GMO White Paper

Aug 2015

Who Ate Joe's Retirement Money?

Sequence Risk and its Insidious Drag on Retirement Wealth

Peter Chiappinelli
Ram Thirukkonda

Summary

Defined Contribution (DC) plan participants are haunted by an invisible risk called sequence risk (sometimes called sequence-of-returns or path dependency risk), that is, getting the "right" returns but in the "wrong" order. Sequence risk in the retirement phase has been studied extensively. Sadly, not as much attention has been paid to sequence risk during the accumulation phase, but it is equally important. Sequence risk rears its head in this way: Even if an individual employee does everything "right" – participates in the plan, defers income religiously, takes full advantage of the company match, and even gets his exact expected return from his investments – he can still fall victim to disappointing final wealth outcomes if the order of those returns works against him. Current models of asset allocation – the most popular being static, or predetermined, target date glidepaths – "know" that sequence risk exists, but behave as if there is nothing that can be done to mitigate it. Valuation-based dynamic allocation, on the other hand, can help soften the bite.

Here's a riddle for you: Who ate almost $300,000 of Joe's retirement money?

Meet some pretend employees, Joe and Jane. They worked for Acme, Inc., and were identical in essentially all aspects of their job, their salary, and their participation in their company's defined contribution retirement plan (a "typical" 401(k)). Here are some key dimensions to consider:

- Identical length of time working: Both started working at 25, and both worked for 40 years.
- Identical starting salary and salary growth.[1]
- Identical deferral rates and identical company match[2] for a combined 9% annual contribution.
- Identical investments: Both invested in a typical target date fund (TDF), which started out with a 90% equity allocation, and glided down to below 50% by retirement age.
- Identical returns: This is a *key* point. During their 40 years of investing, they both earned exactly 5% annualized real (above inflation, or roughly 8% nominal) after fees.

[1] National Association of Colleges and Employers, Average Starting Salary Survey, 2012.

[2] Hewitt Associates, Trends and Experience in 401(k) Plans, 2009. Throughout most of recent history, a 50 cent match per dollar on the first 6% deferral was the most common match formula. Recent surveys have indicated plans are moving toward more generous matches, meaning cash flows will be even more important going forward, as they relate to sequence risk.

Identical in virtually every way. At the end of their 40-year careers, however, Jane had $880,000 in her account, while Joe had $590,000. How is this possible? Who ate $290,000 of Joe's retirement money? How do you explain a 50% gap between these two employees?

The answer? Sequence risk.

Here's what we didn't tell you. Joe started his career in 1954, while Jane started in 1967, 13 years later. So, even though Joe earned the same annualized return as Jane, he earned it in a slightly different sequence, and that made all the difference. Sequence risk – an insidious risk in all DC plans – took a sizeable bite out of his potential retirement nest egg.

Exhibit 1: Same Return, Different Results
Who Ate $290,000 of Joe's Retirement Money?

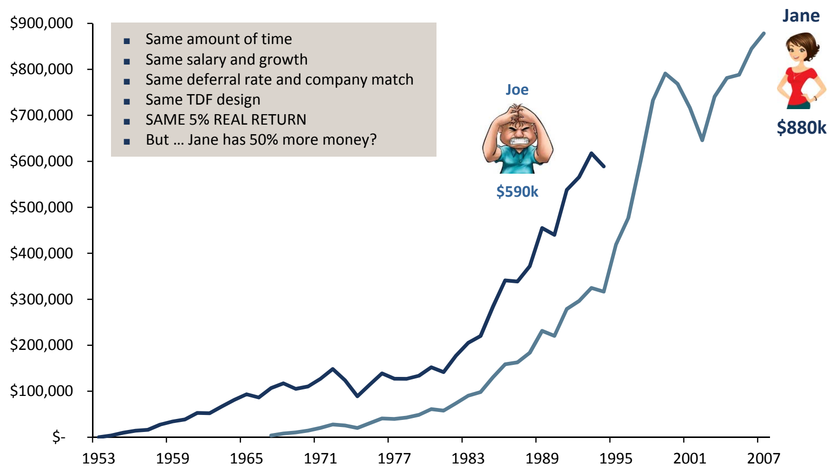

Source: GMO
Assumes that each employee began with a 90% equity/10% bond portfolio, which glided down to a roughly 50% equity/50% bond mix at retirement. Stocks are represented by the S&P 500; bonds are represented by the 10-year U.S. Treasury bond.

Sequence risk has been in the shadows for some time. One reason is that sequence risk is typically not a major concern for traditional Defined Benefit (DB) plans. (See sidebar discussion regarding DB plans on page 11.) Second, when it has been discussed in the academic or investment community, the focus has traditionally been on the *withdrawal* phase of retirement,[3] when cash flows are large. The main thrust of those studies demonstrated that the returns a retiree experiences in the first few years of his/her retirement are extremely important. A significant loss early on, even if it is recouped later, dramatically increases the risk that a retiree will run out of money. Sequence risk is undeniably important in the retirement phase, but most analyses simply start with an assumption that the retiree begins with some large lump sum. But this glosses over the fact that in a DC environment, it takes about 40 years of contributions, matches, and market returns to get to that final lump sum. Sequence risk rears its ugly head wherever cash flows matter – and we know cash flows matter both in the retirement and accumulation phases.

[3] Dr. Wade Pfau, "Sequence Risk vs. Retirement Risk," RetirementResearcher.com, March, 2015; Larry R. Frank, John B. Mitchell, and David M. Blanchett, "Probability-of-Failure-Based Decision Rules to Manage Sequence Risk in Retirement," Journal of Financial Planning, Vol 24, Issue 11, p 44, November 2011; R.G. Stout, "Stochastic Optimization of Retirement Portfolio Asset Allocations and Withdrawals," Financial Services Review, 2008; and Dr. John B. Mitchell, "Withdrawal Rate Strategies for Retirement Portfolios: Preventive Reductions and Risk," presented at the Academy of Financial Services, October, 2009.

This paper tries to demonstrate the importance of sequence risk during the *accumulation* phase. The basic message is this: Even employees, like Joe, who apparently do everything "right" by traditional playbooks – stay in the plan, defer their income religiously, take full advantage of their company match, and even get *the exact expected return from their investments* – can still suffer from the effects of sequence-of-returns risk. That is, they get good returns, but they get them in a bad order (or, more specifically, they get good returns early in their career, and bad returns later when their account balance is higher).

Analysis: quantifying the significance of sequence risk

The Jane and Joe example above is interesting, but it only represents one "run" of history. Another method for measuring the impact is to simulate multiple runs, even thousands, through a stochastic process (akin to Monte Carlo simulations). We can artificially create 20,000 simulations of 40-year runs of history. Before we begin, however, let's remind ourselves that as it relates to returns, employees confront both *investment* risk and *sequence* risk. Investment risk, the variability and distribution of possible returns around the mean, or expected return, is easily observable. Sequence risk, on the other hand, is much more insidious and harder to observe. We need to isolate it in order to see its significance.

We'll start with a simple case, as above, by constructing a very typical target date allocation structure.[4] We ran 20,000 simulations,[5] and the output (see Exhibit 2) is a distribution of both returns (investment risk) and wealth outcomes (a function of both investment risk and sequence risk). The fact that there is a wide distribution of returns and possible wealth outcomes should not be surprising to anybody. Investment risk is well understood.

Many, however, might be surprised by the magnitude of sequence risk, which we can isolate by focusing on those simulated histories where the realized returns are identical. In our 20,000 simulations, for example, there were close to 1,400 instances where the realized returns were 5% real.[6] Fourteen hundred Joes and Janes, if you will. Yet even if we control for a given realized return, *wealth* outcomes are still widely dispersed (see the dotted orange box in Exhibit 2). This is the effect of sequence risk. Each of the 1,400 runs occurred in a unique sequence, some advantageous, some less so. The "luckiest" employee's 5% return netted her over $1,000,000, while her unluckiest colleague, with the identical 5% return, netted $314,000, a massive, almost unimaginable discrepancy.

[4] Our "model" target date glidepath begins with an allocation of 90% equities at ages 25 to 35, glides to 85% equity at age 40, 78% at age 45, 70% at age 50, 63% at age 55, 53% at age 60, and 40% at age 65. This is in line with what major glidepath providers do in the industry.

[5] Key assumptions in the exercise are the following: a) a simple two-asset-class portfolio; b) expected stock return is 6% real, with 18% volatility; c) expected bond return is 2% real, with 5% volatility; d) the correlation between stocks and bonds is zero; e) the portfolio is rebalanced annually; f) the demographic data assumes a starting salary of $43,000 and a 1.1% real rate of growth for 40 years; and g) the deferral rate and company match equate to 9% contributions.

[6] We looked at realized returns between 4.9% and 5.1%, or plus/minus 10 basis points around 5% real.

Exhibit 2: Sequence Risk
Target Date Glidepath: 20,000 Simulations

Source: GMO Simulation

So now let's isolate the orange box and see not just the extremes, but what the entire distribution looks like (see Exhibit 3). It is surprisingly vast. Even the "meat" of the distribution could easily result in a $200,000 to $250,000 gap between experiences. For many, this could be the difference between financial security and financial ruin (i.e., running out of money).[7] The real focus, however, should not be on the middle or the outcomes on the right "tail," what we have called the "Lucky" Cluster. No, the problem is really with the left tail of the distribution, what we are calling the "Unlucky" Cluster of outcomes. Here, the risk of financial ruin is palpable.

Exhibit 3: Simulated Wealth Outcomes
Realized Returns of 5% for 40 Years

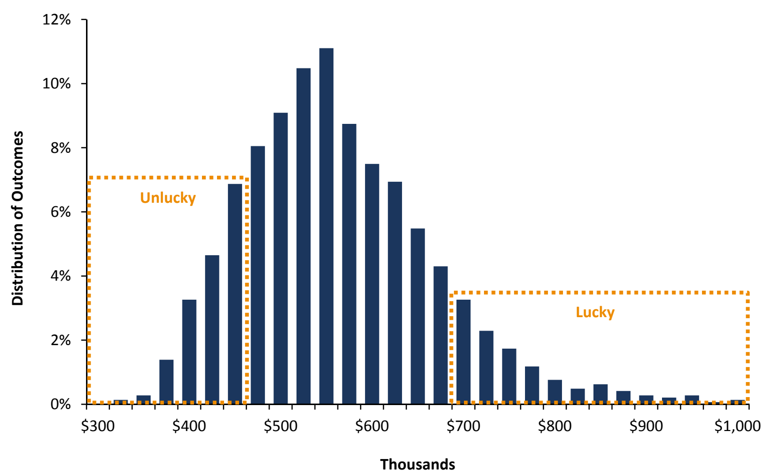

Source: GMO Simulation

[7] Defining risk as "financial ruin" (i.e., running out of money) is addressed at length in a series of papers written by our colleagues. See Ben Inker and Martin Tarlie, "Investing for Retirement: The Defined Contribution Challenge." Also see Jim Sia, "The Road Less Traveled: Minimizing Shortfall and Dynamically Allocating in a DC Plan." Each of these papers can be found at www.gmo.com.

Now, the general intuition behind sequence risk – getting the right returns but in a bad order – is that those unlucky outcomes are likely driven by poor returns or nasty drawdowns in the final phases of someone's career. Let's call this period the "Final 10 Years." And the logic is pretty sound. Early in a career, market returns do not really matter too much for the simple fact that there is not much money in the DC account. Cash flows (employer and employee contributions), on the other hand, are the key driver of growing wealth. As an employee approaches mid-career and beyond, the account base is large enough that cash flows (as a percent of that base) are becoming less important while market returns rise in importance. Exhibit 4 confirms the intuition. The Lucky Cluster (those who experienced the full 40-year 5% annualized returns and ended up with very high wealth) tended to get solid, or even fantastic returns in their Final 10 Years. The Unlucky Cluster, the "Joes" of the world, also earned their 40-year 5% annualized returns overall, but got sub-par and, in many instances, negative returns for their Final 10 Years. Returns in the Final 10 Years have a high correlation with terminal wealth, so it's important to avoid or mitigate drawdowns during any 10-year time frame (remember, ANY 10-year time frame is *somebody's* Final 10 Years).

Exhibit 4: Final Wealth vs. the Final 10 Years - Drawdowns and Bad Returns Matter
(Reminder: All "Dots" Received 5% Annualized Real Returns over the Full 40 Years)

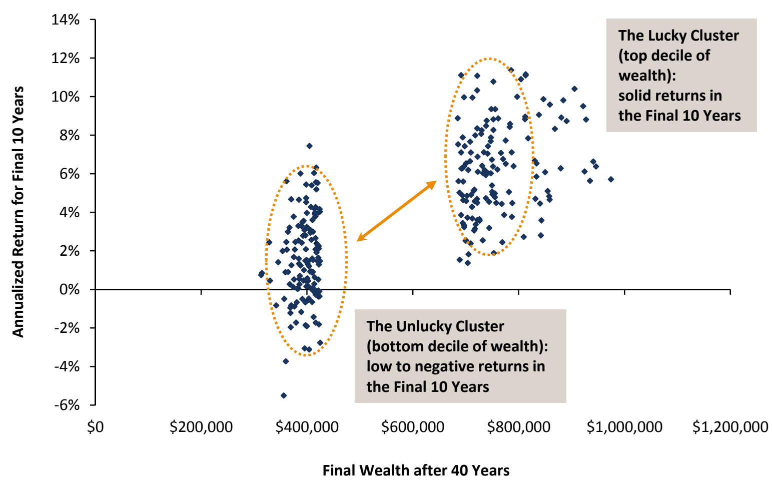

Source: GMO

Static portfolios shrug their shoulders

Okay. Sequence risk – drawdowns and nasty returns late in a career – is a problem. And sequence risk is really unfair to some unlucky employees. So now what?

Unfortunately, there is an entire school of thought that believes sequence risk is just an unfortunate problem. It exists, the argument goes, but there is really nothing one can do about it. In fact, this is an assumption built into every single predetermined glidepath in a static TDF.[8] The assumption underlying their design is that there is no way to forecast future returns. These glidepaths typically

8 Yes, static TDFs do change their allocations through time, but they do so in a predetermined manner. There is no judgment involved. In other words, we know exactly what the equity allocation is going to be 10 years from now or 20 years from now, and the mix will give no consideration whatsoever to current market environments or current valuations of stocks or bonds at that time. All the TDF "knows" is that a participant is of a certain age, and is therefore forced into an x% allocation to equities and 1-x% into bonds.

make no attempt to adjust their asset mix based upon new information. It is one of the cornerstones of Efficient Market Theory that future returns are largely a "random walk." Today's price tells you nothing about future returns. It makes no sense, the argument continues, to adjust a portfolio's asset allocation because future returns are unknowable. In essence, these TDF portfolios shrug their shoulders in the face of sequence risk and say, "Oh, well." It's just a hope and a prayer that your nasty returns don't occur in those Final 10 Years.

Value…might help

We at GMO have a very different belief. Namely, that future returns are the furthest thing from a "random walk"; we can use well-established valuation metrics to forecast returns and help mitigate potential nasty returns in those Final 10 Years. The Cyclically Adjusted P/E ratio (or CAPE) is one such simple metric for determining whether stock markets are cheap, fairly-valued, or expensive[9] (see Exhibit 5). A strategy of owning equities when P/Es were low (i.e., cheap) tended to outperform, while owning equities when P/Es were elevated (i.e., expensive) tended to do quite poorly. There are no useful or reliable tools, we believe, for forecasting returns in the short term, say, one or two years. But we believe the mean-reverting nature of P/Es over longer time horizons means that starting valuations correlate quite strongly with future returns. Exhibit 6 shows that correlations rise to close to 0.6 in the 7- to 10-year range. Therefore, "value" can be quite useful for adjusting a portfolio during anybody's Final 10 Years.

Exhibit 5: S&P 500 Index Returns
From Different Starting Valuation

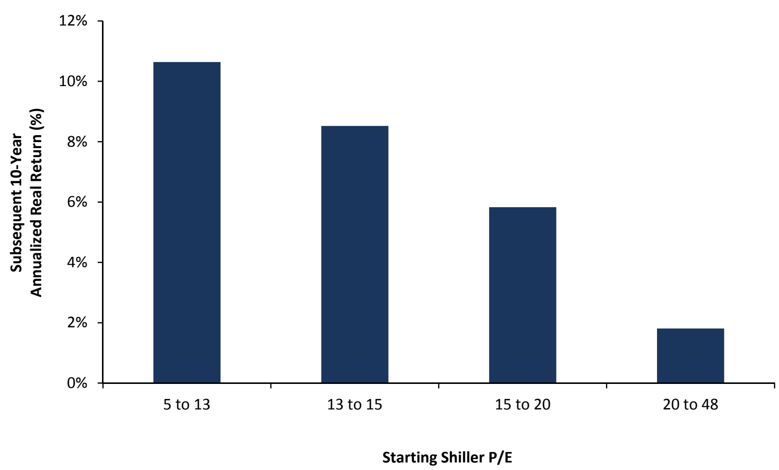

As of 3/31/15
Source: Shiller, GMO
The annualized 10-year forward real returns shown above are calculated using real price, real dividends for the S&P composite index, and cyclically adjusted P/E from the standard Shiller data source (http://www.econ.yale.edu/~shiller/data/ie_data.xls).

[9] See our colleague James Montier's paper, "A CAPE Crusader: A Defence Against the Dark Arts," for a spirited discussion of Shiller's Cyclically Adjusted P/E as a reliable valuation tool. His paper can be found at www.gmo.com.

Exhibit 6: Mean Reversion of P/E Suggests that Future Returns Are Forecastable

Source: GMO. The graph above shows the correlation of simulated expected returns (based on simulated market dynamics) to the return outcomes of Monte Carlo simulations and historical data mapped against return horizons. The historical data uses Shiller data (http://www.econ.yale.edu/~shiller/data.htm) for U.S. stocks, bonds, and inflation from 1881 to 2013. None of the methodologies represent any strategy that is employed by GMO or any of its pooled vehicles.

More importantly, however, and perhaps more relevant, valuation-based sensitivity to the attractiveness of stocks can also help mitigate drawdowns. Cheap assets, as it turns out, tend to hold up a bit better during a market tumult. Stocks typically become cheap for a reason, usually because something bad is happening or is feared it will happen. As a result, valuation multiples drop (i.e., stocks get cheaper). Subsequent to this drop in valuation, if something bad actually does occur, much of the bad news has already been discounted. History bears this out as shown in Exhibit 7, which takes a look at starting CAPE ratios over the last 70 or so years, and tracks how different starting valuation levels performed during drawdown periods. Cheap stock markets suffered drawdowns, of course, but they tended to suffer quite a bit less than expensive markets.

Exhibit 7: Value Can Help by Dampening Drawdowns:
Three-Year Drawdown, from Starting Valuation* (S&P 500, from 1940)

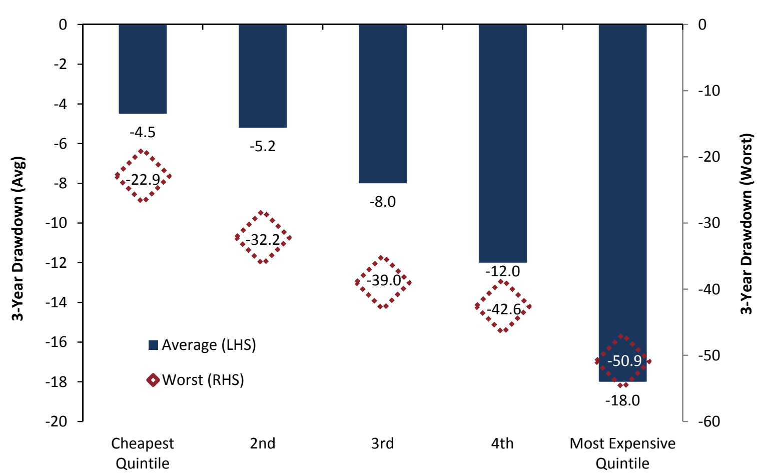

Source: GMO, Hussman Funds, 1940-2012
*10-Year P/E
Cheapness/expensiveness is determined by the Cyclically Adjusted Price/Earnings (CAPE) ratio, otherwise known as Shiller P/E.

Back to the real world

Stochastic modeling is insightful, but at the end of the day, it is just a model of reality. So now let's return to actual capital markets history, but with a twist. With the observations that starting valuation can be a useful tool in forecasting future returns and that owning stocks when they are cheaper tends to win in the long term with the additional benefit of suffering less damaging drawdowns, let's conduct a simple experiment. Suppose we use starting CAPEs as our simple and sole indicator of future stock market returns. Let's go back in history as far as we can and pretend that we are managing the glidepath, but this time, we will manage it *dynamically* – that is, underweighting or overweighting equities and fixed income around the predetermined path, based on one simple valuation metric, which changes through time. (Note that we constrain the shifts to +/- 20%, where applicable.) Basically, if we believed stocks were expensive, we would underweight them, and if stocks appeared cheap, we would overweight them. See Exhibit 8 for a stylized illustration of the type of "freedom" we would have to alter weights around the predetermined path. We have market data that goes back to around 1881, which is the same data set that Robert Shiller uses in his CAPE studies. The key question is: Can we do better than the static (predetermined) glidepath in terms of drawdowns and final wealth?

Exhibit 8: Dynamic Glidepath: +/- 20% "Channel"

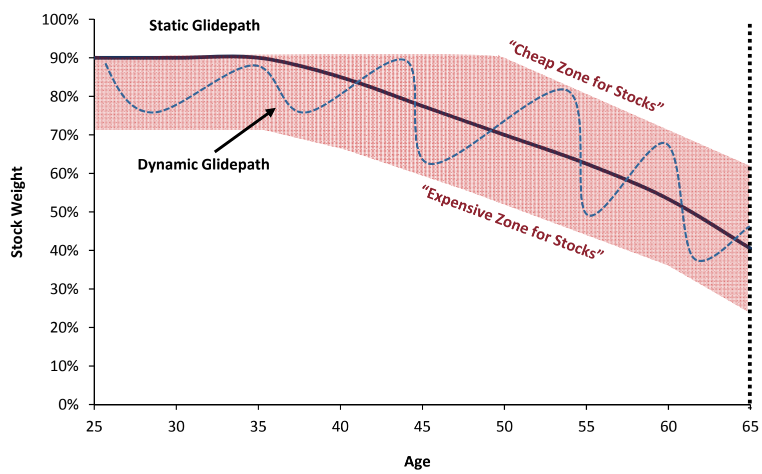

Source: GMO
Cheapness/expensiveness is determined by the Cyclically Adjusted Price/Earnings (CAPE) ratio, otherwise known as Shiller P/E.

Looking at Exhibit 9, the answer is apparently yes. Using a simple valuation tool, we can help mitigate drawdowns and therefore reduce the impact of nasty returns during the Final 10 Years. The Dynamic Glidepath tended to have lesser drawdowns than the static one, in terms of frequency and magnitude. It had lesser drawdowns two-thirds of the time, and the magnitude of its lesser drawdowns was dramatic in many instances.

Exhibit 9: Drawdown During Final 10 Years of Glidepath
Dynamic vs. Static – Which Lost Less and by How Much

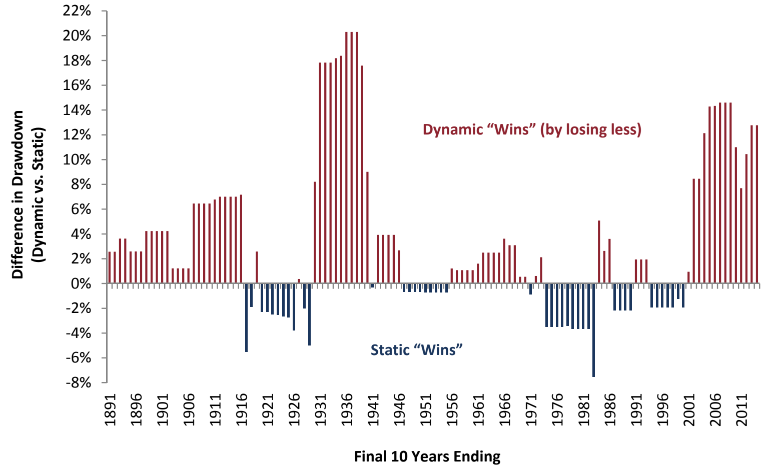

As of 2014
Source: GMO

Rotating into cheap equity markets and away from expensive ones, it turns out, also helps to improve overall wealth effects. Using the same historical data set, we can look at the terminal wealth outcomes for a full 40-year run of history. See Exhibit 10, which shows total wealth outcomes from 1881 to present day, comparing dynamic portfolios to their static counterparts. (Please note that we know these are overlapping periods, but for purposes of discussing sequence risk, the results are still quite relevant.) The Dynamic portfolio tended to have better total wealth outcomes. For the almost 100 runs of history, the Dynamic portfolio ended up with higher wealth three times as often. It is as if the entire distribution, the “cloud” of outcomes (as seen in Exhibit 2), so to speak, was lifted. [Note: To repeat, this improvement came from using one simple valuation tool; it would not be a huge leap to assume that using more asset classes and other valuation-sensitive techniques could improve things further.]

Exhibit 10: Terminal Wealth, Static vs. Dynamic Glidepath

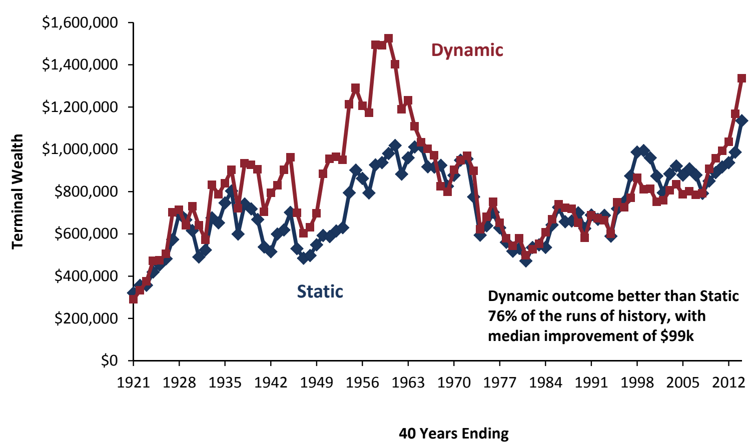

As of 2014
Source: GMO, Shiller

In statistical space, a better way to describe what is happening can be found in Exhibit 11. In this exhibit, you see a literal shift of the distribution to the right. Now, is this truly "solving" for sequence risk? Well, not directly. It's a bit more complicated than that because the benefits show up in a number of ways. First (but of least significance), the average or mean outcome is higher by roughly $100,000. Second, the median (a more meaningful metric) also improved by close to $90,000. The real benefit, however, is that the left tail of the distribution shifts to the right – there are fewer "unlucky Joes" and better outcomes for the few that remain. (It is interesting to note that the standard deviation of wealth outcomes actually increases – the lucky got luckier, so to speak – but this is not a surprising outcome given that a dynamic approach should amplify the outcome for almost everybody, including the already lucky ones.) By tilting the portfolio's allocation away from expensive equity markets (and conversely, tilting the portfolio toward cheap equity markets), we can help mitigate the effects of sequence risk, most importantly, for the unlucky cohorts.

Exhibit 11: Distribution of Wealth Outcomes*
Historical Data, Static vs. Dynamic Glidepaths

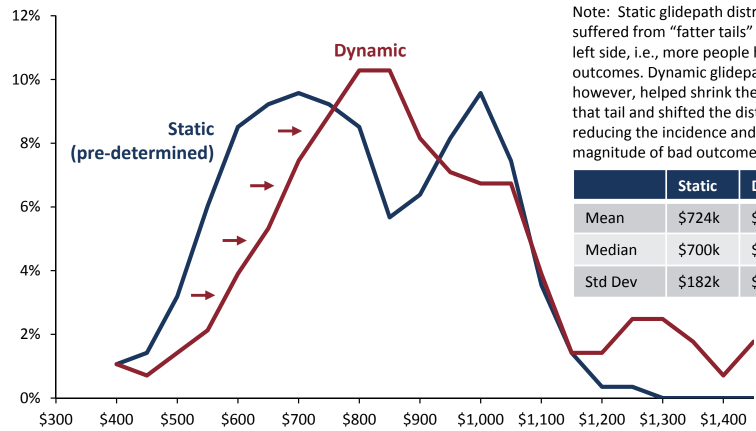

Note: Static glidepath distributions suffered from "fatter tails" on the left side, i.e., more people had bad outcomes. Dynamic glidepaths, however, helped shrink the size of that tail and shifted the distribution, reducing the incidence and magnitude of bad outcomes.

	Static	Dynamic
Mean	$724k	$822k
Median	$700k	$789k
Std Dev	$182k	$250k

As of 2014
Source: GMO
*Using historical returns and valuations. Each wealth outcome was a 40-year "run" of history, using data from 1881. Periods are overlapping.

Conclusion

The story about Jane and Joe is, of course, made up. But assigning real names to imaginary individuals was done on purpose. We are morphing from an era of *pooled* risks in a DB plan, to a DC era of *individually*-owned risks. Sequence risk is just one of these risks, but an important one. With this transition comes a need for heightened sensitivity to the potential for financial ruin; that "left tail" is not some nameless and faceless actuarial cohort: It is a collection of specific individuals with very specific DC account balances. Risk is no longer just the classic definition of "standard deviation of returns" (if it ever was); no, it is much more personal. It is the risk that individuals may not "get" the right returns they need for a healthy retirement, or, more cruelly, they get them, but in the wrong order. This risk of financial ruin is the right prism through which to examine the issue and look for DC solutions. A valuation-sensitive approach to dynamic asset allocation can help mitigate these "new" risks that have been foisted upon the unlucky individual Joes of the world.

DEFINED BENEFIT PLANS AND SEQUENCE RISK… WHAT, ME WORRY?

Hopefully, we've illustrated why sequence risk is a big deal for any DC participant. So why is it NOT a big deal for a traditional, or typical, DB plan? After all, it, too, has to face the sometimes unkind and volatile nature of the capital markets. By and large, however, DB plans are relatively stable pools of assets in terms of cash flows. Yes, DB plans have money coming in from sponsor contributions and money going out to pay beneficiaries, but these cash flows either tend to cancel each other out or, more commonly, their net flow in either direction is de minimis relative to the large underlying asset base. In other words, cash flows don't matter (or at least not that much). And where cash flows don't matter, the order in which a DB plan gets its returns doesn't really matter either. Let's look at a hypothetical DB plan (see table below). Imagine a 10-year series of returns and note the growth of the $100 asset base, which grows to $163, experiencing a 5% annualized rate of return. We call that Return Sequence #1. Now let's take that sequence and reverse it, or flip the order. Let's call that Return Sequence #2. While the path of returns is different, the terminal wealth is identical. Because there were no interim cash flows during those 10 years, the order in which a plan experiences its returns literally does not matter. The chart makes the same point, taking the same annual returns, but jumbling them randomly to create many different sequences. All paths ended up with a 5% annualized return, but their paths did not matter because there were no interim cash flows. Sequence risk for a large, stable pool of assets is not really a worry at all. (Please note that if for some reason a DB plan is experiencing massive cash flows in either direction, perhaps due to large contributions to help raise its funding status, then sequence risk is back on the table.)

Hypothetical DB Plan

Year	Return Sequence #1	Wealth	Return Sequence #2	Wealth
0		$100		$100
1	5.0%	$105	-9.5%	$91
2	-12.0%	$92	6.0%	$96
3	22.0%	$113	20.0%	$115
4	3.5%	$117	14.0%	$131
5	12.0%	$131	-5.0%	$124
6	-5.0%	$124	12.0%	$139
7	14.0%	$141	3.5%	$145
8	20.0%	$170	22.0%	$176
9	6.0%	$180	-12.0%	$155
10	-9.5%	$163	5.0%	$163
10-year Annualized Return	5.0%		5.0%	

10 Years of Returns in Different Orders

Lots of Different Return Paths, but Same Wealth Outcome

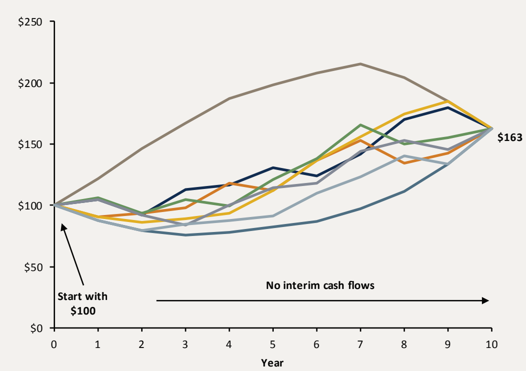

Source: GMO

Peter Chiappinelli: Mr. Chiappinelli is a member of GMO's Asset Allocation team. Prior to joining GMO in 2010, he was an institutional portfolio manager on the asset allocation team at Pyramis Global Advisors, a subsidiary of Fidelity Investments. Previously, he was the director of institutional investment strategy and research at Putnam Investments. Mr. Chiappinelli earned his MBA from The Wharton School at the University of Pennsylvania and his B.A. from Carleton College. He is a CAIA charterholder, and was the founding President of the CAIA Boston chapter. He is a CFA charterholder.

Ram Thirukkonda: Mr. Thirukkonda is a member of GMO's Asset Allocation team. Prior to joining GMO in 2014, he was a quantitative analyst at Batterymarch Financial Management. Previously, he worked at Fidelity Investments, most recently as the Director of Architecture. Mr. Thirukkonda earned his Bachelor of Technology from the Indian Institute of Technology, his MS in Transportation Engineering from MIT, and his MS in SDM focusing on Financial Engineering from MIT. He is a CFA charterholder.
