THE POWER OF REBALANCING

Fact, Fiction and Why it Matters

April 2015

UNCORRELATED ANSWERS®

Adrian Banner, PhD
Chief Executive Officer
Chief Investment Officer

David Schofield
President
International Division

Richard Yasenchak, CFA
Senior Managing Director,
Head of Client Portfolio Management

Key Ideas

It is well understood that rebalancing is a necessary step in restoring a portfolio of volatile assets back to its target weights. Whether it is performed periodically or triggered when actual weightings move too far from target, rebalancing a portfolio will naturally lead to selling assets that have outperformed the portfolio and buying assets that have underperformed the portfolio.

It is much less widely understood that rebalancing can actually be a source of return for the portfolio. Despite the fact that this observation dates back to 1982 [Fernholz and Shay] and has been successfully used to manage portfolios for nearly as long, it has come under considerable attack in the recent past by some academics and practitioners. The main arguments used by these detractors are:

1. There is no return benefit because the portfolio's expected wealth does not increase.

2. The return benefit exists, but is due to diversification, not rebalancing.

3. The return benefit relies on mean-reversion.

These arguments may appear compelling at first glance, but all three are fundamentally flawed. This article explains why.

HARNESS VOLATILITY®

Intech®
JANUS HENDERSON

This page is intentionally left blank.

Typical Return vs. Expected Return

Don't be fooled by the terminology: 'expected' doesn't mean 'likely'.

1. Compound returns are more relevant than expected returns

The meaning of the expected return of a portfolio is often misunderstood. The expected return simply indicates how well the portfolio will perform, on average, in a *single predetermined* period of time. It is closely related to the concept of expected terminal wealth, which is the average value of the investment after a certain period of time (the 'term' of the investment). These measures can be useful for determining the soundness of an investment over a fixed period of time, but can also be highly misleading in the presence of compounding over multiple periods.

The real danger of focusing on expected return comes from overreliance on the simple average. Hundreds of years of research into statistics and probability theory have led to the understanding that the average does not tell the whole story. It is therefore perplexing to see papers published in the 21st century [Chambers and Zdanowicz; Cuthbertson et al.; Qian] explicitly or implicitly promote expected return and volatility as the primary determinants of the desirability of an investment. The real problem with this narrow view is that the median return can be much lower than the average return. This means that it is likely that the typical return experienced by the investor is substantially less than the expected return.

Furthermore, compounding the investment over multiple periods usually exacerbates the difference between the median and the average, so that, after a number of iterations, it becomes less and less likely that the expected (i.e., average) return is achieved. Don't be fooled by the terminology: 'expected' doesn't mean 'likely.'

To illustrate this, consider a speculative investment of $10 that has a 1% chance of hitting a $1,000,000 payoff in one year's time, but a 99% chance of falling to $1 (see Figure 1). The expected return[1] of the investment is a staggering 99,909.9%, but in this case the term 'expected return' is highly misleading: a rational person would probably 'expect' to lose $9 on this

FIGURE 1
OUTCOMES FOR A SPECULATIVE INVESTMENT

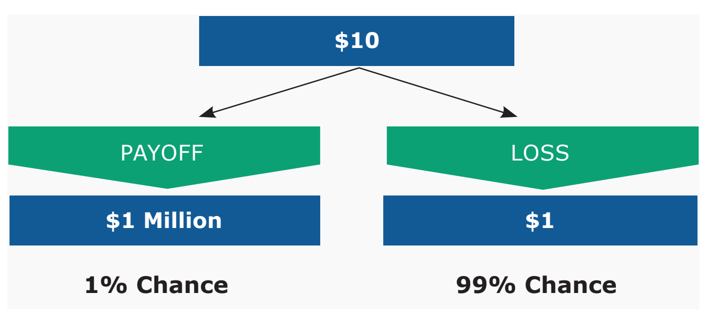

[1] The expected return is computed as follows: E(R) = 0.01 × (1000000/10 – 1) + 0.99 × (1/10 – 1) = 99,909.9%.

investment! It may still seem like a great opportunity since the upside potential is so huge, however very unlikely to be realized.

On the other hand, if you continue to reinvest your proceeds in the same investment, thus subjecting yourself to compounding, you will almost certainly be disappointed in the outcome. Assuming the different periods are independent, the 'expected return' becomes astronomically high over time, but there is also virtually no chance of getting anywhere near it. Indeed, for every five times the investment doesn't pay off, you need to hit the payoff once just to break even. (For example, even if it pays off the first time to give you the million, the next five losses will see you successively having $100K, $10K, $1K, $100 and then the original $10.) No surprise, then, that the compound return[2] of the investment is around −88.5%.

While the example is somewhat extreme in its outcomes, it is nonetheless illustrative of the main point: the expected return is most useful for understanding a reasonably narrow range of outcomes for a fixed term that is known in advance. For an investment with an open-ended time horizon, such as the vast majority of retirement-related investments, the compound return is what really matters. It is therefore appropriate to assess any gain in compound return as a return benefit, even if the expected return stays the same.

2. Rebalancing is vital to preserving the gains from diversification

Suppose you now have access to a hundred uncorrelated investments of the type in the above example. With an initial capital of $1,000, it seems reasonable to place $10 in each of the hundred investments. In fact, the expected return is the same no matter what you do: it is the same staggering but unlikely 99,909.9% we saw in part 1 above, regardless of whether you diversify or put the whole $1,000 in one of the investments. Diversification has no impact on the expected average return. The *compound* return, however, does change. Instead of the awful −88.5% compound return for a single investment, the diversified portfolio has a compound return of 4,186%. This is because there is a greater than 63.3% chance[3] that at least one of the hundred investments hits the big payoff.

If precisely one of the investments does pay off, an immediate rebalance is a great idea. Otherwise, almost all of the capital – $1M out of the $1.000099M, to be precise – remains in the successful investment, and the diversification you started out with has all but vanished. To keep the compound return high, not to mention actually get any compounding at that level of return, it is necessary to rebalance, redistributing most of the capital away from the previously successful investment to the other investments.

It is easy to dismiss rebalancing as simply a tool to preserve diversification, but the relationship between the two is almost symbiotic: you can't have one without the other. Without rebalancing, a portfolio generally tends to become more concentrated in the assets that have performed well, and it is only a matter of time before the diversification benefits erode away. Furthermore, the rebalancing trades themselves have a buy-low, sell-high quality and are profitable on average: in the above example, even though you buy 99 recent losers, the one winner that you sell more than makes up for those losses. The rebalancing has locked in the gains from the diversification, which otherwise might be lost.

Does diversification always improve the compound return? In general, it depends on how diversification is measured. Diversification as a tool for reducing risk has been studied for decades and several ways to measure it have been proposed over the years [Frahm and Wiechers; Choueifaty and Coignard], but only one method provides a direct measurement of the associated increase in compound return. This measure depends on the *difference between* (and not the *ratio of*) the average asset variance and the portfolio variance. To see why this is so, consider the well-known approximate relationship between expected return and compound return for a portfolio:

$$\text{compound return of portfolio} \approx \text{expected return of portfolio} - \frac{1}{2}\,\text{variance of portfolio} \qquad (1)$$

[2] Whenever the term appears in this article, 'compound return' refers to the long-term compound return, in the limit as the number of periods tends to infinity. In many applications, for example when the returns are statistically independent, this equals exp(E(log(1+R))) – 1. In the example, compound return = exp(0.99 log(0.1) + 0.01 log(100000)) – 1 ≈ −88.5%.

[3] This follows from observing that the probability that none of the investments pay off is $0.99^{100} \approx 36.6\%$.

Half the variance of the portfolio is subtracted from the portfolio's expected return in order to arrive at the portfolio's compound return. As the variance of the portfolio appears to detract from the portfolio's compound return, the false conclusion is often drawn that if you reduce the portfolio's variance you will always increase the compound return. This is only true if the reduction in volatility does not also commensurately decrease the expected return of the portfolio.[4]

In the highly unlikely case where all the assets in a portfolio have identical expected returns, the expected return of the portfolio will indeed also be the same (in the absence of trading costs) regardless of the way you structure it. In this particular case, then, it is true that minimizing the variance will also give the highest compound return. It is easy, therefore, to make the mistake of concluding that diversification 'only' improves the compound return because volatility is decreased [Cuthbertson et al.]. But if the expected returns of the individual assets in the portfolio are different, as is generally the case in the real world, one must apply the above relationship between compound and expected return individually to each asset:

(2)

$$\text{compound return of asset} \approx \text{expected return of asset} - \frac{1}{2}\,\text{variance of asset}$$

Since the expected return of the portfolio is simply the weighted average of the expected returns of the assets in the portfolio, it follows that[5] there is a discrepancy between the *compound* return of a portfolio and the weighted average of the *compound* returns of the assets in the portfolio. This gap equals one half the difference between the average asset variance and the portfolio variance, and it represents the boost to the portfolio compound return arising from diversification:

$$\text{increase in portfolio compound return} \approx \frac{1}{2}\left(\text{weighted average asset variance} - \text{portfolio variance}\right)$$

No other formula for diversification shows the return benefit this directly and generally. This formula also demonstrates that when the assets do not all have identical expected returns, minimizing the portfolio variance is not the optimal way to improve diversification: you have to take the individual asset variance into account as well.[6]

It immediately follows that, if the variances of the assets or the portfolio are changing over time, as they inevitably do in the real world, then the diversification benefit will also change. It would naturally make sense to respond to the changing

It is easy to make the mistake of concluding that diversification 'only' improves the compound return because volatility is decreased.

[4] In practice, reducing portfolio volatility without impacting expected return often requires portfolio optimization, which generally results in an increase in diversification.

[5] First, take weighted averages across all the assets in equation (2), then compare with equation (1). The first term on the right-hand side of each equation is the same (expected portfolio return equals weighted average expected return of assets). However, the second term is not, since the portfolio variance is less than the weighted average asset variance, if no short positions are allowed.

[6] If all the assets have identically distributed returns, then the weighted average asset variance is the same no matter what the weighting. In this special case, the only way to increase diversification is evidently to reduce the variance.

landscape by modifying the asset weighting policy to take this into account. In that case, the trading necessary to maintain the portfolio is not strictly speaking 'rebalancing,' since the targets are not static, but are varying. Nonetheless, the terminology is still appropriate since the trading has the same character as traditional rebalancing: on average, you are systematically selling recent winners and reinvesting the profits in recent losers. This has the effect of locking in previous gains and keeping the portfolio more diversified, using the volatility of the individual assets within the portfolio to continue to push up the compound return of the portfolio as a whole.

3. Individual-asset mean-reversion is not necessary

The above formulas are not always intuitive, even if the notion of buying low and selling high does sound reasonable. For this reason, it is common to demonstrate the effect of rebalancing by imagining two risky assets that are perfectly negatively correlated. For example, suppose that each asset has an even chance of rising by 25% or declining by 20% in a year, but due to the negative correlation, if one of the assets rises, then the other one falls. If asset A rises in the first year and falls in the second year, and asset B does the opposite, then the two-year return is zero for each of the two assets as well as for any buy and- hold portfolio of the two assets. Figure 2 (a) depicts a buy and- hold portfolio that is initially equal-weighted, with $100 in each asset.

Rebalancing to equal weights produces a different situation: by selling the recent winner and investing in the recent loser, the portfolio returns 2.5% per year, as is depicted in Figure 2(b). In fact, so long as the portfolio is always rebalanced back to equal weights, then no matter which of asset A and asset B happens to rise, the portfolio always returns 2.5% a year.

Even though Figures 2(a) and 2(b) appear to show mean-reverting assets, the certain 2.5% annual return in the case of rebalancing to equal weights does not depend on this. This is clear in the case shown in Figure 3(b): if asset A rises twice and

FIGURE 2
PERFORMANCE OF PORTFOLIO, MEAN-REVERTING CASE*:
(a) buy-and-hold portfolio performance (b) rebalanced portfolio performance

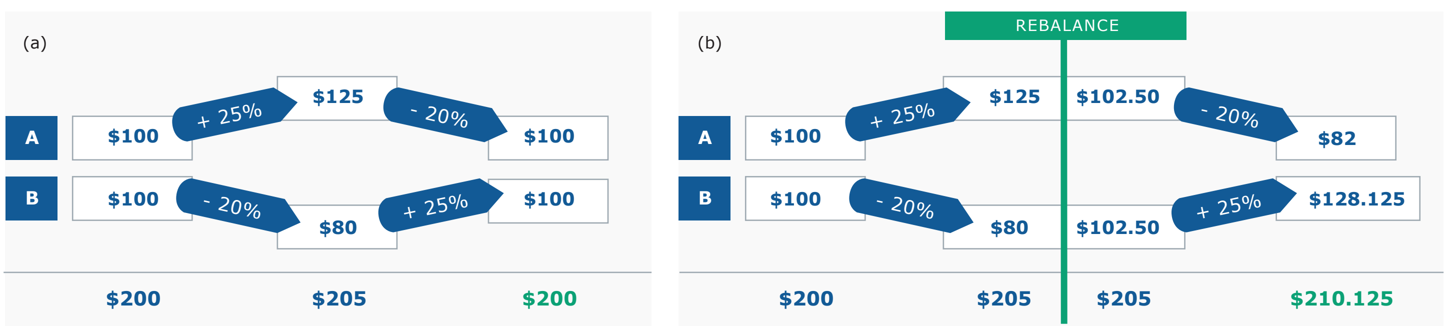

FIGURE 3
PERFORMANCE OF PORTFOLIO, TRENDING CASE*:
(a) buy-and-hold portfolio performance (b) rebalanced portfolio performance

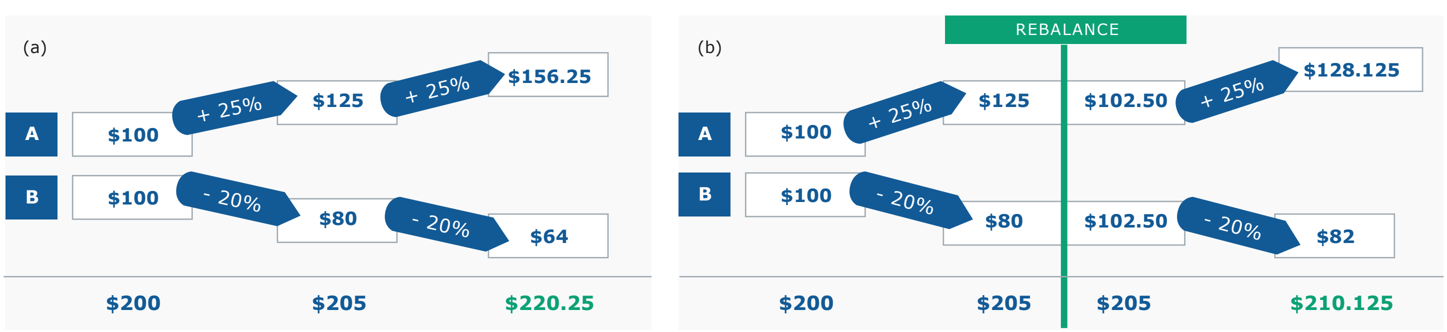

*Hypothetical illustration. Trading costs and other expenses have not been considered.

asset B falls twice, the portfolio value still rises from $205 to $210.125 at the second step: a return of 2.5% a year. This case may be thought of as the 'trending,' or non-mean reverting, case.

How about the non-rebalanced (i.e., buy-and-hold) portfolio? To understand this, it is necessary to look at the two cases separately. In the reversion case (Figure 2), the $200 initial investment ends up at $200, so the rebalanced portfolio is superior. In the trending case (Figure 3), the buy-and-hold portfolio ends up at $220.25, outperforming the rebalanced portfolio. Since these two outcomes are equally likely in this example, and the average is the same whether or not one rebalances,[7] have we gained anything by rebalancing, other than reducing the volatility by narrowing the outcomes from two to one?

Yes! The compound return of the rebalanced portfolio has improved relative to the buy-and-hold one. In the buy-and-hold case, the compound return in the second year is 2.38%, as opposed to 2.5% for the first year.[8] This comes from the drop in diversification, which inevitably follows the failure to rebalance. If this continues for 100 years, the expected terminal wealth is $2,362.74, but the realized wealth is less than this more than 86.6% of the time. The compound return for the entire period is only about 1.24%, compared to a steady 2.5% for the rebalanced portfolio.

This analysis did not assume any mean-reversion of asset prices. Every possibility was included, including the highly unlikely 'ultimate trending' case where one of the assets rises each year for 100 years.[9] The expected terminal wealth is rarely achieved, but when it is, the actual payoff is often huge due to the concentration of wealth in one of the assets. It is true that systematic rebalancing would have prevented this rare opportunity from transpiring. It is worth noting that if one excludes outcomes where one of the two stocks ends up with a market weight greater than 99.99%, the expected terminal wealth of the buy-and-hold portfolio drops to $1,122.18, while the rebalanced portfolio still ends up at $2,362.74. Once again, the notion of expected terminal wealth fails to describe the likely outcome.

Any two-asset example should be thought of as illustrative only. One may retrospectively find a case of two stocks that exhibit mean-reverting behavior [Bouchey et al.], or alternatively a case where one stock dominates the other; these examples cannot possibly prove or disprove anything about the diversification/rebalancing effect any more than a run of five heads in consecutive coin flips 'proves' that the coin is biased.

Real-world portfolios – especially stock portfolios – are usually comprised of many more than two assets. The correlations and variances of these assets will vary over time, but the benefits of diversification and rebalancing persist. If there is extreme concentration in a handful of assets that ends up dominating the others, a buy-and-hold approach could well have worked out better in retrospect. In other circumstances, where no group of assets dominates for long periods, even if none of the individual asset prices are mean-reverting, the collective stability of the market will tend to favor the rebalancing approach.[10] While analysis of this sort of stability is beyond the scope of this article, there is a rich body of literature that indicates that equity markets generally do behave in a stable manner, where no one stock or group of stocks persistently dominates, even though most of the stock prices are not individually mean-reverting. (For example, see [Ichiba et al.].) These are, in fact, ideal circumstances for a rebalancing approach to deliver results.

[7] Rebalancing always produces $210.125 after the second step; the buy-and-hold either produces $200 or $220.25 with equal probability; the average of $200 and $220.25 is indeed $210.125.

[8] The computation is straightforward: since the first period ends with $205, and the second period ends with an equal chance of $200 or $220.25, compound return = $\exp(½ \log(200/205) + ½ \log(220.25/205)) - 1 \approx 2.38\%$.

[9] A corroborating spreadsheet is available upon request from the authors. The experiment can also be analyzed using Monte Carlo simulations, but this is unnecessary since an explicit mathematical analysis is straightforward.

[10] Care must be taken to estimate transaction costs for the scale of assets being traded, then select a trading strategy so that these costs do not overwhelm the expected benefit.

Conclusion

Compound returns are a vital attribute for any investment whose term length is not known in advance, since they are more reflective of *typical* outcomes.

Even if compound returns for each asset are hard to predict in advance, maintaining a diversified portfolio by means of rebalancing increases the compound return of the portfolio as a whole relative to a buy-and-hold portfolio. The formula for this benefit, originally identified by our founder, Dr. E. Robert Fernholz in 1982, specifies the amount of this increased return to the extent that volatilities can be estimated, gross of transaction costs. Rebalancing and diversification are inextricably linked concepts. Mean-reversion of individual asset prices is absolutely not required for this phenomenon to hold, and stability of the overall market structure further reduces the likelihood that a buy-and-hold approach will outperform the diversified/rebalanced approach in the long term.

References

Bouchey, P., Nemtchinov, V., Paulsen, A. and Stein, D.M. (2012) Volatility Harvesting: Why Does Diversifying and Rebalancing Create Portfolio Growth? The Journal of Wealth Management 15(2): 26—35.

Chambers, D.R. and Zdanowicz, J.S. (2014) The Limitations of Diversification Return. The Journal of Portfolio Management Summer 40(4): 65—76.

Choueifaty, Y. and Coignard, Y. (2008) Toward Maximum Diversification. The Journal of Portfolio Management, 34(4): 40—51.

Cuthbertson, K., Hayley, S., Motson, N. and Nitzsche, D. (2015) Diversification Returns, Rebalancing Returns and Volatility Pumping. Available at SSRN: http://ssrn.com/abstract =2311240 or http://dx.doi.org/10.2139/ssrn.2311240

Fernholz, R. and Shay, B. (1982) Stochastic Portfolio Theory and Stock Market Equilibrium. The Journal of Finance, 37: 615—624.

Frahm, G. and Wiechers, C. (2011) On the diversification of portfolios of risky assets. Discussion papers in statistics and econometrics, No. 2/11. Ichiba, T., Papathanakos, V., Banner, A., Karatzas, I. and Fernholz, R. (2011) Hybrid Atlas models. Annals of Applied Probability, 21: 609—644.

Qian, E. E. (2014) To Rebalance or Not to Rebalance: A Statistical Comparison of Terminal Wealth of Fixed-Weight and Buy-and-Hold Portfolios. Available at SSRN: http://ssrn.com/abstract =2402679 or http://dx.doi.org/10.2139/ssrn.240267

Intech® is a specialized global asset management firm that harnesses stock price volatility as a source of excess return and a key to risk control. Founded in 1987 in Princeton, NJ by pioneering mathematician Dr. E. Robert Fernholz, Intech® serves institutional investors across five continents, delivering relative return, low volatility, adaptive volatility and absolute return investment solutions.

Intech Offices

HEADQUARTERS
250 S. Australian Ave., Suite 1800
West Palm Beach, FL 33401
United States of America
+1-561-775-1100

RESEARCH & DEVELOPMENT
One Palmer Square, Suite 441
Princeton, NJ 08542
United States of America
+1-609-497-0443

INTERNATIONAL DIVISION
201 Bishopsgate
London, EC2M 3AE
United Kingdom
+44-20-7818-5600

intechinvestments.com

Intech® — JANUS HENDERSON —