LDI Misapplied
Income Portfolios and Liability-Driven Investing

Morningstar Investment Management LLC
September 2017

Contents

Thomas Idzorek, CFA
Chief Investment Officer—Retirement
+1 312 696-6840
Thomas.Idzorek@morningstar.com

David Blanchett, Ph.D., CFA, CFP®
Head of Retirement Research
David.Blanchett@morningstar.com

Abstract

Liability-driven investing (LDI)—in particular liability-relative optimization—represents a fundamental improvement over more common asset-only portfolio optimization techniques, such as mean-variance optimization. Almost all portfolios exist to pay for some future form of consumption, so liability-relative optimization is almost always more appropriate than asset-only approaches. By considering liability characteristics when solving for the asset allocation, LDI techniques take advantage of the natural hedging quality of certain investments.

Well-meaning, avant-garde practitioners have begun embracing LDI techniques when building portfolios for individual investors (especially retirees) without considering the unique characteristics of the individual's liability or the risk attributes of the assets or cash flows retirees have available to fund the liability (e.g., Social Security retirement benefits). The result is highly conservative portfolios like those used by institutional investors. But these investors tend to have more predictable liabilities, with idiosyncratic risks leveling out across a large population of retirees.

Individual investors are subject to greater idiosyncratic risk in their retirement spending needs, and must also invest appropriately to hedge risk in an LDI context. Importantly, individuals have greater flexibility in retirement spending, in part because retirement expenses aren't a legal liability, and because many have other ways to pay for retirement expenses, including other assets, Social Security retirement benefits, and the ability to work during retirement. After investigating better ways to properly model investor expenses, we find that most investors at or near retirement are likely better served with less conservative, more balanced, and more diversified portfolios that recognize the increased duration associated with living longer.

Introduction

Many of the most sophisticated institutional investors—especially insurance companies and defined-benefit pension plans—have embraced asset allocation techniques that expressly acknowledge that the portfolio has a specific purpose or goal: to pay for an expense or set of ongoing expenses. These ongoing expenses typically are a form of contractual or legal liability. As such, they change the nature of investment risk facing the investor; what matters isn't necessarily the variability of the investment portfolio in isolation, but rather the portfolio's ability to pay for ongoing liabilities. This approach to asset allocation or portfolio formation is referred to as "liability-driven investing," or "LDI." One might characterize LDI as a form of holistic goals-based or outcome-based investing, in the parlance of our times, in which the baseline goal is to produce a flexible, inflation-adjusted consumption stream for life.

Common LDI approaches include the ultra-conservative approach of cash flow matching (matching the timing and size of cash flows from the assets with the required cash flows of the liability), followed by duration matching (matching the interest rate sensitivity of the asset portfolio to that of the liability), and liability-relative optimization, also called surplus optimization. Liability-relative optimization is the most sophisticated and the most flexible of the three approaches. While we have characterized cash flow matching as "ultra-conservative," given that it typically involves a zero-coupon, high-quality bond maturing at the time of the expected cash flow, for an individual with an uncertain or variable consumption pattern (water pipes don't burst on a schedule, even in retirement) the specificity of a cash flow matching investment strategy may actually be relatively risky.

Investors seemingly rediscover the need for liability-driven investing following each financial crisis, with the tech crash bubble in the early 2000s and the 2008 financial crisis being two of the more recent examples. Historically, investors have been most inclined to adopt a liability-driven investing approach when there is legal (or hard) liability and the probability as well as the ramifications for not being able to pay for the liability are significant. Thus, common practitioners of LDI include insurance companies and defined-benefit pension plans, in which there is a contractual, legal obligation to make (nearly non-negotiable) payments. Ironically, LDI may be easier to implement for large insurance companies or defined-benefit plans as the law of large numbers mitigates the type of idiosyncratic events that impact individuals.

More recently, well-meaning practitioners have started using liability-driven investing approaches to aid in the creation of asset allocations for individual investors in which a desired income level represented by ongoing living expenses is treated as an "individual retirement income liability." Importantly and in contrast to the hard legal liabilities faced by insurance companies and defined-benefit pensions plans, the desired income level isn't a true liability—it represents an aspirational target that can be adjusted downward (or upward) if necessary.

We fear that some early adopters of LDI for individuals may have overlooked the unique aspects of the situation faced by an individual investor. These practitioners may have misapplied LDI,[1] which may lead to portfolios that in many cases are overly conservative and thus fail to meet their goal of funding a comfortable retirement. The goal of this paper is to explore the unique aspects of properly applying LDI to the creation of an asset allocation for an individual investor, which, to our knowledge, has not been thoroughly explored in the existing literature.

Liability-Relative Optimization

The theoretical advantage of LDI over the asset-only optimization framework is depicted in Exhibit 1. The top panel represents the asset-only approach and the bottom panel represents the liability-relative approach. On the left side of both panels, the blue line represents the evolving true economic value (net present value) of the liability and is identical in both panels. In the top left graph, we see that the asset-only approach leads to a portfolio of assets with a value represented by the orange line that may not always move in the same direction as the value of liabilities because the portfolio of assets is determined in isolation, with no knowledge of the liability. This, in turn, leads to a portfolio represented by the wavy green line whose health/value (and/or the cost associated with funding the portfolio) can vary significantly over time, as there can be large gaps between the value of the liability and the value of the assets. Pension plans experienced this sort of double whammy after the 2008 global financial crisis, when market declines reduced asset values while plunging interest rates raised the value of liabilities.

In contrast, in the bottom left graph, we see that the liability-relative approach can lead to a portfolio of assets with a value represented by the orange line that is expected to move in unison with the value of the liabilities because the portfolio of asset is determined in a single optimization that is expanded to include the liability as part of the total portfolio. This leads to a total portfolio whose health/value (and/or the cost associated with funding the portfolio) represented by the constant green line that is steadier and more predictable over time as there are fewer mismatches between the value of the assets and the net present economic value associated with the liability.

---

[1] For example, one well-known asset manager has created a glide path that invests nearly 100% in TIPS at retirement, an extremely risk-averse strategy.

Exhibit 1 Improving Total Portfolio Health

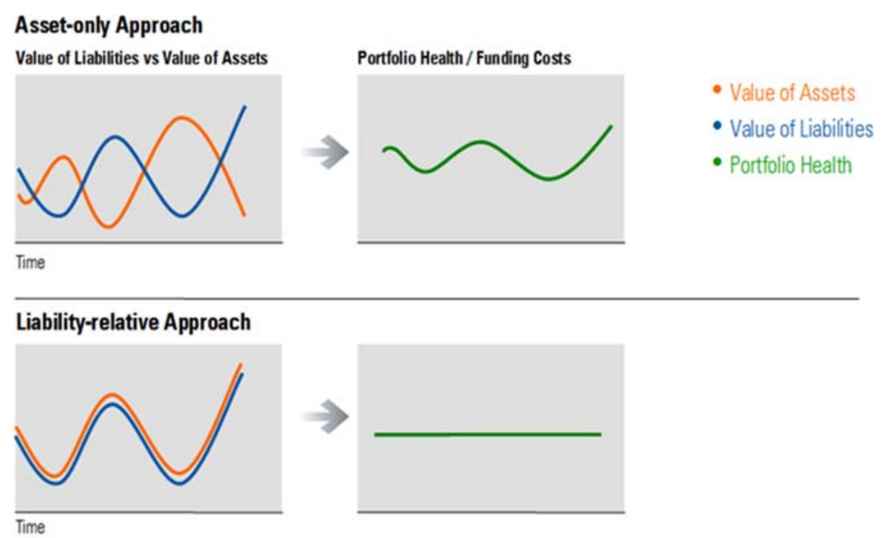

As noted previously, the three most common LDI frameworks are cash flow matching, duration matching, and liability-relative optimization. Of these, cash flow matching and duration matching are generally straightforward and tend to result in what most practitioners would consider ultra-conservative portfolios. In contrast, liability-relative optimization is a more sophisticated framework that builds on Harry Markowitz's mean-variance optimization "asset-only" framework to a total portfolio framework that simultaneously considers the assets, the liabilities, the interaction between assets and liabilities, as well as the size of the assets relative to the size of the liabilities (funding ratio). In contrast with the first two approaches that result in a single, ultra-conservative asset allocation, liability-relative optimization results in a continuum of optimal portfolios, referred to as the "liability-relative efficient frontier," where the most conservative portfolios are often similar to those that would be determined using either a cash flow-matching or duration-matching approach. Meanwhile, the more aggressive mixes on the right side of the liability-relative efficient frontier are similar if not identical to the aggressive mixes from a traditional asset-only frontier. Most importantly, liability-relative optimized portfolios represent a wide range of potential mixes, where the selected allocation is driven by the *risk tolerance* of the investor. Here, risk is appropriately defined as the inability to meet one's goal.

It is unclear exactly when the traditional asset-only Markowitz approach was first extended to include liabilities, but liability-relative optimization has been in use since the late 1970s. Creators and innovators of modern finance theory, such as William F. Sharpe, Jack Treynor, and Roger Ibbotson, advanced tools and techniques to create optimal strategic asset allocations that consider both assets and liabilities. Important works on liability-relative optimization include Siegel and Waring (2004), Waring (2004a, 2004b, 2004c), and Ziemba and Mulvey (1998). These papers focus on the application of liability-relative optimization in a defined-benefit pension setting in which there is a hard, legal liability. Works that explore the application of liability-relative optimization for developing asset allocations for individuals include Idzorek (2008) and Idzorek et al. (2016).

Liability-relative optimization should not be confused with what is often referred to as a traditional "asset-liability study." The historical approach to asset-liability modeling has been to carry out an asset-only, mean-variance optimization followed by a series of Monte Carlo simulations. In the Monte Carlo simulations, typically several mixes from the asset-only efficient frontier are examined with an emphasis on the future distribution of asset values and the cost of funding the liabilities. As Waring (2004c) explains, while these studies often are referred to as "asset-liability studies" the link between the liability and the asset allocation policy decision is tenuous to nonexistent. The portfolio policy decision needs to be considered in a broader, more relevant, and total portfolio context. The true risk for a portfolio that exists to fund a liability is the probability that the cash flows that could be created from the portfolio will be inadequate to pay for the liability. The asset-only or asset-centric approach ignores the natural hedge between the assets and the liabilities of the investor's total portfolio. Asset-only approaches ignore this natural hedge, while liability-relative techniques seek to exploit it for the benefit of the investor and the health of the total portfolio. To put a finer point on it, when the goal is some form of an inflation-adjusted consumption stream, cash or cash equivalents are no longer the "risk-free" asset.

In a liability-relative optimization, we must model the liability. For our purposes, this means finding an asset, or combination of assets, with extremely similar characteristics or sensitivities to market movements whose value would increase or decrease in line with the (net present) value of the expected cash flows that form the liability. The relatively steady stream of cash flows associated with most liabilities resembles the cash flows of a bond. Identifying a liability model enables us to understand the correlation between the change in value of the liability and that of the potential portfolio investments. This in turn helps us set an appropriate market-based discount rate for estimating the net present value of the expected cash flows from the asset portfolio.

A large defined-benefit plan with many participants, in which idiosyncratic events nearly nullify each other, can model its liability simply. In contrast, it's much harder for an individual retiree to model the retirement liability because individual cash flows are more complex and can be dominated by idiosyncratic events, such as unexpected (and unhedged) health expenses or the sudden need to pay for a wedding. Because individual retirement cash flows can be expected to be more volatile than those for a large pension fund, they need to be modeled using a higher discount rate.

Historically, the liability in a liability-relative optimization is modeled as long-term government bonds, since the cash flows associated with a nominal defined-benefit pension plan are similar to those of long-term government bonds. More specifically, those cash flows occur over many years (long duration), are steady, and often made on a nominal basis (not adjusted for inflation). With the introduction of inflation-linked bonds and because many liabilities are linked to inflation (e.g. a defined benefit pension with a cost-of-living adjustment), long duration, inflation-linked bonds are another common proxy for the liability. In addition to nominal vs. inflation-linked characteristics, in the case of a hard or legal liability in which the payments must be made with certainty, seemingly one would want to use the interest rate associated with Treasury bonds or perhaps high-quality corporate bonds that reflect the certainty of the payment. The low discount rate reflects that the payments are steady and they *must* be paid with a high degree of certainty. Additionally, when the consequences of not making the payment are significant one would not want to underestimate the net present value of the cash flow series by using an inappropriately high discount rate. For historical, legal, and accounting reasons, one may need to use a specific discount rate for the purposes of preparing financial statements/disclosures, but when actually setting asset allocation policy it's best to use the most appropriate market-based discount rate.

In the case of a soft liability in which there is more flexibility around payments, one might use the interest rate of a lower-quality bond or perhaps use a combination of high-quality bonds mixed with equities that better reflects the uncertainty and volatility of the soft liability. We will discuss this further, but one might choose different discount rates to reflect different types of expenses—an ultra-safe rate for essential expenses and higher rate for discretionary expenses.

Exhibit 2 depicts a traditional asset-only efficient frontier and liability-relative efficient frontier drawn in traditional asset-only risk and return space. This is the unfortunate standard lens or perspective in which the sole definition of "good" is the return on the asset-only portfolio and "bad" is the variability of that asset portfolio. Notice that in this stylized representation we assume that the investor's liability is model as Treasury inflation-protected securities (TIPS) held short; thus, in the asset-only space the liability is drawn as the mirror image of TIPS (the standard deviation is the same, but the return is now the opposite). For an investor with a liability, the traditional asset-only risk and return space is not

Exhibit 2 Efficient Frontiers in Asset-Only Risk and Return Space

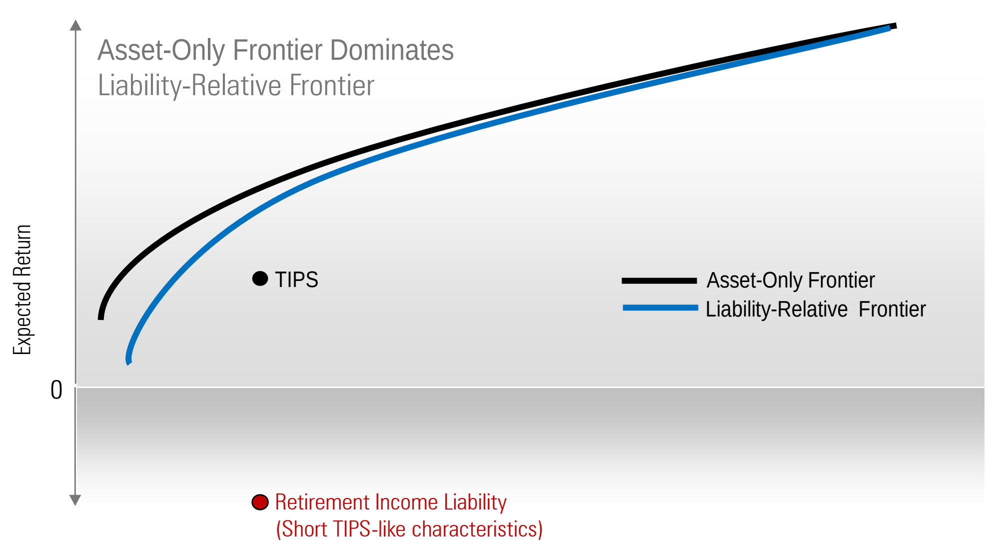

appropriate and has the unintended consequence of creating a myopic view that fails to account for the investor's total portfolio—a portfolio that consists of both assets and liabilities.

In contrast, Exhibit 3 illustrates the more holistic liability-relative risk and return space. The efficient asset allocation mixes that form the two *asset-only* frontiers in Exhibit 2 and Exhibit 3 are identical. Likewise, the efficient asset allocation mixes that form the two *liability-relative* frontiers in Exhibit 2 and Exhibit 3 are identical. What has changed is our perspective or lens with which we are viewing success. The difference in the two lenses or spaces is whether we think about risk and return of a *subset* of the portfolio or the *total* portfolio. In the more applicable and appropriate total portfolio space, the liability-relative frontier now dominates the asset-only frontier. Notice that, in this example, the minimum risk portfolio (analogous to the traditional minimum variance portfolio) has zero expected return and no systematic risk. This is because the minimum risk portfolio is invested 100% in TIPS. Since in this example we assumed that the liability is represented as 100% TIPS held short, we have two positions that perfectly offset one another—an asset portfolio of TIPS and a liability portfolio of TIPS held short. This risk cancelling is what we mean when we said that liability-relative investing seeks to take advantage of the natural hedge between the systematic factors that influence the value of the liability and those that influence the value of the asset portfolio.

Exhibit 3 Efficient Frontiers in Liability-Relative Risk and Return Space

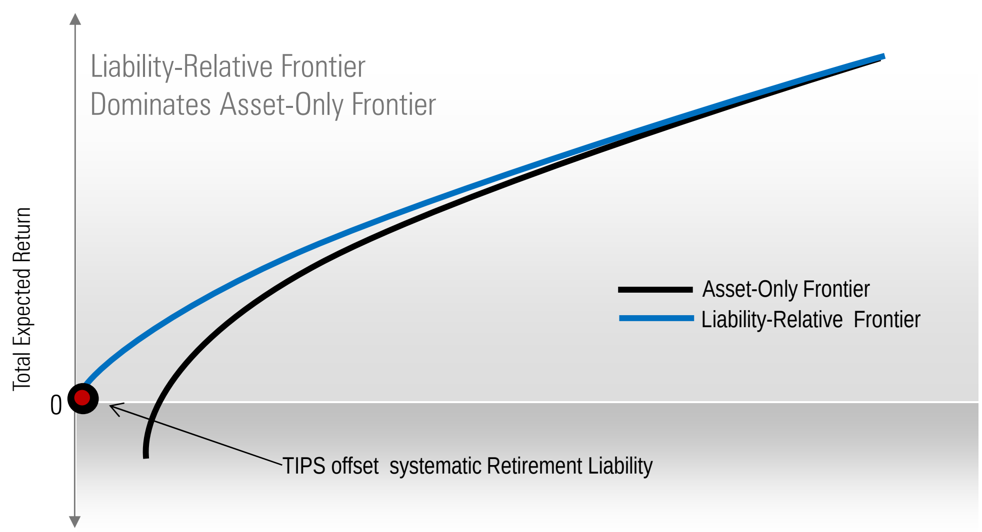

In Exhibit 3, at the far left of the two frontiers there are significant differences in the asset allocation mixes that lead to the different risk and return plot points that form the frontiers. More specifically in this stylized example one could assume the minimum risk asset-only portfolio is invested nearly 100% in cash while the minimum risk liability-relative portfolio is invested 100% in TIPS.[2] Moving from left to right along the frontiers they gradually converge to nearly equivalent asset allocations (and in the absence of any constraints beyond non-negativity at the far right, in both cases 100% is invested in the single asset class with the highest expected return). In other words, for investors with low risk tolerance, liability-relative optimization leads to significant asset allocation differences, while for those with high risk tolerance the asset allocations can be very similar. It is for this reason that you may hear people say that liability-relative optimization is especially pertinent for retirees (which tend to invest more conservatively), while an asset-only approach is sufficient for more aggressive investors. While cash is often thought of as a "safe" investment in an asset-only context, if the liability is linked to inflation and inflation increases dramatically, a portfolio with a significant cash allocation would not be able to pay for the liability while a TIPS-based portfolio would keep up with inflation.

Differences in Efficient Portfolios

In this section, we analyze the asset allocation differences from a series of optimizations that are either based on a traditional asset-only approach or the liability-relative approach. We carry out several liability-relative optimizations, experimenting with different liability models in which the optimizer is constrained to hold a negative (or short position) in one or more asset classes representing the systematic characteristics of the liability. For the liability-relative optimizations, we consider three distinct potential liability models or liability proxies: long-duration inflation-linked bonds (e.g., long TIPS), long-duration nominal bonds, or inflation. At this point, we are not trying to determine the actual nature of the risks associated with a retirement liability (this will be explored in following sections), rather to provide some perspective as to how the optimal portfolio changes in the presence a liability, and to further investigate asset allocation changes when the characteristics of the assumed liability change.

An implicit assumption for those using mean-variance optimization (MVO) is that returns of the assets that form the opportunity set are normally distributed, or follow a "bell curve." Xiong and Idzorek (2011) note that empirically we have seen that most asset classes and portfolios have historical returns that were not normally distributed. Furthermore, investor preferences go beyond mean and variance. For our analysis, we use an enhanced optimization approach focused on minimizing the extreme "left-tail risk" or extreme losses of a return distribution, focusing on the conditional value-at-risk (CVaR). CVaR is closely related to VaR (value-at-risk) and is calculated by taking a probability-weighted average of the possible losses conditional on the loss being equal to or exceeding the specified VaR. Conceptually, this is in the spirit of trying to avoid a permanent loss of capital or, in a liability-relative context, a catastrophic event that would permanently impair the asset portfolio's ability to pay for the liability.

Reflecting more recent advances in the use of higher moment optimization techniques, we used a liability-relative mean-CVaR optimizer in which we model the non-normal qualities of the different asset

---

[2] In a typical asset-only optimization in which cash is included as a risky asset and thus part of the opportunity set, the global minimum variance portfolio is typically 95% or more in cash with very small allocations to a handful of additional diversifying asset classes (see for example the left-hand side of Exhibit 5 below).

classes. There is not a closed-form solution for estimating CVaR. Therefore, the liability-relative mean-CVaR optimization procedure is a form of simulation-based optimization. Based on an expanded set of capital market assumptions that includes forecasts for the mean, standard deviation, skewness, and kurtosis for each option in the opportunity set as well as the correlation matrix, returns are simulated under the assumption that the returns follow a multivariate truncated Lévy flight distribution described by the forecasts. The optimizations are based on 100 sample periods, each consisting of 100 years, where the optimal portfolio allocations are averaged across the optimizations. Additional information about our optimization routine is included in Appendix 1.

When performing the optimizations, we use a relatively large, inclusive opportunity set of asset classes, as presented in Exhibit 4, which includes asset classes representing different durations, credit quality, inflation-hedging characteristics, and different currency exposures. The asset classes and the CMAs associated with them are contained in Appendix 2. While inflation is included in the asset class opportunity set list, it is not considered as an investible asset class, and rather is used as the liability in one of the surplus optimizations.

Exhibit 4 Asset Class Opportunity Set for Optimizations

Asset Class	Proxy
U.S. Large	Morningstar US Large Cap
U.S. SMID	Morningstar US Small-Mid Cap
Developed Markets ex-US	Morningstar DM exUS
Emerging Markets	Morningstar EM
Long-Only Commodity	Morningstar Long-Only Commodity
US Real Estate	Morningstar US REIT
Non-US Real Estate	FTSE EPRA/NAREIT Developed Ex US
US Long-Term Bonds	Morningstar US Short Gov/Corp Bond
US Short-Term Bonds	Morningstar US Long Gov/Corp Bond
US Long-Term TIPS	Bloomberg Barclays US Treasury Infl
US Short TIPS	Bloomberg Barclays U.S. Treasury
Global High Yield	Bloomberg Barclays Global High Yield
Non-US Investment Grade Bonds	Barclays Global Agg Ex USD
Cash	Morningstar Cash
Inflation	IA SBBI US Inflation

Finally, rather than focusing on the risk and return plot points associated with the efficient frontiers, we focus on the asset allocations that form the different efficient frontiers across different risk levels. More specifically we focus on what are called efficient frontier asset allocation area graphs. On an efficient frontier asset allocation area graph, each vertical cross section is an efficient asset allocation mix from the frontier, and moving from left to right up the risk spectrum the vertical cross sections representing different asset allocations evolve to show the range of efficient asset allocations across the entire efficient frontier.

The first efficient asset allocation graph, Exhibit 5, shows the asset allocations that make up a base-case, asset-only traditional efficient frontier. This serves as a baseline for comparison for all the subsequent liability-relative efficient asset allocation area graphs.

Exhibit 5 Asset-Only Efficient Frontier Asset Allocation

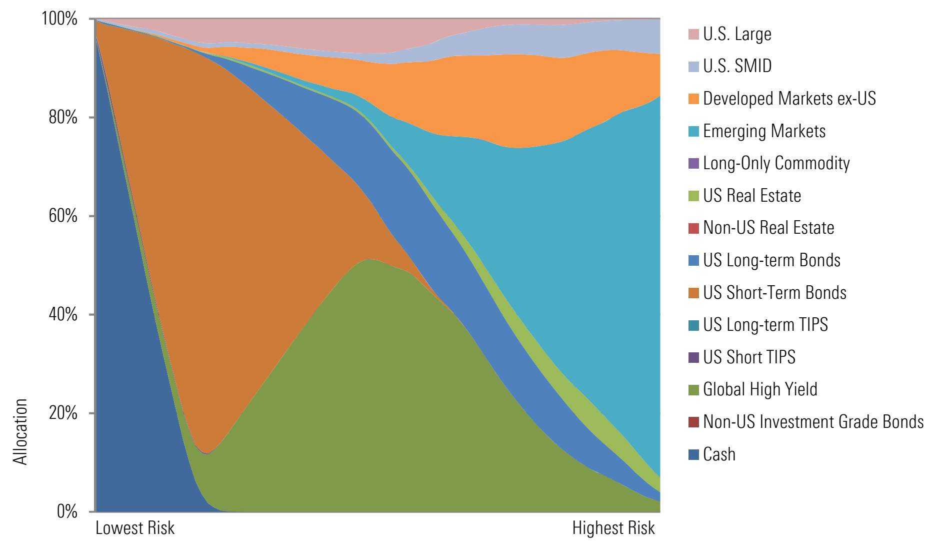

Exhibits 6, 7, and 8 include the asset allocations that make up the liability-relative efficient frontiers in which the liability is modelled as long-duration TIPS, long-duration nominal bonds, and inflation, respectively.

Exhibit 6 Long-Duration TIPS Liability-Relative Efficient Frontier Asset Allocation

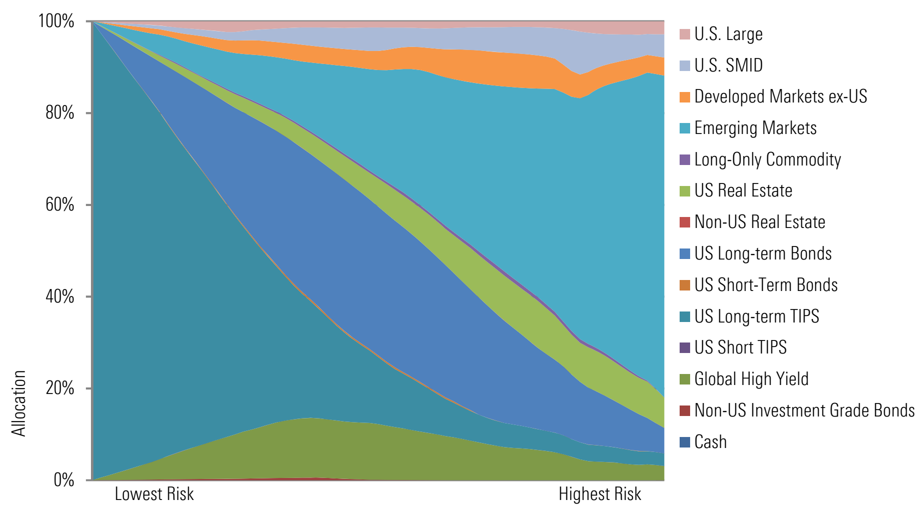

Exhibit 7 Long Nominal Bonds Liability-Relative Efficient Frontier Asset Allocation

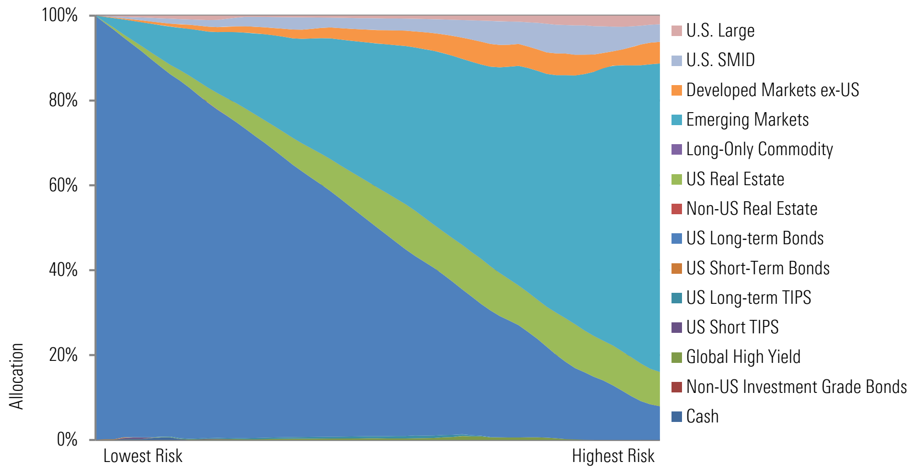

Exhibit 8 Inflation Liability-Relative Efficient Frontier Asset Allocation

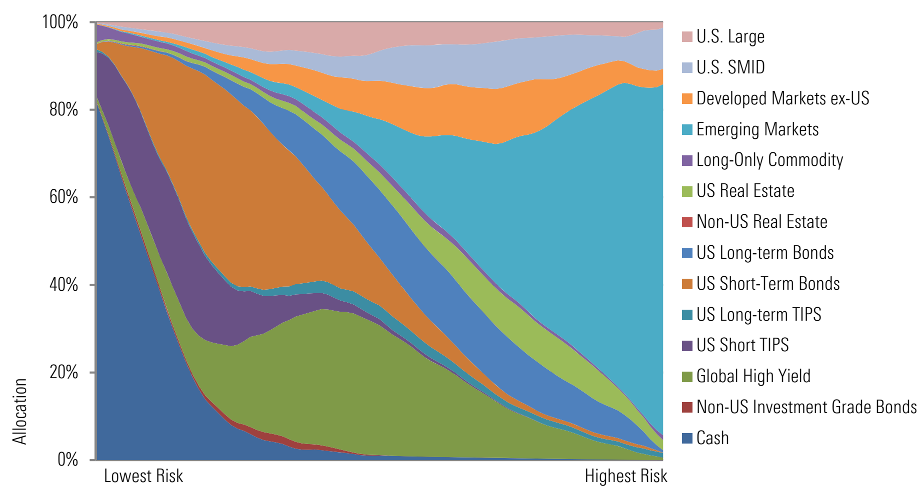

To aid with the comparison across the four asset allocation area graphs, in Exhibit 9 we identified three specific mixes from each of the four area graphs: a Conservative mix (10% equities and 90% fixed income), a Moderate (40% equities and 60% fixed income), and an Aggressive mix (70% equities and 30% fixed income). In the lower panel of Exhibit 9, we calculate various ratios and subtotals to highlight some of the important differences between what one might consider comparable asset allocations given that they have equivalent equity and fixed-income subtotals.

Exhibit 9 Optimization Allocations [%] for Various Portfolios and Liabilities

Asset Class	Conservative 10/90				Moderate 40/60				Aggressive 70/30			
	Asset-Only	Long Bond	Long TIPS	Infl	Asset-Only	Long Bond	Long TIPS	Infl	Asset-Only	Long Bond	Long TIPS	Infl
US Large	5.0	0.5	1.1	4.4	4.2	0.5	1.2	6.0	1.2	1.3	0.7	4.3
US SMID	0.9	0.8	1.1	1.7	3.9	2.5	3.7	7.7	6.2	5.3	14.3	10.8
Developed Mkts ex-US	3.1	0.5	1.4	1.4	15.5	2.4	3.8	9.3	18.6	5.3	6.4	12.6
Emerging Markets	0.5	6.3	4.8	1.0	14.6	27.3	24.7	11.7	39.1	47.7	41.3	33.2
Long-Only Commodity	0.0	0.0	0.2	1.0	0.0	0.0	0.6	1.9	0.0	0.0	1.1	1.0
US Real Estate	0.4	1.8	1.5	1.0	1.8	7.2	6.4	4.4	4.9	10.4	6.8	8.7
Non-US Real Estate	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
US Long-Term Bonds	4.5	89.5	13.1	1.9	16.8	59.3	33.2	15.7	12.0	29.1	18.3	13.3
US Short-Term Bonds	62.4	0.1	0.2	41.0	0.5	0.0	0.3	9.8	0.0	0.0	-0.1	1.9
US Long-Term TIPS	0.0	0.1	70.8	0.7	0.0	0.4	14.2	3.1	0.0	0.1	4.3	1.4
US Short-Term TIPS	0.0	0.0	0.0	18.7	0.0	0.0	0.0	1.1	0.0	0.0	0.0	0.2
Global High Yield	23.1	0.1	5.6	12.9	42.6	0.3	11.9	28.2	18.0	0.9	7.0	12.2
Non-US Invst-Grade Bonds	0.0	0.0	0.2	0.9	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0
Cash	0.0	0.1	0.0	13.4	0.0	0.0	0.0	1.0	0.0	0.0	0.0	0.5
Equity	10.0	10.0	10.0	10.0	40.0	40.0	40.0	40.0	70.0	70.0	70.0	70.0
Fixed Income	90.0	90.0	90.0	90.0	60.0	60.0	60.0	60.0	30.0	30.0	30.0	30.0
Infl-Linked/Total Fixed Income	0.0	0.1	78.7	21.7	0.0	0.7	23.8	7.2	0.0	0.2	14.7	5.3
Long Duration/Total Fixed Inc	5.1	99.6	93.3	2.9	28.0	99.5	79.4	32.0	40.0	97.4	76.8	49.9
Global HY/Total Fixed Income	25.6	0.1	6.2	14.4	71.1	0.5	19.9	47.8	60.0	2.9	23.6	41.3
Non-US Bonds/Total Fixed Inc	10.2	0.1	2.7	6.8	28.4	0.2	8.2	19.2	24.0	1.1	9.5	16.5
Non-US Equity/Total Equity	36.7	69.1	61.5	23.0	75.3	74.4	70.7	51.3	82.4	75.6	67.6	64.8
TIPS+RE+Commodities	0.4	1.9	72.4	21.5	1.8	7.6	21.2	10.5	4.9	10.5	12.2	11.3

As a quick disclaimer, we never advocate that practitioners blindly accept and use asset allocations created from an optimizer; rather, the value from optimizations is in high-level directional comparisons.

For all three splits for comparison, the asset-only allocations are significantly different from the liability-relative optimizations. In general, the asset allocation differences are greater among the fixed-income asset classes.

Notice that the ratio of inflation-linked bonds to total fixed income was greatest for the two liability-relative optimizations in which the liability is either long-duration TIPS or inflation. Next, the ratio of long-duration bonds (both real and nominal) to total fixed income was greatest for the two liability-relative optimizations in which the liability was either long-duration nominal bonds or long-duration TIPS. The ratio of non-U.S. fixed income to total fixed income was greatest for the asset-only optimization, suggesting that liability-relative optimization leads to a more significant home-bias than asset-only optimization.

Focusing on equity allocations, the asset-only optimization results in the largest allocation to non-U.S. equities. Including a dollar-denominated liability consistently reduced the amount allocated to non-U.S. equities, although this home-bias tilt was less prevalent among equities than it was for fixed income.

Finally, the total allocation to real return assets of TIPS, commodities, and real estate, was highest when the liability was long-duration TIPS followed by inflation.

Overall, these optimizations demonstrate primarily that relative to asset-only optimization, liability-relative optimization can lead to significant differences in asset allocation. Additionally, how one chooses to model the liability can also lead to significant differences in asset allocation.

Modeling Retirement Income Liabilities

We have demonstrated that a) moving from asset-only to a liability-relative optimization framework results in significant asset allocation differences, and b) the way one models the liability also leads to significant asset allocation differences. In this section, we attempt to build an asset allocation with the highest likelihood of paying for the *expected* retirement expenses represented by a series of outgoing cash flows. To do this we must attempt to find mark-to-market securities or asset classes that reflect or capture the systematic factors affecting the liability. Admittedly, some portion of retirement expenses should be thought of as idiosyncratic and as such cannot be hedged, but that should not preclude us from attempting to take advantage of the natural hedge that exists between asset classes and the systematic factors that can impact expenses.

A common baseline assumption regarding the retirement income liability (i.e., the expected cash flows in retirement) is that retirees desire to maintain the same, or nearly the same, standard of living throughout retirement. Borrowing from standard economics, the goal is a lifetime of relatively smooth consumption. However, we may view this as a soft liability because, over a longer time horizon, an individual retiree's expenses can be decreased or restructured if necessary. Over the short-term and in certain circumstances, there may be elements of expenses that are difficult to alter, but most retirees have at least some flexibility with respect to spending. Therefore, for the typical retiree, the retirement income liability can be deconstructed into risk-free and risky portions.

The risk-free portion is the "hard" or nondiscretionary aspect of retiree spending. This should be proxied through highly certain bonds with no or low credit risk, such as Treasuries, where the duration would vary based on the expected retirement duration and type. Additionally, one should select either nominal or inflation-linked Treasuries depending on whether the liability is linked to inflation. While a shockingly low number of people purchase income annuities, arguably an inflation-linked income annuity would be an excellent vehicle for defeasing the nondiscretionary portion of spending that simultaneously protects against longevity risk.

The risky portion is the "soft" or discretionary aspect of retiree spending and at some level should attempt to be linked to the types of goods and services in the economy that the retiree will purchase. For example, general equities represent the economy, but this could be refined to potentially overweight real estate, the health care sector, and possibly consumer staples to more accurately reflect the consumption basket of a given retiree. We will provide some perspectives on modeling considerations for these points in the following subsections.

Retiree Consumption

The extent to which retiree consumption will be discretionary or not will vary materially by household—what might be deemed nondiscretionary to one retiree, might be deemed discretionary by another, and vice versa. To better understand the nature of retiree expenditures we use data from the Bureau of Labor Statistic's Consumer Expenditure Survey. We only include households that have no reported wage income (indicating that they are likely retired) and respondents in which the age of the head of household is between 60 and 80 years old.

Exhibit 10 includes information about average household expenditures for four different head-of-household age groups (60-65, 66-70, 71-75, and 76-80) in which there are seven expenditure types (housing, food, transportation, healthcare, entertainment, charity, and other). The majority of expenditures in Exhibit 10 would be considered relatively nondiscretionary (hard) in nature. For example, housing, food, transportation, and healthcare represent approximately 80% of total expenditures. In reality, though, there is likely some degree of flexibility even within these expenditure categories (e.g., food). One notable change in expenditures during retirement is the rising portion of monies allocated to healthcare.

Exhibit 10 Household Expenditures by Age Group

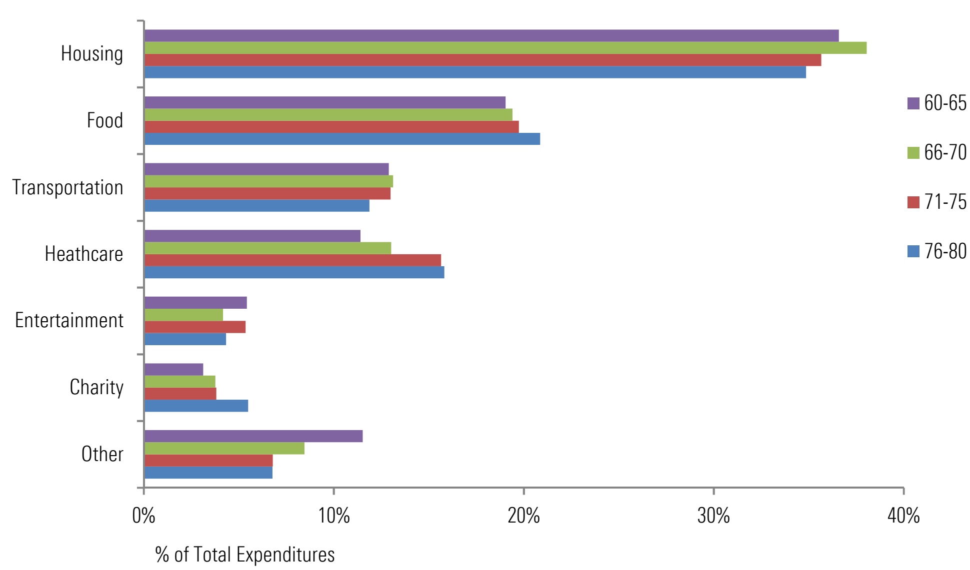

Moving to Exhibit 11, the different colored bars now represent groups based on total expenditures rather than by age (in Exhibit 10). Households that spend less than $10,000 tend to allocate more of their spending to nondiscretionary items. For example, housing, food, transportation, and healthcare represent approximately 70% of total expenditures for households with total expenditures exceeding $50,000 but 85% of total expenditures for households with total expenditures less than $10,000.

Exhibit 11 Household Expenditures by Total Expenditures

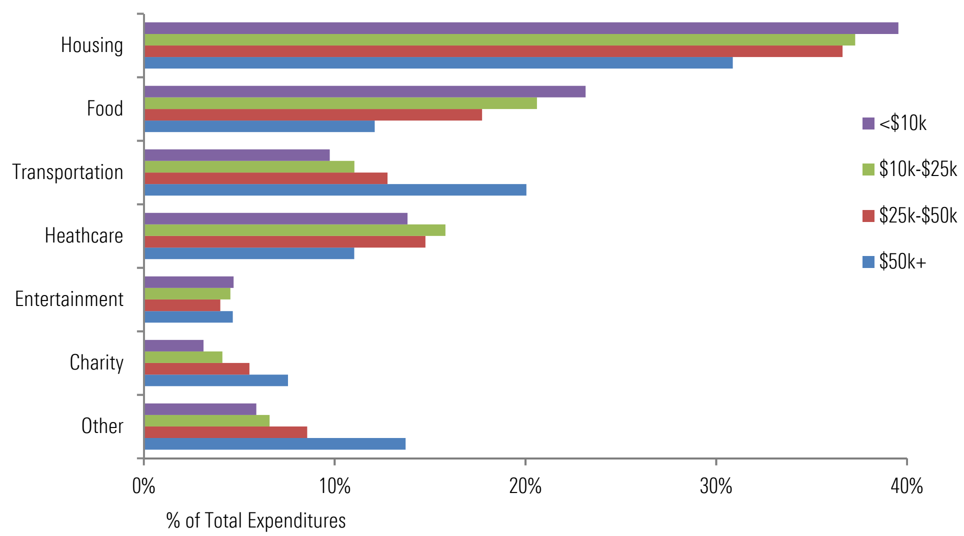

.

Overall, as a rough guide this analysis suggests that approximately 80% of retiree expenditures would be considered nondiscretionary, or "hard" in nature. This would suggest that within a multi-asset liability model, approximately 80% of the liability should be represented by relatively safe assets. Given the significant allocation of expenditures to housing, some type of real estate-focused investment such as REITs would likely be a relevant proxy, as would some general exposure to the U.S. economy, which we assume to be U.S. equities.

Retirement Duration

Another important aspect of the retirement income stream is duration—how long the stream needs to last. In Exhibit 12, the blue line is the conditional duration of a mortality-weighted income stream for life plotted for an investor of a given age and assuming that retirement starts at age 65. This is known as the Macauley duration and represents the weighted average time in years until cash flows occur. Given that most investors should plan to live longer than the life expectancy for the general population, the upper (red) line plots the conditional duration of a mortality-weighted income stream for life assuming a planning horizon that is five years longer than life expectancy. At the typically assumed retirement age of 65, assuming a planning horizon of five years past life expectancy, the duration of the cash flows is

13.3 years. (Later we will discuss the purchase of a single premium immediate annuity at age 65.) And thus, we plot the duration of this assumed purchase, as simply 65 minus the investor's current age.

Exhibit 12 Duration of Retirement Expenses by Current Age and Associated Life Expectancy

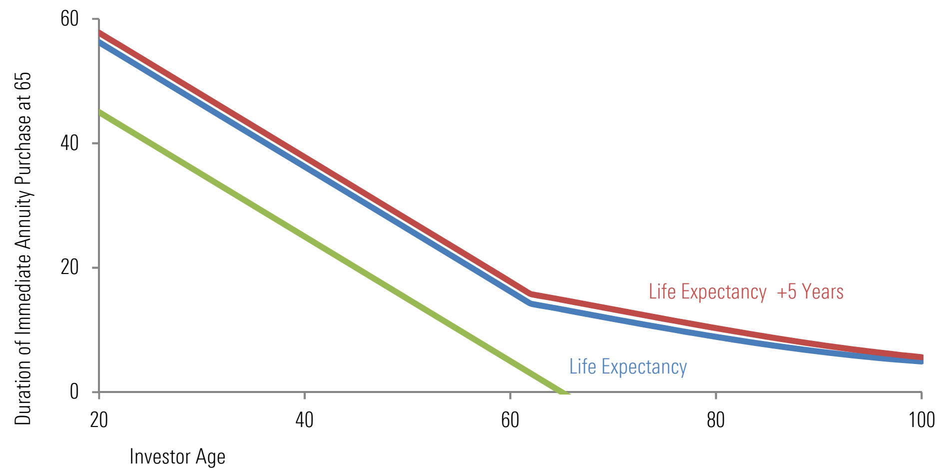

If one accepts that an appropriate liability model should attempt to match the duration of liability (which may not be possible for younger investors given that life expectancy significantly exceeds the duration of available investments), it is important to understand the duration of the asset class or asset classes that

Exhibit 13 Historical Durations for Short-Term and Long-Term Bond Indexes

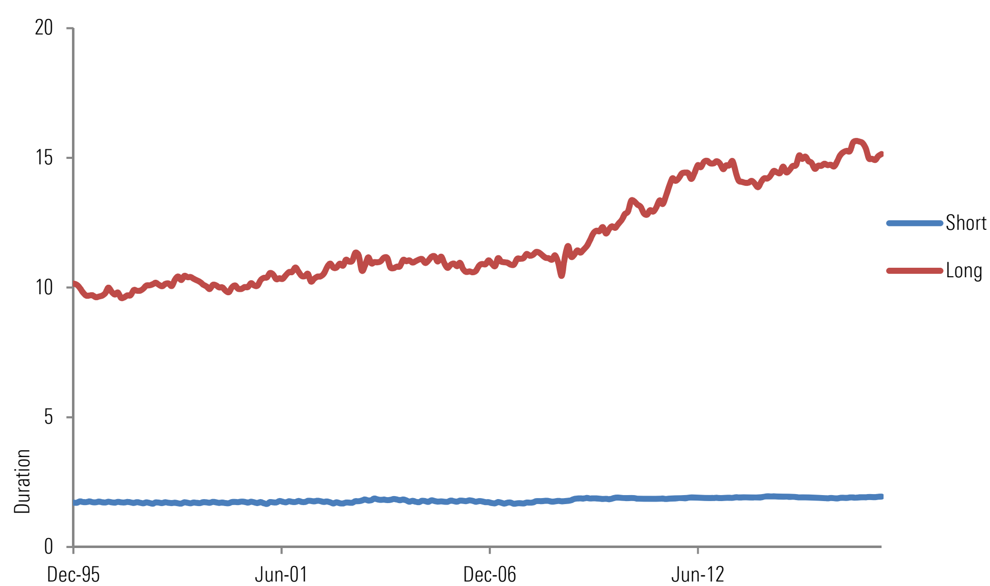

one might use to model the liability. Exhibit 13 plots the historical duration of some of short and long bond indexes since 1995

We see that the duration for short-term bonds has been relatively constant at about two years, while the duration of the long bonds has varied somewhat, with a long-term average of 11 years, and more recently increasing to 15 years as interest rates have declined.

Estimating durations for non-fixed-income asset classes, such as equities, is less straightforward. While it's possible to estimate the implied duration using a dividend discount valuation, the notion of equities having a measurable cash flow duration is a bit controversial (see Blitzer, Dash, and Murphy, 2009, and Agarwal and Liu, 2016). Empirical correlations between realized equity returns and interest rate movements have been quite volatile and even negative at times (see Loomis Sayles, 2016). Therefore, while the true duration of equities is likely relatively long (20-plus years) we ignore these assets classes from duration considerations for optimization purposes.

Overall, Exhibit 12 suggests that durations of retirement cash flows are quite long and, as a result, when creating an asset-based model to represent that liability we should have an evolving model with an equivalent duration.

Nominal vs. Real Retirement Income Need

The retiree income need is often assumed to increase with inflation throughout retirement. To provide some perspective on the actual changes (and volatility) in household spending, we analyzed data from the Health and Retirement Study (HRS), a public use dataset produced and distributed by the University of Michigan, which is then further refined by the RAND Corporation in what is known as the "RAND HRS Data Files." These data files contain panel data that can be used to track the spending of the same households through time.

We reviewed data on *total consumption* for the previous six waves of data releases (2001, 2003, 2005, 2007, 2009, and 2011) as well as *household nondurable spending* and *total spending* for the previous seven waves of data releases (2001, 2003, 2005, 2007, 2009, 2011, and 2013).[3] We limited the analysis to households that considered themselves to be fully retired for the entire test period, with the same household "type," and data for all seven waves. This left us with data on 288 households.

We studied changes in spending for each household between the data releases. For example, between 2001 and 2003 for the 288 households, there were 288 *changes* in real total spending. These 288 changes form a distribution of changes from large negative changes to large positive changes. The distributions of changes were similar across the five waves and are displayed in Exhibit 14. The changes are normally distributed (approximately) with a median change in real spending near zero. The changes associated with the 25th and 75th percentiles were bounded by approximately -11% and +6%,

[3] Total consumption data was not available in the 2013 wave.

respectively, suggesting that 50% of the households experienced relatively small changes between waves. However, the other 50% of the 288 households experienced larger changes in spending, with the 5th and 95th percentiles indicating large changes in wave-over-wave spending. Focusing on the 5th and 95th percentiles and the five distributions in Exhibit 14, for approximately 90% of the households the wave-over-wave change was less than plus or minus 30%.

Exhibit 14 Distribution of Annual Real Changes in Total Household Consumption Across HRS Waves

The median annual real change in total consumption was slightly negative. This is consistent with research by Blanchett (2014), among others, who have observed that retiree spending does not necessarily increase with inflation throughout retirement. While there is certainly a real component to one's preconceived consumption pattern, to suggest that retirement spending must be linked to inflation is not consistent with the empirical evidence on retiree spending. It is unclear how to reconcile an aspirational inflation-adjusted spending pattern with what we observe empirically. In the absence of a clear view, from a liability modelling perspective it would seem to make since to include both inflation-linked bonds as well as nominal bonds.

In Appendix 3, we replicate the analysis of Exhibit 14; however rather than focusing on total consumption, we focus on household nondurable spending and total spending.

Regardless of household spending definition (i.e., total consumption, nondurable spending, or total spending) the average real change in consumption is generally slightly negative. For a subset of households, there are significant variations in spending over time. This provides additional empirical evidence that retiree spending is not entirely fixed (i.e., bond-like) and should not be entirely linked to inflation nor inflation-linked bonds.
With respect to the decision about how to model the fixed portion of the liability using nominal or inflation-linked bonds, actual retiree consumption (or nondurable spending or total spending) has

decreased by approximately 2% per year, on average. While it's not clear to what extent this is by choice or necessity (i.e., are retirees spending less because they choose to or because they have to), it's reasonable to assume that approximately half of the liability is nominal and half of the liability is linked to inflation. The volatility in observed consumption also suggests, similar to the expenditure analysis, that some of the cash flows associated with the retirement liability are in fact relatively volatility and thus should be modeled as a risky asset.

Putting It Together: A Better Picture of the Retiree Liability

Combining the information from the previous subsections suggests that several ingredients need to be combined to mimic or model the systematic factors associated with a typical retiree's cash flows.

- Discretionary Versus Nondiscretionary Spending: Collectively housing, food, transportation, and healthcare (items that we largely think of as nondiscretionary) represent approximately 80% of expenses. This suggests that 80% of the liability model should consist of safe assets.
- Housing: The single largest expenditure is housing, which nearly doubles each of the other expense categories. This suggests that housing/real estate should be included in the liability model.
- Duration: The duration of the expected cash flow series changes with age. The liability model, especially the fixed-income portion of that model, should reflect the changing duration of the cash flow series.
- Nominal Versus Inflation-Adjusted (Real): While we think it is natural to aspire to—and prudent to plan for—an inflation-adjusted income stream, empirically expenditures for retirees have decreased, suggesting that nominal bonds adequately capture expenses. When modeling the liability, it seems reasonable to split the difference.

Based on these observations, Exhibit 15 shows how a potential liability model might evolve over time to match the evolving duration and characteristics of the liability between ages 65 and 85. It consists of 80% fixed income and 20% equities. The 20% in equities is split evenly between U.S. REITs and U.S. equities. On the fixed-income side, there is an even split between nominal vs. inflation-adjusted bonds, and the duration evolves to match the duration of the liability.

Exhibit 15 The Dynamic Retirement Liability

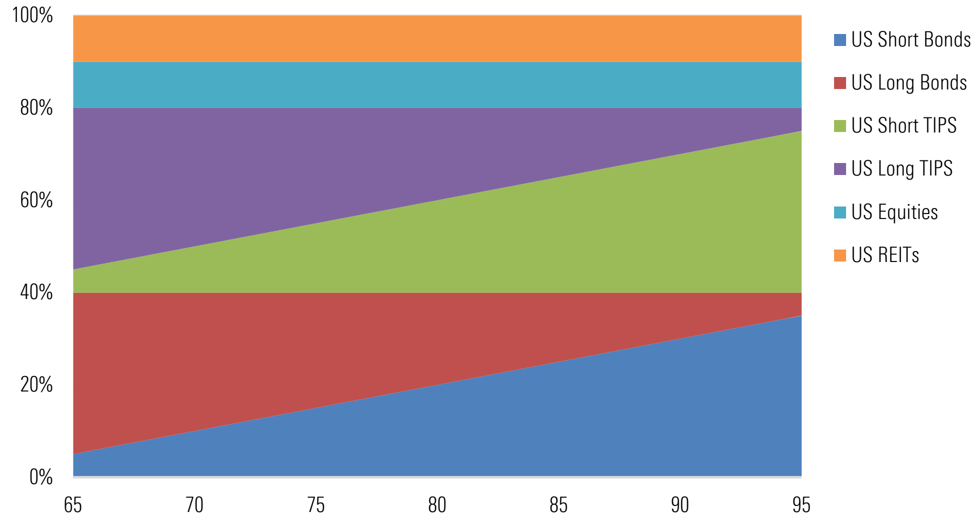

In Exhibit 15, the vertical cross section at age 65 identifies a potential multi-asset liability model for a 65-year-old that consists of 5% U.S. short-term government bonds, 35% U.S. long-term government bonds, 5% U.S. short-term TIPS, 35% U.S. long-term TIPS, 10% U.S. large cap, and 10% U.S. real estate. Based on the enhanced liability model in Exhibit 15 (the far left side corresponding to the 65-year-old), which is arguably a more accurate representation of the systematic factors that influence a retiree's cash flow series, we ran another liability-relative optimization. Exhibit 16 displays the efficient asset allocations from this optimization.

Exhibit 16 Efficient Frontier Asset Allocation Relative to Blended Liability for a 65-Year-Old Retiree

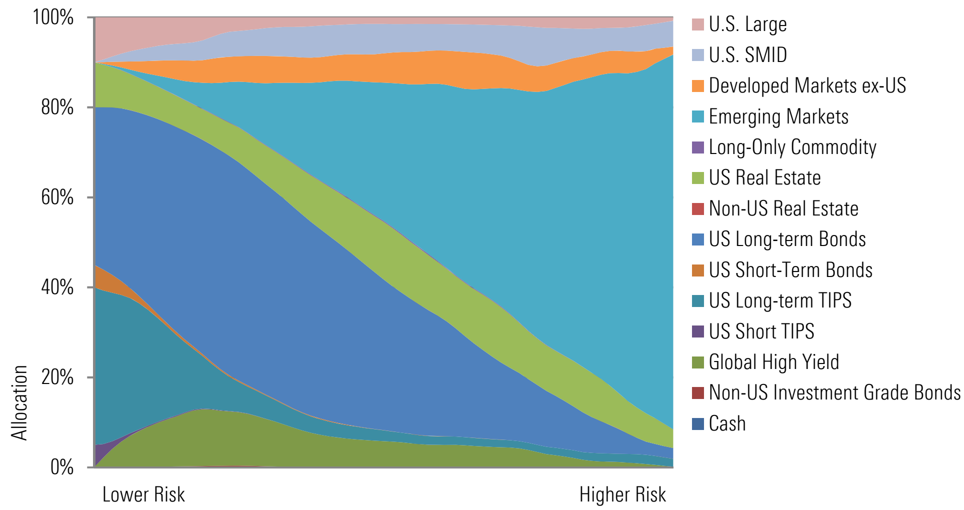

Ignoring the different color schemes, notice that the vertical cross section at the far left of Exhibit 16 perfectly matches the vertical cross section at the far left of the liability model from Exhibit 15. A very risk-averse investor would select this mix. As one's risk tolerance increases, the investor would move to mixes with higher potential return and higher levels of risk relative to the minimum risk asset allocation.

A Total Wealth Perspective

A central tenet of modern portfolio theory is to look beyond individual assets (or liabilities) in isolation and to consider them in a portfolio context. Taking this one more step, the individual's true portfolio and true funding status isn't just the value of their traditional retirement savings relative to their target retirement income. One should consider a holistic view of the retiree's total wealth. For example, approximately nine of 10 American retirees receive Social Security retirement benefits. The existence of Social Security retirement benefits impacts the definition of a truly efficient portfolio.

From a risk perspective, Social Security retirement benefits are similar to long-term TIPS since they provide guaranteed income for life indexed to inflation. Therefore, the optimal portfolio for a retiree with Social Security retirement benefits should be different than in ignorance of those benefits. We demonstrate the impact of considering Social Security retirement benefits on the optimal portfolio through another set of optimizations, the results of which are included in Exhibit 17, where we use the same liability used to create the area graph in Exhibit 16 (5% U.S. short-term government bonds, 35% U.S. long-term government bonds, 5% U.S. short-term TIPS, 35% U.S. long-term TIPS, 10% U.S. large cap, and 10% U.S. real estate) but assume the household has a 40% fixed allocation to long TIPS, representing the Social Security retirement benefits.[4]

[4] According to the Social Security Administration, Social Security retirement benefits represent 34% of the income for elderly Americans. Source: https://www.ssa.gov/news/press/factsheets/basicfact-alt.pdf

Exhibit 17 Efficient Frontier Asset Allocation Relative to Blended Liability for a 65-Year-Old With 40% Social Security

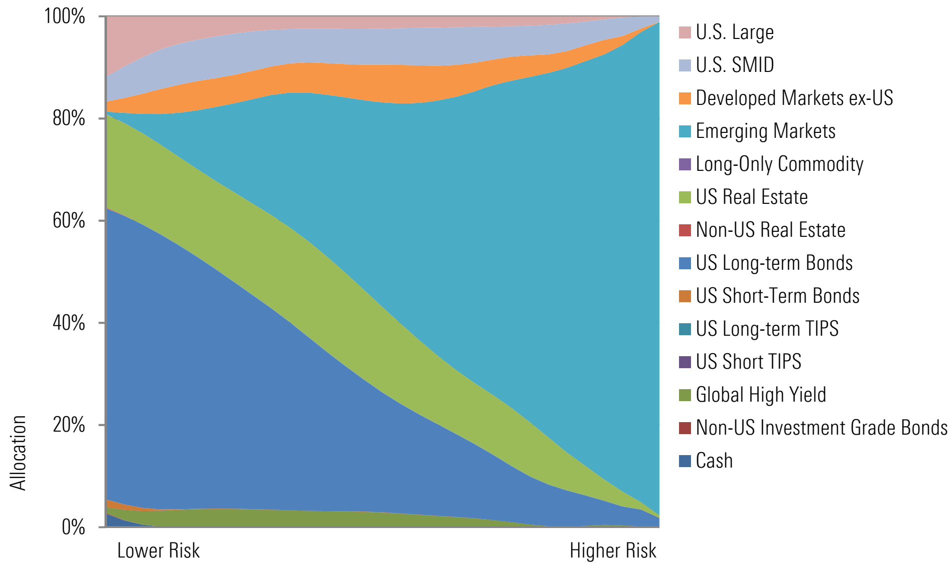

Security retirement benefits are incorporated into the portfolio optimization the allocation TIPS (both short and long) changes to *zero* for even the most conservative portfolios. Therefore, while some academics have contended that TIPS are the optimal hedge against the retirement liability, we find that the true benefit of TIPS is likely significantly less when viewed from a more holistic perspective.

Buying an Annuity at Retirement

While the most popular income strategy among retirees is to slowly sell off a portion of a portfolio, another approach is to purchase a single premium immediate payout annuity upon retirement. With an immediate payout annuity, a lump sum is irrevocably turned over to the annuity provider in exchange for an *insurance contract* that promises to provide the annuity purchaser (the annuitant) income for life. Similar to the cash flows from Social Security, an immediate payout annuity hedges against longevity risk (the risk of outliving one's assets). The amount of income for life that can be purchased with $1 is the annuity payout rate. In addition to the age of the investor at the time of purchase, the annuity payout rate depends heavily on prevailing interest rates.

Preretirement investors that plan to purchase an annuity at retirement face a very different set of circumstances than investors who plan to draw down a portfolio to generate retirement income. More specifically, the annuity is typically purchased at the beginning of retirement, dramatically shortening the investment time horizon and subjecting the individual to the risk of significant changes in payout rates prior to the eventual annuity purchase. Returning to our risk capacity discussion, the shortening of the investor's time horizon reduces the investor's ability to tolerate a significant decrease in the value of the asset that will be used to purchase the annuity. Conversely, once an annuity is purchased, the income floor provided by the annuity coupled with Social Security, provide the investor with risk-taking capacity for any remaining assets.

Looking back to Exhibit 12 in which we plotted the duration of the cash flow of an immediate annuity purchase at age 65, at first blush we see that the duration of that cash flow at age 65 is down to 0. This might be one of the reasons glide paths that stop at an assumed retirement age of 65 and are often ultra-conservative at the point of retirement. This rationale fails to consider how immediate annuities are priced and how payout rates are determined.

From a liability-relative perspective, the purchase of an annuity typically has a known date, such as retirement age. While the specific plan to purchase an annuity remains a soft liability, in this situation the purchase itself should be treated much more like hard liability. The desire to purchase an annuity at retirement results in a slightly different set of concerns than drawing down a portfolio of investments to generate retirement income. The annuity purchase is a *point-in-time* liability while the retirement income goal is more continuous.

Single premium immediate annuity payout rates are strongly tied to prevailing interest rates and therefore an investor (e.g., a pre-retiree) who wants to hedge the future level of guaranteed income that they would be able to purchase (i.e., have more certainty surrounding lifetime income) could purchase a bond portfolio with a similar duration. This approach is similar to the concept of immunization for pension benefits, where a rise in interest rates decreases the value of the bond portfolio, but simultaneously decreases the net present value of the future cash flow stream, and vice versa should interest rates increase.

What is less clear is when and for whom fully immunizing the annuity purchase is called for, whether it be though some type of fixed income instrument or potentially investing in some type of riskier asset (e.g., equities). To better understand this issue, we conducted an analysis that assumed an investor has access to two investment options: equities or a bond portfolio with a duration of 15 years, which is approximately equal to the assumed duration of the annuity (note, the actual duration changes based on prevailing yields). The actual return model and the annuity pricing model are detailed in Appendix 4.

Rather than running a single scenario we vary two key variables: the equity risk premium (i.e., how much stocks outperform bonds, on average), and the assumed risk aversion with respect to the annuity funding.

The equity risk premium is an important assumption since the more stocks outperform bonds the more attractive they are going to be. For the funding risk aversion, we use a constant relative risk aversion (CRRA) utility function where the utility is a function of how well the annuity purchase is funded upon retirement. The risk aversion coefficient captures how the retiree would feel about a funding shortfall, whereby retirees with a higher risk aversion coefficient would be more risk averse (i.e., prefer higher allocations to fixed income).

For the first analysis, we assumed the retiree had no other wealth available for retirement and conducted a 5,000-run Monte Carlo simulation to determine the optimal equity allocation for the various scenarios. The results are included in Exhibit 18.

Exhibit 18 Optimal Portfolio Equity Allocation When Buying an Annuity at Retirement With No Outside Wealth

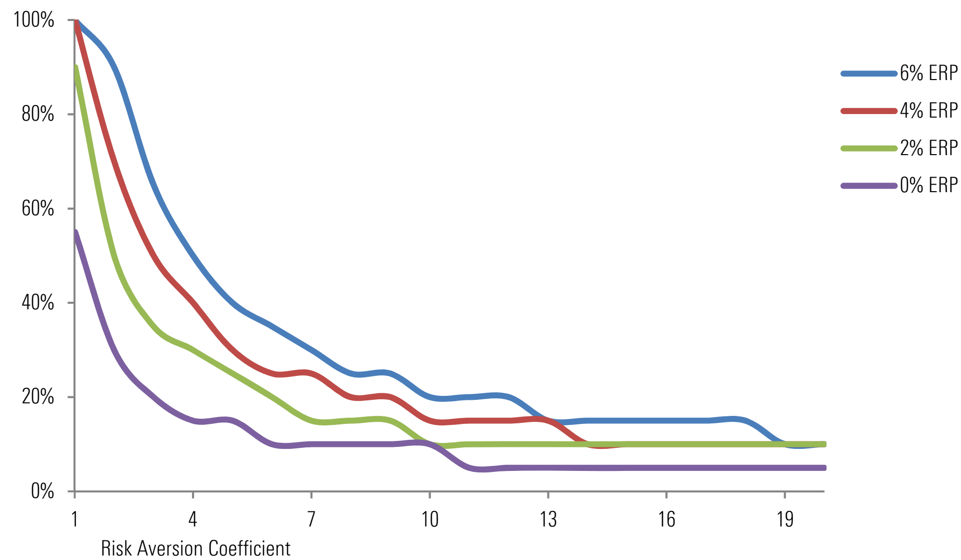

The level of risk aversion and assumed equity risk premium both significantly affect the optimal portfolio allocation. A retiree who expects a relatively low equity risk premium and is relatively risk averse with respect to funding the annuity would allocate a significant amount of the portfolio to bonds (~90%). This could be viewed as an investor engaging in a duration-matching LDI strategy. In constrast, a retiree who is not very risk averse and expects a higher equity risk premium would have a higher allocation to equities (~30%) although would still have a relatively conservative portfolio. The results in Exhibit 17 provide a perspective about how and why a defined-benefit plan may choose to purchase bonds that immunize the liability, especially if the plan sponsor/employer is risk averse.

One issue for the analysis conducted for Exhibit 18 is that it overlooks the fact individual investors (i.e., retirees) have other assets that can be used to fund consumption in retirement, such as Social Security retirement benefits, private pension benefits, home equity, etc. that are not going to be affected by the performance of the market. In other words, as noted previously, while the portfolio is important, performance will not necessarily be the sole driver of income in retirement. Therefore, to provide some perspective about how the optimal equity allocation changes in the presence of outside wealth, we reran the analysis assuming the portfolio value is only half of the retiree's total wealth (i.e., half of the retiree income is expected to come from some other non-market source, such as Social Security retirement benefits, private pension benefits, home equity, etc., and the other half would come from the purchase of an annuity). The results of the analysis are illustrated in Exhibit 19.

Exhibit 19 Optimal Portfolio Equity Allocation When Buying an Annuity at Retirement With Outside Wealth

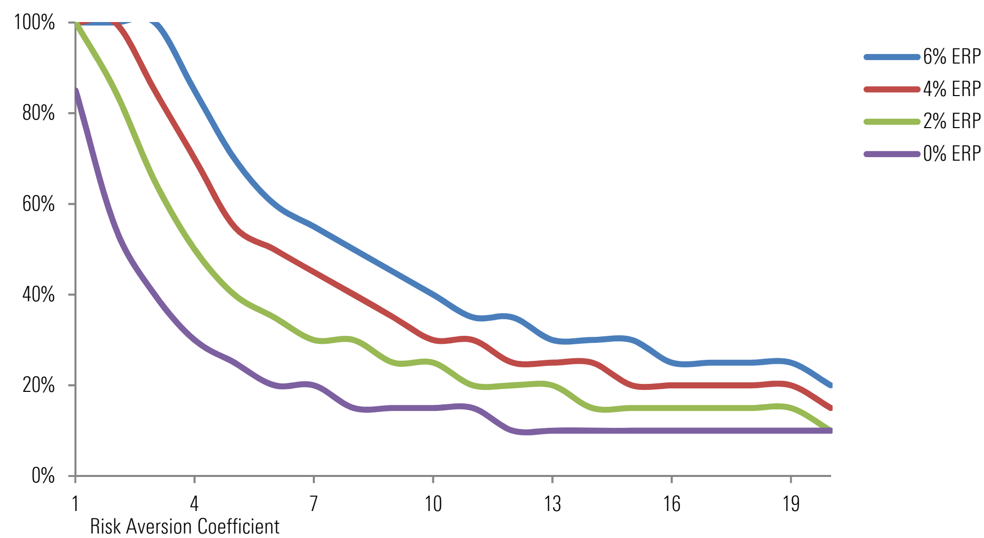

Incorporating the potential existence of outside assets significantly increases the optimal equity allocation. Even a retiree with a relatively high level of risk aversion and assuming a low equity risk premium would select a portfolio with at least 20% equities. A more moderate level of risk aversion coupled with a reasonable equity risk premium would easily suggest a portfolio allocation that has more than 50% equities.

These results are similar to our liability-relative optimized portfolios that incorporate Social Security benefits noted in Exhibit 17. Acknowledging the existence and risk characteristics of other household assets can be used to fund retirement results in a portfolio that is more aggressive than when the portfolio is treated in isolation.

Conclusions

Building efficient retirement income portfolios is complex. It is possible for individual investors and their advisors to borrow some of the liability-driven investing techniques used by pension funds, but it is important to understand how individual liabilities differ from those of institutional investors. For example, retirees have a *soft* liability, since the investor can adjust spending in retirement if needed, while pension funds are legally bound to provide payments at set levels regardless of its funded status. Also, the retiree's portfolio is only one of a number of assets available to fund the retirement liability. Many retirees receive guaranteed pension benefits, such as Social Security retirement benefits, which should change the risk target for the portfolio.

The growing use of LDI is certainly a step in the right direction for retirement investors; however, it is important to proceed with caution. Modeling the wrong liability can result in equally inefficient portfolios, usually ones that are overly conservative. Our analysis suggests that most investors are likely best served with portfolios that are more balanced and diversified in nature.

References

Agarwal, A., & Liu, X. 2016. “LDI: Taking A Holistic, Practical Approach.” Loomis Sayles Research Paper. http://www.loomissayles.com/internet/InternetData.nsf/0/D2648E5DE524236E85257FD2004F5725/$FILE/LDI-Holistic-Approach_FINAL.pdf

Blanchett, D.M. 2014. “Exploring the Retirement Consumption Puzzle.” *Journal of Financial Planning*, Vol. 27, No. 5, P. 34.

Blanchett, D.M., & Straehl, P.U. 2015. “No Portfolio Is an Island.” *Financial Analysts Journal*, Vol. 71, No. 3, P. 15.

Blitzer, D.M., Dash, S., & Murphy, P. 2009. “Equity Duration – Updated Duration of the S&P 500.” Standard & Poor's Research Paper, Sept. 21, 2009. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1476376.

Blattberg, R.C., & Gonedes, N.J. 1974. “A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices.” *Journal of Business*, Vol. 47, No. 2 (April), P. 244.

Clark, P.K. 1973. “A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices.” *Econometrica*, Vol. 41, No. 1 (January), P. 135.

Idzorek, T. 2008. “Lifetime Asset Allocations: Methodologies for Target Maturity Funds.” Ibbotson Associate Research Report. http://corporate.morningstar.com/ib/asp/detail.aspx?xmlfile=1409.xml

Idzorek, T., Tao, J., Xiong, J., & Blanchett, D. 2016. “The Impact of Liability-Relative Optimization and Non-Normal Returns on Asset Allocations.” Morningstar White Paper.

Leibowitz, Martin L. 1987. "Liability Returns: A new look at Asset Allocation." *The Journal of Portfolio Management*. Winter. P. 11-18.

Loomis Sayles. 2016. “LDI: Taking A Holistic, Practical Approach.” Loomis Sayles White Paper.

Mandelbrot, B.B. 1963. “The Variation of Certain Speculative Prices.” *Journal of Business*, Vol. 36, No. 4 (October), P. 394.

Sharpe, William F., and Lawrence G. Tint. 1990. “Liabilities: A new approach.” *The Journal of Portfolio Management*, Winter, P. 5–10.

Siegel, L.B., & Waring, M.B. 2004. “TIPS, the Dual Duration, and the Pension Plan.” *Financial Analysts Journal*, Vol. 60, No. 5 (September-October), P. 52.

Waring, M.B. 2004a. "Liability-Relative Investing." *The Journal of Portfolio Management*, Vol. 30, No. 4 (Summer), P. 8.

Waring, M.B. 2004b. "Liability-Relative Investing II: Surplus Optimization with Beta, Alpha, and an Economic View of the Liability." *The Journal of Portfolio Management*, Vol. 31, No. 1 (Fall), P. 40.

Waring, M.B. 2004c. "Liability-Relative Strategic Asset Allocation Policies." CFA Institute Conference Proceedings: The New World of Pension Fund Management. P. 43.

Xiong, J.X. 2010. "Using Truncated Lévy Flight to Estimate Downside Risk." *Journal of Risk Management in Financial Institutions*, Vol. 3, No. 3 (January), P. 231.

Xiong, J.X., & Idzorek, T.M. 2011. "The Impact of Skewness and Fat Tails on the Asset Allocation Decision." *Financial Analysts Journal*, Vol. 67, No. 2 (March/April), P. 23.

Ziemba, W.T., & Mulvey, J.M., Editors. 1998. *Worldwide Asset and Liability Modeling* (Cambridge, U.K.: Cambridge University Press).

Appendix 1: Additional Information on Liability-Relative Optimization

A key assumption of mean-variance optimization (MVO) is that returns are normally distributed, or follow a "bell curve." Xiong and Idzorek (2011) note that empirically we have seen that most asset classes and portfolios have returns that are not normally distributed. Empirically, the distribution of returns for most asset classes is skewed to the left of the mean, and the tails of the distribution are fatter than the tails of normal distributions. The normal distribution assigns what most people would characterize as meaninglessly small probabilities to extreme events that empirically—due to the skewness and kurtosis of real world asset return distributions—seem to occur approximately 10 times more often than the normal distribution predicts.

There are different models that can be used to account for "fat tails," such as the Lévy stable hypothesis from Mandelbrot (1963), the Student's t-distribution from Blattberg and Gonedes (1974), and the mixture-of-Gaussian-distributions hypothesis from Clark (1973). For our analysis we truncate the tails of the Lévy stable distribution (in order to remove the complications associated with infinite variance), which results in what is known as the truncated Lévy flight (TLF) distribution. Xiong (2010) demonstrated that the TLF model provides an excellent fit for a variety of asset classes, including the mean, asymmetries (skewness), and the thickness of the tails (kurtosis), as well as the minimum and maximum returns determined through the truncation process.

For our analyses, we use an optimization approach focused on minimizing the tail risk of a return distribution, focusing on conditional value-at-risk (CVaR). CVaR is closely related to VaR (value-at-risk) and is calculated by taking a probability-weighted average of the possible losses conditional on the loss being equal to or exceeding the specified VaR.

We fixed the probability level for the CVaR at 5% and take the average of these returns as our definition of risk for the optimization. This approach, which we refer to as *mean-CVaR optimization*, seeks to minimize the downside risk for a target return. This downside risk can be viewed from the lens of an asset-only optimization, or through the lens of both the assets and the liability, i.e., a surplus optimization. In either case, the definition of risk for mean-CVaR is the tail risk, which differs from MVO, since MVO focuses on the entire risk of the portfolio and does not focus exclusively on tail risk, like mean-CVaR.

Appendix 2: Return and Risk Assumptions

Exhibit A2.1 Return and Risk Inputs for Surplus Optimizations

Asset Class	Return [%]	Std Dev	Skew	Kurtosis
U.S. Large	7.41	15.59	-0.46	3.76
U.S. SMID	8.68	20.37	-0.26	3.59
Developed Markets ex-US	8.68	17.40	-0.56	4.51
Emerging Markets	11.87	28.73	-0.44	4.16
Long-Only Commodity	3.81	18.08	-0.30	5.00
US Real Estate	7.60	19.48	-0.74	11.79
Non-US Real Estate	6.16	23.43	-0.97	8.45
US Long-term Bonds	4.63	15.27	0.19	3.18
US Short-Term Bonds	3.12	2.40	0.47	3.40
US Long-term TIPS	3.58	11.07	-0.10	6.88
US Short TIPS	3.25	3.79	-0.25	8.41
Global High Yield	6.67	10.48	-1.90	16.58
Non-US Investment Grade Bonds	3.57	8.11	0.24	3.15
Cash	1.94	1.04	0.69	2.00
Inflation	2.25	3.00	0.33	3.64

Exhibit A2.2 Correlations for Surplus Optimizations

Asset Class	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1 U.S. Large	1.00	0.94	0.90	0.78	0.48	0.74	0.82	-0.08	-0.01	0.10	0.29	0.75	0.36	-0.05	0.11
2 U.S. SMID	0.94	1.00	0.88	0.79	0.48	0.80	0.80	-0.10	-0.02	0.07	0.30	0.78	0.31	-0.07	0.14
3 Developed Markets ex-US	0.90	0.88	1.00	0.90	0.60	0.69	0.92	-0.01	0.12	0.16	0.40	0.83	0.54	0.01	0.13
4 Emerging Markets	0.78	0.79	0.90	1.00	0.62	0.59	0.87	0.05	0.16	0.23	0.45	0.81	0.51	0.06	0.13
5 Long-Only Commodity	0.48	0.48	0.60	0.62	1.00	0.27	0.52	-0.09	0.18	0.18	0.57	0.57	0.49	0.10	0.34
6 US Real Estate	0.74	0.80	0.69	0.59	0.27	1.00	0.75	0.14	0.07	0.22	0.27	0.71	0.36	-0.05	0.08
7 Non-US Real Estate	0.82	0.80	0.92	0.87	0.52	0.75	1.00	0.13	0.18	0.27	0.39	0.84	0.56	0.02	0.08
8 US Long-term Bonds	-0.08	-0.10	-0.01	0.05	-0.09	0.14	0.13	1.00	0.57	0.71	0.27	0.11	0.40	-0.02	-0.30
9 US Short-Term Bonds	-0.01	-0.02	0.12	0.16	0.18	0.07	0.18	0.57	1.00	0.61	0.58	0.21	0.60	0.28	-0.17
10 US Long-term TIPS	0.10	0.07	0.16	0.23	0.18	0.22	0.27	0.71	0.61	1.00	0.65	0.32	0.54	-0.02	-0.06
11 US Short TIPS	0.29	0.30	0.40	0.45	0.57	0.27	0.39	0.27	0.58	0.65	1.00	0.60	0.52	0.13	0.30
12 Global High Yield	0.75	0.78	0.83	0.81	0.57	0.71	0.84	0.11	0.21	0.32	0.60	1.00	0.49	-0.07	0.22
13 Non-US Invt-Grade Bonds	0.36	0.31	0.54	0.51	0.49	0.36	0.56	0.40	0.60	0.54	0.52	0.49	1.00	0.10	0.02
14 Cash	-0.05	-0.07	0.01	0.06	0.10	-0.05	0.02	-0.02	0.28	-0.02	0.13	-0.07	0.10	1.00	0.16
15 Inflation	0.11	0.14	0.13	0.13	0.34	0.08	0.08	-0.30	-0.17	-0.06	0.30	0.22	0.02	0.16	1.00

Appendix 3: Retirement Spending Analysis

Exhibit A3.1 Distribution of Annual Real Changes in Nondurable Spending

Exhibit A3.2 Distribution of Annual Real Changes in Total Household Spending

Appendix 4: Returns Model and Annuity Pricing Model

Returns Model

The first step in the model is to select an initial bond yield (i.e., seed value) for the simulation. This is the bond yield that exists at the beginning of the retirement simulation; based on 10-year U.S. Treasury yields. Given an initial bond yield, the yields for the subsequent years are based on equation 1, where $\varepsilon_{\mathrm{Yld}}$ is an independent white noise that follows a normal distribution with a mean of 0% and a standard deviation of 1.0%. The resulting annual bond yield ($\mathrm{Yld_t}$) is assumed to be bounded between a minimum of 1.0% and a maximum of 10.0%.

$$Yld_t = 0.225\% + 0.955Yld_{t-1} + \varepsilon_Y \qquad [4.1]$$

After we determine the bond yield for a given year we estimate the bond return (r_{bond}) using equation 4.2, where $\varepsilon_{\mathrm{bond}}$ is assumed to have a mean of 0% and standard deviation of 1.0%.

$$r_{bond} = Yld_t + -15(Yld_t - Yld_{t-1}) + \varepsilon_{bond} \qquad [4.2]$$

It is worth noting the 1.0% standard deviation for the error term (ε_{b}) is not the assumed standard deviation for the asset class (bonds in this case), rather the standard deviation for the errors around the regression estimates. The actual standard deviation of bond returns is approximately 12.5%. The actual standard deviation is higher than the standard deviation for the error term because other factors, such as the current yield and the change in yield, are affecting the actual variability of returns. The long-term bond yield is assumed to be 5.0% within this model (i.e., this is the yield that it converges toward throughout the simulation).

Inflation is based on equation 4.3:

$$r_{inf} = 1\% + .4Yld_t + .5(Yld_t - Yld_{t-1}) + \varepsilon_{inf} \qquad [4.3]$$

Stocks are assumed to follow a normal distribution as noted in equation 4.4, where the mean is based on the median yield ($mYld_t$) for a given year plus an equity risk premium (ERP), which varies by scenario. The standard deviation is assumed to be 20% across scenarios.

$$N\ (mYld_t + ERP, 20\%) \qquad [4.4]$$

Annuity Pricing Model

The cost of the annuity upon retirement is determined based on the bond yield at retirement and mortality using equation 4.5:

$$AnnPay = \frac{1-\lambda}{\sum_r^R \frac{q_t}{(1+Yld_r)^r}} \qquad [4.5]$$

λ is the pricing load (i.e., insurance company profit and costs beyond base mortality) which is assumed to be 10%; R is the last year of the simulation, which is equal to 50 years after the age of assumed retirement (r); q_t the probability of the individual surviving to at least to retirement age r and is based on the Society of Actuaries 2012 Immediate Annuity Mortality table; Yld_t is the discount rate, which is the bond yield at the year of retirement.

About Morningstar Investment Management LLC

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

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

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

22 West Washington Street
Chicago, IL 60602 USA