Generating Meaningful Outlooks for Portfolio Performance

Geoff Considine, Ph.D.

The main reason to use a Monte Carlo portfolio planning tool (as opposed to simply looking at historical data) is the ability of a Monte Carlo simulation to generate forward-looking projections of portfolio risk and return that are demonstrably better than looking at trailing performance. Any potential user of a Monte Carlo tool is well advised to ask for some evidence that the tool is better than simply using historical data (i.e. looking at past performance). I recently received an email from an investor who wanted to use a Monte Carlo simulation to assist in his planning, but was concerned about whether these tools were biased towards recent history. This question is one that is not asked frequently enough and inspired me to write this article.

To illustrate how this process works, I will use Quantext Portfolio Planner (QPP). The user runs QPP with recent market history (say three years—but this is a user input) and QPP combines this near-term history with long-term statistics for the risk and return across asset classes in capital markets. This approach (that we call risk-return balancing) yields different (and better) projections of future portfolio performance than simply using historical performance. The projected future risk and return for each asset in the portfolio is combined to calculate a forward-looking assessment of risk and return for the entire portfolio. For a previous article on how this works for a single index investment, the reader may be interested in this study of predicted risk and return for the NASDAQ index (via QQQQ):

http://etf.seekingalpha.com/article/16586

There are two key issues here. First, a model must project the volatility for each asset and for the total portfolio. Second, the model must be able to project expected returns that are reasonable and consistent with risk levels. QPP starts with projecting risk (i.e. volatility) for assets. QPP's volatility projections are routinely benchmarked and are consistent with both long-term history and the implied volatility from options:

http://etf.seekingalpha.com/article/27508

We have also compared the risk and return that QPP projects for broad asset classes and with long-term history and outlooks from major institutional advisors:

http://etf.seekingalpha.com/article/24588

The value of all of this computational technology is the ability to create a meaningful forward-looking estimate of risk and return for a portfolio. We know from a wide range of studies that designing your asset allocation based on what has happened in recent years is a recipe for disaster. Bernstein provides a particularly nice case study of this effect in ***The Intelligent Asset Allocator***. Over periods of longer than one year, markets are generally mean-reverting, which means that the performance of an asset will tend to return to its long-term average return, even though it may drift away from this average for extended periods of time. Further, the average returns that the asset will tend to generate are determined by the assets risk level. Riskier assets will tend to generate a higher long-term average return (small cap stocks, for example) and less risky assets (high-grade bonds, for example) will tend to generate a lower average return. An asset that outperforms relative to its risk level will tend to under-perform in the future to bring the risk and return into line. John Bogle is a strong advocate of mean reversion in markets and he sums up some of the evidence and implications here:

http://www.vanguard.com/bogle_site/sp20020626.html

A key issue in reversion-to-the-mean (RTM) is determining the rational level of return that an asset or asset class tends to revert to. This level is determined by the risk level of the asset. QPP generates outlooks of volatility (risk) and then calculates a projected average return that is consistent with the risk level. This turns out to work remarkably well and also tends to look like RTM.

We have now had enough concepts. Let’s look at an example. I have created a sample portfolio of stocks picked from the Dow Jones Industrial Average components, with equal allocation to each:

Ticker	Percentage of Funds
BA	14.3%
JNJ	14.3%
HON	14.3%
IBM	14.3%
KO	14.3%
MMM	14.3%
XOM	14.2%

Sample blue-chip portfolio

We have a range of sectors here and these are all great companies. Please understand that I am not advocating this portfolio for anyone in particular although I will note, for purposes of disclosure, that I own Johnson and Johnson (JNJ) and Honeywell (HON) in my portfolio.

	Portfolio	BA	JNJ	HON	IBM	KO	MMM	XOM
BA	68%	100%						
JNJ	21%	-15%	100%					
HON	81%	68%	-10%	100%				
IBM	43%	-3%	3%	28%	100%			
KO	59%	29%	22%	29%	10%	100%		
MMM	57%	18%	12%	36%	18%	40%	100%	
XOM	59%	36%	14%	41%	15%	13%	6%	100%

Correlation matrix for portfolio components for 2003-2005

When you look at the correlations in total returns (price appreciation plus reinvested dividends) in the table above, the motivation for the choice of stocks will be clearer. I wanted to choose companies with low correlation to one another so that the portfolio would exploit diversification effectively. The correlations in total returns are quite low, as a group. These correlations are a textbook example of diversification. Coke (KO) has a very low correlation to Exxon (XOM). This makes sense—the business cycle for soft drinks is not related to the economic cycle for oil. Heavy industrials like Boeing (BA) will tend to exhibit low correlation to information services firms like IBM. A company like Johnson and Johnson (JNJ) tends to exhibit low correlation to everything outside of its immediate sector.

The purpose of my analysis is to look at how QPP would project the future performance of this portfolio using three trailing years of data as the input. The portfolio is re-balanced annually. The historical period 1970-1972 was used to predict performance in the period 1973-1975, historical period 1973-1975 was used to predict performance in 1976-1978, etc. I have used all default inputs, including the assumption that the S&P500 will deliver an average return of 8.3% per year with a standard deviation of 15.07%. This assumption will tend to bias the results of this historical analysis downwards, because we know that the S&P500 has delivered an anomalously high equity risk premium over this period. I wanted to run with all default settings, however, to make this study more meaningful.

This analysis really tests the principles of the Monte Carlo simulation because we have a small number of stocks (i.e. substantial non-market risk) and the portfolio does not simply track the S&P500—the R-squared of the portfolio, a measure of how much of the return can be explained by movements in the S&P500, is only 62% for the period from 2003-2005, the most recent period used in this study. This means that only 62% of the return is due to the broader market and the remaining 38% of the return is due to factors specific to these stocks and the correlations between them. Further, one of the questions that come up when people are considering using a Monte Carlo tool is whether it will be able to account for the correlation effects between the assets. Further, these correlations can change in time, and users want to see whether a portfolio tool will be able to generate reasonable results in light of such changes.

Period	Trailing 3-Year Return	Baseline QPP Projected Average Return (using trailing 3 years as input)	Average Annual Return
2003-2005	1.6%	13.9%	14.0%
2000-2002	17.1%	13.5%	1.6%
1997-1999	29.1%	14.5%	17.1%
1994-1996	11.7%	15.2%	29.1%
1991-1993	19.2%	11.0%	11.7%
1988-1990	14.5%	9.6%	19.2%
1985-1987	24.1%	13.6%	14.5%
1982-1984	7.8%	11.9%	24.1%
1979-1981	12.2%	14.3%	7.8%
1976-1978	-0.2%	12.0%	12.2%
1973-1975	20.6%	14.3%	-0.2%

Projected performance of portfolio compared to actual and trailing three-year return for a series of three year periods

In 1970-1972, this portfolio yielded an average return of 20.6% (see the trailing return for the period 1973-1975 above). Using only this period as input to QPP, QPP projected that the period from 1973-1975 would return 14.3%, much lower than the recent three years. The portfolio actually returned -0.2% per year, on average. It is quite simple to see that QPP looks mean-reverting. After a very high-return period (1982-1884, 1994-1996), QPP tends to project lower future returns. After a very low return period (1973-1975, 2000-2002), QPP tends to project higher future returns.

A good statistic for testing the accuracy of predictions (whether it is for the stock market or for the weather) is the Mean Absolute Error (MAE). The MAE is the average difference between the ***predicted value*** and ***what actually happened***, regardless of whether the predicted value is too big or too small. If the predicted average return is 12% and the actual return is 10%, the absolute error is 2%. If the predicted average return 8% and the actual return is 10%, the absolute error is 2%. If we have average these two values, we get the MAE for the two predictions—and the MAE is 2%. It should be clear that this is a much better way to measure the quality of a forecast than simply averaging. In the first case we had an error of -2% and in the second case we had an error of 2%. If we simply averaged the errors, we would have an average error of 0%, but that is not meaningful. A forecast of too high a return does not offset a forecast of too high a return—both lead to sub-optimal decisions.

The Mean Absolute Error (MAE) for annual return from the QPP projections over all of the three-year periods is 6.7%. If you use the average annual return from the trailing three years as your prediction for the next three-year period, your MAE over all of the three years periods is 12.1%. ***This average error in predicted performance is a lot lower in the forward-looking projections from QPP than the errors when you use trailing performance to predict the future performance.***

QPP's projections for the total portfolio return are as good as they are because the model estimates an expected return based on volatility whereas the trailing return in any given three year period is not likely to be a great estimate of the expected future return. There is a very good correspondence between the predicted volatility of the portfolio and the volatility that was observed:

Period	Trailing 3-Year Standard Deviation in Annual Return	QPP Projected Standard Deviation in Return (using trailing 3 years as input)	Actual Standard Deviation in Annual Return
2003-2005	20.0%	17.7%	10.6%
2000-2002	19.7%	20.2%	20.0%
1997-1999	11.0%	20.1%	19.7%
1994-1996	11.5%	19.0%	11.0%
1991-1993	14.0%	17.7%	11.5%
1988-1990	19.3%	16.1%	14.0%
1985-1987	16.9%	20.9%	19.3%
1982-1984	14.8%	17.5%	16.9%
1979-1981	15.1%	22.5%	14.8%
1976-1978	23.8%	20.4%	15.1%
1973-1975	16.2%	21.6%	23.8%

QPP projected portfolio volatility compared to actual volatility and trailing volatility

The results shown here span a wide range of market conditions. The annualized volatility of this portfolio ranged from 10.6% (2003-2005) to 24% (1973-1975). The historical results show the annualized volatility (which is measured by the annualized standard deviation in return) calculated from monthly total returns. ***The correlation between trailing three-year annualized volatility and the observed volatility over the next three years is zero.*** This does not mean that using history provides no way to predict future volatility, however. Using the trailing three years annualized volatility as the prediction for the next three years gives an MAE error of 4.3%--i.e. you would tend to be wrong, on average, by 4.3%. This is much lower than the MAE error when we tried the same thing with average return (MAE error of 12.1%, as discussed above). ***The correlation between the QPP projected volatility and the observed future volatility over the next three years is 55%.*** Even more interesting, if we look at the change in volatility between the trailing three years and the next three year period, we find the QPP is good at predicting changes in portfolio volatility:

Period	QPP Predicted Change in Standard Deviation in Annual Return	Change in Standard Deviation in Annual Return from previous three years
2003-2005	-2.3%	-9.4%
2000-2002	0.4%	0.3%
1997-1999	9.1%	8.8%
1994-1996	7.6%	-0.5%
1991-1993	3.7%	-2.5%
1988-1990	-3.2%	-5.3%
1985-1987	4.0%	2.4%
1982-1984	2.6%	2.0%
1979-1981	7.5%	-0.2%
1976-1978	-3.4%	-8.7%
1973-1975	5.4%	7.6%

QPP projected changes in volatility compared to observed changes in volatility in successive three-year periods

This table is derived from the previous one (17.7%-20.0% = -2.3%). When QPP projects an increase in portfolio volatility relative to the most recent three years, the model is often correct. The correlation between predicted and observed changes in volatility for this portfolio is 78%. This result is generally consistent with financial theory: it is known that volatility is ***relatively*** predictable.

These results show how and why a good forward-looking Monte Carlo tool adds value. QPP projects future volatility reasonably well—certainly better than looking backwards—and these volatility estimates can be used to generate a rational expected return over future years. ***The underlying reason that portfolio performance can be predicted is that risk and return exhibit a consistent relationship over extended periods.*** The expected return is calculated using both the projected volatility and long-term historical statistics—and this is why the results will tend to look mean reverting rather than putting weight on recent performance. If, by contrast, you look at trailing historical performance as the basis of predicting future return and risk, you are likely to have a very poor estimate of future performance. If you optimize to history, say using a mean variance optimizer (MVO), things will get even worse (as Bernstein, among others, has shown). The errors in the projected average return for the portfolio using QPP are far lower than the errors that result from a projection based on trailing historical performance.

It is, of course, fair to ask why one should not look at longer histories, but this is something of a catch-22. The world changes and it does not make sense to argue for using ten years rather than five years, or twenty years rather than ten years. A good Monte Carlo model bridges this problem by combining near-term and long-term data to generate a meaningful outlook in which the general balance of risk and return bears out. Those who do their planning by looking at recent years rather than using a rational world view in which risk and return balance out in the long-term are likely to end up with undesirable results.

So where does this leave us? The portfolio of stocks that we have analyzed in this study is quite predictable, but the mean reversion effects that it exhibits are not simple. Volatility is strongly mean reverting over multi-year periods, and swings in volatility have a direct impact on the probability that you will get a high return or a low return. It is not enough to assume that a high-return period will be followed by a low return period (simple mean reversion) or that you should buy winners (simple momentum). The principles that drive QPP and lead to its good results are the following:

1) Volatility is mean reverting and is therefore (statistically) predictable
2) Risk and return balance out across assets and over multi-year periods
3) Because of (1) and (2), average returns are mean reverting over multi-year periods
4) Estimates of future risk and return must combine long-term data on risk/return balance with near-term data on volatility

Quantext Portfolio Planner shows that there is substantial ***statistical predictability*** in a portfolio of stocks over more than thirty years. This is exactly what investors need from a Monte Carlo planning tool. The tool will not predict that the market will turn down next month. The tool will provide meaningful projections of the average return on a portfolio and the risk in the portfolio that can assist investors in building an asset allocation that meets their needs.

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*