NORTHFIELD | modeling financial markets worldwide

Financial Assets Behaving Badly
The Case of High Yield Bonds

Chris Kantos
Newport Seminar

June 2013

Main Concepts for Today

- The most common metric of financial asset risk is the volatility or standard deviation of return.

- The popularity of standard deviation arises from two very different bases.
  - Standard deviation is the most widely understood statistical measure of the *dispersion of a possible set of outcomes around the expected value*
  - Long term wealth accumulation arising from investment is related to the geometric mean return, but when we typically discuss statistical distributions of anything we measure the arithmetic mean. The difference between the geometric mean and the arithmetic mean is presumed to be equal to one-half the square of the standard deviation (variance).

Main Concepts for Today

- The legitimacy of both uses of standard deviation as a measure depends on the validity of underlying assumptions about the distribution being described.
  - Typically, we assume that the process is a "random walk", resulting in asset returns that are normally distributed, have a consistent mean and volatility, and no serial correlation ("i.i.d", or independently and identically distributed).
  - For most asset classes when we look at returns measured over holding periods of quarters of years or longer, formal statistical tests would not usually reject the hypothesis that a random walk is a sufficiently good description of the return process.
  - For most asset classes we would reject the hypothesis of a random walk process when we measure returns over short periods such as days as the distributions appear to have "fat tails" where extreme events are much more frequent than expected See diBartolomeo (2007).

High Yield Bonds: "Bad Boy" of Asset Classes

- One asset class that clearly shows "bad behavior" is the high yield bond market. Even at the broad market index level, the historically observed distribution of returns has properties that make risk assessment in terms of standard deviation a rather subtle task.

- One might reasonably argue that currently "good estimate" for the annual standard deviation of return to the high yield bond market could range from around 5% to more than 12% depending on what assumptions and sophistication we wish to impart to our analysis.

High Yield Bonds: Empirical Summary Data

- As the basis of our empirical analysis we will use the monthly return history of the Barclay's (formerly Lehman) from July 1, 1983 to December 31st, 2012, a span of near thirty years.
- If we look at the distribution of the monthly returns for the 356 months of data we get the following summary statistics:

Summary Statistics

Mean	0.794387
StDev	2.507523
Median	0.9785
Skew	-0.94826
Kurtosis	8.428941
Mean AbsDev	1.612903

High Yield Bonds: Initial Comments

- Of particular note is the clear evidence of higher moments.
  - The distribution has negative skew (more downside than upside) which should be intuitive for high yield bonds, as the bond either pays the expected stream of cash flow or defaults.
  - The distribution also has a high degree of kurtosis or fat tails, as compared to the normal distribution (kurtosis of a normal = 3).
  - If we make the usual "normal and i.i.d." assumptions, we can annualize the standard deviation by multiplying the monthly value by the square root of twelve, for an annualized value of 8.69%. If asked most market participants for an estimate of the volatility of the returns to this high yield bond index, some value close to 8.69% would be a usual response.

Higher Moments and Volatility

- How must we estimate volatility in the presence of the higher moments?
  - One thing we could do is simply use the more "robust" median and mean-absolute deviation measures for the return distribution. When taken one month at a time, these more robust measures suggest higher returns and lower risks.
- Two possible competing drivers for observed fat tails
  - The first is that the distribution of returns is permanently non-normal and has some other form of stable distribution such as a T-distribution (with low degrees of freedom) or a Gamma distribution.
  - The other possibility is that each moment in time, the distribution of returns is actually normal but the volatility level varies greatly from time to time giving high kurtosis to the total sample. To explore this issue, we simply calculated a 12 month rolling volatility (using the same random walk process) for each possible set of months within the sample period.

Distribution of 12 Month Rolling Volatility

- Mean 7.047272
- Median 5.503421
- Minimum 1.973751
- Maximum 28.60367

- The first thing we should notice is the huge disparity between the minimum value of the 12 month volatility as compared to the maximum value which is roughly fifteen times higher. This suggests kurtosis in the distribution is likely to mostly arise from time variation in the volatility rather than a stable fat tailed distribution, where the standard deviation would be roughly constant through time.

- So if it is our intent of our risk forecast to express the central tendency of the volatility through the next twelve months, a value of roughly 5.5 seems most appropriate, far lower than the 8.69 previously calculated.

Further Discussion

- If we had come to the opposite conclusion that high yield returns were obtained from a fat tailed but stable distribution, we could use the method of Cornish and Fisher (1937) to adjust the mean and standard deviation values to best fit the observed distribution inclusive of higher moments.
- We could also consider the impact of skew and kurtosis on the difference between the geometric and arithmetic means of returns as in Wilcox (2003). Given the standard deviation, skew and kurtosis values (in decimal forms) we can calculate the value of the volatility that would provide the same degree of arithmetic difference between the two forms of the return mean.

$$V \sim (S^2 - (2/3)^* MS^3 + (1/2)^* KS^4)^{.5}$$

- V = adjusted volatility, S = volatility, M = skew, K = excess kurtosis

- Since $M < 0$ and $K > 0$, $V > S$

VaR and CVaR

- Given the negative skew and kurtosis in the return distribution, it may be useful to explore the possibility of using volatility values that are implied from Value at Risk and Conditional Value at Risk estimates
- To estimate the 95% Value at Risk, we simply sort our nearly 30 years of monthly data and find the return value that represents the 95th percentile of outcomes.
  - This is a monthly return of negative 3.27%
  - Under parametric assumptions the equivalent annualized standard deviation value is 6.86%
- Conditional Value at Risk is also called "Expected Shortfall"
  - It is the central tendency value of the events that are worse than the 95% border value
  - This value is negative -5.97%
  - Under parametric assumptions the equivalent annualized standard deviation is 10.05%

The Global Financial Crisis

- Some market participants have also argued that the Global Financial Crisis was an event of very rare and extreme proportion, and is unlikely to reoccur in the foreseeable future.
- To see how this would impact our volatility estimates we simply compute same statistics with the years 2008 and 2009 removed from the sample.

- **Summary Statistics (ex 2008-2009)**

Mean	0.792033
StDev	2.061092
Median	0.9715
Skew	-0.42207
Kurtosis	3.890982
Mean AbsDev	1.416005

"How We Roll" without the GFC

- **Distribution of 12 Month Rolling Volatility (ex 2008-2009)**

Mean	6.227229
Median	5.245626
Minimum	1.973751
Maximum	18.12535

- While the magnitude of skew and kurtosis have been reduced, they remain material. With the GFC period removed, the mean and median of the rolling 12 month volatilities values both decline. With the range between the minimum and the maximum still highly asymmetric around the central tendency, the most reasonable value for a randomly selected 12 month period would be the median of roughly 5.25 (down from 5.5)

Incorporating Serial Correlation

- So far, we have annualized our monthly-based volatility estimates assuming there is no serial correlation in our data, as would be conceptually consistent with a highly liquid market.
  - In the case of the Barclay's index, this assumption is untrue as the market for high yield bonds is relatively illiquid compared to other traded financial assets.
  - For the sample period, the first-order autocorrelation coefficient is .33 ($T > 5$) and is highly statistically significant.
  - This implies that monthly returns to the index are in fact highly predictable, so the actual uncertainty of investors about what returns will occur in the next monthly period is actually far less than is described by the standard deviation (assumes the process is random).

- *The real risk is less because we have very reliable knowledge about what is coming next!*

A Counter Argument

- One can also make a counter argument. Lo, Getmansky and Makarov (2004) argues that for financial markets to operate rationally, they should not be predictable so the financial benefit of any predictability must be offset by trading costs.
- To the extent that trading costs inhibit trading, the changes in observed market prices will lag their true economic values so the volatility we calculate from the return data will understate the true volatility.
- To correct for the degree of autocorrelation in the Barclay's index data, we would need to increase the volatility estimate by 27%.

$$8.69\% * 1.27 = 11.03\%$$

Buy and Hold Investing

- One way to further illustrate this "liquidity driven" dichotomy would be to consider the case of an investor who is aware of the predictability of high yield returns, but decides to trade in or out of the market only once per year no matter what in order to limit possible trading costs.
  - To estimate we simply calculate *annual returns* from one month end to a year later within the sample period. The volatility (standard deviation) of overlapping 12 month returns is 12.63% for the full sample and 10.15% for the sample without the 2008-2009 period.
- Both of these values are larger than the 8.69% value we got originally with our basic sample. This is illustrative of the strong influence of the serial correlation on and the potential high risks that will be encountered by investors who try to "ride out the storm" in event of a downward trend.
  - Higher than the 11.03% derived with the serial correlation correction
  - It should be noted that the degree of liquidity related impact on risk estimation is dependent on the size of the investment portfolio and the likelihood that the investor will decide to exit the market and liquidate all or part of the portfolio
-

Conclusions

- In the presence of higher moments and serial correlation, the estimation of the volatility of a financial asset becomes a complex set of choices.
  - One could rationally estimate an annual volatility measure ranging from around 5% to over 12%
- In the case of the high yield bond market, there are two predominant effects.
  - There is a high degree of variation in the volatility over time. The distribution of 12 month rolling volatility has a large right skew
  - There is a high degree of serial correlation which may be interpreted as either decreasing or increasing risk

References

- diBartolomeo, Dan. “Fat Tails, Tall Tales, Puppy Dog Tails”. Professional Investor, Autumn, 2007.
- Getmansky, M., Andrew Lo and Igor Makarov. “An econometric model of serial correlation and illiquidity in hedge fund returns,” Journal of Financial Economics, 2004, v74(3,Dec), 529-609.
- Wilcox, Jarrod. “Harry Markowitz and the Discretionary Wealth Hypothesis, *Journal of Portfolio Management*, 2003
- Cornish, E., and R. Fisher. 1937. “Moments and Cumulants in the Specification of Distributions.”, *Extrait de la Revue de l’Institute International de Statistique : 1–14.*