DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

Seeking Alpha? It's a Bad Guideline for Portfolio Optimization

Moshe Levy

Hebrew University

Richard Roll

California Institute of Technology

SOCIAL SCIENCE WORKING PAPER #1411

December 2015

Seeking Alpha? It's a Bad Guideline for Portfolio Optimization

Moshe Levy* and Richard Roll**

Abstract

Alpha is the most popular measure for evaluating the performance of both individual assets and funds. The alpha of an asset with respect to a given benchmark portfolio measures the change in the portfolio's Sharpe ratio driven by a marginal increase in the asset's portfolio weight. Thus, alpha indicates which assets should be marginally over/underweighted relative to the benchmark weights, and by how much. This study shows that alpha is actually a bad guideline for portfolio optimization. The reason is that alpha only measures the effects of *infinitesimal* changes in the portfolio weights. For *small but finite* changes, which are those relevant to investors, the optimal weight adjustments are almost unrelated to the alphas. In fact, in many cases the optimal adjustment is in the opposite direction of alpha – it may be optimal to *reduce* the weight of an asset with a positive alpha, and vice versa. Rather than employing alphas as a guideline, one can do much better by direct optimization with the desired constraint on the distance from the benchmark portfolio weights.

Keywords: Alpha, portfolio performance, portfolio optimization, mean-variance analysis.

JEL Classification: G11.

*Jerusalem School of Business, The Hebrew University, Jerusalem 91905, Israel. 972- 2-588-1329. mslm@huji.ac.il.

**California Institute of Technology, 1200 East California Blvd, Pasadena, CA 91125. 626-395-3890. rroll@caltech.edu.

1. Introduction

Alpha is the primary measure for the performance of individual assets as well as mutual funds.¹ It measures excess average return over and above the return expected given the asset's risk exposure. There are two ways to view alpha, and to employ it. The first is as a measure of a fund manager's stock selection and market timing abilities.² The second is as a guideline for the investor wishing to optimize her portfolio. The alpha of an asset, calculated with respect to a given benchmark portfolio, measures the change in the portfolio's Sharpe ratio driven by a marginal increase of the asset's weight in the portfolio. The vector of alphas is thus the direction of marginal adjustment in portfolio-weight space that yields the maximal increase in the portfolio's Sharpe ratio. Alphas tell the investor how to best marginally adjust her portfolio relative to the benchmark: increase the weight of assets with positive alphas, and decrease the weight of assets with negative alphas (and do so proportionately to the absolute size of alpha). Indeed, many investors seek alpha and tilt their portfolios toward assets with positive alphas. This paper is about this second view of alpha, as a tool for portfolio optimization.

While it is recognized that, strictly speaking, alpha is a guideline only for marginal adjustments (see, for example, Dybvig and Ross 1985, and Ferson and Lin 2014), alpha is widely perceived as a good rule-of-thumb also in practical situations with finite portfolio adjustments, hence its extreme popularity. We examine the usefulness of alpha as a guideline for portfolio optimization by examining the increase in the Sharpe ratio obtained by shifting the portfolio weights “in the direction” of the alpha vector. We

¹ See, for example, Jensen (1968), Grinblatt and Titman (1989), Ferson and Schadt (1996), Carhart (1997), Baks, Metrick, and Wachter (2001), Cremers and Petajisto (2009), and Amihud and Goyenko (2013).

² Roll (1978) shows that in this setting performance may be very sensitive to the benchmark employed.

find that while alpha indeed indicates the best way to make an *infinitesimal* adjustment to the portfolio weights (by its mathematical definition), it is *not* useful as a practical guideline, where small but finite adjustments are considered. Rather than adjusting the weights according to alphas, the investor can do much better by directly optimizing the portfolio with the desired constraint on the degree of deviation from the benchmark portfolio weights. When finite adjustments are considered, adjustments according to alpha are not only sub-optimal, but may in fact be in the “wrong direction” – in many cases the weight of an asset with a positive alpha should optimally be *reduced*, and vice versa.

Perhaps surprisingly, these effects take place even when the adjustments relative to the benchmark are small. For instance, if one allows portfolio weights to deviate from the benchmark portfolio weights by only 2% on average, the Sharpe ratio obtained with the alpha-adjusted portfolio is about 30% lower than the Sharpe ratio obtained with the optimized portfolio with the same constraint on the distance from the benchmark. Alphas point in the “wrong direction” for 15% of the assets in this case. For larger deviations from the benchmark these problems become even more severe. Thus, it is hard to see alphas being useful in any practical optimization context.

The reason that alpha is a bad guideline for portfolio optimization is that it is an indication only about the best *infinitesimal* shift in portfolio weights. Once the portfolio is shifted, alphas may change considerably. Thus, making large adjustments in the direction of the original alphas could be very sub-optimal. It is making a large step in the direction of the local gradient in order to search for a global maximum. This is closely related to the well-known fact that a small change in the benchmark portfolio can lead to a big

change in the beta-expected-return relationship, and thus to a big change in the assets' alphas (Roll and Ross 1994). This implies that while making a small shift of the portfolio weights in the direction of the alpha vector increases the portfolio's Sharpe ratio, as one changes the portfolio, the alphas may also change quite dramatically. Therefore, continuing to move "in the direction" of the original alpha vector beyond the initial infinitesimal shift may be far from optimal, as this is no longer "the right direction".²

Theoretical Background

The alpha of an asset relative to some benchmark portfolio is proportional to the derivative of the portfolio's Sharpe ratio with respect to the weight of the asset in the portfolio. To see this, consider a general benchmark portfolio with a weight x_i^B in risky asset i , and a weight of $x_{r_f}^B = 1 - \sum_{i=1}^N x_i^B$ in the risk-free asset. The beta of asset i with respect to this benchmark portfolio is given by:

$$\beta_i = \frac{\text{Cov}(\tilde{r}_i, \tilde{R}_p)}{\sigma_p^2} = \frac{\sum_{j=1}^N x_j \sigma_{i,j}}{\sigma_p^2}, \quad (1)$$

where $\tilde{R}_p$ denotes the benchmark portfolio return, and σ_p denotes its standard deviation.

The alpha of asset i is given by:

$$\alpha_i = \bar{r}_i - [r_f + \beta_i (\bar{R}_p - r_f)]. \quad (2)$$

The derivative of the portfolio's Sharpe ratio with respect to x_i^B is: (3)

$$\frac{\partial \left(\frac{\bar{R}_P - r_f}{\sigma_P} \right)}{\partial x_i^B} = \frac{(\bar{r}_i - r_f)\sigma_P - (\bar{R}_P - r_f) \frac{1}{\sigma_P} \sum_{j=1}^N x_j \sigma_{i,j}}{\sigma_P^2} = \frac{1}{\sigma_P} (\bar{r}_i - [r_f + \beta_i (\bar{R}_P - r_f)]) = \frac{1}{\sigma_P} \alpha_i.^3$$

Thus, the vector of alphas measures the gradient of the Sharpe ratio, i.e. the direction in portfolio weight space that offers the maximal increase in the portfolio's Sharpe ratio.⁴

This is the foundation for the common wisdom asserting that one should overweight assets with positive alphas and underweight assets with negative alphas. Notice, however, that alpha is a local measure. It points to the direction of maximal increase in the Sharpe ratio for an *infinitesimal* shift in the portfolio weights. The goal of this paper is to examine whether alphas offer a good guideline for portfolio optimization when the shift in portfolio weights is small but not infinitesimal.

3. Analysis

Our goal is to examine the increase in the portfolio's Sharpe ratio when the portfolio weights are adjusted in the direction of the alpha vector, and compare it with the increase obtained with the optimal adjustment. We compare the improvements as a

³ Note that both $\bar{R}_P = \sum_{j=1}^N x_j^B \bar{r}_j + \left(1 - \sum_{j=1}^N x_j^B\right) r_f$ and $\sigma_P = \sqrt{\sum_{j=1}^N \sum_{k=1}^N x_j^B x_k^B \sigma_{j,k}}$ are functions of x_i^B . We have: $\partial \bar{R}_P / \partial x_i^B = \bar{r}_i - r_f$ and $\partial \sigma_P / \partial x_i^B = \frac{1}{2\sigma_P} 2 \sum_{j=1}^N x_j^B \sigma_{i,j}$ (where the 2 in the numerator is due to the fact that for each term $x_i^B x_j^B \sigma_{i,j}$ in the double summation there is an equal symmetric term

$x_j^B x_i^B \sigma_{j,i}$). Equation (3) is obtained from the quotient rule $\left(\frac{f}{g} \right)' = \frac{f'g - g'f}{g^2}$. For an alternative proof showing that a marginal increase in the weight of an asset with a positive alpha increases the portfolio's Sharpe ratio, see Theorem 5 in Dybvig and Ross (1985).

⁴ For the tangency portfolio, which has the maximal Sharpe ratio, all the alphas are zero. This is true for any portfolio on the mean/variance efficient frontier, with the appropriate zero-beta rate, see Black (1972) and Roll (1977).

function of the magnitude of the adjustment, i.e. the distance from the benchmark portfolio weights. For very small adjustments, the adjustment in the direction of the alpha vector by definition yields the maximal increase possible, as discussed in the previous section. The interesting comparison is for adjustments that are larger than infinitesimal.

We conduct the analysis with a benchmark portfolio which is the value-weighted portfolio of the 100 largest U.S. stocks as of December 31, 2014 that have return records for the previous decade. The expected returns and covariances for these stocks are taken as their monthly sample values over the January 1, 2005 – December 31, 2014 decade. The risk-free rate is taken as the average monthly T-Bill rate over the same period, which is 0.11%.

There are two important points that should be stressed. First, the results are robust to the data employed. Very similar results are obtained under different specifications: different assets, a different number of assets, different sample periods, etc. The reason for this robustness is that the results are driven by a fundamental mathematical property, which is that the gradient (i.e. the alpha vector) is only a local measure of maximal improvement. Second, our analysis does not deal with sampling errors. The true parameters are assumed to be known. We take the empirical sample values just as a set of reasonable values for the true parameters. Sampling errors are, of course, a central issue in portfolio optimization. However, they are not relevant to the point of this paper: they do not bias the results in the comparison between adjustments in the direction of the alpha vector versus adjustments in other directions. Note that we are considering rather small adjustments, thus, in all cases the weights are not very different than the benchmark

portfolio weights. Sampling errors just add a layer of noise, which we wish to avoid in order to obtain clear results.

Results

When the portfolio is adjusted in the direction of the alpha vector, the adjusted weights are given by:

$$x_i = x_i^B + a\alpha_i, \quad (4)$$

where x_i is the adjusted weight of asset i , x_i^B is the weight of asset i in the benchmark portfolio, α_i is its alpha, and a is a positive constant determining the magnitude of the adjustment.⁵ We report the Sharpe ratio as a function of the deviation from the benchmark, measured by the Euclidean distance in portfolio weight space, normalized by the number of assets:

$$D \equiv \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - x_i^B)^2}. \quad (5)$$

The improvement in the Sharpe ratio as the weights are adjusted in the direction of the alpha vector are shown in Figure 1. The figure also shows the maximal Sharpe ratio that can be obtained under the same constraint on the distance from the benchmark (bold line). This optimal adjustment is found numerically by employing Matlab's *fmincom* constrained optimization algorithm. For small adjustments, the adjustment in the direction of the alpha vector yields the maximal improvement in the Sharpe ratio, as theoretically expected. However, for distances greater than 0.01, a value which implies that the average adjustment in portfolio weights is only about 1%, there is a difference

⁵ As the alphas do not necessarily add up to 0, the x_i 's do not necessarily add up to 1, and we therefore normalize them so that they do. This normalization does not change the adjusted portfolio's Sharpe ratio, as the Sharpe ratio is invariant to leverage.

between the alpha-adjustment and the optimal adjustment. For a distance of 2% ($D=0.02$) the Sharpe ratio of the alpha-adjusted portfolio is about 30% lower than the Sharpe ratio of the optimally adjusted portfolio. This difference increases with the distance D , and becomes very substantial even when the weight adjustments are only in the magnitude of 3%-4%. It is interesting to note that after an initial increase in the Sharpe ratio, continuing to shift weights in the direction of the alpha vector leads to a decrease in the Sharpe ratio. This is because once the portfolio has been shifted, the new alphas, relative to the shifted portfolio, are different from the original alphas, and the original alpha vector no longer points to the “right direction”.

In many of the cases where the adjustment is not infinitesimal, the alpha-adjustment is not only suboptimal, but may actually be in the wrong direction for some of the assets. These are cases where an asset’s alpha with respect to the benchmark is positive, but in the optimally adjusted portfolio its weight should be *reduced* relative to its benchmark portfolio weight, and vice versa. Figure 2 shows the percentage of assets that are optimally adjusted in the opposite direction of alpha, as a function of the distance D .

(Please insert Figures 1, 2 and 3 about here)

In general, when non-infinitesimal adjustments are considered, there is almost no relation between an asset’s alpha and its optimal adjustment. For example, Figure 3 shows the relationship between the assets’ alphas and their optimal adjustments when the distance is $D=0.04$.

These results all imply that alphas are a bad guideline for portfolio optimization: they are almost unrelated to the optimal adjustment, in many cases they actually lead to an adjustment in the wrong direction, and they yield portfolios that are significantly suboptimal relative to the optimal portfolio one could achieve with the same distance from the benchmark.

4. Discussion

The alpha vector points to the direction of maximal increase in the Sharpe ratio for an infinitesimal adjustment in the portfolio weights. The reason that alphas are a bad practical guideline as to how to adjust one's portfolio is that once the portfolio is slightly adjusted, the alphas change, and thus there is a new direction for optimal adjustment. Continuing to go "in the same direction" of the original alpha vector is therefore suboptimal.

If the change in alphas would have been very gradual, the original alpha vector would still be a useful practical guideline, as the optimal direction wouldn't have changed that much. However, this is not the case. Roll and Ross (1994) have shown that betas can be quite sensitive to the portfolio composition. This implies the sensitivity of alphas as well. Indeed, when we directly examine the change in alphas as the portfolio is shifted, this sensitivity is exactly what we find. Figure 4 shows the relationship between the assets' original alpha values, calculated with respect to the benchmark portfolio, and their new alphas, calculated with respect to the new portfolio, which is the benchmark portfolio shifted in the direction of alpha. Panel A shows the relationship for a small shift of $D=0.01$. In this case the new alphas are rather closely related to the original alphas.

However, for a slightly larger shift of $D=0.02$ (recall that this implies that weights are adjusted by only about 2%) there is almost no relation between the original alphas and the new alphas. Thus, the direction of optimal adjustment has completely changed.

(Please insert Figure 4 about here)

5. Conclusion

Alpha is widely accepted as a guideline for portfolio optimization, and has thus become a central performance measure for both individual assets as well as mutual funds. It is widely believed that overweighting assets with positive alphas and underweighting assets with negative alphas optimally increases the portfolio's Sharpe ratio. While this is true for infinitesimal shifts in the portfolio weights, for slightly larger shifts in the weights, of 1% or more, alpha can be a quite misleading guideline. We show that for finite-size adjustments to the portfolio weights one can do much better than shifting the weights in the direction of the alpha vector. For a given distance from the benchmark portfolio, a much larger increase in the Sharpe ratio can be obtained with direct optimization. Alternatively, one can consider modified and heterogeneous alphas that incorporate the investor's utility function and information set (Sharpe 1982, Ferson and Lin 2014). There is almost no relation between traditional alphas and the optimal shift in the portfolio weights. In fact, in many cases the alphas point to the wrong direction: it is optimal to underweight an asset with positive alpha, and vice versa. Thus, alphas do not provide a helpful practical guideline for portfolio optimization.

References

- Amihud, Yakov, and Ruslan Goyenko. "Mutual fund's R2 as predictor of performance." *Review of Financial Studies* (2013): hhs182.
- Baks, Klaas P., Andrew Metrick, and Jessica Wachter. "Should investors avoid all actively managed mutual funds? A study in Bayesian performance evaluation." *The Journal of Finance* 56, no. 1 (2001): 45-85.
- Black, Fischer. "Capital market equilibrium with restricted borrowing." *Journal of Business* (1972): 444-455.
- Carhart, Mark M. "On persistence in mutual fund performance." *The Journal of Finance* 52, no. 1 (1997): 57-82.
- Cremers, KJ Martijn, and Antti Petajisto. "How active is your fund manager? A new measure that predicts performance." *Review of Financial Studies* 22, no. 9 (2009): 3329-3365.
- Dybvig, Philip H., and Stephen A. Ross. "The analytics of performance measurement using a security market line." *The Journal of Finance* 40.2 (1985): 401-416.
- Ferson, Wayne, and Jerchern Lin. "Alpha and performance measurement: the effects of investor disagreement and heterogeneity." *The Journal of Finance* 69.4 (2014): 1565-1596.
- Ferson, Wayne E., and Rudi W. Schadt. "Measuring fund strategy and performance in changing economic conditions." *The Journal of Finance* 51, no. 2 (1996): 425-461.
- Goetzmann, William, Jonathan Ingersoll, Matthew Spiegel, and Ivo Welch. "Portfolio performance manipulation and manipulation-proof performance measures." *Review of Financial Studies* 20, no. 5 (2007): 1503-1546.
- Grinblatt, Mark, and Sheridan Titman. "Portfolio performance evaluation: Old issues and new insights." *Review of Financial Studies* 2, no. 3 (1989): 393-421.
- Jensen, Michael C. "The performance of mutual funds in the period 1945–1964." *The Journal of Finance* 23, no. 2 (1968): 389-416.
- Roll, Richard. "A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory." *Journal of Financial Economics* 4, no. 2 (1977): 129-176.
- Roll, Richard. "Ambiguity when performance is measured by the securities market line." *The Journal of Finance* 33, no. 4 (1978): 1051-1069.

Roll, Richard, and Stephen A. Ross. "On the cross-sectional relation between expected returns and betas." *The Journal of Finance* 49, no. 1 (1994): 101-121.

Sharpe, William F. "Security Codings: Measuring Relative Attractiveness in Perfect and Imperfect Markets." *WF Sharpe PM Cootner, Financial Economics: Essays in Honor of Paul Cootner*, (New Jersey: Prentice-Hall) (1982): 216-229.

Figure 1

The Sharpe ratio as a function of the distance D from the benchmark portfolio (see eq. 5). The light line shows the adjustment in the direction of the alpha vector (as in eq. 4). The bold line shows the optimal adjustment with the same constraint on the distance from the benchmark portfolio. While the alpha-adjustment is optimal for very small adjustments, for adjustments of $D=0.01$ (which implies an average adjustment of about 1% in portfolio weights) or more, one can do much better than the alpha-adjustment. After a small adjustment, continuing to tilt the portfolio in the direction of the alpha vector *reduces* its Sharpe ratio.

Figure 2

The percentage of assets that are optimally adjusted in a direction opposite to their alphas, i.e. assets with positive alphas that are optimally underweighted relative to the benchmark, and vice versa.

Figure 3

The relationship between alphas and the optimal adjustments for a distance of $D=0.04$ from the benchmark portfolio. x_i^* is the weight of asset i in the optimal portfolio under the distance constraint, and hence $\Delta x_i = x_i^* - x_i^B$ is the optimal adjustment relative to the benchmark. Alphas provide almost no clue as to the way the portfolio weights should be optimally adjusted.

Figure 4

The relation between the original alphas (calculated relative to the benchmark portfolio) and the new alphas, calculated relative to the new portfolio, which is the benchmark portfolio adjusted in the direction of the original alphas. For a distance $D=0.01$ alphas

don't change that much (Panel A), but for a slightly larger distance of $D=0.02$ there is almost no relation between the original alphas and the new alphas. Thus, after a small shift the original alphas no longer point in "the right direction".