Stop Guessing
Using Participant Data to Select the Optimal QDIA

Morningstar Investment Management LLC
January 30, 2018

Thomas M. Idzorek, CFA
Chief Investment Officer - Retirement
thomas.idzorek@morningstar.com

David M. Blanchett, Ph.D., CFA
Head of Retirement Research
david.blanchett@morningstar.com

Daniel Bruns
Head of Large Plan Strategy/Solutions
dan.bruns@morningstar.com

Executive Summary
To receive safe harbor regulatory protections, defined contribution plan sponsors must, among other things, follow a prudent process for selecting a qualified default investment alternative. This includes considering the specific demographics of the participants when a plan sponsor is making its QDIA selection - balance/target-risk, target-date (including custom), retirement managed accounts, or a hybrid-QDIA. Unfortunately, the regulations do not provide much of a blueprint on how plan sponsors could and should consider plan demographics within a prudent process. Recent technological advancements, along with increased data availability and quality, enable plan sponsors to stop guessing and move beyond heuristics to use detailed data on individual participants to make a robust, data-driven selection of an appropriate QDIA. This paper presents a framework for using participant data to help determine which type of QDIA is most appropriate for a plan, including which glide path best fits the demographics of the plan participants, when a custom glide path makes sense, what the "pivot" age should be for a hybrid QDIA (target-date funds for younger investors and managed accounts for older investors), and when is retirement managed account necessary to meet the unique demographics of a plan's participants.

Key Takeaways
Plan fiduciaries need to follow a prudent process when selecting a QDIA, and in some circumstances, are required to consider plan demographics when selecting a risk-based QDIA.

To date, there has not been a viable, quantitative method for expressly considering plan demographics. Recent advances in technology coupled with availability of participant data have made it possible to explicitly and individually consider the demographics of plan participants.

We have established a framework for quantifying the cost/benefit of a better-fit QDIA.
As a result, plan sponsors can stop guessing and use participant demographic data to inform:

- Which of the three QDIA options, plus a hybrid combination option, fits the plan best;
- The participant age or other factors that should be used to structure the hybrid QDIA;
- Which glide path fits the plan best; and,
- Whether a custom target-date glide path necessary.

Introduction

Starting in late 2007, the U.S. Department of Labor (DOL) provided plan sponsors conditional fiduciary safe harbor when selecting a default investment option from three types of vehicles: 1) target-risk or balanced funds; 2) target-date funds; or 3) retirement managed accounts. To receive fiduciary liability relief, plan sponsors must follow all ERISA requirements related to the selection of a QDIA (i.e., the plan sponsor must follow a prudent selection process). In certain circumstances this includes specifically considering the demographics of the plan participants when selecting a QDIA.

Unfortunately, the DOL regulation did not provide much guidance on how a plan sponsor could and should consider plan demographics. Furthermore, prior to recent advances, there has not been a practical method beyond grossly simplified heuristics to fulfil this requirement, unless the plan engaged a sophisticated consultant.

A 2015 study by CAPTRUST Advisors found that 45% of plan sponsors selected the proprietary target-date fund series offered by their respective recordkeeper. This suggests to us that many plan sponsors may not be fulfilling their fiduciary duty on this decision. It is simply too big of a coincidence to think that in 45% of cases the most appropriate target-date fund series for a plan sponsor is the one from their recordkeeper.

Looking at the same issue from a different perspective, in the fourth quarter of 2016, Cerulli administered a survey to 801 DC plans in which they asked plan sponsors to identify their primary reason for selecting their current target-date funds: 46% let a consultant/advisor choose it for the plan; 30% based the decision on a familiar brand; 29% went with the recordkeeper's target-date funds series; 24% based the decision on investment performance; and 7% went with the lowest cost.[1] There may be some issues with the construction of this survey question, but, if accurate, these are disturbing responses— they suggest plan sponsors are at times guessing which product is right for their plan. Perhaps more heartening, in the same survey among plan sponsors who selected a custom target-date solution, 79% selected either fiduciary concerns, the need to consider the plan’s unique demographics, the presence of a defined-benefit plan, or the desire to avoid manager concentration as their reason for selecting a custom target-date solution.

Among the three long-term QDIA options, there is a relatively broad consensus among DC plan professionals that retirement managed accounts can create the best outcomes for plan participants, aside from fees. Managed accounts can provide a variety of features beyond the features of target-date funds, such as a personalized investment recommendation, a savings rate recommendation, retirement age guidance, and perhaps most importantly inform the participant if they are on track for a successful retirement. However, this advice tends to come with higher total fees when compared to target-date funds, in which managed account fees typically range from 0 basis points to 60 bps, which is in addition to fund fees. The relative low cost of target-date funds has resulted in target-date solutions capturing the majority of defined contribution plan assets. Lately, there is significant interest in “hybrid” QDIAs in

1 The Cerulli Report: U.S. Retirement Markets 2016. Chapter 10.

which some plan participants (often younger) are defaulted into a target-date fund series and other (often older) plan participants with more complex and less homogeneous circumstances are defaulted into retirement managed accounts.

A retirement managed account engine (i.e., robo-financial planning engine) can be leveraged to determine the ideal, personalized asset allocation recommendation for each participant. Next, individual participant recommendations are aggregated to aid plan sponsors and consultants in the QDIA selection process, providing a data-driven, repeatable, documented process allowing plan sponsors to stop guessing on QDIA selection and finding the best fit solution.

Plans sponsors need to address these types of questions:

1. Which type of QDIA is best for my plan—a target-risk fund, a target-date fund series, retirement managed accounts, or a hybrid?
2. Why and how would I select a hybrid QDIA over the three traditional QDIA options?
3. With a hybrid QDIA, at what age should the transition to managed accounts occur?
4. Which target-date fund series is most appropriate for the plan?
5. When do plan demographics necessitate a custom target-date glide path?
6. How do fees affect these selections?

This paper fits within a segment of the recent and broader "gamma" research that works to quantify the value of more intelligent financial planning decisions, as well as portfolio decisions, as put forth by Blanchett and Kaplan (2017). Other papers that attempt to quantify the value of the various aspects of financial planning include Blanchett and Kaplan (2013), Kinniry et al. (2014), and Envestnet (2015). The literature on quantifying the value of financial planning typically attaches a value ranging from about 20 bps to about 50 bps for "asset-allocation-only" advice, while the value of comprehensive financial planning has been estimated at around 300 bps per year.

In this article, we are far more specific, quantifying the benefit for a unique plan participant and then applying the approach to empower the retirement plan sponsor (or their consultant) to make a data-driven selection of a QDIA.

Increasing Availability of Participant Data

Only a few years ago, the amount of information that employers might know about a given employee was relatively limited compared to what it is today. Even if the information was available in the past, it was extremely challenging to consolidate into a readily useable form. As recordkeepers, custodians, and payroll solutions providers have increased the scope of the data they maintain as well as the flexibility of their systems, it has become much easier to get a relatively complete view of a plan's participants. Additionally, databases and data reporting have improved such that it's never been easier for a plan sponsor to develop a relatively rich demographic picture of their plan participants as depicted in Exhibit 1.

Exhibit 1 Information a Plan Sponsor Should Know About Participants

Available Information	Potential Implication for Asset Allocation Advice
Gender	Estimate life expectancy
Birth Date / Age	Estimate life expectancy and time to retirement
Retirement Date / Age	Determine time to save prior to retirement and length of time to retirement
Start Date / Tenure	Infer other non-plan assets
Salary & Salary Structure	Estimate human capital, risk capacity, and develop an estimate of total wealth
Savings Rate	Estimate human capital, risk capacity, and develop an estimate of total wealth
Plan Balance	Estimate financial capital and develop an estimate of total wealth
Employer Contributions	Estimate human capital and develop an estimate of total wealth
Defined Benefits	Estimate human capital, risk capacity, and develop an estimate of total wealth
Company Stock / Stock Options	Estimate financial capital, risk capacity, and develop an estimate of total wealth
Industry of Employer	Understand nature of human capital and develop a salary curve
Career / Occupation	Understand nature of human capital and develop a salary curve

The information in Exhibit 1 is far more meaningful than the type of information solicited solely from a typical risk-tolerance questionnaire, or RTQ. These RTQs often focus simply on risk preference and a nebulous time horizon and fall short of providing an objective analysis of the investor's total financial situation and capacity for risk. The available data provides a relatively rich picture of the participant's situation and their capacity to take on risk in pursuit of their retirement goal. Additionally, we can couple this participant data with outside survey data, such as the Federal Reserve's Survey of Consumer Finance or our internal managed accounts database of 1.4 million defined contribution plan investors, to relatively accurately infer additional information about the participant. Most importantly, this allows us to estimate Social Security benefits and the total amount of financial assets not in the plan. More specifically, Blanchett and Idzorek (2015) presents a framework for using age, tenure, salary, savings rate, and current balance to infer the target balance for a DC plan participant to compensate for lack of information on outside assets.

In this paper we use a regression model to estimate the amount and risk level of financial assets outside the DC plan. The approach, outlined in detail in Appendix 1, is based on data from the Federal Reserve's 2016 Survey of Consumer Finances.

Leveraging a Retirement Managed Account Engine

Armed with a relatively extensive and accurate picture of an individual participant, this detailed individual participant data is passed to Morningstar's managed accounts engine, which can use the information to arrive at an individualized asset allocation recommendation. The methodology used to determine the appropriate portfolio (i.e., the portfolio assignment approach) is explained in greater detail in Appendix 2. In Exhibit 2 we illustrate this process of taking available data, passing it to an automated retirement managed account engine, and receiving an individualized asset allocation recommendation for a given participant — represented by the blue bubble at the right side of Exhibit 2.

Exhibit 2 Using Individual Participant Data to Arrive at a Recommendation

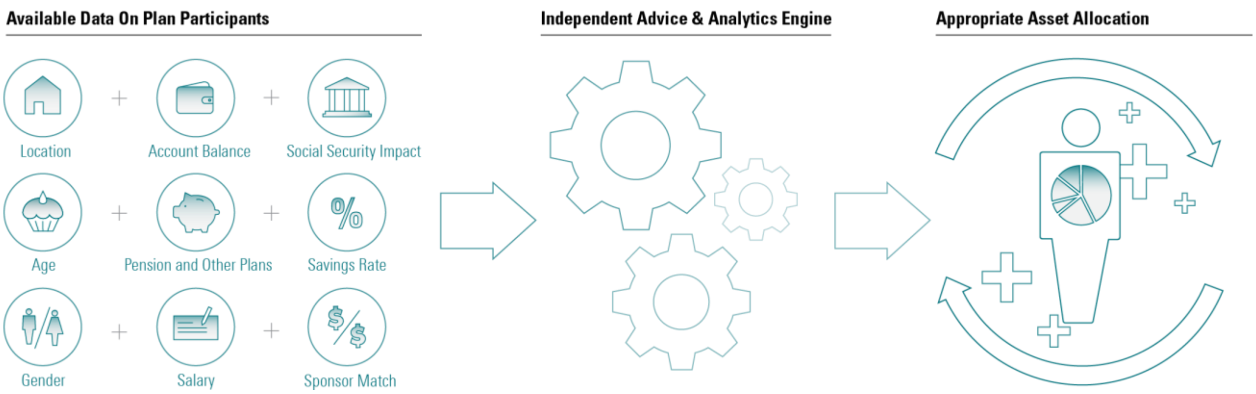

Clearly the quality and completeness of the individual participant data as well as the quality of the "advice" engine used will impact the efficacy of the final asset allocation recommendation. Currently, there are only a handful of full-fledged retirement managed accounts engines capable of carrying out this process in manner that might rival that of a financial planner/advisor, but with the growing number of robo-advisors we suspect this type of capability should be relatively widely available. A comparison of the various advice engines is beyond the scope of this paper, so we move forward under the assumption that retirement managed account engines are highly efficacious and arrive at the "correct" recommendation.

To help visualize the overall analysis, we introduce a relatively new type of chart that we call the "bubble chart." The bubble chart is drawn in the same "space" used to draw a target-date equity glide path in which the vertical axis represents the equity allocation. The horizontal axis typically shows age or in some case the expected calendar year of retirement. Using the concept of a bubble chart, Exhibit 3 illustrates how the process depicted in Exhibit 2 is applied for an individual participant to arrive at a single plot point (i.e., the "bubble") representing the recommended allocation to equities (and hence the overall stock/bond split).

Exhibit 3 Ideal Asset Allocation — For an Individual Investor

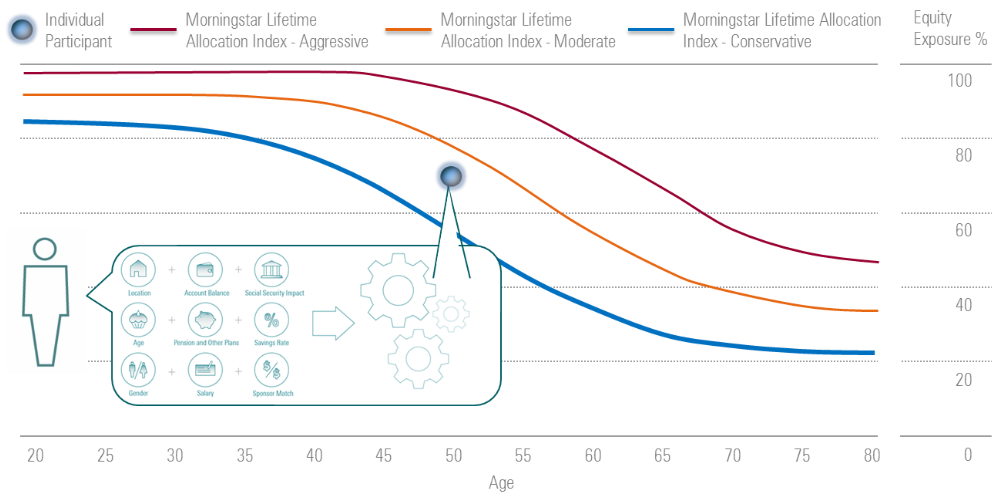

Source for Glide Paths: Morningstar Direct℠ as of June 30, 2016.

In Exhibit 3, we see that the participant in question is 50 years old (the horizontal axis) and the recommended asset allocation coming from the managed account advice engine for this individual is 70% equities (the vertical axis). Exhibit 3 also includes three glide paths — Conservative, Moderate, and Aggressive — of the Morningstar® Lifetime Allocation Indexes℠ to put the bubble into context. The majority of the 60-plus target-date fund family glide paths in the United States fall between the Conservative and Aggressive paths, so the inclusion of these three glide paths clearly shows how an individual plot point might compare to a typical target-date allocation at a particular age.

Moving from the case of a single individual participant to many participants, Exhibit 4 shows all of the individualized asset allocation recommendations for a relatively large plan. Within Exhibit 4, larger dots represent multiple participants of the same age and recommended asset allocation.

Exhibit 4 The Bubble Chart—For All Plan Participants

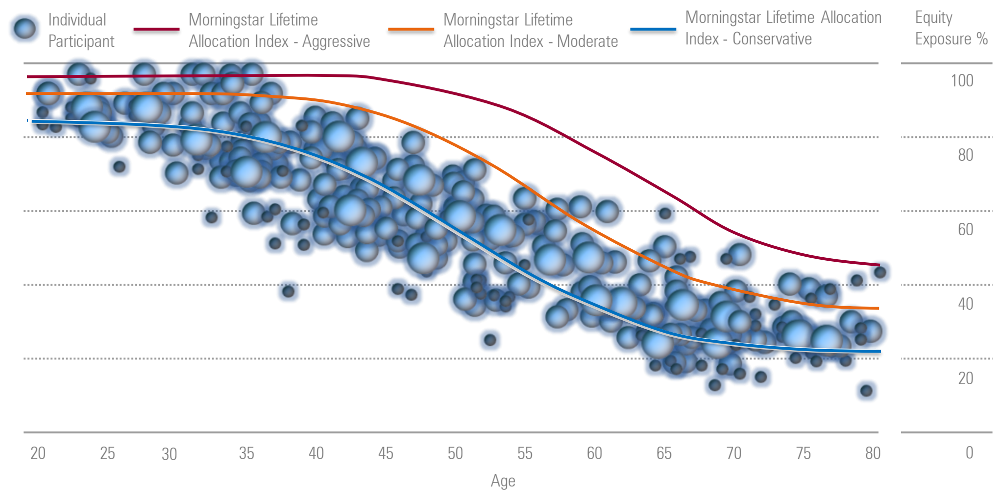

Source for Glide Paths: Morningstar Direct℠ as of June 30, 2016.

It is easy to see how powerful this comprehensive analysis could be for a plan sponsor attempting to select an appropriate QDIA. An implication of Exhibit 4 is that the only type of QDIA capable of a unique asset allocation recommendation for each participant is managed accounts, as any glide path will be too aggressive for some participants and too conservative for others. Moving from left to right along the age spectrum, the bubbles representing the recommended equity allocation follow a typical pattern in which the recommend amount to equity decreases with age. Given this pattern, it follows that a single balanced target-risk fund would be relatively inappropriate for a large percentage of the plan population. We can see that among the three Morningstar® Lifetime Allocation Indexes℠, the Conservative glide path (shown with a white edge to create contrast) fits this particular plan best, suggesting that this particular plan sponsor should probably focus its target-date search on more conservative target-date fund families.

We believe the bubble chart in and of itself is a powerful tool for helping plan sponsors make a much more informed QDIA selection; however, it opens up a host of more rigorous techniques for quantifying appropriateness and the value of a better fit.

Quantifying and Contrasting the Benefit of Different QDIA Options

In this section, we introduce a framework for quantifying the appropriateness—or perhaps more specifically, the "goodness of fit"—for different QDIA options based on the ideal, individualized participant asset allocation recommendations. We use a simplified example consisting of 10 plan participants to introduce this intuitive quantitative framework.

In Exhibit 5, the managed accounts engine has been used to arrive at an individualized asset allocation recommendation for the 10 participants as represented by the 10 blue bubbles. We start by assuming that the QDIA option being evaluated is a 60% equity and 40% fixed-income target-risk (balanced) fund. The orange line at the 60% equity level indicates that regardless of age the target-risk fund has a constant asset allocation, i.e. "one size fits all." For the 10 participants we have indicated the distance in equity allocation percentage points that the ideal individualized asset allocation is from the 60% equity orange line.

Exhibit 5 Quantifying Asset Allocation Misfit vs. Target-Risk (60/40 Balance) Fund

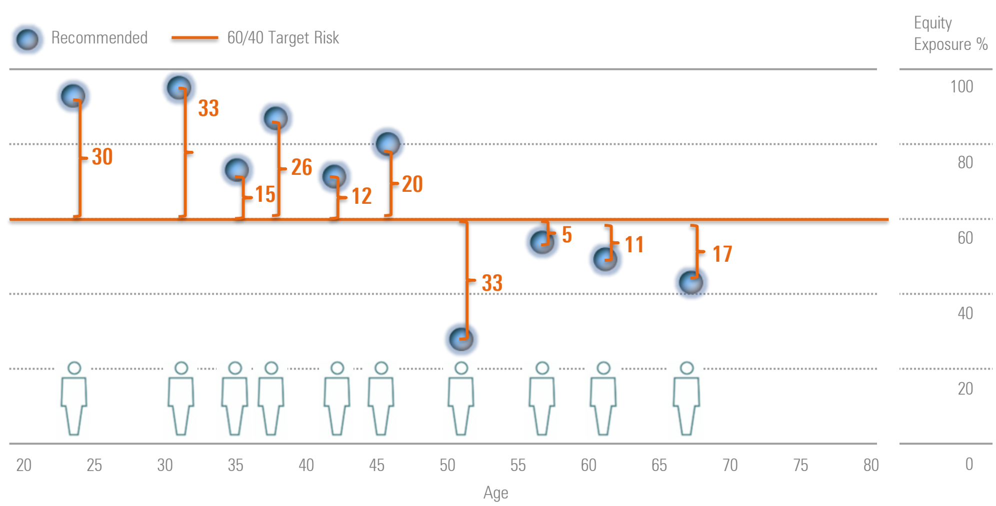

Source: Morningstar, Inc.

Exhibit 6 illustrates how the asset allocation misfits can be quantified. In this case, the median absolute difference in the recommended asset allocation and that of the target-risk (balanced) fund is 18.5 equity percentage points, indicating that the median investor is relatively far from an idealized target asset allocation. Average absolute differences are reported as well.

Exhibit 6 Quantifying Misfit—Target-Risk (60/40 Balanced) Fund

		% Equity			
Participant Number	Age	Recommended	Glide Path	Difference	Absolute Difference
1	23	90	60	30	30
2	32	93	60	33	33
3	35	75	60	15	15
4	37	86	60	26	26
5	43	72	60	12	12
6	45	80	60	20	20
7	51	27	60	–33	33
8	57	55	60	–5	5
9	61	49	60	–11	11
10	67	43	60	–17	17
Average:					**20.2**
Median:					**18.5**

Moving to Exhibit 7, we repeat the analysis, but we now assume the QDIA option under consideration is a target-date fund with a relatively conservative asset allocation, rather than the 60/40 balanced fund. In this example, we use the Conservative glide path from the Morningstar® Lifetime Allocation Index℠ family. Of the 10 participants, only one has a recommended asset allocation that sits roughly on the glide path. The rest of the participants have recommended asset allocations that are either above or below the Conservative glide path. Once again, we note the distance (expressed in equity percentage points) that the recommended asset allocation is from the potential QDIA solution (the Conservative glide path).

Exhibit 7 Quantifying Asset Allocation Misfit vs. Conservative Glide Path

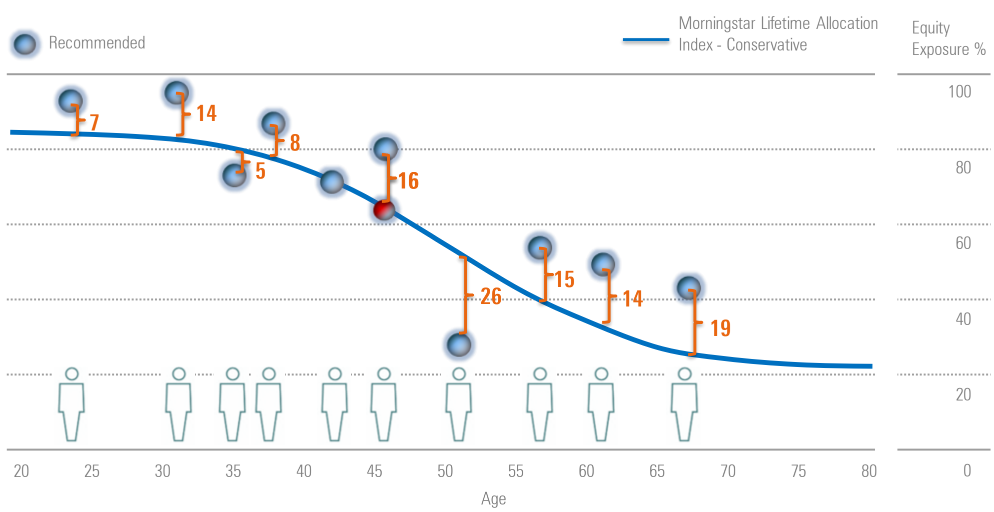

Source: Morningstar, Inc.

In Exhibit 8, we consider the Moderate glide path. In this case, the Moderate glide path seems to fit the majority of these 10 plan participants better than the Conservative glide path (although the seventh participant from the left, age = 52, will not be served particularly well by the Moderate glide path).

Exhibit 8 Quantifying Asset Allocation Misfit vs. Moderate Glide Path

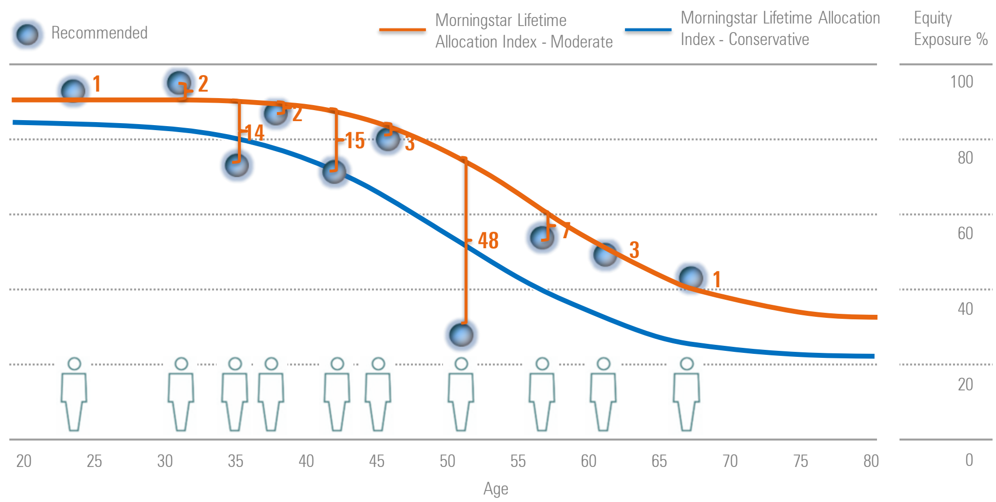

Source: Morningstar, Inc.

Exhibit 9 quantifies the misfits of the Conservative and Moderate glide paths versus the idealized targets. Relative to the misfits associated with the 60/40 target-risk (balanced) fund depicted in Exhibit 6, the misfits in Exhibit 9 are significantly smaller. The average and median absolute distance of the ideal asset allocation for a participant versus that of the glide path is much lower for the Moderate glide path. Looking at the final column of Exhibit 9, the Moderate glide path fits 70% of the participants better while the Conservative glide path fits 30% of the participants better.

Exhibit 9 Quantifying Misfit—Conservative and Moderate Glide Paths

		Conservative Glide Path				Moderate Glide Path				
Participant Number	Age	Recommended % Equity	Glide Path % Equity	Difference	Absolute Difference	Recommended % Equity	Glide Path % Equity	Difference	Absolute Difference	Better Fit Glide Path
1	23	90	83	7	7	90	90	0	0	Moderate
2	32	93	82	11	11	93	90	3	3	Moderate
3	35	75	80	–5	5	75	89	–14	14	Conservative
4	37	86	78	8	8	86	88	–2	2	Moderate
5	43	72	72	0	0	72	87	–15	15	Conservative
6	45	80	64	16	16	80	83	–3	3	Moderate
7	51	27	53	–26	26	27	75	–48	48	Conservative
8	57	55	40	15	15	55	62	–7	7	Moderate
9	61	49	35	14	14	49	52	–3	3	Moderate
10	67	43	24	19	19	43	42	1	1	Moderate
Average					**12.1**				**9.6**	
Median					**12.5**				**3.0**	

We are working with a very simple example, but this technique can be used to compare any number of actual target-date fund families (e.g. those from Vanguard, Fidelity, T. Rowe Price, BlackRock, etc.), helping a plan sponsor make a data-driven decision.

Changing gears slightly, a number of plan sponsors are interested in the use of a custom target-date glide path that is specifically designed to fit the unique demographics of their plan. It is unclear how the various purveyors of custom target-date solutions create tailored glide paths and perhaps it is self-evident, but the concepts behind the bubble chart represent a robust technique for creating a data-driven custom target-date glide path. Custom target-date solutions, which are almost exclusively sought by large plans with thousands of plan participants, tend to naturally result in a smooth, downward sloping glide path. To continue with our simplistic example, we solve for a best-fit glide path under the artificial constraint that the glide path cannot glide upward and must remain smooth. Exhibit 10 illustrates the custom glide path (green line) that fits these 10 participants best.

Exhibit 10 Quantifying Asset Allocation Misfit—A Custom Glide Path

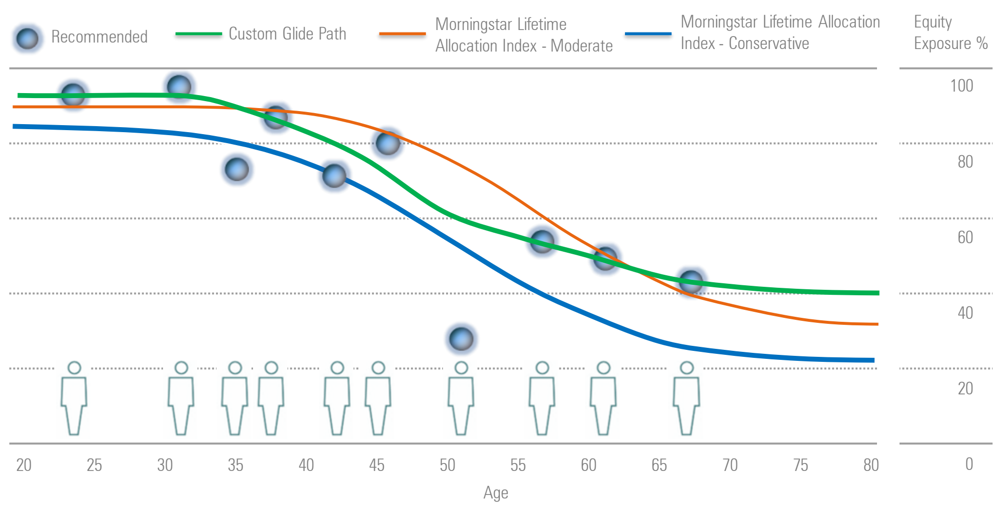

Source: Morningstar, Inc.

Exhibit 11 quantifies the asset allocation misfits associated with the custom glide path, which are considerably lower than the misfits associated with the Conservative and Moderate glide paths displayed in Exhibit 9. For this plan and its hypothetical 10 participants, the custom glide path offers an even better fit.

Exhibit 11 Quantifying Misfit—Custom Glide Path

		% Equity			
Participant Number	Age	Recommended	Glide Path	Difference	Absolute Difference
1	23	90	91	–1	1
2	32	93	91	2	2
3	35	75	88	–13	13
4	37	86	86	0	0
5	43	72	80	–8	8
6	45	80	82	–2	2
7	51	27	58	–31	31
8	57	55	55	0	0
9	61	49	49	0	0
10	67	43	43	0	0
Average					**5.7**
Median					**1.5**

Source: Morningstar, Inc.

Incorporating Fees—Utility-Based Framework

So far, we have avoided the complicating subject of fees. In our examples above, we saw that the Conservative glide path fit the 10 participants better than the static target-risk (60/40 balanced) fund; the Moderate glide path fit the participants better than the Conservative glide path; the custom glide path fit the participants better than the Moderate glide path, and that retirement managed accounts fits the participants perfectly (by definition given the set-up). Presumably, if the options all cost the same, the plan sponsor would always choose retirement managed accounts.

While this may or may not be the case, for the purposes of this example, let's now assume that the Conservative glide path cost more than the static target-risk (60/40 balanced) fund; the Moderate glide path costs more than the Conservative glide path; the custom target-date glide path costs more than the Moderate glide path; and managed accounts costs more than a custom target-date glide path. Not unlike the real world, in our example a more tailored solution costs more. How should a plan sponsor balance the trade-off of higher expenses versus a more appropriate asset allocation fit?

We need a more sophisticated approach that can quantify the trade-off between the difference in fees and the degree to which a given solution fits the plan demographics better. Utility theory provides this type of analysis by formalizing the trade-offs and ultimately determining a precise cost hurdle that makes the investor/participant indifferent to two asset allocation options based on the participants unique risk tolerance/risk aversion. We interrupt our example to present this expected utility-based framework.

In this brief section we include several formulas for those that might be interested, but very little is lost by skipping over the various formulas. Harry Markowitz established the mean-variance optimization framework, which requires three sets of inputs: returns, standard deviations, and pair-wise correlations for the assets in question. When seeking the optimal mix (the utility maximizing mix) for a specific investor, the mean-variance optimization framework takes the following form:

$$U = E(R) - 0.5\lambda\sigma^2$$

where
U = the investor’s utility for the asset allocation;
R = the return of the asset allocation;
λ = the investor’s risk aversion coefficient; and,
σ^2= expected variance of return of the asset allocation.

For a given investor, the optimal asset allocation is the one that maximizes the investor’s utility. Having the ideal asset allocation results in zero utility loss, by definition, since it represents the utility-maximizing asset allocation. An important assumption that we are making is that with an individualized retirement managed account that uses all of the known information about a given participant, the participant will ultimately be assigned to the ideal asset allocation—their personal utility-maximizing asset allocation. This assumption enables us to solve for each investor’s unique risk-aversion coefficient.

When a plan moves to a one-size-fits-all model, most if not all participants will experience some loss in utility relative to an ideal, individualized recommendation. One can calculate the expected loss in utility for an individual participant, which, in turn, can be aggregated to calculate the total expected loss in utility for the entire plan. This enables plan sponsors to leverage a widely used technique from economics and mathematically compare different potential QDIA options. For example, plan sponsors can compare the gain in utility associated with the incremental higher fee of managed accounts to determine if managed accounts justifies the higher fee. Additionally, plan sponsors can determine which glide path minimizes expected utility loss given an individualized assessment of every plan participant. Likewise, one can quantify the expected utility differential between an off-the-shelf glide path versus a potential custom glide path. Another application of this would be helping plan sponsors calibrate the pivot point age for a hybrid QDIA with the highest relative utility. In other words, plan sponsors can stop guessing, and make informed decisions based on a framework steeped in rich theory.

In our case, the retirement managed accounts/advice engine has solved for the optimal stock/bond asset allocation and we treat it as the ideal, utility-maximizing target asset allocation. We can couple this with capital market assumptions and then solve for the implied risk-aversion coefficient that would make the recommended stock/bond asset allocation the utility-maximizing mix. Then, using the implied risk-aversion coefficient for the participant, we can calculate the utility of the corresponding age-based asset allocation of the solution in question. Unless the asset allocation of both the recommended and alternative solutions are identical, the utility of the solution being evaluated will be somewhat lower, allowing for us to calculate the decrease in utility for each participant, which is equivalent to the fee differential that would make an investor indifferent to the two competing solutions. For those interested in some of the mathematical rigor, Exhibit 12 illustrates the framework for comparing the utility of the solution in question to that of the ideal asset allocation.

Exhibit 12 Estimating the Cost/Benefit of Better Fit

In the absence of fees, the ideal solution will always have a higher utility.

$$\underbrace{U_A = E[R_A] - .5\lambda\sigma_A^2}_{\text{Utility of Solution/Portfolio}} \qquad \underbrace{U_B = E[R_B] - .5\lambda\sigma_B^2}_{\text{Utility of Perfect Asset Allocation}}$$

The cost of the perfect solution in basis points is the difference between the two.

$$\text{Cost/Benefit } = [E[R_A\,] - .5\lambda\sigma_A^2] - [E[R_B] - .5\lambda\sigma_B^2]$$

Put differently, it is the "Fee" that satisfies the following.

$$[E[R_A\,] - .5\lambda\sigma_A^2] - [E[R_B] - Fee - .5\lambda\sigma_B^2]\text{=0}$$

The subscript *A* is used to identify the asset allocation solution/portfolio being compared to the ideal or perfect asset allocation, identify by subscript *B*.

Working with a set of capital market assumptions included in Appendix 3, Exhibit 13 shows the loss in utility associated with holding an asset allocation that does not match the ideal or perfect asset allocation. These "utility differentials" should be treated in this case as the annual difference in basis points that makes the participant indifferent to two potential asset allocations. In Exhibit 13, the vertical dimension displays the ideal or perfect asset allocation, while the horizontal dimension displays the actual asset allocation.

Exhibit 13 Loss in Utility, in Basis Points, From Asset Allocation Misfit

Recommeded Asset Allocation (Equity%) \ Asset Allocation of QDIA Being Evaluated (Equity%)	100	90	80	70	60	50	40	30	20	10
100	0	–3	–16	–40	–84	–158	–288	–531	–1,075	–3,025
90	–3	0	–4	–18	–47	–101	–200	–390	–823	–2,390
80	–12	–3	0	–4	–21	–57	–128	–271	–604	–1,830
70	–28	–14	–4	0	–5	–25	–72	–173	–420	–1,345
60	–50	–31	–16	–4	0	–6	–32	–97	–269	–934
50	–78	–55	–35	–18	–5	0	–8	–43	–151	–598
40	–112	–86	–62	–40	–21	–6	0	–11	–67	–336
30	–152	–125	–98	–72	–47	–25	–8	0	–17	–149
20	–199	–170	–140	–112	–84	–57	–32	–11	0	–37
10	–252	–221	–191	–161	–131	–101	–72	–43	–17	0

Notice that when the perfect asset allocation and actual asset allocation match, there is no loss in utility (the diagonal values highlighted in white in Exhibit 13). Intuitively, when the mismatch is larger, the implicit cost as represented by a decrease in utility is more significant. For smaller asset allocation mismatches, the change in utility is relatively symmetrical. However, as the mismatch grows large, the disutility of being underinvested in equities becomes significantly larger than the disutility of an equivalent overinvestment in equities. In other words, from this lens, if a plan sponsor is going to make a mistake when selecting a QDIA, it would be better to select a QDIA that is too aggressive rather than too conservative.

While the figures near the diagonal seem very reasonable, as you may have noticed some of the numbers, especially those in the final column, are extreme, and we would argue disconnected from economic reality. For example, if the perfect asset allocation is 90% in equities and the actual asset allocation is 20% in equities, the loss in utility is 8.23% or 823 bps per year! While this may in fact be the utility loss, it exceeds any reasonable cost that an investor or plan sponsor would need to pay for professional advice to be directed to an appropriate asset allocation. Therefore, we believe it is reasonable to impose a constraint on the maximum loss in utility that corresponds to the cost for such advice. There are a variety of ways to set this maximum, and this flexible framework allows the constraint to be based on the prices a plan may actually pay. For the generalized analysis presented here, we think 100 bps seems like a reasonable level and proceed with this assumption, although larger plans with more pricing power should use a lower number while micro-plans may need to use a slightly higher number—it really comes down to what is available to a given plan sponsor at a given fee level. It is worth noting that this framework and the results in Exhibit 15 are extremely close to what was presented in Blanchett and Kaplan (2017), only here we have explicitly constrained the maximum utility differential.

Before we return to our examples, to assist those that might be seeking a rough heuristic to quantify the advantage of a better fit, we constructed a version of Exhibit 13 using 1% increments (not shown due to space considerations, but a 100 x 100 matrix of 10,000 possible misfits), imposing the 100-basis-point maximum utility loss, and then calculating the differences in basis points associated with various degrees of asset allocation misfit. The 100-basis-point constraint results in a somewhat linear progression in Exhibit 14, for what would typically be an exponential relationship. From Exhibit 14, average asset allocation fit improvements of 1 percentage point are worth cost increases of approximately 1/4th of a basis point; 5 percentage points are worth about 4 basis points, 10 percentage points are worth about 12.5 basis points, etc. These are economically meaningful numbers that plans sponsors should consider when selecting a QDIA for their plan. Depending on the costs of the different options, this could certainly influence the selection of a different target-date glide path or justify the move to a custom target-date solution. For example, if one glide path fits a plan’s participants better by 10 percentage points on average, a plan sponsor should be willing to spend up to approximately 12 additional bps for that better fit. To put it differently, a $100 suit might be the cheapest option, but if it fits terribly it isn’t a good value because it hasn’t served its purpose of making the wearer look good. Exhibit 14 serves as rough guide to how much a plan sponsor should be willing to spend for an improved fit.

Exhibit 14 The Cost (or Benefit) of a Better Average Asset Allocation Fit

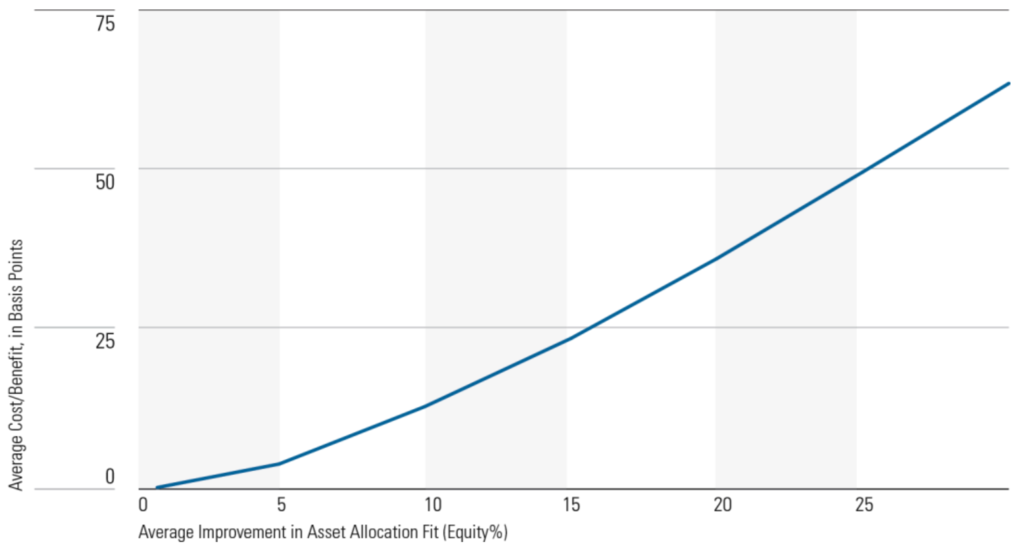

Returning to our 10-person example, the average participant's ideal asset allocation was 20.2 equity percentage points away from the asset allocation of the target (60/40) balanced fund, 12.1 equity percentage points from the Conservative glide path, 9.6 equity percentage points away from the Moderate glide path, and 5.7 equity percentage points away from the custom target-date glide path. Based on Exhibit 14, from an asset allocation improvement perspective, managed accounts is worth approximately 35 basis points more than the target (60/40) balanced fund, about 19 basis points more than the Conservative glide path, about 12 basis point more than the Moderate glide path, and about 5 basis points more than the custom glide path.

In Exhibit 15, working with the same 100 basis points maximum on utility loss, we calculate the participant-level utility differences to compare to the estimates based on Exhibit 14. The left-hand columns contain the unconstrained utility estimates of each of the possible solutions. The four right-hand columns contain the constrained utility differences. Exhibit 15 leads to similar levels of marginal benefits versus those that we approximated using Exhibit 14, although differences in the averages are a bit smaller and specific to this plan.

Exhibit 15 Using Participant-Level Utility to Compare QDIA Options

				Utility of Solution					**Utility vs. Perfect Asset Allocation**			
					Glide Path				Glide Path			
Participant Number	Age	Recommended % Equity	Implied Risk Aversion	Perfect Asset Allocation	Custom	Moderate	Conservative	Target Risk (Balanced)	Custom	Moderate	Conservative	Target Risk (Balanced)
1	23	90	2.77	5.92	5.92	5.92	5.90	5.61	–0.03	0.00	–1.70	–31.13
2	32	93	2.68	6.01	6.01	6.01	5.97	5.65	–0.13	–0.30	–4.05	–36.43
3	35	75	3.34	5.45	5.38	5.37	5.44	5.36	–7.04	–8.17	–1.04	–9.38
4	37	86	2.90	5.79	5.79	5.79	5.77	5.55	0.00	–0.14	–2.32	–24.50
5	43	72	3.48	5.36	5.33	5.26	5.36	5.29	–2.78	–9.78	0.00	–6.26
6	45	80	3.13	5.61	5.60	5.60	5.51	5.45	–0.16	–0.35	–9.99	–15.61
7	51	27	9.71	3.86	2.70	1.07	3.04	2.54	–100.00	–100.00	–81.93	–100.00
8	57	55	4.60	4.82	4.82	4.79	4.69	4.80	0.00	–2.81	–12.91	–1.43
9	61	49	5.18	4.62	4.62	4.62	4.49	4.54	0.00	–0.58	–12.67	–7.82
10	67	43	5.94	4.42	4.42	4.42	4.16	4.21	0.00	–0.07	–26.75	–21.41
Average				**5.19**	**5.06**	**4.88**	**5.03**	**4.90**	**–11.0**	**–12.2**	**–15.3**	**–25.4**

For this example, based on Exhibit 15, a plan sponsor should select managed accounts (the optimal asset allocation solution) as long as its incremental cost is lower than the averages in the final four columns (ignoring the other potential benefits of managed accounts beyond portfolio assignment, which we'll discuss next). The Moderate glide path is about 3.1 bps better than the Conservative glide path. The custom glide path is 1.2 bps better than the Moderate glide path. If retirement managed accounts cost less than 11 bps more than a custom glide path, retirement managed accounts would make sense. In this example, if retirement managed accounts cost more than 11 basis points more than a custom target-date solution, from an asset allocation-only fit perspective, retirement managed accounts might not be justified.

There is one glaring shortfall to this framework: it overlooks the other potential benefits of retirement managed accounts, which can include savings-rate advice, retirement-age advice, tax-location advice, fund-specific monitoring and recommendations from the plan lineup, annuitization advice, Social Security maximization, and drawdown (amount and tax efficiency) advice. Target-risk and target-date solutions don't provide these, and thus, the analysis focuses single factor --- asset allocation advice. As covered in our literature review, broader papers focused on the overall value of financial planning would suggest that these other factors offer real benefits worth paying for, although that discussion is outside of the scope of this paper.

Analyzing Hundreds of Plans and Hundreds of Thousands of Participants

We now move beyond our simplified examples, and look at real world plan data. Working with a filtered data set in which both participants and plan sponsors remained anonymous, we looked at a data set consisting of 333,552 individuals representing 276 401(k) plans. The plans ranged in size from 103 to 35,170 participants. From a filtering perspective, we only include plans with at least 100 participants who are identified as being able to actively participate in the plan. Additionally, we only include plans with at least five participants for each five-year increment from ages 25 to 65.

For each participant, at a minimum, we have data on the participant's age, salary, years of tenure, balance, and savings rate (broken down by Traditional and Roth). We do not have complete information on employer contributions, so we assume a standard employer match of 50% of the first 6% of deferrals for all plans. Data on employer contributions, as well as whether the plan offers a defined benefit plan, would result in greater differences noted in the analysis. Exhibit 16 includes some descriptive statistics across all participants included in the analysis (Panel A) as well as median values by select ages (Panel B).

Exhibit 16 Descriptive Statistics

Panel A: Descriptive Statistics of Test Participant Population

				Deferral %		
Percentile	Age	Years in Plan	Salary ($)	Traditional	Roth	401(k) Balance ($)
5th	27	1.33	22,638	0	0	1,673
25th	36	3.00	40,706	3	0	11,085
50th	46	6.51	64,256	6	0	39,232
75th	56	11.91	100,000	9	0	122,956
95th	65	22.58	210,000	20	6	481,355

Panel B: Median Values by Select Ages

			Deferral %		
Age	Years in Plan	Salary ($)	Traditional	Roth	401(k) Balance ($)
25	1.99	36,100	4	0	5,247
30	3.35	49,846	5	0	13,500
35	4.91	61,224	6	0	27,813
40	6.29	68,189	6	0	38,884
45	7.71	72,800	6	0	51,278
50	8.79	71,997	6	0	64,034
55	9.5	71,700	6	0	78,184
60	9.85	67,808	6	0	87,936
65	10.84	69,090	7	0	77,438

As illustrated back in Exhibit 2, the demographic data for each of the 333,552 plan participants are passed to the retirement managed accounts engine to arrive at an individualized ideal asset allocation recommendation. For our universe of plan participants, Exhibit 17 displays the recommended equity asset allocation percentiles by age.

Exhibit 17 Distribution of Optimal Equity Allocations by Age

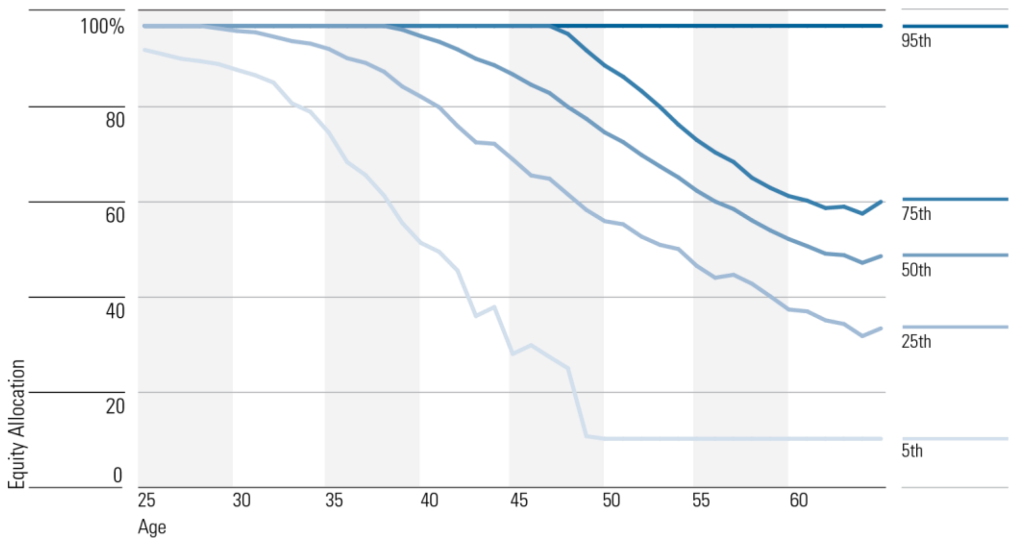

Not surprisingly, the equity asset allocations along the 50th percentile (or median) path are very similar to the Morningstar Moderate glide path. The distribution of equities (the dispersion of the lines associated with the different percentiles) clearly increases with age, offering evidence that younger investors are more homogeneous while older investors are more unique and likely not as well served by a one-size-fits-all solution.

Next, we estimate the utility of the respective ideal equity asset allocation per participant and compare that to the utilities that would be achieved from investing in each of the three Morningstar® Lifetime Allocation Indexes℠ as well as an assumed custom, per plan, target-date fund. For this large-scale analysis, the custom target-date fund is determined by fitting a fourth-order polynomial through the optimal participant allocations within a given plan. We compare the utility differences at the individual participant level and aggregate them at the plan level. Exhibit 18 contains the respective utility distributions for the three Morningstar® Lifetime Allocation Indexes℠ and the plan-specific assumed custom glide path.

Exhibit 18 Distribution of Utility Costs by Selected Glide Path

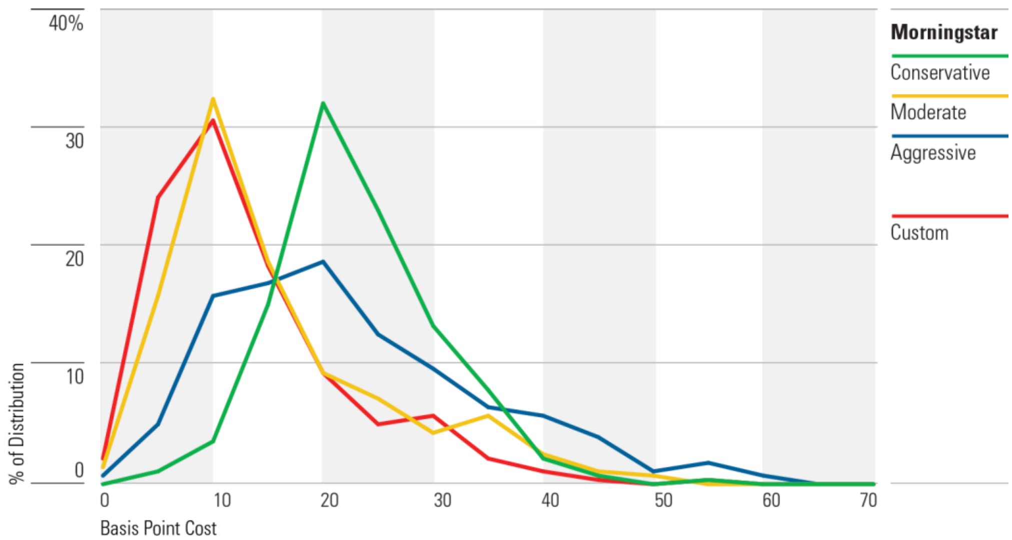

Exhibit 18 is a bit more challenging to interpret than most of our previous figures. Each distribution curve represents all 333,552 plan participants, in which the vertical axis displays the percentage of the distribution curve in question and the horizontal axis represents the basis point loss in utility of the solution in question. In this case, distributions with a mass at the far left (lower numbers) fit the universe of plan participants better. In Exhibit 18, the average and median utility losses relative to the ideal asset allocation are as follows:

- Conservative – Morningstar® Lifetime Allocation Index℠: Average = 26 bps, Median = 25
- Moderate – Morningstar® Lifetime Allocation Index℠: Average = 18, Median = 15
- Aggressive – Morningstar® Lifetime Allocation Index℠: Average = 25, Median = 22
- Respective plan-specific custom glide: Average = 16, Median = 14

Based on this utility framework, if all 333,552 participants belonged to a single plan, managed accounts would lead each participant to the ideal asset allocation and among the non-managed account options custom would be best, followed by Moderate, then Aggressive, and then Conservative.

There is no agreed-upon definition of "best" when determining fit. Should "best" reflect a basic count of the number of participants whose best fit is closest to a given QDIA solution, or should it be based on the average or median utility loss?

Both approaches have their advantages and disadvantages. For the first approach we find that the Conservative, Moderate, and Aggressive are the best for 18%, 75%, and 6%, respectively. For the second approach we find that the Conservative, Moderate, and Aggressive are the best for 8%, 42%, and 51%, respectively. There are relatively significant differences between the two approaches; however, they generally suggest the Moderate Index is likely the best fit for most participants, followed by Aggressive, then Conservative. While it is interesting to look across all 333,552 participants, ultimately the selection of the most appropriate QDIA solution is a plan-level decision.

When we calculated the ideal asset allocation for the 333,552 participants, we kept track of which participants belonged to each of the 276 plans. Thus, for each plan we used their respective participant-level data to construct a potential custom glide path, resulting 276 custom glide paths. After sorting the 276 custom glide paths based on average equity exposures, Exhibit 19 displays the various custom glide paths corresponding to different percentiles, where the 95th percentile represents one of the most aggressive glide paths and the 5th percentile represents one of the most conservative glide paths.

Exhibit 19 Distribution of Optimal Custom Glide Paths Across Paths

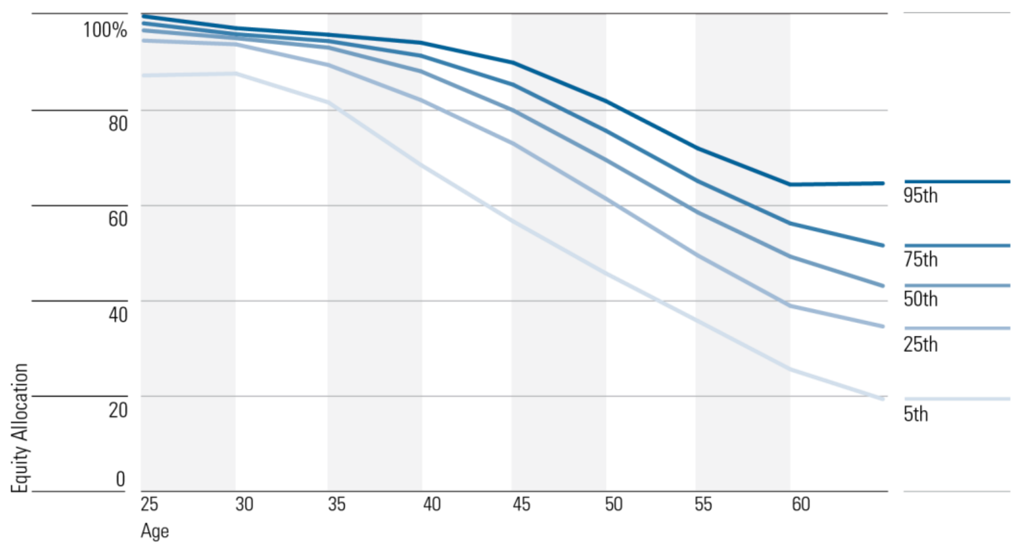

The most important takeaway from Exhibit 19 is that the ideal implied glide paths for different plans can and should differ significantly. For some plans, a conservative glide path will be most appropriate, while for others an aggressive glide path will be most appropriate.

In contrast to Exhibit 17, which looked at all participants regardless of plan, Exhibit 19 focuses on the distribution associated with the implied 276 custom target-date glide paths. In general, these glide paths are less extreme than the equity glide path implied by Exhibit 17, with median glide path from Exhibit 19 coming in somewhat more conservative than the median from Exhibit 17. In both Exhibits 17 and 19, the dispersion of the different percentiles gives merit to the notion of a hybrid QDIA in which a low-cost target-date fund is used for younger investors and a managed accounts solution is used for older investors with more diverse individual circumstances.

Our final analysis is designed to determine the potential pivot-point age for a hybrid QDIA. The setup for this is a bit different than the previous two analyses. We start with the assumption that if retirement managed accounts have the same cost as the target-date option in question, there is basically no need for a hybrid QDIA and the plan should select retirement managed accounts for all participants. For each participant in a given plan, we iteratively increase the cost of managed accounts until the utility is the same as the Moderate glide path. We then look at the average increase in cost that equalizes utility within different age-cohorts (20 to 24, 25 to 29, etc.). As expected, the average increase in cost is lowest for the youngest cohort and increases as we move to the 61 to 65 cohort. Keep in mind different plans pay different amounts for retirement managed accounts. If that cost is higher than the average increase in cost to equalized utility across all of the age cohorts, the plan should select the target-date option rather than the hybrid QDIA. However, if the average increase in cost to equalize utility eventually surpasses the fee associated with managed accounts, the plan has identified the pivot-point age. Recall that this setup continues to focus solely on the benefit associated with a more appropriate asset allocation and ignores most of the potential benefits of managed accounts. Thus, if the incremental cost of managed accounts is at all close to the numbers in Exhibit 20, it would be a strong argument for a plan sponsor to select managed accounts as the QDIA for at least some participants.

Exhibit 20 Distribution of Utility Costs vs. Morningstar Moderate Lifetime Allocation Index by Age Group

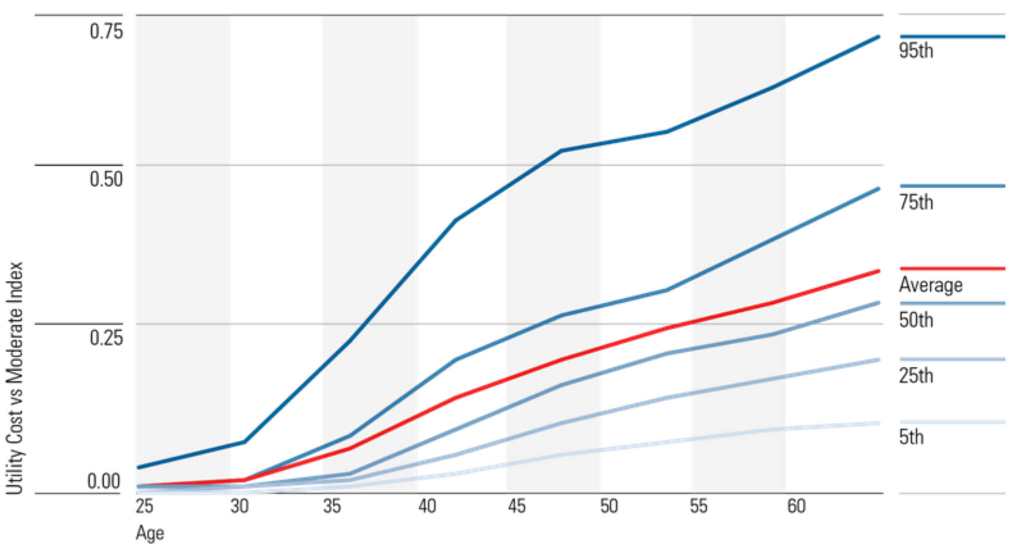

As expected, the value of managed accounts increases with age. Let's work through a hypothetical example to help with the interpretation of the information in Exhibit 20. Let's assume that an investable version of the Moderate glide path is available for 20 bps. A retirement managed account builds individualized, fund-specific portfolios typically using the non-target-date funds on a plan lineup. For ease of comparison, let's also assume that all of the funds on the lineup also cost 20 bps. Finally, let's assume that the incremental retirement managed account fee is 20 bps (recall that this fee can vary from 0 bps to around 80 bps, depending on plan size and myriad other factors). For the average line (red line) in Exhibit 20, this suggest for the 40 to 44 cohort and all younger investors the Moderate target-date solution should be the QDIA, and that for the 45 to 49 cohort and all older investors that retirement managed accounts should be the QDIA.

Rather than relying on the overall and generalized results of Exhibit 20, this analysis should be carried out at the plan level. This enables the data to speak directly to the plan's specific circumstances and the plan sponsor to make an informed decision based on the demographics of their plan and the utility differences of the possible QDIA solutions, and to incorporate the specific price points that are available to the plan.

Conclusions

Prudence requires DC plan sponsors to think carefully about fees and keep costs down, but that doesn't mean selecting the cheapest QDIA, especially if the cheapest solution doesn't fit the plan's population well. Prudence implies finding the right QDIA for the right price. The combination of fit at a low cost produces the best "value" or "utility" for a plan.

Unfortunately understanding fit and the trade-off between fit and cost has been a puzzle for plan sponsors. Providers have offered more customized solutions that should improve fit, typically at higher costs, but plan sponsors were left without much guidance on cost/benefit analysis. How much is too much to spend on customization, or how much customization is prudent—even necessary—for my plan?

This paper offers techniques for comparing the value of different QDIA offerings. By determining the ideal asset allocation for each individual, plan sponsors can use these techniques to calculate the optimal glide path for an entire plan, the optimal QDIA or hybrid QDIA, and the optimal age for a hybrid plan to pivot. A plan sponsor can then use this information to evaluate specific options that are available to a given plan at a given price point, enabling them to make an informed, fact-based decision.

Prudence also requires staying current, and we believe the QDIA selection landscape has changed in three important ways. First, as discussed in this paper, greater availability of participant data calls for a more robust, data-driven QDIA selection process. Second, automated advice solutions have made managed accounts not just attainable, but in many cases a prudent alternative to target-date funds. And third, the rise of the hybrid plan has made managed accounts attractive to a wider set of plans.

We believe the techniques described in this paper will empower plan sponsors and their consultants to stop guessing and use an easy-to-implement, yet rigorous QDIA selection process.

QDIA Selection Analysis

To request an analysis of your plan's demographics, please contact retirement@morningstar.com.

References

Blanchett, David M., and Paul Kaplan. (2013). "Alpha, Beta, and Now… Gamma."
The Journal of Retirement, vol. 1, no. 2:29-45.

Blanchett, David M., and Paul Kaplan. (2017). "The Value of a Gamma-Efficient Portfolio."
Morningstar Investment Management Working Research Paper.

Blanchett, David M., and Philip U. Straehl. (2015). "No Portfolio is an Island."
Financial Analysts Journal. vol, 71, no. 3: 15-33.

Blanchett, David M., and Thomas M. Idzorek. (2015). "The Retirement Plan Effectiveness Score | A Target Balance-Based Measurement and Monitoring System."
Morningstar Investment Management Working Research Paper.

Envestnet. (2015). “Capital Sigma: The Return on Advice.” Envestnet White Paper.

Idzorek, Thomas, Jeremy Stempien, and Nathan Voris. (2011). “Selecting a Target Date Benchmark”.
Morningstar Investment Management Research Paper.
http://corporate.morningstar.com/ib/documents/TargetMaturity/SelectTargetDateBenchmark.pdf

Kinniry, Francis M., Colleen M. Jaconetti, Michael A. DiJoseph, and Yan Zilbering. (2014).
“Putting a value on your value: Quantifying Vanguard Advisor’s Alpha.” Vanguard White Paper.
https://www.vanguard.com/pdf/ISGQVAA.pdf.

Appendices

Appendix 1: Estimating the Value and Risk Level of Non-DC Financial Assets

The Federal Reserve's Survey of Consumer Finances (SCF) contains data that can be used to estimate the value of financial assets held outside of a DC plan and their respective amount. The SCF dataset contains values that would be known to a recordkeeper (or the plan sponsor), such as participant age, salary, DC plan balance, tenure, and savings rate, as well as data that would not be known to the recordkeeper, in particular the value of non-DC financial assets and their respective risk level.

For our analysis, we use data from the 2016 SCF (the most recent version available). We reduce each household replicate to a single value, where the values are aggregated based on the respective household weights. We only include households if the respondent, who is assumed to be the DC plan participant, has a minimum 401(k) balance of $50, is saving at least 1% of salary in a 401(k) plan (i.e., is clearly eligible to participate in the DC plan), and has annual wages of $10,000 or more. These filters reduce the total households for consideration from 6,248 to 1,001.

Outside assets (i.e., non-DC assets) are estimated by subtracting the DC plan balance and all liquid savings accounts, which include checking, savings, and money market accounts, from total financial assets. Liquid assets are excluded since they are assumed to be used to fund ongoing day-to-day consumption expenses and would not generally be considered investible toward retirement. Due to the right skew associated with a number of the variables, we take the natural logarithm of the non-DC assets (the dependent variable for the outside assets regression), as well as participant salary, the DC plan balance (both independent variables), and include these values in the regressions.

We run two sets of ordinary least squares (OLS) regressions. In the first regression, the dependent variable is non-DC financial assets and in the second regression the dependent variable is the risk level of those assets, defined as the percentage of the assets invested in equities. The independent variables are values that would be known to the DC plan recordkeeper, and include age, plan tenure (in years), wages, balance, and total savings rate. For the first regression (total value of non-DC financial assets), in addition to the independent variables noted previously, we also include an interaction variable of tenure and balance and included the square of the natural logarithm of wage for the outside assets regression. For the second regression (risk level of non-DC financial assets), we only include respondents who have non-DC financial assets greater than or equal to $1,000 when estimating the risk level of non-DC financial assets, which reduces the number of from 1,001 to 518. The results for the two regressions are included in Exhibit 21.

Exhibit 21 OLS Regression Results

	Outside Assets	Risk Level (e%)
Intercept	40.837**	–168.390**
Age	0.058**	–0.568**
Plan Tenure	–0.223**	–0.213
ln(Wage)	–7.811**	17.777**
ln(Wage)²	0.411**	—
ln(Balance)	0.210**	3.932**
Deferral%	0.036**	0.266
Ten*ln(Balance)	0.016**	—
R Square	43.96%	11.41%
Observations	1001	518

From Exhibit 21, we see that the Outside Assets OLS regression is relatively robust, with an R-squared of 43.96% with all coefficients statistically significant at the 1% level (and technically coefficients significant at the 0.1% level). This suggests that we can confidently use this model to estimate the value of non-DC assets for participants in DC plans (based on available in-plan DC data). We find that outside assets are likely to be larger for participants who are older, have shorter tenure, larger wages, larger balances, and higher deferral rates. These coefficient signs are consistent with our expectations.

The estimated risk level of outside assets regression is less robust, with an R-squared of 11.41%. However, three of the five coefficients are statistically significant at the 1% level. We find that assets outside the DC plan are likely to be invested more aggressively for participants who are younger, have higher wages, and higher balances.

In Exhibit 22 we provide some guidance as to what our estimate of the non-DC assets would be and the respective size of the respective assets to the DC assets.

Exhibit 22 Estimated Outside Assets and Risk of Outside Assets Using Regression Model

	nonDC		
Age	Assets $	Equity %	Size vs DC Assets %
25	1,773	39.02	33.80
30	3,615	45.39	26.80
35	6,666	48.50	24.00
40	10,613	48.64	27.30
45	16,598	47.85	32.40
50	22,624	45.43	35.30
55	31,669	43.39	40.50
60	41,108	39.86	46.70
65	52,036	36.88	67.20

There are two key takeaways from Exhibit 22. First, the non-DC assets grow both in absolute terms and relative terms (versus DC assets) as age increases. This should not be surprising. The probability (and value) of saving outside the DC plan is likely to increase as people save over time. The median DC plan tenure for a 55-year-old is only approximately 10 years in the dataset. Therefore, if the participant saved before joining the plan (at age 45, which is likely) he or she may have some outside savings (if these monies were not rolled into the DC plan). Second, outside assets tend to be invested relatively conservatively (i.e., have risk levels that tend to be significantly lower than the Morningstar Moderate index). This has important implications with how DC monies should be invested (i.e., more aggressively) and is an important consideration when designing an optimal default.

Appendix 2: Portfolio Assignment Approach

The approach to determining the appropriate risk level for an investor's optimal portfolio allocation (i.e., portfolio assignment) is based on taking a holistic view of an investor's assets. We incorporate the total value and risk attributes of assets that are often overlooked, such as human capital and pension wealth, and use the financial assets in the DC plan as the "completion portfolio" to ensure diversification of the individual's total wealth.

A fundamental part of the total wealth process is modeling and understanding how an individual's wealth changes over the lifecycle. For younger individuals, human capital is typically the dominant household asset. Human capital can be thought of as the mortality-weighted net present value of an individual's future wage income. As individuals age, they tend to save money for retirement, thereby accumulating financial assets (both inside the DC as well as potentially outside it), along with accruing benefits in pension plans (such as Social Security). In other words, over time investors convert a portion of their salary (i.e., human capital) into financial capital by saving and accruing pension benefits, both of which can be used to fund retirement.

Research by Blanchett and Straehl (2015), among others, has noted that human capital is generally a relatively bond-like asset—it usually pays a steady "coupon" in the form of a paycheck, but its risk varies considerably across business cycles, by job skills, as well as the specific occupation and industry of the worker. Because human capital is bond-like and untradeable, a younger investor's financial assets should be invested more aggressively to achieve a more balanced risk level from a total wealth perspective. As the relative value of human capital (as a percentage of total wealth) declines as the individual ages, financial capital should be invested more conservatively to ensure the risk of the total wealth remains balanced throughout the lifecycle. This is the economic rationale underpinning the shape of many glide paths today.

There are two final considerations when determining the optimal risk level for a participant's portfolio. The first is how "on track" that individual/household is for retirement. Within the portfolio assignment process, individuals who are better funded (i.e., have higher funded ratios) can potentially take on more risk in their portfolio based on their target risk level using the total wealth approach. Second, other nonadviseable portfolios (e.g., an IRA or really any monies outside the DC plan) must be considered. For example, if an investor has a large IRA that is invested very aggressively, yet the overall total wealth target risk level is more balanced, the monies in DC plan should be invested more conservatively (and vice versa).

Appendix 3: Capital Market Assumptions for Utility Model

The return and risk assumptions used for this analysis are based on Morningstar Investment Management's 2017 unconditional return forecasts. The proxy for equities is the Russell 3000 and the proxy for fixed income (i.e., cash) is ICE BofAML US 3-Month Treasury Bill Index. We use a cash proxy for fixed income, versus a bond proxy (e.g., aggregate bonds) to better replicate a "risk-free" asset (while equities represent a risky asset). The assumed annual return for equities and fixed income is 9.27% and 3.17%, respectively. The assumed annual standard deviation for equities and income is 15.63% and 1.98%, respectively. The assumed correlation is -0.02.

Acknowledgments

We want to thank Brandon Strom, Carolyn Wilbon, Drew Carter, and John Erdodi for research/design support as well as Nathan Voris, Brian Cosmano, and Jacob Finkel for valuable comments.

About Morningstar Investment Management LLC

The research team within Morningstar's Investment Management group pioneers new investment theories, establishes best practices in investing, and develops new methodologies to enhance a suite of investment services. Published in some of the most respected peer-reviewed journals, the team's award-winning research and patented methodologies are used throughout the industry and are the foundation for many of the Investment Management group's solutions. The team's commitment to ongoing research helps maintain its core competencies in asset allocation, portfolio construction, and holistic automated financial planning.

Morningstar's Investment Management group creates investment solutions that combine its investment research and global resources together with the proprietary data and research of its parent company, Morningstar, Inc. The Investment Management group provides comprehensive retirement, investment advisory, and portfolio management services for financial institutions, plan sponsors, and financial advisors around the world.

For more information, please visit:
http://global.morningstar.com/mim

MORNINGSTAR®

22 West Washington Street
Chicago, IL 60602 USA

The information, data, analyses, and opinions presented herein are provided as of the date written. Past performance is not indicative and not a guarantee of future results. This white paper contains certain forward-looking statements. We use words such as "expects", "anticipates", "believes", "estimates", "forecasts", and similar expressions to identify forward looking statements. Such forward-looking statements involve known and unknown risks, uncertainties and other factors which may cause the actual results to differ materially and/or substantially from any future results, performance or achievements expressed or implied by those projected in the forward-looking statements for any reason. Past performance does not guarantee future results.