ADVISOR PERSPECTIVES

The Dangers of Rebalancing

Michael Edesess
December 7, 2010

Every portfolio should be rebalanced to its targeted asset allocation, we are taught. Indeed, there may be no other precept as routinely and studiously practiced among financial advisors. But does rebalancing either increase expected return or reduce risk? If so, why?

The answers to those questions reveal that it may be prudent to rebalance, but not for the reasons you think.

In the investment field – as in almost no other – it is well known that you can never assume historical patterns will continue. This is because of the arbitrage principle: if prices follow a predictable pattern, investors will arbitrage that pattern out of existence.

Hence, when an investment strategy seems a good idea for whatever reason, somebody often trots out a claim that there's a mathematically-certain reason why it works. You need the mathematical reason to make it a law of nature, immunizing it against the vicissitudes of history and data. Knowing only that it worked in the past, or that it seems like a good idea, isn't enough. It may not work in the future. The advocate of an investment strategy therefore has strong motivation to say – and usually to believe – that there's a mathematical reason why it works, even if there isn't one.

For example, people have advanced mathematical arguments why dollar-cost averaging and so-called fundamental indexing must work. The mathematical arguments are wrong – usually because the claimant hasn't taken the trouble to state clearly what criteria would establish that the strategy "works."

Let's take the case of rebalancing. It is widely believed that there is a mathematical reason why rebalancing should work. This belief is in error – at least, it is in error under all the mathematical models of securities prices that are generally accepted and applied by both the academic and practical investment communities.

But there do seem to be reasons to believe that rebalancing will have benefits. Those reasons are so far only empirical, and more in the nature of anecdotal evidence than clear-cut statistical evidence, and therefore may not be reliable in the future. Nevertheless those reasons point – as do so many things since the financial crisis – toward the need for an improved theory of financial markets, one that takes into account the likelihood and pattern of financial bubbles and crises and derives from it amendments to our models of securities price movements.

The case of rebalancing

Rebalancing began as an afterthought, then became conventional wisdom in no time flat.

The models on which "modern portfolio theory" or MPT were built were – and are – *one-period* models. They apply to one of many decisions in an investor's lifetime, for only a short period of time, such as a year. They ignore all the decisions the investor might make and objectives the investor might have after that.

This may suggest to you that the models have limited value – and you'd be right. The models did, however, at least force the investing world to recognize two principles: (1) it's better to diversify; and (2) you can't get more return without more risk.

Harry Markowitz's mean-variance algorithm appeared to show a portfolio manager how to "optimize" a portfolio's asset mix. But because the model is so sensitive to inputs that they have to be jerry-rigged to produce reasonable outputs, that optimization, even for a single period, doesn't actually work very well in practice. Even if it did, there would be the problem of what to do in the next period, and why. That was taken care of by simply assuming – as an afterthought – that you apply the same optimal asset mix again and again, to every period. Hence, you need to rebalance, because by the time you get to the end of the period your asset mix may not be the same as the one you decided on at the beginning.

This afterthought was fueled by the assumption that the standard deviation of returns over short time periods was an investor's measure of risk. As we'll soon see, that assumption can be severely mistaken. But once that assumption was made, it was natural to assume that an investor wanted to keep his or her risk constant; that was done by rebalancing periodically to the original mix.

How these assumptions can go wrong

But what if your risk measure is the chance of unacceptably low cash flows in 30 years? And what if you control that risk by holding part of your portfolio in 30-year Treasury bonds to provide a secure safety net?

If you rebalance periodically you won't maintain your risk hedge, you'll destroy it. Suppose, for example, you start out with 50% stocks and 50% Treasury bonds, which you will hold to maturity. Suppose the market drops 20%. Rebalancing means you'll have to sell some of the Treasury bonds that provide your safety net.

The problem with a rebalancing strategy is that it is short-sighted, and focuses on the wrong risk measure. Short-term volatility of asset values or returns is certainly a component of risk – but it is far from all of it. For most investors, risk is a long-term, not a short-term, concern. Too much emphasis on rebalancing may divert attention from a true

long-term risk control strategy, such as, for example, one which focuses on protecting against inflation by buying and holding TIPS to maturity.

The mathematical facts about rebalancing under the "usual" mathematical assumptions

Virtually all mathematical derivations and simulations of securities price movements are made under the assumptions of what Eugene Fama called the "weak" random walk. The chief assumption is that securities price movements are independent of each other in non-overlapping time periods. Whatever the price change was last period adds nothing to your knowledge of the probability of price changes in the next time period. The details of the model can vary, but this assumption remains paramount.

Under this assumption there is no expected return benefit to rebalancing. In general, its expected return will be worse than a buy-and-hold strategy if the portfolio assets have different expected returns, and it will be the same as a buy-and-hold strategy if they have the same expected returns.

Another view

Having said that, I take time to address a different view, because respectable parties have voiced it. The view hinges on the observation that – as Roger Gibson, author of "Asset Allocation" states it – "the compound annual return is approximately equal to the simple average return minus one-half of the variance."

This well-known fact seems to be invoked to argue that rebalancing, because it yields a return that is more like a simple average, must be superior to – say – buy-and-hold, because the latter will yield a return that is more like a compound annual return, which will be half a variance less than the simple average.

Unfortunately this argument has not been made in any sort of rigorous way. If it were, I'll wager that it can't be shown to prove anything. The problem is that the simple average of a time series of annual returns is a statistic without a counterpart in the real world. It is not the growth rate of wealth; that is expressed by the compound return. Time series of annual rates of return compound, they don't add. (Unless you're talking about logarithmic returns, but that wouldn't improve the argument.) So why would one want to add them, or take their average? This argument doesn't answer the question: OK, you claim rebalanced returns are better – *but compared to what?*

Empirical evidence

Hence, theoretically – given the usual model assumptions – there is no reason why rebalancing should enhance expected return. But it is always important to look at the data

– even if not to be fooled by it – as it may tell a different story. There appear to be an unusual number of cases in which rebalancing strategies do better than alternatives.

In Gibson's book, as I referenced in a previous article, he exhibits returns over the years 1972-2005 for four risky assets: U.S. stocks, non-U.S. stocks, real estate securities, and commodity-linked securities. There are 11 possible combinations of two or more of the four assets. In five of these 11 combinations, the rebalancing strategy's return was superior to *all* of the compounded returns of the combination's component assets.

This seemed a highly unlikely outcome under the null hypothesis – that is, under the usual mathematical assumptions. So I did a simulation to see how unlikely it was. In fact, it would randomly occur a little less than ten percent of the time under the usual assumptions. (The same simulation showed – as the theoretical mathematics would indicate – that the expected return of the rebalancing strategy was no greater than the expected return of a buy-and-hold strategy.)

It might be easy to dismiss this as a fluke of the data, but William Bernstein apparently observed a similar "rebalancing bonus" in other data in 1997.[1] Furthermore, in his book "Unconventional Success," David Swensen, the celebrated manager of Yale's endowment fund, says:

> As a matter of course, every trading day, Yale estimates the value of each of the components of the endowment. When marketable securities asset classes (domestic equity, foreign developed equity, emerging market equity, and fixed income) deviate from target allocations, the university's investment office takes steps to restore allocations to target levels. In fiscal year 2003, Yale executed approximately $3.8 billion in rebalancing trades, roughly evenly split between purchases and sales. Net profit from rebalancing amounted to approximately $26 million, representing 1.6% return on the 1.6 billion equity portfolio.[2]

We do not know why Yale rebalances or how they do it, or how they measure the benefits, but I can only assume that Swensen has accurately reported them.

Why would there be a benefit to rebalancing?

I have already pointed out that the usual mathematical assumption of independence of rates of return across time periods implies that there is no expected return benefit to rebalancing. Hence, if there is an expected return benefit – as Gibson's, and Bernstein's, and Swensen's data suggest there might be – then the assumption of independence of rates of return across time periods must be violated.

---

[1] "Diversification, Rebalancing, and the Geometric Mean Frontier", by William J. Bernstein and David Wilkinson, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=53503.

[2] David Swensen, *Unconventional Success*, p. 198.

It's certainly not a stretch to believe this. Stock market plunges are often followed by sharp recoveries. Time periods like 1973-1976, 2002-2003, and 2008-2009 illustrate this. The existence of such patterns would appear to violate the arbitrage principle, but it seems that even though investors know that sharp recoveries often follow sharp dips, they're hard-pressed to know at exactly what time they will occur.

As an example of a rebalancing strategy that worked better than a buy-and-hold strategy after such reversals, consider these data from the years 1972-1976, for two assets beginning 1972 in a 50%-50% mix:[3]

	Annual returns			Cumulative annual returns		Buy-and-hold %-ages	
Year ending	**C**	**D**	**rebal**	**buy-hold**	**diff**	**% C**	**% D**
1971	-	-	-	-	-	50.0%	50.0%
1972	8.0%	42.4%	25.2%	25.2%	0.0%	43.1%	56.9%
1973	-15.5%	75.0%	27.5%	30.5%	-3.0%	26.8%	73.2%
1974	-21.4%	39.5%	21.0%	28.0%	-7.0%	17.1%	82.9%
1975	19.3%	-17.2%	15.7%	16.9%	-1.2%	22.9%	77.1%
1976	47.6%	-11.9%	16.1%	13.7%	2.4%	33.3%	66.7%

In this table, C is Real Estate Securities and D is Commodity-Linked Securities. The whipsawing of their annual returns first increased buy-and-hold's advantage dramatically, then decreased it just as dramatically. Notice how in the buy-and-hold case the whipsawing increases the allocation to D from 50 percent to almost 83 percent at the end of 1974, just in time for its decline. The rebalancing strategy avoids this whipsawing.

A possible explanation

The examples suggest – though they don't prove – that one of the justifications often given for rebalancing may have some merit. The justification is that it takes advantage of "mean-reversions," that is, instances where some asset's return takes a swing and then swiftly reverts back again. There's no proof that predictable mean-reversions exist. The arbitrage principle should rule them out.

The arbitrage principle notwithstanding , if there really is autocorrelation – that is, interdependence of returns across time periods, such as swings predictably followed by reversions – then rebalancing's advantage is that it is unaffected by the sequence of returns, and is hence immune to them. Other strategies, such as buy-and-hold, might be whipsawed by them.

Given our recent observations of markets, there may be a connection with the question of how to understand financial bubbles and crises. Financial crises are characterized by

[3] Roger Gibson, *Asset Allocation 4th Edition*, p. 274.

sudden drops following market surges, often followed by sharp recoveries. If we could better understand, as a matter of theory, how and why these happen it could have useful implications for asset management.

More research is needed

More research is needed into whether and why rebalancing should work – especially in the context of a possible next-generation model of securities markets. Most academic research on rebalancing simply begins by assuming that rebalancing is something you want to do without explaining why. For example, a 2003 article in *The Journal of Portfolio Management* by Seth J. Masters starts by saying:[4]

> In general, investors have a target asset allocation that they seek to maintain. When a portfolio's actual allocation deviates from its target allocation, the result will be tracking error compared to the target allocation. The key benefit of rebalancing is to reduce this tracking error.

In other words, you need to rebalance in order to maintain your target allocation. But it doesn't say why you need to maintain your target allocation. True research on rebalancing should go back to square one and try to posit a reason why you should do it in the first place.

What should you do?

What an advisor should do is clear. Consider all of an investor's future goals, then design an investment strategy over the investor's entire horizon to meet them. That strategy will not necessarily keep asset allocation constant. Target date strategies, for example, deliberately change allocations over time. While target date strategies may not have been designed well or explained well to investors, they do illustrate that rebalancing to the original allocation may not always be the best strategy.

Once an advisor has designed a forward-looking strategy, that strategy should be revisited periodically to see if adjustments should be made. The adjustment may or may not be rebalancing to the original target allocation; rather, the allocation may change because the client's goals, time of life, or wealth accumulation path has changed; or because the strategy was intentionally designed to change allocations over time.

In short, an advisor should use analytic methods, not merely apply rebalancing automatically as a shibboleth.

[4] "Rebalancing" by Seth J. Masters, *Journal of Portfolio Management*, Spring 2003, pp 52-57.

Michael Edesess is an accomplished mathematician and economist with experience in the investment, energy, environment, and sustainable development fields. In 2007, he authored a book about the investment services industry titled The Big Investment Lie, published by Berrett-Koehler.

www.advisorperspectives.com

For a free subscription to the Advisor Perspectives newsletter, visit:
http://www.advisorperspectives.com/subscribers/subscribe.php