Quadratic Maximum Loss for Risk Measurement of Portfolios

G. Studer
Prof. H.–J. Lüthi

RiskLab: Technical Report, September 1996

Abstract:

Effective risk management requires adequate risk measurement. A basic problem herein is the quantification of market risks: what is the overall effect on a portfolio if market rates change? The first chapter gives a brief review of the standard risk measure "Value–At–Risk" (VAR) and introduces the concept of "Maximum Loss" (ML) as a method for identifying the worst case in a given scenario space, called "Trust Region". Next, a technique for calculating efficiently ML for quadratic functions is described; the algorithm is based on the Levenberg–Marquardt theorem, which reduces the high–dimensional optimization problem to a one–dimensional root finding.
Following this, the idea of the "Maximum Loss Path" is presented: repetitive calculation of ML for a growing trust region leads to a sequence of worst cases, which form a complete path. Similarly, the paths of "Maximum Profit" (MP) and "Expected Value" (EV) can be determined; the comparison of them permits judgements on the quality of portfolios. These concepts are also applicable to non–quadratic portfolios by using "Dynamic Approximations", which replace arbitrary profit and loss functions with a sequence of quadratic functions.
Finally, the idea of "Maximum Loss Distribution" is explained. The distributions of ML and MP can be obtained directly from the ML and MP paths. They lead to lower and upper bounds of VAR and allow statements about the spread of ML and MP.

Keywords: Risk Measurement — Market Risk — Value At Risk — Maximum Loss Optimization

Contents

Chapter 1

Introduction

1.1 Problem Statement

An increasing number of complex financial instruments is used in today's business. Whereas many quantitative techniques for the analysis of a single instrument are well established, there is still a lack of mathematical methods for estimating the cumulative risk of a portfolio. One particular problem is the measurement of market risks: what loss can a portfolio suffer from if the market rates vary? In this paper we present new concepts, new models and new quantitative techniques to answer this question.
Mathematically, the problem can be formulated as follows: market rates (e.g. commodity prices, foreign exchange rates, equity indices, interest rates) are represented by risk factors $\omega_1, \ldots, \omega_M$, shifted such that $\omega_i = 0$ corresponds to the actual value of market rate i. Risk factors behave randomly: for a time interval of length t (holding period of the portfolio), $\omega = (\omega_1, \ldots, \omega_M)$ are supposed to be multinormally distributed variates with mean 0 and covariance matrix Σ_t (cf. Appendix A).
The effect of a change in the risk factors ω can be determined by revaluating of the portfolio. The *change* in portfolio value — called "profit and loss" (P&L) — is denoted $v(\omega)$; the definitions imply that $v(0) = 0$.

1.2 Value–At–Risk (VAR)

A widely used measure of market risk is "Value–At–Risk" (cf. [RiskMetrics], [Beckström and Campbell]) or VAR, which has an intuitive interpretation, but requires in practice often time consuming calculations. For a fixed confidence level α, VAR can be defined as:

The level of loss
- which is exceeded in $(100 - \alpha)$ percent of all situations
- if the portfolio's value is examined after t days (holding period).

Mathematically, VAR is the α–quantile of the P&L distribution (α is typically chosen 95 % or 99 %). For a linear portfolio (i.e., $v(\omega) = \delta^T \omega$), the normal distribution of the risk factors implies that $v(\omega) \sim \mathcal{N}(0, \delta^T \Sigma_t \delta)$. Therefore, VAR can be written analytically as

$$\text{VAR} = -z_\alpha \sqrt{\delta^T \Sigma_t \delta}, \tag{1.1}$$

where z_α denotes the α–quantile of the standard normal distribution $\mathcal{N}(0, 1)$. But if the portfolio contains one or more nonlinear instruments, there is no analytical expression for VAR. This is due to the fact that it is not possible to calculate the P&L distribution in the general (i.e., nonlinear) case. However, it is possible to calculate numerically the cumulative distribution of a quadratic portfolio $v(\omega) = \frac{1}{2}\omega^T G \omega + \delta^T \omega$, where G is a symmetric $M \times M$ matrix (cf. [Rouvinez], [Ruszczynski] and [Schaefer]). In practice, VAR for nonlinear portfolios is often determined with Monte–Carlo simulation, which is very time demanding.

1.3 Maximum Loss (ML)

The basis of Maximum Loss (ML) is an analysis of worst case scenarios, which is generally more tractable than the calculation of the P&L distribution. ML is defined as follows:

The maximum loss
- that can occur during a period of t days (holding period)
- if the risk factors are restricted to some trust region A_t with probability α
  (i.e., $\Pr(\omega \mid \omega \in A_t) = \alpha$).

At first sight, the statement of VAR is a much stronger than the one of ML: VAR states that the expected level of loss is in α percent of the cases above VAR and in $(1 - \alpha)$ cases below, while ML only says that the loss of α percent of the scenarios is not bigger than some level — no statement is made about the remaining $(1 - \alpha)$ percent of scenarios. This difference is caused by the fact that for determining VAR, the *distribution* of the P&L has to be known, whereas ML is directly defined in the space of the risk factors:

$$\begin{aligned} \text{ML} = \quad & \min \quad v(\omega) \\ & \text{s.t.} \quad \omega \in A_t; \quad \text{where } \Pr(A_t) = \alpha, \end{aligned} \tag{1.2}$$

In contrast to VAR, which depends on parameters t and α, ML has a supplementary degree of freedom called "trust region" A_t: any closed set of risk factors with probability α is a valid trust region. In the sequel, the trust region A_t will always be the ellipsoid defined by

$$A_t = \{\omega \mid \omega^T \Sigma_t^{-1} \omega \leq c_\alpha\}, \tag{1.3}$$

where c_α is the α–quantile of a chi–square distribution with M degrees of freedom (cf. Appendix B). Intuitively, ML is the worst case among all attainable (i.e., lying the trust region A_t) scenarios (cf. figure 1.1). Denote by $\mathrm{ML}(v)$ the solution of (1.2), then it is easy to verify that the following conditions are satisfied for every P&L function $v(\omega)$:

(i) $\mathrm{ML}(v) \geq \inf_{\omega \in \mathbb{R}^M} v(\omega)$,

(ii) $\mathrm{ML}(\mu + \lambda v) = \mu + \lambda \mathrm{ML}(v), \quad \forall \lambda \geq 0, \mu \leq 0$,

(iii) $\mathrm{ML}(\mathbb{1}_A v) \geq \mathrm{ML}(v), \quad \forall A \subset \mathbb{R}^M$,

(iv) $\mathrm{ML}(v_1 + v_2) \geq \mathrm{ML}(v_1) + \mathrm{ML}(v_2)$.

Risk measures which satisfy the conditions (i) — (iv) are called *coherent* (cf. [Artzner et al.]). The subadditivity property (iv) assures that the risk measure respects netting effects; in general, this condition is not satisfied by VAR measures.
In the case of a linear portfolio (i.e., $v(\omega) = \delta^T \omega$), ML can be calculated analytically (cf. [Studer]):

$$\mathrm{ML} \ = -\sqrt{c_\alpha} \sqrt{\delta^T \Sigma_t \delta}. \tag{1.4}$$

This expression is very similar to (1.1); they differ only in the factors z_α and $\sqrt{c_\alpha}$. The fact that $-z_\alpha \geq -\sqrt{c_\alpha}$ implies that VAR $\geq$ ML; [Studer] shows that this relation holds not only linear portfolios, but is true for every P&L function $v(\omega)$ and for every trust region which is a closed set (i.e., not only ellipsoids). But the ML approach not only gives the loss in the worst case, it also identifies the worst case scenario ω^*:

$$\omega^* = -\frac{\sqrt{c_\alpha}}{\sqrt{\delta^T \Sigma_t \delta}} \Sigma_t \delta. \tag{1.5}$$

As will be shown in chapters 3 and 5, an iterative calculation of ML for different confidence levels α gives insights into the portfolio which are more profound than a simple VAR value.

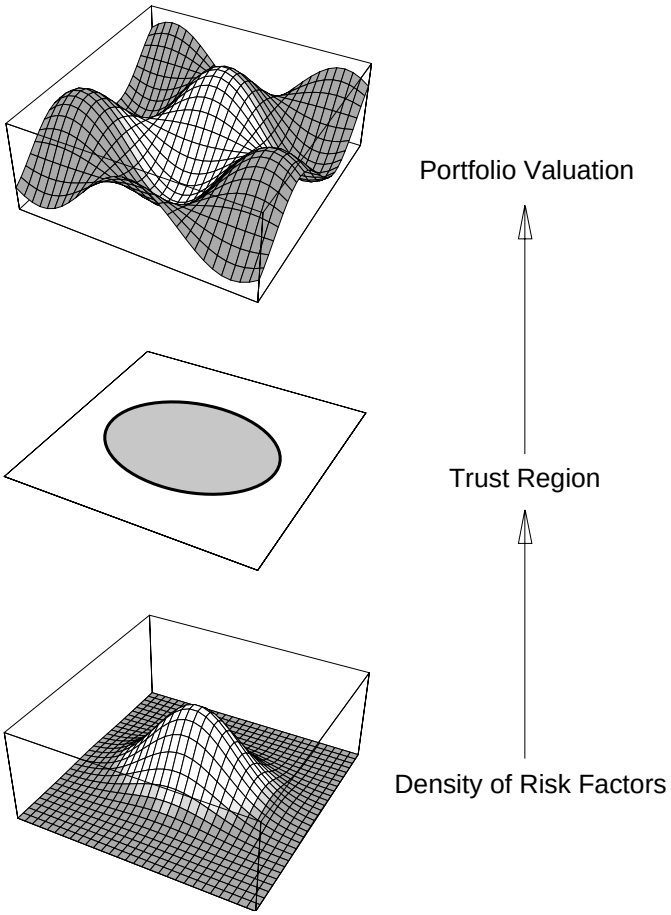

Figure 1.1: Modeling process of Maximum Loss

Chapter 2

Maximum Loss for Quadratic Risk Profiles

As is the case for VAR, there is no analytic expression of ML for quadratic P&L functions $v(\omega) = \frac{1}{2}\omega^T G\omega + g^T\omega$. However, it is possible to calculate *efficiently* the solution to the problem

$$\begin{aligned} \text{ML} \;=\; & \min \quad \frac{1}{2}\omega^T G\omega + g^T\omega \\ & \text{s.t.} \quad \omega^T \Sigma_t^{-1}\omega \le c_\alpha, \end{aligned} \tag{2.1}$$

where G is a symmetric $M \times M$ matrix, g an M–dimensional vector and c_α the α quantile of a chi–square distribution with M degrees of freedom (cf. Appendix B).

2.1 The Levenberg–Marquardt Algorithm

The solution to (2.1) can be calculated numerically with the Levenberg–Marquardt algorithm (cf. [Fletcher]). This algorithm is usually used in restricted step methods (nonlinear optimization). The algorithm cannot directly be applied to problem (2.1): it has first to be reformulated for a spherical trust region. There always exists a Cholesky decomposition

$$\Sigma_t = U^T U, \tag{2.2}$$

because Σ_t is a covariance matrix and therefore positive semidefinite. Writing

$$\omega = U^T\hat{\omega}, \tag{2.3}$$

we get an equivalent formulation to (2.1):

$$\begin{aligned} \text{ML} &= \min \quad \frac{1}{2}\hat{\omega}^T \hat{G}\hat{\omega} + \hat{g}^T\hat{\omega} \\ &\quad \text{s.t.} \quad \hat{\omega}^T\hat{\omega} \leq c_\alpha, \end{aligned} \tag{2.4}$$

where $\hat{g} = Ug$ and $\hat{G} = UGU^T$. Again, the objective function is quadratic, but this time the constraint represents a sphere, centered at the origin. The following theorem is the basis of the solution process:

Theorem 1
$\hat{\omega}$ is a *global* solution to (2.4) if and only if there exists $\nu \in \mathbb{R}$ such that

$$(\hat{G} + \nu\mathbb{1}) \text{ is positive semidefinite,} \tag{2.5}$$

and the following conditions hold:

$$(\hat{G} + \nu\mathbb{1})\hat{\omega} = -\hat{g} \tag{2.6}$$

$$\nu(c_\alpha - \hat{\omega}^T\hat{\omega}) = 0 \tag{2.7}$$

$$\nu \geq 0. \tag{2.8}$$

Moreover, if such a ν exists, then it is unique, and if $(\hat{G} + \nu\mathbb{1})$ is positive definite, then $\hat{\omega}$ is unique.

The proof of this theorem is given in [Fletcher]. The key idea behind the Levenberg–Marquardt algorithm is to make a one–dimensional search for ν. Suppose for a moment we would have an orthonormal basis of eigenvectors of $\hat{G}$: $\mathcal{B} = \{\hat{e}_1, \ldots, \hat{e}_M\}$ with corresponding eigenvalues $\lambda_1 \leq \ldots \leq \lambda_M$ ($\mathcal{B}$ exists because $\hat{G}$ is symmetric). We can express the vectors $\hat{g}$ and $\hat{\omega}$ in this new basis: $\hat{g} = \sum_{i=1}^M \alpha_i\hat{e}_i$ and $\hat{\omega} = \sum_{i=1}^M \beta_i\hat{e}_i$, where the coefficients β_i are the unknowns. Relation (2.6) implies that $\beta_i = -\frac{\alpha_i}{\lambda_i+\nu}, \quad i = 1, \ldots, M$. Therefore, $\hat{\omega}$ can be seen as a function of ν and

$$\|\hat{\omega}(\nu)\|^2 = \sum_{i=1}^M \left(\frac{\alpha_i}{\lambda_i + \nu}\right)^2. \tag{2.9}$$

This is a positive, decreasing function with $\lim_{\nu\to\infty} \|\hat{\omega}(\nu)\|^2 = 0$. Conditions (2.5) and (2.8) imply that $\nu \geq \max(-\lambda_1, 0)$. This situation is represented in figure 2.1. Depending on ν, two cases have to be distinguished:

1. $\nu = 0$: $\hat{G}$ is positive semidefinite by (2.5). From equation (2.6) we conclude that $\hat{\omega} = -\hat{G}^{-1}\hat{g}$ is the worst case we are looking for.

2. $\nu > 0$. Let $\nu_1 < \nu_2$ be two candidates for ν. Equation (2.6) implies that $\hat{\omega}_i = -(\hat{G} + \nu_i\mathbb{1})^{-1}\hat{g}, i = 1, 2$. From (2.9) we know that $\|\hat{\omega}_1\|^2 > \|\hat{\omega}_2\|^2$. Thus, we can apply a bisectionning method to find that value of ν which satisfies (2.7), i.e., $\|\hat{\omega}(\nu)\|^2 = c_\alpha$.

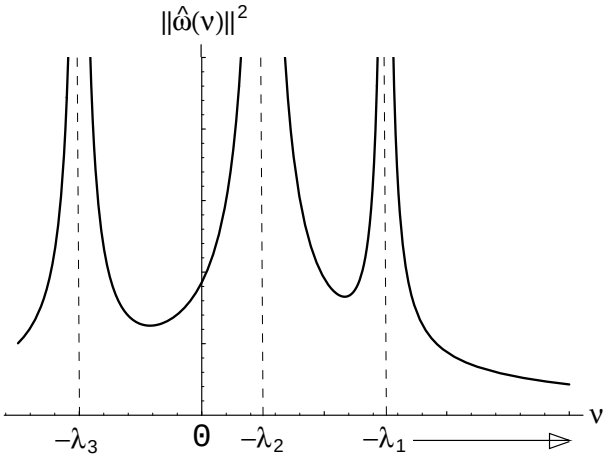

Figure 2.1: Analysis of $\|\hat{\omega}(\nu)\|^2$

Obviously, neither the eigenvectors $\hat{e}_i$ of $\hat{G}$ nor its eigenvalues λ_i are needed to implement such a bisectionning algorithm. In fact, it is possible to implement the algorithm such that it solves the problem in polynomial time:

Theorem 2
The number of arithmetic operations required for calculating the global solution to (2.4) is bounded by a polynomial in

- p: number of digits of the solution $\hat{\omega}$ requested
- L: bit–length of entry data ($\hat{G}, \hat{g}, c_\alpha$: rationals)
- M: dimension of problem.

Elements of the proof are given in [Vavasis]. It is important to note that this is the *only* nonconvex global optimization problem for which a polynomial time algorithm is known. Technical details of the implementation are described in [Finschi].

2.2 Sensitivity Analysis

The effects of small perturbations of the parameters (i.e., G, g, Σ_t, c_α) in problem (2.1) can be calculated analytically. These results will become helpful in chapter 3, when ML is calculated repetitively for slightly modified problems.
Let $u(\omega) = \omega^T \Sigma_t^{-1} \omega - c_\alpha$ be the constraint function of the optimization problem (2.1). Thus, the original optimization problem can alternatively be written as

$$\begin{aligned} \text{ML} &= \min \quad v(\omega) \\ &\text{s.t.} \quad u(\omega) \leq 0. \end{aligned} \tag{2.10}$$

The Lagrangian function of this problem is

$$\mathcal{L}(\omega, \mu) = v(\omega) + \mu u(\omega), \tag{2.11}$$

where μ is a real number. From optimization theory we know that if ω^* is a solution of problem (2.10), then there exists μ^* satisfying the Kuhn–Tucker equations

$$\begin{aligned} \nabla_\omega \mathcal{L}(\omega^*, \mu^*) &= 0 && (2.12) \\ \mu^* u(\omega^*) &= 0 && (2.13) \\ u(\omega^*) &\leq 0 && (2.14) \\ \mu^* &\geq 0, && (2.15) \end{aligned}$$

where ∇_ω denotes the gradient with respect to ω. Since we know ω^* from the Levenberg–Marquardt algorithm (cf. chapter 2.1), it is easy to find μ^*:

1. $u(\omega^*) \neq 0$: equation (2.13) implies $\mu^* = 0$.
2. $u(\omega^*) = 0$. To satisfy equation (2.12) and (2.15) we have to choose

$$\mu^* = \frac{\|G\omega^* + g\|}{2\|\Sigma_t^{-1}\omega^*\|}. \tag{2.16}$$

But μ^* can be interpreted as the marginal contribution (shadow price) of the restriction $u(\omega) \leq 0$. If this restriction is slightly perturbed by $\epsilon > 0$ to become $u(\omega) \leq \epsilon$, then $\omega^*(\epsilon)$ and $\mu^*(\epsilon)$ denote the solutions of the perturbed problem. Equation (2.13) implies that $v(\omega^*(\epsilon)) = \mathcal{L}(\omega^*(\epsilon), \mu^*(\epsilon))$ and therefore

$$\frac{\partial v(\omega^*(\epsilon))}{\partial \epsilon} = \frac{\partial \mathcal{L}(\omega^*(\epsilon), \mu^*(\epsilon))}{\partial \epsilon} = -\mu^*. \tag{2.17}$$

Writing the perturbed restriction as $\omega^T \Sigma_t^{-1} \omega \leq c_\alpha + \epsilon$, equation (2.17) enables us to calculate the effect of a small change Δc_α to ML:

$$\Delta v(\omega^*) = -\mu^* \Delta c_\alpha. \tag{2.18}$$

Of course, small modifications in the objective function $v(\omega)$ can be analyzed with ordinary calculus:

$$\Delta v(\omega^*) = \frac{1}{2} \omega^{*T} \Delta G \omega^* + \Delta g^T \omega^*. \tag{2.19}$$

Chapter 3

Maximum Loss Path

3.1 Expanding Trust Regions

A repetitive calculation of ML gives insights into a portfolio that go far beyond a simple worst case identification. The Levenberg–Marquardt algorithm described in chapter 2.1 allows to determine the value of the maximal loss as well as to identify the worst case scenario. If this calculation is repeated for several confidence levels α, a list of MLs and scenarios is constructed, Table 3.1 shows the results for a real–world portfolio.

α	ML	Risk Factor 1	Risk Factor 2	...	Risk Factor 7
⋮	⋮	⋮	⋮	...	⋮
93 %	-2970.595	-78.33	-38.48	...	-6.69
94 %	-3047.673	-79.71	-39.12	...	-6.80
95 %	-3137.119	-81.30	-39.87	...	-6.93
96 %	-3244.300	-83.17	-40.74	...	-7.07
⋮	⋮	⋮	⋮	...	⋮

Table 3.1: Repetitive calculation of ML for real–world portfolio

The geometric interpretation of this procedure is obvious: increasing α from 0 to some upper limit means expanding the trust region from a single point to the final ellipsoid (cf. figure 3.1).
Such a sequence of scenarios defines a path which starts at the actual value of the risk factors (cf. chapter 1.1) and follows the worst possible route (cf. figure 3.2). To obtain a path that is "smooth", it is necessary to solve the minimization problems many times for different levels of α. However, the results of the sensitivity analysis in chapter 2.2 give first order approximations which allow to reduce the number of minimizations by using some interpolation scheme.

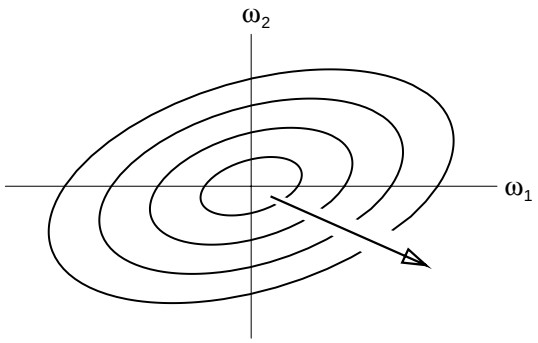

Figure 3.1: Expanding trust regions

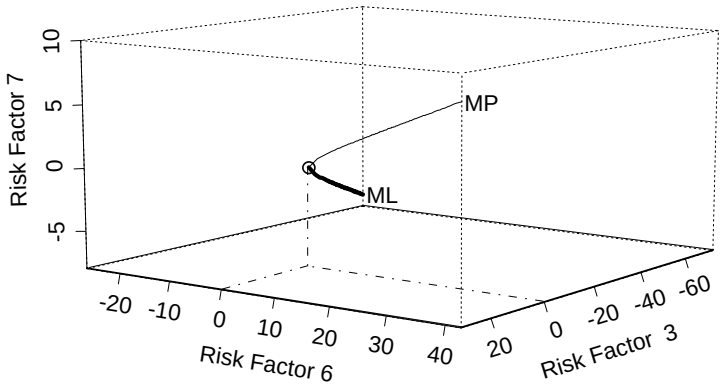

Figure 3.2: ML and MP paths of real–world portfolio

3.2 Maximum Loss On an Ellipsoid

If the matrix G in problem (2.1) is positive definite (i.e., all curvatures are positive), the optimization problem has a unique stationary point which is the *strict* global minimum. Consequently, the path will end at this point as soon as it is reached. This means that as soon as our path reaches this point, it will get stuck there. However, it could be interesting to see how the path would evolve if we restricted the minimization to the *surface* of the expanding trust region. This is equivalent to solve the problem

$$\begin{aligned} \text{ML} \;=\; & \min \quad \frac{1}{2}\omega^T G\omega + g^T\omega \\ & \text{s.t.} \quad \omega^T \Sigma_t^{-1}\omega = c_\alpha, \end{aligned} \tag{3.1}$$

where the inequality has become an equality. By the procedure described in

chapter 2.1, the ellipsoid is transformed to a sphere:

$$\begin{aligned} \text{ML} &= \min \quad \frac{1}{2}\hat{\omega}^T\hat{G}\hat{\omega} + \hat{g}^T\hat{\omega} \\ &\quad \text{s.t.} \quad \hat{\omega}^T\hat{\omega} = c_\alpha. \end{aligned} \tag{3.2}$$

Then the following theorem shows how this problem can be solved using the already known Levenberg–Marquardt algorithm (cf. chapter 2.1):

Theorem 3
For every $\kappa > \max(\lambda_1, 0)$, where λ_1 is the lowest eigenvalue of $\hat{G}$, the two problems

$$\begin{aligned} \min \quad & \frac{1}{2}\hat{\omega}^T(\hat{G} - \kappa\mathbb{1})\hat{\omega} + \hat{g}^T\hat{\omega} \\ \text{s.t.} \quad & \hat{\omega}^T\hat{\omega} = c_\alpha \end{aligned} \tag{3.3}$$

and

$$\begin{aligned} \min \quad & \frac{1}{2}\hat{\omega}^T(\hat{G} - \kappa\mathbb{1})\hat{\omega} + \hat{g}^T\hat{\omega} \\ \text{s.t.} \quad & \hat{\omega}^T\hat{\omega} \leq c_\alpha \end{aligned} \tag{3.4}$$

have an identical solution $\hat{\omega}^*$.

Proof. The matrix $(\hat{G} - \kappa\mathbb{1})$ is not positive semidefinite. To satisfy equation (2.5) of theorem 1, we must have $\nu \geq \kappa > 0$. Hence, $\|\hat{\omega}\|^2 = c_\alpha$ must hold to fulfill condition (2.7). Since ν is unique (cf. theorem 1), it follows that the two problems (3.3) and (3.4) have identical solutions. ∎

The fact that $\hat{\omega}^T\hat{\omega} = c_\alpha$ implies that $\frac{1}{2}\hat{\omega}^T(\hat{G}-\kappa\mathbb{1})\hat{\omega}+\hat{g}^T\hat{\omega} = \frac{1}{2}\hat{\omega}^T\hat{G}\hat{\omega}+\hat{g}^T\hat{\omega}-\frac{1}{2}\kappa c_\alpha$; therefore, $\hat{\omega}^*$ is also solution to (3.2). For applying this theorem, a lower bound of λ_1 must be known, which can be obtained by using Gershgorin discs.
Of course, the ideas developed so far can also be applied to the profit side of the portfolio. Inverting the sign of the objective function of (3.2) leads to the problem

$$\begin{aligned} -\text{MP} &= \min \quad -\frac{1}{2}\hat{\omega}^T\hat{G}\hat{\omega} - \hat{g}^T\hat{\omega} \\ &\quad \text{s.t.} \quad \hat{\omega}^T\hat{\omega} = c_\alpha, \end{aligned} \tag{3.5}$$

and permits to find the path of MP.

3.3 Expected Profit and Loss On an Ellipsoid

Are ML and MP sufficient to judge the quality of a portfolio? The answer is *no*: a study of nothing but the extremes gives no indication about what has to be expected in the "typical" case. Similar to ML and MP, we can also calculate the *expected profit and loss on* the ellipsoid: as before, the ellipsoid is first transformed into a sphere (cf. chapter 2.1), which changes the P&L function to $\hat{v}(\hat{\omega}) = \frac{1}{2}\hat{\omega}^T\hat{G}\hat{\omega} + \hat{g}^T\hat{\omega}$. Appendix C shows that the expected value (EV) on the sphere is

$$\mathrm{E}(\hat{v}(\hat{\omega}) \mid \hat{\omega}^T\hat{\omega} = c_\alpha) = \frac{c_\alpha}{2}\frac{\mathrm{Tr}(\hat{G})}{M}, \tag{3.6}$$

where $\mathrm{Tr}(\hat{G}) = \sum_{i=1}^{M} \hat{G}_{i,i}$. Although this quantity is easy to calculate, it gives important information about the portfolio: plotting the value of (3.6) for increasing α leads to the EV path. Figure 3.3 shows a portfolio whose MP is twice as much as ML — for every value of α. Nevertheless, the expected outcome is almost equal to ML. This emphasizes the fact that a comparison of ML and MP is not sufficient to judge a portfolio's qualities. However, the study of all three paths together allows founded judgments.

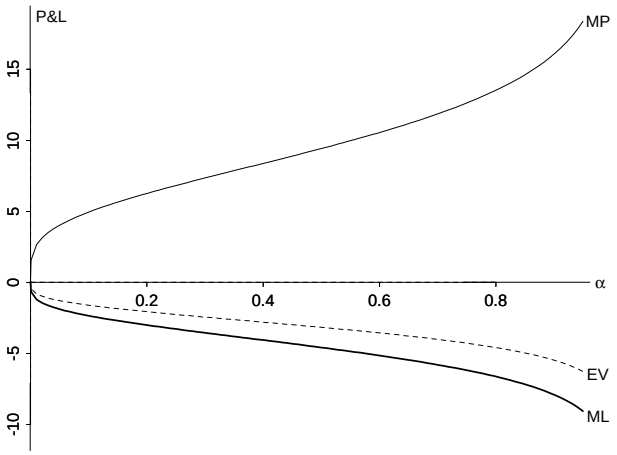

Figure 3.3: Portfolio characterization by ML, MP and EV paths

Chapter 4

Dynamic Approximation

In the previous chapters it has always been assumed that a *quadratic* P&L function $v(\omega)$ was given. Such a function can, for example, be obtained by using Taylor series expansions $v(\omega) = \delta^T\omega + \frac{1}{2}\omega^T, \omega + o(\|\omega\|^2)$. For many financial instruments, local sensitivities $(\delta, ,)$ can be calculated directly from the valuation models (cf. [Hull]). However, such local approximations are only valid for small changes in ω — in our model for short holding periods t or small confidence levels α.

4.1 Quadratic Approximation

For large moves of the risk factors ω, local approximations can lead to large errors. In practice, so–called risk profiles (cf. figure 4.1) are used to analyse the structure of P&L functions $v(\omega)$: a set of scenarios $\mathcal{S} = \{\omega^{(1)}, \ldots, \omega^{(n)}\}$ is defined and the portfolio is fully repriced for each of them separately; the resulting P&Ls are denoted $\mathcal{P} = \{\xi^{(1)}, \ldots, \xi^{(n)}\}$.

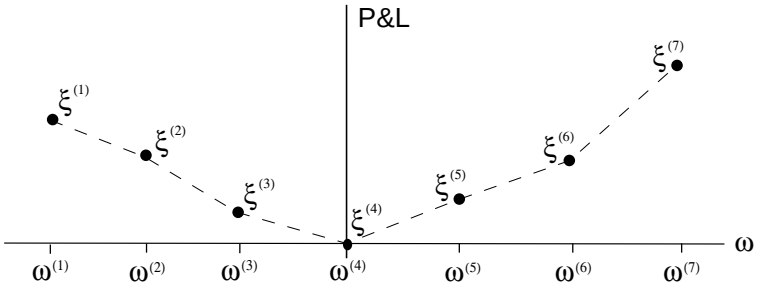

Figure 4.1: Example of one–dimensional risk profile

From the sets $\mathcal{P}$ and $\mathcal{S}$, a quadratic approximation $v(\omega) = \frac{1}{2}\omega^T G\omega + g^T\omega + c$ can be constructed using the method of least squares:

$$\begin{aligned} \min \quad & \sum_{i=1}^{n} \left(\frac{1}{2} \omega^{(i)T} G \omega^{(i)} + g^T \omega^{(i)} + c - \xi^{(i)} \right)^2 \\ \text{s.t.} \quad & G \in \mathbb{R}^{M \times M}, \text{ symmetric} \\ & g \in \mathbb{R}^M \\ & c \in R. \end{aligned} \tag{4.1}$$

The unknowns G, g and c can be determined by solving a set of linear equations; more details are given in Appendix D. Note that the newly introduced constant c can simply be added to the results of the previous chapters.

4.2 Dynamic Weighting

The least squares method of chapter 4.1 leads to quadratic approximations with best fit to the *entire* scenario set $\mathcal{S}$. However, if we seek ML *on* the surface of a specific ellipsoid (i.e., for a fixed level α), we can utilize better approximations by weighting the scenarios adequately: the scenarios $\omega^{(i)}$ lying close to the α–ellipsoid (i.e., $\omega^{(i)T} \Sigma_t^{-1} \omega^{(i)} \approx c_\alpha$) get a higher weight than those lying far away. Hence, for scenario $\omega^{(i)}$ the weight $\vartheta_i(\alpha)$ is set to:

$$\vartheta_i(\alpha) = \frac{1}{1 + \beta \mid \omega^{(i)T} \Sigma_t^{-1} \omega^{(i)} - c_\alpha \mid^\gamma}, \tag{4.2}$$

where $\beta, \gamma > 0$ are parameters which control the smoothness with respect to α. If $\Theta(\alpha)$ denotes the diagonal weighting matrix (i.e., $\Theta_{i,i}(\alpha) = \vartheta_i(\alpha)$; $\Theta_{i,j}(\alpha) = 0$ if $i \neq j$), then the normal equation (D.2) becomes

$$S^T \Theta(\alpha) S x(\alpha) = S^T \Theta(\alpha) \xi, \tag{4.3}$$

where the solution $x(\alpha)$ defines the parameters $G(\alpha), g(\alpha)$ and $c(\alpha)$. Thus, for every value of α, a new P&L function $v(\omega, \alpha) = \frac{1}{2} \omega^T G(\alpha) \omega + g(\alpha)^T \omega + c(\alpha)$ is defined; it is quadratic for α fixed. Hence, the paths of ML, MP and EV can be calculated with the methods described in chapters 3.2 and 3.3. This way, it is possible to analyse portfolios, which are not necessarily quadratic, by using a family of quadratic approximations (cf. figure 4.2). Note that this method produces *exact* results for quadratic functions.

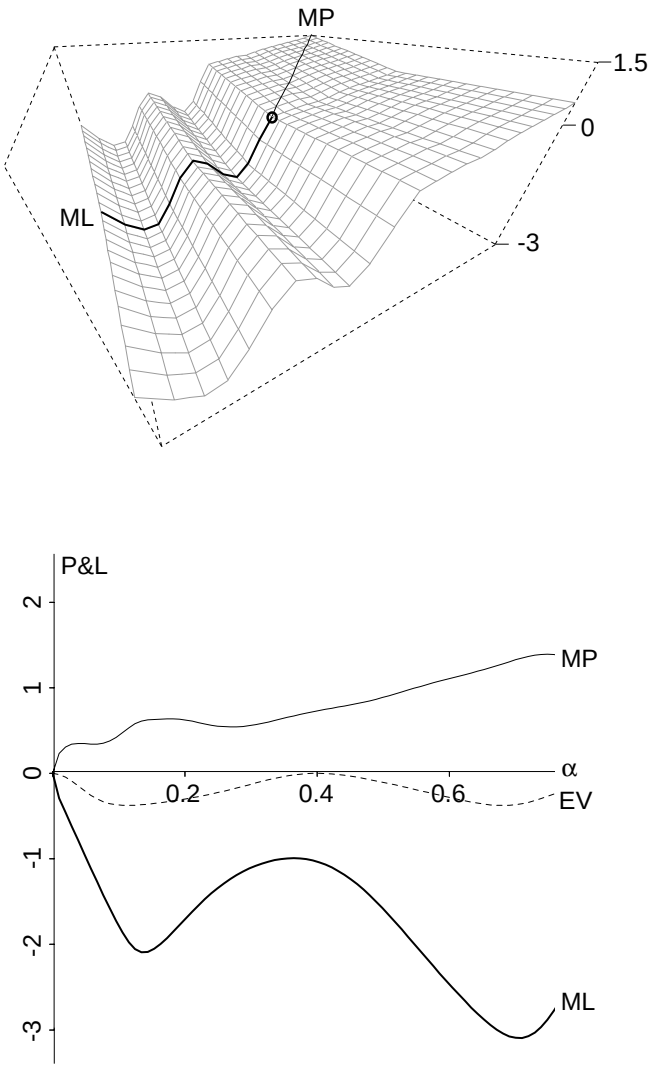

Figure 4.2: Paths for a non–quadratic P&L function

Chapter 5

Maximum Loss Distribution

Using the information contained in the paths, it is possible to determine the distribution of ML and MP. This enables us to answer questions like "With which probability will ML exceed the level of -2400?" or "What is the probability that MP is less than 1300?".

5.1 Shell Model

In this chapter we show how a lower bound of the ML distribution can be obtained using a shell model: the key idea is to dissect the entire space $\mathbb{R}^M$ into a finite, disjoint set of shells $\{S_i\}$. Assume that for the construction of the ML path, the worst cases for $0\% = \alpha^{(0)} < \alpha^{(1)} < \ldots < \alpha^{(N)} = 100\%$ have been calculated. Then, the most natural choice is to take $S_i = \{\omega \mid c_{\alpha^{(i-1)}} \leq \omega^T \Sigma_t^{-1} \omega < c_{\alpha^{(i)}}\}$. Thus, in the transformed problem (2.4), $\mathbb{R}^M$ can be interpreted as a union of thin spherical mantles (cf. figure 5.1).

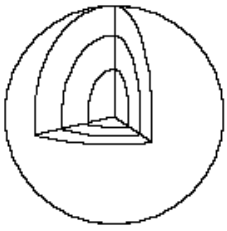

Figure 5.1: Shell model of $\mathbb{R}^3$

Using the natural discretisation implies that shell S_i covers a probability $p_i = \alpha^{(i)} - \alpha^{(i-1)}$ of the space. Moreover, to every shell S_i, the value $m_i = \min\{\mathrm{ML}_{\alpha^{(i-1)}}, \mathrm{ML}_{\alpha^{(i)}}\}$ is assigned; if the discretisation is fine enough, m_i approximates the worst

case in shell S_i. Hence, the pairs (m_i, p_i) define a discrete distribution function of ML.
As a result, it is possible to construct distribution functions of ML and MP for every portfolio. Consequently, the "profit and loss potential" of a portfolio can be quantified: the percentage in which ML dominates MP (i.e., $-\text{ML} > \text{MP}$) is easily found by comparing the two monotonously increasing sequences.

5.2 Boundaries for the Profit and Loss Distribution

The distributions of ML and MP are lower and upper bounds of the distribution of P&L: in the construction of chapter 5.1, all possible P&L values in shell S_i were replaced by the maximal loss m_i in that shell. Thus, $\Pr(v(\omega) \leq y \mid \omega \in S_i) \leq \Pr(m_i \leq y)$ for every shell S_i and, consequently, $\Pr(v(\omega) \leq y) \leq \Pr(\text{ML} \leq y)$. This means that at every level y, the distribution function of ML lies above the distribution function of P&L. Figure 5.2 shows the results for the real–world portfolio; the distribution of P&L was determined with Monte–Carlo simulation.

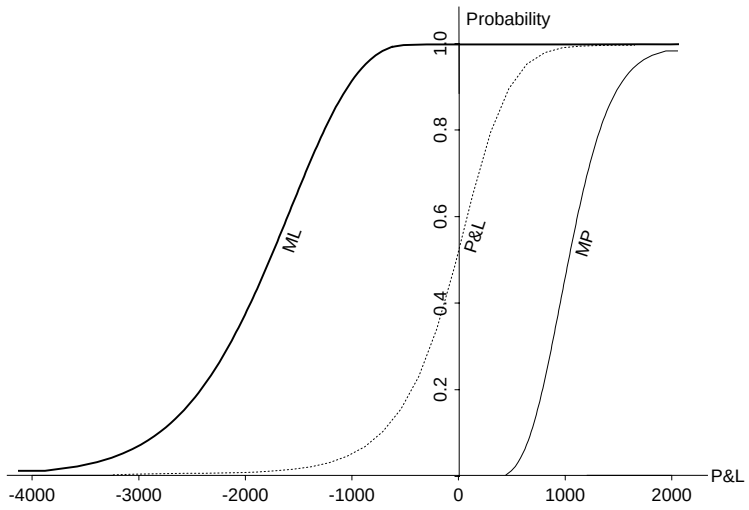

Figure 5.2: ML/MP distributions and true distribution of P&L

In other words, the ML distribution is a lower bound of the P&L distribution; in the same way, the P&L distribution is bounded from above by the MP distribution. Consequently, ML and MP distributions lead to lower and upper bounds of the risk measure "Value–At–Risk" (cf. chapter 1.2), since VAR is by definition a quantile of the P&L distribution. Appendix E presents a different approach which is based on quadratic underestimators and generates bounds that are generally tighter.

Chapter 6

Conclusion

The key idea of Maximum Loss is to determine the worst possible outcome. Mathematically, this means to search for the global minimum of the profit and loss function over a so–called trust region (restricted area of the scenario space). For quadratic P&L functions, ML can be calculated efficiently (i.e., in polynomial time) by use of the Levenberg–Marquardt algorithm.
Repetitive calculation of ML for a set of growing trust regions leads to the Maximum Loss Path. This concept can successfully be applied to non–quadratic portfolios: the key idea is to restrict the analysis to the surface of the trust region and to deal with a new quadratic approximation of the P&L function for each ellipsoid. This way, the worst route can be tracked for almost any kind of portfolio. Similarly, the Maximum Profit Path as well as the path of expected P&L can be constructed. The comparison of the three paths permits to judge the qualities of a portfolio.
The Maximum Loss Distribution can be determined by a shell model. Contrasting the distributions of ML and MP allows to determine to what extent the portfolio's loss potential exceeds the profit potential. Also, ML and MP distributions are lower and upper bounds of the P&L distribution and, therefore, lead to lower and upper bounds of VAR.

Appendix A

Stochastic Models of Market Rate Innovations

The modeling of market rates effects directly the risk measurement model: Market rates r (e.g. commodity prices, equity indices, FX rates) are usually supposed to follow geometric Brownian motion (cf. [Hull]):

$$\frac{dr}{r} = \mu dt + \sigma dz, \tag{A.1}$$

where μ is the drift factor and σ the volatility, t is time and dz is a Wiener process (i.e., $dz \sim \mathcal{N}(0, dt)$). The application of Itô's lemma to the function $\log r$ leads to

$$\log\left(\frac{r+dr}{r}\right) = \left(\mu - \frac{\sigma^2}{2}\right) dt + \sigma dz. \tag{A.2}$$

If the constant drift rate $(\mu - \frac{\sigma^2}{2})dt$ is eliminated, it follows that

$$\log\left(\frac{r+dr}{r}\right) \sim \mathcal{N}(0, \sigma^2 dt), \tag{A.3}$$

which means that the driftless returns are lognormally distributed. A first order Taylor approximation results in

$$\begin{aligned} \log\left(\frac{r+dr}{r}\right) &= \log(r+dr) - \log(r) \\ &\approx \left[\log(r) + dr\frac{1}{r}\right] - \log(r) \\ &= \frac{dr}{r}, \end{aligned} \tag{A.4}$$

which implies that the driftless relative returns are approximately normally distributed:

$$\frac{dr}{r} \sim \mathcal{N}(0, \sigma^2 dt). \tag{A.5}$$

If the value V of a financial instrument depends linearly on r (i.e., $V = Pr$, where P is the position of the instrument), then

$$dV = Pdr = V\frac{dr}{r}. \tag{A.6}$$

Hence, the change in value dV is normally distributed with mean 0 and standard deviation $V\sigma\sqrt{dt}$.
For interest rates, however, the situation is somewhat different: The present value $\tilde{V}$ of a zero coupon bond with maturity n is $\tilde{V} = \tilde{P}(1 + r_n)^{-n}$, where $\tilde{P}$ is the face value of the bond and r_n is the zero coupon rate for a n year investment. It follows that

$$\begin{aligned} d\tilde{V} &= -\frac{n}{1 + r_n}\tilde{P}(1 + r_n)^{-n}dr_n \\ &= -D\tilde{V}dr_n \qquad (A.7) \\ &= -D\tilde{V}r_n\frac{dr_n}{r_n}. \qquad (A.8) \end{aligned}$$

The term $D = \frac{n}{1+r_n}$ is called "modified duration" (cf. [Fabozzi]). If the interest rate r_n is assumed to follow *geometric Brownian motion*, equation (A.8) implies that $d\tilde{V}$ becomes normally distributed:

$$d\tilde{V} \sim \mathcal{N}(0, D^2\tilde{V}^2 r_n^2 \sigma^2 dt). \tag{A.9}$$

However, if r_n is assumed to follow *arithmetic Brownian motion* (i.e., $dr_n = \mu dt + \sigma dz$), then equation (A.7) implies that

$$d\tilde{V} \sim \mathcal{N}(0, D^2\tilde{V}^2 \sigma^2 dt). \tag{A.10}$$

Note that in the case of geometric Brownian motion (A.9), the variance is r_n^2 times the variance of (A.10).

Appendix B

Modeling Trust Regions

Trust regions are primarily defined with the density $f_t(\omega)$ of the risk factors. For centered, multinormally distributed risk factors $\omega_1, \ldots, \omega_M$, the density function is

$$f(\omega) = \frac{1}{(2\pi)^{M/2}\sqrt{\det \Sigma_t}} \exp\left(-\frac{1}{2}\omega^T \Sigma_t^{-1} \omega\right), \tag{B.1}$$

where Σ_t is the covariance matrix. is to find a trust region A_t which covers a probability of α and includes the scenario $\omega = 0$. One possible choice is to search a constant c such that $\Pr\{\omega \mid f_t(\omega) \geq c\} = \alpha$. This makes sense since $f_t(\omega)$ attains its maximum at $\omega = 0$ and leads to the trust region of minimal volume. By eliminating the constants, the problem is reduced to the following: Find c_α such that

$$\Pr(\omega \mid \omega^T \Sigma_t^{-1} \omega \leq c_\alpha) = \alpha. \tag{B.2}$$

A Cholesky decomposition $\Sigma_t = U^T U$ leads to $\omega^T \Sigma_t^{-1} \omega = \omega^T U^{-1} U^{-T} \omega = (U^{-T}\omega)^T (U^{-T}\omega)$. But

$$\begin{aligned} \mathrm{Var}(U^{-T}\omega) &= \mathrm{E}\left[(U^{-T}\omega)(U^{-T}\omega)^T\right] \\ &= U^{-T}\mathrm{Var}(\omega)U^{-1} \\ &= U^{-T}U^T U U^{-1} \\ &= \mathbb{1}. \end{aligned} \tag{B.3}$$

Hence,

$$(U^{-T}\omega) \sim \mathcal{N}(0, \mathbb{1}) \tag{B.4}$$

and

$$\omega^T \Sigma_t^{-1} \omega = \sum_{i=1}^{M} X_i^2, \tag{B.5}$$

where $X_i, i = 1, \ldots, M$, are independent, standard normal variates. Thus, $\sum_{i=1}^{M} X_i^2$ is chi–squared distributed with M degrees of freedom. Consequently, the trust region A_t is defined by:

$$A_t = \{\omega \mid \omega^T \Sigma_t^{-1} \omega \leq c_\alpha\}, \tag{B.6}$$

where c_α is the α quantile of a χ^2_M distribution. This equation represents an ellipsoid centered at the origin.

Appendix C

Conditional Expectation of Profit and Loss On an Ellipsoid

Transforming the ellipsoid to a sphere (cf. chapter 2.1) leads to the P&L function $\hat{v}(\hat{\omega}) = \frac{1}{2}\hat{\omega}^T\hat{G}\hat{\omega} + g^T\hat{\omega}$. The quantity of interest is the conditional expectation $\mathrm{E}(\hat{v}(\hat{\omega}) \mid \hat{\omega}^T\hat{\omega} = c_\alpha)$, where $\hat{\omega} \sim \mathcal{N}(0, \mathbb{1})$ by equations (2.3) and (B.4). Therefore, the density of $\hat{\omega}$ is

$$f(\hat{\omega}) = (2\pi)^{-M/2} \exp(-\frac{1}{2}\sum_{i=1}^{M} \hat{\omega}_i^2). \tag{C.1}$$

For a sphere with radius $\sqrt{c_\alpha}$ we have $\sum_{i=1}^{M} \hat{\omega}_i^2 = c_\alpha$; hence, $f(\hat{\omega} \mid \hat{\omega}^T\hat{\omega} = c_\alpha)$ is constant and does not depend on $\hat{\omega}$. Thus, for the i^{th} component of $\hat{\omega}$ one gets:

$$\begin{aligned}
\mathrm{E}(\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha) &= \frac{1}{2}\mathrm{E}(\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i \geq 0) \\
&\quad + \frac{1}{2}\mathrm{E}(\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i < 0) \\
&= \frac{1}{2}\mathrm{E}(\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i \geq 0) \\
&\quad + \frac{1}{2}\mathrm{E}(-\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i > 0) \\
&= 0.
\end{aligned} \tag{C.2}$$

Similarly, the expectation of the mixed terms $\hat{\omega}_i\hat{\omega}_j, i \neq j$, is

$$\begin{aligned}
\mathrm{E}(\hat{\omega}_i\hat{\omega}_j \mid \hat{\omega}^T\hat{\omega} = c_\alpha) &= \frac{1}{2}\mathrm{E}(\hat{\omega}_i\hat{\omega}_j \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i \geq 0) \\
&\quad + \frac{1}{2}\mathrm{E}(\hat{\omega}_i\hat{\omega}_j \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i < 0)
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2}\mathrm{E}(\hat{\omega}_i\hat{\omega}_j \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i \geq 0) \\
&\quad + \frac{1}{2}\mathrm{E}(-\hat{\omega}_i\hat{\omega}_j \mid \hat{\omega}^T\hat{\omega} = c_\alpha, \hat{\omega}_i > 0) \\
&= 0. \qquad \text{(C.3)}
\end{aligned}
$$

For the quadratic terms, the reasoning is somewhat different. Using the fact that $\hat{\omega}$ lies on a sphere (i.e., $\sum_{j=1}^{M} \hat{\omega}_j^2 = c_\alpha$) leads to

$$
\begin{aligned}
\mathrm{E}(\hat{\omega}_i^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha) &= \mathrm{E}(c_\alpha - \sum_{j\neq i} \hat{\omega}_j^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha) \\
&= \mathrm{E}(c_\alpha \mid \hat{\omega}^T\hat{\omega} = c_\alpha) - \sum_{j\neq i} \mathrm{E}(\hat{\omega}_j^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha) \\
&= c_\alpha - (M-1)\mathrm{E}(\hat{\omega}_i^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha), \qquad \text{(C.4)}
\end{aligned}
$$

hence,

$$
\mathrm{E}(\hat{\omega}_i^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha) = \frac{c_\alpha}{M}. \qquad \text{(C.5)}
$$

Using the results of (C.2), (C.3) and (C.5) simultaneously, the conditional expectation becomes

$$
\begin{aligned}
\mathrm{E}(\hat{v}(\hat{\omega}) \mid \hat{\omega}^T\hat{\omega} = c_\alpha) &= \mathrm{E}(\frac{1}{2}\sum_{i,j=1}^{M} (\hat{G})_{i,j}\hat{\omega}_i\hat{\omega}_j + \sum_{i=1}^{M} \hat{g}_i\hat{\omega}_i \mid \hat{\omega}^T\hat{\omega} = c_\alpha) \\
&= \mathrm{E}(\frac{1}{2}\sum_{i=1}^{M} (\hat{G})_{i,i}\hat{\omega}_i^2 \mid \hat{\omega}^T\hat{\omega} = c_\alpha) \\
&= \frac{1}{2}\sum_{i=1}^{M} (\hat{G})_{i,i}\frac{c_\alpha}{M} \\
&= \frac{c_\alpha}{2}\frac{\mathrm{Tr}(\hat{G})}{M}. \qquad \text{(C.6)}
\end{aligned}
$$

In summary, the expected P&L *on* a sphere only depends on the trace of $\hat{G}$, the radius of the sphere $\sqrt{c_\alpha}$ and the dimension M of the problem. Since $\frac{\mathrm{Tr}(\hat{G})}{M} = \frac{\sum_{i=1}^{M}\lambda_i}{M} = \bar{\lambda}$ (average curvature), we can give an intuitive interpretation of this result: the expectation on a sphere with radius $\sqrt{c_\alpha}$ is $\frac{c_\alpha}{2}$ times the mean curvature of the quadratic function.

Appendix D

Least Squares for Quadratic Approximation

Given a set of scenarios $\mathcal{S} = \{\omega^{(1)}, \ldots, \omega^{(n)}\}$ and their respective P&L values $\mathcal{P} = \{\xi^{(1)}, \ldots, \xi^{(n)}\}$, a quadratic function $v(\omega) = \frac{1}{2}\omega^T G\omega + g^T\omega + c$ has to be found, which best fits the observations. The least squares formulation of this problem is

$$\min_{G,g,c} \sum_{i=1}^{n} \left(\frac{1}{2}\omega^{(i)T} G\omega^{(i)} + g^T\omega^{(i)} + c - \xi^{(i)} \right)^2, \tag{D.1}$$

where $G \in \mathbb{R}^{M\times M}$, symmetric; $g \in \mathbb{R}^M$ and $c \in \mathbb{R}$ are the unknowns. The solution is obtained by solving the normal equation

$$S^T S x = S^T \xi, \tag{D.2}$$

where S is the scenario matrix

$$S = \begin{array}{c} \\ \omega^{(1)} \\ \omega^{(2)} \\ \vdots \\ \omega^{(n)} \end{array} \begin{array}{c} \begin{array}{ccccccccccc} \frac{1}{2}\omega_1^2 & \omega_1\omega_2 & \ldots & \omega_1\omega_M & \frac{1}{2}\omega_2^2 & \omega_2\omega_3 & \ldots & \frac{1}{2}\omega_M^2 & \omega_1 & \ldots & \omega_M \quad 1 \end{array} \\ \begin{pmatrix} \ldots & & & & & & & & & \ldots & 1 \\ \ldots & & & & & & & & & \ldots & 1 \\ & & & & \ddots & & & & & & \vdots \\ \ldots & & & & & & & & & \ldots & 1 \end{pmatrix} \end{array} \tag{D.3}$$

which is a $n \times (\frac{M(M+3)}{2} + 1)$ matrix; row i contains the various products of the components of scenario $\omega^{(i)}$. The vector

$$\xi = \left(\xi^{(1)}, \ldots, \xi^{(n)}\right), \tag{D.4}$$

holds the P&Ls of each scenario. Finally, x is the vector of the unknowns:

$$x = \left(G_{1,1}; G_{1,2}; \ldots; G_{1,M}; G_{2,2}; G_{2,3}; \ldots G_{M,M}; g_1; \ldots g_M; c\right), \tag{D.5}$$

it has $(\frac{M(M+3)}{2}+1)$ elements and contains all elements of G, g and c. If one has knowledge that the portfolio is not sensitive to crosses of the risk factors (i.e., $G_{i,j}=0$ if $i \neq j$), the vector x has only $(2M+1)$ unknowns and the linear system is considerably smaller.

Appendix E

Quadratic Estimators for Boundaries of Profit and Loss Distribution

[Rouvinez] describes a method for calculating numerically the cumulative distribution of a quadratic P&L function. This technique can be used to determine lower bounds of the P&L distribution: given a set of scenarios $\mathcal{S} = \{\omega^{(1)}, \ldots, \omega^{(n)}\}$ with P&Ls $\mathcal{P} = \{\xi^{(1)}, \ldots, \xi^{(n)}\}$, the weights $\psi_i = \exp(-\frac{1}{2}\omega^{(i)T}\Sigma_t^{-1}\omega^{(i)})$, $i = 1, \ldots, n$, are defined, which are proportional to the density of the risk factors. Then, a quadratic underestimator $\underline{v}(\omega) = \frac{1}{2}\omega^T \underline{G}\omega + \underline{g}^T\omega + \underline{c}$ is constructed, which approximates the given scenarios best:

$$\begin{aligned} \min \quad & \sum_{i=1}^{n} \psi_i \mid \underline{v}(\omega^{(i)}) - \xi^{(i)} \mid \\ \text{s.t.} \quad & \underline{v}(\omega^{(i)}) \leq \xi^{(i)}, \quad i = 1, \ldots, n. \end{aligned} \tag{E.1}$$

Setting $\psi = (\psi_1, \ldots, \psi_n)$ and using the notation of Appendix D, the problem (E.1) can be formulated as a linear program:

$$\begin{aligned} \min \quad & d^T x \\ \text{s.t.} \quad & Ax \leq \xi, \end{aligned} \tag{E.2}$$

where $d = A^T\psi$. The solution x defines the unknowns $\underline{G}, \underline{g}$ and $\underline{c}$ and, consequently, the quadratic underestimator $\underline{v}(\omega)$. Hence, the distribution of $\underline{v}(\omega)$ is a lower bound of the true distribution of P&L. Similarly, an upper bound can be calculated by replacing the constraint with $Ax \geq \xi$.
In general, these bounds are tighter than the ML/MP distributions (cf. chapter 5.2). This is due to the fact that ML of a shell counts for *all* M dimensions, whereas quadratic underestimators consider each dimension in the same way.

Bibliography

[Artzner et al.] Ph. Artzner, F. Delbaen, J.–M. Eber, D. Heath. A Characterization of Measures of Risk. Working paper, February 1996.

[Beckström and Campbell] R. Beckström and A. Campbell. An Introduction to VAR. C·ATS Software Inc., Palo Alto, 1995.

[RiskMetrics] J.P. Morgan & Co. RiskMetricsTM - Technical Document. New York, May 1995.

[Fabozzi] F.J. Fabozzi. Bond Markets, Analysis and Strategies, Second Edition. Prentice–Hall, Englewood Cliffs, 1993.

[Finschi] L. Finschi. An Implementation of the Levenberg–Marquardt Algorithm. Technical Report, Zurich, 1996.

[Fletcher] R. Fletcher. Practical Methods of Optimization, Second Edition. John Wiley & Sons, Chichester, 1987.

[Hull] J.C. Hull. Options, Futures and other Derivative Securities. Prentice-Hall, Englewood Cliffs, 1993.

[Rouvinez] Ch. Rouvinez. Mastering Delta Gamma. Risk Magazine, *to appear.*

[Ruszczynski] A. Ruszczynski. Risk estimation for non–central χ^2 outcomes. Technical note, Zurich, June 1996.

[Schaefer] S. M. Schaefer. Risk Assessment and VAR Models. Documentation of the RI$K '95 Workshop on "Implementing The Latest Practical Advances In Making Risk Assessment And VAR Models Work: Effective Risk Aggregation And Compliance With The Basle Proposals". Paris, November, 1995.

[Studer] G. Studer. Value At Risk and Maximum Loss Optimization. RiskLab: Technical Report, Zurich, December 1995.

[Vavasis] S.A. Vavasis. Nonlinear Optimization, Complexity Issues. Oxford University Press, New York, 1991.