Contrasting Client and Manager Perception of Risk: Is Modern Portfolio Theory Harming Investors?

Nathan Eigerman, CFA

Abstract

As generally used and accepted by most practitioners, Modern Portfolio Theory makes simplifying assumptions that are inaccurate with regard to most investors' perception of risk. This can be demonstrated easily and was well appreciated by early users and developers of the theory. Nevertheless, the investment management industry has acted as if MPT's modeling of risk were ideal and unbiased, which has led to sub-optimal products and inadequate risk management solutions.

Introduction

Perhaps the best definition of Modern Portfolio Theory is that it is the practical application of certain financial ideas and models developed in the mid-20th Century. Evolving over decades, and based on the contributions of many authors, MPT lacks both a single inventor and a single moment of invention. Harry Markowitz, who is widely acknowledged to have the strongest claim to being its father, does not himself use the term, preferring simply "portfolio theory" and allegedly quipping that there is nothing modern about it.

As MPT emerged, so did an extensive literature on its limitations. By the 1970s, a range of issues, from MPT's implicit and problematic assumption of static correlations between securities, to difficulties in empirically confirming the Capital Asset Pricing Model, were well understood and well discussed by researchers.[1] Partly in response, alternative variations of MPT were developed, such as Arbitrage Pricing Theory [2] and the Black-Litterman Model.[3] There is even a Post-Modern Portfolio Theory.[4]

Nevertheless, a conventional and mainstream Modern Portfolio Theory has become fundamental to investment management as practiced today. It defines the framework that most professionals use to evaluate investments, and as a result defines the criteria that the industry uses to measure its own success or failure. And although all practitioners may not be conversant with the subtleties of its reasoning, MPT provides the intellectual justification for such prominent features of the investing

[1] For a useful survey of the first half century of MPT, see Edwin J. Elton and Martin J. Gruber, "Modern portfolio theory, 1950 to date", Journal of Banking & Finance 21 (1997) 1743-1759

[2] Ross, Stephen (1976). "The arbitrage theory of capital asset pricing". Journal of Economic Theory 13 (3): 341–360. [3] Black F. and Litterman R.: Asset Allocation Combining Investor Views with Market Equilibrium, Journal of Fixed Income, September 1991, Vol. 1, No. 2: pp. 7-18

[4] Rom, B. M. and K. Ferguson. "Post-Modern Portfolio Theory Comes of Age." Journal of Investing, Winter 1993.

landscape as capitalization weighted indexes, the importance of benchmark relative returns, and Sharpe Ratios.

In the context of the broad use of a framework that is simultaneously understood to have meaningful limitations, it is valid to consider if there are aspects of MPT that have resulted in a sub-optimal output from the investment industry. Put another way, has faith in MPT kept investment professionals from doing the best that they can for their clients?

Measuring Risk

In his seminal 1952 paper "Portfolio Selection"[5] Harry Markowitz derived the principle that at a given level of risk investors will prefer investments with higher expected return and that at a given expected return investors would prefer lower risk. From this simple and appealing dual principle was developed what we now call the Efficient Portfolio Frontier and the critical insight that of all possible portfolios of securities, only a very small number are potentially desirable to an investor. All others are demonstrably and rationally inferior. Much of MPT, the Capital Asset Pricing Model in particular, stems from this construct.

But to build this elegant structure from basic principles, Markowitz and those that followed had to make two, largely implicit, assumptions.[6]

1. All investors perceive and measure risk identically.
2. Risk can be measured as the variance (or standard deviation) of possible outcomes around the mean expected return.

The first of these says nothing of risk tolerance, only that given two or more potential investments all investors will assign the same rank order by risk. The second, which relies on the first, has perhaps broader implications and has been the subject of some uneasiness from the beginning. Using variance or its square root, standard deviation, what would later come to be called volatility in the financial world, as the single definitive measure of risk was understood, even by Markowitz, to be a weak assumption. The 1952 paper suggests only that seeking to minimize risk as measured in this manner would be "reasonable as a working hypothesis and a working maxim" for investors.

Finding measures of risk that are at least as intuitively appealing as variance is not challenging. The probability of significant loss, for example, would strike many as a more direct indicator of the sort of risk that most investors wish to minimize. And there exists an entire category of risk measures that build on the observation that variability in gains is not as unpleasant for investors as variability in losses. These include semi-variance, a variance measure estimated only from negative (left side) observations, value at risk (VaR), essentially an estimate of the worst outcomes at the narrow end of the left tail, and maximum drawdown, the largest peak to trough price movement for an investment over a defined time period.

---

[5] Markowitz, H.M. (March 1952). "Portfolio Selection". The Journal of Finance 7 (1): 77–91

[6] In fairness, although #1 is strongly implied in Markowitz (1952) it is not necessary for the conclusions of the paper. In contrast, it is a fundamental premise of the market pricing work that would follow.

Ultimately, the choice of variance as the solitary measure of risk owes more to practical considerations than theoretical soundness. In the years in which MPT was developed, computing power was scarce and, by 21st Century standards, unimaginably expensive. Asymmetric measures of risk such as semi-variance or VaR might have made more theoretical sense, but were too complex computationally for actual use in portfolio construction.

In his 1990 Nobel Prize Lecture[7] Markowitz made this motivation, and his acknowledgement of the resulting imperfections, clear.

> "So, equipped with databases, computer algorithms and methods of estimation, the modern portfolio theorist is able to trace out mean-variance frontiers for large universes of securities. But, is this the right thing to do for the investor? In particular, are mean and variance proper and sufficient criteria for portfolio choice?
>
> … We seek a set of rules which investors can follow in fact - at least investors with sufficient computational resources. Thus, we prefer an approximate method which is computationally feasible to a precise one which cannot be computed. I believe that this is the point at which Kenneth Arrow's work on the economics of uncertainty diverges from mine. He sought a precise and general solution. I sought as good an approximation as could be implemented."

Like all economic models, MPT includes simplifications designed as useful assumptions rather than perfect analogs, "an approximate method which is computationally feasible." Inevitably, there are cases in which these types of approximations are inaccurate. That fact does not, in principle, invalidate a model. But it does raise questions when the model, useful approximations included, is used as if it were literally true.

A Simple Experiment on Investor Perceptions of Risk

To highlight the inconsistencies between investor perceptions of risk and MPT's modeling of investor perceptions of risk, consider the following simple game. Participants are asked to choose between two alternative schedules of payoffs, each based on the random selection of a number from one to ten. At the lower numbers the participant would lose money and at the higher numbers the participant would gain money. The payoffs are constructed such that the average outcome for both choices is a gain of $400.

[7] Harry M. Markowitz – Nobel Prize Lecture: Foundations of Portfolio Theory, December 7, 1990

	A	B
1	-$1,000	-$600
2	-$500	-$500
3	-$200	-$400
4	$0	$0
5	$0	$0
6	$500	$400
7	$800	$400
8	$1,000	$600
9	$1,400	$1,800
10	$2,000	$2,300
Mean	$400	$400
Standard Deviation	$915	$965

Using informal surveys of investors, financial advisors, and employees of an investment management firm, it has been found that a consistent majority, generally between 70% and 80%, prefer payoff schedule B over A. The reasoning of those that choose B is also fairly consistent. B has a less objectionable worst possible outcome than A, a loss of $600 rather than $1000, and more favorable best outcomes, with gains of $1800 and $2300 on numbers 9 and 10 rather than $1400 and $2000.

And yet, as shown on the last line of the table, schedule A has a lower standard deviation than B. Thus, a majority of individuals directly contradict Markowitz's seminal work when playing this game. Given two potential investments with identical expected (i.e. mean) returns, most subjects prefer the investment with the higher level of risk.

Higher, that is, if standard deviation (or variance) is accepted as the definition of risk. This game, it appears, is a case in which the variance approximation is meaningfully inaccurate. Markowitz's basic principle, that all else equal investors prefer less risk, is undoubtedly sound. But most investors do not perceive B as being more risky than A. Indeed, this experiment demonstrates weakness in both of the implicit assumptions enumerated above. Not only do most investors choose the payoff with higher variance, but different, apparently otherwise rational, investors choose different payoff schedules, meaning that they have different perceptions or measures of risk.

The use of variance or standard deviation as the solitary measure of risk is often justified by its more thoughtful defenders as a convenient approximation that is, in fact, very closely correlated with many of the other more intuitively appealing risk measures. This is a powerful argument when applied to individual equity securities, the investment which many practitioners and academics most readily use as an example of a risky investment. For a single stock, it is very difficult to reject the null hypothesis that returns are normally distributed.[8] And if returns are normally distributed, then the downside risk

[8] To be more precise, that the log of returns are normally distributed.

measures such as semi-variance, VaR, and maximum drawdown are all linear functions of variance or standard deviation.

The flaw in this argument is that, as MPT itself demonstrates, the risk of individual securities is of comparatively little concern to investors. The great majority of investing, be it done by consumers saving for retirement or trustees managing an endowment, is in actively managed portfolios of investments. And in the realm of portfolios, return distributions are less a matter of empirical observations of natural phenomena than of conscious design by portfolio managers.

Buy and Hold Revisited

One of the least frequently questioned doctrines of mainstream investing is known popularly as buy and hold. This is the precept that an investor ought to select a portfolio, classically a capitalization weighted index, and hold that portfolio consistently over as extended a period of time as possible. Arguments in favor of this approach can include advantages such as a reduction in transaction costs and taxes, as well as the often cited, and MPT related, principle that it is impossible to "time" the market. (It might also be observed that buy and hold has a particular attraction for asset managers that continue to receive fee income as long as the customer is invested.)

If risk is defined as standard deviation, buy and hold has a facile justification. Returns compound over time, so as the time period increases expected return grows geometrically.

$$r_t = (1 + r)^t$$

In contrast, the standard deviation of returns grows only with the square root of the time period. That is, the standard deviation of a four year period is just twice that of a one year period.

$$\sigma_t = \sigma\sqrt{t}$$

Thus, assuming that expected annualized return and volatility are constants, as the holding period for an investment increases the expected Sharpe Ratio for the period (r_t/σ_t), MPT's definitive standard of investment desirability, will also increase. Combined with commonplace assumptions about the difficulty of making short-term predictions of returns, this provides considerable justification for buy and hold.

In more simplistic terms, investors are often persuaded to subscribe to a buy and hold strategy with reassurances that although the market has bad as well as good years, there are more good ones than bad ones, so that over a suitably long period returns will average out to a satisfactory level. While that has a strong common sense appeal, investors who examine market history are less likely to be persuaded.

Even over periods as long as 25 years, returns from US equities exhibit meaningful variability. To cite examples from the relatively recent past, the S&P 500 gained an average of 13.04% annually over the 25 years ending 1999, but only 6.81% over the 25 years ending 2011. Investors employing a buy and hold approach with the S&P starting in 1974 would have enjoyed a more than 20-fold gain, while the same

strategy begun twelve years later would have resulted in just a 4-fold gain. Looking farther back, over the 25 years ending 1981 the S&P averaged just 3.94% annually, less than the inflation rate, which averaged 5.02% during the same period.[9]

Further, for many investors, 25 years is an impractically distant time horizon. On a ten year basis, the returns for the S&P appear even more unstable. Of the 76 overlapping ten year periods since 1928, 11 of them, or more than one in seven, exhibit negative average returns.

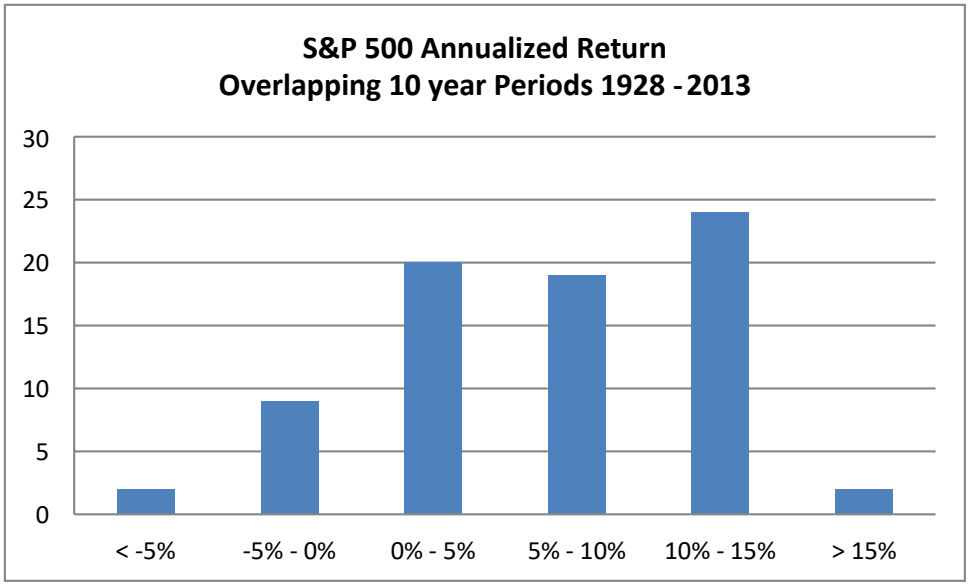

The relative risk reduction afforded by buy and hold, what professionals often consider to be a compelling and obvious benefit, is often elusive from the point of view of investors. At the core of this dissonance between the investment management industry and its customers is a differing perception of risk. If the riskiness of an investment is standard deviation, then the risk benefit of buy and hold is axiomatic. But under other reasonable definitions of risk the effects of buy and hold are much less clear. For example, if risk is defined as the probability of suffering a significant loss, the relationship between holding period and risk is indeterminate.[10] Using this definition, there are situations in which a longer period results in higher risk rather than lower.

[9] All S&P returns are price only, sourced from Standard and Poor's via FactSet. Inflation is based on the US Consumer Price Index.

[10] Classically trained investment professionals may want to convince themselves of this algebraically. For intuition, consider the following question. For which time period is experiencing a loss of 20% or more in the S&P 500 most likely: a day, a week, a month, or a year?

Controlling Risk with Risk Models

The systematic management of risk is arguably the *sine qua non* of serious professional investing, what practitioners might identify as the difference between picking stocks and managing a portfolio. For most professionals the risk in their products is estimated using a multi-factor risk model, a framework that uses reasonable assumptions and sophisticated math to generate precise projections of volatility.

Today's models, including widely used commercially available versions, are generally well constructed and generally quite accurate. But they often suffer from a significant flaw. Although they work well the great majority of the time, they can become significantly inaccurate during times of market stress. This is a broadly recognized phenomenon in the investment world and was particularly evident during the 2008-09 financial crisis. An eloquent argument as to why this might be was offered by Alan Greenspan in March of 2008.

> The most credible explanation of why risk management based on state-of-the-art statistical models can perform so poorly is that the underlying data used to estimate a model's structure are drawn generally from both periods of euphoria and periods of fear, that is, from regimes with importantly different dynamics.
>
> The contraction phase of credit and business cycles, driven by fear, have historically been far shorter and far more abrupt than the expansion phase, which is driven by a slow but cumulative build-up of euphoria. Over the past half-century, the American economy was in contraction only one-seventh of the time. But it is the onset of that one-seventh for which risk management must be most prepared. Negative correlations among asset classes, so evident during an expansion, can collapse as all asset prices fall together, undermining the strategy of improving risk/reward trade-offs through diversification.[11]

Crucially, the comparatively rare periods of fear and dislocation, when risk models tend to fail, are the periods in which investors most value risk management. Conversely, the long periods of relative market calm, when the models work well, are times when investors are relatively unconcerned about their risk exposure. As David Einhorn of Greenlight Capital put it in 2009, writing of Wall Street's pre-crisis VaR estimation methodology, "This is like an air bag that works all the time, except when you have a car accident."[12]

To most investors, a risk model that tends to fail when it is most needed is of questionable value. And yet the investment industry continues to rely on models that do just that. Can it be that the professional investing community is unaware of this weakness in its models?

Once again, the conflict between the actions of the investment management industry and the desires of its clients stems from differing definitions of risk. If risk is the standard deviation of returns, and the goal of a portfolio manager is to minimize that form of risk over time, then current industry practice is ideal. Although risk may be inadequately managed in the small minority of cases when capital markets fall

[11] Greenspan, Alan, "We Will Never Have a Perfect Model of Risk", Financial Times, March 16, 2008.

[12] Nocera, Joe, "Risk Mismanagement" The New York Times, January 2, 2009

dramatically, during the majority of the time the models are a powerful tool in controlling volatility. As a result, the models can be said to enable a significant reduction in the average level of risk over time, when risk is defined as the variance or standard deviation of returns.

Yet to an investor concerned not with risk as described by standard deviation but with the danger of catastrophic loss, risk management during normal times, during Greenspan's "euphoria" years, is largely irrelevant. To this investor, it is risk management during the bad times, the fear dominated years, which matters. While the established paradigm implicitly weights the more common normal periods more heavily, many investors would likely prefer just the opposite. They would prefer a risk measure that emphasized the experience in the worst periods, even if that meant weakening risk management in the good periods.

Many investment managers are adept at reducing the risk in their portfolios when times are good, which is to say that they are skilled at reducing already low levels of risk even lower, but have difficulty controlling higher levels of risk. An investor might prefer an approach that did the reverse, that was poor at controlling risk most of the time but superior at controlling it during a crisis. However, to an investment industry focused on average standard deviation as the measure of risk, such an approach would appear inferior and undesirable.

Is MPT Harmful?

There can be no doubt that the body of work known as Modern Portfolio Theory is one of the great intellectual and financial achievements of the 20th Century. The widespread use of risky financial assets as a store of wealth in the developed world owes a great deal to MPT. And the benefits of systematic and scientific investment management go beyond savers to the global economy as a whole, allowing riskier and more entrepreneurial projects to be undertaken. MPT is not merely elegantly insightful. It has had a powerful practical impact.

In some ways, it is this success that makes MPT dangerous for investment managers. While the limitations of its enabling assumptions may have been well understood by its first users, inevitably as the decades passed those limitations became forgotten. As each new generation of investment professionals joined the business, MPT became increasingly treated more as dicta than as a "working maxim."

In particular with regard to the measurement of risk, MPT has always rested on assumptions that are only generally correct, a fact that was much more obvious in the early years of the theory than it is today. Of course, those assumptions are valid and practical in a great many contexts, as could only be the case given their widespread use. But there are situations in which the assumptions mislead the unwary into less than optimal policies. This paper has highlighted two.

Moreover, because MPT has never been merely a description of how investors behave but also a normative statement about how investors should ideally behave, this particular inaccuracy by assumption has had a wide ranging impact on the investment management business. Investment professionals have, for decades, been designing products that minimize investment risk as roughly

approximated in 1952. Although this was done with the best of intentions, it was not, in retrospect, the best that could have been done for clients.