

Chapter 1

10 Computational Challenges in Finance

Sascha Desmettre and Ralf Korn

Abstract With the growing use of both highly developed mathematical models and complicated derivative products at financial markets, the demand for high computational power and its efficient use via fast algorithms and sophisticated hard- and software concepts became a hot topic in mathematics and computer science. The combination of the necessity to use numerical methods such as Monte Carlo simulation, of the demand for a high accuracy of the resulting prices and risk measures, of online availability of prices, and the need for repeatedly performing those calculations for different input parameters as a kind of sensitivity analysis emphasizes this even more.

In this survey, we describe the mathematical background of some of the most challenging computational tasks in financial mathematics. Among the examples are the pricing of exotic options by Monte Carlo methods, the calibration problem to obtain the input parameters for financial market models, and various risk management and measurement tasks.

1.1 Financial Markets and Models as Sources for Computationally Challenging Problems

With the growing use of both highly developed mathematical models and complicated derivative products at financial markets, the demand for high computational power and its efficient use via fast algorithms and sophisticated hard- and software concepts became a hot topic in mathematics and computer science. The combination of the necessity to use numerical methods such as Monte Carlo (MC) simulations, of the demand for a high accuracy of the resulting prices and risk measures, of online availability of prices, and the need for repeatedly performing those calculations for different input parameters as a kind of sensitivity analysis emphasizes this even more.

S. Desmettre (✉) • R. Korn
Department of Mathematics, TU Kaiserslautern, 67663 Kaiserslautern, Germany
e-mail: desmettre@mathematik.uni-kl.de; korn@mathematik.uni-kl.de

In this survey, we describe the mathematical background of some of the most challenging computational tasks in financial mathematics. Among the examples are the pricing of exotic options by MC methods, the calibration problem to obtain the input parameters for financial market models, and various risk management and measurement tasks.

We will start by introducing the basic building blocks of stochastic processes such as the Brownian motion and stochastic differential equations, present some popular stock price models, and give a short survey on options and their pricing. This will then be followed by a survey on option pricing via the MC method and a detailed description of different aspects of risk management.

1.2 Modeling Stock Prices and Further Stochastic Processes in Finance

Stock price movements in time as reported in price charts always show a very irregular, non-smooth behavior. The irregular fluctuation seems to dominate a clear tendency of the evolution of the stock price over time. The appropriate mathematical setting is that of diffusion processes, especially that of the Brownian Motion (BM).

A one-dimensional BM $W(t)$ is defined as a stochastic process with continuous path (i.e. it admits continuous realizations as a function of time) and

- $W(0) = 0$ almost surely,
- Stationary increments with $W(t) - W(s) \sim \mathcal{N}(0, t - s)$, $t > s \geq 0$,
- Independent increments, i.e. $W(t) - W(s)$ is independent of $W(u) - W(r)$ for $t > s \geq u > r \geq 0$.

A d -dimensional BM consists of a vector $W(t) = (W_1(t), \dots, W_d(t))$ of independent one-dimensional BMs $W_i(t)$. A correlated d -dimensional BM is again a vector of one-dimensional BMs $Z_i(t)$, but with

$$\text{Corr}(Z_i(t), Z_j(t)) = \rho_{ij}$$

for a given correlation matrix ρ .

A simulated path of a one-dimensional BM, i.e. a realization of the BM $W(t)$, $t \in [0, 1]$ is given in Fig. 1.1. It exhibits the main characteristics of the BM, in particular its non-differentiability as a function of time.

In this survey, we will consider a general diffusion type model for the evolution of stock prices, interest rates or additional processes that influence those prices. The corresponding modeling tool that we are using are Stochastic Differential Equations




Fig. 1.1 A path of a Brownian motion

(SDEs) (see [11] for a standard reference on SDEs). In particular, we assume that the price process $S(t)$ of d stocks and an additional, m -dimensional state process $Y(t)$ are given by the SDE

$$\begin{aligned} dS(t) &= \mu(t, S(t), Y(t)) dt + \sigma(t, S(t), Y(t)) dW(t), \quad S(0) = s, \\ dY(t) &= \kappa(t, Y(t)) dt + \nu(t, Y(t)) dW(t), \quad Y(0) = y. \end{aligned}$$

Here, we assume that the coefficient functions μ, σ, κ, ν satisfy appropriate conditions for existence and uniqueness of a solution of the SDE. Such conditions can be found in [11]. Sufficient (but not necessary) conditions are e.g. the affine linearity of the coefficient functions or suitable Lipschitz and growth conditions. Further, $W(t)$ is a k -dimensional BM.

The most popular special case of those models is the Black-Scholes (BS) model where the stock price does not depend on the state process $Y(t)$ (or where formally the state process Y is a constant). We assume that we have $d = k = 1$ and that the stock price satisfies the SDE

$$dS(t) = S(t)(\mu dt + \sigma dW(t)), \quad S(0) = s \quad (1.1)$$

for given constants μ, σ and a positive initial price of s . By the variation of constants formula (see e.g. [13], Theorem 2.54) there exists a unique (strong) solution to the SDE (1.1) given by the *geometric BM*

$$S(t) = s \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W(t)\right).$$

As the logarithm of $S(t)$ is normally distributed, we speak of a *log-normal model*. In this case, we further have

$$\mathbb{E}(S(t)) = s \exp(\mu t).$$

Multi-dimensional generalizations of this example are available for linear coefficient functions $\mu(\cdot), \sigma(\cdot)$ without dependence on the state process $Y(t)$.

A popular example for a stock price model with dependence on an underlying state process is the Stochastic Volatility (SV) model of Heston (for short: *Heston model*, see [10]). There, we have one stock price and an additional state process $v(t)$ that is called the *volatility*. They are given by

$$\begin{aligned} dS(t) &= S(t) \left(\mu dt + \sqrt{v(t)} dW^S(t) \right), \quad S(0) = s, \\ dv(t) &= \kappa(\theta - v(t)) dt + \sigma \sqrt{v(t)} dW^V(t), \quad v(0) = s \end{aligned}$$

with arbitrary constants μ, σ and positive constants κ, θ . Further, we assume

$$\text{corr}(W^S(t), W^V(t)) = \rho$$

for a given constant $\rho \in [-1, 1]$ for the two one-dimensional Brownian motions W^S and W^V . A particular aspect of the volatility process $v(t)$ is that it is non-negative, but can attain the value zero if we have

$$2\theta\kappa \leq \sigma^2.$$

The Heston model is one of the benchmark models in the finance industry that will also appear in further contributions to this book. One of its particular challenges is that the corresponding SDE does not admit an explicit solution. Thus, it can only be handled by simulation and discretization methods, a fact that is responsible for many computational issues raised in this book.

1.3 Principles of Option Pricing

Options are derivative securities as their future payments depend on the performance of one or more underlying stock prices. They come in many ways, plain and simple, and complicated, with many strange features when it comes to determine the actual final payment that their owner receives. As they are a characteristic product of modern investment banking, calculating their prices in an efficient and accurate way is a key task in financial mathematics.

The most popular example of an option is the European call option on a stock. It gives its owner the right (but not the obligation!) to buy one unit of the stock at a predefined future time T (the *maturity*) for an already agreed price of K (the *strike*). As the owner will only buy it when the price of the underlying at maturity $S(T)$ is above the strike, the European call option is identified with the random payment of

$$H = (S(T) - K)^+$$

at time T .

One of the reasons for the popularity of the European call option is that it admits an explicit pricing formula in the BS model, the BS formula

$$c(t, S(t)) = S(t) \Phi \left(\frac{\ln \left(\frac{S(t)}{K} \right) + \left(r + \frac{1}{2} \sigma^2 \right) (T - t)}{\sigma \sqrt{T - t}} \right) - K e^{-r(T-t)} \Phi \left(\frac{\ln \left(\frac{S(t)}{K} \right) + \left(r - \frac{1}{2} \sigma^2 \right) (T - t)}{\sigma \sqrt{T - t}} \right)$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. This formula which goes back to [2] is one of the cornerstones of modern financial mathematics. Its importance in both theory and application is also emphasized by the fact that Myron Scholes and Robert C. Merton were awarded the Nobel Prize in Economics in 1997 for their work related to the BS formula.

The most striking fact of the BS formula is that the stock price drift μ , i.e. the parameter that determines the expected value of $S(t)$, does not enter the valuation formula of the European call option. This is no coincidence, but a consequence of a deep theoretical result. To formulate it, we introduce a riskless investment opportunity, the so-called money market account with price evolution $M(t)$ given by

$$M(t) = e^{rt},$$

i.e. the evolution of the value of one unit of money invested at time $t = 0$ that continuously earns interest payments at rate r .

The financial market made up of this money market account and the stock price of the BS model is called the BS market. There, we have the following result:

Theorem 1 (Option price in the BS model). *The price X_H of an option given by a final payment H with $\mathbb{E}(H^b) < \infty$ for some $b \geq 1$ is uniquely determined by*

$$X_H = \tilde{\mathbb{E}}(e^{-rT} H),$$

where the expectation is taken with respect to the unique probability measure Q under which the discounted stock $\tilde{S}(t) = S(t)/M(t)$ is a martingale. In particular, for the purpose of option pricing, we can assume that we have

$$\mu = r.$$

The reason for this very nice result is the completeness of the BS market, i.e. the fact that every (sufficiently integrable) final payoff H of an option can be created in a synthetic way by following an appropriate trading strategy in the money market account and the stock (see [13] for the full argumentation and the concept of completeness and replication).

In market models where the state process $Y(t)$ has a non-vanishing stochastic component that is not contained in the ones of the stock price, one does not have such a nice result as in the BS setting. However, even there, we can assume that for the purpose of option pricing we can model the stock prices in such a way that their discounted components $\tilde{S}_i(t) = S_i(t)/M(t)$ are martingales. In particular, now and in the following we directly assume that we only consider probability measures P such that we have

$$S_i(0) = \mathbb{E}(S_i(t)/M(t)), \quad i = 1, \dots, d.$$

Thus, all trade-able assets are assumed to have the same expected value for their relative increase in this artificial market. We therefore speak of *risk-neutral valuation*.

As we have now seen, calculating an option price boils down to calculating an expected value of a function or a functional of an underlying stochastic process. For simplicity, we thus assume that the underlying (possibly multi-dimensional) stock price process is given as the unique solution of the SDE

$$dS(t) = \mu(t, S(t))dt + \sigma(t, S(t))dW(t), \quad S(0) = s$$

with $W(t)$ a d -dimensional BM, $s = (s_1, \dots, s_n)'$, and μ, σ being functions satisfying appropriate conditions such that the above SDE possesses a unique (strong) solution. Further, for the moment, we consider a function

$$f: \mathbb{R}^m \rightarrow \mathbb{R}$$

which is non-negative (or polynomially bounded). Then, we can define the (conditional) expectation

$$\mathcal{V}(t, s) = \mathbb{E}^{(t, s)} \left(e^{-r(T-t)} f(S(T)) \right)$$

for a given starting time $t \in [0, T]$ at which we have $S(t) = s$. Of course, we can also replace the function f in the two preceding equations by a functional F that can depend on the whole path of the stock price. However, then the conditional expectation at time t above is in general not determined by only starting in (t, x) . Depending on the functional's form one needs more information of the stock price performance before time t to completely describe the current value of the corresponding option via an appropriate expectation.

However, in any case, to compute this expectation is indeed our main task. There are various methods for computing it. Examples are:

- *Direct calculation of the integral*

$$\mathbb{E}^{(0,s)}(e^{-rT} f(S(T))) = e^{-rT} \int_{\mathbb{R}^m} f(x) h(x) dx$$

if the density $h(\cdot)$ of $S(T)$ (conditioned on $S(0) = s$) is explicitly known.

- *Approximation of the price process $S(t), t \in [0, T]$, by simpler processes – such as binomial trees – $S^{(n)}(t), t \in [0, T]$, and then calculating the corresponding expectation*

$$\mathbb{E}_n^{(0,s^{(n)})} \left(f \left(S^{(n)}(T) \right) \right)$$

in the simpler model as an approximation for the original one (see [15] for a survey on binomial tree methods applied to option pricing).

- *Solution of the partial differential equation* for the conditional expectation $\mathcal{V}(t, s)$ that corresponds to the stock price dynamics. For notational simplicity, we only state it in the one-dimensional case as

$$\mathcal{V}_t(t, s) + \frac{1}{2} \sigma(t, s)^2 \mathcal{V}_{ss}(t, s) + \mu(t, s) \mathcal{V}_s(t, s) - r(t, s) \mathcal{V}(t, s) = 0,$$

$$\mathcal{V}(T, s) = f(T, s)$$

For more complicated option prices depending on the state process $Y(t)$, we also obtain derivatives with respect to y and mixed derivatives with respect to t, s, y .

- *Calculating the expectation via MC simulation*, i.e. simulating the final payoff H of an option N times and then estimating the option price via

$$\mathbb{E}(e^{-rT} H) \approx e^{-rT} \frac{1}{N} \sum_{i=1}^N H^{(i)}$$

where the $H^{(i)}$ are independent copies of H .

We will in the following restrict ourselves to the last method, the MC method. The main reason for this decision is that it is the most flexible of all the methods presented, and it suffers the most from heavy computations, as the number N of simulation runs usually has to be very large.

Before doing so, we will present options with more complicated payoffs than a European call option, so called *exotic options*. Unfortunately, only under very special and restrictive assumptions, there exist explicit formulae for the prices of such exotic options. Typically, one needs numerical methods to price them. Some popular examples are:

- Options with payoffs depending on multiple stocks such as *basket options* with a payoff given by

$$H_{basket} = \left(\frac{1}{d} \sum_{j=1}^d S_j(T) - K \right)^+$$

- Options with payoffs depending on either a finite number of stock prices at different times $t_j \in [0, T]$ such as *discrete Asian options* given by e.g.

$$H_{disc. Asian call} = \left(\frac{1}{d} \sum_{j=1}^d S(t_j) - K \right)^+$$

or a continuous average of stock prices such as *continuous Asian options* given by

$$H_{cont. Asian call} = \left(\frac{1}{T} \int_0^T S(t) dt - K \right)^+$$

- Barrier options that coincide with plain European put or call options as long as certain barrier conditions are either satisfied on $[0, T]$ or are violated such as e.g. a *knock-out-double-barrier call option* with a payoff given by

$$H_{dbkoc} = (S(T) - K)^+ 1_{\{B_1 < S(t) < B_2\}}$$

for constants $0 \leq B_1 < B_2 \leq \infty$

- Options with local and global bounds on payoffs such as *locally and globally capped and floored cliquet options* given by

$$H_{cliquet} = \max \left\{ F, \min \left\{ C, \sum_{j=1}^d \max \left\{ F_j, \min \left\{ C_j, \frac{S(t_j) - S(t_{j-1})}{S(t_{j-1})} \right\} \right\} \right\} \right\}$$

for different time instants $0 \leq t_0 < t_1 < \dots < t_d \leq T$ and constants $F < C, F_j < C_j$

All of those exotic options are tailored to the needs of special customers or markets. As an example, cliquet options are an essential ingredient of modern pension insurance products.

At the end of this section, we will formulate our first computational challenge:

Computational challenge 1: Find a universal framework/method for an efficient calculation of prices of exotic options.

An obvious candidate is the Monte Carlo (MC) method which we are going to present in the next section.

1.4 Monte Carlo Methods for Pricing Exotic Options

MC methods are amongst the simplest methods to compute expectations (and thus also option prices) and are on the other hand a standard example of a method that causes a big computing load when applied in a naive way. Even more, we will show by an example of a simple barrier option that a naive application of the MC method will lead to a completely wrong result that even pretends to be of a high accuracy.

Given that we can generate random numbers which are distributed as the considered real-valued, integrable random variable H , the standard MC method to calculate the expectation $\mathbb{E}(H)$ consists of two steps:

1. Generate N independent, identically distributed copies H_i of H .
2. Estimate $\mu \triangleq \mathbb{E}(H)$ by

$$\hat{\mu}_N = \frac{1}{N} \sum_{i=1}^N H_i.$$

Due to the linearity of the expectation the MC estimator $\hat{\mu}_N$ is unbiased. Further, the convergence of the standard MC method is ensured by the strong law of large numbers. One obtains an approximate confidence interval of level $1 - \alpha$ for μ as (see e.g. [14], Chapter 3)

$$\left[\frac{1}{N} \sum_{i=1}^N H_i - z_{1-\alpha/2} \frac{\sigma}{\sqrt{N}}, \frac{1}{N} \sum_{i=1}^N H_i + z_{1-\alpha/2} \frac{\sigma}{\sqrt{N}} \right].$$

Here, $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of the standard normal distribution and σ is defined via

$$\sigma^2 \triangleq \text{Var}(H).$$

If σ^2 is unknown (which is the typical situation) then it will be estimated by

$$\hat{\sigma}_N^2 = \frac{1}{N-1} \sum_{i=1}^N (H_i - \hat{\mu}_N)^2.$$

σ is then replaced by $\hat{\sigma}_N$ in the MC estimator of the confidence interval for μ . In both cases, the message is that – measured in terms of the length of the confidence interval – the **accuracy of the unbiased MC method is of the order $O(1/\sqrt{N})$** .

This in particular means that we need to increase the number of simulations of H_i by a factor 100 if we want to increase the accuracy of the MC estimator for μ by one order. Thus, we have in fact a very slow rate of convergence.

Looking at the ingredients in the MC method we already see the first challenge of an efficient and robust implementation:

Computational challenge 2: Find an appropriate Random Number Generator (RNG) to simulate the final payments H of an exotic option.

Here, the decision problem is crucial with respect to both performance and accuracy. Of course, the (typically deterministic) RNG should mimic the distribution underlying H as good as possible. Further, as the biggest computational advantage of the MC method is the possibility for parallelization, the RNG should allow a simple way of parallel simulation of independent random numbers.

The standard method here is to choose a suitable RNG that produces good random numbers that are uniformly distributed on $(0, 1)$ and to use the inverse transformation method for getting the right distribution. I.e. let U_i be the i th random number which is uniformly distributed on $(0, 1)$, let F be the desired distribution function of H . Then

$$H_i \triangleq F^{-1}(U_i)$$

has the desired distribution. This method mostly works, in particular in our diffusion process setting which is mainly dominated by the use of the normal distribution. Thus, for the normal distribution one only has to decide between the use of the classical Box-Muller transform or an approximate inverse transformation (see [14], Chapter 2). While the approximate inverse transformation method preserves a good grid structure of the original uniformly distributed random numbers, the Box-Muller transform ensures that even extreme values outside the interval $[-8, 8]$ can occur which is not the case for the approximate inverse method. Having made the decision about the appropriate transformation method, it still remains to find a good generator for the uniformly distributed random numbers U_i . Here, there is an enormous choice. As parallelization is one of the major advantages, the suitability for parallelization is a major issue for deciding on the RNG. Thus, the Mersenne Twister is a favorable choice (see [14], Chapter 2 and [16]).

For a simple standard option with a final payment of $H = f(S(T))$ (such as a European call option) in the Black-Scholes setting, we only have to simulate independent standard normally distributed random variables $Z_i, i = 1, \dots, N$, to obtain

$$H^{(i)} = f\left(se^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}Z_i}\right).$$

However, things become more involved when one either cannot generate the price process $S(t)$ exactly or when one can only simulate a suitably discretized version of the payoff functional.

For the first case, one has to use a discretization scheme for the simulation of the stock price (see [12] for a standard reference on the numerical solution of SDE). The most basic such scheme is the Euler-Maruyama scheme (EMS). To illustrate it, we apply it to a one dimensional SDE

$$dS(t) = \mu(S(t))dt + \sigma(S(t))dW(t), S(0) = s_0.$$

Then, for a step size of $\Delta = T/n > 0$, the discretized process $S^{(\Delta)}(t)$ generated by the EMS is defined by

$$\begin{aligned} S^{(\Delta)}(0) &\triangleq s_0, \\ S^{(\Delta)}(k\Delta) &\triangleq S^{(\Delta)}((k-1)\Delta) + \mu\left(S^{(\Delta)}((k-1)\Delta)\right)\Delta \\ &\quad + \sigma\left(S^{(\Delta)}((k-1)\Delta)\right)\Delta W_k, \quad k = 1, \dots, n. \end{aligned}$$

Here, $\Delta W_k, k = 1, \dots, n$, is a sequence of independent, $\mathcal{N}(0, \Delta)$ -distributed random variables. Between two consecutive discretization points, we obtain the values of $S^{(\Delta)}(t)$ by linear interpolation. The EMS can easily be generalized to a multi-dimensional setting.

If we now replace the original process $S(t)$ by $S^{(\Delta)}(t)$ in the standard MC approach, then we obtain

$$\hat{\mu}_{N,\Delta} \triangleq \frac{1}{N} \sum_{i=1}^N f\left(S_i^{(\Delta)}(T)\right) \xrightarrow{a.s.} \mathbb{E}\left(f\left(S^{(\Delta)}(T)\right)\right) \text{ for } N \rightarrow \infty$$

In particular, this application of the MC that uses the discretized process leads to a biased result. The accuracy of the MC method can then no longer be measured by the variance of the estimator. We have to consider the Mean Squared Error (MSE) to judge the accuracy instead, i.e.

$$\begin{aligned} MSE(\hat{\mu}_{N,\Delta}) &\triangleq \mathbb{E}\left[\left(\hat{\mu}_{N,\Delta} - \mathbb{E}(f(S(T)))\right)^2\right] \\ &= \text{Var}\left(\hat{\mu}_{N,\Delta}\right) + \left(\mathbb{E}(f(S(T))) - \mathbb{E}\left(f\left(S^{(\Delta)}(T)\right)\right)\right)^2 \end{aligned}$$

Thus, the MSE consists of two parts, the MC variance and the so-called discretization bias. We consider this bias a bit more detailed by looking at the convergence behavior of the EMS: Given suitable assumptions on the coefficient functions μ, σ , we have weak convergence of the MSE of order 1 (see e.g. [12]). More precisely, for μ, σ being four times continuously differentiable we have

$$\left|\mathbb{E}(f(S(T))) - \mathbb{E}\left(f\left(S^{(\Delta)}(T)\right)\right)\right| \leq C_f \Delta$$

for four times differentiable and polynomially bounded functions f and a suitable constant C_f not depending on Δ .

With regard to the MSE it is optimal to choose the discretization step size $\Delta = T/n$ and the number of MC simulations N in such a way that both components of the MSE are of the same order. So, given that we have weak convergence of order 1 for the EMS then an MSE of order $\varepsilon^2 = 1/n^2$ can be obtained by the choices of

$$\Delta = O(1/n), \quad N = n^2$$

which lead to an order of $O(n^3)$ measured in the random numbers simulated in total. As this leads to a high computational effort for pricing an option by the standard MC method, we can formulate another computational challenge:

Computational challenge 3: Find a modification of the standard MC method that has an effort of less than $O(n^3)$ for pricing an option including path simulation.

There are some methods now available that can overcome this challenge. Among them are weak extrapolation, the statistical Romberg method and in particular the multi-level MC method which will also play a prominent role in further contributions to this book (see e.g. [14] for a survey on the three mentioned methods).

However, unfortunately, the assumptions on f are typically not satisfied for option type payoffs (simply consider all the examples given in the last section). Further, the assumptions on the coefficients of the price process are not satisfied for e.g. the Heston model.

Thus, in typical situations, although we know the order of the MC variance, we cannot say a lot about the actual accuracy of the MC estimator. This problem will be illustrated by the second case mentioned above where we have to consider the MSE as a measure for accuracy of the MC method, the case where the payoff functional can only be simulated approximately. Let therefore $f(S)$ be a functional of the path of the stock price $S(t)$, $t \in [0, T]$ and $\hat{\mu}_{N,\Delta}$ be a MC estimator based on N simulated stock price paths with a discretization step size for the payoff functional of Δ . Then, we obtain a similar decomposition of the MSE

$$\begin{aligned} MSE(\hat{\mu}_{N,\Delta}) &= \mathbb{E}(\hat{\mu}_{N,\Delta} - \mathbb{E}(f(S)))^2 \\ &= \mathbb{E}\left((\hat{\mu}_{N,\Delta} - \mathbb{E}(f(S;\Delta)))^2\right) + \mathbb{E}\left((\mathbb{E}(f(S)) - \mathbb{E}(f(S;\Delta)))^2\right) \\ &= Var(\hat{\mu}_{N,\Delta}) + bias(\Delta) \end{aligned}$$

where now the bias is caused by the discretization of the payoff functional.

To illustrate the dependence of the accuracy of the MC method on the bias, we look at the problem of computing the price of a one-sided down-and-out barrier call option with a payoff functional given by

$$f(S(t); t \in [0, T]) = (S(T) - K)^+ 1_{S(t) > B \forall t \in [0, T]}.$$

As the one-sided down-and-out barrier call option possesses an explicit valuation formula in the BS model (see e.g. [13], Chapter 4), it serves well to illustrate the effects of different choices of the discretization parameter $\Delta = 1/m$ and the number of MC replications N .

As input parameters we consider the choice of

$$T = 1, r = 0, \sigma = 0.1, S(0) = K = 100, B = 95.$$

We first fix the number of discretization steps m to 10, i.e. we have $\Delta = 0.1$. As we then only check the knock-out condition at 10 time points, the corresponding MC estimator (at least asymptotically for large N) overestimates the true value of the barrier option. This is underlined in Fig. 1.2 where the 95%-confidence intervals do not contain the true value of the barrier option. This, however is not surprising as in this case the sequence of MC estimators converges to the price of the discrete down-and-out call given by the final payoff

$$f(S; N, \Delta) = (S(T) - K)^+ 1_{S(i\Delta T/N) > B \forall i=1, \dots, m}.$$

As a contrast, we now fix the number $N = 100,000$ of simulated stock price paths and consider a varying number of discretization points m in Fig. 1.3. As can be seen from the nearly identical length of the confidence intervals for varying m , the variance of the MC estimator is estimated consistently. Considering the differences




Fig. 1.2 MC estimators with 95%-confidence intervals for the price of a barrier option with fixed time discretization 0, 1 and varying number N of simulated stock price paths




Fig. 1.3 MC estimators with 95%-confidence intervals for the price of a barrier option with varying time discretization $1/m$ for 100,000 stock price paths

of the bias of the different MC estimators from the true value, one can conjecture that the bias behaves as $O(1/\sqrt{m})$, and thus converges at the same speed as the unbiased MC estimator. This example highlights that the order of the convergence of the bias is the critical aspect for the MC method in such a situation. Fortunately, in the case of the barrier options, there are theoretical results by Gobet (see e.g. [9]) that prove the above conjecture of a discretization bias of order 0.5. There are also good modifications of the crude MC method above that produce an unbiased estimator (such as the Brownian bridge method (see e.g. Chapter 5 of [14])), but the effects demonstrated above are similar for other types of exotic options. And moreover, there are not too many results on the bias of the MC estimator for calculating the price of an exotic option.

Thus, in calculating the prices of exotic options by MC methods, we face another computational challenge:

Computational challenge 4: Develop an efficient algorithm to estimate the order of the discretization bias when calculating the price of an exotic option with path dependence by the MC method.

A possibly simple first suggestion is to perform an iterative search in the following way:

1. Start with a rough discretization (i.e. a small number m) and increase the number N of MC simulation runs until the resulting (estimated) variance is below the order of the desired size of the MSE.

2. Increase the number of discretization points by a factor 10 and repeat calculating the corresponding MC estimation with the final N from Step 1 10 times. Take the average over the 10 calculations as an estimator for the option price.
3. Repeat Step 2 until the estimator for the option price is no longer significantly changing between two consecutive steps.

Of course, this is only a kind of simple cooking recipe that leaves a lot of space for improvement. One can also try to estimate the order of the discretization bias from looking at its behavior as a function of the varying step size $1/(10^k m)$.

In any case, not knowing the discretization bias increases the computational effort enormously, if one wants to obtain a trustable option price by the MC method. So, any strategy, may it be more based on algorithmic improvements or on an efficient hardware/software concept, will be a great step forward.

1.5 Risk Measurement and Management

The notion of risk is ubiquitous in finance, a fact that is also underlined by the intensive use of such terms as market risk, liquidity risk, credit risk, operational risk, model risk, just to mention the most popular names. As measuring and managing risk is one of the central tasks in finance, we will also highlight some corresponding computational challenges in different areas of risk.

1.5.1 Loss Distributions and Risk Measures

While we have concentrated on the pricing of single derivative contracts in the preceding sections, we will now consider a whole bunch of financial instruments, a so-called **portfolio** of financial positions. This can be the whole book of a bank or of one of its departments, a collection of stocks or of risky loans. Further, we will not price the portfolio (this would just be the sum of the single prices), but will instead consider the sum of the risks that are inherent in the different single positions simultaneously. What interests us is the potential change, particularly the losses, of the total value of the portfolio over a future time period.

The appropriate concepts for measuring the risk of such a portfolio of financial assets are those of the **loss function** and of **risk measures**. In our presentation, we will be quite brief and refer the reader for more details to the corresponding sections in [17] and [14].

We denote the **value** at time s of the portfolio under consideration by $V(s)$ and assume that the random variable $V(s)$ is observable at time s . Further, we assume that the composition of the portfolio does not change over the period we are looking at.

For a time horizon of Δ the **portfolio loss** over the period $[s, s + \Delta]$ is given by

$$L_{[s, s + \Delta]} \triangleq -(V(s + \Delta) - V(s)).$$

Note that we have changed the sign for considering the differences of the future and the current portfolio value. This is because we are concerned with the possibilities of big losses only. Gains do not play a big role in risk measurement, although they are the main aim of performing the business of a company in general.

Typical time horizons that occur in practice are 1 or 10 days or even a year. As $L_{[s, s + \Delta]}$ is not known at time s it is considered to be a random variable. Its distribution is called the **(portfolio) loss distribution**. We do not distinguish between the *conditional loss* and *unconditional loss* in the following as our objective are computational challenges. We always assume that we perform our computations based on the maximum information available at the time of computation.

As in [17] we will work in units of the fixed time horizon Δ , introduce the notation $V_t \triangleq V(t\Delta)$, and rewrite the loss function as

$$L_{t+1} \triangleq L_{[t\Delta, (t+1)\Delta]} = -(V_{t+1} - V_t). \quad (1.2)$$

Fixing the time t , the distribution of the loss function $L \triangleq L_{t+1}$ for $\ell \in \mathbb{R}$ (conditional on time t) is introduced using a simplified notation as

$$F_L(\ell) \triangleq P(L \leq \ell).$$

With the distribution of the loss function, we are ready to introduce so-called **risk measures**. Their main purpose is stated by Föllmer and Schied in [7] as:

...a risk measure is viewed as a capital requirement: We are looking for the minimal amount of capital which, if added to the position and invested in a risk-free manner, makes the position acceptable.

For completeness, we state:

A **risk measure** ρ is a real-valued mapping defined on the space of random variables (*risks*).

To bring this somewhat meaningless, mathematical definition closer to the above intention, there exists a huge discussion in the literature on reasonable additional requirements that a good risk measure should satisfy (see e.g. [7, 14, 17]).

As this discussion is beyond the scope of this survey, we restrict ourselves to the introduction of two popular examples of risk measures: The one which is mainly used in banks and has become an industry standard is the *value-at-risk*.

The **value-at-risk** of level α (VaR_α) is the α -quantile of the loss distribution of the portfolio:

$$VaR_\alpha(L) \triangleq \inf\{\ell \in \mathbb{R} | P(L > \ell) \leq 1 - \alpha\} = \inf\{\ell \in \mathbb{R} | F_L(\ell) \geq \alpha\},$$

where α is a high percentage such as 95 %, 99 % or 99.5 %.

By its nature as a quantile, values of VaR_α have an understandable meaning, a fact that makes it very popular in a wide range of applications, mainly for the measurement of market risks, but also in the areas of credit risk and operational risk management. VaR_α is not necessarily sub-additive, i.e. the $VaR_\alpha(X + Y) > VaR_\alpha(X) + VaR_\alpha(Y)$ for two different risks X, Y is possible. This feature is the basis for most of the criticism of using value-at-risk as a risk measure. Furthermore, as a quantile, VaR_α does not say anything about the actual losses above it.

A risk measure that does not suffer from these two drawbacks (compare e.g. [1]), and, which is therefore also popular in applications, is the *conditional value-at-risk*:

The **conditional value-at-risk** (or average value-at-risk) is defined as

$$CVaR_\alpha(L) \triangleq \frac{1}{1 - \alpha} \int_\alpha^1 VaR_\gamma(L) d\gamma.$$

If the probability distribution of L has no atoms, then the $CVaR_\alpha$ has the interpretation as the expected losses above the value-at-risk, i.e. it then coincides with the **expected shortfall** or **tail conditional expectation** defined by

$$TCE_\alpha(L) \triangleq \mathbb{E}(L | L \leq VaR_\alpha(L)).$$

As the conditional value-at-risk is the value at risk integrated w.r.t. the confidence level, both notions do not differ remarkably from the computational point of view. Thus, we will focus on the value-at-risk below.

However, as typically the portfolio value V and thus by (1.2) the loss function L depend on a d -dimensional vector of market prices for a very large dimension d , the loss function will depend on the market prices of maybe thousands of different derivative securities. This directly leads us to the first obvious computational challenge of risk management:

Computational challenge 5: Find an efficient way to evaluate the *loss function* of large portfolios to allow for a fast computation of the *value-at-risk*.

1.5.2 Standard Methods for Market Risk Quantification

The importance of the quantification of market risks is e.g. underlined by the popular JPMorgan's Risk Metrics document (see [18]) from the practitioners site or by the reports of the Commission for the Supervision of the Financial Sector (CSSF) (see [19]) from the regulatory point of view. This has the particular consequence that every bank and insurance company have to calculate risk measures, of course for different horizons. While for a bank, risk measures are calculated typically for a horizon of 1–10 days, insurance companies typically look at the horizon of a year.

To make a huge portfolio numerically tractable, one introduces so-called **risk factors** that can explain (most of) the variations of the loss function and ideally reduce the dimension of the problem by a huge amount. They can be log-returns of stocks, indices or economic indicators or a combination of them. A classical method for performing such a model reduction and to find risk factors is a principal component analysis of the returns of the underlying positions.

We do not go further here, but simply assume that the portfolio value is modeled by a so-called **risk mapping**, i.e. for a d -dimensional random vector $\mathbf{Z}_t \triangleq (Z_{t,1}, \dots, Z_{t,d})'$ of risk factors we have the representation

$$V_t = f(t, \mathbf{Z}_t) \tag{1.3}$$

for some measurable function $f : \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}$. Of course, this representation is only useful if the risk factors \mathbf{Z}_t are observable at time t , which we assume from now on. By introducing the **risk factor changes** $(\mathbf{X}_t)_{t \in \mathbb{N}}$ by $\mathbf{X}_t \triangleq \mathbf{Z}_t - \mathbf{Z}_{t-1}$ the portfolio loss can be written as

$$L_{t+1}(X_{t+1}) = -(f(t+1, \mathbf{Z}_t + \mathbf{X}_{t+1}) - f(t, \mathbf{Z}_t)) \tag{1.4}$$

highlighting that the loss is completely determined by the risk factor changes.

In what follows we will discuss some standard methods used in the financial industry for estimating the value-at-risk.

1.5.2.1 The Variance-Covariance Method

The variance-covariance method is some crude, first-order approximation. Its basis is the assumption that risk factor changes X_{t+1} follow a multivariate normal distribution, i.e.

$$X_{t+1} \sim \mathcal{N}_d(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

where $\boldsymbol{\mu}$ is the mean vector and $\boldsymbol{\Sigma}$ the covariance matrix of the distribution.

The second fundamental assumption is that f is differentiable, so that we can consider a first-order approximation L_{t+1}^{lin} of the loss in (1.4) of the form

$$L_{t+1}^{lin}(X_{t+1}) \triangleq -\left(f(t, \mathbf{Z}_t) + \sum_{i=1}^d f_{z_i}(t, \mathbf{Z}_t) X_{t+1,i}\right). \quad (1.5)$$

As the portfolio value $f(t, \mathbf{Z}_t)$ and the relevant partial derivatives $f_{z_i}(t, \mathbf{Z}_t)$ are known at time t , the linearized loss function has the form of

$$L_{t+1}^{lin}(X_{t+1}) = -(c_t + \mathbf{b}'_t \mathbf{X}_{t+1}) \quad (1.6)$$

for some constant c_t and a constant vector \mathbf{b}_t which are known to us at time t . The main advantage of the above two assumptions is that the linear function (1.6) of \mathbf{X}_{t+1} preserves the normal distribution and we obtain

$$L_{t+1}^{lin}(X_{t+1}) \sim \mathcal{N}(-c_t - \mathbf{b}'_t \boldsymbol{\mu}, \mathbf{b}'_t \boldsymbol{\Sigma} \mathbf{b}_t).$$

This yields the following explicit formula:

The **value-at-risk of the linearized loss** corresponding to the confidence level α is given by

$$VaR_\alpha(L_{t+1}^{lin}) = -c_t - \mathbf{b}'_t \boldsymbol{\mu} + \sqrt{\mathbf{b}'_t \boldsymbol{\Sigma} \mathbf{b}_t} \Phi^{-1}(\alpha), \quad (1.7)$$

where Φ denotes the standard normal distribution function and $\Phi^{-1}(\alpha)$ is the α -quantile of Φ .

To apply the value-at-risk of the linearized loss to market data, we still need to estimate the mean vector $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{\Sigma}$ based on the historical risk factor changes $\mathbf{X}_{t-n+1}, \dots, \mathbf{X}_t$ which can be accomplished using standard estimation procedures (compare Section 3.1.2 in [17]).

Remark 1. The formulation of the variance-covariance method based on the first-order approximation L_{t+1}^{lin} in (1.5) of the loss is often referred to as the *Delta-approximation* in analogy to the naming of the first partial derivative with respect to underlying prices in option trading.

Remark 2. Another popular version of the variance-covariance method is the *Delta-Gamma-approximation* which is based on a second-order approximation of the loss function in order to capture the non-linear structure of portfolios that contain a high percentage of options. However, the general advantages and weaknesses of these methods are similar. We therefore do not repeat our analysis for the Delta-Gamma-approximation here.

Merits and Weaknesses of the Method

The main advantage of the variance-covariance method is that it yields an explicit formula for the value-at-risk of the linearized losses as given by (1.7). However, this closed-form solution is only obtained using two crucial simplifications:

1. Linearization (in case of the Delta-approximation) or even a second order approximation (in case of the Delta-Gamma-approximation) is in the fewest cases a good approximation of the risk mapping as given in (1.3), in particular when the portfolio contains many complex derivatives.
2. Empirical examinations suggest that the distribution of financial risk factor returns is leptokurtic and fat-tailed compared to the Gaussian distribution. Thus the assumption of normally distributed risk factor changes is questionable and the value-at-risk of the linearized losses (1.7) is likely to underestimate the true losses.

1.5.2.2 Historical Simulation

Historical simulation is also a very popular method in the financial industry. It is based on the simple idea that instead of making a model assumption for the risk factor changes, one simply relies on the **empirical distribution** of the already observed past data $\mathbf{X}_{t-n+1}, \dots, \mathbf{X}_t$. We then evaluate our portfolio loss function for each of those data points and obtain a set of synthetic losses that would have occurred if we hold our portfolio on the past days $t-1, t-2, \dots, t-n$:

$$\{\tilde{L}_s(\mathbf{X}_s) : s = t-n+1, \dots, t\}. \quad (1.8)$$

Based on these historically simulated loss data, one now estimates the value-at-risk by the corresponding *empirical quantile*, i.e. the quantile of the just obtained historical empirical loss distribution:

Let $\tilde{L}_{n,n} \leq \dots \leq \tilde{L}_{1,n}$ be the ordered sequence of the values of the historical losses in (1.8). Then, the **estimator for the value-at-risk obtained by historical simulation** is given by

$$\text{VaR}_\alpha(\tilde{L}_s) \triangleq \tilde{L}_{[n(1-\alpha)],n},$$

where $[n(1-\alpha)]$ denotes the largest integer not exceeding $n(1-\alpha)$.

Merits and Weaknesses of the Method

Besides being a very easy method, a convincing argument of historical simulation is its independence on distributional assumptions. We only use data that have already appeared, no speculative ones.

From the theoretical point of view, however, we have to assume stationarity of the risk factor changes over time which is also quite a restrictive assumption. And even more, we can be almost sure that we have not yet seen the worst case of losses in the past. The dependence of the method on reliable data is another aspect that can cause problems and can lead to a weak estimator for the value-at-risk.

1.5.2.3 The Monte Carlo Method

A method that overcomes the need for linearization and the normal assumption in the variance-covariance method and that does not rely on historical data is the Monte Carlo (MC) method. Of course, we still need an assumption for the distribution of the future risk factor changes.

Given that we have made our choice of this distribution, the MC method only differs to the historical simulation by the fact that we now simulate our data, i.e. we **simulate** independent identically distributed random future risk factor changes $\tilde{\mathbf{X}}_{t+1}^{(1)}, \dots, \tilde{\mathbf{X}}_{t+1}^{(M)}$, and then compute the corresponding portfolio losses

$$\left\{ \tilde{L}_{t+1}(\tilde{\mathbf{X}}_{t+1}^{(i)}) : i = 1, \dots, M \right\}. \quad (1.9)$$

As in the case of the historical simulation, by taking the relevant quantile of the empirical distribution of the simulated losses we can estimate the value-at-risk:

The **MC estimator for the value-at-risk** is given by

$$\text{VaR}_\alpha(\tilde{L}_{t+1}) \triangleq \inf \{ \ell \in \mathbb{R} \mid \tilde{F}_{t+1}(\ell) \geq \alpha \},$$

(continued)

(continued)

where the empirical distribution function $\tilde{F}_{t+1}(\ell)$ is given by

$$\tilde{F}_{t+1}(\ell) \triangleq \frac{1}{M} \sum_{i=1}^M I_{\{\tilde{L}_{t+1}(\tilde{\mathbf{X}}_{t+1}^{(i)}) \leq \ell\}}.$$

Remark 3 (Some aspects of the MC method).

- (i) Of course, the crucial modeling aspect is the choice of the distribution for the risk factor changes and the calibration of this distribution to historical risk factor change data $\mathbf{X}_{t-n+1}, \dots, \mathbf{X}_t$. This can be a computational challenging problem itself (compare also Sect. 1.6.1 and the chapter by Sayer and Wenzel in this book).
- (ii) The above simulation to generate the risk factor changes is often named the *outer simulation*. Depending on the complexity of the derivatives included in the portfolio, we will need an *inner simulation* in order to evaluate the loss function of the risk factor changes. This means, we have to perform MC simulations to calculate the future values of options in each run of the outer simulation. As this is also an aspect of the historical simulation, we postpone this for the moment and assume that the simulated realizations of the loss distribution given by (1.9) are available.

Merits and Weaknesses of the Method

Of course, the quality of the MC method depends heavily on the choice of an appropriate distribution for the risk factor changes. On the up side, we are not limited to normal distributions anymore. A further good aspect is the possibility to generate as many loss values as one wants by simply choosing a huge value M of simulation runs. This is a clear advantage over the historical simulation where data are limited.

As there is no simplification to evaluate the portfolio, each simulation run will possibly need a huge computational effort, in particular if complicated options are held. On the other hand, this evaluation is then exact given the risk factor changes, which is a clear advantage compared to the variance-covariance method.

1.5.2.4 Challenges When Determining Market Risks

The Choice of a Suitable Risk Mapping

The above three methods have the main problem in common that it is not clear at all how to determine the appropriate risk factors yielding an accurate approximation of

the actual loss. On top of that, their dimension can still be remarkably high. This is a modeling issue and is closely connected to the choice of the function f in (1.3). As already indicated, performing a principal component analysis (compare e.g. [3]) can lead to a smaller number of risk factors which explain the major parts of the market risks. However, the question if the postulated risk factors approximate the actual loss well enough then remains still an issue and translates into the problem of the appropriate choice of the input for the principal component analysis.

The different approaches we explained above each have their own advantages and drawbacks. While the Delta-approximation is usually not accurate enough if the portfolio contains non-linear securities/derivatives, the Delta-Gamma-approximation already performs much better than the Delta-approximation. However, the resulting approximation of the loss function only has a known distribution if we stick to normally distributed risk factors. The most accurate results can be achieved by the MC method but at the cost of a high computational complexity compared to the other methods. The trade-off therein consists of balancing out accuracy and computability. Further, we sometimes have to choose between accuracy and a fast computation which can be achieved via a smart approximation of the loss function (especially with regard to the values of the derivatives in the portfolio). And in the end, the applicability of all methods highly depends on the structure of the portfolio at hand. Also, the availability of computing power can play an important role on the decision for the method to use. Thus, a (computational) challenge when determining market risks is the choice of the appropriate value-at-risk computation method.

(Computational) challenge 6: Given the structure of the portfolio and of the computing framework, find an appropriate algorithm to decide on the adequate method for the computation of the value-at-risk.

Nested Simulation

As already pointed out, in both the historical simulation and in the MC method we have to evaluate the portfolio in its full complexity. This computational challenge carries to extremes, when the portfolio contains a lot of complex derivatives, for which no closed-form price representation is available. In such a case, we will need an *inner MC simulation* in addition to the outer one to compute the realized losses.

To formalize this, assume for notational convenience that the time horizon Δ is fixed, that time $t + 1$ corresponds to time $t + \Delta$, and that the risk mapping $f : \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}$ corresponds to a portfolio of derivatives with payoff functions H_1, \dots, H_K with maturities T_1, \dots, T_K . From our main result Theorem 1.3 we know that the fair time- t price of a derivative is given by the discounted conditional expectation of its payoff function under the risk neutral measure Q (we here assume that our market

satisfies the assumptions of Theorem 1.3). Thus, the risk mapping f at time $t + \Delta$ is given by

$$f(t + \Delta, \mathbf{Z}_t + \tilde{\mathbf{X}}_{t+\Delta}^{(i)}) = \sum_{k=1}^K \tilde{\mathbb{E}} \left[e^{-r(T_k - (t+\Delta))} H_k | \tilde{\mathbf{X}}_{t+\Delta}^{(i)} \right], \quad (1.10)$$

where $\tilde{\mathbb{E}}(\cdot)$ denotes the expectation under the risk neutral measure \mathcal{Q} . For standard derivatives like European calls or puts the conditional expectations in (1.10) can be computed in closed-form (compare again Theorem 1.3). For complex derivatives, however, they have to be determined via MC simulation. This then causes an **inner simulation** as follows that has to be performed for **each** (!!!) realization of the outer simulation:

Inner MC simulation for complex derivatives in the portfolio:

1. Generate N independent realizations $H_k^{(1)}, \dots, H_k^{(N)}$ of the $k = 1, \dots, K$ (complex) payoffs given $\tilde{\mathbf{X}}_{t+\Delta}^{(i)}$.
2. Estimate the discounted conditional expectation of the payoff functions by

$$\tilde{\mathbb{E}} \left[e^{-r(T_k - (t+\Delta))} H_k | \tilde{\mathbf{X}}_{t+\Delta}^{(i)} \right] \approx \frac{1}{N} e^{-r(T_k - (t+\Delta))} \sum_{j=1}^N H_k^{(j)}$$

for $k = 1, \dots, K$.

Remark 4 (Important!).

- (i) The amount of simulation work in the presence of the need for an inner simulation is enormous as the inner simulations have to be redone for each run of the outer simulation. A possible challenge is to find a framework for reusing the simulations in the inner loop for each new outer simulation. A possibility could be to perform the inner simulations only a certain times and then setting up something as an interpolation polynomial for the price of the derivatives as a function of the risk factors.
- (ii) Note further, that for notational simplicity, we have assumed that each derivative in the inner simulation requires the same number N of simulation paths to achieve a desired accuracy for the MC price calculation. This, however, heavily depends on the similarity of the derivatives and the volatility of the underlyings. If the variety of option types in the portfolio is large, substantial savings can be obtained by having a good concept to choose the appropriate number of inner simulation runs per option type.

- (iii) As a minor issue, note that the input for the option pricing typically has to be the price of the underlying(s) at time $t + \Delta$ or even more, the paths of the price of the underlying(s) up to time $t + \Delta$. This input has to be reconstructed from the risk factor changes.
- (iv) Finally, the biggest issue is the load balance between inner and outer simulation. Given only a limited computing time and capacity, one needs a well-balanced strategy. Highly accurate derivative prices in the inner simulation lead to an accurate evaluation of the loss function (of course, conditioned on the correctness of the chosen model for the distribution of the risk factor changes). On the other hand, they cause a big computational effort which then results in the possibility of performing only a few outer simulation runs. This then leads to a poor estimate of the value-at-risk. A high number of outer simulation runs however only allows for a very rough estimation of the derivative prices on the inner run, again a non-desirable effect.

The foregoing remark points in the direction of the probably most important computational challenge of risk management:

Computational challenge 7: Find an appropriate concept for balancing the workload between the inner and outer MC simulation for the determination of the value-at-risk of complex portfolios and design an efficient algorithm that ensures sufficient precision of both the derivative prices in the inner simulation and the MC estimator for the value-at-risk on the outer simulation.

1.5.3 *Liquidity Risk*

Besides the measurement of market risks, another important strand of risk management is the measurement of liquidity risk. We understand thereby liquidity risk as the risk not to be able to obtain needed means of payment or to obtain them only at increased costs. In this article we will put emphasis on liquidity risks which arise in the fund management sector. Fund management focuses in particular on calling risk (liquidity risk on the liabilities side) which is the risk of unexpectedly high claims or claims ahead of schedule as for instance the redemption of shares in a fund. **Liquidity risk in fund management** has gained importance in recent times which manifests itself in European Union guidelines that require appropriate liquidity risk management processes for UCITS (=Undertakings for Collective Investment in Transferable Securities) and AIFMs (=Alternative Investment Funds Managers); compare therefore [20] and [21].

One approach which covers these liquidity risk regulations is to calculate the **peaks over threshold (POT) quantile of the redemptions** of mutual funds. It is well-known (compare e.g. [6]) that the excess distribution can be approximated by

the **generalized Pareto distribution (GPD)** from a certain threshold u . This fact is due to the famous theorem of Pickands, Balkema and de Haan, on which we will give a short mathematical excursion: Define the excess distribution of a real-valued random variable X with a distribution function F as

$$F_u(y) \triangleq P(X - u \leq y | X > u), \text{ where } 0 \leq y < x_F - u,$$

for a fixed right endpoint x_F

$$x_F \triangleq \sup \{x \in \mathbb{R} : F(x) < 1\} \leq \infty, \text{ where } u < x_F.$$

Then we have the following:

Theorem 2 (Pickands, Balkema, de Haan). *There exists an appropriate function $\beta(u)$ such that*

$$\lim_{u \uparrow x_F} \sup_{0 < x < x_F - u} |F_u(x) - G_{\xi, \beta(u)(x)}| = 0,$$

where

$$G_{\xi, \beta}(x) = \begin{cases} 1 - \left(1 + \frac{\xi x}{\beta}\right)^{-1/\xi} & \xi \neq 0, \\ 1 - e^{-x/\beta} & \xi = 0, \end{cases}, x \in D(\xi, \beta) = \begin{cases} [0, \infty) & \xi \geq 0, \\ [0, -\beta/\xi] & \xi < 0, \end{cases}$$

is the generalized Pareto distribution (GPD) with shape parameter $\xi \in \mathbb{R}$ and scale parameter $\beta > 0$.

As a consequence, the excess distribution can be approximated in a similar way by a suitable generalized Pareto distribution as the distribution of a sum can be approximated by the normal distribution. The quantile of the excess distribution then gives a liquidity reserve which is not exceeded by a certain probability p and is called POT quantile. The POT quantile is also referred to as **liquidity-at-risk** and was applied by [22] for the banking sector. Desmettre and Deege [5] then adapted it to the mutual funds sector and provided a thorough backtesting analysis.

The p -quantile of the excess distribution, i.e. the liquidity-at-risk, is given as

$$LaR_p \triangleq u + \frac{\hat{\beta}}{\hat{\xi}} \left(\left(\frac{n}{N_u} (1-p) \right)^{-\hat{\xi}} - 1 \right), \quad (1.11)$$

where N_u is the number of exceedances over the threshold u , n is the sample size, $\hat{\xi}$ is an estimator for the shape parameter and $\hat{\beta}$ is an estimator for the scale parameter of the generalized Pareto distribution.

Thus in order to calculate the liquidity-at-risk, it is necessary to estimate the threshold parameter u , the shape parameter ξ and the scale parameter β of the GPD. The estimation of shape and scale parameter can be achieved using standard maximum likelihood estimators; a procedure for the **estimation of the threshold parameter u** and also its detailed derivation is also given in [5] and is the time-consuming part when computing the liquidity-at-risk as given by (1.11). In what follows we sketch the calibration method and explain how it leads to a computational challenge.

Using well-known properties of the generalized Pareto distribution $G_{\xi,\beta}$, we can conclude that the estimator $\hat{\xi}$ of the scale parameter ξ of the excess distribution (which is approximated by a suitable GPD) is approximately invariant under shifts in the threshold parameter u . Thus a procedure for the determination of the threshold parameter u is given by

Choose the first threshold parameter $u > 0$ such that the estimator $\hat{\xi}$ of the shape parameter ξ of the corresponding GPD is approximately invariant under shifts in the threshold parameter $u > 0$.

The implementation of this method can be sketched as follows (see also [5]):

1. Sort the available data by ascending order and keep a certain percentage of the data.
2. Start with u being the lowest possible threshold and increase it up to the value for which at least k percent of the original data are left. With increasing threshold u truncate the data at the threshold u .
3. Estimate the unknown parameters ξ and β of the GPD by their maximum likelihood estimators for every u from 2.
4. For each u , calculate a suitable deviation measure of the corresponding maximum likelihood estimators $\hat{\xi}(i)_{i=1,\dots,K}$ within a sliding interval.
5. The appropriate threshold u is determined as the threshold which lies in the middle of the interval with the lowest deviation measure. Take the number of exceedances N_u corresponding to this u and the sample size n .
6. The estimates $\hat{\xi}$ and $\hat{\beta}$ are the maximum likelihood estimates which correspond to the threshold u .

The computational challenge now arises when we look at typical data sets. Often, fund redemption data is available over a quite long time horizon such that a time series of a single share class can contain thousands of data points. Moreover, management companies will typically face a large portfolio of share classes which can have a dimension of several hundreds. Combining these two facts we see that a fund manager will have to determine a large amount of estimates for the shape and scale parameter of the Generalized Pareto distribution in order to calibrate the threshold parameter u (compare steps 1–4 of the above algorithm) for a daily

liquidity risk management process of her portfolio. Therefore it is important to have a grip on a fast calibration of the threshold and our next computational challenge can be formulated as

Computational challenge 8: Speed up the calibration of the threshold parameter u for a fast computation of the liquidity-at-risk.

1.5.4 *Intraday Simulation and Calculation*

Up to now we considered daily or yearly time horizons Δ . Nowadays in practice, the demand for so called *intraday calculations and calibrations* is growing, i.e. we face time horizons $\Delta \ll 1$ day and in the extreme the time horizon can have the dimension of a few hours or even 15 and 30 min which represents the time horizon of intraday returns. Especially within times of crises it may be of use to be able to recalibrate all corresponding risk measures of portfolios in order to have as much information as possible. This will allow fund managers to take well-founded decisions. For a concise overview of intraday market risk we refer to [8].

The recalibration and recalculation of the risk measures typically involves a reevaluation of the actual portfolio value as we have for instance seen within the nested simulations of the MC value-at-risk method. Therefore the *intraday evaluation of large portfolios* is also of importance. Summarizing our considerations above we face the computational challenge

Computational challenge 9: Speed up the calculation and calibration of the risk management process of financial firms such that intraday calculations become feasible.

1.6 Further Aspects of Computationally Challenging Problems in Financial Markets

Besides the optimization of MC methods and of risk management calculations, there are various other computational issues in financial mathematics. We will mention only three more, two of them are very important from a practical point of view, the other has big consequences for designing an efficient hardware/software concept:

1.6.1 Calibration: How to Get the Parameters?

Every financial market model needs input parameters as otherwise we cannot calculate any option price or, more general, cannot perform any type of calculation. To highlight the main approach at the derivatives markets to obtain the necessary parameters we consider the BS model. There, the riskless interest rate r can (in principle) be observed at the market. The volatility σ however has to be determined in a suitable way. There are in principle two ways,

- A classical maximum likelihood estimation (or any other conventional estimation technique) based on past stock prices using the fact that the logarithmic differences (i.e. $\ln(S(t_i)/S(t_{i-1}))$, $\ln(S(t_{i-1})/S(t_{i-2}))$, ...) are independent,
- A calibration approach, i.e. the determination of the parameter σ_{imp} which minimizes the squared differences between model and market prices of traded options.

As the second approach is the one chosen at the derivatives markets, we describe it a little bit more detailed. Let us for simplicity assume that at a derivatives market we are currently observing only the prices c_{K_i, T_i} of n call options that are characterized by their strikes K_i and their (times to) maturities T_i . The calibration task now consists of solving

$$\min_{\sigma > 0} \sum_{i=1}^n (c_{K_i, T_i} - c(0, S(0); \sigma, K_i, T_i))^2$$

where $c(0, S(0); K, T)$ denotes the BS formula with volatility $\sigma > 0$, strike K and maturity T . Of course, one can also use a weighted sum as the performance measure to care for the fact that some of these options are more liquidly traded than others.

Note that calibration typically is a highly non-linear optimization problem that even gets more involved if more parameters have to be calibrated. We also recognize the importance of having closed pricing formulae in calibration. If the theoretical prices have to be calculated by a numerical method (say the MC method) then the computational effort per iteration step in solving the calibration problem increases dramatically.

For a much more complicated calibration problem we refer to the work by Sayer and Wenzel in this book.

1.6.2 Money Counts: How Accurate Do We Want Our Prices?

The financial markets are known for their requirement of extremely accurate price calculations. However, especially in the MC framework, a huge requirement for accurate prices increases the computational effort dramatically. It is therefore worth to point out that high accuracy is worthless if the parameter uncertainty (i.e. the

error in the input parameters), the algorithmic error (such as the order of (weak) convergence of the MC method) or the model error (i.e. the error caused by using an idealized model for simulation that will certainly not exactly mimic the real world price dynamics) are of a higher order than the accuracy of the performed computations.

On the other hand, by using a sparse number format, one can speed up the computations and reduce storage capacity by quite a factor. It is therefore challenging to find a good concept for a variable treatment of precision requirements.

For an innovative suggestion of a mixed precision multi-level MC framework we refer to the work by Omland, Hefter and Ritter in this book.

1.6.3 *Data Maintenance and Access*

All mentioned computational methods in this article have in common that they can only be efficiently executed once the data is available and *ready to use*. A good many times, the data access takes as much time as the computations themselves. In general, the corresponding data like market parameters or information about the composition of derivatives and portfolios are stored in large data bases whose maintenance can be time-consuming; for an overview on the design and maintenance of database systems we refer to the textbook of Connolly and Begg [4].

In that regard it is also very useful to thoroughly review the computations that have to be done and to do them in a clever way; for instance a smart approximation of the loss function where feasible may already tremendously accelerate the value-at-risk computations. We thus conclude with the computational challenge

Computational challenge 10: Maintain an efficient data storage and provide an efficient data access.

Acknowledgements Both authors are grateful to the Deutsche Forschungsgemeinschaft for funding within the Research Training Group 1932 “Stochastic Models for Innovations in the Engineering Sciences”. Sascha Desmettre wishes to thank Matthias Deege from IPConcept (Luxemburg) S.A. and Heiko Reiss for useful discussions about the practical challenges of derivative pricing and risk management.

References

1. Acerbi, C., Tasche, D.: On the coherence of expected shortfall. *J. Bank. Financ.* **26**(7), 1487–1503 (2002)
2. Black, F., Scholes, M.: The pricing of options and corporate liabilities. *J. Political Econ.* **81**, 637–654 (1973)

3. Brummelhuis, R., Cordoba, A., Quintanilla, M., Seco, L.: Principal component value at risk. *Math. Financ.* **12**(1):23–43 (2002)
4. Connolly, T.M., Begg, C.E.: *Database Systems: A Practical Approach to Design, Implementation and Management*. Prentice Hall, Upper Saddle River (2005)
5. Desmettre, S., Deege, M.: Liquidity at risk for mutual funds. Preprint available at SSRN. <http://ssrn.com/abstract=2440720> (2014)
6. Embrechts, P., Klüppelberg, C., Mikosh, T.: *Modeling Extremal Events*. Springer, Berlin (1997)
7. Föllmer, H., Schied, A.: *Stochastic Finance: An Introduction in Discrete Time*. de Gruyter, Berlin (2002)
8. Giot, P.: Market risk models for intraday data. *Eur. J. Financ.* **11**(4), 309–324 (2005)
9. Gobet, E.: Advanced Monte Carlo methods for barrier and related exotic options. In: Ciarlet, P., Bensoussan, A., Zhang, Q. (eds.) *Mathematical Modeling and Numerical Methods in Finance. Handbook of Numerical Analysis*, vol. 15, pp. 497–528. Elsevier, Amsterdam (2009)
10. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* **6**(2), 327–343 (1993)
11. Karatzas, I., Shreve, S.E.: *Brownian Motion and Stochastic Calculus*, 2nd edn. Springer, Berlin (1991)
12. Kloeden, P.E., Platen, E.: *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin (1999)
13. Korn, R., Korn, E.: *Option Pricing and Portfolio Optimization. Graduate Studies in Mathematics*, vol. 31. AMS, Providence (2001)
14. Korn, R., Korn, E., Kroisandt, G.: *Monte Carlo Methods and Models in Finance and Insurance. Chapman & Hall/CRC Financial Mathematics Series*. CRC, London (2010)
15. Korn, R., Müller, S.: Binomial trees in option pricing – history, practical applications and recent developments. In: Devroye, L., Karasozen, B., Kohler, M., Korn, R. (eds.) *Recent Developments in Applied Probability and Statistics*, pp. 119–138. Springer, Berlin (2010)
16. Matsumoto, M., Nishimura, T.: Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.* **8**, 3–30 (1998)
17. McNeil, A.J., Frey, R., Embrechts, P.: *Quantitative Risk Management*. Princeton University Press, Princeton/Oxford (2005)
18. RiskMetrics: *RiskMetrics Technical Document*, 4th edn. J.P. Morgan/Reuters, New York (1996)
19. The Commission for the Supervision of the Financial Sector: *CSSF Circular 11/512* (2011)
20. The European Commission: *Commission Directive 2010/43/EU*. Official Journal of the European Union (2010)
21. The European Parliament and the Council of the European Union: *Directive 2011/61/EU of the European Parliament and of the Council*. Official Journal of the European Union (2011)
22. Zeranski, S.: Liquidity at Risk – Quantifizierung extremer Zahlungsstromrisiken. *Risikomanager* **11**(1), 4–9 (2006)