CHARTA LEARNING · LESSON NOTES

The Black-Scholes-Merton Model

FRM Part I

1 Why an option needs a model

- An option is a right to buy or sell at a fixed price
- Its value depends on how far the stock might move
- So we first need a model of the stock price
- Then we can price the option

IN THE LESSON

- Start with a simple question.
- A friend offers you the right, but not the duty, to buy a share one year from now at one hundred dollars.
- What is that right worth today?
- If the share ends the year at one hundred and thirty, you gain thirty.
- If it ends below one hundred, you simply walk away and lose nothing more.
- Your payoff is lopsided, so the answer depends on how far and how wildly the price might move.
- To value the right we need a picture of how a stock price behaves.
- Black, Scholes and Merton gave the market that picture, and then a price.

2 A stock price as a random walk

- Price moves by random jolts around a drift
- Drift = expected return, mu
- Size of jolts = volatility, sigma
- Percentage changes are random (GBM)

IN THE LESSON

- Picture a stock price as a drunk walker on a street.
- Each moment he lurches a little up or down, and nobody can predict the next step.
- But the walker also drifts in one direction on average, that is the expected return, which we call mu.
- And the size of the lurches is the volatility, which we call sigma.
- The Black-Scholes-Merton model says the percentage change in the price follows exactly this pattern.
- The technical name is geometric Brownian motion.
- Geometric means that it is the percentage change that behaves randomly, not the dollar change, so a cheap stock and an expensive stock wobble in proportion.

3 Log returns are normal

$\ln(S_T \div S_0) \sim \text{Normal}$
$\text{mean} = (\mu - \sigma^2 \div 2) \times T$
$\text{variance} = \sigma^2 \times T$

Log return	Log of the price ratio
Expected return, mu	Expected ordinary return per year
Volatility, sigma	Annual volatility
Time in years, T	Years, so variance grows with time

IN THE LESSON

- Here is the key result.
- Take the continuously compounded return over T years, which is the natural log of the final price divided by today's price.
- Under the model it follows a normal distribution, the familiar bell curve.
- Its average is mu minus half sigma squared, all times T.
- Its variance is sigma squared times T.
- Notice the half sigma squared, subtracted.
- The expected return mu is the average ordinary return, but the average compounded return is lower, because ups and downs of equal size do not cancel under compounding.
- Two common slips are to treat mu as the average log return, and to add the half variance instead of subtracting it.

4 Example 1: one year on a hundred dollars

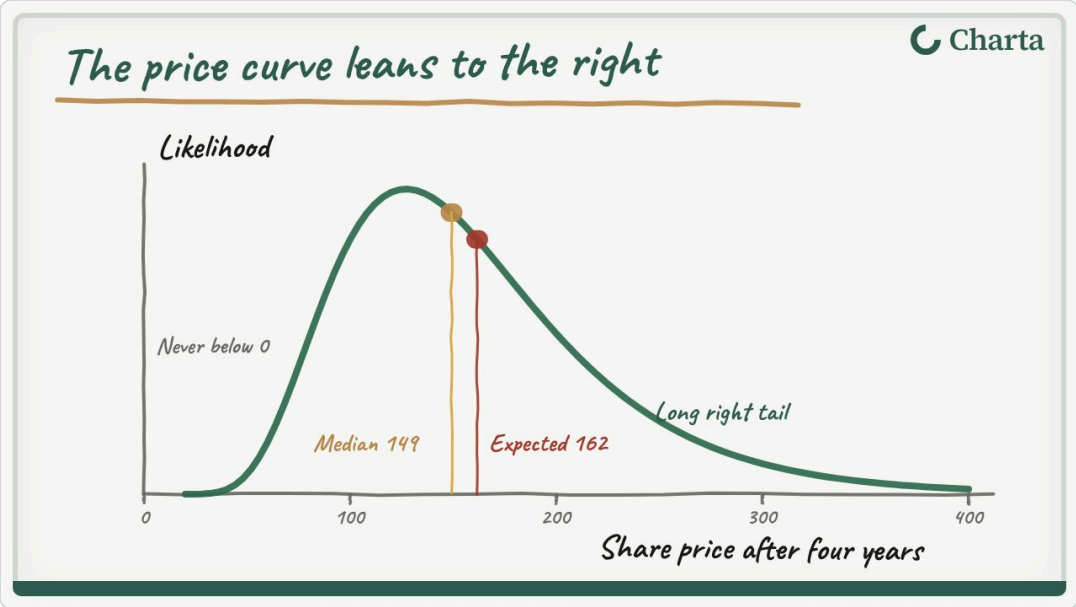
Illustration: price 100, mu 12%, sigma 20%, one year

Mean log return	12% − 20%² ÷ 2	**10%**
Std dev of log return	σ × √1	**20%**
Median price	100 × e^0.10	**110.52**
Expected price	100 × e^0.12	**112.75**

IN THE LESSON

- Let's try it.
- A share costs one hundred dollars.
- The expected return is twelve per cent a year and the volatility is twenty per cent.
- First, the mean log return over one year is twelve per cent minus half of twenty per cent squared.
- Twenty per cent squared is point zero four, half is point zero two, so the mean is ten per cent.
- The standard deviation is twenty per cent.
- Now convert to prices.
- Raise e to the power ten per cent and the typical, or median, price is one hundred and ten dollars fifty two.
- But the expected price uses mu itself: one hundred times e to the point one two, which is one hundred and twelve seventy five.

5 The price curve leans to the right



IN THE LESSON

- Now look at the price itself.
- Because the price equals today's price times e raised to a normal number, the price is called lognormal.
- Here is its shape after four years, using the same example.
- Along the bottom is the share price.
- The curve is not a symmetric bell.
- It leans to the right, with a long tail of large gains, while it can never fall below zero, since e to any power is positive.
- The middle price, the median, is about one hundred and forty nine dollars.
- The expected price is higher, about one hundred and sixty two, because the rare big winners pull the average up.

6 Measuring volatility from history

$u_i = \ln(S_i \div S_{i-1})$
$s^2 = \Sigma(u_i - \bar{u})^2 \div (n - 1)$
$\sigma = s \times \sqrt{252}$

Log return	One period's log return
Number of returns, n	Divide by n minus one, not n
Sample std deviation	Volatility per day
Annualising factor	Root of trading days, not the days

IN THE LESSON

- Volatility cannot be read off a screen, so we estimate it from past prices.
- Record the price at regular intervals, for example daily closes.
- For each interval compute the log return, the natural log of today's price over yesterday's.
- With n such returns, find their average, then the sample variance: add up the squared gaps from that average and divide by n minus one.
- The square root is the standard deviation s, volatility per interval.
- To annualise, multiply s by the square root of the number of intervals in a year.
- With daily data and two hundred and fifty two trading days, multiply by the square root of two hundred and fifty two.
- Do not multiply by two hundred and fifty two itself.

7 Example 2: volatility from six closes

Closes: 100, 101, 99.5, 100.5, 100, 100.9 (five returns)

Mean daily log return	sum of five ÷ 5	**0.179%**
Sample variance	squared gaps ÷ (5 − 1)	**0.000128**
Daily volatility	√0.000128	**1.13%**
Annual volatility	1.13% × √252	**17.96%**

IN THE LESSON

- A tiny example, just to see the steps.
- Six daily closes: one hundred, one hundred and one, ninety nine and a half, one hundred and a half, one hundred, and one hundred point nine.
- That gives five log returns.
- Their mean is about point one eight per cent a day.
- Add up the squared gaps from the mean and divide by four, not five, and the variance is point zero zero zero one two eight.
- The square root is one point one three per cent a day.
- Multiply by the square root of two hundred and fifty two, about fifteen point nine, and the annual volatility is about eighteen per cent.
- Five days is far too short in real life, but the arithmetic is the same.

8 Example 3: realised return and the error

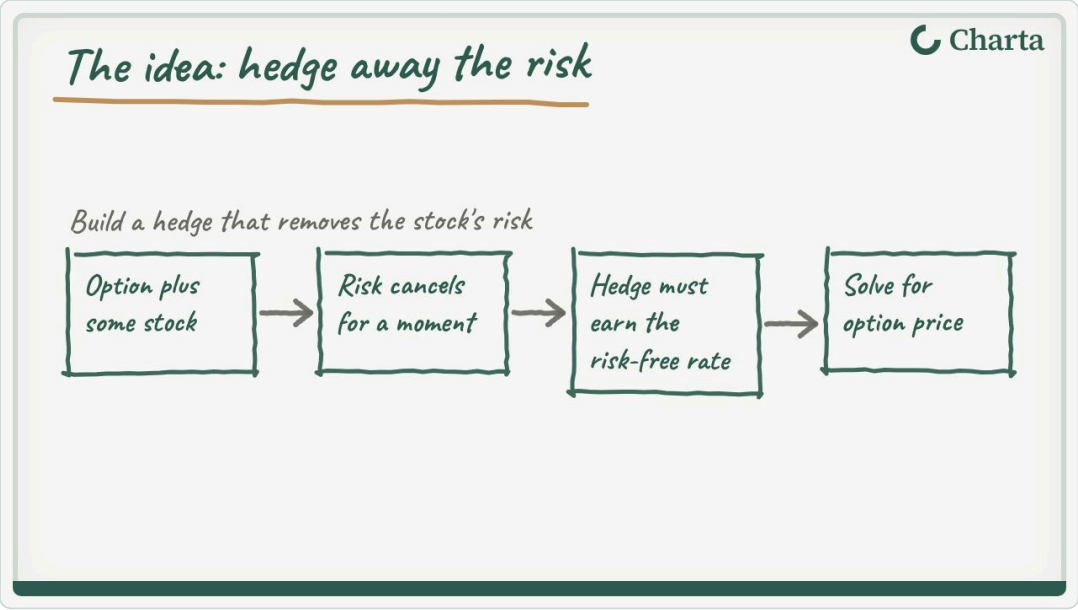
Same data. Then 250 returns at 20% volatility

Mean log return, annual	0.179% × 252	**45.16%**
Half the annual variance	0.5 × 0.0323	**1.61%**
Estimate of μ	45.16% + 1.61%	**46.77%**
Standard error of σ	20% ÷ √(2 × 250)	**0.89%**

IN THE LESSON

- The same data also give the realised return.
- The mean daily log return, point one seven nine per cent, times two hundred and fifty two days, is about forty five per cent a year.
- That estimates mu minus half sigma squared, not mu itself.
- To estimate mu we add half the annualised variance back.
- The annual variance is about point zero three two, so half of it adds one point six per cent, giving roughly forty seven per cent.
- Finally, how reliable is a volatility estimate?
- Its standard error is roughly sigma over the square root of two n.
- With twenty per cent volatility and two hundred and fifty returns, that is about point nine per cent.

9 The idea: hedge away the risk



IN THE LESSON

- Now the heart of the matter, how Black, Scholes and Merton priced an option.
- Their insight was that you can build a portfolio of the stock and a risk free bond that behaves just like the option, at least for a very short moment.
- Look at the picture.
- Buy or sell some stock against the option, so that small price jolts in one cancel the jolts in the other.
- The result has no risk, so it must earn the risk free rate, otherwise there would be free money.
- Solve for the option's price and a remarkable thing happens: the stock's expected return mu disappears from the answer.

10 What the model assumes

- Stock: geometric Brownian motion, constant volatility
- Short selling allowed, no costs or taxes, divisible
- No dividends, no arbitrage, continuous trading
- Constant risk-free rate for all maturities
- So the price ignores mu and risk appetite

IN THE LESSON

- That trick works only under certain conditions, so let's list them.
- The stock follows geometric Brownian motion with constant volatility, so log returns are normal.
- Short selling is allowed, with full use of the cash raised.
- There are no transaction costs or taxes, and securities can be divided into any size.
- The stock pays no dividends during the option's life.
- There is no riskless arbitrage, meaning no free lunch.
- Trading is continuous, so the hedge can be adjusted at every instant.
- And the risk free rate is constant, the same for every maturity.
- Because the hedge removes the stock's risk, neither mu nor investors' appetite for risk affects the price, which is why we may value the option as if investors were risk neutral.

11 Where reality breaks the assumptions

Each assumption has a real-world weak spot

Assumption	Reality
Normal, constant σ	Fat tails, jumps
No dividends	Many pay them
Constant rate	Rates move
Continuous trading	Gaps, closed markets
No costs	Spreads, fees

IN THE LESSON

- The exam likes to ask which assumption fails in practice.
- Here is the usual matching.
- Real returns have fat tails and sudden jumps, so a normal curve understates extreme moves, and volatility is not constant, which shows up as the volatility smile we will meet shortly.
- Most stocks do pay dividends, which we handle with an extension later.
- Interest rates change and differ by maturity.
- Trading is not continuous, since markets close and prices gap overnight.
- And real trading has bid-ask spreads and fees, so a perfect hedge costs money.
- The model is still the standard language of option markets, but traders know where it is wrong.

12 The Black-Scholes-Merton formula

$c = S_0N(d_1) - K \cdot e^{(-rT)} \cdot N(d_2)$
$p = K \cdot e^{(-rT)} \cdot N(-d_2) - S_0N(-d_1)$
$d_1 = [\ln(S_0 \div K) + (r + \sigma^2 \div 2)T] \div (\sigma\sqrt{T})$; $d_2 = d_1 - \sigma\sqrt{T}$

Stock price, S0	Price today
Strike, K	Exercise price, paid at expiry
Risk-free rate, r	Continuous rate
Normal probability, N	Cumulative normal probability

IN THE LESSON

- Here is the famous result for a European option, one exercised only at expiry, on a stock paying no dividends.
- The call value is the stock price times N of d one, minus the strike times e to the minus r T times N of d two.
- The put is the strike, discounted, times N of minus d two, minus the stock price times N of minus d one.
- N is the cumulative normal probability, a number between zero and one.
- And d one is the natural log of stock over strike, plus r plus half sigma squared times T, all divided by sigma root T.
- Then d two equals d one minus sigma root T.
- Keep sigma and T in the same units, years.

13 Reading the formula in words

- First piece: value received if exercised
- Second piece: discounted strike you expect to pay
- $N(d_2)$ = risk-neutral chance the call is exercised
- $N(d_1)$ = delta of the call
- Traps: undiscounted K, wrong sign for puts, units

IN THE LESSON

- Before the numbers, let's read the call formula as a story.
- The first piece, the stock price times N of d one, is what you expect to receive in value if you buy the stock when the option finishes in the money.
- The second piece, the discounted strike times N of d two, is what you expect to pay.
- N of d two has a neat meaning: it is the risk neutral probability that the call ends in the money.
- N of d one is the call's delta, how many dollars the option moves for a one dollar move in the stock.
- The most common errors are forgetting to discount the strike, using N of d where N of minus d is needed for a put, and mixing up units.

14 Example 4: the two inputs, d one and d two

Illustration: S_0 100, K 100, r 5%, σ 20%, T 1 year

ln(S₀ ÷ K)	ln(100 ÷ 100)	**0**
Drift term	(0.05 + 0.02) × 1	**0.07**
d_1	0.07 ÷ (0.20 × 1)	**0.35**
d_2	0.35 − 0.20	**0.15**

IN THE LESSON

- Let's price a call.
- The stock is one hundred dollars, the strike is one hundred, the risk free rate is five per cent, the volatility is twenty per cent, and there is one year to go.
- First d one.
- The natural log of one hundred over one hundred is zero.
- Next, r plus half sigma squared is point zero five plus point zero two, which is point zero seven, times one year.
- So the top of the fraction is point zero seven.
- The bottom is sigma root T, which is point two.
- Divide: d one is point three five.
- Then d two is d one minus point two, which is point one five.

15 Example 4 continued: the call price

Continue: look up N, discount the strike

$N(d_1)$	N(0.35)	**0.6368**
$N(d_2)$	N(0.15)	**0.5596**
Discounted strike	100 × e^−0.05	**95.12**
First piece	100 × 0.6368	**63.68**
Call value	63.68 − 95.12 × 0.5596	**10.45**

IN THE LESSON

- Now look up the normal probabilities.
- N of point three five is point six three six eight.
- N of point one five is point five five nine six.
- The strike discounted for a year is one hundred times e to the minus point zero five, which is ninety five dollars and twelve cents.
- So the first piece is one hundred times point six three six eight, sixty three sixty eight.
- The second piece is ninety five twelve times point five five nine six, which is fifty three twenty three.
- The call is worth sixty three sixty eight minus fifty three twenty three, about ten dollars forty five.
- Notice that the right to buy is worth more than ten per cent of the stock.

16 Example 4 continued: the put and the check

Same inputs: $N(-d_2)$ = 0.4404, $N(-d_1)$ = 0.3632

Put directly	95.12 × 0.4404 − 100 × 0.3632	**5.57**
Put by parity	10.45 + 95.12 − 100	**5.57**
Check on N	0.5596 + 0.4404	**1**
Chance call is exercised	$N(d_2)$, risk-neutral	**56%**

IN THE LESSON

- Now the put with the same strike and expiry.
- We need N of minus d two, which is one minus point five five nine six, point four four zero four, and N of minus d one, point three six three two.
- The put is ninety five twelve times point four four zero four, minus one hundred times point three six three two.
- That is forty one eighty nine, minus thirty six thirty two, about five dollars fifty seven.
- Check with put-call parity: the put equals the call plus the discounted strike minus the stock price.
- Ten forty five plus ninety five twelve minus one hundred is five fifty seven.
- Both routes agree, and the chance of the call finishing in the money is about fifty six per cent.

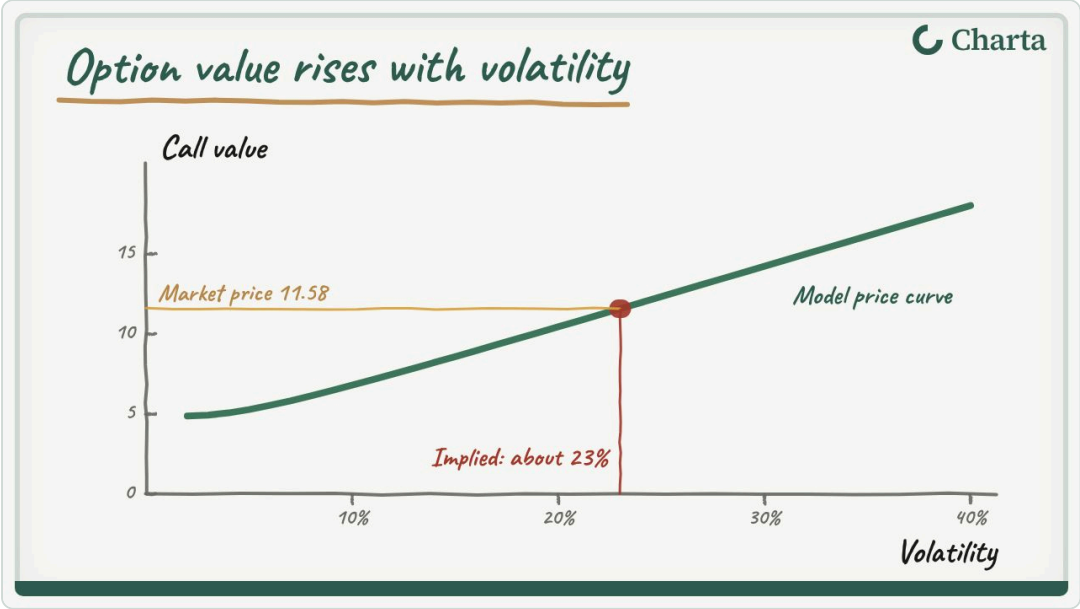
17 Implied volatility

- Observe the option price; solve for the volatility
- Historical volatility looks back; implied looks forward
- Higher option price means higher implied volatility
- Traders quote volatility to compare options

IN THE LESSON

- So far we put volatility in and got a price out.
- Markets work the other way round.
- Every input to the formula can be seen on a screen, the stock price, the strike, the interest rate, the time left, except volatility.
- But the option's price is on the screen.
- Implied volatility is the volatility that, put into the formula, makes the model price equal the market price.
- It is forward looking, because it reflects what traders expect about the future, unlike historical volatility, which looks back.
- A higher option price always means a higher implied volatility.
- Traders often quote options in volatility rather than dollars, because it makes different options easy to compare.

18 Option value rises with volatility



IN THE LESSON

- Here is why a single answer exists.
- Along the bottom is the volatility, and up the side the value of our at the money call, using the same stock, strike, rate and time as before.
- The curve rises steadily and never turns back.
- At twenty per cent the call is worth ten forty five, at thirty per cent fourteen twenty three.
- Suppose the market price is eleven fifty eight.
- Draw a line across at that height, find where it meets the curve, and drop down.
- The answer is about twenty three per cent.
- Because the curve climbs all the way, only one volatility fits, provided the price lies inside the no arbitrage bounds.

19 Example 5: finding it by trial and error

Call: S_0 100, K 100, r 5%, T 1. Market price 11.58

Try 20%	model price	**10.45 low**
Try 30%	model price	**14.23 high**
Try 25%	midpoint	**12.34 high**
Try 22.5%	midpoint	**11.39 low**
Keep narrowing		**about 23.0%**

IN THE LESSON

- The formula cannot be turned around algebraically, so we search.
- Our market price is eleven fifty eight.
- Guess twenty per cent: the model gives ten forty five, too low, so volatility must be higher.
- Guess thirty: fourteen twenty three, too high.
- Halve the gap and try twenty five: twelve thirty four, still too high.
- Try twenty two and a half: eleven thirty nine, a bit low.
- Keep halving, and we settle at about twenty three per cent.
- This halving is called bisection.
- Interpolation draws a straight line between two guesses.
- Newton Raphson is faster: it uses vega, the price's sensitivity to volatility, to decide how big each step should be.

20 Smiles, skews and the VIX

- Real options give different implied volatilities by strike
- Plotted against strike: a smile or a skew
- It shows the market pricing fat tails and jumps
- The VIX is built from index option implied volatilities

IN THE LESSON

- If the Black-Scholes-Merton model were exactly right, every option on the same stock would give the same implied volatility.
- In real markets it does not.
- Implied volatility differs by strike and by maturity.
- Plotted against strike it often curves up at both ends, a smile, or slopes down, a skew, and that is the market telling us that fat tails and jumps are priced in.
- Implied volatility is also used to read the market's mood.
- The VIX is an index built from the implied volatilities of options on a broad stock index, so when fear rises, option prices rise, and the VIX rises with them.

21 American calls and early exercise

- Exercise gives S − K
- Holding is worth at least S − PV of K, which is more
- The call also keeps insurance against a fall
- No dividends: never exercise a call early

IN THE LESSON

- Now American options, which can be exercised at any time before expiry.
- The question is when exercising early beats waiting.
- Take a call on a stock that pays no dividends.
- Exercising gives you only the stock price minus the strike.
- Holding the call is worth at least the stock price minus the present value of the strike, and that is more, because the strike you pay later is worth less today.
- On top of that the call keeps its insurance against a price fall.
- So with no dividends, never exercise an American call early.
- It is worth the same as the European call.
- Dividends change the picture, because the holder of an option does not receive them.

22 Example 6: the dividend test

Call, strike 100, r 5%. First two rows: S 130, no dividend

Exercise now: S − K	130 − 100	**30.00**
Hold, at least: S − PV of K	130 − 95.12	**34.88**
Interest saved, 3 months	100 × (1 − 0.9876)	**1.24**
Dividend 1.00 against 1.24	1.00 ≤ 1.24	**Hold**
Dividend 2.00 against 1.24	2.00 > 1.24	**May exercise**

IN THE LESSON

- Let's make it concrete, with strike one hundred and a five per cent interest rate.
- First, no dividends, stock at one hundred and thirty.
- Exercising now gives thirty.
- Holding is worth at least one thirty minus the discounted strike, ninety five twelve, which is thirty four eighty eight.
- So waiting wins by at least four eighty eight.
- Now add a dividend.
- Early exercise can only make sense just before an ex dividend date, and it is ruled out when the dividend is no more than the interest saved on the strike, which is K times one minus e to the minus r t.
- With three months left that is one twenty four.
- A dividend of one dollar is below it, so hold.
- Two dollars is above, so exercise may be best.

23 American puts are different

- Deep in-the-money put: strike received earns interest
- Exercise early even with no dividends
- A dividend makes waiting more attractive for a put
- Calls: dividends. Puts: interest on the strike

IN THE LESSON

- Puts behave differently.
- Even with no dividends, a deep in the money put may be worth exercising early.
- The reason is that exercising brings in the strike price in cash right now, and cash can earn interest.
- If the stock has fallen almost to zero, the most the put can ever pay is the strike, so waiting only delays receiving it.
- By exercising, you collect the strike today and start earning interest.
- Dividends push the other way for a put.
- When a stock pays a dividend its price drops, which raises the put's intrinsic value, so waiting through the ex dividend date becomes more attractive.
- Remember: calls and dividends, early exercise just before the date.
- Puts, interest on the strike.

24 Stocks that pay a yield

$$c = S_0 e^{-qT} N(d_1) - K \cdot e^{-rT} N(d_2)$$
$$d_1 = [\ln(S_0 \div K) + (r - q + \sigma^2 \div 2)T] \div (\sigma\sqrt{T})$$
$$d_2 = d_1 - \sigma\sqrt{T}$$

Yield, q	Continuous dividend yield
Risk-free rate, r	Risk-free rate
Volatility, sigma	Annual volatility
Time in years, T	Time in years

IN THE LESSON

- The formula extends to other assets by giving the underlying a continuous yield, q, the rate at which it pays out income, expressed as a percentage of its price.
- For a stock with yield q, the call value is the stock price times e to the minus q T times N of d one, minus the strike times e to the minus r T times N of d two.
- Inside d one, the drift becomes r minus q plus half sigma squared.
- The reasoning is simple: the holder of the option misses the payout, so the stock price is reduced by the yield.
- Set q to zero and we are back to the original formula.
- The put is built the same way.

25 Example 7: a three per cent dividend yield

Same inputs as Example 4, but yield q = 3%

d_1	(0.05 − 0.03 + 0.02) ÷ 0.20	**0.20**
d_2	0.20 − 0.20	**0.00**
$S_0e^{-qT} \times N(d_1)$	97.04 × 0.5793	**56.21**
$Ke^{-rT} \times N(d_2)$	95.12 × 0.5000	**47.56**
Call value	56.21 − 47.56	**8.65**

IN THE LESSON

- Same stock, strike, rate, volatility and time as before, but the stock now pays a continuous yield of three per cent.
- In d one, the drift is r minus q plus half sigma squared, which is point zero five minus point zero three plus point zero two, point zero four.
- Divide by point two and d one is point two.
- Then d two is point two minus point two, which is zero.
- Next, the stock price is cut by e to the minus point zero three, to ninety seven oh four.
- Times N of point two, point five seven nine three, gives fifty six twenty one.
- The discounted strike, ninety five twelve times N of zero, point five, is forty seven fifty six.
- The call is eight sixty five, lower than ten forty five without the dividend.

26 Other underlyings, same formula

Pick the right q, then use the yield formula

Underlying	What q is
Stock, yield q	The yield
Currency option	Foreign rate
Futures (Black)	q = r
Known cash dividends	PV off S_0

IN THE LESSON

- The same formula serves several underlyings, once we decide what q means.
- A stock with a continuous yield uses its yield.
- A currency option treats the foreign currency as an asset that pays interest at the foreign risk free rate, so q is the foreign rate, and r is the domestic one.
- A common slip is swapping the two.
- If the dividends are known cash amounts during a European option's life, subtract their present value from the stock price and use the ordinary formula.
- And an option on a futures contract uses Black's model, which is the case where q equals r, because a futures position needs no money upfront and so has no cost of carry.

27 Example 8: an option on a futures price

Futures price F_0 100, K 100, r 5%, σ 20%, T 1 year

d_1	(0 + 0.02) ÷ 0.20	**0.10**
d_2	0.10 − 0.20	**−0.10**
$N(d_1) - N(d_2)$	0.5398 − 0.4602	**0.0797**
Discount factor	$e^{-0.05}$	**0.9512**
Call value	0.9512 × 100 × 0.0797	**7.58**

IN THE LESSON

- For a futures option, Black's model uses the futures price in place of the stock price, and d one has no interest rate term.
- Take a futures price of one hundred, strike one hundred, five per cent interest, twenty per cent volatility and one year.
- The d one is the natural log of one, zero, plus half sigma squared times T, point zero two, divided by point two.
- That is point one.
- And d two is point one minus point two, minus point one.
- N of point one is point five three nine eight and N of minus point one is point four six oh two, so the difference is point oh seven nine seven.
- Multiply by one hundred and by e to the minus point zero five: the call is seven fifty eight.
- Never plug in the spot price here.

28 What a warrant is

- A warrant is a call a company issues on its own stock
- Exercise: company creates new shares, receives the strike
- More shares means existing owners are diluted
- An ordinary call changes no share counts

IN THE LESSON

- The last idea is the warrant.
- A warrant is a call option that a company issues on its own shares.
- Think of it as a coupon, sold to investors, giving them the right to buy a new share at a set price.
- When someone exercises, the company does not buy a share in the market.
- It creates a brand new share and receives the strike price.
- That means there are more shares afterwards, each owning a slightly smaller slice of the firm, so existing shareholders are diluted.
- This is the key difference from an ordinary traded call, where the stock is simply handed from one investor to another and the number of shares does not change.

29 Valuing a warrant

Warrant = N ÷ (N + M) × c
Payoff = N ÷ (N + M) × (S_T − K)
Dilution cost = M × warrant value

Shares, N	Shares already outstanding
Warrants, M	Warrants, one share each
Plain call value	Ordinary call, same terms

IN THE LESSON

- To value a warrant, first value an otherwise identical ordinary call with the Black-Scholes-Merton formula, using today's share price, which already reflects the market's expectation of dilution.
- Then scale it down.
- If N shares are outstanding and M warrants each give the right to one share, the warrant is worth N divided by the quantity N plus M, times the call value.
- The same factor applies to the payoff at exercise: N over N plus M, times the share price just before exercise minus the strike.
- The dilution cost to existing shareholders is the total value of all the warrants, M times the value of one.
- Common errors: using M over N plus M, or applying no factor at all.

30 Example 9: warrants and dilution

N 1,000,000 shares, M 200,000 warrants, K 100, plain call 10.45

Dilution factor	1.0m ÷ (1.0m + 0.2m)	**0.8333**
Value per warrant	0.8333 × 10.45	**8.71**
Dilution cost	200,000 × 8.71	**$1.74m**
Plain call, overstated	200,000 × 10.45	**$2.09m**
Payoff if S_T is 130	0.8333 × (130 − 100)	**25.00**

IN THE LESSON

- Suppose a company has one million shares and has issued two hundred thousand warrants, each for one share, with the strike at one hundred.
- An otherwise identical call is worth ten forty five.
- The factor is one million over one point two million, which is point eight three three three.
- So each warrant is worth point eight three three three times ten forty five, about eight seventy one.
- The dilution cost is two hundred thousand times eight seventy one, about one point seven four million dollars.
- Had we used the plain call, we would have got two point oh nine million, so we would overstate the cost.
- And if the share price just before exercise is one thirty, the payoff per warrant is point eight three three three times thirty, which is twenty five.

31 Recap: the main ideas

- Log returns normal, mean $\mu - \sigma^2 \div 2$; price lognormal
- Volatility: sample std dev × √252 for daily data
- $c = S_0N(d_1) - Ke^{(-rT)}N(d_2)$; implied volatility by search
- Early exercise: calls and dividends; puts and interest
- Yield q, Black's futures model, warrant factor $N \div (N+M)$

IN THE LESSON

- Let's pull it together.
- Stock prices follow geometric Brownian motion, so log returns are normal with mean mu minus half sigma squared times T, and the price is lognormal.
- Historical volatility comes from the sample standard deviation of log returns, annualised by the square root of the periods.
- The model rests on a hedge, which is why mu drops out, and its assumptions fail through fat tails, jumps and dividends.
- The call is S zero N of d one minus the discounted strike times N of d two.
- Implied volatility is found by search.
- Calls on non dividend stocks are never exercised early.
- Add a yield and a warrant dilution factor to extend it.

32 Try three quick questions

1. μ 10%, σ 30%, one year. Mean log return?
   Answer: $10\% - 30\%^2 \div 2 = 5.5\%$.
2. $d_1 = 0.4$, σ = 20%, T = 1. What is d_2?
   Answer: $0.4 - 0.2 = 0.2$.
3. 4m shares, 1m warrants, equivalent call 10. One warrant?
   Answer: $4 \div 5 \times 10 = 8$.

IN THE LESSON

- Now try three quick questions yourself.
- Pause the video, think, then listen.
- Question one.
- A stock has an expected return of ten per cent and a volatility of thirty per cent.
- What is the mean of its one year log return?
- Ten minus half of nine is five and a half per cent.
- Question two.
- A call has d one of point four and a volatility of twenty per cent for one year.
- What is d two?
- Point four minus point two is point two.
- Question three.
- A company has four million shares and one million warrants, and an equivalent call is worth ten dollars.
- What is one warrant worth?
- Four over five times ten is eight dollars.