

Investing for retirement: the defined contribution challenge

Ben Inker and Martin Tarlie | GMO | 08 April 2014

The retirement landscape has changed. Defined benefit plans, the historical workhorse of the retirement system, had the advantage of access to corporate profitability. In the event that financial asset returns fell short of design expectations, this access mitigated the impact on workers' retirement. But, as defined benefit plans have given way to defined contribution (DC) plans, the burden being placed on financial returns in satisfying retirement needs has increased.

Target date funds are rapidly becoming the workhorse for DC plans. These funds have grown substantially in recent years, partly as a result of automatic enrollment made possible by the Pension Protection Act of 2006. By and large, current target date funds resemble the old investment adviser adage that stock weight should be about 110 minus a person's age. While this satisfies the common-sense intuition that, all things being equal, weight in stocks should go down as a person ages, there are a number of problems with this approach. In this paper, we focus on two in particular.

First, the standard solution is inflexible – all things are rarely equal. To address this shortcoming, we introduce a framework based on a common-sense definition of risk, being not having enough wealth in retirement. The goal is not to put investors into yachts, but rather to increase the odds that they have the appropriate level of resources in retirement. Viewing risk this way leads to highly customisable solutions that under certain equilibrium assumptions are consistent with current solutions but offer far more flexibility and insight. Second, the standard solutions do not recognise that expected returns vary over time. We show that dynamic asset allocation – moving your assets – is an essential part of achieving retirement goals.

This paper is divided into two parts. Part I frames the question and explain how our framework leads to flexible, customisable solutions. Part II demonstrates the importance of dynamic allocation.

PART I

1.1 Asking the right question

The most common method for building multi-asset portfolios is based on Modern Portfolio Theory – maximise return for a given level of risk, where risk is return volatility. From the

perspective of the retirement problem, and perhaps more generally, this approach is inadequate. The main problem is that it is asking the wrong question: given a level of risk, i.e., return volatility, which is the portfolio that maximises the expected return?

This is the wrong question because it focuses on returns, not wealth. But returns are only the means to an end, the end being the wealth that is to be consumed throughout retirement. Not only is it the wrong question, but it presupposes the investor has a good reason for choosing a particular level of return volatility. So two investors faced with similar circumstances in terms of current wealth, future income and savings, and future consumption needs may have very different portfolios simply because their attitude toward return volatility differs.

A better approach is to focus on what really matters – wealth. An investor saving for retirement has fairly well-defined needs, both in terms of how much wealth he needs to accumulate and his pattern of consumption in retirement. An investor's portfolio should be driven primarily by the investor's needs and circumstances – what does he need and when does he need it? It should not be a function of his personality. The financial risk to an investor saving for retirement is very simple – it is not having enough wealth. So, the more appropriate question is "which is the portfolio that minimises the expected shortfall of wealth relative to what's needed?"

This definition of risk is central to our framework here. All other things being equal, a person who is more risk averse should either save more or consume less. In contrast, the standard approach gives bad advice. Putting the more risk-averse individual in a less volatile portfolio, one that from a Modern Portfolio Theory (MPT) perspective is considered less risky, without making any compensating savings or consumption adjustments actually increases the wealth risk to that individual in that he is less likely to achieve his wealth needs. A virtue of optimising based on minimising shortfall of wealth is that it is highly customisable and easily able to handle the question of how to invest for a more risk-averse person who expresses his increased risk aversion through, for example, a higher savings rate. This flexibility is a consequence of asking the *right* question.


1.2 Returns versus wealth

To better understand the difference between MPT – a return-focused approach – and the wealth-focused approach that we advocate here, it is helpful to compare the distribution of returns with the distribution of wealth. To a fairly good approximation, returns are normally distributed (Figure 1). While there is plenty of empirical evidence that, at least over shorter horizons, this is not quite true for many asset classes, our problem with the assumption for portfolio construction purposes here is not particularly that returns are fat-tailed or may be slightly skewed in one direction or another. It is rather that, even if returns are normally distributed, the wealth those returns lead to is not.

Figure 1 shows a normal distribution of annual returns for an asset with a 5% return per


annum and a 14% annualised volatility. In a normal distribution, the average is the same as both the median and the mode, the most likely return. Whether you are actually concerned with the average of all of the potential returns, the most likely return (mode), or the return that is in the middle of the distribution (median) is irrelevant – they are all the same.

Figure 1: Normally distributed returns (asset markets)



As returns compound into wealth, however, Figure 1 is no longer relevant. Figure 2 shows the distribution of ending wealth after investing \$1 for 40 years in an asset with the normal return distribution shown above. This distribution is not normal, rather it is log-normal. The shape of the log-normal distribution is profoundly different from that of the normal distribution. The expected value, or mean, of this distribution is the right hand purple vertical line. If you invest for 40 years in an asset with normally distributed returns averaging 5% per annum and an annualised standard deviation of 14%, the average wealth outcome is about \$11. The median outcome, however, is about \$7, and the mode, most likely outcome, is only \$3.4.

Figure 2: Log-normally distributed returns (wealth)



Source: GMO

Expected (mean) values are dominated by the right tail of the distribution – those lucky 40-year periods in which returns happened to average well over 5% real. While those events are rare, they have a big impact on mean wealth. But, for the purposes of saving for retirement, those outcomes are largely irrelevant. If you happen to be lucky enough to have lived and saved during the right period when asset returns were high, it doesn't much matter what your target date allocations were. You'll wind up with more than enough money to retire on.

The more important part of the distribution is the left-hand side – those events when asset markets were not kind and returns were hard to come by. Those are the events where lifetime ruin (i.e., running out of money in retirement) is a real possibility. We believe that the right way to build portfolios for retirement is to focus on how much wealth is needed and when it is needed, with a focus not on maximising expected wealth, but rather on minimising the expected shortfall of wealth from what is needed in retirement.


1.3 The retirement problem – basics

There are two obvious phases of the retirement problem – the accumulation phase, when workers are generating income and investing savings, and the consumption phase, when assets are spent.

Figure 3 shows a simple diagram of the accumulation phase, generated using fairly standard industry assumptions. An employee starts out earning \$43,000 at age 25, with income growing over time at 1.1% above inflation. The contribution rate (i.e., savings relative to

income), starts at 5%, rising to 10% at retirement, and the employer match is 3% of income. This implies an average contribution rate of 10.5% [not dissimilar to the Australian 9% – Ed]. Target wealth is 10 times final annual salary, in this case approximately \$667,000. But, given that cumulative savings total only about \$200,000, it turns out that it will take an average return of about 5.1% real per year to achieve the retirement wealth target.


Figure 3: The accumulation phase



Source: GMO

Chart 4 illustrates the consumption phase. It assumes that the participant spends 50% of final salary every year in retirement, adjusted for inflation – spending of \$33,383. (The assumption of 50% of final salary is a standard one. Implicit is the assumption that government payments will constitute another 30% so that total assumed replacement ratio is 80%.) This amounts to spending a constant 5% of target wealth at retirement. The red line shows the importance of continuing to earn returns in retirement, as a 5% spending rate in the absence of returns consumes the accumulated savings in 20 years. But, it assumes that the retiree lives 30 years beyond retirement. In order to afford this, the retiree needs to earn about 2.8% real per year during the consumption phase.

Figure 4: The consumption phase




Source: GMO

1.4 Expected shortfall

Figure 5 combines the accumulation and consumption phases into one graph. Because we define risk as not having enough money in retirement, our objective is to minimise expected shortfall of wealth after age 65. This concept is illustrated in Figure 5 by the red area – the optimal portfolios minimise expected wealth in this red zone. Minimising wealth in the red zone is equivalent to focusing on the left side of the wealth distribution as discussed in the section "Returns vs. Wealth" above.

Figure 5: Minimising expected shortfall



Source: GMO

In Figure 5 above, the wealth target post-retirement is a solid line rather than the dashed line used in the accumulation phase. This highlights the fact that wealth prior to retirement has an indirect influence on the objective function. Structurally, the objective is to minimise shortfall relative to the wealth target after age 65. Not including wealth prior to retirement in the objective function means that the investor is more tolerant of wealth volatility prior to retirement, leading to portfolios, in equilibrium, that have more weight in stocks for younger investors.


Why is it important to envision the problem in this way? Simply put, it addresses the primary financial risk of not having enough wealth in retirement. Furthermore, if you concentrate on solving this problem, "risk aversion" naturally falls out, rather than having to be guessed at or enquired about as required by MPT. A 25-year-old should invest aggressively because of her circumstances, not because of her personality – there are 40 years until drawdowns really matter for consumption goals.

A 75-year-old should invest more conservatively because of needs and circumstances. Near-term losses cannot necessarily be recovered from the nest egg as it is being consumed. And, to go beyond a static glide path and account for time-varying expected returns (see the section below on Dynamic Allocation), it is essential to have a framework that naturally balances the changing trade-offs between risk and return as investors age.

If we use fairly standard assumptions of 6% real returns for stocks and 2% real returns for bonds, with annualised volatilities of 18% and 5%, respectively, and a correlation of zero – and minimise expected wealth shortfall (ESF) assuming that investors are on the wealth targets illustrated in Figure 5, we can map out the optimal weight in stocks for each age. The blue line in Figure 6 below shows these optimal stock weights. We call this a static ESF glide path because we generate these weights by minimising the expected wealth shortfall assuming that the expected returns for stocks and bonds are constant.

For comparison, we show two additional lines in Figure 6. The yellow line is based on the old investment adviser adage that the stock percentage should be about 110 minus a person's age. The green line is a glide path used by a provider of target date funds. Relative to the ESF portfolio, it looks as if 110–Age is too conservative for almost all ages leading to retirement, and the XYZ path is a little bit too conservative from ages 40 to 60. All three of these glide paths have roughly the same shape, reflecting the basic intuition that the weight in stocks should decline as people age. By and large, the magnitude of the differences between the 110–Age path and the XYZ path are similar to the magnitude of the differences between the XYZ and static ESF paths.

Figure 6: Basic glide path comparison



Source: GMO

1.5 So why bother with ESF?


ESF matters because minimising expected shortfall of wealth provides a powerful conceptual framework that answers the right question in a customisable manner. To illustrate this point, consider Figure 7 where, in addition to the three glide paths shown above, we add a fourth – ABC. This glide path, in red, comes from another provider of target date funds and, while it follows the basic pattern that the weight in stocks falls as people age, it is more conservative in that it aggressively reduces weight in stocks as people age.

So which of these four choices is better?

The question is actually incomplete. We don't really know how XYZ and ABC were constructed. We know neither the assumptions about the plan participants, i.e., "What do they need and when do they need it?", nor do we know what objective, if any, these glide paths satisfy. Furthermore, we also don't know what assumptions were made about asset returns. We discuss this crucial issue in detail below.

But, for the ESF glide path, what we can say is that for a person who has circumstances and needs consistent with the assumptions articulated above regarding income, savings rates, and consumption, the ESF glide path minimises the expected shortfall of wealth – assuming at each age that the person is on wealth target and asset returns are constant. Because minimising shortfall of wealth is such a compelling common-sense objective, given this particular combination of "what and when," we believe the ESF path is better.

Figure 7: More glide paths




Source: GMO

This logic suggests that there may be particular combinations of "what and when" for which ABC, XYZ, or even 110-Age for that matter, are optimal in an expected shortfall sense. But there is simply no way to know. A person shopping for glide paths simply has no basis for choosing among the various possibilities. This points to the power of the ESF framework – given needs and circumstances, what do they need and when do they need it – we can generate optimal portfolios. Crucially, the objective is clear and sensible – to minimise, in expectation, how much wealth falls short of what is needed. Simply put, asking the right question leads to solving the right problem. ESF optimisation therefore provides a powerful basis for choosing portfolios over time.

To further illustrate this point about customisation and how it relates to asking the right question, suppose a particularly risk-averse individual does not have the constitution to tolerate a large equity exposure. Because our solution has no concept of a risk-aversion parameter, the answer, in contrast to that from MPT, is not to simply place this person in a portfolio with low volatility. Doing this, and this alone, will actually increase the risk that he does not have enough wealth in retirement. Instead, this investor has two other choices: (i) either reduce his wealth target, likely an undesirable choice; or (ii) increase his savings rate. In Figure 8, we show the ESF glide path, obtained by assuming as before that expected returns for stocks and bonds are 6% and 2%, respectively, for an individual who chooses the second option and saves at twice the rate assumed in generating the ESF glide path above. So, instead of starting to save at 5% and slowly increasing to 10% at retirement, the investor starts at 10% and slowly increases to 20% at retirement. We see in this case a glide path that has much less weight in stocks prior to retirement. But, during retirement, when the effect of the lower savings rate no longer applies, the glide paths converge.

Figure 8: A high savings rate glide path



Source: GMO

The ESF approach is also vital for correcting a fatal weakness in both the static ESF and the XYZ paths. Both are operating under a bad assumption – both assume that the expected returns are constant over time, when in reality they are anything but. However, given the common-sense objective of minimising shortfall of wealth in retirement, we have a framework that allows us to minimise shortfall of wealth throughout retirement even when expected returns are time varying.


PART II

2.1 Dynamic allocation: move your assets!

The fact that returns are not constant is self-evidently true for bonds, where an investor holding to maturity has an expected return overwhelmingly driven by the yield of the bond when it was purchased. (The only other issue to consider for that investor is the interest rate at which coupons will be reinvested.) Even if you are rolling your bond portfolio periodically (this might be the more relevant point because all DC investors are getting fixed income exposure from a bond fund), starting valuation is the overwhelming driver of subsequent returns.

But, valuation is every bit as important for stocks, and the valuations of stocks has varied hugely over time, as shown in Figure 9, which shows the cyclically-adjusted P/E, or "Shiller P/E," for the S&P500 over time.

Figure 9: US stock market valuation




Sources: Robert Shiller (see endnote 1), GMO analysis

The US stock market has averaged about 16 times normalised earnings over the last 130 years, but there have been times when it has traded far above or below that level. We believe it is the height of folly to assume that a market trading at 45 times normalised earnings, as the S&P 500 was in 2000, can achieve similar returns to one trading at seven times, as it was in 1982, let alone the expected returns of any reasonable glide path.

Stock valuations have been mean-reverting over time, and as a result, have a significant element of predictability to them. Valuation cannot tell us much about what returns will be over a week or a month or a quarter, but over a period of years, the importance of valuation steadily increases. This is illustrated in Figure 10 which shows the correlation between valuation and subsequent stock market returns as the time horizon lengthens from one year to 20 years.

Figure 10: Historical predictability of U.S stock market returns



Sources: Robert Shiller (see endnote 2), GMO analysis

While the correlation of 20% between current valuations and future 1-year returns is not particularly high, the correlation rises steadily with return horizon. For example, starting valuations have had a roughly 60% correlation with future 10-year returns. Given that the retirement savings problem is a 70-year problem, it is an ideal environment to take advantage of the long-term return predictability that comes from valuations. But this means we need a way to answer the question of how a 45-year-old should react if stocks are not priced to deliver 6% over inflation, but 2%, or 1%, or 7%. Mean variance optimisation in its standard form is limited in that it is a single-period optimisation, not to mention the problem raised earlier that it is not clear how to choose the level of return volatility, or equivalently, the level of risk aversion. But minimising expected shortfall, the ESF approach, gives us a natural framework for answering this question.


2.2 Historical simulations

It is perhaps most straightforward to understand this through an example using historical returns. Assume we have a worker who turned 55 in 1965 and to that point was on target for retirement savings. The worker, however, had the misfortune of being in peak savings years during a period in which equity valuations were high and real bond yields were low. The period from the mid-1960s through the 1970s represents some of the worst real returns for both stocks and bonds on record.

A static, or inflexible, glide path would ignore these lofty valuations and ensure the investor had a higher weighting in equities in 1965, when valuations were very high, than in 1974, when valuations were much lower. Figure 11 shows the static glide path and a fully dynamic

one from 1965 to 2005. The fully dynamic stock weight is based on minimising expected shortfall incorporating time-varying expected returns.

Figure 11: A Dynamic ESF case study – stock weights



Sources: Robert Shiller (see endnote 2), GMO analysis. The Dynamic ESF is based on comparing S&P 500 valuations to historical averages and allocating solely between stocks and bonds. Dynamic ESF is shown for illustrative purposes only. GMO does not manage any strategy that implements the Dynamic ESF methodology shown.

At first glance, the dynamic flight path looks nonsensical. It shows no weight in stocks in 1965, when the participant is 55 years old and has another 10 years until retirement but, by 1974, the stock weight rises to 90%, staying at a very high level until the early 1980s despite the fact that the participant is approaching 70 years old, and losses can be devastating. But if the goal is minimising expected shortfall of wealth in retirement, it can make sense to run an aggressive portfolio in retirement if the increase in expected returns is high enough. Furthermore, if the expected return to stocks is actually lower than bonds due to high valuations, such as was the case in 1965, it is hard to see why owning stocks would help at all.


The difference in outcomes between the two strategies is stark, as we can see in Figure 12. The static strategy leaves the participant out of money by 1992, the cost of not taking into account the changing valuations of the stock and bond markets over time, and consuming a constant real dollar amount equal to 5% of target wealth. The dynamic strategy, by contrast, allows the participant to make up for earlier inadequate returns and the money lasts for a full 30-year retirement even though consumption is high.

It is worth pointing out that this is not the best possible outcome for a participant from 1965. The allocations in these charts were not driven by any actual foreknowledge of

subsequent returns, but by very simple value-driven models for estimating expected future returns. Broadly, in our simplified study, which assumes annual rebalancing, as valuation multiples deviate from normal (e.g., a Shiller P/E of 16), younger investors, who have more weight in stocks than older investors, respond more aggressively to these valuation changes. For example, as indicated in Figure 6, when valuations are normal, a 25-year-old will have 90% weight in stocks, but a 65-year-old will have around 40%. If the Shiller P/E rises to 19, the weight in stocks for the 25-year-old will fall to about 45%, a drop of 45 percentage points, whereas the weight in stocks for the 65-year-old will fall to around 20%, a drop of about 20 percentage points.

And, in fact, it turns out that from 1965 to 1975, the dynamic allocation would not have helped relative to a static strategy. Both stocks and bonds did poorly over this period, so dynamic allocation would not have had much immediate impact. It wasn't until the early 1980s that stocks reaped the benefits of the cheap valuations thrown up by the 1973-74 bear market. But, when these benefits accrued to stocks, the impact on a dynamic allocation strategy was profound.

Figure 12: A Dynamic ESF case study – wealth




Sources: Robert Shiller (see endnote 2), GMO analysis

This brings up two points. First, cheap valuations don't guarantee that returns will follow particularly quickly – valuations mean revert slowly, typically reverting one-seventh of the way back to normal every year, meaning that stocks can remain cheap or expensive for very long periods of time. Second, asset allocation in retirement can be at least as important as it is during the accumulation phase.

If we go beyond this single example of starting in 1965 and use every possible starting

period for which we have US stock and bond market data, we can look at the probability of lifetime ruin given data since 1881. Figure 13 shows the amount of money a retiree would have at the age of 95, assuming he was on target at age 55 on the date shown on the horizontal axis. It's worth pointing out that while 1881 feels like a very long time ago, for the purposes of looking at a 70-year problem like retirement savings, it doesn't even provide two non-overlapping time periods. In the interest of maximising the data available, we've shortened this to a 40-year problem by assuming the worker was exactly on target at age 55 when we start the analysis.

Figure 13: Historical simulation of wealth at age 95



Sources: Robert Shiller (see endnote 2), GMO analysis

Both the standard glide path and the optimal static glide path leave plan participants running out of money before the age of 95 about half the time. The dynamic glide path is not perfect either, with a few starting dates in the 1950s and 1960s leaving the participant short. However, the probability of running out of money is about one fifth as high, as we can see in Figure 14.

Figure 14: Probability of lifetime ruin – historical simulation starting age 55


	Probability of lifetime ruin before 95
XYZ flight path	52%
Static shortfall minimisation	50%
Dynamic shortfall minimisation	13%

Sources: Robert Shiller (see endnote 2), GMO analysis. The table above shows, for the methodologies described in the prior pages, the probability of lifetime ruin (i.e., running out of money) before the age of 95. The Full 70-Year Monte Carlo uses simulations assumes \$43,000 starting income at age 25, annual real income growth of 1.1%, beginning savings of 5%, average savings of 9.7%, employer match of 3%, 40 years of savings, and a target retirement wealth of \$667,664. The Monte Carlo Starting at Age 55 uses the same assumptions except assumes the account is on target at age 55 and applies the methodologies from that point. Similarly, the Historical Simulation Starting at Age 55 also assumes the account is on target at age 55 and applies the methodologies from that point, in this case using Shiller Data for U.S stocks, bonds, and inflation from 1881 to 2011 and assumes a saver turns 55 in each year from 1881 to 1972. None of the above methodologies represents any strategy that is employed by GMO or any of its pooled vehicles.

The gap between a 13% and a 50% chance of lifetime ruin is a profound one. And even in the few cases where the dynamic glide path still led to ruin, the money lasted longer than in the static or standard glide path. As a reminder, this is not due to prescience involved in the dynamic process – it is only a simple value methodology that takes advantage of the predictability of stock returns due to mean reverting valuations. So it turns out that no matter when in history you happened to hit retirement, you would have been better off following a dynamic value-driven strategy.

Another way to look at that data is to look at the percent of all observations that have an ending wealth of more than a certain amount. Figure 15 shows the historical simulations in this manner – the vertical axis represents the percent of all observations that have wealth larger than the value on the horizontal axis. As an example, for the dynamic ESF strategy shown in red, approximately 90% of all observations have positive wealth at age 95.

Figure 15: Historical distribution of wealth at age 95



Sources: Robert Shiller (see endnote 2), GMO analysis


Furthermore, at every probability on the vertical axis, the dynamic glide path has more ending wealth at 95 years old than either of the static paths. The average amount of wealth is much higher as well, although this is almost an incidental effect of attempting to minimize the probability of a wealth shortfall.

2.3 Monte Carlo Simulations

The main advantage of showing the data this way is that it allows us to compare historical simulations, which suffer from the problem that we have far less data than we would ideally like, to a Monte Carlo simulation of the stock and bond markets where we can create as many histories as we would like.

One of our goals in building the simulations is to ensure stock valuations are approximately as predictive as they have been historically. Figure 16 shows one way to look at this.

Figure 16: Predictability of stock market returns – historical vs. simulated




Sources: Robert Shiller (see endnote 2), GMO analysis

The blue line in Figure 16 shows the same data as in Figure 10, demonstrating that by using a simple valuation method, the predictability of the stock market increases with time horizon. The red line shows the level of predictability between valuation and returns for our Monte Carlo simulation. We see that the red line is approximately on top of the blue – our simulation builds in approximately the same power for valuation as we've seen historically. (The theoretical underpinning of our simulation methodology and its relation to return predictability and its dependence on time horizon can be found in the Tarlie (2013) paper "[Discount rate dynamics and stock prices](#)").

Our other goal for the simulations is to have their valuation characteristics match those of history. This is illustrated in Figure 17 which shows five simulations of price to trailing 10-year real earnings (P/E10), along with the historical graph. The key point is that the size of the variations of our simulated price multiples match those of the historical series. The historical series has single-digit multiples, e.g., around 1920 and 1980, as well as multiples in the 40s, e.g., around 2000. The simulated series naturally capture the same dynamics as the historical series.

Figure 17: Valuation multiples – historical and simulated




Sources: Robert Shiller (see endnote 3), GMO analysis

This helps comfort us that our Monte Carlo simulations are reasonable and, with much more data, we can get more robust results than even 130 years of data can give. Figure 18 shows the equivalent of Figure 15 for the Monte Carlo simulations. Not surprisingly, the dotted lines, corresponding to the Monte Carlo simulations, are much smoother than the historical simulations. The slightly surprising finding, however, is that history turns out to be an unfortunate special case for 55-year-olds who were on target prior to retirement, where probabilities of lifetime ruin have actually been higher than one might expect. This is simply an unfortunate artifact of this particular, limited, run of history. The gap between dynamic and static glide paths remains, however, with the dynamic glide path having a 5% chance of ruin while the static glide path is four times as high at approximately 20%.

The simulations here – both historical and Monte Carlo – are simplified ones. A value-driven technique has much more potential to improve results when there are more than two asset classes from which to choose. Unfortunately, we don't have long enough data histories for asset classes beyond US stocks and bonds to run a decent historical analysis with a larger set of assets. On the other hand, our simulations assume that the weightings of the retirement portfolio are flexible between zero and 90% for stocks and 10% to 100% for bonds, which is probably beyond what most plan sponsors are willing to contemplate. However, the results are quite robust to changing this assumption in that any flexibility to move away from the fixed flight path will improve outcomes. For example, with a range of +/- 20%, the risk of lifetime ruin is reduced to 26% from 51% in the historical simulation and to 6% from 20% in the Monte Carlo simulation.

Figure 18: Distribution of wealth at age 95
Historical vs. Simulated, Dynamic vs. Static



Sources: Robert Shiller (see endnote 4), GMO Analysis

The structuring of the problem is also a bit oversimplified. Rather than deal with the uncertainty of longevity risk, we simply assumed that participants live 30 years in retirement. This longevity risk should not be ignored, and insurance products such as deferred annuities can be a cost-effective way to reduce the longevity risk that exists for a single plan participant. We have also assumed that investors don't change either their savings or consumption decisions in response to the financial markets.

While we do expect people to adjust their behavior in response to changing circumstances, our focus here is on the role of an investment program, so holding behavior constant is crucial to assessing this. Furthermore, the modeling of investor behavior as it relates to savings and consumption is a difficult and complicated problem that requires more study.

Finally, this paper illustrates the impact of dynamic allocation by focusing on a 55-year-old who is on target prior to retirement. The reason for this focus is, in part, because of the limited amount of historical data. But, the Monte Carlo simulations allow us to look at all ages. And, although we do not present those results here for lack of space, the results for the 55-year-old are representative of what we find for all ages – dynamic allocation is an essential tool.

CONCLUSION

This research brings up several useful points to consider in building retirement portfolios. First, asking the right question is extremely important. The retirement savings problem is

complex, involving two long phases, one of accumulation and the other of consumption. As such, it is essential to go beyond a brute-force approach in designing a glide path to deal with time-varying expected returns and the changing sensitivities, based on needs, as participants age. A common-sense, holistic approach of minimising expected shortfall of wealth in retirement enables the construction of highly customisable portfolios as both investor needs and asset valuations change. The essential point is that because the objective is to minimise a shortfall of wealth, this framework is customisable in a way that ensures that the resulting solution makes sense precisely because it is solving the right problem.

Second, while most plan sponsors focus on building their portfolios appropriately for their employees while they are working, portfolio decisions made during the retirement phase are every bit as important as they are during the accumulation phase – and perhaps even more so. After all, the returns, good or bad, achieved by the retirement portfolio of a 28-year-old will have a limited impact on the ultimate ability of the portfolio to support spending in retirement for the simple reason that they affect only a few years of relatively small contributions. The vast majority of the wealth accumulated at 65 will have been generated off of the contributions made in later years. In retirement, however, the participant has already accumulated all the wealth he will ever have, and returns on that portfolio will have a profound impact on what that person can ultimately spend. Unfortunately, there seems to be much less effort focused on the post-retirement side of things – understandably so for plan sponsors, to be sure, but from a societal perspective, it would be an excellent idea to find a way to ensure that post-retirement portfolios are managed with the same care as they are pre-retirement.

And third, the risk of failure with the traditional glide paths and savings/spending assumptions seems to us to be disturbingly high. Increasing participant savings rates will be crucial to helping ensure retirement success for today's working population, but valuation-aware portfolios, as our research has shown, can make a huge difference in the probability of success. We believe plan sponsors are ignoring an incredibly powerful tool to help their participants if they build their portfolios without taking asset class valuations into account. And, given that workers have very few other effective ways to save for their own retirement, the stakes seem too high to leave such a valuable tool on the table.

Ben Inker, CFA, is co-head of Asset Allocation team at GMO and a member of the GMO Board of Directors. Dr Martin Tarlie, CFA, is a quantitative researcher for GMO's Global Equity team.

This article is abridged and reproduced with permission from [Advisor Perspectives](#).
