ENVESTNET PMC

Arithmetic, Geometric and Other Types of Averages

February 2019

Q | R | G
Quantitative Research Group

ARITHMETIC, GEOMETRIC AND OTHER TYPES OF AVERAGES

JĀNIS ZVINGELIS

1. Types of Average Returns

There are many types of averages used in the applied finance today. There are arithmetic and geometric averages; realized and forecasted averages; real and nominal averages, and others. All of these different types of averages are routinely used in the finance industry (usually when applied to returns), but often without referencing the type of average that is used, which, more often than not, leads to confusion. In this note I will define, compare and gives examples on how and when to use which type of average. Finally, I will also cover several ways of annualizing the average returns.

2. Arithmetic vs Geometric Average Return

Let's start with the two best known and widely used types of averages: arithmetic and geometric averages.

2.1. **Arithmetic.** The popularity of the arithmetic average can be, in part, explained by its simplicity. Let's use r_t to denote the rate of return during the time period t. Then the arithmetic average $\bar{r}_A$ over the investment horizon with T periods is calculated as follows

$$\bar{r}_A \equiv \frac{\sum_{t=1}^{T} r_t}{T} \tag{2.1}$$

The arithmetic average has a very nice property: the longer performance sample one uses to estimate the average, the more "confidence" one can have in the obtained result. To explore this property, I have calculated the arithmetic average return series for the (real) S&P 500 series from January of 1871 to present. In particular, in Figure 1 I have plotted the distributions of averages from non-overlapping time series, where the time series being used in forming the averages vary in their length. Thus, I am comparing the set of all arithmetic average returns over non-overlapping three-, 24-, and 120-month periods. While the overall average across these series is exactly the same (about 0.27 percent of real price return per month), the dispersion of these average returns is vastly different. The most important point to note is that the longer the data series used calculating the arithmetic average, the more

QRG, Envestnet, 35 East Wacker Drive, Suite 2400, Chicago, IL, 60601
Date: November 27, 2019.

concentrated around the overall mean the dispersion of the averages becomes. This phenomenon in statistics is called the "Law of Large Numbers (LLN)"[1], and it is, along with the simplicity of arithmetic average, the main reason for its popularity.

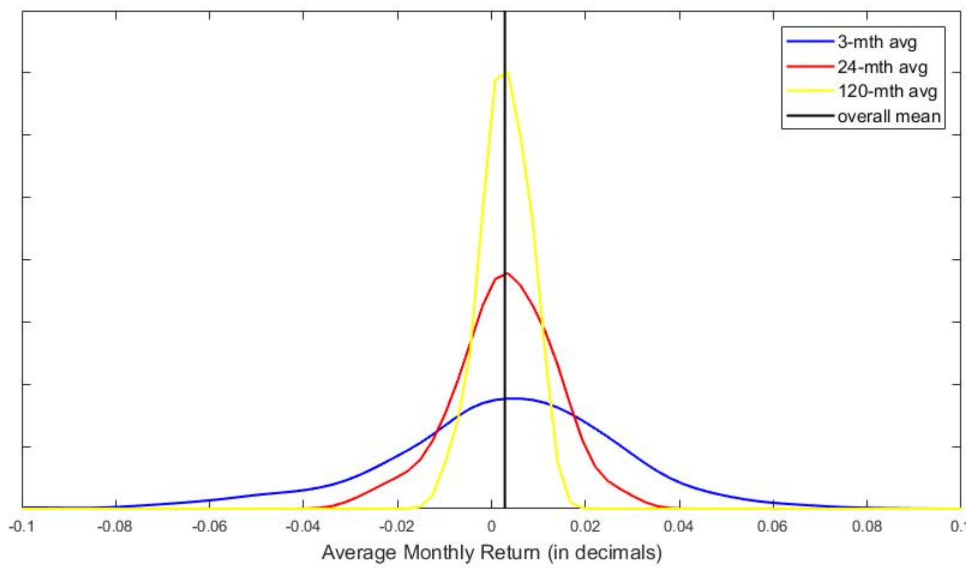

FIGURE 1. Plot of the distribution of the arithmetic average returns over various (non-overlapping) horizons for S&P 500 (1871/01 - 2017/09, real price returns in decimal format). *Source: www.multpl.com & QRG.*

Another application of arithmetic mean return is in forecasting. It turns out that the arithmetic average return can be used as a good predictor (due to the Law of Large Numbers that we mentioned earlier) of the so-called "expected return". "Expected return" can be thought of as an unknown average that we are trying to forecast. However, unlike an arithmetic average, where each observation is weighed by one divided by the number of total observations (see equation 2.1), the "expected value" construction uses the likelihood of a particular value happening as its weight.[2] Because of the property of the arithmetic average is a good predictor

[1]This property (i.e., the variance of the average shrinking, when there are more terms included in the average), when applied to returns, is sometimes mistakenly called the "time diversification". This mistaken argument then is used to justify why holding equity investments make sense, especially over the long term. The argument goes as follows: the longer the holding period (i.e., the more terms there are in the arithmetic average), the lower the variance/volatility of the investment. Hence, investing in higher volatility asset classes (such as equities) makes sense, especially over longer investment horizons. The problem with this argument is that the variance/volatility of the arithmetic mean return does indeed shrink (due to LLN), but *investors are usually not interested with the variances/volatilities of averages of returns, but rather with variances/volatilities of realized return series during some time period.* A more convincing argument behind "time diversification" (i.e., the idea for why holding equities over longer term rather shorter term make sense, was advanced by Campbell & Viceira (2005). It relies on the mean-reversion exhibited by equity returns through time.

[2]There is a lot of confusion regarding the concept of the "expected value" or "expected return", mostly due to the difference in the usage of the word "expected" in statistics and in everyday language. In statistics,

of the "expected return" (the larger the sample used to construct the arithmetic average return, the closer it is to the unknown "expected return"), the terms "arithmetic average return" and "expected return" are used interchangeably, especially in the applied finance.

Finally, while the arithmetic mean return is often used in summarizing a performance of an asset through time, it can result in misleading results, especially during high volatility periods. Table 1 gives a stylized example of a two-period example, where the starting amount and ending amount are the same ($100), but the arithmetic return is positive 25 percent. Clearly, this is a misleading result and points to one of arithmetic return's major weaknesses – inability to correctly summarize performance in high volatility environments. Another type of average return – the geometric average return – does a much better job in these situations, and we now turn to a closer look at this average.

Period	End-of-Period Balance	Period Return
0	$100	
1	$50	-50%
2	$100	100%
arithm.avg.		25%[3]
geom.avg.		0%[4]

TABLE 1. Performance summary with arithmetic and geometric average returns.

2.2. **Geometric.** Let's give a formula that defines the geometric average return. Let's use r_t to denote the rate of return during the time period t. Then the geometric average $\bar{r}_G$ over the investment horizon with T periods is calculated as follows

$$\bar{r}_G \equiv \left(\prod_{t=1}^{T} (1 + r_t) \right)^{1/T} - 1 \tag{2.2}$$

the word "expected" (as in "expected return") refers to a kind of arithmetic average that we just described. The word "expected" in everyday language is used with a connotation of "likely to happen" or "forecasted". As we will see in the next section, it turns out that an "expected wealth", for example, certainly cannot be expected to happen. Because people in finance are not always aware of this distinction, one often hears such phrases as "expected average return", which in statistics would mean (because "expected" again means a kind of an average) an "average average return" – certainly not what the user intended. A better term to use in finance would be "forecasted average return". Importantly, various types of different approaches, besides the arithmetic average return (and its close cousin "expected return"), can be used to construct the forecasted average return. It is therefore important to keep the distinction of "expected return" and "forecasted average return" in mind.

[3] $(-50\% + 100\%)/2 = 25\%$.

[4] $(100/100)^{1/2} - 1 = 0\%$. That is, ending balance divided by the beginning balance, all raised to the power of one over the number of holding periods, minus one.

While not as popular as the arithmetic average return,[5] geometric return has several very useful properties.

First, as demonstrated in table 1, geometric average return is unaffected by the volatility in the return, and therefore accurately captures performance through time. It therefore should be the go-to measure of summarizing financial performance (e.g., manager, account, etc) through time. In fact, the following approximate relationship (see appendix for the proof) holds between the arithmetic and geometric rates of return:

$$\bar{r}_G \approxeq \bar{r}_A - 0.5V(r_t). \tag{2.3}$$

What it says is what we observed in table 1: the higher the volatility of returns, the farther away the geometric average return will be from the arithmetic average return. Also, geometric return will always be no smaller than arithmetic return, and for cases of very small volatility (e.g., cash), the arithmetic and geometric average returns are essentially the same.

Second, geometric return is very useful in forecasting the likely future value of cumulative return or wealth, W_T, defined as $W_T = \prod_{t=1}^{T}(1 + r_t)$, where the sequence of r_t's, $t = 1, \ldots, T$, represents the random sequence of future returns. A common estimator of the future wealth, W_T, is the expected wealth (more on this in the next section), which can be forecasted by the following estimator:[6] It turns out that this estimator is consistently too high. Figure 2 shows long history of (real) price series (so, cumulative wealth) of S&P 500, along with two predictors of this series, one of which is based on arithmetic average return. Note how this estimator gets ever farther away (on the upside) from the series its trying to predict, as the forecast horizon increases. This highlights the problem of using the arithmetic return in the formula $(1 + \bar{r}_A)^T$ to predict W_T.[7]

It turns out that a much better predictor of future values of wealth/cumulative return, W_T is based on geometric average return: $(1 + \bar{r}_G)^T$. As demonstrated in figure 2, this estimator is consistently in the ballpark of the W_T value through time. Appendix also gives a proof that the estimator of future wealth, based of geometric average return ($(1 + \bar{r}_G)^T$), predicts the *median* value of W_T, and, thus, is more appropriate as a predictor, especially for long investment and prediction horizons.

Finally, a note of caution. I sometimes come across a notion that using geometric average returns in simulations is somehow "more conservative" and therefore presumably better. This is misguided on two fronts. First, as the equation 2.3 demonstrates, arithmetic and geometric average returns are related to each other in certain mathematical ways. For this reason they should be thought of as representing the same level of average return, with the only difference that they are *expressed in different units*. The same way as a 1,000 Japanese Yen burger does not become cheaper when translated into dollars (about 9.22 US Dollars at the time of writing this), the average return does not become lower when translated from arithmetic average return to geometric return – just the units are different. Second, behind most simulation engines is

[5] The main reason for this is that the geometric average return does not possess certain nice statistical properties (to be discussed in the next section), mostly because it is a product, rather than a sum, of returns. Products of returns are harder to handle, from a statistical point of view, than sum of returns.

[6] $E(W_T) = E(\prod_{t=1}^{T}(1 + r_t)) = \prod_{t=1}^{T}(1 + E(r_t))$, and an estimator for $E(r_t)$ is the arithmetic average $\bar{r}_A$. Thus, the estimator for W_T is $(1 + \bar{r}_A)^T$.

[7] Appendix also gives a mathematical proof of this and shows that the probability of the $(1+\bar{r}_A)^T$ exceeding W_T approaches one (certainty) as the investment horizon increases.

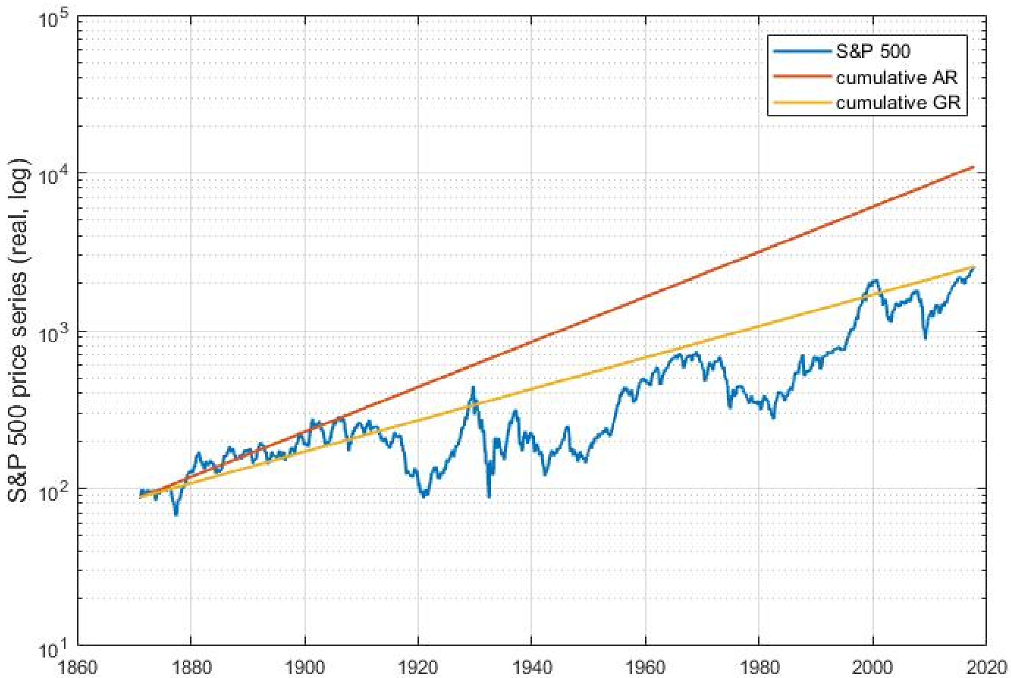

FIGURE 2. Plot of the (real) price series of S&P 500 (1871/01 - 2017/09, in logs). The figure also contains two cumulative return series – one aggregated at the average arithmetic return (about 0.27% monthly real return) and the other at average geometric return (about 0.19% monthly real return), where the averages are calculated over the time period 1871/01 - 2017/09. The starting value of the cumulative return series is the price level of S&P 500 in 1871/01. *Source: www.multpl.com & QRG.*

a statistical concept of "density function" – a way to assign probability to various levels of returns. One of the inputs to these "density functions" is always in terms of arithmetic average return. Thus, it would be a mistake to enter the geometric average return as an input into return simulators, because they are most likely expecting an arithmetic average return input. As a final point of advice – if a more conservative (lower) average return is desired as an input to the simulation engine, it is preferable to reduce the arithmetic average return as needed, while being cognizant of making this reduction, rather than switching to geometric average return units in a logically unsound effort to be "more conservative" with one's capital market estimates.

3. Nominal vs Real Average Returns

In this section we will discuss the inflation adjustment applied to returns. I often comes across the concept of "inflation adjusted" returns and usually it puzzles me, because I cannot figure out whether the inflation adjustment is converting future/past values to today's dollars or the other way around.

A much clearer approach is to calculate returns in either "real" (i.e., today's dollar) terms or "nominal" terms, where the latter refers to the return without any inflation adjustment. Based on this we can define a "nominal average return" (either arithmetic or geometric), $\bar{r}^n_A$ or $\bar{r}^n_G$, which refers to the arithmetic or geometric return in the units that it is observed, without any inflation adjustment. Also, we can define a "real average return" (either arithmetic or geometric), $\bar{r}^r_A$ or $\bar{r}^r_G$, which refers to the arithmetic or geometric return that takes the nominal return and factors out the inflation. Thus, if inflation is positive, as it usually is, the real return will be smaller than the nominal return.

$$\bar{r}^n \approx \bar{r}^r + i, \tag{3.1}$$

where i refers to the rate of annual inflation.

There are only a few financial parameters that can be observed in real terms, with most being quoted in nominal terms. Two that I am aware of are the yields on Treasury Inflation Protected Securities (TIPS) and dividend yields paid out by equities.

4. Annualization

In this section I will briefly cover the topic of annualization. Oftentimes, the rates of return or yields or other quantities (e.g., Sharpe ratios and volatilities) are given in the units underlying the calculations. For example, if I observe historical monthly nominal returns on S&P 500 index and calculate an arithmetic average return from these returns, the arithmetic average return will be given in monthly units. However, what if I want to use my calculated arithmetic average return (in monthly units) to investigate whether year 2018 was an above or below performance year for equities? To carry out this comparison, I will need to annualize my monthly arithmetic average return.

There are generally two types of approaches that can be used to annualize an average return: linear approach (based on the concept of geometric Brownian motion)[8] or compounding approach.

The linear approach to annualizing a return says that if we observe/caculate a return in units that are less than a year (say, months), then to get to an annual quantity we need to multiply the original return by the number of time units that fit in a year. For example, if we calculate the arithmetic average return based on monthly returns and then want to annualize it, we should multiply this average by twelve.

The compounding approach, which is better applicable to the geometric average return is given by the following formula:

$$\bar{r}_{G,a} = (1 + \bar{r}_{G,h})^N - 1,$$

where $\bar{r}_{G,a}$ and $\bar{r}_{G,h}$ refer to annual and higher frequency (e.g., monthly) geometric average rates of return, respectively, and N refers to the number of times that the higher observational frequency fits into a year.

[8]The assumption that a price series follows a geometric Brownian motion is used to obtain the famous Black-Scholes option pricing model. The geometric Brownian motion assumption also implies that $\log\left(\prod_{t=1}^{T}(1+r_t)\right) \sim N(\mu T, \sigma^2 T)$, where r_t is the return observed during a certain time period.

In addition to annualizing returns, we can also annualize standard deviations, variances, covariances, Sharpe ratios and alike. This task is slightly more complicated task, but we will be using the geometric Brownian motion assumption that we referenced above. That is, we will be using the assumption that $\log\left(\prod_{t=1}^{T}(1+r_t)\right) \sim N(\mu T, \sigma^2 T)$. In particular, we will use the assumption that variance of a compounded return log return $\log\left(\prod_{t=1}^{T}(1+r_t)\right)$ is equal to $\sigma^2 T$. That is, the variance is linear in time. Thus, if we calculate a variance that is based on monthly returns and want to annualized it, we would multiply it by 12 – the number of months in a year. Note that the comparable approach to standard deviation/volatility would be to multiply the monthly standard deviation by a *square root* of 12.

Furthermore, when annualizing volatility, especially in the context of calculating Sharpe ratios, we need to be mindful of serial correlation in returns (e.g., negative serial correlation is observed as mean-reversion in returns). This issue is highlighted in Lo (2002), who points out that in the absence of accounting for the serial correlation in volatility calculations, the Sharpe ratio is increasing in the units of standardization. That is, suppose that you calculated a daily return and volatility of a trading strategy and then proceeded to calculate its Sharpe ratio based on these daily values. If you then proceeded to convert this daily Sharpe ratio to a monthly Sharpe ratio or even an annual Sharpe ratio, using the approaches we reviewed above, then the monthly Sharpe ratio would be larger in absolute value than the daily Sharpe ratio, with the annual Sharpe ratio larger still.[9]

5. Application on the Platform

In this section we briefly touch upon the Capital Market Assumptions (CMAs) that are associated with Envestnet's platform. The CMAs that are loaded unto Envestnet's platform are given in *annual nominal arithmetic* average units. Also, the simulation engine that drives teh proposal expects that the CMAs be presented in annual nominal arithmetic average units. However, as described above, if needed, we can easily translate average returns given in other units (e.g., monthly real geometric average) to those in annual nominal arithmetic average units. As in the above sections, let's denote the nominal geometric average return by $\bar{r}_G^n$; real geometric return by $\bar{r}_G^r$; nominal arithmetic average return by $\bar{r}_A^n$; real arithmetic average return by $\bar{r}_A^r$. Then the flowchart below summarizes the relationships given above:

$$\begin{tikzcd}
\bar{r}_A^n \arrow[r, leftrightarrow, "eq.\ 3.1"] \arrow[d, leftrightarrow, "eq.\ 2.3"'] & \bar{r}_A^r \arrow[d, leftrightarrow, "eq.\ 2.3"] \\
\bar{r}_G^n \arrow[r, leftrightarrow, "eq.\ 3.1"'] & \bar{r}_G^r
\end{tikzcd}$$

[9]The reason for this is that the numerator (the average return) of the Sharpe ratios is multiplied by the number of periods in a year, while the denominator (the volatility) is multiplied by the number of periods in a year *squared*. Thus, the original Sharpe ratio (say, daily) is multiplied by a square of number of units of original observation that fit into the period to which you want to standardize the Sharpe ratio (e.g., 20 business days in a month; about 63 business days in a quarter, and about 252 business days in a year). If we don't account serial correlation, in a Sharpe ratio annualization exercise the denominator always grows slower than the numerator, when carrying out the annualization, which leads to misleading Sharpe ratio estimates.

For example, to translate the nominal arithmetic average return, $\bar{r}_A^n$, into real geometric average return, $\bar{r}_G^r$, all we need to do is apply equations 2.3 and 3.1. While the order in which these translating equations (arithmetic to geometric vs nominal to real) are applied matters, in principle, one can safely apply them disregarding the order, as the error in this approximation will not be meaningful.

6. Appendix

6.1. **Arithmetic-to-Geometric Translation Formula.** Let's assume that a gross random return, $1 + r_t$, has a log-normal distribution for all $t = 1, \ldots, T$. What this means is that the natural log of $1 + r_t$ is distributed as a normal random variable with parameters μ and σ: $\log(1 + r_t) \sim N(\mu, \sigma^2)$, which then results in the following formulas for the expected (gross) returns, $E(1 + r_t)$, and variances, $V(1 + r_t)$:

$$E(1 + r_t) = e^{\mu+\sigma^2/2} \tag{6.1}$$

$$V(1 + r_t) = \left(e^{\sigma^2} - 1\right) \cdot e^{2\mu+\sigma^2} \tag{6.2}$$

Solving for μ and σ in the above system of equations we obtain the following solutions (see, for example, de La Grandville (1998)):

$$\mu = \log(1 + E(r_t)) - 0.5 \log\left(1 + \frac{V(1 + r_t)}{(1 + E(r_t))^2}\right) \tag{6.3}$$

$$\sigma^2 = \log\left(1 + \frac{V(1 + r_t)}{(1 + E(r_t))^2}\right) \tag{6.4}$$

Now, let's derive the relationship between the geometric average return and the arithmetic average return:

$$\begin{aligned}
1 + \bar{r}_G = & \left(\prod_{t=1}^{T}(1 + r_t)\right)^{1/T} && \text{(by definition of } \bar{r}_G \text{ in eq. 2.2)} \\
= & \; e^{\log\left(\left(\prod_t^T(1+r_t)\right)^{1/T}\right)} \\
= & \; e^{\sum_t^T \log(1+r_t)/T} \\
\approx & \; e^{E(\log(1+r_t))} && \text{(by the Law of Large Numbers)} \\
= & \; e^{\mu} && \text{(by an assumption } \log(1 + r_t) \sim N(\mu, \sigma^2)) \\
= & \; e^{\log(1+E(r_t)} \cdot e^{-0.5\sigma^2} && \text{(by eq. 6.3 \& 6.4)} \\
\approx & \; (1 + E(r_t)) \cdot (1 - 0.5\sigma^2) && \text{(by Taylor series: } e^x = 1 + x + x^2/2! \ldots) \\
\approx & \; 1 + E(r_t) - 0.5\sigma^2 \\
\approx & \; 1 + E(r_t) - 0.5V(r_t) && \text{(by eq. 6.4 and} \\
& && \text{Taylor series: } \log(1 + x) = x - x^2/2 - \ldots) \\
\approx & \; 1 + \bar{r}_A - 0.5V(r_t) && \text{(by the Law of Large Numbers)}
\end{aligned}$$

Therefore, we obtain the following conversion formula between the arithmetic and geometric mean returns:

$$\bar{r}_G \approx \bar{r}_A - 0.5V(r_t) \tag{6.5}$$

Note that this formula is an approximation and holds best, when A. the number of elements in the average is large; B. individual returns tend to be Normally distributed; C. σ, $E(r_t)$, and $V(r_t)$ are relatively close to zero. In most cases, these assumptions hold, at least partially, and the approximation in equation 6.5 works very well.

6.2. **Predicting Wealth with Expected Wealth.** In this section we will demonstrate that the probability of cumulative return (or wealth) exceeding the expected value of wealth/cumulative returns tends to zero as the investment horizon increases. Also, see sections above as well as figure 2 for a plot supporting this claim.

Let's define W_T to be the cumulative return/wealth at time T from investing one dollar T periods ago: $W_T = \prod_{t=1}^{T}(1+r_t)$. Wealth, W_T, when standing at time zero, is a random variable, since the sequence of future returns $(r_t, t = 1, \ldots, T)$ is random. Oftentimes, researchers are interested in predicting the value of W_T, and an estimator, which looks very reasonable on the face of it, is the expected wealth:

$$\begin{aligned} E(W_T) = &\quad E\left(\prod_{t=1}^{T}(1+r_t)\right) \\ = &\quad \prod_{t=1}^{T}(1+E(r_t)) \qquad \text{(by IID assumption of } r_t\text{'s)} \end{aligned}$$

This estimator seems reasonable, because as we mentioned in previous sections $E(r_t)$ is the best (in some certain mathematical sense) predictor of r_t. It therefore may stand to reason that the expected wealth might be a good predictor for W_T. However, as it turns out, expected wealth is not a good predictor of W_T, and, in fact, the probability of the expected wealth, $E(W_T)$, exceeding the realized wealth, W_T, converges to 1 (certainty) as the investment horizon increases to infinity. In other words, if the investment horizon is long, the realized wealth will almost always fall below the forecasted wealth number of expected wealth.

Why such a seemingly nonintuitive result? What is happening is that W_T is a random variable that is constrained from below by zero (i.e., can't lose more than what's invested), which means that W_T is a skewed distribution, in this case skewed to the positive side. This means that the expected value of this distribution, $E(W_T)$, will always be no smaller (and, in fact, much larger, the more skewed the distribution, i.e., the longer the investment horizon) than the median of this distribution.

Next, let's prove the claim that the probability of expected wealth exceeding the wealth tends to zero with certainty as the investment horizon grows:

$$
\begin{aligned}
P\left(W_t \geq E(W_T)\right) &= P\left(\prod_{t=1}^{T}(1+r_t) \geq E\left(\prod_{t=1}^{T}(1+r_t)\right)\right) && \text{(by definition of } W_T\text{)} \\
&= P\left(\left(e^{\sum_{t=1}^{T}\log(1+r_t)/T}\right)^T \geq \left(e^{\log E(1+r_t)}\right)^T\right) && \text{(by IID assumption on } r_t\text{'s)} \\
&= P\left(\sum_{t=1}^{T}\log(1+r_t)/T \geq \log E(1+r_t)\right) \\
&= P\left(\sum_{t=1}^{T}\log(1+r_t)/T - E\left(\log(1+r_t)\right) \geq \right. \\
&\qquad \left. \underbrace{\log E(1+r_t) - E\left(\log(1+r_t)\right)}_{C}\right) && \text{(} C > 0 \text{ by Jensen's inequality)} \\
&\leq P\left(|\sum_{t=1}^{T}\log(1+r_t)/T - E\left(\log(1+r_t)\right)| \geq C\right) \\
&\leq \frac{1}{C^2}E\left(|\sum_{t=1}^{T}\log(1+r_t)/T - E\left(\log(1+r_t)\right)|^2\right) && \text{by Chebychef's inequality)} \\
&= \frac{1}{C^2}Var\left(\sum_{t=1}^{T}\log(1+r_t)/T\right) \\
&= \frac{1}{C^2}Var\left(\log(1+r_t)\right)/T && \text{(by IID assumption on } r_t\text{'s)} \\
&\xrightarrow{T\to\infty} 0
\end{aligned}
$$

Note that, since the expected return, $E(r_t)$, is predicted by arithmetic mean return (by LLN), the above claim that the expected wealth is a poor predictor of future wealth also applies to the epxpected wealth estimate, where $E(r_t)$ is replaced with the arithmetic mean return, $\bar{r}_A$.

6.3. **Predicting Wealth with Median Wealth.** In this section we explore the claim that a better predictor, compared to the expected wealth analyzed in previous section, of future wealth might be the median wealth.[10] Importantly, median wealth is related to the geometric mean, and we will demonstrate this relationship precisely in this section. Also, see sections above as well as figure 2 for a plot supporting this claim.

Let's assume that the return has a log-normal distribution, i.e., $\log(1 + R_t) \sim N(\mu, \sigma^2)$. Then the cumulative return/wealth ($W_T = \prod_{t=1}^{T}(1+r_t)$) also has a log-normal distribution:[11]

[10] We use results from Hughson, Stutzer, & Yung (2006).

[11] We are assuming the r_t's are IID.

$\log(W_T) \sim N(\mu T, \sigma^2 T)$. Given this, we can write the expression for W_T as follows:

$$W_T = \prod_{t=1}^{T}(1+r_t)$$

$$= e^{\mu T + Z \cdot \sigma \cdot \sqrt{T}}, \tag{6.6}$$

where $Z \sim N(0,1)$. Given the above, it's easily seen that the median wealth can be calculated as follows:[12]

$$\begin{aligned} med(W_T) &= e^{\mu T} && \text{(set } Z = 0\text{, its median, in eq. 6.6)} \\ &\approxeq e^{\left(\sum_{t=1}^{T}\log(1+r_t)/T\right)T} && \text{(by LLN)} \\ &\approxeq \left(\prod_{t=1}^{T}(1+r_t)\right)^{T/T} \\ &\approxeq (1+\bar{r}_G)^T && \text{(by definition of geometric mean return in eq. 2.2)} \end{aligned}$$

Thus, median wealth (or cumulative return) is approximately equal to a geometric return that is accumulated to the power equal to the investment periods.

References

Campbell, John Y., & Luis M. Viceira, 2005, The term structure of the risk-return tradeoff, *Financial Analysts Journal* 61(1), 433–495.

de La Grandville, Olivier, 1998, The long-term expected rate of return: Seting it right, *Financial Analysts Journal* 54:6, 75–80.

Hughson, Eric, Michael Stutzer, & Chris Yung, 2006, The misuse of expected returns, *Financial Analysts Journal* 62(6), 88–96.

Lo, Andrew W., 2002, The statistics of Sharpe ratios, *Financial Analysts Journal* 36–52.

[12] Turns out that the following key expression for a median also holds approximately when the return is not assumed to follow a log-normal distribution (Hughson et al. 2006).