

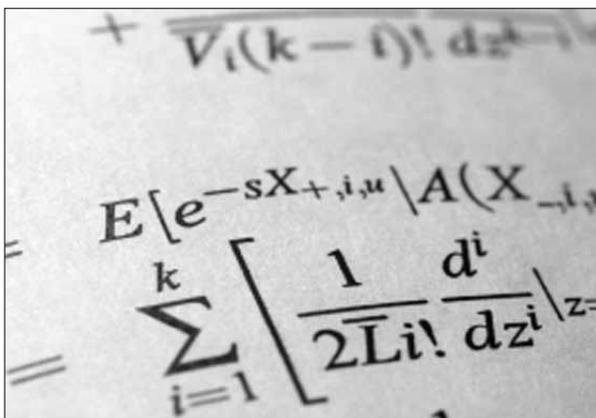


SOCIETY OF ACTUARIES

Article from:

# Risks & Rewards

February 2011 – Issue 57



## THOSE PESKY ARITHMETIC MEANS

By Dick Joss

It has been common in modern academic finance to assume that the arithmetic mean of a series of historical investment return results is the best estimate for next year's investment return on the particular type of investment. The best estimate for a multi-year forecast then becomes this arithmetic mean compounded for the forecasting period.

The rationale for using the arithmetic mean is quite straight forward. It is as if the actual historical results were written on little balls, the balls were placed in a giant rotating bin, and the year's investment return was obtained by drawing a ball out of the bin—just like a lottery drawing that might be seen on television. In this case the best estimate is, in fact, the arithmetic mean of the numbers on the balls in the bin. Sometimes this is referred to as a Monte Carlo simulation. All-in-all the approach seems fairly straight forward.

However, sometimes things that seem straight forward at first glance may turn out to be more complex. Determining arithmetic means is one of those times. For illustration data, consider the small company stock return data from *Ibbotson's Stocks Bonds Bills and Inflation 2008 Yearbook*. This data source provides historical investment returns for an 82-year period—Jan. 1, 1926 through Dec. 31, 2007.

If on Jan. 1, 1926, an investor had invested \$1.00 in the small company stock portfolio described in the *SBBI Yearbook*, the investment would have grown to \$15,091.10 by Dec. 31, 2007. The average annual increase in wealth over the 82-year period is 12.45 percent per year. However, instead of using this geometric mean return as a forecast, the academic finance community often cites the arithmetic mean of the 82 calendar years as the best estimate for next year's return. This number is a much higher 17.08 percent.

What goes unstated, however, is that the arithmetic mean of historical returns varies significantly based on the time period used for the calculation. If instead of using calendar years, Ibbotson had calculated the arithmetic mean using fiscal years ending in October the average return would have been 15.84

percent. If Ibbotson had used fiscal years ending in March, the arithmetic mean would have been 19.35 percent.

These differences could begin to impact investor behavior. Had investors been told that the arithmetic mean of small company stock returns was 15.84 percent instead of 17.08 percent, they might have been less inclined to invest in small company stocks. On the other hand, if the investors had been told that the historical arithmetic mean returns for small company stocks had averaged 19.35 percent instead of 17.08 percent, they might have been more inclined to invest in small company stocks.

Yet there is no particular reason for picking any one of the above arithmetic means over any of the others. The market itself does not rise up on each December 31, wave a big red flag, and shout: "Now is the 'right' time to measure stock market rates of return." The common practice of measuring these rates on a calendar year basis is just that—a common practice. As shown above, this common practice may be providing investors with information that might not be as full and complete as it could be.

Keep in mind one final time that every single one of the above numbers is based on the same set of data. In every case the original dollar invested on Jan. 1, 1926, is still worth \$15,091.10 on Dec. 31, 2007. The only difference between any of the arithmetic mean numbers is the selection of the fiscal year used to group the data.

### ANNUAL ARITHMETIC MEANS BASED ON SHORTER TIME FRAMES

The reason that the above different fiscal year calculations provided different arithmetic means is that for each of the calculations, a significant portion of the data is tossed away. For example, in each of the above Ibbotson small company stock return arithmetic mean calculations, only 82 different values of wealth growth were used (one for each fiscal year end), when 984 (one for each calendar year month) were available. Ninety-two percent of the available data was ignored!

In order to capture this extra data, it is possible to calculate an annual arithmetic mean rate of return in a very different way. Instead of ignoring 92 percent of the data, this process will use all 984 of the Ibbotson monthly wealth values. Using this process, one first takes the arithmetic mean or average of all the monthly returns provided in the *SBBI Yearbook*. This monthly average is then converted into (or expressed as) an annual equivalent rate of return. This process uses all of the listed Ibbotson data, not just a small fraction of it.

For the small company stock returns, the arithmetic mean of the monthly data is 1.3207 percent. This result is simply obtained by adding up all 984 monthly returns and dividing that total by 984. If an investment earns this rate of return each month for 12 months, the rate of return for the year will be 17.05 percent. This process is referred to as converting the monthly rate of 1.3207 percent into an annual equivalent rate of 17.05 percent. Using this process, one could say that the annual arithmetic mean rate of return for small company stocks was 17.05 percent. It is just that this annual arithmetic mean was based on monthly observations.

Some people may be uncomfortable referring to an “annual” arithmetic mean, when the data that was used to calculate it was monthly in nature. To help relieve some concern in this area, think about the answer to the question: How fast are you driving your car? If someone asks you this question while you are actually driving, your normal response is to look at your speedometer and report the speed as something like 45 miles per hour.

By reporting the speed as “45 miles per hour” you are not stating that you actually drove 45 miles, or that you actually spent one hour doing it. You were describing the speed at which you were driving at the time the question was asked using commonly understood terminology. In the same way, describing the monthly returns as 17.05 percent per year uses annual terminology to describe the arithmetic mean of monthly returns.

Because this new annual equivalent rate of return reflects observations from 984 data points (one for each month) instead of 82 data points (one for each year), there may be some comfort to the idea that this number is somehow more accurate because of the increased number of observations.

But anyone who watches a business news program on television is now keenly aware that data is available much more frequently than monthly. Often there is a little window on the television screen which shows the current Dow Jones Industrial Average (DJIA) value. This number changes every few seconds. If one calculated an annual Dow Jones Industrial Average arithmetic mean by just using month-end data, he or she would be throwing away millions of pieces of information as to how the DJIA changed over the course of a year.

This raises a question: Is it possible to calculate an arithmetic mean which captures as much of this data as possible? The surprising answer is yes! It is possible to calculate an arithmetic mean which reflects absolutely every single change in an investment such as the Ibbotson data or the DJIA, or even every single transaction involving an individual share of stock. This calculation process is described below.

### **“ARITHMETIC MEANS” FROM CONTINUOUS GROWTH**

Given that using monthly data to calculate an annual rate of return still means that millions of pieces of information will be lost, the idea of using even smaller time increments, such as a week, a day, or even an hour, might be considered as a possible time period for figuring an arithmetic mean that could then be converted into an annual equivalent rate of anticipated investment growth. The concept of calculating arithmetic means by using more data points and then converting the results to annual equivalents has a good feel about it. Somehow using more and more data gives one the feeling that the result may be more accurate than if just a limited number of observations of the wealth-growth are used.

CONTINUED ON **PAGE 18**

It turns out that figuring out the arithmetic mean of returns taken over a very small time frames is not as difficult as one might think. The first step is to consider how the value of a stock is determined and how this value changes over time. These values and changes are determined by actual investors, just like you.

For example, assume that Allen in Altoona decides to enter into a transaction to buy a particular stock at \$20 per share at exactly 1:00:00 on Oct. 1, 2008. This action by Allen then sets the value of the stock at \$20, until eight seconds later when Bob in Boston agrees to a transaction at \$20.50 per share. Thus, \$20.50 becomes the new value. At exactly 1:00:24 Cheryl in Chicago makes the next transaction for shares at a price of \$20.25. This process continues on and on and on. The stock market in action is truly a marvel to behold.

The first two transactions were eight seconds apart and the next two were 16 seconds apart. This illustrates a significant problem if one tried to calculate an arithmetic mean rate of return based on changes in share value calculated over one second time intervals. What would be the share value for times like 1:00:01, 1:00:02, and 1:00:03 when no transactions occurred?

The answer to this is really quite simple. Just estimate the share price assuming that there is a smooth line that connects every single one of the actual transactions! The smooth line just fills in the gaps between any two actual transactions. Using this new line, a person can estimate a value for the stock for any point in time, even times when no transaction occurred. This process just provides for a reasonable transition from one transaction to the next. This line is what is known as a “continuous and differentiable” curve, and it is easy to show that one can always be created in these types of situations.

While it is not common in modern finance to consider such curves when thinking about or calculating average returns, the concept is really quite logical. The curve passes through the data point for every single transaction of the stock! Not one single transaction is missed. The smooth curve is used only to fill in the gaps between any two successive transactions. And the

gaps are always filled with reasonable stock values, which are in between the two actual transactions at each end of the gap. Assuming that stock price changes may be described by a continuous and differentiable curve, it is now possible to calculate an annual arithmetic mean that reflects absolutely every transaction for the stock. This new annual arithmetic mean is just the annual equivalent of an arithmetic mean of historical results taken over very, very small time frames. Using all of the data, instead of just a minor portion of it, seems to provide additional assurance that one is getting the “right” answer.

The current method of calculating an annual arithmetic mean as shown in the *SBBI Yearbook* uses just 82 data items from the historical record. That’s all, just 82 separate pieces of data. If the process is expanded to monthly, then 984 pieces of information contribute to the calculation process. If the annual arithmetic mean is determined by taking daily snapshots of what was happening, then the calculation would have over 20,000 data elements. But even with daily numbers, the calculations would still miss all the market changes that occur within a single day.

If the calculations were carried out by taking the arithmetic mean of the rates of return over 1/10 of one second time intervals, the total volume of information would be over five billion individual observations of market behavior! It just seems natural that this result would have to be a far more accurate measure of the annual arithmetic mean than the relatively crude 82 element calculation process that is currently used.

There is an amazing answer that results from calculating the arithmetic mean or average of these five billion rates of return calculated over 1/10 of one second time intervals and then expressing this arithmetic mean as an annual rate of return. This amazing answer is that this process yields the geometric mean!

In other words, if the arithmetic mean calculation process used absolutely every single transaction, and then only filled in the gaps between transactions with a reasonable curve with

# // MAYBE IT IS TIME TO REPLACE ALL OF THOSE PESKY ARITHMETIC MEANS WITH A SIMPLE GEOMETRIC ONE. //

values between the two actual transaction values, the annual equivalent of the resulting arithmetic mean is the geometric mean. Using only a limited number of data snapshots from the historical record produces an array of arithmetic means, all of which are larger than the geometric mean.

## INVESTOR IMPACT

The above information now expands the possible range for arithmetic means for Ibbotson small company stock returns. This new range of possible arithmetic means runs from the geometric mean of 12.45 percent, which is based on billions of data point observations, to 19.35 percent, which is based on 82 observations taken as of each March 30. The value of \$10,000 invested for 20 years at these two rates ranges from \$104,518 to \$343,913. Dramatic differences in possible investment forecasts are beginning to show up.

It has been well-documented that workers who had planned on retiring with significant balances in their savings plans have been very disappointed. The cover article of the Oct. 19, 2009, issue of *TIME Magazine* was devoted to the hardships being

faced by workers caught in this problem. To the extent these workers are relying on overly optimistic projections based on traditionally-calculated arithmetic means, perhaps it is time to seriously consider more conservative investment forecasting concepts.

## CONCLUSION

The important point to be taken from this article is that any particular history of investment fund performance has lots of arithmetic means, and as shown above, these means can differ significantly. Any time you see a reference to “the” arithmetic mean of a data set, watch out! The asserted result is heavily dependent on which of the various arithmetic means was used. In particular, while the calendar year arithmetic mean of the 82-year historical small company Ibbotson stock return data is, in fact, 17.08 percent, using an alternative, yet still equally valid approach to calculating an arithmetic mean, this result could have varied anywhere from a low of 12.45 percent to a high of 19.35 percent. And it is particularly significant that the arithmetic mean based on the most data is the 12.45 percent number. Maybe it is time to replace all those pesky arithmetic means with a simple geometric one. **♣**



*Dick Joss, FSA, is retired. He can be contacted at [rrjoss@comcast.net](mailto:rrjoss@comcast.net).*