than ten years of remaining work to 4.7% for those with thirty to forty years of remaining career time, illustrating how divergent retirement preparation may become over time. The reader may also note how the IRR range narrows for those with longer time horizons, interestingly showing the value of long-term compounding and maintaining savings discipline, as well as the tendency for wages (and a higher amount of a worker's replacement ratio not coming from Social Security) to be more disparate among older workers. The "Long-Term Returns: 1926-2016" graphic shows the annualized performance that various asset classes have produced, indicating that the IRRs in the preceding chart have generally been obtainable for long-term investors.

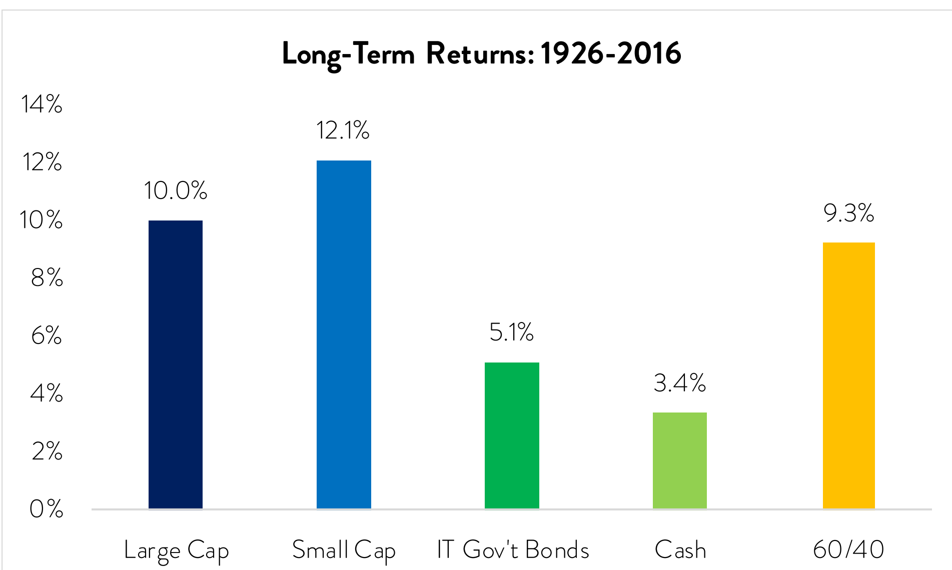

Step Two: Determine the Optimal Glide Path

With the target rates of return being established for each age group in the preceding exercise, the next step is to identify a glide path that can be expected to meet that return in a risk-efficient manner. The "Optimal Glide Path" may have a relatively high or low allocation to stocks for dates far from or close to retirement and follow a steep or moderate, linear or curved, path towards the retirement date's stock allocation. By calculating historical investment returns of various types of glide paths in an iterative manner, the user can determine which particular glide path might work best for their plan. For example, assume that the 30-40 years to retirement age bracket has a target IRR of 7%. What glide path is best for that age bracket? This is tested by calculating the annualized returns of a range of glide paths with varying slope, aggressiveness (equity exposure), and convexity for each rolling sixty-six year period (representing forty years out from retirement to twenty-five years passed retirement) for as far back as we possess historical investment returns (1926). This process can be repeated for each age bracket using the same glide paths and historical returns. The "Glide Path" chart shows some of the data that can be obtained from this exercise using only two sample glide paths. Additional descriptive statistics such as standard deviation, skewness and minimum annualized return per glide path can also be calculated.

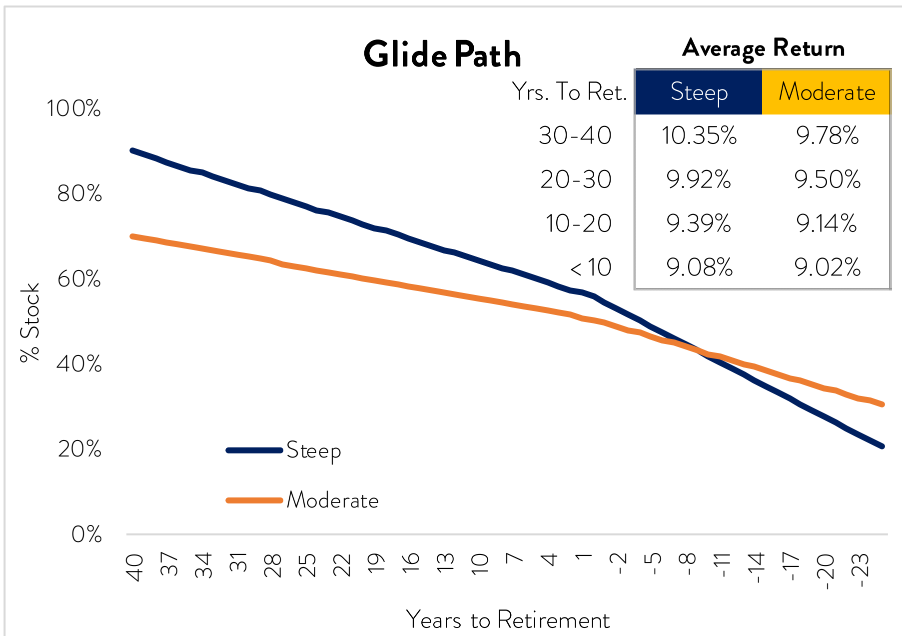

Using allocations to the distinct glide paths mentioned above, a single Optimal Glide Path can be generated for the plan. For example, the Optimal Glide Path may be weighted 50%, 40%, and 10% to steep, moderate, and convex glide paths if this one is estimated to be the best fit for the plan[3]. This paper defines the Optimal Glide Path as one that optimizes the following:

1. A rate of return that consistently exceeds a conservative estimate of each age group's target IRR;
2. A minimal standard deviation of returns; and
3. A maximum (least negative) skewness of returns for participants with less than ten years to retirement[4].

[3] Distinct glide paths are used as a starting point rather than simply using a solver to generate the stock weightings. This is in order to ensure that the output glide path generally resembles and can be compared to the variety of TDFs that are available for investment. Alternatively, a user that does not seek to use construct an optimal glide path out of distinct ones could constrain the solver to ensure that an industry-comparable glide path is produced.

[4] Skewness is a measure of outlier returns relative to average returns. Negative skewness indicates that more extreme outlier returns tend to be below the average return. This is of greatest importance to near-retirement participants who, due to having the highest balances and greatest likelihood of

The solver application in most spreadsheets can be used to construct an Optimal Glide Path with the highest possible, demographically-weighted, score. The graphics below provide an illustration of this. The “Statistic Weighting” in the table below puts the most emphasis on return (70%) due to its fundamental importance.

Statistic	Treatment of Statistic	Statistic Weighting (A)	Statistic per Age Group (B):*				Age Group as % of Plan (C)				Score
			<10 Yrs	10-20	20-30	Up to 40	<10 Yrs	10-20	20-30	Up to 40	ΣΠ A x B x C
Return	Reward	70%									
Standard Deviation	Penalize	20%									
Skewness	Penalize	10%									

**Skewness only applies to the <10 Years Age Group*

The Optimal Glide Path, built from other glide paths, has the highest score.

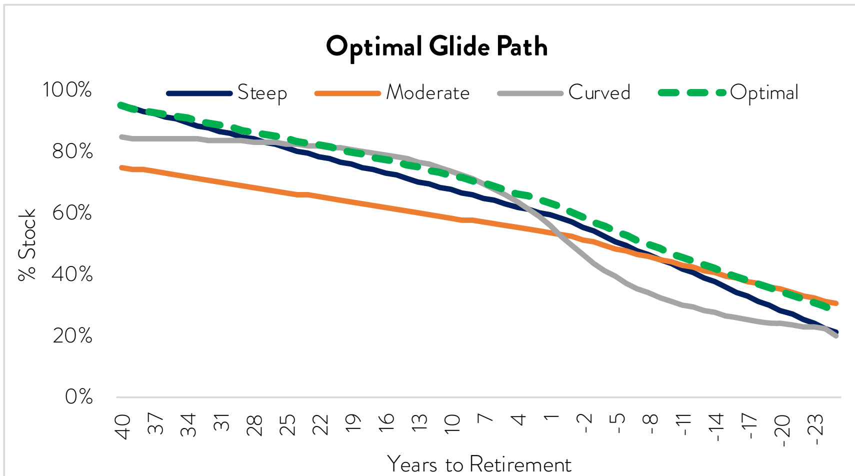

A Sample Exercise

The IRRs calculated earlier in this document are compared to the returns of a variety of glide paths to determine which glide path or combination thereof has produced the highest possible return over the IRR target while generating minimum volatility and skewness (for near-retirees only) in a manner proportional to the plan’s demographics. The Optimal Glide Path shown in the following graphic combines twelve distinct glide paths with varying degrees of slope, aggressiveness, and convexity.

making near-term withdrawals, are most financially sensitive to sequence of returns risk; the risk of making retirement distributions during a period of poor market returns.

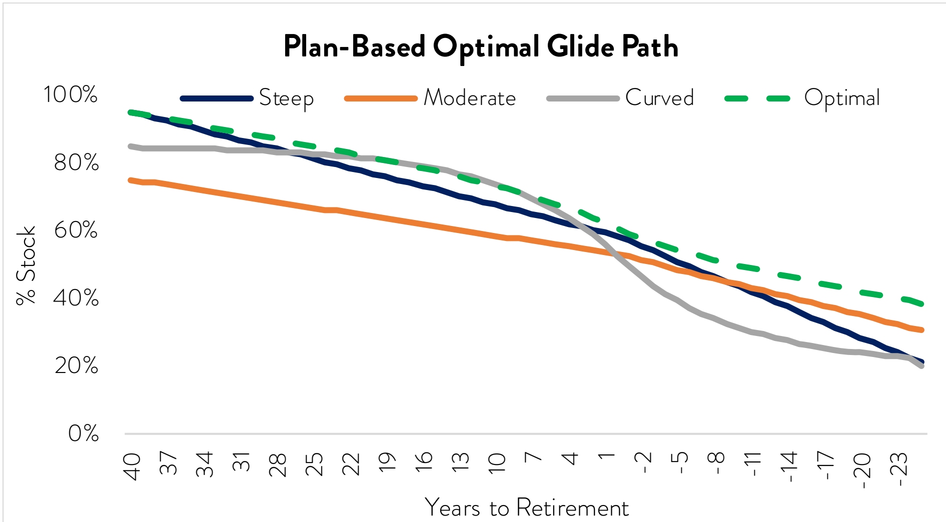

Step Three: Finding Acceptable Products

Once the Optimal Glide Path has been estimated, the next step will be to identify its manner of implementation. Larger plans may consider constructing a custom glide path, which may simply be the Optimal Glide Path, populated with Stock, Fixed Income, and Alternative investment funds of their choice. Instead, most employers will likely select a mutual fund or other pooled vehicle that is readily available. Plan fiduciaries selecting such a vehicle may wish to find one whose score is reasonably close to that of the Optimal Glide Path. Starting with a diverse selection of TDFs from various investment companies which satisfy the plan's other Investment Policy Statement (IPS) criteria (typically concerning expenses, style consistency, investment performance, management stability, and others), the user can then score each TDF relative to the Optimal Glide Path as well as the Category Average glide path. Each TDF should have a score that is comparable to or exceeds that of the Category Average glide path[5], implying that plan fiduciaries have identified a more appropriate glide path than a random draw (of otherwise IPS-appropriate TDFs) would have. The data below compares the stock exposure of the Optimal Glide Path to two fictitious TDFs (A and B) and the Category Average. In this case, TDF A's score was higher and stock exposure tended to be closer to that of the Optimal Glide Path than TDF B and the Category Average. Overall, fiduciaries may view both TDF A and B as plan-appropriate though A is a closer fit.

Difference in % Stock from Optimal Glide Path			
TDF / Metric	**A**	**B**	**Cat. Avg.**
Avg.	-5%	-13%	-10%
Avg. (Abs.)	5%	13%	10%
Range	9%	18%	19%
Score v. Optimal	-0.09%	-0.17%	-0.23%

5 By not being subject to user-constraints used to calculate the Optimal Glide Path, it is possible for a fund company's glide path to have a higher score than the Optimal Glide Path.

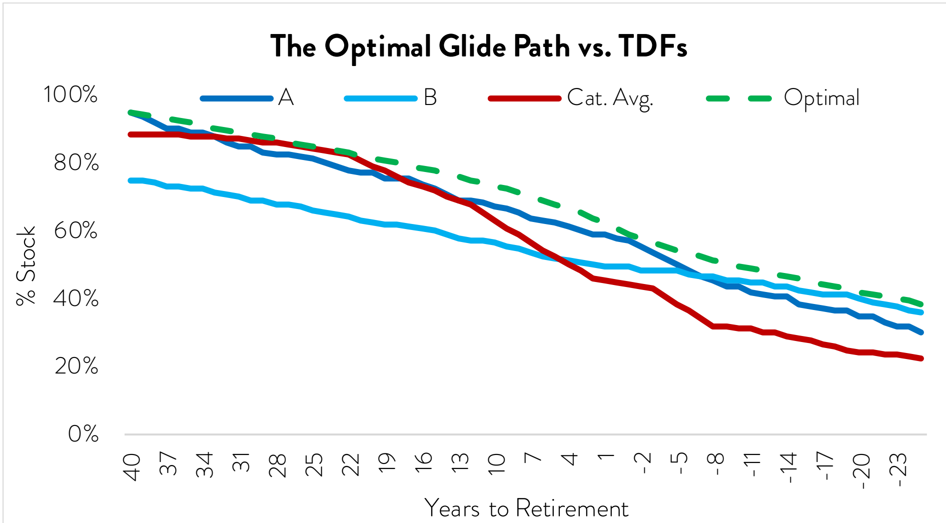

In addition to looking at raw scores produced by this approach, plan sponsors can consider other factors (such as the use of alternative investment strategies) that may give reason for exception from them. Once the appropriateness of a competitive variety of TDFs has been confirmed, plan fiduciaries can then compare them in a manager search document, with this analysis included in the front-end or appendix.

Conclusion

This paper provides plan fiduciaries and consultants with a workable, albeit technical, series of steps to help them follow the DoL's guidance in selecting appropriate TDFs. In addition to seeking to comport with regulatory guidance, this paper is intended to increase information in a fiduciary's selection process by incorporating retirement readiness directly into the glide path selection decision. This added information notwithstanding, it is important to note that although this paper uses optimization techniques, it is not intended to state that the Optimal Glide Path is in fact "perfect", or that TDFs should be selected solely based on their resemblance of the Optimal Glide Path. Rather, this paper is intended to help fiduciaries fulfill their duty, understand their plan(s) better, and avoid meaningful error in TDF selection.

Appendix: IRR Calculation Methodology (Data Adjustments Underlined)

1. Input participants' account balances, genders, and dates of hire into a spreadsheet.
   a. More than one plan: In the case of employers with more than one DC plan (such as a governmental agency offering a 457(b) and 401(a)), balance and contribution amounts can be aggregated.
   b. Prior work and savings history: Employees may have balances from prior employers' retirement plan(s). Omitting this information may cause their balances to appear inaccurately low and lead to excessively high IRRs (especially for near-retirement employees). To mitigate this, the user can assume savings rates for periods prior to current employment, based on each participant's date of hire and current savings levels. The current balance can then be augmented to reflect assumed prior savings and investment returns.

    c. No salary data available: If salary data is not available, the user may refer to the Bureau of Labor Statistic's website, which can be sorted by region and profession, to estimate it.
2. Assume they continue making their current contributions as a percentage of income (adjusted for expected salary growth) until they retire.
    a. Loan repayments: Contributions should include loan repayments. It is assumed that contributions continue as their current percentage of income after the loan is repaid.
    b. Omit rollovers and lump-sum withdrawals
    c. Looking forward: Drawing on the observation of participant atrophy, it is assumed that participants' savings behavior does not change over the remainder of their careers. This includes separated participants. Full vesting is also assumed.
3. Post retirement, assume they make inflation-adjusted withdrawals from their account that are adequate to achieve their "Income % Replacement Ratio" over their life expectancy after incorporating pension and Social Security income.
    a. Estimating pension income: Pension income is usually determined by a formula that multiplies an employee's years of service by a certain percentage (often between 1.5% to 3%) based on their average salary in their last few working years. While benefit determinations are usually substantially more complex in practice, a general expression of the pension formula can be estimated for its applicable employees. More simply, the user may input a conservative estimate about what percentage of retirement income will come from the pension (e.g. 35%). In cases where a pension has been closed to new employees as of a certain date, the estimation can be delineated based on participants' dates of hire.
    b. Estimating Social Security income: A simplified estimation of Social Security that incorporates maximum annual payments, adjustments for early retirement, and the applicable "bends" can help the user to determine how much a participant may expect from Social Security. In practice, exact Social Security estimation is highly complex. In this exercise, the added precision of a more complex estimation is unlikely to materially impact the IRR calculation, given the low portion of retirement income that may be attributable to Social Security. In the case of public sector employers, the user may wish to confirm their participation in Social Security.
4. Compute an IRR that equates the participants' balances plus expected contributions to their expected withdrawals from their account.