

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

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad (1)$$

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

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

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

$$P_t = x_t B_t + y_t S_t. \quad (3)$$

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

$$dP_t = x_t dB_t + y_t dS_t \quad (4)$$

$$= r x_t B_t dt + y_t (\mu S_t dt + \sigma S_t dW_t)$$

$$= (r x_t B_t + y_t \mu S_t) dt + y_t \sigma S_t dW_t. \quad (5)$$

Note that (4) is consistent with our earlier definition¹ 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 .

¹It 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

$$y_t = \frac{\partial C}{\partial S} \quad (6)$$

$$rx_t B_t = \frac{\partial C}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2}. \quad (7)$$

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

$$rS_t \frac{\partial C}{\partial S} + \frac{\partial C}{\partial t} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} - rC = 0, \quad (8)$$

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 $C(S, t) \rightarrow S$ as $S \rightarrow \infty$.

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

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

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

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

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 appear² 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

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

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

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

²The 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

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

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

$$dS_t = (r - q)S_t dt + \sigma S_t dW_t^Q. \quad (12)$$

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

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

In this case the call option price is given by

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

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

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

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 greater³ 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 $\eta_t = (\mu - r)/\sigma$.

³The 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

$$C(S, K, T) := \mathbf{BS}(S, T, r, q, K, \sigma(K, T)) \quad (14)$$

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 always⁴ 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 the⁵ volatility surface for the Eurostoxx 50 index on November 28th, 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 $T \rightarrow \infty$. 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 implementations⁶ 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

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

⁵Note 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.

⁶For 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. For example fix a maturity, T and consider put two options with strikes $K_1 < K_2$. If there is no arbitrage then it must be the case (why?) that $P(K_1) < P(K_2)$ where $P(K_i)$ is the price of the put option with strike K_i . However, if the skew is too steep then we would obtain (why?) $P(K_1) > P(K_2)$.
3. Likewise the *term structure* of implied volatility cannot be too inverted. Otherwise calendar spread arbitrages will exist. This is most easily seen in the case where $r = q = 0$. Then, fixing a strike K , we can let $C_t(T)$ denote the time t price of a call option with strike K and maturity T . Martingale pricing implies that $C_t(T) = E_t[(S_T - K)^+]$. We have seen before that $(S_T - K)^+$ is a Q -submartingale and now standard martingale results can be used to show that $C_t(T)$ must be non-decreasing in T . This would be violated (Why?) if the term structure of implied volatility was too inverted.

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 two principal explanations for the skew.

1. **Risk aversion** which can appear as an explanation in many guises:
 - (a) 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.
 - (b) As markets go down, fear sets in and volatility goes up.
 - (c) Supply and demand. Investors like to protect their portfolio by purchasing out-of-the-money puts and so there is more demand for options with lower strikes.
2. The **leverage effect** which is due to the fact that 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. To see this consider the following:

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

$$\frac{V + \Delta V}{V} = \frac{E + \Delta E}{V} + \frac{D + \Delta D}{V}$$

$$= \frac{E}{V} \left(\frac{E + \Delta E}{E} \right) + \frac{D}{V} \left(\frac{D + \Delta D}{D} \right) \quad (16)$$

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

$$\sigma_E \approx \frac{V}{E} \sigma_V. \quad (17)$$

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. ■

It is interesting to note that there was little or no skew in the market 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.

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

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

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

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

Letting $\Delta K \rightarrow 0$ 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, \dots, 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 two time periods, T_1 and T_2 , of interest and that a non-dividend paying security has risk-neutral distributions given by

$$S_{T_1} = e^{(r-\sigma^2/2)T_1 + \sigma\sqrt{T_1}Z_1} \quad (19)$$

$$S_{T_2} = e^{(r-\sigma^2/2)T_2 + \sigma\sqrt{T_2}(\rho Z_1 + \sqrt{1-\rho^2}Z_2)} \quad (20)$$

where Z_1 and Z_2 are independent $N(0, 1)$ random variables. Note that a value of $\rho > 0$ can capture a *momentum* effect and a value of $\rho < 0$ can capture a mean-reversion effect. We are also implicitly assuming that $S_0 = 1$.

Suppose now that there are two securities, A and B say, with prices $S_t^{(A)}$ and $S_t^{(B)}$ given by (19) and (20) at times $t = T_1$ and $t = T_2$, and with parameters $\rho = \rho_A$ and $\rho = \rho_B$, respectively. Note that the *marginal* distribution of $S_t^{(A)}$ is identical to the marginal distribution of $S_t^{(B)}$ for $t \in \{T_1, T_2\}$. It therefore follows that options on A and B with the same strike and maturity must have the same price. A and B therefore have identical volatility surfaces.

But now consider a knock-in put option with strike 1 and expiration T_2 . In order to knock-in, the stock price at time T_1 must exceed the barrier price of 1.2. The payoff function is then given by

$$\text{Payoff} = \max(1 - S_{T_2}, 0) 1_{\{S_{T_1} \geq 1.2\}}.$$

Question: Would the knock-in put option on A have the same price as the knock-in put option on B ?

Question: How does your answer depend on ρ_A and ρ_B ?

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

$$e^{-rT}K + \text{Call Price} = e^{-qT}S + \text{Put Price}. \quad (21)$$

This of course follows from a simple arbitrage argument and the fact that both sides of (21) equal $\max(S_T, K)$ at time T . Put-call parity is useful for calculating Greeks. For example⁷, it implies that $\text{Vega}(\text{Call}) = \text{Vega}(\text{Put})$ and that $\text{Gamma}(\text{Call}) = \text{Gamma}(\text{Put})$. 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

⁷See below for definitions of vega and gamma.

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 $\text{vega} \times \partial\sigma(K, T)/\partial S$. 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.

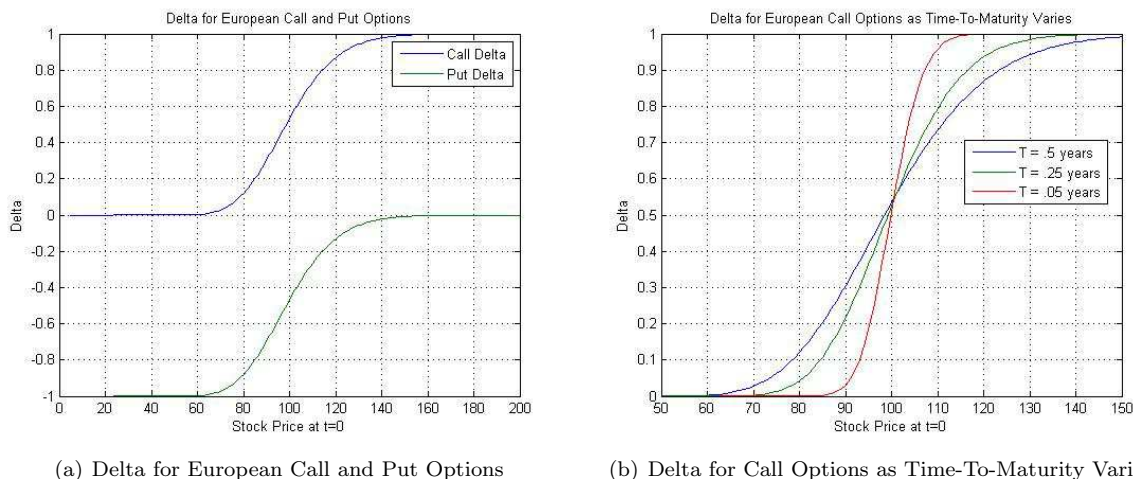


Figure 2: Delta for European Options

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 $\text{delta} = \Phi(d_1) = \text{Prob}(\text{option expires in the money})$. (This is only approximately true when r and q are non-zero.)

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.

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

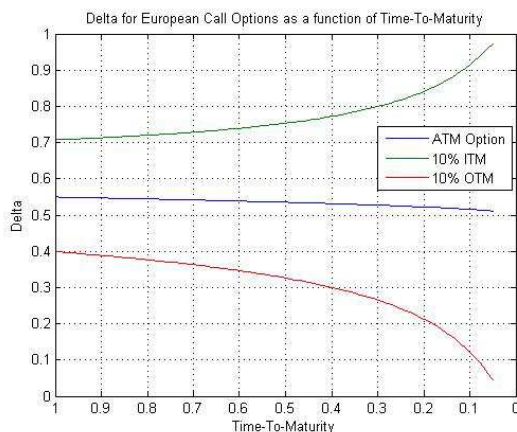


Figure 3: Delta for European Call Options as a Function of Time-To-Maturity

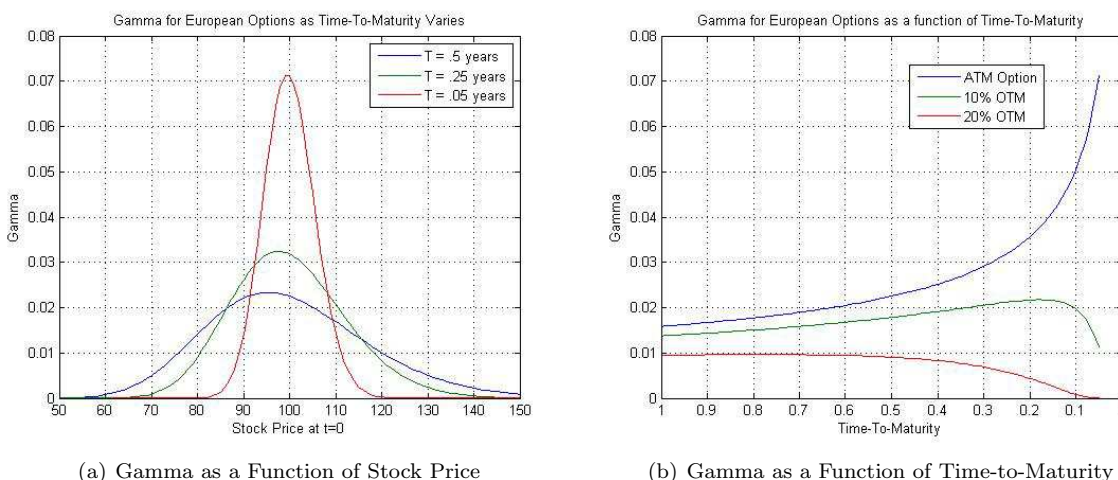


Figure 4: Gamma for European Options

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

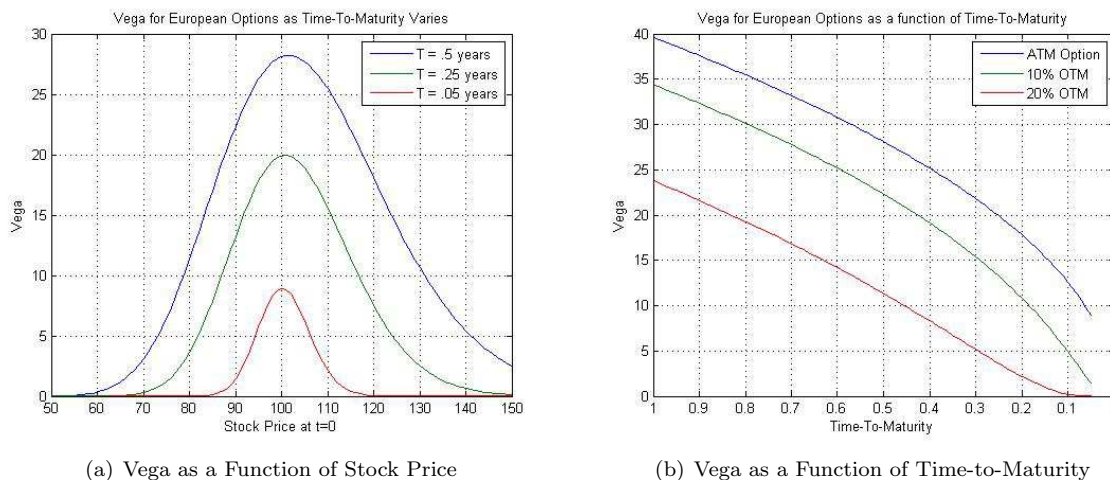


Figure 5: Vega for European Options

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 time-to-maturity.

The theta of a call option satisfies

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




Figure 6: Theta for European Options

In Figure 6(a) we plot theta for three call options of different times-to-maturity 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.

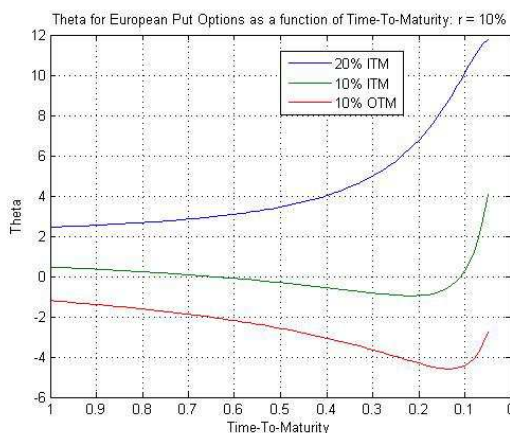


Figure 7: Positive Theta is Possible

The Relationship between Delta, Theta and Gamma

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

$$\frac{\partial P}{\partial t} + (r - q)S \frac{\partial P}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 P}{\partial S^2} = rP. \quad (22)$$

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

$$\theta + (r - q)S\delta + \frac{1}{2} \sigma^2 S^2 \Gamma = rP. \quad (23)$$

Equation (23) 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

$$\theta + \frac{1}{2} \sigma^2 S^2 \Gamma = rP \quad (24)$$

It is clear from (24) 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

$$C(S + \Delta S, \sigma + \Delta \sigma) \approx C(S, \sigma) + \Delta S \frac{\partial C}{\partial S} + \frac{1}{2} (\Delta S)^2 \frac{\partial^2 C}{\partial S^2} + \Delta \sigma \frac{\partial C}{\partial \sigma}$$

$$= C(S, \sigma) + \Delta S \times \delta + \frac{1}{2}(\Delta S)^2 \times \Gamma + \Delta \sigma \times \text{vega}.$$

where $C(S, \sigma)$ is the price of a derivative security as a function⁸ of the current stock price, S , and the implied volatility, σ . We therefore obtain

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

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.

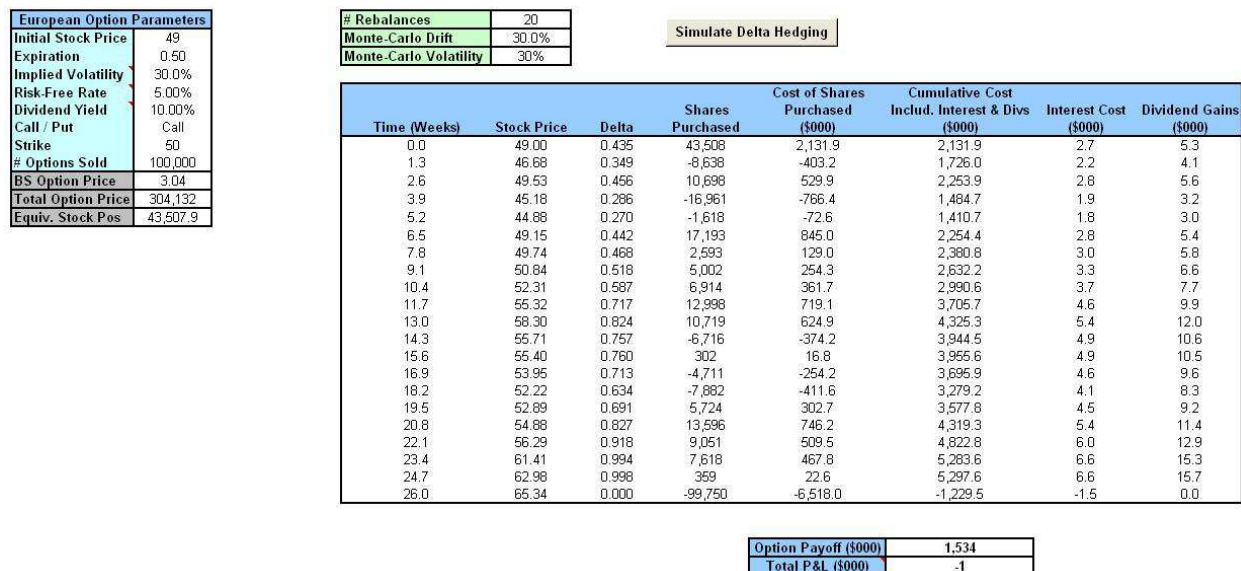


Figure 8: Delta-Hedging in Excel

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

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. Let P_t denote the time t value of the discrete-time self-financing strategy that attempts to replicate the option payoff and let C_0 denote the initial value of the option. The replicating strategy is then given by

$$P_0 := C_0 \quad (25)$$

$$P_{t_{i+1}} = P_{t_i} + (P_{t_i} - \delta_{t_i} S_{t_i}) r \Delta t + \delta_{t_i} (S_{t_{i+1}} - S_{t_i} + q S_{t_i} \Delta t) \quad (26)$$

where $\Delta t := t_{i+1} - t_i$ which we assume is constant for all i , r is the annual risk-free interest rate (assuming per-period compounding) and δ_{t_i} is the Black-Scholes delta at time t_i . This delta is a function of S_{t_i} and some assumed implied volatility, σ_{imp} say. Note that (25) and (26) respect the self-financing condition. 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 $Z \sim \text{N}(0, 1)$. Note the option implied volatility, σ_{imp} , need not equal σ which in turn need not equal the *realized* volatility (when we hedge periodically as opposed to continuously). This has interesting implications for the trading P&L which we may define as

$$\text{P\&L} := P_T - (S_T - K)^+$$

in the case of a short position in a call option with strike K and maturity T . Note that P_T is the terminal value of the replicating strategy in (26). Many interesting questions now arise:

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

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

Question: If $\sigma = \sigma_{imp}$ 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

$$\text{P\&L} = \int_0^T \frac{S_t^2}{2} \frac{\partial^2 V_t}{\partial S^2} (\sigma_{imp}^2 - \sigma_t^2) dt \quad (27)$$

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 often called the **dollar gamma**, as discussed earlier. 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.

Returning to self-financing trading strategy of (25) and (26), note that we can choose any model we like for the security price dynamics. In particular, we are not restricted to choosing geometric Brownian motion and other diffusion or jump-diffusion models could be used instead. It is interesting to simulate these alternative models and to then observe what happens to the replication error in (27) where the δ_{t_i} 's are computed assuming (incorrectly) a geometric Brownian motion price dynamics.

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 easy⁹ to see that the digital price, $D(K, T)$ is given by

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

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

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

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

$$\begin{aligned} D(100, 1) &= \Phi\left(-\frac{\sigma_{atm}}{2}\right) - S_0 \phi\left(\frac{\sigma_{atm}}{2}\right) \times \frac{-0.025}{.1S_0} \\ &= \Phi\left(-\frac{\sigma_{atm}}{2}\right) + .25 \phi\left(\frac{\sigma_{atm}}{2}\right) \\ &\approx .45 + .25 \times .4 \\ &= .55 \end{aligned}$$

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?

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

Example 4 (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. ■

Example 5 (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

$$\frac{dZ_t}{Z_t} = (\mu_y - \mu_x - \rho\sigma_x\sigma_y + \sigma_x^2) dt + \sigma_y dW_t^{(y)} - \sigma_x dW_t^{(x)}. \quad (28)$$

The instantaneous variance of dZ/Z is given by

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

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 $\sigma^2 := (\sigma_x^2 + \sigma_y^2 - 2\rho\sigma_x\sigma_y)$. 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 (28) we can write its dynamics as

$$\frac{dZ_t}{Z_t} = (\mu_y - \mu_x - \rho\sigma_x\sigma_y + \sigma_x^2) dt + \sigma dW_t. \quad (29)$$

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

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

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

$$\begin{aligned} \frac{P_0}{X_0} &= \mathbf{E}_0^{\mathcal{Q}_x} \left[\frac{\max(0, Y_T - X_T)}{X_T} \right] \\ &= \mathbf{E}_0^{\mathcal{Q}_x} [\max(0, Z_T - 1)]. \end{aligned} \quad (30)$$

Equation (29) 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 (30) 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. ■

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 exercises.

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

$$\begin{aligned} C &= Fe^{-r(T-t)}\mathbf{N}(d_1) - Ke^{-r(T-t)}\mathbf{N}(d_2) \\ &= e^{-r(T-t)} \times \text{Expected-Payoff-of-the-Option} \end{aligned} \quad (31)$$

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

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

where the function $\text{Black}(\cdot)$ is defined implicitly by (31). 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 6 (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 $X_0 = S_0 - \text{PV}(\text{Dividends})$ as the "basic" security where

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

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.

Exercises

1. Show that a stock that has a continuous dividend yield of q has risk-neutral dynamics $dS_t = (r - q)S_t dt + \sigma S_t dW_t$ where W_t is a \mathcal{Q} -Brownian motion corresponding to the cash account as the numeraire.
2. Consider a stock that has a continuous dividend yield of q with risk-neutral dynamics $dS_t = (r - q)S_t dt + \sigma S_t dW_t^{\mathcal{Q}}$ where $W_t^{\mathcal{Q}}$ is a \mathcal{Q} -Brownian motion corresponding to the cash account as the numeraire. Show that, as expected,

$$S_0 = \mathbf{E}_0^{\mathcal{Q}} \left[\int_0^T e^{-rt} q S_t dt + e^{-rT} S_T \right]. \quad (32)$$

(You can assume that exchanging the order of integrations in (32) is justified.)

3. Derive the Black-Scholes PDE when the underlying stock has a constant dividend yield of q .
4. Derive the same PDE as in Exercise 3 but this time by using (12) and applying the Feynman-Kac formula to an analogous expression to (11).
5. (a) Use martingale pricing to derive the time t price, $F_t^{(T)}$, of a futures contract for delivery of a stock at time T . You can assume that the Black-Scholes model holds and that the stock pays a dividend yield of q . (You can use your knowledge of futures prices from discrete-time models to justify your answer.)
 (b) Compute the fair price of an option on a futures contract in the Black-Scholes model. You should assume that the futures contract expires at time T and that the option expires at time $\tau < T$. (This is straightforward using the original Black-Scholes formula and your answer from part (a).)
 (c) Confirm directly that the option price you derived in part (b) satisfies the Black-Scholes PDE of Exercise 3.
6. (**Variation on Q3.11 in Back**) Suppose an investor invests in a portfolio with price S and constant dividend yield q . Assume the investor is charged a constant expense ratio α (which acts as a negative dividend) and at date T receives either his portfolio value or his initial investment, whichever is higher. This is similar to a popular type of variable annuity. Letting D denote the number of dollars invested in the contract, the contract pays

$$\max \left(D, \frac{D e^{(q-\alpha)T} S_T}{S_0} \right) \quad (33)$$

at date T . We can rearrange the expression (33) as

$$\begin{aligned} \max \left(D, \frac{D e^{(q-\alpha)T} S_T}{S_0} \right) &= D + \max \left(0, \frac{D e^{(q-\alpha)T} S_T}{S_0} - D \right) \\ &= D + e^{-\alpha T} D \max \left(0, \frac{e^{qT} S_T}{S_0} - e^{\alpha T} \right). \end{aligned}$$

Thus the contract payoff is equivalent to the amount invested plus a certain number of call options written on the gross holding period return $e^{qT} S_T / S_0$. Note that $Z_t := e^{qt} S_t / S_0$ is the date- t value of the portfolio that starts with $1/S_0$ units of the asset (i.e., with a \$1 investment) and reinvests dividends.

Thus, the call options are call options on a non-dividend paying portfolio with the same volatility as S and initial price of \$1.

(a) Compute the date-0 value of the contract to the investor assuming Black-Scholes dynamics for the portfolio. (In practice such a contract would be priced using the appropriate implied volatility surface.)

(b) Create a function (in Matlab, VBA or whatever you prefer) to compute the fair expense ratio; i.e. find α such that the date-0 value of the contract is equal to D . (Hint: You can use $\alpha = 0$ as a lower bound. Because the value of the contract is decreasing as α increases, you can find an upper bound by iterating until the value of the contract is less than D .)

(c) How does the fair expense ratio vary with the maturity, T ?

7. Consider a market with n risky securities with price processes, $(S_t^{(1)}, \dots, S_t^{(n)}) := \mathbf{S}_t$, say. Suppose that the P -dynamics of these securities are driven by m independent Brownian motions so that

$$d\mathbf{S}_t = \mu_t \mathbf{S}_t dt + \Sigma_t d\mathbf{W}_t$$

where \mathbf{W}_t is an $m \times 1$ standard Brownian motion that generates the filtration \mathcal{F}_t , μ_t is an \mathcal{F}_t -adapted $n \times n$ diagonal matrix and Σ_t is an \mathcal{F}_t -adapted $n \times m$ matrix. Assume that there is a cash-account that earns interest at a constant continuously compounded risk-free rate of r . Use Girsanov's Theorem to determine the conditions under which (i) this market is arbitrage-free and (ii) arbitrage-free and complete and (iii) arbitrage-free and incomplete. (Note that there is no loss of generality in assuming that the m Brownian motions are independent. If they were dependent, we could use the Cholesky Decomposition to work with independent Brownian motions. Finally, note that we could also easily allow r to be an adapted process driven by the same set of m Brownian motions.)

8. The current index price is \$100 and the term structure of interest rates is constant at 3%. European call and put option prices of various strikes and maturities are presented below.

	T	.25	.5	1	1.5	
Strike						
60		40.2844	42.4249	50.8521	59.1664	Call Prices
70		30.5281	33.5355	42.6656	51.2181	
80		21.0415	24.9642	34.4358	42.9436	
90		12.2459	16.9652	26.4453	34.7890	
100		5.2025	10.1717	19.4706	27.8938	
110		1.3448	5.4318	14.4225	23.3305	
120		0.2052	2.7647	11.2103	20.7206	
130		0.0216	1.4204	9.1497	19.1828	
140		0.0019	0.7542	7.7410	18.1858	

	T	.25	.5	1	1.5	
Strike						
60		0.0858	2.1546	10.6907	19.3603	Put Prices
70		0.2548	3.1164	12.2087	20.9720	
80		0.6934	4.3962	13.6833	22.2575	
90		1.8232	6.2483	15.3972	23.6629	
100		4.7050	9.3060	18.1270	26.3276	
110		10.7725	14.4171	22.7834	31.3243	
120		19.5582	21.6012	29.2757	38.2744	
130		29.2999	30.1080	36.9195	46.2965	
140		39.2055	39.2929	45.2152	54.8595	

(a) Use put-call parity to determine a piece-wise constant dividend yield implied by the option prices. Does your dividend yield depend on the strikes you choose?

(b) Write a piece of code to determine the Black-Scholes implied volatility for each option and plot the volatility surface.

(In practice, only the bid and ask prices of options are available in the market place and some pre-processing will be necessary to build the volatility surface and calibrate the implied dividends. For example, some options will have very wide bid-offers and are therefore less informative. Moreover, because these options are less liquid it is also the case that these bid-offers may not have been updated as recently as the more liquid options. It is often preferable then to ignore them when building the volatility surface.)

9. Write a computer program that simulates the delta-hedging of a long position in a European option in the Black-Scholes model. Your code should take as inputs the initial stock price S_0 , option expiration T , implied volatility σ_{imp} , risk-free rate r , dividend yield q and strike K as well as whether the option is a call or put. Your code should also take as inputs: (i) the number of re-balancing periods N and (ii) the drift and volatility, μ and σ respectively, of the geometric Brownian motion used to simulate a path of the underlying stock price. Note that σ_{imp} and σ need not be the same. At the very least your code should output the option payoff and the total P&L from holding the option and executing the delta-hedging strategy. (See Figure 8 in the *Black-Scholes and the Volatility Surface* lecture notes for an example where the code was written in VBA with the output in Excel.) Once you have tested your code answer the following questions:

(a) When $\sigma_{imp} = \sigma$ how does the total P&L behave as a function of N ? What happens *on average* if $\sigma_{imp} < \sigma$? If $\sigma_{imp} > \sigma$?

(b) For a fixed N , how does the total P&L behave as $\sigma_{imp} = \sigma$ increases?

(c) How does the drift, μ , affect the total P&L?

(d) Run your code repeatedly for $\sigma_{imp} = 20\%$ and $\sigma = 40\%$ with $S_0 = K = \$50$. Why does the total P&L move about so much? How does the variance of the total P&L depend on the money-ness of the option?

10. Referring to Example 5, suppose asset X pays a continuous dividend yield of q_x . Show that only the strike in (30) needs to be changed in order to obtain the correct option price.

11. Referring again to Example 5, suppose asset Y pays a continuous dividend yield of q_y . Show that (30) is still valid but that we must now assume Z_t pays the same dividend yield.

12. **(Call on the Maximum of 2 Assets)**

Suppose there are two assets with price processes $S_t^{(1)}$ and $S_t^{(2)}$, respectively, that satisfy

$$\begin{aligned} dS_t^{(1)} &= (r - q_1)S_t^{(1)} dt + \sigma_1 S_t^{(1)} dW_t^{(1)} \\ dS_t^{(2)} &= (r - q_2)S_t^{(2)} dt + \sigma_2 S_t^{(2)} dW_t^{(2)} \end{aligned}$$

where $W_t^{(1)}$ and $W_t^{(2)}$ are \mathcal{Q} -Brownian motions with correlation coefficient, ρ . \mathcal{Q} is the EMM corresponding to taking the cash account as numeraire and q_1 and q_2 are the respective dividend yields of the stocks. A call-on-the-max option with strike K and expiration T has a payoff given by

$$\max\left(0, \max\left(S_T^{(1)}, S_T^{(2)}\right) - K\right) = \max\left(0, S_T^{(1)} - K, S_T^{(2)} - K\right).$$

(a) Show that the value of the option at maturity T may be written as

$$xS_T^{(1)} + yS_T^{(2)} - zK$$

where x , y and z are binary random variables taking the values 0 or 1.

(b) By considering numeraires $V_t^{(1)} := e^{q_1 t} S_t^{(1)}$, $V_t^{(2)} := e^{q_2 t} S_t^{(2)}$ and $R_t := e^{rt}$, show that the time 0 value of the option, C say, is given by

$$C = e^{-q_1 T} S_0^{(1)} \text{Prob}^{V_1}(x = 1) + e^{-q_2 T} S_0^{(2)} \text{Prob}^{V_2}(y = 1) - e^{-rT} K \text{Prob}^R(z = 1) \quad (34)$$

where $\text{Prob}^Z(\cdot)$ denotes a risk-neutral probability corresponding to the numeraire Z .

(c) Compute the probabilities in (b) and therefore determine the price of the option.

13. **(Variation on Back Q7.2)** In the Black-Scholes framework, determine the delta-hedging strategy for a call option on a futures contract where the the futures contract is written on a stock that has a continuous dividend yield of q . Without doing any calculations you should be able to tell that as part of the delta hedge you always invest C_t in the cash account at time t where C_t is the time t value of the option. Why is this the case? (This assumes that you use the futures contract and the cash account as your hedging securities rather than the underlying stock and the cash account. The answer to Exercise 5 should be useful in determining the time t position in the underlying futures contract.)

Remark: Here's the way to think about the self-financing condition with futures as the underlying: Let P_t denote the value of the self-financing replicating strategy at time t and assume you hold x_t units of the cash account, B_t , and y_t units of the futures contract at time t . Then

$$P_t = x_t B_t + y_t A_t \quad (35)$$

where A_t is the time t value of the futures contract. (Yes, A_t is in fact identically 0 as we know but let's leave it as A_t for now.) Also let F_t be the time t futures price. If you think about it for just a couple of seconds you'll see that the correct way to view a futures contract is that it is a security that is always worth 0 but that pays out a continuous dividend yield of dF_t . Therefore the self-financing condition applied to (35) yields:

$$dP_t = x_t dB_t + y_t (dA_t + dF_t) \quad (36)$$

which is just the self-financing condition for a dividend paying security. Of course $A_t = 0$ for all t so $dA_t = 0$ so (36) becomes $dP_t = x_t dB_t + y_t dF_t$ which from an economic perspective is clearly correct!

14. **(Idealized Variance-Swap)** A variance-swap with maturity T is a derivative security whose time T payoff is a function of *annualized* realized variance between $t = 0$ and $t = T$. In particular, the purchaser of the variance-swap will receive

$$N_{var} \times (\sigma_{Realized}^2 - \sigma_{Strike}^2) \quad (37)$$

upon expiration at time T where $\sigma_{Realized}^2$ is the annualized realized variance, σ_{Strike}^2 is the strike and N_{var} is the *variance notional* or number of variance units. (The *vega notional* of the variance-swap is defined by $N_{vega} = 2 \times N_{var} \times \sigma_{Strike}$.) The strike, σ_{Strike}^2 , is chosen at time 0 so that the initial value of the variance-swap is zero. Suppose now that the risk-neutral dynamics of a stock price, S_t , satisfy

$$dS_t = (r - q)S_t dt + \sigma_t S_t dW_t \quad (38)$$

where W_t is a \mathcal{Q} -Brownian motion. Then the annualized continuous realized variance is given by

$$\sigma_{Realized}^2 = \frac{1}{T} \int_0^T \sigma_t^2 dt.$$

(a) Show that

$$\frac{1}{T} \int_0^T \sigma_t^2 dt = \frac{2}{T} \left(\int_0^T \frac{dS_t}{S_t} - \ln \left(\frac{S_T}{S_0} \right) \right). \quad (39)$$

Equation (39) implies that the realized leg of a var-swap can be replicated by taking a short position in the *log* contract and by following a dynamic trading strategy that at each time t holds $1/S_t$ shares of the underlying stock.

(b) Use equations (37), (38) and (39) to show that the fair strike of a variance swap satisfies

$$\sigma_{Strike}^2 = \frac{2}{T} \left[(r - q)T - \ln \left(\frac{x}{S_0} \right) - \mathbf{E}_0^Q \left[\ln \left(\frac{S_T}{x} \right) \right] \right]$$

where $x > 0$ is any constant.

(c) Use the mathematical identity

$$-\ln \left(\frac{y}{x} \right) = -\frac{(y-x)}{x} + \int_0^x \frac{1}{k^2} \max(0, k-y) dk + \int_x^\infty \frac{1}{k^2} \max(0, y-k) dk$$

to express the fair strike of the variance-swap in terms of call and put options with expiration T . Simplify your answer by taking x equal to the time T forward value of the stock.

(You have now shown that the fair strike of a variance swap only depends on the prices of vanilla European options with expiration T . Our only model assumptions were (i) that the stock price follow a diffusion as in (38) and that (ii) interest rates were constant. Note that we did **not** assume σ_t in (38) was constant or even deterministic. Of course we did assume that the realized variance was observed continuously. In practice realized variance is accumulated at discrete time intervals, usually via daily or weekly observations. We will address this in the next question.)

15. **(A Variance-Swap in Practice)** In Exercise 14 we considered the continuous-time version of a variance swap. In practice, variance-swap returns are based on discrete (typically daily or weekly) observations. As before, the payoff of a short position in a variance swap satisfies

$$\text{Payoff} = N (K^2 - \sigma_{realized}^2)$$

but now $\sigma_{realized}^2$ is calculated as

$$\sigma_{realized} = 100 \times \sqrt{\frac{A \times \sum_{i=1}^M \left(\ln \frac{S_i}{S_{i-1}} \right)^2}{M}}$$

where M is the number of observation periods and A is the *annualization factor*. For example, if daily returns are used then we typically have $A = 252$. (There are approximately 252 business days in a year. See the sample *term sheet* that is posted on the course web-site.)

(a) Use (i) the approximation $\ln(1+x) \approx x$ for small x and (ii) Taylor's Theorem to show that

$$\sum_{i=1}^M \left(\ln \frac{S_i}{S_{i-1}} \right)^2 \approx 2 \sum_{i=1}^M \frac{1}{S_{i-1}} (S_i - S_{i-1}) - 2 \ln \left(\frac{S_T}{S_0} \right).$$

Show that risk-neutral pricing therefore implies that the fair strike satisfies

$$K^2 \approx 10,000 \times \frac{2}{T} \times \left(rT - \mathbf{E}_0^Q \left[\ln \left(\frac{S_T}{S_0} \right) \right] \right) \quad (40)$$

where r is the risk-free interest rate and $T = M/A$ is the time-to-maturity. (In practice variance-swaps on indices are not "dividend-adjusted" whereas variance-swaps on single stocks are. If a variance-swap is

dividend-adjusted then, assuming a stock goes ex-dividend on date $i - 1$ with dividend d_{i-1} , the variance is calculated with a contribution of $\left(\ln \frac{S_i}{S_{i-1} + d_{i-1}}\right)^2$ instead of $\left(\ln \frac{S_i}{S_{i-1}}\right)^2$. Recall also from Question 4 that the *log* contract in (40) can be replicated using call and put options with maturity T . In particular a variance-swap can be priced using the volatility smile for the maturity of the variance-swap.)

(b) Suppose we want to *mark-to-market* a variance swap that expires at time T and that was initiated at time 0 with a total of M observations and strike K . Assume today is date t and that exactly m observations have already occurred. Let $\sigma_{expected, t}^2 :=$ expected realized variance given the returns up-to time t . Show that

$$\sigma_{expected, t}^2 = \frac{m}{M} \sigma_{0,t}^2 + \frac{M-m}{M} K_{t,T}^2$$

where $K_{t,T}$ is the fair strike at time t for a new variance swap expiring at time T and $\sigma_{0,t}^2$ is the realized variance to date. Hence show that the time t value of the variance-swap, V_t , satisfies

$$\begin{aligned} V_t &= e^{-r(T-t)} N(K^2 - \sigma_{expected, t}^2) \\ &= e^{-r(T-t)} \left[N \frac{m}{M} (K^2 - \sigma_{0,t}^2) + N \frac{(M-m)}{M} (K^2 - K_{t,T}^2) \right] \\ &= \text{Realized P\&L} + \text{Implied P\&L}. \end{aligned}$$

This result is not surprising since it is easily seen that a variance-swap is simply a sum of 1-period variance swaps. Note that this is not true of *volatility* swaps which are therefore much harder to price. (Note also that in practice many-variance-swaps are traded with *caps* on the realized variance. These caps are deep out-of-the-money options on variance.)

(c) Greeks for variance-swaps can be calculated either analytically or by “bumping”, i.e. shifting the parameter by a small amount, recomputing the value of the variance-swap and computing the derivative numerically.

(i) In a Black-Scholes world with a flat volatility surface, does a variance-swap have any delta exposure at the beginning of an observation? Does it have a delta exposure in the middle of an observation?

(ii) What is the daily theta of a variance swap, i.e. the amount you will lose or earn over the next day if the stock price does not move? (You can assume the variance-swap is based on daily observations.)

(iii) Describe at least 2 different methods by which the vega of a variance-swap could be calculated.

(iv) Assuming a constant volatility surface, compute the gamma of a variance-swap at the beginning of an observation. What is the *dollar-gamma* of the variance-swap.

16. Download the Excel spreadsheet *VarSwapUnwind.xls* from the course website where the details of a long variance-swap position can be found. The position was initiated on May 10th 2010 at a strike of 27.2% and with a *vega-notional* of \$200k, i.e. “ $2NK = 200$ ”. The maturity of the variance-swap is May 12th 2011. Today’s date is July 2nd 2010 and the closing price of the underlying index today is 3772.59. The fair value today for a variance-swap that expires on May 12th 2011 is 35.3%. We wish to *unwind* the variance-swap immediately. What is the realized P&L of the variance swap after we unwind it?
17. Use the volatility surface you computed in Exercise 8 to estimate the price of a digital option that pays \$1 in the event that that the stock price in exactly 1 year from now is greater than or equal to \$120. (In order to do this you will need some way of estimating the volatility at non-traded strikes. A convenient way to do this is by fitting a *spline*. This is easy to do in Matlab using just one or two lines of code.)

18. Read the Dividend Swap Primer that can be downloaded from the course web-site. Your goal should be to understand everything in the primer!