A Tale of Three Epiphanies

That Happened and One That Did Not

Dimitry Mindlin, ASA, MAAA, Ph.D.
President
CDI Advisors LLC
dmindlin@cdiadvisors.com

March 13, 2013

Smee: *I've just had an apostrophe!*
Captain Hook: *I think you mean an epiphany*.
Smee: *Lightning has just struck my brain!*
Captain Hook: *Well, that must hurt*.

Hook (1991)

The value of epiphanies extends beyond solving scientific problems. The presence of an epiphany can illuminate the difficulty of a problem and the elegance of its solution. An epiphany can also make an otherwise tedious story about some obscure science much more exhilarating. There is a reason why the Archimedes principle is the most popular fact in hydrostatics, and this reason is the famous Archimedes bathtub epiphany. Any science could use a couple of epiphanies. Three would be even better.

This paper tells a story that spans over half-a-century and contains four parts. Epiphany-like events play an important role in all four parts – their presence is significant in the first three parts; their absence is critical in the fourth. Each part features an attempt by a renowned economist to develop a paradigm-changing approach to optimal portfolio selection. At some points, all of them thought they succeeded. One of them actually did; three did not.

All of them understood the importance of building solid foundations in finance. All of them clearly realized the weakness of the status quo at the time. All of them were dedicated to building scientifically rigorous approaches to prudent investment practices.

This paper discusses these approaches and relationships between them. All these approaches are different, yet remarkably interconnected. All these approaches have great ideas and limitations. Remarkably, these great ideas complement each other across these approaches. The collective wisdom of all these approaches is superior to a simple compilation of them. This is a story of a whole that is larger than the sum of its parts.

The participants of this story need no introduction. They are Harry Markowitz, Fischer Black, Peter Bernstein, and Paul Samuelson.

"Genuine Theory" vs. "Rules of Thumb and Folklore"

"... *the foundations are the most controversial parts of many, if not all, sciences*."

Leonard J. Savage, *The Foundations of Statistics* (1972)

Any science should have solid foundations. In this respect, the science of optimal portfolio selection is no different than any other. What makes optimal portfolio selection somewhat unusual is the fact that its scientific foundations are relatively new.[1] People have been making asset allocation decisions since the advent of assets, yet the foundations of asset allocation as a science were laid down as recently as in the early 1950s.

Harry Markowitz's work led to the creation of Modern Portfolio Theory (MPT). The ideas advanced in Markowitz [1952] eventually became the cornerstones of modern finance and asset management. It is important to note that a self-confessed epiphany led to the formation of MPT.[2]

Sixty years later, countless institutional and individual investors use MPT as the basis for their asset allocation decisions. Many educational programs present MPT as a great example of a scientifically rigorous approach to investing. Overall, MPT has aged gracefully and withstood the test of time.

Peter Bernstein, the guardian of great ideas in finance, once summarized the impact of MPT:

> "*Before Harry Markowitz's 1952 essay on portfolio selection, there was no genuine* ***theory*** *of portfolio construction – there were just rules of thumb and folklore. It was Markowitz who first made risk the centerpiece of portfolio management by focusing on what investing is all about:* ***investing is a bet on an unknown future***."[3]

The key opposing terms here are "*genuine theory*" and "*rules of thumb and folklore*." Not just any theory, but a "genuine" theory – a theory that provides solid foundations for optimal portfolio selection. And not just any "genuine" theory, but a practical theory that justifies prudent investment practices.

Diversification is a good example of the practicability requirement. The principle "do not put all eggs in one basket" has been well-known for a long time. Yet, there had been no "genuine" theory that justified diversification prior to 1952.

Harry Markowitz once remembered one of Peter Bernstein's memorable pronouncements:

> "*For me, the most memorable Peter Bernstein quote is something he said at a conference. At some point in the general discussion he told us, "You don't know what it was like before MPT." Then he described the piecemeal way portfolios used to be put together. His punch line was: Now you have a process*."[4]

So, here are the key conditions of the solid foundations of investing: a practical "genuine" theory that comes with an effective process of optimal portfolio selection. MPT satisfies these conditions, and that is why MPT has been so successful.

Yet, MPT is not perfect. Its portfolio-centric nature limits its applicability and effectiveness. In particular, the original MPT is a portfolio-centric single-period framework applicable to a *hypothetical* investor rather than a *specific* investor with well-defined financial commitments and institutional constraints. In MPT, these specifics play no role in the portfolio optimization process. This feature of MPT is called the *hypothetical investor assumption* is this paper.

There have been several attempts to develop alternative frameworks that would overcome certain problems with MPT in practical applications. There are two major types of these frameworks –

portfolio-centric (the ones that incorporate the hypothetical investor assumption) and investor-specific. This paper is focused on investor-specific frameworks due to the fact that single-period portfolio-centric frameworks designed for hypothetical investors are inadequate for the needs of major segments of institutional and individual investors.

Why does this author believe that this story is relevant today? *This story is relevant because a broadly recognized approach to optimal portfolio selection that is investor-specific, "genuine," and practical still does not exist*. For many investors, the process of optimal portfolio selection is still mostly based on "rules thumb and folklore." Several attempts to develop such an approach are the subject of this story.

First, this paper discusses the Harry Markowitz's epiphany that led to the creation of MPT. Next, the paper discusses Fischer Black's proposal to incorporate the cost of funding retirement benefits into portfolio optimization. Then this paper discusses Peter Bernstein's attempt to refine the concept of policy portfolio. Next, the paper discusses Paul Samuelson's quest to justify evolving policy portfolios. Finally, the paper incorporates the ideas presented in this discussion into an investor-specific framework for optimal portfolio selection. The paper also demonstrates that this framework serves as a solution to the four-decades old Samuelson problem. This author is optimistic that this framework will be eventually recognized as a practical "genuine" approach and become a mainstream methodology for funding financial commitments.

Epiphany Number One: Harry Markowitz, the Early 1950s

"*The epiphany happened sometime one afternoon*
in the business school library of the University of Chicago."

Harry Markowitz (2011)

The first episode of this story takes us to the time when Harry Markowitz was a young aspiring economist trying to get a Ph.D. in economics in the early 1950s. At the time, financial economics as a scientific discipline and systematic approaches to optimal portfolio selection were virtually non-existent. Publications on the subject of investing largely focused on effective ways to identify the assets with the highest expected upside and paid little attention to the interplay between risk and return.

At the same time, incalculable number of investors designed their portfolios following well-known principles "do not put all eggs in one basket" and "nothing ventured, nothing gained." While these principles intuitively made sense to many, there was no "genuine" economic theory that would incorporate these principles and produce a sensible methodology for optimal portfolio selection. In particular, there was no consensus regarding diversification – John Maynard Keynes, for example, was openly skeptical about its benefits.[5]

In this environment, Harry Markowitz was thinking about doing "a dissertation on the stock market." He was studying the workings of financial markets and pursuing his interest in optimal decision-making under uncertainty. On that particular day, Harry Markowitz was reading John

Burr Williams' "Theory of Investment Value" in the business school library and thinking about stock pricing:

> "*Williams says that the value of the stock should be the present value of its future dividends. I said to myself, dividends are uncertain, he must mean the expected value of future dividends*."

It is important to note that the objective of this valuation is to price the stock at the present. The expected value of the present value of future dividends may be a reasonable figure for the value of the stock at the present.

> "*But if you're only interested in the expected value of a security, you must be only interested in the expected value of a portfolio; ...*"

If a portfolio holds multiple securities, the value of the portfolio is the sum of the values of these securities. If the valuation of a single security is limited to its expected value only, then the valuation of a portfolio is limited to its expected value too.

> " *... and if you're only interested in the expected value of a portfolio, the way you maximize that is to put all your eggs in one basket, to put all your money in one security. Now I knew that is not right: everybody knows you're not supposed to put all your eggs in one basket*."

The expected value of a portfolio is obviously an important measurement. But as a sole measurement, it has a critical flaw – it "knows" nothing about the correlations between the assets that comprise the portfolio, and, therefore, "knows" nothing about diversification. The need for another measurement of the future value of the portfolio that "knows" something about diversification was one of the most important ideas that came to light on that day:

> "*I thought of the returns on securities as random variables, and the return on the portfolio as a weighted sum of these random variables. I knew offhand what the expected value of a weighted sum was, but I did not know offhand what the variance of a weighted sum was. So I got a book off the library shelf: Uspensky's "Introduction to Mathematical Probability"; ... I found where it discussed the variance of the weighted sum of correlated random variables; I saw all those covariances — and that was my moment of epiphany! ...* "

The primary object of interest here is portfolio return, which is a random variable. The mean of portfolio return is the first key measurement, but this measurement alone is insufficient. The identification of the second key measurement of portfolio return – variance – is one of the crucial moments of this story.

Once the key measurements (mean and variance) have been identified, the next step is to determine the desirable "direction" for each of these measurements. As a rule, investors prefer *more* money and *less* uncertainty. In other words, investors prefer portfolios with *higher returns* and *lower risk* (*ceteris paribus,* i.e. other things being equal).

As a result, we have two measurements of portfolios return and two goals for optimal portfolio selection – to maximize the mean and minimize the variance. The solution Harry Markowitz proposed was to combine these goals in one dual objective that would lead to a set of optimal portfolios. Conceptually, the nature of the measurements is beside the point. Given the primary object of interest (portfolio return) and two measurements of this object (mean and variance), the dual objective is *to optimize the first measurement given the second and to optimize the second measurement given the first.*

Thus, the desirability of higher returning and lower risk portfolios along with the *ceteris paribus* principle lead to the following dual objective for optimal portfolio selection.

- To maximize the expected portfolio return at a given level of portfolio variance.
- To minimize the portfolio variance at a given level of expected portfolio return.

This dual objective leads to the "mean-variance efficient frontier" of optimal portfolios.

Alternatively, the same efficient frontier can be generated via combining this dual objective into a single objective. Let us define "the risk-adjusted expected return" as follows:

> Risk-Adjusted Expected Return =
> Expected Return –
> Standard Deviation of Return $\times$ Risk Aversion Factor

where the risk aversion factor is a positive constant. Portfolios that minimize the risk-adjusted expected return for all positive values of the risk aversion factor comprise the mean-variance efficient frontier.

Harry Markowitz continues:

> "*I was a budding young economist interested in two quantities: risk and return; so I drew a graph with risk on one axis and return on the other. ... I distinguished between the efficient and the inefficient portfolios. This was the first efficient frontier. Many of the ideas that are basic to portfolio theory happened in that afternoon.*"

Thus, MPT was born.[6]

In this narrative, there is a fascinating twist of logic that makes this story especially captivating. Let us take another look at the Harry Markowitz's description of the epiphany.

> "*Williams says that the value of the stock should be the present value of its future dividends. I said to myself, dividends are uncertain, he must mean the expected value of future dividends. But if you're only interested in the expected value of a security,*

> *you must be only interested in the expected value of a portfolio; if you're only interested in the expected value of a portfolio, the way you maximize that is to put all your eggs in one basket, to put all your money in one security. Now I knew that is not right ...* "

Dear reader, did you notice the shift from the goal of *valuing a portfolio at the present* to the goal of *maximizing the value of a portfolio in the future*? The initial goal of "present valuing" a problem-specific contingent multi-period stream of payments of uncertain timing and magnitude (future dividends) got replaced by the goal of maximizing the future value of a portfolio over a single time period. The problem-specific stream of payments disappeared. As demonstrated in the next section, the goal of "present valuing" of a problem-specific stream of payments would return in about forty years, courtesy of Fischer Black.

This "goal replacement" is not a logical error. Harry Markowitz simply switched from the subject of the security valuation book to his own scientific interests. For this epiphany, the role of the Williams' and Uspensky's books is similar to the role of the Archimedes' bathtub. It is not terribly important what exactly triggered Harry Markowitz's thinking toward MPT. What is important is this epiphany happened, and the rest is history.

Here are the key aspects of MPT that are especially relevant to this story.

- The main accomplishment of MPT is it represents a first-of-its-kind practical "genuine" theory that comes with an effective process of optimal portfolio selection.

- MPT provides valuable guidance for the optimization of stochastic objects. In MPT, the primary object of interest is the investment return of a portfolio of risky assets, which is uncertain by definition. If the problem is to maximize the investment return, then the solution is a *set of efficient portfolios* (efficient frontier). While the investment return is obviously important, investors have vested interests in other important stochastic objects as well. For example, DC plan participants may want to maximize the asset value at retirement and/or sustainable spending in retirement. MPT suggests a sensible methodology for the optimization of these and other stochastic objects.

- MPT disregards most of the essential characteristics of the investor. For DB plan investors, for example, the nature of the plan (public vs. private), funded status, the maturity of the plan (inactive vs. active population), and the level of contributions required to fund the promised benefits play no direct role in the optimization process. For DC plan investors, for example, the age, retirement age, income, saving rates and account balance play no direct role in the optimization process. Other important concerns investors may have – e.g. the inflation and interest rate risks – play no role in the optimization process as well.

This disregard for most of the investor's essential characteristics is not a flaw but an essential feature of MPT. MPT does not require these characteristics as inputs, which to some extent makes MPT applicable to a broad range of investors. MPT works well for a particular investor as long as the investor's goals are expressed in terms of portfolio return.

This disregard is the primary reason many researchers have sought alternative approaches to MPT. The problem with MPT is it does not incorporate the view Peter Bernstein articulated more than half-a-century after Markowitz [1952]: "*the purpose of accumulating capital is to write checks against it some day*."

Investors had been in no rush to adopt MPT as a primary tool for optimal portfolio selection prior to the early 1970s. And this is despite the fact that the ideas of the initial groundbreaking paper Markowitz [1952] were significantly refined and expanded in Markowitz [1959]. It appears that the steep losses in the U.S. stock market in the early 1970s and the enactment of ERISA in 1974 ignited considerable interest in MPT. MPT has been popular among institutional and individual investors ever since.

Epiphany Number Two: Fischer Black, the Early 1990s

"*... how we set up the goals of a pension plan sponsor matters a lot.*"

Fischer Black (1995)

The second episode of this story begins in the mid-1970s when MPT was becoming increasingly popular. Yet, MPT could not answer certain important questions. What if the investor's goals cannot be expressed in terms of portfolio return alone? What if disregarding the investor's nature and financial commitments is not a good idea?

Corporate defined benefit (DB) plans would be a good example of such an investor. Asset allocation decisions of corporate DB plans may have a significant impact on their sponsors. The laws and governing bodies that regulate the benefits and contributions of corporate DB plans should not be ignored. For example, an important feature of these plans is the pension benefit guarantee provided by the Pension Benefit Guarantee Corporation (PBGC).

Not too long after the enactment of ERISA, Sharpe [1976] and Treynor [1977] demonstrated that the PBGC's guarantee gives the sponsor of a DB plan the "PBGC put." The sponsor could, at least in theory and under certain conditions, transfer its pension commitments and assets to the PBGC. To maximize the value of the "PBGC put," the sponsor should fund the DB plan at the lowest level and *invest its assets entirely in stocks*.

Soon afterwards, Black [1980] and Tepper [1981] looked at the problem of optimal portfolio selection for a corporate DB plan from a different angle. Since investing of retirement assets enjoys certain tax benefits, then the presence of different levels of taxation for different types of assets creates a "tax arbitrage." To maximize the value of this "tax arbitrage," the sponsor should fully fund the DB plan and *invest its assets entirely in taxable bonds.*

Obviously, the recommendations of Sharpe/Treynor and Black/Tepper are opposite to each other. It is equally obvious that MPT has nothing to do with the selection these portfolios.

Having advanced seemingly strong arguments for the total bond allocation in Black [1980], Fischer Black may have been startled by the lack of enthusiasm toward the "tax arbitrage" concept and his asset allocation recommendations to corporate DB plan sponsors. The asset allocation recommendations outlined in Sharpe [1976] were equally unpopular.

Despite Fischer Black's persistent efforts to promote his views on the "corporate pension finance" approach, corporate DB plan sponsors largely ignored the approach at the time. And that is despite the fact that corporations are usually receptive to legitimate ways to lower their taxes. It appears that Fischer Black's views on this approach started to change by the late 1980s, which resulted in the publication of Black [1989].

Black [1989] presents a different method of optimal portfolio selection for corporate DB plans. The paper defines two types of "liabilities" – "the narrow liability" and "the broad liability." The narrow liability is defined as a termination-like liability. The broad liability is defined as the present value of all benefits paid to the current, past, and future employees.

The optimal portfolio depends on the plan's view of its pension liability. If the narrow liability is the most appropriate, then pension assets should be invested in bonds. If the broad liability is the most appropriate, then pension assets should be invested in stocks. If some "intermediate" liability (that is between the narrow and the broad liabilities) is the most appropriate, then the best mix should be between the all-bond and all-equity allocations.

In retrospect, it is clear that Black [1989] offered neither a "genuine" theory nor a practical process for optimal portfolio selection. The concept of "pension liability" was not properly defined, and its meaning was unclear in the context of the paper. The concept of "the plan's view of its pension liability" was not specific enough to serve as a basis for a "genuine" theory.

Soon after Black [1989], several authors demonstrated that pension accounting liabilities could be incorporated in the process of optimal portfolio selection. Sharpe-Tint [1990] introduced the concept of *surplus optimization* that utilized "surplus returns" in lieu of the "asset-only" portfolio returns. In the early 1990s, a series of papers by Martin Leibowitz and his colleagues presented further examples of the utilization of pension accounting liabilities in the development of optimal portfolios for corporate DB plans.[7] While these attempts enjoyed some success, the popularity of asset/"pension accounting liability" management was nowhere near the popularity of MPT.

Then came Fischer Black's final "bombshell" message to pension practitioners entitled "The Plan Sponsor's Goal" (Black [1995]).[8] Black [1995] is a short and stunningly controversial paper.

This paper is remarkable for several reasons. First, Black [1995] does not limit itself to corporate pension plans and discusses pension plans in general. Second, Black [1995] ignores pension accounting liabilities, government guarantees, and the "tax arbitrage." Third, Black [1995] does not mention the soon-to-become-popular Black-Litterman model – an alternative technique to incorporate the investor's views in the process of optimal portfolio selection.[9] The last noninclusion is especially noteworthy since Fischer Black and Robert Litterman developed this model roughly at the same time as Fischer Black was developing the ideas of Black [1995].[10]

It appears that Fischer Black discovered an entirely different perspective on the subject. This author conjectures (or, in all honesty, speculates) that Fischer Black may have had an epiphany-like event at some point in the early 1990s. The contrast between Fischer Black's new perspective and his earlier views was so stark, that two letters to the *Financial Analysts Journal* written in 2008 suggested that somebody else wrote some parts of Black [1995]. (This suggestion turned out to be incorrect.)[11]

Black [1995] starts with a series of questions.

> "*How should defined benefit pension plan sponsors choose an investment strategy for their pension funds? How should they allocate the assets in their funds among broad asset classes such as stocks, bonds, and real estate? How should they diversify within and across asset classes? How should their plans' liabilities affect their investment strategies? Should the indefinite life of such plans make a difference?*"

Black [1995] recognizes that "*we can imagine many answers to our investment policy questions*" and gives a short review of these answers:

> "*Some plan sponsors want to maximize expected return subject to limits on the amounts invested in specific asset classes. Some want to maximize expected return on the pension fund assets subject to a given risk of the assets. ... Some sponsors look at their liabilities. They choose assets to match the character of their liabilities, or they try to maximize expected return for given risk relative to their liabilities. ... Plan sponsors may define their liabilities narrowly or broadly. ... All these points of view lead to different answers to our investment policy questions. All of them seem legitimate.*"

Having acknowledged the legitimacy of these approaches, which include MPT, surplus optimization and his own views circa 1989, Fischer Black summarily dismisses all of them!

> "*My view of the pension plan suggests an approach that differs from all of these.*"

Black [1995] proposes two momentous principles of retirement investing. The first principle is *it is necessary to spell out the plan sponsor's goal* that would guide the process of optimal portfolio selection. The title of the paper – "The Plan Sponsor's Goal" – highlights the importance of this principle. Black [1995] hammers it in further:

> "*... how we set up the goals of a pension plan sponsor matters a lot.*"

The last sentence of the paper is a pointed question that emphasizes this principle as the final takeaway:

> "*Which goal makes the most sense to you?*"

On the surface, the principle "spell-out-the-goal-first" sounds almost trivial. Of course, one has to spell out the problem before starting to look for a solution. Beneath the surface, however, the idea that a "genuine theory" of optimal portfolio selection should contain a clearly enunciated objective is imperative. To this day, this principle is far too often missing in practical work.

Now it is clear why the "answers to our investment policy questions" did not satisfy Fischer Black's requirements. The plan sponsor's goals implied by these approaches were problematic:

> "*I question the usual formulation of the plan sponsor's goals.*"

However, Black [1995] goes much further than indicating problems in other approaches. Black [1995] proposes a novel formulation of the plan sponsor's goal "that differs from all of these":

> "*... a plan sponsor may want to choose an investment strategy to minimize the present value of future contributions to the plan.*"

The goal of minimizing the present value of future contributions is the second momentous principle proposed in Black [1995]. This goal is as groundbreaking as it is controversial.

The logic behind this principle is simple. If a portfolio produces the lowest present value of future contributions, then this portfolio produces the highest present value of the sponsor *ceteris paribus*. In other words, lower labor cost means higher value of the sponsor. From this perspective, Fischer Black indicates that *this is a portfolio selection problem*:

> "*The plan sponsor could choose assets that have the highest possible present values.*"

At the same time, the paper still contains certain elements of "corporate pension finance," and Fischer Black brings up one its cornerstones (twice!):

> "*The present value of $100 in stocks is equal to the present value of $100 in bonds.*"

From that perspective, Fischer Black indicates that *this is not a portfolio selection problem*:

> "*Diversification does not affect present value in any direct way. ... Similarly, the allocation of a plan's assets between debt and equity investments does not directly affect the assets' present values.*"

Clearly, Black [1995] presents two conflicting perspectives. The mindset in which the present value of $100 invested in any security is always $100 sharply disagrees with the mindset in which there exist "assets that have the highest possible present values." One cannot have it both ways. This disagreement goes to the heart of the general concept of present value.

Loosely speaking, present value is the current worth of a series of payments. Present value of a series of payments given a series of investment returns is defined as the asset value at the present that is required to make all payments in the series. A present value is calculated via discounting the payments by the investment returns.

Think of $100 invested in a security and the future value of this $100 in *N* years. What is the present value of this future value? In other words, how much should be invested in the security to accumulate this future value in *N* years with certainty? The answer is obviously $100. In the discounting procedure, the same return factors in the numerator and denominator simply cancel out. The nature of the security is immaterial.

This is the logic behind the statement "the present value of $100 invested in any security is $100." In this portfolio-centric framework, the issue of optimal portfolio selection is relevant only in the context of taxation. This is one of the main messages of Black [1980].

This portfolio-centric framework has stringent limitations – the framework neither incorporates the investor's financial commitments nor allows future risks. Black [1995] is correct to note that "*diversification does not affect present value*" in this framework. But then this present value is always $100 and cannot be minimized, which contradicts the main message of Black [1995].

To find the "*assets that have the highest possible present values*," these limitations should be relaxed. The expanded framework would incorporate the investor's financial commitments directly and recognize that many investors fund their commitments via investing in risky assets. Portfolio selection impacts future risks and returns and, therefore, impacts present values. As a result, present values can be optimized in the expanded framework. However, as far as the ability to fund future financial commitments is concerned, $100 invested in stocks is not the same as $100 invested in bonds, which contradicts "corporate pension finance."[12]

What we have here is a spectacular clash between Fischer Black circa the early 1980s and Fischer Black circa the early 1990s.

This clash is unresolved in Black [1995]. Furthermore, Fischer Black sheds little light on the calculation of the present value of future contributions. He offers even less guidance about the ways to minimize this present value. Fischer Black's excruciatingly difficult personal situation at the time of the paper's submission may have been the primary reason for these and other "imperfections."[13]

Black [1995] is a short two page paper that does not even attempt to define all concepts it discusses. Fischer Black may have thought that delivering the main message on time – "*how we set up the goals of a pension plan sponsor matters a lot*" – was more important than refining all the details.[14]

This paper does not aspire to resolve all controversies in Black [1995]. However, this paper does attempt to illuminate the great ideas in Black [1995]. Black [1995] offers useful indications about Fischer Black's vision even though a certain level of conjecturing is unavoidable. The rest of this section discusses possible developments of these ideas.[15]

Let us start with the calculation of the present value of future contributions. It appears that this calculation involves certain complications related to the regulations that govern pension contributions. However, these regulations primarily impact the timing of these contributions, not their present value. Therefore, a simple approximation removes these complications.

Let us assume that the existing assets and the future contributions will be used exclusively to make the plan's future benefit payments. Then

Existing Assets +
Present Value of Future Contributions =
Present Value of Future Payments

Since the existing assets is a constant, then the problem of minimization of the present value of future contributions is equivalent to the problem of minimization of the *present value of future payments* (PVFP) to plan participants. The concept of PVFP is much more transparent.

The next issue is the nature of these present values. One approach to calculate PVFP is to estimate the market value of a portfolio of tradable assets whose payouts match the benefit payments. The discounting procedure would utilize the market interest rates consistent with the riskiness of the benefit payments. This approach would be consistent with the "corporate pension finance" framework, but inconsistent with the goal of minimizing PVFP. A market value cannot be minimized at the present. It is what it is.

A different approach would be to recognize the stochastic nature of present values for the investors that endeavor to fund their financial commitments via investing in risky assets. Stochastic present values were well-known in the early 1990s; their properties were presented in textbooks and academic publications.[16] Stochastic present values offer powerful analytical tools for the cost-risk analysis of retirement plans.[17] Stochastic present values can be optimized similar to the (stochastic) portfolio returns discussed in the prior section.

Did Fischer Black favor the first approach, the second approach, or something entirely different? Black [1995] is silent about this issue. This author's conjecture is Fischer Black was much closer to the second approach than to the first one because the impracticality of the minimization of present values in the first approach was too obvious. Assuming that Fischer Black envisioned stochastic present values, how would he minimize them?

To answer this question, let us think of the area where Fischer Black's achievements are renowned – option pricing. The Black-Scholes formula calculates the value of an option as the *mean of the stochastic present value* of the option payments taken over all future economic scenarios. “*Why not use the same criterion to make pension fund investment decisions*?”

Let us assume for a moment that PVFP could be valued similar to Fischer Black's view on the valuation of options. Then “the present value of benefit payments” could be viewed as the mean of PVFP taken over all future economic scenarios. This expected value critically depends on the future risks and returns.

A very risky strategy may produce very low long-term returns in a substantial segment of economic scenarios. A very safe strategy may produce stable but still inadequate returns in most scenarios. Both strategies may generate higher expected PVFP. The challenge is to find the right trade-off between risky and safe strategies that would generate the lowest expected PVFP.

Yet, Fischer Black clearly understood the limitations of expected values:

> “*A very risky strategy with high present value, though, may hurt the plan participants. We should only consider strategies with low or moderate risk.*”

In other words, "a very risky strategy" should be avoided regardless of the resulting expected values, even if it would produce a "high present value" of the sponsor. But how would you identify a “very risky strategy" without a measurement of risk?

Dear reader, does this discussion remind you a young economist named Harry Markowitz sitting in the business school library of the University of Chicago more than four decades before Black [1995]? Harry Markowitz was looking at the expected value of a stochastic object and thinking that the expected value alone was insufficient. Under the assumptions outlined above, Fischer Black would be in a similar position.[18]

Now would be a great time to pull "a Harry Markowitz" on the problem of minimizing PVFP as follows. The expected value alone is insufficient; there is a need for a second risk measurement, e.g. the variance of PVFP. "Very risky" strategies would be quantitatively defined, and the trade-off between a very low expected value and a very high risk would be transparent.

The next step is to determine the desirable "direction" for each of these measurements – investors want lower expected present value and lower risk. Given these two measurements and their desirable directions, the solution is to optimize the first measurement given the second and/or to optimize the second measurement given the first. The dual objective of minimizing the expected value given risk and minimizing risk given the expected value leads to the concept of "*cost-risk efficient frontier*." The analogy with MPT is clear.

Alternatively, the same efficient frontier can be generated via combining this dual objective into a single objective. Let us define "the risk-adjusted expected cost" as follows.

Risk-Adjusted Expected Cost =
Expected PVFP +
Standard Deviation of PVFP $\times$ Risk Aversion Factor

where the risk aversion factor is a positive constant. Portfolios that minimize the risk-adjusted expected cost for all positive values of the risk aversion factor comprise the cost-risk efficient frontier. Again, the analogy with MPT is clear.

It did not happen in Black [1995]. This extraordinary paper remains little-known to this day. Black [1995] may have been too controversial, too inconsistent, too short, as well as may not have presented a practical "genuine" theory. Or, perhaps, "*Fischer was, as usual, too far ahead of his potential clients*."[19]

Epiphany Number Three: Peter Bernstein, the Early 2000s

"... *the purpose of accumulating wealth is to spend money at some future point for some purpose ...* "

Peter Bernstein (2003)

The third episode of this story begins in the early 2000s. By that time, most institutional investors had embraced the concept of *policy portfolio* defined as the existing portfolio and the portfolios the investor expects to use in the future (e.g. the same portfolio). Many investors implemented elaborate procedures for the development of policy portfolios. MPT and the capital market assumptions associated with it had become the key components of these procedures.

Many investors had come to utilize MPT in a manner that the original MPT did not necessarily assert. Namely, the capital market assumptions and resulting optimal portfolios had become "long-term" – a precondition the original MPT as a *single*-period model did not require. The utilization of the "long-term" version of MPT was based on the following two assumptions.

Assumption A. A policy portfolio must be specified at the present.
Assumption B. All portfolios in a policy portfolio are the same.

Most retirement plans had assumed that the plan assets would be regularly rebalanced to the same portfolio, and this process would continue indefinitely. This "target" portfolio had been equity-centric in most cases ranging between 50/50 to 70/30 [stocks/bonds]. The 60/40 portfolio had been by far the most popular.

Enter Peter Bernstein, a broadly recognized authority on the workings of capital markets and the history of finance. Peter Bernstein's works – "Capital Ideas," "Against the Gods," the semi-monthly publication "Economics and Portfolio Strategy" and many others – had become required reading for many practitioners and academics alike.

Peter Bernstein offered his full support to the idea of a stationary long-term 60/40 policy portfolio in an essay named "The 60/40 Solution" published in early 2002.[20] He observed that the 60/40 policy portfolio was very popular "before the great bull market of the '90s," but lost some of its popularity in the 1990s. Many investors moved away from the "stodgy" 60/40 and sought out more aggressive allocations in pursuit of higher returns "as stocks roared ahead."

After the earth-shaking events of the early 2000s, Peter Bernstein was in favor of returning to the good old 60/40. The last sentence of the essay left no doubt that Peter Bernstein fully embraced the idea that there should be a single portfolio that would serve as a long-term target allocation.:

> "*I propose restoring 60/40 to its rightful place as the center of gravity of asset allocation for long-term investors*."

Incredibly, Peter Bernstein's view on the subject of policy portfolios became diametrically opposite a year later. This turnaround was so abrupt, that this author conjectures (or speculates) that Peter Bernstein may have had an epiphany-like event in 2002.

The March 1, 2003 issue of "Economic and Portfolio Strategy" was entitled "Are Policy Portfolios Obsolete?" (Bernstein [2003A]). Peter Bernstein's answer to the question in the title is an astonishingly unequivocal "yes":

> "*My response to the question in our title is* ***policy portfolios are obsolete***."

Bernstein [2003A] questions the quintessence of the policy portfolio development process many investment committees employed. The "long-term" version of MPT is the core of the process:

> "*A policy portfolio is a broad statement of an investment committee's* ***long-term*** *views of the relative attractiveness of risk in available asset classes. More technically, a policy portfolio is the result of an optimization process involving long-run expected returns, volatilities and covariances in order to maximize future long-run returns at the chosen level of exposure to risk.*"

These long-term views are the critical component of the process:

> *"... long-term is the key to the concept, theoretically and operationally.*"

But the development of long-term views is exceedingly difficult. In numerous publications, Peter Bernstein had been adamant about his mantra "*we do not know what the future holds*." Bernstein [2003A] accentuates that the fundamental nature of long-term views is problematic:

> *"The concept of the long run is vulnerable on both empirical and philosophical grounds.*"

If long-term views are problematic, then the policy portfolios that critically depend on these views are problematic as well.

Form the technical and operational standpoints, the key inputs to the optimizer (e.g. the risk premium, correlations, etc.) are of paramount importance. Yet, Peter Bernstein would question the productivity of an investment committee that spent considerable amount of time arguing about the size of the risk premium or whether the correlation between stocks and bonds should be 0.2 or -0.2. Is long-term return forecasting the primary mission of investment committees?

Furthermore, Peter Bernstein and Robert Arnott published a paper entitled "What Risk Premium Is 'Normal'?" in 2002.[21] The paper claims that "*the long-term forward-looking risk premium ... may well be near zero, perhaps even negative*." Bernstein [2003A] reiterates this point:

> "... *in today's markets, the expected equity risk premium over the long run (say, the next ten years) is too low to justify a rigid adherence, through thick and thin, to a portfolio with a policy allocation of more than fifty percent in equity instruments.*"

Thus, Bernstein [2003A] rejects the abovementioned *Assumption A*, *Assumption B*, and the broadly utilized "long-term" version of MPT.

It is important to note that the main message of Bernstein [2003A] is its rejection of *Assumption A* (expected portfolios should be pre-specified). Bernstein [2003A] is silent about *Assumption B* (all expected portfolios are the same), which is rejected simply because the already rejected *Assumption A* is a precondition to *Assumption B*. This observation is important for the next section of this paper.

In general, when Peter Bernstein spoke, the investment community listened. Bernstein [2003A] jump-started a broad discussion about the concept of policy portfolio and related issues. In a later paper, Peter Bernstein mentioned the spirited discussion Bernstein [2003A] ignited:

> "... *the image of the policy portfolio as the lodestar of asset allocation has long been accepted without question among investment officers and investment committees in the vast majority of foundations, pension funds, and endowment funds.*
>
> *At least until now. Nothing I have ever written and spoken about has created such waves as my attack on the policy portfolio concept. One senses the investment community was waiting for somebody to point out that the emperor of asset allocation has no clothes, and it was I who turned out to be the innocent child.*"[22]

The direction of this discussion, however, took an abrupt turn a few months later. In August, 2003, Peter Bernstein dropped another "bombshell." In a paper entitled "Which Policy Do You Mean?" (Bernstein [2003B]), Peter Bernstein continued the discussion of policy portfolios. It appears that a paper Peter Bernstein received about ten days before the publication of Bernstein [2003B] – entitled "Solving the Investor's Problem" written by Kevin Kneafsey (Kneafsey [2003]) – compeled a complete turnaround. Did Kneafsey [2003] trigger another epiphany-like event? Regardless of the answer to this question, the main conclusion of Bernstein [2003B] was the opposite of the main conclusion of Bernstein [2003A]: "... *a policy portfolio is essential*."

Bernstein [2003B] commends the ideas presented in Kneafsey [2003] and reiterates them mostly unchanged but from a somewhat different perspective. In particular, Bernstein [2003B] emphasizes the importance of setting the right goals:

> "*... all of us have been debating the wrong questions. The issues we have been discussing are not the issues that matter.* ***Instead of asking whether "market timing should replace a policy portfolio, we should be asking which "policy" a policy portfolio is designed to fulfill***."

Those "all of us" who have been debating the wrong questions would well-advised to remember Black [1995]. Fischer Black's question "*Which goal makes the most sense to you?*" sounds very much like Peter Bernstein's "*What is the investor's problem?*"

The answer to this question Bernstein [2003B] suggests is the following:

> "*The investor's problem is to fund a stream of liabilities. ... the policy is to provide with the highest probability of being able to pay for the groceries when the time comes*."

In other words, Bernstein [2003B] suggests that *the probability of funding future payments* is a valuable measurement. As discussed in the previous section, Black [1995]] suggests that *the mean of the present value of future payments* (PVFP) a valuable measurement as well. Are these measurements related in any way?

In fact, they are closely related. The probability of funding is equal to the probability of PVFP being less than the existing asset value.[23] Consequently, both Black [1995] and Bernstein [2003B] utilize measurements of PVFP, which should be the primary object of interest.

Now would be a great time to pull "a Harry Markowitz" when the primary object of interest is PVFP rather than portfolio return. Bernstein [2003B] suggests one measurement – the probability of funding future payments. One measurement is insufficient, there is a need for a second one, e.g. the expected PVFP as effectively proposed in Black [1995]. The desirable "directions" for these measurements are higher probability and lower expected value. Given these two measurements and their desirable directions, the solution is to optimize the first measurement given the second and/or to optimize the second measurement given the first. This dual objective leads to the "cost-risk efficient frontier" ...

It did not happen in the Bernstein [2003A]–Kneafsey [2003]–Bernstein [2003B] discussion. This discussion may have become "the single most profound change in investment thinking since Markowitz's work on portfolio theory," but it did not happen either. The discussion offered neither "a genuine" theory nor a practical process for optimal portfolio selection.

Overall, the messages Fischer Black and Peter Bernstein sent to the investment community were very similar: the conventional approach is inadequate, and there is a need for innovative thinking about the concept of policy portfolio. Their works contain valuable guidance to the future developments of the concept of policy portfolio.

Throughout his illustrious career, Peter Bernstein may have changed his mind about a couple of things. But he resolutely maintained his long-term views on the concept of "long-term":

> "*The long run is an impenetrable mystery. It always has been.*"[24]

The Epiphany That Did Not Happen: Paul Samuelson, From the 1960s to the 2000s

> "*I had not proved that you should do something different over time. That has been a recurring problem ...*"
>
> Paul Samuelson (2005)

When Peter Bernstein stated that "policy portfolios are obsolete," he rejected *Assumption A* (expected portfolios should be pre-specified) and necessarily by implication *Assumption B* (all expected portfolios are the same). When Peter Bernstein stated that "a policy portfolio is essential," *Assumption A* got reinstated.

But what about *Assumption B*? In this discussion, *Assumption B* has received little attention. This assumption has been accepted without question among investment officers and investment committees, has it not? In general, should optimal policy portfolios rebalance to the same allocation? Is there a "genuine" theory that supports the assumption of stationary policy portfolios?

In fact, such a theory has been in existence since the 1960s. Paul Samuelson demonstrated that an optimal policy portfolio is stationary under the assumptions of constant relative risk aversion and random-walk returns:[25] As Paul Samuelson put it in an interview,

> "*... you do the same thing when you have one year to go … as you could do if you had 100,000 years to go*."[26]

This result in no way satisfied Paul Samuelson. Despite its scientific rigorousness, Paul Samuelson called this result "*a defiance of folk wisdom and casual introspection*" and "*felt that the first pass at the problem was a failure*."[27]

What was "the problem" that "the first pass" failed to solve? "Folk wisdom and casual introspection" suggested the principle "more equities for the young, less equities for the old," and many investors accepted this principle "without question." Financial economists, however, could not present a "genuine theory" that would justify this principle. The development of such a theory is called *the Samuelson problem* in this paper.[28]

The solution to the Samuelson problem would be a practical framework that would rationalize evolving policy portfolios. The framework Paul Samuelson utilized for "the first pass at the problem" failed to prove that optimal policy portfolios should evolve. Hence, it was "a failure":

> "*I had not proved that you should do something different over time*."[29]

Paul Samuelson had returned to this problem over and over again over the years:

> "*That has been a recurring problem, of course …*"[30]

Paul Samuelson's strong desire to rationalize evolving policy portfolios is reflected in the titles of some of his papers. Twenty years after Samuelson [1969], Samuelson wrote a paper named "A Case At Last for Age-Phased Reduction in Equity" – note the priceless "At Last" in the title (Samuelson [1989]). Another "At Last" appeared in the title of section 7 in Samuelson [1991] – "Intuition At Last Vindicated." Yet another "At Last" appeared in the title of Samuelson [1992] – "At Last, a Rational Case for Long-Horizon Risk Tolerance and for Asset-Allocation Timing?"

Moreover, at the conclusion of Samuelson [1991], Paul Samuelson proclaimed,

> "*My three-decade search for confirmation of enhanced risk tolerance among long horizon investors has … achieved success*."

These papers, however, did not successfully conclude Paul Samuelson's "three-decade search" for a solution to the problem, even though they did rationalize evolving policy portfolios under certain conditions. Paul Samuelson appeared to understand the limitations of his solutions and modestly claimed that "some very considerable progress has been made."

A glimpse of Paul Samuelson's thinking about the role of scientific foundations appeared in a delightful essay entitled "What Makes for a Beautiful Problem in Science?"[31] Discussing the history of the calculus of variations, Paul Samuelson wrote:

> "*All this is good enough for the brilliant eighteenth century. But by the nineteenth it was a scandal that a rigorous mathematical theory was still not known*."[32]

Paul Samuelson may have felt that the reliance on "rules of thumb and folklore" was no longer acceptable, and "it was a scandal that a rigorous mathematical theory" that rationalized evolving policy portfolios was still not known. Paul Samuelson may have thought that closing all logical gaps in the problem of stationary policy portfolios was essential.

Remarkably, other participants of this story may have given Paul Samuelson valuable guidance to solving this problem. Fischer Black may have advised to scrutinize the primary objective. Peter Bernstein may have suggested an objective different from the maximization of expected utility. For instance, "the purpose of accumulating capital is to write checks against it some day," and the primary goal of optimal portfolio selection should be to maximize the likelihood that all these checks are funded. Harry Markowitz may have suggested to select two appropriate measurements, build an efficient frontier, and see if optimal policy portfolios were stationary.

Of course, it is a bit silly to presume that suggestions made at different times would be useful for Paul Samuelson's quest. Nonetheless, it is fascinating to ponder the way Paul Samuelson may have responded. This author conjectures that Paul Samuelson's response to the Fischer Black's question regarding investment goals would likely be that the right goal for policy portfolio optimization is the maximization of the investor's expected utility.

While the goal of maximizing expected utility is popular in academic literature, this popularity is far from universal. For instance, Roy [1952] contains the following memorable statement:[33]

> "*In calling in a utility function to our aid, an appearance of generality is achieved at the cost of a loss of practical significance and applicability in our results. A man who seeks advice about his actions will not be grateful for the suggestion that he maximise expected utility.*"[34]

If the goal is in question, then the solution is in question as well. Maximizing expected utility gives one set of optimal policy portfolios, and pursuing some other goal may give another set. Are there other sensible investment goals out there? Yes, undoubtedly. One of these goals is Fischer Black's goal of minimizing "the present value of future contributions." Another one is Peter Bernstein's goal of achieving "the highest probability of being able to pay for the groceries when the time comes." The next section demonstrates how the ideas discussed in prior sections lead to a solution to the Samuelson problem. The Appendix contains a simplified numerical example of the solution.

Paul Samuelson never had an epiphany that led to a solution to this problem and did not finalize his decades-long quest to build a framework that rationalized evolving policy portfolios.

Commitment Driven Investing

... the liability structure must dominate
the management of the assets."

Peter Bernstein (2003)

This section gives a very concise presentation of the basics of an optimal portfolio selection framework that incorporates the ideas discussed in prior sections.

Let us start with the key premise of MPT. MPT identifies the primary object of interest – portfolio return. Then MPT specifies two key measurements of portfolio return – expected return and variance of return. Based on these measurements, MPT generates optimal portfolios.

The main problem with this approach is reflected in Peter Bernstein's assertion "the liability structure must dominate the management of the assets." This "liability structure" plays no part in MPT. Regardless of the investor's "liability structure," MPT performs one-period optimization of the *future* asset value assuming that the asset value is known at the *present*.

For an investor with a non-trivial "liability structure," the challenge is exactly the opposite. The future values ("the liability structure") are given. The investor's biggest challenge is at the present – how much to contribute and how to allocate the assets. Fischer Black's objective of minimizing the present value of future contributions and Peter Bernstein's objective of maximizing the likelihood of funding "the liability structure" are consistent with this challenge.

Before we go any further in the development of the new paradigm, we need to go through a short exercise in semantics. The term "liability structure" likely means future payments. If this is the case, this choice of words is somewhat unfortunate. This author believes that the term "liability" is generally inappropriate to indicate a future payment.

"Liability" is an accounting, actuarial, and legal term whose meaning critically depends on the context the term is used. Traditionally, the term "liability" in accounting and actuarial reporting is reserved for *present values* of future payments, not the payments themselves. To use the term "liability" for both present values of future payments and future payments is at the very least undesirable. In a discussion related to scientific foundations and "genuine" theories, the use of the same term for two fundamentally different objects is next to unacceptable.

This author uses the term "financial commitment" (or just "commitment") for future payments and the term "liability" for their accounting and actuarial present values. The commitment of a DB plan is the benefit payments promised to the plan participants. The commitment of a DC plan participant may be the lifetime post-retirement income equal to a pre-specified percentage of the pre-retirement income. The term "commitment" better reflects this situation than the term "liability." If the plan participant fails to fund the desired income, no one is "liable" for anything. The term "liability" possesses too many undesirable connotations, and the term "commitment" is much more appropriate in many cases.

If this convention is accepted, two statements discussed in the paper get rephrased as follows:

Old (P. Bernstein):	"*The liability structure* must dominate the management of the assets."
New:	"*The financial commitment* must dominate the management of the assets."

Old (P. Bernstein):	"The investor's problem is to fund a *stream of liabilities*."
New:	"The investor's problem is to fund a *financial commitment*."

Here are the basic attributes of a framework that incorporates the ideas discussed in this paper.

- The investor's goal is to fund a financial commitment.
- The investor's primary risk is a *shortfall event* defined as the failure of the investor's existing assets and future contributions to fund the commitment.
- Prudent risk management requires a comprehensive set of risk measurements. These measurements include, but are not limited to, shortfall probability, size and volatility.[35]
- The primary object of interest is the *stochastic present value of the commitment* (SPVC).
- The selection of key measurements results in various versions of CDI. In particular, the selection of the mean and variance of SPVC leads to the "*mean-variance*" version; the selection of the "disastrous" cost and shortfall probability leads to the "*safety-first*" version.
- Optimal policy portfolios in CDI form a *cost-risk efficient frontier*.

- The presence of a financial commitment expands the conventional two-dimensional "Risk–Return" paradigm to the three-dimensional paradigm "Commitment–Cost–Risk" called *the funding triangle*. The funding triangle is optimized according to the principle “given two vertices, optimize the third.”

To emphasize the central role financial commitments play in this framework, the framework is called *Commitment Driven Investing* (CDI). The CDI framework includes, among other useful features, a powerful and flexible portfolio optimization engine. Mindlin [2009B] presents a general introduction to CDI as an economic theory.

The rest of this section presents a succinct demonstration that the CDI framework is a solution to the Samuelson problem. In other words, the CDI framework generally generates *evolving* optimal policy portfolios. Incidentally, a solution to this problem was found in 2009, the year Paul Samuelson passed away at age 94. This section presents just a sketch of the solution; a detailed discussion is presented in Mindlin [2009A]. A simplified numerical example of the solution is presented in the Appendix.

Let us think of an investor with a financial commitment to fund (e.g. a DB plan committed to funding the promised benefit payments or a DC plan participant committed to accumulating a pre-specified asset value at retirement). The goal is to optimize the investor's policy portfolio.

To specify the optimization objectives, we could follow Fischer Black's suggestion to minimize the present value of future contributions *and* Peter Bernstein's suggestion to maximize the probability that the commitment will be funded. In other words, the investor's goal is dual – to lower future contributions and improve the ability to fund the commitment at the same time. Given this dual optimization objective, the solution is to optimize the first measurement given the second and/or to optimize the second measurement given the first.

It is important to note that most investors revisit the issue of asset allocation regularly, and this practice should be reflected in the framework. Therefore, we should assume that the funding problem involves multiple asset allocation decisions.

This assumption has profound consequences. The choice of the policy portfolio made at the present is not binding in any way when the investor makes an asset allocation decision next time. Future portfolio selections cannot be mandated today. Rather, future portfolio selections should be viewed as today's *rational expectations* of the investor's behavior in the future. As a result, we have a "series" of investors – the investor at the present and the future hypothetical investors, which are the original investor's ageing "clones."[36]

What we have here is a set of "players," a set of actions (portfolio selections), and preferences (the lower cost the better). These components define a *strategic game*.[37] Under common conditions of rationality, the optimal strategy (policy portfolio) should form a *Nash equilibrium* solution to the problem of funding the commitment.[38]

Getting back to the Samuelson problem, the key question is whether the portfolios that form a Nash equilibrium solution should be the same. Generally, the answer is “no.” In particular, *a Nash equilibrium policy portfolio would not generally have the same portfolios even under the assumptions of constant risk tolerance and random-walk portfolio returns*. The missing component in the Samuelson framework may have been the concept of Nash equilibrium.

The wisdom of Fischer Black's emphasis on identifying sensible goals is on display here. Maximizing expected utility gives one set of optimal policy portfolios, and maximizing the ability to fund one’s financial commitment gives another. Dear reader, which goal makes the most sense to you?

Accepting "Without Question"?

"*... the image of the policy portfolio*
as the lodestar of asset allocation
has long been accepted without question ..."

Peter Bernstein (2003)

One of the key lessons of this story is the importance of asking the right questions. Even a seemingly innocuous question may illuminate issues of utmost importance. In many cases, a good question is already a major part of a good solution.[39]

Yet, good questions are not necessarily uncontroversial. Let us take another look at the observation Peter Bernstein made in 2007:

> "*Before Harry Markowitz’s 1952 essay on portfolio selection, there was no genuine **theory** of portfolio construction – there were just rules of thumb and folklore. It was Markowitz who first made risk the centerpiece of portfolio management ...*"

Dear reader, imagine that you had told a *pre-1952* investment manager that his approach to investing was based merely on "rules of thumb and folklore." Imagine also that you had told the manager the approach did not manage risk properly. You would likely find yourself on the receiving end of a strongly worded disagreeing tirade. Of course the manager's approach was based on a good theory and solid risk management. They all believed so at the time.

Obviously, Peter Bernstein had the benefit of more than half-a-century of retrospection. Besides, in all fairness, the development of "genuine" theories was hardly this investment manager's job. The manager was doing his best under the circumstances.

Now, dear reader, imagine that you are raising similar issues in a conversation with an investment manager *today*. Imagine that you have told the manager that his approach to investing is based merely on "rules of thumb and folklore" and lacks proper risk management. Similar to the prior imaginary conversation, you would likely find yourself on the receiving end of a strongly worded disagreeing tirade. Of course the manager's approach is based on a good theory and solid risk management. They all believe so.

But should we wait another half-a-century for a future expert to tell us – retrospectively – that we were "debating the wrong questions" and did not have a "genuine" theory of, for example, retirement investing? We should try to ask the right questions regarding certain practices accepted "without question" now, should we not?

Today, the practices accepted "without question" are in abundance. Some of these practices are very common, yet the foundations of these practices may contain plenty of "rules of thumb and folklore." For the reader's amusement, here is a partial list of these practices and related questions.

- For a DB plan, it is considered advantageous to match asset and actuarial and/or accounting liabilities. At the same time, DB plans should invest their assets in the best interest of plan participants. Is this asset-"liability" matching in the best interests of plan participants? Does asset-"liability" matching have anything to do with the safety of benefits and the cost of funding these benefits? Is there a "genuine" theory that demonstrates the benefits of asset-"liability" matching?

- The Liability Driven Investing (LDI) framework is getting increasingly popular among corporate DB plans. A common objective in the LDI framework is to improve the plan's funded status. Do policy portfolios generated by the LDI framework maximize the probability that the promised benefits will be paid? Do these policy portfolios minimize the cost of funding of the promised benefits? Is LDI in perfect compliance with ERISA?[40] What is the goal of LDI? Is there a "genuine theory" that serves as the foundation of LDI? Does anybody know what the letter "L" in the term "LDI" stands for?[41]

- Target date funds utilize evolving policy portfolios ("glide paths"). Do these glide paths maximize the participants' standard of living in retirement? Do these glide paths maximize the likelihood that a pre-specified standard of living will be funded? What is the goal of optimal glide path design? Is there a "genuine" theory of optimal glide path design?[42]

Dear reader, this author would like to encourage you to ask similar probing questions. You may discover that the wisdom of certain "rules of thumb and folklore" currently accepted without question is greatly exaggerated.

Conclusion

"*The emperor has no clothes!*"

Hans Christian Andersen (1837)

This paper is dedicated to Harry Markowitz, Fischer Black, Peter Bernstein, Paul Samuelson, and their steadfast commitment to the development of solid foundations in finance. This paper endeavors to illustrate one of the most important messages of their works – prudent investment practices require disciplined systematic approaches. "Rules of thumb and folklore" will not work in the long run.

When it comes to solid foundations and "genuine" theories, existing optimal portfolio selection methodologies contain major imperfections. The theoretical foundations of portfolio selection methodologies utilized by most DB plans do not withstand even a cursory examination. The theory of optimal glide path design for DC plan participants is in its infancy. Asset-"liability" analysis for foundations and endowments is virtually non-existent. MPT and its portfolio-centric modifications is essentially the only "genuine" theory available to these investors today.

One of this paper's objectives is to illuminate a decades-long quest for an investor-specific approach to optimal portfolio selection. As discussed in this paper, Harry Markowitz succeeded in developing a "genuine" practical *portfolio-centric* approach to optimal portfolio selection. Several attempts to develop a "genuine" practical *investor-specific* approach to optimal portfolio selection have not succeeded. Some of these approaches were not exactly "genuine" theories. Some other approaches were "genuine" but impractical theories. The quest for an investor-specific "genuine" practical approach to optimal portfolio selection continues to this day.

Yet, it is insufficient to point out that "the emperor" of optimal portfolio selection "has no clothes." It is also necessary to offer the emperor appropriate attire. Harry Markowitz not only tailored great attire for the emperor, but also set high standards for future tailors to follow. This attire has been gradually wearing off, but several attempts to offer better sartorial products have not satisfied the emperor's stringent requirements so far.

The attempts to create such an approach discussed in this paper have produced great ideas. Some of these ideas have been incorporated into the Commitment Driven Investing (CDI) framework. This author is optimistic that CDI will eventually be recognized as a practical "genuine" investor-specific theory and become a mainstream approach to funding financial commitments.

Finally, it is fair to inquire whether this author has had an epiphany that led to the formation of CDI. The answer is an unequivocal "no." This author has always viewed investing as a means to fund financial commitments and recognized the theoretical and practical advantages of stochastic present values. Moreover, I am not aware of any other "genuine" approach to funding financial commitments, thank you very much.

APPENDIX. Evolving Optimal Policy Portfolios in the CDI Framework

This section presents a simplified numerical example that demonstrates that the CDI framework generates *evolving optimal policy portfolios* and, therefore, is a solution to the Samuelson problem. The desire for simplicity and accessibility was the primary motivation behind the choice of this example and underlying assumptions. All calculations for this example fit in a modest size Excel worksheet.

Think of an investor with the commitment to make a payment of $100 in two years. The investor has to make two portfolio selections – one at the present and one at the beginning of the second year. Therefore, the investor's policy portfolio has two portfolios, and the goal is to optimize this policy portfolio. We assume that the investor uses two assets *A* and *B*; their risk/return assumptions are presented in *Exhibit 1*.

Exhibit 1

	Expected Return Geometric	Expected Return Arithmetic	Standard Deviation
Asset *A*	7.00%	8.03%	15.00%
Asset *B*	2.00%	2.005%	1.00%

We also assume that the correlation between assets *A* and *B* is zero, portfolio return factors are lognormally distributed, and portfolio returns in years 1 and 2 are independent.

Let us consider this funding problem in the "mean-variance" version of the CDI framework. As was discussed earlier in the paper, this version is analogous to MPT. For the reader's convenience, here is the definition of the *risk-adjusted expected return* that is instrumental in the development of the classic mean-variance efficient frontier.

Risk-Adjusted Expected Return =
Expected Return –
Standard Deviation of Return $\times$ Risk Aversion Factor

where the risk aversion factor is a positive constant. Portfolios that minimize the risk-adjusted expected return for all positive values of the risk aversion factor comprise the mean-variance efficient frontier.

The "mean-variance" version of the CDI framework employs a similar technique with one major exception. The key random variable – portfolio return – gets replaced by another random variable – the stochastic present value of the commitment, which is called *Required Assets* (*RA*) and defined as follows:

$$RA = \frac{100}{(1+R_1)(1+R_2)}$$

where R_1 and R_2 are the portfolio returns in years 1 and 2 correspondingly.

Similar to the risk-adjusted expected return, the *risk-adjusted expected cost* is defined as follows:

Risk-Adjusted Expected Cost =
Expected RA +
Standard Deviation of RA $\times$ Risk Aversion Factor

where the risk aversion factor is a positive constant.[43] Portfolios that minimize the risk-adjusted expected cost for all positive values of the risk aversion factor comprise the cost-risk efficient frontier.[44]

As a first step, we minimize the risk-adjusted expected cost for the risk aversion factor value of 0.3. The results are shown in *Exhibit 2*.

Exhibit 2

	Year 1	Year 2
Risk Aversion Factor	0.3	–
The Lowest Risk-Adjusted Expected Cost	94.16	–
Asset *A*	78%	78%
Asset *B*	22%	22%

As expected, the portfolios in year 1 and 2 are the same 78/22. It is important to note that optimization is performed only once, and the investor's risk tolerance in year 2 plays no role in the optimization procedure. Therefore, the risk aversion factor and risk-adjusted expected cost are not specified in year 2. In other words, the portfolio in year 2 is effectively mandated – the second "player" is not allowed to take any action.

Now let us assume that the investor's risk tolerance in year 2 does not disappear and the risk aversion factor has (for simplicity) the same value of 0.3. We have two "players," their actions, and preferences. The Nash equilibrium solution to the problem of funding the commitment is constructed via the process of *backward induction*.[45] In other words, the year 2 portfolio is generated first, followed by the year 1 portfolio.

Specifically, the stochastic present value of the commitment RA_2 for year 2 is defined as follows:

$$RA_2 = \frac{100}{1+R_2}$$

where R_2 is the portfolio return in years 1. The 42/58 portfolio minimizes the risk-adjusted expected cost for year 2 (97.76). Next, the stochastic present value of the commitment RA for years 1 and 2 is defined as follows:

$$RA = \frac{100}{(1+R_1)(1+R_2^*)}$$

where R_1 is the return for a not-yet-determined portfolio in year 1 and R_2^* is the return for the 42/58 portfolio in year 2. The 63/37 portfolio minimizes the risk-adjusted expected cost at the present (94.44). These results are summarized in *Exhibit 3*.

Exhibit 3

	Year 1	Year 2
Risk Aversion Factor	0.3	0.3
The Lowest Risk-Adjusted Expected Cost	94.44	97.76
Asset *A*	63%	42%
Asset *B*	37%	58%

Thus, in the CDI framework, optimal policy portfolios do not have the same portfolios even under the assumptions of constant risk tolerance and random-walk portfolio returns.

REFERENCES

Arnott, R., Bernstein, P. [2002]. What Risk Premium Is "Normal"? *Financial Analysts Journal*, March/April, 2002, 64–85.

Aumann, R. J. [1995]. Backward Induction and Common Knowledge of Rationality, *Games and Economic Behavior*, 8, 1995.

Bernstein, P. [2003A]. Are Policy Portfolios Obsolete? *Economics and Portfolio Strategy*, March 1, 2003.

Bernstein, P. [2003B]. Which Policy Do You Mean? *Economics and Portfolio Strategy*, August 15, 2003.

Bernstein, P. [2007]. Capital Ideas Evolving, *John Willey & Sons*, 2007.

Black, F. [1980]. The Tax Consequences of Long-Run Pension Policy, *Financial Analysts Journal*, July/August 1980, 21–28.

Black, F. [1989]. Should You Use Stocks To Hedge Your Pension Liability? *Financial Analysts Journal*, January/February 1989, 10–12.

Black, F. [1995]. The Plan Sponsor's Goal, *Financial Analysts Journal*, July-August 1995, 6-7.

Black, F., Litterman, R. [1992]. Global Portfolio Optimization, *Financial Analysts Journal*, September/October, 1992, 28–43.

Kellison, S. [2009]. The Theory of Interest, Third Edition, *McGraw-Hill Irwin*, 2009.

Kneafsey, K. [2003]. Solving the Investor's Problem, *Investment Insights*, The Investment Research Journal from Barclays Global Investors, Volume 6, Issue 5, August 3, 2003.

Leibowitz, M., Bader, L., Kogelman, S. [1996]. Return Targets and Shortfall Risks, Studies in Strategic Asset Allocation, *Irwin*, 1996.

Loeb, J. [1950]. The Battle for Investment Survival, *John Wiley and Sons*, New York, 1950.

Mehrling, P. [2005]. Fischer Black and the Revolutionary Idea of Finance, *Wiley*, 2005.

Markowitz, H. [1952]. Portfolio Selection, *Journal of Finance*, Vol. 7, No. 1, 1952.
Markowitz, H. [1959]. Portfolio Selection: Efficient Diversification of Investments, *John Wiley and Sons*, New York, 1959.
Merton, R. C., [1969]. Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case, *Review of Economics and Statistics* 51, 247-257, 1969.
Mindlin, D. [2009A]. Commitment Driven Investing and Time Diversification, *CDI Advisors Research*, CDI Advisors LLC, 2009, http://www.cdiadvisors.com/papers/CDIandTimeDiversification090406.pdf.
Mindlin, D. [2009B]. Commitment Driven Investing: An Introduction, *CDI Advisors Research*, CDI Advisors LLC, 2009, http://www.cdiadvisors.com/papers/CDIIntroduction091109.pdf.
Mindlin, D. [2009C]. The Case for Stochastic Present Values, *CDI Advisors Research*, CDI Advisors LLC, 2009, http://www.cdiadvisors.com/papers/CDITheCaseforStochasticPV.pdf.
Mindlin, D. [2009D]. "The Plan Sponsor's Goal" Revisited, *CDI Advisors Research*, CDI Advisors LLC, 2009, http://www.cdiadvisors.com/papers/CDIThePlanSponsorsGoalRevisited091124.pdf.
Mindlin, D. [2010A]. Principles of Optimal Glide Path Design, *CDI Advisors Research*, CDI Advisors LLC, 2010, http://www.cdiadvisors.com/papers/CDIOptimalGlidepathPrinciples.pdf.
Mindlin, D. [2010B]. Is LDI Legal?, *CDI Advisors Research*, CDI Advisors LLC, 2010, http://www.cdiadvisors.com/papers/CDIIsLDILegal.pdf
Mindlin, D. [2011]. LDI vs. CDI: Bitter Foes or Best Friends?, *CDI Advisors Research*, CDI Advisors LLC, 2011, http://www.cdiadvisors.com/papers/CDILDIBitterFoesOrBestFriends.pdf.
Osborne, M. J., [2004]. An Introduction to Game Theory, *Oxford University Press*, 2004.
Rabin, M., Thaler, R. [2001]. Anomalies: Risk Aversion, *Journal of Economic Perspectives*, 15, 219-32, 2001.
Roy, A. [1952]. Safety First and the Holding of Assets, *Econometrica*, Vol. 20, July, 1952.
Samuelson, P. [1969]. Lifetime portfolio selection by dynamic stochastic programming, *The Review of Economics and Statistics* 51, 239-246, 1969.
Samuelson, P. [1970]. What Makes for a Beautiful Problem in Science? *Journal of Political Economy*, December, 1970.
Samuelson, P. [1973]. Proof That Properly Discounted Present Values of Assets Vibrate Randomly, *The Bell Journal of Economics and Management Science*, 369-374, Autumn, 1973.
Samuelson, P. [1989]. A Case At Last for Age-Phased Reduction in Equity, *Proc. Natl. Acad. Sci. USA*, Vol. 86, 9048-9051, 1989.
Samuelson, P. [1991]. Long Run Risk Tolerance When Equity Return Are Mean Regressing: Pseudoparadoxes and Vindication of "Businessman's Risk", In Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin, Edited by W. C. Brainard, W. D. Nordhaus, H. W. Watts , *MIT Press*, 1991.
Samuelson, P. [1992]. At Last, a Rational Case for Long-Horizon Risk Tolerance and for Asset-Allocation Timing? Active Asset Allocation, edited by Arnott, R.D. and Fabozzi, F.J., *Probus Publishing Company*, 1992.
Sharpe, W. F. [1976]. Corporate pension funding policy, *Journal of Financial Economics*, 3: 183-93, 1976.
Tepper, I. [1981]. Taxation and corporate pension policy, *Journal of Finance*, 32:627-38, 1981.
Treynor, J. [1977]. The principles of corporate pension finance, *Journal of Finance*, 32:627-38, 1977.

ENDNOTES

[1] Unusual, but not unique. The foundations of probability theory were laid down in the 1930s even though people have been making probability based inferences since ... well, always.
[2] In a recent interview, Harry Markowitz gave a detailed account of his thinking at the moment of the epiphany. The interview was published in the *Journal Of Investment Management*, Fourth Quarter, 2011.
[3] See Bernstein [2007], page *xii*. Throughout this paper, emphasis in all quotes is presented as in the original.
[4] Peter L. Bernstein Special Commemorative Section, *Journal of Portfolio Management*, Summer 2009.
[5] In 1934, John Maynard Keynes, in a letter to a business associate, wrote the following:

> "*As time goes on, I get more and more convinced that the right method in investment is to put fairly large sums into enterprises which one thinks one knows something about and in the management of which one thoroughly believes. It is a mistake to think that one limits one's risk by spreading too much between enterprises about which one knows little and has no reason for special confidence*."

More than half a century after this letter and four decades after Markowitz [1952], Warren Buffett expressed similar sentiments and included this quote in his letter to the shareholders of Berkshire Hathaway Inc. in 1992 (see at http://www.berkshirehathaway.com/letters/1991.html).

Loeb [1950] held similar views:

> "*Once you obtain confidence, diversification is undesirable; diversification [is] an admission of not knowing what to do and an effort to strike an average*."

[6] It should be mentioned that Roy [1952] advanced somewhat similar ideas but was published several months after Markowitz [1952]. The Safety-First approach proposed in Roy [1952] utilized the "disastrous" return and the probability of exceeding this return, which can be considered as two measurements of portfolio return. The principle "optimize the first measurement given the second and to optimize the second measurement given the first" is perfectly applicable to this approach. This principle would lead to the Safety-First efficient frontier.
[7] These developments are presented in Leibowitz-Bader-Kogelman [1996].
[8] Sadly, the paper was truly final – Fischer Black passed away around the time of the publication of Black [1995].
[9] See Black-Litterman [1992].
[10] The key ideas of Black [1995] appeared in internal memos at Goldman Sachs in the early 1990s.
[11] The September-October 2008 issue of the *Financial Analyst Journal* (*FAJ*) contains two letters regarding Black [1995]. The first letter written by C. Kerwin hypothesized that the abstract of the article "*must have been written in error by someone at the FAJ*." In the second letter, R. Ennis, the editor of the *FAJ* at the time, agreed with Mr. Kerwin's hypothesis that "*then-editors either added the abstract or modified its wording*." This hypothesis did not withstand a close scrutiny. The May/June 2009 issue of the *FAJ* contains a letter written by Prof. P. Mehrling (Barnard College, Columbia University) and this author that demonstrated that the article in its entirety is the work of Fischer Black. Mr. Ennis published an elegantly worded retraction in the same issue.
[12] Actually, $100 in stocks is not the same as $100 in bonds even in the portfolio-centric framework. For example, does anybody believe that a derivative-based downside protection for the all-bonds and all-stocks portfolios would have the same cost?
[13] See endnote 8.
[14] It appears that Fischer Black had anticipated that he would have been unable to fully participate in the process of editing Black [1995]. In the letter to the editor that accompanied the submission of Black [1995], Fischer Black wrote, "*I would like to publish this, though I may not be around to make any changes the referee may suggest. If I'm not, and if it seems roughly acceptable, could you publish it as is with a note explaining the circumstances?*" (Emanuel Derman, *Financial Analysts Journal*, May/June, 2006, p. 75-76.)
[15] See Mindlin [2009D] for a detailed analysis of Black [1995].
[16] For example, see Kellison [1991], chapter 10.
[17] See Mindlin [2009C] for more details.

[18] How did young Harry Markowitz avoid all these complications with present values? He started with expected present values as well, did he not? Yes, he did, but Harry Markowitz switched from a *present value* based multi-period framework to a *future value* based single-period framework. As discussed in the previous section, these frameworks are fundamentally different.
[19] See Mehrling [2005], p. 222.
[20] See Peter L. Bernstein, The 60/40 Solution, *Bloomberg Personal Finance*, January/February 2002.
[21] See Arnott-Bernstein [2002].
[22] See Bernstein [2003B].
[23] Specifically, if the existing asset value is equal to the *P*th percentile of the present value of future payments, then the probability of funding is equal to *P*. See Mindlin [2009C] for more details.
[24]P. Bernstein, The Flight of the Long Run, *Financial Times*, February 25, 2009, http://www.ft.com/cms/s/0/267900e0-0360-11de-b405-000077b07658.html#axzz1rs7J3Pla.
[25] See Samuelson [1969]. Robert Merton got the same result in the continuous-time case, see Merton [1969].
[26] From the interview conducted for the History of Finance project sponsored by the American Finance Association in 2005. This interview is quoted according to the transcript posted at http://www.afajof.org/association/historyfinance.asp.
[27] "*... a defiance of folk wisdom and casual introspection*" is quoted from Samuelson [1989]. "*I felt that the first pass at the problem was a failure*" is quoted from the Interview.
[28] It should be mentioned that this problem is closely related to the "time diversification problem" – whether or not time "diversifies" the riskiness of equities. One's views on the riskiness of equities depends on one's preferred risk measurements. For instance, the volatility of *annualized* returns may *decrease* with time; the volatility of *compounded* returns may *increase* with time; the *likelihood* of loss may *decrease* with time; the *magnitude* of loss may *increase* with time. There is no universal definition of risk, so different market participants may perceive the (lack of) benefits of time diversification differently.

If one believes that time "diversifies" the riskiness of equities, then the young by definition have more time and should hold more equities than the old. However, this logic does not necessarily work in reverse. If one does not believe that time "diversifies" the riskiness of equities, then the question about the optimal allocation of equities is still open. It may still be a good idea for the young to hold more equities regardless of the time diversification properties of equities or lack thereof. Or it may not.
[29] See the 2005 interview, page 2, see endnote 26.
[30] See the 2005 interview, page 2, see endnote 26.
[31] See Samuelson [1970].
[32] The calculus of variations became a rigorous mathematical theory largely due to the work of Karl Weierstrass.
[33] See endnote 6. Remarkably, Roy [1952] re-emerges in the area seemingly unrelated to MPT.
[34] Other economists have questioned the used of expected utilities as well. For example, Rabin-Thaler [2001] argues that "*… it is time for economists to recognize that expected utility is an ex-hypothesis, so that we can concentrate our energies on the important task of developing better descriptive models of choice under uncertainty*."
[35] There is a somewhat hidden assumption that we may want to spell out here. We assume that the investor endeavors to fund the commitment via investing future contributions in *risky assets*. Due to the presence of risky assets in the investor's policy portfolio, the funding of the commitment is not certain, and the probability of shortfall is normally greater than zero.
[36] The investor and the investor's ageing "replicas" is called *the investor's cohort* in Mindlin [2009A].
[37] For example, see Osborne [2004].
[38] See Mindlin [2009A] for more details.
[39] As an example, think of the main message of Black [1980] – corporate pension plans should have all-bond policy portfolios. From the "corporate pension finance" perspective, this policy portfolio makes sense because it maximizes the value of the "tax arbitrage." But is the plan sponsor's primary goal to maximize the value of the "tax arbitrage"? Do you think your ERISA attorney would be happy to see this goal in your investment policy statement? Suddenly the utility of the all-bond policy portfolio is not so obvious.
[40]See Mindlin [2010B] for more details regarding the issues of LDI compliance with ERISA.
[41] Mindlin [2011] discusses the problems with LDI and the advantages of the CDI framework over LDI.

---

[42] Mindlin [2010A] discusses the principles of optimal glide path design and demonstrates that Nash equilibrium glide paths maximize the participants' standard of living in retirement.

[43] Note that Standard Deviation of Return $\times$Risk Aversion Factor is *subtracted* from Expected Return in the calculation of Risk-Adjusted Expected Return. At the same time, Standard Deviation of *RA* $\times$Risk Aversion Factor is *added* to Expected *RA* in the calculation of Risk-Adjusted Expected Cost. This distinction is a reflection of the directional difference between MPT and CDI. In MPT, it is desirable to have *high* returns. In CDI, it is desirable to have *low* cost.

[44] See Kellison [2009] or Mindlin [2009C] for the calculations of the moments of *RA*.

[45] R. Aumann, the 2005 Nobel laureate in economics, proved that, under certain conditions, common knowledge of rationality implies backward induction. See Aumann [1995].

Important Information