IEOR E4602: Quantitative Risk Management — Spring 2016

A Brief Review of Derivatives Pricing & Hedging

In these notes we briefly describe the martingale approach to the pricing of derivatives securities. While most readers are probably more familiar with the dynamic replication approach for pricing derivatives this approach only works in *complete market* settings. It is therefore less general than the martingale approach. Moreover, the approach taken in practice for pricing derivatives has more in common with the martingale approach. After briefly describing the basic ideas of the martingale approach we will apply it to the pricing of call options in the binomial model and discuss how the Black-Scholes price is obtained in the limit as the number of periods goes to infinity.

We will also discuss the weaknesses of the Black-Scholes model, i.e. the geometric Brownian motion model, 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.

1 Martingale Pricing Theory for Multi-Period Models

We assume there are $N+1$ securities, m possible states of nature and that the true probability measure is denoted by $P = (p_1, \ldots, p_m)$ with each $p_i > 0$. We assume that the investment horizon is $[0, T]$ and that there are a total of T trading periods. Securities may therefore be purchased or sold at any date t for $t = 0, 1, \ldots, T-1$. We use $S_t^{(i)}$ to denote the time t price of the i^{th} security for $i = 0, \ldots, N$. Figure 1 below shows a typical multi-period model with $T = 2$ and $m = 9$ possible states. The manner in which information is revealed as time elapses is clear from this model. For example, at node $I_1^{4,5}$ the available information tells us that the true state of the world is either ω_4 or ω_5. In particular, no other state is possible at $I_1^{4,5}$.

Figure 1

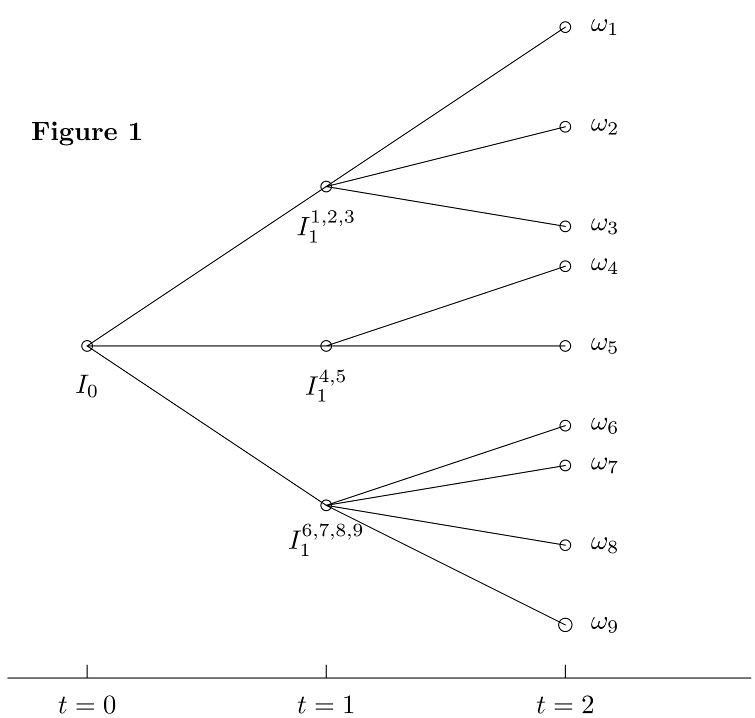

Note that the multi-period model is composed of a series of single-period models. At date $t = 0$ in Figure 1, for example, there is a single one-period model corresponding to node I_0. Similarly at date $t = 1$ there are three possible one-period models corresponding to nodes $I_1^{1,2,3}$, $I_1^{4,5}$ and $I_1^{6,7,8,9}$, respectively. The particular one-period model that prevails at $t = 1$ will depend on the true state of nature. Given a probability measure, $P = (p_1, \ldots, p_m)$, we can easily compute the conditional probabilities of each state. In Figure 1, for example, $\mathbf{P}(\omega_1 | I_1^{1,2,3}) = p_1/(p_1 + p_2 + p_3)$. These conditional probabilities can be interpreted as probabilities in the corresponding single-period models. For example, $p_1 = \mathbf{P}(I_1^{1,2,3} | I_0)\ \mathbf{P}(\omega_1 | I_1^{1,2,3})$.

Trading Strategies and Self-Financing Trading Strategies

We first need to introduce the concepts of a *trading strategy* and a *self-financing trading strategy*. We will assume initially that none of the securities pay "dividends" or have intermediate cash-flows.

Definition 1 *A* **predictable** *stochastic process is a process whose time t value, X_t say, is known at time $t - 1$ given all the information that is available at time $t - 1$.*

Definition 2 *A* **trading strategy** *is a vector, $\theta_t = (\theta_t^{(0)}(\omega), \ldots, \theta_t^{(N)}(\omega))$, of* predictable *stochastic processes that describes the number of units of each security held just* before *trading at time t, as a function of t and ω.*

For example, $\theta_t^{(i)}(\omega)$ is the number of units of the ith security held[1] between times $t - 1$ and t in state ω. We will sometimes write $\theta_t^{(i)}$, omitting the explicit dependence on ω. Note that θ_t is known at date $t - 1$ as we insisted in Definition 2 that θ_t be predictable. In our financial context, 'predictable' means that θ_t cannot depend on information that is not yet available at time $t - 1$.

Example 1 (Constraints Imposed by Predictability of Trading Strategies)

Referring to Figure 1, it must be the case that for all $i = 0, \ldots, N$,

$$\begin{aligned}
\theta_2^{(i)}(\omega_1) &= \theta_2^{(i)}(\omega_2) = \theta_2^{(i)}(\omega_3) \\
\theta_2^{(i)}(\omega_4) &= \theta_2^{(i)}(\omega_5) \\
\theta_2^{(i)}(\omega_6) &= \theta_2^{(i)}(\omega_7) = \theta_2^{(i)}(\omega_8) = \theta_2^{(i)}(\omega_9).
\end{aligned}$$

∎

Exercise 1 *What can you say about the relationship between the $\theta_1^{(i)}(\omega_j)$'s for $j = 1, \ldots, m$?*

Definition 3 *The* **value process**, *$V_t(\theta)$, associated with a trading strategy, θ_t, is defined by*

$$V_t = \begin{cases} \sum_{i=0}^N \theta_1^{(i)} S_0^{(i)} & \text{for } t = 0 \\ \\ \sum_{i=0}^N \theta_t^{(i)} S_t^{(i)} & \text{for } t \geq 1. \end{cases}$$

Definition 4 *A* **self-financing** *(s.f.) trading strategy is a strategy, θ_t, where changes in V_t are due entirely to trading gains or losses, rather than the addition or withdrawal of cash funds. In particular, a self-financing strategy satisfies*

$$V_t = \sum_{i=0}^N \theta_{t+1}^{(i)} S_t^{(i)} \quad \text{for } t = 1, \ldots, T - 1.$$

Definition 4 states that the value of an s.f. portfolio just before trading or *re-balancing* is equal to the value of the portfolio just after trading, i.e., no additional funds have been deposited or withdrawn.

[1] If $\theta_t^{(i)}$ is negative then it corresponds to the number of units sold short.

Exercise 2 *Show that if a trading strategy, θ_t, is s.f. then the corresponding value process, V_t, satisfies*

$$V_{t+1} - V_t \;=\; \sum_{i=0}^{N} \theta_{t+1}^{(i)} \left(S_{t+1}^{(i)} - S_t^{(i)} \right). \tag{1}$$

Exercise 2 states that the changes in the value of the portfolio (that follows a s.f. trading strategy) are due to capital gains or losses and are not due to the injection or withdrawal of funds. Note that we can also write (1) as

$$dV_t = \theta_t^\top \, dS_t,$$

which anticipates the continuous-time definition of an s.f. trading strategy. We can now the concepts of arbitrage, attainable claims and completeness.

Arbitrage

Definition 5 *We define a* **type A arbitrage** *opportunity to be a self-financing trading strategy, θ_t, such that $V_0(\theta) < 0$ and $V_T(\theta) = 0$. Similarly, a* **type B arbitrage** *opportunity is defined to be a self-financing trading strategy, θ_t, such that $V_0(\theta) = 0$, $V_T(\theta) \geq 0$ and $\mathrm{E}_0^P[V_T(\theta)] > 0$.*

Attainability and Complete Markets

Definition 6 *A* **contingent claim**, *C, is a random variable whose value at time T is known at that time given the information available then. It can be interpreted as the time T value of a security.*

Definition 7 *We say that the contingent claim C is* **attainable** *if there exists a self-financing trading strategy, θ_t, whose value process, V_T, satisfies $V_T = C$.*

Note that the value of the claim, C, in Definition 7 must equal the initial value of the replicating portfolio, V_0, if there are no arbitrage opportunities available. We can now define completeness.

Definition 8 *We say that the market is* **complete** *if every contingent claim is attainable. Otherwise the market is said to be* **incomplete**.

We also need the definition of a *numeraire*:

Definition 9 *A* **numeraire security** *is a security with a strictly positive price at all times, t.*

It is often convenient to express the price of a security in units of a chosen numeraire. For example, if the nth security is the numeraire security, then we define

$$\overline{S}_t^{(i)}(\omega_j) \;:=\; \frac{S_t^{(i)}(\omega_j)}{S_t^{(n)}(\omega_j)}$$

to be the date t, state ω_j price (in units of the numeraire security) of the i^{th} security. We say that we are *deflating* by the nth or numeraire security. Note that the deflated price of the numeraire security is always constant and equal to 1.

Definition 10 *The* **cash account** *is a particular security that earns interest at the risk-free rate of interest. In a single period model, the date $t = 1$ value of the cash account is $1 + r$ (assuming that \$1 had been deposited at date $t = 0$), regardless of the terminal state and where r is the one-period interest rate that prevailed at $t = 0$.*

In practice, we often deflate by the cash account *if* it exists. Note that deflating by the cash account is then equivalent to the usual process of discounting. We will use the zeroth security with price process, $S_t^{(0)}$, to denote the cash account whenever it is available. With our definitions of a numeraire security and the cash account remaining unchanged, we can now define what we mean by an **equivalent martingale measure** (EMM), or set of **risk-neutral probabilities**.

Equivalent Martingale Measures (EMMs)

We assume again that we have in mind a specific numeraire security with price process, $S_t^{(n)}$.

Definition 11 *An* **equivalent martingale measure** *(EMM), $Q = (q_1, \ldots, q_m)$, is a set of probabilities such that*

1. *$q_i > 0$ for all $i = 1, \ldots, m$.*
2. *The deflated security prices are martingales. That is*

$$\overline{S}_t^{(i)} \;:=\; \frac{S_t^{(i)}}{S_t^{(n)}} \;=\; \mathrm{E}_t^Q\left[\frac{S_{t+s}^{(i)}}{S_{t+s}^{(n)}}\right] \;=:\; \mathrm{E}_t^Q\left[\overline{S}_{t+s}^{(i)}\right]$$

for $s, t \geq 0$, for all $i = 0, \ldots, N$, and where $\mathrm{E}_t^Q[.]$ denotes the expectation under Q conditional on information available at time t. (We also refer to Q as a set of **risk-neutral probabilities***.)*

1.1 Absence of Arbitrage $\equiv$ Existence of EMM

We begin with two important propositions.

Proposition 1 *If an equivalent martingale measure, Q, exists, then the* deflated *value process, V_t, of any self-financing trading strategy is a Q-martingale.*

Proof: Let θ_t be the self-financing trading strategy and let $\overline{V}_{t+1} := V_{t+1}/S_{t+1}^{(n)}$ denote the deflated value process. We then have

$$\begin{aligned}
\mathrm{E}_t^Q\left[\overline{V}_{t+1}\right] &= \mathrm{E}_t^Q\left[\sum_{i=0}^{N} \theta_{t+1}^{(i)} \overline{S}_{t+1}^{(i)}\right] \\
&= \sum_{i=0}^{N} \theta_{t+1}^{(i)} \mathrm{E}_t^Q\left[\overline{S}_{t+1}^{(i)}\right] \\
&= \sum_{i=0}^{N} \theta_{t+1}^{(i)} \overline{S}_t^{(i)} \\
&= \overline{V}_t.
\end{aligned}$$

We can then apply the tower property of conditional expectations to see that $\mathrm{E}_t^Q\left[\overline{V}_{t+s}\right] = \overline{V}_t$ for any $s \geq 0$. This completes the proof. ∎

Remark 1 *Note that Proposition 1 implies that the deflated price, $\overline{V}_t$, of any attainable security can be computed as the Q-expectation of the terminal deflated value of the security.*

Proposition 2 *If an equivalent martingale measure, Q, exists, then there can be no arbitrage opportunities.*

Proof: The proof follows almost immediately from Proposition 1. ∎

We can now now state the principal result for multi-period models, assuming as usual that a numeraire security exists.

Theorem 3 *(***Fundamental Theorem of Asset Pricing: Part 1***)*

In the multi-period model there is no arbitrage if and only if there exists an EMM, Q.

We will not prove Theorem 3 but note that one direction follows immediately from Proposition 2. The other direction can be proved quite easily using linear programming (LP) duality techniques. We could, for example, use LP duality to first prove the result for one-period models. We could then extend the result to multi-period models by noting that each multi-period model is composed of a series of embedded one-period models and then constructing the multi-period EMM from the various one-period EMMs etc.

1.2 Complete Markets $\equiv$ Existence of a Unique EMM

In this subsection we assume that there is no arbitrage so that an EMM is guaranteed to exist. The following proposition makes intuitive sense and is straightforward to prove.

Proposition 4 *The market is complete if and only if every embedded one-period model is complete.*

We then have the following theorem.

Theorem 5 *(***Fundamental Theorem of Asset Pricing: Part 2***)*

Assume there exists a security with strictly positive price process and that there are no arbitrage opportunities. Then the market is complete if and only if there exists exactly one risk-neutral martingale measure, Q.

As was the case with Theorem 3 we could prove Theorem 5 by first proving it for one-period models and then building up to multi-period models in a straightforward manner.

1.3 Dividends and Intermediate Cash-Flows

Thus far, we have assumed that none of the securities pay intermediate cash-flows. An example of such a security is a dividend-paying stock. This is not an issue in the single period models since any such cash-flows are captured in the date $t = 1$ value of the securities. For multi-period models, however, we sometimes need to explicitly model these intermediate cash payments. All of the results that we have derived in these notes still go through, however, as long as we make suitable adjustments to our price processes and are careful with our bookkeeping. In particular, deflated **cumulative gains** processes rather than deflated security prices are now Q-martingales. The cumulative gain process, G_t, of a security at time t is equal to value of the security at time t plus accumulated cash payments that result from holding the security until time t.

Consider our discrete-time, discrete-space framework where a particular security pays dividends. Then if the model is arbitrage-free there exists an EMM, Q, such that

$$\overline{S}_t \;=\; \mathrm{E}_t^Q\left[\sum_{j=t+1}^{t+s} \overline{D}_j \;+\; \overline{S}_{t+s}\right] \tag{2}$$

where D_j is the time j dividend that you receive if you hold one unit of the security, and S_t is its time t **ex-dividend** price. This result is easy to derive using our earlier results. All we have to do is view each dividend as a separate security with S_t then interpreted as the price of the *portfolio* consisting of these individual securities as well as a security that is worth S_{t+s} at date $t + s$. The definitions of complete and incomplete markets are unchanged and the associated results we derived earlier still hold when we also account for the dividends in the various payoff matrices. For example, if θ_t is a self-financing strategy in a model with dividends then V_t, the corresponding value process, should satisfy

$$V_{t+1} - V_t \;=\; \sum_{i=0}^{N} \theta_{t+1}^{(i)} \left(S_{t+1}^{(i)} + D_{t+1}^{(i)} - S_t^{(i)}\right). \tag{3}$$

Note that the time t dividends, $D_t^{(i)}$, do not appear in (3) since we assume that V_t is the value of the portfolio just after dividends have been paid. This interpretation is consistent with taking S_t to be the time t ex-dividend price of the security.

The various definitions of complete and incomplete markets, state prices, arbitrage etc. are all unchanged when securities can pay dividends. As mentioned earlier, the First Fundamental Theorem of Asset Pricing now states that deflated cumulative gains processes rather than deflated security prices are now Q-martingales. The second fundamental theorem goes through unchanged.

The two fundamental theorems of asset pricing are the cornerstone of derivatives pricing theory. Moreover, given that they hold for multi-period models it should be no surprise that they hold more generally general for continuous-time models although some additional technical assumptions are also required then. The first derivatives pricing models, i.e. the Black-Scholes model and the (equivalent) discrete-time binomial model, were complete-market models where all derivative securities could be dynamically replicated. It should be clear (why?) that the dynamic replication and martingale approaches yield the same price for any derivative security that is replicable. When a derivative security is not replicable – such securities must exist in incomplete market models – a unique arbitrage-free price cannot be computed. (This is simply another way of saying the derivative cannot be replicated by an s.f. trading strategy.) But the second fundamental theorem does tell us nonetheless that, in the absence of arbitrage, there exists infinitely many EMMs that can be used to construct arbitrage-free prices.

1.4 Derivatives Pricing in Practice

Derivatives pricing in practice often proceeds as follows: security price and / or other state variable[2] are modeled directly under an EMM, $Q(\phi)$, where ϕ is some parameter vector. All derivative securities (replicable or not) are then priced using (2) with $Q = Q(\phi)$. Proposition 2 then guarantees that within this model there can be no arbitrage. A specific value of ϕ is chosen by *calibrating* the model. That is we choose ϕ so that the model prices of *liquid* securities coincide with the prices of those securities that we can see in the market-place. In practice we typically cannot perform a perfect calibration so instead we solve a least-squares problem. In particular, we minimize (over ϕ) the (weighted) sum-of-squared differences between model and market prices of liquid securities. The resulting calibrated model can then be used for hedging purposes or to obtain model prices of exotic or less liquid securities.

---

2 The Binomial Model

The binomial model is a discrete-time, discrete space model that describes the price evolution of a single risky stock[3] that does not pay dividends. If the stock price at the beginning of a period is S then it will either increase to uS or decrease to dS at the beginning of the next period. In the model below we have set $S_0 = 100$, $u = 1.06$ and $d = 1/u$.

The binomial model assumes that there is also a cash account available that earns risk-free interest at a gross rate of R per period. We assume R is constant[4] and that the two securities (stock and cash account) may be purchased or sold short. We let $B_k = R^k$ denote the value at time k of \$1 that was invested in the cash account at date 0.

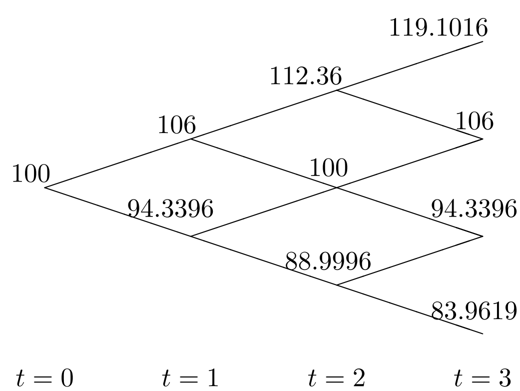

[2]In term-structure models, for example, we typically model the dynamics of interest-rates rather than bond prices. These interest-rates are the *state variables* in these models.

[3]Binomial models are often used to model commodity and foreign exchange prices, as well as dividend-paying stock prices and interest rate dynamics.

[4]This assumption may easily be relaxed.

2.1 Martingale Pricing in the Binomial Model

We have the following result which we will prove using martingale pricing. It is also possible to derive this result using replicating[5] arguments.

Proposition 6 *The binomial model is arbitrage-free if and only if*

$$\boxed{d \;<\; R \;<\; u.} \tag{4}$$

Proof: The first fundamental theorem of asset pricing states that there is no arbitrage in any of the embedded one-period models at time t if and only if there exists a q satisfying $0 < q < 1$ such that

$$\begin{aligned} \frac{S_t}{R^t} &= \mathrm{E}_t^Q\left[\frac{S_{t+1}}{R^{t+1}}\right] \\ &= q\frac{uS_t}{R^{t+1}} \;+\; (1-q)\frac{dS_t}{R^{t+1}}. \end{aligned} \tag{5}$$

Solving (5), we find that $q \;=\; (R-d)/(u-d)$ and $1-q \;=\; (u-R)/(u-d)$. The result now follows since each of the embedded one-period models in the binomial model are identical. ■

Note that the q we obtained in the above Proposition was both unique and node independent. Therefore the binomial model itself is arbitrage-free and complete[6] if (4) is satisfied and we will always assume this to be the case. We will usually use the cash account, B_k, as the numeraire security so that the price of any security can be computed as the discounted expected payoff of the security under Q. Thus the time t price of a security[7] that is worth X_T at time T (and does not provide any cash flows in between) is given by

$$X_t \;=\; B_t \, \mathrm{E}_t^Q\left[\frac{X_T}{B_T}\right] \;=\; \frac{1}{R^{T-t}} \, \mathrm{E}_t^Q[X_T]. \tag{6}$$

The binomial model is one of the workhorses of financial engineering. In addition to being a complete model, it is also *recombining*. For example, an up-move followed by a down-move leads to the same node as a down-move followed by an up-move. This recombining feature implies that the number of nodes in the tree grows linearly with the number of time periods rather than exponentially. This leads to a considerable gain in computational efficiency when it comes to pricing *path-independent* securities.

Example 2 (Pricing a Call Option)

Compute the price of a European call option on the security of Figure 1 with expiration at $T = 3$, and strike $K = 95$. Assume also that $R = 1.02$.

Solution: First, we find $q = \frac{R-d}{u-d} = \frac{1.02-1.06^{-1}}{1.6-1.06^{-1}} = 0.657$ which is the Q-probability of an up-move. If C_0 denotes the date 0 price of the option then (6) implies that it is given by

$$C_0 \;=\; \frac{1}{R^3}\, \mathrm{E}_0^Q[C_T] \;=\; \frac{1}{R^3}\, \mathrm{E}_0^Q[\max(0,\; S_3 - 95)]. \tag{7}$$

At this point, there are two possible ways in which we can proceed:

(i) Compute the Q-probabilities of the terminal nodes and then use (7) to determine C_0. This method does not bother to compute the intermediate prices, C_t.

(ii) Alternatively, we can work backwards in the lattice one period at a time to find C_t at each node and at each time t. This procedure is sometimes referred to as *backwards recursion* and is no more than a simple application of dynamic programming.

---

[5]Indeed, most treatments of the binomial model use replicating arguments to show that (4) must hold in order to rule out arbitrage.

[6]We could also have argued completeness by observing that the matrix of payoffs corresponding to each embedded one-period model has rank 2 which is equal to the number of possible outcomes.

[7]X_t is also the time t value of a self-financing trading strategy that replicates X_T.

Stock Price				European Option Price			
			119.10				24.10
		112.36	106.00			19.22	11.00
	106.00	100.00	94.34		14.76	7.08	0.00
100.00	94.34	89.00	83.96	11.04	4.56	0.00	0.00
t=0	t=1	t=2	t=3	t=0	t=1	t=2	t=3

For example in the *European Option Payoff* table above, we see that $14.76 = \frac{1}{R}\left(q(19.22) + (1-q)(7.08)\right)$, i.e., the value of the option at any node is the discounted expected value of the option one time period ahead. This is just restating the Q-martingale property of discounted security price processes. We find that the call option price at $t = 0$ is given by $\$11.04$. ∎

Example 3 **(A Counter-Intuitive Result)**

Consider the same option-pricing problem of Example 2 except that we now take $R = 1.04$. We then obtain a European call option price of 15.64 as may be seen from the lattices given below. Note that this price is *greater* than the option price, 11.04, that we obtained in Example 2.

Stock Price				European Option Price			
			119.10				24.10
		112.36	106.00			21.01	11.00
	106.00	100.00	94.34		18.19	8.76	0.00
100.00	94.34	89.00	83.96	15.64	6.98	0.00	0.00
t=0	t=1	t=2	t=3	t=0	t=1	t=2	t=3

This observation seems counterintuitive: after all, we are dealing only with positive cash flows, the values of which have not changed, i.e. the option payoffs upon expiration at $t = 3$ have not changed. On the other hand, the interest rate that is used to discount cash flows has increased in this example and so we might have expected the value of the option to have *decreased*. What has happened? (Certainly this situation would never have occurred in a deterministic world!)

First, from a purely mechanical viewpoint we can see that the risk-neutral probabilities have changed. In particular, the risk-neutral probability of an up-move, $q = (R-d)/(u-d)$, has *increased* since R has increased. This means that we are more likely to end up in the higher-payoff states. This increased likelihood of higher payoffs more than offsets the cost of having a larger discount factor and so we ultimately obtain an increase in the option value.

This, however, is only one aspect of the explanation. It is perhaps more interesting to look for an intuitive explanation as to why q should increase when R increases. You should think about this! ∎

2.2 Calibrating the Binomial Model

In continuous-time models, it is often assumed that security price processes are *geometric Brownian motions*. In that case we write $S_t \sim \text{GBM}(\mu, \sigma)$ if

$$S_{t+s} = S_t \, e^{(\mu - \sigma^2/2)s \, + \, \sigma(B_{t+s} - B_t)} \tag{8}$$

where B_t is a standard Brownian motion. Note that this model has the nice property that the gross return, $R_{t,t+s}$, in any period, $[t, t+s]$, is independent of returns in earlier periods. In particular, it is independent of S_t. This follows by noting

$$R_{t,t+s} = \frac{S_{t+s}}{S_t} = e^{(\mu-\sigma^2/2)s \, + \, \sigma(B_{t+s}-B_t)}$$

and recalling the independent increments property of Brownian motion. It is appealing[8] that $R_{t,t+s}$ is independent of S_t since it models real world markets where investors care only about returns and not the absolute price level of securities. The binomial model has similar properties since the gross return in any period of the binomial model is either u or d, and this is independent of what has happened in earlier periods.

We often wish to *calibrate* the binomial model so that its dynamics match that of the geometric Brownian motion in (8). To do this we need to choose u, d and p, the real-world probability of an up-move, appropriately. There are many possible ways of doing this, but one of the more common choices[9] is to set

$$p_n = \frac{e^{\mu T/n} - d_n}{u_n - d_n} \tag{9}$$

$$u_n = \exp(\sigma\sqrt{T/n}) \tag{10}$$

$$d_n = 1/u_n = \exp(-\sigma\sqrt{T/n}) \tag{11}$$

where T is the expiration date and n is the number of periods. (This calibration becomes more accurate as n increases.) Note then, for example, that $\mathrm{E}[S_{i+1}|S_i] = p_n u_n S_i + (1-p_n) d_n S_i = S_i \exp(\mu T/n)$, as desired.

We will choose the gross risk-free rate per period, R_n, so that it corresponds to a continuously-compounded rate, r, in continuous time. We therefore have

$$R_n = e^{rT/n}. \tag{12}$$

Remark 2 *Recall that the true probability of an up-move, p, has no bearing upon the risk-neutral probability, q, and therefore it does not* directly *affect how securities are priced. From our calibration of the binomial model, we therefore see that μ, which enters the calibration only through p, does not impact security prices. On the other hand, u and d depend on σ which therefore does impact security prices. This is a recurring theme in derivatives pricing.*

Remark 3 *We just stated that p does not* directly *affect how securities are priced. This means that if p should suddenly change but S_0, R, u and d remain unchanged, then q, and therefore derivative prices, would also remain unchanged. This seems very counter-intuitive but an explanation is easily given. In practice, a change in p would generally cause one or more of S_0, R, u and d to also change. This would in turn cause q, and therefore derivative prices, to change. We could therefore say that p has an* indirect *effect on derivative security prices.*

2.3 Convergence of the Binomial Model to Black-Scholes

The Black-Scholes formula for the price of a call option on a non-dividend paying security with initial price S_0, strike K, time to expiration T, continuously compounded interest rate r, and volatility parameter σ, is given by

$$C(S_0, T) = S_0 N(d_1) - K e^{-rT} N(d_2) \tag{13}$$

where

$$d_1 = \frac{\log(S_0/K) + (r+\sigma^2/2)T}{\sigma\sqrt{T}},$$

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

[8] More sophisticated models will sometimes allow for *return predictability* where $R_{t,t+s}$ is not independent of S_t. Even then, it is still appropriate to model returns rather than absolute security values.

[9] We write p_n, u_n and d_n to emphasize that their values depend explicitly on the number of periods, n, for a fixed expiration, T.

and $N(\cdot)$ is the CDF of a standard Normal random variable.

One way to obtain the Black-Scholes formula is to consider the sequence of binomial models, M_n, that are parameterized by (9), (10), (11) and (12). Letting $n \to \infty$ in this sequence it can be shown via a version of the Central Limit Theorem that the call option price in the binomial model will converge to (13).

---

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

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

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[10] 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.

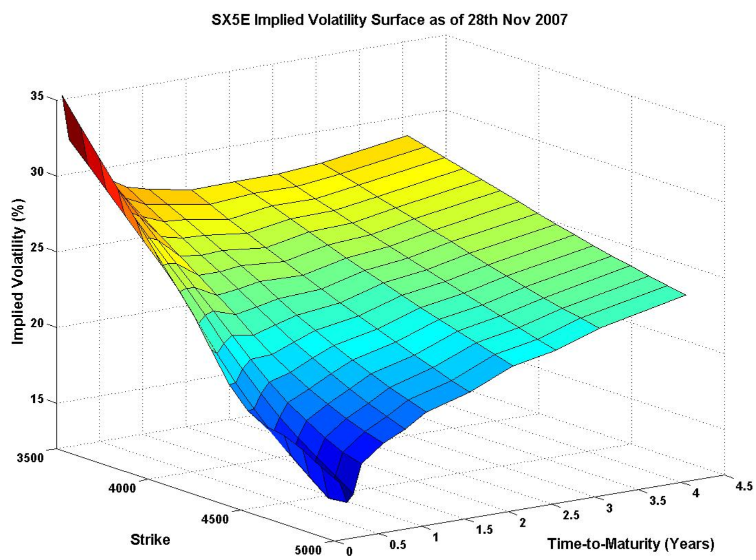

Figure 1: The Volatility Surface

In Figure 1 above we see a snapshot of the[11] 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

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

[11] 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.

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[12] 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 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 $C_t(T) = \mathsf{E}_t[(S_T - K)^+]$. 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

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

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

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

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

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

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

Example 4 (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

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

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

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

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 5 (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

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

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

where Z_1 and Z_2 are independent $\mathsf{N}(0,1)$ 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

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

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

$$e^{-rT}\,K + \text{ Call Price } = e^{-qT}\,S + \text{ Put Price.} \tag{21}$$

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 example[13], it implies that Vega(Call) = Vega(Put) and that Gamma(Call) = Gamma(Put). It is also extremely useful for **calibrating**

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

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 $\times\ \partial\sigma(K,T)/\partial S$.

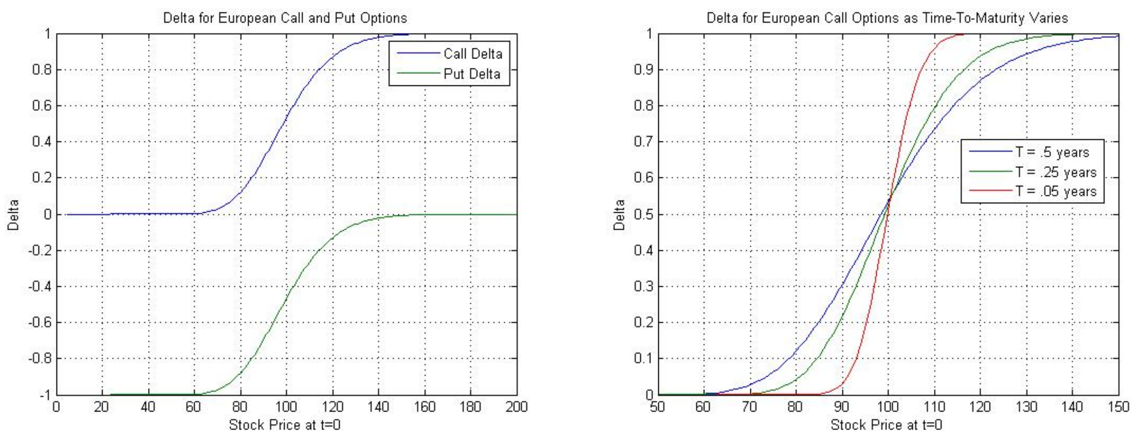

(a) Delta for European Call and Put Options

(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 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

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

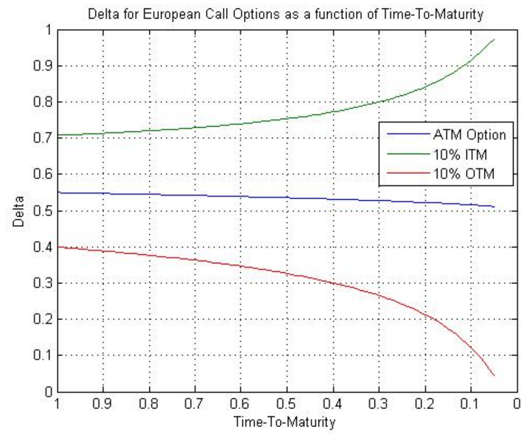

Figure 3: Delta for European Call Options as a Function of Time-To-Maturity

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

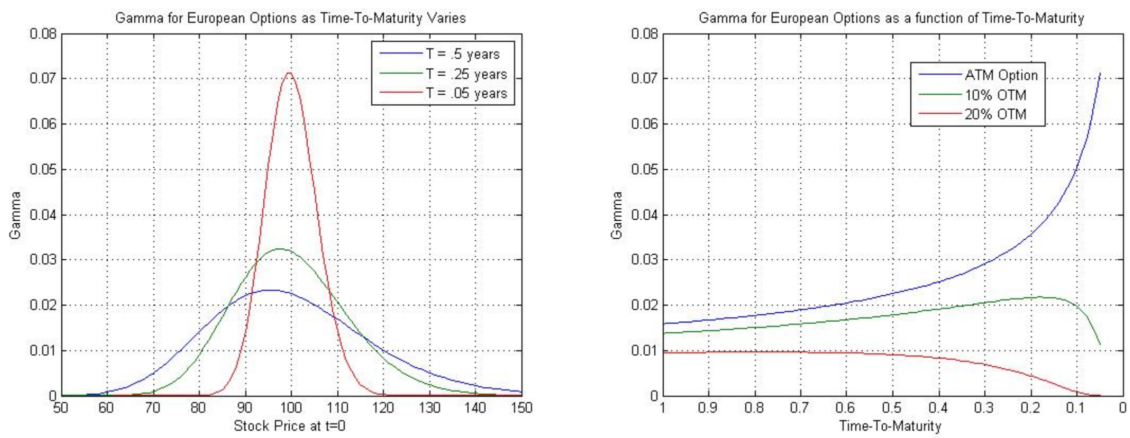

(a) Gamma as a Function of Stock Price

(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

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

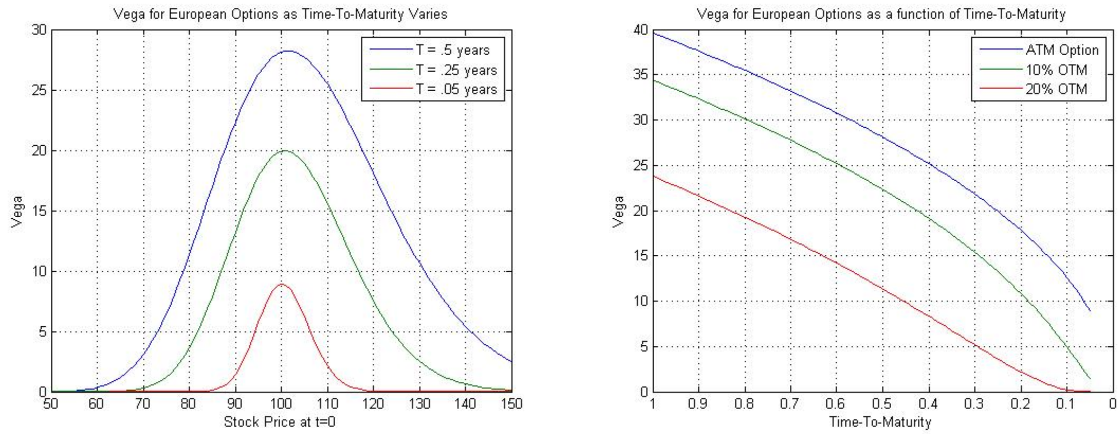

(a) Vega as a Function of Stock Price

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

Figure 5: Vega for European Options

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 time-to-maturity.

The theta of a call option satisfies

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

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.

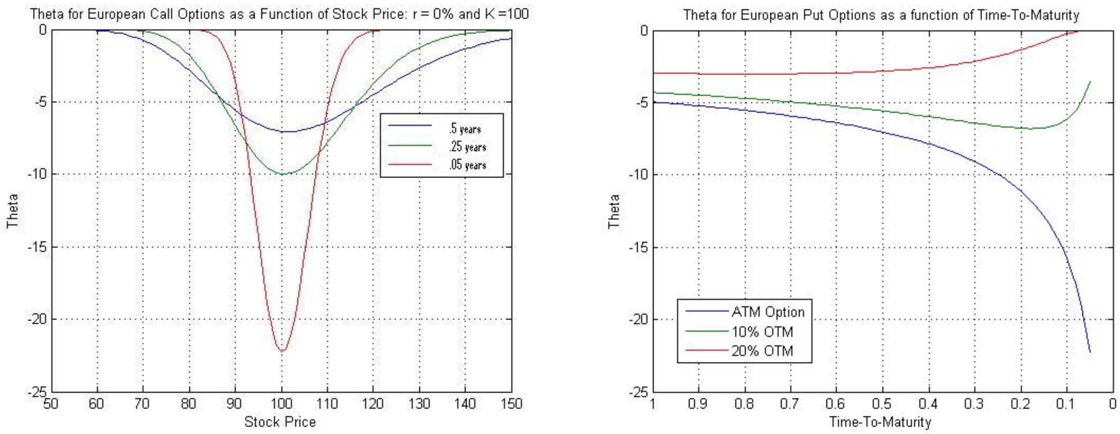

(a) Theta as a Function of Stock Price (b) Theta as a Function of Time-to-Maturity

Figure 6: Theta for European Options

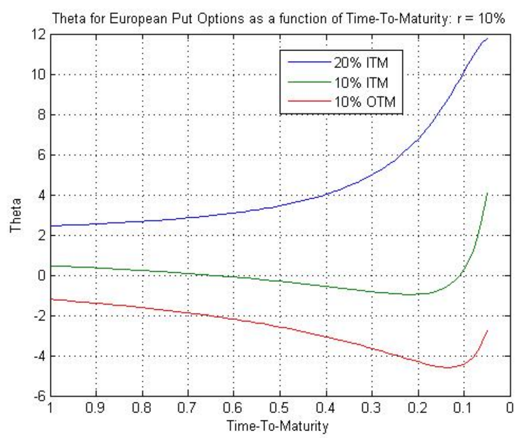

Figure 7: Positive Theta is Possible

The Relationship between Delta, Theta and Gamma

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

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

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

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

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

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

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

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

where $C(S,\sigma)$ is the price of a derivative security as a function[14] 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**[15] exactly via **delta-hedging**. The idea behind delta-hedging is to continuously re-balance an s.f. portfolio that continuously trades the underlying stock and cash account so that the portfolio always has a delta equal to the delta of the option being hedged. Of course in practice we cannot hedge continuously and so instead we hedge periodically. Periodic or discrete hedging then results in some *replication error*.

Delta-hedging proceeds as follows. The portfolio begins with an initial cash value of C_0 which is the time $t = 0$ value of the position or option we want to hedge. Then, and in each subsequent period, the portfolio is re-balanced in such a way that at time t we are long δ_t units of the stock where δ_t is the time t delta of the option. Any remaining cash is invested in the cash account and therefore earns at the risk-free rate of interest. The stock position accrues dividends according to the dividend yield, q. Let P_t denote the time t value of this discrete-time trading strategy. The value of the corresponding portfolio then evolves according to

$$P_0 := C_0 \tag{25}$$

$$P_{t_{i+1}} = P_{t_i} + (P_{t_i} - \delta_{t_i} S_{t_i})\, r\Delta t + \delta_{t_i}\left(S_{t_{i+1}} - S_{t_i} + qS_{t_i}\Delta t\right) \tag{26}$$

where $\Delta t := t_{i+1} - t_i$ which we assume is constant for all i and r is the annual risk-free interest rate (assuming per-period compounding). Note that δ_{t_i} is a function of S_{t_i} and some assumed implied volatility, σ_{imp} say. We also note that (25) and (26) respect the self-financing condition. Stock prices are simulated assuming $S_t \sim \text{GBM}(\mu, \sigma)$ so that

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

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

[15] In fact one way to obtain the Black-Scholes formula is via a replication argument that leads to the so-called Black-Scholes PDE. The solution of this PDE is the Black-Scholes formula.

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 (when we hedge periodically as opposed to continuously). This has interesting implications for the trading P&L which we may define as

$$\begin{aligned}\mathsf{P\&L} &:= P_T - C_T \\ &= P_T - (S_T - K)^+\end{aligned}$$

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 (where the underlying S_t never *jumps*) satisfies

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

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. This is actually the situation in practice – we don't know the true market dynamics but we often hedge using the Black-Scholes model.

---

6 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. Most of these models are incomplete models and the pricing philosophy then proceeds as described in Section 1.4.

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.