Derivative Securities

Options on interest-based instruments: pricing of bond options, caps, floors, and swaptions. The most widely-used approach to pricing options on caps, floors, swaptions, and similar instruments is Black's model. We discuss how this model works, why it works, and when it is appropriate. The main alternative to Black's model is the use of a suitable interest rate tree (or a continuous-time model of the risk-free interest rate dynamics). We discuss briefly the calibration of a binomial tree – essentially, a special case of the Black-Derman-Toy model.

****

Black's model. My discussion of Black's model and its applications follows mainly chapter 20 of Hull, augmented by some examples from Clewlow and Strickland.

The essence of Black's model is this: consider an option with maturity T, whose payoff $\phi(V_T)$ is determined by the value V_T of some interest-related instrument (a discount rate, a term rate, etc). For example, in the case of a call $\phi(V_T) = (V_T - K)_+$. Black's model stipulates that

(a) the value of the option today is its discounted expected payoff.

No surprise there – it's the same principle we've been using all this time for valuing options on stocks. If the payoff occurs at time T then the discount factor is $B(0, T)$ so statement (a) means

$$\text{option value} = B(0, T)E_*[\phi(V_T)].$$

We write E_* rather than E_{RN} because in the stochastic interest rate setting this is *not* the risk-neutral expectation; we'll explain why E_* is different from the risk-neutral expectation later on. For the moment however, we concentrate on making Black's model computable. For this purpose we simply specify that (under the distribution associated with E_*)

(b) the value of the underlying instrument at maturity, V_T, is lognormal; in other words, V_T has the form e^X where X is Gaussian.

(c) the mean $E_*[V_T]$ is the forward price of V (for contracts written at time 0, with delivery date T).

We have not specified the variance of $X = \log V_T$; it must be given as data. It is customary to specify the "volatility of the forward price" σ, with the convention that

$$\log V_T \text{ has standard deviation } \sigma\sqrt{T}.$$

Notice that the Gaussian random variable $X = \log V_T$ is fully specified by knowledge of its standard deviation $\sigma\sqrt{T}$ and the mean of its exponential $E_*[e^X]$, since if X has mean m then $E_*[e^X] = \exp(m + \frac{1}{2}\sigma^2 T)$.

Most of the practical examples involve calls or puts. For a call, with payoff $(V_T - K)_+$, hypothesis (b) gives

$$E_*[(V_T - K)_+] = E_*[V]N(d_1) - KN(d_2)$$

where

$$d_1 = \frac{\log(E_*[V_T]/K) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\log(E_*[V_T]/K) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.$$

This is a direct consequence of the lemma we used long ago (in Section 5) to evaluate the Black-Scholes formula. Using hypotheses (a) and (c) we get

$$\text{value of a call} = B(0,T)[F_0 N(d_1) - KN(d_2)]$$

where F_0 is the forward price of V today, for delivery at time T, and

$$d_1 = \frac{\log(F_0/K) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\log(F_0/K) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.$$

These formulas are nearly identical to the ones we obtained in Sections 8 and 9 for pricing options on foreign currency rates, options on stocks with continuous dividend yield, and options on futures. The only apparent difference is the discount factor: in the constant interest rate setting of Sections 8 and 9 it was e^{-rT}; in the present stochastic interest rate setting it is $B(0,T)$.

It is by no means obvious that Black's formula is correct in a stochastic interest rate setting. We'll give the honest justification a little later. But here is a rough, heuristic justification. Since the value of the underlying security is stochastic, we may think of it as having its own lognormal dynamics. If we treat the risk-free rate as being constant then Black's formula can certainly be used. Since the payoff takes place at time T, the only reasonable constant interest rate to use is the one for which $e^{-rT} = B(0,T)$, and this leads to the version of Black's formula given above.

Black's model applied to options on bonds. Here is an example, taken from Clewlow and Strickland (section 6.6.1). Let us price a one-year European call option on a 5-year discount bond. Assume:

- The current term structure is flat at 5 percent per annum; in other words $B(0,t) = e^{-.05t}$ when t is measured in years.
- The strike of the option is 0.8; in other words the payoff is $(B(1,5) - 0.8)_+$ at time $T = 1$.
- The forward bond price volatility σ is 10 percent.

Then the forward bond price is $F_0 = B(0,5)/B(0,1) = .8187$ so

$$d_1 = \frac{\log(.8187/.8000) + \frac{1}{2}(0.1)^2(1)}{(0.1)\sqrt{1}} = 0.2814, \quad d_2 = d_1 - \sigma\sqrt{T} = 0.2814 - 0.1\sqrt{1} = .1814$$

and the discount factor for income received at the maturity of the option is $B(0,1) = .9512$. So the value of the call option now, at time 0, is

$$.9512[.8187N(.2814) - .8N(.1814)] = .0404.$$

Black's formula can also be used to value options on coupon-paying bonds; no new principles are involved, but the calculation of the forward price of the bond must take into account the coupons and their payment dates; see Hull's Example 20.1.

One should avoid using the same σ for options with different maturities. And one should never use the same σ for underlyings with different maturities. Here's why: suppose the option has maturity T and the underlying bond has maturity $T' > T$. Then the value V_t of the underlying is known at both $t = 0$ (all market data is known at time 0) and at $t = T'$ (all bonds tend to their par values as t approaches maturity). So the variance of V_t vanishes at both $t = 0$ and $t = T'$. A common model (if simplified) model says the variance of V_t is $\sigma_0^2 t(T' - t)$ with σ_0 constant, for all $0 < t < T'$. In this case the variance of V_T is $\sigma_0^2 T(T' - T)$, in other words $\sigma = \sigma_0\sqrt{T' - T}$. Thus σ depends on the time-to-maturity $T' - T$. In practice σ – or more precisely $\sigma\sqrt{T}$ – is usually inferred from market data.

Black's model applied to caps. A cap provides, at each coupon date of a bond, the difference between the payment associated with a floating rate and that associated with a specified cap rate, if this difference is positive. The ith caplet is associated with the time interval (t_i, t_{i+1}); if $R_i = R(t_i, t_{i+1})$ is the term rate for this interval, R_K is the cap rate, and L is the principal, then the ith caplet pays

$$L \cdot (t_{i+1} - t_i) \cdot (R_i - R_K)_+$$

at time t_{i+1}. Its value according to Black's formula is therefore

$$B(0, t_{i+1})L\Delta_i t[f_i N(d_1) - R_K N(d_2)].$$

Here $\Delta_i t = t_{i+1} - t_i$; $f_i = f_0(t_i, t_{i+1})$ is the forward term rate for time interval under consideration, defined by

$$\frac{1}{1 + f_i\Delta_i t} = \frac{B(0, t_{i+1})}{B(0, t_i)};$$

and

$$d_1 = \frac{\log(f_i/R_K) + \frac{1}{2}\sigma_i^2 t_i}{\sigma_i\sqrt{t_i}}, \quad d_2 = \frac{\log(f_i/R_K) - \frac{1}{2}\sigma_i^2 t_i}{\sigma_i\sqrt{t_i}} = d_1 - \sigma_i\sqrt{t_i}.$$

The volatilities σ_i must be specified for each i; in practice they are inferred from market data. The value of a cap is obtained by adding the values of its caplets.

A floor is to a cap as a put is to a call: using the same notation as above, the ith floorlet pays

$$L\Delta_i t(R_K - R_i)_+$$

at time t_{i+1}. Its value according to Black's formula is therefore

$$B(0, t_{i+1})L\Delta_i t[R_K N(-d_2) - f_i N(-d_1)]$$

where d_1 and d_2 are as above. The value of a floor is obtained by adding the values of its floorlets.

Here's an example, taken from Section 20.3 of Hull. Consider a contract that caps the interest on a 3-month, \$10,000 loan one year from now; we suppose the interest is capped at 8% per annum (compounded quarterly). This is a simple caplet, with $t_1 = 1$ year and $t_2 = 1.25$ years. To value it, we need:

- The forward term rate for a 3-month loan starting one year from now; suppose this is 7% per annum (compounded quarterly).
- The discount factor associated to income 15 months from now; suppose this is .9220.
- The volatility of the 3-month forward rate underlying the caplet; suppose this is 0.20.

With this data, we obtain

$$d_1 = \frac{\log(.07/.08) + \frac{1}{2}(0.2)^2(1)}{0.2\sqrt{1}} = -0.5677, \quad d_2 = d_1 - 0.2\sqrt{1} = -0.7677$$

so the value of the caplet is, according to Black's formula,

$$(.9220)(10,000)(1/4)[.07N(-.5677) - .08N(-.7677)] = 5.19 \text{ dollars.}$$

Problem 3 of HW6 is very much like the preceding example, except that the necessary data is partly hidden in financial jargon. Here's some help interpreting that problem. The relevant term rate is LIBOR 3-month rate, 9 months from now. The statement that "the 9-month Eurodollar futures price is 92" implies (if we ignore the difference between futures and forwards) that the present 3-month forward term rate for borrowing 9 months from now is 8% per annum. The statement that "the interest rate volatility implied by a 9-month Eurodollar option is 15 percent per annum" gives $\sigma = .15\%$.

Black's model applied to swaptions. A swaption is an option to enter into a swap at some future date T (the maturity of the option) with a specified fixed rate R_K. To be able to value it, we must first work a bit to represent its payoff.

Let R_{swap} be the par swap rate at time T, when the option matures. If $t_1 < \ldots < t_N$ are the coupon dates of the swap and $t_0 = T$ then R_{swap} is characterized (see Problem 2b of HW6) by

$$\sum_{i=1}^{N} B(T, t_i) R_{\text{swap}}(t_i - t_{i-1})L = (1 - B(T, t_N))L$$

where L is the notional principal. Moreover the left hand side is the value at time T of the fixed payments at rate R_{swap} while the right hand side is the value of the variable payments. Suppose the swaption gives its holder the right to pay the fixed rate R_K and receive the floating rate. Then it will be in the money if $R_{\text{swap}} > R_K$, and in that case its value to the holder at time T is

$$\begin{aligned} V_{\text{float}} - V_{\text{fixed}} &= (1 - B(T, t_N))L - \sum_{i=1}^{N} B(T, t_i) R_K (t_i - t_{i-1}) L \\ &= \sum_{i=1}^{N} B(T, t_i) R_{\text{swap}} (t_i - t_{i-1}) L - \sum_{i=1}^{N} B(T, t_i) R_K (t_i - t_{i-1}) L \\ &= (R_{\text{swap}} - R_K) \sum_{i=1}^{N} B(T, t_i)(t_i - t_{i-1}) L. \end{aligned}$$

The ith term is the payoff of an option on R_{swap} with maturity T and cash flow

$$L(t_i - t_{i-1})(R_{\text{swap}} - R_K)_+$$

received at time t_i. Black's formula gives the time-0 value of this option as

$$B(0, t_i) L (t_i - t_{i-1}) [F_{\text{swap}} N(d_1) - R_K N(d_2)]$$

where F_{swap} is the forward swap rate and

$$d_1 = \frac{\log(F_{\text{swap}}/R_K) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\log(F_{\text{swap}}/R_K) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.$$

The forward swap rate is obtained by taking the definition of the par swap rate, given above, and replacing $B(T, t_i)$ by the forward rate $F_0(T, t_i) = B(0, t_i)/B(0, T)$ for each i. To get the value of the swap itself we sum over all i:

$$\text{value of swap} = LA[F_{\text{swap}} N(d_1) - R_K N(d_2)] \quad \text{where } A = \textstyle\sum_{i=1}^{N} B(0, t_i)(t_i - t_{i-1}).$$

Here's an example, taken from Clewlow and Strickland section 6.6.1. Suppose the yield curve is flat at 5 percent per annum (continuously compounded). Let us price an option that matures in 2 years and gives its holder the right to enter a one-year swap with semiannual payments, receiving floating rate and paying fixed term rate 5 percent per annum. We suppose the volatility of the forward swap rate is 20% per annum.

The first step is to find the forward swap rate F_{swap}. It satisfies

$$\sum_{i=1}^{2} \frac{B(0, t_i)}{B(0, T)} F_{\text{swap}}(1/2) = \left(1 - \frac{B(0, t_2)}{B(0, T)}\right)$$

with $T = 2$, $t_1 = 2.5$, and $t_2 = 3.0$. Since the yield curve is flat at 5% compounded continuously, we have

$$\frac{B(0, 2.5)}{B(0, 2)} = e^{-(.05)(.5)} = .9753, \quad \frac{B(0, 3)}{B(0, 2)} = e^{-(.05)(1)} = .9512$$

and simple arithmetic gives $F_{\text{swap}} = .0506$, in other words 5.06%. Now

$$d_1 = \frac{\log(.0506/.0500) + \frac{1}{2}(0.2)^2(2)}{0.2\sqrt{2}} = 0.1587, \quad d_2 = d_1 - 0.2\sqrt{2} = -.0971,$$

and

$$\sum_{i=1}^{2} B(0,t_i)(t_i - t_{i-1}) = \tfrac{1}{2}(e^{-(.05)(2.5)} + e^{-(.05)(3)}) = .8716,$$

so the value of the swaption is

$$.8716L[.0506N(.1587) - .05N(-.0971)] = .0052L$$

where L is the notional principal of the underlying swap.

****

When and why is Black's model correct? Black's model is widely-used and appropriate for pricing European-style options on bonds, and analogous instruments such as caps, floors, and swaptions. It has two key advantages: (a) simplicity, and (b) directness. By simplicity I mean not that Black's model is easy to understand, but rather that it requires just one parameter (the volatility) to be inferred from market data. By directness I mean that we model the underlying instrument directly – the basic hypothesis of Black's model is the lognormal character of the underlying.

The main alternative to Black's model is the use of an interest-rate tree (or a continuous-time analogue thereof). Such a tree models the risk-neutral interest-rate process, which can then be used to value bonds of all types and maturities, and options of all types and maturities on these bonds. Interest-rate trees are not "simple" in the sense used above: to get started we must calibrate the entire tree to market data (e.g. the yield curve). And they are not "direct" in the sense used above: we are modeling the risk-neutral interest rate process, not the underlying instrument itself; thus there are two potential sources of modeling error: one in modeling the value of the underlying instrument, the other in modeling how the option's value depends on that of the underlying instrument.

The simplicity and directness of Black's model are also responsible for its disadvantages. Black's model must be used separately for each class of instruments – we cannot use it, for example, to hedge a cap using bonds of various maturities. For consistent pricing and hedging of multiple instruments one must use a more fundamental model such as an interest rate tree. Another restriction of Black's model: it can only be used for European-style options, whose maturity date is fixed in advance. Many bond options permit early exercise – sometimes American-style (permitting exercise at any time) but more commonly Bermudan (permitting exercise at a list of specified dates, typically coupon dates). Black's model does not allow for early exercise. Trees are much more convenient for this purpose, since early exercise is easily accounted for as we work backward in the tree.

Now we turn to the question of *why* Black's model is correct. The explanation involves "change of numeraire". (The following is a binomial-tree version of Hull's section 19.5.) The word numeraire refers to a choice of units.

Up to now our numeraire has been cash (dollars). Its growth as a function of time is described by the money-market account introduced in Section 9. The money-market account has balance is $A(0) = 1$ initially, and its balance evolves in time by $A_{\text{next}} = e^{r\delta t} A_{\text{now}}$. We are accustomed to finding the value f of a tradeable instrument (such as an option) by working backward in the tree using the risk-neutral probabilities. At each step this amounts to

$$f_{\text{now}} = e^{-r\delta t}[q f_{\text{up}} + (1-q) f_{\text{down}}]$$

where q and $1-q$ are the risk-neutral probabilities of the up and down states. As we noted in Section 9, this can be expressed as

$$f_{\text{now}}/A_{\text{now}} = E_{\text{RN}}[f_{\text{next}}/A_{\text{next}}],$$

and it can be iterated in time to give

$$f(t)/A(t) = E_{\text{RN}}[f(t')/A(t')] \quad \text{for } t < t'.$$

This is captured by the statement that "$f(t)/A(t)$ is a martingale relative to the risk-neutral probabilities."

But sometimes the money-market account is not the convenient comparison. In fact we may use *any* tradeable security as the numeraire – though when we do so we must also change the probabilities. Indeed, for any tradeable security g there is a choice of probabilities on the tree such that

$$\frac{f_{\text{now}}}{g_{\text{now}}} = \left[q_* \frac{f_{\text{up}}}{g_{\text{up}}} + (1-q_*) \frac{f_{\text{down}}}{g_{\text{down}}} \right].$$

This is an easy consequence of the two relations

$$f_{\text{now}} = e^{-r\delta t}[q f_{\text{up}} + (1-q) f_{\text{down}}] \quad \text{and} \quad g_{\text{now}} = e^{-r\delta t}[q g_{\text{up}} + (1-q) g_{\text{down}}],$$

which hold (using the risk-neutral q) since both f and g are tradeable. A little algebra shows that these relations imply the preceding formula with

$$q_* = \frac{q g_{\text{up}}}{q g_{\text{up}} + (1-q) g_{\text{down}}}.$$

(The value of q_* now varies from one binomial subtree to another, even if q was uniform throughout the tree.) Writing E_* for the expectation operator with weight q_*, we have defined q_* so that

$$f_{\text{now}}/g_{\text{now}} = E_*[f_{\text{next}}/g_{\text{next}}].$$

Iterating this relation gives (as in the risk-neutral case)

$$f(t)/g(t) = E_*[f(t')/g(t')] \quad \text{for } t < t';$$

in other words "$f(t)/g(t)$ is a martingale relative to the probability associated with E_*." In particular

$$f(0)/g(0) = E_*[f(T)/g(T)]$$

where T is the maturity of an option we may wish to price.

Let us apply this result to explain Black's formula. For simplicity we focus on options whose maturity T is also the time the payment is received. (This is true for options on bonds, not for caplets or swaptions – but the modification needed for caplets and swaptions is straightforward.) The convenient choice of g is then

$$g(t) = B(t, T).$$

Since $g(T) = 1$ this choice gives

$$f(0) = g(0)E_*[f(T)] = B(0, T)E_*[f(T)].$$

We shall apply this twice: once with f equal to the value of the underlying instrument, what we called V_t on page one of these notes; and a second time with f equal to the value of the option. The first application gives

$$E_*[V_T] = V_0/B(0, T)$$

and we recognize the right hand side as the *forward price* of the instrument. For this reason the probability distribution associated with this E_* is called *forward risk-neutral.* The second application gives

$$\text{option value} = B(0, T)E_*[\phi(V_T)]$$

where $\phi(V_T)$ is the payoff of the option – for example $\phi(V_T) = (V_T - K)_+$ if the option is a call.

This explains Black's formula, except for one crucial feature: the hypothesis that V_T is lognormal with respect to the distribution associated with E_* (the forward-risk-neutral distribution). This is of course only asserted in the continuous-time limit, and only if the risk-neutral interest rate process is itself lognormal. The assertion is most easily explained using continuous-time (stochastic differential equation) methods, and we will not attempt to address it here.

****

Interest rate trees. We have already discussed the limitations of Black's model. Consistent pricing and hedging of diverse interest-based instruments requires a different approach. So does the pricing of American or Bermudan options, which permit early exercise.

A typical alternative is the use of a binomial tree. We explained briefly how this works in Section 9, where we discussed how to pass from the tree to the various discount factors $B(0, t)$. Valuing options on the tree is also easy (just work backward). So is hedging (each binomial submarket is complete, so a risky instrument can be hedged using any pair of zero-coupon bonds). These topics are discussed very clearly in Chapter 15 of Jarrow & Turnbull and I recommend reading them there.

To make this a practical alternative, however, we must say something about how to calibrate the tree. Jarrow and Turnbull aren't very clear on this; for an excellent treatment see the

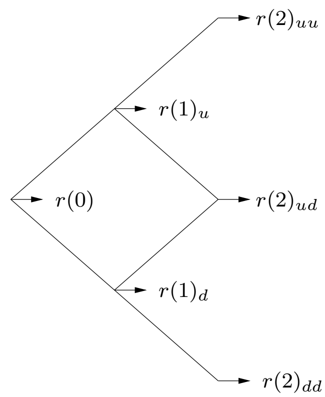

presentation of the Black-Derman-Toy model in Chapter 8 of Clewlow and Strickland. To give the general flavor, I'll discuss just the simplest version – calibration of the tree to the yield curve, with constant volatility – for a two-period tree of the type discussed in Section 9 (see the figure).

The basic ansatz is this:

- at time 0: $r(0) = a_0$;
- at time 1: $r(1)_u = a_1 e^{\sigma\sqrt{\delta t}}$ and $r(1)_d = a_1 e^{-\sigma\sqrt{\delta t}}$;
- at time 2: $r(2)_{uu} = a_2 e^{2\sigma\sqrt{\delta t}}$, $r(2)_{ud} = a_2$, and $r(2)_{dd} = a_2 e^{-2\sigma\sqrt{\delta t}}$.

More generally: at time j, the possible values of $r(j)$ are $a_j u^k d^{j-k}$ with $u = e^{\sigma\sqrt{\delta t}}$, $d = e^{-\sigma\sqrt{\delta t}}$, and k ranging from 0 to j. The parameter σ is the volatility of the spot rate; we assume it is known (e.g. from market data) and constant. The parameters a_0, a_1, a_2, etc. represent a time-dependent drift in the spot rate (more precisely $\mu_i = (1/\delta t)\log(a_i)$ is the drift in the spot rate, since $a_i = e^{\mu_i \delta t}$). The task of calibration is to find the drift parameters a_0, a_1, etc. from the market observables, which are $B(0,1)$, $B(0,2)$, $B(0,3)$, etc.

We proceed inductively. Getting started is easy: $B(0,1) = e^{-r(0)\delta t}$ so $a_0 = r(0)$ is directly observable. To determine a_1 let $Q(1)_u$ be the value at time 0 of the option whose value at time 1 is 1 in the up state and 0 in the down state; let $Q(1)_d$ be the value at time 0 of the option whose value at time 1 is 1 in the down state and 0 in the up state. Their values are evident from the tree:

$$Q(1)_u = \tfrac{1}{2} e^{-r(0)\delta t}, \quad Q(1)_d = \tfrac{1}{2} e^{-r(0)\delta t}.$$

They are useful because examination of the tree gives

$$B(0,2) = Q(1)_u e^{-r(1)_u \delta t} + Q(1)_d e^{-r(1)_d \delta t}.$$

The left hand side is known, while the right hand side depends on a_1; so this equation determines a_1 (it must be found numerically – no analytical solution is available).

Let's do one more step to make the scheme clear. To determine a_2 we define $Q(2)_{uu}$, $Q(2)_{ud}$, and $Q(2)_{dd}$ to be the values at time 0 of options worth 1 and the indicated time-2 node (the

uu node for $Q(2)_{uu}$, etc.) and worth 0 at the other time-2 nodes. Their values are evident from the tree:

$$\begin{aligned}
Q(2)_{uu} &= \tfrac{1}{2}e^{-r(1)_u\delta t}Q(1)_u \\
Q(2)_{ud} &= \tfrac{1}{2}e^{-r(1)_u\delta t}Q(1)_u + \tfrac{1}{2}e^{-r(1)_d\delta t}Q(1)_d \\
Q(2)_{dd} &= \tfrac{1}{2}e^{-r(1)_d\delta t}Q(1)_d.
\end{aligned}$$

Notice that a_1 enters this calculation but a_2 does not. Now observe that

$$B(0,3) = Q(2)_{uu}e^{-r(2)_{uu}\delta t} + Q(2)_{ud}e^{-r(2)_{ud}\delta t} + Q(2)_{dd}e^{-r(2)_{dd}\delta t}.$$

The left hand side is known, while the right hand side depends on a_2; so this equation determines a_2 (solving for it numerically).

The general idea should now be clear: at each new timestep j we must find the $j+1$ values of $Q(j)_{xx}$; the tree gives us an formula based on the values of $Q(j-1)_{xx}$ and the time $j-1$ interest rates $r(j-1)_{xx}$. Then $B(0,j+1)$ can be expressed in terms of the various $Q(j)_{xx}$ and the time-j interest rates $r(j)_{xx}$, giving a nonlinear equation to solve for a_j. This procedure can easily be turned into an implementable algorithm. The only thing I haven't explained is how to index the nodes systematically, leading to general formulas for $Q(j)_{xx}$ and $B(j)_{xx}$. This is left as an exercise – or you can find it explained in section 8.4 of Clewlow and Strickland.