Continuous Time Finance Notes, Spring 2004 – Section 2, *Jan. 28, 2004*
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course Continuous Time Finance.

In Section 1 we discussed how Girsanov's theorem and the martingale representation theorem tell us how to price and hedge options. This short section makes that discussion concrete by applying it to (a) options on a stock which pays dividends, and (b) options on foreign currency. We close with a brief discussion of Siegel's paradox. For topics (a) and (b) see also Baxter and Rennie's sections 4.1 and 4.2 – which are parallel to my discussion, but different enough to be well worth reading and comparing to what's here. For a discussion of Siegel's paradox and some related topics, see chapter 1 of the delightful book *Puzzles of Finance: Six Practical Problems and their Remarkable Solutions* by Mark Kritzman (J. Wiley & Sons, 2000, available as an inexpensive paperback).

****

Options on a stock with dividend yield. You probably already know from Derivative Securities how to price an option on a stock with continuous dividend yield q. If the stock is lognormal with volatility σ and the risk-free rate is (constant) r then the "risk-neutral process" is $dS = (r-q)S\,dt + \sigma S\,dw$, and the time-0 value of an option with payoff $f(S_T)$ and maturity T is $e^{-rT}E_{\mathrm{RN}}[f(S_T)]$. From the SDE we get $dE_{\mathrm{RN}}[S]/dt = (r-q)E_{\mathrm{RN}}[S]$, so $E_{\mathrm{RN}}[S](T) = e^{(r-q)T}S_0$. This is the *forward price*, i.e. the unique choice of k such that a forward with strike k and maturity T has initial value 0. (Proof: apply the pricing formula to $f(S_T) = S_T - k$.) When the option is a call, we get an explicit valuation formula using the fact that if X is lognormal with mean $E[X] = F$ and volatility s (defined as the standard deviation of $\log X$), then

$$E[(X-K)_+] = FN(d_1) - KN(d_2) \tag{1}$$

with

$$d_1 = \frac{\ln(F/K) + s^2/2}{s}, \quad d_2 = \frac{\ln(F/K) - s^2/2}{s}.$$

There is of course a similar formula for a call (easily deduced by put-call parity).

What does it mean that the "risk-neutral process is $dS = (r-q)S\,dt + \sigma S\,dw$?" What happens when the dividends are paid at discrete times? We can clarify these points by using the framework of Section 1. Remember the main points: (a) there is a unique "equivalent martingale measure" Q, obtained by an application of Girsanov's theorem; and (b) this Q is characterized by the property that the value V_t of any tradeable asset satisfies $V_t/B_t = E[V_T/B_T \mid \mathcal{F}_t]$ where B_t is the value of a risk-free money-market account, i.e. it solves $dB = rB\,dt$ with $B(0) = 1$. Put differently: V/B is a Q-martingale.

Suppose the subjective stock process is $dS = \mu S\,dt + \sigma S\,dw$, and it pays dividends at constant rate q. If we hold the stock, we receive the dividend yield as well. So the stock itself is not a tradeable asset; but the stock *with dividends reinvested* is a tradeable. The value of this asset is $X_t = S_t D_t$ where $dD = qD\,dt$. Since the SDE for D has no dw term, Ito's formula becomes the ordinary Leibniz rule $dX = SdD + DdS$. Thus, after a bit of manipulation, the SDE for X is

$$dX = (\mu + q)X\,dt + \sigma X\,dw \tag{2}$$

A similar calculation gives the SDE satisfied by its discounted value $Y = X/B$:

$$dY = (\mu + q - r)Y\,dt + \sigma Y\,dw.$$

The preceding SDE's all use the original w, which is a Brownian motion in the original, subjective probability.

Now we apply Girsanov's theorem. Remember how it works: Girsanov changes the drift but not the volatility. The risk-neutral measure Q has the property that Y is a Q-martingale. So the SDE for Y must be

$$dY = \sigma Y\,d\tilde{w}$$

where $\tilde{w}$ is a Q-Brownian motion. To get the SDE for S, we observe that recall that $S_t = Y_t B_t / D_t$ so

$$dS = (r - q)S\,dt + \sigma S\,d\tilde{w}.$$

This is the "risk-neutral process" alluded to above, and the "risk-neutral" expectations E_{RN} are simply to be taken using the measure Q.

It was not important for the preceding calculations that the interest rate (or the dividend rate) be constant. For example, q could be a function of S. If r is function of time however the option value is not $e^{-rT} E_{\text{RN}}[f(S_T)]$ but $\exp\left(-\int_0^T r(s)\,ds\right) E_{\text{RN}}[f(S_T)]$. (If r is random the discount factor must be brought inside the expectation; but normally the randomness of r would be independent of, or at least not fully correlated with, that of S; we'll discuss problems with more than one source of randomness soon.) Conclusion: the equation for the risk-neutral process is actually quite general; the specialization to constant volatility, constant risk-free rate, and constant dividend rate is needed only if we desire an explicit formula for the value of the option.

What if the dividends are paid at discrete times? Suppose at each time $T_1, T_2, \ldots$ the stock pays a dividend equal to fraction q of the stock price. Then S_t is lognormal between the dividend dates (say, $dS = \mu S\,dt + \sigma S\,dw$) but at each dividend date S must decrease by exactly the value of the dividend, i.e. at T_j it jumps from S_{T_j} to $(1-q)S_{T_j}$. But as above, the stock itself is not a tradeable; rather, the convenient tradeable is the stock with all dividends reinvested. Its value X now satisfies $dX = \mu X\,dt + \sigma X\,dw$. The risk-neutral Q is defined by the property that $Y = X/B$ is a martingale; a calculation parallel to the one done above shows that

$$dX = r\,X dt + \sigma X\,d\tilde{w}$$

where $\tilde{w}$ is a Q-martingale. Now, if N dividend dates occur between time 0 and time T, then the final-time stock price is $S_T = (1-q)^N X_T$. Thus under the risk-neutral probability S_T is lognormal, with mean $(1-q)^N e^{rT} S_0$ and volatility $\sigma\sqrt{T}$. We can again get explicit prices for calls by using the formula (1).

****

Options on an exchange rate. You probably also learned in Derivative Securities how to price an option on a foreign currency rate. Here's the basic setup: suppose the US dollar risk-free rate is r; the British pound risk-free rate is q, and the dollar value of one pound is lognormal, i.e. the exchange rate C_t with units dollars/pound satisfies $dC = \mu C\,dt + \sigma C\,dw$. To a dollar investor the pound looks like a "stock with continuous dividend yield q." So no further work is needed: from our calculation with continuous dividend yield, the risk-free process is

$$dC = (r - q)C\,dt + \sigma C\,d\tilde{w}. \tag{3}$$

Here $\tilde{w}$ is a Q-Brownian motion, where Q is the dollar investor's risk-neutral measure. A typical application would be to price an option giving its holder the right to buy a British pound at time T for K dollars. Its payoff (to the dollar investor) is $(C_T - K)_+$. Since C_T is lognormal under the risk-neutral probability, with mean $e^{(r-q)T}C_0$ and volatility $\sigma\sqrt{T}$, we get an explicit price from (1).

But now a new issue arises. What about an investor who thinks in pounds? He can do a similar calculation of course. Will his calculation be consistent with that of the dollar investor? We expect so, since an inconsistency would lead to arbitrage. But let's check the consistency explicitly.

We start by spelling out the pound investor's calculation. His money-market discount factor is D_t not B_t (here $dD = qD\,dt$ and $dB = rB\,dt$). His exchange rate is $1/C$ not C. By Ito, if $dC = \mu C\,dt + \sigma C\,dw$ then

$$d(1/C) = -C^{-2}dC + C^{-3}dCdC = (-\mu + \sigma^2)(1/C)\,dt - \sigma(1/C)\,dw.$$

What is the pound investor's risk-neutral measure? Well, it isn't Q! We find it by considering the obvious stochastic tradeable: the value in pounds of the dollar money-market account. Arguing as for (2), its value $\overline{X} = BC^{-1}$ satisfies

$$d\overline{X} = (r - \mu + \sigma^2)\overline{X}\,dt - \sigma\overline{X}\,dw$$

and the associated discounted value $\overline{Y} = \overline{X}/D$ satisfies

$$d\overline{Y} = (r - q - \mu + \sigma^2)\overline{Y}\,dt - \sigma\overline{Y}\,dw.$$

Thus pound-investor's risk-neutral measure $\overline{Q}$ has the property that

$$d\overline{Y} = -\sigma\overline{Y}\,d\overline{w}$$

where $\overline{w}$ is a $\overline{Q}$-Brownian motion. In terms of $\overline{w}$ the SDE for $C^{-1} = D\overline{Y}/B$ is

$$d(1/C) = (q - r)(1/C)\,dt - \sigma(1/C)\,d\overline{w} \tag{4}$$

and by Ito's lemma the SDE for C is therefore

$$dC = (r - q + \sigma^2)C\,dt + \sigma C\,d\overline{w}.$$

Comparing this with (3) we see that

$$d\tilde{w} = d\overline{w} + \sigma\, dt.$$

Recall that $\overline{w}$ is a $\overline{Q}$-Brownian motion, while $\tilde{w}$ is a Q-Brownian motion. Thus $\overline{Q}$ is indeed different from Q. From Girsanov's theorem (with $\overline{Q}$ and $\overline{w}$ in place of P and w) we know the relation between them:

$$\frac{dQ}{d\overline{Q}} = e^{-\sigma\overline{w}(T) - \frac{1}{2}\sigma^2 T} \tag{5}$$

with the usual convention that $\overline{w}(0) = \tilde{w}(0) = 0$.

OK, now we're ready to check for consistency. Consider an option whose payoff in dollars is f. (For example: an option to buy a pound for K dollars at time T, whose payoff is $f = (C_T - K)_+$.) Note that its payoff in pounds is $C_T^{-1} f$. To the dollar investor its value at time 0 is

$$e^{-rT} E_Q[f] \quad \text{dollars.}$$

To the pound investor its value at time 0 is

$$e^{-qT} E_{\overline{Q}}[C_T^{-1} f] \quad \text{pounds.}$$

To show the prices are consistent, we must verify that

$$C_0 e^{-qT} E_{\overline{Q}}[C_T^{-1} f] = e^{-rT} E_Q[f].$$

Since

$$E_Q[f] = E_{\overline{Q}}\left[\frac{dQ}{d\overline{Q}} f\right]$$

it will be sufficient to show that

$$C_0 e^{-qT} C_T^{-1} = e^{-rT} \frac{dQ}{d\overline{Q}}. \tag{6}$$

But solving the SDE (4) we have

$$C_T^{-1} = C_0^{-1} e^{(q - r - \frac{1}{2}\sigma^2)T - \sigma\overline{w}(T)},$$

which makes the left hand side of (6) explicit in terms of $\overline{w}(T)$. Equation (5) makes the right hand side explicit. Examining the resulting expressions we find that the desired consistency condition (6) is indeed valid.

****

It is worth digressing to mention *Siegel's paradox.* Recall that under the dollar investor's risk-neutral process

$$dC = (r - q)C\, dt + \sigma C\, d\tilde{w},$$

and that the forward rate (in dollars per pound) is the expected risk-neutral mean

$$F = E_Q[C_T] = e^{(r-q)T} C_0.$$

The pound investor has a similar calculation: under his risk-neutral process

$$d(1/C) = (q - r)(1/C)\,dt - \sigma(1/C)\,d\overline{w}$$

so his forward rate (in pounds per dollar) is

$$\overline{F} = E_{\overline{Q}}[C_T^{-1}] = e^{(q-r)T}C_0^{-1}.$$

They are consistent: $\overline{F} = 1/F$, as must be the case.

This might at first seem surprising, since the function $1/x$ is convex, which implies by Jensen's inequality that

$$(E[C])^{-1} < E[C^{-1}]$$

when both means are taken with respect to the same probability distribution and C is not deterministic. But the means defining F and $\overline{F}$ involve two different probability distributions: Q and $\overline{Q}$.

What about the subjective probability distribution P. Is it plausible that $E_P[C_T] = e^{(r-q)T}C_0$? By now we know the answer: it's possible, though unlikely, since this holds only for a particular choice of the drift μ. What would the financial implication be? Well, consider the following investment strategy: (a) borrow X dollars at time 0, at interest rate r; (b) convert them to X/C_0 pounds; (c) invest these pounds at interest rate q; (d) liquidate the position at time T. The amount the investor realizes by liquidation is $X[(C_T/C_0)e^{qT} - e^{rT}]$. So his expected outcome is 0 exactly if $E_P[C_T] = e^{(r-q)T}C_0$, in other words *if the forward exchange rate is an unbiased estimate of the future spot exchange rate.*

Siegel's paradox is the observation that the forward exchange rate cannot, in general, be an unbiased estimate of the future spot exchange rate. More precisely: this cannot be true simultaneously for the dollar investor and the pound investor, since $(E_P[C])^{-1} < E_P[C^{-1}]$. If the property holds for the dollar investor, then it is false for the pound investor. This is also clear from our calculation of the risk-neutral measures Q and $\overline{Q}$: if $\sigma \neq 0$ then they are different, so they can't both be equal to the subjective probability.