ADVISOR PERSPECTIVES

The Academic Failure to Understand Rebalancing

August 16, 2016

by Michael Edesess

Perhaps the most universally accepted investing principle is to periodically rebalance one's portfolio. Advisors have been drilled that rebalancing results in some combination of improved performance and reduced risk. Unfortunately, this precept is the byproduct of imperfect mathematics; the benefits of rebalancing are far smaller than what advisors have come to believe.

The best technical article on portfolio rebalancing – by a very wide margin – was not published in a finance journal. It was written by A. J. Wise and published in the British Actuarial Journal in 1996. Meanwhile, dozens of articles on rebalancing in finance journals are filled with mathematics that is riddled with obvious flaws. The conclusions most of these articles draw from their mathematics are not only trivial but also so vaguely stated that it is difficult to discern exactly what the conclusions are and what their relationship is to the practical issue of portfolio rebalancing.

Why does academic finance fill its articles with mathematical sound and fury that signifies so little? What is wrong with the field of "mathematical finance" that it can produce such inferior work, published in peer-reviewed journals, with so little relevance for anyone who wishes to know whether portfolio rebalancing provides a benefit?

I will explore these questions after discussing the results of the academic finance papers and comparing them with Wise's results.

The simple results of the finance papers

It is challenging to examine the academic finance papers on rebalancing[1] in a non-technical publication, given the arcane nature and impenetrability of those papers' mathematics and language. Therefore, I will paraphrase them in a simple way. Rest assured that this paraphrased version is completely analogous to the analyses in the papers themselves.

Let's say someone offers you a bet on the toss of a fair coin (50-50 chance of heads or tails). If it comes up heads, your bet will be doubled (100% return); if tails, your bet will be halved (–50% return). In other words, a dollar bet on the coin toss will become two dollars if it comes up heads, but only 50 cents if it comes up tails.

Suppose you bet a dollar on the first coin toss and let it ride. If the coin is tossed some number of times – say, n times – and it comes up heads m of those times, a little easy math shows that your initial dollar will become two raised to the m power times one half raised to the (n – m) power dollars (that is, $\$2^m 0.5^{n-m}$). If there were 10 tosses in total and m was six (six heads), your initial dollar will become $\$2^6 0.5^4$, or four dollars.

But if the 10 tosses produced only five heads, you'll end up with the same one dollar you started with, for a zero rate of return.

In the long run, the law of large numbers says that the proportion of heads will gradually become one half. So in the long run, m will come closer and closer to half of n. That means that the longer you play the game, the more likely it is that you'll just wind up with the same dollar you started with and realize only a zero rate of return.

Kelly's trick

However, you can "trick" this game, using a strategy devised by John Kelly in 1956. The first time you play you bet only half your dollar and keep the other 50 cents in reserve. Win or lose, you play the same strategy each time – you "rebalance" your portfolio. Each time you play, you place 50% of your portfolio on the bet and keep the other 50% in cash. For example, if you win the first toss you'll have \$0.50+0.50x2 = \$1.50. So on the second toss you'll put 75 cents in reserve and bet the other 75 cents.

It can now be shown using easy math that after n tosses, m of which are heads, your portfolio will have a number of dollars equal to $3^n/2^{2n-m}$. In the long run, m will get closer and closer to one half of n. So this number will converge to $\$3^n/2^{1.5n}$, which is much greater than one dollar.

For example, if you play 20 times and m, the number of heads, is the expected 10, then you'll have $\$3^{20}/2^{1.5x20}$, which is $3.25 – much more than the $1 you would have had if you had bet 100% of your portfolio every time. It's also greater than what you would have had if you had initially bet only half of your dollar and then let it ride – a 50/50 buy-and-hold strategy. That strategy would have yielded, after 20 tosses, half of which were heads, an end result of $1 also.

It would seem that the strategy of rebalancing on every toss to a 50-50 cash/bet mix outperforms both the 100% bet strategy and the 50-50 buy-and-hold strategy.

As we shall see, however, this is woefully incomplete; yet it is pretty much all that the finance articles say. Virtually all of the most-cited finance papers on the subject bury this simple result under layers of stochastic differential calculus, accompanied by suggestive but inadequately defined terminology such as "diversification return," "volatility pumping" and "excess return." Then they claim or imply, based on this result alone, that rebalancing is the superior strategy.

But it is only a small part of the picture

There's much, much more to the issue than this result. Rebalancing is not the free lunch it appears to be. Using the Kelly criterion gives up a lot of the upside.

First of all, m may become half of n in the "long run" – that is, as n goes to infinity – but it does nothing of the kind in any finite series of tosses. On any (even) number n of tosses, it is unlikely that m will be exactly half of n. And as n grows larger, it becomes more and more unlikely.[2]

That makes a big difference. For example, for 20 tosses of the coin, if it comes up heads exactly half the time, as I've shown, you'll have $3.25 with the rebalancing strategy but only $1 with either the 100% risk strategy or the 50-50 buy-and-hold strategy. The probability this will happen – that the coin will come up heads exactly half the time – is only 17.6%.

But if the coin comes up not exactly half of n – not exactly 10 times – but 11 or more times, the picture changes radically. The probability of 11 or more heads is 41.2%, more than twice the probability of exactly 10 heads. But if the coin comes up heads 11 or more times, the average ending value with the rebalancing strategy will be only $23 while the average ending value with the 50-50 buy-and-hold strategy will be $105 – more than four and a half times as much – and with the 100% buy-and-hold strategy it will be $210.

This advantage to buy-and-hold on the upside is not balanced by a disadvantage at the other end. On average for all possible numbers of heads m, the ending value with rebalancing will be $11, while with 50-50 buy-and-hold it will be $43, and for 100% buy-and-hold it is $87. The rebalancing strategy is, in this particular example, inferior on average to the other strategies, even though it is superior if the median result of 10 heads occurs.

Hence, if you were to play the game 1,001 times, you would end up with $87,000 with the 100% buy-and-hold strategy, $43,000 with the 50-50 buy-and-hold strategy, but only $11,000 with the rebalancing strategy. However, your median result – the 501st ranked result out of the 1,001 – would have been only $1 with either 100% buy-and-hold or 50-50 buy-and-hold, but $3.25 with rebalancing.

What the finance papers have done, in essence, in a very opaque way, is to focus only on the medians of the probability distributions of the rebalancing and buy-and-hold strategies – the case of m equals half of n. The result of m being equal to half of n is the median result. But it is only one point in the distribution. The finance articles give little if any attention to the rest of it, and yet they proclaim – or at least imply by using suggestive terminology like "excess return" – that rebalancing produces a premium return over buy-and-hold. This is true only at the median; it is not true at the average.

The analyses and results of A. J. Wise

The Wise article is by far the most mathematically sophisticated and well-specified of all the journal articles I have found. It produces the most meaningful practical results as well. And its inferences about rebalancing of real portfolios are in complete agreement with the conclusions of its own math, unlike, believe it or not, some of the finance articles.

Yet in the many articles on rebalancing published in finance journals, I was able to find only two citations of the Wise article, and those citations are very much in passing. One merely notes that the CRP (continuously rebalanced portfolio) formula appeared in Wise; the other, in a footnote, notes that Wise (1996) discusses dynamic strategies and the intuition behind them. None cite the rich mathematics and comprehensive results found in Wise.

Wise focused not only on the median of the distribution of returns with rebalancing or buy-and-hold, nor on any other single summary statistic, but on the whole distribution. By means of a piece of mathematics that made me green with envy when I read it, he deduced that a rebalancing strategy will beat a buy-and-hold strategy about two thirds of the time when the constituent assets in the portfolio have the same expected return. When buy-and-hold beats rebalancing, however, it beats it by a much larger margin, so that the average returns to rebalancing and to buy-and-hold, in the equal expected returns case, are the same.

I reached the same conclusions in a 2014 article, but not with Wise's mathematical elegance. Rather, I observed them from simulations. (I had not yet discovered Wise's article at the time.) I then confirmed this conclusion using historical stock price data, as did Wise. If the expected returns on the constituent assets are not the same, then of course the probability that buy-and-hold will beat rebalancing becomes greater. Wise's formulas provide a convenient way to assess that probability.

Wise then went on to derive an interesting formula for the probability that buy-and-hold will beat rebalancing in the very long run – that is, as time goes to infinity. The result is that if the average difference between constituent portfolio assets' expected returns is small, less than about one percent, then rebalancing will, in the very long run – that is, at eternity – surely beat buy-and-hold; but if the average difference is greater than that, buy-and-hold will surely beat rebalancing.

Median versus average

The results of the finance papers are merely a trivial corollary of Wise's work, and of mine. If rebalancing beats buy-and-hold (by a small margin) two thirds of the time, that implies that the median result with rebalancing is greater than the median result with buy-and-hold. This is the simple result that the finance articles derive, though it is buried under all of their (wholly unnecessary) stochastic differential calculus and opaque terminology. While in the case of equal expected returns on constituent assets the median for rebalancing will be greater than the median for buy-and-hold, the average returns with rebalancing and with buy-and-hold will be equal.

The results stated in the previous paragraph apply only if the assets in the portfolio all have the same expected return. This is not the usual situation. When the constituent assets do not have the same expected rate of return, the results with a buy-and-hold strategy will, on average, beat the results with a rebalancing strategy. And if the expected returns on the constituent assets are different enough, the median return with a buy-and-hold strategy will be greater than the median return with rebalancing as well.

Why are the finance articles so much less revealing?

It is a sad commentary, in my view, that the finance articles take so little note of the Wise article, which produces in relatively few pages and concise mathematics so much more in terms of valid and useful results than all of the finance articles put together. This may be a measure of how isolated the academic finance field is from external critical review – and from any ability to take advantage of work outside its immediate profession. This isolation may also account for the self-congratulatory attitudes (evidenced by a seemingly blind willingness to cite references to one another's papers[3]) of many in academic finance and its adherents and disciples.

The saddest thing is that this profession has, by its arcane mathematics and language and smug self-congratulations, insulated itself from proper scrutiny and thus convinced so many other practitioners that it is indeed sophisticated and that it performs miracles with investment results. As we know so well from the overall statistics, the latter is patently untrue.

Since getting involved with the financial industry years ago, I have found that it most often uses mathematics not as a means to solve problems or derive answers, the way it is used in the sciences and technology, but as a sales tool, to make it appear that it was used to solve problems or derive answers – and to impress and baffle. This is not necessarily done deliberately; it is a habit that is easy to slide into, if one joins the profession and attempts to mimic one's peers.

This sets a relatively low bar for achievement – and that achievement is very richly rewarded. All that is necessary is to make the mathematics look, to someone who doesn't examine it with a great deal of care, as though it can be and is used to solve problems and derive answers. If that criterion is met, you can get published in a top peer-reviewed journal, whether or not the math actually can be so used, and whether or not it is well-specified and produces the conclusions that are stated or implied in the article.

All that is really necessary is that it look like it can.

The bottom line

The bottom line can be readily gleaned from Wise's article and my two previous articles on rebalancing.[4] Contrary to common belief and to the misguided conclusions of most of the articles in academic finance journals, it is that rebalancing offers no “free lunch,” either in terms of enhanced return or reduced risk. The choice of rebalancing as an investment discipline as compared with an alternative such as buy-and-hold is simply a risk-return tradeoff – though one that is a little more subtle than most.[5] At the more common middling levels of performance of portfolio assets, or when portfolio assets perform similarly, rebalancing will slightly outperform buy-and-hold. But in less common cases (about one in three) when one or more of the assets performs very well, buy-and-hold will markedly outperform rebalancing. This choice may not be easy to make; either of the choices may be regarded as satisfactory. But it is a far cry from saying that it has been proven that rebalancing is something you absolutely must do.

Michael Edesess, a mathematician and economist, a senior research fellow with the Centre for Systems Informatics Engineering at City University of Hong Kong, chief investment strategist of Compendium Finance and a research associate at EDHEC-Risk Institute. In 2007, he authored a book about the investment services industry titled The Big Investment Lie, *published by Berrett-Koehler. His new book,* The Three Simple Rules of Investing, *co-authored with Kwok L. Tsui, Carol Fabbri and George Peacock, was published by Berrett-Koehler in June 2014.*

PARTIAL LIST OF REFERENCES

Booth, David G., and Fama, Eugene F. "Diversification returns and asset contributions." *Financial Analysts Journal* 48.3 (1992).

Cuthbertson, Keith, et al. "What Does Rebalancing Really Achieve?." *International Journal of Finance & Economics* (2016).

Dempster, Michael AH, Igor V. Evstigneev, and Klaus R. Schenk-Hoppé. "Volatility-induced financial growth." *Quantitative Finance* 7.2 (2007): 151-160.

Fernholz, E. Robert. "Stochastic portfolio theory." *Stochastic Portfolio Theory*. Springer New York, 2002. 1-24.

Fernholz, Robert, and Brian Shay. "Stochastic portfolio theory and stock market equilibrium." *The Journal of Finance* 37.2 (1982): 615-624.

Kuhn, Daniel, and David G. Luenberger. "Analysis of the rebalancing frequency in log-optimal portfolio selection." *Quantitative Finance* 10.2 (2010): 221-234.

Luenberger, David. "Investment science: International edition." *OUP Catalogue* (2009).

Mulvey, John M., and Woo Chang Kim. "Multistage financial planning models: Integrating stochastic programs and policy simulators." *Stochastic Programming*. Springer New York, 2010. 257-275.

Mulvey, J., and W. Kim. *Constantly rebalanced portfolios-is mean-reverting necessary*. Technical report, working paper Princeton University, Princeton, NJ, USA, 2008.

Wise, A. J. "The investment return from a portfolio with a dynamic rebalancing policy." *British Actuarial Journal* 2.04 (1996): 975-1001.

[1] See partial list of references at the end of this article. A more complete set of references can be found in Cuthbertson et al (2016). I am grateful to Joe Tomlinson for a comment on APViewpoint that led me to explore this body of work.

[2] In fact, a previous statement, “The longer you play the game, the more likely it is that you’ll just wind up with the same dollar you started with,” is not strictly true. It is an example of how a seemingly correct conclusion can be faulty if it is not checked rigorously. It is also intrinsic to the flaws in the academic articles on rebalancing.

[3] For example Mulvey and Kim (2008) state, citing Fernholz (1982), “we describe the performance advantages of the fixed-mix rule over a static, buy- and-hold perspective. This rule generates greater return than the static model by means of rebalancing.” This statement is far more than Fernholz proved. Fernholz’s choice of terminology, however, such as labeling a particular arbitrary difference between two parameters “excess growth,” apparently led Mulvey and Kim to their conclusion.

[4] Does Rebalancing Really Pay Off?? and Does Rebalancing Reduce Risk?

[5] For those comfortable with comparing probability distributions, the tradeoff is depicted in Figure 1 of my article Does Rebalancing Reduce Risk?.

