



## Risk and Robustness in Portfolio Construction A Path to Long-Term Value Creation

*Alas, the theory is elegant but flawed, as anyone who lived through the booms and busts of the 1990s can now see. The old financial orthodoxy was founded on two critical assumptions in Bachelier’s key model: Price changes are statistically independent, and they are normally distributed. Theory suggests that over that time, there should be fifty-eight days when the Dow moved more than 3.4 percent; in fact, there were 1,001. Theory predicts six days of index swings beyond 4.5 percent; in fact, there were 366. And index swings of more than 7 percent should come once every 300,000 years; in fact, the twentieth century saw forty-eight such days. Truly, a calamitous era that insists on flaunting all predictions. Or, perhaps, our assumptions are wrong.*

*-Benoit B. Mandelbrot*

Portfolio construction<sup>1</sup> is fundamental to boosting returns, diversifying exposure, and managing risk. This is true whether you are an institutional allocator, sophisticated asset manager or individual investor. In this paper, we contrast “Conventional” and “Robust” portfolio construction techniques to demonstrate how Conventional methods fall short in the face of non-linear risk dynamics like changing correlations and volatility. Below we offer several Robust alternatives that can be used collectively to address risk and support the goal of maximizing value creation over time.

### Conventional

We define Conventional methods to include the following:

Conventional approach	Objective
Maximize Sharpe ratio	Boost returns
Control volatility	Manage tail risk
Overlay rules/limits	Manage exposures

These approaches tend to be simple, efficient and easy to monitor. They rely on historical data to model portfolio outcomes. And while no one would argue with the objectives, having a simplistic reliance on historical data to model forward outcomes is, at best, extremely fragile and, at worst, an exercise is chasing one’s proverbial tail.

### Sharpe Pitfalls

Consider the example of Sharpe Ratio maximization, which is mathematically equivalent to optimizing the mean-variance tradeoff. Mean-variance analysis measures an asset’s risk or variance

---

<sup>1</sup>Portfolio construction as discussed herein is the amalgamation of both offensive (investment selection) and defensive (risk management) considerations.

<b>Pros</b>	<b>Cons</b>
Computationally efficient	Overfits in-sample
Statistically tractable	Unrealistic assumptions

relative to the asset’s likely return. This optimization equally weights positive and negative volatility. It is an attractive approach because it seemingly uses a simple and verifiable means to find the best return-versus-risk tradeoff, but the problem is how it is achieved. Conventional implementation of this strategy is extremely dependent on the dataset (time series) being used to estimate “optimal” portfolio holdings. This is not simply a restatement of the common idea that past performance does not guarantee future results. This dependency means that small changes to future market dynamics can result in disproportional underperformance. This fragility is known, though perhaps under-appreciated.

Maximizing Sharpe ratio is equivalent to Ordinary Least Squares (OLS) regression, and the fragility of Sharpe ratio optimization to forward-looking market dynamics is akin to the well-studied issues with overfitting a regression to a particular time series. Conventional Sharpe-ratio maximization or mean-variance construction inherits the classic problems of regression<sup>2</sup>: overfitting the data, estimating with low power, and not adjusting for model selection.

*Illustrating the Problem*

Examine an allocation (long and short) across the five benchmarks shown in Table 1 below over the period 2000 to 2019. Using OLS to get the maximum-Sharpe portfolio amongst constituents, Table 2 shows the deterioration of the Sharpe Ratio resulting from the overfitting of this procedure.

Table 1: Five benchmarks used for illustrative allocation

S&P 500	SPY US Equity
Commodities	SPGSCITR Index
Crude Oil	CL1 Comdty
Currencies	DXY Curncy
Treasury Bonds	IEF US Equity

Table 2: Sharpe Ratios for OLS Allocation

	<b>2000-2016</b>	<b>2017-2019</b>
In-sample	1.31	2.85
Out-of-sample		<b>-0.18</b>

Specifically, Table 2 highlights that while in-sample Sharpe ratios were strong across both subperiods, estimating over 2000-2016 to allocate over 2017-2019 led to terrible performance and a negative Sharpe ratio. Worse yet, it occurs during a period where market returns were extremely strong, as

<sup>2</sup>Some of these problems could be categorized as follows:

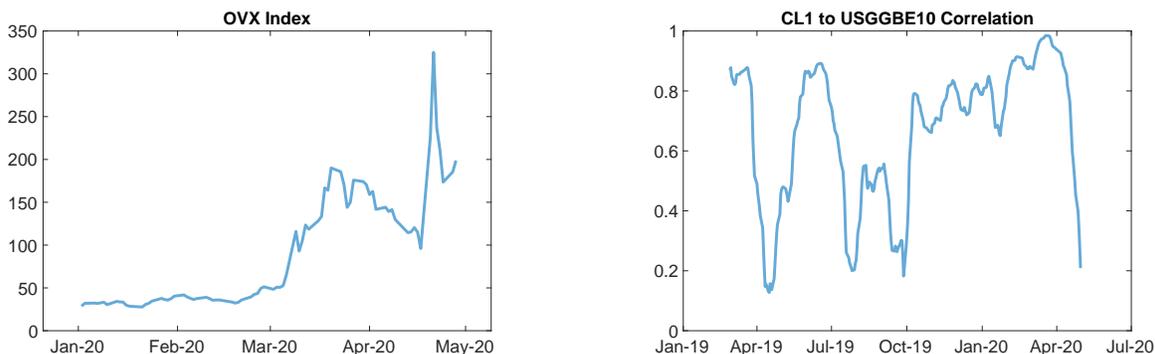
- Overfitting. Estimates include substantial long-short positions meant to provide, diversification and balance risk. But with even small changes to the correlation structure, risk is exacerbated, not diversified.
- Low power. Given correlated data, estimation struggles to determine which positions are significant.
- Selection. Optimization loads into every security available, even when marginally unimportant.

indicated by the in-sample Sharpe for 2017-2019 of 2.85. The problem lies in the fragility of “optimal” estimates. The allocations are heavily dependent on the specific covariance matrix, which is used to make sure various exposures net out. When there are even natural levels of change in these covariances, the optimizer relies on a certain netting of risk exposures, but finds that risk is exacerbated, not diversified, by the linear nature of these models.

### *Non-Linearity Exists*

Linearity, in this context, assumes that the relationship between any two securities can be expressed using a straight-line method. Conventional methods using mean-variance optimization may be optimal with respect to linear models of risk exposure, but the dynamics (changing correlations) of market returns make this assumption tenuous even in simple cases. While linear factor models have proven to be useful tools for portfolio analysis, the assumption of a linear relationship between factor values and expected return is quite restrictive. Specifically, the use of linear models assumes that each factor affects the return independently and hence, they ignore the possible interaction between various factors. Take for example, a traditional long-short equity fund. Even without holding derivative positions, the dynamics of changing correlations, volatility and tail-risk mean that the fund’s holdings will not behave linearly over time. Extraneous asset prices have a dramatic impact.

Over-reliance on static correlations can be problematic to portfolio outcomes. Consider how the recent extremes in oil prices have impacted a range of assets. The first chart below shows the crude oil volatility index (OVX) during the past year has seen a low-high range from 24 to 325. The second chart shows the rolling 20-day correlation of WTI Crude Oil (CL1) and 10-year US Inflation Breakeven Index (USGGBE10). What these charts highlight is the fragility of relationships. As a result of spikes in volatility, any assumed diversification or hedge benefit becomes extremely tenuous. If you extrapolate this example across a broader portfolio the results are even more unstable.



### *Dynamic Risk Measures Are Not Enough*

One common misunderstanding is that this fragility can be solved exclusively with time-varying estimates rather than static estimates. While dynamic measures of risk, including covariance matrices, are necessary, they are not sufficient. Below in Table 3, we examine the performance of these conventional methods with time-varying risk measurements. The dynamic version of the Conventional (OLS) strategy produces a positive Sharpe ratio, unlike what we saw in the static version in Table 2 above. However, it is significantly less than what a Robust covariance estimation can achieve.

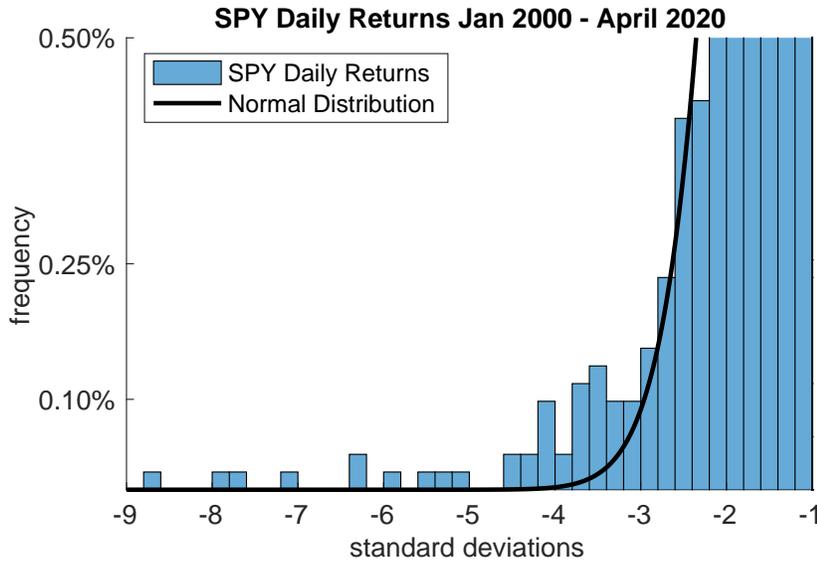
Table 3: Performance Comparison

	Sharpe	Sortino	Beta	Skew	Max DD
Conventional (OLS)	0.3620	0.5017	0.1979	<b>(-0.7438)</b>	0.1432
Robust Covariance	0.5335	0.7526	0.2897	<b>(-0.3817)</b>	0.1154
Tail Risk	0.5306	0.7496	0.2741	<b>(-0.3499)</b>	0.1133

*Average Imperfections*

In addition to fragile Sharpe ratios, conventional optimization techniques attempt to manage tail risk as a by-product of managing volatility. By tail risk, we are referring to a portfolio’s worst losses, whether measured through skewness, maximum drawdown, or other statistics. The first row of Table 3 (above) shows the performance statistics of the conventional “OLS” approach. Skewness is negative and large in magnitude, while maximum drawdown is higher relative to more robust methods.

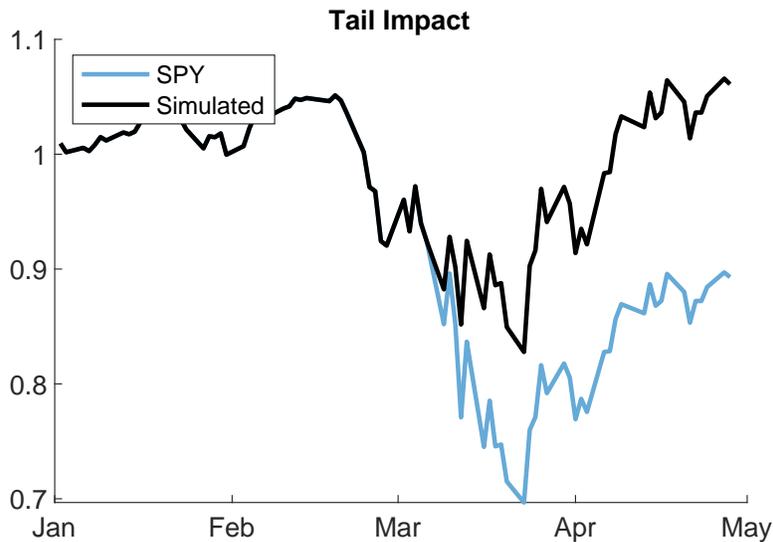
These results reflect the fact that tail-risk is more than a by-product of volatility, and it is thus underestimated. The key fallacy in tying volatility to tail-risk is assuming a Gaussian (normal) distribution for asset returns even though they are known to be strongly non-Gaussian. Note below the difference in tail outcomes between a normal return distribution and the actual return distribution of the S&P over the past 20 years.



*“It’s so much easier to work with that [Gaussian] curve, because everybody knows the properties of that curve, and can make calculations to eight decimal places using that curve. But the only problem is that curve is not applicable to behavior in markets, and people find that out periodically.”*

**- Warren Buffett**

A mean-variance framework which models returns according to a normal distribution leads to an underestimation of tail risk. As volatility and tail risk are not the same, any attempts to manage tail risk by controlling volatility will lead to greater tail risk than anticipated. This is particularly problematic because the distinction reveals itself at the worst times, often exacerbating losses (magnified in leveraged strategies) rather than providing the intended protections. Left tail events lead to a much larger drag on compound returns than localized volatility moves. As a simple example, consider the impact of a left tail event on SPY as shown in the graph below. We compare the performance of SPY in the March 2020 to a risk managed version of SPY which removes excess tail risk. As the figure shows, conventional methods fail to capture excess tail risk leading to substantially lower performance.



It is clear from this illustration that the mitigation of tail events is key to maximizing long-term portfolio performance. And while known, the larger challenge becomes achieving this outcome in a repeatable and efficient way.

#### *Constraints Come with a Cost*

Often the failures of conventional portfolio management are patched over by constraints. It is common for sub-portfolios to be optimized within buckets by asset class or some other categorization. Conventional methods do not diversify in a way that accounts for changing correlation and tail-risk, which leads to poor risk management. Forcing the optimization to handle sub-portfolios ensures that the resulting allocation is spread over a number of these buckets. But this approach gives up optimization of how one bucket may influence another. For strategies seeking to exploit opportunities across sectors, factors, or asset classes, this myopic approach is costly, especially if the buckets are highly correlated. It also introduces a whole new problem: deciding on how to bucket the space, updating this to market dynamics, and figuring out how to aggregate the sub-portfolios into a final portfolio.

In addition to constraining the application of the optimization, it is common to constrain the answers that come out of the model. Again, we have discussed that conventional portfolio techniques may prioritize in-sample Sharpe ratios too highly. This means that the portfolio may be unrealistic given considerations of position sizes, factor exposures, and other risk metrics.

Methods	Implementation
Robust covariance	Ridge Regression
Tail-risk management	Weighted Least Squares
Managed factor exposure	Constrained Regression

## Robust Alternatives

Acknowledging that correlations change, non-linear risk dynamics exist, and volatility is not the best measure or control for tail risk, we present three Robust portfolio construction tools which address the limitations of conventional methods. Each is illustrated with a simple implementation.

### *Modeling the Unknown Through Robust Covariance*

The biggest problem with classic optimization is that the covariance matrix of highly correlated assets is estimated poorly, which makes the allocation extremely sensitive to changes in the correlation structure or mean returns. Consider that a covariance matrix for  $n$  assets is estimating  $(1/2)(n - 1)n$  parameters. Thus, a portfolio of 10 longs and 10 shorts has 45 different dynamic correlation pairs. Robust portfolio construction uses a covariance matrix that better accounts for the estimation uncertainty and results in allocations that do not disproportionately underperform given changes in return dynamics. Finding a robust covariance matrix means more than updating it through time. While dynamic updating is helpful, it does not eliminate the fragility that Sharpe maximization inherits from Conventional OLS, but simply focuses it on a narrower (and hopefully more relevant) slice of historic data.

A Robust approach reformulates the problem to account not just for the volatility of the data, but also the uncertainty of the estimation. The result is less fine tuning for in-sample optimality and more flexibility to achieve better out-of-sample forecasting and risk management. Below, we implement this with a robust covariance matrix. Unlike a conventional approach which uses a typical covariance matrix to engineer diversification, the robust covariance matrix does assume that the historic data is noisy and estimates a less specific covariance matrix. This does not attempt the same fine-tuning of diversification and hedging; and thus, it avoids learning the wrong lessons from the sample data. Instead, it takes a more flexible approach, which focuses on learning what is likely to show up in the future rather than what showed up in the past. We implement this with Ridge Regression below, which can be viewed within a larger set of machine-learning tools meant to prioritize flexible modeling and focus on out-of-sample performance. In fact, the robust covariance matrix can also be derived from Bayesian statistics as the optimal way to combine the noise and signal in the historic data. From either of these perspectives, the robust covariance matrix is the result of accounting for estimation uncertainty and changing market relationships.

Table 4 shows that there is much less out-of-sample deterioration. Out-of-sample results remain strong with a Sharpe of 0.56. Contrast this to classic mean-variance optimization, where we found a Sharpe ratio of -0.18 referenced in Table 2.

Table 4: Sharpe Ratios for Robust Covariance Matrix

	2000-2016	2017-2019
In-sample	0.83	1.04
Out-of-sample		0.56

## Catch the Tails

A key limitation of Sharpe-Ratio maximization can be seen in the fact that OLS penalizes variance, or equivalently volatility, without accounting for tail-risk. Even in the simplest setting of normally distributed returns, volatility is not a sufficient statistic for tail risk due to time-varying covariances and forecasts. And given the distinct non-normality of returns, trying to account for tail-risk through volatility is potentially disastrous. One route to dealing with this is to use conventional portfolio allocation methods, and then overlay some rules or limits to manage tail risk. The problem with this approach is that step 1 (optimizing to volatility) and step 2 (limiting tail risk) may be in conflict. A better approach is to address the problem in reverse by directly scoring tail-risk in the objective function. This means scoring potential allocations for return-per-tail-risk rather than return-per-volatility within the portfolio construction stage rather than ex-post. Mathematically this introduces new complexity in estimation relative to the conventional approach (OLS regression). However, the consistency in portfolio design and improved performance in tail-events is worth the mathematical and statistical complexity.

As an example, consider a solution that manages Conditional Value-at-Risk (CVaR) instead of volatility. We return to the allocation problem with the five securities listed above in Table 1. To avoid conflating the forecasting performance with the risk-management, we keep to a simple expected return framework that uses expanding sample averages. To emphasize that the issue does not arise simply from static estimates, the strategy is rebalanced monthly using daily data.

Table 5 shows the summary statistics for a portfolio using the Conventional OLS strategy against portfolios with various robust improvements. All are constructed with the five securities in Table 1. The first two rows compare the Conventional OLS and robust covariance matrix approaches, but with dynamic estimates. Again, we see that the robust covariance estimator greatly improves performance—not just with regard to Sharpe ratio, but also tail-risk as measured in skewness and maximum drawdown. In the row titled Tail Risk, note how optimizing for tail-risk directly shows out-of-sample improvement over classic methods. It not only improves tail risk, but also improves on the Sharpe ratio—the very thing that the conventional approach was supposed to be optimizing.

Table 5: Portfolio Comparison

	Sharpe	Sortino	Beta	Skew	Max DD
Conventional (OLS)	0.3620	0.5017	0.1979	<b>(-0.7438)</b>	0.1432
Robust Covariance	0.5335	0.7526	0.2897	<b>(-0.3817)</b>	0.1154
Tail Risk	0.5306	0.7496	0.2741	<b>(-0.3499)</b>	0.1133
Constrained Exposures	0.7025	0.9949	0.1467	<b>(-0.2632)</b>	0.1011

## Integrated Constraints Lead to Further Improvements

The robust covariance matrix and tail-risk objective each substantially improved Sharpe ratio, Sortino ratio and tail-risk, while slightly raising the market beta relative to the Conventional strategy. For this, a third approach is useful. The Conventional approach of managing exposure limits allows for unbounded factor exposures and positions according to a logic that these are means to an end. Given the fragility of the solution, these are problematic, so a common approach is to layer on ad-hoc limits

after the optimization. Like with Tail Risk, separating exposure constraints from construction leads to inconsistencies and suboptimal performance. Instead, our third method uses constrained optimization to integrate the factor exposure limits into the broader objective, allowing for better optimization and robustness. One reason the two-step, inconsistent approach has been popular is computational efficiency. But with modern constrained optimization methods, it is more feasible than ever to deal with this directly in portfolio construction. As an example, we return to our five-benchmark portfolio and optimize according to Managed Exposures. Table 5 above shows this procedure outperforms the conventional OLS regression in all phases: return-risk ratios, factor exposure, and tail risk.

Too often, portfolio construction is based on maximizing Sharpe ratios, managing volatility, and layering on risk constraints ex-post. Portfolio construction with these conventional methods pursues worthwhile objectives but makes significant compromises that lead to fragility and underperformance out-of-sample. Portfolio optimization using robust methods is better suited to achieve diversification and tail risk management. As seen in the five-benchmark examples, each of these approaches leads to improved performance ratios, factor exposures, and tail risk. And even better, these methods can be used together if done carefully, to compound improvements to conventional portfolio analysis and pave a more direct path to long term wealth creation.

#### **Disclosures:**

This report has been prepared by the team at Racon Capital Partners LLC (the “Investment Manager”). This communication does not constitute an offer for investment. Offer for investments may only be made by means of an applicable confidential private offering memorandum (“CPOM”) which contains important information (including investment objective, policies, risk factors, fees, tax implications and relevant qualifications), and only in those jurisdictions where permitted by law, by permission of the Investment Manager’s senior management following the completion of an investor suitability analysis. Any investment in the funds managed by the Investment Manager involves a high degree of risk and may not be suitable for all investors. Before making an investment decision, potential investors are advised to carefully read the applicable CPOM, the limited partnership agreement (if applicable), and the related subscription documents, and to consult with their tax, legal and financial advisors. While all of the information prepared in this document is believed to be accurate, the Investment Manager makes no express warranty as to its completeness or accuracy, nor can it accept responsibility for errors appearing in the document. The manager relied on certain third-party sources of information when completing this report. The details, parameters and descriptions found in this document may be modified in the future by the Investment Manager in any manner. The information in this report is provided for informational purposes only and should be not be relied upon or construed as creating formal investment rules or requirements. In the case of any inconsistency between the descriptions or terms in this document and the CPOM, the CPOM shall control. Opinions expressed by the author(s) of this article do not necessarily reflect those of the Investment Manager or its affiliates. Examples in this report are not intended to indicate performance that may be expected to be achieved. Past performance is not indicative or a guarantee of future results. No assurance can be made that profits will be achieved or that substantial or complete losses will not be incurred. The Investment Manager, its affiliates and employees do not provide tax, accounting or legal advice. Please consult your tax, legal or accounting advisor regarding such matters.

#### **Racon Capital Partners LLC**

234 W. Florida St. — Suite 700 — Milwaukee, WI 53204

Phone 414.877.1270 — [www.raconcapital.com](http://www.raconcapital.com)

© Racon Capital Partners LLC