A dynamic sponsor-aware LDI framework for de-risking corporate pension funds

Section 1: Introduction

The pension crisis has brought defined benefit (DB) pension plans under increased scrutiny. There is clear evidence [1] that pure asset maximization strategies result in excess funded status volatility during periods of market stress. Liability-Driven Investing (LDI), on the other hand, has presented itself as a promising alternative. Defined-benefit plans are pension plans that must meet predefined benefit payments to retirees. LDI addresses this objective by focusing on liability risk management [2]. One popular LDI strategy is the funding ratio glide path, which shifts a pension's asset allocation into less risky portfolios as funded status milestones are reached. Although intuitively attractive, this approach suffers from the drawback that de-risking points are often rigidly defined. In this paper, we incorporate the dynamic de-risking features of glide paths into a portfolio optimization framework that does not rely on predefined de-risking triggers. We further enhance our strategy to incorporate the risk of the plan sponsor in portfolio allocation.

In the absence of liabilities (and additional assumptions on investor utility and return distributions), one may achieve optimal asset allocation by applying the mean-variance optimization of Markowitz [3,4]. However, in the presence of a stream of liabilities, direct application of the mean-variance portfolio theory is no longer appropriate, as it does not take into account the risk of becoming underfunded. One widely employed strategy for LDI is the surplus management approach (Sharpe and Tint [5]), which defines surplus as the value of assets net of liabilities, and conducts mean-variance optimization over the surplus return. This approach incorporates the correlation between the asset and the liability into the objective function. As Ang et al. point out in [6], a drawback to the Sharpe-Tint method is that it handles the downside (underfunding) and upside (overfunding) risks in a symmetric fashion.

In practice, the upside risk is of little concern to pension sponsors (due to unfavorable taxes on withdrawing surplus assets) while the focus is mostly on the shortfalls. To address this issue, Ang et al. proposed a variant of the Sharpe and Tint model that only penalizes liability deviation on the downside. They modeled the shortfall as a put option whose end-of-period payout is given by the discrepancy between future liability and asset value, or zero in the case of a surplus. Mathematically, this payout is given at the end of the period by $\max(L_1 - A_1, 0)$. The Ang model obtains portfolio weights through maximization of a modified objective function:

$$\arg\max_w \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) - \frac{c}{A_0} P(w, L_0, A_0),$$

where w represents the equity weight, $r_A = wr_E + (1-w)r_B$ represents returns on pension assets, r_E and r_B are equity and bond assets which make up the asset portfolio, and $P(w, L_0, A_0)$ is the value of the put option at funding level A_0/L_0. In the above, λ represents the risk aversion coefficient in the classical mean-variance setting, and c is an additional risk parameter that can be tailored to the plan sponsor's preferences. The attractiveness of this model is that it produces endogenous risk aversion characteristics such that the optimal equity weight changes according to funded ratio level. As the funded ratio increases towards 100%, the non-intrinsic value of the put option also increases and leads to higher effective risk aversion relative to liabilities. Thus, this optimization problem includes a policy similar to industry glide path offerings, where tracking error to liabilities decreases as funded ratio increases.

We develop an enhanced LDI methodology by including a risk aversion parameter that varies with time, and a term that penalizes the co-movement of the equity returns of the plan sponsor with plan asset returns. The additional covariance penalty term allows our solution to account for the implicit short put position on the funded ratio that a plan sponsor assumes when promising to pay its beneficiaries. This put position has taken on an increased importance given the Pension Protection Act of 2006 (PPA) [7] where underfunded pensions are given a 7-year horizon to reach 100% funding. Given that the put option value is countercyclical (as funded ratios typically decline in recessions due to negative returns of equity assets and falling interest rates), the co-movement with plan assets of the sponsor should be a key determinant in the asset allocation of a particular plan.

The time-varying aspect of our coefficient allows us to set a desired time horizon to reach full funding and vary the asset allocation based on time elapsed. We demonstrate that our proposed product improves the distribution of positive funding outcomes, and does so with less funded ratio risk relative to alternative formulations. We establish this using both historical backtests and Monte Carlo simulations along with appropriate statistical tests.

The remainder of the exposition is arranged as follows: In Section 2, we outline the specifications and methodology behind our product. Section 3 presents and discusses our results, both quantitatively and qualitatively, and Section 4 concludes.

Section 2: Specifications

Our model produces portfolio weights as the output of the optimization

$$\arg\max_{w} \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) - cov(r_A, r_S) - \frac{c(\tau)}{A_0} P(w, L_0, A_0),$$

In our objective function, $cov(r_s, r_A)$ represents the covariance of the equity returns of the plan sponsor with that of the portfolio assets. Recall that the Capital Asset Pricing Model (CAPM) states that $r_s = r_f + \beta(r_E - r_f) + \epsilon$. We use this to isolate the systematic component of the covariance and simplify the corresponding penalty term. This term enables us to penalize situations when both the plan sponsor and the plan assets are under duress. Additionally, we wish to penalize a firm that possesses more systematic risk than the market, so we set the penalty term to be centered around the $\beta = 1$ level. That is to say, for a $\beta = 1$ firm, there is no additional penalty, but larger (smaller) β will, all else equal, adopt less (more) equity-heavy asset allocations. Thus, we replace the covariance term $cov(r_A, r_S)$ by the following expression:

$$(\beta - 1)cov(r_A, r_E) = (\beta - 1)(w\sigma_E^2 + (1-w)\sigma_E\sigma_B\rho_{EB})$$

The overall effect of this term minimizes the likelihood that a plan sponsor will need to make contributions when it is most costly to the firm, i.e. when firm assets as well as plan assets are depressed in value.

Additionally, we have made $c(\tau)$ a time-varying risk aversion parameter for the shortfall option. The desired behavior is captured by the functional form:

$$c(\tau) = \begin{cases} c_0, & \frac{A_0}{L_0} \geq 1 \\ c_0 \cdot \frac{\tau}{T}, & \frac{A_0}{L_0} < 1 \end{cases}$$

where $\tau = T - t$ is the time remaining, T is the pre-specified time horizon, and t is the elapsed time.

The risk aversion coefficient attached to the shortfall put decreases as time passes at a constant rate. This implies that the optimal holdings for underfunded plans approach the mean-variance weights as the end of the investment horizon nears. The reasoning for taking on this additional risk stems from the fact that if the plan is underfunded with little time left, a more aggressive asset allocation will be required in order to meet goals. In other words, those plan sponsors facing significant penalties should be less averse to shortfalls as the time left to reach full funding dwindles.

The time-varying component of risk aversion in our model may appear to undertake additional portfolio risk near a seemingly arbitrary deadline. However, statutory demands of the PPA necessitate reaching full funding in a 7 year time period, with the sponsor being required to contribute any shortfall at the end of this period. Given these requirements, we believe it is reasonable to manage our strategy around a given time horizon.

To solve the optimization problem, we assume the following dynamics for the equity and the bond respectively:

$$\frac{dE}{E} = \mu_E dt + \sigma_E dW_t^E, \qquad \frac{dB}{B} = \mu_B dt + \sigma_B dW_t^B, \qquad dW_t^E dW_t^B = \rho \mathrm{dt}.$$

The value of the asset portfolio at end of a single period is determined by weights to stocks, w, and bonds, (1 - w), chosen at time 0:

$$A_1 = wA_0 \exp\left(\left(\mu_E - \frac{\sigma_E^2}{2}\right) + \sigma_E \mathrm{W}_1^E\right) + (1-w)A_0 \exp\left(\left(\mu_B - \frac{\sigma_B^2}{2}\right) + \sigma_B W_1^B\right).$$

The dynamics of the liabilities are also assumed to be lognormal with drift μ_L and standard deviation σ_L (we also assume correlation between assets and liabilities):

$$\frac{dL}{L} = \mu_L dt + \sigma_L dW_t^L \quad , \qquad L_1 = L_0 \exp\left(\left(\mu_L - \frac{\sigma_L^2}{2}\right) + \sigma_L W_1^L\right).$$

Valuation of the shortfall put option presents a challenge since the quantity $L_1 - A_1$ is not lognormally distributed. Although it is trivial to compute the price using Monte Carlo simulation, the resulting penalty (as a function of w and A_0) will not be smooth enough to optimize using any method other than computationally expensive brute-force grid searches or simulation-based methods. Objective functions that are too rough or jagged will appear to have multiple local optima and will cause optimization routines to terminate prematurely at sub-optimal points. Thus, as in Ang et al., we employ the approximation frameworks set forth by Margrabe [8] and Venkatramanan & Alexander [9]. The payout of the shortfall is represented as that of a spread option, which is then approximated by a compound exchange option for which there is a closed-form valuation. This allows us to construct an objective function that is sufficiently smooth for use with efficient optimization routines; in this work, we use the interior point method.

Dynamic Strategy and Rebalancing Policy

Given the asset allocation methodology outlined above, we apply our framework over a pre-specified time horizon. We calculate optimal portfolio weights at annual intervals and implement the following dynamic strategy:

1. Determine starting asset allocation based on an initial optimization.
2. At the end of the following month, calculate the change in funded ratio based on the respective returns of assets and liabilities.
3. If the end-of-month funded ratio drops below a minimum specified level, assume the plan sponsor contributes to bring the funded ratio back to minimum threshold. We measure the cumulative contribution amount for benchmarking purposes.
4. Repeat steps 2 and 3, calculating new optimal holdings only at annual intervals.

We make the assumption of a minimum funded ratio barrier for two reasons. First, we would like to maximize the probability of success: at extremely low funding levels, the amount of growth needed to reach full funding becomes practically insurmountable. The second reason is that there may be statutory or plan sponsor-related minimum funding requirements; in the case of U.S. law, the PPA requires plans deemed "at-risk" to have more stringent shortfall covering requirements.

Beyond the parameters required in our optimization procedure, the following additional parameters are required for our strategy:

- $\boldsymbol{F_0}$: Initial funded ratio
- $\boldsymbol{F_{MIN} < F_0}$: Minimum permissible funded ratio barrier (either by statute or sponsor policy)
- $\boldsymbol{T}$: Target horizon to reach full funding (generally 5-10 years given PPA constraints)
- $\boldsymbol{\beta}$: CAPM Beta of the corporate pension sponsor (used as a proxy for firm/asset systematic risk)
- $\boldsymbol{\lambda, c_0, k}$: Risk aversion parameter constants

A graphical presentation of a simulated path of the funded ratio for a 10-year horizon is shown below.

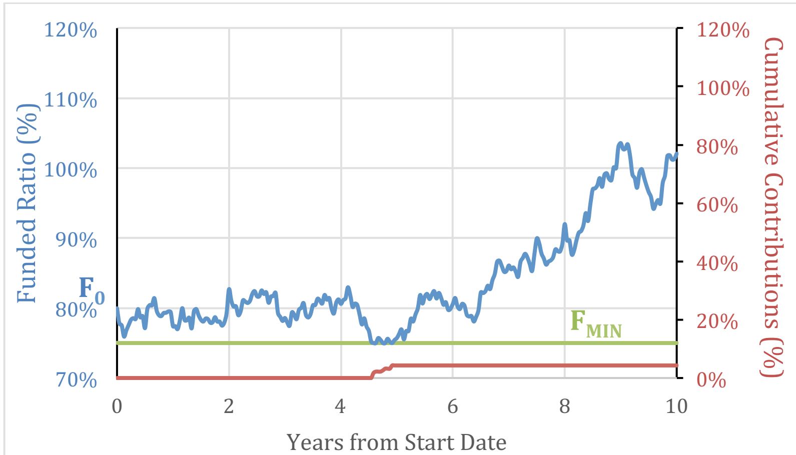

Section 3: Historical Backtest and Simulation Results

In determining the efficacy of our strategy relative to other approaches, we employ both historical backtesting and Monte Carlo simulation. We identify several key criteria that can be used to quantify the relative success of our approach:

- High terminal funded ratio
- Low funding ratio volatility
- Small probability of underfunding at completion
- Minimal cumulative sponsor contributions
- Moderate portfolio turnover

As proxies for our asset and liability processes, we select the following indexes for use in testing: the S&P 500 Index for equities, the Barclays Capital Long Corporate Index for fixed income, and the Citigroup Pension Liability Index to represent the liability. Our motivation for using a long corporate index instead of the more popular Barclays Aggregate is that we want a high duration fixed income asset in order to hedge the even higher duration liability index as much as possible. We assume risk, return and correlation characteristics for our assets given by the 2015 JP Morgan Asset Management Long-Term Capital Markets Assumptions [10]. These forward-looking, "through-the-cycle" values are intended to estimate risk and returns for asset allocators. For liabilities, we assume excess return and risk relative to a JP Morgan long duration treasury index assumptions. However, in modeling the correlation of liabilities to assets, we assume historically realized correlations using CPLI data.

	Annual Return	**Annual Volatility**
S&P 500 Index (r_E)	7.5%	14.75%
Barclays Long Credit Index (r_B)	5%	9.75%
Liabilities (CPLI) (r_L)	5.5%	12.50%

Correlations	S&P 500	Long Credit	CPLI
S&P 500 Index	1	0.25	0.2
Long Credit Index	0.25	1	0.98
Liabilities (CPLI)	0.2	0.98	1

We use these assumptions during all periods in both simulations and backtests. To benchmark our strategy against alternative procedures, we compare our strategy against the following three alternatives: mean-variance portfolio optimization (MVO), Sharpe

Tint surplus optimization (ST), and the Ang et al. framework (Ang). The objective functions from the respective procedures are reproduced below.

$$\boldsymbol{MVO}\text{: } \arg\max_w \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) \qquad \boldsymbol{ST}\text{: } \arg\max_w \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) - \frac{\lambda}{2}\frac{L_0}{A_0} cov(r_A, r_L)$$

$$\boldsymbol{Ang}\text{: } \arg\max_w \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) - \frac{c}{A_0} P(w, L_0, A_0)$$

$$\boldsymbol{Our\ Proposal}\text{: } \arg\max_w \mathrm{E}(r_A) - \frac{\lambda}{2} var(r_A) - (\beta - 1) \cdot cov(r_A, r_E) - \frac{c(\tau)}{A_0} P(w, L_0, A_0)$$

We use the following parameter assumptions for our simulations and backtests:

Parameter	**Value**
Initial Funded Ratio	85%
Min Funded Ratio Barrier	75%
Investment Horizon	10 years
Sponsor CAPM Beta	1.0
λ, c_0	4.0, 2.0

Our choice of $\boldsymbol{\lambda = 4.0}$ results in mean-variance weights of 51% to equities and 49% to fixed income, weights that represent a reasonable asset allocation for a non-LDI pension plan. Our choice of $\boldsymbol{c_0 = 2.0}$ represents significant shortfall risk aversion and is sourced from the Ang et al. paper.

Monte-Carlo Simulation

Below we display a histogram and cumulative distribution function of ending funded ratio outcomes across 100,000 simulations of each methodology. We use the same set of random numbers as inputs, thus allowing us to compare outcomes given similar inputs.

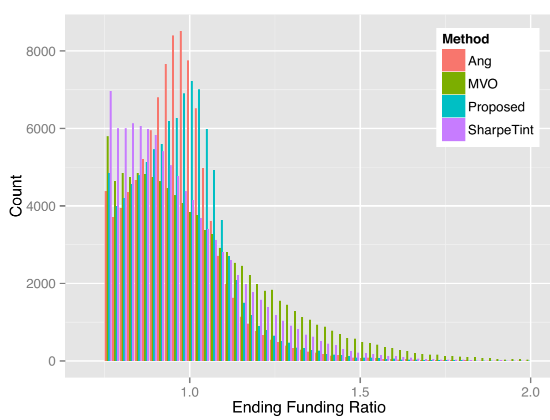

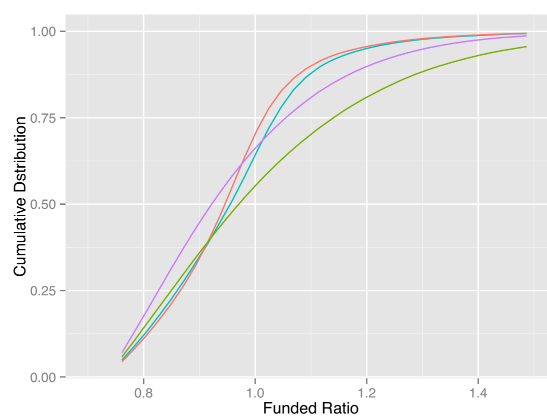

The blue items represent results using our approach. Relative to the Ang et al. approach, our methodology shifts the distribution of funded ratio to the right, with more mass occurring at outcomes at or above full funding. This is a result of our time-varying shortfall aversion parameter that we specify, where underfunded plans will re-risk as time decreases. Relative to the mean-variance and Sharpe and Tint approaches, our strategy reduces mass at both tails, placing more emphasis on outcomes near full funding as we set out to do.

We display summary statistics of the results of our simulations below. Note that since the mean-variance weights do not change, the corresponding turnover is 0.

	MVO	ST	Ang	Proposed ($\beta = 1$)
	Ending Funded Ratio			
Mean	1.0333	0.9726	0.9644	0.9720
SD	0.2281	0.1772	0.1330	0.1404
	Underfunded at Ending Time Horizon			
Mean	0.5333	0.6418	0.6651	0.6053
SD	0.0016	0.0015	0.0015	0.0015
	Volatility of Funded Ratio			
Mean	0.0944	0.0712	0.0627	0.0656
SD	0.0573	0.0420	0.0300	0.0309
	Cumulative Contribution			
Mean	0.0869	0.0570	0.0367	0.0380
SD	0.1567	0.1168	0.0691	0.0702
	Annualized Portfolio Turnover			
Mean	0	0.0065	0.0651	0.0675
SD	0	0.0020	0.0206	0.0242

Given that our methodology builds upon the Ang et al. framework, we conducted statistical tests to compare the outcomes. We compare differences in ending funded ratio using paired two-sample t-tests, and differences in probability of end time underfunding using McNemar's chi-squared test for correlated samples [11]. The results and interpretation of our tests are displayed below.

	Our Approach vs. Ang	Interpretation
Ending Funded Ratio (A/L)*	121.22***	Mean of our ending A/L is higher
Underfunded at end time**	5192.97***	Probability of ending underfunded is lower for our method
Volatility of Funded Ratio*	267.95***	Our funded ratio volatility is higher
Cumulative Contribution*	118.19***	We have higher cumulative contributions

*Differences are tested by paired two-sample t tests
** Differences are tested by McNemar's chi-squared test
*** p-value < 0.01

Overall, relative to the Ang et al. method, we improve funding ratio outcomes at the cost of additional volatility and contributions. This is a conscious decision on our part to make our strategy add equity risk as we approach the time horizon if the plan is underfunded. The assumed beta of 1.0 implies a moderate tolerance to making contributions; below we will investigate the effect of varying the beta parameter.

Historical Backtest

We complement our Monte Carlo simulations by conducting a backtesting procedure using historical asset and liability returns. While the Monte Carlo simulation procedure allows us to compare distributions of outcomes, it comes at the cost of a strong assumption of normally distributed returns given constant parameters. We thus adapt our optimization framework to historical data, using monthly realized returns for the S&P 500, Barclays Long Credit Index, and Citigroup Pension Liability Index. Backtest results are calculated across rolling 10-year window horizons using data from January 1995 to January 2015.

Across the 120 (overlapping) ten year windows in which we backtest, we generate a histogram of ending funded ratios as well as a cumulative distribution function. The cumulative distribution function suggests that our methodology results in a higher proportion of outcomes that occur at higher ending funding levels. Our strategy accomplishes this with a modest 8.14% annualized turnover, as compared to 11.03% using the Ang method. Across methodologies, funded ratio goals are typically not met in the historical backtest, which is due to a combination of mixed equity returns and a large compression in long-term rates over the sample, which greatly boosted CPLI performance. Overall, the backtest results support our methodology over alternatives, and do so in a challenging investment environment, both in regard to asset performance and liability growth.

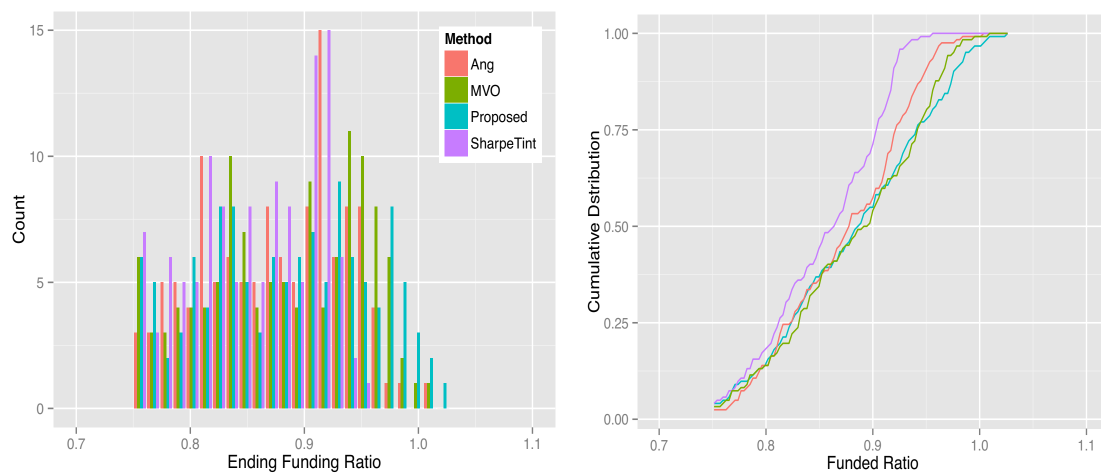

Effects of Differing CAPM Beta Assumptions

Having compared our methodology to alternative approaches, we now illustrate the effects of varying the CAPM beta within our strategy. We run 100,000 simulations using three CAPM beta assumptions of $\beta = 0.5, \beta = 1.0, \beta = 1.5$, and compute both sample statistics in addition to a Box-and-Whisker plot showing the distribution of cumulative contributions for each assumption. Based on our model design, we expect to decrease required contributions for higher beta firms.

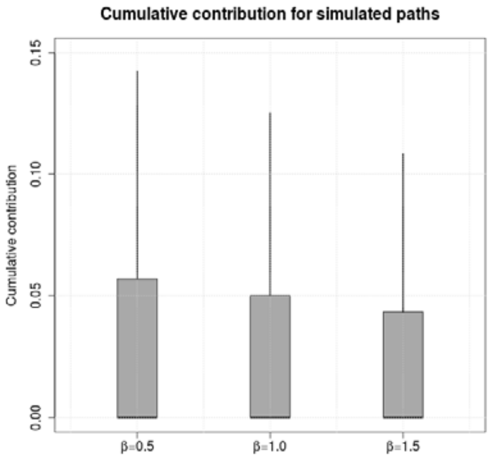

Proposal	*ß* = 0.5	*ß* = 1	*ß* = 1.5
	Ending Funded Ratio		
Mean	0.9830	0.9720	0.9611
SD	0.1485	0.1404	0.1330
	Underfunded at Ending Time Horizon		
Mean	0.5693	0.6053	0.6452
SD	0.0016	0.0015	0.0015
	Volatility of Funded Ratio		
Mean	0.0693	0.0656	0.0620
SD	0.0333	0.0309	0.0288
	Cumulative Contribution		
Mean	0.0419	0.0380	0.0341
SD	0.0759	0.0702	0.0646
	Annualized Portfolio Turnover		
Mean	0.0766	0.0675	0.0592
SD	0.0254	0.0242	0.0224

The edge of the grey box represents the 75^{th} percentile cumulative contribution outcome and the end of the black line represents the simulation with the highest cumulative contribution amount. As expected, our plot shows that as CAPM beta increases, the distribution of cumulative contributions narrows significantly. Thus, as beta increases, the expected amount of contributions that will need to be made decreases monotonically. This is exactly the behavior we set out to achieve through inclusion of the sponsor covariance term.

Looking at ended funded ratios, we see that the average ending funded ratio also decreases monotonically as beta increases. This is a consequence of having less equity risk all else equal as sponsor beta increases. Higher sponsor beta also leads to a decrease in funded status volatility in our simulations. This can be interpreted as implying a lower cost to a hypothetical put option that insures the sponsor against funding shortfalls. Although not pursued in this study, the cost to the sponsor can be quantified if we were to implement a dynamic hedging strategy for downside protection.

Section 4: Conclusion

We put forth a model that effectively de-risks a corporate pension plan relative to liabilities while still prioritizing full-funding goals. Our model builds upon the work of Ang et al., yet enhances the framework in ways that improve the distribution of outcomes. Our improvements focus on two key areas: incorporating the sponsor's systematic risk in asset allocation, and shortfall risk aversion as a function of time horizon. From the plan sponsor's perspective, our deliberate modifications to the optimization framework minimize the chance of shortfall and the associated costs to the firm sponsor when it is obligated to contribute.

Through simulation and testing, we show that our method results in a higher proportion of outcomes most favorable to plan sponsors while reducing outcomes that are either grossly overfunded or underfunded. We also apply our methodology to actual historical data where we find similarly positive results. Lastly, we illustrate the flexibility of our strategy to different plan sponsor systematic risk levels, and show how our framework minimizes contributions by those sponsors least able to contribute in market stress scenarios.

Finally, it is worth noting that our framework is flexible across alternative assumptions. Our single period framework implicitly assumes that very near term cash payouts are managed separately, and we also do not assume any ability to forecast time-varying asset returns. However, these additional conditions can be adopted to form a multi-period optimization problem. End users can supply custom return generating and liability-hedging portfolios with time-varying capital market assumptions. The framework is also adaptable to different initial funding ratio and time horizon assumptions.

References

1. Corporate Pension Funding Study, http://us.milliman.com/Solutions/Products/Corporate-Pension-Funding-Study/

2. Amenc, N., F. Goltz, L. Martellini, and V. Milhau, 2010, New frontiers in benchmarking and liability-driven investing, EDHEC-Risk Institute.

3. Markowitz, H., 1952, Portfolio Selection, *Journal of Finance*, 7, 77-91.

4. Markowitz, H., 1959, Portfolio selection: efficient diversification of investments. 1959. *Basil Blackwall*, New York.

5. Sharpe, W. F., and L. G. Tint, 1990, Liabilities – A New Approach, *Journal of Portfolio Management*, Winter, 5-10.

6. Ang, A., B. Chen and S. Sundaresan, 2013, Liability Driven Investment with Downside Risk, *Journal of Portfolio Management*, 40, 71-87.

7. U.S. Code 430. Minimum funding standards for single-employer defined benefit pension plans. http://www.law.cornell.edu/uscode/text/26/430

8. Margrabe, W., 1978, The Value of an Option to Exchange one Asset for Another, *Journal of Finance*, 33, 177-86.

9. Venkatramanan, A., and Alexander,C., 2011, Closed Form Approximation for Spread Options, *Applied Mathematical Finance*, 18, 5, 447-472.

10. JP Morgan Long-Term Capital Market Return Assumptions 2015, https://am.jpmorgan.com/es/institutional/ltcmra

11. McNemar, Q., 1947, Note on the sampling error of the difference between correlated proportions or percentages, *Psychometrika* 12 (2): 153-157.