"Truthiness," "Mathiness" and the Costs they Impose on Your Clients' Assets

March 22, 2016

by Michael Edesess

The comedian Stephen Colbert coined the word "truthiness" to mean something that "feels right in the gut" but lacks empirical or theoretical support. Its counterpart in the realm of finance is "mathiness," where academicians or marketers present seemingly rigorous mathematical proofs of assertions that, upon closer inspection, have little or no basis in reality. If the investment products you use rely on mathiness and not provable math, you and your clients will ultimately pay a price.

Mathematical sophistication in the investment industry is a sham. Most mathematical articles in peer-reviewed academic finance journals suffer from lack of model specificity, failing either to properly define terms or to specify exactly what they are trying to prove, or drawing conclusions in ordinary language that aren't warranted by the mathematics. In these respects, they resemble more what economist Paul Romer has called mathiness than real mathematics. Even when the articles are mathematically sophisticated and do not have these failings, investment industry apostles tout, too strongly, conclusions based on them that cannot be found in the articles themselves.

I will illustrate this with a particular line of mathematical finance that has been pursued at great length in the literature (perhaps following up with others in future articles). To emphasize the positive, I will use two examples of sound mathematics applied to this topic, one by a physicist and the other by two mathematicians. However salutary – and rare – may be the contributions that these works make to the process of academic thought experimentation, nevertheless, like so many similar articles, they have, as the authors themselves reveal, no immediate practical implications for investors.

The mean and the median

It is well known that the topmost echelons of the wealth and income spectrum in the United States have experienced higher rates of real income growth in recent decades than the lower echelons. For example, the real after-tax incomes of the top 1% of earners rose by 138% since 1978, while those of the bottom 90% rose by only 15%.

Yet the only summary statistic typically cited for the economy as a whole is the growth of GDP per capita, the aggregate of all incomes divided by the number of recipients. The GDP per capita over the same period grew by 82%; this is the mean of the distribution of all income growth rates. Income growth at this high a rate was, however, experienced by only a small minority of the population. The median of the distribution was much lower.

But now suppose that each income recipient could experience the income growth rate of a different U.S. citizen every year. One year in a hundred, you'll get a raise equal to the annualized equivalent of the 138% earned by the top 1% since 1978; 90 years in a hundred you'll get a raise equal to the annualized equivalent of 15% over that period, or about 0.4%.

What kind of income growth would you experience over a long period of time? It would be a growth rate that is close to the median of all growth rates, or close to zero, not to the mean of 82%.

Which growth rate is more appropriate to use, the mean or the median – the one-year average of all income growth rates, or the long-term time-average?

The answer, of course, is that it depends on the context – the purpose for which you are using it. And, is it necessary to use only one of these averages when other information about the distribution of income growth rates is available too?

Yet this little question – that is, which one should you use, the mean growth rate or the median? – has been the underlying subject of scores and probably hundreds of finance articles. If this is not the underlying subject, it is difficult to discern what is. Yet most of these articles obscure this simple question under layers of obtuse and hard-to-interpret mathematics. Then, they draw conclusions that are unrelated to, or well beyond, what the mathematics implies.

The simple math that subsumes all of these articles

Growth rates can be expressed as being continuously compounded, periodically compounded or not compounded at all. For example, at a 6% annual interest rate, $1,000 will grow in two years to $1,126.50 if the 6% annual rate is continuously compounded; $1,126.15 if it is monthly compounded; $1,123.60 if it is compounded annually; and $1,120.00 if it is not compounded at all (that is, if the two-year uncompounded interest rate is 12%).

If the rate is continuously compounded, then rates add over time. For example if the continuously-compounded rates are 6% in the first year and 5% in the second, the two-year continuously-compounded rate is 11%, or 5.5% if annualized.

If each continuously-compounded annual growth rate is generated randomly from some probability distribution, then over a long period of time, the annual average rate will eventually converge with certainty, as time goes to infinity, to the mean of the distribution. If the continuously-compounded rate has a symmetric distribution, this number will also be its median, and the long-run time-averaged rate will converge to the median as well.

But, if you convert these continuously-compounded growth rates to annually-compounded rates (by exponentiation), the time average rate will converge to the median, and the mean won't equal the median anymore. This results in the observation, by virtually hundreds of finance articles, that if the continuously-compounded rates are normally distributed, then the annually-compounded rates will be log-normally distributed, and the mean and median will differ by half the square root of the variance.

There's nothing magical about that – it can be easily obtained by completing the square in an integration (at least easily for a competent freshman in a calculus course). It's a consequence of the relationship between the normal and log-normal distributions, no more and no less. Yet in the finance literature, it is repeatedly derived by using differential stochastic calculus, a method that is completely unnecessary, not to mention hazardous because it involves the use of "infinitesimals" (quantities that, by mathematical definition, are smaller than any known quantity). Practical investment inferences are sometimes mysteriously attributed to this difference – inferences such as "diversification bonus" or "rebalancing bonus" – that it in no way implies.

Ole Peters and the solution to the St. Petersburg paradox

The St. Petersburg paradox is an extreme form of the difference in aggregate growth rates that is illustrated by the income growth example above. This paradox was first put forward by the mathematician Nicolaus Bernoulli in 1713 and has intrigued people ever since.

In the St. Petersburg paradox, a game is played in which a coin is tossed repeatedly. If it comes up heads on the first toss, the player receives a dollar, and the game ends. If it comes up tails on the first toss but heads on the second, the player receives $2. If it comes up tails twice and then heads, the player receives $4. If tails three times and then heads, the reward is $8. And so forth. If the coin comes up tails 24 times and then heads, the player will receive 2 to the 24th power, or more than $16 million – but the chance that this will occur is less than one in 16 million.

The mathematically calculated "expected payoff" for this game is an infinite number of dollars. Should the player then be willing to pay an infinite amount of money to play? No – no real player is willing to pay more than about $10, because the chances are that the game will pay off and end after no more than the first four coin-tosses, returning the player only $1, $2, $4 or $8 for her $10 bet.

Ole Peters, a physicist who is a fellow of the London Mathematical Laboratory and an External Professor at the Santa Fe Institute, points out in an article on the paradox that economists explain it by using utility theory. They argue that the utility of $16 million is not two million times the utility of $8, but less than that. Therefore the "expected utility" of the game might be only equal to the utility of an initial bet of $10.

Peters calls the probability-weighted average of all payoffs from playing the game – which is what the expected payoff is – the "ensemble average." He argues that the explanation of the St. Petersburg paradox doesn't require utility theory, and also that restrictions economists place on their utility functions are unnecessary if you use the time-averaged return on the game instead of the ensemble average – that is, the long-run result of playing the game repeatedly over time.

In the problem about the distribution of income growth rates, Peters would call the annual GDP growth per capita the "ensemble average." By contrast, the "time average" is the annualized average growth over time, if a person could experience in succession the income growth rates of each and every person. The ensemble average is the mean of the distribution, or what probabilists call the "expected value," while the time average for a continuous probability distribution – over long periods of time – converges to the median, not the mean.

Peters has applied these theoretical concepts to finance and economics. In an extended Skype interview, I asked Peters what practical implications his work has for investing. He said that, fortunately, he doesn't have to worry about that. His work – as it relates to investing at least – remains in the theoretical realm.

Peters is a thoughtful person with an interest in wide-ranging problems to which his work might apply. He has applied it to atmospheric physics, as well as to economics and mathematical games. But he is correct in suggesting that his work does not, or does not yet, have any clear application to the world of practical investing.

Even so, it would not surprise me to hear that next week some investment manager has claimed to have a market-beating strategy based on the work of Ole Peters.

Pal and Wong and the long-run decomposition of returns

Two mathematicians at the University of Washington, Soumik Pal and Ting-Kam Leonard Wong, have written a paper titled "Energy, Entropy, and Arbitrage." I include this paper not only because its authors explore much the same math as Peters, which relates to long-run average returns, but because Pal and Wong's paper is a fine example of how mathematics in the finance field ought to be done.[1] Peer reviewers in the finance field should read this paper very carefully to see what proper mathematics looks like. If they did, perhaps they wouldn't approve the sloppy math that runs rampant in finance journals.

In Pal and Wong's paper, everything is defined precisely and meticulously, as is the passage from one step to the next. The paper sidesteps the use of differential stochastic calculus, thus avoiding the highly questionable conclusions that even papers that it cites create. For example, some papers cited by Pal and Wong blithely and tacitly assume the existence of what can only be called – and often is explicitly called – "continuous rebalancing." As Pal and Wong themselves remark, continuous rebalancing is an impossibility; real-world portfolios can be rebalanced only at finite intervals. It is not at all clear – and it is doubtful – that rebalancing at ever-smaller intervals would converge to a probability distribution similar to one that is derived in some of these articles, but in a much more cavalier and poorly specified manner than the work of Pal and Wong.

Pal and Wong, nevertheless – following in the path of these debatable precedents – produce a paper that, because it is so careful and so well-specified, merits serious consideration of its content. In it, the authors try to decompose long-run investment return into three components, which they call volatility return (defined to be always positive), dispersion return and drift return. The first component they also call energy, and the latter two "relative entropy" because of a close resemblance of their formulas to the formula for entropy in information theory.

The upshot was stated to me in an email from Wong, in which he says that the question is "what makes the rebalancing portfolio out/under-perform the buy-and-hold portfolio? Our answer is that it is the tradeoff of two quantities: relative entropy and energy (excess growth rate, aka 'market volatility'). The rebalancing portfolio outperforms *if and only if* 'energy' is larger than relative entropy."

In other words, they do not reach any firm conclusion, such as that rebalancing is always better, or even better on average in some sense. They are exploring the circumstances under which it is better and asking questions like, as Wong wrote to me, "i) What path properties are needed for outperformance," and "ii) Does the market satisfy i)".

A liability of their approach – so far as the possibility for practical implications is concerned – is that it is an exploration of the relationships among the components of investment return in the long run, meaning as time approaches infinity. It is doubtful that the limit of the average return over time as time goes to infinity has any practical implications for real-world investing, no matter what the results. As time goes to infinity – as I mentioned earlier – the time-averaged return converges with certainty to the median of its distribution. But what does that imply for practical investing? Over any finite time period the probability distribution of final wealth, and the income that can be derived from it, will be spread over a wide range of possible values. Over no finite time period do the wealth or income possibilities collapse to a single number, as the time-averaged rate of return does in eternity. What, then, can be the practical implications?

Unfortunately, Pal and Wong's very tentative exploratory work has already appeared in marketing documents as support for the idea that rebalancing provides a bonus – without specification even as to whether that means a bonus "on average." Such are the exigencies of marketing in the highly competitive and absurdly high-profit investment management world.

Why do people go along with this?

Imagine if climate change research were ridden with bad and poorly specified mathematics, and even when the math was well done, research program managers wildly exaggerated its results in order to argue for more research funds – and that those research managers pulled in incomes of millions of dollars a year.

Suppose that researchers found that temperatures might increase by 5°F in a thousand years if CO_2 emissions weren't stabilized within 100 years[2], but program managers relayed the results as having found that temperatures would increase by 7°F in 20 years if research didn't find ways to stabilize emissions in 5 years – and that spending at least one percent of GDP on climate change research was necessary to accomplish this.

If this were discovered, there would be a furor.

If the shenanigans so described really were going on in the field of climate research, it wouldn't mean that all research in the field was bad, or that there was no research that might be useful. But it might mean that too much of it was being done, and it would certainly mean that too much was being paid for it by those with a stake in the results – and that a massive scam was being perpetrated on those who pay.

Yet there has been no furor as purveyors of investment management and investment products draw on mathematical finance research, much, though not all, of which is ill-specified and poorly executed – and then wildly exaggerate the findings of that research in order to amass more assets under management or advisement, for percentage-of-assets fees.

This is the situation in the money management industry now. And what is shameful is that not only the industry itself, but the academic finance profession is a party to this scam.

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 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] I am grateful to Paul Bouchey for introducing me to Mr. Wong as well as for an extended email interchange.

[2] Not the real numbers.