

The Black-Scholes formula and its applications. This Section deduces the Black-Scholes formula for a European call or put, as a consequence of risk-neutral valuation in the continuous time limit. Then we discuss the delta, gamma, vega, theta, and rho of a portfolio, and their significance for hedging. At the end, we briefly discuss the volatility smile and skew. The material covered here is very important, since it concerns the practical impact of the Black-Scholes-Merton theory. These notes are far from a complete treatment; Chapters 18 and 19 of Hull’s 8th edition are strongly recommended for additional discussion.

We’ll do two passes, as usual. First we consider options on a non-dividend-paying underlying with lognormal dynamics. Then we consider options on a forward price with lognormal dynamics. As we’ll see, the latter case is actually simpler and more general, because the forward price is a martingale under the risk-neutral measure no matter what the value of the risk-free rate (and regardless of whether the underlying pays a dividend). Thus we could alternatively have started with options on a forward price, then deduced the results for options on a non-dividend-paying stock price (and options on a dividend-paying underlying, like a foreign currency rate) from that.

All the mathematics in this section uses probability and calculus to derive conclusions from the results obtained in Section 4. We won’t be introducing any new financial assumptions or arbitrage arguments (but remember that the results in Section 4 were based on such arguments). Later in the semester we’ll derive results similar to those of Section 4 in other settings. The analysis in this section won’t have to be redone – it will permit us to immediately deduce the prices and hedges of puts and calls in those settings too.

The Black-Scholes formula for a European call or put. The upshot of Section 4 is this: the value at time t of a European option with payoff $f(s_T)$ is

$$V(f) = e^{-r(T-t)} E_{\text{RN}}[f(s_T)].$$

Here $E_{\text{RN}}[f(s_T)]$ is the expected value of the price at maturity with respect to a special probability distribution – the risk-neutral one. For a non-dividend-paying stock with lognormal dynamics, this distribution is determined by the property that

$$s_T = s_t \exp \left[\left(r - \frac{1}{2} \sigma^2 \right) (T - t) + \sigma \sqrt{T - t} Z \right]$$

where s_t is the spot price at time t and Z is Gaussian with mean 0 and variance 1. Equivalently: $\log[s_T/s_t]$ is Gaussian with mean $(r - \frac{1}{2}\sigma^2)(T - t)$ and variance $\sigma^2(T - t)$.

The value of the option can be evaluated for any payoff f by numerical integration. But for puts and calls we can do better, by obtaining explicit expressions in terms of the “cumulative distribution function”

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du.$$

($N(x)$ is the probability that a Gaussian random variable with mean 0 and variance 1 has value $\leq x$.) The explicit formulas have enormous advantages over numerical integration: besides being easy to evaluate, they permit us to see quite directly how the value and the hedge portfolio depend on strike price, spot price, risk-free rate, and volatility.

It's sufficient, of course, to consider $t = 0$. Let

$$\begin{aligned} c[s_0, T; K] &= \text{value at time 0 of a European call with strike } K \\ &\quad \text{and maturity } T, \text{ if the spot price is } s_0; \\ p[s_0, T; K] &= \text{value at time 0 of a European put with strike } K \\ &\quad \text{and maturity } T, \text{ if the spot price is } s_0. \end{aligned}$$

The explicit formulas are:

$$\begin{aligned} c[s_0, T; K] &= s_0 N(d_1) - K e^{-rT} N(d_2) \\ p[s_0, T; K] &= K e^{-rT} N(-d_2) - s_0 N(-d_1) \end{aligned}$$

in which

$$\begin{aligned} d_1 &= \frac{1}{\sigma\sqrt{T}} \left[\log(s_0/K) + (r + \frac{1}{2}\sigma^2)T \right] \\ d_2 &= \frac{1}{\sigma\sqrt{T}} \left[\log(s_0/K) + (r - \frac{1}{2}\sigma^2)T \right] = d_1 - \sigma\sqrt{T}. \end{aligned}$$

To derive these formulas we use the following result. (The Lemma toward the end of Section 4 was a special case.)

Lemma: Suppose X is Gaussian with mean μ and variance σ^2 . Then for any real numbers a and k ,

$$E \left[e^{aX} \text{ restricted to } X \geq k \right] = e^{a\mu + \frac{1}{2}a^2\sigma^2} N(d)$$

with $d = (-k + \mu + a\sigma^2)/\sigma$.

Proof: The left hand side is defined by

$$E \left[e^{aX} \text{ restricted to } X \geq k \right] = \frac{1}{\sigma\sqrt{2\pi}} \int_k^\infty e^{ax} \exp \left[\frac{-(x - \mu)^2}{2\sigma^2} \right] dx.$$

Complete the square:

$$ax - \frac{(x - \mu)^2}{2\sigma^2} = a\mu + \frac{1}{2}a^2\sigma^2 - \frac{[x - (\mu + a\sigma^2)]^2}{2\sigma^2}.$$

Thus

$$E \left[e^{aX} | X \geq k \right] = e^{a\mu + \frac{1}{2}a^2\sigma^2} \cdot \frac{1}{\sigma\sqrt{2\pi}} \int_k^\infty \exp \left[\frac{-[x - (\mu + a\sigma^2)]^2}{2\sigma^2} \right] dx.$$

If we set $u = [x - (\mu + a\sigma^2)]/\sigma$ and $\kappa = [k - (\mu + a\sigma^2)]/\sigma$ this becomes

$$\begin{aligned} e^{a\mu + \frac{1}{2}a^2\sigma^2} \cdot \frac{1}{\sqrt{2\pi}} \int_\kappa^\infty e^{-u^2/2} du &= e^{a\mu + \frac{1}{2}a^2\sigma^2} [1 - N(\kappa)] \\ &= e^{a\mu + \frac{1}{2}a^2\sigma^2} N(d) \end{aligned}$$

where $d = -\kappa = (-k + \mu + a\sigma^2)/\sigma$.

We note in passing the following corollary: if X is Gaussian with mean μ and variance σ^2 , then the probability that $X \geq k$ is $N(d)$ with $d = (\mu - k)/\sigma$. This is precisely the assertion of the Lemma when $a = 0$. This result permits us, for example, to calculate the probability that an option will be “in the money” at maturity. (Note that the answer depends on μ . So you’ll get different answers using the “subjective” versus “risk-neutral” probabilities.)

We now apply the Lemma to price a European call. Our task is to evaluate

$$e^{-rT} \int_{-\infty}^{\infty} (s_0 e^x - K)_+ \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[-\frac{(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx.$$

The integrand is nonzero when $s_0 e^x > K$, i.e. when $x > \log(K/s_0)$. Applying the Lemma with $a = 1$ and $k = \log(K/s_0)$ we get

$$e^{-rT} \int_k^{\infty} s_0 e^x \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[-\frac{(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx = s_0 N(d_1);$$

applying the Lemma again with $a = 0$ we get

$$e^{-rT} \int_k^{\infty} K \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[-\frac{(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx = K e^{-rT} N(d_2);$$

combining these results gives the formula for $c[s_0, T; K]$.

The formula for the value of a European put can be obtained similarly. Or – easier – we can derive it from the formula for a call, using put-call parity:

$$\begin{aligned} p[s_0, T; K] &= c[s_0, T; K] + K e^{-rT} - s_0 \\ &= K e^{-rT} [1 - N(d_2)] - s_0 [1 - N(d_1)] \\ &= K e^{-rT} N(-d_2) - s_0 N(-d_1). \end{aligned}$$

For options with maturity T and strike price K , the value at any time t is naturally $c[s_t, T - t; K]$ for a call, $p[s_t, T - t; K]$ for a put.

Hedging. We know how to hedge in the discrete-time, multiperiod binomial tree setting: the payoff is replicated by a portfolio consisting of $\Delta = \Delta(0, s_0)$ units of stock and a (long or short) bond, chosen to have the same value as the derivative claim. At time δt the stock price changes to $s_{\delta t}$ and the value of the hedge portfolio changes by $\Delta(s_{\delta t} - s_0)$. The new value of the hedge portfolio is also the new value of the option, so

$$\Delta(0, s_0) = \frac{\text{change in value of option from time 0 to } \delta t}{\text{change in value of stock from time 0 to } \delta t}.$$

The replication strategy requires a self-financing trade at every time step, adjusting the amount of stock in the portfolio to match the new value of Δ .

In the real world prices are not confined to a binomial tree, and there are no well-defined time steps. We cannot trade continuously. So while we can pass to the continuous time limit for the value of the option, we must still trade at discrete times in our attempts to replicate it. Suppose, for simplicity, we trade at equally spaced times with interval δt . What to use for the initial hedge ratio Δ ? Not being clairvoyant we don't know the value of the stock at time δt , so we can't use the formula given above. Instead we should use its continuous-time limit:

$$\Delta(0, s_0) = \frac{\partial(\text{value of option})}{\partial(\text{value of stock})}.$$

There's a subtle point here: if the stock price changes continuously in time, but we only rebalance at discretely chosen times $j\delta t$, then we cannot expect to replicate the option perfectly using self-financing trades. Put differently: if we maintain the principle that the value of the hedge portfolio is equal to that of the option at each time $j\delta t$, then our trades will no longer be self-financing. We will address this point soon, after developing the continuous-time Black-Scholes theory. We'll show then that (if transaction costs are ignored) the expected cost of replication tends to 0 as $\delta t \rightarrow 0$. (In practice transaction costs are *not* negligible; deciding when, really, to rebalance, taking into account transaction costs, is an important and interesting problem – but one beyond the scope of this course.)

For the European put and call we can easily get formulas for Δ by differentiating our expressions for c and p : at time T from maturity the hedge ratio should be

$$\Delta = \frac{\partial}{\partial s_0} c[s_0, T; K] = N(d_1)$$

for the call, and

$$\Delta = \frac{\partial}{\partial s_0} p[s_0, T; K] = -N(-d_1)$$

for the put. The “hard way” to see this is an application of chain rule: for example, in the case of the call,

$$\frac{\partial}{\partial s_0} c = N(d_1) + s_0 N'(d_1) \frac{\partial d_1}{\partial s} - K e^{-rT} N'(d_2) \frac{\partial d_2}{\partial s}.$$

But $d_2 = d_1 - \sigma\sqrt{T}$, so $\partial d_1/\partial s = \partial d_2/\partial s$; also $N'(x) = \frac{1}{\sqrt{2\pi}} \exp[-x^2/2]$. It follows with some calculation that

$$s_0 N'(d_1) \frac{\partial d_1}{\partial s} - K e^{-rT} N'(d_2) \frac{\partial d_2}{\partial s} = 0,$$

so finally $\partial c/\partial s_0 = N(d_1)$ as asserted. There is however an easier way: differentiate the original formula expressing the value as a discounted risk-neutral expectation. Passing the derivative under the integral, for a call with strike K :

$$\Delta = \frac{\partial}{\partial s_0} e^{-rT} \int_{-\infty}^{\infty} (s_0 e^x - K)_+ \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[\frac{-(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx$$

$$\begin{aligned}
&= e^{-rT} \int_{-\infty}^{\infty} \frac{\partial(s_0 e^x - K)_+}{\partial s_0} \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[\frac{-(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx \\
&= e^{-rT} \int_{\log(K/s_0)}^{\infty} e^x \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[\frac{-(x - [r - \sigma^2/2]T)^2}{2\sigma^2 T}\right] dx \\
&= N(d_1).
\end{aligned}$$

Notice that since $0 < N(d_1) < 1$ and $-1 < -N(-d_1) < 0$, the delta of a call lies between 0 and 1 while the delta of a put lies between 0 and -1 .

Since we have explicit formulas for the value of a put or call, we can differentiate them to learn the dependence on the underlying parameters. Some of these derivatives have names:

Definition	Call	Put
delta = $\frac{\partial}{\partial s_0}$	$N(d_1) > 0$	$-N(-d_1) < 0$
gamma = $\frac{\partial^2}{\partial s_0^2}$	$\frac{1}{s_0\sigma\sqrt{2\pi T}} \exp(-d_1^2/2) > 0$	$\frac{1}{s_0\sigma\sqrt{2\pi T}} \exp(-d_1^2/2) > 0$
theta = $\frac{\partial}{\partial(-T)}$	$-\frac{s_0\sigma}{2\sqrt{2\pi T}} \exp(-d_1^2/2) - rKe^{-rT}N(d_2) < 0$	$-\frac{s_0\sigma}{2\sqrt{2\pi T}} \exp(-d_1^2/2) + rKe^{-rT}N(-d_2) \begin{matrix} > \\ < \end{matrix} 0$
vega = $\frac{\partial}{\partial\sigma}$	$\frac{s_0\sqrt{T}}{\sqrt{2\pi}} \exp(-d_1^2/2) > 0$	$\frac{s_0\sqrt{T}}{\sqrt{2\pi}} \exp(-d_1^2/2) > 0$
rho = $\frac{\partial}{\partial r}$	$TKe^{-rT}N(d_2) > 0$	$-TKe^{-rT}N(-d_2) < 0$.

These formulas apply at time $t = 0$; the formulas applicable at any time t are similar, with T replaced by $T - t$. These are obviously useful for understanding how the value of the option changes with time, volatility, etc. But more: they are useful for designing improved hedges. For example, suppose a bank sells two types of options on the same underlying asset, with different strike prices and maturities. As usual the bank wants to limit its exposure to changes in the stock price; but suppose in addition it wants to limit its exposure to changes (or errors in specification of) volatility. Let $i = 1, 2$ refer to the two types of options, and let n_1, n_2 be the quantities held of each. (These are negative if the bank sold the options.) The bank naturally also invests in the underlying stock and in risk-free bonds; let n_s and n_b be the quantities held of each. Then the value of the bank's initial portfolio is

$$V_{\text{total}} = n_1 V_1 + n_2 V_2 + n_s s_0 + n_b.$$

We already know how the stock and bond holdings should be chosen if the bank plans to replicate (dynamically) the options: they should satisfy

$$V_{\text{total}} = 0$$

and

$$n_1 \Delta_1 + n_2 \Delta_2 + n_s = 0.$$

Notice that the latter relation says $\partial V_{\text{total}}/\partial s_0 = 0$: the value of the bank's holdings are insensitive (to first order) to changes in the stock price.

If we were dealing in just one option there would be no further freedom: we would have two homogeneous equations in three variables n_1, n_s, n_b , restricting their values to a line – so that n_1 determines n_s and n_b . That's the situation we're familiar with. But if we're dealing in two (independent) options then we have the freedom to impose one additional linear equation. For example we can ask that the portfolio be insensitive (to first order) to changes in σ by imposing the additional condition

$$n_1 \text{Vega}_1 + n_2 \text{Vega}_2 = 0.$$

Thus: by selling the two types of assets in the proper proportions the bank can reduce its exposure to change or misspecification of volatility.

If the bank sells three types of options then we have room for yet another condition – e.g. we could impose first-order insensitivity to changes in the risk-free rate r . And so on. It is not actually necessary that the bank use the underlying stock as one of its assets. Each option is *equivalent* to a portfolio consisting of stock and risk-free bond; so a portfolio consisting entirely of options and a bond position will function as a hedge portfolio so long as its total Δ is equal to 0.

Replication requires dynamic rebalancing. The bank must change its holdings at each time increment to set the new Δ to 0. In the familiar, one-option setting this was done by adjusting the stock and bond holdings, keeping the option holding fixed. In the present, two-option setting, maintaining the additional condition $\text{Vega}_{\text{total}} = 0$ will require the ratio between n_1 and n_2 to be dynamically updated as well, i.e. the bank will have to sell or buy additional options as time proceeds.

Options on a forward rate. Recall from Section 4: for a European option on a forward price with payoff $f(\mathcal{F}_T)$, the value at time t is

$$V(f) = e^{-r(T-t)} E_{\text{RN}}[f(\mathcal{F}_T)];$$

Moreover under the risk-neutral probability distribution

$$\mathcal{F}_T = \mathcal{F}_t \exp \left[-\frac{1}{2} \sigma^2 (T-t) + \sigma \sqrt{T-t} Z \right]$$

where \mathcal{F}_t is the forward price at time t and Z is Gaussian with mean 0 and variance 1. Equivalently:

$$\mathcal{F}_T = \mathcal{F}_0 e^X$$

where X is Gaussian with mean $-\frac{1}{2} \sigma^2 (T-t)$ and variance $\sigma^2 (T-t)$.

For a call (payoff $(\mathcal{F}_T - K)_+$) or a put (payoff $(K - \mathcal{F}_T)_+$) we can derive explicit valuation formulas using the same approach as we used for options on a non-dividend-paying stock. Focusing on time 0 as usual, the result is:

$$\begin{aligned} c[\mathcal{F}_0, T; K] &= e^{-rT} [\mathcal{F}_0 N(d_1) - KN(d_2)] \\ p[\mathcal{F}_0, T; K] &= e^{-rT} [KN(-d_2) - \mathcal{F}_0 N(-d_1)] \end{aligned}$$

where

$$\begin{aligned} d_1 &= \frac{1}{\sigma\sqrt{T}} \left[\log(\mathcal{F}_0/K) + \frac{1}{2}\sigma^2 T \right] \\ d_2 &= \frac{1}{\sigma\sqrt{T}} \left[\log(\mathcal{F}_0/K) - \frac{1}{2}\sigma^2 T \right] = d_1 - \sigma\sqrt{T}. \end{aligned}$$

For variety, instead of writing everything out in terms of integrals as we did earlier, let's derive the value of the call using the notation of the Lemma proved near the beginning of this section:

$$c[\mathcal{F}_0, T; K] = e^{-rT} \left\{ \mathcal{F}_0 E[e^X \text{ restricted to } X \geq k] - KE[1 \text{ restricted to } X \geq k] \right\}$$

where $k = \ln(K/\mathcal{F}_0)$ and X is Gaussian with mean $-\frac{1}{2}\sigma^2 T$ and variance $\sigma^2 T$. By the Lemma,

$$E[e^X \text{ restricted to } X \geq k] = N(d_1) \quad \text{and} \quad E[1 \text{ restricted to } X \geq k] = N(d_2).$$

Our assertion about the value of the call is an immediate consequence. The argument for the put is similar (or one can use put-call parity).

As noted already in Section 4, the valuation of options on a forward price has the pleasant feature that the risk-free rate is almost irrelevant – it enters only through the discount factor e^{-rT} out front. *This is true even if the underlying pays a dividend at constant rate d .* (A forward price is always a martingale under the risk-neutral measure. Indeed, the proof in Section 2 that $\mathcal{F}_{\text{now}} = q\mathcal{F}_{\text{up}} + (1 - q)\mathcal{F}_{\text{down}}$ applies regardless whether the underlying pays a dividend or not. To use the results from Section 4 we must assume the forward price is lognormal. This will be true, from Section 1, if the underlying is lognormal and dividends are paid at a constant rate, since then $F_t = e^{(r-d)(T-t)} s_t$.)

Our previous formulas for calls and puts on a non-dividend-paying stock can be deduced from the ones just obtained for options on a forward price, by simply substituting $\mathcal{F}_0 = e^{rT} s_0$. (Exercise: check this.) Similarly, the value of an option on an underlying with constant dividend yield d is obtained by substituting $\mathcal{F}_0 = e^{(r-d)T} s_0$.

The calculation of “the Greeks” for options on a forward price is no more difficult than for options on a non-dividend-paying stock. (Indeed, one set of formulas can be obtained from

the other via chain rule.) For the record:

Definition	Call	Put
delta = $\frac{\partial}{\partial \mathcal{F}_0}$	$e^{-rT} N(d_1) > 0$	$-e^{-rT} N(-d_1) < 0$
gamma = $\frac{\partial^2}{\partial \mathcal{F}_0^2}$	$\frac{e^{-rT}}{\mathcal{F}_0 \sigma \sqrt{2\pi T}} \exp(-d_1^2/2) > 0$	$\frac{e^{-rT}}{\mathcal{F}_0 \sigma \sqrt{2\pi T}} \exp(-d_1^2/2) > 0$
theta = $\frac{\partial}{\partial(-T)}$	$-\frac{e^{-rT} \mathcal{F}_0 \sigma}{2\sqrt{2\pi T}} \exp(-d_1^2/2)$ $+re^{-rT}[\mathcal{F}_0 N(d_1) - KN(d_2)]$	$-\frac{e^{-rT} \mathcal{F}_0 \sigma}{2\sqrt{2\pi T}} \exp(-d_1^2/2)$ $+re^{-rT}[KN(-d_2) - \mathcal{F}_0 N(-d_1)]$
vega = $\frac{\partial}{\partial \sigma}$	$\frac{e^{-rT} \mathcal{F}_0 \sqrt{T}}{\sqrt{2\pi}} \exp(-d_1^2/2) > 0$	$\frac{e^{-rT} \mathcal{F}_0 \sqrt{T}}{\sqrt{2\pi}} \exp(-d_1^2/2) > 0$
rho = $\frac{\partial}{\partial r}$	$-Te^{-rT}[\mathcal{F}_0 N(d_1) - KN(d_2)] < 0$	$-Te^{-rT}[KN(-d_2) - \mathcal{F}_0 N(-d_1)] < 0.$

Our previous comments on hedging using the Greeks applies equally to options on forwards.

Further discussion of “the Greeks”. Most of the following applies equally to options on a forward price or options on a stock. (Where it’s important to choose, we focus on options on a forward price.)

- All our formulas for the Greeks are from the viewpoint of the option buyer. The option seller’s Greeks have the same magnitudes but the opposite sign. Greeks for portfolios of options can be computed as sums of the Greeks for each individual option position in the portfolio.
- Gamma, being the derivative of delta with respect to price, is a measure of how much the delta hedge will change with changes in the price. But more important, it is also a measure of how much the price of a delta hedged option will change with *jumps* in the forward price. (Practically speaking, a jump can be viewed as a price move that takes place too quickly to allow rehedging). As you can see from Hull’s figure 18.7 and the associated text, the larger the change in the delta, the greater the impact of a sudden move that is not hedged.
- Notice that the sign of a call delta is opposite from the sign of a put delta. The buyer of a call has a position that will benefit from a rise in the forward price, while the buyer of a put has a position that will benefit from a fall in the forward price.
- By contrast, notice that the sign of a call gamma is the same as the sign of a put gamma, the sign of a call vega is the same as the sign of a put vega, and the first (generally dominant) of the two terms in a call theta is the same as the sign of the first of the two terms in a put theta. Purchasers of options, whether calls or puts, benefit from greater price movement and longer time to option expiry, while sellers of

options, whether calls or puts, are hurt by greater price movement and longer time to option expiry. An equivalent way of seeing this is that a call can be transformed into a put (and vice versa) by a forward (through put-call parity) and that the forward, not being an option, is not sensitive to volatility and has relatively weak dependence on the time to option expiry.

- In fact, the term $\exp(-d_1^2/2)$ appears in all three of these derivatives - but multiplied by different constants. Time to option expiry appears in the numerator of the constant multiplier for vega and in the denominator of the constant multiplier for gamma and theta. This means that a longer-term option will tend to have a higher vega and lower gamma and theta than a shorter-term option. So a change in volatility will have a greater impact on longer-term options than shorter-term ones, but a jump in prices will have a greater impact on shorter-term options than on longer-term ones. This allows you to answer the popular options "brain teaser": How do you construct a portfolio that is positive vega and negative gamma? Just buy long-term options and sell short-term options in the right proportions.
- When an option is at-the-money (strike = forward), its delta is roughly 50% and its vega and gamma reach their maxima (see Hull figures 18.3, 18.9, and 18.11).
- Theta is defined as the derivative with respect to a decrease in time to option expiry. This is the standard market convention, since the primary concern is with a "decay" in option value as time passes. (The market convention is often to show the theta as the derivative with respect to a decrease of one business day, in which case the value returned by the formula needs to be divided by the number of business days in a year, about 252). The first (generally dominant) term of the theta represents the impact of shorter time on the impact of the option's volatility (this is the part of the theta due to the derivative with respect to T inside the $N(d_1)$ and $N(d_2)$ terms). Since T always appears with a factor involving σ , the first term of theta closely resembles the vega, which is the derivative with respect to σ . Shorter time to option expiry has the same impact as a decrease in volatility and so is always negative to the option purchaser. The second term represents the impact of a shorter time until the option payout is received. This is the part of the theta due to the derivative of T in the e^{-rT} term and so is closely related to the rho, which is the derivative with respect to r . It is always positive, reflecting the fact that a shorter time to option payout results in a less deeply discounted present value.
- The rho is negative, reflecting the fact that a higher risk-free rate will result in a more deeply discounted present value of the option payout.
- A classic interview question: A trading desk sells a 1 year option at an implied volatility of 16% (that is, the price is equal to the Black-Scholes formula price with 16% as the input for volatility). The desk delta-hedges the option until maturity. Realized volatility over the year turns out to be 15%. Has the desk made money? Your first instinct should be to say that the desk did make money, since the negative vega of a sold option indicates that positive profits will result from a decline in volatility. And, if volatility is constant during the year in question, your answer will be correct.

But if volatility varies from one part of the year to another, we must be concerned with sensitivity to volatility over a short time period, not just the entire period, and sensitivity to increased volatility over a short period is the gamma, the sensitivity of the option price to price jumps in the forward (in some sense, the vega represents an “integral” of all the different gamma exposures over shorter time periods). If the periods during the year in which volatility is higher than the average coincide with periods when gamma is higher (due to the forward price being close to the strike), and periods in which the volatility is lower than the average coincide with periods in which gamma is lower (due to the forward price being significantly higher or lower than the strike), the desk could lose money even though the volatility averaged 15%.

How is this used in practice? We’ve discussed hedging based on the Black-Scholes pricing formula and its sensitivities to the various parameters (“the Greeks”). Now let’s step back a bit and ask how plausible this strategy is in practice.

First, let’s review why this matters. As we stated in the opening lecture, the vast majority of derivatives users are not interested in performing an arbitrage trade – they just want to buy or sell an option because they want to achieve the payoff profile it offers. The very last thing most option buyers would want to do is perform a dynamic hedge of the option they just bought – this would cancel out the very payoff profits they were looking for. But all options users should care about the plausibility of arbitraging options through dynamic hedging, because it is only the possibility of a few arbitrage traders actually successfully engaging in such activity that allows all users to find reasonable prices by utilizing Black-Scholes theory. And having reasonable price parameters is valuable for all options buyers and sellers – it lets them know whether they are being offered a fair price; it lets them estimate the price at which they could liquidate an existing position (very important for financial reporting purposes); and it lets them estimate the riskiness of their positions by calculating sensitivity to changes in underlying variables.

Ever since Black and Scholes (and Merton) first proposed the theory we’ve been studying, there has been controversy – in both the academic and business worlds – about the adequacy of this no-arbitrage-based theory. Some of the objections raised should not be taken too seriously:

- The assumption of a lognormal distribution of forward prices has been questioned. In fact, as already noted, almost no market practitioners believe the lognormal distribution is the right one, but there are relatively easy fixes for this, which we’ll discuss shortly.
- The original Black-Scholes argument was stated in terms of a cash instrument underlying, rather than a forward price, and utilized an assumption of constant risk-free rates. But as we’ve seen, we can price options on a forward rate; then only the discount rate for lending with maturity T enters the formula. So the assumption of constant risk-free rates is not important.

- We started with binomial trees. But this was just a matter of convenience. We'll soon develop a more robust, continuous-time version of the theory (corresponding to the framework originally used by Black, Scholes, and Merton).

There are, however, two more serious objections:

1. The dynamic hedging argument assumes that hedging can be performed continuously. This is, of course, not achievable in practice and, in fact, cannot even be closely approximated – even if you could re hedge once every minute, the resulting transaction costs (in the real world, where bid-asked spreads do exist) would make the dynamic hedging strategy prohibitively expensive.
2. Our theory assumes that we know the volatility. In fact we don't; moreover, any study of historical price data will show that lognormal statistics are an imperfect approximation, and volatility is difficult to estimate with any accuracy. The very fact that options traders calculate their sensitivity to changes in volatility (vega) shows that no one takes this assumption seriously.

Price = 49, Interest rate = 5%, Dividend rate = 0, Forward price \approx 50
 Strike = 50
 Volatility = 20%
 Time to maturity = 20 weeks (.3846 years)
 Drift rate = 13%
 Option price = \$240,000 for 100,000 shares

Frequency of Reheding	<u>Performance Measure (ratio of Standard Deviation to Cost of Option)</u>			
	Stop-Loss	Delta Hedge		
		No Vol of Vol	10% Vol of Vol	33% Vol of Vol
5 weeks	102%	43%	44%	57%
4 weeks	93%	39%	41%	52%
2 weeks	82%	26%	29%	45%
1 week	77%	19%	22%	47%
1/2 week	76%	14%	18%	43%
1/4 week	76%	9%	14%	38%
Limit as frequency goes to 0	76%	0%	11%	40%
With no hedging, performance measure is 130%				

Table 1: Performance of dynamic hedging strategies (courtesy of Steve Allen)

And yet, many trading desks engage daily in what certainly looks like dynamic hedging of options positions, furiously changing their hedges based on deltas calculated using the Black-Scholes framework. How can we explain the practical utility of this apparently defective theory? Here are some suggestions:

1. Arbitrage does not have to be perfect to be effective. A trading strategy that reduces the uncertainty of the outcome to a relatively small window will still attract traders

who are seeking a low-risk return. We can't expect that such arbitrage will result in a single definite price for the option, but it will constrain the option price to within a narrow range. Column 3 of Table 1, which extends the results from tables 18.1 and 18.4 in Hull, shows that even relatively infrequent rehedging can still reduce a very large portion of the uncertainty of profit and loss.

2. Arbitrage traders in options, with almost no exceptions, engage in arbitrage trades on a large number of closely related options simultaneously. Aggregating deltas across trades, which are mixtures of option purchases and option sales, and of puts and calls, results in a lot of netting of required delta hedges. Changes in delta hedging, and hence transaction costs, can be kept relatively small, particularly when combined with the less frequent hedging discussed in the preceding point.

However, the uncertainty of volatility is still a major concern. Columns 4 and 5 of Table 1 show that when we introduce uncertainty of volatility, much more frequent hedging is required to achieve a given reduction in uncertainty; moreover, when the volatility evolves randomly there are certain floor levels of uncertainty that even continuous-time hedging cannot penetrate. A reasonably accurate summary of the situation is that dynamic hedging can almost completely eliminate uncertainty of results due to the final price of the underlying forward but cannot do anything to eliminate uncertainty of results due to uncertainty of volatility.

To deal with this situation, market participants must be able to quickly see the volatility assumption that is implied by a given option price. This is known as the *implied volatility* of the option price. We'll worry about how to calculate it later, but certainly we know how to check that it's been calculated correctly – just plug it into the Black-Scholes formula and see if the actual option price comes out. Implied volatility is an important concept in working with options. Market participants rarely talk about options prices; they almost always talk about implied volatilities and leave the actual option price as a detail to be worried about by the “back office.”

Working with implied volatility and the dynamic delta hedging calculations of Black-Scholes theory as their key tools, arbitrage traders typically proceed as follows:

1. If the implied volatility of an option is very high, an arbitrage trader will sell the option and use dynamic hedging to make sure that his P&L is only dependent on his judgment about volatility and not on the price level of the underlying forward. Conversely, if the implied volatility of an option is very low, an arbitrage trader will buy the option and dynamically hedge it to make sure that his P&L is only dependent on his judgment about volatility and not on the price level of the underlying forward. The net effect of the actions of such traders is to keep implied volatilities within reasonable bounds.
2. If two different options based on the same underlying have different implied volatilities, an arbitrage trader will buy the option with lower implied volatility, sell the option with higher implied volatility, and dynamically hedge the resulting position to try to reduce P&L sensitivity to the price level of the underlying forward. The arbitrage trader will still have to make a judgment about the degree to which the difference in

implied volatility reflects mispricing; it could instead have reasonable economic sources (for example: the actual evolution being different from the Black-Scholes assumption of lognormality, e.g. due to changes in volatility over time). But the net effect is that actions of arbitrage traders keep implied volatilities between options on the same underlying forward reasonably close to one another. The closer the options are in terms of strike levels and option expiry dates, the more we can expect the actions of arbitrage traders to keep their implied volatilities close together. This allows all market participants to estimate the implied volatility for an option whose price they can't readily observe by interpolation from the implied volatilities of options with nearby strikes and expiry dates they can observe.

Years of experience in options markets have proven to the satisfaction of most market participants that the Black-Scholes theory, when used in the ways we've just outlined, is sufficiently accurate to achieve good results. It is also possible to use Monte Carlo simulation to get good intuition as to how this works (this is done in the class Risk Management class; if you are interested in seeing some results along those lines, look at Chapter 9 of Steve Allen's book *Financial Risk Management*).

The volatility skew and smile. In practice, the implied volatility depends on the strike. It tends to be higher at strikes far above or far below the current forward price (Hull, figure 19.1), leading to the term *volatility smile*. Its graph (as a function of the strike) is usually asymmetric: the implied volatility is typically higher at very low strikes than at very high ones (Hull, figure 19.3), leading to the term *volatility skew*. We'll discuss briefly some explanations for these effects. A fuller discussion can be found in Hull's Chapter 19 and in section 9.6.2 of Allen's *Financial Risk Management* (John Wiley & Sons, 2003).

One type of explanation is based on the observed statistics of forward price dynamics. Like most financial time series, forward prices have fatter tails than one gets from a lognormal distribution. This can be explained (or modelled) as being due to stochastic volatility or the existence of jumps. (In a stochastic volatility model, the volatility is not constant but is itself a random variable; in a jump-diffusion the price can change abruptly at random times). Either effect (or a combination of the two) leads to fatter tails than a lognormal distribution. To see why fat tails produce an implied volatility smile, consider a call with strike far above the current price. For it to be in the money at maturity, the price must go up a lot. In the Black-Scholes framework, this is more likely if the volatility is larger. In the real world, fat tails make it more likely than the lognormal model predicts; thus fat tails have the same effect on the price as increasing the volatility. The same argument applies to deep out-of-the money puts. If the fatness of the tails is asymmetric then this argument predicts a skew as well as a smile.

A second type of explanation involves supply and demand. There are far more investors in stocks than short-sellers of stock. An investor with a long stock position (e.g. a mutual fund) would commonly buy protection against large decreases in the stock price by purchasing "protective puts" with strike well below the current price. Similarly, an investor with a short

position would commonly purchase protective calls with strikes well above the current price. But since most investors have long positions, there is much more demand for protective puts than for protective calls. This can create, at least temporarily, higher implied volatilities for low strikes than for high strikes. Now, if arbitrage traders could completely rely on the original Black-Scholes model (which assumes constant volatility) they would always be willing to sell the call protection demanded at low strikes and buy options at high strikes against them, bringing the two implied volatilities back into line. But, arbitrage traders cannot rely on the constant-volatility Black-Scholes model. As we've discussed, an arbitrage trader would only take this combination of positions if he believes the implied volatility difference is sufficiently out-of-line with likely price behavior to give reasonable assurance of a profit. So arbitrage trading places a brake on the impact of investor demand on volatility skew, but it does not entirely eliminate the impact of this effect.

Computing implied volatility. There is no closed form formula for implied volatility. Implied volatility is, however, easy to calculate numerically using Newton's method.

Recall that in general, Newton's method for solving an equation $f(x) = c$ is an iterative scheme. Starting with an initial guess x_0 , it produces better estimates x_1, x_2 , etc, by the rule

$$x_{n+1} = x_n + \frac{c - f(x_n)}{f'(x_n)}.$$

The logic is that by Taylor expansion, $f(x) \approx f(x_n) + f'(x_n)(x - x_n)$. If we set this linear approximation equal to c and solve for x , we get the formula for x_{n+1} .

To calculate implied volatility we apply this scheme with x =volatility, $f(x)$ =option pricing formula, and c =observed option price. Note that we have a formula for $f'(x)$: this is the vega of the option. Let's do an example. Suppose the forward price now is $F_0 = 57.7$. Consider a call with strike $K = 55$ and maturity $T = 1$. Assume $r = 0$. If the market value of the call is 7.6273, what is the implied volatility?¹

- Let's start with a guess of $\sigma = .2$. We use our formulas to get a call value of 5.9664 and vega of 21.7290. So our estimate of σ was too low. Newton's method gives the revised estimate $.2 + (7.6273 - 5.9664)/21.7290 = 0.2764$.
- If $\sigma = .2764$ our valuation formula gives a call value of 7.6361 and a vega of 21.9282. Another iteration of Newton's method gives the revised estimate $.2764 + (7.6273 - 7.6361)/21.9282 = 0.2760$.
- If $\sigma = .2760$, our valuation formula gives a call value of 7.6273. So we're done: the implied volatility is .2760.

¹This example was corrected 10/18/2012.