

What do we know about investor irrationality?

Michael Edesess | EDHEC–Risk Institute | 04 July 2017

It has become conventional wisdom that individuals underperform the investments they invest in, due to their irrational investment behavior. For the creation and propagation of this conventional wisdom, we have DALBAR to thank. Now that Wade Pfau has shown that DALBAR's research is likely to be worthless because [it calculates its numbers wrong](#), it is time to question whether the conventional wisdom has even a scintilla of meaning.

That wisdom says, in particular, that investors underperform the very investments that they invest in, because they get into and out of them at the wrong times. Many investment professionals think that this is obviously true, based on their observation of anecdotal evidence. But like the options day-trader who tells people how successful he is, they may remember the observations they are motivated to remember, and forget or fail to notice the rest. Investment professionals have a clear motivation to believe that individuals are irrational and need their help.

THE HYPOTHESIS

The assumption is that investors panic and get out of the market after it drops, then become greedy and get back in after it rises. This is assumed to be the worst possible timing for the investor.

The respected investment research firm Morningstar Inc., building on what was erroneously thought to be DALBAR's methodology, has done its own calculations to test this. In a recent *Advisor Perspectives* [article](#), Morningstar's head of retirement research, David Blanchett, stated his hypothesis clearly: "If mutual fund investors on average made "smart" market-timing trades, there would most likely be inflows into equity funds at market bottoms and outflows at market tops. What we see, though, is effectively the opposite, where net equity mutual fund flows are positive after the market does well and negative after the market does poorly".

However, timing market tops and bottoms correctly is all but impossible. As I have shown in [a previous article](#), it is quite possible for an investor to exit the market in a panic after a sharp drop, and reenter it again after a sharp rise, and still wind up buying at lower prices than when s/he sold. In fact, if market movements truly followed a mathematical random walk, it would be equally likely that this behavior would result in a buy–low–sell–high scenario as a buy–high–sell–low one. The reason for this is that momentum – a continuation of the downward or upward trend – would be as likely as a reversal.

THE EVIDENCE ADVANCED IN FAVOUR OF THE HYPOTHESIS

Building, as I say, on what was thought to be DALBAR's methodology, Morningstar calculates the time-weighted return on a dollar invested in a mutual fund or funds, and compares it with the dollar-weighted return (i.e. the internal rate of return, IRR) calculated from the cash flow into or out of the fund and its beginning and ending values. The time-weighted return is called the investment return, and the dollar-weighted return is called the investor return. The difference between these two returns, if the investor return is less than the investment return – as it often is – is assumed to be due to investors' poor timing.

But why? How can we be sure that the difference between investment and investor return is actually a measure of the quality of the investor's timing?

We can't be sure of this – or at least not without further, and much more careful and thoughtful research.

I believe this is yet another example of a practice that is all too prevalent in the finance field, in the peer-reviewed academic literature as much as in the rumor mill – namely, the tendency to leap to a preconceived conclusion based on a mathematical observation that is not actually warranted by the math. We don't really know that if the investor return, so measured, is less than the investment return, it means that the investor timed the market badly.

Normally, one would compare two investment strategies by comparing the ending wealth resulting from applying one strategy as compared with the other. But to do that, the cash flows have to be the same in both cases.

Comparing investor and investment return is comparing results with two very different cash flows. It is not even a comparison of ending wealth in both cases, but of rates of return that are calculated in very different ways and have very different meanings.

I'll come back to this comparison, and the difficulty of determining what it really means, later.

But first, I will describe an effort I made to address the fundamental hypothesis.

A TEST OF "POOR" AND "GOOD" MARKET TIMING

The assumed "bad timing" that costs investors their investment returns is their presumed habit of pulling out of the market after a drop, and getting back in after a rise.

The reverse of this – presumably, "good timing" – would be to get into the market after a drop, and to exit after a rise. Is it possible to design a research project to compare the "bad timing" of the previous paragraph with the "good timing" of this one?

That is what I have attempted to do. Let us say, for example, that the poor timer exits the market after it drops 20% or more over a two-month period, and subsequently reenters after it rises 50% over a period of months, while the good timer does precisely the opposite. Both begin by being fully invested in the market. When not in the market, the investor is assumed to invest in short-term T-bills.

The problem with trying to do the comparison with these thresholds is that one of the two investors – the "good timer" or the "bad timer" – will be invested in the market for more months than the other one. That will tend to bias the results in favor of the one that is in the market for longer, and complicate the comparison.

To overcome this difficulty in making the comparison, I used the following methodology. Let's call the threshold for the poor timer getting out of the market – down 20% in the example above – the "drop" and the threshold for getting back in – up 50% in my example – the "rise". What I did was to try two-month drops from –10% to –25% in half-percent increments, and subsequent rises from 50% to 125% in one percent increments, until I found a combination resulting in the "poor timer" being in the market for the same number of months, on average, as the "good timer". For the 50-year, monthly-rolling periods from 1926 through 2016, the combination of drop and rise that produced this result was a drop of 15.5% and a rise of 87%.

It would seem intuitively that getting out after the market dropped and getting back in after it rose would not be as good a strategy as the opposite. And, indeed, my results do seem to confirm that, at least initially. The poor timer's average return was 7.9%, while the good timer's average return was 8.7%. The good timer beat the poor timer in 82% of the 50-year periods. The annualised stock market return averaged over all those 50-year periods was 11.2%. Therefore a better strategy than either that of the "poor timer" or "good timer" would have been to stay in the market the whole time. (I used the CRSP monthly total cap-weighted market returns for the market, and one-month T-bills as the return when out of the market.)

But I also tried two other tests.

One was to run the two strategies over the whole 91-year period from 1926 to 2016. To make it so that the two strategies were in the market for the same number of months, I had to use different drop and rise thresholds: –14% for the drop and +64% for the rise. The results were similar and more pronounced. The poor timer had a 6.1% annualised return while the good timer's was 7.5%. The market's annualised return was 9.8%.

Then I tried another test, with the aim of overcoming the fact that the 50-year rolling periods underweight the earlier and latter parts of the full 91-year period, both times when the market experienced pronounced ups and downs.

The method was to take the last 599 months of the 91-year period and append those in front of the beginning of the period, then run the same test as the first one for 50-year

monthly-rolling periods within the full 1691 months. This results in every month of the 1926–2016 interval being represented the same number of times as every other month.

For this test, the threshold drop and rise had to change again. Now the threshold drop had to be -16% and the rise +103% in order to make the average number of months in the market the same for both good timer and poor timer. Perhaps this change explains the change in the results. In this case the "poor timer" had an average return of 8.6% while the "good timer" had an average return of 6.4%. The average market return was 10.3%. The supposedly good-timing "rational" investor beat the "irrational" investor only 5% of the time.

The results were therefore inconclusive.

NOW WHAT ABOUT THAT "INVESTOR RETURN"?

My true intent in running these studies was to explore the meaning of "investor return", as defined by Morningstar: the dollar-weighted return on flows into and out of a portfolio of assets (in this case the US stock market). Often, one gains insight into the meaning of a mathematical measure in finance by "operationalising" it – that is, by trying to apply it to real cash flows and balances.

"Investor return" fails the test. When I calculated investor return for the 50-year periods, it made little sense and seemed to have little use for purposes of the comparison to which it is often applied. A simple example will show why.

Consider a "good" investor who started with \$1,000 and was extremely good at timing the stock market over the years 1998 through 2016. This investor's strategy was to get out of the market after a one-year drop greater than 20%, then to get back in after the market had risen more than 55% over subsequent years. Figure 1 shows the results.

Figure 1: Calculation of dollar-weighted rate versus time-weighted rate

Year	Stock Return	T-bill Return	Investor Balance	Cum Mkt Rtn Since Exit	In/Out of Market	Flow In/Out of Market
			\$1,000			\$1,000
1998	24.30%	4.78%	\$1,243	24.30%		\$0
1999	25.22%	4.56%	\$1,556	55.65%	Out	-\$1,556
2000	-11.54%	5.76%	\$1,646			\$0
2001	-11.16%	3.78%	\$1,708			\$0
2002	-21.14%	1.63%	\$1,736		In	\$1,736
2003	31.62%	1.02%	\$2,285	31.62%		\$0
2004	11.97%	1.20%	\$2,559	47.37%		\$0
2005	6.15%	2.96%	\$2,716	56.44%	Out	-\$2,716
2006	15.49%	4.79%	\$2,846			\$0
2007	5.80%	4.67%	\$2,979			\$0
2008	-36.67%	1.47%	\$3,023		In	\$3,023
2009	28.80%	0.10%	\$3,894	28.80%		\$0
2010	17.72%	0.12%	\$4,584	51.63%		\$0
2011	0.74%	0.04%	\$4,617	52.75%		\$0
2012	16.15%	0.06%	\$5,363	77.41%	Out	-\$5,363
2013	35.15%	0.03%	\$5,365			\$0
2014	11.58%	0.02%	\$5,366			\$0
2015	-3.36%	0.01%	\$5,366			\$0
2016	13.59%	0.19%	\$5,376			\$0
Ann Rtn:	6.70%	1.94%	Ann Rtn: 9.3%		IRR:	18.9%

When out of the market, the investor was invested in T-bills. The column headed "Investor Balance" shows the investor's cumulative balance. The last column on the right shows the investor's flows in and out of the market at the end of each year.

We can see on the bottom line that the investor was indeed a very good timer, obtaining an annualised rate of return of 9.3% (ending balance of \$5,376 divided by beginning balance of \$1,000, and annualised). This exceeded the market's return of 6.7% itself ("investment return") by 2.6 percentage points.

But what are we to make of the last figure on the lower right? This is the "investor return" on the market, the dollar-weighted return on all flows into and out of the stock market. This appears to show that the premium that the investor received for good timing was 12.2% ("investor return" of 18.9% minus the market's investment return of 6.7%), not a mere 2.6%. Which one is right?

To answer this, we need to answer another question, which has been kicked around in a discussion on *APViewpoint*. What does the IRR, i.e. the "investor return" implicitly assume is the investor's alternative investment – the one that the investor invests in when not invested in the asset in question.

The answer is that the "investor return" implicitly assumes that the alternative investment is one that yields the same ROR as the calculated IRR – in this case 18.9%.

This is, at least in this case, obviously an absurd assumption. But it is why the "investor return" came out so much higher than the "investment return".

THE BOTTOM-BOTTOM LINE

The comparison of "investor return" to "investment return" has not been thought through carefully enough to warrant using their difference as a measure of "poor timing" or "good timing". Its use to draw such conclusions is an example of a common practice in finance – found at least as much in peer-reviewed academic literature as in the rumor mill – namely, to leap to a conclusion that is supposedly based on a piece of mathematics, without adequate consideration of whether the mathematics actually warrants that conclusion. Often, the conclusion is one that has already been preconceived by the researcher.

There are other problems as well with the conclusion that investors underperform through bad timing, which I have taken up in [other articles](#).

But this one, I believe, should seal the resolution that a comparison of "investor return" with "investment return" should not be used to draw any implications, unless and until its meaning is clarified through further research.

Noteworthy is that my research project described above clearly shows, if nothing else, that in the three cases I ran, staying in the market the whole time beat both the "poor-timing" and the "good-timing" strategies. To the extent that an investment professional's counsel is simply to refrain from trying to time the market, it may well add value.



Michael Edesess, a mathematician and economist, is a senior research fellow with the Centre for Systems Informatics Engineering at City University of Hong Kong, chief strategist of Compendium Finance, adviser to mobile financial planning software company Plynty.

This article is abridged and reproduced with permission from [Advisor Perspectives](#).
