m3f22l16tex

Lecture 16 11.11.2016

Proof of the Completeness Th. (concluded).

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 equivalent martingale measure, different from P^*. So incompleteness implies non-uniqueness of equivalent martingale measures. //

Martingale Representation. To say that every contingent claim can be replicated means that every P^*-martingale (where P^* is the risk-neutral measure, which is unique) can be written, or *represented*, as a martingale transform (of the discounted prices) by the replicating (perfect-hedge) trading strategy H. In stochastic-process language, this says that all P^*-martingales can be *represented* as martingale transforms of discounted prices. Such Martingale Representation Theorems hold much more generally, and are very important. For the Brownian case, see VI and [RY], Ch. V.

Note. In the example of Chapter I, we saw that the simple option there could be replicated. More generally, in our market set-up, *all* options can be replicated – our market is *complete*. Similarly for the Black-Scholes theory below.

§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**: §2). If the market is *viable* (arbitrage-free), equivalent martingale measures P^* *exist*.

2 (**Completeness Theorem**: §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 Pricing 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 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 $T_{n+1} := S_{n+1}/S_n$. The filtration is $\mathcal{F}_0 = \{\emptyset, \Omega\}$ (trivial σ-field), $\mathcal{F}_T = \mathcal{F} = 2^\Omega$ (power-set of Ω: all subsets of Ω), $\mathcal{F}_n = \sigma(S_1, \cdots, S_n) = \sigma(T_1, \cdots, T_n)$. For $\omega = (\omega_1, \cdots, \omega_N) \in$

Ω, $P(\{\omega_1, \cdots, \omega_N\}) = P(T_1 = \omega_1, \cdots, T_N = \omega_N)$, so knowing the pr. measure P (i.e. 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

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

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 §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 equivalent mg measure 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).