

Risk Management Primer

Pete Benson

pbenson@umich.edu



Department of Mathematics

530 Church Street, 2082C East Hall

Ann Arbor, MI 48109-1043, USA

Chapter 1

What is risk management?

I started writing this to expand on what students learn about risk management during their Master's studies in Quantitative Finance and Risk Management at the University of Michigan. The term *risk management* has many interpretations, so we will narrow this to two relevant subjects.

When people speak of risk management in finance, they are typically talking about one of two different sets of questions:

1. Given a trade that contains some form of undesirable risk, can we efficiently mitigate that risk?
2. What is the risk of our portfolio, where is the risk coming from, and what if anything might we do about it?

The first we might call trade risk management, and the second *portfolio risk management*. You might think of these as a reductionist versus a holistic view of risk management. Both are important. We will focus on portfolio risk management, but we must understand the distinction.

1.1 Trade risk management

Suppose a USD-based trader believes AAPL is undervalued, and buys Apple stock. He does not want to be exposed to changes that can be attributed to movements in the market. While believes the price of Apple is low, he recognizes that if the entire market falls, AAPL is likely to fall as well, though perhaps not as much. He could **short** the SP500 index just enough to remove his exposure to the market.

Or perhaps our AAPL investor doesn't want to be exposed to the risk that AAPL price falls, regardless of what happens in the market. He can buy put options on AAPL so that if the price falls, the put options cover all losses on the stock.

Another trader works for a London firm, and also believes AAPL is undervalued. This trader buys AAPL by selling GBP for USD, then buying the

AAPL stock. She now has currency risk: AAPL may rise, but if the GBP/USD rate falls, she may not see any profits when she sells the stock. If she expects to sell the AAPL stock in one year (receiving USD), she can sell USD one year forward via FX futures contracts.

1.2 Portfolio risk management

A risk manager is responsible for tracking the **market risk** of a firm's portfolio, using statistics such as **VaR** and **expected shortfall**. Given their capitalization, the firm has chosen a 5% quantile one-day VaR limit of USD 20m, with expected shortfall 30m. This means that the probability of losing more than 20m in one day should be no more than 5%, and that in the event that the loss is greater than 20m, it will on average be 30m. When the risk manager arrives in the morning, she checks the risk reports, and notes that the firm's expected shortfall is USD 32m. Looking further, she notes that the swap desk is responsible for USD 8m of the expected shortfall, but the desk is not supposed to have more than 6m. Drilling in, she notes volatility has increased for some parts of the yield curve, and there are several large trades that are not suitably hedged. She pays a visit to the head of the swap desk to discuss.

Chapter 2

Measuring Market Risk

2.1 A very brief history

Market risk is the risk of financial loss due to changes in security prices. To manage this risk, you must first measure it. The measure is a statistic summarizing the forecast of a **nonstationary** process out to some given horizon.

The P&L (profits and losses) between now and the horizon is represented by a random variable. Risk management selects one or more models for the distribution of that random variable, and characterizes the distribution with descriptive statistics. Markowitz's landmark 1952 paper *Portfolio Selection* used standard deviation as the measure. As the variety of securities increased, including options with very asymmetric payoffs, other risk measures became more popular. In 1994, J.P. Morgan put forth RiskMetrics with VaR suggested as a better measure of risk than standard deviation. VaR has its own drawbacks which we will discuss, and expected shortfall is today generally preferred.

2.2 A portfolio with a single risk factor

A risk factor is a random variable that affects portfolio value. It might represent the return on a stock, commodity, currency, or other security. It might be an interest rate, an implied volatility, or some kind of index. As we increase the number of factors, the complexity of our model increases. So we start with a portfolio whose value is driven by a single factor: the return on a stock.

Imagine you hold 100 shares of a stock trading at USD 200/share. Your simple portfolio is worth USD 20,000. Tomorrow at this time it will have a different value, assuming the market is open in the next 24 hours. The price of the stock may go up or down, and you will realize a profit or loss, accordingly. The P&L is a random variable, which we will call P . Here are some ways we can model P :

1. **Parametric:** Assume a distribution for P , (e.g. normal), and estimate the parameters for the distribution.

2. **Historical simulation:** Assume that the relative change in AAPL price tomorrow will be exactly like a randomly selected day from an historical period (say, the past year).
3. **Monte Carlo simulation:** Assume a distribution for each of the factors affecting the value of your stock (for now, just the stock price itself). Generate many market scenarios, and revalue your portfolio in each scenario. The distribution of scenario P&L is an approximation of the distribution of P .

Each of these models requires historical pricing of AAPL stock. Ordinarily, you want at least several months of data. For simplicity, however, we will use 4 days. Suppose the stock had these prices:

date	price
20180906	200
20180905	198
20180904	194
20180903	206

Rather than modeling the price directly, we will model the return R on the stock price, based on the observed returns from our dataset. For a price that goes from p to p' , we might compute the observed return r via

$$p' = p(1 + r) \tag{2.1}$$

which gives $r = \frac{p'}{p} - 1$. Often, we will prefer to treat the return like a continuously compounded rate:

$$p' = pe^r \tag{2.2}$$

giving us $r = \ln \frac{p'}{p}$. In practice, if returns are small ($< 10\%$), the two methods give similar results. As an example, the return from 20180905 to 20180906 is $\ln \frac{200}{198} \approx \frac{200}{198} - 1 \approx 0.01$. We add a column for the return to our table:

date	price	return
20180906	200	0.01
20180905	198	0.02
20180904	194	-0.06
20180903	206	

Note that there is no return for 20180903, since we don't know the previous price. In general, if you have $n + 1$ historical prices, you will have n returns.

We will use this data for each of our three modeling methods.

Parametric model

We assume the return R is normally distributed, and we estimate the standard deviation (aka *volatility*) σ from the observed returns. For reasons that become

clear when we have a more complex portfolio, we want P&L , P , to also be normally distributed. Hence, we use equation 2.1 to compute returns.

What about the mean return? Forecasting return is especially difficult. Moreover, since our risk horizon is typically quite short (from 1 day to perhaps 10 days), whatever expected return there might be is dominated by the variance. Hence, we take the mean return to be zero, so the variance estimate is

$$\hat{\sigma}^2 = \frac{1}{n} \sum r_i^2 \quad (2.3)$$

where n is the number of observed returns.

2.1 When computing variance, why would we divide by n rather than $n - 1$?

2.2 For our simple dataset with three returns, what is the standard deviation of the returns?

For the parametric model, the P&L of a **position** is the return times the value v of the position, i.e. $P = vR$. Likewise, $\Sigma_P = v\sigma_R$.

2.3 What is the standard deviation of P&L , Σ_P , for our portfolio?

2.4 How many standard deviations from the mean is the 5% quantile for a normal distribution?

2.5 What is the 5% quantile P&L for our single stock portfolio?

VaR numbers are always quoted with a quantile (e.g. 5%) or confidence level (e.g. 95%), and a horizon (e.g. one day). Also, because VaR is about forecasting potential losses, a loss is reported as a positive number.

2.6 What is the one day 5% quantile VaR for our portfolio of one stock?

2.7 If instead of long, we were short 100 shares, will that affect the VaR?

2.8 What do you think could be drawbacks of the parametric model?

Historical simulation

Historical simulation assumes that between now and the horizon, returns will be exactly like a period of the same length randomly selected from the historical record. For our simple dataset, this means the return tomorrow on today's portfolio value will be either 0.01, 0.02, or -0.06, each with probability 1/3.

2.9 What is the 50% quantile P&L for our portfolio under historical simulation?

2.10 What is the 50% quantile P&L for our portfolio under the parametric approach?

Consider a continuous random variable X with distribution function F .

2.11 What is the distribution of $F(X)$?

If X is observed n times, we can sort the observations yielding *order statistics* $X_{(1)} \leq X_{(2)} \leq X_{(3)} \dots \leq X_{(n)}$. Let $U = F(X)$. Can we say something about the order statistics on the quantiles $U_{(1)} \leq U_{(2)} \leq U_{(3)} \dots \leq U_{(n)}$?

2.12 Write a Monte Carlo simulation in Python to estimate the expected value of $U_{(k)}$ given n draws.

2.13 Determine the distribution of $U_{(k)}$, and calculate its mean.

2.14 Given n historical returns, what quantiles can we report?

2.15 How do you think we should handle other quantiles?

2.16 What do you think could be drawbacks of the historical simulation model?

For the purposes of exposition, having just three observations is convenient, but not representative of reality. You can download historical data from a variety of sources. Yahoo Finance (<https://finance.yahoo.com/>) is a quick way to get started.

For the following problems, download a year's worth of historical price data from Yahoo Finance for Apple for 2017-09-01 to 2018-09-01. Use closing prices for computing log returns. Build a Google doc or Excel spreadsheet to answer the questions. Assume the current price for Apple is yesterday's closing price.

2.17 What is the standard deviation of the percent returns?

2.18 What is the standard deviation of the log returns?

2.19 If you had 100 shares of AAPL at USD 220/share, what is your one day, 95% confidence level VaR, using the parametric approach?

2.20 Take a look at your spreadsheet's PERCENTILE() function. Test it on a set of data (e.g. 1, 2, 4). Does it give the answers you would expect for the 25% quantile? If not, can you figure out how to make it give the "correct" answer?

2.21 If you had 100 shares of AAPL at USD 220/share, what is your one day, 95% confidence level VaR, using historical simulation?

2.22 In Python, implement historical simulation on 100 shares of AAPL. Inputs should be a filename of historical data, a number of shares, current price, and a list of quantiles. Use `pandas` to read your input. Output is VaR for each of the quantiles.

2.23 What happens to VaR when your position is short instead of long?

2.24 Implement `expected shortfall` in your Python program. What is the expected shortfall for long 100 shares AAPL? What is the expected shortfall for short 100 shares AAPL?

Monte Carlo simulation model

Monte Carlo shares some characteristics with the parametric and historical simulation methods.

Like parametric, Monte Carlo is based on an underlying normal distribution. Parametric, however, often assumes P&L is normally distributed, whereas Monte Carlo must assume log returns are normal. For our single stock, we assume log return $R \sim N(0, \sigma^2)$, where σ is computed from the log returns.

Like historical simulation, we arrive at statistics about the distribution by looking at a set of observations. Historical, however, is nonparametric, and has a fixed number of simulations equal to the number of observed returns. Monte Carlo can offer an unlimited number of simulations,

2.25 Why might we prefer Monte Carlo simulation over the parametric approach for computing risk statistics?

2.26 Why might we prefer Monte Carlo simulation over historical simulation for computing risk statistics?

2.27 Why model log returns rather than percentage returns as normal?

2.28 Build a simulation in a spreadsheet with 999 simulations of your portfolio (hint: google `RAND()` and `NORMSINV()` for spreadsheets). Make the standard deviation (i.e. the volatility) an input you can adjust. Compute 5% quantile VaR.

2.29 What is 5% quantile VaR in your spreadsheet model if you are `short` 100 shares? In the limit, as you increase the number of simulations, should it matter whether you are long or short?

2.30 What happens to VaR on the short portfolio as you increase volatility to high levels (e.g. 80%)?

2.31 In Python, implement Monte Carlo simulation of AAPL. Inputs should be a filename of historical data, a number of shares, current price, a list of quantiles, and number of simulations. Output is VaR and expected shortfall for each of the quantiles. Note that your results will vary each time unless you set the random seed at the beginning of your program.

2.32 Using numpy and pyplot, graph the histogram of your simulated P&L.

2.3 A portfolio with two or more factors

Adding a second random variable to our problem means that we must think of their joint distribution. We will eventually generalize to n factors, but the ideas can be easier to illustrate in the case where $n = 2$.

We are going to keep the security model simple by just adding a second stock. Note, however, that even a single security might involve multiple factors. Modeling a portfolio of a single stock option can require factors for stock, interest rate, and implied volatility.

Parametric model

Markowitz specified the relationship between factors via a covariance matrix. Let $\vec{R} = \{R_1, \dots, R_n\}$ be vector of random variables representing percent returns on factors, with covariance matrix Σ . We assume we can assign a weight w to each factor, representing how much of that factor is in our portfolio. For a stock, the position size is the market value of the position, which is the number of shares times the share price. Let these position weights be denoted $\vec{w} = \{w_1, w_2, \dots, w_n\}$. Then we can compute the portfolio variance

$$\sigma^2 = \vec{w}'\Sigma\vec{w} \quad (2.4)$$

For the following exercises, retrieve price data for Google stock from Yahoo Finance. Use the same range of dates you used for Apple. You will use both stocks.

2.33 In your spreadsheet, compute the variance of each stock, and the covariances.

2.34 With 500 shares of AAPL at USD 220/share, and 100 shares of GOOG at USD 1200/share, what is your portfolio's standard deviation?

2.35 What happens to the portfolio standard deviation if you are short 100 shares of GOOG? Why?

Historical simulation

Extending historical simulation to handle more factors is more straightforward than the parametric model.

2.36 Make a copy of your AAPL simulation spreadsheet, and add a P&L column for GOOG. Provide cells for changing the size of your positions.

Appendices

Appendix A

Additional reading

1. [A general approach to calculating VaR without volatilities and correlations](#). This describes a method for multivariate normal Monte Carlo on large numbers of factors without requiring decomposing a covariance matrix. This method is used today by [MSCI RiskMetrics RiskManager](#).
2. [RiskMetrics™—Technical Document](#). This is the 1996 revision of the original 1994 technical document. Provides a good background on portfolio risk measurement.
3. [Return to RiskMetrics](#). This is the 2004 update to the RiskMetrics Technical Document. Stronger emphasis on Monte Carlo,