

The Fallacy behind Investor versus Fund Returns (and why DALBAR is dead wrong)

July 19, 2016

by Michael Edesess

It has become accepted, conventional wisdom that investors underperform their investments by timing those investments badly. This conclusion has been arrived at by combining (1) a desired result (for financial professionals whose compensation is contingent on accumulating assets); with (2) perceived “common sense” (perhaps inculcated by the desirability of that result); and (3) incautious, cavalier application of an inadequately-analyzed, poorly-understood mathematical measure.

As I will explain, there is no way to determine if investors underperform the mutual funds they own, either because of bad timing or for any other reason.

Why is it a desirable result for investment professionals?

A company named DALBAR has found success by publishing a report each year titled “Quantitative Analysis of Investment Behavior” (QAIB). For \$775, it allows an advisor to buy the right to use all of the QAIB’s charts and data in its marketing materials.

Of course, this report invariably shows that the investor makes serious errors and needs help. If it didn’t, DALBAR would sell no copies of its report to financial advisors and would quickly go out of business. This fact itself – together with the fact that it is difficult to find out the exact details of DALBAR’s methodology[1] – should be enough to provoke skepticism about the QAIB’s perennial result, which is that mutual fund investors underperform their investments by a large margin.

DALBAR’s recurring result has been turned into conventional wisdom, something to be drilled into investors’ minds repeatedly, in various ways, to keep them aware of their behavioral investing foibles. For example, Figure 1 below shows a widely circulated cartoon by Carl Richards, illustrating how investors repeatedly buy at market tops and sell at market bottoms.

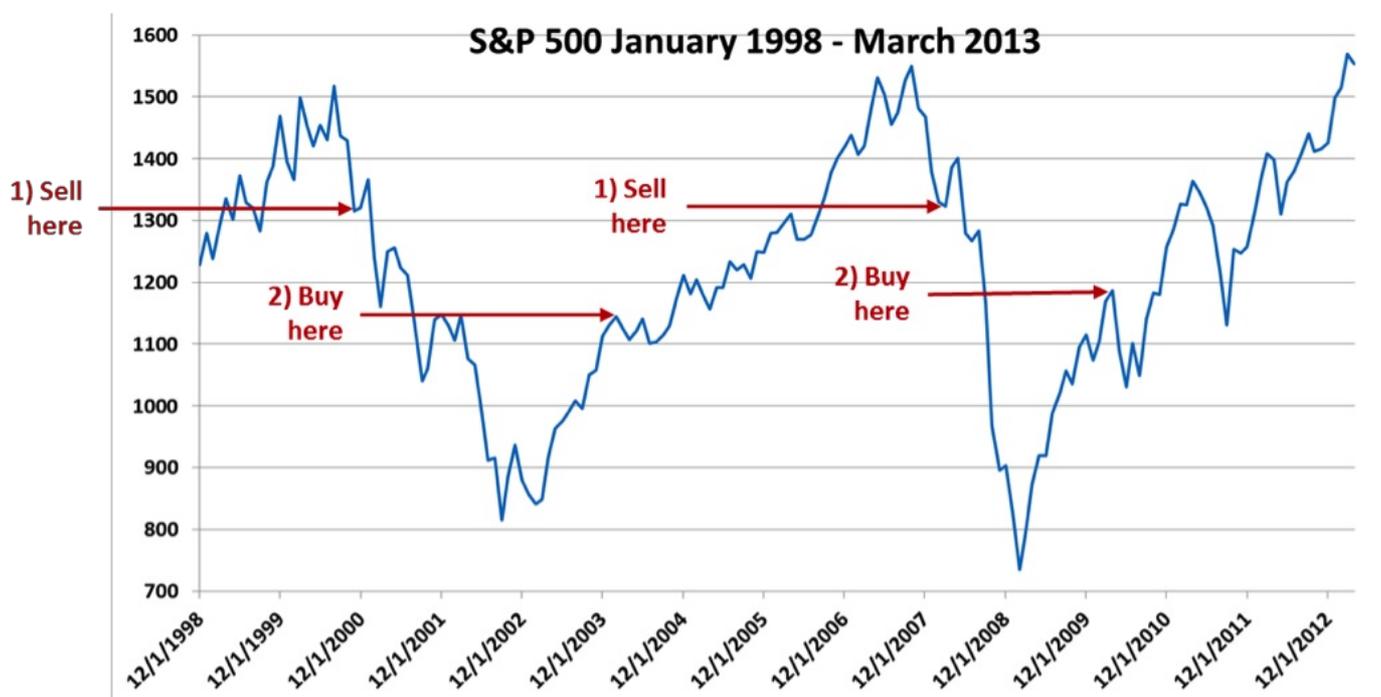
Figure 1. Carl Richards’ investor-fail diagram



Wouldn't it be wonderful to have a client who invested as reliably as that? All you'd have to do is buy when the investor sold and sell when the investor bought, and you'd earn a tremendous return.

But investors don't buy precisely at market tops and sell precisely at market bottoms. Perhaps they do sell on market panic and buy on exuberance, but that's a long way from selling at bottoms and buying at tops. Figure 2 shows my hypothetical example of a more likely scenario of panic-selling and greed-buying, one in which the buys and sells are not so perfectly market-timed but are more realistically random.

Figure 2. More realistic "irrational" buying and selling



In Figure 2, notice that the “irrational” investor sells in panic after a drop, and buys on anticipatory exuberance after a rise. In fact, in each case in this depiction, the investor sells, disadvantageously, at a local minimum and buys at a local maximum.

And yet this particular investor has in each case sold at a higher price than she bought. So in this particular arbitrarily-constructed case, our investor beat the market. Of course, the long-term investor would, in all probability, do better by not trying to time the stock market – which means being out of the market for periods of time – but by staying in the market persistently at the level of allocation chosen.

Nevertheless, the scenario depicted by Figure 1 will be much more attractive than the one in Figure 2 to financial advisors, who can use it to show their clients the mistakes they are likely to make if they are not held to a discipline by their wiser and more-experienced advisors. Hence, the value of the DALBAR report to financial and investment advisors.

The incautious, cavalier use of an inadequately-understood mathematical measure

As I have often pointed out, mathematical measures and analyses are frequently used in the investment finance field with little true depth of analysis or understanding. They are used to leap to a desired conclusion, or a conclusion that seems self-evident, but that is not in fact fully supported by the math. This failure to draw conclusions that are actually supported by the mathematics holds true even – perhaps especially – in many articles published in top financial journals.

Let us now go into some depth scrutinizing the mathematical measure that is used to show that investors underperform their investments. I shall refer to Morningstar’s methodology rather than to DALBAR’s because Morningstar’s is more transparent and less motivated by the need to sell a particular result.

Morningstar director of manager research Russel Kinnel describes Morningstar’s investor return methodology in its *Mind the Gap 2016* report (may require registration).

Kinnel says, “Generally, investor returns fall short of a fund’s stated time-weighted returns because, in the aggregate, people tend to buy after a fund has gained value and sell after it has lost value.” He adds, “That fact plays out for all types of investors – not just fund investors. Studies show stock investors and pension fund managers do the same thing.”

Aggregate investor return, he explains, is defined by the internal rate of return – also known as the dollar-weighted return – for each mutual fund, taking into account cash inflows and outflows to the fund. See Box 1 for a simple example.

Box 1. Investor Return

Here is a simple example of investor return. Suppose a mutual fund began a two-year period with \$1 billion. Suppose, for simplicity’s sake, that during the two-year period a single cash inflow of \$400 million occurred precisely at the end of the first year.

If the fund’s (time-weighted) rate of return were +10% in the first year and -10% in the second year, then it would have ended the first year worth \$1.1 billion. After the influx of the \$400 million it would have been worth \$1.5 billion. But after the second year’s -10% return it would have been worth \$1.35 billion.

The fund’s two-year time-weighted rate of return is its first year return of +10% compounded with its second year return of -10%, which comes to minus one percent for the two years, annualizing to about minus one half a percent.

The investors’ dollar-weighted return, on the other hand, is -2.1%. This is the annualized internal rate of return that can be derived from the fact that the fund began with \$1 billion, experienced an inflow of \$400 million at year-end 1, and ended year 2 with \$1.35 billion. (To see why, consider that \$1 billion invested for a year at -2.1% return is \$979 million; after adding the year-end inflow of \$400 million to get a fund value of \$1.379 billion and applying the -2.1% return again over year 2, the fund at year-end 2 will be worth \$1.35 billion.)

If instead of an inflow of \$400 million at year-end 1, there had been an outflow of \$400 million, the situation would change. Now the fund would be worth \$700 million at the end of year 1 after the outflow and worth \$630 million at the end of year 2 after the -10% return. The dollar-weighted rate of return for this scenario is plus 1.85% instead of minus 2.1% – better than the time-weighted rate instead of worse.

From the example in Box 1, the meaning of investor returns is obvious. If the cash flow of \$400 million went *into* the fund before the second year’s downturn, then the investor’s performance, her “investor return,” is

worse than the fund's performance. However, if the cash flow came *out of* the fund before the second year's downturn, then the investor's performance is better than the fund's performance.

Indeed, at a simple level the difference between the dollar-weighted rate of return and the time-weighted rate of return appears to measure how salutary the timing of cash inflows and outflows was.

But there are problems.

There will always be a substantial difference between the two returns for all investors in the aggregate

Let us consider a realistic scenario of all investors in the aggregate for a particular time period. Suppose that for the 20 years from 1990 to 2009, the typical investor began with \$10,000 in savings and saved a net additional \$1,000 at the beginning of each year. (If there were 100 million investors, this would be an aggregate of \$1 trillion invested at the beginning of 1990 and \$100 billion net savings in each of the 20 years.)

In this scenario, and assuming each investor was invested in the S&P 500 for the years 1990 through 2009, the average investor would have ended 2009 with \$89,298 in savings. The investor's internal rate of return – her “investor return” – calculated based on beginning savings of \$10,000, additional annual savings of \$1,000 and ending assets of \$89,298, would have been 7.52%.

However, the annualized time-weighted rate of return on the S&P 500 for that time period was 8.20%. The investor's investor return “underperformed” the S&P 500's return by almost 70 basis points per year.

Where is there any investor attempt to time the S&P 500 in this scenario? The investor merely saves on a regular annual schedule. Any timing “failure” or “success” is a result of the sequence of returns that the S&P 500 experienced during the time period, not of the investor's effort to time them.

If one is tempted to fault the investor's timing because the investor invested the whole time in the large-cap U.S. stock market, one should rethink. The rate of return for *all* investments – the entire potential investable global investment portfolio – varies from year to year. So even if the average investor invested in the fully diversified global market portfolio, the investor's investor return would be substantially different from the global market portfolio's time-weighted return.

The result is even more pronounced for other time periods. Consider the 10-year time period 1993-2002 inclusive, and suppose the average investor had begun with \$10,000 and made a net investment of an additional \$1,000 each year. Then the annualized investor return would have been 8.25% while the S&P 500's annualized time-weighted return was 9.33%, more than a full percentage greater than the investor's return – not as a result of poor investor timing, but as a result of unfortunate timing beyond the control of the investor.

Yet in Morningstar's report all investor-investment return differences are attributed to investor timing

In spite of the fact that investor-investment return differences can occur independently of good or bad investor timing, in the Morningstar report all investor-investment return differences are viewed through the lens of good or bad investor timing. The investor-investment return difference is analyzed for each type of fund or market sector in light of what may have caused the investor to time badly in the case of investor-investment return shortfalls, or to time well in the case of investor-investment return surpluses. But, clearly, other factors were at work as well that had nothing to do with investor timing. No recognition is made of this fact.

Thus, a complete leap has been made from a too-facile interpretation of the mathematical measure, investor return, and its difference from time-weighted investment return, to a conventional wisdom about the meaning of this difference. Without a better analysis of the characteristics of the two returns and the meaning of the difference, however, we do not really know how to interpret it.

Another very fishy thing about the measure

Recall that Kinnel notes in his report that “Generally, investor returns fall short of a fund's stated time-weighted returns because, in the aggregate, people tend to buy after a fund has gained value and sell after it has lost value.”

How is this possible? For every sale or purchase of a mutual fund, there are one or more counterparties who bought or sold the same amount of the securities in the fund. If the panic-seller or exuberance-buyer timed

the fund's movement badly, the counterparties must have timed the movements well – and for every trade there is an equal and opposite counterparty transaction. Then how could the aggregate of all investors have timed the movements badly?

Consider a moment in time at which investors in a mutual fund are net redeemers of the fund – presumably, according to the now-conventional wisdom, in all likelihood making a mistake by timing their redemptions badly. The orphaned securities are sold by the fund and are bought by counterparties. If the redemptions were timed badly, the trades by the counterparties must be timed well. Who are these superior timers? Are they investors who invest in individual securities rather than in mutual funds?

But Kinnel, in his Morningstar report, says, “Studies show stock investors and pension fund managers do the same thing [buy after a fund or security has gained value and sell after it has lost value].” Where, then, are the good timers? Surely, if one party to a transaction is a bad timer, the counterparty must be a good timer – either consciously or unconsciously; in either case a measure of timing should show that one side timed well while the other timed badly.

What can it mean that all investors time the market badly, traders and their counterparties alike, notwithstanding their precisely opposite timing?

And how can such firm conclusions be drawn from a measure about which – as far as I can tell – these questions have not even been asked, let alone answered, by those who draw the conclusions?

This is just one more example of the failure to properly analyze and understand the measure being used, before leaping to a conclusion about its meaning. This is unfortunately not atypical of standard industry practice, which is to jump to a preconceived or opportunistically desirable conclusion about an often arbitrarily-defined mathematical measure.

This is not science and it is not logic. It is hasty interpretation of mathematical or empirical results in order to feed the need for marketing arguments or talking points.

The upshot in this particular case is that we do not really know that the typical or average investor underperforms because of bad market timing. This new conventional wisdom has quickly become deeply rooted. It will be hard to uproot, but uprooted it should be.

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.

[1] In response to an inquiry to DALBAR as to their mathematical methodology, the company sent me a two-page document containing an example that is much simpler and less explanatory than the one in Box 1. A follow up email requesting more complete and precise specifications of the calculations resulted in this reply: “This is the most detailed description that is available to the public. Anything beyond what is in this document we consider proprietary in nature.”