

FMO6 — Web:

<https://tinyurl.com/yca1oqk6> Polls:
<https://pollev.com/johnarmstron561>

Lecture 7

Dr John Armstrong

King's College London

August 22, 2020

Modern Portfolio Theory

Modern Portfolio Theory

Theory developed by H. Markowitz in 1950's. Not very modern!

- There are n assets whose returns over some fixed time period T are random variables R_i .

-

$$R_i = \frac{\text{price of asset } i \text{ at time } T - \text{current price of asset } i}{\text{current price of asset } i}$$

- The random variables R_i are assumed to have a known mean μ and covariance matrix Σ .
- We wish to invest a fixed amount P over the time period T in a fixed portfolio of the assets.
- The portfolio is determined by choosing the weights w_i to assign to each asset with $\sum w_i = 1$.
- We are allowed to short sell so the w_i may take any real values.

Optimization problem

- Suppose we wish to have an expected return on our portfolio of r , what portfolio should we choose to minimize the standard deviation of the portfolio?
- If \vec{w} is the vector of weights, then the return of the portfolio has a distribution with mean:

$$\vec{\mu}^T \vec{w}$$

and variance:

$$\vec{w}^T \Sigma \vec{w}$$

- So problem can be stated as:
 - Minimize $\vec{w}^T \Sigma \vec{w}$
 - Subject to the constraints $\sum_i w_i = 1$ and $\vec{\mu}^T \vec{w} = r$.
 - Σ is a positive definite matrix, $\vec{\mu}$ is a vector.

Quadratic optimization problem

- Minimize

$$\frac{1}{2} \vec{x}^T H \vec{x} + f^T \vec{x}$$

- Optionally subject to constraints of the form:

$$A \vec{x} \leq \vec{b}$$

$$A' \vec{x} = \vec{b}'$$

$$\vec{l} \leq \vec{x} \leq \vec{u}$$

For some matrices A and A' and vectors \vec{b} , \vec{b}' , \vec{l} and \vec{u} .

- We write $\vec{x} \leq \vec{y}$ if every entry of the vector x is less than or equal to the corresponding entry in y .
- The vectors \vec{l} and \vec{u} may contain the values $-\infty$ and $+\infty$ respectively.
- To solve this problem simply use `quadprog`

MATLAB quadprog

For a problem as defined on the previous slide with $A' = Aeq$ and $b' = beq$

```
[x, fVal, exitFlag]=quadprog(H,f,A,b,Aeq,beq,lb,ub,x0,options);
```

- x is the optimal value of the vector
- $fVal$ is the value of $\frac{1}{2}\vec{x}^T H \vec{x} + \vec{f}^T \vec{x}$
- `exitFlag` indicates whether the optimization worked. You must check this!
- You can use the empty array `[]` where a constraint is not needed.
- The `options` parameter allows you to fine tune the optimization if you wish. It can be omitted.
- `quadprog` standard for “quadratic programming”

Markowitz problem as quadratic optimization

- The condition $\sum_i w_i = 1$ and $\vec{\mu}^T \vec{w} = r$ can be rewritten $A' \vec{w} = \vec{b}'$ where

$$A' = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ \mu_1 & \mu_2 & \mu_3 & \dots & \mu_n \end{pmatrix}, \quad b' = \begin{pmatrix} 1 \\ r \end{pmatrix}$$

- Thus the Markowitz problem is just a special case of quadratic programming.

MATLAB implementation

```
function [ ret, variance, w ] = markowitzOptimizeRet(...  
r, mu, sigma )  
  
H = sigma;  
n = size(sigma,1);  
f = zeros(n,1);  
Aeq = [ ones(1,n); mu];  
beq = [1; r];  
  
[w,~,exitFlag] = quadprog(H,f,[],[],Aeq,beq,[],[],[]);  
assert(exitFlag>0);  
ret = mu * w;  
variance = w' * sigma * w;  
  
end
```


Remarks

- If you run the code on the previous slide, MATLAB prints out a lot of messages and a warning.
- To switch off the messages you can customize the options
- To switch off the warning you use the `warning` command to indicate that you aren't interested in the specific warning message.

Quieter MATLAB implementation

```
function [ ret, variance, w ] = markowitzOptimizeRet( ...
    r, mu, sigma )
warning('off', 'optim:quadprog:WillRunDiffAlg');

H = sigma;
n = size(sigma,1);
f = zeros(n,1);
Aeq = [ ones(1,n); mu];
beq = [1; r];

options = optimset('quadprog');
options = optimset(options,'Display','off');
[w,~,exitFlag]=...
    quadprog(H,f,[],[],Aeq,beq,[],[],[],options);
assert(exitFlag>0);
ret = mu * w;
variance = w' * sigma * w;

end
```

Creating options

- To create a set of options first call `optimset` passing in a single parameter, the name of the optimization routine you will call. This returns an `options` object.

```
options = optimset('quadprog');
```

- Call `optimset` a second time to modify the options. This time you should pass in your `options` object and then a set of key/value pairs that indicate what options you would like to set. We're setting the option `Display` to `off`.

```
options = optimset(options, 'Display', 'off');
```

The function `optimset` returns the new modified options object.

Using the options

- You can now pass this customized set of options object to the optimization function.

```
[w,~,exitFlag]=...  
quadprog(H,f,[],[],Aeq,beq,[],[],[],options);
```

- The reason for this procedure is to make sure that you choose the default values for all options except the ones you wish to customize.

Available options

When optimizing in MATLAB you can use options to set things like:

- How much information to print during the calculation (`Display`)
- The actual algorithm to use (`Algorithm`)
- How accurate the answer needs to be
- The maximum number of steps to perform of the algorithm before giving in
- etc.
- Consult the MATLAB documentation for details of the available options and their settings.

Portfolio optimization example

- We wish to select an optimal portfolio of FTSE 100 stocks.
- The file `ukx.xlsx` contains data downloaded from Bloomberg for the FTSE 100 index.
- It contains weekly prices for each of entry of the FTSE 100 since the year 2000
 - For the vast majority of FTSE 100 companies, we have a complete series of data
 - For a couple of companies there are missing entries. We will exclude those from our Portfolio.
- For the remaining companies, we'd like to compute the weekly returns and hence estimate the sample covariance matrix Σ and the mean weekly vector $\vec{\mu}$ for the FTSE 100.
- We then assume that over the next week, the stock returns follow a multivariate normal distribution with parameters $\vec{\mu}$ and Σ .

Reading the Excel data

```
function data = ukxData( nSecurities )

% Read the raw data from the excel file
bloombergData = xlsread( 'ukx.xlsx', 'A3:KP736' );

% n is the number of securities that we've read
n = min( nSecurities, floor(size(bloombergData,2)/3));

% Now eliminate the empty columns and any data
% for a security where
% we don't have a full history of returns
index = 1;
for i=1:n
    col = bloombergData(:,(i-1)*3+1);
    if (~isnan(col(end)))
        data(:,index)=col;
        index = index+1;
    end
end
end
```

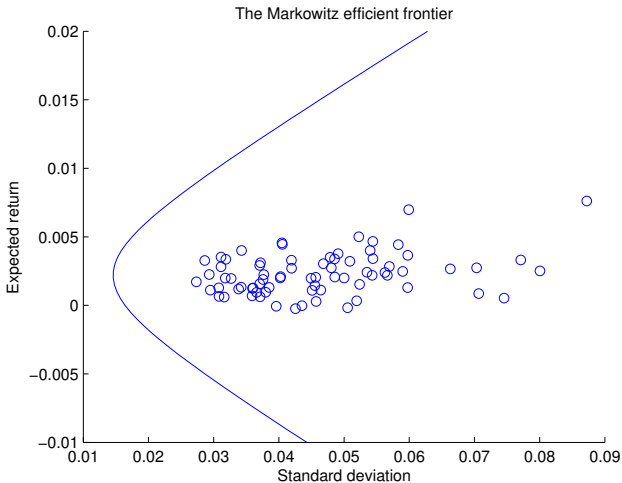
Scatter plot of UKX returns and sds

```
data = ukxData( nSecurities );
returns = (data(2:end,:) - data(1:end-1,:))./data(1:end-1,:);

% We assume that historical returns allow us to estimate
% expected return (mu) and covariance (sigma)
sigma = cov(returns);
mu = mean( returns );

% Plot a scatter plot of the return for each constituent
sds = sqrt( diag( sigma ) );
scatter( sds, mu );
xlabel('Standard deviation');
ylabel('Expected return');
```


Efficient frontier using all stocks

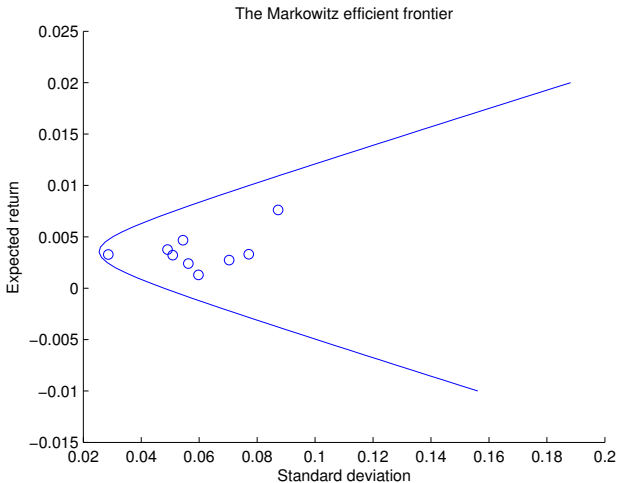


- Given an expectation r , there is a corresponding minimum standard deviation.
- As r varies, this traces out the efficient frontier.
- The points represent the performance of individual stocks in the FTSE 100
- They all lie inside the efficient frontier.

Plotting the efficient frontier

```
% Now compute the efficient frontier and plot it
r = -0.01:0.0005:0.02;
frontierX = zeros(length(r),1);
frontierY = zeros(length(r),1);
for i=1:length(r);
    [ret,var]=markowitzOptimizeRet( r(i), mu, sigma );
    frontierX(i)=sqrt(var);
    frontierY(i)=ret;
end
hold on;
plot( frontierX, frontierY );
title('The Markowitz efficient frontier');
hold off;
```

Efficient frontier with 10 stocks



Remarks on Markowitz theory

- The Markowitz model only makes sense if you think standard deviation is a good measure of risk. For normal distributions, this is uncontroversial. For more complex distributions more general utility optimization makes more sense. This is why normally distributed returns are often assumed.
- For large markets, the assumptions are very unrealistic
 - Precisely known drift and covariance matrix for returns
 - Unlimited buying and selling
- It tends to produce unreasonable portfolios that exploit quirks in the data.
 - Example: suppose X and Y are two highly correlated stocks with the same price. Markowitz optimization might construct a portfolio consisting of 1000 units of X and sell 999 units of Y .
- Despite its shortcomings, it is enormously influential and is where the idea of portfolio optimization originated.

Key Skills for the exam

- Given a financial problem similar to the Markowitz problem, can you write down a formal mathematical statement of the problem?
- Given a mathematical statement of a programming problem, can you write down the matrices you would need to call a MATLAB function like `quadprog`?


Risk Management - How to measure risk for non-normal distributions

VaR

- The definition of Value at Risk requires you to choose a model for security prices. Once this model is chosen:

Definition

The 5%, 10-day Value at Risk of a position is the 95-th percentile of the loss distribution over a 10-day time horizon.



Definition

The $p\%$, d -day Value at Risk of a position is the $(100 - p)$ -th percentile of the loss distribution over a d -day time horizon.

Informally this is the same thing as:

Definition

The $p\%$, d -day Value at Risk of a position is the maximum amount of money you lose in the $(100 - p)\%$ best case scenarios over a d -day time horizon.

Typical values for p are 5% and 1%. Typical values for d are 1-day, 1-month and 1-year.

One way banks are regulated is via the “Basel accords”. Basel II said capital requirement constraints should be set in terms of VaR so VaR became the industry standard risk figure. Basel III uses expected shortfall which we discuss in the next video.

Danger

- What we call the $p\%$ VaR some people call the $(100 - p)\%$ VaR.
- For us the higher the VaR, the riskier the position. Some people change the sign so that the lower the VaR, the riskier the position.
- I've adopted the conventions of Föllmer and Schied "Stochastic Finance: an introduction in discrete time". The same conventions are used by Wikipedia MOSTLY.

Formal definition

Definition

Let X be the random variable denoting the profit of a portfolio after d -days

$$\text{VaR}_\alpha(X) := F_{-X}^{-1}(1 - \alpha)$$

where F_{-X}^{-1} is the quantile function of the loss $-X$.

Where

Definition

The quantile function or inverse distribution function of a random variable Y is

$$F_Y^{-1}(\alpha) = \inf\{y \in \mathbb{R} : F_Y(y) \geq \alpha\}$$

VaR methodologies

- Monte Carlo VaR
- Historic VaR
- Parametric VaR

Monte Carlo VaR

Monte Carlo VaR algorithm

Algorithm

- *Simulate a large number of price paths of the underlyings in the \mathbb{P} measure over an d -day horizon.*
- *For each price path, compute the value of your position on day d .*
- *Hence compute the loss each day d .*
- *Read off the $(100 - p)\%$ percentile of the loss distribution.*
- *This gives you an approximation to the $p\%$, d -day VaR.*

Example

- We will assume that the stock price follows the Black-Scholes model with parameters μ and σ .
- We will compute the risk for a portfolio consisting of a single call option with a known strike K and maturity T .
- We need to be able to compute the price of the portfolio at the start time and after d -days in each scenario. We will use the Black-Scholes formula to compute the price. So we are assuming there is a known volatility σ that applies to our simulation.
- To price more general portfolios, just replace the code marked with the comment "price Portfolio".

MATLAB implementation

```
function [ var ] = monteCarloVar( ...
    percent,days, ...
    strike, maturity, ...
    spot, drift, volatility, ...
    riskFreeRate, ...
    nPaths)

t = days/365;
startPrice = blackScholesCallPrice( ...
    strike, maturity, ...
    spot,riskFreeRate, volatility ); % price portfolio at start
stockPrices = generateBSPaths( ...
    t, spot, drift, volatility,...
    nPaths, 1);
endPrices = blackScholesCallPrice( strike, maturity-t, ...
    stockPrices,riskFreeRate,volatility ); % price portfolio
profits = endPrices - startPrice;
var = prctile( -profits, 100-percent );

end
```


Monte Carlo VaR confidence interval

- Suppose we have computed the $p\%$ d -day VaR using m samples.
- We want to compute a $q\%$ confidence interval.
- Let $\alpha = p/100$, so $0 \leq \alpha < 1$. $\beta = q/100$.
- Let m be the number of samples created.
- Define:

$$\frac{r}{100} = \alpha + \sqrt{\frac{\alpha(1-\alpha)}{m}} N^{-1} \left(\frac{1-\beta}{2} \right)$$

$$\frac{s}{100} = \alpha - \sqrt{\frac{\alpha(1-\alpha)}{m}} N^{-1} \left(\frac{1-\beta}{2} \right)$$

- With $q\%$ confidence, the $p\%$ VaR figure lies between the $r\%$ VaR estimate and the $s\%$ VaR estimate.
- Justification: see FM05.

Monte Carlo VaR Summary

Calculate value at risk by simulating price changes and reading off the desired percentile.

- Good points
 - Highly flexible, you can use any pricing model and compute VaR for any security that you know how to price.
 - Intuitive and easy to program assuming you can already simulate your stochastic model and price your securities.
- Bad points
 - Slow. We may have to run tens of thousands of scenarios and then price the security in each scenario. This may require an additional Monte Carlo simulation if the security is exotic. Thus we can easily need to generate hundreds of millions of scenarios in total.
 - Subjective. Depends heavily upon the choice of model to generate prices, but which model is correct?
 - Only as good as your model. If your model doesn't have fat tails, you will underestimate risk.

Historic VaR

- Monte Carlo VaR relies on having a good model of log-returns on risk factors.
- Unfortunately there is no general agreement on what the right model is.
- So why not use historic log-returns and assume that a similar pattern will be repeated.
- This approach is called historic VaR.
- Note: it uses a different \mathbb{P} model than Monte Carlo VaR. Historic VaR is not an approximation to Monte Carlo VaR, it is a completely different figure.

Definition

The log-return of an asset from time t_1 to time t_2 is

$$r_{t_1, t_2} = \log \left(\frac{P_{t_2}}{P_{t_1}} \right)$$

where P_t is the price of the asset at time t .

Definition

The relative-return of an asset from time t_1 to time t_2 is

$$\frac{P_{t_2} - P_{t_1}}{P_{t_1}}.$$

Warning: the log-return of an asset is NOT the log of the relative-return of an asset.

- Over short time periods: log-return \approx relative-return.
- The advantage of log-returns is:

$$r_{t_1, t_3} = r_{t_1, t_2} + r_{t_2, t_3}$$

- Using log-returns also marries well with geometric Brownian motion which assumes log-returns are normally distributed.

Historical VaR Algorithm

Algorithm

- Suppose that we have m days of historical log-returns for N underlying risk factors. $r_i^{(a)}$ is the log-return of risk factor a on day i .
- For each day i in our sample we generate a simulated d -day log-return by the formula:

$$r_i^{(a)} \sqrt{d}$$

- We use these simulated log-returns to simulate the values of our risk factor on day d
- We then compute the historic VaR by reading off the desired percentile.

Discussion

- We have used the familiar scaling properties of geometric Brownian motion and have assumed these hold for the risk factor under consideration.
- You could look at non overlapping historic d -day log-return values instead of doing this scaling, but that will not be practical for larger d due to the lack of historic data.
- A typical underlying risk-factor is a stock. We have a portfolio of options on the stock, but we only need to simulate the underlying risk factor.

MATLAB code

```
function [ var ] = historicVar( ...
    percent,days, ...
    strike, maturity, ...
    spot, historicPrices, ...
    riskFreeRate, volatility)

dailyLogReturns = diff( log( historicPrices ) );
dDayLogReturns = dailyLogReturns .* sqrt( days );
stockPrices = exp(dDayLogReturns) * spot;

t = days/365;
startPrice = blackScholesCallPrice( ...
    strike, maturity, ...
    spot,riskFreeRate, volatility ); % price portfolio at start
endPrices = blackScholesCallPrice( strike, maturity-t, ...
    stockPrices,riskFreeRate,volatility ); % price portfolio
profits = endPrices - startPrice;
var = prctile( -profits, 100-percent );

end
```


Historic VaR summary

If you can compute Monte Carlo VaR and have access to historic data, computing historic VaR is easy.

- Good points
 - Historic VaR is less subjective than Monte Carlo VaR.
 - Historic VaR may show fat tails behaviour, this is an important component of risk.
 - As simple to calculate as Monte Carlo Var
- Bad points
 - Relies on the assumption that the past will predict the future
 - Only quite crude calculations possible due to lack of historic data.

Parametric VaR

- Parametric VaR is a method for quickly calculating approximate VaR figures.
- It is also called the Delta-Normal method.
- It only works well for positions that depend approximately linearly upon the returns of the underlying. So it works well for stock portfolios but poorly for derivative portfolios.

Parametric VaR

- Let V be the price of a security that depends upon underlying risk factors $S^{(1)}, S^{(2)}, \dots, S^{(n)}$. In our examples these underlying risk factors will be stock prices.
- Define $z^{(a)} = \log S^{(a)}$.
- Suppose that the log-returns of $S^{(a)}$ are normally distributed with mean $\mu^{(a)} = 0$. Suppose the covariance matrix of the log-returns is given by Σ .

- Define

$$\delta^{(a)} = \frac{\partial V}{\partial z^{(a)}} = S^{(a)} \frac{\partial V}{\partial S^{(a)}}$$

- In other words $\delta^{(a)}$ is the sensitivity of the security price to the log-return of $S^{(a)}$.

Parametric VaR

- We have the following Taylor series for V .

$$\begin{aligned}V(t, z^{(1)}, \dots, z^{(n)}) &= V_0 + \frac{\partial V}{\partial t} t + \sum_{i=1}^n \frac{\partial V}{\partial z^{(i)}} r^{(i)} + \dots \\ &= V_0 + \frac{\partial V}{\partial t} t + \sum_{i=1}^n \delta^{(i)} r^{(i)} + \dots\end{aligned}$$

Where r is the log-return.

- Note that the risk factors scale proportional to \sqrt{t} so, the $\frac{\partial V}{\partial t}$ term is negligible compared to the terms involving $r^{(i)}$.
- Given the assumption that the log-returns are normally distributed, V is approximately normally distributed, and we can read off the mean (0) and standard deviation in terms of Σ the covariance matrix of the $r^{(i)}$.

Parametric VaR cont

- Therefore to order \sqrt{d} , the distribution for the change in V over a d -day period is normally distributed with mean 0 and standard deviation:

$$\sqrt{\frac{d\delta^T \Sigma \delta}{365}}$$

Here δ is the vector with components $\delta^{(a)}$.

- Therefore the $p\%$ d -day VaR can be approximated by

$$N^{-1} \left(\frac{100 - p}{100} \right) \sqrt{\frac{d\delta^T \Sigma \delta}{365}}$$

Parametric VaR algorithm

Thus we have a simple first order algorithm for estimating VaR.

Algorithm

- *Compute the sensitivities of V*

$$\delta^{(a)} = S^{(a)} \frac{\partial V}{\partial S^{(a)}}$$

Note that in the case where S is a stock, this is $S\Delta$.

- *Use the formula*

$$\text{VaR}_\alpha \approx N^{-1}(1 - \alpha) \sqrt{\frac{d\delta^T \Sigma \delta}{365}}$$

Parametric VaR Summary

- Good points:
 - It's extremely quick and easy to calculate
- Bad points:
 - For nonlinear securities such as call options which are nearly at the money (i.e. when the current stock price is near the strike price), it is unlikely to be very accurate.
 - It is tied firmly to the multivariate normal distribution so will never account for fat tails.

Analytic formula

- Risk figures are usually computed for portfolios with many risk factors. For simple portfolios with one risk factor it can be possible to compute VaR analytically.
- Let $C_{BS}(S_t, K, t, T, \mu, \sigma, r)$ be the Black-Scholes pricing formula for a call.
- C_{BS} is monotonic in the stock price S_t hence the 5th percentile of the call price on day d occurs when the stock price is at the 5th percentile.
- So the Var between time 0 and time t of a call option with maturity T is

$$\text{VaR}_\alpha = C_{BS}(F_{S_t}^{-1}(\alpha), K, t, T, \mu, \sigma, r) - C_{BS}(S_0, K, 0, T, \mu, \sigma, r)$$

Which VaR methodology should I use?

- Given all these options, which VaR methodology should you use?
- All of them.
 - Use Monte Carlo for a potentially more accurate figure that reflects your beliefs about the market
 - Use Historic VaR to backtest your Monte Carlo prices - do they really seem plausible?
 - Use parametric VaR for rapid intra day recalculation of VaR.
 - (Or innovate: this is an opportunity for your bank to be able to allocate money to the right place, so its a chance to gain a commercial edge)
- The exception is the analytic formula. For complex portfolios analytic VaR calculations are not possible so this mostly a mathematical curiosity.

Common Errors

Error 1

- Historic VaR figures and Monte Carlo VaR will usually be different. This is because they use different probability models.
 - The model for Monte Carlo is a subjective and is chosen by the risk manager.
 - In historic VaR calculations, the probability model is unknown, but we assume the historic data is a random sample from this unknown probability model.

Common Errors

Error 2

- We have not discussed how to formulate the probability model for Monte Carlo VaR calculations. Typically risk managers will perform a statistical analysis of historic data and develop a model based on this analysis. For example, they may fit a GARCH (1,1) model to the historic data. Note that this is NOT the same thing as a historic VaR calculation. Historic VaR is calculated in a very specific way that does not depend upon choosing any specific statistical model.
- Use the techniques from FM05 to fit a time series model to historical data. These methods are not examinable on this course.

How to measure risk

- Is VaR a sensible way to measure risk?
- It is widely used in industry. But does that mean it is sensible?
- In an attempt to answer this, in a famous paper Artzner, Delbaen, Eber and Heath (1998) put together a mathematical theory of coherent risk measures which they believed shows VaR is NOT sensible.

Definition

Let $\rho : L^0(\Omega) \rightarrow \mathbb{R}$ be a function that sends the random variable representing the payoff of a portfolio to some measure of the risk of that portfolio, then ρ is said to be subadditive if

$$\rho(X_1 + X_2) \leq \rho(X_1) + \rho(X_2)$$


Value at risk is not sub-additive. This suggests using value at risk as a risk-measure may discourage diversification.

Expected Shortfall

Informal "business" definition:

Definition

For portfolios with continuous loss distribution, the $p\%$ d -day Expected Shortfall is the expected loss in the worst $p\%$ of cases.



Expected Shortfall - formal definition

Formal mathematical definition (they essentially match for continuous distributions):

Definition

For $\alpha = \frac{p}{100} \in (0, 1)$

$$ES_{\alpha}(X) = \frac{1}{\alpha} \int_0^{\alpha} VaR_{\gamma}(X) d\gamma$$

- Expected Shortfall is also known as:
 - CVaR
 - average value at risk

Some of my online past papers may use the term CVaR but I'm trying to switch my notes to the more common Expected Shortfall.

Computing Expected Shortfall

- Monte Carlo Expected Shortfall. Simulate the underlyings and compute the average loss of the worst $p\%$ of scenarios.
- Historic Expected Shortfall. Use the same algorithm as for Monte Carlo Expected Shortfall but use historic returns to compute possible scenarios for the underlyings as is done in Monte Carlo VaR.
- You can compute parametric (a.k.a. delta normal) expected shortfall by approximating with a normal distribution. However, since the main interest in expected shortfall is in how it handles fat tails this perhaps isn't such a useful idea.

Expected Shortfall properties

Advantages

- It is quite easy to understand and explain to non-specialists
- It is a sub-additive
- It doesn't completely ignore the risk in the tail
- It has some nice convexity properties which make it easy to use in optimizations

Disadvantages

- It is more computationally expensive to calculate than VaR.
- It uses expectations in the tail, so it is effectively “risk neutral” in the tail
- It requires us to model the tail accurately, but that is the part of the model we understand least well.

Basel III requires banks to use Expected Shortfall.

My own views

- The real theoretical problem with Value at Risk is that it ignores the risk in the tail of the distribution. This is what causes its failure to be sub-additive but I think this is the real issue.
- In my view, Expected Shortfall is problematic because it is risk-neutral in the tail. It is the expected loss in the tail, not an expected utility in the tail. I believe it should be risk-averse in the tail.
- The risk-neutrality of Expected Shortfall comes from the positive homogeneity property. I believe one should consider risk-measures based on expected utilities.
- This is a controversial view. Some authors (e.g. Acerbi and Scandolo) argue that risk measures based on expected utility are too hard for business users to understand.

VaR advantages and disadvantages

Advantages

- It is easy to understand and explain to non-specialists

Disadvantages

- It ignores risk in the tail
- It is not sub-additive risk measure so may discourage diversification
- It may not curb prevent excessive risk-seeking behaviour for trader's with S-shaped utilities (see Armstrong, Brigo "Basel risk limits will not curb rogue traders", The Banker)

ES advantages and disadvantages

Advantages

- It is reasonably easy to understand and explain to non-specialists
- It is a sub-additive risk measure

Disadvantages

- It is risk neutral in the tail
- It is more expensive to compute than Monte Carlo VaR
- The value depends upon the details of your model for the tail of the distribution. This is where you have least knowledge.
- It may not curb prevent excessive risk-seeking behaviour for trader's with S-shaped utilities (see Armstrong, Brigo "Basel risk limits will not curb rogue traders", The Banker)

Basel II advantages and disadvantages

Advantages

- Market transparency seems good on the face of it.
- If everyone used a different methodology it would be hard to evaluate banks, so consistency seems a good idea.
- Regulation to prevent banks taking excessive risk seems necessary to prevent the need for government bail outs.

Disadvantages

- The methodology for evaluating risk in Basel II is overly simplistic. In particular the evaluation of credit risk is highly suspect.
- If everyone is using the same methodology to compute risk and operating under similar constraints, this might lead to herd behaviour and hence crashes.
- Any methodology to compute risk may lead to over confidence that risks are being properly accounted for. We just can't

Basel III

Basel III is a response to the financial crisis

- Increase the amount of shareholder capital required for given levels of risk
- Leverage ratio limit on ratio of tier I capital to total assets
- Liquidity coverage ratio
- Use Expected Shortfall rather than VaR
- ...

Note that simple log normal models would not be considered sufficiently conservative by modern standards as they underestimate risk in the tails. Models such as GARCH, student t-GARCH and from extreme value theory are likely to be preferred.

Summary

- Risk management is a subtle topic
- VaR and ES are popular risk-figures in practice because they are required by the Basel accords.
- Other risk-measures exist that many mathematicians like such as the entropic risk measure.

Some references:

- Jorion: "Value at Risk"
- Follmer and Schied: "Stochastic Finance: an introduction in discrete time"
- Artzner, Delbaen, Eber and Heath: "Coherent measures of Risk"
- Armstrong, Brigo: "Rogue traders versus value-at-risk and expected shortfall"