so-called bid-ask spread, i.e., the difference between the lower rate at which the foreign currency can be sold to market makers and the higher rate at which it can be bought from them. (This concept will be familiar to anyone who has exchanged currency abroad.) Depending on the particular currency pair, the rate is agreed upon today while the actual transfer takes place either in two business days (US dollar/yen, euro/US dollar, pound/US dollar) or in one business day (US dollar/Canadian dollar). Instantaneous FXRs (as well as stock and bond prices) are determined by market forces (through the "invisible hand" envisioned by Adam Smith) and reflect their intrinsic values, as well as considerations of supply and demand. FXRs constantly fluctuate around their moving equilibrium values. Their most important characteristics are rates of return on investments in foreign currency (yearly, monthly, daily, hourly, etc.). Rates of return are random and have to be treated via statistical methods for studying time series. To give the reader an idea of what can be expected, we just mention that the distribution of daily returns for the foreign exchange rate USD/DEM over a period of ten years from 1986 to 1996 has volatility 0.11, skewness -0.1, kurtosis 5, and no daily deviations exceeded five standard deviations, while for the S&P 500 Index over the same period the distribution of returns has volatility 0.16, skewness -5, kurtosis 111, and that one (five) daily deviation (deviations) exceeded ten (five) standard deviations. The distribution of daily returns for the FXR is reasonably close to Gaussian, while for the S&P 500 Index the corresponding distribution is strongly non-Gaussian. In most cases it is not necessary to explain the observed FXRs and rates of return on investments in foreign currencies in fundamental terms; the main objective is to develop a model for pricing derivative instruments in terms of the underlying ones and to solve the asset management problem. Even though, in general, the distribution of the daily returns for underlying instruments is non-Gaussian, it is frequently assumed to be Gaussian for practical purposes. Surprisingly, more often than not, this approach produces satisfactory results.

Typical behavior of FXRs is illustrated in Figure 1.1.

1.5 Derivatives: forwards, futures, calls, puts, and all that

Exchange of currencies at spot rates serves only the most obvious and the most immediate needs of market participants. Their more sophisticated needs are met by derivative instruments. The basic types of forex and equity derivatives are forward and futures contracts, and calls and puts; their fixed income counterparts are known as forward rate agreements, Euros, caps, and floors.

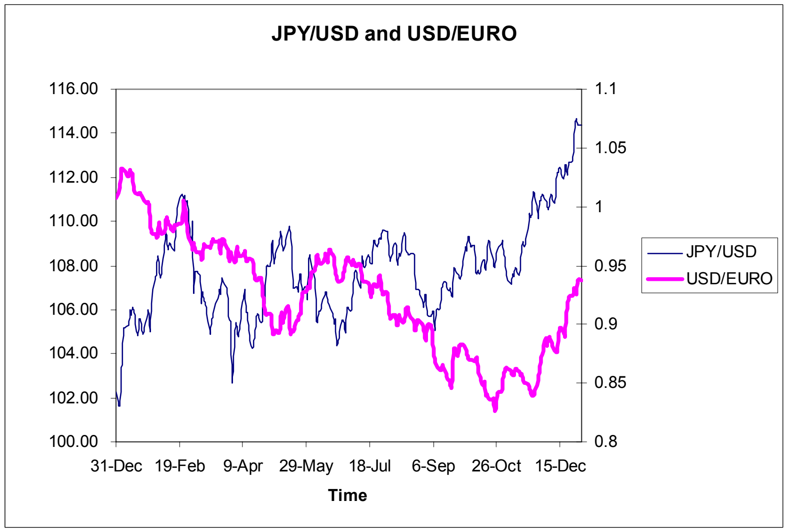

Figure 1.1: The behavior of JPY/USD (left scale) and USD/Euro (right scale) over a representative one-year period. (From 12/31/99 to 12/31/2000.)

In a nutshell, a forward (and to a large degree futures) contract on a financial instrument imposes an obligation on the buyer (seller) to buy (sell) this instrument for a predetermined strike price at a predetermined maturity time. The settlement either occurs at maturity (forward contracts) or continuously (futures contracts). Theoretically, one of the parties has to pay an upfront premium in order to enter into a forward contract with a given strike. However, in practice, the strike is determined in such a way that the initial value of the claim is zero, so that entering into a forward contract does not require any initial investment from the parties involved. The party that buys (sells) an instrument forward expects its price to be above (below) the predetermined price at maturity. Broadly speaking, the forward and futures prices are expectations of the spot price of the corresponding instrument at maturity (with respect to appropriate probability measures). Finding the correct forward FXR is one of the basic problems of financial engineering. It is remarkable that this rate can be determined without knowing any details about the behavior of the underlying FXR *per se*.

Options on a financial instrument are different from forward and futures contracts because they give the buyer of an option the right to buy (call) or to sell (put) the underlying instrument and impose on the seller of an option the obligation to sell (call) or to buy (put) this instrument at a predetermined strike price at maturity (European options) or at any moment between the inception of the option and its maturity (American options). Since the buyer of an option receives a right and its seller accepts an obligation, options have to be purchased for an upfront premium with the subsequent settlement at (or before) maturity. The buyer of an option cannot lose more than the premium paid, while the losses of the seller can, in certain cases, be unlimited. The holder of a call (put) option will benefit from the rise (fall) of the price of the underlying instrument. Knowing prices of forward and futures contracts, and calls and puts with different strikes and maturities, we can obtain detailed probabilistic information about the future spot price of the underlying. Pricing and risk-managing derivatives is one of the most important objectives of financial engineering which we discuss in detail below. We emphasize that it is necessary for market makers to know fair prices of derivatives because they need to quote these prices (with bid-ask spread) before being told if they are going to be a seller or a buyer of the corresponding derivative. This situation is not dissimilar to the one discussed in the old mathematical problem: how two persons should divide a cake in such a way that each one receives a piece which is perceived to be greater or equal to half the cake. The answer is that the first person (the market maker) divides the cake in two parts, while the second person (the investor) chooses the part which he thinks is bigger. [1] The way the market maker earns a living is via the bid-ask spread. Under normal circumstances the fair price is sandwiched between the bid and ask prices. There are at least three reasons for an investor to buy a derivative instrument: (A) for protection against unfavorable forex changes; (B) for leveraging his market views; (C) for speculation.

We consider forward contracts and European calls and puts, which are called plain vanilla instruments, and denote the prices of these claims at time t for the spot FXR equal S by $FO(t,S,T,K)$, $C(t,S,T,K)$, $P(t,S,T,K)$, where the second pair of arguments emphasizes their dependence on the strike K and maturity T. When $t = T$ we have

$$FO(T,S,T,K) = S - K, \tag{1.2}$$

$$C(T,S,T,K) = \max\{S-K,0\} \equiv (S-K)_+, \tag{1.3}$$

[1] The reader is encouraged to solve a similar problem for N persons.

$$P(T, S, T, K) = \max\{K - S, 0\} \equiv (K - S)_+, \tag{1.4}$$

respectively. Here and below $x_+ = \max\{x, 0\}$. It is clear that for a forward contract the payoff can be both positive and negative while for calls and puts it is always positive, so that forwards can be both assets and liabilities while calls and puts are always assets for the buyer who cannot lose more than the option premium. For a forward contract, we need to find the strike K such that $FO(0, S, K, T) = 0$, while for calls and puts strikes are defined externally. Payoffs (1.2) - (1.4) are shown in Figure 1.2

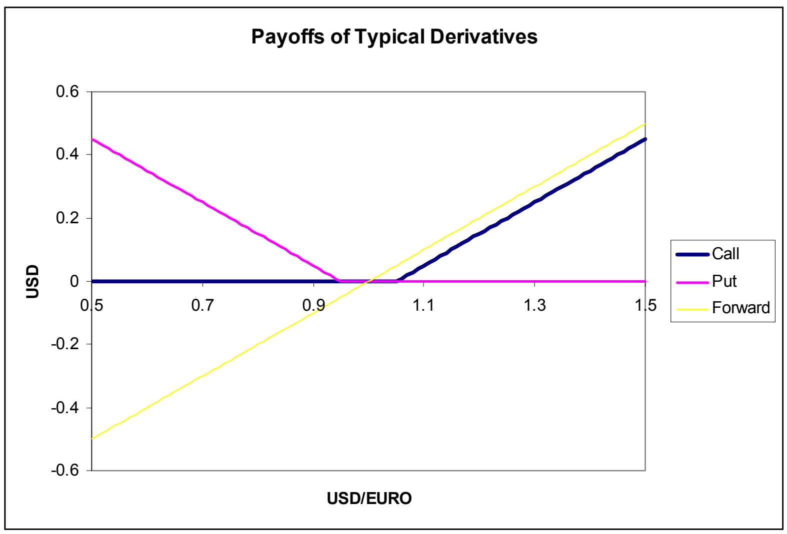

Figure 1.2: Representative payoffs of the call, put, and forward contract.

Even without knowing much about the dynamics of the underlying FXR we can say a lot about the valuation of forwards, calls and puts. First, we can find the fair forward strike K, which we denote by $F_{0,T}$. If we sell a forward contract, we agree to deliver one unit of foreign currency at maturity in exchange for $F_{0,T}$ units of domestic currency. To avoid the exposure to random fluctuation of the FXR, we have to buy one foreign zero coupon bond maturing at T. Accordingly, we have to borrow $e^{-r^1 T} S_0$ units of domestic currency which we need to repay in the amount of $e^{-r^1 T} S_0 / e^{-r^0 T}$ at time T. At maturity our bond produces one unit of foreign currency which we deliver to the buyer and receive $F_{0,T}$ units of domestic currency which we have to pay back to the lender of domestic currency, so that

$$F_{0,T} = e^{\left(r^0 - r^1\right)T} S_0. \tag{1.5}$$

This important relation, known as the *interest rate parity theorem,* establishes the fair forward FXR because of the perfect reversibility of the entire procedure which we used to establish it. (This reversibility is similar, in more than one respect, to the reversibility in thermodynamics.) We emphasize that the forward FXR depends only on the spot FXR, and domestic and foreign bond prices. A foreign currency is called a discount currency if $B^1_{0,T} < B^0_{0,T}$, and a premium currency otherwise. For discount currencies the forward FXR is lower than the spot one, while for premium currencies the opposite is true. The reader should think about this seemingly simple statement at some length because it is not obvious. For an arbitrary (nonequilibrium) strike, we can write

$$FO(t, S, T, K) = e^{-r^1(T-t)} S - e^{-r^0(T-t)} K.$$

Comparison of expressions (1.1), (1.5) shows that if a domestic investor uses the forward FXR rather than the spot FXR for repatriating his profits at time T, the rate of return on foreign investment becomes deterministic and vanishes. (No pain, no gain.)

We use the forward rate $F_{0,T}$ in order to classify calls and puts. We say that calls and puts with $K = F_{0,T}$ are at the money forward (ATMF), calls (puts) with $K < F_{0,T}$ ($K > F_{0,T}$) are in the money forward (ITMF), while calls (puts) with $K > F_{0,T}$ ($K < F_{0,T}$) are out of the money forward (OTMF).

To find the bounds for the price of calls and puts we can use the comparison principle for two portfolios which says that the portfolio which pays more at maturity should be more expensive at any time t before maturity. It is clear that the portfolio consisting of a European call with strike K and of K domestic bonds pays at maturity no less than a portfolio consisting of a foreign bond. This is true at any time t before maturity. At the same time the portfolio consisting of a call pays at maturity no more than the portfolio consisting of one unit of foreign currency. Accordingly, we can conclude that

$$\left(e^{-r^1(T-t)} S - e^{-r^0(T-t)} K\right)_+ \leq C(t, S, T, K) \leq S. \tag{1.6}$$

Similarly,

$$\left(e^{-r^0(T-t)} K - e^{-r^1(T-t)} S\right)_+ \leq P(t, S, T, K) \leq K. \tag{1.7}$$

It is interesting to note that in the special case when $e^{-r^1(T-t)} = 1$ (so that the foreign interest rates is equal to zero) we have

$$(S-K)_+ \leq C(t,S,T,K) \leq S,$$

i.e., the call price never falls below its intrinsic value. Similarly, when $e^{-r^0(T-t)} = 1$ (so that the domestic interest rates is equal to zero) the put price does not fall below its intrinsic value,

$$(K-S)_+ \leq P(t,S,T,K) \leq K.$$

It is easy to establish relations between European calls and puts. It follows from contract definition that at maturity we have

$$C(T,S,T,K) - P(T,S,T,K) = S - K = FO(T,S,T,K). \tag{1.8}$$

The pricing problem for European options is linear, i.e., any linear combination of its solutions is a solution, too. Accordingly, equation (1.8) is valid for any t, so that

$$C(t,S,T,K) - P(t,S,T,K) = e^{-r^1(T-t)}S - e^{-r^0(T-t)}K.$$

This relation is known as *put-call parity*.

In addition to put-call parity there exists one more important relation between calls and puts which we call put-call symmetry. The buyer of a call has the right to buy one euro (the notional amount) at the rate of K dollars/euro. Alternatively, he has the right to sell K dollars at the rate of $1/K$ euros/dollar. The right to buy euros is worth $C^0(0,S,T,K)$ dollars. (Here and below we use superscripts to show the currency in which the corresponding option is valued.) The right to sell dollars is worth $KP^1(0,1/S,T,1/K)$ euros. Since buying euros is equivalent to selling dollars the corresponding amounts have to coincide once they are expressed in the same currency (domestic, say). Accordingly,

$$C^0(0,S,T,K) = SKP^1\left(0,\frac{1}{S},T,\frac{1}{K}\right).$$

This relation expresses put-call symmetry. Since C^0 and P^1 are homogeneous functions of degree one, we can rewrite the above relation in the form

$$C^0(0,S,T,K) = P^1(0,K,T,S),$$

which is probably more elegant but definitely less useful. To switch from P^0 to P^1 one needs to interchange the domestic and foreign interest rates, $r^0 \to r^1$, $r^1 \to r^0$.

American calls and puts are determined by the same terminal payoffs as their European counterparts but can be executed at any time between the inception and maturity of the option; their prices are denoted by $C'(t,S,T,K)$, $P'(t,S,T,K)$. Since American options can be executed at any time t, $0 \leq t \leq T$, their prices cannot fall below their intrinsic values, additionally, they are always more expensive that the corresponding European options. Accordingly,

$$\max\{(S-K)_+, C(t,S,T,K)\} \leq C'(t,S,T,K) \leq S, \tag{1.9}$$

$$\max\{(K-S)_+, P(t,S,T,K)\} \leq P'(t,S,T,K) \leq K. \tag{1.10}$$

When $e^{-r^1(T-t)} = 1$, the prices of American and European calls coincide,

$$C(t,S,T,K) = C'(t,S,T,K),$$

when $e^{-r^0(T-t)} = 1$, the prices of American and European puts coincide,

$$P(t,S,T,K) = P'(t,S,T,K).$$

We cannot expect that *put-call parity* is preserved for American options since the early exercise feature destroys linearity. Instead, weaker inequalities can be proved, namely

$$e^{-r^1(T-t)}S - K \leq C^{'}(t,S,T,K) - P^{'}(t,S,T,K) \leq S - e^{-r^0(T-t)}K.$$

At the same time *put-call symmetry* is preserved since both the pricing equations and the constraints are invariant with respect to switching countries:

$$C'^0(0,S,T,K) = SKP'^1\left(0,\frac{1}{S},T,\frac{1}{K}\right).$$

This formula expresses the symmetry between a euro call (i.e. the right to buy euros) and a dollar put (i.e. the right to sell dollars). This relation becomes more transparent when we take into account the fact that dollars grow at the domestic risk-free rate r^0, while euros grow at the foreign risk-free rate r^1.

It is difficult to say more about option prices without constructing an adequate model for the behavior of S. We discuss different possibilities throughout the book.

Certain linear combinations of European calls and puts with the same maturity are particularly popular in the market place. These are strangles, straddles, and risk reversals which are discussed in Chapter 9.

In addition to options mentioned above, several others are important in the forex context. The reason is that European and American calls and puts can be rather expensive, so that their cheaper versions can be attractive to investors.

One popular way of making them cheaper, is to take a view on the distribution of the future FXR and to finance the purchase of one option by selling another one. For instance, an investor who thinks that the foreign currency will be traded above a certain level K_1 but below some other level K_2, where $K_1 < K_2$, can buy a call with strike K_1 and sell a call with strike K_2, thus creating the so-called call spreads with payoffs of the form

$$\text{payoff} = \begin{cases} 0, & S < K_1 \\ S - K_1, & K_1 \leq S < K_2 \\ K_2 - K_1, & K_2 \leq S \end{cases} .$$

Another approach is to add various barrier features to ordinary calls and puts. For instance, a down-and-out call (up-and-out put) disappears if the FXR hits a certain level below (above) the strike level. Even though the corresponding option provides less protection than its European or American analogue, it can be attractive to investors having specific market views, or those who can adjust to low (high) FXRs. Other options with mild barrier features which are frequently traded in practice are the so-called timers and faders (also known as time trades) with payoffs depending on the amount of time the FXR spends above (or below) a certain barrier.

In addition, there are many investors whose needs are better served by the so-called lookback calls (puts) which give the buyer the right to buy (sell) foreign currency at the best FXR observed between the inception and maturity of the corresponding option, and Asian options which give the right to buy or sell currency at the average FXR. We also study the so-called passport options. In contrast to all other options considered in this book, passport options are written on the trading account rather than on the underlying FXR. The buyer of such an options has the right to keep all the profit and is compensated for any loss generated by buying and selling a specific amount of foreign currency between the inception of the option and its maturity. Passport options require active participation of the buyer in the trading process, so that their value depends not only on the evolution of the FXR but also on the buyer's strategy. Barrier, lookback, Asian, and passport options are examples of the so-called path-dependent options whose payoff depends on the entire trajectory of the FXR between the inception and maturity of an option rather than on the terminal value of the FXR. Since it is more difficult to price and hedge such options, below we spend considerable effort in order to show how it can be done.

Yet another class of path-dependent options which gained popularity in the past few years includes the so-called volatility and variance swaps and options on realized volatility and variance. As their name suggests, these options are written on the volatility or variance of returns on foreign currency. For example, the buyer of a variance swap receives some fixed amount and has to pay the realized variance of returns on foreign currency multiplied by some notional amount. Since all other options depend on volatility, options on volatility are interesting not only in their own right but also as important hedging instruments for a portfolio of options.

1.6 References and further reading

The key original contributions mentioned in this chapter are as follows (in alphabetical order): Arrow (1971), Bachelier (1900), Black (1976), Black, Derman and Toy (1990), Black and Karasinski (1991), Black and Scholes (1973), Boness (1964), Boyle (1977), Brace, Gatarek and Musiela (1997), Brennan and Schwartz (1978), Cox, Ingersoll and Ross (1985), Cox, Ross and Rubinstein (1979), Debreu (1959), Delbaen and Schachermayer (1994), Garman and Kohlhagen (1983), Harrison and Kreps (1979), Harrison and Pliska (1981), Heath, Jarrow and Morton (1992), Ho and Lee (1986), Hull and White (1990), Lintner (1965), Margrabe (1978), Markowitz (1990), Merton (1973), Modigliani and Miller (1958), Mossin (1966), de Pinto (1771), Ross (1976), Rubinstein and Reiner (1991), Samuelson (1965), Schwartz (1977), Sharpe (1985), Sprenkle (1961), and Vasicek (1977). Details on empirical distributions of returns of forex and other asset classes are discussed by many authors, for example, by Duffie and Pan (1997). An overview of the subjects covered in this chapter can be found in Lipton-Lifschitz (1999), and Lipton (2000 a).

There are many (some would say too many) books on financial engineering. Here are the books which the present author finds particularly useful for his purposes: Avellaneda and Laurence (2000), Baxter and Rennie (1996), Dixit and Pindyck (1994), Dothan (1990), Duffie (1996), Cox and Rubinstein (1985), Fama and Miller (1972), Huang and Litzenberger (1988), Hull (2000), Hunt and Kennedy (2000), Ingersoll (1987), Jarrow and Rudd (1983), Karatzas and Shreve (1998), Lamberton and Lapeyre (1996), Luenberger (1998), Malkiel (1990), Merton (1990), Musiela and Rutkovsky (1997), Nielsen (2000), Pliska (1997), Shimko (1992), Shiryaev (1999), Wilmot *et al.* (1993), Wilmott *et al.* (1995), and Zang (1997).