$$\frac{\mathrm{Im}\chi_{\perp}(\boldsymbol{q},\omega_E)}{\omega_N} = \frac{Dq^2}{(Dq^2)^2+\omega_E^2}\frac{\omega_E}{\omega_N} . \tag{3.36}$$

The ratio of both Larmor frequencies is independent of the magnetic field (and equal to the ratio of the inverse electronic and nuclear masses), and will not be considered further. The sum over $\boldsymbol{q}$ in the first fraction on the right-hand side crucially depends on dimension: in 1D, one obtains $\sim \omega_E^{-1/2}$ and in 3D, an ω_E-independent result for small ω_E. Converting to magnetic fields, one finds

$$\frac{1}{T_1 T} \sim \begin{cases} \text{const.} & (3D) \\ 1/\sqrt{H_0} & (1D) . \end{cases} \tag{3.37}$$

The experimental results are shown in Fig. 3.10. For ambient pressure, curve (a), they show a wide range of fields where the electronic spin diffusion is indeed 1D. Only at small fields does one observe a crossover to a field-independent relaxation rate typical for 3D diffusion, (3.37). The idea behind this crossover is the following. Even in a rather 1D band structure, the electrons will have a small but finite chance of tunneling to a neighboring chain. They will thus have a finite lifetime $\tau_{\perp}$ on one chain. This lifetime will cut off the influence of their diffusive motion on spin-relaxation because, due to the locality of the hyperfine interaction, the nucleus will no longer see the electronic spin. The 1D limit then corresponds to $\tau_{\perp} \to \infty$ while the 3D limit is $\tau_{\perp} \to 0$. The lifetime of a spin on a chain is estimated to be $\tau_{\perp} \sim 8 \times 10^{-12}$s at 300 K from this experiment [41].

4. The Black–Scholes Theory of Option Prices

We now turn to the determination of the prices of derivative securities such as forwards, futures, or options in the presence of fluctuations in the price of the underlying. Such investments for speculative purposes are risky. Bachelier's work on futures already shows that for *relative* prices, even the deterministic movements of the derivative are much stronger than those of the bond, and it seems clear that an investment into a derivative is then associated with a much higher risk (see also Bachelier's evaluation of success rates) than in the underlying security, although the opportunities for profit would also be higher.

Derivative prices depend on certain properties of the stochastic process followed by the price of the underlying security. Remember from Chap. 2 that options are some kind of insurance: the price of an insurance certainly depends on the frequency of occurrence of the event to be insured. We therefore introduce the standard model of stock prices, as used in textbooks of quantitative finance [10], [12]–[16] and place this model in a more general context of stochastic processes.

4.1 Important Questions

Based on these models, we will discuss some of the important questions which are listed below.

- What determines the price of a derivative security?
- What is the role of the return of the underlying security, i.e., the drift in its price?
- What are the appropriate stochastic processes to model financial time series? Are they independent of the assets considered?
- How can we classify stochastic processes?
- How can we calculate with stochastic variables?
- What is geometric Brownian motion? Is it different from Bachelier's model?
- What is the risk of an investment in a derivative?
- What is the price of risk?
- Can risk in financial markets be eliminated? At what cost?

- Can option pricing be related to diffusion? What would be different from standard diffusion problems?
- How can we calculate option prices in ideal markets? What is different in real markets?
- What are "The Greeks"?
- How do traders represent the deviations of traded option prices from those calculated in idealized models?
- How are derivative prices related to the expected payoff of the derivative?
- What is the difference in pricing European and American-style options?
- Can options be created synthetically?
- What is a volatility index, and how is it constructed?

The important achievement of Black and Scholes [42] and Merton [43] was to answer almost all of these questions, at least for a certain idealized market. While of course one can take a speculative position in a derivative involving a big risk, Black, Merton, and Scholes show that the risk can be eliminated in principle by a hedging strategy, i.e., by an investment in another security correlated with the derivative, so as to offset all or part of the price variations. For options, there is a dynamic hedging strategy by which *the risk can be eliminated completely.* At the same time, the possibility of hedging the risk allows to one fix a fair price of an option: it is determined by the expected payoff for the holder and the cost of the hegde, and no *additional risk premium* is necessary on options in idealized markets. Although their assumptions are not necessarily realistic, this is a benchmark result which earned Merton and Scholes the 1997 Nobel Prize in Economics, Black having died meanwhile. For forwards and futures, a static hedge, implemented at the time of writing, is sufficient.

Here, we only present the theoretical framework established in finance [10]. Of course, this heavily draws on the assumption of a random walk followed by financial time series. While we have discussed random walks in finance and physics in the previous chapter quite generally, we will specify in detail the model used by economists. More advanced and more speculative proposals for derivative pricing and hedging will be discussed later in Chap. 7. Also, we will limit our discussion to the most basic derivatives (forwards, futures, and European options): they are sufficient to illustrate the main principles. The methods developed here can then be applied, with only minor extensions, to more complicated instruments [10]. Hull's book [10] also contains much more information on practical aspects, and is highly recommended for reading.

4.2 Assumptions and Notation

4.2.1 Assumptions

Here, we summarize the main economic assumptions underlying the work of Black, Merton, and Scholes, as well as much related work on derivative pricing

and financial engineering. More specific assumptions on the stochastic process followed by the underlying security will be developed in Sect. 4.4. We assume:

- a complete and efficient market;
- zero transaction costs;
- that all profits are taxed in a similar way, and that consequently, tax considerations are irrelevant;
- that all market participants can lend and borrow money at the same risk-free interest rate r;
- that all market participants use all arbitrage possibilities;
- continuous compounding of interest, i.e., an amount of cash y accumulates interest as $y(T) = y(t)\exp[r(T-t)]$;
- that short selling with full profits is allowed;
- that there are no payoffs such as dividends, from the underlying securities (*we* shall make this assumption here to simplify matters; it is not realistic, and payoffs can be incorporated into derivative pricing schemes [10]).

4.2.2 Notation

Here we list the most important symbols used in the following chapters:

- T ... time of maturity of a derivative
- t ... present time
- S ... price of the underlying security
- K ... delivery price in a forward or futures contract
- f ... value of a long position in a forward or futures contract
- F ... price of forward contract
- r ... risk-free interest rate
- C ... price of a call option
- P ... price of a put option
- X ... strike price of the option.

4.3 Prices for Derivatives

Some price considerations are independent of the fluctuations of the price of the underlying securities. These are the forward prices and futures prices because they are binding contracts to both parties, and can be perfectly, and statically, hedged. (There are some restrictions to this statement for futures because they can be traded on exchanges.) We shall treat them first. Also some price limits for options can be derived without knowing the stochastic process of the underlying securities. An accurate calculation, however, requires this knowledge and will be deferred to Sect. 4.5.

4.3.1 Forward Price

We claim that the price of a forward contract on an underlying without payoff, such as dividends, is

$$F(t) = S(t)\exp[r(T-t)] \ . \tag{4.1}$$

Notice that this is the price *today* of the contract with maturity T. It is just the spot price with accumulated risk-free interest, and is *independent* of any historical or future drift in the price S of the underlying! We prove this equation in two different ways, in order to illustrate the methods of proofs often used in finance.

First Proof

We prove (4.1) by contradiction, relying on a "no arbitrage" argument. Assume first that $F(t) > S(t)\exp[r(T-t)]$. Then, at time t, an investor can borrow an amount of cash S and use it to buy the underlying at the spot price $S(t)$. At the same time, he goes short in the forward. This involves no cost because the forward is just a contract carrying the obligation to deliver the underlying at maturity. At maturity T, the credit must be reimbursed with interest accrued, i.e., there is a cash flow $-S(t)\exp[r(T-t)]$. The underlying is now sold under the terms of the forward contract, which results in a cash flow $F(T)$, the (yet) undetermined forward price. However, $F(T) = F(t)$, because the price of the forward has been fixed at the time of writing of the contract, and there are no trading opportunities. The total cash flow is therefore $F(t) - S(t)\exp[r(T-t)] > 0$, and a riskless profit can be made. This is contrary to the assumption of no arbitrage opportunities.

For the opposite assumption, $F(t) < S(t)\exp[r(T-t)]$, an investor can generate a riskless profit $S(t)\exp[r(T-t)] - F(t)$ by (i) taking the long position in the forward at t, (ii) short-selling the underlying asset at t, giving a cash flow $+S(t)$, (iii) investinging this money at the risk-free rate r at t, (iv) buying back the underlying asset at T under the terms of the forward contract, resulting in a cash flow $-F(T) = -F(t)$, and (v) getting back $S(t)\exp[r(T-t)]$ from his risk-free cash investment. Consequently, the only price compatible with the absence of arbitrage possibilities is (4.1).

Second Proof

The idea here is to construct two portfolios out of the three assets: forward, underlying and cash. These two portfolios carry the same risk, and their value at some instant of time can be shown to be equal.

Portfolio A contains a long position in the forward with a value $f(t)$, and an amount of cash $K\exp[-r(T-t)]$. At time T, this will be worth K. Portfolio B contains one underlying asset. At maturity T, the long position of the forward is used to acquire the asset, and both portfolios are worth the same because the delivery price K must be spent and both portfolios contain one asset. Moreover, both portfolios carry the same risk for all times because the long position in the forward *necessarily* receives the asset at maturity. Hence both portfolios have the same value for all times, i.e.,

$$f(t) + K\exp[-r(T-t)] = S(t) \ . \tag{4.2}$$

Now, the forward price can be fixed to the delivery price $F(t) = K$ by requiring that the net value of the long position at the time of writing is zero, i.e., that a fair contract for both parties is written. $f(t) = 0$ in (4.2) directly leads back to (4.1).

While these results may look trivial, they are indeed noteworthy:

- The prices of forwards and (to some extent, to be specified below) futures can be fixed at the time of writing the contract. They do not depend on the future evolution of the price of the underlying, up to maturity. Of course, a forward contract entered at a time $t' > t$, when the price of the underlying has changed to $S(t')$, will have a different price $F(t')$, determined again by (4.1). As the second proof makes clear, the "forward price" F actually is the delivery price of the underlying asset at maturity. It is not a price reflecting the intrinsic value of the contract. Unlike for the options to be discussed later, this intrinsic value is zero. The reason is that the outcome is certain: the underlying asset is delivered at maturity.
- In the above proofs, this fact was used to calculate the forward price in terms of the price of the underlying. A position in the forward, or in the underlying asset, carries a risk, connected to the price variations of the underlying asset. However, this risk can be hedged away statically (i.e., once and for all): for a long position in the forward, one can go short in the underlying, and for a short position in the forward, a long position in the underlying asset will eliminate the risk completely. This allows another interpretation of the forward price (4.1): in such a portfolio with a perfect hedge, there is no longer any risk. In the absence of arbitrage opportunities, it only can earn the risk-free interest rate r. This is precisely what (4.1) states.

4.3.2 Futures Price

Futures are distinguished from forwards mainly by being standardized, tradable instruments. If the interest rates do not vary during the period of the contract, the futures price equals the forward price. The prices are different, however, when interest rates vary. These differences are introduced by details of the trading procedures. For a forward, there is no cash flow for either party until maturity, where it will be settled. For futures, margin accounts (where a fixed fraction of the liabilities of a derivative portfolio is deposited for security) must be opened with the broker, and balanced daily. The money flowing in and out of these margin accounts in the case of a futures contract can then be invested, resp. must have been liquidated, at current market conditions, i.e., based on interest rates that may be different from those at the time the contract was entered. This gives different prices for forwards and futures. Empirically, however, the differences seem to be rather small [10].

4.3.3 Limits on Option Prices

The forward and future prices for contracts written today are independent of the details of the price history of the underlying, such as the drift or variance of the price. This is not so for options, and for accurate price calculations a knowledge of the important parameters of the price variations of the underlying is necessary. This will be developed in Sect. 4.5.1 below. On the other hand, it is fairly simple to obtain certain limits to be obeyed by option prices without knowing the price fluctuations of the underlying. If not stated otherwise, we will always consider European type options.

Upper Limits

A call option, by construction, can never be worth more than the underlying security. Therefore

$$C(t) \leq S(t) . \tag{4.3}$$

The value of a put option can never exceed the strike price

$$P(t) \leq X . \tag{4.4}$$

If one of these inequalities is violated, an arbitrageur can make riskless profit by buying the stock and selling the option (call), or simply selling the option (put). For a European put, a more stringent condition can be given because the strike price is also fixed in the future, and can be discounted from maturity to the present date

$$P(t) \leq X \exp[-r(T-t)] . \tag{4.5}$$

Lower Limits

To determine the lower limits of a call price, we construct two portfolios: A contains one call at price C and $X \exp[-r(T-t)]$ in cash; B contains one stock. At maturity, B is worth $S(T)$. If $S(T) > X$, the call in A is exercised, and A is worth $S(T)$ (X is used to buy the stock). If $S(T) < X$, the call option expires worthless, and portfolio A is worth X. The value of A is therefore $\max[S(T), X] \geq S(T)$, the value of B. This is valid for all times because the value of both portfolios depends only on the same source of uncertainty, the evolution of the stock price S. Consequently,

$$C(t) \geq \max\{S(t) - X \exp[-r(T-t)], 0\} . \tag{4.6}$$

The equivalent relation for a put,

$$P(t) \geq \max\{X \exp[-r(T-t)] - S(t), 0\} , \tag{4.7}$$

can be derived in a similar way, using one portfolio (C) containing the put option and the stock, and another (D) with $X \exp[-r(T-t)]$ in cash.

These limits, together with a sketch of the dependence of option prices on those of the underlying, are shown in Figs. 4.1 (call) and 4.2 (put). The arrows in Figs. 4.1 and 4.2 indicate how the curve is displaced, resp. distorted,

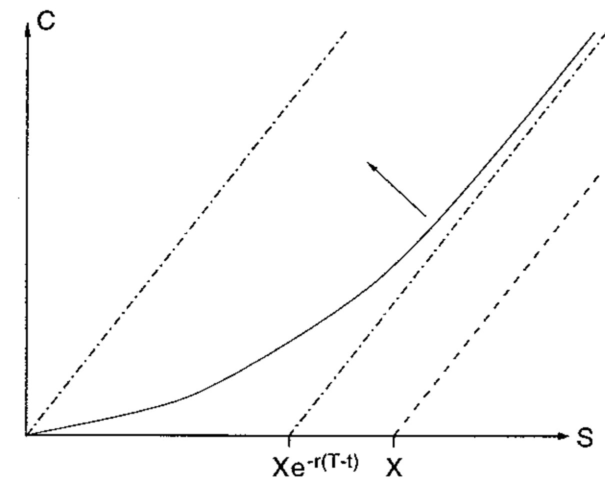

Fig. 4.1. Price limits for call options. The *curved* line sketches a realistic price curve. The *arrow* marks the direction of displacement of the curve when r, $T-t$, or the volatility (standard deviation) σ of the stock price increase

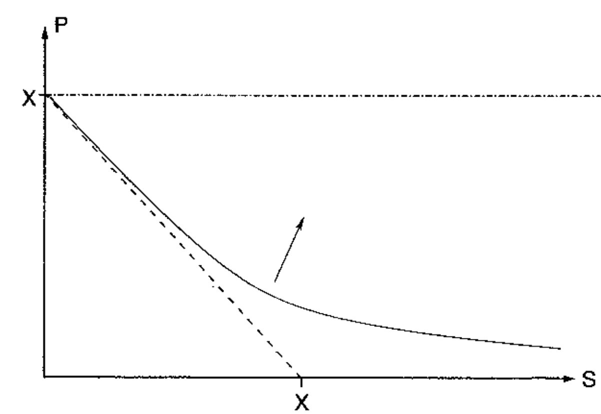

Fig. 4.2. Price limits for put options

when the interest rate r, the time to maturity $T - t$, or the volatility of the underlying stock (measured by the standard deviation σ of the stock price) change. An empirical investigation on 58 US stocks from August 1976 to June 1977, discussed by Hull [10], finds that the lower limits for calls, (4.6) and Fig. 4.1, were violated in 1.3% of the quotations. Out of these, 29% were corrected on the next quote while 71% were smaller than applicable transaction costs. Therefore no arbitrage was possible despite these limit violations.

Another important relation, put–call parity, can be derived by comparing the portfolios A and C:

$$C(t) + X \exp[-r(T-t)] = P(t) + S(t) . \tag{4.8}$$

This equation does not rely on any specific assumption on the options or on the prices of the underlying and therefore provides a rather stringent test on the correct operation (complete and efficient) of the markets. The empirical study cited by Hull [10] finds occasional violations of put–call parity on a 15-minute time scale.

Checking put–call parity simply from newspaper quotes may be more involved, as shown by the following example with options traded in late 1998 on the EUREX exchange. At $t = 1998/10/21$, call and put options on the Bayer stock with nominal maturity December 1998, i.e., $T = 1998/12/18$ and a strike price of $X =$ DM 65, were quoted $C =$ DM 2.38 and $P =$ DM 5.50. Bayer was quoted $S(1998/10/21) =$ DM 61.25. Assuming then $r = 3\%$ p.a., and $T - t = (1/6)y$, one has (in DM)

$$2.38 + 65 \exp(-0.005) = 67.05 \neq 66.75 = 5.50 + 61.25 . \tag{4.9}$$

It is not clear, however, that this is an actual violation of put–call parity. In particular, the assumption on r has been made ad hoc with rates relevant for savings accounts of a private consumer, and may not correspond to the market situation for institutional investors. Assuming put–call parity and calculating backwards, would give $r(T-t) = 0.01$, i.e., twice as much as used above, and then would certainly indicate an interest rate much higher than 3% p.a.

4.4 Modeling Fluctuations of Financial Assets

The question about the appropriate modeling of financial time series may well be answered differently by academics and practitioners. The basic approach taken by academics, and more generally all people with a skeptical attitude towards the financial markets, goes back to Bachelier and assumes some kind of random walk, or stochastic process. Essentially, this will be the attitude adopted in this book. Some aspects of random walks have been discussed

in Chap. 3. Others will be introduced below, together with a more general summary of important facts on stochastic processes.

Among the practitioners, traders and analysts classified as "chartists", practicing "technical analysis", would not share this opinion. This group of operators attempts to distinguish recurrent patterns in financial time series and tries to make profit out of their observation. The citation from Malkiel's book *A Random Walk Down Wall Street* reproduced in Chap. 1 testifies to this, as well as numerous books on technical analysis at different levels. However, the issue of correlations in financial time series is nontrivial. We shall discuss simple aspects in Sect. 5.3.2, but subtle aspects are still the subject of ongoing research. It has to be taken seriously because technical analysis is alive and well on the markets, and one therefore must conclude that some money can be earned this way, and that certain correlations indeed exist in financial data, perhaps even introduced by a sufficient number of traders following technical analysis even on purely random samples. Systematic studies of the profitability of technical analysis reach controversial conclusions, however [31].

4.4.1 Stochastic Processes

Classic references on stochastic processes are Cox and Miller, and Lévy [32]. There are two excellent books by J. Honerkamp, concerned with, or touching upon, stochastic processes [44], and presenting a more physics-oriented perspective.

We say that a variable with an unpredictable time evolution follows a *stochastic process.* The changes of this variable are drawn from a probability distribution according to some specified rules. *One* distinction of stochastic processes is made according to whether time is treated as a continuous or a discrete variable, and whether the stochastic variable is continuous or discrete. We will be rather sloppy on this distinction here.

Stochastic processes are described by the specification of their dynamics and of the probability distribution functions from which the random variables are taken. The dynamics is usually given by a stochastic difference equation such as, e.g.,

$$x(t+1) = x(t) + \varepsilon(t) \tag{4.10}$$

where x is the stochastic variable and ε is a random variable whose probability distribution must be specified, or by differential equations such as

$$\dot{x}(t) = ax(t) + b\varepsilon(t) , \tag{4.11}$$
$$\dot{x}(t) = ax(t) + bx(t)\varepsilon(t) . \tag{4.12}$$

Equation (4.11) describes "additive noise" because the random variable is added to the stochastic variable, and (4.12) describes "multiplicative noise".

Next, we must specify the probability distribution function of $\varepsilon(t)$, e.g.,

$$p(\varepsilon, t) = \frac{1}{\sqrt{2\pi\sigma^2 t}} \exp\left(-\frac{\varepsilon^2}{2\sigma^2 t}\right) . \tag{4.13}$$

Correlations in a stochastic process can be described either in its defining equation, e.g., by a dependence on earlier times [cf., e.g., the various autoregressive processes (4.44), (4.46) and (4.48) below], or by the *conditional probability*

$$p\left[x(t_1) = x_1 | x(t_0) = x_0, x(t_{-1}) = x_{-1}, \ldots\right] , \tag{4.14}$$

which measures the probability that the variable x takes the value x_1 at t_1 *provided that* x_0 has been observed at t_0 *and* x_{-1} at t_{-1}, etc. For a continuous variable, the conditional probability density $p[\ldots]\mathrm{d}x_1$ measures the probability that at t_1, $x_1 \le x \le x_1 + \mathrm{d}x_1$, provided x_0 has been observed at t_0, etc. The *unconditional probability* (or marginal probability) of observing x_1 at t_1, independently of earlier realizations of x, is then

$$\begin{aligned} &p\left[x(t_1) = x_1\right] \\ &= \int \mathrm{d}x_0 \mathrm{d}x_{-1} \ldots p\left[x(t_1) = x_1, x(t_0) = x_0, x(t_{-1}) = x_{-1}, \ldots\right] \quad (4.15\mathrm{a}) \\ &= \int \mathrm{d}x_0 \mathrm{d}x_{-1} \ldots p\left[x(t_1) = x_1 | x(t_0) = x_0, x(t_{-1}) = x_{-1}, \ldots\right] \\ &\quad \times\ p(x_0, t_0) p(x_{-1}, t_{-1}) \ldots \quad (4.15\mathrm{b}) \end{aligned}$$

where $p[\ldots]$ on the right-hand side of (4.15a) is the *joint probability* which measures the probability of observing x_1 at t_1 and x_0 at t_0, etc. It is related to the conditional probability (4.14) by the second equality (4.15b).

A stochastic process is stationary if

$$p(x, t) = p(x) , \tag{4.16}$$

and it is a martingale stochastic process if

$$E(x_1 | x_0, x_{-1}, \ldots) = \int \mathrm{d}x_1 x_1 p\left[x(t_1) = x_1 | x(t_0) = x_0, x(t_{-1}) = x_{-1}, \ldots\right] = x_0 , \tag{4.17}$$

where E is the expectation value conditioned on earlier observations x_0, x_{-1}, etc.

We now discuss a few important stochastic processes.

Markov Processes

For a Markov process, the next realization only depends on the present value of the random variable. There is no longer-time memory. For $\ldots t_{-2} \le t_{-1} \le t_0 \le t_1 \le \ldots$, a Markov process satisfies

$$p\left[x(t_1) = x_1 | x(t_0) = x_0, x(t_{-1}) = x_{-1}, \ldots\right] = p\left[x(t_1) = x_1 | x(t_0) = x_0\right] . \tag{4.18}$$

Markov processes obey the Chapman–Kolmogorov–Smoluchowski equation (3.10), derived by Bachelier [6].

For Markov processes in continuous time, one can take a short-time limit of the conditional probability distributions

$$p(x, t | x', t') \to \delta(x - x') \quad \text{for} \quad t \to t' , \tag{4.19}$$

and expand it around this limit in first order in $t - t'$:

$$p(x, t | x', t') \approx \left[1 - a(x, t)(t - t')\right] \delta(x - x') + (t - t') w(x, x', t) , \tag{4.20}$$

where $a(x, t)$ and $w(x, x', t)$ are expansion coefficients. $a(x, t)$ is the reduction, in first order in the time difference, of the initial "certainty", i.e., the weight of $\delta(x - x')$ due to the widening of the conditional probability distribution, and $w(x, x', t)$ quantifies precisely this effect in first order in $t - t'$. Inserting this expansion into the Chapman–Kolmogorov–Smoluchowski equation (3.10), one obtains the master equation

$$\frac{\partial p(x, t)}{\partial t} = \int \mathrm{d}x' w(x, x', t) p(x', t) - \int \mathrm{d}x' w(x, x', t) p(x, t) . \tag{4.21}$$

The first term on the right-hand side describes transitions $x' \to x$ at t, and the second term transitions $x \to x'$. We have made an integro-differential equation from the original convolution equation. In special situations, the master equation may reduce to a partial differential equation, the Fokker–Planck equation [37], which will be discussed in later in Chap. 6.

In finance, Markov processes are consistent with an efficient market. If this were not so, technical analysis would allow one to produce above-average profits. Conversely, to the extent that technical analysis generates consistent profits above the market return, the assumption of a Markov process for financial time series must be questioned.

The Wiener Process

The Wiener process, often also called the Einstein–Wiener process, or Brownian motion, is a particular Markov process with continuous variable and continuous time. It was formulated for the first time by Bachelier [6], and discussed on an elementary level in Sect. 3.2.2. If the stochastic variable is called z, its two important properties are:

1. Consecutive Δz are statistically independent.
2. Δz is given, for a small but finite time interval Δt, and for an infinitesimal interval $\mathrm{d}t$, by

$$\Delta z = \varepsilon \sqrt{\Delta t} \tag{4.22}$$

$$\mathrm{d}z = \varepsilon \sqrt{\mathrm{d}t} . \tag{4.23}$$

ε is drawn from a normal distribution

$$p(\varepsilon) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\varepsilon^2}{2}\right) \tag{4.24}$$

with zero mean and unit variance.

The passage from a Wiener process in discrete time to one in continuous time is illustrated in Fig. 4.3.

The conditions for a Wiener process are stronger than for a general Markov process, in that it uses independent, identically distributed (abbreviated: IID) random variables. Being independent, the correlations of the random numbers ε are

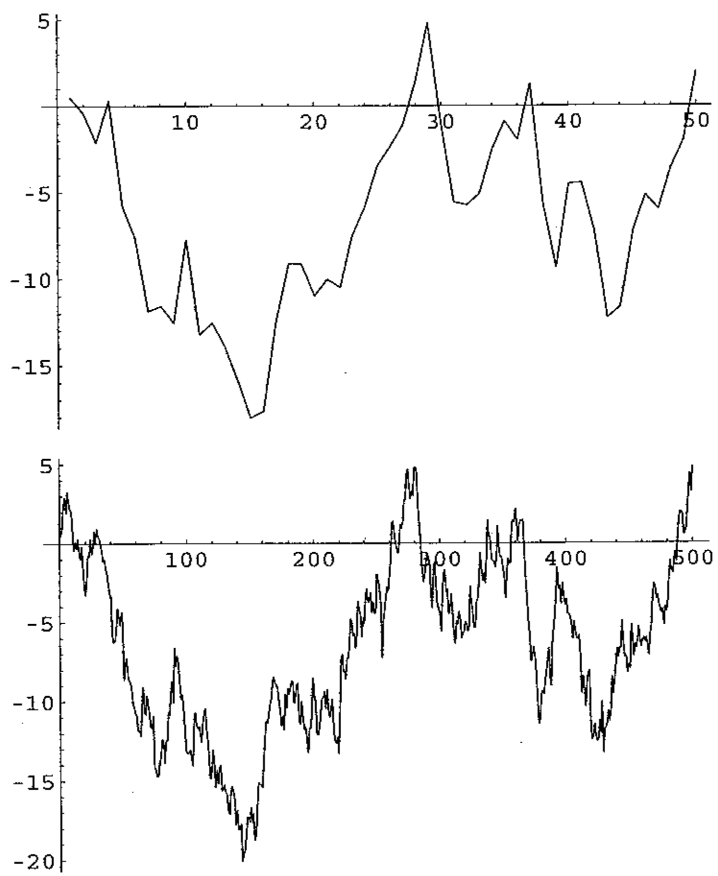

Fig. 4.3. Passage from discrete time to continuous time for a Wiener process. The increments were drawn from a normal distribution with zero mean and unit variance

$$\langle \varepsilon(t)\varepsilon(t')\rangle = \sigma^2 \begin{cases} \delta_{t,t'} \\ \delta(t-t') \end{cases} \tag{4.25}$$

where σ^2 is the variance of the underlying normal distribution. Its noise spectrum is

$$F(\omega) = \int_{-\infty}^{\infty} \mathrm{d}\tau \langle \varepsilon(t)\varepsilon(t+\tau)\rangle \mathrm{e}^{i\omega\tau} = \sigma^2 \,. \tag{4.26}$$

It is independent of frequency, and therefore "white noise". Often, this is also written as

$$\varepsilon(t) \in \mathrm{WN}(0,\sigma^2) \,. \tag{4.27}$$

W characterizes the random variables as "white noise", N denotes "normally distributed", and the arguments are the mean and variance. A stochastic process with an additive white noise term describes algebraic Brownian motion. Notice that some authors, e.g., Hull [10], prefer to take the standard deviation instead of the variance, as the second argument of WN in (4.27).

Equation (4.23) may seem very surprising for those who are not familiar with stochastic processes. It is to be interpreted in the sense of mean square fluctuations, resp. expectation values. A more detailed argument goes as follows. Let a stochastic process be defined by the differential equation

$$\frac{\mathrm{d}z(t)}{\mathrm{d}t} = \varepsilon(t) \tag{4.28}$$

where $\varepsilon(t)$ is the random variable. Then, the change $\mathrm{d}z(t)$ of the random variable z in an infinitesimal time interval $\mathrm{d}t$ is given by integration

$$\mathrm{d}z(t) = \int_t^{t+\mathrm{d}t} \mathrm{d}t' \varepsilon(t') \,. \tag{4.29}$$

For a nonstochastic variable, this integral would be trivial and given by $\mathrm{d}z = \varepsilon(t)\mathrm{d}t$. That this can't hold for a stochastic variable is clear from taking the expectation values of (4.29)

$$\langle \mathrm{d}z(t)\rangle = \int_t^{t+\mathrm{d}t} \mathrm{d}t' \langle \varepsilon(t')\rangle = 0 \,. \tag{4.30}$$

On the other hand, the expectation value of $(\mathrm{d}z)^2$ becomes

$$\langle \mathrm{d}z(t)\mathrm{d}z(t)\rangle = \int_t^{t+\mathrm{d}t} \mathrm{d}t_1 \mathrm{d}t_2 \langle \varepsilon(t_1)\varepsilon(t_2)\rangle = \sigma^2 \int_t^{t+\mathrm{d}t} \mathrm{d}t_1 = \sigma^2 \mathrm{d}t \,. \tag{4.31}$$

For the second equality, we have used (4.25), and the third equality obtains in the usual way because σ^2 is a nonstochastic quantity. These expectation values are consistent with $\mathrm{d}z = \varepsilon\sqrt{\mathrm{d}t}$, (4.23).

For a Wiener process, the *expectation value* of the stochastic variable in a small time interval vanishes

$$E(\Delta z) \equiv \langle \Delta z \rangle = \int_{-\infty}^{\infty} \mathrm{d}(\Delta z)\, \Delta z\, p(\Delta z) = 0\,. \tag{4.32}$$

Its *variance* is linear in Δt,

$$\mathrm{var}(\Delta z) = \int_{-\infty}^{\infty} \mathrm{d}(\Delta z)(\Delta z)^2 p(\Delta z) = \Delta t\,, \tag{4.33}$$

and its *standard deviation* behaves as

$$\sqrt{\mathrm{var}(\Delta z)} = \sqrt{\Delta t}\,. \tag{4.34}$$

Finite time intervals T may be considered as being composed of many small intervals ($T = N\Delta t$ fixed, as $N \to \infty$ and $\Delta t \to 0$), each of which corresponds to one time step of a Wiener process. For sums of normally distributed quantities, the mean values and variances are additive:

$$\langle z(T) - z(0) \rangle \quad = 0\,, \tag{4.35}$$

$$\mathrm{var}[z(T) - z(0)] = T\,, \tag{4.36}$$

and the standard deviation is $\sqrt{T}$.

The Wiener process may be generalized by superposing a drift $a\,\mathrm{d}t$ onto the stochastic process $\mathrm{d}z$

$$\mathrm{d}x = a\,\mathrm{d}t + b\,\mathrm{d}z\,. \tag{4.37}$$

For this *generalized Wiener process,* we have

$$\langle x(T) - x(0) \rangle \quad = aT\,, \tag{4.38}$$

$$\mathrm{var}[x(T) - x(0)] = b^2 T\,. \tag{4.39}$$

This generalized Wiener process is shown in Fig. 4.4.

A further generalization is the Itô process where the drift term and prefactor of the stochastic component depend on the random variable [$a \to a(x,t)$, $b \to b(x,t)$], i.e.,

$$\mathrm{d}x = a(x,t)\mathrm{d}t + b(x,t)\mathrm{d}z\,, \tag{4.40}$$

and $\mathrm{d}z = \varepsilon\sqrt{\mathrm{d}t}$ describes a Wiener process. The Itô process will play an important role in the standard model for stock prices.

Other Important Processes

For completeness, we discuss some more important stochastic processes or classification criteria.

1. **Self-similar stochastic processes** with index, or Hurst exponent, H are defined by

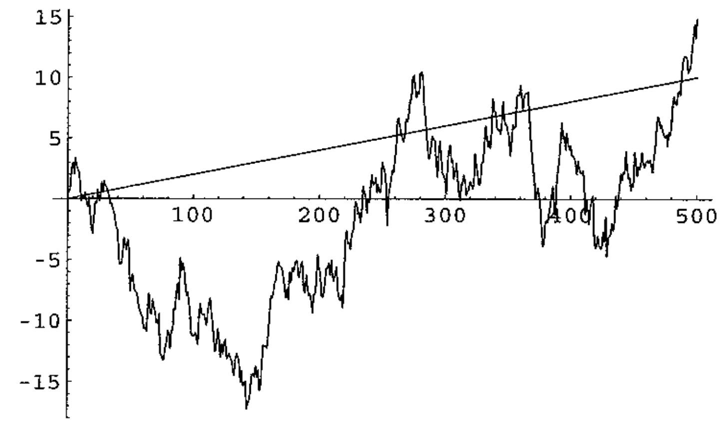

Fig. 4.4. The generalized Wiener process. The straight line shows the drift superposed on the data in the bottom panel of Fig. 4.3

$$p\,[x(at)] = p\left[a^H x(t)\right] \quad \text{with } a > 0\,. \tag{4.41}$$

A rescaling of time leads to a change in length scale, and there is no intrinsic scale associated with this process. Such a process violates (4.16) and therefore cannot be stationary. Brownian motion (cf. above) is self-similar with $H = 1/2$. However, the converse is not true: There are non-Gaussian stochastic processes with independent increments but $H = /2$ [45, 46].

2. In **fractional Brownian motion,** introduced by Mandelbrot [47], the random variables are not uncorrelated, and therefore describe "colored noise". The construction is done starting from ordinary Brownian motion $\mathrm{d}z \equiv \mathrm{d}z(t)$, (4.23), and a parameter H satisfying $0 < H < 1$. Then fractional Brownian motion of exponent H is essentially a moving average over $\mathrm{d}z(t)$ in which past increments of $z(t)$ are weighted by a power-law kernel $(t-s)^{H-1/2}$. Mandelbrot and van Ness define fractional Brownian motion of exponent H, $B_H(t)$, as [47]

$$B_H(t) = B_H(0) + \frac{1}{\Gamma(H+\frac{1}{2})}\left\{\int_{-\infty}^{0}\left[(t-s)^{H-1/2} - (-s)^{H-1/2}\right]\mathrm{d}z(s) + \int_0^t (t-s)^{H-1/2}\mathrm{d}z(s)\right\} \tag{4.42}$$

$B_H(0)$ is an arbitrary initial starting position. For $H = 1/2$, fractional Brownian motion reduces to ordinary Brownian motion.
The ranges $H < 1/2$ and $H > 1/2$ are very different. For $H < 1/2$, the paths look less ragged than ordinary Brownian motion, and the variations are "antipersistent" (positive variations preferentially followed by

negative ones). $H > 1/2$ is the "persistent" regime, i.e., there are positive correlations, and the paths are significantly rougher than Brownian motion. Notice that the paths of fractional Brownian motion are continuous but not differentiable.

3. **Lévy processes** are treated in greater detail in Sect. 5.4. The IID random variable $\varepsilon(t)$ is drawn from a stable Lévy distribution. Unlike the Gaussian distribution, Lévy distributions decay as power laws

$$L_\mu(x) \sim \frac{\mu A^\mu}{|x|^{1+\mu}}, \quad |x| \to \infty . \tag{4.43}$$

They are stable, i.e., form-invariant under addition, when $0 < \mu < 2$. Large events being more probably by orders of magnitude than under a Gaussian, the corresponding stochastic process possesses frequent discontinuities.

4. **Autoregressive processes** are non-Markovian. The equation of motion contains memory terms which depend on past values of variables. The equation

$$x(t) = \sum_{k=1}^{p} \alpha_k x(t-k) + \varepsilon(t) + \sum_{k=1}^{q} \beta_k \varepsilon(t-k) \tag{4.44}$$

describes an autoregressive, moving average, ARMA(p,q), process. It depends on the past p realizations of the stochastic variable x, and on the past q values of the random number ε. ARMA(p,q) processes can be interpreted as stochastically driven oscillators and relaxators [44].
Variants thereof, the ARCH and GARCH processes, are important in econometrics and finance [13]. The acronyms stand for autoregressive [process with] conditional heteroscedasticity, and generalized autoregressive [process with] conditional heteroscedasticity. Heteroscedasticity means that the variance of the process is not constant but depends on random variables. To be specific, an ARCH(q) process [48] is defined by (4.22) with

$$\varepsilon(t) \quad \in \mathrm{WN}\left[0, \sigma^2(t)\right] \tag{4.45}$$

$$\sigma^2(t) = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon^2(t-i) , \tag{4.46}$$

and a GARCH(p,q) process [49] by (4.22) with

$$\varepsilon(t) \quad \in \mathrm{WN}\left[0, \sigma^2(t)\right] \tag{4.47}$$

$$\sigma^2(t) = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon^2(t-i) + \sum_{i=1}^{p} \beta_i \sigma^2(t-i) . \tag{4.48}$$

In both cases, the random variable is drawn from a normal distribution with zero mean and a *time-dependent* variance $\sigma^2(t)$ which depends on the last q realizations of the random variable ε and, for the GARCH(p,q) process, in addition on the last p values of the variance σ^2.

4.4.2 The Standard Model of Stock Prices

Bachelier modeled stock or bond prices by a random walk superimposed on a constant drift (with the exception of the liquidation days where coupons were detached from the bonds, or the maturity dates of the futures where a prolongation fee had to be paid eventually). The drift was further eliminated from the problem by considering the equivalent martingale process as the fundamental variable, i.e., a Wiener process with zero mean and a variance increasing linearly in time.

There are two problems with this proposal:

1. The stock or bond prices in the model may become negative, in principle, when the changes $\Delta S(T)$ accumulated over a time interval T exceed the starting price $S(0)$. While this is not likely in practical situations, it should be a point of concern, in principle.
2. In Bachelier's model, the profit of an investment into a stock with price S over a time interval T is

$$\langle S(T) - S(0) \rangle = \frac{\mathrm{d}S}{\mathrm{d}t} T , \tag{4.49}$$

where $\mathrm{d}S/\mathrm{d}t$ is the drift which was assumed fixed and independent of S. More important than the *profit*, for an investor, will be the *return* on his capital invested. An investor will require that the return of an investment will be independent of the price of the asset (in other words, if a return of 15% p.a. is required when a stock is at \$40, it will also be required at \$65). This can be written as

$$\mathrm{d}S = \mu S \mathrm{d}t , \tag{4.50}$$

giving $S(t) = S_0 \mathrm{e}^{\mu t}$ where μ is the return rate, and $\mu \Delta t$ the return over a time interval Δt. This has consequences for the risk of an investment, measured by the standard deviation or – in financial contexts – volatility of asset prices. (Being careful, one should distinguish between variances accumulated over certain time intervals, or variance rates, entering the stochastic differential equations, resp. the corresponding quantities for the standard deviations.) A reasonable requirement is that the variance of the returns $\mu = \mathrm{d}S/\mathrm{d}t$ should be independent of S, i.e., that the uncertainty on reaching the 15% return discussed above, is the same regardless of whether the stock price is at \$40 or 80\$. This implies that, over a time interval Δt

$$\sigma^2 \Delta t = \mathrm{var}\left(\frac{\Delta S}{S}\right) \tag{4.51}$$

is independent of the stock price, or that

$$\mathrm{var}(S) = \sigma^2 S^2 \Delta t . \tag{4.52}$$

These requirements suggest that the asset price can be represented as an Itô process

$$dS = \mu S dt + \sigma S dz \ , \ \text{resp.} \quad \frac{dS}{S} = \mu dt + \sigma dz = \mu dt + \sigma\varepsilon\sqrt{dt} \qquad (4.53)$$

with instantaneous drift and standard deviation rates μ and σ. In other words,

$$\frac{dS}{S} \in \text{WN}(\mu dt, \sigma^2 dt) \ , \qquad (4.54)$$

i.e., dS/S is drawn from a normal distribution with mean μdt and standard deviation $\sigma\sqrt{dt}$. Concerning (4.53) and (4.54), notice that

$$\frac{dS}{S} \neq d \ln S \quad \text{for stochastic variables.} \qquad (4.55)$$

The process (4.53) is referred to as *geometric Brownian motion.* S follows a stochastic process subject to multiplicative noise. It avoids the problem of negative stock prices, and apparently is in better agreement with observations.

Notice that the model of stock prices following geometric Brownian motion (4.53) must be considered as a hypothesis which has to be checked critically, and not as an established and universal theory. A critical comparison to empirical market data will be given in Chap. 5. For a superficial comparison, Fig. 4.5 shows the chart of the Commerzbank through the year 1997. This chart is not primarily shown for supportive purposes. More intended to inspire caution, it demonstrates the enormous variety of behavior encountered even for a single blue chip stock, which contrasts with the simplicity of the postulated standard model (4.53). While a priori the parameters μ and σ of the standard model are taken as constants, Fig. 4.5 suggests that this may be a valid approximation – if ever – only over limited time spans. The annualized volatility is 33.66%, and the drift during this year is $\mu = 82\%$. As is apparent from the figure, μ and σ in practice depend on time, and on shorter time scales in the course of the year they may be rather far from the values cited. Analyses taking μ and σ constant will only have a finite horizon of application. This observation has been an important motivation for the study of the ARCH and GARCH processes discussed in Sect. 4.4.1. Due to its simplicity, and the fundamental insights it allows, we shall use the model of geometric Brownian motion in the remainder of this chapter to develop a theory of option pricing. To do so, we must know, however, some properties of functions of stochastic variables.

4.4.3 The Itô Lemma

If we assume that the price process of a financial asset follows a stochastic process, the process followed by a derivative security, such as an option, will

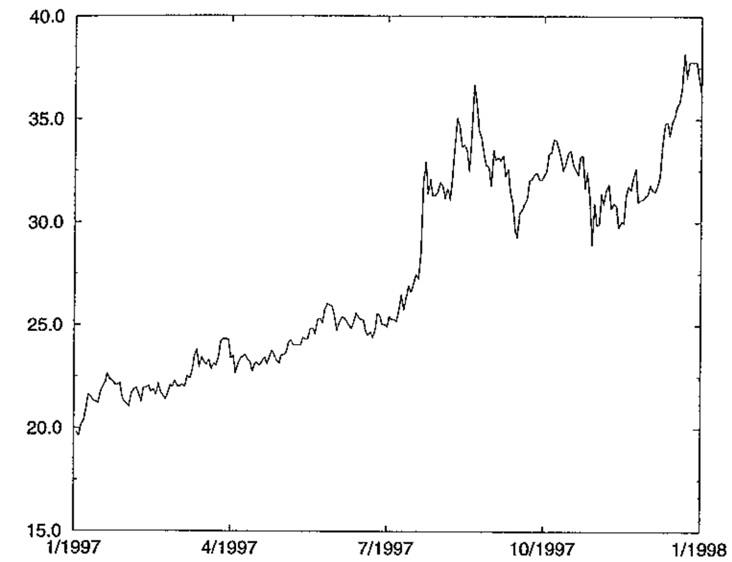

Fig. 4.5. Chart of the Commerzbank share from 1/1/1997 to 31/12/1997. The price has been converted to Euros. The volatility is $\sigma = 33.66\%$

(i) again be stochastic, and (ii) be a function of the price of the underlying. We therefore must know the properties of functions of stochastic variables.

An important result here, and the only one we need for future development, is a lemma due to Itô. Let $x(t)$ follow an Itô process, (4.40),

$$dx = a(x,t)dt + b(x,t)dz = a(x,t)dt + b(x,t)\varepsilon\sqrt{dt} \ . \qquad (4.56)$$

Then, a function $G(x,t)$ of the stochastic variable x and time t also follows an Itô process, given by

$$dG = \left(\frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} b^2 \frac{\partial^2 G}{\partial x^2} \right) dt + b \frac{\partial G}{\partial x} dz \ . \qquad (4.57)$$

The drift of the Itô process followed by G is given by the first term on the right-hand side in parentheses, and the standard deviation rate is given by the prefactor of dz in the second term.

There is a handwaving way to motivate the different terms in (4.57). We attempt a Taylor expansion of $G(x + dx, t + dt)$ about $G(x,t)$ to first order in dt. The first order expansion in dx produces the first and the last terms on the right-hand side of (4.57), and the first order expansion in dt produces the second term. Stopping the expansion at this stage would not be consistent, however, because dx contains a terms proportional to $\sqrt{dt}$, shown explicitly in (4.56). The second-order expansion in dx therefore produces another contribution of first order in $\mathrm{d}t$, the third term on the right-hand side of (4.57). That this term

$$\frac{1}{2}b^2\frac{\partial^2 G}{\partial x^2}\varepsilon^2\mathrm{d}t$$

is nonstochastic, and given correctly in (4.57), can be shown in a spirit similar to the argument in Sect. 4.4.1. Take the expectation value of $\varepsilon^2\mathrm{d}t$

$$\langle\varepsilon^2\mathrm{d}t\rangle = \langle\varepsilon^2\rangle\mathrm{d}t = \mathrm{d}t\ , \tag{4.58}$$

where the last equality follows from $\varepsilon \in \mathrm{WN}(0,1)$. On the other hand, its variance,

$$\mathrm{var}(\varepsilon^2\mathrm{d}t) = \langle\varepsilon^4\mathrm{d}t^2\rangle - \langle\varepsilon^2\mathrm{d}t\rangle^2 = \left(\langle\varepsilon^4\rangle - 1\right)\mathrm{d}t^2 \tag{4.59}$$

tends to zero more quickly than the mean, as $\mathrm{d}t \to 0$. Consequently, $\varepsilon^2\mathrm{d}t$ represents a sharp variable.

A full proof of this lemma is the subject of stochastic analysis and will not be given here. Applications will be given in the following sections.

4.4.4 Log-normal Distributions for Stock Prices

We now derive the probability distribution for the stock prices, based on the assumption of geometric Brownian motion. To do that, we start from the stochastic differential equation (4.53) for the price changes

$$\mathrm{d}S = \mu S\mathrm{d}t + \sigma S\mathrm{d}z\ , \tag{4.60}$$

and apply the Itô lemma with $G(S,t) = \ln S(t)$ [remember (4.55)!]

$$\frac{\partial G}{\partial S} = \frac{1}{S}\ ,\quad \frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2}\ ,\quad \frac{\partial G}{\partial t} = 0 \quad \Rightarrow \tag{4.61}$$

$$\mathrm{d}G = \left(\mu - \frac{\sigma^2}{2}\right)\mathrm{d}t + \sigma\mathrm{d}z\ . \tag{4.62}$$

With μ = const. and σ = const., $\ln S$ follows a generalized Wiener process with an effective drift $\mu - \sigma^2/2$ and standard deviation rate σ. Notice that both S and G are affected by the same source of uncertainty: the stochastic process $\mathrm{d}z$. This will become important in the next section, where S and G will represent the prices of the underlying and the derivative securities, respectively. [As is clear from (4.53), $\mathrm{d}S/S$ also follows, under the same assumptions, a generalized Wiener process, however with an unrenormalized drift μ. This illustrates (4.55). The consequences will be discussed below.]

If t denotes the present time, and T some future time, the probability distribution of $\ln S$ will be a normal distribution with mean and variance

$$\langle\ln S\rangle = \left(\mu - \frac{\sigma^2}{2}\right)(T-t) \tag{4.63}$$

$$\mathrm{var}(\ln S) = \sigma^2(T-t)\ , \tag{4.64}$$

i.e.,

$$p(\ln S_T/S_t) = \frac{1}{\sqrt{2\pi\sigma^2(T-t)}}\exp\left(-\frac{\left[\ln\left(\frac{S_T}{S_t}\right) - \left(\mu - \frac{\sigma^2}{2}\right)(T-t)\right]^2}{2\sigma^2(T-t)}\right)\ . \tag{4.65}$$

The stock price changes themselves are then distributed according to a log-normal distribution [use $p(\ln S_T/S_t)\mathrm{d}\ln S_T/S_t = \tilde{p}(S_T)\mathrm{d}S_T$]

$$\begin{aligned}\tilde{p}(S_T) &= \frac{1}{S_T}p\left(\ln\frac{S_T}{S_t}\right)\\ &= \frac{1}{\sqrt{2\pi\sigma^2(T-t)}}\frac{1}{S_T}\exp\left(-\frac{\left[\ln\left(\frac{S_T}{S_t}\right) - \left(\mu - \frac{\sigma^2}{2}\right)(T-t)\right]^2}{2\sigma^2(T-t)}\right)\ .\end{aligned} \tag{4.66}$$

This distribution is shown in Fig. 4.6.

Using this distribution and the substitution $S_T/S_t = \exp(\omega)$, we find that the expectation value of S_T evolves as

$$\langle S_T\rangle = \int_0^\infty \mathrm{d}S_T\ S_T\,\tilde{p}(S_T) = S_t\exp[\mu(T-t)]\ , \tag{4.67}$$

and its variance as

$$\mathrm{var}(S_T) = S_t^2\exp[2\mu(T-t)]\{\exp[\sigma^2(T-t)] - 1\}\ . \tag{4.68}$$

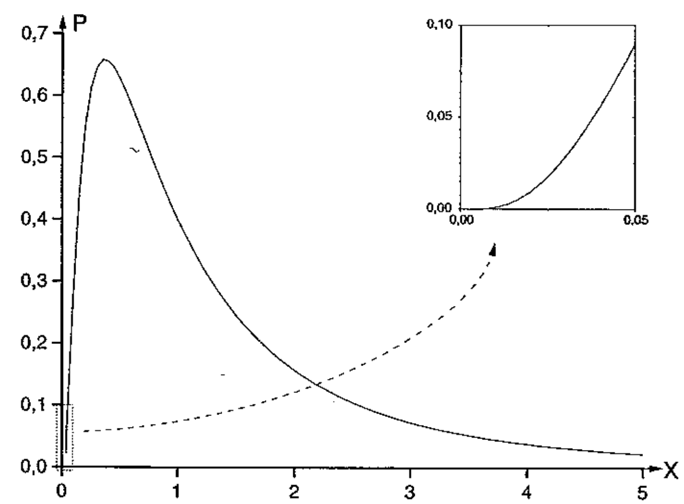

Fig. 4.6. The log-normal distribution $\tilde{p}(S)$

Observe that the expectation value of S_T grows with a rate μ, $\ln\langle S_T\rangle \propto \mu(T-t)$, in line with the definition of μ as the expectation value of the rate of return. Notice, however, that from (4.63), the expectation value of $\ln S$ grows with a different rate $\mu - \sigma^2/2$. The two different results correspond to two different situations where return rates are measured. Equation (4.53) shows that μ is the average of the return rate over a short time interval. The expectation value of the stock price grows with the average return rate over short time intervals. On the other hand, if one takes an actual investment with a specific return rate history with the same average, and calculates its, say yearly, return, this will be less than the average of the yearly returns determined on the way. For a specific example, assume an average growth rate of 10% p.a. over four years. Then, the expected price of the stock after four years is $S_T = S_t(1.1)^4 = 1.464S_t$. Now assume that the actual growth rates in the four years are $\mu_1 = 5\%$, $\mu_2 = 12\%$, $\mu_3 = 13\%$, $\mu_4 = 10\%$. Then $S_T = 1.05 \times 1.12 \times 1.13 \times 1.1S_t = 1.462S_t$, and the actual rate of return over the four years is only 9.5% p.a. If many such investments at a given average return rate μ are considered and their returns are averaged over, the average rate of return will converge to $\mu - \sigma^2/2$. Moreover, the binomial theorem $(1+x)(1-x) = 1 - x^2 \le 1$ shows that the average short-term growth rate can only be reached in the absence of randomness ($x = 0$), and that the general conclusion is independent of the particular realization assumed in the example. Of course, this is common experience of any investor who determines the return of his investments.

Another way of looking at the different return rates is to notice that, due to the skewness of the log-normal distribution, the rather frequent small prices from negative returns are less weighted in the expectation value than the less frequent very high prices from positive returns. Few very high profits count more in the expectation value than the same number of almost total losses, while the opposite is true for an actual investment history with the same short-time return rate.

4.5 Option Pricing

4.5.1 The Black–Scholes Differential Equation

We now turn to the pricing of options, and the hedging of positions involving options. Investments in options are usually considered to be risky, significantly more risky than investments into stocks or bonds. This is because of the finite time to maturity, the high volatility of options (significantly higher than the volatility of its underlying), and the possibility of a total loss of the invested capital for the long position, and losses even potentially unlimited for the short position, in the case of unfavorable market movements (cf. the discussion in Sect. 2.4, and Figs. 2.1 and 2.2). With f the price of an option ($f = C, P$, for call and put options, respectively), we have

$$\sqrt{\mathrm{var}\,(\Delta f)} = \frac{\partial f}{\partial S}\sqrt{\mathrm{var}\,(\Delta S)}\,, \quad \sqrt{\mathrm{var}\left(\frac{\Delta f}{f}\right)} = \frac{S}{f}\frac{\partial f}{\partial S}\sqrt{\mathrm{var}\left(\frac{\Delta S}{S}\right)} \qquad (4.69)$$

for the volatility of the option in terms of the volatility of the underlying. Figs. (4.1) and (4.2) show that $\partial f/\partial S < 1$ in general. While the volatility of the option *prices* is smaller than that of the prices of the underlyings, the volatility of the option *returns* described by the second equation in (4.69) is much higher than that of the returns of their underlyings because the option prices usually are much lower than the prices of the underlyings, $S/f \gg 1$.

Moreover, the writer of an option engages a liability when entering the contract, while the holder has a freedom of action depending on market movement, i.e., an insurance: buy or not buy (sell or not sell) the underlying at a fixed price, in the case of a call (put) option. The question then is: What is the risk premium for the writer of the option, associated with the liability taken over? Or what is the price of the insurance, the additional freedom of choice for the holder? What is the value of the asymmetry of the contract?

These questions were answered by Black and Scholes [42] and Merton [43], and the answer they came up with, under the assumptions specified in Sect. 4.2.1 and developed thereafter, i.e., geometric Brownian motion, is surprising: *There is no risk premium required for the option writer!* The writer can entirely eliminate his risk by a dynamic and self-financing hedging strategy using the underlying security only. The price of the option contract, the value for the long position, is then determined completely by some properties of the stock price movements (volatility) and the terms of the option contract (time to maturity, strike price). For simplicity, and because we are interested only in the important qualitative aspects, we shall limit our discussion to European options, mostly calls, and ignore dividend payments and other complications. For other derivatives or more complex situations, the reader should refer to the literature [10, 12]–[15].

The main idea underlying the work of Black, Merton, and Scholes [42, 43] is that it is possible to form a *riskless* portfolio composed of the option to be priced and/or hedged, and the underlying security. Being riskless, it must earn the risk-free interest rate r, in the absence of arbitrage opportunities. The formation of such a riskless portfolio is possible because, and only because, at any instant of time the option price f is correlated with that of the underlying security. This is shown by the solid lines in Figs. 4.1 and 4.2, which sketch the possible dependences of option prices on the prices of the underlying. The dependence of the option price on that of the underlying is given by $\Delta = \partial f/\partial S$ which, of course, is a function of time. In other words, both the stock and the option price depend on the same source of uncertainty, resp. the same stochastic process: the one followed by the the stock price. Therefore the stochastic process can be eliminated by a suitable linear combination of both assets.

To make this more precise, we take the position of the writer of a European call. We therefore form a portfolio composed of

1. a short position in one call option,
2. a long position in $\Delta = \partial f/\partial S$ units of the underlying stock. Notice that Δ fluctuates with the stock price, and a continuous adjustment of this position is required.

The stochastic process followed by the stock is assumed to be geometric Brownian motion, (4.53),

$$\mathrm{d}S = \mu S \mathrm{d}t + \sigma S \mathrm{d}z \,. \tag{4.70}$$

A priori, we do not know the stochastic process followed by the option price. We know, however, that it depends on the stock price, and therefore, we can use Itô's lemma, (4.57),

$$\mathrm{d}f = \left(\frac{\partial f}{\partial S}\mu S + \frac{\partial f}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} \right) \mathrm{d}t + \frac{\partial f}{\partial S}\sigma S \mathrm{d}z \,. \tag{4.71}$$

The value of our portfolio is

$$\Pi = -f + \frac{\partial f}{\partial S} S \,, \tag{4.72}$$

and it follows the stochastic process

$$\mathrm{d}\Pi = -\mathrm{d}f + \frac{\partial f}{\partial S}\mathrm{d}S = \left(-\frac{\partial f}{\partial t} - \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} \right) \mathrm{d}t \,. \tag{4.73}$$

Notice that the stochastic process $\mathrm{d}z$, the source of uncertainty in the evolution of both the stock and the option prices, no longer appears in (4.73). Moreover, the drift μ of the stock price has disappeared, too. Eliminating the risk from the portfolio also eliminates the possibilities for profit, i.e., the risk premium $\mu > r$ associated with an investment into the underlying security alone (an investor will accept putting his money in a risky asset only if the return is higher than for a riskless asset). The portfolio being riskless, it must earn the risk-free interest rate r,

$$\mathrm{d}\Pi = r\Pi \mathrm{d}t = r \left(-f + \frac{\partial f}{\partial S} S \right) \mathrm{d}t \,. \tag{4.74}$$

Equating (4.73) and (4.74), we obtain

$$\frac{\partial f}{\partial t} + rS\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf \,, \tag{4.75}$$

the Black–Scholes (differential) equation. This is a linear second-order partial differential equation of parabolic type. Its operator structure is very similar to the Fokker–Planck equation in physics or the Kolmogorov equation in mathematics (two different names for the same equation) [37]. There are two differences, however: (i) the sign of the term corresponding to the diffusion

constant is negative, and (ii) this is a differential equation for a (at present rather arbitrary) function f while the Fokker–Planck equation usually refers to a differential equation for a normalized distribution function $p(x,t)$ whose norm is conserved in the time evolution. (For the use of Fokker–Planck equations in the statistical mechanics of capital markets, see Chap. 6).

For a complete solution to the Black–Scholes equation, we still have to specify the boundary or initial conditions. Unlike physics, here we deal with a final value problem. At maturity $t = T$, we know the prices of the call and put options, (4.6) and (4.7),

$$\left. \begin{array}{l} \mathrm{Call}: \; f = C = \max(S - X, 0) \\ \mathrm{Put}: \; f = P = \max(X - S, 0) \end{array} \right\} \; t = T \,. \tag{4.76}$$

The solution of this final value problem, (4.75) and (4.76) will be given in the next section. Notice that for second-order partial differential equations, the number and type of conditions (initial, final, boundary) required for a complete specification of the solution depends on the type of problem considered. For diffusion problems such as (3.20), (3.30), or (4.75), a single initial or final condition is sufficient.

Stock prices change with time. Keeping the portfolio riskless in time therefore requires a *continuous* adjustment of the stock position $\Delta = \partial f/\partial S$, as it varies with the stock price. It is clear that this can only be done in the idealized markets considered here, and subject to the assumptions specified earlier. Transaction costs, e.g., would prevent a continuous adjustment of the portfolio, and immediately make it risky. The same applies to credit costs incurred by the adjustments. In practice, therefore, a riskless portfolio will usually not exist, and there will be a finite risk premium on options (often determined empirically by the writing institutions).

The important achievement of Black, Merton, and Scholes was to show that, in idealized markets, the risk associated with an option can be hedged away completely by an offsetting position in a suitable quantity Δ of the underlying security (this hedging strategy is therefore called Δ-hedging), and that no risk premium need be asked by the writer of an option. The hedge can be maintained dynamically, and is self-financing, i.e., does not generate costs for the writer. Of course, this is an approximation in practice because none of the assumptions on which the Black–Scholes equation is based, are fulfilled. This will be discussed in Chap. 5. Despite this limitation, it allows fundamental insights into the price processes for derivatives, and we now proceed to solve the equation.

4.5.2 Solution of the Black–Scholes Equation

The following solution of (4.75) essentially follows the original Black–Scholes article [42], and consists in a reduction to a 1D diffusion equation with special boundary conditions. (This may not be too surprising: Fisher Black held a degree in physics.)

We substitute

$$f(S,t) = \mathrm{e}^{-r(T-t)} y(u,v) \,, \tag{4.77}$$

$$u = \frac{2\rho}{\sigma^2}\left(\ln\frac{S}{X} + \rho[T-t]\right) \,,$$

$$v = \frac{2}{\sigma^2}\rho^2(T-t) \,, \quad \rho = r - \frac{\sigma^2}{2} \,. \tag{4.78}$$

Then, the derivatives $\partial f/\partial S$, $\partial^2 f/\partial S^2$, and $\partial f/\partial t$ are expressed through $\partial y/\partial u$, $\partial y/\partial v$, etc., and $y(u,v)$ satisfies the 1D diffusion equation

$$\frac{\partial y(u,v)}{\partial v} = \frac{\partial^2 y(u,v)}{\partial u^2} \,. \tag{4.79}$$

The boundary conditions (4.76) for a call option translate into

$$y(u,0) = \begin{cases} 0 & u < 0 \\ X\left(\mathrm{e}^{u\sigma^2/2\rho} - 1\right) & u \geq 0 \,. \end{cases} \tag{4.80}$$

Diffusion equations are solved by Fourier transform in the spatial variable(s)

$$y(u,v) = \int_{-\infty}^{\infty} \mathrm{d}q \mathrm{e}^{\mathrm{i}qu} y(q,v) \,, \tag{4.81}$$

reducing (4.79) to an ordinary differential equation in v with the solution

$$y(q,v) = y(q,0)\exp\left(-q^2 v\right) \,. \tag{4.82}$$

$y(q,0)$, formally, is given by the Fourier transform of the boundary conditions (4.80) which, however, should NOT be performed explicitly. The trick, instead, is to transform the solution (4.82) back to u-variables, giving a convolution integral

$$y(u,v) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{d}w \, y(w,0) f(u-w) \quad \text{with} \quad f(x) = \sqrt{\frac{\pi}{v}}\exp\left(-\frac{x^2}{4v}\right) \,. \tag{4.83}$$

Another substitution $z = (w-u)/\sqrt{2v}$ almost gives the final result

$$y(u,v) = \frac{X}{\sqrt{2\pi}}\int_{-u/\sqrt{2v}}^{\infty} \mathrm{d}z \mathrm{e}^{-z^2/2}\left[\exp\left(\frac{\sigma^2}{2\rho}\left\{\sqrt{2v}z + u\right\}\right) - 1\right] \,. \tag{4.84}$$

The only task remaining is to complete the square in the exponent, and insert all substituted quantities. This gives the Black–Scholes equation for a European call option (remember that the boundary conditions for a call have been used in the derivation)

$$C(S,t) \equiv f(S,t) = SN(d_1) - X\mathrm{e}^{-r(T-t)}N(d_2) \,. \tag{4.85}$$

The equivalent solution for a European put option is

$$P(S,t) = X\mathrm{e}^{-r(T-t)}N(-d_2) - SN(-d_1) \,. \tag{4.86}$$

$N(d)$ is the cumulative normal distribution

$$N(d) = \frac{1}{\sqrt{2\pi}}\int_{-d}^{\infty} \mathrm{d}x \mathrm{e}^{-x^2/2} \,, \tag{4.87}$$

and its two arguments in (4.85) are given by

$$d_1 = \frac{\ln\frac{S}{X} + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \,, \tag{4.88}$$

$$d_2 = \frac{\ln\frac{S}{X} + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \,. \tag{4.89}$$

Clearly, $S \equiv S(t)$. The behavior of $C(S)$ is sketched in Fig. 4.1 as the solid line, and the equivalent put price is sketched in Fig. 4.2. The time evolution of a call price, as given by the Black–Scholes equation (4.85), is displayed in Fig. 4.7. In that figure, all parameters have been kept fixed, and only time elapses. We therefore monitor the *time value* of the options. The intrinsic value is given by $S(t) - X$, i.e., the payoff if the option was exercised today. While the intrinsic value fluctuates with the evolution of the stock price,

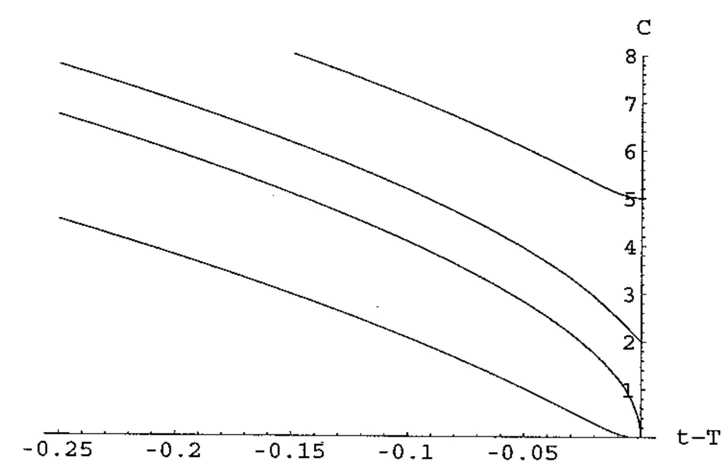

Fig. 4.7. Time evolution of the price of a European call option as a function of time *before* maturity in years. Fixed stock price $S = 100$, interest rate $r = 6\%/y$, and volatility $\sigma = 30\%/\sqrt{y}$ have been assumed. The curves represent different strike prices $X = 95, 98, 100, 105$ from top to bottom, i.e., the options are in the money (top two lines), at the money, and out of the money, respectively

the time value always decreases. It measures the probability left at time t for a favorable stock price movement to occur before maturity T. It varies strongest for options at the money, and less for options far in or out of the money.

There are a few interesting limiting cases of (4.85). If $S \gg X$, the option is exercised almost certainly. In this case, it will become equivalent to a forward contract with a delivery price X. If $S \gg X$, $d_1, d_2 \to \infty$, and $N(d_{1,2}) \to 1$. The Black–Scholes equation then reduces to

$$f(S,t) = S - X\mathrm{e}^{-r(T-t)} . \tag{4.90}$$

This was precisely the expression for the value of the long position in a forward contract derived earlier, (4.2). In that problem, the delivery price was to be fixed so that the value of the contracts for both parties came out to $f = 0$. Here, the strike price of the option is fixed from the outset, and f therefore represents the intrinsic value of the long position in the option, which has become equivalent to a forward by the assumption $S \gg X$. Notice that S must be exponentially large compared to X for our derivation to hold.

If $\sigma \to 0$, the stock becomes almost riskless. In (4.85), two different cases must be considered. If $\ln(S/X) + r(T-t) > 0$, $d_{1,2} \to \infty$, $N(d_i) \to 1$, and (4.90) continues to hold. If, on the other hand, $\ln(S/X) + r(T-t) < 0$, $d_{1,2} \to -\infty$, $N(d_i) \to 0$, and $f(S,t) \to 0$. Putting both cases together,

$$C(S,t) \equiv f(S,t) = \max(S - X\mathrm{e}^{-r(T-t)}, 0) . \tag{4.91}$$

If on the other hand, the stock is almost riskless, it will grow from S to $S_T = S\mathrm{e}^{r(T-t)}$ in the time interval $T-t$ almost deterministically. The value of the option at maturity is $\max(S_T - X, 0)$, and a factor $\exp[-r(T-t)]$ must be applied to discount this value to the present day, showing that (4.91) gives a consistent result also in this limit.

The different terms in (4.85) have an immediate interpretation if the term $\exp[-r(T-t)]$ is factored out:

1. $N(d_2)$ is the probability for the exercise of the option, $P(S_T > X)$, *in a risk-neutral world* (cf. below), i.e., where the actual drift of a financial time series can be replaced by the risk-free rate r.
2. $XN(d_2)$ is then the strike price times the probability that it will be paid, i.e., the expected amount of money to be paid under the option contract.
3. $SN(d_1)\exp[r(T-t)]$ is the expectation value of $S_T\Theta(S_T - X)$ in a risk-neutral world, i.e., the expected payoff under the option contract.
4. The difference of this term with $XN(d_2)$ then is the profit expected from the option. The prefactor $\exp[-r(T-t)]$ factored out discounts that profit, realized at maturity T, down to the present day t. The option price is precisely this discounted difference.

This interpretation is consistent with the capital asset pricing model which deals with the relation of risk and return in market equilibrium. It states that the expected return on an investment is the discounting rate which one must apply to the profit expected at maturity, in order to obtain the present price. In our interpretation of (4.85), one would just read this sentence from the backwards.

For an option, no specific risk premium is necessary. The entire risk is contained in the price of the underlying security, and can be hedged away.

Because of their importance, we reiterate some statements made in earlier sections, or implicitly contained therein:

1. The construction of a risk-free portfolio is possible only for Itô–Wiener processes.
2. Because of the nonlinearity of $f(S)$, $\partial f/\partial S$ is time-dependent.
3. The portfolio is risk-free only instantaneously. In order to keep it risk-free over finite times, a continuous adjustment is required.
4. Beware of calculating the option price by a naïve expectation value of the profit, and discounting such as

$$\mathrm{e}^{-r(T-t)} \int_0^\infty p_{\mathrm{hist}}(S_T)(S_T - X)\Theta(S_T - X)$$
$$= \langle \max(S_T - X, 0)\rangle_{\mathrm{hist}} \neq C(S,t) , \tag{4.92}$$

using the historical (recorded) distribution of prices $p_{\mathrm{hist}}(S)$. This will give the wrong result! Such a calculation will give too high a price for the option because p_{hist} is based on a stochastic process with the historic drift μ which ignores the possibility of hedging and overestimates the risk involved in the option position. This will be discussed further in the next section.

We have just discussed the simplest option contract possible, a European call option. The equivalent pricing formulae for a put option can be derived straightforwardly by the reader: they only differ in the boundary condition (4.76) used in the solution of the Black–Scholes differential equation. Many generalizations are possible, such as for options on dividend paying stocks, currencies, interest rates, indices or futures, combi or exotic options, etc. The interested reader is referred to the finance literature [10, 12]–[15] for discussions using similar assumptions as made here (geometric Brownian motion, etc.).

Also path integral methods familiar from physics may be useful [50]. In fact, one can solve the Black–Scholes equation (4.75) by noting the similarity to a time-dependent Schrödinger equation. Time, however, is imaginary, $\tau = it$, identifying the problem as one of quantum statistical mechanics rather than one of zero-temperature quantum mechanics corresponding to real times. The "Black–Scholes Hamiltonian" entering the Schrödinger equation then becomes

$$H_{BS} = -\frac{\sigma^2}{2}\frac{\partial^2}{\partial x^2} + \left(\frac{\sigma^2}{2} - r\right)\frac{\partial}{\partial x} = \frac{p^2}{2m} + \frac{\mathrm{i}}{\hbar}\left(\frac{\sigma^2}{2} - r\right)p \tag{4.93}$$

$$\text{with} \quad x = \ln S\,, \quad p = -\mathrm{i}\hbar\frac{\partial}{\partial x}\,, \quad \text{and} \quad m = \frac{\hbar^2}{\sigma^2}\,.$$

The Black–Scholes equation (4.85) is then obtained by evaluating the path integral using the appropriate boundary conditions (4.76). This method can also be generalized to more complicated problems such as option pricing with a stochastically varying volatility $\sigma(t)$ [51]. That such a method works is hardly surprising from the similarity between the Black–Scholes and Fokker–Planck equations. For the latter, both path-integral solutions, and the reduction to quantum mechanics, are well established [37]. We will use the path integral method in Chap. 7 to price and hedge options in market situations where some of the assumptions underlying the Black–Merton–Scholes analysis are relaxed.

4.5.3 Risk-Neutral Valuation

As mentioned in Sect. 4.5.1, eliminating the stochastic process in the Black–Scholes portfolio as a necessary consequence also eliminates the drift μ of the underlying security. μ, however, is the only variable in the problem which depends on the risk aversion of the investor. The other variables, $S, T-t, \sigma$ are independent of the investor's choice. (Given values for these variables, an operator will only invest his money, e.g., in the stock if the return μ satisfies his requirements.) Consequently, the solution of the Black–Scholes differential equation does not contain any variable depending on the investor's attitude towards risk such as μ, cf. (4.85).

One can therefore assume any risk preference of the agents, i.e., any μ. In particular, the assumption of a risk-neutral (risk-free) world is both possible and practical. In such a world, all assets earn the risk-free interest rate r. The solution of the Black–Scholes found in a risk-neutral world is also valid in a risky environment (our solution of the problem above takes the argument in reverse). The reason is the following: in a risky world, the growth rate of the stock price will be higher than the risk-free rate. On the other hand, the discounting rate applied to all future payoffs of the derivative, to discount them to the present day value, then changes in the same way. Both effects offset each other.

Risk-neutral valuation is equivalent to assuming martingale stochastic processes for the assets involved (up to the risk-free rate r). Equation (4.92) shows that simple expectation value pricing of options, using the historical probability densities for stock prices $p_{\text{hist}}(S)$, does not give the correct option price. In other words, if an option price was calculated according to (4.92), arbitrage opportunities would arise. On the other hand, intuition would suggest that some form of expectation value pricing of a derivative should be possible: the present price of an asset should depend on the expected future cash flow it generates.

Indeed, even in the absence of arbitrage, expectation value pricing is possible, but at a price: a price density $q(S)$ different from the historical density $p_{\text{hist}}(S)$ must be used [52]. This is the consequence of a theorem which states that under certain conditions (which we assume to be fulfilled), for a stochastic process with a probability density $p_{t,T}(S_T)$ for S_T, and conditional densities including the information available up to t, $p_{t,T}(S_T|S_t, S_{t-1}, S_{t-2}, \ldots)$, there is an equivalent martingale stochastic process described by a different probability $q_{t,T}(S_T)$, such that in the absence of arbitrage opportunities, the price of an asset with a payoff function $h(S_T)$ is given by a discounted expectation value using $q_{t,T}$

$$f(t) = \mathrm{e}^{-r(T-t)} \int_{-\infty}^{\infty} \mathrm{d}S_T\, h(S_T) q_{t,T}(S_T)\,. \tag{4.94}$$

As an example, for a call option, the payoff function is $h(S_T) = \max(S_T - X, 0)$ and, with the correct probability density for the equivalent martingale process, involving the risk-free rate r instead of the drift μ of the underlying, the price

$$C(T) = \mathrm{e}^{-r(T-t)} \int_{-\infty}^{\infty} \mathrm{d}S_T\, \max(S_T - X, 0) q_{t,T}(S_T) \tag{4.95}$$

will produce the Black–Scholes solution (4.85). Also, the discounted stock price is an equivalent martingale

$$S_t = \mathrm{e}^{-r(T-t)} \int_{-\infty}^{\infty} \mathrm{d}S_T S_T q_{t,T}(S_T)\,. \tag{4.96}$$

Using equivalent martingales, expectation value pricing for financial assets is possible. Martingales are tied to the notion of risk-neutral valuation.

4.5.4 American Options

The valuation of American options employs the same general risk-neutral framework as for European options. In principle, a riskless hedge of the option position is possible by holding a suitable quantity of the underlying asset. A short position in one American call option still is hedged by a long position in Δ shares of the underlying – the difference to European options is in the numerical value of Δ. The valuation therefore can be based on equivalent martingale processes, with the risk-free rate r as the drift. However, the possibility of early exercise introduces significant complexity and prevents an exact analytic solution.

The basic principle for the valuation of an American option can be illustrated easily. Assume first that time is a discrete variable, $t_i = i\Delta t$, $i = 0, \ldots, N$, $\Delta t = T/N$, where T is the maturity of the option. An American option then can be exercised at any t_i. For geometric Brownian motion, the probability distributions (4.65) and (4.66) are obtained with the trivial replacements $t \to t_i$ and $S_t \to S_i$. The transition probability (conditional probability density) for an elementary time step of the equivalent martingale process in the risk-neutral world, for geometric Brownian motion becomes

$$q_{t_{i-1},t_i}(S_i) \equiv q(S_i, t_i \mid S_{i-1}, t_{i-1}) = \frac{1}{\sqrt{2\pi\sigma^2 \Delta t}} \exp\left(- \frac{\left[\ln\left(\frac{S_i}{S_{i-1}}\right) - \left(r - \frac{\sigma^2}{2}\right)\Delta t\right]^2}{2\sigma^2 \Delta t} \right) . \tag{4.97}$$

One time step before expiry, at t_{N-1}, it is advantageous to exercise the option if its immediate payoff exceeds its value on the assumption of holding it to maturity,

$$h(S_{N-1}) > f(t_{N-1}) , \tag{4.98}$$

where $h(S_i)$ is the payoff function, and $f(t_i)$ is the value of the option, cf. (4.94). To be specific, an American call with payoff $h(S_i) = \max(S_i - X, 0)$ should be exercised at t_{N-1} when

$$S_{N-1} - X > C(t_{N-1}) \tag{4.99}$$

with $C(t_{N-1})$ given by using the discretized version of (4.95). This argument can be iterated backward in time because for an American option, no particular significance is attached to the time of maturity. Consequently, at time t_{i-1}, early exercise is advantageous when the payoff received immediately exceeds the value of the option derived from holding it until the next possibility to exercise, i.e. t_i. The early exercise condition is

$$h(S_{i-1}) > \mathrm{e}^{-r\Delta t} \int_{-\infty}^{\infty} \mathrm{d}S_i \, h(S_i) q_{t_{i-1},t_i}(S_i) . \tag{4.100}$$

The right-hand side has been taken from (4.94) and rewritten for a single time step. For an American call, we get

$$S_{i-1} - X > \mathrm{e}^{-r\Delta t} \int_{-\infty}^{\infty} \mathrm{d}S_i \, \max(S_i - X, 0) q_{t_{i-1},t_i}(S_i) , \tag{4.101}$$

in analogy to (4.95). The option at $t = t_0$ then is priced, and hedged, by iterating the problem backward from maturity, t_T, to $t = t_0$, and taking the continuum limit of time, $\Delta t \to 0$, $N \to \infty$ with $T = N\Delta t$ fixed. Of course, a closed solution of this problem is impossible because for every possible price S_i, a decision on early exercise must be taken at each step i.

A variety of approximate solutions has been developed, all suffering from drawbacks though. Monte Carlo simulations are an obvious choice. Random price increments are drawn from a normal distribution (in the case of geometric Brownian motion) to simulate the price history of the underlying, and the average over many runs is taken when ensemble properties are required.

While Monte Carlo simulations in principle give the desired answer, they are computationally inefficient because the errors on averages over finitely many realizations decrease rather slowly. For plain vanilla options, the use of binomial trees provides an alternative. In a binomial tree, price increments have fixed modulus ΔS, i.e. only $\pm\Delta S$ are allowed. This restriction gives enough simplification to make calculations for plain vanilla options practical. However, for exotic, path-dependent options, the discretization of the price increments is an undesirable feature.

General arguments suggest that American call options should never be exercised early in the absence of dividend payments. Dividend payments have not been considered for European options, and will not be discussed here for American options. The role of dividend payments in option pricing, hedging, and exercise is discussed in the standard financial literature [10].

4.5.5 The Greeks

The derivatives of option prices with respect to the parameters and variables upon which the option price depends, play important roles in trading and hedging strategies. Most of them are labelled by greek letters. Collectively, they are called "the Greeks".

We already encountered one of the Greeks, Delta, and its application in hedging, when setting up the riskless Black–Scholes portfolio in (4.72). There, a short position in a call option was combined with a long position in

$$\Delta_C = \frac{\partial C}{\partial S} \tag{4.102}$$

units of the underlying resulting in a portfolio which was riskless agains infinitesimal variations of the price of the underlying, all other things remaining constant. Similarly, the Delta for a put option is

$$\Delta_P = \frac{\partial P}{\partial S} . \tag{4.103}$$

The definition of Delta, as well as that of the other Greeks is valid for all options. For European options described by the Black–Scholes equations (4.85) and (4.86), we can evaluate Delta explicitly as

$$\Delta_C = N(d_1) , \quad \Delta_P = N(d_1) - 1 , \tag{4.104}$$

where $N(d_1)$ and d_1 are defined in (4.87) and (4.88). Its dependence on the price of the underlying, for different times to maturity, is shown in Fig. 4.8.

Delta describes the dollar variation of an option when the price of the underlying changes by one dollar. More important to investors is the leverage of an option, defined as the percentage variation of the option when the price of the underlying varies by one percent. This quantity is given by

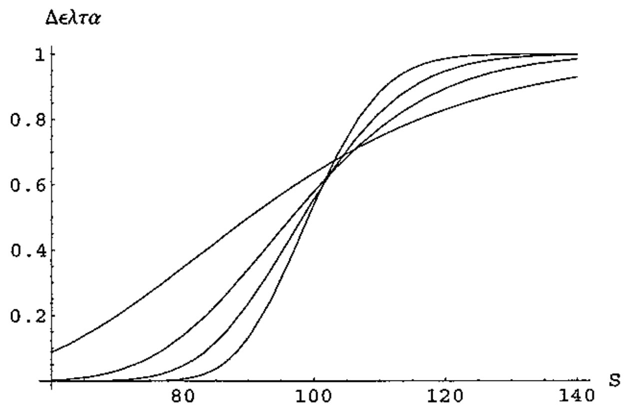

Fig. 4.8. Delta of a European call option described by the Black–Scholes equation as a function of the price of the underlying, for times to maturity of one, two, four and twelve months, from bottom to top at the left margin. The other parameters are $r = 6\%/y$ and $\sigma = 30\%/\sqrt{y}$ as in Fig. 4.7

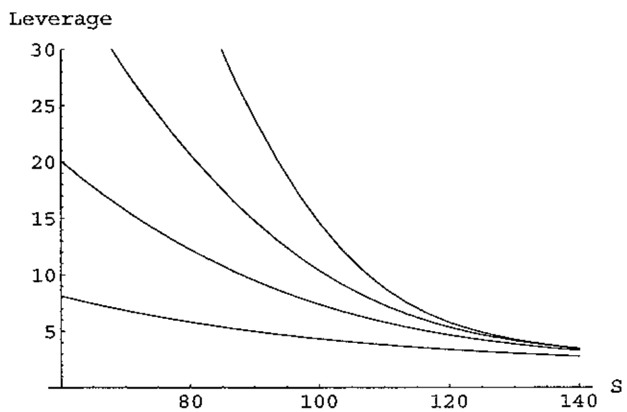

Fig. 4.9. Leverage of a European call option described by the Black–Scholes equation as a function of the price of the underlying, for times to maturity of one, two, four and twelve months, from top to bottom. The other parameters are $r = 6\%/y$ and $\sigma = 30\%/\sqrt{y}$ as in Fig. 4.7

$$\frac{S}{C}\frac{\partial C}{\partial S} \quad \text{and} \quad \frac{S}{P}\frac{\partial P}{\partial S}$$

for call and put options, respectively. The dependence of the leverage on the price of the underlying is displayed in Fig. 4.9 for a European call option. Quite generally, out-of-the money options possess a higher leverage than in-the-money options, and the leverage of a call option decreases when the price of the underlying increases. The downside risk of an option therefore always

is superior to its upside chances. Also, all other things remaining constant, the leverage of an option increases when the time to maturity decreases. As a consequence of these two observations, speculative investments in options are advisable only when the investor holds a strong view on the price movement of the underlying, and on the time scale over which this price movement is realized.

The sensitivity of the option price with respect to time to maturity is expressed by Theta,

$$\Theta_C = \frac{\partial C}{\partial t}\,, \quad \Theta_P = \frac{\partial P}{\partial t}\,. \tag{4.105}$$

For European call and put options described by the Black–Scholes equation, we have

$$\Theta_{C,P} = -\frac{S\sigma}{2\sqrt{2\pi(T-t)}}\mathrm{e}^{-d_1^2/2} \mp rX\mathrm{e}^{-r(T-t)}N(\pm d_2)\,. \tag{4.106}$$

The upper signs apply for a call option, the lower signs for a put. The dependences of Theta on the price of the underlying and on time to maturity is shown in Fig. 4.10. Theta diverges for an at-the-money option when the time to expiration goes to zero. Theta tends towards a finite value when the option is in the money, i.e., in such a case, the loss in value of the call is linear in time shortly before expiration. Theta converges to zero for an out-of-the money call, i.e., such an option has lost all of its value already some time before expiration. Notice that, at least for the European call considered here, the schematic figures in Hull's book [10] seem to indicate an incorrect behavior close to maturity.

Gamma captures the curvature in the derivative prices with respect to the underlying and is defined as

$$\Gamma_C = \frac{\partial^2 C}{\partial S^2}\,, \quad \Gamma_P = \frac{\partial^2 P}{\partial S^2}\,. \tag{4.107}$$

In the Black–Scholes framework,

$$\Gamma_C = \Gamma_P \equiv \Gamma = \frac{1}{S\sigma\sqrt{2\pi(T-t)}}\mathrm{e}^{-d_1^2/2}\,. \tag{4.108}$$

The dependence on the price of the underlying has the same functional form as the probability density function of a lognormal distribution. The dependence on time to maturity is more interesting and shown in Fig. 4.11. When an option expires at the money, Gamma diverges. Gamma tends towards zero, on the other hand, both for options in and out of the money. This behavior is easily understood by considering the payoff profiles of call and put options shown in Fig. 2.1. At expiry, there is a discontinuity in slope in the option payoff at $S = X$. In and out of the money, on the other hand, the payoffs are linear in the price of the underlying.

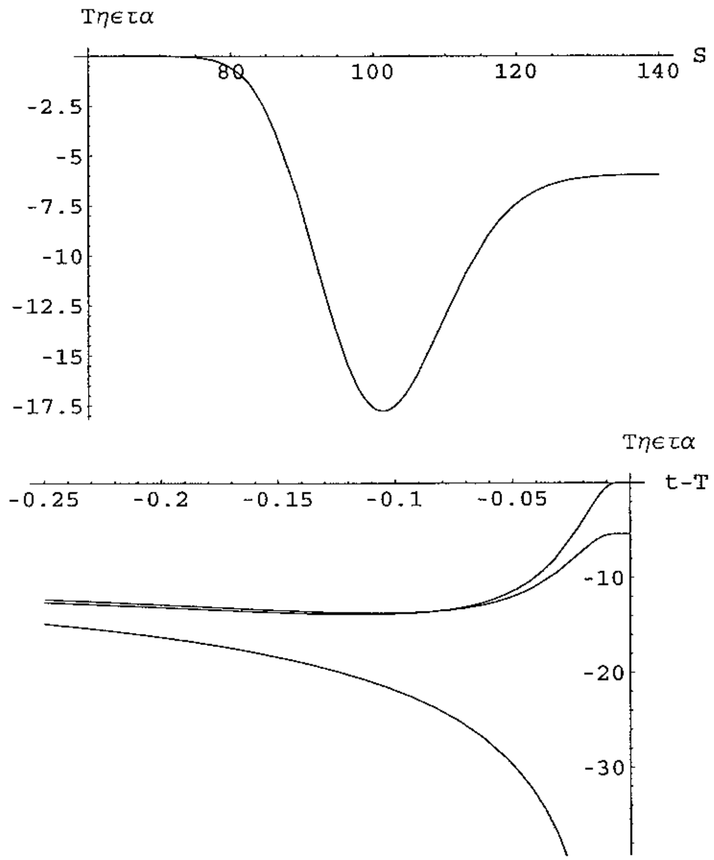

Fig. 4.10. Theta for European call options. The upper panel displays the dependence on the price of the underlying ($X = 100$, $r = 6\%/y$, $\sigma = 30\%/\sqrt{y}$, $T-t = 2m$. The lower panel shows the dependence on time to maturity for $S = 100$ and strike prices $X = 110$ (*top curve*, out of the money), $X = 90$ (*middle curve*, in the money), and $X = 100$ (*bottom curve*, at the money)

The sensitivity of the price of an option with respect to a variation in volatility is important, too. This derivative is called Vega, and is defined as

$$\mathcal{V}_C = \frac{\partial C}{\partial \sigma}\,, \quad \mathcal{V}_P = \frac{\partial P}{\partial \sigma}\,. \tag{4.109}$$

Vega is the same for call and put options. When the Black–Scholes equation applies, we have

$$\mathcal{V} = \frac{S\sqrt{T-t}}{\sqrt{2\pi}}\mathrm{e}^{-d_1^2/2}\,. \tag{4.110}$$

The variation of Vega with the price of the underlying is S^2 times the lognormal probability density function. For an option at the money, the dependence

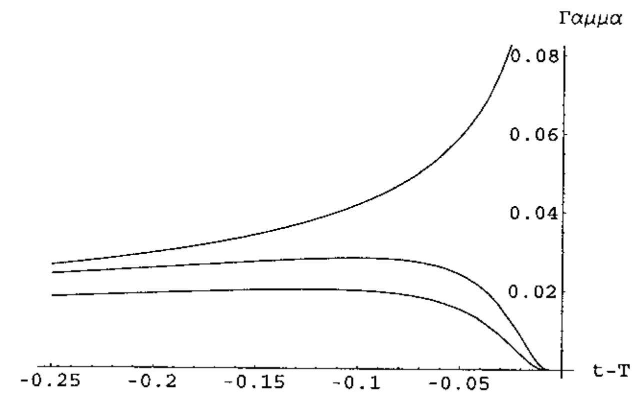

Fig. 4.11. Gamma European call options described by the Black–Scholes equation as a function of time to expiration. The parameters are $S = 100$, $r = 6\%/y$ and $\sigma = 30\%/\sqrt{y}$ and $X = 100$ (*top curve*, at the money), $X = 90$ (*middle curve*, in the money) and $X = 110$ (*bottom curve*, out of the money)

on time to maturity is

$$\mathcal{V} \sim \sqrt{T-t} \quad \text{as} \quad T-t \to 0 \quad (S = X)\,. \tag{4.111}$$

For options in and out of the money,

$$\mathcal{V} \sim \sqrt{T-t}\,\mathrm{e}^{-1/(T-t)} \quad \text{as} \quad T-t \to 0 \quad (S \neq X)\,. \tag{4.112}$$

Except for a different power-law prefactor, this behavior is similar to that shown for Gamma in Fig. 4.11.

Finally, a parameter Rho

$$\mathcal{R}_C = \frac{\partial C}{\partial r}\,, \quad \mathcal{R}_P = \frac{\partial P}{\partial r} \tag{4.113}$$

measures the sensitivity of the prices of call and put options against variations of the risk-free interest rate r. In a Black–Scholes world,

$$\mathcal{R} = \pm X(T-t)\mathrm{e}^{-r(T-t)}N(\pm d_2)\,, \tag{4.114}$$

where the upper and lower signs apply to calls and puts, respectively.

We will come back to Vega later in Sect. 4.5.8 on volatility indices. The use of the Greeks in hedging option positions is discussed in Chap. 10 on risk management.

4.5.6 Synthetic Replication of Options

When the risk-free Black–Scholes portfolio was set up for a short position in a European call option with price C in Sect. 4.5.1, a long position in

$\Delta_C = \partial C/\partial S$ units of the underlying S was added to form a riskless portfolio Π_r

$$\Pi_r = -C + \Delta_C S\,. \tag{4.115}$$

The portfolio consisting of the short option position and the long position in the underlying is exactly equivalent to a long position in a riskless asset of value Π_r. We can transform (4.115) into

$$C = -\Pi_r + \Delta_C S\,. \tag{4.116}$$

A long call position is equivalent to a short position of value Π_r in a riskless asset and a long position in Δ_C units of the underlying of the call, priced at S.

For a short position in a European put option, the risk-free Black–Scholes portfolio is

$$\Pi_r = -P + \Delta_P S = -P - |\Delta_P| S\,. \tag{4.117}$$

The short put position is hedged by a short position in $|\Delta_P|$ units of the underlying, as $\Delta_P < 0$. A long position in a put option then is equivalent to

$$P = -\Pi_r - |\Delta_P| S\,, \tag{4.118}$$

i.e., to a short position of value Π_r in a risk-free asset and another short position in $|\Delta_P|$ units of the underlying.

These equivalences are general and do not assume the validity of the Black–Scholes model. Only the numerical values of Δ_C and Δ_P depend on the price dynamics of the underlying, and on the exercise features of the options. Also, they are not limited to call and put options. The important message is that *any option can be created synthetically by a suitable combination of a position in a riskless asset and another position in the underlying.* This is a result of great practical importance. Whenever an investor wishes to take a position in an option which is not available in the market, he can synthetically replicate the option by taking positions in a risk-free asset and in the underlying. Many portfolio managers and risk managers use this technique to implement their trading and hedging strategies when standard options are not available.

4.5.7 Implied Volatility

Writing the option price in (4.85) symbolically as $C_{\mathrm{BS}}(S,t;r,\sigma;X,T)$, most parameters of the Black–Scholes equation can be observed directly either in the market, or on the option contract under consideration. S and t are independent variables, X and T contract parameters, and r and σ market resp. asset parameters. The volatility σ stands out in that it cannot be observed directly. At best, it can be estimated from historical data on the underlying – a procedure which leaves many questions unanswered.

For a variety of reasons which are the principal motivation of the remainder of this book, the traded prices of options usually differ from their Black–Scholes prices. This is shown in Fig. 4.12 for a series of European calls on the DAX with a lifetime of one month to maturity. The horizontal axis "moneyness", $m = X/S$, represents the dimensionless ratio of strike price over underlying. For comparison, the Black–Scholes solution is also displayed as solid lines. The upper line uses a volatility of $35\%y^{-1/2}$, while the lower one takes $20\%y^{-1/2}$. Under the assumptions of the Black–Scholes theory and geometric Brownian motion, a single value of the volatility should be sufficient to describe the entire series of call options, and the prices should fall on one of the solid lines. Figure 4.12 rejects this hypothesis for real-world option markets.

In the absence of an accurate *ab initio* estimation of the volatility, a rough and pragmatic procedure consists in taking the traded prices for granted and invert the Black–Scholes equation (4.85) for the *implied volatility* σ_{imp} [10]

$$C_{\mathrm{market}}(S,t;r,\sigma;X,T) \equiv C_{\mathrm{BS}}(S,t;r,\sigma_{\mathrm{imp}};X,T)\,. \tag{4.119}$$

The idea is to pack all factors leading to deviations from Black–Scholes theory, independently of their origin, into the single parameter σ_{imp}. Volatility, anyway, is difficult to estimate *a priori*. For the series of options used in Fig. 4.12, the implied volatilites are shown in Fig. 4.13. Apparently, there are deviations of traded option prices from a Black–Scholes equation which depend on the contract to be priced. In this representation, they turn into an implied volatility which explicitly depends on the moneyness of the options. In a purist perspective, implied volatility adds nothing new to the theory of option

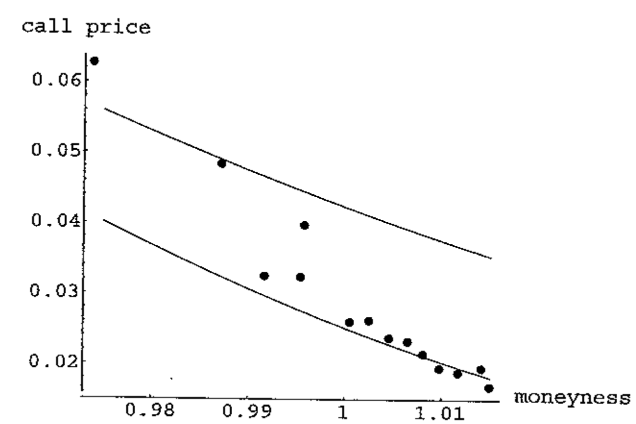

Fig. 4.12. Prices of a series of European call options on the DAX index with one month to maturity, given in units of the index value, against moneyness X/S (*dots*). The two *solid lines* represent the dependence of the Black–Scholes solutions on moneyness with two volatilities $\sigma = 35\%y^{-1/2}$ (*top*) and $\sigma = 20\%y^{-1/2}$ (*bottom*)

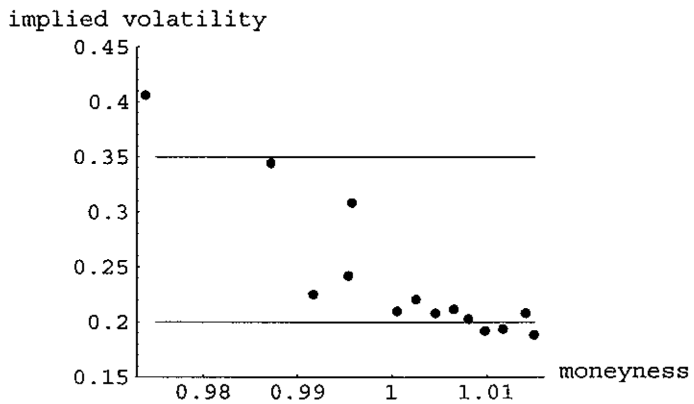

Fig. 4.13. Implied volatilities of a series of European call options on the DAX index with one month to maturity, against moneyness X/S (*dots*) in $\%y^{-1/2}$. Geometric Brownian motion and the Black–Scholes theory take volatility independent of the option contract to be priced. The two *solid lines* mark the contract-independent volatilities used to generate the *solid lines* in Fig. 4.12

pricing, and might even lead to confusion. However, it is a simple transformation of option prices and therefore is an observable on equal footing with the prices. This is similar to physics: When temperature is measured, the basic observable most often is an electric current or voltage drop, or height of a mercury column, etc., which then is transformed into a temperature reading with a suitable calibration. Also, implied volatility is the standard language of derivatives traders and analysts to describe option markets.

The generic shapes of implied volatilities against moneyness are shown in Fig. 4.14. Apparently, a pure smile was characteristic of the US option markets before the 1987 october crash [53]. Ever since, it has become a rather smirky structure. The aim of market models more sophisticated than geometric Brownian motion and of option pricing theories beyond Black–Merton–Scholes, can be restated as to correctly describe implied volatility smiles.

When a series of options with the same strike price but different maturities is analyzed, a term structure (maturity dependence) of the implied volatility is obtained in complete analogy to its moneyness dependence. The volatility smile turns into a two-dimensional implied volatility surface. Figure 4.15 shows a series of cuts through an implied volatility surface of European call options. Unlike Fig. 4.13, these curves do not represent market observations but are the results of a model calculation. Superficially, the one-month curve is not dissimilar to the empirical data, suggesting that theoretical models indeed may be capable of correctly describing option markets. Attempts to fit volatility smiles for a fixed time to maturity usally employ quadratic functions with different parameters for in- and out-of-the-money options, to account for the systematic asymmetry [53].

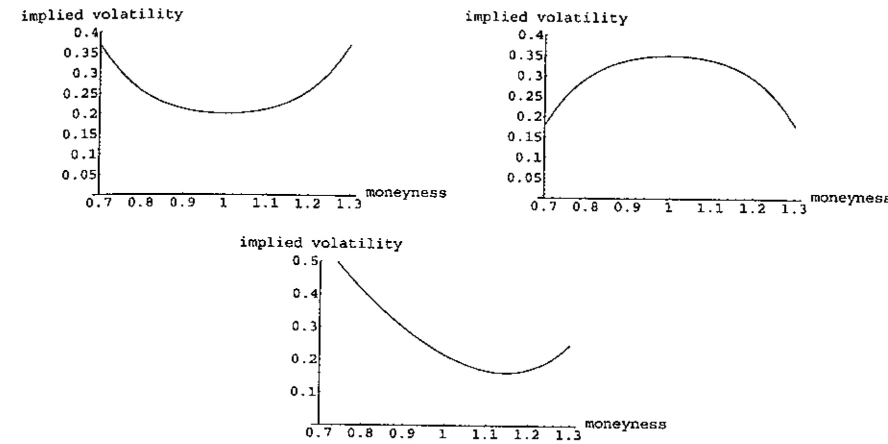

Fig. 4.14. Sketches of implied volatilities against moneyness. Three generic shapes can be observed: a smile (*top left*), a frown (*top right*) and a smirk resp. skewed smile (*bottom*). In equity markets, the smirk is observed most frequently. Often, the term "volatility smile" includes all three shapes

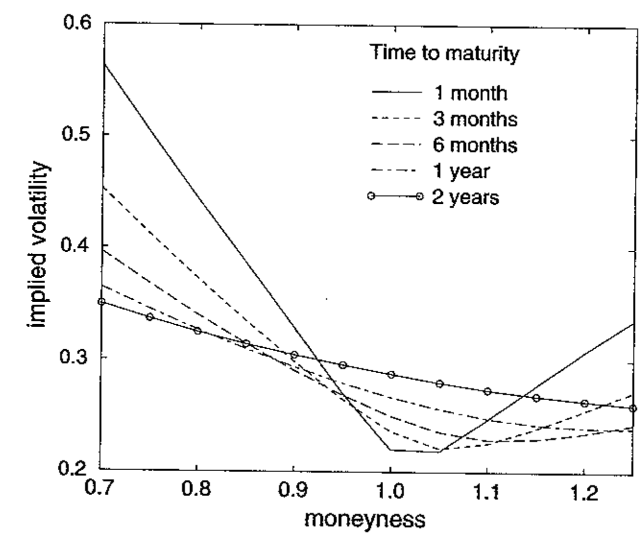

Fig. 4.15. Term structure and moneyness dependence of the implied volatility of a series of European call options based on a model calculation

With reference to the subsequent chapters where we will develop an in-depth description of financial markets, the main handles on the volatility smiles and term structures are:

- The ordinate scale is determined by the average volatility of the market/model.

- Smiles or frowns are the consequence deviations of the actual return distributions, especially in their wings, from the Gaussian assumed in geometric Brownian motion.
- The skew in the implied volatility is the consequence either of a skewness (asymmetry) of the return distribution of the underlying or of return–volatility correlations.
- The term structure of the volatility smiles is determined by the time scales (or time-scale-free behavior) of the important variables in the problem.

Figures 4.12 and 4.13 show the option prices and implied volatilites of DAX options on one particular trading day. Both quantities show an interesting dynamics when studied with time resolution. The price of a specific option, of course, possesses a dynamics because of the variation in the price of the underlying. When the prices of a series of options are represented in terms of moneyness, however, these variations are along the price curve $C(X/S)$ once the effects of changing time to maturity are eliminated, and should not lead to dynamical variations of the price curve itself. Additional dynamics may come, e.g., from the increasing autonomy of option markets which are increasingly driven by demand and supply, in addition to the price movements of the underlying [54]. One can analyze this dynamics of $\sigma_{\mathrm{imp}}(m)$ almost at the money, $m \approx 1$. When, e.g., the time series of $\sigma_{\mathrm{imp}}(1-\delta) - \sigma_{\mathrm{imp}}(1)$ and $\sigma_{\mathrm{imp}}(1) - \sigma_{\mathrm{imp}}(1+\delta)$ are plotted against time, there are long periods where both stochastic time series are strongly correlated, and other shorter periods where their correlation is weak [53]. The former correspond to almost rigid shifts of the smile patterns while the latter appear in periods where the smile predominantly changes shape. Both time series can be modeled as AR(1) processes which describes an implied volatility with a mean-reversion time of about 30 days, comparable to the time to maturity of liquid options.

This line of research can be carried much further by studying the dynamical properties of a two-dimensional implied volatility surface with coordinates moneyness (m) and time to maturity ($T-t$) [54]. Implied volatilities are strongly correlated across moneyness and time to maturity, cf. above, which suggests a description in terms of surface dynamics. A practical aspect are trading rules for volatility prediction based on implied volatility. The "sticky moneyness" rule predicts that the implied volatility surface tomorrow is the same as that today at constant moneyness and time to maturity. The "sticky strike" rule stipulates that the implied volatility tomorrow is the same as today at constant strike and constant maturity (i.e. absolute quantities).

Volatility surfaces can be generated for various series of liquid options such as calls and puts on the S&P500, the FTSE, or the DAX. With a generalization of principal component analysis – a technique widely used in image processing – the implied volatility surfaces can be described as fluctuating random surfaces driven by a small number of dominant eigenmodes. These eigenmodes parameterize the shape fluctuations of the surface. Their fluctuating prefactors describe the amplitude of surface variations. The first

eigenmode which accounts for about 80% of the daily variance of the implied volatility surface is a flat sheet in $\sigma_{\mathrm{imp}} - m - (T-t)$ space, almost independent of $T-t$ and with a small positive slope in m. This mode essentially has the same properties as the time series discussed in the second preceding paragraph. It is also negatively correlated with the price of the underlyings, i.e. contributes to a "leverage effect" to be discussed in Sect. 5.6.3. The second eigenmode changes sign at the money and is positive for $m > 1$ and negative for $m < 1$. A positive variation of this mode increases the volatilities of out-of-the money calls and descreases it for out-of-the-money puts. It contributes to the skewness of the risk-neutral distributions (when thinking backwards from implied volatility to risk-neutral measures) and, due to its slope in $T-t$, to the term structure. It also possesses the dynamics of a mean-reverting AR(1)-process. The third mode is a butterfly modes which changes the convexity of the implied volatility surface. It leads to a fattening of the tails of the risk-neutral distributions, cf. the mechanistic rules listed above [54].

This dynamics can be cast in a low-dimensional factor model

$$X(t;m,T-t) \equiv \ln \sigma_{\mathrm{imp}}(t;m,T-t) = X(0;m,T-t) + \sum_{k=1}^{d} x_k(t) f_k(m,T-t) \,. \tag{4.120}$$

f_k is one of the d dominant eigenfunctions of the principal component decomposition. They are time-independent and describe the spatial variation of the fluctuations. The dynamics comes from the randomly fluctuating prefactors $x_k(t)$ which, according to the findings above, can be modeled as Ornstein-Uhlenbeck processes

$$\mathrm{d}x_k(t) = -\lambda_k \left[x_k(t) - \bar{x}_k\right] \mathrm{d}t + \gamma_k \mathrm{d}z_k \,. \tag{4.121}$$

λ_k is the rate of mean reversion and $\bar{x}_k$ is the average of the k^{th} eigenmode. The stochastic increments $\mathrm{d}z_k$ are uncorrelated and may be drawn from a Gaussian (consistent with the lognormal distribution of implied volatilites, cf. below) or a more general distribution. The ranks of the fluctuating expansion coefficients $x_k(t)$ in (4.120) are ordered according to their variances γ_k^2 which measures the amplitude of the fluctuations they impart on $\sigma_{\mathrm{imp}}(t;m,T-t)$. The dynamics of the implied volatility surfaces analyzed above can be faithfully represented by three factors $x_1(t) \ldots x_3(t)$ [54].

4.5.8 Volatility Indices

Volatility is the most important and least accessible quantity in option theory. Volatility can be inferred either from historical time series [estimate σ in (4.53)] or from implied volatility of options by inverting the Black–Scholes equation as in (4.119). For derivative markets, the second method is preferrable because the information is derived directly from derivative

5. Scaling in Financial Data and in Physics

The Black–Scholes equation for option prices is based on a number of hypotheses and assumptions. Subsequent price changes were assumed to be statistically independent, and their probability distribution was assumed to be the normal distribution. Moreover, the risk-free interest rate r and the volatility σ were assumed constant (in the simplest version of the theory). In this chapter, we will examine financial data in the light of these assumptions, develop more general stochastic processes, and emphasize the parallels between financial data and physics beyond the realm of Brownian motion.

5.1 Important Questions

We will be interested, among others, in answering the following important questions:

- How well does geometric Brownian motion describe financial data? Can the apparent similarities between financial time series and random walks emphasized in Sect. 3.4.1 be supported quantitatively?
- What are the empirical statistics of price changes?
- Are there stochastic processes which do not lead to Gaussian or log-normal probability distributions under aggregation?
- Is there universality in financial time series, i.e., do prices of different assets have the same statistical properties?
- Are financial markets stationary?
- Are real markets complete and efficient, as assumed by Bachelier?
- Why is the Gaussian distribution so frequent in physics?
- What are Lévy flights? Are they observable in nature?
- Are there correlations in financial data?
- How can we quantify temporal correlations in a financial time series?
- How can we quantify cross-correlations between various asset price histories?

Before discussing in detail the stochastic processes underlying real financial time series, we address the stationarity of financial markets.

5.2 Stationarity of Financial Markets

Geometric Brownian motion underlying the Black–Scholes theory of option pricing works with constant parameters: the drift μ and volatility σ of the return process, and the risk-free interest rate r are assumed independent of time. Is this justified? And is the dynamics of a market the same irrespective of time? That is, are the rules of the stochastic process underlying the return process time-independent?

For a practical option- pricing problem with a rather short maturity, say a few months, the estimation of the Black–Scholes parameters should pose no problem. For an answer to the questions posed above, on longer time scales, we will investigate various time series of returns. The following quantities will be of interest:

- The time series of (logarithmic) returns of an asset priced at $S(t)$ over a time scale τ

$$\delta S_\tau(t) = \ln\left(\frac{S(t)}{S(t-\tau)}\right) \approx \frac{S(t) - S(t-\tau)}{S(t-\tau)} \,. \tag{5.1}$$

- The time series of returns normalized to zero mean and unit variance

$$\delta s_\tau(t) = \frac{\delta S_\tau(t) - \langle \delta S_\tau(t)\rangle}{\sqrt{\langle [\delta S_\tau(t)]^2\rangle - \langle \delta S_\tau(t)\rangle^2}} \,, \tag{5.2}$$

  where the expectation values are taken over the entire time series under consideration.

We first examine the time series of DAX daily closes from 1975 to 2005 shown in Fig. 1.2. The daily returns $\delta S_{1\mathrm{d}}(t)$ derived from the data up to 5/2000 are shown in Fig. 5.1. At first sight, the return process looks stochastic with zero mean. The impressive long-term growth of the DAX up to 2000 and sharp decline thereafter, emphasized in Fig. 1.2, here show up in a small, almost invisible positive resp. negative mean of the return, of much smaller amplitude, however, than the typical daily returns. We also clearly distinguish periods with moderate (positive and negative) returns, i.e., low volatility (more frequent in the first half of the time series) from periods with high (positive and negative) returns, i.e., high volatility (more frequent in the second half of the time series). The main question is if data like Fig. 5.1 are consistent with a description, and to what accuracy, in terms of a simple stochastic process with constant drift and constant volatility. Or, to the contrary, do we have to take these parameters as time dependent, such as in the ARCH(p) or GARCH(p,q) models of Sect. 4.4.1? Or, worse even, do the constitutive functional relations of the stochastic process change with time?

As a first, admittedly superficial test of stationarity, we now divide the DAX time series into seven periods of approximately equal length, and evaluate the average return and volatility in each period. The result of this evaluation is shown in Table 5.1. The central column shows the increase resp.

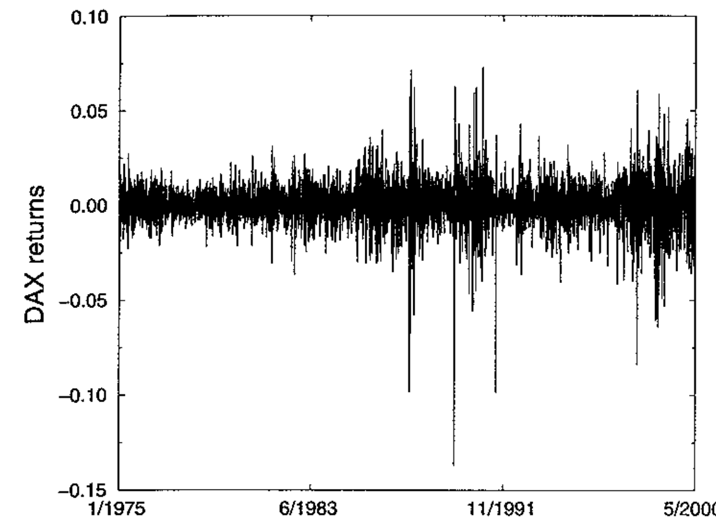

Fig. 5.1. Time series of daily returns of the DAX German blue chip index from 1975 to 2000. Analysis courtesy of Stephan Dresel based on data provided by Deutsche Bank Research

Table 5.1. Average return $\langle \delta S_{1\mathrm{d}}(t)\rangle$ and volatility σ of the DAX index in seven approximately equally long periods from January 2, 1975, to December 31, 2004. Analysis courtesy of Stephan Dresel based on data provided by Deutsche Bank Research supplemented by data downloaded from Yahoo, `http://de.finance.yahoo.com`

Period	Return $[\mathrm{d}^{-1}]$	Volatility $[\mathrm{d}^{-1/2}]$
02.01.1975–15.03.1979	0.00028	0.0071
16.03.1979–10.06.1983	0.00021	0.0078
13.06.1983–03.09.1987	0.00072	0.0104
04.09.1987–02.12.1991	0.00002	0.0155
03.12.1991–14.02.1996	0.00042	0.0091
16.02.1996–05.05.2000	0.00106	0.0149
08.05.2000–31.12.2004	−0.00049	0.0184

decrease of the average returns with time, which is responsible for the increasing slope of the DAX index in Fig. 1.2. The average return increases by a factor of three to four from 1975 to 2000, and decreases to even become negative in the drawdown period from 2000 to 2005. The rather low value in the fourth period is due to the October crash in 1987 right after the beginning of our period, and another crash in 1991. The last column shows the volatilities which also increase with time. The volatility is particularly big after 2000.

In the six periods up to May 5, 2000, we now subtract the average return from the daily returns and then divide by the standard deviation, in order to obtain a process with mean zero and standard deviation unity. Figure 5.2 shows the probability distributions of the returns normalized in this way, in the six periods. Except for a few points in the wings, the six distributions do not deviate strongly from each other. One therefore would conclude that the rules of the stochastic process underlying financial time series do not change with time significantly, and that most of the long-term evolution of markets can be summarized in the time dependence of its parameters.

Notice, however, that, strictly speaking, this finding invalidates geometric Brownian motion as a model for financial time series because μ and σ were assumed constant there. On the other hand, if such time dependences of parameters only are important on sufficiently long time scales (which we have not checked for the DAX data), one might take a more generous attitude, and consider geometric Brownian motion as a candidate for the description of the DAX on time scales which are short compared to the time scale of variations of the average returns or volatilities. Physicists take a similar attitude,

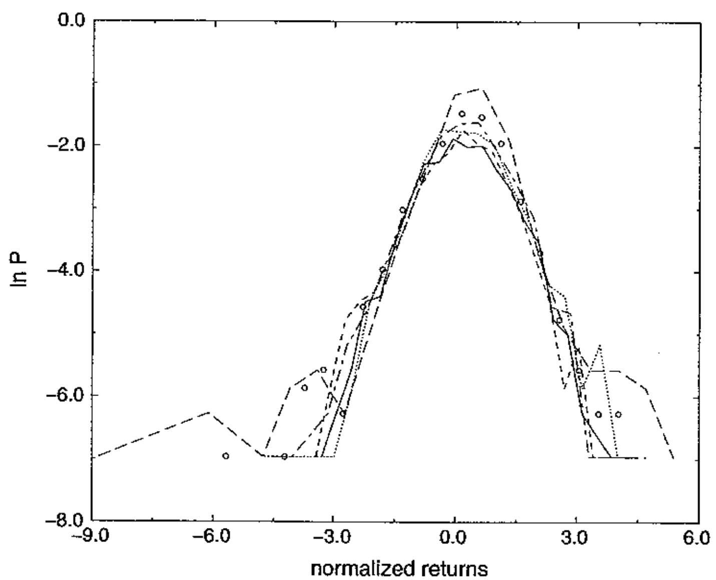

Fig. 5.2. Probability distributions of normalized daily returns of the DAX German blue chip index in the six equally long periods from 1975 to 2000. The normalization procedure is explained in the text and the parameters are summarized in Table 5.1. *Solid line*: period 1, *dotted line*: period 2, *dashed line*: period 3, *long-dashed line*: period 4, *dot-dashed line*: period 5, *circles*: period 6. Analysis courtesy of Stephan Dresel based on data provided by Deutsche Bank Research

e.g., with temperature, in systems slightly perturbed away from equilibrium. While being an equilibrium property in the strict sense, one may introduce local temperatures in an inhomogeneous system on scales that are small with respect to those over which the temperature gradients vary appreciably.

Returning to the probability distributions of the DAX returns, Fig. 5.3 shows the probability distributions of three periods (1, 4, and 6) displaced for clarity. Period 1 is not clearly Gaussian although its tails are not very fat, a fact that we qualitatively reproduce in periods 2 and 3. The distributions of periods 4, 5 (not shown), and 6 do possess rather fat tails whose importance, however, changes with time. In the DAX sample, period 4 including the October crash in 1987, and some more turmoil in 1990 and 1991, clearly has the fattest tails. One therefore should be extremely careful in analyzing market data from very long periods. Markets certainly change with time, and there may be more time dependence in financial time series than just a slow variation of average returns and volatilities. As Fig. 5.3 suggests, even the shape of the probability distribution might change with time.

These complications have not been studied systematically, and are ignored in the following discussion. Depending on the underlying time scales, they may or may not affect the conclusions of the various studies we review. We first proceed to a critical examination of geometric Brownian motion.

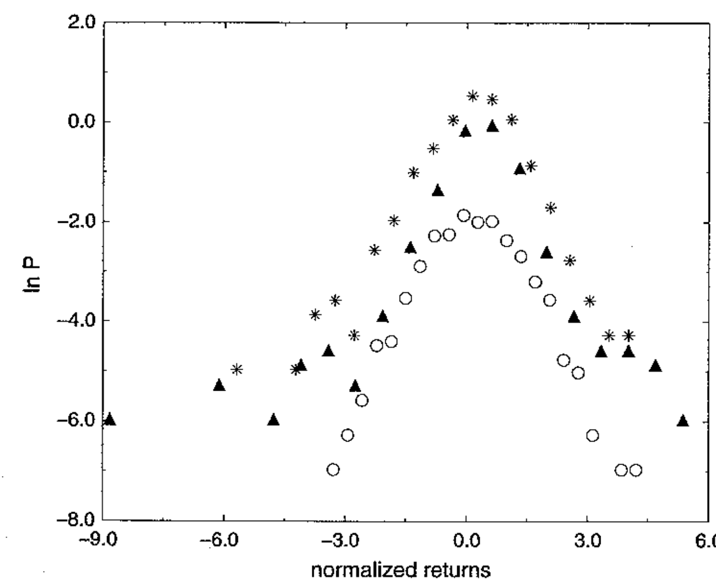

Fig. 5.3. Probability distributions (vertically displaced for clarity) of normalized daily returns of the DAX German blue chip index in the periods 1 (*open circles*), 4 (*filled triangles*), and 6 (*stars*) specified in Table 5.1. Analysis courtesy of Stephan Dresel based on data provided by Deutsche Bank Research

5.3 Geometric Brownian Motion

Geometric Brownian motion makes two fundamental hypotheses on a stochastic process:

1. Successive realizations of the stochastic variable are statistically independent.
2. Returns of financial markets, or relative changes of the stochastic variable, are drawn from a normally distributed probability density function, i.e., the probability density function of the stochastic variable, resp. prices, is log-normal.

Here, we examine these properties for financial time series.

5.3.1 Price Histories

Figure 5.4 shows three financial time series which we shall use to discuss correlations: the S&P500 index (top), the DEM/US$ exchange rate (center), and the BUND future (bottom) [17]. The BUND future is a futures contract on long-term German government bonds, and thereby a measure of long-term interest-rate expectations. The data range from November 1991 to February 1995.

Figure 5.5 gives a chart of high-frequency data of the DAX taken on a 15-second time interval. The history is a combination of data collected in a purpose-built database of German stock index data [59, 60] at the Department of Physics, Bayreuth University, and data provided by an economics database at Karlsruhe University [61].

5.3.2 Statistical Independence of Price Fluctuations

A superficial indication of statistical independence of subsequent price fluctuations was given by the comparison of our numerical simulations based on an IID random variable to the DAX time series, shown in Figs. 1.3 and 3.7. The overall similarity between the simulation of a random walk and the daily closing prices of the DAX would support such a hypothesis. Notice, however, that the DAX is an index composed of 30 stocks, and correlations in the time series of individual stocks may be lost due to averaging. Also, correlations may well persist on time scales smaller than one day.

The question of correlations has a different emphasis for the statistician or econometrician, and for a practitioner. Academics ask for any kind of dependence in time series. Practitioners will more frequently inquire if possible dependences can be used for generating above-average profits, and if successful trading rules can be built on such correlations. Despite what has been said in the preceding paragraph, the apparent importance of technical analysis suggests that there may indeed be tradable though subtle correlations.

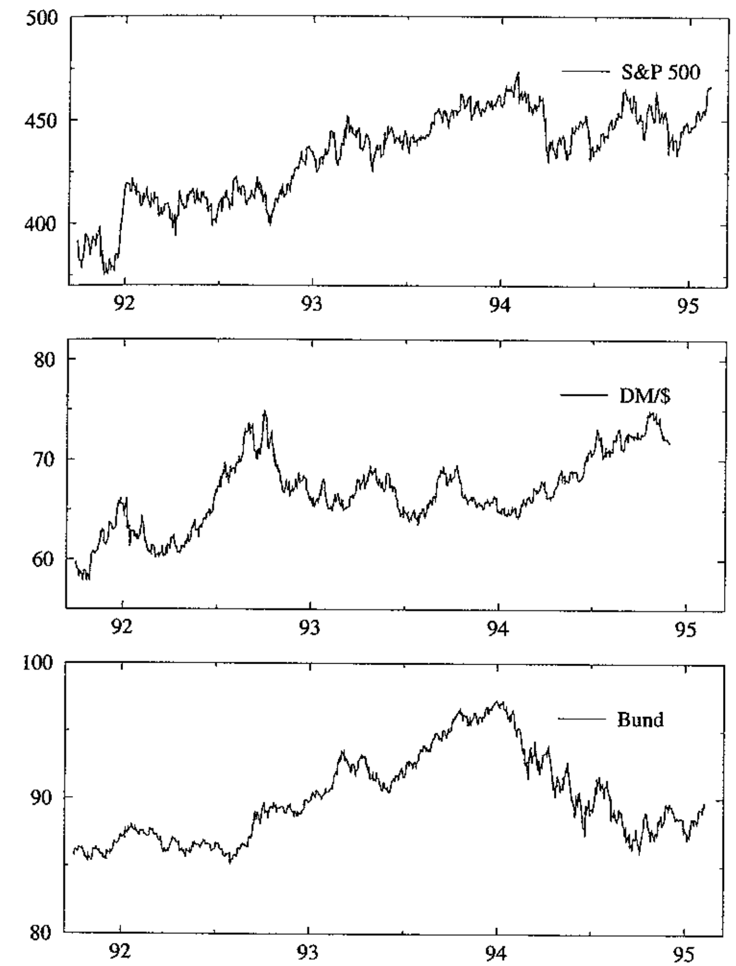

Fig. 5.4. Three financial time series from November 1991 to February 1995: the S&P500 index (*top*), the DEM/US$ exchange rate (*center*), and the BUND futures (*bottom*). From J.-P. Bouchaud and M. Potters: *Théorie des Risques Financiers*, by courtesy of J.-P. Bouchaud. ©1997 Diffusion Eyrolles (Aléa-Saclay)

Correlation Functions

We now analyze correlation functions of returns on a fixed time scale τ, $\delta S_\tau(t)$, (5.1). The autocorrelation function of this quantity is

$$C_\tau(t-t') = \frac{1}{D\tau}\langle [\delta S_\tau(t) - \langle \delta S_\tau(t)\rangle] [\delta S_\tau(t') - \langle \delta S_\tau(t')\rangle]\rangle \; , \tag{5.3}$$

where

$$D\tau = \mathrm{var}\,[\delta S_\tau(t)] \; , \tag{5.4}$$

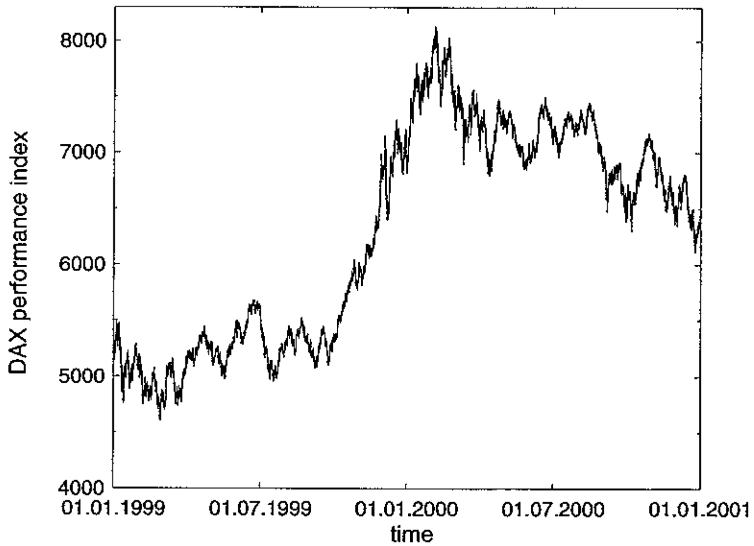

Fig. 5.5. Chart of the DAX German blue chip index during 1999 and 2000. Data are taken on a 15-second time scale. From S. Dresel: *Modellierung von Aktienmärkten durch stochastische Prozesse*, Diplomarbeit, Universität Bayreuth, 2001, by courtesy of S. Dresel

to emphasize the similarity to diffusion. Using (5.2), we also have

$$C_\tau(t-t') = \langle \delta s_\tau(t) \delta s_\tau(t') \rangle \,. \tag{5.5}$$

For statistically independent data, we have $C_\tau(t-t') = 0$ for $t \neq t'$ (at least in the limit of very large samples).

Figure 5.6 shows the autocorrelation functions of the three assets represented in Fig. 5.4 with price changes evaluated on a τ = 5-minute scale [17]. For time lags below 30 minutes, there are weak correlations above the 3σ level. Above 30-minute time lags, correlations are not significant.

When errors are random and normally distributed (a standard assumption), the standard deviation determines the confidence levels as

$$\left. \begin{array}{r} \sigma : 32\% \\ 2\sigma : 5\% \\ 3\sigma : 0.2\% \\ 10\sigma : 2 \times 10^{-23} \end{array} \right\} = 2 \int_A^\infty P(S) \mathrm{d}S \,. \tag{5.6}$$

Under a null hypothesis of vanishing correlations, 32% of the data may randomly lie outside a 1σ corridor, or 0.2% of the data may be outside a 3σ corridor.

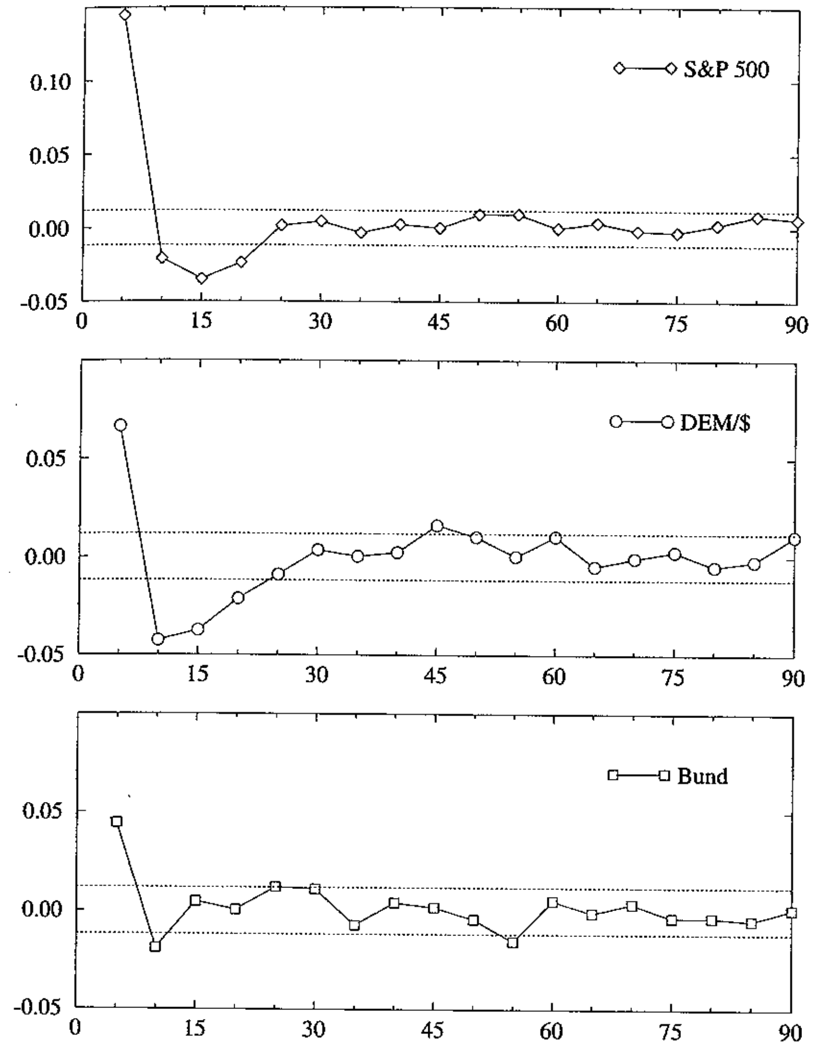

Fig. 5.6. Autocorrelation functions of the S&P500 index (*top*), the DEM/US$ exchange rate (*center*), and the BUND future (*bottom*), over a time scale $\tau = 5$ minutes. The horizontal scale is the time separation $t - t'$ in minutes. The horizontal dotted lines are the 3σ confidence levels. From J.-P. Bouchaud and M. Potters: *Théorie des Risques Financiers*, by courtesy of J.-P. Bouchaud. ©1997 Diffusion Eyrolles (Aléa-Saclay)

In Fig. 5.6 for time lags above 30 minutes, the (null) hypothesis of statistically independent price changes therefore cannot be rejected for the three assets studied. The non-random deviations out of the 3σ corridor for smaller time lags, on the other hand, indicate non-vanishing correlations in this range. Consistent with this is the finding that no correlations significant on the 3σ

level can be found for the same assets when the time scale for price changes is increased to $\tau = 1$ day [17].

More precise autocorrelation functions can be obtained from the DAX high-frequency data [59, 60]. Figure 5.7 shows the autocorrelation function $C_{15''}(t-t')$ of this sample together with 3σ error bars. Correlations are positive with a short 53-second correlation time and negative (overshooting) with a longer 9.4-minute correlation time. The remarkable feature of Fig. 5.7 is, however, the small weight of these correlations! The solid line represents a fit of the data to a function

$$C_{15''}^{\text{fit}}(t-t') = 0.89\delta_{t,t'} + 0.12\mathrm{e}^{-|t-t'|/53''} - 0.01\mathrm{e}^{-|t-t'|/9.4'}, \qquad (5.7)$$

implying that the data are uncorrelated to almost 90%, even at a 15-second time scale. Bachelier's postulate is satisfied remarkably well.

The delta-function contribution at zero time lag is also present, although with a smaller prefactor, in a study based on 1-minute returns in the S&P500 index [62], although only positive correlations with a correlation time of 4 minutes and no overshooting to negative correlations at longer times are found there. A strong zero-time-lag peak and overshooting to negative

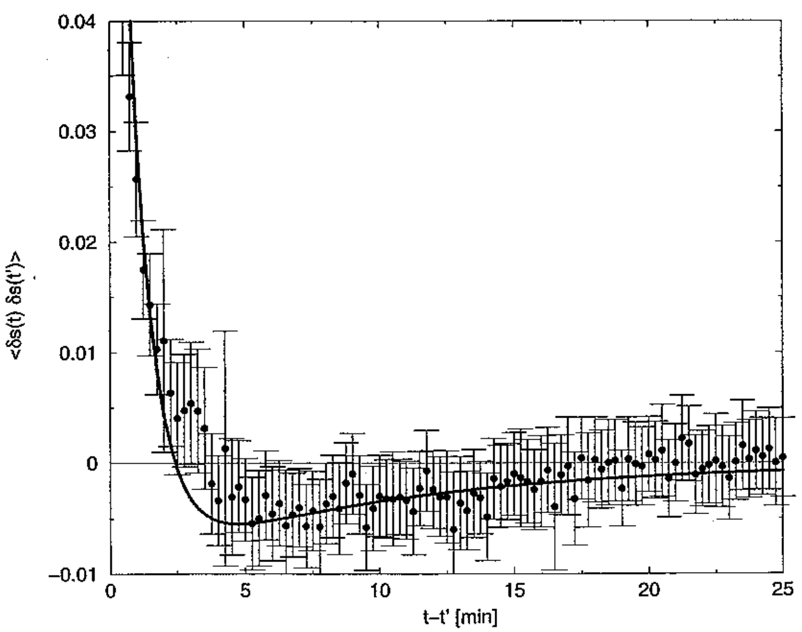

Fig. 5.7. Linear autocorrelation function $C_{15''}(t-t')$ for 15-second DAX returns (*dots*) with 3σ error bars. The solid line is a fit to (5.7) and demonstrates that the data are almost uncorrelated. From S. Dresel: *Modellierung von Aktienmärkten durch stochastische Prozesse,* Diplomarbeit, Universität Bayreuth, 2001, by courtesy of S. Dresel

correlations at about 15 minutes are also visible in 1-minute data from the Hong Kong Hang Seng stock index [63].

That subsequent price changes are essentially statistically independent is not a new finding. It was established, based on time-series analysis, back in 1965 by Fama [64] (and before, perhaps, by others). In the next section, we shall discuss another interesting aspect of Fama's work.

Filters

Fama's work was motivated by Mandelbrot's objections (to be discussed below in Sect. 5.3.3) to the standard geometrical Brownian motion model of price changes of financial assets, Sect. 4.4.2. In the course of his criticism, Mandelbrot also pointed to the "fallacies of filter trading" [33]. Filters were invented by Alexander [65] and were trading rules purported to generate above-average profits in stock market trading.

An x%-filter works like this: if the relative daily price change of an asset $\Delta S/S > x\%$ after a local minimum, then buy the stock and hold until $\Delta S/S < -x\%$ after a local maximum. At this point, sell the stock and simultaneously go short until $\Delta S/S > x\%$ after another local minimum. Close out the short position and go long at the same time, etc. If filters are successful, more successful than, e.g., a naïve buy-and-hold strategy, there must be non-trivial correlations in the stock market.

Fama conducted a systematic investigation of such filters on all Dow Jones stocks from late 1957 to September 1962 [64]. Important results of his study are summarized in Table 5.2. The comparison with simple buy-and-hold is rather negative. Even ignoring transaction costs, only 7 out of the 30 Dow Jones stocks generated higher profits by filter trading than by buy-and-hold. Filter trading, however, involves frequent transactions, and when transaction costs are included, buy-and-hold was the better strategy for all 30 stocks, leading Fama to the conclusion: "From the trader's point of view, the independence assumption of the random-walk model is an adequate description of reality" [64].

Notice that Fama's investigation addresses correlations in the time series of individual stocks, as well as the practical aspects. We now turn to the statistics of price changes.

5.3.3 Statistics of Price Changes of Financial Assets

Early "tests" of the statistics of price changes did not reveal obvious contradictions to a (geometrical) Brownian motion model. Bachelier himself had conducted empirical tests of certain of his calculations, and of the underlying theory of Brownian motion [6]. Within the uncertainties due to the finite (small) sample size, there seemed to be at least consistency between the data and his theory. The problem we remarked on in Sect. 3.2.3, that price changes