CDI: pushing beyond the LDI horizon?

Julia Bonafede, Steven Foresti and **Dimitry Mindlin** of Wilshire Associates argue the thinking behind the dominant trend of liability driven investing is fundamentally flawed

We've all seen the light; asset allocation decisions should be expanded from an asset-only framework to one that considers benefit and spending commitments. This is an uncontroversial view that is now almost universally shared. In fact, a dissenting opinion would be difficult to find. However, as the sun rises over this consensus view and sheds light on the dawn of a new era in asset allocation, many investors and their managers have been blinded by the light.

Opinions regarding the role that financial commitments should play in the asset allocation process differ. Much of the ambiguity revolves around the overused term "liability". Is the generically named "liability" a series of future benefit commitments or an accounting measurement of its present value? If it is accounting based, is it ABO, PBO, VBO, TBO, HBO, UFO, ELO, BTO, Tivo or something else all together? Perhaps it's time to review the historical context that has led to this renaissance in liability-based investing, look at the objectives of some of the narrow liability driven investing (LDI) offerings that have been presented as solutions, and put forth views on an appropriate and comprehensive approach to "commitment driven investing" (CDI).

Historical background

Many proponents of LDI solutions point to the recent experience by defined benefit (DB) pension plans as evidence that asset allocations are broken and a fix is needed. The centerpiece of evidence, the so called smoking gun, is the historical volatility of funding ratios. The funding ratio for pension plans is an accounting observation that compares a plan's "assets" to its "liabilities". In this calculation, "assets" are represented by either the market value or actuarial value of assets; while "liabilities" reflect a present value measurement of a plan's benefit commitments. For example, the steep decline in the aggregate funding ratio of corporate DB plans within the S&P 500 Index, from 125% in 2000 to a mere 83% in 2002, paints a picture of pension funding instability. How can funding levels deteriorate by so much in just two years? The argument put forth is that this roller coaster ride of funding levels is a clear sign that assets are not being structured appropriately. Since a pension plan is designed to fund future benefits, assets should be invested relative to that "liability benchmark" and should, therefore, track it more closely, keeping funding ratios stable over time.

While this argument seems completely reasonable on the surface, it has several major shortcomings. The LDI explanation begins and ends very logically, but takes a sneaky turn in the middle. Here is a brief summary of how it often goes:

1. The goal of a pension plan is to fund future liabilities;
2. Assets should be structured to increase the likelihood of meeting these liabilities;
3. Therefore, the true benchmark for asset performance is the fund's liability benchmark;
4. The solution is a "liability driven" investment strategy that allocates assets to mimic the characteristics and behaviour of a plan's liabilities (as defined in accounting statements), thereby lowering risk.

Points one and two are right on target as they recognise a plan's commitments, but notice how this future stream of promised benefits was quickly simplified into a single accounting value represented by a "liability benchmark" in point three. So what's the big deal? Why does the seemingly innocuous term "liability" create such a problem when substituted for "commitment" in the debate? The term "commitment" captures the very nature, complexity and economic reality of a pension plan: funding a stream of future benefits that are usually not known with complete certainty (i.e. they are contingent on other factors, such as inflation). Introducing the term "liability" moves the discussion from the realities of a pension fund to the fascinating world of accounting standards. We have now arrived at the crux of the issue. If one accepts that assets should be managed to behave like "liabilities" and these "liabilities" behave like bonds, where does that drive our investment decision? What asset behaves like a bond? A bond, of course.

The current LDI landscape

As it stands now, the main shortcoming of the typical LDI solution is that it is product driven and the product, being bond-like, over-emphasises interest rate risk. There are two problems with this simplification. First, while there is nothing wrong with investing in bonds – even a very high concentration of bonds – as they offer a very high level of expected risk-adjusted return, they provide relatively low levels of absolute return. Accepting low levels of return, albeit at a lower level of expected asset risk, will increase the expected cost of running a pension plan and will result in larger, but more stable contributions. More importantly, measuring pension liabilities like bonds is an unnecessary and incorrect oversimplification that can introduce significant risks to plans structured according to this philosophy. A primary driver of bond market volatility is interest rate behaviour, so a pension plan that is valued like a nominal bond is assumed to have similar reactions to changes in interest rates, but this assumption is dubious at best.

Interest rate movements have a dramatic effect on bond valuations. When rates rise, bond prices go down; when rates fall, bond prices go up. Since pension accounting liabilities are calculated via a discounting procedure that is closely related to a relevant yield curve, the same relationship exists between pension accounting liabilities and the relevant yield curve. The primary goal of most LDI approaches is to manage interest rate risk. In other words, the strategies are designed to ensure that a plan's assets will move similarly in both direction and size to movements in accounting liabilities that result from fluctuating interest rates. The argument concludes that if assets and "liabilities" are moving in tandem, future funding levels and contributions will be much more stable and predictable. While this may be true, the lower levels of expected returns from an immunisation strategy suggest that this stable level of future contributions will likely be at a much higher level than could be expected otherwise. But, as we mentioned above, these higher costs are only one factor in the story. The larger concern is the problem of managing assets against accounting targets rather than managing them to meet the economic reality of having sufficient funds to pay future benefit commitments. The typical LDI approach, by over-emphasising interest rate risk, causes plans to be structured to fulfil an accounting artifact, rather than their pension obligations.

In reality, the typical commitment made to employees through a defined benefit pension plan has little, if anything at all, to do with changes in interest rates. Benefit formulas are based on factors such as years of service and ending salary. A typical plan participant will find nothing in their plan package materials to suggest that their benefit calculation is tied to the level of interest rates. Whether rates are 5%, 10% or 20%, when the participant retires, their benefit stream will be indifferent.

The one (indirect) connection between a participant's benefit and interest rates is based on the impact that inflation has on both of these values. Interest rates are directly affected by inflation as lenders demand a yield that is expected to protect them from a dollar's loss of value to inflation over time. Therefore, higher than expected inflation pressures yields higher, while lower than expected inflation accommodates lower interest rates. Similarly, wage inflation and general inflation are very highly correlated, so high levels of inflation lead to high levels of wages, a factor in the typical benefit calculation, which ultimately lead to higher levels of benefits. This triangular relationship between interest rates, unexpected inflation and future benefit streams is an under-appreciated risk in off-the-shelf LDI approaches.

To highlight this risk, imagine an environment of rising inflation with no end in sight, which we'll refer to as the "Now what?" scenario. Higher than expected wage inflation leads to higher than expected benefit commitments, which, to be met, require larger than expected assets. Now let's consider our LDI portfolio and evaluate how it might behave during this rising inflation environment. Thanks to our LDI immunisation strategy, our funding level stays very well insulated during this period. Certainly our fund's value has likely declined as higher interest rates, brought on by rising inflation, drove bond prices down, but the good news is that our accounting liabilities also dropped significantly in value, so we look fine on paper. But how do we look off paper; in the real world where we've made commitments to real people? Wages have gone up beyond the levels we expected, so future benefit payments are higher than we had previously anticipated. If assets are down and promises are up, why is it we still look ok on paper? Because of the accounting bias. Accounting measurements of the pension commitment went down because they are now being discounted by the higher level of current interest rates that were brought about by high levels of expected inflation. The economic reality, however, is that we now owe more, not less, to our participants.

This discussion is not to suggest that other asset mixes would have gone up in a period of

rising inflation, but by recognising the role that inflation has in a pension plan's risk, a more comprehensive risk management approach to the asset allocation process would have revealed the valuable contribution of investments that mitigate a portion of that risk and could have insulated the value of assets from the ravages of inflation. The LDI approach, on the other hand, ignored this possibility and subjected the plan to a 'double whammy': falling assets with rising commitments. As a result, a product that was promoted as "the low risk strategy" actually delivers a "double whammy" on the downside and a "double bonus" on the upside, which are not the attributes of a low risk hedging strategy.

Commitment driven investing

If the benefit stream doesn't behave like a bond, and we believe that asset liability models should produce asset allocation policies that fund present and future benefit payments, how should it all work?

Taking another look at a pension plan's goals may help:

• The goal of a pension plan is to fund its commitments;

• Assets should be managed to increase the likelihood of meeting these commitments.

Therefore, the true benchmark for asset performance should be one that satisfies the objectives of the constituents of the plan, i.e. the participants and shareholders/taxpayers. In which case:

1. The main objective of plan participants is to maximise the safety of promised retirement benefits;

2. The main objective of shareholders or taxpayers is to fund promised retirement benefits at the lowest cost.

Thus, it stands to reason that the role of the policy portfolio is to minimise the cost of running the plan at a given level of benefit security or, stated differently, maximise the probability of meeting benefit commitments at a given level of cost. The centerpiece of our analysis should therefore be the actual benefit stream that is calculated based on the eligibility requirements and benefit formula in place for a particular plan. All actuarial and accounting liabilities start out with this object. All liabilities represent measurements of this object.

Any liability satisfies a particular objective and reveals a particular property of the benefit stream, but may be inadequate for another objective that requires the analysis of other properties of the stream. Outside of the area of compliance with relevant regulations, there is no need to restrict ourselves to a particular discounting procedure for the stream before we've determined the policy portfolios we wish to consider. We analyse the benefit stream in its entirety – nothing is lost, nothing is hidden.

Recall the four-point "LDI explanation" we discussed earlier. The lapse in its logic is based on the unfortunate fact that the term "liability" has two entirely different meanings – it represents the benefit stream in points one and two and the present value of the benefit stream in point three. To avoid this kind of confusion, we have to assign different terms to the benefit stream and its present values. Going forward, the stream of benefit payments promised to plan members is a "pension commitment". The term "liability" is reserved for conventional actuarial and accounting reports. Let's reframe our "LDI discussion" in the context of our new wording conventions.

1. The goal of a pension plan is to fund pension commitments;

2. Assets should be managed to maximise the likelihood that the money will be readily available whenever needed;

3. Equally important, assets should be managed to minimise the cost of providing pension benefits;

4. The solution is a "commitment driven" investment strategy that allocates assets to simultaneously minimise the risk and cost of the plan.

This is the foundation of CDI. It is important to note that framing LDI in this manner does not mean that the health of financial statements should be sacrificed in the short or medium term. Less volatile fixed income strategies can, and in certain instances should, be employed to help mitigate financial statement volatility. The difference between the investment products currently being offered to plan sponsors and the CDI framework is that we try to satisfy both goals without destabilising the plan in the long run.

What's in a measurement?

We measure certain and uncertain objects in the past, present and future. The primary reason for the existence of various measurements of all kinds of things is, outside of sheer unadulterated curiosity, that measurements help us make better decisions. Some objects allow precise measurements as "known values". As an example of a "known value" measurement, the price of a zero-coupon Treasury bond that pays US$100 in ten years is a precise measurement of a certain future event (getting $100 in ten years).

Other objects are inherently uncertain and do not allow direct and precise measurements. For example, the future investment return on a portfolio of risky assets is uncertain. However, given the assumptions about expected returns, standard deviations, and correlations between asset classes, we can measure the riskiness of the portfolio. The standard deviation of the portfolio return is a risk measurement – it is a measurement of the uncertainty of the underlying object (the future value of the portfolio). This risk measurement is very useful in an optimisation procedure that allows us to distinguish between diversifiable and non-diversifiable risks and identify efficient portfolios.

In a variety of endeavours, we measure not only tangible known objects, but uncertain future ones as well. As far as the defined benefit system is concerned, pension plans face numerous risks that must be measured and managed. Unfortunately, neither accounting nor actuarial valuations contain risk measurements. Instead, the uncertainties of pension commitments and asset returns are depicted as single "known values".

The need for better risk measurements and risk management tools is greater than ever. The principles of CDI are well-suited to producing the measurements needed for the efficient management of cost and risk. The valuation that produces these new risk measurements and conventional actuarial/accounting valuations are fundamentally different.

Therefore, this unique valuation deserves a new name, and we call it asset-liability valuation (ALV). As mentioned before, the term "liability" here doesn't specify any particular single concept but does indicate that financial commitments are taken into account directly.

Asset-liability valuation (ALV)

ALV is an analytical framework that provides important pension risk measurements. One of the strengths of ALV is that it "pre-experiences" all possible investment outcomes implied in capital market assumptions by analysing asset return and risk with the periodicity of the benefit stream. In other words, capital market assumptions are used to measure the probability of all portfolio mixes meeting future benefit commitments at a given level of cost.

As shown in Chart 1, this distribution of assets required to meet future commitments will contain information about the best case and worst case scenarios based on experiencing these market returns and the probability that the plan will need to accumulate additional funds to meet this obligation.

By setting the discount rate as a random variable, we now have a near-infinite number of possible investment outcomes derived from the combinations of asset classes and total returns implied by our capital market assumptions. This is not a simulation, as we can narrow down this investment universe by identifying optimal asset mixes through a deterministic closed-form process.

Chart 1: Distribution of required assets to fund commitments

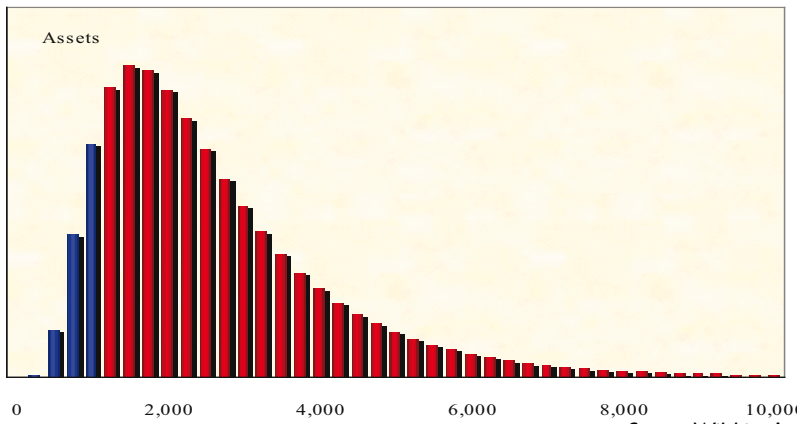

Source: Wilshire Associates

Chart 2: Cost-risk efficient frontier

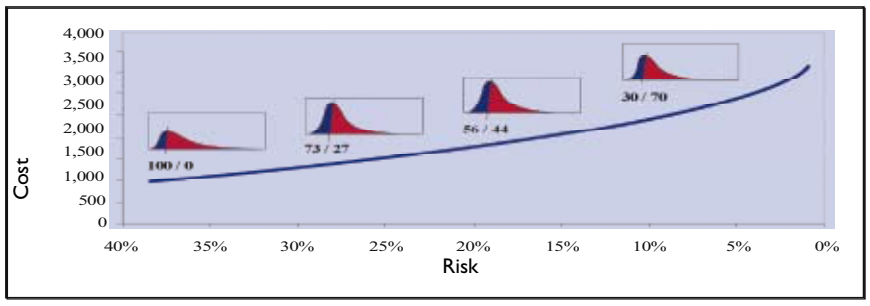

U.S. Equities	0.62	0.46	0.35	0.18	0.15
U.S. Bonds	0.00	0.09	0.19	0.35	0.38
Int Equities	0.38	0.27	0.21	0.12	0.11
TIPS	0.00	0.18	0.25	0.35	0.37
Geometric Return	0.0815	0.0742	0.0692	0.0598	0.0582
StDev of Return	0.1676	0.1234	0.0985	0.0611	0.0560
Arithmetic Return	0.0941	0.0812	0.0737	0.0616	0.0597
RAMean	2,325	2,409	2,542	2,947	3,039
RAStDev	1,545	1,135	952	700	666

Source: Wilshire Associates

Optimisation, not simulation

Unlike asset-only approaches, which require only forecasts of asset behaviour, a comprehensive asset-liability valuation process analyses more information to generate efficient portfolios. Instead of forming policies based on expected return and risk, we are looking for policies that meet our dual objectives of guaranteeing the safety of benefits and minimising the risk of having to make contributions.

This additional information includes a correlation matrix of asset returns, standard deviations of return and one final, but important, ingredient: a second correlation matrix that describes the relationship between asset class returns and the expected inflation characteristics of the particular plan's benefit stream. The first three components are generic to the properties of each asset class used. The last component is unique to each plan sponsor and directs the optimiser to produce asset mixes that satisfy our two objectives.

ALV leads to better information on which to make allocation decisions. As shown in Chart 2, we can now plot the cost-risk efficient frontier.

The cost-risk frontier defines risk as the probability of a successful outcome (i.e. assets are sufficient to fund benefits), versus the traditional efficient frontier that measures risk as the standard deviation of expected return. The y-axis on the cost-risk frontier describes the present value of assets required to fund a particular plan's benefit commitment. The x-axis describes the range of probabilities that the given asset mix will fund the obligation. The intersection of the points on each axis creates the frontier. Each intersection is a policy mix that has its own distribution of outcomes. The policy mix has been produced to minimise the extreme values, or tails, of negative outcomes. The negative outcomes are indicated in red.

The final feature of the cost-risk frontier is that there is a specific policy mix on the frontier where the cost (level of contributions) is too high to justify investing in the policy mix on the frontier. At that point it is more cost effective to immunise a portion of the portfolio.

Summarising the argument, CDI versus LDI

The concept of LDI has received a tremendous amount of attention recently. While the core catalyst fuelling this focus is to be embraced, namely that the allocation of pension assets should be driven by a plan sponsor's commitments, we have serious concerns regarding the specific direction most LDI approaches have taken. The vast majority of LDI solutions over-simplify the role of asset allocation decisions by limiting their primary objective to an analysis of accounting measurements as the basis for determining the liability. As a consequence, these approaches routinely overemphasise interest rate risk by valuing benefit commitments like bonds. As we discussed, these strategies, which are put forth as reducing risk, can in fact expose a plan sponsor to risks that are more directly in conflict with their core objective – maximising the safety of meeting future commitments.

In contrast, the allocation of pension assets is at its core an exercise in risk management that requires a more comprehensive analysis. Our approach to LDI, which we more appropriately label commitment driven investing, recognises this reality in a more robust framework. By focusing on a plan's true objectives, maximising the safety of benefits and minimising the cost of funding those benefits, and by presenting allocation alternatives in the context of cost versus risk, CDI provides decision makers with vital information in the most straightforward way possible.

■ *Julia Bonafede is senior managing director and Steven Foresti and Dimitry Mindlin are both managing directors at Wilshire Associates in the US*

● This article is an abridged version of the Wilshire Consulting paper *Commitment-Driven Investing (CDI): LDI as we "C" it.* For the full paper, go to www.wilshire.com/BusinessUnits/Consulting/Investment/CDI_LDI.pdf