ull3b.tex **Week 3: pm, 10.10.2018**

Proof of the Completeness Theorem. $\Rightarrow$: Assume viability and completeness. Then for any $\mathcal{F}_N$-measurable random variable $h \geq 0$, there exists an admissible (so self-financing) strategy H replicating h: $h = V_N(H)$. As H is self-financing, by IV.1

$$h/S_N^0 = \tilde{V}_N(H) = V_0(H) + \Sigma_1^N H_j.\Delta\tilde{S}_j.$$

We know by the No-arbitrage Theorem (IV.2) that an equivalent martingale measure P^* exists; we have to prove uniqueness. So, let P_1, P_2 be two such EMMs. For $i = 1, 2$, $(\tilde{V}_n(H))_{n=0}^N$ is a P_i-martingale. So,

$$E_i[\tilde{V}_N(H)] = E_i[V_0(H)] = V_0(H),$$

since the value at time zero is non-random ($\mathcal{F}_0 = \{\emptyset, \Omega\}$). So

$$E_1[h/S_N^0] = E_2[h/S_N^0].$$

Since h is arbitrary, E_1, E_2 have to agree on integrating all non-negative integrands. Taking negatives and using linearity: they agree on non-positive integrands also. Split into positive and negative parts: they have to agree on all integrands. Now E_i is expectation (i.e., integration) w.r.t. the measure P_i, and measures that agree on integrating all integrands must coincide. So $P_1 = P_2$. That is, EMMs (which *exist*, by the NA Theorem) are *unique*. //

Before proving the converse, we prove a lemma. Recall that an admissible strategy is SF with all values non-negative. The Lemma shows that the non-negativity of contingent claims extends to all values of any SF strategy replicating it – i.e., this gives equivalence of admissible and SF replicating strategies. [SF: isolated from external wealth; admissible: actually worth something. These sound similar; the Lemma shows they are the same here. So we only need one term; we use SF as it is shorter.]

Lemma. In a viable market, any attainable h (i.e., any h that can be replicated by an admissible strategy H) can also be replicated by a self-financing strategy H.

Proof. If H is self-financing and P^* is an equivalent martingale measure under which discounted prices $\tilde{S}$ are P^*-martingales (an EMM – such *exist* by the

NA Theorem, IV.2), $\tilde{V}_n(H)$ is also a P^*-martingale, being the martingale transform of $\tilde{S}$ by H (see IV.1). So

$$\tilde{V}_n(H) = E^*[\tilde{V}_N(H)|\mathcal{F}_n] \qquad (n = 0, 1, \cdots, N).$$

If H replicates h, $V_N(H) = h \geq 0$, so discounting, $\tilde{V}_N(H) \geq 0$, so the above equation gives $\tilde{V}_n(H) \geq 0$ for each n. Thus *all* the values at each time n are non-negative – not just the final value at time N – so H is admissible. //

Proof of the Theorem (continued). $\Leftarrow$ (*not examinable)*: Assume the market is viable but incomplete: then there exists a non-attainable $h \geq 0$. By the Lemma, we may confine attention to self-financing strategies H (which will then automatically be admissible). By the Proposition of IV.1, we may confine attention to the risky assets $S^1, \cdots, S^d$, as these suffice to tell us how to handle the bank account S^0.

Call $\tilde{\mathcal{V}}$ the set of random variables of the form

$$U_0 + \Sigma_1^N H_n.\Delta\tilde{S}_n$$

with U_0 $\mathcal{F}_0$-measurable (i.e. deterministic) and $((H_n^1, \cdots, H_n^d))_{n=0}^N$ predictable; this is a vector space. Then by above, the discounted value h/S_N^0 does not belong to $\tilde{\mathcal{V}}$, so $\tilde{\mathcal{V}}$ is a *proper* subspace of the vector space $\mathbb{R}^\Omega$ of all random variables on Ω. Let P^* be a probability measure equivalent to P under which discounted prices are martingales (an EMM: such exist by the NA Theorem, IV.2). Define the scalar product

$$(X, Y) \mapsto E^*[XY]$$

on random variables on Ω. Since $\tilde{\mathcal{V}}$ is a proper subspace, by Gram-Schmidt orthogonalisation there exists a non-zero random variable X orthogonal to $\tilde{\mathcal{V}}$. That is,

$$E^*[X] = 0.$$

Write $\|X\|_\infty := \max\{|X(\omega)| : \omega \in \Omega\}$, and define P^{**} by

$$P^{**}(\{\omega\}) = \Big(1 + \frac{X(\omega)}{2\|X\|_\infty}\Big) P^*(\{\omega\}).$$

By construction, P^{**} is equivalent to P^* (same null-sets – actually, as $P^* \sim P$ and P has no non-empty null-sets, neither do P^*, P^{**}). As X is non-zero,

P^{**} and P^* are *different*. Now

$$\begin{aligned} E^{**}[\Sigma_1^N H_n.\Delta\tilde{S}_n] &= \Sigma_\omega P^{**}(\omega)\Big(\Sigma_1^N H_n.\Delta\tilde{S}_n\Big)(\omega) \\ &= \Sigma_\omega\Big(1+\frac{X(\omega)}{2\|X\|_\infty}\Big)P^*(\omega)\Big(\Sigma_1^N H_n.\Delta\tilde{S}_n\Big)(\omega). \end{aligned}$$

The '1' term on the right gives $E^*[\Sigma_1^N H_n.\Delta\tilde{S}_n]$, which is zero since this is a martingale transform of the E^*-martingale $\tilde{S}_n$. The 'X' term gives a multiple of the inner product

$$(X, \Sigma_1^N H_n.\Delta\tilde{S}_n),$$

which is zero as X is orthogonal to $\tilde{\mathcal{V}}$ and $\Sigma_1^N H_n.\Delta\tilde{S}_n \in \tilde{\mathcal{V}}$. By the Martingale Transform Lemma, $\tilde{S}_n$ is a P^{**}-martingale since H (previsible) is arbitrary. Thus P^{**} is a second EMM, different from P^*. So incompleteness implies non-uniqueness of equivalent martingale measures. //

§4. The Fundamental Theorem of Asset Pricing; Risk-neutral valuation.

We summarise what we have learned so far. We call a measure P^* under which discounted prices $\tilde{S}_n$ are P^*-martingales a *martingale measure*. Such a P^* equivalent to the true probability measure P is called an *equivalent martingale measure*. Then
1 (**No-Arbitrage Theorem**: IV2). If the market is *viable* (arbitrage-free), equivalent martingale measures P^* *exist*.
2 (**Completeness Theorem**: IV.3). If the market is *complete* (all contingent claims can be replicated), equivalent martingale measures are *unique*. Combining:

Theorem (Fundamental Theorem of Asset Pricing, FTAP). In a complete viable market, there exists a unique equivalent martingale measure P^* (or Q).

Let h (≥ 0, $\mathcal{F}_N$-measurable) be any contingent claim, H an admissible strategy replicating it:

$$V_N(H) = h.$$

As $\tilde{V}_n$ is the martingale transform of the P^*-martingale $\tilde{S}_n$ (by H_n), $\tilde{V}_n$ is a

P^*-martingale. So $V_0(H)(= \tilde{V}_0(H)) = E^*[\tilde{V}_N(H)]$. Writing this out in full:

$$V_0(H) = E^*[h/S_N^0].$$

More generally, the same argument gives $\tilde{V}_n(H) = V_n(H)/S_n^0 = E^*[(h/S_N^0)|\mathcal{F}_n]$:

$$V_n(H) = S_n^0 E^*[\frac{h}{S_N^0}|\mathcal{F}_n] \qquad (n = 0, 1, \cdots, N).$$

It is natural to call $V_0(H)$ above the *value* of the contingent claim h at time 0, and $V_n(H)$ above the value of h at time n. For, if an investor *sells* the claim h at time n for $V_n(H)$, he can follow strategy H to replicate h at time N and clear the claim. To sell the claim for *any other amount* would provide an arbitrage opportunity (as with the argument for put-call parity). So this value $V_n(H)$ is the *arbitrage price* (or more exactly, *arbitrage-free price* or *no-arbitrage price*); an investor selling for this value is *perfectly hedged.*

We note that, to calculate prices as above, we need to know only
(i) Ω, the set of all possible states,
(ii) the σ-field $\mathcal{F}$ and the filtration (or information flow) $(\mathcal{F}_n)$,
(iii) the EMM P^* (or Q).
We do **NOT** need to know the underlying probability measure P – only its null sets, to know what 'equivalent to P' means (actually, in this model, only the empty set is null).

Now option pricing is our central task, and for pricing purposes P^* is vital and P itself irrelevant. We thus may – and shall – focus attention on P^*, which is called the *risk-neutral* probability measure. *Risk-neutrality* is the central concept of the subject. The concept of risk-neutrality is due in its modern form to Harrison and Pliska [HP] in 1981 – though the idea can be traced back to actuarial practice much earlier. Harrison and Pliska call P^* the *reference measure*; other names are *risk-adjusted* or *martingale measure.* The term 'risk-neutral' reflects the P^*-martingale property of the risky assets, since martingales model fair games.

To summarise, we have the

Theorem (Risk-Neutral Valuation Formula). In a complete viable market, arbitrage-free prices of assets are their discounted expected values under the risk-neutral (equivalent martingale) measure P^* (or Q). With payoff h,

$$V_n(H) = (1+r)^{-(N-n)} E^*[V_N(H)|\mathcal{F}_n] = (1+r)^{-(N-n)} E^*[h|\mathcal{F}_n].$$

§5. European Options. The Discrete Black-Scholes Formula.

We consider the simplest case, the Cox-Ross-Rubinstein *binomial model* of 1979; see [CR], [BK]. We take $d = 1$ for simplicity (one risky asset, one riskless asset or bank account); the price vector is (S_n^0, S_n^1), or $((1+r)^n, S_n)$, where

$$S_{n+1} = \begin{cases} S_n(1+a) & \text{with probability } p, \\ S_n(1+b) & \text{with probability } 1-p \end{cases}$$

with $-1 < a < b$, $S_0 > 0$. So writing N for the expiry time,

$$\Omega = \{1+a, 1+b\}^N,$$

each $\omega \in \Omega$ representing the successive values of $T_{n+1} := S_{n+1}/S_n$, $n = 0, 1, \cdots, N-1$. The filtration is $\mathcal{F}_0 = \{\emptyset, \Omega\}$ (trivial σ-field), $\mathcal{F}_T = \mathcal{F} = 2^\Omega$ (power-set of Ω: class of all subsets of Ω), $\mathcal{F}_n = \sigma(S_1, \cdots, S_n) = \sigma(T_1, \cdots, T_n)$. For $\omega = (\omega_1, \cdots, \omega_N) \in \Omega$, $P(\{\omega_1, \cdots, \omega_N\}) = P(T_1 = \omega_1, \cdots, T_N = \omega_N)$, so knowing the probability measure P (equivalently, knowing p) means we know the distribution of $(T_1, \cdots, T_N)$.

For $p^* \in (0,1)$ to be determined, let P^* correspond to p^* as P does to p. Then the discounted price $(\tilde{S}_n)$ is a P^*-martingale iff

$$E^*[\tilde{S}_{n+1}|\mathcal{F}_n] = \tilde{S}_n, \qquad E^*[(\tilde{S}_{n+1}/\tilde{S}_n)|\mathcal{F}_n] = 1,$$

$$E^*[T_{n+1}|\mathcal{F}_n] = 1 + r \qquad (n = 0, 1, \cdots, N-1),$$

since $S_n = \tilde{S}_n(1+r)^n$, $T_{n+1} = S_{n+1}/S_n = (\tilde{S}_{n+1}/\tilde{S}_n)(1+r)$. But

$$E^*[T_{n+1}|\mathcal{F}_n] = (1+a).p^* + (1+b).(1-p^*)$$

is a weighted average of $1+a$ and $1+b$; this can be $1+r$ iff $r \in [a,b]$. As P^* is to be *equivalent* to P and P has no non-empty null-sets, $r = a, b$ are excluded. Thus by IV.2:

Lemma. The market is viable (arbitrage-free) iff $r \in (a,b)$.

Next, $1+r = (1+a)p^* + (1+b)(1-p^*)$, $r = ap^* + b(1-p^*)$: $r - b = p^*(a-b)$:

Lemma. The EMM exists, is unique, and is given by

$$p^* = (b-r)/(b-a).$$

Corollary. The market is complete.

Now $S_N = S_n \Pi_{n+1}^N T_i$. By the Fundamental Theorem of Asset Pricing, the price C_n of a call option with strike-price K at time n is

$$\begin{aligned} C_n &= (1+r)^{-(N-n)} E^*[(S_N - K)_+ | \mathcal{F}_n] \\ &= (1+r)^{-(N-n)} E^*[(S_n \Pi_{n+1}^N T_i - K)_+ | \mathcal{F}_n]. \end{aligned}$$

Now the conditioning on $\mathcal{F}_n$ has no effect – on S_n as this is $\mathcal{F}_n$-measurable (known at time n), and on the T_i as these are independent of $\mathcal{F}_n$. So

$$\begin{aligned} C_n &= (1+r)^{-(N-n)} E^*[(S_n \Pi_{n+1}^N T_i - K)_+] \\ &= (1+r)^{-(N-n)} \Sigma_{j=0}^{N-n} \binom{N-n}{j} p^{*j} (1-p^*)^{N-n-j} (S_n (1+a)^j (1+b)^{N-n-j} - K)_+; \end{aligned}$$

here j, $N-n-j$ are the numbers of times T_i takes the two possible values $1+a, 1+b$. This is the *discrete Black-Scholes formula* of Cox, Ross & Rubinstein (1979) for pricing a European call option in the binomial model. The European put is similar – or use put-call parity (I.3).

To find the (perfect-hedge) strategy for replicating this explicitly: write

$$c(n,x) := \Sigma_{j=0}^{N-n} \binom{N-n}{j} p^{*j} (1-p^*)^{N-n-j} (x(1+a)^j (1+b)^{N-n-j} - K)_+.$$

Then $c(n,x)$ is the undiscounted P^*-expectation of the call at time n given that $S_n = x$. This must be the value of the portfolio at time n if the strategy $H = (H_n)$ replicates the claim:

$$H_n^0 (1+r)^n + H_n S_n = c(n, S_n)$$

(here by previsibility H_n^0 and H_n are both functions of $S_0, \cdots, S_{n-1}$ only). Now $S_n = S_{n-1} T_n = S_{n-1}(1+a)$ or $S_{n-1}(1+b)$, so:

$$\begin{aligned} H_n^0 (1+r)^n + H_n S_{n-1}(1+a) &= c(n, S_{n-1}(1+a)) \\ H_n^0 (1+r)^n + H_n S_{n-1}(1+b) &= c(n, S_{n-1}(1+b)). \end{aligned}$$

Subtract:

$$H_n S_{n-1}(b-a) = c(n, S_{n-1}(1+b)) - c(n, S_{n-1}(1+a)).$$

So H_n in fact depends only on $S_{n-1}, H_n = H_n(S_{n-1})$ (by previsibility), and

Proposition. The perfect hedging strategy H_n replicating the European call option above is given by

$$H_n = H_n(S_{n-1}) = \frac{c(n, S_{n-1}(1+b)) - c(n, S_{n-1}(1+a))}{S_{n-1}(b-a)}.$$

Notice that the numerator is the difference of two values of $c(n,x)$ with the larger value of x in the first term (recall $b > a$). When the payoff function $c(n,x)$ is an increasing function of x, as for the European call option considered here, this is non-negative. In this case, the Proposition gives $H_n \geq 0$: the replicating strategy does not involve short-selling. We record this as:

Corollary. When the payoff function is a non-decreasing function of the final asset price S_N, the perfect-hedging strategy replicating the claim does not involve short-selling of the risky asset.

§6. Continuous-Time Limit of the Binomial Model.

Suppose now that we wish to price an option in continuous time with initial stock price S_0, strike price K and expiry T. We can use the work above to give a discrete-time approximation, where $N \to \infty$. Given $\rho \geq 0$, the instantaneous interest rate in *continuous* time, define r by

$$r := \rho T/N : \qquad e^{\rho T} = \lim_{N\to\infty}(1 + \frac{\rho T}{N})^N = \lim_{N\to\infty}(1+r)^N.$$

Here r, which tends to zero as $N \to \infty$, represents the interest rate in *discrete* time for the approximating binomial model.
For $\sigma > 0$ fixed (σ^2 is to be a variance in continuous time, which will correspond to the *volatility* of the stock), define a, b by

$$\log((1+a)/(1+r)) = -\sigma/\sqrt{N}, \qquad \log((1+b)/(1+r)) = \sigma/\sqrt{N}$$

(a, b both go to zero as $N \to \infty$). We now have a sequence of binomial models, for each of which we can price options as in IV.5. We shall show that the pricing formula converges as $N \to \infty$ to a limit (we identify this later with the continuous Black-Scholes formula of Ch. VI); see e.g. [BK], 4.6.2.

Lemma. Let $(X_j^N)_{j=1}^N$ be iid with mean μ_N satisfying

$$N\mu_N \to \mu \qquad (N \to \infty)$$

and variance $\sigma^2(1+o(1))/N$. If $Y_N := \Sigma_1^N X_j^N$, then Y_N converges in distribution to normality:

$$Y_N \to Y = N(\mu, \sigma^2) \qquad (N \to \infty).$$

Proof. Use characteristic functions (CFs): since Y_N has mean $\mu_N = \mu(1+o(1))/N$ and variance as given, it also has second moment $\sigma^2(1+o(1))/N$. So it has CF

$$\phi_N(u) := E\exp\{iuY_N\} = \Pi_1^N E\exp\{iuX_j^N\} = [E\exp\{iuX_1^N\}]^N$$

$$= (1 + \frac{iu\mu}{N} - \frac{1}{2}\frac{\sigma^2 u^2}{N} + o(\frac{1}{N}))^N \to \exp\{iu\mu - \frac{1}{2}\sigma^2 u^2\} \qquad (N \to \infty).$$

This is the CF of the normal law $N(\mu, \sigma^2)$. The result follows, since convergence of CFs implies convergence in distribution by Lévy's continuity theorem for CFs ([W], §18.1). //

We can apply this to pricing the call option above:

$$C_0^{(N)} = (1 + \frac{\rho T}{N})^{-N} E^*[(S_0 \Pi_1^N T_n - K)_+]$$

$$= E^*[(S_0\exp\{Y_N\} - (1 + \frac{\rho T}{N})^{-N} K)_+], \qquad (*)$$

where

$$Y_N := \sum\nolimits_1^N \log(T_n/(1+r)).$$

Since $T_n = T_n^N$ above takes values $1+b, 1+a$, $X_n^N := \log(T_n^N/(1+r))$ takes values $\log((1+b)/(1+r))$, $\log((1+a)/(1+r)) = \pm\sigma/\sqrt{N}$ (so has second moment σ^2/N). Its mean is

$$\mu_N := \log\Big(\frac{1+b}{1+r}\Big)(1-p^*) + \log\Big(\frac{1+a}{1+r}\Big)p^* = \frac{\sigma}{\sqrt{N}}(1-p^*) - \frac{\sigma}{\sqrt{N}}p^* = (1-2p^*)\sigma/\sqrt{N}$$

(we shall see below that $1 - 2p^* = O(1/\sqrt{N})$, so the Lemma will apply). Now (recall $r = \rho T/N = O(1/N)$)

$$a = (1+r)e^{-\sigma/\sqrt{N}} - 1, \qquad b = (1+r)e^{\sigma/\sqrt{N}} - 1,$$

so $a, b, r \to 0$ as $N \to \infty$, and

$$1 - 2p^* = 1 - 2\frac{(b-r)}{(b-a)} = 1 - 2\frac{[(1+r)e^{\sigma/\sqrt{N}} - 1 - r]}{[(1+r)(e^{\sigma/\sqrt{N}} - e^{-\sigma/\sqrt{N}})]}$$

$$= 1 - 2\frac{[e^{\sigma/\sqrt{N}} - 1]}{[e^{\sigma/\sqrt{N}} - e^{-\sigma/\sqrt{N}}]}.$$

Now expand the two $[\cdots]$ terms above by Taylor's theorem: they give

$$\frac{\sigma}{\sqrt{N}}(1 + \frac{1}{2}\frac{\sigma}{\sqrt{N}} + \cdots), \qquad \frac{2\sigma}{\sqrt{N}}(1 + \frac{\sigma^2}{6N} + \cdots).$$

So, cancelling $\sigma/\sqrt{N}$,

$$1 - 2p^* = 1 - \frac{2(1 + \frac{1}{2}\frac{\sigma}{\sqrt{N}} + \cdots)}{2(1 + \frac{\sigma^2}{6N} + \cdots)} = -\frac{1}{2}\frac{\sigma}{\sqrt{N}} + O(1/N):$$

$$N\mu_N = N.\frac{\sigma}{\sqrt{N}}.(-\frac{1}{2}\frac{\sigma}{\sqrt{N}} + O(1/N)) \to \mu := -\frac{1}{2}\sigma^2 \qquad (N \to \infty).$$

So the Lemma applies, with $\mu = -\frac{1}{2}\sigma^2$. In $(*)$, we have $Y_N \to Y$ in distribution and $(1 + \frac{\rho T}{N})^{-N} \to e^{-\rho T}$ as $N \to \infty$. This suggests that

$$C_0^{(N)} \to E[(S_0 e^Y - e^{-\rho T}K)_+],$$

where E is the expectation for the distribution of Y, which is $N(-\frac{1}{2}\sigma^2, \sigma^2)$. This can be justified, by standard properties of convergence in distribution (see e.g. [W]). So if $Z := (Y + \frac{1}{2}\sigma^2)/\sigma$, $Z \sim N(0,1)$, $Y = -\frac{1}{2}\sigma^2 + \sigma Z$, and

$$C_0^{(N)} \to \int_{-\infty}^{\infty} [S_0 \exp\{-\frac{1}{2}\sigma^2 + \sigma x\} - e^{-\rho T}K]_+ \frac{e^{-\frac{1}{2}x^2}}{\sqrt{2\pi}} dx \qquad (N \to \infty).$$

To evaluate the integral, note first that $[...] > 0$ where

$$S_0 \exp\{-\frac{1}{2}\sigma^2 + \sigma x\} > e^{-\rho T}K, \qquad -\frac{1}{2}\sigma^2 + \sigma x > \log(K/S_0) - \rho T:$$

$$x > [\log(K/S_0) + \frac{1}{2}\sigma^2 - \rho T]/\sigma = c, \quad \text{say.}$$

So writing $\Phi(x)$ for the standard normal distribution function,

$$C_0 = S_0 \int_c^\infty e^{-\frac{1}{2}\sigma^2}.\exp\{-\frac{1}{2}x^2 + \sigma x\}dx/\sqrt{2\pi} - Ke^{-\rho T}[1 - \Phi(c)].$$

The remaining integral is

$$\int_c^\infty \exp\{-\frac{1}{2}(x-\sigma)^2\}dx/\sqrt{2\pi} = \int_{c-\sigma}^\infty \exp\{-\frac{1}{2}u^2\}du/\sqrt{2\pi} = 1 - \Phi(c-\sigma).$$

So the option price is given as a function of the initial price S_0, strike price K, expiry T, interest rate ρ and variance σ^2 by

$$C_0 = S_0[1 - \Phi(c-\sigma)] - Ke^{-\rho T}[1 - \Phi(c)], \qquad c = [\log(K/S_0) + \frac{1}{2}\sigma^2 - \rho T]/\sigma.$$

Notation. To compare with our later work, we now make two changes:
(i) We replace ρ by r: r is our preferred letter for the interest rate, in discrete or continuous time depending on context.
(ii) We replace σ^2 by $\sigma^2 T$; so σ^2 is now the *variance per unit time*. Its square root, σ, is called the *volatility* of the stock. Now $c - \sigma$, c are $c_\pm$, where

$$c_\pm := [\log(K/S_0) - (r \pm \frac{1}{2}\sigma^2)T]/\sigma\sqrt{T}.$$

The result extends immediately to give the price of the option at time $t \in (0,T)$, by replacing T by $T-t$, S_0 by S_t.

We re-write the formula above in this notation. Using the symmetry of the normal distribution, $1 - \Phi(c_\pm) = \Phi(-c_\pm) = \Phi(d_\pm)$, say, where $d_\pm := -c_\pm$:

Theorem (Black-Scholes Formula, 1973). The price of the European call option is

$$c_t = S_t\Phi(d_+) - e^{-r(T-t)}K\Phi(d_-), \quad d_\pm := [\log(S/K) + (r \pm \frac{1}{2}\sigma^2)(T-t)]/\sigma\sqrt{T-t}.$$

We shall return to this in Chapter VI, where we re-derive it by continuous-time *methods* (Brownian motion and Itô calculus) – better, as this is a continuous-time *result*. The same argument (or put-call parity) gives:

Theorem (Black-Scholes Formula). The price of the European put is

$$p_t = Ke^{-r(T-t)}\Phi(-d_-) - S_t\Phi(-d_+).$$

Note. 1. The proof above starts from a binomial distribution and ends with a normal distribution. The binomial distribution is that of a sum of independent Bernoulli random variables. That sums (equivalently, averages) of independent random variables with finite means and variances gives a normal limit is the content of the Central Limit Theorem or CLT (the *Law of Errors*, as physicists would say). The particular form of the CLT used here – normal approximation to the binomial – is the *de Moivre-Laplace limit theorem*.

The picture for this is familiar. The Binomial distribution $B(n, p)$ has a histogram with $n + 1$ bars, whose heights peak at the mode and decrease to either side. For large n, one can draw a smooth curve through the histogram. The curve looks like a normal density curve (with the appropriate location and scale, i.e. mean and variance). The result proved above, and the classical de Moivre-Laplace limit theorem, say that this is exactly right.

2. The Cox-Ross-Rubinstein binomial model above goes over in the limit to the geometric Brownian motion model of VI.1.

3. For similar derivations of the discrete Black-Scholes formula and the passage to the limit to the continuous Black-Scholes formula, see e.g. [CR], §5.6.

4. One of the most striking features of the Black-Scholes formula is that it does **not** involve the mean rate of return μ of the stock - only the riskless interest-rate r and the volatility of the stock σ. Mathematically, this reflects the fact that the change of measure involved in the passage to the risk-neutral measure involves a change of drift. This eliminates the μ term; see Ch. VI.

5. The Black-Scholes formula involves the parameter σ (where σ^2 is the variance of the stock per unit time), called the *volatility* of the stock. In financial terms, this represents how sensitive the stock-price is to new information – how 'volatile' the market's assessment of the stock is. This volatility parameter is very important, *but* we do not know it; instead, we have to *estimate* the volatility for ourselves. There are two approaches:

(a) *historic volatility*: here we use Time Series methods to estimate σ from past price data. The more variable past prices, the more volatile the stock price is, we can estimate σ in this way given enough data.

(b) *implied volatility*: match observed option prices to theoretical option prices. For, the price we see options traded at tells us what the *market* thinks the volatility is (estimating volatility this way works because the dependence is monotone; see later).

If the Black-Scholes model were perfect, these two estimates would agree (to within sampling error). But discrepancies can be observed, which shows the imperfections of our model.

6. Volatility estimation is a major topic, both theoretically and in practice. We return to this in IV.7.3-4 below and VI.7.5-8. But looking ahead:
(i) trading is itself one of the major causes of volatility;
(ii) *options like volatility* [i.e., option prices go up with volatility].
Recalling Ch. I, this shows that volatility is a 'bad thing' from the point of view of the real economy (uncertainty about, e.g., future material costs is nothing but a nuisance to manufacturers), but a 'good thing' for financial markets (trading increases volatility, which increases option prices, which generates more trade ...) – at the cost of increased instability.
7. Volatility is important, not only at company (microeconomic) level, as with options on stock as here, but also at national (macroeconomic) level. The major indexes – FTSE ("footsie": Financial Times Stock Exchange Index, UK), Dow-Jones (US), DAX (Germany) etc. – themselves show volatility, in response to major geopolitical/geofinancial changes, crises etc.

§7. More on European Options

1. *Bounds.*

We use the notation above. We also write c, p for the values of European calls and puts, C, P for the values of the American counterparts.

Obvious upper bounds are $c \leq S, C \leq S$, where S is the stock price (we can buy for S on the market without worrying about options, so would not pay more than this for the right to buy). For puts, one has correspondingly the obvious upper bounds $p \leq K, P \leq K$, where K is the strike price: one would not pay more than K for the right to sell at price K, as one would not spend more than one's maximum return. For lower bounds:
$c_0 \geq \max(S_0 - Ke^{-rT}, 0)$.
Proof. Consider the following two portfolios:
I: one European call plus Ke^{-rT} in cash; II: one share. Show "I $\geq$ II".
$p_0 \geq \max(Ke^{-rT} - S_0, 0)$.
Proof. By above and put-call parity.