Alternative Index Weighting and the Impact on Portfolio Risk

Mean-variance efficiency has been sought since Markowitz's seminal work on portfolio optimization in 1952, but generally speaking, attaining such efficiency critically depends on assumptions involved in modeling.[1] With the growing popularity of exchange-traded products that provide packaged index solutions, it is important to assess the benefits and pitfalls of alternative weighting mechanisms that seek efficiency, many of which purport to be the right solution for the sophisticated investor.

The recent move in the investment community towards other approaches to portfolio management includes some that may sidestep return forecasting altogether and simply focus on volatility. The "mean" in "mean-variance optimization" (MVO) is completely ignored. In many cases, investor interest in these approaches is driven by a desire to capture the "low beta anomaly," which proposes that low beta stocks have historically exhibited positive alpha.[2] Among others, these approaches include:

1. Minimum Variance portfolios (MV)
2. Low Volatility portfolios (LV)
3. Equal Weight portfolios (EW)
4. Equal Risk Weight portfolios (EQR)

We will include Market Cap Weighting for some of our analysis.

In our analysis, we will focus exclusively on the risk characteristics of the different approaches. We propose a simple problem: an investor has a fixed amount of capital to be invested in a fixed number of stocks. The entire amount of capital is invested, and there is no short-selling. In other words, each stock has a weight in the portfolio between 0% and 100%, and the sum total of the weights in the portfolio equals 100%. How the investor invests the capital depends on the method of choice, and each has a different utility optimization function, or none at all.

Our research broadly shows that (1) the assumptions required to make several of these solutions "mean-variance efficient" are questionable, (2) the range of embedded risk among popular approaches may vary far more widely than investors may be aware of, and (3) investors may not be getting out of some exchange-traded products what they promise, e.g. low volatility portfolios may not, in fact, be low volatility at all.

[1] Markowitz, H. (1952). "Portfolio Selection." *Journal of Finance*. Vol. 7, No. 1, pp. 77-91.
[2] Frazzini and Pedersen (2010) argue that this anomaly stems from investors who cannot use leverage seeking to achieve it by selecting high beta stocks, causing them to be overbought.

In particular, we establish a risk ordering of five different index weighting approaches ranging from Minimum Variance (MV) on the low end to Equal Weight (EW) on the high end:

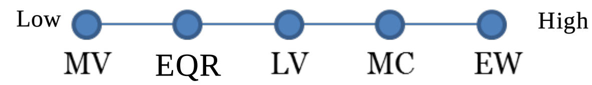

MV is Minimum Variance, EQR is Equal Risk Weighting, LV is Low Volatility, MC is Market Cap Weighting and EW is Equal Weighting.

Furthermore, we establish that the risk contributions of component stocks (and therefore concentration risk) vary quite dramatically across different approaches, with Equal Risk Weight (EQR) showing the least variation across component stocks to Minimum Variance (MV) exhibiting the highest:

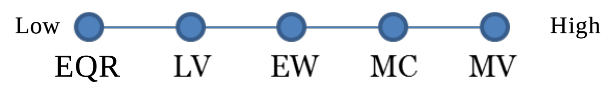

To set the stage for what follows, it is constructive to briefly review the mean-variance framework as well as each of the approaches examined in this paper.

Mean-Variance Efficiency

The thinking behind mean-variance optimization is straightforward: Investors seek higher returns and lower volatility of those returns. They try to maximize returns and minimize risk. Consider a scenario where each unit of expected return gave the investor one unit of "utility" and each unit of risk, measured in variance detracted from the investor's utility. For a given amount of utility, the choices of portfolios available to the investor are infinite, but only one portfolio would provide the investor with the greatest utility for a given level of risk or a given expected return. This set of portfolios, collectively, is known as the efficient frontier, and all portfolios on the frontier are considered "mean-variance efficient" (or simply "efficient"). In other words, they each provide the highest possible expected return relative to risk.

Efficient portfolios exist in theory, but investors must attempt to approximate them, and they are the result of some algorithm. Put expected returns and risk characteristics of a set of stocks (volatilities, correlations) into a model, run the model, and out comes the portfolio weights. In order for a particular solution to be actually mean-variance efficient in real life, the data that goes into the model and/or the model itself must have assumptions that turn out to be true. If the model requires correlations, the correlation estimate has to be accurate. If it requires returns, you must forecast actual returns.

Note that when we say "efficient" we are still describing a rather stylized world, even before we consider constraints. In particular, there is no measurement error, no time-varying parameters, etc., and this is not to say that efficient portfolio will outperform other schemes, nor does it guarantee that *ex post* the portfolio will have been mean-variance efficient. This is an important point, for even when one portfolio has a lower *ex ante* volatility, for example, there is no guarantee that this portfolio will not have higher volatility than that of a higher expected volatility portfolio.

The models we present all seek to either simplify the optimization (or even eliminate it) and/or reduce or eliminate estimation error in the parameters for the model. Expected (excess) returns are very difficult to estimate, as are covariance matrices. Methods that eliminate some of the estimation may trade off some accuracy in the parameter estimates for accuracy in the optimization.

Below we examine five models and specify the assumptions that would have to hold true for each model to generate portfolios that were actually efficient.

Market Capitalization (MC)

Description

The Market capitalization (MC) weighting scheme, as the name suggests, weights all of the assets in the portfolio according to size, as measured by the outstanding market capitalization of each asset. No optimization is required to construct the MC portfolio. If, among other assumptions, all investors were mean variance motivated, had the same information set and all stock prices fully reflected that information, then each investor would hold the individual portfolio that maximized their utility, and the aggregate market would therefore be mean variance efficient.

Constraints Required for Efficiency

Several constraints are required, as outlined in Haugen and Baker (1991),[3] including that "all investors agree about the risk and expected return for all securities, all investors can short-sell all securities, no investor's return is exposed to federal or state income tax liability, [and] the investment set for all investors holding any security in the index is restricted to the securities in the capitalization-weighted index." In addition all investors must be able to both borrow and lend at the risk free rate, must be rational, and must have no transaction costs.

Comments

Given the practical reality that many of the assumptions are violated, it is likely that stocks are mispriced, and therefore MC schemes tend to overweight overvalued stocks that have had outsized historical returns, while underweighting undervalued stocks. In addition, it has been suggested that because many investors do not have access to leverage, they access it through choosing high beta stocks, therefore causing these stocks to be overvalued on average.

Minimum Variance (MV)

Description

For the Minimum Variance (MV) portfolio, suppose the investor is extremely risk averse and is only interested in the portfolio that, *ex ante*, provides the lowest risk, as measured by portfolio standard deviation or variance. There are several important features of this portfolio optimization process, which is a constrained version of the MVO to note. The investor can either (a) set the risk parameter infinitely high, in which case the expected excess returns no longer enter the equation, or (b) set each of the expected returns equal to one another. By design, the marginal contribution to risk of each asset is the same, since that is effectively the optimization condition for the minimum variance portfolio. This means that incrementally changing the weight of any stock in the portfolio would cause the same increase to the portfolios expected risk.

Constraints Required for Efficiency

In order for this approach to actually generate an efficient portfolio, the expected returns of all assets in the portfolio must be the same. In any other case, there is a portfolio with the same expected variance, but a higher expected return.

Comments

MV portfolios tend to be concentrated, not surprisingly, in low volatility stocks. In the extreme case where no two stocks are correlated, minimum variance weights will be inversely proportional to variance.[4,5] This may capture the "low volatility" anomaly (that lower volatility stocks have often outperformed), but unless this anomaly is absolute, there is a more efficient portfolio.

Equal Weight

Description

Equal (dollar) weighting schemes don't require optimization. If an investor has $X and N stocks under consideration, each stock gets a $X/N position in the portfolio. Believe it or not, this is a constrained version of the basic mean-variance optimization.[6]

Constraints Required for Efficiency

In a stylized world where the expected return of each asset is proportional to the sum of its covariances to all assets in the portfolio, the result of a mean-variance optimization will yield equal weights on each of the assets. There isn't much of an

[3] Haugen, R. and Baker, N. (1991). "The efficient market inefficiency of capitalization-weighted stock portfolios." *Journal of Portfolio Management.*

[4] http://www.edhec-risk.com/latest_news/featured_analysis/RISKArticle.2011-11-15.4803?newsletter=yes.

[5] Note that this is different than the Low Volatility model which uses the inverse of volatility, not variance.

[6] The sum of the log weights is constrained to be n•ln(n).

intuitive explanation for this assumption; it is simply a mathematical relationship which would equate these two models.

Comments

The benefit of simplicity of this approach is offset by its resulting high volatility relative to other approaches, and it is extremely unlikely to be efficient.

Low Volatility (LV)

Description

The low volatility (LV) model seeks to overweight stocks with low volatility and underweight stocks with high volatility. The model construction is as follows: all of the stocks under consideration are ordered from the highest volatility to the lowest volatility, and then the weights are distributed such that the weights are proportional to the reciprocal of volatility (i.e. 1/volatility) for each stock.

For example, a stock with 15 vol would have double the weighting of a stock with a vol of 30. The volatility of each individual stock position is therefore equal, but the model makes no statement about portfolio volatility since correlations, and by extension, covariances are not inputs into the model.

Constraints Required for Efficiency

In order for the low volatility solution to be mean-variance efficient, the dataset requires expected returns that are proportional to volatility and zero correlations across all assets in the portfolio. If there is any information about asset correlations, this model would not be efficient.

Comments

It should be clear from the description of the model that there is actually no optimization of any kind vis-à-vis portfolio volatility. In this sense, LV shares a feature of the Equal Weight model: fewer parameters to estimate. Whereas a full covariance matrix requires N*(N-1)/2 parameters, LV models require only N parameters. Offsetting this, however, is lack of optimization in the LV framework. There is no formal utility function relating to portfolio volatility or portfolio return and the LV solution will almost certainly lie off the efficient frontier.

Equal Risk Weight

Description

Equal Risk Weight portfolios (often called Risk Parity) match the total contribution that each asset makes to the overall risk of the portfolio. Consider the addition of two stocks to a given portfolio: Stock A is entirely uncorrelated to any asset already in the portfolio, while Stock B is highly correlated to the holdings in the portfolio. All else equal (the size of the position, the volatility of the stock), the addition of Stock B will raise the portfolio risk more than the addition of Stock A. Each *contributes* differently to the overall risk of the portfolio. In the Equal Risk Weight (EQR) model, the investor runs an optimization based on the covariance matrix of all stocks under consideration such that the total contribution to overall risk that each position makes is the same. This is in contrast to a minimum variance portfolio, where the *marginal* risk contribution is equal. The solution may yield weights whereby a low volatility stock that is highly correlated to the portfolio may contribute the same as a high volatility stock with a low correlation to the portfolio.

Constraints Required for Efficiency

Since the EQR model yields the LV model if all correlations are constant, the conditions for LV efficiency are sufficient for EQR efficiency, but not necessary.

Comments

In the case where all of the correlations are identical, the solution to the EQR model mimics the solution to the LV model. In other words, the LV solution explicitly ignores correlation, or more specifically, ignores any variation in correlation: the EQR solution will be equivalent to the Low Volatility solution if and only if all correlations are assumed to be equal. In addition it has been shown that the EQR portfolio will always lie between the Minimum Variance portfolio and the Equal weight portfolio.[7]

Analysis

Much has been written on each of the models we have described above. Our interest is in determining the impact on portfolio risk of each of the models in our model set. In particular, we are interested in answering two important questions:

1. What is the risk ordering, in terms of *ex ante* portfolio volatility for the models in question?
2. Which model yields the largest variation in risk contributions (idiosyncratic risk) and which yields the smallest?

Methodology

To determine the answers to these questions, we perform the following analysis:

1. We take all of the stocks that make up the S&P 500 Index (as of 07/02/2013).
2. We randomly draw N stocks from the list for which we have at least one year of daily returns. N is the portfolio size, and we perform the analysis for N=5, 10, 15, 20, 25, 30, 50, 75, 100.
3. We calculate the historical covariance matrix for the selected stocks.
4. We also capture the market capitalization of each of selected stocks (for the MC portfolio).

[7] Maillard, S., Roncalli, T. and Teiletche, J. "On the Properties of Equally-Weighted Risk Contributions Portfolios." (2010) *Journal of Portfolio Management.*

5. We calculate the portfolio weights for the five models under investigation: EQR, EW, LV, MC, MV. Only the EQR and the MV models require optimization.
6. For a given portfolio size (N), we repeat this procedure 500 times.

Results: Portfolio Volatility

For each model, we have 500 iterations for each portfolio size. In Figure 1, we plot the average (expected) portfolio volatility for each N for each of the models for a given portfolio size, assuming the covariance matrix forms our expectations of forward-looking volatility. This tells for a given size of portfolio, randomly selected from the S&P500, which approach will have the lowest portfolio volatility.

We observe, not surprisingly, that with each of the models under review the average portfolio volatility is decreasing with the size of the portfolio (i.e. the number of stocks in the portfolio). As expected, the two models that don't incorporate any volatility information, EW and MC, have the highest average portfolio volatility. The minimum variance (MV) portfolios tend to have the sharpest descent relative to portfolio size, while the LV portfolio tends to have the slightest decline, as measured by the difference between average volatility for N=100 versus average volatility for N=5. This observation can serve as a proxy for the impact of diversification through increasing the size (i.e. the number of stocks) in the portfolio. While EQR and LV start out roughly equivalent, the EQR model seems to decline in overall portfolio volatility faster than the LV model, more so at larger portfolio sizes. This makes sense, given EQR's incorporation of correlations and LV's assumption that all assets are uncorrelated.

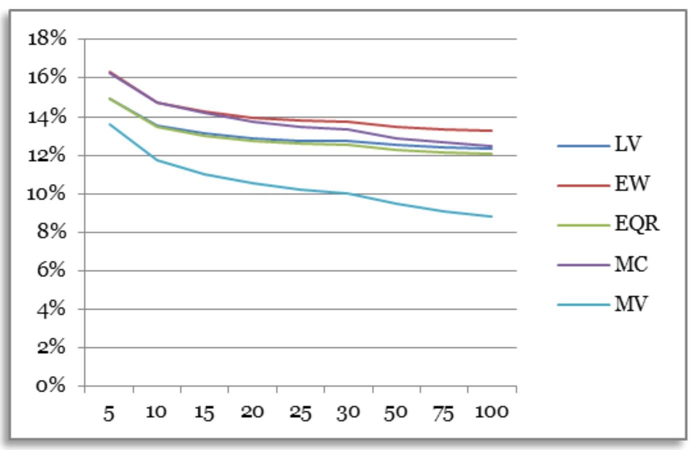

Figure 1: Average Portfolio Volatility versus Portfolio Size

Focusing in on the case where the portfolio contains 100 stocks randomly selected form the S&P 500, we plot the portfolio volatilities from each of the 500 iterations rank ordered on a per-model basis in Figure 2.

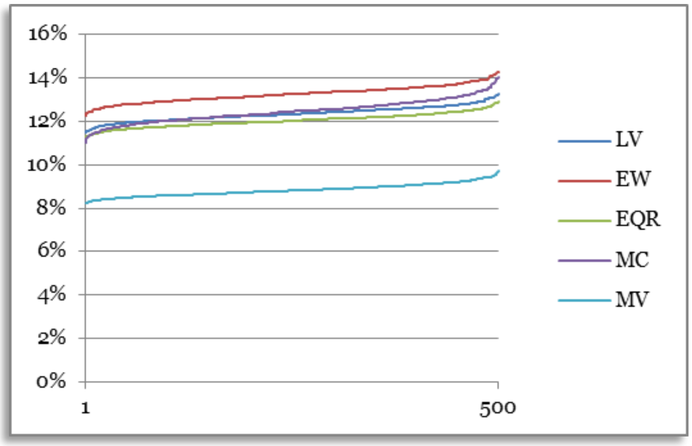

Figure 2: Ordered Portfolio Volatility for N=100

In other words, we rank order the MV portfolio volatilities, the MC, portfolio volatilities, etc. What we see is quite illuminating. Needless to say, the MV volatilities are always the lowest volatilities. What we also observe is that the distribution of EQR portfolio volatilities is the next lowest. We can further confirm that the EQR volatility is, in every instance, lower than its LV counterpart, as shown in Figure 3 (where the red line is the 45-degree line).

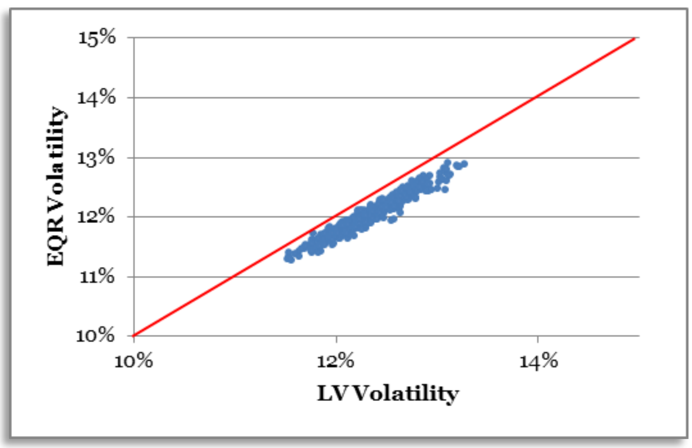

Figure 3: EQR Portfolio Volatility v. LV Portfolio Volatility for N=100

In other words, for portfolios of size 100, every portfolio had an *ex ante* volatility that was lower in the EQR solution than in the LV solution. More broadly, every iteration had a confirmed ordering from highest volatility to lowest volatility of EW, LV, EQR, and MV, with MC mostly more volatile than LV, EQR and MV and less volatile than EW. On occasion the MC model performed outside this range, as seen in Figure 4.

On a related note, it has been shown that under certain expected return conditions, maximum drawdown is directly related to portfolio volatility. Therefore, under those conditions, the risk ordering solved for in this section is the same as the maximum drawdown ordering for the same portfolios.

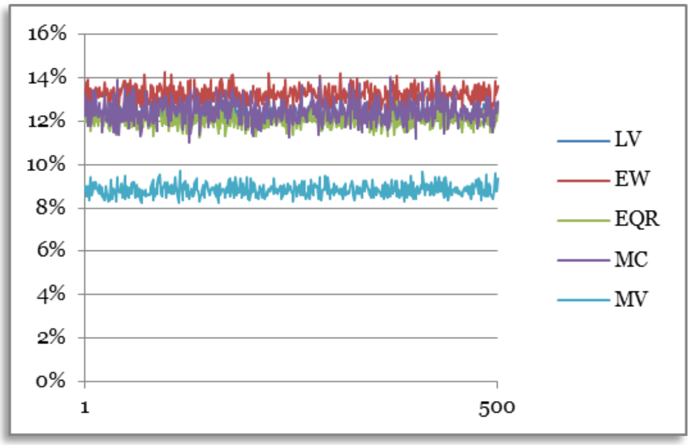

Figure 4: Portfolio Volatilities for N = 100.

Results: Risk Contributions

The EQR methodology highlights another way to look at the inclusion of each position in a portfolio, namely its total contribution to overall portfolio risk. This matters because it indicates how much we may be concentrating our risk to idiosyncratic factors which are not reflected ex-ante. We know that the EQR portfolio is designed to equate risk contributions, but how do the other models fare with respect to this metric?

In Figure 5 we plot the ratio of the highest risk contribution stock to the lowest risk contribution stock for portfolios of 100 stocks randomly selected from the S&P 500. We only plot four of the five strategies under consideration, because the MV portfolio would have such an extreme ratio that it is literally off the charts. We also see from the figure that the MC portfolio can have risk contributions that are orders of magnitude higher for some stocks than for others. In one iteration, the highest contribution was about 500 times higher than the lowest.

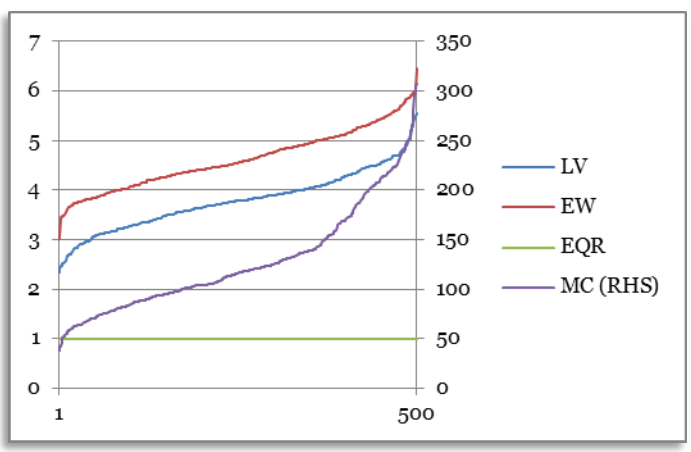

Figure 5: Ratio of the Highest Risk Contribution to the Lowest Risk Contribution (Ordered) for N = 100.

The LV, EW and EQR lie more in the same range, though not without variation: by definition, the EQR model has the lowest ratio in every iteration, a constant of one. EW contribution ratios range from 4 to more than 9, while LV contribution ratios range from about 2.5 to 5.5.

Rather than look at the extremes, we also look at the ratios of the highest risk contribution to the mean risk contribution for each portfolio in Figure 6. We see a similar story, where now once again we see large ratios for MC and MV, and smaller ratios for EW, LV and EQR. On average, the EW ratio tends to be a bit more than 2, and the LV ratio tends to be just shy of 1.5. Again, by definition, the EQR ratio is 1.

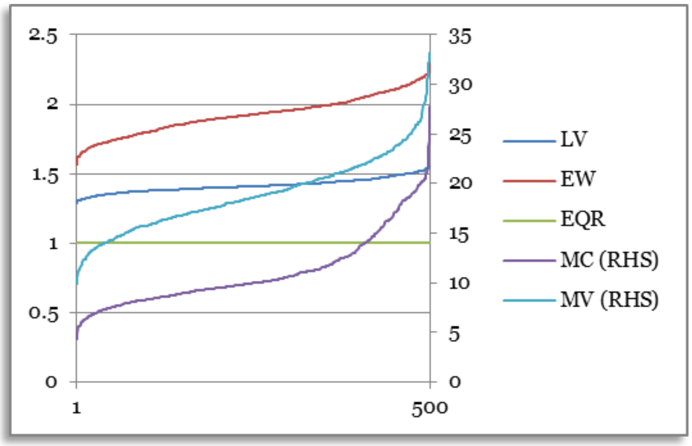

Figure 6: Ratio of the Highest Risk Contribution to the Mean Risk Contribution (Ordered) for N = 100.

Conclusion

The efficient portfolio is difficult to achieve in practice. The standard methodology of market-capitalization weighting has been shown to be sub-optimal, and behavioral considerations lead to anomalies like the low-volatility and low-beta anomalies. With the challenges of expected return estimation and parameter estimation in general, mean-variance optimization can be as difficult in practice as it is elegant in theory.

Investors need a pragmatic methodology that takes optimization considerations into account. We believe that the Equal Risk Weighting (EQR) approach is the best approach, one that seeks to minimize variance but spreads the risk of the portfolio equally across the constituents in the portfolio. EQR importantly takes account of correlation variations, unlike Low Volatility, Market Capitalization and Equal Weighting, and as our results show, strikes a good balance between overall risk of the portfolio and risk concentration among holdings.

General Disclaimer

Exposure to an asset class is available through investable instruments based on an index. It is not possible to invest directly in an index. There is no assurance that investment products based on the index will accurately track index performance or provide positive investment returns.

This document does not constitute an offer of services in jurisdictions where VelocityShares or its affiliates do not have the necessary licenses. VelocityShares receives compensation in connection with licensing its indices to third parties.

All information provided by VelocityShares is impersonal and not tailored to the needs of any person, entity or group of persons. VelocityShares is not an investment advisor, and VelocityShares make no representation regarding the advisability of investing in any such investment fund or other vehicle. A decision to invest in any such investment fund or other vehicle should not be made in reliance on any of the statements set forth in this document. Prospective investors are advised to make an investment in any such fund or other vehicle only after carefully considering the risks associated with investing in such funds, as detailed in an offering memorandum or similar document that is prepared by or on behalf of the issuer of the investment fund or other vehicle. Inclusion of a security within an index is not a recommendation by VelocityShares to buy, sell, or hold such security, nor is it considered to be investment advice.

VelocityShares is not a tax advisor. A tax advisor should be consulted to evaluate the tax consequences of making any particular investment decision.

VelocityShares does not guarantee the accuracy and/or completeness of any VelocityShares index, any data included therein, or any data from which it is based, and VelocityShares shall have no liability for any errors, omissions, or interruptions therein. VelocityShares makes no warranties, express or implied, as to results to be obtained from use of information provided by VelocityShares and used in this service, and VelocityShares expressly disclaims all warranties of suitability with respect thereto. While VelocityShares has obtained information believed to be reliable, VelocityShares shall not be liable for any claims or losses of any nature in connection with information contained in this document, including but not limited to, lost profits or punitive or consequential damages, even if it is advised of the possibility of same. These materials have been prepared solely for informational purposes based upon information generally available to the public from sources believed to be reliable. VelocityShares makes no representation with respect to the accuracy or completeness of these materials, the content of which may change without notice. The methodology involves rebalancing and maintenance of the indices that are made periodically during each year and may not, therefore, reflect real-time information.

VelocityShares and its affiliates provide a wide range of services to, or relating to, many organizations, including issuers of securities, investment advisers, broker-dealers, investment banks, other financial institutions and financial intermediaries, and accordingly may receive fees or other economic benefits from those organizations, including organizations whose securities or services they may recommend, rate, include in model portfolios, evaluate or otherwise address.