Robust by Design: A New Approach to Strategic Asset Allocation

Keith R. Collier, CFA
Director, Asset Allocation Research BNY Mellon Investor Solutions

Today's institutional investors face an increasingly complex set of challenges that often require customized strategic policy portfolios tailored to their specific needs. To address these challenges, BNY Mellon Investor Solutions created a proprietary process that is unique in the industry: Robust Strategic Asset Allocation (RSAA).

How Robust SAA is Different

Recent investment industry trends have seen an increasing move toward model portfolio solutions that support scalability, open architecture and a platform for customization. The common goal of many of these investment portfolio offerings is to maximize overall utility (typically return) with the least amount of risk. While this is a reasonable goal, there are limitations to this line of portfolio construction, and sophisticated institutional clients often face more complex challenges. In a novel and bold step, our portfolio design methodology goes beyond simply balancing utility and risk-appetites. Our approach, which we call Robust SAA (RSAA), systematically designs custom portfolios targeting any number of investment objectives with resilience to a range of market conditions.

We Start with Three Goals in Mind

Our RSAA process places great emphasis on both the structural advantages and asymmetrical risks of the underlying asset classes. Additionally, our methodology attempts to achieve scalability and systemization, while still allowing for customized portfolio nuances. In summary, our RSAA process targets three goals:

1. **Robustness:** Identify strategic allocations that are relatively immune, at the portfolio level, to forecast errors or annual changes in the Capital Market Assumptions (CMAs), rather than optimizing around point-forecasts for the component asset classes;
2. **Risk Mitigation**: Explicit consideration of potential tail risk scenarios and shortfall probabilities;
3. **Customization**: Incorporate multiple client-driven objectives (such as return, drawdown, yield, etc.) and custom constraints on portfolio allocation weights.

Robust Rather Than Optimal

A key differentiator of our process is a preference for "robust" rather than "optimal" portfolios. We define a *robust portfolio* as the allocation that minimizes the adverse performance shortfall (for a given probability threshold) versus all other feasible portfolio candidates considered.

Because CMAs are influenced by numerous factors over their forecast horizon, they are never intended to perfectly predict the future and will always result in some margin of error between the "forecast" and the "actual." Ideally, we want our SAA portfolio weights to be relatively stable in the face of modest CMA changes year to year, while at the same time remaining near-optimal according to Modern Portfolio Theory (MPT). In practice, we achieve this stability and resilience by comparing "clouds" of many possible outcomes that contain the yet unknown outcome, rather than optimizing to a single-point forecast that will most likely miss the mark. By searching for the most resilient portfolio solutions within hundreds of possibilities – instead of a single "hoped-for" scenario – our robust portfolio solution is less impacted by any arbitrary shift or error in the point-forecast or "average" expectation.

To address the inherent error in CMAs, we use a quantitative method called perturbation to introduce "noise" into the forecasted numbers – thereby simulating a range of uncertainty in the scenarios that might occur. We then apply this technique en masse to generate a whole distribution of possibilities across a bounded range of plausible outcomes for each performance characteristic (return, volatility, yield, correlation) of every asset class in our CMA lineup.

By perturbing the point estimates into a cloud of uncertainty, we can build and compare multiple iterations of "candidate portfolio" solutions. These candidate portfolios are supplied by a high-throughput random portfolio generation algorithm that only supplies portfolio candidates that already satisfy any custom allocation constraints desired for the mandate. We then evaluate and compare each candidate's expected performance at various tail thresholds across the range of potential outcomes, and choose the candidate with the least adverse performance for the specified probability. Because distributions contain so much more information (in terms of data points), comparing distributions is obviously more robust and less error-sensitive than relying on a single representative value. By evaluating portfolio performance under adverse tail thresholds (at 20%, or even 5% likelihood), BNY Mellon Investor Solutions can create custom asset allocations to weather even highly suboptimal market environments.

Distribution of Scenario Outcomes for Two Methods of Portfolio Design

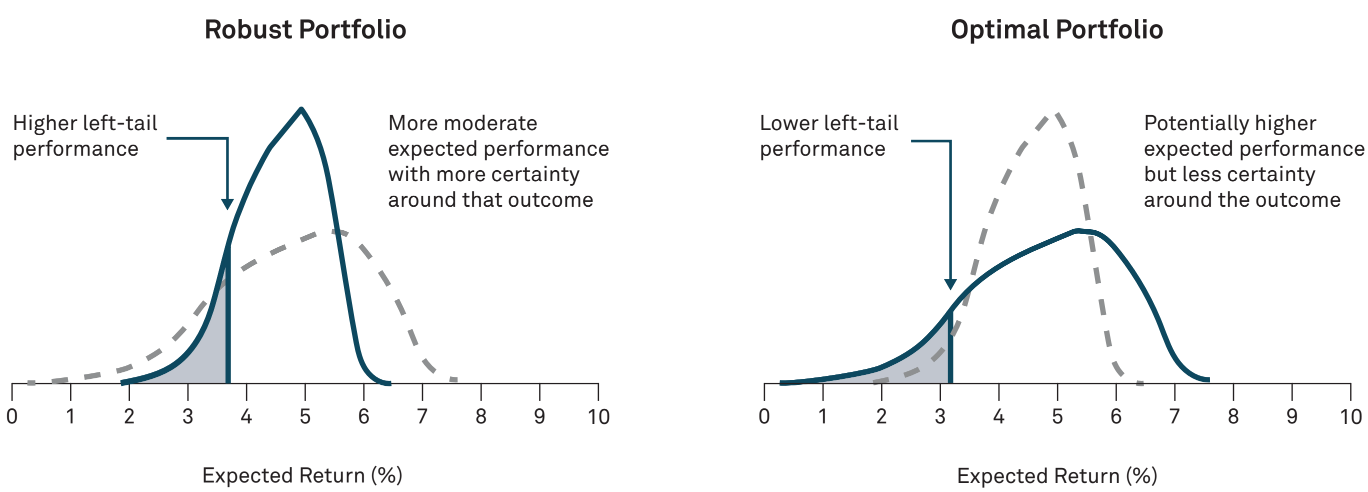

Source: BNY Mellon Investor Solutions. Charts are provided for illustrative purposes and are not indicative of the past or future performance of any BNY Mellon product. Projections or forecasts regarding future events, targets or expectations, are subject to change. There is no assurance that such events or expectations will be achieved, and actual results may be significantly different.

Using Machine Learning to Find the Most Robust Portfolio

To rapidly identify the most robust portfolio, we utilize a machine learning search and optimization technique called Simulated Annealing ("annealing"). This is an iterative method that explores and considers varying degrees of suboptimal solutions in the pursuit of discovering a practical, near-optimal solution – within a finite number of trials.

The efficiency of annealing stems from its hybridization of two opposing and complimentary search strategies: stochastic search and gradient descent. Stochastic search is roughly analogous to randomly testing multiple solutions and settling upon the most optimal solution over a finite number of trials. In contrast, gradient descent is a "greedy" algorithm that only accepts a challenger candidate solution if its value is superior to the existing "best" solution – by even a small margin.

Since the annealing algorithm is an iterative process, it repeatedly performs a simple set of computational steps over hundreds, or even thousands, of cycles. Early in the process, each iteration allows the algorithm to explore explicitly sub-optimal candidate portfolios on the premise that they may eventually lead to robust solutions. As the annealing search process continues, the algorithm becomes less exploratory and more pragmatic – only favoring allocation changes that yield immediately obvious performance improvement. The overall effect is the gradual "refinement" of an acceptable but arbitrary initial portfolio into a highly robust and stable portfolio solution when the process completes.

Customization is the Key

The ability of our process to consider different dimensions of clients' objectives and constraints is a key differentiator of this methodology from conventional industry practice.

Many model portfolio solution providers today still use traditional mean variance optimization (MVO) or related risk-based methodologies. However, a widely known weakness of MVO is its sensitivity to input parameters (expected return, volatility, and correlation forecasts). This means that the optimal portfolio allocation weights can change dramatically even if forecast inputs change only marginally. In practice, this is quite suboptimal from a model stability standpoint and creates a great deal of spurious turnover and transaction costs.

Additionally, MVO's "best" portfolio is always based on optimized return per unit of risk (Sharpe ratio). This is fine if the client only cares about the two objectives of return and volatility. But what if yield or turnover matters – or even drawdown – for three, four, or more objectives? Unlike MVO, our approach allows us to consider an unlimited number of portfolio utility objectives, and we can even custom weight these objectives in cases where an unequal preference between trade-offs is desired. In addition to investment objectives, as stated earlier, any number of absolute or relative asset-level and asset group-level allocation weight constraints can be applied to the candidate search set to incorporate any behavioral or commercial realities the client wishes to impose.

The Robust SAA Process

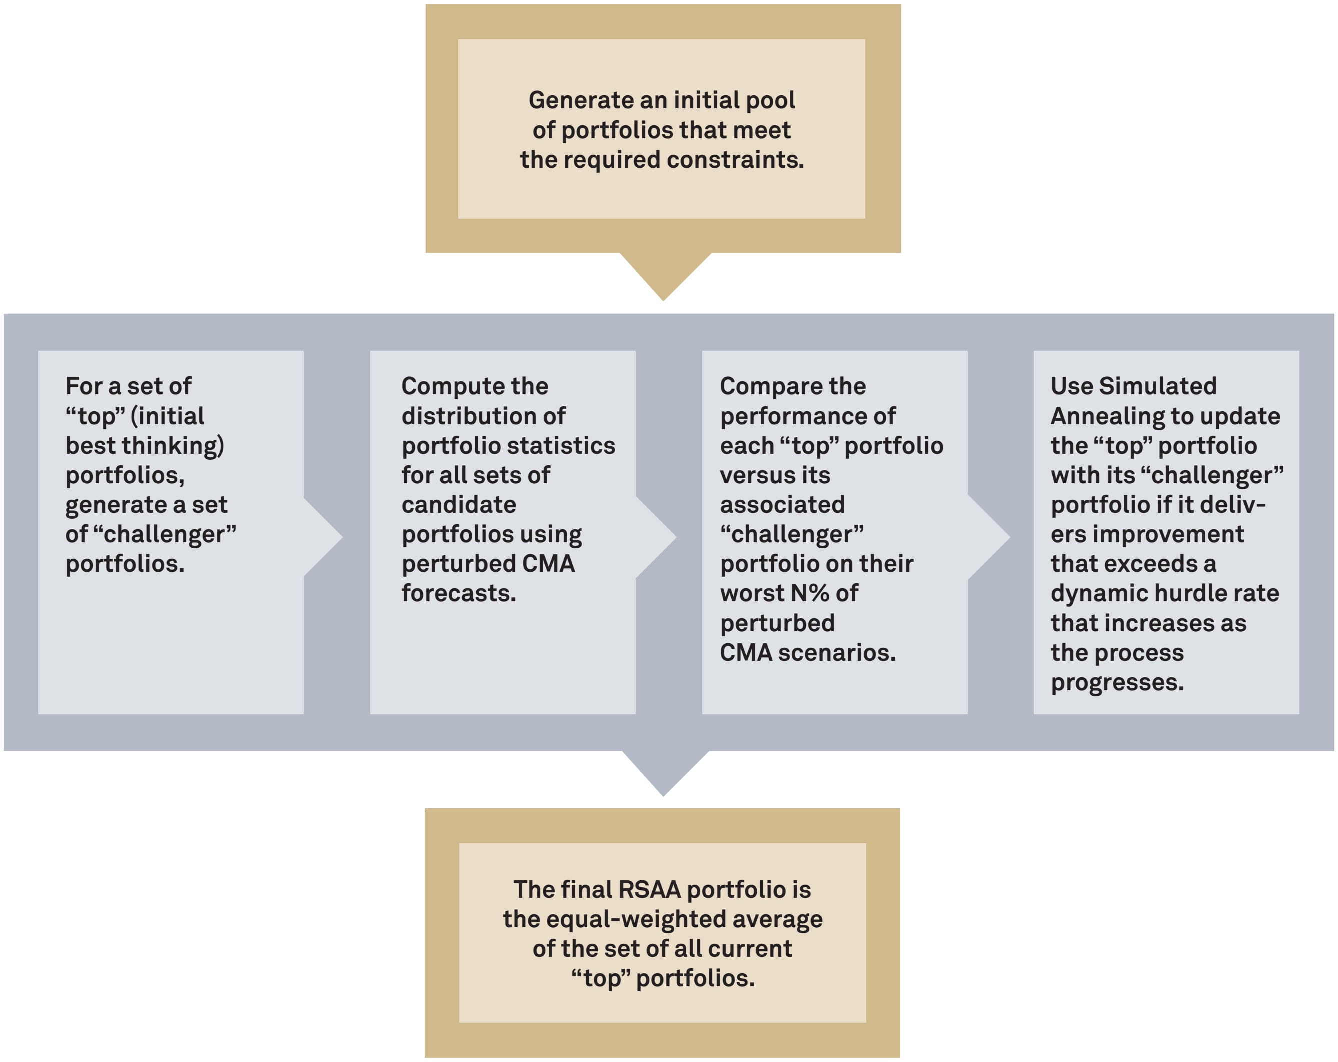

Benefits to Your Investment Program

Given today's volatile and dynamic investment environment – with the shifting sands of interest rates, inflation, and growth – investors may need to confront the possibility of higher uncertainty in their future investment outcomes. Small year-to-year forecast deviations and high-impact "Black Swan" events are both inescapable realities of asset allocation, not remote contingencies that can be reduced or eliminated with better predictions. Our RSAA process acknowledges these realities with a sober and pragmatic analysis of a wider risk landscape. Even in times of high uncertainty, some allocation choices are still clearly better than others.

In summary, our RSAA process elegantly delivers three key features: resilience to adverse market conditions with asymmetric outcomes, allowance for multiple competing investment objectives, and controls for customized allocation constraints. We believe that by reducing tail risks and improving reliability, we can craft a customized strategic investment program that helps your organization better achieve its long-term financial goals.

Read our full white paper:

https://im.bnymellon.com/us/en/documents/manual/perspectives/designing-robust-strategic-asset-allocations.pdf

Keith Collier, CFA, is Director of Asset Allocation Research for BNY Mellon Investor Solutions. In this role, he oversees the team's research agenda, process improvement, and strategic initiatives. His quantitative investment background covers systematic asset allocation, model portfolio design, data science/machine learning, product management, and digital strategy. Keith joined the firm in 2015, and has more than thirteen years of quantitative investment experience.

Prior to joining BNY Mellon, he worked at a boutique hedge fund startup and Legg Mason's quantitative affiliate, Batterymarch Financial Management, where he designed the firm's first ETF-based tactical asset allocation strategy. Prior to his investment career, Keith worked in real estate development and technical sales and marketing.

Keith received an MBA degree from Cornell University. He also earned Bachelor's and Master's degrees in Architectural Engineering from Penn State University and is a graduate of the Schreyer Honors College. He is a CFA charterholder and a member of the CFA Institute, CFA Society New York, and the Society of Quantitative Analysts.

Investor Solutions

Our mission is to leverage BNY Mellon's strengths as a leading asset management, wealth management and investment servicing firm to bring Outsourced Chief Investment Officer (OCIO) and advisory solutions to institutional investors worldwide. We serve endowments, foundations, retirement plans, family offices, governments and financial intermediaries.

BNY MELLON