January 2019

BUILDING A BETTER ASSET ALLOCATION PORTFOLIO

Introducing a Momentum Factor to Modern Portfolio Theory

Executive Summary

While the traditional approach to Modern Portfolio Theory has historically relied on long-term averages to build diversified portfolios, we are proposing a new model – one that continually focuses on the realities of an ever changing global market.

In this paper, we suggest that effective diversification comes not from the *number* of positions that are held, but from the *current correlation* of the underlying positions. We will also examine a new model we call TrueDiversification that attempts to maximize the benefits of diversification by continually targeting a portfolio with the optimal risk, return, and correlation metrics. And we will examine the methodology behind the process that systematically and continuously updates portfolio characteristics to match current market conditions. In short, we present a new model for asset allocation investors that we believe will provide a dramatically improved investment experience.

Daniel Snover
Vice President, Deputy CIO

William Davis
Principal, Chief Marketing Officer

Contents

Introduction

Investors should seek the benefits of diversification when selecting an investment portfolio. In 1952, Harry Markowitz illustrated how diversification can reduce risk without sacrificing returns in his Modern Portfolio Theory (MPT).

Based on the basic tenets of Markowitz's Theory, traditional asset allocation models took the investment scene by storm in the early 2000s. The stated goal of the offerings was to capture diversification benefits of a multi-asset strategy while eliminating the company specific risk that harmed so many during the tech-bust of 2001-2003.

Regrettably, the 60-year old MPT approach relies on static, backward-looking assumptions to inform broad asset allocation decisions. Said another way, MPT keeps investors locked into investment decisions determined by historical results, which, by definition, will not have adapted for present shifting market conditions.

As a result, investors have been underwhelmed by the performance of traditional asset allocation portfolios, in large part due their failure to provide protection in volatile down markets. The clearest example, of course, is the Financial Crisis of 2008, when even conservative allocation strategies fell almost as much as the overall equity market. Worse, perhaps, the same portfolios lagged considerably during the subsequent equity bull market.

Worthy investment strategies offer a plan to limit downside with strong gain potential, not strong loss potential with limited upside.

In pursuit of an elegant solution to this messy problem, we have found that the estimation of decision-making inputs can be improved by acknowledging a persistency factor in asset prices, otherwise known as momentum. We are suggesting here that the addition of a momentum factor to the allocation process allows for mean-variance optimization that effectively aids in creating the optimal allocation of portfolio assets.

This paper is will outline the TrueDiversification process used to construct the Fund Architects Blueprint Portfolios. Starting with Mean Variance Optimization as the decision-making framework, several modifications and assumptions are made to create a unique model for maximizing the benefits of diversification.

*****

Portfolio Framework

Markowitz's **Modern Portfolio Theory** described a process called *mean-variance optimization (MVO)*, a mathematical framework for assembling a portfolio of assets such that the expected return is maximized for a given level of risk. It is a formalization of diversification in investing, the idea that owning different kinds of financial assets is less risky than owning only one type.

The theory's key insight is that an asset's risk and return should not be assessed by itself, but by how it contributes to a portfolio's overall risk and return.

The basic risk and return trade-off can be seen in a two-asset portfolio. The inputs are the expected return and standard deviation of each asset, and the correlation matrix between each asset:

Asset	Expected Return	Standard Deviation	Correlation Matrix	
			Asset 1	*Asset 2*
Asset 1	8%	5%	1	-1
Asset 2	15%	15%	-1	1

To illustrate the benefits of diversification, it's easiest to assume the two assets in the portfolio have perfect negative correlations to each other. If the return of Asset 1 is 5% over a given period, it is expected that the return of Asset 2 be -5% over that same period.

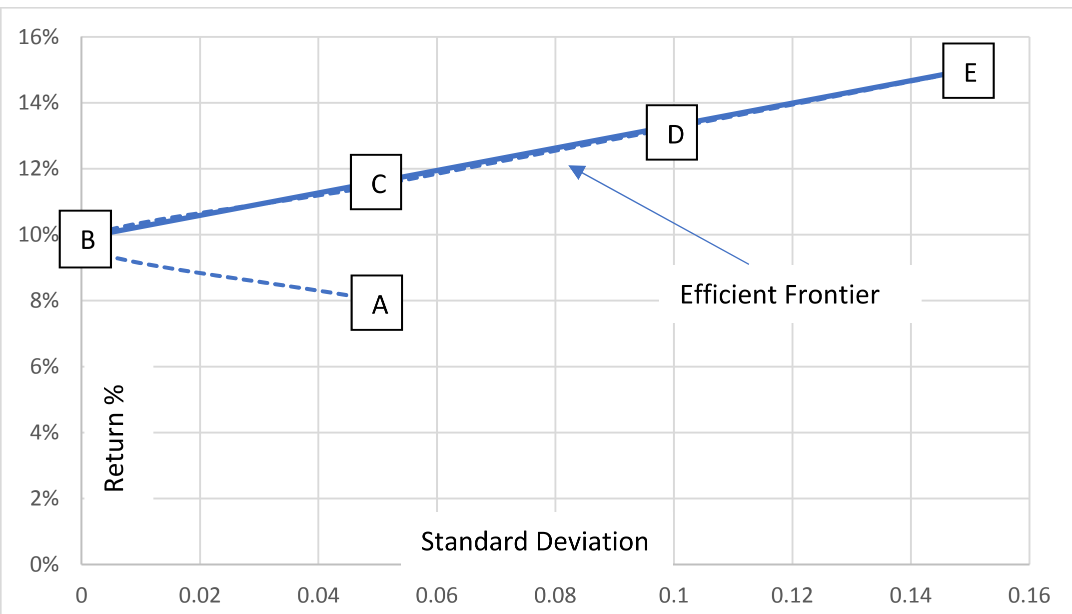

Starting from Portfolio A, which is 100% invested in Asset 1, we begin to include Asset 2 at 25% increments. At Portfolio B, we have not only increased the expected return, but *the expected variance has dropped to zero*. This is only possible when two assets are perfectly negatively correlated, and Portfolio B is therefore assumed to have no risk.

The **efficient frontier** runs along the solid line from Portfolio B to Portfolio E, with Portfolio B known as the "minimum variance portfolio". Rational investors will prefer a portfolio along the efficient frontier, although which one depends on the individual's utility or tolerance for risk.

In this example, the highest returning portfolio, Portfolio E, is completely invested in the highest returning asset.

While it is possible to decrease the risk below the lowest risk asset through diversification, it is not possible to increase the return above the highest returning asset.

Keep in mind that several key assumptions have been made here, including:

- The risk of a portfolio is measured by the variability of its returns. Standard deviation is the statistical metric used for risk.
- The standard deviation of a portfolio can be reduced by the inclusion of assets whose correlations are less than 1.
- Investors are risk averse. This means that for two portfolios with the same return, an investor will prefer the portfolio with less risk.
- Investors are rational. This means that for two portfolios with the same level of risk, investors will prefer the portfolio with higher returns.
- Analysis is based on a single period model of investment.

Theory vs. Reality: Issues with the application of Modern Portfolio Theory

The *efficient frontier* is widely accepted as the standard theoretical model for investment decision making. Should we be surprised that the model is seldom used by the investment community? Not really. There is actually a fairly negative connotation to MPT and MVO in practice given empirical evidence that they're both ineffective in their current form.

The shortcomings of mean-variance optimization become obvious when applied. These shortcomings might be described best by Richard Michaud (1989) as an "error-maximization" device. He says mean-variance optimization "overweights those securities that have large estimated returns, negative correlations and small variances. These securities are, of course, the ones most likely to have large estimation errors." This phenomenon leads to instability of the model and can shake the confidence of portfolio managers in the process.

A classic case, perhaps, of academic theory versus real-world application?

We all know that a model is only as good as the estimation of its inputs. Previous applications of MVO use long-term averages for estimations, which is problematic as short-term realities of these factors can vary greatly from their long-term averages. **We instead advocate the use of rolling estimation horizons,** which have historically been associated with return persistence -- the momentum factor -- which is well established in financial asset prices (see Faber 2007, Hurst 2012, Antonacci 2011 and Asness 2014).

Using Momentum to Improve Mean-Variance Optimization Inputs

The simplest way to demonstrate the benefits of adding momentum factors is to stick with our two-asset example. We will use U.S. Large Cap Equities, as measured by the S&P500 Index ETF (SPY), and U.S. Treasuries, as measured by the U.S. Long Term Treasury Bond ETF (TLT), which illustrates the classic tradeoff of the negatively correlated stock and bond markets.

Step 1: Use Long-Term Averages to Create an Efficient Frontier

We begin with an example using long-term (15-year) metrics for these two asset classes:

Asset	15 Yr. Avg. Return	15 Yr. Avg. Standard Deviation	Correlation Matrix	
			Asset 1	*Asset 2*
U.S. Equities (SPY)	8%	13%	1	-0.55
U.S. Treasuries (TLT)	6%	12%	-0.55	1

We can then calculate portfolio combinations of SPY and TLT. For this example, we use 10% increases of these weightings:

Weight(SPY)	Weight(TLT)	Portfolio Risk	Portfolio Return	Return/Risk Ratio
100.00%	0.00%	13.00%	8.00%	0.6154
90.00%	10.00%	11.33%	7.80%	0.6883
80.00%	20.00%	9.81%	7.60%	0.7744
70.00%	30.00%	8.52%	7.40%	0.8681
60.00%	40.00%	7.58%	7.20%	0.9497
50.00%	50.00%	7.12%	7.00%	0.9826
40.00%	60.00%	7.24%	6.80%	0.9387
30.00%	70.00%	7.92%	6.60%	0.8337
20.00%	80.00%	9.02%	6.40%	0.7097
10.00%	90.00%	10.41%	6.20%	0.5954
0.00%	100.00%	12.00%	6.00%	0.5000

The risk and return of each portfolio are then plotted on a graph to create an efficient frontier:

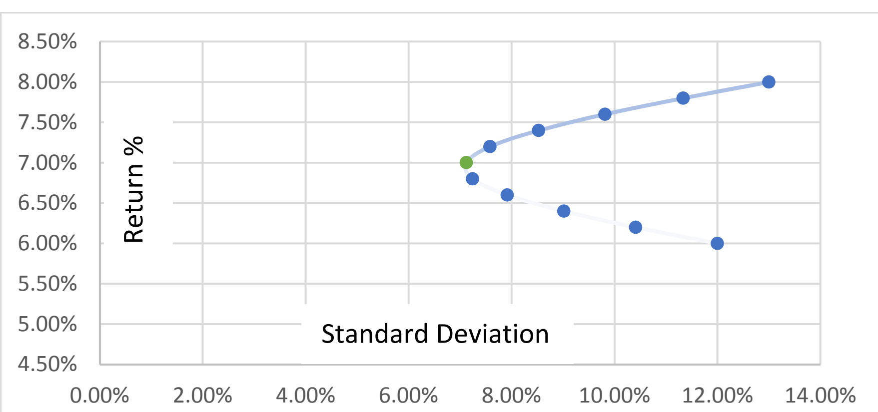

Notice that the risk -- standard deviation -- of the portfolio can be reduced below that of the standard deviation of either component. This outcome is owed to the negative correlation between the asset classes which produces the diversification benefits.

Traditionally, asset allocation portfolios were created along the efficient frontier for investors of various risk tolerances. For our part, we advocate the use of a portfolio with the *highest ratio of return-to-risk*, known as the "Maximum Sharpe Portfolio" or the "Optimal Portfolio" (we discuss later how to alter the Optimal Portfolio to meet the needs of investors with different risk tolerances).

To illustrate how adding a momentum factor to mean-variance optimization improves portfolio metrics, we must first look at an example of a traditional asset allocation model. For this example we buy and hold the Optimal Portfolio using the long term averages above, which is the portfolio with 50% SPY and 50% TLT.

Step 2: Buy-and-Hold the Optimal Portfolio

Traditional asset allocation portfolios buy-and-hold fixed percentages of various asset classes over time. Holding the Optimal Portfolio (50% TLT and 50% SPY) over the past 15 years (1/1/2004 through 9/30/2019) produces the following results:

	50 SPY /50TLT	**SPY Only**
Return (%)	8.23	8.91
Std Dev (%)	7.89	13.18
Max Drawdown (%)	-20.81	-50.89
Sharpe Ratio	0.88	0.62

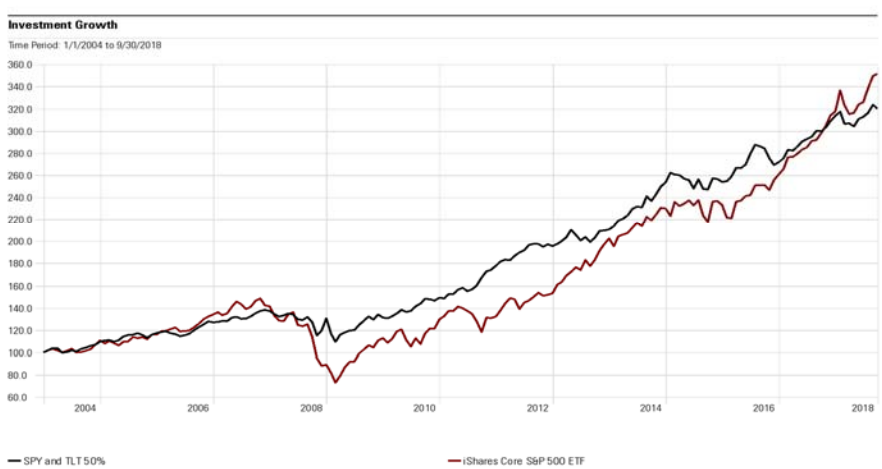

It's important to note that the return of the combined portfolio of assets cannot be improved above that of the highest returning asset.

<u>The diversification benefits come from a reduction in risk.</u> While the 'SPY Only' portfolio has a higher return than the combined 50% SPY and 50% TLT portfolio, the combined portfolio has almost the same return with nearly half the risk and less than half the maximum drawdown.

Step 3: Adding a Momentum Factor

To show how adding a momentum factor to the estimation of MVO inputs can be an improvement, we simply swap out the long-term averages of risk, return, and correlations for SPY and TLT and use a three-month period to calculate these metrics. An efficient frontier is then calculated on a monthly basis, and the portfolio is rebalanced to the Optimal Portfolio (again, the portfolio with the highest return-to-risk ratio) each month.

For example, if we were to create an efficient frontier and choose an Optimal Portfolio in September of 2008, the previous three months (June, July, and August of 2008) would be used to calculate the inputs:

May 31, 2008 through August 31, 2008				
			Correlation Matrix	
Asset	**3 Mo. Return**	**3 Mo. Standard Deviation (Annualized)**	*Asset 1*	*Asset 2*
SPY	-7.89%	25.60%	1.00%	-0.52%
TLT	5.00%	13.18%	-0.52%	1.00%

Weight (SPY)	Weight (TLT)	Portfolio Risk	Portfolio Return	Return/Risk Ratio
100.00%	0.00%	25.60%	-7.89%	-0.3082
90.00%	10.00%	22.87%	-6.60%	-0.2886
80.00%	20.00%	20.24%	-5.31%	-0.2624
70.00%	30.00%	17.75%	-4.02%	-0.2267
60.00%	40.00%	15.45%	-2.73%	-0.1769
50.00%	50.00%	13.46%	-1.45%	-0.1073
40.00%	60.00%	11.94%	-0.16%	-0.0131
30.00%	70.00%	11.06%	1.13%	0.1025
20.00%	80.00%	10.99%	2.42%	0.2204
10.00%	90.00%	11.75%	3.71%	0.3160
0.00%	100.00%	13.18%	5.00%	0.3793

In this example, the portfolio with highest return-to-risk ratio in September 2008 would be 100% in the U.S. Long Term Treasury Bond (TLT).

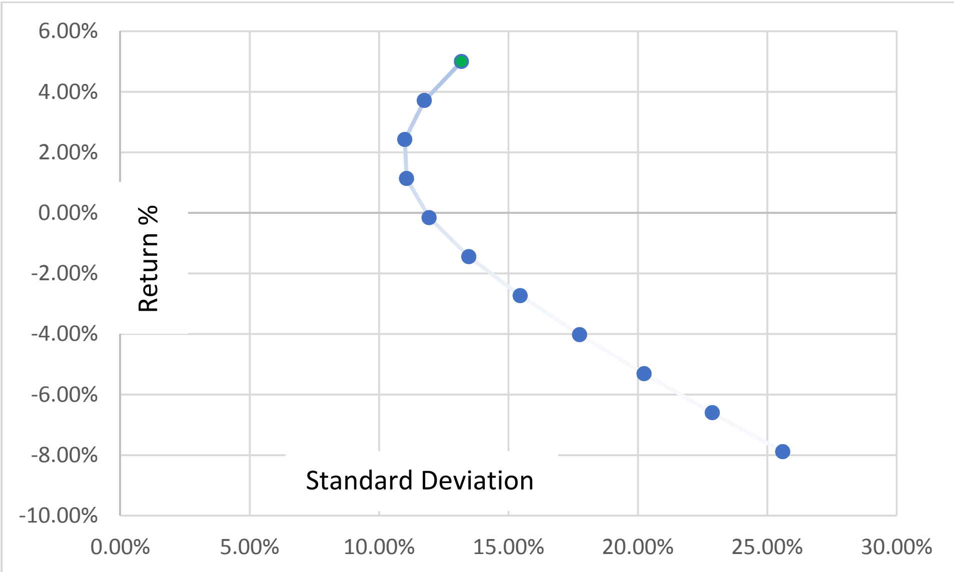

If this process were to be repeated, with the previous trailing three months used for estimation of inputs and the portfolio rebalanced monthly to the portfolio with the highest return-to-risk ratio, the results would have been as follows:

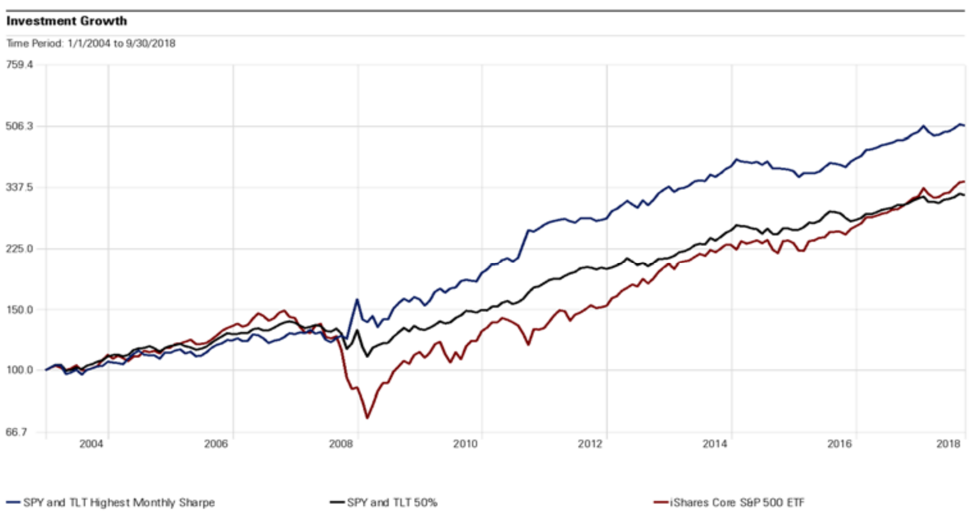

	Monthly 'Optimal Portfolio'	50 SPY /50 TLT	SPY Only
Return (%)	11.67	8.23	8.91
Std Dev (%)	10.62	7.89	13.18
Max Drawdown (%)	-16.84	-20.81	-50.89
Sharpe Ratio	0.97	0.88	0.62

Remarkably, the return of the portfolio rebalanced monthly to the Optimal Portfolio was higher than that of the S&P500. It can be proved, in other words, that ***returns of a more dynamic portfolio can be improved above that of its highest returning constituent.*** This contrasts with a more traditional static portfolio whose diversification benefits are limited to a reduction in risk. Also notable are improvements in the Max Drawdown and Sharpe Ratio of the portfolio rebalanced monthly.

The concept of adding a momentum factor to mean-variance optimization is the key insight into the TrueDiversification process that drives the Fund Architects Blueprint Portfolios. We believe the benefits are derived from an improvement in the estimation of inputs.

TrueDiversification: Adding Constraints to Traditional MVO

To create the complete TrueDiversification process, several constraints must be made to the traditional mean-variance optimization framework described above. Keep in mind that the intention of the TrueDiversification process is to maximize the benefits of diversification by continually targeting a portfolio with the optimal risk, return, and correlation metrics given current market conditions.

Constraint One: A Carefully Defined Investable Universe

We take a multi-asset approach

By defining an investable universe that includes multiple asset classes, we can benefit from (1) those classes' lack of correlation, and (2) their participation in global financial markets. Neither benefit is available when the investable universe is limited to a single index like the S&P 500.

We take advantage of the global financial markets

While the market capitalization of the S&P 500 ($25 trillion) represents about half of the total global investable stock market, as measured by the MSCI ACWI All Cap ($50 trillion), equity is only a small part of a marketable global investable universe that includes enormous fixed income and commodities markets. Only a strategy that tracks multiple asset classes and is unbound by geography can tap the potential of the vast global financial market.

We insist on the availability of non-correlated assets

Diversification benefits come from varying correlations between assets. It is our position that effective diversification comes not from the *number* of positions that are held, but from the *current correlation* of the underlying positions. A more stable return distribution demands a broad list of global asset classes.

Think about how many strategies limit their investable universe to U.S. Large Cap Stocks. The issue is that during a market panic, stocks in the S&P500 are likely to fall together. For example, during the Financial Crisis (defined from January 2008 through the bottom of March 9, 2009), even the top performing quintile of stocks in the S&P500 experienced more than a 30% drawdown over that period. Even a manager skilled enough to pick the highest performing stocks over that period would still have seen their investors lose a third of assets.

We define the TrueDiversification investable universe

- U.S. Large Cap Stocks
- U.S. Mid Cap Stocks
- U.S. Small Cap Stocks
- U.S. Real Estate
- U.S. Long Term Treasury Bonds
- Emerging Markets Equity
- European Equity
- Cash
- Gold

Constraint Two: A Measured Momentum Interval

Defining the time period used in the estimation of inputs is of utmost importance. The TrueDiversification process uses a combination of shorter-term (three month) and medium term (six month) intervals. If the interval is too short, one week to one month for example, then the process will react to normal market 'noise'. Our objective is for the process to avoid reacting to short-term market fluctuations that don't materialize into a longer-term trend. If the interval period is too long, more than one year for example, the process doesn't identify trends quickly enough.

It's encouraging that in testing, the intervals between three to six months had little significance. That is to say risk and return was not greatly affected by choosing a specific day or interval in this window. This insight is significant for the robustness of model.

Constraint Three: No Short Selling

Short selling was excluded from the process as it exponentially increases the number of input estimations, and creates operational difficulties since there are not inverse investment vehicles readily available for the investable universe. More importantly, if the asset classes included in the process have low to negative correlation with each other, then short-selling is not necessary.

Constraint Four: Asset Weightings Capped at 40%.

Given our intention to maximize diversification benefits, it is necessary to always hold at least two asset classes. It is therefore necessary to cap the individual asset class weightings below 100%.

In testing, it was discovered that holding three to four asset classes typically produced the most favorable results. As market conditions change and momentum trends reverse, the additional sources of return provided a smoother return distribution. It also makes intuitive sense that an asset allocation portfolio should not be majorly comprised of one asset class.

The individual asset class weightings are thus capped at 40%. This constraint ensures the portfolios will always have three to four asset classes, and that no asset class has a majority weighting. Importantly, it also allows for concentrated positions that can capitalize on current trends.

Constraint Five: Monthly Rebalancing to the Optimal Portfolio

Traditional asset allocation portfolios rebalance at certain intervals, typically quarterly or annually. This practice is to reduce exposure to positions with outsized gains and use the proceeds to increase exposure to underperforming assets, the concept being that assets trends change over time.

The TrueDiversification process rebalances monthly to the Optimal Portfolio for the same reasons.

Testing was performed using multiple rebalance periods. If the portfolios were rebalanced too frequently, for example daily or weekly, risk and return metrics were not improved but turnover (which negatively impacts transaction costs and taxes) increased. If the rebalance frequency was too long, the process did not react quickly enough to changing market conditions.

The main concern with a monthly rebalance is that the process is not able to protect against intra-month market crashes. In reality, the TrueDiversification process is intended to respond to clearly defined trends. In the rare occasion of something like a "flash crash", it is likely that any built-in selling mechanism would cause the portfolios to sell at exactly the wrong time.

Using TrueDiversification to Identify the Optimal Portfolio

The completed TrueDiversification process, which uses mean-variance optimization from Modern Portfolio Theory as the decision-making framework and applies the above listed constraints, can now be used in practice. The first step of the process is to identify an efficient frontier by creating portfolios of every possible asset class combinations.

Here is a sample monthly efficient frontier from the results of the TrueDiversification process:

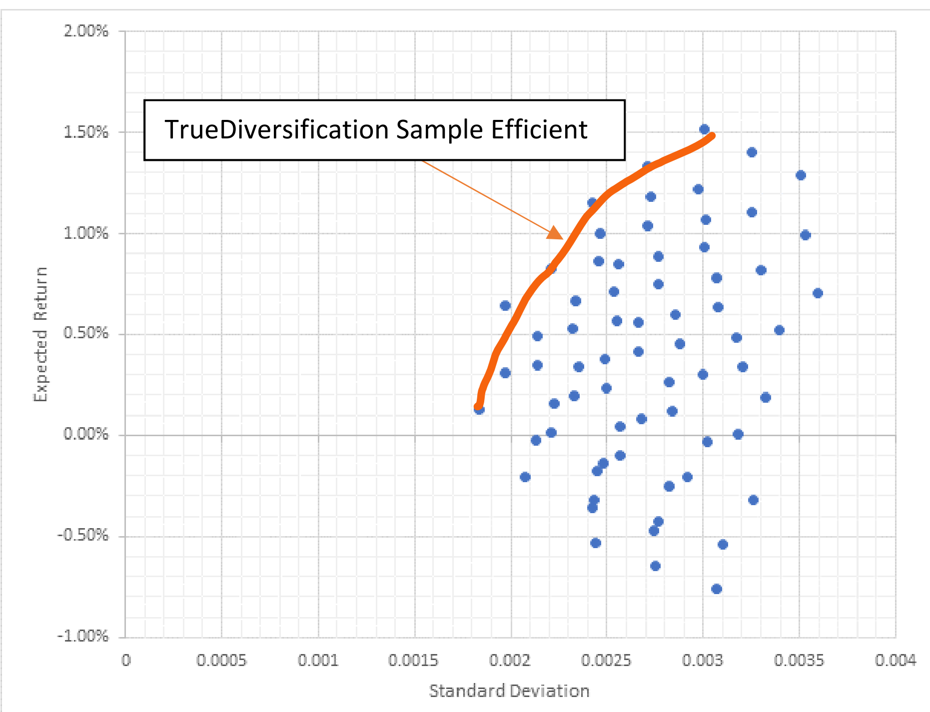

The shape of this efficient frontier changes each month. Occasionally, it will be steep, indicating that for a marginal increase in risk there can be meaningful performance gains, or the curve may be flat, indicating that the portfolio would require large additional risk to increase return potential. The Optimal Portfolio will therefore be at a different point along the efficient frontier each month, depending on the shape of the curve.

This efficient frontier is not static. Since we are using momentum indicators to assist with the estimation of risk, return and correlations, the shape of the curve will vary with changing market conditions. The purpose of this exercise is to identify the Optimal Portfolio each month, which is the portfolio with the highest ratio of return to risk.

Every investor should want the Optimal Portfolio since it is the most efficient portfolio along the efficient frontier.

While all portfolios along the efficient frontier may have the highest return for a given level of risk, each portfolio along the frontier will have a different return to risk ratio. This observation is evidenced by the changes in slope of the efficient frontier along the curve.

The traditional approach to allocating investors across the risk spectrum is to choose a portfolio along the efficient frontier that matches the investors' tolerance for risk, as measured by standard deviation. *We can say with great confidence that this conventional approach does not maximize the potential of the frontier*, as it gives each investor a different risk/return trade-off.

How, then, can we provide the Optimal Portfolio to investors with different risk tolerances? The solution is borrowed from the concept of the Capital Allocation Line.

Using the Optimal Portfolio to Construct Blueprint Portfolios

A solution to providing all investors with the same risk/return trade-off was proposed by Markowitz through the Capital Allocation Line (CAL), also known as the Capital Market Line (CML). The CAL is simply a series of portfolios created by the combination of a risk-free asset and the Optimal Portfolio. For investors with a risk tolerance less than the Optimal Portfolio, an allocation to the risk-free asset is included at increasing percentages. For investors with greater tolerance for risk than the Optimal Portfolio, a higher percentage of the Optimal Portfolio is purchased by borrowing against the risk-free asset, which creates leverage.

The CAL is always formed at the tangent point between the intersection of (1) the risk-free rate at '0' risk on the x axis, and (2) the Optimal Portfolio from the efficient frontier. This tangent point gives the CAL the highest slope possible. Since the slope of the CAL is measured as return over risk, the CAL will therefore have the highest possible return to risk ratio at the tangent point. Furthermore, since the CAL is a straight line, meaning the rate of change of the slope of the CAL is constant, every investor along the CAL will have the same return to risk ratio.

The area between the CAL and the efficient frontier (shaded below) represents the amount of borrowing (leverage to increase risk) or lending (a bond portfolio to decrease risk). *It also represents the amount of improvement to either return or risk as compared to the respective portfolio along the efficient frontier*.

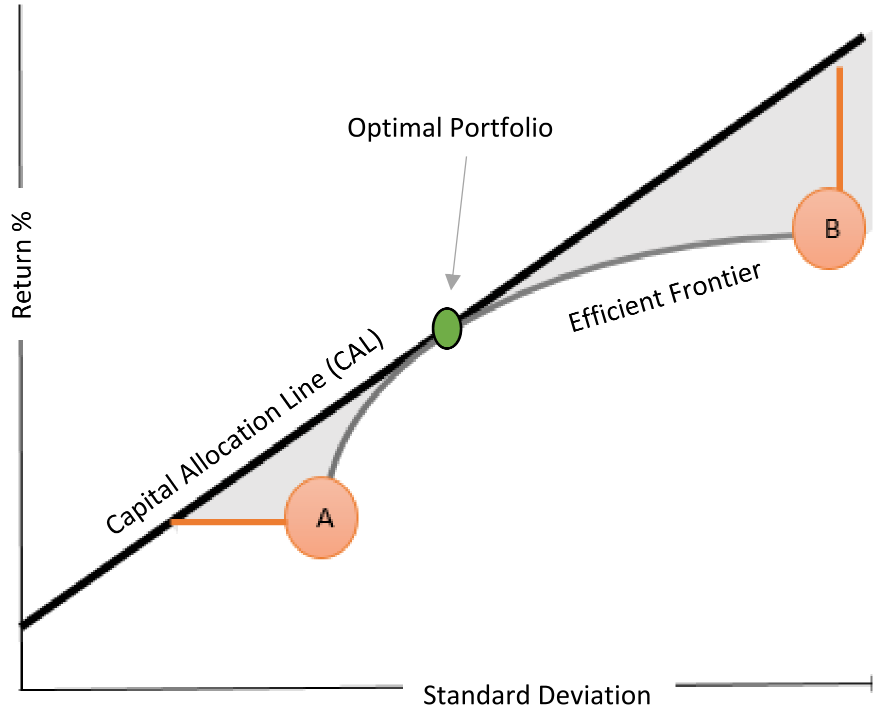

A defensive investor would traditionally be allocated along the efficient frontier at the minimum variance portfolio (Portfolio A). *This same investor can realize a material reduction in risk, with the same expected return, by instead allocating along the Capital Allocation Line.*

An Aggressive Investor (Portfolio B) can improve return for the same level of risk.

Improvement beyond the efficient frontier, as measured by an increase in return for the same level of risk or a reduction of risk for the same level of return, will depend on where the portfolio lies along the efficient frontier, as well as the shape of the efficient frontier.

In practice, a bond portfolio is purchased instead of the risk-free rate, and leverage can be achieved through the purchase of 2x leveraged ETFs.

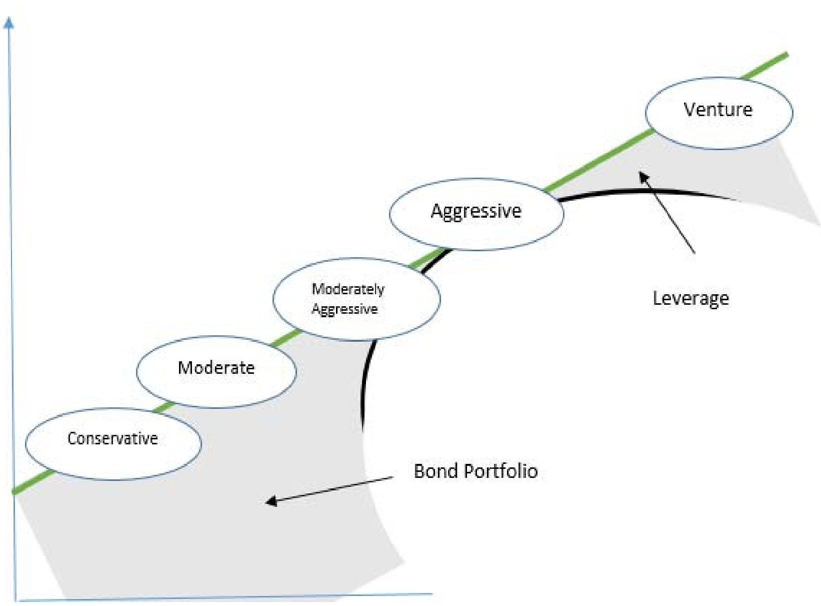

Fund Architects Blueprint Portfolio	Optimal Portfolio	Bond Portfolio
Venture	200%	0%
Aggressive	100%	0%
Moderately Aggressive	75%	25%
Moderate	50%	50%
Conservative	25%	75%

Conclusion

The concept of adding a momentum factor to mean-variance optimization is the key insight into the TrueDiversification process. By acknowledging a persistency factor in asset prices, the estimation of inputs is improved to the extent a more dynamic portfolio can produce returns higher than that of its highest returning constituent.

All of this contrasts with a more traditional static portfolio whose diversification benefits are limited to a reduction in risk.

Further, by continually targeting a portfolio with the optimal risk, return, and correlation metrics based on current market conditions, TrueDiversification's "Optimal Portfolio" comes closer to actually capturing the benefits of diversification suggested by Modern Portfolio Theory more than 60 years ago.

The TrueDiversification process can be accessed through the Blueprint Portfolios. In all, we find compelling evidence this unique approach to asset allocation will provide investors a dramatically improved investment experience.

*****

References

Antonacci, G, 2011, Optimal Momentum: A Global Cross Asset Approach, SSRN 1833722

Asness, C.S., Frazzini, A., Israel, R. and Moskowitz, T.J., 2014, Fact, Fiction and Momentum Investing. Journal of Portfolio Management, Fall 2014 (40th Anniversary Issue); Fama-Miller Working Paper. SSRN 2435323

Faber, M. T., 2007, A Quantitative Approach to Tactical Asset Allocation, Journal of Wealth Management, Spring 2007. Update (2013) as SSRN 962461.

Hurst, B., Ooi, Y.H., and Pedersen, L.H., 2012, A Century of Evidence on Trend-Following Investing, working paper, AQR Capital Management. 21

Markowitz, H. M. , 1952, Portfolio Selection, Journal of Finance, 7(1), 77–91.

Kondor, Papp and Caccioli, 2017, Analytic solution to variance optimization with no short selling. SSRN 2899431

W. Keller, A. Butler, and I. Kipnis, 2015, Momentum and Markowitz: a Golden Combination. SSRN 2606884

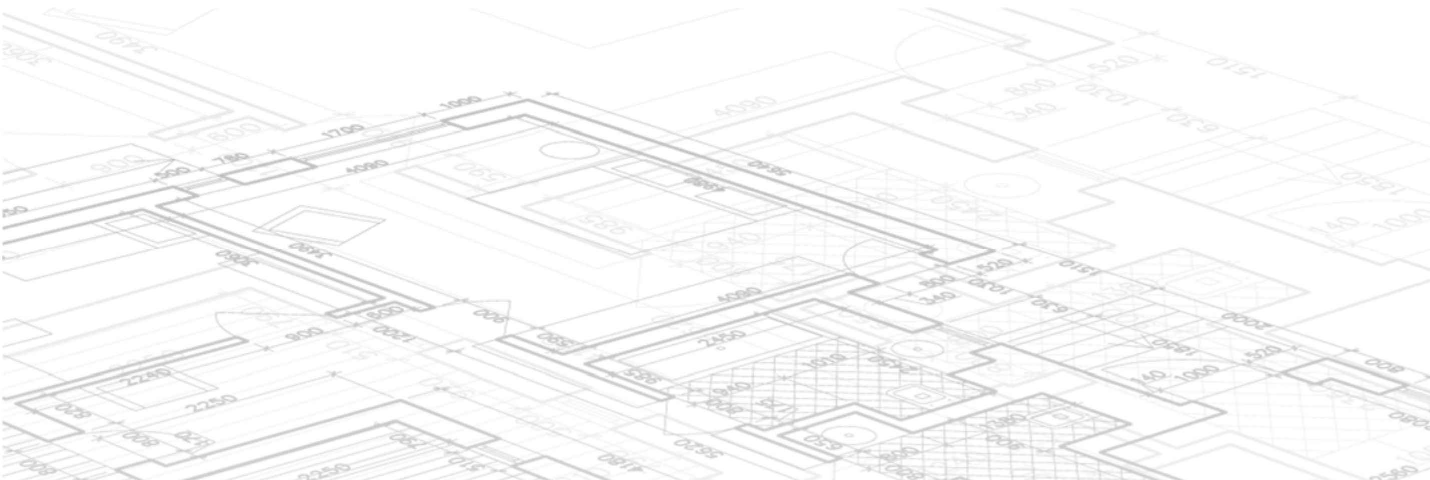

FAWP-0119