The Cost of Volatility

Geoff Considine, Ph.D.

In a recent article, I discussed why it is important to pay attention to risk measures like volatility and Beta:

http://www.quantext.com/WhyVolatilityMatters.pdf

There is another way to examine the effective cost of volatility for investors: ***volatility drag***. In financial modeling, we often look at average returns as the metric of growth rate, but there is another measure: Compounded Annual Growth Rate (CAGR), which is sometimes referred to as annualized return (Note: this terminology can be confusing because annualized return does not necessarily refer to CAGR, though it is often treated this way). The difference between the average annual return and CAGR is determined by volatility (as measured by the standard deviation in return)—and this difference is important. For a brief discussion of the difference between CAGR and annual return, I like this online calculator at *MoneyChimp.com*:

http://www.moneychimp.com/features/market_cagr.htm

CAGR for a period is calculated by calculating the total return over the period using the starting price and ending price (adjusted for dividends and splits) and then calculating the compounded growth rate that would get you from the starting price to the ending price in that period of time:

http://www.investopedia.com/terms/c/cagr.asp

Consider how this differs from a calculation of the average annual return over ten years—for the average annual return, you simply average the returns from each of the ten years. Why does this matter? If you have a stock that goes from $10 to $20 in a year, you have a 100% gain for that year. In the next year, let's say that the stock drops from $20 to $5---that's a 75% loss. The average return over these two years is (100%-75%)/2 = 12.5% per year. Average annual return is always higher than CAGR because of the volatility—

and this effect is called ***volatility drag***. An approximate equation that lets you calculate the volatility drag in terms of the Standard Deviation in annual return is:

$$(1 + \text{Average Return})^2 - (\text{Standard Deviation})^2 = (1 + \text{CAGR})^2$$

Average annual return is only equal to CAGR when standard deviation goes to zero. So why don't we just look at everything in terms of CAGR rather than average annual return? One reason is that the statistical models for risk and return are more naturally written in terms of average return and standard deviation in return. From my perspective, average annual return and standard deviation in return are the fundamental variables—and CAGR is derived from these.

Volatility drag can be very important to investors. Consider the following calculations for CAGR vs. average annual return for the past ten years (through the end of May 2007) for a number of Dow components: Alcoa (AA), Boeing (BA), Disney (DIS), Honeywell (HON), Hewlett Packard (HPQ), IBM, and Coca-Cola (KO).

Ticker	Average Annual Return	CAGR	Difference
AA	16.3%	10.2%	-6.0%
BA	12.8%	8.6%	-4.2%
DIS	7.8%	4.0%	-3.8%
HON	12.1%	6.0%	-6.2%
HPQ	17.3%	9.6%	-7.7%
IBM	14.5%	9.9%	-4.6%
KO	2.2%	-0.8%	-3.0%

Trailing ten-year average annual return vs. CAGR

If these results for the 'cost of volatility' look bad, there are two factors that make it worse. First, recent years have seen very low overall volatility in the stock market—which will tend to mask these effects in the near term. In other words, you have not felt these effects in the most recent years. Second, these are 'blue chip' stocks—DJIA components. These are far from the most volatile companies out there. Over the past ten years, the impact of volatility on the total growth in an investment in Alcoa (AA) has been -6% per year. Over that same period, this 'cost of volatility' has been -6.2% per year for Honeywell (HON). For Disney (DIS), the cost of volatility has been -3.8 per year.

As an aside, when you look at historical performance statistics, you need to be sure what you are looking at—CAGR or average annual return—and some data sources do not make this clear.

Personally, I find the volatility drag calculation less useful than simply looking at the average and standard deviations of investments and how they impact the percentiles of being able to fund my future income---but these are simply different ways to look at the problem. If we look at the trailing ten-year statistics for a portfolio equally allocated between these stocks, we can see that this portfolio has delivered reasonable returns, but with considerable volatility:

Historical Data	
Start: 6/1/1997	End: 5/31/2007
Average Annual Return	Standard Deviation (Annual)
11.9%	21.5%

Trailing ten-year average annual return and standard deviation in return

The 'volatility drag' on this portfolio is above 2% per year (using the formula shown earlier).

If you read my articles, you will find that I tend to advocate looking for an equity portfolio with average return that is about equal to the standard deviation through the use of better diversification:

http://www.quantext.com/RiskReturn.pdf

If we could find a portfolio with a standard deviation of 11.9% and an annual return of 11.9%, the 'volatility drag' would be 0.6%---less than one third of the volatility drag on the equally-allocated portfolio between these Dow components. Maximizing return relative to standard deviation will drive down volatility drag.

What should concern investors is (1) getting the highest average return for the level of volatility in their portfolio and (2) the ultimate impact of volatility upon the ability to fund future income. The effects of compounding that create 'volatility drag' are fully accounted for in Monte Carlo simulations like Quantext Portfolio Planner (QPP) and Quantext Retirement Planner (QRP). In addressing these two primary concerns, volatility drag is managed as well.

Now let's get down to why all of this should be of immediate concern. I have been writing about the fact that volatility has been running at very low levels over recent years. Looking at charts of an index that measures volatility called VIX gave me an idea about how to illustrate this. VIX and what it means are discussed in this article:

http://www.quantext.com/MarketRisk2.pdf

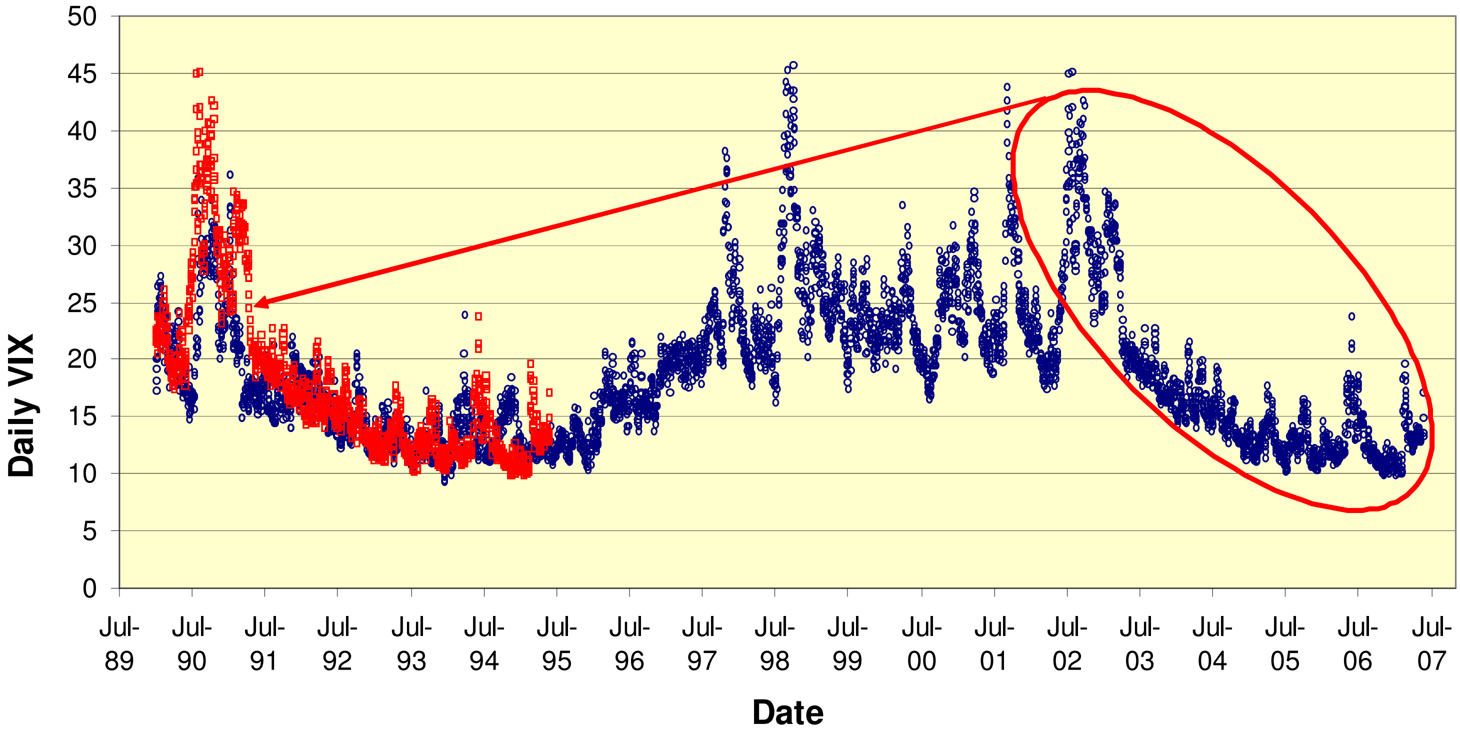

Historical daily VIX—comparing recent years to the early 1990's

The chart above shows historical VIX data going from June 7, 2007 back to about the start of 1990. If I take the period in the ellipse above (the recent years) and simply overlay it on the early 1990's, the trajectory of VIX matches really well (red vs. blue). These data show that we are in a situation that is quite similar to the early-to-mid 1990's---and it is easy to see what happened next. This is, of course, a qualitative comparison but this kind of volatility surge is reasonably to be expected. As I showed in the article about VIX linked above, standard deviations in return scale quite linearly with VIX. If VIX doubles, we can expect to see standard deviation in annual return double over current levels—which will in turn lead to considerably higher volatility drag, depending on your portfolio. If you want to be able to forecast CAGR reasonably well, you must have outlooks for both average returns and volatility in returns---both of which are provided by tools like Quantext Portfolio Planner (QPP) and Quantext Retirement Planner (QRP).

As I mentioned earlier, I do not look at CAGR on an operational basis. I know that volatility has a cost and I can see it in the Monte Carlo probabilities of being able to fund my future income from my investments (which accounts for volatility drag implicitly)—and these calculations are provided by QPP and QRP. Still, it can be useful to look at CAGR rather than average annual returns in really grasping the importance of volatility for planning.

Disclosure: the author owns AA and HON

Quantext Portfolio Planner *is a Monte Carlo portfolio management tool. Extensive case studies, as well as access to a free extended trial, are available at http://www.quantext.com/gpage3.html*