Journal Of Investment Management, Vol. 23, No. 2, (2025), pp. 33–45

JOIM
www.joim.com

MARKOWITZ WEALTH MANAGEMENT TO PENSION PLANS

Sanjiv Das*

Transcript of a talk presented at the Spring JOIM Conference honoring Harry M. Markowitz on March 24–26, 2024, at the Rady School of Management, UCSD.

To begin, I'll mention a little bit of history, then I want to segue into a new area of work that I started with Harry Markowitz, and I continue to work on till today. I began working with Harry close to two decades ago and my first introduction to Harry was actually by Gifford who sat me down next to Harry at dinner one day. We spent the next hour or so talking about computer science, which was so much fun, and he spoke to me a lot about Simscript, which you've heard mentioned already today. I went back and read the Simscript paper and it's a fascinating paper because there have been lots of related developments since then. In many ways Simscript was a forerunner to the simulations of environments that people are using for reinforcement learning today. You're basically running a simulation environment in which agents react to a changing landscape (i.e., a state space) in real time. And there's an open-source library called "gymnasium" (originally "gym") developed by

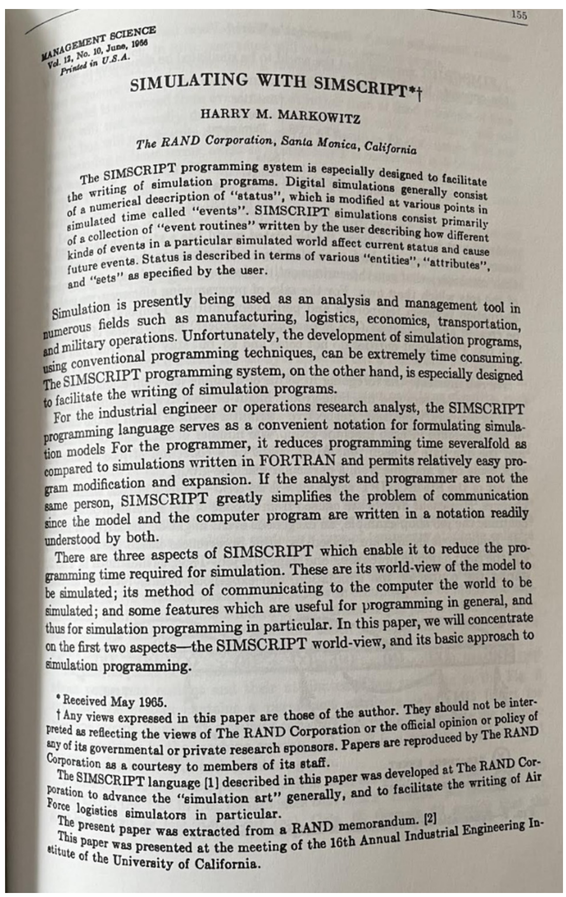

155

MANAGEMENT SCIENCE
Vol. 12, No. 10, June, 1966
Printed in U.S.A.

SIMULATING WITH SIMSCRIPT*†

HARRY M. MARKOWITZ

The RAND Corporation, Santa Monica, California

The SIMSCRIPT programming system is especially designed to facilitate the writing of simulation programs. Digital simulations generally consist of a numerical description of "status", which is modified at various points in simulated time called "events". SIMSCRIPT simulations consist primarily of a collection of "event routines" written by the user describing how different kinds of events in a particular simulated world affect current status and cause future events. Status is described in terms of various "entities", "attributes", and "sets" as specified by the user.

Simulation is presently being used as an analysis and management tool in numerous fields such as manufacturing, logistics, economics, transportation, and military operations. Unfortunately, the development of simulation programs, using conventional programming techniques, can be extremely time consuming. The SIMSCRIPT programming system, on the other hand, is especially designed to facilitate the writing of simulation programs.

For the industrial engineer or operations research analyst, the SIMSCRIPT programming language serves as a convenient notation for formulating simulation models For the programmer, it reduces programming time severalfold as compared to simulations written in FORTRAN and permits relatively easy program modification and expansion. If the analyst and programmer are not the same person, SIMSCRIPT greatly simplifies the problem of communication since the model and the computer program are written in a notation readily understood by both.

There are three aspects of SIMSCRIPT which enable it to reduce the programming time required for simulation. These are its world-view of the model to be simulated; its method of communicating to the computer the world to be simulated; and some features which are useful for programming in general, and thus for simulation programming in particular. In this paper, we will concentrate on the first two aspects—the SIMSCRIPT world-view, and its basic approach to simulation programming.

* Received May 1965.

† Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

The SIMSCRIPT language [1] described in this paper was developed at The RAND Corporation to advance the "simulation art" generally, and to facilitate the writing of Air Force logistics simulators in particular.

The present paper was extracted from a RAND memorandum. [2]

This paper was presented at the meeting of the 16th Annual Industrial Engineering Institute of the University of California.

*I am grateful to my co-authors on their collaborations that are summarized in the work in this presentation: Aviva Casanova, Daniel Ostrov, Anand Radhakrishnan, Deep Srivastav, Jonathan Scheid, Meir Statman, Wylie Tollette, and of course Harry Markowitz.

OpenAI that pretty much does what Simscript does, and of course a little bit more given the needs of reinforcement learning. Harry was therefore already talking about things that are important today in reinforcement learning, which is now used in so many different ways.

The second interaction I had with Harry was a few years later when Harry, Meir Statman, and I ended up on the investment committee of a wealth management company and through that process there was this back-and-forth debate, as you can well imagine, between Harry the theorist, and Statman the behaviorist, and the twain never could meet, you know. There was fundamental disagreement about many things. But eventually that slowly dissolved into a nice consensus. And along the way Harry gave me this paper to read. If you want to learn about what influenced his thinking, reading this paper titled "Trains of Thought" would be insightful.

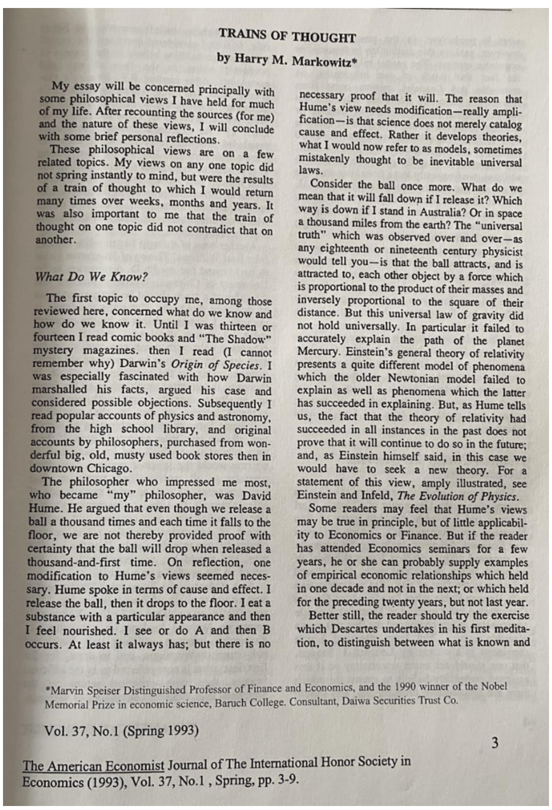

TRAINS OF THOUGHT

by Harry M. Markowitz*

My essay will be concerned principally with some philosophical views I have held for much of my life. After recounting the sources (for me) and the nature of these views, I will conclude with some brief personal reflections.

These philosophical views are on a few related topics. My views on any one topic did not spring instantly to mind, but were the results of a train of thought to which I would return many times over weeks, months and years. It was also important to me that the train of thought on one topic did not contradict that on another.

What Do We Know?

The first topic to occupy me, among those reviewed here, concerned what do we know and how do we know it. Until I was thirteen or fourteen I read comic books and "The Shadow" mystery magazines. then I read (I cannot remember why) Darwin's *Origin of Species*. I was especially fascinated with how Darwin marshalled his facts, argued his case and considered possible objections. Subsequently I read popular accounts of physics and astronomy, from the high school library, and original accounts by philosophers, purchased from wonderful big, old, musty used book stores then in downtown Chicago.

The philosopher who impressed me most, who became "my" philosopher, was David Hume. He argued that even though we release a ball a thousand times and each time it falls to the floor, we are not thereby provided proof with certainty that the ball will drop when released a thousand-and-first time. On reflection, one modification to Hume's views seemed necessary. Hume spoke in terms of cause and effect. I release the ball, then it drops to the floor. I eat a substance with a particular appearance and then I feel nourished. I see or do A and then B occurs. At least it always has; but there is no necessary proof that it will. The reason that Hume's view needs modification—really amplification—is that science does not merely catalog cause and effect. Rather it develops theories, what I would now refer to as models, sometimes mistakenly thought to be inevitable universal laws.

Consider the ball once more. What do we mean that it will fall down if I release it? Which way is down if I stand in Australia? Or in space a thousand miles from the earth? The "universal truth" which was observed over and over—as any eighteenth or nineteenth century physicist would tell you—is that the ball attracts, and is attracted to, each other object by a force which is proportional to the product of their masses and inversely proportional to the square of their distance. But this universal law of gravity did not hold universally. In particular it failed to accurately explain the path of the planet Mercury. Einstein's general theory of relativity presents a quite different model of phenomena which the older Newtonian model failed to explain as well as phenomena which the latter has succeeded in explaining. But, as Hume tells us, the fact that the theory of relativity had succeeded in all instances in the past does not prove that it will continue to do so in the future; and, as Einstein himself said, in this case we would have to seek a new theory. For a statement of this view, amply illustrated, see Einstein and Infeld, *The Evolution of Physics*.

Some readers may feel that Hume's views may be true in principle, but of little applicability to Economics or Finance. But if the reader has attended Economics seminars for a few years, he or she can probably supply examples of empirical economic relationships which held in one decade and not in the next; or which held for the preceding twenty years, but not last year.

Better still, the reader should try the exercise which Descartes undertakes in his first meditation, to distinguish between what is known and

*Marvin Speiser Distinguished Professor of Finance and Economics, and the 1990 winner of the Nobel Memorial Prize in economic science, Baruch College. Consultant, Daiwa Securities Trust Co.

Vol. 37, No.1 (Spring 1993)

3

The American Economist Journal of The International Honor Society in Economics (1993), Vol. 37, No.1 , Spring, pp. 3-9.

It's a short paper about how modern philosophers actually influenced Harry's own thinking and it's worth reading many times. I thoroughly enjoyed this paper, and I strongly recommend it for anybody who wants to look at influences on Harry—I think it will change your own thinking too.

Eventually the viewpoints of the behaviorist (Statman) and the rational optimizer (Harry) ended up in an amalgamation, which was actually quite mathematical and simple, and I'll talk a little bit about that today as well. When we got there, Harry sent me his book as we worked together to finalize these ideas. It's a pretty good deal I had over the years—I would send Harry my papers, and he would send me his books.

We ended up writing this paper for optimization of mental accounts and another follow up paper on wealth management. The main result in the paper "Portfolio Optimization with Mental Accounts" (JFQA 2010) is that it shows how to map mental accounting structure from behavioral portfolio

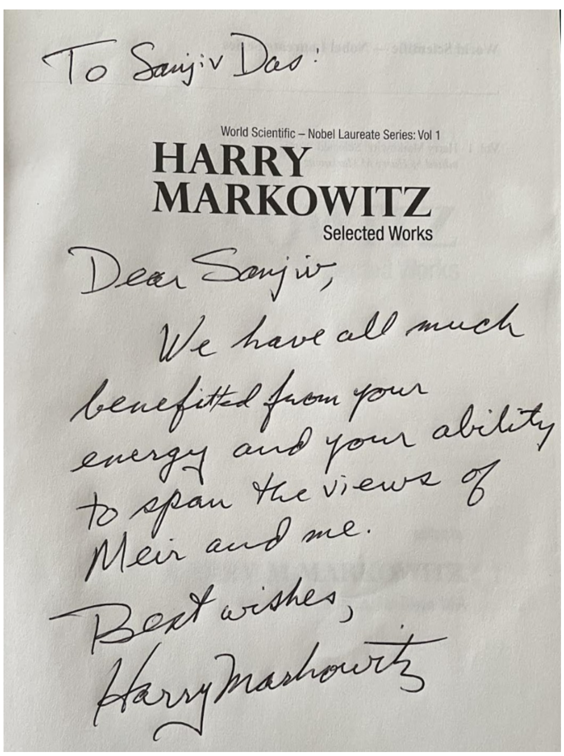

theory to mean variance theory. And thus, we were able to look at goals-based wealth management in the light of mean-variance theory. I'll talk a little bit about both these things today and then try and extend these ideas to what they mean for pension management. The two papers with Harry on the subject are shown on the slide below:

Here's a related series of papers with many co-authors that followed the work with Harry, and

Connecting Mean-Variance Optimization with Behavioral Optimization

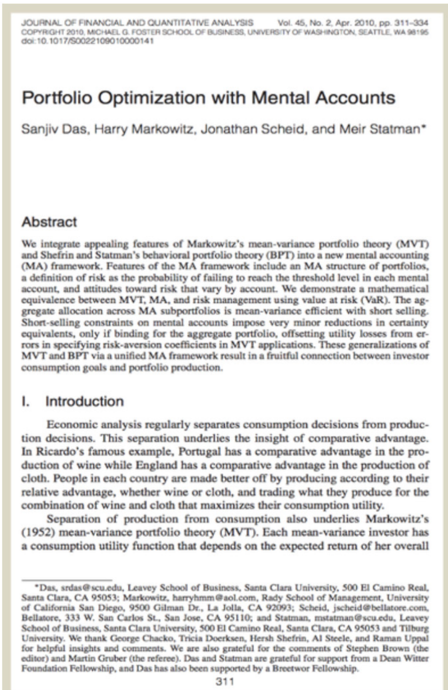
JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 45, No. 2, Apr. 2010, pp. 311–334
COPYRIGHT 2010, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109010000141

Portfolio Optimization with Mental Accounts

Sanjiv Das, Harry Markowitz, Jonathan Scheid, and Meir Statman*

Abstract

We integrate appealing features of Markowitz's mean-variance portfolio theory (MVT) and Shefrin and Statman's behavioral portfolio theory (BPT) into a new mental accounting (MA) framework. Features of the MA framework include an MA structure of portfolios, a definition of risk as the probability of failing to reach the threshold level in each mental account, and attitudes toward risk that vary by account. We demonstrate a mathematical equivalence between MVT, MA, and risk management using value at risk (VaR). The aggregate allocation across MA subportfolios is mean-variance efficient with short selling. Short-selling constraints on mental accounts impose very minor reductions in certainty equivalents, only if binding for the aggregate portfolio, offsetting utility losses from errors in specifying risk-aversion coefficients in MVT applications. These generalizations of MVT and BPT via a unified MA framework result in a fruitful connection between investor consumption goals and portfolio production.

I. Introduction

Economic analysis regularly separates consumption decisions from production decisions. This separation underlies the insight of comparative advantage. In Ricardo's famous example, Portugal has a comparative advantage in the production of wine while England has a comparative advantage in the production of cloth. People in each country are made better off by producing according to their relative advantage, whether wine or cloth, and trading what they produce for the combination of wine and cloth that maximizes their consumption utility.

Separation of production from consumption also underlies Markowitz's (1952) mean-variance portfolio theory (MVT). Each mean-variance investor has a consumption utility function that depends on the expected return of her overall

*Das, srdas@scu.edu, Leavey School of Business, Santa Clara University, 500 El Camino Real, Santa Clara, CA 95053; Markowitz, harryhmm@aol.com, Rady School of Management, University of California San Diego, 9500 Gilman Dr., La Jolla, CA 92093; Scheid, jscheid@bellatore.com, Bellatore, 333 W. San Carlos St., San Jose, CA 95110; and Statman, mstatman@scu.edu, Leavey School of Business, Santa Clara University, 500 El Camino Real, Santa Clara, CA 95053 and Tilburg University. We thank George Chacko, Tricia Doerksen, Hersh Shefrin, Al Steele, and Raman Uppal for helpful insights and comments. We are also grateful for the comments of Stephen Brown (the editor) and Martin Gruber (the referee). Das and Statman are grateful for support from a Dean Witter Foundation Fellowship, and Das has also been supported by a Breetwor Fellowship.

311

Portfolios for Investors Who Want to Reach Their Goals While Staying on the Mean–Variance Efficient Frontier

Sanjiv Das, Harry Markowitz, Jonathan Scheid, Meir Statman
The Journal of Wealth Management Fall 2011, 14 (2) 25 - 31
DOI: 10.3905/jwm.2011.14.2.025

We can map MV mathematics to
(1) Mental accounts in behavioral portfolio theory
(2) Goal probabilities in goals-based wealth management (GBWM)

From MVT to GBWM

1. "Augmenting the Funded Ratio: New Metrics for Liability Based Plans" 2024, (Sanjiv Das, Daniel Ostrov, Anand Radhakrishnan, Deep Srivastav, Wylie Tollete) [PDF]

2. "Lifestyle, Longevity, and Legacy Risks with Annuities in Retirement Portfolio Decumulation" 2023, (Sanjiv Das, Daniel Ostrov, Anand Radhakrishnan, Deep Srivastav), *Journal of Wealth Management* [PDF]. *[Examines the tradeoff between consumption, leaving a legacy, modulated by life expectancy. Should we ever buy annuities?]*

3. "Efficient Goal Probabilities: A New Frontier" 2023, (Sanjiv Das, Daniel Ostrov, Anand Radhakrishnan, Deep Srivastav), *Journal of Investment Management*. [PDF]. *[How to generate Pareto-optimal goal probability frontiers in goals-based wealth management.]*

4. "Optimal Goals-Based Investment Strategies For Switching Between Bull and Bear Markets," (2022), (Sanjiv Das, Daniel Ostrov, Aviva Casanova, Anand Radhakrishnan, Deep Srivastav), *Journal of Wealth Management*, Spring. [PDF]. [JWM]. *[Goals-based wealth management with regime changes. Unless regime cognizance is high, a single regime approach may be better.]*

5. "Combining Investment and Tax Strategies for Optimizing Lifetime Solvency under Uncertain Returns and Mortality" (2021), (Sanjiv Das, Dan Ostrov, Aviva Casanova, Anand Radhakrishnan, Deep Srivastav), *Journal of Risk and Financial Management*, v14(7), 285. [Paper]. *[Goals-based wealth management combined with tax optimization across taxable and tax-deductible accounts, fully accounting for longevity risk.]*

6. "Dynamic Optimization for Multi-Goals Wealth Management." 2021 (Sanjiv Das, Dan Ostrov, Anand Radhakrishnan, Deep Srivastav), *Journal of Banking and Finance*, v140, 106192. [JBF], [DOI]. [PDF]. *[Goals-based wealth management with multiple, concurrent, and partial goals.]*

7. "Dynamic Portfolio Allocation in Goals-Based Wealth Management", (2020), (Sanjiv Das, Dan Ostrov, Anand Radhakrishnan, and Deep Srivastav), *Computational Management Science*, v17, June, 613-640, DOI:10.1007/s10287-019-00351-7. [PDF]. *[Goals-based wealth management implemented with dynamic programming outperforms target date funds.]*

8. "A New Approach to Goals-Based Wealth Management" (2018), (Sanjiv Das, Dan Ostrov, Anand Radhakrishnan, Deep Srivastav), *Journal of Investment Management*, v16(3), 1-27. [Paper]. *[New framework and mathematics for implementing goals-based wealth management, where risk is redefined as the likelihood of not meeting one's goals. Winner of the Harry Markowitz Award for Best Paper.]*

the one I'll get to eventually as the main paper for today is the one at the top (in red).

We are all familiar with the Markowitz mean-variance problem. You minimize variance of a portfolio for a given level of expected return. Below is a very small variation on that problem, where you maximize the mean portfolio return, minus some function of the risk (variance), and gamma (γ) is a risk aversion coefficient. The solution to this problem is easy to derive and is shown below.

If somebody could tell you what their gamma (γ) is, then solving this problem tells you where their portfolio should reside on the efficient frontier. The solution is a function of your risk aversion coefficient now, which you probably don't know.

MV Math with Risk Aversion

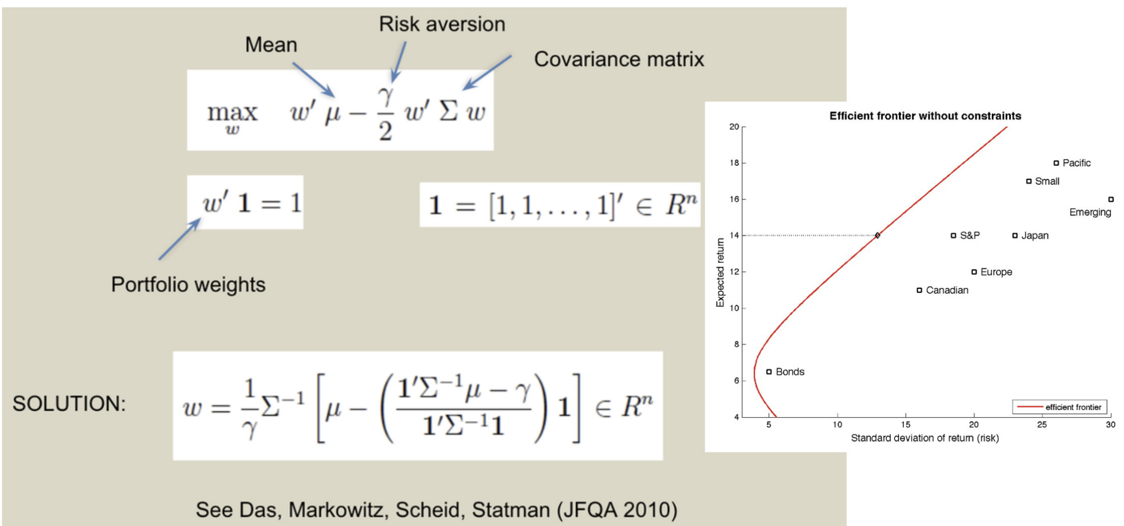

Connecting MVT with Behavioral Portfolio Theory (BPT)

The Behavioral Portfolio Theory (BPT) optimization is as follows:

Based on Shefrin and Statman (JFQA, 2000)

aspiration/greed — fear

$$\max_{\mathbf{w}} \mathbf{w}^\top \mu, \qquad s.t. \quad \text{Prob}[r \le H] \le \alpha$$

For normal returns r, the constraint may be stated explicitly as

$$H \le \mathbf{w}^\top \mu + \Phi^{-1}(\alpha)[\mathbf{w}^\top \mathbf{\Sigma} \mathbf{w}]^{1/2}$$

The easy way to solve for the optimal weights is to search for the value of γ such that the optimal $\mathbf{w}(\gamma)$ satisfies the optimality equation above:

$$\mathbf{w}(\gamma) = \frac{1}{\gamma}\Sigma^{-1}\left[\mu - \left(\frac{\mathbf{1}^\top\Sigma^{-1}\mu - \gamma}{\mathbf{1}^\top\Sigma^{-1}\mathbf{1}}\right)\mathbf{1}\right] \in \mathcal{R}^n$$

and satisfies the constraint:

$$H \le \mathbf{w}(\gamma)^\top \mu + \Phi^{-1}(\alpha)[\mathbf{w}(\gamma)^\top \mathbf{\Sigma} \mathbf{w}(\gamma)]^{1/2}$$

And the idea is to try to back it out. You actually need something with a behavioral view. And after a lot of debate, downside risk was something we all could agree on. We ended up with this formulation of greed versus fear, where you're maximizing the objective function subject to a downside constraint. The downside constraint aims to keep the probability of the portfolio's return (r) falling below a threshold (H) to no more than α, i.e., $Pr[r \le H] \le \alpha$.

Here, H is the threshold value and alpha (α) the probability that portfolio return (r) falls below the threshold. Now, if you expand the right-hand side ("fear") constraint equation, assuming some distribution for r, (it doesn't have to be normal, as in this case), you get the second equation above (the constraint). You can drop that equation into the solution we had from the previous problem, because we know that the portfolio weights $w(\gamma)$ are a function of your risk aversion. So if we substitute the last equation above into the model and make that an equality and drop $w(\gamma)$ in there, you can solve for γ (numerically, that is). So you provide H, which is the lower threshold of return you're willing to live with, and the probability (α) with which you do not want to go below the threshold return. The solution tells you which point to be on the efficient frontier. It also tells you your risk aversion γ. This complements the Markowitz solution by incorporating downside risk.

It turns out you can then do mental accounting because you can have different accounts with different thresholds H and probabilities α. For example, retirement, education, and bequest mental accounts. You don't know what your risk preferences are for each of those accounts, you could very well be very risky with your bequest portfolio, but very risk averse with your retirement portfolio. All you have to specify is what your downside risk is, and the probability with which you don't want to go below that. As shown here.

For example, in the retirement subportfolio below, I don't want to lose more than 10% with a probability of more than 5%. The solution is translated back into a risk aversion (γ) of 3.79. You can do the same thing for the education and bequest subportfolios. You can aggregate across all these three, because as long as you're allowing short selling, we know that a combination of portfolios on the efficient frontier will end up on the efficient frontier as well. Thus, mental account aggregation is also possible, and we did some experiments

Mental Account Subportfolio Scenarios

Risk aversion:	$\gamma = 3.7950$	$\gamma = 2.7063$	$\gamma = 0.8773$	60:20:20 mix
	Retirement Sub-portfolio	Education Sub-portfolio	Bequest Sub-portfolio	Aggregate Portfolio
Threshold (H)	Prob$[r < H]$	Prob$[r < H]$	Prob$[r < H]$	Prob$[r < H]$
-25.00%	0.00	0.01	0.15	0.03
-20.00%	0.01	0.03	0.17	0.05
-15.00%	0.02	0.05	0.20	0.08
-10.00%	0.05	0.09	0.23	0.12
-5.00%	0.11	0.15	0.26	0.18
0.00%	0.20	0.23	0.30	0.25
5.00%	0.34	0.33	0.33	0.33
10.00%	0.49	0.45	0.37	0.42
15.00%	0.65	0.57	0.41	0.52
20.00%	0.79	0.68	0.45	0.62
25.00%	0.89	0.78	0.49	0.71
Mean return	10.23%	12.18%	26.35%	13.84%
Std. deviation	12.30%	16.57%	49.13%	20.32%

Since all sub portfolios lie on the efficient frontier, the aggregate portfolio also does.

But if we choose long only portfolios, this breaks, but analysis shows that the slippage is minor.

to see if you impose short selling constraints, then how far away are you from the optimal? We compared aggregation of subportfolios with and without short selling and the difference turned out to be pretty small.

This is the first step, which is, we can map downside risk into mean variance theory and make sure we sit on the frontier while capturing the behavioral aspect of downside risk as well (and detecting one's γ). The second step is that we want to introduce goals-based wealth management (GBWM) into this framework. The basic notion of GBWM is diagrammed below.

You have an evolution equation like the one above, where you go from initial wealth $W(0)$ up to some wealth $W(t)$ at horizon t, and Z is a shock term. The mean growth is modulated by μ and the volatility parameter is σ. You can rearrange it so that you have μ on the left-hand side, everything else on the right-hand side. Say, in 10 years, I want my wealth to go from \$400,000 to \$500,000. That's my goal. And I want to do that with an 80% probability. The equation can be used to draw a locus of points μ, σ that will satisfy that particular equation for a Z corresponding to 80% under the normal distribution, in this case (see the plot on the bottom right). We call this the goal level probability curve (GLPC). If you increase that probability from 80% to 90%, you rotate the GLPC up to the left (see the bottom left plot below). Take your mean variance efficient frontier and superimpose it on the GLPCs and find the GLPC that is tangent to the efficient frontier (we will see this in ensuing plots).

The geometry of this problem can be solved in closed-form using constrained optimization.

We rearrange the earlier equation to have Z on the left-hand side, and Z is a proxy for the probability with which you meet your goal. You can maximize Z subject to the constraint that you remain on the efficient frontier. To maximize Z, we take this function for the efficient frontier above, and all you have to do is solve for a single variable, which is μ. Because once you have the μ, you can use the efficient frontier to get the corresponding σ as well. It turns out the solution comprises a cubic

Mapping MV Theory to GBWM

$$\tilde{W}(t) = W(0)e^{\left(\mu-\frac{\sigma^2}{2}\right)t+\sigma\sqrt{t}Z}$$

$$\mu = \frac{1}{2}\sigma^2 + \frac{z_0}{\sqrt{t}}\sigma + \frac{1}{t}\ln\left(\frac{W(t)}{W(0)}\right).$$

where z_0 is defined so that the Target Probability equals $\Phi(z_0)$, with $\Phi(z)$ being the cumulative distribution function (CDF) for a standard normal

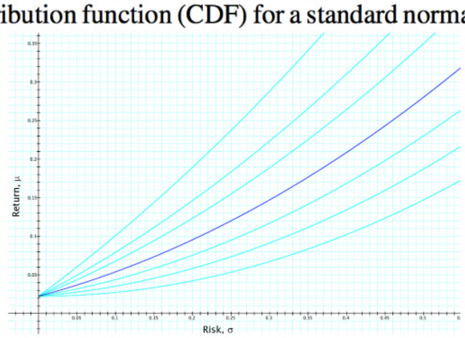

Example for Goals:
1) Investment Tenure: 10 years.
2) Initial Wealth: \$400,000
3) Target Wealth: \$500,000
4) Target Probability: 80%

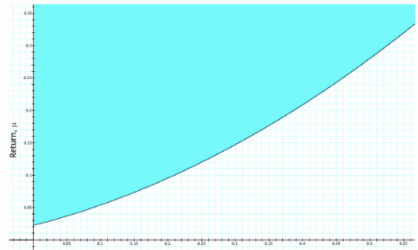

Figure 3: A Goal Region Example. The goal region (in light blue) in this figure corresponds to a base case where an investor has an Investment Tenure of 10 years, an Initial Wealth of \$400,000, a Target Wealth of \$500,000, and a Target Probability of 80%. The boundary of the goal region is the 80% Goal Probability Level Curve (GPLC), which is defined as the curve for which any (σ, μ) pair above the curve corresponds to a portfolio with an 80% or higher chance of the investor attaining their Target Wealth at the end of their investment tenure.

Optimizing Goal Probability -- The Cubic

Efficient frontier: $\sigma = \sqrt{a\mu^2 + b\mu + c}$,

$$\begin{aligned} a &= \mathbf{h}^\top \Sigma \mathbf{h} \\ b &= 2\mathbf{g}^\top \Sigma \mathbf{h} \\ c &= \mathbf{g}^\top \Sigma \mathbf{g}, \end{aligned} \qquad \begin{aligned} \mathbf{g} &= \frac{l\Sigma^{-1}\mathbf{O} - \mathbf{k}\Sigma^{-1}\mathbf{M}}{lm - k^2} \\ \mathbf{h} &= \frac{m\Sigma^{-1}\mathbf{M} - \mathbf{k}\Sigma^{-1}\mathbf{O}}{lm - k^2} \end{aligned} \qquad \begin{aligned} k &= \mathbf{M}^\top \Sigma^{-1} \mathbf{O} \\ l &= \mathbf{M}^\top \Sigma^{-1} \mathbf{M} \\ m &= \mathbf{O}^\top \Sigma^{-1} \mathbf{O}. \end{aligned}$$

$$\mu = \frac{1}{2}\sigma^2 + \frac{z_0}{\sqrt{t}}\sigma + \frac{1}{t}\ln\left(\frac{W(t)}{W(0)}\right). \longrightarrow \quad z(\sigma,\mu) = \frac{1}{\sigma}\left(\left(\mu - \frac{\sigma^2}{2}\right)\sqrt{t} - \frac{1}{\sqrt{t}}\ln\left(\frac{W(t)}{W(0)}\right)\right)$$

Max $z(\sigma,\mu)$ s.t. Remaining on the Efficient Frontier : $g(\sigma,\mu) = a\mu^2 + b\mu + c - \sigma^2 = 0$

Solve for the optimal portfolio, i.e., "mu". The solution to the Lagrangian problem is a cubic equation in μ:

$$c_3\mu^3 + c_2\mu^2 + c_1\mu + c_0 = 0,$$

$$\begin{aligned} c_3 &= a^2 \\ c_2 &= \frac{3ab}{2} \\ c_1 &= ac + \frac{b^2}{2} - b - \frac{2a}{t}\ln\left(\frac{W(t)}{W(0)}\right) \\ c_0 &= \frac{bc}{2} - 2c - \frac{b}{t}\ln\left(\frac{W(t)}{W(0)}\right) \end{aligned}$$

The Optimal Goal Probability Point

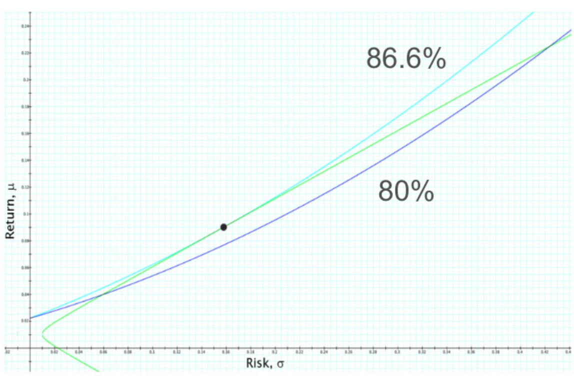

Initial wealth is W(0) = 100, target wealth is W(t) = 125, t = 5 years, and the optimal goal probability is 72%.

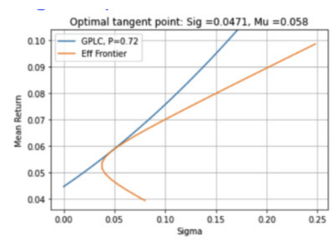

Initial wealth is W(0) = 100, target wealth is W (t) = 150, t = 6 years, and the optimal goal probability is 51%.

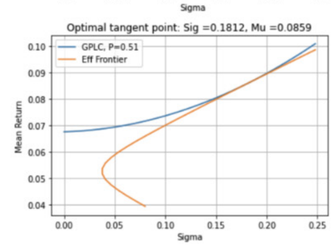

equation in one variable μ. And we were really surprised—we didn't expect this cubic to show up. But it does. And then the question was, can you solve it? And it turns out you can actually solve it. We dusted off Cardano's formula to get the solution. You get three roots for μ. Two of them are negative and one is positive. It's obvious which one is to be picked. And once you've got that, you've solved the problem. Visually, see the exhibits above.

The light green line is the efficient frontier. The dark blue curve is for the 80% goal probability. And if you rotate that line up, you get 86.6% probability GLPC (light blue curve), which is the tangential GLPC. And you also find the point on the efficient frontier you need to be at!

There are two examples above on the right. (1) You want to reach a goal in five years going from 100 to a goal of 125, and you'll tangent at a point where you get a goal probability of 72% and a low risk portfolio. (2) In the second case, you have a little more aggressive goal (of 150) over a slightly longer period of time (6 years). And you need to take a bigger risk by being higher up on the frontier and a lower goal probability of 51%.

That's the static solution. Now, you might have multiple goals over a large number of time periods. And you want to solve this problem dynamically, as shown next.

We can solve the same problem dynamically with infusions and multiple goals. I'm not going through many of the details here, but you set up a giant grid to do your dynamic programming. The transition probabilities depend on which portfolio you choose. It becomes a pretty interesting engineering problem. And of course, you go ahead and solve your standard Bellman equation. I did try to solve it using Bob Merton's frameworks in those famous 1969 and 1971 papers, but the set up here is not amenable to that treatment.

Let's look at the goals (top right below). Here's an investor who's got a 60-year horizon, with many, many different goals. At the top are the ones that are really important goals, the ones at the bottom are less important and more or less discretionary. You can see the first goal, for example, is mortgage payments from one year to 25 years. That means you have 25 sequential goals that you have to meet over time. Now, what we did was we threw in some utility weights for these different goals, and the idea was to dynamically optimize this so that you maximize the total utility you picked up across all these goals across time, deciding at each point in time, which goals to take (or not) and which portfolio to hold.

But utilities aren't important. Once you solve this problem optimally (this table represents almost 300 different goals at a time), and the solution can be expressed in terms of the probabilities of achieving all these goals, given a dynamically optimal strategy for portfolio choice and goal-taking. The top goals having higher utility weights would obviously be given preference, because at every point in time you have a choice, when you're taking some goals and using up money and not meeting more important goals in the future, and

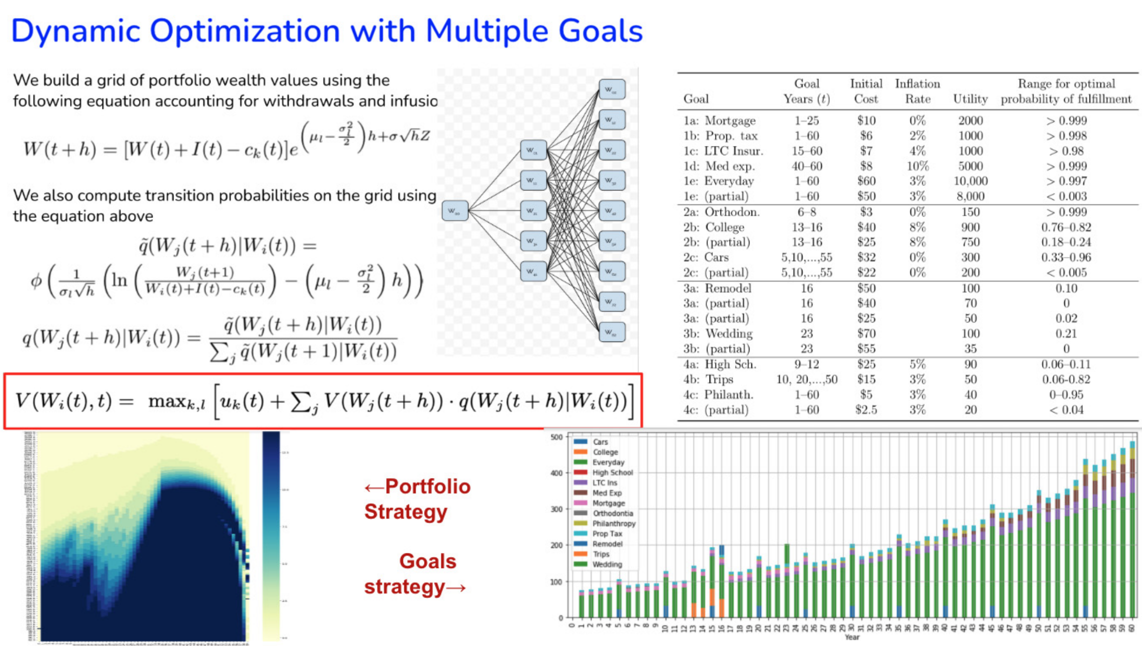

Goal	Goal Years (t)	Initial Cost	Inflation Rate	Utility	Range for optimal probability of fulfillment
1a: Mortgage	1–25	$10	0%	2000	> 0.999
1b: Prop. tax	1–60	$6	2%	1000	> 0.998
1c: LTC Insur.	15–60	$7	4%	1000	> 0.98
1d: Med exp.	40–60	$8	10%	5000	> 0.999
1e: Everyday	1–60	$60	3%	10,000	> 0.997
1e: (partial)	1–60	$50	3%	8,000	< 0.003
2a: Orthodon.	6–8	$3	0%	150	> 0.999
2b: College	13–16	$40	8%	900	0.76–0.82
2b: (partial)	13–16	$25	8%	750	0.18–0.24
2c: Cars	5,10,...,55	$32	0%	300	0.33–0.96
2c: (partial)	5,10,...,55	$22	0%	200	< 0.005
3a: Remodel	16	$50		100	0.10
3a: (partial)	16	$40		70	0
3a: (partial)	16	$25		50	0.02
3b: Wedding	23	$70		100	0.21
3b: (partial)	23	$55		35	0
4a: High Sch.	9–12	$25	5%	90	0.06–0.11
4b: Trips	10, 20,...,50	$15	3%	50	0.06-0.82
4c: Philanth.	1–60	$5	3%	40	0–0.95
4c: (partial)	1–60	$2.5	3%	20	< 0.04

so on. This is what dynamic programming takes care of. If you have an investor look at the college goal, for example, and she says, you know, there's only a 76–82% chance of my kids college being paid for. I want to be more sure of that and I'm going to push that 900 utility weight up to 2,000, and all you would see happening is that the probability for achieving the college goals goes up and that of some of the other goals obviously will go down. The optimal trading strategy is shown bottom left. We use different portfolios that go from low risk to high risk.

It's a very large-scale problem, if you do not code it properly, it takes time and it took 20 minutes to run in our first attempts to program it. We eventually got it down to about 15 seconds using various algorithmic and programming tricks. We have a bunch of portfolios from which to choose every rebalancing period, and they may not even be efficient. But of course you want them to be efficient. The whole collection of portfolios will be available to choose from and dynamic programming tells you exactly which one to use depending on where the portfolio is at any point in time.

In the graphic above you see the various inputs to the problem, i.e., candidate portfolios, goals with utilities and goal costs. These are fed into the algorithm to decide the optimal portfolio as a function of portfolio wealth and time so as to maximize total utility from all goals taken, which are also optimally determined. Given the solution, we can work out the probabilities of achieving each goal under the dynamic strategy. For the same problem, if we change the utility values, we will get a different solution, which is another set of goal probabilities. Each solution is a point in n dimensional space, where n is the number of goals. The collection of such points for all combinations of utilities is called the Efficient Goal Probability Frontier (EGPF), a new acronym that we introduced in this work. See examples of the EGPF when $n = 2$ and $n = 3$ (middle column below). In the case where $n = 2$, we can see how the EGPF traces out the various optimal goal probabilities for two goals. In the top right plot below, you can also see how the EGPF shifts if the initial wealth is modified. And at the bottom right, if $n > 3$, we can represent the results as a table.

Efficient Goal Probability Frontier (EGPF)

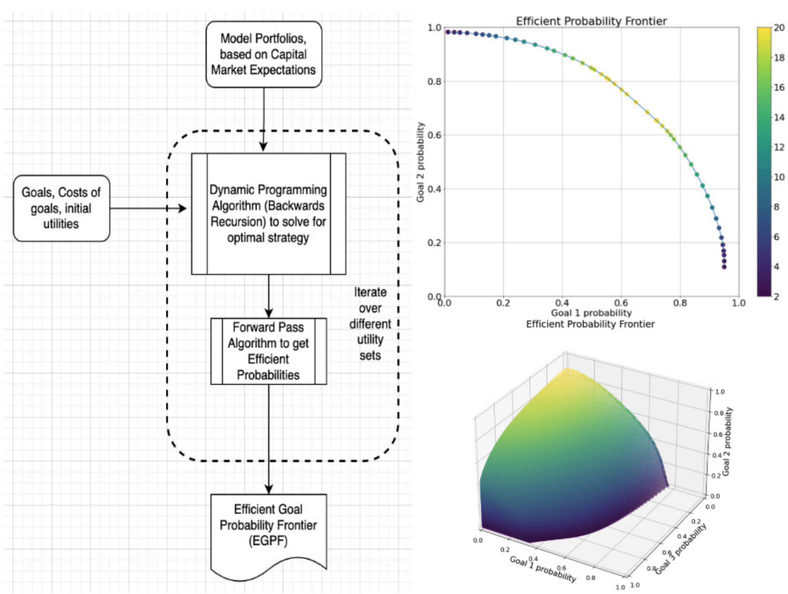

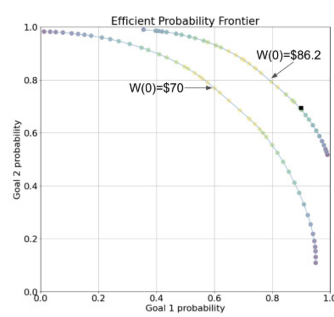

	G	C	U	t	cumProbs	indProbs
0	CAR	7	100	10	0.716683	0.716683
1	HOME	9	150	10	0.842696	0.215903
2	HOME	20	300	10	0.626793	0.626793
3	COLLEGE	10	125	15	0.859657	0.066750
4	COLLEGE	20	250	15	0.792906	0.069156
5	COLLEGE	30	375	15	0.723751	0.083274
6	COLLEGE	40	500	15	0.640476	0.640476
7	RETIRE	50	500	30	0.967559	0.098502
8	RETIRE	90	1000	30	0.869056	0.869056

This framework has now been implemented and is seeing investor traction. It has been branded as Goals Optimization Engine (GOE), shown above.

The nice thing about this is it gives you a completely different view of risks. Normally risk is the standard deviation of returns of the portfolio, and some of the higher moments that we often look at. But here the risk is the chance of not meeting your goals. And this is slightly different, because you might actually minimize variance, for example, but then completely kill off your chances of reaching your goals. There's a sweet spot somewhere in between where you actually balance these two things off. And so that's the tension we're trying to resolve here with an algorithm that we can use for both these risks at the same time.

What I now want to talk about is, how do you take this framework into the world of pensions? Not reaching your goals is the equivalent of the probability of a liability based plan (LBP) failing to meet its obligations. you can take everything we do for an individual and actually port it over quite nicely to what you do for pension plans as well, just thinking now of risk in terms of probabilities of not meeting plan obligations. We may think of LBPs in the following taxonomy, ranging from LDI (low risk) to ALM (high risk), shown below

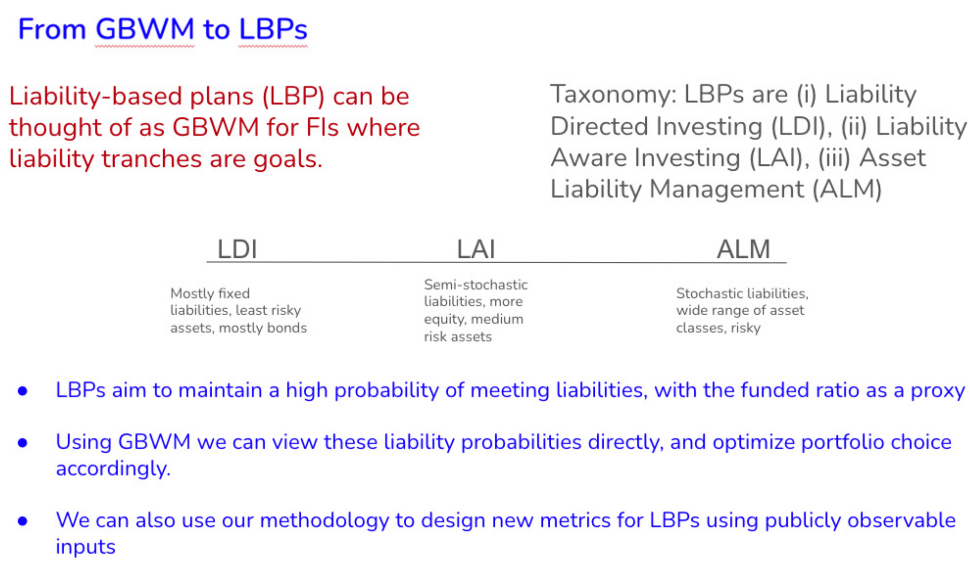

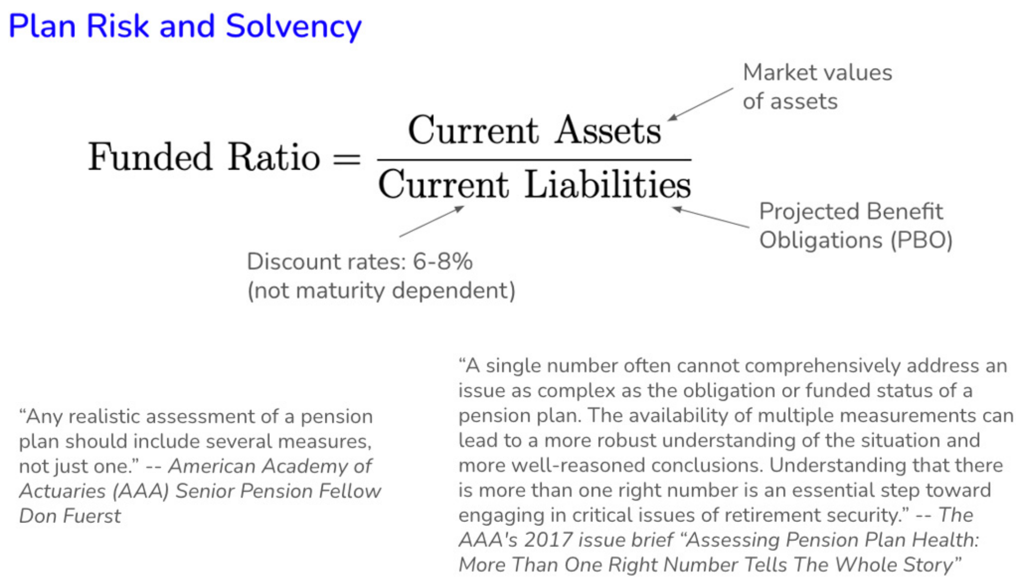

as a continuum. We will look at goal probability-like metrics for LBPs using the ideas in GBWM.

LBP solvency is usually measured by the Funded Ratio shown above. The numerator is the value of plan assets and the denominator is the PBO, the liabilities discounted at discount rates that are usually in the 6–8% range, taken to be a proxy for the expected return on plan assets. It is widely accepted that the funded ratio has many deficiencies, and that more metrics are needed. In this talk, we present additional, new metrics for LBPs, based on the ideas in GBWM.

So we're going to propose three new metrics, shown below.

The SAM and FAM are based on the multiple (usually $\alpha > 1$) of current portfolio assets needed to ensure solvency, so the reciprocal of this multiple acts as a measure of solvency. These measures are based on dynamically optimal trading strategies and hence, account for uncertainty.

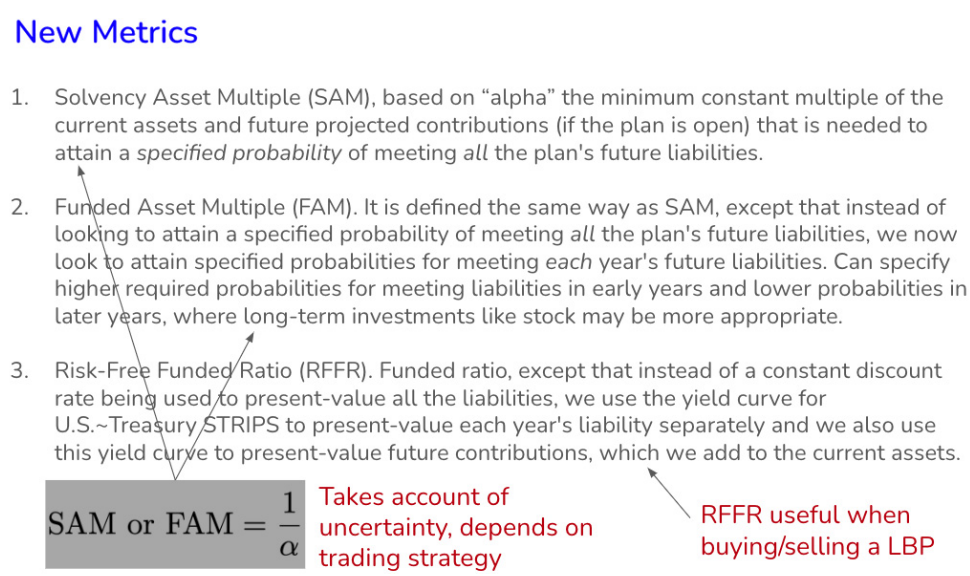

Regulations, management, and metrics

1. The PPA (Pension Protection Act of 2006) and FASB Statement 158 require plan sponsors to fund any shortfalls over certain time horizons and report their liabilities and funding status directly on their balance sheets.
2. Common term structures: (i) the IRS Discount Curve, which is used for the PPA, (ii) Citigroup Pension Discount Curve, used for FASB 158. Yields of high-quality (investment grade) US dollar corporate bonds; reflect interest rate risk and credit spread risk.
3. Sharpe and Tint (1990), Ezra (1991), and Delong et al. (2008) -- maximize the funded ratio minus a constant multiple of surplus volatility.
4. Ang et al. (2013) incorporates downside risk, adds a penalty proportional to the value of an ITM option when the funded ratio drops below one.
5. Upbin et al. (2012) presents a series of benchmark LDI indexes with target durations between six and sixteen years, against which pension fund performance may be benchmarked.

Dynamic Optimization

1. Stockton et al. (2008) approaches LDI by looking to meet a plan's liabilities while constraining the volatility of the funded ratio over time.
2. Delong et al. (2008); Cox et al. (2013) minimize the expected value of the funding variation, where ϱ is a discount rate and U(t) stands for the unfunded liabilities in year t, by determining optimal asset weights and contribution value. $E\left[\sum_{t=1}^{\infty}\left(\frac{U(t)}{(1+\rho)^t}\right)^2\right]$
3. Chang and Cheng, 2002, Josa- Fombellida and Rincoń-Zapatero (2008) considers optimal management of an aggregated dynamic pension fund, with multiple classes of workers with stochastic salaries. Minimize the cost of contributions and maximize the utility of the final surplus, measured as the relative fund level with respect to the mean salary, which is a proxy for the replacement rate.
4. Josa-Fombellida et al. (2018), when the plan is initially underfunded, the objective is to minimize E[(X(T))2], the expected square of the negative surplus, X, at the fund's horizon time, T. But when the plan is initially overfunded, the objective is to maximize the expected power law utility of the positive surplus, X, 1−ξ at the fund's horizon, where ξ > 0 is the risk aversion coefficient. $E\left[\frac{(X(T))^{1-\xi}}{1-\xi}\right]$

The RFFR metric delivers what the plan solvency or fundedness is for purposes of total immunization. Several papers have proposed alternate metrics and objective functions for optimization, and a few are shown above.

For example, in terms of optimization Bill Sharpe and others did a whole bunch of work in the 90's. They recognized that the funded ratio doesn't account for uncertainty at all. While discounting liabilities, the nature of the assets is not really accounted for except in the discount rate, which is mostly within a small range. Different approaches are taken, and our work here complements these shown above.

To summarize, there are many deficiencies of the funded ratio, enumerated in the graphic below. The new metrics we propose account for the volatility of the assets and liabilities, the timing of cashflows, and use dynamic strategies to determine plan fundedness.

There's a lot more in the paper. But really the idea here was to introduce three new metrics: SAM, FAM, and RFFR by taking the ideas from goal-based optimization, thinking of the liability tranches as goals, and using the same sorts of ideas to get probabilities for meeting these tranche obligations, and then converting these probabilities into risk metrics that might be more interesting to look at simply because they are taking stochastic assets and stochastic liabilities, stochastic inflation, and things like that into account. The funded ratio doesn't really advise on a portfolio strategy. All the variables we use here are completely open. We can extend this problem now to many more features like a stochastic number of people in the

Limitations of the Funded Ratio

1. Inaccurate discount rates only affect liabilities, with no balancing effect on the assets in the funded ratio calculation.
2. Using a discount rate only reflects the expected return (of the assets) and not the volatility of a LDI fund.
3. Liabilities, but not assets, have additional sources of uncertainty approximated by constant values.
4. The funded ratio does not reflect well the differential importance of the liabilities across due dates.
5. Using one discount rate for all horizons may be problematic.
6. Future contributions and new liabilities are not addressed in the funded ratio.
7. Augmenting the funded ratio to include new liabilities can also cause issues.
8. The funded ratio does not take into account the backstop abilities of the employer/provider.
9. The funded ratio gives no useful information on how the plan should be invested.

We need to account for the **volatility** of assets *and* liabilities	Account for the **timing** of cashflows	Use **dynamic** strategies

Closing words

"Trains of Thought" (American Economist, 1993)

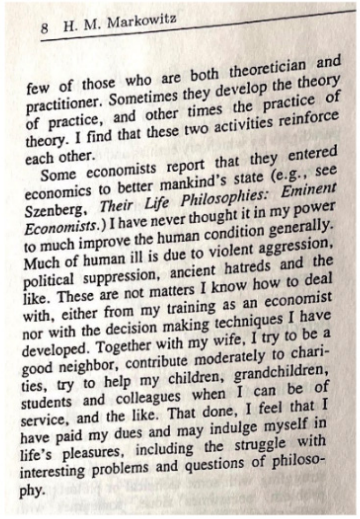

8 H. M. Markowitz

few of those who are both theoretician and practitioner. Sometimes they develop the theory of practice, and other times the practice of theory. I find that these two activities reinforce each other.

Some economists report that they entered economics to better mankind's state (e.g., see Szenberg, *Their Life Philosophies: Eminent Economists*.) I have never thought it in my power to much improve the human condition generally. Much of human ill is due to violent aggression, political suppression, ancient hatreds and the like. These are not matters I know how to deal with, either from my training as an economist nor with the decision making techniques I have developed. Together with my wife, I try to be a good neighbor, contribute moderately to charities, try to help my children, grandchildren, students and colleagues when I can be of service, and the like. That done, I feel that I have paid my dues and may indulge myself in life's pleasures, including the struggle with interesting problems and questions of philosophy.

JOIM's motto: Bridging Theory and Practice

fund, the social security problem, for example, and so on using reinforcement learning.

So to summarize:

(1) The funded ratio is simple but has deficiencies—it does not incorporate a dynamic view of the LBP, among others.
(2) We introduced SAM, FAM, RFFR, i.e., three additional metrics that account for uncertainty.
(3) The funded ratio does not advise on a portfolio strategy for LBPs—dynamic strategies surely help.
(4) The variables needed for these new metrics are publicly available and can easily be specified in standard regulation.
(5) Extensions using RL for large scale problems like social security are the subject of future work.

I wanted to end with the last paragraph of Harry's article, "Trains of Thought."

Harry always talked about both theory and practice. In fact, the *Journal of Investment Management* has the motto of bridging theory and practice. And yes, in the paragraph above, it says there's a bunch of people in finance, sometimes they develop the theory of practice, but other times the practice of theory. These two activities reinforce each other. And if you read that last paragraph you can see in a nutshell what motivated Harry.

I thought I'd end with that only because, you know, pension planning, at least to some extent, is dealing with bigger life problems and how to get to a solution, something Harry cared deeply about.

End