PORTFOLIO OPTIMIZATION FOR EFFICIENT STOCK PORTFOLIOS:

APPLICATIONS AND DIRECTIONS

ReSolve
ASSET MANAGEMENT

TABLE OF CONTENTS

Contents

Introduction	3
Summary of findings	4
The Right Tool for the Job	4
Alternatives to Cap Weighting	6
Evolution of Passive Investing	6
Optimal Stock Portfolios	8
Capitalization Weighted Portfolio	9
Risk Efficient Portfolio Decision Trees	12
Data and Methodology	13
Results	14
Leveraging efficient portfolios	17
Bootstrap Tests	18
Portfolio Concentration	21
Conclusion	25

INTRODUCTION

With record-breaking and accelerating flows to cap weighted stock indices there has never been a more urgent need to fundamentally explore the concept of passive investing and the "market" portfolio.

We start from the assumption that investors seek the greatest return for the least amount of risk. While many investors would describe risk as "probability of loss" or perhaps "probability of not achieving financial objectives", we take the position that more conventional measures like volatility and beta are directionally consistent with these objectives. Moreover, volatility and beta have the advantage of being directly observable.

As such, we assert that investors want "efficient" portfolios, which may be expected to produce the highest ratio of excess returns to risk. We will measure the realized efficiency of portfolios in several ways to address different measures of "risk".

The purpose of this study is to demonstrate why the theoretical foundations for conventional notions of "passive" investing – most notably market-capitalization weighted stock portfolios – are profoundly inconsistent with observed market behavior. Quite simply, the market cap weighted portfolio is shockingly inefficient relative to other reasonable alternatives.

We leverage the toolset brought to bear in our "Portfolio Optimization Machine" framework to explore alternative "risk-efficient" methods for equity investing. Like capitalization-weighted portfolios, other risk-efficient portfolios express simple views about relationships between risk and return. However, certain alternative approaches express views that are entirely consistent with the actual relationships between risk and returns observed in stocks over the past half-century. These methods have delivered excess risk-adjusted performance consistent with expectations.

Our approach is motivated by the objectives of a professional investor, so we prioritized empirical reality over academic theory. This is critical to our objectives because the historical record has broadly failed to confirm most dominant theories about the relationships between risk and return in stocks.

To ensure even very large allocators can employ the methods under investigation at scale, we limit our analysis to large capitalization, liquid securities[1]. We also take steps to limit turnover and tax consequences by trading overlapping monthly portfolios with annual holding periods.

We source data from the Center for Research in Securities Prices (CRSP) that included daily prices and total returns for all currently listed and delisted U.S. equity securities from 1960 through the end of 2019. We chose our start date based on the number of missing data points in daily return data, which is corrosive to estimates of covariance.

[1] Stocks were ranked by market cap as of December 31st of each year and the top 1000 stocks by market capitalization at the end of year t were considered eligible in year t+1

Our study covers two broad themes: (1) Risk-efficient optimizations to form mean-variance optimal alternatives to naive market capitalization or equal weighted portfolios without active views; and (2) Optimal factor portfolios formed by taking factor scores as active views in mean-variance optimizations.

This article explores the first theme in detail. We have also performed extensive analysis on the second theme, which we will document in a follow-on article.

SUMMARY OF FINDINGS

We find that optimal risk-efficient portfolios formed from liquid large-cap securities with overlapping annual rebalancing have produced very economically and statistically significant performance improvements over naive portfolios in the sample period. Optimal portfolios and naive portfolios produced similar returns, but appropriately specified optimal portfolios produced their returns with substantially less risk.

Investors in minimum variance type portfolios experienced less than half the volatility and just over two-thirds the maximum drawdown of the value-weighted portfolio without sacrificing returns. In addition, minimum portfolios were almost 50% more efficient than the equal weight portfolio as measured by Sharpe ratios[2].

When we scaled exposure (borrowing at LIBOR) to the concentration constrained minimum variance portfolio by 1.3x[3] the strategy produced 12.12% annualized return, compared with 9.86% produced by the cap weighted portfolio, with materially lower volatility, average drawdowns, conditional value at risk, and maximum drawdown.

THE RIGHT TOOL FOR THE JOB

Investors typically gravitate to passive investing because of lower costs in terms of fees and expected taxes, and to minimize active decision making where they have a perceived disadvantage.

In reality, passive investors in the cap weighted portfolio do not abdicate views entirely. Rather, they avoid having to express views on each individual security by collapsing their views to a single assumption about the relationship between risk and return. These views follow from the Capital Asset Pricing Model (CAPM), which asserts that the only risk that is compensated in markets is market risk, which produces the Equity Risk Premium[4].

Under the CAPM investors price investments so that expected returns for a stock will be proportional to the stock's beta, where beta is a stock's sensitivity to market excess returns. In theory, stocks with a beta greater than 1 would be expected to

2 An important benefit of more efficient portfolios is the ability to scale portfolio exposure using leverage to achieve higher returns with market-like risk. However, most investors can't borrow at the risk-free government rate. Rather, most funds and institutions must borrow at a rate that accounts for overnight credit risk. For this reason, all excess returns and performance statistics that incorporate a risk-free rate are reported net of LIBOR. In addition, we assume investors would pay LIBOR rates on leverage so the results we report should have been generally achievable by funds and institutions.

3 Traditional 40-act mutual funds have the ability to lever portfolios by up to 30 percent without seeking exemptions from the SEC.

4 See Treynor, J. L. (1962). "Toward a Theory of Market Value of Risky Assets."; Sharpe, W. F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk."; and Lintner, J. (1965b). "Securities Prices, Risk, and Maximal Gains from Diversification."

produce higher returns than the market to compensate for the higher risk. If this risk-efficient relationship holds then market-cap weighted stock portfolios are mean-variance optimal.

Sadly, in reality market beta has had no relationship with stock returns in the historical record, so cap weighted portfolios are risk-***in***efficient.

Perhaps surprisingly, leading academics acknowledged this flaw in CAPM almost forty years ago. Fama and French demonstrated a lack of relationship between market beta and returns in Table 1. (not shown) of their seminal 1992 paper on the "Cross-Section of Stock Returns". They showed that for stocks with similar market capitalizations, low risk stocks have had the same return as high risk stocks. In Figure 1 we plot stock returns as a function of their ex post beta by decile using updated data from 1962 through 2017 [5]. While CAPM assumes an upward sloping line, the relationship has a slope that is statistically indistinguishable from zero.

This is an insurmountable challenge for the CAPM and leads to the inescapable conclusion that market cap weighted portfolios are inefficient. Many investors are over-exposed to market beta despite the fact that this exposure has never produced higher returns.

A review of the literature[6] makes clear that there is no relationship between risk or return - measured as either beta or volatility - for international or emerging market stocks either.

Figure 1: Average monthly excess return vs. ex-post beta of large cap (top 40% by market cap) U.S. equity portfolios, sorted by beta deciles: 1962 – 2017.

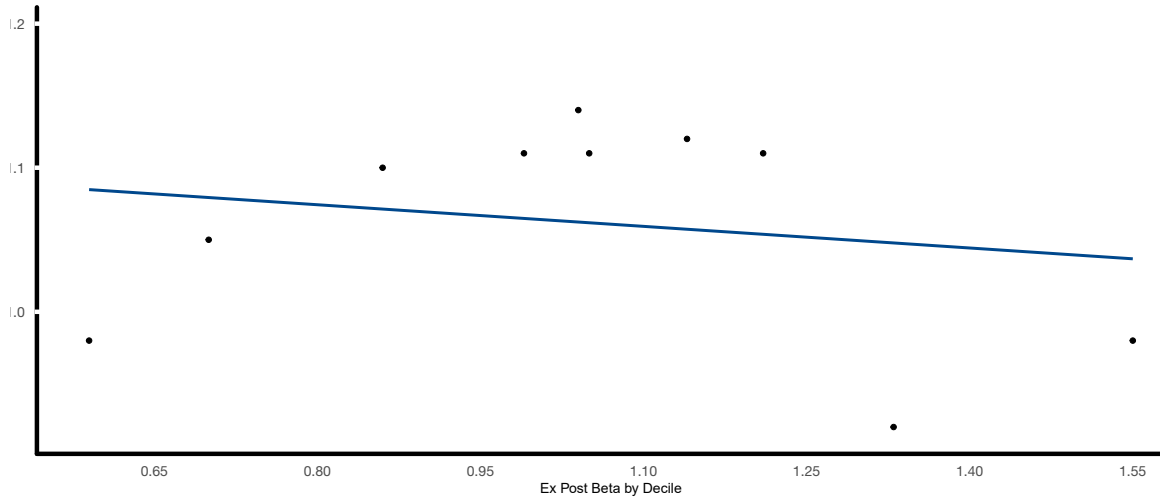

Source: Calculations by ReSolve Asset Management. Data for U.S. decile portfolios sorted on beta from Molyboga and Qian (unpublished 2018).

5 This study was conducted separately by a different team with data through 2017.

6 See "The Volatility Effect" and "Is the Relation between Volatility and Expected Stock Returns Positive, Flat or Negative?" by Pim van Vliet and the team at Robeco, and "Betting Against Betting Against Beta" by Novy Marx.

The death of the Capital Asset Pricing Model prompts an all-important question - If CAPM doesn't hold, and the market cap-weighted portfolio is inefficient, what is the optimal passive stock portfolio?

ALTERNATIVES TO CAP WEIGHTING

A minimally contentious definition of "passive investing" might emphasize the quality that a passive investment strategy requires minimal trading. From this perspective any investment approach where a portfolio is formed at inception and never touched would be considered "passive".

A passive market portfolio would theoretically purchase all risky assets at inception and never trade them. The constitution of the portfolio would ebb and flow over time commensurate with changes in relative market values. The value of many assets would go to zero while a small number of assets would grow from zero to emerge as victors, consistent with the creative destruction encouraged by capitalist competition.

In practice, investable implementations of the passive market portfolio require some trading to dispose of zero-value assets; purchase interests in new entities (IPOs); absorb secondary offerings from existing firms (SEOs); and, re-invest cash distributions. Investible market portfolios also must adjust for business splits/divestitures; company acquisitions and; re-domiciliations. These are non-trivial challenges to the concept of passive investing.

Pedersen[7] shows that maintenance of exposure to the CRSP US equity universe and Russell 2000 index require about twelve percent and forty percent annual turnover, respectively. In addition, Pedersen found that investors who purchase the entire US stock market (CRSP) and do not participate in any IPOs, SEOs, or share repurchases and do not reinvest any dividends, will gradually own a smaller and smaller fraction of the market. After just ten years a typical buy-and-hold investor will own just 60-70% of the stock market.

On the practical assumption that investors will in fact rebalance to account for IPOs, SEOs and reinvested distributions, perhaps the most virtuous property of the market portfolio is that it will always hold positive exposure to the victors in the capitalist ecosystem. This may be the most advantageous property of market-cap oriented investing.

EVOLUTION OF PASSIVE INVESTING

For many years passive investing was unambiguously synonymous with market cap weighting. In the most recent decade the definition has evolved so that Investopedia now describes passive investing as "the opposite of active management in which a fund's manager(s) attempt to beat the market with various investing strategies and buying/selling decisions of a portfolio's securities. Passive management is also referred to as 'passive strategy,' 'passive investing,' or 'index investing.' "

In this context, passive investing is for all practical purposes defined by what it is not, i.e. it is not active management, where active management is defined as active security selection.

7 Pedersen, Lasse Heje. 2018. Sharpening the Arithmetic of Active Management. Financial Analysts Journal 74 (1): 21-36

Once we concede that portfolios that deviate from total market-cap weighting can be passive, the concept of active versus passive becomes nebulous. In fact, it's even hard to draw lines when we agree on market-cap weighting. Is an S&P 500 index fund passive? After all, this index ignores over eighty percent of all listed companies. What about a fund that tracks the Russell 2000 small-cap index? It is market-cap weighted, but represents less than twenty percent of total market value.

It's a trivial step to extend the passive classification to other forms of indexing. In fact it's common for practitioners to refer to strategies that invest in virtually any "index" product as passive. When strategies that take extreme positions in a small number of stocks - as many so-called "factor" or "smart beta" based index funds do - are considered passive investments the term starts to lose all meaning.

At root, the advantage of the CAPM was that it abstracted away the necessity of having active views on the expected returns of individual securities by assuming that returns are a tight linear function of securities' covariance with the market, i.e. their beta. Under this condition the market-cap weighted portfolio would have the highest expected return while minimizing portfolio volatility through diversification.

While we demonstrated that this relationship between risk and return does not hold in practice, we take the position that the fundamental ambition of the CAPM - abstracting away the necessity of active views on expected returns - is the cornerstone of passive investing.

Given tribal attachments to the term "passive" we'll sidestep the drama and shift our discussion back to risk-efficient portfolios in general. Remember, the cap weighted portfolio is just one attempt at creating a risk-efficient portfolio; albeit one that expresses an invalid relationship between risk and return.

Our goal is to identify risk-efficient "market portfolios" that investors can own instead of the inefficient cap-weighted portfolio. A risk efficient market portfolio would have the following characteristics:

1. Avoids active views on expected stock returns for individual stocks

In this context, active views are defined as idiosyncratic views on stocks that imply certain stocks have higher(lower) expected returns due to fundamental characteristics. It is the idiosyncratic nature of active views that distinguishes active management from "risk-efficient" management. The latter implies that stocks may have different expected returns, but that those differences are systematically explained by a linear risk model (like the CAPM).

2. Expresses a consistent relationship between risk and return

The CAPM expresses a consistent relationship between risk and return; excess return is proportional to a stock's beta to the market. It's just that this relationship is not validated empirically. Rather, stocks have consistently exhibited the property that returns are independent of risk (beta or volatility). This is a consistent relationship with a slope of zero.

3. Seeks to diversify across the independent sources of risk and return in the market

A market portfolio should be designed to be representative of the many drivers of economic production. The cap weighted portfolio holds all stocks and therefore has exposure to every public company. However, there are reasons to believe the cap weighted portfolio is not very diversified.

There are a variety of methods to create risk-efficient portfolios. Each method is optimal under different kinds of expected risk-return relationships. We explore some popular risk-efficient portfolio construction approaches below.

OPTIMAL STOCK PORTFOLIOS

Our paper, The Portfolio Optimization Machine was motivated by a desire to link portfolio construction with appropriate views on the relationship between risks and returns. It prescribes a roadmap that any investor can use to form optimal portfolios based on their own assumptions regarding any investment universe.

Figure 2 lists passive optimization methods and describes 1) the objective of each optimization; 2) the parameters that are required as inputs; 3) what the optimization expresses about the relationship between risk and return; and 4) what the portfolio typically looks like when it's formed.

Figure 2: Summary of passive portfolio methods.

	Objective	Parameters	Optimality Conditions	Portfolio Characteristics
Equal Weight	Equlize capital weights	Number of Assets	Similar returns and covariances	Minimal concentration, highly sensitive to investment universe
Inverse Volatility	Weights proportional to inverse volatility	Asset volatilities	Similar Sharpe ratios and homogeneous correlations	Low concentration, highly sensitive to investment universe
Capitalization Weight	Weights proportional to market cap	Covariance with market portfolio	Similar Treynor ratios	Low Concentration, sensitive to investment universe
Equal Risk Contribution	Equalize risk contributions	Covariance matrix	Similar marginal Sharpe ratios	Low concetration, moderately sensitive to covariances, moderately sensitive to investment universe
Maximum Diversification	Maximize portfolio diversification	Covariance matrix	Similar Sharpe ratios	Sometimes concentrated, sensitive to correlations, moderately sensitive to volatilities, robust to investment universe
Minimum Variance	Minimize portfolio volatility	Covariance matrix	Similar returns	Sometimes concentrated, sensitive to covariances, robust to investment universe
Mean-Variance	Maximize return subject to risk tolerance	Expected returns + Covariance matrix	NA	Sometimes concentrated, sensitive to return and covariance estimates, robust to investment universe

Source: ReSolve Asset Management. For illustrative purposes only.

Below we describe the risk-efficient portfolio construction methods in Figure 2 in more detail. Again, different portfolio construction techniques are appropriate in different circumstances that are informed by expected relationships between risk and return. Only one of these optimizations will be optimal for a unique combination of investment universe and assumptions about risk and return.

> **Critically, if we have strong confidence in our assumptions we can know ex ante which method should produce the most efficient portfolios. In this context, empirical simulations serve as hypothesis tests rather than data-mining exercises.**

In the next section we will state our hypothesis about which of these approaches is likely to be most efficient in the formation of stock market portfolios.

Note that in the following formulas w is a vector of portfolio weights with length n equal to the number of stocks under consideration. σ is a vector of estimated standard deviations and Σ denotes the estimated covariance matrix. Superscript ^{T} implies a transpose.

Equal Weight Portfolio

$$w = \frac{1}{n}$$

Intuitively, the equal weight portfolio is optimal when an investor knows nothing about the characteristics of the underlying investments. It follows that this is equivalent to assuming that all of the assets have the same expected volatilities, correlations and returns.

These assumptions may be perfectly legitimate in certain circumstances where there is little data available and/or the assets are near substitutes for one another. For example, stocks in the same sector with similar market capitalizations.

Inverse Volatility Portfolio

$$w = \frac{1/\sigma}{\sum_{i=1}^{n} 1/\sigma_i}$$

An inverse volatility weighted portfolio is constructed so that assets are weighted in inverse proportion to their relative volatilities, so it requires volatility estimates. It is optimal when asset returns are expected to have a close linear relationship with volatility - in other words, when assets are expected to have the same Sharpe ratio.

Importantly, the inverse volatility portfolio is only optimal when assets have similar pairwise correlations. If some assets are highly correlated while other assets are uncorrelated or negatively correlated the inverse volatility portfolio will be sub-optimal. Since inverse volatility weighting doesn't account for correlations it is less effective for optimizing on diverse universes like futures or global asset classes.

Inverse volatility weighted portfolios have the comforting quality of owning all of the assets, though highly volatile assets will necessarily have more modest weightings.

CAPITALIZATION WEIGHTED PORTFOLIO

We touched on capitalization weighted portfolios above, where each asset is held in proportion to its market value.
Market cap weighted portfolios are mean-variance optimal only if investments are expected to produce returns in proportion to their market beta. In other words cap weighting schemes assume stocks have equal Treynor ratios.

Market cap weighted portfolios can be extremely concentrated. For example just 25 stocks in the S&P 500 - 5% of stocks - currently account for over a quarter of the portfolio and the lowest weighted 250 stocks account for less than 12% of portfolio weight.

Risk Parity Portfolios

Equal Risk Contribution (ERC) and Maximum Diversification (MAX-DIV) portfolios typically fall under the category of "risk parity" though they optimize on slightly different objectives.

$$w = arg\min \frac{1}{2} w^T \cdot \Sigma \cdot w - \frac{1}{n} \sum_{i=1}^{n} \ln(w_i)$$

Note that the $arg\min$ terms describes the objective to minimize the value of the function.

As its name suggests, equal risk contribution optimization seeks portfolio weights that ensure all assets contribute an equal amount of volatility to the portfolio after accounting for diversification. The portfolio is optimal only on the condition that portfolio constituents are expected to produce returns in proportion to their total risk contribution to the portfolio, and when assets have diverse correlation relationships.

In its long-only form the ERC portfolio will always hold all assets in positive weight, though assets with high volatility and/or high correlation with other assets in the portfolio will have less weight, and vice versa. As such, the ERC optimization is appropriate for universes with diverse assets and where an investor is concerned with total risk rather than systematic sources of risk.

$$w = arg\max \frac{w \times \sigma}{\sqrt{(w^T \cdot \Sigma \cdot w)}}$$

The Maximum Diversification optimization seeks portfolio weights which maximize the Diversification Ratio, and is only optimal when securities are expected to have equal Sharpe ratios. In this way it is similar to the Inverse Volatility portfolio. However, where the Inverse Volatility portfolio imposes the view that securities all have similar pairwise correlations, the Maximum Diversification optimization accounts for diverse correlations.

What does it mean to maximize diversification, and what is the Diversification Ratio?

Consider a simple case of two markets with similar volatility of 10% annualized. If they are perfectly correlated then an equally weighted portfolio of the two markets will have the same volatility as holding either market on its own - 10%. There is no benefit from diversification. If they are perfectly uncorrelated then an equally weighted portfolio will have a volatility of 7.1% because of diversification.

You can measure the amount of diversification in a portfolio by dividing the portfolio volatility if all of the investments were perfectly correlated by the portfolio volatility after accounting for diversification. In the case of our two asset portfolio the ratio would be 10% / 7.1% or 1.41. This is the portfolio's Diversification Ratio, and the Maximum Diversification optimization seeks the portfolio weights that maximize this value.

Let's go one step further with our two asset example.

Intuitively if the assets are perfectly correlated they are near-perfect substitutes for one another, so holding both assets in equal weight represents just one bet. Commensurately, if they are perfectly uncorrelated they would represent two independent bets.

Luckily, there is a quick way to determine the number of uncorrelated sources of return - independent bets - in a portfolio: you square the Diversification Ratio.

For proof, recall that the diversification ratio of our two asset portfolio where the assets were perfectly uncorrelated was 1.41. The square of 1.41 is 2.

The Maximum Diversification optimization is equivalent to maximizing the portfolio Sharpe ratio, but where expected returns are shrunk to equal their respective volatilities. The resulting portfolio can often be quite concentrated as it will only load on assets that represent truly uncorrelated directions of risk.

Minimum Variance

$$w = arg\min w^T \cdot \Sigma \cdot w$$

The Minimum Variance optimization is also equivalent to a maximum Sharpe ratio optimization, but where expected returns are all shrunk to the same value. Since expected returns are all the same, the sole objective of the optimization is to minimize portfolio volatility.

The optimization has two levers to pull to minimize portfolio volatility: 1) it can load up on low volatility securities and 2) it can seek assets with low correlation.

Notably, the minimum variance portfolio expresses similar expectations about the relationship between risk and return as the equal weight portfolio. That is that all returns are equal and independent of risk.

However, the minimum variance optimization takes advantage of cross-sectional differences in volatility and correlations to minimize portfolio volatility.

For completeness, in our analysis below we also examine the character of a relatively new optimization method: the Hierarchical Risk Parity portfolio [HRP].

Despite its name, the HRP optimization has the objective of minimizing portfolio volatility (i.e. NOT creating risk parity) by sequentially-weighting clusters of stocks in an inverse variance hierarchy. This is designed to avoid the matrix inversion step that has the potential to corrupt traditional convex optimizations. We would encourage readers interested in the method to dig deeper in Building Diversified Portfolios that Outperform Out-of-Sample by Marcos Lopez de Prado.

Low Volatility

We vacillated on whether to include a low volatility tilt portfolio in our analysis since it is not a risk-efficient approach, but we ultimately decided to include it for two reasons:

1. Readers who are less familiar with minimum variance portfolios may conflate minimum variance with low volatility tilt portfolios and this may be a source of confusion.
2. Readers who are familiar with the literature on low volatility factor portfolios may be curious about how minimum variance portfolios diverge in character from low volatility portfolios.

We followed the methodology used by S&P Dow Jones Indices. Portfolios were formed by ranking stocks at each rebalance on trailing 1-year daily volatility and holding the bottom 20 percent of stocks. Eligible stocks were then weighted by inverse volatility.

RISK EFFICIENT PORTFOLIO DECISION TREES

The Portfolio Optimization Machine paper includes decision trees to guide investors' choice of optimal portfolio method given the nature of the investment universe under consideration, and what they believe they can estimate. Figure 3 illustrates how an investor might choose to construct a portfolio when they believe all of the securities have the same expected return, but where they are not confident about their ability to estimate covariances. Unsurprisingly, the investor is guided towards the equal weight portfolio.

Figure 3: Portfolio optimization decision tree for stock portfolios when returns are independent of risk and we are naive about covariances.

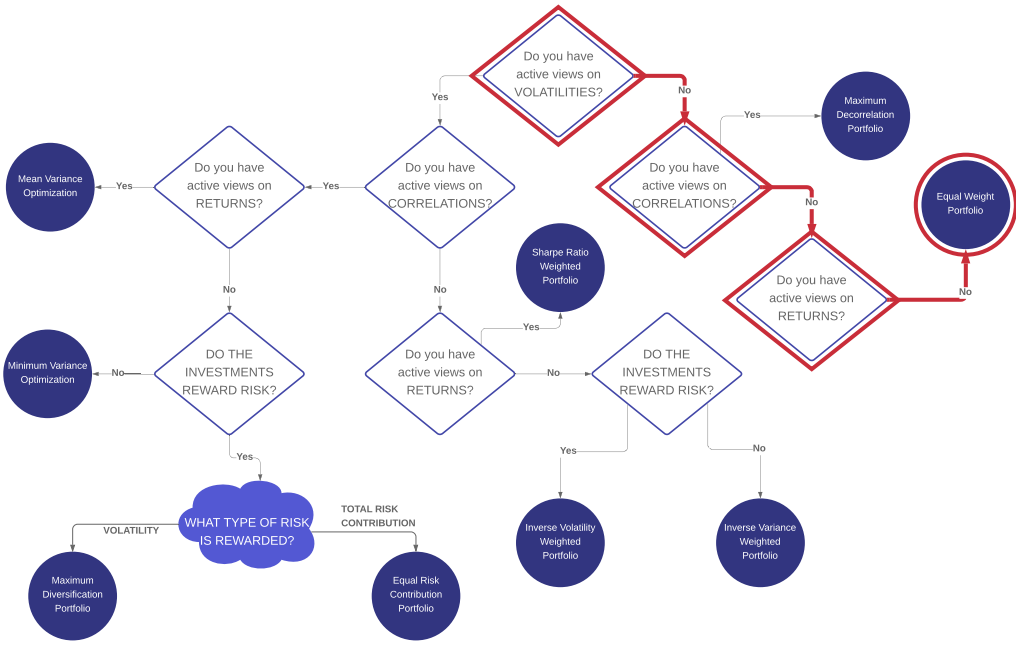

Source: ReSolve Asset Management. For illustrative purposes only.

In slight contrast, the example in Figure 4 below maps the decision path to the optimal portfolio method on the assumption that stock returns are independent of risk (both beta and volatility), and that we can reasonably estimate covariances. Following the red arrows to express our views on volatilities, correlations and returns we find that the minimum variance portfolio should be mean-variance optimal.

Figure 4: Portfolio optimization decision tree for stock portfolios when returns are independent of risk and we can estimate covariances.

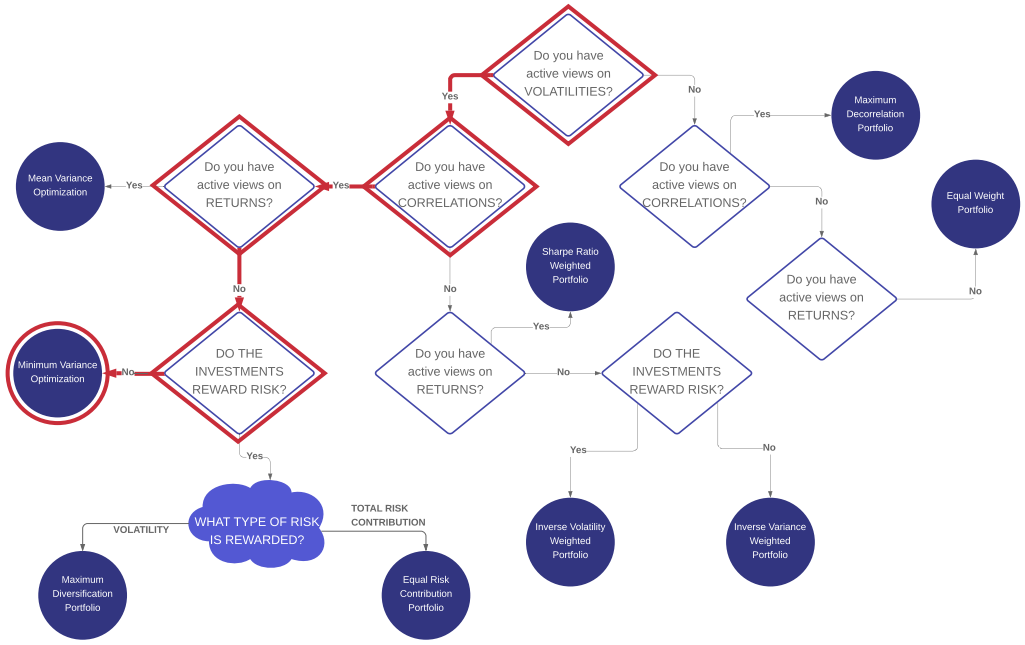

Source: ReSolve Asset Management. For illustrative purposes only.

DATA AND METHODOLOGY

We now present simulations to test our hypothesis about the most efficient way to form stock market portfolios.

Data was sourced from the Center for Securities Research and Prices (CRSP) daily stock tables. Specifically, we limited our analysis to large capitalization securities in the top 1000 stocks by market cap at the end of each year. Stocks were ranked by market cap as of December 31st of each year and the top 1000 stocks by market capitalization at the end of year t were considered eligible in year t+1.

At each monthly portfolio formation date we eliminated from consideration stocks with extreme historical variance above(below) the 0.99(0.01) percentile of all eligible stocks on that date. As such, our investment universe averaged 980 stocks at each rebalance date.

Portfolios were formed on the last trading day of the month and traded at the close five days later. We lagged trading by five days to ensure timely data could be effectively collected and processed in advance of trading, and to minimize the potential for look-ahead bias in any step. We traded twelve overlapping “tranches” so that 1/12th of the portfolio was rebalanced each

month and then held for the next 12 months. This step was taken to minimize expected turnover and commensurate slippage and transaction costs, and to minimize the impact of taxes on short-term gains.

Consistent with our objective to evaluate strategies that are coherent, investable alternatives for investors seeking broad stock market exposure, we have constrained portfolios to be long-only with weights that sum to one.

Covariances were estimated with an Exponentially Weighted Moving Average (EWMA) of the following form:

$$cov(X,Y) = \sum_{i=1}^{n} \alpha_i \cdot (x_i - \overline{x}) \cdot (y - \overline{y})$$

where:

$$\alpha_{i+1} = \lambda \cdot \alpha_i$$

and λ was set to approximate a center of mass of approximately 1 year. To ensure stable matrix inversion and in observation of known biases in sample covariance estimates we applied simple Ledoit-Wolf shrinkage[8] with a shrinkage parameter of 50 percent.

Portfolios formed by optimization often lead to highly concentrated portfolios. This may not represent an economic challenge given that our universe is constrained to large liquid stocks, but many investors would prefer to hold more securities for a variety of reasons. Investors and regulators also impose position level concentration constraints. To address this potential issue we also ran these optimizations with position limits of 1.5%, 2.5% and 5%.

RESULTS

Before we investigate the relative merits of more sophisticated portfolio construction methods we should set expectations about what we should observe. We base these expectations on the relationships between risk and return observed in the actual historical record.

Recall from Figure 1 that stock returns have historically exhibited no relationship between returns and beta, and other research finding no relationships between stock returns and volatility. If there has been a relationship, the slope has been inverted so that higher risk stocks have produced somewhat lower returns than lower risk stocks.

From Figure 3 and Figure 4 we can derive the optimal portfolio construction method on the condition that stock returns are independent of risk. Where we ignore forecasts of covariance, the equal weighted portfolio is optimal. If we incorporate covariance estimates the minimum variance portfolio is expected to be most efficient in terms of maximizing the expected Sharpe ratio.

Indeed, since 1965 the Equal Weight portfolio of large capitalization stocks has dominated the value-weighted portfolio in terms of both gross excess returns and risk-adjusted returns.

8 See Honey I Shrunk the Sample Covariance Matrix

Figure 5: Market cap versus equal weighted large cap stocks (top 30% by market cap) 1957-2018

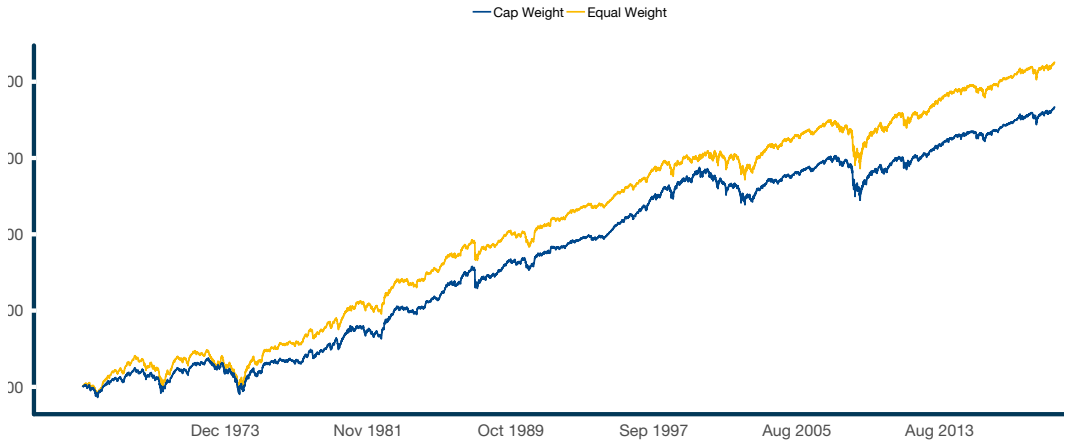

Source: ReSolve Asset Management. Data from Ken French top 30% by ME. For illustrative purposes only. Past performance is not indicative of future results.

When we expand our view to include optimal methods that account for covariance estimates and explicitly emphasize diversification, it's clear that optimal portfolios are competitive with naive portfolios in terms of compound returns without the use of leverage.

From Table 2 we observe that optimal and naive strategies produced comparable performance, but optimal strategies produced their returns with significantly less risk. Investors in minimum variance type portfolios experienced less than half the volatility and just over two-thirds the maximum drawdown of the value- weighted portfolio without sacrificing returns. In addition, minimum variance portfolios were almost 50% more efficient than the equal weight portfolio as measured by Sharpe ratios.

By convention most academic papers assume investors can borrow and lend at the same rate as the U.S. government. But most investors can't borrow at the risk-free government rate. Rather, funds and institutions must borrow at a rate that accounts for overnight credit risk. For this reason, all risk-adjusted ratios are reported in excess of LIBOR to account for true estimated borrowing costs available to large institutions and funds.

Table 2: Performance statistics

	Cap Weighted	Equal Weight	Low Vol	Inverse Vol	ERC	HRP	MAX-DIV	MAX-DIV Max 1.5%	MAX-DIV Max 2.5%	MAX-DIV Max 5.0%	MIN-VAR	MIN-VAR Max 1.5%	MIN-VAR Max 2.5%	MIN-VAR Max 5.0%
Annualized Return	9.87%	11.54%	12.00%	11.97%	11.52%	11.80%	9.10%	9.49%	9.29%	9.14%	9.46%	10.38%	9.94%	9.58%
Annualized Volatility	15.70%	15.70%	10.60%	14.10%	12.60%	11.60%	8.60%	9.20%	8.80%	8.60%	6.80%	7.60%	7.20%	6.90%
Sharpe Ratio	0.41	0.50	0.72	0.57	0.59	0.66	0.56	0.57	0.57	0.57	0.74	0.78	0.78	0.75
Max Drawdown	-55.50%	-58.80%	-46.30%	-55.70%	-52.00%	-48.20%	-40.90%	-43.10%	-41.60%	-40.90%	-38.20%	-39.60%	-39.30%	-38.30%
Avg Drawdown	-5.40%	-4.60%	-2.90%	-3.90%	-3.60%	-3.20%	-3.60%	-3.50%	-3.50%	-3.60%	-2.30%	-2.30%	-2.20%	-2.20%
Beta to Cap Weighted	1.00	0.97	0.59	0.87	0.76	0.70	0.47	0.51	0.48	0.47	0.36	0.41	0.38	0.36
Annual Alpha to Cap Weighted	0.00%	1.70%	3.90%	2.50%	2.60%	3.10%	1.80%	2.00%	2.00%	1.90%	2.70%	3.30%	3.10%	2.80%
Return to Ulcer Ratio	0.70	0.94	1.36	1.10	1.18	1.27	1.01	1.05	1.04	1.03	1.37	1.46	1.42	1.39

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

The Ulcer Index captures both the depth and the duration of drawdowns and is a good measure of how much anxiety a hypothetical investor would have endured from investing in a strategy. It is the geometric average of all instances where an investor's wealth dropped below an all-time high, and accounts for both the depth and duration of wealth impairments.

$$R_i = \frac{P - MAX(P)}{MAX(P)}$$

and

$$Ulcer = \sqrt{\frac{R_1^2 + R_2^2 + \cdots + R_N^2}{N}}$$

The Cap Weighted strategy produced 0.7 units of return per "Ulcer" unit while low volatility and minimum variance portfolios approached 1.37 units of return per "Ulcer" unit, suggesting optimal portfolios were about 2× more attractive in terms of "gains" per unit of "pain".

Of note, different portfolio methods expressed exposures to a surprisingly diverse set of risks. Strategy correlations ranged from 0.81 to 1 with Cap Weighted and MIN-VAR strategies exhibiting the lowest pairwise correlations. Portfolios with concentration constraints exhibited higher correlations because the constraints act to shrink portfolios toward the equal weight portfolio.

We were interested to see that MIN-VAR and Low VOL portfolios had a correlation of about 0.92 - 0.94, which was in the middle of the range across the various methods. More constrained MIN-VAR methods exhibited greater exposure to lower volatility stocks since they were constrained in their ability to seek diversification. These are our first clues that while MIN-VAR and Low VOL strategies have produced similar performance in the Russell 1000 sample, the performance has been derived from different sources. We will explore these differences in greater detail below and in future articles.

Figure 6: Correlations between levered strategies

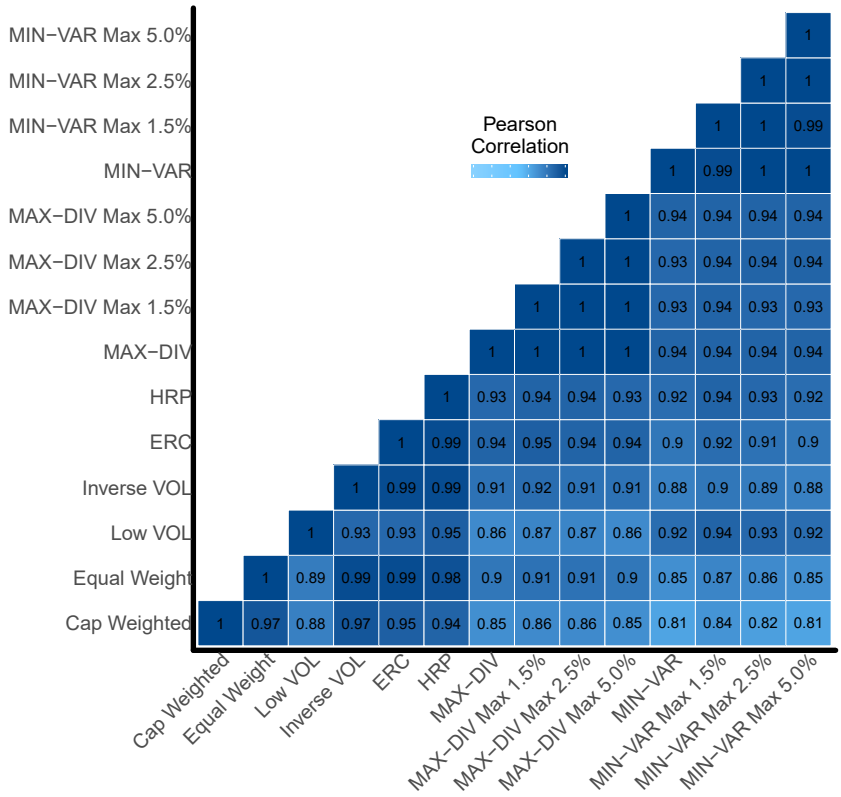

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

LEVERAGING EFFICIENT PORTFOLIOS

An important benefit of more efficient portfolios is the ability to scale portfolio exposure using leverage to achieve higher returns with market-like risk.

In the 1960s Bill Sharpe described why a rational investor would prefer to own the most efficient portfolio and then borrow or lend (at the "risk-free" rate) against this portfolio to achieve return or risk objectives. An investor who can tolerate greater risk and requires a higher return will borrow funds to invest greater than 100 percent of their wealth in the efficient portfolio. Conversely, an investor with less tolerance for risk or lower return objectives would only invest a portion of their wealth in the efficient portfolio and hold the balance in risk-free instruments. This is the concept of the Capital Market Line.

Consistent with this premise, we also simulated the performance of each investment strategy when scaled to 130 percent exposure. Leverage of this magnitude is readily available in 40-Act mutual fund structures and to most institutional investors. Again, our objective is to analyze and evaluate investable strategies with assumptions that reflect real-world constraints. For this reason, we have reported all excess returns and performance statistics net of LIBOR rather than the more traditional T-bills. In addition, we assume investors would pay LIBOR rates on leverage so the results we report should have been generally achievable by funds and institutions.

Table 3: Performance statistics for levered portfolios

	Cap Weighted	Equal Weight	Low Vol	Inverse Vol	ERC	HRP	MAX-DIV	MAX-DIV Max 1.5%	MAX-DIV Max 2.5%	MAX-DIV Max 5.0%	MIN-VAR	MIN-VAR Max 1.5%	MIN-VAR Max 2.5%	MIN-VAR Max 5.0%
Annualized Return	9.86%	11.50%	14.16%	13.89%	13.40%	13.82%	10.39%	10.88%	10.62%	10.43%	10.93%	12.12%	11.55%	11.08%
Annualized Volatility	15.70%	15.70%	13.80%	18.40%	16.30%	15.00%	11.20%	11.90%	11.50%	11.20%	8.90%	9.90%	9.30%	8.90%
Sharpe Ratio	0.41	0.50	0.72	0.57	0.59	0.66	0.56	0.57	0.57	0.57	0.74	0.78	0.78	0.75
Max Drawdown	-55.50%	-58.80%	-57.20%	-67.10%	-63.20%	-59.20%	-50.60%	-53.10%	-51.40%	-50.70%	-47.70%	-49.20%	-48.80%	-47.70%
Avg Drawdown	-5.40%	-4.60%	-4.10%	-5.50%	-5.00%	-4.50%	-5.20%	-5.20%	-5.20%	-5.20%	-3.40%	-3.40%	-3.40%	-3.40%
Beta to Cap Weighted	1.00	0.97	0.77	1.13	0.99	0.90	0.61	0.66	0.63	0.61	0.46	0.53	0.49	0.47
Annual Alpha to Cap Weighted	0.00%	1.70%	5.00%	3.20%	3.30%	4.10%	2.40%	2.60%	2.50%	2.40%	3.60%	4.30%	4.00%	3.70%
Return to Ulcer Ratio	0.70	0.93	1.19	0.96	1.02	1.11	0.86	0.90	0.88	0.87	1.15	1.26	1.20	1.18

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

Table 3 shows that optimal portfolios can achieve impressive returns with very manageable leverage. Low VOL and MIN-VAR style portfolios produced 2-4 percentage points of annual excess returns and 4-5 percentage points of annual alpha when regressed on Equal Weight returns, with similar maximum drawdowns and volatility.

BOOTSTRAP TESTS

Table 2 and Table 3 show that optimal methods outperformed the naive equal weight portfolio in the historical sample. But the historical sample represents just one instantiation from a distribution of possible return trajectories. Investors are faced with the practical question of which method is most likely to outperform going forward. Specifically, investors are concerned with the probability and magnitude of positive and negative outcomes.

While there are parametric tests available to construct confidence regions on returns, Sharpe ratios and other performance measures, these approaches fall short when time series of returns are not independent and identically distributed ("iid"). There is a large volume of financial literature asserting that the time series of stock returns are not independent. In particular, researchers have made the case for clustering of excess volatility, short-term and long-term mean-reversion in prices, and intermediate-term momentum. If stock returns are dependent as this research suggests, standard parametric tests will yield biased statistics.

Bootstrapping is a form of Monte Carlo analysis where monthly returns are drawn randomly from the distribution of actual realized returns. Each time a return is drawn it is placed back in the distribution so that it can be drawn again. The advantage of bootstrap over parametric tests is that it preserves the empirical shape of the return distribution rather than assuming returns conform to traditional shapes like Gaussian or Poisson distributions.

Traditional bootstraps draw returns randomly so that the next draw is completely independent of the previous draw. This approach is appropriate when returns are iid since it eliminates any relationships that might exist between return observations,

such as autocorrelation effects. The stationary block bootstrap test[9] is an alternative testing procedure that produces valid standard errors and confidence regions in the presence of either iid or weakly dependent observations. Random blocks of observations are chosen where each block contains a random number of rows chosen from a geometric distribution. The average number of observations per block is chosen in order to minimize the mean squared error between the bootstrap variance and the variance of the original sample[10].

We performed a block bootstrap test with 10 thousand samples. For each sample we chose random blocks of rows (with replacement) from the original data so that the total number of rows in each sample was equal to the number of rows in the original data. At each sample draw we calculated performance metrics including Sharpe ratios, annualized returns and maximum drawdowns, as well as excess returns and maximum drawdowns relative to the equal weight strategy.

For each performance metric we calculated the percentage of samples where a strategy produced more attractive performance than all the other strategies. We performed pairwise comparisons across 14 strategies, which produced 171 unique comparisons. Figure 7 through Figure 9 describe the results of these comparisons for Sharpe ratios, annualized returns, and ratios of returns to maximum drawdowns (MAR ratios).

Each cell in a matrix describes the percentage of bootstrap samples where the strategy corresponding to the row on the y-axis produced better results than the strategy at the intersecting column on the x-axis. For example in Figure 7 the cell in the bottom left corner quantifies the percentage of bootstrap samples where the minimum variance portfolio constrained to 1.5% maximum position size (MIN-VAR Max 1.5%) produced a higher Sharpe ratio than the cap weighted portfolio. In this case the minimum variance portfolio produced a higher Sharpe ratio than the equal weight portfolio in 100% of samples.

The matrices are color coded so that results in excess of 80% are highlighted and darker colors correspond to higher probabilities.

We used unlevered returns to compare Sharpe ratios and MAR ratios, but levered returns to compare annualized returns.

9 See D. N. Politis and J. P. Romano (1994): The Stationary Bootstrap. Journal of the American Statistical Association 89, 1303–1313.

10 See [ftp://ftp.stat.math.ethz.ch/Research-Reports/52a.ps.gz] for details on on this approach.

Figure 7: Bootstrap probability that y-axis optimizations have higher Sharpe ratios than x-axis optimizations.

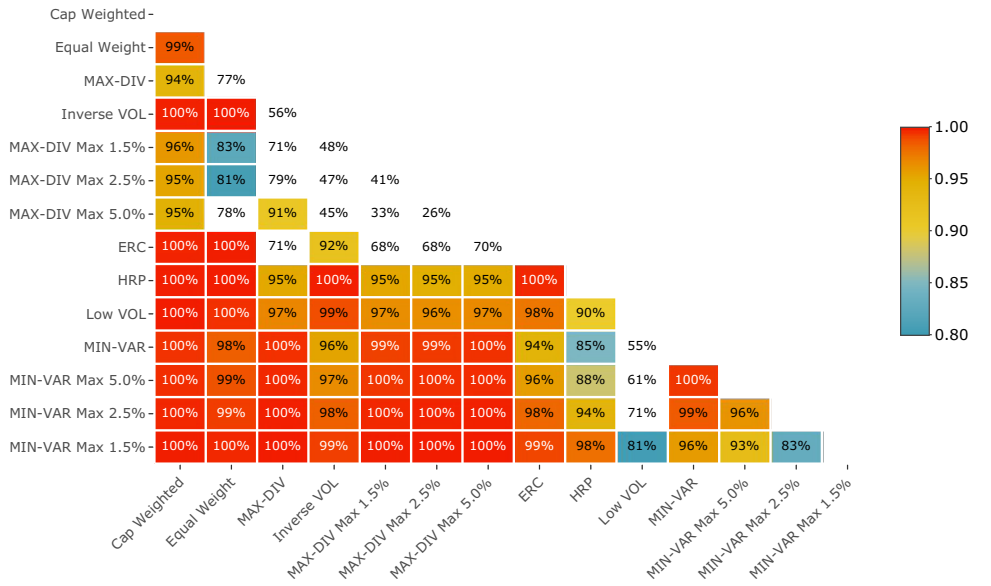

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

The majority of highlighted cells in the Cap Weighted column reflects the fact that most of the strategies produced a higher Sharpe ratio than the Cap Weighted Weight strategy at least 80% of the time. We observe similar outcomes for MAX-DIV, Inverse VOL, ERC and HRP strategies, where they are universally dominated by Low VOL and MIN-VAR methods. We see no clear pattern of dominance between Low VOL or MIN-VAR; they are both equally attractive in our sample. However, this does not mean the two approaches are the same as we shall discover below.

Figure 8: Bootstrap probability that y-axis optimizations have higher annualized returns than x-axis optimizations.

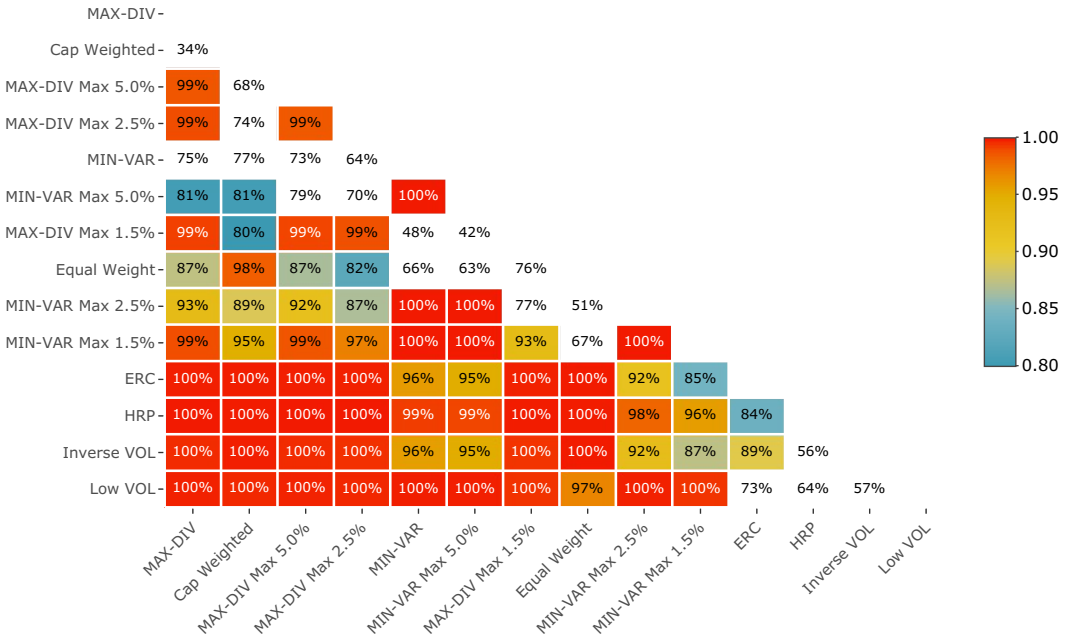

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

Figure 9: Bootstrap probability that y-axis optimizations have higher annualized return / maximum drawdown ratios than x-axis optimizations.

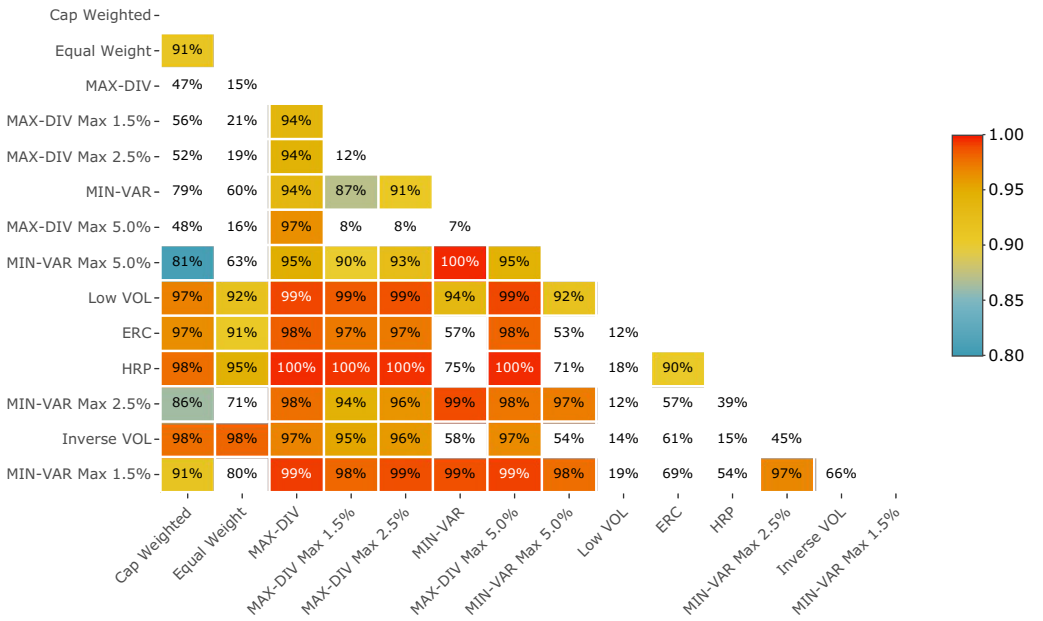

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

PORTFOLIO CONCENTRATION

By design, some optimization methods will create more concentrated portfolios than others.

For example, equal-weight, inverse volatility and equal risk contrubution portfolios will hold positive weights in all stocks at all times by definition. While some stocks will earn greater weight than others in the latter optimizations, security weights will usually be relatively evenly distributed.

On the other hand, mean-variance type optimizations like minimum variance and maximum diversification often produce portfolios that hold a relatively small number of positions relative to the number of available securities. This happens because the optimizers will only continue to add securities to the extent that new securities increase the portfolio Sharpe ratio (or reduce variance) at the margin[11].

To measure concentration we invoke the Inverse Herfindahl Index (HHI), which effectively measures the Euclidean distance between the portfolio weights and the weights of an equal-weighted portfolio. This makes sense since the equal weighted portfolio is least concentrated by definition. The inverse HHI is calculated as the square of the sum of the weights divided by the sum of the squared weights:

$$Inverse.Herfindahl = \frac{sum(w)^2}{sum(w^2)}$$

11 Maximum Diversification has the objective of maximizing expected Sharpe ratio, but where an asset's volatility is a proxy for the expected return. So while the optimization technically will continue to add assets until adding new positions no longer increases the Diversification Ratio, the Diversification Ratio is actually a Sharpe ratio under the stated assumptions.

Figure 10: Rolling implied number of holdings (inverse Herfindahl index) since inception of Russell 1000

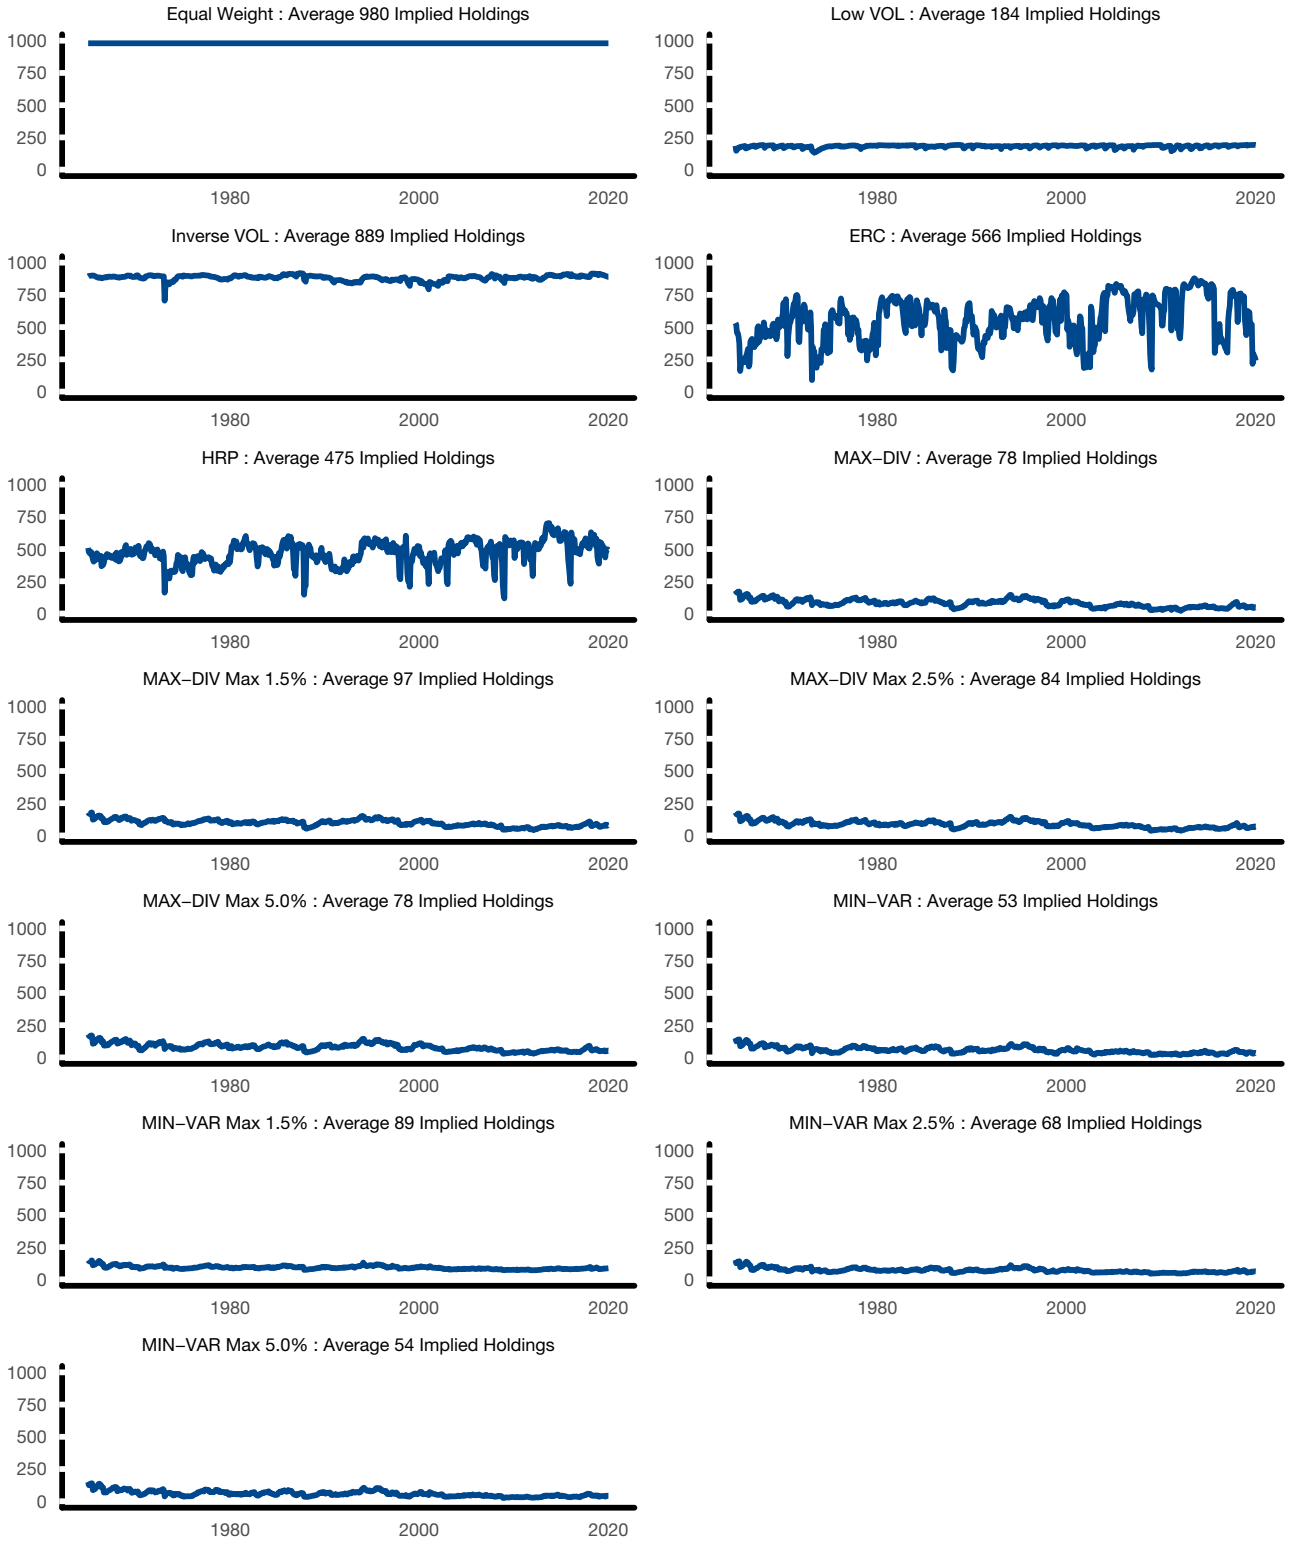

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

Importantly, a portfolio can appear to be highly concentrated but exhibit substantial diversification, and vice versa. In fact, more diversified portfolios are often highly concentrated, as can be seen in Figure 10 and Figure 11. The MIN-VAR portfolio is a good example. Note that the portfolio averages just 56.2 equivalent equally weighted holdings as measured by the inverse HHI, which many investors would perceive as highly concentrated.

However, with an average of 29.8 independent bets, the MIN-VAR portfolio was second only to MAX-DIV portfolios in terms of diversification. This is due to the fact that the MIN-VAR optimization seeks to lower portfolio volatility by aggregating lowly correlated stocks. In contrast, the Equal Weight portfolio, which is the least concentrated by definition, averaged just 11.9 independent directions of risk.

Figure 11: Rolling implied number of independent bets (Diversification Ratio2)

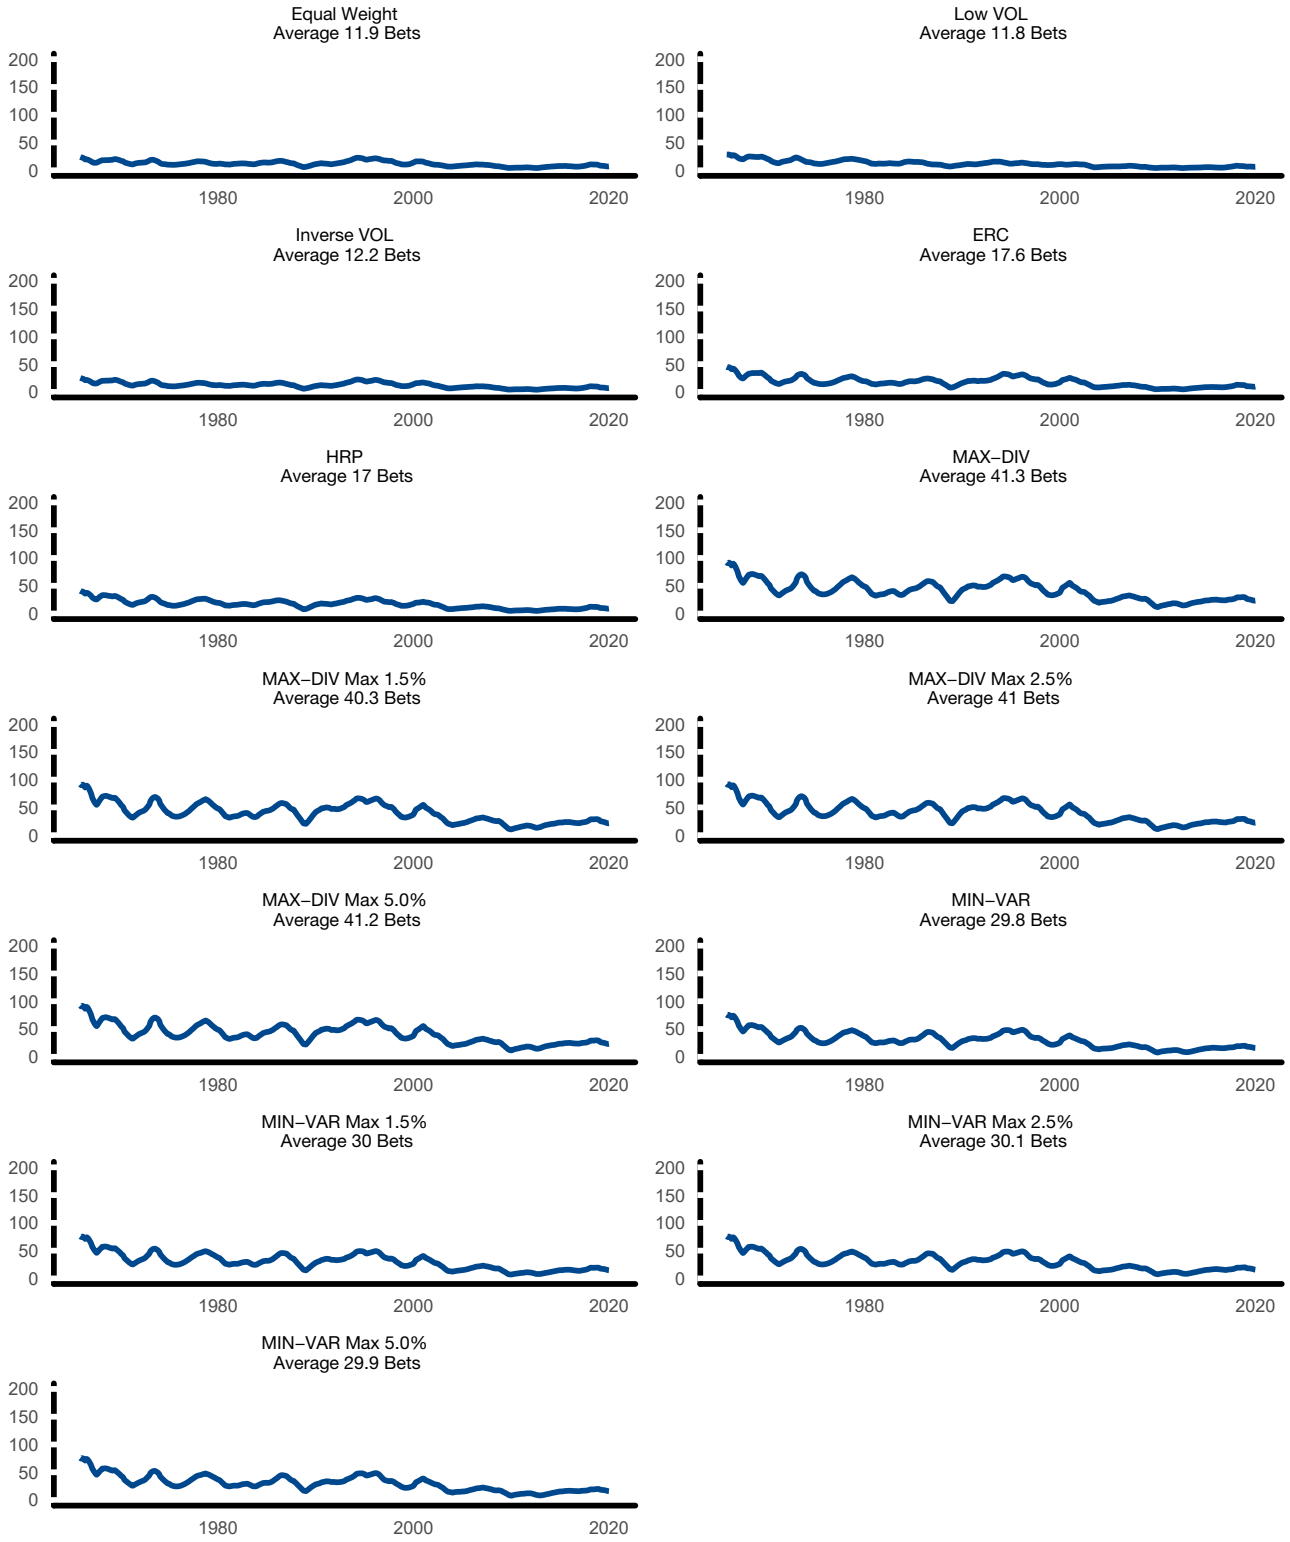

Source: Data from Center for Research in Securities Prices (CRSP). Analysis by ReSolve Asset Management. SIMULATED RESULTS. Past performance is not indicative of future performance.

CONCLUSION

We reported on the death of the Capital Asset Pricing Model and asked an all-important question - If CAPM doesn't hold and the market cap weighted portfolio is inefficient, what is the optimal risk-efficient stock market portfolio?
We proposed that an optimal market portfolio should exhibit the following three qualities:

1. Ignore active views on expected stock returns for individual stocks
2. Expresses a consistent relationship (or no relationship) between risk and return
3. Explicitly seek to diversify across the independent sources of risk in the market

Motivated by the objectives of practitioners we set out to analyze the performance of optimal versus naïve portfolios on a universe of large, liquid U.S. stocks.

Consistent with the observation that stock returns are independent of risk, we surmised that the minimum variance portfolio would produce the most efficient returns.

In simulation, sensibly constrained minimum variance portfolios delivered returns that were competitive with the equal weight portfolio, but with substantially less risk. These results were confirmed in bootstrap tests, where the Sharpe ratio of minimum variance portfolios dominated the Sharpe ratio of naïve portfolios in over 95% of samples.

We also scrutinized the concentration and diversification characteristics of each portfolio. Portfolios formed via mean-variance type optimization – MIN-VAR and MAX-DIV type portfolios – exhibited high levels of concentration. However, when viewed through the lens of diversification these portfolios derived returns from a much larger number of independent sources of risk, suggesting they do a better job of spanning the economic drivers of market returns.

In summary, optimal portfolios have produced economically and statistically significant outperformance in the modern period relative to naïve portfolios. This outcome is robust to practical trading considerations and can be deployed at scale by large institutions.

Disclaimer:

Confidential and proprietary information. The contents hereof may not be reproduced or disseminated without the express written permission of ReSolve Asset Management Inc. ("ReSolve").

ReSolve is registered as an Investment Fund Manager in Ontario, Quebec and Newfoundland and Labrador, and as a Portfolio Manager and Exempt Market Dealer in Ontario, Alberta, British Columbia and Newfoundland and Labrador. ReSolve is also registered as a Commodity Trading Manager in Ontario and Derivatives Portfolio Manager in Quebec. Additionally, ReSolve is a non-resident Registered Investment Adviser with the SEC. ReSolve is also registered with the Commodity Futures Trading Commission as a Commodity Trading Advisor and a Commodity Pool Operator. This registration is administered through the National Futures Association ("NFA"). Certain of ReSolve's employees are registered with the NFA as Principals and/or Associated Persons of ReSolve if necessary or appropriate to perform their responsibilities. ReSolve has claimed an exemption under CFTC Rule 4.7 which exempts ReSolve from certain part 4 requirements with respect to offerings to qualified eligible persons in the U.S.

PAST PERFORMANCE IS NOT NECESSARILY INDICATIVE OF FUTURE RESULTS.

General information regarding hypothetical performance and simulated results. Past results are not necessarily indicative of future results. It is expected that the simulated performance will periodically change as a function of both refinements to our simulation methodology and the underlying market data. These results are based on simulated or hypothetical performance results that have certain inherent limitations. Unlike the results in an actual performance record, these results do not represent actual trading. Also, because these trades have not actually been executed, these results may have under- or over-compensated for the impact, if any, of certain market factors, such as lack of liquidity. Simulated or hypothetical trading programs in general are also subject to the fact that they are designed with the benefit of hindsight. No representation is being made that any account or fund will or is likely to achieve profits or losses similar to those being shown.

These materials do not purport to be exhaustive and although the particulars contained herein were obtained from sources ReSolve believes are reliable, ReSolve does not guarantee their accuracy or completeness. The contents hereof does not constitute an offer to sell or a solicitation of interest to purchase any securities or investment advisory services in any jurisdiction in which such offer or solicitation is not authorized.

Forward-Looking Information. The contents hereof may contain "forward-looking information" within the meaning of the Securities Act (Ontario) and equivalent legislation in other provinces and territories. Because such forward-looking information involves risks and uncertainties, actual performance results may differ materially from any expectations, projections or predictions made or implicated in such forward-looking information. Prospective investors are therefore cautioned not to place undue reliance on such forward-looking statements. In addition, in considering any prior performance information contained herein, prospective investors should bear in mind that past results are not necessarily indicative of future results, and there can be no assurance that results comparable to those discussed herein will be achieved. The contents hereof speaks as of the date hereof and neither ReSolve nor any affiliate or representative thereof assumes any obligation to provide subsequent revisions or updates to any historical or forward-looking information contained herein to reflect the occurrence of events and/or changes in circumstances after the date hereof.