



The False Promise of Target Date Funds

By

David N. Esch, Ph.D. and Robert O. Michaud
New Frontier Advisors, LLC
Boston, MA 02110

Forthcoming in the *Journal of Indexes*, January/February 2014

Abstract

We show that, contrary to popular belief, investment glide paths with decreasing allocations to risky investments over time do not intrinsically reduce the volatility of wealth at a target date compared to other portfolio management plans. In spite of the derisking claim, risk from all investments, including those in earlier years, propagates forward to final wealth. Since volatility is not decreased, it remains a question of interest whether a glide path can increase expected or median wealth the end of the investment cycle. We implement a Monte Carlo simulation to examine the impact of various simple glide paths on the probability distribution of retirement wealth over a wide range of market outcome scenarios. Glide paths are matched on a simple risk measure to provide a fair comparison. For every glide path there is a fixed portfolio investment with the same risk at target date. The glide path with the best expected and/or median wealth at target date is sensitive to return distribution assumptions and contribution patterns, but regardless, equal risk glide paths generally provide similar wealth performance. For investors with an initial endowment, a fixed portfolio investment is superior to any glide path. Even for investors with predictable contributions over time, the benefit of a glide path is dubious. The relative performance of a particular glide path relative to other *a priori* equivalently risky portfolios is largely determined by the timing of portfolio returns. We believe that target date fund investors may have a false sense of risk control, and may be especially at risk if market performance does not match the assumptions that justified that choice of glide path. We further believe that it is bad policy to allow these funds to replace proper fiduciary review for naïve or otherwise uninformed investors.

I. Introduction

Since the U. S. Pension Protection Act of 2006, target date funds (TDFs) have surged in popularity as retirement investment strategies, largely because they are simple to administer, qualify as default investment alternatives (QDIA) for automatic enrollment in retirement plans and are perceived to lower risk at retirement by gradually decreasing allocation to high-risk assets. Market share of the 401K industry in TDFs has grown from 5% to 13% from 2006-2011, and the percentage of 401K participants holding TDFs has grown from 19% to 39%. Total assets in TDF mutual funds have grown from 71 billion at the end of 2005 to 481 billion in total net assets at the end of 2012 (Investment Company Institute 2013).

This current confidence in TDFs may be misplaced. We show that every glide path can be matched to many others with the same risk characteristics at a target date, including an unchanging portfolio and even paths increasing in equity exposure over time. Furthermore, although specific glide paths can be selected to maximize either expected or median wealth, the "best" glide path for a given risk level at retirement is sensitive to both mean-variance and contribution pattern assumptions and does not meaningfully increase the retirement wealth over nearby alternatives. Across a broad set of assumptions, while the optimal strategy is sometimes a gently decreasing glide path, the particular glide path is unlikely to be identifiable under realistic uncertainty about the inputs, and the benefits of this mathematical optimality are dubious in practical terms compared to other nearly optimal glide paths. In particular, a fixed portfolio over the investment period provides a transparent indication of its true risk characteristics, while being the best strategy for lump sum investment and often nearly optimal for other investment plans. For the purposes of simulation experiment presented here we consider the fixed portfolio as truly fixed and unchanging through the duration of the glide path, but in practice this type of investment could be adjusted to match changing risk appetites or markets, whereas a glide path needs to be followed through to attain its targeted risk level. The ability to cross-sectionally match a risk target is important, especially since market conditions through the entire investment horizon are unlikely to match those at the time when saving starts.

While each TDF has its own details as to how the portfolio is adjusted as retirement approaches, they all rely on a planned shift in holdings over time. This plan is generally set into motion at the onset of saving and acts as a "cruise control" or "auto-pilot," with relatively few interventions. There is no theory or even consensus among providers as to the optimal glide path, but invariably these funds are built on the assumption that riskier assets should have greater exposures at first, and as the target date nears, allocations should be shifted toward lower-risk securities. The perception is that the glide path will create a more favorable probability distribution for retirement wealth. The process of trading towards increasingly low-risk model portfolios is often referred to in the industry as "derisking" the portfolio. However, there is wide disagreement over what constitutes proper life cycle risk aversion, and indeed, some advocate for less risk in early investment years.¹

¹ This is partially based on the observation that early investment experiences affects lifetime investment behavior. An early experience of loss can lead to prohibitive risk aversion later in life.

In spite of the derisking claim, risk from all investments, including those in earlier years, propagates forward to final wealth. It is doubtful that a predetermined schedule of risk targeting, ignoring information updates, is a good or even acceptable way to optimize retirement wealth through uncertain markets. In order to achieve its intended risk characteristics at retirement, a glide path must be followed to completion and cannot be adjusted mid-path. Good investments depend on more than the number of years until retirement. New information must be recognized and acted upon to provide proper fiduciary service to investors. Through the savings period of a retirement plan, personal circumstances change and it seems unwise to leave retirement assets unattended in a locked one-size-fits-all plan which completely ignores both market and investor.

A better default investment for long-term investors without opinions or views about markets may be a strategic diversified portfolio close to the composition of the market, such as a 60% equity and 40% fixed income portfolio. Target-risk portfolios are more transparent about their risk exposures and offer expanded possibilities of adding value by seeking improved diversification across risk factors and asset classes that may be unattainable in a glide path. The costs and implementation of funds for many target dates may put constraints on the ability to cross-sectionally diversify a portfolio.²

This article sets out to reveal the relationship between glide paths and the return distribution at the target date through a simulation experiment. Our hypothesis is that predetermined portfolios up to the retirement date provide scant advantages over other possible plans, given market uncertainty. Rigorous glide paths amount to bets on particular market outcomes. We hypothesize that the "best" glide path will depend strongly on any observed patterns in market returns, and therefore that a specific choice of glide path is a bet that market conditions will change accordingly.

To provide evidence that our simulations are consistent with market realities, we also analyze glide paths on historical return data, show that outcomes are time-dependent across glide paths, and that the choice of glide paths has far less impact on retirement wealth than the overall risk tolerance through the entire investment history. Indeed, the historical example suggests an investor in a typical glide path has both a higher standard deviation and a lower median wealth than a comparable fixed risk investor.

The article is organized as follows: in section 2, we describe the main simulation experiment. In section 3, we present some details about the methodology in computing the simulations and statistics. In section 4, we present the simulation results and the results of our historical analysis, and in section 5, we test the robustness of our conclusions to changes in some of the assumptions of the experiment and address some other concerns about the experiment. Section 6 concludes the paper. Some additional technical analysis of the matching on standard deviation is in the Appendix.

² A typical family of TDFs may have as many as 12 dates in five year increments and an income fund for investors in retirement.

II. The Simulation Experiment

Our simulation experiment illustrates how glide paths affect retirement wealth and how the selection of the best glide path relies on results which are unknowable *a priori*, offering a false sense of risk control to the investor.³ Of course TDF products on the market are far more intricate and contain detailed glide paths for other features than just stock/bond ratios, but nevertheless we find the simple example presented here compelling. Adding complexity to the glide path only adds to the specificity of the assumed future performance of the markets.

Investor contributions are difficult to generalize. A typical investor may increase saving through the years before retirement as income allows. Since erratic contributions further hinder the value of a glide path, we assume a steady contribution, growing by a constant inflation rate, over a 40-year period. This would match a retirement investor saving, for example, from age 25 to 65, making the maximum matched contribution to a 401K plan over the course of a typical lifetime career.⁴

For all the simulations, we assume two assets with contrasting risk levels. These are meant to correspond to a typical stock and bond index. For our base case, we assume a stationary return distribution roughly matching the historical performance of U. S. stocks and bonds. We describe the mathematical details of the return assumptions in section 3 and we explore the impact of varying return assumptions and outcomes further in section 4.

We compute retirement wealth amounts for many thousands of simulations of market return histories and investment patterns. Each glide path implies a wealth amount at target date under a simulated history. Glide paths can be appropriately selected to fairly compare simulated target date wealth characteristics. Since the usual selling point of TDFs is risk control, we match glide paths on risk at target date, with risk measured by standard deviation of retirement wealth.⁵ We assess the suitability and quality of the match with other statistics calculated from the retirement wealth simulations, including the median as well as the lower and upper 1%, 5%, 10%, and 25% quantiles. We require these additional checks because wealth distributions are often highly right skewed and no one measure of risk totally captures the investment implications of investment risk policies. To further justify our comparisons of glide paths matched on standard deviation, we explore the target date wealth distributions of the matched glide paths in the Appendix using graphical and numerical summaries. We find that for our capital market assumptions the observed differences in skewness are not large enough to outweigh the marked differences in other criteria

³ Under restricted assumptions, optimal glide paths can be solved for analytically. We believe the illustrations of the Monte Carlo results are more persuasive since the most strict and unrealistic assumptions can be tested and varied.

⁴ We need not test the impact of varying this assumption too greatly, since the impact of an unexpected outstanding contribution in a particular year is mathematically the same in terms of retirement wealth as the impact of an outstanding portfolio return, and outstanding returns do occur with considerable frequency in the simulations. Because our experiment considers a wide range of outcome scenarios mathematically equivalent to erratic investment patterns, we ignore the pattern of contributions as an experimental variable without loss of generality.

⁵ We tried matching on other risk measures such as Value at Risk (VaR) and conditional value at risk (CVaR), but standard deviation provided the best comparisons, especially as presented in the graphical analysis of the appendix.

for strategy selection, most notably expected retirement wealth, and that our basic conclusions are generally reliable.⁶

Within each family of matched glide paths, there is always a winner with regard to either median or mean target date wealth, but these attributes peak only weakly at the winning glide path, so there is little compelling reason to prefer them. Under an individual history, the winning glide path has a good chance of losing to another glide path in spite of its slightly better expected return. It should be stressed that a good performance in one glide path simulation does not guarantee a good performance for other glide paths in the same simulated history. Individual histories show a remarkable variation in terms of the best performing glide path of the family, often with either extreme exhibiting the best performance. We compare performance of glide paths under specific histories in section 4.

In addition to the base case, we present results for investors with an initial endowment fully invested at the start and no additional contributions. For stationary returns, a fixed investment over the time horizon dominates any glide path in all risk and return measures for the distribution. This is exactly because a consistent exposure to risk over time is the optimal path. Simulation results with the initial endowment contribution pattern are entirely consistent with this theoretical result.

We then test the sensitivity of the winning glide paths to varying market trends by running the simulations under different capital market assumptions. Market return assumptions which favor either stocks or bonds earlier or later in the investment period will tend to have winning glide paths which allocate more portfolio weight to the winning asset per time period. This sensitivity, along with the differences in glide path performance under individual simulations, means that choosing a particular glide path is equivalent to a bet that the returns will take a particular path.

III. Methodology

The base case was designed with simple but reasonably plausible assumptions, in order to observe the pure effect on retirement wealth of several assumptions and strategies, most notably the

⁶ In the single lump sum case, it is convenient to note that the wealth distributional properties of any glide path investment policy are identical to those of a single fixed risk policy over the investor's investment horizon. Wealth distributions result from compounding return over time. A constant compounding of return over an investor's investment horizon generally leads to highly right skewed wealth distributions. In this case a terminal wealth risk criterion will necessarily have to deal with the often misleading character of the mean for describing investor risk. This is because the mean of the income of 99 paupers and one millionaire may often inadequately reflect a summary of a wealth distribution. As a consequence the median of terminal wealth is often the criterion of choice to describe compound return that long-term investor experience. Michaud (1981, 2003) showed that the median of compound terminal wealth is often well approximated by the n th power of the mean of the geometric mean distribution. This framework is analytically convenient for understanding the distributional properties of the lump sum glide path wealth distribution. In the more general case of 401K investing, the investor is often encouraged to provide periodic additions to investments resulting in a wealth distribution that is the sum of many lump sum compounding return wealth distributions. There is no simple analytical framework for analyzing such wealth distributions and simulation remains the only solution. We report median wealth results but find that, for our assumptions, the mean of retirement wealth moves nearly parallel to the median and generally provides similar conclusions in comparing the glide paths.

glide paths. All of the examples are far simpler than any typical real-world retirement investment plan, but sufficiently complex to illustrate several important concepts. We assume a simple portfolio with two assets, intended to stand for a stock fund and a bond fund. At first we assume an independent and identical distribution for each time period. This distribution was selected to approximately match long term total returns of typical stock and fixed income funds in the US: expected annual returns of 10% and 3% and standard deviations of 17% and 5% for stocks and bonds, respectively, and a correlation of 0.1 between the two assets.⁷ Our base case is matched to a static portfolio of the two assets through the entire investment period, choosing a portfolio of 40% stocks and 60% bonds as a base case.⁸ These simulation parameters allow for a wide range of outcomes, including negative returns for both stocks and bonds, occasional extreme downside returns, and both bull and bear markets. Although the expected returns match the historical nominal and not the real returns for these assets, the conclusion is exactly the same for either choice – there is a wide range of glide paths with the same risk level all with roughly equivalent retirement wealth distributions. The terminal wealth for the real returns case would be less than for the nominal returns, but the conclusions are unaffected.

For the purpose of maintaining a comparable scale across experiments, we arbitrarily set the initial investment amount at 10, corresponding to an initial \$10,000 contribution to the plan, growing by 2% for each yearly contribution. Of course a different starting amount would simply change the scales on the graphs, but not the conclusions.

We solve for the matching glide paths using numerical grid search and bisection algorithms. We then can compare the various glide paths with the expected or median target date wealth, or any other statistic, to investigate the impact of changing the glide path, with the retirement risk amount held constant. Among the glide paths, there is a winner in terms of expected or median wealth for all the cases we studied.

After examining the results with these baseline assumptions in the next section, we systematically vary the assumptions and examine how glide path results are affected.

IV. Results

Figure 1 is a plot for the base case simulation, overlaying the family of matched glide paths for the 40/60 static investment plan with summaries of their target wealth distributions for each glide path. The simulations were performed with the base assumptions in Section 3. The glide paths were matched to a standard deviation of target date wealth of 953.68 – the same as the 40% stocks, 60% bonds static investment plan. For example, the leftmost linear glide path begins with 98.04% stocks, ends with 0% stocks and results in the same standard deviation of target date wealth as the 40/60 static plan. Starting with a higher stock percentage will yield a retirement

⁷ We simulate returns using a Student t distribution with 5 degrees of freedom. We make no claims that this model for returns is a perfect representation of real stock and bond returns, only that this distribution provides suitably fat tails to represent a broad array of realistic market scenarios, including many cases of substantial downside risk.

⁸ The robustness of the results to various stock and bond mixes was of particular concern. We tried several different stock/bond mixes, and found that a 40% stocks, 60% bonds provided a nice range for comparison because glide paths generally match to a lower stock ratio than their average stock ratio over time, but the analysis is extremely similar for 50/50 or 60/40 stock/bond ratios. We present a 60/40 analysis in Figure 6.

wealth with higher standard deviation for any linear glide path. The rightmost glide path begins with 0% stocks, ends with 59.48% stocks and also results in the same standard deviation of retirement wealth. Schematically, every such linear glide path is shown by an arrow in the plot, with its tail and head corresponding to starting and ending stock allocations given in the legend on the right hand side of the plot. The heavy blue line shows the expected retirement wealth under each glide path, and the other roughly parallel curved lines show the quantiles of retirement wealth as annotated in the legend below the plot. Corresponding wealth values are shown on the left vertical axis. The static 40/60 investment over the period is marked with a red vertical rule, and corresponds to an arrow of zero length. The mean target date wealth attains its maximum within the matched family with a gently descending glide path, starting at 53.24% stocks and ending at 32.38% stocks at retirement, although its mean is only slightly better than the other glide paths'. The median target wealth is maximized for an even more gently descending glide path that starts at 48.29% stocks and ends at 35.31% stocks. Typical industry glide paths which start with a large allocation to equities and end with a small one fare somewhat worse than either the optimal glide path or the static 40/60 investment.

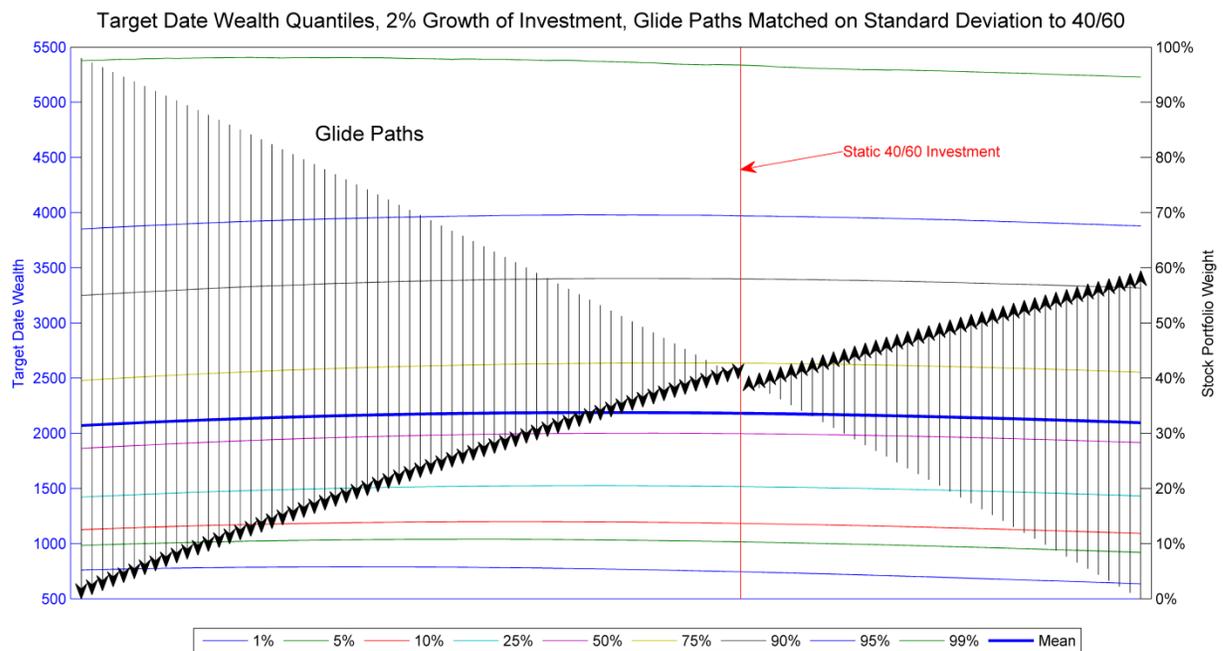


Figure 1: A range of glide paths, shown by black arrows, matched to a 40/60 static investment path. Each black arrow represents a glide path. The direction of each arrow is the direction of time, and its percentage of stocks at beginning and end of its glide path can be read from the tail and head of the arrow's position relative to the right vertical axis. Each glide path's expected return (mean) along with the labeled quantiles can be read by corresponding the curves at each arrow's horizontal position with the left vertical axis. Of the 101 glide paths shown on this chart, the mean target date wealth is maximized for glide path 50 (53.24% to 32.38%) and the median is maximized for glide path 55 (48.29% to 35.31%).

In spite of the relative consistency of target date wealth distributions across matched glide paths, choosing any particular glide path can have a large impact on retirement wealth under a particular history of returns. Figure 2 shows the wealth outcomes of 20 individual simulated histories, for the same glide paths of Figure 1. Clearly there is a wide range of outcomes, and the best performing glide path is nearly impossible to discern from this picture. We note that while the static 40/60

portfolio is not usually the highest return under any particular scenario, it is never the lowest. A particular glide path provides no guarantee of strong investment performance relative to other glide paths, and the more extreme glide paths are more exposed to extreme outcomes near the peak of their equity exposure.

Another notable case is where only one contribution is invested over the entire period. Here, the sequence of returns is of no consequence and every return has equal impact on wealth at the final date. The static investment is optimal from a theoretical perspective, and the simulation results in Figure 3 agree with the theoretical results. Glide paths can only be preferable to static portfolios based on patterns of investments or patterns of returns.

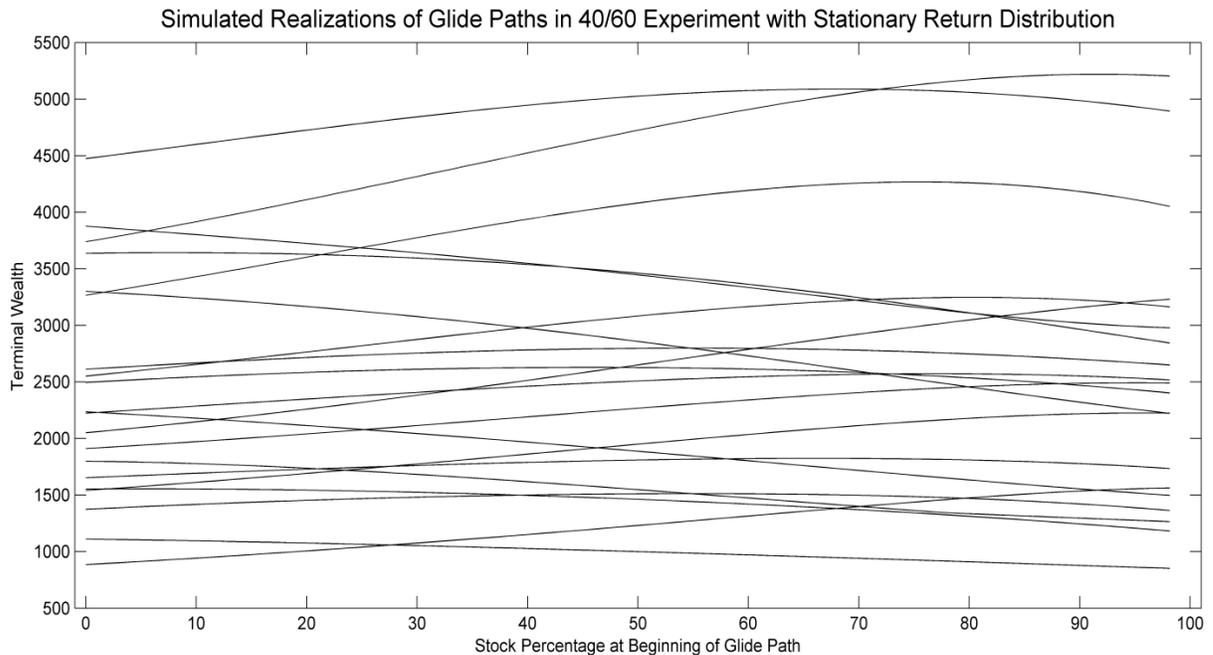


Figure 2: 20 simulated histories of returns, applied to the glide paths from Figure 1.

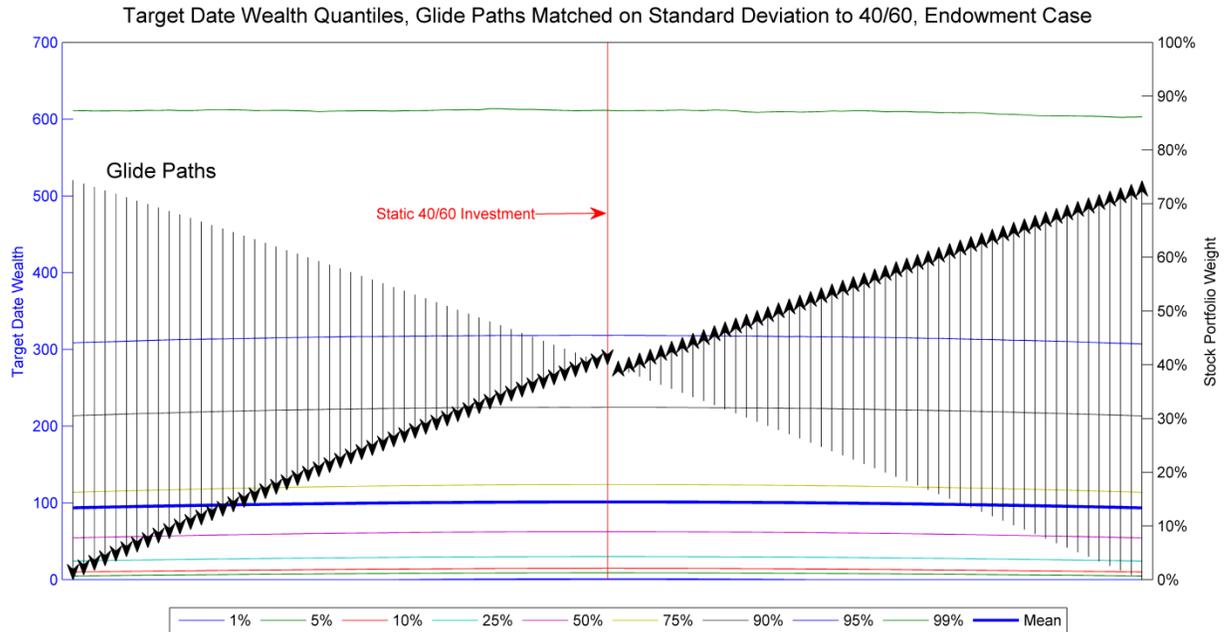


Figure 3: Matched Glide Paths for the 40/60 portfolio for an endowment investor, where no contributions are made after the first one. The static investment is optimal for both target date mean and median wealth in this case.

A. Robustness to changes in the return distribution assumptions

We first investigate changes in assumptions which motivate different choices of glide paths. First, we present two other return distribution scenarios. These are admittedly contrived, but demonstrate the sensitivity of glide path wealth to outperformance of an asset class early or late in the savings period. These return distribution assumptions are different from the base case only in the expected return of the equity asset. The first variation is a decreasing expected equity return, linearly from 8% to 2% annualized return over the investment period. The second case is a reversal of the first, with expected stock returns increasing linearly from 2% to 8% over the savings period.

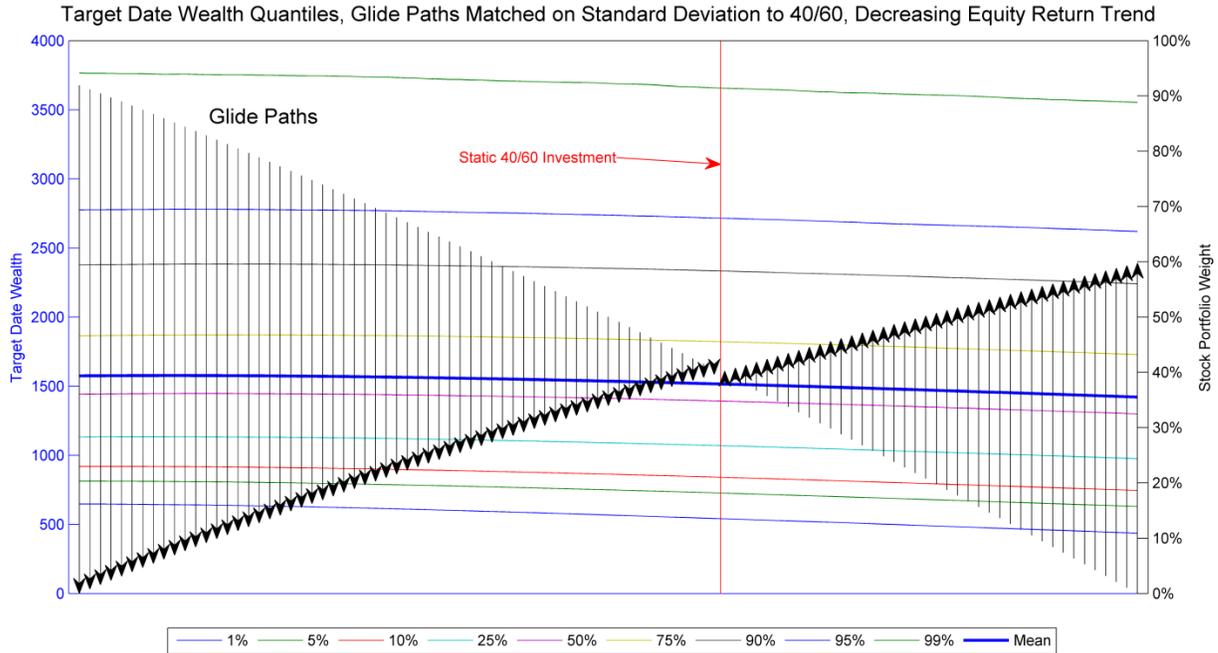


Figure 4: The simulation of Figure 1, modified so that expected stock returns fall gradually from 8% to 2% over the investment period. The decreasing glide paths become more attractive under this scenario. Of the 101 glide paths in this chart, the mean target date wealth is maximized for glide path 10 (85.21% to 6.95%) and the median target date wealth is maximized for glide path 13 (82.88% to 9.17%).

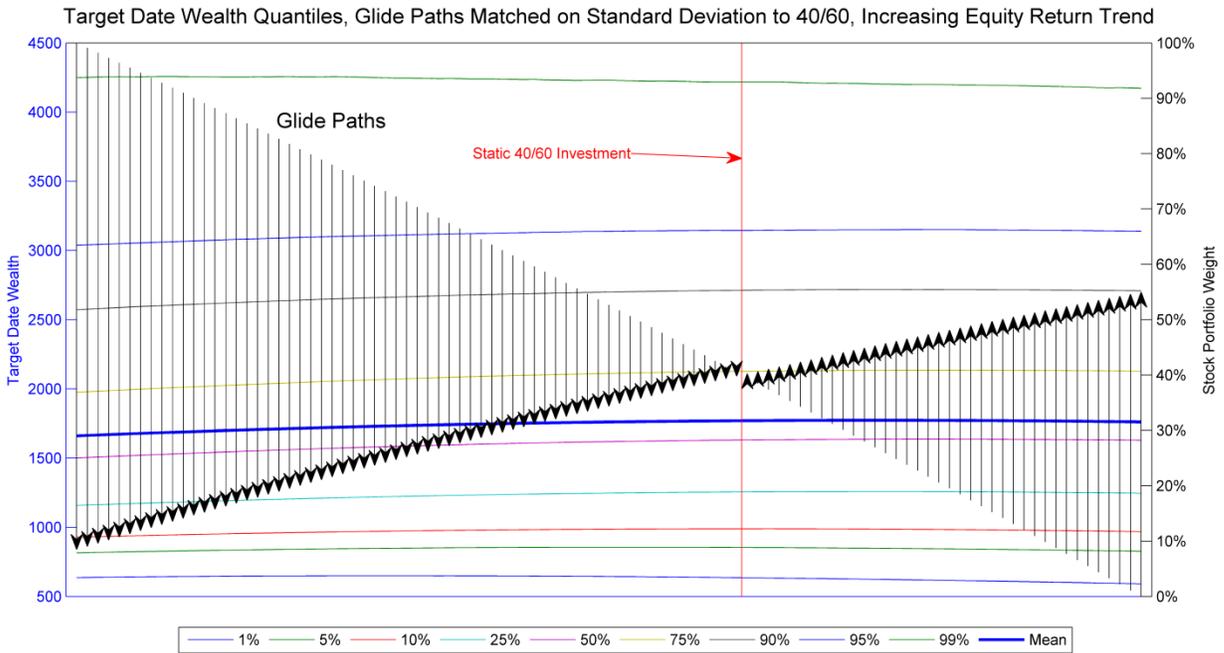


Figure 5: The simulation of Figure 1, modified so that expected stock returns rise gradually from 2% to 8%. This return trend shifts favor toward glide paths with increasing equity exposures over the investment period. Of the 101 glide paths in this chart, the mean target date wealth is maximized for glide path 76 (26.96% to 45.31%) and the median target date wealth is maximized for glide path 82 (20.61% to 47.75%).

The performance charts for these two cases are shown in Figures 4 and 5. The increasing returns scenario favors an ascending glide path, whereas the descending returns scenario favors a descending one. This is completely consistent with the intuition that the investor wants exposure to better performance when it is likelier to happen.

Clearly, different trends over the investment period support different glide path choices. The optimal glide path is quite sensitive to subtle variations in return assumptions and investment changes over time. Perhaps these examples are not inconsistent with the realities of many potential TDF investors. Along with the evidence in Figure 2, we find these modifications to the base simulation provide a persuasive case that choosing a particular glide path amounts to a bet on the relative performance of the constituent assets over time.

Because of space limitations, we do not show other experiments which tested the changing risk patterns to allow for more volatility, reducing expected equity and/or bond returns, allowing for a "turbulent" period of several years at either the beginning, middle, or end of the investment period, or any of several other modifications to the assumptions. Although the optimal glide paths are affected by most of these changes, none of them strongly favor any glide path and our overall conclusions remain unaffected by these sensitivity analyses.

B. Historical Confirmation

In general, we prefer simulation to an historical test, since history is only one realization of many possibilities. However, an examination of similar strategies over history may provide reassuring validation of the simulation framework. For this purpose, we use inflation-adjusted stock and bond returns from 1928 to 2012.⁹ Glide paths pose an additional problem for testing due to the lack of available independent 40-year windows on which to test. International markets don't help either—few markets have more than one 40-year observation, and international markets are correlated. We chose not to use the 85 observations in a single series due to the asymmetry in the importance of the data—the middle years from 1967 to 1973 would be represented in 40 experiments, whereas the first and last year in the sample would only be used once. To overcome this, we extend the data beyond 2012 reusing the period beginning with 1928 so that the 85 annual observations are used with equal importance. We note that history has been favorable to stocks, and while we believe the example is instructive, we do not believe more equities results in more wealth for all future scenarios.

We apply history to an annual investment starting at \$10,000 and increasing by 2% per year¹⁰ to be consistent with our simulations. We then analyze two static investments, 40% and 60% equities, and corresponding glide paths chosen to be superficially risk-equivalent, from 80% to 0% and from 100% to 20% equities with the remainder invested in bonds. This analysis corroborates our simulation experiments—glide paths have no special ability to reduce risk at their target date. If anything, history is harsher than our simulations on extreme glide paths. For example, compared to the 40/60 portfolio, a glide path from 80% to 0% equities results in both lower mean/median wealth at 65, as well as greater standard deviation and an inferior worst period. Someone who

⁹ Ten year treasury returns and CPI data are available from the Federal Reserve database (FRED). Equity returns are available through CRSP.

¹⁰ The total real investment over 40 years is \$604,020.

invested in a 40/60 portfolio from 1941 to 1981 – its worst 40-year period – would see real wealth increase only to \$675,000, whereas the worst case for the glide path investor is from 1940 to 1980 and results in a real wealth of \$632,000. The results for the four cases are summarized below.

	Mean	Median	Standard Deviation	Worst Period
40% Equities	\$1,580,000	\$1,483,000	\$501,000	\$675,000
80->0 glide path	\$1,387,000	\$1,273,000	\$547,000	\$632,000
60% Equities	\$1,941,000	\$2,020,000	\$555,000	\$913,000
100->20 glide path	\$1,729,000	\$1,559,000	\$565,000	\$919,000

V. Discussion

The marketing pitch for glide path funds may implicitly endorse faulty logic which assumes that risk taken earlier in an investor's lifecycle is absent from the retirement wealth distribution. The metaphor for a glide path is an airplane flight – cruising at altitude is like a higher-risk portfolio, and the smooth landing represents a gradual transition to a lower risk portfolio. However, the airplane metaphor is misleading because it implies a known destination and deterministic outcomes of cockpit controls. Our investigations show that a glide path with decreasing exposure to equities does not offer the claimed protection from risk over other investment schedules, and essentially bets on the timing of the relative performance of assets. Furthermore, the optimal glide path is sensitive to the pattern of contributions.¹¹

The choice which matters most, and has the most impact on the retirement wealth distribution, is the overall risk appetite of the investor. A 60/40 stock/bond ratio is riskier than a 40/60 mix, but the 60/40 portfolio is also rewarded with better expected and median returns in our simulations. The chart for a 60/40 investment is shown in Figure 6. The scale of the wealth distributions now extends to much greater values. The standard deviation of retirement wealth is approximately 2,042, as compared with 953 for the 40/60 portfolio. The increase in risk comes with a substantial reward, as the range of expected retirement wealth amounts across glide paths is well outside the similar range for the 40/60 investment. However, the lowest quantiles of the 60/40 matched cases do worse than those of the matched 40/60 cases, so there is risk associated with the reward of greater median and expected retirement wealth. The main conclusion from the comparison is that the overall choice of risk level, proxied here by stock/bond ratio, has far more impact on the retirement wealth outcome than the choice of glide path. It is well known in finance that the ratio of stocks to bonds is the most important determiner of portfolio performance (e. g. Brinson et al 1986, 1991), and our simulations add further confirming evidence to this hypothesis.

¹¹ Perhaps the strongest criticism of glide path investing is that few investors follow the steady contribution plan that glide paths assume they will be making.

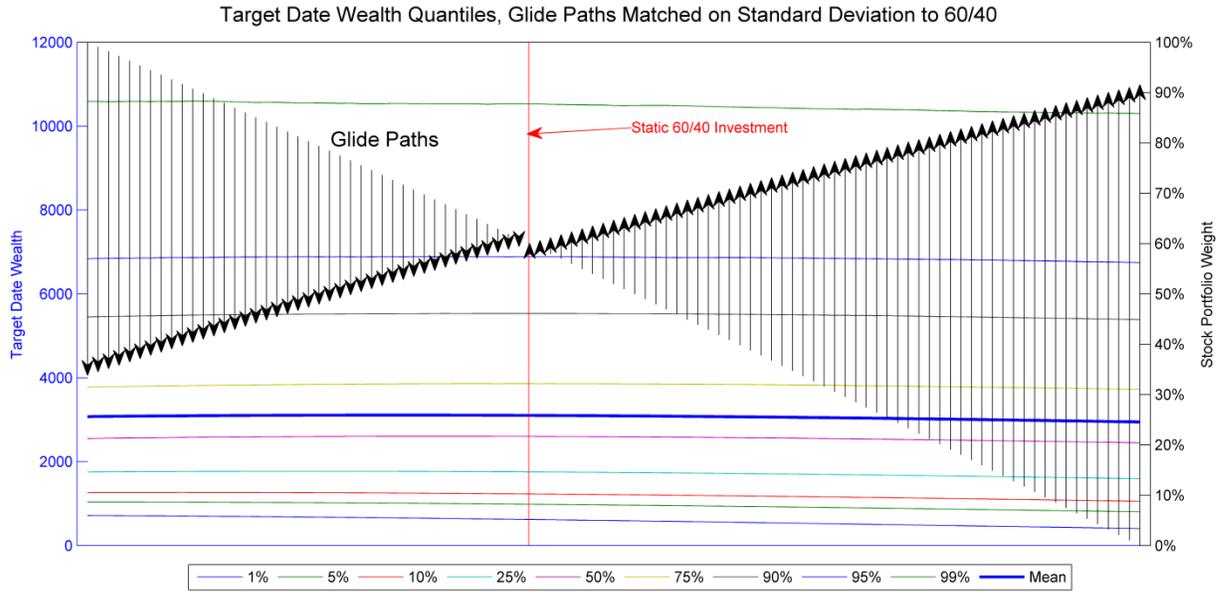


Figure 6: The experiment of Figure 1, repeated for a riskier 60/40 static portfolio. All conclusions remain intact from before, but the scale of the retirement wealth and its risk are greater. Of the 101 glide paths in this chart, the mean target date wealth is maximized for glide path 30 (72.67% to 52.25%) and the median target date wealth is maximized for glide path 34 (68.78% to 54.68%).

We have provided evidence, through our simulation experiment, that the optimal glide path at a particular risk level is highly dependent on assumptions of both market returns and investment contribution pattern, but there are many glide paths going in both directions that match this optimal one in risk characteristics and nearly match in target date wealth distributions. Investors are likely better served by choosing a default risk approximating the market with a 60/40 or 50/50 stock/bond ratio, rather than predetermining a glide path that may not be a good fit later in the investment period when either market or personal circumstances change.

VI. Conclusions

We tested a variety of return assumptions, investor savings schemes, risk tolerances, using simulated and historical data. We examined both average statistics (such as mean and standard deviation) as well as percentiles (such as the median or 95% lower bound) to measure risk and return. In none of these comparisons did we find compelling evidence for outstanding superiority of a particular glide path. Instead, we found a fixed stock/bond ratio investment to be surprisingly robust to potentially uncertain market conditions. We are concerned that retirement investors, who may have little information or expertise about investing, are being steered into a false sense of security investing in glide path funds. We have presented evidence that even a simple glide path is no guarantee of lower risk at retirement, including an historical analysis showing the 80% to 0% glide path worsened both risk and return compared to a static 40% stock strategy – the standard deviation of wealth being \$47,000 higher and the mean wealth being \$193,000 lower. The highly sloped glide paths popular in many target date funds fare little better in our simulations.

Default glide paths correspond to a particular risk tolerance, which may not match the investor's circumstances or preferences faithfully, especially over time. The true level of risk taken over the

course of savings is somewhat obfuscated by the glide path, and the very fact that it can be matched in risk characteristics with many other paths without greatly harming the retirement wealth distribution shows the arbitrariness of the glide path selection. Offering glide path funds as a qualified default alternative may give a false impression that fiduciary review has been completed. We believe this is bad policy and may constrain cross-sectional suitability of the portfolio.

Empirical evidence from real-world TDF investors also bears out the possibility that glide path funds are not always matched well to market climates – TDFs can lose value just like any other funds if they are exposed to assets with falling prices.¹² Losing one's retirement savings as the target date approaches is a shame, especially if a prudent review of the portfolio or greater diversification could have prevented some of the loss. For investors sticking to a long term and consistent investment plan, TDFs have dubious value. For those investing a lump sum, they're inferior to an equivalent target risk portfolio.

¹² E. g. Silverblatt 2009

Appendix

I. Concerns about matching on standard deviation

Wealth distributions tend to exhibit positive skewness, with a few wealthy individuals exerting disproportionate influence on the mean and standard deviation. If the matched glide paths shown, for example, in figures 1, 3, 4, 5, and 6 have non-matching skewness, the comparisons may be invalid. To investigate whether the standard deviation matched comparisons are valid, we performed both a visual exploratory analysis and a computation of skewness values for all the glide paths. The glide paths shown in Figure 1 have computed skewness values from 1.97 for the maximally descending glide path to 1.49 for the maximally ascending glide path. The static 40/60 portfolio has a skewness of 1.53, and the glide path with maximum expected return has a skewness of 1.58. There are clearly differences in skewness among the glide paths. However, it is unclear if they are substantial enough to invalidate the comparisons since they are not great enough to clearly sway an investor's preference, given the differences in expected returns. To investigate further, we construct several diagnostic plots showing the simulated retirement wealth amounts of the maximally ascending, static 40/60, and maximally descending glide paths. These plots appear in figure 7. They are arranged in a matrix, with all horizontal and vertical axes aligned. On the diagonal of the matrix appear labels and a histogram showing the simulated retirement wealth amounts for the labeled glide path. Each off-diagonal axis plots the two glide path retirement wealth simulations on the corresponding diagonal labels against each other. The upper three plots are matched on simulated outcomes – each point on the plot underwent the same history of returns under both glide paths. The lower three plots are matched on distribution quantiles, i. e. the simulated wealth amounts for each glide path are sorted before plotting. These quantile-quantile plots in the lower three axes demonstrate any differences in skewness among the plotted elements. If one distribution is more skewed, the blue path of quantiles curves away from the red diagonal rule in the direction of the more skewed sample. For example, it can be seen that the maximally descending glide path is slightly more skewed than either of the others, matching the analysis from computing the numerical sample skewness from the simulations.

Simulated Terminal Wealth Distribution Comparison for Multiple 40/60 Glide Paths

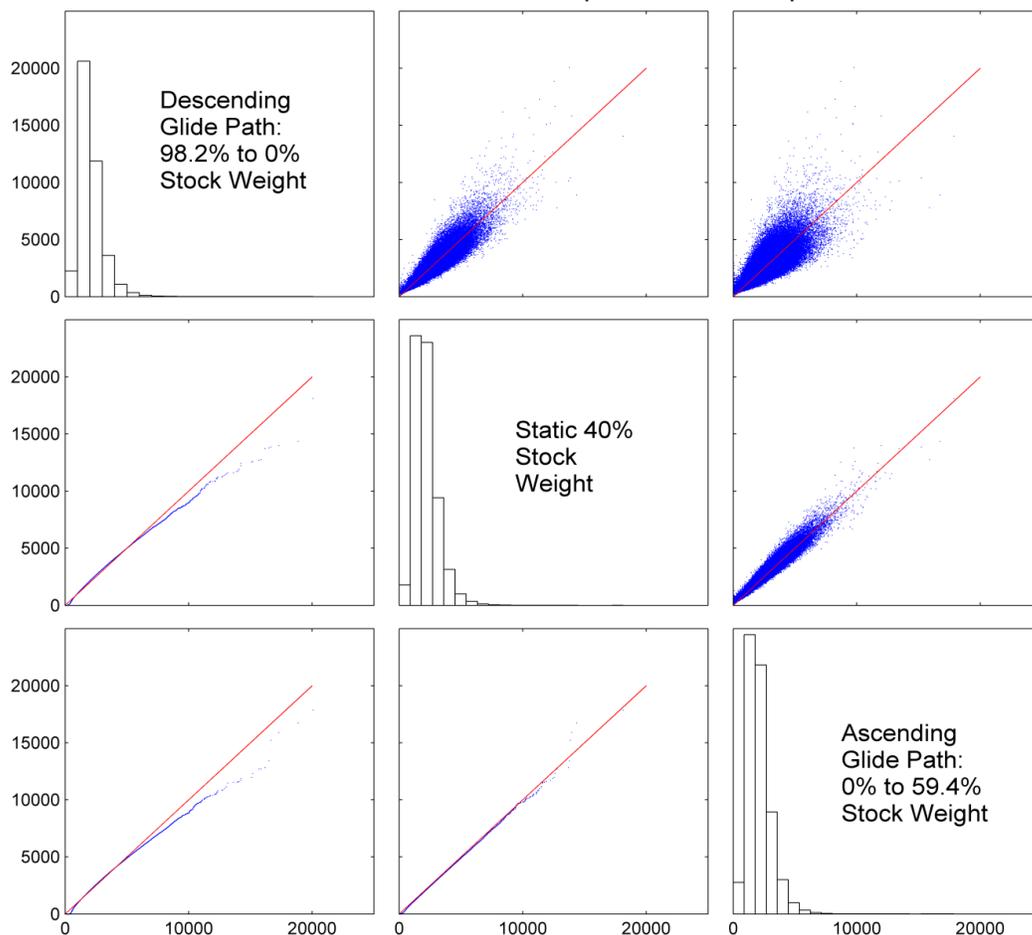


Figure 7: Some comparisons of the final wealth distributions of 3 glide paths from Figure 1, as annotated. The lower panels plot quantiles of wealth within glide paths against each other, and the upper panels plot simulated outcomes directly, not sorted by quantile.

The matrix of plots does show slight differences among the three cases, but given that the majority of these differences occur in extreme cases and the expected return is substantially different for these cases, it is our opinion that the differences in skewness are far outweighed by the major differences in expected return. The skewness differences shown in either the numerical or the graphical analysis here are not great enough to be of concern. Other risk measures were considered for matching as well, but none matched as consistently as standard deviation and they all fared much worse in as far as overall differences in final wealth distribution.

Figure 7 also shows the lack of perfect correlation in terms of wealth outcome offered by the different glide paths. The upper three off-diagonal plots fan out considerably from the diagonal lines representing equality between horizontal and vertical axes. The shape of the clouds of points mean that many good outcomes under one glide path, within a particular history, would have been bad outcomes had another glide path been selected.

References

- Brinson, G., Hood, R., and Beebower, G. 1986. "Determinants of Portfolio Performance." *Financial Analyst Journal* 42(4): 39-44.
- Brinson, G., Singer, B., and Beebower, G. 1991. "Determinants of Portfolio Performance II: An Update." *Financial Analysts Journal* 47(3): 40-48.
- Investment Company Institute. 2013. 2013 Investment Company Fact Book: A Review of Trends and Activities in the U.S. Investment Company Industry. 53rd edition.
http://www.ici.org/pdf/2013_factbook.pdf
- Markowitz, H. M., & Usmen, N. 1996a. The likelihood of various stock market return distributions, Part 1: Principles of inference. *Journal of Risk and Uncertainty*, 13(3), 207-219.
- Markowitz, H. M., & Usmen, N. 1996b. The likelihood of various stock market return distributions, Part 2: Empirical results. *Journal of Risk and Uncertainty*, 13(3), 221-247.
- Michaud, R. 1981. "Risk Policy and Long Term Investment." *Journal of Financial and Quantitative Analysis*. 16(2), 147-167.
- Michaud, R. 2003. "A Practical Framework for Portfolio Choice." *Journal of Investment Management*, 1(2), 1-16.
- Michaud, R. and Michaud, R. 2008a. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Oxford University Press, New York. 1st ed. 1998, originally published by Harvard Business School Press, Boston.
- Michaud, R. and Michaud, R. 2010. Target Date Funds Aren't a Panacea. *Investment News*, May 2, 2010.
- Silverblatt, R., 2009. Target Date Funds Go Under the Microscope. *Fund Observer*, US News and World Report, Oct. 30, 2009. <http://money.usnews.com/money/blogs/Fund-Observer/2009/10/30/target-date-funds-go-under-the-microscope>