Maximizing Personal Surplus: Liability-Driven Investment for Individuals

Michael Ashton

Abstract

To date, the financial literature has focused on very simple algorithms designed to improve the solution to the two-part challenge of determining the optimal portfolio asset-allocation strategy and determining the maximum sustainable withdrawal rate for retirees. Most research, for example the well-known "Trinity Study" of Cooley, Hubbard, and Walz, pursues the asset-allocation problem by maximizing long-run asset growth subject to a withdrawal rule and a given acceptable probability of remorse (a.k.a. shortfall). However, the Liability-Driven Investing (LDI) thought process improves the approach by seeking instead to maximize return for a given level of volatility of the portfolio surplus, rather than optimizing on the basis of the volatility of the assets themselves; by more closely matching assets and liabilities, the sensitivity of the strategy to unexpected returns, risks, and correlations is greatly decreased. I updated the Trinity Study to incorporate inflation-indexed bonds and then illustrate how the LDI thought process may be applied to individual investors.

Introduction

The two-part challenge of determining the optimal portfolio asset-allocation strategy and determining the maximum sustainable withdrawal rate for retirees during their "golden years" has received a fair amount of attention over the last two decades, and for good reason. Retirees face a complicated linear programming problem whose solution is non-intuitive and beyond the ken of most retirees. It is a key set of decisions to which financial advisors, aided by financial theory, should be able to add considerable value for their clients.

To date, the financial literature has focused on very simple algorithms designed to improve the decision-making process, and these algorithms have value. However, we can do better than these simple approaches, if the problem is phrased correctly.

The two-part planning issue is often thought of as (1) "How can I maximize return in the long run," and (2) "Given the riskiness of this portfolio, how do I set withdrawal policy[1] to ensure I actually reach the long run" with some trade-off of portfolio riskiness and withdrawal policy riskiness taken into account?

A typical approach, used for example in Bergen (1994) and treated more thoroughly by Cooley, Hubbard and Walz (1998) in the famous "Trinity Study," is to use historical sampling methods to determine the range of outcomes that would historically have resulted from a particular combination of asset allocation and withdrawal policies. For example, Cooley et. al. established that given a portfolio mix of 75 percent stocks and 25 percent bonds and a withdrawal rate of 6 percent of the initial portfolio value for a 30-year holding period (over the historical interval covered by the study), the portfolio would have failed 32 percent of the time for, conversely, a 68 percent success rate.[2]

But this approach, unfortunately, is limited by the same factors that are its strength. Using a lengthy data series covering a wide range of market environments greatly improves upon an approach that uses historical mean returns as the benchmark for what can be withdrawn. But it also limits the possible outcomes to those observed in the historical data set. To some extent, this is a difficult problem to overcome since our history is, after all, the only one we have, but it may obscure real risks to the plan that are not diagnosed only because they haven't happened yet. Thus, using this approach, the Trinity authors can conclude, "…the presence of common stocks provides upside potential and holds the promise of higher suitable exchange rates."

---

[1] A withdrawal policy describes how the investor will draw on the portfolio over time. It is usually phrased as a proportion of the original portfolio value, and may be considered either a level nominal dollar amount or adjusted for inflation (a real amount).

[2] The Trinity study includes several tables that calculate success rates ignoring inflation. However, since (a) the asset returns themselves have components that reflect inflation; for example, the portion of the bond yield that reflects inflation compensation, and (b) the client's spending pattern is generally related at least somewhat to inflation, those success rates are not particularly useful in practice. However, the authors also calculate success rates incorporating an inflation adjustment, and it is a result from that table (which is replicated as Table 1 in this paper) that I cite here.

Are we sure of this? The conclusion is only possible because during the limited history available, especially as of 1998, *most* sufficiently long periods have seen stocks appreciate. But this is clearly not a guarantee; there is no reason to suppose, for example, that it is *impossible* for stocks to decline in real terms for 40 consecutive years (however *improbable* that is). We cannot determine from the historical record the chances of such an occurrence, unless we wish to declare returns to be Gaussian with a true mean and true variance similar to the historical mean and variance—and that particular state of the world is one that has actually been judged on the basis of considerable evidence to be fairly unlikely.

This may sound like mere theoretical critique with no obvious prescription for improvement. Perhaps we can improve modeling approaches, but aren't we after all limited to making assumptions of some kind about mean returns, variances and correlations *regardless* of the approach we take? Yes, but a healthy skepticism that the future will look much like the past can lead us to prefer solutions that **depend less** on such assumptions.

Fundamentally, retirement planning involves financing a relatively low-volatility stream of real and nominal expenses[3] with a high-volatility stream of returns. Scott, Sharpe, and Watson (2008) argue from theory that payout rules are inherently inefficient because "they attempt to finance a constant, non-volatile spending plan using a risky, volatile investment strategy." If this were a pension fund we would say the liabilities are roughly fixed in real space and the assets are very volatile. There is, in fact, a developing literature that concerns how to jointly consider both the asset mix and the spending requirement. That process is called Liability-Driven Investing, or LDI.

What Is LDI?

The basic goal for these plans is to invest today in such a way as to ensure that future liabilities can be covered. The pension fund is not supposed to maximize its real assets, subject to risk constraints based on the variability of the asset portfolio; it is supposed to maximize the pension's economic surplus, subject to constraints on the variability of the surplus.

This makes a difference because the volatility of the present value of the liability stream is a component of the volatility of the funding status, which is itself a component of the volatility of firm value. Moreover, because the firm is not guaranteed to exist forever, but the plan is supposed to last for a long time regardless of the firm's status, the period of relevant volatility to the plan is a lot shorter than the plan's investment horizon.

The mandate to maximize economic surplus subject to its variability recognizes that the pension fund is not just a pool of assets, but a pool of liabilities that are to be funded with assets. The mission of the plan sponsor is to consider these jointly. This has spawned an industry of "Liability-Driven Investment" research and sell-side services.

[3] This is generally true if the "tails" are explicitly hedged with insurance. For example, in the absence of long-term care insurance the real expenses of a retiree may become quite volatile towards life's end, but with such insurance these tails are mitigated.

But if LDI is the right approach for a fund, then it may also be the right approach for an individual. Let us consider how we can look at the individual's retirement issue in a similar context

How Is LDI Applicable to Individuals?

We must be careful to think about an individual investor's liabilities correctly. The liabilities of a pension or post-employment benefits fund are generally contractual or semicontractual in nature. That is, there is a promise by the plan sponsor to pay a stream of future income to the beneficiary, or to provide a stream of services of a certain type (as is the case with postretirement medical plans). In the individual's case, not all of the liabilities are contractual. Some, of course, are: the home mortgage is a nominal liability with known characteristics. So, too, are car payments, alimony or child support payments, and other such things. But an individual also has noncontractual liabilities that are no less crucial to the planning exercise: "optional" fixed costs like life- or long-term-care insurance premiums, as well as other costs of living that tend to rise with inflation such as food, clothing, utilities, entertainment, travel and the like. Moreover, investors may want to either explicitly fund a bequest or to treat part of the portfolio as essentially carved out for that purpose, and this may be nominal or real in nature.

Of course, there are some important differences between a pension fund's mandate and that of an individual investor. One of the main ones has already been mentioned: the fact that a pension fund generally has contractual terms that make its liabilities relatively clear, while an individual investor generally does not. A pension fund also benefits from actuarial averaging; its exposure is to a *general* increase in longevity, for example, or a spike in average medical costs. An individual by contrast is exposed to the randomness of a single spin of the wheel of fortune when it comes to his own longevity or the possibility of large medical bills due to his own poor health (see footnote 4). And beyond the mandate, of course, individuals have access to a narrower set of investment options since (for example) most investors cannot access the over-the-counter swaps market, the credit default swap market, or any of several other markets that are mostly reserved for institutional investors. In general, though, this has less effect in an LDI framework because most of those markets have little utility for hedging an individual's exposures (e.g., CDS) and the democratization of risky product exposures—for example, through structured notes or levered ETF products—allow the replication of many of these investment options although admittedly at a higher price.

Why does the LDI mandate construction make sense in the case of an individual as well as for a pension plan, endowment or foundation? By reducing the mismatch between assets and liabilities, we reduce the importance of the assumption that the experienced returns, variances, and correlations will be like those of the historical data set. At the limit, if assets and liabilities are precise mirror images of one another, then the returns, variances and correlations between assets will be completely irrelevant. This is the ultimate comfort level: "I don't know what the future holds, and I don't care."

On the other hand, *failing* to consider how well assets and liabilities match can lead to planning absurdities. Scott, Sharpe and Watson (2008) cite a tool provided by Vanguard (2008) that recommends a 3.75, 4.75 or 5.25 percent rate of withdrawal based on whether the investor is in a conservative, moderate, or aggressive equity allocation, respectively. In other words, the response to a **riskier asset mix** is to create a **higher-cost spending mix**. *This is akin to borrowing money for a bigger house than you can afford, because house prices should go up in the long run.*

There is another, subtle advantage to an approach of risk-adjusted surplus maximization, rather than asset maximization, and it involves investor behavior. Consider two cases: in case A, the pension fund (or individual investor) enjoys a 100 percent gain in its portfolio, followed by a 50 percent loss. In case B, the 50 percent loss comes first and is followed by a 100 percent gain. Obviously, the asset portfolio has the same terminal value in both cases. However, in case A there is a much greater chance that the investor's spending pattern has increased than in case B, so that the actual realized surplus will probably be smaller in the former case.

This problem is especially pronounced in the case of public pensions. An overfunded pension seldom stays that way for long, as politicians seize on the opportunity to add more benefits; conversely, benefits are almost never *cut* symmetrically for an underfunded pension. Consequently, and significantly, **funding status variance induces negative funding status expectation**. This "political ratchet" effect has an analog in individual investors, who are likely to respond to runs of good luck by increasing withdrawals, whatever the master plan says. Accordingly, managing the asset/liability match so as to decrease swings in funding status is likely to reduce the behavior-management challenge.

Approach

The methodology I use to model the value of liability-driven investing for individuals is a Monte Carlo simulation based on an estimate of steady-state, long-run, asset-class returns, using historical inflation and asset-class volatilities and correlations based on the 1978 to 2008 period.[4] I included headline inflation, inflation-indexed bonds, nominal bonds and equities.[5] Nominal bonds and equities are typical inclusions in such an exercise, but it is necessary to include TIPS (more generally, inflation-indexed bonds) because neither of those other two asset classes protects against inflation, the largest exposure for most retirees. Note also that since I am only using a subset of the universe of securities available to both individual and institutional investors, the conclusions are applicable to both classes of investor.

With the exception of the results summarized in Table 2 (and which exception is explained later), I used the following long-term nominal return assumptions:

- CPI: 3.1 percent[6]
- Inflation-Indexed Bonds: 5.1 percent[7]
- Nominal bonds: 5.1 percent[8]
- Equities: 7.6 percent[9]

The precision of these return assumptions is less important for the illustration than the general ranking of them and the correlations between them.

Another simplification that I made for the sake of modeling, which is generally consistent with withdrawal policy literature, is to ignore the taxation of income and investment returns. Clearly, this sacrifices considerable realism but it dramatically simplifies the comparability of the subjects. In modeling the surplus for a flesh-and-blood client, tax implications in the structure of the portfolio and the tax-advantaged or tax-free vehicles used must certainly be considered, as well as the possibility that the status of these vehicles and the tax laws generally may change over time.

---

4 This methodology can be easily extended to incorporate additional asset classes, such as commodity indices, not included here.

5 Headline inflation data was obtained from the Bureau of Labor Statistics at http://www.bls.gov/cpi. Inflation-indexed bond returns used the Barclays Inflation-Linked Bond index from 1997 to 2008, and simulated returns prior to that. Nominal bond returns were represented by the Lehman/Barclays Aggregate Bond index. Equity returns were the S&P 500 total returns with dividends reinvested.

6 The actual 1926 to 2009 arithmetic average.

7 Assumed 2 percent real yield, plus expected inflation.

8 Over the long run, additional credit spread should be balanced by defaults if credit is fairly priced. Higher returns indicated by some indices may be the result of survivor bias that neglects the negative tails associated with defaults in the portfolio, or it might indicate an actual, systematic mispricing of credit. I assume, for simplicity, that the pricing is efficient.

9 Consists of 2 percent real growth plus inflation plus 2.5 percent dividend yields; see (Cornell & Arnott, 2008).

Given these parameters, I simulated multiple random "histories" of various lengths for all included asset classes and economic variables; for each "history" I computed the portfolio performance of a given asset allocation regime, assuming (costless) annual rebalancing to policy weights.

For each year in each history, I calculated the nominal income and expenses, given the inflation outcome produced, and then computed the year-end assets as the year-beginning assets minus expenses, plus revenues, plus or minus asset returns. At the end of the run, I recorded the total assets and adjusted for the simulated price index to get "2010 dollars."

Then, taking in aggregate all of the runs for a particular asset-allocation regime, I computed the mean and standard deviation of the distribution of terminal outcomes associated with that investment and withdrawal policy. I also calculated the proportion of terminal outcomes that resulted in "portfolio failure," defined as the condition of ending wealth being less than zero.

Note that in each case, the "withdrawal policy" is treated as a liability such as it would appear in an LDI setting: it represents a series of future negative inflation-adjusted cash flows.

With this approach, I attempted the following:

1. I studied a range of "portfolio success rates" and produced a chart that is analogous to Table 3 in the Trinity Study, "Inflation-Adjusted Portfolio Success Rates: 1926 to 1995."[10]
2. I updated this chart to incorporate inflation-linked bonds (which asset class was not available for the period covered by the Trinity Study).
3. I considered one particular, illustrative case in which the situation of the retiree is not best modeled by simple withdrawal rates.

Results and a conclusion, based on these results, follow.

[10] For a more accurate comparison to the Cooley et. al. results, I used historical average (simple, uncompounded) return rates from the 1926 to 1995 period for equities (12.5 percent) and inflation (3.2 percent) and the median return rate for bonds (4.0 percent). The reason I used the median rate for bonds is that the mean is skewed very high (5.7 percent) because two outlier points in 1982 (+43.8 percent) and 1985 (+30.90 percent). The other 68 values lie between -8.1 percent and +19.9 percent. Using the mean, rather than the median, implies much higher success rates for bond-heavy portfolios than the Trinity authors found, implying that the median in this case is a better fit to the historical data.

Results

Table 1 is a replica of Chart 3 from the Trinity Study (Cooley, Hubbard, & Walz, 1998). It shows inflation-adjusted portfolio success rates for contiguous historical payout periods between 1926 and 1995.

Table 2, alongside it, gives the results from my simulation for the same subset of assets, but a very different methodology—the one described above.

Table 3 presents the differences between these two charts (Chart 2 – Chart 1).

It is apparent that the methodology I am using produces results that are broadly similar to those produced by the historical survey method the Trinity authors used. It is clear, for example, that for spartan withdrawal policies, portfolio exhaustion is rare while for prodigal withdrawal policies, portfolio exhaustion is typical. Where there are differences, it largely represents the fact that the simulation data are "smoother" than the historical data. This close fit occurred despite the fact that the sample variance of returns, which I took to be constant over the Trinity interval of 1926 to 1995, was not actually constant over the interval. Moreover, the correlation structure between the two asset classes and inflation is, of course, not stationary, and there are some mean-reversion tendencies in real markets that are not incorporated in my method. Nevertheless, the close match between Table 1 and Table 2 supports the value of the methodology.

The care that was required to make Table 2 mirror as close as possible the appearance of Table 1 points out both strengths and weaknesses of the historical survey method. One big weakness is that using historical outcomes (or historical distributions, as I did when I tried to replicate their results numerically) involves using distribution characteristics that are likely to be wrong on a forward-looking basis. A large part of the equity return from 1926 to 1995, for example, was the result of multiple expansion; the 12.5 percent return to equities compared to 3.2 percent inflation is certainly not the likely, nor even the naïve, estimate of future returns. It would be very aggressive to assume that return distribution would be replicated in the ensuing period, even on average. In later Tables, I use a 7.6 percent expected return on equities, computed as described in footnote 10.

One strength of the historical survey method is that it will capture some of the "fat tail" outcomes that using a lognormal distribution will not; that fact, however, is also a weakness, for while that method captures some of the "fat tail" events, it doesn't capture ones *that did not actually happen* during the survey period.

In other words, it is the peculiar configuration of our actual history that produces the high portfolio success rates—especially for equity-heavy portfolios—that Trinity authors (and others) find but, more importantly I think, highlights the danger of relying too much on grids of this type. As noted earlier, the historical sampling method can only sample from existing history, which can promote over-optimism based on the range of economic/investing environments actually observed—sort of a financial anthropic principle. We would like to reduce as much as is practical this reliance on future conditions being quite like past conditions.

TABLE 1
Inflation-Adjusted Portfolio Success Rates: 1926 to 1995
(Source: Cooley, Hubbard, & Walz, 1998)

	Withdrawal Rate as a % of Initial Portfolio Value:									
Payout Period	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%
100% Stocks										
15 Years	100	100	100	91	79	70	63	55	43	34
20 Years	100	100	88	75	63	53	43	33	29	24
25 Years	100	100	87	70	59	46	35	30	26	20
30 Years	100	95	85	68	59	41	34	34	27	15
75% Stocks/25% Bonds										
15 Years	100	100	100	95	82	68	64	46	36	27
20 Years	100	100	90	75	61	51	37	27	20	12
25 Years	100	100	85	65	50	37	30	22	7	2
30 Years	100	98	83	68	49	34	22	7	2	0
50% Stocks/50% Bonds										
15 Years	100	100	100	93	79	64	50	32	23	13
20 Years	100	100	90	75	55	33	22	10	0	0
25 Years	100	100	80	57	37	20	7	0	0	0
30 Years	100	95	76	51	17	5	0	0	0	0
25% Stocks/75% Bonds										
15 Years	100	100	100	89	70	50	32	18	13	7
20 Years	100	100	82	47	31	16	8	4	0	0
25 Years	100	93	48	24	15	4	2	0	0	0
30 Years	100	71	27	20	5	0	0	0	0	0
100% Bonds										
15 Years	100	100	100	71	39	21	18	16	14	9
20 Years	100	90	47	20	14	12	10	2	0	0
25 Years	100	46	17	15	11	2	0	0	0	0
30 Years	80	20	17	12	0	0	0	0	0	0

Note: Numbers rounded to the nearest whole percentage. The number of overlapping 15-year payout periods from 1926 to 1995, inclusively, is 56; 20-year periods, 51; 25-year periods, 46; 30-year periods, 41.

TABLE 2
Inflation-Adjusted Portfolio Success Rates:
Monte Carlo Simulation

	Withdrawal Rate as a % of Initial Portfolio Value:									
Payout Period	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%
100% Stocks										
15 Years	100	99	95	92	85	78	68	58	49	42
20 Years	98	96	91	82	74	64	52	42	35	30
25 Years	96	92	84	75	64	56	46	36	28	23
30 Years	96	88	81	73	58	50	40	32	26	17
75% Stocks/25% Bonds										
15 Years	100	99	98	94	85	74	63	51	38	31
20 Years	99	96	89	81	70	56	47	32	22	16
25 Years	98	95	85	71	59	47	34	23	16	8
30 Years	96	91	78	67	52	36	28	17	12	8
50% Stocks/50% Bonds										
15 Years	100	100	98	93	83	68	53	38	23	17
20 Years	100	98	91	78	60	42	28	16	9	5
25 Years	98	93	79	61	42	27	15	8	5	2
30 Years	97	89	69	49	29	17	10	4	3	1
25% Stocks/75% Bonds										
15 Years	100	100	98	91	74	54	31	16	10	4
20 Years	100	98	81	58	37	17	8	3	1	0
25 Years	99	86	63	33	16	7	2	1	0	0
30 Years	94	71	47	21	8	4	1	0	0	0
100% Bonds										
15 Years	100	99	92	73	49	25	11	5	1	1
20 Years	98	84	58	28	13	4	1	0	0	0
25 Years	88	60	29	11	3	1	0	0	0	0
30 Years	74	36	15	4	1	1	0	0	0	0

Note: Numbers rounded to the nearest whole percentage.

TABLE 3
Illustration of Differences Between Monte Carlo Simulation Results (Chart 2) and Historical Study Results (Chart 1)

Withdrawal Rate as a % of Initial Portfolio Value: Payout Period	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%
100% Stocks										
15 Years	0	-2	-5	1	6	8	5	3	6	8
20 Years	-2	-5	3	7	11	11	9	9	6	6
25 Years	-4	-8	-3	5	5	10	11	6	2	3
30 Years	-4	-7	-4	5	-1	9	6	-2	-1	2
75% Stocks/25% Bonds										
15 Years	0	-1	-2	-1	3	6	-1	5	2	4
20 Years	-1	-4	-1	6	9	5	10	5	2	4
25 Years	-2	-5	0	6	9	10	4	1	9	6
30 Years	-4	-7	-5	-1	3	2	6	10	10	8
50% Stocks/50% Bonds										
15 Years	0	0	-2	0	4	4	3	6	0	4
20 Years	0	-2	1	3	5	9	6	6	9	5
25 Years	-2	-7	-1	4	5	7	8	8	5	2
30 Years	-3	-7	-7	-2	12	12	10	4	3	1
25% Stocks/75% Bonds										
15 Years	0	0	-2	2	4	4	-1	-2	-3	-4
20 Years	0	-2	-1	11	6	1	0	-1	1	0
25 Years	-2	-7	15	9	1	3	0	1	0	0
30 Years	-6	0	20	1	3	4	1	0	0	0
100% Bonds										
15 Years	0	-1	-8	2	10	4	-7	-11	-13	-9
20 Years	-2	-6	11	8	-1	-8	-9	-2	0	0
25 Years	-12	14	12	-4	-8	-1	0	0	0	0
30 Years	-6	16	-2	-8	1	1	0	0	0	0

Note: Numbers rounded to the nearest whole percentage. Darker shading indicates a larger absolute difference. Green indicates Chart 2 values are higher; red indicates Chart 1 values are higher.

In Table 4, I update the prior example using the forward-looking return estimates and incorporate an asset class that was not included in the original Trinity study (or in most other withdrawal rate studies): inflation-indexed bonds.[11] To limit the size of the chart, I have chosen one investing horizon (25 years) at which all paths are measured. I assume that the investor, who has need for an inflation-linked withdrawal (by assumption), first decides on his concentration in inflation-indexed bonds (TIPS, specifically) and then decides how to allocate the remainder of his portfolio.

[11] Because inflation-indexed bonds were not available over the 1926 to 1995 interval and limiting the data set to 1997 to 2009 would risk too many idiosyncrasies in the correlation matrix, I used the correlation matrix that produced Table 2, but with 1997 to 2009 correlations between the asset classes and inflation-indexed bonds.

The top section of Table 4, in which the allocation to TIPS is zero, corresponds to the relevant "25 Years" line from Table 2 associated with the stocks/bonds mix shown in the left-most column, but has been recalculated based on the different return assumptions described above. The next section, in which the TIPS allocation is 20 percent, shows the concentrations in stocks and bonds assuming that the remaining portfolio is allocated 100/0, 75/25, 50/50, 25/75, or 0/100. For example, the second row of that section, labeled "60%/20%", indicates that the scenarios in that row correspond to a portfolio in which 20 percent is allocated to TIPS, and the *remainder* is split 75 to 25 percent. This produces a portfolio of 20 percent TIPS, 60 percent stocks and 20 percent bonds.

TABLE 4
Inflation-Adjusted Portfolio Success Rates: Monte Carlo Simulation
(TIPS, Stocks, Bonds)

Stocks/ Bonds	Withdrawal Rate as a % of Initial Portfolio Value: 3%	4%	5%	6%	7%	8%	9%	10%	11%	12%
TIPS Allocation: 0%										
100%/0%	85	69	52	41	28	21	16	12	6	3
75%/25%	91	72	58	40	26	19	10	6	4	2
50%/50%	95	82	60	39	23	10	4	3	1	1
25%/75%	97	84	58	32	15	5	2	1	0	0
0%/100%	96	74	47	23	8	2	1	0	0	0
TIPS Allocation: 20%										
80%/0%	90	76	59	43	26	17	11	6	2	2
60%/20%	94	81	60	42	24	15	8	3	2	1
40%/40%	99	86	62	36	17	6	2	1	0	0
20%/60%	100	88	60	26	9	4	1	0	0	0
0%/80%	97	84	51	22	5	2	0	0	0	0
TIPS Allocation: 40%										
60%/0%	96	81	59	41	20	13	6	2	1	0
45%/15%	98	89	65	38	17	6	2	1	0	0
30%/30%	100	91	64	30	14	3	0	0	0	0
15%/45%	100	93	60	24	6	1	0	0	0	0
0%/60%	100	89	50	20	4	1	0	0	0	0
TIPS Allocation: 60%										
40%/0%	100	91	64	34	12	5	1	0	0	0
30%/10%	100	95	68	28	10	2	0	0	0	0
20%/20%	100	96	67	25	5	0	0	0	0	0
10%/30%	100	97	63	18	2	0	0	0	0	0
0%/40%	100	94	54	12	1	0	0	0	0	0
TIPS Allocation: 80%										
20%/0%	100	97	70	22	3	0	0	0	0	0
15%/5%	100	98	70	18	1	0	0	0	0	0
10%/10%	100	97	63	14	2	0	0	0	0	0
5%/15%	100	97	58	12	1	0	0	0	0	0
0%/20%	100	96	58	10	1	0	0	0	0	0
All-TIPS Portfolio										
N/A	100	94	52	11	1	0	0	0	0	0

Note: Numbers rounded to the nearest whole percentage.

An elementary observation is that as the concentration of TIPS in the portfolio increases, the probabilities of portfolio success at low withdrawal rates increases while the probabilities of portfolio success at high withdrawal rates declines. This should not be terribly surprising. Clearly, for any reasonable time horizon, a 10 percent withdrawal rate is only sustainable if the investor takes extraordinary risks *and* draws the lucky part of this distribution. The "success rates" shown for high withdrawal rates are somewhat misleading, because they do not distinguish between portfolios that fail toward the end of the planning period from ones that fail spectacularly, closer to the beginning of the planning period. Each is recorded as a single portfolio failure, but the former circumstance is clearly preferable.

Another way, then, to summarize the effect of the addition of TIPS to the portfolio is to look at the range of potential outcomes over the entire planning period. Chart 1 shows two "cones" that correspond to two different strategies. For each cone, the upper line corresponds to the 90^{th} percentile outcome for that strategy and portfolio, at each point in time; the lower line corresponds to the 10^{th} percentile outcome; the dashed line represents the median. Put another way, the cones represent a trimmed-range of outcomes for the two strategies, over a 25-year time period.[12] The two strategies illustrated are (1) a 75 to 25 percent mix of stocks and bonds and (2) a 80 percent TIPS, 5 percent stocks, 15 percent bonds mix. In both cases, the investor starts with $1,000 and a 5 percent withdrawal rate is assumed.

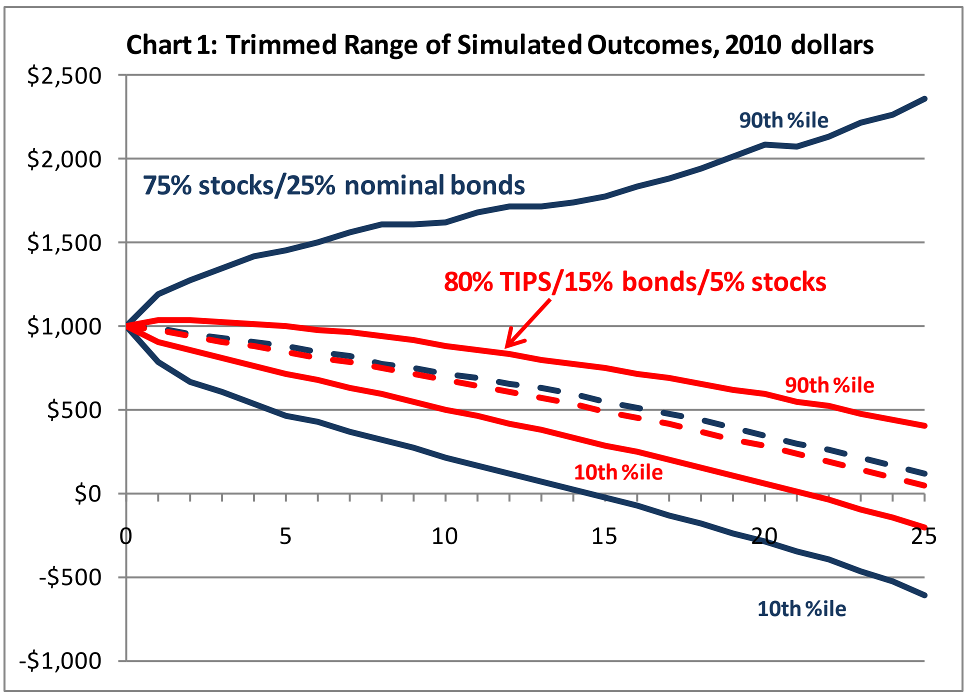

[12] Each line, however, doesn't represent a single simulation. Each point on the top line, for example, might come from a different simulation that happened to be doing very well at that particular moment. It would be a very peculiar outcome indeed to have a portfolio that gained (or lost) in such a linear fashion over *any* simulated history!

According to Table 4, both of these portfolio allocation/withdrawal strategies have roughly a 58 percent chance of portfolio success (note how closely the median lines track), but as Chart 1 shows, this number does not convey the relative riskiness of the two approaches. In the case of the equity/bonds mix, the investor stands a 10 percent chance of running out of money as soon as the 15th year: bankrupt, with 10 years left on the horizon. In the case of the TIPS-concentrated strategy, this unfortunate possibility doesn't enter at the 10 percent level until the 22nd year of the simulation. Table 5 illustrates this measure of portfolio failure and shows why the higher "success rates" of the equity-heavy portfolio at high withdrawal rates may be deceiving. Consider an 8 percent withdrawal rate. Table 4 tells us that a 100 percent equity allocation (top line) should "succeed" 21 percent of the time, compared to 0 percent for an All-TIPS portfolio (bottom line). But Table 5 cautions that by year 8, 10 percent of the equity portfolios have failed while with the All-TIPS portfolio you are reasonably confident of funds holding out at least until year 12. Even at a low 3 percent withdrawal rate, there is a 10 percent chance of an all-equity portfolio failing in the first 19 years.

TABLE 5
Years to 10% Portfolio Failure Rate: Monte Carlo Simulation
(TIPS, Stocks, Bonds)

	Withdrawal Rate as a % of Initial Portfolio Value:									
Bonds	3%	4%	5%	6%	7%	8%	9%	10%	11%	12%
TIPS Allocation: 0%										
100%/0%	19	15	13	11	9	8	7	7	6	6
75%/25%	24	18	15	12	11	10	8	8	7	7
50%/50%	*	22	17	14	12	10	10	9	8	7
25%/75%	*	24	19	16	13	11	10	9	8	8
0%/100%	*	24	19	16	13	11	10	9	8	8
TIPS Allocation: 20%										
80%/0%	24	19	14	13	11	9	9	8	7	7
60%/20%	*	22	17	15	12	11	9	8	8	7
40%/40%	*	25	19	16	13	11	10	9	8	8
20%/60%	*	*	20	17	14	12	11	9	9	8
0%/80%	*	*	20	16	14	12	11	9	9	8
TIPS Allocation: 40%										
60%/0%	*	22	17	14	12	10	9	8	8	7
45%/15%	*	24	19	16	13	11	10	9	8	8
30%/30%	*	*	20	16	14	12	11	10	9	8
15%/45%	*	*	21	17	14	12	11	10	9	8
0%/60%	*	*	20	17	14	12	11	10	9	8
TIPS Allocation: 60%										
40%/0%	*	*	19	16	14	12	10	9	8	8
30%/10%	*	*	21	17	14	12	11	10	9	8
20%/20%	*	*	21	17	15	13	11	10	9	8
10%/30%	*	*	22	17	15	13	11	10	9	8
0%/40%	*	*	21	17	14	12	11	10	9	8
TIPS Allocation: 80%										
20%/0%	*	*	21	17	15	13	11	10	9	8
15%/5%	*	*	21	17	15	13	11	10	9	8
10%/10%	*	*	22	17	15	13	11	10	9	8
5%/15%	*	*	21	17	15	13	11	10	9	8
0%/20%	*	*	21	17	14	13	11	10	9	8
All-TIPS Portfolio										
N/A	*	*	21	17	14	12	11	10	9	8

Note: Numbers rounded to the nearest whole year Asterisk (*) indicates portfolio failure rate is less than 10% at the planning horizon (25y).

Of course, the more conservative allocation also misses out on the 90th-percentile upside. But as observed by Scott, Sharpe and Watson (2009), the upside variance of the equity-concentrated mix does not "pay for" the downside variance because those surpluses likely have considerably less salience to the 90-year-old investor than do the shortfalls. Scott et. al., observe that this is an inefficient use of investing dollars, and if the investor could sell some of that upside variance to buy downside protection, he likely should. As the TIPS-centric cone illustrates, the investor essentially *can* pursue such a strategy.

The tight range of outcomes from the TIPS-centric strategy does not result merely from a narrow range of possible nominal returns for TIPS securities. Indeed, from 1998 to 2008 the standard deviation of annualized performance of TIPS in nominal terms was around 5.5 percent, which was actually more volatility than displayed by the Lehman/Barclays Aggregate, which had a 3.8 percent standard deviation of annual returns (although equities' risk was 21.1 percent over that time period). To illustrate that the variance of TIPS returns, compared to equity returns, was not the main driver of this improvement, I simulated the two portfolios, keeping the original mean returns for each asset while assuming that the standard deviation for all three asset classes was equal to the 5.5 percent of TIPS. The results are summarized in Table 5.

TABLE 5
Dispersion of Outcomes at 25th Year (90th – 10th Percentiles)

	Original Std. Deviations	Uniform Std. Dev = 5.5%
a. 75% Equity, 25% Bonds	$2,971	$1,308
b. 80% TIPS, 15% Bonds, 5% Equity	$605	$543
Proportion dispersion (b/a)	20%	42%

Removing differences in the variances of the different asset classes accounts for only about one quarter (22/80 percent) of the difference in the dispersion of terminal outcomes. The implication is that the narrow range of after-inflation outcomes experienced by the investor in case 2 derives in large part from the superior matching of the TIPS investment flows to the inflation-adjusted withdrawal demand of the investor. In other words, this investment outcome is driven partly by the implied **liabilities** of the investor. This style of result—tighter ranges of potential outcomes for the net asset-liability portfolio around the expected outcome—is a signature of a successful LDI-based strategy.

However, not all demands on the portfolio by the retiree will necessarily be real (that is, inflation-adjusted) demands. How do the outcomes change when the retiree's liabilities change?

An Example of LDI Applied to Retirement

Consider a retiree with $1mm in cash, who wishes to fund his retirement for the next 25 years. Initially, we will assume that all of the retiree's expenses are real; that is, they all are expected to increase generally with inflation over time. The question is, what is the *minimum* amount of spending the investor should be able to sustain with a low-risk investment strategy?

The risk-minimizing strategy clearly would be to buy an inflation-adjusted annuity, if they were generally available. If the real rate curve were flat at 2 percent, this investor could secure an annuity that would pay $51,220 in today's dollars[13] every year for the next 25 years. This is the maximum withdrawal rate that is guaranteed to never fail (over a 25-year horizon).

Such a product is not generally offered, but we can reasonably approximate this payout by holding a ladder of TIPS bonds weighted appropriately. Table 4 shows that a slightly coarser approach of simply holding the TIPS index produces, at a 5 percent real withdrawal rate, a 52 percent chance of portfolio success. This percentage near 50 percent indicates the median outcome is approximately break-even.[14]

Now, let us consider the same retiree, but this time we posit that he has committed nominal payouts (perhaps, for long-term care or life insurance policies) of $42,500 per year for the next 25 years. This, however, constitutes the bulk of his expenditure. The risk-minimizing strategy for this investor, of course, is different from that of the first investor. If nominal rates are at 5.1 percent then he can buy a 25-year annuity of $42,500 per annum with no inflation adjustment for $593,035, and with the balance of $406,965 he can purchase a 25-year real annuity that pays $20,845 in today's dollars annually. This 2.085 percent of original assets is the maximum withdrawal rate that is guaranteed not to fail over a 25-year horizon, conditional on its being purchased along with a 5.1 percent fixed annuity. Again, as an alternative to the annuities the retiree could approximate the payout by holding a ladder of nominal bonds to fund the nominal liabilities and a ladder of TIPS to fund the real withdrawal rate. In our example, the investor would hold these two bond portfolios in a ratio of approximately 59:41 percent.

These two investors have identical initial endowments of $1mm, but different liability structures, and clearly the risk-minimizing portfolio is decidedly different. What happens when we compare the success rates of the same portfolio sets for these different investors?

Table 6 is analogous to Table 4, except that I've modeled the second investor's combined nominal-and-real liability stream rather than the first investor's real-only liability stream (I have also narrowed the range of withdrawal rates, because very high withdrawal rates almost always produce portfolio failure). The withdrawal rate modeled is the inflation-adjusted amount that is withdrawn *in addition* to the $42,500 fixed withdrawal.

---

13 $PVAF_{2\%,25}$=19.5236.

14 Since the TIPS index is not a maturity-matched ladder of TIPS, there is some variance around this median point that may well be unacceptable to our investor; thus, in practice an investor pursuing the maximum-safe-withdrawal strategy will either need to withdraw slightly less than 5 percent per annum or carefully construct a true annuity-like portfolio.

TABLE 6
Inflation-Adjusted Portfolio Success Rates for Investor With Nominal Liabilities

Stocks/ Bonds	Withdrawal Rate: 0%	1%	2%	3%	4%
TIPS Allocation: 0%					
100%/0%	84	71	54	41	30
75%/25%	93	76	56	40	23
50%/50%	97	83	57	36	18
25%/75%	100	91	60	27	7
0%/100%	100	94	48	14	2
TIPS Allocation: 20%					
80%/0%	91	74	56	39	26
60%/20%	96	79	59	38	20
40%/40%	100	89	64	31	12
20%/60%	100	96	64	17	3
0%/80%	100	97	49	8	1
TIPS Allocation: 40%					
60%/0%	96	85	58	39	20
45%/15%	99	90	64	34	13
30%/30%	100	96	64	25	6
15%/45%	100	99	64	12	1
0%/60%	100	99	53	6	0
TIPS Allocation: 60%					
40%/0%	100	93	62	33	10
30%/10%	100	96	67	26	5
20%/20%	100	99	65	17	1
10%/30%	100	99	63	6	0
0%/40%	100	100	54	4	0
TIPS Allocation: 80%					
20%/0%	100	98	66	16	1
15%/5%	100	99	64	12	0
10%/10%	100	99	57	6	0
5%/15%	100	99	58	5	0
0%/20%	100	98	52	4	0
All-TIPS Portfolio					
N/A	100	95	50	6	0

Note: Numbers rounded to the nearest whole %.

Let's compare the success rates for the two investors in the case at the margin. Chart 2 plots the portfolio success rates for an all-fixed-income portfolio divided in various ways between nominal debt and inflation-indexed debt (TIPS) for the two investors at real withdrawal rates of 5 percent (for the investor with all-real liabilities) and 2 percent real + $42,500 (for the investor with a mix of real and nominal liabilities). You can see that the former investor does consistently better as we move to higher concentrations of TIPS, while the mixed-liabilities investor does better in the middle range and worse at the ends.[15]

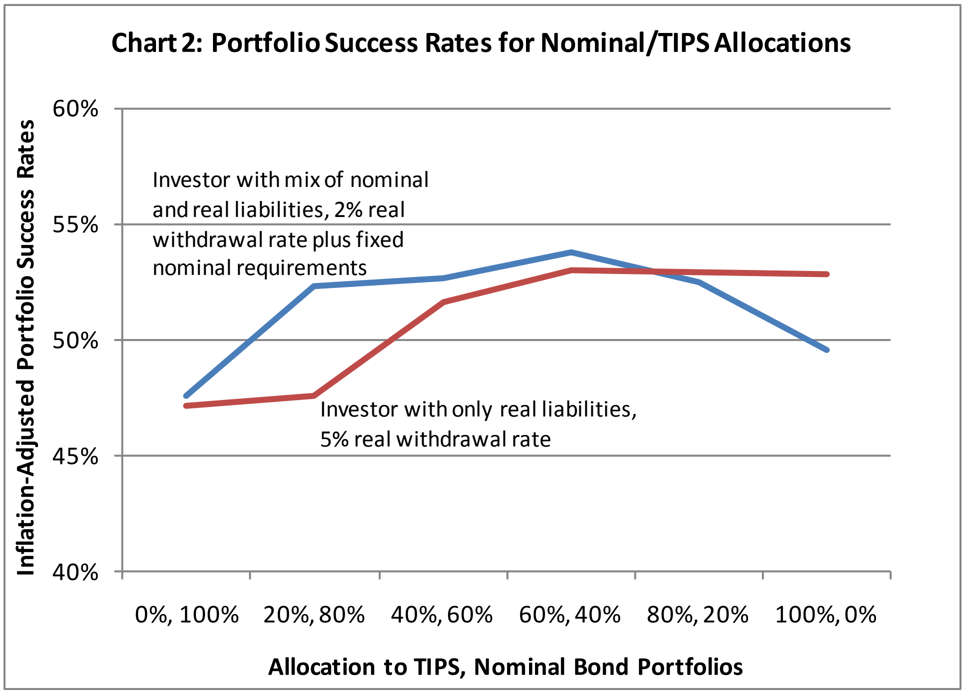

Both lines above represent cases near the median of the expected return distributions. When we move away from the middle of the bell curve into the tails, the difference in performance can be more stark. Chart 3 shows the portfolio success rates for the same two investors under slightly more conservative withdrawal rates. It turns out that choosing an asset portfolio to more-closely match liabilities is not *less* important if the withdrawal rate is more conservative, but *even more important.*

Take a moment to consider why that is true. When our expected outcome is well above the "safe" level, failure is more remote. Portfolio failure now requires the **joint** occurrence of (a) poorly-matched assets and liabilities combined with (b) bad luck. As noted earlier, the primary reason for matching assets and liabilities closely is to minimize the impact of bad draws from the return distribution.

So, in Chart 3, the investor with all-real liabilities assures a high probability of success if he combines a conservative withdrawal rate with a liability-matched portfolio. But if this investor holds a mostly nominal portfolio, this admits a new possibility that a high-inflation environment (against which he is no longer hedged) will lead to portfolio failure. The investor with two different sorts of liability streams has some natural diversification as a result, so while he does worse with heavy concentrations

15 The success rates around 50 percent do not conflict with the prior assertion that the portfolios can be guaranteed success if annuitized at rates of 5.122 percent and 2.085 percent respectively. Even a small amount of volatility around the 'maximum safe annuitized withdrawal rate' will result in a distribution that is half "successful" and half "failures," but if the volatility of the surplus is small then the "failures" will be reasonably small failures – see Chart 1.

of either TIPS or nominal bonds, he is never as exposed as the all-real-liabilities investor who holds all nominal bonds.

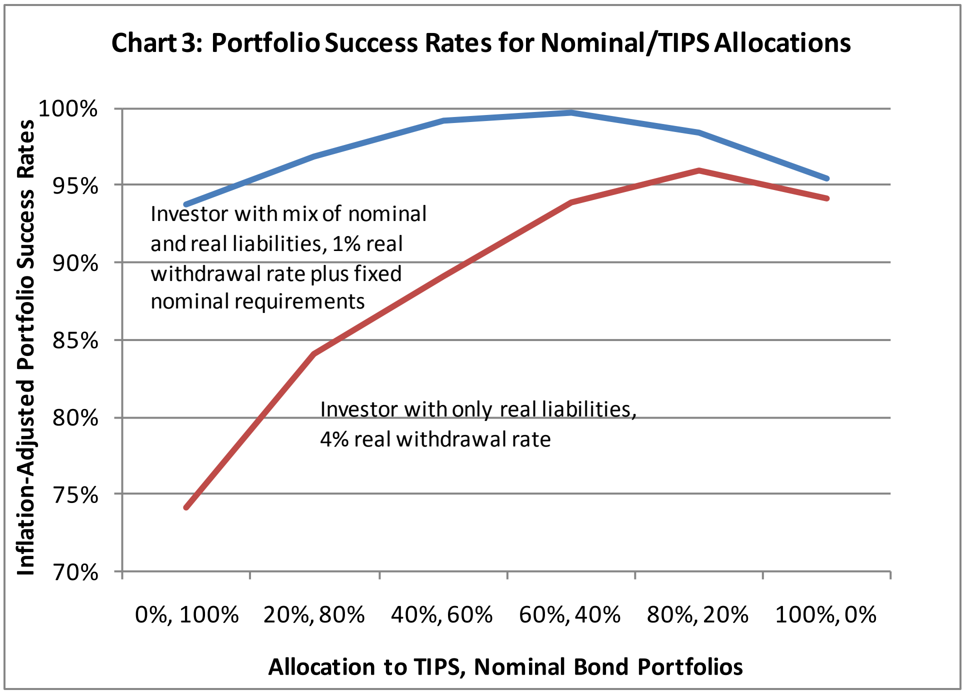

Let us advance one step further and examine the case of the investor with a mix of real and nominal liabilities a little more closely.

Consider an investor who has expenses that generally rise with inflation, but has a **one-time nominal payment** that is due in 15 years (say, a bequest he has made to his alma mater, or a balloon payment due on his mortgage) of $1,250,000. Our investor with $1,000,000 in assets is still solvent, because (at a 5.1 percent nominal discount rate) that liability is only worth $593,000 today. We assume that the investor also plans to withdraw $10,000 (that is, 1 percent of the initial portfolio value), adjusted for inflation, each year.

We will examine two decision-making processes. In the first, the investor simply makes an allocation between nominal and real bonds and maintains that allocation for the next 25 years. In the second, the investor allocates according to his initial liabilities and maintains that allocation until the 15th year. After the nominal payment is made, the investor switches to an all-real bonds allocation since his remaining liabilities are all inflation-sensitive.

Chart 4 shows the results. The fire-and-forget investor does as well as the LDI-aware investor if he selects a portfolio between 40 and 70 percent TIPS, but worse if he chooses a portfolio with an allocation to real bonds outside of that range. The fire-and-forget investor never does meaningfully *better*.

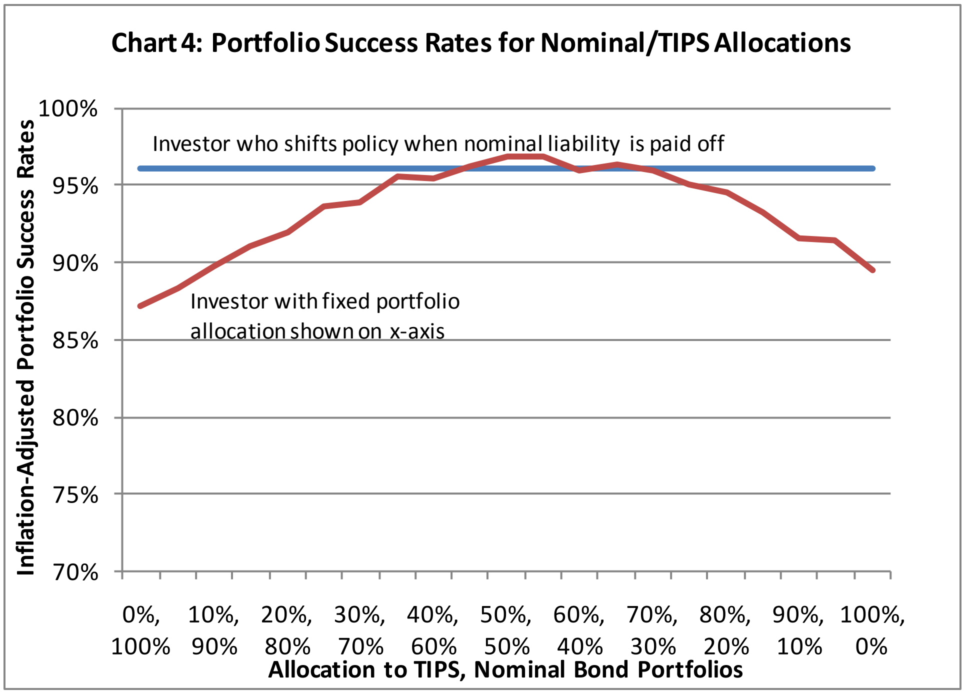

This example seems contrived, because it **is** contrived. But consider one implication. It would seem that the argument for being LDI-aware is not particularly strong here, because the fire-and-forget investor would probably choose something around 40 percent TIPS/60 percent nominal simply based on his initial liabilities. However, in this example the investor's real and nominal demands never changed except on schedule, and the big nominal liability was some years out (so that the remaining years of the simulation had less time to accentuate errors). Reality is somewhat messier, and so the important observation to make here is that the LDI-aware investor never did measurably worse than the naïve investor. In a less-contrived, more realistic situation, that investor will likely be more comfortable (and therefore, perhaps, able to invest surpluses in a more return-seeking than risk-minimizing fashion).

Conclusion

Traditionally, retirement planning is conducted the way most investing has been conducted for a long time: the goal is treated as maximizing assets.

More recently, the retirement planning goal has been rephrased as maximizing nominal or real annuitized income, whether explicitly annuitized or self-annuitized through the use of "maximum sustainable withdrawal rules." This of course reduces to the same goal: maximize the asset pool at the time of annuitization, or continue to attempt to grow the asset pool with annuitization as a fall-back if balances decline dangerously.

With this construct, it is natural to pursue the asset-allocation decision as maximizing long-run asset growth subject to a withdrawal rule and a given acceptable probability of remorse (a.k.a. shortfall). As we have seen over the last decade-plus; however, the long run can become intolerably long *even for an institutional investor deemed to have a perpetual planning life.* Consequently, such investors have been increasingly focused on controlling the risk associated with surplus volatility (that is, the volatility of the *difference* between their assets and their liabilities). Such planning is driven by two important observations, although rarely made explicit:

1. Reducing the mismatch between the volatile assets and the usually less-volatile liabilities implies that long-term planning depends less on assumptions about mean returns, variances and correlations because errors in these assumptions operate only on the mismatch, rather than on the whole asset base.
2. Reducing the mismatch actually has positive expectation in that it reduces the "political ratchet" that occurs when the presence of surpluses stimulates increased liabilities (because additional promises are made) while the presence of deficits does not produce a mirror-image reduction in liabilities.

Both of these observations are also valid in the case of individual investors. Individual investors benefit from relying less on the accuracy of their assumptions about the probable future states of the world, and individual investors, if anything, probably suffer more from the behavioral "ratchet effect." This suggests that liability-driven investing (LDI) concepts should be considered by the individual investor and his or her advisors.

Planning in the context of LDI requires recognizing the character of the liabilities, and that means more than answering the simple question of whether the liabilities are real or nominal in nature. The advisor should also take into account the timing of the liabilities and the timing and character of any other assets and annuities (e.g., Social Security). If the client has need for a particular nominal or real sum at the planning horizon (for example, to fund bequests), it is crucial to include it since this zero-coupon liability can have a significant effect on the optimal portfolio. I illustrated in Chart 4 that an LDI-aware strategy is difficult to dominate.

Expanding the methodology further, we can also differentiate between liabilities meant to fund transportation, housing, or medical care expenses and model them as being linked to the appropriate inflation sub-index. Whether this will result in improved portfolio efficiency is open to question, because there are no assets currently on the market that specifically track these inflation subcomponents, but some assets (for example, the family home) may track better against a different index.

Even with all of these improvements, an investor is hardly invulnerable. The investor remains exposed to changes in asset values that are not driven by the same macro factors driving his liabilities (for example, the bankruptcy of a particular bond holding), to changes in tax regulations and legislated changes in retirement entitlements, and to personal risks like extended illness or the happy accident of longevity beyond the planned horizon. Some of these can be hedged with insurance products, but some simply demand the investor maintain an extra margin of safety. Liability-driven investing techniques will sharply reduce some of the largest macro portfolio risks, but the importance of consulting with an experienced advisor who is familiar with these techniques, throughout the pre- and postretirement periods, will continue to grow in importance as the expected period of retirement continues to lengthen.

Bibliography

Bengen, W. P. (1994). Determining Withdrawal Rates Using Historical Data. *Journal of Financial Planning , 7* (4), 171-180.

Cooley, P. L., Hubbard, C. M., & Walz, D. T. (1998). Retirement Savings:Choosing a Withdrawal Rate That Is Sustainable. *American Association of Individual Investors Journal , 20* (2), 16-21.

Cornell, B., & Arnott, R. (2008). *CRA Insights: Credit Crisis. The Basic Speed Law for Capital Markets Returns.* CRA International.

Scott, J. S., Sharpe, W. F., & Watson, J. G. (2009). The 4% Rule - At What Price? *Journal Of Investment Management* (Third Quarter 2009), Available at SSRN: http://ssrn.com/abstract=1115023.

Vanguard. (2008). *"Planning & Education: Retirement: Managing Your Retirement: Tap Your Assets: Determine how much you can withdraw."* Retrieved Feb 26, 2008, from (cited in and accessed by Scott, Sharpe, and Watson 2008): https://personal.vanguard.com/us/planningeducation/retirement/PEdRetTapDetermineWDContent.jsp.

Michael Ashton, CFA, is the Managing Principal at Enduring Investments LLC, a boutique inflation markets consulting and investment management firm. He can be reached at *m.ashton@enduringinvestments.com*.