Lecture Notes in Finance 2 (MiQE/F, MSc course at UNISG)

Paul Söderlind[1]

10 December 2013

[1]University of St. Gallen. *Address:* s/bf-HSG, Rosenbergstrasse 52, CH-9000 St. Gallen, Switzerland. *E-mail:* Paul.Soderlind@unisg.ch. Document name: Fin2MiQEFAll.TeX

Contents

12 Interest Rate Calculations

Main references: Elton, Gruber, Brown, and Goetzmann (2010) 21–22 and Hull (2009) 4
Additional references: McDonald (2006) 7; Fabozzi (2004); Blake (1990) 3–5; and Campbell, Lo, and MacKinlay (1997) 10

12.1 Interest Rate Conventions

Suppose we borrow one unit of currency (that is, the face value of the loan is 1) that should be repaid with interest rate m periods later. The payment in period m is then the face value (of 1) plus the interest, so the payment in m is

$$\text{payment} = [1 + Y(m)]^m \tag{12.1}$$
$$= \exp[my(m)] \tag{12.2}$$
$$= 1 + mZ(m), \tag{12.3}$$

where $Y(m)$ is the *effective interest rate*, $y(m)$ the *continuously compounded interest rate* and $Z(m)$ is the *simple interest rate*.

Remark 12.1 *(The transformation from one type of rate to the other) We have*

$$y(m) = \ln[1 + Y(m)] \text{ and } y(m) = \ln[1 + mZ(m)]/m,$$
$$Y(m) = \exp[y(m)] - 1 \text{ and } Y(m) = [1 + mZ(m)]^{1/m} - 1$$
$$Z(m) = \{[1 + Y(m)]^m - 1\}/m \text{ and } Z(m) = \{\exp[y(m)] - 1\}/m.$$

The different interest rates (effective, continuously compounded and simple) are typically very similar, except for very high rates. See Figure 12.2 for an illustration.

12.2 Zero Coupon (discount or bullet) Bonds

Suppose a bond without dividends costs $B(m)$ in t and gives one unit of account in $t+m$ (the trade date index t is suppressed to simplify notation—in case of potential confusion,

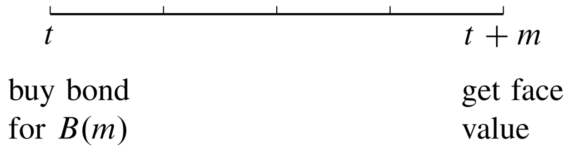

Figure 12.1: Timing convention of zero coupon bond

we can write $B_t\,(m)$). See Figure 12.1 for an illustration.

The gross return (payoff divided by price) from investing in this bond is $1/B(m)$, since the face value is normalized to unity.

$$\frac{1}{B\,(m)} = [1 + Y(m)]^m \text{ , or} \tag{12.4}$$

$$Y(m) = B(m)^{-1/m} - 1. \tag{12.5}$$

Another way to think of this is that if we invest the amount $B(m)$ by buying one bond, then after m periods we get $B(m)$ times the interest rate, that is, $B(m)\,[1 + Y(m)]^m = 1$. In practice, bond quotes are typically expressed in percentages (like 97) of the face value, whereas the discussion here effectively uses the fraction of the face value (like 0.97).

The relation between the rate and the price is clearly non-linear—and depends on the time to maturity (m): short rates are more sensitive to bond price movements than long rates. Conversely, prices on short bonds are less sensitive to interest rate changes than prices on long bonds. See Figure 12.2 for an illustration.

In terms of the continuously compounded rate, we have

$$\frac{1}{B\,(m)} = \exp\left[my\,(m)\right] \text{, or} \tag{12.6}$$

$$y(m) = -\ln B(m)/m. \tag{12.7}$$

Example 12.2 *(Effective and continuously compounded rates) Let the period length be a year (which is the most common convention for interest rates). Consider a six-month bill so $m = 0.5$. Suppose $B\,(m) = 0.95$. From (12.4) we then have that*

$$\frac{1}{0.95} = [1 + Y(0.5)]^{0.5} \text{, so } Y(0.5) \approx 0.108, \text{ and } y(0.5) \approx 0.103.$$

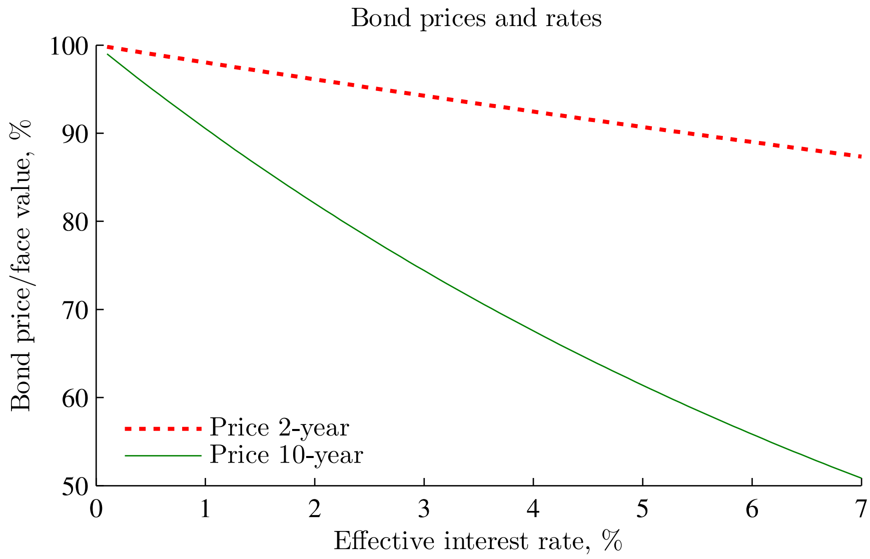

Figure 12.2: Interest rate vs. bond price

Example 12.3 *(Bond price changes vs interest rate changes) Suppose that, over a split second (so the time to maturity is virtually unchanged), the log bond price changes by* $\Delta \ln B$*, then (12.7) says that the change in the interest rate is*

$$\Delta y(m) = -\Delta \ln B(m)/m.$$

Inverting gives

$$\Delta \ln B(m) = -m\Delta y(m).$$

For instance, if the price of a 10-year bond decreases from 0.95 to 0.86 we get that the interest rate increases by

$$-\ln(0.86/0.95)/10 = 0.01,$$

that is, from 0.5% to 1.5%. Similarly, as the rate increases with 1%, the log price changes by

$$10 \times 0.01 = 0.1.$$

Some fixed income instruments (in particular inter bank loans, LIBOR/EURIBOR) are quoted in terms of a simple interest rate. The "price" of a deposit that gives unity at

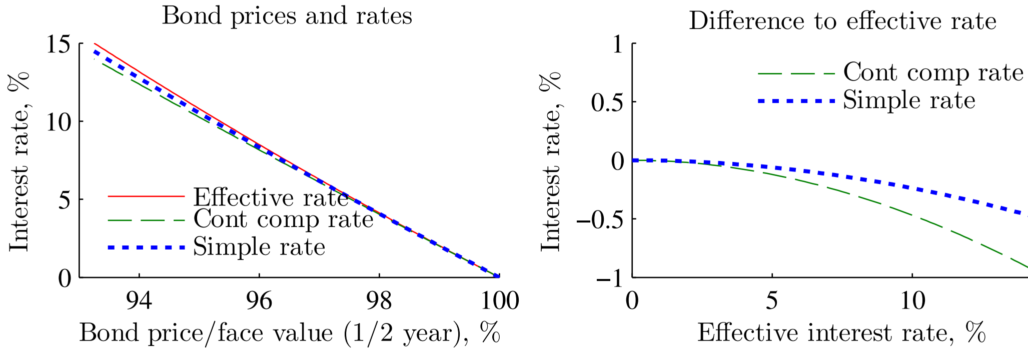

Figure 12.3: Different types of interest rates

maturity is then

$$B\,(m) = \frac{1}{1 + mZ(m)}, \text{ or} \tag{12.8}$$

$$Z(m) = \frac{1/B\,(m) - 1}{m}. \tag{12.9}$$

12.3 Forward Rates

12.3.1 Implied Forward Rates

A forward contract written in t stipulates buying at $t + m$, a discount bond that pays one unit of account at time $t + n$—see Figure 12.5 for an illustration. An arbitrage argument (see Figure 12.6) shows that the forward price must satisfy

$$\text{forward price} = B(n)/B(m). \tag{12.10}$$

Proof. (of (12.10)) In period t, buy one bond maturing in $t + n$ at the cost of $B(n)$ and sell $B(n)/B(m)$ bonds maturing in $t + m$ at the value of $B(n)$: the net investment in t is zero. In $t + m$, pay the principal of the maturing bonds at the cost $B(n)/B(m)$—this is the net investment in $t + m$. The payoff in $t + n$ is one. The forward contract has the same payoff in $t + n$ and must therefore specify the same net investment in $t + m$, the forward price: $B(n)/B(m)$. ∎

Buying a forward contract is effectively an investment from $t + m$ to $t + n$, that

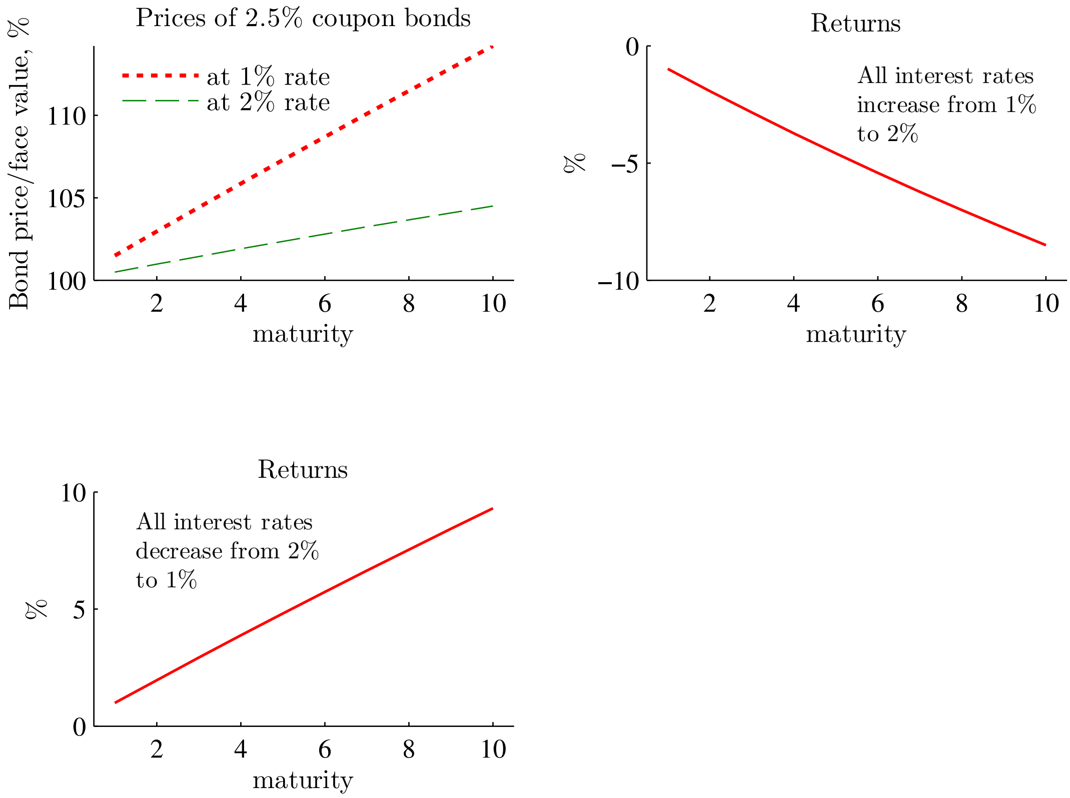

Figure 12.4: Gains and losses at interest rate changes

is, over $n - m$ periods. The gross return (which happens to be known already in t) is $1/[B(n)/B(m)]$. We define a per period effective rate of return, a forward rate, $F(m,n)$, analogous with (12.4)

$$\frac{1}{B(n)/B(m)} = [1 + F(m,n)]^{n-m}\,. \tag{12.11}$$

Notice that $F(m,n)$ here denotes a forward rate, not a forward price. This is the rate of return over $t + m$ to $t + n$ that can be guaranteed in t. By using the relation between bond prices and yields (12.4) this expression can be written

$$F(m,n) = \left[\frac{B(m)}{B(n)}\right]^{1/(n-m)} - 1 = \frac{[1 + Y(n)]^{n/(n-m)}}{[1 + Y(m)]^{m/(n-m)}} - 1. \tag{12.12}$$

See Figure 12.7 for an illustration.

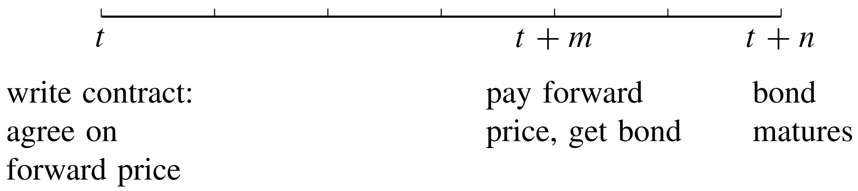

Figure 12.5: Timing convention of forward contract

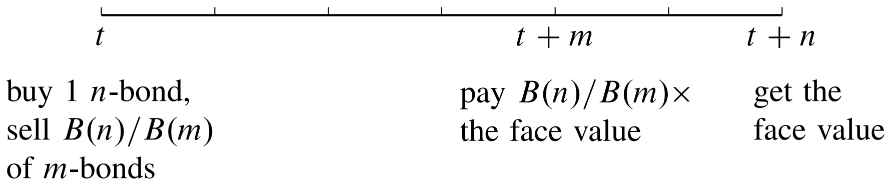

Figure 12.6: Synthetic forward contract

Remark 12.4 *(Alternative way of deriving the forward rate*) Rearrange (12.12) as*

$$[1+Y(m)]^m\,[1+F(m,n)]^{n-m}=[1+Y(n)]^n\,.$$

This says that compounding $1+Y(m)$ *over m periods and then* $1+F(m,n)$ *for* $n-m$ *periods should give the same amount as compounding the long rate,* $1+Y(n)$*, over n periods.*

Split up the time until n into n/h intervals of length h (see Figure 12.8). Then, the n-period spot rate equals the geometric average of the h-period forward rates over t to $t+n$

$$\begin{aligned}1+Y(n)&=[1+F(0,h)]^{h/n}\times[1+F(h,2h)]^{h/n}\times\ldots\times[1+F(n-h,n)]^{h/n}\\&=\prod\nolimits_{s=0}^{n/m-1}\{1+F[sh,(s+1)h]\}^{h/n}\,.\end{aligned}\qquad(12.13)$$

This means that the forward rate can be seen as the "marginal cost" of making a loan

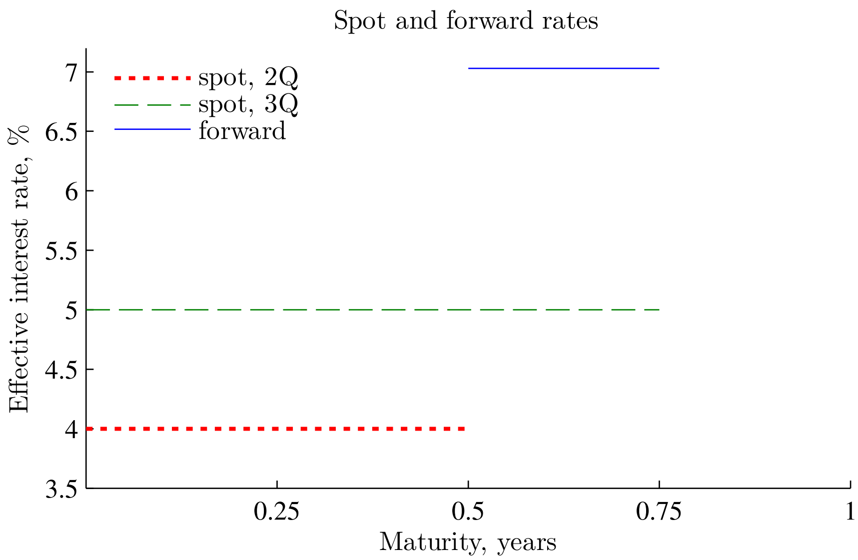

Figure 12.7: Spot and forward rates

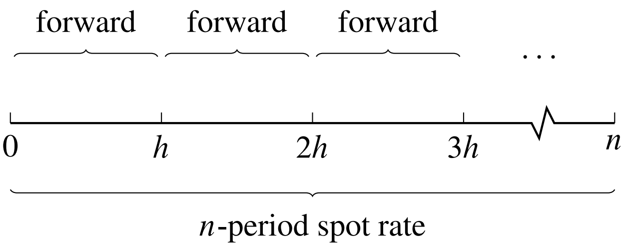

Figure 12.8: Forward contracts for several future periods

longer. See Figure 12.9 for an illustration.

Proof. (of (12.13)) Let $n = 2m$ and use (12.11) for forward contracts between 0 to m and m to $2m$

$$\frac{1}{B(m)/B(0)} = [1 + F(0,m)]^m \text{ and } \frac{1}{B(2m)/B(m)} = [1 + F(m,2m)]^m .$$

Multiply and simplify to get

$$\frac{1}{B(n)} = [1 + F(0,m)]^m \times [1 + F(m,2m)]^m .$$

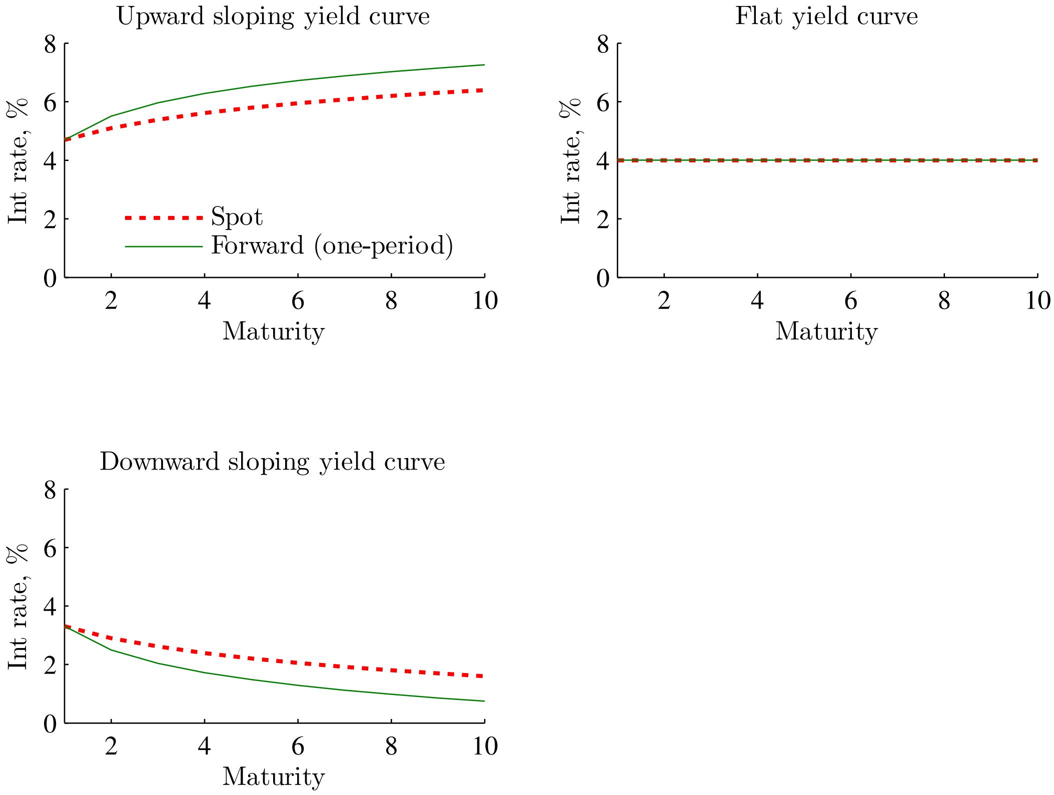

Figure 12.9: Spot and forward rates

Raise to the power of $1/n$ to get the interest rate

$$1 + Y(n) = [1 + F(0, m)]^{m/n} \times [1 + F(m, 2m)]^{m/n} .$$

■

Example 12.5 *(Forward rate) Let $m = 0.5$ (six months) and $n = 0.75$ (nine months), and suppose that $Y(0.5) = 0.04$ and $Y(0.75) = 0.05$. Then (12.12) gives*

$$[1 + F(0.5, 0.75)]^{0.75-0.5} = \frac{(1 + 0.05)^{0.75}}{(1 + 0.04)^{0.5}},$$

which gives $F(0.5, 0.75) \approx 0.07$. See Figure 12.7 for an illustration.

Example 12.6 *(Forward rate) Let the period length be a year. Let $m = 1$ (one year) and*

$n = 2$ (two years), and suppose that $Y(1) = 0.04$ and $Y(2) = 0.05$. Then (12.12) gives

$$F(1,2) = \frac{(1+0.05)^2}{(1+0.04)^1} - 1 \approx 0.06.$$

Example 12.7 *(Spot as average forward rate) In the previous example, (12.13) gives, using $F(0,1) = Y(1)$,*

$$1.04^{1/2} 1.06^{1/2} \approx 1.05,$$

which indeed equals $1 + Y(2)$.

Remark 12.8 *(Forward Rate Agreement) An* FRA *is an over-the-counter contract that guarantees an interest rate during a future period. The FRA does not involve any lending/borrowing—only compensation for the deviation of the future interest rate (typically LIBOR) from the agreed forward rate. An FRA can be emulated by a portfolio of zero-coupon bonds, similarly to a forward contract.*

12.3.2 Continuously Compounded and Simple Forward Rates

Taking logs of $1 + F(m,n)$ in (12.12) we get the continuously compounded forward rate

$$f(m,n) = \frac{1}{n-m} \ln \frac{B(m)}{B(n)} = \frac{ny(n) - my(m)}{n-m}. \tag{12.14}$$

Conversely, the n-period (continuously compounded) spot rate equals the average (continuously compounded) forward rate (take logs of 12.13)

$$y(m) = \frac{h}{n} \sum\nolimits_{s=0}^{n/h-1} f[sh, (s+1)h]. \tag{12.15}$$

A simple forward rate (used on interbank markets) is defined as

$$\frac{1}{B(n)/B(m)} = 1 + (n-m) Z^f(m,n), \text{ so} \tag{12.16}$$

$$Z^f(m,n) = \frac{1}{n-m}\left[\frac{B(m)}{B(n)} - 1\right] = \frac{nZ(n) - mZ(m)}{(n-m)[1 + mZ(m)]}. \tag{12.17}$$

12.3.3 Instantaneous Forward Rates

The instantaneous forward rate, $f(m)$, is defined as the limit when the maturity date of the bond approaches the settlement date of the forward contract, $n \to m$. This can be

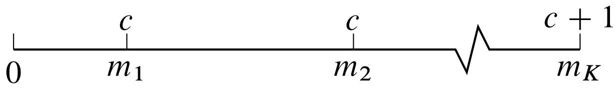

Figure 12.10: Timing convention of coupon bond

thought of as a forward "overnight" rate m periods ahead in time. From (12.14) it is

$$
\begin{aligned}
f(m) &= \lim_{n \to m} f(m,n) && (12.18)\\
&= \lim_{n \to m} \frac{n-m}{n-m} y(n) - \lim_{n \to m} \frac{m\left[y(m) - y(n)\right]}{n-m} \\
&= y(m) + m \frac{dy(m)}{dm}. && (12.19)
\end{aligned}
$$

Conversely, the average of the forward rates over t to $t+n$ is the spot rate, which we see by integrating (12.19) to get

$$
y(n) = \frac{1}{n} \int_0^n f(s) ds. \qquad (12.20)
$$

Equations (12.19) and (12.20) show that the difference between the forward and spot rates, $f(n) - y(n)$, is proportional to the slope of the yield curve.

Proof. (of (12.20)) Integrating the first term on the right hand side of (12.19) over $[0, n]$ gives $\int_0^n y(s)ds$. Integrating (by parts) the second term on the right hand side of (12.19) over $[0, n]$, $\int_0^n s \frac{dy(s)}{ds} ds$, gives $ny(n) - \int_0^n y(s)ds$. Adding the two terms gives $ny(n)$. ■

12.4 Coupon Bonds

12.4.1 Bond Basics

Consider a bond which pays coupons, c, for K periods ($t + m_1, t + m_2$,..), and one unit of account (the "face" or "par" value) in the last period $t + m_K$—see Figure 12.10.

The coupon bond is, in fact, a portfolio of zero coupon bonds: c maturing in $t + m_1$, c in $t + m_2$,..., and 1 in $t + m_K$. The price of the coupon bond, $B(K, c)$, must therefore equal the price of the portfolio

$$B\,(K,c) = \textstyle\sum_{k=1}^{K} B(m_k)c + B(m_K) \tag{12.21}$$

where $B(m)$ is the price of a zero coupon bond. Using the relation between (zero coupon) bond prices and yields in (12.4), this can also be written

$$B\,(K,c) = \textstyle\sum_{k=1}^{K} \frac{c}{[1+Y(m_k)]^{m_k}} + \frac{1}{[1+Y(m_K)]^{m_K}}. \tag{12.22}$$

Example 12.9 *(Coupon bond price) Suppose* $B(1) = 0.95$ *and* $B(2) = 0.90$. *The price of a bond with a 6% annual coupon with two years to maturity is then*

$$1.01 \approx 0.95 \times 0.06 + 0.90 \times 0.06 + 0.90.$$

Equivalently, the bond prices imply that $Y(1) \approx 5.3\%$ *and* $Y(2) \approx 5.4\%$ *so*

$$1.01 \approx \frac{0.06}{1.053} + \frac{0.06+1}{1.054^2}.$$

Example 12.10 *(Coupon bond price at par) A 9% (annual coupons) Suppose* $B(1) = 1/1.06$ *and* $B(2) = 1/1.091^2$. *The price of a bond with a 9% annual coupon with two years to maturity is then*

$$\frac{0.09}{1.06} + \frac{0.09}{1.091^2} + \frac{1}{1.091^2} \approx 1.$$

This bond is (approximately) sold "at par", that is, the bond price equals the face (or par) value (which is 1 in this case).

If we knew all the spot interest rates, then it would be easy to calculate the correct price of the coupon bond. The special (admittedly unrealistic) case when all spot rates are the same (flat yield curve) is interesting since it provides good intuition for how coupon bond prices are determined. In partcular, if the next coupon payment is one period ahead ($m_k = k$), then (12.22) becomes

$$B\,(K,c) = 1 + \frac{c-Y}{Y}\left[1-(1+Y)^{-K}\right], \tag{12.23}$$

where Y is the (common) spot rate. The term in square brackets is positive (for $K > 0$), so when the interest rate (which then equals the yield to maturity, see below) is below the coupon rate, then the bond price is above the face value (since $c - Y > 0$)—and vice versa. See Figure 12.11 for an illustration.

However, the situation is typically the reverse: we know prices on several coupon bonds (different maturities and coupons), and want to calculate the spot interest rates that are compatible with them. This is to *estimate the yield curve*. The implied zero coupon bonds prices is often called the *discount function*.

Remark 12.11 *(STRIPS, Separate Trading of Registered Interest and Principal of Securities) A coupon bond can be split up into its embedded zero coupon bonds—and traded separately. STRIPS are therefore zero coupon bonds.*

12.4.2 Yield to Maturity

The effective *yield to maturity* (also called redemption yield), θ, on a coupon bond is the internal rate of return which solves

$$B(K, c) = \sum_{k=1}^{K} \frac{c}{(1+\theta)^{m_k}} + \frac{1}{(1+\theta)^{m_K}}, \tag{12.24}$$

where the bond pays coupons, c, at $m_1, m_2, ..., m_K$ periods ahead. This equation can be solved (numerically) for θ. Quotes of bonds are typically the yield to maturity or the price. For a *par bond* (the bond price equals the face value, here 1), the yield to maturity equals the coupon rate. For a zero coupon bond, the yield to maturity equals the spot interest rate.

Example 12.12 *(Yield to maturity) A 4% (annual coupon) bond with 2 years to maturity. Suppose the price is 1.019. The yield to maturity is 3% since it solves*

$$1.019 \approx \frac{0.04}{1+0.03} + \frac{0.04}{(1+0.03)^2} + \frac{1}{(1+0.03)^2}.$$

Example 12.13 *(Yield to maturity of a par bond) A 4% (annual coupon) par bond (price of 1)with 2 years to maturity. The yield to maturity is 4% since*

$$\frac{0.04}{1+0.04} + \frac{0.04}{(1+0.04)^2} + \frac{1}{(1+0.04)^2} = 1$$

Example 12.14 *(Yield to maturity of a portfolio) A 1-year discount bond with a ytm (effective interest rate) of 7% has the price* $1/1.07$ *and a 3-year discount bond with a ytm of*

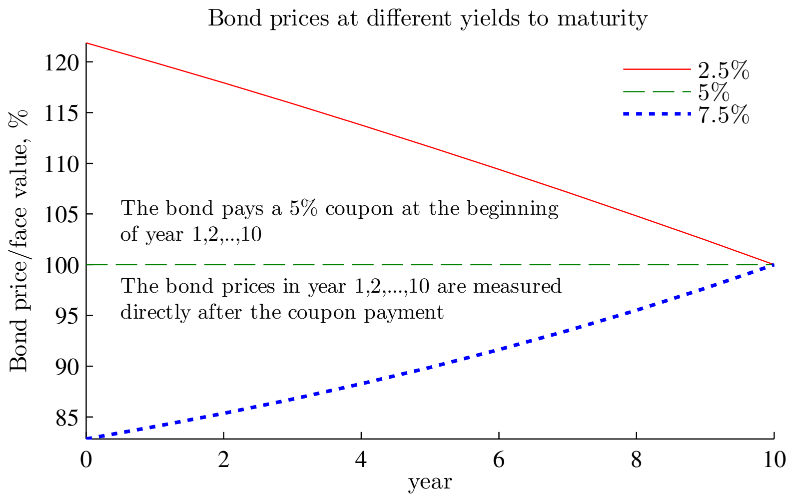

Figure 12.11: Bond price and yield to maturity

10% has the price $1/1.1^3$. *A portfolio with one of each bond has a ytm*

$$\frac{1}{1.07} + \frac{1}{1.1^3} = \frac{1}{1+\theta} + \frac{1}{(1+\theta)^3}, \textit{ with } \theta \approx 0.091.$$

This is clearly not the average ytm of the two bonds. It would be, however, if the yield curve is flat.

Note that the yield to maturity is just a convention. In particular, it does not provide a measure of the return to an investor who buys the bond and keeps it until maturity—unless the investor can reinvest all coupons to the same interest rate as the ytm of the bond. If the yield curve is flat, then that can be done by forward contracts or swaps. If the yield curve is not flat, then the rates at which the coupons are reinvested are typically different from the ytm of the bond—so the overall return of holding the bond will also be.

12.4.3 The Return of Holding a Coupon Bond until Maturity

To calculate the buy-and-hold (until maturity) return of a coupon bond we need to specify how the coupons are reinvested. One useful assumption is that the coupons are reinvested via forward contracts. This means that the investor buys the bond now and receives nothing until maturity—as if he/she had bought a zero-coupon bond. Indeed, no-arbitrage arguments show that the return (from now to maturity) is indeed the spot interest on a zero-coupon bond.

Proof. (Buy-and-hold return on a coupon bond, simple case) Consider a 3-period coupon bond. From (12.22), the price of the bond is

$$B\,(K, c) = B(1)c + B(2)c + B(3)c + B(3).$$

From (12.11), we know that the forward contract for the first coupon has the gross return (until maturity) $1/[B(3)/B(1)]$ and that the forward contract for the second coupon has the cross return (until maturity) $1/[B(3)/B(2)]$. The value of the reinvested coupons and the face value at maturity is then

$$\frac{B(1)}{B(3)}c + \frac{B(2)}{B(3)}c + c + 1.$$

Dividing by the first equation (the investment) gives $1/B(3)$ so the return on buying and holding (and reinvesting the coupons) this coupon bond is the same as the 3-period spot interest rate. (The extension to more periods is straightforward.) ■

Example 12.15 *(Yield to maturity versus return) Suppose also that the spot (zero coupon) interest rates are 4% for one year to maturity and 9% for 2 years to maturity. Notice that the forward rate (between year 1 and 2) is 14.24%. A 3% coupon bond with 2 years to maturity must have the price*

$$\frac{0.03}{1.04} + \frac{0.03 + 1}{1.09^2} \approx 0.8958.$$

The yield to maturity is 8.91% since

$$0.8958 \approx \frac{0.03}{1 + 0.0891} + \frac{0.03 + 1}{(1 + 0.0891)^2}.$$

However, the value of the bond at maturity, if the coupon is reinvested by a forward contract, is

$$0.03 \times (1 + 0.1424) + 0.03 + 1 \approx 1.0643,$$

so the gross return over two years is approximately $1.0643/0.8958$. *Annualized (* $\sqrt{1.0643/0.8958}$*) this becomes 1.09 so the effective annual return is 9%—just like the 2-year spot rate.*

12.4.4 Calculating the Yield to Maturity*

Remark 12.16 *(Calculating θ in a simple case) If m_k in (12.24) is the integer k, then subtracting B from both sides of (12.24) gives a K^{th} order polynomial in $\Theta = 1/(1+\theta)$,*

$$0 = -B + \sum_{k=1}^{K-1} c\Theta^k + (c+1)\,\Theta^K,$$

where all coefficients except one are positive. There is then only one positive real root, Θ_1. Many software packages contain routines for finding roots of polynomials. Once that is done, pick the only positive real root, Θ_1, and calculate the yield as $\theta = (1-\Theta_1)/\Theta_1$.

Remark 12.17 *(Calculating θ in the simplest case) If the bond price, B, is unity, then the bond is sold "at par." If also m_k in (12.24) is the integer k (as in the previous remark), then $\theta = c$.*

Example 12.18 *(Par bond) A 9% (annual coupons) 2-year bond with a yield to maturity of 9%, and exactly two years to maturity has the price*

$$\frac{0.09}{1+0.09} + \frac{0.09}{(1+0.09)^2} + \frac{1}{(1+0.09)^2} = 1.$$

Remark 12.19 *(Newton-Raphson algorithm for solving (12.24)) It is straightforward to use a Newton-Raphson algorithm to solve (12.24). It is then useful to note that the derivative is*

$$\frac{dB(\theta)}{d\theta} = -\sum_{k=1}^{K} \frac{m_k c}{(1+\theta)^{m_k+1}} - \frac{m_K}{(1+\theta)^{m_K+1}}.$$

The Newton-Raphson algorithm is based on a first order Taylor expansion of the bond price equation

$$B(\theta_1) = B(\theta_0) + \frac{dB(\theta_0)}{d\theta}\,(\theta_1 - \theta_0)\,.$$

Set the left hand side equal to the observed price, B, guess a values of θ and call it θ_0; then solve for θ_1 as $\theta_1 = \theta_0 + [B - B(\theta_0)]\,/\frac{dB(\theta_0)}{d\theta}$. θ_1 is probably a better guess of θ than θ_0. Improve by repeating this updating as $\theta_2 = \theta_1 + [B - B(\theta_1)]\,/\frac{dB(\theta_1)}{d\theta}$, and so forth until θ_n converges.

Remark 12.20 *(Bisection method for solving (12.24)) The bisection method is a very simple (no derivatives are needed) and robust way to solve for the yield to maturity. First, start with a lower (θ_L) and higher (θ_H) guess of the yield which are known to bracket the* *true value, that is, $B(\theta_H) \leq B \leq B(\theta_L)$ where B is the observed bond price and $B(\theta)$ is the value according to (12.24). Recall that $B(\theta)$ is decreasing in θ. Second, calculate the bond price at the average of the two guesses: $B[(\theta_L + \theta_H)/2]$. Third, replace either θ_L or θ_H according to: if $B[(\theta_L + \theta_H)/2] \geq B$ (so the midpoint $(\theta_L + \theta_H)/2$ is below the true yield) then replace θ_L by $(\theta_L + \theta_H)/2$ (a higher value), but if $B[(\theta_L + \theta_H)/2] < B$ then replace θ_H by $(\theta_L + \theta_H)/2$ (a lower value). Fourth, iterate until $\theta_L \approx \theta_H$.*

Example 12.21 *(Bisection method). The first couple of iterations for a 2-year bond with a 4% coupon and a price of 1.019 are (see also Figure 12.12)*

Iteration	θ_L	θ_H	$(\theta_L + \theta_H)/2$	$B[(\theta_L + \theta_H)/2]$
1	0	0.05	0.0250	1.0289
2	0.025	0.05	0.0375	1.0047
3	0.025	0.0375	0.03125	1.0167
4	0.025	0.03125	0.028125	1.0228
5	0.028125	0.03125	0.029687	1.0197

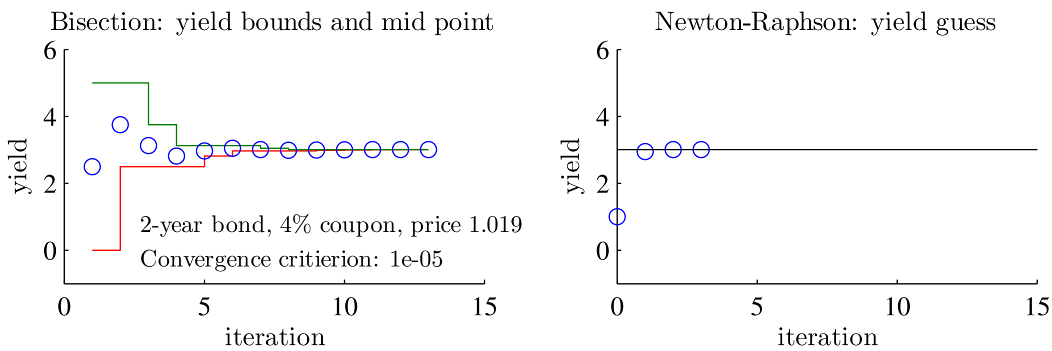

Figure 12.12: Bisection method to calculate yield to maturity

12.4.5 Par Yield

A par yield for is the coupon rate at which a bond would trade at par (that is, have a price equal to the face value). Setting $B(K, c) = 1$ in (12.21) and solving for the implied coupon rate gives

$$c = \frac{1}{\sum_{k=1}^{K} B(m_k)} \left[1 - B(m_K)\right], \text{ or} \tag{12.25}$$

$$= \frac{1}{\sum_{k=1}^{K} \frac{1}{[1+Y(m_k)]^{m_k}}} \left[1 - \frac{1}{[1+Y(m_K)]^{m_K}}\right]. \tag{12.26}$$

Typically, this is very similar to the zero coupon rates.

Example 12.22 *Suppose* $B(1) = 0.95$ *and* $B(2) = 0.90$*. We then have*

$$1 = (0.95 + 0.9)c + 0.9, \text{ so } c = \frac{1}{0.95 + 0.9}(1 - 0.9) \approx 0.054.$$

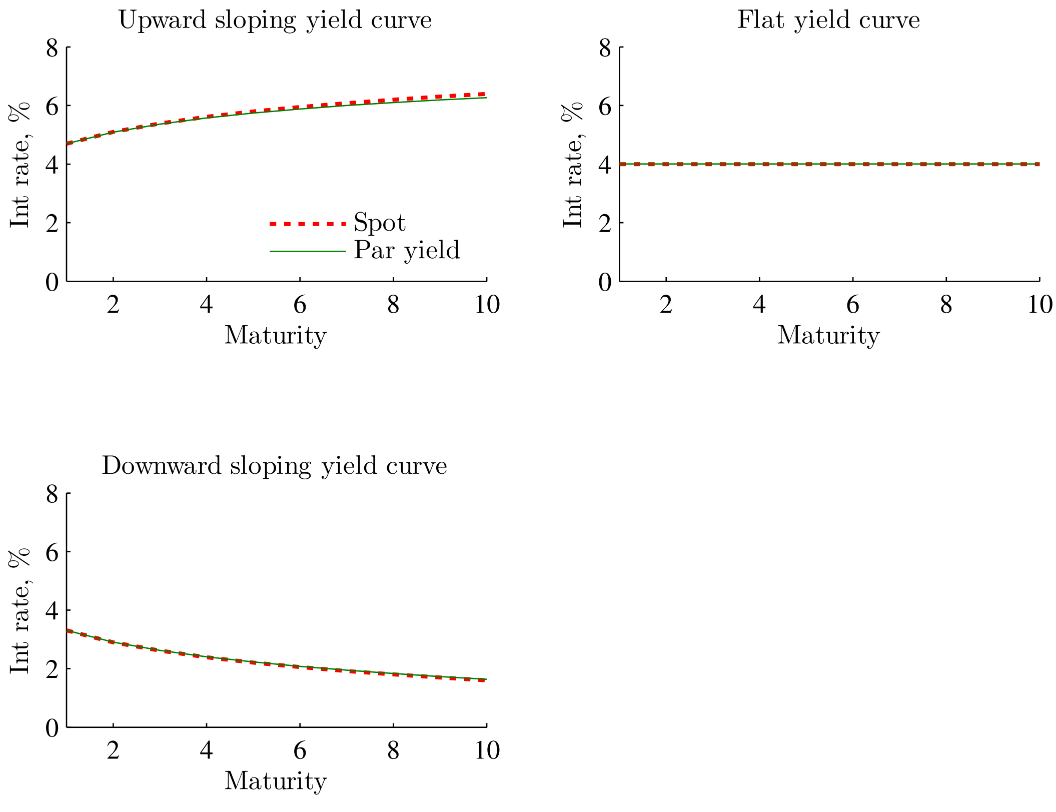

Figure 12.13: Spot and par yield curve

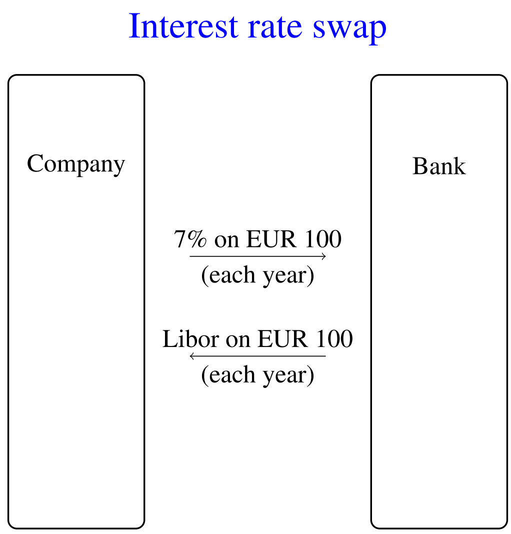

Figure 12.14: Interest rate swap

12.5 Swap and Repo

12.5.1 Swap

A swap contract involves a sequence of payment over the life time (maturity) of the contract: for each tenor (that is, sub period, for instance a quarter) it pays the floating market rate (say, the 3-month Libor) in return for a fixed *swap rate*. Split up the time until maturity n into n/h intervals of length h—see Figure 12.15. In period sh, the swap contract pays

$$h[Z_{(s-1)h}(h) - R] \tag{12.27}$$

where $Z_{(s-1)h}(h)$ is the short (floating) simple h-period interest rate in $(s-1)h$ and R is the (fixed) swap rate determined in t (as part of the swap contract).

The issuer can lock in the floating rate payments by a sequence of forward rate agreements that pay the floating rate in return for the forward rate. In this way the swap contract becomes riskfree so its present value must be zero. This implies that the swap rate must therefore be (assuming no default or liquidity premia)

$$R = \frac{1}{h}\frac{1 - B(n)}{\sum_{s=1}^{n/h} B(sh)}, \tag{12.28}$$

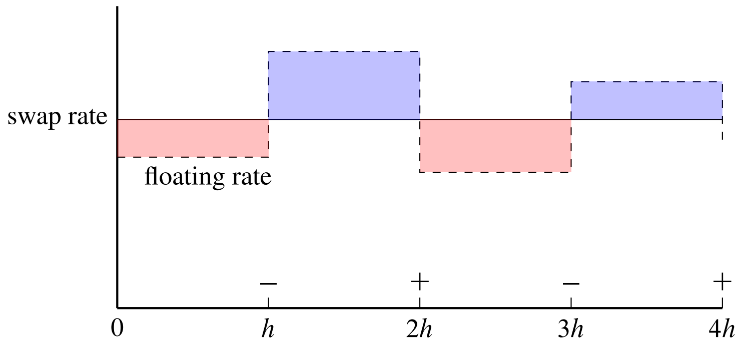

(The party receiving the floating rate pays the fixed swap rate)
(The net payments are marked by + or −)

Figure 12.15: Timing convention of interest rate swap

which is proportional to the par yield in (12.25).

Example 12.23 *(Swap rate) Consider a one-year swap contract with quarterly periods* ($n = 1, h = 1/4$)*. (12.28) is then*

$$R = 4\frac{1 - B(1)}{B(1/4) + B(1/2) + B(3/4) + B(1)}.$$

With the bond prices (0.99,0.98,0.97,0.96) we have

$$R = 4\frac{1 - 0.96}{0.99 + 0.98 + 0.97 + 0.96} \approx 4.1\%.$$

An *Overnight Indexed Swap* (OIS) is a swap contract where the floating rate is tied to an index of floating rates (for instance, federal funds rates in the U.S., EONIA in Europe—which is a weighted average of all overnight unsecured interbank lending transactions). Since the OIS has very little risk (as the face value or notional never changes hands—only the interest payment is risked in case of default), it is little affected by interbank risk premia. The quote is in terms of the fixed rate (called the swap rate, quoted a simple interest rate)—which typically stays close to secured lending rates like repo rates.

Proof. (of (12.28)) Notice that a simple forward rate for an investment from sh to $(s+1)h$ is

$$Z^f[sh, (s+1)h] = \frac{1}{h}\left[\frac{B(sh)}{B[(s+1)h]} - 1\right].$$

We can therefore write the present value of (12.27) as

$$PV = \sum_{s=1}^{n/h} B(sh) \left\{ \left[\frac{B[(s-1)h]}{B(sh)} - 1 \right] - hR \right\}.$$

Since it is riskfree (assuming no default and liquidity premia) the PV should be zero (or else there are arbitrage opportunities), which we rearrange as

$$hR \sum_{s=1}^{n/h} B(sh) = \sum_{s=1}^{n/h} B(sh) \left[\frac{B[(s-1)h]}{B(sh)} - 1 \right]$$
$$hR \sum_{s=1}^{n/h} B(sh) = 1 - B(n),$$

where we have used the fact that $B(0) = 1$. Finally, solve for hR to get (12.28). ■

12.5.2 Repo

A *Repo* (Repurchase agreement) is a way of borrowing against a collateral. Suppose bank A sells a security to bank B, but there is an agreement that bank A will buy back the security at some fixed point in time (the next day, after a week, etc.)—at a price that is predetermined (or decided according to some predetermined formula). This means that bank A gets a loan against a collateral (the asset)—and pays an interest rate (final buy price/initial sell price minus one). See Figure 12.16. Bank B is said to have made a reverse repo. Another way to think about the repo is that bank A has made a sale of the security, but also acquired a forward contract on it (the position of bank B is just the reverse). The repo clearly means that bank B has "borrowed" the security—which can then be sold to someone else. This is a way of shortening the security, so the repo rate is low if there is a demand for shortening the security. A *haircut* (of 3%, say) means that the collateral (security) has market value that is 3% higher than the price agreed in the repo. This provides a safety margin to the lender—since the market price of the security could decrease over the life span of the repo.

Example 12.24 *(Long-short bond portfolio). First, buy bond X and use it as collateral in a repo (the repo borrowing finances the purchase of the bond). Second, enter a reverse repo where bond Y is used as collateral and sell the bond (selling provides cash for the repo lending).*

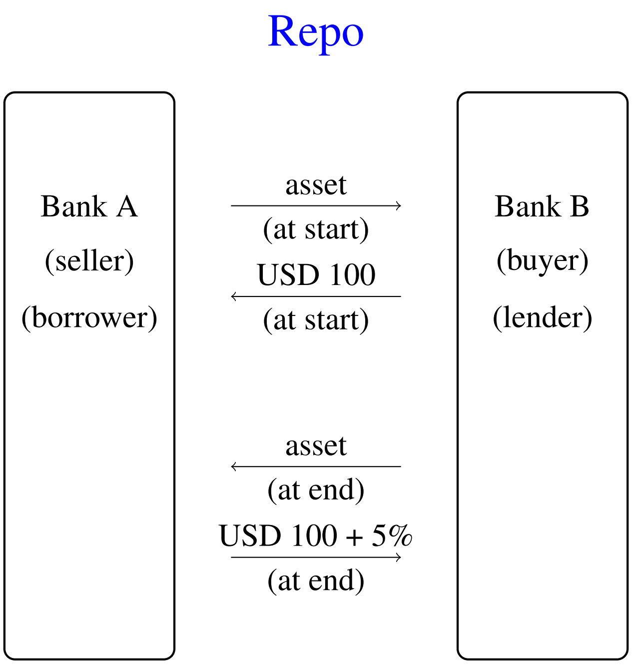

Figure 12.16: Repo

12.6 Estimating the Yield Curve

The (zero coupon) spot rate curve is of particular interest: it helps us price any bond or portfolio of bonds—and it has a clear economic meaning ("the price of time").

In some cases, the spot rate curve is actually observable—for instance from swaps and STRIPS. In other cases, the instruments traded on the market include some zero coupon instruments (bills) for short maturities (up to a year or so), but only coupon bonds for longer maturities. This means that the spot rate curve needs to be calculated (or estimated). This section describes different methods for doing that.

12.6.1 Direct Calculation of the Yield Curve ("Bootstrapping")

We can sometimes calculate large portions of the yield curve directly from asset prices. The idea is to calculate a short yield first (from a bill/bond with short time to maturity) and then use this to calculate the yield for the next (longer) bond, and so on.

For instance, suppose we have a one-period coupon bond, which by (12.21) must have the price

$$B[1, c(1)] = B(1)[c(1) + 1], \tag{12.29}$$

where we use $c(1)$ to indicate the coupon value of this particular bond. The equation immediately gives the one-period discount function value, $B(1)$. Suppose we also have a two-period coupon bond, which pays the coupon $c(2)$ in $t + 1$ and $t + 2$ as well as the principal in $t + 2$, with the price (see (12.21))

$$B[2, c(2)] = B(1)c(2) + B(2)[c(2) + 1]. \tag{12.30}$$

The two period discount function value, $B(2)$, can be calculated from this equation since it is the only unknown. We can then move on to the three-period bond,

$$B[3, c(3)] = B(1)c(3) + B(2)c(3) + B(3)[c(3) + 1] \tag{12.31}$$

to calculate $B(3)$, and so forth. Finally, we can use (12.4) to transform these zero coupon bond prices to spot interest rates.

Remark 12.25 *(Numerical calculation of the bootstrap) Equations (12.29)–(12.31) can clearly be written*

$$\begin{bmatrix} B[1, c(1)] \\ B[2, c(2)] \\ B[3, c(3)] \end{bmatrix} = \begin{bmatrix} c(1) + 1 & 0 & 0 \\ c(2) & c(2) + 1 & 0 \\ c(3) & c(3) & c(3) + 1 \end{bmatrix} \begin{bmatrix} B(1) \\ B(2) \\ B(3) \end{bmatrix},$$

which is a recursive (triangular) system of equations.

Example 12.26 *(Bootstrapping) Suppose we know that* $B(1) = 0.95$ *and that the price of a bond with a 6% annual coupon with two years to maturity is 1.01. Since the coupon bond must be priced as*

$$0.95 \times 0.06 + B(2) \times 0.06 + B(2) = 1.01,$$

we can solve for the price of a two-period zero coupon bond as $B(2) \approx 0.90$*. The spot interest rates are then* $Y(1) \approx 0.053$ *and* $Y(2) \approx 0.054$.

Unfortunately, the bootstrap approach is tricky to use. First, there are typically gaps between the available maturities. On way around that is to interpolate. Second (and quite the opposite), there may be several bonds with the same maturity but with different coupons/prices, so it hard to calculate a unique yield curve. This could be solved by forming an average across the different bonds or by simply excluding some data.

12.6.2 Estimating the Yield Curve with Regression Analysis

If we attach some random error to the bond prices in (12.21), then that equation looks very similar to regression equation: the coupon bond price is the dependent variable; the coupons are the regressors, and the discount function values are the coefficients to estimate—perhaps with OLS. This is a way of overcoming the second problem discussed above since multiple bonds with the same maturity, but different coupons, are just additional data points in the estimation.

The first problem mentioned above, gaps in the term structure of available bonds, is harder to deal with. If there are more coupon dates than bonds, then we cannot estimate all the necessary zero coupon bond prices from data (fewer data points than coefficients). The way around this is to decrease the number of parameters that need to be estimated by postulating that the price of a discount bond, $B(m)$, is a linear combination of some J predefined functions of maturity, $g_1(m)$,..., $g_J(m)$,

$$B(m) = 1 + \textstyle\sum_{j=1}^{J} a_j g_j(m), \tag{12.32}$$

where $g_j(0) = 0$ since $B(0) = 1$ (the price of a bond maturing today is one).

Once the $g_j(m)$ functions are specified, (12.32) is substituted into (12.21) and the j coefficients a_1,..., a_j are estimated by minimizing the squared pricing error (see, for instance, Campbell, Lo, and MacKinlay (1997) 10).

One possible choice of $g_j(m)$ functions is a polynomial, $g_j(m) = m^j$. Another common choice is to make the discount function a spline (see McCulloch (1975)).

Example 12.27 *(Quadratic discount function) With a quadratic discount function*

$$B(m) = a_0 + a_1 m + a_2 m^2,$$

we get

$$
\begin{aligned}
B\,(K,c) &= \textstyle\sum_{k=1}^{K}\left(a_0 + a_1 m_k + a_2 m_k^2\right) c + \left(a_0 + a_1 m_K + a_2 m_K^2\right) \\
&= a_0\,(Kc+1) + a_1\left(c\textstyle\sum_{k=1}^{K} m_k + m_K\right) + a_2\left(c\textstyle\sum_{k=1}^{K} m_k^2 + m_K^2\right).
\end{aligned}
$$

The a_0, a_1, and a_2 can be estimated by OLS if we have data on at least three bonds. This method can, however, lead to large errors in the fitted yields (if not the prices). See Figure 12.17 for an example.

Example 12.28 *(Cubic discount function). With a cubic discount function*

$$
B(m) = a_0 + a_1 m + a_2 m^2 + a_3 m^3,
$$

we get

$$
B\,(K,c) = a_0\,(Kc+1) + a_1\left(c\textstyle\sum_{k=1}^{K} m_k + m_K\right) + a_2\left(c\textstyle\sum_{k=1}^{K} m_k^2 + m_K^2\right) + a_3\left(c\textstyle\sum_{k=1}^{K} m_k^3 + m_K^3\right).
$$

12.6.3 Estimating a Parametric Forward Rate Curve

Yet another approach to estimating the yield curve is to start by specifying a function for the instantaneous forward rate curve, and then calculate what this implies for the discount function. (These will typically be complicated and not satisfy the simple linear structure in (12.32).)

Let $f(m)$ denote the instantaneous forward rate with time to settlement m. The *extended Nelson and Siegel forward rate function* (Svensson (1995)) is

$$
f(m;b) = \beta_0 + \beta_1 \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \frac{m}{\tau_1}\exp\left(-\frac{m}{\tau_1}\right) + \beta_3 \frac{m}{\tau_2}\exp\left(-\frac{m}{\tau_2}\right), \quad (12.33)
$$

where $b = (\beta_0, \beta_1, \beta_2, \tau_1, \beta_3, \tau_2)$ is a vector of parameters (β_0, τ_1 and τ_2 must be positive, and $\beta_0 + \beta_1$ must also be positive—see below). The original Nelson and Siegel function sets $\beta_3 = 0$. Note that in either case

$$
\begin{aligned}
&\lim_{m\to 0} f(m;b) = \beta_0 + \beta_1, \text{ and} \\
&\lim_{m\to \infty} f(m;b) = \beta_0,
\end{aligned}
$$

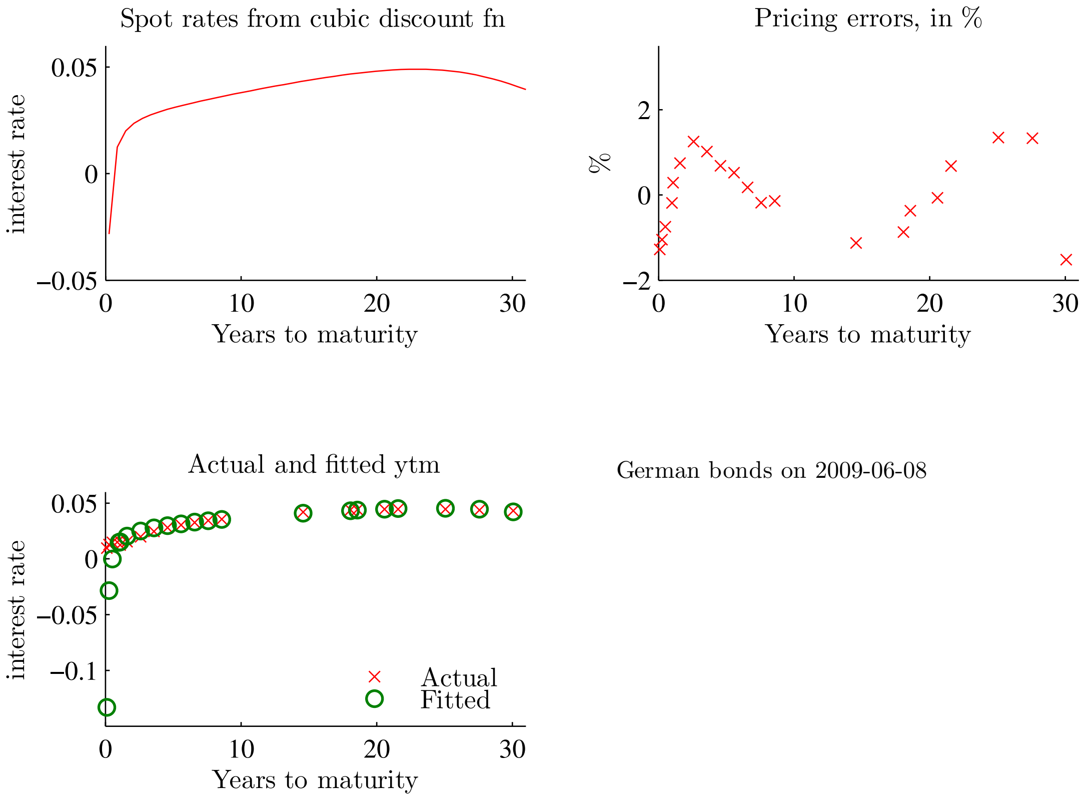

Figure 12.17: Estimated yield curves

so $\beta_0 + \beta_1$ corresponds to the current very short spot interest rate (an overnight rate, say) and β_0 to the forward rate with settlement very far in the future (the asymptote).

The spot rate implied by (12.33) is (integrate as in (12.20) to see that)

$$\begin{aligned} y(m;b) = \beta_0 + \beta_1 \frac{1-\exp(-m/\tau_1)}{m/\tau_1} + \beta_2 \left[\frac{1-\exp(-m/\tau_1)}{m/\tau_1} - \exp\left(-\frac{m}{\tau_1}\right) \right] \\ + \beta_3 \left[\frac{1-\exp(-m/\tau_2)}{m/\tau_2} - \exp\left(-\frac{m}{\tau_2}\right) \right]. \end{aligned} \tag{12.34}$$

One way of estimating the parameters in (12.33) is to substitute (12.34) for the spot rate in (12.6), and then minimize the sum of the squared price errors (differences between actual and fitted prices), perhaps with $1/\sqrt{\text{duration}}$ as the weights (a practice used by many central banks). Alternatively, one could minimize the sum of the squared yield errors (differences between actual and fitted yield to maturity). See Figures 12.18–12.20

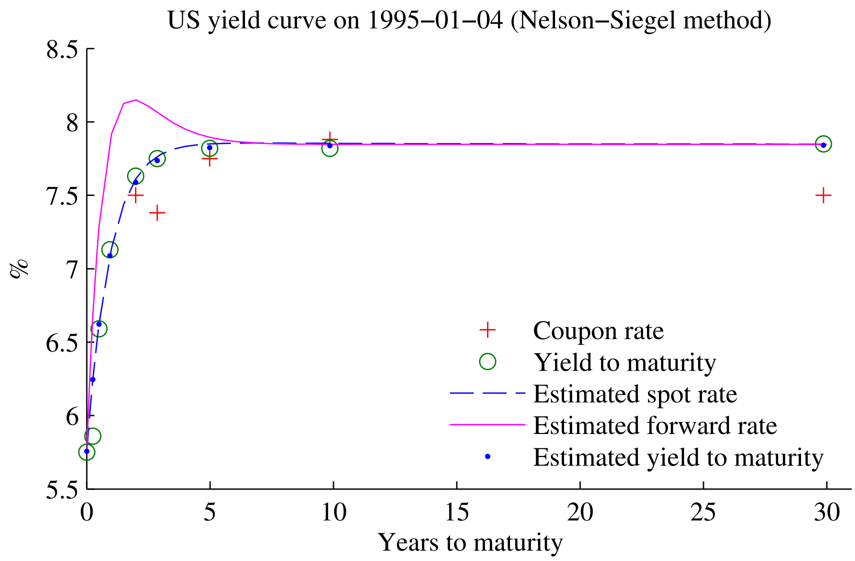

Figure 12.18: Estimated US yield curve, Nelson-Siegel method

for illustrations.

12.6.4 Par Yield Curve

When many bonds are traded at (approximately) par, the par yield curve (12.25) can be obtained by just plotting the coupon rates. In practice, the yield to maturity is used instead (to partly compensate for the fact that the bonds are only approximately at par)—and the gaps (across maturities) are filled by interpolation. (Recall that for a par bond, the yield to maturity equals the coupon rate.) This is basically the way the Constant Maturity Treasury yield curve, published by the US Treasury, is constructed.

12.6.5 Swap Rate Curve

The swap rates for different maturities can also be used to construct a yield curve.

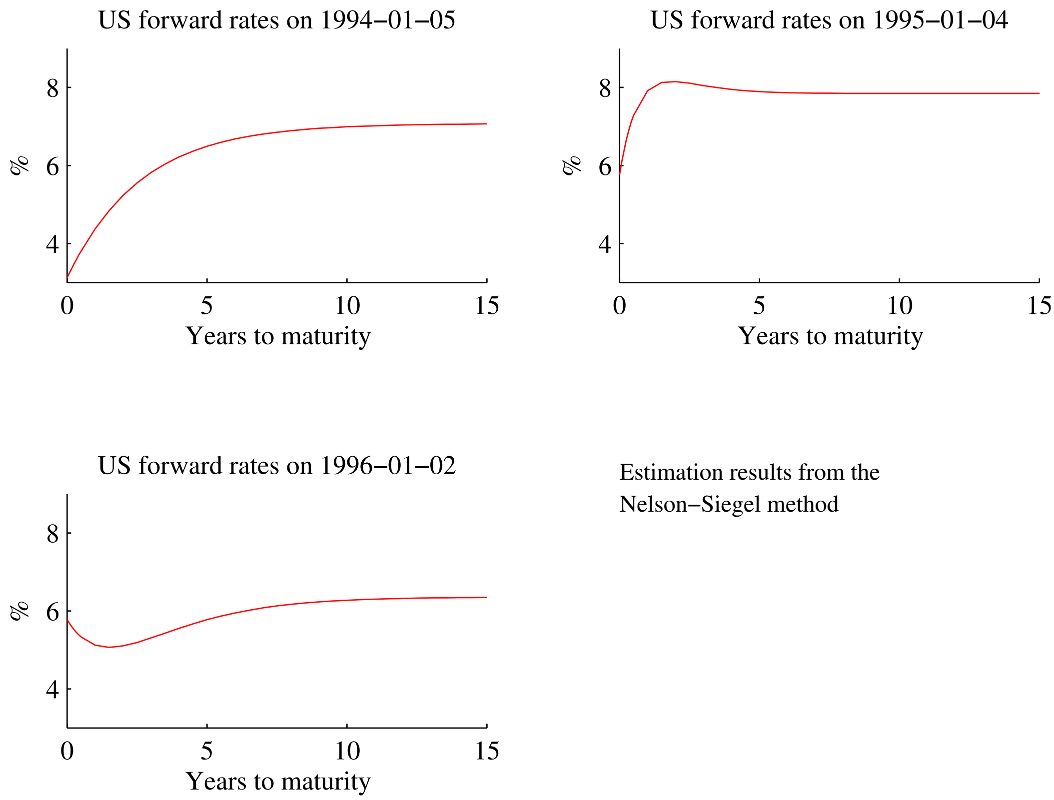

Figure 12.19: Estimated US yield curve, Nelson-Siegel method

12.7 Conventions on Important Markets*

12.7.1 Compounding Frequency

Suppose the interest rate r is compounded 2 times per year. This means that the amount B invested at the beginning of the year gives $B(1 + r/2)$ after six months—which is reinvested and therefore gives $B(1 + r/2)(1 + r/2)$ after another six months (at the end of the year). To make this payoff equal to unity (as we have used as our convention) it must be the case that the bond price $B = 1/(1 + r/2)^2$. By comparing with the definition of the effective interest rate (with annual compounding) in (12.4) we have

$$\frac{1}{B} = \left(1 + \frac{r}{2}\right)^2 = 1 + Y, \tag{12.35}$$

where Y is the annual effective interest rate.

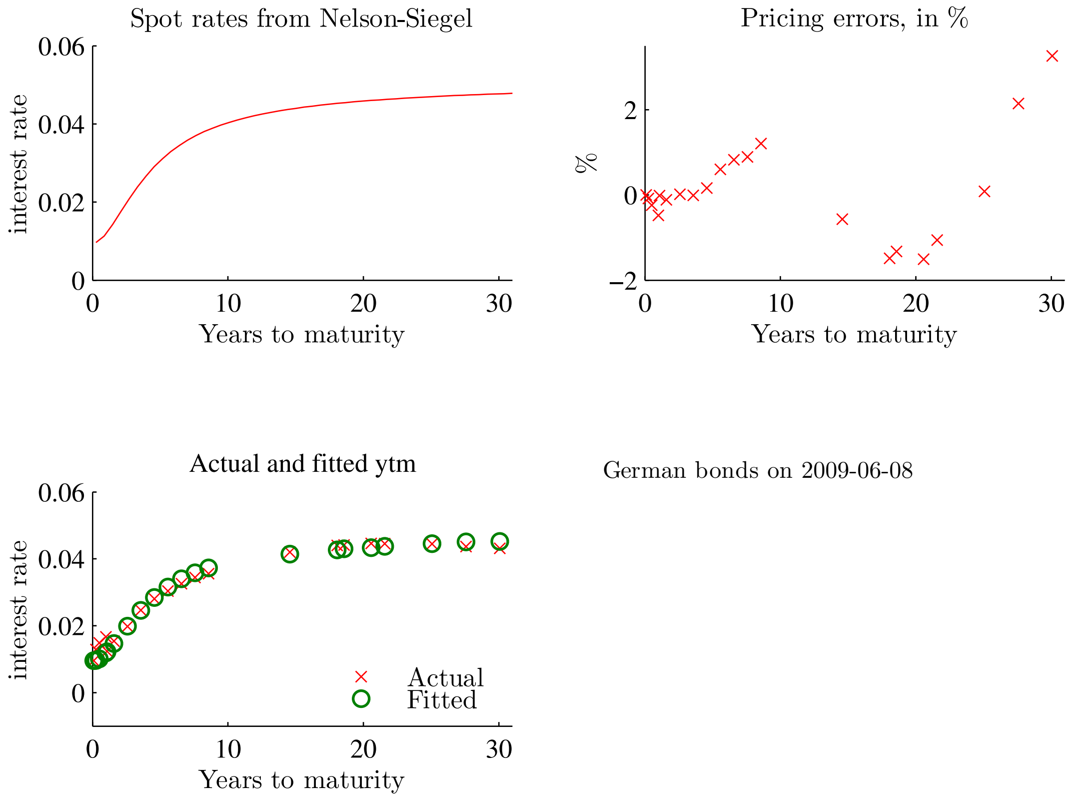

Figure 12.20: Estimated yield curves

This shows how we can transform from semi-annual compounding to annual compounding (and vice versa).

More generally, with compounding n times per year, we have

$$\frac{1}{B} = \left(1 + \frac{r}{n}\right)^n = 1 + Y. \tag{12.36}$$

12.7.2 US Treasury Notes and Bonds

The convention for *US Treasury notes and bonds* (issued with maturities longer than one year) is that coupons are paid semi-annually (as half the quoted coupon rate), and that yields are semi-annual effective yields. (This applies also to most as well as for most US corporate bonds and UK Treasury bonds.)

However, both are quoted on an annual basis by multiplying by two. The quoted *yield*

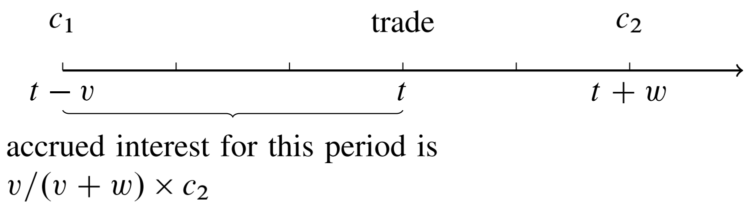

Figure 12.21: Accrued interest

to maturity, ϕ, solves

$$B(K, c) = \textstyle\sum_{k=1}^{K} \frac{c/2}{(1 + \phi/2)^{n_k}} + \frac{1}{(1 + \phi/2)^{n_K}}, \tag{12.37}$$

where the bond pays coupons $c/2$, at $n_1, n_2, ..., n_K$ half-years ahead. By using (12.35), the yield quoted, ϕ, can be expressed in terms of an annual effective interest rate.

Example 12.29 *A 9% US Treasury bond (the coupon rate is 9%, paid out as 4.5% semi-annually) with a yield to maturity of 7%, and one year to maturity has the price*

$$\frac{0.09/2}{1 + 0.07/2} + \frac{0.09/2}{(1 + 0.07/2)^2} + \frac{1}{(1 + 0.07/2)^2} = 1.019.$$

From (12.35), we get that the yield to maturity rate expressed as an annual effective interest is $(1 + 0.035)^2 - 1 \approx 0.071$.

Accrued Interest on US Bonds

The quotes of bond prices (as opposed to yields) are not the full price (also called the dirty price, invoice price, or cash price) the investor actually pays. Instead, it is the "clean price" that is quoted, which is the full price less the accrued interest:

full price = quoted price + accrued interest.

The buyer of the bond (buying in t) will typically get the next coupon (trading is "cum-dividend"). The accrued interest is the faction of that next coupon that has been accrued during the period the seller owned the bond. It is calculated as

accrued interest = next coupon $\times$ days since last coupon/182.5.

See Figure 12.21.

12.7.3 US Treasury Bills

Discount Yield

US Treasury bills have no coupons and are issued in 3, 6, 9, and 12 months maturities—but the time to maturity does of course change over time. They are quoted in terms of the (banker's) *discount yield*, $Y_{db}(m)$, which satisfies

$$B(m) = 1 - mY_{db}(m), \text{ where } m = \text{days}/360, \text{ so} \tag{12.38}$$

$$Y_{db}(m) = [1 - B(m)]/m. \tag{12.39}$$

Notice the convention of $m = $days$/360$.

From (12.4) and (12.38) it is the clear that the effective interest rate and the continuously compounded interest rates can be solved as

$$Y(m) = [1 - mY_{db}(m)]^{-1/m} - 1 \tag{12.40}$$

$$y(m) = -\ln[1 - mY_{db}(m)]/m. \tag{12.41}$$

Example 12.30 *A T-bill with 44 days to maturity and a quoted discount yield of 6.21% has the price* $1-(44/360)\times 0.0621 \approx 0.992$*. The effective interest rate is* $[1-(44/360)\times 0.0621]^{-360/44} - 1 \approx 6.43\%$.

12.7.4 LIBOR and EURIBOR

The LIBOR (London Interbank Offer Rate) and the EURIBOR (Euro Interbank Offered Rate) are the simple interest rate on a short term loan without coupons. It is quoted as a simple annual interest rate, using a "actual/360" day count—with the exception of pounds which are quoted "actual/365." This means that borrowing one dollar for 150 days at a 6% LIBOR requires the payment of $0.06 \times 150/360$ dollars in interest at maturity. Rescaling to make the payment at maturity equal to unity (which is the convention used in these lecture notes), the loan must be $1/(1 + 0.06 \times 150/360)$—which is the "price" of a deposit that gives unity in 150 days.

12.7.5 European Bond Markets

The major continental European bond markets (in particular, France and Germany) typically have annual coupons and the accrued interest is calculated according to the "actual/actual" convention, that is, as

accrued interest = next coupon $\times$ days since last coupon/365 (or 366).

(The computation is slightly more complicated for the UK and the Scandinavian countries, since they have ex-dividend periods.)

12.8 Inflation-Indexed Bonds*

Reference: Deacon and Derry (1998)

Consider an inflation-indexed coupon bond issued in t, which has both coupons and principal adjusted for inflation up to the period of payment (this is called "capital indexed," which is the most common type). Let P_t be the value of the relevant price index in period t. The coupon payments are cP_{t+m_1-l}/P_{t-l} at $t+m_1$, cP_{t+m_2-l}/P_{t-l} at $t+m_2$, and so forth—and also the principal is paid as P_{t+m_K-l}/P_{t-l} in $t + m_K$.

The lag factor l is the *indexation lag*. There are two reasons for this lag. First, the convention on many markets is that the bond price is quoted disregarding accrued interest (clean price). The typical case is as follows. The next coupon payment is m_1 periods ahead. The buyer of the bond in t will get this coupon (trading is "cum-dividend"). The full price the buyer pays to the seller in t is therefore

full price = quoted price + accrued interest,

where the accrued interest is typically the coupon payment times the fraction of this coupon period that has already passed. To pay this accrued interest, we have to know the next coupon payment, that is, cP_{t+m_1-l}/P_{t-l}; in t we must know the price level in $t + m_1 - l$. This mean that $l \geq m_1$ must always hold: the indexation lag must be at least as long as the time between coupon payments (six months in the UK).

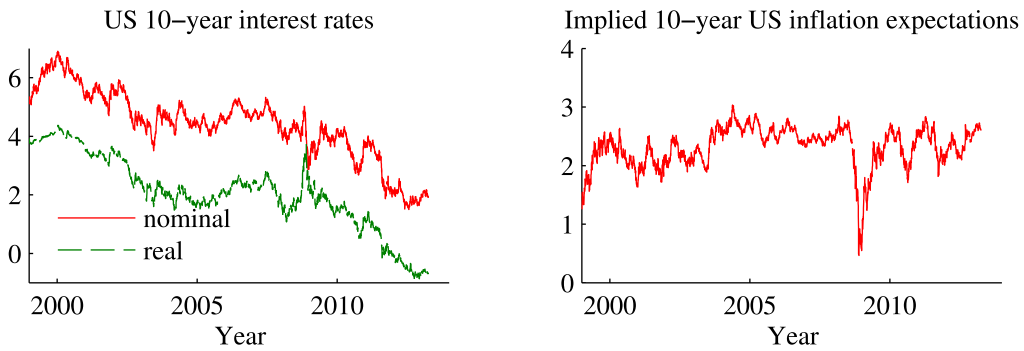

Figure 12.22: US nominal and real interest rates

Second, it takes time to calculate and publish price indices. Suppose we learn to know P_s in $s+k$. This means that the indexation lag must be an additional k periods, $l \geq m_1 + k$, so it uses a known price level. For instance, in the UK, the indexation lag is 8 months.

To simplify matters in the rest of this section, suppose the indexation lag is zero. Use (12.21), modified to allow for different coupons, to price the inflation-indexed bond. To further simplify, suppose that bonds do not have any riskpremia (clearly a strong assumption), so that the bond price equals the discounted expected payoffs

$$B\left(K,c\right) = \sum_{k=1}^{K} \frac{c \operatorname{E}_t P_{t+m_k}/P_t}{\left[1+Y(m_k)\right]^{m_k}} + \frac{\operatorname{E}_t P_{t+m_K}/P_t}{\left[1+Y(m_K)\right]^{m_K}}. \tag{12.42}$$

The Fisher equation is

$$\left[1+Y(m)\right]^m = \left[1+R(m)\right]^m \frac{\operatorname{E}_t P_{t+m}}{P_t}, \tag{12.43}$$

where R is the real interest rate. It splits up the gross nominal return in the bond into a gross real return and gross inflation rate. Notice that the Fisher equation assumes that there is no risk premia, which is a strong assumption.

Use (12.43) to rewrite (12.42) as

$$\begin{aligned} B\left(K,c\right) &= \sum_{k=1}^{K} \frac{c \operatorname{E}_t P_{t+m_k}/P_t}{\left[1+R(m_k)\right]^{m_k} \operatorname{E}_t P_{t+m_k}/P_t} + \frac{\operatorname{E}_t P_{t+m_K}/P_t}{\left[1+R(m_K)\right]^{m_K} \operatorname{E}_t P_{t+m_K}/P_t} \\ &= \sum_{k=1}^{K} \frac{c}{\left[1+R(m_k)\right]^{m_k}} + \frac{1}{\left[1+R(m_K)\right]^{m_K}} \end{aligned} \tag{12.44}$$

With a set of inflation-indexed bonds, we could therefore estimate a *real yield curve*, that is, how $R(m)$ depends on m. If the Fisher equation indeed holds, then the difference between a nominal interest rate and a real interest rate can be interpreted as a measure of the market's inflation expectations (often called the "break-even inflation rate").

12.9 Other Instruments

12.9.1 Collateralized Debt Obligations

CDO is a repackaging of a set of assets ("collaterals," typically bonds) where the claims (payouts) are tranched (have different priorities)

CDOs are created for two main reasons. First, it is a way for the issuer (typically a bank), to "package and sell off,":that, it sa way to shrink the balance sheet for the bank (securitisation) but still earn a fee. Second, a CDO transform risky bonds to (a) some safe bonds and (b) some very risky ones. This open up new possibilities for investors. For instance, it may allow risk averse investors (including pension funds) to invest into the safe tranches, while they would otherwise not dare (or be allowed to) invest into the original bonds.

It is clear that the correlation of the defeaults of the bonds in the CDO? The idea of tranching (in particular, to regard the senior tranche as safe) depends on the assumption that not all underlying debts default at the same time. Underestimating the correlation can lead to serious overpricing of the senior trances—as was often the case just before the financial crisis 2008–9.

Another important aspect of the CDO is whether the originator (bank) holds the junior trance or not. If it does, then it has the incentives to screen the borrowers/monitor the loans, otherwise not.

12.9.2 Credit Default Swaps

A credit default swap is an insurance against default on a bond (eg. , Greece). In fact, if you hold a portfolio of one risky bond and a CDS on it, then you effectively own a riskfree bond. The other way around is to buy one riskfree bond and issue a CDS, which gives effectively the same as owning the risky bond. This simple observation is the key to understanding how the CSD ("insurance") premium is determined.

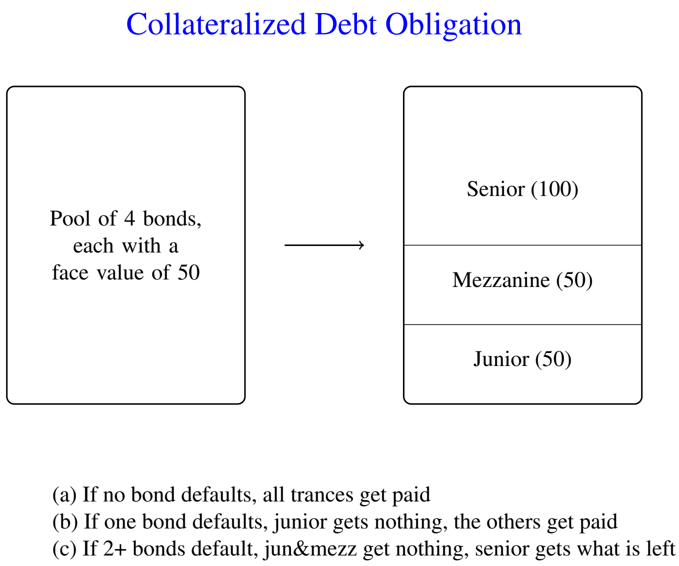

Figure 12.23: Collateralized Debt Oblication

year	Probability of survival through year t	Probability of default in year t	Expected spread payment	Expected payment from insurance	Expected PV of net payment
1	0.98	0.02	$0.98s$	0.02×0.6	$0.98s - 0.012$
2	0.95	0.03	$0.95s$	0.03×0.6	$0.95s - 0.018$
Sum					$1.93s - 0.03$

Table 12.1: Example of the payment flows of a 2-year CDS with an assumed recovery rate of 0.4 and a riskfree interest rate of zero. The CDS spread is denoted s.

A More Details on Bond Conventions

A.1 Bond Equivalent Yields on US Bonds

The financial press typically quotes a bond equivalent yield for T-bills—in an attempt to make the yields comparable. The bond equivalent yield is the coupon (and yield to maturity) of a par bond that would give the same yield as the T-bill. For a T-bill with at most half a year to maturity, this gives a simple interest rate, but for longer T-bills the expression is more complicated.

We first analyze *a T-bill with more than half a year to maturity*. Consider a coupon

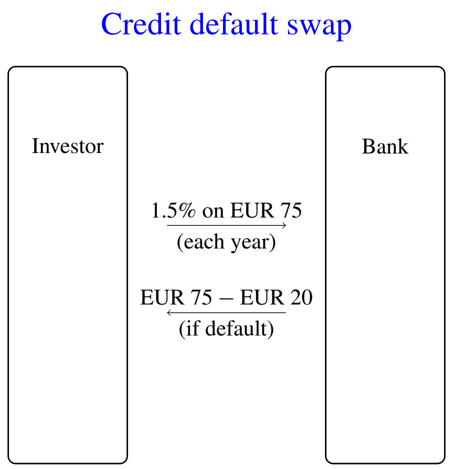

Figure 12.24: Credit default swap

bond with face value B (which equals the current price of the T-bill), semi-annual coupon $c/2$ and the same yield to maturity. Since the coupon and the yield to maturity are the same, the "clean price" of the bond (the price to pay if the seller gets to keep the accrued interest on the first coupon payment) equals the face value (here B): it is traded at par. Notice that the latter means that the buyer gets the following fraction of the next coupon payment (which is $B \times c/2$): the fraction of a half year until the next coupon payment (or (days to next coupon)$\times 2/365$).

When the T-bill has more than half a year to maturity, then the bond has two coupon payments left (including the maturity). At maturity, the owner will have the following: *(i)* the principal plus final coupon, $B(1 + c/2)$; *(ii)* the part of the first coupon that belongs to the current owner, $d = B \times 2n \times c/2$, where $n =$(days to next coupon)$/365$; and *(iii)* the interest on d when reinvested at the semi-annual rate $c/2$ for half a year, $d \times c/2$.

To get the same return as on the T-bill, the owner of the coupon bond must get a value of one at maturity (the return is then $1/B$), or

$$1 = B[1 + c/2 + 2n \times c/2 \times (1 + c/2)]. \tag{A.1}$$

Solving for c gives the bond equivalent yield

$$c = \frac{\sqrt{2n/B + 1/4 - n + n^2} - n - 1/2}{n}. \tag{A.2}$$

Example A.1 *A T-bill with 212 days to maturity and a quoted discount yield of 5.9% has*

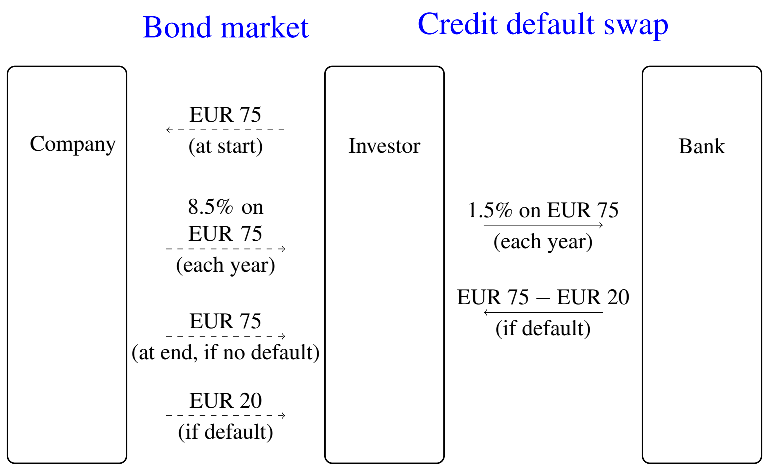

Figure 12.25: Credit default swap

the price $1 - (212/360) \times 0.059 \approx 0.965$*. There must be* $212 - 182 = 30$ *days to the next coupon payment, so* $n = 30/365$*. The bond equivalent yield is the* c *such that*

$$c = \frac{\sqrt{2\,(30/365)\,/0.965 + 1/4 - (30/365) + (30/365)^2} - (30/365) - 1/2}{(30/365)} \approx 6.2\%$$

Remark A.2 *If we define* $h = (\text{days to maturity})/365$*, then* $n = h - 1/2$ *and we can rearrange (A.2) as*

$$c = \frac{2\sqrt{h^2 + (2h-1)\,(1/B-1)} - 2h}{2h-1}.$$

This is the expression in McDonald (2006) Appendix 7.A and Blake (1990) 4.2.

We now apply the same logic to a *T-bill with at most half a year to maturity*. The bond then only has the final coupon left (which is split with the previous owner), and the face value (which is not split). In particular, there is no reinvestment. In this case, (A.1) simplifies to

$$1 = B(1 + 2n \times c/2). \tag{A.3}$$

Solving for c (and using the fact that $n = h = (\text{days to maturity})/365$) gives

$$c = \frac{1/B - 1}{h} \text{ or} \tag{A.4}$$

$$B = \frac{1}{1 + h \times c}. \tag{A.5}$$

Example A.3 *A T-bill with 44 days to maturity and a quoted discount yield of 6.21% has the price* $1-(44/360)\times 0.0621 \approx 0.992$*. The bond equivalent yield is the* c *such that*

$$0.992 = \frac{1}{1+\frac{44}{360}c} \text{ or } c = 6.6\%.$$

Remark A.4 *There are two other, but equivalent, expressions for the bond equivalent yield for maturities of at most half a year (see, for instance, McDonald (2006) Appendix 7.A). The first is*

$$c_1 = \frac{1-B}{B}\frac{1}{m}.$$

Substituting for B *using (A.5) shows that* $c_1 = c$*. The second is*

$$c_2 = \frac{365 \times Y_{db}}{360 - Y_{db} \times days}.$$

Substituting for Y_{db} *using (12.39) shows that* $c_2 = c_1 = c$*.*

Bibliography

Blake, D., 1990, *Financial market analysis*, McGraw-Hill, London.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The econometrics of financial markets*, Princeton University Press, Princeton, New Jersey.

Deacon, M., and A. Derry, 1998, *Inflation-indexed securities*, Prentice Hall Europe, Hemel Hempstead.

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

Fabozzi, F. J., 2004, *Bond markets, analysis, and strategies*, Pearson Prentice Hall, 5th edn.

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

McCulloch, J., 1975, "The tax-adjusted yield curve," *Journal of Finance*, 30, 811–830.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

Svensson, L., 1995, “Estimating forward interest rates with the extended Nelson&Siegel method,” *Quarterly Review, Sveriges Riksbank*, 1995:3, 13–26.

13 Bond Portfolios

Main references: Elton, Gruber, Brown, and Goetzmann (2010) 21–22 and Hull (2009) 4
Additional references: McDonald (2006) 7

13.1 Duration: Definitions

The "duration" of a coupon bond is used to analyse how the bond price will change in response to changes in the yield curve. This section gives the definitions of the most commonly used duration measures.

Recall that the yield to maturity, θ, of a coupon bond satisfies

$$\begin{aligned} B(K, c) &= \frac{c_1}{(1+\theta)^{m_1}} + \frac{c_2}{(1+\theta)^{m_2}} + \ldots + \frac{c_K}{(1+\theta)^{m_K}} \\ &= \sum_{k=1}^{K} \frac{c_k}{(1+\theta)^{m_k}}, \end{aligned} \tag{13.1}$$

where the bond pays c_k at m_k periods from now. The principal is included in the last "coupon" payment. We allow the payments to differ between periods—to simplify the notation and to be able to treat a bond portfolio in the same way as an ordinary bond.

The derivative of a coupon bond price with respect to its yield to maturity is

$$\frac{dB(K, c)}{d\theta} = -\frac{1}{1+\theta} \sum_{k=1}^{K} m_k \frac{c_k}{(1+\theta)^{m_k}}. \tag{13.2}$$

This measures the sensitivity of the bond price to a small change in the yield to maturity.

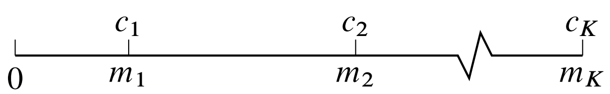

Figure 13.1: Timing convention of bond portfolio

The *dollar duration*, $D^{\$}$, is typically defined as this derivative times minus one

$$D^{\$} = -\frac{dB\left(K,c\right)}{d\theta} \tag{13.3}$$

$$= \frac{1}{1+\theta}\sum_{k=1}^{K} m_k \frac{c_k}{(1+\theta)^{m_k}}. \tag{13.4}$$

The change of the bond price, $\Delta B\left(K,c\right)$, due to a small change in the yield, $\Delta\theta$, is approximately

$$\Delta B\left(K,c\right) \approx \frac{dB\left(K,c\right)}{d\theta} \times \Delta\theta \tag{13.5}$$

$$= -D^{\$} \times \Delta\theta \tag{13.6}$$

This says that an increase in the interest rate (more precisely, the yield to maturity, θ) translates into a decrease in the bond price—and more so if the duration ($D^{\$}$) is long.

Example 13.1 *(Price of discount bond: exact and approximation) Consider a 5-year zero coupon bond with an effective interest rate of 2% or 3%. Its price at the two different interest rates would be*

$$\frac{1}{(1+0.02)^5} = 0.9057$$
$$\frac{1}{(1+0.03)^5} = 0.8626.$$

If the interest rate changes from 2% to 3%, the price change is $0.8626 - 0.9057 = -0.043$. *The dollar duration at 2% is*

$$D^{\$} = \frac{5}{1+0.02} \times 0.9057 \approx 4.44,$$

so the approximate price change according to (13.6) is

$$-4.44 \times (0.03 - 0.02) = -0.044,$$

which is close to the exact number.

It is common to divide the dollar duration by the bond price, $B\left(K,c\right)$, to get the *adjusted (or modified) duration*, D^a,

$$D^a = D^{\$} \frac{1}{B\,(K, c)}. \tag{13.7}$$

By dividing both sides of (13.6) by the bond price and using the definition of the adjusted duration we see that the relative (percentage) change of the bond price due to a small change in the yield is approximately

$$\frac{\Delta B\,(K, c)}{B\,(K, c)} \approx -D^a \times \Delta\theta \tag{13.8}$$

It is also common to multiply the dollar duration by $(1+\theta)/B\,(K, c)$ to get *Macaulay's duration*, D^M,

$$D^M = D^{\$} \frac{1+\theta}{B\,(K, c)} \tag{13.9}$$

$$= \sum\nolimits_{k=1}^{K} m_k \frac{c_k}{(1+\theta)^{m_k}\, B\,(K, c)}. \tag{13.10}$$

By multiplying both sides of (13.6) by $(1+\theta)/B\,(K, c)$ and using the definition of Macaulay's duration we see that the relative (percentage) change of the bond price due to a small relative (percentage) change in the yield is approximately

$$\frac{\Delta B\,(K, c)}{B\,(K, c)} \approx -D^M \times \frac{\Delta\theta}{1+\theta}. \tag{13.11}$$

The term last term, $\Delta\theta/(1+\theta)$, is the relative change in the gross yield—since $\Delta\theta = \Delta(1+\theta)$.

Example 13.2 *(Relative price change of discount bond) Using the same figures as in Example 13.1, Macaulay's duration is 5 (same as the time to maturity). The approximate relative price change according to (13.11) is*

$$-5 \times \frac{0.03 - 0.02}{1 + 0.02} = -0.049.$$

In contrast, the exact relative change is $(0.8626 - 0.9057)/0.9057 = -0.048$.

Notice that Macaulay's duration is a weighted average of the time to the coupon (and face) payments $(m_1, m_2, ..., m_K)$. The weight of m_k is $c_k/[(1+\theta)^{m_k}\, B\,(K, c)]$, so the weights sum to unity and they are clearly the percentage of the bond price accounted for by the respective coupon (or principal) payments. Macaulay's duration is therefore an average "time to payment" of the bond. For instance, for a zero coupon bond, Macaulay's duration is the time to maturity (set $c = 0$ in (13.10)). For bonds with coupons, Macaulay's duration is less than the time to maturity—and this effect is more pronounced at high coupon rates and at high yields to maturity. This is illustrated in Figure 13.2.

In any case, the different duration measures are fairly similar, so we can typically think of the dollar duration to represent "average time to payment."

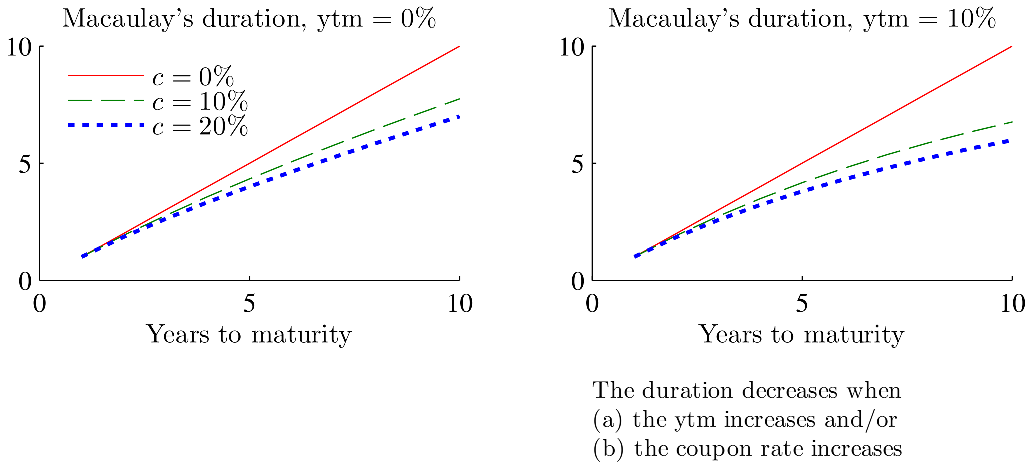

Figure 13.2: Macaulay's duration

Remark 13.3 *(Duration of a zero coupon bond) For a zero-coupon bond with a face value of unity and maturity of K, the price is $B = 1/(1+y)^K$, where y is the yield to maturity. The duration measures are*

$$D^{\$} = \frac{K}{1+y} B$$
$$D^{a} = \frac{K}{1+y}, \text{ and}$$
$$D^{M} = K.$$

Example 13.4 *(Duration) Consider a 4% (annual) coupon bond with 2 years to maturity. Suppose the price is 1.019. The the yield to maturity is 3% since it solves*

$$1.019 \approx \frac{0.04}{1+0.03} + \frac{1.04}{(1+0.03)^2}.$$

The dollar duration is

$$D^{\$} = \frac{1}{1.03}\left[\frac{0.04}{1.03} + 2\frac{1.04}{1.03^2}\right] \approx 1.94,$$

so the adjusted duration and Macaulay's duration are

$$\begin{aligned} D^a &= 1.94\frac{1}{1.019} \approx 1.90 \\ D^M &= 1.94\frac{1.03}{1.019} \approx 1.96. \end{aligned}$$

Example 13.5 *(Duration of a zero coupon bond) A two-period zero coupon bond with price 0.94 has a ytm equal to 0.03, since*

$$0.94 \approx \frac{1}{1.03^2}.$$

The dollar duration is

$$\frac{1}{1.03}2\frac{1}{1.03^2} \approx 1.83,$$

and Macaulay's duration is

$$\frac{1.03}{1/1.03^2} \times \frac{1}{1.03}2\frac{1}{1.03^2} = 2.$$

The duration of a portfolio of bonds is easily calculated from the durations of the individual bonds—at least if the bonds have the same ytm. This is summarised in the next proposition.

Proposition 13.6 *(Duration of a portfolio) If the yield to maturities of bond i and j are the same, then a portfolio of both bonds has the dollar duration $D_i^{\$} + D_j^{\$}$ and the Macaulay's duration $B_i/(B_i+B_j)D_i^M + B_j/(B_i+B_j)D_j^M$ (the value weighted average of the different Macaulay's durations). If the ytms are different, this does not hold.*

Proof. (Duration of a portfolio*) The first part is intuitive since the dollar duration of a coupon bond is considered "correct"—and it uses the same ytm for all the coupons.).

For the second part, multiply the dollar duration $D_i^\$ + D_i^\$$ by the ytm and divide by the portfolio value $(B_i + B_j)$. This is Macaulay's duration of the portfolio. Now, rewrite by using $D^\$ = BD^M/(1+\theta)$ to get the result in the proposition. ■

Example 13.7 *(Duration of a portfolio, same ytm) A 1-year discount bond with a ytm (effective interest rate) of 10% has the price* $1/1.1$ *and a 3-year discount bond with a ytm of 10% has the price* $1/1.1^3$*.The dollar duration and Macaulay's durations are*

$$\begin{aligned}
\text{1-year bond: } D^\$ &= \frac{1}{1.1^2} \approx 0.83 \text{ and } D^M = 1 \\
\text{3-year bond: } D^\$ &= \frac{3}{1.1^4} \approx 2.05 \text{ and } D^M = 3.
\end{aligned}$$

A portfolio with one of each bond has a price equal $B_p = 1/1.1 + 1/1.1^3$ *and a ytm*

$$B_p = \frac{1}{1+\theta} + \frac{1}{(1+\theta)^3}, \text{ with } \theta = 0.1.$$

The duration and Macaulay's duration of the portfolio are then

$$\begin{aligned}
D^\$ &= \frac{1}{1.1}\left[\frac{1}{1.1} + 3\frac{1}{1.1^3}\right] \approx 2.88, \\
D^M &= D^\$ \frac{1.1}{B_p} \approx 1.90.
\end{aligned}$$

Compare with

$$\begin{gathered}
0.83 + 2.05 \approx 2.88 \text{ and} \\
\frac{1/1.1}{B_p} + \frac{1/1.1^3}{B_p} 3 \approx 1.90,
\end{gathered}$$

which are the same.

Example 13.8 *(Duration of a portfolio, different ytm) A 1-year discount bond with a ytm (effective interest rate) of 7% has the price* $1/1.07$ *and a 3-year discount bond with a ytm of 10% has the price* $1/1.1^3$*.The dollar duration and Macaulay's durations are*

$$\begin{aligned}
\text{1-year bond: } D^\$ &= \frac{1}{1.07^2} \approx 0.87 \text{ and } D^M = 1 \\
\text{3-year bond: } D^\$ &= \frac{3}{1.1^4} \approx 2.05 \text{ and } D^M = 3.
\end{aligned}$$

A portfolio with one of each bond has a price $B_p = 1/1.07 + 1/1.1^3$ *and a ytm*

$$B_p = \frac{1}{1+\theta} + \frac{1}{(1+\theta)^3}, \text{ with } \theta \approx 0.091.$$

The duration and Macaulay's duration of the portfolio are then

$$D^\$ = \frac{1}{1.091}\left[\frac{1}{1.091} + 3\frac{1}{1.091^3}\right] \approx 2.96,$$
$$D^M = D^\$\frac{1.091}{B_p} \approx 1.91.$$

Compare with

$$0.87 + 2.05 \approx 2.92 \text{ and}$$
$$\frac{1/1.07}{B_p} + \frac{1/1.1^3}{B_p}3 \approx 1.89,$$

which are slightly different.

13.2 Duration Matching

13.2.1 Basic Idea

Suppose we want to hedge against price movements of a bond portfolio or a liability stream. (This is also called immunisation.) The portfolio can be thought of as a coupon bond (with a possibly complicated set of coupons), so the previous formulas apply. One was of doing that would be to use a (potentially large) set of futures—to match every cash flow of the bond, but that may well be both difficult and costly (transaction costs). Duration matching is the other extreme: finding a single instrument to use in the hedging.

Example 13.9 *(Why a liability is not hedged by putting its present value on a bank account) Suppose your liability stream is that you owe 150 next year and 250 the year after that. Suppose all interest rates are 5%. The present value of your liability stream is*

$$\frac{150}{1.05} + \frac{250}{1.05^2} = 369.615.$$

Put 369.615 in a bank account. Next day, the interest rate decreases to 4%. The present

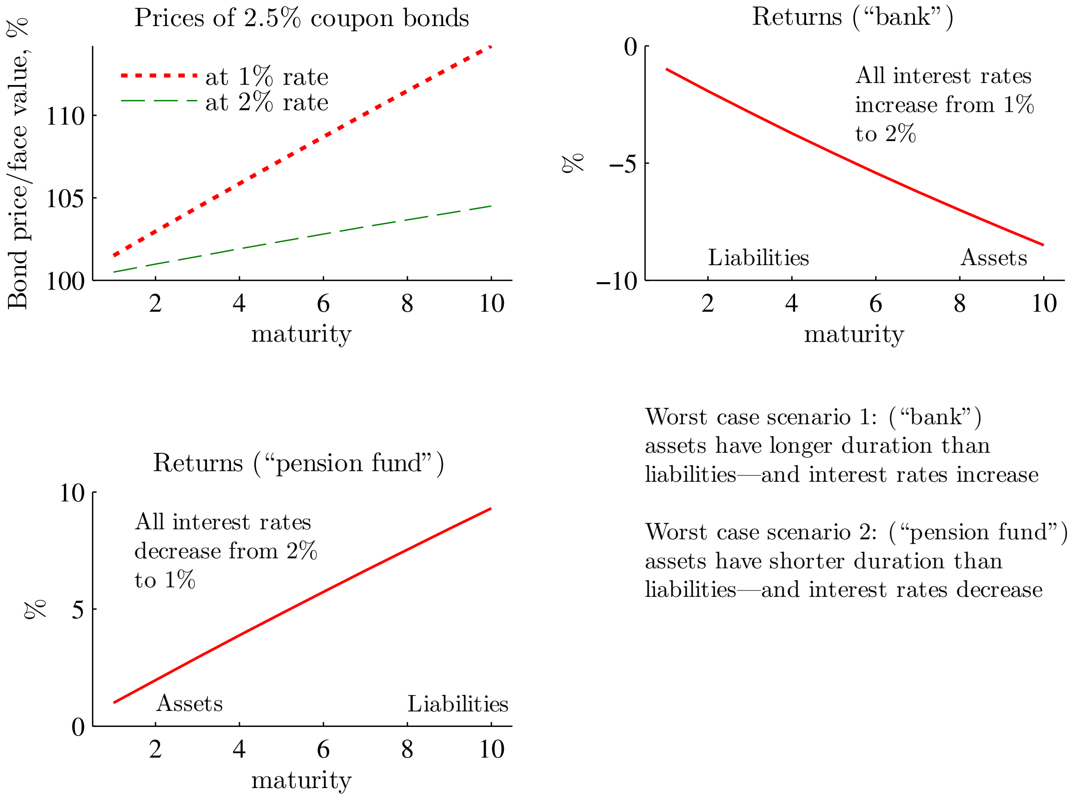

Figure 13.3: Gains and losses at interest rate changes

value (market value) of the liability stream is

$$\frac{150}{1.04} + \frac{250}{1.04^2} = 375.37,$$

which is more than you put away in the bank (369.615).

Example 13.10 *(continued) Another perspective on the same problem. With* 369.615 *in the bank and 4% interest rate, you have*

$$369.615 \times 1.04 = 384.4 \textit{ at end of year 1.}$$

Use 150 to cover the cover the payment—leaving 234.4. After another year (still 4% rate) you have

$$234.4 \times 1.04 = 243.776 \textit{ at end of year 2.}$$

	5% interest rate		5% → 4% interest rate on day 1			
	bank account		bank account		3.9962 1.6-year bonds	
	Flow	Balance	Flow	Balance	Flow	Balance
Day 0	369.615	369.615	369.615	369.615	369.615	369.615
Day 1		369.615		369.615		375.313
Year 1	18.481 − 150	238.096	14.785 − 150	234.400	15.012 − 150	240.326
Year 2	11.905 − 250	0.001	9.376 − 250	−6.224	9.613 − 250	−0.061

Table 13.1: Effect of interest rate decrease on the possibility to finance a liability. With the 1.6-year bonds, it is assumed that we sell them after the interest rate decrease and put the balance on a bank account.

Not enough to cover the 250. See Table 13.1.

A liability is the same as being short one unit of a bond (portfolio) with price B_L and dollar duration $D_L^\$$. We will hedge this portfolio by buying h units of a bond portfolio with price B_H and dollar duration $D_H^\$$. The value of the overall position is then

$$V = hB_H - B_L. \tag{13.12}$$

(To keep the notation simple, we use the short hand notation B to indicate a bond price.)

Often, the balance, $-V$, is held in a money market account (which has a zero duration) to make the initial value of the V and money market account zero

$$0 = \underbrace{h_t B_{H,t} + M_t}_{\text{hedging portfolio}} - B_{L,t}, \text{ so } M_t = -h_t B_{H,t} + B_{L,t}. \tag{13.13}$$

The subsequent value of the money market account will change as the interest rates are collected/paid and the gains and losses of the bonds accumulate.

Remark 13.11 *(Overall portfolio value over several subperiods*) Start by creating a portfolio with a zero initial value*

$$0 = h_t B_{H,t} - B_{L,t} + M_t, \text{ so } M_t = 0 - h_t B_{H,t} + B_{L,t},$$

where M_t is the amount held in a money market account (almost zero duration) with an *interest rate y_t. In $t + s$ (say, after a day so $s = 1/365$), this portfolio is worth*

$$V_{t+s} = h_t(B_{H,t+s} + c_{H,t+s}) - (B_{L,t+s} + c_{L,t+s}) + M_t(1 + y_t)^s,$$

where $c_{H,t+s}$ and $c_{L,t+s}$ are the coupon payments (or any other cash flows) and the bond prices are measured after coupons. After rebalancing in $t + s$, we need h_{t+s} units of bond H and we are still short one bond L, so the balance on the money market account is

$$M_{t+s} = V_{t+s} - h_{t+s}B_{H,t+s} + B_{L,t+s}.$$

This is similar to the expression for M_t in the first equation, except that V_{t+s} may be non-zero. The value of the portfolio in $t + 2s$ is computed as in the second equation, but with subscripts advanced one period.

Using the approximate relation of the bond price change (13.6) we have that the change of value (due to a sudden change in the interest rates) of the overall position is

$$\begin{aligned} \Delta V &\approx h\frac{dB_H}{d\theta} \times \Delta\theta_H - \frac{dB_L}{d\theta} \times \Delta\theta_L &\quad (13.14)\\ &= -hD_H^{\$} \times \Delta\theta_H + D_L^{\$} \times \Delta\theta_L. &\quad (13.15) \end{aligned}$$

where the durations are dollar durations. Alternatively, in terms of Macaulay's duration (recall that $D^M = D^{\$}(1 + \theta)/B$) , this can be written

$$\Delta V \approx -hD_H^M B_H \times \frac{\Delta\theta_H}{1 + \theta_H} + D_L^M B_L \times \frac{\Delta\theta_L}{1 + \theta_L}. \tag{13.16}$$

Several of the hedging approaches discussed below) assume that $\Delta\theta_H = \Delta\theta_L$, that is, a parallel shift of the yield curve. The weakness of that assumption is also discussed below.

13.2.2 Naive Hedging (Flat Yield Curve)

Suppose we just set $h = B_L/B_H$. Combining with (13.16) and assuming a flat yield curve that shifts up/down gives

$$\frac{\Delta V}{B_L} \approx \left(D_L^M - D_H^M\right) \times \frac{\Delta\theta}{1 + \theta}. \tag{13.17}$$

For instance, suppose interest rates decrease ($\Delta\theta < 0$) and the duration of the liability is longer than of the hedge bond ($D_L^M > D_H^M$). Then, the portfolio will lose money. The reason is that the value of the liability goes up more than the value of the hedge bond—as longer bonds are more sensitive to interest rate changes than short bonds. To overcome this problem, we need to invest more than proportionally into the hedge bond.

Remark 13.12 *(Effect yield of curve shift on a bank) A bank typically has liabilities with short duration (deposits, inter-bank lending) and assets with long duration (loans to companies and households), so $D_L^M < D_H^M$. Equation (13.17) shows that an increase in the interest rate level will hurt the bank: the assets decrease more than the liabilities. This can also be phrased as follows: the bank has fixed incomes from the loans it has made, but it now needs to refinance itself (deposits and inter-bank loans) at a higher cost.*

13.2.3 Duration Hedging (Parallel Shift of the Yield Curve)

Suppose we instead choose

$$h = \frac{D_L^M}{D_H^M} \times \frac{B_L}{B_H}, \text{ so} \tag{13.18}$$

$$hB_H = \frac{D_L^M}{D_H^M} \times B_L, \tag{13.19}$$

where the second equation shows the *amount* (dollars) invested into the hedge bond.

Combine (13.16) and the hedge ratio (13.18) to get that the relative change in the portfolio value is approximately

$$\frac{\Delta V}{B_L} \approx D_L^M \times \left(\frac{\Delta\theta_L}{1+\theta_L} - \frac{\Delta\theta_H}{1+\theta_H} \right), \tag{13.20}$$

which shows that the overall portfolio will not change ($\Delta V \approx 0$) if the yields change equally much. See Figure 13.4 for an illustration.

For instance, if the two bonds have the same Macaulay's duration, then the positions in the two bonds are worth the same ($hB_H = B_L$), making the value of the V portfolio zero (initially). More generally, we invest more into the hedge bond H if bond L has a longer duration, et vice versa. The intuition is that a the price of a bond with long duration is more sensitive to a yield curve shift than the price of a short bond. Therefore, to hedge

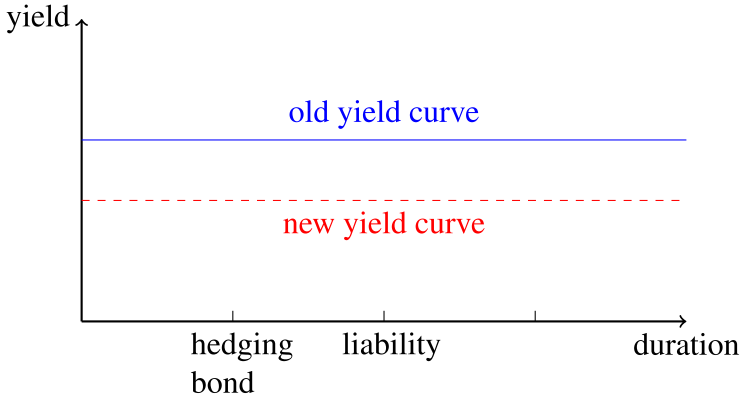

The relative change of overall portfolio value is

$$D_L^M \times \left(\frac{\Delta\theta_L}{1+\theta_L} - \frac{\Delta\theta_H}{1+\theta_H}\right), \text{ eg. } 3 \times (-0.01 - -0.01) = 0$$

Figure 13.4: Parallell shift of yield curve

a bond with a long duration (as bond L) we need to buy more of the bond with a short duration (bond H). See Figure 13.5 for an illustration.

Remark 13.13 *(Using the dollar duration instead*) If we set $\Delta V = 0$ in (13.15), then*

$$h = \frac{D_L^\$}{D_H^\$} \times \frac{\Delta\theta_L}{\Delta\theta_H}.$$

If we assume that both yields change equally much, then

$$h = \frac{D_L^\$}{D_H^\$}.$$

Example 13.14 *(Duration hedging of the liability in Example 13.9) Instead of putting* 369.615 *on a bank account, buy 1.6-year bonds for this amount—since Macaulay's duration of the liability is 1.6 years. According to (13.19) we should then choose the number of 1.6 year bonds bough (h) so that the investment equals the value of the liability* (369.615)*. The price of each bond (with face value 100) is* $100/1.05^{1.6}$*, so you buy h of them*

$$h\frac{100}{1.05^{1.6}} = 369.615 \text{ or } h = 3.9962.$$

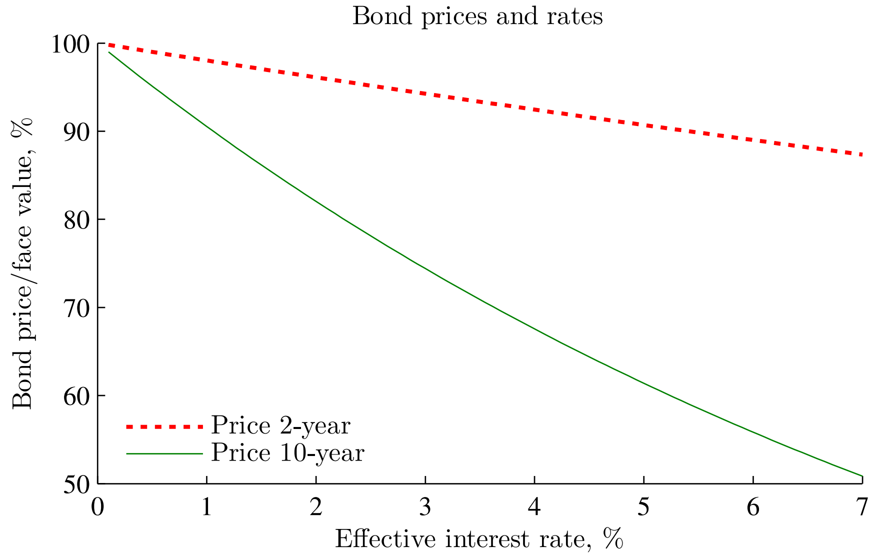

Figure 13.5: Interest rate vs bond prices

The value of this position after the interest rate has changed to 4% *is*

$$3.9962 \times \frac{100}{1.04^{1.6}} = 375.313,$$

which is almost the same as the PV of the liability stream after the interest rate chance. You could therefore sell your bond and put the money in a bank account. It would be enough to pay the liabilities—if there were no further interest changes...

13.2.4 Problem 1: Approximation Error

The formula for the price change (13.6) is only exact for infinitesimal yield changes—and the approximation error is likely to be large when the yield changes are.

The formula is really a first-order Taylor approximation of the form

$$\Delta B \approx \frac{dB}{d\theta} \times \Delta\theta. \tag{13.21}$$

Obviously, a second-order Taylor approximation is more precise. It would be

$$\Delta B \approx \frac{dB}{d\theta} \times \Delta\theta + \frac{1}{2}\frac{d^2 B}{d\theta^2} \times (\Delta\theta)^2. \tag{13.22}$$

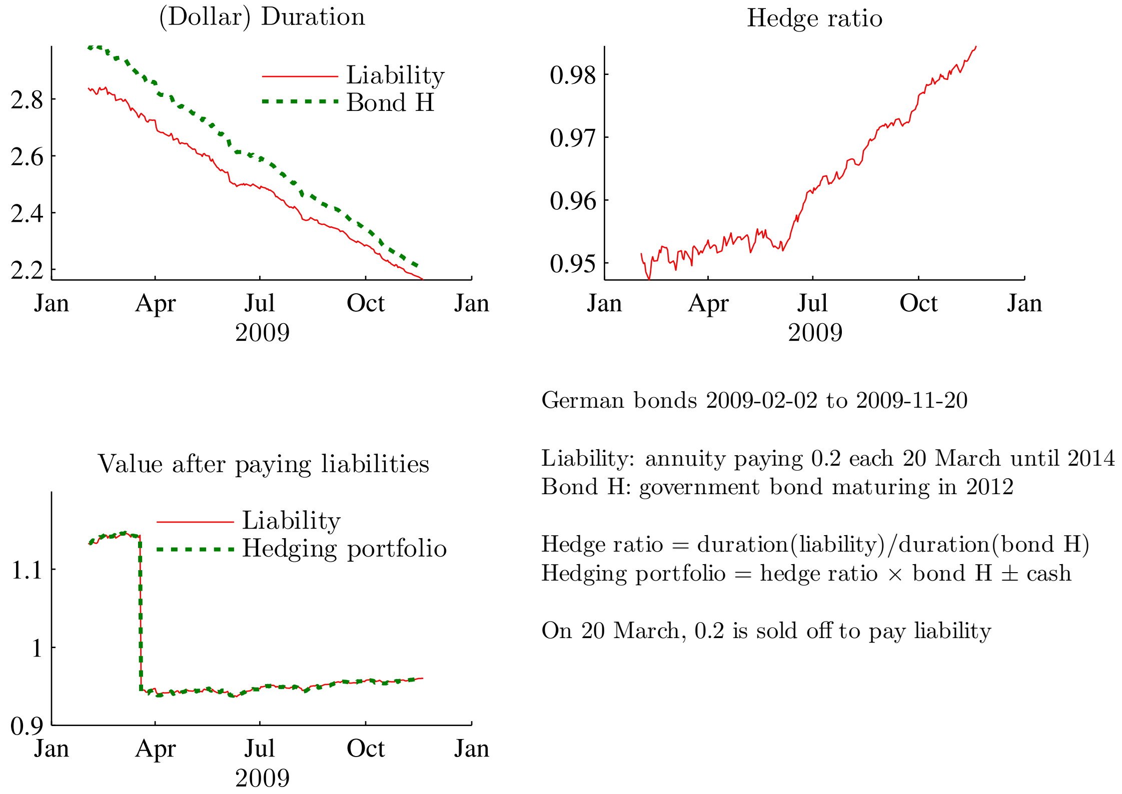

Figure 13.6: Duration hedging

where the last term includes the second derivative of the bond price with respect to the yield to maturity. The second derivative is easily calculated to be

$$\frac{d^2 B\,(K,c)}{d\theta^2} = \textstyle\sum_{k=1}^{K} m_k(m_k+1)\dfrac{c_k}{(1+\theta)^{m_k+2}}. \tag{13.23}$$

Dividing (13.22) by the bond price and using (13.8) gives

$$\frac{\Delta B}{B} \approx -D^a \times \Delta\theta + \frac{1}{2}C \times (\Delta\theta)^2, \tag{13.24}$$

where C (often called “convexity”) is the second derivative in (13.22) divided by the bond price. It can be shown that the convexity is positive, but decreasing in the coupon rate — for a given ytm and maturity. (The convexity is actually increasing in the coupon rate for a given ytm and modified duration.) See Figure 13.5 for an illustration of the non-linear

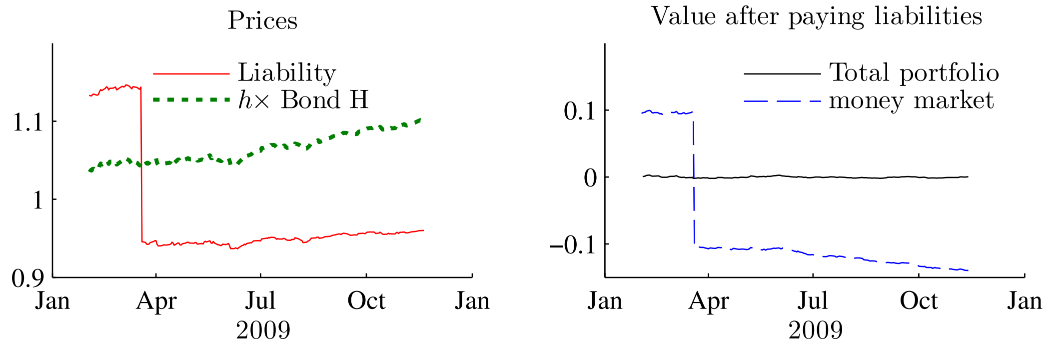

Figure 13.7: Duration hedging

effect.

By choosing the hedging bond (portfolio) so that it has a similar convexity to the bond to be hedged may make the hedge more precise.

Example 13.15 *(Convexity) The convexity of the bond in Example 13.4 is*

$$C = \frac{1}{1.019}\left[1(1+1)\frac{0.04}{1.03^3} + 2(2+1)\frac{1.04}{1.03^4}\right] \approx 5.51.$$

For a zero coupon bond in Example 13.5 (which has the same ytm and maturity), the convexity is

$$C = \frac{1}{1/1.03^2}\left[2(2+1)\frac{1}{1.03^4}\right] = \frac{6}{1.03^2} \approx 5.66.$$

13.2.5 Problem 2: Changing Cash Flows

The duration measures assume that the times when the coupons and the face value are paid are unaffected by the yield change. That is true for many instruments (like most government bonds), but not for callable bonds—and effectively not for bonds whose risk premium depends on the interest rate level as most corporate bonds do (as the interest rate level affects the default risk).

13.2.6 Problem 3: Yield Curve Changes vs. Changes in Yields to Maturity

The probably most important problem with using duration for hedging is that the hedge ratio in (13.25) depends on how the yields change—and that is not known when we construct the hedging portfolio.

The ideal case for duration hedging is when the yields (to maturity) move in parallel. This will be the case, for instance, if the yield curve is flat (across maturities)—and the only movements are parallel shifts up and down. In reality, most movements in the yield curve are parallel, but changes in slope and curvature are not uncommon either. Often the short interest rates move more (in response to news) than long rates.

Equation (13.20) shows how the value of the overall portfolio depends on the yields of the liability and the hedge bond. For instance, suppose the yield curve changes from being flat to being downward sloping and the hedging bond has shorter duration than the liability. In this case, the overall portfolio loses value. The reason is that the value of the hedging portfolio increases less (the yield decreases less) in price than the liability. See Figure 13.8 for an illustration.

To overcome this problem, the hedge ratio should be (set $\Delta V = 0$ in (13.16))

$$h = \frac{D_L^M}{D_H^M} \times \frac{B_L}{B_H} \times \frac{\Delta\theta_L/(1+\theta_L)}{\Delta\theta_H/(1+\theta_H)}. \tag{13.25}$$

This is indeed the same as (13.18) if the yield curve is flat and all changes are parallel shifts. However, the relative frequencies of the yield curve movements (level, slope, curvature) seem to change over time (according to business cycle conditions and monetary policy regime). This suggests that the ability of a simple duration matching to provide a hedge is different in different time periods and different markets.

Explicit models of how the entire yield curve moves in response to a small number of factors have implications for how the two yields will change—which may vary across instruments and time. This would allow us to also model how $\Delta\theta_L/\Delta\theta_H$ in (13.25) would react to the drivers of the yield curve—and this provide a more precise hedge ratio.

Bibliography

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

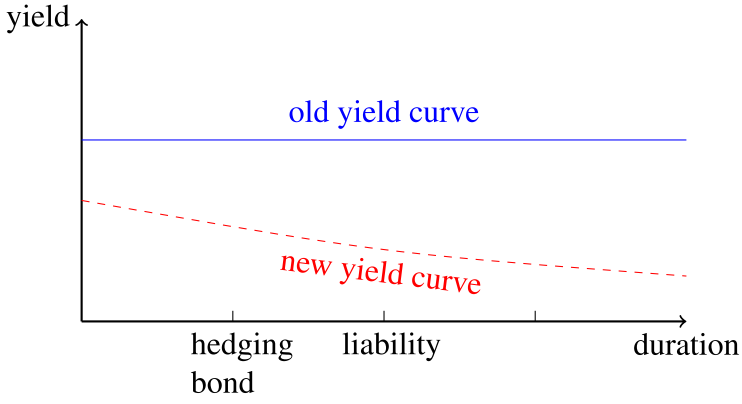

Relative change of overall portfolio value is

$D_L^M \times \left(\frac{\Delta\theta_L}{1+\theta_L} - \frac{\Delta\theta_H}{1+\theta_H}\right)$,
eg. $3 \times (-0.01 - -0.005) = -0.015$

Figure 13.8: Change of yield curve slope—and effect on hedging

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

14 Interest Rate Models

Main references: Elton, Gruber, Brown, and Goetzmann (2010) 21–22 and Hull (2009) 4
Additional references: McDonald (2006) 7

14.1 Yield Curve Models

14.1.1 Some Basic Facts

Yield curves (in the US and most other developed countries) tend to have the following features (see Figure 14.1 for some examples).

First, most of the time, the yield curve is upward sloping. This is only consistent expectations hypothesis if short rates are expected to be higher in the future. This means that short rates should (most of the time) be increasing over time—which contradicts empirical evidence. It is more likely that long rates tend to be high because of risk premia. See Figure 14.3 for an illustration.

Second, the yield curve changes over time. It is common to describe the movements in terms of three "factors": level, slope, and curvature. One way of measuring these factors is by defining

$$
\begin{aligned}
\text{Level}_t &= y_{10y} \\
\text{Slope}_t &= y_{10y} - y_{3m} \\
\text{Curvature}_t &= \left(y_{2y} - y_{3m}\right) - \left(y_{10y} - y_{2y}\right). \qquad (14.1)
\end{aligned}
$$

This means that we measure the level by a long rate, the slope by the difference between a long and a short rate—and the curvature (or rather, concavity) by how much the medium/short spread exceeds the long/medium spread. For instance, if the yield curve is hump shaped (so y_{2y} is higher than both y_{3m} and y_{10y}), then the curvature measure is positive. In contrast, when the yield curve is U-shaped (so y_{2y} is lower than both y_{3m} and y_{10y}), then the curvature measure is negative. See Figure 14.4 for an example.

An alternative is to use principal component analysis. See the Remarks below, and

Figure 14.5 for an example.

Most evidence on US data suggest (see, for instance, Cochrane (2001) 19) that changes in the level dominate—perhaps accounting for 80–90% of the total variation in yields. The slope comes second (perhaps accounting for 10%), and hump third (accounting for a few percent). Similar results are found by principal component analysis.

Remark 14.1 *(Principal component analysis*) The first (sample) principal component of the zero mean $N \times 1$ vector $z_t = Y_t - \bar{Y}$ is $w_1' z_t$ where w_1 is the eigenvector associated with the largest eigenvalue of $\Sigma = \text{Cov}(z_t)$. This value of w_1 solves the problem $\max_w w'\Sigma w$ subject to the normalization $w'w = 1$. This eigenvalue equals $\text{Var}(w_1' z_t) = w_1' \Sigma w_1$. The jth principal component solves the same problem, but under the additional restriction that $w_i' w_j = 0$ for all $i < j$. The solution is the eigenvector associated with the jth largest eigenvalue (which equals $\text{Var}(w_j' z_t) = w_j' \Sigma w_j$). This means that the first K principal components are those (normalized) linear combinations that account for as much of the variability as possible—and that the principal components are uncorrelated ($\text{Cov}(w_i' z_t, w_j' z_t) = 0$)). Dividing an eigenvalue with the sum of eigenvalues gives a measure of the relative importance of that principal component (in terms of variance). If the rank of Σ is K, then only K eigenvalues are non-zero.*

Remark 14.2 *(Principal component analysis 2*) Let W be NxN matrix with w_i as column i. We can the calculate the $Nx1$ vector of principal components as $pc_t = W' z_t$. Since $W^{-1} = W'$ (the eigenvectors are orthogonal), we can invert as $z_t = W pc_t$, that is, $Y_t = \bar{Y} + W pc_t$. The w_i vector (column i of W) therefore shows how the different elements in Y_t change as the ith principal component changes. As an example, with three yields, and the ith principal component in the column vector $[w_i^{(1)}, w_i^{(2)}, w_i^{(3)}]'$, we have*

$$\begin{bmatrix} pc1_t \\ pc2_t \\ pc3_t \end{bmatrix} = \begin{bmatrix} w_1^{(1)} & w_2^{(1)} & w_3^{(1)} \\ w_1^{(2)} & w_2^{(2)} & w_3^{(2)} \\ w_1^{(3)} & w_2^{(3)} & w_3^{(3)} \end{bmatrix}' \begin{bmatrix} Y_t(1) - \bar{Y}_t(1) \\ Y_t(3) - \bar{Y}_t(3) \\ Y_t(5) - \bar{Y}_t(5) \end{bmatrix} \text{ and}$$

$$\begin{bmatrix} Y_t(1) \\ Y_t(3) \\ Y_t(5) \end{bmatrix} = \begin{bmatrix} \bar{Y}_t(1) \\ \bar{Y}_t(3) \\ \bar{Y}_t(5) \end{bmatrix} + \begin{bmatrix} w_1^{(1)} & w_2^{(1)} & w_3^{(1)} \\ w_1^{(2)} & w_2^{(2)} & w_3^{(2)} \\ w_1^{(3)} & w_2^{(3)} & w_3^{(3)} \end{bmatrix} \begin{bmatrix} pc1_t \\ pc2_t \\ pc3_t \end{bmatrix}.$$

For instance, the second eigenvector (second column of the W matrix) shows how each of the yields "react" to the second principal component.

Remark 14.3 *(Diebold-Li approach*) Diebold and Li (2006) use the Nelson-Siegel model for an m-period interest rate as*

$$y(m) = \beta_0 1 + \beta_1 \frac{1-\exp(-m/\tau_1)}{m/\tau_1} + \beta_2 \left[\frac{1-\exp(-m/\tau_1)}{m/\tau_1} - \exp\left(-\frac{m}{\tau_1}\right) \right],$$

and set $\tau_1 = 1/(12 \times 0.0609)$. *For a given trading date, construct the factors (the terms multiplying the beta coefficients) for each bond. Then, run a regression of the cross-section of yields on these factors—to estimate the beta coefficients. Repeat this for every trading day—and plot the three time series of the coefficients. See Figure 14.6 for an example. The results are very similar to the factors calculated directly from yields (cf. Figure 14.4).*

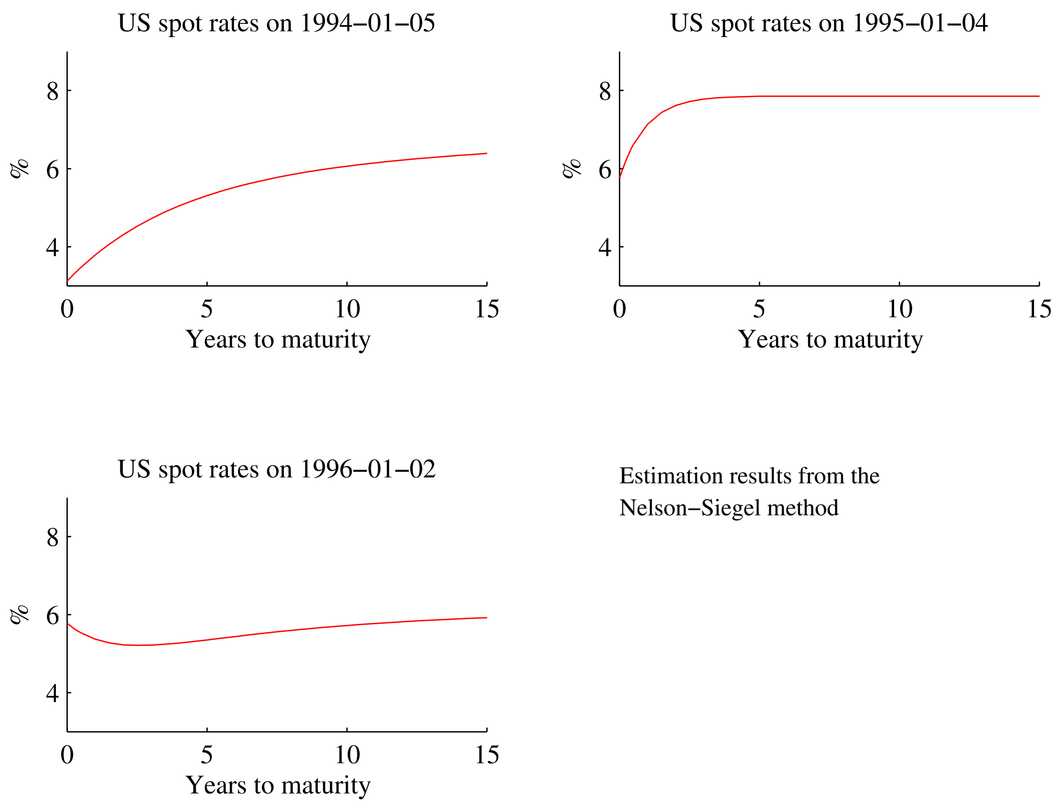

Figure 14.1: Estimated yield curve, Nelson-Siegel method

Interest rates are strongly related to business cycle conditions, so it often makes sense to include macro economic data in the modelling. See Figure 14.7 for how the term

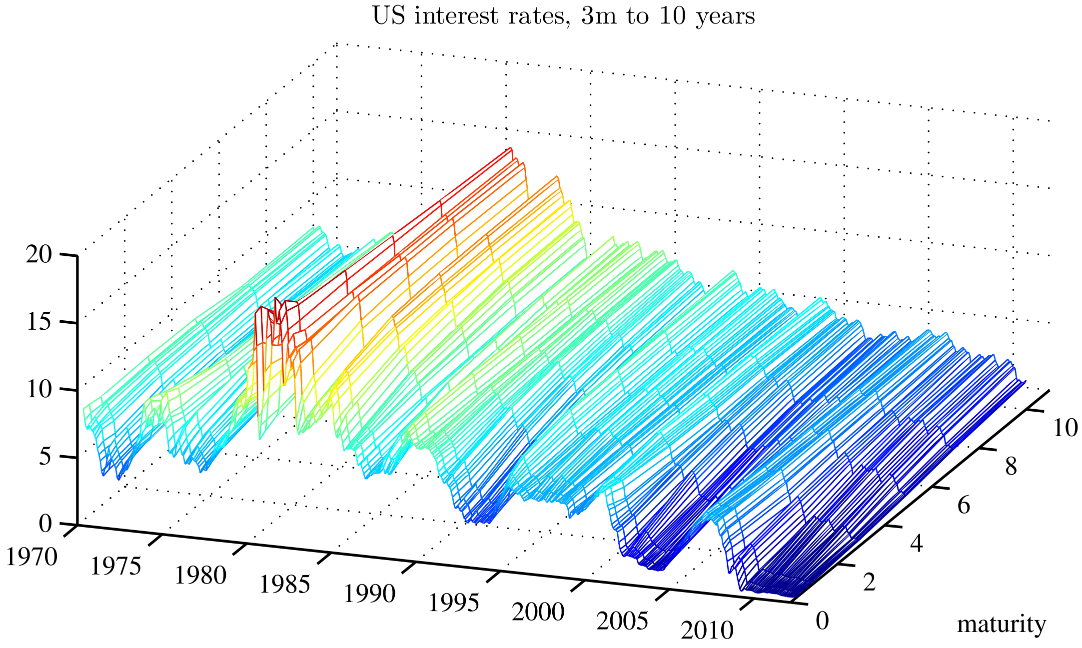

Figure 14.2: US yield curves

spreads are related to recessions: term spreads typically increase towards the end of recessions. The main reason is that long rates increase before short rates. There is also a tendency for the term spreads to be very small (or even negative) at the beginning of recessions.

14.1.2 The Expectations Hypothesis of Interest Rates

The expectations hypothesis of interest rates says that long bonds have no, or possibly constant, risk premia. In that case, forward interest rates can be interpreted as expected future short interest rates. The evidence on the expectations hypothesis is mixed (see Section 14.3.4), so it can only be thought of as a rough (although convenient) approximation.

To illustrate how the expectations hypothesis works, it is easiest to work with continuously compounded interest rates. Recall that a continuously compounded interest rate, $y_t(m)$, satisfies

$$\frac{1}{B_t(m)} = \exp\left[m y_t(m)\right], \text{ or } y_t(m) = -\ln B_t(m)/m, \tag{14.2}$$

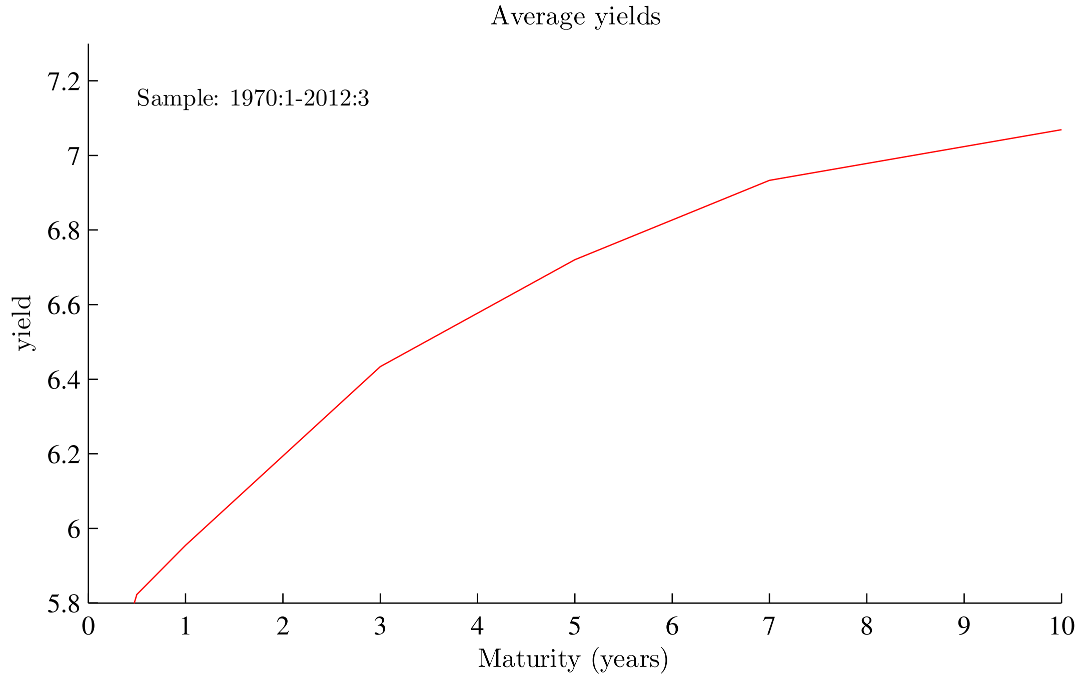

Figure 14.3: Average US yield curve

where $B_t(m)$ is the price (in t) of a zero-coupon bond which matures in $t + m$.

The expectations hypothesis says that the current long rate equals the expected future average short interest rate. Split up the time until m into m intervals of length 1. Then, the expectations hypothesis says that the m-period spot rate equals the geometric average of the 1-period short rates over t to $t + m$

$$y_t(m) = \lambda(m) + \frac{1}{m}[y_t(1) + \mathrm{E}_t\, y_{t+1}(1) + \mathrm{E}_t\, y_{t+2}(1) + \ldots]. \tag{14.3}$$

If $\lambda(m) = 0$, then the pure expectations are said to hold. Hence, the expectations hypothesis (although not in its pure form) allows for constant risk premia.

See Figure 14.8 for an illustration.

When the short interest rate is denoted r_t, then (14.3) can be written

$$y_t(m) = \lambda(m) + \frac{1}{m}[r_t + \mathrm{E}_t\, r_{t+1} + \mathrm{E}_t\, r_{t+2} + \ldots]. \tag{14.4}$$

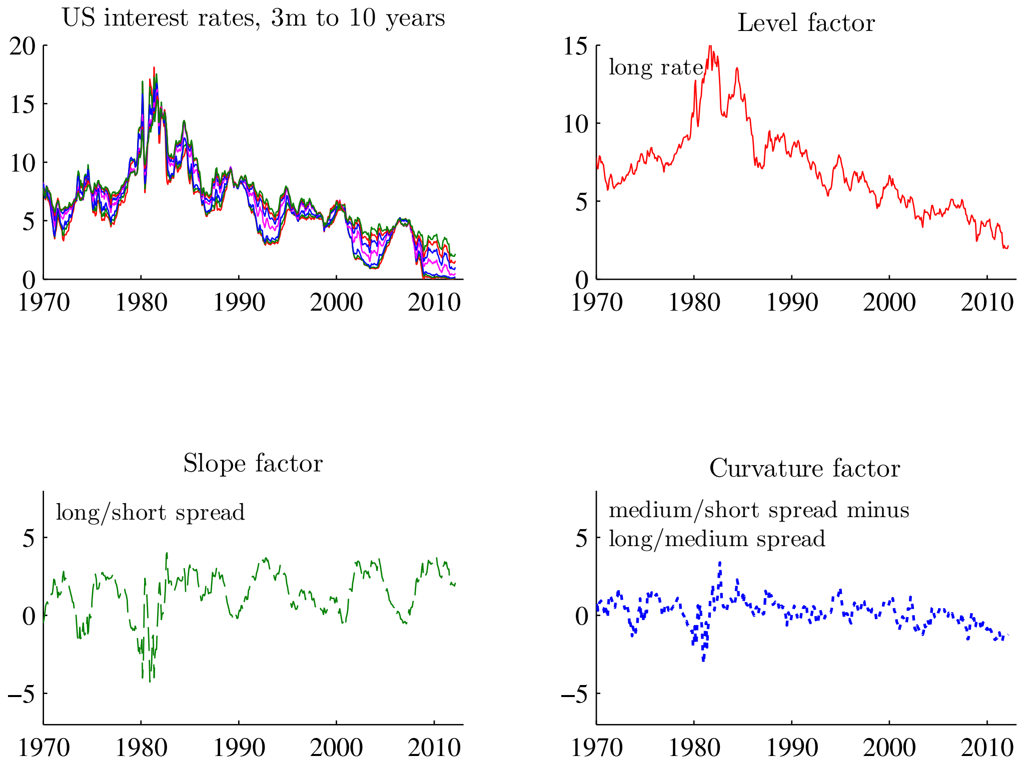

Figure 14.4: US yield curves: level, slope and curvature

14.1.3 Risk Premia

There are several reasons for why bonds should have risk premia. First, the real return of a long bond is very sensitive to inflation changes—probably more than equity. Bonds are therefore likely to have inflation risk premia. Second, long bonds are risky for investors who don't intend to keep them until maturity—and will therefore have term premia. Third, some bonds are not traded much (for instance, off-the-run bonds and many index-linked bonds)—so they are likely to have liquidity premia.

14.1.4 A Simple One-Factor Model: The Vasicek Model

The Vasicek model assumes that the state variable (the short interest rate) is an AR(1). The specification typically involves shifting the mean of the process to allow for a risk (term) premia. To simplify, I will crudely assume that there are some unspecified constant premia (the expectations hypothesis). (The more general formulation derives the risk premia in

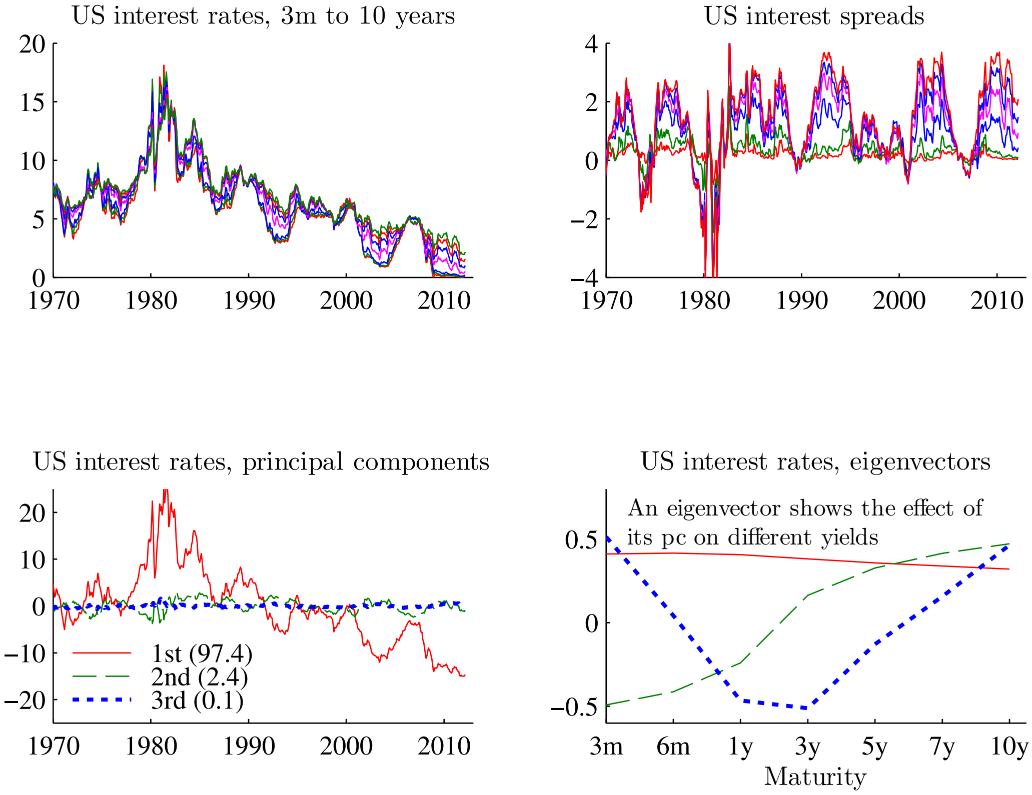

Figure 14.5: US yield curves and principal components

terms of the mean reversion and volatility of the short rate.)

To simplify the notation, let r_t be the state variable (short rate). It follows an AR(1)

$$r_{t+1} - \mu = \rho (r_t - \mu) + \varepsilon_{t+1}, \tag{14.5}$$

where μ is the mean.

Remark 14.4 *(Alternative formulation of (14.5)) The process is sometimes specified in terms of changes as*

$$r_{t+1} - r_t = a (\mu - r_t) + \varepsilon_{t+1}.$$

Clearly, this can be written

$$r_{t+1} - \mu = (1 - a)(r_t - \mu) + \varepsilon_{t+1},$$

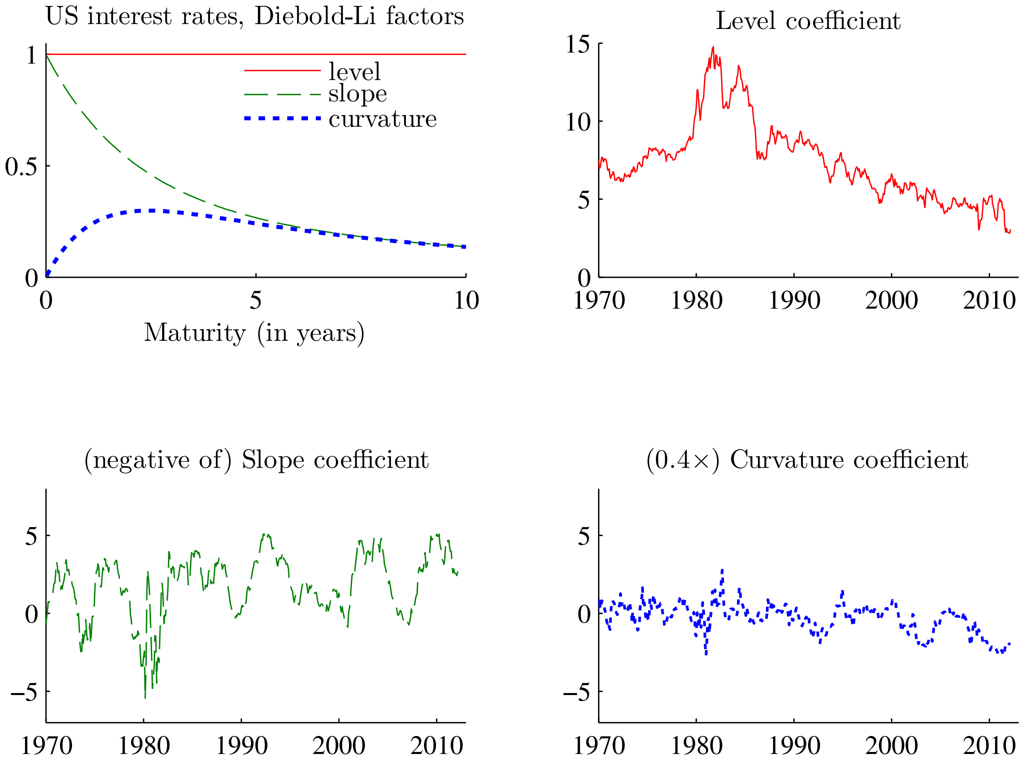

Figure 14.6: US yield curves: level, slope and curvature, Diebold-Li approach

which is of the same form as before. With $0 < a < 1$ *(that is, with* $0 < \rho < 1$*) the process is mean reverting.*

The forecast for $t + m$ is

$$\mathrm{E}_t\, r_{t+m} = (1 - \rho^m)\, \mu + \rho^m r_t. \tag{14.6}$$

We now assume that the expectations hypothesis of interest rates holds. Using this in (14.4) gives the long interest rate. For instance, the $m = 2$ rate is

$$\begin{aligned} y_t(2) &= a(2) + \frac{1}{2}\left[r_t + (1 - \rho)\, \mu + \rho r_t\right] \\ &= a(2) + \mu\,(1 - \rho)/2 + r_t\,(1 + \rho)/2. \end{aligned} \tag{14.7}$$

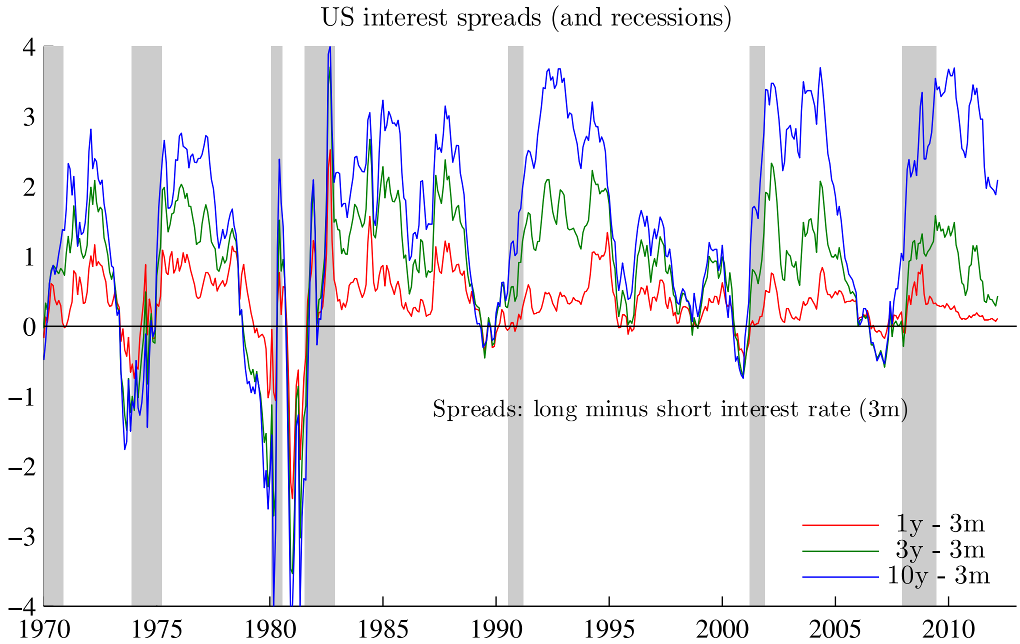

Figure 14.7: US term spreads (over a 3m T-bill)

The general expression for a maturity of m is

$$\begin{aligned} y_t(m) &= a(m) + b(m)r_t, \text{ where} \\ a(m) &= \lambda(m) + \mu\,[1 - b(m)] \text{ where} \\ b(m) &= (1 + \rho + \ldots + \rho^{m-1})/m = (1 - \rho^m)/[(1 - \rho)m]. \end{aligned} \tag{14.8}$$

In this model, all movements of the yield curve are driven by the state variable (the short rate), so it is a *one-factor model*. However, the shifts of the yield are parallel if $\rho = 1$ (the

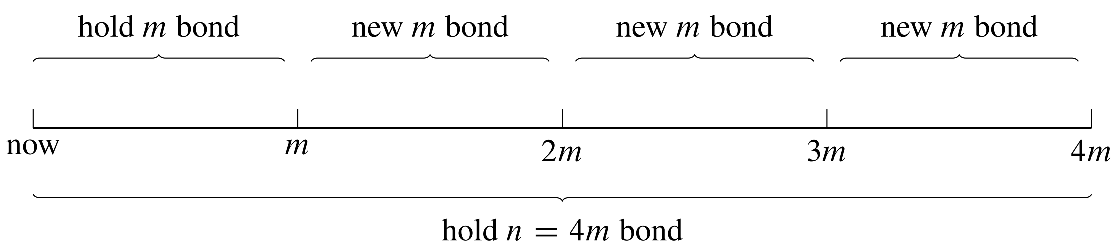

Figure 14.8: Timing for expectations hypothesis

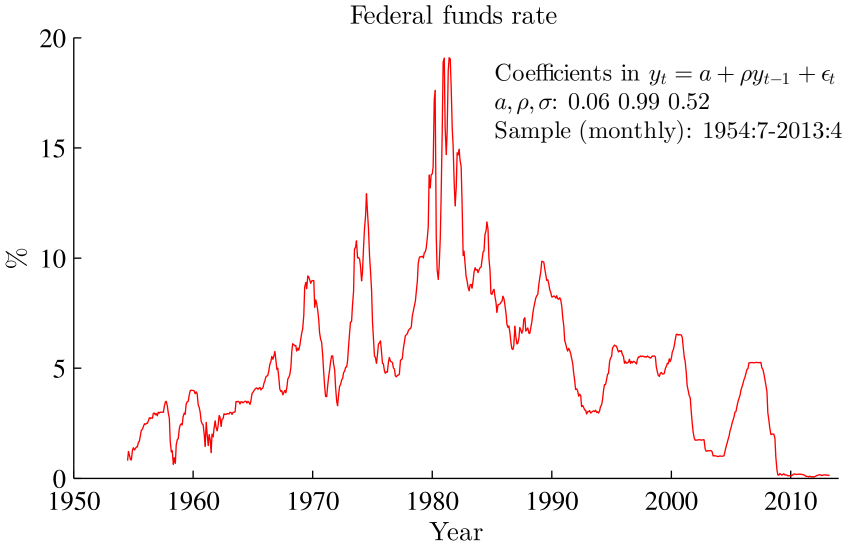

Figure 14.9: Federal funds rate, monthly data

random walk model) since then $b(m) = 1$ in (14.8), so we get

$$y_t(m) = \lambda(m) + r_t, \text{ if } \rho = 1. \tag{14.9}$$

See Figures 14.11–14.12 for an illustration. In particular, notice that the main driver of the long rates is the average expected future short rates. Since the model has a mean-reverting short rate (unless $\rho = 1$), expected future short rates (and therefore the current long rates) are always closer the the mean than the current short rate.

Example 14.5 *(Vasicek model) For $\rho = 0.9$ and $\mu = 0.05$, (14.8) gives (assuming no risk premia)*

$$\begin{bmatrix} y_t(1) \\ y_t(2) \\ y_t(3) \\ y_t(4) \end{bmatrix} \approx \begin{bmatrix} 0 \\ 0.0025 \\ 0.0048 \\ 0.007 \end{bmatrix} + \begin{bmatrix} 1 \\ 1.9/2 \\ 2.71/3 \\ 3.44/4 \end{bmatrix} r_t.$$

14.1.5 The Vasicek Model: Hedging a Bond

The Vasicek model allows us to calculate (or rather estimate) the proper way of *hedging a bond*. In particular, it can be used to calculate the yield changes in a duration hedge.

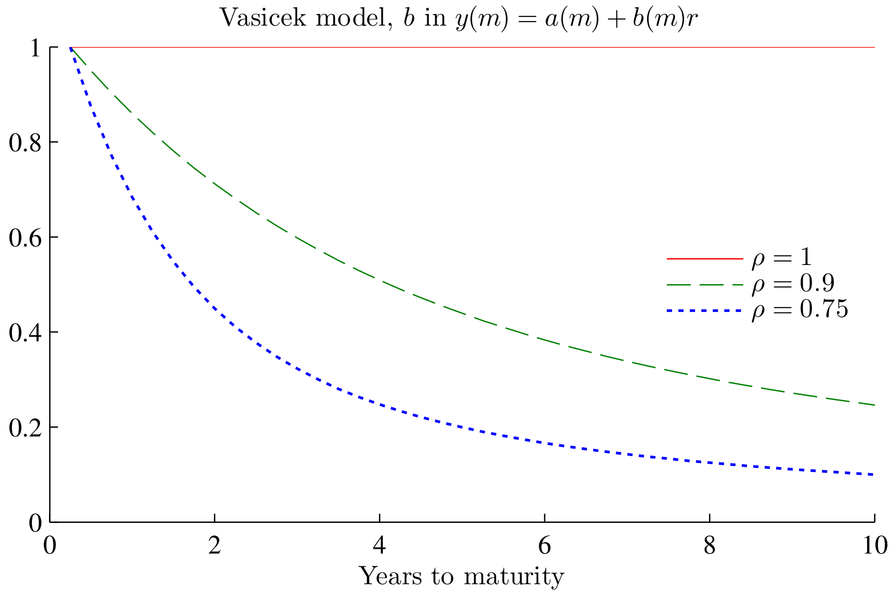

Figure 14.10: Slope in Vasicek model

Recall that the value (dollars) invested into a hedge bond is (H) relative to the value of the liability (L) is,

$$hB_H/B_L = \frac{D_L^M}{D_H^M} \times \frac{\Delta\theta_L/(1+\theta_L)}{\Delta\theta_H/(1+\theta_H)} \tag{14.10}$$

where D_i^M is Macaulay's duration, θ_i the yield to maturity and B_i the price of bond i.

To illustrate that, consider zero coupon bonds (it's the simplest example). In this case, Macaulay's duration equals the duration (m) and the ytm equals the spot rate. Using the results from the Vasicek model (14.8), we get

$$\Delta y_t(m) = b(m)\Delta r_t. \tag{14.11}$$

For the bonds H and L we have

$$\Delta\theta_H = b(H)\Delta r_t \tag{14.12}$$

$$\Delta\theta_L = b(L)\Delta r_t. \tag{14.13}$$

We also know that H and L are the Macaulay's durations (since the bonds here are zero

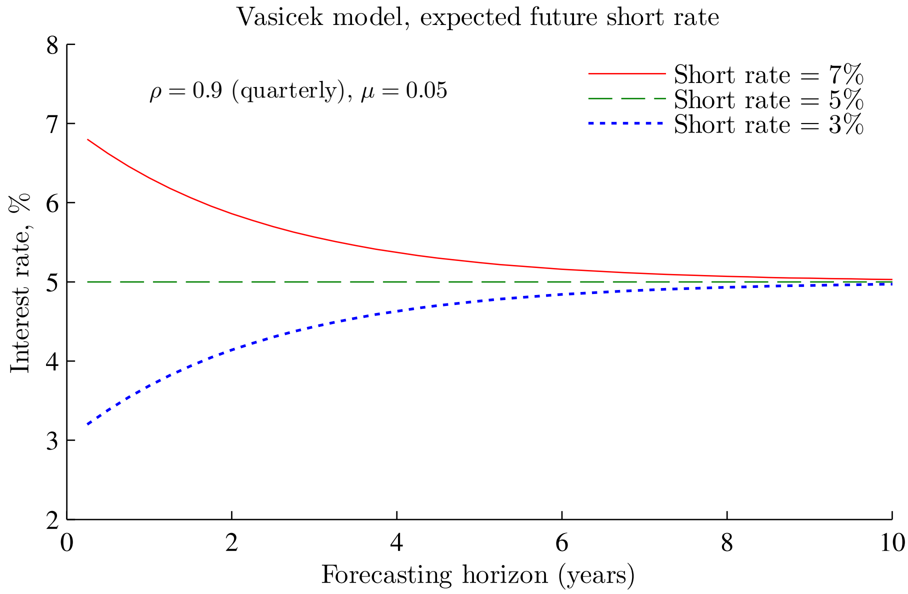

Figure 14.11: Expected future short rate in Vasicek model, for different intial short rates

coupon bonds), so by combing with (14.10) we have

$$\begin{aligned} hB_H/B_L &= \frac{L}{H} \times \frac{b(L)\Delta r_t}{b(H)\Delta r_t} \times \frac{1+\theta_H}{1+\theta_L} \\ &= \frac{L \times b(L)}{H \times b(H)} \times \frac{1+\theta_H}{1+\theta_L}. \end{aligned} \tag{14.14}$$

Also the the (θ_H, θ_L) terms could be rewritten by using (14.8), but that would create a long expression.

The key point is that the proper hedging depends on the first ratio in (14.14), $Lb(L)/[Hb(H)]$, which measures the sensitivity of the two bond prices to changes on the short rate. This ratio includes both the relative durations (here L/H) and how the yields are affected by the change of the state variable (Δr). To understand the former, recall that the bond price and the continuously compounded interest rate are related according to $1/B(m) = \exp[my(m)]$. Therefore

$$\Delta \ln B(m) = -m\Delta y(m). \tag{14.15}$$

For instance, if the interest rate on a 10-year bond increases with 0.01, then the bond price decreases by 0.1. In contrast, the relative response to changes in the state variable

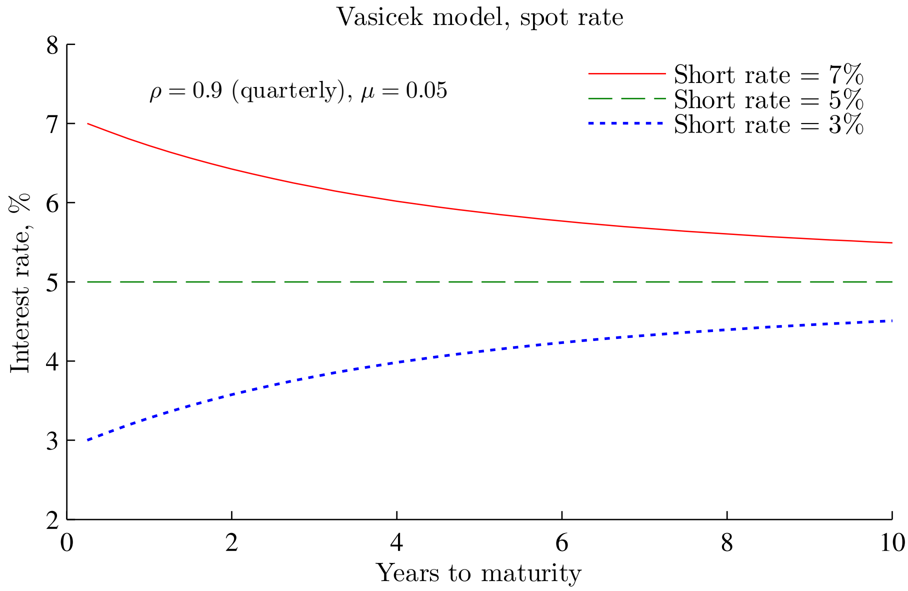

Figure 14.12: Vasicek model, spot rates for different intial short rates

$(b(L)/b(H))$ are related to the persistence of the state variable. If the state variable is a random walk, then this ratio is 1, but if the state variable is strongly mean reverting (ρ a lot lower than 1), then this ratio depends of which bond has the longer maturity.

If the liability has a longer maturity than the hedging bond ($L > H$), then the relative value invested into the hedge bond (14.14) is affected by two factors: *(a)* $L/H > 1$ so we invest more into the hedge bond because the price of a long bond (L) reacts more to a yield change; *(b)* but on the other hand $b(L)/b(H) \leq 1$ since long interest rates typically reacts less to changes in r. Only with a random walk model ($\rho = 1$) when all yield curve movements are parallel shifts do we have $b(L)/b(H) = 1$, otherwise the ratio is below 1.

For bonds with coupons (or other portfolios of zero coupon bonds), the ratio in (14.14) typically has to been calculated numerically, but that is straightforward.

Example 14.6 *(Hedging in a Vasicek model) If we want to hedge the 4-period bond by using the 2-period bond in Example 14.5, then (14.14) is*

$$hB_H/B_L = \frac{3.44}{1.9} \times \frac{1+\theta_H}{1+\theta_L} \approx 1.81 \times \frac{1+\theta_H}{1+\theta_L}.$$

If the yields are similar, then we get a a value around 1.81. Comparing with just using the ratio of durations ($4/2 = 2$), the Vasicek model suggests a somewhat smaller position in the (shorter) hedge bond. The reason is that the interest rate on the hedge bond reacts more (to changes in the short rate r) than the interest rate on the liability. In contrast, if there was no mean reversion ($\rho = 1$), then $hB_H/B_L = 2$.

All one-factor models (not least the Vasicek model) imply that all yields are perfectly correlated (there is a common single driving force) and only fairly limited yield curve movements are possible. *Multi-factor models* overcome most of those limitations. For instance, the model in Nelson and Siegel (1987) is a two-factor model.

14.2 Interest Rates and Macroeconomics

This section outlines three (not mutually exclusive) macroeconomic approaches to modelling the yield curve.

14.2.1 The Fisher Equation and Index-Linked Bonds

Let π_{t+n} be the one period inflation rate over t to $t+n$, that is,

$$\pi_{t+n} = \ln(P_{t+n}/P_t). \tag{14.16}$$

Also, let $y_t^r(n)$ be the one period continuously compounded real interest rate (an interest rate measured in goods).

The Fisher equation (here in the form of continuously compounded rates) says that the nominal interest rate includes compensation both for inflation expectations, $\mathrm{E}_t\, \pi_{t+n}$,the real interest rate, $y_t^r(n)$, and possibly a constant (across time) risk premium, $\psi(n)$,

$$y_t(n) = \mathrm{E}_t\, \pi_{t+n} + y_t^r(n) + \psi(n). \tag{14.17}$$

Example 14.7 *(Fisher equation) Suppose the nominal interest rate is $y(n) = 0.07$, the real interest rate is $y^r(n) = 0.03$, and the nominal bond has no risk premium ($\psi = 0$), then the expected inflation is $\mathrm{E}_t\, \pi_{t+n} = 0.04$.*

The same type of relation holds for forward rates. The Fisher equation suggests a framework for analysing nominal interest rates in terms of real interest rates and inflation

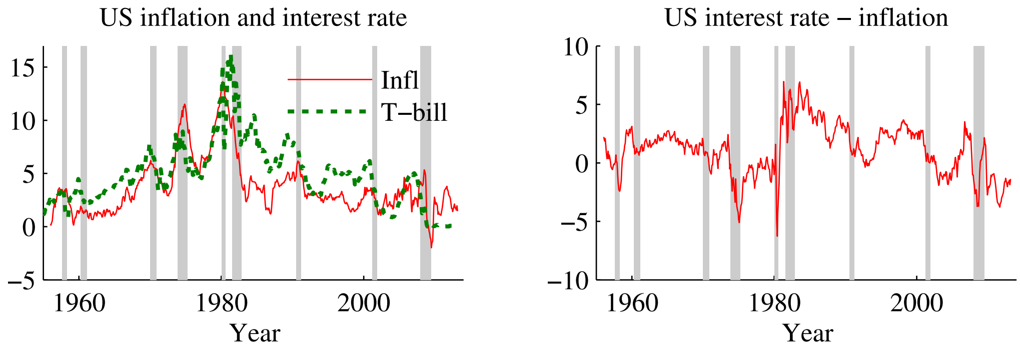

Figure 14.13: US inflation and 3-month interest rate

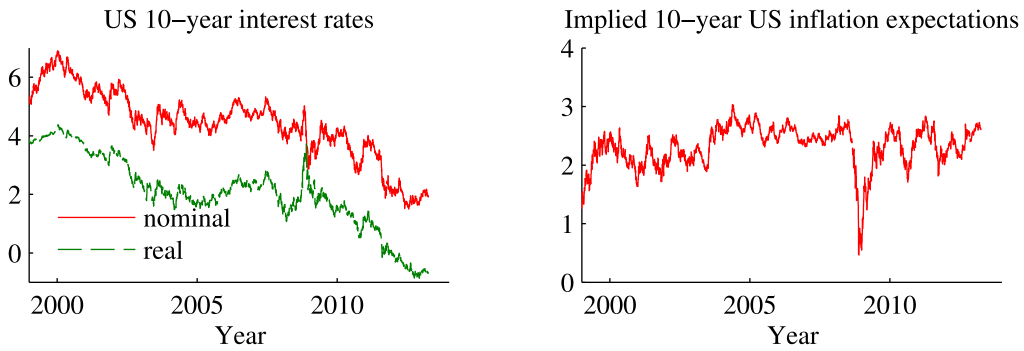

Figure 14.14: US nominal and real interest rates

expectations. This is commonly used for long rates. Information about real interest rates can be elicited from *index-linked bonds*, that is, bonds which give automatic compensation for actual inflation.

Empirical results typically indicate that there are non-trivial movements in the real interest rate and/or risk premia—especially for short forecasting horizons. This holds also when inflation expectations, as measured by surveys, are used as the dependent variable. Inflation expectations seems to vary by less than the interest rate. It is therefore not straightforward to extract inflation expectations from nominal interest rates.

The Fisher equation could also be embedded in a macro model to construct a sophisticated (and complicated) model of the yield curve. This involves using macro theory/empirics to model how real interest rates and inflation expectations (for different ma-

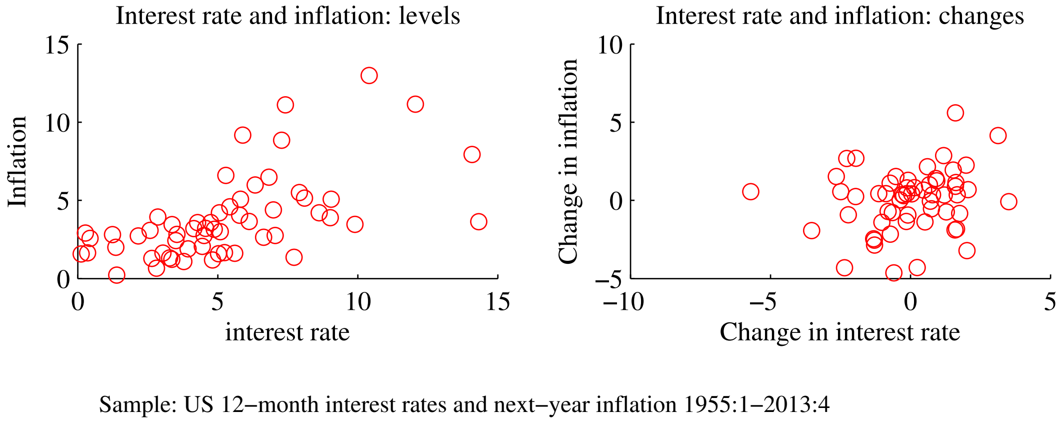

Figure 14.15: US nominal interest rates and subsequent inflation

turities) depend on the state of the economy.

14.2.2 The Expectations Hypothesis of Interest Rates

The expectations hypothesis of interest rates says that long interest rates equal an average of expected future short rates, possibly with a constant (across time, not maturities) risk premium as in (14.3). Alternatively, that forward rates equal expected future spot rates.

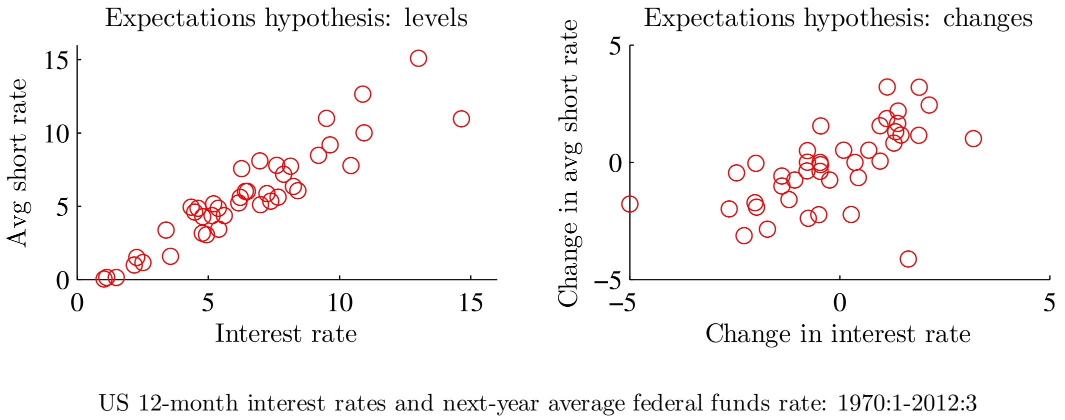

Figure 14.16: US 12-month interest and average federal funds rate (next 12 months)

The expectations hypothesis is often used to calculate implied "forecasts" of future short interest rates. For instance, suppose the central bank increases its policy rate (typically a very short rate, at most a week or two). This is likely to affect also longer interest rates, but how is another matter. Let us consider a few different cases. For simplicity we assume that risk premia are unaffected by this move in the policy rate.

First, one possibility is that only the very short interest rates change, and that all longer interest rates stay unchanged. This would happen if the policy move was well anticipated.

Second, another possibility is that most long interest rates increase. Under the expectations hypothesis of interest rates the interpretation is that the market now expects high short interest rates also in the future. That is, that the central bank will not reverse its policy action in the foreseeable future. If we are willing to assume that the real interest rate was not affected by the policy move, then one possible interpretation is that the central bank has received information about a long-lasting inflation pressure.

Third, and finally, short rates may increase, but really long interest rates decrease. A common interpretation of this scenario is that the central bank has become more inflation averse. It therefore raises the policy rate to bring down inflation. If the market believes that it will succeed, then it follows that it will eventually be possible to lower interest rates (when inflation and inflation expectations are lower).

The expectations hypothesis has been tested many times, typically by an ex post linear regression (realized interest rates regressed on lagged forward rates). The results often reject the expectations hypothesis, but the results depend on how the test is done. It is not clear, however, if the rejection is due to systematic risk premia or to fairly small samples (compared to the long swings in interest rates). The expectations hypothesis gets more support when survey data on interest rate expectations is used instead on realized interest rates.

14.2.3 Uncovered Interest Rate Parity

Uncovered interest rate parity says that the difference between a domestic and foreign interest rate equals the expected depreciation plus a constant (across time, not maturities) risk premium

$$y_t(m) - y_t^*(m) = \mathrm{E}_t\, s_{t+m} - s_t + \varphi_m, \tag{14.18}$$

where $y_t^*(m)$ is a (continuously compounded) foreign interest rate, and s_t is the logarithm of the exchange rate (number of domestic currency units per foreign currency unit). If this condition hold with a zero risk premium, then the expected return from investing in foreign bonds and then buy domestic currency equals the known return from investing in domestic bonds.

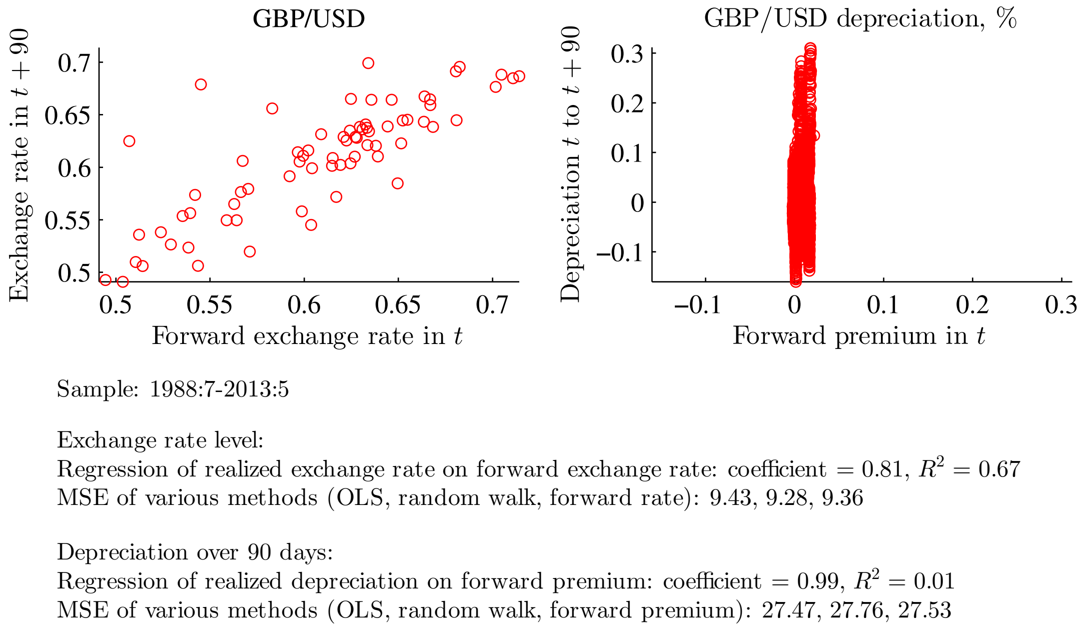

Figure 14.17: GBP/USD spot and forward exchange rates

Empirical evidence suggests that there might be large movements in the risk premia over time (or that there have been systematic surprises in historical samples).

14.2.4 A New-Keynesian Model of Monetary Policy

Monetary policy is a crucial part of the macroeconomic picture these days, so it is important to understand how monetary policy is formed. It has not always been this way: there are long periods when many countries adopted a very simple (or so it seemed) monetary policy by pegging the currency to another currency. Macroeconomic policy was then synonymous with fiscal policy. Recently, the roles have changed.

Modern macro models are often smaller than the older macroeconometric models and they pay more attention to both the supply side of the economy and the role of expectations. These models try to capture the key elements in the way central banks (and most other observers) reason about the interaction between inflation, output, and monetary policy.

In these models, inflation depends on expected future inflation (some prices are set today for a long period and will therefore be affected about expectations about future costs and competitors' prices), lagged inflation, and a "Phillips effect" where an *output gap* (output less trend output) affects price setting via demand pressure. For instance, inflation (π_t) is often modelled as

$$\pi_t = \alpha \, \mathrm{E}_t \pi_{t+1} + \beta \pi_{t-1} + \phi x_t + \varepsilon_{\pi t}, \tag{14.19}$$

where x_t is the output gap and $\varepsilon_{\pi t}$ can be interpreted as "cost push" shocks (wage demands, oil price shocks). This equation can be said to represent the supply side of the economy and it is typically derived from a model where firms with some market power want to equate marginal revenues and marginal costs, but choose to change prices only gradually.

The demand side of the economy is modelled from consumers' savings decision, where the trade off between consumption today and tomorrow depends on the real interest rates. Simplifying by setting consumption equal to output we get something like the following equation for the output gap

$$x_t = x_{t-1} - \gamma(i_t - \mathrm{E}_t \pi_{t+1}) + u_t, \tag{14.20}$$

where i_t is the nominal interest rate (set by the central bank) and u_t is a shock to demand. Note that the expected *real* interest rate affects demand (negatively).

In some cases, the real exchange rate is added to both (14.19) and (14.20), capturing price increases on imported goods and foreign demand for exports, respectively. The exchange rate is then linked to the rest of the model via an assumption of uncovered interest rate parity (that is, expected exchange rate depreciation equals the interest rate differential).

Some of the important features of this simple model are: *(i)* inflation expectations matter for today's inflation (think about wage inflation), *(ii)* the instrument for monetary policy, the short interest rate i_t, can ultimately affect inflation only via the output gap; *(iii)* it is the real, not the nominal, interest rate that matters for demand.

To make the model operational, two more things must be added: the monetary policy (the way the interest rate is set) and the expectations in (14.19)–(14.20) must be specified.

It is common to assume that the central bank has some instrument rule like the famous

"Taylor rule"

$$i_t = \theta_0 + 0.5x_t + 1.5\pi_t + v_t. \tag{14.21}$$

The residual v_t is a "monetary policy shock," which picks up factors left out of the model (for instance, the central bank's concern for the banking sector or simply changes in the central bank's objectives). This simple reaction function has been able to track US monetary policy fairly well over the last decade or so. Another approach to find a policy rule is to assume that the central bank has some loss function that it minimizes by choosing an policy rule. This loss function is often a weighted average of the variance of inflation and the variance of the output gap. The policy rule is the solution of the minimization problem, and can often look more complicated than the Taylor rule. However, there is one interesting special case. Suppose the central bank wants to minimize the (unconditional) variance of inflation. The formal optimization problem is then

$$\min_{i_t} \text{Var}(\pi_t), \text{ subject to (14.19) and (14.20).} \tag{14.22}$$

The solution is then that the interest rate should be set so that actual inflation is zero (here the mean) in every period. If the model is changed so there is a time lag between the interest rate decision and its effect on inflation (for instance, by letting inflation in (14.19) react to x_{t-1} instead of x_t), then the interest rate should be set so that the conditional expectation of next period's inflation is zero (the mean), $\mathrm{E}_t\, \pi_{t+1} = 0$. This type of "rule" is used in much of the monetary policy debate.

The expectations in (14.19)–(14.20) can be handled in many ways. The perhaps most straightforward way is to assume that the expectations about the future equal the current value of the same variable (a "random walk"). A more satisfactory way is to use survey data on inflation expectations. Finally, many model builders assume that expectations are "rational" (or "model consistent") in the sense that the expectation equals the best guess we could do under the assumption that the model is correct. This latter approach typically requires a sophisticated way of solving the model (as the model both generates the best guesses and depends on them).

14.3 Forecasting Interest Rates

14.3.1 Forecasting Monetary Policy or Inflation?

There is a two-way causality: inflation and the real economy (which depend on the real interest rate) affect monetary policy, and monetary policy can surely affect inflation and the real economy. This makes it difficult to analyse and forecast interest rates. However, for short term forecasting, the emphasis is typically on forecasting the next monetary policy move. Long run forecasting relies more on understanding the determinants of real interest rates and inflation, which depends on the general business cycle prospects, but also on the long run stance of monetary policy ("tough on inflation or not?").

14.3.2 Interest Rate Forecasts by Analysts

Kolb and Stekler (1996) use a semi-annual survey of (12 to 40) professional analysts' interest rate forecasts published in Wall Street Journal. The (6 months ahead) forecasts are for the 6-month T-bill rate and the yield on 30-year government bonds. The paper studies four questions, and I summarize the findings below.

1. Q. Is the distribution of the forecasts (across forecasters) at any point in time symmetric? (Analyzed by first testing if the sample distribution could be drawn from a normal distribution; if not, then checking asymmetry (skewness).) A. Yes, in most periods. (The authors argues why this makes the median forecast a meaningful representation of a "consensus forecast.")

2. Q. Are all forecasters equally good (in terms of ranking of (absolute?) forecast error)? A. Yes for the 90-day T-bill rate; No for the long bond yield.

3. Q. Are some forecasters systematically better (in terms of absolute forecast error)? (Analyzed by checking if the absolute forecast error is below the median more than 50% of the time) A. Yes.

4. Q. Do the forecasts predict the direction of change of the interest rate? (Analyzed by checking if the forecast gets the sign of the change right more than 50% of the time.) A. No.

14.3.3 Market Positions as Interest Rate Forecasts

Hartzmark (1991) has data on daily futures positions of large traders on eight different markets, including futures on 90-day T-bills and on government bonds. He uses this data to see if the traders changed their position in the right direction compared to realized prices (in the future) and if they did so consistently over time.

The results indicate that these large investors in T-bills and bond futures did no better than an uninformed guess of the direction of change of the bill and bond prices. He gets essentially the same results if the size of the change in the position and in the price are also taken into account.

There is of course a distribution of how well the different investors do, but it looks much like one generated from random guesses (uninformed forecasts). The investors change places in this distribution over time: there is very little evidence that successful investors continue to be successful over long periods.

14.3.4 Long Rates as Forecasts of Future Short Rates: The Expectations Hypothesis

The expectations hypothesis has been tested many times, typically by an ex post linear regression (realized interest rates regressed on lagged forward rates). The results typically reject the expectations hypothesis. It is not clear, however, if the rejection is due to systematic risk premia or to fairly small samples (compared to the long swings in interest rates). The expectations hypothesis gets more support when survey data on interest rate expectations is used instead on realized interest rates.

14.4 Risk Premia on Fixed Income Markets

There are many different types of risk premia on fixed income markets.

Nominal bonds are risky in real terms, and are therefore likely to carry *inflation risk premia*. Long bonds are risky because their market values fluctuate over time, so they probably have *term premia*. Corporate bonds and some government bonds (in particular, from developing countries) have *default risk premia*, depending on the risk for default. Interbank rates may be higher than T-bill of the same maturity for the same reason (see the TED spread, the spread between 3-month Libor and T-bill rates) and illiquid bonds may carry *liquidity premia* (see the spread between off-the run and on-the-run bonds).

Figures 14.18–14.21 provide some examples.

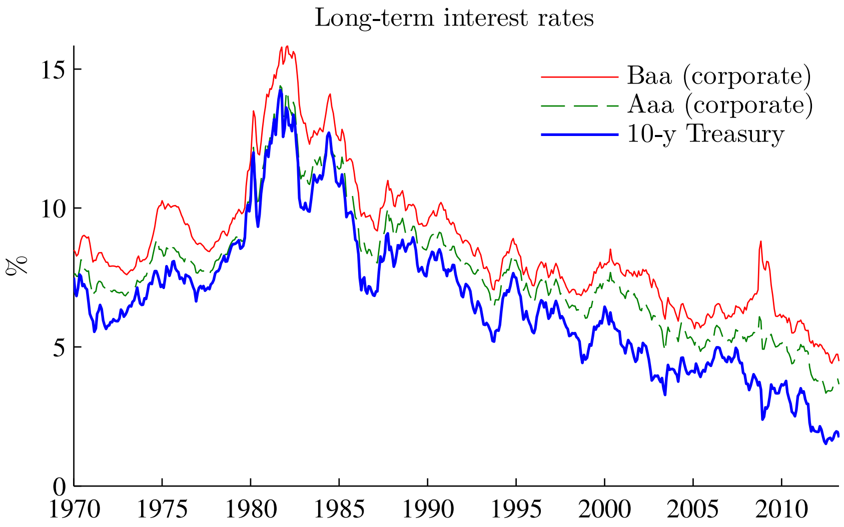

Figure 14.18: US interest rates

Bibliography

Cochrane, J. H., 2001, *Asset pricing*, Princeton University Press, Princeton, New Jersey.

Diebold, F. X., and C. Li, 2006, "Forecasting the term structure of government yields," *Journal of Econometrics*, 130, 337–364.

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

Hartzmark, M. L., 1991, "Luck versus forecast ability: determinants of trader performance in futures markets," *Journal of Business*, 64, 49–74.

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

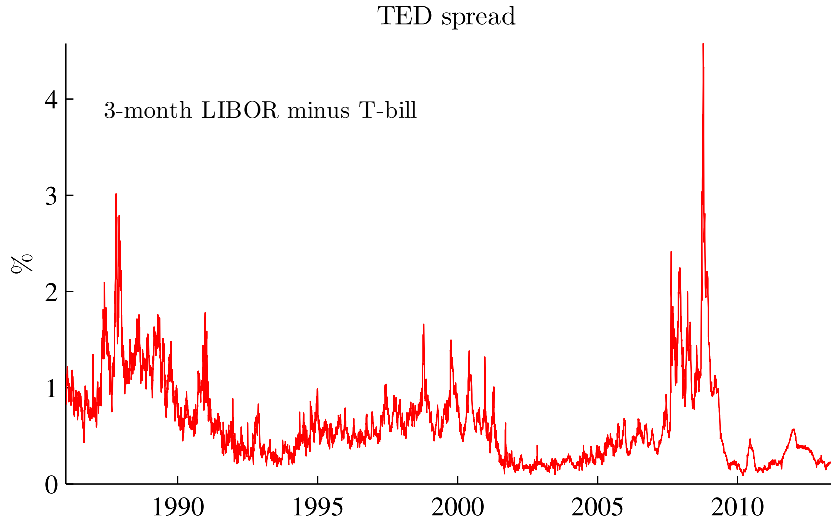

Figure 14.19: TED spread

Kolb, R. A., and H. O. Stekler, 1996, “How well do analysts forecast interest rates,” *Journal of Forecasting*, 15, 385–394.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

Nelson, C., and A. Siegel, 1987, “Parsimonious modeling of yield curves,” *Journal of Business*, 60, 473–489.

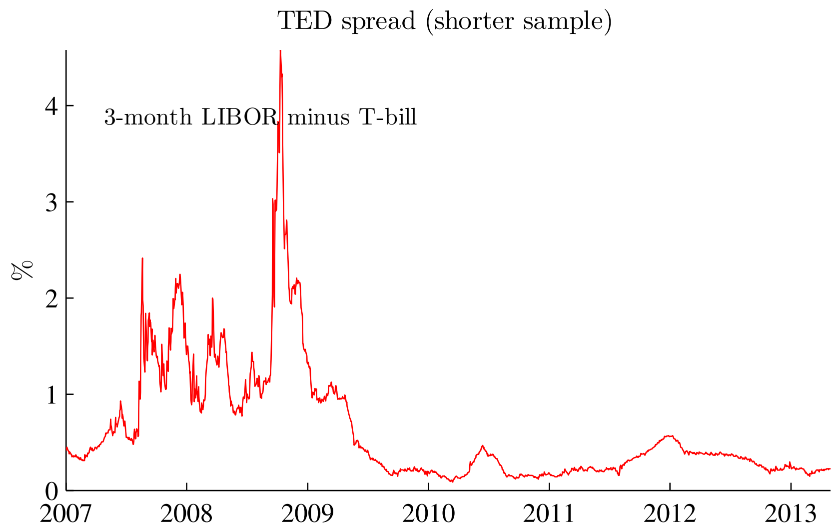

Figure 14.20: TED spread recently

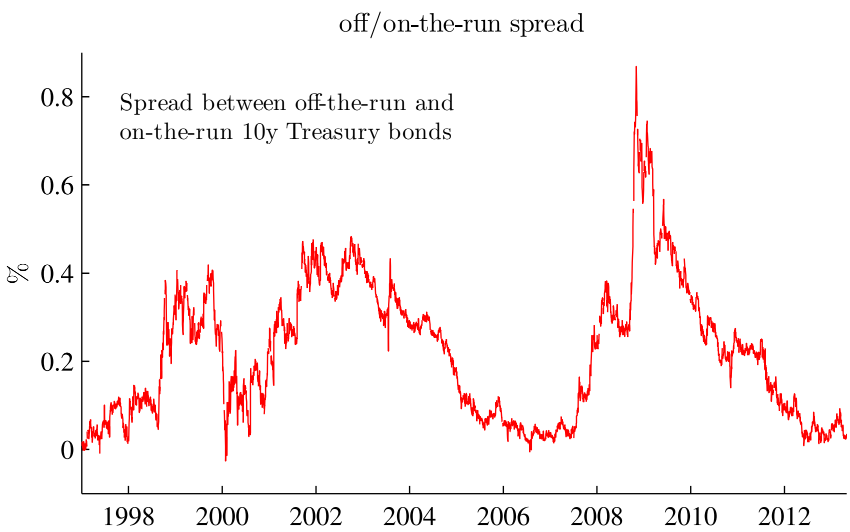

Figure 14.21: Off-the-run liquidity premium

15 Basic Properties of Futures and Options

Main References: Elton, Gruber, Brown, and Goetzmann (2010) 23–24 and Hull (2009) 5 and 8–9
Additional references: McDonald (2006) 9–12

15.1 Derivatives

Derivatives are assets whose payoff depend on some underlying asset (for instance, the stock of a company). The most common derivatives are futures contracts (or similarly, forward contracts) and options. Sometimes, options depend no directly on the underlying, but on the price of a futures contract on the underlying. See Figure 15.1.

Derivatives are in zero net supply, so a contract must be issued (a short position) by someone for an investor to be able to buy it (long position). For that reason, gains and losses on derivatives markets sum to zero.

15.2 Forward and Futures

15.2.1 Present Value

Forward prices play an important role in simplifying option analysis, so we first discuss the forward-spot parity.

The present value of one unit paid m periods into the future must be the price of a bond, $B(m)$, maturing at the same time. We therefore have that the present value of Z is

$$\begin{aligned}
\text{Present value}_m(Z) &= B_t(m)Z, \text{ or} && (15.1)\\
&= [1 + Y_t(m)]^{-m} Z, \text{ or} && (15.2)\\
&= e^{-my_t(m)}Z, && (15.3)
\end{aligned}$$

where $Y_t(m)$ is effective spot interest rate, and $y_t(m)$ is the continuously compounded interest rate ($y_t\ (m) = \ln[1 + Y_t(m)]$).

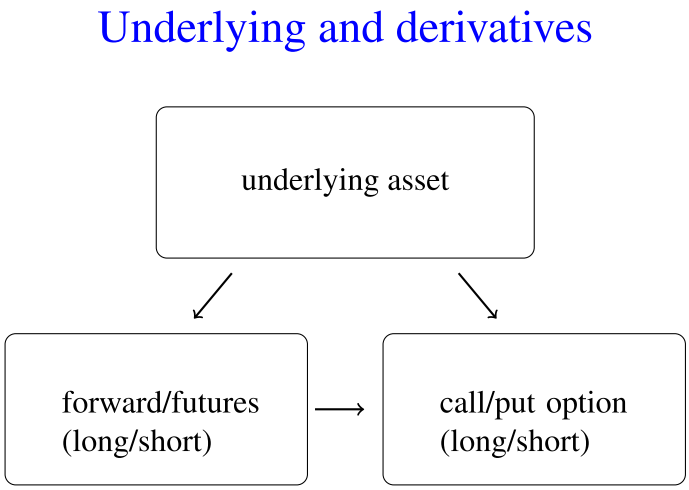

Figure 15.1: Derivatives on an underlying asset

Example 15.1 *(Present value) With* $y_t(m) = 0.05$ *and* $m = 3/4$ *we have the present value* $e^{-0.05\times 3/4}Z \approx 0.963Z$.

15.2.2 Definition of a Forward Contract

A forward contract specifies (among other things) which asset that should be delivered at the expiration and what the price is then (the forward price). See Figure 15.2 for an illustration. The forward (and also a futures, see below) are zero sum games: the profit of the buyer is the loss of the seller (or vice versa).

The profit (payoff) of a forward contract at expiration is very straightforward. Let S_{t+m} be the price (on the spot market) of the underlying asset at expiration (in $t+m$). Then, for the *buyer* of a forward contract

$$\text{payoff of a forward contract} = S_{t+m} - F. \tag{15.4}$$

The owener of the forward contract pays F to get the asset, sells it immediately on spot market for S_{t+m}. See Figure 15.3. Similarly, the payoff for the *seller* of a forward contract is $F - S_{t+m}$ (she buys the asset on spot market for S_{t+m}, gets F for asset according to the contract). This sums to *zero*.

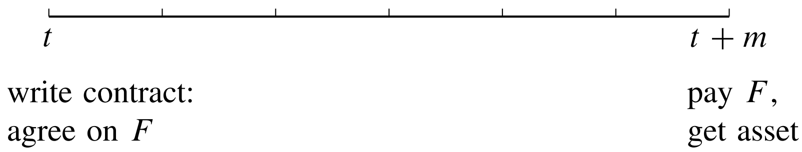

Figure 15.2: Timing convention of forward contract

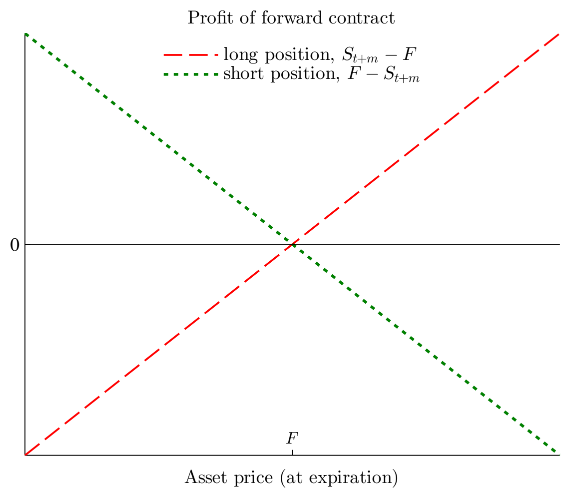

Figure 15.3: Profit (payoff) of forward contract (at expiration)

15.2.3 Forward-Spot Parity

Proposition 15.2 *(Forward-spot parity, no dividends) The present value of the forward price, $F_t(m)$, contracted in t (but to be paid in $t+m$) on an asset without dividends equals the spot price:*

$$e^{-my_t(m)} F_t(m) = S_t, \text{ so} \tag{15.5}$$

$$F_t(m) = e^{my_t(m)} S_t. \tag{15.6}$$

The intuition is that the forward contract is like buying the underlying asset on credit—$e^{-my_t(m)} F_t(m)$ can be thought of as a prepaid forward contract.

Proof. (of Proposition 15.2) Portfolio A: enter a forward contract, with a present value

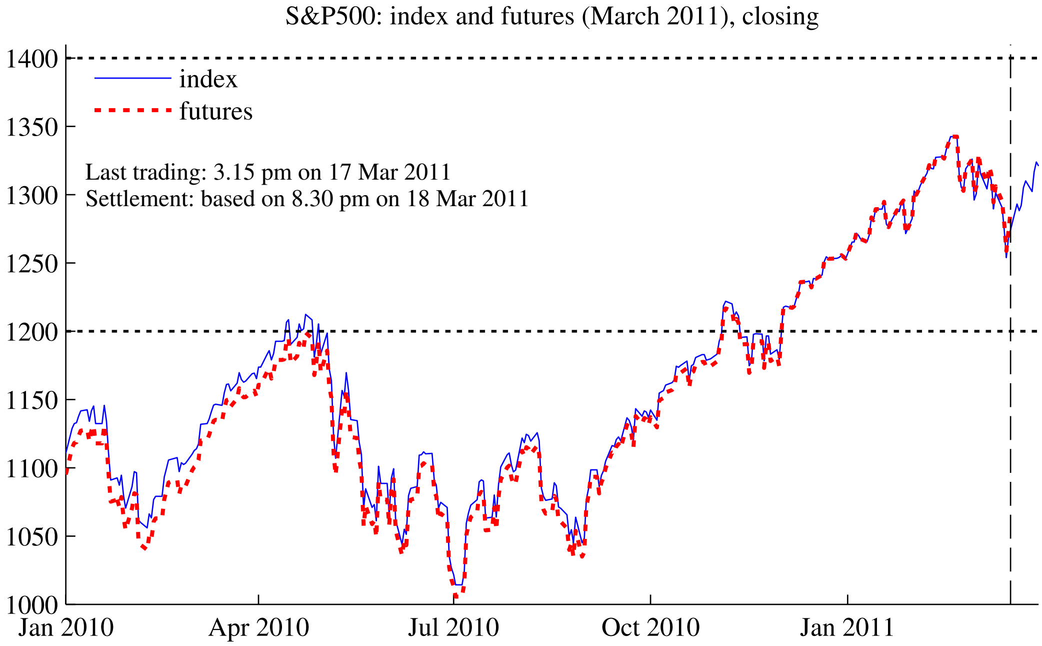

Figure 15.4: S&P500 index level and futures

of $e^{-my}F$. Portfolio B: buy one unit of the asset at the price S. Both portfolios give one asset at expiration, so they must have the same costs today. ■

Proposition 15.3 *(Forward-spot parity, discrete dividends) Suppose the underlying asset pays the dividend D_i at m_i $(i = 1, ..., n)$ periods into the future (but before the expiration date of the forward contract). The dividends must be known already in t. The forward price then satisfies*

$$e^{-my_t(m)}F_t(m) = S_t - \sum\nolimits_{i=1}^{n} e^{-m_i y_t(m_i)} D_i, \text{ so} \tag{15.7}$$

$$F_t(m) = e^{my_t(m)}S_t - e^{my_t(m)} \sum\nolimits_{i=1}^{n} e^{-m_i y_t(m_i)} D_i \tag{15.8}$$

The last term of (15.7) is the sum of the present values of the dividend payments. The intuition is that the forward contract does not give the right to these dividends so its value is the underlying asset value stripped of the present value of the dividends.

Proof. (of Proposition 15.3) Portfolio A: enter a forward contract, with a present value of $e^{-my}F$. Portfolio B: buy one unit of the asset at the price S and sell the rights to the known dividends at the present value of the dividends. Both portfolios give one asset at expiration, so they must have the same costs today. ■

Proposition 15.4 *(Forward-spot parity, continuous dividends) When the dividend is paid continuously as the rate δ (of the price of the underlying asset), then*

$$e^{-my_t(m)}F_t(m) = S_t e^{-\delta m}, \text{ so} \tag{15.9}$$

$$F_t(m) = S_t e^{m[y_t(m)-\delta]} \tag{15.10}$$

Proof. (of Proposition 15.4) Portfolio A: enter a forward contract, with a present value of $e^{-my}F$. Portfolio B: buy $e^{-\delta m}$ units of the asset at the price $e^{-\delta m}S$, and then collect dividends and reinvest them in the asset. Both portfolios give one asset at expiration, so they must have the same costs today. ■

Remark 15.5 *(Forward-spot parity, currencies) Investing in foreign currency effectively means investing in a foreign interest bearing instrument which earns the continuous interest rate ("dividend") $y_t^*(m)$. Use $\delta = y_t^*(m)$ in (15.9).*

Remark 15.6 *Figure 15.4 provides an example of how the futures price (on SP 500), the intrinsic value of the option and the option price developed over a year. Notice how the futures prices converges to the index level at expiration of the futures. Before it can deviate because of delayed payment (+) and no part in dividend payments (−).*

15.2.4 Value of an Old Forward Contract

Consider a forward contract that expires in $t + m$, but where the contract was written at some earlier point in time ($\tau < t$) and specified a forward price of F_τ. The value of this contract in t is

$$\text{Value of old forward contract} = e^{-ym}(F_t - F_\tau), \tag{15.11}$$

where F_t is today's forward price on the same underlying asset (and same expiration date). The intuition is that an owner of an old (τ) forward contract can short sell a new forward contract (t) and thereby cancel all risk—and stand to win $F_t - F_\tau$ at expiration. The present value of this is (15.11).

Proof. (15.11) An investor sells (issues) a forward contract in t. At expiration, this will give $F_t - S_{t+m}$, where S_{t+m} is the price of the underlying asset at expiration. If she buys an old forward contract for the price V_t, the payoff of that is $S_{t+m} - F_\tau$ at expiration. Hence, the total portfolio has the payoff $F_t - F_\tau$, which is riskfree. There is a arbitrage opportunity unless the price of the old contract is $V_t = e^{-ym}(F_t - F_\tau)$. ■

15.2.5 Forwards versus Futures

A forward contract is typically a private contract between two investors—and can therefore be tailor made. A futures contract is similar to a forward contract (write contract, get something later), but is typically traded on an exchange—and is therefore standardized (amount, maturity, settlement process). The settlement is either cash settlement or physical settlement. The latter does not work for synthetical assets like equity indices.

Another important difference is that a forward contract is settled at expiration, whereas a futures contract is settled daily (*marking-to-market*), which essentially means that gains and losses (because of prices changes) are transferred between issuer and owner daily—but kept at the at an interest bearing account at the exchange. The counter parties have to post and initial margin—and the marking-to-market then adds to/subtracts from this margin. If the amount decreases below a certain level (maintainance margin), then a margin call is issues to the inestor—informing him/her to add cash to the margin account. If interest rates change randomly over time (and they do), the rate at which the money on the margin account is invested at (refinanced) will therefore be different from the rate when the futures was issued. This risk of this happening is reflected in the futures price. The proposition below show that, if (hypothetically) the interest rate path was non-stochastic, then the forward and futures prices would be the same. In practice, the difference between forward and futures prices is typically small.

Proposition 15.7 *(Forward vs. futures prices, non-stochastic interest rates) The forward and futures prices would be the same if the interest rate only changed in a non-stochastic way.*

Proof. (of Proposition 15.7) To simplify the notation, let $t = 0$ and $m = 2$. Also, let r_s be the continuously compounded one-day interest rate and f_s be the futures price. Strategy A: have e^{r_0} long futures contracts on (the end of) day 0, increase it to $e^{r_0+r_1}$ on day 1. Provided we reinvest the settlements in one-day bills, we have

Day (s)	Position	Settlement	End-value of reinvested settlement
0	e^{r_0}	0	0
1	$e^{r_0+r_1}$	$e^{r_0}(f_1-f_0)$	$e^{r_0}(f_1-f_0)e^{r_1}$
2	0	$e^{r_0+r_1}(f_2-f_1)$	$e^{r_0+r_1}(f_2-f_1)$

The end-value of strategy A is therefore $e^{r_0+r_1}(f_2-f_0)$, which equals $e^{r_0+r_1}(S_2-f_0)$ since the value at expiration is the value of the underlying asset. Strategy B: be long $e^{r_0+r_1}$ forward contracts, which gives a payoff on day 2 of $e^{r_0+r_1}(S_2-F_0)$. Both strategies take on exactly the same risk, so the prices must be the same: $f_0=F_0$. (The proof relies on knowing r_1 already on day 0.) ■

Example 15.8 *(Margin account) Margin account of a buyer (holder) of a futures contract (maintenance margin = 0.75×initial margin) could be as follows (assuming a zero interest rate):*

Day	*Futures price*	*Daily gain*	*Posting of margin*	*Margin account*
0	100		4	4
1	99	−1		3
2	97	−2	2	3
3	99	2		5

Notice that the overall profit is the difference of what has been put into the margin account ($4+2$*) and the final balance (5), that is,* -1*. This is also the* cumulative *daily gain (*$-1-2+2=-1$*). With marking to market this is all that happens: no payment of the futures price and no delivery of the underlying asset. However, it is equivalent to what happen without marking to market, since at expiration, the gain is* $99-100=-1$ *(futures=underlying at expiration).*

15.3 Introduction to Options

15.3.1 Definition of European Calls and Puts

A European *call* option contract traded in t may stipulate that the buyer of the contract has the right to buy one unit of the underlying asset from the issuer of the option on the

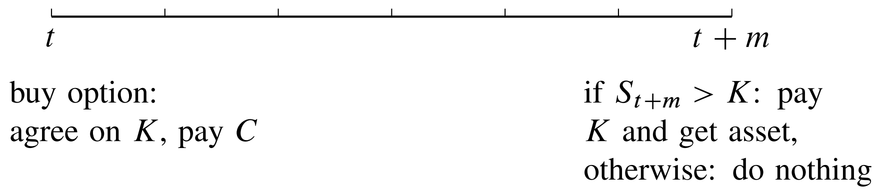

Figure 15.5: Timing convention of a European call option contract

expiration date $t + m$ at the strike price K. See Figure 15.5 for the timing convention.

The payoff at exercise is zero or, if larger, the price of the underlying asset, S_{t+m}, minus the strike price

$$\text{call payoff}_{t+m} = \max\left(0, S_{t+m} - K\right). \tag{15.12}$$

Clearly, an owner of a call option benefits from an increase in the price of the underlying asset (exercise the right to buy for K and sell asset at a higher price). The payoff of the original seller of the option (the option writer who has a short option position) is the mirror image of the buyer's payoff: the buyer's gain is the writer's loss: a *zero sum game*. This is true both for the payoff at exercise (15.12) as well as the profit

$$\text{call profit}_{t+m} = \text{call payoff}_{t+m} - \text{call price}_t. \tag{15.13}$$

See Figure 15.6 for an illustration.

A buyer of an option does not have to post any margin, but a seller does. The reason is that a default of the seller would create a loss for the option buyer if the option is worth exercising. In contrast, the default of the buyer cannot create a loss for the seller.

A *put* option instead gives the buyer of the contract the right to sell one unit of the underlying asset. The put price is here denoted by P. An owner of a put option benefits from a decrease in the price of the underlying asset (buy the asset cheaply and exercise the right to sell for K). The payoff is

$$\text{put payoff}_{t+m} = \max\left(0, K - S_{t+m}\right). \tag{15.14}$$

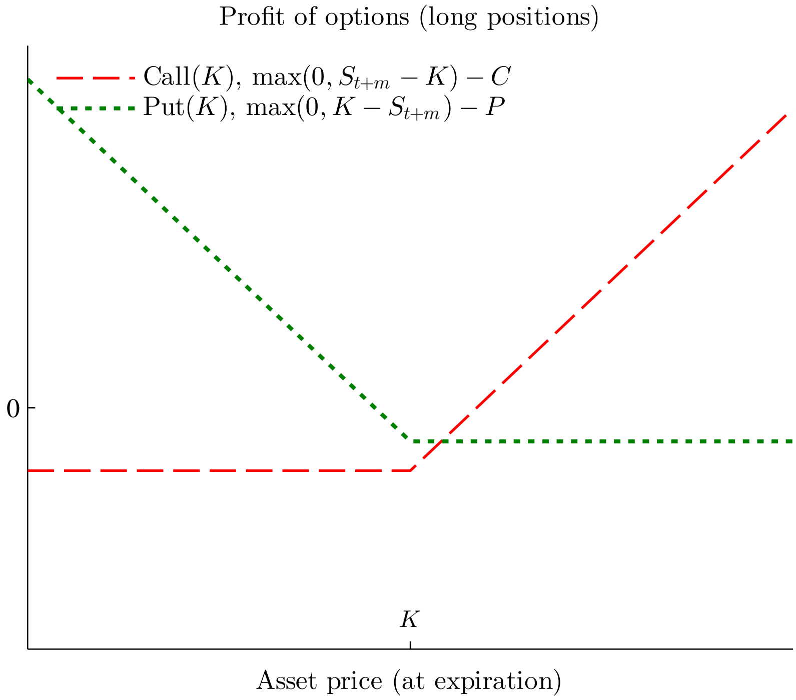

Figure 15.6: Profit of options, long positions

An option that would be profitable to exercise now is called *in-the-money*; an option that would be unprofitable to exercise is called *out-of-the-money*—and an option that would just break even is called at-the-money.

Figures 15.8–15.10 illustrates the trade intensity of options with different strike prices (but same expiration and underlying asset). Most of the trade is close to expiration, and there is a seasonality pattern related to rolling over the investment from other (expired) options.

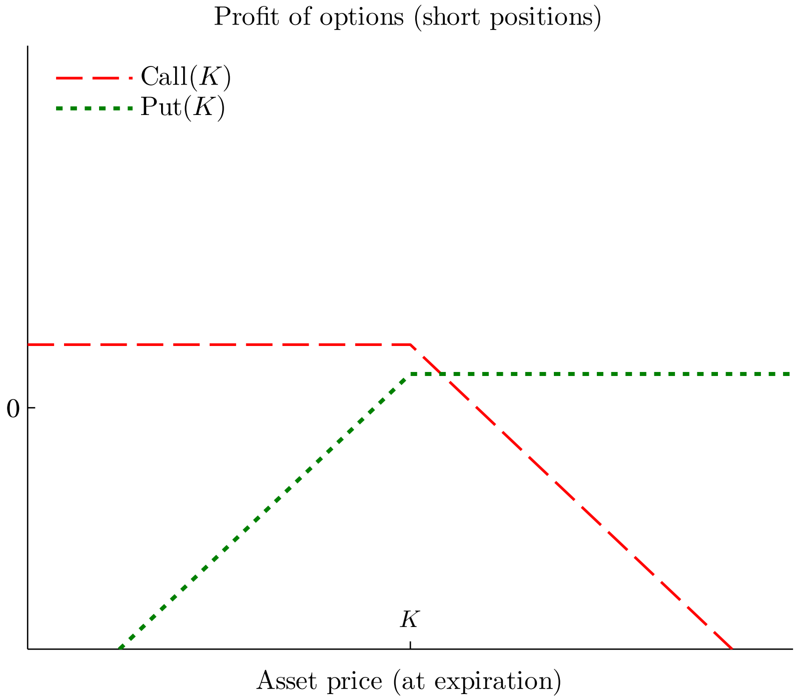

Figure 15.7: Profit of options, short positions

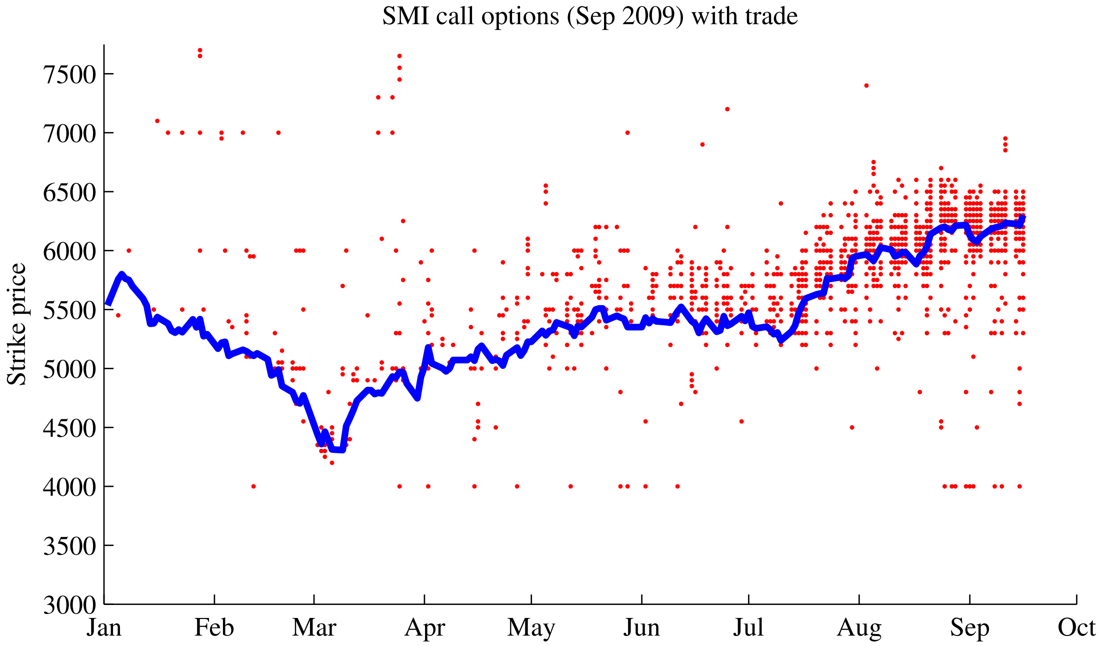

Figure 15.8: Traded options

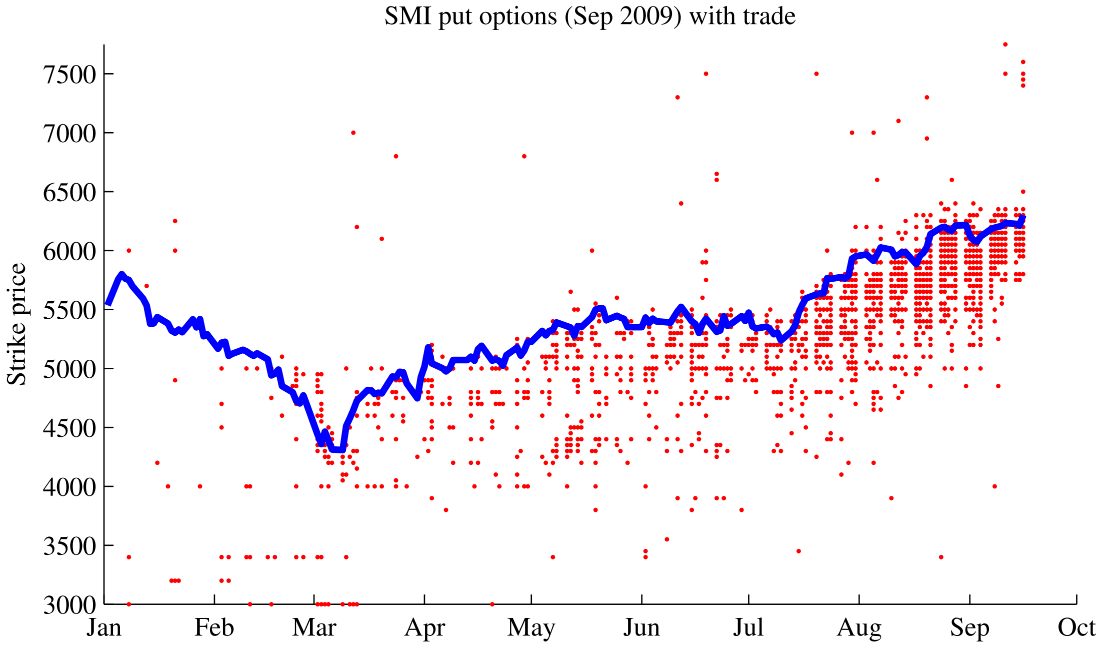

Figure 15.9: Traded options

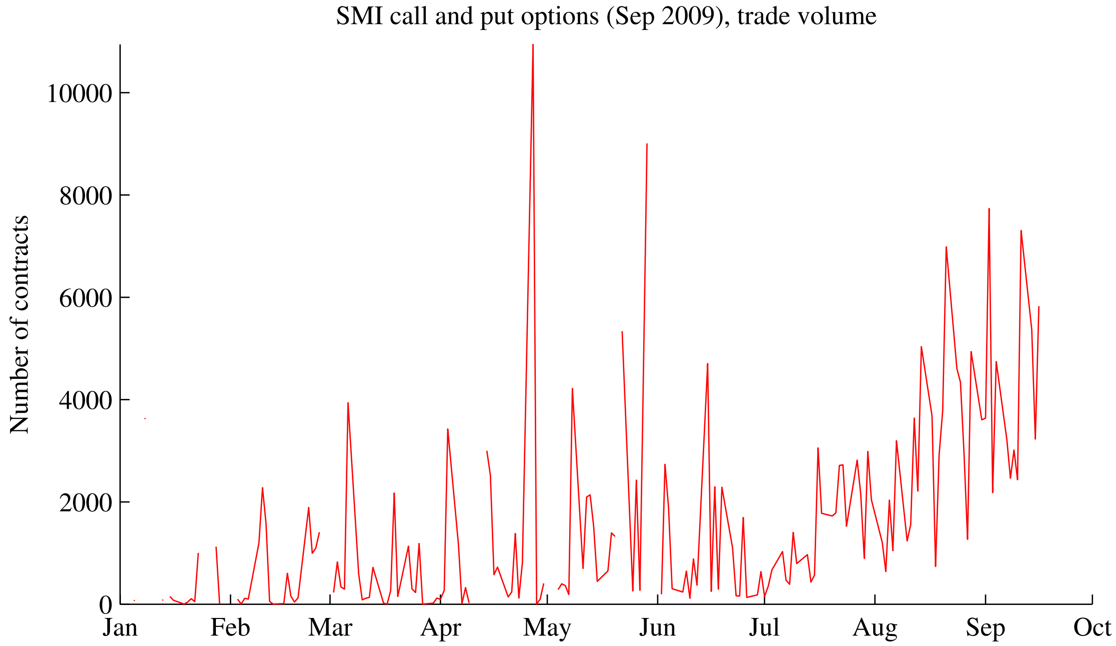

Figure 15.10: Option trade volume

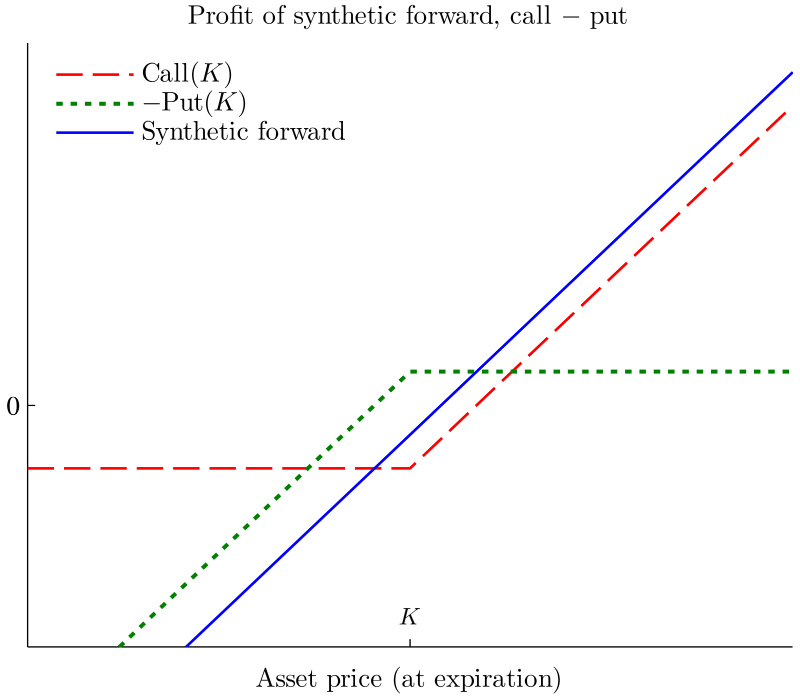

Figure 15.11: Profit of an option portfolio that replicates a forward contract

15.3.2 Financial Engineering

Replicating a Forward

Options markets are often very liquid—and are therefore useful for constructing replicating portfolios. The portfolio Call(K) - Put(K) for $K = F$ (the forward price) replicates a forward contract, so it is a synthetic forward. (Actually, Call(K) - Put(K) for any K will replicate a forward–but with an intial payment. See the "put-call parity" below.) Clearly, we can then replicate a short position in a forward contract by selling such a portfolio. See Figure 15.11.

Portfolio Insurance

A *protective put* is a combination of a put and a position in the underlying asset. This allows the owner to capture the upside of the price movement (of the underlying), at the same time as insuring against the downside. This is indeed very similar to just buying a call option. See Figure 15.12.

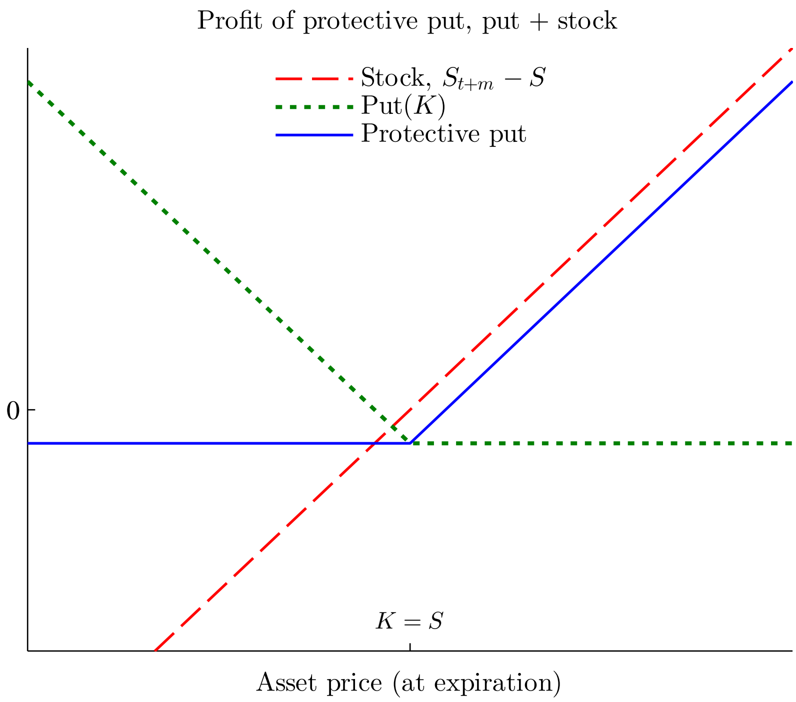

Figure 15.12: Profit of an option portfolio than insures a stock

Betting on Large Changes

An option is a bet on a change in a specific direction. Option portfolios can be constructed to instead make a bet on a large change in either direction (that is, high volatility): a *straddle* is Call(K) + Put(K), and a *strangle* is Call(K_2) + Put(K_1) where $K_1 < K_2$. See Figure 15.13.

Betting on a Large Price Decrease 2

A variation on the synthetic short forward is the *collar*: $-$Call(K_2) + Put(K_1) where $K_1 < K_2$. It also looks like a short position in a forward contract, except that the payoff is flat between the strike prices. Clearly, this is betting on a large price decrease. Selling a collar (or *reversal*) is instead a bet on a large price increase.

A collar (reversal) can be used to hedge a long (short) position in the underlying asset, except that there is no hedge between the strike prices. It provides insurance outside the strike prices. See Figure 15.14.

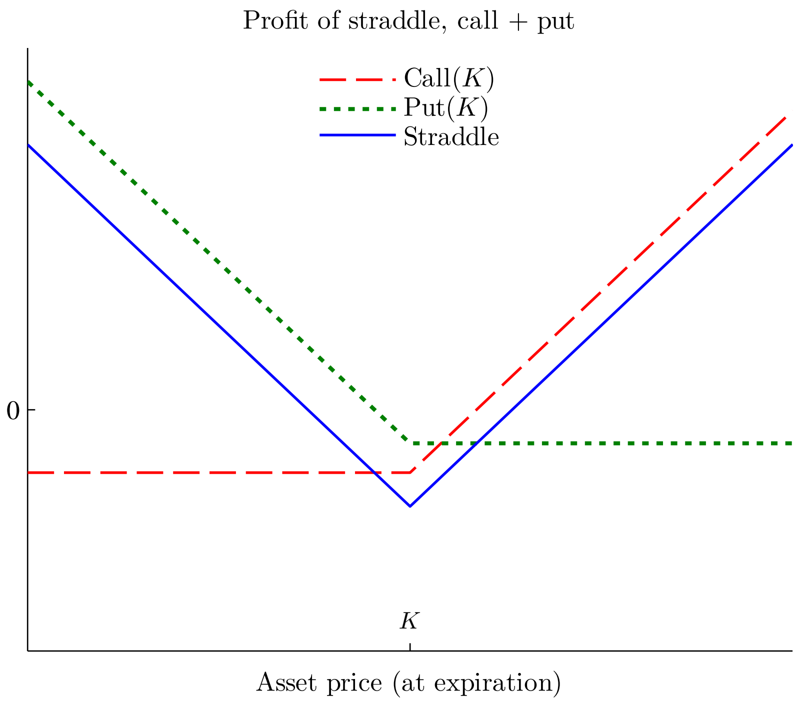

Figure 15.13: Profit of an option portfolio than bets on volatility

Betting On a Small Price Increase

To bet on a small increase in the price of the underlying asset we can use a *bull spread*: Call(K_1) - Call(K_2) where $K_1 < K_2$. This portfolio has flat payoffs outside the strike prices, but a a payoff that increases with the underlying asset between them. Selling a bull spread creates a *bear spread*, which is a bet on a small decrease of the underlying price. (These spreads can also be constructed by combing puts.) See Figure 15.14.

15.3.3 Options Are Risky Assets

The buyer always stands the risk of getting a zero payoff, that is, a return of -100%. For instance, the net return on a European call option is

$$\text{return on call}_{t+m} = \frac{\max\left(0, S_{t+m} - K\right)}{C} - 1, \tag{15.15}$$

where C is the call option price. Whenever the option isn't exercised, the whole investment is lost (and the return is -100%).

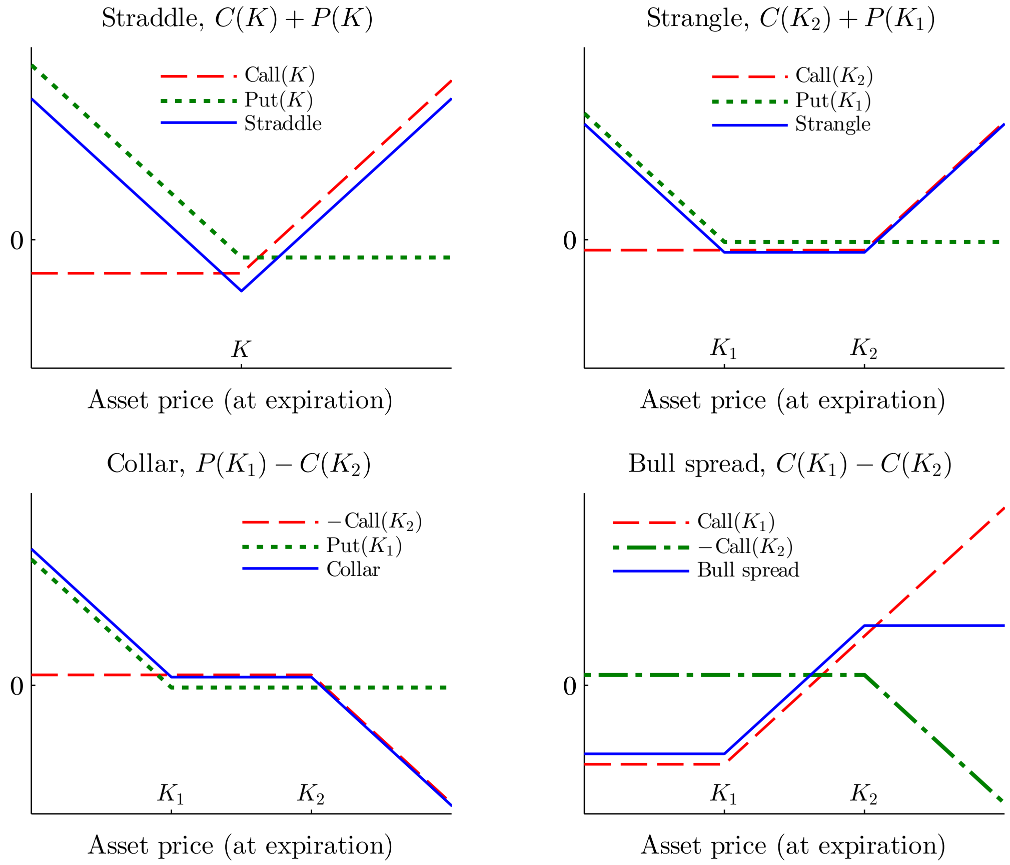

Figure 15.14: Profits of option portfolios

It is the clear that option returns cannot be normally (or even lognormally) distributed: the density function has a spike at -100% (whose probability is the same as the probability of $S_{t+m} \leq K$). This means, that we cannot motivate "mean-variance" pricing of options by referring to a normal distribution of the return. (This does not rule out mean-variance pricing, which could be motivated by, for instance, mean-variance preferences.)

Since options are exposed to risk factors, they can be used to hedge risk, that is, to create an "insurance." For instance, an owner of the underlying asset can hedge by buying a put option. This guarantees that he/she always gets at least the strike price.

Similarly, an investor who has short-sold the underlying asset (borrowed the asset by someone and then sold it) can hedge by buying a call option. This puts a limit (the strike price) on how much he/she will have to pay for the asset when it is time to turn it back to

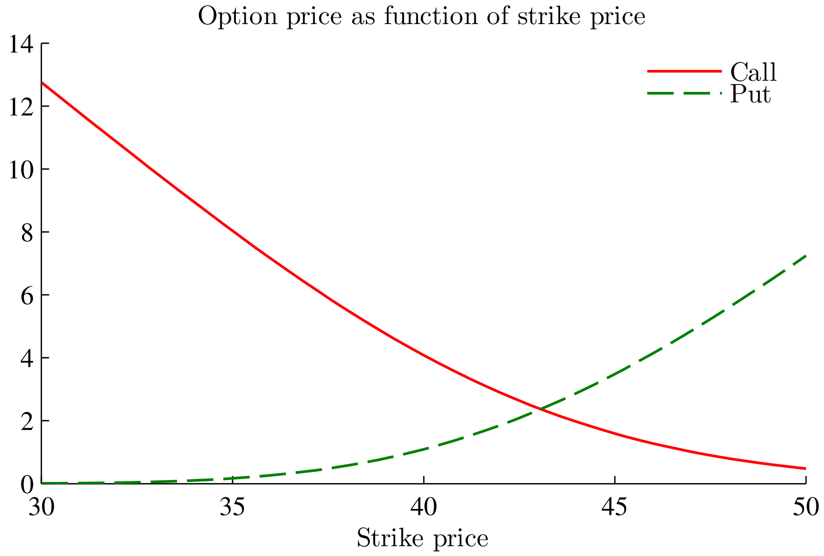

Figure 15.15: Option price as a function of the strike price

the lender.

15.3.4 Basic Properties of Option Prices

Options prices depend on many things, but there are some fairly general results

First, call option prices are decreasing in the strike price, while put options prices are increasing in the strike price. See Figure 15.15 for an illustration.

The intuition is that a higher strike price means that an owner of a call option will have to pay more in case of exercise—and there is also a lower chance of exercise. This is illustrated in Figure 15.16.

Second, both call and put option prices are typically increasing in the dispersion of the distribution of the future price of the underlying asset. The intuition for the second result is that a wider dispersion increases the probability of a really high price of the underlying asset (which is good). Of course, it also increases the probability of a really low asset price, but that is of less concern is the call option payoff is bounded below at zero. This is illustrated in Figures 15.17–15.18.

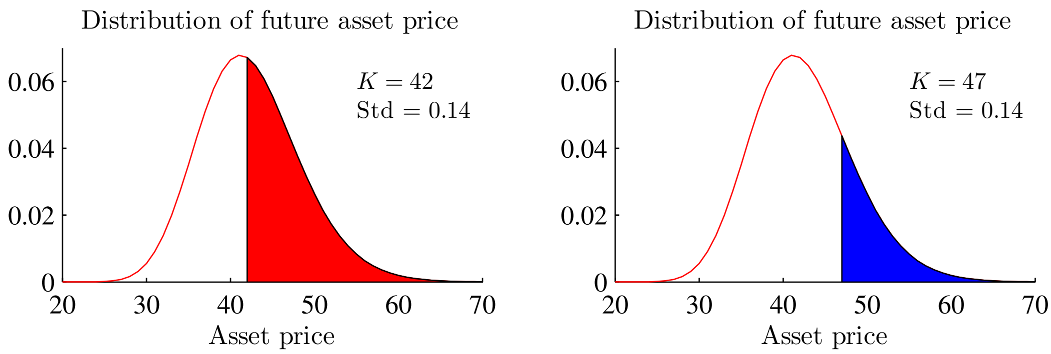

Figure 15.16: Distribution of future stock price

15.3.5 Definition of American Calls and Puts

An American option is like a European option, except that it can be exercised on any day before or on the expiration date. This means that an American option has more rights than a European option and is therefore worth at least as much

$$C_A \geq C_E \text{ and } P_A \geq P_E. \tag{15.16}$$

If the (American) option is exercised, then the immediate payoff is $\max(0, S - K)$ for a call and $\max(0, K - S)$ for a put. This means, that the options prices must (at any point in time) obey

$$\begin{aligned} C_A &\geq \max(0, S - K) \\ P_A &\geq \max(0, K - S). \end{aligned} \tag{15.17}$$

(It can be shown that the first of these expressions is true also for European style options, but the second need not be.)

We will later demonstrate the following results. First, if there are no dividends, then it is never optimal to exercise an American call option early (such a call option will have the same price as a European call option), but it can still be optimal to exercise an American put option early. Second, if there are dividends, then the American call option should only be exercised just prior to the dividend payments, while an American put should perhaps also be exercised also at other times. See Figures 15.19–15.20.

Remark 15.9 *Figures 15.4 and 15.19–15.21 provide an example of how the futures price*

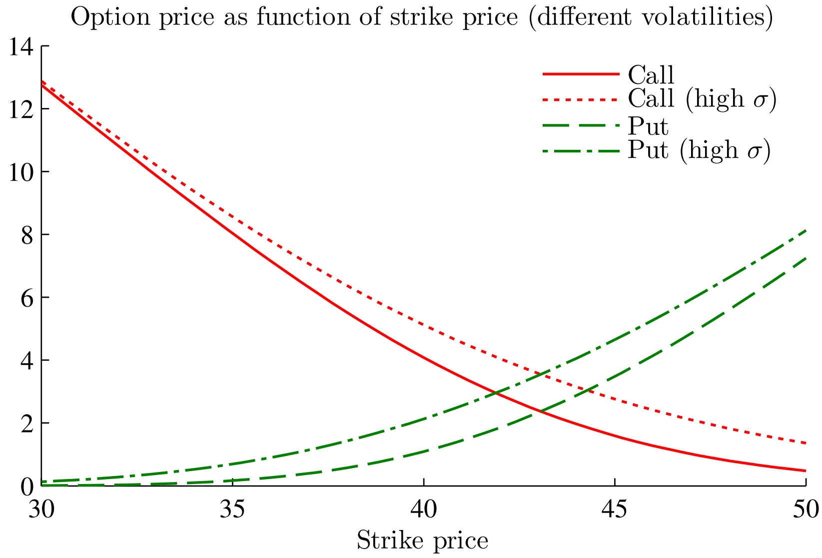

Figure 15.17: Option price as a function of the strike price

(on SP 500), the intrinsic value of the option and the option price developed over a year. Notice how the futures prices converges to the index level at expiration of the futures. Before it can deviate because of delayed payment (+) and no part in dividend payments (−). Also notice that even options with zero intrinsic value (zero payoff if exercised now) can have fairly high option prices—at least if the time to expiration is long, but it converges to zero the expiration date gets closer.

15.4 Put-Call Parity for European Options

There is a tight link between European call and put prices. If you know one of them (and the forward price), then you can easily calculate that the other must be. The following proposition is more precise.

Proposition 15.10 *(Put-call parity for European options) The put-call parity for European options is*

$$C - P = e^{-my}(F - K), \tag{15.18}$$

where $e^{-my}(F - K)$ is the present value of the forward price minus the strike price.

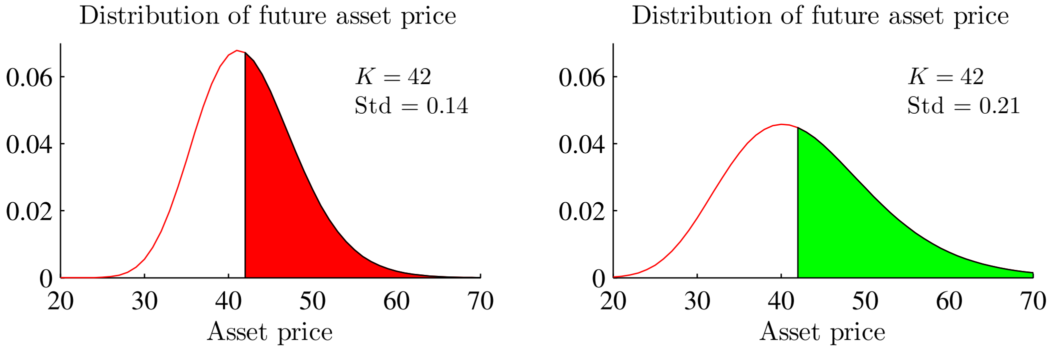

Figure 15.18: Distribution of future stock price

Time subscripts and indicators of maturity have been suppressed to make the notation a bit easier. The parity holds irrespective of whether the underlying asset has dividends or not (since the expression uses the forward price). Its practical importance is that it allows us to use two of the assets to replicate the third asset. For instance, we can combine a call option and a forward contract to replicate a put option, or buy a call and sell a put to replicate a forward contract. Transaction costs can cause (relatively small) deviations from the parity condition.

See Figure 15.11 for an illustration.

Proof. (of Proposition 15.10) Buy one call option and sell one put option, both with the strike price K. This will with certainty give one asset at maturity at the price K. The present value of the cost is $C - P + e^{-my}K$. The same is achieved by entering a forward contract—the present value of the cost is $e^{-my}F$. ■

This formula is very general, but a few special cases are of particular interest. First, when the underlying asset pays no dividends, then (15.18) together with (15.5)–(15.9) give

$$C - P = S - e^{-my}K \text{ if no dividends,} \tag{15.19}$$

$$C - P = S - \sum\nolimits_{i=1}^{n} e^{-m_i y_t(m_i)} D_i - e^{-my}K \text{ if dividends,} \tag{15.20}$$

$$C - P = Se^{-\delta m} - e^{-my}K \text{ if continuous dividend rate } \delta. \tag{15.21}$$

Example 15.11 *(Put-call parity)* $S = 42, m = 1/2, y = 5\%, K = 38$. *If* $C = 5.5$ *for*

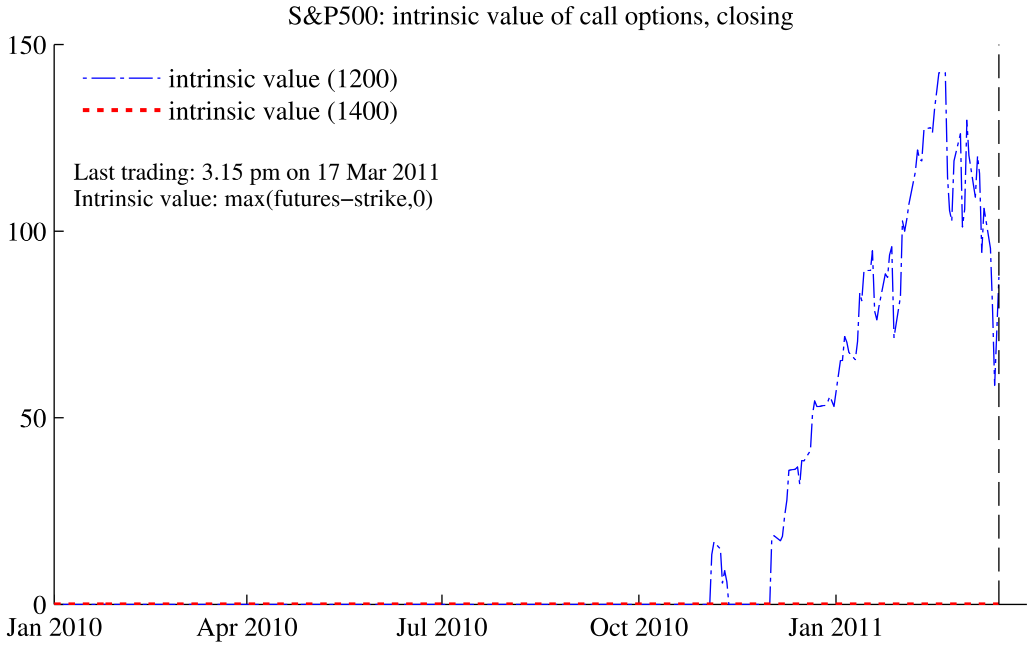

Figure 15.19: Intrinsic value of S&P500 options

an underlying asset without dividends, then (15.19) gives

$$5.5 - P = 42 - e^{-0.5 \times 0.05} 38 \text{ } or \text{ } P \approx 0.56.$$

15.5 Pricing Bounds and Convexity of Pricing Functions

15.5.1 Pricing Bounds for (European and American) Call Options

The price of both American or European call option must satisfy the following restrictions

$$C \leq e^{-my} F \leq S \tag{15.22}$$

$$0 \leq C \tag{15.23}$$

$$e^{-my}(F - K) \leq C \tag{15.24}$$

The motivations are basically as follows (the intuition based on European options, but the results extend to American options as well). First, a call option with a zero strike

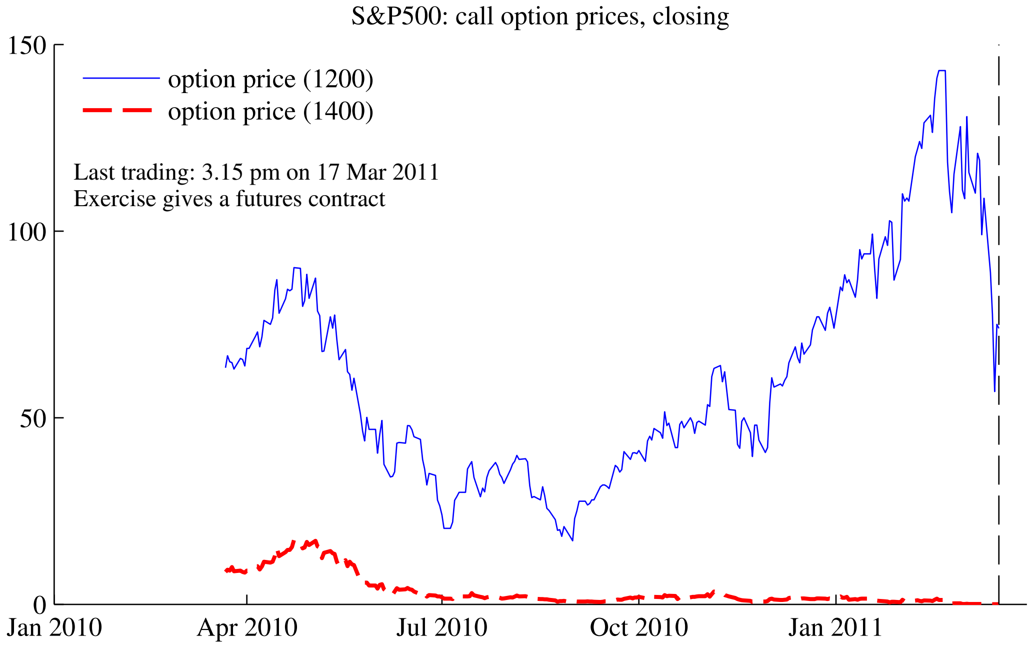

Figure 15.20: S&P500 options

price ($K = 0$) would be the same as owning a prepaid forward contract (which is worth as much or less than the underlying asset). Whenever the strike price is higher, the call price is lower. Second, the call option gives rights, not obligations: its price value cannot be negative. Third, the lowest possible value of a put option is zero, so the put-call parity (15.18) immediately gives that the call price must exceed the present value of $F - K$. (See below for an alternativ proof.) Transaction costs can cause (relatively small) failures of the bounds.

Combining the bounds, we get

$$C \leq e^{-my} F \leq S \tag{15.25}$$

$$C \geq \max[0, e^{-my}(F - K)]. \tag{15.26}$$

In particular, for a financial asset without dividends (until expiration of the option), we have $C \leq S$ and $C \geq \max(0, S - e^{-my} K)$. See Figures 15.22, 15.23 and 15.25 for illustrations.

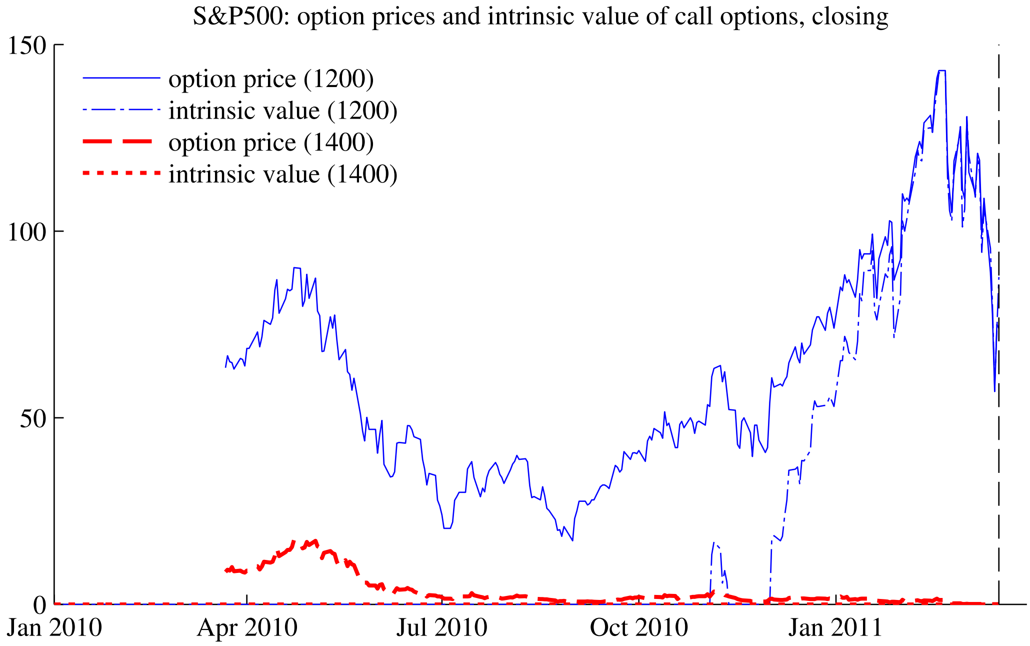

Figure 15.21: S&P500 options

Remark 15.12 *(The price pattern in Figure 15.23) At very low strike prices, it is almost certain that the option will be exercised at expiration. The present value of the cost is $C - e^{-ym}K$. This must be the same as the present value of a a forward contract, $e^{-ym}F$, since it achieves the same. Combining gives $C = e^{-my}(F - K)$. At very high strike prices, the probability of exercise is almost zero—so the option price is too.*

Example 15.13 *(Pricing bounds for call option) Using the same parameters as in Example 15.11, we get the following bounds*

$$4 \leq C \leq 42.$$

In this case, the second bound ($0 \leq C$) is superfluous.

Proof. (of (15.24)) Portfolio A: one European call option and $e^{-my}K$ on a bank account. At expiration, this portfolio is worth S_{t+m} if the option is exercised, and K otherwise: $\max(S_{t+m}, K)$. Portfolio B: one prepaid forward contract, which is worth S_{t+m}

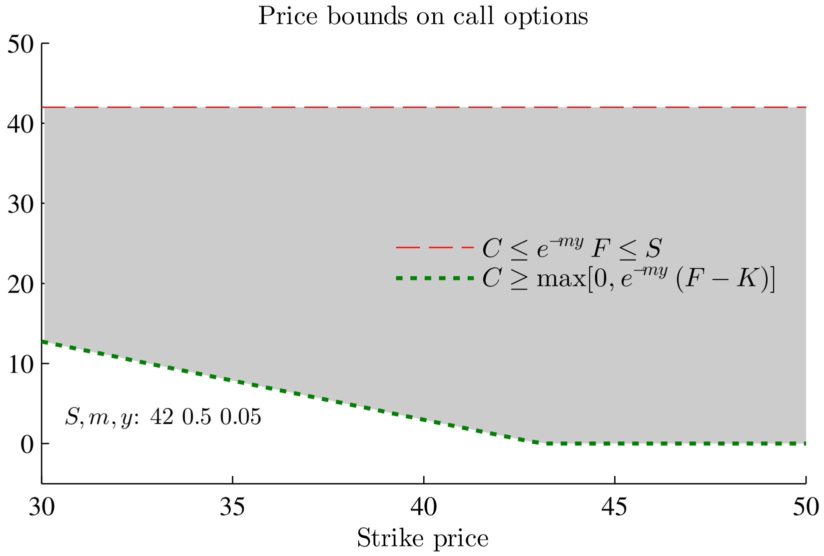

Figure 15.22: Call option price bounds as a function of the strike price

at expiration. (Since you pay $e^{-my}F$ now, there is no payment at expiration.) Clearly, portfolio A is always worth more at expiration, so it must also be worth more right now: $C_E + e^{-my}K \geq e^{-my}F$. Rearrange to get (15.24). Since $C_A \geq C_E$, the bound holds also for an American call option. ■

15.5.2 Pricing Bounds for (European and American) Put Options

The price of American or European call option must satisfy the following restrictions

$$P_E \leq e^{-my}K \text{ and } P_A \leq K \tag{15.27}$$

$$0 \leq P_E \text{ and } 0 \leq P_A \tag{15.28}$$

$$e^{-my}(K - F) \leq P_E \text{ and } K - S \leq P_A. \tag{15.29}$$

See Figure 15.24.

The motivations are as follows. First, the payoff from a put option is $\max(K - S, 0)$, so the maximum value is the strike price (when $S = 0$). For the European put, this payoff occurs at maturity, so the maximum value is the present value of the strike price. Second, the put option gives rights, not obligations: its price value cannot be negative. Third, the lowest possible value of a call option is zero, so the put-call parity (15.18) immediately

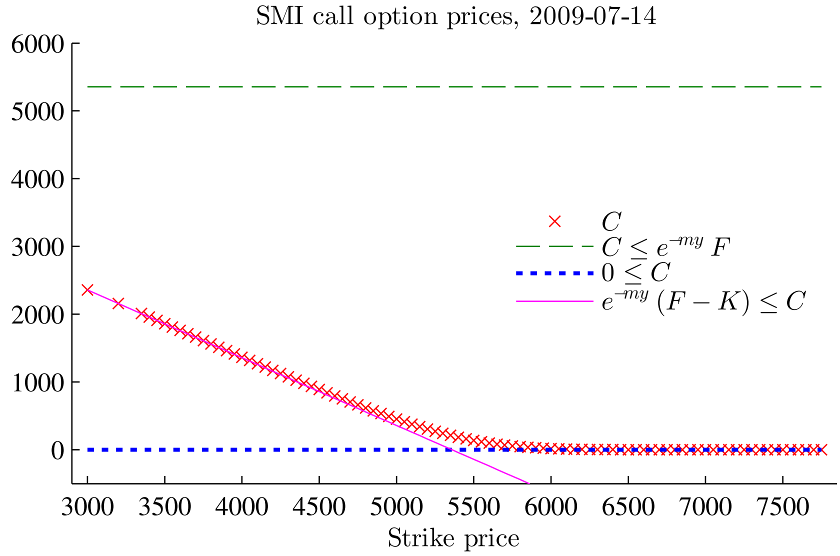

Figure 15.23: Prices and bounds for SMI options

gives that the European put price must exceed the present value of $K - F$. (See below for an alternativ proof.) In contrast, the American put can be exercised now so its value must be at least as high as the intrinsic value.

Proof. (of (15.29)) Portfolio A: one European put option and a prepaid forward contract. At expiration, this portfolio is worth K if the option is exercised, and S_{t+m} otherwise: $\max(K, S_{t+m})$. Portfolio B: $e^{-my}K$ on a bank account, which is worth K at expiration. Clearly, portfolio A is always worth more at expiration, so it must also be worth more right now: $P_E + e^{-my}F \geq e^{-my}K$. Rearrange to get (15.29). Since $P_A \geq P_E$, the bound holds also for an American put option. ■

15.5.3 Prices of Call Options for Different Strike Prices

Suppose we have American or European call options with different strike prices, $K_1 < K_2$. We the have the following price relations

$$C(K_2) - C(K_1) \leq 0 \tag{15.30}$$

$$\frac{C(K_2) - C(K_1)}{K_2 - K_1} \geq -1 \tag{15.31}$$

$$C\left[\lambda K_1 + (1-\lambda)K_2\right] \leq \lambda C(K_1) + (1-\lambda)C(K_2), \text{ for } 0 \leq \lambda \leq 1. \tag{15.32}$$

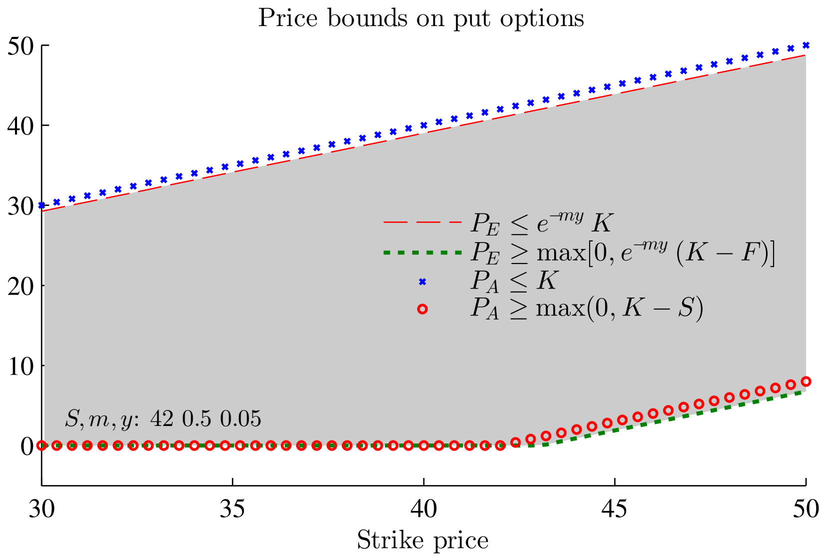

Figure 15.24: Put option price bounds as a function of the strike price

The first relation says that the call option price is decreasing in the strike price. The intuition is that a higher strike price means that an owner of a call option will have to pay more in case of exercise—and there is also a lower chance of exercise. The second relation says that change is smaller than the change in the strike price. The third relation says that the relation is convex. If these relations do not hold, then there are arbitrage opportunities (see the proofs below).

In other words, these three conditions say that we have the following partial derivatives (if they exist) of the call option price function

$$-1 \leq dC(K)/dK \leq 0 \text{ and } dC^2(K)/dK^2 \geq 0. \tag{15.33}$$

This means that the call option price is decreasing in the strike price, but slower than the strike price itself, but that the curve flattens out at high strike prices.

See Figure 15.15 for an illustration.

Proof. (of (15.30)) If (15.30) was not true, so $C(K_2) > C(K_1)$, then a bull spread (buy $C(K_1)$ and sell $C(K_2)$), would have a negative price ($C(K_1) - C(K_2) < 0$). However, the payoff of a bull spread is

$$\max(0, S - K_1) - \max(0, S - K_2) = \begin{cases} 0 & \text{if } S \le K_1 \\ S - K_1 & \text{if } K_1 < S \le K_2 \\ K_2 - K_1 & \text{if } K_2 < S. \end{cases}$$

This would give a non-negative payoff for a negative asset price, which creates arbitrage opportunities. ■

Proof. (of (15.31)) If (15.31) was not true, so $C(K_1) - C(K_2) \ge K_2 - K_1$, then we can sell a bull spread (sell $C(K_1)$ and buy $C(K_2)$) and invest the proceeds in a T-bill (zero investment). The payoff at expiration (m period later) is then

$$\max(0, S - K_2) - \max(0, S - K_1) = \underbrace{[C(K_1) - C(K_2)]e^{rm}}_{> K_2 - K_1} + \begin{cases} 0 & \text{if } S \le K_1 \\ -(S - K_1) & \text{if } K_1 < S \le K_2 \\ -(K_2 - K_1) & \text{if } K_2 < S. \end{cases}$$

In either case, there is a positive profit (recall that the initial investment is zero), which creates arbitrage opportunities. ■

Proof. (of (15.32)) Let $\bar{K} = \lambda K_1 + (1 - \lambda) K_2$. If (15.32) was not true, so $C(\bar{K}) > \lambda C(K_1) + (1 - \lambda) C(K_2)$, then we can sell $C(\bar{K})$ and buy $\lambda C(K_1) + (1 - \lambda) C(K_2)$ (zero investment). The payoff at expiration (m period later) is then

$$\lambda \max(0, S - K_1) - \max(0, S - \bar{K}) + (1 - \lambda) \max(0, S - K_2)$$

$$= \begin{cases} 0 & = 0 & \text{if } S \le K_1 \\ \lambda(S - K_1) & = \lambda(S - K_1) & \text{if } K_1 < S \le \bar{K} \\ \lambda(S - K_1) - (S - \bar{K}) & = (1 - \lambda)(S - K_2) & \text{if } \bar{K} < S \le K_2 \\ \lambda(S - K_1) - (S - \bar{K}) + (1 - \lambda)(S - K_1) & 0 & \text{if } K_2 < S, \end{cases}$$

where the second column uses the definition of $\bar{K}$. All payoffs are non-negative, and some are positive. Since the initial investment is zero, this creates arbitrage opportunities. ■

15.6 Early Exercise of American Options

This section discusses early exercise of American options. There are some cases where we can exclude early exercise, so the American option is priced as a European option. In other cases, we cannot exclude early exercise—but we may still be able to say something about when early exercise is likely. More precise answers will require building a model for the pricing. Clearly, the answer is then model dependent.

15.6.1 Early Exercise of American Call Options (No Dividends)

American call options on an asset without dividends (until expiration of the option) are not exercised early. The following proposition is more precise.

Proposition 15.14 *(No early exercise, American call, no dividends) An American call option on an asset without dividends should never be exercised early (if the interest rate is positive). It therefore has the same price as a European call option.*

See Figure 15.25 for an illustration of the fact that early exercise is not profitable since $C_A \geq C_E > \max(0, S - K)$.

Suppose that you are pretty sure that price of the underlying will drop tomorrow. The above argument shows that you should still not exercise the call option, but it might be sensible to sell the option today. If we exercised early, then we would effectively through away the downside protection inherent in the call option and be left with the underlying asset and also pay the strike price now instead of later—neither of which is good (and which a potential buyer of the call option would be willing to pay for).

Proof. (of Proposition 15.14) Notice from (15.23)–(15.24) that, as long as the interest rate is positive, $C_A > \max(0, S - K)$ for any underlying asset without dividends (since $S = e^{-my}F$ and $K > e^{-my}K$). Exercising now, which gives $\max(0, S - K)$, is therefore never optimal since the market price of the option is higher. ■

Proof. (of Proposition 15.14, alternative) Consider an investor who is willing to keep the underlying asset until tomorrow (at least). Clearly, such investors must exist, or else the underlying asset would not be worth its current price. Portfolio A: one American call option and $e^{-y/365}K$ on a bank account. Tomorrow, this portfolio is worth S_{t+1} if the option is exercised, and K otherwise: $\max(S_{t+1}, K)$. Portfolio B: one unit of the underlying asset, which is worth S_{t+1} tomorrow. Clearly, portfolio A is always worth more tomorrow, so it must also be worth more right now: $C_A + e^{-y/365}K \geq S$. Rearrange gives $C_A \geq S - e^{-y/365}K$, and we also know that $C_A \geq 0$, so $C_A > \max(S - K, 0)$ as long as the interest rate is positive. If you are not an investor who is willing to keep the underlying asset until tomorrow, then you should sell the option to such an investor. ■

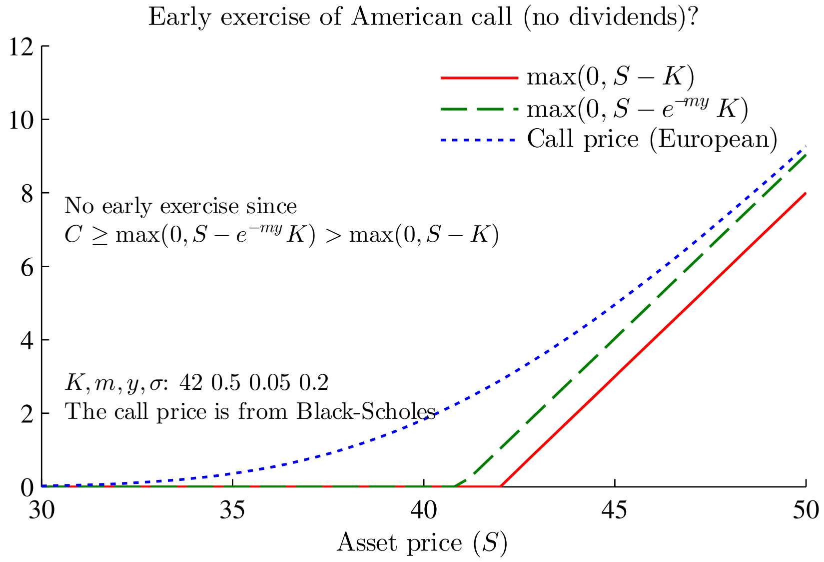

Figure 15.25: Early exercise of American call option (no dividends)

15.6.2 Early Exercise of American Put Options (No Dividends)

American put options on an asset without dividends (until expiration of the option) may be exercised early. The following proposition is more precise.

Proposition 15.15 *(Early exercise, American put, no dividends) An American put option on an asset without dividends could be exercised early. However, we can rule out early exercise when $P_E > \max(0, K - S)$, since $P_A \geq P_E$ then implies that selling the option is better than exercising. From the put-call parity for European options, we notice that $P_E > \max(0, K - S)$ happens when $C_E > (1 - e^{-my})K$. For instance, this is always the case if the interest rate is zero—so there is no early exercise. This holds when $K < S$ (the put is out of the money), when K is slightly above S (the put is in the money, but not much), but not when the put is deep in the money. Hence, early exercise is only possible when the asset price is very low compared to the strike price.*

See Figure 15.26 for an illustration of the fact that early exercise is not profitable (since $P_E > \max(0, K - S)$) for high asset prices, but might be so for low asset prices (since $P_E < \max(0, K - S)$ means that $P_A < \max(0, K - S)$ is possible). Clearly, the proposition relies on having information about a European put price—or a good model of

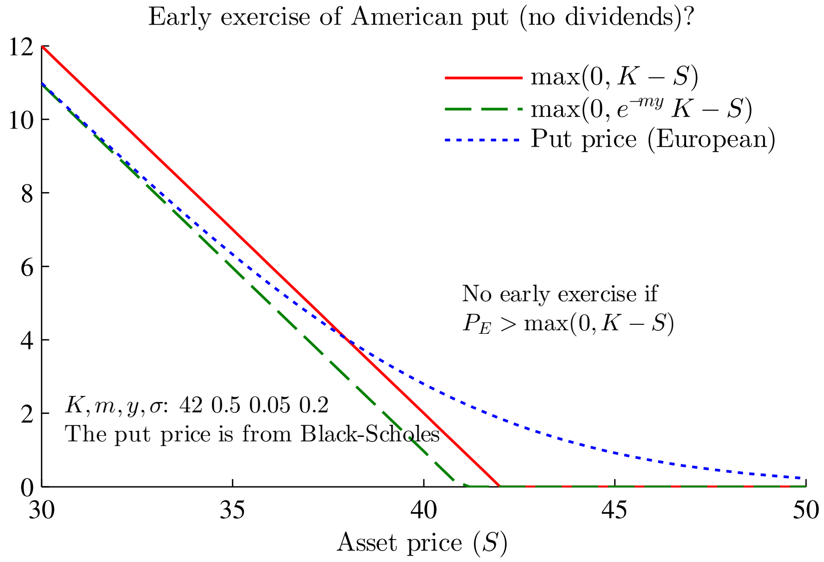

Figure 15.26: Early exercise of American put option (no dividends)

what the price should be. If we do not have that information, the proposition is not very useful—except in telling us that early exercise is more likely if the asset price is low and the interest rate high.

Since there is no early exercise at high asset prices, the American put is similar to a European put—so their prices are also similar. In contrast, at low asset prices early exercise is probable, so the American put is worth more than the European put. See Figures 15.27–15.28 for an illustration, based on a numerical solution for the price on an American put option. The first figure shows in which nodes early exercise is optimal: at low asset prices. The second picture illustrates how the price is related to the European put price and a upper boundary (to be discussed later).

Example 15.16 *(Early exercise of American put option?) When the underlying asset goes bankrupt, then $S = 0$ and it is known that it will stay at $S = 0$. Exercising the American put option now gives K, whereas waiting until expiration has a present value of $e^{-my}K$ (which is lower): early exercise is optimal.*

Example 15.17 *(Early exercise of American put option?) Using the same parameters as*

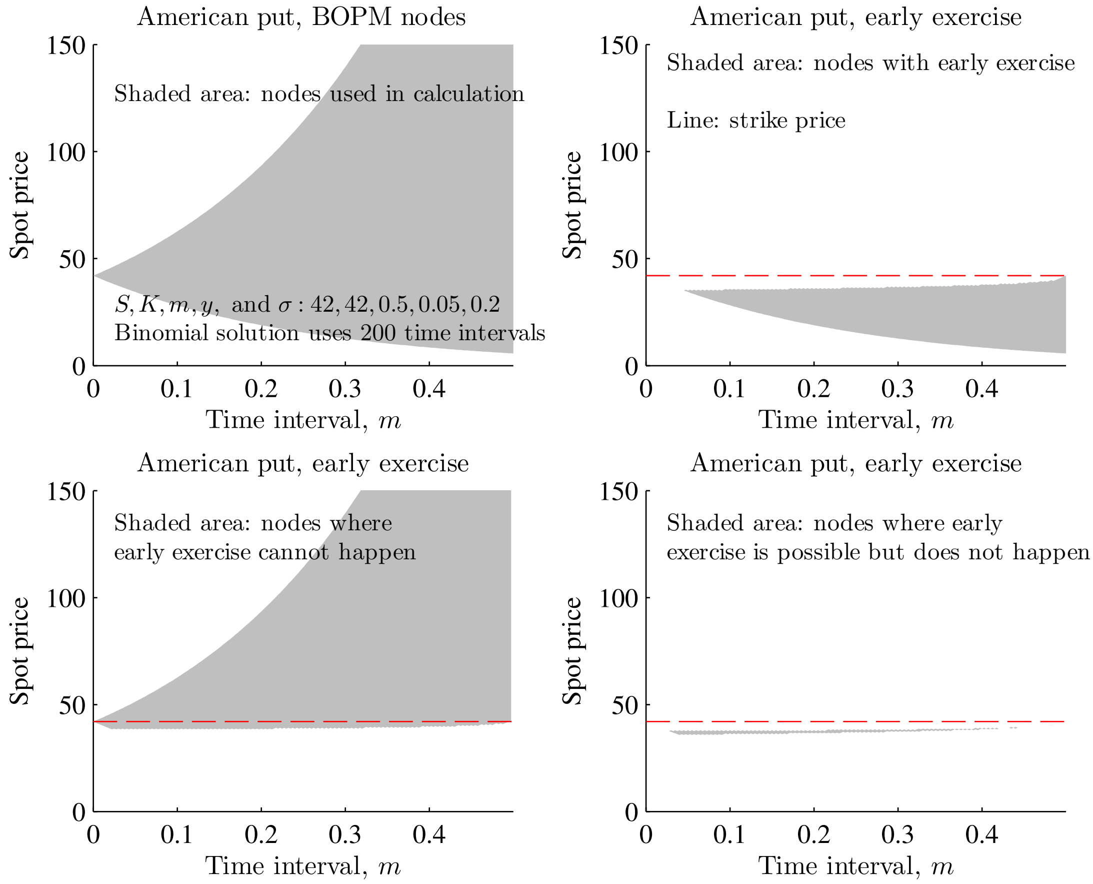

Figure 15.27: Numerical solution of an American put price (no dividends)

in Example 15.11, we have that $C_E > (1 - e^{-my})K$ is satisfied since

$$5.5 > (1 - e^{-1/2 \times 0.05})38 = 0.94,$$

so there is no early exercise of the American put option. The reason is that we from the put-call parity for European options (15.19) and the fact $P_A \geq P_E$ then have

$$P_A \geq P_E = \underbrace{C_E}_{5.5} + K - S - \underbrace{(1 - e^{-my})K}_{0.94},$$

so selling the put option (getting P_A) gives the same as exercising ($K - S$) plus at least $5.5 - 0.94 = 4.56$, so selling gives more than exercising. If, for some reason, we instead have $y = 35\%$ (so $(1 - e^{-my})K = (1 - e^{-1/2 \times 0.35})38 = 6.1$) but the same prices, then

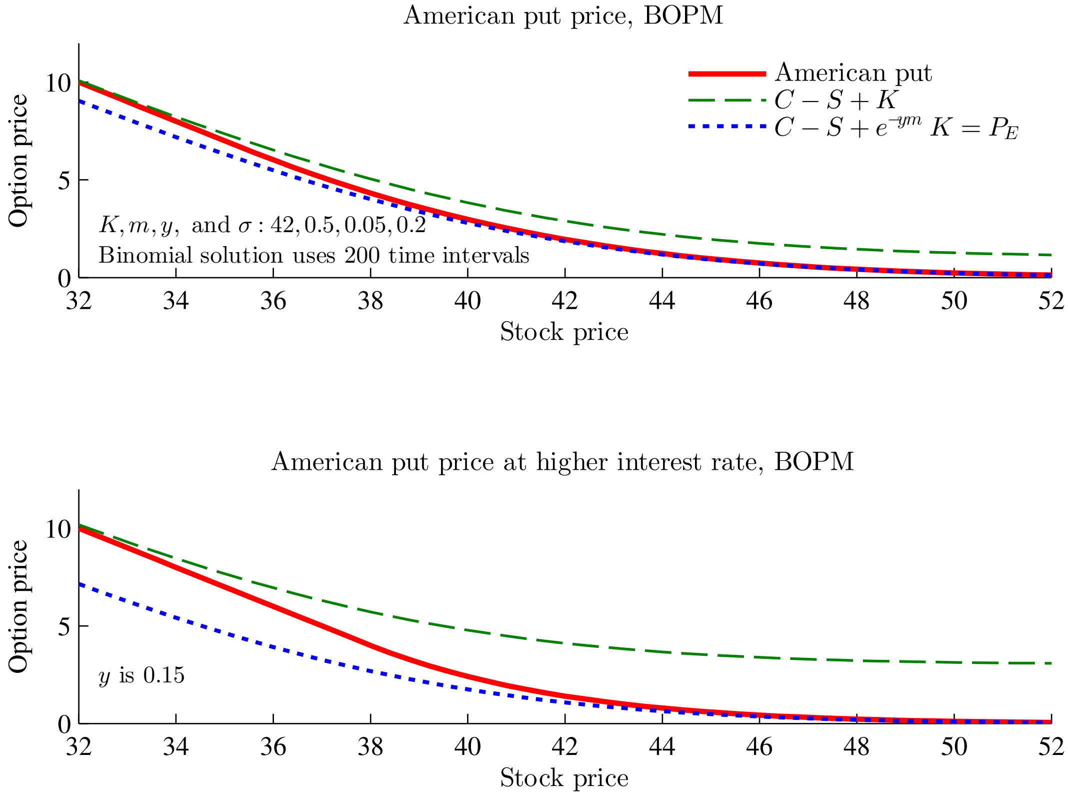

Figure 15.28: Numerical solution of an American put price (no dividends)

we would perhaps get early exercise. In particular, the expression above would say

$$P_A \geq P_E = \underbrace{C_E}_{5.5} + K - S - \underbrace{(1 - e^{-my})K}_{6.1},$$

so selling the put gives the same as exercising $(K - S)$ *plus at least* $5.5 - 6.1 = -0.6$*, so it's not sure that selling is better than exercising (could be or not).*

Proof. (of Proposition 15.15) To avoid early exercise, selling (getting P_A) should be more profitable than exercising (getting $K - S$), $P_A > K - S$. Put-call parity for European options (15.19) says

$$P_E = C_E + K - S - (1 - e^{-my})K.$$

If

$$C_E > (1 - e^{-my})K,$$

then $P_A \geq P_E > K - S$ so selling is better than exercising. This means that there is no early exercise if the European call price is high (high asset price compared to strike price), the strike price is low, or if the discounting until expiration is low (low interest rate or small time to expiration). For instance, with a zero interest rate, $P_A \geq C_E + K - S$, so there is never early exercise as long as $C_E > 0$. If these conditions are not satisfied, we cannot rule out early exercise. ■

15.6.3 Early Exercise of American Call and Put Options (Dividends)

American call and put options on an asset with dividends (until expiration of the option) may be exercised early. The following propositions are more precise.

Proposition 15.18 *(Early exercise, American call, dividends) An American call option on an asset with dividends could be exercised early, especially just before a dividend payment and when the option is deep in-the-money (low strike price/high asset price). Conversely, there is no early exercise if* $(1 - e^{-my})K > \sum_{i=1}^{n} e^{-m_i y_t(m_i)} D_i$*, that is, with a high strike price and low present value of the dividends.*

Example 15.19 *(Early exercise, American call, dividends?) Suppose there is one dividend payment one month ahead:* $D_1 = 0.95$ *at* $m_1 = 4/12$*. If we use the same parameters as in Example 15.11, we then have*

$$(1 - e^{-1/2 \times 0.05})38 = 0.94 > e^{-4/12 \times 0.05} 0.95 = 0.93,$$

so we can rule out early exercise. However, if the dividend payment is at $m_1 = 1/12$*, then we cannot.*

Proof. (of Proposition 15.18) To avoid early exercise, selling (getting C_A) should be more profitable than exercising (getting $S - K$), $C_A > S - K$. Put-call parity for European options (15.20) says

$$C_E = S - K - \sum_{i=1}^{n} e^{-m_i y_t(m_i)} D_i + (1 - e^{-my})K + P_E.$$

If

$$(1 - e^{-my})K > \sum_{i=1}^{n} e^{-m_i y_t(m_i)} D_i,$$

and $P_E \geq 0$ (always true), then $C_A \geq C_E > S - K$: selling is better than early exercise. Hence, there is no early exercise if the present value of dividends is low, the strike price is high or if the discounting until expiration is large (high interest rate or long time to expiration). In the opposite case, we cannot rule out early exercise. ■

Proposition 15.20 *(Early exercise, American put, dividends) Early exercise is possible...*

15.7 Put-Call Relation for American Options

There is no put-call parity for American options. However, pricing bounds (based on the values of European options) can be derived.

Proposition 15.21 *(Put-call, American option, no dividend) For an American option on an asset without dividends, the put price must be inside the interval*

$$\underbrace{C_A - S + e^{-my} K}_{P_E} \leq P_A \leq \underbrace{C_A}_{C_E} - S + K. \tag{15.34}$$

The lower boundary is the European put price from (15.19). The reason is that the American and European call options have the same prices (the American call option on an asset without dividends is never exercised early—see Section 15.6). The upper bound is very similar, except that it involves the strike price, not its present value. Clearly, when the interest rate is low, then the interval is narrow—and with a zero interest rate it collapses to the put-call parity of European options. (The latter corresponds to the fact that an American put option on an asset without dividends is never exercised early if the interest rate is zero, see Section 15.6).

See Figures 15.28 and 15.29 for illustrations.

Example 15.22 *(Bounds for an American put option) Using the same parameters as in Example 15.11, we get the following bounds for an American put option (no dividends)*

$$0.56 \leq P_A \leq 5.5 - 42 + 38 = 1.5.$$

Proof. (of Proposition 15.21) The lower boundary is the European put price (since $C_A = C_E$ when there are no dividends) and it is always true that $P_A \geq P_E$.

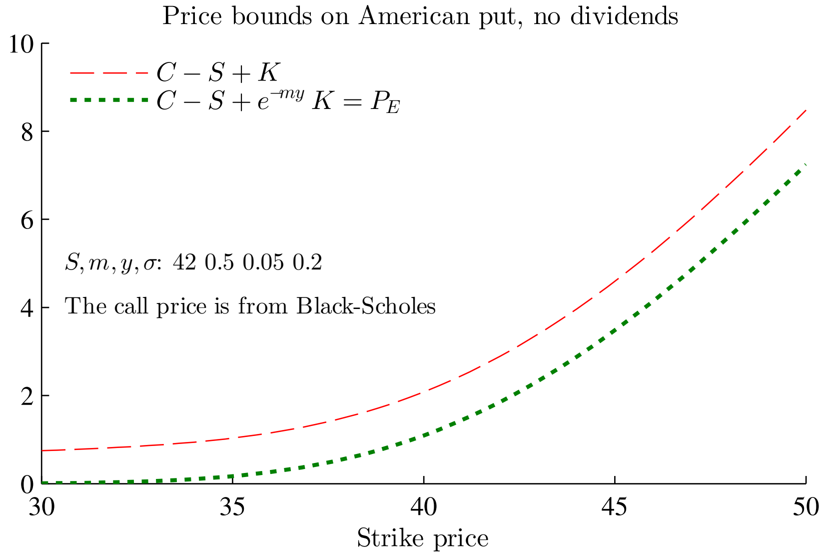

Figure 15.29: Option price as a function of the strike price

The upper boundary follows from the following argument where we compare two portfolios. Portfolio A: one call option with strike price K plus a deposit of K. Portfolio B: one put option plus one underlying asset. If the put option is held until expiration (the call is not exercised early), then portfolio A will be worth $\max(0, S_m - K) + e^{my}K$ in period m (where m is date of expiration), and portfolio B will be worth $\max(0, K - S_m) + S_m$, so portfolio A is worth (weakly) more. If, instead, the put is exercised earlier ($l < m$), then portfolio A will be worth $C_{A,l} + e^{ly}K$ in period l, and portfolio B will be worth $K - S_l + S_l = K$, so portfolio A is worth (weakly) more. In period 0 ($0 \leq l < m$) we don't know when/if the early exercise of the put will happen—but we know that in either case A portfolio will then be worth more than a portfolio B: portfolio A must therefore be worth (weakly) more than B already in 0: $C_{A,0} + K \geq P_{A,0} + S_0$, which is the upper bound in (15.34). ■

Proposition 15.23 *(Put-call, American option, dividends) With dividends, the upper boundary in (15.34) is changed by adding the present value of the dividend stream*

$$C_A - S + e^{-my}K \leq P_A \leq C_A - S + K + \sum\nolimits_{i=1}^{n} e^{-m_i y_t(m_i)} D_i. \tag{15.35}$$

Notice that the lower boundary is not equal to the European put price anymore (since $C_A \geq C_E$ and the present value of the dividends is not added). Together this means that the interval is wider with dividends than without dividends.

Proof. (of Proposition 15.23) The lower boundary follows from the following argument. Buy one call option, lend $e^{-my}K$, and sell one asset—the total value is $C_A + e^{-my}K - S$, which is the left hand side of (15.35). If the call is exercised prior to expiry, the payoff is $S - K + e^{-my}K - S = (e^{-my} - 1)K < 0$ which must be less than the value of the put whose value is nonnegative. If no early exercise, then the payoff at expiration is $\max(0, S - K) + K - S = \max(0, K - S)$ which is the same as the put payoff.

The upper boundary is a bit trickier, so we leave it for now. ■

Bibliography

Cochrane, J. H., 2001, *Asset pricing*, Princeton University Press, Princeton, New Jersey.

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

16 The Binomial Option Pricing Model

Main references: Elton, Gruber, Brown, and Goetzmann (2010) 23 and Hull (2009) 11
Additional references: McDonald (2006) 9–12

16.1 Overview of Option Pricing

There are basically two ways to model option prices: by some sort of factor model (like CAPM) or by a no-arbitrage argument. The latter is clearly much more precise, so it is typically preferred—when it works. These notes focus on a particularly simple case: when the underlying asset follows a binomial process.

16.2 The Basic Binomial Model

The binomial model (where the change of the price of the underlying asset only can take two values) is very stylized, but it is useful for establishing the key ideas of option pricing. It can also be transformed into a realistic model by cumulating many (short) subperiods. In the limit (as the subperiods became very many/very short) it converges to the well-known Black-Scholes model.

16.2.1 Binomial Process for the Stock Price

The binomial tree for the underlying asset starts at the price S and has probability q of moving to Su in the next period and a probability of $1-q$ of moving to Sd. This is illustrated in Figure 16.1. These probabilities are the true ("natural") probabilities. If we denote the price today by S_t and in the next period by S_{t+h}, then we have

$$S_{t+h}/S_t = \begin{cases} u & \text{with probability } q \\ d & \text{with probability } 1-q. \end{cases} \tag{16.1}$$

Remark 16.1 *(Mean and variance of a binomial process) The mean of a (shifted) binomial process like (16.1) is* $qu + (1-q)d$ *and the variance is* $q(1-q)(u-d)^2$.

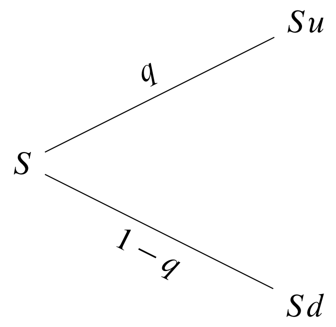

Figure 16.1: Natural binomial process for S

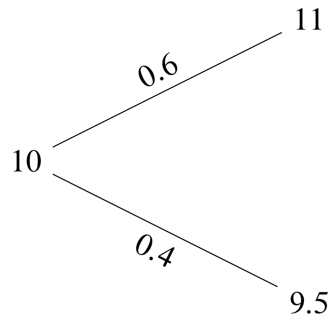

Figure 16.2: Natural binomial process for S

Example 16.2 *(Binomial process) Suppose* $S = 10, u = 1.1, d = 0.95$*, and* $q = 0.6$*. Then, the process has a 60% probability of increasing from 10 to 11 and a 40% probability of decreasing to 9.5. See Figure 16.2. This gives an expected relative change of* $0.6 \times 1.1 + 0.4 \times 0.95 = 1.04$ *and a variance of the relative change of* $0.6 \times 0.4 \times (1.1 - 0.95)^2 = 0.0054$.

16.2.2 No-Arbitrage Pricing of a Derivative

Consider a derivative asset that will be worth f_u in case we end up at Su and f_d if we end up at Sd—see Figure 16.1. Notice that f_u is the notation for the value (price) of the derivative in the up state (it should *not* be read as f times u). As an example, the derivative could be a call option with strike price K, so if the next time period is the time expiration, then

$$f_u = \max(Su - K, 0) \text{ and } f_d = \max(Sd - K, 0). \tag{16.2}$$

Alternatively, it could be a forward contract (and the next time period is the time of expiration), so

$$f_u = Su - F \text{ and } f_d = Sd - F. \tag{16.3}$$

We next use a no-arbitrage argument to derive what today's price of the derivative (denoted f) must be. In doing so, we take it for granted that

$$u > e^{yh} > d. \tag{16.4}$$

If this condition is not satisfied, then there are a trivial arbitrage opportunities. For instance, if $e^{yh} > u$, then we could short the stock and buy bonds: this would guarantee a positive payoff for a zero investment (an arbitrage possibility).

Example 16.3 *(European call option) With the parameters in Example 16.2, equation (16.2) shows that a European call option with strike price of 10 has*

$$f_u = \max(11 - 10, 0) = 1 \textit{ and } f_d = \max(9.5 - 10, 0) = 0,$$

while a strike price of 9 gives

$$f_u = \max(11 - 9, 0) = 2 \textit{ and } f_d = \max(9.5 - 9, 0) = 0.5.$$

Step 1: Construct a Riskfree Portfolio

We now construct the following portfolio

$$\begin{aligned} &\Delta \text{ of the underlying asset, and} \\ &-1 \text{ of the derivative,} \end{aligned} \tag{16.5}$$

where will pick the value of Δ to make the portfolio riskfree.

The payoff of the portfolio at expiry is $\Delta Su - f_u$ in the "up" state and $\Delta Sd - f_d$ in the "down" state. To make the portfolio riskfree, Δ must be such that the payoff is the same in both cases

$$\begin{aligned} \Delta S u - f_u &= \Delta S d - f_d \text{, so} \\ \Delta &= \frac{f_u - f_d}{S\,(u-d)}. \end{aligned} \tag{16.6}$$

With this choice of Δ (also called the "delta hedge") the portfolio is riskfree and must therefore have the same return as the riskfree rate.

Example 16.4 *(European call option) Continuing Example 16.3 we get*

$$\Delta = \frac{1-0}{10(1.1-0.95)} = \frac{2}{3} \text{ for } K = 10, \text{ and } \Delta = \frac{2-0.5}{10(1.1-0.95)} = 1 \text{ for } K = 9.$$

Step 2: Make the Return of the Portfolio Equal to the Riskfree Rate

The present value of our riskfree portfolio is $e^{-yh}\,(\Delta S u - f_u)$, where y is the interest rate per year and h is the length of the time interval. (The present value is also equal to $e^{-yh}\,(\Delta S d - f_d)$, but that is the same, as discussed before.) Since the portfolio is riskfree, this present value must be equal to the cost of the portfolio, $\Delta S - f$,

$$e^{-yh}\,(\Delta S u - f_u) = \Delta S - f. \tag{16.7}$$

Remark 16.5 *(16.7) is clearly the same as saying that the (safe) log return on the portfolio equals yh*

$$yh = \ln\left(\frac{\Delta S u - f_u}{\Delta S - f}\right).$$

Solve for the (current) price of the derivative, f, and use the value of Δ from (16.6) that ensures that the portfolio is riskfree

$$f = \Delta S\left(1 - e^{-yh}u\right) + e^{-yh} f_u \tag{16.8}$$

$$= \frac{f_u - f_d}{u-d}\left(1 - e^{-yh}u\right) + e^{-yh} f_u \tag{16.9}$$

$$= e^{-yh}\left[p f_u + (1-p)\, f_d\right] \text{ with } p = \frac{e^{yh} - d}{u-d} \tag{16.10}$$

$$= e^{-yh}\,\mathrm{E}^*(\text{future payoff of derivative}) \tag{16.11}$$

Equation (16.9) shows what the price of the derivative must be—and is written in terms of the possible outcomes and the interest rate. Notice that neither probabilities (of the

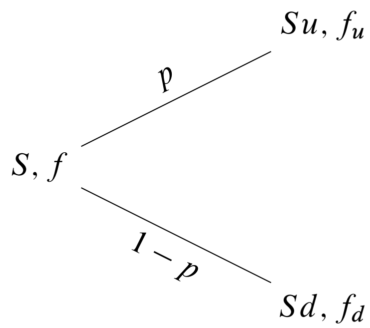

Figure 16.3: Risk neutral binomial process for S and f

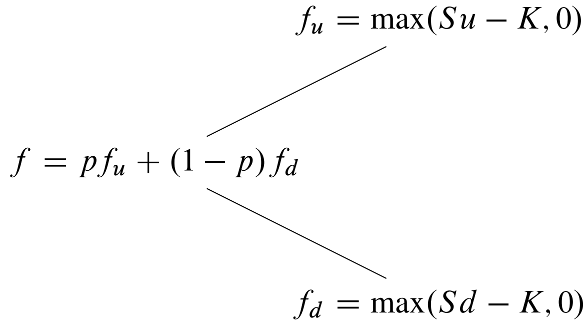

Figure 16.4: Solving for a call option price, zero interest rate

different outcomes), nor risk preferences enter this expression—since we have used a no-arbitrage argument to price this derivative. This works (that is, we can construct a riskfree portfolio) because we have as many (relevant) assets (riskfree and underlying risky asset) as there are possible outcomes (up or down).

Equations (16.9)–(16.11) are alternative ways to write the price of the derivative. Equation (16.10) shows that the current price of the derivative is the discounted value (e^{-yh}) times what seems as an expectation of the payoff of the derivative. This expression is quite useful since we can think of p as a "risk neutral probability"—although it is not a probability in the usual sense: it is just a convenient construction. Notice that p does not depend on which derivative asset (with the same underlying asset) we consider. Under the restrictions in (16.4), $0 < p < 1$, as any "probability" should be. This interpretation is highlighted in (16.11), where E^* stands for the expectations according to the *risk-neutral distribution* (more about that later).

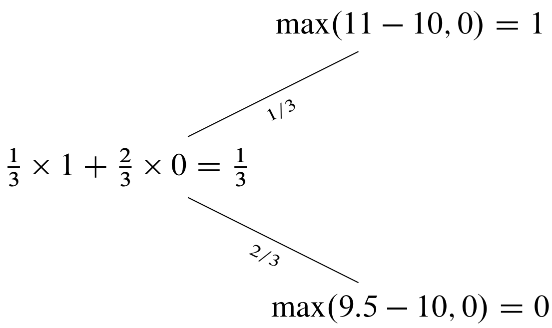

Figure 16.5: Numerical example of call option price, zero interest rate

Example 16.6 *(European call option) Continuing Example 16.3 and assuming that* $y = 0$*, equation (16.10) gives the price of a call option with strike price 10 as*

$$\begin{aligned} f &= e^{-0}\left[p1 + (1-p)\,0\right] \text{ with } p = \frac{1-0.95}{1.1-0.95} = 1/3 \\ &= 1/3. \end{aligned}$$

See Figure 16.5. For the call option with a strike price of 9, we get

$$f = e^{-0}\left[(1/3) \times 2 + (2/3) \times (1/2)\right] = 1.$$

16.2.3 Applying the No-Arbitrage Pricing on Different Derivatives

This section discusses how we apply (16.10) on some special derivatives.

Consider *the underlying asset itself.* It is clearly a (trivial) derivative with $f_u = Su$ and $f_d = Sd$. According to (16.10) the current price of the underlying asset should be

$$S = e^{-yh}\left[pSu + (1-p)\,Sd\right]. \tag{16.12}$$

This looks (again) like a discounted expected future payoff.

Example 16.7 *(The underlying asset itself) Continuing Example 16.6, equation (16.12) gives*

$$S = e^{-0}\left[(1/3) \times 11 + (2/3) \times 9.5\right] = 10.$$

A *forward contract* has a zero current price (nothing is paid until expiry), and the payoff at expiry is $f_u = Su - F$ in the up state (the value of the underlying asset minus the forward price) and $f_d = Sd - F$ in the down state. Using this in (16.10) gives

$$0 = e^{-yh} \left[p \left(Su - F \right) + (1 - p) \left(Sd - F \right) \right], \text{ so} \tag{16.13}$$

$$F = pSu + (1 - p) Sd. \tag{16.14}$$

This shows that the mean of the risk neutral distribution equals the forward price. Combining (16.12) and (16.14) clearly gives the spot-forward parity, $F = e^{yh} S$.

A *riskfree asset* can also be priced by this method. The only way an asset can be riskfree in this setting is if $f_u = f_d$. We then get a zero hedge ratio (Δ) and (16.10) gives

$$f = e^{-yh} f_u, \tag{16.15}$$

which is the discounted value of the (sure) payoff.

An "*Arrow-Debreu asset*" (a sort of theoretical derivative often used in asset pricing models) for the "up" pays off one unit in the up state and zero otherwise ($f_u = 1$ and $f_d = 0$). This is also a so-called "cash-or-nothing" call option provided the up state means that the option is in the money ($Su > K$). From (16.10) we have

$$f = e^{-yh} p. \tag{16.16}$$

16.2.4 Replicating (and Hedging) a Derivative

The no-arbitrage argument in (16.6) was based on the fact that a portfolio of Δ of the underlying asset and of -1 of the derivative replicated a bond.

This argument can be turned around to replicate the derivative by holding the following portfolio

$$\begin{aligned} &\Delta \text{ of the underlying asset, and} \\ &-e^{-yh} \left(\Delta Su - f_u \right) \text{ bills.} \end{aligned} \tag{16.17}$$

The payoff of this portfolio in the up state is $\Delta Su - (\Delta Su - f_u) = f_u$ and in the down state it is $\Delta Sd - (\Delta Sd - f_d) = f_d$ (since $\Delta Su - f_u = \Delta Sd - f_d$). This replicates the derivative's payoff. We can therefore hedge a short position in the derivative by holding Δ of the underlying asset ("delta hedging").

Example 16.8 *(Replicating the call option) For the call option with a strike price of 10 and with a zero interest rate, we have (see Example 16.4)* $\Delta = 2/3$ *and*

$$-e^{-yh}\left(\Delta S u - f_u\right) = -1\left(\frac{2}{3} \times 10 \times 1.1 - 1\right) = -6\frac{1}{3}$$

16.2.5 Where is the Risk Premium?

We have used a no-arbitrage method to price the derivative. It works since the derivative is a redundant asset: it can be replicated by a portfolio of the underlying asset and a riskfree asset—and therefore must have the same price as this portfolio. This does not mean, however, that the option is in itself riskfree. In fact, options are typically very risky and therefore carry large risk premia. It may seem as if the pricing formula (16.10) is free from the preference parameters that would determine the risk premium. Not correct. The pricing formula contains the current asset price (through f_u and f_d) which is indeed affected by preference parameters.

The easiest way to see this is perhaps to recall that we can replicate the portfolio by holding a portfolio of the underlying asset and bills, see (16.17). Clearly, this portfolio will incorporate a risk premium—and so must the derivative.

The only case without a risk premium is when the derivative payoffs are unrelated to the asset price—so the derivative is actually a safe asset as in (16.15).

16.3 Interpretation of the Riskneutral Probabilities

When the underlying asset has a *positive risk premium*, then the expected return (over t to $t + h$) is higher than the risk free rate

$$\frac{\mathrm{E}_t\, S_{t+h}}{S_t} > e^{yh} \text{ (with positive risk premium).} \tag{16.18}$$

From the spot-forward parity for an asset without dividends, we know that the the discounted forward price equals the spot price

$$e^{-yh}F = S_t. \tag{16.19}$$

Combine to get

$$\frac{\mathrm{E}_t\, S_{t+h}}{S_t} > \frac{F}{S_t} \text{ (with positive risk premium).} \tag{16.20}$$

The binomial process implies that the expected value of the future asset price is

$$\mathrm{E}_t\, S_{t+h} = qSu + (1-q)\, Sd, \tag{16.21}$$

where q is the natural probability of the up state. At the same time, the risk-neutral expected value equals the forrward price (see (16.14))

$$F = pSu + (1-p)\, Sd. \tag{16.22}$$

Use (16.21)–(16.22) in (16.20) to get

$$q > p \text{ (with positive risk premium).} \tag{16.23}$$

One interpretation is that a risk neutral investor would be happy with a lower probability of the up state (and thus a lower expected return), than a risk averse investor.

In contrast, when the risk premium on the underlying asset is zero, then

$$q = p \text{ (with no risk premium).} \tag{16.24}$$

This means that if the underlying asset is priced by *"risk-neutral investors,"* then p equals the true probabilities—suggesting the name "risk-neutral probability."

Example 16.9 *(Natural versus risk neutral probability) With the parameters in Example 16.2*

$$\frac{\mathrm{E}_t\, S_{t+h}}{S} = 0.6 \times \frac{11}{10} + (1-0.6) \times \frac{9.5}{10} = 1.04.$$

With $y = 0$, $F = S = 10$, so $F/S = 1$. In this case, the underlying asset indeed has a positive risk premium—and $q = 0.6$ while $p = 1/3$.

Proof. (of (16.23)) Use (16.21)–(16.22) in (16.20)

$$qu + (1-q)\, d > pu + (1-p)\, d,$$

if the positive risk premium is positive. Subtract d from both sides and notice that $u-d > 0$. ■

16.4 Numerical Applications of the Binomial Model

16.4.1 How to Construct a Tree for the Asset Price

We now discuss how to construct a binomial tree with many small time steps—so that it mimics the behaviour of the asset price process.

The binomial distribution converges to a normal distribution as we chop up a given time to expiration into smaller and smaller time steps—and the normal distribution is fully described by the mean and variance. It is therefore common practice to construct the binomial tree to match the mean and variance of the underlying series.

Suppose the log price of the underlying asset is a random walk with drift

$$\begin{aligned} \ln S_{t+h} - \ln S_t &= \mu h + \varepsilon_{t+h}, \text{ with} \\ \mathrm{E}\,\varepsilon_{t+h} &= 0 \text{ and } \mathrm{Var}(\varepsilon_{t+h}) = \sigma^2 h. \end{aligned} \tag{16.25}$$

This means that data sampled as $t, t+h, t+2h, ...$has the mean and variance

$$\mathrm{E}(\ln S_{t+h} - \ln S_t) = \mu h \tag{16.26}$$

$$\mathrm{Var}(\ln S_{t+h} - \ln S_t) = \sigma^2 h \tag{16.27}$$

For instance, if we measure periods in years, then $h = 1/52$ corresponds to weekly data. Expressing the moments in terms of the sampling interval h helps relating to the binomial model—and to compare results across different sampling intervals.

If we approximate this price process with the binomial model (16.1), then the log price process becomes

$$\ln S_{t+h} - \ln S_t = \begin{cases} \ln u & \text{with probability } q \\ \ln d & \text{with probability } 1-q. \end{cases} \tag{16.28}$$

(Notice that (16.1) says that $S_{t+h}/S_t = u$ with probability q. Just take logs to get the results here.) The binomial process implies that the mean and variance of the asset price change are therefore (see Remark 16.1)

$$\mathrm{E}(\ln S_{t+h} - \ln S_t) = q \ln u + (1-q) \ln d, \tag{16.29}$$

$$\mathrm{Var}(\ln S_{t+h} - \ln S_t) = q(1-q)(\ln u - \ln d)^2. \tag{16.30}$$

There are three parameters (u, d, and q) which can be chosen to match the two moments (mean and variance) in (16.26)–(16.27), so we can make one arbitrary choice. The following is a common approach.

First, for any u and d (not yet decided), pick q to match the mean drift over a time step of size h (which is μh, see (16.25)), that is,

$$q \ln u + (1-q) \ln d = \mu h, \text{ so} \tag{16.31}$$

$$q = \frac{\mu h - \ln d}{\ln u - \ln d}. \tag{16.32}$$

Second, pick u and d to match the variance over a time step of size h (which is $\sigma^2 h$, see (16.25)), that is,

$$q(1-q)(\ln u - \ln d)^2 = \sigma^2 h. \tag{16.33}$$

Use (16.32) to substitute for q and simplify

$$(\mu h - \ln d)(\ln u - \mu h) = \sigma^2 h. \tag{16.34}$$

There are several ways to proceed from here, but the most common is approach of Cox, Ross, and Rubinstein (1979) where

$$u = e^{\sigma\sqrt{h}} \text{ and } d = e^{-\sigma\sqrt{h}}. \tag{16.35}$$

Using this on the left hand side of (16.34) gives

$$(\mu h + \sigma\sqrt{h})(\sigma\sqrt{h} - \mu h) = \sigma^2 h - \mu^2 h^2. \tag{16.36}$$

This clearly does not fit the volatility exactly (compare with the right hand side of (16.34)), but the approximation improves quickly as h decreases (the second order term h^2 vanishes fast). There are other ways to construct the binomial tree, but they have similar properties.

Notice that once we have the values of u and d, the pricing of derivatives does not use the natural probability of the up state (q).

16.4.2 Multiperiod Trees

The binomial model is very useful for numerical calculations of the option price. In such numerical applications, the time to expiry is divided into many small time steps, and it is

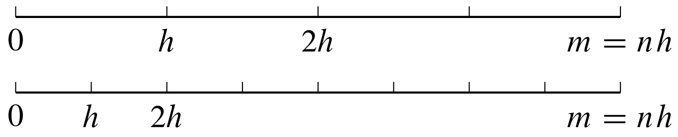

Figure 16.6: Two different time steps with same time to expiration m

assumed that the price of the underlying asset can make an up or down movement in each subinterval—and that the no-arbitrage portfolio is rebalanced every time step. The logic is that the binomial pricing model applies to each up/down branching, so we can combine many such branchings into a tree. Of course, the size of the up and down movements (u and d in the previous analysis), as well as the discounting, is scaled by the number of subintervals.

Let m be the time to expiration of the derivative. With n short time intervals, the length of each interval is $h = m/n$.

Example 16.10 *(Parameters to binomial tree) If the time to maturity of the option is* 20 *(trading) days, then* $m = 20/252$ *(assuming 250 trading days per year). If we want* $n = 100$ *short time intervals (5 per day) in the binomial tree, then* $h = (20/252)/100 \approx 0.00079$*.*

The Cox, Ross, and Rubinstein (1979) approach implies (from using (16.10) and (16.35))

$$
\begin{aligned}
& u = e^{\sigma\sqrt{h}}, d = e^{-\sigma\sqrt{h}}, && (16.37)\\
& p = (e^{yh} - d)/(u - d), \text{ and} \\
& \text{and discounting by } e^{-yh}.
\end{aligned}
$$

Notice that we must keep h small enough so (16.4) holds ($u > e^{yh} > d$, to rule arbitrage opportunities), that is,

$$
e^{\sigma\sqrt{h}} > e^{yh} > e^{-\sigma\sqrt{h}}, \tag{16.38}
$$

which requires $\sqrt{h} < \sigma/y$.

As we start to change the interval length (h), it may no longer correspond to the sampling frequncy of data used in estimating the mean and variance of data in (16.26)–

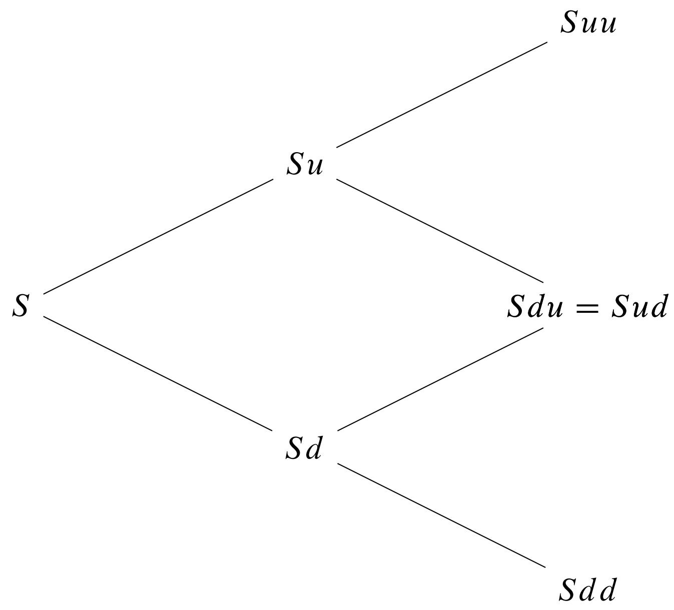

Figure 16.7: Binomial tree for underlying asset ($n = 2$)

(16.27). That is not a problem. Rather, use data to extract the σ^2 (and perhaps μ) and then use that in the binomial tree (16.37). In other words, the h in the sampling frequency and the h in the binomial tree need not be the same. (Just more notational recycling.)

Example 16.11 *(Parameters to binomial tree) Suppose weekly data on a changes in a log stock price has a variance of 0.4, that is $\sigma^2/52 = 0.4$. Then, $\sigma^2 = 20.8$ is the annualised variance. If the binomial tree uses $h = (20/252)/100$, then the $\sigma\sqrt{h}$ used in (16.37) is* $\sqrt{20.8}\sqrt{(20/252)/100} \approx 0.128$.

Figure 16.7 is an illustration of a binomial tree with two subintervals. This tree has only three final nodes, since $Sud = Sdu$—it is "recombining," which is very useful to keep the number of nodes manageable (when we have many time steps). The corresponding prices of the derivative are illustrated in Figure 16.8.

Example 16.12 *(European call option) For a European call option with strike price K and three months (0.25 years) to expiration, the nodes for two steps ($n = 2$, so the length*

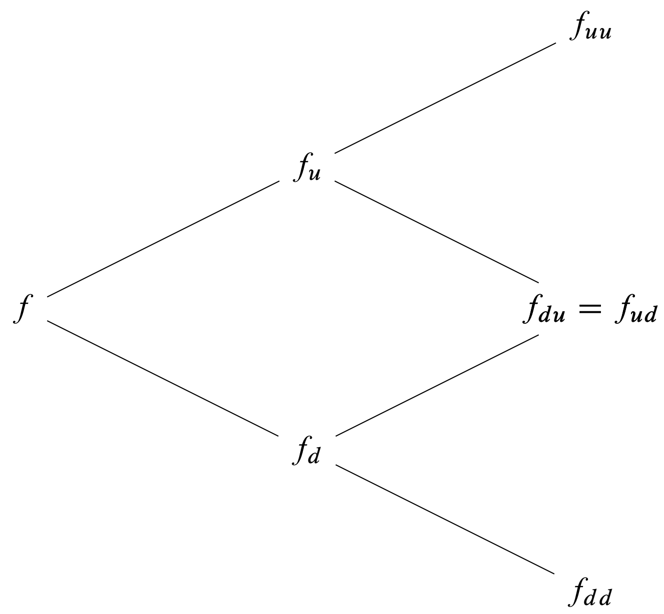

Figure 16.8: Binomial tree for derivative ($n = 2$)

of each time interval is $0.25/2 = 1/8$ *long) in Figure 16.9 are*

$$f = e^{-y/8}[pf_u+(1-p)f_d], \begin{bmatrix} f_u = e^{-y/8}[pf_{uu} + (1-p)f_{ud}] \\ f_d = e^{-y/8}[pf_{du} + (1-p)f_{dd}] \end{bmatrix}, \text{ and } \begin{bmatrix} f_{uu} = \max(Suu - K, 0) \\ f_{ud} = \max(Sud - K, 0) \\ f_{dd} = \max(Sdd - K, 0) \end{bmatrix}$$

where $p = (e^{y/8} - d)/(u - d)$. *Notice that the calculation begins at the end (right) and works backwards towards the start of the tree (left).*

Example 16.13 *(European put option) The tree for a European put option is the same as for a European call option, except for the end nodes. With two steps, as in Figure 16.10, we have*

$$\begin{bmatrix} f_{uu} = \max(K - Suu, 0) \\ f_{ud} = \max(K - Sud, 0) \\ f_{dd} = \max(K - Sdd, 0) \end{bmatrix}$$

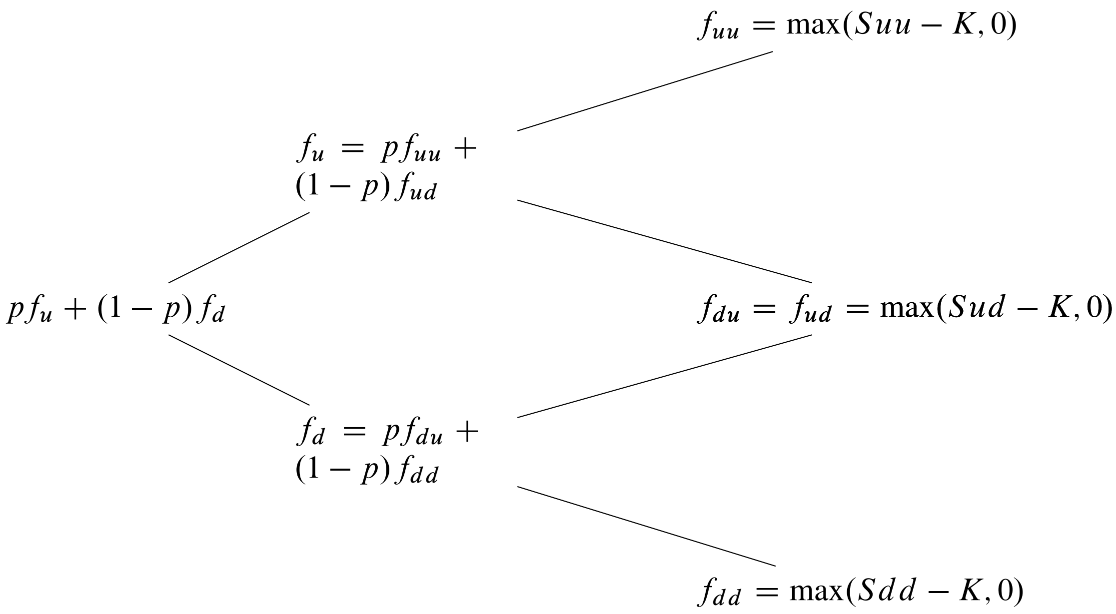

Figure 16.9: Binomial tree for European call option ($n = 2$), zero interest rate

16.4.3 Using a Binomial Tree for Pricing American Options

The binomial tree we have used so far assumes that the derivative is "alive" until the end of the period. This is not necessarily the case for American options, so the approach needs to be modified to handle the possibility of early exercise.

The option value is then the maximum of the exercise value and the value if keeping the option "alive." The latter is defined in the same way as in (16.10). Together this gives the price of the derivative as

$$f = \max(\text{value if exercised now, } e^{-yh}\,[pf_u + (1-p)\,f_d]), \tag{16.39}$$

where p is defined as before (in (16.10)). For instance, a two-step tree for an American put option would have

$$f = \max(K - S, e^{-yh}\,[pf_u + (1-p)\,f_d]), \text{ where} \tag{16.40}$$

$$f_u = \max(K - Su, 0) \text{ and } f_d = \max(K - Sd, 0). \tag{16.41}$$

Example 16.14 *(An American put option) With an American put option with strike price*

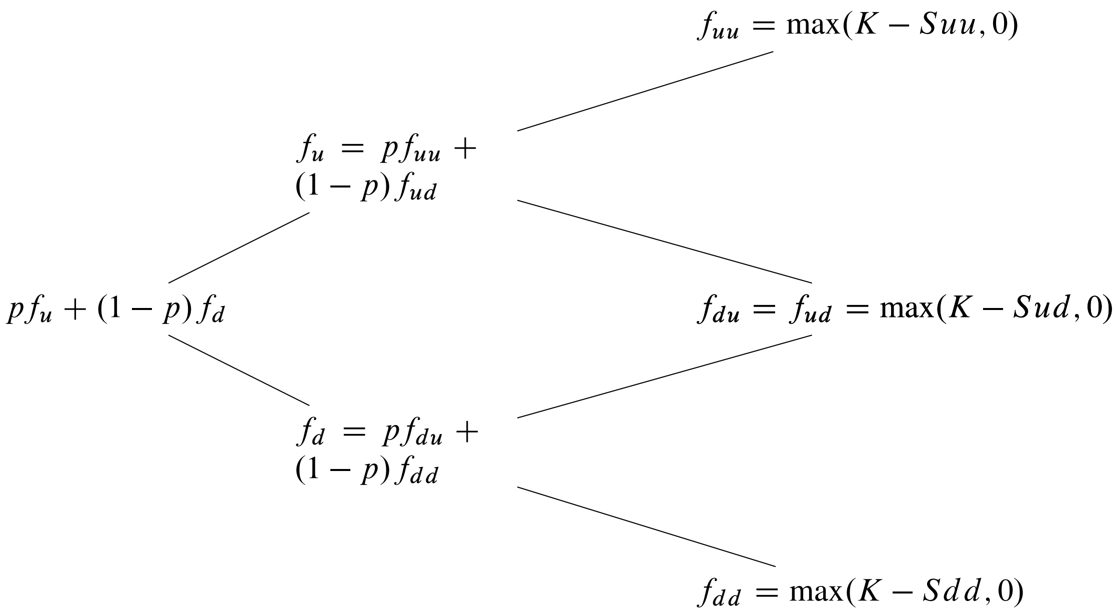

Figure 16.10: Binomial tree for a European put option ($n = 2$), zero interest rate

K and six months (0.5 years) to expiration the nodes for two steps ($n = 2$, so the length of each time interval is $0.5/2 = 1/4$ long) in Figure 16.11, we must account for the possibility of an early exercise. At each node, the option value is the maximum of the value if exercised (K minus the asset price) and the value if kept "alive" (denoted f^a below) The latter is the discounted risk-neutral expected value of the option value next period—just like for a European option. We therefore have

$$f = \max(K - S, f^a), \text{ where } f^a = e^{-y/4}[pf_u + (1-p)f_d]$$

$$f_u = \max(K - Su, f_u^a), \text{ where } f_u^a = e^{-y/4}[pf_{uu} + (1-p)f_{ud}]$$
$$f_d = \max(K - Sd, f_d^a), \text{ where } f_d^a = e^{-y/4}[pf_{du} + (1-p)f_{dd}]$$

$$f_{uu} = \max(K - Suu, 0)$$
$$f_{ud} = \max(K - Sud, 0)$$
$$f_{dd} = \max(K - Sdd, 0),$$

$$f_{uu} = \max(K - Suu, 0)$$

$$f_u = \max(pf_{uu} + (1-p)f_{ud}, K - Su)$$

$$\max(pf_u + (1-p)f_d, K - S)$$

$$f_{du} = f_{ud} = \max(K - Sud, 0)$$

$$f_d = \max(pf_{du} + (1-p)f_{dd}, K - Sd)$$

$$f_{dd} = \max(K - Sdd, 0)$$

Figure 16.11: Binomial tree for an American put option ($n = 2$), zero interest rate

where $p = (e^{y/4} - d)/(u - d)$*. As always, the calculation begins at the end and works backwards down the tree.*

Figure 16.12 illustrates the solution for an American put option on an asset without dividends. Notice that the American put price exceeds the European put price—and more so at low asset prices ($S < K$ is necessary, but not sufficient for early exercise) and high interest rates, that is, when it is likely that the option will be exercised early.

Figure 16.13 illustrates the calculations of the American put price for one current value of the underlying asset. The shaded areas show the location of the nodes (future prices of the underlying asset) that are used in the calculation—and at which nodes that early exercise will happen.

16.4.4 A Binomial Tree with Continuous Dividends*

It is straightforward to construct another tree that allows for continuous dividends.

Suppose dividends are paid at the continuous *rate* δ. Let the up and down movements in the asset price reflect the ex-dividend price, and assume that any dividends are reinvested in the stock.

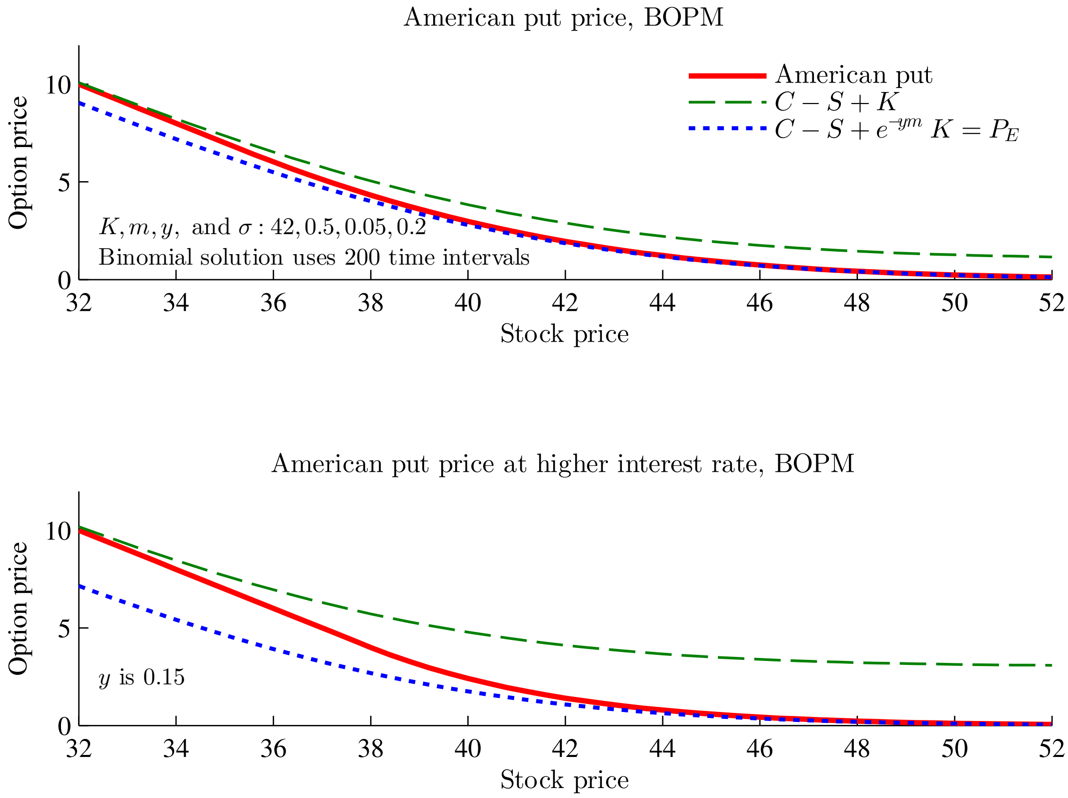

Figure 16.12: Numerical solution of an American put price

First, to construct a riskfree portfolio, hold Δ of the underlying asset and -1 of the derivative. The payoff of the portfolio at expiry is $\Delta S e^{\delta h} u - f_u$ in the "up" state and $\Delta S e^{\delta h} d - f_d$ in the "down" state. The $e^{\delta h}$ factor comes from reinvestment. To make the portfolio riskfree the delta must be

$$\Delta = \frac{f_u - f_d}{S e^{\delta h} (u - d)}. \tag{16.42}$$

Second, to make the return of the portfolio equal to the riskfree rate, we set the present value of our riskfree portfolio equal to the cost of the portfolio

$$e^{-yh} \left(\Delta S e^{\delta h} u - f_u \right) = \Delta S - f. \tag{16.43}$$

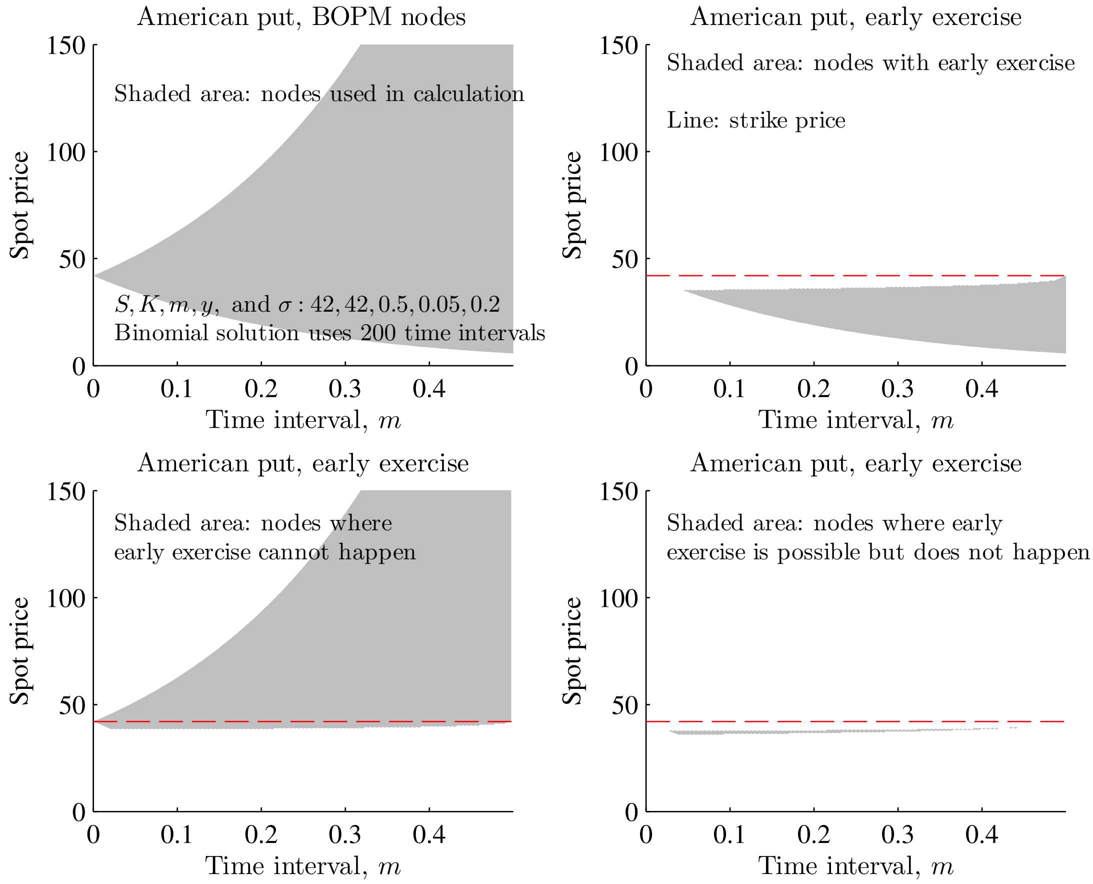

Figure 16.13: Numerical solution of an American put price

Use (16.42) and rearrange as

$$f = \Delta S \left(1 - e^{(\delta - y)h} u\right) + e^{-yh} f_u \tag{16.44}$$

$$= \frac{f_u - f_d}{e^{\delta h}(u - d)} \left(1 - e^{(\delta - y)h} u\right) + e^{-yh} f_u \tag{16.45}$$

$$= e^{-yh} \left[p f_u + (1 - p) f_d\right] \text{ with } p = \frac{e^{(y-\delta)h} - d}{u - d}. \tag{16.46}$$

With this new definition of p, the rest of the computations are as in the case without dividends. In particular, the drift of the asset price does not matter, so u and d can be chosen as before, for instance, as in (16.35).

Remark 16.15 *(Risk neutral drift with continuous dividends) With continuous dividends,* *the risk-neutral expected value is* $\mathrm{E}_t^p\, S_{t+h}/S_t = e^{(y-\delta)h}$*, so the drift is* $(y-\delta)h$ *over the short time interval* h*.*

Bibliography

Cochrane, J. H., 2001, *Asset pricing*, Princeton University Press, Princeton, New Jersey.

Cox, J. C., S. A. Ross, and M. Rubinstein, 1979, "Option pricing: a simplified approach," *Journal of Financial Economics*, 7, 229–263.

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

17 The Black-Scholes Model and the Distribution of Asset Prices

Main references: Elton, Gruber, Brown, and Goetzmann (2010) 23 and Hull (2009) 13 and 17
Additional references: McDonald (2006) 9–13 and Cox, Ross, and Rubinstein (1979)

17.1 The Black-Scholes Model

17.1.1 The Basic Black-Scholes Model without Dividends

Assume that the change over a short interval (between t and $t + h$) in the log asset price is an iid process

$$\ln S_{t+h} - \ln S_t = \mu h + \varepsilon_{t+h}, \text{ with } \varepsilon_{t+h} \sim iid\ N(0, \sigma^2 h). \tag{17.1}$$

(Notice that $\ln S_{t+h} - \ln S_t = \ln S_{t+h}/S_t$, which is also used below.) This implies that, based on the information in period 0, the logarithm of the stock price in period m, S_m, is normally distributed

$$\ln S_m \sim N(\ln S + \mu m, \sigma^2 m), \tag{17.2}$$

where S is the current asset price (the subscript is dropped to reduce clutter). For instance, if there is no volatility ($\sigma^2 = 0$), then $S_m = e^{\mu m} S$. See Figure 17.1 for an illustration.

Proof. (of (17.2)) Notice that (17.1) implies that

$$\begin{aligned}
\ln S_1 &= \ln S_0 + \mu + \varepsilon_1 \text{ and} \\
\ln S_2 &= \ln S_1 + \mu + \varepsilon_2 = (\ln S_0 + \mu + \varepsilon_1) + \mu + \varepsilon_2.
\end{aligned}$$

Since $\mathrm{E}_0\, \varepsilon_1 = 0$ and $\mathrm{E}_0\, \varepsilon_2 = 0$, the conditional means are $\mathrm{E}_0 \ln S_1 = \ln S_0 + \mu$ and $\mathrm{E}_0 \ln S_2 = \ln S_0 + 2\mu$. The conditional variances are just the variances of the forecast errors (the ε part), so $\mathrm{Var}_0(\ln S_1) = \mathrm{Var}(\varepsilon_1) = \sigma^2$ and $\mathrm{Var}_0(\ln S_2) = \mathrm{Var}(\varepsilon_1) + \mathrm{Var}(\varepsilon_2) = 2\sigma^2$. ∎

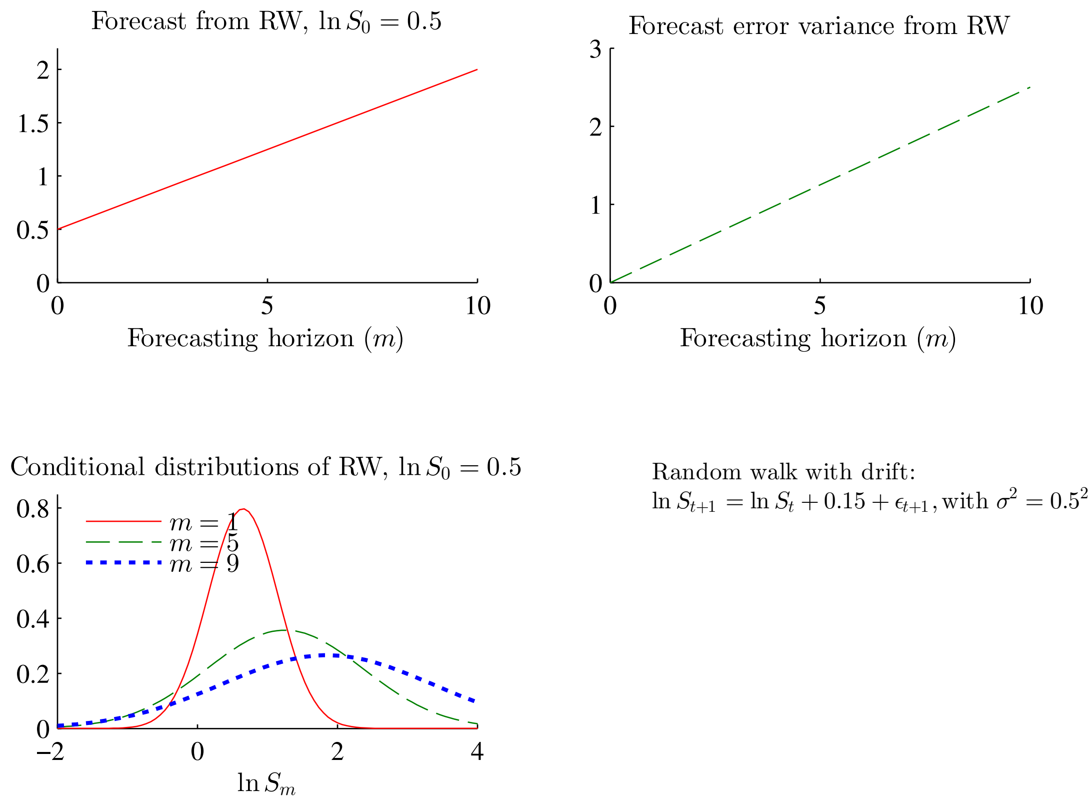

Figure 17.1: Condititional distribution from random walk with drift

If we take the proper limit as the time interval h goes towards zero, then we have a Brownian motion for the log asset price ($d \ln S_t = \mu dt + \sigma d W_t$, where $d W_t$ are the increments to a Wiener process).

A hedging/no arbitrage argument similar to the binomial model then leads to the Black-Scholes formula (for an asset without dividends) where the European call option price is

$$C = S\Phi\left(d_1\right) - e^{-ym} K\Phi\left(d_2\right), \text{ where} \tag{17.3}$$

$$d_1 = \frac{\ln(S/K) + \left(y + \sigma^2/2\right) m}{\sigma\sqrt{m}} \text{ and } d_2 = d_1 - \sigma\sqrt{m}. \tag{17.4}$$

In this formula, $\Phi(d)$ denotes the probability of $x \leq d$ when x has an $N(0, 1)$ distribution (that is, the distribution function value at d).

The call option price is increasing in the asset price, volatility, time to maturity and

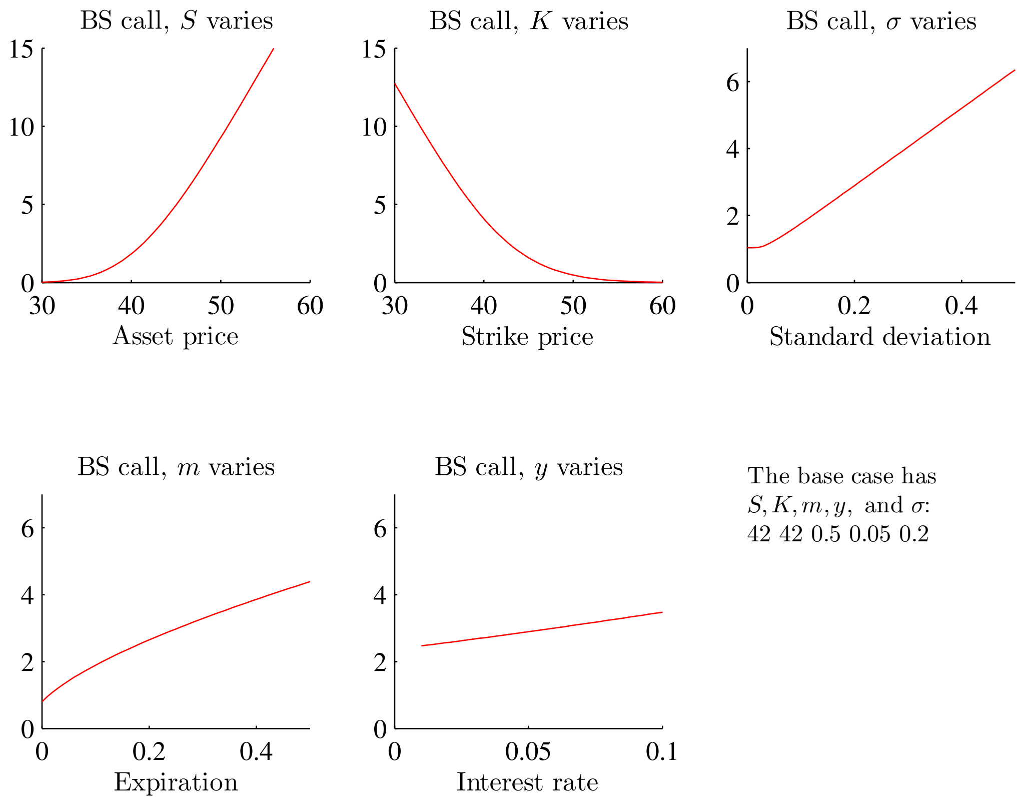

Figure 17.2: Call option price, Black-Scholes model

the interest rate, but decreasing in the strike price. See Figure 17.2.

Remark 17.1 *(*Black-Scholes formula when $\sigma = 0$) From (17.4)* $\lim_{\sigma\to 0} d_1 = \lim_{\sigma\to 0} d_2 = \infty$ *if* $e^{ym}S \geq K$ *and* $-\infty$ *otherwise. Therefore,* $\lim_{\sigma\to 0} \Phi(d_1) = \lim_{\sigma\to 0} \Phi(d_2) = 1$ *if* $e^{ym}S \geq K$ *and* 0 *otherwise. The Black-Scholes call option price at* $\sigma = 0$ *is therefore* $\max(S - e^{-ym}K, 0)$.

Remark 17.2 *(*Call option price when $\sigma = 0$, version 2) When the underlying asset is riskfree ($\sigma = 0$), then its return (denoted μ in (17.2)) must equal the riskfree rate y, so the value of the underlying asset is $e^{ym}S$ at expiration. Consider the following trading strategy. If $\sigma = 0$ and $e^{ym}S \geq K$, then buying a call option gives the underlying asset for sure at the price K (the option will be exercised). The total present value of the contract (including the exercise) is then $C + e^{-ym}K$, where the strike price is discounted since it is paid at expiration. Alternatively, you buy the underlying asset today, at the price S. These* *two ways of acquiring the underlying asset must cost the same, so* $C + e^{-ym}K = S$*, or* $C = S - e^{-ym}K$*. Essentially,* C *is positive since the option allows you to postpone (part of) the payment. In contrast, If* $\sigma = 0$ *and* $e^{ym}S < K$*, then the option will never be exercised, so it is worthless. This is the same as for the Black-Scholes formula. Another perspective is that with* $\sigma = 0$*, then we know that the underlying asset is worth* $e^{ym}S$ *at expiration, so the present value of the known call payoff is* $e^{-ym}\max(e^{ym}S - K, 0)$*, which is still the same as in the previous remark.*

Remark 17.3 *(*Black-Scholes formula when* $m = 0$*) From (17.4)* $\lim_{\sigma\to 0} d_1 = \lim_{\sigma\to 0} d_2 = \infty$ *if* $S \geq K$ *and* $-\infty$ *otherwise. Therefore,* $\lim_{\sigma\to 0} \Phi(d_1) = \lim_{\sigma\to 0} \Phi(d_2) = 1$ *if* $S \geq K$ *and* 0 *otherwise. The Black-Scholes call option price at* $m = 0$ *is therefore max*$(S - K, 0)$.

17.1.2 The Black-Scholes Model with Dividends

Consider a European option on an underlying asset that pays (continuous or discrete) dividends before expiration. Then, the Black-Scholes formula is not correct. It may seem as if dividends would just affect the mean drift in (17.1), and therefore not affect the option price—but this is wrong. The basic reason is that buying the underlying asset now is different from knowing that you will get the asset at the expiration of the option, since you get the dividends if you hold the asset.

To handle this, we could apply the BS formula on a forward contract (expiring on the same day as the option) instead. Let a prepaid forward contract (present value of forward price), worth $e^{-ym}F$, play the role of the underlying asset in (17.1). This gives the BS formula (17.3)–(17.4) but with $e^{-ym}F$ substituted for S

$$C = e^{-ym}F\Phi(d_1) - e^{-ym}K\Phi(d_2), \text{ where} \tag{17.5}$$

$$d_1 = \frac{\ln(F/K) + (\sigma^2/2)m}{\sigma\sqrt{m}} \text{ and } d_2 = d_1 - \sigma\sqrt{m}. \tag{17.6}$$

This is *Black's model* which has many applications.

For instance, for an asset with a continuous dividend rate of δ, the forward-spot parity says $F = Se^{(y-\delta)m}$. In this case (17.5)–(17.6) can also be written

$$C = e^{-\delta m} S\Phi(d_1) - e^{-ym} K\Phi(d_2), \text{ where} \tag{17.7}$$

$$d_1 = \frac{\ln(S/K) + (y - \delta + \sigma^2/2)m}{\sigma\sqrt{m}} \text{ and } d_2 = d_1 - \sigma\sqrt{m}. \tag{17.8}$$

When the asset is a currency (read: foreign money market account) and δ is the foreign interest rate, then this is the "Garman-Kolhagen" formula.

Remark 17.4 *(*Approximation of option price) A Taylor approximation gives that the call option price close to $F = K$ and $\sigma = 0$ is $C \approx e^{-ym} F\sigma\sqrt{m/(2\pi)}$.*

Remark 17.5 *(*Practical hint: code for Black's model with a forward price) Suppose you have a computer code for the BS model (17.3)—(17.4) which takes the inputs (S, K, y, m, σ). To use that code for Black's model (17.5)–(17.6), substitute $(F, 0)$ for (S, y) and multiply the results by e^{-ym}.*

Remark 17.6 *(*Practical hint: code for BS model with continuous dividends) Suppose you have a computer code for the BS model (17.3)—(17.4) which takes the inputs (S, K, y, m, σ). To use that code for Black's model (17.5)–(17.6), substitute $e^{-\delta m}S$ for S.*

Remark 17.7 *(*Practical hint: finding the dividend rate) If you don't know what the dividend rate is, use the forward-spot parity, $F = Se^{(y-\delta)m}$, to calculate it as $\delta = y - \ln(F/S)/m$.*

17.1.3 Implied Volatility: A Measure of Market Uncertainty

The Black-Scholes formula contains only one unknown parameter: the variance $\sigma^2 m$ in the distribution of $\ln S_m$ (see 17.2). With data on the option price, spot and forward prices, the interest rate, and the strike price, we can solve for the standard deviation σ (see from Figure 17.2 that the option price and the volatility have a monotonic relation). The term σ is often called the *implied volatility*—and it is often used as an indicator of market uncertainty about the future asset price, S_m. It can be thought of as an annualized (provided a period is defined as a year) standard deviation. See Figure 17.3 for an empirical example.

Note that we can solve for one implied volatility for each available strike price. If the Black-Scholes formula is correct, that is, if the assumption in (17.1) is correct, then these

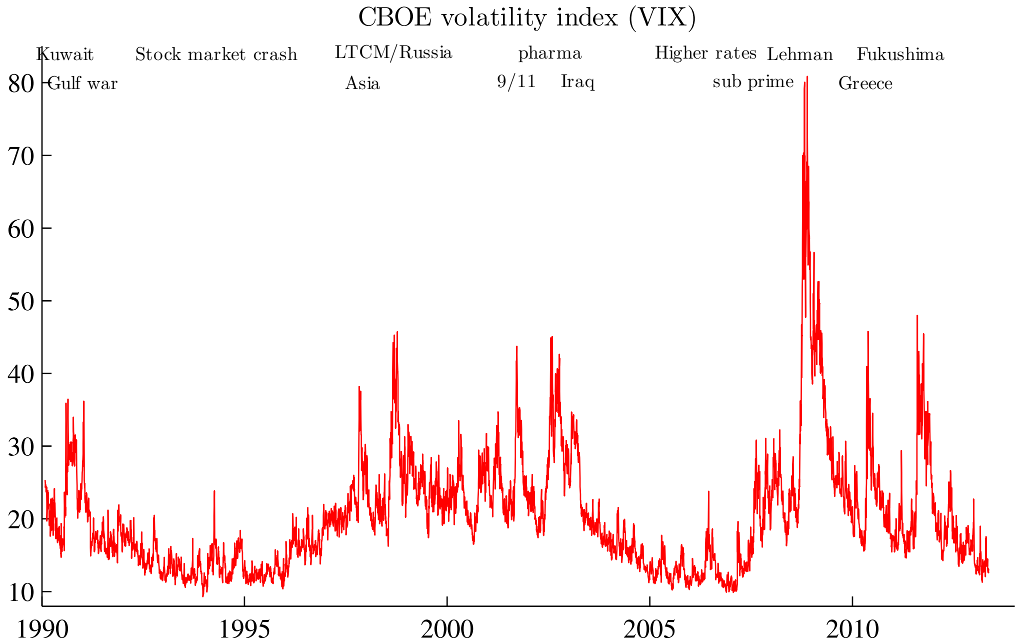

Figure 17.3: CBOE VIX, summary measure of implied volatilities (30 days) on US stock markets

volatilities should be the same across strike prices. On currency markets, we often find a volatility "smile" (volatility is a U-shaped function of the strike price). One possible explanation is that the (perceived) distribution of the future asset price has relatively more probability mass in the tails ("fat tails") than a normal distribution has. On equity markets, we often find a volatility "smirk" instead, where the volatility is very high for very low strike prices. This is often interpreted as that investors are willing to pay a lot for put options that protect them from a dramatic fall in the stock price. One possible explanation is thus that the distribution has more probability mass than a normal distribution at very low stock prices (negative skewness). See Figure 17.4 for an example.

Remark 17.8 *(*Practical hint: starting value for finding σ) From Remark 17.4 we get a starting guess of $\sigma \approx C/[e^{-ym}F\sqrt{m/(2\pi)}]$. Alternatively, it is often recommended to use the starting value $\sigma = \sqrt{|\ln(F/K)|\,2/m}$.*

Remark 17.9 *(*Practical hint: bisection method for finding σ) The bisection method is a very simple (no derivatives are needed) and robust way to solve for the implied volatility. First, start with a lower (σ_L) and higher (σ_H) guess of the yield which are known to*

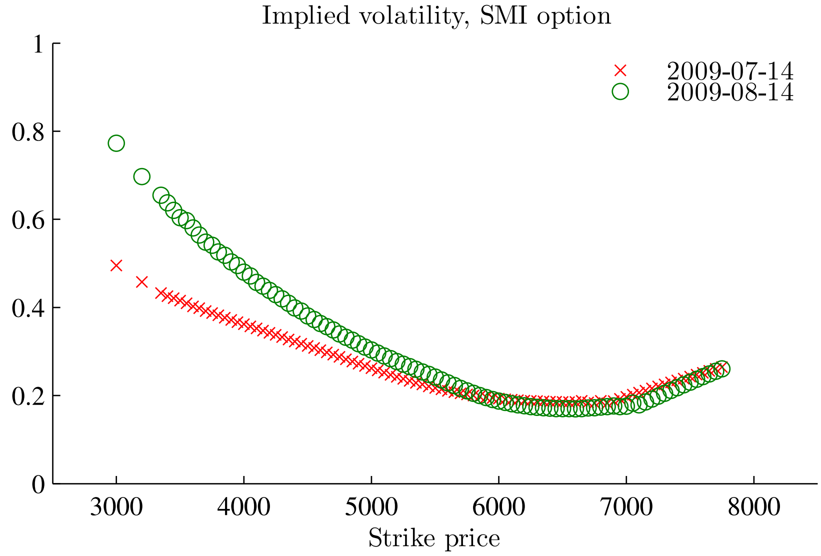

Figure 17.4: Implied volatilities of SMI options, selected dates

bracket the true value, that is, $C(\sigma_L) \leq C \leq C(\sigma_H)$ where C is the observed call option price and $C(\sigma)$ denotes the Black-Scholes formula (as a function of σ). Recall that $C(\sigma)$ is increasing in σ. Second, calculate the option price at the average of the two guesses: $C[(\sigma_L + \sigma_H)/2]$. Third, replace either σ_L or σ_H according to: if $C[(\sigma_L + \sigma_H)/2] \leq C$ (so the midpoint is below the true volatility) then replace σ_L by $(\sigma_L + \sigma_H)/2$ (a higher value), but if $C[(\sigma_L + \sigma_H)/2] > C$ then replace σ_H by $(\sigma_L + \sigma_H)/2$ (a lower value). Fourth, iterate until $\sigma_L \approx \sigma_H$. See Figure 17.6.

17.2 Convergence of the BOPM to Black-Scholes

This section demonstrates that the option prices from the BOPM converges to the prices from the Black-Scholes model. See Figure 17.7 for an illustration.

We know that the risk neutral pricing of a European call option is

$$C = e^{-ym}\,\mathrm{E}^*\max(0, S_m - K), \tag{17.9}$$

where E^* denotes the expectation according to the risk neutral distribution.

For the Black-Scholes model, the (physical) normal distribution for the log asset price

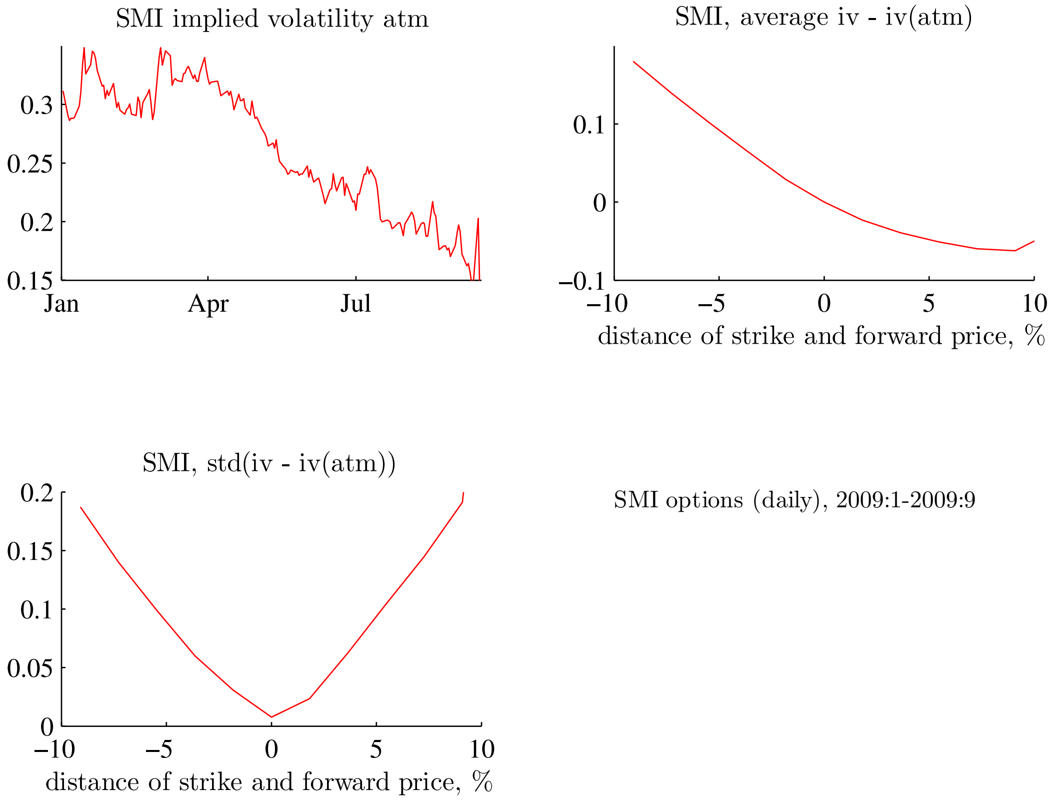

Figure 17.5: Implied volatilities

(17.2) implies that the risk neutral distribution of $\ln S_m$ is

$$\ln S_m \sim^* N(\ln S + ym - \sigma^2 m/2, \sigma^2 m), \tag{17.10}$$

where S is the current asset price. This distribution has a lower mean than the physical distribution, but the same variance.

To see that this is indeed is the risk neutral distribution, recall that the mean of the risk neutral distribution should equal the forward price. In fact, with (17.10) we get

$$\mathrm{E}^* S_m = Se^{ym} = F. \tag{17.11}$$

The second equality comes from the forward-spot parity—and the together this shows that the risk neutral expectation equals the forward price.

Proof. (that (17.11) has a mean equal to the forward rate) Recall that $\mathrm{E}[\exp(x)] =$

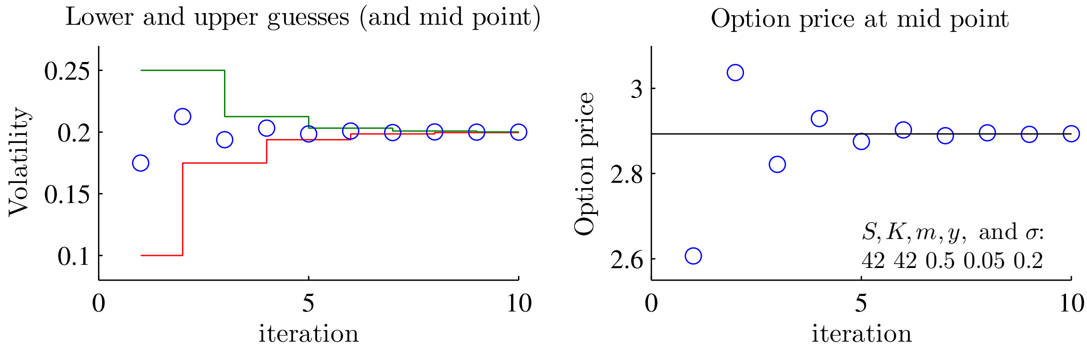

Figure 17.6: Bisection method for finding the implied volatility σ

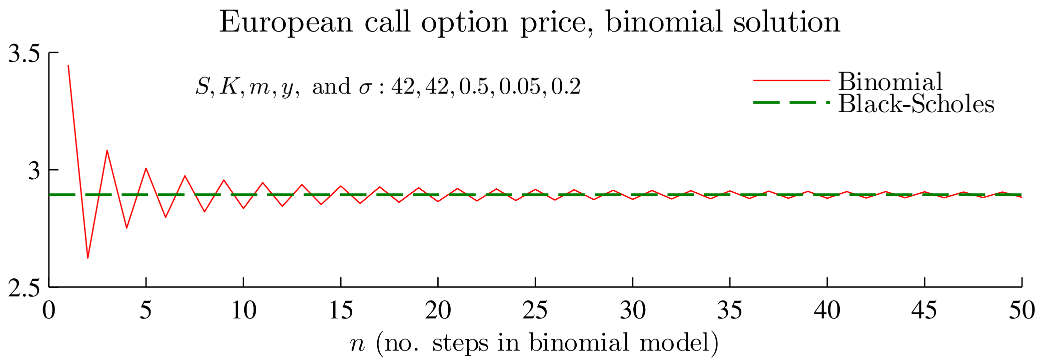

Figure 17.7: Convergence of the binomial price to the Black-Scholes price

$\exp(\mu + s^2/2)$, if $x \sim N(\mu, s^2)$. Applying on the distribution in (17.10) gives $\mathrm{E}^*[\exp(\ln S_m)] = \exp(\ln S + ym) = S^{ym}$, which equals the forward price (by the forward-spot parity). ■

For the binomial option pricing model (BOPM) we have that, in the risk neutral binomial tree for the asset price, implies the following binomial process for the *log asset price*

$$\ln S_{t+h} - \ln S_t = \begin{cases} \ln u & \text{with probability } p \\ \ln d & \text{with probability } 1 - p. \end{cases} \tag{17.12}$$

(In the risk neutral binomial tree $S_{t+h} = S_t u$ with probability p and $S_t d$ with probability $1 - p$. This implies (17.12) for the logs.) The parameters u, d and p all depend on the time step length h in such a way that we match the mean and variance of the price series. In fact, if they are chosen so that the mean and variance of $\ln S_{t+h} - \ln S_t$ are (at least in the limit) proportional to h.

Clearly, the binomial tree means that we reach $\ln S_m$ by adding n steps of the kind in (17.12). To save clutter, assume the current time period is zero. Then,

$$\ln S_m = \ln S_0 + \sum\nolimits_{i=1}^{n} [\ln S_{ih} - \ln S_{(i-1)h}]. \tag{17.13}$$

I demonstrate the convergence in two steps: first, that the binomial distribution converges to a normal distribution; and second that both distributions have the same mean and variance in the limit.

17.2.1 The Central Limit Theorem at Work

If we can show that the risk neutral distribution implied by the binomial model converges (as the number of time steps increase, keeping time to expiration constant) to a normal distribution, then it is plausible that the Black-Scholes model can be thought of as the limit of the binomial model.

The Black-Scholes model is based on normally distributed changes of log prices. In the binomial model, the log price changes can only take two values, but the sum of many such changes will converge to a normally distributed variable as the number of time steps increases—providing the step size decreases. This may seem counter intuitive since central limit theorems apply to samples averages (times the square root of the sample size), not to sums. However, the rescaling of the log price changes as the number of time steps increases, means that the sum is effectively a (scaled) sample average—so a CLT indeed applies.

See Figure 17.8 for an example of how the distribution converges.

Proposition 17.10 *If u, d and p in the binomial process (17.12) are such that the mean and variance of $\ln S_{t+h} - \ln S_t$ are proportional to h, then the distribution converges to a normal distribution as the number of time steps n increases, keeping the maturity m constant (so $h = m/n$).*

Remark 17.11 *(The Lindeberg-Lévy central limit theorem) If x_i is independently and identically distributed with $\mathrm{E}\, x = 0$ and $\mathrm{Var}(x_i) = \sigma < \infty$, then,*

$$\sqrt{n}\frac{1}{n}\sum\nolimits_{i=1}^{n} x_i = \sum\nolimits_{i=1}^{n} x_i/\sqrt{n} \xrightarrow{d} N(0, \sigma^2).$$

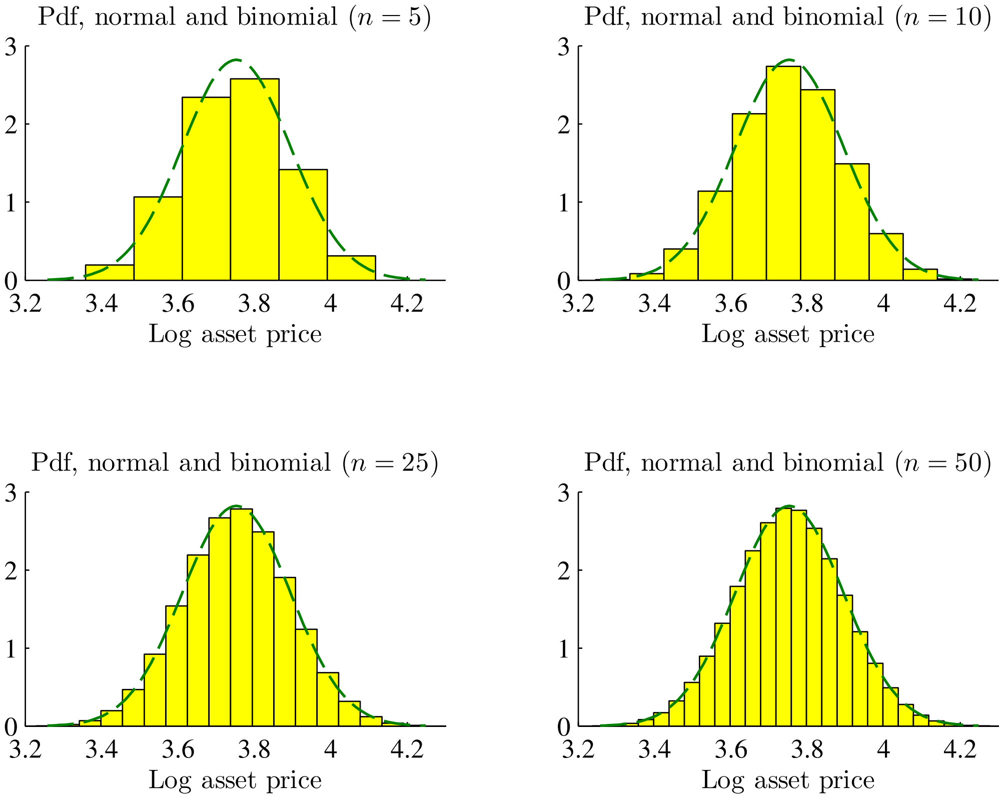

Figure 17.8: Convergence of the binomial model to the Black-Scholes model

Proof. (of Proposition 17.10) The binomial model (17.12)–(17.13) means that we can write the log price change over a time step of length as $\varepsilon_i\sqrt{h} + \mu h$, where ε_i is an iid zero mean random variable with variance σ^2

$$\varepsilon_i\sqrt{h} + \mu h = \ln S_{ih} - S_{(i-1)h}, \text{ where } \mathrm{E}\,\varepsilon_i = 0 \text{ and } \mathrm{Var}(\varepsilon_i) = \sigma^2.$$

Clearly, $\mathrm{Var}(\varepsilon_i\sqrt{h})$ is proportional to h and the means are all zero, as required. For instance, for the first time interval we have $\varepsilon_1\sqrt{h} + \mu h = \ln S_h - \ln S_0$. Since we take n steps (of length $h = m/n$) to get from from period 0 to m, we can write the change in the log price (from 0 to m) as the sum of $\sqrt{h}\varepsilon_i + \mu h$ from $i = 1$ to n

$$\begin{aligned}\ln S_m - \ln S_0 &= \sqrt{h}\textstyle\sum_{i=1}^{n}\varepsilon_i + nh\mu \\ &= \sqrt{m}\textstyle\sum_{i=1}^{n}\varepsilon_i/\sqrt{n} + \mu m,\end{aligned}$$

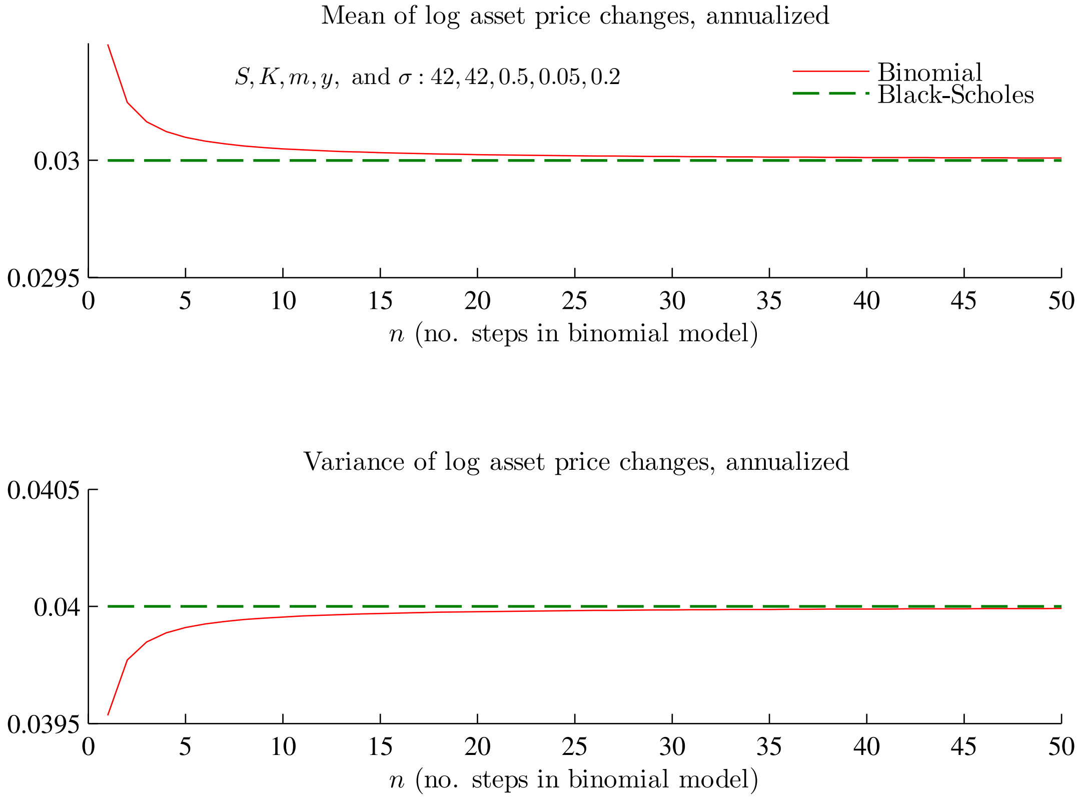

Figure 17.9: Convergence of the binomial mean and variance

where we have used $h = m/n$. The first term on the right hand (except the constant m) is of the same form as in the central limit theorem in Remark 17.11. The second term is just a constant. ■

17.2.2 Convergence of the Mean and Variance

This section demonstrates that the mean and variance of the binomial distribution converges to the same values as in the risk neutral distribution of the Black-Scholes model (17.10). See Figure 17.9 for an illustration.

Proposition 17.12 *(Moments of CRR steps) In the Cox, Ross, and Rubinstein (1979) tree, the parameters in (17.12) are*

$$\ln u = \sigma\sqrt{h}, \ln d = -\sigma\sqrt{h} \textit{ and } p = (e^{yh} - d)/(u - d).$$

As $n \to \infty$, but $h = m/n$ we have (since the price changes are independent) the following results for the sum of them

$$\mathrm{E} \ln S_m - \ln S_0 = m(y - \sigma^2/2) \textit{ and } \mathrm{Var}(\ln S_m - \ln S_0) = m\sigma^2.$$

This the same as in the risk neutral distribution of the Black-Scholes model.

Proof. (of Proposition 17.12) Recall that the mean and variance of $\ln S_{t+h} - \ln S_t$ are $p \ln u + (1-p) \ln d$ and $p(1-p)(\ln u - \ln d)^2$ respectively. Since the terms in (17.13) are uncorrelated, the mean and the variance of the sum are $n\,\mathrm{E}(\ln S_{t+h} - \ln S_t)$ and $n\,\mathrm{Var}(\ln S_{t+h} - \ln S_t)$. Substitute for u, d and p and take the limits of as $n \to \infty$, but $h = m/n$. (This is straightforward, but slightly messy, calculus.) ■

17.3 The Probabilities in the BOPM and Black-Scholes Model*

The price of a European (call or put) option calculated by the binomial model converges to the Black-Scholes price as the number of subintervals increases (keeping the time to expiration constant, so the subintervals become shorter). This is illustrated in Figure 17.7.

Both the binomial option pricing model (BOPM) and the Black-Scholes model imply that the call option price can be written as the discounted risk neutral expected payoff (17.9), which we can write as

$$C = e^{-ym} \int_K^\infty (S_m - K)\, f^*(S_m) dS_m, \tag{17.14}$$

where $f^*(S_m)$ is the risk neutral density function of the asset price at expiration (S_m). We can clearly rewrite this expression as

$$\begin{aligned} C &= e^{-ym}\,\mathrm{E}^*(S_m - K | S_m > K) \Pr{}^*(S_m > K) & (17.15)\\ &= e^{-ym}\,\mathrm{E}^*(S_m | S_m > K) \Pr{}^*(S_m > K) - e^{-ym} K \Pr{}^*(S_m > K). & (17.16)\end{aligned}$$

The first term is (the present value of) the expected asset price conditional on exercise, times the probability of exercise. The second term is (the present value of) the strike price times the probability of exercise.

The discussion below demonstrates that these probabilities are the same (in the limit) in the BOPM and the Black-Scholes models.

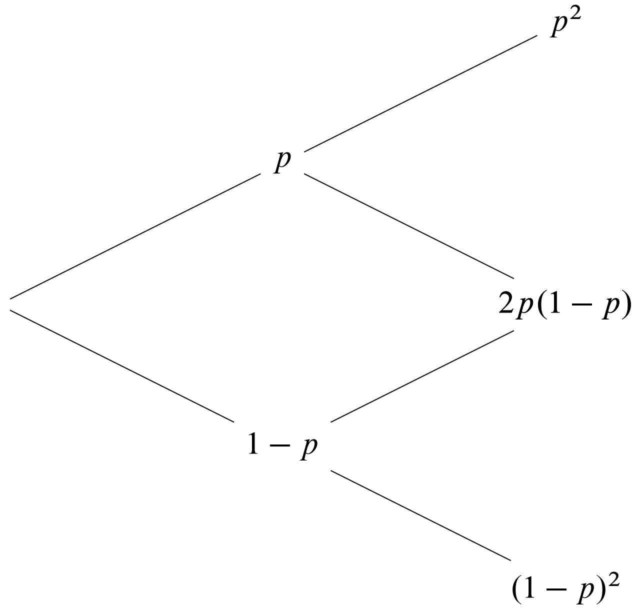

Figure 17.10: Probabilities of different nodes in a binomial tree

17.3.1 The Probabilities in the Binomial Tree*

To understand the binomial model a bit better, consider a binomial tree with 2 subintervals ($n = 2$) of length h as illustrated in Figures 17.10–17.11.

The price of the call option is the discounted risk neutral expected value of the value in the next period

$$C = e^{-yh}[pC_u + (1-p)C_d], \begin{bmatrix} C_u = e^{-yh}[pC_{uu} + (1-p)C_{ud}] \\ C_d = e^{-yh}[pC_{du} + (1-p)C_{dd}] \end{bmatrix}, \text{ and } \begin{bmatrix} C_{uu} = \max(Suu - K, 0) \\ C_{ud} = \max(Sud - K, 0) \\ C_{dd} = \max(Sdd - K, 0) \end{bmatrix} \quad (17.17)$$

where $p = (e^{y/h} - d)/(u - d)$.

Remark 17.13 *(Probabilities for the final nodes) With two trials ($n = 2$), the probabili-*

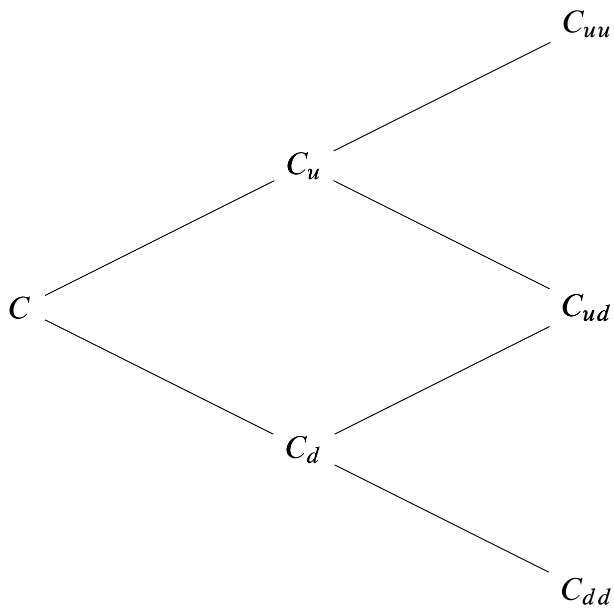

Figure 17.11: Binomial tree for derivative ($n = 2$)

ties for the final nodes are

$$\begin{aligned}
\Pr(uu) &= p^2 \\
\Pr(ud) &= 2p(1-p) \\
\Pr(dd) &= (1-p)^2.
\end{aligned}$$

Combining (and using $2h = m$)

$$C = e^{-ym}\left[p^2 \max(Suu - K, 0) + 2p(1-p)\max(Sud - K, 0) + (1-p)^2 \max(Sdd - K, 0)\right], \tag{17.18}$$

which expresses the call option price as the discounted risk-neutral expectation of the option payoff.

Suppose only $Suu > K$, that is, it is only at the up and up branch, uu, that we exercise. Then

$$\begin{aligned}
C &= e^{-ym} p^2 \left(Suu - K\right) \\
&= e^{-ym} \underbrace{Suu}_{\mathrm{E}^p(S_m|S_m>K)} \underbrace{p^2}_{\Pr^p(uu)} - e^{-ym} K \underbrace{p^2}_{\Pr^p(uu)}.
\end{aligned} \tag{17.19}$$

The first term is the (discounted value of) the risk-neutral expected value of the asset price, conditional on being so high that we exercise the call option, times the risk neutral probability of that event. The second term is the (discounted value of) the strike price times the risk neutral probability of exercise. This clearly has the same form as (17.16). This extends to n steps, except that the expressions for the probabilities are more complicated.

Remark 17.14 *(Bernoulli and binomial distributions) The random variable X can only take two values: 1 or 0, with probability p and $1-p$ respectively. This gives* $\mathrm{E}(X) = p$ *and* $\mathrm{Var}(X) = p(1-p)$*. After n independent trials, the number of successes (y) has the binomial pdf, $n!/[y!(n-y)!]p^y(1-p)^{n-y}$ for $y = 0, 1, ..., n$. This gives* $\mathrm{E}(Y) = np$ *and* $\mathrm{Var}(Y) = np(1-p)$*. To find the probability of at least z successes, sum the pdf over $y = z, z+1, z+2, \ldots$*

17.3.2 The Probabilities in the Black-Scholes Model*

The following remark is useful for the proofs further on.

Remark 17.15 *(Properties of a lognormal distribution) Let $x \sim N(\mu, s^2)$ and define $k_0 = (\ln K - \mu)/s$. First,* $\Pr[\exp(x) > K] = \Phi(-k_0)$*. Second,* $\mathrm{E}[\exp(x)|\exp(x) > K] = \exp\left(\mu + s^2/2\right)\Phi(s - k_0)/\Phi(-k_0)$*. (To prove this, just integrate.)*

Proposition 17.16 *(Riskneutral probability of $S_m > K$) The $\Phi(d_2)$ term in the Black-Scholes formula (17.3)–(17.4) is the risk-neutral probability that $S_m > K$.*

Proposition 17.17 *($S\Phi(d_1)$ in Black-Scholes) The $S\Phi(d_1)$ term in the Black-Scholes formula (17.3)–(17.4) is (the present value of) the expected asset price conditional on exercise, times the probability of exercise, that is, the first term in (17.16).*

Proof. (of Proposition 17.16) The risk neutral probability of $\ln S_m$ is $N(\ln S + ym - \sigma^2 m/2, \sigma^2 m)$. To calculate the probability $\Pr[S_m > K] = \Phi(-k_0)$, notice that k_0 is

$$k_0 = \frac{\ln K - \overbrace{\left(\ln S + ym - \sigma^2 m/2\right)}^{\text{mean}}}{\underbrace{\sigma\sqrt{m}}_{\text{std}}}.$$

Clearly, $-k_0$ is then the same as the argument d_2 in (17.4)

$$d_2 = \frac{\ln(S/K) + \left(y - \sigma^2/2\right) m}{\sigma\sqrt{m}}.$$

■

Proof. (of Proposition 17.17) First, the first term in (17.16) can be written

$$FirstTerm = e^{-ym} \exp\left(\mu + s^2/2\right) \Phi(s - k_0),$$

since the two $\Phi(-k_0)$ terms cancel. Clearly,

$$\begin{aligned} \mu + s^2/2 &= \ln S + ym, \\ s - k_0 &= \sigma\sqrt{m} - \frac{\ln K - \left(\ln S + ym - \sigma^2 m/2\right)}{\sigma\sqrt{m}} = d_1, \end{aligned}$$

where the last line follows from comparing with (17.4). We can therefore write $FirstTerm$ as $S\Phi(d_1)$, since the $e^{-ym}e^{ym}$ term cancels. This is the same as in the Black-Scholes formula. ■

17.4 Hedging an Option

This section discusses how we can hedge a European call option. The setting might be that we have written such an option, but we do not want to carry the risk.

The derivatives in the following remark will be useful.

Remark 17.18 *(The "Greeks") The derivatives of the Black-Scholes formula for an asset*

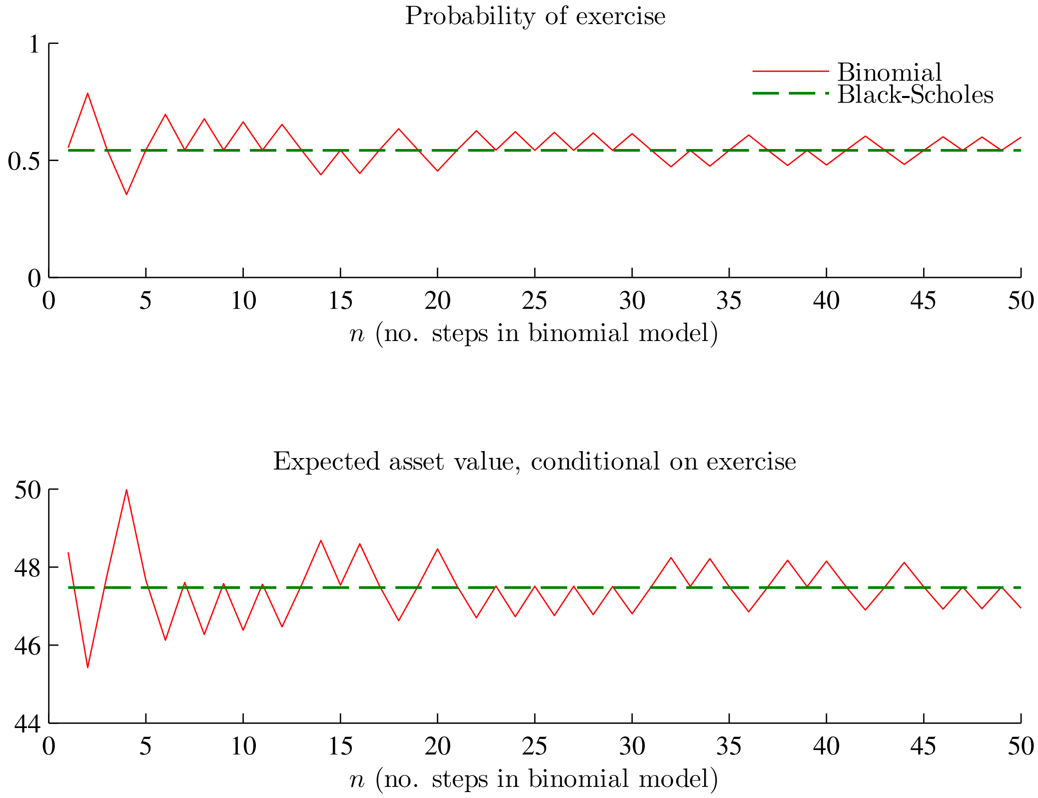

Figure 17.12: Convergence of the binomial model to the Black-Scholes model

with continuos dividends (17.7)–(17.8) are

$$\begin{aligned}
\Delta &= \frac{\partial C}{\partial S} = e^{-\delta m}\Phi\left(d_1\right) \\
\Gamma &= \frac{\partial^2 C}{\partial S^2} = \frac{e^{-\delta m}\phi\left(d_1\right)}{S\sigma\sqrt{m}} \\
\theta &= \frac{\partial C}{\partial t} = -\frac{\partial C}{\partial m} = \delta S^{-\delta m}\Phi(d_1) - yKe^{-ym}\Phi(d_2) - \frac{1}{2\sqrt{m}}e^{-\delta m}S\phi(d_1)\sigma \\
vega &= \frac{\partial C}{\partial \sigma} = Se^{-\delta m}\phi\left(d_1\right)\sqrt{m} \\
\rho &= \frac{\partial C}{\partial y} = mKe^{-ym}\Phi(d_2),
\end{aligned}$$

where $\phi()$ is the standard normal probability density function (the derivative of $\Phi()$). See Figures 17.13–17.14.

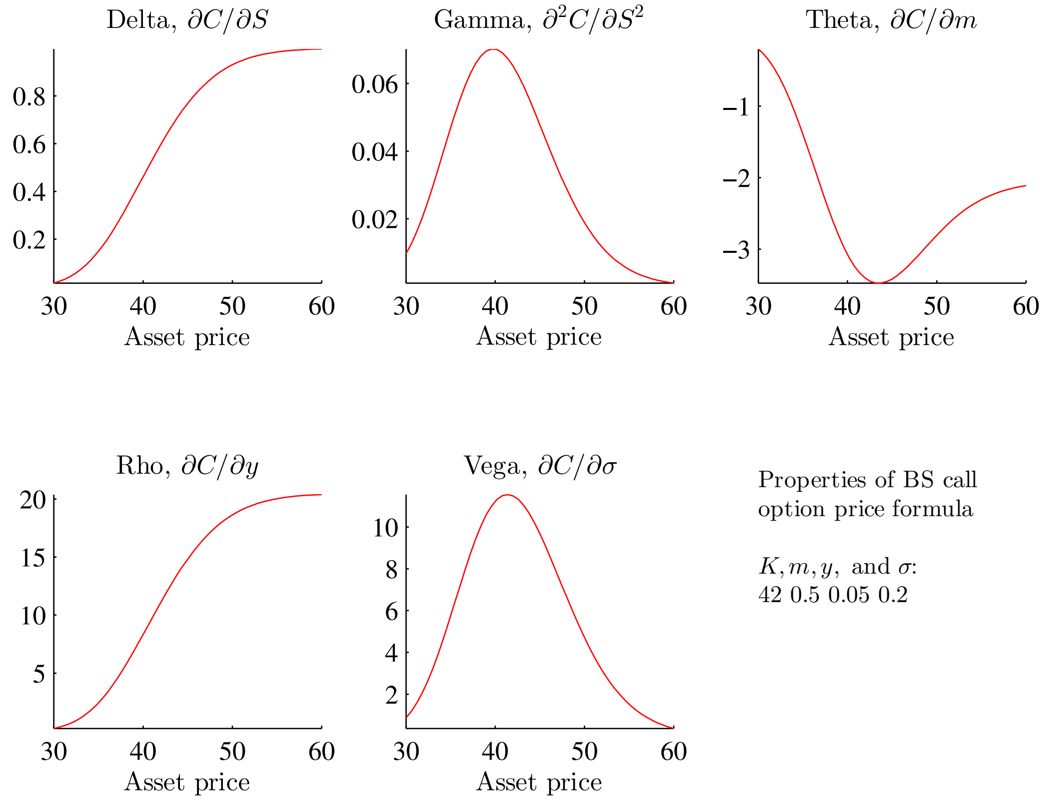

Figure 17.13: The Greeks in the Black-Scholes model as a function of asset price

17.4.1 Delta Hedging

Consider a portfolio with Δ_t of the underlying asset (the hedging portfolio) and short one call option. The value of the overall position is

$$V_t = \Delta_t S_t - C_t. \tag{17.20}$$

Assume that only the price of the underlying asset can change (clearly not true, but at least a starting point for the analysis). A first-order Taylor approximation of the call option price is

$$C_{t+h} - C_t \approx \frac{\partial C_t}{\partial S} (S_{t+h} - S_t). \tag{17.21}$$

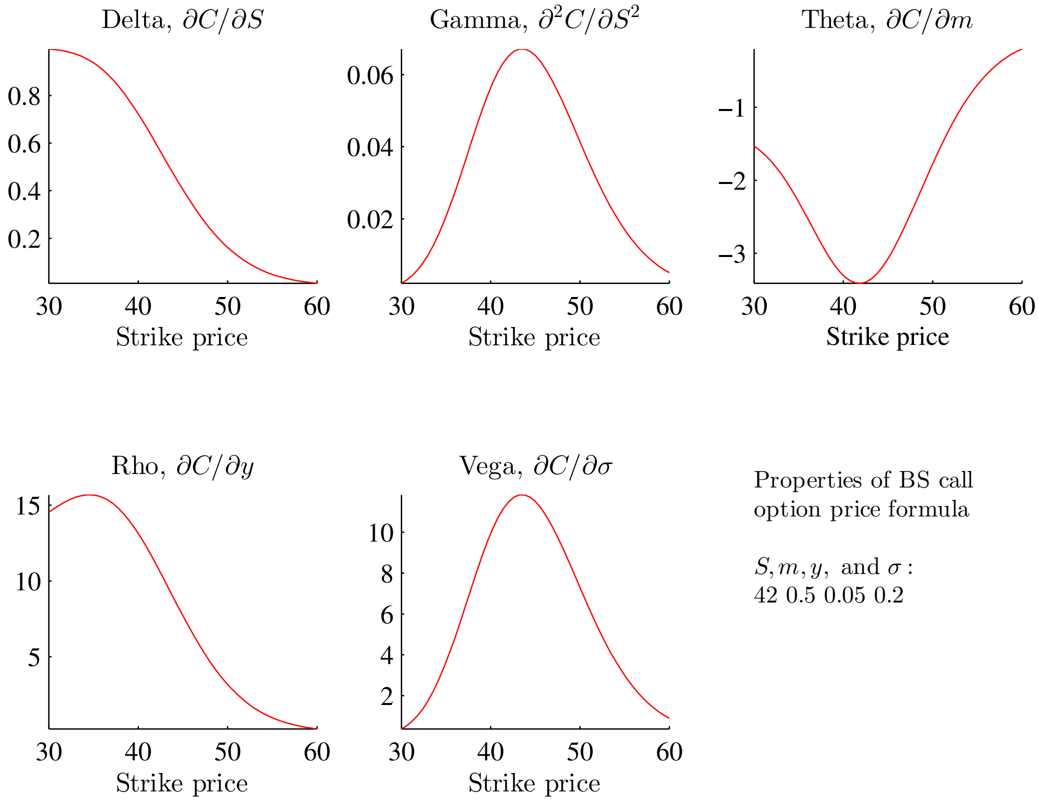

Figure 17.14: The Greeks in the Black-Scholes model as a function of strike price

Use (17.21) to approximate the change of the value of the overall portfolio as

$$\begin{aligned} V_{t+h} - V_t &= \Delta_t (S_{t+h} - S_t) - C_{t+h} - C_t \\ &\approx \Delta_t (S_{t+h} - S_t) - \frac{\partial C_t}{\partial S} (S_{t+h} - S_t) \\ &\approx 0 \text{ if } \Delta_t = \frac{\partial C_t}{\partial S}. \end{aligned} \tag{17.22}$$

This is a *delta hedge*. Clearly, the delta is likely to change from period to period, so the portfolio needs to be frequently rebalanced.

In practice, the hedging portfolio (also called the replicating portfolio) also includes a negative position in a short-term money market account ($M_t < 0$)—so the overall portfolio has a zero (initial) value

$$0 = \underbrace{\Delta_t S_t + M_t}_{\text{hedging portfolio}} - C_t, \text{ so } M_t = -\Delta_t S_t + C_t. \tag{17.23}$$

This means that we finance the purchase of the underlying asset with the proceeds from the selling the option and from borrowing. See Remark 17.19 for details on how the portfolio value changes over time.

Remark 17.19 *(Overall portfolio value over several subperiods*) Start by creating a hedge portfolio with a zero initial value as in (17.23). In $t + h$ (say, after one day so $h = 1/365$), this portfolio is worth (this is the marking-to-market)*

$$V_{t+h} = \Delta_t(D_{t+h} + S_{t+h}) + M_t e^{y_t h} - C_{t+h},$$

where the underlying pays a dividend ($D_{t+h} = 0$ if no dividends), the prices are measured after dividends and y_t is the interest rate. In $t + h$ we need Δ_{t+h} units of the underlying asset (value $\Delta_{t+h} S_{t+h}$). Since we already own Δ_t of the underlying asset, this means that we must withdraw an additional $(\Delta_{t+h} - \Delta_t)S_{t+h}$ from the money market account. On that account, we have since last period $M_t e^{y_t h} + \Delta_t D_{t+h}$ (old holdings with interest plus the dividends we received in cash), so our holdings in $t + h$ (after having rebalanced the holdings) is

$$M_{t+h} = M_t e^{y_t h} + \Delta_t D_{t+h} - (\Delta_{t+h} - \Delta_t)S_{t+h}.$$

The value of the overall portfolio in $t + 2h$ (marking-to-market) is computed as in the first equation, but with subscripts advanced one period. See Figure 17.16 for an illustration. In that figure, "m-to-m" stands for the marking-to-market stage (first equation in this remark) and "rebalancing" for the stage after rebalancing the portfolio (second equation in this remark).

In the Black-Scholes model for an asset with dividends, the delta is

$$\Delta = \frac{\partial C}{\partial S} = e^{-\delta m} \Phi(d_1), \tag{17.24}$$

where d_1 is given by (17.8), see Remark 17.18. Without dividends, just set $\delta = 0$. From

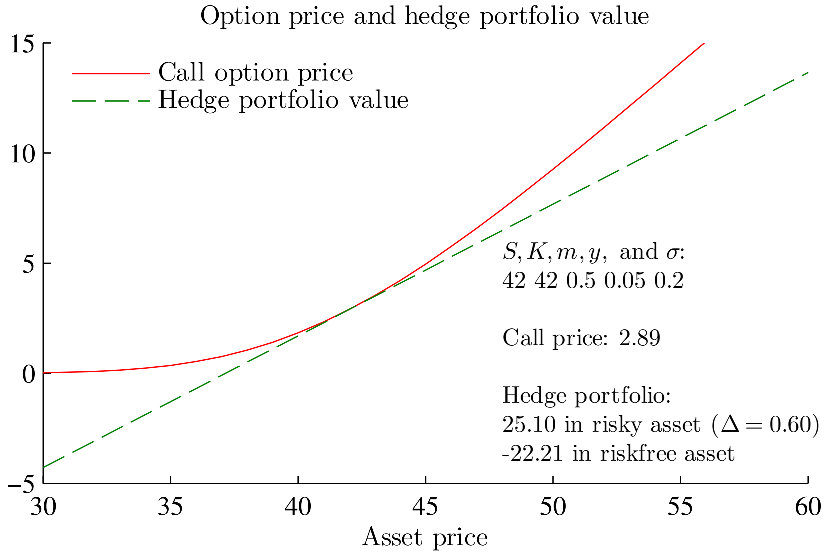

Figure 17.15: Delta hedging as an approximation of the option price

the put-call parity, it is clear that the delta of a put option is

$$\frac{\partial P}{\partial S} = e^{-\delta m}[\Phi(d_1) - 1] = -e^{-\delta m}\Phi(-d_1), \tag{17.25}$$

which is negative. The second equality follows from the symmetry of the normal distribution.)

See Figures 17.13–17.14 for an illustration of how Δ (and other derivatives) depend on the strike and underlying price. In particular, notice that $0 \leq \Delta \leq 1$ and that Δ increasing in the price of the underlying asset. Intuitively, an option that is deep out of the money will not be very sensitive to the asset price—since the chance of exercising is so low. Conversely, an option that is deep in the money moves almost in tandem with the asset price, since it will almost for sure be exercised.

See Figure 17.15 for how the hedging portfolio approximates the call option price (as well for numbers for the positions in the hedge). Also, see Figures 17.17–17.18 for an example of how a delta hedge works on real data.

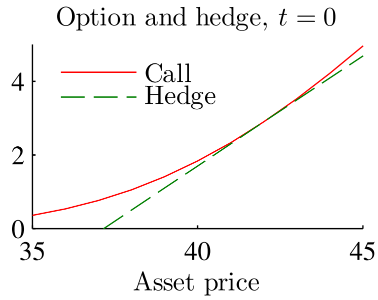

$S, K, m, y,$ and σ:
42 42 0.5 0.05 0.2

	S_0	Δ_0	$-C_0$	$\Delta_0 S_0$	M_0	V_0
Day 0	42.00	0.60	-2.89	25.10	-22.21	0.00
	S_1	Δ_0	$-C_1$	$\Delta_0 S_1$	$M_0 e^{yh}$	V_1
Day 1 (m-to-m)	43.00	0.60	-3.51	25.70	-22.22	-0.02
	S_1	Δ_1	$-C_1$	$\Delta_1 S_1$	M_1	V_1
Day 1 (rebalance)	43.00	0.66	-3.51	28.40	-24.91	-0.02

Figure 17.16: Delta hedging over time

17.4.2 Delta-Gamma Hedging*

Delta hedging can be imprecise if the price of the underlying asset changes much. A second-order Taylor approximation of the option price gives

$$C_{t+h} - C_t \approx \Delta_t (S_{t+h} - S_t) + \frac{1}{2}\Gamma_t (S_{t+h} - S_t)^2, \text{ where } \Delta_t = \frac{\partial C_t}{\partial S} \text{ and } \Gamma_t = \frac{\partial^2 C_t}{\partial S^2}. \tag{17.26}$$

The Δ and Γ of the Black-Scholes model are given in Remark 17.18, see Figures 17.13–17.14 for illustrations.

This movement can be hedged by holding v_t of the underlying asset and w_t of other option. Let Δ_t^* and Γ_t^* be the delta and gamma of this other option. A second-order Taylor approximation of the value of this portfolio (denoted U_t) is

$$U_{t+h} - U_t \approx v_t (S_{t+h} - S_t) + w_t \Delta_t^* (S_{t+h} - S_t) + \frac{1}{2} w_t \Gamma_t^* (S_{t+h} - S_t)^2. \tag{17.27}$$

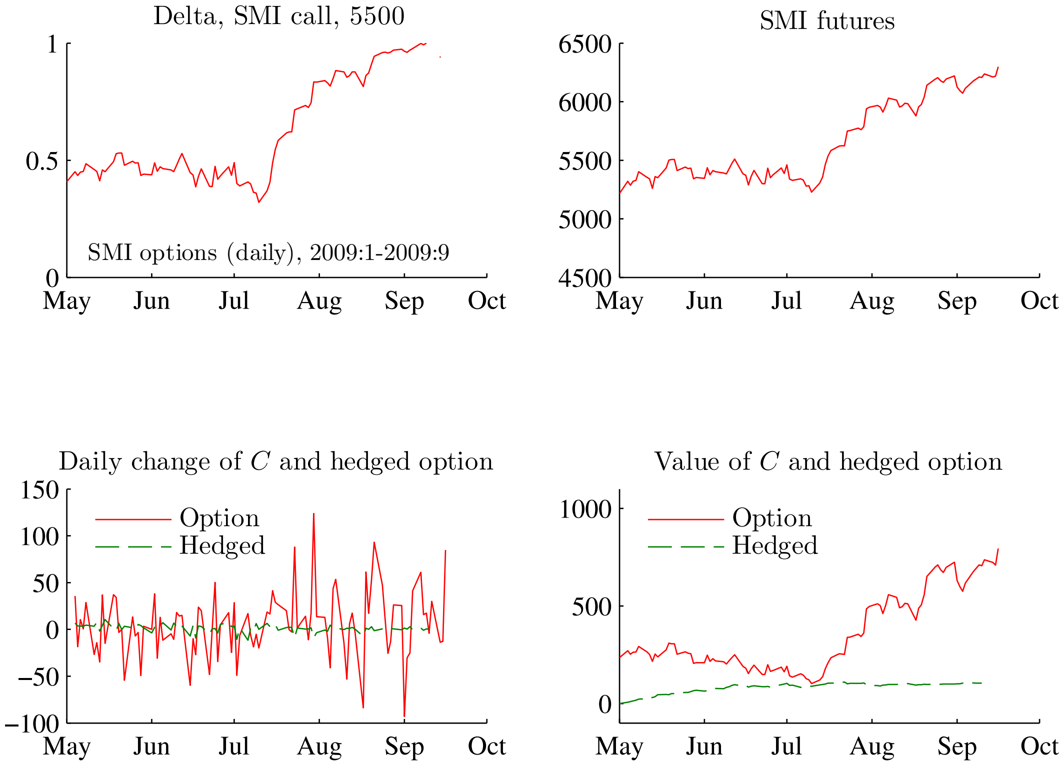

Figure 17.17: Delta hedging an SMI call option

Subtracting (17.26) from (17.27)

$$(U_{t+h}-U_t)-(C_{t+h}-C_t)\approx\underbrace{\left(v_t+w_t\Delta_t^*-\Delta_t\right)}_{A_t}(S_{t+h}-S_t)+\underbrace{\left(w_t\Gamma_t^*-\Gamma_t\right)}_{B_t}\frac{1}{2}(S_{t+h}-S_t)^2. \tag{17.28}$$

By first choosing w_t to make the B_t term zero and then v_t to make the A_t term zero, we get a hedge. This clearly gives

$$w_t=\Gamma_t/\Gamma_t^*\text{, and} \tag{17.29}$$

$$v_t=\Delta_t-\left(\Gamma_t/\Gamma_t^*\right)\Delta_t^*. \tag{17.30}$$

Example 17.20 *(Delta-gamma hedging) Suppose* $(\Delta_t,\Gamma_t)=(0.5,0.07)$ *and* $(\Delta_t^*,\Gamma_t^*)=(0.3,0.03)$, *then* $B_t=w_t0.03-0.07=0$ *requires* $w_t=2.33$ *and* $A_t=v_t+2.33\times0.3-0.5=0$ *requires* $v_t=-0.2$. *Clearly, this is quite different from a delta hedge (which has*

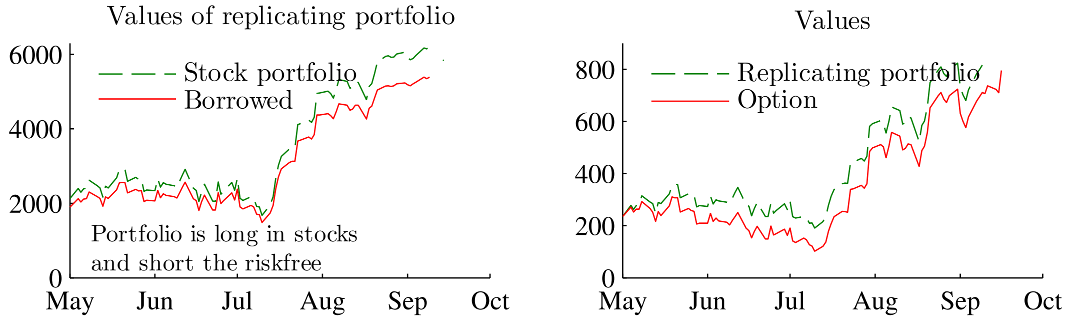

Figure 17.18: Delta hedging an SMI call option

$v_t = 0.5$ *and* $w_t = 0$*). Here, the lower sensitivity (gamma) of the second option to the quadratic term, means that the hedge portfolio includes a lot of the second option. As a consequence, it becomes overexposed to the linear term, which is compensated for by a short position in the underlying asset.*

17.4.3 Delta-Vega Hedging*

The volatility of financial markets seem to change over time. To account for that, a first-order Taylor approximation of the call option price with respect to the underlying and the volatility is

$$C_{t+h} - C_t \approx \frac{\partial C_t}{\partial S}(S_{t+h} - S_t) + \frac{\partial C_t}{\partial \sigma}(\sigma_{t+h} - \sigma_t). \tag{17.31}$$

Consider a more sophisticated hedge portfolio with Δ_t of the underlying asset, v_t of a financial instrument that moves in tandem with the volatility (for instance a futures on the VIX) and short one call option. The value of the overall position is

$$V_t = \Delta_t S_t + v_t \sigma_t - C_t. \tag{17.32}$$

By comparing (17.31) and (17.32) it is clear that the value of the overall portfolio would be immune (to a first order approximation) against movements in both the underlying and the volatility if $\Delta_t = \partial C_t / \partial S$ and $v_t = \partial C_t / \partial \sigma$. Although the Black-Scholes model is inconsistent with the idea of a time-varying volatility, a starting point could still be to use the derivatives according to that model.

17.5 Estimating Riskneutral Distributions*

We have seen that the price of a derivative is a discounted risk-neutral expectation of the derivative payoff, see (17.9).

In the Black-Scholes model, this risk-neutral distribution is that $\ln S_m$ is normally distributed as in (17.2) except that the mean is different (this is the difference between the natural and the risk-neutral distribution). However, risk neutral distributions can be derived from other assumptions than those in the Black-Scholes model, and (17.9) would still be valid. For instance, it holds in the binomial model, whose distribution is not normal (unless we make the time steps very many and small). Alternatively, we could construct a binomial tree where the time steps have different volatilities (this is often done to fit the yield curve)—and even in the limit (with many and small time steps) the distribution would be non-normal. Once again, the Black-Scholes formula would not be exact, but (17.9) would still be true.

Example 17.21 *(Call prices, three states) Suppose that S_m only can take three values: 90, 100, and 110; and that the risk neutral probabilities for these events are: 0.5, 0.4, and 0.1, respectively. We consider three European call option contracts with the strike prices 89, 99, and 109. From (17.9) their prices are (if $y = 0$)*

$$\begin{aligned} C\ (K = 89) &= 0.5(90 - 89) + 0.4(100 - 89) + 0.1(110 - 89) = 7 \\ C\ (K = 99) &= 0.5 \times 0 + 0.4(100 - 99) + 0.1(110 - 99) = 1.5 \\ C\ (K = 109) &= 0.5 \times 0 + 0.4 \times 0 + 0.1(110 - 109) = 0.1. \end{aligned}$$

With prices on several options with different strike prices (but otherwise identical), it is possible to estimate the risk-neutral distribution.

Example 17.22 *(Extracting probabilities) Suppose we observe the option prices in Example 17.21, and want to use these to recover the probabilities. We know the possible states, but not their probabilities. Let $\Pr(x)$ denote the probability that $S_m = x$. From Example 17.21, we have that the option price for $K = 109$ equals*

$$\begin{aligned} C\ (K = 109) &= 0.1 \\ &= \Pr(90) \times 0 + \Pr(100) \times 0 + \Pr(110)(110 - 109), \end{aligned}$$

which we can solve as $\Pr(110) = 0.1$. *We now use this in the expression for the option price for* $K = 99$

$$\begin{aligned} C\,(K = 99) &= 1.5 \\ &= \Pr(90) \times 0 + \Pr(100)(100 - 99) + 0.1(110 - 99), \end{aligned}$$

which we can solve as $\Pr(100) = 0.4$. *Since probabilities sum to one, it follows that* $\Pr(90) = 0.5$.

A common approach is to make an assumption about the form of the distribution, for instance, that it is mixture of two normal distributions. The parameters of this distribution are then chosen (estimated) by minimizing the sum (across strike prices) of squared differences between observed and predicted prices. (This is like the minimization problem behind the least squares method in econometrics.) This allows the possibility to pick up skewed (downside risk different from upside risk?) and even bimodal distributions.

Remark 17.23 *Figure 17.19 shows some data and results (assuming a mixture of two normal distributions) for German bond options around the announcement of the very high money growth rate on 2 March 1994.*

Remark 17.24 *Figures 17.20–17.22 show results for the CHF/EUR exchange rate around the period of active (Swiss) central bank interventions on the currency market.*

Remark 17.25 *(Robust measures of the standard deviation and skewness) Let* P_α *be the* α*th quantile (for instance, quantile 0.1) of a distribution. A simple robust measure of the standard deviation is just the difference between two symmetric quantile,*

$$\text{Std} = P_{1-\alpha} - P_\alpha,$$

where it is assumed that $\alpha < 0.5$. *Sometimes this measure is scaled so it would give the right answer for a normal distribution. For instance, with* $\alpha = 0.1$, *the measure would be divided by 2.56 and for* $\alpha = 0.25$ *by 1.35.*
One of the classical robust skewness measures was suggested by Hinkley

$$Skew = \frac{(P_{1-\alpha} - P_{0.5}) - (P_{0.5} - P_\alpha)}{P_{1-\alpha} - P_\alpha}.$$

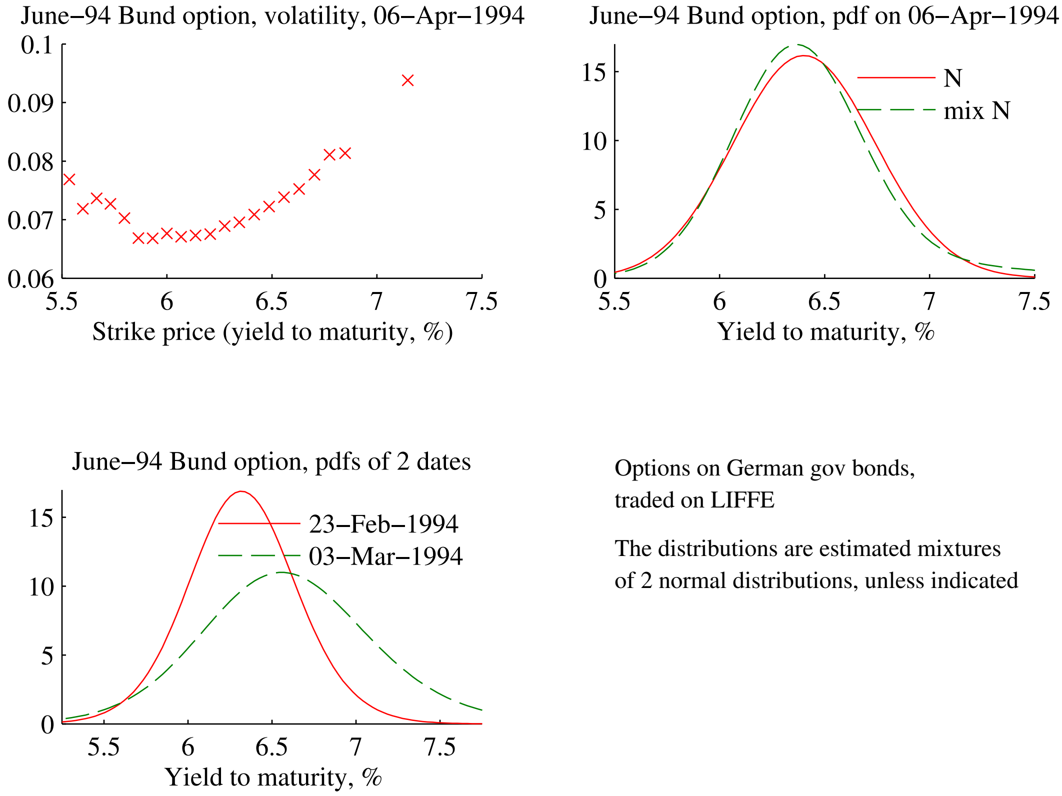

Figure 17.19: Bund options 23 February and 3 March 1994. Options expiring in June 1994.

This skewness measure can only take on values between -1 *(when* $P_{1-\alpha} = P_{0.5}$*) and* 1 *(when* $P_{\alpha} = P_{0.5}$*). When the median is just between the two percentiles (*$P_{0.5} = (P_{1-\alpha} + P_{\alpha})/2$*), then it is zero.*

Bibliography

Cox, J. C., S. A. Ross, and M. Rubinstein, 1979, "Option pricing: a simplified approach," *Journal of Financial Economics*, 7, 229–263.

Elton, E. J., M. J. Gruber, S. J. Brown, and W. N. Goetzmann, 2010, *Modern portfolio theory and investment analysis*, John Wiley and Sons, 8th edn.

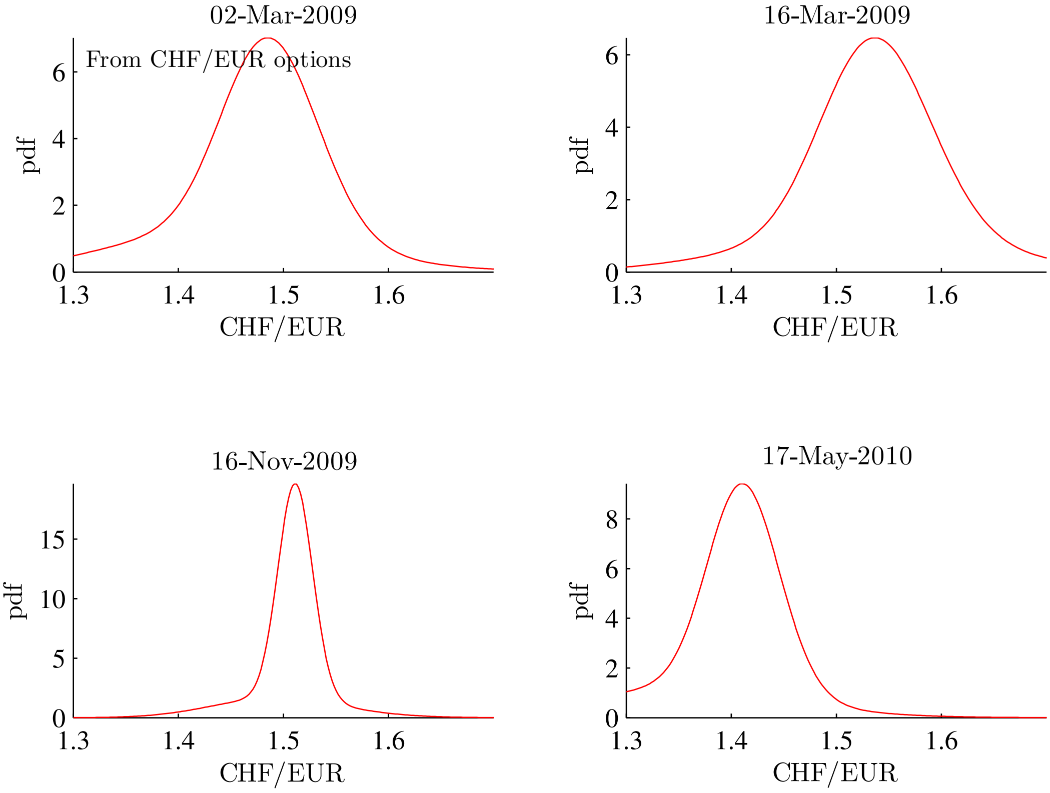

Figure 17.20: Riskneutral distribution of the CHF/EUR exchange rate

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

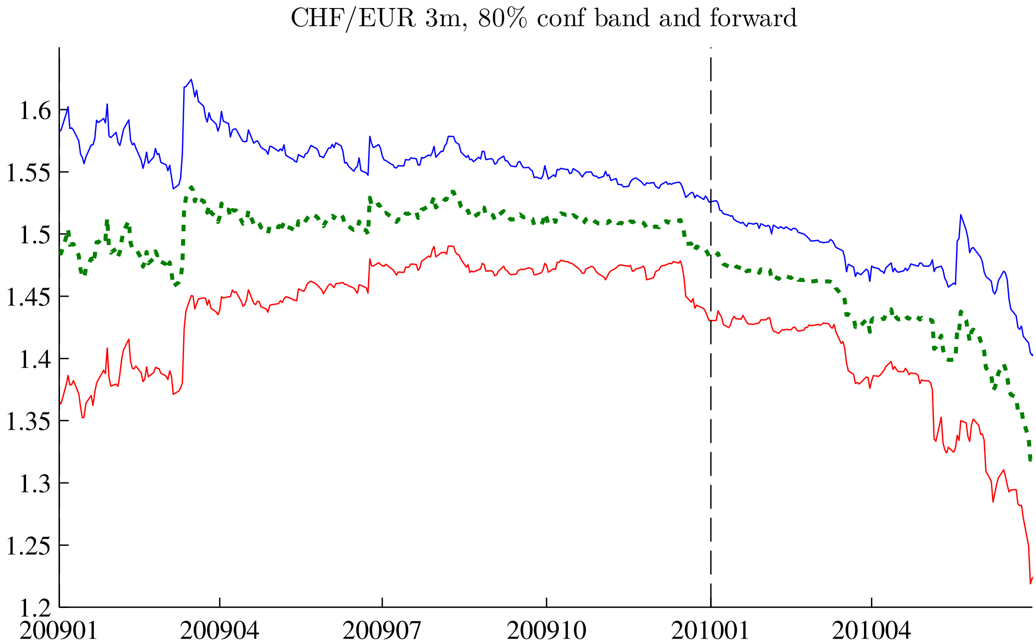

Figure 17.21: Riskneutral distribution of the CHF/EUR exchange rate

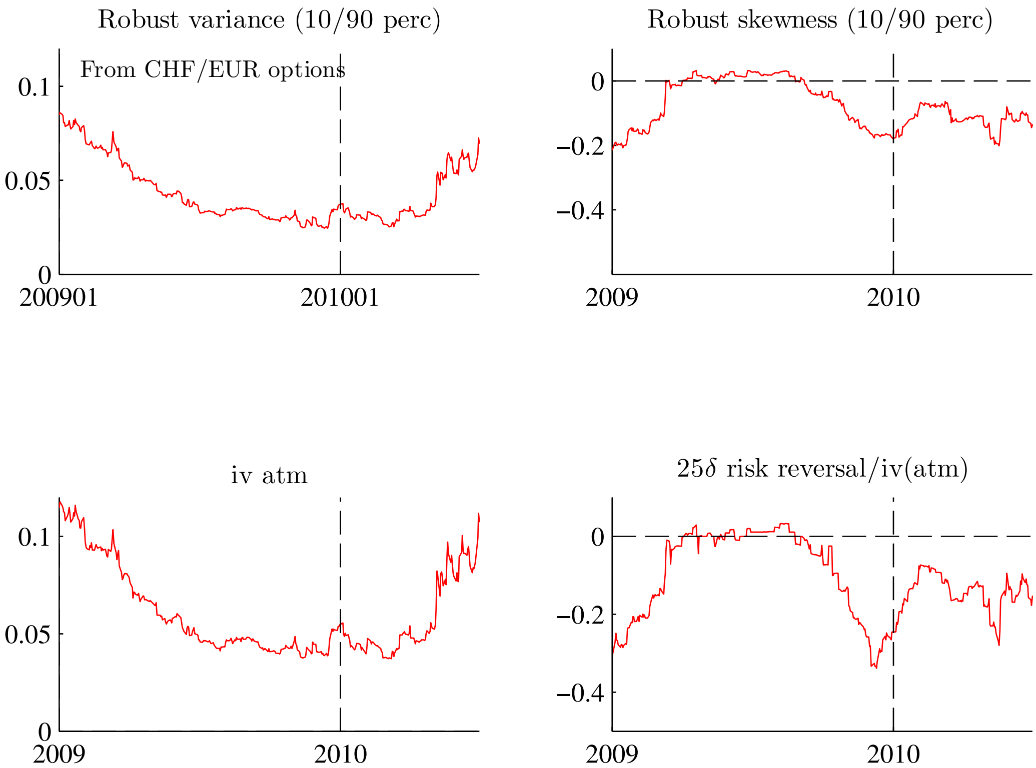

Figure 17.22: Riskneutral distribution of the CHF/EUR exchange rate

18 FX and Interest Rate Options*

Main references: Hull (2009) 15 and Wystup (2006)

18.1 Forward Contract on a Currency

Proposition 18.1 *(Forward-spot parity, continuous dividends) When the dividend is paid continuously as the rate δ (of the price of the underlying asset), then*

$$e^{-ym}F = Se^{-\delta m}, \tag{18.1}$$

where m is the time to expiration of the forward contract, y is continuously compounded interest rate (until expiration), F the forward price (to be paid at expiration) and S is the current asset price.

Investing in foreign currency effectively means investing in a foreign interest bearing instrument which earns the continuous interest rate ("dividend"), denoted δ in (18.1).

Consider a forward contract that expires in $t+m$, but where the contract was written at some earlier point in time ($\tau < t$) and specified a forward price of F_τ. The value of this contract in t is

$$W_t = e^{-ym}(F_t - F_\tau). \tag{18.2}$$

18.2 Summary of the Black-Scholes Model

For an asset with a continuous dividend rate of δ, the forward-spot parity says $F = Se^{(y-\delta)m}$. In this case the Black-Scholes formula can be written

$$C = e^{-\delta m}S\Phi(d_1) - e^{-ym}K\Phi(d_2)\text{, where} \tag{18.3}$$

$$d_1 = \frac{\ln(S/K) + (y - \delta + \sigma^2/2)m}{\sigma\sqrt{m}} \text{ and } d_2 = d_1 - \sigma\sqrt{m}. \tag{18.4}$$

When the asset is a currency (read: foreign money market account) and δ is the foreign interest rate, then this is the "Garman-Kolhagen" formula. The sensitivity to the underlying asset price is

$$\frac{\partial C}{\partial S} = e^{-\delta m}\Phi(d_1), \tag{18.5}$$

where d_1 is given by (18.4). Without dividends, just set $\delta = 0$. From the put-call parity, it is clear that the sensitivity of a put option is

$$\frac{\partial P}{\partial S} = e^{-\delta m}[\Phi(d_1) - 1] = -e^{-\delta m}\Phi(-d_1), \tag{18.6}$$

which is negative. The second equality follows from the symmetry of the normal distribution.)

Black's model for an option on a forward contract is

$$C = e^{-ym}F\Phi(d_1) - e^{-ym}K\Phi(d_2), \text{ where} \tag{18.7}$$

$$d_1 = \frac{\ln(F/K) + (\sigma^2/2)m}{\sigma\sqrt{m}} \text{ and } d_2 = d_1 - \sigma\sqrt{m}. \tag{18.8}$$

For an asset with a continuous dividend, this d_1 is the same as in (18.4). The sensitivity of the call option price to the forward price is

$$\frac{\partial C}{\partial F} = e^{-ym}\Phi(d_1), \tag{18.9}$$

where d_1 is given by (18.8). Similarly, the sensitivity of a put option is

$$\frac{\partial P}{\partial F} = e^{-ym}[\Phi(d_1) - 1] = -e^{-ym}\Phi(-d_1). \tag{18.10}$$

18.3 Hedging

18.3.1 Hedging with the Underlying Asset

Consider a portfolio with Δ_t of the underlying asset (the hedging portfolio) and short one call option. The value of the overall position is

$$V_t = \Delta_t S_t - C_t. \tag{18.11}$$

This portfolio has an almost stable value if

$$\Delta_t = \frac{\partial C_t}{\partial S}. \tag{18.12}$$

18.3.2 Hedging with a Forward Contract

Consider a portfolio with Δ_t of a forward contract and short one call option. The value of the overall position is

$$V_t = \Delta_t W_t - C_t, \tag{18.13}$$

where W_t is the value of an old forward contract as in (18.2). This portfolio is almost stable if

$$\Delta_t = e^{ym} \frac{\partial C_t}{\partial F}. \tag{18.14}$$

Proof. Assume that only the forward price can change (clearly not true, but at least a starting point for the analysis). A first-order Taylor approximation of the call option price is

$$C_{t+h} - C_t \approx \frac{\partial C_t}{\partial F} (F_{t+h} - F_t).$$

Use to approximate the change of the value of the overall portfolio as

$$\begin{aligned}
V_{t+h} - V_t &= \Delta_t (W_{t+h} - W_t) - C_{t+h} - C_t \\
&\approx \Delta_t e^{-ym} (F_{t+h} - F_t) - \frac{\partial C_t}{\partial F} (F_{t+h} - F_t) \\
&\approx 0 \text{ if } \Delta_t = e^{ym} \frac{\partial C_t}{\partial F}.
\end{aligned}$$

■

18.4 FX Options: Put or Call?

Buying one currency entails selling another. It should therefore come as no surprise that a call option on a currency is also a put option on the other currency. To be precise, the option prices are related according to

$$C_d(\text{strike} = K) = S_t K P_f(\text{strike} = 1/K). \tag{18.15}$$

On the left hand side, C_d is the domestic price of a call option on the foreign currency—with the strike price (K) is expressed in the domestic currency. On the right hand side, S_t is the current exchange rate (domestic price of one unit of the foreign currency), and P_f is the foreign price of a put option on the domestic currency—with the strike price ($1/K$).

In particular, we can rewrite the expression as

$$P_f(\text{strike} = 1/K) = \frac{C_d(\text{strike} = K)}{S_t K}, \tag{18.16}$$

which expresses the price (measured in foreign currency) of the put on the domestic currency.

Example 18.2 *Let* $C_d = £0.01$ *for an option on US dollars and the strike price is £0.6 (to get one dollar). If the current exchange rate is £0.58 (per dollar), then the dollar price of a put option on GBP with a strike price of* $1/0.6$ *dollars per GBP is* $0.01/(0.58 \times 0.6) =$ \$0.0287.

Remark 18.3 *(Option price quoted in which currency?) In practice, it is important to consider which currency the option price is quoted in. For instance, most options involving the USD have option prices quoted in USD, while most options involving the EUR (unless also the USD is involved) have prices quoted in EUR.*

Proof. (of (18.15)) The payoff of a call option (denominated in the domestic currency) on foreign currency with strike price K is

$$\max(0, S_{t+m} - K),$$

where K is the strike price and S_{t+m} is the exchange rate at expiration—both expressed as the domestic price of one unit of foreign currency (for instance, GBP 0.6 per USD). The payoff is clearly expressed in the domestic currency. In contrast, the payoff of a put option (denominated in the foreign currency) on the domestic currency (with strike price $1/K$) has the payoff

$$\max(0, 1/K - 1/S_{t+m}),$$

which is clearly expressed in the foreign currency. Notice that both options are exercised when $S_{t+m} > K$. In fact, these options are identical, except for a scaling factor and the currency denomination. To see that, consider buying K of the foreign denominated options and then convert the payoff to the domestic currency (multiply by S_{t+m})

$$S_{t+m} K \max(0, 1/K - 1/S_{t+m}) = \max(0, S_{t+m} - K),$$

which is clearly the same as for the first option. For that reason, buying K of the foreign currency denominated put options should have the same price (when measured in domestic currency—multiply by S_t) as the domestically denominated call option. ■

18.5 FX Options: Risk Reversals and Strangles

Options on the FX (exchange rate) markets are often sold (on the OTC market) as special portfolios (consisting of straddles, risk-reversals and strangles) and quoted in terms of the implied volatilities. Apart from these conventions, options on exchange rates are no different from options on other assets (but, remember that currencies carry "dividends" since holding a currency in practice means holding a money market account in that currency).

A *delta-neutral straddle*, that is, a long position in a call and also in a put. To make it delta-neutral (with respect to the spot), we need

$$\frac{\partial C}{\partial S} + \frac{\partial P}{\partial S} = 0, \tag{18.17}$$

which from (18.5)–(18.6) gives (with d_1 defined by (18.8) or equivalently (18.4))

$$d_1 = 0, \text{ that is, } K_{atm} = F e^{\sigma_{atm}^2 m/2}. \tag{18.18}$$

This straddle is typically quoted in terms of the implied volatility (σ_{atm}) of an option at K_{atm}. A higher value of the straddle indicates more overall uncertainty. See Figure 18.1 for illustrations of the profits of different option portfolios.

A *25-delta risk reversal* is a portfolio of one call option with a strike price K_2 such that the delta is 0.25 and short one put option with a strike price K_1 such that the delta is -0.25. Both options are out of the money so the strike price for the put is lower than the forward price, which in turn is lower than the strike price of the call ($K_1 < F < K_2$). The risk reversal is typically quoted as the difference of the two implied volatilities

$$rr = \sigma_2 - \sigma_1, \tag{18.19}$$

where σ_2 and σ_1 are the implied volatilities of the options with strike prices K_2 and K_1 respectively (notice that, by the put-call parity, a put and a call with the same strike price have the same implied volatility). A higher value of the risk reversal indicates beliefs of an increase in the underlying—so it captures skewness of the exchange rate distribution.

A *25-delta strangle* has a long position in the 25-delta call and also in the 25-delta put. A *25-delta butterfly* is a portfolio that is long one 25-delta straddle and short one delta-neutral straddle. It is typically quoted as the average implied volatility of the K_2 and K_1 options (call and put, respectively) minus the at-the-money volatility

$$bf = \frac{\sigma_2 + \sigma_1}{2} - \sigma_{atm}. \tag{18.20}$$

An increase in bf signals a belief in fatter tails, so it captures kurtosis. Notice that a proportional increase of all volatilities does not change bf (it is “vega” neutral).

With the quotes on the risk reversal (18.19) and the butterfly (18.20), we can solve for the implied volatilities σ_1 and σ_2 as

$$\begin{aligned}\sigma_1 &= bf + \sigma_{atm} - rr/2 \\ \sigma_2 &= bf + \sigma_{atm} + rr/2.\end{aligned} \tag{18.21}$$

It is straightforward to invert the formulas for the deltas to derive what the strike prices are. If we use the convention that the deltas are with respect to the spot price, then by setting $\partial C/\partial S = \Delta$ (say, $\Delta = 0.25$) in (18.5) to derive the strike price K_2 and $\partial P/\partial S = -\Delta$ in (18.6) to derive the strike price K_1 we get the following strike prices (using $F = Se^{(y-\delta)m}$)

$$\begin{aligned}K_1 &= F \exp[\sigma_1\sqrt{m}\Phi^{-1}(e^{\delta m}\Delta) + m\sigma_1^2/2] \\ K_2 &= F \exp[-\sigma_2\sqrt{m}\Phi^{-1}(e^{\delta m}\Delta) + m\sigma_2^2/2],\end{aligned} \tag{18.22}$$

Clearly, by changing setting $\Delta = 0.25$, we get the results for a 25-delta risk reversal.

See Figure 18.2 for how the strike prices are calculated and Figure 18.3 for an empirical illustration.

Example 18.4 *(σ_{atm}, rr and bf on 1 April 2005, 1-month EUR/GBP) For this particular date and contract σ_{atm} was* 4.83%*, the 25 delta risk reversal was* 0.18% *and the 25 delta strangle (really, a 25 delta butterfly) was* 0.15%*. (See Wystup (2006), tables 1.7–9.) This gives*

$$\begin{aligned}\sigma_1 &= 0.15 + 4.83 - 0.18/2 = 4.89 \\ \sigma_2 &= 0.15 + 4.83 + 0.18/2 = 5.07.\end{aligned}$$

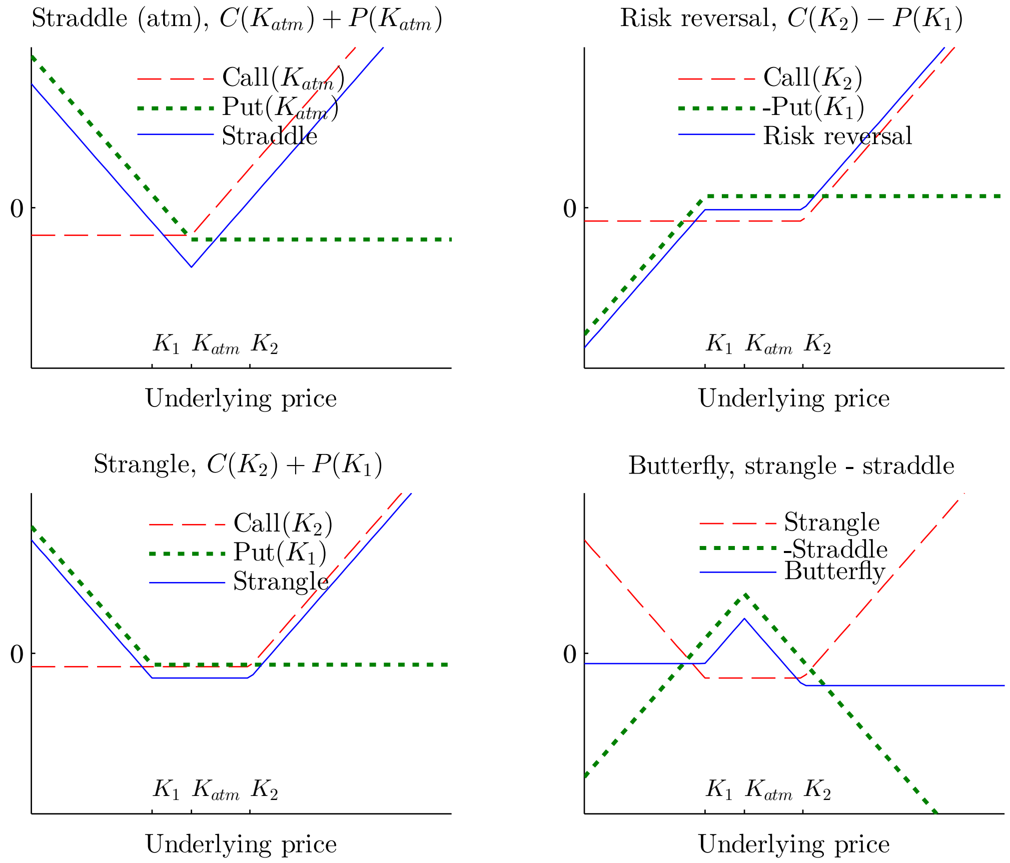

Figure 18.1: Profits diagrams for FX option portfolios

The spot exchange rate was 0.6859 (the price of one EUR, in terms of GBP) and the 1-month interest rates were 4.87 in the UK and 2.10 in the euro zone, so the forward rate was $F = 0.6859 \times \exp[(0.0487 - 0.0210)/12] \approx 0.6875$. *This gives* $K_1 = 0.6811$, $K_{atm} = 0.6876$ *and* $K_2 = 0.6941$.

Remark 18.5 *(Deltas with respect to the forward price*) If we instead use the convention that the deltas are with respect to the forward price, then* K_{atm} *in (18.18) is unchanged, but* Δ *is substituted for* Δ *in (18.22). Both conventions are used. (The forward deltas are more common for options with long time expiration and for emerging market currencies.)*

Proof. (of (18.18)) If we use spot deltas, then (18.5)–(18.6) give

$$\frac{\partial C}{\partial S} + \frac{\partial P}{\partial S} = e^{-\delta m} \Phi\left(d_1\right) - e^{-\delta m} \Phi\left(-d_1\right) = 0,$$

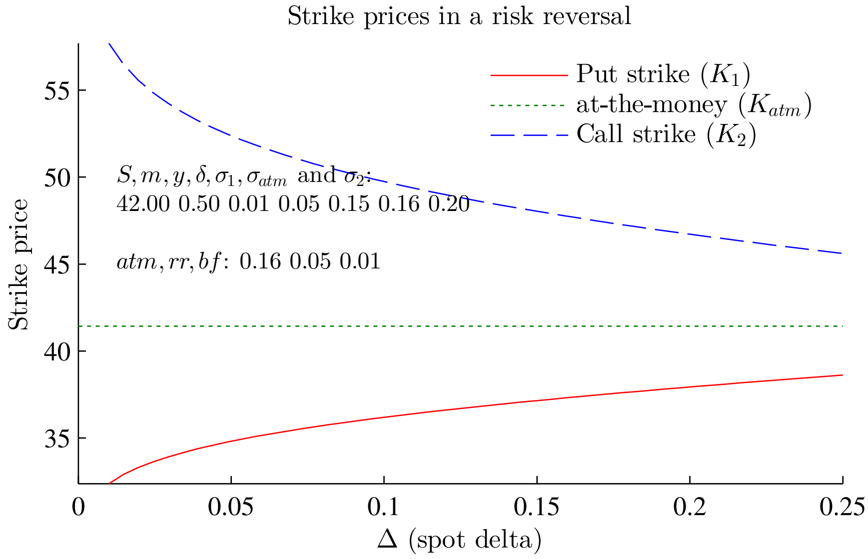

Figure 18.2: Strike prices in a risk reversal

which requires $d_1 = 0$. With d_1 defined by (18.4) we have

$$\ln K = \ln S + (y - \delta + \sigma^2/2)m = \ln F + (\sigma^2/2)m$$

If we instead use forward deltas, use (18.12) and (18.9)–(18.10) and set to zero

$$\frac{\partial C}{\partial F} + \frac{\partial P}{\partial F} = e^{ym}\left[e^{-ym}\Phi(d_1) - e^{-ym}\Phi(-d_1)\right] = 0,$$

which still requires $d_1 = 0$ (and d_1 is the same in (18.8) and (18.4)). ■

Proof. (of (18.22)) If we use the spot delta, then set (18.5) equal to 0.25

$$0.25 = e^{-\delta m}\Phi(d_1)\text{, so we need } d_1 = \Phi^{-1}\left(e^{\delta m}0.25\right)$$

With d_1 given by (18.4) we get

$$\ln K_2 = \ln F + (\sigma^2/2)m - \sigma\sqrt{m}\Phi^{-1}(e^{\delta m}0.25),$$

since $\ln F = \ln S + (y - \delta)m$. Instead, if we use the forward delta from using (18.9) in

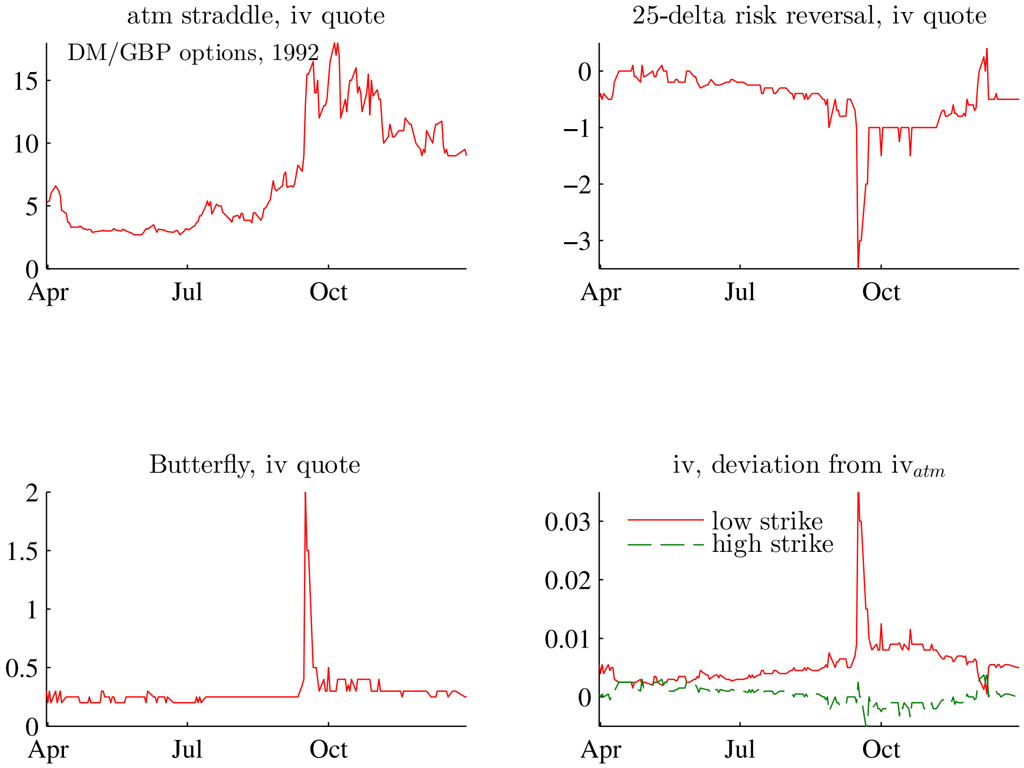

Figure 18.3: DM/GDP options, 1992

(18.14)

$$
\begin{aligned}
0.25 &= e^{ym}\frac{\partial C_t}{\partial F} = \Phi\left(d_1\right), \text{ so} \\
d_1 &= \Phi^{-1}(0.25).
\end{aligned}
$$

With d_1 given by (18.8) we get

$$\ln K_2 = \ln F + (\sigma^2/2)m - \sigma\sqrt{m}\Phi^{-1}(0.25).$$

The calculations for the strike prices K_1 for the put are similar. ■

18.6 FX Options: Implied Volatility for Different Deltas

Another way to quote FX option prices is to list the implied volatility for different strike prices—but where the strike prices are expressed as deltas. For instance, $\Delta = (-0.25, 0, 0.25)$. Often, these are labelled "$25\Delta P$", atm, and "$25\Delta C$", where $25\Delta P$ stands for the strike price where a put has a delta of -0.25, atm stands for the strike price at the money, and $25\Delta C$ is the strike price where a call has a delta of 0.25.

Typically, the atm strike price is as in (18.18), while the "$25\Delta P$" strike price is calculated as K_1 in (18.22) by setting $\Delta = 0.25$ and the "$25\Delta C$" strike price is calculated as K_2 in (18.22) by setting $\Delta = 0.25$. Other deltas are similar.

Remark 18.6 *(Premium-adjusted deltas) When the option price is quoted in the foreign currency, then the deltas reported do not correspond to (18.18) and (18.22). See Wystup (2006) for more details.*

18.7 Options on Interest Rates: Caps and Floors

Options on bonds are basically no different from options on equity, although bonds typically pay "dividends" (the coupons). For instance, a call option on a bond gives the right to buy the bond (at the expiration of the option) at the strike price.

Options on interest rates are also very similar, but often have a more complicated structure. A *caplet* is a call option that protects against higher interest rates (typically a floating 3-month market rate or similar). Let Z_{t+s} be the (annualized) market interest rate for a loan between $t + s$ and $t + s + m$ and let Z_K be the (annualized) cap rate. The payoff in $t + s + m$ (notice: paid at the end of the borrowing period) is

$$\max[0, m(Z_{t+s} - Z_K)]. \tag{18.23}$$

The second term is the interest rate cost for a loan (with a face value of unity) between $t + s$ and $t + s + m$ according to the market rate minus the same cost according to the cap rate. Clearly, buying such an option is a way to make sure that interest rate paid on a loan will not exceed the cap rate. If settled at $t + s$ the payoff is just the discounted value

$$\frac{\max[0, m(Z_{t+s} - Z_K)]}{1 + mZ_{t+s}}. \tag{18.24}$$

The payoff in (18.24) can be rewritten as

$$(1 + mZ_K) \max\left(0, \frac{1}{1 + mZ_K} - B_{t+s}(m)\right) \tag{18.25}$$

Notice that the max() term defines the payoff of a put option on an m-period bond in $t + s$ (whose value turns out to be $B_{t+s}(m) = 1/(1 + mZ_{t+s})$)—with a strike price of of $1/(1 + mZ_K)$. The caplet is therefore proportional to a put option on a bond.

Proof. (of (18.25)) Multiply and divide (18.24) by $(1 + mZ_K)$ and rearrange

$$\begin{aligned}&(1 + mZ_K) \max\left[0, \frac{mZ_{t+s} - mZ_K}{(1 + mZ_{t+s})(1 + mZ_K)}\right] \\ &= (1 + mZ_K) \max\left(0, \frac{1}{1 + mZ_K} - \frac{1}{1 + mZ_{t+s}}\right).\end{aligned}$$

Notice that $B_{t+s}(m) = 1/(1 + mZ_{t+s})$. ■

We can apply the Black's formula (18.7)–(18.8) to price the caplet by assuming that a forward contract on either Z_{t+s} or (somewhat less often) B_{t+s} has a lognormal distribution. (These two assumptions are not compatible, since the latter is the same as assuming that $1 + mZ_{t+s}$ has a lognormal distribution.)

Remark 18.7 *(Simple interest rates) If Z is a simple interest rates, then of a zero-coupon bond that gives unity at maturity is*

$$B\,(m) = \frac{1}{1 + mZ(m)}, \text{ or } Z(m) = \frac{1/B\,(m) - 1}{m}.$$

A simple forward rate for the period s to $s + m$ periods in the future is defined as

$$Z^f(s, s + m) = \frac{1}{m}\left[\frac{B(s)}{B(s + m)} - 1\right].$$

A forward rate (determined t) for the future investment period $t + s$ to $t + s + m$, denoted Z^f, clearly coincides with the market rate in $t + s$. We can therefore apply Black's formula to the underlying mZ^f by assuming that it is lognormally distributed—and using the strike "price" mZ_K. However, we need to discount by $\exp[-(s + m)y]$ instead of $\exp(-sy)$ since the payoff (18.23) is paid in $t + s + m$ (not in $t + s$). The value of this caplet is therefore

$$\text{Caplet}(s, m; \sigma, Z_K) = me^{-(s+m)y}[Z^f \Phi(d_1) - Z_K \Phi(d_2)], \text{ where} \tag{18.26}$$

$$d_1 = \frac{\ln(Z^f/Z_K) + (\sigma^2/2)s}{\sigma\sqrt{s}} \text{ and } d_2 = d_1 - \sigma\sqrt{s}, \tag{18.27}$$

where σ is the (annualized) volatility of the log forward rate.

An *interest rate cap* is a portfolio of different caplets which protects the owner over several *tenors* (subperiods). Typically, the first caplet is deleted (as there is no uncertainty about what the short rate is today) and the last payment is done on the maturity date n. Therefore, the tenors are $[m, 2m]$, $[2m, 3m]$ and so forth until the last one which is $[n-m, n]$ so there are $n/m-1$ caplets. (The start/end of a tenor is called a reset/settlement date.) For instance, a 1-year cap on the 3-month Libor consists of 3 caplets. See Figure 18.4 for an illustration. (The cap could also be scheduled to start at a later date.)

If we apply the same volatility to all caplets ("flat volatilities"), then the price of a cap (according to the Black-Scholes model) starting now and ending in n, is

$$\text{Cap}(n, m; \sigma, Z_K) = \sum\nolimits_{i=1}^{n/m-1} \text{Caplet}(im, m; \sigma, Z_K). \tag{18.28}$$

Caps are often quoted in terms of the implied volatility (σ) that solves this equation—meaning that there is one implied volatility per cap contract, but it may differ across cap rates ("strike prices") and maturities. (If the cap is scheduled to start S periods ahead, instead of now, then im should be replaced by $S + im$.)

Example 18.8 *(1-year Cap starting now, 3-month tenors) Let* $n = 1$ *(1-year cap) and* $m = 1/4$ *(3-month tenors). The payoffs are based on the difference between the 3-month Libor and the cap rate at the beginning of the tenors (*$1/4, 2/4, 3/4$*), but are paid one quarter later. Equation (18.28) is therefore*

$$Cap(1, 1/4; \sigma, Z_K) = Caplet(1/4, 1/4; \sigma, Z_K) + Caplet(2/4, 1/4; \sigma, Z_K) + Caplet(3/4, 1/4; \sigma, Z_K).$$

Floorlets and *floors* are similar to caplets and caps, except that they pays off when the interest goes below the cap rate.

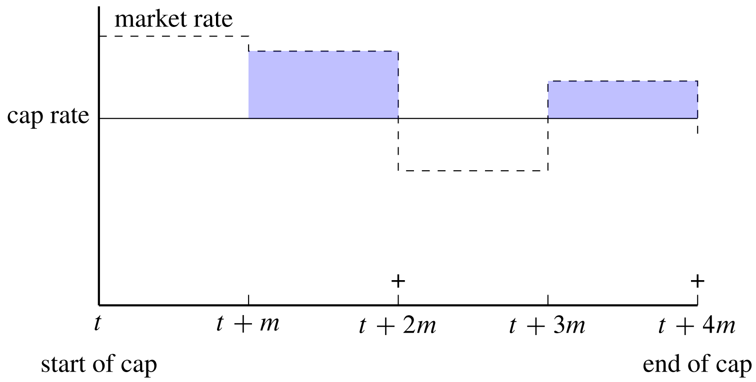

(The time of payments are marked by +)
(No payment before $t+2m$)

Figure 18.4: Interest rate cap

Bibliography

Hull, J. C., 2009, *Options, futures, and other derivatives*, Prentice-Hall, Upper Saddle River, NJ, 7th edn.

Wystup, U., 2006, *FX Options and Structured Products*, Wiley.

19 Trading Volatility

Reference: Gatheral (2006) and McDonald (2006)

More advanced material is denoted by a star (*). It is not required reading.

19.1 The Purpose of Trading Volatility

By using option portfolios (for instance, straddles) it is possible to create a position that is a bet on volatility—and is (in principle) not sensitive to the direction of change of the underlying. See Figure 19.1 for an illustration.

Volatility, as an asset class, has some interesting features. In particular, returns on the underlying asset and volatility are typically negatively correlated: very negative returns

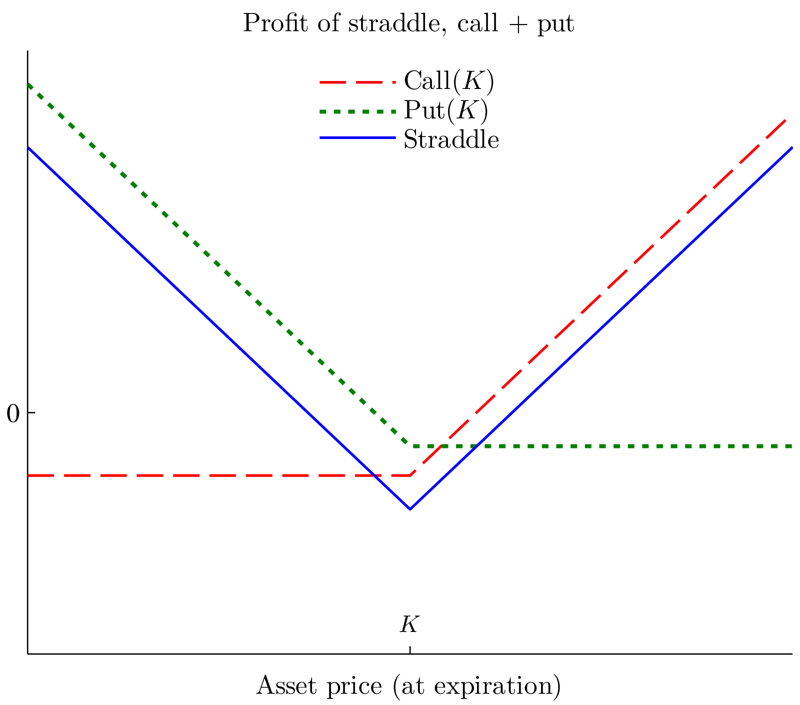

Figure 19.1: Profit of straddle

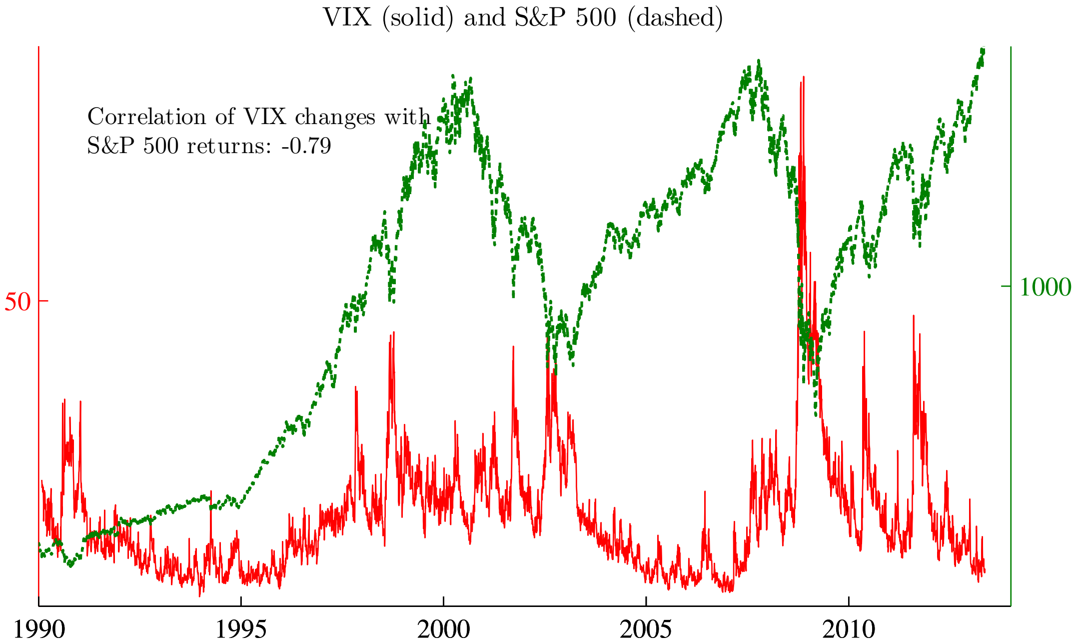

Figure 19.2: S&P 500 and VIX

are typically accompanied by increases in future actual volatility as well as beliefs about higher future volatility (as priced into options). See Figure 19.2 for an illustration, where changes in the VIX are taken to proxy the one-day holding return on a straddle.

There are several ways of trading volatility: straddles (and other option portfolios), futures (and options) on the VIX, as well as volatility (and variance) swaps.

19.2 VIX and VIX Futures

The VIX is an index of volatility, calculated from 1-month options on S&P 500. It used to be calculated as an average of implied volatilities, but since 2003 the calculation is more complicated (the old series is now called VXO). It can be shown (although it is a bit tricky) that the VIX is a very good approximation to the square root of the variance swap rate (see below) for a 30-day contract. There are also futures contracts on VIX with payoff

$$\text{VIX futures payoff}_{t+m} = VIX_{t+m} - \text{futures price}_t. \tag{19.1}$$

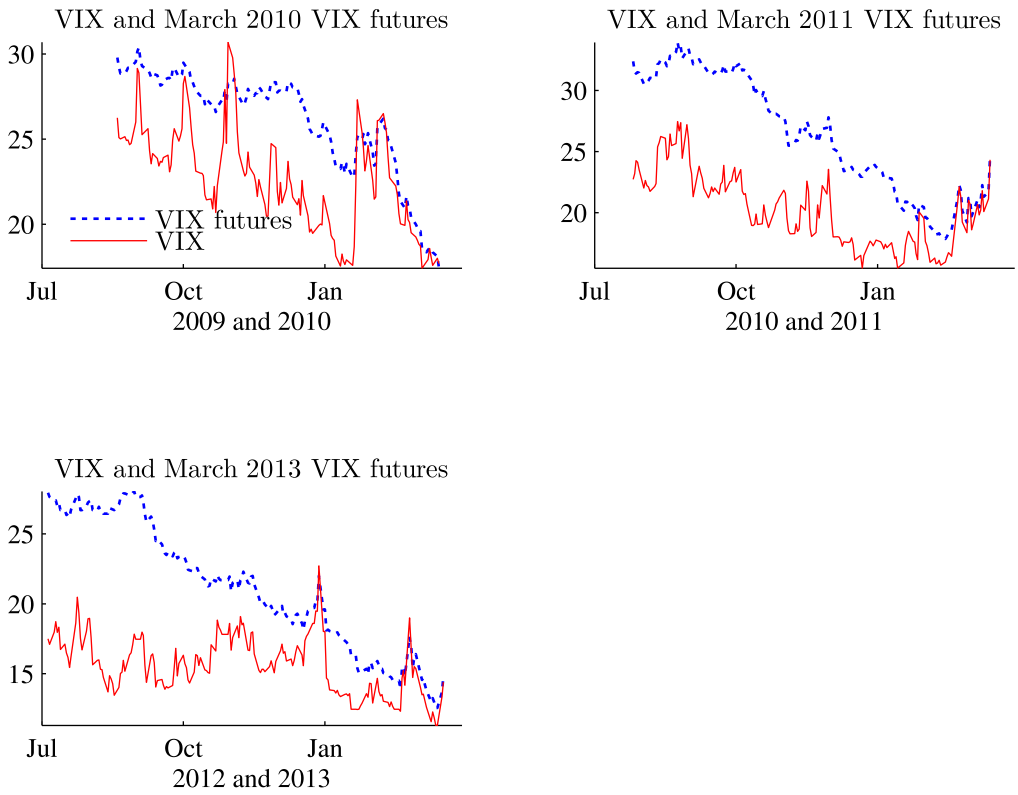

Figure 19.3: VIX and futures contract on VIX

Notice that VIX_{t+m} is really a guess of what the volatility will be during the month after $t+m$, so the futures contract pays off when the expected volatility (in $t+m$) is higher than what was thought in t.

See Figures 19.3–19.4 for an empirical illustration. Notice that the futures prices indicate that volatility is mean reverting: high VIX levels are associated with negative spreads (the futures is lower than the current VIX). This indicates that market participants believe that volatility will settle down.

Remark 19.1 *(Calculation of VIX) Let F be the forward price, $\Delta K_i = (K_{i+1} - K_{i-1})/2$ and let K_0 denote the first strike price below F. Then, the VIX is calculated as*

$$VIX^2 = \frac{2}{m}\exp(ym)\sum_{K_i \le K_0}\frac{\Delta K_i}{K_i^2}P(K_i) + \frac{2}{m}\exp(ym)\sum_{K_i > K_0}\frac{\Delta K_i}{K_i^2}C(K_i) - \frac{1}{m}(F/K_0 - 1)^2,$$

where m is the time to expiration (around 1/12), y the interest rate, $P()$ the put price and

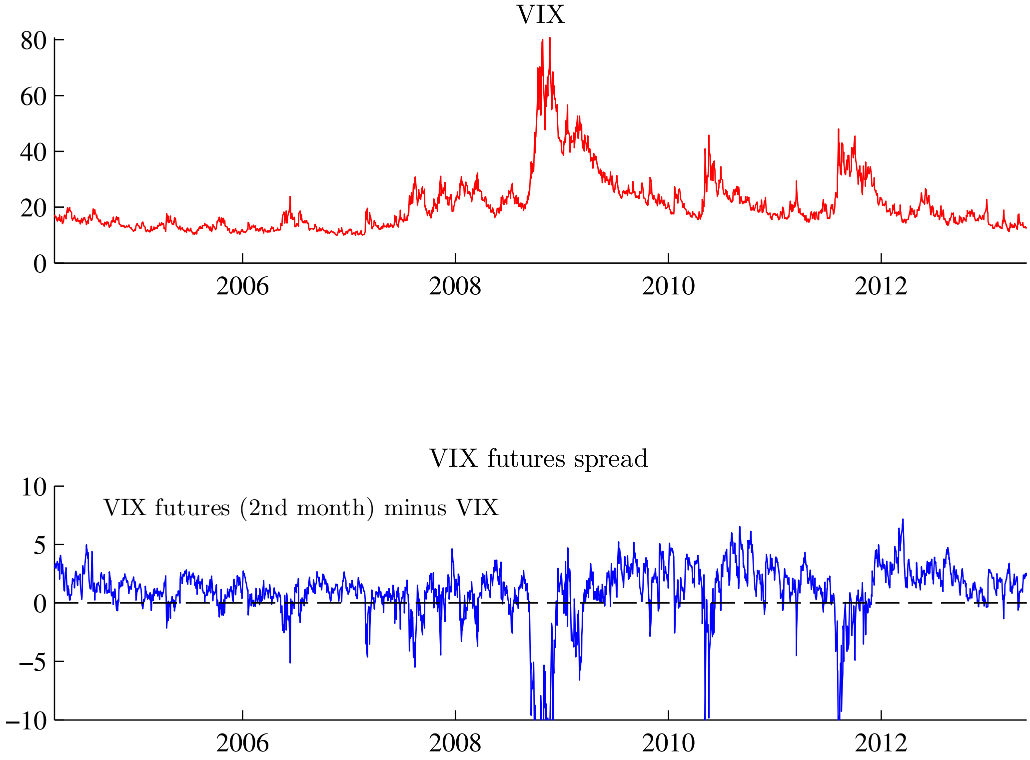

Figure 19.4: VIX futures spread

$C()$ the call price.

19.3 Variance and Volatility Swaps

Instead of investing in straddles, it is also possible to invest in *variance swaps*. Such a contract has a zero price in inception (in t) and the payoff at expiration (in $t+m$) is

$$\text{Variance swap payoff}_{t+m} = \text{realized variance}_{t+m} - \text{variance swap rate}_t, \tag{19.2}$$

where the variance swap rate (also called the strike or forward price for) is agreed on at inception (t) and the realized volatility is just the sample variance for the swap period. Both rates are typically annualized, for instance, if data is daily and includes only trading days, then the variance is multiplied by 252 or so (as a proxy for the number of trading days per year).

A *volatility swap* is similar, except that the payoff it is expressed as the difference between the standard deviations instead of the variances

$$\text{Volatility swap payoff}_{t+m} = \sqrt{\text{realized variance}_{t+m}} - \text{volatility swap rate}_t, \qquad (19.3)$$

If we use daily data to calculate the realized variance from t until the expiration(RV_{t+m}), then

$$RV_{t+m} = \frac{252}{m}\sum\nolimits_{s=1}^{m} R_{t+s}^2, \qquad (19.4)$$

where R_{t+s} is the net return on day $t+s$. (This formula assumes that the mean return is zero—which is typically a good approximation for high frequency data. In some cases, the average is taken only over $m-1$ days.)

Notice that both variance and volatility swaps pays off if actual (realized) volatility between t and $t+m$ is higher than expected in t. In contrast, the futures on the VIX pays off when the expected volatility (in $t+m$) is higher than what was thought in t. In a way, we can think of the VIX futures as a futures on a volatility swap (between $t+m$ and a month later).

Since VIX^2 is a good approximation of variance swap rate for a 30-day contract, the return can be approximated as

$$\text{Return of a variance swap}_{t+m} = (RV_{t+m} - VIX_t^2)/VIX_t^2. \qquad (19.5)$$

Figures 19.5–19.6 illustrate the properties for the VIX and realized volatility of the S&P 500. It is clear that the return of a variance swap (with expiration of 30 days) would have been negative on average. (Notice: variance swaps were not traded for the early part of the sample in the figure.) The excess return (over a riskfree rate) would, of course, have been even more negative. This suggests that selling variance swaps (which has been the speciality of some hedge funds) might be a good deal—except that it will incur some occasional really large losses (the return distribution has positive skewness). Presumably, buyers of the variance swaps think that this negative average return is a reasonable price to pay for the "hedging" properties of the contracts—although the data does not suggest a very strong negative correlation with S&P 500 returns.

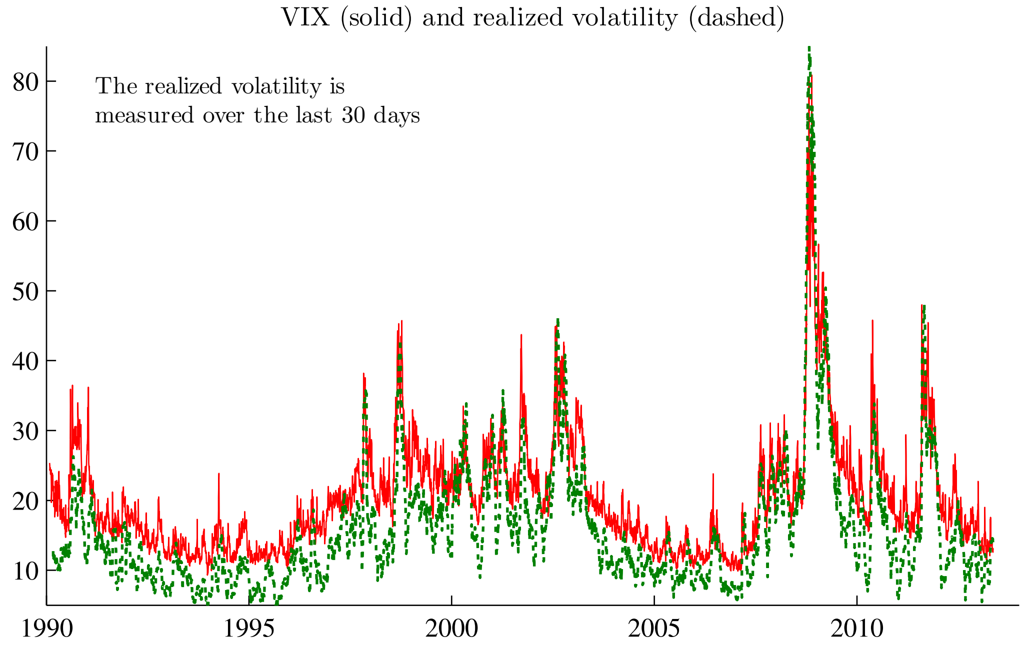

Figure 19.5: VIX and realized volatility (variance)

Bibliography

Gatheral, J., 2006, *The volatility surface: a practitioner's guide*, Wiley.

McDonald, R. L., 2006, *Derivatives markets*, Addison-Wesley, 2nd edn.

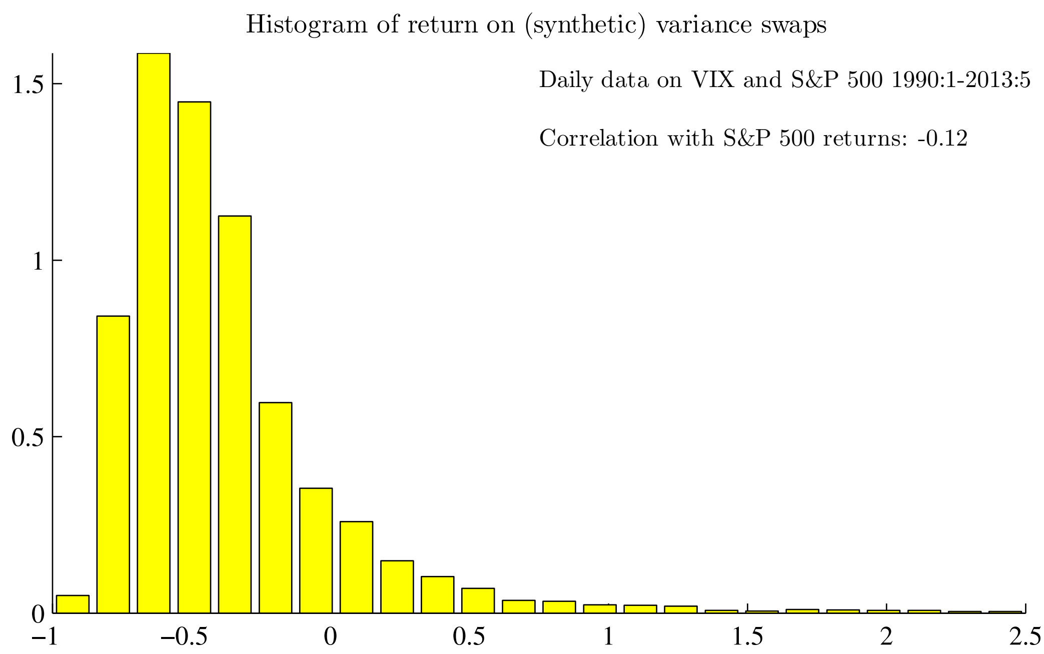

Figure 19.6: Distribution of return from investing in variance swaps