3 Option Pricing

We are now going to apply our continuous-time methods to the pricing of nonstandard securities. In particular, we will consider the class of derivative securities known as options in this section of the notes. We will derive the famous Black-Scholes option pricing model in this section, for which Scholes recently won the Nobel prize (Black died a few years ago). The original paper by Black and Scholes is the supplementary reading that goes with this lecture note. The paper is a model of how all good papers should be written. There is not a lot of extraneous information in there - it's pretty short. However, its influence has been enormous. They seem to have maximized the ratio of economics to math in their paper. If you do not know much about options, you may want to talk to me about other readings.

3.1 The Black-Scholes Model

We will assume that stock prices follow geometric Brownian motion,

$$\frac{dS}{S} = \alpha dt + \sigma dW \tag{64}$$

We also know that a particular call option's price is a function of time and its underlying stock price,

$$C = C(S,t). \tag{65}$$

Since the call option's value is a function of time and the stock price, we know that it will follow a diffusion process. We posit (and later verify) that the call option's value follows geometric Brownian motion,

$$\frac{dC(S,t)}{C} = \hat{\alpha} dt + \hat{\sigma} dW. \tag{66}$$

Geometric Brownian motion makes sense for the call price because call prices cannot be negative. Now Using Ito's lemma we can find an expression for dC,

$$dC = C_S dS + C_t dt + \frac{1}{2} C_{SS} dS^2, \tag{67}$$

and we can take the expectation of dC,

$$E[dC] = C_S \alpha S dt + C_t dt + \frac{1}{2} C_{SS} S^2 \sigma^2 dt. \tag{68}$$

In order to verify that the call option's price follows Geometric Brownian motion we need to solve for $\hat{\alpha}$ and $\hat{\sigma}$. If we can find expressions for both of these terms then the call price really does follow GBM. We solve for the terms using our expression for dC and $E[dC]$. Beginning with $\hat{\alpha}$,

$$\hat{\alpha} = \frac{E[dC]}{Cdt} = \frac{C_S \alpha S + C_t + \frac{1}{2} C_{SS} S^2 \sigma^2}{C}. \tag{69}$$

The term for $\hat{\sigma}$ is

$$\hat{\sigma} = \frac{dC - E[dC]}{CdW} = \frac{C_S(dS - E[dS]) + \frac{1}{2} C_{SS}(dS^2 - E[dS^2])}{CdW}, \tag{70}$$

which simplifies to

$$\hat{\sigma} = \frac{C_S S \sigma}{C}. \tag{71}$$

So we have shown that the call option price follows GBM with particular values for $\hat{\alpha}$ and $\hat{\sigma}$. Now we want to think about a portfolio, or a position that involves putting w dollars in the call option and putting $(1 - w)$ in the stock. The stochastic proces that will describe the value of our position, V, through time is another GBM process,

$$\frac{dV}{V} = w \frac{dC}{C} + (1 - w) \frac{dS}{S}. \tag{72}$$

Using the processes for dC and dS, we can be more specific about the process driving the value of our position,

$$\frac{dV}{V} = [w\hat{\alpha} + (1-w)\alpha]dt + [w\hat{\sigma} + (1-w)\sigma]dW. \tag{73}$$

We choose a value for w that eliminates all the risk in our position in the next instant by setting the coefficient on the dW term above to zero,

$$w^* = \frac{\sigma}{\sigma - \hat{\sigma}}. \tag{74}$$

Choosing this particular w gives us the risk-free process

$$\left.\frac{dV}{V}\right|_{w^*} = [w^*\hat{\alpha} + (1-w^*)\alpha]dt = rdt \tag{75}$$

We have apparently done a very curious thing. By carefully choosing w, we can make the coefficient on the dW term go to zero. If we set up our position just right, we can eliminate all of the risk in the portfolio. This is called a dynamic hedge. The risk-free position that we create will only be risk-free over the next instant - to maintain a risk-free position, we will have to constantly adjust w as both S and C move around through time. This is the continuous-time analog to one node of the binomial tree that we saw in the previous set of notes. The continuous-time hedge is the basis for most arbitrage pricing in continuous time. To effectively hedge continuously you must have very small transactions costs. Some researchers derive no-arbitrage results with transactions costs - they find that rather than arriving at a unique no-arbitrage price, they get a range of possible option values that don't admit arbitrage. This is a highly techincal research area that is dominated by statisticians and mathematicians.

We have shown how to convert our position to a risk-free security through a dynamic hedge. Since this position is riskless over the next instant, dt, its payoff must be equal to the payoff of a risk-free security over the next instant (by the law of one price),

$$\frac{\sigma}{\sigma - \hat{\sigma}}\hat{\alpha} - \frac{\hat{\sigma}}{\sigma - \hat{\sigma}}\alpha = r. \tag{76}$$

Doing a little bit of algebra, we can rearrange this to be,

$$\hat{\alpha} - r = \frac{\hat{\sigma}}{\sigma}(\alpha - r). \tag{77}$$

Substituting the expressions for $\hat{\alpha}$ and $\hat{\sigma}$ into this expression,

$$\frac{C_S \alpha S + C_t + \frac{1}{2}C_{SS}\sigma^2 S^2}{C} - r = \frac{\frac{C_S S \sigma}{C}}{\sigma}(\alpha - r), \tag{78}$$

which further simplifies to

$$C_S \alpha S + C_t + \frac{1}{2}C_{SS}\sigma^2 S^2 - rC = C_S S(\alpha - r). \tag{79}$$

Dropping $C_S S \alpha$ from both sides,

$$C_S S r + C_t + \frac{1}{2}C_{SS}\sigma^2 S^2 - rC = 0, \tag{80}$$

which is the equation that we have been solving for, a famous differential equation known as the heat transfer equation. Its solution was derived many years ago.

One interesting feature of this differential equation is that it does not depend on investor preferences in any direct way. In particular, it is not a function of the expected stock return, α. This is important because it provides some intuition for the "risk-neutral" pricing results that we will derive later. Since the option value does not depend on attitudes about risk, it is the same regardless of the level of risk aversion in the economy. It is even the same in an economy characterized by risk-neutral agents. In a risk-neutral economy, the growth rate of the stock return is just the risk-free rate.

If we take the expected value of the option payoff after assuming that $\alpha = r_f$, we should get the right option price. We will save this exercise for another portion of the notes.

For the Black-Scholes formula, we solve the heat transfer equation subject to the following boundary conditions,

$$\begin{aligned} \text{Boundary conditions}: \; C(0,t) &= 0 \\ C(S,T) &= \max[S-X,0] \\ C(S,t) &\leq S. \end{aligned} \tag{81}$$

Unfortunately, you will just have to take it on faith (for now) that the solution to the heat transfer equation is actually the Black-Scholes formula. For reference purposes, the equation can be written as

$$C(S,t) = S\mathcal{N}(d_1) - Xe^{-r\tau}\mathcal{N}(d_2), \tag{82}$$

where

$$d_1 = \frac{\ln(S/X) + (r + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \quad d_2 = d_1 - \sigma\sqrt{\tau}, \tag{83}$$

and τ is the option's time to maturity, $\mathcal{N}(\cdot)$ is the cumulative normal distribution, and X is the strike price.

This is not the last time that we will see the Black-Scholes formula. We will derive it again with "change of measure" methods when we talk about the absence of arbitrage in continuous-time. Let's talk about the model a little more right now to get a little more understanding of how it works.

3.2 Solving for Delta

An important quantity for the Black-Scholes model is what is commonly called Δ, or the sensitivity of the option's call price to changes in the value of the option's underlying security. As we stated previously, in the Black-Scholes framework, a call option's Δ is just equal to $\mathcal{N}(d_1)$. We should prove this rigorously. We start by taking the derivative of the Black-Scholes call option price,

$$\Delta = \frac{\partial C}{\partial S} = \mathcal{N}(d_1) + S\frac{\partial \mathcal{N}(d_1)}{\partial S} - Xe^{-r\tau}\frac{\partial \mathcal{N}(d_2)}{\partial S}. \tag{84}$$

Remembering that we can express d_1 as,

$$d_1 = \frac{\ln(S/X)}{\sigma\sqrt{\tau}} + \frac{r}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau}, \tag{85}$$

we can take the derivative of d_1 with respect to S,

$$\frac{\partial d_1}{\partial S} = \frac{1}{S\sigma\sqrt{\tau}}. \tag{86}$$

Now using the chain rule,

$$\frac{\partial \mathcal{N}(d_1)}{\partial S} = \frac{\partial \mathcal{N}(d_1)}{\partial d_1}\frac{\partial d_1}{\partial S} = \frac{f(d_1)}{S\sigma\sqrt{\tau}}. \tag{87}$$

Since d_2 is just d_1 minus $\sigma\sqrt{\tau}$, a similar result holds for $\mathcal{N}(d_2)$,

$$\frac{\partial \mathcal{N}(d_2)}{\partial S} = \frac{f(d_1 - \sigma\sqrt{\tau})}{S\sigma\sqrt{\tau}}. \tag{88}$$

In both of these expressions, $f(\cdot)$ is just the standard normal distribution,

$$f(x) = \frac{e^{-X^2/2}}{\sqrt{2\pi}}. \tag{89}$$

Returning to our expression for Δ, we can now state that

$$\Delta = \frac{\partial C}{\partial S} = \mathcal{N}(d_1) + \frac{f(d_1)}{\sigma\sqrt{\tau}} - \frac{Xe^{-r\tau}}{S}\frac{f(d_1 - \sigma\sqrt{\tau})}{\sigma\sqrt{\tau}}. \tag{90}$$

Therefore, our result ($\Delta = \mathcal{N}(d_1)$) will hold as long as

$$f(d_1) = \frac{Xe^{-r\tau}}{S} f(d_1 - \sigma\sqrt{\tau}). \tag{91}$$

To see that this condition is true, we express the ratio of $f(d_1)$ to $f(d_2)$ as

$$\frac{e^{-d_1^2/2}}{e^{-(d_1 - \sigma\sqrt{\tau})^2/2}} = \frac{Xe^{-r\tau}}{S}. \tag{92}$$

Next, we express the denominator as

$$e^{-(d_1 - \sigma\sqrt{\tau})/2} = e^{-(d_1^2 - 2d_1\sigma\sqrt{\tau} + \sigma^2\tau)/2}, \tag{93}$$

which yields the condition

$$e^{-d_1\sigma\sqrt{\tau} + \sigma^2\tau/2} = \frac{Xe^{-r\tau}}{S}. \tag{94}$$

Finally, substituting the definition of d_1 into our condition,

$$e^{-\ln(S/X) - (r + \sigma^2/2)\tau + \sigma^2\tau/2} = \frac{Xe^{-r\tau}}{S}, \tag{95}$$

which we can verify as true. Thus, we have shown that

$$\Delta = \frac{\partial C}{\partial S} = \mathcal{N}(d_1). \tag{96}$$

What do we use this Δ term for? For one thing, Δ tells us how to perform our continuous-time hedge. As we stated in the notes on the binomial tree pricing model, we always want to hold Δ shares of stock for each written call option in a risk-free portfolio. As the components of d_1 change, Δ will change. We can update our position by constantly checking if we have Δ shares of stock per short call.

A second reason for caring about Δ is that it tells us about the expected return on an option. To see this, we will define another quantity known as omega,

$$\Omega = \frac{\partial C}{\partial S}\frac{S}{C} = \frac{S\mathcal{N}(d_1)}{S\mathcal{N}(d_1) - Xe^{-r\tau}\mathcal{N}(d_2)} > 1. \tag{97}$$

Omega is the elasticity of the call option's price to changes in the value of the call's underlying security. Thus, If the underlying asset's value changes by 3% because the market moves 3%, the call option's price should move by Ω times 3%. In other words, it can be shown (you will basically show in a homework) that

$$\beta_C = \Omega\beta_S, \tag{98}$$

where β_C is the call option's market beta and β_S is the stock's market beta. Since Ω is always greater than one (as demonstrated above), the expected return of a call option written on an asset with a positive beta is always higher than the expected return of the underlying asset in a Black-Scholes/CAPM world. There are, of course, much more general things that can be said about option returns. However, the intuition behind options returns arguments is the Black-Scholes/CAPM case.

3.3 CAPM-Based Derivation

We can, in fact, use our results about beta to derive the Black-Scholes formula in a different way. If we define the risk premium on the market portfolio to be γ, then the CAPM stipulates that

$$\begin{aligned} E\left[\tfrac{dS}{S}\right] &= rdt + \gamma\beta_S dt \\ E\left[\tfrac{dC}{C}\right] &= rdt + \gamma\beta_C dt. \end{aligned} \tag{99}$$

Using (98), we can express the expected change in the call price as,

$$E[dC] = rCdt + \gamma C_S S\beta_S dt. \tag{100}$$

Of course, from Ito's lemma we also know that dC can be expressed as

$$dC = C_S dS + C_t dt + \frac{1}{2} C_{SS} dS^2, \tag{101}$$

so the expected change in the call price is,

$$E[dC] = C_S[rS + \gamma S \beta_S]dt + C_t dt + \frac{1}{2} C_{SS} S^2 \sigma^2 dt \tag{102}$$

Setting equations (100) and (102) equal to eachother and dividing through by dt yields the differential equation,

$$C_S Sr + C_t + \frac{1}{2} C_{SS} \sigma^2 S^2 - rC = 0, \tag{103}$$

which is equivalent to (80).

So what is the point of this derivation? It is nice to see that the Black-Scholes model is entirely consistent with the continuous-time CAPM. This makes sense - they basically make the same assumptions. This also gives some intuition about risk-neutral pricing, which we haven't really discussed yet. While it is true that there exists a probability measure under which all security prices are just discounted expected payoffs, that probability measure is not arbitrarily chosen. For the CAPM and risk-neutral pricing to be compatible, it must be that the probability measure assigns a higher probability to bad states of the world than the "true" subjective probability. In a CAPM world, bad states of the world correspond to low market returns.

We have spent a great deal of time discussing option pricing. This may seem unwarranted because options are just one type of financial instrument available to market participants. However, Black and Scholes point out at the end of their article that lots of securities actually have option-like characteristics. Consider, for example, the stock of a company that has some level of debt outstanding that must be paid in a lump sum. If the company plans to liquidate immediately after paying off its debt then the company's stock can be thought of as an option on the value of the firm. If the firm is sufficiently profitable then the option will expire in the money and stockholders will be paid. Otherwise, all the firm's value will go to bondholders and the option (equity) will expire worthless.

3.4 Homework Problems

1. Assume that returns are driven by a k-factor model, so that

$$dS_i/S_i = \alpha_i dt + \beta_{i1}[df_{1t}/f_{1t}] + \ldots + \beta_{ik}[df_{kt}/f_{kt}] + \sigma_i dW_i \tag{104}$$

   and assume that each factor follows geometric Brownian motion with zero drift and a volatility of one. Let the Wiener process of each factor (dW_{fj}) be independent of the corresponding process of each other factor and let it be independent of dW_i for all i. Assume also all of the assumptions that make the Black-Scholes formula valid. Show that the jth factor beta of a European call option on stock i is equal to Ω_i times the stock's factor beta, β_{ij}.

2. Suppose that a stock pays out a continuous dividend at rate $\lambda S dt$. Derive the differential equation for the option's value under this assumption and write down what the option's value should be. It should look something like the Black-Scholes formula.

3. Verify by direct calculation that the differential equation (80), subject to the boundary conditions (81) and the constraint that $S > 0$ is solved by the Black-Scholes formula, (82).