Preferred Portfolios:
An Improved Blueprint to Construct Multi Strategy Portfolios

Lars N. Kestner
March 2020

Summary

This paper introduces a blueprint for combining individual strategies into a robust portfolio using five principles.

» Continued research is paramount to replenish individual strategies as performance wanes over time

» Identifying and classifying strategies correctly based on their dynamic return characteristics leads to temporal validity in portfolio performance

» Employing strategies on a diverse set of markets increases the opportunity set available

» Understanding the benefits and limitations of diversification guides manager decisions for research and resource allocation

» Sufficient time is required to give portfolios the opportunity to prosper from the skill in the investment program

Together these five elements build the basis of Statistically Applied Trading (Saτ).

STATISTICALLY APPLIED TRADING
www.satquant.com
New York

Introduction

The original title for this paper was *Perfect Portfolios*. But there is no such thing as perfection in financial markets as noise dominates signal, and strategy performance is anything but constant over time. As managers, the best we can do is to anticipate what could change and construct robust portfolios that weather these changes. Portfolios may be optimal at one point in time, but not next quarter, let alone next year. An investment manager's goal is to create temporal validity in the portfolio construction process, not just portfolios that perform at one point in time or market regime. This paper introduces a blueprint for combining diverse strategies into portfolios that balance exposures by identifying similarities and differences using creative new means. We will think about what a strategy is designed to accomplish, categorizing each based on their structural similarities, rather than indiscriminately optimizing on historical correlations. We will show where to focus effort in future research once we define our expectations and compare these to our research results. We construct our portfolios with two major building blocks: a diverse set of strategies and a wide range of markets to apply these strategies on. The strategies will be clustered into four categories based on their characteristics. Markets will be split into nine classes. We will also show why it is of great importance to give these well-constructed multi strategy portfolios time to perform as investors may not receive the returns they deserve without proper patience. We will not detail specific quantitative strategies here. Rather, we will concentrate on introducing a framework named Statistically Applied Trading (Saτ) that builds portfolios of strategies into a complete investment program.

The contents of this paper are applicable to a wide variety of investment professionals. Multi strategy managers can use the blueprint to rethink how various strategies fit into a portfolio. Investment advisors can employ the framework to better understand the relationships of their clients' exposures across a wide range of asset classes. For investors in risk premia portfolios, the framework will contribute on how to combine return streams to create a balanced all-weather portfolio as more alternative investment options become available. Simply looking at historical return and covariance matrices without understanding the basis behind these outputs will cause investor disappointment precisely when proper diversification is needed.

Quantitative trading strategies

Why quantitative? Over the last two decades, behavioral finance has gained credibility as an explanation for certain market anomalies. Anchoring, confirmation bias, loss aversion, and overconfidence bias are among the many cognitive biases that have been discovered as reasons why individuals make certain investment decisions. This prior research has suggested that behaviors resulting from cognitive biases are large enough to push edge from one side of a trade to another. But if most seasoned investment professionals agree with that conclusion, why do inefficiencies persist? Are we unable to conquer our own human shortcomings with the assistance of process driven investment decisions? Perhaps opportunities stay because of algorithm aversion. Economists from The Wharton School [documented](#) that people were less trusting of successful algorithms after seeing these models occasionally fail. It may be that the individual mind is not able to put aside emotions when making investment decisions.

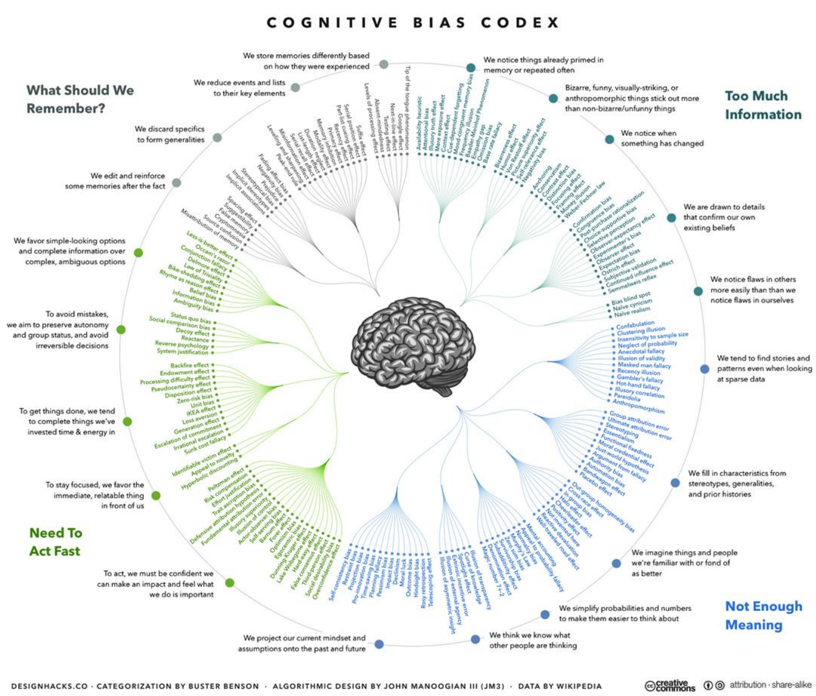

The Cognitive Bias Codex splits cognitive biases into similar groupings.

Source: https://upload.wikimedia.org/wikipedia/commons/6/65/Cognitive_bias_codex_en.svg

Regardless of why inefficiencies exist, how to exploit them is best completed quantitatively[1]. The benefit of quantitative trading is that we can experiment with predictive models before risking money. Even once a program is trading, we can form new opinions and change our actions as new data and analysis suggest we do so. Unlike a game of

[1] We use the terms quantitative, algorithmic, statistical, and systematic interchangeably to characterize a trading strategy that is rules based.

blackjack, which is stationary with constant rules, markets and strategy performance do change over time. Hypothesis formulation allows us to perform real time experiments to determine optimal allocation decisions dynamically. Only through daily research and experimentation can we constantly improve our process, making it robust over time.

Saτ looks to create portfolios that have temporal validity[2]. Temporal validity is a concept borrowed from the social sciences which considers the outcome of experiments if they were repeated over time. Results of a study would not have temporal validity if the outcome could not be replicated with small changes in the study's environment. This concept fits nicely with what we are trying to accomplish thematically. Instead of solving for the optimal now or in the past, Saτ creates portfolios designed to withstand unpredicted changes and even shocks to the system. We believe the best way to achieve temporal validity is to avoid trying to predict the trading environment. Our job is to sift through the noise with a systematic process, find small edges, and craft them together into a powerful combination of return streams. These portfolios should perform under almost all market conditions, making the need to predict the market less relevant.

Projecting Strategy Performance over Time

Performance comes in all shapes and sizes. One of our favorite papers on market efficiency hypothesizes that inefficiencies are like a rubber band: the further you stretch them, the faster they snap back.

The signals with the highest Sharpe ratios[3] have fast performance decay and should not persist with the same strength over time as they are discovered by more market participants. Many of these strategies profit from true mispricings, opportunities in nascent markets with high barriers to entry, or utilize higher frequency strategies with capacity constraints. Eventually, other market participants solve the pricing mistake, find ways to enter a market, or game theory their way to remove some of the excess edge. On the other side, lower performing signals can remain for years, but the rewards to harvesting these are also much lower. A majority of these signals have been traded for years and have been discussed in academic literature. Examples of such strategies are equity factors, time series momentum, and currency carry.

[2] Temporal validity is sometimes referred to as temporal external validity.

[3] In this paper we use Sharpe ratio as a general measure of reward per unit of risk. While we are cognizant of its shortcomings for strategies with significant skewness or convexity, we leave that discussion for later writings.

To illustrate this point, consider a graph where we scatter the expected instantaneous Sharpe ratio versus the time since the signal was discovered. The highest performing signals begin at the top left of the graph but decay as they are discovered by other participants or exploited to capacity. Lower performing signals can remain sticky as the benefit is not large enough to attract new money or at least balanced by old money leaving to chase better opportunities. Note the curvature on the graph. High performance strategies decay faster than low performance as the rewards to trading are more profitable.

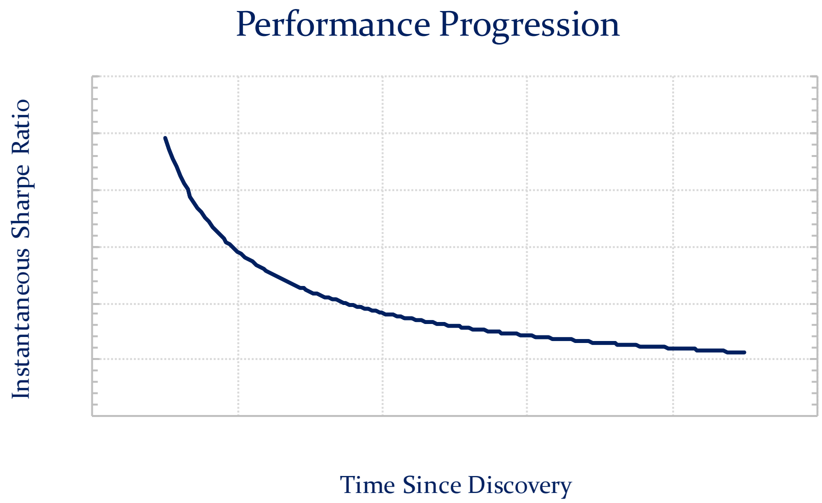

Source: Saт calculations

The instantaneous Sharpe ratio of a strategy is graphed versus the time elapsed since discovery. Strategies are in a continual progression from top left to bottom right.

On the top left are partner funded strategies usually constrained by capacity. The middle is often the domain of hedge funds and funded partnerships charging management and incentive fees. Strategies in the bottom right are the realm of asset gatherers with names such as smart beta or risk premia who typically charge management fees only.

When developing strategies, a manager needs to estimate which quadrant a strategy resides in. Be honest. If a strategy is truly in the top left, expect performance to decay quickly. If a strategy is in the bottom right, you might be able to retain your edge for years, but that smaller edge could leave you vulnerable to an unlucky performance draw that yields losses.

Saт believes that strategies are in a continual progression from top left to bottom right. Strategies are discovered and degrade. Research is key to replenishing a constantly declining signal strength. Managers who refuse to innovate and are using the same strategies as 10 years ago should acknowledge they are likely in the bottom right of the graph, and certainly not in the top left. Saτ believes it is new research and constant improvement that maintains the performance of the portfolio, creating temporal validity in investment results.

Building blocks

To create a multi strategy portfolio, managers combine individual strategies together, balancing characteristics such as expected returns and expected correlations. While measurements such as historical correlations are a reasonable start to determine similarities of strategies, Saτ believes there are smarter, more advanced means to classify ideas.

In the graph below, we use a technique called hierarchal agglomerative clustering (HAC) to view similarities and differences of 18 assets classes and strategies[4,5] based on their historical returns. While the specific math behind HAC is beyond the scope of this paper, in general, the technique begins with each data object belonging to its own cluster, successively merging them by similarity until there is a single group containing all the data objects. The HAC output is shown most often as a dendrogram where each combination of successive groups is represented by a horizontal bar.

Dendrogram

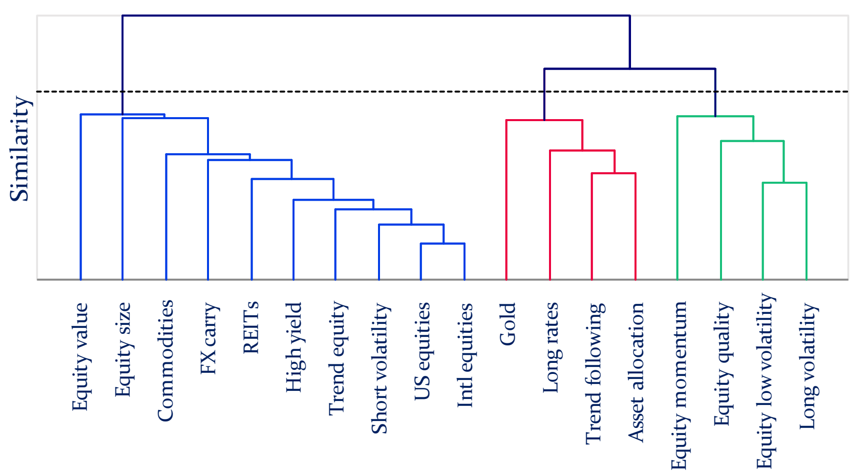

Source: Bloomberg, Saτ calculations

Dendrogram graph displaying the output of the clustering technique. Strategies in blue could be considered “risk-on” while strategies in red and green could be considered market neutral.

In general, the clustering technique does a satisfactory job grouping strategies by like characteristics. Most of the blue strategies on the left perform better with calm, rising equity markets. The red grouping is

[4] These return series were chosen because each is a disclosed investment process that can be fully replicated. A true multi strategy portfolio might not be as transparent to investors.

[5] Equity value=DJTMNSV Index, Equity size=DJTMNSS Index, Commodities=SPGSCITR Index, FX carry=FXCTG10 Index, REITs=IYR Equity, High yield=HYG Equity, Trend equity=NRROMOT Index, Short volatility=SPVXSPI Index, US equities=SPY Equity, Intl equities=EFA Equity, Gold=GLD Equity, Long rates=TLT Equity, Trend following=NEIXCTA Index, Asset allocation=EEJPUS5E Index, Equity momentum=DJTMNMO Index, Equity quality=DJTMNQU Index, Equity low volatility=DJTMNAB Index, Long volatility=SPVXSP Index

ETFs data was used where possible to reduce asynchronicity of returns.

primarily composed of macro strategies such as long rates, long gold, or rule-based trend following and asset allocation strategies. The green grouping is a combination of long/short equity factor strategies along with a long volatility strategy that would benefit from increases in perceived risk in the financial system. If we summarized results from this analysis, we would likely say that most blue strategies would be considered "risk-on", and most of the strategies in the red and green categories would be considered market neutral.

Market neutral is an interesting concept that we can achieve using various trading techniques. We could pair a long stock portfolio selected using fundamental characteristics with a short stock portfolio of opposite characteristics. Depending on the fundamental metric, the resulting return stream could have very little relationship to broad equity market performance. Separately, we might also use time series analysis to determine when to buy or sell the S&P 500. This method could also be considered market neutral if it longs and shorts an equal amount of the market over time. Over a long time horizon, its average beta is zero. However, the characteristics of our two market neutral strategies are very different. The long/short method is a real diversifier while the time series momentum method selects a perfectly correlated long or short return stream that only averages to zero correlation over long horizons. Not all market neutral strategies are constructed to be truly market neutral.

The previous classification process is a good beginning but is not sophisticated enough to recognize changes in relationships over time. We need an alternative classification system that is able identify these dynamic relationships. Saτ completes this requirement by recognizing four categories to classify strategies: **Boosters**, **Defenders**, **Diversifiers**, and **Selectors**. Boosters are strategies that have a positive relationship to risky assets. That is, they generally perform better when equities rise and suffer as equities fall. Defenders' returns are opposite – performing better as equities fall and worse as equities rise. Diversifiers have much less consistent correlation to risky assets as they spread longs and shorts symmetrically. At any point in time, Selectors can have a specific performance bias with respect to risky assets, but that direction changes over time as signals switch signs from long to short and back, or the composition of the portfolio takes on characteristics that are no longer market neutral compared to "risk-on" asset classes. This property leads Selectors to have more variation in correlation with benchmarks such as equity indices compared to other classes of strategies.

Four Strategy Classifications within Saτ

Saτ recognizes four categories to classify strategies: **Boosters**, **Defenders**, **Diversifiers**, and **Selectors**.

To pinpoint which category a strategy belongs to, we will compare two calculations: an average correlation to our benchmark ($\bar{\rho}$) and the variability of our correlation calculation to the benchmark ($\sigma_{correlation}$). Using our 17 sample markets identified earlier as an example of this classification process[6], we first calculate rolling 52-week return correlations for each of the strategies with returns on a broad US equity benchmark spanning between 2009 and 2019. Next, we estimate the variability of the correlation between each market and the benchmark by calculating the volatility of the correlation series over the entire date series.

When we graph the average correlation versus the volatility of correlation, we find breakpoints that are helpful to classify strategies. Those with larger positive $\bar{\rho}$ are "risk-on" Boosters, while strategies with more negative $\bar{\rho}$ are "risk-off" Defenders. Strategies with mildly positive or negative $\bar{\rho}$ but smaller $\sigma_{correlation}$ are Diversifiers. Strategies with mildly positive or negative $\bar{\rho}$ but larger $\sigma_{correlation}$ are Selectors.

[6] The 18th market, the S&P 500, is removed and becomes our benchmark.

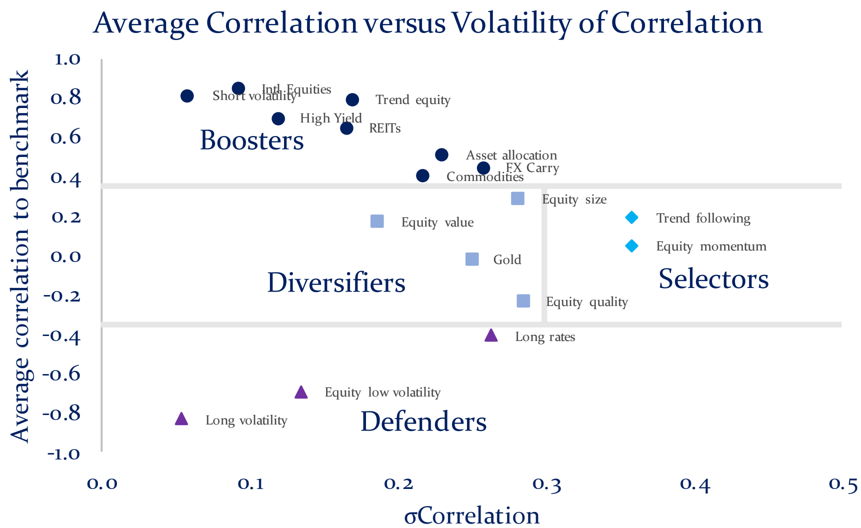

Chart of average correlation to the benchmark versus the volatility of the correlation to the benchmark. The combination of each helps us determine the classification of the strategy.

Source: Bloomberg, Saт calculations

Being able to categorize a strategy is helpful, but of greater importance is the stability in the behavior of each type. Below we graph the average correlation of the four strategy types to the S&P 500. The correlation of Boosters to the S&P 500 is positive throughout the sample studied. Defenders' correlation is negative throughout. Diversifiers meander near zero while the correlation of Selectors is much more variable in both directions. All these follow expectations based on how we classified the strategy type.

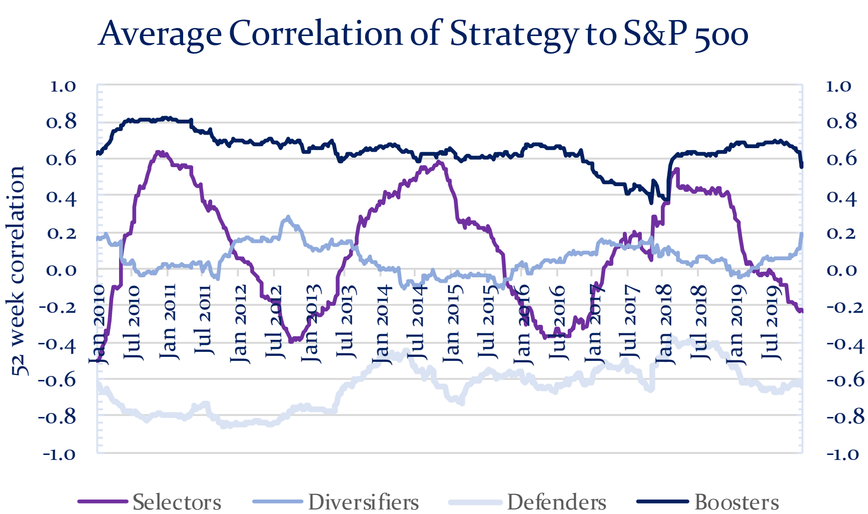

For all strategies in a specific classification, the average correlation to the benchmark is calculated on a rolling basis.

Source: Bloomberg, Saт calculations

Determining where a strategy should be classified is part experience and part science. Our experience comes from decades of designing systematic trading strategies coupled with the expertise of managing risk in derivative markets. We recognize how position signals will behave under different circumstances and market environments. The science part comes from utilizing sophisticated mathematical techniques to confirm what our market experience suggests.

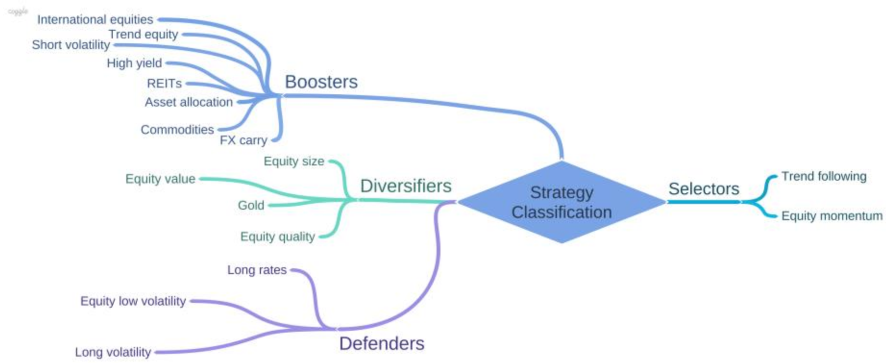

Source: Bloomberg, Saτ calculations

A mind map visualizes the output of the classification process.

The process of combining strategies based on their classification will depend on the ultimate goals of the portfolio. A market neutral portfolio needs a balanced amount of risk between Boosters and Defenders. If the consistency of market neutrality is a priority, portfolios should not be weighted heavily towards Selectors as these exposures can switch from risk-on to risk-off based on the underlying rules and composition of the individual strategies. Individual managers should have a specific target risk profile in mind for net exposure and variability of exposure in the portfolio construction process.

Classifying potential trading strategies is central to designing a portfolio, but we also need to diversify our trading over many markets and asset types. As we will see later, genuine diversification benefit requires low strategy correlation. The best means to achieve low correlation is to expand the range of markets well beyond those in the first page of the financial press. Saτ considers nine types of markets to apply strategies on.

Nine market types within Saτ

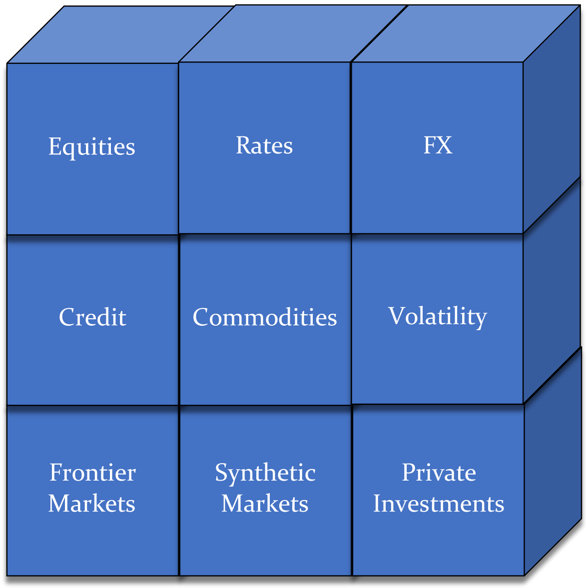

There are nine market types that employ strategies within Saτ.

While each class of markets in our cube is not completely orthogonal to the others, we believe that this method of organizing trading instruments is a practical and effective way to view where risk is being taken.

After classifying ideas into one of four categories of trading strategies and deciding which of the nine classes of markets to apply each strategy to, we have taken a step towards building Saτ.

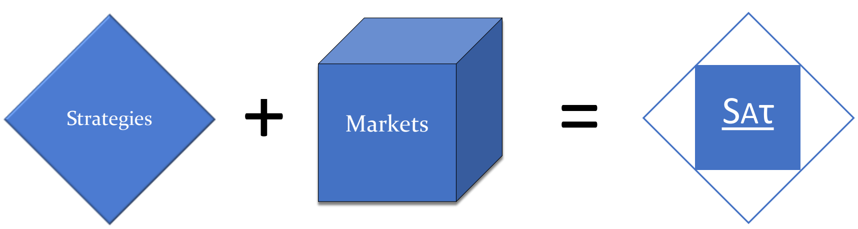

When we combine Strategies plus Markets, we have taken a step towards building Saτ.

Benefits and Limits of Diversification

When creating an investment portfolio of any kind, the initial plan should include estimating the portfolio's expected return and volatility. Once the manager has designed individual strategies to maximize their individual expected returns and crafted methods to control for volatility, the focus turns to combining strategies together to define the portfolio. Diversification across non-correlated return streams is the foremost value creator in portfolio management. Expected returns can be maintained while decreasing portfolio risk as we add strategies that zig when others zag. But the limits of these benefits are bounded, subject to specific parameters that enter the calculation of volatility. To the multi strategy manager, an obvious question is "how many strategies do we need to be diversified"? While more is better than less, the incremental value to adding strategies might be smaller than one thinks. In a recent conversation, a colleague phrased this thought as "the miscalculation of diversification"[7].

There are mathematical limits of diversification that managers and investors need to be aware of. These benefits might be reached with fewer strategies than expected. Consider an example where we create portfolios with a varied number of strategies. For simplicity, each strategy has an identical expected return of 4% per annum with annualized volatility of 10%, yielding a Sharpe ratio of 0.4. Below is a graph of portfolio Sharpe ratios as we change two variables: a. the number of strategies in the portfolio, and, b. the correlation among strategies in the portfolio[8].

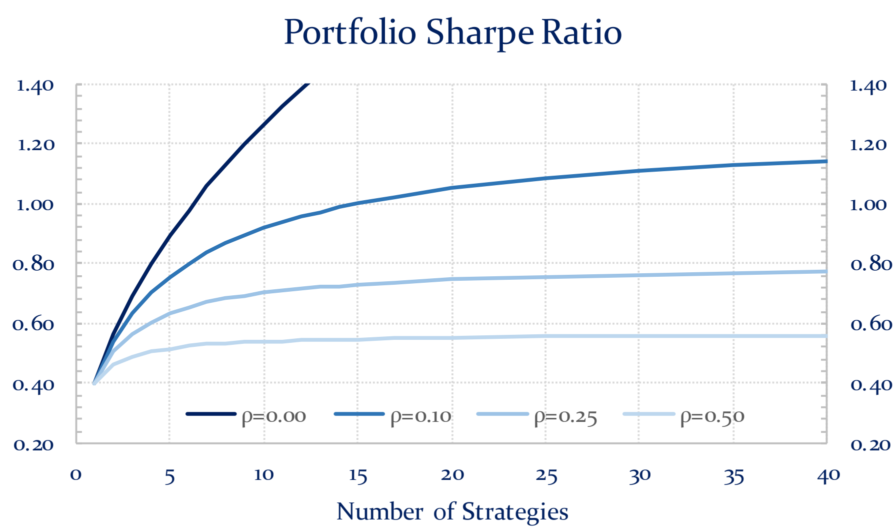

Source: Saτ calculations

The Sharpe ratios of diversified portfolios are calculated as the number of strategies is increased. Unless strategies are completely uncorrelated with all others, most of the benefits of diversification occur within approximately 20 strategies.

The incremental benefits of adding strategies is very dependent on the correlation between them. For any correlation less than 1, the portfolio's reward to risk improves with each strategy added, albeit in diminishing amounts. The rate at which improvement diminishes also depends on the correlation among strategies. Suppose we identified 200 strategies and created an equal weight portfolio of these strategies. The Sharpe ratio of the equally weighted 200-strategy portfolios is 5.66, 1.24, 0.79, and 0.56 if each strategy's correlation to all others were ρ=0.00, ρ=0.10, ρ=0.25, and ρ=0.50 respectfully. For reference, we next

[7] Maybe it was a colleague. Maybe it was my 10 year old daughter. Regardless, it was a clever comment.

[8] The volatility of a portfolio whose individual asset volatility and correlation among assets is equivalent can be calculated as:

$$\sigma_{portfolio} = \sqrt{\frac{\sigma^2}{n} + \frac{(n-1)}{n}\rho\sigma^2}$$

solve for the number of strategies needed where we receive an overwhelming 80% of the diversification benefits[9] of the 200-strategy portfolio. We see that this number can be quite low. If correlation among strategies is ρ=0.50, 80% of the diversification benefits occur at only 8 strategies. At correlation of ρ=0.25, that 80% benefit is reached at 12 strategies. Even with very minimal correlation of ρ=0.10 among strategies, 23 strategies get us to 80% benefit. This exercise shows an important result: even if a portfolio's strategies are largely uncorrelated, most of the diversification benefits are achieved once we have approximately 20 strategies[10].

Turning the analysis around to focus on expectations, consider a multi strategy manager who is targeting expected excess returns equal to that of expected volatility (i.e. a portfolio Sharpe ratio of 1.0). The manager can achieve this result in different ways depending on the number of strategies traded, the strength of performance of each strategy, and correlation among strategies created. As an example, two solutions that lead to a portfolio Sharpe of 1.0 are creating 5 strategies each with Sharpe of 0.63 with a ρ=0.25 correlation among them, or 40 strategies each with a Sharpe ratio of 0.35 with ρ=0.10 correlation among them. Managers need to ask themselves if this combination is realistic based on researched results. And if research finds you short of that goal, what path is more likely to achieve the desired result? Should you focus on taking existing strategies and improving them, or, should you spend incremental time focusing on new, largely uncorrelated signals once you have extracted as much value as you can out of existing signals? It is important to understand what you need as a manager to meet your goals, and to be aware of what may be out of your reach. The number of truly orthogonal bets that scale in size is likely smaller than you think. The tradeoff above can help managers evaluate the relative benefit between focusing on improving current strategies versus searching for new strategies with minimal correlation to existing strategies.

9 We calculate the percentage benefit when increasing strategies as:

$$\frac{SR_n - SR_1}{SR_{200} - SR_1}$$

where SR_x is the Sharpe ratio of a portfolio with *x* number of strategies

10 Coincidentally, this same point was made on a popular financial blog in the weeks leading to this paper being published.

Ten combinations to achieve a Sharpe ratio of 1.00

Number of Strategies	Individual Sharpe Ratio	Correlation among Strategies	Portfolio Sharpe Ratio
2	0.74	0.10	1.00
2	0.79	0.25	1.00
5	0.53	0.10	1.00
5	0.63	0.25	1.00
10	0.44	0.10	1.00
10	0.57	0.25	1.00
20	0.38	0.10	1.00
20	0.54	0.25	1.00
40	0.35	0.10	1.00
40	0.52	0.25	1.00

Source: Saт calculations

We present ten combinations that vary the number of strategies, the performance of each strategy, and the correlation among strategies to achieve a Sharpe ratio of 1.00.

The Value of Time in the Portfolio Process

Perhaps the most underestimated ingredient for portfolio success is time. Even as we craft truly robust investment portfolios, we remain subject to luck and random variability in performance. The only way to transform skill into returns is time. Having enough time for the investment program to achieve its goals and compound value is paramount.

Let us consider an investment program that has a true expected return of 7.50% per annum with annualized volatility of 7.50% per annum. Most investors should consider this as good performance, particularly if the returns are uncorrelated to asset class benchmarks. When 100 random paths of monthly returns over 10 years are simulated, the average path is quite positive, but variation does occur around expectations. The ending wealth in this simulation assumes that investors remained with the program over all 10 years.

Suppose we change our investors' behavior. Using the exact same set of returns as above, investors now redeem if either of the two following conditions are met: a. 12 month trailing returns are less than -7.50% or b. peak to trough drawdowns are more than 10.00%. The redemption decision now has a large impact on ending wealth. We see the difference in the simulated wealth curves and cumulative distributions of ending wealth on the next page. The redemption decision ended many of the positive wealth outcomes earlier than ideal, creating a drag on average performance. By following these redemption rules, investors missed out on more than 25% of the 10 year investment returns had they simply left the funds invested. Even more glaring, more than 50% of the managers faced a redemption over the 10 year simulation, despite no changes to the underlying expected return and volatility of the portfolio.

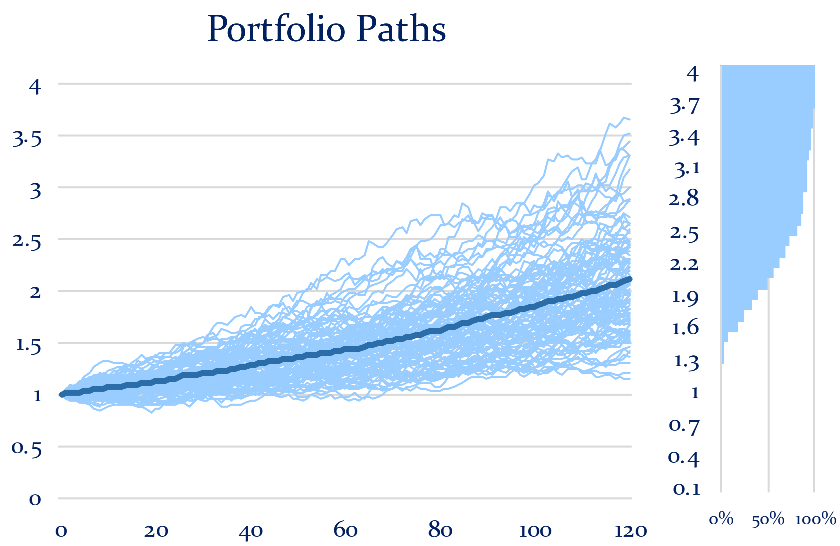

Source: SaT calculations

100 paths are simulated over 10 years with each light blue line representing one path. The dark blue line is the average of the 100 simulations. The adjacent chart is the cumulative distribution of ending wealth based on the 100 simulations.

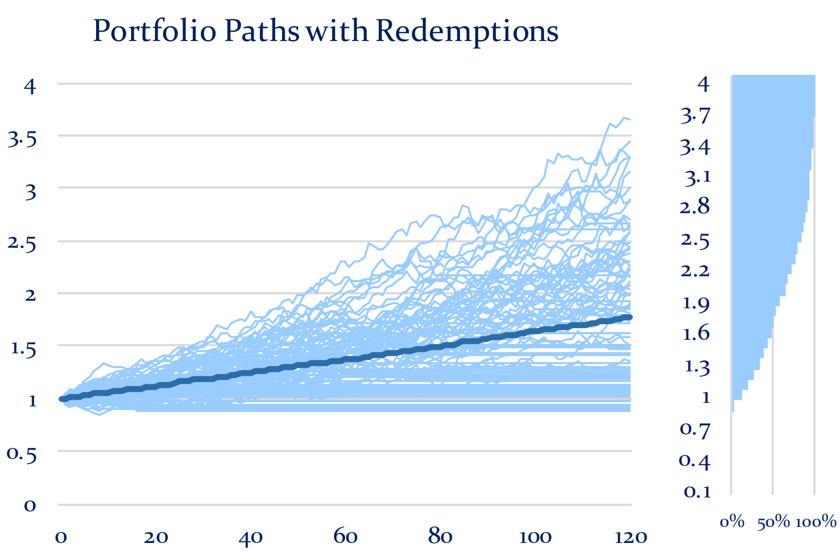

Source: SaT calculations

This second simulation tracks the same 100 paths, but with the two redemption rules applied. The dark blue line showing the average ending wealth is less steep than the average without redemption rules. While the redemption funds could be deployed elsewhere to reduce that gap, the simulation shows the shortcomings of exiting investments when the expected outcome has not changed.

This illustrates that persistence is a very important requirement for portfolio success. While investors need to scrutinize their managers' results for true changes to the underlying return distribution and adjust allocations according, they also need persistence to realize the benefits of a well-constructed investment program.

Conclusion

This paper introduces the core framework to combine strategies into a robust portfolio. We believe that successful multi strategy portfolios are created with five elements: continual research to update sources of returns, quantitative trading strategies to utilize this research, trading across multiple unrelated markets, proper diversification among strategy classifications, and patience in order to benefit from the skill of the portfolio. Continued research is paramount to replace seasoned strategies with new, novel return streams. Quantitative strategies can be classified into one of four categories to more accurately describe their characteristics. Nine varieties of markets provide the widest range of available diversification. And once portfolios have been crafted, it is of great importance to give the investment program time to succeed. When the elements of research, strategies, markets, diversification, and time are applied collectively and correctly, we have built the blueprint of Statistically Applied Trading (Saτ).