Portfolio Sharpening

Patrick Burns*

21st September 2003

Abstract

We explore the effective gain or loss in alpha from the point of view of the investor due to the volatility of a fund and its correlations to other asset classes. Fund managers and investors can be guided by this to increase the utility that is ultimately delivered to the investor. In this analysis of investor utility the Sharpe ratio is shown to be misleading and the tracking error has no role at all. A new class of funds—called "hyperpassive"—is suggested which are similar to traditional index funds, but which aim to deliver a comparable expected return with less volatility than the benchmark. It is also shown that the optimal allocation to additional asset classes can be surprisingly high when the correlations are small.

1 Introduction

The job of a fund manager is to provide a service to investors. A traditional measure of the quality of a fund manager is the Sharpe ratio. While the Sharpe ratio is valid evidence of the skill of the manager, it is not a good measure of the value of the manager's fund to the investor. Indeed, we will see that a negative Sharpe ratio could be good for the investor, and a positive Sharpe ratio could be bad.

The idea of portfolio sharpening is to mold the manager's portfolio in such a way as to be of more value to the investor. In essence, to get more "Sharpe ratio" out of the same level of skill.

In what follows we do a small amount of manipulation of the investor's utility function, and then discuss the implications. Similar to [Waring and Siegel, 2003] we assume that the investor has settled on an asset allocation. So the manager has a set fraction of the investor's portfolio, and there is a benchmark for the manager's portion. Specifically we are interested in the investor's choice—as illustrated in Figure 1—between putting a fraction w of the portfolio into a replication of the benchmark or using the manager for that same fraction.

*This report can be found in the working papers section of the Burns Statistics website http://www.burns-stat.com/.

Figure 1: The investor's choice between a passive and an active strategy for the portion in question, where V is the total value of the portfolio.

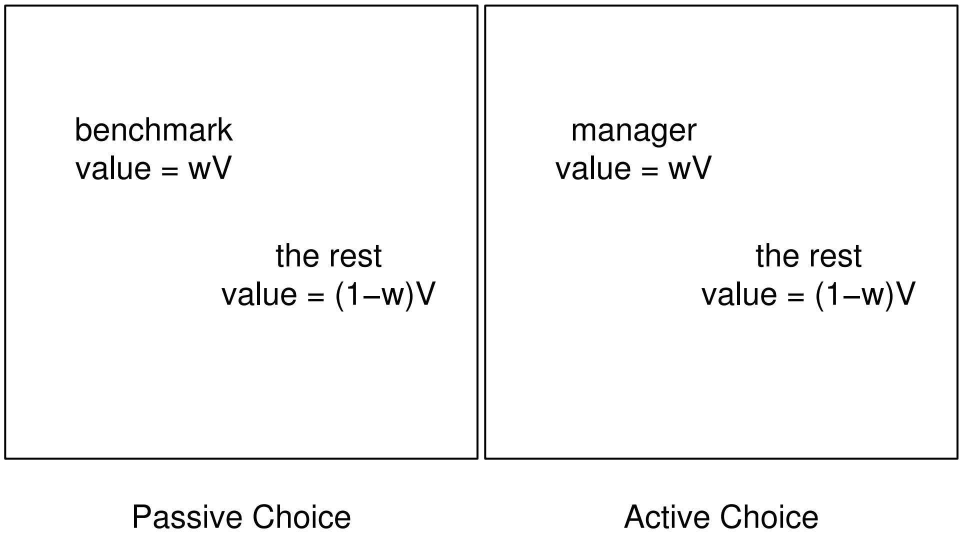

Section 2 presents the utility for the investor's decision. The risk aversion parameter is discussed in Section 3 and an attempt is made to get a sense of the values that investors might have. Section 4 develops the concept of the alpha proxy—the measure of the effect on alpha due to volatilities and correlations—and provides some indication of its size. Hyperpassive funds are introduced in Section 5, while active management is discussed in Section 6. Section 7 looks at adding new asset classes to an asset allocation. Section 8 concludes.

2 The Investor's Utility

Assume that the investor has a mean-variance utility with risk aversion λ. While mean-variance utility is subject to criticism—especially if the portfolio includes options—it is quite a reasonable first approximation. Though the mathematics of what follows would change (undoubtedly towards more complex) with a different utility function, the implications should remain largely intact.

The general mean-variance utility is

$$U = \mu - \lambda\sigma^2 \tag{1}$$

where μ is the expected return of the portfolio and σ^2 is the portfolio variance.

First we look at the case where the investor uses the benchmark instead of the manager. The portfolio can be broken into two pieces—the benchmark and the rest of the portfolio. The utility in this case will be:

$$U_b = w\mu_b + (1-w)\mu_o - \lambda[w^2\sigma_b^2 + 2w(1-w)\sigma_b\sigma_o\rho_{bo} + (1-w)^2\sigma_o^2] \quad (2)$$

Here μ_b is the expected return of the benchmark, μ_o is the expected return of the rest of the portfolio, w is the fraction of the total portfolio invested in the benchmark, σ_b is the volatility of the benchmark, σ_o is the volatility of the rest of the portfolio, and ρ_{bo} is the correlation between the benchmark and the rest of the portfolio.

Equation 3 is exactly the same as Equation 2 except the "b" subscripts are replaced by "m" to mean manager—that is, this is the utility when the benchmark is replaced by the manager's portfolio.

$$U_m = w\mu_m + (1-w)\mu_o - \lambda[w^2\sigma_m^2 + 2w(1-w)\sigma_m\sigma_o\rho_{mo} + (1-w)^2\sigma_o^2] \quad (3)$$

Taking the difference of these, we get:

$$U_m - U_b = w(\mu_m - \mu_b) - \lambda w^2(\sigma_m^2 - \sigma_b^2) - 2\lambda w(1-w)\sigma_o(\sigma_m\rho_{mo} - \sigma_b\rho_{bo}) \quad (4)$$

Let

$$\alpha = \mu_m - \mu_b$$

and recall that w is a constant, hence it can be divided out of the utility. So we can revise Equation 4 to be:

$$\Delta U^* = \alpha - \lambda w(\sigma_m^2 - \sigma_b^2) - 2\lambda(1-w)\sigma_o(\sigma_m\rho_{mo} - \sigma_b\rho_{bo}) \quad (5)$$

Equation 5 is the one of interest. This gives the change in utility from hiring the manager instead of tracking the benchmark—the value of the manager to the investor.

Ideally ΔU^* is optimized. But above all the manager should ensure that it is positive. If ΔU^* is negative, then the manager is harming the investor rather than helping. Since each of the three terms in Equation 5 can be negative or positive, it is possible for the manager to add value for the investor even with a negative alpha. In the same light, the manager can degrade the investor's portfolio even with a positive alpha.

We will return to ΔU^* after we look at the risk aversion parameter.

3 The Risk Aversion Parameter

The risk aversion parameter is unchanged when expected returns and the variance are annualized since both the expected returns and the variance are multiplied by the frequency. However, the risk aversion is dependent on whether fractions or percent are used. When switching to percent, the expected returns are multiplied by 100 and the variance is multiplied by 10,000. So the risk aversion parameter is divided by 100.

In all that follows the risk aversion parameter is expressed as if both the expected returns and the variance matrix are in terms of fractions.

The two-asset problem is simple enough that we can easily get a formula for the optimal weight. In addition to the risk aversion parameter there are just five quantities that enter—the expected returns μ_1 and μ_2, the standard deviations σ_1 and σ_2 and the correlation ρ. The process is to take the derivative of the utility with respect to one of the weights and set that to zero. With some algebraic manipulation, we get the optimal weight for asset 1 to be:

$$w_1^* = \frac{\mu_1 - \mu_2}{2\lambda(\sigma_1^2 - 2\sigma_1\sigma_b\rho + \sigma_2^2)} + \frac{-\sigma_1\sigma_2\rho + \sigma_2^2}{\sigma_1^2 - 2\sigma_1\sigma_2\rho + \sigma_1^2} \quad (6)$$

To satisfy the constraint that weights are between 0 and 1, w_1^* should be taken as zero whenever Equation 6 yields a negative number, and taken as 1 when it gives a number greater than 1. We can characterize the two terms on the right-hand side of Equation 6 as a risk-reward term and a pure diversification term. Note that if there is no difference in the expected returns, then the risk aversion parameter drops out.

Hypothetical equity-bond data gives a sense of the values that the risk aversion parameter might take on. In our first scenario, equity has an expected return of 0.08 and a volatility of 0.2, while bonds have an expected return of 0.05 and a volatility of 0.1. The correlation between equities and bonds is varied from 0 to 50%. Figure 2 shows the fraction of an optimal portfolio that is put into equities for a range of risk aversions.

We will be particularly interested in the meaning of a risk aversion of 1. One characterization of a risk aversion of 1 is that it gives a 50–50 allocation between two assets when the expected returns are 0.08 and 0.05, and the volatilities are 0.2 and 0.1, no matter what the correlation.

Figure 3 shows the fraction put into equity as the expected return of equities varies from 0.05 to 0.10 with the correlation fixed at 10%. The other values are as before: the expected return of bonds is 0.05, and the volatilities are 0.2 and 0.1.

4 The Alpha Proxy

We can rewrite Equation 5 as

$$\Delta U^* = \alpha + \lambda\xi$$

where

$$\xi = -w(\sigma_m^2 - \sigma_b^2) - 2(1-w)\sigma_o(\sigma_m\rho_{mo} - \sigma_b\rho_{bo}) \quad (7)$$

Figure 2: Optimal weight of equity as correlation varies in the hypothetical scenario where the expected returns are 0.08 for equities and 0.05 for bonds.

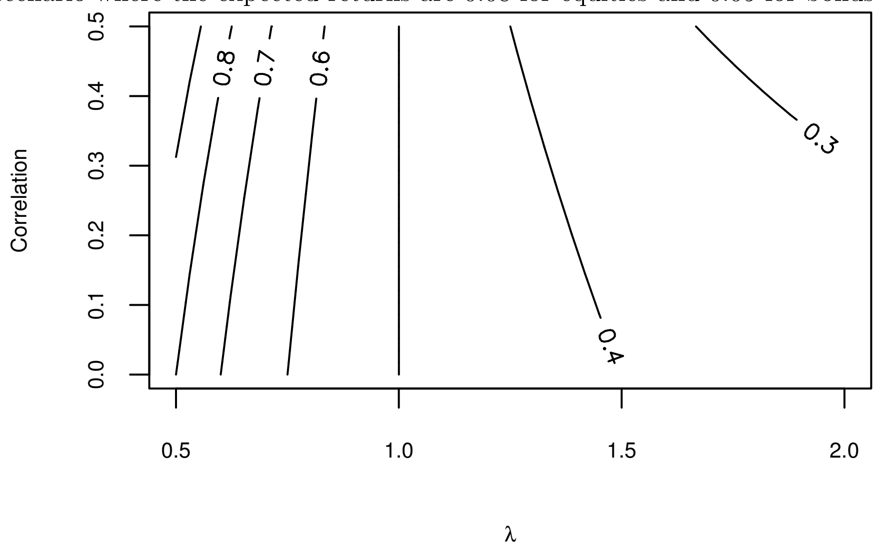

Figure 3: Optimal weight of equity as the equity expected return varies in the hypothetical scenario with correlation equal to 0.10.

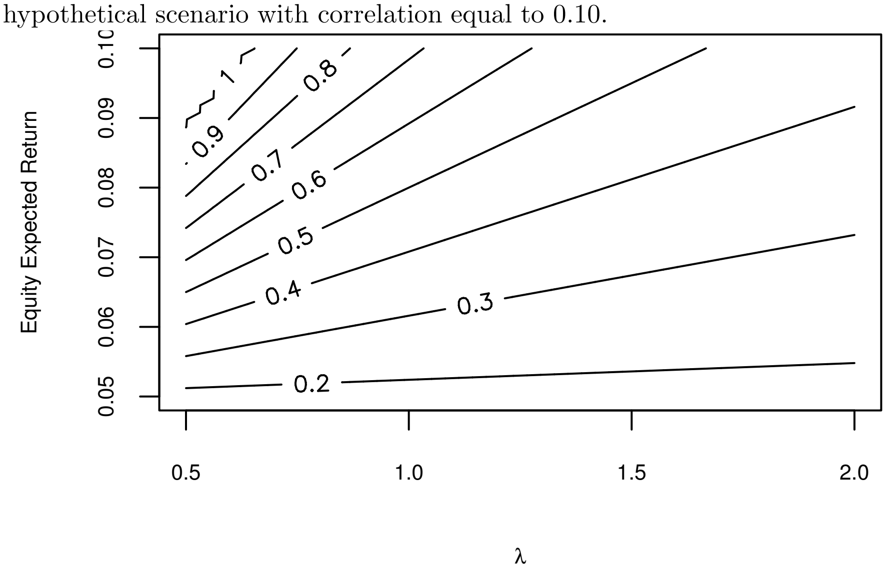

Figure 4: Alpha proxy (in basis points) with weight and manager volatility varied. The remaining four variables are each at 0.20.

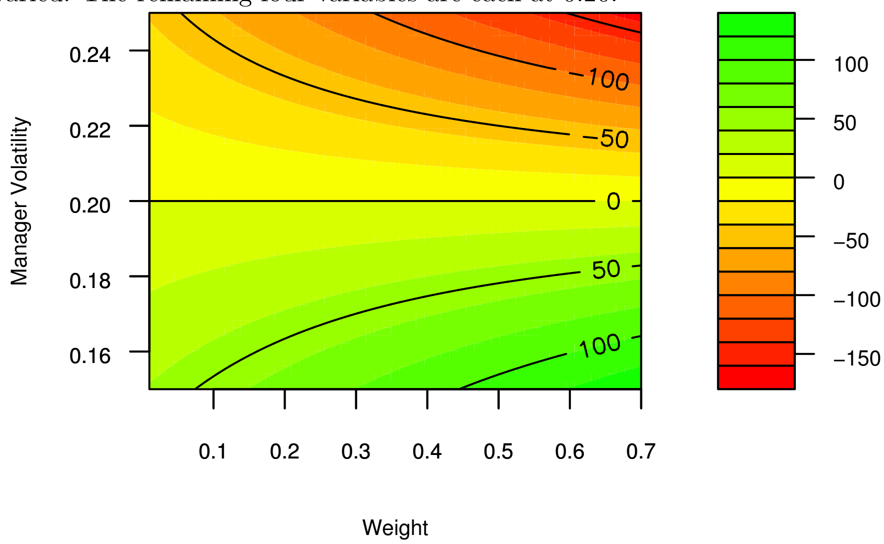

We call the quantity ξ the *alpha proxy*. ξ represents the gain or loss in alpha caused by the volatility and correlation of the manager's portfolio for an investor with risk aversion parameter equal to 1. The alpha proxy can be multiplied by the risk aversion parameter to get the equivalent alpha for investors whose risk aversion is not 1.

The alpha proxy is a function of six quantities—three volatilities, two correlations and the weight. We can use contour plots to look at two-dimensional slices. In all of the plots shown here the four variables that are not varied in the plot are set to 0.20. Figure 4 shows that the volatility of the manager portfolio becomes more important as the weight increases. The correlation of the manager portfolio with the remainder of the investor's portfolio becomes more important as the weight decreases, as is shown in Figure 5. Figure 6 shows the alpha proxy as the manager volatility and correlation both change.

Since a paper can not possibly present all of the plots that are of interest, the code that created Figures 4, 5 and 6 is available in the public domain section of the Burns Statistics website. The code runs in the R language [Ihaka and Gentleman, 1996] which can be obtained for free from http://www.r-project.org.

5 Hyperpassive Funds

Positive alphas are very hard to generate while volatilities are relatively easy to control. The second term of Equation 5 is positive if the manager's volatility is

Figure 5: Alpha proxy (in basis points) with weight and manager correlation varied. The remaining four variables are each at 0.20.

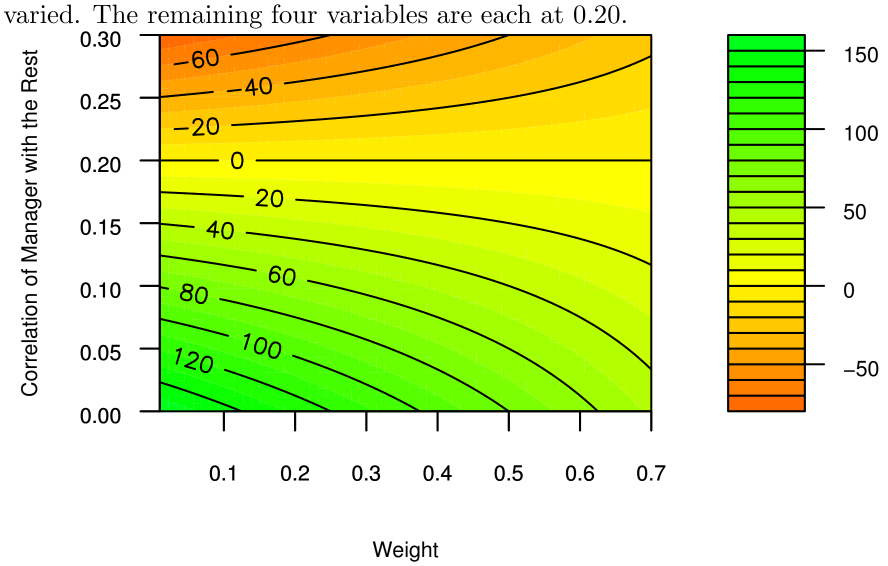

Figure 6: Alpha proxy (in basis points) with manager volatility and correlation varied. The remaining four variables are each at 0.20.

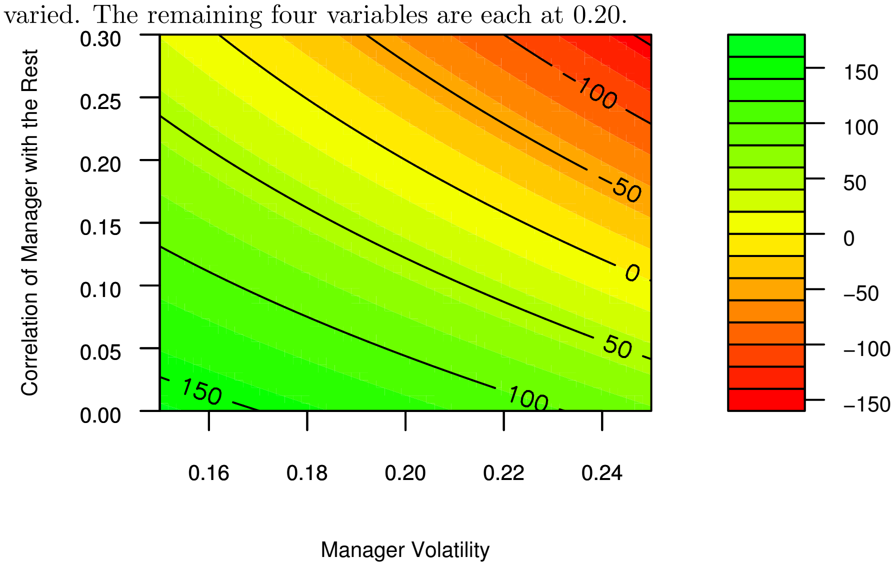

less than that of the benchmark. The same may well be true for the third term.

This observation suggests a class of funds that track a benchmark but have lower volatility—hyperpassive funds. At least one fund management company already has such a fund. Hyperpassive funds could be run as cheaply as conventional index funds—perhaps more cheaply—but would have more value to investors. Operationally the fund could be created and maintained by minimizing the volatility of the fund, possibly using a tracking error constraint. Some control on correlations could be implemented as well.

An index fund needs to invest in essentially the entire index, and is forced to rebalance whenever the benchmark is modified. Since a hyperpassive fund would not have such tight constraints, it could have an advantage over an index fund both because it avoids forced trading and because it need not hold all of the assets in the index. Neither of these funds has the expense of searching for positive alpha.

An argument against hyperpassive funds would be that part of the return of the benchmark is due to the risk premium of its volatility, so any portfolio with smaller volatility will have a lower expected return. However, it seems fairly extreme to suppose that all benchmarks are on the efficient frontier of the universe of assets available to the manager.

6 Active Management

Active managers should provide as much utility to their investors as possible. As we have seen, this is not just a matter of generating a positive alpha. The volatility of the manager's portfolio and its correlation with the rest of the investor's portfolio can have a significant impact on the effective alpha that is delivered to the investor.

To best serve the investor, the manager should perform an optimization from the investor's point of view. Ultimately the manager would need estimates of the covariances between the rest of the portfolio and each asset in the manager's universe. A history—possibly approximate—of the returns of the remainder of the portfolio may suffice to get the covariances. The manager would also need the fraction of the portfolio in each of the two portions, and the investor's risk aversion (make sure of the definition of the risk aversion parameter—there is sometimes a 2 involved).

If such information is not available or the fund is not dedicated to a single investor, then an approximate answer is required. One obvious approximation is to create what might be typical for the rest of the portfolio and guess the average weight and average risk aversion. With these an optimization can be done. Another approach is to use the alpha proxy as a guide for what volatilities and correlations should look like.

As we saw in Figures 4 and 5, the manager's volatility is very important when the weight is large, and the manager's correlation with the rest of the portfolio is most important when the weight is small. Since a manager is unlikely to have a large fraction of an investor's portfolio, it is the difference in covariances that is likely to dominate. This means that fund managers should attempt to control the correlation of their portfolio with what is likely to be in the rest of an investor's portfolio as well as controlling volatility. For example an equity manager should aim for small (possibly negative) correlation with fixed income. An equity growth manager might attempt small correlation with both fixed income and equity value funds.

It is important to note that tracking error does not appear in Equation 5. The tracking error involves the correlation between the manager's portfolio and the benchmark, but that is nowhere to be found. While volumes have been written about tracking error, this analysis suggests that it is largely irrelevant.

7 The Investor Allocation

Until now we have assumed that the investor's allocation was fixed, and the onus was on the manager to do as well as possible. Let's briefly glance in the other direction—at how the investor might alter the asset allocation. The alpha proxy suggests that adding asset classes with low correlation to the existing classes may be useful.

We'll explore this by looking at one hypothetical situation with two asset classes. The first class has expected return 0.08 and volatility 0.2, while the second class has expected return 0.05 and volatility 0.1. The correlation between them is 0.1. This is obviously trying to mimic equities and bonds. We'll suppose that our investor has a 60% allocation to the first class. An optimizer (or inverting Equation 6) tells us that the risk aversion parameter is about 0.77.

We now wonder how much of a third asset class the investor will want. The third class has an expected return of 0.05 and a volatility of 0.2. This is not an attractive class—it has the same expected return as the lower of the first two, but has the same volatility as the more volatile. Perhaps most of us would be disinclined to invest in this class at all. We can see how the correlation of this class with the other two affects the optimal allocation.

Figure 7 shows the allocation that our hypothetical investor should give to the new asset under the simplifying assumption that the correlations of the new asset to the old ones are equal. We see that we should have some investment in the new asset whenever the correlation is less than 15%. When the correlation is 10%, our investor should have more than 4% invested in the new asset.

Figure 8 shows the allocation to the new asset when its correlations to the old assets are not necessarily the same. The allocation is much more sensitive to the correlation with the first asset. That is, the investor will want the new asset class to have small correlation with the high-return, high-volatility asset. The correlation with the lower volatility asset plays a more minor role.

Figure 7: Allocation to the new asset when correlations with the old assets are equal.

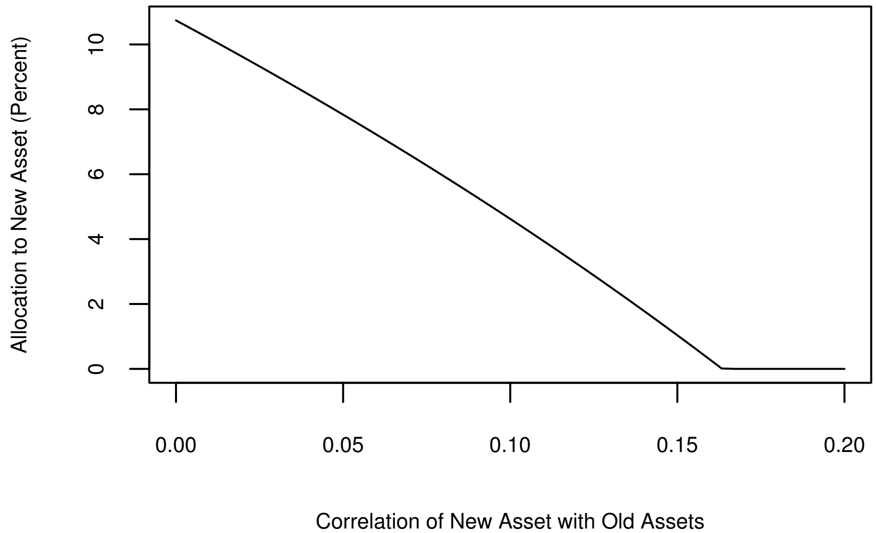

Figure 8: Allocation (in percent) of the new asset as its correlations vary with the old assets.

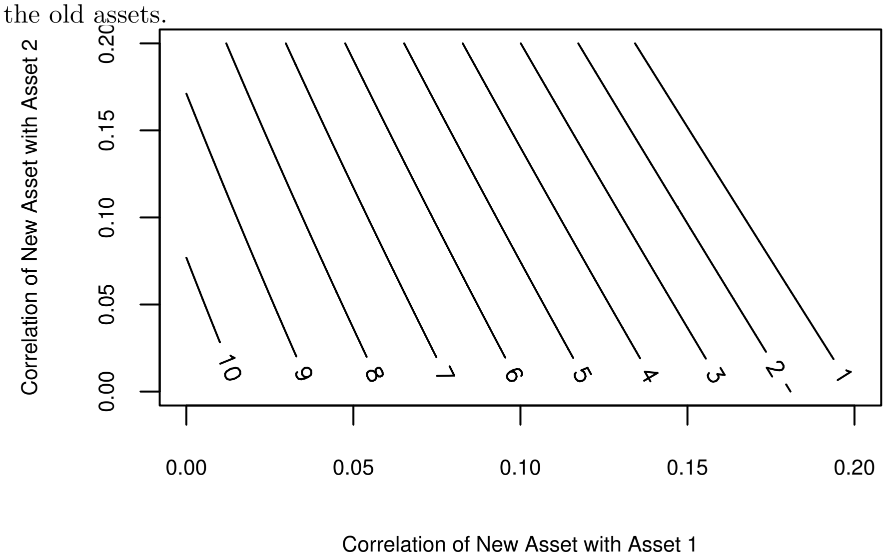

8 Discussion

When considering the utility that is actually delivered to the investor, there are currently no signposts available that say how well a manager is doing. From this perspective we should pay much more attention to the volatility of a manager's portfolio and its correlations to other asset classes. The roles of the Sharpe ratio and especially tracking error should be reduced.

From an active manager's point of view this should be good news. Volatilities and correlations are much easier to control (and verify) than alphas. So with minor effort active managers should be able to increase their value to the investors—to sharpen their portfolio from the perspective of the investor.

Action points for investors are to adjust the criteria on which active managers are judged, and to invest in hyperpassive funds. Additionally there is a suggestion that allocations should be increased to funds—such as market neutral funds (or market negative as in [Burns, 2003])—that have low correlation with other asset classes. The suggestion is confirmed by [Amenc and Martellini, 2002]. Gaining utility through diversification is generally an easier route than searching out extraordinary alpha generators.

References

[Amenc and Martellini, 2002] Amenc, N. and Martellini, L. (2002). Portfolio optimization and hedge fund style allocation decisions. *The Journal of Alternative Investments*, 5(2):7–20.

[Burns, 2003] Burns, P. (2003). Does my beta look big in this? Working paper, Burns Statistics, http://www.burns-stat.com/.

[Ihaka and Gentleman, 1996] Ihaka, R. and Gentleman, R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5(3):299–314.

[Waring and Siegel, 2003] Waring, M. B. and Siegel, L. B. (2003). The dimensions of active management. *The Journal of Portfolio Management*, 29(3):35–51.