April 2018 | Putnam Perspectives

The glide path less traveled

A critical examination of assumptions, outcomes, and glide path specification

Brett Goldstein, CFA
Co-Chief Investment Officer, Global Asset Allocation

Key takeaways

This paper defines key assumptions, derived from empirical data, to deliver highly accurate and representative simulations of participant outcomes.

The paper constructs an alternative glide path design that we believe more efficiently balances major retirement savings risks — longevity/shortfall, return sequence, and inflation — than the industry average glide path.

The paper demonstrates that an alternative glide path — one that begins with higher equity allocations than the industry average, follows a steeper slope of asset shifts, and arrives at a risk-parity solution at the point of retirement — has less variability in outcomes than the industry average.

Executive summary

Successful glide path design must marry accurate estimation of quantitative inputs with efficient risk management. Using the most representative set of assumptions for the broadest segment of the plan participant demographic, we employ simulations to relate these assumptions to outcomes for the industry average glide path.

Our philosophy on glide path construction is that the main goal of a glide path is to provide adequate wealth throughout retirement. For this reason, our simulations target whether or not an individual would have enough retirement income to reach life expectancy age with a non-zero portfolio balance.

Longevity risk and shortfall risk are primarily the result of not accumulating enough retirement savings, either as a result of insufficient savings or inadequate risk-taking. Sequence risk, on the other hand, is the result of taking too much risk, allowing the order of market returns to damage a portfolio that is in the process of de-risking. These risks are constantly at odds with each other, and a well-designed glide path must balance these risks appropriately throughout the glide path. We also consider the impact of inflation risk, but we find little compelling evidence that this should play a major role in glide path design, aside from using real returns, rather than nominal returns, in our simulations.

Despite high success rates for the industry average glide path, we find the simulated results to suggest potential inefficiencies in approaching the retirement savings problem. The industry average glide path does too little to combat longevity/shortfall risk early in the glide path, and it fails to fully defend against sequence risk late in the glide path.

We provide an alternative glide path specification, which is steeper than the industry average. Our alternative is fully invested in equity at the start of the glide path, has a higher equity allocation for 25 years, and then more steeply reduces the equity allocation to arrive at a risk-parity solution at maturity.

The industry average, relative to our calculated alternative, has a slightly higher probability of success, along with higher mean and median values at retirement. But those gains are the result of more variability in the distribution of portfolio values at retirement. Much of the higher variability manifests with higher portfolio values on the “right” half of the distribution, where retirees will easily be able to meet their need in retirement, but with that higher variability also comes a higher probability of landing a plan participant in the far “left tail” of the distribution — a poor tradeoff in our opinion.

For this reason, we find solace in the lower variability of outcomes that result from the alternative portfolio. Much of the probability distribution is clustered around an acceptable outcome, and even if the retirement portfolio value comes up short, it is likely to be only a little short. In these scenarios, slight changes in retirement spending can easily make up for the savings shortfall. Only the most significant deficits, which are more likely to occur in the higher-risk glide path, leave unfortunate retirees with few options.

Finally, it is worth considering the outcomes outside the scope of a binary simulation. In reality, failure is not binary, and no matter the outcome, plan participants must make the best of their situation.

Introduction

Retirement planning is a daunting task. The complex relationships — and the uncertainty — between assumptions such as capital markets returns, life expectancy, savings rate, income replacement rate, and real income growth rate are more than enough to make a plan participant anxious. And while existing analysis has yielded some general rules to follow, we believe there is a dearth of research identifying the right assumptions to use and relating these assumptions back to the construction of a glide path.

We recognize that assumptions might change over time. Annual Morningstar Target-Date Fund Landscape reports[1] show that U.S. target-date mutual fund assets have grown to almost $900 billion at the end of 2017 from a base of less than $100 billion in 2005. Throughout this period of rapid growth, capital markets and investment management in general, and retirement products in particular, have changed drastically. In fact, at the start of that period plan participants could not be auto-enrolled nor was there the ability to auto-escalate savings rates, both of which were enabled by the Pension Protection Act (PPA), passed in 2006. The PPA has proven to be a watershed moment in the development of retirement savings, and it has empowered plan sponsors and plan participants to better prepare for a successful retirement.

In the following paper, we seek to determine what the right assumptions are for glide path design *today*. Once our set of baseline assumptions has been selected, we estimate outcomes and relate these assumptions to the risks that must be addressed by the glide path. Next, we consider where the investment manager, as it relates to the shape of the glide path, has the greatest ability to improve the outcomes for the plan participants. Finally, we identify additional considerations that are difficult to simulate but important to incorporate into retirement planning.

Background and philosophy[2]

Our philosophy of glide path construction is that the main goal of a glide path is to provide adequate wealth throughout retirement. We therefore define "success" as the ability of a plan participant to reach life expectancy with a positive portfolio balance. Thus, we are indifferent with respect to higher portfolio values if the probability of reaching life expectancy with a positive portfolio balance is essentially unchanged.

The assumptions driving glide path construction are of critical importance, so they should represent the greatest proportion of the population that will engage with a glide path. In this regard, we seek out the central tendency of a given distribution when defining our glide path assumptions, as this should logically be the basis for the strategic glide path. We expect that deviations from these assumptions will be minimized overall by using their central tendencies and will have a limited impact on the outcome of a plan participant, recognizing, of course, that a glide path is neither intended nor suited to meet the needs of the individuals in the far tails of a distribution.

1 Bruns et al. *The Glide Path Selection Problem*. Morningstar. February, 2016.
Holt et al. *2017 Target-Date Fund Landscape*. Morningstar. April 20, 2017.

2 This paper assumes some familiarity with glide path research. For background on these topics, please see "Optimizing the glide path for a smoother landing in retirement" https://www.putnam.com/literature/pdf/DC976.pdf.

It is also important to recognize that some aspects of retirement planning are outside the control of the glide path and of the portfolio manager, such as plan sponsor match incentives, participant savings rates, and income or debt outside of retirement accounts. No matter the choice of glide path, individuals will never have enough wealth at retirement if they save too little. Instead, we must build the glide path to address the risks that *can* be helped by an effective glide path, and we have employed Monte Carlo simulations to allow us to compare the costs and benefits of the various options.

Assumptions

As stated earlier, the goal of a glide path is to provide adequate investment such that retirees do not consume all of their savings in their lifetime. Although we think of this as a singular goal, there are actually two embedded goals that must be considered separately: building enough wealth during the accumulation phase before retirement and retaining enough wealth during the distribution phase in retirement.

These two goals, of course, are not independent. So, while it is helpful to define a target for "enough" wealth at retirement, it must be understood that true success can only be measured by a non-zero portfolio value at the end of retirement or at full life expectancy. While it seems reasonable to attempt to calculate a probability of "success" based on portfolio value at retirement, this is not a fair metric as "failures" at retirement are not necessarily failures overall, and missing a target value by a small amount is very different from missing it by a lot.

In order to create a glide path that takes into account all of these considerations, the investment manager should combine market-based empirical research with simulations under different sets of assumptions. We also believe that the glide path should only change if there is a major adjustment to any long-term assumptions. It is important that the glide path not change on a regular basis, or in response to small shifts in assumptions, because plan participants should be able to understand and anticipate how their assets are invested. Instead, it should be the role of the investment manager to make these dynamic shifts. After all, a passive manager that makes regular changes to a glide path based on changing assumptions is essentially making active decisions.

Key assumptions for the simulations include market returns, retirement age, life expectancy, savings rate, income replacement rate, and real income growth rate.

Capital market assumptions

First, since a glide path should be based on long-term assumptions, we derive capital market assumptions using a comprehensive history of liquid financial markets.[3] As highlighted earlier, any short-term fluctuations in expected returns for stocks and bonds are best expressed as a dynamic allocation choice by the manager and not by redefining the glide path.

Second, in addition to using long-term assumptions for stocks and bonds, we also considered the best way to account for inflation. Inflation is a significant risk posed to retirees and those saving for retirement, and the impact of inflation must be carefully weighed when crafting a glide path. For this reason, we use real expected returns in our simulations. Below, we also include a comparison to excess returns, which tend to be closely related to real returns, since cash rates generally track inflation over time.

---

3 Stocks = S&P 500 TR (backfilled with IA SBBI S&P 500 TR USD prior to 2/29/1988), Bonds = Bloomberg Barclays U.S. Treasury 7–10 Year TR (backfilled with fitted blend of IA SBBI IT Govt. TR USD and IA SBBI LT Govt. TR USD prior to 2/29/92), Cash = Bloomberg Barclays U.S. Treasury Bills 1–3 Month TR (backfilled with IA SBBI U.S. 30-Day T-Bill TR USD prior to 2/29/92), and Inflation = U.S. CPI Urban Consumers NSA (backfilled with IA SBBI U.S. Inflation prior to 2/28/47). All data is updated through 12/31/17.

Figure 1. Historical asset class returns, 1926–2017

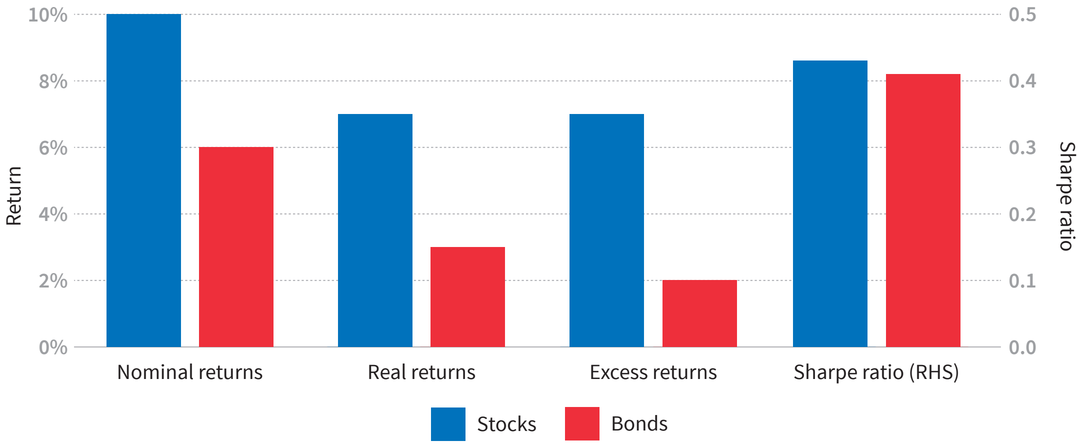

Source: Putnam.

Based on almost a century of capital markets data, we assume real returns of 6% and 2% for stocks and bonds, respectively, which corresponds to approximately a 0.4 Sharpe ratio. Intuitively, it makes sense that these asset classes should have similar risk-adjusted returns, as measured by Sharpe ratio, and this is also borne out by the data, as seen in Figure 1.

Retirement age and life expectancy

We use 67 for the retirement age, as this is the age at which any plan participant born after 1959 will be able to claim full Social Security Administration (SSA) benefits.[4] We assume an accumulation phase that is 45 years long. Our accumulation phase is this long because we see no need to alter the portfolio allocation more than 45 years prior to retirement, but we encourage plan participants to invest prior to the implied starting age of 22.

Assuming a life expectancy of 92, we obtain a distribution period of 25 years. According to calculations derived from the SSA Actuarial Life Table[5] (2014 data was the most recent data available at the time of publication), a 67-year-old male has approximately a 15% probability of outliving this life expectancy, and a 67-year-old female has approximately a 25% probability of outliving this life expectancy. Thus, a plan participant has roughly a 20% probability of outliving the life expectancy used in our assumptions. We also note that the probability of living one additional year declines rapidly after this point.

4 https://www.ssa.gov/planners/retire/agereduction.html.

5 https://www.ssa.gov/oact/STATS/table4c6.html.

Figure 2. Life expectancies are derived from Social Security Administration data

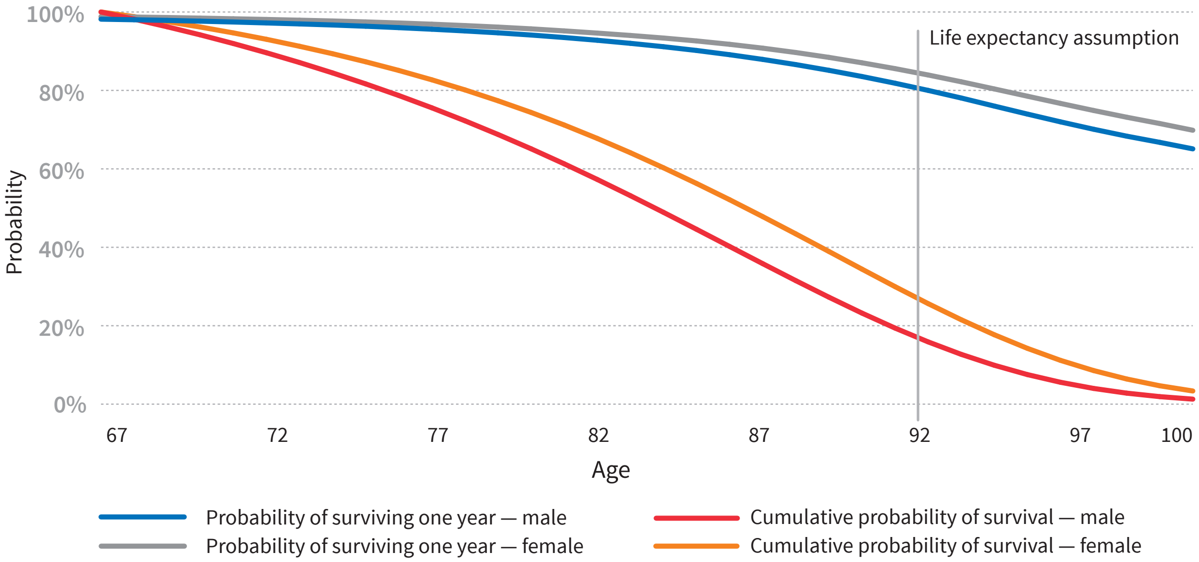

Source: Putnam, using SSA data.

Probability of surviving one year:	Cumulative probability of survival:
Probability of living one full year, conditional on current age	Probability of survival at a given age, conditional on reaching retirement age (67 years)

Savings rate

The savings rate assumption is also informed by empirical data, although in this case we rely on survey data. In an effort to minimize confirmation bias and sampling error, we consider multiple data sources.

The first source, the widely followed Retirement Confidence Survey[6] published by the Employee Benefit Research Institute (EBRI), finds that the total median savings rate (including employer match) for those workers currently saving for retirement is 10%. Interestingly enough, those surveyed estimated a median required savings rate of 16%, which could imply a willingness or intention to save more in the future.

The other source is the Lifetime Income Score, published by Empower Retirement (an affiliate of Putnam Investments) and conducted through their research partner, Brightwork Partners, and used with their permission. This survey, entitled the "Lifetime Income Study 6" (LIS),[7] is a comprehensive analysis of the behavior and expectations of over 4,000 working Americans. The survey tracks their savings behavior and estimates their overall ability to meet their financial needs in retirement.

According to the LIS, among households saving for retirement, the median total household retirement savings rate is 11%, with estimates ranging from 10% to 12% since 2011. If we include all households, and do not limit the sample to those saving for retirement, the median total household savings rate is 9%, with estimates ranging from 9% to 10% since 2011.

6 https://www.ebri.org/surveys/rcs/2017/.

7 Lifetime Income Study 6. Empower Retirement, conducted through Brightwork Partners. March 31, 2016.

As we mentioned at the outset, the goal of this paper is to identify what assumptions should be used on a forward-looking basis. While these surveys are helpful for understanding past behavior, we would be remiss if we completely ignored where savings rates might trend from here. For this, we use an age-based cohort analysis of the LIS data, since Millennials have had a materially different retirement savings experience than other cohorts (as a result of PPA passage in 2006), and they will define the trends moving forward.

Encouragingly, the youngest cohort (ages 18 to 29) has the highest absolute employee-only savings rate, with a 6.5% median savings rate, including individuals who make no contributions at all. There has also been a sharp uptick in participants who are auto-enrolled and auto-escalated, which helps to boost overall savings rates. As a cohort, the Millennials again lead the way, with more than half of the youngest cohort taking advantage of these features.

These findings suggest that 10% is our best estimate of the median combined (employee + employer) savings rate on a forward-looking basis. There is also reason to be optimistic that this rate might increase in the future, but we do not find this evidence to be convincing enough to argue for a higher savings rate assumption at this time.

One concern about savings rates we often hear is that many workers nationwide are already living paycheck to paycheck and cannot possibly increase savings rates, much less attain a 10% savings rate. We categorically reject this argument. First, we are hard-pressed to find any evidence that on the whole, median savings rates for plan participants are materially below 10%.

Moreover, we find evidence that those workers who do not save enough currently are likely capable of saving more. Specifically, we turn to the work of Benartzi and Thaler[8], who pioneered the use of behavioral economics applications for retirement savings. In their 2004 paper[9], they discuss the implementation of an auto-escalation program at a manufacturing company plagued by low savings rates. For those employees who met with a financial planner and agreed to participate in the deferred auto-escalation program, which increased participant savings rates by 3% each year coinciding with pay increases (so as not to decrease take-home pay), they observed an increase of more than 10% over four cycles. At the end of the study, those workers who joined their program increased average savings rates to 13.6% from only 3.5% before the start of the program. Other workers, who met with a financial planner but did not participate, saw savings rates decline to 5.9% from 6.1% initially. This is compelling evidence that unwillingness to save, rather than inability to save, is the greater impediment to higher savings rates in many cases.

Income replacement rate

Retirees typically rely on two sources of income in retirement: Social Security Administration (SSA) benefits and retirement savings. First, we will estimate the total required income replacement for a retiree. Then, we can calculate an estimated contribution from SSA benefits, which can be used to back into the ultimate burden on the retirement savings portfolio.

An SSA research paper[10] on income replacement ratios found that the median replacement ratio in the first two years of retirement was 73.5% based on shared household income. Beyond the first two years, median replacement ratios were found to decline, eventually falling below 60% more than six years after retirement.

8 Richard Thaler won the 2017 Nobel Prize in Economics for contributions to behavioral economics.

9 Benartzi and Thaler. *Save More Tomorrow™: Using Behavioral Economics to Increase Employee Saving.* 2004.

10 https://www.ssa.gov/policy/docs/ssb/v72n3/v72n3p37.html.

A Morningstar paper,[11] however, suggests the SSA paper might miss rising income replacement needs much later in life. Since annual real percentage change in consumption is negative on average, total real spending will decrease somewhat throughout retirement, but the overall decrease can be less so than implied by the SSA analysis due to flat-to-increasing costs as retirees move far beyond their retirement date. As a result, they caution that while a "rule of thumb" of 70%–80% is reasonable, true requirements can be anywhere from 54%–87% based on individual circumstances.

In contrast to these studies, the EBRI survey found that retirees reported higher-than-expected expenses in retirement. Almost half of retirees (47%) indicated that healthcare costs were higher than expected, with respondents more than three times more likely to face higher-than-expected healthcare costs than lower-than-expected healthcare costs. More than a third of retirees (37%) found non-healthcare costs to be higher than expected, making them four times more likely to face higher-than-expected non-healthcare costs than lower-than-expected non-healthcare costs.

Given the propensity for declining expenditures in retirement, but with the understanding that evidence also indicates higher-than-expected expenses, we choose to use a constant withdrawal rate throughout retirement. Since our withdrawal rate will remain constant, we select the lower end of the estimated range, as we are concerned that a higher constant rate could potentially result in unintended bias. With uncertainty around life expectancy and only a modestly declining income replacement rate in the studies we read, we believe this is a reasonable tradeoff to support our final assumption of a 75% income replacement rate.

Once the overall income replacement rate has been determined, SSA benefits must be estimated. Using median household income and assuming a lifetime income stream similar to that which is estimated in our simulations (further detail in the next section), SSA benefits will contribute approximately 35% of ending salary. This would reduce the income needed from the retirement portfolio to only a 40% replacement rate.

Figure 3. Income replacement rate from Social Security, by income level

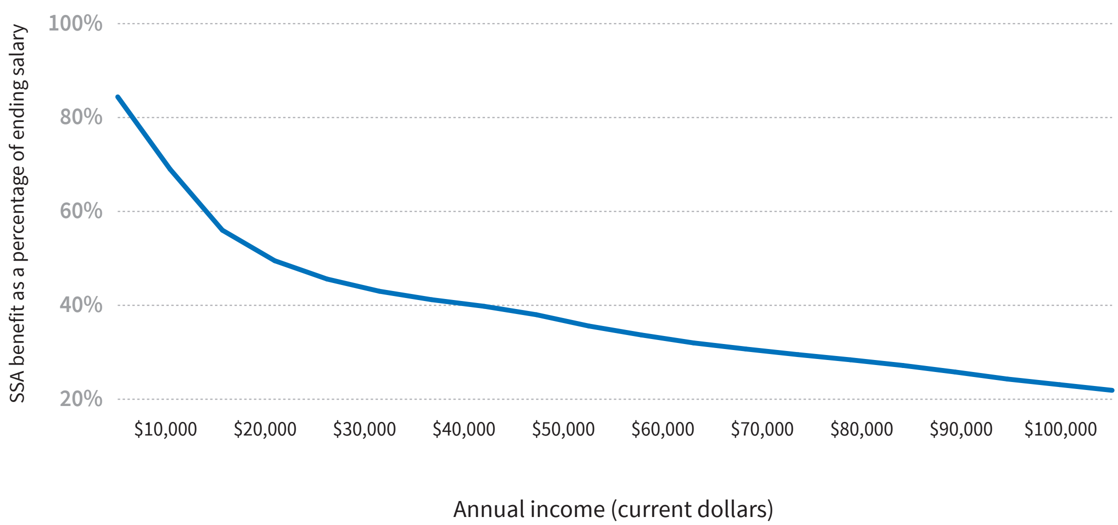

Source: Putnam.

11 Blanchett, David. *Estimating the True Cost of Retirement*. Morningstar. November 5, 2013.

The inclusion of Social Security also somewhat dampens the impact of the income replacement rate assumption, as individuals who will realize a lower income replacement rate in the form of SSA benefits will be higher earners, who are more likely to require lower income replacement, according to the SSA paper cited above. Of course, this all assumes that the plan participant begins to claim SSA benefits at full retirement age, but participants can claim up to a 24% increase in the amount of their benefits by postponing benefit claims until age 70.[12]

One important caveat to our SSA benefits calculation is that we assume full SSA benefits distributions persist without modification in the future. Currently, SSA benefits can be paid in full thanks to trust fund reserves that sustain deficits. The summary of the SSA Trustees Report[13] estimates these trust funds will run out in 2034, resulting in a 23% reduction to SSA benefits paid out at that time.

A recent book[14] addressing the state of retirement savings in America identifies a return to solvency for Social Security as one of three pillars necessary to ensure a successful retirement for workers. Cuts to social insurance programs have been difficult to implement historically, and in the event action is ultimately taken, it is typically phased in slowly, as evidenced by the increase in full retirement age for SSA benefits. Therefore, we deem it unlikely that workers will truly experience a "Social Security cliff" in 2034, and we choose to base our assumption on SSA benefits as currently promised. We also note that the risk of assuming full SSA benefits is dampened by the fact that younger workers appear to fall on the high end of our savings rate assumption, so lower income replacement from SSA benefits could be somewhat tempered by higher income replacement from retirement savings.

Finally, as a check on the studies using observed data, we also calculated a theoretical after-tax replacement rate. We attempt to isolate the impact of retirement on after-tax income, without making assumptions about changes in various deductions or the impact of state taxes on retirement income distributions. Before-retirement "net" income available to a plan participant is therefore calculated with respect to taxable income, federal income tax paid, FICA tax paid, and retirement savings. Post-retirement "net" income is calculated as a function of SSA benefits, retirement savings distributions, and federal income tax paid. These calculations were performed for the same set of incomes used in the SSA benefit calculations above, and we found that a 75%–80% replacement rate in retirement would generally result in a full replacement of after-tax income from a plan participant's working years.

Real income growth rate

In order to accurately estimate wage growth on an individual basis, we again turn to SSA data. While there is readily available data on median income through time, this is a flawed measure of wage growth for our purposes. Median income is a measure of the distribution of income at a given time, and therefore can be affected by shifting demographics and other external factors. For example, median income for the entire population can remain unchanged over time, even as workers age and advance in their careers, providing positive real income growth on an individual level. To accurately represent an individual saving for retirement, our analysis requires a longitudinal growth rate, which follows individuals from the start of their careers to retirement.

---

12 https://www.ssa.gov/planners/retire/delayret.html.

13 https://www.ssa.gov/oact/trsum/.

14 Reynolds, Robert. *From Here to Security: How Workplace Savings Can Keep America's Promise*. 2018.

Our analysis finds a positive but decreasing real income growth rate for working individuals.[15] This assumption diverges from many existing studies, which typically use a constant real (or nominal) growth rate. Unfortunately, while a constant real growth rate is simple to apply, it is far removed from the reality of empirical longitudinal real growth rates, introducing significant bias into simulations and leading to unrealistically low success probabilities. Based on our analysis, we select a 3% exponential real growth rate with a 10-year half-life to determine the assumed path of real wages.

We found additional evidence of a decreasing real growth rate in a N.Y. Fed blog post that attempts to estimate real wage growth for U.S. workers.[16] This analysis, which uses Current Population Survey (CPS) data from the Bureau of Labor Statistics (BLS), makes remarkably similar conclusions about the trajectory of real wages throughout a worker's lifetime.[17] It goes on to analyze various demographic cohorts within the population, but the general shape of real wage growth appears very similar across cohorts.

Our income growth rate assumption attempts to find a happy medium between assuming a constant real growth rate, which will bias success probabilities lower, and assuming rapidly decreasing real income growth curves such as those found in the N.Y. Fed analysis, which could be more prone to overstating success probabilities. The most important piece of the assumption, in our opinion, is that it is *not* a constant growth rate, and there is a wealth of evidence to support this conclusion.

Outcomes and sensitivities

Outcomes

Having defined all the relevant assumptions, we can now investigate how they drive prospective outcomes. For simplicity, our investigation will use the industry average glide path, as reported by Morningstar.

Figure 4. Industry average glide path

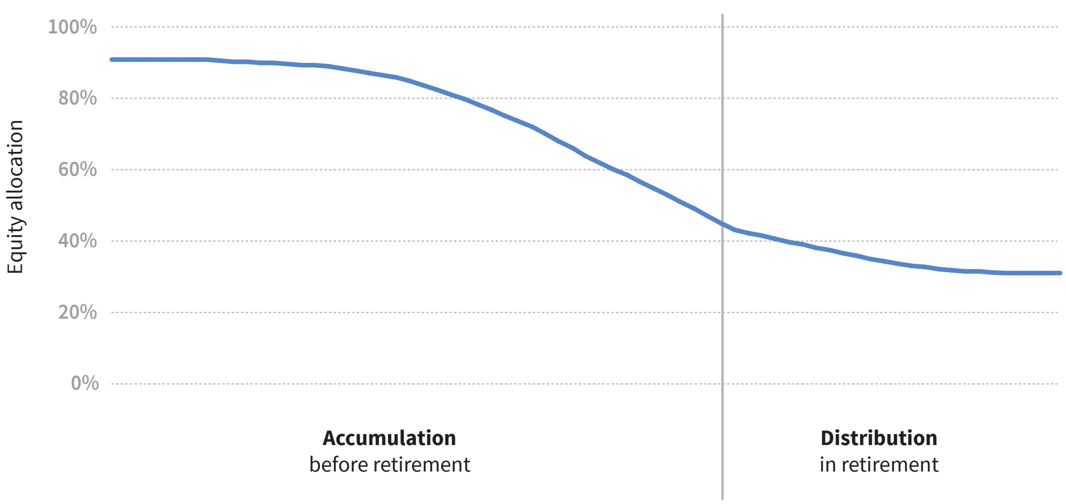

Source: Putnam.

15 A detailed explanation of our data analysis can be found in Appendix 1.

16 http://libertystreeteconomics.newyorkfed.org/2016/09/us-real-wage-growth-fast-out-of-the-starting-blocks.html.

17 See Appendix 2 for a comparison of our analysis and the N.Y. Fed analysis.

Using our baseline set of assumptions, we run a Monte Carlo simulation with 10,000 random samples and we calculate a 95.7% probability of success for the industry average glide path. This is our best estimate of the true probability of success based on the assumptions we use, but it is still just a point estimate. In the table below, we provide summary statistics about the distribution of outcomes observed in our simulations.

Figure 5. Estimates of portfolio balances at retirement and at life expectancy

Retirement portfolio balance (multiple of ending salary, real terms)

	Mean	Median
At retirement	12.0	11.3
At life expectancy	10.2	9.4

Source: Putnam.

While the outcomes based on the industry average glide path are informative, we find the greatest benefit comes from understanding how variability in assumptions drives changes in outcomes. Below, we provide detail about each individual assumption used, and later we relate the sensitivities back to costs and benefits of alternative glide path specifications.

Sensitivities

Capital market assumptions have a predictable impact on the results: Higher returns result in higher success probabilities, and lower returns result in lower success probabilities. We believe there is a high bar to overrule almost a century of empirical data, but in the event that real returns are structurally lower for the long term, with all else equal, the glide path would have to be adjusted higher to compensate.

Redefining milestone ages, such as retirement age or life expectancy, also has a predictable impact on success probabilities. Later retirement ages allow for greater savings and a longer period over which to realize capital appreciation, both of which increase success probabilities. Longer life expectancy, on the other hand, results in decreasing success probabilities as the burden on the retirement portfolio increases.

The savings rate is the single most important determinant of success probability for a plan participant. Like the assumptions discussed above, it also has a predictable effect on success probability. What is different for the savings rate assumption, however, is the magnitude of the impact. For example, we find that a 1% decrease in the savings rate requires more than a 10% increase in the equity allocation *throughout the entire glide path* to counteract the effect on success probability. This is a sobering fact — it means that for plan participants that save considerably less than 10%, there is no amount of equity that can consistently provide adequate savings to meet the full income replacement rate. Low savings rates are extremely costly for retirement planning, and the risks borne from inadequate savings can be only marginally ameliorated by a glide path.

Income replacement also plays a role in retirement success. The less income that needs to be provided by a retirement portfolio, the more likely it will be able to provide sufficient income for a considerable length of time. For a given set of assumptions, one way of looking at the impact on outcomes is to observe the probability of success purely as a function of the ratio of the income replacement rate to the savings rate. If the ratio of the assumed income replacement is too high relative to the savings rate, the glide path cannot make up for insufficient savings.

To isolate and visualize this impact, below we show the relationship for the industry average glide path, varying only income replacement relative to savings rate. Our baseline assumption can be seen with an x-value of 4, which is the 40% income replacement (net of SSA benefits) divided by the 10% savings rate.

Figure 6. Probability of success as a function of the ratio of income replacement rate to savings rate

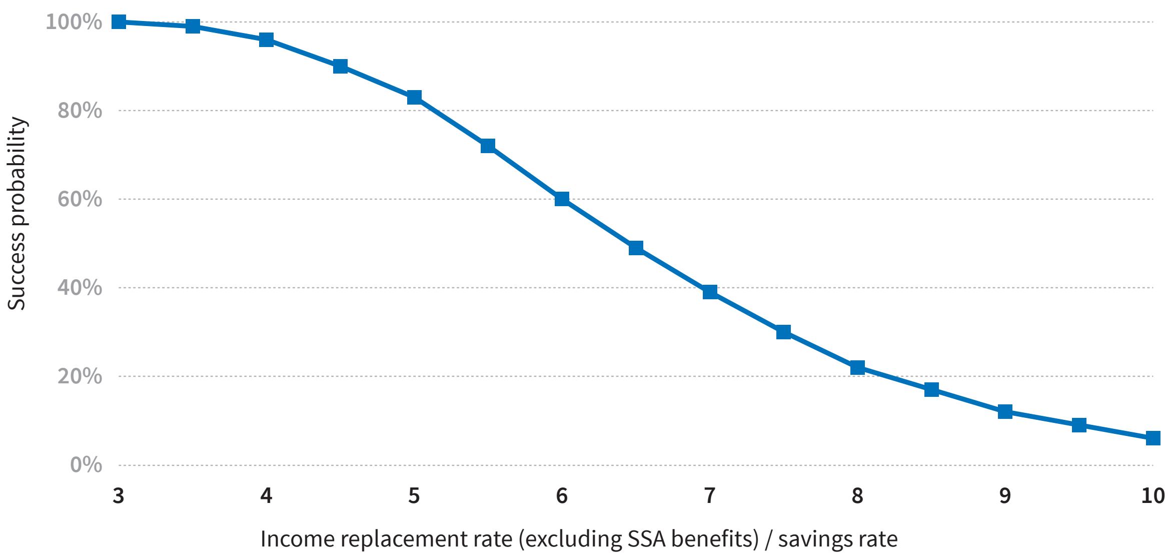

Source: Putnam.

Finally, the real income growth rate plays a role in sensitivities since it shapes the contributions to the retirement portfolio. As income grows faster and earlier, the success probability for a plan participant increases, all else remaining equal. If, instead of an exponentially decreasing real income growth rate, we chose to use a constant growth rate of 1.4%, the median compound annual growth rate (CAGR) calculated in the appendix, the success probability falls to 84.3% — a significant drop. Given that this alternative assumption is not an accurate depiction of reality, we are not concerned about the drop, but it is telling that what appears to be a minor reshaping of the lifetime income stream can have a major impact on the conclusions that can be drawn from a simulation.

Managing retirement risks

Understanding how assumptions affect outcomes is only one part of glide path design — it is merely used to establish a baseline. More important is the next step, which uses the chosen set of assumptions to determine how and when to adjust the shape of the glide path to manage the risks faced by plan participants. The major risks faced by plan participants are sequence risk, longevity risk, and inflation risk.

Sequence risk

Sequence risk is the risk posed to the retirement portfolio by the path of capital market returns. While a single lump-sum deposit into an investment account is not affected by the timing of returns (assuming no rebalancing), a retirement portfolio is affected by the timing of returns because it is subject to recurring cash flows and systematic reallocations.

Sequence risk increases as the size of the portfolio increases and the time to retirement decreases. Sequence risk is dangerous because it essentially locks in realized losses. Avoiding a significant drawdown to the retirement portfolio near the retirement date is paramount. When these losses occur in close proximity to retirement, they have the potential to impair the retirement portfolio to such an extent that it may no longer be able to meet the needs of the plan participant throughout retirement.

In addition, sequence risk also increases as the volatility of the maturity portfolio increases. The higher the volatility of the portfolio, the greater the probability of realizing a significant drawdown, and therefore, the greater the probability of imposing significant permanent losses on the portfolio.

In order to meet the spending needs of the plan participant, however, the portfolio must still be fully invested and cannot simply sit in cash to counterbalance sequence risk, as this will likely fail to provide an adequate real return to sustain the portfolio throughout retirement. Instead, the portfolio must be invested in the most efficient manner possible, where any volatility is well compensated by capital markets.

The goal of the maturity portfolio, therefore, should be to earn high risk-adjusted returns, as measured by the Sharpe ratio. This necessitates a portfolio that has equally balanced risk between stocks and bonds, which is sometimes referred to as the risk-parity portfolio.[18] The greater focus on efficiency is a significant difference from the early stages of the glide path, where capital appreciation is favored, even if it results in a less efficient portfolio.

Figure 7. Asset allocation differs from risk allocation because of the greater risk of equities

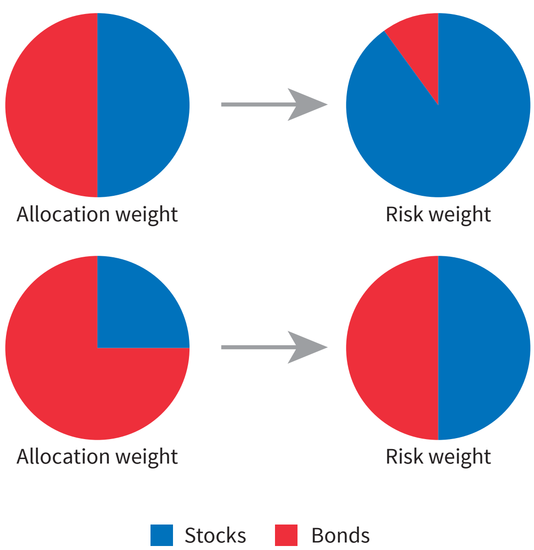

Source: Putnam.

18 Schoen, Robert. *Parity strategies and maximum diversification.* 2013. (link: https://s349713815306bb38.jimcontent.com/download/version/1446305380/module/8369748395/name/whitepaper_parity_strategies.pdf).

Sequence risk should be of the utmost concern to an investment manager, as this risk will impact *all* plan participants that will depend on their portfolio for retirement income. Furthermore, plan participants have little power to control sequence risk on their own, as they have no control over the timing of capital market returns. In fact, those who are most successful and plan best for retirement, by maintaining high savings rates and high portfolio values, will be most susceptible to sequence risk as they near retirement. Nevertheless, those participants who did not save enough while working and who might not be able to support a full 75% income replacement in retirement are also the least able to afford to lose part of their portfolio due to unfortunate timing of drawdowns.

Finally, since all participants will face some sequence risk in a traditional glide path that decreases equity allocations through time, the manager must pay special attention to when this risk is greatest for plan participants. Participants are assumed to require withdrawals each year in retirement to provide income, so it is preferable that they are not also subject to glide path reallocations as well. The combination of withdrawals and reallocations could potentially result in significant selling of assets just after a market drawdown, thereby hindering the ability of the portfolio to recover.

When the timing of this impairment comes *after* retirement, the ability of individuals to adjust their situation is also limited. By concluding glide path reallocations prior to retirement, the retirement portfolio has more time to recover from drawdowns, and at the same time participants likely have more options available to adjust to their current situation (for example, by working for a few more years). This is an example of where it is important to think of a glide path in a holistic, applied view, even after considering the outcomes of various simulations.

Longevity risk

Longevity risk is probably the most significant fear for most plan participants — the fear of living longer than their retirement portfolio can support. The thought of being many years into retirement and seeing an account with a zero balance is a much more tangible concern than sequence risk impairing a portfolio early in retirement or inflation slowly eating away at purchasing power over time. Perhaps this powerful fear causes people to overestimate their own longevity risk.

According to the EBRI study, 68% of retirees surveyed believed it was “very likely” or “somewhat likely” that they would live at least until age 85, despite the fact that based on the life expectancy calculations above, men have only a 44% probability and women have only a 56% probability of living to that age, conditional on reaching retirement age.

Even if longevity risk was better understood by plan participants, however, it would still only impact a small proportion of the population. Based on our life expectancy calculations, we can reasonably expect about 20% of plan participants to live to our life expectancy assumption of 92 years of age. Not all of these individuals will face longevity risk. Some, if they are reasonably represented by our other assumptions, will still have adequate wealth to support many more years of retirement spending.

In our simulations, we found the median portfolio value at retirement was 11.3 times ending salary. Using a conservative, risk-balanced portfolio in retirement, we also found that the median balance at life expectancy was 9.4 times ending salary (in real terms), which would provide for adequate retirement income well beyond any reasonable life expectancy for a 92-year old plan participant. So, while 20% is the highest possible proportion of individuals that could face longevity risk, the expected proportion of participants at risk is considerably lower. In fact, at full life expectancy, we observe only 6.4% of simulations resulting in non-zero balances below the assumed real ending salary (equivalent to 2.5 years of withdrawals). The joint probability in this case of living beyond 92 *and* having less than 2.5 years' worth of withdrawals in a retirement account is just over 1%.

Longevity risk affects participants unevenly

In addition, longevity itself does not affect all plan participants equally. One reason is that there is evidence of heritability of longevity.[19] Hence, while we can inform our life expectancy assumption based on SSA data, each individual participant might be able to more accurately determine his or her own longevity risk based on genetic predisposition. The resulting information asymmetry, from the perspective of the investment manager, means it is difficult to incorporate this potential information into glide path design.

Furthermore, the SSA has also found evidence that socioeconomic status plays a role in determining life expectancy.[20] If the individuals who are most likely to outlive our assumed life expectancy are also more likely to save more or require lower income replacement in retirement as a result of their higher socioeconomic status, then the impact of longevity risk on our assumed glide path will again be muted.

Finally, we must consider how best to mitigate the impact of longevity risk. Longevity risk can be addressed by increasing portfolio contributions, increasing portfolio growth, or decreasing portfolio liabilities, potentially by hedging the risk directly.

By saving more, plan participants can greatly increase the probability of a successful outcome, no matter how long they might live. By working longer, plan participants can increase contributions, increase portfolio growth, and decrease the amount of time their portfolio will need to provide income in retirement. If longevity risk is a major concern, they have the ability to purchase a Longevity Income Annuity (LIA), which directly hedges longevity risk by providing guaranteed deferred income. And with the passing of a rule by the IRS and Treasury[21] in 2014, retirees can even use their retirement savings to purchase this "longevity insurance" if they want to combat longevity risk.

The only way the investment manager can attempt to address longevity risk, however, is by increasing equity allocation throughout the glide path. Unfortunately, doing so is fraught with risk, as it necessarily increases sequence risk for all plan participants, especially later in the glide path. The impact of such a reallocation is also limited as simulations show that a small increase in savings rate is far more effective than even a large increase in equity across the entire glide path. If longevity risk is to be addressed directly, it is best addressed by the plan participant, who has better information and better tools to do so.

19 See Herskind, et al (1996) https://www.ncbi.nlm.nih.gov/pubmed/8786073, and Kerber, et al (2001) https://www.ncbi.nlm.nih.gov/pubmed/11253150.

20 Waldron (2007) https://www.ssa.gov/policy/docs/ssb/v67n3/v67n3p1.html.

21 https://www.federalregister.gov/documents/2014/07/02/2014-15524/longevity-annuity-contracts.

Inflation risk

Inflation risk is the risk that the purchasing power of a retiree's savings is eroded by rising prices. SSA benefits are, for the most part, insulated from inflation risk because SSA benefits receive a cost-of-living adjustment (COLA) according to increases in the Consumer Price Index for Urban Wage Earners and Clerical Workers (CPI-W).

Another aspect of inflation risk is that the consumption basket for a retiree might differ from that of the CPI calculation by, for example, including a higher exposure to healthcare inflation. While this exposure is difficult to hedge directly, Health Savings Accounts (HSAs) provide at least some relief for individuals in qualified high-deductible plans facing high healthcare costs. These accounts allow for greater tax advantages than retirement accounts on savings for qualified healthcare expenses. Given the preferential treatment of these accounts, most advisors suggest that workers first maximize their employer match in retirement accounts, and then direct contributions to their HSA before making additional contributions to their retirement accounts.

Aside from this unique way of addressing rising healthcare costs, we have found that the best hedge for inflation as a whole is a fully invested portfolio. As noted earlier, the capital markets assumptions use real returns for the asset classes to be included in the portfolio, and so have already accounted for the impact of inflation to some degree.

Another strategy for portfolio construction is to include assets that will do well in an inflationary environment. The typical assets cited are commodities, TIPS, and REITS. Unfortunately, all of these assets come with additional risks that might make them a poor fit for a static allocation.

Commodities, for example, as represented by the S&P GSCI Index, have delivered a Sharpe ratio of less than half that of stocks and bonds since the index inception in 1970. Even worse, the index did not go live until May 1991, and since then the index has realized a negative excess return (as of December 31, 2017). Consequently, while inflation has decreased the purchasing power of a dollar over that time, commodities exposure has done little to hedge that risk.

Treasury inflation-protected securities (TIPS), while they are directly tied to CPI in the form of principal adjustments, might still suffer in a rising-rate environment. For example, if breakeven inflation[22] is static but real yields increase, TIPS will likely suffer. TIPS were first issued by the U.S. Treasury in 1997, so there is limited history from which to draw quantitative conclusions, but over that time period, they have been dominated by interest-rate risk. Given their high correlation to Treasuries, it is not clear that TIPS warrant a strategic allocation in a glide path portfolio.

Finally, REITS are dominated by equity risk. While they also demonstrate some sensitivity to interest-rate risk, it is hard to argue that they deserve a strategic allocation in a retirement portfolio more so than any other equity sector. In fact, one could reasonably argue that energy sector stocks (since they are correlated to commodities) or healthcare sector stocks (since they are a better fit for retiree spending) would be a better static inflation hedge. Nevertheless, we hesitate to recommend static, permanent sector allocations, as it is extremely hard to justify from a quantitative perspective and we believe that sector tilts are better left to a dynamic allocation process around an unbiased benchmark.

22 Breakeven inflation is the difference between the yield of a nominal fixed-rate bond and an inflation-linked bond.

Plan participants should want a portfolio that will effectively protect them from inflation risk, but they should not be forced to suffer the lower risk-adjusted returns or hidden risks that often come along with hedging inflation. This is a strong argument in favor of active management, which allows for dynamic asset allocation that invests in these asset classes only when inflation will pose a significant threat to the portfolio, and therefore should provide better risk-adjusted returns.

Designing the glide path and focusing on outcomes

Because the risks to plan participants will vary throughout time, we will specifically address the nature of glide path design in pieces.

Defining the maturity portfolio

Once the maturity portfolio is being drawn upon, market-driven drawdowns have a significant and lasting impact on the portfolio. At maturity, when capital preservation takes precedence over capital appreciation, it makes sense to sacrifice higher expected returns for better diversification within the portfolio. In this sense, sequence risk is the primary risk to be managed when nearing retirement, and therefore the maturity portfolio must be built above all to withstand this risk.

In any case where stocks and bonds have equivalent risk-adjusted returns, which is our null hypothesis, a portfolio with equal risk contributions from each asset class will provide the highest Sharpe ratio possible. The optimal portfolio in this sense is the risk parity portfolio, or in the case of a stock/bond glide path, approximately a 25/75 portfolio (assuming stocks are three times as volatile as bonds).

We also stress test this portfolio for different levels of withdrawal rates in retirement. Using 100 random samples of returns, which are then reordered 100 times, giving us effectively 10,000 samples, we test 21 different static-allocation portfolios with different stock and bond allocations to make sure the portfolio is robust to sequence risk in retirement. What we find is that the risk-parity portfolio, highlighted in red in Figure 8, is adequate for any reasonable, generally accepted withdrawal rate, which is typically assumed to be 4%–5%.[23] Note that our simulations earlier assumed a 40% income replacement rate and had a median portfolio balance of 11.3 times ending salary at retirement, implying a 3.5% withdrawal rate in retirement.[24]

Figure 8. The risk parity portfolio is adequate for generally accepted withdrawal rates

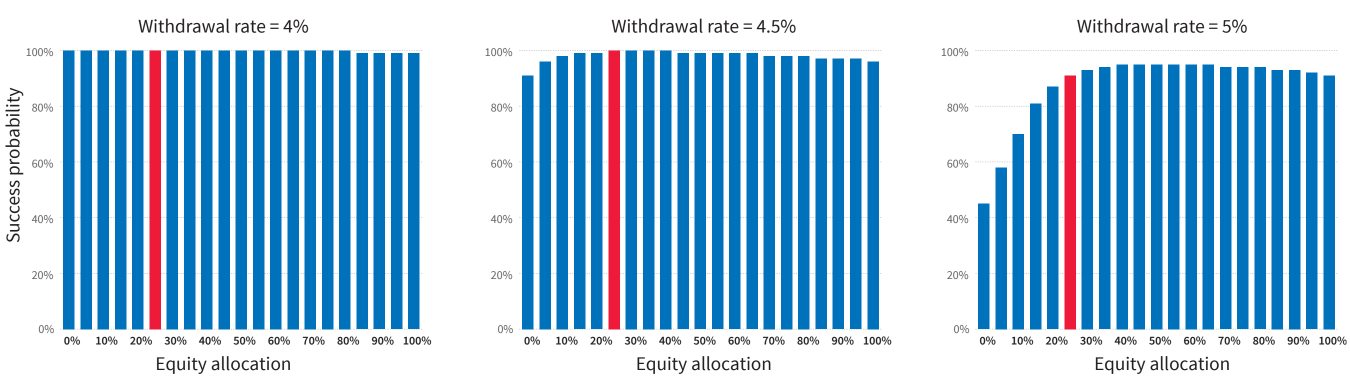

Source: Putnam.

23 www.retailinvestor.org/pdf/Bengen1.pdf.

24 See Appendix 3 for further detail relating withdrawal rates to income replacement rates.

Based on these charts, it would appear that a maturity portfolio with a higher equity allocation could also be acceptable, but we caution that this specific example only takes into consideration the retirement period, so it would therefore overlook the additional sequence risk experienced in the years prior to retirement. Given the limited marginal impact on the probability of success in the post-retirement period from a higher equity allocation, it is a poor tradeoff to increase sequence risk immediately prior to retirement by choosing a higher equity allocation for the maturity portfolio.

Defining the starting portfolio

Once the maturity portfolio has been defined, the rest of the glide path must be determined in order to manage the full set of risks an individual will face in planning for retirement. The maturity portfolio is built to control sequence risk, but the starting portfolio should be focused on longevity risk. In order to combat the risk of a shortfall in retirement savings, this portfolio should be built for growth.

Since sequence risk necessitates a relatively low equity allocation at maturity, the glide path would benefit from starting with a high equity allocation initially. This starting point allows for greater growth during the early years of the contribution phase. We call this the Aggressive Accumulation phase. For those participants with 40+ years to retirement and a relatively low portfolio balance, sequence risk is minimal and temporary volatility of the portfolio is of little concern relative to the need to grow the assets from a low base. Sequence risk only poses a threat to the success of the plan participant when the portfolio value is large and reallocations or withdrawals are occurring. Since these conditions are not present early in the glide path, sequence risk is of minimal concern for the portfolio during the Aggressive Accumulation phase.

Connecting the dots

As stated, at the start of the glide path, longevity/shortfall risk is the primary concern, but at the end of the glide path, sequence risk is the focus. Everything in between is about efficiently managing these risks and transitioning from managing one risk to the other.

During the first part of the contribution phase, which we referred to above as the Aggressive Accumulation phase, the glide path must prioritize the equity allocation in order to maximize the total wealth of the portfolio. As the individual moves closer to retirement, sequence risk grows, and the equity allocation must be reduced, while still allowing for portfolio growth to help attain the required portfolio wealth at retirement. We call this the Conservative Accumulation phase.

When retirement is far enough away, an equity drawdown will not have a lasting impact because the portfolio has time to recover. But as retirement approaches, the probability of the portfolio being permanently impaired due to an equity drawdown increases significantly as the investment horizon decreases.

Figure 9. Equity return variability is greater for shorter investment horizons

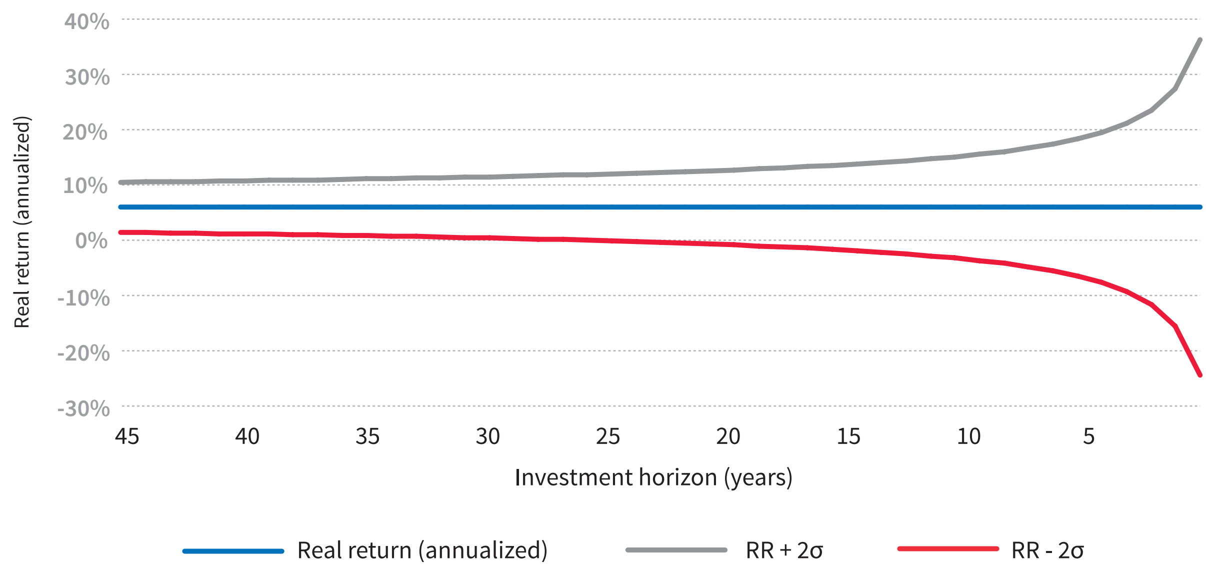

Source: Putnam.

Because the probability of such an equity drawdown increases *at an increasing rate*, an efficient portfolio should theoretically also decrease the equity allocation at an increasing rate, without going so far as to result in undue sequence risk immediately prior to retirement.

Modifying outcomes

The industry average glide path, based on our simulations above, yields a notably high success rate. However, after gaining a more thorough understanding of the risks faced throughout the glide path, it appears that the industry average glide path probably does too little to combat longevity risk at the start of the glide path, and it does too little to protect against sequence risk near retirement. This is evident in the low starting equity allocation and the high ending equity allocation, which fails to reach a risk-parity portfolio by retirement and is forced to continue de-risking beyond the expected retirement date. Instead of addressing the different risks at the right times, the industry average glide path attempts to strike a balance of managing multiple risks at all points along the glide path, creating a potential inefficiency.

As an alternative to the industry average glide path, we create another glide path that is fully invested in equity at the start, reaches a risk parity portfolio at retirement, and approximately tracks the path of the industry average in between, having the exact same allocation as the industry average at year 25 (20 years before retirement). Obviously, this alternative glide path will have a steeper slope overall. The goal is to observe how outcomes change as we adjust the investment of the portfolio to better manage the risks that affect plan participants. We do not necessarily advocate for the use of this specific glide path, rather we present it to illuminate how changes to the shape of the glide path can endure risks and modify outcomes.

Figure 10. An alternate glide path with a steeper slope than the industry average

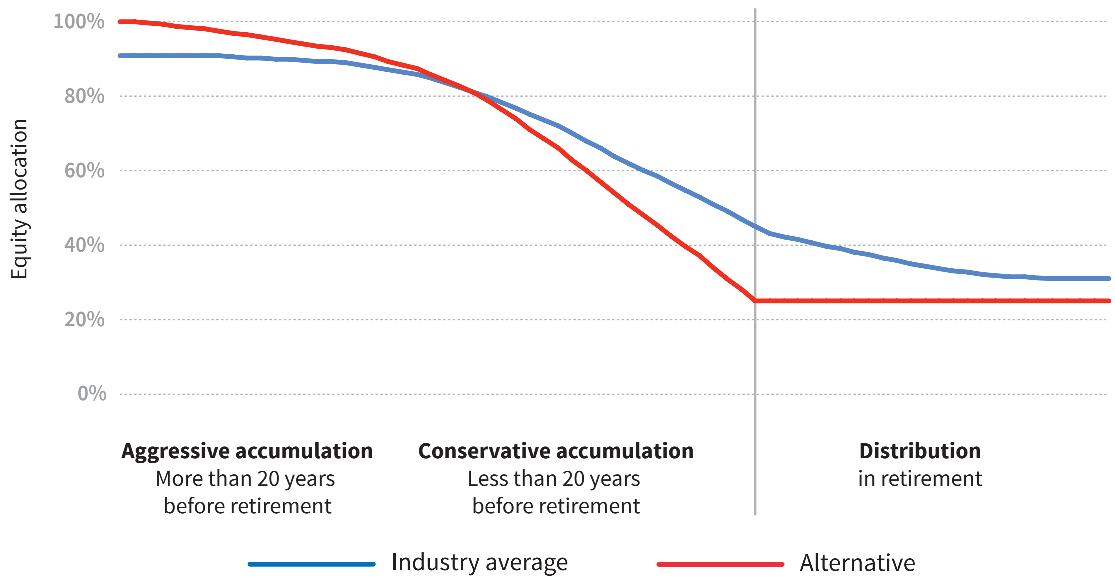

Source: Putnam.

Since this glide path is more conservative where portfolio values are greatest, it is reasonable to expect that it will have a lower simulated success probability. In fact, we observe that the overall success probability of the alternative glide path decreases to 92.9%. But this cost means other tangible benefits for the plan participant.

The industry average, relative to our calculated alternative, has a higher probability of success, along with higher mean and median values at retirement. But those gains are the result of more variability in the distribution of portfolio values at retirement. Much of the higher variability manifests as higher portfolio values on the right half of the distribution, where retirees will easily be able to meet their needs in retirement, but with that higher variability also comes a higher probability of landing a plan participant in the far left tail of the distribution.

Figure 11. Comparing simulated outcomes of the industry average and alternative glide paths

Retirement portfolio balance (multiple of ending salary, real terms)

	Mean	Median	Standard deviation
Industry average	12.0	11.3	4.0
Alternative	11.3	10.7	3.6

Source: Putnam.

A retirement portfolio that has ten times ending salary, assuming it needs to provide the 40% income replacement outside of SSA benefits as calculated above, will have a withdrawal rate of 4%. With only a 4% withdrawal rate, the retirement portfolio simply needs to outpace inflation (managing sequence risk, of course) in order to deliver 25 years of retirement income.

Alternatively, we could define a target range of outcomes that we deem to be acceptable. With a target 4% withdrawal rate, this could be a portfolio value at retirement that allows for a 3%–5% withdrawal rate. Evaluating the outcomes of the two glide paths, we find that the alternative glide path has a higher probability of delivering an outcome in this range (59.8% versus 55.8% for the industry average). It seems counterintuitive to us to risk a high probability of retirement success in order to increase this ending portfolio value, especially when it also increases the probability of a major failure.

Figure 12. Comparing outcome distributions of the industry average and alternative glide paths

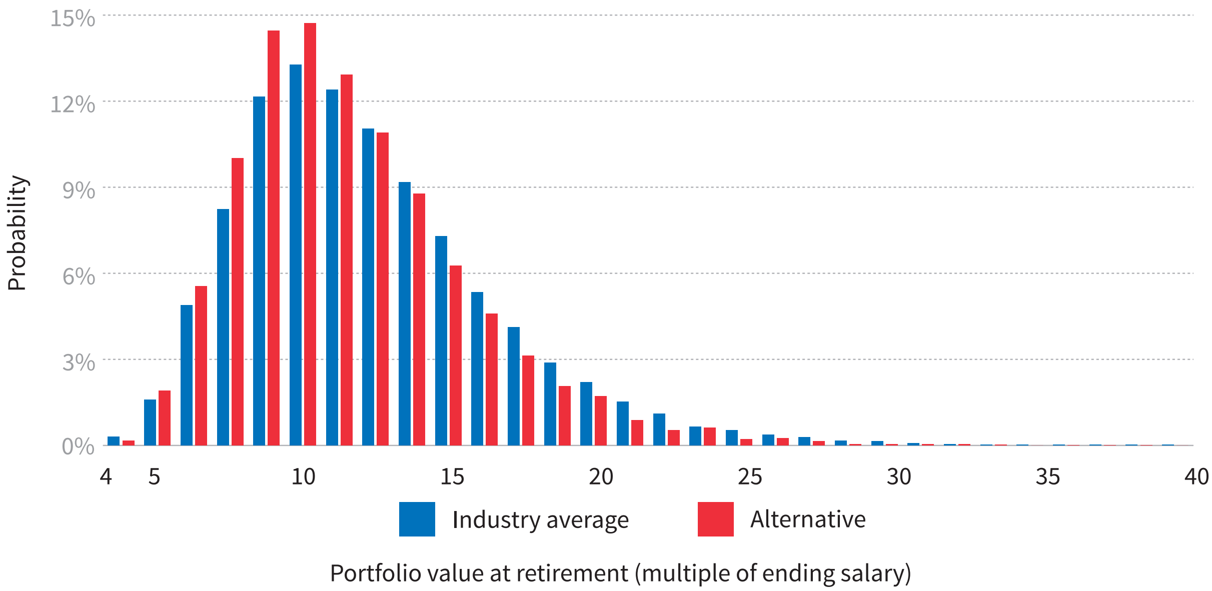

Source: Putnam.

Finally, it is worth considering the outcomes outside the scope of a binary simulation. In reality, failure is not binary, and no matter the outcome, plan participants must make the best of their situation. For this reason, we find solace in the lower variability of outcomes that result from the alternative portfolio. Much of the probability distribution is clustered around an acceptable outcome, and even if retirement savings comes up short, it is likely to be only a little short. In these scenarios, slight changes in retirement spending can easily make up for the savings shortfall. Only the large magnitude misses, which are more likely to occur in the higher-risk glide path, leave unfortunate retirees with very few options.

Other considerations

Although simulations can be extremely informative, they can also overlook key aspects of planning for retirement. Realistically, people don't plan for retirement based on an arbitrary age chosen when they are in their 20s. More likely, they plan for retirement based on a combination of current age, current total wealth, and expected retirement spending needs.

Planning for retirement

A conservative portfolio allocation in the decade before retirement better allows plan participants to plan for retirement. In this case, it is reasonable to trade the slight statistical advantage of a riskier glide path just before and into retirement for the pragmatic application of retirement planning. A lower volatility portfolio gives individuals less uncertainty when planning for retirement, which is an important part of keeping their best interests in mind.

Additional income/debt

This analysis does not account for revenue from other sources. If there are any other assets whatsoever, these assets could provide additional income in retirement that has not been accounted for. This additional income could be in the form of after-tax savings outside of retirement accounts, home equity, inheritance, etc. The result is a material increase in the probability of success as the above analysis would have effectively underestimated the true savings rate for the individual. While not included in our assumptions, any additional income would either increase success probabilities or allow for more conservative glide path allocations.

Alternatively, as additional income outside the scope of our simulations improves retirement outcomes, debt could provide a headwind to a successful retirement. For the most part, unsecured debt poses the greatest risk, since the greatest source of secured debt is mortgage debt, which tends to help individuals build home equity. We do not find evidence that unsecured household debt is likely to have a large enough impact to consider its impact on glide path specification, and we find evidence that fears related to student loan debt are likely overblown.[25]

Conclusion

There is no perfect glide path. It is impossible to say a glide path is optimal, even under a certain set of assumptions. Instead, managers must first identify the most representative set of assumptions for the participant population, and then make hard decisions around the costs and benefits of managing the numerous risks facing participants at various points along the glide path.

Glide path design might be as much art as it is science, but that is not to discount the importance of the quantitative aspect. Without a detailed understanding of how assumptions relate to outcomes, glide path specification is simply a guessing game. By providing a deeper awareness of the relationships above, we hope to elucidate the complex interactions of assumptions and risks. The ultimate goal is that all market participants — from investment managers to consultants to plan sponsors to the participants themselves — are empowered with better understanding and are able to make better informed decisions that lead to more successful outcomes.

With greater perception of how and where managers can have the most significant impact, glide path design is an integral part of a successful retirement. But it is still only one piece. Plan participants must save enough. Plan sponsors must encourage participation. Consultants must guide understanding and advocate best practices. Even politicians can play their part, with common-sense reforms and innovative solutions that make Social Security solvent and incentivize retirement savings. Ultimately, when all the pieces come together, it is the responsibility of the investment manager to shepherd each and every plan participant to the most successful retirement experience achievable.

25 See Appendix 4 for more detail concerning student loan debt.

Appendix 1

Real income growth rate

Using the SSA Benefits and Earnings Public-Use File,[26] we are able to access a random sample of 1% of the SSA's Master Beneficiary Record in order to generate more than 470,000 observations of longitudinal earnings history. After removing outliers and placing simple constraints on the observations, we retain approximately 35% of our initial sample (more than 165,000 observations).

When cleaning the data, we select only individual observations with at least 20 years of non-zero income, of which at least five years must fall into the third quintile of observed income. We consider the first year of true earnings to be the first observation in the series to fall within the third quintile, and we consider the last year of true earnings to be the last observed non-zero income that is greater than or equal to the average non-zero income for the individual over the entire time period. Thus, none of our samples start from an artificially low base, as the result of part-time work early in an individual's life, nor do they include extremely low observations in retirement as a result of post-retirement supplemental income. Instead, we have selected only individual income streams with enough observations to be meaningful, while also making sure they are not affected by outlier earnings within the series of income values.

It should also be noted that incomes in this data are effectively capped by the SSA taxable maximum (the "tax max"). This has had a different impact on taxable earnings through time.[27] Since about 1980, less than 10% of workers have been affected by the tax max, but for a period of time before that it affected a greater proportion of workers, as evidenced by the drift of the third quintile earnings as a percent of the tax max in the graph below. Nevertheless, we found no evidence that our CAGR calculations are at all biased as a result of the changing level of the tax max through time.

Figure 13. The impact of the "tax max" has drifted through time

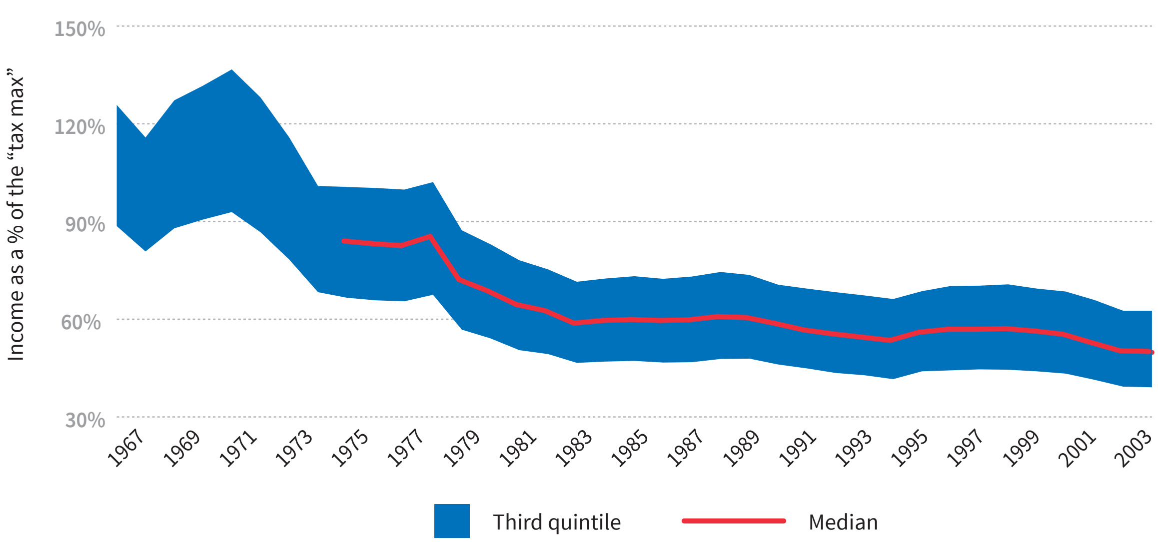

Source: Putnam, based on SSA data.

Within our cleaned sample, we find a median longitudinal real growth rate of 1.4%. See Figure 14 for selected percentile data reported.

26 https://www.ssa.gov/policy/docs/microdata/earn/.

27 https://www.ssa.gov/policy/docs/policybriefs/pb2011-02.html.

Figure 14. Distribution of nominal and real income CAGR

Percentile	1	5	10	25	50	75	90	95	99
Nominal CAGR	0.87%	2.51%	3.24%	4.46%	5.84%	7.13%	8.27%	8.96%	10.40%
Real CAGR	-3.36%	-1.74%	-1.03%	0.11%	1.39%	2.59%	3.58%	4.15%	5.14%

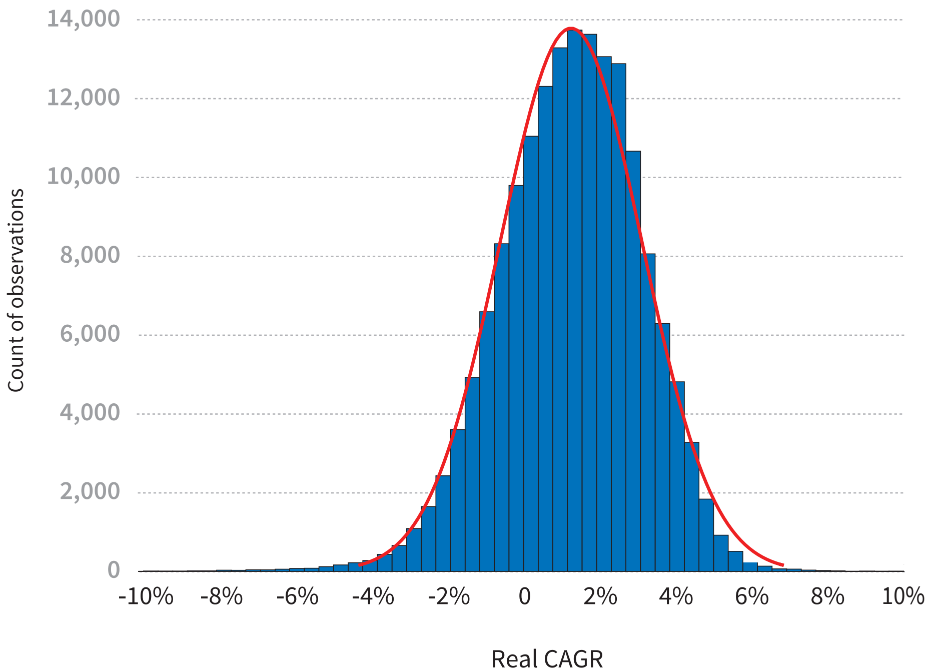

Source: Putnam, based on SSA data.

What is interesting about the distribution, however, is that real income growth is generally front-loaded. When dividing each individual series in half, we find that the median real income growth of the second half of the series is quite often negative.

It should be noted that neither half of the sample should demonstrate any obvious bias as a result of unusual earnings observations, as we have already removed outlier earnings when we cleaned the data.

Figure 15. Dividing the series reveals that real income growth is generally front-loaded

Percentile	1	5	10	25	50	75	90	95	99
Nominal CAGR	0.87%	2.51%	3.24%	4.46%	5.84%	7.13%	8.27%	8.96%	10.40%
Real CAGR	-3.36%	-1.74%	-1.03%	0.11%	1.39%	2.59%	3.58%	4.15%	5.14%
Real CAGR (1st half)	-1.70%	0.94%	1.81%	2.84%	4.42%	6.25%	7.97%	9.37%	14.37%
Real CAGR (2nd half)	-8.73%	-5.32%	-4.30%	-2.90%	-1.10%	0.93%	2.37%	3.15%	4.96%

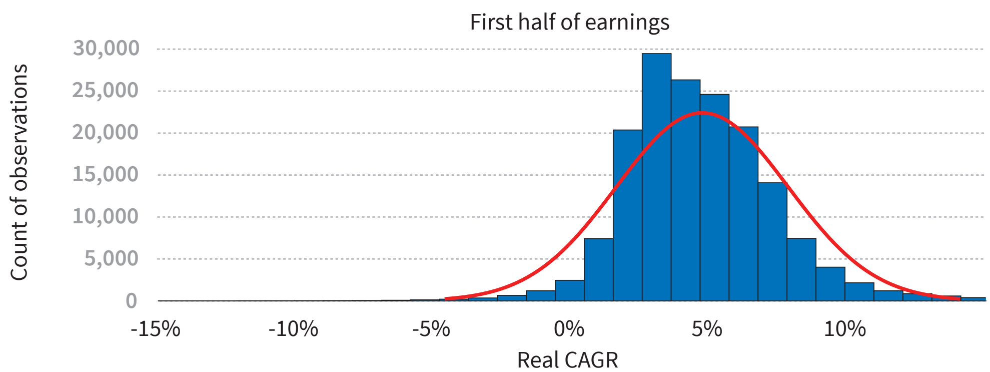

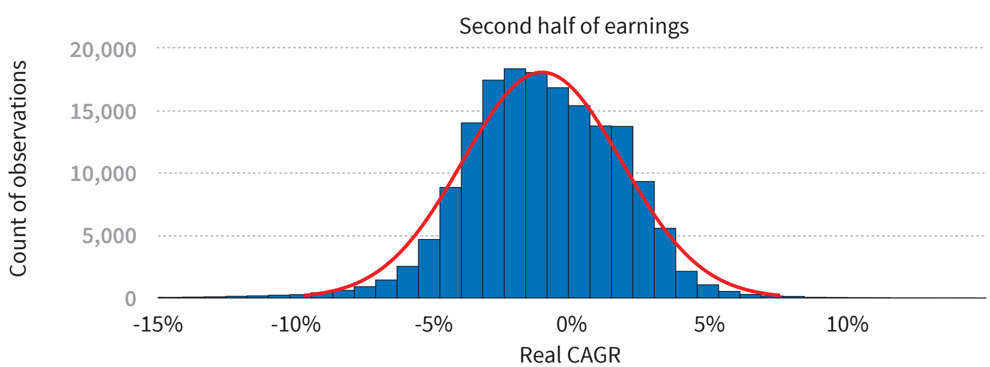

Source: Putnam, based on SSA data.

Appendix 2

Real income growth rate comparison with N.Y. Fed analysis

In Figure 16, we include a comparison of the N.Y. Fed findings and our real wage growth assumption, using their cohort for white males born in the 1950s and having a high school diploma. Since the raw data was not included in the blog post, we note that their data is approximate based on the graphics included in their post.

Figure 16. Comparing our assumptions with N.Y. Fed analysis of real wage growth

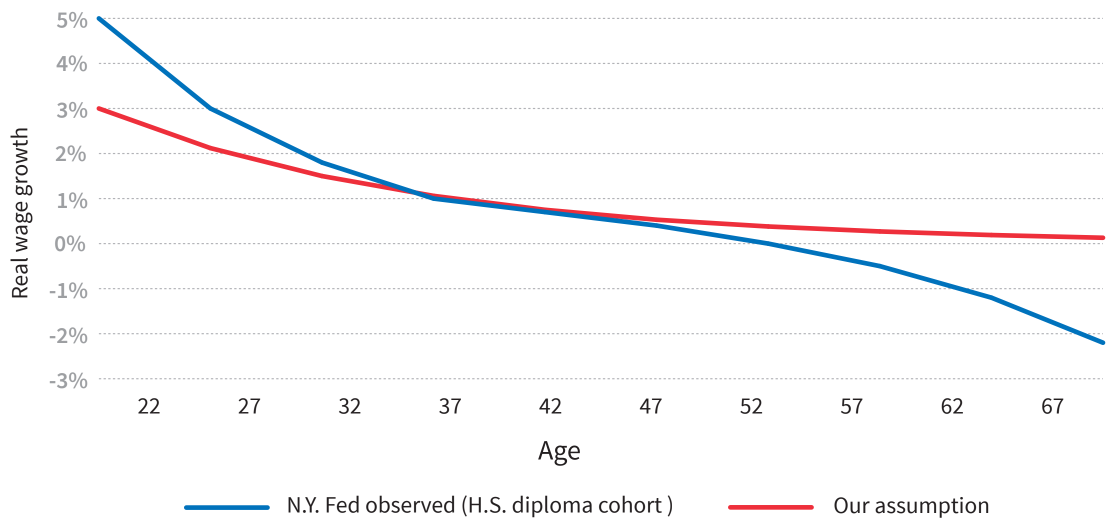

Sources: Putnam and New York Federal Reserve Bank.

Obviously, our assumption differs from the Fed analysis in two ways — we use a lower assumed growth rate toward the beginning and a higher growth rate at the end. There are two reasons for this.

First, the N.Y. Fed arrive at their income growth rate estimates by calculating an estimated level of real wages for each cohort and deriving an imputed real wage growth rate from this series. This means the Fed assumption is essentially a *weighted average* real growth rate, while ours is a *median* real growth rate. For glide path research, we believe the median is a more appropriate measure, and therefore we are comfortable with our divergence from the Fed assumption at the tails.

Second, we recognize that real income growth rates have a significant impact on the outcomes in our simulations, and small changes in this assumption can have a large impact on success probabilities. Whereas some assumptions can have a dampening effect on outcomes, the real income growth rate assumption has a compounding effect on outcomes. For example, recall that our assumption for SSA benefits is based roughly on median household income. Since the proportion of ending income that is replaced by Social Security will be lower for higher-income individuals, we understand that this is a rough estimate. But we also know that higher-income individuals tend to save more, on average, and so the ultimate impact of our Social Security income assumption on the final success probability is dampened in the end.

In the case of real income growth, higher real income at the beginning of one's career and lower real income at the end will lead to considerably higher success probabilities. This is because savings (derived from a savings rate applied to real income) is front-loaded, allowing for greater compounding of portfolio returns. Then, by assuming real income growth decreases at the end of a career, the income replacement rate required for a successful outcome is less taxing on this larger portfolio. For this same reason, a constant real growth rate acts in the opposite way: Savings contributions are delayed considerably and income growth is back-loaded, which results in unrealistically low success probabilities.

Figure 17. Comparison of different rates of real growth in wages over time

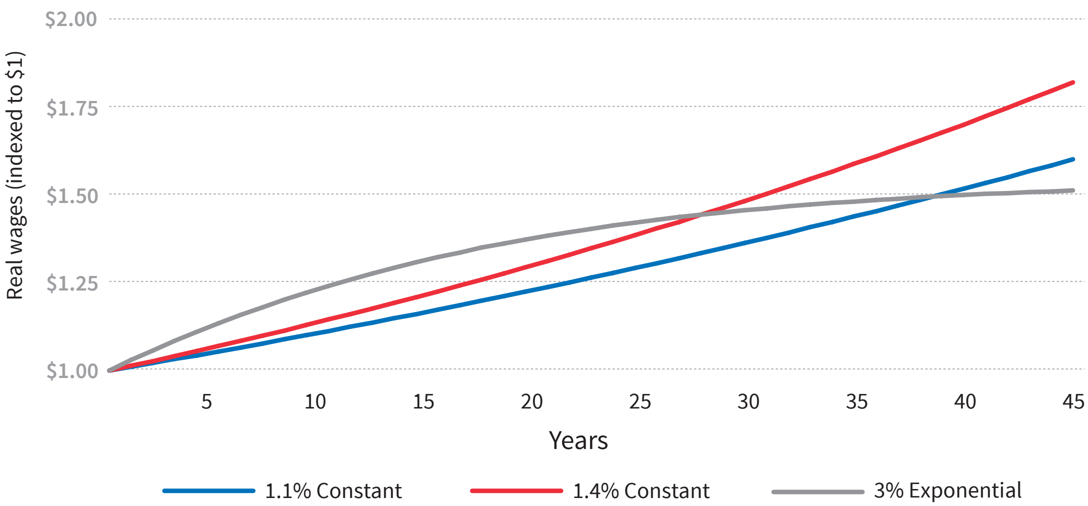

Source: Putnam.

We can also view this effect in terms of the total dollar contributions to the retirement portfolio. Figure 18 does not take into account the increased net present value of portfolio contributions; it merely calculates the total dollar contributions that result from a constant savings rate applied to real wages. Despite the fact that ending salary will be lowest with a 3% exponential real growth rate, as compared with the constant 1.1% and 1.4% real growth rates, the total dollar contributions end up similar to the 1.4% constant real growth rate scenario.

Figure 18. Comparison of total dollar contributions at different growth rates

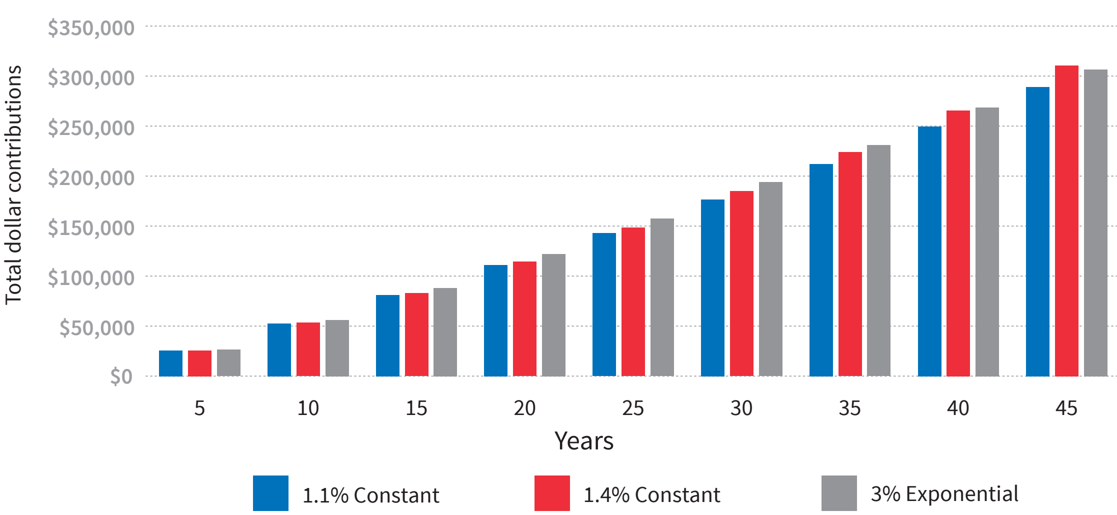

Source: Putnam.

Appendix 3

Income replacement rate versus sustainable withdrawal rate

Another important consideration is the relationship between income replacement relative to ending salary and the effective withdrawal rate of the post-retirement portfolio. Generally, 4%–5% is considered to be a sustainable withdrawal rate for a successful retirement. We do not explicitly target a certain withdrawal rate, but it is still important to consider the importance of the withdrawal rate as it relates to the portfolio value at retirement.

If we assume an income replacement rate (net of SSA benefits), we can calculate a target portfolio value at retirement assuming a given target withdrawal rate. We can use the following identity to calculate the table below:

$$ES * IR = RW = RP * WR \quad \Rightarrow \quad \frac{RP}{ES} = \frac{IR}{WR}$$

$$ES = Ending\ salary$$
$$IR = Income\ replacement\ rate$$
$$RW = Retirement\ withdrawal\ amount$$
$$RP = Retirement\ portfolio\ value$$
$$WR = Withdrawal\ rate$$

Ending salary multiplier

		Withdrawal rate			
		4.0%	**4.5%**	**5.0%**	**5.5%**
Income replacement rate	***35%***	8.75	7.78	7.00	6.36
	40%	10.00	8.89	8.00	7.27
	45%	11.25	10.00	9.00	8.18
	50%	12.50	11.11	10.00	9.09
	55%	13.75	12.22	11.00	10.00

Looking purely at the success or failure of the glide path at retirement, however, is only marginally informative. Since "failures" at retirement that fail by only a slight margin still have a high probability of providing adequate income in retirement, the true success probability is significantly understated using the probability calculated at retirement, as seen in the table below (calculations reflect the industry average glide path).

Success target	Portfolio target (salary multiple)	(A) Pre-retirement success probability	(B) Post-retirement success probability	(A*B) Estimated joint probability	True success probability
5% withdrawal	8x	86.2%	89.7%	77.3%	95.7%
4.75% withdrawal	8.42x	82.5%	98.0%	80.8%	95.7%
4.5% withdrawal	8.89x	77.5%	99.7%	77.3%	95.7%
4% withdrawal	10x	64.1%	100.0%	64.1%	95.7%

In addition, we can also plot the success and failure measures at retirement and at life expectancy to better depict the distinction between the two. For example, taking a target 4.5% withdrawal rate in retirement, we plot the two distributions in Figure 19.

Figure 19. Comparison of distribution at retirement and at life expectancy

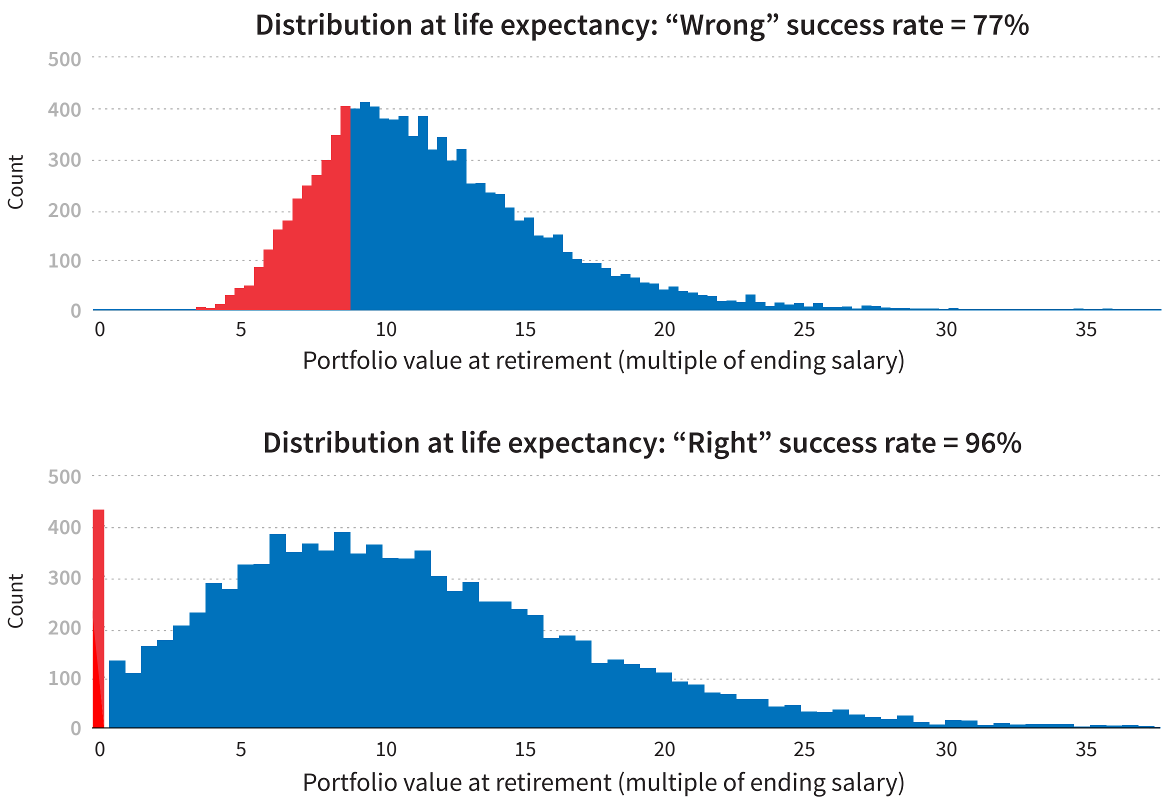

Source: Putnam.

In reality, with a 4.5% target withdrawal rate, many of the simulations will result in a far lower withdrawal rate when providing the 40% income replacement in retirement, as a result of their much higher portfolio values at retirement. In addition, those below the target portfolio value (8.89x for a 4.5% withdrawal rate) will have a high probability of successfully providing retirement income, despite being forced to withdraw more than 4.5% of their portfolio in their first year of retirement. Thus, while the relationship between income replacement rate and withdrawal rate is an important consideration, we focus on income replacement and the probability of success when reaching life expectancy as a more robust metric to inform decisions about glide path design.

Appendix 4

Student loan debt

While total household debt decreased from 2008 to 2016, student loan debt doubled over the period.[28] Some key observations include:

- Individuals under the age of 25 make up approximately half of all originations.[29]
- Balances held by individuals under age 30 account for less than one third of outstanding balances.
- Of the 44 million student loan borrowers in 2016, 5% had loan balances in excess of $100,000. This cohort represents 30% of the total outstanding balance of student loan debt.
- The cohort with the highest outstanding balances also has the lowest default rates of any cohort of student loan borrowers.

The N.Y. Fed analysis also observes that payment progress has slowed for more recent graduates, primarily due to higher borrowing amounts, while noting the effect could also be from the uptick in participation in Income-Based Repayment plans. If payment progress is slower for graduates, either because they are contributing their marginal dollar after debt service to retirement savings or because they are likely to receive some amount of debt forgiveness in the future, then this should be an encouraging sign for retirement success. After all, borrowers should not prepay a student loan and earn only the interest rate on the loan if they can take advantage of an employer match, which is almost certain to guarantee a higher return, even ignoring the potential income tax savings.

The N.Y. Fed briefing also finds that college education is associated with markedly higher homeownership rates, regardless of debt status. While student debt results in a marginal decrease in the homeownership rate, it generally has a smaller impact than the question of whether or not an individual attended college at all. According to the N.Y. Fed study, the homeownership rate for college non-grads with debt rises above the homeownership rate for those who did not attend college by age 27, and the spread continues to increase, rising to a 7% spread (35% versus 28%) by age 33. For those students who graduate, the divergence is even greater, with the homeownership rate for college grads with debt in excess of 52% by age 33.

It stands to reason that if these individuals can manage higher rates of homeownership, they should also be able to save for retirement while continuing to service their student loan debt. Thus, we believe the growth in student loan debt is unlikely to impair the ability of individuals to adequately save for retirement.

28 https://www.newyorkfed.org/press/pressbriefings/household-borrowing-student-loans-homeownership.

29 https://www.newyorkfed.org/medialibrary/media/newsevents/mediaadvisory/2015/Student-Loan-Press-Briefing-Presentation.pdf.

Bloomberg Barclays U.S. Treasury Bills 1–3 Month Total Return Index — The Bloomberg Barclays U.S. Treasury Bills 1–3 Month Total Return Index is an unmanaged index that measures the performance of one- to three-month maturity U.S. Treasury bills.

Bloomberg Barclays U.S. Treasury 7–10 Year Total Return Index — The Bloomberg Barclays U.S. Treasury 7–10 Year Total Return Index measures the performance of all public U.S. government obligations with maturities of seven to ten years.

Consumer Price Index (CPI) — The Consumer Price Index (CPI) is a measure of the average change over time in the prices paid by urban consumers for a market basket of consumer goods and services.

IA SBBI IT Government Total Return Index — The Ibbotson Stocks, Bonds, Bills, and Inflation Intermediate-Term Government Total Return Index is an unmanaged index that measures the performance of five-year maturity U.S. Treasury bonds, and includes reinvestment of income.

IA SBBI LT Government Total Return Index — The Ibbotson Stocks, Bonds, Bills, and Inflation Long-term Government Total Return Index is an unmanaged index that measures the performance of twenty-year maturity U.S. Treasury bonds, and includes reinvestment of income.

IA SBBI S&P 500 Index — The Ibbotson Stocks, Bonds, Bills, and Inflation S&P 500 Total Return Index is an unmanaged index of common stock performance.

IA SBBI U.S. Inflation Index — The Ibbotson Stocks, Bonds, Bills, and Inflation U.S. Inflation Index measures the average monthly change in the price for goods and services paid by consumers.

IA SBBI U.S. 30-Day T-Bill Total Return Index — The Ibbotson U.S. 30-Day Treasury Bill Total Return Index is an unmanaged index that measures the performance of one-month maturity U.S. Treasury bills.

S&P 500 Total Return Index — The S&P 500 Total Return Index is an unmanaged index of common stock performance.

Visit putnam.com for continuing market updates, expert insights, and investment commentaries.

Find us

Brett Goldstein, CFA, is Co-Chief Investment Officer, Global Asset Allocation (GAA), and a member of Putnam's Operating Committee. In his role as Co-CIO of GAA, he directs the overall strategy and positioning of Putnam's GAA products. He is responsible for the research and implementation of risk and portfolio construction methods across GAA products. In addition, Mr. Goldstein works extensively with retirement glide path research and GAA target-date funds. His work also contributes to Putnam's Capital Markets Outlook.

Mr. Goldstein has been in the investment industry since he joined Putnam in 2010. Previously at Putnam, he was a Portfolio Manager, an Analyst, and an Investment Associate in the GAA group.

Mr. Goldstein earned an M.P.S. in Statistics, a B.S. in Biometry and Statistics, and a B.S. in Finance from Cornell University.

A world of investing.®

Putnam INVESTMENTS

Putnam Retail Management

Putnam Investments | 100 Federal Street | Boston, MA 02110 | putnam.com

II103_DC 333656 4/23