Continuous Time Finance Notes, Spring 2004 – Section 1. *1/21/04*
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course Continuous Time Finance.

This section discusses risk-neutral pricing in the continuous-time setting, for a market with just one source of randomness. In doing so we'll introduce and/or review some essential tools from stochastic calculus, especially the *martingale representation theorem* and *Girsanov's theorem*. For the most part I'm following Chapter 3 of Baxter & Rennie.

****

Let's start with some general orientation. Our focus is on the pricing of derivative securities. The most basic market model is the case of a lognormal asset with no dividend yield, when the interest rate is constant. Then the asset price S and the bond price B satisfy

$$dS = \mu S\,dt + \sigma S\,dw, \quad dB = rB\,dt$$

with μ, σ, and r all constant. We are also interested in more sophisticated models, such as:

(a) Asset dynamics of the form $dS = \mu(S,t)S\,dt + \sigma(S,t)S\,dw$ where $\mu(S,t)$ and $\sigma(S,t)$ are known functions. In this case S is still *Markovian* (the statistics of dS depend only on the present value of S). Almost everything we do in the lognormal case has an analogue here, except that explicit solution formulas are no longer so easy.

(b) Asset dynamics of the form $dS = \mu_t S\,dt + \sigma_t S\,dw$ where μ_t and σ_t are stochastic processes (depending only on information available by time t; more technically: they should be $\mathcal{F}_t$-measurable where $\mathcal{F}_t$ is the sigma-algebra associated to w). In this case S is *non-Markovian*. Typically μ and σ might be determined by separate SDE's. Stochastic volatility models are in this class (but they use two sources of randomness, i.e. the SDE for σ makes use of another, independent Brownian motion process; this makes the market incomplete.)

(c) Interest rate models where r is not constant, but rather random; for example, the spot rate r may be the solution of an SDE. (An interesting class of path-dependent options occurs in the modeling of mortgage-backed securities: the rate at which people refinance mortagages depends on the history of interest rates, not just the present spot rate.)

(d) Problems involving two or more sources of randomness, for example options whose payoff depends on more than one stock price [e.g. an option on a portfolio of stocks; or an option on the max or min of two stock prices]; quantos [options involving a random exchange rate and a random stock process]; and incomplete markets [e.g. stochastic volatility, where we can trade the stock but not the volatility].

My goal is to review background and to start relatively slowly. Therefore I'll focus in this section on the case when there is just one source of randomness (a single, scalar-valued Brownian motion).

Option pricing on a binomial tree is easy. The subjective probability is irrelevant, except that it determines the stock price tree. The risk-neutral probabilities (on the same tree!) are determined by

$$E_t^{\mathrm{RN}}[e^{-r\delta t}s_{t+\delta t}] = s_t.$$

If the interest rate r is constant, then the price of a contingent claim with payout $f(S_T)$ at time T is obtained by taking discounted expected payoff

$$V_0 = e^{-rT}E^{\mathrm{RN}}[f(S_T)].$$

The binomial tree method works even if interest rates are not constant. They can even be random (they can vary from node to node; all that matters is that the interest earned starting from any node be the same whether the stock goes up or down from that node). But if interest rates are not constant then we must discount correctly; and if they are random then the discounting must be inside the expectation:

$$V_0 = E^{\mathrm{RN}}\left[e^{-\int_0^T r(s)\,ds}f(S_T)\right]$$

If the payout is path-dependent (or if the subjective process is non-Markovian, i.e. its SDE has history-dependent coefficients) then we cannot use a recombining tree; in this case the binomial tree method is OK in concept but hard to use in practice.

Recall what lies behind the binomial-tree pricing formula: the binomial-tree market is complete, i.e. every contingent claim is replicatable. The prices given above are the initial cost of a replicating portfolio. So they are forced upon us (for a given tree) by the absence of arbitrage. (My Derivative Securities notes demonstrated this "by example," but see Chapter 2 of Baxter & Rennie for an honest yet elementary proof.)

The major flaw of this binomial-tree viewpoint: the market is not really a binomial tree. Moreover, a similar argument with a trinomial tree would not give unique prices (since a trinomial market is not complete), though it's easy to specify a trinomial model which gives lognormal dynamics in the limit $\delta t \to 0$. Similar but more serious: what to do when there are two stocks that move independently? (Using a binomial approximation for each gives an incomplete model, since each one-period subtree has four branches but just three tradeables.) Thus: if the market model we'd really like to use is formulated in continuous time, it would be much better to formulate the theory as well in continuous time (viewing the time-discrete models as numerical approximation schemes).

In the continuous-time setting, options can be priced using the Black-Scholes PDE. The solution of the Black-Scholes PDE tells us how to choose a replicating portfolio (namely: hold $\Delta = \frac{\partial V}{\partial S}(S_t, t)$ units of stock at time t); so the price we get using the Black-Scholes PDE is forced upon us by the absence of arbitrage, just as in the analysis of binomial trees. Let's be clear about what this means: a continuous-time *trading strategy* is a pair of stochastic processes, ϕ_t and ψ_t, each depending only on information available by time t; the associated (time-dependent) "portfolio" holds ϕ_t units of stock and ψ_t units of bond at time t. Its value at time t is thus

$$W_t = \phi_t S_t + \psi_t B_t.$$

(I call this W not V, because it represents the investor's wealth at time t as he pursues this trading strategy.) It is self-financing if

$$dW = \phi\, dS + \psi\, dB.$$

Our assertion is that if $V(S,t)$ solves the Black-Scholes PDE, and its terminal value $V(S,T)$ matches the option payoff, then the associated trading strategy

$$\phi_t = \frac{\partial V}{\partial S}(S_t, t), \quad \psi_t = (V(S_t,t) - \phi_t S_t)/B_t$$

is self-financing, and its value at time t is $V(S_t,t)$ for every t. (The assertion about its value is obvious; the fact that it's self-financing is a consequence of Ito's formula applied to $V(S_t,t)$, together with the Black-Scholes PDE and the fact that $dB = rBdt$.) Thus $V(S_0, 0)$ is the time-0 value of a trading strategy that replicates the option payoff at time T.

We can recognize the option value as its "risk-neutral expected discounted payoff" by noticing that the Black-Scholes PDE is linked to a suitably-defined "risk neutral" diffusion by the Feynman-Kac formula. Indeed, suppose V solves the Black-Scholes PDE and S solves the stochastic differential equation $dS = rS\,dt + \sigma S\,dw$; for simplicity let's suppose the interest rate r is constant. Then $e^{r(T-t)}V(S_t,t)$ is a martingale, since its stochastic differential is $e^{r(T-t)}$ times $-rV dt + V_S dS + (1/2)V_{SS}\sigma^2 S^2 dt + V_t = \sigma S V_S\, dW$. So the expected value of $e^{r(T-t)}V(S_t,t)$ at time T (the risk-neutral expected payoff) equals the value of this expression at time 0, namely $e^{rT}V(S_0,0)$, as asserted. (The preceding argument works, with obvious modifications, even if r is time-dependent or S-dependent. In that setting the discount factor $\exp(r(T-t))$ must be replaced by $\exp(\int_t^T r\,ds)$.)

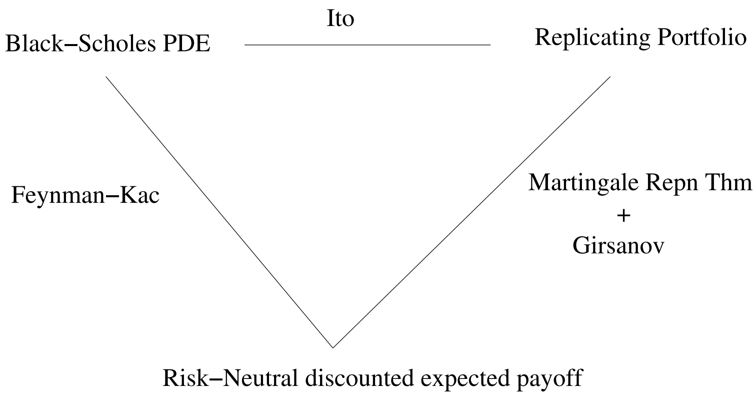

The major shortcoming of this Black-Scholes PDE viewpoint is this: it is limited to Markovian evolution laws ($dS = \mu S\,dt + \sigma S\,dw$, where μ and σ may depend on S and t, but not on the full history of events prior to time t) and path-independent options.

Beyond this rather practical restriction, there's also a conceptual issue. The probabilistic viewpoint is in many ways more natural than the PDE-based one; for example, it gives the analogue of our binomial tree discussion; it provides the basis for Monte Carlo simulation.

Moreover, for some purposes (e.g. understanding change of numeraire) the probabilistic framework is the only one that's really clear. So it's natural to seek a continuous-time understanding directly analogous to what we achieved with binomial trees. (It was also natural to postpone this till now, since such an understanding requires sufficient command of stochastic calculus.)

In summary: consider the triangle shown in the Figure, which shows three alternative ways of thinking about the pricing of derivative securities in a continuous-time framework. The course Derivative Securities focused mainly on two legs of the triangle (connecting the Black-Scholes PDE with the replicating portfolio, and with a suitable discounted expected risk-neutral payoff). Now we'll seek an equally clear understanding of the third leg.

****

The main tools we need from stochastic calculus are the *martingale representation theorem* and *Girsanov's theorem*. We now discuss the former, for diffusions involving a single source of randomness, i.e. diffusions of the form $dY = \mu_t\, dt + \sigma_t\, dw$, where μ_t and σ_t may be random (but depend only on information available at time t). Recall that an $\mathcal{F}_t$-adapted stochastic process M_t is called a *martingale* if

(a) it is integrable, i.e. $E[|Y_t|] < \infty$ for all t, and

(b) its conditional expectations satisfy $E[M_t|\mathcal{F}_s] = M_s$ for all $s < t$.

The expectation is of course with respect to a measure on path space. When M solves an SDE of the form $dM = \mu\, dt + \sigma\, dw$ where w is Brownian motion then the expectation we have in mind is the one associated with the underlying Brownian motion (and the sigma-algebra $\mathcal{F}_t$ is the one generated by this Brownian motion). Soon we'll discuss Girsanov's theorem, which deals with changing the measure (and, as a result, changing the stochastic differential equation). Then we must be careful which measure we're using – calling M a P-martingale if conditions (a) and (b) hold when the expectation is taken using measure P. But for now let's suppose there's only one measure under discussion, so we need not specify it.

The integrability condition is important to a probabilist (it's easy to construct processes satisfying (b) but not (a) – these are called "local martingales" – and many of the theorems we want to use have counterexamples if M is just a local martingale). However the processes of interest in finance are always integrable, and my purpose is to get the main ideas without unnecessary technicality, so I won't fuss much over integrability.

There are two different ways of constructing martingales defined for $0 \le t \le T$. One is to take an $\mathcal{F}_T$-measurable random variable X (e.g. an option payoff – which is now permitted to be path-dependent, i.e. to depend on all information up to time T) and take its conditional expectations:

$$M_t = E[X|\mathcal{F}_t]. \tag{1}$$

This is a martingale due to the "tower property" of conditional expectations. The second method of constructing martingales is to solve a stochastic differential equation of the form

$$dM = \phi_t\, dw \tag{2}$$

where ϕ_t is $\mathcal{F}_t$-adapted (i.e. it depends only on information available by time t).

Of course in considering (b), we need some condition on ϕ to be sure M is integrable. An obvious sufficient condition is that $E[\int_0^T \psi_s^2 \, ds] < \infty$, since this assures us that $E[M^2(t)] < \infty$. In finance (2) typically takes the form $dM = \sigma_t M \, dw$. In this case one can show that the "Novikov condition" $E\left[\exp\left(\frac{1}{2}\int_0^T \sigma_s^2 \, ds\right)\right] < \infty$ is sufficient to assure integrability. This fact is not easy to prove with full rigor; but it is certainly intuitive, since the SDE $dM = \sigma_t M \, dw$ has the explicit solution

$$M_t = M_0 e^{\int_0^t \sigma_s \, dw - \frac{1}{2}\int_0^t \sigma_s^2 \, ds}$$

(To see this: apply Ito's formula to the right hand side to verify that the process defined by this formula satisfies $dM = \sigma_t M \, dw$.)

In its simplest form, the martingale representation theorem says that *any martingale can be expressed in the form (2).* Moreover the associated ϕ_t is unique. Thus, for example, a martingale created via conditional expectations as in (1) can alternatively be described by an SDE.

We need a slightly more sophisticated version of the martingale representation theorem. Rather than expressing M in terms of the Brownian motion w, we'll need to represent it in terms of *another* martingale, say N. The feasibility of doing this is obvious: if N and M are both martingales, then our simplest form of the martingale representation theorem says they can both be represented in the form (2):

$$dM = \phi^M \, dw \quad \text{and} \quad dN = \phi^N \, dw.$$

So eliminating dw gives a representation of the form

$$dM = \psi_t \, dN \tag{3}$$

with $\psi = \phi^M/\phi^N$. This representation is the more sophisticated version we wanted. Of course we cannot divide by 0: we assumed, in deriving (3), that $\phi^N \neq 0$ with probability 1. Like our simpler version (2), the representation (3) is unique, i.e. the density ψ is uniquely determined by M and N.

****

The martingale representation theorem is all we need to connect risk-neutral expectations with self-financing portfolios. This connection works even for path-dependent options, and even for stock processes whose drift and volatility are random (but $\mathcal{F}_t$-measurable). It also permits the risk-free rate to be random. To demonstrate the strength of the method we work at this level of generality. Thus our market has a stock whose price satisfies

$$dS = \mu_t S \, dt + \sigma_t S \, dw, \quad S(0) = S_0$$

and a bond whose value satisfies

$$dB = r_t B \, dt, \quad B(0) = 1$$

where μ_t, σ_t, and r_t are all $\mathcal{F}_t$-measurable. Our goal is to price an option with payoff X, where X is any $\mathcal{F}_T$-measurable random variable.

To get started, we need one more assumption. We suppose there is a measure Q on path space such that the ratio S_t/B_t is a Q-martingale. This is the "risk-neutral measure." The whole point of Girsanov's theorem is to prove existence of such a Q (and to describe it). For now, we simply assume Q exists.

Here's the story: given such a Q, we'll show that every payoff X is replicatable. Moreover if V_t is the value at time t of the replicating portfolio then V_t is characterized by

$$V_t/B_t = E_Q[X/B_T|\mathcal{F}_t]. \tag{4}$$

In particular, the initial cost of the replicating portfolio is the expected discounted payoff $V_0 = E_Q[X/B_T]$. This is therefore the value of the option.

The proof is surprisingly easy. Let V_t be the process defined by (4). Then V_t/B_t is a Q-martingale, by (1) applied to X/B_T. Also S_t/B_t is a Q-martingale, by hypothesis. Therefore by the martingale representation theorem there is a process ϕ_t such that

$$d(V/B) = \phi_t\, d(S/B).$$

We shall show that the trading strategy which holds ϕ_t units of stock and $\psi_t = (V_t - \phi_t S_t)/B_t$ units of bond at time t is self-financing and replicates the option.

Easiest first. It replicates the option because $V_t = \phi_t S_t + \psi_t B_t$ (by the choice of ψ) and $V_T = X$ (by (4), noting that X/B_T is $\mathcal{F}_T$ measurable so taking its conditional expectation relative to $\mathcal{F}_T$ doesn't change it). Thus $V_T = X$, which is exactly what we mean when we say it replicates the option.

OK, now a little work. We must show that the proposed trading strategy is self-financing, i.e. that $dV = \phi_t\, dS + \psi_t\, dB$. Recall from Ito's lemma that when X and Y are diffusions, $d(XY) = X\,dY + Y\,dX + dX\,dY$. Also note that if the SDE for one of the two processes has no "dw" term then $dX\,dY = 0$ and the formula becomes $d(XY) = X\,dY + Y\,dX$. We apply this twice: since $V = (V/B)B$ and $d(V/B) = \phi\, d(S/B)$ we have

$$dV = d((V/B)B) = B\phi\, d(S/B) + (V/B)\, dB; \tag{5}$$

and since $S = (S/B)B$ we have

$$\phi\, dS = B\phi\, d(S/B) + \phi(S/B)\, dB. \tag{6}$$

Now, from the definition of ψ we have

$$\psi\, dB = [(V/B) - \phi(S/B)]\, dB. \tag{7}$$

Adding the last two equations gives

$$\phi\, dS + \psi\, dB = B\phi\, d(S/B) + (V/B)\, dB.$$

Comparing this with (5), we conclude that the portfolio is self-financing. The proof is now complete.

Question: where did we use the hypothesis that $dB = r_t B dt$? Answer: when we applied Ito's formula in (5) and (6). Remarkably, however, this hypothesis was not necessary! The assertion remains true (though the proof needs some adjustment) even if B is a volatile diffusion process. Thus B need not be a the value of a bond! We'll return to this when we discuss *change of numeraire*. But let's check now that the preceding argument works even if B is a volatile diffusion (i.e. if it solves an SDE with a nonzero dw term). In this case (5) must be replaced by

$$dV = d((V/B)B) = B\phi\, d(S/B) + (V/B)\, dB + d(V/B)\, dB$$

and (6) by

$$\phi\, dS = B\phi\, d(S/B) + \phi(S/B)\, dB + \phi\, d(S/B)\, dB.$$

Equation (7) remains unchanged:

$$\psi\, dB = [(V/B) - \phi(S/B)]\, dB.$$

Adding the last two equations, and remembering that $\phi\, d(S/B) = d(V/B)$, we conclude as before that $dV = \phi\, dS + \psi\, dB$, so the portfolio is still self-financing.

Our argument tells us how to replicate (and therefore price) any option, assuming only the existence of a "risk-neutral measure" Q with respect to which S/B is a martingale. We have in effect shown that our simple market (with one source of randomness) is complete. It also follows that the value P_t of *any* tradeable must be such that P_t/B_t is a Q-martingale (for the same measure Q). Indeed, we can replicate the payoff P_T starting at time t at cost $B_t E_Q[P_T/B_T|\mathcal{F}_t]$. To avoid arbitrage, this had better be the market price of the same payoff, namely P_t. Thus $P_t/B_t = E_Q[P_T/B_T|\mathcal{F}_t]$, as asserted.

****

Now what about the existence of Q? This is the point of Girsanov's theorem. Restricted to the case of a single Brownian motion, it says the following. Consider a measure P on paths, and suppose w is a P-Brownian motion (i.e. the induced measure on $w_t - w_s$ is Gaussian for each $s < t$, with mean 0 and variance $t - s$, and $w_t - w_s$ is independent of $w_s - w_r$ for $r < s < t$.). Consider in addition an adapted process γ_t, and the associated martingale

$$M_t = e^{-\int_0^t \gamma_s\, dw - \frac{1}{2}\int_0^t \gamma_s^2\, ds}.$$

(As noted earlier in these notes, γ_t must satisfy some integrability condition to be sure M_t is a martingale; the Novikov condition $E_P\left[\exp\left(\frac{1}{2}\int_0^T \gamma_s^2\, ds\right)\right] < \infty$ is sufficient.) Finally, suppose we're only interested in behavior up to time T. Then the measure Q defined by

$$E_Q[X] = E_P[M_T X] \quad \text{for every } \mathcal{F}_T\text{-measurable random variable } X$$

(in other words $dQ/dP = M_T$) has the property that

$$\tilde{w}_t = w_t + \int_0^t \gamma_s\, ds \quad \text{is a } Q\text{-Brownian motion.} \tag{8}$$

Moreover, we have the following formula for conditional probabilities taken with respect to Q:

$$E_Q[X_t|\mathcal{F}_s] = E_P[(M_t/M_s)X_t|\mathcal{F}_s] = M_s^{-1}E_P[(M_tX_t|\mathcal{F}_s] \tag{9}$$

whenever X is $\mathcal{F}_t$-measurable and $s \le t \le T$.

It's easy to state the theorem, but not so easy to get one's head around it. The discussion in Baxter & Rennie is very good: they explain in particular how this generalizes our familiar binomial-tree change of measure from the "subjective probabilities" to the "risk-neutral" ones. I also recommend the short Section 3.1 of Steele (on importance sampling, for Gaussian random variables), to gain intuition about change-of-measure in the more elementary setting of a single Gaussian random variable. Baxter & Rennie give an honest proof of (8) for the special case γ_t = constant in Section 3.4 (note that the solution of Problem 3.9 is in the back of the book). Avellaneda & Laurence give essentially the same argument in their Section 9.3, but they present it in greater generality (for nonconstant γ_t). Moreover Avellaneda & Laurence explain the formula (9) for conditional probabilities in the appendix to their Chapter 9. I recommend reading all these sources to gain an understanding of the theorem and why it's true. Here I'll not repeat that material; rather, I simply make some comments:

(a) In finance, the main use of Girsanov's theorem is the following: suppose S solves the SDE $dS = \alpha_t\,dt + \beta_t\,dw$ where w is a P-Brownian motion. Then it also solves the SDE $dS = (\alpha_t - \beta_t\gamma_t)\,dt + \beta_t\,d\tilde{w}$ where $d\tilde{w}$ is a Q-Brownian motion. Indeed, (8) can be restated as $d\tilde{w} = dw + \gamma_t\,dt$, and with this substitution the two SDE's are identical.

(b) Why is there a measure Q such that S/B is a Q-martingale? Well, if S and B solve SDE's then so does S/B (by Ito's lemma). Suppose this SDE has the form $d(S/B) = \alpha_t\,dt + \beta_t\,dw$. Then we can eliminate the drift by choosing $\gamma_t = \alpha_t/\beta_t$. In other words, applying Girsanov with this choice of γ_t gives $d(S/B) = \beta_t\,d\tilde{w}$ with $\tilde{w}$ a Q-Brownian motion. Thus S/B is a Q-martingale.

(c) Our statement of Girsanov's theorem selected (arbitrarily) an initial time 0. This choice should be irrelevant, since Brownian motion is a Markov process (so it can be viewed as starting at any time s with initial value w_s). Also, our statement selected (arbitrarily) a final time T, which should also be irrelevant. These choices are indeed unimportant, as a consequence of (i) the fact that M_t is a P-martingale, and (ii) relation (9), combined with the observation that

$$M(t)/M(s) = e^{-\int_s^t \gamma_\rho\,dw - \frac{1}{2}\int_s^t \gamma_\rho^2\,d\rho}.$$

Bottom line: we have shown in great generality that there is a "risk-neutral" measure Q with respect to which S/B is a martingale, and using it we have shown how to replicate (and therefore price) an arbitrary (even path-dependent) contingent claim.

What good is this? First of all it tells us that Monte-Carlo methods can be used even for path-dependent options and non-Markovian diffusions (whose SDE's have coefficients depending on prior history). But besides that, it is even useful in relatively simple settings (e.g. constant drift, volatility, and interest rates) for mastering concepts like change-of-numeraire. We'll discuss such applications next. (If you want to read ahead, look at Chapter 4 of Baxter & Rennie.)