Regularization methods for inverse problems

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

José Alberto Orozco Rodríguez

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Prof. Fadil Santosa, Adviser

March, 2011

Acknowledgements

I feel extremely fortunate to have met Prof. Fadil Santosa. I first was introduced to many interesting and relevant applications in industry in a course that I took with him. Since then, he has supported me in so many different ways. His encouragement was decisive in my successful attempts at the written examinations at a time when, otherwise, I would have opted for a different path. I am very glad I stayed on this path, and that I have benefited from the many exciting opportunities that derive from Prof. Santosa's great creativity, whether in the form of a summer internship, a rewarding dissertation, or an internship leading to the job I wanted. My deepest gratitude to Prof. Santosa for his great generosity and invaluable mentorship.

I would also like to thank Prof. Gilad Lerman for all of his advise while I was in his research group. Had my dissertation subject been better aligned with his interests, I would have loved to have him as my co-adviser. Thank you to many members of the staff of the Math department for their support, in particular, to Prof. Paul Garrett and the admissions committee. I am also very thankful to my colleagues, especially, to Mike Aschenbeck and Ivan Merev for their help and friendship. Very special thanks to Fanhuan Zhou, whose solid work I am now building on.

I would like to thank my wife, Marie, who has made many sacrifices in order for us to be together. Her efforts have been superhuman, taking care of her career and our two little daughters while I was at the last stages of the dissertation. Thank you to Valeria and Adriana, too, for all of the joy and happiness they bring to our lives.

Abstract

Many applications in industry and science require the solution of an inverse problem. To obtain a stable estimate of the solution of such problems, it is often necessary to implement a regularization strategy. In the first part of the present work, a multiplicative regularization strategy is analyzed and compared with Tikhonov regularization. In the second part, an inverse problem that arises in financial mathematics is analyzed and its solution is regularized.

Tikhonov regularization for the solution of discrete ill-posed problems is well documented in the literature. The L-curve criterion is one of a few techniques that are preferred for the selection of the Tikhonov parameter. A more recent regularization approach less well known is a multiplicative regularization strategy, which unlike Tikhonov regularization, does not require the selection of a parameter. We analyze a multiplicative regularization strategy for the solution of discrete ill-posed problems by comparing it with Tikhonov regularization aided with the L-curve criterion.

We then proceed to analyze the stability of a method for estimating the risk-neutral density (RND) for the price of an asset from option prices. RND estimation is an inverse problem. The method analyzed first applies the principle of maximum entropy, where the maximum entropy solution (MES) corresponds to the estimated RND. Next, it provides an effective characterization of the constraint qualification (CQ) under which the MES can be computed by solving the dual problem, where an explicit function in finitely many variables is minimized. In our analysis, we show that the MES is stable under parameter perturbation, but the parameters are unstable under data perturbation. When noisy data are used, we show how to project the data so that the CQ is satisfied and the method can be used. To stabilize the method, we use Tikhonov regularization and choose the penalty parameter via the L-curve method. We demonstrate with numerical examples that the method becomes then much more stable to perturbation in data. Accordingly, we perform a convergence analysis of the regularized solution.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

A typical situation that occurs in many applications in science and industry is that the properties of interest of a system cannot be measured directly. It is then necessary to obtain estimates from some other data that are only indirectly related to the desired properties. In all applications, the data are measured quantities with uncertainties that cannot be controlled. Hence, an estimate will only be reliable if it is stable under data perturbations, which is often not the case. A problem where small noise in the data may lead to enormous errors in the estimate is called ill-posed.

When an application calls for the solution of an ill-posed problem, it is possible, nonetheless, to obtain a stable estimate, for example, by solving instead a judicious approximation to the original problem. Methods that estimate a stable solution of an ill-posed problem are called regularization methods.

A well established regularization method nowadays, which has been extensively analyzed, is the so-called Tikhonov regularization method, introduced independently by Tikhonov [1] and Phillips [2] almost half a century ago. As with several other regularization methods, Tikhonov regularization introduces a parameter which is not part of the original problem. The theory underlying this method is developed under the assumption that the regularization parameter that has been introduced can be selected appropriately. Therefore, in order to estimate a stable solution of an ill-posed problem using Tikhonov regularization, one has to first implement a parameter selection method.

A method that has succeeded in numerous applications to select an acceptable parameter is the L-curve method [3], which has been used for about two decades. This

and other methods such as Generalized Cross Validation or the discrepancy principle have been widely analyzed.

Parameter selection methods are usually based on the minimization of certain functional or the location of a root of certain equation. This typically requires iterative methods that might be computationally intensive. As a response to this possibly undesirable feature, other regularization methods that do not require the selection of a parameter have been proposed. Such is the case of a multiplicative regularization strategy that has been used in perhaps dozens of applications throughout the last decade (see for example [4, 5, 6, 7, 8, 9]).

In the first part of this thesis, we analyze a multiplicative regularization strategy in the setting of the so-called discrete ill-posed problems. Here, we want to solve the matrix problem $Ax = b$ or $\min \|Ax - b\|_2$ which usually arises after discretizing a problem. A common source of discrete ill-posed problems is the discretization of Fredholm integral equations of the first kind which model inverse problems that arise in a great variety of important applications in science and engineering.

In addition to the theoretical aspects covered in the first part of the thesis, we solve a simple numerical example previously considered in the literature with both the multiplicative regularization strategy and Tikhonov regularization aided with the L-curve method. The numerical example presented is based on the problem of image deblurring. However, its formulation corresponds to the general formulation of discrete ill-posed problems. Hence, the numerical results in conjunction with the theoretical results presented allow us to reach conclusions about the multiplicative regularization strategy in the setting of the discrete ill-posed problems in general.

In the second part of the thesis, we analyze the stability of a method for estimating the risk-neutral density (RND) for the price of an asset from option prices. Typically, the RND is an element in L^1, and its estimation from option prices is an inverse problem. Different methods have been proposed to solve it. The method analyzed here provides a constraint qualification (CQ) under which the problem of estimating the RND as the maximizer of an entropy functional is reduced to the problem of minimizing a cost function whose solution is an element in $\mathbb{R}^n$. We benefit from the approach of this method, for it provides the alternative to analyze the stability of optimizers in $\mathbb{R}^n$ rather than in L^1. We also benefit from the characterization of the CQ when using the

data, and we show that it suffices to regularize the optimizers of the dual problem in $\mathbb{R}^n$ in order to obtain a stable solution (RND estimate) in L^1. In the method analyzed, the estimated RND corresponds to the so-called maximum entropy solution (MES). Hence, throughout the thesis, we always refer to the estimated RND as the MES.

The MES depends on a finite number of parameters which in turn depend on the available option prices (the data). In our analysis, we show that the MES is stable to perturbations in the parameters. More specifically, given a parameter vector $\boldsymbol{\lambda}$ and a perturbed one $\bar{\boldsymbol{\lambda}}$ with corresponding MESs $p(\boldsymbol{\lambda})$ and $p(\bar{\boldsymbol{\lambda}})$, we find a Lipschitz constant L such that $\|p(\boldsymbol{\lambda}) - p(\bar{\boldsymbol{\lambda}})\| \leq L\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|$ for some appropriate norms. We also show that, for the parameters to be stable under data perturbation, the Hessian of a cost function required by the method must be well-conditioned. Nonetheless, we illustrate with numerical examples that this Hessian is typically ill-conditioned, and the ill-conditioning is more severe as more data are used.

When noisy data are used, the CQ provided is usually not satisfied and the method cannot be used. We show how the original noisy data can be replaced with its projection onto a convenient subset so that the CQ is satisfied. Our numerical examples show that the projected noisy data may have less noise than the original noisy data, but ill-conditioning of the Hessian of the cost function is still present.

For this Hessian to become well-conditioned, we show that it suffices to add to the cost function an appropriate term which includes a new parameter. Ill-conditioning of the Hessian can be interpreted as ill-posedness of the minimization problem. Hence, addition of the new term to the cost function can be viewed as Tikhonov regularization of the ill-posed minimization problem. A MES obtained by minimizing the modified cost function is then called a regularized MES, the additional term in the cost function is referred to as the penalty term, and the new parameter is called the regularization parameter. We show that, using the L-curve method, an adequate regularization parameter can be selected. We demonstrate with numerical examples that the modified method is much more stable to perturbation in data.

We finish the second part of the thesis by supplying the regularization procedure with an analysis that shows that the regularized MES enjoys the classical convergence property.

The thesis is organized as follows. Chapter 2 is a review of the concepts of ill-posedness and regularization, as well as the L-curve method. The image deblurring problem is also described here. The review is mainly based on the monograph by Vogel on computational methods for inverse problems [10]. We finish chapter 2 by contributing a remark on the equivalence of the L-curve method and the efficient frontier.

Chapter 3 starts by introducing a multiplicative regularization strategy. The results of our analysis are then presented. We start the analysis by comparing the solutions obtained by the multiplicative regularization strategy with those obtained by Tikhonov regularization aided with the L-curve criterion. We then address the issue of determining the existence of a solution for a multiplicative strategy. We finish chapter 3 by presenting a condition for the local convexity of the cost function used in the multiplicative regularization strategy.

Chapter 4 presents a brief account of the literature of RND estimation. The problem that must be solved when the principle of maximum entropy is applied to this inverse problem is stated. Finally, the method analyzed to solve this problem is also described in this chapter.

Chapter 5 provides a calculation of a Lipschitz constant that shows the stability of the MES under parameter perturbation. The link between the ill-conditioning of a Hessian required by the method and the instability of the parameters under data perturbation is established in chapter 6. Numerical examples with noiseless, noisy, and projected noisy data are also presented in this chapter.

The regularization of the ill-posed problem and the use of the L-curve method in this setting is treated in chapter 7. Regularized MESs of the numerical examples as well as a convergence analysis are also found in this chapter. The thesis ends in chapter 8 with a discussion of the results obtained in both parts of the thesis.

Chapter 2

Review of ill-posedness and regularization

2.1 Ill-posedness

Hadamard [11] introduced the notion of well-posedness by stating that a well-posed problem has a solution which is unique and depends continuously on the data. More precisely, we have the following definition.

Definition 2.1.1. *[10] Let $K : \mathcal{H}_1 \to \mathcal{H}_2$. An operator equation,*

$$K(f) = g, \tag{2.1}$$

or problem, is said to be well posed if

(i) for each $g \in \mathcal{H}_2$, there exists $f \in \mathcal{H}_1$ such that (2.1) holds, in which case f is called a solution of (2.1);

(ii) the solution f is unique; and

(iii) if $K(f_*) = g_*$ and $K(f) = g$, then $f \to f_*$ whenever $g \to g_*$.

If any of the conditions, (i)–(iii), are not met, the problem (2.1) is called ill-posed.

Ill-posed problems arise in a great variety of applications such as signal processing and image restoration, biomedical imaging, mineral and petroleum prospecting, or stock

price distribution estimation, to mention just a few. Details of these and many other ill-posed problems are presented, for example, in [12, 13, 14, 15, 16] and the references therein.

If $\mathcal{H}_1$ and $\mathcal{H}_2$ in definition 2.1.1 are infinite-dimensional and K is a compact linear operator, then it can be shown that problem (2.1) is ill-posed. More precisely,

Theorem 2.1.2. *[10] Let $K : \mathcal{H}_1 \to \mathcal{H}_2$ be a compact linear operator, and let $\mathcal{H}_1$ and $\mathcal{H}_2$ be infinite-dimensional. If* Range(K) *is infinite-dimensional, then conditions* (i) *and* (iii) *of definition 2.1.1 are violated. If* Range(K) *has finite dimension, then condition* (ii) *is violated.*

We next present the problem of image deblurring which serves as an example for a better grasp of the concepts of ill-posedness and regularization. Our numerical examples in later sections will be based on this particular problem.

2.2 Image deblurring

In optical imaging, a physical object and its corresponding image are considered. Typically, a gray-scale image is represented by a function $g : \mathbb{R}^2 \to \mathbb{R}$, where $g(x, y)$ represents the light intensity of the point located at (x, y) as recorded by an image recording device. Let us denote by $f : \mathbb{R}^2 \to \mathbb{R}$ the actual light source intensity or object.

The blurring problem can then be described by the equation

$$g(x, y) = \iint_D k(x - x', y - y') f(x', y') dx' dy'. \tag{2.2}$$

It states that a *blurred continuous image* g is obtained by convolving the *light source intensity* or *object* f, which is defined on a rectangle $D \subset \mathbb{R}^2$, with a *blurring kernel* k, also called *point spread function* (PSF). A kernel that characterizes long-time average blurring effects of atmospheric turbulence on light propagation is the Gaussian [10, 17].

An application of interest is to solve the inverse problem. Namely, given the kernel k and the blurred image g, estimate the source f. The inverse problem given by (2.2) can be viewed as $(\mathcal{K}f)(x, y) = g(x, y)$. Smoothness of k leads to compactness of $\mathcal{K}$ (see e.g. [10]), so by theorem 2.1.2, problem (2.2) is ill-posed.

The one-dimensional version of model (2.2) is

$$g(x) = \int_0^1 k(x - x')f(x')dx' := (\mathcal{K}f)(x), \qquad 0 < x < 1. \tag{2.3}$$

The one-dimensional version of the Gaussian kernel is

$$k(x) = C \exp\left(-\frac{x^2}{2\gamma^2}\right), \tag{2.4}$$

where C and γ are positive parameters.

The blurred image g can be calculated at any value of x on $(0, 1)$ by using standard numerical quadrature. If we use midpoint quadrature n subintervals of the same size, the midpoints are $x_i = (i - 1/2)h$ for $i = 1, \ldots, n$, where $h = 1/n$. Suppose we evaluate g at the midpoints x_i. Then $g(x_i)$ can be calculated accurately by $K\mathbf{f}$, where

$$[K]_{ij} := hC \exp\left(-\frac{((i-j)h)^2}{2\gamma^2}\right), \qquad 1 \le i, j \le n, \tag{2.5}$$

and

$$[\mathbf{f}]_i := f(x_i), \qquad 1 \le i \le n.$$

This discretized version of (2.3) could be used to solve the inverse problem in a straightforward manner by simply inverting K if it is nonsingular. The problem is that the image recording device does not really record $g(x_i)$, but only $d_i := g(x_i) + \eta_i$, where η_i is noise that cannot be controlled. Then, as n becomes larger, K becomes more ill-conditioned, increasing the magnitude of $K^{-1}\boldsymbol{\eta}$ enormously. Since we cannot detach this term from $\mathbf{f}$, this approach ends up estimating $K^{-1}\boldsymbol{\eta}$ rather than $\mathbf{f}$, which has much smaller magnitude.

2.3 Tikhonov regularization

A common approach that still uses the discrete linear system $K\mathbf{f} = \mathbf{d}$, but circumvents direct matrix inversion is called regularization by filtering. We consider the discrete data model

$$\mathbf{d} = K\mathbf{f}_{\text{true}} + \boldsymbol{\eta} \tag{2.6}$$

with *noise level*

$$\delta := \|\boldsymbol{\eta}\| \ge 0, \tag{2.7}$$

where $\|\cdot\|$ denotes the Euclidean norm, $\mathbf{f}_{\text{true}}$ corresponds to the true discretized light source intensity, $\boldsymbol{\eta}$ represents errors in the data, and δ is known as the error level. Let us assume that K is an invertible matrix with real entries. Then K has a *singular value decomposition* (SVD) [18]

$$K = U\,\mathrm{diag}(s_i)V^T, \tag{2.8}$$

with strictly positive decreasing *singular values* s_i. The *right singular vectors*, or column vectors $\mathbf{v}_i$ of V, and the *left singular vectors*, or column vectors $\mathbf{u}_i$ of U, satisfy

$$\mathbf{u}_i^T\mathbf{u}_j = \delta_{ij}, \qquad \mathbf{v}_i^T\mathbf{v}_j = \delta_{ij}, \tag{2.9}$$

$$K\mathbf{v}_i = s_i\mathbf{u}_i, \qquad K^T\mathbf{u}_i = s_i\mathbf{v}_i. \tag{2.10}$$

The function δ_{ij} denotes the usual Kronecker delta. Thus, $U^T = U^{-1}$ and $V^T = V^{-1}$.

By properties (2.9), (2.10), we have

$$K^{-1}\mathbf{d} = V\,\mathrm{diag}\left(s_i^{-1}\right)U^T\mathbf{d} = \mathbf{f}_{\text{true}} + \sum_{i=1}^{n} s_i^{-1}\left(\mathbf{u}_i^T\boldsymbol{\eta}\right)\mathbf{v}_i. \tag{2.11}$$

Now, small deviations from the true data due to errors lead to small changes in the singular values, but those changes may be of the same order as that of some small singular values, which can result in big changes in the estimation of $K^{-1}\mathbf{d}$, as there is division by the singular values s_i in (2.11). This is known as *instability*. A common practice to deal with this instability is to multiply the s_i^{-1}'s in (2.11) by a *regularizing filter function* $w_\alpha(s_i^2)$ such that the product $w_\alpha(s^2)s^{-1} \to 0$ as $s \to 0$. In this way, one gets rid of small singular components of $K^{-1}\mathbf{d}$ which correspond to small singular values, obtaining an approximation to $\mathbf{f}_{\text{true}}$ parametrized by α which has the form

$$\mathbf{f}_\alpha = V\,\mathrm{diag}\left(w_\alpha\left(s_i^2\right)s_i^{-1}\right)U^T\mathbf{d} \tag{2.12}$$

We cannot get rid of too many singular components as that would lead to a poor estimation. To obtain an estimation with an acceptable degree of accuracy, we must keep singular components corresponding to large singular values. This is achieved by requiring that $w_\alpha(s^2) \approx 1$ for large values of s^2. A natural example of such a filter function would be

$$w_\alpha(s^2) = \begin{cases} 1, & \text{if } s^2 > \alpha; \\ 0, & \text{if } s^2 \le \alpha. \end{cases} \tag{2.13}$$

The resulting approximation $\mathbf{f}_\alpha$ is known, not surprisingly, as the *truncated* SVD (TSVD) solution to $K\mathbf{f} = \mathbf{d}$. Another example of a filter function which is commonly used is the *Tikhonov filter* function

$$w_\alpha\left(s_i^2\right) = \frac{s^2}{s^2+\alpha}. \tag{2.14}$$

One of the main problems that one faces with this approach is the selection of an appropriate regularization parameter α. On one hand, when α is very small, the singular components describing the main features are preserved, but likewise, part of the noise is not filtered, causing oscillations in $\mathbf{f}_\alpha$. On the other hand, when α is large, the noise components are successfully suppressed, but most components with relevant features of the solution are also filtered out, so that $\mathbf{f}_\alpha$ becomes too smooth. In addition, for very large ill-conditioned systems, it is often not practical to calculate their corresponding SVD required for the regularization by filtering.

Fortunately, solution (2.12) with w_α as in (2.14), which is known as the Tikhonov solution, has an alternate variational representation which might be easier to compute. We state this result as proposition next.

Proposition 2.3.1. *Solution* (2.12) *with* w_α *as in* (2.14)*, has the alternate variational representation*

$$\mathbf{f}_\alpha = \arg\min_{\mathbf{f}\in\mathbb{R}^n} \|K\mathbf{f}-\mathbf{d}\|^2 + \alpha\|\mathbf{f}\|^2. \tag{2.15}$$

The proof of this is left as an exercise in [10] and can be obtained by direct calculations.

Besides the possible advantage of being easier to compute, representation (2.15) offers the possibility to replace the penalty term $\|\mathbf{f}\|^2$ with other penalty term that incorporate a priori information. For example, for the image deblurring problem, if it is known that the image has sharp edges, $\|\mathbf{f}\|^2$ can be replaced with a penalty functional that penalizes highly oscillatory solutions while still allowing sharp discontinuities. It has been shown in the literature [19, 20, 21, 22, 23] that the *total variation*, $TV(\mathbf{f}) = \sum_{i=1}^{n-1} |f_{i+1} - f_i|$, has that property.

The test problem we will take from the literature is described in pages 1–3, and 12 of Vogel [10]. We simulate the data by first solving the direct problem with true source

intensity

$$f_{\text{true}} = \begin{cases} 0.75, & 0.1 < x < 0.25; \\ 0.25, & 0.3 < x < 0.32; \\ \sin^4(2\pi x), & 0.5 < x < 1; \\ 0 & \text{otherwise.} \end{cases}$$

For the kernel (2.4), $\gamma = 0.03$ and $C = 1/(\gamma\sqrt{2\pi})$ were used. Figure 2.1 shows $\mathbf{d}$ with a noise level $\|\boldsymbol{\eta}\|$ of 5%, that is, $\|\boldsymbol{\eta}\| = 0.05\|K\mathbf{f}\|$.

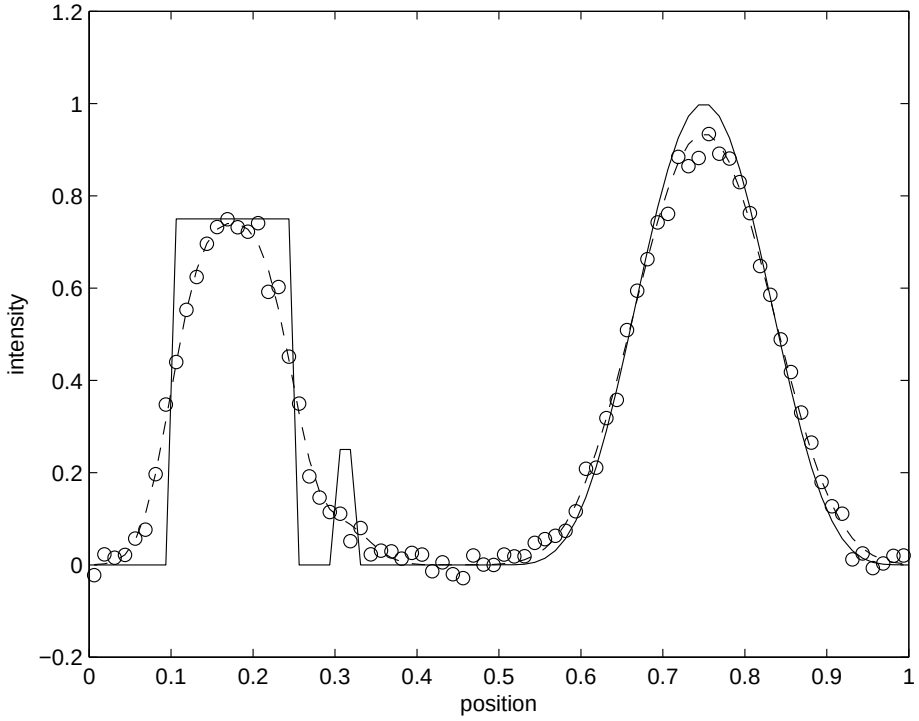

Figure 2.1: *Graph of a one-dimensional image. The source intensity f is represented by the solid line, the blurred image g is represented by the dashed line, and the noisy data $\mathbf{d}$ is represented by circles.*

This test problem does have sharp edges, so total variation for the penalty functional would be appropriate. However, we analyze the case where the Tikhonov penalty functional $\|\mathbf{f}\|^2$ is used, and this analysis cannot be easily extended when $TV(\mathbf{f})$ is used

instead. Hence, we will use the penalty term $\|B\mathbf{f}\|^2$, where

$$B := \begin{pmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & & 1 \\ & & & -1 \end{pmatrix}. \tag{2.16}$$

The difference between the term $TV(\mathbf{f})$ and the term $\|B\mathbf{f}\|^2$ is that $TV(\mathbf{f})$ approximates the L^1 norm of the derivative of f, whereas $\|B\mathbf{f}\|^2$ approximates the square of the L^2 norm.

2.4 Error analysis

A question that arises at this point is whether the regularization parameter can be selected in a manner that guarantees convergence as the error level goes to zero. The following analysis answers that question.

Let $R_\alpha := V \operatorname{diag}(w_\alpha(s_i^2)s_i^{-1})U^T$. Thus, by (2.12) we have $\mathbf{f}_\alpha = R_\alpha \mathbf{d}$. Let

$$\mathbf{e}_\alpha := \mathbf{f}_\alpha - \mathbf{f}_{\text{true}} = \mathbf{e}_\alpha^{\text{trunc}} + \mathbf{e}_\alpha^{\text{noise}}, \tag{2.17}$$

called the *regularized solution error*, where

$$\mathbf{e}_\alpha^{\text{trunc}} := R_\alpha K \mathbf{f}_{\text{true}} - \mathbf{f}_{\text{true}} = \sum_{i=1}^{n} (w_\alpha(s_i^2) - 1)(\mathbf{v}_i^T \mathbf{f}_{\text{true}})\mathbf{v}_i, \tag{2.18}$$

called the *solution truncation error due to regularization*, and

$$\mathbf{e}_\alpha^{\text{noise}} := R_\alpha \boldsymbol{\eta} = \sum_{i=1}^{n} w_\alpha(s_i^2) s_i^{-1} (\mathbf{u}_i^T \boldsymbol{\eta})\mathbf{v}_i, \tag{2.19}$$

called the *noise amplification error*.

Theorem 2.4.1. *[10] For both the TSVD filter* (2.13) *and the Tikhonov filter* (2.14)*, the regularization parameter α can be selected in a manner that guarantees that both these errors converge to zero as the error level $\delta \to 0$.*

For a proof, see [10], pages 6–7.

2.5 The L-curve method

The proof to the proposition above gives a parameter selection rule when the error level (2.7) is known. In practice, the knowledge of the error level is not available, so other parameter selection methods have to be implemented. Figure 2.2 illustrates the importance of the appropriate selection of the regularization parameter. The plots correspond to different parameters α selected for

$$\mathbf{f}_\alpha = \arg\min_{\mathbf{f}\in\mathbb{R}^n} \|K\mathbf{f} - \mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2 \tag{2.20}$$

applied to the data shown in figure 2.1. With little regularization, that is, with small α, not enough noise is suppressed and $\mathbf{f}_\alpha$ is highly oscillatory. If α is too large, then important singular components are also suppressed, so that $\mathbf{f}_\alpha$ becomes too smooth.

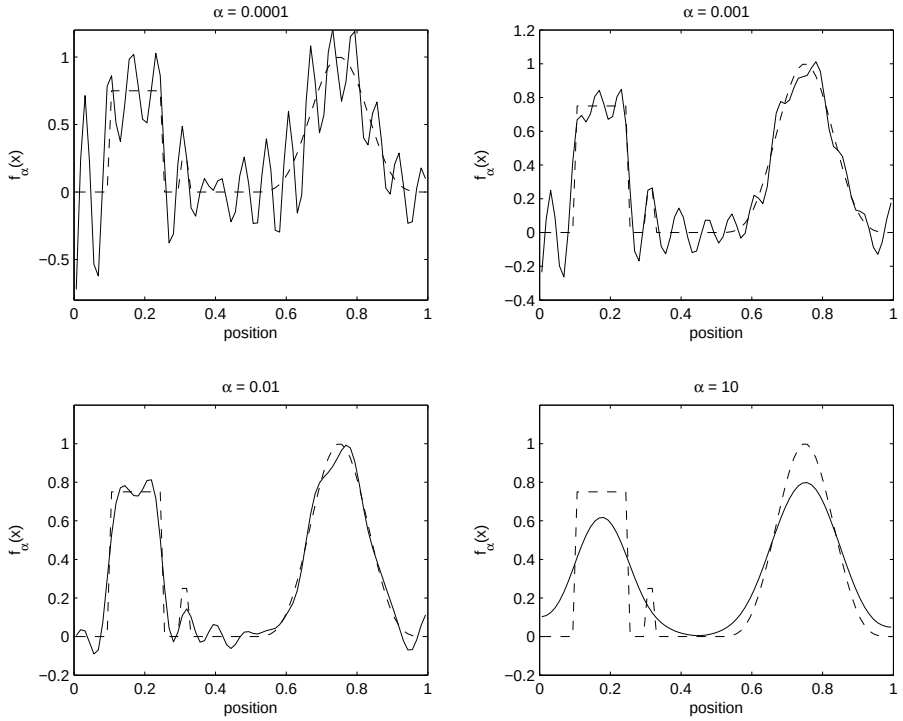

Figure 2.2: *Tikhonov regularized solutions* $\mathbf{f}_\alpha$ *for different values of* α*. From left to right and top to bottom,* α *is 0.0001, 0.001, .01 and 10.*

Because we know the true solution $\mathbf{f}_{\text{true}}$ for our example, we can easily pick a good

value for α given the plots below, namely $\alpha = 0.01$. Methods for the selection of the regularization parameter in certain optimal way without knowledge of the true solution nor the error level are numerous. Vogel [10] lists and briefly explains some of the most well-known methods, such as the discrepancy principle, the Generalized Cross Validation (GCV) method and the L-curve method. Hansen and O'Leary [3] give a detailed explanation on how to implement the L-curve method when dealing with discrete ill-posed problems.

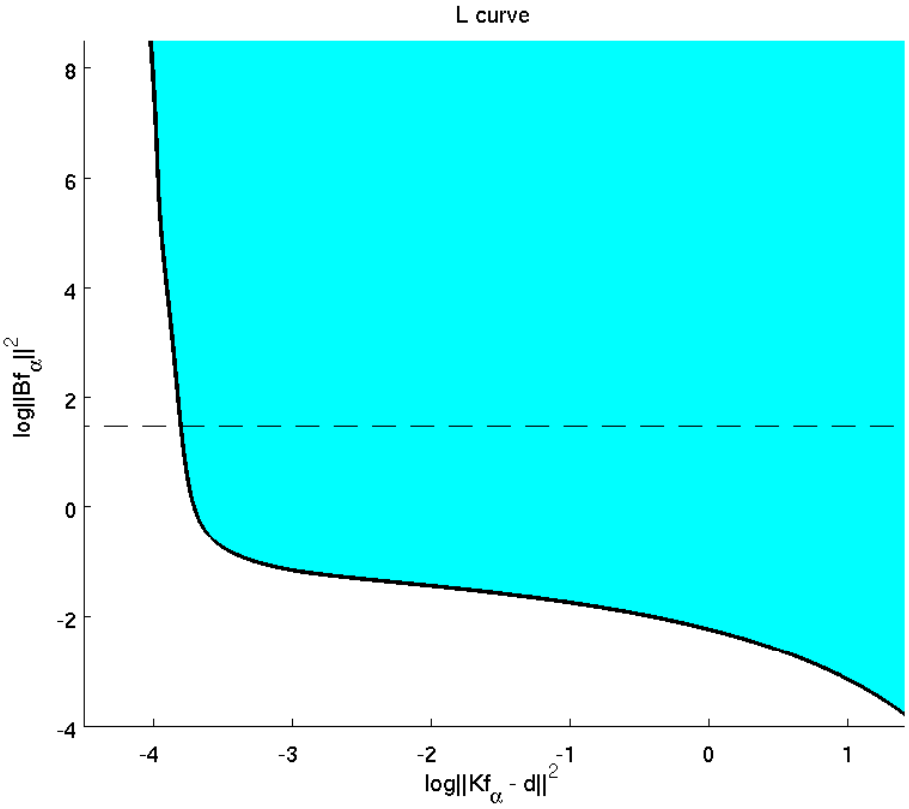

Figure 2.3: *L-curve. The solid line is a parametric plot of* $(\log ||K\mathbf{f}_\alpha - \mathbf{d}||^2, \log ||B\mathbf{f}_\alpha||^2)$ *known as the L-curve. Here,* K*,* B*,* $\mathbf{d}$*, and* $\mathbf{f}_\alpha$ *are given by* (2.5)*,* (2.16)*,* (2.6)*, and* (2.20) *respectively. A noise level of 5% was used, and* $[\alpha_{min}, \alpha_{max}] = [10^{-7}, 10^2]$.

The *L-curve* is a parametric plot of $(\log \|K\mathbf{f}_\alpha - \mathbf{d}\|^2, \log \|B\mathbf{f}_\alpha\|^2)$. For many problems, this plot has an L-shape. The solid line shown in figure 2.3 corresponds to the L-curve of the deblurring problem with K, B, $\mathbf{f}$, and $\mathbf{d}_\alpha$ given by (2.5), (2.16), (2.6), and (2.20)

respectively. In that case, a noise level of 5% was used, and a range $[\alpha_{\min}, \alpha_{\max}] = [10^{-7}, 10^{2}]$ was chosen. The L-curve method chooses the value of α for which $(\log \|K\mathbf{f}_\alpha - \mathbf{d}\|^2, \log \|B\mathbf{f}_\alpha\|^2)$ is exactly on the *corner* of the L-curve. The reason for this is that, at the horizontal part of the curve, the regularization errors dominate, whereas at the vertical part, the perturbation errors dominate. Hence, the corner is a point where both the regularization and the perturbation errors are balanced. Hansen and O'Leary [3] argue that, for discrete ill-posed problems, it is more appropriate to plot in a log-log scale. They prove that the corner is particularly pronounced in this scale.

By computing the curvature κ of the L-curve as a function of α and finding its maximum, it is possible to locate the corner. Given a parametric curve $(X(\alpha), Y(\alpha))$, its curvature function is

$$\kappa(\alpha) := \frac{X'(\alpha)Y''(\alpha) - X''(\alpha)Y'(\alpha)}{(X'(\alpha)^2 + Y'(\alpha)^2)^{3/2}}.$$

For the curve $(X(\alpha), Y(\alpha)) = (\log \|K\mathbf{f}_\alpha - \mathbf{d}\|^2, \log \|\mathbf{f}_\alpha\|^2)$, it is possible to obtain explicit representations for the first and second derivatives of X and Y with respect to α when $\mathbf{f}_\alpha$ is given by (2.12) with Tikhonov filter (2.14). In many other situations, however, the values of X and Y are available only at a limited number of values of α and no explicit representations for the first and second derivatives of X and Y are available. In those cases, we wish not to compute too many points $(X(\alpha), Y(\alpha))$ on the L-curve, as the computational effort required for each point is usually significant.

Convergence properties of the regularization parameter selected by the L-curve method were studied by Hanke [24] and Vogel [25] back in the time when the method was relatively new. The results suggested that the maximum curvature may not be the optimal feature of the L-curve for selecting the regularization parameter. On the other hand, comparisons of parameter selection methods that have followed in a variety of applications in the last fifteen years have often concluded that the L-curve method with maximum curvature location outperforms the other selection methods. Some of the most recent comparisons supporting this conclusion include [26, 27, 28]. GCV is the other selection method that often has been the method of choice in these comparisons.

2.6 Equivalent formulation of Tikhonov regularization

As it was seen above, it is possible to use different penalty functionals in the variational representation of Tikhonov regularization (2.15) other than $\|\mathbf{f}\|^2$. Let us focus only in the case when $\|\mathbf{f}\|^2$ is replaced with $\|B\mathbf{f}\|^2$, for some matrix B. That is,

$$\mathbf{f}_\alpha = \arg\min_{\mathbf{f}\in\mathbb{R}^n} \|K\mathbf{f}-\mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2.$$

It has been noticed in the literature that the penalty term could be added as a constraint, rather than as part of the cost function. This gives, for instance, the formulation

$$\mathbf{f}_c = \arg\min_{\mathbf{f}\in\mathbb{R}^n} \|K\mathbf{f}-\mathbf{d}\|^2 \qquad \text{s.t.} \qquad \|B\mathbf{f}\|^2 = c. \tag{2.21}$$

Golub et al. [29] state the equivalent formulation with inequality constraint $\|B\mathbf{f}\|^2 \leq c$. They mention that it can be shown that, under certain conditions for c, the solution $\mathbf{f}_\alpha$ to (2.20) is identical to the solution $\mathbf{f}_c$ to (2.21) with inequality constraint for an appropriately chosen c.

When equality constraint is considered as in (2.21), we show how to calculate c given α.

Proposition 2.6.1. *Given $\alpha > 0$, and a solution $\mathbf{f}_\alpha$ of problem (2.20) with well conditioned matrix B, there exists $c > 0$ such that $\mathbf{f}_\alpha$ is a solution of problem (2.21). The converse is also true.*

Proof. Let $\alpha > 0$, and let $\mathbf{f}_\alpha$ be a minimizer of

$$T(\mathbf{f}) := \|K\mathbf{f}-\mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2.$$

Then

$$\nabla T(\mathbf{f}_\alpha) = 2K^T(K\mathbf{f}_\alpha - \mathbf{d}) + 2\alpha B^T B\mathbf{f}_\alpha = \mathbf{0},$$

or

$$(K^TK + \alpha B^TB)\mathbf{f}_\alpha = K^T\mathbf{d}.$$

Since the square matrix $K^TK + \alpha B^TB$ is positive definite, it is invertible, and so

$$\mathbf{f}_\alpha = (K^TK + \alpha B^TB)^{-1}K^T\mathbf{d}.$$

Now, let $c := \|B(K^TK + \alpha B^TB)^{-1}K^T\mathbf{d}\|^2$. Then $(\mathbf{f}, \lambda) = (\mathbf{f}_\alpha, \alpha)$ is a solution of the system

$$\begin{aligned} 2K^T(K\mathbf{f} - \mathbf{d}) + 2\lambda B^TB\mathbf{f} &= \mathbf{0}, \\ \|B\mathbf{f}\|^2 &= c, \end{aligned}$$

or

$$\begin{aligned} \nabla_{\mathbf{f}}\mathcal{L}(\mathbf{f}_\alpha, \alpha) &= \mathbf{0}, \\ \|B\mathbf{f}_\alpha\|^2 &= c, \end{aligned}$$

where

$$\mathcal{L}(\mathbf{f}, \lambda) := \|K\mathbf{f} - \mathbf{d}\|^2 + \lambda(\|B\mathbf{f}\|^2 - c)$$

is the Lagrangian for the problem

$$\min_{\mathbf{f}\in\mathbb{R}^n} \|K\mathbf{f} - \mathbf{d}\|^2 \qquad \text{s.t.} \qquad \|B\mathbf{f}\|^2 = c.$$

Since $\nabla^2_{\mathbf{f},\mathbf{f}}\mathcal{L}(\mathbf{f}_\alpha, \alpha) = 2(K^TK + \alpha B^TB)$ is positive definite, $\mathbf{f}_\alpha$ is a minimizer of $\|K\mathbf{f} - \mathbf{d}\|^2$ such that $\|B\mathbf{f}_\alpha\|^2 = c$.

Conversely, let $c > 0$, and let $\mathbf{f}_c$ be a minimizer of $\|K\mathbf{f} - \mathbf{d}\|^2$ such that $\|B\mathbf{f}_c\|^2 = c$. Then there is a Lagrange multiplier $\lambda_c \neq 0$ such that

$$\nabla_{\mathbf{f}}\mathcal{L}(\mathbf{f}_c, \lambda_c) = 2K^T(K\mathbf{f}_c - \mathbf{d}) + 2\lambda_c B^TB\mathbf{f}_c = 0,$$

and

$$\mathbf{w}^T\nabla^2_{\mathbf{f},\mathbf{f}}\mathcal{L}(\mathbf{f}_c, \lambda_c)\mathbf{w} = 2\mathbf{w}^T(K^TK + \lambda_c B^TB)\mathbf{w} \geq 0,$$

for $\mathbf{w}$ such that $(B^TB\mathbf{f}_c)^T\mathbf{w} = 0$. But $\mathbf{w}^TK^TK\mathbf{w} \geq 0$ and $\mathbf{w}^TB^TB\mathbf{w} > 0$ for any $\mathbf{w}$, so necessarily $\lambda_c > 0$. Hence, for $\alpha = \lambda_c$,

$$\nabla T(\mathbf{f}_c) = 2K^T(K\mathbf{f}_c - \mathbf{d}) + 2\alpha B^TB\mathbf{f}_c = 0,$$

and

$$\nabla^2 T(\mathbf{f}_c) = 2(K^TK + \alpha B^TB)$$

is positive definite. Therefore, $\mathbf{f}_c$ is a minimizer of

$$T(\mathbf{f}) = \|K\mathbf{f}_c - \mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2,$$

where $\alpha = \lambda_c$. □

Just as the L-curve arises by plotting the penalty term $\|B\mathbf{f}_\alpha\|^2$ from (2.20) versus the residual $\|K\mathbf{f}_\alpha - \mathbf{d}\|^2$ and varying the value of α, we can plot the constraint term $\|B\mathbf{f}_c\|^2$ from (2.21) versus the residual $\|K\mathbf{f}_c - \mathbf{d}\|^2$ and vary the value of c. For example, suppose that we are solving problem (2.21) for the constraint $\log \|B\mathbf{f}\|^2 = 1.5$. The pair $(\log \|K\mathbf{f} - \mathbf{d}\|^2, \log \|B\mathbf{f}\|^2)$ will be anywhere on the intersection of the dashed line and the shaded region shown in figure 2.3, but only $(\log \|K\mathbf{f}_c - \mathbf{d}\|^2, \log \|B\mathbf{f}_c\|^2)$, where $\mathbf{f}_c$ solves problem (2.21), will be on the frontier of the shaded region. Hence, the curve $(\log \|K\mathbf{f}_c - \mathbf{d}\|^2, \log \|B\mathbf{f}_c\|^2)$ obtained for a range of values c in $[c_{\min}, c_{\max}]$ is called the *efficient frontier*. The proposition says, in other words, that the L-curve and the efficient frontier are the same curve.

Chapter 3

A multiplicative regularization strategy

3.1 Introduction to multiplicative regularization

We saw earlier from figure 2.1 that the determination of a good Tikhonov regularization parameter is crucial. For image deblurring, total variation (TV) regularization is preferred, but a parameter selection is required as well. It is pointed out by many workers that the parameter selection is the most difficult task in the regularization of this kind of inverse problems such as image deblurring. There is an approach, a *multiplicative regularization strategy* which does not require the previous selection of a parameter. This technique incorporates a regularizing term as a multiplicative constraint, and was originally introduced by van den Berg et al. [4] and by Abubakar and van den Berg [5].

We should notice that Tikhonov regularization can be applied directly to ill-posed problems such as (2.3), and not just to the discrete version (2.6) of problem (2.3). The discretization required for the numerical solution can be done after regularization. The cost functional in the continuous setting analogous to (2.12) would be

$$\|\mathcal{K}f - g\|^2 + \alpha\|f\|^2,$$

where now $\|\cdot\|$ denotes the L^2 norm.

Abubakar et al. [8] analyze the multiplicative strategy in the continuous setting. For the image deblurring problem in two dimensions, they propose the minimization of

the cost functional

$$\|\mathcal{K}f - g\|^2(\|b\nabla f\|^2 + \delta^2\|b\|^2),$$

where $b(\mathbf{x})$ is a positive weight and where δ^2 is a positive parameter that assures that the regularization factor does not vanish. They devise an algorithm where the parameter δ^2 is controlled by the optimization process itself. Multiplicative regularization strategies for ill-posed problems have been used in a variety of applications in the last decade (see for example [4, 5, 6, 7, 8, 9]).

In this chapter, we consider a multiplicative strategy applied to the discrete ill-posed problem (2.6) of the form

$$\mathbf{f}_* = \arg\operatorname{loc}\min_{\mathbf{f}\in\mathbb{R}^n-\{\mathbf{0}\}} \|K\mathbf{f} - \mathbf{d}\|^2\|B\mathbf{f}\|^2. \tag{3.1}$$

Here, we look for a local minimizer other than $\mathbf{f} = \mathbf{0}$. This expression could be seen as the most direct way of forming a parameter-free multiplicative cost functional with the factors being the residual and the penalty term of (2.20). In this chapter, we will call $\mathbf{f}_*$ a *multiplicative regularized solution*.

Notice that if we take the logarithm of the cost function in (3.1), then we could alternatively minimize $\log\|K\mathbf{f} - \mathbf{d}\|^2 + \log\|B\mathbf{f}\|^2$. This representation shares with Tikhonov regularization the property of an added penalty term. The difference being that this alternative representation penalizes the logarithm of the residual term with the logarithm of a seminorm of the solution sought.

As a side note, a functional similar to $\|K\mathbf{f} - \mathbf{d}\|^2\|B\mathbf{f}\|^2$, but dependent on a regularization parameter, has already been analyzed in the past by Regińska [30]. She relates the L-curve criterion to minimization of the function

$$\Phi_\beta(\alpha) := \|K\mathbf{f}_\alpha - \mathbf{d}\|^\beta\|\mathbf{f}_\alpha\|,$$

where $\beta > 0$ is a real parameter.

3.2 Tikhonov parameter for the multiplicative regularized solution

A natural question at this stage is whether the point $(\log\|K\mathbf{f}_* - \mathbf{d}\|^2, \log\|B\mathbf{f}_*\|^2)$ lies on the L-curve for a multiplicative regularized solution $\mathbf{f}_*$ given by (3.1). The following

proposition answers this question in the affirmative way.

Proposition 3.2.1. *Let* $\mathbf{f}_* \in \mathbb{R}^n$ *be a nonzero minimizer of*

$$M(\mathbf{f}) := \|K\mathbf{f} - \mathbf{d}\|^2 \|B\mathbf{f}\|^2$$

such that $\|K\mathbf{f} - \mathbf{d}\|^2 > 0$ *and* $\|B\mathbf{f}\|^2 > 0$. *Assume that* B *is nonsingular. Then there exists* $\alpha > 0$ *such that* $\mathbf{f}_*$ *is a minimizer of*

$$T_\alpha(\mathbf{f}) := \|K\mathbf{f} - \mathbf{d}\|^2 + \alpha \|B\mathbf{f}\|^2$$

Proof. Let $\mathbf{f}_*$ be as in the proposition. Then

$$\frac{1}{2}\nabla M(\mathbf{f}_*) = \|B\mathbf{f}_*\|^2 K^T(K\mathbf{f}_* - \mathbf{d}) + \|K\mathbf{f}_* - \mathbf{d}\|^2 B^T B\mathbf{f}_* = \mathbf{0}.$$

Let

$$\alpha := \frac{\|K\mathbf{f}_* - \mathbf{d}\|^2}{\|B\mathbf{f}_*\|^2}.$$

Then

$$\mathbf{0} = K^T(K\mathbf{f}_* - \mathbf{d}) + \alpha B^T B\mathbf{f}_* = \frac{1}{2}\nabla T_\alpha(\mathbf{f}_*).$$

Moreover, since $\alpha > 0$ and B is nonsingular, $\frac{1}{2}\nabla^2 T_\alpha(\mathbf{f}_*) = K^T K + \alpha B^T B$ is positive definite. Hence, $\mathbf{f}_*$ is a minimizer of $T_\alpha(\mathbf{f})$ for $\alpha = \frac{\|K\mathbf{f}_* - \mathbf{d}\|^2}{\|B\mathbf{f}_*\|^2}$. □

Knowing that there is a Tikhonov parameter α for which a multiplicative regularized solution $\mathbf{f}_*$ coincides with a Tikhonov regularized solution $\mathbf{f}_\alpha$, we now can estimate how far $\mathbf{f}_\alpha$ is from the solution $\mathbf{f}_\alpha^c$ picked by the L-curve criterion. We estimate this in the case where B in (2.20) is the identity matrix.

Proposition 3.2.2. *Let* $\mathbf{f}_* \in \mathbb{R}^n$ *be a nonzero minimizer of*

$$M(\mathbf{f}) := \|K\mathbf{f} - \mathbf{d}\|^2 \|\mathbf{f}\|^2$$

such that $\|K\mathbf{f} - \mathbf{d}\|^2 > 0$. *By proposition* (3.2.1), *let* $\alpha^* \in [\alpha_{min}, \alpha_{max}]$ *be such that* $\mathbf{f}_* = \mathbf{f}_{\alpha^*}$ *is a minimizer of* $T_{\alpha^*}(\mathbf{f})$, *where*

$$T_\alpha(\mathbf{f}) := \|K\mathbf{f} - \mathbf{d}\|^2 + \alpha \|\mathbf{f}\|^2,$$

Let $\alpha^c \in [\alpha_{min}, \alpha_{max}]$ *be such that* $(\log \|K\mathbf{f}_{\alpha^c} - \mathbf{d}\|^2, \log \|\mathbf{f}_{\alpha^c}\|^2)$ *is the L-curve corner and* $\mathbf{f}_{\alpha^c}$ *is a minimizer of* $T_{\alpha^*}(\mathbf{f})$*. Then*

$$\|\mathbf{f}_{\alpha^*} - \mathbf{f}_{\alpha^c}\| \leq C(K, [\alpha_{min}, \alpha_{max}]) \cdot |\alpha^* - \alpha^c|,$$

where

$$C(K, [\alpha_{min}, \alpha_{max}]) := \left[\sum_{i=1}^{n} \frac{\sigma_i^2 (\mathbf{u}_i^T \mathbf{d})^2}{(\sigma_i^2 + \alpha_{min})^4} \right]^{\frac{1}{2}},$$

and $K = [\mathbf{u}_1, \ldots, \mathbf{u}_n] \operatorname{diag}(\sigma_i) [\mathbf{v}_1, \ldots, \mathbf{v}_n]^T$ *is the SVD of* K*.*

Proof. Let $K = U \operatorname{diag}(\sigma_i) V^T$ be the SVD of K. It follows by proposition 2.3.1 that

$$\begin{aligned}
\|\mathbf{f}_{\alpha^*} - \mathbf{f}_{\alpha^c}\|^2 &= \left\| V \left[\operatorname{diag}\left(\frac{\sigma}{\sigma^2 + \alpha^*} \right) - \operatorname{diag}\left(\frac{\sigma}{\sigma^2 + \alpha^c} \right) \right] U^T \mathbf{d} \right\|^2 \\
&= \left\| \sum_{i=1}^{n} \sigma_i (\mathbf{u}_i^T \mathbf{d}) \left(\frac{1}{\sigma_i^2 + \alpha^*} - \frac{1}{\sigma_i^2 + \alpha^c} \right) \mathbf{v}_i \right\|^2 \\
&= \sum_{i=1}^{n} \sigma_i^2 (\mathbf{u}_i^T \mathbf{d})^2 \left[\frac{\alpha^c - \alpha^*}{(\sigma_i^2 + \alpha^*)(\sigma_i^2 + \alpha^c)} \right]^2 \\
&\leq \left[\sum_{i=1}^{n} \frac{\sigma_i^2 (\mathbf{u}_i^T \mathbf{d})^2}{(\sigma_i^2 + \alpha_{min})^4} \right] (\alpha^* - \alpha^c)^2.
\end{aligned}$$

□

This result gives an estimate of $\|\mathbf{f}_* - \mathbf{f}_{\alpha^c}\|$ in terms of the distance between the Tikhonov parameters α^* and α_c. Now, estimating $|\alpha^* - \alpha_c|$ would be desirable, but we do not see a direct way of doing this analytically. Nonetheless, we can show numerically that this distance is small for our test problem for a wide range of noise levels.

Now, Elden [31] shows that $\min_{\mathbf{f}} \|K\mathbf{f} - \mathbf{d}\|^2 + \alpha \|B\mathbf{f}\|^2$ with B being a band matrix, can be transformed into the standard form $\min_{\mathbf{f}} \|\tilde{K}\tilde{\mathbf{f}} - \tilde{\mathbf{d}}\|^2 + \alpha \|\tilde{\mathbf{f}}\|^2$. The vectors $\mathbf{f}$ and $\tilde{\mathbf{f}}$ are related by

$$\tilde{\mathbf{f}} = B\mathbf{f}, \qquad \mathbf{f} = A_1 \tilde{\mathbf{f}} + A_2 \mathbf{d},$$

where A_1 and A_2 depend on K and B. See [31] for their construction. Then, if $\mathbf{f}_{\alpha^*}$ is a minimizer of $\|K\mathbf{f} - \mathbf{d}\|^2 \|B\mathbf{f}\|^2$ and $\mathbf{f}_{\alpha^c}$ is a minimizer of $\|K\mathbf{f} - \mathbf{d}\|^2 + \alpha \|B\mathbf{f}\|^2$, we have that

$$\|\mathbf{f}_{\alpha^*} - \mathbf{f}_{\alpha^c}\| \leq C \cdot |\alpha^* - \alpha^c|,$$

where now C depends on K, B, and $[\alpha_{min}, \alpha_{max}]$.

To implement the L-curve method for our test problem, we used MATLAB. We first found a value α_{min} such that the solution $\mathbf{f}_\alpha$ looked very oscillatory. Then, we found α_{max} such that $\mathbf{f}_\alpha$ looked very smooth. This guaranteed that the interval $[\alpha_{min}, \alpha_{max}]$ would contain α^c selected by the L-curve criterion. We then divided the interval into 20 spaces such that $\log(\alpha_{min}) = \log(\alpha_1), \log(\alpha_2), \ldots, \log(\alpha_{21}) = \log(\alpha_{max})$ were equally spaced. Then we used finite differences to calculate the curvature $\kappa(\alpha)$ and find its maximum. We repeated the method for a new much smaller interval $[\alpha_{min}, \alpha_{max}]$ containing the maximizer found. We took the maximizer of the curvature in the second round as α^c. Since our test problem has only 80 data points, we could find $\mathbf{f}_{\alpha^c}$ for hundreds of simulations with the methodology described in a matter of seconds.

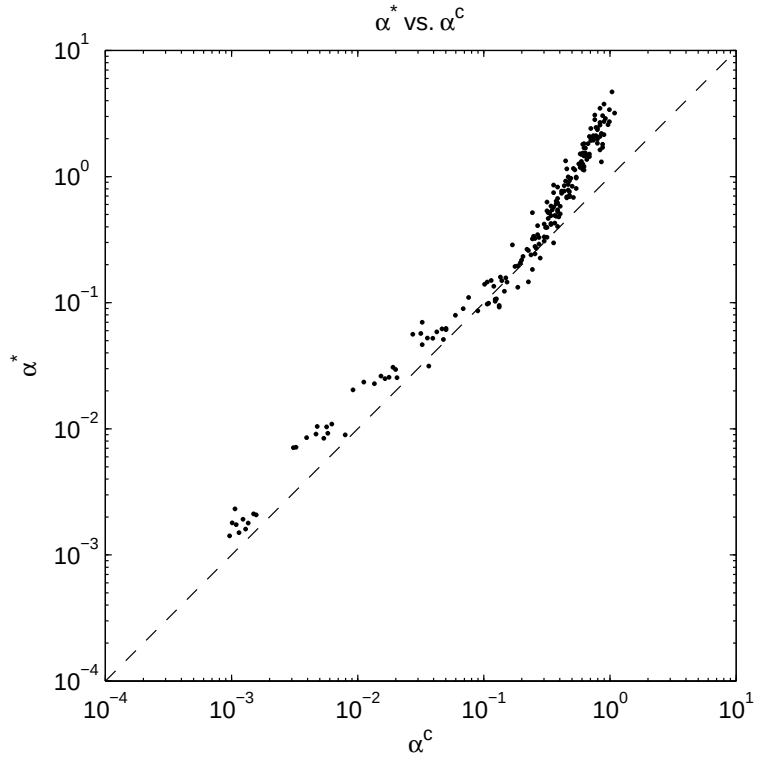

Figure 3.1: *Scatter plot in log-log scale of α^* versus α^c for the image deblurring problem shown in figure 2.1. 200 pairs (α^c, α^*) for different noise levels are shown. There are 10 noise levels equal to 1% of $\|K\mathbf{f}_{true}\|$, 10 noise levels equal to 2% of $\|K\mathbf{f}_{true}\|$, and so on up to 10 noise levels equal to 20%.*

For the multiplicative strategy, we used the built-in function `fminunc` from `MATLAB`. We supplied the gradient and the Hessian of the multiplicative cost function. We used the data $\mathbf{d}$ as the initial guess. The solutions $\mathbf{f}_*$ for hundreds of simulations were also found in a matter of seconds.

We ran 200 simulations of the deblurring problem shown in figure (2.1). Each simulation had different noise $\boldsymbol{\eta}$. We started with noise level $\delta = 1\%$ of $\|K\mathbf{f}_{true}\|$. After 10 simulations, we increased δ to 2% of $\|K\mathbf{f}_{true}\|$, and so on until 20% of $\|K\mathbf{f}_{true}\|$. Then we solved each simulation with both Tikhonov regularization using the L-curve method and with the multiplicative regularization strategy. We found the Tikhonov parameter α^* corresponding to each multiplicative regularized solution and plotted them against the α^cs chosen by the L-curve criterion. Figure 3.1 shows the resulting scatter plots. We can see that α^* tends to be larger than α^c, but only slightly, which shows that $|\alpha^* - \alpha_c|$ is small.

3.3 Existence of a local minimizer for the multiplicative cost functional

We saw that the Tikhonov parameters corresponding to the L-curve criterion and the multiplicative regularization strategy are close for noise levels under 20%. However, when we run simulations for noise levels greater than 20%, the minimizer found for the multiplicative strategy is the zero vector. We confirm this for hundreds of simulations with noise level between 20 and 40%. Only about one to two in ten simulations with a noise level of 21% yield a nonzero local minimizer. This suggests that the existence of a nonzero local minimizer might not be guaranteed after certain noise level. In fact, notice that, as the noise level increases, so does $\|\mathbf{d}\|^2$, and this term does play a role in the minimization with respect to $\mathbf{f}$ of

$$\|K\mathbf{f} - \mathbf{d}\|^2\|B\mathbf{f}\|^2 = \|K\mathbf{f}\|^2\|B\mathbf{f}\|^2 - 2(K\mathbf{f})^T\mathbf{d}\|B\mathbf{f}\|^2 + \|\mathbf{d}\|^2\|B\mathbf{f}\|^2.$$

On the contrary, the term $\|\mathbf{d}\|^2$ is irrelevant in the minimization of

$$\|K\mathbf{f} - \mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2 = \|K\mathbf{f}\|^2 - 2(K\mathbf{f})^T\mathbf{d} + \|\mathbf{d}\|^2 + \alpha\|B\mathbf{f}\|^2.$$

To be more precise, recall that $\mathbf{d}$ is given by (2.6), so that the multiplicative cost

function is

$$\begin{aligned}\|K\mathbf{f}-(K\mathbf{f}_{\text{true}}+\boldsymbol{\eta})\|^2\|B\mathbf{f}\|^2 &= \|K(\mathbf{f}-\mathbf{f}_{\text{true}})\|^2\|B\mathbf{f}\|^2 \\ &- 2(K(\mathbf{f}-\mathbf{f}_{\text{true}}))^T\boldsymbol{\eta}\|B\mathbf{f}\|^2+\|\boldsymbol{\eta}\|^2\|B\mathbf{f}\|^2.\end{aligned}$$

Since $\boldsymbol{\eta}$ is uncorrelated with $K(\mathbf{f}-\mathbf{f}_{\text{true}})$, the inner product $(K(\mathbf{f}-\mathbf{f}_{\text{true}}))^T\boldsymbol{\eta}$ is almost surely small. Hence, only for low to moderate noise level $\|\boldsymbol{\eta}\|$ is it possible to balance $-2(K(\mathbf{f}-\mathbf{f}_{\text{true}}))^T\boldsymbol{\eta}\|B\mathbf{f}\|^2$ with $\|K(\mathbf{f}-\mathbf{f}_{\text{true}})\|^2\|B\mathbf{f}\|^2+\|\boldsymbol{\eta}\|^2\|B\mathbf{f}\|^2$. But as $\|\boldsymbol{\eta}\|$ becomes larger, $-2(K(\mathbf{f}-\mathbf{f}_{\text{true}}))^T\boldsymbol{\eta}\|B\mathbf{f}\|^2$ remains small, whereas $\|\boldsymbol{\eta}\|^2\|B\mathbf{f}\|^2$ becomes larger and dominates the other terms. Any minimizer of $\|K\mathbf{f}-\mathbf{d}\|^2\|B\mathbf{f}\|^2$ is then given by a minimizer of $\|\boldsymbol{\eta}\|^2\|B\mathbf{f}\|^2$, which is in the null space of B. In our case, that is only the zero minimizer.

We now devise a method to determine whether the cost function $M(\mathbf{f})=\|K\mathbf{f}-\mathbf{d}\|^2\|B\mathbf{f}\|^2$ has a nonzero local minimizer for given data $\mathbf{d}$. If $M(\mathbf{f})$ has a nonzero local minimizer $\mathbf{f}_*$, then

$$\frac{1}{2}\nabla M(\mathbf{f}_*)=\|B\mathbf{f}_*\|^2K^T(K\mathbf{f}_*-\mathbf{d})+\|K\mathbf{f}_*-\mathbf{d}\|^2B^TB\mathbf{f}_*=\mathbf{0}.$$

It is not possible to determine directly the solution $\mathbf{f}_*$ to this equation and only an iterative optimization method can find it. However, by dividing $\frac{1}{2}\nabla M(\mathbf{f}_*)$ by $\|B\mathbf{f}_*\|^2$, we note that we can alternatively state that $\mathbf{f}_*$ satisfies the pair of equations

$$K^T(K\mathbf{f}_*-\mathbf{d})+\alpha B^TB\mathbf{f}_* = \mathbf{0}, \tag{3.2}$$

$$\frac{\|K\mathbf{f}_*-\mathbf{d}\|^2}{\|B\mathbf{f}_*\|^2} = \alpha, \tag{3.3}$$

for certain $\alpha>0$. Notice that (3.2) is

$$\frac{1}{2}\nabla\left\{\|K\mathbf{f}_*-\mathbf{d}\|^2+\alpha\|B\mathbf{f}_*\|^2\right\}=0,$$

the necessary condition for a Tikhonov minimizer. For this equation, we know it is possible, at least theoretically, to solve directly for $\mathbf{f}_*$ by inverting the matrix $K^TK+\alpha B^TB$. Let us denote that solution by $\mathbf{f}_\alpha$. That is, $\mathbf{f}_\alpha=(K^TK+\alpha B^TB)^{-1}(K^T\mathbf{d})$. Then we can evaluate $\|K\mathbf{f}_\alpha-\mathbf{d}\|^2/\|B\mathbf{f}_\alpha\|^2$.

The previous argument naturally gives rise to a method for determining the existence of a nonzero local minimizer: If we can find α_1 such that $\|K\mathbf{f}_{\alpha_1}-\mathbf{d}\|^2/\|B\mathbf{f}_{\alpha_1}\|^2>\alpha_1$ and

α_2 such that $\|K\mathbf{f}_{\alpha_2} - \mathbf{d}\|^2/\|B\mathbf{f}_{\alpha_2}\|^2 < \alpha_2$, then we know there must exist α between α_1 and α_2 such that $\|K\mathbf{f}_\alpha - \mathbf{d}\|^2/\|B\mathbf{f}_\alpha\|^2 = \alpha$. Then $\mathbf{f}_\alpha$ is a stationary point of $M(\mathbf{f})$. If the Hessian $\nabla^2 M(\mathbf{f}_\alpha)$ is positive definite, then $\mathbf{f}_\alpha$ is a nonzero local minimizer. On the other hand, if, for every α in certain interval $[\alpha_{\min}, \alpha_{\max}]$, the difference $\|K\mathbf{f}_\alpha - \mathbf{d}\|^2/\|B\mathbf{f}_\alpha\|^2 - \alpha$ has always the same sign, then we conclude that there is no nonzero local minimizer.

Using the range of values $[\alpha_{\min}, \alpha_{\max}]$ that were used for the L-curve criterion, we can plot, say, the function

$$\phi(\alpha) = \frac{\|K\mathbf{f}_\alpha - \mathbf{d}\|^2}{\alpha\|B\mathbf{f}_\alpha\|^2}.$$

Figure 3.2 shows the graph of this function for data $\mathbf{d}$ with noise level of 5 and 30%

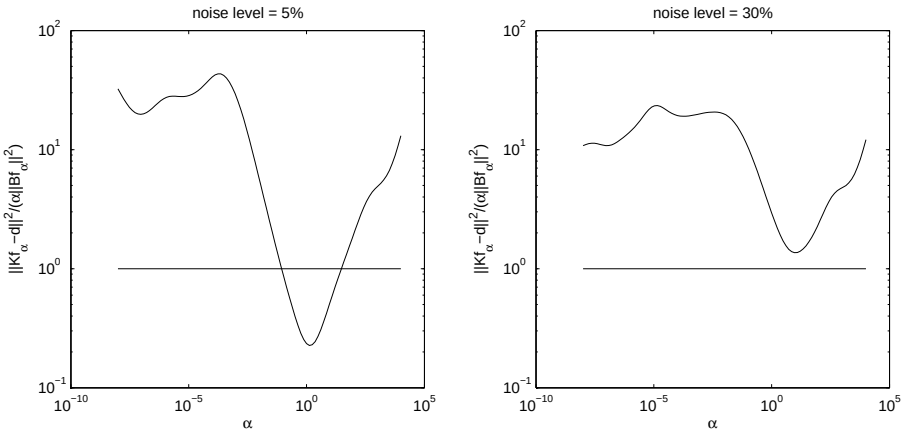

Figure 3.2: *Graph of* $\|K\mathbf{f}_\alpha - \mathbf{d}\|^2/(\alpha\|B\mathbf{f}_\alpha\|^2)$. *The left subplot shows a graph for data* $\mathbf{d}$ *with noise level of 5%. The right subplot shows a graph for data* $\mathbf{d}$ *with noise level of 30%.*

respectively. A graph similar to the first one is obtained for all those problems where a nonzero local minimizer was found. A graph similar to the second one is obtained for all those problems whose minimizer found was the zero vector. This rules out the possibility of an initial guess that was closer to zero than to a nonzero local minimizer and confirms the nonexistence of nonzero local minimizers in those cases.

From the graph with noise level of 5%, we see that there are two values of α which yield a stationary point $\mathbf{f}_\alpha$, but $\nabla^2 M(\mathbf{f}_\alpha)$ is positive definite only at the point corresponding to the smaller α, so there is only one nonzero local minimizer. The situation is the same for all the problems where a nonzero local minimizer was found.

We would like to visualize the behavior of the cost function $M(\mathbf{f})$, which obviously

cannot be plotted against $\mathbf{f}$. However, we can plot

$$\Phi(t) = M(t\mathbf{f}_0 + (t-1)\mathbf{f}_1)$$

against $t \in \mathbb{R}$ for a given $\mathbf{f}_0$ and $\mathbf{f}_1$. For example, let $\mathbf{f}_1$ be the zero vector, and let $\mathbf{f}_0$ be the minimizer of a cost function $M(\mathbf{f})$ where the data have 5% noise level. Then, when $t = 1$, the graph will show the value of the local minimum $M(\mathbf{f}_*)$; and when $t = 0$, it will show the value of the global minimum $M(\mathbf{0})$, which is zero.

Keeping $\mathbf{f}_0$ and $\mathbf{f}_1$ fixed, let us also plot $\Phi(t)$ where the data have 15, 25, and 35% noise level. Figure 3.3 shows the plots obtained for all four cases (5, 15, 25, and 35% noise level).

The graph for a 5% noise level shows a very small value for $M(\mathbf{f})$ at the nonzero minimizer. This value increases as the noise level increases, until, for 35%, clearly there is no nonzero minimizer. For a 25% noise level, $\Phi(t)$ appears to have a local minimimum. However, $M(\mathbf{f})$ has no local minimum in this case. We could imagine the graph of $M(\mathbf{f})$ against $\mathbf{f}$ as a sort of "curved slide" in $\mathbb{R}^{n+1}$ with a cross section such as the plot shown, but where it is possible to "slide" all the way down to zero by surrounding the "hump" which is in the middle.

3.4 A local convexity condition for the multiplicative cost functional

It is clear from the previous sections that the multiplicative cost functional $M(\mathbf{f})$ is not convex. It is then necessary to check, given $\mathbf{f}_0 \in \mathbb{R}^n$, whether the Hessian $\nabla^2 M(\mathbf{f}_0)$ is positive definite. In the previous section, for example, we checked for positive definiteness at the stationary points of $M(\mathbf{f})$ to determine which point was a minimizer. In this section, we derive a condition for the convexity of $M(\mathbf{f})$ at a given $\mathbf{f}$ which reduces the calculation of positive definiteness of the Hessian $\nabla^2 M(\mathbf{f})$ to the calculation of the minimum eigenvalue of $B^T B$ and the calculation of some inner products. When B is the identity matrix, then the result derived here reduces the positive definiteness calculation to just the calculation of some inner products.

The proof of the main result of this section uses a lemma that we state and prove first.

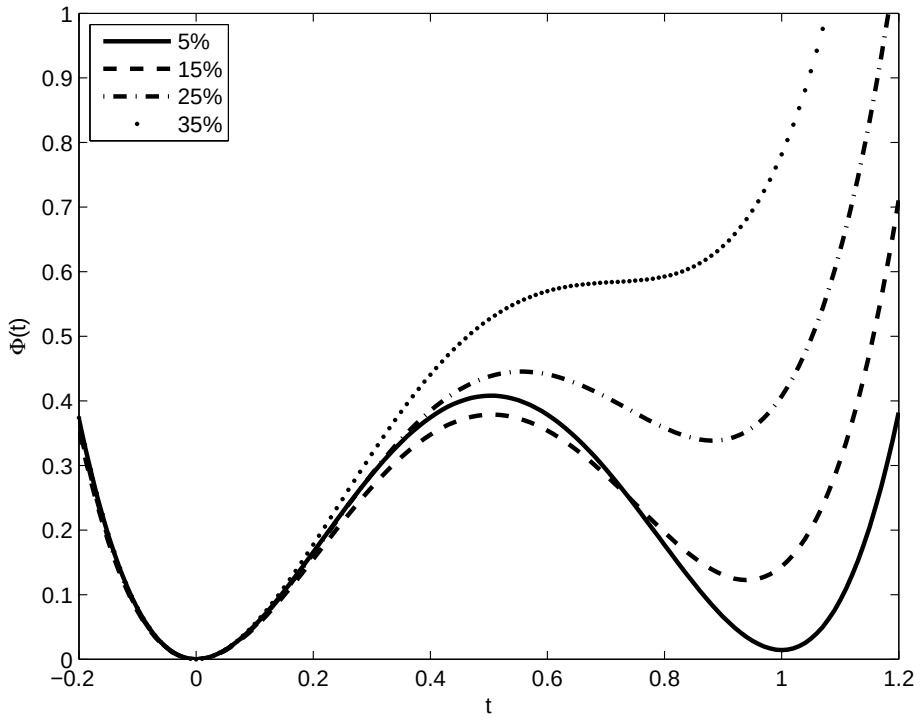

Figure 3.3: *Graphs of* $\|K(t\mathbf{f}_\alpha) - \mathbf{d}\|^2 \|B(t\mathbf{f}_\alpha)\|^2$ *versus* t *for noise levels of 5, 15, 25, and 35%.*

Lemma 3.4.1. *If* $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ *are nonzero vectors that are not parallel, then the matrix* $A := \mathbf{x}\mathbf{y}^T + \mathbf{y}\mathbf{x}^T$ *has rank 2 with eigenvalues*

$$\lambda_1 = \mathbf{x}^T\mathbf{y} + \|\mathbf{x}\|\|\mathbf{y}\|, \qquad \lambda_2 = \mathbf{x}^T\mathbf{y} - \|\mathbf{x}\|\|\mathbf{y}\|, \tag{3.4}$$

and corresponding eigenvectors

$$\mathbf{v}_1 = \|\mathbf{y}\|\mathbf{x} + \|\mathbf{x}\|\mathbf{y}, \qquad \mathbf{v}_2 = \|\mathbf{y}\|\mathbf{x} - \|\mathbf{x}\|\mathbf{y}. \tag{3.5}$$

Consequently, one eigenvalue is positive and the other one is negative.

Proof. The matrix A is $[y_1\mathbf{x} + x_1\mathbf{y} \quad \dots \quad y_n\mathbf{x} + x_n\mathbf{y}]$, so any column of A is a linear combination of $\mathbf{x}$ and $\mathbf{y}$. Conversely, $\mathbf{x}$ or $\mathbf{y}$ is the linear combination of $y_i\mathbf{x} + x_i\mathbf{y}$ and $y_j\mathbf{x} + x_j\mathbf{y}$ for some $1 \le i, j \le n$. Hence, A has rank 2. Its eigenvectors are then of the form $\alpha\mathbf{x} + \beta\mathbf{y}$. That is, $A(\alpha\mathbf{x} + \beta\mathbf{y}) = \lambda(\alpha\mathbf{x} + \beta\mathbf{y})$ for some $\alpha, \beta, \lambda \in \mathbb{R}$. Expanding

both sides,

$$[\alpha(\mathbf{x}^T\mathbf{y}) + \beta(\mathbf{y}^T\mathbf{y})]\mathbf{x} + [\alpha(\mathbf{x}^T\mathbf{x}) + \beta(\mathbf{x}^T\mathbf{y})]\mathbf{y} = \lambda\alpha\mathbf{x} + \lambda\beta\mathbf{y}.$$

By equating coefficients, we obtain the system

$$\begin{aligned}(\mathbf{x}^T\mathbf{y} - \lambda)\alpha + (\mathbf{y}^T\mathbf{y})\beta &= 0,\\ (\mathbf{x}^T\mathbf{x})\alpha + (\mathbf{x}^T\mathbf{y} - \lambda)\beta &= 0,\end{aligned}$$

which has a nontrivial solution (α, β) only if $(\mathbf{x}^T\mathbf{y} - \lambda)^2 = (\mathbf{x}^T\mathbf{x})(\mathbf{y}^T\mathbf{y})$. Solving this last equation for λ gives the eigenvalues (3.4). Substituting λ into the system and solving for (α, β) gives the eigenvectors (3.5). The sign of each eigenvalue follows from Cauchy-Schwarz inequality. □

To simplify the notation of the proposition, one definition is in order. Given matrices A_1 and A_2 of sizes $m_1 \times n_1$ and $m_2 \times n_2$ respectively, define the functional $\langle\cdot,\cdot\rangle_{A_1,A_2} : \mathbb{R}^{m_1} \times \mathbb{R}^{m_2} \to \mathbb{R}$ by

$$\langle \mathbf{v}_1, \mathbf{v}_2\rangle_{A_1,A_2} := (A_1^T\mathbf{v}_1)^T(A_2^T\mathbf{v}_2), \qquad \text{for all } \mathbf{v}_1 \in \mathbb{R}^{m_1}, \mathbf{v}_2 \in \mathbb{R}^{m_2}. \tag{3.6}$$

When $A_1 = A_2$, the functional above induces a norm in $\mathbb{R}^{m_1}$, namely

$$\|\mathbf{v}\|_A := \sqrt{\langle \mathbf{v}, \mathbf{v}\rangle}_{A,A} = \sqrt{\mathbf{v}^T A A^T \mathbf{v}}, \qquad \text{for all } \mathbf{v} \in \mathbb{R}^m. \tag{3.7}$$

Proposition 3.4.2. *Let K be an $m \times n$ matrix and let B be a well conditioned $p \times n$ matrix. Let $\mathbf{d} \in \mathbb{R}^m$, and let $\mathbf{f} \in \mathbb{R}^n$ be a nonzero vector such that $\|K\mathbf{f} - \mathbf{d}\|^2 > 0$. Let β_{min} be the smallest eigenvalue of the positive definite matrix B^TB. If*

$$\frac{1}{2}\beta_{min}\|K\mathbf{f} - \mathbf{d}\|^2 + \langle K\mathbf{f} - \mathbf{d}, B\mathbf{f}\rangle_{K,B} - \|K\mathbf{f} - \mathbf{d}\|_K\|B\mathbf{f}\|_B > 0, \tag{3.8}$$

(cf. (3.6), (3.7)) then the Hessian of $M(\mathbf{f}) := \|K\mathbf{f} - \mathbf{d}\|^2\|B\mathbf{d}\|^2$ is positive definite.

Proof. The Hessian of $M(\mathbf{f})$ is

$$\begin{aligned}\nabla^2 M(\mathbf{f}) &= 4[K^T(K\mathbf{f} - \mathbf{d})](B^TB\mathbf{f})^T + 2\|B\mathbf{f}\|^2K^TK\\ &+ 4(B^TB\mathbf{f})[K^T(K\mathbf{f} - \mathbf{d})]^T + 2\|K\mathbf{f} - \mathbf{d}\|^2B^TB.\end{aligned}$$

Or, for later convenience,

$$\begin{aligned}\frac{1}{4}\nabla^2 M(\mathbf{f}) &= \frac{1}{2}\|K\mathbf{f}-\mathbf{d}\|^2 B^T B + \frac{1}{2}\|B\mathbf{f}\|^2 K^T K \\ &+ [K^T(K\mathbf{f}-\mathbf{d})](B^T B\mathbf{f})^T + (B^T B\mathbf{f})[K^T(K\mathbf{f}-\mathbf{d})]^T.\end{aligned}$$

If $\mathbf{f}$ is a nonzero vector such that $\|K\mathbf{f}-\mathbf{d}\|^2 > 0$, and B is well conditioned, then $P := \frac{1}{2}\|K\mathbf{f}-\mathbf{d}\|^2 B^T B$ is positive definite and $Q := \frac{1}{2}\|B\mathbf{f}\|^2 K^T K$ is positive semidefinite. Now, let $\mathbf{x} := K^T(K\mathbf{f}-\mathbf{d})$ and $\mathbf{y} := B^T B\mathbf{f}$. Then the last two terms of the equation above are $R := \mathbf{x}\mathbf{y}^T + \mathbf{y}\mathbf{x}^T$. By lemma 3.4.1, R has eigenvalues $\lambda_1 \geq 0 = \ldots = 0 \geq \lambda_2$.

To find the condition required for $\nabla^2 M(\mathbf{f})$ to be positive definite, consider a nonzero $\mathbf{w} \in \mathbb{R}^n$. Then

$$\frac{1}{4}\mathbf{w}^T\nabla^2 M(\mathbf{f})\mathbf{w} = \mathbf{w}^T P\mathbf{w} + \mathbf{w}^T Q\mathbf{w} + \mathbf{w}^T R\mathbf{w}.$$

Let $\{\mathbf{v}_i\}$ be an orthonormal basis of eigenvectors for P with corresponding eigenvalues γ_i. Then

$$\begin{aligned}\mathbf{w}^T P\mathbf{w} &= \sum_i \alpha_i \mathbf{v}_i^T P \sum_j \alpha_j \mathbf{v}_j = \sum_i \alpha_i \mathbf{v}_i^T \sum_j \alpha_j \gamma_j \mathbf{v}_j = \sum_i \alpha_i^2 \gamma_i \\ &\geq \gamma_{\min} \sum_i \alpha_i^2 = \gamma_{\min} \sum_i \alpha_i \mathbf{v}_i^T \sum_j \alpha_j \mathbf{v}_j = \gamma_{\min}\mathbf{w}^T\mathbf{w},\end{aligned}$$

for some $\alpha_i \in \mathbb{R}$.

Now, if $\lambda_2 > -\gamma_{\min}$, then the eigenvalues of $R + \gamma_{\min} I$ are $\lambda_1 + \gamma_{\min} \geq \gamma_{\min} = \ldots = \gamma_{\min} \geq \lambda_2 + \gamma_{\min} > 0$ so that $R + \gamma_{\min} I$ is positive definite. Hence, for every nonzero $\mathbf{w} \in \mathbb{R}^n$,

$$\begin{aligned}0 &< \mathbf{w}^T(R + \gamma_{min} I)\mathbf{w} = \mathbf{w}^T R\mathbf{w} + \gamma_{min}\|\mathbf{w}\|^2 \\ &\leq \mathbf{w}^T R\mathbf{w} + \mathbf{w}^T P\mathbf{w} \leq \mathbf{w}^T R\mathbf{w} + \mathbf{w}^T P\mathbf{w} + \mathbf{w}^T Q\mathbf{w} \\ &= \frac{1}{4}\mathbf{w}^T\nabla^2 M(\mathbf{f})\mathbf{w},\end{aligned}$$

so that $\nabla^2 M(\mathbf{f})$ is positive definite.

If $\beta_{\min}$ is the smallest eigenvalue of $B^T B$, notice that then

$$\gamma_{\min} = \frac{1}{2}\beta_{\min}\|K\mathbf{f}-\mathbf{d}\|^2;$$

and by lemma 3.4.1,

$$\lambda_2 = [K^T(K\mathbf{f}-\mathbf{d})]^T(B^T B\mathbf{f}) - \|K^T(K\mathbf{f}-\mathbf{d})\|\|B^T B\mathbf{f}\|.$$

Using definitions (3.6) and (3.7) and substituting these last two equations into $\lambda_2 > -\gamma_{\min}$ gives the inequality of the proposition. □

It is important to emphasize that the convexity condition just derived is only sufficient. That is, if (3.8) holds at a given $\mathbf{f}$, then the Hessian $\nabla^2 M(\mathbf{f})$ is guaranteed to be positive definite, but $\nabla^2 M(\mathbf{f})$ could be positive definite even if $\mathbf{f}$ does not satisfy (3.8). Figure 3.4 illustrates this situation with a plot of both (one fourth of) the minimum eigenvalue of $\nabla^2 M(\mathbf{f}_t)$ and the left hand side of inequality (3.8) for a range of values $\mathbf{f}_t$. We choose the range of values $\mathbf{f}_t$ in the same way as in the previous section, namely, $\mathbf{f}_t = t\mathbf{f}_0 + (1-t)\mathbf{f}_1$, with $\mathbf{f}_0$ being the minimizer of $M(\mathbf{f})$, and $\mathbf{f}_1$, the zero vector. To simplify calculations, we choose B as the identity matrix.

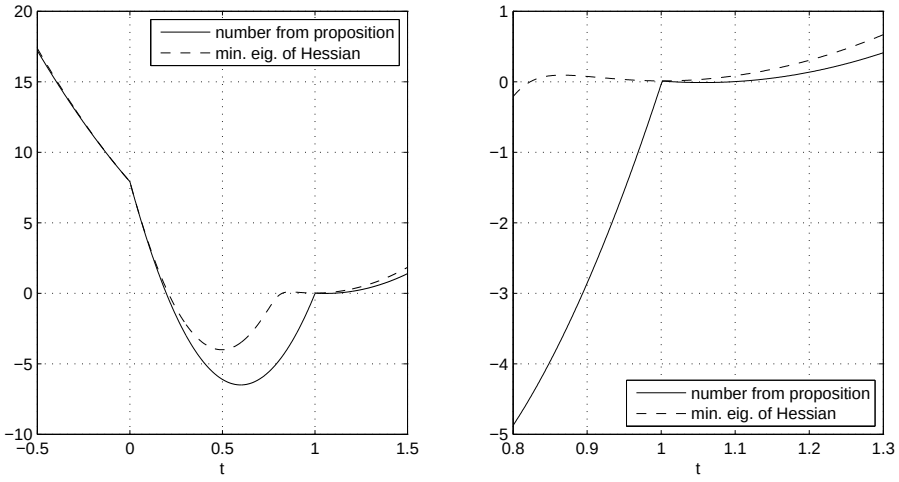

Figure 3.4: *Minimum eigenvalue of* $\nabla^2 M(\mathbf{f}_t)$ *and the left hand side of* (3.8) *plotted against* t*, where* $\mathbf{f}_t = t\mathbf{f}_0 + (1-t)\mathbf{f}_1$*. For the left subplot,* $-.5 \le t \le 1.5$*. For the right subplot,* $.8 \le t \le 1.2$*.*

From the figure, we can see that $M(\mathbf{f}_t)$ is positive definite for roughly $t \le 0.2$ and $t \ge 0.82$. The right hand side of inequality (3.8)is positive for that range of values of t, except for $0.82 \le t \le 1$.

Chapter 4

Risk-neutral density estimation

4.1 Introduction to risk-neutral density estimation

We now turn our attention to the estimation of risk-neutral densities from option prices, which is an inverse problem that arises in financial mathematics.

The link between option prices and the risk-neutral density (RND) for the price of an asset was established in 1976 when Cox and Ross [32] showed that the valuation of an option with expiration date T amounts to the calculation of its discounted expected payoff under the risk-neutral probability distribution of the asset price at time T. Since option prices are usually available for a variety of strike prices while the RND is unknown, it is the inverse problem of estimating the RND from the available option prices that needs to be addressed. Hidden in the option prices are the market beliefs that are driven by political and economic events. Investment professionals can assess these market beliefs by means of the RND. In addition, the RND estimated from option prices corresponding to a particular underlying asset and expiration date can be used to valuate any other exotic derivatives with the same underlying and expiry.

Almost twenty years after the relationship between option prices and the RND was pinpointed, a variety of methods to solve the inverse problem were quickly developed when, in the mid-1990s, both powerful computers and large option databases were made available to the public. Some of the methods were based on the risk-neutral valuation equation mentioned above, but many others were based on the equation proposed by Breeden and Litzenberger [33], which states that the RND is proportional to the second

derivative of the call (put) price viewed as a function of the strike price.[1] To date, the risk-neutral valuation equation and the second derivative equation remain the core of RND estimation from option prices. Current research is mainly focused on improving the features of both the methods and the estimates by resorting to mathematical tools that had not been employed in the past.

One of the approaches used in the estimation of the RND is to apply the principle of maximum entropy. Entropy optimization for the estimation of a probability distribution from information on a finite number of its moments is a Bayesian method of statistical inference well established throughout the sciences [35]. Buchen and Kelly [36] first applied this idea to estimate the RND. In their approach, the entropy is a functional on L^1 whose maximizer is taken to be the RND. The risk-neutral valuation equations are the moments of the RND which serve as constraints. The main motivation to apply the principle of maximum entropy in this context is that the RND obtained is the least prejudiced estimate in the sense that it is least committal with respect to unknown or missing information (see [36] and the references therein). One of its main advantages over other approaches is that this method produces RNDs that are strictly positive.[2]

On the other hand, the Hessian that appears in the estimation of the RND in this method (see [36, 38]) becomes more ill-conditioned as the number of available option prices increases.

Research that has followed in this stream of work includes that of Stutzer [39], who also estimates the RND via the principle of maximum entropy, but instead of using option prices, he uses historical data of the underlying asset. Avellaneda et al. [40] calibrate volatility surfaces to a prescribed set of option prices via the maximum entropy principle. Neri and Schneider [41] also estimate the RND from call option prices via the maximum entropy principle, but they additionally include constraints based on binary options, which have been available to the public only very recently. Also very recently, in order to avoid the ill-conditioned Hessian mentioned above, Rompolis [42] does not directly use the set of option prices in the constraints when applying the principle of

[1] For an account and comparison of most of the RND estimation methods developed by 2003, see Jackwerth's monograph [34], which also includes details on risk neutrality and distinguishes the RND from the actual distribution.

[2] Some popular methods estimate RNDs that may result in negative probabilities. For a concise review of the mathematical ideas behind some of the most popular methods, see the appendix B of Bondarenko [37].

maximum entropy. Instead, he uses the non-central moments of the RND implied by this set. Borwein et al. [43] provide an effective characterization of the constraint qualification (CQ) under which the problem, as proposed by Buchen and Kelley [36] or Avellaneda [38], reduces to optimizing an explicit function in finitely many variables.

A routine procedure to maximize the entropy, which is subject to a set of constraints, is to use the theory of Lagrange multipliers. In the case of RND estimation, one "differentiates" the Lagrangian with respect to density functions $p \in L^1$ [36, 40, 38]. Borwein et al. [43] circumvent this differentiation by using the theory of convex duality and partially finite convex programming.

4.2 The maximum entropy problem

We want to infer, from some available data, the *risk-neutral probability distribution* $p(x)$ for the *price* x of an *asset* at a set *future time* T. The available data are the *current price* d_1 of the asset, the prices $d_2, \ldots, d_m$ of a finite number of *European call options*, as well as the associated strike prices $k_1, \ldots, k_m$, where $k_1 = 0$. The options are assumed to be sorted so that $0 = k_1 < k_2 < \cdots < k_m$. The interval I of possible prices for the asset is of the form $[0, K)$, where K is either some fixed positive number or $K = +\infty$. Moreover, we assume risk-neutrality; that is, all assets have the same expected return, so that the price of a synthetic asset is the present value of its expected payoff.

According to the maximum-entropy principle, $p(x)$ should be estimated as the minimizer of the following primal problem (see [43]):

$$\left|\begin{array}{lrcl} \text{minimize} & \mathcal{I}_h(p) & := & \displaystyle\int_I h(p(x))dx \\ \\ \text{s.t.} & 1 & = & \displaystyle\int_I p(x)dx \\ \\ & d_j & = & \displaystyle DC(T)\int_I c_j(x)p(x)dx \end{array}\right. \tag{4.1}$$

where

$$c_j(x) = (x - k_j)^+ = \max\{0, x - k_j\} \tag{4.2}$$

is the payoff of the jth option, $h : \mathbb{R} \to \mathbb{R}$ is a convex function representing the entropy functional, and $DC(T)$ is the *riskless discount factor* up to time T.

Given the expected return r, $DC(T)$ is often taken to be equal to e^{-rT}. Throughout the thesis we assume, without loss of generality, that $DC(T) = 1$. We also assume that $p \in L^1(I)$. The main result will be stated for the entropy functional

$$h(t) := \begin{cases} t \log t - t, & \text{if } t > 0; \\ 0, & \text{if } t = 0; \\ +\infty, & \text{if } t < 0; \end{cases} \tag{4.3}$$

which is known as the *Boltzmann-Shannon entropy*.

The minimizer of problem (4.1) is referred to as the maximum entropy solution (MES). Strictly speaking, the maximum entropy principle in this context gives rise to the maximization of $\mathcal{I}_{-h}(p) = -\mathcal{I}_h(p)$.

Buchen and Kelley [36] and Avellaneda [38] further extend problem (4.1) to the case where the *Kullback's cross-entropy function*,

$$\int_I p(x) \log \left(\frac{p(x)}{q(x)} \right) dx,$$

is minimized. In this way, if it is known a priori that the asset price distribution should look like $q(x)$, this information can be incorporated into the method.[3] However, the derivation of an explicit dual function of finitely many variables can only be carried out when the Boltzmann-Shannon entropy is used [43]. In this and the subsequent chapters, we restrict ourselves to the use of (4.3) in problem (4.1).

4.3 A method to calculate the maximum entropy solution

We summarize in three steps the method described by Borwein et al. [43] to find a solution to problem (4.1) with h given by (4.3).

STEP 1. Calculate matrices required for the constraint qualification (CQ).

[3] For details on the principle of minimum cross entropy, see Cover and Thomas [44].

Let $\mathbf{u}$ be a vector of length $m-1$ whose entries are all ones, and let N be the $(m-1)\times(m-1)$ matrix given by

$$N := \begin{pmatrix} k_2-k_1 & \cdots & k_m-k_1 \\ & \ddots & \vdots \\ & & k_m-k_{m-1} \end{pmatrix}.$$

Let B and B_K be the $(m-1)\times m$ matrices given by

$$B := \begin{pmatrix} 1 & & -1 \\ & \ddots & \vdots \\ & 1 & -1 \end{pmatrix}, \qquad B_K := \begin{pmatrix} 1 & & -\frac{K-k_1}{K-k_m} \\ & \ddots & \vdots \\ & 1 & -\frac{K-k_{m-1}}{K-k_m} \end{pmatrix}.$$

STEP 2. Verify the CQ under which the problem has a unique solution.

In the case of $K=+\infty$, verify whether $\mathbf{d}=(d_1,\ldots,d_m)^T$ satisfies

$$d_m>0, \quad N^{-1}B\mathbf{d}>\mathbf{0}, \quad \text{and} \quad \langle N^{-1}B\mathbf{d},\mathbf{u}\rangle<1. \tag{4.4}$$

In the case of $0<K<\infty$, verify whether $\mathbf{d}$ satisfies

$$d_m>0, \quad N^{-1}B_K\mathbf{d}>\mathbf{0}, \quad \text{and} \quad \langle N^{-1}B_K\mathbf{d},\mathbf{u}\rangle<1. \tag{4.5}$$

STEP 3. Calculate the maximum entropy solution.

If the CQ is satisfied, then problem (4.1), with h given by (4.3), has a unique solution, namely

$$p_{\boldsymbol{\lambda}^*}(x)=\frac{1}{Z(\boldsymbol{\lambda}^*)}e^{\sum_{i=1}^m \lambda_i^* c_i(x)}, \tag{4.6}$$

where $\boldsymbol{\lambda}^*=(\lambda_1^*,...,\lambda_m^*)^T\in\mathbb{R}^m$ is the unique solution to the problem

$$\min\left\{\log Z(\boldsymbol{\lambda})-\sum_{i=1}^m \lambda_i d_i\right\}. \tag{4.7}$$

Here,

$$Z(\boldsymbol{\lambda}):=\int_0^K \exp\left[\sum_{i=1}^m \lambda_i (x-k_i)^+\right]dx=\sum_{j=1}^m e^{-\nu_j}\frac{e^{\mu_j k_{j+1}}-e^{\mu_j k_j}}{\mu_j} \tag{4.8}$$

in which $k_{m+1}:=K$, $\nu_j:=\sum_{i=1}^j \lambda_i k_i$, and $\mu_j:=\sum_{i=1}^j \lambda_i$. The expression $\mu_j^{-1}(\exp(\mu_j k_{j+1})-\exp(\mu_j k_j))$ is understood to be $k_{j+1}-k_j$ when $\mu_j=0$.

This completes the method. An important fact also shown by Borwein et al. [43], which we use during the analysis, is that the minimizer $\boldsymbol{\lambda}^*$ is unique and has the property that $\sum_{i=1}^{m} \lambda_i^* < 0$ whenever $K = +\infty$. Notice also that the integrand in (4.8) is always positive, so $Z(\boldsymbol{\lambda}) > 0$.

Buchen and Kelley [36] and Avellaneda [38] also arrive to (4.6) and (4.7) using the theory of Lagrange multipliers. The CQ, however, is only obtained within the framework of partially finite programming. Originally, the CQ states that for the MES to exist, the data must lie in the relative interior of the set of points $\mathbf{d}$ such that, given p with $\mathcal{I}_h(p)$ finite, $d_j = \int_I c_j(x)p(x)dx$ [43]. This set is an open polyhedral, and (4.4) ((4.5)) is an effective characterization of that set. Note, on the other hand, that the feasible set of problem (4.1) is not relatively open. In fact, there can be feasible boundary points that fail to satisfy the CQ for which the analysis via the Lagrange multipliers will fail [43].

Chapter 5

Stability of the MES under parameter perturbation

5.1 Continuous dependence of the MES on $\boldsymbol{\lambda}$

Avellaneda [38] and Buchen and Kelley [36] show that

$$\begin{aligned}(\log Z(\boldsymbol{\lambda}))_{\lambda_i,\lambda_j} &= \frac{Z_{\lambda_i,\lambda_j}}{Z} - \frac{Z_{\lambda_i} Z_{\lambda_j}}{Z^2} \\ &= \mathrm{Cov}_{p_{\boldsymbol{\lambda}}}\left[c_i(x), c_j(x)\right];\end{aligned}$$

and since covariance matrices are nonnegative definite, the objective function in (4.7), $\log Z(\boldsymbol{\lambda}) - \sum_{i=1}^{m} \lambda_i d_i$ is convex. Moreover, strict convexity follows if the m payoff functions c_j are linearly independent. Avellaneda [38] further proves via convex duality that, assuming linear independence of the payoff functions, stability of the minimizer of the objective in (4.7) is guaranteed. However, as Buchen and Kelley [36] observe, linearly independent payoff functions can be close enough to linear dependence in some sense when strike prices are close to each other. This causes the Hessian of $\log Z(\boldsymbol{\lambda}) - \sum_{i=1}^{m} \lambda_i d_i$ to be poorly conditioned so that the corresponding minimizer may become unstable to perturbations in the data. We analyze this issue in more detail in the next chapter.

In this chapter, we show that the MES is stable to perturbations of the minimizers of the objective in (4.7). Hence, it suffices to regularize the dual problem (i.e. the problem of solving (4.7)) in order to obtain a MES that is stable to perturbations in the data.

In fact, it is clear that the MES, which is available analytically as a function of the minimizers $\boldsymbol{\lambda}$ of the dual problem (4.7), is differentiable with respect to $\boldsymbol{\lambda}$. Hence, continuous dependence of the MES on $\boldsymbol{\lambda}$ is guaranteed. Nonetheless, for a more thorough stability analysis, we find a Lipschitz constant L such that for MESs $p_{\boldsymbol{\lambda}}$ and $p_{\bar{\boldsymbol{\lambda}}}$ given as in (4.6), we have that $\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 \leq L\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|$.[1]

5.2 A Lipschitz constant

Before calculating a Lipschitz constant L such that that $\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 \leq L\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|$, we state and prove the following result.

Proposition 5.2.1. *Let $\boldsymbol{\lambda} \in \mathbb{R}^m$ and define $p_{\boldsymbol{\lambda}} : [0, K) \to \mathbb{R}$ by*

$$p_{\boldsymbol{\lambda}}(x) = \frac{1}{Z(\boldsymbol{\lambda})} e^{\sum_{i=1}^m \lambda_i c_i(x)}, \tag{5.1}$$

where either $K < +\infty$ or $K = +\infty$, $Z(\boldsymbol{\lambda})$ is given by (4.8), and $c_i(x)$ is given by (4.2). Given $p_{\boldsymbol{\lambda}}$ and $p_{\bar{\boldsymbol{\lambda}}}$,

$$\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 \leq \left[\frac{2}{Z(\boldsymbol{\lambda})} \sum_{j=1}^m M_j\right] \left\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\right\|_1, \tag{5.2}$$

where, for $j = 1, \ldots, m-1$ and arbitrary K, or $j = m$ if K is finite,

$$M_j = e^{y_j}(w_j + k_j)(k_{j+1} - k_j),$$

$$y_j = \max\{-\nu_j + \mu_j w_j, -\bar{\nu}_j + \bar{\mu}_j w_j\}, \tag{5.3}$$

$$\nu_j := \sum_{i=1}^j \lambda_i k_i, \quad \mu_j := \sum_{i=1}^j \lambda_i, \quad \bar{\nu}_j := \sum_{i=1}^j \bar{\lambda}_i k_i, \quad \bar{\mu}_j := \sum_{i=1}^j \bar{\lambda}_i, \tag{5.4}$$
$$w_j := \arg\max \left|e^{-\nu_j + \mu_j x} - e^{-\bar{\nu}_j + \bar{\mu}_j x}\right|, k_j \leq x \leq k_{j+1},$$

k_j are the strike prices, $k_{m+1} := K$; and for $K = +\infty$,

$$M_m := |\beta|^{-2} e^{\alpha + \beta k_m}(2|\beta| k_m + 1),$$

with $\alpha := \max\{-\nu_m, -\bar{\nu}_m\}$, $\beta := \max\{\mu_m, \bar{\mu}_m\}$.

[1] In the case where $K = +\infty$, adding the condition that $\sum_{i=1}^m \lambda_i < 0$ guarantees that $p_{\boldsymbol{\lambda}}$ is in L^1 [43]. In addition, from (4.8), we see that $\int_0^K p_{\boldsymbol{\lambda}}(x)dx$ is indeed equal to 1.

Proof. We first prove the case where K is finite. By definition,

$$\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 = \int_0^K \left| \frac{1}{Z(\boldsymbol{\lambda})} e^{\sum_{i=1}^m \lambda_i c_i(x)} - \frac{1}{Z(\bar{\boldsymbol{\lambda}})} e^{\sum_{i=1}^m \bar{\lambda}_i c_i(x)} \right| dx.$$

Now we factor the common denominator out of the integral and add the term $0 = Z(\bar{\boldsymbol{\lambda}})e^{\bar{\boldsymbol{\lambda}}^T \mathbf{c}(x)} - Z(\bar{\boldsymbol{\lambda}})e^{\bar{\boldsymbol{\lambda}}^T \mathbf{c}(x)}$ inside the absolute value, where $\mathbf{c}(x) = (c_1(x), \ldots, c_m(x))^T$. Then we use the triangle inequality. Next, we replace $Z(\boldsymbol{\lambda})$ in the numerator with its definition (4.8) and use the fact that the absolute value of the integral is no greater than the integral of the absolute value. We then arrive to

$$\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 \leq \frac{2}{Z(\boldsymbol{\lambda})} \int_0^K \left| e^{\boldsymbol{\lambda}^T \mathbf{c}(x)} - e^{\bar{\boldsymbol{\lambda}}^T \mathbf{c}(x)} \right| dx. \tag{5.5}$$

Using (4.2), and the notation given by (5.4), we further notice that

$$\int_0^K \left| e^{\boldsymbol{\lambda}^T \mathbf{c}(x)} - e^{\bar{\boldsymbol{\lambda}}^T \mathbf{c}(x)} \right| dx = \sum_{j=1}^m \int_{k_j}^{k_{j+1}} \left| e^{-\nu_j + \mu_j x} - e^{-\bar{\nu}_j + \bar{\mu}_j x} \right| dx. \tag{5.6}$$

Let $\varphi_j(x) := e^{-\nu_j + \mu_j x} - e^{-\bar{\nu}_j + \bar{\mu}_j x}$ for all x. From calculus, we see that $|\varphi(x)|$ has one local maximum if $\mu_j \neq \bar{\mu}_j$, or none if $\mu_j = \bar{\mu}_j$. Hence, $|\varphi(x)| \leq |\varphi(w_j)|$ on the interval $[k_j, k_{j+1}]$, where w_j is either k_j, or k_{j+1}, or the local maximizer if it is between k_j and k_{j+1}. Therefore,

$$I_j := \int_{k_j}^{k_{j+1}} \left| e^{-\nu_j + \mu_j x} - e^{-\bar{\nu}_j + \bar{\mu}_j x} \right| dx \leq |\varphi_j(w_j)| \, (k_{j+1} - k_j). \tag{5.7}$$

Now, we use the fact that, for $a < b$, $\frac{|e^b - e^a|}{|b-a|} \leq \frac{d}{dx} e^x \big|_{x=b}$ to obtain

$$I_j \leq |-\nu_j + \mu_j w_j - (-\bar{\nu}_j + \bar{\mu}_j w_j)| \, e^{y_j} (k_{j+1} - k_j)$$

(cf. (5.3)). Back in terms of $\boldsymbol{\lambda}$ and $\bar{\boldsymbol{\lambda}}$, this leads to

$$I_j \leq \sum_{i=1}^m \left| \bar{\lambda}_i - \lambda_i \right| (w_j + k_j) e^{y_j} (k_{j+1} - k_j). \tag{5.8}$$

Estimate (5.2) follows from (5.5), (5.6) and (5.8).

To prove the case where $K = +\infty$ it remains to estimate I_m. Let $\alpha := \max\{-\nu_m, -\bar{\nu}_m\}$, $\beta := \max\{\mu_m, \bar{\mu}_m\}$, $\bar{\alpha} := \min\{-\nu_m, -\bar{\nu}_m\}$, $\bar{\beta} := \min\{\mu_m, \bar{\mu}_m\}$. It follows that $|\varphi_m(x)| \leq$

$e^{\alpha+\beta x} - e^{\bar{\alpha}+\bar{\beta}x}$. After direct integration of the right hand side and the use of the triangle inequality, we have that

$$I_m \leq \frac{1}{|\bar{\beta}|}\left|e^{\alpha+\beta x} - e^{\bar{\alpha}+\bar{\beta}k_m}\right| + \left|\frac{1}{\bar{\beta}} - \frac{1}{\beta}\right| e^{\alpha+\beta k_m}.$$

Much in the same fashion as after (5.7), we obtain

$$I_m \leq |\beta|^{-2}\, e^{\alpha+\beta k_m}(2\,|\beta|\,k_m + 1)\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|_1.$$

□

We now derive the Lipschitz constant from the above result.

Corollary 5.2.1. *Consider $p_{\boldsymbol{\lambda}}$ and $p_{\bar{\boldsymbol{\lambda}}}$ as in* (5.1)*, but with the additional assumption that $\boldsymbol{\lambda}, \bar{\boldsymbol{\lambda}} \in [-a, a]^m$ for fixed $a > 0$. If $K = +\infty$, assume also that $\mu_m, \bar{\mu}_m \leq -b$ for fixed $b > 0$. Then $\|p_{\boldsymbol{\lambda}} - p_{\bar{\boldsymbol{\lambda}}}\|_1 \leq L\|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|_1$ with*

$$L := \frac{2}{Z_{min}} \sum_{j=1}^{m} L_j, \tag{5.9}$$

where

$$Z_{min} := \sum_{j=1}^{m} e^{a\sum_{i=1}^{j} k_i} \frac{e^{-ajk_j} - e^{-ajk_{j+1}}}{aj}, \tag{5.10}$$

$$L_j := e^{a(\sum_{i=1}^{j} k_i + jk_{j+1})}\,(k_{j+1} + k_j)\,(k_{j+1} - k_j)\,,$$

for $j = 1, \ldots, m-1$ and $j = m$ if K is finite, and

$$L_m := b^{-2} e^{a(\sum_{i=1}^{m} k_i + mk_m)}\,(2amk_m + 1)$$

if $K = +\infty$.

Proof. L given by (5.9) follows from (5.2) if $Z(\boldsymbol{\lambda}) \geq Z_{min}$ and $M_j \leq L_j$. It follows directly from the assumptions that $M_j \leq L_j$ for all j and arbitrary K. To prove that $Z(\boldsymbol{\lambda}) \geq Z_{min}$ given by (5.10), notice that

$$Z(\boldsymbol{\lambda}) = \sum_{j=1}^{m} \int_{k_j}^{k_{j+1}} \exp\left[\sum_{i=1}^{j} \lambda_i (x - k_i)\right] dx. \tag{5.11}$$

Let $g_j(\lambda_1, \ldots, \lambda_j) = \sum_{i=1}^{j} \lambda_i(x - k_i)$ with $k_j < x < k_{j+1}$. Then $\frac{\partial g_j}{\partial \lambda_i} = x - k_i > x - k_j > 0$. That is, g_j is an increasing function of λ_i. Hence, when $\boldsymbol{\lambda} \in [-a, a]^m$, $\exp\left[\sum_{i=1}^{j} \lambda_i(x - k_i)\right] \geq \exp\left[\sum_{i=1}^{j}(-a)(x - k_i)\right]$ for $k_j < x < k_{j+1}$. Direct integration of the latter term for each j and (5.11) lead to (5.10). □

This result confirms the stability of the MES to perturbations on $\boldsymbol{\lambda}$, but in addition, the Lipschitz constant (5.9) shows that stability remains even if more option prices with strike prices closer to each other are used. This is because of the proportionality of L_j to $(k_{j+1} - k_j)$.

Chapter 6

Instability of the dual minimizer under data perturbation

6.1 The dual cost function

Problem (4.7) is known as the dual of problem (4.1) [43]. In this section we view the dual functional as a function D of $\boldsymbol{\lambda}$ and $\mathbf{d}$,

$$D(\boldsymbol{\lambda}, \mathbf{d}) := \log Z(\boldsymbol{\lambda}) - \mathbf{d}^T \boldsymbol{\lambda}. \tag{6.1}$$

Given data $\mathbf{d}^*$, the minimizer $\boldsymbol{\lambda}^*$ of D is unique [43].

At $\boldsymbol{\lambda}^*$ the gradient vanishes, that is,

$$h(\boldsymbol{\lambda}^*, \mathbf{d}^*) := \nabla_{\boldsymbol{\lambda}} D(\boldsymbol{\lambda}^*, \mathbf{d}^*) = \frac{1}{Z(\boldsymbol{\lambda}^*)} \nabla Z(\boldsymbol{\lambda}^*) - \mathbf{d}^* = \mathbf{0}. \tag{6.2}$$

If we assume that the Jacobian $\nabla_{\boldsymbol{\lambda}} h$ of h with respect to $\boldsymbol{\lambda}$ is nonsingular at $(\boldsymbol{\lambda}^*, \mathbf{d}^*)$, then, by the implicit function theorem, there exist open sets $\mathcal{N}_{\boldsymbol{\lambda}} \subset \mathbb{R}^m$ and $\mathcal{N}_{\mathbf{d}} \subset \mathbb{R}^m$ containing $\boldsymbol{\lambda}^*$ and $\mathbf{d}^*$, respectively, and a continuous function $\boldsymbol{\lambda} : \mathcal{N}_{\mathbf{d}} \to \mathcal{N}_{\boldsymbol{\lambda}}$ such that $\boldsymbol{\lambda}^* = \boldsymbol{\lambda}(\mathbf{d}^*)$ and $h(\boldsymbol{\lambda}(\mathbf{d}), \mathbf{d}) = 0$ for all $\mathbf{d} \in \mathcal{N}_{\mathbf{d}}$. Moreover,

$$[\nabla_{\boldsymbol{\lambda}} h(\boldsymbol{\lambda}^*, \mathbf{d}^*)]^T \nabla \boldsymbol{\lambda}(\mathbf{d}^*) + \nabla_{\mathbf{d}} h(\boldsymbol{\lambda}^*, \mathbf{d}^*) = 0, \tag{6.3}$$

[45] (See, for example, the appendix of reference [46]).

Hence, given data $\mathbf{d}$ and $\mathbf{d}^\delta$ with $\|\mathbf{d}^\delta - \mathbf{d}\| \leq \delta$, and corresponding minimizers $\boldsymbol{\lambda}(\mathbf{d})$ and $\boldsymbol{\lambda}(\mathbf{d}^\delta)$ of D, as long as it is true that $\nabla_{\boldsymbol{\lambda}} h(\boldsymbol{\lambda}, \mathbf{d})$ is nonsingular, we can speak of a

first order Taylor approximation

$$\boldsymbol{\lambda}(\mathbf{d}^\delta) - \boldsymbol{\lambda}(\mathbf{d}) = [\nabla \boldsymbol{\lambda}(\mathbf{d})]^T \left(\mathbf{d}^\delta - \mathbf{d}\right) + \|\mathbf{d}^\delta - \mathbf{d}\| E(\mathbf{d}, \delta),$$

where $E(\mathbf{d}, \delta) \to 0$ as $\delta \to 0$; from which it follows that

$$\left\|\boldsymbol{\lambda}(\mathbf{d}^\delta) - \boldsymbol{\lambda}(\mathbf{d})\right\| = O\left(\left\|[\nabla \boldsymbol{\lambda}(\mathbf{d})]^T\right\| \|\mathbf{d}^\delta - \mathbf{d}\|\right). \tag{6.4}$$

From (6.3), we can solve for the Jacobian $\nabla \boldsymbol{\lambda}(\mathbf{d})$ of $\boldsymbol{\lambda}(\mathbf{d})$,

$$\nabla \boldsymbol{\lambda}(\mathbf{d}) = -\left\{[\nabla_{\boldsymbol{\lambda}} h(\lambda, \mathbf{d})]^T\right\}^{-1} \nabla_{\mathbf{d}} h(\lambda, \mathbf{d}).$$

From (6.2), we see that $\nabla_{\mathbf{d}} h(\lambda, \mathbf{d})$ is equal to the negative of the identity matrix, so that

$$[\nabla \boldsymbol{\lambda}(\mathbf{d})]^T = [\nabla_{\boldsymbol{\lambda}} h(\lambda, \mathbf{d})]^{-1}. \tag{6.5}$$

Furthermore, the Jacobian $\nabla_{\boldsymbol{\lambda}} h(\boldsymbol{\lambda}, \mathbf{d})$ of h is the Hessian $\nabla^2_{\boldsymbol{\lambda}} D(\boldsymbol{\lambda}, \mathbf{d})$ of D,

$$\nabla_{\boldsymbol{\lambda}} h(\boldsymbol{\lambda}, \mathbf{d}) = \nabla^2_{\boldsymbol{\lambda}} D(\boldsymbol{\lambda}, \mathbf{d}). \tag{6.6}$$

Combining (6.5) and (6.6) and substituting into (6.4), we obtain

$$\left\|\boldsymbol{\lambda}(\mathbf{d}^\delta) - \boldsymbol{\lambda}(\mathbf{d})\right\| = O\left(\left\|\left[\nabla^2_{\boldsymbol{\lambda}} D(\boldsymbol{\lambda}, \mathbf{d})\right]^{-1}\right\| \|\mathbf{d}^\delta - \mathbf{d}\|\right).$$

Thus, the norm of the inverse of the Hessian of the dual objective, if it exists, can be used as an indicator of how far apart $\boldsymbol{\lambda}(\mathbf{d}^\delta)$ and $\boldsymbol{\lambda}(\mathbf{d})$ might be. Given a specific numerical example, we can compute the Hessian of the dual objective. A numerically noninvertible Hessian, or an invertible Hessian with an inverse whose norm is large, would be evidence of the ill-conditioning of the dual of the option-maximum entropy problem.

In the following section we analyze some numerical examples to observe the behavior of both the aforementioned Hessian and the MESs. The explicit function provided by Borwein et al.'s method allows for the symbolic calculation of the Hessian, which can then be supplied to the Newton's method for the search of the minimizer of the unconstrained dual problem. For reference, we state the Hessian of D,

$$\nabla^2_{\boldsymbol{\lambda}} D(\boldsymbol{\lambda}, \mathbf{d}) = -\frac{1}{Z^2} \nabla Z \, (\nabla Z)^T + \frac{1}{Z} \nabla^2 Z,$$

where

$$Z = Z(\boldsymbol{\lambda}(\mathbf{d})), \qquad \nabla Z = \left[\frac{\partial Z}{\partial \lambda_i}\right]_{m\times 1}, \qquad \nabla^2 Z = \left[\frac{\partial^2 Z}{\partial \lambda_l \partial \lambda_i}\right]_{m\times m}.$$

The partial derivatives of Z of first and second order with respect to λ_i are explicitly as follows,

$$\frac{\partial Z}{\partial \lambda_i} = \sum_{j=i}^{m} e^{-\nu_j} \mu_j^{-1} \left[-e^{\mu_j k_{j+1}} \left(k_i - k_{j+1} + \mu_j^{-1} \right) + e^{\mu_j k_j} \left(k_i - k_j + \mu_j^{-1} \right) \right],$$

$$\begin{aligned} \frac{\partial^2 Z}{\partial \lambda_l \partial \lambda_i} = & \sum_{j=max\{i,l\}}^{m} e^{-\nu_j} \mu_j^{-1} \\ & \times \left\{ e^{\mu_j k_{j+1}} \left[\left(k_i - k_{j+1} + \mu_j^{-1} \right) \left(k_l - k_{j+1} + \mu_j^{-1} \right) + \mu_j^{-2} \right] \right. \\ & \left. - e^{\mu_j k_j} \left[\left(k_i - k_j + \mu_j^{-1} \right) \left(k_l - k_j + \mu_j^{-1} \right) + \mu_j^{-2} \right] \right\}. \end{aligned}$$

6.2 Analysis of numerical examples

It is a nontrivial task to model stock price behavior. The most celebrated model is that by Black and Scholes, which is discussed on any reference dedicated to derivative securities (see e.g. [47]). They model stock price behavior as geometric Brownian motion [47]. This model requires two parameters: the expected return r and the volatility σ. Given the price d_1 of a stock today, this model implies that the probability distribution of the stock price x after T years is log-normal [47]. More specifically,

$$\log x \sim N \left[\log d_1 + \left(r - \frac{\sigma^2}{2} \right) T, \sigma^2 T \right],$$

where $N(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2. Typical values of σ range from 0.15 to 0.60 [47].

For our numerical examples, all we need is a specific model which will give a prescribed RND to solve the direct problem of valuating European call options. We assumed a stock price distribution given as above. We considered a current stock price of $d_1 = 40$ dollars, a time in the future of $T = 0.5$ years, an expected return of $r = 0.16$, and two different values of volatility, namely, $\sigma = 0.5$ and $\sigma = 0.2$. Figure 6.1 shows the probability distributions in both cases. We use the value σ as a label to identify each distribution.

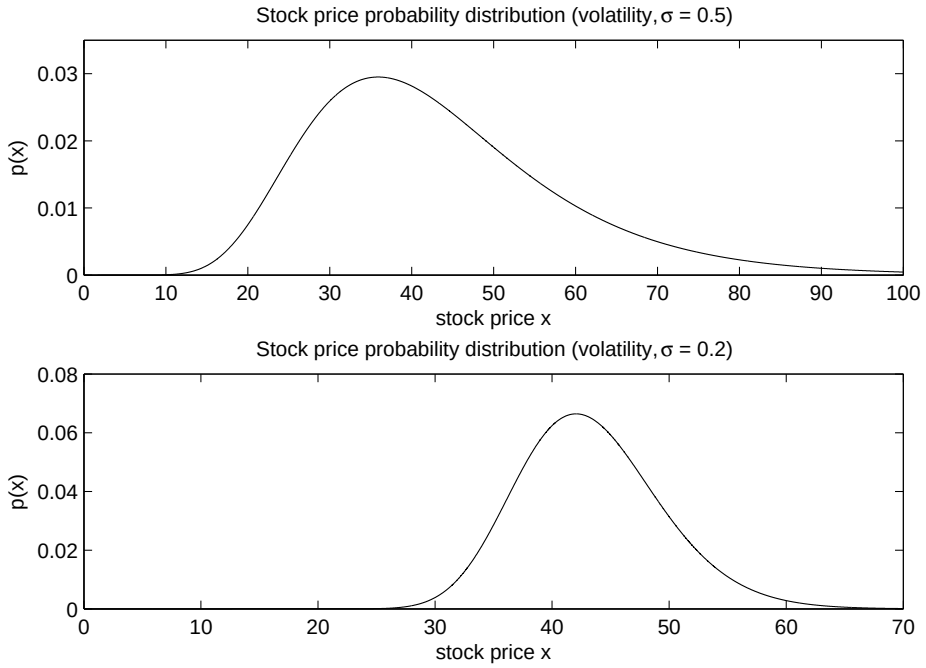

Figure 6.1: *Log-normal stock price probability distributions obtained when modeling the stock price behavior as geometric Brownian motion. The values used for these distributions are a current stock price of* $d_1 = 40$ *dollars, a time in the future of* $T = 0.5$ *years, and an expected return of* $r = 0.16$*. The upper subplot shows the distribution when the volatility of the Brownian motion is* $\sigma = 0.5$*. The lower subplot shows the distribution when the volatility is* $\sigma = 0.2$*.*

Given a set of different strike prices k_j and a prescribed stock price density function $p(x)$, the prices d_j of European call options are calculated as the risk neutral expected payoffs (see the second constraint of the option-maximum entropy problem (4.1)),

$$d_j = e^{-rT} \int_0^\infty c_j(x) p(x) dx. \tag{6.7}$$

The first four examples that we considered were generated from the density function where the volatility is $\sigma = 0.5$, and vectors of strike prices that are as follows:

$$\mathbf{k} = [0, 20, 40, 60, 80],$$
$$\mathbf{k} = [0, 10, 20, \ldots, 70, 80],$$
$$\mathbf{k} = [0, 5, 10, \ldots, 75, 80],$$
$$\mathbf{k} = [0, 4, 8, \ldots, 76, 80].$$

With these, we obtained, respectively, 5, 9, 17, and 21 data d_j. The volatility σ and the number m of data suffice then to label each example. Figure 6.2 shows the values of the option prices plotted against their respective strike prices.

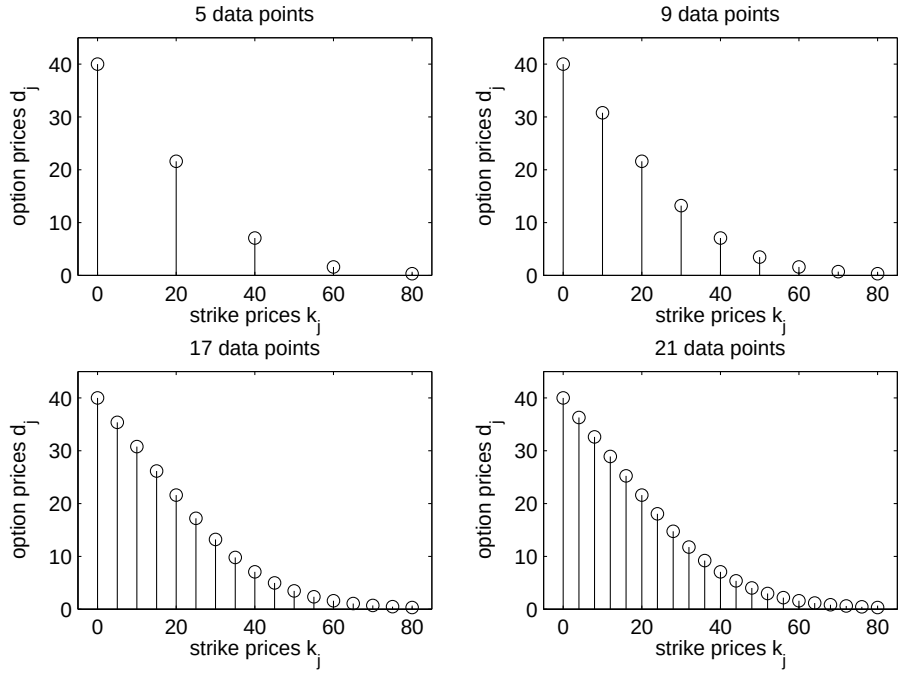

Figure 6.2: *Plots of option prices d_j versus their respective strike prices k_j. The prices d_j are calculated according to equation* (6.7)*, where $r = 0.16$ and $T = 0.5$. The stock price distribution $p(x)$ used is shown in the upper subplot of figure 6.1, where $\sigma = 0.5$.*

When the volatility is $\sigma = 0.2$ the vectors of strike prices used were as follows:

$$\mathbf{k} = [0, 25, 30, 40, 50, 60],$$
$$\mathbf{k} = [0, 25, 30, 35, \ldots, 55, 60],$$
$$\mathbf{k} = [0, 25, 27.5, 30, \ldots, 57.5, 60],$$
$$\mathbf{k} = [0, 25, 26, 28, 30, \ldots, 58, 60].$$

The respective number of data were 6, 9, 16, and 20, and the option prices are shown in figure 6.3.

When the second density function is assumed, prices on the range $(0, 25)$ are unlikely. Hence, no strike prices are considered on that range. In fact, for the second density

function, option prices based upon strike prices in $(0, 25)$ do not satisfy the constraint qualification provided by Borwein et al. [43]

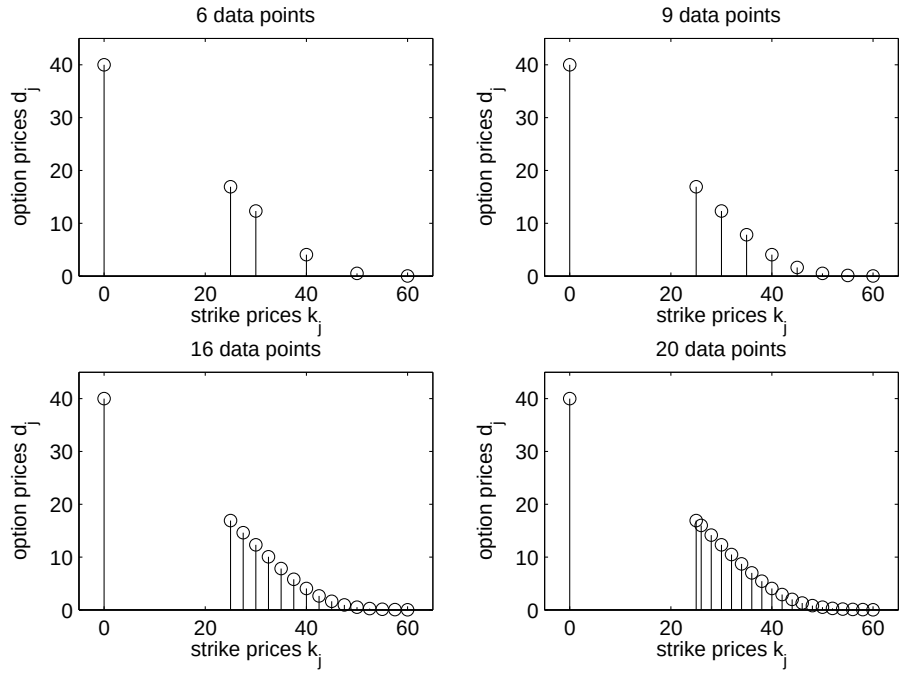

Figure 6.3: *Plots of option prices d_j versus their respective strike prices k_j. The prices d_j are calculated according to equation* (6.7)*, where $r = 0.16$ and $T = 0.5$. The stock price distribution $p(x)$ used is shown in the lower subplot of figure 6.1, where $\sigma = 0.2$.*

Once the option prices were calculated, we obtained the maximum entropy solution given by eq. (4.6) by minimizing the dual objective given in problem (4.7). For this minimization problem we used `MATLAB`'s built-in function `fminunc`. We first used a quasi-Newton line search method to get closer to the minimizer. Then we switched to Newton's method, which requires the inverse of the Hessian of the dual objective. We accepted a numerical solution until the decrease in the value of the objective function was less than 10^{-12}. Under this criterion, a numerical solution was found only for those examples with 9 or less data points. For the examples with more than 9 data points, the Hessian became singular to numerical precision before the method was able to find an acceptable solution. In those cases, we accepted as a minimizer the solution obtained a few iterations before numerical singularity of the Hessian occurred. Figures 6.4 and 6.5

show the maximum entropy solutions found for $\sigma = 0.5$ and $\sigma = 0.2$, respectively. Note the presence of "spikes" near $x = 0$ in both figures 6.4 and 6.5. These spikes reach a high value and appear in the bottom two subfigures of figure 6.4 as well, although their supports are narrower. We believe this is an artifact of instability since they disappear under regularization (see the next chapter).

Let us denote the Hessian of the dual objective by H. Table 6.1 shows the 2-norm of the inverse of the Hessian evaluated at the accepted numerical minimizers. Only the first of all of these examples shows a small norm of the inverse, which suggests a stable solution. As m increases, the solution becomes more unstable until it is impossible to be found numerically.

$\sigma = 0.5$		$\sigma = 0.2$	
m	$\|H^{-1}\|_2$	m	$\|H^{-1}\|_2$
5	0.6423	6	4.0461e+03
9	204.9889	9	4.0637e+03
17	4.1309e+05	16	2.3522e+04
21	1.8326e+09	20	1.5791e+08

Table 6.1: *2-norm of the inverse of the Hessian obtained for each of the eight examples. The first (third) column shows the number m of data points used to recover the first (second) stock price distribution plotted in figure 6.1. The second (fourth) column shows the 2-norm of the inverse of the Hessian H of the dual objective evaluated at the minimizers obtained when recovering the first (second) stock price distribution.*

So far we have used only noiseless data. However, in the real world, there are inherent uncertainties, which can be interpreted as noise in the data. For option prices, we usually have only the bid and the ask prices, and these are at best accurate to a few decimal places, so there is always a difference between the estimated option price (say, the midpoint between bid and ask price) and the unavailable true value. There could be additional unknown sources of error. For such reasons, we would like to quantify how much the MESs are altered when using noisy data instead of noiseless data. Let $\mathbf{d}$ denote the noiseless data, and let $\bar{\mathbf{d}} := \mathbf{d} + \boldsymbol{\eta}$ denote the noisy data, where $\boldsymbol{\eta}$ is some noise. Let $\boldsymbol{\lambda}$ and $\bar{\boldsymbol{\lambda}}$ denote the minimizers of the dual obtained when using data $\mathbf{d}$ and $\bar{\mathbf{d}}$ respectively. For our numerical examples, we chose the noise so that the ratio $\|\mathbf{d} - \bar{\mathbf{d}}\|_2/\|\mathbf{d}\|_2$ was 5%. Componentwise, $|d_j - \bar{d}_j|/|d_j|$ oscillated between 0 and 20%.

We considered the examples with 16 or more data points. The top subplots of figures 6.6 and 6.7 show $(\bar{d}_j - d_j)/d_j$ plotted against the strike prices in the examples considered when $\sigma = 0.5$ and $\sigma = 0.2$, respectively.

It turns out that, for some values of j, $\bar{d}_j$ does not satisfy the CQ (cf. (4.4)). Even when the ratio $\|\mathbf{d} - \bar{\mathbf{d}}\|_2/\|\mathbf{d}\|_2$ is as low as 10^{-4}, some of the data fail to satisfy the CQ. When using the dual objective with these nonadmissible data, the minimization algorithm runs into difficulties, typically being trapped in a local minimum, and finding minimizers $\bar{\boldsymbol{\lambda}}$ with $\sum_{i=1}^{m} \bar{\lambda}_i > 0$ which in turn yield solutions $p_{\bar{\boldsymbol{\lambda}}}$ that are not even in L^1.

σ	number of data points	$\dfrac{\|\mathbf{d} - \bar{\mathbf{d}}\|_2}{\|\mathbf{d}\|_2}$	$\dfrac{\|p_{\boldsymbol{\lambda}} - p_{\tilde{\boldsymbol{\lambda}}}\|_1}{\|p_{\boldsymbol{\lambda}}\|_1}$
0.5	17	5.00%	111.52%
0.5	21	5.00%	114.58%
0.2	16	5.00%	104.51%
0.2	20	5.00%	90.18%

Table 6.2: *Error levels of the data and the MESs. The first column indicates the assumed volatility σ. The second column shows the number of data points. The third column shows the error level of the noisy data relative to the size of the noiseless data. The fourth column shows the error level of the MES $p_{\tilde{\boldsymbol{\lambda}}}$ with noisy data relative to the size of the MES $p_{\boldsymbol{\lambda}}$ with noiseless data.*

One possibility to overcome this difficulty would be to omit some data points until the remaining data satisfy the CQ. If we do this with our examples, where the original number of available data points is between 16 and 21, then we end up using only 10 to 11 data points, so a significant amount of available information is lost. A much better approach is to project the data onto the open polyhedral described by the CQ (4.4) (or (4.5) if K is finite). Given the noisy data $\bar{\mathbf{d}}$, and the matrices N, B, and $\mathbf{u}$ introduced in **STEP 1** of the method above, the projected data $\tilde{\mathbf{d}}$ can be obtained by solving the constrained linear least-squares problem of the form

$$\min_{\tilde{\mathbf{d}}} \|\tilde{\mathbf{d}} - \bar{\mathbf{d}}\|_2^2 \text{ such that } \begin{cases} N^{-1}B\tilde{\mathbf{d}} & \geq \epsilon\mathbf{u}, \\ \mathbf{u}^T N^{-1} B\tilde{\mathbf{d}} & \leq 1 - \epsilon, \end{cases}$$

where $\epsilon > 0$.

That is, we seek a vector $\tilde{\mathbf{d}}$ of data constrained to the open set described by the CQ which is as close as possible to the original data. We introduce ϵ since inequality constraints in linear programs cannot be strict. For our numerical examples, $\epsilon = 10^{-4}$ was suitable. The constraint $d_m > 0$ is usually satisfied, so we omitted it from the program. By using this preliminary step, we were able to use all of the available data. Local minima are removed and the iterations approach the unique minimizer, but ill-conditioning of the Hessian of the objective is still present. Hence, we again accepted as a minimizer the solution obtained a few iterations before numerical singularity of the Hessian occurred. Table 6.2 shows the error level $\|p_{\boldsymbol{\lambda}} - p_{\tilde{\boldsymbol{\lambda}}}\|_1$ of the MES $p_{\tilde{\boldsymbol{\lambda}}}$ with the projected noisy data $\tilde{\mathbf{d}}$ relative to the size of the MES $p_{\boldsymbol{\lambda}_\alpha}$ with noiseless data $\mathbf{d}$. The bottom two subplots of figures 6.6 and 6.7 show the graphs of the MESs $p_{\bar{\boldsymbol{\lambda}}}$ for each example when $\sigma = 0.5$ and $\sigma = 0.2$, respectively.

Projection of the noisy data onto the open polyhedral described by the CQ provide an additional benefit, as the error levels in the projected data are consistently and significantly reduced compared to the error levels in the original noisy data of our numerical examples. This is reported in table (6.3). As a consequence, this shows empirically that such projection of the data can be used advantageously as a preliminary step in any method which extracts information from option prices.

σ	m	$\frac{\|\mathbf{d} - \bar{\mathbf{d}}\|_2}{\|\mathbf{d}\|_2}$	$\frac{\|\mathbf{d} - \tilde{\mathbf{d}}\|_2}{\|\mathbf{d}\|_2}$	correction
0.5	17	5.00%	3.48%	30.44%
0.5	21	5.00%	0.83%	83.48%
0.2	16	5.00%	3.11%	37.85%
0.2	20	5.00%	2.93%	41.47%

Table 6.3: *Error level of the noisy data $\bar{\mathbf{d}}$ compared to the error level of the projection $\tilde{\mathbf{d}}$ of the noisy data onto the open polyhedral described by the CQ. The first column indicates the assumed volatility σ. The second column shows the number m of available noisy data points. The third column shows the error level of the noisy data $\bar{\mathbf{d}}$. The fourth column shows the error level of the projected noisy data $\tilde{\mathbf{d}}$. The fifth column shows how much the error level of the noisy data is decreased after projecting the noisy data.*

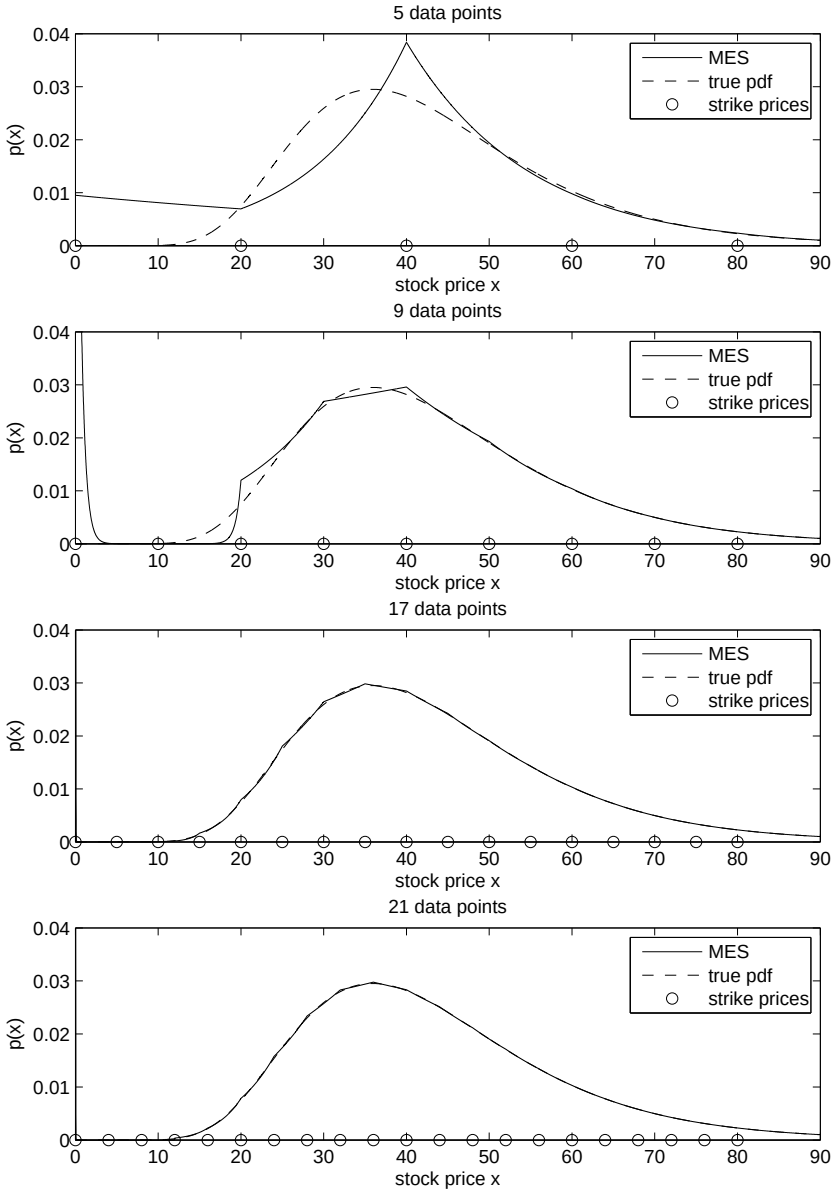

Figure 6.4: *Maximum entropy solutions to recover the first stock price distribution (where the volatility is $\sigma = 0.5$). The solutions are represented by the solid line, the true distribution is represented by the dashed line, the circles indicate the strike prices. From top to bottom, the solutions use 5, 9, 17, and 21 data points, respectively.*

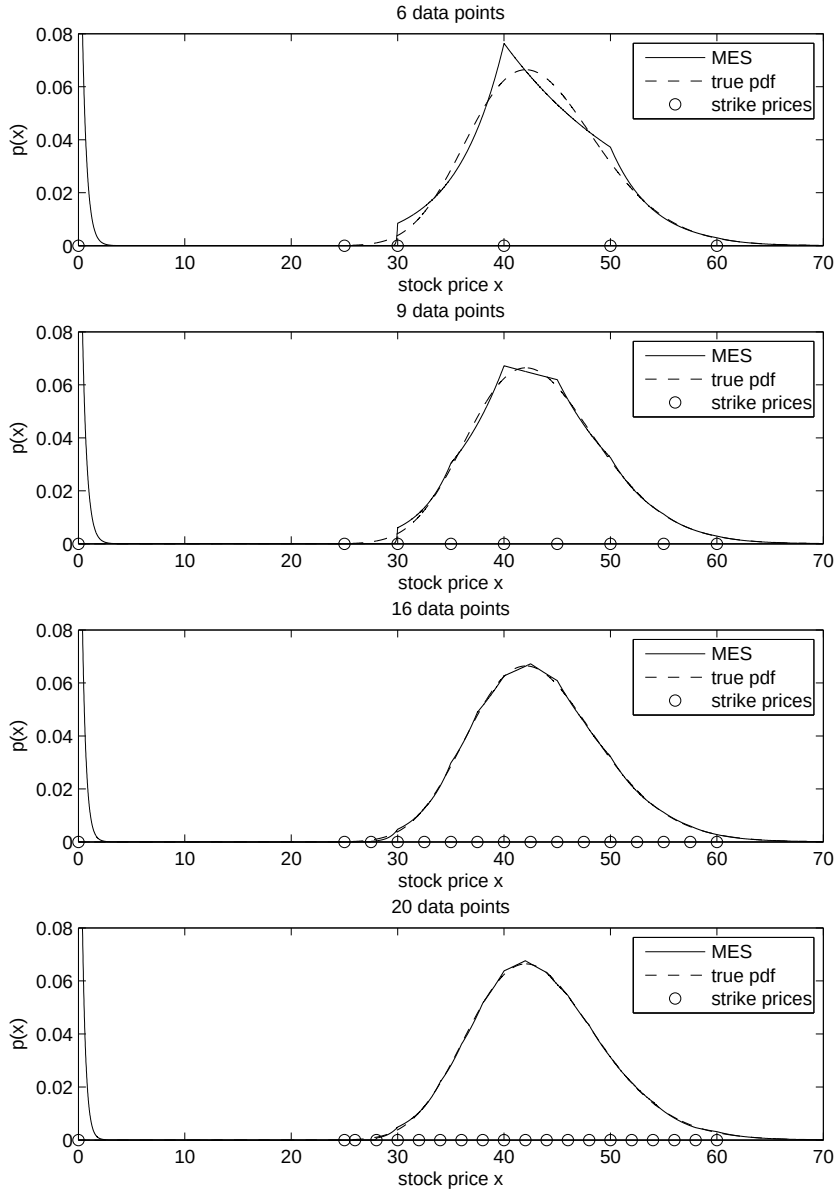

Figure 6.5: *Maximum entropy solutions to recover the second stock price distribution (where the volatility is $\sigma = 0.2$). The solutions are represented by the solid line, the true distribution is represented by the dashed line, the circles indicate the strike prices. From top to bottom, the solutions use 6, 9, 16, and 20 data points, respectively.*

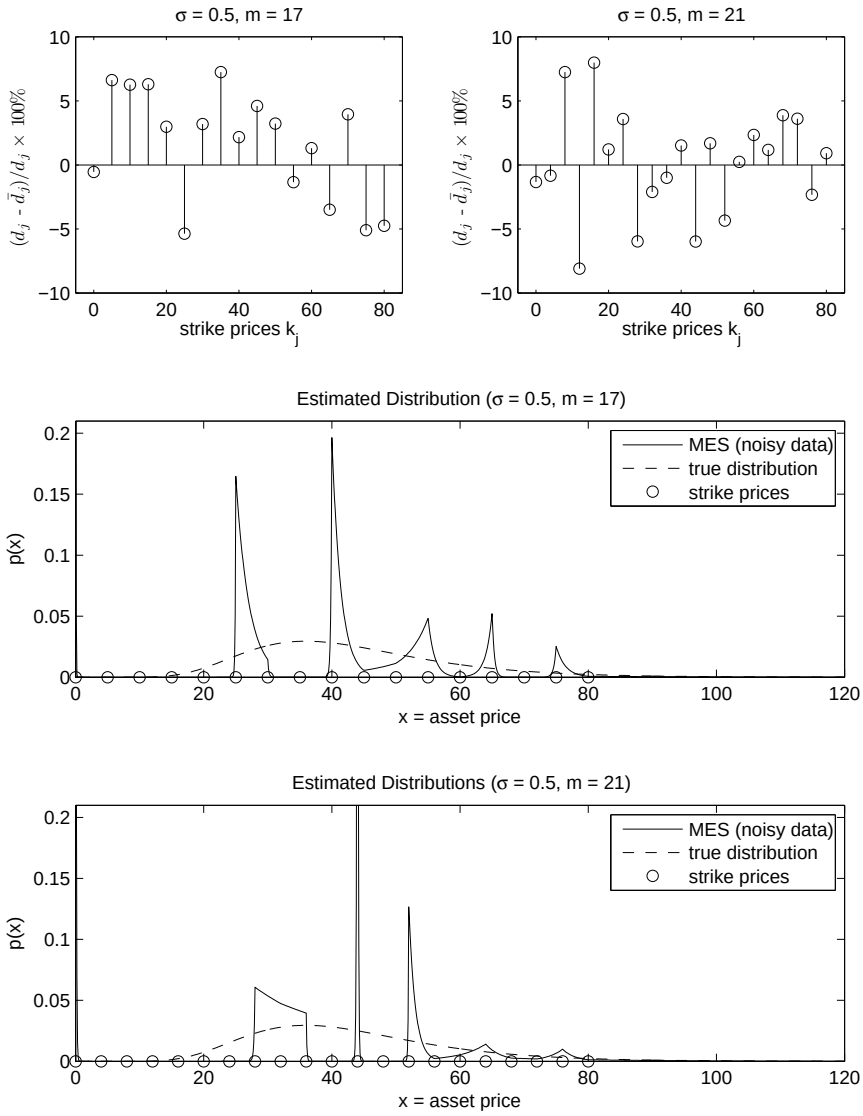

Figure 6.6: *Data with noise added and corresponding MESs. The top left (right) subplot shows* $(\bar{d}_j - d_j)/d_j$ *plotted against strike prices* k_j *of the example with 17 (21) data points. The middle (bottom) subplot shows the MES obtained for the example with 17 (21) data points. The MES is represented by the solid line. The true distribution is represented by the dashed line.*

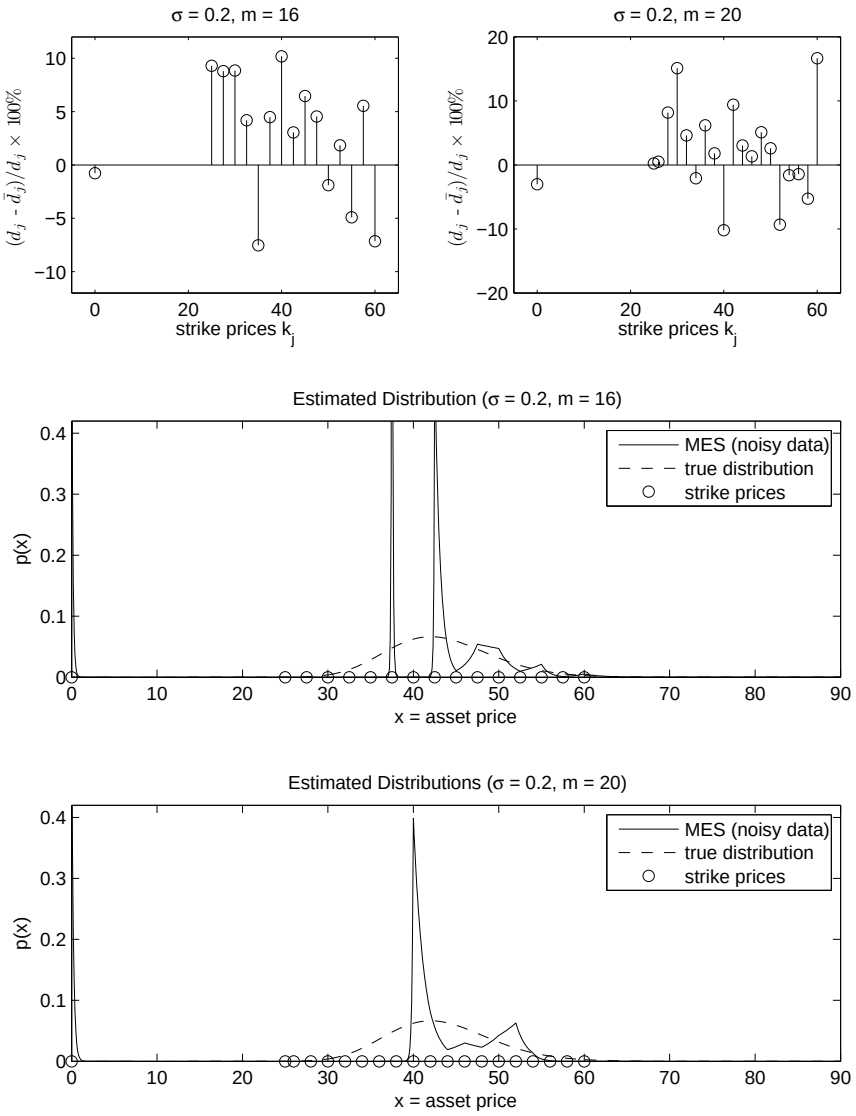

Figure 6.7: *Data with noise added and corresponding MESs. The top left (right) subplot shows* $(\bar{d}_j - d_j)/d_j$ *plotted against strike prices* k_j *of the example with 16 (20) data points. The middle (bottom) subplot shows the MES obtained for the example with 16 (20) data points. The MES is represented by the solid line. The true distribution is represented by the dashed line.*

Chapter 7

Regularization of the dual problem

7.1 Tikhonov regularization of the dual problem

The problem

$$\min_{\boldsymbol{\lambda}\in\mathbb{R}^m} D(\boldsymbol{\lambda}, \mathbf{d}), \tag{7.1}$$

where

$$D(\boldsymbol{\lambda}, \mathbf{d}) = \log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d},$$

is ill-posed. As we saw in the previous chapter, the norm of the inverse of the Hessian of the cost function D is an indicator of the instability of the minimizer $\boldsymbol{\lambda}^*$ of D. Table 6.1 shows that this norm is usually very large. As a consequence, even when the errors in the data are small, the errors in the estimated distributions can be very large. This is evident in figures 6.6 and 6.7.

Since problem (7.1) is ill-posed, it must be regularized. This can be done by adding the term $\alpha\|\boldsymbol{\lambda}\|_2^2$ with $\alpha > 0$ to $D(\boldsymbol{\lambda}, \mathbf{d})$. We know from previous chapters that this is Tikhonov regularization, with the difference that the penalty term is now added to a cost function and not to one side of an equation to be solved. Tikhonov himself used this procedure to regularize ill-posed minimization problems (see e.g. [12]). The regularized solution corresponds then to the solution of

$$\min_{\boldsymbol{\lambda}\in\mathbb{R}^m} D_\alpha(\boldsymbol{\lambda}, \mathbf{d}),$$

where

$$D_\alpha(\boldsymbol{\lambda}, \mathbf{d}) := \log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d} + \alpha \|\boldsymbol{\lambda}\|_2^2$$

and $\alpha > 0$.

In this way, even when the Hessian $\nabla_{\boldsymbol{\lambda}} D$ of D is singular, the Hessian $\nabla_{\boldsymbol{\lambda}} D_\alpha = \nabla_{\boldsymbol{\lambda}} D + 2\alpha I$ (where I denotes the identity matrix of appropriate size) is always nonsingular. Given the convexity of $D(\boldsymbol{\lambda}, \mathbf{d})$ with respect to $\boldsymbol{\lambda}$ and the convexity of $\|\boldsymbol{\lambda}\|_2^2$, it follows that $D_\alpha(\boldsymbol{\lambda}, \mathbf{d})$ is convex as well, so the new objective attains a minimum.

Typically, given a minimizer $\boldsymbol{\lambda}_\alpha$ of D_α and a minimizer $\boldsymbol{\lambda}^*$ of D, the value $D(\boldsymbol{\lambda}_\alpha, \mathbf{d})$ should not be much larger than $D(\boldsymbol{\lambda}^*, \mathbf{d})$, whereas $\|\boldsymbol{\lambda}_\alpha\|_2^2$ could be much smaller than $\|\boldsymbol{\lambda}^*\|_2^2$. It is then said that minimization of D_α penalizes minimizers of D whose Euclidean norm is large. How is the estimate $p_{\boldsymbol{\lambda}^*}$ affected by this penalization? In other words, what is the difference between the true MES $p_{\boldsymbol{\lambda}^*}$ and the *regularized MES* $p_{\boldsymbol{\lambda}_\alpha}$? We answer this question next.

Notice from figures 6.6 and 6.7 that the MESs present very high spikes at some strike prices k_j. At those, $\lim_{x \to k_j^-} p'_{\boldsymbol{\lambda}}(x) \gg \lim_{x \to k_j^+} p'_{\boldsymbol{\lambda}}(x)$. We show that this gap between the left and right derivatives at k_j is related to the size of $\boldsymbol{\lambda}$. Taking natural logarithms on both sides of (5.1), we have that

$$\log p_{\boldsymbol{\lambda}}(x) = -\log Z(\boldsymbol{\lambda}) + \sum_{i=1}^{m} \lambda_i c_i(x).$$

Differentiating with respect to x, this gives

$$\frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)} = \sum_{i=1}^{m} \lambda_i c_i'(x),$$

whose limits as x approaches k_j from left and right are, respectively,

$$\lim_{x \to k_j^-} \frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)} = \sum_{i=1}^{j-1} \lambda_i, \qquad \lim_{x \to k_j^+} \frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)} = \sum_{i=1}^{j} \lambda_i,$$

for $j = 2, \ldots, m$. For $j = 1$, only the right derivative is defined so, for convenience, we will define the left derivative at 0 as 0. Hence,

$$\left[\frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)}\right]_{k_j} := \left| \lim_{x \to k_j^+} \frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)} - \lim_{x \to k_j^-} \frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)} \right| = |\lambda_j|\,.$$

Therefore,

$$\|\boldsymbol{\lambda}\|_2^2 = \sum_{i=2}^{m} \left[\frac{p'_{\boldsymbol{\lambda}}(x)}{p_{\boldsymbol{\lambda}}(x)}\right]_{k_i}^2.$$

Thus, it can be said that Tikhonov regularization of the dual problem (7.1) penalizes MESs with large mean square jumps in their logarithmic derivatives at the strike prices.

7.2 Selection of the regularization parameter α

In order to select a regularization parameter α, we must first determine a range $[\alpha_{min}, \alpha_{max}]$ of possible values for such parameter. This is typically done simply by trial and error. For example, we can try $\alpha = 10^{-4}$, 10^{-3}, and so on until we find α large enough so that the Hessian of $D_\alpha(\boldsymbol{\lambda}, \mathbf{d})$ becomes sufficiently well-conditioned. This gives a way of determining α_{min}. On the other hand, if α is too large, the penalty functional $\|\boldsymbol{\lambda}\|_2^2$ will have too much weight, the components of the minimizer $\boldsymbol{\lambda}_\alpha$ will be too small, and the graph of $p_{\boldsymbol{\lambda}\alpha}$ will be too flat. By avoiding that situation we find α_{max}.

There are a variety of methods which attempt to select an appropriate value of α in the range $[\alpha_{min}, \alpha_{max}]$. Vogel [10] lists and briefly explains some of the most well-known methods, including the L-curve method, which is the one we implemented. Hansen and O'Leary [3] give a detailed explanation on how to implement the L-curve method when dealing with discrete ill-posed problems. In discrete ill-posed problems, Tikhonov regularization chooses the solution x_α that solves the minimization problem

$$\min_x \|Ax - b\|_2^2 + \alpha\|x\|_2^2,$$

where A is a rectangular matrix, b is a vector of known data, and x is the vector to be recovered. The parametric plot of $(\|Ax_\alpha - b\|_2^2, \|x_\alpha\|_2^2)$ has an L-shape. The L-curve method chooses the value of α for which $(\|Ax_\alpha - b\|_2^2, \|x_\alpha\|_2^2)$ is exactly on the corner of that L-curve.

In our case, the L-curve is the parametric plot of $(\log Z(\boldsymbol{\lambda}_\alpha) - \boldsymbol{\lambda}_\alpha^T \mathbf{d}, \|\boldsymbol{\lambda}_\alpha\|_2^2)$. Hansen and O'Leary [3] argue that, for discrete ill-posed problems, it is more appropriate to plot in a log-log scale. This is not the case for our problem, where often the log-log scale of the original plot is no longer an L-curve. Moreover, the x-coordinate of our L-curve is the original dual objective, whose value can be negative. This prevents the use of a log-log scale.

σ	m	$\frac{\|\mathbf{d}-\bar{\mathbf{d}}\|_2}{\|\mathbf{d}\|_2}$	$\frac{\|p_{\boldsymbol{\lambda}_\alpha}-p_{\tilde{\boldsymbol{\lambda}}_\alpha}\|_1}{\|p_{\boldsymbol{\lambda}_\alpha}\|_1}$	$\sigma(\boldsymbol{\lambda})$	$\sigma(\boldsymbol{\lambda}_\alpha)$	$\sigma(\tilde{\boldsymbol{\lambda}}_\alpha)$	$\frac{\|\sigma(\boldsymbol{\lambda}_\alpha)-\sigma(\tilde{\boldsymbol{\lambda}}_\alpha)\|}{\|\sigma(\boldsymbol{\lambda}_\alpha)\|}$
0.5	17	5.00%	9.60%	0.64	0.64	0.58	8.84%
0.5	21	5.00%	4.55%	0.64	0.62	0.62	0.64%
0.2	16	5.00%	13.18%	0.45	0.42	0.45	6.71%
0.2	20	5.00%	16.49%	0.45	0.42	0.49	17.22%

Table 7.1: *Volatility estimation. The first column indicates the assumed volatility σ. The second column shows the number m of data points. The third column shows the error level of the noisy data $\bar{\mathbf{d}}$ relative to the size of the noiseless data $\mathbf{d}$. The fourth column shows the error level of the regularized MES $p_{\tilde{\boldsymbol{\lambda}}_\alpha}$ with noisy data relative to the size of the regularized MES $p_{\boldsymbol{\lambda}_\alpha}$. The fifth column shows the volatility $\sigma(\boldsymbol{\lambda})$ recovered from the MES $p_{\boldsymbol{\lambda}}(x)$ with noiseless data $\mathbf{d}$. The sixth column shows the volatility $\sigma(\boldsymbol{\lambda}_\alpha)$ recovered from the regularized MES $p_{\boldsymbol{\lambda}_\alpha}(x)$ with noiseless data $\mathbf{d}$. The seventh column shows the volatility $\sigma(\tilde{\boldsymbol{\lambda}}_\alpha)$ recovered from the regularized MES $p_{\tilde{\boldsymbol{\lambda}}_\alpha}(x)$ with the projected noisy data $\tilde{\mathbf{d}}$. The eighth column shows the error level of $\sigma(\tilde{\boldsymbol{\lambda}}_\alpha)$ relative to the size of $\sigma(\boldsymbol{\lambda}_\alpha)$.*

By computing the curvature κ of the L-curve as a function of α and finding its maximum, it is possible to locate the corner. Given a parametric curve $(X(\alpha), Y(\alpha))$, its curvature function is

$$\kappa(\alpha) := \frac{X'(\alpha)Y''(\alpha) - X''(\alpha)Y'(\alpha)}{(X'(\alpha)^2 + Y'(\alpha)^2)^{3/2}}.$$

In our case, $X(\alpha) = \log Z(\boldsymbol{\lambda}_\alpha) - \boldsymbol{\lambda}_\alpha^T \mathbf{d}$ and $Y(\alpha) = \|\boldsymbol{\lambda}_\alpha\|_2^2$. There is no explicit representation for the first or second derivatives of X and Y with respect to α, and we only have the values of X and Y at a limited number of values of α. Furthermore, we wish not to compute too many points $(X(\alpha), Y(\alpha))$ on the L-curve, as the computational effort required for each point is significant. Fortunately, even though the interval $[\alpha_{min}, \alpha_{max}]$ might be large, in practice, the interval $[\log(\alpha_{min}), \log(\alpha_{max})]$ is not that large. Hence, enough points in that range can be computed so that $\kappa(t)$ with $t = \log(\alpha)$ can be evaluated approximately at the available points by numerical differentiation. For each of our examples, we used 20 values of α in $[\alpha_{max}, \alpha_{min}]$ with $\log(\alpha_i)$ equally spaced, and selected the value of α at which κ was the largest. By letting $i^* := \arg\max\{\kappa(\log(\alpha_i)), i = 1, 2, \ldots, 20\}$, it was seen that the graphs of the MESs using α_{i^*}, or α_{i^*-1}, or α_{i^*+1} were all very similar, so a more accurate location of the corner of the L-curve seemed unnecessary.

With α selected, we found the minimizers $\boldsymbol{\lambda}_\alpha$ of the regularized dual objective $D_\alpha(\boldsymbol{\lambda}, \mathbf{d})$ and the corresponding regularized MESs $p_{\boldsymbol{\lambda}_\alpha}$ for the cases with 16 or more data. We did this for both the noiseless data $\mathbf{d}$ and the projected noisy data $\tilde{\mathbf{d}}$ that were introduced in the previous section. The resulting MESs are shown in figures 7.1 — 7.8. Note that the spike near $x = 0$ has disappeared under proper regularization. Each MES is accompanied with its respective L-curve, curvature of the L-curve, and MESs obtained with other choices of the parameter α. Table 7.1 shows the error level $\|p_{\boldsymbol{\lambda}_\alpha} - p_{\tilde{\boldsymbol{\lambda}}_\alpha}\|_1$ of the regularized MES $p_{\tilde{\boldsymbol{\lambda}}_\alpha}$ with the projected noisy data $\bar{\mathbf{d}}$ relative to the size of the regularized MES $p_{\boldsymbol{\lambda}_\alpha}$ with noiseless data $\mathbf{d}$. A comparison of tables 6.2 with 7.1 shows that the error level of the regularized MESs with noisy data is much smaller than that of the original MESs with noisy data.

We are comparing the estimated distributions by means of the L^1-norm. Another quantity which we can also compare, and is always of interest to investment professionals, is volatility. Our probability distributions were simulated under the assumption of geometric Brownian motion for the stock price behavior. A current stock price d_1, an expiration date T, an expected return r, and a volatility σ completely determine the probability distribution $p(x)$. Inversely, if we know $p(x)$ and the parameters other than σ, then we can recover

$$\sigma = \left(\log \left(1 + \frac{\mathrm{Var}(x)}{\mathrm{E}(x)^2} \right) T^{-1} \right)^{\frac{1}{2}},$$

where

$$\mathrm{E}(x) := \int_0^\infty x p(x) dx, \qquad \mathrm{Var}(x) := \int_0^\infty (x - \mathrm{E}(x))^2 \, p(x) dx.$$

Using the original distributions from figure 6.1 we can verify that $\sigma = 0.5$ and $\sigma = 0.2$ respectively. Using the formula above, let us denote by $\sigma(\boldsymbol{\lambda})$, $\sigma(\boldsymbol{\lambda}_\alpha)$ and $\sigma(\tilde{\boldsymbol{\lambda}}_\alpha)$ the recovered volatilities when using the MES $p_{\boldsymbol{\lambda}}(x)$ with noiseless data $\mathbf{d}$, the regularized MES $p_{\boldsymbol{\lambda}_\alpha}(x)$ with noiseless data $\mathbf{d}$, and the regularized MES $p_{\tilde{\boldsymbol{\lambda}}_\alpha}(x)$ with the projected noisy data $\tilde{\mathbf{d}}$, respectively. We report these values in table 7.1 as well. The last column shows that the error level of the volatility from the regularized MESs is comparable with the error level of the regularized MESs theirselves. On the other hand, there is a large discrepancy between the volatility recovered from the MESs and the actual volatility, which is more than doubled when $\sigma = 0.2$. However, this discrepancy is caused by the nature of the original MES and not by the regularization procedure,

as it can be seen by comparing the first and fifth columns of table 7.1. Specifically, the high value of $p_{\boldsymbol{\lambda}}(x)$ at $x = 0$ causes the integrand $(x - \mathrm{E}(x))^2 p_{\boldsymbol{\lambda}}(x)$ to be significant for small values of x so that $Var(x)$ is much larger for the MES than what it is for the lognormal distribution (where $p(0) = 0$).

We can use regularized MESs to estimate stock price distributions with real data. In figures 7.9 — 7.11 we show the estimated stock price distributions of a `Microsoft` stock whose price on March 11, 2010, was 29.18 dollars. The expiration dates are, respectively, April 16, 2010, January 20, 2012. The strike prices, ordered by increasing chronological order of expiration, were the following:

$$\mathbf{k} = [0, 12.5, 15, 17.5, 19, 20, 21, \ldots, 34, 35, 36, 40],$$
$$\mathbf{k} = [0, 2.5, 5, \ldots, 30, 32.5, 35, 40, 45],$$
$$\mathbf{k} = [0, 12.5, 15, 17.5, \ldots, 35, 37.5, 40, 45].$$

The number m of data points used were 23, 18, and 14, respectively. The option prices were taken to be the average of the bid and the ask. The riskless discount factor $DC(T)$ (see equation (4.1)) was set equal to 1, which is good enough for the purposes of the thesis. It is possible to obtain more accurate estimates of $DC(T)$ which will be just slightly less than 1.

7.3 Convergence analysis of the regularized minimizer

When a regularization procedure is implemented, the regularized solution must have the property that, as the error level in the data approaches 0, the errors in the regularized solution converge to zero as well. The analysis of this property is known as convergence analysis. We show in this section that a minimizer $\boldsymbol{\lambda}_\alpha$ of $D_\alpha(\boldsymbol{\lambda}, \mathbf{d})$ has this property.

The cost function D depends nonlinearly on $\boldsymbol{\lambda}$. Engl et al. [13] have considered the problem of solving for $\mathbf{x}$ in $F(\mathbf{x}) = \mathbf{y}$, where F is a nonlinear function. When the problem is ill-posed, they analyze the convergence of the regularized solution. We deal, however, with a minimization problem, so the approach is slightly different. To show convergence of the regularized minimizer, we start by following the approach by Engl et al., and then we benefit from the special form of the cost function which depends linearly

on $\mathbf{d}$. We should point out that some work on the regularization of nonlinear problems has been done since the book by Engl et al. [13] was published in 1996, including work done in Linz using Bregman divergence (see e.g. [48]). However, such work deals with operators $F : X \to Y$ where X and Y are infinite dimensional Banach spaces. In our case we just have $F : \mathbb{R}^m \to \mathbb{R}$.

Theorem 7.3.1. *Let $F : \mathbb{R}^m \to \mathbb{R}$ be a continuous function. Let $\mathbf{d}$, $\mathbf{d}^\delta \in \mathbb{R}$ with $\|\mathbf{d} - \mathbf{d}^\delta\| \leq \delta$ and let $\alpha(\delta)$ be such that $\delta < \alpha(\delta) \to 0$ as $\delta \to 0$.*[1] *Let $\left\{\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\}$ be any sequence such that $\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}$ is a minimizer of*

$$F(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}^{\delta_k} + \alpha_k \|\boldsymbol{\lambda}\|^2, \tag{7.2}$$

where $\delta_k \to 0$ as $k \to \infty$ and $\alpha_k := \alpha(\delta_k)$. Then every such sequence $\left\{\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\}$ has a convergent subsequence. The limit of every convergent subsequence is a minimizer of

$$F(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}. \tag{7.3}$$

If in addition, the minimizer $\boldsymbol{\lambda}^*$ of (7.3) is unique, then

$$\lim_{\delta \to 0} \boldsymbol{\lambda}_{\alpha(\delta)}^{\delta} = \boldsymbol{\lambda}^*.$$

Proof. Let α_k and $\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}$ be as above and let $\boldsymbol{\lambda}^*$ be a minimizer of (7.3). Then, by definition of $\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}$,

$$\begin{aligned} F\left(\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right) - {\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}}^T \mathbf{d}^{\delta_k} + \alpha_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|^2 &\leq F\left(\boldsymbol{\lambda}^*\right) - {\boldsymbol{\lambda}^*}^T \mathbf{d}^{\delta_k} + \alpha_k \left\|\boldsymbol{\lambda}^*\right\|^2 \\ &= F\left(\boldsymbol{\lambda}^*\right) - {\boldsymbol{\lambda}^*}^T \mathbf{d} + {\boldsymbol{\lambda}^*}^T \left(\mathbf{d} - \mathbf{d}^{\delta_k}\right) + \alpha_k \left\|\boldsymbol{\lambda}^*\right\|^2. \end{aligned} \tag{7.4}$$

By using the definition of $\boldsymbol{\lambda}^*$ on the right hand side of the above equation,

$$\begin{aligned} &F\left(\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right) - {\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}}^T \mathbf{d}^{\delta_k} + \alpha_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|^2 \\ &\quad \leq F\left(\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right) - {\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}}^T \mathbf{d} + {\boldsymbol{\lambda}^*}^T \left(\mathbf{d} - \mathbf{d}^{\delta_k}\right) + \alpha_k \left\|\boldsymbol{\lambda}^*\right\|^2, \end{aligned}$$

or

$$\alpha_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|^2 \leq {\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}}^T \left(\mathbf{d}^{\delta_k} - \mathbf{d}\right) + {\boldsymbol{\lambda}^*}^T \left(\mathbf{d} - \mathbf{d}^{\delta_k}\right) + \alpha_k \left\|\boldsymbol{\lambda}^*\right\|^2.$$

[1] *For example, $\alpha(\delta) := \sqrt{\delta}$ has this property.*

By Cauchy-Schwarz inequality and the definition of $\mathbf{d}^{\delta}$,

$$\begin{aligned}\alpha_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|^2 &\leq \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| \left\|\mathbf{d}^{\delta_k} - \mathbf{d}\right\| + \|\boldsymbol{\lambda}^*\| \left\|\mathbf{d} - \mathbf{d}^{\delta_k}\right\| + \alpha_k \|\boldsymbol{\lambda}^*\|^2 \\ &\leq \delta_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| + \delta_k \|\boldsymbol{\lambda}^*\| + \alpha_k \|\boldsymbol{\lambda}^*\|^2,\end{aligned}$$

or

$$\alpha_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|^2 - \delta_k \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| \leq \delta_k \|\boldsymbol{\lambda}^*\| + \alpha_k \|\boldsymbol{\lambda}^*\|^2.$$

Let $x := \|\boldsymbol{\lambda}^*\|$, $y_k := \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\|$, so that the last inequality becomes

$$\alpha_k y_k^2 - \delta_k y_k \leq \delta_k x + \alpha_k x^2.$$

If k is such that $y_k^2 < y_k$. Then $y_k = \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| < 1$. Otherwise, k is such that $y_k^2 \geq y_k$, in which case

$$\begin{aligned}\alpha_k y_k - \delta_k y_k &\leq \alpha_k y_k^2 - \delta_k y_k \\ &\leq \delta_k x + \alpha_k x^2 \\ &\leq (\delta_k + \alpha_k) \max\left\{x, x^2\right\},\end{aligned}$$

or

$$\left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| \leq \frac{\alpha_k + \delta_k}{\alpha_k - \delta_k} \max\left\{\|\boldsymbol{\lambda}^*\|, \|\boldsymbol{\lambda}^*\|^2\right\}.$$

Since $\alpha(\delta) > \delta$ as $\delta \to 0$, $\frac{\alpha(\delta)+\delta}{\alpha(\delta)-\delta} = 1 + \frac{2\delta}{\alpha(\delta)-\delta} \to 1$ as $\delta \to 0$. Hence,

$$\lim_{k\to\infty} \left\|\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\| \leq \max\left\{1, \|\boldsymbol{\lambda}^*\|, \|\boldsymbol{\lambda}^*\|^2\right\}.$$

Thus, $\left\{\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\}$ is bounded. Therefore, there exist an element $\boldsymbol{\lambda}^\dagger \in \mathbb{R}^m$ and a subsequence again denoted by $\left\{\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}\right\}$ such that

$$\lim_{k\to\infty} \boldsymbol{\lambda}_{\alpha_k}^{\delta_k} = \boldsymbol{\lambda}^\dagger.$$

Hence, taking the limit in inequality (7.4) as $k \to \infty$, we have, by continuity of F, that

$$F\left(\boldsymbol{\lambda}^\dagger\right) - {\boldsymbol{\lambda}^\dagger}^T \mathbf{d} \leq F\left(\boldsymbol{\lambda}^*\right) - {\boldsymbol{\lambda}^*}^T \mathbf{d} \leq F\left(\boldsymbol{\lambda}^\dagger\right) - {\boldsymbol{\lambda}^\dagger}^T \mathbf{d}.$$

Thus,

$$F\left(\boldsymbol{\lambda}^\dagger\right) - {\boldsymbol{\lambda}^\dagger}^T \mathbf{d} = F\left(\boldsymbol{\lambda}^*\right) - {\boldsymbol{\lambda}^*}^T \mathbf{d}.$$

Therefore, $\boldsymbol{\lambda}^\dagger$ also minimizes (7.3). If $\boldsymbol{\lambda}^*$ is unique, convergence of $\boldsymbol{\lambda}_{\alpha(\delta)}^{\delta}$ follows since every sequence has a subsequence converging toward $\boldsymbol{\lambda}^*$. □

Corollary 7.3.1. *Let* $\mathbf{d}_\delta$, $\alpha(\delta)$ *and* α_k *be as in the theorem. Let* $\boldsymbol{\lambda}_{\alpha_k}^{\delta_k}$ *be a minimizer of*

$$\log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}^{\delta_k} + \alpha_k \|\boldsymbol{\lambda}\|^2,$$

where $Z(\boldsymbol{\lambda})$ *is given by* (4.8)*. Then*

$$\lim_{\delta \to 0} \boldsymbol{\lambda}_{\alpha(\delta)}^{\delta} = \boldsymbol{\lambda}^*,$$

where $\boldsymbol{\lambda}^*$ *is the minimizer of*

$$\log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}.$$

Proof. Recall that the expression $\mu_j^{-1}(\exp(\mu_j k_{j+1}) - \exp(\mu_j k_j))$ is understood to be $k_{j+1} - k_j$ when $\mu_j = 0$. It can then be easily verified that the function $Z(\boldsymbol{\lambda})$ is continuous. Hence, so is $\log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}$. Borwein et al. [43] show that $\log Z(\boldsymbol{\lambda}) - \boldsymbol{\lambda}^T \mathbf{d}$ has a unique minimizer. Therefore, the result follows from the theorem. □

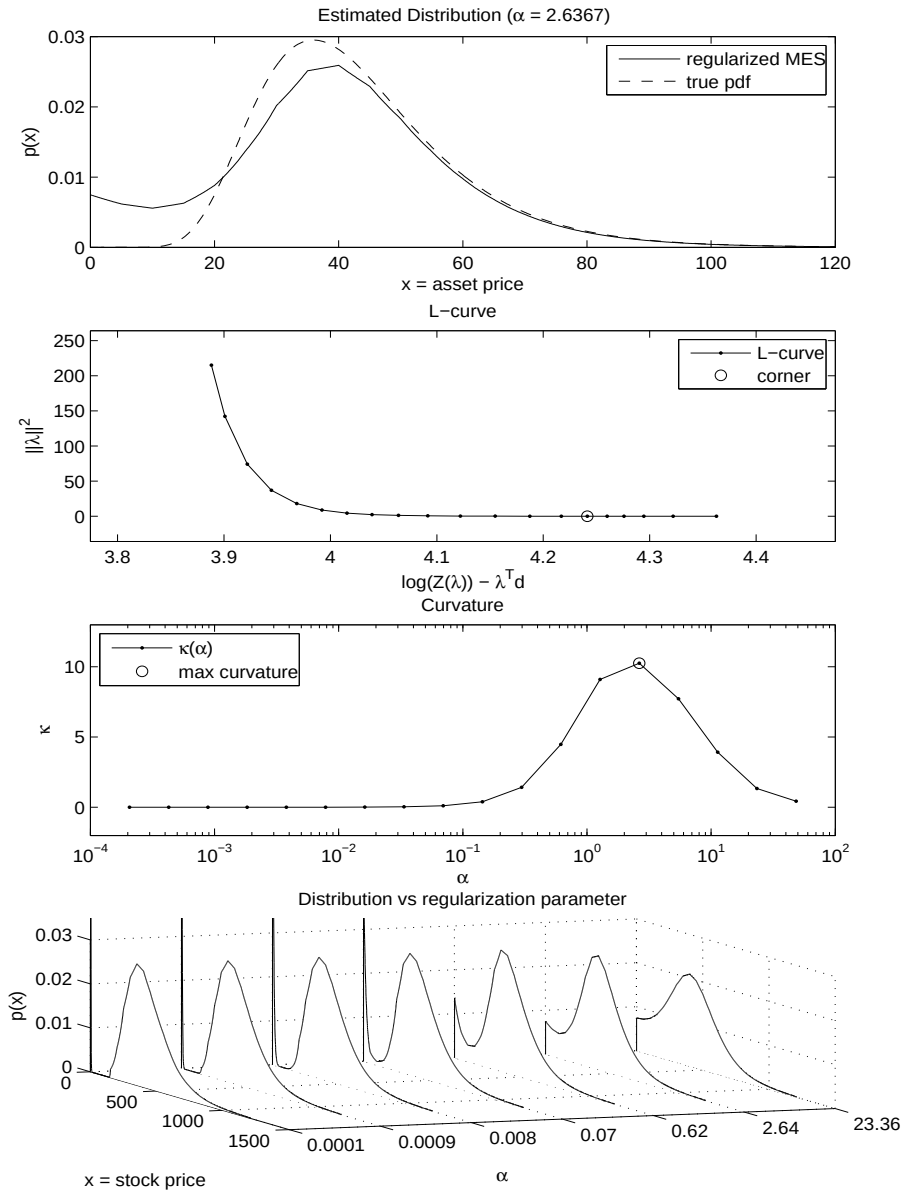

Figure 7.1: *Regularized MESs with noiseless data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 17 noiseless data points* $\mathbf{d}$ *when* $\sigma = 0.5$. *The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

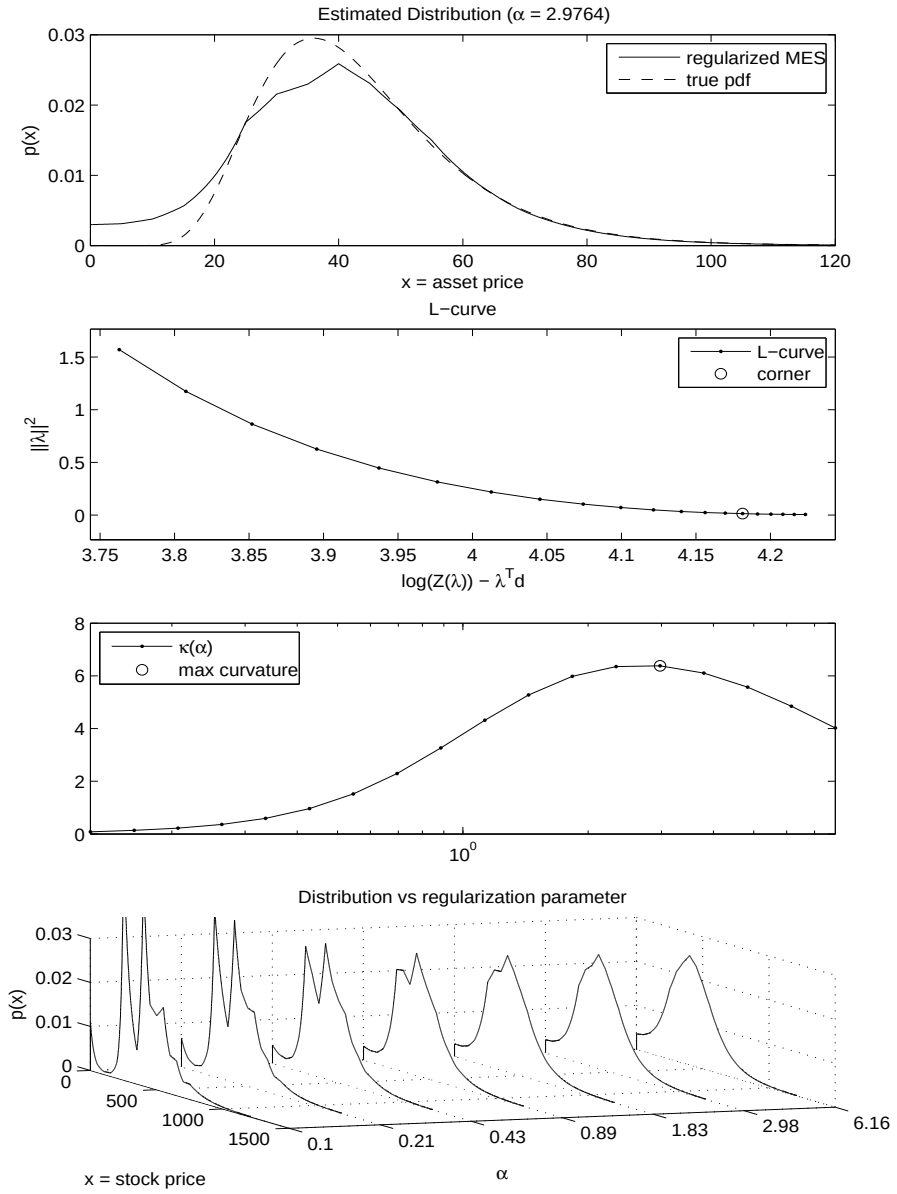

Figure 7.2: *Regularized MESs with noiseless data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 17 noisy data points* $\mathbf{d}$ *when* $\sigma = 0.5$ *(see the first subplot of figure 6.6). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α*.*

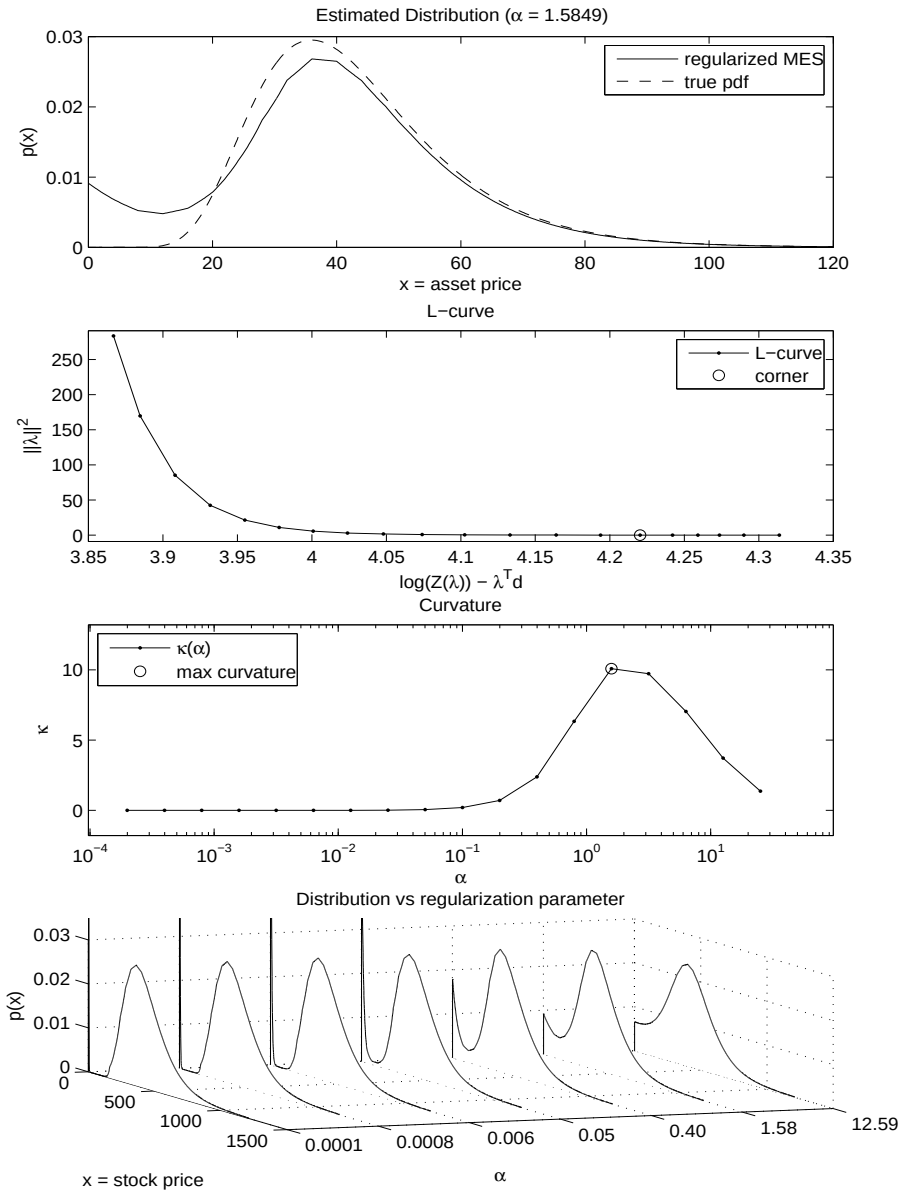

Figure 7.3: *Regularized MESs with noiseless data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 21 noiseless data points* $\mathbf{d}$ *when* $\sigma = 0.5$ *(see the fourth subplot of figure 6.2). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

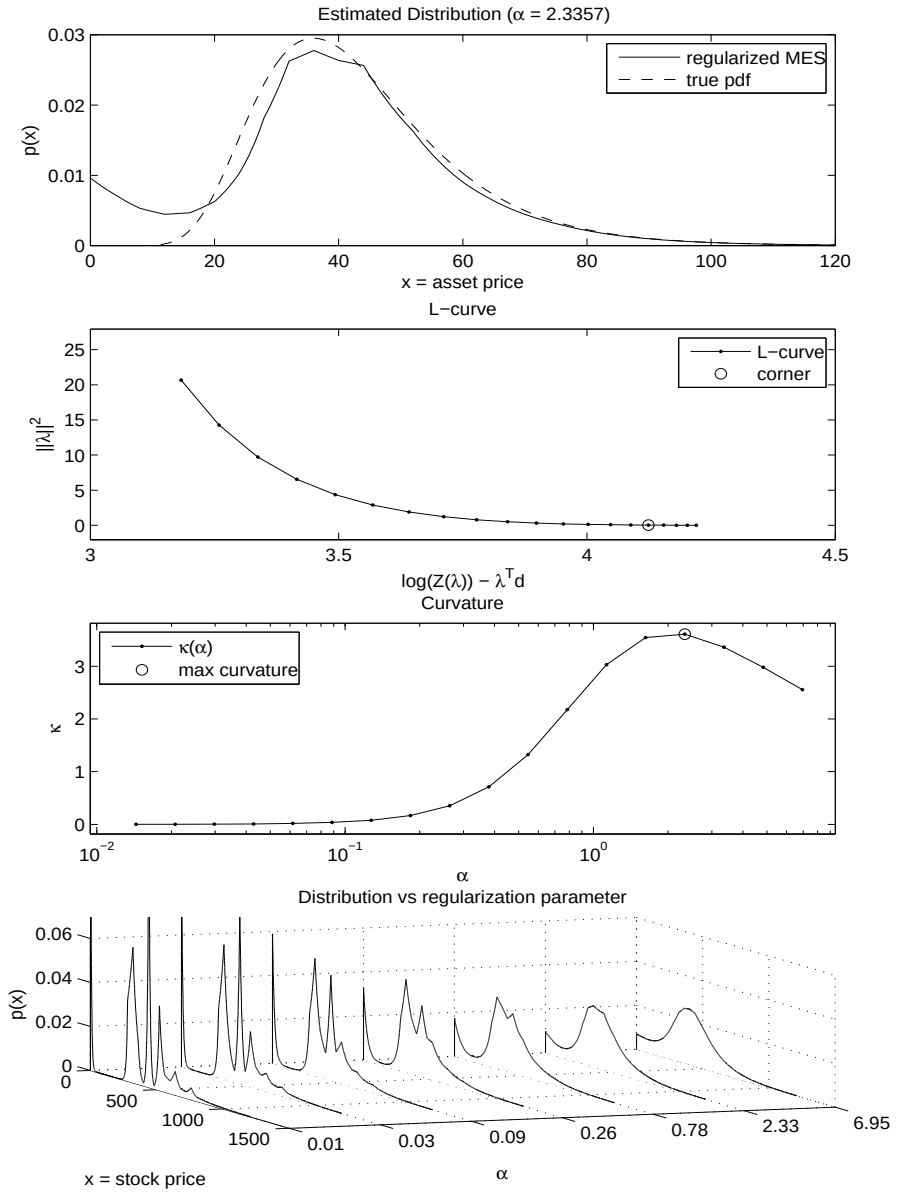

Figure 7.4: *Regularized MESs with noisy data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 21 noisy data points* $\mathbf{d}$ *when* $\sigma = 0.5$ *(see the second subplot of figure 6.6). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

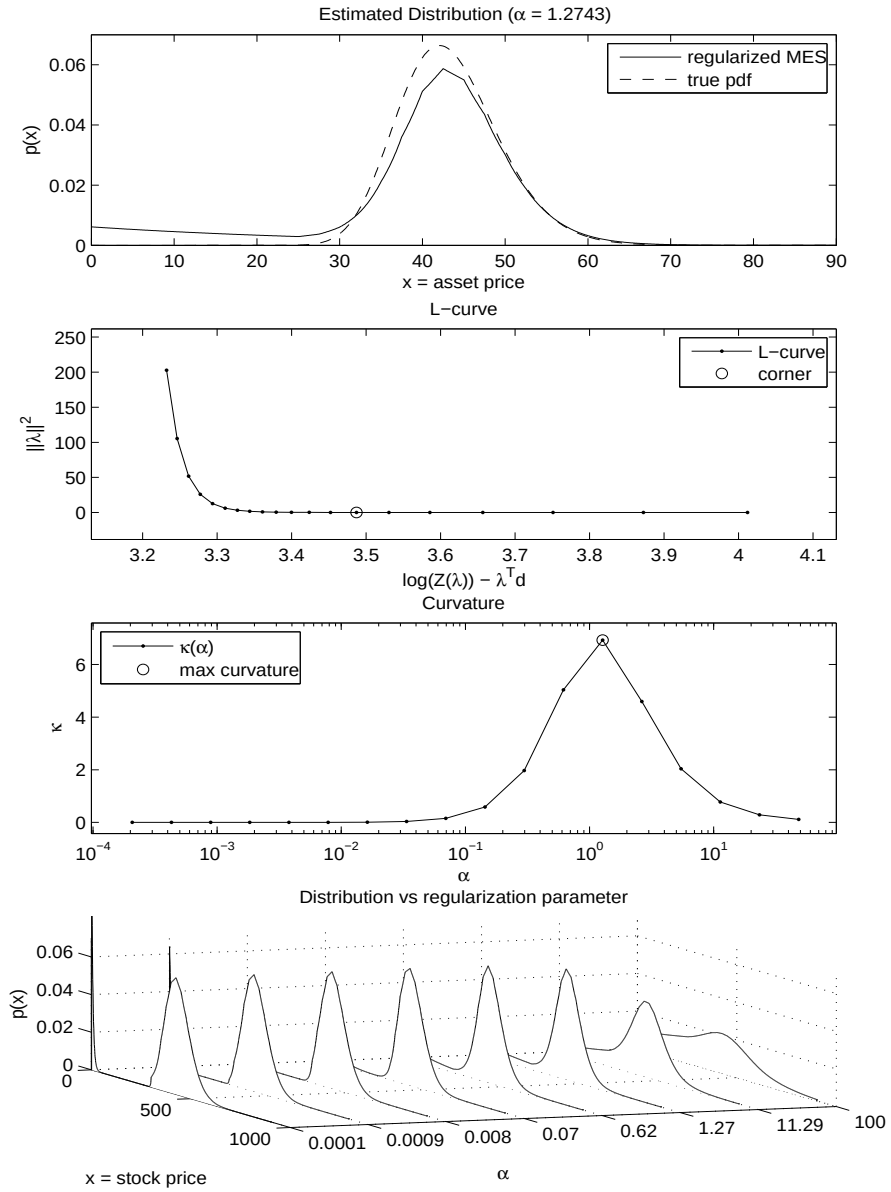

Figure 7.5: *Regularized MESs with noiseless data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 16 noiseless data points* $\mathbf{d}$ *when* $\sigma = 0.2$ *(see the third subplot of figure 6.3). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

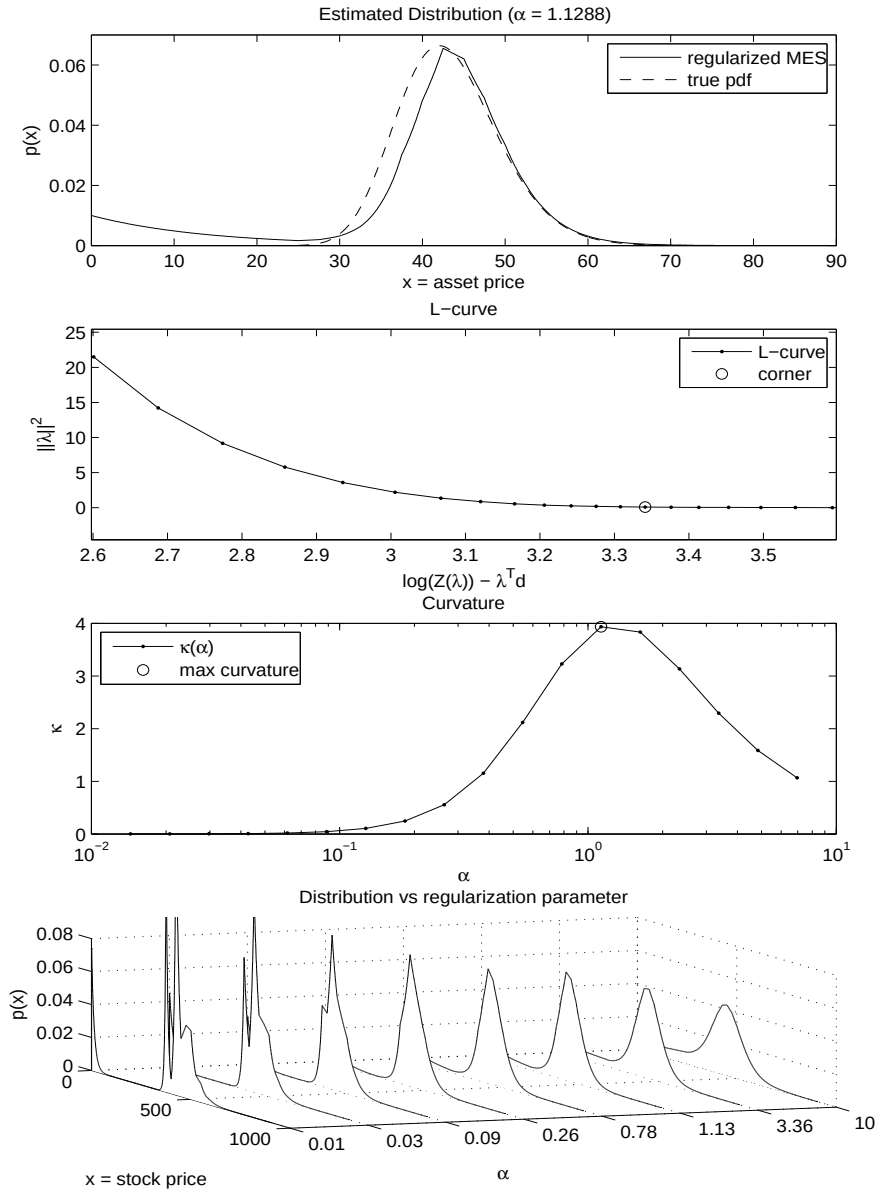

Figure 7.6: *Regularized MESs with noisy data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 16 noisy data points* $\mathbf{d}$ *when* $\sigma = 0.2$ *(see the first subplot of figure 6.7). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

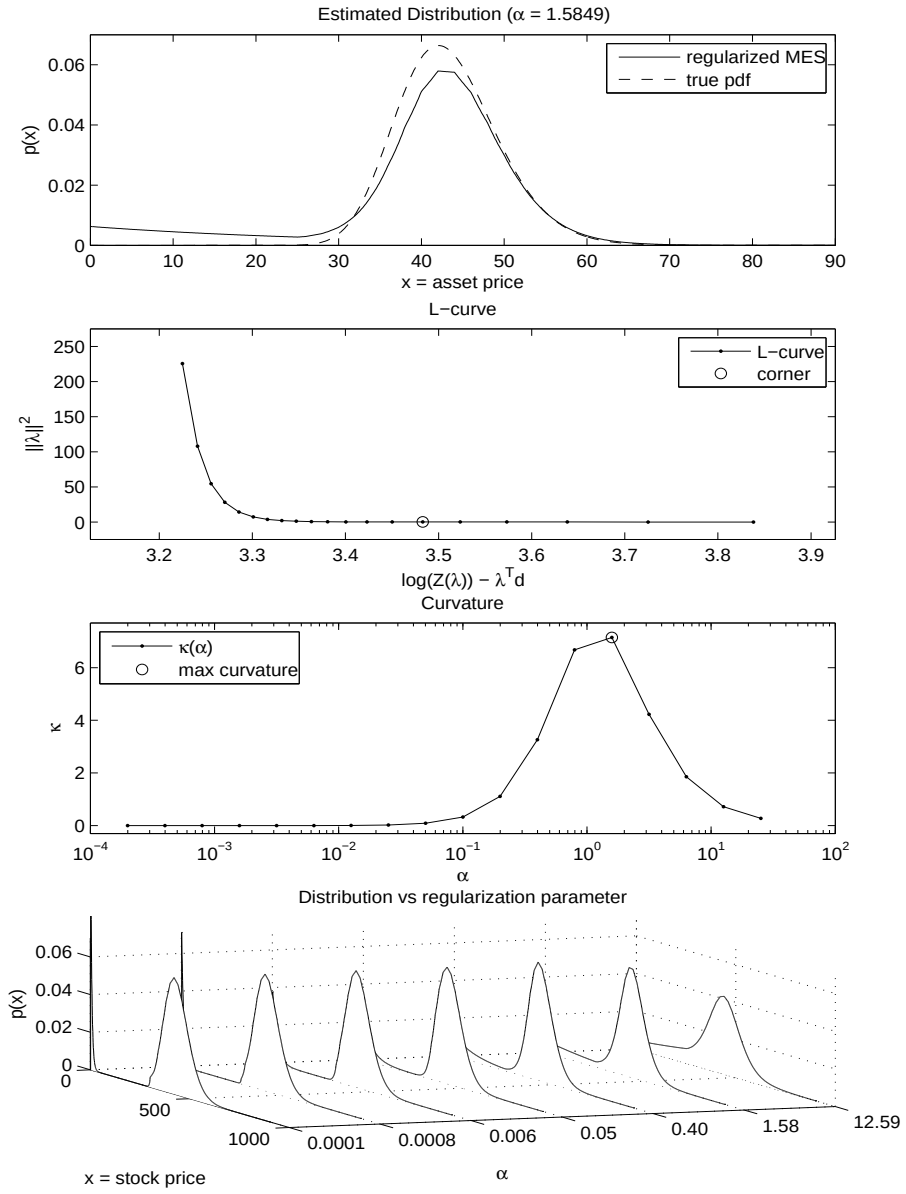

Figure 7.7: *Regularized MESs with noiseless data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 20 noiseless data points* $\mathbf{d}$ *when* $\sigma = 0.2$ *(see the fourth subplot of figure 6.3). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α*.*

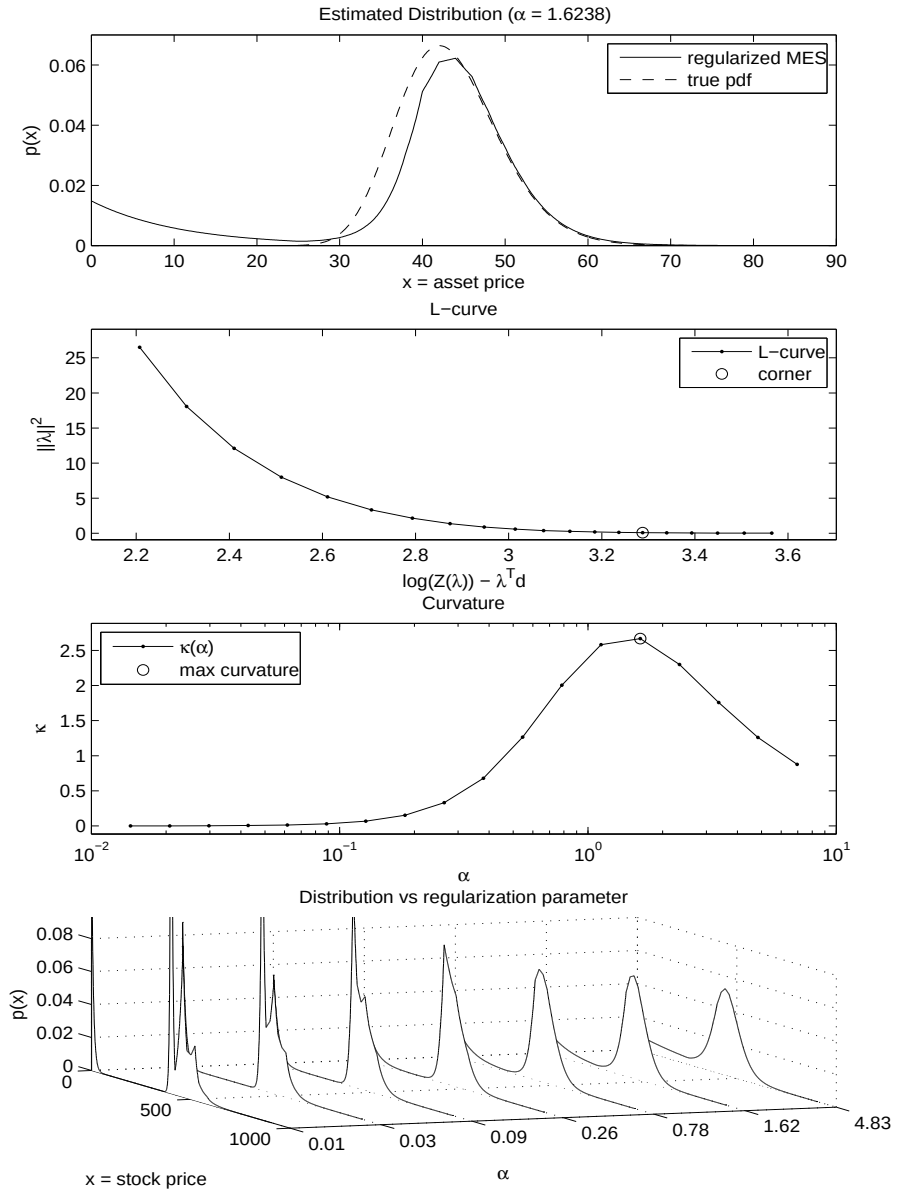

Figure 7.8: *Regularized MESs with noisy data. The solid line of the first subplot corresponds to the graph of the estimated stock price distribution using the 20 noisy data points* $\mathbf{d}$ *when* $\sigma = 0.2$ *(see the second subplot of figure 6.7). The dashed line corresponds to the graph of the true distribution. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α*.*

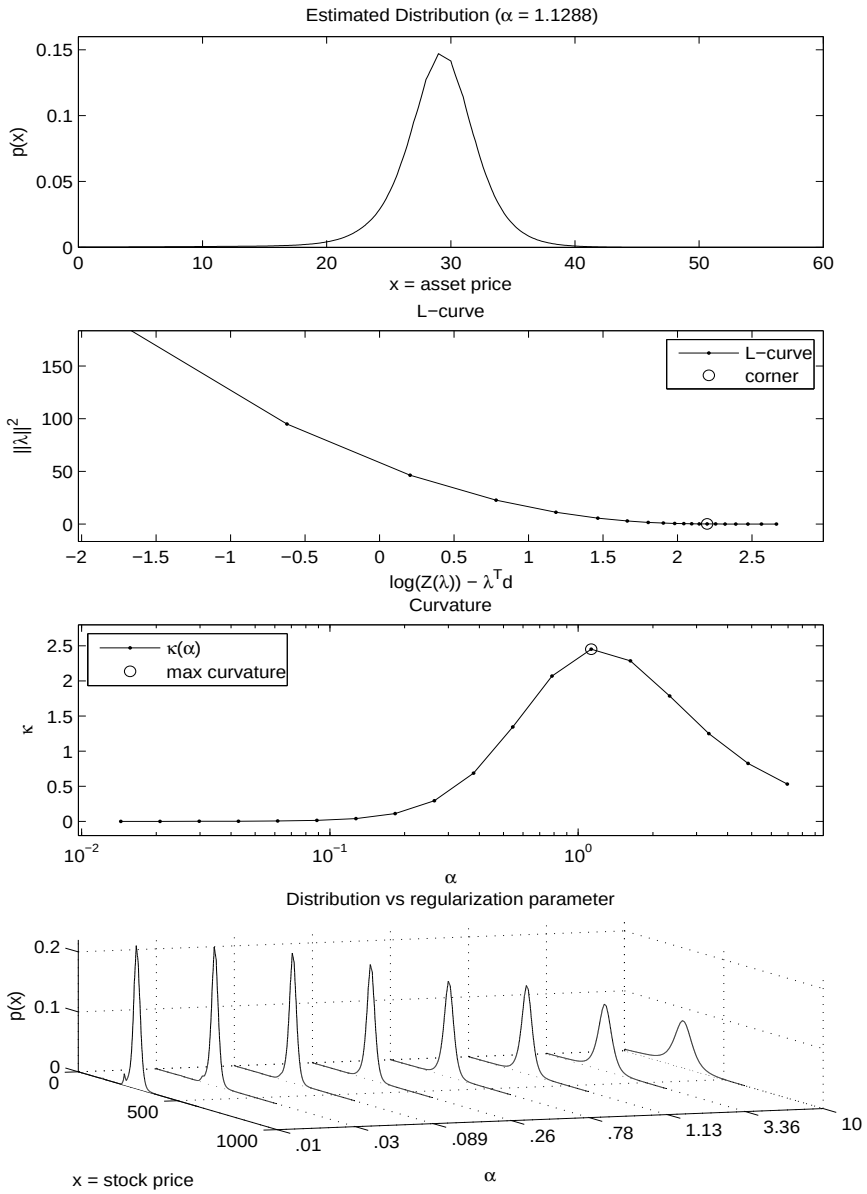

Figure 7.9: *Regularized MESs from* `Microsoft` *call option prices. The first subplot shows the graph of the estimated stock price distribution using the* `Microsoft` *call option prices from March 11, 2010, with expiration on April 16, 2010. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

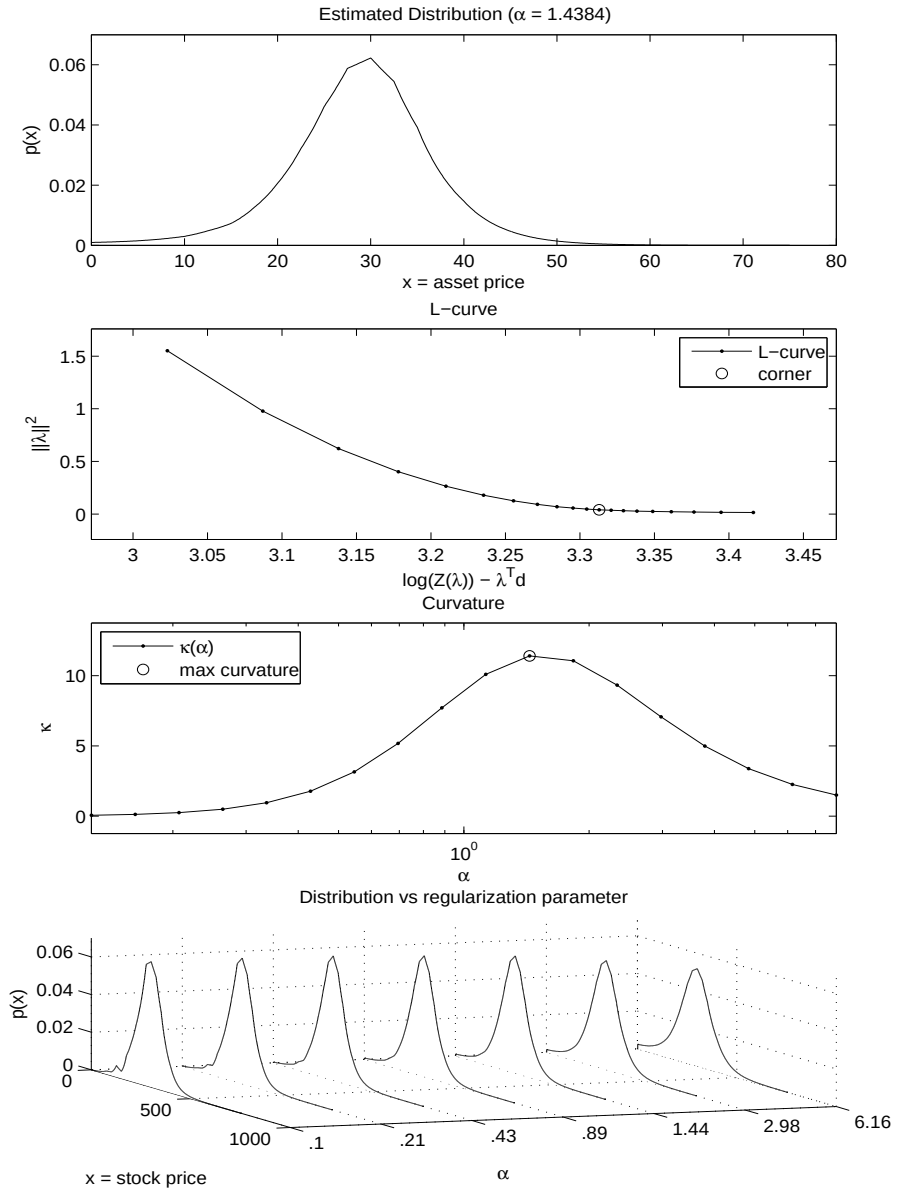

Figure 7.10: *Regularized MESs from* `Microsoft` *call option prices. The first subplot shows the graph of the estimated stock price distribution using the* `Microsoft` *call option prices from March 11, 2010, with expiration on January 21, 2011. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

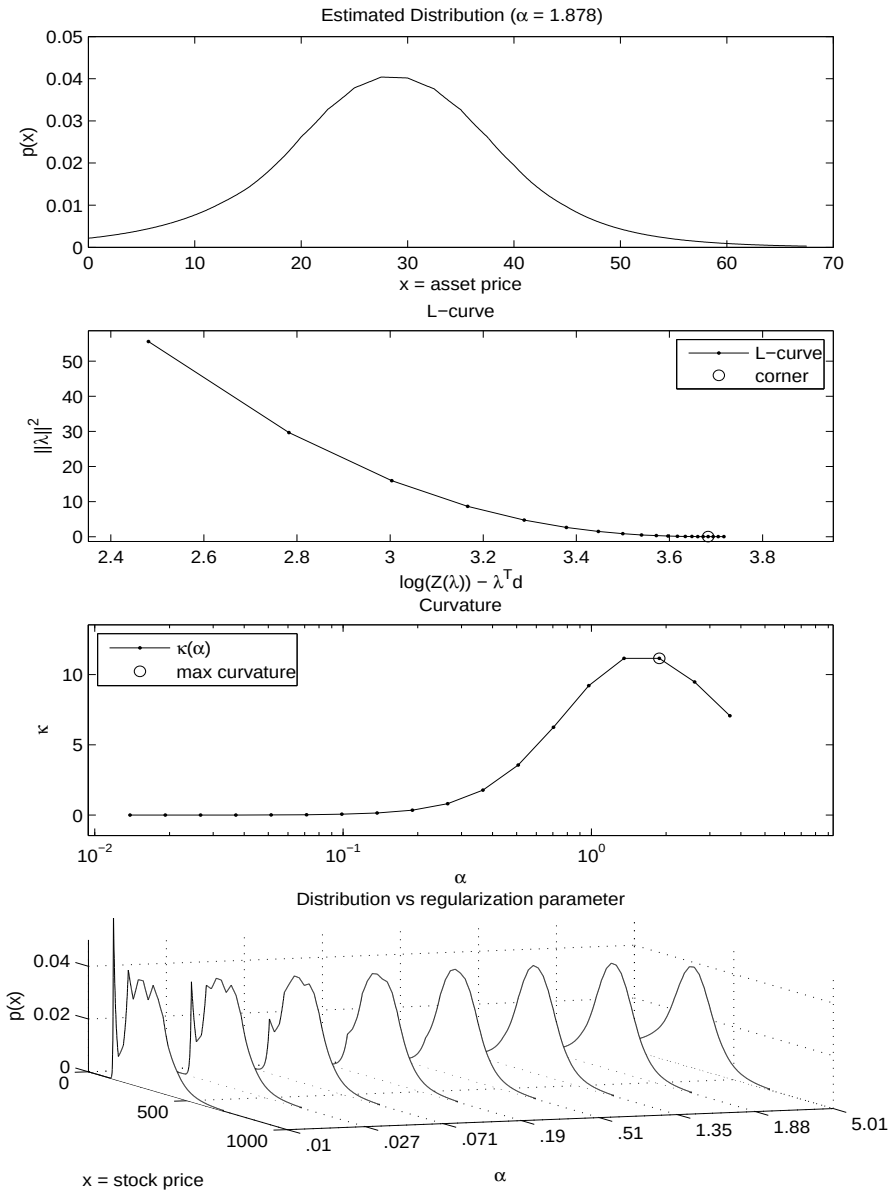

Figure 7.11: *Regularized MESs from* `Microsoft` *call option prices. The first subplot shows the graph of the estimated stock price distribution using the* `Microsoft` *call option prices from March 11, 2010, with expiration on January 20, 2012. The second subplot shows the corresponding L-curve. The third subplot shows the corresponding curvature. The fourth subplot shows different stock price distributions given some values of* α.

Chapter 8

Discussion

The main purpose of the first part of this thesis was to analyze a multiplicative regularization strategy. The strategy analyzed has the property that no regularization parameter is introduced. Regularization strategies where a parameter is introduced are routinely used in the applications, so we compared the multiplicative regularization strategy with one of the well known regularization methods that make use of a parameter, namely Tikhonov regularization. We reviewed Tikhonov regularization in the context of the image deblurring problem, as well as a well known method for the selection of a Tikhonov parameter, namely the L-curve criterion.

While reviewing the L-curve criterion, it was shown that the L curve and the efficient frontier are the same curve, even though they depend on parameters that appear in different regularization formulations. After introducing a multiplicative regularization strategy, it was shown that there exists a Tikhonov parameter for which the Tikhonov regularized solution is the same as the multiplicative regularized solution. This means that, when a multiplicative regularized solution $\mathbf{f}_*$ exists, the point $(\log \|K\mathbf{f}_* - \mathbf{d}\|^2, \log \|B\mathbf{f}_*\|^2)$ does lie on the L-curve or efficient frontier.

We showed that the distance between the regularized solution with Tikhonov parameter chosen by the L-curve criterion and the multiplicative regularized solution is proportional to the distance between their corresponding Tikhonov parameters. Hundreds of simulations of the deblurring problem with noise levels between 1 and 20% show that the corresponding Tikhonov parameters are close to each other, although the parameter corresponding to the multiplicative regularized solution tends to be slightly

larger. This implies that the solutions given by each method are very close to each other, but the solution given by the multiplicative strategy tends to be slightly smoother.

When the noise level is larger than 20%, we saw in the simulations that the solution given by multiplicative strategy was the zero vector. An analysis of the multiplicative cost function shows that the multiplicative cost function may not have nonzero local minimizers, which are regularized solutions given by the strategy. We devised a method that confirms the nonexistence of nonzero local minimizers when the noise level is large.

The multiplicative cost function is not convex everywhere, but it has to be convex at a neighborhood of the local minimizer if it exists. To determine this convexity, it is necessary to check the positive definiteness of the Hessian of the cost function. We exploited the special structure of the Hessian to give a sufficient condition much simpler than directly calculating its eigenvalues. However, since the condition is only sufficient, the Hessian may be positive definite even if the condition does not hold.

We conclude from the analysis in the first part that, unlike Tikhonov regularization aided with the L-curve criterion, the applicability of the multiplicative strategy presented here depends on the size of the noise level. Nonetheless, in most applications one deals with data where the noise level is small, so a multiplicative regularization is possible. In those situations, we know that the regularized solution obtained will be similar to the solution obtained via Tikhonov regularization with the L-curve criterion. The computational cost required by the multiplicative strategy is necessarily lower as no parameter selection is needed.

In the second part of the thesis, we studied the problem of estimating the probability density function of an asset from its European call option prices. Our approach was to investigate the dual formulation of Borwein et al. We found that while the dual problem has a unique solution, and that we can establish that the probability density function is stable under parameter perturbation, it is not stable to perturbation in the data.

This fact is demonstrated in numerical simulations. In the case of noiseless data, figure 6.4 shows that, as the number of data used is increased, the maximum entropy solution (MES) closely resembles the true distribution. Indeed, in the last subplot of figure 6.4, where 21 data points were used, the MES and the true distribution are almost identical, except for the "spike" near $x = 0$. On the other hand, we had to stop the Newton iterations after only a few steps because of the inherent instability, reflected by

the ill-conditioning of the Hessian as we demonstrate in table 6.1.

When noise is added to the data, the MES is clearly unstable. Several issues arise. Data with small noise may not satisfy the constraint qualification (4.4)–(4.5). In these cases, only a limited number of the available data do satisfy the constraint qualification. We are then forced to modify the original noisy data if we want to apply the method with all of the available data. We accomplish this by projecting the data onto the open polyhedral described by the constraint qualification. As an interesting consequence, we observe that this projection may reduce the noise of the data, which suggests that this strategy may be used advantageously as a preliminary step at any application which extracts information from option prices. Nonetheless, even when using the projected noisy data and judiciously applying a stopping rule, the resulting probability distributions are far from the true distributions; see figures 6.6.

The instability of the minimizers of the dual problem is overcome by adding a penalty functional to the dual objective $D(\boldsymbol{\lambda}, \mathbf{d})$, as it is done in Tikhonov regularization. A sufficiently large regularization parameter α may be enough to guarantee numerical well-conditioning of the Hessian of the new objective $D_\alpha(\boldsymbol{\lambda}, \mathbf{d})$, but a more subtle selection of α is necessary in order to obtain MESs that do recover stock price distributions that are reasonable, with no spikes present. We achieve this by means of the L-curve method.

Under Tikhonov regularization and using the L-curve method for choosing the penalty parameter, we demonstrate that the method is much more stable to perturbation in data, figures 7.1 and 7.2. Table 7.1 shows that, under proper regularization, the MES is stable to data errors in simulations. As a final numerical example, we used published options data to estimate the price distribution of `Microsoft` stock, figure 7.11.

An analysis further shows that the regularized MES enjoys the classical convergence property.

References

[1] A.N. Tikhonov. Regularization of incorrectly posed problems. *Soviet Mathematics*, 4:1624–1627, 1963.

[2] D.L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *Journal of the Association for Computing Machinery*, 9:84–97, 1962.

[3] C. Hansen and D.P. O'Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14:1487–1503, 1993.

[4] P.M. van den Berg, A.L. van Broekhoven, and A. Abubakar. Extended contrast source inversion. *Inverse Problems*, 15:1325–1344, 1999.

[5] A. Abubakar and P.M. van den Berg. Total variation as a multiplicative constraint for solving inverse problems. *IEEE Transactions on Image Processing*, 10:1384–1392, 2001.

[6] A. Abubakar, P.M. van den Berg, and J.J. Mallorqui. Imaging of biomedical data using a multiplicative regularized contrast source inversion method. *IEEE Transactions on Microwave Theory and Techniques*, 50:1761–1770, 2002.

[7] A. Abubakar and P.M. van den Berg. Two- and three- dimensional algorithms for microwave imaging and inverse scattering. *Journal of Electromagnetic Waves and Applications*, 17:209–231, 2003.

[8] A. Abubakar, P.M. van den Berg, T.M. Habashy, and H. Braunisch. A multiplicative regularization approach for deblurring problems. *IEEE Transactions on Image Processing*, 13:1524–1532, 2004.

[9] A. Abubakar, P.M. van den Berg, and T.M. Habashy. An integral equation approach for 2.5-dimensional forward and inverse electromagnetic scattering. *Geophysics Journal International*, 165:744–762, 2006.

[10] C.R. Vogel. *Computational Methods for Inverse Problems*, volume 23 of *Frontiers in Applied Mathematics*. SIAM, Philadelphia, 2002.

[11] J. Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. Yale University Press, New Haven, 1923.

[12] A.N. Tikhonov and V.Y. Arsenin. *Solutions of Ill-Posed Problems*. Scripta Series in Mathematics. Winston, Washington, 1977.

[13] H. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, Boston, 1996.

[14] A. Kirsch. *An Introduction to the Mathematical Theory of Inverse Problems*. Springer-Verlag, New York, 1996.

[15] V.A. Morozov. *Regularization Methods for Ill-Posed Problems*. CRC Press, Boca Raton, FL, 1993.

[16] F. Natterer. *The Mathematics of Computerized Tomography*. Wiley, New York, 1986.

[17] M.C. Roggemann and B. Welsh. *Imaging Through Turbulence*. CRC Press, Boca Raton, FL, 1996.

[18] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 3rd. edition, 1996.

[19] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[20] D. Dobson and F. Santosa. Recovery of blocky images from noisy and blurred data. *SIAM Journal on Applied Mathematics*, 56:1181–1198, 1996.

[21] C.R. Vogel and M.E. Oman. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17:227–238, 1996.

[22] C.R. Vogel and M.E. Oman. A fast, robust algorithm for total variation based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7:813–824, 1998.

[23] R.H. Chan, T.F. Chan, and C.K. Wong. Cosine transform based preconditioners for total variation deblurring. *IEEE Transactions on Image Processing*, 8:1472–1478, 1999.

[24] M. Hanke. Limitations of the L-curve method in ill-posed problems. *BIT Numerical Mathematics*, 36:287–301, 1996.

[25] C.R. Vogel. Non-convergence of the L-curve regularization parameter selection method. *Inverse Problems*, 12:535–547, 1996.

[26] J.F.P.J. Abascal, S.R. Arridge, R.H. Bayford, and D.S. Holder. Comparison of methods for optimal choice of the regularization parameter for linear electrical impedance tomography of brain function. *Physiological Measurement*, 29:1319–1334, 2008.

[27] T. Correia, A. Gibson, M. Schweiger, and J. Hebden. Selection of regularization parameter for optical topography. *Journal of Biomedical Optics*, 14:1–11, 2009.

[28] F. Luan, C. Lee, J.H. Choi, and H.K Jung. A comparison of regularization techniques for magnetoencephalography source reconstruction. *IEEE Transactions on Magnetics*, 46:3209–3212, 2010.

[29] G.H. Golub, P.C. Hansen, and D.P. O'Leary. Tikhonov regularization and total least squares. *SIAM Journal on Matrix Analysis and Applications*, 21:185–194, 1999.

[30] T. Regińska. A regularization parameter in discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 17:740–749, 1996.

[31] L. Elden. Algorithms for the regularization of ill-conditioned ill-posed problems. *BIT*, 17:134–145, 1977.

[32] J. Cox and S. Ross. The valuation of options for alterantive stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.

[33] D. Breeden and R. Litzenberger. Prices of state contingent claims implicit in option prices. *Journal of Business*, 51:621–652, 1978.

[34] J.C. Jackwerth. *Option-Implied Risk-Neutral Distributions and Risk Aversion*. Research Foundation of AIMR, Charlottesville, VA, 2004.

[35] D.M. Lin and E.K. Wong. A survey on the maximum entropy method and parameter spectral estimation. *Physics Reports*, 193:41–135, 1990.

[36] P.W. Buchen and M. Kelley. The maximum entropy distribution of an asset inferred from option prices. *Journal on Financial and Quantitative Analysis*, 31:143–159, 1996.

[37] O. Bondarenko. Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics*, 116:85–112, 2003.

[38] M. Avellaneda. The minimum-entropy algorithm and related methods for calibrating asset-pricing models. In *Proceedings of the International Congress of Mathematicians*, volume 3, pages 545–563, Berlin, 1998. Documenta Mathematica.

[39] M. Stutzer. A simple nonparametric approach to derivative security valuation. *Journal of Finance*, 51:1633–1652, 1996.

[40] M. Avellaneda, C. Friedman, R. Holmes, and D. Samperi. Calibrating volatility surfaces via relative entropy minimization. *Applied Mathematical Finance*, 4:37–64, 1997.

[41] C. Neri and L. Schneider. Maximum entropy distributions inferred from option portfolios on an asset. preprint (2009), available at `http://arxiv.org/pdf/0903.4542`.

[42] L.S. Rompolis. Retrieving risk-neutral densities from European option prices based on the principle of maximum entropy. *Journal of Empirical Finance*, pages 1–20, 2010.

[43] J.M. Borwein, P. Maréchal, and R. Choksi. Probability distributions of assets inferred from option prices via the principle of maximum entropy. *SIAM Journal on Optimization*, 4:464–478, 2003.

[44] J. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley, New York, 1991.

[45] J. Nocedal. Private communication. 2009.

[46] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, New York, 2nd edition, 2006.

[47] J.C. Hull. *Options, Futures, and Other Derivatives*. Prentice-Hall, New Jersey, 6th edition, 2006.

[48] E. Resmerita and R.S. Anderssen. Joint additive Kullback-Leibler residual minimization and regularization for linear inverse problems. *Mathematical Methods in the Applied Sciences*, 30:1527–1544, 2007.