

Introduction

If you have used or read articles about Monte Carlo portfolio planning tools, you may have wondered how a user can determine whether the predicted risk in the stock market as a whole and for individual sectors is correct. Monte Carlo models are supposed to help you plan by accounting for risk in your portfolio, but how do you know that their projections of risk are reasonable? Monte Carlo models are only as good as the inputs. If a model predicts market volatility that is too low, the tools will make overly optimistic projections about the income you can expect from a specific savings level and portfolio allocation. If the predicted market volatility is too high, the model will predict that you need to save more than you actually do to ensure a specific level of future income. The three critical inputs that Monte Carlo models use are (1) expected rates of return, (2) Beta and (3) projected volatility for each portfolio component.

The best way to test a portfolio model to determine whether its volatility projections make sense is to look at where options on the portfolio components are trading in the market. Options are financial instruments that give you the right to purchase or sell stock at a specific price (called the strike price). The value of an option with a specific strike price is determined by the assumed underlying volatility on a stock or fund and on the time until expiration. The same assumptions about future markets used to drive a Monte Carlo portfolio model also provide the key inputs for valuing options. If the inputs to a Monte Carlo portfolio simulation model yield option values that are consistent with the values at which the options are trading in the market, you can have confidence that the projected future portfolio risk in the Monte Carlo model is consistent with the way that the market is assessing risk. Comparing options values driven by a Monte Carlo model to the level at which options are trading in the market is a standard tool in corporate portfolio risk management applications. This process is called *Marking to Market* or MTM. If you have options that are trading with expirations going out three years, these options values tell you what the market thinks future volatility will be. Many companies

require that their portfolio simulation tools are consistent with current option quotes and the process of ensuring this is MTM.

Comparisons between options values in a Monte Carlo model and options values trading in the market can be performed for individual stocks, but for many investors it may be of greatest interest to look at options on Exchange Traded Funds (or ETF's) to determine whether the Monte Carlo model's assessment of risk associated with the market as a whole and for individual sectors is close to the future volatility implied by the prices of options.

This paper examines the market projections for risk generated using the Quantext Retirement Planner (QRP). QRP is a Monte Carlo personal portfolio management tool. More information is available at <http://www.quantext.com>.

A Motivating Case Study

When making portfolio projections, we are often interested in looking at projected risk (or volatility) associated with the market as a whole. The portfolio projections from QRP or any other portfolio management tool will depend quite strongly on the projected future volatility in the S&P500 and in other market indices. Without exception, the Monte Carlo tools for personal portfolio management that I have looked at (including Financial Engines and a number of others) do not provide any basis for the user to have confidence that projected future market risk is at all realistic. In QRP, there are always two benchmarks that are available to the user. One is to show how projected risk relates to historical risk over a user-specified period of time. The other basis is to show that projected risk or volatility in a stock or fund is consistent with where the market is trading options---marking the level of volatility to market levels.

The simplest way to examine this issue is to look at an investment in SPY, an S&P500 ETF. We have entered SPY as an element of a sample portfolio in QRP and we have selected the three years through the end of September 2005 (as this is the most recent

with where options are trading. QRP uses the first component of the list of holdings (see chart above) in options calculations. If you had employee stock options, you would enter your employer's stock in this position. Because we are just checking the value of options from simulated future markets, we have entered SPY in this position. We now look at the implied values of options using the input parameters from above:

<i>Retirement Planning Report</i>			
Prepared For: John Doe		Preparation Date: 10/27/2005	
Page 8: Stock Options			
Options for: SPY			Recent Share Price: \$118.60
www.quantext.com			
Strike Price	Shares	Expiration	Current Expected Value Per Share
\$120.00	1	01/20/06	\$3.04
\$120.00	1	03/17/06	\$4.25
\$120.00	1	09/15/06	\$7.31
\$120.00	1	12/15/06	\$8.60
\$120.00	1	12/21/07	\$13.13

Figure 2: Sample options for SPY

We have used all of the default settings in QRP, along with the inputs from Figure 1. On the stock options screen (excerpted in Figure 2), we have entered a series of expiration dates for options from early in 2006 through the end of 2007. We obtained these quotes from *Yahoo! Finance*. Options trade for one expiration date in each month, so we are looking at the options for January 2006, March 2006, September 2006, December 2006 and finally—looking out as far as we can—December 2007. We are using a strike price of \$120 per share and SPY is currently trading at \$118.60, so these options are close to the current market level. Recall that a call option gives us the right to buy a share of SPY at the strike price of \$120 at any time between today and expiration. The more volatile the market expects the stock to be, the higher the price of the call options and vice versa.

The *current expected value per share* in Figure 2 shows the fair value for these options calculated using the parameters that drive QRP. We now compare these calculated values to the levels at which the market is trading these options.




Figure 3: Prices of SPY call options: default inputs

The market values of SPY call options are very close to the values calculated by QRP for each expiration date going out more than two years into the future. These results show that the projected volatility used in QRP for the S&P500 is consistent with the volatility that is implied by the prices of options trading in the market.

We will now look at another case—the NASDAQ 100 ETF (QQQQ). The NASDAQ 100 is comprised of the 100 largest non-financial firms that trade on NASDAQ. When we perform the same set of calculations for call options with a strike of \$38 (QQQQ closed at \$38.03 at the time of this writing), we find that the agreement between QRP’s projections and the market is also very close (Figure 4).

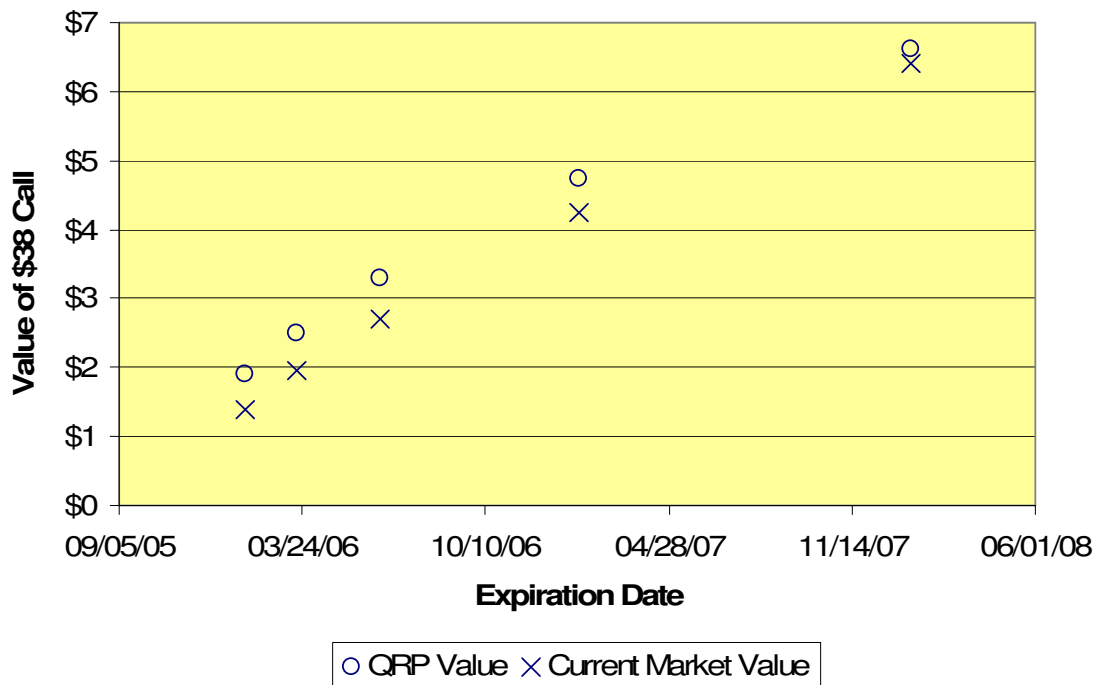


Figure 4: Prices of QQQQ call options: default values

Having looked at the S&P500 index and the more volatile high-Beta NASDAQ 100 index, we can perform the same check on a low-Beta sector ETF. This is important because the total risk (volatility) associated with a stock or fund is a combination of risk due to the market (systematic risk) and risk that is specific to that particular stock or fund (non-systematic risk). When we look at QQQQ (with a Beta of 160%) we are primarily emphasizing the representation of systematic risk, but when we look at a low Beta fund, we can test the representation of non-systematic risk. For this purpose, we have chosen to look at the Utilities Select Sector ETF (ticker: XLU). The *Risk-Return Balanced* value of Beta calculated for XLU is 28% (Figure 1), but the annual Standard Deviation (SD) in return is still higher than the market as a whole, so this is a good test case for non-systematic risk. Once again, we calculate the values of a series of call options using default input values. The agreement between QRP’s calculated values and the market values is very good (Figure 5).

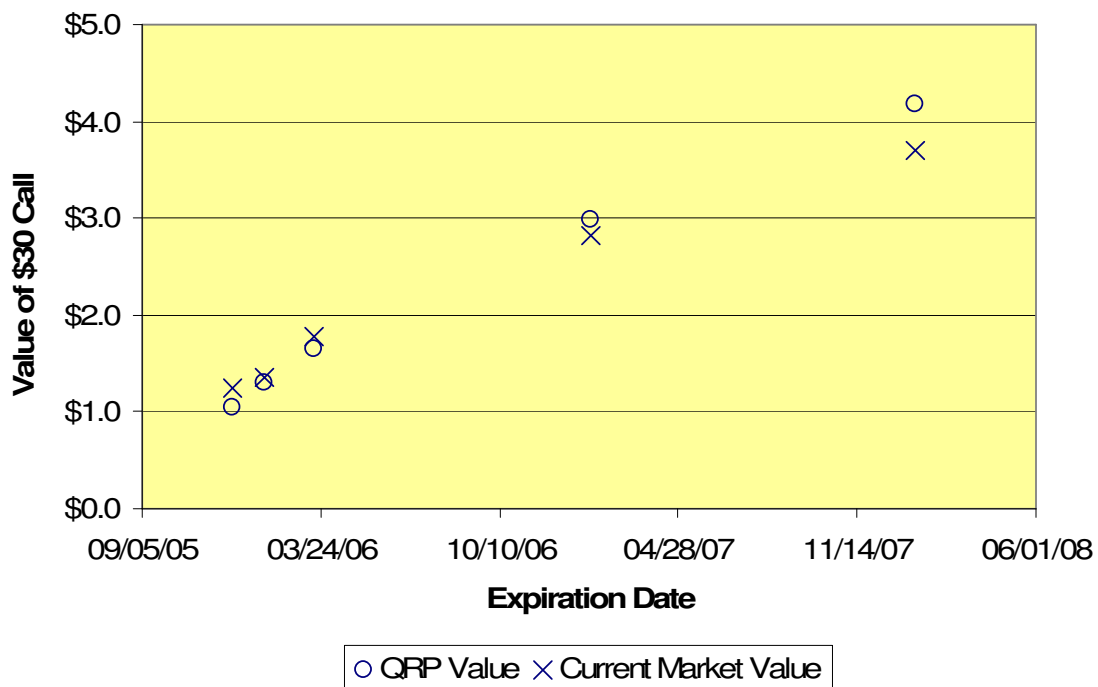


Figure 5: Values of XLU options: default values

Analyzing the Total Portfolio

The comparison of individual options prices on these broad market and sector ETF's are a powerful validation of the risk projections in the underlying indices. Another very important feature of a portfolio management tool, however, is to demonstrate that the total portfolio risk for various allocations into these individual indices is consistent with the market as a whole. This is also easily shown in QRP.

QRP allows the user to specify the projected future average return and standard deviation of return for the market as a whole (we use the S&P500 index). For the calculations shown earlier in this paper, we have used default market inputs that assume an average return of 8.3% per year, with a standard deviation of 15.07% per year, values that are lower than the most recent decades, but are consistent with longer market histories (Figure 6).

Years	Period	Average Annual Return	Standard Deviation in Annual Return
100	1905-2004	8.00%	18.95%
90	1915-2005	8.56%	19.29%
80	1925-2004	8.87%	19.33%
70	1935-2004	8.67%	17.73%
60	1945-2004	8.61%	16.68%
50	1955-2004	7.65%	15.85%
40	1965-2004	6.56%	16.19%
30	1975-2004	9.52%	15.54%
20	1985-2004	10.68%	16.35%
10	1995-2004	10.59%	19.80%

Raw Data: Website of Robert Schiller at Yale (<http://www.econ.yale.edu/~shiller/data.htm>)

Figure 6: Average Return and Standard Deviation of Annual Return for the S&P500 over various historical periods.

These projected risk and return figures used in QRP for the market as a whole are consistent with current thinking that the market will yield returns that are lower than the 10% per year that we have seen over the most recent two decades. Over the three years of historical data used as input for QRP, the market has yielded a high average annual return and a fairly low standard deviation (low volatility) by comparison (bottom of Figure 1). When we adjust the input market data for QRP to be consistent to that for the most recent three years, we can compare the Monte Carlo calculated *portfolio* Beta, average return, and standard deviation of return to the actual market values over the past three years (Figure 7). We have also used the Historical Data from Figure 1 as input rather than the risk-return balanced data. We have specified a sample portfolio allocated 30% to QQQQ, 40% to SPY, and 30% to XLU. The Beta for this portfolio over the past three years (*Historical Beta* in Figure 7) is 94.16% as compared to the simulated Beta of 95.18%. The historical average annual return and standard deviation of annual return (*Historical Data* table in Figure 7) are very close to the Monte Carlo simulated values (*Portfolio Stats* in Figure 7). The close agreement shows that the simulated values of these ETF's are correlated in a way that is consistent with their historical behavior. QRP uses the historical period to calibrate the correlations (derived from Beta) between

portfolio sectors. These results for average return and standard deviation of return are not projections into the future, however. For that, we will use risk-return balanced data for input and our long-term projections of 8.3% per year of average return for the market as a whole and a standard deviation of 15% per year.

Fund or Stock Ticker	Beta	Standard Deviation (Annual)	Alpha (Annual)	Check
QQQQ	160.5%	20.55%	1.02%	OK
SPY	96.1%	10.96%	2.24%	OK
XLU	27.9%	11.91%	20.08%	OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
Annual Return = Beta x (Annual Return on S&P500)+Alpha				
<i>Note: This definition for Alpha has rolled in the risk-free rate of return</i>				
			Portfolio Stats	
Fund Name	Percentage of Funds	Average Annual Return	Average Annual Return	Standard Deviation (Annual)
QQQQ	30.0%	23.88%	20.79%	11.55%
SPY	40.0%	15.94%		
XLU	30.0%	24.16%	Historical Data	
-	0.0%	0.00%	Start:	End:
-	0.0%	0.00%	9/30/2002	9/30/2005
-	0.0%	0.00%	Average Annual Return	Standard Deviation (Annual)
-	0.0%	0.00%	20.84%	11.18%
-	0.0%	0.00%	Historical Beta: 94.96%	
Sums to	100.0%			
Simulated Portfolio Beta: 95.17%			Market Index (S&P500)	
			Average Annual Return	Standard Deviation (Annual)
			14.30%	11.06%

Figure 7: Market matched input data for S&P500 in past three years

When we go to the default projections, we obtain very different portfolio results from those in Figure 7.

Fund or Stock Ticker	Beta	Standard Deviation (Annual)	Alpha (Annual)	Check
QQQQ	160.5%	24.21%	1.31%	OK
SPY	96.1%	14.50%	1.85%	OK
XLU	27.9%	16.24%	8.37%	OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
				OK
Annual Return = Beta x (Annual Return on S&P500)+Alpha				
<i>Note: This definition for Alpha has rolled in the risk-free rate of return</i>				
			Portfolio Stats	
Fund Name	Percentage of Funds	Average Annual Return	Average Annual Return	Standard Deviation (Annual)
QQQQ	30.0%	14.51%	11.50%	15.09%
SPY	40.0%	9.75%		
XLU	30.0%	10.82%		
-	0.0%	0.00%	Historical Data	
-	0.0%	0.00%	Start: 9/30/2002	End: 9/30/2005
-	0.0%	0.00%	Average Annual Return	Standard Deviation (Annual)
-	0.0%	0.00%	20.84%	11.18%
-	0.0%	0.00%	Historical Beta: 94.96%	
Sums to	100.0%			
Simulated Portfolio Beta: 95.06%			Market Index (S&P500)	
			Average Annual Return	Standard Deviation (Annual)
			8.30%	15.07%

Figure 8: Portfolio Projection

When we examine a forward projection for this portfolio using risk-return balanced data (as we used for valuing the options in the earlier sections), the projected total average

annual return and standard deviation of return is quite different from what we have seen in the most recent three years (Figure 8). From the earlier tests, we know that the underlying volatility for each of these ETF's yields options values that are very close to the market as a whole. One might take issue with the projected average annual returns in each of these ETF's and that is why the parameters for Alpha can be adjusted to be different from the values automatically calculated by QRP (Figure 8). Changing Alpha for a stock or fund will change the average annual return without changing volatility which means that the user can change Alpha without impacting the MTM results for options. Through various tests, we generally find that the risk-return balanced parameters are reasonable. A projected annual return of 11.5% with a standard deviation of a bit more than 15% is a reasonably consistent with the historical relationships between risk and return. If anything, it is probably a bit optimistic and a conservative user might reduce the levels of Alpha.

A useful additional feature of QRP's full Monte Carlo approach is that you can look at projected risk and return on specific time horizons. While Figure 8 shows annual return and standard deviation in return, you can specify any time horizon that is of interest. When the user inputs a 90-day horizon, for example (Figure 9), he/she can see the projected median, average, and other percentiles for portfolio value (in this case a starting value of \$100,000), gain or loss, and return. With the portfolio in Figure 8, there is a 5% chance of losing at least 9.7% of the portfolio value in 90 days, for example (Figure 9). This type of projection would be a great deal less meaningful without the MTM comparisons to the options markets.

Portfolio Value:	\$100,000	comments	
Time Horizon (days):	90		
Portfolio Beta:	95.06%		

Percentile	Portfolio Value	Gain / Loss	Return
1%	\$86,164	-\$13,836	-13.8%
5%	\$90,261	-\$9,739	-9.7%
10%	\$92,970	-\$7,030	-7.0%
15%	\$94,894	-\$5,106	-5.1%
20%	\$96,507	-\$3,493	-3.5%
25%	\$97,599	-\$2,401	-2.4%
30%	\$98,889	-\$1,111	-1.1%
35%	\$100,126	\$126	0.1%
40%	\$101,035	\$1,035	1.0%
45%	\$102,027	\$2,027	2.0%
50%	\$102,989	\$2,989	3.0%
55%	\$103,941	\$3,941	3.9%
60%	\$104,793	\$4,793	4.8%
65%	\$105,729	\$5,729	5.7%
70%	\$106,859	\$6,859	6.9%
75%	\$107,836	\$7,836	7.8%
80%	\$109,021	\$9,021	9.0%
85%	\$110,442	\$10,442	10.4%
90%	\$112,120	\$12,120	12.1%
95%	\$115,421	\$15,421	15.4%
99%	\$120,273	\$20,273	20.3%
Average	\$102,835	\$2,835	2.8%

Figure 9: Risk – Return for portfolio from Figure 7 for 90 day horizon

Bringing it All Together

Monte Carlo models for retirement planning or other portfolio management applications have a number of important features. The most significant of these is the ability to account for investment risk in planning. Without accounting for the volatility in returns on investments, portfolio planning will tend to be far too optimistic. Determining savings rates and allocations without accounting for risk is like assuming that all investments are risk-free and many would-be retirees have paid a heavy price for this failure in planning. Monte Carlo models are rapidly becoming standard of practice in financial planning, but there is a substantial risk to users if the inputs that drive these models are not realistic. In

the Monte Carlo models that I have used, the argument is typically that the risk-return characteristics used as inputs to the models are derived from historical values in some fashion, but details are scant. Most investors don't really want to know the details, of course, so this may be okay. On the other hand, bad input assumptions can essentially negate the value of Monte Carlo models. Since the portfolio projections depend on the inputs, it is very important to make sure that those inputs project a reasonable future market. By using Mark-to-Market tests for QRP, we can demonstrate that the projected volatility associated with individual stocks, funds, and ETF's is consistent with the market's assessment of volatility—at least as far out as options trade. This is considered an important validation in professional risk management applications and it provides significant confidence that a model is not way off base for personal portfolio planning. Certainly the market's assessment of future volatility for the market as a whole and various sectors (as reflected in ETF options) is not a perfect forecast but it is a consensus outlook from the market---and it does not make sense to ignore this source of information.

Performing mark-to-market comparisons between QRP's Monte Carlo output and options trading in the market is an important part of demonstrating the validity of QRP's analytics as a basis for portfolio planning. I have not seen any similar validation for other Monte Carlo portfolio models, which may lead the user to ask where these volatility estimates come from. These volatility estimates are the key driver for Monte Carlo portfolio projections, so it makes sense to look at how these inputs can be 'sanity checked.'