m3a22l20tex

Lecture 20 24.10.2015

American puts (ctd).

5. Treat these values as 'terminal node values', and fill in the values one time-step earlier by repeating Step 4 for this 'reduced tree'.

6. Iterate. The value of the American put at time 0 is the value at the root - the last node to be filled in. The 'early-exercise region' is the node set where the early-exercise value is the higher; the rest is the 'continuation region'.

Note. The above procedure is simple to describe and understand, and simple to programme. It is laborious to implement numerically by hand, on examples big enough to be non-trivial. Numerical examples are worked through in detail in [H1], 359-360 and [CR], 241-242.

Mathematically, the task remains of describing the *continuation region* - the part of the tree where early exercise is not optimal. This is a classical *optimal stopping problem*. No explicit solution is known (and presumably there isn't one). We will, however, connect the work above with that of III.7 [L13] on the Snell envelope. Consider the pricing of an American put, strike price K, expiry N, in discrete time, with discount factor $1 + r$ per unit time as earlier. Let $Z = (Z_n)_{n=0}^N$ be the payoff on exercising at time n. We want to price Z_n, by U_n say (to conform to our earlier notation), so as to avoid arbitrage; again, we work backwards in time. The recursive step is

$$U_{n-1} = \max(Z_{n-1}, \frac{1}{1+r} E^*[U_n|\mathcal{F}_{n-1}]),$$

the first alternative on the right corresponding to early exercise, the second to the discounted expectation under P^*, as usual. Let $\tilde{U}_n = U_n/(1 + r)^n$ be the discounted price of the American option. Then

$$\tilde{U}_{n-1} = \max(\tilde{Z}_{n-1}, E^*[\tilde{U}_n|\mathcal{F}_{n-1}]) :$$

$(\tilde{U}_n)$ is the *Snell envelope* (III.7) of the discounted payoff process $(\tilde{Z}_n)$. So:

(i) a P^*-supermartingale,

(ii) the smallest supermartingale dominating $(\tilde{Z}_n)$,

(iii) the solution of the optimal stopping problem for $\tilde{Z}$.

Note. One can use the Snell envelope to prove Merton's theorem (equivalence of American and European calls) without using arbitrage arguments. For details see e.g. [BK, Th. 4.7.1 and Cor. 4.7.1].

P-measure and P^*- (or $Q-$) measure.

We use P and P^* in the above, as E and E^* are convenient, but P and Q when the emphasis is on Q, for brevity.

The measure P, the *real* (or real-world) probability measure, models the uncertainty driving prices, which are indeed uncertain, thus allowing us to bring mathematics to bear on financial problems. But P is difficult to get at directly. By contrast, Q is more accessible: the *market* tells us about Q, or more specifically, *trading* does. In addition, trading also tells us about the *volatility* σ, via implied volatility, which we can infer from observing the prices at which options are traded. So Q is certainly more accessible than P. There is thus a sense in which it is Q, rather than P, which is the more real.

It is as well to bear all this in mind when looking at specific problems, particularly numerical ones. Now that we know the CRR binomial-tree model, which gives us the Black-Scholes formula in discrete time (and hence also, by the limiting argument above, the Black- Scholes formula in continuous time, the main result of the course), we can recognise the 'one-period, up or down' model ($/SFr in I.8 L5, price of gold in Problems 5), though clearly artificial and stylised, as a workable 'building block' of the whole theory. Because P itself does not occur in the Black-Scholes formula(e), from a purely financial point of view there is little need to try to construct more realistic, and so more complicated, models of P. Instead, one can exploit what one can infer about Q, which does occur in Black-Scholes, from seeing the prices at which options trade.

From the economic point of view, it is the real world, the real economy, and so the real probability measure P, that matters. The 'Q-measure-eye view of the world' has a degree of artificiality, in so far as options do. One can eat food, and needs to. One can't eat options.

Where we are.

The course splits neatly into three parts: Ch. I, II [L 1-10] on background, Ch. III, IV [L 11-20] on discrete time, and Ch. V, VI [L 20-30] on continuous time. We have already seen the main ideas – and proved nearly everything seen so far. In V, VI we gain the tremendous power of Itô (stochastic) calculus (calculus is our most powerful weapon, in mathematics and science!), and the ability to work in continuous time. What we lose is the ability to prove so much and to see what is happening so clearly and so concretely.

Chapter V. STOCHASTIC PROCESSES IN CONTINUOUS TIME

§1. Filtrations; Finite-Dimensional Distributions

The underlying set-up is as before, but now *time is continuous rather than discrete*; thus the time-variable will be $t \geq 0$ in place of $n = 0, 1, 2, \ldots$. The information available at time t is the σ-field $\mathcal{F}_t$; the collection of these as $t \geq 0$ varies is the filtration, modelling the information flow. The underlying probability space, endowed with this filtration, gives us the stochastic basis (filtered probability space) on which we work,

We assume that the filtration is *complete* (contains all subsets of null-sets as null-sets), and *right-continuous*: $\mathcal{F}_t = \mathcal{F}_{t+}$, i.e.

$$\mathcal{F}_t = \cap_{s>t} \mathcal{F}_s$$

(the 'usual conditions' – right-continuity and completeness – in Meyer's terminology).

A stochastic process $X = (X_t)_{t \geq 0}$ is a family of random variables defined on a filtered probability space with X_t $\mathcal{F}_t$-measurable for each t: thus X_t is known when $\mathcal{F}_t$ is known, at time t.

If $\{t_1, \cdots, t_n\}$ is a finite set of time-points in $[0, \infty)$, $(X_{t_1}, \cdots, X_{t_n})$, or $(X(t_1), \cdots, X(t_n))$ (for typographical convenience, we use both notations interchangeably, with or without ω: $X_t(\omega)$, or $X(t, \omega)$) is a random n-vector, with a distribution, $\mu(t_1, \cdots, t_n)$ say. The class of all such distributions as $\{t_1, \cdots, t_n\}$ ranges over all finite subsets of $[0, \infty)$ is called the class of all *finite-dimensional distributions* of X. These satisfy certain obvious consistency conditions:

(i) deletion of one point t_i can be obtained by 'integrating out the unwanted variable', as usual when passing from joint to marginal distributions,

(ii) permutation of the t_i permutes the arguments of the measure $\mu(t_1, \cdots, t_n)$ on $\mathbb{R}^n$.

Conversely, a collection of finite-dimensional distributions satisfying these two consistency conditions arises from a stochastic process in this way (this is the content of the DANIELL-KOLMOGOROV Theorem: P. J. Daniell in 1918, A. N. Kolmogorov in 1933).

Important though it is as a general existence result, however, the Daniell-Kolmogorov theorem does not take us very far. It gives a stochastic process X as a random function on $[0, \infty)$, i.e. a random variable on $\mathbb{R}^{[0,\infty)}$. This is a vast and unwieldy space; we shall usually be able to confine attention to much smaller and more manageable spaces, of functions satisfying regularity conditions. The most important of these is *continuity*: we want to be able to realise $X = (X_t(\omega))_{t \geq 0}$ as a random *continuous* function, i.e. a member of $C[0, \infty)$; such a process X is called *path-continuous* (since the map $t \to X_t(\omega)$ is called the sample path, or simply path, given by ω) – or more briefly, *continuous*. This is possible for the extremely important case of *Brownian motion* (below), for example, and its relatives. Sometimes we need to allow our random function $X_t(\omega)$ to have jumps. It is then customary, and convenient, to require X_t to be *right-continuous with left limits* (rcll), or càdlàg (continu à droite, limite à gauche) – i.e. to have X in the space $D[0, \infty)$ of all such functions (the *Skorohod space*). This is the case, for instance, for the *Poisson process* and its relatives.

General results on realisability – whether or not it is possible to *realise*, or obtain, a process so as to have its paths in a particular function space – are known, but it is usually better to *construct* the processes we need directly on the function space on which they naturally live.

Given a stochastic process X, it is sometimes possible to improve the regularity of its paths without changing its distribution (that is, without changing its finite-dimensional distributions). For background on results of this type (separability, measurability, versions, regularization, ...) see e.g. Doob's classic book [D].

The continuous-time theory is technically much harder than the discrete-time theory, for two reasons:
(i) questions of path-regularity arise in continuous time but not in discrete time,
(ii) *uncountable* operations (like taking sup over an interval) arise in continuous time. But measure theory is constructed using *countable* operations: uncountable operations risk losing measurability.

Filtrations and Insider Trading

Recall that a filtration models an information flow. In our context, this is the information flow on the basis of which market participants – traders, investors etc. – make their decisions, and commit their funds and effort. All this is information in the *public* domain – necessarily, as stock exchange prices are publicly quoted.