Continuous Time Finance Notes, Spring 2004 – Section 3, *Feb. 4, 2004*
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course Continuous Time Finance.

Announcement: There will be no class on Wednesday, February 11.

These notes wrap up our treatment of the probabilistic framework for option pricing. We conclude the discussion of problems with one source of randomness by discussing the "market price of risk" and the use of alternative numeraires (leading to consideration of equivalent martingale measures other than the risk-neutral one). Then we discuss the analogous theory for complete markets with multiple sources of randomness. Finally we discuss two important examples where there are two sources of randomness: (a) exchange options, and (b) quanto options. Except for exchange options, this material is covered in Sections 4.4-4.5 and 6.1-6.4 of Baxter & Rennie (however my discussion of quantos is organized differently). The corresponding part of Hull (5th edition) is Chapter 21 (my discussion of exchange options is from there). Hull is relatively terse, but well worth reading.

****

The market price of risk. We continue a little longer the hypothesis that the market has just one source of randomness (a scalar Brownian motion w). Recall what we achieved in Section 1: let r_t be the risk-free rate, and B the associated money-market instrument (characterized by $dB = r_t B\,dt$ with $B(0) = 1$); let S be tradeable, and assume S solves an SDE of the form $dS = \mu_t S\,dt + \sigma_t S\,dw$. We showed that if S_t/B_t is a martingale relative to Q then any payoff X at time T can be replicated by self-financing portfolio which invests in just S and B, and the time-t value of this portfolio V_t is characterized by $V_t = B_t E_Q[X/B_T \mid \mathcal{F}_t]$.

We also showed that Q exists (and is unique), by an application of Girsanov's theorem. Reviewing this: we need S/B to be a Q-martingale. Under the original probability we have

$$d(S/B) = (\mu_t - r_t)(S/B)\,dt + \sigma_t(S/B)\,dw.$$

According to Girsanov's theorem changing the measure (without changing the sets of measure zero) can change the drift but not the volatility. So Q is characterized by

$$d(S/B) = \sigma_t(S/B)\,d\tilde{w}$$

where $\tilde{w}$ is a Q-Brownian motion. Comparing these equations we see that

$$dw + \frac{\mu_t - r_t}{\sigma_t}\,dt = d\tilde{w}.$$

The ratio

$$\lambda_t = \frac{\mu_t - r_t}{\sigma_t}$$

is called the *market price of risk*, by analogy with the Capital Asset Pricing Model (notice that λ is the ratio of excess return to volatility). It is important because it determines Q in terms of P, namely:

$$dQ/dP = \exp\left(-\int_0^T \lambda_t\,dw - \tfrac{1}{2}\int_0^T \lambda_t^2\,dt\right).$$

This Q is called the "risk-neutral measure."

Finally, we observed that if $\overline{S}$ is *any* tradeable in this market, its value $\overline{S}_t$ must have the property that $\overline{S}_t/B_t$ is a Q-martingale (for the *same* risk-neutral measure Q). This shows that *all tradeables must have the same market price of risk.*

It's enlightening to give another proof that all tradeables have the same market price of risk. Suppose S and $\overline{S}$ are distinct tradeables, with

$$dS = \mu_t S\,dt + \sigma_t S\,dw, \quad d\overline{S} = \overline{\mu}_t\overline{S}\,dt + \overline{\sigma}_t\overline{S}\,dw.$$

Consider the following trading strategy: at time t, hold $\overline{\sigma}_t\overline{S}$ units of S and $-\sigma_t S$ units of $\overline{S}$. The value of the associated portfolio is

$$V = \overline{\sigma}_t\overline{S}S - \sigma_t S\overline{S} = (\overline{\sigma}_t - \sigma_t)S\overline{S}.$$

The portfolio is risk-free, since its increments are

$$dV = \overline{\sigma}_t\overline{S}\,dS - \sigma_t S\,d\overline{S} = (\overline{\sigma}_t\mu_t - \sigma_t\overline{\mu}_t)S\overline{S}\,dt.$$

So it must rise at the risk-free rate: $dV = r_t\,V\,dt$, which gives

$$(\overline{\sigma}_t\mu_t - \sigma_t\overline{\mu}_t) = (\overline{\sigma}_t - \sigma_t)r_t$$

Algebraic rearrangement gives

$$\frac{\overline{\mu}_t - r_t}{\overline{\sigma}_t} = \frac{\mu_t - r_t}{\sigma_t},$$

which is the desired conclusion. This argument has the advantage of being in some sense more elementary. However it does not reveal the true meaning of the market price of risk – which is that it expresses the risk-neutral measure Q in terms of the original measure P.

****

Equivalent martingale measures. We observed in Section 1 (the top half of page 7) that the theory just summarized works equally well when B is replaced by any tradeable. In other words, the "numeraire" can be any tradeable N. To price options using N in place of B, we need a measure $\overline{Q}$ such that S/N is a $\overline{Q}$-martingale. Then the value V of payoff X has the property that V/N is a $\overline{Q}$-martingale. And any tradeable $\overline{S}$ has the property that $\overline{S}/N$ is a $\overline{Q}$-martingale. In particular, the value at time 0 of payoff X is

$$V_0 = N_0 E_{\overline{Q}}[X/N_T]. \tag{1}$$

The replicating portfolio uses S and N. If N is not the risk-free money-market fund B then $\overline{Q}$ is different from Q; it is called "forward risk-neutral with respect to N." To see why, observe that the contract to exchange S for k units of N at time T is valueless at time 0 when $k = E_{\overline{Q}}[S/N_T]$. Thus: the $\overline{Q}$ mean of S/N_T is the forward price of S, in units of N.

The risk-neutral measure Q, and more generally the forward risk-neutral measure $\overline{Q}$ associated with any tradeable N, are called *equivalent martingale measures.* We can price options using any of them; the choice is up to us, and is a matter of convenience.

For markets with a single source of randomness the main application of (1) is the justification of Black's formula. In that case $N_t = B(t,T)$ is the value at time t of a zero-coupon bond worth 1 dollar at time T, and (1) becomes

$$V_t = B(0,T)E_{\overline{Q}}[X].$$

Thus, by using the measure $\overline{Q}$ rather than Q, we can arrange that the discount factor be outside the expectation.

We'll give another application of (1) below, to the pricing of an exchange option. That example however involves two sources of randomness. So this is a convenient time to discuss how the theory generalizes to problems with multiple sources of randomness.

****

Complete markets with n sources of randomness. The most important examples involve two sources of randomness, for example an "exchange option" (whose payoff is $(S-\overline{S})_+$, where S and $\overline{S}$ are two stocks) or a "quantos call" (whose payoff in dollars is $(S-K)_+$, where S is the stock price of British Petroleum in pounds). Later we'll specialize to these cases. But to get started let's suppose there are n sources of randomness, and n (independent) tradeables $S^1, \ldots, S^n$. Each individually satisfies an SDE of the form

$$dS_i = \mu_i S_i\,dt + \sigma_i S_i\,dw_i \tag{2}$$

however the Brownian motions w_i may (and often will) be *correlated*. There are two (ultimately equivalent) ways to handle this. The more straightforward one is to express everything in terms of n independent Brownian motions $z_1, \ldots, z_n$:

$$dS_i = \mu_i S_i\,dt + S_i \sum_j \sigma_{ij}\,dz_j \tag{3}$$

where $\Sigma = (\sigma_{ij})$ is an $n \times n$ matrix. (Our n tradeables are "independent" if the matrix Σ is invertible.) Notice that (3) is equivalent to (2) when σ_i is the length of the ith row of Σ and $dw_i = \sigma_i^{-1} \sum_j \sigma_{ij} dz_j$. (Notice that w_i is a Brownian motion!). Sometimes however it is convenient to stick with the correlated Brownian motions w_i. In that case, when we apply Ito's formula, we must replace the shorthand $dw\,dw = dt$ by

$$dw_i\,dw_j = \rho_{ij}\,dt$$

with

$$\rho_{ij} = \sigma_i^{-1}\sigma_j^{-1} \sum_k \sigma_{ik}\sigma_{jk}\, dt.$$

Notice that along the diagonal $\rho_{ii} = 1$, consistent with the fact that each w_i is individually a Brownian motion process.

The n-factor version of the martingale representation theorem is the direct analogue of the 1-factor version: if M is a martingale (for the σ-algebra $\mathcal{F}_t$ generated by our n Brownian motions) then there exist adapted processes $\phi_1, \ldots, \phi_n$ such that

$$dM = \sum_j \phi_j dz_j.$$

If $N_1, \ldots, N_n$ are also martingales and $dN_j = \sum_k \psi_{jk} dz_k$ then (provided the matrix ψ_{jk} is invertible) we deduce the existence of representation

$$dM = \sum_j \phi'_j dN_j.$$

These representations are unique.

The n-factor version of Girsanov's theorem is also directly analogous to the one-factor version. For any adapted, vector-valued process $\gamma = (\gamma_1, \ldots, \gamma_n)$ there is a measure Q such that

$$\tilde{z} = z + \int_0^t \gamma\, ds$$

is an n-dimensional Brownian motion under Q (i.e. each component is an independent Q-Brownian motion). Its density relative the the original measure P (with respect to which z is an n-dimensional Brownian motion) is

$$dQ/dP = \exp\left(-\int_0^T \sum_i \gamma_i\, dz_i - \tfrac{1}{2} \int_0^T |\gamma|^2\, dt \right).$$

Notice that the relation between $\tilde{z}$ and z can be written as

$$d\tilde{z} = dz + \gamma dt.$$

Our discussion of option pricing carries over to the n-factor setting with essentially no change. We won't repeat everything, but let characterize the risk-neutral measure Q. To price options using the tradeables S_i described by (3) we need S_i/B to be a martingale under Q. Under the original probability we have

$$d(S_i/B) = (\mu_i - r)(S_i/B)\, dt + (S_i/B) \sum_j \sigma_{ij}\, dz_j.$$

According to Girsanov's theorem changing the measure (without changing the sets of measure zero) can change the drift but not the volatility. So Q is characterized by

$$d(S_i/B) = (S_i/B) \sum_j \sigma_{ij}\, d\tilde{z}_j.$$

Thus we need

$$d\tilde{z} = dz + \Sigma^{-1}(\mu - r1)\,dt$$

using vector notation, with $\Sigma = (\sigma_{ij})$ and $1 = (1, \ldots, 1)$. The vector-valued process

$$\lambda = \Sigma^{-1}(\mu - r1)$$

is again called the *market price of risk*. It characterizes the risk-neutral measure, which is unique. So as in the one-factor case, the market price of risk is independent of the particular choice of n independent tradeables used to value options and construct replicating portfolios.

A brief digression about incomplete markets. We assumed the existence of n independent tradeables so we could solve uniquely for λ and Q. In some settings however there are fewer independent tradeables than sources of randomness. A simple example would be an option whose payoff depends the values of two underlyings $S_1(T)$ and $S_2(T)$, if only S_1 is tradeable. A more fundamental example is any *stochastic volatility* model, since volatility is never tradeable (at least, not in the literal sense). In such settings the linear system giving $d\tilde{z}$ in terms of dz is underdetermined (we have n unknowns, but only as many equations as we have independent tradeables). Therefore there is more than one possible measure Q, i.e. more than one possible market price of risk. It is still true that every tradeable must have the *same* market price of risk (see Hull's Appendix 21B for a proof). So it is still true that all tradeables must be priced using the same measure Q. But this market is not complete; the measure Q cannot be determined from absence of arbitrage alone; most payoffs cannot be replicated. Pricing and hedging in an incomplete market requires an entirely different viewpoint. We'll discuss some incomplete market models in the last 1/3 of the semester.

Returning now to the complete setting: As in the one-factor case, we need not use the money-market account as our numeraire; rather, we can use any tradeable N. The associated equivalent martingale measure has the property $\overline{Q}$ has the property that S/N is a $\overline{Q}$-martingale for every tradeable S. In particular the value at time 0 of an option with payoff X is $V_0 = N_0 E_{\overline{Q}}[X/N_T]$.

What must we do in practice to value and hedge an option in this multifactor setting? If the risk-free rate r and the volatilities σ_{ij} are constant then the risk-neutral process

$$dS_i = rS_i\,dt + S_i \sum_j \sigma_{ij}\,d\tilde{z}_j$$

is easy to understand: the vector $(\log S_1, \ldots, \log S_n)$ has a multivariate normal distribution. One can therefore value any European option by a suitable integration using the multivariate normal distribution. In greater generality – if σ_{ij} depend only on the state variables S_i and time – one can value options by solving the multivariate analogue of the Black-Scholes PDE. The hedge is, as usual, determined by using $\partial V/\partial S_i$ units of the ith stock.

In some important special cases we can get more explicit valuation formulas. The rest of this section is devoted to two examples of this kind.

****

Exchange options. Suppose now there are two sources of randomness, and we are interested in two stocks S_1 and S_2. Let's use the representation (2), and let ρ be the correlation between w_1 and w_2:

$$dw_1\, dw_2 = \rho\, dt.$$

Consider the "exchange option," which gives the holder the right to exchange S_1 for S_2 at time T. Its payoff is $(S_2 - S_1)_+$.

The key to valuing this option is to use the equivalent martingale measure $\overline{Q}$ which is forward risk-neutral with respect to S_1. It has the property that S/S_1 is a $\overline{Q}$-martingale for every tradeable S. The value it assigns to our option at time 0 is

$$V_0 = S_1(0) E_{\overline{Q}}[(S_2/S_1 - 1)_+]$$

since

$$\frac{(S_2 - S_1)+}{S_1} = (S_2/S_1 - 1)_+.$$

By Ito's formula, the SDE satisfied by $X = S_2/S_1$ (under the original measure) is

$$dX = S_2 d(S_1^{-1}) + S_1^{-1} dS_2 + dS_2\, dS_1^{-1}.$$

Since

$$d(S_1^{-1}) = -S_1^{-2}\, dS_1 + S_1^{-3}\, dS_1\, dS_1 = (\sigma_1^2 - \mu_1) S_1^{-1}\, dt - \sigma_1 S_1^{-1}\, dw_1$$

a bit of manipulation gives

$$dX = (\text{stuff})\, dt - \sigma_1 X\, dw_1 + \sigma_2 X\, dw_2.$$

Under $\overline{Q}$ the process must be drift-free, with the same volatility; therefore

$$dX = -\sigma_1 X\, d\tilde{w}_1 + \sigma_2 X\, d\tilde{w}_2$$

where $\tilde{w}_1$ and $\tilde{w}_2$ are independent $\overline{Q}$-Brownian motions. Now assume that σ_1, σ_2 and ρ are constant, and observe that $\sigma_1 \tilde{w}_1 + \sigma_2 \tilde{w}_2$ has mean value 0 and variance $\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2$. A moment's reflection reveals that it is in fact $\hat{\sigma}$ times a $\overline{Q}$ Brownian motion, where $\hat{\sigma} = (\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2)^{1/2}$. So $X_T = S_2(T)/S_1(T)$ is lognormal under $\overline{Q}$, with mean value $S_2(0)/S_1(0)$ and volatility $\hat{\sigma}$. By Equation (1) of Section 2 we conclude that

$$E_{\overline{Q}}[(S_2/S_1 - 1)_+] = \frac{S_2(0)}{S_1(0)} N(d_1) - N(d_2)$$

where

$$d_1 = \frac{\ln[S_2(0)/S_1(0)] + \frac{1}{2}\hat{\sigma}^2 T}{\hat{\sigma}\sqrt{T}}, \quad d_2 = \frac{\ln[S_2(0)/S_1(0)] - \frac{1}{2}\hat{\sigma}^2 T}{\hat{\sigma}\sqrt{T}}.$$

The value of the exchange option is therefore

$$V_0 = S_2(0) N(d_1) - S_1(0) N(d_2).$$

The options with payoff $\min(S_1, S_2)$ and $\max(S_1, S_2)$ are easily priced using this result. Indeed, we have

$$\min(S_1, S_2) = S_2 - (S_2 - S_1)_+$$

and

$$\max(S_1, S_2) = S_1 + (S_2 - S_1)_+$$

so the value of the min payoff is $S_2(0) - V_0$ and the value of the max payoff is $S_1(0) + V_0$.

****

Quantos. A quanto option is an option on a stock price quoted in a foreign currency. For example, if S is the price of a stock in British pounds, a a quanto call on this underlying with strike K has payoff $(S_T - K)_+$ *dollars*. Clearly we are again dealing with two sources of randomness: the stock price S (in pounds) and the exchange rate C (expressed as dollars per pound). It is again convenient to use viewpoint (2): we suppose S and C satisfy

$$dS = \mu_S S\,dt + \sigma_S S\,dw_S$$

and

$$dC = \mu_C C\,dt + \sigma_C C\,dw_C$$

and we suppose the Brownian motions w_S and w_C have correlation ρ. As in Section 2, the risk-free rate in pounds is u and the risk-free rate in dollars is r, so the pound money market account D and the dollar money market account B satisfy $dD = uD\,dt$ and $dB = rB\,dt$.

The pound investor sees tradeables D (the pound money market account), S (the stock), and BC^{-1} (the dollar money market account, valued in pounds). His risk-neutral measure – let's call it $\overline{Q}$, reserving the notation Q for the dollar investor's risk-neutral measure – has the property that S/D and B/CD are $\overline{Q}$-martingales. We could work out what $\overline{Q}$ is if we wanted to. But we don't need to, so we won't.

The dollar investor sees tradeables CD (the pound money market account, valued in dollars), CS (the stock, valued in dollars), and B. His risk-neutral measure Q has the property that CD/B and CS/B are Q-martingales. Since the volatilities are unchanged when we pass from the subjective measure to $\overline{Q}$ or Q, the fact that CD/B is a Q-martingale tells us that

$$d(CD/B) = \sigma_C(CD/B)\,d\tilde{w}_C$$

where $\tilde{w}_C$ is a Q-Brownian motion.

We're interested in an option whose dollar payoff is a function of S_T (for example a quanto call, with payoff $f(S_T) = (S_T - K)_+$). So the crucial question is: what is the risk-neutral process satisfied by S? Since changing the measure leaves the volatility invariant, it must take the form

$$dS = \alpha S\,dt + \sigma_S S\,d\tilde{w}_S$$

where $\tilde{w}_S$ is a Q-Brownian motion, whose correlation with $\tilde{w}_C$ is ρ, i.e. the same as the correlation between w_S and w_C. Our task is to find α. Our tool is the fact that CS/B is a Q-martingale. Since $CS/B = (S/D)(CD/B)$, Ito's formula gives

$$\begin{aligned} d(SC/B) &= (S/D)\,d(CD/B) + (CD/B)\,d(S/D) + d(S/D)\,d(CD/B) \\ &= (SC/B)\sigma_C\,d\tilde{w}_C + (SC/B)[(\alpha - u)\,dt + \sigma_S S\,d\tilde{w}_S] + (SC/B)\sigma_S\sigma_C\,d\tilde{w}_S\,d\tilde{w}_C \\ &= (SC/B)(\alpha - u + \rho\sigma_S\sigma_C)\,dt + (SC/B)\sigma_C\,d\tilde{w}_C + (SC/B)\sigma_S\,d\tilde{w}_S \end{aligned}$$

since $\tilde{w}_S$ and $d\tilde{w}_S\,d\tilde{w}_C = \rho\,dt$. For this to be a Q-martingale the drift must vanish, so

$$\alpha = u - \rho\sigma_S\sigma_C.$$

Thus if r is constant, the time-0 value (in dollars) of the quanto with payoff $f(S_T)$ (dollars) is

$$e^{-rT}E_Q[f(S_T)]. \tag{4}$$

If σ_C, σ_S, and ρ are constant then the distribution of S_T (viewed using Q) is lognormal with mean value

$$E_Q[S_T] = e^{(u-\rho\sigma_S\sigma_C)T}S_0$$

(this is the quanto forward rate, i.e. the choice of k for which payoff $S_T - k$ has initial value 0) and volatility $\sigma_S\sqrt{T}$. For a quanto put or a quanto call, we easily deduce a Black-Scholes-like formula (for example: use eqn 1 of Section 2 for the case of a call).