CHAPTER 1

Foreign Exchange Derivatives

The FX derivatives market consists of FX swaps, FX forwards, FX or currency options, and other more general derivatives. FX structured products are either standardized or tailor-made linear combinations of simple FX derivatives including both vanilla and exotic options, or more general structured derivatives that cannot be decomposed into simple building blocks. The market for structured products is restricted to the market of the necessary ingredients. Hence, typically there are mostly structured products traded in the currency pairs that can be formed between USD, JPY, EUR, CHF, GBP, CAD and AUD. In this chapter we start with a brief history of options, followed by a technical section on vanilla options and volatility, and deal with commonly used linear combinations of vanilla options. Then we will illustrate the most important ingredients for FX structured products: the first and second generation exotics.

1.1 LITERATURE REVIEW

While there are tons of books on options and derivatives in general, very few are dedicated specifically to FX options. After the 2008 financial crisis, more such books appeared. Shamah [118] is a good source to learn about FX markets with a focus on market conventions, spot, forward, and swap contracts, and vanilla options. For pricing and modeling of exotic FX options I (obviously) suggest Hakala and Wystup's *Foreign Exchange Risk* [65] or its translation into Mandarin [68] as useful companions to this book. One of the first books dedicated to *Mathematical Models for Foreign Exchange* is by Lipton [92]. In 2010, Iain Clark published *Foreign Exchange Option Pricing* [28], and Antonio Castagna one on *FX Options and Smile Risk* [25], which both make a valuable contribution to the FX derivatives literature. A classic is Alan Hicks's *Managing Currency Risk Using Foreign Exchange Options* [76]. It provides a good overview of FX options mainly from the corporate's point of view. An introductory book on *Options on Foreign Exchange* is by DeRosa [38]. The *Handbook of Exchange Rates* [82] provides a comprehensive compilation of articles on the FX market structure, products, policies, and economic models.

1.2 A JOURNEY THROUGH THE HISTORY OF OPTIONS

The very first options and futures were traded in ancient Greece, when olives were sold before they had reached ripeness. Thereafter the market evolved in the following way.

16th century Ever since the 15th century, tulips, which were desired for their exotic appearance, were grown in Turkey. The head of the royal medical gardens in Vienna, Austria, was the first to cultivate those Turkish tulips successfully in Europe. When he fled to Holland because of religious persecution, he took the bulbs along. As the new head of the botanical gardens of Leiden, Netherlands, he cultivated several new strains. It was from these gardens that avaricious traders stole the bulbs to commercialize them, because tulips were a great status symbol.

17th century The first futures on tulips were traded in 1630. As of 1634, people could buy special tulip strains by the weight of their bulbs – the bulbs had the same value as gold. Along with the regular trading, speculators entered the market and the prices skyrocketed. A bulb of the strain, "Semper Octavian," was worth two wagonloads of wheat, four loads of rye, four fat oxen, eight fat swine, twelve fat sheep, two hogsheads of wine, four barrels of beer, two barrels of butter, 1,000 pounds of cheese, one marriage bed with linen, and one sizable wagon. People left their families, sold all their belongings, and even borrowed money to become tulip traders. When in 1637 this supposedly risk-free market crashed, traders as well as private individuals went bankrupt. The Dutch government prohibited speculative trading; the period became famous as Tulipmania.

18th century In 1728, the West India and Guinea Company, the monopolist in trading with the Caribbean Islands and the African coast, issued the first stock options. These were options on the purchase of the French island of Sainte-Croix, on which sugar plantings were planned. The project was realized in 1733 and paper stocks were issued in 1734. Along with the stock, people purchased a relative share of the island and the valuables, as well as the privileges and the rights of the company.

19th century In 1848, 82 businessmen founded the Chicago Board of Trade (CBOT). Today it is the biggest and oldest futures market in the entire world. Most written documents were lost in the great fire of 1871; however, it is commonly believed that the first standardized futures were traded as of 1860. CBOT now trades several futures and forwards, not only treasury bonds but also options and gold.

In 1870, the New York Cotton Exchange was founded. In 1880, the gold standard was introduced.

20th century

- In 1914, the gold standard was abandoned because of the First World War.
- In 1919, the Chicago Produce Exchange, in charge of trading agricultural products, was renamed the Chicago Mercantile Exchange. Today it is the most important futures market for the Eurodollar, foreign exchange, and livestock.
- In 1944, the Bretton Woods System was implemented in an attempt to stabilize the currency system.
- In 1970, the Bretton Woods System was abandoned for several reasons.

- In 1971, the Smithsonian Agreement on fixed exchange rates was introduced.
- In 1972, the International Monetary Market (IMM) traded futures on coins, currencies and precious metal.
- In 1973, the CBOE (Chicago Board of Exchange) firstly traded call options; four years later it added put options. The Smithsonian Agreement was abandoned; the currencies followed managed floating.
- In 1975, the CBOT sold the first interest rate future, the first future with no "real" underlying asset.
- In 1978, the Dutch stock market traded the first standardized financial derivatives.
- In 1979, the European Currency System was implemented, and the European Currency Unit (ECU) was introduced.
- In 1991, the Maastricht Treaty on a common currency and economic policy in Europe was signed.
- In 1999, the Euro was introduced, but the countries still used cash of their old currencies, while the exchange rates were kept fixed.

21st century In 2002, the Euro was introduced as new money in the form of cash.

FX forwards and options originate from the need of corporate treasury to hedge currency risk. This is the key to understanding FX options. Originally, FX options were not speculative products but hedging products. This is why they trade over the counter (OTC). They are tailored, i.e. cash flow matching currency risk hedging instruments for corporates. The way to think about an option is that a corporate treasurer in the EUR zone has income in USD and needs a hedge to sell the USD and to buy EUR for these USD. He would go long a forward or a EUR call option. At maturity he would exercise the option if it is in-the-money and receive EUR and pay USD. FX options are by default delivery settled. While FX derivatives were used later also as investment products or speculative instruments, the key to understanding FX options is corporate treasury.

1.3 CURRENCY OPTIONS

Let us start with a definition of a currency option:

Definition 1.3.1 *A* Currency Option Transaction *means a transaction entitling the Buyer, upon Exercise, to purchase from the Seller at the Strike Price a specified quantity of Call Currency and to sell to the Seller at the Strike Price a specified quantity of Put Currency.*

This is the definition taken from the *1998 FX and Currency Option Definitions* published by the International Swaps and Derivatives Association (ISDA) in 1998 [77]. This definition was the result of a process of standardization of currency options in the industry and is now widely accepted. Note that the key feature of an option is that the holder has a right to exercise. The definition also demonstrates clearly that calls and puts are equivalent, i.e. a call on one currency is always a put on the other currency. The definition is designed for a treasurer, where an actual cash flow of two currencies is triggered upon exercise. The definition also shows that the terms *derivative* and *option* are not synonyms. Derivative is a much wider term for financial transactions that depend on an underlying traded instrument. Derivatives include forwards, swaps, options, and exotic options. But not any derivative is also an option. For a currency option there is always a holder, the buyer after buying the option, equipped with the right to exercise, and upon exercise a cash flow of two pre-specified currencies is triggered. Anything outside this definition does not constitute a currency option. I highly recommend reading the 1998 ISDA definitions. The text uses legal language, but it does make all the terms around FX and currency options very clear and it is the benchmark in the industry. It covers only put and call options, options that are typically referred to as *vanilla* options, because they are the most common and simple products. The definition allows for different exercise styles: European for exercise permitted only at maturity, American for exercise permitted at any time between inception and maturity, as well as Bermudan for exercise permitted as finitely many pre-specified points in time. Usually, FX options are European. If you don't mention anything, they are understood to be of European exercise style. Features like cash settlement are possible: in this case one would have to make the call currency amount the net payoff and the put currency amount equal to zero. There are a number of exotic options, which we will cover later in this book, that still fit into this framework: in particular, barrier options. While they have special features not covered by the 1998 ISDA definitions, they still can be considered currency option transactions. However, variance swaps, volatility swaps, correlation swaps, combination of options, structured products, target forwards, just to mention a few obvious transactions, do not constitute currency option transactions.

1.4 TECHNICAL ISSUES FOR VANILLA OPTIONS

It is a standard in the FX options market to quote prices for FX options in terms of their implied volatility. The one-to-one correspondence between volatilities and options values rests on the convex payoff function of both call and put options. The conversion firmly rests on the Black-Scholes model. It is well known in the financial industry and academia that the Black-Scholes model has many weaknesses in modeling the underlying market properly. Strictly speaking, it is inappropriate. And there are in fact many other models, such as local volatility or stochastic volatility models or their hybrids, which reflect the dynamics much better than the Black-Scholes model. Nevertheless, as a basic tool to convert volatilities into values and values into volatilities, it is the market standard for dealers, brokers, and basically all risk management systems. This means: good news for those who have already learned it – it was not a waste of time and effort – and bad news for the quant-averse – you need to deal with it to a certain extent, as otherwise the FX volatility surface and the FX smile construction will not be accessible to you. Therefore, I do want to get the basic math done, even in this book, which I don't intend to be a quant book. However, I don't want to scare away much of my potential readership. If you don't like the math, you can still read most of this book.

We consider the model *geometric Brownian motion*

$$dS_t = (r_d - r_f)S_t\, dt + \sigma S_t\, dW_t \tag{1}$$

for the underlying exchange rate quoted in FOR-DOM (foreign-domestic), which means that one unit of the foreign currency costs FOR-DOM units of the domestic currency. In the case of EUR-USD with a spot of 1.2000, this means that the price of one EUR is 1.2000 USD. The notion of *foreign* and *domestic* does not refer to the location of the trading entity but only to this quotation convention. There are other terms used for FOR, which are *underlying*, *CCY1*, *base*; there are also other terms for DOM, which are *base*, *CCY2*, *counter* or *term*, respectively. For the quants, DOM is also considered the *numeraire* currency. I leave it to you to decide which one you wish to use. I find "base" a bit confusing, because it refers sometimes to FOR and sometimes to DOM. I also find "CCY1" and "CCY2" not very conclusive. The term "numeraire" does not have an established counterpart for FOR. So I prefer FOR and DOM. You may also stick to the most liquid currency pair EUR/USD, and think of FOR as EUR and DOM as USD.

We denote the (continuous) foreign interest rate by r_f and the (continuous) domestic interest rate by r_d. In an equity scenario, r_f would represent a continuous dividend rate. Note that r_f is *not* the interest rate that is typically used to discount cash flows in foreign currency, but is the (artificial) foreign interest rate that ensures that the forward price calculated in Equation (9) matches the market forward price. The volatility is denoted by σ, and W_t is a standard Brownian motion. The sample paths are displayed in Figure 1.1. We consider this standard model not because it reflects the statistical properties of the exchange rate (in fact, it doesn't) but because it is widely used in practice and front-office systems and mainly serves as a tool to

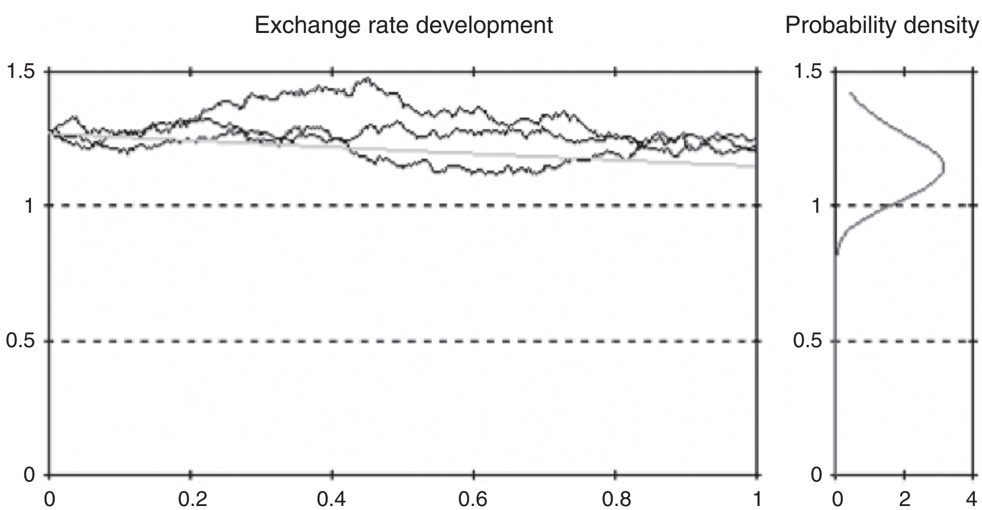

FIGURE 1.1 Simulated paths of a geometric Brownian motion. The distribution of the spot S_T at time T is log-normal. The light gray line reflects the average spot movement.

communicate prices of vanilla call and put options and switch between quotations in price and in terms of implied volatility. Currency option prices are commonly quoted in terms of volatility in the sense of this model. Model (1) is sometimes referred to as the Garman-Kohlhagen model [54]. However, all that happened there was adding the foreign interest rate r_f to the Black-Scholes model [15]. For this reason Model (1) is generally and in this book referred to as the Black-Scholes model.

Applying Itô's rule to $\ln S_t$ yields the following solution for the process S_t

$$S_t = S_0 \exp\left\{\left(r_d - r_f - \frac{1}{2}\sigma^2\right) t + \sigma W_t\right\}, \tag{2}$$

which shows that S_t is log-normally distributed, more precisely, $\ln S_t$ is normal with mean $\ln S_0 + (r_d - r_f - \frac{1}{2}\sigma^2)t$ and variance $\sigma^2 t$. Further model assumptions are:

1. There is no arbitrage.
2. Trading is frictionless, no transaction costs.
3. Any position can be taken at any time, short, long, arbitrary fraction, no liquidity constraints.

The payoff for a vanilla option (European put or call) is given by

$$F = [\phi(S_T - K)]^+, \tag{3}$$

where the contractual parameters are the strike K, the expiration time T and the type ϕ, a binary variable which takes the value +1 in the case of a call and −1 in the case of a put. The symbol x^+ denotes the positive part of x, i.e., $x^+ \stackrel{\Delta}{=} \max(0, x) \stackrel{\Delta}{=} 0 \vee x$. We generally use the symbol $\stackrel{\Delta}{=}$ to *define* a quantity. Most commonly, vanilla options on foreign exchange are of *European style*, i.e. the holder can only exercise the option at time T. *American style options*, where the holder can exercise any time, or *Bermudan style options*, where the holder can exercise at selected times, are not used very often except for *time options*, see Section 2.1.19.

1.4.1 Valuation in the Black-Scholes Model

In the Black-Scholes model the value of the payoff F at time t if the spot is at x is denoted by $v(t, x)$ and can be computed either as the solution of the *Black-Scholes partial differential equation* (see [15])

$$v_t - r_d v + (r_d - r_f) x v_x + \frac{1}{2}\sigma^2 x^2 v_{xx} = 0, \tag{4}$$

$$v(T, x) = F \tag{5}$$

or equivalently (*Feynman-Kac Theorem*) as the discounted expected value of the payoff-function

$$v(x, K, T, t, \sigma, r_d, r_f, \phi) = e^{-r_d \tau} \mathbb{E}[F]. \tag{6}$$

This is the reason why basic financial engineering is mostly concerned with solving partial differential equations or computing expectations (numerical integration). The result is the *Black-Scholes formula*

$$v(x, K, T, t, \sigma, r_d, r_f, \phi) = \phi e^{-r_d\tau}[f\mathcal{N}(\phi d_+) - K\mathcal{N}(\phi d_-)]. \tag{7}$$

The result of this formula is the value of a vanilla option in USD for one unit of EUR nominal. We abbreviate

$$x : \text{current price of the underlying,}$$

$$\tau \stackrel{\Delta}{=} T - t\text{: time to maturity,} \tag{8}$$

$$f \stackrel{\Delta}{=} \mathbb{E}[S_T|S_t = x] = xe^{(r_d - r_f)\tau}\text{: forward price of the underlying,} \tag{9}$$

$$\theta_\pm \stackrel{\Delta}{=} \frac{r_d - r_f}{\sigma} \pm \frac{\sigma}{2}, \tag{10}$$

$$d_\pm \stackrel{\Delta}{=} \frac{\ln\frac{x}{K} + \sigma\theta_\pm\tau}{\sigma\sqrt{\tau}} = \frac{\ln\frac{f}{K} \pm \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}, \tag{11}$$

$$n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}t^2} = n(-t) \text{ normal density,} \tag{12}$$

$$\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\,dt = 1 - \mathcal{N}(-x) \text{ normal distribution function.} \tag{13}$$

We observe that some authors use d_1 for d_+ and d_2 for d_-, which requires extra memory and completely ruins the beautiful symmetry of the formula.

The Black-Scholes formula can be derived using the integral representation of Equation (6)

$$\begin{aligned} v &= e^{-r_d\tau}\mathbb{E}[F] \\ &= e^{-r_d\tau}\mathbb{E}[[\phi(S_T - K)]^+] \\ &= e^{-r_d\tau}\int_{-\infty}^{+\infty}\left[\phi\left(xe^{(r_d - r_f - \frac{1}{2}\sigma^2)\tau + \sigma\sqrt{\tau}y} - K\right)\right]^+ n(y)\,dy. \end{aligned} \tag{14}$$

Next one has to deal with the positive part and then complete the square to get the Black-Scholes formula. A derivation based on the partial differential equation can be done using results about the well-studied *heat equation*. For valuation of options it is very important to ensure that the interest rates are chosen such that the forward price (9) matches the market, as otherwise the options may not satisfy the put-call parity (41).

1.4.2 A Note on the Forward

The *forward price* f is the pre-agreed exchange rate which makes the time zero value of the *forward contract* with payoff

$$F = S_T - f \tag{15}$$

equal to zero. It follows that $f = \mathbb{E}[S_T] = xe^{(r_d - r_f)T}$, i.e. the forward price is the expected price of the underlying at time T in a risk-neutral measure (drift of the geometric Brownian motion is equal to cost of carry $r_d - r_f$). The situation $r_d > r_f$ is called *contango*, and the situation $r_d < r_f$ is called *backwardation*. Note that in the Black-Scholes model the class of forward price curves is quite restricted. For example, no seasonal effects can be included. Note that the post-trade value of the forward contract after time zero is usually different from zero, and since one of the counterparties is always short, there may be settlement risk of the short party. A *futures contract* prevents this dangerous affair: it is basically a forward contract, but the counterparties have to maintain a *margin account* to ensure the amount of cash or commodity owed does not exceed a specified limit.

1.4.3 Vanilla Greeks in the Black-Scholes Model

Greeks are derivatives of the value function with respect to model and contract parameters. They are important information for traders and have become standard information provided by front-office systems. More details on Greeks and the relations among Greeks are presented in Hakala and Wystup [65] or Reiss and Wystup [107]. Initially there was a desire to use Greek letters for all these mathematical derivatives. However, it turned out that since the early days of risk management many higher order Greeks have been added whose terms no longer reflect Greek letters. Even vega is not a Greek letter but we needed a Greek sounding term that starts with a "v" to reflect volatility and Greek doesn't have such a letter. For vanilla options we list some of them now.

(Spot) Delta.

$$\frac{\partial v}{\partial x} = \phi e^{-r_f \tau} \mathcal{N}(\phi d_+) \tag{16}$$

This spot delta ranges between 0% and a *discounted* ±100%. The interpretation of this quantity is the amount of FOR the trader needs to buy to delta hedge a short option. So for instance, if you sell a call on 1 M EUR, that has a 25% delta, you need to buy 250,000 EUR to delta hedge the option. The corresponding forward delta ranges between 0% and ±100% and is symmetric in the sense that a 60-delta call is a 40-delta put, a 75-delta put is a 25-delta call, etc. I had wrongly called it "driftless delta" in the first edition of this book.

Forward Delta.

$$\phi \mathcal{N}(\phi d_+) \tag{17}$$

The interpretation of forward delta is the number of units of FOR of forward contracts a trader needs to buy to delta hedge a short option. See Section 1.4.7 for a justification.

Future Delta.

$$\phi e^{-r_d \tau} \mathcal{N}(\phi d_+) \tag{18}$$

Gamma.

$$\frac{\partial^2 v}{\partial x^2} = e^{-r_f \tau} \frac{n(d_+)}{x\sigma\sqrt{\tau}} \tag{19}$$

The interpretation of gamma is the change of delta as spot changes. A high gamma means that the delta hedge must be adapted very frequently and will hence cause transaction costs. Gamma is typically high when the spot is near a strike of a barrier, generally whenever the payoff has a kink or more dramatically a jump. Trading systems usually quote gamma as a *traders' gamma*, using a 1% *relative* change in the spot price. For example, if gamma is quoted as 10,000 EUR, then delta will increase by that amount if the spot rises from 1.3000 to 1.3130 = 1.3000 · (1 + 1%). This can be approximated by $\frac{\partial^2 v}{\partial x^2} \cdot \frac{x}{100}$.

Speed.

$$\frac{\partial^3 v}{\partial x^3} = -e^{-r_f \tau} \frac{n(d_+)}{x^2\sigma\sqrt{\tau}} \left(\frac{d_+}{\sigma\sqrt{\tau}} + 1 \right) \tag{20}$$

The interpretation of speed is the change of gamma as spot changes.

Theta.

$$\begin{aligned} \frac{\partial v}{\partial t} = & -e^{-r_f \tau} \frac{n(d_+) x\sigma}{2\sqrt{\tau}} \\ & + \phi[r_f x e^{-r_f \tau} \mathcal{N}(\phi d_+) - r_d K e^{-r_d \tau} \mathcal{N}(\phi d_-)] \end{aligned} \tag{21}$$

Theta reflects the change of the option value as the clock ticks. The *traders' theta* that you spot in a risk management system usually refers to a change of the option value in one day, i.e. the traders' theta can be approximated by $365 \frac{\partial v}{\partial t}$.

Charm.

$$\frac{\partial^2 v}{\partial x \partial \tau} = -\phi r_f e^{-r_f \tau} \mathcal{N}(\phi d_+) + \phi e^{-r_f \tau} n(d_+) \frac{2(r_d - r_f)\tau - d_- \sigma\sqrt{\tau}}{2\tau\sigma\sqrt{\tau}} \tag{22}$$

Color.

$$\frac{\partial^3 v}{\partial x^2 \partial \tau} = -e^{-r_f \tau} \frac{n(d_+)}{2x\tau\sigma\sqrt{\tau}} \left[2r_f \tau + 1 + \frac{2(r_d - r_f)\tau - d_- \sigma\sqrt{\tau}}{2\tau\sigma\sqrt{\tau}} d_+ \right] \tag{23}$$

Vega.

$$\frac{\partial v}{\partial \sigma} = xe^{-r_f\tau}\sqrt{\tau}n(d_+) \tag{24}$$

Trading and risk management systems usually quote vega as a *traders' vega*, using a 1% *absolute* change in the volatility. For example, if vega is quoted as 4,000 EUR, then the option value will increase by that amount if the volatility rises from 10% to 11% = 10% + 1%. This can be approximated by $\frac{\partial v}{\partial \sigma} \cdot 100$.

Volga.

$$\frac{\partial^2 v}{\partial \sigma^2} = xe^{-r_f\tau}\sqrt{\tau}n(d_+)\frac{d_+d_-}{\sigma} \tag{25}$$

Volga is also sometimes called *vomma* or *volgamma* or *dvega/dvol*. Volga reflects the change of vega as volatility changes. Traders' volga assumes again a 1% absolute change in volatility.

Vanna.

$$\frac{\partial^2 v}{\partial \sigma \partial x} = -e^{-r_f\tau}n(d_+)\frac{d_-}{\sigma} \tag{26}$$

Vanna is also sometimes called *dvega/dspot*. It reflects the change of vega as spot changes. Traders' vanna assumes again a 1% relative change in spot. The origin of the term vanna is not clear. I suspect it goes back to an article in *Risk* magazine by Tim Owens in the 1990s, where he asked "Wanna lose a lot of money?" and then explained how a loss may occur if second order Greeks such as vanna and volga are not hedged.

Volunga.

$$\frac{\partial^3 v}{\partial \sigma^3} = \frac{\text{vega}}{\sigma^2}\left((d_+d_-)^2 - d_+^2 - d_+d_- - d_-^2\right) \tag{27}$$

This is actually not a joke.

Vanunga.

$$\frac{\partial^3 v}{\partial x \partial \sigma^2} = \frac{\text{vega}}{\sigma^2 x\sqrt{\tau}}\left(d_+ + d_+d_- - d_+d_-^2\right) \tag{28}$$

This one isn't a joke either.

Rho.

$$\frac{\partial v}{\partial r_d} = \phi K\tau e^{-r_d\tau}\mathcal{N}(\phi d_-) \tag{29}$$

$$\frac{\partial v}{\partial r_f} = -\phi x\tau e^{-r_f\tau}\mathcal{N}(\phi d_+) \tag{30}$$

Trading and risk management systems usually quote rho as a *traders' rho*, using a 1% *absolute* change in the interest rate. For example, if rho is quoted as 4,000 EUR, then the option value will increase by that amount if the interest rate rises from 2% to 3% = 2% + 1%. This can be approximated by $\frac{\partial v}{\partial \rho} \cdot 100$. **Warning:** FX options always involve two currencies. Therefore, there will be two interest rates, a domestic interest rate r_d, and a foreign interest rate r_f. The value of the option can be represented in both DOM and FOR units. This means that you can have a change of the option value in FOR as the FOR rate changes, a change of the value of the option in FOR as the DOM rate changes, a change of the value of the option in DOM as the FOR rate changes, and a change of the value of the option in DOM as the DOM rate changes. Some systems add to the confusion as they list one rho, which refers to the change of the option value as the *difference of the interest rates* changes, and again possibly in both DOM and FOR terms.

Dual Delta.

$$\frac{\partial v}{\partial K} = -\phi e^{-r_d \tau} \mathcal{N}(\phi d_-) \tag{31}$$

The non-discounted version of the dual delta, also referred to as the forward dual delta, also represents the risk-neutral exercise probability of the option.

Dual Gamma.

$$\frac{\partial^2 v}{\partial K^2} = e^{-r_d \tau} \frac{n(d_-)}{K\sigma\sqrt{\tau}} \tag{32}$$

Dual Theta.

$$\frac{\partial v}{\partial T} = -v_t \tag{33}$$

Dual Greeks refer to changes of the option value as contractual parameters change. This has no application in market risk management, because the contractual parameters are fixed between counterparts and cannot be changed on the way. However, the dual Greeks contribute a lot to understanding of derivatives. The dual gamma (on the strike space) for example – up to a discount factor – is identical to the probability density of the underlying exchange rate.

1.4.4 Reoccurring Identities

$$\frac{\partial d_\pm}{\partial \sigma} = -\frac{d_\mp}{\sigma} \tag{34}$$

$$\frac{\partial d_\pm}{\partial r_d} = \frac{\sqrt{\tau}}{\sigma} \tag{35}$$

$$\frac{\partial d_\pm}{\partial r_f} = -\frac{\sqrt{\tau}}{\sigma} \tag{36}$$

$$xe^{-r_f\tau}n(d_+) = Ke^{-r_d\tau}n(d_-) \tag{37}$$

$$\mathcal{N}(\phi d_-) = I\!P[\phi S_T \geq \phi K] \tag{38}$$

$$\mathcal{N}(\phi d_+) = I\!P\left[\phi S_T \leq \phi \frac{f^2}{K}\right] \tag{39}$$

The *put-call parity* is a way to express the trivial equation $x = x^+ - x^-$ in financial terms and is the relationship on the payoff level

$$\text{call} - \text{put} = \text{forward}$$
$$(S_T - K)^+ - (K - S_T)^+ = S_T - K, \tag{40}$$

which translates to the value functions of these products via

$$v(x, K, T, t, \sigma, r_d, r_f, +1) - v(x, K, T, t, \sigma, r_d, r_f, -1) = xe^{-r_f\tau} - Ke^{-r_d\tau}. \tag{41}$$

A forward contract that is constructed using a long call and a short put option is called a *synthetic forward*.

The *put-call delta parity* is

$$\frac{\partial v(x, K, T, t, \sigma, r_d, r_f, +1)}{\partial x} - \frac{\partial v(x, K, T, t, \sigma, r_d, r_f, -1)}{\partial x} = e^{-r_f\tau}. \tag{42}$$

In particular, we learn that the absolute values of a spot put delta and a spot call delta are not exactly adding up to 100%, but only to a positive number $e^{-r_f\tau}$. They add up to one approximately if either the time to expiration τ is short or if the foreign interest rate r_f is close to zero. The corresponding forward deltas do add up to 100%.

Whereas the choice $K = f$ produces identical values for call and put, we seek the *delta-symmetric strike* or *delta-neutral strike* K_+ which produces absolutely identical deltas (spot, forward or future). This condition implies $d_+ = 0$ and thus

$$K_+ = fe^{+\frac{\sigma^2}{2}\tau}, \tag{43}$$

in which case the absolute spot delta is $e^{-r_f\tau}/2$. In particular, we learn that always $K_+ > f$, i.e., there can't be a put and a call with identical values *and* deltas. Note that the strike K_+ is usually chosen as the middle strike when trading a straddle or a butterfly. Similarly the dual-delta-symmetric strike $K_- = fe^{-\frac{\sigma^2}{2}T}$ can be derived from the condition $d_- = 0$. Note that the delta-symmetric strike K_+ also maximizes gamma and vega of a vanilla option and is thus often considered a center of symmetry.

1.4.5 Homogeneity based Relationships

We may wish to measure the value of the underlying in a different unit. This will obviously affect the option pricing formula as follows:

$$av(x, K, T, t, \sigma, r_d, r_f, \phi) = v(ax, aK, T, t, \sigma, r_d, r_f, \phi) \text{ for all } a > 0. \tag{44}$$

Differentiating both sides with respect to a and then setting $a = 1$ yields

$$v = xv_x + Kv_K. \tag{45}$$

Comparing the coefficients of x and K in Equations (7) and (45) leads to suggestive results for the delta v_x and dual delta v_K. This *space-homogeneity* is the reason behind the simplicity of the delta formulas, whose tedious computation can be saved this way.

Time Homogeneity We can perform a similar computation for the time-affected parameters and obtain the obvious equation

$$v(x, K, T, t, \sigma, r_d, r_f, \phi) = v\left(x, K, \frac{T}{a}, \frac{t}{a}, \sqrt{a}\sigma, ar_d, ar_f, \phi\right) \text{ for all } a > 0. \tag{46}$$

Differentiating both sides with respect to a and then setting $a = 1$ yields

$$0 = \tau v_t + \frac{1}{2}\sigma v_\sigma + r_d v_{r_d} + r_f v_{r_f}. \tag{47}$$

Of course, this can also be verified by direct computation. The overall use of such equations is to generate double checking benchmarks when computing Greeks. These homogeneity methods can easily be extended to other more complex options.

Put-Call Symmetry By *put-call symmetry* we understand the relationship (see [9, 10, 19] and [24])

$$v(x, K, T, t, \sigma, r_d, r_f, +1) = \frac{K}{f} v\left(x, \frac{f^2}{K}, T, t, \sigma, r_d, r_f, -1\right). \tag{48}$$

The strike of the put and the strike of the call result in a geometric mean equal to the forward f. The forward can be interpreted as a *geometric mirror* reflecting a call into a certain number of puts. Note that for at-the-money options ($K = f$) the put-call symmetry coincides with the special case of the put-call parity where the call and the put have the same value.

Rates Symmetry Direct computation shows that the *rates symmetry*

$$\frac{\partial v}{\partial r_d} + \frac{\partial v}{\partial r_f} = -\tau v \tag{49}$$

holds for vanilla options. In fact, this relationship holds for all European options and a wide class of path-dependent options as shown in [107].

Foreign-Domestic Symmetry One can directly verify the *foreign-domestic symmetry* as relationship

$$\frac{1}{x}v(x, K, T, t, \sigma, r_d, r_f, \phi) = Kv\left(\frac{1}{x}, \frac{1}{K}, T, t, \sigma, r_f, r_d, -\phi\right). \tag{50}$$

This equality can be viewed as one of the faces of put-call symmetry. The reason is that the value of an option can be computed in units of domestic currency as well as in units of foreign currency. We consider the example of S_t modeling the exchange rate of EUR/USD. In New York, the call option $(S_T - K)^+$ costs $v(x, K, T, t, \sigma, r_{usd}, r_{eur}, 1)$ USD and hence $v(x, K, T, t, \sigma, r_{usd}, r_{eur}, 1)/x$ EUR. This EUR-call option can also be viewed as a USD-put option with payoff $K\left(\frac{1}{K} - \frac{1}{S_T}\right)^+$. This option costs $Kv(\frac{1}{x}, \frac{1}{K}, T, t, \sigma, r_{eur}, r_{usd}, -1)$ EUR in Frankfurt, because S_t and $\frac{1}{S_t}$ have the same volatility. Of course, the New York value and the Frankfurt value must agree, which leads to (50). We will also learn later that this symmetry is just one possible result based on *change of numeraire*.

1.4.6 Quotation Conventions

Quotation of the Underlying Exchange Rate Equation (1) is a model for the exchange rate. The quotation is a permanently confusing issue, so let us clarify this here. The exchange rate means how many units of the *domestic* currency are needed to buy one unit of *foreign* currency. For example, if we take EUR/USD as an exchange rate, then the default quotation is EUR-USD, where USD is the domestic currency and EUR is the foreign currency. The term *domestic* is in no way related to the location of the trader or any country. It merely means the *numeraire* currency. The terms *domestic*, *numeraire*, *currency two* or *base currency* are synonyms, as are *foreign*, *currency one* and *underlying*. Some market participants even refer to the foreign currency as the base currency, one of the reasons why I prefer to avoid the term base currency altogether. Throughout this book we denote with the slash (/) the currency pair and with a dash (-) the quotation. The slash (/) does *not* mean a division. For instance, EUR/USD can also be quoted in either EUR-USD, which then means how many USD are needed to buy one EUR, or in USD-EUR, which then means how many EUR are needed to buy one USD. There are certain market standard quotations listed in Table 1.1.

Trading Floor Language We call one million a *buck*, one billion a *yard*. This is because a billion is called “milliarde” in French, German and other languages. For the British pound one million is also often called a *quid*.

TABLE 1.1 Standard market quotation of major currency pairs with sample spot prices.

Currency pair	Default quotation	Sample quote
GBP/USD	GBP-USD	1.6000
GBP/CHF	GBP-CHF	2.2500
EUR/USD	EUR-USD	1.3000
EUR/GBP	EUR-GBP	0.8000
EUR/JPY	EUR-JPY	135.00
EUR/CHF	EUR-CHF	1.2000
USD/JPY	USD-JPY	108.00
USD/CHF	USD-CHF	1.0100

Certain currencies also have names, e.g. the New Zealand dollar NZD is called a *Kiwi*, the Australian dollar AUD is called *Aussie*, the Canadian dollar CAD is called *Loonie*, the Scandinavian currencies DKK, NOK (*Nokkies*) and SEK (*Stockies*) are collectively called *Scandies*.

Exchange rates are generally quoted up to five relevant figures, e.g. in EUR-USD we could observe a quote of 1.2375. The last digit "5" is called the *pip*, the middle digit "3" is called the *big figure*, as exchange rates are often displayed in trading floors and the big figure, which is displayed in bigger size, is the most relevant information. The digits left of the big figure are known anyway. If a trader doesn't know these when getting to the office in the morning, he may most likely not have the right job. The pips right of the big figure are often negligible for general market participants of other asset classes and are highly relevant only for currency spot traders. To make it clear, a rise of USD-JPY 108.25 by 20 pips will be 108.45 and a rise by 2 big figures will be 110.25.

Cable Currency pairs are often referred to by nicknames. The price of one pound sterling in US dollars, denoted by GBP/USD, is known by traders as the *cable*, which originates from the time when a communications cable under the Atlantic Ocean synchronized the GBP/USD quote between the London and New York markets. So where is the cable?

I stumbled upon a small town called Porthcurno near Land's End on the southwestern Cornish coast and by mere accident spotted a small hut called the "cable house" admittedly a strange object to find on a beautiful sandy beach. Trying to find Cornish cream tea I ended up at a telegraphic museum, which had all I ever wanted to know about the cable (see the photographs in Figure 1.2). Telegraphic news transmission was introduced in 1837, typically along the railway lines. Iron was rapidly replaced by copper. A new insulating material, gutta-percha, which is similar to rubber, allowed cables to function under the sea, and as Britain neared the height of its international power, submarine cables started to be laid, gradually creating a global network of cables, which included the first long-term successful trans-Atlantic cable of 1865 laid by the Great Eastern ship.

FIGURE 1.2 The Cable at Porthcurno, in the telegraphic museum and on the beach near the cable house.

The entrepreneur of that age was John Pender, founder of the Eastern Telegraph Company. He had started as a cotton trader and needed to communicate quickly with various ends of the world. In the 1860s telegraphic messaging was the new and only way to do this. Pender quickly discovered the value of fast communication. In the 1870s, an annual traffic of around 200,000 words went through Porthcurno. By 1900, cables connected Porthcurno with India (via Gibraltar and Malta), Australia and New Zealand. The cable network charts of the late 1800s reflect the financial trading centers of today very closely: Tokyo, Sydney, Singapore, Mumbai, London, New York.

Fast communication is ever so important for the financial industry. You can still go to Porthcurno and touch the cables. They have been in the sea for more than 100 years, but they still work. However, they have been replaced by fiber glass cables, and communications have been extended by radio and satellites. Algorithmic trading relies on getting all the market information within milliseconds.

The word “cable” itself is still used as the GBP/USD rate, reflecting the importance of fast information.

Crosses Currency pairs not involving the USD such as EUR/JPY are called a *cross* because it is the cross rate of the more liquidly traded USD/JPY and EUR/USD. If the cross is illiquid, such as ILS/MYR, it is called an illiquid cross. Spot transactions would then happen in two steps via USD. Options on an illiquid cross are rare or traded at very high bid-offer spreads.

Quotation of Option Prices Values and prices of vanilla options may be quoted in the six ways explained in Table 1.2.

TABLE 1.2 Standard market quotation types for option values. In the example we take FOR = EUR, DOM = USD, $S_0 = 1.2000$, $r_d = 3.0\%$, $r_f = 2.5\%$, $\sigma = 10\%$, $K = 1.2500$, $T = 1$ year, $\phi = +1$ (call), notional = 1,000,000 EUR = 1,250,000 USD. For the pips, the quotation 291.48 USD pips per EUR is also sometimes stated as 2.9148% USD per 1 EUR. Similarly, the 194.32 EUR pips per USD can also be quoted as 1.9432% EUR per 1 USD.

Name	Symbol	Value in units of	Example
domestic cash	**d**	DOM	29,148 USD
foreign cash	**f**	FOR	24,290 EUR
% domestic	**% d**	DOM per unit of DOM	2.3318% USD
% foreign	**% f**	FOR per unit of FOR	2.4290% EUR
domestic pips	**d pips**	DOM per unit of FOR	291.48 USD pips per EUR
foreign pips	**f pips**	FOR per unit of DOM	194.32 EUR pips per USD

The Black-Scholes formula quotes **d pips**. The others can be computed using the following instruction.

$$\textbf{d pips} \overset{\times\frac{1}{S_0}}{\longrightarrow} \%f \overset{\times\frac{S_0}{K}}{\longrightarrow} \%d \overset{\times\frac{1}{S_0}}{\longrightarrow} \textbf{f pips} \overset{\times S_0 K}{\longrightarrow} \textbf{d pips} \tag{51}$$

Delta and Premium Convention The spot delta of a European option assuming the premium is paid in DOM is well known. It will be called *raw spot delta* δ_{raw} now. It can be quoted in either of the two currencies involved. The relationship is

$$\delta_{raw}^{reverse} = -\delta_{raw}\frac{S}{K}. \tag{52}$$

The delta is used to buy or sell spot in the corresponding amount in order to hedge the option up to first order. The raw spot delta, multiplied by the FOR nominal amount, represents the amount of FOR currency the trader needs to buy in order to delta hedge a short option. How do we get to the reverse delta? It rests firmly on the symmetry of currency options. A FOR call is a DOM put. Hence, buying FOR amount in the delta hedge is equivalent to selling DOM amount multiplied by the spot S. The negative sign reflects the change from buying to selling. This explains the negative sign and the spot factor. A right to buy 1 FOR (and pay for this K DOM) is equivalent to the right to sell K DOM and receive for that 1 DOM. Therefore, viewing the FOR call as a DOM put and applying the delta hedge to one unit of DOM (instead of K units of DOM) requires a division by K. Now read this paragraph again and again and again, until it clicks. Sorry.

For consistency the premium needs to be incorporated into the delta hedge, since a premium in foreign currency will already hedge part of the option's delta risk. In a stock options context such a question never comes up, as an option on a stock is always paid in cash, rather than paid in shares of stock. In foreign exchange, both currencies are cash, and it is perfectly reasonable to pay for a currency option in either DOM or FOR currency. To make this clear, let us consider EUR-USD. In any financial markets model, $v(x)$ denotes the value or premium in USD of an option with 1 EUR notional, if the spot is at x, and the raw delta v_x denotes the number of EUR to buy to delta hedge a short position of this option. If this raw delta is negative, then EUR have to be sold (silly but hopefully helpful remark for the non-math freak). Therefore, xv_x is the number of USD to sell. If now the premium is paid in EUR rather than in USD, then we already have $\frac{v}{x}$ EUR, and the number of EUR to buy has to be reduced by this amount, i.e. if EUR is the premium currency, we need to buy $v_x - \frac{v}{x}$ EUR for the delta hedge or equivalently sell $xv_x - v$ USD. This is called a *premium-adjusted delta* or delta with premium included.

The same result can be derived by looking at the risk management of a portfolio whose accounting currency is EUR and risky currency is USD. In this case spot is $\frac{1}{x}$ rather than x. The value of the option – or in fact more generally of a portfolio of derivatives – is then $v\left(\frac{1}{x}\right)$ in USD, and $v\left(\frac{1}{x}\right)\frac{1}{x}$ in EUR, and the change of the portfolio value in EUR as the price of the USD measured in EUR is

$$\begin{aligned}
\frac{\partial}{\partial \frac{1}{x}} \frac{v\left(\frac{1}{x}\right)}{x} &= \frac{\partial}{\partial x} \frac{v\left(\frac{1}{x}\right)}{x} \frac{\partial x}{\partial \frac{1}{x}} \\
&= \frac{v_x\left(\frac{1}{x}\right)x - v\left(\frac{1}{x}\right)}{x^2}\left(\frac{\partial \frac{1}{x}}{\partial x}\right)^{-1} \\
&= \frac{xv_x - v}{x^2}\left(-\frac{1}{x^2}\right)^{-1} \\
&= -[xv_x - v]. \qquad (53)
\end{aligned}$$

We observe that both the trader's approach deriving delta from the premium and the risk manager's approach deriving delta from the portfolio risk arrive at the same number. Not really a surprise, is it?

The premium-adjusted delta for a vanilla option in the Black-Scholes model becomes

$$\begin{aligned}
-[xv_x - v] &= -\left[\phi x e^{-r_f\tau}\mathcal{N}(\phi d_+) - \phi[xe^{-r_f\tau}\mathcal{N}(\phi d_+) - Ke^{-r_d\tau}\mathcal{N}(\phi d_-)]\right] \\
&= -\phi K e^{-r_d\tau}\mathcal{N}(\phi d_-) \qquad (54)
\end{aligned}$$

in USD, or $-\phi e^{-r_d\tau}\frac{K}{x}\mathcal{N}(\phi d_-)$ in EUR. If we sell USD instead of buying EUR, and if we assume a notional of 1 USD rather than 1 EUR (= K USD) for the option, the premium-adjusted delta becomes just

$$\phi e^{-r_d\tau}\mathcal{N}(\phi d_-). \qquad (55)$$

If you ever wondered why delta uses $\mathcal{N}(d_+)$ and not $\mathcal{N}(d_-)$, which is really not fair, you now have an answer: both these terms are deltas, and only the FX market can really explain what's going on:

- $\phi e^{-r_f\tau}\mathcal{N}(\phi d_+)$ is the delta if the premium is paid in USD,
- $\phi e^{-r_d\tau}\mathcal{N}(\phi d_-)$ is the delta if the premium is paid in EUR.

In FX options markets there is no preference between the two, as a premium can always (well, normally always) be paid in either currency. The premium-adjusted delta is therefore also related to the dual delta (31).

Default Premium Currency Quotations in FX require some patience because we need to first sort out which currency is domestic, which is foreign, what is the notional currency of the option, and what is the premium currency. Unfortunately this is not symmetric, since the counterpart might have another notion of domestic currency for a given currency pair. Hence in the professional inter bank market there is a generic notion of delta per currency pair. Table 1.3 provides a short overview. Details on all currency pairs can be found in your risk management system (if you have a good one). Essentially there are only four currency pairs with a premium paid in domestic currency by default. All other pairs use premium adjustment.

Example of Delta Quotations We consider two examples in Table 1.4 and Table 1.5 to compare the various versions of deltas that are used in practice.

TABLE 1.3 Default premium currency for a small selection of currency pairs. LHS currency pairs assume premium paid in USD (domestic currency), RHS assume premium paid in foreign currency.

Premium-unadjusted	Premium-adjusted
EUR/USD	USD/CAD
GBP/USD	EUR/GBP
AUD/USD	USD/JPY
NZD/USD	EUR/JPY
	USD/BRL
	USD/CHF
	EUR/CHF
	USD/ILS
	USD/SGD
	EUR/TRY

TABLE 1.4 1y EUR call USD put strike $K = 0.9090$ for a EUR-based bank. Market data: spot $S = 0.9090$, volatility $\sigma = 12\%$, EUR rate $r_f = 3.96\%$, USD rate $r_d = 3.57\%$. The raw delta is 49.15% EUR and the value is 4.427% EUR.

Delta ccy	Prem ccy	FENICS	Formula	Delta
% EUR	EUR	lhs	$\delta_{raw} - P$	44.72
% EUR	USD	rhs	δ_{raw}	49.15
% USD	EUR	rhs [flip F4]	$-(\delta_{raw} - P)S/K$	−44.72
% USD	USD	lhs [flip F4]	$-(\delta_{raw})S/K$	−49.15

TABLE 1.5 1y EUR call USD put strike $K = 0.7000$ for a EUR-based bank. Market data: spot $S = 0.9090$, volatility $\sigma = 12\%$, EUR rate $r_f = 3.96\%$, USD rate $r_d = 3.57\%$. The raw delta is 94.82% EUR and the value is 21.88% EUR.

Delta ccy	Prem ccy	FENICS	Formula	Delta
% EUR	EUR	lhs	$\delta_{raw} - P$	72.94
% EUR	USD	rhs	δ_{raw}	94.82
% USD	EUR	rhs [flip F4]	$-(\delta_{raw} - P)S/K$	−94.72
% USD	USD	lhs [flip F4]	$-\delta_{raw}S/K$	−123.13

1.4.7 Strike in Terms of Delta

Since $v_x = \Delta = \phi e^{-r_f\tau}\mathcal{N}(\phi d_+)$ we can retrieve the strike as

$$K = x \exp\left\{-\phi\mathcal{N}^{-1}(\phi\Delta e^{r_f\tau})\sigma\sqrt{\tau} + \sigma\theta_+\tau\right\}. \tag{56}$$

Forward Delta I had labeled $\phi\mathcal{N}(\phi d_+)$ as forward delta in (17) and interpreted it as the number of units of FOR of forward contracts a trader needs to buy to delta hedge a short option. Here is justification: translating the above hedge ratio into calculus means that we need to compute (for a call)

$$\frac{\partial v_{\text{option}}}{\partial v_{\text{forward}}}. \tag{57}$$

The value of a forward contract v_{forward} for an agreed forward exchange rate K is obviously

$$\phi[xe^{-r_f\tau} - Ke^{-r_d\tau}], \tag{58}$$

(the Black-Scholes formula with the normal probabilities set equal to one, because a forward is always transacted at maturity), and its derivative with respect to spot x is

$$\frac{\partial v_{\text{forward}}}{\partial x} = \phi e^{-r_f\tau}. \tag{59}$$

Using the chain rule and the derivative of the inverse function, we obtain for the forward delta

$$\begin{aligned}\frac{\partial v_{\text{option}}}{\partial v_{\text{forward}}} &= \frac{\partial v_{\text{option}}}{\partial x}\frac{\partial x}{\partial v_{\text{forward}}}\\ &= \phi e^{-r_f\tau}\mathcal{N}(\phi d_+)\left(\frac{\partial v_{\text{forward}}}{\partial x}\right)^{-1}\\ &= \phi e^{-r_f\tau}\mathcal{N}(\phi d_+)\left(e^{-r_f\tau}\right)^{-1}\\ &= \phi\mathcal{N}(\phi d_+),\end{aligned} \tag{60}$$

and Bob's your uncle. High-school calculus – no university degree needed. Needless to say, there can also be a premium-adjusted forward delta which is (for the sake of completeness)

$$\phi \frac{K}{f} \mathcal{N}(\phi d_-). \tag{61}$$

By now you are hopefully aware that in FX options markets the notion of delta is highly important. It is imperative to develop a habit to ask which notion of delta is applied when traders or quants talk about it: with or without premium adjustment, spot or forward, nominal currency and the quotation of the respective currency pair. And all of this will expand even further once other models beyond Black-Scholes are applied and smile is included.

1.4.8 Volatility in Terms of Delta

The mapping $\sigma \mapsto \Delta = \phi e^{-r_f \tau} \mathcal{N}(\phi d_+)$ is not one-to-one. The two solutions are given by

$$\sigma_\pm = \frac{1}{\sqrt{\tau}} \left\{ \phi \mathcal{N}^{-1}(\phi \Delta e^{r_f \tau}) \pm \sqrt{(\mathcal{N}^{-1}(\phi \Delta e^{r_f \tau}))^2 - \sigma \sqrt{\tau}(d_+ + d_-)} \right\}. \tag{62}$$

Thus using just the delta to retrieve the volatility of an option is not advisable.

1.4.9 Volatility and Delta for a Given Strike

The determination of the volatility and the delta for a given strike is an iterative process involving the determination of the delta for the option using at-the-money volatilities in a first step and then using the determined volatility to re-determine the delta and to continuously iterate the delta and volatility until the volatility does not change more than $\epsilon = 0.001\%$ between iterations. More precisely, one can perform the following algorithm. Let the given strike be K.

1. Choose σ_0 = at-the-money volatility from the volatility matrix.
2. Calculate $\Delta_{n+1} = \Delta(\text{call}(K, \sigma_n))$.
3. Take $\sigma_{n+1} = \sigma(\Delta_{n+1})$ from the volatility matrix, possibly via a suitable interpolation.
4. If $|\sigma_{n+1} - \sigma_n| < \epsilon$, then quit, otherwise continue with step 2.

In order to prove the convergence of this algorithm we need to establish convergence of the recursion

$$\begin{aligned} \Delta_{n+1} &= e^{-r_f \tau} \mathcal{N}(d_+(\Delta_n)) \\ &= e^{-r_f \tau} \mathcal{N}\left(\frac{\ln(S/K) + (r_d - r_f + \frac{1}{2}\sigma^2(\Delta_n))\tau}{\sigma(\Delta_n)\sqrt{\tau}} \right) \end{aligned} \tag{63}$$

for sufficiently large $\sigma(\Delta_n)$ and a sufficiently smooth volatility smile surface. We must show that the sequence of these Δ_n converges to a fixed point $\Delta^* \in [0, 1]$ with a fixed volatility $\sigma^* = \sigma(\Delta^*)$.

This proof has been carried out in the thesis by Borowski [18] and works like this. We consider the derivative

$$\frac{\partial \Delta_{n+1}}{\partial \Delta_n} = -e^{-r_f \tau} n(d_+(\Delta_n)) \frac{d_-(\Delta_n)}{\sigma(\Delta_n)} \cdot \frac{\partial}{\partial \Delta_n} \sigma(\Delta_n). \tag{64}$$

The term

$$-e^{-r_f \tau} n(d_+(\Delta_n)) \frac{d_-(\Delta_n)}{\sigma(\Delta_n)}$$

converges rapidly to zero for very small and very large spots, being an argument of the standard normal density n. For sufficiently large $\sigma(\Delta_n)$ and a sufficiently smooth volatility surface in the sense that $\frac{\partial}{\partial \Delta_n} \sigma(\Delta_n)$ is sufficiently small, we obtain

$$\left| \frac{\partial}{\partial \Delta_n} \sigma(\Delta_n) \right| \stackrel{\Delta}{=} q < 1. \tag{65}$$

Thus for any two values $\Delta_{n+1}^{(1)}$, $\Delta_{n+1}^{(2)}$, and continuously differentiable smile surface we obtain

$$|\Delta_{n+1}^{(1)} - \Delta_{n+1}^{(2)}| < q |\Delta_n^{(1)} - \Delta_n^{(2)}|, \tag{66}$$

due to the mean value theorem. Hence the sequence Δ_n is a contraction in the sense of the fixed point theorem of Banach. This implies that the sequence converges to a unique fixed point in $[0, 1]$, which is given by $\sigma^* = \sigma(\Delta^*)$.

1.4.10 Greeks in Terms of Deltas

In foreign exchange markets the moneyness of vanilla options is always expressed in terms of deltas and prices are quoted in terms of volatility. This makes a ten-delta call a financial object as such independent of spot and strike. This method and the quotation in volatility makes objects and prices transparent in a very intelligent and user-friendly way. At this point we list the Greeks in terms of deltas instead of spot and strike. Let us introduce the quantities

$$\Delta_+ \stackrel{\Delta}{=} \phi e^{-r_f \tau} \mathcal{N}(\phi d_+) \text{ spot delta}, \tag{67}$$

$$\Delta_- \stackrel{\Delta}{=} -\phi e^{-r_d \tau} \mathcal{N}(\phi d_-) \text{ dual delta}, \tag{68}$$

which we assume to be given. From these we can retrieve

$$d_+ = \phi \mathcal{N}^{-1}(\phi e^{r_f \tau} \Delta_+), \tag{69}$$

$$d_- = \phi \mathcal{N}^{-1}(-\phi e^{r_d \tau} \Delta_-). \tag{70}$$

Interpretation of Dual Delta The dual delta introduced in (31) as the sensitivity with respect to strike has another – more practical – interpretation in a foreign exchange setup. We have seen in Section 1.4.5 that the domestic value

$$v(x, K, \tau, \sigma, r_d, r_f, \phi) \tag{71}$$

corresponds to a foreign value

$$v\left(\frac{1}{x}, \frac{1}{K}, \tau, \sigma, r_f, r_d, -\phi\right) \tag{72}$$

up to an adjustment of the nominal amount by the factor xK. From a foreign viewpoint the delta is thus given by

$$\begin{aligned} & -\phi e^{-r_d\tau}\mathcal{N}\left(-\phi\frac{\ln\left(\frac{K}{x}\right)+\left(r_f-r_d+\frac{1}{2}\sigma^2\tau\right)}{\sigma\sqrt{\tau}}\right) \\ &= -\phi e^{-r_d\tau}\mathcal{N}\left(\phi\frac{\ln\left(\frac{x}{K}\right)+\left(r_d-r_f-\frac{1}{2}\sigma^2\tau\right)}{\sigma\sqrt{\tau}}\right) \\ &= \Delta_-, \end{aligned} \tag{73}$$

which means the dual delta is the delta from the foreign viewpoint. This is again in line with its interpretation of a premium-adjusted delta as in Equation (55). We will see below that foreign rho, vega and gamma do not require knowing the dual delta. We will now state the Greeks in terms of $x, \Delta_+, \Delta_-, r_d, r_f, \tau, \phi$.

Value.

$$v(x, \Delta_+, \Delta_-, r_d, r_f, \tau, \phi) = x\Delta_+ + K\Delta_-\frac{e^{-r_f\tau}n(d_+)}{e^{-r_d\tau}n(d_-)} \tag{74}$$

(Spot) Delta.

$$\frac{\partial v}{\partial x} = \Delta_+ \tag{75}$$

Forward Delta.

$$\frac{\partial v}{\partial v_f} = e^{r_f\tau}\Delta_+ \tag{76}$$

Gamma.

$$\frac{\partial^2 v}{\partial x^2} = e^{-r_f\tau}\frac{n(d_+)}{x(d_+ - d_-)} \tag{77}$$

Taking a trader's gamma (change of delta if spot moves by 1%) additionally removes the spot dependence, because

$$\Gamma_{trader} = \frac{x}{100}\frac{\partial^2 v}{\partial x^2} = e^{-r_f\tau}\frac{n(d_+)}{100(d_+ - d_-)} \tag{78}$$

Speed.

$$\frac{\partial^3 v}{\partial x^3} = -e^{-r_f\tau}\frac{n(d_+)}{x^2(d_+ - d_-)^2}\left(2d_+ - d_-\right) \tag{79}$$

Theta.

$$\frac{1}{x}\frac{\partial v}{\partial t} = -e^{-r_f\tau}\frac{n(d_+)(d_+ - d_-)}{2\tau} + \left[r_f\Delta_+ + r_d\Delta_-\frac{e^{-r_f\tau}n(d_+)}{e^{-r_d\tau}n(d_-)}\right] \tag{80}$$

Charm.

$$\frac{\partial^2 v}{\partial x\partial\tau} = -\phi r_f e^{-r_f\tau}\mathcal{N}(\phi d_+) + \phi e^{-r_f\tau}n(d_+)\frac{2(r_d - r_f)\tau - d_-(d_+ - d_-)}{2\tau(d_+ - d_-)} \tag{81}$$

Color.

$$\frac{\partial^3 v}{\partial x^2\partial\tau} = -\frac{e^{-r_f\tau}n(d_+)}{2x\tau(d_+ - d_-)}\left[2r_f\tau + 1 + \frac{2(r_d - r_f)\tau - d_-(d_+ - d_-)}{2\tau(d_+ - d_-)}d_+\right] \tag{82}$$

Vega.

$$\frac{\partial v}{\partial\sigma} = xe^{-r_f\tau}\sqrt{\tau}n(d_+) \tag{83}$$

Volga.

$$\frac{\partial^2 v}{\partial\sigma^2} = xe^{-r_f\tau}\tau n(d_+)\frac{d_+d_-}{d_+ - d_-} \tag{84}$$

Vanna.

$$\frac{\partial^2 v}{\partial\sigma\partial x} = -e^{-r_f\tau}n(d_+)\frac{\sqrt{\tau}d_-}{d_+ - d_-} \tag{85}$$

Rho.

$$\frac{\partial v}{\partial r_d} = -x\tau\Delta_-\frac{e^{-r_f\tau}n(d_+)}{e^{-r_d\tau}n(d_-)} \tag{86}$$

$$\frac{\partial v}{\partial r_f} = -x\tau\Delta_+ \tag{87}$$

Dual Delta.

$$\frac{\partial v}{\partial K} = \Delta_- \qquad (88)$$

Dual Gamma.

$$K^2 \frac{\partial^2 v}{\partial K^2} = x^2 \frac{\partial^2 v}{\partial x^2} \qquad (89)$$

Dual Theta.

$$\frac{\partial v}{\partial T} = -v_t \qquad (90)$$

As an important example we consider vega.

Vega in Terms of Delta The mapping $\Delta \mapsto v_\sigma = xe^{-r_f\tau}\sqrt{\tau}n(\mathcal{N}^{-1}(e^{r_f\tau}\Delta))$ is important for trading vanilla options. Observe that this function does not depend on r_d or σ, just on r_f. Quoting vega in % foreign will additionally remove the spot dependence. This means that for a moderately stable foreign term structure curve, traders will be able to use a moderately stable vega matrix. For $r_f = 3\%$ the vega matrix is presented in Table 1.6. There are traders who know this table by heart. Just as motivation in case you don't have any plans for tonight. Usually, a vega hedge uses an option with the same maturity, so the table is in fact also independent of r_f. The vega hedging default product is an at-the-money straddle with the same maturity. However, note that vega is mostly hedged on a portfolio level, rather than on individual transactions. Individual vega hedging would be performed on big size tickets.

TABLE 1.6 Vega in terms of delta for the standard maturity labels and various deltas. It shows that one can vega hedge a long 9M 35 delta call with 4 short 1M 20 delta puts.

Mat/Δ	50%	45%	40%	35%	30%	25%	20%	15%	10%	5%
1D	2	2	2	2	2	2	1	1	1	1
1W	6	5	5	5	5	4	4	3	2	1
1W	8	8	8	7	7	6	5	5	3	2
1M	11	11	11	11	10	9	8	7	5	3
2M	16	16	16	15	14	13	11	9	7	4
3M	20	20	19	18	17	16	14	12	9	5
6M	28	28	27	26	24	22	20	16	12	7
9M	34	34	33	32	30	27	24	20	15	9
1Y	39	39	38	36	34	31	28	23	17	10
2Y	53	53	52	50	48	44	39	32	24	14
3Y	63	63	62	60	57	53	47	39	30	18

1.4.11 Settlement

Standard textbooks dealing with the valuation of options deal with only two times, one for the beginning and one for the end of the deal. In practice, the story is a bit more advanced and deals with the dates listed in Figure 1.3.

We are going to use the following notation:

- T_h *horizon date*, represents the date on which the derivative is evaluated. In many cases it represents today.
- T_{hs} *horizon spot date*, two business days after the horizon date.
- T_e *expiry date*. For path independent options the payoff depends on the quoted spot or forward on this date.
- T_{es} *expiry spot date*, two business days after the expiry date.
- T_d *delivery date*, represents the date on which cash flows implied by the option payoff will be settled.

For corridors, faders, target forwards and other derivatives transactions whose payoff depends on a fixing schedule, additional fixing dates are introduced:

- T_f *fixing date*, on this date it is decided whether the underlying is inside a specified range.
- T_{fs} *fixing spot date*, two business days after the fixing date.

In general we have

$$T_h \overset{2bd}{\leq} T_{hs} \leq T_f \overset{2bd}{\leq} T_{fs} \leq T_e \overset{2bd}{\leq} T_{es} \leq T_d. \tag{91}$$

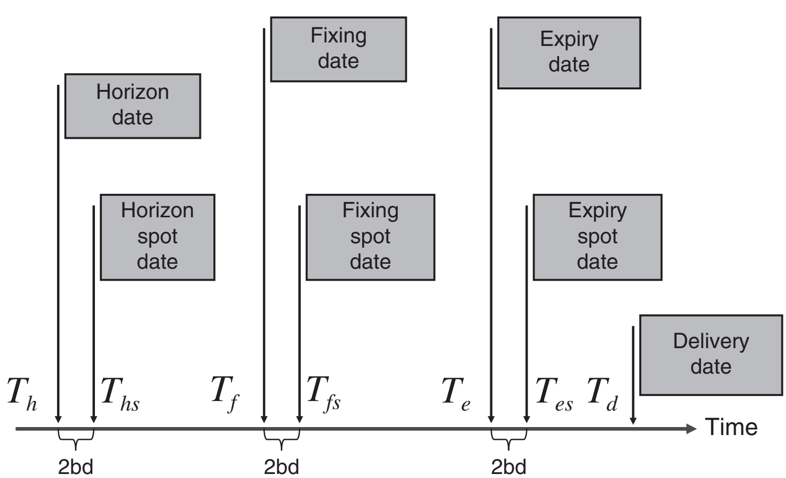

FIGURE 1.3 Relevant dates for trading options. The spot dates are usually two business days (2bd) after the horizon, fixing or expiry date.

The payoff of FX options depends either on the quoted (2bd) spot S_{T_e} at the expiry date T_e, or the forward F_{T_{es},T_d} from that date to a specific delivery date T_d. In many cases the delivery date T_d corresponds to the expiry spot date T_{es}, where we have $F_{T_{es},T_d} = S_{T_e}$.

The constant volatility used in the model (see Equation (1)) corresponds to the time period $T_h = 0$ to $T_e = T$. This is the time interval that is relevant for the exchange rate risk. The risk starts right after the transaction and ends as soon as the option is exercised or expires.

All interest rates r_d and r_f are assumed to correspond to the respective period in their context, for example: the term $e^{-r_d(T_d - T_{hs})}$ represents the factor used for discounting the payoff, using the domestic (forward) interest rate from the spot date T_{hs} to the delivery date T_d.

Remark 1.4.1 *In rare cases the delivery date can be* before *the expiry spot date, that is $T_{es} > T_d$. However, it can never be before the expiry date T_e. For instance, in Canada, the delivery date is usually only one business day after the expiry date. One can have a delivery on the same day upon special request.*

The Black-Scholes Model for the Actual Spot The standard approach is to assume a Black-Scholes model as in Equation (1)

$$\frac{d\hat{S}_t}{\hat{S}_t} = (r_d - r_f)\, dt + \sigma\, dW_t \tag{92}$$

for the underlying, where $\hat{S}_t$ means the price of the underlying at time t.

Instead of modeling this zero-day spot $\hat{S}$ we need to model the *2bd spot* S, which is usually quoted in FX markets.

Assuming no arbitrage opportunities this leads to the relationships

$$\begin{aligned} \hat{S}_t &= S_t e^{-(r_d - r_f)(T_{hs} - t)}, \\ \hat{S}_{T_e} &= S_{T_e} e^{-(r_d - r_f)(T_{es} - T_e)}. \end{aligned} \tag{93}$$

Assuming the exponents in (93) are deterministic, the quoted spot at some future time T_e satisfies

$$S_{T_e} = S_t \exp\left[(r_d - r_f)(T_{es} - T_{hs}) - \tfrac{1}{2}\sigma^2 (T_e - t) + \sigma W_{T_e - t}\right], \tag{94}$$

which follows directly from (93) and (2).

Cash Settlement In case of *cash settlement* the seller of the option pays a cash amount depending on the payoff formula and the quoted spot S_{T_e} to the holder. By default, the cash arrives in the holder's account on the expiry spot date T_{es}, but in general on the delivery date T_d. For example, in the case of a vanilla quoted in FOR-DOM, the DOM cash amount

$$\left(\phi\left(S_{T_e}-K\right)\right)^+ \tag{95}$$

is paid, where as usual $\phi = \pm 1$ for calls and puts respectively. The option will usually be exercised if $\phi S_{T_e} > \phi K$.

The value of the vanilla option at time $t = T_{hs}$ is then given by

$$\begin{aligned} v_t^C &= I\!E\left[e^{-r_d(T_d-T_{hs})}\left(\phi\left(S_{T_e}-K\right)\right)^+ \middle| S_t\right] \\ &= I\!E\left[e^{-r_d(T_d-T_{es})}e^{-r_d(T_{es}-T_{hs})}\left(\phi\left(S_{T_e}-K\right)\right)^+ S_t\right] \\ &= e^{-r_d(T_d-T_{es})}v_t^{PV}, \end{aligned} \tag{96}$$

where v_t^{PV} denotes the value of a vanilla in the Black-Scholes model published in textbooks – compare with Equation (7). In particular, we have

$$d_\pm^C = \frac{\ln\frac{S_t}{K} + (r_d - r_f)(T_{es} - T_{hs}) \pm \frac{1}{2}\sigma^2(T_e - t)}{\sigma\sqrt{T_e - t}}, \tag{97}$$

and

$$\begin{aligned} v_t^C &= \phi e^{-r_d(T_d-T_{hs})}\left(e^{(r_d-r_f)(T_{es}-T_{hs})}S_t\mathcal{N}(\phi d_+^C) - K\mathcal{N}(\phi d_-^C)\right) \\ &= \phi e^{-r_d(T_d-T_{es})}\left(e^{-r_f(T_{es}-T_{hs})}S_t\mathcal{N}(\phi d_+^C) - e^{-r_d(T_{es}-T_{hs})}K\mathcal{N}(\phi d_-^C)\right). \end{aligned} \tag{98}$$

For other options this works similarly.

Delivery Settlement In the case of *delivery settlement* the cash amounts of the two involved currencies FOR and DOM are physically exchanged on the delivery date T_d. Therefore the intrinsic value is given by

$$\left(\phi\left(F_{T_{es},T_d}-K\right)\right)^+ \tag{99}$$

$$= \left(\phi\left(S_{T_e}e^{(r_d-r_f)\cdot(T_d-T_{es})}-K\right)\right)^+, \tag{100}$$

where the second equality is meant to hold at expiry date T_e, when the rates are known. The option will typically be exercised if

$$\phi F_{T_{es},T_d} > \phi K, \tag{101}$$

which is equivalent to

$$\phi S_{T_e} > \phi K e^{-(r_d-r_f)(T_d-T_{es})}. \tag{102}$$

The value of a vanilla option at time $t = T_{hs}$ is then given by

$$v_t^D = \mathbb{E}\left[e^{-r_d(T_d-T_{hs})}\left(\phi\left(F_{T_{es},T_d} - K\right)\right)^+ \middle| S_t\right] \tag{103}$$

$$= \phi e^{-r_d(T_d-T_{hs})}\left(e^{(r_d-r_f)(T_d-T_{hs})}S_t\mathcal{N}(\phi d_+^D) - K\mathcal{N}(\phi d_-^D)\right)$$

$$= \phi e^{-r_f(T_d-T_{hs})}S_t\mathcal{N}(\phi d_+^D) - \phi e^{-r_d(T_d-T_{hs})}K\mathcal{N}(\phi d_-^D),$$

$$d_\pm^D = \frac{\ln\frac{S_t}{K} + (r_d - r_f)(T_d - T_{hs}) \pm \frac{1}{2}\sigma^2(T_e - t)}{\sigma\sqrt{T_e - t}}.$$

Options with Deferred Delivery Options in FX OTC markets are typically delivery-settled, as corporates do have cash to exchange. Options are bought for the purpose of investment, or speculation as well, and in that case a cash settlement is more suitable. In this section we deal with delivery-settled options, where the delivery date is deferred.

In the case of an option with delivery settlement often it becomes important to the corporate treasurer to have a settlement date significantly after the expiration date of the contract. The default is two business days, but corporates sometimes wish to delay the settlement up to one year or even further. The reason is that a decision to exchange amounts in different currencies in the future is often taken much earlier than the actual payment time, just like in the case of *compound* or *installment* options.

This means that the option under consideration is not an option on the FX *spot* but rather an option on the FX *forward*, which is sometimes called *compound on forward*.

To be concrete, let us derive the formula for the deferred delivery-settled vanilla call. We let the current time be T_h, the expiration time be T_e, and the delivery time be T_d. The buyer of a deferred-delivery call with strike K has the right to enter into a forward contract with strike K at time T_e, which is then delivered at time T_d. In a Black-Scholes model framework with constant interest rates, the forward price at time t maturing at time T_d is given by the random variable

$$f_t(T_d) = S_t e^{(r_d-r_f)(T_d-t)}, \tag{104}$$

in particular, at time zero,

$$f_0(T_d) = S_0 e^{(r_d-r_f)T_d}, \tag{105}$$

is the current forward price. We let T_d be fixed and view t as the variable. Using Itô's formula,[1] we see that the forward price satisfies

$$df_t(T_d) = \sigma f_t(T_d)dW_t, \tag{106}$$

hence it is a martingale. In a risk-neutral valuation approach, the value of a call on $f_t(T_d)$ is given by

$$v(0) = e^{-r_dT_d}\mathbb{E}[(f_t(T_d) - K)^+]. \tag{107}$$

[1] $df(S_t) = f'(S_t)\, dS_t + \frac{1}{2}f''(S_t)\,(dS_t)^2$

In order to compute this, we notice that we can use the existing Black-Scholes Equation (7) for the special case $r_d = r_f$ (due to Equation (106)) and $S_0 = f_0(T_d)$, which is

$$\begin{aligned} v(0) &= e^{-r_d T_d}\left[f_0(T_d)\mathcal{N}(d_+) - K\mathcal{N}(d_-)\right] \\ &= S_0 e^{-r_f T_d}\mathcal{N}(d_+) - Ke^{-r_d T_d}\mathcal{N}(d_-), \end{aligned} \tag{108}$$

$$\begin{aligned} d_\pm &= \frac{\ln\frac{f_0(T_d)}{K} \pm \frac{1}{2}\sigma^2 T_e}{\sigma\sqrt{T_e}} \\ &= \frac{\ln\frac{S_0}{K} + (r_d - r_f)T_d \pm \frac{1}{2}\sigma^2 T_e}{\sigma\sqrt{T_e}}. \end{aligned} \tag{109}$$

This calculation works similarly for all European style path-independent options. The basic procedure is to reuse existing formulas for options as a function of the spot for a zero drift $r_d = r_f$ and replace the spot variable by the forward.

1.4.12 Exercises

Vega Maximizing Strike Show that the delta-symmetric strike K_+ maximizes vega.

Gamma Maximizing Strike Show that the delta-symmetric strike K_+ maximizes gamma.

Forward Delta Normally, when selling an option, the trader buys delta times the notional of the option contract of the underlying in the spot market. However, in many situations, such as trading long-term contracts in-the-money, it turns out to be more suitable to replace this spot hedge by a forward hedge. This means that instead of buying the underlying, the trader buys a forward contract on the underlying with the same maturity as the option.

(i) The *forward delta* is defined to be the number of units of forward contracts to buy so the total delta of the short option and the long forward contract is zero. Show that the forward delta is equal to

$$\phi\mathcal{N}(\phi d_+). \tag{110}$$

(ii) Let us fix the initial notional of the forward traded to hedge and call it

$$\phi\tilde{\mathcal{N}}(\phi d_+). \tag{111}$$

Keeping this quantity fixed, write down the value of the portfolio P of the short vanilla option and the long forward contract. Verify that the total initial spot delta is zero.

(iii) Argue why replacing the spot hedge with a forward hedge does not change the vega and gamma positions.

(iv) Show that the portfolio P has a foreign rho $\frac{\partial P}{\partial r_f}$ of zero.

(v) Compute the domestic rho of the portfolio $\frac{\partial P}{\partial r_d}$ and argue why this is often considerably smaller than the domestic rho of a vanilla contract.

Overall these considerations show that forward hedging also takes care of most of the interest rate risk, which is why, in principle, it is the better way to hedge. However, spot markets are normally more liquid than forward markets, and bid-ask spreads are smaller, which prevents traders from doing forward hedges most of the time. A common practice is to always do a spot hedge at inception of the trade and replace it by a forward hedge if the underlying contract has a high sensitivity to the carry (difference of interest rates). This can be done by an FX swap.

Vega-Delta Implied volatility is sometimes quoted in terms of the spot delta in the Black-Scholes model, rather than in terms of the strike, particularly in foreign exchange markets. The reason is that vega

$$\frac{\partial v}{\partial \sigma} = x e^{-r_f \tau} \sqrt{\tau} n(d_+) \tag{112}$$

does not depend on the domestic rate r_d and not on the volatility σ, when quoted in terms of delta. Prove these two statements.

Premium-Adjusted Delta If the underlying security is an exchange rate, then the option premium can be paid in both the domestic and the foreign currency. Therefore, we must re-think our notion of delta, which represents the number of units of foreign currency to buy when selling a vanilla option. This works fine if the premium is paid in domestic currency. However, if the premium is paid in foreign currency, applying a standard delta hedge would be wrong, because the trader would end up having too many units of foreign currency in his book. We define the *premium-adjusted spot delta* as

$$x \cdot \frac{\partial}{\partial x}\left(\frac{v}{x}\right). \tag{113}$$

(i) Show that the premium-adjusted spot delta can be written as

$$\Delta_S - \frac{v}{x}, \tag{114}$$

where Δ_S is the standard vanilla delta, which is not premium-adjusted. Argue intuitively, why this is the correct quantity.

(ii) Similarly, we define the *premium-adjusted forward delta* as

$$\Delta_{Fpa} \triangleq \frac{x \cdot \frac{\partial}{\partial x}\left(\frac{v}{x}\right)}{\frac{\partial v_f}{\partial x}}, \tag{115}$$

which is the number of forward contracts to buy when the premium is paid in foreign currency. Show that the premium-adjusted forward delta is given by

$$\phi \frac{K}{f} \mathcal{N}(\phi d_-). \tag{116}$$

(iii) Show that the strike K_- that generates a premium-adjusted forward delta position of zero of a straddle (call and put with strike K_-) is given by

$$K_- = f e^{-\frac{1}{2}\sigma^2 \tau}. \tag{117}$$

Such a straddle is used to hedge vega without affecting the delta position.

(iv) Show that in the case of spot deltas without premium adjustment the relationship of strike and delta is monotone and show how to retrieve the strike from a given spot delta.

(v) Show that in the case of premium-adjusted forward deltas the relationship of strike to delta is no longer guaranteed to be monotone. In particular, show that there is another strike that generates the same delta as $\check{K}$ if

$$\sigma > \sqrt{\frac{2}{\pi \tau}}. \tag{118}$$

This ambiguity would pop up for 30-year options for volatilities above 14.6%, which is not unrealistic, even in FX markets.

Deferred Delivery Driven by Forward Discuss the implications of a deferred delivery on the value of a vanilla call with a fixed strike. What happens to the value if the forward curve is downward sloping? What are the main risk factors involved if the deferral period is very long? How about correlation between the spot and the interest rates?

Greeks in a Binomial Tree Model Implement an N-step binomial tree model that allows the value and delta of European style call and put options to be computed.

1. Compute the values and deltas of put and call options for the following sample contract data: strike $K = 1.4500$, maturity $\tau = 1$ year. Use $N = 120$ time steps, so a single time step will be $\Delta_t = \frac{\tau}{N}$. Let the market data be given by $S_0 = 1.5000$, $r = 5\%$ p.a., $r_f = 4\%$ p.a., volatility $\sigma = 10\%$. For the implementation use the formulas

$$u = e^{+\sigma\sqrt{\Delta_t}}, \tag{119}$$

$$d = e^{-\sigma\sqrt{\Delta_t}}, \tag{120}$$

$$\tilde{p} = \frac{e^{(r-r_f)\Delta_t} - d}{u - d}. \tag{121}$$

For the discounting per time step you need to use $e^{-r\Delta_t}$.

2. Examine the convergence behavior of the call option value and delta for a strike of $K = 1.4500$ and maturity of $T = 1$ year empirically. This means computing the sequences of results for values of N ranging from 1 to 200. Again, a single time step will be $\Delta_t = \frac{T}{N}$. Let the market data be given by $S_0 = 1.5000$, $r = 5\%$ p.a., $r_f = 0$, volatility $\sigma = 10\%$. Plot your results in a graph taking the values for N on the x-axis and the value and delta on the y-axis. Describe the convergence behavior you see, explain, why it possibly looks the way it looks, and suggest and try methods to improve the convergence.

1.5 VOLATILITY

Volatility is the *annualized standard deviation of the log-returns*. It is *the* crucial input parameter in determining the value of an option. Hence, the crucial question is where to derive the volatility from. If no active option market is present, the only source of information is estimating the historic volatility. This would give some clue about the *past*. In liquid currency pairs volatility is often a traded quantity on its own, quoted by traders, brokers, and real-time data pages. These quotes reflect market participants' views about the *future*.

Since volatility normally does not stay constant, option traders are highly concerned with hedging their volatility exposure. Hedging vanilla options' vega is comparatively easy because vanilla options have convex payoffs, whence vega is always positive, i.e. the higher the volatility, the higher the price. Let us take for example a EUR-USD market with spot 1.2000, USD- and EUR rate at 2.5%. A three-month at-the-money call with 1 M EUR notional would cost 29,000 USD at a volatility of 12%. If the volatility now drops to a value of 8%, then the value of the call would be only 19,000 USD. This monotone dependence is not guaranteed for non-convex payoffs, as we illustrate in Figure 1.4.

1.5.1 Historic Volatility

We briefly describe how to compute the historic volatility of a time series

$$S_0, S_1, \ldots, S_N \tag{122}$$

of daily data. First, we create the sequence of log-returns

$$r_i = \ln \frac{S_i}{S_{i-1}}, \quad i = 1, \ldots, N. \tag{123}$$

Then, we compute the average log-return

$$\bar{r} = \frac{1}{N} \sum_{i=1}^{N} r_i, \tag{124}$$

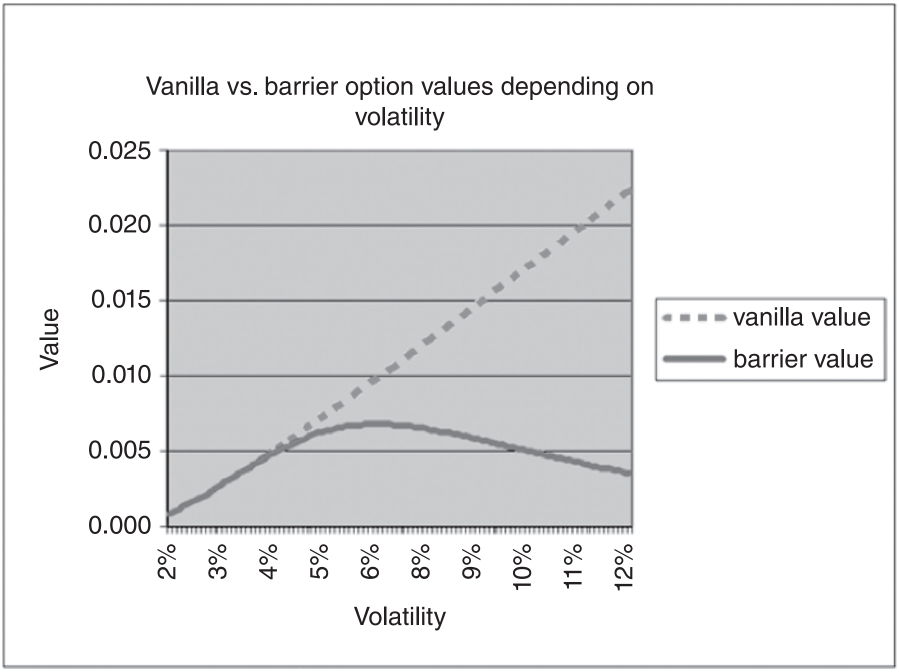

FIGURE 1.4 Dependence of the value of a vanilla call and a reverse knock-out call on volatility. The vanilla value is monotone in the volatility, whereas the barrier value is not. The reason is that as the spot gets closer to the upper knock-out barrier, an increasing volatility would increase the chance of knock-out and hence decrease the value.

their variance

$$\hat{\sigma}^2 = \frac{1}{N-1} \sum_{i=1}^{N} (r_i - \bar{r})^2, \tag{125}$$

and their standard deviation

$$\hat{\sigma} = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (r_i - \bar{r})^2}. \tag{126}$$

The annualized standard deviation, which is the volatility, is then given by

$$\hat{\sigma}_a = \sqrt{\frac{B}{N-1} \sum_{i=1}^{N} (r_i - \bar{r})^2}, \tag{127}$$

where the *annualization factor B* is given by

$$B = \frac{N}{k} d, \tag{128}$$

and k denotes the number of calendar days within the time series and d denotes the number of calendar days per year.

Assuming normally distributed log-returns, we know that $\hat{\sigma}^2$ is χ^2-distributed. Therefore, given a confidence level of p and a corresponding error probability $\alpha = 1 - p$, the p-confidence interval is given by

$$\left[\hat{\sigma}_a\sqrt{\frac{N-1}{\chi^2_{N-1;1-\frac{\alpha}{2}}}}, \hat{\sigma}_a\sqrt{\frac{N-1}{\chi^2_{N-1;\frac{\alpha}{2}}}}\right], \tag{129}$$

where $\chi^2_{n;p}$ denotes the p-quantile of a χ^2-distribution[2] with n degrees of freedom.

As an example let us take the 256 ECB fixings of EUR-USD from 4 March 2003 to 3 March 2004 displayed in Figure 1.5. We get $N = 255$ log-returns. Taking $k = d = 365$, we obtain

$$\bar{r} = \frac{1}{N}\sum_{i=1}^{N} r_i = 0.0004166,$$

$$\hat{\sigma}_a = \sqrt{\frac{B}{N-1}\sum_{i=1}^{N}(r_i - \bar{r})^2} = 10.85\%,$$

and a 95% confidence interval of [9.99%, 11.89%].

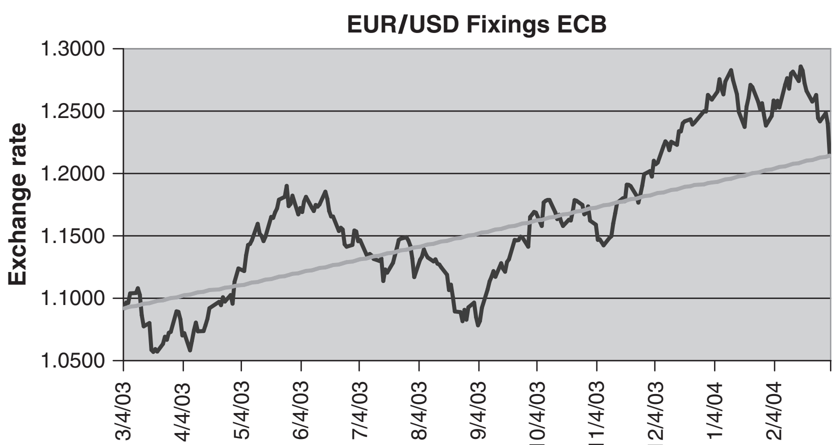

FIGURE 1.5 ECB fixings of EUR-USD from 4 March 2003 to 3 March 2004 and the line of average growth.

[2]Values and quantiles of the χ^2-distribution and other distributions can be computed on the internet, e.g. at http://eswf.uni-koeln.de/allg/surfstat/tables.htm or via EXCEL function CHIINV.

Notice that the quality of the return estimate is horrendously bad. All that matters is the slope of the line connecting the start point and the end point. Therefore, the return estimate is completely determined by the choice of these two points and can hence be considered arbitrary. This is good news for asset managers as they will always be able to find a time interval in a historic data series with positive return. Furthermore, the fact that returns are impossible to estimate justifies more than 95% of all jobs in asset management. It is also why Markowitz doesn't work but Black-Scholes does.

Quality of the volatility estimate is substantially better than the return estimate. In risk-neutral pricing (assuming it works), returns can be inferred from money market rates and do not need to be estimated from historic time series. The remaining uncertainty is volatility, which in turn has become a liquid tradable quantity in many currency pairs. If there are no options in a currency pair, then volatility can only be estimated from historic data.

1.5.2 Historic Correlation

As in the preceding section we briefly describe how to compute the historic correlation of two time series,

$$x_0, x_1, \dots, x_N,$$
$$y_0, y_1, \dots, y_N,$$

of daily data. First, we create the sequences of log-returns

$$X_i = \ln \frac{x_i}{x_{i-1}}, \quad i = 1, \dots, N,$$
$$Y_i = \ln \frac{y_i}{y_{i-1}}, \quad i = 1, \dots, N. \tag{130}$$

Then, we compute the average log-returns

$$\overline{X} = \frac{1}{N} \sum_{i=1}^{N} X_i,$$
$$\overline{Y} = \frac{1}{N} \sum_{i=1}^{N} Y_i, \tag{131}$$

their variances and covariance

$$\hat{\sigma}_X^2 = \frac{1}{N-1} \sum_{i=1}^{N} (X_i - \overline{X})^2, \tag{132}$$

$$\hat{\sigma}_Y^2 = \frac{1}{N-1} \sum_{i=1}^{N} (Y_i - \overline{Y})^2, \tag{133}$$

$$\hat{\sigma}_{XY} = \frac{1}{N-1}\sum_{i=1}^{N}(X_i - \overline{X})(Y_i - \overline{Y}), \tag{134}$$

and their standard deviations

$$\hat{\sigma}_X = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(X_i - \overline{X})^2}, \tag{135}$$

$$\hat{\sigma}_Y = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(Y_i - \overline{Y})^2}. \tag{136}$$

The estimate for the correlation of the log-returns is given by

$$\hat{\rho} = \frac{\hat{\sigma}_{XY}}{\hat{\sigma}_X \hat{\sigma}_Y}. \tag{137}$$

This correlation estimate is often not very stable, yet often it is the only available information. More recent work by Jäkel [51] treats robust estimation of correlation. We will revisit FX correlation risk in Section 1.9.2.

1.5.3 Volatility Smile

The Black-Scholes model assumes constant volatility throughout. However, market prices of traded options imply different volatilities for different maturities and different deltas. We start with some technical issues on how to imply the volatility from vanilla options.

Retrieving the Volatility from Vanilla Options Given the value of an option. Recall the Black-Scholes formula in Equation (7). We now look at the function $v(\sigma)$, whose derivative (vega) is

$$v'(\sigma) = xe^{-r_f\tau}\sqrt{\tau}n(d_+). \tag{138}$$

The function $\sigma \mapsto v(\sigma)$ is

1. strictly increasing,
2. concave up for $\sigma \in [0, \sqrt{2|\ln f - \ln K|/\tau})$,
3. concave down for $\sigma \in (\sqrt{2|\ln f - \ln K|/\tau}, \infty)$

and also satisfies

$$v(0) = [\phi(xe^{-r_f\tau} - fe^{-r_d\tau})]^+, \tag{139}$$

$$v(\infty, \phi = 1) = xe^{-r_f\tau}, \tag{140}$$

$$v(\sigma = \infty, \phi = -1) = Ke^{-r_d\tau}, \tag{141}$$

$$v'(0) = xe^{-r_f\tau}\sqrt{\tau}/\sqrt{2\pi}I\!I_{\{f=K\}}. \tag{142}$$

In particular the mapping $\sigma \mapsto v(\sigma)$ is invertible. However, the starting guess for employing Newton's method should be chosen with care, because the mapping $\sigma \mapsto v(\sigma)$ has a saddle point at

$$\left(\sqrt{\frac{2}{\tau}\left|\ln\frac{f}{K}\right|}, \phi e^{-r_d\tau}\left\{f\mathcal{N}\left(\phi\sqrt{2\tau\left[\ln\frac{f}{K}\right]^+}\right) - K\mathcal{N}\left(\phi\sqrt{2\tau\left[\ln\frac{K}{f}\right]^+}\right)\right\}\right), \tag{143}$$

as illustrated in Figure 1.6.

To ensure convergence of Newton's method, we are advised to use initial guesses for σ on the same side of the saddle point as the desired implied volatility. The danger is that a large initial guess could lead to a negative successive guess for σ. Therefore one should start with small initial guesses at or below the saddle point. For at-the-money options, the saddle point is degenerate for a zero volatility and small volatilities serve as good initial guesses.

The problem is well known and has been studied in detail by Jäckel [80]. He uses a clever rearrangement of the variables in combination with Newton's and Halley's method along with a good initial guess. Li and Lee show that the calculation speed can be substantially improved by *Successive Over-Relaxation (SOR)* [90, 91].

Market Data Now that we know how to imply the volatility from a given value, we can take a look at the market. We take EUR/GBP at the beginning of April 2005. The at-the-money volatilities for various maturities are listed in Table 1.7. We observe that implied volatilities are not constant but depend on the time to maturity of the option

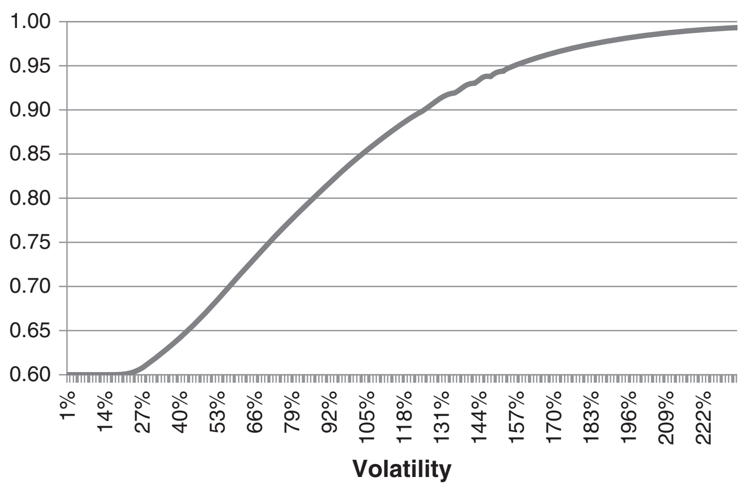

FIGURE 1.6 Value of a European call in terms of volatility with parameters $x = 1$, $K = 1$, $\tau = 5$, $r_d = 20\%$, $r_f = 0\%$. The saddle point is at $\sigma = 60\%$. The value starts at the value of the forward contract 0.5981 USD per EUR and converges to 1.0000 EUR, which is the (foreign discounted) value of the call currency amount expressed in USD.

TABLE 1.7 EUR/GBP implied volatilities in % for at-the-money vanilla options. Source: BBA (British Bankers' Association), http://www.bba.org.uk.

Date	Spot	1 week	1 month	3 month	6 month	1 year	2 year
1-Apr-05	0.6864	4.69	4.83	5.42	5.79	6.02	6.09
4-Apr-05	0.6851	4.51	4.88	5.34	5.72	5.99	6.07
5-Apr-05	0.6840	4.66	4.95	5.34	5.70	5.97	6.03
6-Apr-05	0.6847	4.65	4.91	5.39	5.79	6.05	6.12
7-Apr-05	0.6875	4.78	4.97	5.39	5.79	6.01	6.10
8-Apr-05	0.6858	4.76	5.00	5.41	5.78	6.00	6.09

as well as on the current time. This shows that the Black-Scholes assumption of a constant volatility is not fully justified looking at market data. We have a *term structure of volatility* as well as a stochastic nature of the term structure curve as time passes.

Besides the dependence on the time to maturity (term structure) we also observe different implied volatilities for different degrees of moneyness. This effect is called the *volatility smile*. The term structure and smile together are called a *volatility matrix* or *volatility surface*, if it is graphically displayed. Various possible reasons for this empirical phenomenon are discussed by Bates, among others, for example in [10].

In foreign exchange options markets implied volatilities are generally quoted and plotted against the deltas of out-of-the-money call and put options. This allows market participants to ask various partners for quotes on a 25-delta call, which is spot independent. The actual strike will be set depending on the spot once the trade is close to being finalized. The at-the-money option has a specific strike. In foreign exchange options markets, there are essentially three variants. For retail products, strike is commonly understood to be at the current spot. In the interbank market, strike is equal to the forward rate for options with tenors above one year or for emerging markets currency pairs. Setting the strike equal to the outright forward rate is equivalent to the *value* of the call and the put being equal. For tenors up to one year in standard currency pairs, strike is chosen to make the *delta* of the call and put equal. Other types of *at-the-money* are discussed in Section 1.5.6. The delta of a vanilla option in the case of strike equal to the outright forward rate is

$$\frac{\partial v}{\partial x} = \phi e^{-r_f \tau} \mathcal{N}\left(\phi \tfrac{1}{2}\sigma\sqrt{\tau}\right), \tag{144}$$

for a small volatility σ and short time to maturity τ, a number near $\phi 50\%$. This is why an at-the-money option is often called a 50-delta option. This is acceptable as a rough approximation; however, it is incorrect. In particular, it is incorrect for long-term vanilla options. Further market information is the implied volatilities for puts and calls with a delta of $\phi 25\%$. Other or additional implied volatilities for $\phi 10\%$ deltas are also quoted. Volatility matrices for more delta pillars are usually interpolated. Which at-the-money version is applied is convention-based and may change over time. A first source of information can be your risk management system, provided it is good and current.

Symmetric Decomposition Generally in Foreign Exchange, volatilities are decomposed into a *symmetric* part of the smile reflecting the *convexity* and a *skew-symmetric* part of the smile reflecting the *skew*. The way this works is that the market quotes *risk reversals (RR)* and *butterflies (BF)* or strangles – see Sections 1.6.2 and 1.6.4 for the description of the products and Figure 1.7 for the payoffs. Note that the terms risk reversal and butterfly/strangle are used both as names for trading strategies and numbers indicating the symmetric decomposition of the volatility smile. Not easy for newcomers. Here we are talking about the respective *volatilities* and *volatility differences* and about how to use these to construct the volatility smile and eventually price all vanilla options across the range of deltas. Sample quotes are listed in Tables 1.8 and 1.9. The relationship between risk reversal and strangle/butterfly quotes and the volatility smile are graphically explained in Figure 1.8.

FIGURE 1.7 The risk reversal (upper payoff) is a skew-symmetric product, the butterfly (lower payoff) is a symmetric product.

TABLE 1.8 EUR/GBP 25 delta risk reversal in %. Source: BBA (British Bankers' Association). This means that for example on 4 April 2005, the 1-month 25-delta EUR call was priced with a volatility of 0.15% higher than the 25-delta EUR put. At that moment the market apparently favored calls indicating a market sentiment in an upward movement.

Date	Spot	1 month	3 month	1 year
1-Apr-05	0.6864	0.18	0.23	0.30
4-Apr-05	0.6851	0.15	0.20	0.29
5-Apr-05	0.6840	0.11	0.19	0.28
6-Apr-05	0.6847	0.08	0.19	0.28
7-Apr-05	0.6875	0.13	0.19	0.28
8-Apr-05	0.6858	0.13	0.19	0.28

TABLE 1.9 EUR/GBP 25 delta butterfly in %. Source: BBA. This means that for example on 4 April 2005, the 1-month 25-delta EUR call and the 1-month 25-delta EUR put are on average quoted with a volatility of 0.15% higher than the 1-month at-the-money calls and puts. Using the quotes in Table 1.7 and Table 1.8, the result is that the 1-month 25-delta EUR call is quoted with a volatility of 4.88% + 0.15% + 0.075% and the 1-month 25-delta EUR put is quoted with a volatility of 4.88% + 0.15% – 0.075%. Note that some market participants (including BBA) also use the term "strangle" as a synonym for the butterfly quote. The vocabulary is not consistent, but the meaning has to be inferred from the context.

Date	Spot	1 month	3 month	1 year
1-Apr-05	0.6864	0.15	0.16	0.16
4-Apr-05	0.6851	0.15	0.16	0.16
5-Apr-05	0.6840	0.15	0.16	0.16
6-Apr-05	0.6847	0.15	0.16	0.16
7-Apr-05	0.6875	0.15	0.16	0.16
8-Apr-05	0.6858	0.15	0.16	0.16

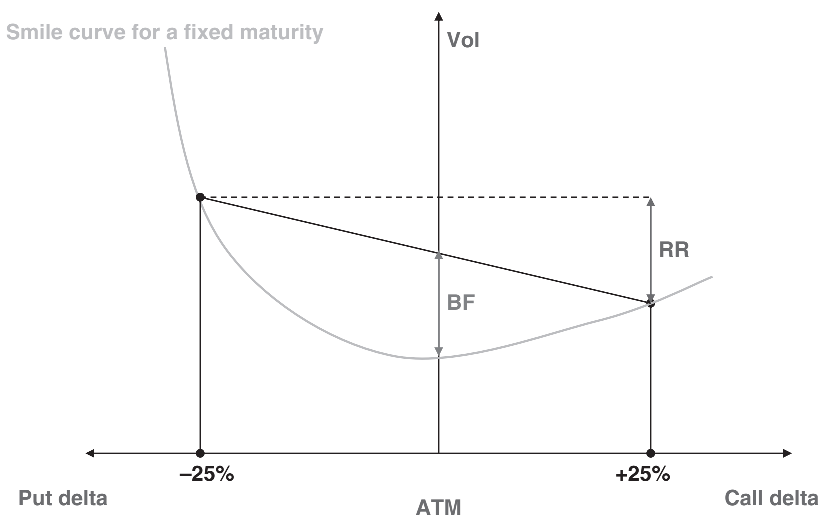

FIGURE 1.8 Risk reversal and butterfly in terms of volatility for a given FX vanilla option smile.

Generally, there are several ways the FX volatility smile can be constructed from quotes for at-the-money, risk reversals and butterflies/strangles. In the next two paragraphs I outline the basic two versions, one being the smile version, which works purely on the level of implied volatilities, the other one being the brokers' version, which requires a detour via values of strangles, in which the construction is harder than just solving a set of linear equations in volatility.

Smile Version In the smile version, the relationships between risk reversal quoted in terms of volatility (RR) and butterfly/strangle (BF) quoted in terms of volatility are defined by

$$\mathrm{RR} = \sigma_+ - \sigma_-, \tag{145}$$

$$\mathrm{BF} = \frac{\sigma_+ + \sigma_-}{2} - \mathrm{ATM}, \tag{146}$$

and therefore, equivalently, the volatilities of 25-delta calls and puts are

$$\sigma_+ = \mathrm{ATM} + \mathrm{BF} + \tfrac{1}{2}\mathrm{RR}, \tag{147}$$

$$\sigma_- = \mathrm{ATM} + \mathrm{BF} - \tfrac{1}{2}\mathrm{RR}, \tag{148}$$

where $\mathrm{ATM} = \sigma_0$ denotes the at-the-money volatility of both put and call, see Table 1.7, σ_+ the volatility of an out-of-the-money call (usually $25\,\Delta$) and σ_- the volatility of an out-of-the-money put (usually $25\,\Delta$). Our sample market data is given in terms of RR and BF. Translated into implied volatilities of vanilla options we obtain the data listed in Table 1.10 and illustrated in Figure 1.9.

Brokers' Version In the more common brokers' version of the smile, (146) is replaced by

$$\sigma_{Str} + \sigma_{Str} \stackrel{\Delta}{=} 2(\mathrm{ATM} + \mathrm{BF}) = \sigma_+ + \sigma_- \tag{149}$$

$$\mathrm{Call}(K_{C25}, \sigma_{Str}) + \mathrm{Put}(K_{P25}, \sigma_{Str}) = \mathrm{Call}(K_{C25}, \sigma_+) + \mathrm{Put}(K_{P25}, \sigma_-), \tag{150}$$

where the 25-delta call strike K_{C25} and 25-delta put strike K_{P25} are calculated with the – yet to be found – 25-delta call volatility σ_+ and 25-delta put volatility σ_-. The non-linear (150) means that the LHS value of the 25-delta strangle is calculated as the sum of a call and a put value, both of which are wrong, because they are valued with the at-the-money volatility rather than with the volatilities from the actual FX smile. The RHS value adds up to the correct value, where also both summands are correct values. The smile version is much easier to solve because it is linear, and can

TABLE 1.10 EUR/GBP implied volatilities in % of 4 April 2005. Source: BBA. They are computed based on the market data displayed in Tables 1.7, 1.8 and 1.9 using Equations (147) and (148).

Maturity	25 delta put	at-the-money	25 delta call
1M	4.955	4.880	5.105
3M	5.400	5.340	5.600
1Y	6.030	5.990	6.295

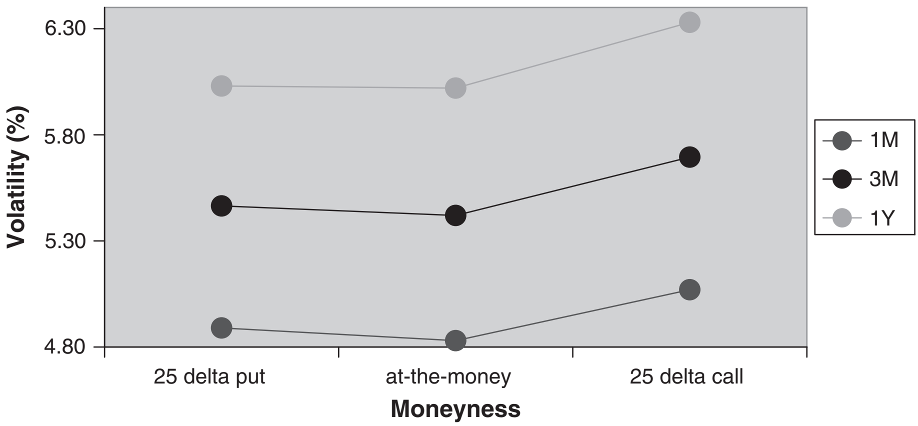

FIGURE 1.9 Implied volatilities for EUR-GBP vanilla options as of 4 April 2005. Source: BBA. Note that the dots are the only market input. Their interpolation and extrapolation are part of the volatility smile construction.

be considered a good approximation, especially if risk reversals are small. Note that the risk reversal (145) stays the same, and is not translated to the level of values of the building blocks. The strikes of the market strangle can also be determined using the at-the-money volatility σ_0. This is easier, but the delta of the strangle will no longer be zero after it is booked.

You might – legitimately – ask the question: why is this made so complicated? As a matter of legacy, brokers liked to quote just one number, which is the market strangle volatility σ_{Str}, to be interpreted approximately as the average volatility for 25-delta calls and puts. Figure 1.10 illustrates this.

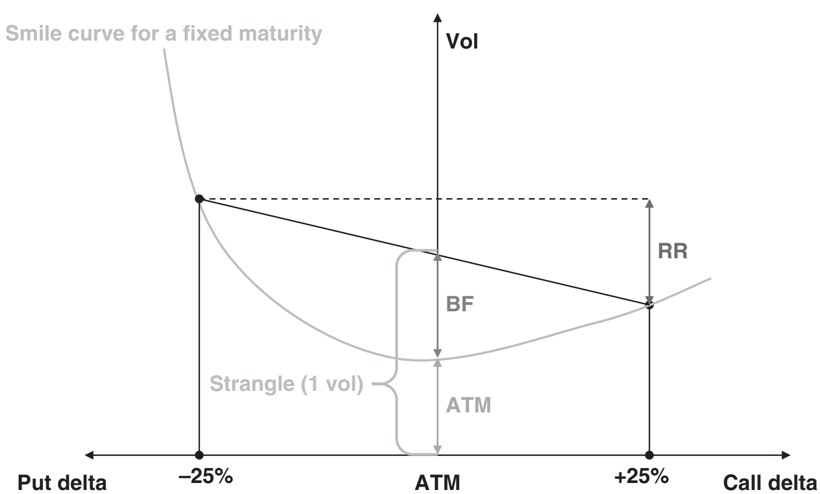

FIGURE 1.10 Relationship of risk reversal, butterfly, and market strangle volatility (the 1-vol strangle).

FX smile construction has been studied in detail in a number of papers by Dimitri Reiswich and myself, see [108, 109, 110, 111], and we will revisit the topic in [141] for quantitative matters. Bloomberg's method has been documented in [62]. An advanced parametric construction of the FX implied volatility surface is treated by Damghani in [35].

1.5.4 At-The-Money Volatility Interpolation

We denote by σ the volatility, which is the *spot volatility* for the time interval from horizon to expiry. The corresponding spot variance is denoted by $\sigma^2\tau$, where τ is the time difference between expiry and horizon. Corresponding forward volatilities σ_f and forward variances apply to a time interval starting later than horizon. The interpolation of at-the-money volatilities takes into account the effect of reduced volatility on weekends and on days closed in the global main trading centers London or New York and the local market, e.g. Tokyo for JPY trades. The change is done for the one-day forward volatility. You may apply a reduction in the one-day forward variance of 25% for each London and New York closed day. For local market holidays you may use a reduction of 25%, where local holidays for EUR are ignored. Weekends can be accounted for by a reduction to 15% variance. The variance on trading days is adjusted to match the volatility on the pillars of the ATM–volatility curve exactly. Obviously, the reduction percentages are arbitrary and different traders may have different opinions about these.

The procedure starts from the two pillars t_1, t_2 surrounding the date t_r in question. The ATM forward volatility for the period is calculated based on the consistency condition

$$\sigma^2(t_1)(t_1 - t_0) + \sigma_f^2(t_1, t_2)(t_2 - t_1) = \sigma^2(t_2)(t_2 - t_0), \tag{151}$$

which means that total variance over a time interval from t_0 to t_2 is the sum of all the variances of the sub-intervals, in this case the spot variance and the forward variance. Therefore,

$$\sigma_f(t_1, t_2) = \sqrt{\frac{\sigma^2(t_2)(t_2 - t_0) - \sigma^2(t_1)(t_1 - t_0)}{t_2 - t_1}}. \tag{152}$$

For each day the factor is determined, and from the constraint that the sum of one-day forward variances matches exactly the total variance, the factor for the enlarged one-day business variances $\alpha(t)$ with t business day is determined.

$$\sigma^2(t_1, t_2)(t_2 - t_1) = \sum_{t=t_1}^{t_r} \alpha(t)\sigma_f^2(t, t+1) \tag{153}$$

The variance for the period is the sum of variances to the start and sum of variances to the required date.

$$\sigma^2(t_r) = \sqrt{\frac{\sigma^2(t_1)(t_1 - t_0) + \sum_{t=t_1}^{t_r} \alpha(t)\sigma_f^2(t, t+1)}{t_r - t_0}} \tag{154}$$

1.5.5 Volatility Smile Conventions

The volatility smile is quoted in terms of delta and one at-the-money pillar. We recall that there are several notions of delta, namely

- spot delta $e^{-r_f\tau}N(d_+)$,
- forward delta $N(d_+)$,

and there is the premium which might be included in either delta. It is important to specify the notion that is used to quote the smile. There are four different deltas concerning plain vanilla options.

1.5.6 At-The-Money Definition

There is one specific at-the-money pillar in the middle. There are numerous notions for the meaning of *at-the-money (ATM)*.

Value parity: choose the strike such that call value = put value

Delta parity: choose the strike such that delta call = − delta put

Fifty delta: call delta = 50% and put delta = 50%

Maximum vega: choose the strike that maximizes vega of a call (or put)

Maximum gamma: choose the strike that maximizes gamma of a call (or put)

ATM spot: set strike equal to spot

Maximum time value: choose the strike that maximizes the time value of the call (or put)

Moreover, the notions involving delta use different versions of delta, namely either spot, forward, and premium included or excluded.

Obviously, put–call parity implies that *value parity* is equivalent to setting the strike equal to the outright forward rate. Therefore, *value parity* is also referred to as ATM forward. The *50-delta* notion is flawed, as for a fixed strike we cannot have both the call *and* the put have a spot delta of 50% (except in the case where $r_f = 0$). However, this flawed notion that an ATM option has a 50-delta is surprisingly widespread. Probably caused by herds of backyard academics writing books and papers about options ignoring interest and dividend rates. The 50-delta notion is only well defined when we use forward deltas and in that case becomes equivalent to *delta parity*. This notion is also referred to as *ATM delta-neutral*. I leave it to you to determine the details of which notions of ATM are equivalent. For risk management systems and common FX pricing tools there are two essential notions: *ATM forward* ($K = f$) and *ATM delta-neutral* ($K = K_\pm$) as derived in Equation (43) in Section 1.4.4. Entering “ATM” into the box for the strike in a pricing tool for EUR/USD options up to one-year tenor is expected to produce K_+ (K_- respectively for USD/JPY).

1.5.7 Interpolation of the Volatility on Fixed Maturity Pillars

Now let me provide a short introduction into the handling of FX implied volatility market data – especially their inter- and extrapolation across delta space and time.

We discuss a low-dimensional Gaussian kernel approach as an example showing several advantages over usual smile interpolation methods, for example cubic splines.

Before the discussion of specific interpolation methods it is recommended to take a step backwards and remember Rebonato's well-known statement of implied volatility as the wrong number in the wrong formula to obtain the right price [105]. So the explanatory power of implied volatilities for the dynamics of a stochastic process remains limited. Implied volatilities give a lattice on which marginal distributions can be constructed. However, even using many data points to generate marginal distributions, forward distributions, and extremal distributions – determining the prices of compound and barrier products, for example – cannot be uniquely defined by implied volatilities – see Tistaert *et al.* [127] for a discussion of this.

The attempt to capture FX smile features can lead into two different general approaches.

Parameterization One possibility for expressing smile or skew patterns is just to capture them as the calibration parameter set of an arbitrary stochastic volatility or jump diffusion model which generates the observed market implied volatilities. However, as spreads are rather narrow in liquid FX options markets, it is preferable to exactly fit the given input volatilities. This automatically leads to an interpolation approach.

Pure Interpolation As an introduction we would like to pose four requirements for an acceptable volatility surface interpolation:

1. Smoothness in the sense of continuous differentiability. Especially with respect to the possible application of Dupire-style local volatility models it is crucial to construct an interpolation which is at least C^2 (twice continuously differentiable). This becomes obvious when looking at the expression for the local volatility in this context:

$$\sigma_t^{local}(S(t)) = \left(2 \frac{\frac{\partial Call(S,t;K,T)}{\partial T} + \frac{\partial Call(S,t;K,T)}{\partial K}}{K^2 \frac{\partial^2 Call(S,t;K,T)}{\partial K^2}} \right)^{\frac{1}{2}}.$$

   Note in addition that local volatilities can directly be extracted from delta-based FX volatility surfaces, i.e. the Dupire formula can alternatively be expressed in terms of delta. See Hakala and Wystup [66] for details.
2. Absence of oscillations, which is guaranteed if the sign of the curvature of the surface does not change over different strike or delta levels.
3. Absence of arbitrage possibilities on single smiles of the surface as well as absence of calendar arbitrage.
4. A reasonable extrapolation available for the interpolation method.

A classical interpolation method widely spread are cubic splines. They attempt to fit surfaces by fitting piecewise cubic polynomials to given data points. They are specified by matching their second derivatives at each intersection. While this ensures the required smoothness by construction, it does not prevent oscillations – which directly leads to the danger of arbitrage possibilities – nor does it define how to extrapolate the smile. We therefore introduce the concept of a slice kernel volatility surface – as described in Hakala and Wystup [66] – as an alternative:

Definition 1.5.1 (Slice Kernel) *Let $(x_1, y_1), (x_2, y_2) \cdots, (x_n, y_n)$ be n given points and $g : I\!R \mapsto I\!R$ a smooth function which fulfills*

$$g(x_n) = y_n, \forall n = 1, \ldots, n. \tag{155}$$

A smooth interpolation is then given by

$$g(x) \triangleq \frac{1}{\Phi_\lambda(x)} \sum_{i=1}^{N} \alpha_i K_\lambda(\| x - x_i \|), \tag{156}$$

where

$$\Phi_\lambda(x) \triangleq \sum_{i=1}^{N} K_\lambda(\| x - x_i \|) \tag{157}$$

and

$$K_\lambda(u) \triangleq \exp\left\{-\frac{u^2}{2\lambda^2}\right\}. \tag{158}$$

The described kernel is also called a Gaussian kernel. The interpolation reduces to the determination of the α_i which is straightforward via solving a linear equation system. Note that λ remains a free smoothing parameter which also impacts the condition of the equation system. At the same time it can be used to fine-tune the extrapolation behavior of the kernel.

The idea behind this approach is as follows. The parameters which solve the interpolation conditions hyperref[interpolbed] (156) are $\alpha_1, \ldots, \alpha_n$. The parameter λ determines the "smoothness" of the resulting interpolation g and should be fixed according to the nature of the points (x_n, y_n). If these points yield a smooth surface, a "large" λ might yield a good fit, whereas in the opposite case when for neighboring points x_k, x_n the appropriate values y_k, y_n vary significantly, only a small λ, that means $\lambda << \min_{n,k} \| X_k - X_n \|$, can provide the needed flexibility.

For the set of delta pillars of $10\%, 25\%, ATM, -25\%, -10\%$ one can use $\lambda = 25\%$ for a smooth interpolation.

Normally the slice kernel produces reasonable output smiles based on a maximum of seven delta-volatility points. Then it fulfills all above mentioned requirements: it is C^∞, does not create oscillations, passes typical no-arbitrage conditions as they are posed by Gatheral [55], for example, and finally has an inherent extrapolation method.

In time direction one might connect different slice kernels by linear interpolation of the variances for same deltas. This also normally ensures the absence of calendar arbitrage, for which a necessary condition is a non-decreasing variance for constant moneyness F/K (see also Gatheral [55] for a discussion of this).

Figure 1.11 displays the shape of a slice kernel applied to a typical FX vol surface constructed from 10- and 25-delta volatilities, and the ATM volatility.

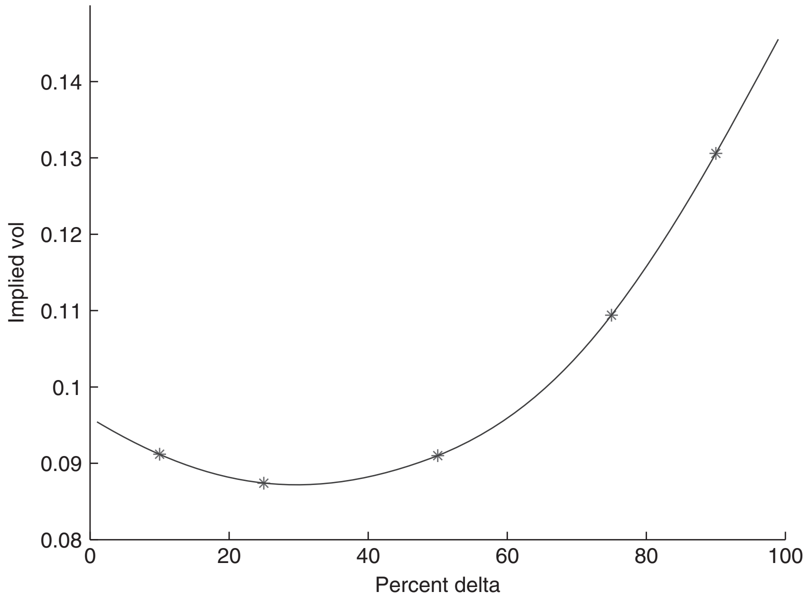

FIGURE 1.11 Kernel interpolation to generate an FX volatility smile.

Critical Judgment Kernels have been used in the industry going back to the 1990s. I had discussed slice kernels in [66] and [137]. However, kernels are by no means the only way to interpolate. Other methods include splines, stochastic volatility-based approaches (Heston, SABR), vanna-volga-based approaches. Malz' parabolic approach is discussed in the exercises in Section 1.5.12. The so-called stochastic volatility-inspired (SVI) approach reported by Gatheral [55] appears to be one of the most promising. A clever way of illustrating its practical implementation can be found in Zeliade [144]. We come back to the issue of interpolation and extrapolation on the smile in more detail in [141].

1.5.8 Interpolation of the Volatility Spread between Maturity Pillars

Interpolation of smile curves between maturity pillars is another issue. One can come up with numerous ideas, which we revisit in [141]. A simple method to interpolate the volatility spread to ATM uses the interpolation of the spread on the two surrounding maturity pillars for the initial Black–Scholes delta of the option. The spread is interpolated using square root of time where $\tilde{\sigma}$ is the volatility spread,

$$\tilde{\sigma}(t) = \tilde{\sigma}_1 + \frac{\sqrt{t} - \sqrt{t_1}}{\sqrt{t_2} - \sqrt{t_1}}(\tilde{\sigma}_2 - \tilde{\sigma}_1). \tag{159}$$

The spread is added to the interpolated ATM volatility as calculated above.

An example of a complete volatility surface with interpolation is shown in Figure 1.12.

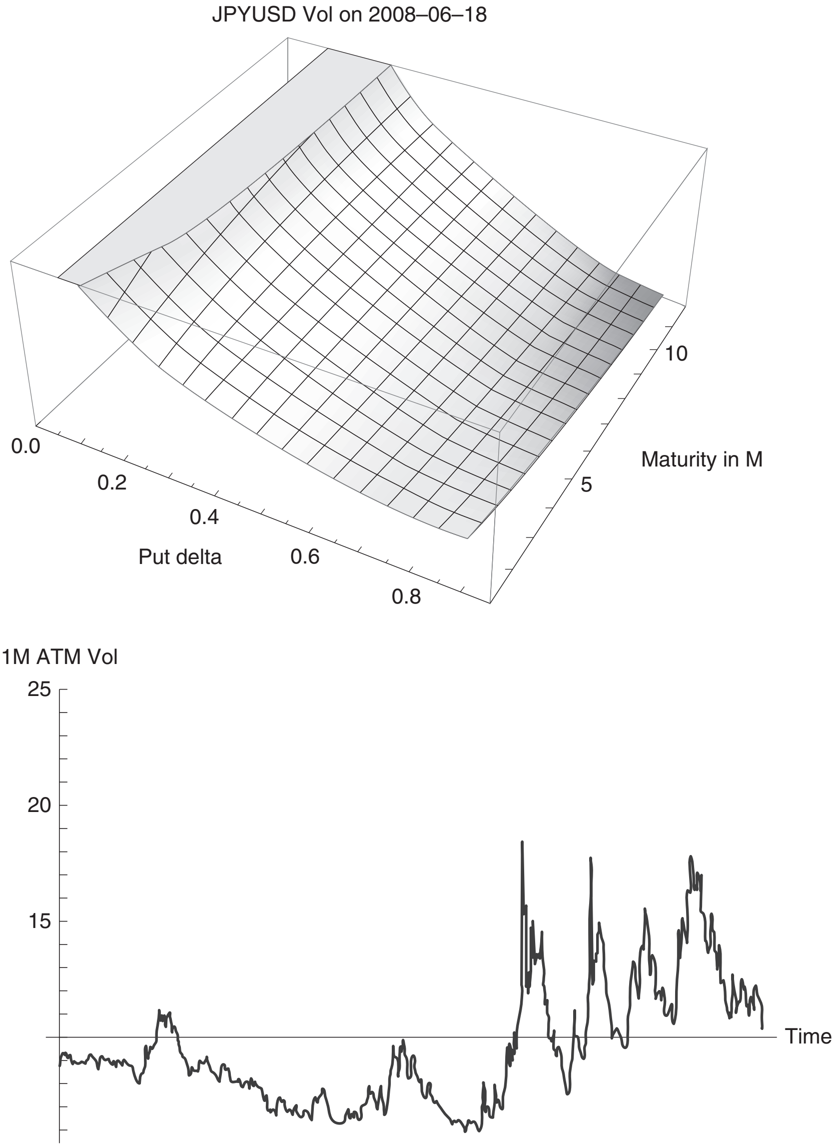

FIGURE 1.12 USD/JPY volatility surface on the delta space up to one-year tenor and historic ATM volatilities.

1.5.9 Volatility Sources

1. BBA, the British Bankers' Association, used to provide historic smile data for all major currency pairs in spreadsheet format at www.bba.org.uk.
2. Olsen Data (www.olsendata.com) can provide tic data of historic spot rates, from which the historic volatilities can be computed.

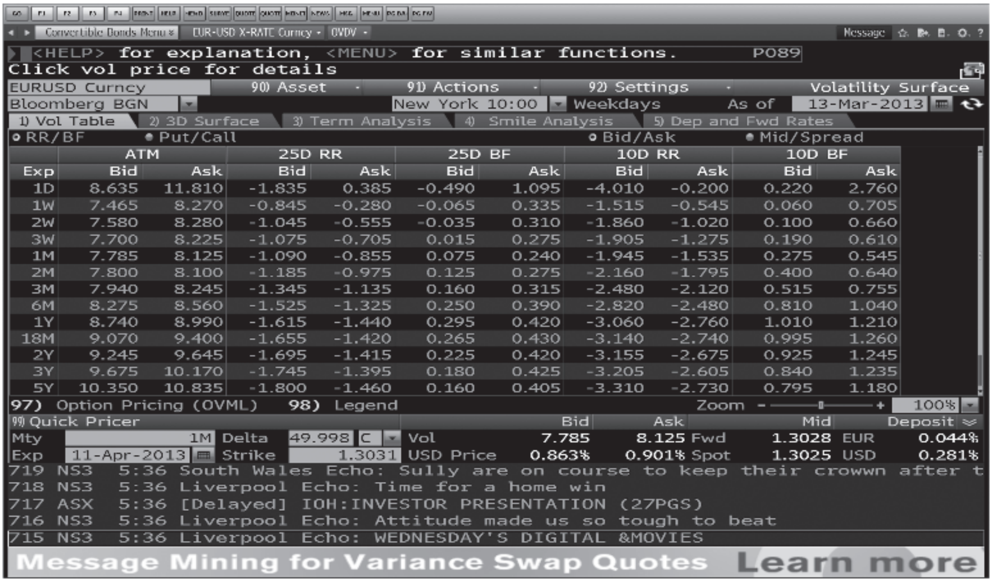

	ATM		25D RR		25D BF		10D RR		10D BF	
Exp	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask
1D	8.635	11.810	-1.835	0.385	-0.490	1.095	-4.010	-0.200	0.220	2.760
1W	7.465	8.270	-0.845	-0.280	-0.065	0.335	-1.515	-0.545	0.060	0.705
2W	7.580	8.280	-1.045	-0.555	-0.035	0.310	-1.860	-1.020	0.100	0.660
3W	7.700	8.225	-1.075	-0.705	0.015	0.275	-1.905	-1.275	0.190	0.610
1M	7.785	8.125	-1.090	-0.855	0.075	0.240	-1.945	-1.535	0.275	0.545
2M	7.800	8.100	-1.185	-0.975	0.125	0.275	-2.160	-1.795	0.400	0.640
3M	7.940	8.245	-1.345	-1.135	0.160	0.315	-2.480	-2.120	0.515	0.755
6M	8.275	8.560	-1.525	-1.325	0.250	0.390	-2.820	-2.480	0.810	1.040
1Y	8.740	8.990	-1.615	-1.440	0.295	0.420	-3.060	-2.760	1.010	1.210
18M	9.070	9.400	-1.655	-1.420	0.265	0.430	-3.140	-2.740	0.995	1.260
2Y	9.245	9.645	-1.695	-1.415	0.225	0.420	-3.155	-2.675	0.925	1.245
3Y	9.675	10.170	-1.745	-1.395	0.180	0.425	-3.205	-2.605	0.840	1.235
5Y	10.350	10.835	-1.800	-1.460	0.160	0.405	-3.310	-2.730	0.795	1.180

FIGURE 1.13 Bloomberg page OVDV quoting currency option volatilities.

3. Bloomberg provides both implied volatilities (OVDV) and historic volatilities, see Figure 1.13.
4. SuperDerivatives (www.superderivatives.com) besides being an internet pricing platform also provides FX volatility surfaces. An example of such a surface is presented in Figure 1.14.
5. Reuters pages such as FXMOX, SGFXVOL01, and others are commonly used and contain mostly implied volatilities. JYSKEOPT is a common reference for volatilities of Scandinavia (scandie-vols). NMRC has some implied volatilities for precious metals. An example of a EUR/USD volatility surface is presented in Figure 1.15. The common RICs (Reuters Instrument Codes) used for EUR-USD are EUR= for the spot, EUR6M= for the 6-month outright forward, EUR6MD= and USD6MD= for the 6-month money market deposit rates in EUR and USD, EUR6MO= for the 6-month ATM volatility, EUR6MRR=, EUR6MR10=, EUR6MBF=, EUR6MB10= for 6-month 25-delta and 10-delta risk reversals and butterflies. The entries of the EURVOLSURF page are available as RIC via EUR20P3M=R for the 6-month 20-delta EUR put volatility; the different contributing brokers can be addressed directly by stating EUR6MRR=FI for GFI (FENICS), =ICAP, =TPI for Tullett, =R for a recalculation performed by Reuters. Good luck!
6. Among the common brokers we find ICAP, GFI, or Tullett Prebon – see Figure 1.16 for an example of USD-JPY volatilities.
7. Telerate pages such as 4720 deliver implied volatilities.
8. Cantorspeed 90 also provides implied volatilities.

superderivatives DGX FX IR CM EQ CD

Volatility Surface Term Structure | Refresh | Calculate F2 | Chart

Negative 25ΔRR means favor EUR put

Click on the strike to open the pricing page

Trade date: Fri, 24 Jul 2015 Currency pair: EUR INR Spot: 69.98

	25ΔRR	25ΔFly		5 Δ	10 Δ	15 Δ	20 Δ	25 Δ	30 Δ	ATM	30 Δ	25 Δ	20 Δ	15 Δ	10 Δ	5 Δ
1W	0.520	0.275	Vol	11.58	11.26	11.10	11.00	10.92	10.87	10.910	11.28	11.45	11.61	11.82	12.12	12.68
			Strike	68.238	68.675	68.952	69.166	69.346	69.505	70.049	70.632	70.810	71.013	71.259	71.585	72.117
1M	0.785	0.273	Vol	10.77	10.57	10.48	10.42	10.39	10.36	10.500	10.98	11.17	11.37	11.62	11.96	12.59
			Strike	66.718	67.554	68.098	68.523	68.883	69.204	70.312	71.571	71.962	72.410	72.957	73.688	74.899
2M	0.884	0.308	Vol	10.89	10.69	10.59	10.53	10.50	10.47	10.625	11.16	11.38	11.61	11.88	12.27	12.96
			Strike	65.730	66.861	67.600	68.177	68.668	69.106	70.614	72.384	72.938	73.577	74.360	75.415	77.181
3M	0.998	0.349	Vol	11.28	11.06	10.96	10.90	10.86	10.83	11.000	11.61	11.85	12.11	12.42	12.84	13.62
			Strike	64.739	66.176	67.115	67.851	68.476	69.035	70.951	73.277	74.011	74.861	75.908	77.328	79.733
6M	1.239	0.409	Vol	11.72	11.45	11.32	11.24	11.19	11.17	11.400	12.14	12.43	12.73	13.09	13.60	14.53
			Strike	63.192	65.268	66.621	67.681	68.581	69.382	72.123	75.693	76.841	78.177	79.839	82.127	86.100
9M	1.439	0.451	Vol	12.04	11.68	11.51	11.41	11.35	11.33	11.619	12.45	12.79	13.11	13.50	14.05	15.08
			Strike	62.429	65.055	66.754	68.078	69.197	70.187	73.565	78.217	79.737	81.508	83.723	86.801	92.253
1Y	1.635	0.493	Vol	12.35	11.88	11.65	11.51	11.42	11.41	11.750	12.68	13.06	13.39	13.79	14.34	15.40
			Strike	61.528	64.676	66.690	68.248	69.556	70.701	74.591	80.219	82.088	84.258	86.976	90.767	97.563
2Y	2.038	0.575	Vol	12.48	11.95	11.70	11.55	11.46	11.46	11.924	13.05	13.50	13.90	14.38	15.05	16.33
			Strike	60.684	65.157	68.038	70.272	72.151	73.792	79.280	88.400	91.555	95.270	100.012	106.813	119.568
				5 Δ	10 Δ	15 Δ	20 Δ	25 Δ	30 Δ	ATM	30 Δ	25 Δ	20 Δ	15 Δ	10 Δ	5 Δ

FIGURE 1.14 SuperDerivatives displaying EUR/INR option volatilities. The figures in boxes are meant to be read as market input; all other figures are calculated by SuperDerivatives' proprietary interpolation and extrapolation methods.

Quote: EURVOLSURF

EURVOLSURF EUR VOLSURFACE

Tue 18 Dec 2007 13:24 GMT Standard Time Logical Displays<0#EURVOL=R>

	10DPUT	15DPUT	20DPUT	25DPUT	30DPUT	35DPUT	40DPUT	45DPUT	ATM	45DCALL	40DCALL
1W	11.150	11.056	10.951	10.825	10.671	10.499	10.324	10.161	10.025	9.926	9.863
2W	10.975	10.840	10.704	10.563	10.416	10.269	10.126	9.993	9.875	9.778	9.706
3W	10.763	10.574	10.393	10.225	10.075	9.943	9.826	9.721	9.625	9.538	9.469
1M	10.825	10.460	10.122	9.838	9.625	9.474	9.365	9.280	9.200	9.113	9.038
6W	10.826	10.373	9.956	9.611	9.364	9.199	9.090	9.010	8.934	8.845	8.766
2M	10.600	10.207	9.843	9.537	9.312	9.153	9.041	8.955	8.875	8.788	8.714
3M	10.725	10.303	9.914	9.588	9.348	9.182	9.066	8.980	8.900	8.813	8.739
4M	10.662	10.240	9.848	9.518	9.271	9.097	8.973	8.881	8.799	8.715	8.647
5M	10.365	9.999	9.656	9.360	9.128	8.952	8.819	8.716	8.628	8.549	8.492
6M	10.075	9.756	9.455	9.188	8.968	8.793	8.653	8.541	8.450	8.374	8.326
9M	10.500	9.993	9.526	9.138	8.857	8.668	8.543	8.454	8.375	8.287	8.216
1Y	10.525	10.001	9.517	9.112	8.817	8.614	8.478	8.382	8.300	8.215	8.151
2Y	10.475	9.930	9.428	9.013	8.715	8.518	8.394	8.314	8.250	8.183	8.135

FIGURE 1.15 Reuters displaying EUR/USD option volatilities. It is not obvious to me (but hopefully to Reuters and its subscribers) which figures are used as market input and how the rest of the figures are interpolated or extrapolated.

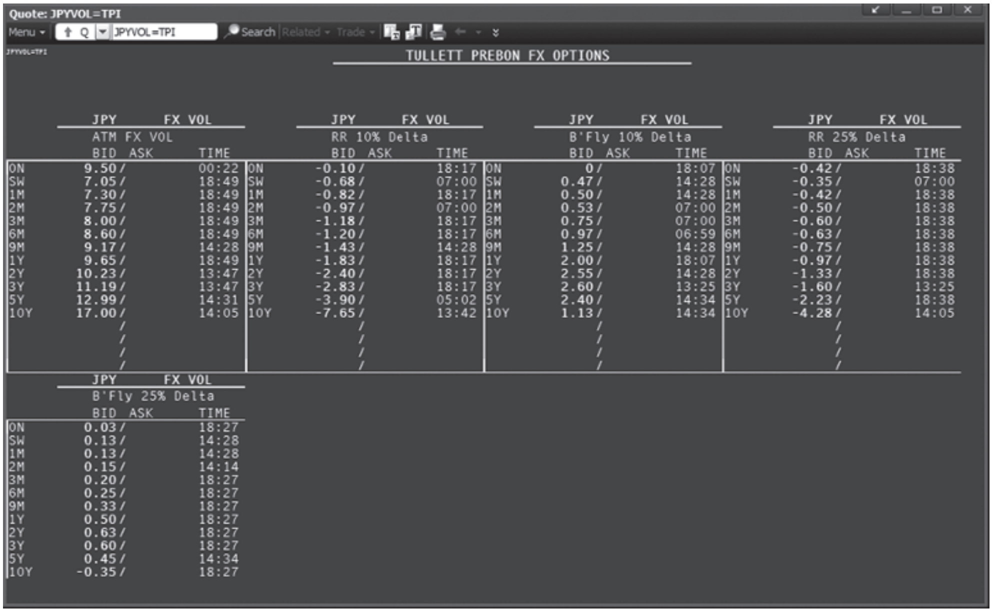

Quote: JPYVOL=TPI

TULLETT PREBON FX OPTIONS

	JPY FX VOL ATM FX VOL BID ASK	TIME	JPY FX VOL RR 10% Delta BID ASK	TIME	JPY FX VOL B'Fly 10% Delta BID ASK	TIME	JPY FX VOL RR 25% Delta BID ASK	TIME	JPY FX VOL B'Fly 25% Delta BID ASK	TIME
ON	9.50 /	00:22	-0.10 /	18:17	0 /	18:07	-0.42 /	18:38	0.03 /	18:27
SW	7.05 /	18:49	-0.68 /	07:00	0.47 /	14:28	-0.35 /	07:00	0.13 /	14:28
1M	7.30 /	18:49	-0.82 /	18:17	0.50 /	14:28	-0.42 /	18:38	0.13 /	14:28
2M	7.75 /	18:49	-0.97 /	07:00	0.53 /	07:00	-0.50 /	18:38	0.15 /	14:14
3M	8.00 /	18:49	-1.18 /	18:17	0.75 /	07:00	-0.60 /	18:38	0.20 /	18:27
6M	8.60 /	18:49	-1.20 /	18:17	0.97 /	06:59	-0.63 /	18:38	0.25 /	18:27
9M	9.17 /	14:28	-1.43 /	14:28	1.25 /	14:28	-0.75 /	18:38	0.33 /	18:27
1Y	9.65 /	18:49	-1.83 /	18:17	2.00 /	18:07	-0.97 /	18:38	0.50 /	18:27
2Y	10.23 /	13:47	-2.40 /	18:17	2.55 /	14:28	-1.33 /	18:38	0.63 /	18:27
3Y	11.19 /	13:47	-2.83 /	18:17	2.60 /	13:25	-1.60 /	13:25	0.60 /	18:27
5Y	12.99 /	14:31	-3.90 /	05:02	2.40 /	14:34	-2.23 /	18:38	0.45 /	14:34
10Y	17.00 /	14:05	-7.65 /	13:42	1.13 /	14:34	-4.28 /	14:05	-0.35 /	18:27

FIGURE 1.16 Tullett Prebon quoting USD-JPY volatilities of 14 April 2014 in terms of at-the-money, risk reversals and butterflies.

The way market data is organized can be best seen in *Volmaster*, as shown in Figure 1.17.

1.5.10 Volatility Cones

Volatility cones visualize whether current at-the-money volatility levels for various maturities are high or low compared with a recent history of these implied volatilities. This may provide information on whether it is currently advisable to buy volatility or sell volatility, i.e. to buy vanilla options or to sell vanilla options. We fix a time horizon of historic observations of mid-market at-the-money implied volatility and look at the maximum, the minimum that traded over this time horizon, and compare this with the current volatility level. Since long-term volatilities tend to fluctuate less than short-term volatility levels, the chart of the minimum and the maximum typically looks like a part of a cone. We illustrate this in Figure 1.18 based on the data provided in Table 1.11.

1.5.11 Stochastic Volatility

Stochastic volatility models are very popular in FX options, whereas *jump diffusion models* can be considered as the cherry on the cake and are left to currency pairs with high jump risk. The most prominent reason for the popularity is very simple: FX volatility *appears* stochastic, as *is* shown for instance in Figure 1.19. Treating stochastic volatility in detail here is way beyond the scope of this book. A more recent overview

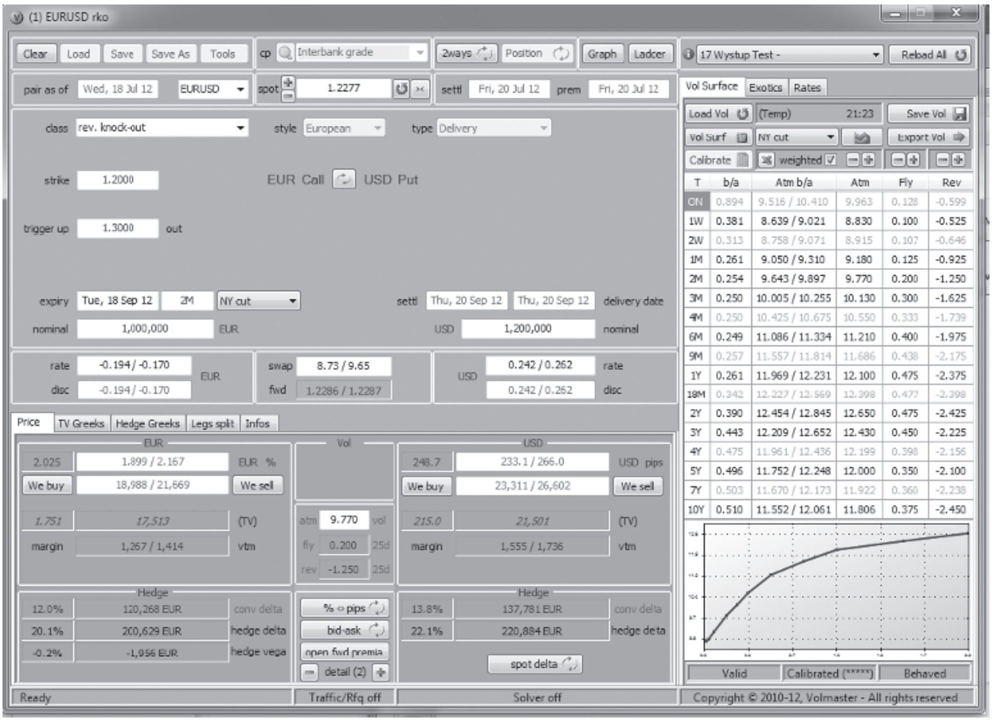

FIGURE 1.17 Volmaster single leg pricing screen with market input data on the right. "Fly" and "Rev" represent butterfly and risk reversal respectively.

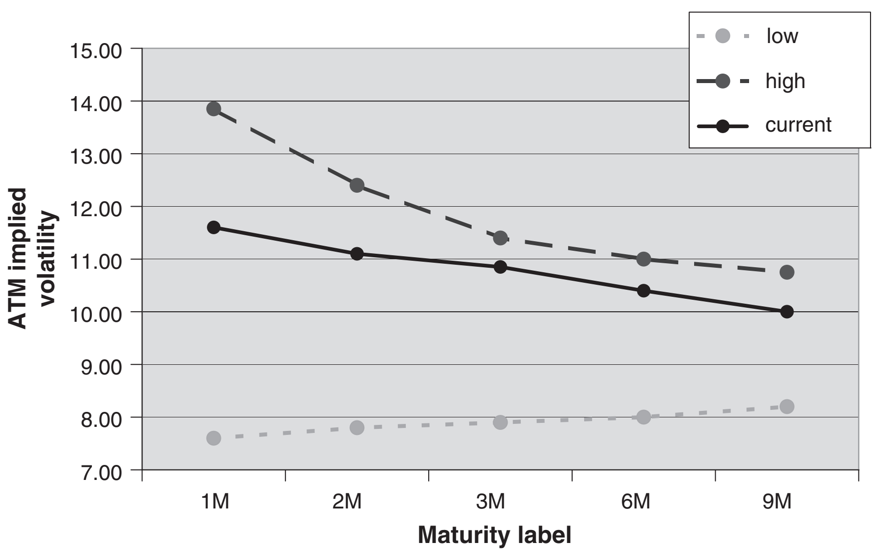

FIGURE 1.18 Example of a volatility cone in USD-JPY for a 6-month time horizon from 6 September 2003 to 24 February 2005.

TABLE 1.11 Sample data of a volatility cone in USD-JPY for a 6-month time horizon from 6 September 2003 to 24 February 2005.

Maturity	Low	High	Current
1M	7.60	13.85	11.60
2M	7.80	12.40	11.10
3M	7.90	11.40	10.85
6M	8.00	11.00	10.40
12M	8.20	10.75	10.00

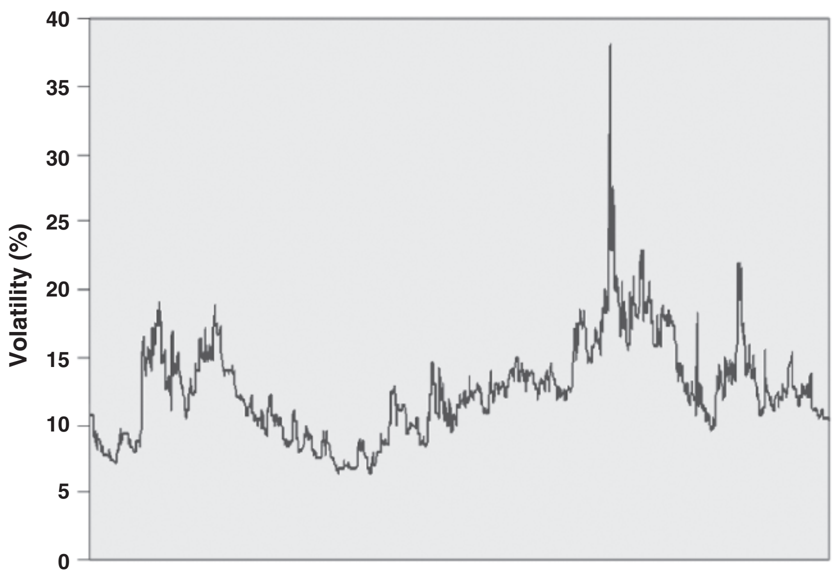

FIGURE 1.19 Historic implied volatilities for USD-JPY 1-month vanilla at-the-money options for the period 1994–2000.

can be found in the article *The Heston Model and the Smile* in Weron and Wystup [132] and subsequently by Janek *et al*. [81].

1.5.12 Exercises

Converting Risk Reversals and Strangles into Volatilities For the market data in Tables 1.7, 1.8 and 1.9 determine a smile matrix for at-the-money and the 25-deltas. Also compute the corresponding strikes for the three pillars or moneyness. We assume the continuous EUR interest rate to be $r_f = 3\%$ and the continuous GBP interest rate to be $r_d = 5\%$. The at-the-money notion here is spot delta without premium adjustment for the delta-neutral straddle.

Malz Parabolic Interpolation One way of generating a quick (and semi-dirty) volatility smile from ATM, risk reversals, and butterflies is to fit the three pieces of information to a parabola on the delta space. Assume you want to generate a parabolic function

$$\sigma(\Delta_f) = A + B(\Delta_f - 50\%) + C(\Delta_f - 50\%)^2, \tag{160}$$

where the variable is the forward delta Δ_f, the market input parameters are ATM, RR, and BF (in %), and σ is the desired parabolic interpolation function. Determine A, B, and C as functions of ATM, RR, and BF quoted in the smile version (not the brokers' version), and hereby derive the formula by Malz [95].

Slope of the Malz-Smile on the Strike Space Using the Malz parabola [95]

$$\sigma(\Delta_f) = A + B(\Delta_f - 50\%) + C(\Delta_f - 50\%)^2 \tag{161}$$

of the previous exercise, calculate the slope of the smile on the strike space in the Black-Scholes model, i.e. determine the function $\sigma'(K)$. Show that for the ATM delta-neutral strike, the windmill-adjustment (which is vanilla vega times $\sigma'(K)$) converges to

$$-\frac{\text{RR}}{\pi\sigma_{\text{ATM}}}, \tag{162}$$

as time to maturity turns to zero. As usual RR denotes the risk reversal. Assume that RR and the ATM volatility σ_{ATM} are constant over time.

Volatility Cones Using the historic data, generate a volatility cone for USD-JPY.

At-the-money Delta for Long-Term Options It is often believed that an at-the-money (in the sense that the strike is set equal to the forward) vanilla call has a delta near 50%. What can you say about the delta of a 15-year at-the-money USD-JPY call if USD rates are at 5%, JPY rates are at 1%, and the volatility is at 11%?

Vega Hedging ATM Calls Suppose you have just sold a 2Y ATM call. How many 1Y ATM calls do you need to buy to be vega neutral? You may assume the ATM convention to be delta-neutral.

Which Forward Rate Does the Smile Surface Use? Suppose you are provided with a EUR/USD volatility smile matrix as in Reuters VOLSURF, and for 6 months (182 days) you see a 25-delta put volatility 12.191% for strike 1.0273, 10-delta put volatility 14.147% for strike 0.9591: which forward rate and EUR money market rate are used? You may assume that delta is spot delta, premium excluded.

1.6 BASIC STRATEGIES CONTAINING VANILLA OPTIONS

Linear combinations of vanillas are quite well known and have been explained in several textbooks, including the one by Spies [121]. Therefore, we will restrict our attention in this section to the most basic strategies.

1.6.1 Call and Put Spread

A call spread is a combination of a long and a short call option. It is also called *capped call*. The motivation to do this is the fact that buying a simple call may be too expensive and the buyer wishes to lower the premium. At the same time he does not expect the underlying exchange rate to appreciate above the strike of the short call option.

The call spread entitles the holder to buy an agreed amount of a currency (say EUR) on a specified date (maturity) at a pre-determined rate (long strike) as long as the exchange rate is above the long strike at maturity. However, if the exchange rate is above the short strike at this time, the holder's profit is limited to the spread as defined by the short and long strikes (see example below). Buying a call spread provides protection against a rising EUR with full participation in a falling EUR. The holder has to pay a premium for this protection. The holder will typically exercise the option at maturity if the spot is above the long strike.

Advantages

- Protection against stronger EUR/weaker USD
- Low-cost product
- Maximum loss is the premium paid

Disadvantages

- Protection is limited when the exchange rate is above the short strike at maturity

The buyer has the chance of full participation in a weaker EUR/stronger USD. However, in case of very high EUR at maturity the protection works only up to the higher strike.

For example, a company wants to buy 1 M EUR. At maturity:

1. If $S_T < K_1$, it will not exercise the option. The overall loss will be the option's premium. But instead the company can buy EUR at a lower spot in the market.
2. If $K_1 < S_T < K_2$, it will exercise the option and buy EUR at strike K_1.
3. If $S_T > K_2$, it will buy the 1 M EUR at a rate $K_2 - K_1$ below S_T.

Example A company in the EUR-zone wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD. The company wishes to be able to buy EUR at a lower spot rate if the EUR weakens on the one hand, but on the other hand be protected against a stronger EUR. The vanilla call is too expensive, but the company does not expect a large upward movement of the EUR.

In this case a possible form of protection that the company can use is to buy a call spread as for example listed in Table 1.12. The payoff and effective final exchange rate are exhibited in Figure 1.20.

TABLE 1.12 Example of a call spread.

Spot reference	1.1500 EUR-USD
Company buys	EUR call USD put with lower strike
Company sells	EUR call USD put with higher strike
Maturity	1 year
Notional of both call options	EUR 1,000,000
Strike of the long call option	1.1400 EUR-USD
Strike of the short call option	1.1800 EUR-USD
Premium	USD 20,000.00
Premium of the long EUR call only	USD 63,000.00

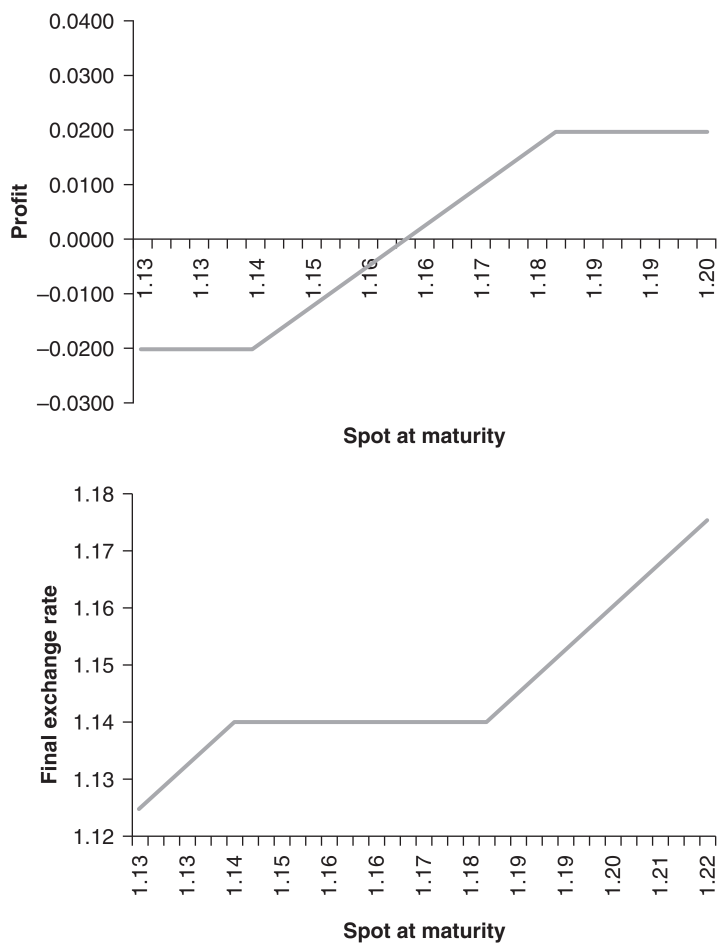

FIGURE 1.20 Profit & loss and final exchange rate of a call spread.

- If the company's market expectation is correct, it can buy EUR at maturity at the strike of 1.1400.
- If the EUR-USD exchange rate is below the strike at maturity, the option expires worthless. However, the company would benefit from a lower spot when buying EUR.
- If the EUR-USD exchange rate is above the short strike of 1.1800 at maturity, the company can buy the EUR amount 400 pips below the spot. Its risk is that the spot at maturity is very high.

The EUR seller can buy a EUR put spread in a similar fashion.

Critical Assessment The Call Spread lowers the cost of the protection against a rising EUR for the treasurer, but fails to protect the extreme risk. Viewed as an insurance it covers small losses but fails to cover potential big losses. One can use it, but we would want to be sure the treasurer understands the consequences. The situation is different for a different client type: the investor, i.e. the market participant *without* the underlying cash flow. The investor buying a Call Spread merely waives participation in an extreme rise of the underlying market. This is why Call Spreads are commonly used in private banking and retail banking.

Ratio Call Spread A variation of a Call Spread is a *Ratio Call Spread*. The treasurer/investor sells more than one Call with the higher strike, e.g. two Calls. The Call Spread becomes *leveraged* with a leverage of 2. The incentive is to generate a strategy that is even cheaper than the Call Spread, ideally in fact zero cost. One can achieve this by lowering the higher strike or by increasing the leverage. The problem will then be that the position of a long Ratio Call Spread can become negative if spot increases substantially. Let me tell you the story of a Turkish trader who set up a speculative position in USD-TRY in 2008. The trader used a highly leveraged ratio call spread on a margined account following a trade idea that the Turkish Lira would depreciate over the coming months, but not become weaker than 1.6000 USD-TRY, a level the trader had chosen for the higher strike, as illustrated in Figure 1.21.

It turned out that the trader had underestimated the Gamma/Vega exposure with rising spot, which consequently led to a high VaR. The trader received repeated margin calls, could eventually not meet them, and the bank had force-closed his position. The trader filed a claim against the bank stating that the close-out was unfair, and that eventually the Lira did what he had predicted.

The claimant being an experienced SuperDerivatives user alleged he had never heard of smile. The effect of smile on the valuation is illustrated in Figure 1.22. What went wrong? At the trade's inception, the value of the position was close to zero, because the leverage of the Ratio Call Spread was designed like that. Furthermore, since the vega and gamma of the long Call and the vega and gamma of the short Call neutralized each other, the trader was deceived by spotting a zero-cost and zero-risk strategy in his portfolio. Then, ignoring the smile effect and underestimating the change of the vega/gamma position with a rising USD-TRY spot led to unexpected continued margin calls. It was impossible to navigate this position through the October 2008 financial crisis with a limited amount of cash.

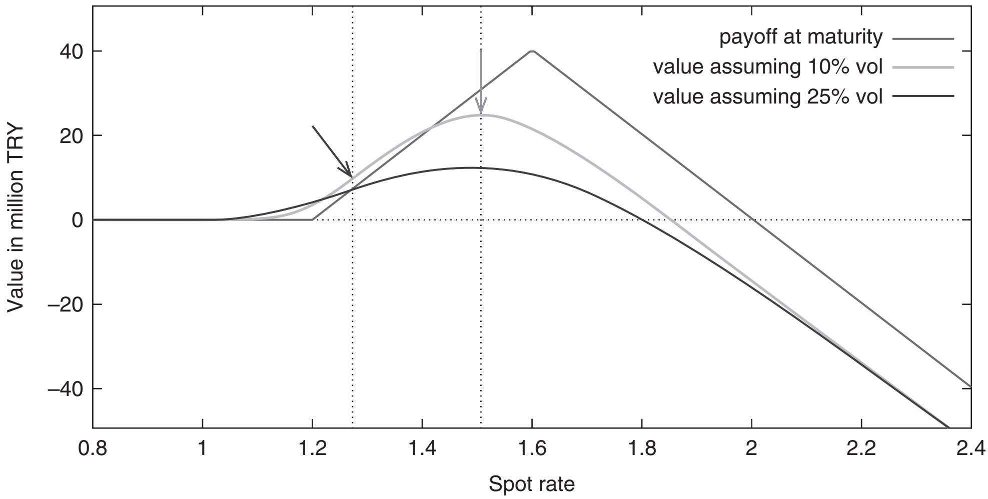

FIGURE 1.21 Position of a Ratio Call Spread reflecting the view of a rise and sharp landing of USD-TRY at about 1.6000 in 2008.

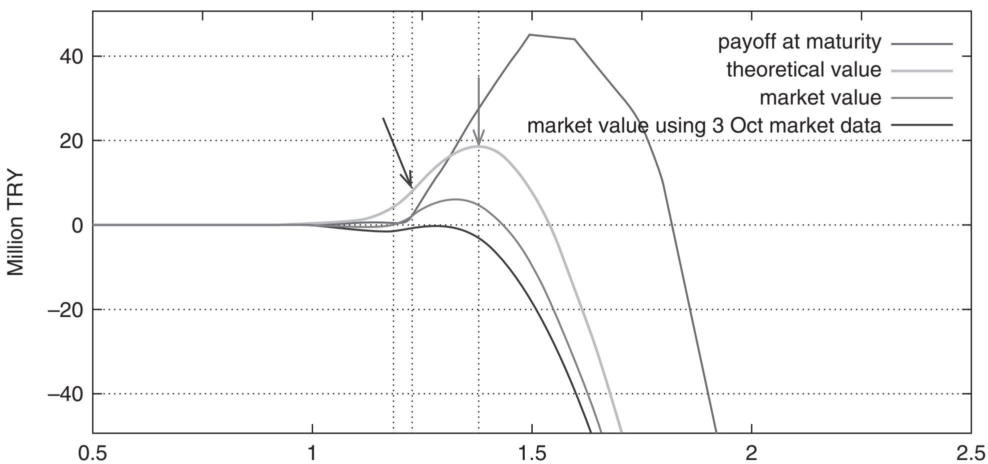

FIGURE 1.22 Smile effect in a ratio call spread in USD-TRY.

Loss Calculation In order to calculate the loss of the position, which was delta-hedged on top of everything for a reason I failed to understand, the claimant's dodgy expert calculated the loss amount by constructing a "zero gamma – long theta" portfolio.

- Consider the change of value of the option book V:

$$V(S + \delta S, t + \delta t) = V(S, t) + \Delta\ \delta S + \frac{1}{2}\Gamma\ \delta S^2 + \Theta\ \delta t + o(\delta t, \delta S^2)$$

- Then the change of the delta-hedged option portfolio $P(S,t) \overset{\Delta}{:=} V(S,t) + h\,S$ becomes

$$P(S+\delta S, t+\delta t) = P(S,t) + (h+\Delta)\delta S + \frac{1}{2}\Gamma\ \delta S^2 + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2).$$

- The P&L at time $t+\delta t$ is

$$\begin{aligned} &P(S+\delta S, t+\delta t) - P(S,t) \\ &= (h_t+\Delta)\delta S + \frac{1}{2}\Gamma\ \delta S^2 + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2) \end{aligned} \tag{163}$$

- The standard delta hedge $h_t := -\Delta$ yields

$$\text{P\&L at time } (t+\delta t) = \frac{1}{2}\Gamma\ \delta S^2 + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2)$$

The dodgy expert had argued that the spot reference for the delta hedge was unknown because the delta hedge had been executed during the day, and that therefore one must use the average of the previous day's end of day spot and the current day's end of day spot as a proxy for the spot reference for the intra-day delta hedge. While this may sound reasonable to an outsider, it is obvious that this average approach uses the information of the future spot, and is therefore a crystal ball approach. Some basic calculations make this very clear. Using the dodgy expert's "average hedge"

$$h_t = -\frac{\Delta_{t+\delta t} + \Delta_t}{2} = -\frac{V'(S+\delta S, t+\delta t) + V'(S,t)}{2},$$

the fact that $\Delta_t = V'(S,t)$, and Equation (163) we obtain the P&L at time $t+\delta t$ as

$$\begin{aligned} &\delta S(h_t + \Delta_t) + \frac{1}{2}\delta S^2\ \Gamma + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2) \\ &= \delta S \frac{-\Delta_{t+\delta t} - \Delta_t + 2\Delta_t}{2} + \frac{1}{2}\delta S^2\ \Gamma + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2) \\ &= -\delta S \frac{\Delta_{t+\delta t} - \Delta_t}{2} + \frac{1}{2}\delta S^2\ \Gamma + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2) \\ &= -\delta S \frac{1}{2} V''(S,t)\delta S + \frac{1}{2}\delta S^2\ \Gamma + \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2), \\ &= \Theta\ \delta t + \mathrm{o}(\delta t, \delta S^2). \end{aligned}$$

The dodgy expert introduced a systematic error in the P&L by artificially removing the entire gamma risk. The claim was based on artificially generated money. The mistakes were spotted by the court rather quickly and the Turkish trader withdrew his claim. As a conclusion I would like to reiterate that even simple vanilla structures can cause surprises and losses, which are unpleasant for all parties involved. Issues like leverage and short options should always be carefully discussed with the buy-side to avoid such surprises and losses.

1.6.2 Risk Reversal

Very often corporates seek so-called zero-cost strategies to hedge their international cash flows. Since buying a call requires a premium, the buyer can sell another option to finance the purchase of the call. A popular liquid product in FX markets is the risk reversal or collar or range forward. The term *cylinder* is also used as a synonym for the Risk Reversal, or more often actually refers to a more general form of a Risk Reversal to distinguish it from the standard case. A risk reversal is a combination of a long call and a short put. It entitles the holder to buy an agreed amount of a currency (say EUR) on a specified date (maturity) at a pre-determined rate (long strike) assuming the exchange rate is above the long strike at maturity. However, if the exchange rate is below the strike of the short put at maturity, the holder is obliged to buy the amount of EUR determined by the short strike. Therefore, buying a risk reversal provides full protection against rising EUR. The holder will typically exercise the option only if the spot is above the long strike at maturity. The risk on the upside is financed by a risk on the downside. Since the risk is reversed, the strategy is named Risk Reversal.

Advantages

- Full protection against stronger EUR/weaker USD
- Can be structured as a zero-cost strategy

Disadvantages

- Participation in weaker EUR/stronger USD is limited to the strike of the sold put

For example, a company wants to sell 1 M USD. At maturity T:

1. If $S_T < K_1$, it will be obliged to sell USD at K_1. Compared with the market spot the loss can be large. However, compared with the outright forward rate at inception of the trade, K_1 is usually only marginally worse.
2. If $K_1 < S_T < K_2$, it will not exercise the call option. The company can trade at the prevailing spot level.
3. If $S_T > K_2$, it will exercise the option and sell USD at strike K_2.

Example A company wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD. The company wishes to be fully protected against a stronger EUR. But it finds that the corresponding plain vanilla EUR call is too expensive and would prefer a zero-cost strategy by financing the call with the sale of a put.

In this case a possible form of protection that the company can use is to buy a risk reversal, for example as indicated in Table 1.13 and exhibited in Figure 1.23.

If the company's market expectation is correct, it can buy EUR at maturity at the strike of $K_2 = 1.2250$. The risk is when the EUR-USD exchange rate is below the strike of $K_1 = 1.0775$ at maturity, the company is obliged to buy 1 M EUR at the rate of 1.0775. $K_2 = 1.2250$ is the guaranteed worst case, which can be used as a budget rate.

TABLE 1.13 Example of a Risk Reversal.

Spot reference	1.1500 EUR-USD
Company buys	EUR call USD put with higher strike
Company sells	EUR put USD call with lower strike
Maturity	1 year
Notional of both options	EUR 1,000,000
Strike of the long Call option	1.2250 EUR-USD
Strike of the short Put option	1.0775 EUR-USD
Premium	EUR 0.00

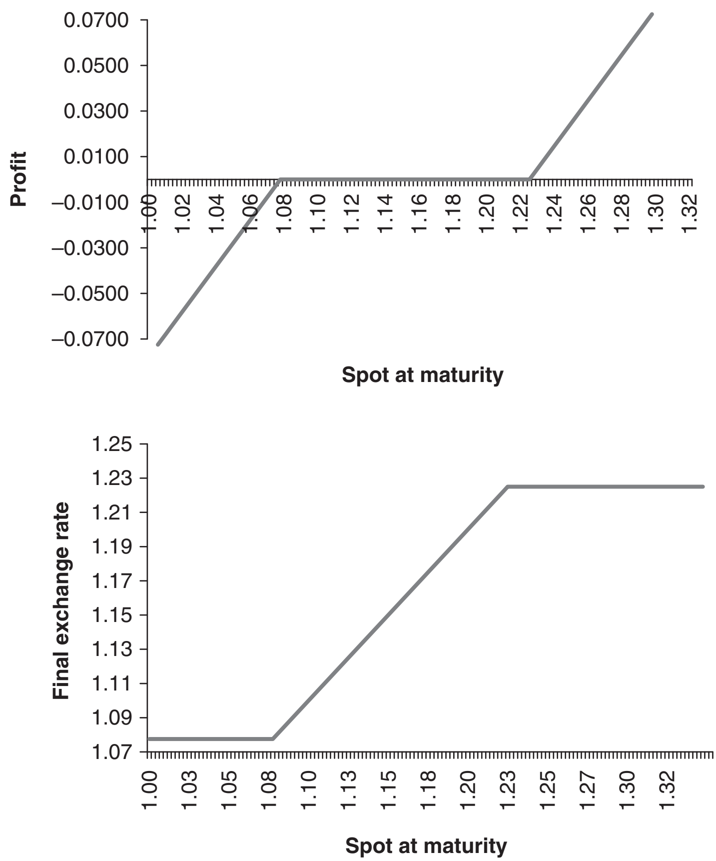

FIGURE 1.23 Payoff and final exchange rate of a risk reversal.

Critical Assessment The risk of the spot ending below the put strike is a risk only if the strategy was traded by an investor *without* the underlying cash flow. For a treasurer using the Risk Reversal as a hedge for an existing cash flow, the situation is not risky at all; in fact the treasurer can normally still buy the EUR at a rate lower than the outright forward rate he could have agreed on at inception of the trade. A common misconception is to judge the quality of a strategy based on the prevailing spot reference and compare the strategy with the "do nothing; wait for a better spot" strategy. This is not fair because the decision about how to hedge has to be taken at inception of the trade when the future spot is still unknown. The Risk Reversal is one of the most popular hedging strategies for corporate treasurers, which is why the brokers' market quote risk reversals directly, rather than stand-alone out-of-the-money option volatilities – see Section 1.5.3 on the symmetric decomposition of the smile, where risk reversal is a market quote for the skew in the smile curve. The risk-warehousing sell-side uses risk reversals as a vanna hedge, i.e. a hedge of the skew of the volatility surface – see the exercises in Section 1.6.9.

Risk Reversal Flip As a variation of the standard risk reversal, we consider the following trade on EUR/USD spot reference 1.2400 with a tenor of two months.

1. Long 1.2500/1.1900 risk reversal (long 1.2500 EUR call, short 1.1900 EUR put)
2. If 1.3000 trades before expiry, it flips into a 1.2900/1.3100 risk reversal (long 1.2900 EUR put, short 1.3100 EUR call)
3. Zero premium

The corresponding view is that EUR/USD looks bullish and may break on the upside of a recent trading range. However, a runaway higher EUR/USD setting a new all-time high within two months looks unlikely. However, if EUR/USD overshoots to 1.30, then it will likely retrace thereafter.

The main thrust is to long EUR/USD for zero cost, with a safe cap at 1.30. So the initial risk is EUR/USD below 1.19. If 1.30 is breached, then all accrued profit from the 1.25/1.19 risk reversal is lost, and the maximum risk becomes spot levels above 1.31. Therefore, this trade is not suitable for EUR bulls who feel there is scope above 1.30 within two months. On the other hand, this trade is suitable for those who feel that if spot overshoots to 1.30, then it will retrace down quickly. For early profit taking: with two weeks to go and spot at 1.28, this trade should be worth approximately 0.84 % EUR. However, with that spot the maximum profit occurs at the trade's maturity.

Composition Clearly, this risk reversal flip is rather a proprietary trading strategy than a corporate hedging structure.

The composition is presented in Table 1.14. The options used are standard barrier options, see Section 1.7.3.

TABLE 1.14 Example of a Risk Reversal flip.

client buys	1.2500 EUR call up-and-out at 1.3000
client sells	1.1900 EUR put up-and-out at 1.3000
client buys	1.2900 EUR put up-and-in at 1.3000
client sells	1.3100 EUR call up-and-in at 1.3000

1.6.3 Straddle

A straddle is a combination of a put and a call option with the same strike. At inception of the trade the strike is usually the delta-neutral at-the-money strike $K_{\pm}$ as in Equation (43). It entitles the holder to buy an agreed amount of a currency (say EUR) on a specified date (maturity) at a pre-determined rate (strike) if the exchange rate is above the strike at maturity. Alternatively, if the exchange rate is below the strike at maturity, the holder is entitled to sell the amount at this strike. Buying a straddle provides participation in both an upward and a downward movement where the direction of the rate is unclear. The buyer has to pay a premium for this product.

Advantages

- Full participation in or protection against market movement or increasing volatility
- Maximum loss is the premium paid

Disadvantages

- Expensive product
- Not suitable as directional hedge and consequently not suitable for hedge accounting of spot positions as it should be clear if the client wants to sell or buy EUR

Potential profits of a long straddle arise from movements in the spot and also from increases in implied volatility. If the spot moves, the call or the put can be sold before maturity with profit. Conversely, if a quiet market phase persists the strategy is unlikely to generate much revenue.

Figure 1.24 shows the payoff of a long straddle. The payoff of a short straddle looks like the straddle below a seesaw on a children's playground, which is where the name straddle originated.

For example, a company buys a Straddle with a nominal of 1 M EUR. At maturity T:

1. If $S_T < K$, it would sell 1 M EUR at strike K.
2. If $S_T > K$, it would buy 1 M EUR at strike K.

Example A company wants to benefit from its view that the EUR-USD exchange rate will move far from a specified strike (Straddle's strike). In this case a possible product to use is a Straddle as presented in Table 1.15 for example.

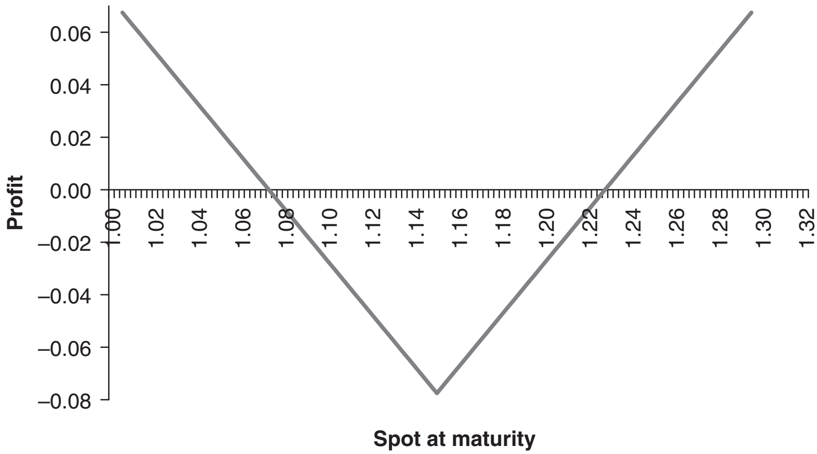

FIGURE 1.24 Profit and loss of a long Straddle.

TABLE 1.15 Example of a Straddle, valued at a volatility of about 7.8%.

Spot reference	1.1500 EUR-USD
Company buys	EUR call USD put
Company buys	EUR put USD call
Maturity	1 year
Notional of both the options	EUR 1,000,000
Strike of both options	1.1500 EUR-USD
Premium	USD 77,500.00

- If the spot rate is above the strike at maturity, the company can buy 1 M EUR at the strike of 1.1500.
- If the spot rate is below the strike at maturity, the company can sell 1 M EUR at the strike of 1.1500.

The break even points are 1.0726 for the put and 1.2274 for the call. If the spot is between the break even points at maturity, then the company will make an overall loss.

Applications and Assessment The straddle is not a product for corporate treasury. Buyers are typically investors. It is a vehicle to go long or short vega, i.e. the main application on the sell-side is volatility risk management of the at-the-money volatility or levels of the volatility surface. In the exercises in Section 1.5.12 you can verify that the delta-neutral straddle does in fact maximize the vega position.

1.6.4 Strangle

A strangle is a combination of an out-of-the-money put and call option with two different strikes. At inception, the strikes are typically 25-delta. It entitles the holder to buy an agreed amount of currency (say EUR) on a specified date (maturity) at a pre-determined rate (call strike), if the exchange rate is above the call strike at maturity. Alternatively, if the exchange rate is below the put strike at maturity, the holder is entitled to sell the amount at this strike. Buying a strangle provides full participation in a strongly moving market, where the direction is not clear. The buyer has to pay a premium for this product.

Advantages

- Full participation in or protection against a highly volatile exchange rate or increasing volatility
- Maximum loss is the premium paid
- Cheaper than the straddle

Disadvantages

- Expensive product
- Not suitable for directional hedging and consequently not suitable for hedge accounting as it should be clear if the client wants to sell or buy EUR

As in the straddle the chance of the strangle lies in spot movements. If the spot moves significantly, the call or the put can be sold before maturity with profit. Conversely, if a quiet market phase persists, the strategy is unlikely to generate much revenue. Figure 1.25 shows the profit and loss diagram of a long strangle.

For example, a company buys a strangle with a nominal of 1 M EUR. At maturity T:

1. If $S_T < K_1$, it would sell 1 M EUR at strike K_1.
2. If $K_1 < S_T < K_2$, it would not exercise either of the two options. The overall loss will be the strategy's premium.
3. If $S_T > K_2$, it would buy 1 M EUR at strike K_2.

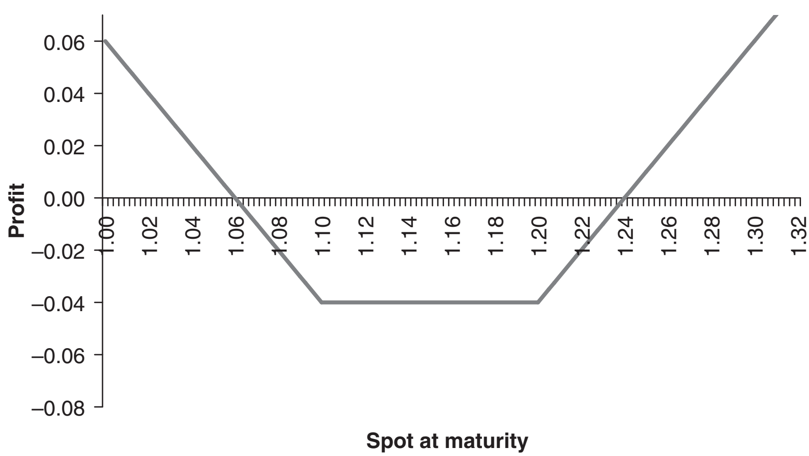

FIGURE 1.25 Profit and loss of a Strangle.

TABLE 1.16 Example of a Strangle, valued at a volatility of about 8%.

Spot reference	1.1500 EUR-USD
Company buys	EUR call USD put
Company buys	EUR put USD call
Maturity	1 year
Notional of both the options	EUR 1,000,000
Put strike	1.1000 EUR-USD
Call strike	1.2000 EUR-USD
Premium	USD 40,000.00

Example A company wants to benefit from its view that the EUR-USD exchange rate will move far from two specified strikes (strangle's strikes). In this case a possible product to use is a strangle as listed in Table 1.16, for example.

- If the spot rate is above the call strike at maturity, the company can buy 1 M EUR at the strike of 1.2000.
- If the spot rate is below the put strike at maturity, the company can sell 1 M EUR at the strike of 1.1000.

However, the risk is that if the spot rate is between the put strike and the call strike at maturity, both options expire worthless.

The break even points are 1.0600 for the put and 1.2400 for the call. If the spot is between these points at maturity, then the company makes an overall loss.

Applications and Assessment The strangle is not a product for corporate treasury. Buyers are typically investors. It is a vehicle to go long or short volga, i.e. the main application on the sell-side is volatility risk management of the convexity of the volatility surface. In the exercises in Section 1.6.9 you can verify that a delta-neutral straddle does not have much vanna, but a strangle does. This explains why a risk-warehousing sell-side needs strangles as a hedging instrument for second order risk on the volatility surface. In fact, the importance of the strangle can be observed by the brokers' market quoting a *one-vol-strangle* or a *market strangle* as a strategy rather than the building blocks. For details on the market strangle revisit Section 1.5.3.

1.6.5 Butterfly

A long butterfly is a combination of a long strangle and a short straddle. Buying a long butterfly provides participation where a highly volatile exchange rate condition exists. The buyer typically *receives* a premium for this product.

Advantages of a Short Butterfly

- Limited participation in or protection against market movement or increasing volatility
- Maximum loss is the premium paid
- Cheaper than the straddle

Disadvantages of a Short Butterfly

- Limited profit
- Not suitable for directional hedging and consequently not suitable for hedge accounting as it should be clear if the client wants to sell or buy EUR

If the spot will remain volatile, a short butterfly can be bought back before maturity with profit. Conversely, if a quiet market phase persists, a short butterfly strategy is unlikely to be bought back early. Figure 1.26 shows the profit and loss diagram of a long and a short butterfly.

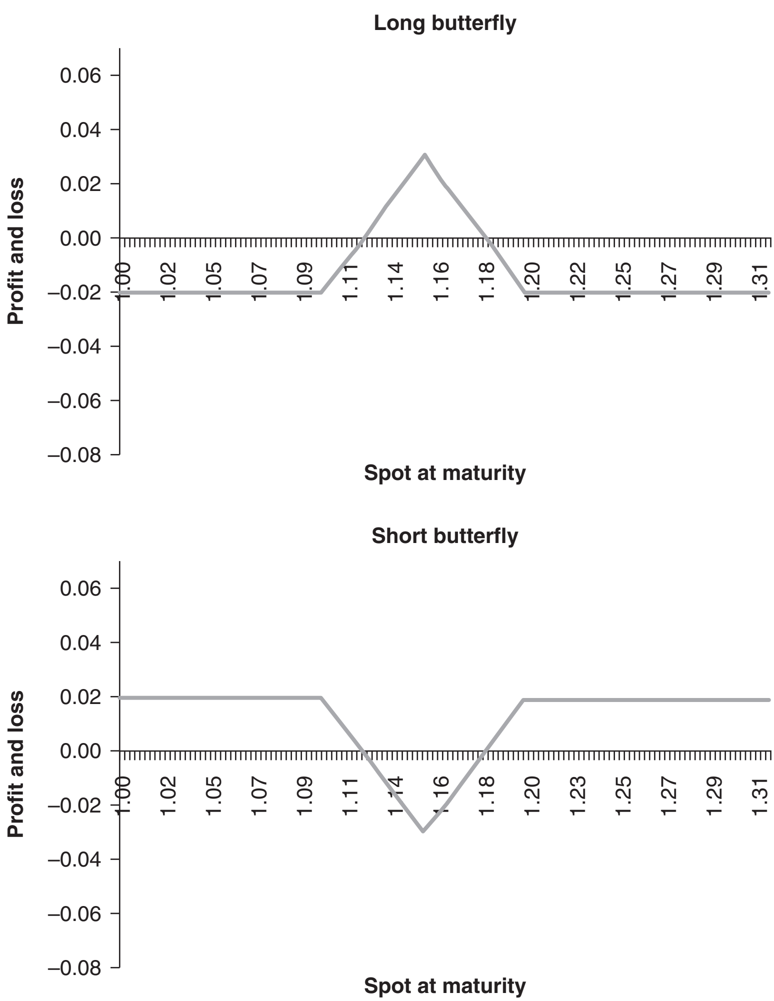

FIGURE 1.26 Profit and loss of a long and short butterfly.

For example, a company trades a short butterfly with a nominal of 1 M EUR and strikes $K_L < K_0 < K_U$. At maturity T:

1. If $S_T < K_L$, it would sell 1 M EUR at a rate $K_0 - K_L$ higher than the market.
2. If $K_L < S_T < K_0$, it would sell 1 M EUR at strike K_0.
3. If $K_0 < S_T < K_U$, it would buy 1 M EUR at strike K_0.
4. If $S_T > K_U$, it would buy 1 M EUR at a rate $K_U - K_0$ less than the market.

Example A company wants to benefit from its view that the EUR-USD exchange rate will remain volatile from a specified strike (the middle strike K_0).

In this case a possible product to use is a short butterfly as listed in Table 1.17 for example.

- If the spot rate is between the lower and the middle strike at maturity, the company can sell 1 M EUR at the strike of 1.1500.
- If the spot rate is between the middle and the higher strike at maturity, the company can buy 1 M EUR at the strike of 1.1500.
- If the spot rate is above the higher strike at maturity, the company will buy EUR 500 points below the spot.
- If the spot rate is below the lower strike at maturity, the company will sell EUR 500 points above the spot.

Applications and Assessment In principle, one can argue what is a long butterfly and what is a short butterfly. The version presented above is the common market standard, although it appears unusual to receive a premium when going long a strategy. The idea behind it is to go long convexity on the smile curve by taking a long position in a butterfly. Just like a strangle, a butterfly serves as a volga hedging strategy for the risk-warehousing sell-side. The butterfly is often preferred to the strangle because the premium is lower, delta is close to zero, vega is partially netted, and the volga of a butterfly and a strangle is identical, since at-the-money straddles do not have (much) volga. One can even go one step further and increase the notional of the strangle to generate a vega-neutral butterfly. This construction will then become a pure volga hedge without delta, gamma, or vega and is referred to as a *vega-weighted butterfly*.

TABLE 1.17 Example of a short butterfly.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional of all options	EUR 1,000,000
Lower strike K_L	1.1000 EUR-USD
Middle strike K_0	1.1500 EUR-USD
Upper strike K_U	1.2000 EUR-USD
Premium	USD 30,000.00

Alternative Construction of a Butterfly A long butterfly can also be constructed by going

- long a call struck at the lower strike,
- short two calls struck at the middle strike,
- long a call struck at the upper strike.

This construction shows that a butterfly is a market equivalent of the second derivative,

$$\text{call}(K_L) - 2\text{call}(K_0) + \text{call}(K_U), \tag{164}$$

which again underlines that a long butterfly reflects a position of long convexity. As a limiting case as difference of the strikes goes to zero, the butterfly can be viewed as an *Arrew-Debreu security*, or the probability density of the terminal spot price. Therefore, the long butterfly should have a non-negative value. If the long butterfly price is negative, it means that we have a negative probability density of the final spot distribution, which is referred to as *butterfly arbitrage*.

Note that the premium of a long butterfly constructed via call options is different from the premium of a long butterfly constructed via strangle and straddle. However, the shape of the payoff as well as the Greeks gamma, vega, vanna, and volga are the same for both constructions of the butterfly.

1.6.6 Condor

A long condor is a combination of a long strangle with far away strikes (e.g. 10-delta) and a short strangle with closer strikes (e.g. 25-delta). Buying a short condor provides participation in a large market movement up or down. A long condor provides participation in a low-volatility market. The buyer of a long condor typically *receives* a premium for this product.

Advantages of a Short Condor

- Limited participation in or protection against larger market movement or increasing volatility
- Maximum loss is the premium paid
- Cheaper than the short butterfly

Disadvantages of a Short Condor

- Limited profit
- Not suitable for directional hedging and consequently not suitable for hedge accounting as it should be clear if the client wants to sell or buy EUR

If the spot becomes more volatile, a short condor can be bought back before maturity with profit. Conversely, if a quiet market phase persists, a short condor strategy is unlikely to be bought back early. Figure 1.27 shows the profit and loss diagram of a

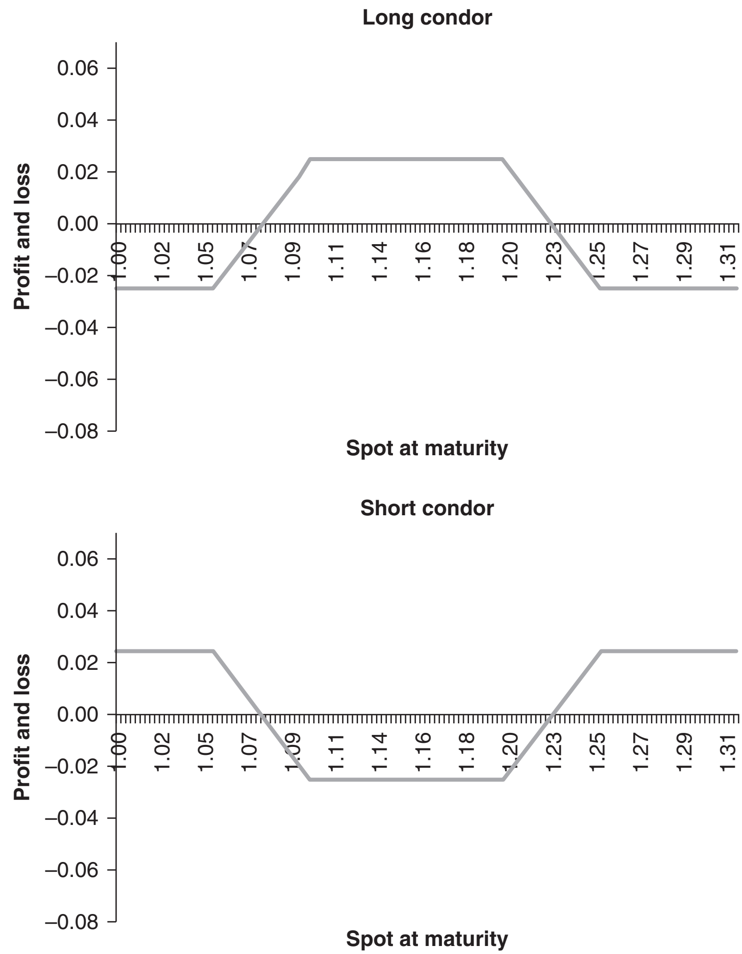

FIGURE 1.27 Profit and loss of a long and short condor, valued at a volatility of about 10%.

long and a short condor. The wide wing span of the south American condor bird is the origin of the name of the trading strategy.

For example, a company trades a short condor with a nominal of 1 M EUR and strikes $K_1 < K_2 < K_3 < K_4$. At maturity T:

1. If $S_T < K_1$, it would sell 1 M EUR at a rate $K_2 - K_1$ higher than the market.
2. If $K_1 < S_T < K_2$, it would sell 1 M EUR at strike K_2.
3. If $K_2 < S_T < K_3$, all options are out-of-the-money and no actions would follow from the strategy.
4. If $K_3 < S_T < K_4$, it would buy 1 M EUR at strike K_3.
5. If $S_T > K_4$, it would buy 1 M EUR at a rate $K_4 - K_3$ less than the market.

TABLE 1.18 Example of a short condor.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional of all options	EUR 1,000,000
Lower down strike K_1	1.0500 EUR-USD
Higher down strike K_2	1.1000 EUR-USD
Lower up strike K_3	1.2000 EUR-USD
Higher up strike K_4	1.2500 EUR-USD
Premium	USD 25,000.00

Example A company wants to benefit from its view that the EUR-USD exchange rate will become very volatile.

In this case a possible product to use is a short condor as listed in Table 1.18 for example.

- If the spot rate is between the lower and the higher down strike at maturity, the company can sell 1 M EUR at the strike of 1.1000.
- If the spot rate is between the lower and the higher up strike at maturity, the company can buy 1 M EUR at the strike of 1.2000.
- If the spot rate is above the higher up strike at maturity, the company will buy EUR 500 points below the spot.
- If the spot rate is below the lower down strike at maturity, the company will sell EUR 500 points above the spot.

Applications and Assessment As in the butterfly, one can argue what is a long condor and what is a short condor. The version presented above is the common market standard, although it appears unusual to receive a premium when going long a strategy. The idea behind it is to go long outer wing convexity on the smile curve by taking a long position in a butterfly. Obviously, this is a rather specialized market view.

Alternative Construction of a Condor A short condor can also be constructed by going

- long a call spread with lower and higher up strike,
- long a put spread with lower and higher down strike.

There is also an alternative construction of a long condor, where the buyer goes long a call at the lowest strike, short two calls at two different medium strikes, and long a call at the upper far strike, which can be remembered by going *long the wings short the belly*.

1.6.7 Seagull

A long seagull call strategy is a combination of a long call with a center strike, a short call with a higher strike, and a short put with a lower strike. It is similar to a risk reversal and typically trades at zero cost at inception. It entitles its holder to purchase an agreed amount of a currency (say EUR) on a specified date (maturity) at a pre-determined long call strike if the exchange rate at maturity is between the long call strike and the short call strike (see below for more information). If the exchange rate is below the short put strike at maturity, the holder must buy this amount in EUR at the short put strike. Buying a seagull call strategy provides good protection against a rising EUR.

Advantages

- Good protection against stronger EUR/weaker USD
- Better strikes than in a risk reversal
- Zero-cost product

Disadvantages

- Maximum loss depends on spot rate at maturity and can be arbitrarily large

As in a call spread, the protection against a rising EUR is limited to the interval from the long call strike and the short call strike. The biggest risk is a large upward movement of EUR.

Figure 1.28 shows the payoff and final exchange rate diagram of a seagull call. Rotating the payoff clockwise by about 45 degrees shows the shape of a flying seagull.

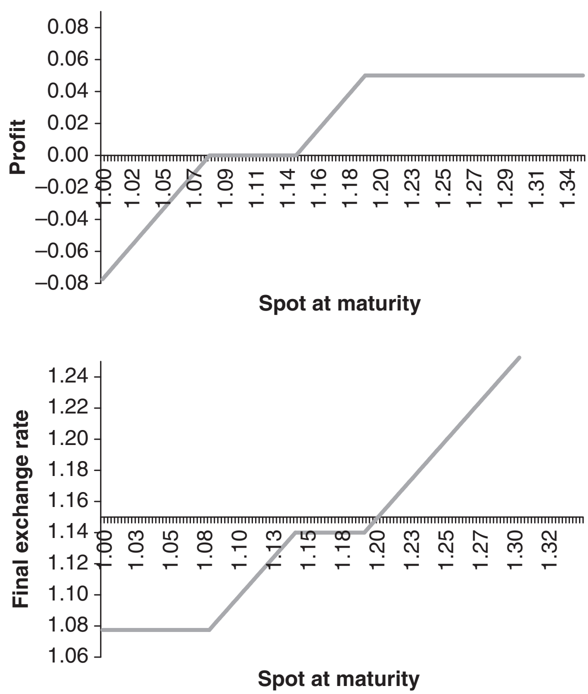

FIGURE 1.28 Profit (equivalent to the payoff for a zero-value structure) and final exchange rate of a seagull call.

Note that the diagram shows linear dependencies. This is correct only if the notional is in foreign currency (EUR). If the principal amount to be sold is in domestic currency (USD), then the relationships are no longer exactly linear.

For example, a company wants to sell 1 M USD and buy EUR. At maturity T:

1. If $S_T < K_1$, the company must sell 1 M USD at rate K_1.
2. If $K_1 < S_T < K_2$, all involved options expire worthless and the company can sell USD in the spot market.
3. If $K_2 < S_T < K_3$, the company would buy EUR at strike K_2.
4. If $S_T > K_3$, the company would sell 1 M USD at a rate $K_2 - K_1$ less than the market.

Example A company wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD but not a large upward movement of the EUR. The company wishes to be protected against a stronger EUR and finds that the corresponding plain vanilla is too expensive and would prefer a zero-cost strategy and is willing to limit protection on the upside.

In this case a possible form of protection that the company can use is a seagull call as presented in Table 1.19 for example.

- If the company's market expectation is correct, it can buy EUR at maturity at the strike of 1.1400.
- If the EUR-USD exchange rate will be above the short call strike of 1.1900 at maturity, the company will sell USD at 500 points less than the spot.
- If the EUR-USD exchange rate is below the strike of 1.0775 at maturity, it will have to sell 1 M USD at the strike of 1.0775.

Critical Assessment Which part of the seagull is risky depends on the client type. For a treasurer *with* the underlying cash flow, the risky part is the lack of protection against large spot moves up. The same problem as in the case of a call spread applies. For an investor *without* the underlying cash flow, the short put is the driver of risk. One can remove the risk for both if we alter the seagull structure: the EUR buyer USD seller could buy a EUR call and finance it by selling a EUR put spread. This construction would remove the risk on the downside for the investor because the potential loss is limited. It would also fix the lack of missing upside protection for the treasurer. However, this version of a seagull generates less attractive levels.

TABLE 1.19 Example of a seagull call.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional	USD 1,000,000
Company buys	EUR call USD put strike 1.1400
Company sells	EUR call USD put strike 1.1900
Company sells	EUR put USD call strike 1.0775
Premium	USD 0.00

1.6.8 Calendar Spread

A calendar spread is a strategy comprised of a long option with a longer maturity and a short option with a shorter maturity. Typically, both options are at-the-money. For at-the-money options the longer dated option is more expensive than the shorter dated option, whence the buyer has to pay a premium for the calendar spread. The trader of a calendar spread reflects his view of the slope of the volatility on the tenor space. Negative calendar spread prices indicate a *calendar arbitrage*. A calendar spread of options with different moneyness is called a *diagonal spread*.

More generally, calendar spreads can be differences of derivatives or structured products with two different maturity dates. For example, one can trade a rolling strategy of calendar spreads of variance swaps as a crash protection.

1.6.9 Exercises

Seagull For EUR/GBP spot ref 0.7000, volatility 8%, EUR rate 2.5%, GBP rate 4%, and flat smile, find the strike of the short EUR put for a 6-month zero-cost seagull put, where the strike of the long EUR put is 0.7150, the strike of the short call is 0.7300, and the desired sales margin is 0.1% of the GBP notional. What is the value of the seagull put after three months if the spot is at 0.6900 and the volatility is at 7.8%?

Straddle Vanna At-the-money straddles are commonly used as vega hedge instruments. Since vega of a vanilla option is concentrated around ATM, explain why vanna is not also concentrated around ATM. To clarify: vega attains its maximum if the strike is chosen as the one that makes a straddle delta-neutral. Vega is the first derivative of the option's value as volatility changes. Vanna is the second derivative of the option's value as spot and volatility change. As a mixed second derivative it has two interpretations:

Interpretation 1: Vanna = ∂ vega /∂ spot

Interpretation 2: Vanna = ∂ delta /∂ vol

1. To understand why vega peaks ATM and vanna OTM, plot vanilla vega on the spot space.
2. Calculate vanna for an at-the-money straddle as a formula in the Black-Scholes model.
3. Following interpretation 2, plot a vanilla (call) delta for a small volatility and a big volatility.

Straddle Volga Calculate volga for an at-the-money straddle.

Butterfly Premium Difference By how much does the premium of a long butterfly constructed via call options as in Equation (164) differ from the premium of a long butterfly constructed via strangle and straddle?

Short Gamma Long Vega Find a strategy (in the sense of a linear combination) of vanilla options that is short gamma and long vega (in the Black-Scholes model). Explain why this is not possible for single vanilla option. Hint: revisit Equation (19) for gamma and Equation (24) for vega of a vanilla option in the Black-Scholes model.

1.7 FIRST GENERATION EXOTICS

For the sake of example we consider EUR/USD – the most liquidly traded currency pair in the foreign exchange market. Internationally active market participants are always subject to changing foreign exchange rates. To hedge this exposure an immense variety of derivatives transactions is traded worldwide. Besides vanilla (European style put and call) options, the so-called first generation exotics have become standard derivative instruments.

1.7.1 Classification

The term *first generation exotic* does not refer to a clearly defined set of derivatives contracts, especially not in a legal sense. However, it is universally agreed that Foreign Exchange transactions (spot and forward contracts) and vanilla options are not in the set. It is also universally agreed that flip-flop-kiko-tarns and correlation swaps are not in the set either. We can then classify first generation exotics by:

Time of Introduction: Here we consider the history and the time when certain contracts first traded.

Existence of Standardized Deal Confirmations: We would classify a transaction as first generation exotic if there exists a standardized deal confirmation template, such as the ones provided by ISDA.

Replicability: We would classify a transaction as first generation exotic if it can be statically or semi-statically replicated or approximated by spot, forward, and vanilla option contracts.

Trading Volume: We would classify a transaction as first generation exotic if its trading volume is sufficiently high (and the transaction is not a spot, forward, or vanilla option).

There can also be other approaches to classify first generation exotics. I would like to point out that a first generation exotic does not necessarily need to be a currency option. For example, a flexi forward can be considered a first generation exotic in terms of both timing and standardization, but is clearly not an option. A variance swap can be considered a first generation exotic in terms of both standardization and replicability, but is clearly not an option because there is no right to exercise. Classification by trading volume would change the set of first generation exotics over time and is consequently not suitable for classification purposes. The various classifications would generate overlaps as well as differences. One could certainly argue to label barrier options as first generation exotic because they would satisfy all of the above: timing, standardization, replicability, and volume. For Asian options, the timing criterion would make them first generation as they started trading in Tokyo in 1987, but there was – even in 2016 – no standardized deal confirmation provided by ISDA. Power options satisfy timing and replicability, but not standardization or trading volume. This leads to the effect that the transition between the generations is not strict and can depend on the person you ask and classification they have in mind. A clean approach to classification could be sticking to the standardization, which would classify barrier options and touch products, as well as variance and volatility swaps as first generation exotic, based on the existing ISDA definitions and their supplements. The question of which transaction is standardized can then be viewed in light of ISDA's *Barrier Option Supplement* [78], which appeared in 2005. ISDA has extended the 1998 *FX and Currency Option Definitions* [77] to the range of touch products and single and double barrier options, including time windows for barriers. These are (a) options that knock in or out if the underlying hits a barrier (or one of two barriers) and (b) all kind of touch products: a one-touch [no-touch] pays a fixed amount of either USD or EUR if the spot ever [never] trades at or beyond the touch level and zero otherwise. Double one-touch and no-touch contracts work the same way but have two barriers. More on barrier options can be found in Section 1.7.3. The ISDA *Barrier Option Supplement* contains all the relevant definitions required to confirm these transactions by standardized short templates. It is clearly defined what a *barrier event* or a *determination agent* is. However, for purposes of classification, the product range covered by this ISDA supplement is not necessarily viewed as equivalent by all market participants. Moreover, the set of first generation exotics would then change each time ISDA publishes a new supplement. My personal preference is to classify the set of first generation exotics by the time of introduction in the market. This is reflected in this section.

1.7.2 European Digitals and the Windmill Effect

In this section we discuss the digital options along with the questions

1. How can we price digital options with smile?
2. Is the implied volatility of a digital option the same as the implied volatility of the corresponding vanilla option?
3. What is the windmill effect?

Digital Options (European) digital options pay off

$$v(T;S_T) = \mathbb{I}_{\{\phi S_T \geq \phi K\}} \text{ domestic paying,} \tag{165}$$

$$w(T;S_T) = S_T \mathbb{I}_{\{\phi S_T \geq \phi K\}} \text{ foreign paying.} \tag{166}$$

In the domestic paying case the payment of the fixed amount is in domestic currency, whereas in the foreign paying case the payment is in foreign currency. We obtain for the theoretical value functions

$$v(t;x) = e^{-r_d \tau} \mathcal{N}(\phi d_-), \tag{167}$$

$$w(t;x) = x e^{-r_f \tau} \mathcal{N}(\phi d_+), \tag{168}$$

of the digital options paying one unit of domestic and paying one unit of foreign currency respectively.

The question is how we can use the existing smile for vanilla options to read off a suitable volatility that we can plug into (167) to get a smile-adjusted value for the digital. In particular, can we take the same volatility as for the vanilla with strike K? The answer can be found by looking at the static replication and its associated cost.

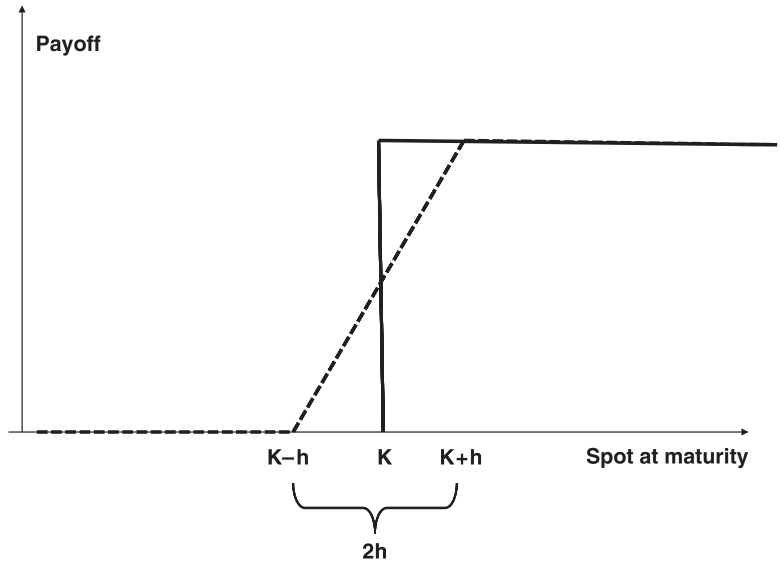

FIGURE 1.29 Replicating a digital call with a vanilla call spread.

Replication of Digital Options An obvious attempt to replicate a domestic digital is the call spread illustrated in Figure 1.29.

In the limit we have

$$\text{digital}(K) = \lim_{h\to 0} \frac{\text{vanilla}(K-h) - \text{vanilla}(K+h)}{2h} \tag{169}$$

However, since the number $\frac{1}{2h}$ corresponds to the (foreign!) notional of the vanilla options to trade, there are practical limitations we need to approximate the digital with a call spread with finite notional. To be on the safe side, the replication can be built as a super-replication with the upper strike chosen as the strike of the digital. This is too expensive to be used for pricing. For pricing we go for the symmetric compromise and choose one strike to be lower than the strike of the digital and the other one higher. The practical limitation for the difference of the two strikes is two pips. This is equivalent to a factor of 5000 to compute the notional of the vanilla options. In practice one would mostly take a larger difference or equivalently a smaller notional, say a factor of 50, for the notional multiplier. In this case the two strikes are two big figures apart. Consequently, we need to think very carefully about which volatilities to choose for the pricing. Taking the same volatility for the digital as for the corresponding vanilla would mean that we would price the options in the replicating portfolio with a flat volatility. Since the smile is not constant, this could produce a significant error. We should take the market volatilities for the replication to find a good price for the digital with smile. The mismatch is caused by the windmill effect and is illustrated in Figure 1.30.

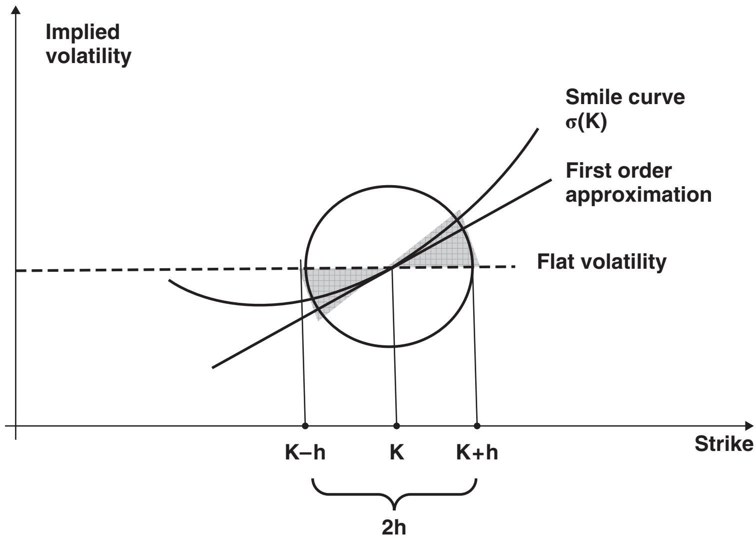

FIGURE 1.30 Windmill effect. The shaded gray areas show the mis-pricing. Working with a flat volatility is not sufficient. A first order approximation is attained using the windmill adjustment and will be considerably better. Working on the smile curve exactly is achieved by the call spread replication.

We conclude from Equation (169) that the value of the domestic-paying digital call is generally, and independent of any model, given by

$$\text{digital}(K) = -\frac{\partial}{\partial K}\text{vanilla}(K). \tag{170}$$

This helps us to get a much better approximation of the smile-adjusted value of the digital call. Let us abbreviate the vanilla call value function by $v(K, \sigma)$ and notice that, in a valuation with smile, the volatility σ is itself a function of K. The rest is a consequent application of calculus 101:

$$\begin{aligned}\text{digital}(K) &= -\frac{\partial}{\partial K}v(K, \sigma(K)) \\ &= -v_K(K, \sigma(K)) - v_\sigma(K, \sigma(K))\sigma'(K).\end{aligned} \tag{171}$$

The first term is the dual delta of the vanilla, i.e. $e^{-r_d\tau}\mathcal{N}(\phi d_-)$ as in (167). The second term is the windmill adjustment. It depends on the vanilla vega at the given strike and the slope of the smile on the strike space. The vega impact is maximal for the ATM (delta-neutral) strike and the contribution decreases from that ATM strike in both directions. Vanilla vega is always positive. The slope of the smile on the strike space means that we need to consider the implied volatility as a function of strike, i.e. strike is on the x-axis and implied volatility on the y-axis. This slope can be positive or negative. For a strongly down-skewed market, the slope is negative, whence the windmill adjustment is positive, and the digital call (put) will be valued higher (lower) when smile is taken into consideration. In short: down-skew leads to marking the digital calls up and the digital puts down; up-skew leads to marking the digital calls down and the digital puts up. This effect is highly important, as the market price of the digital translates to the market price of European knock-out, the one-touch, and hence the market price of the reverse knock-out and therefore into all structured products that contain these first generation exotics as building blocks. Note that the slope of the implied volatility curve is the local slope around the strike under consideration. Its sign is not necessarily equal to the sign of the 25-delta-risk reversal. We provide an example in Table 1.20. The windmill-adjustment is analytically tractable for a parametric smile interpolation – see Formula (162) in Section 1.5.12.

Volatility Implied by Digital Options With the windmill-adjustment it is obvious that digital options cannot be priced with the same volatility as the corresponding vanilla. Technically, it is possible to retrieve the volatility from a digital option's price. This is equivalent to retrieving the volatility from a given delta. It boils down to a quadratic equation with two solutions, similar to Equation (62). The volatility implied by the digital call price listed in Table 1.20 is based on this result. Note that the implied volatility of about 22% is very different from the smile volatility 15%, so in terms of volatility, the windmill effect is all the more visible.

Drift Sensitivity of Digital Options Looking at the payoff of digital options again shows that their value is highly drift sensitive. This means that we need to be very careful about which values for the rates and forward points to use for the valuation. For instance, selling a digital call would require buying foreign currency in the delta hedge. If we hedge delta with a forward contract (or even if we just pretend to do this), then we need to use the offer side of the forward points.

Using the same market data as in Table 1.20, we see that the forward points come out to be −597 mid market, implying a value of 30.89% DOM of a three-year digital call. The offer side of the forward points will be around 558 and the digital call value 31.22% DOM. On a notional of 1 M domestic currency, this difference accounts for 3,325 units of domestic currency. The trader selling the digital call would lose this amount if he sold the digital call priced with mid market forward points. Hence, always make sure you use the quote of the forward points consistent with your hedge: offer side for selling digital calls, bid side for selling digital puts.

TABLE 1.20 Windmill-adjustment for a digital call paying one unit of domestic currency. Contract data: Time to maturity $\tau = 186/365$, strike $K = 1.4500$, market data spot $S = 1.4000$, $r_d = 2.5\%$, $r_f = 4.0\%$, volatilities $\sigma(K) = 15.0000\%$, $\sigma(1.4499) = 15.0010\%$, $\sigma(1.4501) = 14.9990\%$, $h = 0.0001$.

Digital value without smile	$-v_K(K, \sigma(K)) = e^{-r_d\tau}\mathcal{N}(\phi d_-)$	0.322134
Value of the replication	$\frac{1}{2h}[v(K-h, \sigma(K-h)) - v(K+h, \sigma(K+h))]$	0.358975
Windmill-adjustment	$-v_\sigma(K, \sigma(K))\sigma'(K)$	0.036845
Digital value with smile	$-v_K(K, \sigma(K)) - v_\sigma(K, \sigma(K))\sigma'(K)$	0.358978
Implied volatility		22.005%

Domestic and Foreign Paying Digitals We leave it as an exercise for you to replicate a foreign-paying digital statically with vanillas and domestic-paying digitals. This perfect static replication will tell you all about the theoretical value and the smile-adjusted value of the foreign-paying digital.

Applications of Digital Options European digitals as a stand-alone transaction is mainly a vehicle of speculators. Trading desks and institutional investors could use digitals to hedge themselves against unexpected market moves, such as the EUR-CHF plunge on 15 January 2015. As a building block it comes up frequently in structured forwards (e.g. in Section 2.1.10) and structured FX-linked swaps (e.g. in Section 2.5.3). Conceptually, with the digital call one trades the probability distribution function directly.

European Knock-Out Barrier Options (EKO) A further application of digital options is the construction of European barrier options with payoff

$$\text{EKO} = [\phi(S_T - K)]^+ I\!I_{\{\eta S_T > \eta B\}}, \tag{172}$$

where K denotes the strike, ϕ the put-call indicator, $\eta = +1$ for a lower barrier B and $\eta = -1$ for an upper barrier B. Note that for European knock-outs, the barrier is always in-the-money, i.e. $B > K$ for the call and $B < K$ for the put. This payoff can be replicated statically with vanillas and digitals (exercise!), and therefore all we need to know about the market price follows from the vanilla smile and the windmill-adjustment. Since EKOs are building blocks of a dynamic replication of target forwards (see Section 2.2), most of what you need to know about target forwards goes back to understanding the EKO, which in turn goes back to the digital and the windmill effect. I would like to stress at this point that understanding the features of the volatility smile for vanilla options can help you understand the smile-adjusted price of many exotics and structured products.

1.7.3 Barrier Options

Barrier Options clearly belong to the first generation exotic options. The definitions of their terms have been standardized by the 2005 ISDA *Barrier Option Supplement* [78]. An entire book devoted to *FX Barrier Options* has been written by Dadachanji [34].

Knock-out Call Option (American Style Barrier) A knock-out call option entitles the holder to purchase an agreed amount of a currency (say EUR) on a specified expiration date at a pre-determined rate called the strike K provided the exchange rate never hits or crosses a pre-determined barrier level B. However, there is no obligation to do so. Buying a EUR knock-out call provides protection against a rising EUR if no knock-out event occurs between the trade date and expiration date while enabling full participation in a falling EUR. The holder has to pay a premium for this protection. The holder will typically exercise the option only if at expiration time the spot is above the strike and can exercise the option only if the spot has failed to touch the barrier between the trade date and expiration date (American style barrier) or if the spot on the expiry date does not touch or cross the barrier (European style barrier) – see Figure 1.31. We display the profit and the final exchange rate of an up-and-out call in Figure 1.32.

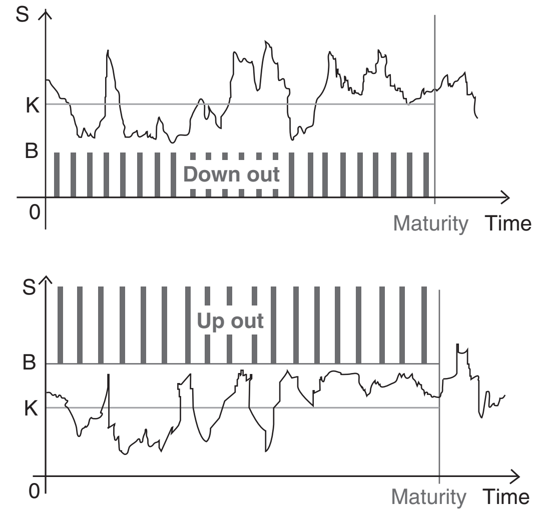

FIGURE 1.31 *Down-and-out* American barrier: if the exchange rate is never at or below B between the trade date and maturity, the option can be exercised. *Up-and-out* American barrier: if the exchange rate is never at or above B between the trade date and maturity, the option can be exercised.

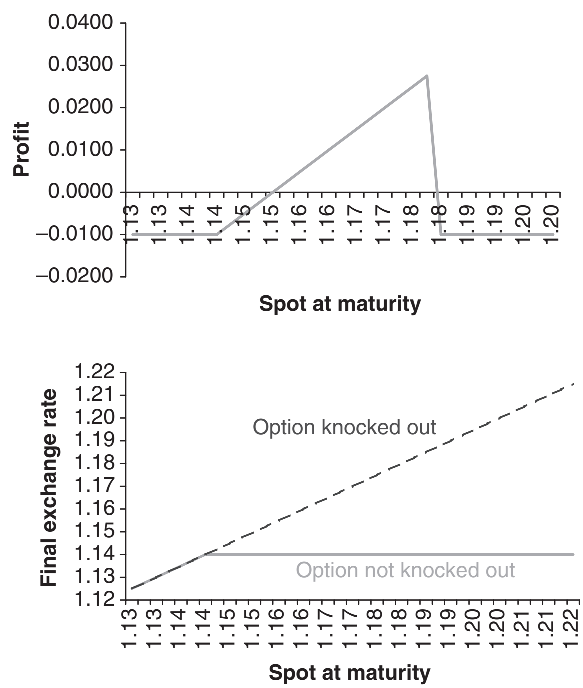

FIGURE 1.32 *Up-and-out* American barrier option payoff (above) and final exchange rate (below).

Advantages

- Cheaper than a plain vanilla
- Conditional protection against stronger EUR/weaker USD
- Full participation in a weaker EUR/stronger USD

Disadvantages

- Option may knock out, i.e. the hedge is lost if the option is used to hedge FX risk
- Premium has to be paid

For example, a company wants to sell 1 M USD and buy EUR at time T. If as usual S_t denotes the exchange rate of EUR-USD at time t then at maturity

1. if $S_T < K$, the company would not exercise the option,
2. if $S_T > K$ and if S has respected the conditions pre-determined by the barrier, the company would exercise the option and sell 1 M USD at strike K.

Example A company wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD. The company wishes to be able to buy EUR at a lower spot rate if the EUR becomes weaker on the one hand, but on the other hand be protected against a stronger EUR, and finds that the corresponding vanilla EUR call USD put is too expensive and is prepared to take more risk.

In this case a possible form of protection that the company can use is to buy a EUR knock-out call option as listed in Table 1.21 for example.

If the company's market expectation is correct, then it can buy EUR at maturity at the strike of 1.1500.

If the EUR-USD exchange rate touches the barrier at least once between the trade date and maturity the option will expire worthless.

Types of Barrier Options Generally the payoff of a standard knock-out option can be stated as

$$[\phi(S_T - K)]^+ I\!I_{\{\eta S_t > \eta B,\, 0 \le t \le T\}}, \tag{173}$$

where $\phi \in \{+1, -1\}$ is the usual put/call indicator and $\eta \in \{+1, -1\}$ takes the value +1 for a lower barrier (down-and-out) or −1 for an upper barrier (up-and-out). Compare

TABLE 1.21 Example of an up-and-out call.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional	EUR 1,000,000
Company buys	EUR call USD put
Strike	1.1500 EUR-USD
Up-and-out American barrier	1.3000 EUR-USD
Premium	EUR 12,553.00

this path-dependent payoff with the payoff (172) of the European barrier option. The corresponding *knock-in* options become alive only if the spot ever trades at or beyond the barrier between trade date and expiration date. Naturally,

$$\text{knock-out} + \text{knock-in} = \text{vanilla}, \tag{174}$$

which means that a portfolio containing both a knock-out and a knock-in option with otherwise identical contractual parameters can be considered economically equivalent to a vanilla option. Furthermore, we distinguish (see Figure 1.33):

(Regular) knock-out (KO): the barrier is out-of-the-money.

Reverse knock-out (RKO): the barrier is in-the-money. This version of a knock-out/knock-in is also referred to as *kick-out/kick-in*.

Strike out: the barrier is at the strike.

Losing a regular barrier option due to the spot hitting the barrier is not critical when used as a hedge for FX risk management. If the spot drops and a down-and-out EUR call USD put knocks out, then the hedge is lost, but at hitting time EUR is cheaper than at inception of the trade. The treasurer can trade another FX forward at a lower exchange rate at hitting time and lock in a better FX rate at no extra cost. Losing a reverse barrier option due to the spot hitting the barrier is more painful since the treasurer already has accumulated a positive intrinsic value, loses the hedge and is faced with a very expensive

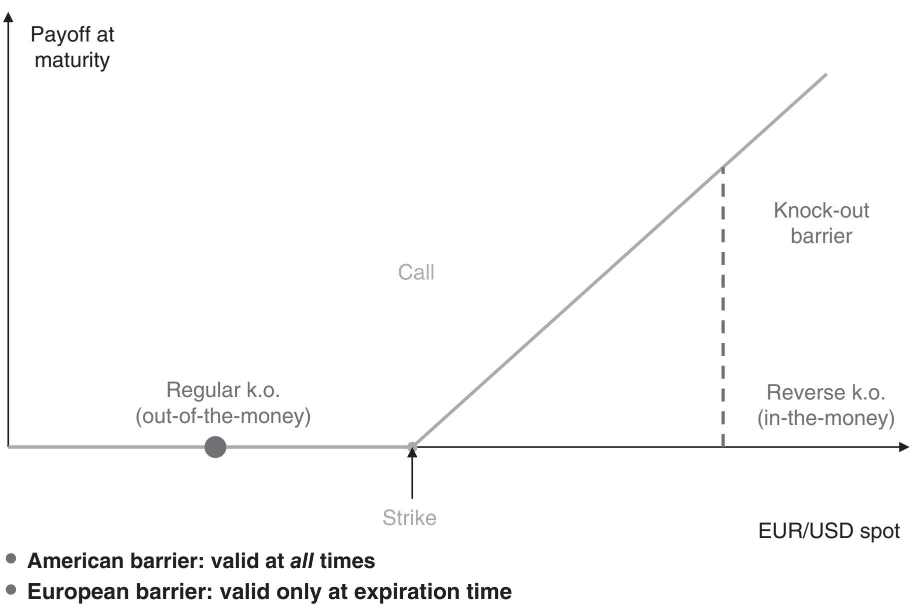

FIGURE 1.33 Barrier option terminology: regular barriers are out-of-the-money, reverse barriers are in-the-money.

EUR at hitting time. As a treasurer I would think twice before using a reverse knock-out as a stand-alone element of a currency hedge. In fact, reverse barrier options are used more commonly as part of structured FX forwards rather than as stand-alone elements.

This means that there are in total 16 different types of barrier options, call or put, in or out, up or down, regular or reverse (knock or kick).

Theoretical Value of Barrier Options For the standard type of FX barrier options a detailed derivation of values and Greeks can be found in [65].

Foreign–Domestic Symmetry for Barrier Options For a standard knock-out barrier option we let the value function be

$$v(S, r_d, r_f, \sigma, K, B, T, t, \phi, \eta), \tag{175}$$

where B denotes the barrier and the variable η takes the value +1 for a lower barrier and −1 for an upper barrier. With this notation at hand, we can state our FOR-DOM symmetry as

$$v(S, r_d, r_f, \sigma, K, B, T, t, \phi, \eta) = v(1/S, r_f, r_d, \sigma, 1/K, 1/B, T, t, -\phi, -\eta) \cdot S \cdot K. \tag{176}$$

Note that the rates r_d and r_f have been purposely interchanged. This implies that if we know how to price barrier contracts with upper barriers, we can derive the formulas for lower barriers. This symmetry is contractual and hence model-independent.

Barrier Option Terminology This paragraph is based on Hakala and Wystup [64].

American vs. European – Traditionally barrier options are of American style, which means that the barrier level is active during the entire duration of the option: any time between today and maturity the spot hits the barrier, the option becomes worthless. If the barrier level is active only at maturity, the barrier option is of European style and can in fact be replicated by a vertical spread and a digital option – see Section 1.7.2.

Single, double and outside barriers – Instead of taking just a lower or an upper barrier one could have both if one feels confident that the spot will remain in a range for a while. In this case besides vanillas, constant payoffs at maturity are popular – they are called range binaries. If the barrier and strike are in different exchange rates, the contract is called an outside barrier option or double currency barrier option. Such options traded a few years ago with the strike in USD/DEM and the barrier in USD/FRF taking advantage of the imbalance between implied and historic correlation between the two currency pairs.

Rebates – For knock-in options an amount R is paid at expiration by the seller of the option to the holder of the option if the option failed to kick in during its lifetime. For knock-out options an amount R is paid by the seller of the option to the holder of the option if the option knocks out. The payment of the rebate is either at maturity or at the first time the barrier is hit. Including such rebate features makes hedging easier for reverse barrier options and serves as a consolation for the holder's disappointment in case of a knock-out or a missing knock-in. The rebate part of a barrier option can be completely separated from the barrier contract and can in fact be traded separately, in which case it is called a one-touch (digital) or hit binary (in the knock-out case) and no-touch (in the knock-in case). We treat the touch contracts in detail in Section 1.7.2.

Determination of knock-out event – *Barrier events* are best described in ISDA's *Barrier Option Supplement* [78]. We highlight some of the relevant issues in Section 1.7.4.

How the Barrier is Monitored (Continuous vs. Discrete) and how this Influences the Value How often and when exactly do you check whether an option has knocked out or kicked in? This question is not trivial and should be clearly stated in the deal. The intensity of monitoring can create any price between a standard barrier and a vanilla contract. The standard for barrier options is continuous monitoring. Any time the exchange rate hits the barrier the option is knocked out. An alternative is to consider just daily/weekly/monthly currency fixings which makes the knock-out option more expensive because chances of knocking out are smaller (see Figure 1.34). A detailed discussion of the valuation of discrete barriers can be found in [57]. Discretely monitored barrier options are not very popular in FX markets, mainly because a trader is exposed to a fixing risk and you do know that all fixings are manipulated, don't you?

The Popularity of Barrier Options

- They are less expensive than vanilla contracts: in fact, the closer the spot is to the barrier, the cheaper the knock-out option. Any price between zero and the vanilla premium can be obtained by taking an appropriate barrier level, as we see

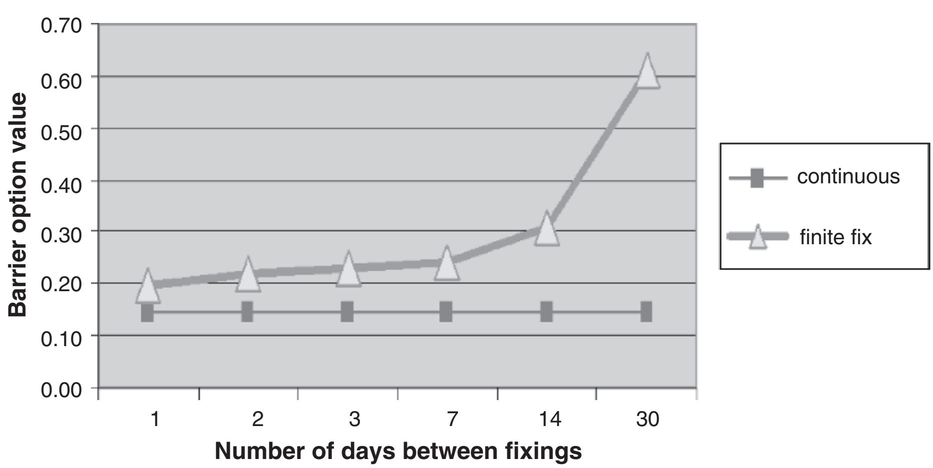

FIGURE 1.34 Comparison of a discretely and a continuously monitored knock-out barrier option.

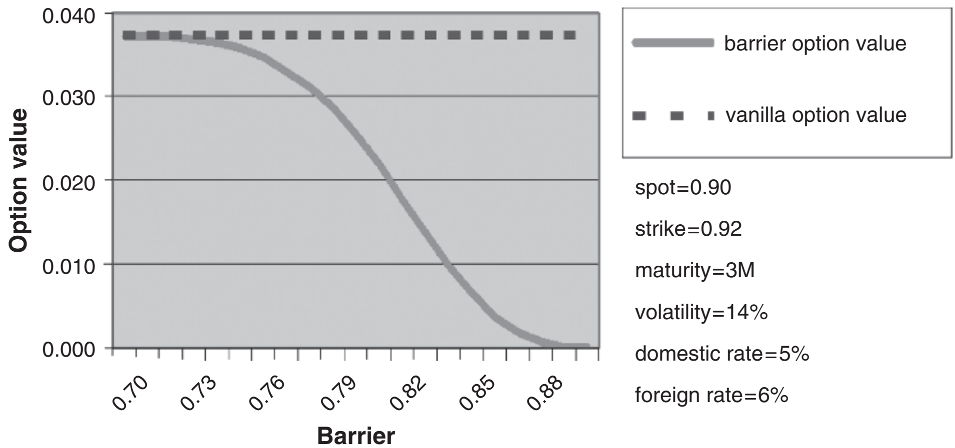

FIGURE 1.35 Comparison of a vanilla put and a down-and-out put. As the barrier moves far away from the current spot, the barrier option behaves like a vanilla option. As the barrier moves close to the current spot, the barrier option becomes worthless.

in Figure 1.35. One must be aware, however, that the lower the value of a barrier option, the more likely it will knock out.

- They allow foreign exchange risk exposure to be designed to customers' special needs. Instead of lowering the premium one can increase the nominal coverage of the vanilla contract by adding a barrier. Several customers feel sure about exchange rate levels not being hit during the next month which could be exploited to lower the premium. Others really want to cover their exchange rate exposure only if the market moves drastically which can be reflected by a knock-in option.
- The reduced cost allows another hedge of foreign exchange risk exposure if the first barrier option happens to knock out.
- The contract is easy to understand if one knows about vanillas.
- Many pricing and trading systems include barrier options in their standard.
- Pricing and hedging barriers in the Black-Scholes model are well understood and most pricing tools use closed-form solutions for the theoretical value, which allow fast and stable implementation; some even closed-form solutions in more advanced systems. Most of the big players have a stochastic-local volatility model in place to calculate values and Greeks of barrier options.
- Barrier options are standard building blocks in structured FX forwards – see for example the *shark forward* in Section 2.1.6.

Barrier Option Crisis 1994–1996 In the currency market barrier options became popular in 1994. The exchange rate between USD and DEM was then between 1.50 and 1.70. Since the all time low before 1995 was 1.3870 at September 2 1992 there were a lot of down and out barrier contracts written with a lower knock-out barrier of 1.3800. The sudden fall of the US dollar at the beginning of 1995 was unexpected and the 1.3800 barrier was hit at 10:30 am on March 29 1995 and fell even more to its all time low of

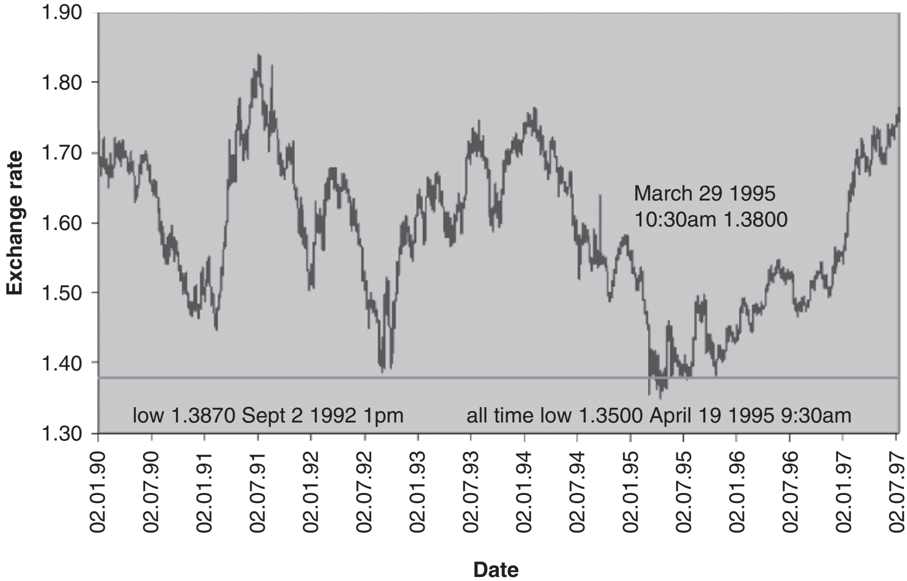

FIGURE 1.36 Barrier had lost popularity in 1994–1996 when USD-DEM had dropped below its historic low.

1.3500 at 9:30 am on April 19 1995. Numerous barrier option holders were shocked that a knock-out event was something that can really happen (see Figure 1.36). The shock lasted for more than a year and barrier options were unpopular for a while until many market participants had forgotten the event. Events like this often lead traders to question using exotics. Complicated products can in fact lead to unpleasant surprises. However, to cover foreign exchange risk reflecting individual market views at minimal cost requires exotic options. Often they appear as an integral part of an investment portfolio. The number of market participants who understand the advantages and pitfalls is growing steadily.

Risk Management of Barrier Options There are many ways to risk manage barrier options for a trading desk. Most typical is *hedging* the spot risk in the underlying and the vega and gamma risk using vanilla options. In order to understand the driving risk factors and the impact of the vanilla volatility surface on the price of barrier options, it is helpful to also consider *replication* of barrier options with vanilla options, either as a *static (buy-and-hold) replication* of the payoff or a *semi-static replication*, where the replicating portfolio must be unwound if and when the spot hits the barrier. We consider both hedging and replication. Since payoffs with discontinuities as we find them in touch contracts and reverse- and double-knock-out barrier options generate exploding Greeks, smoothing or face-lifting of payoffs is applied in practice before Greeks are used for hedging.

Hedging For barrier contracts a delta and vega hedge is most common. Obviously, both delta and vega hedges must be readjusted during the lifetime; however, this is normally conducted on a portfolio level. A vega hedge can be done using two vanilla options or more. In the example we consider a three-month up-and-out put with strike 1.0100 and barrier 0.9800. The vega minimizing hedge consists of 0.9 short three-month 50-delta calls and 0.8 long two-month 25-delta calls. Spot reference for EUR/USD is 0.9400 with rates 3.05% and 6.50% and volatility 11.9% – see Figure 1.37. This up-and-out put knocks out in-the-money, so is generally referred to as a reverse knock-out (RKO). Typically, the standard RKO put would be with a barrier below strike and the spot above the barrier. This is a less common example of a RKO put with a barrier above initial spot.

Replication of the Regular Knock-Out We start by semi-statically hedging regular barriers with a risk reversal as indicated in Figure 1.38. We choose the put strike to be zero for spot at the barrier and then recalculate the price of the risk reversal under current spot. The problem is, of course, that the value of the risk reversal at knock-out time does not need to be zero, and we do not know the market value of the risk reversal at hitting time. In fact, considering hedging a down-and-out call with a spot approaching the barrier and assuming a down trend in the spot triggers a down trend in the risk reversals, the calls will tend to be even cheaper than the

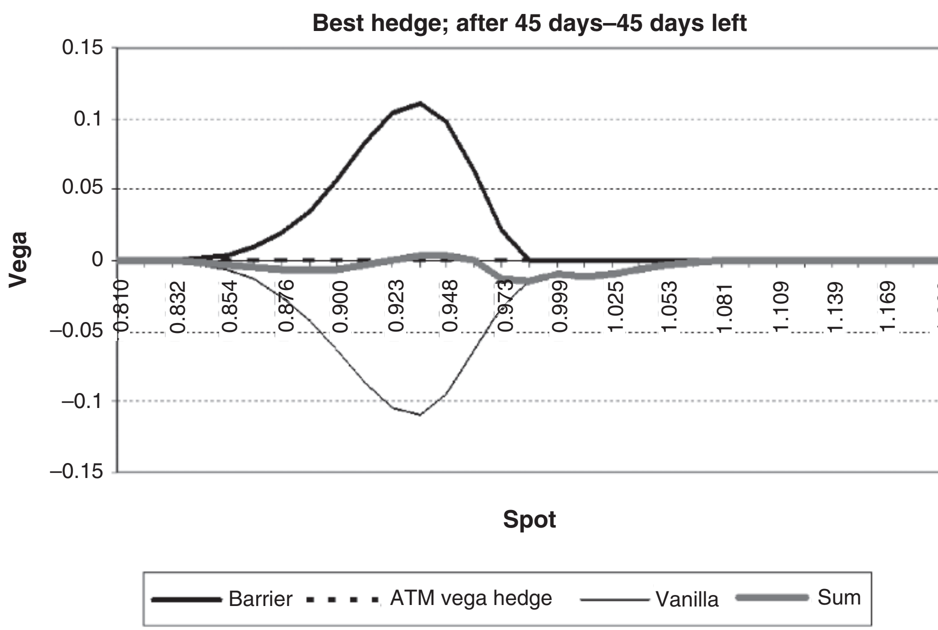

FIGURE 1.37 Vega depending on spot of an up-and-out put and a vega hedge consisting of two vanilla options.

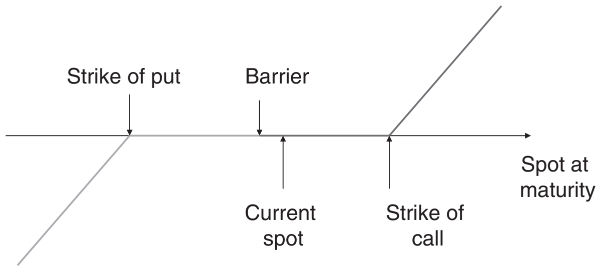

FIGURE 1.38 Semi-static replication of the regular knock-out with a risk reversal. A short down-and-out call is hedged by a long call with the same strike and a short put with a strike chosen in such a way that the value of the call and put portfolio is zero if the spot is at the barrier.

puts, so to unwind the hedge one would get less for the call one could sell and pay more than expected for the put to be bought back. Quantitatively, one can estimate the probability distribution of the first hitting time or its expected value under a pre-specified model. In the Black-Scholes model, it is in fact known in closed form. Furthermore, one can come up with an educated guess of the future risk reversal at hitting time, for example by running a regression analysis on historic relative returns versus relative risk reversals (relative meaning in relation to the at-the-money volatility). And if you believe in your model for the underlying and also believe in any sort of predictive power of historic financial data, then you can determine a put strike, that generates an unwind cost closer to zero. However, at the end of the day, the future risk reversal is unknown and the trading desk has to take a view on it. Future skew is the only risk that a risk reversal cannot replicate and this feature justifies the existence of the regular knock-out option. In fact, treasurers hardly ever use regular knock-out options to hedge their FX (directional) risk, but would much rather use a risk reversal. A view on the future skew at hitting time is more likely to be taken by a hedge fund. Therefore, many traders and researchers like to think of stochastic skew models taking exactly this effect into account. In summary, the regular knock-out can be seen essentially as a risk reversal, can be semi-statically replicated with a risk reversal. This replication is not used much in practice though: the treasurer would not replicate but replace the regular knock-out with the risk reversal, and the risk managing trading desk would hedge using the Greeks. The semi-static replication helps us understand the regular knock-out in full detail and can be used for valuation.

Semi-Static Replication of Reverse Knock-Out Options The perfect static replication of the reverse knock-out uses a portfolio of regular knock-out and single touch contracts (exercise). The regular knock-out in turn is essentially a risk reversal plus future skew risk, and the touch is approximately two European digitals. The European digital in turn can be viewed as a vanilla call spread with high notional amounts. Therefore, the risk drivers of the reverse knock-out are well understood. Several authors claim that reverse knock-out barrier options can be replicated semi-statically with a portfolio of vanilla options. These approaches are problematic if the hedging portfolio has to be unwound at hitting time, since volatilities for the vanillas may have changed between the time the replicating portfolio is composed and the first hitting time. Moreover, the occasionally high nominals and low deltas can cause a high price for the replicating portfolio. The approach by Maruhn and Sachs in [96] appears most promising. To understand the key problem and risk driver, note that a first order approximation of the reverse knock-out call is a ratio call spread with infinite leverage – see Section 1.6.1, and remember the disaster of the USD-TRY trader. This tells it all: high notionals that do not work in practice and the infinitely many at-the-money calls to buy back at first hitting time are driven by the future at-the-money volatility level. This exposure is very high because of the large notional of the short call.

Hedging the Reverse Knock-Out Reverse barrier options have extremely high values for delta, gamma, and theta when the spot is near the barrier and the time is close to expiry – see for example the delta in Figure 1.39. This is because the intrinsic value of the option jumps from a positive value to zero when the barrier is hit. In such a situation a simple delta hedge is impractical. However, there are ways to tackle this undesirable state of affairs by moving the barrier or more systematically applying valuation subject to portfolio constraints such as limited leverage – see Schmock *et al.* in [117]. The idea is to keep the *contractual barrier* for the contract and to determine the barrier event but

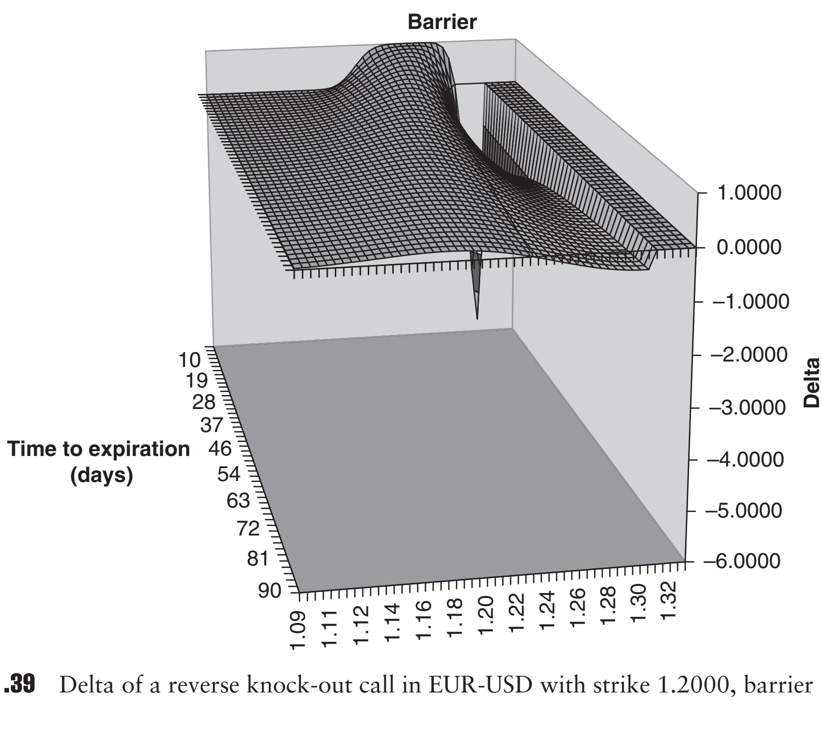

FIGURE 1.39 Delta of a reverse knock-out call in EUR-USD with strike 1.2000, barrier 1.3000.

use a *shadow barrier* to derive the value and the Greeks. A more advanced approach is called *barrier bending*, where the shadow barrier's distance from the contractual barrier is a function of time to maturity [126].

How Large Barrier Contracts Affect the Market Let's take the example of a reverse up-and-out call in EUR/USD with strike 1.2000 and barrier 1.3000. A trading desk delta-hedging a short position with nominal EUR 10 M has to buy 10 million times delta EUR, which is negative in this example. As the spot goes up to the barrier, delta reduces in size, requiring the hedging trading desk to sell more and more EUR. This can influence the market since steadily offering EUR slows down the spot movement towards the barrier and can in extreme cases prevent the spot from crossing the barrier. This is illustrated in Figure 1.40.

On the other hand, if the hedging trading desk runs out of breath or the upward market movement cannot be stopped by the delta-hedging institution, then the option knocks out and the hedge is unwound. Then suddenly more EUR are in demand whence the upward movement of an exchange rate can be accelerated once a large barrier contract in the market has knocked out. Situations like this happened to the USD-DEM spot in the early 1990s (see Figure 1.36), where many reverse knock-out puts have been written by banks, as traders are telling.

The reverse situation occurs when the bank hedges a long position, in which case EUR has to be bought when the spot approaches the barrier. This can cause an accelerated movement of the exchange rate towards the barrier and a sudden halt once the barrier is breached.

In modern compliance-dominated markets, regulators have actually started asking banks to ensure that by delta-hedging their own options positions they need to have a policy in place that prevents them from moving the market. While I know what to do

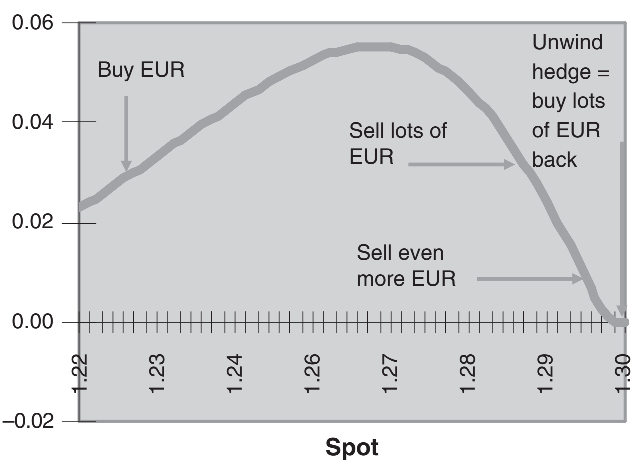

FIGURE 1.40 Delta hedging a short reverse knock-out call. The y-axis shows the USD value per EUR notional. Recall that the slope of the value function of a long contract represents the delta hedge to employ.

to reduce exploding Greeks, I fail to understand how one can ensure that markets are not affected by trading. Regulation should be regulated!

1.7.4 Touch Contracts

Now we take a more detailed look at the pricing of one-touch contracts – often also called (American style) binary or digital options, hit options, or rebate options. They trade as listed and over-the-counter products.

The touch-time is the first time the underlying trades at or beyond the touch level. The barrier determination agent, who is specified in the contract, determines the touch- or barrier event. The *Foreign Exchange Committee* recommends to the foreign exchange community a set of best practices for the barrier options and touch market. In the next stage of this project, the Committee is planning on publishing a revision of the *International Currency Options Master Agreement (ICOM) User Guide* to reflect the new recommendations.[3] Some key features are

- Determination whether the spot has breached the barrier must be due to actual transactions in the foreign exchange markets.
- Transactions must occur between 5:00 a.m. Sydney time on Monday and 5:00 p.m. New York time on Friday.
- Transactions must be of commercial size. In liquid markets, dealers generally accept that commercial size transactions are a minimum of 3 M USD.
- The barrier event determination agent may use cross-currency rates to determine whether a barrier has been breached in respect of a currency pair that is not commonly quoted.

The barrier or touch-level is usually monitored continuously over time. A further contractual issue to specify is the time of the payment of the one-touch. Typically, the notional is paid at the delivery date, which is two business days after the maturity. Another common practice is two business days after the barrier event (first hitting time, as we say in probability theory). In FX markets the former is the default.

Classification of One-Touch Contracts I was once asked whether the one-touch is an option. At first glance my reaction was that it is an option, because it is a derivative with a non-negative payoff. However, when going back to ISDA's 1998 **FX and Currency Option Definitions**, it appeared that there are two ways to view the one-touch:

Right to Receive: One can squeeze the one-touch into ISDA's 1998 FX and Currency Option Definitions and set the call currency amount equal to the one-touch notional, the put currency amount equal to zero, furthermore American style. Considering the option as a right to exercise would not work here because the holder would always exercise immediately and claim his call currency amount. Therefore, one must additionally agree on cash settlement and on a spot reference to define clearly in which case the option will be

[3] For details see http://www.ny.frb.org/fxc/fxann000217.html

automatically exercised. Once you think about the legal terms, it is no longer that obvious that the one-touch is actually a currency option.

Obligation to Pay: Indeed, an alternative way to consider the one-touch as a legal contract is to confirm it as an obligation of the seller to pay a fixed amount in case of a barrier event. Such deal confirmations have been used and clearly take the one-touch outside the class of currency options.

This doesn't really matter for the financial engineer, trader, or structurer, and if you find this subtle distinction irrelevant, then fair enough. However, it matters for clients who have agreed to trade only currency options and forwards. A payment obligation is clearly not in this set.

Applications of One-Touch Contracts Market participants of a rather speculative nature like to use one-touch contracts to reflect a view on a rising or falling exchange rate. In fact, nowadays, even individuals can trade them over the internet on a whole bunch of binary trading platforms. Hedging-focused clients often buy one-touch contracts as a rebate, so they receive a payment as a consolation if the strategy they believe in does not work. One-touch contracts also often serve as parts of structured products designed to enhance a forward rate or an interest rate.

Theoretical Value of the One-Touch In the standard Black-Scholes model for the underlying exchange rate of EUR/USD,

$$dS_t = S_t[(r_d - r_f)dt + \sigma dW_t], \tag{177}$$

where t denotes the running time in years, r_d the USD interest rate, r_f the EUR interest rate, σ the volatility, W_t a standard Brownian motion under the risk-neutral measure, the payoff is given by

$$R\mathbb{I}_{\{\tau_B \le T\}}, \tag{178}$$

$$\tau_B \stackrel{\Delta}{=} \inf\{t \ge 0 : \eta S_t \le \eta B\}. \tag{179}$$

This type of contract pays a domestic cash amount R USD if a barrier B is hit any time before the expiration time. We use the binary variable η to describe whether B is a lower barrier ($\eta = 1$) or an upper barrier ($\eta = -1$). The stopping time τ_B is called the first hitting time. The contract can be viewed either as the rebate portion of a knock-out barrier option or as an American cash-or-nothing digital. In FX markets it is usually called a *(single) one-touch (option)*, *one-touch-digital* or *hit binary*. The modified payoff of a *(single) no-touch (option)*, $R\mathbb{I}_{\{\tau_B \ge T\}}$, describes a rebate which is being paid if a knock-in option has not knocked in by the time it expires and can be valued similarly simply by exploiting the identity

$$R\mathbb{I}_{\{\tau_B \le T\}} + R\mathbb{I}_{\{\tau_B > T\}} = R. \tag{180}$$

We will further distinguish the cases

$$\omega = 0, \text{ rebate paid at hit,}$$
$$\omega = 1, \text{ rebate paid at end.}$$

It is important to mention that the payoff is one unit of the base currency. For a payment in the underlying currency EUR, one needs to exchange r_d and r_f, replace S and B by their reciprocal values, and change the sign of η.

For the one-touch we will use the abbreviations

- T: expiration time (in years)
- t: running time (in years)
- $\tau \stackrel{\Delta}{=} T - t$: time to expiration (in years)
- $\theta_{\pm} \stackrel{\Delta}{=} \frac{r_d - r_f}{\sigma} \pm \frac{\sigma}{2}$
- $S_t = S_0 e^{\sigma W_t + \sigma\theta_- t}$: price of the underlying at time t
- $n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2}$
- $\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt$
- $\vartheta_- \stackrel{\Delta}{=} \sqrt{\theta_-^2 + 2(1-\omega)r_d}$
- $e_{\pm} \stackrel{\Delta}{=} \frac{\pm \ln \frac{x}{B} - \sigma\vartheta_- \tau}{\sigma\sqrt{\tau}}$

We can describe the value function of the one-touch as a solution to a partial differential equation setup. Let $v(t,x)$ denote the value of the option at time t when the underlying is at x. Then $v(t,x)$ is the solution of

$$v_t + (r_d - r_f)xv_x + \frac{1}{2}\sigma^2 x^2 v_{xx} - r_d v = 0, t \in [0,T], \eta x \geq \eta B, \tag{181}$$

$$v(T,x) = 0, \eta x \geq \eta B, \tag{182}$$

$$v(t,B) = Re^{-\omega r_d \tau}, t \in [0,T]. \tag{183}$$

The theoretical value of the one-touch turns out to be

$$v(t,x) = Re^{-\omega r_d \tau} \left[\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} \mathcal{N}(-\eta e_+) + \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} \mathcal{N}(\eta e_-) \right] \tag{184}$$

Note that $\vartheta_- = |\theta_-|$ for rebates paid at end ($\omega = 1$).

Greeks We list some of the sensitivity parameters of the one-touch here, as they seem hard to find in the existing literature, but many people have asked me for them, so here we go.

Delta

$$v_x(t,x) = -\frac{Re^{-\omega r_d \tau}}{\sigma x}\left\{\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}}\left[(\theta_- + \vartheta_-)\mathcal{N}(-\eta e_+) + \frac{\eta}{\sqrt{\tau}} n(e_+)\right]\right.$$
$$\left. + \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}}\left[(\theta_- - \vartheta_-)\mathcal{N}(\eta e_-) + \frac{\eta}{\sqrt{\tau}} n(e_-)\right]\right\} \tag{185}$$

Gamma can be obtained using $v_{xx} = \frac{2}{\sigma^2 x^2}[r_d v - v_t - (r_d - r_f) x v_x]$ and turns out to be

$$v_{xx}(t,x) = \frac{2Re^{-\omega r_d \tau}}{\sigma^2 x^2}\cdot \tag{186}$$
$$\left\{\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}}\mathcal{N}(-\eta e_+)\left[r_d(1-\omega) + (r_d - r_f)\frac{\theta_- + \vartheta_-}{\sigma}\right]\right.$$
$$+\left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}}\mathcal{N}(\eta e_-)\left[r_d(1-\omega) + (r_d - r_f)\frac{\theta_- - \vartheta_-}{\sigma}\right]$$
$$+\eta\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} n(e_+)\left[-\frac{e_-}{2\tau} + \frac{r_d - r_f}{\sigma\sqrt{\tau}}\right]$$
$$\left. +\eta\left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} n(e_-)\left[\frac{e_+}{2\tau} + \frac{r_d - r_f}{\sigma\sqrt{\tau}}\right]\right\}.$$

Theta

$$v_t(t,x) = \omega r_d v(t,x) + \frac{\eta Re^{-\omega r_d \tau}}{2\tau}\left[\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} n(e_+) e_- - \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} n(e_-) e_+\right]$$
$$= \omega r_d v(t,x) + \frac{\eta Re^{-\omega r_d \tau}}{\sigma \tau^{(3/2)}}\left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} n(e_+)\ln\left(\frac{B}{x}\right). \tag{187}$$

The computation exploits the identities (209), (210), and (211) derived below.

Vega requires the identities

$$\frac{\partial \theta_-}{\partial \sigma} = -\frac{\theta_+}{\sigma} \tag{188}$$

$$\frac{\partial \vartheta_-}{\partial \sigma} = -\frac{\theta_- \theta_+}{\sigma \vartheta_-} \tag{189}$$

$$\frac{\partial e_\pm}{\partial \sigma} = \pm\frac{\ln\frac{B}{x}}{\sigma^2\sqrt{\tau}} + \frac{\theta_- \theta_+}{\sigma \vartheta_-}\sqrt{\tau} \tag{190}$$

$$A_{\pm} \stackrel{\Delta}{=} \frac{\partial}{\partial\sigma}\frac{\theta_- \pm \vartheta_-}{\sigma} = -\frac{1}{\sigma^2}\left[\theta_+ + \theta_- \pm \left(\frac{\theta_-\theta_+}{\vartheta_-} + \vartheta_-\right)\right] \tag{191}$$

and turns out to be

$$\begin{aligned} v_\sigma(t,x) = {} & Re^{-\omega r_d \tau} \cdot \\ & \left\{ \left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} \left[\mathcal{N}(-\eta e_+) A_+ \ln\left(\frac{B}{x}\right) - \eta n(e_+)\frac{\partial e_+}{\partial\sigma}\right] \right. \\ & \left. + \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} \left[\mathcal{N}(\eta e_-) A_- \ln\left(\frac{B}{x}\right) + \eta n(e_-)\frac{\partial e_-}{\partial\sigma}\right] \right\}. \end{aligned} \tag{192}$$

Vanna uses the identity

$$d_- = \frac{\ln\frac{B}{x} - \sigma\theta_-\tau}{\sigma\sqrt{\tau}} \tag{193}$$

and turns out to be

$$\begin{aligned} v_{\sigma x}(t,x) = {} & \frac{Re^{-\omega r_d \tau}}{\sigma x} \cdot \\ & \left\{ \left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} \left[\mathcal{N}(-\eta e_+) A_+ \left(-\sigma - (\theta_- + \vartheta_-)\ln\left(\frac{B}{x}\right)\right) \right.\right. \\ & \left. - \frac{\eta n(e_+)}{\sqrt{\tau}}\left(d_-\frac{\partial e_+}{\partial\sigma} + A_+ \ln\left(\frac{B}{x}\right) - \frac{1}{\sigma}\right)\right] \\ & + \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} \left[\mathcal{N}(\eta e_-) A_- \left(-\sigma - (\theta_- - \vartheta_-)\ln\left(\frac{B}{x}\right)\right) \right. \\ & \left.\left. + \frac{\eta n(e_-)}{\sqrt{\tau}}\left(d_-\frac{\partial e_-}{\partial\sigma} - A_- \ln\left(\frac{B}{x}\right) + \frac{1}{\sigma}\right)\right]\right\} \end{aligned} \tag{194}$$

Volga uses the identities

$$g = \frac{1}{\sigma^2\vartheta_-}\left[-\theta_+^2 - \theta_-^2 - \theta_+\theta_- + \frac{\theta_-^2\theta_+^2}{\vartheta_-^2}\right] \tag{195}$$

$$\frac{\partial^2 e_\pm}{\partial\sigma^2} = \mp\frac{2\ln\left(\frac{B}{x}\right)}{\sigma^3\sqrt{\tau}} + g\sqrt{\tau} \tag{196}$$

$$\frac{\partial A_\pm}{\partial\sigma} = \frac{\theta_+ + \theta_-}{\sigma^3} - \frac{2A_\pm \pm g}{\sigma} \tag{197}$$

and turns out to be

$$
\begin{aligned}
v_{\sigma\sigma}(t,x) = Re^{-\omega r_d \tau}\cdot \quad (198)\\
\Bigg\{ \left(\frac{B}{x}\right)^{\frac{\theta_- + \vartheta_-}{\sigma}} & \Bigg[\mathcal{N}(-\eta e_+) \ln\left(\frac{B}{x}\right) \left(A_+^2 \ln\left(\frac{B}{x}\right) + \frac{\partial A_+}{\partial \sigma} \right) \\
& - \eta n(e_+) \left(2\ln\left(\frac{B}{x}\right) A_+ \frac{\partial e_+}{\partial \sigma} - e_+ \left(\frac{\partial e_+}{\partial \sigma}\right)^2 + \frac{\partial^2 e_+}{\partial \sigma^2} \right) \Bigg] \\
+ \left(\frac{B}{x}\right)^{\frac{\theta_- - \vartheta_-}{\sigma}} & \Bigg[\mathcal{N}(\eta e_-) \ln\left(\frac{B}{x}\right) \left(A_-^2 \ln\left(\frac{B}{x}\right) + \frac{\partial A_-}{\partial \sigma} \right) \\
& + \eta n(e_-) \left(2\ln\left(\frac{B}{x}\right) A_- \frac{\partial e_-}{\partial \sigma} - e_- \left(\frac{\partial e_-}{\partial \sigma}\right)^2 + \frac{\partial^2 e_-}{\partial \sigma^2} \right) \Bigg] \Bigg\}
\end{aligned}
$$

Touch Probability The risk-neutral probability of knocking out is given by

$$
\begin{aligned}
\mathbb{P}[\tau_B \leq T] &= \mathbb{E}\left[\mathbb{I}_{\{\tau_B \leq T\}}\right] \\
&= \frac{1}{R} e^{r_d T} v(0, S_0). \quad (199)
\end{aligned}
$$

The touch probability is a non-discounted one-touch unit price for a one-touch paying domestic currency. This relationship is generic, i.e. independent of the model. However, it is not independent of the measure. The market price of a one-touch is derived under a risk-neutral measure. This must be taken into consideration when deriving touch probabilities from market prices of the one-touch.

Properties of the First Hitting Time τ_B As derived in [119], for example, the first hitting time

$$\tilde{\tau} \stackrel{\Delta}{=} \inf\{t \geq 0 : \theta t + W(t) = x\} \quad (200)$$

of a Brownian motion with drift θ and hit level $x > 0$ has the density

$$\mathbb{P}[\tilde{\tau} \in dt] = \frac{x}{t\sqrt{2\pi t}} \exp\left\{ -\frac{(x-\theta t)^2}{2t} \right\} dt, \quad t > 0, \quad (201)$$

the cumulative distribution function

$$\mathbb{P}[\tilde{\tau} \leq t] = \mathcal{N}\left(\frac{\theta t - x}{\sqrt{t}}\right) + e^{2\theta x}\mathcal{N}\left(\frac{-\theta t - x}{\sqrt{t}}\right), \quad t > 0, \quad (202)$$

the Laplace-transform

$$\mathbb{E}e^{-\alpha\tilde{\tau}} = \exp\left\{ x\theta - x\sqrt{2\alpha + \theta^2} \right\}, \quad \alpha > 0,\ x > 0, \quad (203)$$

and the property

$$\mathbb{P}[\tilde{\tau} < \infty] = \begin{Bmatrix} 1 & \text{if } \theta \geq 0 \\ e^{2\theta x} & \text{if } \theta < 0 \end{Bmatrix} \tag{204}$$

For upper barriers $B > S_0$ we can now rewrite the first passage time τ_B as

$$\begin{aligned}\tau_B &= \inf\{t \geq 0 : S_t = B\} \\ &= \inf\left\{t \geq 0 : W_t + \theta\_t = \frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)\right\}.\end{aligned} \tag{205}$$

The density of τ_B is hence

$$\mathbb{P}[\tilde{\tau_B} \in dt] = \frac{\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)}{t\sqrt{2\pi t}} \exp\left\{-\frac{\left(\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right) - \theta\_t\right)^2}{2t}\right\}, \quad t > 0. \tag{206}$$

Derivation of the Value Function Using the density (206) the value of the paid-at-end ($\omega = 1$) upper rebate ($\eta = -1$) option can be written as the following integral:

$$\begin{aligned}v(T, S_0) &= Re^{-r_d T}\mathbb{E}\left[\mathbb{I}_{\{\tau_B \leq T\}}\right] \\ &= Re^{-r_d T}\int_0^T \frac{\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)}{t\sqrt{2\pi t}} \exp\left\{-\frac{\left(\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right) - \theta\_t\right)^2}{2t}\right\} dt.\end{aligned} \tag{207}$$

To evaluate this integral, we introduce the notation

$$e_{\pm}(t) \stackrel{\Delta}{=} \frac{\pm\ln\frac{S_0}{B} - \sigma\theta\_t}{\sigma\sqrt{t}} \tag{208}$$

and list the properties

$$e\_(t) - e_+(t) = \frac{2}{\sqrt{t}}\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right), \tag{209}$$

$$n(e_+(t)) = \left(\frac{B}{S_0}\right)^{-\frac{2\theta\_}{\sigma}} n(e\_(t)), \tag{210}$$

$$\frac{\partial e_{\pm}(t)}{\partial t} = \frac{e_{\mp}(t)}{2t}. \tag{211}$$

We evaluate the integral in (207) by rewriting the integrand in such a way that the coefficients of the exponentials are the inner derivatives of the exponentials using properties (209), (210), and (211).

$$\begin{aligned}
&\int_0^T \frac{\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)}{t\sqrt{2\pi t}} \exp\left\{-\frac{\left(\frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)-\theta_- t\right)^2}{2t}\right\} dt \\
&= \frac{1}{\sigma}\ln\left(\frac{B}{S_0}\right)\int_0^T \frac{1}{t^{(3/2)}} n(e_-(t))\, dt \\
&= \int_0^T \frac{1}{2t} n(e_-(t))[e_-(t) - e_+(t)]\, dt \\
&= -\int_0^T n(e_-(t))\frac{e_+(t)}{2t} + \left(\frac{B}{S_0}\right)^{\frac{2\theta_-}{\sigma}} n(e_+(t))\frac{e_-(t)}{2t}\, dt \\
&= \left(\frac{B}{S_0}\right)^{\frac{2\theta_-}{\sigma}} \mathcal{N}(e_+(T)) + \mathcal{N}(-e_-(T)).
\end{aligned} \tag{212}$$

The computation for lower barriers ($\eta = 1$) is similar.

Semi-Static Replication and Smile Effect The one-touch can be valued in many models, including stochastic-local volatility models and jump diffusion models. The purpose of this is to include the smile effect and the dynamics of the underlying consistently into the valuation of the one-touch. In order to see the driving risk and really understand it, we may again turn to a semi-static replication. We start by analyzing the touch probability.

Touch Probability and Moneyness Probability The touch probability for a touch level B above initial spot S_0 can be rewritten as

$$\text{touch-probability} = \mathbb{P}[M_T \geq B], \tag{213}$$

$$M_T \stackrel{\Delta}{=} \max_{0\leq t\leq T} S_t. \tag{214}$$

The event $\{M_T \geq B\}$ can be split into a case where the final spot S_T is also at or above B and another case where it is below B. We continue as

$$\begin{aligned}
\mathbb{P}[M_T \geq B] &= \mathbb{P}[M_T \geq B, S_T \geq B] + \mathbb{P}[M_T \geq B, S_T < B] \\
&= \mathbb{P}[S_T \geq B] + \mathbb{P}[M_T \geq B, S_T < B] \\
&= \mathbb{P}[S_T \geq B] + \mathbb{P}[S_T \geq B].
\end{aligned} \tag{215}$$

In the first summand we have dropped the event $\{M_T \geq B\}$ because it is implied by the event $\{S_T \geq B\}$. The second summand collapses to $\mathbb{P}[S_T \geq B]$ because of the *reflection principle*. It states that if a spot path starts at a level below B and ends at a level above B, then the first hitting time must occur before T and from the first hitting time, there is a spot path reflected at the barrier level that ends above B, and now comes the crucial part: and that the reflected path is equally likely. Now the reflection principle as we know it from probability theory can be proved under the assumption of a *symmetric* Brownian motion. As we introduce an up or down drift, this will clearly fail. For practical matters, we may work with an assumption of a symmetric FX spot model and view the result of the reflection principle as a very handy rule of thumb:

$$\text{one-touch} \approx 2 \cdot \text{digitals}. \tag{216}$$

The closer the spot price to a perfect symmetry, the better the approximation. So for a market with high swap points as in emerging markets, this approximation is not very good, but still provides a first orientation. I always check the European digital price when I see or need to quote a one-touch price. As for the intuition: if the spot is on the barrier at inception, then the one-touch price is 100%, and the digital price is around 50%, because there is a 50/50 chance for the spot to end up at or above the barrier at maturity. In this extreme case, we see how important the assumption about symmetry is.

Semi-Static Replication To semi-statically replicate the one-touch we follow the recipe and buy two digitals, obviously with the same barrier, same maturity, same pay currency, and same notional. If the spot doesn't hit the barrier, everything is zero, so our replication is perfect. If the spot hits the barrier, we sell the two digitals at first hitting time. Both European digitals are then at-the-money spot, and their value is consequently approximately 50%. Since we have two, the total value adds up to 100%, and Bob's your uncle. Again, we see that the symmetry is crucial. However, we can get a very good smile-adjusted value of the European digital using the windmill effect. This shows us that the slope of the smile on the strike space at the ATM point drives the risk of the one-touch. It is the term structure of this future slope that matters, and which is unknown at inception. Now, would you actually do this replication?

Quotation Conventions and Bid-Ask Spreads If the payoff is at maturity, the undiscounted value of the one-touch is the touch probability under the risk-neutral measure. The market standard is to quote the price of a one-touch in percent of the payoff, a number between 0 and 100%. The market value of a one-touch depends on the theoretical value (TV) of the above formula, the smile adjustment (either derived by our semi-static replication thought experiment or from quantifying other hedge cost such as vanna and volga explained in Section 4.1), and the bid-ask spread. The spread in turn depends on the currency pair and the client. For interbank trading, spreads are usually between 2% and 4% for liquid currency pair – see Section 4.2 for details. Bid and offer quotes are usually rounded to the next quarter.

Two-Touch A two-touch pays one unit of currency (either foreign or domestic) if the underlying exchange rate hits both an upper and a lower barrier during its lifetime. This can be structured using basic touch contracts in the following way. The long two-touch with barriers L and H is equivalent to

1. a long single one-touch with lower barrier L,
2. a long single one-touch with upper barrier H,
3. a short double one-touch with barriers L and H.

This is easily verified by looking at the possible cases.

If the order of touching L and H matters, then the above hedge no longer works, but we have a new product, which can be valued, for example, using a finite-difference grid or Monte Carlo Simulation.

Double-No-Touch The payoff

$$I\!I_{\{L\leq \min_{[0,T]} S_t < \max_{[0,T]} S_t \leq H\}} \tag{217}$$

of a double-no-touch (DNT) is in units of domestic currency and is paid at maturity T. The lower barrier is denoted by L, the higher barrier by H.

Derivation of the Value Function To compute the expectation, let us introduce the stopping time

$$\tau \stackrel{\Delta}{=} \min\{\inf\{t \in [0,T] | S_t = L \text{ or } S_t = H\}, T\} \tag{218}$$

and the notation

$$\tilde{\theta}_\pm \stackrel{\Delta}{=} \frac{r_d - r_f \pm \frac{1}{2}\sigma^2}{\sigma} \tag{219}$$

$$\tilde{h} \stackrel{\Delta}{=} \frac{1}{\sigma} \ln \frac{H}{S_t} \tag{220}$$

$$\tilde{l} \stackrel{\Delta}{=} \frac{1}{\sigma} \ln \frac{L}{S_t} \tag{221}$$

$$\theta_\pm \stackrel{\Delta}{=} \tilde{\theta}_\pm \sqrt{T-t} \tag{222}$$

$$h \stackrel{\Delta}{=} \tilde{h}/\sqrt{T-t} \tag{223}$$

$$l \stackrel{\Delta}{=} \tilde{l}/\sqrt{T-t} \tag{224}$$

$$y_\pm \stackrel{\Delta}{=} y_\pm(j) = 2j(h-l) - \theta_\pm \tag{225}$$

$$n_T(x) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi T}} \exp\left(-\frac{x^2}{2T}\right) \tag{226}$$

$$n(x) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} \tag{227}$$

$$\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt. \tag{228}$$

On $[t, \tau]$, the value of the double-no-touch is

$$v(t) = I\!E^t \left[e^{-r_d(T-t)} I\!I_{\{L \le \min_{[0,T]} S_t < \max_{[0,T]} S_t \le H\}} \right], \tag{229}$$

on $[\tau, T]$,

$$v(t) = e^{-r_d(T-t)} I\!I_{\{L \le \min_{[0,T]} S_t < \max_{[0,T]} S_t \le H\}}. \tag{230}$$

The joint distribution of the maximum and the minimum of a Brownian motion can be taken from [113] and is given by

$$I\!P \left[\tilde{l} \le \min_{[0,T]} W_t < \max_{[0,T]} W_t \le \tilde{h} \right] = \int_{\tilde{l}}^{\tilde{h}} k_T(x)\, dx \tag{231}$$

with

$$k_T(x) = \sum_{j=-\infty}^{\infty} \left[n_T(x + 2j(\tilde{h} - \tilde{l})) - n_T(x - 2\tilde{h} + 2j(\tilde{h} - \tilde{l})) \right]. \tag{232}$$

Hence the joint density of the maximum and the minimum of a Brownian motion with drift $\tilde{\theta}$, $W_t^{\tilde{\theta}} \stackrel{\Delta}{=} W_t + \tilde{\theta} t$, is given by

$$k_T^{\tilde{\theta}}(x) = k_T(x) \exp\left\{ \tilde{\theta} x - \tfrac{1}{2}\tilde{\theta}^2 T \right\}. \tag{233}$$

We obtain for the value of the double-no-touch on $[t, \tau]$

$$\begin{aligned} v(t) &= e^{-r_d(T-t)} I\!E I\!I_{\{L \le \min_{[0,T]} S_t < \max_{[0,T]} S_t \le H\}} \\ &= e^{-r_d(T-t)} I\!E I\!I_{\{\tilde{l} \le \min_{[0,T]} W_t^{\tilde{\theta}_-} < \max_{[0,T]} W_t^{\tilde{\theta}_-} \le \tilde{h}\}} \\ &= e^{-r_d(T-t)} \int_{\tilde{l}}^{\tilde{h}} k_{(T-t)}^{\tilde{\theta}_-}(x) dx \qquad (234) \\ &= e^{-r_d(T-t)} \qquad (235) \\ &\quad \cdot \sum_{j=-\infty}^{\infty} \left[e^{-2j\theta_-(h-l)} \left\{ \mathcal{N}\left(h + y_-\right) - \mathcal{N}\left(l + y_-\right) \right\} \right. \\ &\quad \left. - e^{-2j\theta_-(h-l) + 2\theta_- h} \left\{ \mathcal{N}\left(h - 2h + y_-\right) - \mathcal{N}\left(l - 2h + y_-\right) \right\} \right] \end{aligned}$$

and on $[\tau, T]$

$$v(t) = e^{-r_d(T-t)} I\!I_{\{L \le \min_{[0,T]} S_t < \max_{[0,T]} S_t \le H\}}. \tag{236}$$

Of course, the value of the double-one-touch (DOT) on $[t, \tau]$ is given by

$$e^{-r_d(T-t)} - v(t), \tag{237}$$

following the fact that a DNT and a DOT in combination are economically equivalent to the zero-coupon bond.

Greeks We take the space to list some of the sensitivity parameters I have been asked about frequently.

Vega

$$\begin{aligned} v_\sigma(t) = \frac{e^{-r_d(T-t)}}{\sigma} \cdot \sum_{j=-\infty}^{\infty} \\ \Big\{ e^{-2j\theta_-(h-l)} & \big[2j(h-l)(\theta_+ + \theta_-)\left\{\mathcal{N}\left(h+y_-\right) - \mathcal{N}\left(l+y_-\right)\right\} \\ & + n(h+y_-)(-h-y_+) - n(l+y_-)(-l-y_+)\big] \\ - e^{-2j\theta_-(h-l)+2\theta_- h} & \big[2(\theta_+ + \theta_-)(j(h-l)-h)\left\{\mathcal{N}\left(-h+y_-\right) - \mathcal{N}\left(l-2h+y_-\right)\right\} \\ & + n(-h+y_-)(h-y_+) - n(l-2h+y_-)(-l+2h-y_+)\big]\Big\} \end{aligned} \tag{238}$$

Vanna

$$v_{\sigma S_t}(t) = \frac{e^{-r_d(T-t)}}{S_t \sigma^2 \sqrt{T-t}} \cdot \sum_{j=-\infty}^{\infty} \left\{ e^{-2j\theta_-(h-l)}(T_1 - T_2) - e^{-2j\theta_-(h-l)+2\theta_- h}(T_3 + T_4 - T_5) \right\} \tag{239}$$

$$T_1 = n(h+y_-)\left\{1 - 2j(h-l)(\theta_+ + \theta_-) - (h+y_-)(h+y_+)\right\} \tag{240}$$

$$T_2 = n(l+y_-)\left\{1 - 2j(h-l)(\theta_+ + \theta_-) - (l+y_-)(l+y_+)\right\} \tag{241}$$

$$\begin{aligned} T_3 = {} & 2(\theta_+ + \theta_-)\left[-2\theta_- j(h-l) + 2\theta_- h + 1\right] \\ & \cdot \left\{\mathcal{N}(-h+y_-) - \mathcal{N}(l-2h+y-)\right\} \end{aligned} \tag{242}$$

$$\begin{aligned} T_4 = {} & n(-h+y_-)\left\{-2\theta_-(h-y_+) + 2(\theta_+ + \theta_-)(j(h-l)-h) \right. \\ & \left. + (h-y_-)(h-y_+) - 1\right\} \end{aligned} \tag{243}$$

$$\begin{aligned} T_5 = {} & n(l-2h+y_-)\left\{-2\theta_-(-l+2h-y_+) + 2(\theta_+ + \theta_-)(j(h-l)-h) \right. \\ & \left. + (-l+2h-y_-)(-l+2h-y_+) - 1\right\} \end{aligned} \tag{244}$$

Volga

$$v_{\sigma\sigma}(t) = \frac{e^{-r_d(T-t)}}{\sigma^2} \cdot \sum_{j=-\infty}^{\infty} \left\{ e^{-2j\theta_-(h-l)}(T_1 + T_2) - e^{-2j\theta_-(h-l)+2\theta_- h}(T_3 + T_4) \right\} \tag{245}$$

$$T_1 = (2j(\theta_+ + \theta_-)(h - l) - 3)\left\{2j(h - l)(\theta_+ + \theta_-)\left[\mathcal{N}(h + y_-) - \mathcal{N}(l + y_-)\right]\right\}$$
$$+ (4j(\theta_+ + \theta_-)(h - l) - 1)\left[n(h + y_-)(-h - y_+) - n(l + y_-)(-l - y_+)\right] \quad (246)$$
$$T_2 = n(h + y_-)(h + y_-)\left[1 - (h + y_+)^2\right] - n(l + y_-)(l + y_-)\left[1 - (l + y_+)^2\right] \quad (247)$$
$$T_3 = (2(\theta_+ + \theta_-)(j(h - l) - h) - 3)\left\{2(\theta_+ + \theta_-)(j(h - l) - h)\right.$$
$$\left.\cdot \left[\mathcal{N}(-h + y_-) - \mathcal{N}(l - 2h + y_-)\right]\right\} + (4(\theta_+ + \theta_-)(j(h - l) - h) - 1)$$
$$\cdot \left[n(-h + y_-)(h - y_+) - n(l - 2h + y_-)(-l + 2h - y_+)\right] \quad (248)$$
$$T_4 = n(-h + y_-)(h - y_-)\left[(h - y_+)^2 - 1\right]$$
$$- n(l - 2h + y_-)(-l + 2h - y_-)\left[(-l + 2h - y_+)^2 - 1\right] \quad (249)$$

Applications of Double-Touch Contracts A double-no-touch is a very liquid instrument. Going long a DNT reflects a view of spot staying quiet over the lifetime of the contract. Long a DNT means short vega, so it is a contract one can buy and be vega short. DNT and DOT serve as rebates to double-knock-out and double-knock-in options. They also appear as building blocks in structured FX forwards and yield enhancing deposits like wedding cakes, towers, and onions – see Section 2.4.6.

1.7.5 Compound and Installment

Compound Options A compound call (put) option gives its holder the right to exercise and upon exercise buy (sell) a vanilla option called the daughter option for a pre-specified price called the strike of the mother option. It works in a similar way to a vanilla call, but allows the holder to pay the premium of the call option spread over time. A first payment is made on inception of the trade. On the following payment day the holder of the compound call can decide to turn it into a plain vanilla call, in which case he has to pay the second part of the premium, or terminate the contract by simply not paying any more.

Advantages

- Full protection against stronger EUR/weaker USD
- Maximum loss is the premium paid
- Initial premium required is less than in the vanilla call
- Easy termination process, especially useful if future cash flows are uncertain

Disadvantages

- Premium required (not a zero-cost structure)
- More expensive than the vanilla call

Example A company wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD. The company wishes to be able to buy EUR at a lower spot rate if the EUR becomes weaker on the one hand, but

TABLE 1.22 Example of a compound call option.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional	USD 1,000,000
Company buys	EUR call USD put strike 1.1500
Premium per half year of the compound	USD 23,000.00
Premium of the vanilla call	USD 40,000.00

on the other hand be fully protected against a stronger EUR. The future income in USD is uncertain but will be under review at the end of the next half year.

In this case a possible form of protection that the company can use is to buy a EUR compound call option with two equal semi-annual premium payments as illustrated in Table 1.22 for example.

The company pays 23,000 USD on the trade date. After a half year, the company has the right to buy a plain vanilla call. To do this the company must pay another 23,000 USD.

Of course, besides not paying the premium, another way to terminate the contract is always to sell it in the market or to the seller. So if the option is not needed but deep in-the-money, the company can take profit from paying the premium to turn the compound into a plain vanilla call and then sell it.

If the EUR-USD exchange rate is above the strike at maturity, then the company can buy EUR at maturity at a rate of 1.1500.

If the EUR-USD exchange rate is below the strike at maturity, the option expires worthless. However, the company would benefit from being able to buy EUR at a lower rate in the market.

Variations of Compound Options

Distribution of payments The payments do not have to be equal. However, the rule is that the more premium is paid later, the higher the total premium. The cheapest distribution of payments is to pay the entire premium in the beginning, which corresponds to a plain vanilla call.

Exercise style Both the mother and the daughter of the compound option can be European and American style. The market default is European style.

Compound strategies One can think of a compound option on any structure, as for instance a compound put on a knock-out call or a compound call on a structured forward.

Forward Volatility The daughter option of the compound requires knowing the volatility for its lifetime, which starts on the exercise date T_1 of the mother option and ends on the maturity date T_2 of the daughter option. This volatility is not known at inception of the trade, so the only proxy traders can take is the forward volatility $\sigma(T_1, T_2)$ for this time interval. In the Black-Scholes model the consistency equation for the forward volatility is given by Equation (152).

The more realistic way to look at this unknown forward volatility is that the fairly liquid market of forward volatility-sensitive derivatives could be taken to back out the forward volatilities since this is the only unknown. These derivatives include forward start options and forward volatility agreements.

In a market with smile, the payoff of the compound option can be approximated by a linear combination of vanillas, whose market prices are known. For the payoff of the compound option itself we can take the forward volatility as in Equation (152) for the at-the-money value and the smile of today as a proxy. More details on this can be found in Schilling [116] for example. The actual forward volatility, however, is a trader's view and with any luck can be taken from other market prices.

Installment Options This section is based on Griebsch *et al.* – see [61].

An installment call option generalizes the idea of a compound call: it allows the holder to pay the premium of the call option in more than two installments spread over time. A first payment is made at inception of the trade. On the following payment days the holder of the installment call can decide to prolong the contract, in which case he has to pay the second installment of the premium, or to terminate the contract by simply not paying any more. After the last installment payment the contract turns into a plain vanilla call. We illustrate two scenarios in Figure 1.41.

Example A company wants to hedge receivables from an export transaction in USD due in 12 months' time. It expects a stronger EUR/weaker USD. The company wishes to be able to buy EUR at a lower spot rate if the EUR becomes weaker on the one hand, but on the other hand be fully protected against a stronger EUR. The future income in USD is uncertain but will be under review at the end of each quarter.

In this case a possible form of protection that the company can use is to buy a EUR installment call option with four equal quarterly premium payments as illustrated in Table 1.23 for example.

The company pays 12,500 USD on the trade date. After one quarter, the company has the right to prolong the installment contract. To do this the company must pay another 12,500 USD. After six months, the company has the right to prolong the contract and must pay 12,500 USD in order to do so. After nine months the same decision has to be taken. If on one of these three decision days the company does not pay, then the contract terminates. If all premium payments are made, then the contract turns into a plain vanilla EUR call.

Of course, besides not paying the premium, another way to terminate the contract is always to resell it in the market. So if the option is not needed but deep in-the-money,

TABLE 1.23 Example of an installment call.

Spot reference	1.1500 EUR-USD
Maturity	1 year
Notional	USD 1,000,000
Company buys	EUR call USD put strike 1.1500
Premium per quarter of the installment	USD 12,500.00
Premium of the vanilla call	USD 40,000.00

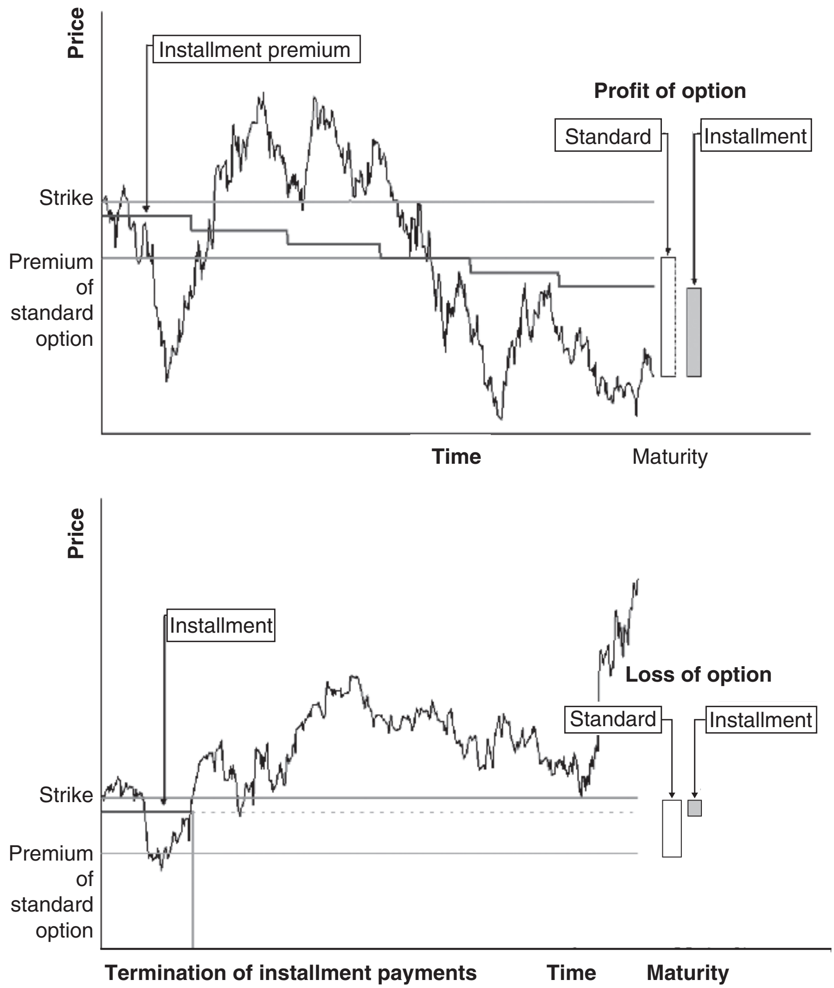

FIGURE 1.41 Comparison of two scenarios of an installment option. The top graph shows a continuation of all installment payments until expiration. The graph below shows a scenario where the installment option is terminated after the first decision date.

the company can take profit from paying the premium to prolong the contract and then sell it.

If the EUR-USD exchange rate is above the strike at maturity, then the company can buy EUR at maturity with a rate of 1.1500.

If the EUR-USD exchange rate is below the strike at maturity, the option expires worthless. However, the company would benefit from being able to buy EUR at a lower rate in the market.

Compound options can be viewed as a special case of installment options, and the possible variations of compound options apply analogously to installment options.

Reasons for Trading Compound and Installment Options We observe that in the buy-and-hold-to-maturity scenario, compound and installment options are always more expensive than buying a vanilla, sometimes substantially more expensive. So why are people buying them? The number one reason is an *uncertainty* about a future cash flow in a foreign currency. If the cash flow is certain, then buying a vanilla is in principle the better deal. An exception may be the situation that a treasurer has a budget constraint, i.e. limited funds to spend for foreign exchange risk. With an installment he can then split the premium over time. The main issue is, however, if a treasurer has to deal with an uncertain cash flow, and buys a vanilla instead of an installment, and then is faced with a far out-of-the-money vanilla at time T_1, then selling the vanilla does not give him as much as the savings between the vanilla and the sum of the installment payments.

The Theory of Installment Options This book is not concerned primarily with valuation of options. However, we do want to give some insight into selected topics that come up very often and are of particular relevance to foreign exchange options and have not been published in books so far. We will now take a look at the valuation, the implementation of installment options, and the limiting case of a continuous flow of premium payments.

Valuation in the Black-Scholes Model The intention of this section is to obtain a closed-form formula for the n-variate installment option in the Black-Scholes model. For the cases $n = 1$ and $n = 2$, the Black-Scholes formula and Geske's compound option formula (see [56]) are already well known.

We consider an exchange rate process S_t modeled by a geometric Brownian motion,

$$S_{t_2} = S_{t_1} \exp((r_d - r_f - \sigma^2/2)\Delta t + \sigma\sqrt{\Delta t}Z) \quad \text{for } 0 \le t_1 \le t_2 \le T, \tag{250}$$

where $\Delta t = t_2 - t_1$ and Z is a standard normal random variable independent of the past of S_t up to time t_1.

Let $t_0 = 0$ be the installment option inception date and $t_1, t_2, \ldots, t_n = T$ a schedule of decision dates in the contract on which the option holder has to pay the premiums $k_1, k_2, \ldots, k_{n-1}$ to keep the option alive. To compute the price of the installment option, which is the upfront payment V_0 at t_0 to enter the contract, we begin with the option payoff at maturity T

$$V_n(s) \stackrel{\Delta}{=} [\phi_n(s - k_n)]^+ \stackrel{\Delta}{=} \max[\phi_n(s - k_n), 0],$$

where $s = S_T$ is the price of the underlying asset at T and as usual $\phi_n = +1$ for a call option, $\phi_n = -1$ for a put option.

At time t_i the option holder can either terminate the contract or pay k_i to continue. Therefore by the risk-neutral pricing theory, the holding value is

$$e^{-r_d(t_{i+1}-t_i)}\mathbb{E}[V_{i+1}(S_{t_{i+1}}) \mid S_{t_i} = s], \quad \text{for } i = 0, \ldots, n-1, \tag{251}$$

where

$$V_i(s) = \begin{cases} \left[e^{-r_d(t_{i+1}-t_i)}\mathbb{E}[V_{i+1}(S_{t_{i+1}}) \mid S_{t_i} = s] - k_i\right]^+ & \text{for } i = 1, \ldots, n-1 \\ V_n(s) & \text{for } i = n \end{cases}. \tag{252}$$

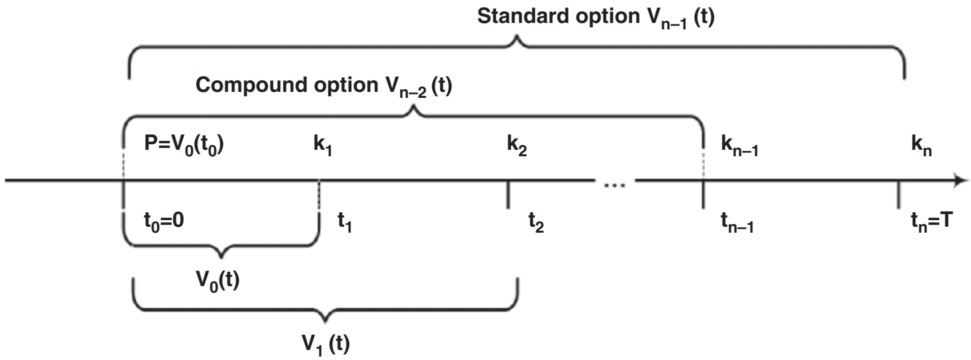

FIGURE 1.42 Lifetime structure of the options with value V_i for the i-th Option.

Then the unique arbitrage-free value of the initial premium is

$$P \stackrel{\Delta}{=} V_0(s) = e^{-r_d(t_1-t_0)}\mathbb{E}[V_1(S_{t_1}) \mid S_{t_0} = s]. \tag{253}$$

Figure 1.42 illustrates this context.

One way of pricing this installment option is to evaluate the nested expectations through multiple numerical integration of the payoff functions via backward iteration. Alternatively, one can derive a solution in closed form in terms of the n-variate cumulative normal.

The Curnow and Dunnett Integral Reduction Technique Denote the n dimensional multivariate normal integral with upper limits $h_1, \ldots, h_n$ and correlation matrix $R_n \stackrel{\Delta}{=} (\rho_{ij})_{i,j=1,\ldots,n}$ by $\mathcal{N}_n(h_1, \ldots, h_n; R_n)$, and the univariate normal density function by $n(\cdot)$. Let the correlation matrix be non-singular and $\rho_{11} = 1$.

Under these conditions Curnow and Dunnett [33] derived the following reduction formula for multivariate normal integrals:

$$\mathcal{N}_n(h_1, \ldots, h_n; R_n) = \int_{-\infty}^{h_1} \mathcal{N}_{n-1}\left(\frac{h_2-\rho_{21}y}{(1-\rho_{21}^2)^{1/2}}, \ldots, \frac{h_n-\rho_{n1}y}{(1-\rho_{n1}^2)^{1/2}}; R_{n-1}^*\right) n(y)dy,$$

$$R_{n-1}^* \stackrel{\Delta}{=} (\rho_{ij}^*)_{i,j=2,\ldots,n},$$

$$\rho_{ij}^* \stackrel{\Delta}{=} \frac{\rho_{ij}-\rho_{i1}\rho_{j1}}{(1-\rho_{i1}^2)^{1/2}(1-\rho_{j1}^2)^{1/2}}. \tag{254}$$

A Closed-Form Solution for the Value of an Installment Option Heuristically, the formula which is given in the theorem below has the structure of the Black-Scholes formula in higher dimensions, namely $S_0\mathcal{N}_n(\cdot) - k_n\mathcal{N}_n(\cdot)$ minus the later premium payments $k_i\mathcal{N}_i(\cdot)$ $(i = 1, \ldots, n-1)$. This structure is a result of the integration of the vanilla option payoff, which is again integrated minus the next installment, which in turn is integrated with the following installment and so forth. By this iteration the vanilla payoff is integrated with respect to the normal density function n times and the i-payment is integrated i times for $i = 1, \dots, n-1$.

The correlation coefficients ρ_{ij} of these normal distribution functions contained in the formula arise from the overlapping increments of the Brownian motion, which models the price process of the underlying S_t at the particular exercise dates t_i and t_j.

Theorem 1.7.1 *Let $\vec{k} = (k_1, \dots, k_n)$ be the strike price vector, $\vec{t} = (t_1, \dots, t_n)$ the vector of the exercise dates of an n-variate installment option, and $\vec{\phi} = (\phi_1, \dots, \phi_n)$ the vector of the put/call indicators of these n options.*

The value function of an n-variate installment option is given by

$$\begin{aligned}
V_n(S_0, M, \vec{k}, \vec{t}, \vec{\phi}) = {} & e^{-r_f t_n} S_0 \phi_1 \cdot \dots \cdot \phi_n \\
& \times \mathcal{N}_n \left[\frac{\ln \frac{S_0}{S_1^*} + \mu^{(+)} t_1}{\sigma\sqrt{t_1}}, \frac{\ln \frac{S_0}{S_2^*} + \mu^{(+)} t_2}{\sigma\sqrt{t_2}}, \dots, \frac{\ln \frac{S_0}{S_n^*} + \mu^{(+)} t_n}{\sigma\sqrt{t_n}}; R_n \right] \\
& - e^{-r_d t_n} k_n \phi_1 \cdot \dots \cdot \phi_n \\
& \times \mathcal{N}_n \left[\frac{\ln \frac{S_0}{S_1^*} + \mu^{(-)} t_1}{\sigma\sqrt{t_1}}, \frac{\ln \frac{S_0}{S_2^*} + \mu^{(-)} t_2}{\sigma\sqrt{t_2}}, \dots, \frac{\ln \frac{S_0}{S_n^*} + \mu^{(-)} t_n}{\sigma\sqrt{t_n}}; R_n \right] \\
& - e^{-r_d t_{n-1}} k_{n-1} \phi_1 \cdot \dots \cdot \phi_{n-1} \\
& \times \mathcal{N}_{n-1} \left[\frac{\ln \frac{S_0}{S_1^*} + \mu^{(-)} t_1}{\sigma\sqrt{t_1}}, \frac{\ln \frac{S_0}{S_2^*} + \mu^{(-)} t_2}{\sigma\sqrt{t_2}}, \dots, \frac{\ln \frac{S_0}{S_{n-1}^*} + \mu^{(-)} t_{n-1}}{\sigma\sqrt{t_{n-1}}}; R_{n-1} \right] \\
& \vdots \\
& - e^{-r_d t_2} k_2 \phi_1 \phi_2 \mathcal{N}_2 \left[\frac{\ln \frac{S_0}{S_1^*} + \mu^{(-)} t_1}{\sigma\sqrt{t_1}}, \frac{\ln \frac{S_0}{S_2^*} + \mu^{(-)} t_2}{\sigma\sqrt{t_2}}; \rho_{12} \right] \\
& - e^{-r_d t_1} k_1 \phi_1 \mathcal{N} \left[\frac{\ln \frac{S_0}{S_1^*} + \mu^{(-)} t_1}{\sigma\sqrt{t_1}} \right],
\end{aligned} \tag{255}$$

where S_i^* *(i = 1, … , n) is to be determined as the spot price* S_t *for which the payoff of the corresponding i-installment option (i = 1, … , n) is equal to zero and* $\mu^{(\pm)}$ *is defined as* $r_d - r_f \pm \frac{1}{2}\sigma^2$.

The correlation coefficients in R_i *of the i-variate normal distribution function can be expressed through the exercise dates* t_i,

$$\rho_{ij} = \sqrt{t_i/t_j} \text{ for } i,j = 1, \dots, n \text{ and } i < j. \tag{256}$$

The proof is established with Equation (254). Formula (255) has been independently derived by Thomassen and van Wouwe in [125].

Valuation of Installment Options with the Algorithm of Ben-Hameur, Breton, and François The value of an installment option at time t is given by the snell envelope of the discounted payoff processes, which is calculated with the dynamic programming method used by the algorithm of Ben-Hameur, Breton, and François below. Their original work in [14] deals with installment options with an additional exercise right at each installment date. This means that at each decision date the holder can either exercise, terminate, or continue.

We examine this algorithm now for the special case of zero value in case it is exercised at $t_1,\dots,t_{n-1}$. The difference between the above mentioned types of installment options consists in the (non-)existence of an exercise right at the installment dates, but this does not change the algorithm in principle.

Model Description The algorithm developed by Ben-Hameur, Breton, and François approximates the value of the installment option in the Black-Scholes model, which is the premium P paid at time t_0 to enter the contract.

The exercise value of an installment option at maturity t_n is given by $V_n(s) \stackrel{\Delta}{=} \max[0, \phi_n(s - k_n)]$ and zero at earlier times. The value of a vanilla option at time t_{n-1} is denoted by $V_{n-1}(s) = e^{-r_d \Delta t} I\!E[V_n(s) \mid S_{t_{n-1}} = s]$. At an arbitrary time t_i the holding value is determined as

$$V_i^h(s) = e^{-r_d \Delta t} I\!E[V_{i+1}(S_{t_{i+1}}) \mid S_{t_i} = s] \text{ for } i = 0, \dots, n-1, \tag{257}$$

where

$$V_i(s) = \begin{cases} V_0^h(s) & \text{for } i = 0, \\ \max[0, V_i^h(s) - k_i] & \text{for } i = 1, \dots, n-1, \text{ (DP)}. \\ V_n^e(s) & \text{for } i = n. \end{cases} \tag{258}$$

The function $V_i^h(s) - k_i$ is called net holding value at t_i, for $i = 1, \dots, n-1$, which is shown in Figure 1.43.

The option value is the holding value or the exercise value, whichever is greater. The value function V_i, for $i = 0, \dots, n-1$, is unknown and has to be approximated. Ben-Hameur, Breton, and François propose an approximation method, which solves the above dynamic programming (DP)-equation (258) in a closed form for all s and i.

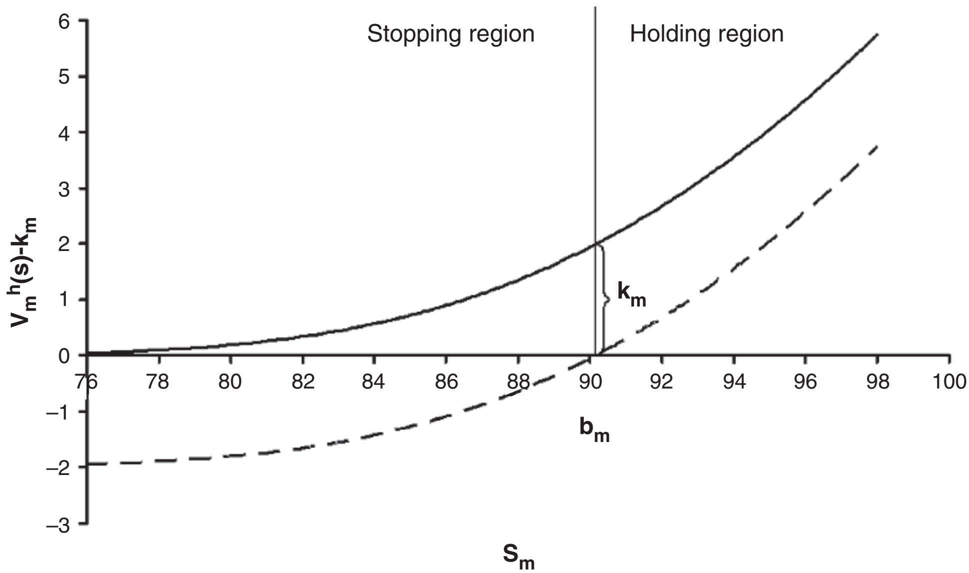

FIGURE 1.43 The holding value shortly before t_3 for an installment option with 4 rates is shown by the solid line. The positive slope of this function is less than 1 and the function is continuous and convex. The net holding value of an installment call option $V^h_m(s) - k_m$ for $(s > 0)$ and a decision time m is presented by the dashed line. This curve intersects the x-axis in the point, where it divides the stopping region and the holding region. The value function is zero in the stopping region $(0, b_i)$ and equal to the net holding value in the holding region $[b_i, \infty)$, where b_i is a threshold for every time t_i, which depends on the parameters of the installment option.

Valuation of Installment Options with Stochastic Dynamic Programming The idea of the above mentioned authors is to partition the positive real axis into intervals and approximate the option value through piecewise linear interpolation. Let $a_0 = 0 < a_1 < \cdots < a_p < a_{p+1} = +\infty$ be points in $\mathbb{R}_0^+ \cup \{\infty\}$ and $(a_j, a_{j+1}]$for $j = 0, \ldots, p$ a partition of $\mathbb{R}_0^+$ in (p+1) intervals.

Given approximations $\tilde{V}_i$ of option values V_i at supporting points a_j at the i-th step (at the beginning of the algorithm, at T, this is provided through the input values), this function is piecewise linearly interpolated by

$$\hat{V}_i(s) = \sum_{i=0}^{p} (\alpha^i_j + \beta^i_j s) \mathbb{I}_{\{a_j < s \le a_{j+1}\}}, \tag{259}$$

where $\mathbb{I}$ is the indicator function. The local coefficients of this interpolation in step i, the y-axis intercepts α^i_j and the slopes β^i_j, are obtained by solving the following linear equations,

$$\tilde{V}_i(a_j) = \hat{V}_i(a_j) \text{ for } j = 0, \ldots, p-1. \tag{260}$$

For $j = p$, one chooses

$$\alpha_p^i = \alpha_{p-1}^i \text{ and } \beta_p^i = \beta_{p-1}^i. \tag{261}$$

Now it is assumed that $\hat{V}_{i+1}$ is known. This is a valid assumption in this context because the values $\hat{V}_{i+1}$ are known from the previous step. The mean value (257) is calculated in step i through

$$\begin{aligned}\tilde{V}_i^h(a_k) &= e^{-r_d\Delta t}\mathbb{E}[\hat{V}_{i+1}(S_{t_{i+1}})|S_{t_i} = a_k] \\ &\overset{(259)}{=} e^{-r_d\Delta t}\sum_{j=0}^{p}\alpha_j^{i+1}\mathbb{E}\left[\mathbb{I}_{\left\{\frac{a_j}{a_k}<e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\leq\frac{a_{j+1}}{a_k}\right\}}\right] \\ &\quad + \beta_j^{i+1}a_k\mathbb{E}\left[e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\mathbb{I}_{\left\{\frac{a_j}{a_k}<e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\leq\frac{a_{j+1}}{a_k}\right\}}\right],\end{aligned} \tag{262}$$

where $\mu \triangleq r_d - r_f - \sigma^2/2$ and $\tilde{V}_i^h$ denotes the approximated holding value of the installment option. Define

$$x_{k,j} \triangleq \frac{\ln\left(\frac{a_j}{a_k}\right) - \mu\Delta t}{\sigma\sqrt{\Delta t}}, \tag{263}$$

so for $k = 1, \ldots, p$ and $j = 0, \ldots, p$ the first mean values in Equation (262), namely

$$A_{k,j} \triangleq \mathbb{E}\left[\mathbb{I}_{\left\{\frac{a_j}{a_k}<e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\leq\frac{a_{j+1}}{a_k}\right\}}\right] \tag{264}$$

can be expressed as

$$(264) = \begin{cases}\mathcal{N}(x_{k,1}) & \text{for } j = 0, \\ \mathcal{N}(x_{k,j+1}) - \mathcal{N}(x_{k,j}) & \text{for } 1 \leq j \leq p-1, \\ 1 - \mathcal{N}(x_{k,p}) & \text{for } j = p,\end{cases} \tag{265}$$

and similarly

$$B_{k,j} \triangleq \mathbb{E}\left[a_k e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\mathbb{I}_{\left\{\frac{a_j}{a_k}<e^{\mu\Delta t+\sigma\sqrt{\Delta t}z}\leq\frac{a_{j+1}}{a_k}\right\}}\right] \tag{266}$$

can be expressed in the following way

$$(266) = \begin{cases}a_k\mathcal{N}(x_{k,1} - \sigma\sqrt{\Delta t})e^{(r_d-r_f)\Delta t} & \text{for } j = 0, \\ a_k[\mathcal{N}(x_{k,j+1} - \sigma\sqrt{\Delta t}) - \mathcal{N}(x_{k,j} - \sigma\sqrt{\Delta t})]e^{(r_d-r_f)\Delta t} & \text{for } 1 \leq j \leq p-1, \\ a_k[1 - \mathcal{N}(x_{k,p} - \sigma\sqrt{\Delta t})]e^{(r_d-r_f)\Delta t} & \text{for } j = p,\end{cases} \tag{267}$$

where $n(z) \triangleq \frac{1}{\sqrt{2\pi}}e^{-\frac{z^2}{2}}$ and $\mathcal{N}$ denotes the cumulative normal distribution function.

In the simplifying notation (265) and (267) the points a_i $(i = 1, \dots, p)$ can be understood as the quantiles of the log-normal distribution. These are not chosen directly but are calculated as the quantiles of (e.g. equidistant) probabilities of the log-normal distribution. Thereby the supporting points lie closer together, in areas where the modification rate of the distribution function is great. The number a_k in Equation (263) is the given exchange rate at time t_i and therefore constant. In the implementation it requires an efficient method to calculate the inverse normal distribution function. One possibility is to use the Cody-Algorithm taken from [21].

An Algorithm in Pseudo Code For a better understanding let us sketch the procedure described in Section 1.7.5 in the form of an algorithm. The algorithm works according to the dynamic programming principle backwards in time, based on the values of the exercise function of the installment option at maturity T at pre-determined supporting points a_j. Through linear connection of these points an approximation of the exercise function can be obtained. The exercise function at maturity is the payoff function of the vanilla option, which is constant up to the strike price K and in the region behind (i.e. $\geq K$) it is linear. The linear approximation at maturity T is therefore exact, except on the interval $K \in (a_l, a_{l+1})$, in case K does not correspond to one of these supporting points. For this reason the holding value of this linear approximation is calculated by the means of $A_{k,j}$ and $B_{k,j}$ from above. The transition parameters $A_{k,i}$ and $B_{k,j}$ can be calculated before the first iteration, because only values which are known in the beginning are required. The advantage of this approach is that the holding value needs to be calculated only at the supporting points a_j and because of linearity, the function values for all s are obtained. The values of the holding value at a_j are used again as approximations of the exercise values at time t_{n-1} and it proceeds as in the beginning. The output of the algorithm is the value of the installment option at time t_0.

A Description in Pseudo Code First the a_k are generated as quantiles of the distribution of the price at maturity of the exchange rate S_T and can be approximated by the Cody-Algorithm, for example.

1. Calculate $q_1, \dots, q_p$-quantiles of the standard normal distribution via the inverse distribution function.
2. Calculate $a_1, \dots, a_p$-quantiles of the log-normal distribution with mean $\log S_0 - \mu T$ and variance $\sigma\sqrt{T}$ by

$$\exp(q_i \sigma \sqrt{T} + \log S_0 + \mu T) = a_i.$$

In pseudo code the implementation of the theoretical consideration of Section 1.7.5 can be worked out in the following way. The principle of the backward induction is realized as a for-loop that counts backwards from $n - 1$ to 0.

1. Calculate $\hat{V}_n(s)$ for all s, using (259), i.e. calculate all α_i^n, β_i^n for $i = 0, \dots, p$.
2. For $j = n$ to 1
   a. Calculate $\tilde{V}_{j-1}^h(a_k)$ for a_k $(k = 1, \dots, p)$ in closed form using (262).
   b. Calculate $\tilde{V}_{j-1}(a_k)$ for $k = 1, \dots, p$ using (DP) with $\tilde{V}_{j-1}^h(a_k)$ for $V_{j-1}^h(a_k)$.

c. Calculate $\hat{V}_{j-1}(s)$ for all $s > 0$ using (259), i.e. calculate all α_i^{j-1}, β_i^{j-1} for $i = 1, \ldots, k$. Unless $j - 1$ is already equal to zero, calculate $\hat{V}_{j-1}(s)$ for $s = S_0$ and break the algorithm.
d. Substitute $j \leftarrow j - 1$.

Repeat these steps until $\hat{V}_0(S_0)$ is calculated, which is the value of the installment option at time 0.

This algorithm works with equidistant installment dates, constant volatility, and constant interest rates. Constant volatility and interest rates are assumptions of the applied Black-Scholes Model, but the algorithm would be extendable for piecewise constant volatility and interest rate as functions of time, with jumps at the installment dates. The interval length Δt in the calculation can be replaced in every period by arbitrary $t_{i+1} - t_i$. Furthermore the computational time could be decreased by omitting smaller supporting points in the calculation as soon as one of them generates a zero value in the maximum function.

Installment Options with a Continuous Payment Plan Let $g = (g_t)_{t\in[0,T]}$ be the stochastic process describing the discounted net payoff of an installment option expressed as multiples of the domestic currency. If the holder stops paying the premium at time t, the difference between the option payoff and premium payments (all discounted to time 0) amounts to

$$g(t) = \begin{cases} e^{-r_d T}(S_T - K)^+ \mathbf{1}_{(t=T)} - \frac{p}{r_d}(1 - e^{-r_d t}) & \text{if } r_d \neq 0 \\ (S_T - K)^+ \mathbf{1}_{(t=T)} - pt & \text{if } r_d = 0 \end{cases}, \tag{268}$$

where K is the strike. Given the premium rate p, there is a unique no-arbitrage premium P_0 to be paid at time 0 (supplementary to the rate p) given by

$$P_0 = \sup_{\tau\in\mathcal{T}_{0,T}} \mathbb{E}_Q(g_\tau). \tag{269}$$

Ideally, p is chosen as the *minimal* rate such that

$$P_0 = 0. \tag{270}$$

Note that P_0 from (269) can never become negative as it is always possible to stop payments immediately. Thus, besides (270), we need a minimality assumption to obtain a unique rate. We want to compare the installment option with the American contingent claim $f = (f_t)_{t\in[0,T]}$ given by

$$f_t = e^{-r_d t}(K_t - C_E(T - t, S_t))^+, t \in [0, T], \tag{271}$$

where $K_t = \frac{p}{r_d}\left(1 - e^{-r_d(T-t)}\right)$ for $r_d \neq 0$ and $K_t = p(T - t)$ when $r_d = 0$. C_E is the value of a standard European call. Equation (271) represents the payoff of an American put on a European call where the variable strike K_t of the put equals the part of the installments *not* to be paid if the holder decides to terminate the contract at time t. Define by $\tilde{f} = (\tilde{f}_t)_{t\in[0,T]}$ a similar American contingent claim with

$$\tilde{f}(t) = e^{-r_d t}\left[(K_t - C_E(T-t,S_t))^+ + C_E(T-t,S_t)\right], \quad t \in [0,T]. \tag{272}$$

As the process $t \mapsto e^{-r_d t} C_E(T-t,S_t)$ is a Q-martingale we obtain that

$$\sup_{\tau\in\mathcal{T}_{0,T}} \mathbb{E}_Q(\tilde{f}_\tau) = C_E(T,s_0) + \sup_{\tau\in\mathcal{T}_{0,T}} \mathbb{E}_Q(f_\tau). \tag{273}$$

The following theorem has been proved in [60] using earlier results of El Karoui *et al.* in [47].

Theorem 1.7.2 *An installment option is the sum of a European call plus an American put on this European call, i.e.*

$$\underbrace{P_0 + p\int_0^T e^{-r_d s}\,ds}_{\textit{total premium payments}} = C_E(T,s_0) + \sup_{\tau\in\mathcal{T}_{0,T}} \mathbb{E}_Q(f_\tau)$$

1.7.6 Asian Options

This section is joint work with Silvia Baumann, Marion Linck, Michael Mohr, and Michael Seeberg.

Asian Options are options on the average usually of spot fixings and are very popular and common hedging instruments for corporates. Average options belong to the class of path dependent options. The Term *Asian Options* comes from their origin in the Tokyo office of Bankers Trust in 1987.[4] The payoff of an Asian Option is determined by the path taken by the underlying exchange rate over a fixed period of time. We distinguish the four cases listed in Table 1.24 and compare values of average price options with vanilla options in Figure 1.44.

Variations of Asian options refer particularly to the way the average is calculated.

TABLE 1.24 Types of Asian options for $T_0 \le t \le T$, where $[T_0, T]$ denotes the time interval over which the average is taken, K denotes the strike, S_T the spot price at expiration time and A_T the average.

Product name	**Payoff**	**Product name**	**Payoff**
Average price call	$(A_T - K)^+$	Average strike call	$(S_T - A_T)^+$
Average price put	$(K - A_T)^+$	Average strike put	$(A_T - S_T)^+$

[4] see https://en.wikipedia.org/wiki/Asian_option

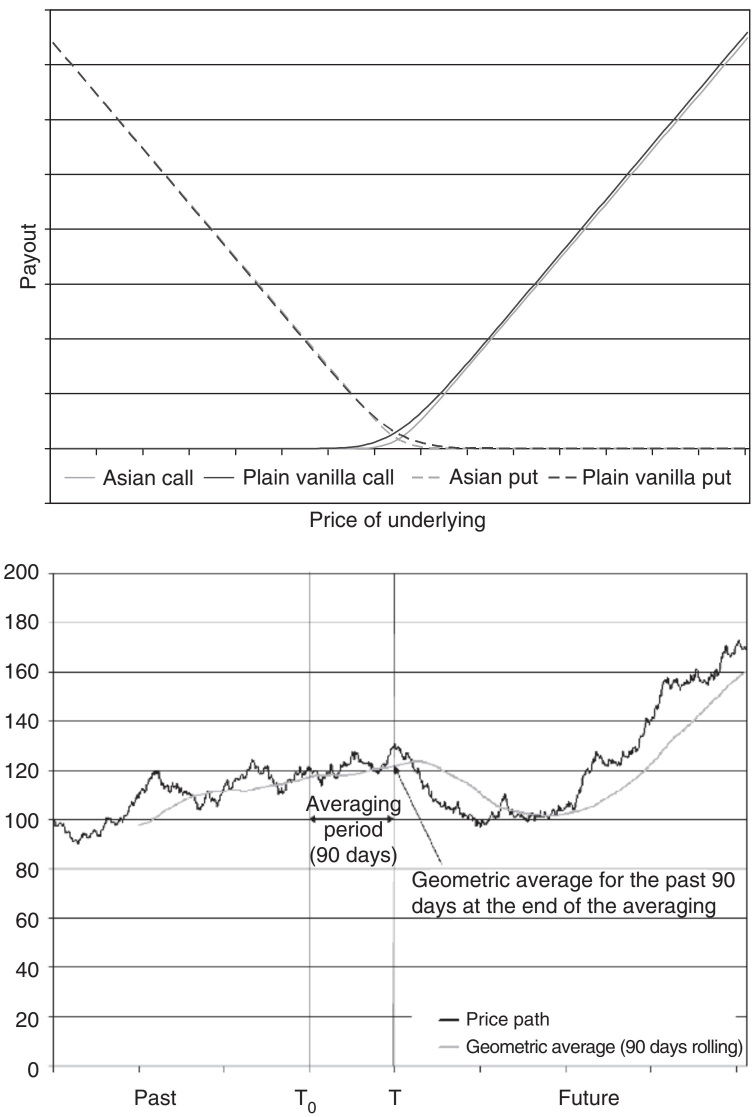

FIGURE 1.44 Above: comparing the value of average price options with vanillas, we see that average price options are cheaper. The reason is that averages are less volatile and hence less risky. Below: ingredients for average options: a price path, 90-days rolling price average (here: geometric), and an averaging period for an option with 90-days maturity.

Kind of average We find geometric, arithmetic, or harmonic average of prices. Harmonic averaging originates from a payoff in *domestic* currency and will be treated in Section 1.8.7.

Time interval We need to specify the period over which the prices are taken. The end of the averaging interval can be shorter than or equal to the option's expiration date, the starting value can be any time before. In particular, after an average option is traded, the beginning of the averaging period typically lies in the past, so that parts of the values contributing to the average are already known.

Sampling style The market generally uses discrete sampling, like daily fixings. In the literature we often find continuous sampling.

Weighting Different weights may be assigned to the prices to account for a non-linear, i.e. skewed, price distribution – see Hansen and Jorgensen [69], pp. 1116–1117, and the example below under 3.

Variations The wide range of variations covers also the possible right for early exercise, Asian options with barriers.

Asian Options are applied in risk management, especially for currencies for the following reasons.

1. Protection against rapid price movements or manipulation in thinly traded underlyings at maturity, i.e. reduction of significance of reference prices through averaging.
2. Reduction of hedging cost through
   - the lower fair value compared with regular options since an average is less volatile than single prices, and
   - to achieve a similar hedging effect as with vanilla options, buying a chain of such options – obviously a more expensive strategy.
3. Adjustment of option payoff to payment structure of the firm
   - Average price options can be used to hedge a stream of (received) payments (e.g. a USD average call can be bought to hedge the ongoing EUR revenues of a US-based company). Different amounts of the payments can be reflected in flexible weights, i.e. the prices related to higher payments are assigned a higher weight than those related to smaller cash flows when calculating the average.
   - With average strike options the strike price can be set at the average of the underlying price – a helpful structure in volatile or hardly predictable markets.

Valuation The valuation approaches developed differ depending on the specific characteristics considered, for example averaging method, option style, etc. In the following, we present the value formula for a *European geometric average price call*. In the sequel two common approaches to evaluating arithmetic average price options are introduced. Henderson and Wojakowski prove the symmetry between average price options and average strike options in [72], allowing the use of the more established fixed-strike valuation methods to price *floating strike Asian options*. Asian options were the first ones traded in history, where *pre-trade valuation* and *post-trade valuation* are different.

Geometric Average Options Kemna and Vorst [86] derive a closed form solution for geometric average price options in a geometric Brownian motion model

$$dS_t = S_t[(r_d - r_f)dt + \sigma dW_t]. \tag{274}$$

A geometric average price call pays $(A_T - K)^+$, where A_T denotes the geometric average of the foreign exchange rate. In the discrete case, A_T is calculated as

$$A_T \stackrel{\Delta}{=} \sqrt[n+1]{\prod_{i=0}^{n} S_{t_i}}, \tag{275}$$

in the continuous case as

$$A_T \stackrel{\Delta}{=} \exp\left\{\frac{1}{T-T_0}\int_{T_0}^{T} \log S_t \, dt\right\}. \tag{276}$$

The random variable $\int_t^T W(u)\, du$ is normally distributed with mean zero variance

$$\Sigma^2 \stackrel{\Delta}{=} \frac{T^3}{3} + \frac{2t^3}{3} - t^2 T \tag{277}$$

for any $t \in [T_0, T]$. This can be calculated following the instructions in Shreve's lecture notes [119]. Therefore, the geometric average of a log-normally distributed random variable is log-normally distributed. In the continuous case, the distribution parameters can be derived as

$$\log A_T \sim \mathcal{N}\left[\tfrac{1}{2}\left(r_d - r_f - \tfrac{1}{2}\sigma^2\right)(T-T_0) + \log S_0; \tfrac{1}{3}\sigma^2(T-T_0)\right]. \tag{278}$$

The interesting feature of these terms is that the average has half the drift and one third of the variance of the spot price. In the Black-Scholes model the value of the option can be computed as the expected payoff under the risk-neutral probability measure. Using the money market account $e^{-r_d(T-T_0)}$ as numeraire leads to the value of the continuously sampled geometric Asian fixed strike call,

$$C_{G-Asian} = I\!\!E\left[e^{-r_d(T-T_0)}(A_T - K) I\!I_{\{A_T > K\}}\right], \tag{279}$$

where we observe that the remaining computation works just like a vanilla. In order to derive a useful general result we need to generalize the payoff of the continuously sampled geometric Asian fixed strike option to

$$[\phi(A(-s,T) - K)]^+, \tag{280}$$

$$A(-s,T) \stackrel{\Delta}{=} \exp\left\{\frac{1}{T+s}\int_{-s}^{T} \log S(u)\, du\right\}, s \geq 0. \tag{281}$$

This definition includes the case where parts of the average are already known, which is important to value the option after it has been written (post-trade valuation).

With the abbreviations

- T for the expiration time (in years),
- s for the time before valuation date (in years), for which the values and average of the underlying are known,
- K for the strike of the option,
- ϕ taking the values +1 or −1 if the option is a call or a put respectively,
- $\alpha \stackrel{\Delta}{=} \frac{T}{T+s} \in [0, 1]$,
- $\theta_\pm \stackrel{\Delta}{=} \frac{r_d - r_f}{\sigma} \pm \frac{\sigma}{2}$,
- $S_t = S_0 e^{\sigma W_t + \sigma\theta\_t}$ for the price of the underlying at time t,
- $d_\pm \stackrel{\Delta}{=} \frac{\ln \frac{S_0}{K} + \sigma\theta_\pm T}{\sigma\sqrt{T}}$,
- $n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2}$,
- $\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt$,
- $\text{Vanilla}(S_0, K, \sigma, r_d, r_f, T, \phi) = \phi \left(S_0 e^{-r_f T} \mathcal{N}(\phi d_+) - K e^{-r_d T} \mathcal{N}(\phi d_-) \right)$,
- $H \stackrel{\Delta}{=} \exp\left\{ -\frac{\alpha T}{2} \left(r_d - r_f + \frac{\sigma^2}{2}[1 - \frac{2\alpha}{3}] \right) \right\}$,

the value of the continuously sampled geometric Asian fixed strike call or put is then given by

$$\text{Asiangeo}(S_0, K, T, s, \sigma, r_d, r_f, \phi) = e^{(\alpha-1) r_d T} H \left(\frac{S_0}{A(-s, 0)} \right)^{\alpha - 1}$$

$$\text{Vanilla}\left(S_0, \frac{K}{H} \left(\frac{S_0}{A(-s, 0)} \right)^{1-\alpha}, \frac{\alpha\sigma}{\sqrt{3}}, \alpha r_d, \alpha r_f, T, \phi \right). \tag{282}$$

This way a geometric Asian option with fixed strike can be viewed as a version of a vanilla option with the same spot and time to maturity but different parameters such as notional, strike, volatility, and interest rates. We observe in particular that as time to maturity becomes smaller, the known part of the average becomes more prominent, α tends to zero, and hence the volatility of the auxiliary vanilla option tends to zero. Moreover, the properties known for the function vanilla carry over to the function Asiangeo. Greeks can also be derived from this relation.

Let us now consider the case where averaging starts after T_0, i.e. the payoff is changed to

$$[\phi(A(t, T) - K)]^+, \tag{283}$$

$$A(t, T) \stackrel{\Delta}{=} \exp\left\{ \frac{1}{T - t} \int_t^T \log S(u)\, du \right\}, t \in [0, T]. \tag{284}$$

Then the value becomes

$$\begin{aligned}
&\text{Asiangeowindow}(S_0, K, T, t, \sigma, r_d, r_f, \phi) \\
&= H \text{ Vanilla}\left(S_0, \frac{K}{H}, \frac{\Sigma\sigma}{(T-t)\sqrt{T}}, r_d, r_f, T, \phi\right), \qquad (285)\\
&H \stackrel{\Delta}{=} \exp\left\{-\frac{\sigma\theta_-}{2}(T-t) - \frac{\sigma^2}{2}(t - \frac{\Sigma}{T-t})\right\}. \qquad (286)
\end{aligned}$$

Derivation of the Value Function First we consider the call without history ($s = 0$). We rewrite the geometric average as

$$\begin{aligned}
A(0,T) &= \exp\left\{\frac{1}{T}\int_0^T \log S(u)\, du\right\} \\
&= S_0 \exp\left\{\frac{\sigma}{2}\theta_- T + \frac{\sigma}{T}\int_0^T W(u)\, du\right\} \qquad (287)
\end{aligned}$$

and compute the value function as

$$\begin{aligned}
&\text{Asiangeo}(S_0, K, T, 0, \sigma, r_d, r_f, \phi) \\
&= e^{-r_d T}\mathbb{E}[(A(0,T) - K)^+] \\
&= e^{-r_d T}\int_{-\infty}^{+\infty}\left(S_0 \exp\left\{\frac{\sigma}{2}\theta_- T + \sigma\sqrt{\frac{T}{3}}x\right\} - K\right)^+ n(x)\, dx \\
&= S_0 e^{-r_f T} e^{-\frac{T}{2}(r_d - r_f + \frac{\sigma^2}{6})}\mathcal{N}\left(\frac{\ln\frac{S_0}{K} + \frac{\sigma}{2}\theta_- T}{\sigma\sqrt{\frac{T}{3}}} + \sigma\sqrt{\frac{T}{3}}\right) \\
&\quad - K e^{-r_d T}\mathcal{N}\left(\frac{\ln\frac{S_0}{K} + \frac{\sigma}{2}\theta_- T}{\sigma\sqrt{\frac{T}{3}}}\right), \qquad (288)
\end{aligned}$$

which leads to the desired result. The analysis for the put option is similar. We obtain

$$\text{Asiangeo}(S_0, K, \sigma) = \sqrt{\frac{S_0}{K_+}}\text{vanilla}\left(S_0, K\sqrt{\frac{K_+}{S_0}}, \frac{\sigma}{\sqrt{3}}\right), \qquad (289)$$

where K_+ is the delta-neutral straddle strike as in Equation (43); however, in the case of the Asian option calculated with $\frac{\sigma}{\sqrt{3}}$ instead of σ. For $s > 0$ (real history) note that

$$A(-s,T) = A(-s,0)^{1-\alpha}A(0,T)^{\alpha}. \qquad (290)$$

The first factor of this product is non-random at time 0, hence the value of a call with history is given by

$$\begin{aligned}
&\text{Asiangeo}(S_0, K, T, s, \sigma, r_d, r_f, \phi) \qquad (291)\\
&= e^{-r_d T}\mathbb{E}[(A(-s,T) - K)^+]\\
&= e^{-r_d T}A(-s,0)^{1-\alpha}\mathbb{E}\left[\left(A(0,T)^\alpha - \frac{K}{A(-s,0)^{1-\alpha}}\right)^+\right]\\
&= e^{-r_d T}\int_{-\infty}^{+\infty}\left(S_0^\alpha \exp\left\{\frac{\alpha\sigma}{2}\theta_T + \alpha\sigma\sqrt{\frac{T}{3}}x\right\} - \frac{K}{A(-s,0)^{1-\alpha}}\right)^+ n(x)\,dx.
\end{aligned}$$

It is now an easy exercise to complete this calculation.

Arithmetic Average Options Since the distribution of the arithmetic average of log-normally distributed random variables is not normal, a closed form solution for the frequently used arithmetic average price options is not immediately available. Some of the approaches to solve this valuation task are

1. Numerical approaches, e.g. Monte Carlo simulations work well, as one can take the geometric Asian option as a highly correlated control variate. Taking a PDE approach is equally fast as Večeř has shown how to reduce the valuation problem to a PDE in one dimension in [133].
2. Modifications of the geometric average approach.
3. Approximations of the density function for the arithmetic average – see [88] on p. 430.

For instance, Turnbull and Wakeman (see [131]) develop an approximation of the density function by defining an alternative distribution for the arithmetic average with moments that match the moments of the true distribution, similar to that in Section 1.9.2. One can also match the cumulants up to fourth order: mean, variance, skew, and kurtosis. The adjusted mean and variance are finally plugged into the general Black-Scholes formula. Lévy states in [88] that considering only the first two moments delivers acceptable results for typical ranges of volatility and simultaneously reduces the complexity of the Turnbull and Wakeman approach. Hakala and Perissé show in [65] how to include higher moments. We apply a Monte Carlo simulation of price paths to value arithmetic average price options. To improve the quality of the results, we take geometric average options with similar specifications as control variate – see [86], p. 124. For variance reduction techniques see [58], pp. 414–418. For further suggestions on the implementation of pricing models see, for example, [29], pp. 118–123. We show in Table 1.25 that the results are close to the analytical approximations provided by Turnbull and Wakeman as well as Lévy.

Sensitivity Analysis We analyze now the sensitivities of the values with respect to various input parameters and compare them with vanilla options. Throughout we will use the parameters $K = 1.2000$, $S_0 = 1.2000$, $r_d = 3\%$, $r_f = 2.5\%$, $\sigma = 10\%$, $T - T_0 = 3$ months

TABLE 1.25 Values of average options. Input parameters are $K = 1.2000$, $S_0 = 1.2000$, $\sigma = 20\%$, $r_d = 3\%$, $r_f = 2.5\%$, $T - T_0 = 90$ days = 90/365 years, 90 observations (implying a time step of 0.002739726 years), 10,000 price paths in the Monte Carlo simulation. The arithmetic average options are average price options. All values are in domestic pips.

Method	Ar. call	Ar. put	Geo. price call	Geo. price put	Geo. strike call	Geo. strike put
Analytical	—	—	271.19	273.63	295.21	248.19
Monte Carlo	295.92	251.95	290.53	256.44	295.38	244.62
With control variate	276.57	269.14	—	—	—	—
Turnbull/Wakeman	276.36	269.02	—	—	—	—
Lévy	276.36	269.02	—	—	—	—

(91 days). The similarity of vanilla and average options, and the effects from averaging prices, which already dominated the derivation of the value formula, are reflected in the *Greeks* as well. Both option types react in the same direction to parameter changes and differ only in the quantity of the option value change. This holds especially for delta, gamma, and vega. These sensitivities, which are related to the underlying, represent best the properties of average options, i.e. initially, the option is very sensitive to price changes in the underlying. Delta, gamma, and vega have accordingly high values. With decreasing time to maturity, the impact of single prices on the final payoff diminishes, delta stabilizes, and gamma approaches zero, see [101], pp. 63–64. Figure 1.45 illustrates the similarity between vanilla and average price options with respect to delta and gamma.

For the same level of volatility in the underlying, average options have a lower vega compared with vanilla options because fluctuations of the underlying price are smoothed by the average. Note that the lower the volatility, the smaller the value difference between average and vanilla options, see Figure 1.46.

Since single prices – especially at maturity – influence the payoff of average options less significantly than for vanilla options, time, i.e. the chance of a finally favorable performance, plays a less important role in determining the value of average options, leading to a lower theta. The interest rate sensitivity rho of average options is smaller than for vanilla options.

Risk Management With the sensitivity analysis in mind, the question arises as to how the writer of an average option should deal with the risks of a short position.

Dynamic Hedging For a call position, for instance, one way is hedging with an investment in the underlying that is funded by borrowing. The delta of the option suggests how many units of the foreign currency have to be bought. Since delta changes over time, as in a vanilla option, the amount invested in the underlying has to be adjusted frequently. From the risk analysis it can be inferred that average options are easier to hedge than vanilla options, in particular the delta of average options stabilizes over time. Accordingly, the scope of required re-balancing of the hedge and the related transaction cost decrease over time. The costs of the hedge include interest payments as well as commissions and bid-ask spreads due at every re-balancing transaction. See [128] and

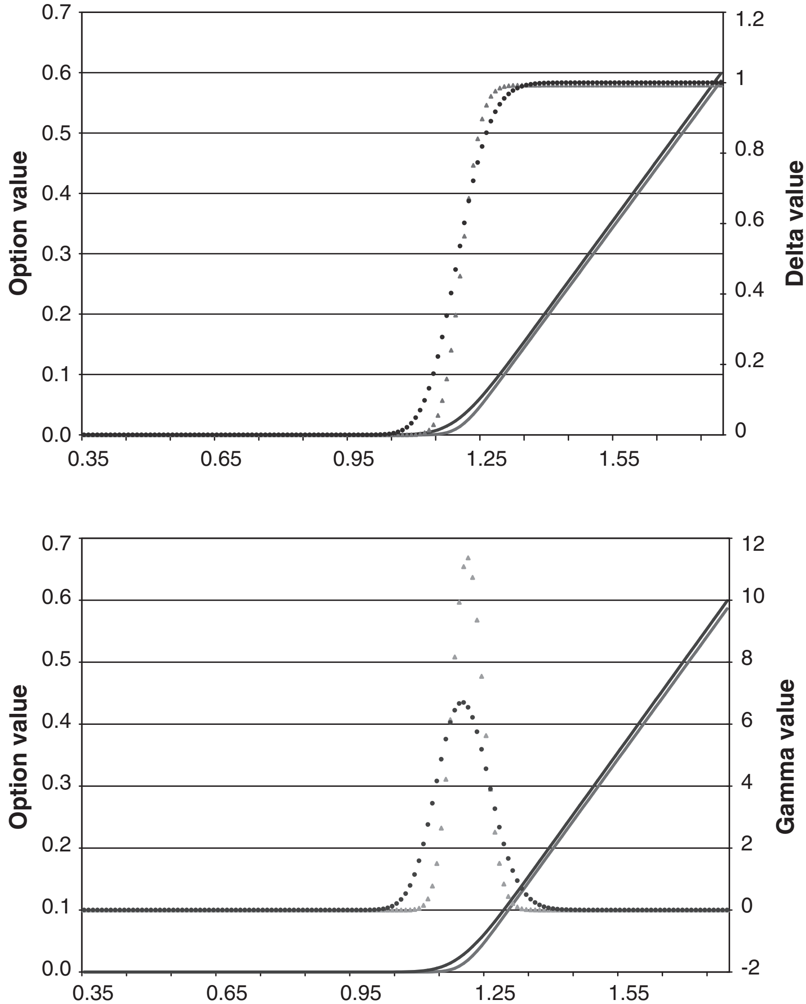

FIGURE 1.45 Above: option values and delta depending on the underlying price; Asian value is lower than vanilla value, Asian delta is smoother than vanilla delta; below: option values and gamma depending on the underlying price; Asian value is lower than vanilla value, Asian gamma is lower than vanilla gamma.

[101] for empirical analysis on the cost of dynamic and static hedging. Dynamic hedging neutralizes the delta exposure inherent in the option position. The volatility exposure can be hedged with vanilla contracts, typically ATM straddles.

Static Replication Alternatively, a static replication involving vanilla options can be set up. The position remains generally unchanged until maturity of the average option. Vanilla options are traded in liquid markets at relatively small bid-ask spreads. Furthermore, not only the delta risk but also the gamma and volatility exposure can be reduced

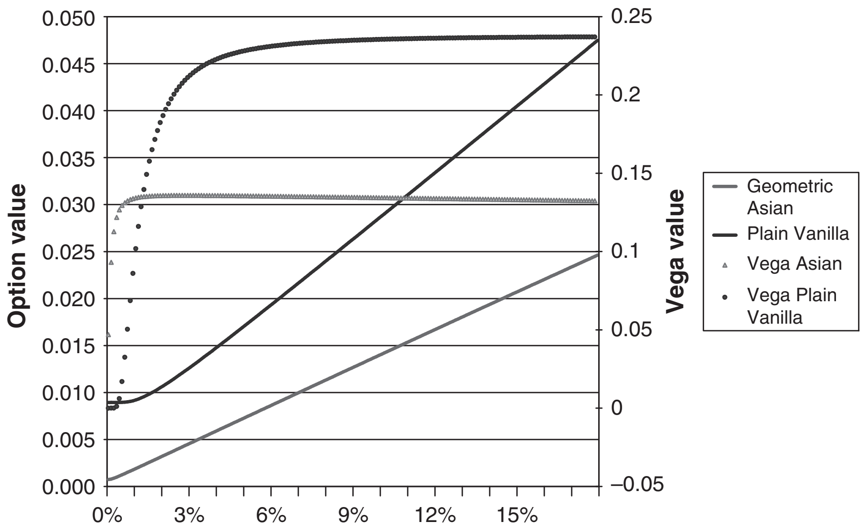

FIGURE 1.46 Option values and vega depending on volatility for at-the-money options.

with an option based replication. Static replication with vanilla options has therefore become common market practice, see [89]. For instance, Lévy suggests in [89] as a rule of thumb choosing a vanilla call with a similar strike as a short average price call and an expiration that is one-third of the averaging period of the exotic, based on the appearance of the factor $\frac{T}{3}$ in Equation (288). As the graph below of Figure 1.47 shows, the sensitivities of the short average price call are at their highest levels in the first third of the averaging period. Hedging with options only during this most critical time period already significantly reduces the sensitivity of the position to underlying price changes. Simultaneously, choosing vanilla calls with shorter maturity saves hedging costs. Nevertheless, this approach leaves the option writer with an open position for the remaining time to maturity unless she decides to build up a new replicating portfolio (semi-static replication). Since the stabilized delta in the later lifetime of the average option reduces the re-balancing effort, a dynamic hedge could be an alternative to a renewed replication with vanilla options.

1.7.7 Lookback Options

This section is joint work with Silvia Baumann, Marion Linck, Michael Mohr, and Michael Seeberg.

Lookback options are, like Asian options, path dependent. At expiration the holder of the option can "look back" over the lifetime of the option and exercise based upon the optimal underlying value (extremum) achieved during that period. Thus, lookback options (like Asians) avoid the problem of European options that the underlying performed favorably throughout most of the option's lifetime but

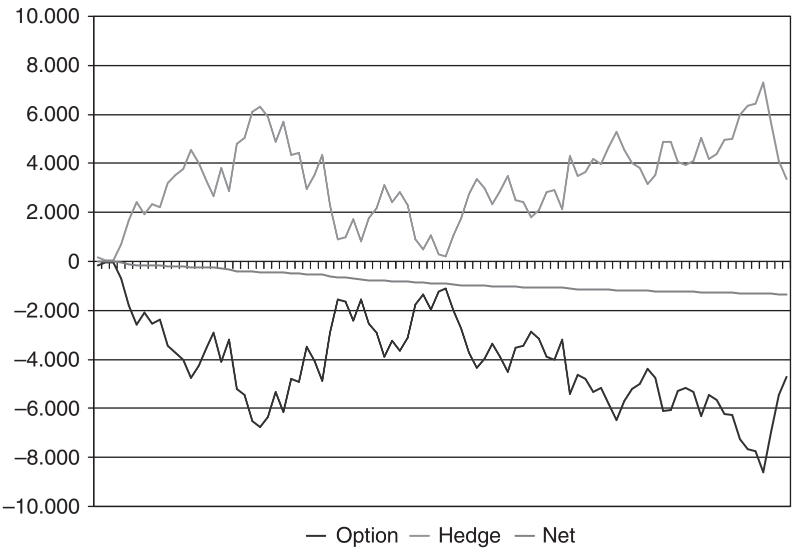

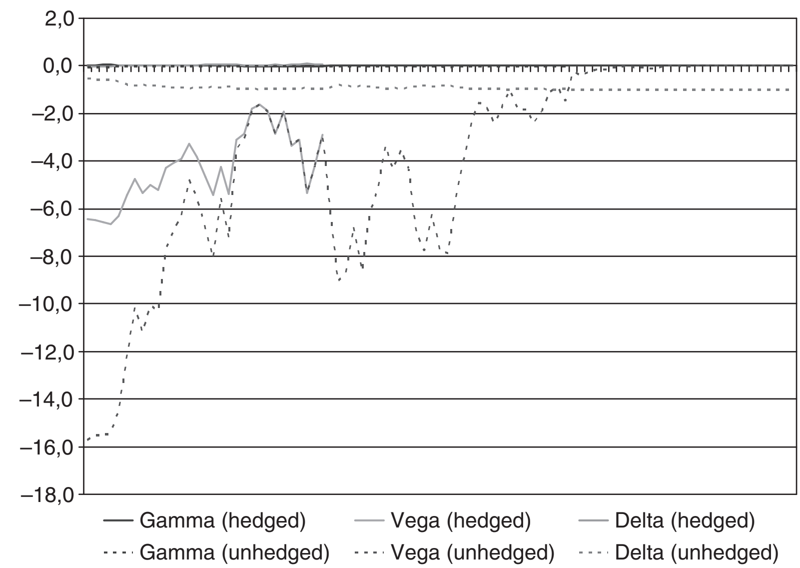

FIGURE 1.47 Graph above: dynamic hedging: performance of option position and hedge portfolio; graph below: static hedging: comparison of hedged and unhedged "Greek" exposure. For both, sample prices were generated randomly.

moves into a non-favorable direction towards maturity. Moreover (unlike American options), lookback options optimize the market timing because the investor gets – by definition – the most favorable underlying price. As summarized in Table 1.26 lookback options can be structured in two different types with the extremum representing either

TABLE 1.26 Types of lookback options. The contract parameters T and X are the time to maturity and the strike price respectively, and S_T denotes the spot price at expiration time. Fixed strike lookback options are also called hindsight options.

Payoff	Lookback type	Parameter
$M_T - S_T$	floating strike put	$\phi = -1, \overline{\eta} = -1$
$S_T - m_T$	floating strike call	$\phi = +1, \overline{\eta} = +1$
$(M_T - X)^+$	fixed strike call	$\phi = +1, \overline{\eta} = -1$
$(X - m_T)^+$	fixed strike put	$\phi = -1, \overline{\eta} = +1$

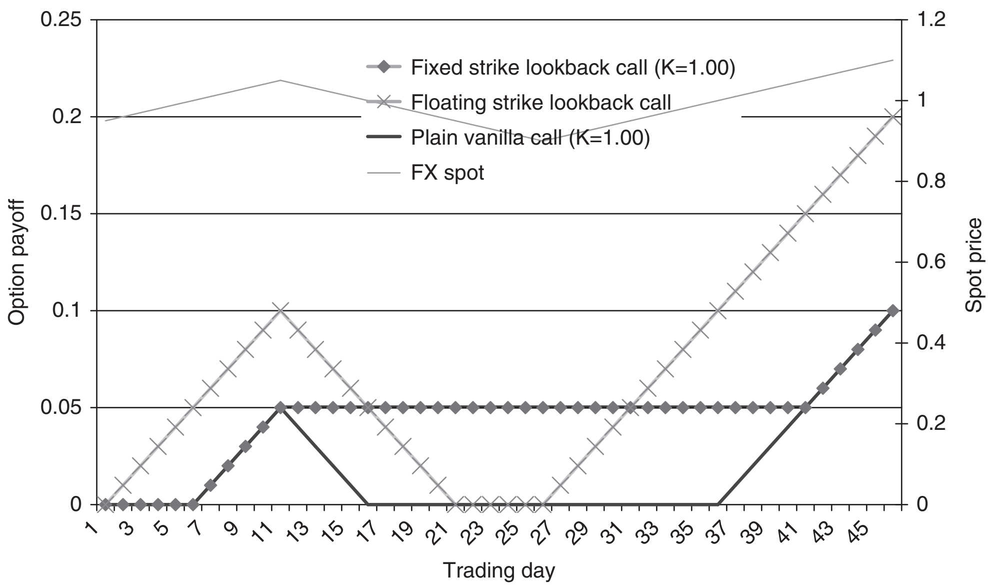

FIGURE 1.48 Payoff profile of lookback calls (sample underlying price path, 20 trading days).

the strike price or the underlying value. Figure 1.48 shows the development of the payoff of lookback options depending on a sample price path. In detail we define

$$M_{t,T} \stackrel{\Delta}{=} \max_{t \le u \le T} S(u), \tag{292}$$

$$M_T \stackrel{\Delta}{=} M_{0,T}, \tag{293}$$

$$m_{t,T} \stackrel{\Delta}{=} \min_{t \le u \le T} S(u), \tag{294}$$

$$m_T \stackrel{\Delta}{=} m_{0,T}. \tag{295}$$

Variations of lookback options include *partial lookback options*, where the monitoring period for the underlying is shorter than the lifetime of the option. Conze and Viswanathan [31] present further variations, such as *limited risk* and *American lookback options*. Lookback options are not traded much with corporate treasurers, as they are too expensive; they are mainly used by speculators, see [27]. An often cited strategy is building *lookback straddles* paying

$$M_{t,T} - m_{t,T}, \tag{296}$$

(also called *range* or *hi-lo option*), a combination of lookback put(s) and call(s) that guarantees a payoff equal to the observed range of the underlying asset. In theory, Garman pointed out in [53] that lookback options can also add value for risk managers because floating (fixed) strike lookback options are good means to solve the timing problem of market entries (exits), see [75]. For instance, a minimum strike call is suitable to avoid missing the best exchange rate in currency linked security issues. However, this right is very expensive. Since one buys a guarantee for the best possible exchange rate observed in a time interval, lookback options are generally way too expensive and hardly ever trade. Exceptions are performance notes, where lookback and average features are mixed, e.g. performance notes paying say 50% of the best of 36 monthly average gold price returns.

Valuation As in the case of Asian options, closed form solutions exist only for specific products – in this case basically for any lookback option with continuously monitored underlying value. We consider the example of the floating strike lookback call. Again, the value of the option is given by

$$\begin{aligned} v(0,S_0) &= \mathbb{E}\left[e^{-r_d T}(S_T - m_T)\right] \\ &= S_0 e^{-r_f T} - e^{-r_d T}\mathbb{E}\left[m_T\right]. \end{aligned} \tag{297}$$

In the standard Black-Scholes model (1), the value can be derived using the reflection principle and results in

$$\begin{aligned} v(t,x) = \phi \Bigg\{ & x e^{-r_f\tau}\mathcal{N}(\phi b_1) - K e^{-r_d\tau}\mathcal{N}(\phi b_2) + \frac{1-\eta}{2}\phi e^{-r_d\tau}[\phi(R-X)]^+ \\ & + \eta x e^{-r_d\tau}\frac{1}{h}\left[\left(\frac{x}{K}\right)^{-h}\mathcal{N}(-\eta\phi(b_1 - h\sigma\sqrt{\tau})) - e^{(r_d-r_f)\tau}\mathcal{N}(-\eta\phi b_1)\right]\Bigg\}. \end{aligned} \tag{298}$$

This value function has a removable discontinuity at $h = 0$ where it turns out to be

$$\begin{aligned} v(t,x) = \phi \Bigg\{ & x e^{-r_f\tau}\mathcal{N}(\phi b_1) - K e^{-r_d\tau}\mathcal{N}(\phi b_2) + \frac{1-\eta}{2}\phi e^{-r_d\tau}[\phi(R-X)]^+ \\ & + \eta x e^{-r_d\tau}\sigma\sqrt{\tau}\left[-b_1\mathcal{N}(-\eta\phi b_1) + \eta\phi n(b_1)\right]\Bigg\}. \end{aligned} \tag{299}$$

The abbreviations we use are

$$t\text{: running time (in years)}, \tag{300}$$

$$x \stackrel{\Delta}{=} S_t\text{: known spot at time of evaluation}, \tag{301}$$

$$\tau \stackrel{\Delta}{=} T - t\text{: time to expiration (in years)}, \tag{302}$$

$$n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2}, \tag{303}$$

$$\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt, \tag{304}$$

$$h \stackrel{\Delta}{=} \frac{2(r_d - r_f)}{\sigma^2}, \tag{305}$$

$$K \stackrel{\Delta}{=} \begin{cases} R & \text{floating strike lookback} \quad (X \leq 0) \\ \overline{\eta}\min(\overline{\eta}X, \overline{\eta}R) & \text{fixed strike lookback} \quad (X > 0) \end{cases}, \tag{306}$$

$$\eta \stackrel{\Delta}{=} \begin{cases} +1 & \text{floating strike lookback} \quad (X \leq 0) \\ -1 & \text{fixed strike lookback} \quad (X > 0) \end{cases}, \tag{307}$$

$$b_1 \stackrel{\Delta}{=} \frac{\ln \frac{x}{K} + (r_d - r_f + \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}, \tag{308}$$

$$b_2 \stackrel{\Delta}{=} b_1 - \sigma\sqrt{\tau}. \tag{309}$$

Note that this formula basically consists of that for a vanilla call (first two terms) plus another term. Conze and Viswanathan also show closed form solutions for fixed strike lookback options and the variations mentioned above in [31]. Heynen and Kat develop equations for *partial fixed and floating strike lookback options* in [74]. For those preferring the PDE approach of deriving formulas, we refer to [59]. For most practical matters, where we have to deal with fixings and lookback features in combination with averaging, the only reasonable valuation technique is Monte Carlo simulation.

Example We list some sample results in Table 1.27.

Sensitivity Analysis

delta.

$$v_x(t,x) = \phi \left\{ e^{-r_f\tau}\mathcal{N}(\phi b_1) + \eta e^{-r_d\tau}\frac{1}{h} \cdot \left[\left(\frac{x}{K}\right)^{-h} \mathcal{N}(-\eta\phi(b_1 - h\sigma\sqrt{\tau}))(1-h) - e^{(r_d - r_f)\tau}\mathcal{N}(-\eta\phi b_1) \right] \right\} \tag{310}$$

TABLE 1.27 Sample valuation results for lookback options. For the input data we used spot $S_0 = 0.8900$, $r_d = 3\%$, $r_f = 6\%$, $\sigma = 10\%$, $\tau = \frac{1}{12}$, running min = 0.9500, running max = 0.9900, $m = 22$. We find the analytic results in the continuous case in agreement with the ones published in Haug [71]. We can also reproduce the numerical results for the discretely sampled floating strike lookback put contained in Nahum [98].

Payoff	Analytic model	Continuous
$M_T - S_T$	0.0231	0.0255
$S_T - m_T$	0.0310	0.0320
$(M_T - 0.99)^+$	0.0107	0.0131
$(0.97 - m_T)^+$	0.0235	0.0246

At $h = 0$ this simplifies to

$$v_x(t,x) = \phi \left\{ e^{-r_f\tau} \mathcal{N}(\phi b_1) + \eta e^{-r_d\tau} \left[\sigma\sqrt{\tau} \left[-b_1 \mathcal{N}(-\eta\phi b_1) + \eta\phi n(b_1) \right] - \mathcal{N}(-\eta\phi b_1) \right] \right\} \tag{311}$$

gamma.

$$v_{xx}(t,x) = \frac{2e^{-r_f\tau}}{x\sigma\sqrt{\tau}} n(b_1) - \phi\eta e^{-r_d\tau} \frac{1-h}{x} \mathcal{N}(-\phi\eta(b_1 - h\sigma\sqrt{\tau})) \tag{312}$$

theta. We can use the Black-Scholes partial differential equation to obtain theta from value, delta, and gamma.

vega.

$$v_\sigma(t,x) = \phi\eta x e^{-r_d\tau} \frac{2}{\sigma} \left[\left(\frac{x}{K}\right)^{-h} \mathcal{N}(-\eta\phi(b_1 - h\sigma\sqrt{\tau})) \left(\frac{1}{h} + \ln\frac{x}{K}\right) - e^{(r_d - r_f)\tau} \frac{1}{h} \mathcal{N}(-\eta\phi b_1) \right] \tag{313}$$

At $h = 0$ this simplifies to

$$v(t,x) = \phi\eta x e^{-r_d\tau} \sqrt{\tau} \left[-\sigma\sqrt{\tau} b_1 \mathcal{N}(-\eta\phi b_1) + 2\eta\phi n(b_1) \right] \tag{314}$$

Discrete Sampling In practice, one cannot take the average over a continuum of exchange rates. The standard is to specify a *fixing calendar* and take only a finite number of fixings into account. Suppose there are m equidistant sample points left until expiration at which we evaluate the extremum. In this case the value can be determined by an approximation described by Broadie *et al.* [20]. We set

$$\beta_1 = 0.5826 = -\zeta(1/2)/\sqrt{2\pi}, \tag{315}$$

$$a = e^{\phi\beta_1\sigma\sqrt{\tau/m}}, \tag{316}$$

and obtain for fixed strike lookback options

$$v(t, x, r_d, r_f, \sigma, R, X, \phi, \overline{\eta}, m) = v(t, x, r_d, r_f, \sigma, aR, aX, \phi, \overline{\eta})/a, \tag{317}$$

and for floating strike lookback options

$$v(t, x, r_d, r_f, \sigma, R, X, \phi, \overline{\eta}, m) = av(t, x, r_d, r_f, \sigma, R/a, X, \phi, \overline{\eta}) - \phi(a-1)xe^{-r_f\tau}. \tag{318}$$

One interesting observation is that when the options move deep in-the-money and have the same strike price, lookback options and vanilla options have the same value, except for extreme risk parameter inputs. This can be explained recalling that a floating strike lookback option has an exercise probability of 1 and buys (sells) at the minimum (maximum). When the strike price of a vanilla option equals the extremum of the exotic and is deep in-the-money, the holder of the option will also buy (sell) at the extremum with a probability very close to 1. Moreover, recall that the floating strike lookback option consists of a vanilla option and an additional term. Garman names this term a *strike-bonus option*, see [53]. It can be considered as an option that has an increased payoff whenever a new extremum is reached. When the underlying price moves very far away from the current extremum, the strike-bonus option has almost zero value.

The structure of the Greeks delta, rho, theta, and vega is similar for lookback and vanilla calls. Nonetheless, the intensity of these sensitivities against changes differs, see Figure 1.49.

Close to or at the money, lookback calls have a significantly lower *delta* than their vanilla counterparts. The reason is that the strike-bonus option in the lookback call has a negative delta when the underlying value is close to the current extremum and a delta next to zero when it is far in-the-money. Intuitively, the lower lookback delta is explained by the fact that the closer the current exchange rate is to the extremum, the more likely it is that the payoff of the lookback option remains unchanged, which is different for vanilla options where the payoff changes with every spot movement. Note that whenever a new extremum is achieved, the payoff for a lookback option equals zero and remains unchanged until the underlying value moves into the adverse direction.

Floating strike lookback options have a lower *rho* than vanilla options (with equal strikes at the time of observation), which can be explained by the fact that the option holder needs to pay more up-front and thus has a lower principal profiting from favorable interest rate movements. As a rule of thumb, a floating strike lookback option is worth twice as much as a vanilla option. The longer the time to maturity, the more intensively floating strike lookback options react compared with vanilla options.

The higher *theta* for lookback options reflects the fact that the optimal value achieved to date is “locked in” and the longer the time to maturity, the higher the chance to lock in an even better extremum.

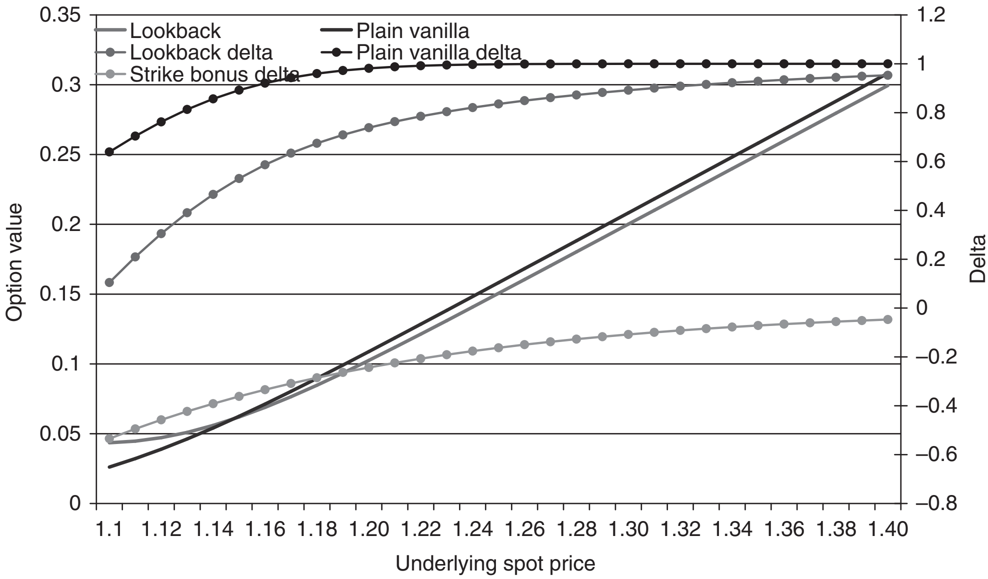

FIGURE 1.49 Vanilla and lookback call (left-hand scale) with deltas (right-hand scale) using $\min_t S_t = 1.00$. The lookback delta equals the sum of the delta of a vanilla option plus the delta of a strike bonus option.

Regarding *vega*, lookback options show a stronger reaction than regular options. The higher the volatility of the underlying, the higher the probability of reaching a new extremum. Moreover, having "locked in" this new extremum the option value can benefit even more from the higher chance of adverse price movements.

As pointed out by Taleb in [123], one particularly interesting risk parameter is the lookback *gamma* since it is *one-sided*, while the vanilla gamma changes symmetrically for up and down movements of the underlying, see Figure 1.50. A lookback option always has its maximum gamma at the extremum which can move over time. Vanilla options, however, have their maximum gamma at the strike price. The *lookback gamma asymmetry* indicates that gamma risk cannot be consistently (statically) hedged with vanilla options. The fact that gamma is considerably higher for lookback options implies that a frequent re-balancing of the hedging portfolio and hence high transaction costs are likely, see [32].

Semi-Static Replication Due to the maximum (minimum) function that allows the strike price to change there exists no buy-and-hold static replication for floating strike lookback options. Instead, a *semi-static rollover strategy* can be applied, see [53]. As can be read from the value in Equation (298), we can replicate parts of a lookback with a vanilla option. Whenever the maximum (minimum) changes, the writer of the option buys a new put (call) struck at the current market price and sells the old put (call). However, this does not work without costs. While the new put (call) is at-the-money, the old put (call) is out-of-the-money at the time of the sale and hence returns less money

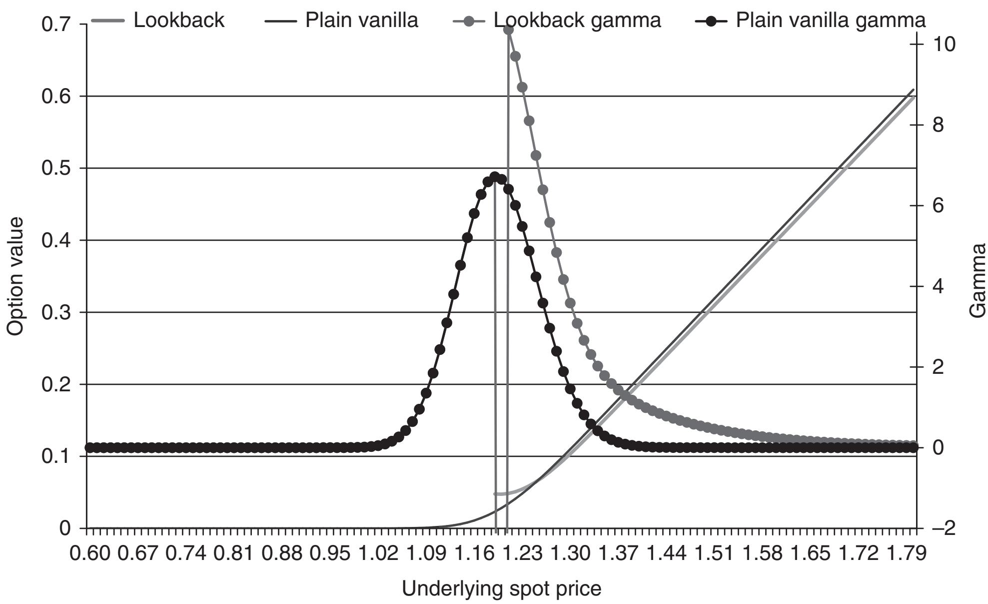

FIGURE 1.50 Value (left-hand scale) and gamma (right-hand scale) of an at-the-money floating strike lookback and a vanilla call.

than the amount necessary to purchase the new option. We encounter vanilla option bid-ask spreads, and smile risk. The strike-bonus option returns exactly the money that is needed for the rollover (in theory, and assuming no bid-offer spreads). This approach, however, is rather theoretical, since strike-bonus options are hardly available in the market.

In practice, floating strike lookback options are usually approximated by a straddle, see Section 1.6.3. Cunningham and Karumanchi also explain a dynamic replication strategy for fixed strike lookback options in [32]. The straddle to use is a combination of a vanilla put and a vanilla call, which have a term to maturity equal to that of the lookback option to be replicated, and a strike equal to the maximum (minimum) achieved by the underlying. At maturity T, the call (put) of this straddle becomes worthless since the strike is below (above) the terminal stock price S_T. The remaining put (call) exactly satisfies the obligation of the lookback option (see Figure 1.51). Over the lifetime of the option, the strike price of the straddle needs to be adapted if the current exchange rate S_t rises above (falls below) the current maximum (minimum). Regarding the intrinsic value, the holder of the hedging portfolio will not lose money since for instance the intrinsic value lost by the put will be exactly gained by the call. However, the deltas of the two options differ, not only in their sign. In addition, attempting to create a replicating portfolio with zero delta, the trader has to buy a certain number of puts per one call. Figure 1.51 shows that for the latter two reasons this is not a self-financing replicating strategy. Note that the strategy would not be self-financing even if the straddle was not adapted for a zero delta of the position. Specifying a *re-hedge threshold* and a maximum number of trades per period can help to balance the risk taken with

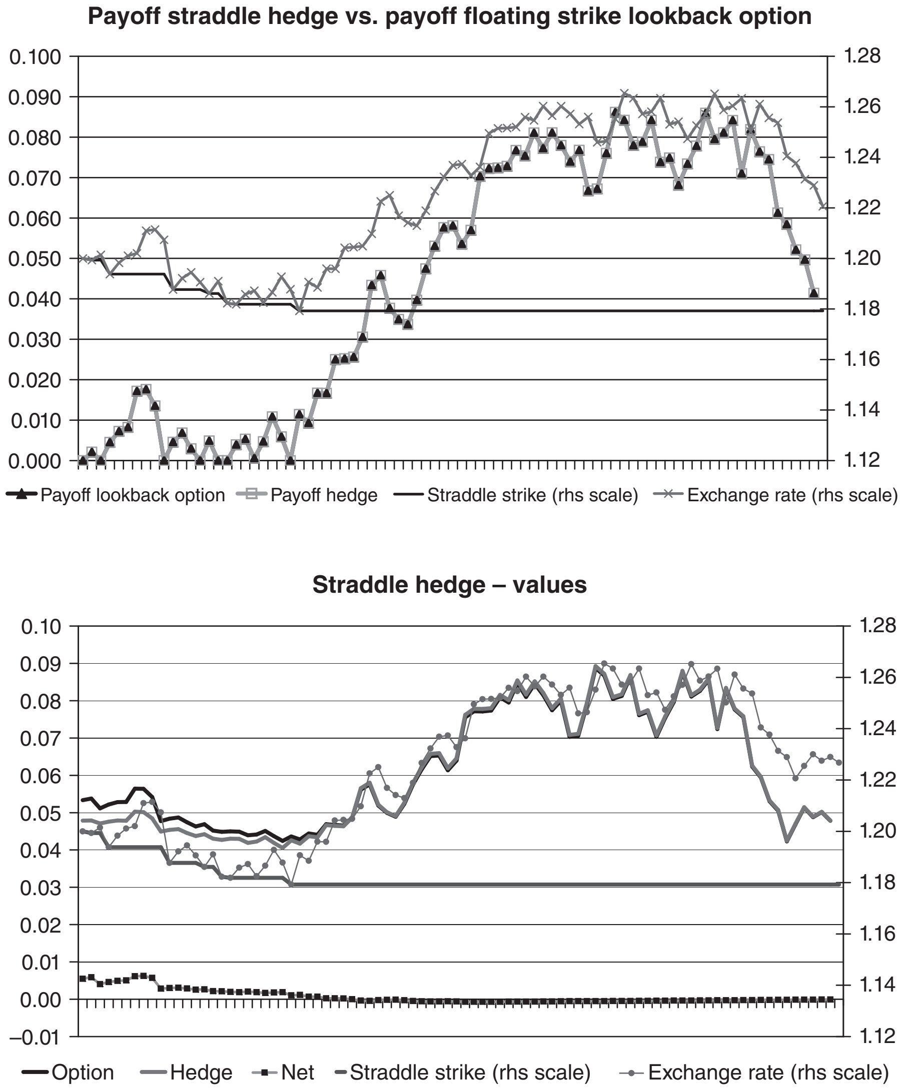

FIGURE 1.51 Comparison of the payoffs of a floating strike lookback option and a vanilla straddle (graph above) and the values of the positions (graph below) for a random time path (exchange rate and straddle strike on the rhs scale).

transaction and administrative costs. We refer to this replication as semi-static since the basic idea of the replication is that of a static one: initialize the hedge and wait until maturity. However, due to the changes of the extrema, the static replication has to be adapted – a characteristic which is usually associated with dynamic hedging. Apart from this *semi-static* replication, a *dynamic hedge* using spot and money market is also possible and more commonly applied. Due to the risk parameters, especially gamma and vega, which are difficult to hedge, the hedge appears to deviate considerably in value relative to the option over time.

1.7.8 Forward Start, Ratchet, and Cliquet Options

A forward start vanilla option is just like a vanilla option, except that the strike is set on a future date t. It pays off

$$[\phi(S_T - K)]^+, \tag{319}$$

where K denotes the strike and ϕ takes the values $+1$ for a call and -1 for a put. The strike K is set as αS_t at time $t \in [0, T]$. Very commonly α is set to 1.

Advantages

- Protection against spot market movement and against increasing volatility
- Buyer can lock in current volatility level
- Spot risk easy to hedge

Disadvantages

- Protection level not known in advance

The Value of Forward Start Options Using the abbreviations

- x for the current spot price of the underlying,
- $\tau \stackrel{\Delta}{=} T - t$,
- $F_s \stackrel{\Delta}{=} \mathbb{E}[S_s | S_0] = S_0 e^{(r_d - r_f)s}$ for the outright forward of the underlying,
- $\theta_\pm \stackrel{\Delta}{=} \frac{r_d - r_f}{\sigma} \pm \frac{\sigma}{2}$,
- $d_\pm \stackrel{\Delta}{=} \frac{\ln \frac{x}{K} + \sigma\theta_\pm \tau}{\sigma\sqrt{\tau}} = \frac{\ln \frac{f}{K} \pm \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}$,
- $d_\pm^\alpha \stackrel{\Delta}{=} \frac{-\ln \alpha + \sigma\theta_\pm \tau}{\sigma\sqrt{\tau}}$,
- $n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2} = n(-t)$,
- $\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt = 1 - \mathcal{N}(-x)$,

we recall the value of a vanilla option in Equation (7),

$$v(x, K, T, t, \sigma, r_d, r_f, \phi) = \phi e^{-r_d \tau}[f\mathcal{N}(\phi d_+) - K\mathcal{N}(\phi d_-)]. \tag{320}$$

For the value of a forward start vanilla option in a constant-coefficient geometric Brownian motion model we obtain

$$\begin{aligned} v &= e^{-r_d t}\mathbb{E}v(S_t, K = \alpha S_t, T, t, \sigma, r_d, r_f, \phi) \qquad (321)\\ &= \phi e^{-r_d T}[F_T \mathcal{N}(\phi d_+^\alpha) - \alpha F_t \mathcal{N}(\phi d_-^\alpha)]. \end{aligned}$$

Noticeably, the value computation is easy here because the strike K is set as a *multiple* of the future spot. If we were to choose to set the strike as a constant *difference* of the future spot, the integration would not work in closed form, and we would have to use numerical integration.

The crucial pricing issue here is that one needs to know the volatility, which is the *forward volatility*, i.e. the volatility that will materialize at the future time t for a maturity $T - t$. It is not obvious from the market which proxy to take for this forward volatility. The standard is to use Equation (152).

Greeks The Greeks are the same as for vanilla options after time t, when the strike has been set. Before time t they are given by

(Spot) Delta.

$$\frac{\partial v}{\partial S_0} = \frac{v}{S_0} \tag{322}$$

Gamma.

$$\frac{\partial^2 v}{\partial x^2} = 0 \tag{323}$$

Theta.

$$\frac{\partial v}{\partial t} = r_f v \tag{324}$$

Vega.

$$\frac{\partial v}{\partial \sigma} = -\frac{e^{-r_d T}}{\sigma}[F_T n(d_+^\alpha)d_-^\alpha - \alpha F_t n(d_-^\alpha)d_+^\alpha] \tag{325}$$

Rho.

$$\frac{\partial v}{\partial r_d} = \phi e^{-r_d T}\alpha F_t(T-t)\mathcal{N}(\phi d_-^\alpha) \tag{326}$$

$$\frac{\partial v}{\partial r_f} = -Tv - \phi e^{-r_d T}\alpha F_t(T-t)\mathcal{N}(\phi d_-^\alpha) \tag{327}$$

Example We consider an example in Table 1.28.

Reasons for Trading Forward Start Options The key reason for trading a forward start is trading the forward volatility without any spot exposure. In quiet market phases with low volatility, buying a forward start is cheap. Keeping a long position will allow participation in rising volatility, independent of the spot level. In recent years, forward start options have become more popular, especially among institutional clients who hedge their volatility exposure. They use *forward volatility agreements (FVAs)*, see

TABLE 1.28 Value and Greeks of a forward start vanilla in USD on EUR/USD – spot of 0.9000, $\alpha = 99\%$, $\sigma = 12\%$, $r_d = 2\%$, $r_f = 3\%$, maturity $T = 186$ days, strike set at $t = 90$ days.

	Call	Put		Call	Put
Value	0.0251	0.0185			
Delta	0.0279	0.0206	vega	0.1793	0.1793
Gamma	0.0000	0.0000	rho_d	0.1217	−0.1052
Theta	0.0007	0.0005	rho_f	−0.1329	0.0950

Section 1.8.9, which is pretty much a re-branding of the good old forward start option or a combination of them as a forward start straddle. The key reason for sales is that for notes there is typically a subscription phase, where a product such as a dual currency investment is announced and investors can subscribe for it during the *subscription phase*. This phase can be for example four weeks. During this time notional is collected and the investor subscribes to a product whose strike will be fixed at the end of the subscription phase.

Variations Forward start options can be altered in all kind of ways: they can be of American style, they can come with a deferred delivery or deferred premium, they can have barriers or appear as a strip.

A strip of forward start options is generally called a *cliquet*.

A *ratchet* consists of a series of forward start options, where the strike for the next forward start option is set equal to the spot at maturity of the previous.

1.7.9 Power Options

This section is joint work with Silvia Baumann, Marion Linck, Michael Mohr, and Michael Seeberg.

For power options, the vanilla option payoff function $[\phi(S_T - K)]^+$ is adjusted by raising the entire function or parts of the function to the n-th power, see e.g. Zhang [146]. The result is a non-linear profile with the potential of a higher payoff at maturity with a greater leverage than standard options. Interestingly, many people don't know that numbers smaller than 1 tend to get smaller when squared. If the exponent n is exactly 1, the option is equal to a vanilla option. We distinguish between *asymmetric* and *symmetric* power options. Their payoffs in comparison with vanilla options are illustrated in Figure 1.52 and Figure 1.53.

Asymmetric Power Options With an asymmetric power option, the underlying S_T and strike K of a standard option payoff function are individually raised to the n-th power,

$$[\phi(S_T^n - K^n)]^+. \tag{328}$$

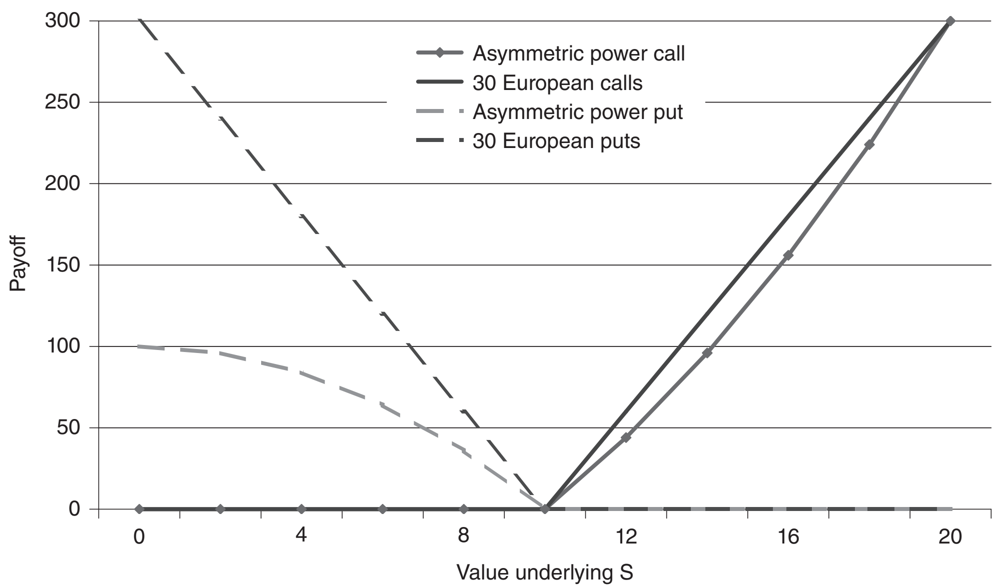

FIGURE 1.52 Payoff of asymmetric power options vs. vanilla options, using $K = 10$, $n = 2$.

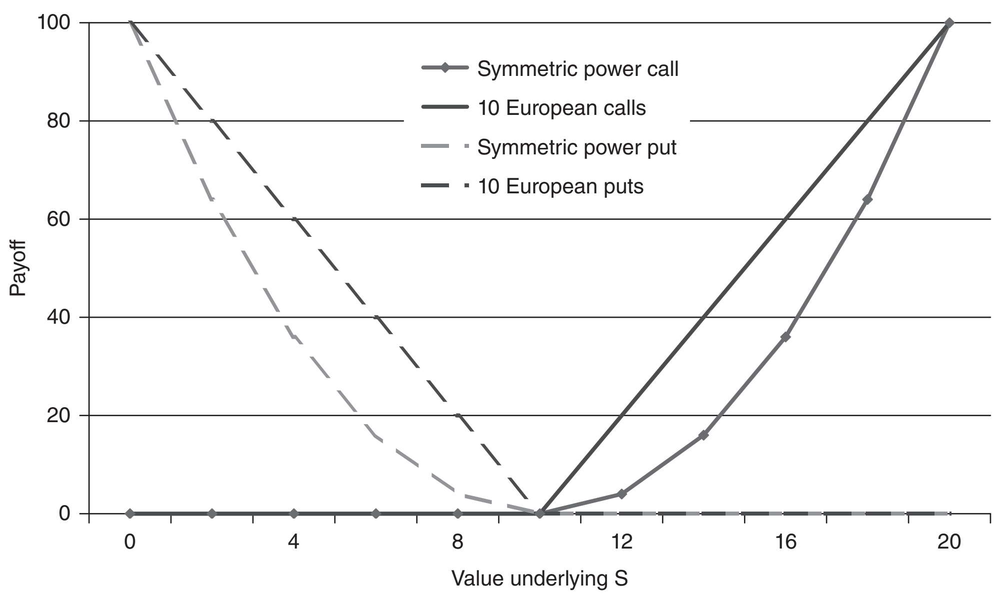

FIGURE 1.53 Payoff of symmetric power options vs. vanilla options, using $K = 10$, $n = 2$.

Figure 1.52 illustrates why this option type is called *asymmetric*. With increasing S_T, the convex call payoff grows exponentially. Given the limited and fixed profit potential of $K^2 = 10^2$, the concave put payoff decreases exponentially. It requires 30 vanilla options to replicate the call payoff if the underlying S_T moves to 20.

Symmetric Power Options In the symmetric type, the entire vanilla option payoff is raised to the n-th power,

$$[[\phi(S_T - K)]^+]^n, \tag{329}$$

see [129]. Figure 1.53 insinuates naming this option type *symmetric*, since put and call display the same payoff shape. Here, 10 vanilla options suffice to replicate the symmetric power option if the underlying S_T moves to 20.

Combining a symmetric power call and put as in Figure 1.53 leads to a symmetric power straddle, which pays

$$|S_T - K|^n. \tag{330}$$

Reasons for Trading Power Options Power options are often equipped with a payoff cap C to limit the short position risk as well as the option premium for the buyer. For example, the payments of the short position at $n = 3$ for $K = 10$ shoots to 2375(125) for the asymmetric (symmetric) power call if S_T moves to 15. Even with cap, the highly leveraged payoff motivates speculators to invest in the product that demands a considerably higher option premium than a vanilla option. Power options are mostly popular in the listed derivatives and retail market, due to their high leverage and due to their mere name. There exist more advanced power-like derivatives, see for example the multiplicity power option in Section 1.8.1. Furthermore, a *self-quanto* option comes up in retail markets if the vanilla payoff $\phi(S_T - K)^+$ is meant to pay in foreign currency, which means that it must pay

$$\phi(S_T - K)^+ S_T \tag{331}$$

in domestic currency, such that a division of the final payoff by the final spot price S_T will yield the vanilla payoff in foreign currency. This is a version of a symmetric power option, which is why some market participants refer to self-quanto options as power options. Quanto options are treated in Section 1.7.10. Besides the obvious reasons one can think of the following additional motives.

1. Hedging future levels of implied volatility. Vega, which is volatility risk, is extremely difficult to hedge as there is no directly observable measure available, see [106]. A power straddle is an effective instrument to do so as it preserves the volatility exposure better than a vanilla straddle when the price of the underlying moves significantly as shown in Section 1.7.9 on sensitivities to risk parameters.
2. Through their exponential, non-linear payoff, power options can hedge non-linear price risks. An example is an importer earning profits merely through a percentage mark-up on imported products. An exchange rate change will lead to a price change, which in turn may affect demand volumes. The importer faces a risk of non-linearly decreasing earnings, see [70].

3. With very large short positions in vanilla options, a re-balancing of a dynamic delta hedge of this short gamma position may require such massive buying (selling) of the underlying that this impacts the price of the underlying, which in turn requires further hedge adjustments and may "pin" the underlying to the strike price, see p. 37 in [129]. To *smooth this pin risk*, option sellers propose a *soft strike option* with a similarly smooth and continuous payoff curvature as power options. As we will show in the hedging analysis of this section, this payoff curvature can be effectively replicated using vanilla options with different strike prices. The *diversified* range of strikes then softens any effect of a move in the underlying price. For details on *soft strike options* see [129], p. 37 and [70], p. 51.

Valuation of the Asymmetric Power Option The value can be written as the discounted expected value of the payoff under the risk-neutral measure. Using the domestic discount factor $e^{-r_d T}$ yields

$$\text{asymmetric power option value } v_{aPC} = e^{-r_d T} \mathbb{E}\left[\phi(S_T^n - K^n) \mathbb{I}_{\{\phi S_T > \phi K\}}\right]. \tag{332}$$

As K is a constant, S_T is the only random variable which simplifies the equation to

$$v_{aPC} = \phi e^{-r_d T} \mathbb{E}\left[S_T^n \mathbb{I}_{\{\phi S_T > \phi K\}}\right] - \phi e^{-r_d T} K^n \mathbb{E}\left[\mathbb{I}_{\{\phi S_T > \phi K\}}\right]. \tag{333}$$

The expectation of an indicator function is just the probability that the event $\{S_T > K\}$ occurs. In the Black-Scholes model, S_T is log-normally distributed and evolves according to a geometric Brownian motion (1). Itô's Lemma implies that S_T^n is also a geometric Brownian motion following

$$dS_t^n = \left[n(r_d - r_f) + \tfrac{1}{2} n(n-1)\sigma^2\right] S_t^n \, dt + n\sigma S_t^n \, dW_t. \tag{334}$$

Solving the differential equation and calculating the expected value in Equation (333) leads to the desired closed form solution

$$v_{aPC} = \phi e^{-r_d T} \left[f^n e^{\frac{1}{2} n(n-1)\sigma^2 T} \mathcal{N}(\phi d_+^n) - K^n \mathcal{N}(\phi d_-)\right], \tag{335}$$

$$f \stackrel{\Delta}{=} S_0 e^{(r_d - r_f)T},$$

$$d_- \stackrel{\Delta}{=} \frac{\ln \frac{f}{K} - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}},$$

$$d_+^n \stackrel{\Delta}{=} \frac{\ln \frac{f}{K} + \left(n - \frac{1}{2}\sigma^2\right) T}{\sigma\sqrt{T}}.$$

Valuation of the Symmetric Power Option Due to the binomial term $(S_T - K)^n$ the general value formula derivation for the symmetric version looks more complicated. That is why a more intuitive approach is taken and the valuation logic is shown based on the asymmetric option discussed above. Taking the example of $n = 2$ the difference between asymmetric and symmetric call is

$$[S_T^2 - K^2] - [S_T^2 - 2S_TK + K^2] = 2K(S_T - K). \tag{336}$$

The symmetric version for $n = 2$ is thus exactly equal to the asymmetric power option minus $2K$ vanilla options. This way pricing and hedging the symmetric power option becomes a structuring exercise, see Figure 1.54. Tompkins and Zhang both discuss the more complicated derivation of the general formula for symmetric power options in [129] and [146]. Tompkins also presents a formula for symmetric power straddles for $n = 2$.

Sensitivity Analysis Looking at the *Greeks* of asymmetric power options compared to vanilla options, the exponential elements of power options are well reflected in the exposures. This is especially true for delta and gamma, as can be seen in Figure 1.55, but is also valid for theta and vega. The power option rhos are very similar to the vanilla version.

Contrary to the asymmetric power option, the symmetric power option sensitivities exhibit new features that cannot be found with vanilla options, namely extreme delta values and a gamma that resembles the plain vanilla delta. When combined in a straddle this creates a *constant gamma exposure*, see Figure 1.56.

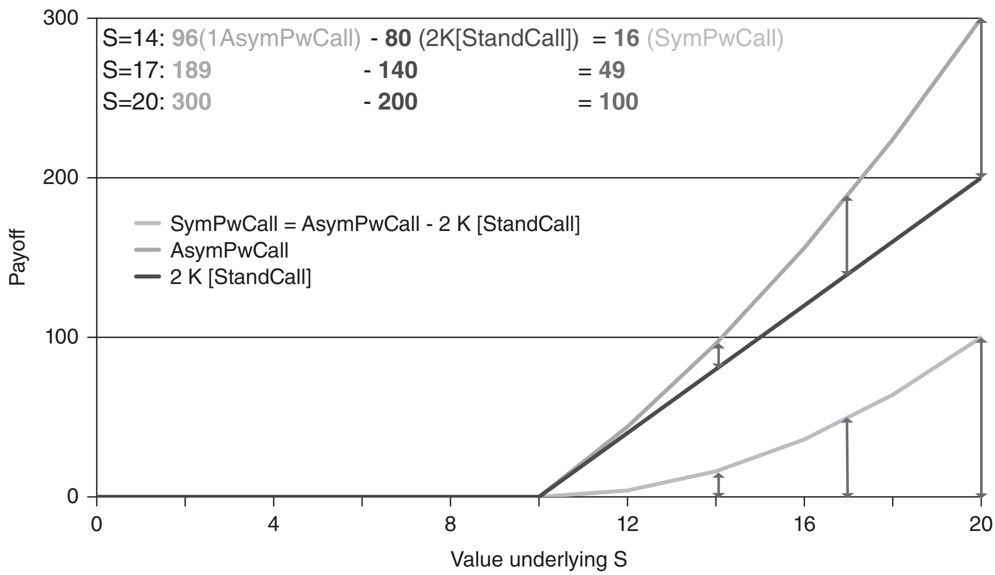

FIGURE 1.54 Symmetric power call replicated with asymmetric power and vanilla calls, using $K = 10$, $n = 2$.

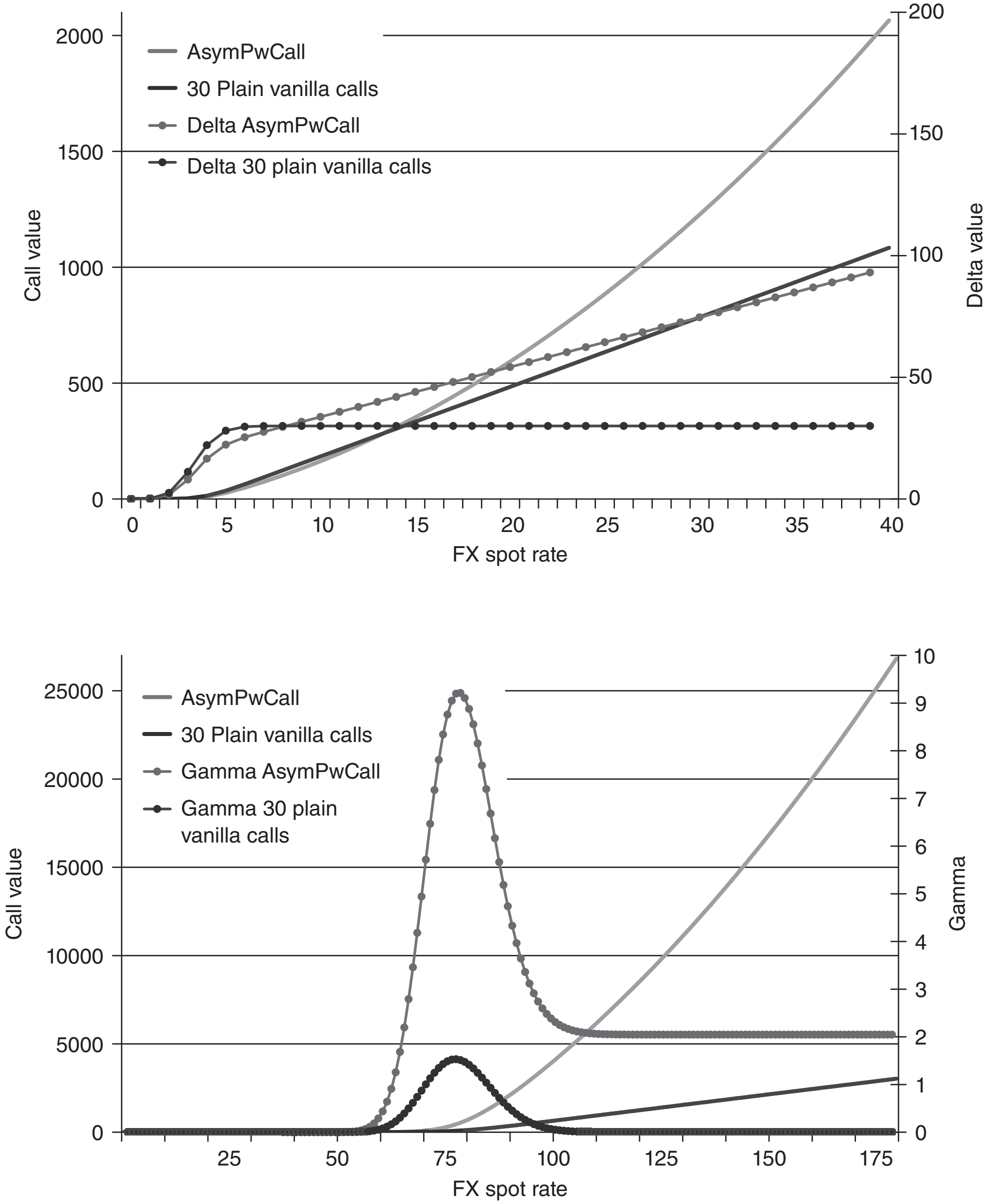

FIGURE 1.55 Asymmetric power call and vanilla call value, delta (lhs) and gamma (rhs) on the spot space, using $K = 10$, $n = 2$, $\sigma = 20\%$, $r_d = 5\%$, $r_f = 0\%$, $T = 90$ days.

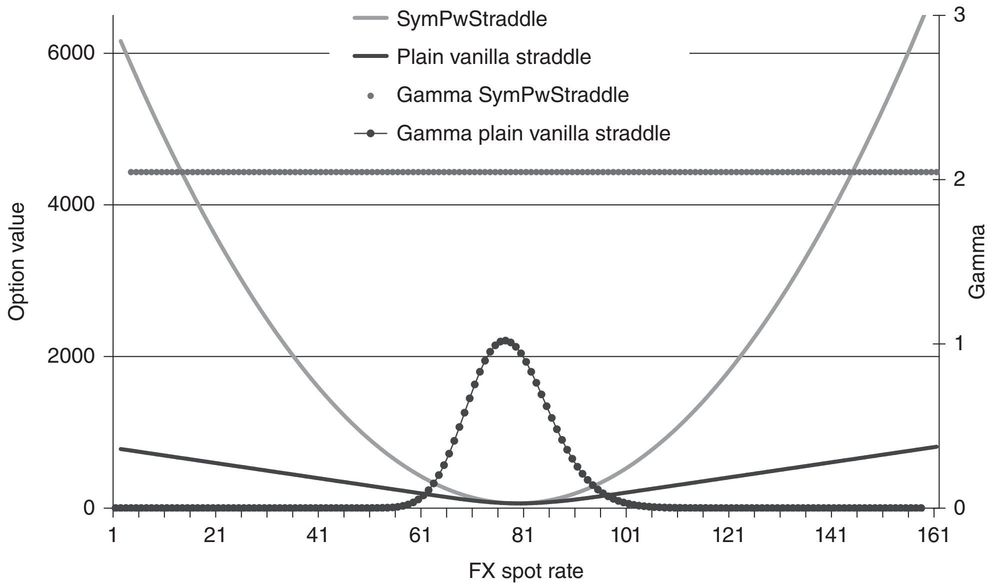

FIGURE 1.56 Gamma exposure of a symmetric power versus vanilla straddle, using $K = 80$ (at-the-money), $n = 2$, $\sigma = 20\%$, $r_d = 5\%$, $r_f = 0\%$, $T = 90$ days.

At the same time, if the underlying increases significantly, the symmetric power straddle preserves the exposure to volatility, whereas the vanilla straddle value becomes more and more invariant to the volatility input. Therefore, the power straddle is useful to hedge implied volatility, see Figure 1.57. This feature is similar to a variance swap explained in Section 1.8.8.

Replication with Vanilla Options The insights from the option payoffs, valuation, and the sensitivity analysis provide an effective static replication strategy for both asymmetric and symmetric power options. The respective call values are considered as an example.

Approximate Static Replication The continuous curvature of a power option can be approximated piecewise, adding up linear payoffs of vanilla options with different strike prices, see [129].

The symmetric power call for $n = 2$ can be replicated statically, as explained in the pricing section, just an asymmetric power call less $2K$ vanilla options, see Table 1.29.

The piecewise linear approximation with vanilla options is a super-replication, whence the value of the super-replicating portfolio is a natural upper boundary for the value of the symmetric power option as it overestimates the option value, see Table 1.30. The complexity of a static replication increases enormously with higher values of n. For the above example, a package of 25499 (3439) vanilla options is required to replicate one asymmetric (symmetric) power call. Overall, the static replication works very well, as can be seen in Figure 1.58. The static super-replication of power options with vanilla options takes into account the smile correctly, whence the value of the static

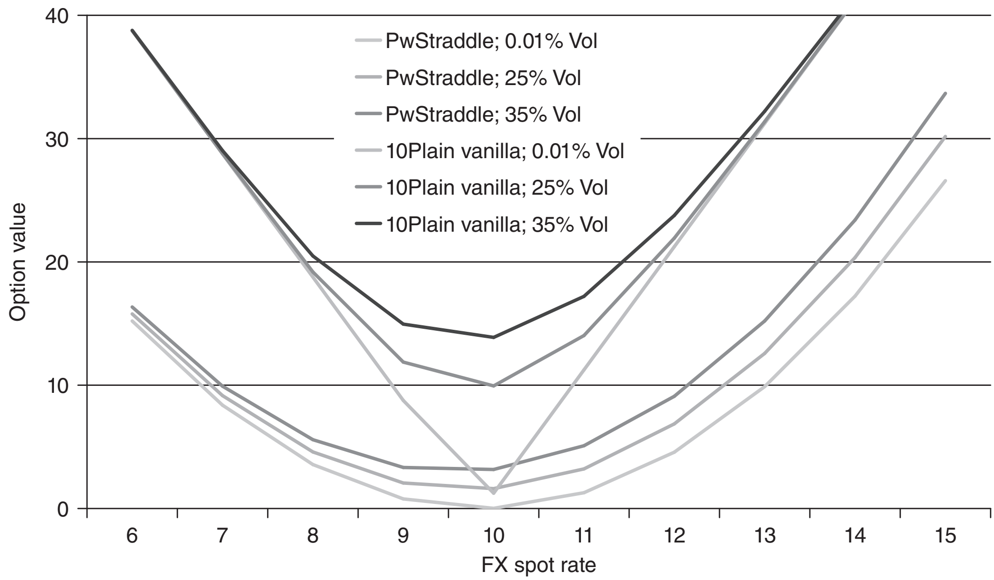

FIGURE 1.57 Vega exposure of a symmetric power versus vanilla straddle, using $K = 10$ (at-the-money), $n = 2$, $\sigma = 20\%$, $r_d = 5\%$, $r_f = 0\%$, $T = 90$ days.

TABLE 1.29 Static replication for the asymmetric power call, using $K = 20$, $n = 2$. For the symmetric power call the $2K$ standard calls need to be removed.

Underlying price		10	11	12	13	14	15	16	17	18	19	20
Asym. power call		0	21	44	69	96	125	156	189	224	261	300
Vanilla call package	Sum	0	21	44	69	96	125	156	189	224	261	300
Package components	Strike											
$2K$ standard calls	10		20	40	60	80	100	120	140	160	180	200
One standard call	10		1	2	3	4	5	6	7	8	9	10
Two standard calls	11			2	4	6	8	10	12	14	16	18
Two standard calls	12				2	4	6	8	10	12	14	16
Two standard calls	13					2	4	6	8	10	12	14
Two standard calls	14						2	4	6	8	10	12
Two standard calls	15							2	4	6	8	10
Two standard calls	16								2	4	6	8
Two standard calls	17									2	4	6
Two standard calls	18										2	4
Two standard calls	19											2

super-replicating portfolio can serve as a market value of the power option. All the sensitivities can be correctly derived from the sensitivities of the super-replicating portfolio. Since even far-out strikes need to be taken into consideration, the precision of the smile on the wings (extrapolation) is crucial. The power option shares this feature with the variance swap.

TABLE 1.30 Asymmetric power call replication versus formula value, using $K = 10$ (at-the-money), $n = 2$, $\sigma = 15\%$, $r_d = 5\%$, $r_f = 0\%$, $T = 90$ days.

Asym. power call	**Formula value**	**11.91**
Vanilla call package	Sum	12
Package components	Strike	
2K standard calls	10	11.0354
One standard call	10	0.55177
Two standard calls	11	0.33257
Two standard calls	12	0.06928
Two standard calls	13	0.01034
Two standard calls	14	0.00116
Two standard calls	15	0.00010
Two standard calls	16	7.55E-06
Two standard calls	17	4.74E-07
Two standard calls	18	2.64E-08
Two standard calls	19	1.33E-09

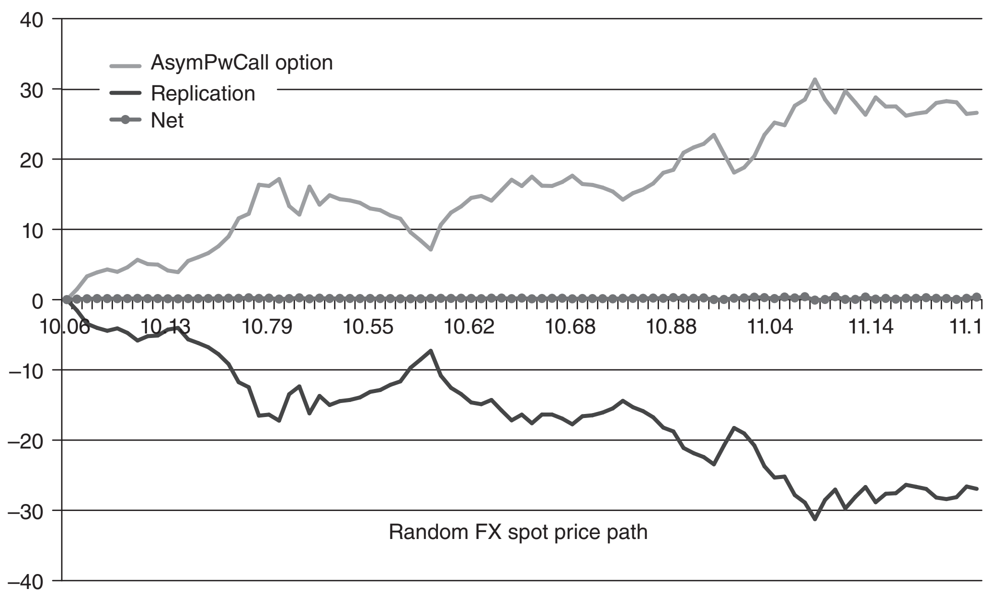

FIGURE 1.58 Static replication performance of an asymmetric power call, using $K = 10$, $n = 2$, $\sigma = 150\%$, $r_d = 5\%$, $r_f = 0\%$, $T = 90$ days.

Dynamic Hedging A dynamic hedge involves setting up and managing a position in the underlying currencies that offsets any value change in the option position. Usually the difficulty of dynamic hedging lies in second order, that is in gamma, and in vega risk. As the symmetric power straddle has a constant gamma, this simplifies delta hedging activities. In practice, however, it is common for the risk-warehousing desk to first set up a static replication and then hedge the Greeks of the residual position dynamically.

1.7.10 Quanto Options

A quanto option can be any cash-settled option, whose payoff is converted into a third currency at maturity at a pre-specified rate, called the *quanto factor*. There can be quanto plain vanilla, quanto barriers, quanto forward starts, quanto corridors, etc. The valuation theory is covered for example in [120] and [65]. This section is based on [138].

FX Quanto Drift Adjustment We take the example of a gold contract with underlying XAU/USD in XAU-USD quotation that is quantoed into EUR. Since the payoff is in EUR, we let EUR be the numeraire or domestic or base currency and consider a Black-Scholes model

$$\text{XAU-EUR: } dS_t^{(3)} = (r_{EUR} - r_{XAU})S_t^{(3)}\,dt + \sigma_3 S_t^{(3)}\,dW_t^{(3)}, \tag{337}$$

$$\text{USD-EUR: } dS_t^{(2)} = (r_{EUR} - r_{USD})S_t^{(2)}\,dt + \sigma_2 S_t^{(2)}\,dW_t^{(2)}, \tag{338}$$

$$dW_t^{(3)} dW_t^{(2)} = \rho_{23}\,dt, \tag{339}$$

where we use an (implicit) plus sign in front of the correlation because both $S^{(3)}$ and $S^{(2)}$ have the same base currency (DOM), which is EUR in this case. The scenario is displayed in Figure 1.59. The actual underlying is then

$$\text{XAU-USD: } S_t^{(1)} = \frac{S_t^{(3)}}{S_t^{(2)}}. \tag{340}$$

Using Itô's formula, we first obtain

$$\begin{aligned} d\frac{1}{S_t^{(2)}} &= -\frac{1}{(S_t^{(2)})^2}\,dS_t^{(2)} + \tfrac{1}{2}\cdot 2\cdot\tfrac{1}{(S_t^{(2)})^3}\,(dS_t^{(2)})^2 \\ &= (r_{USD} - r_{EUR} + \sigma_2^2)\frac{1}{S_t^{(2)}}\,dt - \sigma_2\frac{1}{S_t^{(2)}}\,dW_t^{(2)}, \end{aligned} \tag{341}$$

and hence

$$\begin{aligned} dS_t^{(1)} &= \frac{1}{S_t^{(2)}}\,dS_t^{(3)} + S_t^{(3)}\,d\frac{1}{S_t^{(2)}} + dS_t^{(3)}\,d\frac{1}{S_t^{(2)}} \\ &= \frac{S_t^{(3)}}{S_t^{(2)}}(r_{EUR} - r_{XAU})\,dt + \frac{S_t^{(3)}}{S_t^{(2)}}\sigma_3\,dW_t^{(3)} \\ &\quad + \frac{S_t^{(3)}}{S_t^{(2)}}(r_{USD} - r_{EUR} + \sigma_2^2)\,dt - \frac{S_t^{(3)}}{S_t^{(2)}}\sigma_2\,dW_t^{(2)} - \frac{S_t^{(3)}}{S_t^{(2)}}\rho_{23}\sigma_2\sigma_3\,dt \\ &= (r_{USD} - r_{XAU} + \sigma_2^2 - \rho_{23}\sigma_2\sigma_3)S_t^{(1)}\,dt + S_t^{(1)}(\sigma_3\,dW_t^{(3)} - \sigma_2\,dW_t^{(2)}). \end{aligned}$$

Since $S_t^{(1)}$ is a geometric Brownian motion with volatility σ_1, we introduce a new Brownian motion $W_t^{(1)}$ and find

$$dS_t^{(1)} = (r_{USD} - r_{XAU} + \sigma_2^2 - \rho_{23}\sigma_2\sigma_3)S_t^{(1)}\,dt + \sigma_1 S_t^{(1)}\,dW_t^{(1)}. \tag{342}$$

Now Figure 1.59 and the *law of cosine* imply

$$\sigma_3^2 = \sigma_1^2 + \sigma_2^2 - 2\rho_{12}\sigma_1\sigma_2, \tag{343}$$

$$\sigma_1^2 = \sigma_2^2 + \sigma_3^2 - 2\rho_{23}\sigma_2\sigma_3, \tag{344}$$

which yields

$$\sigma_2^2 - \rho_{23}\sigma_2\sigma_3 = \rho_{12}\sigma_1\sigma_2. \tag{345}$$

As explained in the *currency triangle* in Figure 1.59, ρ_{12} is the correlation between XAU-USD and EUR-USD, whence $\rho \stackrel{\Delta}{=} -\rho_{12}$ is the correlation between XAU-USD and USD-EUR (i.e. the correlation between the currency pairs FOR-DOM and DOM-QUANTO). Inserting this into Equation (342), we obtain the usual formula for the drift adjustment

$$dS_t^{(1)} = (r_{USD} - r_{XAU} - \rho\sigma_1\sigma_2)S_t^{(1)}\,dt + \sigma_1 S_t^{(1)}\,dW_t^{(1)}. \tag{346}$$

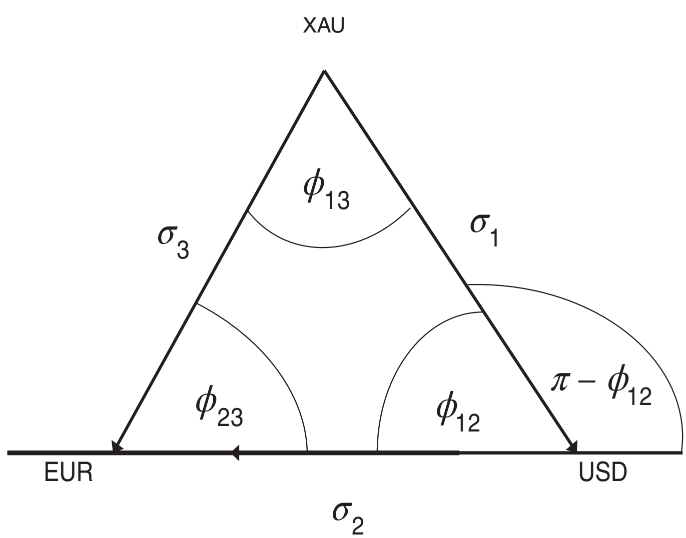

FIGURE 1.59 XAU-USD-EUR FX quanto triangle. The arrows point in the direction of the respective base currencies. The length of the edges represents the volatility. The cosine of the angles $\cos\phi_{ij} = \rho_{ij}$ represents the correlation of the currency pairs $S^{(i)}$ and $S^{(j)}$, *if* both $S^{(i)}$ and $S^{(j)}$ have the same base currency (DOM). If the base currency (DOM) of $S^{(i)}$ is the underlying currency (FOR) of $S^{(j)}$, then the correlation is denoted by $-\rho_{ij} = \cos(\pi - \phi_{ij})$.

This is the risk-neutral pricing process that can be used for the valuation of any derivative depending on $S_t^{(1)}$ which is quantoed into EUR.

Extensions to Other Models The previous derivation can be extended to the case of term structure of volatility and correlation. However, introduction of volatility smile would distort the relationships. Nevertheless, accounting for smile effects is important in real market scenarios. To do this, one could, for example, capture the smile for a multi-currency model with a *weighted Monte Carlo technique* as described by Avellaneda *et al.* in [5]. This would still allow the previous result to be used.

Quanto Vanilla Common among foreign exchange options is a quanto plain vanilla paying

$$Q[\phi(S_T - K)]^+, \tag{347}$$

where K denotes the strike, T the expiration time, ϕ the usual put-call indicator taking the value +1 for a call and −1 for a put, S the underlying in FOR-DOM quotation, and Q the quanto factor from the domestic currency into the quanto currency. We let

$$\tilde{\mu} \stackrel{\Delta}{=} r_d - r_f - \rho\sigma\tilde{\sigma}, \tag{348}$$

be the *adjusted drift*, where r_d and r_f denote the risk-free interest rates of the domestic and foreign underlying currency pair respectively, $\sigma = \sigma_1$ the volatility of this currency pair, $\tilde{\sigma} = \sigma_2$ the volatility of the currency pair DOM-QUANTO, and

$$\rho = \frac{\sigma_3^2 - \sigma^2 - \tilde{\sigma}^2}{2\sigma\tilde{\sigma}} \tag{349}$$

the correlation between the currency pairs FOR-DOM and DOM-QUANTO in this quotation. Furthermore we let r_Q be the risk-free interest rate of the quanto currency. With the same principles as in [139] we can derive the formula for the value as

$$v = Qe^{-r_Q T}\phi[S_0 e^{\tilde{\mu}T}\mathcal{N}(\phi d_+) - K\mathcal{N}(\phi d_-)], \tag{350}$$

$$d_\pm = \frac{\ln\frac{S_0}{K} + \left(\tilde{\mu} \pm \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}, \tag{351}$$

where $\mathcal{N}$ denotes the cumulative standard normal distribution function and n its density.

Quanto Forward Similarly, we can easily determine the value of a quanto forward paying

$$Q[\phi(S_T - K)], \tag{352}$$

where K denotes the strike, T the expiration time, ϕ the usual long-short indicator, S the underlying in FOR-DOM quotation, and Q the quanto factor from the domestic currency into the quanto currency. Then the formula for the value can be written as

$$v = Qe^{-r_Q T}\phi[S_0 e^{\tilde{\mu}T} - K]. \tag{353}$$

This follows from the vanilla quanto value formula by taking both the normal probabilities to be one. These normal probabilities are exercise probabilities under some measure. Since a forward contract is always exercised, both these probabilities must be equal to 1.

Quanto Digital A European style quanto digital pays

$$Q\mathbb{I}_{\{\phi S_T \geq \phi K\}}, \tag{354}$$

where K denotes the strike, S_T the spot of the currency pair FOR-DOM at maturity T, ϕ takes the values +1 for a digital call and −1 for a digital put, and Q is the pre-specified conversion rate from the domestic to the quanto currency. The valuation of European style quanto digitals follows the same principle as in the quanto vanilla option case. The value is

$$v = Qe^{-r_Q T}\mathcal{N}(\phi d_-). \tag{355}$$

We provide an example of a European style digital put in USD/JPY quanto into EUR in Table 1.31.

Hedging of Quanto Options Dynamic hedging of quanto options can be done by running a multi-currency options book. All the usual Greeks can be hedged. Delta hedging is done by trading in the underlying spot market. An exception is the *correlation risk*, which can only be hedged with other derivatives depending on the same correlation. This is

TABLE 1.31 Example of a quanto digital put. The buyer receives 100,000 EUR if at maturity the ECB fixing for USD-JPY (computed via EUR-JPY and EUR-USD) is below 108.65. Terms were created on January 12 2004 with the following market data: USD-JPY spot ref 106.60, USD-JPY ATM vol 8.55%, EUR-JPY ATM vol 6.69%, EUR-USD ATM vol 10.99% (corresponding to a correlation of −27.89% for USD-JPY against JPY-EUR), USD rate 2.5%, JPY rate 0.1%, EUR rate 4%.

Notional	100,000 EUR
Maturity	3 months (92 days)
European style barrier	108.65 USD-JPY
Theoretical value	71,555 EUR
Fixing source	ECB

generally not possible in other asset classes. In FX the correlation risk can be translated into a vega position as shown in [135] or in Section 1.9.2 on foreign exchange basket options. We illustrate this approach for quanto plain vanilla options now.

Vega Positions of Quanto Plain Vanilla Options Starting from Equation (350), we obtain the sensitivities

$$\frac{\partial v}{\partial \sigma} = QS_0 e^{(\tilde{\mu}-r_Q)T}\left[n(d_+)\sqrt{T} - \phi\mathcal{N}(\phi d_+)\rho\tilde{\sigma}T\right],$$

$$\frac{\partial v}{\partial \tilde{\sigma}} = -QS_0 e^{(\tilde{\mu}-r_Q)T}\phi\mathcal{N}(\phi d_+)\rho\sigma T,$$

$$\frac{\partial v}{\partial \rho} = -QS_0 e^{(\tilde{\mu}-r_Q)T}\phi\mathcal{N}(\phi d_+)\sigma\tilde{\sigma}T,$$

$$\begin{aligned}\frac{\partial v}{\partial \sigma_3} &= \frac{\partial v}{\partial \rho}\frac{\partial \rho}{\partial \sigma_3}\\ &= \frac{\partial v}{\partial \rho}\frac{\sigma_3}{\sigma\tilde{\sigma}}\\ &= -QS_0 e^{(\tilde{\mu}-r_Q)T}\phi\mathcal{N}(\phi d_+)\sigma\tilde{\sigma}T\frac{\sigma_3}{\sigma\tilde{\sigma}}\\ &= -QS_0 e^{(\tilde{\mu}-r_Q)T}\phi\mathcal{N}(\phi d_+)\sigma_3 T\\ &= -QS_0 e^{(\tilde{\mu}-r_Q)T}\phi\mathcal{N}(\phi d_+)\sqrt{\sigma^2+\tilde{\sigma}^2+2\rho\sigma\tilde{\sigma}}T.\end{aligned}$$

Note that the computation is standard calculus and repeatedly uses Identity (37). The understanding of these Greeks is that σ and $\tilde{\sigma}$ are both risky parameters, independent of each other. The third independent source of risk is either σ_3 or ρ, depending on what is more likely to be known. This shows exactly how the three vega positions can be hedged with plain vanilla options in all three legs, provided there is a liquid vanilla options market in all three legs. In the example with XAU-USD-EUR the currency pairs XAU-USD and EUR-USD are traded; however, there is a very illiquid vanilla market in XAU-EUR. Therefore, the correlation risk remains unhedgeable. Similar statements would apply for quantoed stocks or stock indices. However, in FX, there are situations with all legs being hedgeable, for instance EUR-USD-JPY.

The signs of the vega positions are not uniquely determined in all legs. The FOR-DOM vega is smaller than the corresponding vanilla vega in the case of a call and positive correlation or put and negative correlation, larger in the case of a put and positive correlation or call and negative correlation. The DOM-Q vega takes the sign of the correlation in the case of a call and its opposite sign in the case of a put. The FOR-Q vega takes the opposite sign of the put-call indicator ϕ.

We provide an example of pricing and vega hedging scenario in Table 1.32, where we notice, that dominating vega risk comes from the FOR-DOM pair, whence most of the risk can be hedged.

Applications Standard applications are performance linked deposits or notes explained in Section 2.4.2 or participation notes explained in Section 2.6. Any time the performance of an underlying asset needs to be converted into the notional currency invested,

TABLE 1.32 Example of a quanto plain vanilla.

		Data set 1	Data set 2	Data set 3
FX pair	FOR-DOM	XAU-USD	XAU-USD	XAU-USD
Spot	FOR-DOM	800.00	800.00	800.00
Strike	FOR-DOM	810.00	810.00	810.00
Quanto	DOM-Q	1.0000	1.0000	1.0000
Volatility	FOR-DOM	10.00%	10.00%	10.00%
Quanto volatility	DOM-Q	12.00%	12.00%	12.00%
Correlation	FOR-DOM – DOM-Q	25.00%	25.00%	−75.00%
Domestic interest rate	DOM	2.0000%	2.0000%	2.0000%
Foreign interest rate	FOR	0.5000%	0.5000%	0.5000%
Quanto currency rate	Q	4.0000%	4.0000%	4.0000%
Time in years	T	1	1	1
1=call −1=put	FOR	1	−1	1
Quanto vanilla option	value	30.81329	31.28625	35.90062
Quanto vanilla option	vega FOR-DOM	298.14188	321.49308	350.14600
Quanto vanilla option	vega DOM-Q	−10.07056	9.38877	33.38797
Quanto vanilla option	vega FOR-Q	−70.23447	65.47953	−35.61383
Quanto vanilla option	correlation risk	−4.83387	4.50661	−5.34207
Quanto vanilla option	vol FOR-Q	17.4356%	17.4356%	8.0000%
Vanilla option	value	32.6657	30.7635	32.6657
Vanilla option	vega	316.6994	316.6994	316.6994

and the exchange rate risk is with the seller, we need a quanto product. Naturally, an underlying like gold, which is quoted in USD, would be a default candidate for a quanto product, when the investment is in a currency other than USD. This shows that quanto features play a key role in asset management whenever the investor wants to protect his investment against foreign exchange rate risk. Quanto options are embedded in many investment strategies with underlying assets being stocks or stock indices. Quanto options in foreign exchange markets do not trade as much because FX is not viewed as an interesting asset class by many asset managers. This may change in the future. On the treasury side, quanto options do sometimes come up in currency related swaps, as explained in Section 2.5.7.

1.7.11 Exercises

Foreign Digital Value via Static Replication Derive Equation (168) using a static replication of the foreign digital by domestic digital and vanilla options.

Foreign Digital Value via Change of Measure Derive Equation (168) by calculating the expected value of the payoff (166) and a change of measure.

Compound Consider a EUR-USD market with spot at 1.2500, EUR rate at 2.5%, USD rate at 2.0%, volatility at 10.0%, and the situation of a treasurer expecting to receive 1 M USD in one year, that he wishes to change into EUR at the current spot rate of

1.2500. In six months he will know whether the company gets the definite order. Compute the price of a vanilla EUR call USD put in EUR. Alternatively compute the price of a compound with two thirds of the total premium to be paid at inception and one third to be paid in six months. Do the same computations if the sales margin for the vanilla is 1 EUR per 1,000 USD notional and for the compound is 2 EUR per 1,000 USD notional. After six months the company ends up not getting the order and can waive its hedge. How much would it get for the vanilla if the spot is at 1.1500, at 1.2500 and at 1.3500? Would it be better for the treasurer to own the compound and not pay the second premium? How would you split up the premium for the compound to persuade the treasurer to buy the compound rather than the vanilla? (After all, there is more margin to earn.)

Perpetual One-Touch Replication Find the fair price and a semi-static replication of a *perpetual one-touch*, which pays one unit of the domestic currency if the barrier $H > S_0$ is ever hit, where S_0 denotes the current exchange rate. How about payment in the foreign currency? How about a *perpetual no-touch*? These thoughts are developed further to a *vanilla-one-touch duality* by Peter Carr [22].

Perpetual Double-One-Touch Find the value of a *perpetual double-one-touch*, which pays a rebate R_H, if the spot reaches the higher level H before the lower level L, and R_L, if the spot reaches the lower level first. Consider as an example the EUR-USD market with a spot of S_0 at time zero between L and H. Let the interest rates of both EUR and USD be zero and the volatility be 10%. The specified rebates are paid in USD. There is no finite expiration time, but the rebate is paid whenever one of the levels is reached. How would you replicate a short position (semi-) statically?

Strike-Out Replication and Impact of Jumps A call (put) option is the right to buy (sell) one unit of foreign currency on a maturity date T at a pre-defined price K, called the strike price. A knock-out call with barrier B is like a call option that becomes worthless if the underlying ever touches the barrier B at any time between inception of the trade and its expiration time. Let the market parameters be spot $S_0 = 120$, all interest rates be zero, volatility $\sigma = 10\%$. In a liquid and jump-free market, find the value of a one-year *strike-out*, i.e. a down-and-out knock-out call, where $K = B = 100$.

Suppose now that the spot price movement can have downward jumps but the forward price is still constant and equal to the spot (since we assume zero interest rates). How do these possible jumps influence the value of the knock-out call?

The solution to this problem is used for the design of *turbo notes*, see Section 2.6.4.

Strike-Out Call Vega What is the vega profile as a function of spot for a strike-out call? What can you say about the sign of vega?

Double-No-Touch with Notional in Foreign Currency Given Equation (235), which represents the theoretical value (TV) of a double-no-touch in units of domestic currency, where the payoff currency is also domestic, let us denote this function by

$$v^d(S, r_d, r_f, \sigma, L, H), \tag{356}$$

where the superscript d indicates that the payoff currency is domestic. Using this formula, prove that the corresponding value in domestic currency of a double-no-touch paying one unit of *foreign* currency is given by

$$v^f(S, r_d, r_f, \sigma, L, H) = Sv^d\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right). \tag{357}$$

Assuming you know the sensitivity parameters of the function v^d, derive the following corresponding sensitivity parameters for the function v^f,

$$\frac{\partial v^f}{\partial S} = v^d\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right) - \frac{1}{S}\frac{\partial v^d}{\partial S}\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right), \tag{358}$$

$$\frac{\partial^2 v^f}{\partial S^2} = \frac{1}{S^3}\frac{\partial^2 v^d}{\partial S^2}\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right),$$

$$\frac{\partial v^f}{\partial \sigma} = S\frac{\partial v^d}{\partial \sigma},$$

$$\frac{\partial^2 v^f}{\partial \sigma^2} = S\frac{\partial^2 v^d}{\partial^2 \sigma},$$

$$\frac{\partial^2 v^f}{\partial S \partial \sigma} = S\frac{\partial v^d}{\partial \sigma}\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right) - \frac{1}{S}\frac{\partial^2 v^d}{\partial S \partial \sigma}\left(\frac{1}{S}, r_f, r_d, \sigma, \frac{1}{H}, \frac{1}{L}\right),$$

$$\frac{\partial v^f}{\partial r_d} = \frac{\partial v^d}{\partial r_f},$$

$$\frac{\partial v^f}{\partial r_f} = \frac{\partial v^d}{\partial r_d}.$$

Static Replication of DNT with DKO Suppose your front-office application for double-no-touch contracts is out of order, but you can use double-knock-out options. Replicate statically a double-no-touch paying one unit of domestic currency using two double-knock-out options.

Static Replication of a FOR-Paying Double-No-Touch The nominal amounts of the respective double-knock-out options that statically replicate the double-no-touch depend on the currency in which the payoff is settled. In the case of a EUR-USD double-no-touch paying one unit of USD (domestic currency), the nominal amounts of the double-knock-out call and put must be chosen to be both $\frac{1}{H-L}$ (which currency?). Find a static replication of a double-no-touch paying one unit of EUR (foreign currency).

Static Replication of a FOR-Paying Double-No-Touch with One Double-Knock-Out Option Can you statically replicate a double-no-touch paying one unit of foreign currency using just *one* double-knock-out option?

Gold Price Return vs. EUR-USD Return The price of an ounce of gold is quoted in USD. If the price of Gold drops by 5%, but the price of Gold in EUR remains constant, determine the change of the EUR-USD exchange rate.

Vega Profile of Double-No-Touch Suppose you are long a double-no-touch. Draw the possible vega profiles as a function of the spot and discuss the possible scenarios.

Vanilla Vega Hedging Suppose you know the vega of a two-month at-the-money vanilla. By what factor is the vega of a four-month at-the-money vanilla bigger? How does this look for a five-year vanilla in comparison with a 10-year vanilla? You may assume a delta-neutral straddle notion of at-the-money.

One-Touch Replication with Digitals Suppose the exchange rate S follows a Brownian motion without drift and constant volatility. How can you replicate semi-statically a single-one-touch with European digital call or put options? Hint: Use the reflection principle. Which risk remains unhedged in your semi-static replication?

Static Replication of European Barrier Options with Vanilla and Digital Options Given vanillas and digitals, how can you statically replicate European style barrier options?

Self-Quanto Option Determine the Black-Scholes value of a CHF-EUR *self-quanto option* with strike K, which is cash-settled in CHF at maturity. As an example you may consider the payoff

$$N \cdot \frac{(K - S_T)^+}{S_T} \tag{359}$$

paid in CHF, where N is a CHF notional amount, K a strike in EUR-CHF, and S_T the spot price in EUR-CHF at maturity time T.

Implementing the View of a Rising USD/JPY Suppose a client believes very strongly that USD/JPY will reach a level of 120.00 or higher in three months' time. With a current spot level of 110.00, volatility of 10%, JPY rate of 0%, USD rate of 3%, find a contract reflecting the client's view and create a term sheet for the client explaining chances and risks.

Gamma of a Symmetric Power Straddle Prove that a symmetric power straddle has a constant gamma. What does this imply for delta, vega, rho (domestic and foreign), and theta?

Static Replication of RKO with KO and Touch Contracts Let NT(B) and OT(B) denote the value of a no-touch and a one-touch with barrier B respectively, both paid at the end. Let KOPut(K, B) and KOCall(K, B) denote the value of a regular knock-out put and call with strike K and barrier B respectively. Let SOPut(K) and SOCall(K) denote the value of a strike-out put and call with strike K and barrier K respectively. Finally, let

RKOPut(K, B) and RKOCall(K, B) denote the value of a reverse knock-out put and call with strike K and barrier B respectively. How can you replicate statically reverse knock-outs using touch contracts, strike-outs and regular knock-outs? Support your answer with a suitable figure. This implies in particular that the market prices for reverse knock-outs can be implied from the market prices of touch contracts and regular barrier options.

Non-Deliverable Self-Quanto Forward Consider the value function (353) of a quanto forward. Derive the value function of a (non-deliverable) *self-quanto forward*, for example in a EUR-USD market, with payoff $S_T - K$ in EUR rather than USD. If the amount he receives is negative, then the client pays.

No-Touch Value with Infinite Volatility Derive the limiting value in the Black-Scholes model of a no-touch with either an upper barrier or a lower barrier as volatility goes to infinity, and justify your answer intuitively.

1.8 SECOND GENERATION EXOTICS (SINGLE CURRENCY PAIR)

In this section we present an overview of some of the most common second generation exotics, of which some are exotic options and others are more similar to forwards and swaps. Many of the transactions are considered second generation because they are multi-currency transactions with correlation risk (and are not simple quantos). These include outside barrier options, spread and exchange options, basket options, best-of and worst-of options. Other second generation exotics are considered volatility trades, including the variance swap, volatility swap, and forward volatility agreement. The correlation swap is a correlation trade as its name indicates. Other second generation exotics comprise options or strategies with path-dependent notional amounts (corridors, faders) or some barrier options with special features. There are also path-independent second generation exotics on one underlying. As an example we consider the multiplicity power option.

1.8.1 Multiplicity Power Options

The basic idea of a power option (see Section 1.7.9) can be further extended. An extreme case of selling options in private banking was this: BNPP offered something called a "multiplicity" before the 2008 financial crisis. An investor sells a USD Put/JPY Call with strike X to the bank presumably for some huge premium expressed as a yield on some yield-enhancing note or swap. The strike X is taken as $\frac{K}{S^2}$, where S is the USD-JPY spot and K would be set somewhere around initial spot cubed, S^3, to make the initial strike below ATM. Thus as spot S falls, the strike X of the USD Put shoots up. Furthermore, X was capped at 600. Figure 1.60 illustrates the payoff of the multiplicity compared with a vanilla put. The risk the investor takes is substantial but well hidden in the nested definition of the powers. Hard to imagine that these products were very popular and traded a lot in Asia. The industry can gain back a lot of trust by explaining the risks appropriately before the product trades.

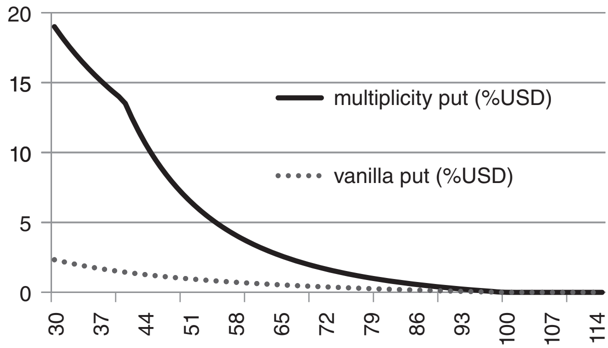

FIGURE 1.60 Payoff of a multiplicity power put compared with a vanilla put with strike $K = 100$ and cap of 600. USD-JPY final spot on the x-axis.

1.8.2 Corridors/Range Accruals

A European *corridor* or *range accrual (RAC)* entitles its holder to receive a pre-specified amount of a currency (say EUR) on a specified date (maturity) proportional to the number of fixings inside a range between the start date and maturity. The buyer has to pay a premium for this product.

Advantages

- High leverage product, high profit potential
- Can take advantage of a quiet market phase
- Easy to price and to understand

Disadvantages

- Not suitable for the long term
- Expensive product
- Price spikes and large market movements can lead to loss

Figure 1.61 shows a sample scenario for a corridor. At delivery, the holder receives $\frac{n}{N}\cdot$ notional, where n is the number of fixings inside the range and N denotes the maximum number of fixings.

Types of Corridors

European style corridor. The corridor is *resurrecting*, i.e. all fixings inside the range count for the accumulation, even if some of the fixings are outside. Given a *fixing schedule* $\{S_{t_1}, S_{t_2}, \ldots, S_{t_N}\}$ the payoff can be specified by

$$\text{notional} \cdot \frac{1}{N} \sum_{i=1}^{N} I\!I_{\{S_{t_i} \in [L,H]\}}, \tag{360}$$

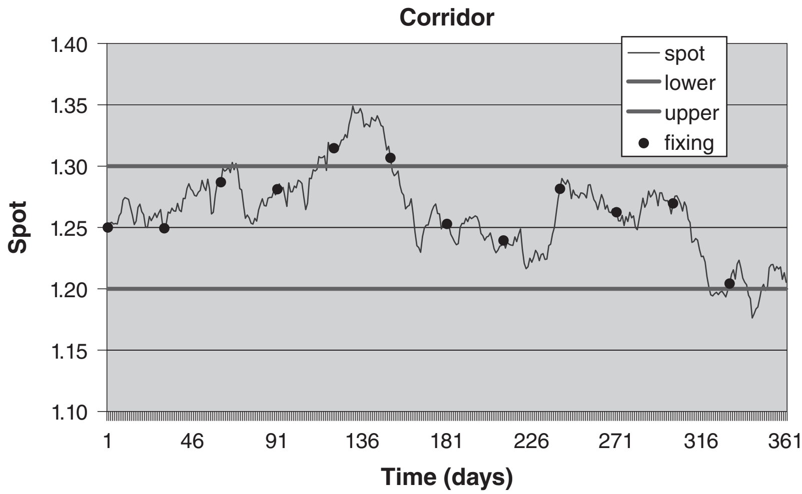

FIGURE 1.61 Example of a corridor or range accrual with spot 1.2500, domestic interest rate 3.00%, foreign interest rate 2.75%, volatility 10%, for a maturity of 1 year with 12 monthly fixings indicated by the dots. The range is 1.2000–1.3000. In a resurrecting corridor, the investor would accumulate 10 out of 12 fixings. In a non-resurrecting corridor, the investor would accumulate 4 out of 12 fixings as the fifth is outside the range.

where N denotes the total number of fixings, L the lower barrier, H the higher barrier.

American style corridor. This is *non-resurrecting*, i.e. only fixing dates count for the accumulation that occur before the first fixing is outside the range. The holder of the corridor keeps the accumulated amount. Introducing the stopping time

$$\tau \stackrel{\Delta}{=} \min\{t : S_{t_i} \notin (L, H)\}, \tag{361}$$

the payoff can be specified by

$$\text{notional} \cdot \frac{1}{N} \sum_{i=1}^{N} I\!I_{\{S_{t_i} \in [L,H]\}} I\!I_{\{t_i < \tau\}}. \tag{362}$$

American style corridor with continuously observed knock-out. This is an American style corridor, where all of the accumulated amount is lost once the exchange rate trades at or outside the range. This is equivalent to a double-no-touch. The payoff can be specified by

$$\text{notional} \cdot \frac{1}{N} \sum_{i=1}^{N} I\!I_{\{S_{t_i} \in [L,H]\}} I\!I_{\{L < \min_{0 \le t \le T} S_t \le \max_{0 \le t \le T} S_t < H\}}, \tag{363}$$

where T denotes the maturity time. This type of corridor can be generalized as the range for the fixings does not need to be identical to the range for the continuously observed knock-out condition.

American style corridor with discrete knock-out. This is like an American style corridor where the knock-out occurs when the fixing is outside the range for the first time, i.e. we use the stopping time in Equation (361) and replace the payoff by

$$\text{notional} \cdot \frac{1}{N} \sum_{i=1}^{N} I\!I_{\{S_{t_i} \in [L,H]\}} I\!I_{\{\tau > T\}}. \tag{364}$$

In this case the holder receives either the full notional or nothing, so it is very similar to a double-no-touch.

Forward start corridor. In this type, which can be European or American as before, the range will be set relative to a future spot level, see also Section 1.7.8.

Example An investor wants to benefit from his view that the EUR-USD exchange rate will be often between two levels during the next 12 months. In this case he may consider buying a European corridor as presented in Table 1.33 for example.

If the investor's market expectation is correct, then he will receive 1 M EUR at delivery, twice the initial premium.

Explanations

Fixings are *official* exchange rate sources such as from the European Central Bank, the Federal Reserve Bank or private banks, which take place on each business day. For details on the impact on pricing see Becker and Wystup [13].

Fixing source is the exact source of the fixing, for example Reuters pages ECB37, WMRSPOT11, or BFIX on Bloomberg.[5]

Fixing schedule requires a start date, an end date, and a frequency such as daily, weekly, or monthly. It can also be customized. Since there are often disputes about holidays, it is advisable to specify any fixing schedule explicitly in the deal confirmation. A common way is to agree on the open days in the TARGET system.

TABLE 1.33 Example of a European corridor. To compare, the premium for the same corridor in American style would be 100,000 EUR.

Spot reference	1.1500 EUR-USD
Notional	1,000,000 EUR
Maturity	1 year
European style corridor	1.1000–1.18000 EUR-USD
Fixing schedule	monthly
Fixing source	ECB37
Premium	500,000 EUR

[5] Also exhibited on http://www.bloomberg.com/markets/currencies/fx-fixings

Composition and Applications Obviously, a European style corridor is a sum of digital call spread options (or equivalently digital put spread options). The only modification is that the expiration times are the fixing times and the delivery time is the same for all digital options and in fact deferred. Furthermore the digital payout in a corridor is usually fixing based, whereas a stand-alone digital may be exercised based on the usual NY or Tokyo cut.

Similarly, an American style corridor is a sum of double-barrier digitals with deferred delivery. We refer the reader to the exercises to work out the details.

Corridors occur very often as part of structured products such as a *range accrual forward* explained in Section 2.1.10 or a *corridor deposit* explained in Section 2.4.4.

1.8.3 Faders

Fader options are options whose nominal is directly proportional to the number of fixings inside or outside a pre-defined range. A *fade-in option* has a progressive activation of the nominal. In a *fade-out option* the concept of a progressive activation of the nominal is changed to a progressive deactivation. The term fader is sometimes used in a more general way to describe transactions with fade-in or fade-out notional amounts. These transactions do not need to be options but can be combinations of options or structured products. We discuss as an example the fade-in put option, whose characteristics are the pre-defined range and the associated fixing schedule with the maximal number of fixing being M. For each fixing date with the fixing inside the pre-defined range, the holder of a fade-in put option holds a contract, which at maturity is economically equivalent to a vanilla put option with the total notional multiplied by

$$\frac{\text{number of fixings inside the range}}{M}. \tag{365}$$

Buying a fade-in put option provides protection against falling EUR and allows full participation in a rising EUR. The holder has to pay a premium for this protection. He will typically exercise the option only if at maturity the spot is below the strike. The seller of the option receives the premium but is exposed to market movements and would need to hedge his exposure accordingly.

Advantages

- Protection against weaker EUR/stronger USD
- Premium not as high as for a plain vanilla put option
- Full participation in a favorable spot movement

Disadvantages

- Selling amount dependent on market movements between inception and maturity
- No guaranteed worst case exchange rate for the full notional

Example for the Computation of the Notional We explain this product with a EUR Put-USD Call with strike K, which has two ranges and six fixings, in Figure 1.62.

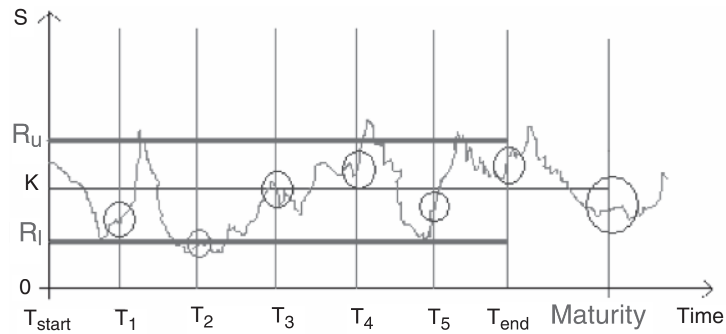

FIGURE 1.62 Notional of a fade-in put. At T_{end}, the holder would be entitled to sell $\frac{5}{6} \cdot 1$ M EUR, where 5 is the number of fixings between the lower and the upper level R_l and R_u on a resurrecting basis (here $n = 5$ because at T_2 the spot fixing is below the lower level). The total number of fixings inside the range will be known only at T_{end}. Hence, the notional of the put will be known only at T_{end}.

At maturity, the fade-in put works like a vanilla put. The holder would typically exercise the option and sell $\frac{5}{6} \cdot 1$ M EUR at the strike K if the spot is below the strike. If the spot ends up above the strike, the holder would let the option expire. For an investor client type, the maximum potential loss of the buyer is the fader's initial premium. For a treasurer client type, it can be worse in case he has used the fade-in put to protect the full notional of 1 M EUR. His worst case is that all fixings are outside the range. In this case he will have lost the initial premium and in the extreme worst case will receive nothing for 1 M EUR, assuming a limiting worst case of the EUR-USD spot going to infinity.

Example A company wants to hedge receivables from an export transaction in EUR due in 12 months' time. It expects a weaker EUR/stronger USD. The company wishes to be able to sell EUR at a higher spot rate if the EUR becomes stronger on the one hand, but on the other hand be protected against a weaker EUR. The company finds the corresponding vanilla EUR put/USD call too expensive and is prepared to take more risk. The treasurer believes that EUR/USD will not trade outside the range 1.1000–1.2000 for a significantly long time.

In this case a possible form of protection that the company can use is to buy a EUR fade-in put option, as presented in Table 1.34 for example.

- If the EUR-USD exchange rate is below the strike at maturity, then the company can sell EUR at maturity at the strike of 1.1600.
- If the EUR-USD exchange rate is above the strike at maturity, the company would let the option expire. However, the company will benefit from a higher spot when selling EUR.

The biggest risk is that all EUR-USD fixings are outside the range and the spot at maturity is low. In this case the company would need to sell EUR at the prevailing low market

TABLE 1.34 Example of a fade-in put. In comparison the corresponding vanilla put costs 50,000.00 EUR.

Spot reference	1.1500 EUR-USD	Strike	1.1600 EUR-USD
Company buys	EUR put USD call	Lower level	1.0000 EUR-USD
Fixing schedule	Monthly	Upper level	1.2000 EUR-USD
Maturity	1 year	Premium	EUR 6,000.00
Notional amount	EUR 1,000,000	Vanilla premium	EUR 50,000.00

TABLE 1.35 Example of a fade-in forward.

Spot reference	1.1500 EUR-USD	Strike	1.0000 EUR-USD
Company buys	EUR-USD forward	Lower level	1.0000 EUR-USD
Fixing schedule	Monthly	Upper level	1.1800 EUR-USD
Maturity	1 year	Premium	EUR 9,000.00
Notional amount	EUR 1,000,000	Fade-in call premium	EUR 27,000.00

spot price. Therefore, the company should have a risk policy in place that triggers an action as soon as EUR-USD drops below the lower level. Such an action could be to trade a forward contract for the rest of the unprotected notional to prevent further losses. The company should be aware that this is not a buy-and-hold strategy.

Variations Besides puts, there are fade-in calls or fade-in forwards, see Table 1.35 or the live trade in Table 2.3 in Section 2.1.4. Also more exotic types of faders can be created by taking exotic transactions and let them fade in or out.

Faders often have an additional knock-out range just like corridors, see Section 1.8.2. One then classifies faders into *resurrecting*, *non-resurrecting, keeping the accrued amount*, and *non-resurrecting losing parts or all of the accrued amount*. Faders are most popularly applied in structuring *accumulative forwards*, see Section 2.1.11.

1.8.4 Exotic Barrier Options

Digital Barrier Options Just like barrier options, which are calls or puts with knock-out or knock-in barriers, one can consider digital calls and puts with additional American style knock-out or knock-in barriers. Knowing the digitals, we can derive the knock-in digitals from the knock-out digitals. The knock-out digitals can be viewed as the limiting case of tight knock-out call spread with very high notionals and can to some extent be approximated by those – with the usual practical limitations. It is clear from this approximation that Greeks are expected to take extreme values and change their signs in different regions of spot and time. The motivation for such products is to reduce the cost in taking a view in a market event.

Window Barriers Barriers need not be active for the entire lifetime of an option. Window barrier options are extensions of (digital) call or put options with barriers where the

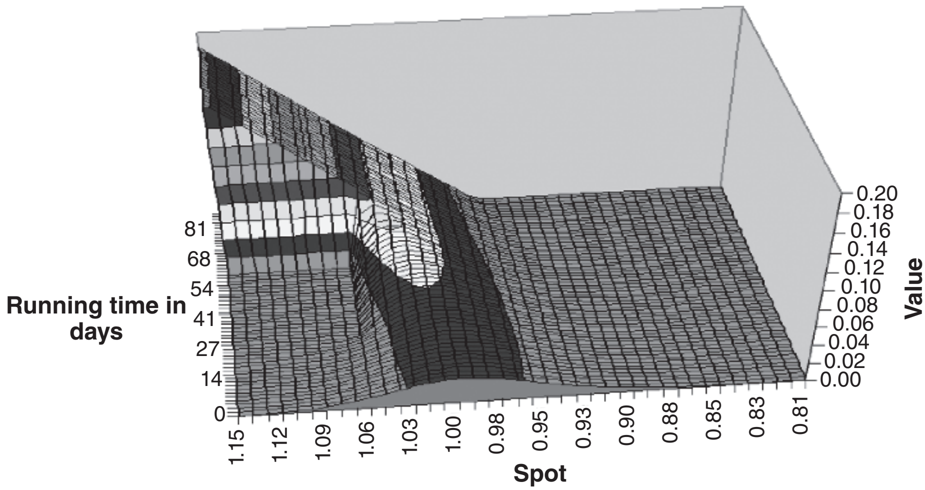

FIGURE 1.63 Value function $v(t,x)$ of an up-and-out call option with window barrier active only for the second month, with strike $K = 0.9628$, knock-out barrier $B = 1.0590$, and maturity 3 months. We used the interest rates $r_d = 6.68\%$, $r_f = 5.14\%$, volatility $\sigma = 11.6\%$, and $R = 0$.

barriers are active during a period of time which is shorter than the whole lifetime of the option – for example only the first three months from an option with six months' maturity. One can specify arbitrary time ranges with piecewise constant barrier levels or even non-constant barriers. Window barrier options are the most exotic options that still fit and can be confirmed with the definitions of the *ISDA 2005 Barrier Option Supplement* [78]. See Figure 1.63 for the value function of a window barrier option. Linear and exponential barriers are useful if there is a high drift in the exchange rate caused, for example by a high interest rate differential (high swap points).

Step and Soft Barriers In case of a knock-out event, a client might argue: "Come on, the spot only crossed the barrier for a very short moment, can't you make an exception and not let my option knock out?" This is a very common concern: how to get protection against price spikes. Such a protection is certainly possible, but surely has its price. One way is to measure the time the spot spends opposite the knock-out barrier and let the option knock out gradually. For instance, one could agree that the option's nominal decreases by 10% for each day the exchange rate fixing is opposite the barrier. Barriers can be constant, linear, or exponential functions of time. Continuously observed barrier contracts of this type are referred to as *occupation time derivatives*. There are even closed form solutions for the value of occupation time derivatives in the Black-Scholes model and some jump-diffusion models. While the occupation time of a (geometric) Brownian motion at a specific level is well understood and defined, it would be difficult, if not impossible, to specify a continuous occupation time of a traded spot. Consequently, occupation time derivatives do not trade. The discrete version of these are called faders and are explained in Section 1.8.3.

Fluffy Barriers Protection against price spikes can be achieved by having a spot spend time *sufficiently long* beyond the barrier. Another way to define a knock-out event is based on the spot going *sufficiently far* beyond the barrier. This feature can be structured by a fluffy barrier contract, where a payoff is kept constant, but the notional to which the payoff is applied depends on how far the spot goes beyond a pre-specified barrier. For instance, with a knock-out on the upside, one can specify a first and second barrier level and let the notional be proportional to the ratio of the difference of level 2 and the maximum spot and the difference of level 2 and level 1. If the spot goes beyond level 1 and also reaches level 2, then the total notional and the entire contract terminate worthless. If the spot does not hit level 1, then there is no knock-out. If the spot goes through level 1 and then to the middle of level 1 and level 2, then 50% of the notional of the contract is gone. The more common version is one with a discrete gradual knock-out: for instance one can specify a barrier range of 2.20 to 2.30 where the option loses 25% of its nominal when 2.20 is breached, 50% when 2.25 is breached, 75% when 2.275 is breached, and 100% when 2.30 is breached.

Parisian and Parasian Barriers Another way to get price spike protection is to let the option knock out only if the spot spends a certain pre-specified length of time opposite the barrier – either in total (Parasian) or in a row (Parisian). Clearly the plain barrier option is the least expensive, followed by the Parasian, then the Parisian barrier option, and finally the corresponding vanilla contract. See Figure 1.64. The name Parisian probably originates from Société Générale trading such contracts in Paris. However, there are other theories. Valuation is typically done by Monte Carlo simulation, although there are PDE methods available. The Parasian knock-out event is equivalent to a *counter*-based early termination in a target forward, see Section 2.2.3.

Resettable Barriers This is a way to give the holder of a barrier option a chance to reset the barrier during the life of the option n times at pre-specified N decision times in the future ($N \geq n$). This kind of extra protection also makes the barrier option more

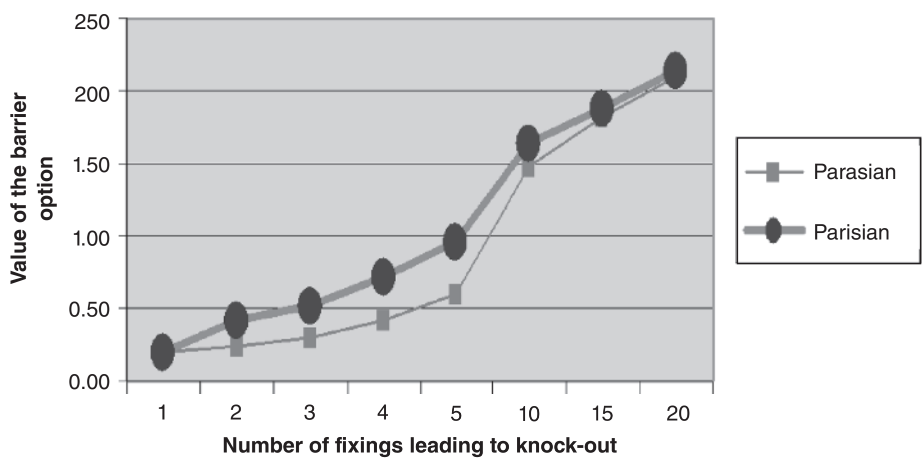

FIGURE 1.64 Comparison of Parisian and Parasian barrier option values.

expensive. Obviously, this contract requires many specifications, in particular to which new barrier level a barrier can be reset. A common way to do it is to reset the barrier to x% beyond the current spot fixing, but there can be more complex reset functions or even rights of choice for the holder. Resettable barriers occur mostly in complex retail certificates, but are also used as building blocks for corporate hedging strategies.

Transatlantic Barrier Options For transatlantic barrier options one barrier is of American style, the other one of European style. Naturally, the European style barrier is in-the-money, the American style barrier usually out-of-the-money. Therefore, there are essentially two versions,

1. a call with strike K, a European style up-and-out $H > K$, and an American style down-and-out at $L \leq K$,
2. a put with strike K, a European style down-and-out $L < K$, and an American style up-and-out $H \geq K$.

The motivation for such products is of course the lower cost in comparison with vanilla or single barrier options on the one hand and the fear of price spikes and a resulting preference for European style barriers on the other hand.

The pricing and hedging is comparatively easy provided we have regular and digital barrier options available as basic products. Then we can structure the transatlantic barrier option just like in Equation (23), with an additional out-of-the-money knock-out barrier.

Knock-In-Knock-Out Options Knock-In-Knock-Out Options (KIKOs) are barriers with both a knock-out and a knock-in barrier. However, it is not as simple, because there are three fundamentally different types:

1. The knock-out can happen *any time*.
2. The knock-out can happen only *after* the knock-in.
3. The knock-out can happen only *before* the knock-in.

The first one is the market standard, but when dealing one should always clarify which type of knock-in-knock-out is agreed upon. For example, let the lower barrier L be a knock-out barrier and the upper barrier H be a knock-out barrier. Standard type 1 KIKO can be exercised only if L is never touched *and* H has been touched at least once. This can be replicated by standard barrier options via

$$\text{KIKO}(L, H) = \text{KO}(L) - \text{DKO}(L, H). \tag{366}$$

Therefore, pricing and hedging of this KIKO is no more complicated than pricing and hedging of knock-out options.

The second type is a special case of a *knock-in on strategy* contract. Any structure can be equipped with a *global* knock-in barrier, that has to be touched before the structure becomes alive. Knock-out events in the structure are active only *after* the structure knocks in. This is a product of its own and requires an individual valuation, pricing, and hedging approach.

In the third type of KIKO a knock-out can happen only before the knock-in. Once the option is knocked in, the knock-out barrier is no longer active. This is also a product of its own and requires an individual valuation, pricing, and hedging approach.

James Bond Range As James Bond can only live twice, the *James Bond range* is a double-no-touch type contract. Given an upper barrier H and a lower barrier L, it pays one unit of currency if the spot remains inside (L, H) at all times until expiry T, or if the spot hits L the spot thereafter remains in a new range to be set around L, or similarly if the spot hits H the spot thereafter remains in a new range to be set around H. The contract is also called *tolerant double-no-touch*.

1.8.5 Pay-Later Options

A pay-later option is a vanilla option whose premium is paid only if the option is in-the-money or is exercised at the expiration time. If the spot is not in-the-money, the holder of the option would normally not exercise the option, and will end up not having paid anything. However, if the spot is in-the-money, the holder of the option has to pay the option premium, which will then be noticeably higher than the plain vanilla. For this reason pay-later options are not traded very often. Note that the payment of the premium is conditional. The pay-later option is not to be confused with a vanilla option whose premium is (unconditionally) deferred to its maturity date.

Advantages

- Full protection against spot market movement
- Premium is paid only if the option ends up in-the-money
- Premium is paid only at maturity

Disadvantages

- More expensive than a plain vanilla
- Credit risk for the seller as payoff can be negative

The Valuation for the Pay-Later Option The payoff of a pay-later option is defined as

$$\left[\phi(S_T - K) - P\right] \mathbb{I}_{\{\phi S_T \geq \phi K\}} \tag{367}$$

and illustrated in Figure 1.65. As usual, the binary variable ϕ takes the value +1 for a call and −1 for a put, K the strike in units of the domestic currency, and T the expiration time in years. The *price* P of the pay-later option is paid at time T, but it is set at time zero in such a way that the time zero *value* of the above payoff is zero. Carefully notice the difference between price and value. After the option is written, the price P does not change any more.

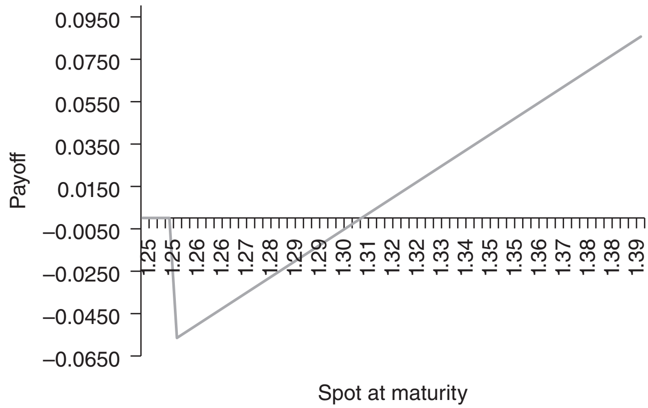

FIGURE 1.65 Payoff of a pay-later EUR call USD put. We use the market input spot $S_0 = 1.2000$, volatility $\sigma = 10\%$, EUR rate $r_f = 2\%$, USD rate $r_d = 2.5\%$, strike $K = 1.2500$, time to maturity $T = 0.5$ years. The vanilla value is 0.0158 USD, the digital value is 0.2781 USD, the resulting pay-later price is 0.0569 USD, which is substantially higher than the plain vanilla value. Consequently the break-even point is at 1.3075, which is quite far off. For this reason pay-later type structures do not trade very often.

We denote the current spot by x and the current time by t and define furthermore the abbreviations

$$n(t) \stackrel{\Delta}{=} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2}, \tag{368}$$

$$\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\, dt, \tag{369}$$

$$\tau \stackrel{\Delta}{=} T - t, \tag{370}$$

$$f = xe^{(r_d - r_f)\tau}, \tag{371}$$

$$d_\pm \stackrel{\Delta}{=} \frac{\log \frac{f}{K} \pm \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}, \tag{372}$$

$$\text{vanilla}(x, K, T, t, \sigma, r_d, r_f, \phi) = \phi e^{-r_d\tau}[f\mathcal{N}(\phi d_+) - K\mathcal{N}(\phi d_-)], \tag{373}$$

$$\text{digital}(x, K, T, t, \sigma, r_d, r_f, \phi) = e^{-r_d\tau}\mathcal{N}(\phi d_-). \tag{374}$$

The formulas of vanilla and digital options have been derived in Sections 1.4 and Section 1.7.2 respectively.

The payoff can be rewritten as

$$\left[\phi(S_T - K)\right]^+ - P\mathbb{I}_{\{\phi S_T \geq \phi K\}}, \tag{375}$$

whence the value of the pay-later option in the Black-Scholes model

$$dS_t = S_t \left[(r_d - r_f)dt + \sigma \, dW_t\right] \tag{376}$$

is easily read as

$$\begin{aligned} \text{paylater}(x, K, P, T, t, \sigma, r_d, r_f, \phi) = \text{vanilla}(x, K, T, t, \sigma, r_d, r_f, \phi) \\ - P \cdot \text{digital}(x, K, T, t, \sigma, r_d, r_f, \phi). \end{aligned} \tag{377}$$

In particular, this leads to a quick implementation of the value and all the Greeks having the functions vanilla and digital at hand. This relationship between the pay-later option, the vanilla, and the digital is generic, and not model dependent.

For the *pay-later price* setting

$$\text{paylater}(x, K, P, T, 0, \sigma, r_d, r_f, \phi) = 0 \tag{378}$$

yields

$$P = \frac{\text{vanilla}(x, K, T, 0, \sigma, r_d, r_f, \phi)}{\text{digital}(x, K, T, 0, \sigma, r_d, r_f, \phi)} \tag{379}$$

$$= \text{vanilla}(x, K, T, 0, \sigma, r_d, r_f, \phi) \frac{e^{r_d T}}{\mathcal{N}(\phi d_-)}. \tag{380}$$

This can be interpreted as follows. The value P is like the value of a vanilla option, except that

- we must pay interest $e^{r_d T}$, since the premium is due only at time T and
- the premium needs to be paid only if the option is in-the-money or is exercised, which is why we divide by the (risk-neutral) probability that the option is exercised $\mathcal{N}(\phi d_-)$.

We observe that the pay-later option can be viewed as a *structured product*. All we need are vanilla and digital options. The structurer will easily replicate a short pay-later with a long vanilla and a short digital. We learn that several types of options can be composed from existing ones, which is the actual job of structuring. This way it is also straightforward to determine a market price, given a vanilla market.

Variations Pay-later options are an example of the family of contingent or deferred payment options. We can also simply defer the payment of a vanilla without any conditions on the moneyness. Another variation is paying back the vanilla premium if the spot stays inside some range, see the exercises in Section 2.1.24. Naturally, the pay-later effect can be extended beyond vanilla options to all kind of options.

1.8.6 Step Up and Step Down Options

The step option is an option where the strike of the option is readjusted at predefined fixing dates, but only if the spot is more favorable than that of the previous fixing date. The step option can either be a plain vanilla option or a single barrier option. The concept of a progressive *step up* or *step down* could be changed also to a progressive *step up* or *step down* for a forward rate.

1.8.7 Options and Forwards on the Harmonic Average

Consider a schedule of observation times $t_1, \dots, t_n$ of some underlying. Options and forwards on the arithmetic average

$$\frac{1}{n}\sum_{i=1}^{n} S(t_i) \tag{381}$$

have been analyzed and traded for some time, see Section 1.7.6. The geometric average

$$\sqrt[n]{\prod_{i=1}^{n} S(t_i)} \tag{382}$$

has often been used as control variate for pricing the arithmetic average, whose distribution in a multiplicative model like Black-Scholes is cumbersome to deal with. The *harmonic average*

$$\frac{n}{\sum_{i=1}^{n} \frac{1}{R(t_i)}} \tag{383}$$

comes up if a client wants to exchange an amount of *domestic* currency into the *foreign* currency at an average rate of the currency pair FOR-DOM, e.g. wants to exchange USD into EUR at a rate which is an average of observed EUR-USD rates. In this case USD is the domestic currency and we need to actually look at the exchange rate of $R = 1/S$ in DOM-FOR quotation in order to allow the domestic currency as a notional amount. As in the case of standard Asian contracts, there can be forwards and options on the harmonic average, both with fixed and floating strike. We treat one possible example in the next section. Options and forwards on the harmonic average are also sometimes referred to as *Australian* derivatives, possibly because Australia is on the other side of the equator. The fact that a harmonic average is used is often not explicitly shown in the systems. A domestic notional amount and cash settlement in the foreign currency are strong indications that harmonic averaging is performed behind the curtains. The valuation of harmonic average contracts can be related to arithmetic average contracts by a clever change of numeraire, see Večeř [134].

Harmonic Asian Swap We consider a EUR-USD market with spot reference 1.0070, swap points for time T_1 of −45, swap points for time $T_2 > T_1$ of −90. As a contract specification, the client buys N USD at the daily average of the period of one month before T_1, denoted by A_1. Then the client sells N USD at the daily average of a period of one month before T_2, denoted by A_2. The payoff in EUR of this structure (cash settled two business days after T_2) is

$$\frac{N}{A_2} - \frac{N}{A_1}. \tag{384}$$

To replicate this using the fixed strike Asian forward we can decompose it as follows:

1. We sell to the client the payoff $1 - \frac{1}{A_1}$ (using strike 1 by default) with notional N.
2. We buy from the client the payoff $1 - \frac{1}{A_2}$ (using strike 1 by default) with notional N.

On a notional of $N = 5$ M USD this could have a theoretical value of 23,172 EUR. This is what the sell-side should charge the client in addition to overhedge and sales margin. One problem is that the structure is very transparent for the client. If we take the forward for mid February, we have −45 swap points, for mid June −90 swap points. This means that the client would know that in a first order approximation he owes the bank 45 swap points, which is

$$5 \text{ M USD} \cdot 0.0045 = 22{,}500 \text{ EUR}.$$

If the swap ticket requires entering a strike, one can use 1.0000 in both tickets, but this value does not influence the value of the swap.

1.8.8 Variance and Volatility Swaps

A variance swap is a contract that pays the difference of a pre-determined fixed variance (squared volatility), which is usually determined in such a way that the trading price is zero, and a realized historic annualized variance, which can be computed only at maturity of the trade. Therefore, the variance swap is an ideal instrument to hedge volatility exposure, a need for funds and institutional clients. Of course, one can hedge vega with vanilla options or straddles, but this is then also subject to spot movements and time decay of the hedge instruments. The variance swap also serves as a tool to take a view on volatility. Variance and volatility swaps were standardized by ISDA in 2013 [79].

Advantages

- Insurance against changing volatility levels
- Independence of spot
- Zero-cost product
- Fixed volatility (break-even point) easy to approximate as average of smile

Disadvantages

- Difficult to understand
- Many details in the contract to be set
- Variance harder to capture than volatility
- Volatility swaps harder to price than variance swaps

Example Suppose the one-month implied volatility for EUR/USD at-the-money options is close to its one-year historic low. This can easily be noticed by looking at *volatility cones*, see Section 1.5.10. Suppose further that you are expecting a period of higher volatility during the next month. You are looking for a zero-cost strategy, where you would profit from a rising volatility, but you are ready to encounter a loss otherwise. In this case a suitable strategy to trade is a variance or volatility swap. We consider an example of a variance swap in Table 1.36.

To make this clear we consider the following two scenarios with possible fixing results listed in Table 1.37 and Figure 1.66.

- If the realized variance is 0.41% (corresponding to a volatility of 6.42%), then the market was quieter than expected and you need to pay 10 M USD · (0.85% − 0.41%) = 44,000 USD.
- If the realized variance is 1.15% (corresponding to a volatility of 10.7%), then your market expectation turned out to be correct and you will receive 10 M USD · (1.15% − 0.85%) = 30,000 USD.

A volatility swap trades

$$\sqrt{\frac{B}{N-1}\sum_{i=1}^{N}(r_i-\bar{r})^2} \tag{385}$$

TABLE 1.36 Example of a variance swap in EUR-USD. The quantity r_i is called the log-return from fixing day $i-1$ to day i and the average log-return is denoted by $\bar{r}$. The notation %% means a multiplication with 0.0001. It is also sometimes denoted as $\%^2$.

Spot reference	1.0075 EUR-USD
Notional M	USD 10,000,000
Start date	19 November 2002
Expiry date	19 December 2002
Cash settlement	23 December 2002
Fixing period	Every weekday from 19-Nov-02 to 19-Dec-02
Fixing source	ECB fixings $F_0, F_1, \ldots, F_N$
Number of fixing days N	23 (32 actual days)
Annualization factor B	262.3 = 23/32·365
Fixed strike K	85.00%% corresponding to a volatility of 9.22%
Payoff	$M \cdot$ (realized variance − K)
Realized variance	$\frac{B}{N-1}\sum_{i=1}^{N}(r_i-\bar{r})^2; \bar{r}=\frac{1}{N}\sum_{i=1}^{N} r_i; r_i=\ln\frac{F_i}{F_{i-1}}$
Premium	none

TABLE 1.37 Example of two variance scenarios in EUR-USD. The left column shows a possible fixing set with a lower realized variance, the right column a scenario with a higher variance.

Date	Fixing (low vol)	Fixing (high vol)	Date	Fixing (low vol)	Fixing (high vol)
19/11/02	1.0075	1.0075	6/12/02	0.9953	1.0037
20/11/02	1.0055	1.0055	9/12/02	0.9966	0.9962
21/11/02	1.0111	1.0111	10/12/02	0.9986	0.9986
22/11/02	1.0086	1.0086	11/12/02	1.0003	0.9907
25/11/02	1.0027	1.0027	12/12/02	0.9956	1.0018
26/11/02	1.0019	1.0067	13/12/02	0.9981	1.0000
27/11/02	1.0033	0.9997	16/12/02	0.9963	0.9963
28/11/02	1.0096	1.0113	17/12/02	1.0040	1.0040
29/11/02	1.0077	1.0062	18/12/02	1.0045	1.0017
2/12/02	1.0094	1.0094	19/12/02	1.0085	1.0114
3/12/02	1.0029	0.9999	variance	0.41%	1.15%
4/12/02	1.0043	1.0043	volatility	6.42%	10.70%
5/12/02	0.9977	0.9977			

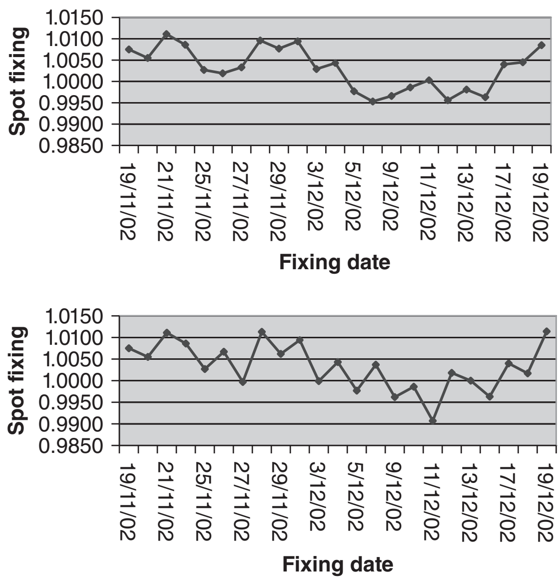

FIGURE 1.66 Comparison of scenarios for a low variance (graph above) and a higher variance (graph below).

against a fixed volatility, which is usually determined in such a way that the trading price is zero. Since the square root is not a linear function of the variance, this product is more difficult to price than a standard variance swap. For details on pricing and hedging we refer to Carr and Lee [23], as well as the paper on *more than you ever wanted to know about volatility swaps* [37]. As a rule of thumb, the fixed variance or volatility to make the contract worth zero is the average of the volatilities in the volatility smile matrix for the maturity under consideration as for the variance swap there exists a static hedging portfolio consisting of vanilla options with the same maturity. In principle, the replicating portfolio of a variance swap consists of a portfolio of vanilla options approximating the payoff $\ln S_T$. This requires options with arbitrary far away strikes. Therefore, the approximation is quite sensitive to implied volatilities for very small and very big deltas. The variance swap price can be viewed as an indicator of how an FX options desk handles the extrapolation of the volatility smile on the far wings. Aggressive market makers tend to ignore the extreme strikes. In fact, one may ask how a variance swap differs from its replication. Options in the replication with extremely far away strikes would be required if there is a jump in the underlying exchange rate. The correctly priced variance swap would quantify the jump risk, whereas the replication for all practical matters would not, simply because options with extremely far away strikes are not tradable, or market prices are not available. Note that in FX markets, other than in equity markets, variance and volatility swaps normally do not have caps.

Forward Variance Swap In a standard variance swap, the first spot fixing is at inception of the trade or two business days thereafter. However, there may be situations where a client needs to hedge a forward volatility exposure that originates from a compound, installment, forward start, cliquet, or other exotic option with a significant forward volatility dependence. We will illustrate now how to structure a forward variance swap, where the first fixing is at some time in the future, using standard variance swaps. Let there be J fixings in the initial period and M fixings in the second period. The total number of fixings is hence $M+J$. We can then split the payoff

$$\frac{B}{M-1}\sum_{i=J+1}^{J+M}(r_i-\overline{r})^2-K \tag{386}$$

into the two parts

$$\begin{aligned}
&\frac{B}{M-1}\sum_{i=1}^{J+M}(r_i-\overline{r})^2-K-\left[\frac{B}{M-1}\sum_{i=1}^{J}(r_i-\overline{r})^2-0\right]\\
&=\frac{C}{J+M-1}\sum_{i=1}^{J+M}(r_i-\overline{r})^2-K-\left[\frac{D}{J-1}\sum_{i=1}^{J}(r_i-\overline{r})^2-0\right]
\end{aligned} \tag{387}$$

and find as the only solution for the numbers C and B

$$\begin{aligned}
C&=\frac{(J+M-1)B}{M-1},\\
D&=\frac{(J-1)B}{M-1}.
\end{aligned} \tag{388}$$

Modifications When computing the variance of a random variable X whose mean is small, we can take the second moment $\mathbb{E}X^2$ as an approximation of the variance

$$\mathbf{var}(X) = \mathbb{E}X^2 - (\mathbb{E}X)^2. \tag{389}$$

Following this idea and keeping in mind that the average of log-returns of FX fixings is indeed often close to zero, the variance swap is sometimes understood as a second moment swap rather than an actual variance swap. To clarify, traders specify in their dialogue whether the product is *mean subtracted* or not. We have presented here the variance swap with the mean subtracted.

1.8.9 Forward Volatility Agreements (FVAs)

Besides variance and volatility swaps, the forward volatility agreement is another instrument to trade volatility. In this case the forward volatility is the focus. We have already seen in Section 1.7.8 that the value of a forward start option essentially depends on the forward volatility σ_f, the one that applies between fixing time T_f in the future and the maturity time T_e. Equation (152) establishes the relationship between the forward volatility and the spot volatilities σ_{T_f} and σ_{T_e}. In a forward volatility agreement party A agrees to buy a strategy or a single option from party B. The strategy will start on the future fixing time T_f. However, the premium to pay for the strategy on the fixing value date T_{fd} is calculated with a fixed forward volatility σ_f. The prevailing spot volatility on the fixing date will most likely be different. Party A locks in the volatility using the forward volatility agreement.

Typical strategies to trade are ATM call and put vanilla options, but the most common is the ATM straddle (see Section 1.6.3). The reason to trade ATM strategies is to focus on the ATM volatility and its view on the term structure, while keeping smile effects aside. As ATM volatility, the usual FX conventions apply: delta-neutral straddle strikes $K_\pm$ as in Equation (43). Forward volatility agreements typically trade at the *par volatility* σ_f as in Equation (152). Essentially party A pays a par-vol premium for a future market spot vol product. Therefore, if the future market spot volatility turns out to be higher than the par volatility at inception, party A makes a profit; conversely, if market spot volatility turns out to be lower, then party A makes a loss. This transaction makes the forward volatility indirectly tradable.

Traded FVA Example We consider the terms of a traded forward volatility agreement in Table 1.38. Essentially, one can view a forward volatility agreement as a re-branding of the forward start option or combination of forward starting options as a straddle. In fact, it is common that forward volatility agreement is confirmed as a *forward setting currency option transaction* or *forward start straddle*. The difference is that usually, in a forward start option, the premium is paid at the horizon spot date, whereas in a forward volatility agreement the premium is paid at the fixing spot date. Generally, forward start options allow more flexibility in setting the strike different from ATM, whereas forward volatility agreements typically work with the ATM strike. Another difference is that in a forward volatility agreement, the premium will be calculated with the prevailing market data (spot, interest, and forward rates) and the pre-specified fixed volatility, whereas in a forward start option, the current forward interest rates are used.

TABLE 1.38 Example of a traded forward volatility agreement in EUR/GBP.

Trade date	29 November 2007
Strategy	European EUR Call/GBP Put and EUR Put/GBP Call
Notional amount	EUR 365,000,000.00
Business days	London & Any TARGET Settlement day
Business day Convention	Following
Expiration date	26 November 2008
Expiration time	10.00 a.m. (local time in New York)
Settlement date	28 November 2008
Strike price	delta neutral with fixed volatility of 7.30%, the prevailing EUR/GBP Spot Rate, Forward and Deposit Rate at the Fixing Time on the Fixing Date, as determined by the Calculation Agent in its sole discretion
Delta neutral straddle	the Strike Price where the Premium Currency delta of the Call and the Premium Currency delta of the Put are equal and opposite and thus sum to zero
Fixing time	10.00 a.m. (local time in New York)
Fixing date	28 May 2008
Fixing rate	Spot Rate at the Fixing Time on the Fixing Date
Premium	to be determined by using 7.30% Volatility
Premium payment date	30 May 2008

GBP/USD Forward Volatility Agreement Example We consider an example in the version of an ATM straddle. Market data (possibly on trade date T_h 7 April 2016): spot $S_0 = 1.4000$, spot date T_s 7 April 2016, 6 M ATM volatility $\sigma_{T_f} = 13.736\%$, 6 M GBP money market $r_f = 0.482\%$, 6 M USD money market $r_d = 0.667\%$, 6 M GBP-USD forward rate $F_{T_f} = 1.40136$, 6 M RR −4.377% (favoring GBP puts), 6 M BF 0.395%, 12 M ATM volatility $\sigma_{T_e} = 12.866\%$, 12 M GBP money market $r_f = 0.439\%$, 12 M USD money market $r_d = 0.731\%$, 12 M GBP-USD forward rate $F_{T_f} = 1.40421$, 12 M RR −4.055% (favoring GBP puts), 12 M BF 0.435%. Contract data: expiry T_e 5 April 2017, delivery T_d 7 April 2017, fixing data T_f 5 Oct 2016 (183 days), fixing spot date T_{fd} 7 Oct 2016 (365 days), notional GBP 1 M. The par volatility is

$$\sigma_f = \sqrt{\frac{\sigma_{T_e}^2(T_e - T_s) - \sigma_{T_f}^2(T_f - T_s)}{T_e - T_f}} = \sqrt{\frac{12.866\%^2 365 - 13.736\%^2 183}{182}} = 11.93\%. \tag{390}$$

A bid and offer par volatility could be 11.25%–12.60%, assuming a bid offer spread of 1.30%. The working spot volatility spreads are 0.533% for 6 M and 0.417% for 12 M. In this case the buyer would pay a premium calculated for a straddle with strike K_+ and the prevailing money market and spot rates at the fixing date and a volatility of 12.60%. The exact premium is not known at the trade date. The buyer then holds a straddle. With a fixed volatility of 8% (instead of the par volatility), the initial bid and offer price would be GBP 18,400–26,000. Obviously, the FVA has no delta or gamma before the fixing date. Initial vega would be GBP 5,638.

1.8.10 Exercises

European Style Corridor Starting with the value for digital options, derive exactly the value of a European style corridor in the Black-Scholes model. Discuss how to find a market price based on the market of vanilla options.

Fade-Out Call How would you structure a *fade-out call* that starts with a nominal amount of M? As the exchange rate evolves, the notional will be decreased by $\frac{M}{N}$ for each of the N fixings that is outside a pre-defined range.

Fader Payoff As for the corridors in Section 1.8.2, write down the exact payoff formulas for the various variations of faders in Section 1.8.3.

Fade-In Forward Describe a possible client view that could lead to trading a fade-in forward in Table 1.35.

Static Replication of the Tolerant Double-No-Touch As a variation of the James Bond range in Section 1.8.4, we consider barriers A, B, C, D as illustrated in Figure 1.67.

A rather *tolerant double-no-touch* knocks out after the second barrier is touched or crossed. How would you replicate it statically using standard barrier and touch contracts?

Pay-Later with Premium in Foreign Currency The pay-later value in Equation (379) is measured in units of domestic currency. Does this change if the premium is specified to be paid in foreign currency? If no, argue why. If yes, specify how.

Pay-Later Digital Derive the pay-later value of a digital option.

Pay-Later Call Spread Derive the pay-later value of a call spread.

Pay-Later Up-and-Out Call How would you structure an up-and-out call whose premium is paid only if the spot is in-the-money at the expiration time and if it has not knocked out?

Chooser Option A *chooser option* lets the buyer decide at expiration time whether he wants to exercise a call with strike K or a put with strike K. Discuss how to find a market price and how to statically hedge it. (Hint: straddle.) Moreover, if the decision as to which of the options to take is made at time t strictly before the expiration time T, how would you price and hedge the chooser?

FIGURE 1.67 Nested double-no-touch ranges.

1.9 SECOND GENERATION EXOTICS (MULTIPLE CURRENCY PAIRS)

There are also a number of derivatives involving multiple currency pairs. They are referred to as *multi-currency* or *rainbow* derivatives. Some serve as FX hedging instruments for corporates, others occur as hedging instruments for institutions and hedge funds, and some in private banking.

1.9.1 Spread and Exchange Options

A spread option compensates a spread in exchange rates and pays off

$$\left[\phi\left(aS_T^{(1)} - bS_T^{(2)} - K\right)\right]^+. \tag{391}$$

This is a European spread put ($\phi = -1$) or call ($\phi = +1$) with strike $K > 0$ and the expiration time in years T. We assume without loss of generality that the weights a and b are positive. These weights are needed to make the two exchange rates comparable, as USD-CHF and USD-JPY differ by a factor of the size of 100. A standard for the weights are the reciprocals of the initial spot rates, i.e. $a = \frac{1}{S_0^{(1)}}$ and $b = \frac{1}{S_0^{(2)}}$.

Spread options are not traded very often in FX markets. If they are they are usually cash-settled. Exchange options come up more often as they entitle the owner to exchange one currency for another, which is very similar to a vanilla option, which is reflected in the valuation formula.

In the two-dimensional Black-Scholes model

$$dS_t^{(1)} = S_t^{(1)}\left[\mu_1 dt + \sigma_1 dW_t^{(1)}\right], \tag{392}$$

$$dS_t^{(2)} = S_t^{(2)}\left[\mu_2 dt + \sigma_2 dW_t^{(2)}\right], \tag{393}$$

$$\mathbf{Cov}\left[W_t^{(1)}, W_t^{(2)}\right] = \rho t, \tag{394}$$

with positive constants σ_i denoting the annual volatilities of the i-th foreign currency, ρ the instantaneous correlation of their log-returns, r the domestic risk-free rate, and risk-neutral drift terms

$$\mu_i = r - r_i, \tag{395}$$

where r_i denotes the risk-free rate of the i-th foreign currency, the value is given by (see [104])

$$\text{spread} = \int_{-\infty}^{+\infty} \text{vanilla}\left(S(x), K(x), \sigma_1\sqrt{1-\rho^2}, r, r_1, T, \phi\right) n(x)\, dx \tag{396}$$

$$S(x) \stackrel{\Delta}{=} aS_0^{(1)} e^{\rho\sigma_1\sqrt{T}x - \frac{1}{2}\sigma_1^2\rho^2 T} \tag{397}$$

$$K(x) \stackrel{\Delta}{=} bS_0^{(2)} e^{\sigma_2\sqrt{T}x + \mu_2 T - \frac{1}{2}\sigma_2^2 T} + K. \tag{398}$$

Notes

1. The integration can be done numerically, e.g. using the Gauß-Legendre algorithm with integration limits −5 and 5. The function vanilla (European put and call) can be found in Section 1.4.
2. The integration can be done analytically if $K = 0$. This is the case of *exchange options*, the right to exchange one currency for another.
3. To compute Greeks one may want to use homogeneity relations as discussed in [107].
4. In a foreign exchange setting, the correlation can be computed in terms of known volatilities. This can be found in Section 1.9.2.

Derivation of the Value Function We use Equation (7) for the value of vanilla options along with the abbreviations thereafter.

We rewrite the model in terms of independent new Brownian motions $W^{(1)}$ and $W^{(2)}$ and get

$$S_T^{(1)} = S_0^{(1)} \exp\left[(\mu_1 - \tfrac{1}{2}\sigma_1^2)T + \sigma_1 \rho W_T^{(2)} + \sigma_1 \sqrt{1-\rho^2}\, W_T^{(1)}\right], \tag{399}$$

$$S_T^{(2)} = S_0^{(2)} \exp\left[(\mu_2 - \tfrac{1}{2}\sigma_2^2)T + \sigma_2 W_T^{(2)}\right]. \tag{400}$$

This allows us to write $S_T^{(1)}$ in terms of $S_T^{(2)}$, i.e.

$$S_T^{(1)} = \exp\left[\hat{\mu}_1 + \frac{\sigma_1 \rho}{\sigma_2}(\ln S_T^{(2)} - \hat{\mu}_2) + \sigma_1 \sqrt{1-\rho^2}\, W_T^{(1)}\right], \tag{401}$$

$$\hat{\mu}_i \stackrel{\Delta}{=} \ln S_0^{(i)} + (\mu_i - \tfrac{1}{2}\sigma_i^2)T, \tag{402}$$

which shows that given $S_T^{(2)}$, $\ln S_T^{(1)}$ is normally distributed with mean and variance

$$\mu = \hat{\mu}_1 + \frac{\sigma_1 \rho}{\sigma_2}(\ln S_T^{(2)} - \hat{\mu}_2), \tag{403}$$

$$\sigma^2 = \sigma_1^2(1-\rho^2)T. \tag{404}$$

We recall from the derivation of the Black-Scholes formula for vanilla options that (and in fact, for $\rho = 0$ this *is* the Black-Scholes formula)

$$\begin{aligned} &\mathbb{E}[(\phi(S_T^{(1)} - K))^+] \\ &= \phi\left[e^{\mu + \frac{\sigma^2}{2}} \mathcal{N}\left(\phi \frac{-\ln K + \mu + \sigma^2}{\sigma}\right) - K\mathcal{N}\left(\phi \frac{-\ln K + \mu + \sigma^2}{\sigma}\right)\right], \end{aligned} \tag{405}$$

which allows to compute the value of a spread option as

$$e^{-rT}\mathbb{E}[(\phi(aS_T^{(1)} - bS_T^{(2)} - K))^+] \tag{406}$$

$$= a\mathbb{E}\left[e^{-rT}\mathbb{E}\left[\left(\phi(S_T^{(1)} - (\frac{b}{a}S_T^{(2)} + \frac{K}{a}))\right)^+ \middle| S_T^{(2)}\right]\right] \tag{407}$$

$$= a \cdot \mathbb{E}\left[\text{vanilla}\left(S_0^{(1)} \exp\left\{\frac{\sigma_1\rho}{\sigma_2}(\ln S_T^{(2)} - \hat{\mu}_2) - \frac{1}{2}\sigma_1^2\rho^2 T\right\},\right.\right.$$
$$\left.\left.\frac{b}{a}S_T^{(2)} + \frac{K}{a}, \sigma_1\sqrt{1-\rho^2}, r, r_1, T, \phi\right)\right]$$

$$= \int_{\infty}^{\infty} \text{vanilla}\left(aS_0^{(1)} \exp\left\{\sigma_1\rho\sqrt{T}x - \frac{1}{2}\sigma_1^2\rho^2 T\right\},\right. \tag{408}$$
$$\left. b\exp\{\sigma_2\sqrt{T}x + \hat{\mu}_2\} + K, \sigma_1\sqrt{1-\rho^2}, r, r_1, T, \phi\right) n(x)\, dx$$

$$= \int_{-\infty}^{+\infty} \text{vanilla}\left(S(x), K(x), \sigma_1\sqrt{1-\rho^2}, r, r_1, T, \phi\right) n(x)\, dx. \tag{409}$$

Example We consider the example in Table 1.39. An investor or corporate believes that EUR/USD will outperform GBP/USD in 6 months. To make the exchange rates comparable we first normalize both exchange rates by dividing by their current spot and then want to reward the investor by one pip for each pip the normalized EUR/USD will be more than 20 pips higher than normalized GBP/USD.

1.9.2 Baskets

This section is joint work with Jürgen Hakala and appeared first in [67].

In many cases corporate and institutional currency managers are faced with an exposure in more than one currency. Generally these exposures would be hedged using individual strategies for each currency. These strategies are composed of spot transactions, forwards, and in many cases options on a single currency. Nevertheless, there are instruments that include several currencies, and these can be used to build a multi-currency strategy that is almost always more cost effective than the portfolio

TABLE 1.39 Example of a spread option.

	EUR	GBP	USD rate	3%
Spot in USD	1.2000	1.8000	Correlation	20%
Interest rates	2%	4%	Maturity	0.5 years
Volatility	10%	9%	Strike	0.0020
Weights	1/1.2000	1/1.8000	Value	0.0375 USD

of the individual strategies. As a prominent example we consider now basket options in detail.

Protection with Currency Baskets Basket options are derivatives based on a common base currency, say EUR, and several other risky currencies. The option is actually written on the basket of risky currencies. A basket option in Foreign Exchange markets is a European option granting its holder the right to exercise and upon exercise the holder pays/receives a portfolio of put currency amounts and receives/pays a call currency amount. In case of cash settlement this is economically equivalent to paying the difference between the basket cutoff value and the strike, if positive, for a basket call, or the difference between strike and basket value, if positive, for a basket put respectively at maturity. The risky currencies have different weights in the basket to reflect the details of the exposure.

For example, one can write down the payoff at maturity T of a basket call on two currencies USD and JPY as

$$\max\left(a_1\frac{S_1(T)}{S_1(0)} + a_2\frac{S_2(T)}{S_2(0)} - K, 0\right), \tag{410}$$

where $S_1(t)$ denotes the exchange rate of EUR-USD and $S_2(t)$ denotes the exchange rate of EUR-JPY at time t, a_i the corresponding weights, and K the basket strike. A basket option protects against a drop in both currencies at the same time. Individual options on each currency cover some cases that are not protected by a basket option, which occurs if one exchange rate falls more than the other exchange rate rises. In this case the basket would be out-of-the-money, but one of the individual options would be in-the-money. This is indicated by the shaded triangular areas in Figure 1.68, and that is why the portfolio of individual options would cost more than a basket. However, for the corporate the basket provides sufficient protection against the joint risk. If one exchange rate falls more than the other one rises, the corporate treasurer could buy the cheaper currency in the spot market and would make some extra profit, which may feel good, but this effect is just a potential bonus for which the treasurer would have paid when buying individual options rather than a basket.

Valuation of Basket Options in the Black-Scholes Model Basket options should be priced in a consistent way with plain vanilla options. In the Black-Scholes model we assume a log-normal process for the individual correlated basket constituents. A decomposition into uncorrelated constituents of the exchange rate processes

$$dS_i = \mu_i S_i\, dt + S_i \sum_{j=1}^{N} \Omega_{ij}\, dW_j \tag{411}$$

is the basis for pricing. Here μ_i denotes the difference between the foreign and the domestic interest rate of the i-th currency pair, dW_j the j-th component of independent Brownian increments. The covariance matrix is given by

$$C_{ij} = (\Omega\Omega^T)_{ij} = \rho_{ij}\sigma_i\sigma_j. \tag{412}$$

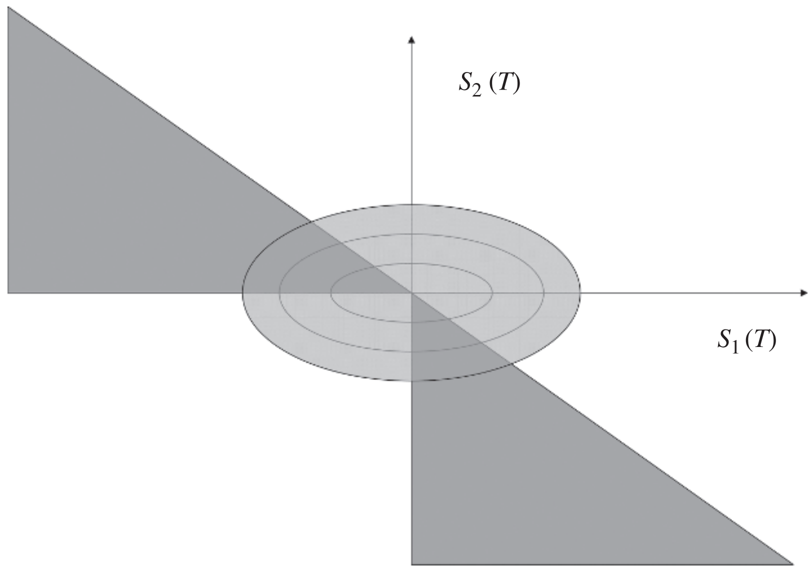

FIGURE 1.68 Protection with a basket option in two currencies. The ellipsoids connect the points that are reached with the same probability assuming that the forward prices are at the center.

Here σ_i denotes the volatility of the i-th currency pair and ρ_{ij} the correlation coefficients.

Exact Method. Starting with the uncorrelated components the pricing problem is reduced to the N-dimensional integration of the payoff. This method is accurate but rather slow for more than two or three basket components.

A Simple Approximation via Moment Matching assumes that the basket spot itself is a log-normal process with drift μ and volatility σ driven by a Wiener process $W(t)$,

$$dS(t) = S(t)[\mu\ dt + \sigma\ dW(t)] \tag{413}$$

with solution

$$S(T) = S(t)e^{\sigma W(T-t)+\left(\mu-\frac{1}{2}\sigma^2\right)(T-t)}, \tag{414}$$

given we know the spot $S(t)$ at time t. It is a fact that the sum of log-normal processes is not log-normal, but as a crude approximation it is certainly a quick method that is easy to implement. In order to price the basket call, the drift and the volatility of the basket spot need to be determined. This is done by matching the first and second moment of the basket spot with the first and second moment of the log-normal model for the basket spot. The moments of log-normal spot are

$$\mathbb{E}[S(T)] = S(t)e^{\mu(T-t)}, \tag{415}$$

$$\mathbb{E}[S(T)^2] = S(t)^2e^{(2\mu+\sigma^2)(T-t)}. \tag{416}$$

We solve these equations for the drift and volatility,

$$\mu = \frac{1}{T-t} \ln \left(\frac{\mathbb{E}[S(T)]}{S(t)} \right), \tag{417}$$

$$\sigma = \sqrt{\frac{1}{T-t} \ln \left(\frac{\mathbb{E}[S(T)^2]}{S(t)^2} \right)}. \tag{418}$$

In these formulas we now use the moments for the basket spot,

$$\mathbb{E}[S(T)] = \sum_{j=1}^{N} \alpha_j S_j(t) e^{\mu_j (T-t)}, \tag{419}$$

$$\mathbb{E}[S(T)^2] = \sum_{i,j=1}^{N} \alpha_i \alpha_j S_i(t) S_j(t) e^{\left(\mu_i + \mu_j + \sum_{k=1}^{N} \Omega_{ki} \Omega_{jk}\right)(T-t)}. \tag{420}$$

The value is given by the well-known Black-Scholes-Merton formula for plain vanilla call options,

$$v(0) = e^{-r_d T} \left(f \mathcal{N}(d_+) - K \mathcal{N}(d_-) \right), \tag{421}$$

$$f = S(0) e^{\mu T}, \tag{422}$$

$$d_\pm = \frac{\ln \frac{f}{K} \pm \frac{1}{2}\sigma^2 T}{\sigma \sqrt{T}}, \tag{423}$$

where $\mathcal{N}$ denotes the cumulative standard normal distribution function and r_d the domestic interest rate.

A More Accurate and Equally Fast Approximation. The previous approach can be taken one step further by introducing one more term in the Itô-Taylor expansion of the basket spot, which results in

$$v(0) = e^{-r_d T} \left(F \mathcal{N}(d_1) - K \mathcal{N}(d_2) \right), \tag{424}$$

$$F = \frac{S(0)}{\sqrt{1-\lambda T}} e^{\left(\mu - \frac{\lambda}{2} + \frac{\lambda \sigma^2}{2(1-\lambda T)}\right) T}, \tag{425}$$

$$d_2 = \frac{\sigma - \sqrt{\sigma^2 + \lambda \left(\left(1 + \frac{\lambda}{1-\lambda T} \right) \sigma^2 T - 2 \ln \frac{F\sqrt{1-\lambda T}}{K} \right)}}{\lambda \sqrt{T}}, \tag{426}$$

$$d_1 = \sqrt{1-\lambda T} d_2 + \frac{\sigma \sqrt{T}}{\sqrt{1-\lambda T}}. \tag{427}$$

The new parameter λ is determined by matching the third moment of the basket spot and the model spot. For details see [65]. Most remarkably, this major improvement in the accuracy requires only a marginal additional computation effort.

Correlation Risk Correlation coefficients between market instruments are usually not obtained easily. Either historical data analysis or implied calibrations need to be done. Implying the correlation from a traded instrument, however, will not produce the correlation, but the worst case assumption about correlation that somebody else has made, who does not know either what the correlation is. However, in the foreign exchange market the cross instrument is sometimes traded as well, for the example above the USD-JPY spot and options are traded, and the correlation in the Black-Scholes model can be determined from this contract. In fact, denoting the volatilities as in the tetrahedron in Figure 1.69, we obtain formulas for the correlation coefficients in terms of known market implied volatilities

$$\rho_{12} = \frac{\sigma_3^2 - \sigma_1^2 - \sigma_2^2}{2\sigma_1\sigma_2}, \tag{428}$$

$$\rho_{34} = \frac{\sigma_1^2 + \sigma_6^2 - \sigma_2^2 - \sigma_5^2}{2\sigma_3\sigma_4}. \tag{429}$$

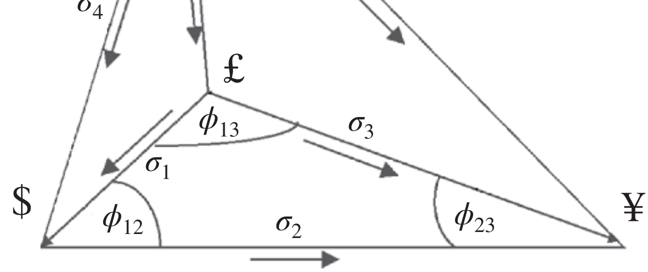

FIGURE 1.69 Relationship between volatilities σ (edges) and correlations ρ (cosines of angles) in a tetrahedron with four currencies and six currency pairs. The arrows mark the market standard quotation direction, i.e. in EUR-USD the base currency is USD and the arrow points to USD.

This method also allows hedging correlation risk by trading FX implied volatility. For details see [65]. While these relationships hold in a log-normal model for the at-the-money volatilities, De Col and Kuppinger [30] extend them to pricing multi-dimensional FX derivatives via stochastic local correlations. Zerolis explains how to picture volatility and correlation in [145].

Pricing Basket Options with Smile The previous calculations are all based on the Black-Scholes model with constant market parameters for rates and volatility. This can all be made time-dependent and can then include the term structure of volatility. If we wish to include the smile in the valuation, then we can either switch to a more appropriate model or perform a Monte Carlo simulation where the probabilities of the exchange rate paths are computed in such a way that the individual vanilla prices are correctly determined. This *weighted Monte Carlo approach* has been discussed by Avellaneda *et al*. in [5].

Practical Example We want to find out how much one can save using a basket option. We take EUR as a base currency and consider a basket of three currencies: USD, GBP, and JPY. We list the contract data and the amount of option premium one can save using a basket call rather than three individual call options in Table 1.40 and the market data in Table 1.41.

TABLE 1.40 Sample contact data of a EUR call basket put. The value of the basket is noticeably less than the value of three vanilla EUR calls.

Contract data	Strikes	Weights	Single option prices
EUR/USD	1.1390	33.33%	4.94%
EUR/GBP	0.7153	33.33%	2.50%
EUR/JPY	125.00	33.33%	3.87%
Sum		100%	3.77%
Basket price			2.90%

TABLE 1.41 Sample market data of 21 October 2003 of four currencies: EUR, GBP, USD, and JPY. The correlation coefficients are implied from the volatilities based on Equation (428) for the triangles and Equation (429) for the tetrahedra.

Vol	Spot	Correlation ccy pair	GBP/USD	USD/JPY	GBP/JPY	EUR/USD	EUR/GBP	EUR/JPY
8.80	1.6799	GBP/USD	1.00	−0.49	0.42	0.72	−0.15	0.29
9.90	109.64	USD/JPY	−0.49	1.00	0.59	−0.55	−0.21	0.41
9.50	184.17	GBP/JPY	0.42	0.59	1.00	0.09	−0.35	0.70
10.70	1.1675	EUR/USD	0.72	−0.55	0.09	1.00	0.58	0.54
7.50	0.6950	EUR/GBP	−0.15	−0.21	−0.35	0.58	1.00	0.42
9.80	128.00	EUR/JPY	0.29	0.41	0.70	0.54	0.42	1.00

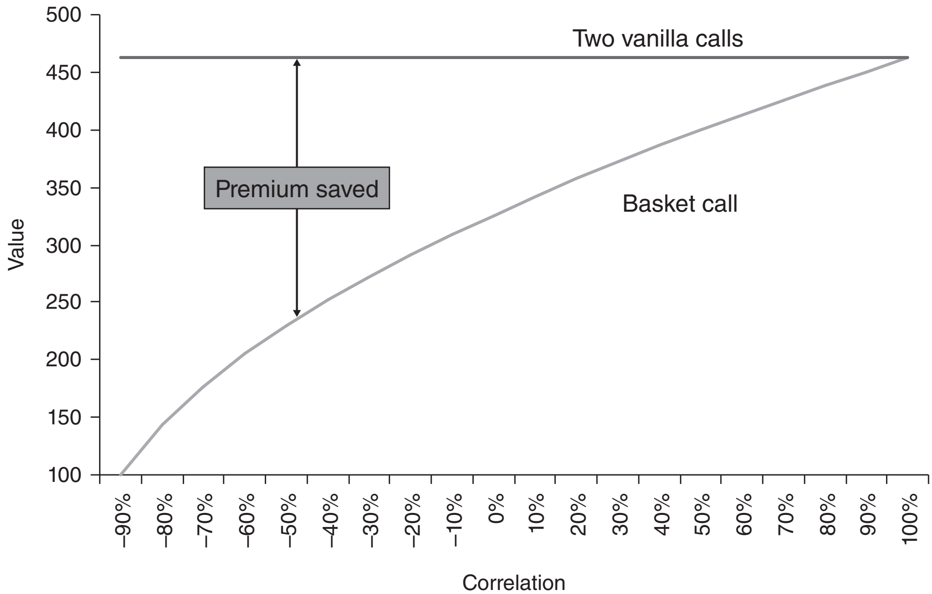

FIGURE 1.70 Amount of premium saved in a basket of two currencies compared with two single vanillas as a function of correlation: the smaller the correlation, the higher the premium savings effect.

The amount of premium saved essentially depends on the correlation of the currency pairs. In Figure 1.70 we take the parameters of the previous scenario, but restrict ourselves to the currencies USD and JPY. Note that graph is – while legally and mathematically correct – a sales slide. (Why?)

Conclusions Many corporate clients are exposed to multi-currency risk. One way to turn this fact into an advantage is to use multi-currency hedge instruments. We have shown that basket options are convenient instruments protecting against exchange rates of most of the basket components changing in the same direction. A rather unlikely market move of half of the currencies' exchange rates in opposite directions is not protected by basket options, but when taking this residual risk into account the hedging cost is reduced substantially. Note that for a treasurer with the underlying cash flow in the various currencies there is actually no risk. Risk comes in if a hedge fund sells the vanilla portfolio and buys the basket to generate net positive premium at inception and if then its correlation trade turns against the trader, and one of the vanilla options ends up deep in-the-money, whereas the basket is out-of-the-money. Another example of how to use currency basket options as part of a note is discussed in Section 2.6.2.

1.9.3 Outside Barrier Options

Outside barrier options are options in one currency pair with one or several barriers or window barriers in another currency pair. In general form the payoff can be written as

$$\left[\phi\left(S_T - K\right)\right]^+ I\!I_{\{\min_{0\le t\le T}(\eta R(t))>\eta B\}}. \tag{430}$$

This extends a European put or call with strike K by a knock-out barrier H in a second currency pair, called the *outer* currency pair. As usual, the binary variable ϕ takes the value +1 for a call and −1 for a put and the binary variable η takes the value +1 for a lower barrier and −1 for an upper barrier. I am still puzzled why outside barrier options trade, but this is definitely some funky correlation trade. On the other hand, why not?

Valuation We will now take a look at the valuation of the outside barrier option in the Black-Scholes model. The derivation is part of the exercises in integration, which you may or may not enjoy, but skipping it does not slow you down in reading the rest of the book. We let the positive constants σ_i denote the annual volatilities of the i-th asset or foreign currency, ρ the instantaneous correlation of their log-returns, r the domestic risk-free rate, and T the expiration time in years. In a risk-neutral setting the drift terms μ_i take the values

$$\mu_i = r - r_i \tag{431}$$

where r_i denotes the risk-free rate of the i-th foreign currency. Knock-in outside barrier options values can be obtained by the standard relationship *knock-in plus knock-out = vanilla*.

In the standard two-dimensional Black-Scholes model

$$dS_t = S_t\left[\mu_1 dt + \sigma_1 dW_t^{(1)}\right], \tag{432}$$

$$dR_t = R_t\left[\mu_2 dt + \sigma_2 dW_t^{(2)}\right], \tag{433}$$

$$\mathbf{Cov}\left[W_t^{(1)}, W_t^{(2)}\right] = \sigma_1\sigma_2\rho t, \tag{434}$$

Heynen and Kat derive the value in [74].

$$\begin{aligned} V_0 = {} & \phi S_0 e^{-r_1 T}\mathcal{N}_2(\phi d_1, -\eta e_1; \phi\eta\rho) \\ & - \phi S_0 e^{-r_1 T}\exp\left(\frac{2(\mu_2 + \rho\sigma_1\sigma_2)\ln(H/R_0)}{\sigma_2^2}\right)\mathcal{N}_2(\phi d_1', -\eta e_1'; \phi\eta\rho) \\ & - \phi K e^{-rT}\mathcal{N}_2(\phi d_2, -\eta e_2; \phi\eta\rho) \\ & + \phi K e^{-rT}\exp\left(\frac{2\mu_2\ln(H/R_0)}{\sigma_2^2}\right)\mathcal{N}_2(\phi d_2', -\eta e_2'; \phi\eta\rho), \end{aligned} \tag{435}$$

$$d_1 = \frac{\ln(S_0/K) + (\mu_1 + \sigma_1^2)T}{\sigma_1\sqrt{T}}, \tag{436}$$

$$d_2 = d_1 - \sigma_1\sqrt{T}, \tag{437}$$

$$d_1' = d_1 + \frac{2\rho\ln(H/R_0)}{\sigma_2\sqrt{T}}, \tag{438}$$

$$d_2' = d_2 + \frac{2\rho \ln(H/R_0)}{\sigma_2\sqrt{T}}, \tag{439}$$

$$e_1 = \frac{\ln(H/R_0) - (\mu_2 + \rho\sigma_1\sigma_2)T}{\sigma_2\sqrt{T}}, \tag{440}$$

$$e_2 = e_1 + \rho\sigma_1\sqrt{T}, \tag{441}$$

$$e_1' = e_1 - \frac{2\ln(H/R_0)}{\sigma_2\sqrt{T}}, \tag{442}$$

$$e_2' = e_2 - \frac{2\ln(H/R_0)}{\sigma_2\sqrt{T}}. \tag{443}$$

The bi-variate standard normal distribution $\mathcal{N}_2$ and density functions n_2 are defined by

$$n_2(x, y; \rho) \stackrel{\Delta}{=} \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x^2 - 2\rho xy + y^2}{2(1-\rho^2)}\right), \tag{444}$$

$$\mathcal{N}_2(x, y; \rho) \stackrel{\Delta}{=} \int_{-\infty}^{x} \int_{-\infty}^{y} n_2(u, v; \rho)\, du\, dv. \tag{445}$$

Greeks For the Greeks, most of the calculations of partial derivatives can be simplified substantially by the homogeneity method described by Reiss and Wystup in [107], which states, for instance, that

$$V_0 = S_0 \frac{\partial V_0}{\partial S_0} + K \frac{\partial V_0}{\partial K}. \tag{446}$$

We list some of the sensitivities for reference.

delta (inner spot)

$$\begin{aligned}\frac{\partial V_0}{\partial S_0} &= \phi e^{-r_1 T} \mathcal{N}_2(\phi d_1, -\eta e_1; \phi\eta\rho) \\ &\quad - \phi e^{-r_1 T} \exp\left(\frac{2(\mu_2 + \rho\sigma_1\sigma_2)\ln(H/R_0)}{\sigma_2^2}\right) \mathcal{N}_2(\phi d_1', -\eta e_1'; \phi\eta\rho)\end{aligned} \tag{447}$$

dual delta (inner strike)

$$\begin{aligned}\frac{\partial V_0}{\partial K} &= -\phi e^{-rT} \mathcal{N}_2(\phi d_2, -\eta e_2; \phi\eta\rho) \\ &\quad + \phi e^{-rT} \exp\left(\frac{2\mu_2 \ln(H/R_0)}{\sigma_2^2}\right) \mathcal{N}_2(\phi d_2', -\eta e_2'; \phi\eta\rho)\end{aligned} \tag{448}$$

gamma (inner spot)

$$\frac{\partial^2 V_0}{\partial S_0^2} = \frac{e^{-r_1 T}}{S_0 \sigma_1 \sqrt{T}} \left[n(d_1)\mathcal{N}\left(\frac{-\phi\rho d_1 - \eta e_1}{\sqrt{1-\rho^2}}\right) \right.$$
$$\left. - \exp\left(\frac{2(\mu_2 + \rho\sigma_1\sigma_2)\ln(H/R_0)}{\sigma_2^2}\right) n(d_1')\mathcal{N}\left(\frac{-\phi\rho d_1' - \eta e_1'}{\sqrt{1-\rho^2}}\right)\right] \quad (449)$$

The standard normal density function n and its cumulative distribution function $\mathcal{N}$ are defined in (473) and (480). Furthermore, we use the relations

$$\frac{\partial}{\partial x}\mathcal{N}_2(x, y; \rho) = n(x)\mathcal{N}\left(\frac{y - \rho x}{\sqrt{1-\rho^2}}\right), \quad (450)$$

$$\frac{\partial}{\partial y}\mathcal{N}_2(x, y; \rho) = n(y)\mathcal{N}\left(\frac{x - \rho y}{\sqrt{1-\rho^2}}\right). \quad (451)$$

dual gamma (inner strike) Again, the homogeneity method described in [107] leads to the result

$$S^2 \frac{\partial^2 V_0}{\partial S_0^2} = K^2 \frac{\partial^2 V_0}{\partial K^2}. \quad (452)$$

1.9.4 Best-of and Worst-of Options

Options on the maximum or minimum of two or more exchange rates can be defined by their payoffs in their simple version as

$$\left[\phi\left(\eta \min(\eta S_T^{(1)}, \eta S_T^{(2)}) - K\right)\right]^+. \quad (453)$$

This payoff resembles a European put or call with expiration time T in years on the minimum ($\eta = +1$) or maximum ($\eta = -1$) of the two underlying exchange rates $S_T^{(1)}$ and $S_T^{(2)}$ with strike K. As usual, the binary variable ϕ takes the value +1 for a call and −1 for a put.

Valuation in the Black-Scholes Model In the two-dimensional Black-Scholes model

$$dS_t^{(1)} = S_t^{(1)}\left[\mu_1 dt + \sigma_1 dW_t^{(1)}\right], \quad (454)$$

$$dS_t^{(2)} = S_t^{(2)}\left[\mu_2 dt + \sigma_2 dW_t^{(2)}\right], \quad (455)$$

$$\mathbf{Cov}\left[W_t^{(1)}, W_t^{(2)}\right] = \sigma_1\sigma_2\rho t, \quad (456)$$

we let the positive constants σ_i denote the volatilities of the i-th foreign currency, ρ the instantaneous correlation of their log-returns, r the domestic risk-free rate. In a risk-neutral setting the drift terms μ_i take the values

$$\mu_i = r - r_i, \tag{457}$$

where r_i denotes the risk-free rate of the i-th foreign currency.

The value was published originally by Stulz in [122] and happens to be

$$\begin{aligned} & v(t, S_t^{(1)}, S_t^{(2)}, K, T, r_1, r_2, r, \sigma_1, \sigma_2, \rho, \phi, \eta) \\ & \quad = \phi \Bigg[S_t^{(1)} e^{-r_1\tau} \mathcal{N}_2(\phi d_1, \eta d_3; \phi\eta\rho_1) + S_t^{(2)} e^{-r_2\tau} \mathcal{N}_2(\phi d_2, \eta d_4; \phi\eta\rho_2) \\ & \qquad - K e^{-r\tau} \left(\frac{1-\phi\eta}{2} + \phi\eta \mathcal{N}_2(\eta(d_1 - \sigma_1\sqrt{\tau}), \eta(d_2 - \sigma_2\sqrt{\tau}); \rho) \right) \Bigg], \end{aligned} \tag{458}$$

$$\sigma^2 \triangleq \sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2, \tag{459}$$

$$\rho_1 \triangleq \frac{\rho\sigma_2 - \sigma_1}{\sigma}, \tag{460}$$

$$\rho_2 \triangleq \frac{\rho\sigma_1 - \sigma_2}{\sigma}, \tag{461}$$

$$\tau \triangleq T - t, \tag{462}$$

$$d_1 \triangleq \frac{\ln(S_t^{(1)}/K) + (\mu_1 + \frac{1}{2}\sigma_1^2)\tau}{\sigma_1\sqrt{\tau}}, \tag{463}$$

$$d_2 \triangleq \frac{\ln(S_t^{(2)}/K) + (\mu_2 + \frac{1}{2}\sigma_2^2)\tau}{\sigma_2\sqrt{\tau}}, \tag{464}$$

$$d_3 \triangleq \frac{\ln(S_t^{(2)}/S_t^{(1)}) + (r_1 - r_2 - \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}, \tag{465}$$

$$d_4 \triangleq \frac{\ln(S_t^{(1)}/S_t^{(2)}) + (r_2 - r_1 - \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}. \tag{466}$$

The bivariate standard normal distribution and density functions $\mathcal{N}_2$ and n_2 are defined in Equation (445) and Equation (444). I let you enjoy deriving the Greeks in the exercises.

Variations Options on the maximum and minimum can be generalized in various ways. For instance, they can be quantoed or have individual strikes for each currency pair. We consider some examples.

Multiple strike option. This variation of best-of/worst-of options deals with individual strikes, i.e. they pay off

$$\max_i \left[0; M_i(\phi(S_T^{(i)} - K_i))\right]. \tag{467}$$

Quanto best-of/worst-of options. These options come up naturally if an investor wants to participate in several exchange rate movements with a payoff in a currency other than the base currency.

Barrier best-of/worst-of options. One can also add knock-out and knock-in features to all the previous types discussed.

Application in Re-Insurance Suppose you want to protect yourself against a weak USD compared with several currencies for a period of one year. As USD seller and buyer of EUR, GBP, and JPY you need simultaneous protection against all three rising against the USD. Of course, you can buy three put options, but if you need only one of the three, then the premium can be considerably reduced, as shown in Table 1.42. We can imagine a situation like this if a re-insurance company insures ships in various oceans. If a ship sinks near the coast of Japan, the client will have to be paid an amount in JPY. The re-insurance company is long USD and assumes only one ship at most will sink in one year so is ready to take the residual risk of more than one sinking.

Since the accidents can occur any time, all options are of American style, i.e. they can be exercised any time. The holder of the option can choose the currency pair to exercise. Hence, he can decide for the one with the highest profit, even if the currency of accident is a different one. It would be difficult to incorporate and hedge this event insurance into the product, whence the protection needs to assume the worst case scenario that is still acceptable to the re-insurance company. For example, if the re-insurance company needs GBP and the spots at exercise time are at EUR/USD = 1.1200, USD/JPY = 134.00, and GBP/USD = 1.6400, you will find both the EUR and GBP constituents in-the-money. However, exercising in GBP would pay a net of 613,496.93 USD, exercising in EUR pays 6,666,666.67 USD. The client would then exercise in EUR, buy the desired GBP in the EUR/GBP spot market, and keep the rest of the EUR.

TABLE 1.42 Example of a triple strike best-of call (American style) with 100 M USD notional and one year maturity. Compared with buying vanilla options one saves 800,000 USD or 20%. All premiums are in USD.

Currency pair	Spot	Strikes	Vanilla premium	Best-of premium
EUR/USD	0.9750	1.0500	1.4 M	
USD/JPY	119.00	110.00	1.7 M	
GBP/USD	1.5250	1.6300	0.9 M	
		Total in USD	4.0 M	3.2 M

Application in Corporate and Private Banking Just like a *dual currency deposit* described in Section 2.4.1, one can use a worst-of put to structure a *multi-currency deposit* with a coupon even higher. We refer the reader to the exercises. This idea mostly comes up when volatilities are low and therefore selling options does not generate sufficient yield enhancement.

1.9.5 Other Multi-Currency Options

Generally, for multi-currency derivatives, there are no bounds to being creative. You get all colors and flavors. For this reason, multi-currency options are also sometimes referred to as *rainbow options*. Most of them have been more popular in equity markets.

Quanto Exotics In foreign exchange options markets cash-settle options can have payoffs in a currency different from the underlying currency pair. For instance, a USD/JPY call is designed to be paid in EUR, where the exchange rate for EUR/JPY is fixed upfront. Surely such quanto features can be applied to exotics as well. Same principles for quanto options apply as explained in Section 1.7.10.

Madonna, Pyramid, Montain Range, and Himalaya Options There are in fact derivatives where one might (rightfully) ask if this is meant to be serious or not. A few examples are contained in but are not limited to the following list.

Madonna option. This one pays the *Euclidian distance*,

$$\max\left[0; \sqrt{\sum_i (S_T^{(i)} - K_i)^2}\right]. \tag{468}$$

Pyramid option. This one pays the *maximum norm*,

$$\max\left[0; \sum_i |S_T^{(i)} - K_i| - K\right]. \tag{469}$$

Mountain range and Himalaya option. This type of option comes in various flavors and is rather popular in equity markets, so we will not discuss them here. A reference is the thesis by Mahomed [94].

In the next section we return to the real world of currency derivatives and deal with the correlation swap.

1.9.6 Correlation Swap

In a *correlation swap* two parties trade a fixed notional amount multiplied by the difference of a fixed correlation ρ_{fixed} and a prevailing historic correlation ρ_{historic} calculated by Equation (137). The payoff is simply

$$\text{correlation swap payoff} = \text{notional} \cdot (\rho_{\text{fixed}} - \rho_{\text{historic}}). \tag{470}$$

The fixed correlation that makes the transaction worth zero is called the *par correlation rate* or *fair correlation rate*.

Example of a Traded Correlation Swap On 25 June 2007 a correlation swap traded, terminating on 25 September 2007 on a notional of USD 19,563,090.00. The observation period was from and including the trade date to and including the termination date. The fixed rate was +82%, the floating rate based on prevailing historic correlation calculated by Equation (137). The first currency pair was GBP/NZD and the second one USD/NZD. As spot rate references the parties agreed on the source Reuters page **WMRSPOT11** using New York business days. Calculation date was 26 September 2007, and settlement date 27 September 2007. *Settlement differential* means the fixed rate minus the floating rate. The Settlement Differential is to be expressed as a percentage which may be a positive number or a negative number. Here is how the lawyers deal with minus signs: *Settlement Amount* means:

1. if the Settlement Differential is positive, then the Settlement Amount is an amount equal to the notional amount multiplied by the Settlement Differential and such amount shall be paid by the fixed rate payer to the floating rate payer on the settlement date;
2. if the Settlement Differential is negative, then the Settlement Amount is an amount equal to the notional amount multiplied by the Settlement Differential and such amount shall be paid by the floating rate payer to the fixed rate payer on the settlement date.

This traded as a zero-cost product. Obviously, the fixed rate payer has a small potential maximum gain of 18% of the USD notional and a much larger maximum potential loss of 182%. The fixed rate payer takes a view in USD and GBP both becoming stronger against NZD in the next three months. I am not sure on which analysis such a view is taken. Notice that the maximum potential loss exceeds the notional amount. So watch out: correlations can be negative!

Extensions One can put bounds on the best and/or worst case correlation. There are also correlation options instead of swaps.

1.9.7 Exercises

Exchange Option Compute the integral of the spread option for the special case of a zero strike in Equation (396) to get a closed form solution for the exchange option.

TABLE 1.43 Sample market ATM volatilities of four currencies: EUR, GBP, USD, and CHF.

Ccy pair	Volatility	Ccy pair	Volatility
GBP/USD	9.20%	EUR/USD	10.00%
USD/CHF	11.00%	EUR/GBP	7.80%
GBP/CHF	8.80%	EUR/CHF	5.25%

Implied Correlation Compute the correlation coefficients implied from the volatilities in Table 1.43. What are the upper and lower limits for the EUR/USD volatility to guarantee all correlation coefficients being contained in the interval $[-1, +1]$, assuming all the other volatilities are fixed?

Outside Barrier Option Value Formula Derive the value function of the outside barrier option in the Black-Scholes model, see Equation (435). We start with a triple integral. We treat the up-and-out call as an example. The value of an outside up-and-out call option is given in Section 24 in Shreve's lecture notes [119] by the integral

$$V_0 \triangleq \frac{e^{-rT}}{\sqrt{T}} \int_{\hat{m}=0}^{\hat{m}=m} \int_{\hat{b}=-\infty}^{\hat{b}=\hat{m}} \int_{\tilde{b}=-\infty}^{\tilde{b}=\infty} F(\hat{b}, \tilde{b}) n\left(\frac{\tilde{b}}{\sqrt{T}}\right) f(\hat{m}, \hat{b})\, d\tilde{b}\, d\hat{b}\, d\hat{m}, \tag{471}$$

where the payoff function F, the normal density function n, the joint density function f, and the parameters m, b, $\hat{\theta}$, γ are defined by

$$F(\hat{b}, \tilde{b}) \triangleq \left(S_0 e^{\gamma\sigma_2 T + \rho\sigma_2\hat{b} + \sqrt{1-\rho^2}\sigma_2\tilde{b}} - K\right)^+ \tag{472}$$

$$n(t) \triangleq \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2}, \tag{473}$$

$$f(\hat{m}, \hat{b}) \triangleq \frac{2(2\hat{m} - \hat{b})}{T\sqrt{2\pi T}} \exp\left\{-\frac{(2\hat{m} - \hat{b})^2}{2T} + \hat{\theta}\hat{b} - \tfrac{1}{2}\hat{\theta}^2 T\right\}, \tag{474}$$

$$m \triangleq \frac{1}{\sigma_1} \ln \frac{L}{Y_0}, \tag{475}$$

$$b \triangleq \frac{1}{\sigma_2} \ln \frac{L}{S_0}, \tag{476}$$

$$\hat{\theta} \triangleq \frac{r}{\sigma_1} - \frac{\sigma_1}{2}, \tag{477}$$

$$\gamma \triangleq \frac{r}{\sigma_2} - \frac{\sigma_2}{2} - \rho\hat{\theta}. \tag{478}$$

The goal is to write the above integral in terms of the bi-variate normal distribution function (445). For easier comparison we use the mapping of the notation in Table 1.44.

TABLE 1.44 Relating the notation of Heynen and Kat to the one by Shreve.

Heynen/Kat	Shreve	Heynen/Kat	Shreve
S_0	S_0	H	L
R_0	Y_0	K	K
σ_1	σ_2	μ_1	$r - \frac{\sigma_2^2}{2}$
σ_2	σ_1	μ_2	$r - \frac{\sigma_1^2}{2}$

The solution can be obtained by taking the following steps:

(a) Use a change of variables to prove the identity

$$\int_{-\infty}^{x} \mathcal{N}(az+B)n(z)\,dz = \mathcal{N}_2\left(x, \frac{B}{\sqrt{1+a^2}}; \frac{-a}{\sqrt{1+a^2}}\right), \tag{479}$$

where the cumulative normal distribution function $\mathcal{N}$ is defined by

$$\mathcal{N}(x) \stackrel{\Delta}{=} \int_{-\infty}^{x} n(t)\,dt. \tag{480}$$

A probabilistic proof is presented in [52].

(b) Extend the identity (479) to

$$\int_{-\infty}^{x} e^{Az}\mathcal{N}(az+B)n(z)\,dz = e^{\frac{A^2}{2}}\mathcal{N}_2\left(x-A, \frac{aA+B}{\sqrt{1+a^2}}; \frac{-a}{\sqrt{1+a^2}}\right). \tag{481}$$

(c) Change the order of integration in Equation (471) and integrate the $\hat{m}$ variable.

(d) Change the order of integration to make $\tilde{b}$ the inner variable and $\hat{b}$ the outer variable. Then use the condition $F(\hat{b}, \tilde{b}) \geq 0$ to find a lower limit for the range of $\tilde{b}$. This will enable you to skip the positive part in F and write Equation (471) as a sum of four integrals.

(e) Use (479) and (481) to write each of these four summands in terms of the bi-variate normal distribution function $\mathcal{N}_2$.

(f) Compare your result with the one provided by Heynen and Kat using Table 1.44.

Inside Barrier Option as Special Case of an Outside Barrier Option Derive the value function of an inside barrier option by viewing it as a special case of the outside barrier option in Equation (435).

TABLE 1.45 Sample short time series of two spots.

Spot 1	Spot 2
100	0.8
110	0.7
120	0.6
130	0.5

Greeks of Best-of/Worst-of Options Derive the Greeks of the best-of/worst-of value function in Equation (458). Hint: most of the calculations of partial derivatives can be simplified substantially by considering homogeneity properties described in [107].

Correlation Shocker Given the two time series of spot prices in Table 1.45, calculate the correlation (as it would be used in a correlation swap).