

Fun with Numbers: *Or, Why All the Managers Are Above Average*

"Figures don't lie, but liars can figure."
-Old, old truism

"There are three kinds of lies: Lies, damned lies, and statistics."
-Attributed to Benjamin Disraeli

Introduction

Investors who review marketing pitches from money managers may wind up feeling as if they have stepped into Lake Woebegone: Just as with all the children of Garrison Keillor's fictional village, all the money managers are above average. Or at least, that's what would appear to be the case, based on the track records presented.

Since it's not mathematically possible for everyone to be above average, investors may wonder what's going on. The answer is that there are multiple ways to present numbers. Investors should assume that money manager marketing materials will naturally choose the presentation that leaves the most favorable impression of how they've done.

This article explores how money managers can shape investors' perceptions of performance, through the way they present their track records. The main topics covered include:

- I. *Gross or net*
- II. *Price-change-only versus total return*
- III. *Hypothetical, back-tested, or simulated, versus live-money returns*
- IV. *Annualizations*
- V. *Audits and performance standards*
- VI. *Take-aways: What investors should do*

I. Gross or Net

The first way managers can shade performance data is the simplest: Just omit fees. That is, report the gross-of-fees returns generated by *the product*, rather than the net-of-fees return retained by *the client*.

The usual rationale given for this omission is that different clients are charged different fee rates. For example, the fee for a \$10 million investment, as a percentage of the amount invested, is often lower than the fee for a \$10,000 investment. Large accounts get fee breaks, usually at pre-specified breakpoints.

However, gross-of-fees returns are often seen in reports of individual accounts provided to their account-holders. Managers in this circumstance certainly know how much they are charging the investor. They're just choosing not to make it transparent.

Please note that publicly-traded mutual fund returns are always quoted net of management fees: The embedded fee is deducted prior to the return calculation. However, in the case of funds that carry a load (sales fee), the load paid by each particular investor will not be reflected in performance calculations provided on public databases.

Gross return quotations arise typically in marketing pitches and performance reports for products that are sold as separate accounts (where the account-holder owns securities directly, rather than a share of a commingled pool), and sometimes for private equity funds and other alternative investments.

Sometimes, rather than omit net returns altogether, managers provide both gross and net returns, but do so in a way that draws investors' attention to the gross return series. For example, gross returns may be shown in **large, bold** typeface, while net returns are shown in smaller, lighter typeface.

Inclusion or exclusion of fees can have a significant effect on the performance impression conveyed, especially in a generally low-return environment. In the case of separate accounts investing in equities, a typical fee might be 1.0% - 1.5% of assets per year. In the case of alternatives, it will generally be meaningfully higher.

So, how can an investor know whether returns are quoted gross or net? First, look carefully at the marketing information or performance report at issue. It should state whether performance is gross or net of fees—though this information may be in a footnote, in very small print.

If the information isn't provided, *ask*. Anyone who hesitates to answer this question should not be entrusted with your money.

Taxable individual investors should also consider another potential deduction from gross returns, which is taxes. A full discussion of taxation of investment products is outside the scope of this article. However, investors should be aware that some investment products, typically vehicles holding equities for long periods of time, will generate a return stream characterized primarily as capital gains, while other products, such as high-turnover trading strategies and those employing short-selling, will generate returns that are mostly ordinary income and/or short-term gains, which are taxed at higher rates.

II. Price Change vs. Total Return

The provision of gross returns instead of net returns is intended to make a manager's *absolute* returns (the amount of gain generated) look better than they are. This next statistical ploy is intended to make a manager's *relative* returns (returns compared to a benchmark) look better than they are.

NASDAQ Price Index versus Total Return

	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	Ending Value / \$1 Invested	Ann'd Return
Price Index	9.52%	9.81%	(40.54%)	43.89%	16.91%	(1.80%)	15.91%	38.32%	13.40%	5.73%	7.50%	\$2.441	9.33%
Total Return	10.28%	19.55%	(40.03%)	45.32%	18.02%	(0.83%)	17.45%	40.12%	14.75%	6.96%	8.87%	\$2.735	10.58%
Difference	0.76%	0.74%	0.51%	1.43%	1.11%	0.97%	1.54%	1.80%	1.35%	1.23%	1.37%		1.25%

Total returns on investments are composed of two parts: An income return, from interest and/or dividends, and a capital gain or loss, from the change in the price of a security between the time it is purchased and the time it is sold. Managers can make their relative returns look better by comparing their total returns to the price-change-only return on a benchmark.

The NASDAQ index return seems to be used most often in such a calculation, perhaps on the theory that the NASDAQ index is dominated by high-growth companies that pay low or no dividends. However, even the NASDAQ includes a reasonable component of income (i.e., dividends). Note also that income returns from dividends are much more stable than capital gains from price changes.

A manager who used a NASDAQ price-change return as a comparator was automatically handicapping the benchmark by an annualized 1.25% over the ten years ending in 2016. That could easily make the difference between appearing to beat, or appearing to lag, the benchmark (see table below).

Other common equity indexes typically include a greater income component than does the NASDAQ. For example, as of early March of 2017, the annualized dividend yield on the S&P 500 was about 2.5%.

How can investors tell which calculation is being used? First, any benchmark nomenclature such as "S&P price" or "NASDAQ index" should prompt the question: "Is this index calculated on a total-return or price-change basis?"¹ Second, check the footnotes and small print; these, also, should tell you the answer.

¹ If the benchmark name is followed by the abbreviation "TR", this means the calculation is for the Total Return, i.e., it includes income.

III. Hypothetical, Back-Tested, or Simulated vs. "Live-money" Returns

The use of gross versus net returns or total returns versus price-change-only indexes are two methods of shading the truth of what happened. Hypothetical, back-tested, or simulated returns take the issue to a different level: They are returns that never actually happened at all. They are simply a result that theoretically *might* have happened.

Hypothetical, back-tested, or simulated results are an answer to the marketer's question of how to sell a newly-created financial product that has no track record. In some cases, a new product may be created by, for example, combining and repackaging existing products. In other cases, a new product may be created based on a financial theory of what might generate attractive returns.

In the former case, creating a hypothetical return is straightforward: Simply blend the track records of the existing products in the proportions expected for the new product, and show that as a hypothetical return history. As one might suppose, the components of such "new" products generally have performed well—they wouldn't be chosen otherwise—so that the new product appears to have a strong record.

In the case of new products based on financial theories, statistical modeling is generally involved. The question the model addresses is, How would a portfolio based on this theory have performed, if it had actually existed?

Unfortunately, many of these models have an element of circularity: They mine data to see what actually *did* work, and then create models that would have captured and exploited the winning factors².

The problem is that life, and markets, change. Consider, for example, a model based on U.S. capital markets data from, say, 2000 to 2007. Such a model might wind up favoring securities of companies that were highly-leveraged participants in the real estate and financial services industries, because those stocks were the big winners over the period in question. But an investor who loaded up on such securities in 2008 would have invested in a portfolio that was on the edge of disaster, as the financial crisis was about to strike.

The term "live money" refers to portfolios that are *not* hypothetical or modeled: Live money is actual money, invested in actual products, earning an actual return, by actually doing whatever the manager has promised to do.

Existing U.S. mutual funds, for example represent live money. If you go on the Morningstar website (www.morningstar.com) and look up the return for, say, the Vanguard Dividend Growth Fund (ticker VDIGX) for calendar 2016, you will find that the return was 7.53%, and the average asset base was nearly \$31 Billion. Had you been invested in the fund as of January 1, 2016, and held it throughout the year, you would have earned this return—as millions of investors, in fact, did. There was nothing hypothetical about it.

Investors who may be tempted by marketing claims made for modeled or hypothetical portfolios should be aware that, to use the old cliché, there tends to be a pretty big slip between the cup and the lip. When a hypothetical or modeled strategy gathers actual money and actually invests—that is, when it goes live-- the resulting returns tend to be meaningfully lower than the modeled returns had been.

"Meaningfully lower" can mean several percentage points of return per year—and it can also include blow-ups that permanently impair investor capital. Investment history is replete with examples of leveraged trading strategies that looked wonderful on paper, and perhaps had initially strong results with modest amount of live money—but then crashed, after assets grew, markets changed, and/or leverage exceeded prudent levels.

² There are many possible sources of weakness in financial models. A full discussion of this topic is outside the scope of this article.

IV. Annualizations

The use of annualized returns is perhaps the most widespread form of marketing “fun with numbers”. With enough creativity, and at least a few good returns along the way, marketers can perform magic-- or something close to it-- in terms of the way investors perceive investment track records.

We’ll get to an example in the minute. But first: What, exactly, is an annualized return, and how do you calculate it?

Conceptually, an annualized return is the single-point return that would equate a beginning value with the related ending value, given the stream of actual returns that occurred. Technically, it is referred to as a geometric (multiplicative) average return, and is differentiated from an arithmetic (additive) average return.

A simple example will illustrate the difference between multiplicative and additive averages. Suppose an investment starts at \$100, then rises +50% in year 1 and subsequently falls -50% in Year 2. What is the return?

On an arithmetic-return basis, the return is zero: $(+50\% - 50\%) / 2 \text{ years} = 0$

However, money does not grow on an arithmetic basis. It grows on a multiplicative basis. In this scenario:

$\$100 * (1+0.5) = \150 at the end of Year 1 $* (1-0.5) = \$75$ at the end of Year 2³

So, the investor here would wind up with an ending value of \$75, having started with a value of \$100. The cumulative loss over the two years would be -\$25, or -25%.


Now we need to convert that two-year cumulative return into the annualized return that would have produced the same result. This is done by exponentiating (raising to the power of) the end result by $(1 / \text{by the number of periods})$ and then subtracting the 1, the placeholder for the original investment. The up-caret symbol (^) represents exponentiation.

Here:

$$0.75^{(1/2)} - 1 = -13.4\%$$

That is, a \$100 investment that lost -13.4% in year 1, to a value of \$86.40, which subsequently lost another -13.4% in Year 2, would end at \$75—precisely the cumulative value obtained above.

+ 50% Return Followed By - 50% Return ≠ 0!



The general formula for an annualized return is:
 $CAGR = ((\text{Ending Value} / \text{Beginning Value})^{(1/n)}) - 1$

Where:

CAGR = Compound Annualized Growth Rate
 n = Number of compounding periods

³The 1 in this equation is a placeholder, representing the original investment.

Now let's make this more interesting, with a real example.

Suppose an investment manager were faced with having to market a product with a track record that, for the last seven calendar years, looks like this:

Last year	2 yrs ago	3 yrs ago	4 yrs ago	5 yrs ago	6 yrs ago	7 yrs ago
26.1%	57.7%	(7.4%)	(2.9%)	(27.6%)	61.6%	(18.1%)

This is part of the actual track record of a real investment product. It would be a tough sell based on that set of numbers: They display enormous volatility—way up! way down!-- and four of the seven calendar years have produced negative absolute returns.

This is where annualization generally comes in handy. There are two standard types: Rolling-period (typically three or five years) annualizations, and annualizations to an endpoint. The following chart illustrates, conceptually, what would be included in each type.

Annualized rolling three-years

	Last yr	2 yrs ago	3 yrs ago	4 yrs ago	5 yrs ago
Ending Last yr	X	X	X		
Ending 2 yrs ago		X	X	X	
Ending 3 yrs ago			X	X	X

Annualized to an endpoint

	Last yr	2 yrs ago	3 yrs ago	4 yrs ago	5 yrs ago
One year	X				
Three years	X	X	X		
Five years	X	X	X	X	X

So let's apply these concepts to the actual return stream given above.

For example, to get the three-year annualized return ending last year, we would

(1) Calculate the cumulative return by multiplying (1+each year's return):

$$= (1+.261) * (1+.577) * (1+-.074) = \mathbf{1.841}$$

so \$1 grew by 84.1% over the three years.

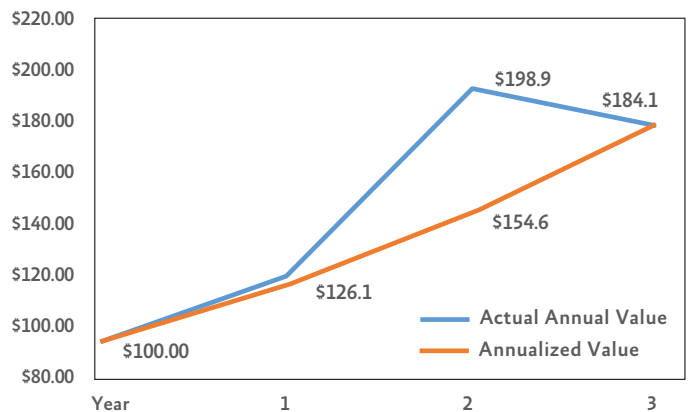
⁴ This is as far back as we can go. The rolling three years ended 5 years ago includes the calendar years 5, 6, and 7 years ago. Seven years ago was the endpoint of the track record.

(2) Raise the calculated cumulative number to the power of (1/ # of years) and subtract 1, which is the original dollar, to get the annualized three-year return:

$$= 1.841^{(1/3)} - 1 = \mathbf{22.6\%}$$

So, the annualized three-year return, ending last year, was 22.6%.

Annual and Annualized Returns: Same Outcome, Different Paths



Let's go back to the problem confronting the investment marketer with the volatile track record to massage. Suppose she tries, first, to smooth out the record by using annualizations based on rolling three-year returns. Here's what the results would look like:

Rolling Three Yrs Ending	Calendar Year Return Data Included							Annualized
	Last yr	2 yrs ago	3 yrs ago	4 yrs ago	5 yrs ago	6 yrs ago	7 yrs ago	
Last yr	26.1%	57.7%	(7.4%)					22.6%
2 yrs ago		57.7%	(7.4%)	(2.9%)				12.3%
3 yrs ago			(7.4%)	(2.9%)	(27.6%)			(13.3%)
4 yrs ago				(2.9%)	(27.5%)	61.6%		4.3%
5 yrs ago					(27.5%)	61.6%	(18.1%)	(1.4%)

That looks better than a track record showing a -27.6% loss on one year, but it probably still doesn't satisfy the marketer: Two out of the five rolling three-year periods still show a loss.

And therefore the marketer, noticing that the last two years have been strong, decides to calculate annualized returns for all periods ending last year. This is the result:

Calendar Year Return Data Included

Annualized Returns for Periods Ending Last Year	Last yr	2 yrs ago	3 yrs ago	4 yrs ago	5 yrs ago	6 yrs ago	7 yrs ago	Annualized
1 yr	26.1%							26.1%
2 yrs	26.1%	57.7%						41.0%
3 yrs	26.1%	57.7%	(7.4%)					22.6%
4 yrs	26.1%	57.7%	(7.4%)	(2.9%)				15.6%
5 yrs	26.1%	57.7%	(7.4%)	(2.9%)	(27.5%)			5.3%
6 yrs	26.1%	57.7%	(7.4%)	(2.9%)	(27.5%)	61.6%		13.1%
7 yrs	26.1%	57.7%	(7.4%)	(2.9%)	(27.5%)	61.6%	(18.1%)	8.0%

Voilà. A scary-looking track record characterized by enormous volatility and steep interim losses has, through the math of annualization, been transformed into something much tamer-looking, with every calculated return showing as a gain. This is because the calculations overlap: The recent strong period (last year and the year before) is counted in each one.

So what is wrong with this? The annualized returns for the periods ending last year are correctly calculated. They are accurate. They do present what happened.

The problem is that they mislead investors about what can be expected in terms of performance patterns. In order to get to the long term, investors have to live through, and hold through, the short term. Many investors, faced with three consecutive years of losses, which cumulatively would have cost -35% of their investment value (as happened here in the period between five years ago and three years ago), would have sold out long before the investment rebounded, and so would not have captured the positive returns of the last two years.

Any investor presented with performance data showing only annualized returns should ask for the individual period (e.g., calendar-year) returns. And anyone who sees performance data showing only annualized return data to the same endpoint can bet the endpoint was strong, and is masking weaker prior performance⁵.

So, what is the investment of which the above represents a track record? Nothing exotic: It's the Vanguard Emerging Markets Index fund, for the period from 1998 to 2004. The author originally used this track record for a 2005 program of instruction for bankers. It's proven useful ever since.

For anyone tempted to say that such an extreme could not happen again: It could, and almost certainly will. We just don't know when, or to which investment or category. This is why the prudent investor diversifies, avoids excessive leverage, and keeps a long-term perspective.

V. Audits and performance standards

Two more aspects of quality-assurance for investment track records are audits and performance standards.

The first is straightforward. Investment track records—particularly anything that looks suspiciously positive—should be verified by a reputable, independent third-party accounting firm as auditor. Much investment heartache could be avoided in this way. The Ponzi scheme run by Bernie Madoff is a spectacular case in point, but it is only the most visible and publicized of recent cases.

Any auditing firm whose name is unfamiliar to a prospective investor should prompt an internet search, at the least. Any investment product sold by a firm that lacks an independent auditor should be avoided altogether.

Performance standards are more subtle. The nuances may require some study. Briefly, however, the words investors should look for are “GIPS-verified”. They indicate that performance data has been prepared in compliance with the Global Investment Performance Standards and verified by an independent third party. These are voluntary standards, created and promulgated by the Investment Performance Council of the CFA Institute⁶.

⁵ This writer vividly recalls a marketing pitch that hit a trifecta of Fun With Numbers: The track record emphasized the annualized ten-year premium to the market; on a gross-of-fees basis; the entire ten-year purported premium had been “earned” via a 2500 basis-point (25 percentage point) outperformance in one historical year—and that year was a simulation.

⁶ Chartered Financial Analyst Institute. The CFA Institute describes itself as “a global community of investment professionals working to build an investment industry where investors’ interests come first, financial markets function at their best, and economies grow.” For more information, please see <https://www.cfainstitute.org>.

These standards have been in effect since 2006, with predecessor standards in effect since 1999. They require, among other things, that returns be presented on a time-weighted basis⁷; that any composites presented be logical and meaningful⁸; that disclosure be made as to whether returns are gross or net of fees; and that no “cherry-picking” of the best products or performance years occur.

Not all firms present GIPS-compliant track records, and GIPS compliance does not assure the quality of an investment strategy. However, this is a meaningful step in the right direction.

VI. Take-Aways: What’s an Investor to Do?

Here are some steps an investor can take, when reviewing money manager marketing documents:

First—At the risk of sounding like a pre-millennial broken record: **Read the fine print.** Read the footnotes; read the endnotes; read the disclaimers. This is where you’re likely to find the truth.

Second—Develop, and stick with, your own preferred format for reviewing performance track records, and insist that managers provide it in that format, so that you can make “apples-to-apples” comparisons.

At BWA, for example, we require managers to provide us their performance in the form of quarterly returns, net of fees. These data then become inputs to the proprietary analytical tool we use to screen managers. We have chosen quarterly returns as our comparison basis because the quarterly period reflects our reporting cycle and is the shortest period over which we think manager performance comparisons can reasonably be made. We use net returns because this reflects what the *investor* (as opposed to the *investment*) receives.

Third—Do your own homework and come to your own conclusions. Don’t simply accept whatever you are told.

For example, at BWA, we make a practice of requiring money managers to submit their marketing materials

to us well in advance of any meetings. This allows us time for thoughtful review and analysis before each meeting. We strive always to be prepared with the questions *we* want to address, rather than being in the position of simply responding to whatever we are being told.

Wherever possible, we also use independent sources, such as Morningstar, to provide data on investments. We do not meet with wholesalers⁹ of mutual funds and other publicly-traded investment products, as their perspective is unavoidably biased by their financial interest in selling their products. When we do meet with financial salespeople, such as representatives of private equity funds and other non-public products, we engage in an extensive independent due diligence process, including checking with public and private information sources.

Fourth—Make sure the track records you are being shown, and the financial products you are being offered, come from upstanding firms with clean audits from reputable auditing firms and a rigorous internal operations and compliance culture. And look for GIPS compliance.

Summing Up: Trust, but Verify

At BWA, our approach can be summed it up in four words: **Be skeptical and vigilant!** The marketers of Wall Street are exceptionally adept at making numbers tell almost any story they want (or need) to concoct. The proverb often used by President Reagan when discussing relations with Russia is at the core of BWA’s value proposition: We verify the basis of managers’ performance claims. If we find the evidence of a disciplined and repeatable investment process along with a demonstrated manager skill set, we then begin a rigorous due diligence review process. Our earlier white paper – *How We Invest* – explains the nuts and bolts of our approach in detail. In the end, the numbers we typically rely on are “un-spun” in order to understand thoroughly the potential opportunity -- as well as the risks it may carry.

⁷ See the Appendix to this article for a discussion of the difference between time-weighted and dollar-weighted returns.

⁸ So that a composite does not mix, say, accounts investing in U.S. equities with those investing in emerging markets bonds. These are two different categories of investments, and should go into separate composites.

⁹ In this context, a wholesaler means someone employed by an investment product manufacturer, such an investment bank or money-management firm, to sell products to financial advisor intermediaries like BWA.

About the Author



Elizabeth P. Anderson, CFA, is the founder of Beekman Wealth Advisory LLC, a boutique financial consultancy that provides outsourced CIO (Chief Investment Officer) services and custom wealth education to a select number of ultra-high net worth families.

Publication Date: MAY 2017
 © 2017, Beekman Wealth Advisory, LLC

Appendix

Time-Weighted Versus Dollar-Weighted Returns

Time-weighted returns describe what happened to an investment generically, assuming dollars were invested at the beginning of the time period at issue and simply left there. Dollar-weighted returns (also called IRRs, or Internal Rates of Return) describe what happened to a *particular investor's* investment, based on when she actually invested or withdrew dollars. They can be quite different.

Consider the following example:

Three different investors are considering the same investment. Each investor has \$100 to invest, for a period of two years.

Investor #1 has done her diligence, and is convinced the investment is worth pursuing. She invests her entire \$100 at the outset of the two years, and leaves it to compound.

Investor #2 has decided to dollar-cost average. He will invest \$50 of his \$100 at the beginning of Year 1, and then the other \$50 at the beginning of Year 2.

Investor #3 prefers to wait and see what happens before fully committing her dollars. She will invest \$10 at the outset of this investment period, and then add the remaining \$90 if all goes well.

The investment goes on to gain +30% in year 1 and lose -10% in Year 2.

Here are the outcomes:

	Year 1			Year 2			Dollar-Wtd Ann'd Return
	Investment	Return	Value	Investment	Return	Value	
Investment		30%			-10%		8.2%
Investor #1	\$100	30%	\$130	\$130+0	-10%	\$117.0	8.2%
Investor #2	\$50	30%	\$65	\$65+\$50	-10%	\$103.5	1.7%
Investor #3	\$10	30%	\$13	\$13+\$90	-10%	\$92.7	-3.7%

In each case, the time-weighted return—the return to the investment itself—is the same: An annualized rate of +8.2%. However, the only investor whose personal, dollar-weighted return matches this is Investor #1—the investor who invested her money and left it.

Investors #2 and 3 had less than all of their money invested during the stronger period (Year 1), but all of their money invested during the weaker period (Year 2), so that their personal, dollar-weighted returns fell below the returns of the investment itself. Investor #3 actually lost money, even though the investment itself was profitable over the full two years.

In actual practice, investors' dollar-weighted returns tend to lag the time-weighted returns of their investments because of poor timing decisions: Investors tend to put more money into investments that have already risen, and to withdraw money from investments that have already fallen.