FMO6 — Web: https://tinyurl.com/ycaloqk6 Polls: https://pollev.com/johnarmstron561

Revision Lecture

Dr John Armstrong

King's College London

August 22, 2020

Types of options

We can categorize options in the following ways

- A European option's payoff depends only on the price of the stock at maturity
- An Asian option's payoff depends upon the average price of the stock
- An American option's payoff depends on the price of the stock at the time that the buyer *chooses* to exercise the option.

We can also categorize options by the payoff function:

- Vanilla Put and call options I hope you know!
- A digital call option has value 1 if the stock price is greater than the strike, 0 otherwise.
- A digital put option has value 1 if the stock price is less than the strike, 0 otherwise.

Barriers

We can also add path dependence with barriers. Each barrier has two parameters: is it "up" or "down" and is it "in" or "out".

- Is it a knock in or a knock out barrier? The value is 0 if it hits the barrier for knock out options. The value is 0 if it never hits the barrier for knock in options.
- The direction of the barrier is up or down. A down and out option pays 0 if the price ever goes down below the barrier. And up and out option pays 0 if the price ever goes up above the barrier.

In total there are four different kinds of barriers. Thus as far as the options we price in this course are concerned there are three dimensions to an option:

- Its ethnicity: (European, Asian, American)
- Its payoff function: (Put, Call, Digital Put, Digital Call)
- Its barriers: (Up and Out, Down and Out, Up and In, Down and In)

Question 1 - Numerical Integration

- **Question** (i) State the rectangle rule for integrating a real valued function $f : [a, b] \to R$ defined on a closed interval. [20%]

- **Question** (i) State the rectangle rule for integrating a real valued function $f : [a, b] \to R$ defined on a closed interval. [20%]
- **Answer** L N be an integer and define

$$h = \frac{b-a}{N}$$

$$x_n = a + \left(n - \frac{1}{2}\right) h \qquad n = 1, 2, \ldots N$$

Then the rectangle rule approximation for $\int_a^b f$ is:

$$\sum_{n=1}^{N} f(x_n) * h$$

- **Question** (ii) Write the MATLAB code to integrate e^{-x^2} over the interval $[0, 1]$ using the rectangle rule. [30%]

- **Question** (ii) Write the MATLAB code to integrate e^{-x^2} over the interval $[0, 1]$ using the rectangle rule. [30%]
- **Answer**

```
N = 1000;
a = 0;
b = 1;
h = (b-a)/N;
total = 0.0;
for n=1:N
   xN = a + (n-0.5)*h;
   fxN = exp( -xN^2 );
   total = total+h*fxN;
end
```

- **Question** (iii) Name three other numerical integration techniques that you could use to evaluate this integral and sketch a log-log plot of their convergence as the number of steps increases. Your plot should indicate how rounding errors on a digital computer limit the maximum accuracy. [30%]

- **Question** (iii) Name three other numerical integration techniques that you could use to evaluate this integral and sketch a log-log plot of their convergence as the number of steps increases. Your plot should indicate how rounding errors on a digital computer limit the maximum accuracy. [30%]
- **Answer** The trapezium rule, Simpson's rule and Monte Carlo integration. These converge at the rate $O(n^{-2})$, $O(n^{-4})$ and $O(n^{-\frac{1}{2}})$ respectively.

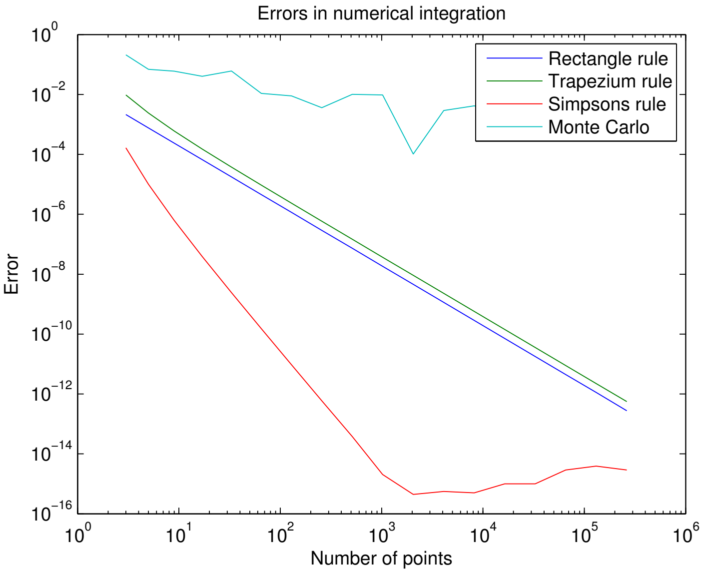
Errors in numerical integration
Error
Number of points
Rectangle rule
Trapezium rule
Simpsons rule
Monte Carlo
10^{0}
10^{-2}
10^{-4}
10^{-6}
10^{-8}
10^{-10}
10^{-12}
10^{-14}
10^{-16}
10^{0}
10^{1}
10^{2}
10^{3}
10^{4}
10^{5}
10^{6}

- **Question** (iv) Explain how pricing options by Monte Carlo simulation can be interpreted in terms of numerical integration. [20%]

- **Question** (iv) Explain how pricing options by Monte Carlo simulation can be interpreted in terms of numerical integration. [20%]
- **Answer** In risk neutral pricing, one computes the price of an option as a (discounted) expectation in the risk neutral measure. By definition, of expectation this can be seen as an integral:

$$\text{Price} = e^{-rT} \int_0^\infty (\text{Payoff given S}) \times \text{q(S)} \,\mathrm{d}S$$

In this equation $q(S)$ is the p.d.f. of the stock price at time T in the risk neutral measure. We can use the cumulative distribution function to make a change of variables $u = \text{cdf}(S)$ where u takes values between 0 and 1.

$$\text{Price} = e^{-rT} \int_0^1 (\text{Payoff given u}) \mathrm{d}u$$

Evaluating this integral by Monte Carlo integration is precisely equivalent to Monte Carlo simulation.

Integration methods

What else could be asked about integration methods?

Integration methods

What else could be asked about integration methods?

- How do you perform a substitution to evaluate an integral?
- How could you integrate a general function? (i.e. passing functions using `@`)
- Why is Monte Carlo usually preferred for high dimensional integrals?
- ...

Question 2 - Delta hedging

- **Question** (i) Write a MATLAB function to simulate stock price paths that follow the Black–Scholes model with given parameters [30%]

- **Question** (i) Write a MATLAB function to simulate stock price paths that follow the Black–Scholes model with given parameters [30%]
- **Answer**

```
function [ S, times ] = generateBSPaths( ...
    T, S0, mu, sigma,nPaths, nSteps  )

dt = T/nSteps;
logS0 = log( S0);
W = randn( nPaths, nSteps );
dlogS = (mu-0.5*sigma^2)*dt + sigma*sqrt(dt)*W;
logS = logS0 + cumsum( dlogS, 2);
S = exp(logS);
times = dt:dt:T;

end
```

- **Question** (ii) Suppose that a trader writes a call option at the Black–Scholes price and then performs discrete time delta hedging up to the maturity of the option. They rebalance their portfolio at time points $\{0, \delta t, 2\delta t, \ldots, T\}$. Any money not invested in the stock is invested in a bank account which grows at the risk free rate r.
  - a Write down the difference equations for the number of assets held at each time point. [40%]
  - b Sketch a histogram of the expected profit and loss of this hedging strategy [10%]

Question:

a Write down the difference equations for the number of assets held at each time point. [40%]

Question:

a. Write down the difference equations for the number of assets held at each time point. [40%]

Answer: At each time point i, let $n_i, b_i, S_i, \Delta i$ denote the quantity of stock held, the bank balance, the stock price and the Black–Scholes delta of the stock at time point i, let b_i denote the bank balance at time point i. Let S_i denote the stock price at time point i, let Δ_i denote the delta of the option at time point i as computed by the Black–Scholes formula. Let P denote the Black–Scholes price at time 0. We have:

$$n_0 = \Delta_0$$

$$b_0 = P - n_0 S_0$$

At subsequent times:

$$n_i = \Delta_i$$

$$b_i = e^{r\delta t} b_{i-1} + (n_{i-1} - n_i) S_i$$

At maturity we can compute the profit and loss as:

$$\text{PnL} = n_N S_N + b_N - \max\{S_N - K, 0\}$$

Where $N = T/\delta t$ and K is the strike of the option.

- **Question:** (b) Sketch a histogram of the expected profit and loss of this hedging strategy [10%]

- **Question:** (b) Sketch a histogram of the expected profit and loss of this hedging strategy [10%]
- **Answer:** Assuming zero transaction costs and that the stock follows the black scholes model, it should have mean 0.

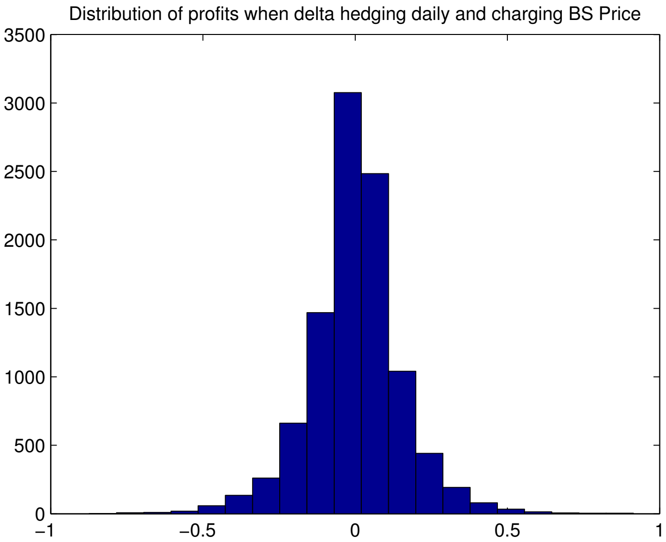
Distribution of profits when delta hedging daily and charging BS Price
3500
3000
2500
2000
1500
1000
500
0
–1
–0.5
0
0.5
1

- **Question:** (iv) Explain briefly what is meant by gamma hedging and explain why a trader might choose to gamma hedge an exotic option. [20%]

- **Question:** (iv) Explain briefly what is meant by gamma hedging and explain why a trader might choose to gamma hedge an exotic option. [20%]
- **Answer:** Due to transaction costs, you should not rehedge too often. This means that in practice purchasing options at market prices is a more cost effective way to hedge the risk of an exotic option than to simply trade in the underlying. To gamma hedge an exotic option one would purchase the stock and another liquidly traded option to maintain a portfolio that is approximately delta and gamma neutral. Such a portfolio would not need to be rebalanced as often as an option that was merely delta neutral - this is because the portfolio would be hedged against both first and second order changes in the underlying. As a result the transaction costs of the strategy are likely to be less. In addition, since the portfolio is gamma neutral, the overall risk figures such as VaR for the portfolio would be lower and so the trader's risk manager may as a result allow the trader to take a larger position.

Hedging

What else could be asked about hedging?

- Rate of convergence
- Pseudo code for implementation
- Formulae for gamma hedging
- ...

Question 3 - Finite Difference Methods

- **Question:** (i) Draw a table summarizing the numerical methods for risk neutral option pricing that were taught in this course and indicated which of these methods can be used to price the following types of option:
  - (a A European call option
  - (b An American put option
  - (c An Asian call option
  - (d An up-and-out option

  [30%]

Answer:

	Finite Difference	1-d numerical Integration	Monte Carlo Simulation
European Call	Yes	Yes	Yes
American Put	Yes	No	No
Asian Call	No	No	Yes
Up-and-Out Option	Yes	No	Yes

- **Question:** (ii) When pricing a European put option by the finite difference method, what boundary conditions would you use? [20%]

- **Question:** (ii) When pricing a European put option by the finite difference method, what boundary conditions would you use? [20%]
- **Answer:** When the stock price is much higher than the strike, a European put is worth approximately 0, so I would use the condition $V = 0$ along the top boundary. When the stock price is near 0, the put is worth approximately $e^{-r(T-t)}K$ (i.e. the discounted final strike), so I would use the condition $V = e^{-r(T-t)}K$ along the bottom boundary.

- **Question:** Recall that the Black–Scholes PDE is

$$V_t + \frac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0$$

where subscripts denote partial differentiation. Use this to derive the difference equations that must be solved to price a put option by the explicit finite difference method. [30%]

- **Question:** Recall that the Black–Scholes PDE is

$$V_t + \frac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0$$

  where subscripts denote partial differentiation. Use this to derive the difference equations that must be solved to price a put option by the explicit finite difference method. [30%]
- **Answer:** The stencil for the explicit finite difference method is:

For the explicit method, we use the following estimate for V_t

$$V_t \approx \frac{V_{i,j} - V_{i-1,j}}{\delta t}$$

We use the following estimate for V_S

$$V_S \approx \frac{V_{i,j+1} - V_{i,j-1}}{2\delta S}$$

And for V_{SS}

$$V_{SS} \approx \frac{V_{i,j+1} - 2V_{i,j} + V_{i,j-1}}{\delta S^2}$$

Hence:

$$\frac{V_{i,j} - V_{i-1,j}}{\delta t} + \frac{1}{2}\sigma^2 S_j^2 \frac{V_{i,j+1} - 2V_{i,j} + V_{i,j-1}}{\delta S^2} + rS_j \frac{V_{i,j+1} - V_{i,j-1}}{2\delta S} - rV_{i,j} = 0$$

$$V_{i-1,j} = V_{i,j} + \delta t \left(\frac{1}{2}\sigma^2 S_j^2 \frac{V_{i,j+1} - 2V_{i,j} + V_{i,j-1}}{\delta S^2} + rS_j \frac{V_{i,j+1} - V_{i,j-1}}{2\delta S} - rV_{i,j} \right)$$

In addition we have boundary conditions:

$$V_{i,j_{\mathrm{min}}} = e^{-r(T-t_i)}K$$

$$V_{i,j_{\mathrm{max}}} = 0$$

And initial conditions.

$$V_{i_{\mathrm{max}},j} = \max\{K - S_j, 0\}$$

In these formula $V_{i,j}$ is our estimate for the option price at point (i,j) in our discretization. S_j is the stock price corresponding to the value j and t_i is the time corresponding to i.
The top boundary condition is an approximation, we need the maximum value of S in our grid to be chosen so that the option is unlikely to be in the money. A value of $S_{\mathrm{max}} = e^{-(r-\frac{1}{2}\sigma^2)T+4\sigma\sqrt{T}}K$ would be a reasonable choice.

- **Question:** (iv) How do the difference equations change when pricing an American put option? [20%]

- **Question:** (iv) How do the difference equations change when pricing an American put option? [20%]
- **Answer:** The bottom boundary condition becomes $V_{i,0} = K$ but the top boundary condition does not change. Define $\overline{V}_{i-1,j}$ to be the term on the right hand side of the difference equation for a European put, then the corresponding difference equation for an American put is:

$$V_{i-1,j} = \max\{\overline{V}_{i-1,j}, \max\{K - S_j, 0\}\}$$

Finite differences

What else could be asked about finite differences?

Finite differences

What else could be asked about finite differences?

- Implicit scheme
- Transformation to heat equation
- Boundary conditions for calls
- Barrier options
- Rate of convergence
- ...
- Not: Crank-Nicolson scheme (retake only, I think this is too fiddly for exam conditions myself)
- Not: American options by implicit method

Break

- You might want to look at the links on Keats

Question 4 - Interesting stochastic processes

- **Question:** (i) What is meant by a pseudo square root of a positive definite symmetric matrix A? [10%]

- **Question:** (i) What is meant by a pseudo square root of a positive definite symmetric matrix A? [10%]
- **Answer:** It is a matrix U such that $UU^T = A$.

- **Question:** (i) What is meant by a pseudo square root of a positive definite symmetric matrix A? [10%]
- **Answer:** It is a matrix U such that $UU^T = A$.
- **Question:** (ii) What is meant by the Cholesky decomposition of a positive definite symmetric matrix A? [10%].

- **Question:** (i) What is meant by a pseudo square root of a positive definite symmetric matrix A? [10%]
- **Answer:** It is a matrix U such that $UU^T = A$.
- **Question:** (ii) What is meant by the Cholesky decomposition of a positive definite symmetric matrix A? [10%].
- **Answer:** It is the unique lower triangular pseudo square root of A which is lower triangular and has positive entries on the diagonal.

- **Question:** (iii) Compute the Cholesky decomposition of the matrix:

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

[30%]

- **Question:** (iii) Compute the Cholesky decomposition of the matrix:

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

[30%]

- **Answer:** Let U be the Cholesky decomposition. Write

$$U = \begin{pmatrix} \alpha & 0 \\ \beta & \gamma \end{pmatrix}$$

We require $UU^T = A$ i.e.

$$\begin{pmatrix} \alpha & 0 \\ \beta & \gamma \end{pmatrix} \begin{pmatrix} \alpha & \beta \\ 0 & \gamma \end{pmatrix} = \begin{pmatrix} \alpha^2 & \alpha\beta \\ \alpha\beta & \beta^2 + \gamma^2 \end{pmatrix} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

So $\alpha = 1$ (since we also require α is positive) and $\beta = \rho$. So $\gamma = \sqrt{1 - \rho^2}$.

- **Question:** (iv) Explain how you would generate a sample of random variables X_1 and X_2 from a two dimensional multivariate normal distribution with mean 0 and covariance matrix:

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

[20%]

- **Question:** (iv) Explain how you would generate a sample of random variables X_1 and X_2 from a two dimensional multivariate normal distribution with mean 0 and covariance matrix:

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

[20%]

- **Answer:** First generate a $2 \times N$ matrix M of independent normally distributed random variables with mean 0 and standard deviation 1. We can consider each column of UM to be a sample from the desired distribution, giving N samples in total.

- **Question:** Describe how you can simulate stock prices in discrete time that approximately follow the Heston model:

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sqrt{\nu_t} S_t \mathrm{d}W_t^S$$

$$\mathrm{d}\nu_t = \kappa(\theta - \nu_t)\mathrm{d}t + \xi\sqrt{\nu_t}\mathrm{d}W_t^\nu$$

where W_t^S and W_t^ν are Wiener processes with correlation ρ, S_t is the stock price at time t, ν_t is the volatility process and all other terms are constants. [30%]

- **Question:** Describe how you can simulate stock prices in discrete time that approximately follow the Heston model:
$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sqrt{\nu_t} S_t \mathrm{d}W_t^S$$
$$\mathrm{d}\nu_t = \kappa(\theta - \nu_t)\mathrm{d}t + \xi\sqrt{\nu_t}\mathrm{d}W_t^\nu$$
where W_t^S and W_t^ν are Wiener processes with correlation ρ, S_t is the stock price at time t, ν_t is the volatility process and all other terms are constants. [30%]
- **Answer:** Let δt be a chosen time step. Using the answer to part (iv) one can generate random variables X_t^1 and X_t^2 such that $\sqrt{\delta t} X_t^i$represent the increments of the Wiener proceses. Define a discrete set of approximations to the stock price process by the difference equations:
$$S_t = S_{t-1} + \mu S_{t-1}\delta t + \sqrt{\nu_{t-1}} S_{t-1}\sqrt{\delta t} X_t^1$$
$$\nu_{t+1} = \nu_t + \kappa(\theta - \nu_t)\mathrm{d}t + \xi\sqrt{\nu_t \delta t} X_t^2$$
So given initial conditions S_0 and ν_0 one can use these equations to simulate stock prices. This is the Euler scheme for numerically approximating the stochastic differential equations.

More interesting processes

What else could be asked?

- Show that the generated random variables have the desired covariance matrix?
- What would you expect the error of the simulation to be?
- ...

Question 5 - Risk Measures

- **Question:** What is meant by VaR and CVaR

- **Question:** What is meant by VaR and CVaR
- **Answer:** The VaR of an portfolio (or a single security) at a given percentage $p\%$ and time horizon t is the maximum loss that of the portfolio in the $(100 - p)\%$ best case scenarios over that time horizon. CVaR is the expected loss of the $p\%$ worst case scenarios over that time horizon.

- **Question:** (ii) What is the sub-additivity property of a coherent risk measure? Show that VaR is not sub-additive.

- **Question:** (ii) What is the sub-additivity property of a coherent risk measure? Show that VaR is not sub-additive.
- **Answer:** A risk measure ρ that associates risk measurement to possible portfolios is said to be sub-additive if $\rho(A+B) \leq \rho(A) + \rho(B)$ where A and B are two portfolios and $A+B$ is the porfolio obtained by combining A and B. Let A be a portfolio consisting of -1 digital put option on a stock that we expect to pay out 1 4% of the time and 0 96% of the time. Let B be a portfolio consisting of -1 digital call option on the same stock that we expect to pay out 1 4% of the time and 0 the rest of the time. Let a and b the current value of the two portfolios. The 5% VaR of A is $-a$. The 5% VaR of B is $-b$. The 5% VaR of $A+B$ is $-a-b+1$. Thus VaR is not sub-additive.

- **Question:** (iii) Write a difference equation you could use to simulate a stock price that follows the Black–Scholes model. [20%]

- **Question:** (iii) Write a difference equation you could use to simulate a stock price that follows the Black–Scholes model. [20%]
- **Answer:** Use the difference equation

$$s_t = s_{t-1} + (\mu - \frac{1}{2}\sigma^2)\delta t + \sigma\sqrt{\delta t}\epsilon_t$$

to simulate the log of the stock price and then compute the stock price $S_t = exp(s_t)$ in discrete time with time interval δt. Here ϵ_t is a sequence of independent normally distributed random variables with mean 0 and standard deviation 1.

- **Question:** (iv) Describe how you could use the results of such a simulation to estimate the VaR of a call option on a stock.

- **Question:** (iv) Describe how you could use the results of such a simulation to estimate the VaR of a call option on a stock.
- **Answer:** Suppose we wish to estimate the $p\%$ T VaR. Use the difference equation to simulate a large number of scenarios for stock prices at time T. One need only simulate using a single time step $\delta t = T$. One can then use the Black Scholes formula to price the call option at time T in each scenario. By computing the $(100 - p)\%$ percentile of the loss distribution one can estimate the VaR.

- **Question:** (v) Explain briefly how you could go about testing the results of this calculation.

- **Question:** (v) Explain briefly how you could go about testing the results of this calculation.
- **Answer:** It is actually quite easy to compute an analytic formula for the VaR of a call option. This is because the Black Scholes formula is an increasing function of stock price. First compute the p%-percentile of the possible stock price which is easy to do since we know stock prices are log normally distributed. Then plug this number into the Black Scholes formula to find the p% percentile of the option price at time T. One can then compare this analytic formula with the Monte Carlo price.
  Any sensible test you had come up with would have been credit. This is just *my* answer.

VaR and CVaR

What else might be asked?

VaR and CVaR

What else might be asked?

- What are the pros and cons of VaR and CVaR?
- What is parameteric VaR?
- What is historic VaR?
- What are the pros and cons of different kinds of VaR?
- What is the exponentially weighted moving average?
- ...

Bonus Question 1 - MATLAB programming

- **Question:** (i) (a) State the Monte Carlo integration rule for a function $f : [a, b] \to \mathbb{R}$ defined on a closed interval. [20%]

- **Question:** (i) (a) State the Monte Carlo integration rule for a function $f : [a, b] \to \mathbb{R}$ defined on a closed interval. [20%]
- **Answer:** Choose a sample size N. Pick N points x_n from the uniform distribution over $[a, b]$. Then the Monte Carlo estimate for the integral is:

$$\frac{b-a}{N} \sum_{i=1}^{N} f(x_i)$$

- **Question:** (i) (b) Write the MATLAB code to integrate e^{-x^2} over the interval $[0, 1]$ using Monte Carlo integration. [30%]

- **Question:** (i) (b) Write the MATLAB code to integrate e^{-x^2} over the interval $[0, 1]$ using Monte Carlo integration. [30%]
- **Answer:**

```
N=10000;
x = rand(1,N);
integral = 1/N * sum( exp( -x.^2 ));
```

- **Question:** (ii) The Box–Muller algorithm is an algorithm to generate independent normally distributed random numbers with mean 0 and standard deviation 1. One first generates uniformly distributed random numbers U_1 and U_2 between 0 and 1. One then defines $Z_1 = R\cos(\theta)$, $Z_2 = R\sin(\theta)$ where $R^2 = -2\log U_1$ and $\theta = 2\pi U_2$. (a) Write a function `boxMuller` which takes a parameter n and returns a $2 \times n$ sample if independent normally distributed random numbers generated by the Box–Muller algorithm. [20%]

- **Question:** (ii) The Box–Muller algorithm is an algorithm to generate independent normally distributed random numbers with mean 0 and standard deviation 1. One first generates uniformly distributed random numbers U_1 and U_2 between 0 and 1. One then defines $Z_1 = R\cos(\theta)$, $Z_2 = R\sin(\theta)$ where $R^2 = -2\log U_1$ and $\theta = 2\pi U_2$. (a) Write a function `boxMuller` which takes a parameter n and returns a $2 \times n$ sample if independent normally distributed random numbers generated by the Box–Muller algorithm. [20%]
- **Answer:**

```
function ret=boxMuller( n )
U1 = rand(1,n);
U2 = rand(1,n);
R = sqrt( -2 * log( U1 ));
theta = 2 * pi * U2;
ret = [ R.*cos( theta ); R.*sin(theta) ];
end
```

- **Question:** (b) How would you test this function? [10%]

- **Question:** (b) How would you test this function? [10%]
- **Answer:** I would write a unit test that confirms that for a large sample the mean, standard deviation and correlation of the generated random variables are approximately 0,1 and 0 respectively. To ensure that the test always passes I would seed the random number generator at the start of the test.

- **Question:** (c) Using the MATLAB function `chol` or otherwise, show how you would generate a sample from a two dimensional multivariate normal distribution with mean 0 and covariance matrix
$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$
[20%]

- **Question:** (c) Using the MATLAB function `chol` or otherwise, show how you would generate a sample from a two dimensional multivariate normal distribution with mean 0 and covariance matrix

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

[20%]

- **Answer:**

```
N = 1000;
rho = 0.1;
x = boxMuller( N );
c = chol( [ 1 rho; rho 1 ], 'lower' );
y = c*x;
```

Each column of matrix y gives a sample from the desired distribution

MATLAB programming?

What else could be asked

- Implement practically anything in MATLAB — only a few line of MATLAB are likely to be required if it is something you've never seen before.
- If you need to use a MATLAB function and can't remember its name, make up something sensible and appropriate. You could make a remark on what the function you've invented does.
- Don't "cheat". I won't give credit if you just use a MATLAB function that answers the problem in full. For example it is clear in the above question that using `randn` is unacceptable in the first part of the question.

Bonus Question 2 - Optimization

- **Question:** You believe that the 5 stocks will have annual returns that follow a multivariate normal distribution with mean vector μ and covariance matrix Σ. You have \$1000000 to invest in these stocks and wish to achieve an expected return of 10% over the year. You wish to select a static portfolio, i.e. you must buy and hold. Express the problem of selecting the portfolio that meets these requirements with the minimum standard deviation as a quadratic programming problem [30%]

- **Question:** You believe that the 5 stocks will have annual returns that follow a multivariate normal distribution with mean vector μ and covariance matrix Σ. You have \$1000000 to invest in these stocks and wish to achieve an expected return of 10% over the year. You wish to select a static portfolio, i.e. you must buy and hold. Express the problem of selecting the portfolio that meets these requirements with the minimum standard deviation as a quadratic programming problem [30%]

- **Answer:** Let w denote the 5-vector of weights of each stock held in your portfolio. The expected return is given by $\mu.w$. The variance of the return is $w^T \Sigma w$. The condition that w are weight vectors is $\sum w = 1$. Thus the problem is:

$$\text{Minimize}\, w^T \Sigma w$$

$$\text{Subject to}\, \mu^T w = 0.1$$

$$\text{and}\, i^T w = 1$$

Where i is a vector of ones.

- **Question:** Explain what is meant by the efficient frontier and sketch its expected shape. Indicate in the same diagram how portfolios consisting of investments in a single stock would perform. [20%]
- **Question:** If one plots the expected return of a portfolio against the standard deviation of the portfolio, the efficient frontier is the curve given by the minimum standard deviation portfolio for a given expected return. For unconstrained Markowitz optimization it will be a hyperbola. The points representing individual stocks will all be contained within the area on the right of the hyperbola as shown (there should be five points - I've been lazy)

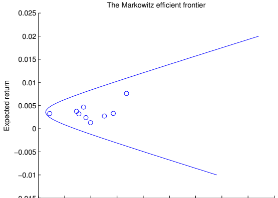

- **Question:** Suppose that we do not believe that the stocks have normally distributed returns, but that the 5 stocks follow some specific stochastic process. Explain how you could use Monte Carlo simulation to find the optimal static portfolio in terms of a utility function u. [30%]
- **Answer:** We wish to find the quantities q_i that maximize

$$E(u(q_i P(1)_i))$$

where $P(1)_i$ is the payoff of stock i subject to:

$$\sum q_i P(0)_i = 1000000$$

where $P(0)_i$ is the initial price of stock i. Use Monte Carlo to simulate a large number of possible scenarios in the P measure and write P_i^α for the payoff of stock i in scenario α. We approximate our problem as minimizing

$$-\sum_\alpha u(q_i P_i^\alpha)$$

subject to the same constraint.

Continued from previous slide

One should introduce a new variable $q'_i = \frac{P(0)_i q_i}{1000000}$ so that the problem is well-scaled. We now have minimize:

$$-\sum_{\alpha} u\left(1000000 q'_i \frac{P_i^{\alpha}}{P_i^{0}}\right)$$

$$\sum q'_i = 1$$

This is a constrained convex optimization problem that can be solved using `fmincon`.

- **Question:** You decide instead to pursue a dynamic investment strategy. Investment strategy S1 is to, once a week, invest all your money in the stock that had the most return in the previous week. Investment strategy S2 is to, once a week invest all your money in the stock that had the least return in the previous week. Assuming the stocks follow a known stochastic process and you have a known utility function u, how could you devise a trading strategy that is guaranteed to be at least as good as strategies S1 and S2? [20%]
- **Answer:** The same algorithm can be used as for the previous part of the question. Calculate the payoffs and costs of following strategies S1 and S2 and then find the optimal convex combination of these strategies. It will be at least as good as the strategies S1 and S2 taken individually.

Bonus Question 3 - More Optimization

A trader has P units of cash and wishes to invest in a stock and a risk free bond to maximize their expected utility at time T. Their utility function is:

$$u(x) = \begin{cases} \ln(x) & \text{if } x > 0 \\ -\infty & \text{otherwise} \end{cases}$$

The trader believes the stock follows geometric Brownian motion:

$$\mathrm{d}S_t = S_t(\mu \, \mathrm{d}t + \sigma \, \mathrm{d}W_t)$$

The bond has interest rate r. At time 0 the trader invests an amount Q of their wealth in stock and the rest in bonds.

- **Question:** Write the expected utility as an integral [40%]
- **Answer:** By Ito's lemma, the log of the stock price s_t follows:

$$\mathrm{d}s_t = (\mu - \frac{1}{2}\sigma^2)\mathrm{d}t + \sigma\,\mathrm{d}W_t)$$

So s_T is normally distributed with mean $s_0 + (\mu - \frac{1}{2}sigma^2)T$ and standard deviation $\sigma\sqrt{T}$. If the trader has initial wealth P then there portfolio consists of Q quantities of the stock and $P - QS_0$ of the bond. Thus the payoff is

$$Q\exp(s_T) + (P - QS_0)e^{rT}$$

Hence the expected utility is:

$$\int_{-\infty}^{\infty} \ln(Q\exp(s_T) + (P - QS_0)e^{rT})p(s_T)\,\mathrm{d}s_T$$

where $p(s_T)$ is the p.d.f. of the normal distribution with mean and standard deviation as above.

Continued from the previous slide

- Introducing $x_t = \text{normcdf}(s_T, s_0 + (\mu - \frac{1}{2}\sigma^2)T, \sigma\sqrt{T})$ we can write this as:
$$\int_0^1 \ln(Q\exp(\text{norminv}(x, s_0 + (\mu - \frac{1}{2}\sigma^2)T, \sigma\sqrt{T}) + (P - QS_0)e^{rT})\, dx_T$$
where `normcdf` and `norminv` are the cumulative normal distribution function and its inverse for specified mean and variance.

Note that I've performed a small trick here. I find it hard to remember the formula for the p.d.f. of the lognormal distribution, so I used Ito's Lemma to transform the equation to one only involving the normal distribution. Even so I end up with an integral with infinite limits which might be fiddly to integrate by Monte Carlo, so I transform it to an integral with finite limits which is now trivial to integrate by Monte Carlo. The code will be just the standard MATLAB for simulating stock prices and then computing the expected utility.

- **Question:** Write the MATLAB code to compute this integral by a Monte Carlo method [30%]
- **Answer:**

```
function e = computeExpectedUtility(...
    P, Q, r, S0, mu, sigma, T, N )
x = rand(N,1);
s = log(S0) + ...
    (mu-0.5*sigma^2)*T + sigma*sqrt(T)*norminv(x);
u = log(Q*exp(s) + (P-Q*S0)*exp(r*T));
e = mean(u);
end
```

- **Question:**State a variance reduction technique you could use to improve the rate of convergence of the Monte Carlo method [10%]
- **Answer:**You could use antithetic sampling. (If you had time to spare you might explain briefly what this means, but the question doesn't seem to actually be asking for you to do this)

- **Question:** $u(x)$ takes the value $-\infty$ when x is negative. What trading constraint does this imply? [10%]
- **Answer:** This means that any possibility of bankruptcy is not acceptable. For the specific problem this means that one can't short the stock.

- **Question:** How could you use MATLAB to find the optimal value of Q?
- **Question:** You could use *fminunc* to find the optimal value of Q. Note that you would need to ensure that the same random numbers were used in each simulation. One could do this by seeding the random number generator. Better yet, use the rectangle rule or some other deterministic integration method to implement `computeExpectedUtility`!

Good Luck!