Black-Scholes and the Volatility Surface

When we studied discrete-time models we used martingale pricing to derive the Black-Scholes formula for European options. It was clear, however, that we could also have used a replicating strategy argument to derive the formula. In this part of the course, we will use the replicating strategy argument in continuous time to derive the Black-Scholes *partial differential equation*. We will use this PDE and the Feynman-Kac equation to demonstrate that the price we obtain from the replicating strategy argument is consistent with martingale pricing.

We will also discuss the weaknesses of the Black-Scholes model, i.e. geometric Brownian motion, and this leads us naturally to the concept of the volatility surface which we will describe in some detail. We will also derive and study the Black-Scholes Greeks and discuss how they are used in practice to hedge option portfolios. We will also derive Black's formula which emphasizes the role of the forward when pricing European options. Finally, we will discuss the pricing of other derivative securities and which securities can be priced uniquely given the volatility surface. Change of numeraire / measure methods will also be demonstrated to price exchange options.

1 The Black-Scholes PDE

We now derive the Black-Scholes PDE for a call-option on a non-dividend paying stock with strike K and maturity T. We assume that the stock price follows a geometric Brownian motion so that

$$dS_t = \mu S_t \, dt \; + \; \sigma S_t \, dW_t \tag{1}$$

where W_t is a standard Brownian motion. We also assume that interest rates are constant so that \$1 invested in the cash account at time 0 will be worth $B_t := \$\exp(rt)$ at time t. We will denote by $C(S,t)$ the value of the call option at time t. By Itô's lemma we know that

$$dC(S,t) = \left(\mu S_t \frac{\partial C}{\partial S} + \frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2}\right) \, dt \; + \; \sigma S_t \frac{\partial C}{\partial S} \, dW_t \tag{2}$$

Let us now consider a *self-financing* trading strategy where at each time t we hold x_t units of the cash account and y_t units of the stock. Then P_t, the time t value of this strategy satisfies

$$P_t = x_t B_t + y_t S_t. \tag{3}$$

We will choose x_t and y_t in such a way that the strategy replicates the value of the option. The self-financing assumption implies that

$$\begin{aligned} dP_t &= x_t \, dB_t \; + \; y_t \, dS_t && (4)\\ &= r x_t B_t \, dt + y_t \left(\mu S_t \, dt \; + \; \sigma S_t \, dB_t\right) \\ &= \left(r X_t B_t + y_t \mu S_t\right) \, dt \; + \; y_t \sigma S_t \, dW_t. && (5) \end{aligned}$$

Note that (4) is consistent with our earlier definition[1] of self-financing. In particular, any gains or losses on the portfolio are due entirely to gains or losses in the underlying securities, i.e. the cash-account and stock, and **not** due to changes in the holdings x_t and y_t.

[1] It is also worth pointing out that the mathematical definition of self-financing is obtained by applying Itô's Lemma to (3) and setting the result equal to the right-hand-side of (4).

Returning to our derivation, we can equate terms in (2) with the corresponding terms in (5) to obtain

$$
\begin{aligned}
y_t &= \frac{\partial C}{\partial S} && (6)\\
r x_t B_t &= \frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2}. && (7)
\end{aligned}
$$

If we set $C_0 = P_0$, the initial value of our self-financing strategy, then it must be the case that $C_t = P_t$ for all t since C and P have the same dynamics. This is true by construction after we equated terms in (2) with the corresponding terms in (5). Substituting (6) and (7) into (3) we obtain

$$
r S_t \frac{\partial C}{\partial S} + \frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} - rC = 0, \tag{8}
$$

the **Black-Scholes** PDE. In order to solve (8) boundary conditions must also be provided. In the case of our call option those conditions are: $C(S,T) = \max(S-K,0)$, $C(0,t) = 0$ for all t and $C(S,t) \to S$ as $S \to \infty$.

The solution to (8) in the case of a call option is

$$
C(S,t) = S_t \Phi(d_1) - e^{-r(T-t)} K \Phi(d_2) \tag{9}
$$

$$
\begin{aligned}
\text{where } d_1 &= \frac{\log\left(\frac{S_t}{K}\right) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\\
\text{and } d_2 &= d_1 - \sigma\sqrt{T-t}
\end{aligned}
$$

and $\Phi(\cdot)$ is the CDF of the standard normal distribution. One way to confirm (9) is to compute the various partial derivatives, then substitute them into (8) and check that (8) holds. The price of a European put-option can also now be easily computed from put-call parity and (9).

The most interesting feature of the Black-Scholes PDE (8) is that μ does not appear[2] anywhere. Note that the Black-Scholes PDE would also hold if we had assumed that $\mu = r$. However, if $\mu = r$ then investors would not demand a premium for holding the stock. Since this would generally only hold if investors were risk-neutral, this method of derivatives pricing came to be known as *risk-neutral pricing*.

---

2 Martingale Pricing

We can easily see that the Black-Scholes PDE in (8) is consistent with martingale pricing. Martingale pricing theory states that deflated security prices are martingales. If we deflate by the cash account, then the deflated stock price process, Y_t say, satisfies $Y_t := S_t/B_t$. Then Itô's Lemma and Girsanov's Theorem imply

$$
\begin{aligned}
dY_t &= (\mu - r) Y_t \, dt + \sigma Y_t \, dW_t\\
&= (\mu - r) Y_t \, dt + \sigma Y_t \, (dW_t^{\mathcal{Q}} - \eta_t \, dt)\\
&= (\mu - r - \sigma\eta_t) Y_t \, dt + \sigma Y_t \, dW_t^{\mathcal{Q}}.
\end{aligned}
$$

where $\mathcal{Q}$ denotes a new probability measure and $W_t^{\mathcal{Q}}$ is a $\mathcal{Q}$-Brownian motion. But we know from martingale pricing that if $\mathcal{Q}$ is an equivalent martingale measure then it must be the case that Y_t is a martingale. This then implies that $\eta_t = (\mu - r)/\sigma$ for all t. It also implies that the dynamics of S_t satisfy

$$
\begin{aligned}
dS_t &= \mu S_t \, dt + \sigma S_t \, dW_t\\
&= r S_t \, dt + \sigma S_t \, dW_t^{\mathcal{Q}}. && (10)
\end{aligned}
$$

Using (10), we can now derive (9) using martingale pricing. In particular, we have

$$
C(S_t, t) = \mathbf{E}_t^{\mathcal{Q}}\left[e^{-r(T-t)} \max(S_T - K, 0)\right] \tag{11}
$$

[2]The discrete-time counterpart to this observation was when we observed that the true probabilities of up-moves and down-moves did not have an impact on option prices.

where

$$S_T \;=\; S_t e^{(r-\sigma^2/2)(T-t)+\sigma W_T^{\mathcal{Q}}}$$

is log-normally distributed. While the calculations are a little tedious, it is straightforward to solve (11) and obtain (9) as the solution.

Dividends

If we assume that the stock pays a continuous dividend yield of q, i.e. the dividend paid over the interval $(t,\ t+dt]$ equals $qS_t dt$, then the dynamics of the stock price satisfy

$$dS_t \;=\; (r-q)S_t\, dt \;+\; \sigma S_t\, dW_t^{\mathcal{Q}}. \tag{12}$$

In this case the *total gain process*, i.e. the capital gain or loss from holding the security plus accumulated dividends, is a $\mathcal{Q}$-martingale. The call option price is still given by (11) but now with

$$S_T \;=\; S_t e^{(r-q-\sigma^2/2)(T-t)+\sigma W_T^{\mathcal{Q}}}.$$

In this case the call option price is given by

$$C(S,t) \;=\; e^{-q(T-t)}S_t\Phi(d_1) \;-\; e^{-r(T-t)}K\Phi(d_2) \tag{13}$$

$$\text{where}\quad d_1 \;=\; \frac{\log\left(\frac{S_t}{K}\right)+(r-q+\sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}$$

$$\text{and}\quad d_2 \;=\; d_1-\sigma\sqrt{T-t}.$$

Exercise 1 *Follow the argument of the previous section to derive the Black-Scholes PDE when the stock pays a continuous dividend yield of q.*

Feynman-Kac

We have already seen that the Black-Scholes formula can be derived from either the martingale pricing approach or the replicating strategy / risk neutral PDE approach. In fact we can go directly from the Black-Scholes PDE to the martingale pricing equation of (11) using the Feynman-Kac formula.

Exercise 2 *Derive the same PDE as in Exercise 1 but this time by using (12) and applying the Feynman-Kac formula to an analogous expression to (11).*

While the original derivation of the Black-Scholes formula was based on the PDE approach, the martingale pricing approach is more general and often more amenable to computation. For example, numerical methods for solving PDEs are usually too slow if the number of dimensions are greater[3] than three. Monte-Carlo methods can be used to evaluate (11) regardless of the number of state variables, however, as long as we can simulate from the relevant probability distributions. The martingale pricing approach can also be used for problems that are non-Markovian. This is not the case for the PDE approach.

The Black-Scholes Model is Complete

It is worth mentioning that the Black-Scholes model is a complete model and so every derivative security is attainable or replicable. In particular, this means that every security can be priced uniquely. Completeness follows from the fact that the EMM in (10) is unique: the only possible choice for η_t was $\eta_t = (\mu - r)/\sigma$.

[3] The Black-Scholes PDE is only two-dimensional with state variable t and S. The Black-Scholes PDE is therefore easy to solve numerically.

3 The Volatility Surface

The Black-Scholes model is an elegant model but it does not perform very well in practice. For example, it is well known that stock prices jump on occasions and do not always move in the smooth manner predicted by the GBM motion model. Stock prices also tend to have fatter tails than those predicted by GBM. Finally, if the Black-Scholes model were correct then we should have a flat *implied volatility surface*. The volatility surface is a function of strike, K, and time-to-maturity, T, and is defined implicitly

$$C(S,0) := \mathbf{BS}\left(S, T, r, q, K, \sigma(K,T)\right) \tag{14}$$

where $C(S,K,T)$ denotes the current market price of a call option with time-to-maturity T and strike K, and $\mathbf{BS}(\cdot)$ is the Black-Scholes formula for pricing a call option. In other words, $\sigma(K,T)$ is the volatility that, when substituted into the Black-Scholes formula, gives the market price, $C(S,K,T)$. Because the Black-Scholes formula is continuous and increasing in σ, there will always[4] be a unique solution, $\sigma(K,T)$. If the Black-Scholes model were correct then the volatility surface would be flat with $\sigma(K,T) = \sigma$ for all K and T. In practice, however, not only is the volatility surface not flat but it actually varies, often significantly, with time.

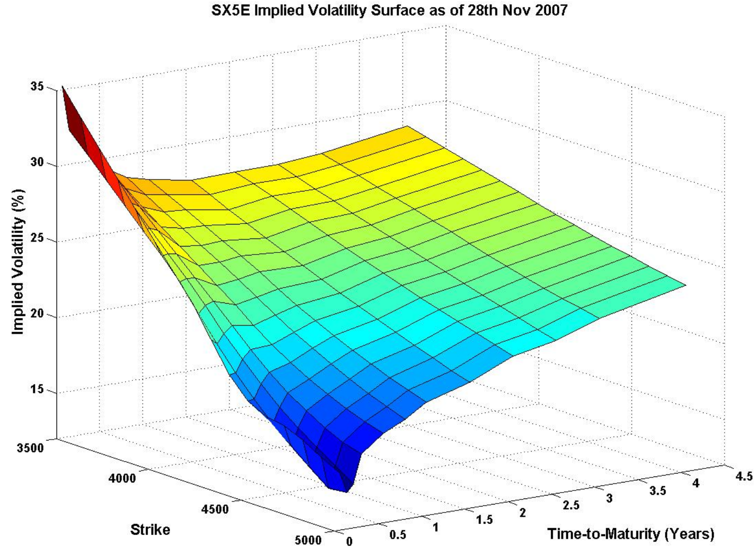

Figure 1: The Volatility Surface

In Figure 1 above we see a snapshot of the[5] volatility surface for the Eurostoxx 50 index on November 28^{th}, 2007. The principal features of the volatility surface is that options with lower strikes tend to have higher implied volatilities. For a given maturity, T, this feature is typically referred to as the volatility **skew** or **smile**. For a given strike, K, the implied volatility can be either increasing or decreasing with time-to-maturity. In general, however, $\sigma(K,T)$ tends to converge to a constant as $T \to \infty$. For T small, however, we often observe an inverted volatility surface with short-term options having much higher volatilities than longer-term options. This is particularly true in times of market stress.

It is worth pointing out that different implementations[6] of Black-Scholes will result in different implied volatility surfaces. If the implementations are correct, however, then we would expect the volatility surfaces to be very

[4] Assuming there is no arbitrage in the market-place.

[5] Note that by put-call parity the implied volatility $\sigma(K,T)$ for a given European call option will be also be the implied volatility for a European put option of the same strike and maturity. Hence we can talk about "the" implied volatility surface.

[6] For example different methods of handling dividends would result in different implementations.

similar in shape. Single-stock options are generally American and in this case, put and call options will typically give rise to different surfaces. Note that put-call parity does not apply for American options.

Clearly then the Black-Scholes model is far from accurate and market participants are well aware of this. However, the language of Black-Scholes is pervasive. Every trading desk computes the Black-Scholes implied volatility surface and the Greeks they compute and use are Black-Scholes Greeks.

Arbitrage Constraints on the Volatility Surface

The shape of the implied volatility surface is constrained by the absence of arbitrage. In particular:

1. We must have $\sigma(K,T) \geq 0$ for all strikes K and expirations T.
2. At any given maturity, T, the skew cannot be too steep. Otherwise butterfly arbitrages will exist.
3. Likewise the *term structure* of implied volatility cannot be too inverted. Otherwise calendar spread arbitrages will exist.

In practice the implied volatility surface will not violate any of these restrictions as otherwise there would be an arbitrage in the market. These restrictions can be difficult to enforce, however, when we are "bumping" or "stressing" the volatility surface, a task that is commonly performed for risk management purposes.

Why is there a Skew?

For stocks and stock indices the shape of the volatility surface is always changing. There is generally a skew, however, so that for any fixed maturity, T, the implied volatility *decreases* with the strike, K. It is most pronounced at shorter expirations. There are several explanations for the skew:

1. Stocks do not follow GBM with a fixed volatility. Markets often jump and jumps to the downside tend to be larger and more frequent than jumps to the upside.
2. Risk aversion: as markets go down, fear sets in and volatility goes up.
3. Supply and demand. Investors like to protect their portfolio by purchasing out-of-the-money puts. This is another form of risk aversion.
4. The total value of company assets, i.e. debt + equity, is a more natural candidate to follow GBM. If so, then equity volatility should increase as the equity value decreases. This is known as the leverage effect. (See below for further explanation.)

Interestingly, there was little or no skew before the Wall street crash of 1987. So it appears to be the case that it took the market the best part of two decades before it understood that it was pricing options incorrectly.

The Leverage Effect

Let V, E and D denote the total value of a company, the company's equity and the company's debt, respectively. Then the fundamental accounting equations states that

$$V = D + E. \tag{15}$$

Equation (15) is the basis for the classical structural models that are used to price risky debt and credit default swaps. Merton (1970's) recognized that the equity value could be viewed as the value of a call option on V with strike equal to D.

Let ΔV, ΔE and ΔD be the change in values of V, E and D, respectively. Then $V + \Delta V = (E + \Delta E) + (D + \Delta D)$ so that

$$\begin{aligned} \frac{V+\Delta V}{V} &= \frac{E+\Delta E}{V} + \frac{D+\Delta D}{V} \\ &= \frac{E}{V}\left(\frac{E+\Delta E}{E}\right) + \frac{D}{V}\left(\frac{D+\Delta D}{D}\right) \end{aligned} \tag{16}$$

If the equity component is substantial so that the debt is not too risky, then (16) implies

$$\sigma_V \approx \frac{E}{V}\,\sigma_E$$

where σ_V and σ_E are the firm value and equity volatilities, respectively. We therefore have

$$\sigma_E \approx \frac{V}{E}\,\sigma_V. \tag{17}$$

Example 1 (The Leverage Effect)

Suppose, for example, that $V = 1$, $E = .5$ and $\sigma_V = 20\%$. Then (17) implies $\sigma_E \approx 40\%$. Suppose σ_V remains unchanged but that over time the firm loses 20% of its value. Almost all of this loss is borne by equity so that now (17) implies $\sigma_E \approx 53\%$. σ_E has therefore increased despite the fact that σ_V has remained constant. This is due to the **leverage effect** which helps explain the presence of skew in equity option markets. ■

What the Volatility Surface Tells Us

To be clear, we continue to assume that the volatility surface has been constructed from European option prices. Consider a **butterfly** strategy centered at K where you are:

1. long a call option with strike $K - \Delta K$
2. long a call with strike $K + \Delta K$
3. short 2 call options with strike K

The value of the butterfly, B_0, at time $t = 0$, satisfies

$$\begin{aligned}
B_0 &= C(K-\Delta K, T) - 2C(K,T) + C(K+\Delta K, T) \\
&\approx e^{-rT}\ \text{Prob}(K - \Delta K \le S_T \le K + \Delta K) \times \Delta K/2 \\
&\approx e^{-rT}\ f(K,T) \times 2\Delta K \times \Delta K/2 \\
&= e^{-rT}\ f(K,T) \times (\Delta K)^2
\end{aligned}$$

where $f(K,T)$ is the probability density function (PDF) of S_T evaluated at K. We therefore have

$$f(K,T) \quad\approx\quad e^{rT}\ \frac{C(K-\Delta K, T) - 2C(K,T) + C(K+\Delta K, T)}{(\Delta K)^2}. \tag{18}$$

Letting $\Delta K \to 0$ in (18), we obtain

$$f(K,T) = e^{rT}\ \frac{\partial^2 C}{\partial K^2}.$$

The volatility surface therefore gives the **marginal** risk-neutral distribution of the stock price, S_T, for any time, T. It tells us **nothing** about the **joint** distribution of the stock price at multiple times, $T_1, \ldots, T_n$.

This should not be surprising since the volatility surface is constructed from European option prices and the latter only depend on the marginal distributions of S_T.

Example 2 (Same marginals, different joint distributions)

Suppose there are 2 stocks, A and B, and that they are both initially priced at \$1. There are two time periods, T_1 and T_2. The stock prices for $t = T_1$ and $t = T_2$ satisfy

$$\begin{aligned}
S_t^{(A)} &= e^{Z_t^{(A)}} \\
S_t^{(B)} &= e^{Z_t^{(B)}}.
\end{aligned}$$

Suppose now that $Z_{T_1}^{(A)}$ and $Z_{T_2}^{(A)}$ are independent $\mathsf{N}(0,1)$ random variables. Suppose also that $Z_{T_1}^{(B)} = Z_{T_1}^{(A)}$ and that

$$Z_{T_2}^{(B)} = \alpha Z_{T_1}^{(A)} + \sqrt{1-\alpha^2}\, Z_{T_2}^{(A)}.$$

Then each $Z_t^{(A)}$ and $Z_t^{(B)}$ has an $\mathsf{N}(0,1)$ distribution and so each stock has the same log-normal *marginal* distribution at times T_1 and T_2. It therefore follows that options on A and B with the same strike and maturity must have the same price.

But now consider an *at-the-money (ATM) knockout put option* with strike $= 1$, barrier at 1.1 and expiration at T_2. The payoff function is then given by

$$\text{Payoff} \;=\; \max\left(1 - S_{T_2},\; 0\right)\; 1_{\{S_t \leq 1.1 \text{ for all } t:\, 0 \leq t \leq T\}}.$$

Question: Would the knockout option on A have the same price as the knockout on B?

Question: How does your answer depend on α?

Question: What does this say about the ability of the volatility surface to price barrier options? ∎

---

4 The Greeks

We now turn to the sensitivities of the option prices to the various parameters. These sensitivities, or the **Greeks** are usually computed using the Black-Scholes formula, despite the fact that the Black-Scholes model is known to be a poor approximation to reality. But first we return to put-call parity.

Put-Call Parity

Consider a European call option and a European put option, respectively, each with the same strike, K, and maturity T. Assuming a continuous dividend yield, q, then put-call parity states

$$e^{-rT}\, K + \text{ Call Price } \;=\; e^{-qT}\, S + \text{ Put Price.} \tag{19}$$

This of course follows from a simple arbitrage argument and the fact that both sides of (19) equal $\max(S_T, K)$ at time T. Put-call parity is useful for calculating Greeks. For example[7], it implies that Vega(Call) = Vega(Put) and that Gamma(Call) = Gamma(Put). It is also extremely useful for **calibrating dividends** and constructing the volatility surface.

The Greeks

The principal Greeks for European call options are described below. The Greeks for put options can be calculated in the same manner or via put-call parity.

Definition: The **delta** of an option is the sensitivity of the option price to a change in the price of the underlying security.

The delta of a European call option satisfies

$$\text{delta} \;=\; \frac{\partial C}{\partial S} \;=\; e^{-qT}\, \Phi(d_1).$$

This is the usual delta corresponding to a volatility surface that is *sticky-by-strike*. It assumes that as the underlying security moves, the volatility of the option does **not** move. If the volatility of the option did move then the delta would have an additional term of the form vega $\times\ \partial\sigma(K,T)/\partial S$. In this case we would say that the volatility surface was *sticky-by-delta*. Equity markets typically use the sticky-by-strike approach when computing deltas. Foreign exchange markets, on the other hand, tend to use the sticky-by-delta approach. Similar comments apply to gamma as defined below.

[7] See below for definitions of vega and gamma.

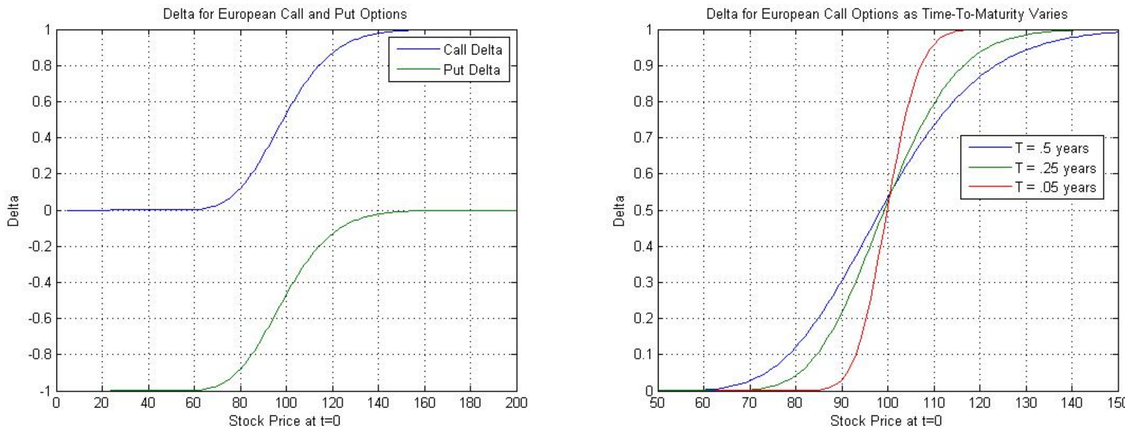

(a) Delta for European Call and Put Options (b) Delta for Call Options as Time-To-Maturity Varies

Figure 2: Delta for European Options

By put-call parity, we have $\text{delta}_{\text{put}} = \text{delta}_{\text{call}} - e^{-qT}$. Figure 2(a) shows the delta for a call and put option, respectively, as a function of the underlying stock price. In Figure 2(b) we show the delta for a call option as a function of the underlying stock price for three different times-to-maturity. It was assumed $r = q = 0$. What is the strike K? Note that the delta becomes steeper around K when time-to-maturity decreases. Note also that delta $= \Phi(d_1) =$ Prob(option expires in the money). (This is only approximately true when r and q non-zero.)

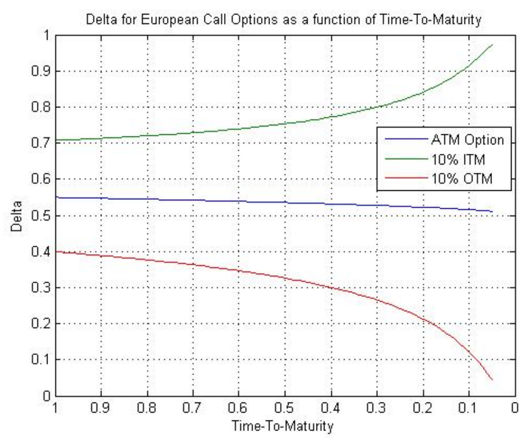

Figure 3: Delta for European Call Options as a Function of Time-To-Maturity

In Figure 3 we show the delta of a call option as a function of time-to-maturity for three options of different *money-ness*. Are there any surprises here? What would the corresponding plot for put options look like?

Definition: The **gamma** of an option is the sensitivity of the option's delta to a change in the price of the underlying security.

The gamma of a call option satisfies

$$\text{gamma} \;=\; \frac{\partial^2 C}{\partial S^2} \;=\; e^{-qT}\,\frac{\phi(d_1)}{\sigma S\sqrt{T}}$$

where $\phi(\cdot)$ is the standard normal PDF.

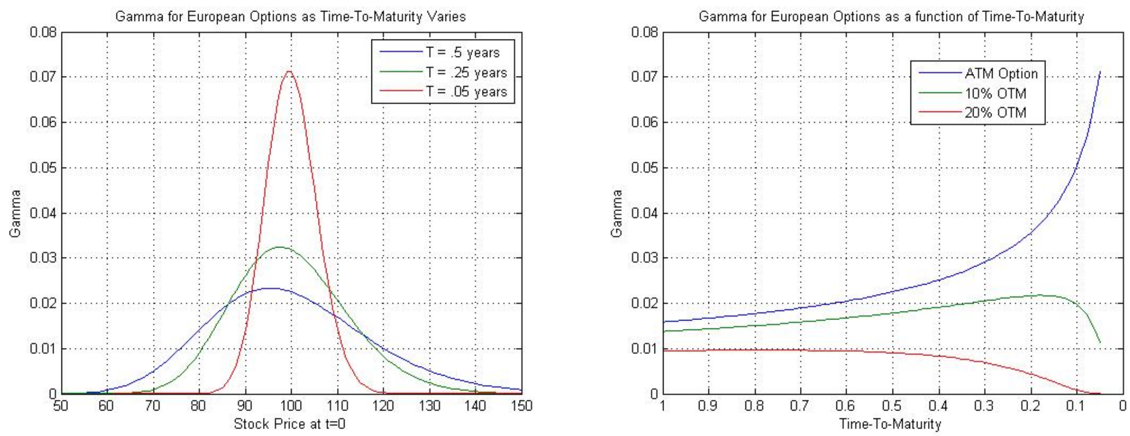

(a) Gamma as a Function of Stock Price (b) Gamma as a Function of Time-to-Maturity

Figure 4: Gamma for European Options

In Figure 4(a) we show the gamma of a European option as a function of stock price for three different time-to-maturities. Note that by put-call parity, the gamma for European call and put options with the same strike are equal. Gamma is always positive due to option **convexity**. Traders who are long gamma can make money by gamma *scalping*. Gamma scalping is the process of regularly re-balancing your options portfolio to be delta-neutral. However, you must pay for this long gamma position up front with the option premium. In Figure 4(b), we display gamma as a function of time-to-maturity. Can you explain the behavior of the three curves in Figure 4(b)?

Definition: The **vega** of an option is the sensitivity of the option price to a change in volatility.

The vega of a call option satisfies

$$\text{vega} \;=\; \frac{\partial C}{\partial \sigma} \;=\; e^{-qT} S\sqrt{T}\;\phi(d_1).$$

In Figure 5(b) we plot vega as a function of the underlying stock price. We assumed $K = 100$ and that $r = q = 0$. Note again that by put-call parity, the vega of a call option equals the vega of a put option with the same strike. Why does vega increase with time-to-maturity? For a given time-to-maturity, why is vega peaked near the strike? Turning to Figure 5(b), note that the vega decreases to 0 as time-to-maturity goes to 0. This is consistent with Figure 5(a). It is also clear from the expression for vega.

Question: Is there any "inconsistency" to talk about vega when we use the Black-Scholes model?

Definition: The **theta** of an option is the sensitivity of the option price to a *negative* change in

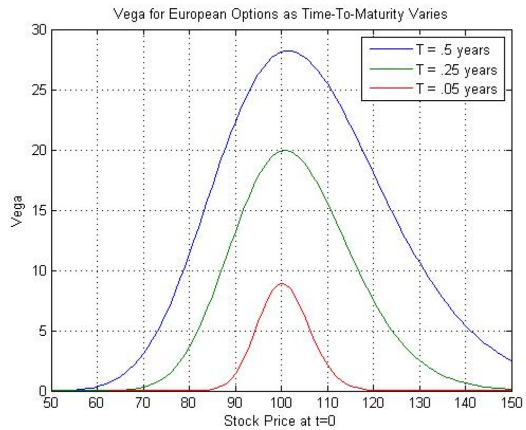

(a) Vega as a Function of Stock Price

(b) Vega as a Function of Time-to-Maturity

Figure 5: Vega for European Options

time-to-maturity.

The theta of a call option satisfies

$$\text{theta} \;=\; -\frac{\partial C}{\partial T} \;=\; -e^{-qT}S\phi(d_1)\frac{\sigma}{2\sqrt{T}} \;+\; qe^{-qT}S\mathsf{N}(d_1) \;-\; rKe^{-rT}\mathsf{N}(d_2).$$

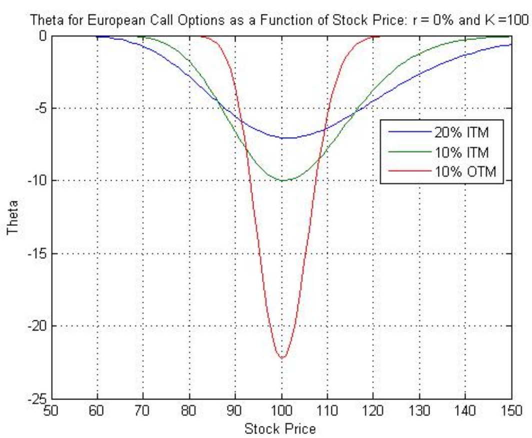

(a) Theta as a Function of Stock Price

(b) Theta as a Function of Time-to-Maturity

Figure 6: Theta for European Options

In Figure 6(a) we plot theta for three call options of different money-ness as a function of the underlying stock price. We have assumed that $r = q = 0\%$. Note that the call option's theta is always negative. Can you explain why this is the case? Why does theta become more negatively peaked as time-to-maturity decreases to 0?

In Figure 6(b) we again plot theta for three call options of different money-ness, but this time as a function of time-to-maturity. Note that the ATM option has the most negative theta and this gets more negative as time-to-maturity goes to 0. Can you explain why?

Options Can Have Positive Theta: In Figure 7 we plot theta for three put options of different money-ness as a function of time-to-maturity. We assume here that $q = 0$ and $r = 10\%$. Note that theta can be positive for in-the-money put options. Why? We can also obtain positive theta for call options when q is large. In typical scenarios, however, theta for both call and put options will be negative.

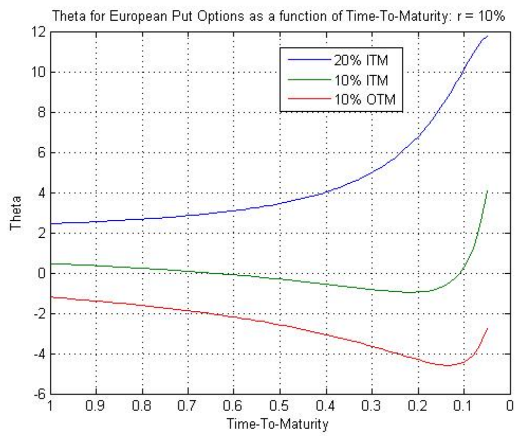

Figure 7: Positive Theta is Possible

The Relationship between Delta, Theta and Gamma

Recall that the Black-Scholes PDE states that any derivative security with price P_t must satisfy

$$\frac{\partial P}{\partial t} + (r - q)S\frac{\partial P}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 P}{\partial S^2} = rP. \tag{20}$$

Writing θ, δ and Γ for theta, delta and gamma, we obtain

$$\theta + (r - q)S\delta + \frac{1}{2}\sigma^2 S^2 \Gamma = rP. \tag{21}$$

Equation (21) holds in general for any portfolio of securities. If the portfolio in question is *delta-hedged* so that the portfolio $\delta = 0$ then we obtain

$$\theta + \frac{1}{2}\sigma^2 S^2 \Gamma = rP \tag{22}$$

It is clear from (22) that any gain from gamma is offset by losses due to theta. This of course assumes that the correct implied volatility is assumed in the Black-Scholes model. Since we know that the Black-Scholes model is wrong, this observation should only be used to help your intuition and not taken as a "fact".

Delta-Gamma-Vega Approximations to Option Prices

A simple application of Taylor's Theorem says

$$\begin{aligned} C(S + \Delta S, \sigma + \Delta\sigma) \quad &\approx \quad C(S,\sigma) + \Delta S\frac{\partial C}{\partial S} + \frac{1}{2}(\Delta S)^2\frac{\partial^2 C}{\partial S^2} + \Delta\sigma\frac{\partial C}{\partial \sigma} \\ = C(S,\sigma) \quad &+ \quad \Delta S \times \delta + \frac{1}{2}(\Delta S)^2 \times \Gamma + \Delta\sigma \times \text{vega}. \end{aligned}$$

where $C(S, \sigma)$ is the price of a derivative security as a function[8] of the current stock price, S, and the implied volatility, σ. We therefore obtain

$$\text{P\&L} \quad = \quad \delta\Delta S + \frac{\Gamma}{2}(\Delta S)^2 + \text{vega } \Delta\sigma$$

[8]The price may also depend on other parameters, in particular time-to-maturity, but we suppress that dependence here.

$$= \text{ delta P\&L } + \text{ gamma P\&L } + \text{ vega P\&L}$$

When $\Delta\sigma = 0$, we obtain the well-known *delta-gamma* approximation. This approximation is often used, for example,in historical **Value-at-Risk** (VaR) calculations for portfolios that include options. We can also write

$$\begin{aligned} \text{P\&L} &= \delta S \left(\frac{\Delta S}{S}\right) + \frac{\Gamma S^2}{2}\left(\frac{\Delta S}{S}\right)^2 + \text{vega } \Delta\sigma \\ &= \text{ESP} \times \text{Return} + \$\text{ Gamma} \times \text{Return}^2 + \text{vega } \Delta\sigma \end{aligned}$$

where ESP denotes the **equivalent stock position** or **"dollar" delta**.

---

5 Delta Hedging

In the Black-Scholes model with GBM, an option can be **replicated** exactly by **delta-hedging** the option. In fact the Black-Scholes PDE we derived earlier was obtained by a delta-hedging / replication argument. The idea behind delta-hedging is to re-balance the portfolio of the option and stock continuously so that you always have a new delta of zero. Of course it is not practical in to hedge continuously and so instead we hedge periodically. Periodic or discrete hedging then results in some *replication error*. Consider Figure 8 below which displays a screen-shot of an Excel spreadsheet that was used to simulate a delta-hedging strategy.

European Option Parameters	
Initial Stock Price	49
Expiration	0.50
Implied Volatility	30.0%
Risk-Free Rate	5.00%
Dividend Yield	10.00%
Call / Put	Call
Strike	50
# Options Sold	100,000
BS Option Price	3.04
Total Option Price	304,132
Equiv. Stock Pos	43,507.9

# Rebalances	20
Monte-Carlo Drift	30.0%
Monte-Carlo Volatility	30%

Simulate Delta Hedging

Time (Weeks)	Stock Price	Delta	Shares Purchased	Cost of Shares Purchased ($000)	Cumulative Cost Includ. Interest & Divs ($000)	Interest Cost ($000)	Dividend Gains ($000)
0.0	49.00	0.435	43,508	2,131.9	2,131.9	2.7	5.3
1.3	46.68	0.349	-8,638	-403.2	1,726.0	2.2	4.1
2.6	49.53	0.456	10,698	529.9	2,253.9	2.8	5.6
3.9	45.18	0.286	-16,961	-766.4	1,484.7	1.9	3.2
5.2	44.88	0.270	-1,618	-72.6	1,410.7	1.8	3.0
6.5	49.15	0.442	17,193	845.0	2,254.4	2.8	5.4
7.8	49.74	0.468	2,593	129.0	2,380.8	3.0	5.8
9.1	50.84	0.518	5,002	254.3	2,632.2	3.3	6.6
10.4	52.31	0.587	6,914	361.7	2,990.6	3.7	7.7
11.7	55.32	0.717	12,998	719.1	3,705.7	4.6	9.9
13.0	58.30	0.824	10,719	624.9	4,325.3	5.4	12.0
14.3	55.71	0.757	-6,716	-374.2	3,944.5	4.9	10.6
15.6	55.40	0.760	302	16.8	3,955.6	4.9	10.5
16.9	53.95	0.713	-4,711	-254.2	3,695.9	4.6	9.6
18.2	52.22	0.634	-7,882	-411.6	3,279.2	4.1	8.3
19.5	52.89	0.691	5,724	302.7	3,577.8	4.5	9.2
20.8	54.88	0.827	13,596	746.2	4,319.3	5.4	11.4
22.1	56.29	0.918	9,051	509.5	4,822.8	6.0	12.9
23.4	61.41	0.994	7,618	467.8	5,283.6	6.6	15.3
24.7	62.98	0.998	359	22.6	5,297.6	6.6	15.7
26.0	65.34	0.000	-99,750	-6,518.0	-1,229.5	-1.5	0.0

Option Payoff ($000)	1,534
Total P&L ($000)	-1

Figure 8: **Delta-Hedging in Excel**

Mechanics of the Excel spreadsheet

In every period, the portfolio is re-balanced so that it is delta-neutral. This is done by using the delta of the options portfolio to determine the total stock position. This stock position is funded through borrowing at the risk-free rate and it accrues dividends according to the dividend yield. The timing of the cash-flows is ignored when calculating the hedging P&L.

Stock prices are simulated assuming $S_t \sim \text{GBM}(\mu, \sigma)$ so that

$$S_{t+\Delta t} = S_t e^{(\mu - \sigma^2/2)\Delta t + \sigma\sqrt{\Delta t} Z}$$

where $Z \sim \mathsf{N}(0,1)$. Note the option implied volatility, σ_{imp}, need not equal σ which in turn need not equal the **realized** volatility. This has interesting implications for the trading P&L and many questions arise.

Question: If you sell options, what typically happens the total P&L if $\sigma < \sigma_{\text{imp}}$?

Question: If you sell options, what typically happens the total P&L if $\sigma > \sigma_{\text{imp}}$?

Question: If $\sigma = \sigma_{\text{imp}}$ what typically happens the total P&L as the number of re-balances increases?

Some Answers to Delta-Hedging Questions

Recall that the price of an option increases as the volatility increases. Therefore if realized volatility is higher than expected, i.e. the level at which it was sold, we expect to lose money on average when we delta-hedge an option that we sold. Similarly, we expect to make money when we delta-hedge if the realized volatility is lower than the level at which it was sold.

In general, however, the payoff from delta-hedging an option is **path-dependent**, i.e. it depends on the price path taken by the stock over the entire time interval. In fact, we can show that the payoff from continuously delta-hedging an option satisfies

$$\text{P\&L} = \int_0^T \frac{S_t^2}{2} \frac{\partial^2 V_t}{\partial S^2} \left(\sigma_{imp}^2 - \sigma_t^2\right) \, dt$$

where V_t is the time t value of the option and σ_t is the realized instantaneous volatility at time t.

The term $\frac{S_t^2}{2}\frac{\partial^2 V_t}{\partial S^2}$ is sometimes called the **dollar gamma**. It is always positive for a call or put option, but it goes to zero as the option moves significantly into or out of the money.

---

6 Pricing Exotics

In this section we will discuss the pricing of three exotic securities: (i) a *digital* option (ii) a *range accrual* and (iii) an *exchange* option. The first two can be priced using the implied volatility surface and so their prices are not *model dependent*. We will price the third security using the Black-Scholes framework. While this is not how it would be priced in practice, it does provide us with an opportunity to practice change-of-measure methods.

Pricing a Digital Option

Suppose we wish to price a digital option which pays \$1 if the time T stock price, S_T, is greater than K and 0 otherwise. Then it is easy[9] to see that the digital price, $D(K,T)$ is given by

$$\begin{aligned} D(K,T) &= \lim_{\Delta K \to 0} \frac{C(K,T) - C(K+\Delta K,T)}{\Delta K} \\ &= -\lim_{\Delta K \to 0} \frac{C(K+\Delta K,T) - C(K,T)}{\Delta K} \\ &= -\frac{\partial C(K,T)}{\partial K}. \end{aligned}$$

In particular this implies that digital options are uniquely priced from the volatility surface. By definition, $C(K,T) = C_{BS}(K,T,\sigma_{BS}(K,T))$ where we use $C_{BS}(\cdot,\cdot,\cdot)$ to denote the Black-Scholes price of a call option as a function of strike, time-to-maturity and volatility. The chain rule now implies

$$D(K,T) = -\frac{\partial C_{BS}(K,T,\sigma_{BS}(K,T))}{\partial K}$$

[9] This proof is an example of a **static replication** argument.

$$
\begin{aligned}
&= -\frac{\partial C_{BS}}{\partial K} - \frac{\partial C_{BS}}{\partial \sigma_{BS}} \frac{\partial \sigma_{BS}}{\partial K} \\
&= -\frac{\partial C_{BS}}{\partial K} - \text{vega} \times \text{skew}.
\end{aligned}
$$

Example 3 **(Pricing a digital)** Suppose $r = q = 0$, $T = 1$ year, $S_0 = 100$ and $K = 100$ so the digital is at-the-money. Suppose also that the skew is 2.5% per 10% change in strike and $\sigma_{atm} = 25\%$. Then

$$
\begin{aligned}
D(100, 1) &= \Phi\left(-\frac{\sigma_{atm}}{2}\right) - S_0\ \phi\left(\frac{\sigma_{atm}}{2}\right) \times \frac{-.025}{.1S_0} \\
&= \Phi\left(-\frac{\sigma_{atm}}{2}\right) + .25\ \phi\left(\frac{\sigma_{atm}}{2}\right) \\
&\approx .45 + .25 \times .4 \\
&= .55
\end{aligned}
$$

Therefore the digital price $= 55\%$ of notional when priced correctly. If we ignored the skew and just the Black-Scholes price using the ATM implied volatility, the price would have been 45% of notional which is significantly less than the correct price. ∎

Exercise 3 *Why does the skew make the digital more expensive in the example above?*

Pricing a Range Accrual

Consider now a 3-month range accrual on the Nikkei 225 index with range $13,000$ to $14,000$. After 3 months the product pays $X\%$ of notional where

$$X = \%\text{ of days over the 3 months that index is inside the range}$$

e.g. If the notional is $\$10M$ and the index is inside the range 70% of the time, then the payoff is $\$7M$.

Question: Is it possible to calculate the price of this range accrual using the volatility surface?

Hint: Consider a portfolio consisting of a pair of digital's for each date between now and the expiration.

Pricing an Exchange Option

Suppose now that there are two *non-dividend-paying* securities with dynamics given by

$$
\begin{aligned}
dY_t &= \mu_y Y_t\ dt\ +\ \sigma_y Y_t\ dW_t^{(y)} \\
dX_t &= \mu_x X_t\ dt\ +\ \sigma_x X_t\ dW_t^{(x)}
\end{aligned}
$$

so that each security follows a GBM. We also assume $dW_t^{(x)} \times dW_t^{(y)} = \rho\ dt$ so that the two security returns have an instantaneous correlation of ρ.

Let $Z_t := Y_t/X_t$. Then Itô's Lemma (check!) implies

$$\frac{dZ_t}{Z_t} = \left(\mu_y - \mu_x - \rho\sigma_x\sigma_y + \sigma_x^2\right)\ dt\ +\ \sigma_y dW_t^{(y)}\ -\ \sigma_x dW_t^{(x)}. \tag{23}$$

The instantaneous variance of dZ/Z is given by

$$
\begin{aligned}
\left(\frac{dZ_t}{Z_t}\right)^2 &= \left(\sigma_y dW_t^{(y)}\ -\ \sigma_x dW_t^{(x)}\right)^2 \\
&= \left(\sigma_x^2 + \sigma_y^2 - 2\rho\sigma_x\sigma_y\right)\ dt
\end{aligned}
$$

Now define a new process, W_t as

$$dW_t = \frac{\sigma_y}{\sigma}\, dW_t^{(y)} \;-\; \frac{\sigma_x}{\sigma}\, dW_t^{(x)}$$

where $\sigma^2 := \left(\sigma_x^2 + \sigma_y^2 - 2\rho\sigma_x\sigma_y\right)$. Then W_t is clearly a continuous martingale. Moreover,

$$\begin{aligned}(dW_t)^2 &= \left(\frac{\sigma_y\, dW_t^{(y)} \;-\; \sigma_x\, dW_t^{(x)}}{\sigma}\right)^2 \\ &= dt.\end{aligned}$$

Hence by Levy's Theorem, W_t is a Brownian motion and so Z_t is a GBM. Using (23) we can write its dynamics as

$$\frac{dZ_t}{Z_t} \;=\; \left(\mu_y - \mu_x - \rho\sigma_x\sigma_y + \sigma_x^2\right)\, dt \;+\; \sigma dW_t. \tag{24}$$

Consider now an exchange option expiring at time T where the payoff is given by

$$\text{Exchange Option Payoff} \;=\; \max\,(0,\, Y_T - X_T)\,.$$

We could use martingale pricing to compute this directly and explicitly solve

$$P_0 \;=\; \mathbf{E}_0^{\mathcal{Q}}\left[e^{-rT}\, \max\,(0,\, Y_T - X_T)\right]$$

for the price of the option. This involves solving a two-dimensional integral with the bivariate normal distribution which is possible but somewhat tedious.

Instead, however, we could price the option by using asset X_t as our numeraire. Let $\mathcal{Q}_x$ be the probability measure associated with this new numeraire. Then martingale pricing implies

$$\begin{aligned}\frac{P_0}{X_0} &= \mathbf{E}_0^{\mathcal{Q}_x}\left[\frac{\max\,(0,\, Y_T - X_T)}{X_T}\right] \\ &= \mathbf{E}_0^{\mathcal{Q}_x}\left[\max\,(0,\, Z_T - 1)\right]. \end{aligned} \tag{25}$$

Equation (24) gives the dynamics of Z_t under our original probability measure (whichever one it was), but we need to know its dynamics under the probability measure $\mathcal{Q}_x$. But this is easy. We know from Girsanov's Theorem that only the drift of Z_t will change so that the volatility will remain unchanged. We also know that Z_t must be a martingale and so under $\mathcal{Q}_x$ this drift must be zero.

But then the right-hand side of (25) is simply the Black-Scholes option price where we set the risk-free rate to zero, the volatility to σ and the strike to 1.

Exercise 4 *If asset X pays a continuous dividend yield of q_x then show that only the strike in (25) needs to be changed.*

Exercise 5 *If asset Y pays a continuous dividend yield of q_y then show that (25) is still valid but that now we must assume Z_t pays the same dividend yield.*

Pricing Other Exotics

Perhaps the two most commonly traded exotic derivatives are barrier options and variance-swaps. In fact at this stage these securities are viewed as more semi-exotic than exotic. As suggested by Example 2, the price of a barrier option cannot be priced using the volatility surface as the latter only defines the marginal distributions of the stock prices. While we could use Black-Scholes and GBM with some constant volatility to determine a price, it is well known that this leads to very inaccurate pricing. Moreover, a rule employed to determine the constant volatility might well lead to arbitrage opportunities for other market participants.

It is generally believed that variance swaps can be priced uniquely from the volatility surface. However, this is only true for variance-swaps with maturities that are less than two or three years. For maturities beyond that, it is probably necessary to include stochastic interest rates and dividends in order to price variance swaps accurately. Variance-swaps will be studied in detail in the assignments.

7 Dividends, the Forward and Black's Model

Let $C = C(S, K, r, q, \sigma, T-t)$ be the price of a call option on a stock. Then the Black-Scholes model says

$$C = Se^{-q(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2)$$

where

$$d_1 = \frac{\log(S/K) + (r - q + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}},$$

$d_2 = d_1 - \sigma\sqrt{T-t}$. Let $F := Se^{(r-q)(T-t)}$ so that F is the time t forward price for delivery of the stock at time T. Then we can write

$$\begin{aligned} C &= Fe^{-r(T-t)}\mathsf{N}(d_1) - Ke^{-r(T-t)}\mathsf{N}(d_2) \qquad (26) \\ &= e^{-r(T-t)} \times \text{Expected-Payoff-of-the-Option} \end{aligned}$$

where

$$d_1 = \frac{\log(F/K) + (T-t)\sigma^2/2}{\sigma\sqrt{T-t}},$$

$d_2 = d_1 - \sigma\sqrt{T-t}$.

Note that the option price now only depends on F, K, r, σ and $T - t$. In fact we can write the call price as

$$C = \mathsf{Black}(F, K, r, \sigma, T-t).$$

where the function Black$(\cdot)$ is defined implicitly by (26). When we write option prices in terms of the forward and not the spot price, the resulting formula is often called Black's formula. It emphasizes the importance of the forward price in establishing the price of the option. The spot price is only relevant in so far as it influences the forward price.

Dividends and Option Pricing

As we have seen, the Black-Scholes formula easily accommodates a continuous dividend yield. In practice, however, dividends are discrete. In order to handle discrete dividends we could convert them into dividend yields but this can create problems. For example, as an ex-dividend date approaches, the dividend yield can grow arbitrarily high. We would also need a different dividend yield for each option maturity. A particularly important problem is that delta and the other Greeks can become distorted when we replace discrete dividends with a continuous dividend yield.

Example 4 **(Discrete dividends)**

Consider a deep in-the-money call option with expiration 1 week from now, a current stock price $= \$100$ and a \$5 dividend going ex-dividend during the week. Then

$$\begin{aligned} \text{Black-Scholes delta} &= e^{-qT}\ \Phi(d_1) \\ &\approx e^{-qT} \\ &= e^{-(.05\times 52)/52} \\ &= 95.12\% \end{aligned}$$

But what do you think the real delta is? ∎

Using a continuous dividend yield can also create major problems when pricing American options. Consider, for example, an American call option with expiration T on a stock that goes ex-dividend on date $t_{div} < T$. This is the only dividend that the stock pays before the option maturity. We know the option should only ever be exercised at either expiration or immediately before t_{div}. However, if we use a continuous dividend yield, the pricing algorithm will never "see" this ex-dividend date and so it will never exercise early, even when it is optimal to do so.

There are many possible solutions to this problem of handling discrete dividends. A common solution is to take $X_0 = S_0 - \text{PV(Dividends)}$ as the "basic" security where

$$\text{PV(Dividends)} = \text{present value of dividends going ex-dividend between now and option expiration.}$$

This works fine for European options (recall that what matters is the forward). For American options, we could, for example, build a binomial lattice for X_t. Then at each date in the lattice, we can determine the stock price and account properly for the discrete dividends, determining correctly whether it is optimal to early exercise or not. In fact this was the subject of a question in an earlier assignment.

---

8 Extensions of Black-Scholes

The Black-Scholes model is easily applied to other securities. In addition to options on stocks and indices, these securities include currency options, options on some commodities and options on index, stock and currency futures. Of course, in all of these cases it is well understood that the model has many weaknesses. As a result, the model has been extended in many ways. These extensions include jump-diffusion models, stochastic volatility models, local volatility models, regime-switching models, garch models and others.

One of the principal uses of the Black-Scholes framework is that is often used to quote derivatives prices via implied volatilities. This is true even for securities where the GBM model is clearly inappropriate. Such securities include, for example, caplets and swaptions in the fixed income markets, CDS options in credit markets and options on variance-swaps in equity markets.