NORTHFIELD | modeling financial markets worldwide

Estimating an Investor's Volatility/Return Tradeoff: The Answer is Always Six

Dan diBartolomeo

Webinar March 2020

Introduction

- A constant but vexing question among investors is “How much incremental return must I expect to justify a particular increase in the risk of my portfolio?”.
- In this webinar, we will derive a “rule of thumb” that efficiently describes the optimal return-volatility tradeoff parameter for a wide range of practical cases.
- Since seminal paper of Levy and Markowitz (1979, *Journal of Finance*), the most popular way to describe investor objectives has been the mean-variance utility function.
- The alternative objective of Sharpe Ratio has also be shown to be equivalent to mean-variance under certain circumstances as described by DeGroot and Plantinga (2001, *Journal of Performance Measurement*).

The Problem with Mean-Variance

- However, many investors find statistical variance an unintuitive measure and so prefer to think of the tradeoff between return and risk with the unit of risk being standard deviation
  - Or a scalar of standard deviation such as VaR or CVaR, see Cornish and Fisher (1937)
- Most investors are simply unable to numerically express their mean variance tradeoff parameter ("lambda" or it's percentage reciprocal which Northfield terms "RAP") with any confidence.
  - I know that if I had asked my grandmother about her "mean-variance risk aversion" she would have slapped me.
- While the properties of any particular mean-variance objective can be mapped to mean-standard deviation using the "chain rule" of calculus, a different and pragmatic approach to this question can be derived from the methodology in Wilcox (*Journal of Portfolio Management*, 2003).

Extending the Discretionary Wealth Hypothesis

- In this process, we assume that the investor has chosen their current level of portfolio risk so as to bound a "worst case scenario".
- From this boundary condition we can infer the mean variance risk tolerance as a *scalar function of the current portfolio expected return and volatility*.
- Algebraic simplification leads to the conclusion that for a wide range of situations and asset allocations, the optimal tradeoff parameter between *incremental expected return and incremental volatility is typically about one sixth.*
  - For example, a 2% increase in portfolio volatility is reasonably justified by a .33% (2/6) increase in expected return The method works identically across absolute risk, tracking error, or any desired blend of the two risk measures.

Growth Optimal Investors

- Let's start with a "growth optimal" investor whose *only objective* is maximize their long term geometric mean return. The usual objective function is

$$U = R - \lambda * \sigma^2$$

For a growth optimal investor who only cares about maximizing the future geometric mean return, λ =.5 or ½ assuming all units are in decimals. My preference is to remove the potential for $\lambda = 0$ so I'll rewrite this with the tradeoff parameter in the denominator and also convert to percentages

$$U = R - \sigma^2 / T$$

For the growth optimal case T = 200.

- The question is what is the appropriate <u>risk tolerance</u> T for a more risk averse investor, as almost all are.

The Maximum Loss Fraction

- Now let's work through a simple example, where our existing portfolio has R = 8 and σ = 12
  - Implicit in choosing a portfolio of a particular risk (e.g. 12%) is the idea that I don't want to put all my money at risk, just some of it (Wilcox's Discretionary Wealth Hypothesis, JPM, 2003). A reasonable expectation for the **maximum loss fraction** of the portfolio value would be something like

$M = (Z * \sigma) - R$, where Z is your choice Z-score of the worst case scenario (say something like 3.5)

$M = 3.5 * 12 - 8 = 34$

- So we're only willing to put 34% (.34) of the portfolio at risk (implicit T= 200) which means that the other part of the portfolio must be riskless (implicit T = 0)

$T = .34 * 200 + (1 - .34) * 0 = 68$

Getting To Approximately Six

- If we divide T through by σ, we obtain 68/12 = 5.67 (about six). So now we can express our objective as

  $$U = R - \sigma^2/(5.67 * \sigma)$$

  Which simplifies to

  $$U = R - \mathbf{(1/5.67)} * \sigma$$

  So our tradeoff between return and standard deviation is about 1/6. For a broad range of empirical cases, the 1/6 relationship holds rather nicely.

- For tracking error cases rather than absolute volatility, I tend to use Z = 3, R = 0, which **renders exactly 1/6.**

On the Value of Portfolio Construction

- A recent *Journal of Portfolio Management* publication of my colleague Jason MacQueen addresses the issues of portfolio construction for smart beta ETF portfolios
  - https://jpm.pm-research.com/content/46/2/64.abstract
  - This material was also presented at Boston QWAFAFEW in February

- Within the paper, the concept of the mean-standard deviation tradeoff appears in formulae using the Greek letter Ψ

- What is new in this presentation is that this coefficient takes a value of approximately 1/6 for a broad range of portfolio construction decisions.

Optimizing With Higher Moments

- We use a more elaborate version of this trick “trick” in the latest version of the Northfield Optimizer that adjusts for higher moments of asset return distributions.
  - If higher moments are present the *effective volatility* of each asset may be higher or lower.
  - Depending on degree of diversification of the portfolio and time horizon of the risk estimate you can calculate what portion of the increments (up or down) in volatility at the asset level work through to the portfolio level.
  - Once you have that you can use the (about 1/6) relation to adjust the input alphas (higher risk means you want a smaller bet so you move the alpha toward zero, lower risk means you’re willing to take a bigger bet so you move the alpha away from zero).
  - The alpha increments are scaled to the magnitude of the increments in risk caused by the higher moments. This solution allows us to keep the economic intuition of mean-variance while sensibly allowing for higher moments.

Inference About the Coronavirus Emergency

- In our recent research release about the coronavirus emergency, we use the "six" rule to infer how long investors believe the heightened risks of the pandemic conditions will persist.
- To assess very short term risk, Northfield has maintained our US Short Term risk model *since 1997*. In this approach a statistical factor model is adjusted daily for changes in the implied volatility of options traded on equities in the US.
  - A mathematical process then maps the day to day changes in security level volatility across the factors of covariance so that changes in market conditions are applied to all securities, not just equities on which options are traded. Security coverage includes all non-US equities traded in on US exchanges in ADR form, so the model does the majority of most large publicly traded firms globally.
  - Mathematical details of the model are presented in https://www.northinfo.com/Documents/534.pdf which was subsequently published in 2005. Although volatility levels are presented in the usual annualized units, the time horizon for the risk forecast is the next trading day.

Recent Volatility of the S&P 500

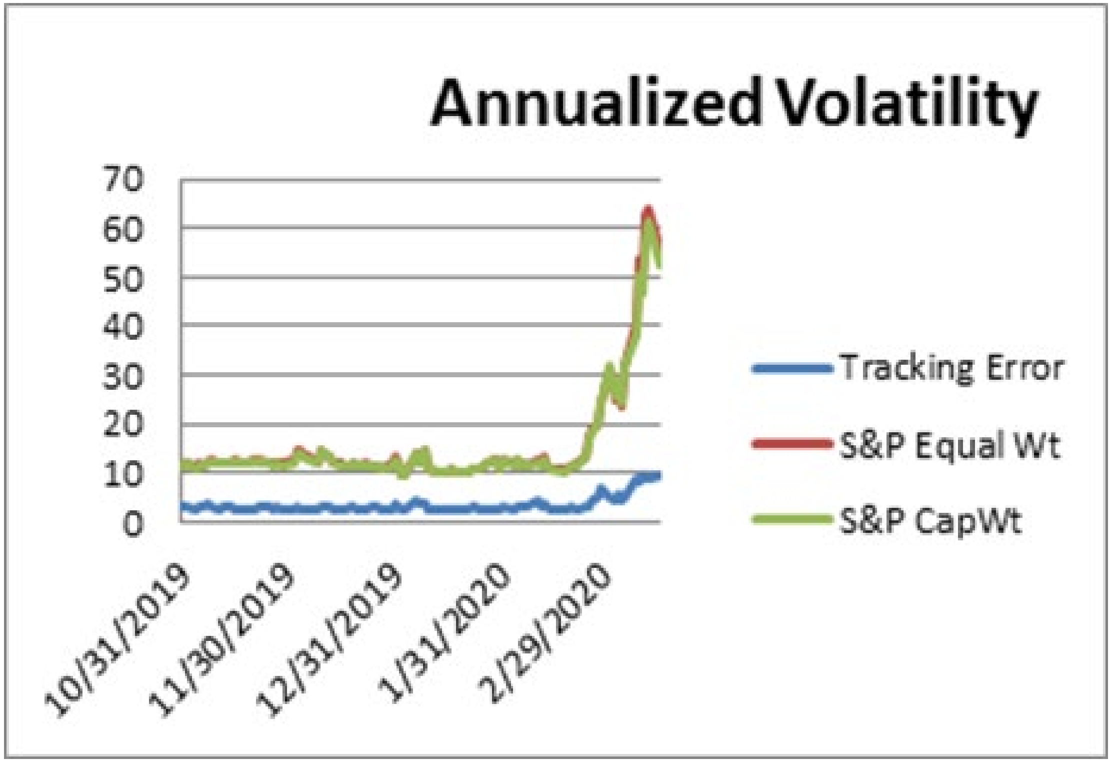

The chart above shows the expected volatility of the S&P 500 on an equal weighted basis, the S&P 500 on the conventional capitalization weighted basis, and the tracking error of the two portfolios. The time period is from the end of October 2019 through to March 16th, 2020, at a daily frequency.

Change in Volatility November 17 – March 16

- The chart shows that the volatility values for both series increased roughly five fold over the sample period, starting around 12% in November.
- Peak annualized volatility values of over 60% were observed on March 13th but decreased to 52% for cap weighted index and 55% on the equal weighted index by March 16th.
- As is common in crisis periods, correlations have increased with the *expected* correlation of the two portfolios going from .968 at the start to .984 at the end.
- Over the period of the sample, the expected volatility of the S&P 500 increased by roughly 40 percentage points. We assert that rational investors would respond by increasing their required rate of return by 6.7% per annum (40/6) for the period of heightened risk. The denominator scalar of about six is derived from the implied risk boundary of an investor's original risk level.

Impact on Market Valuation

- Let's assume that at the start of the sample period investors believed that the S&P 500 was fairly valued with an expected (and required) total return of 6%. Of this return, 2% was then known dividend yield and 4% expected growth rate. In a Gordon dividend discount model we obtain:
- P(start) = 1 / (.06 -.04) or $50 per $1 of dividends (2% yield)
- If we now add increment the required return for the heightened expected risk we get:
- P (now) = 1 / (.127 - .04) = $11.49
- So if investors believed the massively heighted risks would be permanent, the S&P 500 should have fallen by 78%, not the observed roughly 20%.

Long Term Pricing of Risk

- To get the 20% drop to make sense, the *long term* required rate of annual return must increase from 6% to 6.5%.
- Using our "rule of six", this means that investors are pricing the market as if the long term expected volatility increased by 3% (.5 * 6) or from 12% to 15%.
- If we assume that the median survival time for a traded firm in the USA is 20 years (consistent with diBartolomeo, *Journal of Investing*, 2010) we can infer how long investors expected the heightened risks to last.
  - The underlying model in this paper is based on the contingent claims concept from Merton (JOF, 1974). Essentially it argues that stockholders have two options that lenders don't have. One is a call option on the assets of the firm which can be exercised by paying off the firm's debt. The other option is a put option associated with the limited legal liability of shareholders. If the assets of a firm fall sufficiently, the shareholders can walk away transferring the assets to lenders.

Implied Length of the Pandemic

- Using the mathematical property that variances are additive, we can set the problem up as follows:

$$(15)^2 = X * (52)^2 + (1-X) * (12)^2$$

$$X = .031$$

Assuming the 20 year horizon from the JOI paper, we get that the implied length of the pandemic is .031 * 20 years or about *seven months.*

The "horizon blending" function of Northfield's risk applications has used a similar concept for many years as described in https://www.northinfo.com/documents/779.pdf

Conclusions

- While exact translations from a known mean-variance objective function can be translated to mean-standard deviation space, the "rule of six" provide a quick, approximate solution that is intuitive to most investors, over a wide span of practical cases.
  - Using the method of Cornish and Fisher (1937), comparable problems with VaR and CVaR are accommodated.

- The derivation of the rule from the Discretionary Wealth Hypothesis (Wilcox 2003) is hopefully clear as presented.

- We have presented three examples of where this rule of thumb may be of immediate use to investors, including making an inference about the aggregate market expectation for the time length of the coronavirus emergency.