An Introduction to Value at Risk

Research Paper 001

March 28, 2013

info@northstarrisk.com

An Introduction to Value at Risk

In this paper we provide an introduction to value at risk. We describe three basic models, and explain how to assess the effectiveness of value at risk models through backtesting. In a subsequent paper we will explore more additional models, and assess their performance using actual historical data.

What is Value at Risk?

Value at risk (VaR) is one of the most widely used risk measures in finance. VaR was popularized by J.P. Morgan in the 1990s. The executives at J.P. Morgan wanted their risk managers to generate one statistic at the end of each day, which summarized the risk of the firm's entire portfolio. What they came up with was VaR.

If the 95% VaR of a portfolio is $100, then we expect the portfolio will lose $100 or less in 95% of the scenarios, and lose $100 or more in 5% of the scenarios. We can define VaR for any confidence level, but 95% has become an extremely popular choice in finance. The time horizon also needs to be specified for VaR. On trading desks, with liquid portfolios, it is common to measure the one-day 95% VaR. In other settings, in which less liquid assets may be involved, time frames of up to one year are not uncommon.

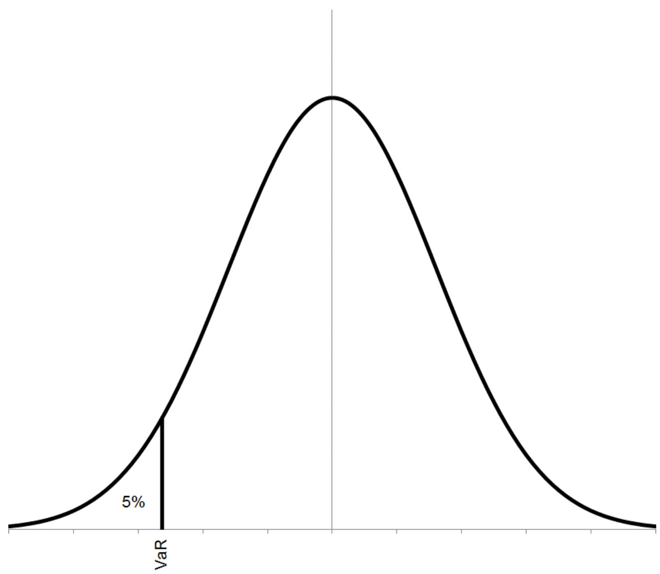

Figure 1

Figure 1 provides a graphical representation of VaR at the 95% confidence level. The figure shows the distribution of returns for a portfolio. 5% of the distribution is to the left of the VaR level, and 95% is to the right.

We can define VaR formally in terms of the loss to our portfolio, L. L is simply the negative of the return to our portfolio. If the return of our portfolio is −$600, then the loss, L, is +$600. For a given confidence level, γ, then, we can define value at risk as:

$$P\left[L \geq \mathrm{VaR}_{\gamma}\right] = 1 - \gamma$$

Equation 1

If the one-day 95% VaR of a portfolio is $400 this means, then γ euqals 95%, and there is a 5%, 1- γ = 5%, probability that the portfolio will *lose* $400 or more on any given day.

If an actual loss exceeds the predicted VaR threshold, that event is known as an exceedance. Another assumption of VaR models is that exceedance events are uncorrelated with each other. In other words, if our VaR measure is set at a one-day 95% confidence level, and there is an exceedance event today, then the probability of an exceedance event tomorrow is still 5%. An exceedance event today has no impact on the probability of future exceedance events. More generally, the probability of an exceedance conditional on all available information should equal the unconditional probability of an exceedance. In other words, an exceedance event should be no more likely to occur on a Tuesday, the day after the market is up, or when the level of risk is high. We return to these assumptions at the end of the paper when we discuss backtesting.

VaR has become extremely popular in risk management. The appeal of VaR is its simplicity. Because VaR can be calculated for any portfolio, it allows us to easily compare the risk of different portfolios. Because it boils risk down to a single number, VaR provides us with a convenient way to track the risk of a portfolio over time. Finally, the concept of VaR is intuitive, even to those not versed in statistics. Because it is so popular, VaR has come under a lot of criticism. As we will see, VaR is not a perfect risk measure, but it is incredibly versatile. While some of the criticism is justified, much of the criticism is misplaced.

Delta-Normal VaR

One of the simplest and easiest ways to calculate VaR is to make what are known as delta-normal assumptions. For any underlying asset, we assume that the log returns are normally distributed, and we approximate the returns of any option based on its delta-adjusted exposure. For portfolios, the delta-normal model assumes that the relationships between securities can be fully described by their correlation.

The delta-normal assumptions make it very easy to calculate VaR statistics even with limited computing power. This made delta-normal models a popular choice when VaR models were first introduced. Predictably, the results of such a simple model were often disappointing. A good risk manager would often be able to compensate for these deficiencies, but the basic model presented an easy target for critics. Delta-normal models are rarely used in practice today, but are still an excellent starting point when learning about VaR models. By understanding the pros and cons of the delta-normal

 info@northstarrisk.com

model we will be better able to understand the pros and cons of more complex models. Unfortunately many people outside of risk management believe that delta-normal models are still widely used in practice, or believe that the shortcomings of these simple models are somehow inherent to all risk models.

To calculate the delta-normal VaR of a security, we start by calculating the standard deviation of returns for the security, or, in the case of an option, for the returns of the option's underlying. For non-option securities, we then multiply the return standard deviation by the absolute market value or notional of our position to get the position's standard deviation. For options we multiply by the absolute delta-adjusted exposure. The delta adjusted exposure being the value of the underlying multiplied by the option's delta. We then multiply the position's standard deviation by an appropriate factor based on the inverse of the standard normal distribution (e.g. −1.64 for 95% VaR).

Notice that we have not said anything about the expected return. In practice, most VaR models assume that the distribution of return has a mean of zero. This is almost always a very reasonable assumption at short horizons. At longer horizons this assumption may no longer be reasonable. Some practitioners will also assume that the time decay for options is also zero. While this assumption may also be valid in many situations, it can fail even over short time horizons. In what follows, unless stated otherwise, assume security returns have zero mean but include theta in calculating VaR.

Historical VaR

Another very simple model for estimating VaR is historical simulation or the historical method. In this approach we calculate VaR directly from past returns. For example, suppose we want to calculate the 1-day 95% VaR for an equity using 100 days of data. The 95th percentile corresponds to the least worst of the worst 5% of returns. In this case, because we are using 100 days of data, the VaR simply corresponds to the 5th worst day. If we have the following 100 returns, sorted from lowest to highest:

	Rank	**Return**
Worst	1	-4.00%
	2	-3.62%
	3	-3.57%
	4	-3.52%
	5	**-3.37%**
	6	-3.24%
	7	-3.14%
	8	-3.07%
	9	-2.92%
	…	…
	99	3.79%
Best	100	4.08%

Table 1

The 95th percentile VaR in this case corresponds to −3.37%, or, dropping the negative sign, we would say that our 1-day 95% VaR is a loss of 3.37%.

For an infinitely-lived underlying asset, the historical approach could not be easier. For derivatives, such as equity options, or other instruments with finite lifespans, such as bonds, it is slightly more complicated. For a derivative, we do not want to know what the actual return series was, we want to know what the return series would have been had we held exactly the same derivative in the past. For example, suppose we own an at-the-money put with two days until expiry. 250 days ago, the option would have had 252 days until expiry, and it may have been far in or out of the money. We do *not* want to know what the return would have been for the option with 252 to expiry, we want to know what the return *would have been* for an at-the-money put with two days to expiry, given conditions in the financial markets 250 days ago. Similarly, for a bond with 30 days to expiry, for risk purposes, we are interested in what the return of a bond with 30 days to maturity *would have been* 250 days ago, *not* what the return of a bond with 280 days to maturity *was*. These constant maturity series, or back-cast series, are quite common in finance. The easiest way to calculate the back-cast series for an option would be to use a delta approximation. If we currently hold a put with a delta of −30%, and the underlying return 250 days ago was 5%, then our back-cast return for that day would be −1.5% , −1.5% = −30% × 5%. A more accurate approach would be to fully reprice the option, taking into account not just changes in the underlying, but time-decay, changes in implied volatility and changes to the risk-free rate. Just as we could approximate option returns using delta, we could approximate bond returns using DV01, but a more accurate approach would be to fully reprice the bond based on changes in the relevant interest rates and credit spreads.

One advantage of historical VaR is that it is extremely simple to calculate. Another advantage is that it is easy to explain to non-risk professionals. Most financial professionals will be used to looking at cumulative return charts. The returns used to create a cumulative return chart are the same returns used to calculate historical VaR. If there is ever a question about the validity of a historical VaR calculation it is easy enough to pull up a chart of historical returns to look for a potential source of error.

The delta-normal approach is an example of what we call a parametric model. We say that the model is parametric because it is based on a mathematically defined, or parametric, distribution (in this case, the normal distribution). By contrast the historical approach is non-parametric. We have not made any assumptions about the distribution of historical returns. There are advantages and disadvantages to both approaches. The historical approach easily reproduces all the quirks that we see in historical data, changing standard deviation, skewness, kurtosis, jumps, etc. Developing a parametric model that reproduces all of the observed features of financial markets can be very difficult. At the same time, models based on distributions often make it easier to draw general conclusions. In the case of the historical approach, it may not be easy to tell if a VaR forecast is the result of a particularly unusual set of input returns.

Monte Carlo Simulation

Monte Carlo simulations are widely used throughout finance, and they can be a very powerful tool for calculating VaR. As an example of how we would calculate VaR using a Monte Carlo simulation, imagine we have a position in gold, and we believe that the daily log returns of gold are normally distributed with a mean of 0.01% and a standard deviation of 1.40%. To calculate the VaR of this position, we could generate 1,000 draws from a normal distribution with a mean of 0.01% and a standard deviation of 1.40%, convert the log returns into standard returns, and then sort the returns from lowest to highest. If we are interested in our 95% VaR, we simply select the 50th worst return from the list. For this set up, the Monte Carlo simulation is very straightforward, but it is also inefficient. Because the log returns are normally distributed, we know that the 5th percentile is −1.64 standard deviations below the mean, corresponding to a log return of −2.29% = 0.01% −1.64 × 1.40%.

The real power of Monte Carlo simulations is in more complex settings, where instruments are non-linear, prices are path dependent, and distributions do not have well defined inverses. Mote Carlo simulations can also be useful when the relationships between securities are more complicated.

Monte Carlo simulations also make it easy to calculate multi-period VaR. In the preceding example, if instead of being interested in the 1-day VaR, we wanted to know the 4-day VaR, we would simply generate four 1-day log returns, using the same distribution as before, and add them together to get one 4-day return. We could repeat this process 1,000 times, generating a total of 4,000 1-day returns. As with the 1-day example, in this particular situation, there are more efficient ways to calculate the VaR statistic. That said, it is easy to imagine how multiday scenarios could quickly become very complex. What if your policy was to reduce your position by 50% every time you suffered a loss in excess of 3%? What if returns exhibited positive serial correlation, with positive excess returns more likely to be followed by positive excess returns, and negative excess returns more likely to be followed by negative excess returns?

Monte Carlo simulations are usually based on parametric distributions, but we could also use non-parametric methods, randomly sampling from historical returns. Continuing with our gold example, if we had 500 days of returns for gold, and we wanted to calculate the 4-day VaR, we would randomly pick a number from 1 to 500, and select the corresponding historical return. We would do this four times, to create one 4-day return. We can repeat this process, generating as many 4-day returns as we desire. The basic idea is very simple, but there are some important details to keep in mind. First, generating multi-period returns this way involves what we call sampling with replacement. Pretend that the first draw from our random number generator is a 10, and we select the 10th historical return. We don't remove that return before the next draw. If, on the second draw, our random number generator produces 10 again, then we select the same return. If we end up pulling 10 four time in a row, then our 4-day return will be composed of the same 10th return repeated four times. Even though we only have 500 returns to start out with, there are 500^4, or 62.5 billion, possible 4-day returns that we can generate this way. This method of estimating parameters using sampling with replacement is often referred to as bootstrapping. The second detail that we need to pay attention to is serial correlation. We can only generate multi-period returns in the way just describe if single-period returns are independent of each other. For example, suppose that the standard deviation of gold has gone through long periods of high volatility followed by long periods of low volatility, and we believe our historical data accurately reflects this starting with 250 days of low volatility followed by 250 days of high volatility. If we randomly select returns with replacement, then the probability of getting a draw from the high volatility period is 1/2 each time. If our random numbers are generated independently then there is only $1/16 = (1/2)^4$ chance of drawing four returns in a row from the high period, whereas historically the probability was much closer to 1/2 (except for the transition in the middle of the sample, where we switched from low to high volatility, low volatility days were always followed by low volatility days, and high volatility days were always followed by high volatility days). A simple solution to this problem: instead of generating a random number from 1 to 500, generate a random number from 1 to 497, and then select four successive returns. If our random number generator generates 125, then we create our 4-day return from returns 125, 126, 127, and 128. While this method will capture any serial dependence between periods, it greatly reduces the number of possible returns. In this case, the number of possible 4-day returns is reduced from 62.5 billion to 497, and effectively reduces the Monte Carlo simulation to the historical simulation method.

Of the three methods we have considered so far, Monte Carlo simulations are generally considered to be the most flexible. Their major drawback is speed. As computers get faster and faster, the speed of Monte Carlo simulations is becoming less of an issue. Still in some situations — a trading desk that require real-time risk number, for example —this speed issue may still rule out the use of Monte Carlo simulations.

Hybrid VaR

For the historical method or historical simulation, all of the data points are given equal weight. In practice, because market risk tends to changes over time, it might make sense to give more weight to more recent data. One very simple way to do this is to apply exponentially decreasing weight to the historical data. For example, with a decay factor of 0.99, we would apply a weight of 1.00 to the most recent return, 0.99 to the second most recent, 0.99^2 to third, and so on.

This general approach, using historical returns with decreasing weights, is often called the hybrid approach because it combines aspects of standard historical simulation and weighted parametric approaches; see, for example, Allen, Boudoukh, and Saunders (2004).

Suppose we have 100 returns, sorted from worst to best, as before. In the case of the historical simulation, we found the VaR by moving down the table until we got to the fifth data point. For the hybrid approach we simply move down the table until we get to 5% of the total weights:

	Rank	*t*	Return	Weight	% of Total Weight	Cumulative % Weight
Worst	1	50	-4.00%	0.61	0.95%	0.95%
	2	40	-3.62%	0.55	0.86%	1.82%
	3	9	-3.57%	0.40	0.63%	2.45%
	4	48	-3.52%	0.59	0.94%	3.38%
	5	37	-3.37%	0.53	0.84%	4.22%
	6	**18**	**-3.24%**	**0.44**	**0.69%**	**4.91%**
	7	**54**	**-3.14%**	**0.63**	**0.99%**	**5.91%**
	8	63	-3.07%	0.69	1.09%	7.00%
	9	8	-2.92%	0.40	0.63%	7.62%
	…	…	…	…	…	…
	11	100	-2.80%	1.00	1.58%	10.60%
	…	…	…	…	…	…
	99	41	3.79%	0.55	0.87%	99.32%
Best	100	16	4.08%	0.43	0.68%	100.00%

Table 2

In this case, we get to 5% of the total weight between the sixth and seventh returns. At this point there are two approaches to deciding the VaR. The more conservative approach is to take the sixth return, –3.24%. The alternative is to interpolate between the sixth and seventh returns, to come up with –3.23%. Unless there is a strong justification for choosing the interpolation method, the conservative approach is recommended.

The hybrid approach is fairly easy to implement, computationally efficient, and is underpinned by familiar historical returns. While only slightly more complicated than the delta-normal and historical simulation methods, the hybrid approach often produces much more reliable estimates of VaR.

Backtesting

An obvious concern when using VaR is choosing the appropriate confidence interval. As mentioned, 95% has become a very popular choice in risk management. In some settings there may be a natural choice for the confidence level, but most of the time the exact choice is arbitrary.

A common mistake for newcomers is to choose a confidence level that is too high. Naturally, a higher confidence level sounds more conservative. A risk manager who measures one-day VaR at the 95% confidence level will, on average, experience an exceedance event every 20 days. A risk manager who measures VaR at the 99.9% confidence level expects to see an exceedance only once every 1,000 days. Is an event that happens once every 20 days really something that we need to worry about? It is tempting to believe that the risk manager using the 99.9% confidence level is concerned with more serious, riskier outcomes, and is therefore doing a better job.

The problem is that, as we go further and further out into the tail of the distribution, we become less and less certain of the shape of the distribution. In most cases, the assumed distribution of returns for our portfolio will be based on historical data. If we have 1,000 data points, then there are 50 data points to back up our 95% confidence level, but only one to back up our 99.9% confidence level. As with any distribution parameter, the variance of our estimate of the parameter decreases with the sample size. One data point is hardly a good sample size on which to base a parameter estimate.

A related problem has to do with backtesting. Good risk managers should regularly backtest their models. Backtesting entails checking the predicted outcome of a model against actual data. Any model parameter can be backtested.

In the case of VaR, backtesting is easy. Each period can be viewed as a Bernoulli trial. In the case of one-day 95% VaR, there is a 5% chance of an exceedance event each day, and a 95% chance that there is no exceedance. Because exceedance events are expected to be independent, over the course of *n* days the distribution of exceedances follows a binomial distribution:

$$P[K = k] = \binom{n}{k} p^k (1 - p)^{n-k}$$

Equation 2

In this case, *n* is the number of periods that we are using to backtest, *k* is the number of exceedances, and (1 − *p*) is our confidence level.

As an example, pretend we have been forecasting the 1-day 95% VaR for a portfolio for the past 100 days. The probability of getting exactly 4 exceedance would be:

$$P[K = k] = \binom{100}{4} 0.05^4 (0.95)^{96} = 17.81\%$$

Table 3 shows the probability of getting between 1 and 12 exceedance over 100 days for a 1-day 95% VaR. As we might expect, the most likely outcome is 5 exceedances. What might not be so obvious is that the probability of getting 4 or 6 exceedances is not that much less.

The last two columns of Table 3 show the probability of seeing *n* or less, or *n* or more exceedances over 100 days. There is actually a 23.40% probability of seeing 7 or more exceedances, so we should not be too surprised if we see 7 exceedances. On the other hand, there is less than a 1% chance of seeing no exceedances, or 12 or more. If we see no exceedances, or 12 exceedances, we are either experiencing highly unusual markets, or our model is not working like it should. Seeing 0 or 12 exceedances should be a clear red flag to a risk manager, and lead to further investigation.

n	**P[*X*=*n*]**	**P[*X*≤*n*]**	**P[*X*≥*x*]**
0	0.59%	0.59%	100.00%
1	3.12%	3.71%	99.41%
2	8.12%	11.83%	96.29%
3	13.96%	25.78%	88.17%
4	17.81%	43.60%	74.22%
5	18.00%	61.60%	56.40%
6	15.00%	76.60%	38.40%
7	10.60%	87.20%	23.40%
8	6.49%	93.69%	12.80%
9	3.49%	97.18%	6.31%
10	1.67%	98.85%	2.82%
11	0.72%	99.57%	1.15%
12	0.28%	99.85%	0.43%

Table 3

As stated at the beginning of the paper, another assumption of the VaR model is that the probability of an exceedance event, conditional on all available information, is equal to the unconditional probability of an exceedance. For example suppose we looked back over several years of data and observed that 30% of exceedances occurred on Tuesday and only 10% occurred on Thursdays. In theory, an exceedance should be equally likely on any given day of the week. In the long run, we would expect to see the exceedances equally allocated to each of the five business days, or 20% to each. Of course the 30%/10% split between Tuesday and Thursday could just be a fluke. However, if the split turned out to be statistically significant, we would need to make an adjustment to our model.

A common problem with VaR models is serial correlation in exceedances. The probability of an exceedance today, given that an exceedance occurred yesterday, should be no higher or lower than on any other day. A simple way to test for this serial correlation is to count the number of exceedances, and then count the number of times there is an exceedance on the day after an exceedance. The distribution of the day-after exceedances should also follow a binomial distribution. If we are forecasting a 1-day 95% VaR, and there have been 40 exceedances, then we would expect that 2 of those exceedances occurred after another of the 40 exceedances, $2 = 40 \times 5\%$.

Another common problem with VaR models is the tendency for the probability of exceedances to vary with the level of risk. This may seem counterintuitive, but exceedances should be no more likely to occur when risk is high than when risk is low. If we are measuring 95% VaR then there is always a 5% chance of an exceedance. To test that there is no correlation between exceedances and the level of risk we could divide our sample into high and low risk days. Just as before, we could test the significance of the number of exceedances in each subset using the binomial distribution.

Conclusion

In this paper we introduced three basic value at risk models, the delta-normal model, the historical model, and the hybrid model. We also established criteria for assessing the effectiveness of VaR models through backtesting. The first two models, the delta-normal model and the historical model, tend not to perform very well, but they do serve as a starting point for building more complex models, such as the hybrid model. Taken together, the three models start to suggest some of the issues that we will need to consider when building more complex VaR models, including non-normal returns, the impact of derivatives, and changing volatility. In a subsequent paper we will examine the effectiveness of these models and others using historical data.

info@northstarrisk.com

References

Allen, Linda, Jacob Boudoukh, and Anthony Saunders. 2004. *Understanding Market, Credit, and Operational Risk: The Value at Risk Approach*. Malden, MA: Blackwell Publishing.

 info@northstarrisk.com