Terminal Wealth Under Leverage

Erik Forseth[1], Tom Feng[2], Ed Tricker[3]

Abstract
In this research note, we explore the effect of volatility on compound returns. We start with the so-called *volatility drag*, which refers to the well-known – but sometimes counterintuitive – difference between the arithmetic and geometric mean returns of a portfolio. We then present some results that show the effect of leverage on terminal wealth.

Keywords
Volatility, Compounding, Leverage

[1] *Senior Quantitative Research Analyst*
[2] *Co-Director of Quantitative Research*
[3] *Chief Investment Officer, Quantitative Strategies*

1. Volatility Drag

Suppose we have a \$100 portfolio. As we watch its value fluctuate each day, we notice an asymmetry. One day the portfolio falls by 10%, and the next, it rises by 10%. However, we haven't wound up where we started, as our ending balance is now

$$\$100 - \$\,(0.1 \times 100) + \$\,(0.1 \times 90) = \$99.$$

This phenomenon has been referred to as *volatility drag* in reference to the way that losses must be made up by even bigger gains to tread water [Spitznagel (2017)]. Of course, there's nothing nefarious about this situation – we've simply described a basic fact of mathematics, that *geometric* returns are not equivalent to *arithmetic* returns[1].

Next, imagine that we decide to take on more risk, and move our money into a similar investment with 3x volatility. We've just seen that consecutive $\pm 10\%$ moves left us with 99% of our initial investment, for a 1% total loss. We might naively expect then that our 3x levered fund will lose 3% over the same sequence of moves. This is not the case! Each 10% move in the underlying corresponds to a 30% move in the new investment, and hence consecutive 30% moves result in a $1.3 \times 0.7 = 0.91$ final multiplier (or a 9% loss).

Despite this example, a key point of this note is that higher volatility is not inherently a bad thing. Importantly, we should not use the term leverage synonymously with "adding volatility," as leverage scales both the mean as well as the variance [2]. Indeed, leverage can dramatically improve the terminal wealth distribution for the long term investor. However, it can also have interesting and non-obvious effects on the shape of that distribution. While there does come a point beyond which leverage is counterproductive, that point depends on the Sharpe ratio of the investment and is often higher than one might think.

[1] The arithmetic mean return in our example is $\sum_i^n r_i = 0$, while the geometric mean return – which is relevant for compounding – is $\sqrt[n]{\prod_i^n (1 + r_i)} - 1 = -0.005$.

[2] Unlike in the example above, where the second case had higher variance for the same mean.

2. A More General Example

We've just described a scenario for which volatility resulted in a meaningful difference in performance. Now we generalize by considering the effect of volatility on the distribution of outcomes for a long-horizon investor.

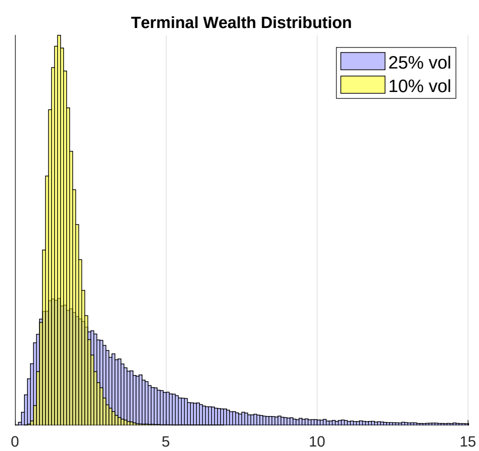

Figure 1. Ten year terminal wealth distributions for $SR = 0.5$ at two volatility targets. Each histogram shows the relative frequency of terminal wealth values over many independent realizations of the stochastic returns process.

Suppose that we have access to some investment whose Sharpe ratio SR is known. That is, assume we know the average annualized return μ of the investment relative to its

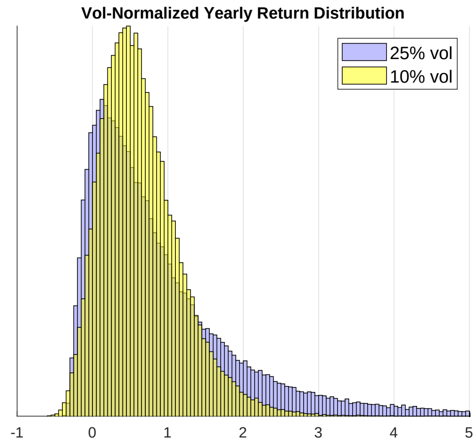

Figure 2. Distribution of vol-normalized annual returns for a 0.5 Sharpe strategy.

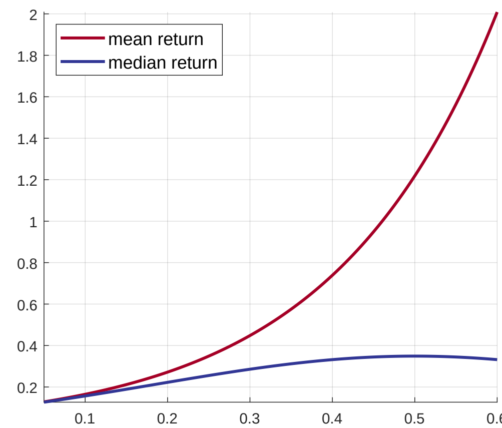

Figure 3. Mean and median annualized returns for a 0.5 Sharpe strategy as a function of volatility. Note that the median peaks at $\sigma = 0.5$.

annualized volatility σ. It's reasonable to ask how changes in σ should be expected to influence the range of outcomes for this investment. Figure 1 shows the distribution of terminal wealth values after a ten-year period for an investment with $SR = 0.5$, at two different annualized volatility targets.

We clearly see that the investment with the higher volatility tends to lead to a higher terminal wealth on average.

It's also interesting to ask how each investment fares on a 'volatility-normalized' basis. That is, what kind of yearly compound return *per unit of annual volatility* can we expect? Figure 2 shows this result.

We observe that the higher volatility investments has a higher vol-normalized annual return (9.96% vs 6.49%), predominantly due to a large right-tail. These two pictures suggest that, except when we are strongly risk averse or do not have a long-term view, we should generally prefer the distribution of outcomes indicated by the 25% vol investment.

3. Distribution of Terminal Wealth

It is somewhat remarkable how leverage seems to alter the terminal wealth distribution. Increasing the volatility has increased the probability of outlier realizations, as we can see from the growing right tail. It's also clear from this tail that the *mean* of the distribution increases with volatility[3]. On the other hand, it doesn't seem so clear that the distribution's *median* will follow suit. The median is a more robust measure of central tendency. It is perhaps a better representation of what the typical investor should expect to realize, particularly over short-term investment horizons.

By extending our experiment's scope, we see that as volatility increases, so does the gap between mean and median outcomes. Figure 3 shows the mean and median of the annualized compound return as a function of volatility. The mean grows exponentially, but the median displays a peak. Beyond a certain point, higher volatility pushes the median return toward zero! To further understand what's happening here, we can describe the mathematics behind these results.

3.1 A Model for Returns

We start with the standard assumption of a lognormal distribution for wealth; that is, the logarithm of wealth follows a normal distribution. Compound percentage returns are modeled by a discrete-time version of the geometric Brownian Motion process:

$$\frac{dS}{S} = \mu dt + \sigma dW,$$

where S represents wealth, μ the average return, σ the volatility, and W a $\mathcal{N}(0,1)$ a white noise process. The Sharpe ratio of the process is $SR = \mu/\sigma$. Given some initial investment S_0, the terminal wealth at time t is given by

$$S_t = S_0 \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma W_t\right].$$

We're now equipped to understand the divergence between the mean and median that we'd previously observed [Rabault (2014)]. The compound return after T periods is

$$\begin{aligned} R_T &= \prod_t^T \exp\left(\mu - \frac{1}{2}\sigma^2 + \varepsilon_t\right) \\ &= \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right) T + \sum_t^T \varepsilon_t\right]. \end{aligned} \quad (1)$$

[3] As we'll see shortly, the mean realization technically depends only on μ. However, for fixed Sharpe ratio, increasing σ is equivalent to increasing μ. Once again, we're really talking here about increasing *leverage*, i.e. scaling both μ and σ.

The mean of the compound return is

$$E\left[R_T\right] = \exp\left(\mu T\right).$$

For a fixed Sharpe ratio, μ scales with σ as $\mu = \sigma SR$. Hence, the mean compound return grows exponentially with σ, as we see clearly in Figure 3. Because the exponential function is convex, we see that the mean is lifted by a small number of large positive outcomes.

What about the median outcome? The median realization of the noise term $\sum_t \varepsilon_t$ is zero, and so the compound return follows

$$\text{med}(R_T) \sim \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right) T\right].$$

It should be immediately apparent that while volatility has no impact on the mean compound return, it can lower the median compound return.

3.2 Optimal Leverage Under Logarithmic Utility

In the limit that margin requirements are small relative to total capital, it is convenient to specify a desired risk level. Applying leverage amounts to choosing a multiplier k by which to scale the returns and volatility, $\hat{\mu} = k\mu$ and $\hat{\sigma} = k\sigma$. Under a certain choice of *utility function*[4], there is a simple connection between the optimal k and the Sharpe ratio of the investment. The *Kelly criterion*[5] invokes a *logarithmic utility* function; that is, the investor should aim to maximize the growth rate of the logarithm of her terminal wealth. The log wealth of the levered investment is given by

$$L_t \equiv \log(S_t) = \log(S_0) + \left(k\mu - \frac{1}{2}k^2\sigma^2\right) t + k\sigma W_t,$$

and its growth rate is then

$$G_t \equiv \frac{dL_t}{dt} = k\mu - \frac{1}{2}k^2\sigma^2.$$

Maximizing G_t with respect to leverage k, it is easy to find

$$k_{\text{opt}} = \frac{\mu}{\sigma^2} = \frac{SR}{\sigma}.$$

That is, the optimal leverage is equal to the Sharpe ratio divided by the risk. Note that when $\sigma = SR$, we find $k_{\text{opt}} = 1$; in other words, no leverage is needed. This implies that the "optimal volatility" for a portfolio is theoretically equal to its Sharpe ratio. For example, our strategy in the beginning had a Sharpe ratio of 0.5 and a volatility of 10%. We now see that the optimal leverage for this strategy is 5x, i.e. running the strategy at 50% volatility. This coincides with the peak in Figure 3.

Logarithmic utility is known to lead to rather aggressive policies. Hence, there exists a class of "fractional Kelly" strategies [Ziemba (2003)] for which the investor allocates a fraction f of his portfolio to a log utility (Kelly) strategy, and the remainder $1 - f$ to some risk-free asset.

[4] In order to optimize a portfolio, the investor needs to specify some criterion by which an outcome can be considered optimal. A utility function is a mathematical description of the investor's preferences about risk versus reward.

[5] Described by John Kelly in 1956 [Kelly Jr. (2011)], the Kelly criterion is commonly thought of in the context of gambling. Say there is a biased coin which lands on heads with probability p. A bettor with knowledge of p should wager $2p - 1$ on heads at each coin flip in order to maximize her long-term winnings. (A negative bet is a bet on tails.)

4. Conclusion

For the long-term investor, leverage can be a useful tool for maximizing wealth. The investor should target leverage by weighing his/her risk appetite versus his/her expectations about the return of the investment. The investor should also be aware that leverage can have profound consequences for the distribution of terminal outcomes. The average compound return – which always grows with increasing risk – can be misleading, and the investor should also weigh other information, such as the median.

References

J. Kelly Jr. A new interpretation of information rate. In *The Kelly Capital Growth Investment Criterion*, chapter 3, pages 25–34. 2011.

G. Rabault. The volatility-drag. https://investmentmath.com/finance/2014/03/04/volatility-drag.html, 2014. Accessed: 2020-08-04.

M. Spitznagel. Commentary: Not all risk mitigation is created equal. *Pensions& Investments*, 2017. Accessed: 2020-03-03.

W. Ziemba. The stochastic programming approach to asset-liability and wealth management. 01 2003.

Legal Disclaimer