NEW RESULTS IN OPTIMIZATION WITH FUNCTIONAL CONSTRAINTS

BY GÁBOR RUDOLF

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Operations Research

Written under the direction of

Professor Farid Alizadeh

and approved by

New Brunswick, New Jersey

October, 2009

ABSTRACT OF THE DISSERTATION

New Results in Optimization with Functional Constraints

by Gábor Rudolf

Dissertation Director: Professor Farid Alizadeh

We consider optimization problems featuring nonnegativity constraints on functions associated with the decision variables. The first part of the thesis is devoted to the topic of optimization over the cone of positive polynomials. We present a method for constructing non-negative spline approximations to the arrival rate of a non-homogenous Poisson process based on observed arrival data, along with numerical results and a comparison to previous approaches. Our results are obtained by formulating the problem as a semidefinite program; we explore the theoretical obstacles to a more direct method by proving that only a constant number of linearly independent bilinear optimality conditions exist for cones of positive polynomials, regardless of dimension.

In the second part we look at optimization with second order stochastic dominance constraints. Here the functional inequalities appear naturally, featuring the integral of the distribution function of a random variable defined by the decision variables. We develop new duality results as well as cutting plane methods that are shown to perform well on a class of portfolio optimization problems. Finally we point out an interesting connection, arising as part of our duality considerations, between the theory of measures with given marginals and network feasibility. This results in a new proof of Strassen's Theorem from its trivial discrete case.

Acknowledgements

I would like to thank the members of my dissertation committee. I am particularly indebted to Professor Andrzej Ruszczyński, who initiated and guided much of the research presented in the second half of this thesis.

The results of this thesis are based on joint papers with Farid Alizadeh, Andrzej Ruszczyński, Jonathan Eckstein, Nilay Noyan and Dávid Papp. The relevant publications are cited in each chapter.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis we present various approaches to analyzing and solving optimization problems of the following general form:

$$\begin{aligned} \max \quad & c(x) \\ \text{s.t.} \quad & \varphi_x(t) \geq 0 \ \ \forall t \in \mathcal{T} \\ & x \in \mathcal{X}, \end{aligned} \tag{1.1}$$

where $\mathcal{X}$ and $\mathcal{T}$ are arbitrary sets, $c : \mathcal{X} \to \mathbb{R}$ is a given objective function, and the map $\varphi : \mathcal{X} \to \mathbb{R}^{\mathcal{T}}$ assigns a real-valued function $\varphi(x) = \varphi_x$ to every value of the variable x. A large variety of interesting problems can be cast in this setting, including the following examples, the first and third of which we are later going to examine in detail:

- The set $\mathcal{X}$ consists of real vectors $\mathbf{x}$ to which we assign a polynomial (or spline) function $\varphi_{\mathbf{x}}$ with coefficients $\mathbf{x}$. The functional constraint in (1.1) expresses the nonnegativity of said polynomial (or spline) over the set $\mathcal{T}$.

- The set $\mathcal{X}$ consists of real-valued random variables X to which we assign $\varphi_X = F_Y - F_X$, where F_X is the distribution function of X, and F_Y is the distribution function of some 'reference' random variable Y. If we set $\mathcal{T} = \mathbb{R}$, the functional constraint expresses the first order stochastic dominance relation $X \succeq_{(1)} Y$.

- In the previous example replace the definition of φ_X by

  $$\varphi_X(t) = \int_{-\infty}^{t} F_Y(s) - F_X(s)\, ds.$$

  The functional constraint now expresses the second order dominance relation $X \succeq_{(2)} Y$.

If both $\mathcal{X}$ and $\mathcal{T}$ are subsets of finite dimensional real vector spaces, (1.1) becomes equivalent to the well-known *semi-infinite programming* problem with a finite number of variables and infinitely many constraints. Although general purpose solution methods exist (for an exhaustive overview see [RR98]), the special structure of the actual problems encountered in our computational work can always be exploited to provide an equivalent description of the feasible sets with finite set of convex constraints.

1.1 Overview of results

In Chapter 2 we describe a method to approximate the arrival rate function of a non-homogeneous Poisson process based on observed arrival data. We estimate the function by cubic splines, using an optimization model based on the maximum likelihood principle. A critical feature of the model is that the splines are constrained to be everywhere nonnegative. These constraints are enforced by using a characterization of nonnegative polynomials by positive semidefinite matrices. We also describe versions of our model that allow for periodic arrival rate functions and input data of limited time precision. We formulate the estimation problem as a convex nonlinear program, and solve it with standard nonlinear optimization packages. Numerical results using both an actual record of e-mail arrivals over a period of sixty weeks, and artificially generated data sets are provided, along with a performance comparison to previous approaches. We also present a cross-validation procedure for determining an appropriate number of spline knots to model a set of arrival observations.

While the previous application shows the usefulness of optimization over the cone of polynomials, solving such problems via semidefinite programming (SDP) has a few drawbacks. Firstly, currently available SDP packages (such as SeDuMi, [Stu01]) do not support nonlinear objective functions, despite the fact that concave maximization subject to SDP constraints appears to be a tractable problem. Secondly, transforming the problem to an SDP leads to a quadratic increase in the number of variables.

These issues motivate us to look for a direct approach to solving these problems. Efficient primal-dual interior point methods for linear programming, second order cone programming and SDP make use of complementary slackness conditions which are bilinear in the primal and dual variables. Unfortunately, while such conditions always exist for symmetric cones, this turns out not to be the case for cones of positive polynomials. In Chapter 3 we examine some well-known non-symmetric cones, including the cone of positive polynomials over the real line and its dual, the closure of the moment cone. We show that the number of linearly independent bilinear optimality conditions is constant (4 in the case of positive polynomials over $\mathbb{R}$), regardless of the dimension of the cones.

In Chapter 2 we explored non-negative splines as a tractable and rich class of functions suitable for applications which require non-parametric approximation subject to functional constraints. However, such constraints also arise naturally in other contexts. In particular, stochastic dominance constraints require inequalities to hold between distribution functions, or integrals of distribution functions, of random variables.

In Chapter 4 we develop a new form of duality theory, featuring measures on the product of the probability space and the real line, for stochastic optimization problems with second order stochastic dominance constraints. We present two formulations involving small numbers of variables and exponentially many constraints: primal and dual. The dual formulation reveals connections between dominance constraints, generalized transportation problems, and the theory of measures with given marginals. Both formulations lead to classes of cutting plane methods. Finite convergence of both methods is proved in the case of finitely many events. Numerical results for a portfolio problem are provided.

Strassen's Theorem about the existence of measures with given marginals ([Str65, Theorem 6]) is fundamental to the duality developed in Chapter 4. For the discrete case explored in our computational work this theorem specializes to a simple and well-known network feasibility condition. Surprisingly, using discretization techniques this elementary combinatorial result allows us to derive the continuous case of Strassen's Theorem. This alternative proof is presented in Chapter 5.

1.2 Sources

Chapter 2 is based on [AENR08]. The main results of Chapter 3 have first appeared in [NRA05], but the current simplified presentation is taken from [NRAP] (which also contains some recent results not included in this thesis). Chapter 4 has been published as [RR08], while Chapter 5 first appears here.

Since each chapter is self-contained, we have opted to everywhere keep the original notation from the sources mentioned above.

Chapter 2

Arrival Rate Approximation by Nonnegative Cubic Splines

2.1 Introduction

This chapter describes a method for constructing a smooth approximation of the arrival rate of a nonhomogeneous Poisson process, based on observed arrival data. This estimation problem has applications in most situations where nonhomogeneous Poisson processes can be used, including the management of database replicas [GE01, GES03] or website mirrors. Our basic methodology may also be applied to a variety of other statistical function-estimation problems.

The essence of the problem is as follows: we observe an arrival process for a time interval $[\underline{t}, \overline{t}]$, and observe a vector of arrival times $\mathbf{t} = (t_1, \ldots, t_n)$, where $\underline{t} \leq t_1 < \cdots < t_n \leq \overline{t}$. We believe that these arrivals have been generated by a nonhomogeneous Poisson process with unknown arrival rate function $\lambda : [\underline{t}, \overline{t}] \to \mathbb{R}_+$, and wish to estimate the smooth nonnegative function $\lambda(\cdot)$ based on the data $\mathbf{t}$. We use cubic splines to model $\lambda(\cdot)$, enforcing nonnegativity by taking advantage of a representation of nonnegative polynomials using semidefinite matrices. This approach leads to optimization problems with convex objective functions, linear constraints, and semidefinite cone constraints; because of their special form, the semidefinite constraints may also be viewed as second order cone constraints.

Abstractly, our problem consists of choosing a function $\lambda(\cdot)$ from some set of candidates Λ in some way that best fits the data $\mathbf{t}$; note that we consider the number of arrivals n to be part of the data embodied in $\mathbf{t}$. Approaches to such problems may be classified as either *parametric* or *nonparametric*, depending on the dimension of Λ. In parametric approaches, Λ is explicitly finite dimensional, with a form typically determined by application-specific modeling. In nonparametric approaches, Λ is an infinite-dimensional set, but the estimation procedure picks an estimate for $\lambda(\cdot)$ from a finite-dimensional subset $\overline{\Lambda} \subset \Lambda$ which it determines using the data $\mathbf{t}$.

To find $\lambda(\cdot)$ within $\overline{\Lambda}$, the nonparametric procedure solves one or more finite-dimensional optimization problems whose size depends on the data $\mathbf{t}$. In parametric approaches, by contrast, the dimensionality of the optimization problem is fixed by the choice of parametric model, without reference to the observed data $\mathbf{t}$. In this chapter, we take the nonparametric approach and determine the number of spline knots, and thus the dimension of our optimization problem, by a cross-validation procedure; see §2.4.1 below.

One of the simplest nonparametric approaches is to let Λ be the set of piecewise constant functions over $[\underline{t}, \overline{t}]$. In this case, we choose a set of *knots* $a_0, a_1, \ldots, a_m$, with $\underline{t} = a_0 < a_1 < \cdots < a_m = \overline{t}$, and require that $\lambda(\cdot)$ be constant on each interval $[a_{i-1}, a_i)$. Once the a_i are fixed, the assumed Poisson process becomes homogeneous within each interval, with an easily estimated arrival rate. Even in the special case of evenly spaced knots, it can be shown that setting m sufficiently large allows one to approximate to arbitrary precision any function one could reasonably consider for $\lambda(\cdot)$ (technically, any square integrable function). In practice, m should not be chosen too large relative to the amount of input data, so as to avoid overfitting.

Another standard nonparametric approach is to assume $\lambda(\cdot)$ is piecewise linear and continuous. Here, we again choose knots $\underline{t} = a_0 < a_1 < \cdots < a_m = \overline{t}$, but assume that $\lambda(\cdot)$ is affine on each interval $[a_{i-1}, a_i]$, and continuous at all the interval boundaries a_i. Once the a_i have been fixed, this approach yields a convex optimization problem with linear equality and inequality constraints, solvable by standard methods of mathematical programming; see for example [TT90] for the formulation of a closely related problem. It can be shown that a large enough m allows arbitrarily close approximation of a wide range of continuous functions on $[\underline{t}, \overline{t}]$ (technically, any absolutely continuous function with a square integrable derivative).

For values of m appropriate to realistic amounts of sample data, such simple piecewise-constant or piecewise-linear models tend to produce estimates of $\lambda(\cdot)$ having abrupt changes in the arrival rate or its derivative. Poisson arrival models typically arise from the aggregate behavior of large numbers of independent actors; arrival rate functions with large jumps or sharp nondifferentiabilities are appropriate when there are discrete events that simultaneously affect all actors, or a significant subset of actors. Examples of such events include the start or close of trading on a stock exchange, or maintenance shutdown of a web or e-mail server. Arrival rate estimation is this kind of environment should involve either *a priori* knowledge of the times of such events, or a data sample sufficiently large to make them clearly observable.

By contrast, we focus on constructing smoother estimates of process arrival rates. Smooth arrival rate functions are appropriate when there are time-correlated but imperfectly synchronized changes in the population of actors giving rise to the Poisson process. Consider, for example, a 24-hour support center for a business-oriented product. The incidence of calls to the support center should increase at the beginning of the normal work day; however, not all customers will start work at precisely 9 AM, as would be implicitly assumed by a $\lambda(\cdot)$ function with a discontinuous step increase at that time. Smoothly varying $\lambda(\cdot)$ functions provide more realistic and internally consistent models for such cases.

Here, we therefore choose the set Λ to contain only functions of a certain smoothness. *Splines* are natural and standard tools to use in such situations: for a spline of order $k \geq 0$, we again select knots $\underline{t} = a_0 < a_1 < \cdots < a_m = \bar{t}$, and assume that $\lambda(\cdot)$ is a polynomial of degree k on each interval. That is, with the convention $0^0 = 1$,

$$\lambda(t) = p^{(i)}(t) = \sum_{\ell=0}^{k} p_\ell^{(i)} (t - a_{i-1})^\ell \tag{2.1}$$

whenever $t \in [a_{i-1}, a_i]$, where the $p_\ell^{(i)}$, $i = 1, \ldots, m$, $\ell = 0, \ldots, k$, are real coefficients linearly constrained so that $\lambda(\cdot)$ has continuous derivatives of order up to $k - 1$; see §2.3. Note that the piecewise-constant model is essentially the special case $k = 0$, and the piecewise-linear model is the special case $k = 1$. It can be shown that, even with the restriction of evenly spaced knots, choosing m large enough allows an order-k spline to approximate to arbitrary precision any member of the *Sobolev-Hilbert* space $\mathcal{H}^{(k)}$, consisting of functions whose derivatives of order smaller than k exist and are absolutely continuous, and whose derivative of order k is square integrable; see for example [Wah90], [TT90], or [Sch81]. Thus, unless m is fixed independently of the data $\mathbf{t}$, this approach should be considered nonparametric. Much of our analysis applies to splines of arbitrary order, but we focus particularly on the cubic spline case $k = 3$—see for example [dB78]—as cubic splines possess a desirable combination of simplicity and versatility.

Any reasonable choice of $\lambda(\cdot)$ should be nonnegative throughout $[\underline{t}, \bar{t}]$; in fact, in §2.3 we show that without such a restriction, maximum-likelihood estimation of $\lambda(\cdot)$ may become an ill-posed problem. In terms of the coefficients $p_\ell^{(i)}$ of the spline polynomials, however, the condition that $\lambda(t) \geq 0$ for all $t \in [\underline{t}, \bar{t}]$ may not take an immediately obvious form. In our approach, we enforce nonnegativity by applying a characterization of nonnegative polynomials related to semidefinite programming (SDP) and second-order cone programming (SOCP), both areas of intensified research activity in the mathematical programming community over the last decade. The theory of nonnegative polynomials is at least a century old; the text by [KS66] develops a complete version of this theory from which the semidefinite characterization of the cone of nonnegative univariate polynomials and nonnegative univariate trigonometric polynomials is easily derived. [Nes00] has extended this SDP representation to the class of possibly multivariate functions that are weighted sums of squares of other functions. Using the characterization described in Karlin and Studden, we formulate the maximum-likelihood estimation of a nonnegative cubic spline function $\lambda(\cdot)$ from the observed arrival data $\mathbf{t}$. Our maximum-likelihood approach is also easily modified to use aggregated input data: instead of a vector of exact arrival times $\mathbf{t}$, we have a set of counts of arrivals during some arbitrary set of disjoint time intervals. In either the exact or aggregated case, the resulting nonlinear optimization problem is convex, with a unique global optimum. With some adjustments to improve the model's scaling properties (see §2.3.3), the optimal $\lambda(\cdot)$ can be found routinely by nonlinear programming software packages such as KNITRO [NW03], LOQO [BVS02], and IPOPT [WB06], all of which are available through NEOS servers [CMM98, Dol01, GM97].

We should mention at this point an alternative approach to enforcing nonnegativity that does not require SDP/SOCP methods. Suppose one can identify a basis $b_1(\cdot), \ldots, b_{\tilde{m}}(\cdot)$ of span$(\overline{\Lambda})$ such that each $b_i(\cdot)$ is nonnegative throughout $[\underline{t}, \bar{t}]$. Then, for any scalars $\alpha_1, \ldots, \alpha_{\tilde{m}} \geq 0$, we have $\alpha_1 b_1(t) + \cdots + \alpha_{\tilde{m}} b_{\tilde{m}}(t) \geq 0$ for all $t \in [\underline{t}, \bar{t}]$. Thus, it is possible to enforce nonnegativity by defining $\lambda(t) = \alpha_1 b_1(t) + \cdots + \alpha_{\tilde{m}} b_{\tilde{m}}(t)$ and letting $\alpha_1, \ldots, \alpha_{\tilde{m}}$ be the optimization decision variables, under the simple constraints $\alpha_1, \ldots, \alpha_{\tilde{m}} \geq 0$. *B-splines* [dB78, Sch81] are an example of this technique, forming a natural nonnegative basis for the space of spline functions.

The drawback of these approaches is that they impose tighter constraints than are actually required. For example, suppose $\overline{\Lambda}$ consists of the nonnegative splines over $[\underline{t}, \bar{t}]$ with knots $\underline{t} = a_0 < a_1 < \cdots < a_m = \bar{t}$. Then $\overline{\Lambda}$ is a convex cone, but not a polyhedral cone. Taking nonnegative linear combinations of a basis $b_1(\cdot), \ldots, b_{\tilde{m}}(\cdot)$ for span$(\overline{\Lambda})$, on the other hand, only yields a polyhedral cone, which must hence be *strictly* contained in $\overline{\Lambda}$. For a very simple example of such a situation, consider the nonnegative basis $\{1, 1-t, t^2\}$ for all quadratic polynomials on $[0, 1]$. The set of polynomials of the form $\alpha_1 + \alpha_2(1 - t) + \alpha_3 t^2$, where $\alpha_1, \alpha_2, \alpha_3 \geq 0$, does not contain the nonnegative polynomial $1 + t + t^2 = 2 - (1 - t) + t^2$. To characterize nonnegative functions in a linear space, neither nonnegativity of the basis functions nor of the coefficients is necessary. Our SDP/SOCP techniques, without increasing the number of parameters, use convex nonlinear constraints to work directly and exactly with cones of nonnegative splines.

Because solutions to the nonlinear optimization models we formulate can be found efficiently, it is practical to solve a collection of related problems using different numbers of knots m and different subsets of the sample data. This capability allows us to determine an appropriate value of m by using a cross-validation procedure we describe in §2.4.1. The basic structure of this procedure is fairly generic and could be applied to a variety of related models.

In summary, the principal contribution of our work is to point out that, with careful use of the theory of nonnegative polynomials and techniques from SDP/SOCP and computational nonlinear programming, estimation of arrival functions by nonnegative cubic splines, including determination of an appropriate number of knots, is practical. There is no need to restrict the search by constraints that are strictly stronger than function nonnegativity.

In the §2.2 and §2.3, respectively, we present the necessary background information on nonhomogeneous Poisson processes, maximum likelihood estimators, cubic splines and nonnegative polynomials, ending with the formulation of our estimation problem as a convex optimization model. In the course of the development, we also discuss variants of the problem in which the arrival data are aggregated, the splines are constrained to be periodic functions, or both. §2.3.4 outlines our procedure for determining the most appropriate number of spline knots. Next, §2.4 addresses validating our approach and demonstrating its practicality, and gives more details of our cross-validation procedure. Our numerical experiments use both a real-world e-mail arrival dataset and artificially generated datasets. In §2.4.3, we also provide numerical comparisons of our methodology to some existing approaches to arrival rate estimation whose basic properties are described in §2.1.2. Finally, §2.5 presents some conclusions and possible directions for future work.

In the remainder of this section, we briefly describe how our techniques can be extended to related function estimation problems, and then compare our approach to other arrival rate estimation techniques.

2.1.1 Related function estimation problems

Regression and shape constraints.

In regression problems, we are given a set of points $x_1, \ldots, x_n$ in an interval $[\underline{x}, \overline{x}]$, and possibly noisy observations $y_1, \ldots, y_n$ of respective function values of $f(x_1), \ldots, f(x_n)$. From this information, we wish to estimate the unknown function $f(\cdot)$. Various error models are possible for the y_j, depending on the application: one simple case is the standard homoscedastic normal-error model $y_j = f(x_j) + \epsilon_j$, where $\epsilon_1, \ldots, \epsilon_n$ are independent normal random variables with mean 0 and standard deviation σ. We wish to find the function $f(\cdot)$ that best fits the data $(x_1, y_1), \ldots, (x_n, y_n)$; various notions of fit can be used, depending on the error model. For example, applying the maximum likelihood principle to the normal-error model above results in the least-squares fitting criterion min $\sum_{i=1}^{n} (y_i - f(x_i))^2$. A common parametric approach to this problem is to select some known functions $f_1(\cdot), \ldots, f_m(\cdot) : [\underline{x}, \overline{x}] \to \mathbb{R}$, and adopt the model $f(x) = \theta_1 f_1(x) + \cdots + \theta_m f_m(x)$. With a least-squares fitting criterion, estimation of $(\theta_1, \ldots, \theta_m)$ then reduces to a multiple linear regression problem, which may be solved by classical computational techniques.

Several nonparametric approaches are also common. For example, one could use a piecewise-constant or piecewise-linear model for $f(\cdot)$, with a variable number of knots m determined from the input data $(x_1, y_1), \ldots, (x_n, y_n)$. The piecewise-linear approach is described in [TT90]. Nonparametric spline regression models of order $k > 1$ have also been extensively investigated; see for example [Wah90]. Wahba's work presents two ways to avoid overfitting: a penalty approach and a cross-validation method similar to the techniques described in this chapter.

The most significant contribution of our methodology to regression applications arises in the case of regression subject to *shape constraints* that restrict the form of the $f(\cdot)$. Such constraints can be applied in either the parametric or nonparametric case. For example, we may wish $f(\cdot)$ to be nonnegative throughout $[\underline{x}, \overline{x}]$, to be either nondecreasing or nonincreasing, or to be either convex or concave. All these constraints are examples of the general form $Lf(x) \geq 0$ for all $x \in [\underline{x}, \overline{x}]$, where L is some linear functional operator:

- For $f(\cdot)$ to be nonnegative throughout $[\underline{x}, \overline{x}]$, we set L to the identity mapping.
- For $f(\cdot)$ to be nondecreasing, we let $L = D$, the differential operator. For nonincreasing $f(\cdot)$, we let $L = -D$
- For $f(\cdot)$ to be convex, we let $L = D^2$, the second derivative operator. For concave $f(\cdot)$, we let $L = -D^2$.

The methodology described in this chapter could easily be adapted to perform regression to spline functions under such shape constraints. With further modification, it could also be applied in other contexts, such as regression to trigonometric polynomials under shape constraints.

Density or distribution estimation.

In the common and fundamental problem of density estimation, we are given a sample $x_1, \ldots, x_n$ of observed values of some continuous random variable X, and wish to estimate either its probability density function $f_X(\cdot)$ or its cumulative distribution function $F_X(\cdot)$ over $[\underline{x}, \overline{x}]$, where $\underline{x} \leq x_1, \ldots, x_n \leq \overline{x}$. As in arrival rate estimation, $f_X(\cdot)$ must be everywhere nonnegative; analogously, $F_X(\cdot)$ must be nondecreasing. If it is acceptable for the estimate of $f_X(\cdot)$ or $F_X(\cdot)$ to be discontinuous or have discontinuous derivatives, then we may employ piecewise-constant or piecewise-linear nonparametric models, with the number of breakpoints m determined appropriately from the sample data $\mathbf{x} = (x_1, \ldots, x_n)$. As with arrival rate estimation, the piecewise-constant model leads to a closed-form estimation formula, while the piecewise-linear model leads to a relatively straightforward convex programming problem, as in [Sco76] and [STT80]. If a smoother estimate of $f_X(\cdot)$ or $F_X(\cdot)$ is desired, then the problem may be solved by techniques similar to this chapter's, with some minor variations: for $f_X(\cdot)$, the maximum likelihood objective function is simpler, taking the form $\max \sum_{j=1}^{n} \ln f_X(x_j)$, and one must also require $\int_{\underline{x}}^{\overline{x}} f_X(x)\, dx = 1$, which may be enforced by a linear constraint on the spline coefficients. For the cumulative distribution, we must impose the nondecreasing shape constraint $DF_X(x) \geq 0$ for all $x \in [\underline{x}, \overline{x}]$, as well as $F_X(\overline{x}) = 1$.

2.1.2 Comparison to other Poisson rate estimation methods

First, we note that it is possible to replace the problem of estimating the arrival rate $\lambda(\cdot)$ of a nonhomogeneous Poisson process with the equivalent problem of estimating its *mean value function* $\overline{N}(\cdot)$, defined by

$$\overline{N}(t) = E\big[N(t)\big] = \int_0^t \lambda(t)\,dt,$$

where $N(t)$ denotes the number of arrivals up to time t. The requirement that $\lambda(\cdot)$ be nonnegative is equivalent to the constraint that $\overline{N}(\cdot)$ be nondecreasing.

In a series of related papers, Kuhl, Damerdji, Wilson, Johnson, Bhairgond, and Sumant have studied Poisson arrival rate estimation from various perspectives. The work described in [KDW98], [KSW06], and [KW01] uses a multiresolution procedure involving nested periodicities; at each resolution, the authors fit a normalized and aggregated version of the observed accumulated arrival process $N(\cdot)$ with a polynomial function (but not a spline). On the other hand, [KDW98] and [KWJ97] assume a particular parametric model, namely

$$\lambda(t) = \lambda(t;\, \mathbf{p}, \boldsymbol{\gamma}, \boldsymbol{\omega}, \boldsymbol{\phi}) = \exp\left(\sum_{i=0}^{k_1} p_i t^i + \sum_{j=1}^{k_2} \gamma_j \sin(\omega_j t + \phi_j) \right), \tag{2.2}$$

where $\mathbf{p} \in \mathbb{R}^{k_1}$ and $\boldsymbol{\gamma}, \boldsymbol{\omega}, \boldsymbol{\phi} \in \mathbb{R}^{k_2}$. This formula includes periodic effects of unknown period and known shape; by contrast, our principal computational experiments on e-mail data in §2.4.1 below model periodic effects of known period but unknown shape, an approach we consider more likely to be useful in practice for a given number of parameters. The main difference between [KDW98] and [KWJ97] is that the latter work uses a maximum likelihood approach, while the former applies a least squares method to $\overline{N}(\cdot)$.

Note that the exponential function in (2.2) assures positivity of $\lambda(\cdot)$ without the need for any explicit constraints on the parameters. This technique is convenient for modeling multiplicative effects, but may also have some drawbacks: it makes it difficult to model additive effects, and the arrival intensity can never be exactly zero. Qualitatively, the rapid growth of the exponential produces $\lambda(\cdot)$ estimates that tend to have sharp peaks and shallow valleys; see for example Figure 2.4.3 below. The corresponding log likelihood function is concave in terms of the parameters $\mathbf{p}$ and $\boldsymbol{\gamma}$, but it is nonconcave in $\boldsymbol{\omega}$ and $\boldsymbol{\phi}$, leading to a computationally difficult estimation problem for the parameters.

[KB00] employ wavelets to estimate $\lambda(\cdot)$, adapting the nonnegative wavelet estimation of general density functions proposed by [WS98]. In this approach, the arrival rate is expressed as a nonnegative linear combination of nonnegative wavelets, with the coefficients obtained by simple recurrence formulas. As with the nonnegative basis technique discussed above, the drawback of such methods is that they impose more restrictions on $\lambda(\cdot)$ than are actually needed to obtain nonnegativity: in fact, neither nonnegativity of the basis wavelets nor of the wavelet coefficients are required to produce a nonnegative function. On the other hand, the Walter-Shen approach results in closed form formulas, or at worst unconstrained optimization, and is therefore computationally much simpler than our constrained convex nonlinear programming approach.

In comparison to the work described above, one deficiency of our current work is that it does not explicitly combine periodic behavior with a long term trend; we either model an arbitrary function over $[\underline{t}, \overline{t}]$, or a function that has a certain known periodicity. In principle, there are various ways our approach could be augmented to combine periodic behavior with long term trend effects; we leave the investigation of the these alternatives for future research.

To gauge the relative effectiveness of our approach, §2.4.3 applies it to simulated data generated from the model (2.2). Overall, our methodology compares favorably to the results reported in [KB00] and [KDW98].

2.2 Nonhomogeneous Poisson Likelihood Functions

Our models' objective functions are based on the maximum likelihood principle. Consequently, we begin by presenting the maximum likelihood functions arising from the nonhomogeneous Poisson model. For the moment, assume that we are trying to select an estimated arrival rate function $\lambda(\cdot)$ from some arbitrary given set $\overline{\Lambda}$. Later, we refine our models for specific forms of $\overline{\Lambda}$ based on nonnegative splines.

2.2.1 Individual arrival data

Given t_{j-1}, a choice of $\lambda(\cdot)$ assigns, through the nonhomogeneous Poisson model, a probability density to t_j, namely

$$f_j(t_j, \lambda) = \lambda(t_j) \exp\left(-\int_{t_{j-1}}^{t_j} \lambda(t)\, dt\right).$$

Fixing $t_0 = \underline{t}$, the joint probability density of $(t_1, \dots, t_n)$ is then

$$f(\mathbf{t}, \lambda) = \prod_{i=1}^{n} f_j(t_j, \lambda) = \prod_{i=1}^{n} \lambda(t_j) \exp\left(-\int_{t_{j-1}}^{t_j} \lambda(t)\, dt\right).$$

To make use of the information that there were no arrivals during $(t_n, \bar{t}]$, we multiply $f(\mathbf{t}, \lambda)$ by the probability that a Poisson random variable with mean $\int_{t_n}^{\bar{t}} \lambda(t)\, dt$ has the value 0, yielding

$$\tilde{f}(\mathbf{t}, \lambda) = \exp\left(-\int_{t_n}^{\bar{t}} \lambda(t)\, dt\right) f(\mathbf{t}, \lambda) = \exp\left(-\int_{t_n}^{\bar{t}} \lambda(t)\, dt\right) \prod_{i=1}^{n} \lambda(t_j) \exp\left(-\int_{t_{j-1}}^{t_j} \lambda(t)\, dt\right).$$

The maximum likelihood principle suggests that we choose $\lambda \in \overline{\Lambda}$ to maximize $\tilde{f}(\mathbf{t}, \lambda)$, or equivalently $L(\mathbf{t}, \lambda) = \ln \tilde{f}(\mathbf{t}, \lambda)$. Using the convention $\ln(0) = -\infty$, this *log likelihood function* is

$$L(\mathbf{t}, \lambda) = \sum_{j=1}^{n} \left(\ln \lambda(t_j) - \int_{t_{j-1}}^{t_j} \lambda(t)\, dt\right) - \int_{t_n}^{\bar{t}} \lambda(t)\, dt = \sum_{j=1}^{n} \ln \lambda(t_j) - \int_{\underline{t}}^{\bar{t}} \lambda(t)\, dt. \tag{2.3}$$

2.2.2 Aggregated arrival data

In many practical situations, we may not have *exact* arrival time information. Instead, we may only have access to data of the following aggregated form: given some times $q_0 < q_1 < \cdots < q_\nu$, we know the number of arrivals n_j in each interval $(q_{j-1}, q_j]$, but not the exact arrival times within these intervals. Here, we apply the maximum likelihood principle in the following way: an arrival rate function $\lambda : [q_0, q_\nu] \to \mathbb{R}_+$ and the Poisson model assign a probability of

$$P(n_j, q_{j-1}, q_j, \lambda) = \frac{1}{n_j!} \left(\int_{q_{j-1}}^{q_j} \lambda(t)\, dt\right)^{n_j} \exp\left(-\int_{q_{j-1}}^{q_j} \lambda(t)\, dt\right)$$

to the occurrence of n_j arrivals in $(q_{j-1}, q_j]$. Letting $\mathbf{n} = (n_1, \dots, n_\nu)$ and $\mathbf{q} = (q_0, \dots, q_\nu)$, the joint probability of the arrival pattern $\mathbf{n}$ is

$$P(\mathbf{n}, \mathbf{q}, \lambda) = \prod_{j=1}^{\nu} P(n_j, q_{j-1}, q_j, \lambda).$$

Again, the maximum likelihood principle suggests choosing $\lambda \in \overline{\Lambda}$ to maximize $P(\mathbf{n}, \mathbf{q}, \lambda)$, or equivalently $L_d(\mathbf{n}, \mathbf{q}, \lambda) = \ln P(\mathbf{n}, \mathbf{q}, \lambda)$. Simplifying L_d, we obtain

$$L_d(\mathbf{n}, \mathbf{q}, \lambda) = \sum_{j=1}^{\nu} \left(n_j \ln\left(\int_{q_{j-1}}^{q_j} \lambda(t)\, dt \right) - \ln n_j! \right) - \int_{q_0}^{q_\nu} \lambda(t)\, dt. \tag{2.4}$$

Note that the terms $\ln n_j!$ are independent of $\lambda(\cdot)$, and may therefore be ignored when performing the optimization $\max_{\lambda \in \overline{\Lambda}} L_d(\mathbf{n}, \mathbf{q}, \lambda)$. Our experimental data in Section 2.4 are in this form.

2.2.3 Periodic arrival rate functions

In many situations, it is reasonable to assume that the arrival rate follows a repetitive periodic pattern. For example, events tied to the rhythm of the work week should exhibit a repeating weekly arrival pattern. Formally, the assumption of such a pattern means that we restrict λ to take the form $\lambda(t) = \omega(t \,(\mathrm{mod}\, T))$ for some period $T > 0$ and $\omega : [0, T) \to \mathbb{R}_+$. For simplicity, assume that we are given the arrival times $0 < t_1 < \ldots < t_n < cT$ for the time period $[0, cT]$, where c is a positive integer. Rewriting the log-likelihood function (2.3) in terms of $\omega(\cdot)$, one obtains

$$L(\mathbf{t}, \omega) = \sum_{j=1}^{n} \ln \omega\big(t_j \,(\mathrm{mod}\, T)\big) - c \int_0^T \omega(t)\, dt. \tag{2.5}$$

For aggregated data in the interval $[q_0, q_\nu] = [0, cT]$, we may similarly rewrite the log-likelihood function (2.4) in terms of ω as

$$L_d(\mathbf{n}, \mathbf{q}, \omega) = \sum_{j=1}^{\nu} \left[n_j \ln\left(\int_{q_{j-1}}^{q_j} \omega\big(t \,(\mathrm{mod}\, T)\big)\, dt \right) - \ln n_j! \right] - c \int_0^T \omega(t)\, dt. \tag{2.6}$$

Supposing that we have some arbitrary set $\overline{\Omega}$ from which we wish to select $\omega(\cdot)$, the maximum likelihood estimation problems for pointwise or aggregated data are then respectively

$$\max_{\omega \in \overline{\Omega}} \{L(\mathbf{t}, \omega)\} \qquad \text{and} \qquad \max_{\omega \in \overline{\Omega}} \{L_d(\mathbf{n}, \mathbf{q}, \omega)\}.$$

Below, we consider the specific cases of these formulations in which $\overline{\Omega}$ arises from periodic cubic spline functions.

2.3 Nonnegative Cubic Splines

We now consider a specific choice for the set $\overline{\Lambda}$ from which to select the arrival rate functions $\lambda(\cdot)$. As suggested in §2.1, we fix real numbers $\underline{t} = a_0 < a_1 < \cdots < a_m = \bar{t}$, and let $\overline{\Lambda}$ be the set of nonnegative cubic splines with knots $a_0, a_1, \ldots, a_m$. The $4m$ real numbers $p_\ell^{(i)}$, $i = 0, \ldots, m-1$, $\ell = 0, \ldots, 3$, determine such a spline. With $k = 3$, (2.1) reduces to

$$\lambda(t) = p^{(i)}(t) = \sum_{\ell=0}^{3} p_\ell^{(i)} \big(t - a_{i-1}\big)^\ell \qquad \forall\, t \in [a_{i-1}, a_i], \tag{2.7}$$

with the convention $0^0 = 1$. The $p_\ell^{(i)}$ are the main decision variables in our optimization model. To be a spline, the resulting function must be continuous and have continuous derivatives up to order $k - 1 = 2$, meaning that it must obey, for $i = 1, \ldots, m-1$, and $d_i = a_i - a_{i-1}$, the linear equations

$$p_0^{(i+1)} = p_0^{(i)} + p_1^{(i)} d_i + p_2^{(i)} d_i^2 + p_3^{(i)} d_i^3 \tag{2.8}$$

$$p_1^{(i+1)} = p_1^{(i)} + 2p_2^{(i)} d_i + 3p_3^{(i)} d_i^2 \tag{2.9}$$

$$2p_2^{(i+1)} = 2p_2^{(i)} + 6p_3^{(i)} d_i. \tag{2.10}$$

In our numerical experiments, we use evenly spaced knots $a_i = \underline{t} + (i/m)(\bar{t} - \underline{t})$, in which case $d_i = (\bar{t} - \underline{t})/m$ for all $i = 1, \ldots, m$. This choice of equidistant knots is simply for convenience, and is not a fundamental requirement in our approach. Finally, we determine m using the cross-validation procedure described in §2.3.4 and §2.4.1.

Consider now the periodic models of §2.2.3, in which we have $\lambda(t) = \omega\big(t \,(\mathrm{mod}\, T)\big)$. To apply a spline model in this case, given knots $0 = a_0 < a_1 < \cdots < a_m = T$, we define ω analogously to (2.7). However, in order for $\lambda(\cdot)$ to be continuous and twice continuously differentiable at $T, 2T, \ldots, cT$, we must also have $\omega(T) = \omega(0)$, $\omega'(T) = \omega'(0)$, and $\omega''(T) = \omega''(0)$. Therefore, in terms of the spline coefficients $p_\ell^{(i)}$, we must have

$$p_0^{(1)} = p_0^{(m)} + p_1^{(m)} d_m + p_2^{(m)} d_m^2 + p_3^{(m)} d_m^3 \tag{2.11}$$

$$p_1^{(1)} = p_1^{(m)} + 2p_2^{(m)} d_m + 3p_3^{(m)} d_m^2 \tag{2.12}$$

$$2p_2^{(1)} = 2p_2^{(m)} + 6p_3^{(m)} d_m. \tag{2.13}$$

Functions satisfying these conditions, together with (2.8)-(2.10) for $i = 1, \ldots, m$, are called *periodic splines* with period T. We choose the set $\overline{\Omega}$ of candidates for $\omega(\cdot)$ to be the set of nonnegative m-knot periodic splines over $[0, T]$. As in the nonperiodic case, we use equidistant knots for our numerical experiments, and select the number of knots m through a cross-validation procedure.

We now turn to the question of how to constrain the $p_\ell^{(i)}$ so that the resulting arrival rate function is nonnegative throughout $[\underline{t}, \bar{t}]$. In §2.4.2, we empirically demonstrate that nonnegativity is an essential requirement: for a small data set of e-mail arrivals, without the nonnegativity assumption, the computed arrival rate does in fact take negative values. However, there is even a more fundamental problem. Without the nonnegativity assumption, the likelihood function corresponding to a cubic spline arrival rate (or any polynomial spline of order $k \geq 1$) may be unbounded. Specifically, maximizing the likelihood function (2.3) in the absence of such a nonnegativity constraint is in general ill-posed, even in a finite-dimensional linear functional space. To see how such ill-posedness can occur, consider what happens when the feasible region $\overline{\Lambda}$ for $\lambda(\cdot)$ is a cone and

$$\exists\, \mu(\cdot) \in \overline{\Lambda}: \quad \mu(t_j) > 0 \;\; \forall\, j = 1, \ldots, n, \quad \int_{\underline{t}}^{\bar{t}} \mu(t)\, dt \;\leq\; 0. \tag{2.14}$$

Now consider setting $\lambda(t) = \alpha\mu(t)$, where $\alpha > 0$; since $\overline{\Lambda}$ is a cone, such a $\lambda(\cdot)$ always lies in $\overline{\Lambda}$. By increasing α, it is easily confirmed that one can make the likelihood function (2.3) arbitrarily large. For the periodic situation (2.5), analogy to (2.14) suggests the condition that $\overline{\Omega}$ is a cone and

$$\exists\, \widetilde{\mu}(\cdot) \in \overline{\Omega}: \quad \widetilde{\mu}\big(t_j\, (\mathrm{mod}\, T)\big) > 0 \;\; \forall\, j = 1, \ldots, n, \quad \int_0^T \widetilde{\mu}(t)\, dt \;\leq\; 0. \tag{2.15}$$

Much as in the nonperiodic case, setting $\omega(t) = \alpha\widetilde{\mu}(t)$ and letting $\alpha \to \infty$ then results in an unbounded objective. Indeed, for periodic cubic splines, (2.15) is satisfied when $m = 3$ and $0 < t_j\, (\mathrm{mod}\, T) < T/3$ for all j. Similar phenomena can occur with the aggregate-data likelihood functions (2.4) and (2.6) if there are observation intervals with $n_j = 0$. Thus, it is in general necessary to require nonnegativity simply to obtain a well-defined problem, and it is therefore imperative to consider how to characterize the polynomials that are nonnegative over each segment $[a_{i-1}, a_i]$.

2.3.1 Characterization of nonnegative polynomial splines

The most important results of the well-established theory of nonnegative polynomials and associated dual moment cones are covered in the text of [KS66]. These topics are related to *semidefinite programming* (SDP) and its special case called *second order cone programming* (SOCP). For a general reference on semidefinite programming, see [WSV00] and its extensive bibliography. For a survey of SOCP, see [AG03]. Here, we present a quick review of these topics and their connection to nonnegative polynomials and splines.

Recall that a real $r \times r$ symmetric matrix $\mathbf{X}$ is *positive semidefinite*, written $\mathbf{X} \succcurlyeq 0$, if for every real r-vector $\boldsymbol{v}$, we have $\boldsymbol{v}^\top \mathbf{X} \boldsymbol{v} \geq 0$. It is clear from the definition that the set $\mathcal{M}_r$ of $r \times r$ real symmetric positive semidefinite matrices is a closed convex cone. Semidefinite programming is simply the optimization of linear functions over affine transformations and affine preimages of $\mathcal{M}_r$.

Given some inner product $\langle \cdot, \cdot \rangle$, the *dual cone* of any convex cone $\mathcal{K}$ is defined as:

$$\mathcal{K}^* = \{\boldsymbol{z} \mid \langle \boldsymbol{x}, \boldsymbol{z} \rangle \geq 0 \quad \text{for all } \boldsymbol{x} \in \mathcal{K}\}.$$

For symmetric matrices, we number rows and columns starting with 0, and employ the inner product

$$\langle \mathbf{A}, \mathbf{B} \rangle = \mathbf{A} \bullet \mathbf{B} = \sum_{i=0}^{r-1} \sum_{j=0}^{r-1} \mathbf{A}_{ij} \mathbf{B}_{ij}. \tag{2.16}$$

With this inner product, it is well known that the cone of positive semidefinite symmetric matrices is self-dual, that is, $\mathcal{M}_r = \mathcal{M}_r^*$.

If we represent a polynomial $p_0 + p_1 x + \cdots + p_n x^k$ by its vector of coefficients $\mathbf{p} = (p_0, p_1, \ldots, p_k)$, then it is clear that the nonnegative polynomials over the interval $[a, b]$ comprise a closed convex cone, since any nonnegative combination of such polynomials remains nonnegative over $[a, b]$; we denote this cone by $\mathcal{P}_{k+1}(a, b)$. Using the standard inner product for $\mathbb{R}^{k+1}$, the dual cone $\mathcal{P}_{k+1}^*(a, b)$ of $\mathcal{P}_{k+1}(a, b)$ consists of vectors that are moments of some nondecreasing function of bounded variation over $[a, b]$; see [KS66]. More precisely, if $\mathcal{F}[a, b]$ is the set of probability measures over $[a, b]$, then

$$\mathcal{P}_{k+1}^*(a, b) = \left\{ (c_0, \ldots, c_k) \;\middle|\; \exists\, \alpha > 0,\ F \in \mathcal{F}[a, b] \ :\ c_i = \alpha \int_a^b t^i \, dF(t),\ i = 0, \ldots, k \right\}.$$

$\mathcal{P}^*_{k+1}(a,b)$ is called the *moment cone* over $[a,b]$. The key property of $\mathcal{P}^*_{k+1}(a,b)$ is that it is *SD-representable* [Nem99]. This property simply means that the moment cone can be expressed as an affine preimage of the positive semidefinite cone $\mathcal{M}_r$ for some integer $r > 0$.

Specifically, if k is even, let $h = k/2$ and define the square matrices

$$\underline{\mathbf{H}}_k = (c_{i+j})_{ij}, \qquad 0 \le i,j \le h \tag{2.17}$$

$$\overline{\mathbf{H}}_k = \big((a+b)c_{i+j+1} - c_{i+j+2} - abc_{i+j}\big)_{ij}, \qquad 0 \le i,j \le h-1. \tag{2.18}$$

If k is odd, let $h = \lfloor k/2 \rfloor$, and define the $(h+1)\times(h+1)$ matrices

$$\underline{\mathbf{H}}_k = \big(c_{i+j+1} - ac_{i+j}\big)_{ij}, \qquad 0 \le i,j \le h \tag{2.19}$$

$$\overline{\mathbf{H}}_k = (bc_{i+j} - c_{i+j+1})_{ij}, \qquad 0 \le i,j \le h. \tag{2.20}$$

In either case, $\underline{\mathbf{H}}_k$ and $\overline{\mathbf{H}}_k$ are *Hankel matrices*, that is, the value of entry (i,j) is determined by $i+j$ (implying they are symmetric). From the analysis in [KS66], [DS97], and [Nes00], $\underline{\mathbf{H}}_k$ and $\overline{\mathbf{H}}_k$ are related to the moment cone as follows:

$$(c_0, c_1, \ldots, c_k) \in \mathcal{P}^*_{k+1}(a,b) \quad \Longleftrightarrow \quad \underline{\mathbf{H}}_k \succcurlyeq 0 \text{ and } \overline{\mathbf{H}}_k \succcurlyeq 0. \tag{2.21}$$

For any $h \ge 0$, let $\mathbf{E}^h_\ell$ be the $(h+1)\times(h+1)$ matrix given by

$$\left(\mathbf{E}^h_\ell\right)_{ij} = \begin{cases} 1, & i+j = \ell \\ 0, & i+j \neq \ell, \end{cases} \qquad 0 \le i,j \le h.$$

Then $\mathbf{E}^h_0, \mathbf{E}^h_1, \ldots, \mathbf{E}^h_{2h}$ form a basis for the space of $(h+1)\times(h+1)$ Hankel matrices.

Using (2.21), we can now characterize the cone $\mathcal{P}^*_{k+1}(a,b)$ and its dual $\mathcal{P}_{k+1}(a,b)$. The details differ depending on whether k is even or odd, and thus whether we employ (2.17)-(2.18) or (2.19)-(2.20).

When k is odd.

Letting $h = \lfloor k/2 \rfloor$ and rewriting (2.19) and (2.20) in terms of the basis elements $\mathbf{E}^h_0, \ldots, \mathbf{E}^h_{2h} = \mathbf{E}^h_{k-1}$, we have

$$\underline{\mathbf{H}}_k = -c_0 a\mathbf{E}^h_0 + c_1(\mathbf{E}^h_0 - a\mathbf{E}^h_1) + c_2(\mathbf{E}^h_1 - a\mathbf{E}^h_2) + \cdots + c_{k-1}(\mathbf{E}^h_{k-2} - a\mathbf{E}^h_{k-1}) + c_k\mathbf{E}^h_{k-1}$$

$$\overline{\mathbf{H}}_k = c_0 b\mathbf{E}^h_0 + c_1(b\mathbf{E}^h_1 - \mathbf{E}^h_0) + \cdots + c_{k-1}(b\mathbf{E}^h_{k-1} - \mathbf{E}^h_{k-2}) - c_k\mathbf{E}^h_{k-1}.$$

Therefore, by (2.21), the cone $\mathcal{P}^*_{k+1}(a,b)$ consists of all vectors $(c_0, c_1, \ldots, c_k)$ satisfying

$$-c_0 a\mathbf{E}^h_0 + c_1(\mathbf{E}^h_0 - a\mathbf{E}^h_1) + \cdots + c_{k-1}(\mathbf{E}^h_{k-2} - a\mathbf{E}^h_{k-1}) + c_k\mathbf{E}^h_{k-1} \succcurlyeq 0 \tag{2.22}$$

$$c_0 b\mathbf{E}^h_0 + c_1(b\mathbf{E}^h_1 - \mathbf{E}^h_0) + \cdots + c_{k-1}(b\mathbf{E}^h_{k-1} - \mathbf{E}^h_{k-2}) - c_k\mathbf{E}^h_{k-1} \succcurlyeq 0. \tag{2.23}$$

Since $\mathcal{P}_{k+1}(a,b)$ is closed and convex, we have $\mathcal{P}^{**}_{k+1}(a,b) = \mathcal{P}_{k+1}(a,b)$, so $\mathcal{P}_{k+1}(a,b)$ is simply the dual of the cone described by (2.22)-(2.23). To characterize this dual cone, we associate symmetric positive semidefinite matrices $\mathbf{X}$ and $\mathbf{Y}$ with (2.22) and (2.23), respectively. These matrices play much the same role as Lagrange multipliers in general nonlinear programming, except that they must be matrices of the same shape as the two sides of the semidefinite inequalities (2.22)-(2.23), that is, both $\mathbf{X}$ and $\mathbf{Y}$ are $(h+1)\times(h+1)$ symmetric matrices. Using the inner product defined in (2.16), we then argue that $(p_0, p_1, \ldots, p_k)$ is in $\mathcal{P}_{k+1}(a,b)$ whenever there exist $\mathbf{X}, \mathbf{Y} \succcurlyeq 0$ such that

$$\begin{aligned}
p_0 &= -a\mathbf{E}^h_0 \bullet \mathbf{X} + b\mathbf{E}^h_0 \bullet \mathbf{Y}\\
p_1 &= (\mathbf{E}^h_0 - a\mathbf{E}^h_1) \bullet \mathbf{X} + (b\mathbf{E}^h_1 - \mathbf{E}^h_0) \bullet \mathbf{Y}\\
p_2 &= (\mathbf{E}^h_1 - a\mathbf{E}^h_2) \bullet \mathbf{X} + (b\mathbf{E}^h_2 - \mathbf{E}^h_1) \bullet \mathbf{Y}\\
&\vdots\\
p_\ell &= (\mathbf{E}^h_{\ell-1} - a\mathbf{E}^h_\ell) \bullet \mathbf{X} + (b\mathbf{E}^h_\ell - \mathbf{E}^h_{\ell-1}) \bullet \mathbf{Y}\\
&\vdots\\
p_k &= \mathbf{E}^h_{k-1} \bullet \mathbf{X} - \mathbf{E}^h_k \bullet \mathbf{Y}.
\end{aligned}$$

Formally, this result may be obtained by applying the following fact: suppose the cone $C \subset \mathbb{R}^r$ is defined by $C = \{\boldsymbol{u} \in \mathbb{R}^r \mid \boldsymbol{A}\boldsymbol{u} \in \mathcal{K}\}$, where $\boldsymbol{A}$ is an $s \times r$ matrix and $\mathcal{K} \subset \mathbb{R}^s$ is some closed convex cone; then $C^* = \{\boldsymbol{A}^\top \boldsymbol{w} \mid \boldsymbol{w} \in \mathcal{K}^*\}$. This result follows from the separating hyperplane property of convex sets and is subsumed, for example, by [Roc70, Corollary 16.3.2].

When k is even.

In this case, we can let $h = k/2$ and apply an analysis similar to (2.17)-(2.18) and (2.21), resulting in the characterization that $(p_0, \dots, p_k) \in \mathcal{P}_{k+1}(a, b)$ if and only if

$$\begin{aligned}
p_0 &= \mathbf{E}_0^h \bullet \mathbf{X} - ab\mathbf{E}_0^{h-1} \bullet \mathbf{Y} \\
p_1 &= \mathbf{E}_1^h \bullet \mathbf{X} + \left((a+b)\mathbf{E}_0^{h-1} - ab\mathbf{E}_1^{h-1}\right) \bullet \mathbf{Y} \\
p_2 &= \mathbf{E}_2^h \bullet \mathbf{X} + \left(-\mathbf{E}_0^{h-1} + (a+b)\mathbf{E}_1^{h-1} - ab\mathbf{E}_2^{h-1}\right) \bullet \mathbf{Y} \\
&\vdots \\
p_\ell &= \mathbf{E}_\ell^h \bullet \mathbf{X} + \left(-\mathbf{E}_{\ell-2}^{h-1} + (a+b)\mathbf{E}_{\ell-1}^{h-1} - ab\mathbf{E}_\ell^{h-1}\right) \bullet \mathbf{Y} \\
&\vdots \\
p_k &= \mathbf{E}_k^h \bullet \mathbf{X} - \mathbf{E}_{k-2}^{h-1} \bullet \mathbf{Y} \\
\mathbf{X} &\succcurlyeq 0 \\
\mathbf{Y} &\succcurlyeq 0,
\end{aligned}$$

where the symmetric positive semidefinite matrices $\mathbf{X}$ and $\mathbf{Y}$ have dimension $(h+1) \times (h+1)$ and $h \times h$, respectively.

Cubic polynomials and shifted representations.

Here, it is convenient to represent a polynomial over $[a, b]$ by $p(x) = p_0 + p_1(x - a) + p_2(x - a)^2 + \cdots + p_n(x - a)^n$. In this case, $p(x)$ is nonnegative over $[a, b]$ if and only if $p_0 + p_1 t + p_2 t^2 + \cdots + p_n t^n$ is nonnegative over $[0, b - a]$, so the representations given above can be modified by replacing a with zero and b with $d = b - a$.

In particular, consider the cone $\mathcal{P}$ of cubic polynomials $p(x) = p_0 + p_1(x - a) + p_2(x - a)^2 + p_3(x - a)^3$ that are nonnegative over $[a, b]$. Consider (2.22)-(2.23) specialized to the case $k = 3$ and thus $h = \lfloor 3/2 \rfloor = 1$. Replacing $a \leftarrow 0$ and $b \leftarrow d$, we conclude that the vector (c_0, c_1, c_2, c_3) is in the dual cone $\mathcal{P}_3^*(0, d)$ if and only if

$$\begin{pmatrix} c_1 & c_2 \\ c_2 & c_3 \end{pmatrix} \succcurlyeq 0 \quad \text{and} \quad \begin{pmatrix} dc_0 - c_1 & dc_1 - c_2 \\ dc_1 - c_2 & dc_2 - c_3 \end{pmatrix} \succcurlyeq 0$$

Specializing the analysis of §2.3.1 to $k = 3$, we observe that the cubic polynomial $p_0 + p_1(x -$

$a) + p_2(x-a)^2 + p_3(x-a)^3$ is nonnegative on $[a,b]$ whenever there are 2×2 matrices

$$\mathbf{X} = \begin{pmatrix} x & y \\ y & z \end{pmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{pmatrix} s & v \\ v & w \end{pmatrix}$$

satisfying

$$\begin{aligned}
p_0 &= d(\mathbf{E}_0^1 \bullet \mathbf{Y}) && \Longleftrightarrow && p_0 = ds \\
p_1 &= \mathbf{E}_0^1 \bullet \mathbf{X} + (d\mathbf{E}_1^1 - \mathbf{E}_0^1) \bullet \mathbf{Y} && \Longleftrightarrow && p_1 = x + 2dv - s \\
p_2 &= \mathbf{E}_1^1 \bullet \mathbf{X} + (d\mathbf{E}_2^1 - \mathbf{E}_1^1) \bullet \mathbf{Y} && \Longleftrightarrow && p_2 = 2y + dw - 2v \\
p_3 &= \mathbf{E}_2^1 \bullet \mathbf{X} + -\mathbf{E}_2^1 \bullet \mathbf{Y} && \Longleftrightarrow && p_3 = z - w \\
\mathbf{X} &\succcurlyeq 0 && \Longleftrightarrow && x, z \geq 0, \quad \mathrm{Det}(\mathbf{X}) = xz - y^2 \geq 0 \\
\mathbf{Y} &\succcurlyeq 0 && \Longleftrightarrow && s, w \geq 0, \quad \mathrm{Det}(\mathbf{Y}) = sw - v^2 \geq 0.
\end{aligned}$$

In this case, because $\mathbf{X}$ and $\mathbf{Y}$ are 2×2, the positive semidefiniteness constraints $\mathbf{X}, \mathbf{Y} \succcurlyeq 0$ can be reformulated as the linear and quadratic constraints $x, z, s, w \geq 0$, $xz - y^2 \geq 0$, and $sw - v^2 \geq 0$. This form of the constraints makes it possible to use standard nonlinear programming software. Incidentally, we note that the quadratic constraints may be considered to be second order cone programming (SOCP) constraints. Second Order Cone Programming (SOCP) involves optimization over affine transformations and affine preimages of the second order, or *Lorentz*, cone, namely

$$\mathcal{Q}_{r+1} = \left\{ (x_0, x_1, \ldots, x_r) \mid x_0 \geq \sqrt{x_1^2 + \cdots + x_r^2} \right\}$$

Specifically, when $x, z, s, w \geq 0$, we may express the quadratic inequalities $xz - y^2 \geq 0$ and $sw - v^2 \geq 0$ as

$$\left(\frac{x+z}{2}\right)^2 \geq \left(\frac{x-z}{2}\right)^2 + y^2 \quad \Longleftrightarrow \quad \begin{pmatrix} \frac{x+z}{2} \\ \frac{x-z}{2} \\ y \end{pmatrix} \in \mathcal{Q}_3, \tag{2.24}$$

$$\left(\frac{s+w}{2}\right)^2 \geq \left(\frac{s-w}{2}\right)^2 + v^2 \quad \Longleftrightarrow \quad \begin{pmatrix} \frac{s+w}{2} \\ \frac{s-w}{2} \\ v \end{pmatrix} \in \mathcal{Q}_3. \tag{2.25}$$

2.3.2 Practical optimization models

The nonnegativity of a cubic spline with coefficients $\{p_\ell^{(i)}\}$ and knots $a_0, \dots, a_m$ is equivalent to the nonnegativity of each polynomial $p^{(i)}(t) = \sum_{\ell=0}^{3} p_\ell^{(i)}(t - a_{i-1})^\ell$ over the respective interval $[a_{i-1}, a_i]$. Summarizing the results of the last section, we thus have

Theorem 1 $p^{(i)}(t) = p_0^{(i)} + p_1^{(i)}(t - a_{i-1}) + p_2^{(i)}(t - a_{i-1})^2 + p_3^{(i)}(t - a_{i-1})^3 \geq 0$ *for all* $t \in (a_{i-1}, a_i)$ *if and only if there exist* $x_i, y_i, z_i, s_i, v_i, w_i \in \mathbb{R}$ *such that, for* $d_i = a_i - a_{i-1}$,

$$p_0^{(i)} = d_i s_i \tag{2.26}$$

$$p_1^{(i)} = x_i - s_i + 2d_i v_i \tag{2.27}$$

$$p_2^{(i)} = 2y_i - 2v_i + d_i w_i \tag{2.28}$$

$$p_3^{(i)} = z_i - w_i \tag{2.29}$$

$$x_i z_i \geq y_i^2 \tag{2.30}$$

$$s_i w_i \geq v_i^2 \tag{2.31}$$

$$x_i, z_i, s_i, w_i \geq 0. \tag{2.32}$$

Theorem 1 completes our characterization of the sets $\overline{\Lambda}$ and $\overline{\Omega}$ of (periodic) nonnegative splines over which we wish to maximize the likelihood functions (2.3), (2.4), (2.5), or (2.6). First, consider (2.5). For $j = 1, \dots, n$, define $i_j = \operatorname{argmin}_{i=1,\dots,m}\{a_i : a_i \geq t_j\,(\mathrm{Mod}\,T)\}$. Inserting the definition of a periodic cubic spline into (2.5), we obtain the likelihood function

$$L(\mathbf{t}, \mathbf{p}) = \sum_{j=1}^{n} \ln\left[\sum_{\ell=0}^{3} p_\ell^{(i_j)}\big(t_j\,(\mathrm{Mod}\,T) - a_{i_j-1}\big)^\ell\right] - c\sum_{i=1}^{m}\sum_{\ell=0}^{3} p_\ell^{(i)}\frac{d_i^{\ell+1}}{\ell+1}. \tag{2.33}$$

Here and below, we let $\mathbf{p}$ represent the $4m$-vector of spline coefficients $p_\ell^{(i)}$, $\ell = 0, \dots, 3$, $i = 1, \dots, m$. The constraints on $\mathbf{p}$ are just the spline continuity and differentiability restrictions (2.8)-(2.13) and the nonnegativity constraints summarized in Theorem 1, where we introduce additional variables $x_i, y_i, z_i, s_i, v_i, w_i$, $i = 1, \dots, m$. Thus, we have the optimization problem

$$
\begin{aligned}
\max \quad & \sum_{j=1}^{n} \ln\left[\sum_{\ell=0}^{3} p_\ell^{(i_j)}\left(t_j\,(\mathrm{Mod}\,T) - a_{i_j-1}\right)^\ell\right] - c\sum_{i=1}^{m}\sum_{\ell=0}^{3} p_\ell^{(i)}\frac{d_i^{\ell+1}}{\ell+1} \\
\text{s.t.} \quad & \text{Constraints (2.8)-(2.10) hold,} \qquad i = 1,\dots,m-1 \\
& \text{Constraints (2.11)-(2.13) hold} \\
& \text{Constraints (2.26)-(2.32) hold,} \qquad i = 1,\dots,m.
\end{aligned}
\tag{2.34}
$$

The non-periodic version of the problem—maximizing (2.3) over $\overline{\Lambda}$—would amount to deleting "$(\mathrm{Mod}\,T)$", "c", and constraints (2.11)-(2.13).

Note that the objective function of (2.34) is concave in the decision variables, since all its terms are either linear in $\mathbf{p}$ or logarithms of linear functions of $\mathbf{p}$. The constraints (2.8)-(2.10) and (2.11)-(2.13) are also linear in the decision variables $\mathbf{p}$. As for the spline nonnegativity constraints (2.26)-(2.32), it is clear that (2.26)-(2.29) and (2.32) are linear in the decision variables $\mathbf{p}$ and $x_i, y_i, z_i, s_i, v_i, w_i$, $i = 1, \dots, m$. This leaves only the constraints (2.30)-(2.31), which are nonlinear, but describe a convex region, since they are equivalent to the semidefinite constraints $\mathbf{X}, \mathbf{Y} \succcurlyeq 0$ and SOCP constraints of the form (2.24)-(2.25). Thus, (2.34) describes the maximization of a concave objective function over a convex region. As a result, any local maximum is necessarily global. Except for its concave nonlinear objective function, the problem is an SOCP. General results of [NN94] imply that it is tractable to optimize (2.34) by standard interior point methods.

The same observations also apply when the data are aggregated, and the likelihood function take the form (2.6). In this case, the constraints are identical to those of (2.34), as are the linear terms $\sum_{i=1}^{m}\sum_{\ell=0}^{3} p_\ell^{(i_j)} d_i^{\ell+1}/(\ell+1)$ in the objective function. The logarithmic objective terms corresponding to $\sum_{j=1}^{\nu} n_j \ln\left(\int_{q_{j-1}}^{q_j} \omega(t\,(\mathrm{mod}\,T))\,dt\right)$ in (2.6) take a slightly different form than the logarithmic terms in (2.33), and may in general be quite complicated. Clearly, however, the arrival rate function $\lambda(t) = \omega\big(t\,(\mathrm{mod}\,T)\big)$ is a linear function of the parameters $\mathbf{p}$, and consequently the definite integrals $\int_{q_{j-1}}^{q_j} \lambda(t)\,dt$ are also linear in $\mathbf{p}$. Thus, when viewed as a function of $\mathbf{p}$, (2.6) has the same basic form as (2.33)—a sum of linear functions and logarithms of linear functions—and is therefore concave. We conclude that the maximum likelihood problem with aggregated data has the same general properties as (2.34).

A relatively simple case of the aggregate model is when the aggregation points $q_0, \ldots, q_\nu$ are chosen to align within spline segments, that is, for each $j = 1, \ldots, \nu$, there exists an $i(j) \in \{1, \ldots, m\}$ and an integer $r(j) \geq 0$ such that $[q_{j-1}, q_j] \subseteq [r(j)T + a_{i(j)-1}, r(j)T + a_{i(j)}]$. Then, dropping the constant terms $\ln n_j!$, we obtain the model

$$\begin{aligned}
\max \quad & \sum_{j=1}^{\nu} n_j \ln\left(\sum_{\ell=0}^{3} p_\ell^{(i(j))} \frac{(q_j - q_{j-1})^{\ell+1}}{\ell+1}\right) - c \sum_{i=1}^{m} \sum_{\ell=0}^{3} p_\ell^{(i)} \frac{d_i^{\ell+1}}{\ell+1} \\
\text{s.t.} \quad & \text{Constraints (2.8)-(2.10) hold,} \qquad i = 0, \ldots, m-1 \\
& \text{Constraints (2.11)-(2.13) hold} \\
& \text{Constraints (2.26)-(2.32) hold,} \qquad j = 1, \ldots, m.
\end{aligned} \tag{2.35}$$

The nonperiodic case of maximizing (2.4) is obtained by deleting the constant c and constraints (2.11)-(2.13).

While we have focused on cubic splines, it is clear that by employing the theory described in §2.3.1 and appropriately generalizing the constraints (2.8)-(2.10) and (2.11)-(2.13), our methodology could readily be extended to higher-order or quadratic splines. The main computational difference would be that for splines of order higher than 3, the positive semidefinite constraints $\mathbf{X}, \mathbf{Y} \succcurlyeq 0$ would not reduce so simply to nonlinear functional constraints like (2.30)-(2.31), so one would have to maximize a concave objective over a set of linear and semidefinite constraints. However, the basic convexity properties of the problem would still hold. From now on, we consider only cubic splines.

2.3.3 Improved scaling

Although models like (2.34) and (2.35) are in principle readily solvable to global optimality by standard nonlinear programming tools, they have scaling properties that can make them numerically challenging. Specifically, if the problem is formulated so that the d_i are significantly different in magnitude from 1, the scaling of the optimal spline coefficients $p_\ell^{(i)}$ can vary dramatically with ℓ. For example, if d_i is on the order of 0.01, then $p_3^{(i)}$ can easily be on the order of 10^6 times larger than $p_0^{(i)}$, and sometimes more. For higher-order splines, this effect would be even more pronounced. In practice, these scaling issues can make models like (2.34) and (2.35) difficult to solve: of the solvers we had available, only KNITRO [BHN99, NW03] was able to solve models of this form to optimality.

Fortunately, these scaling difficulties may be overcome by expressing the spline in a different way. In the spline construction we have used so far, we translate the $[0, d_i]$ portion of the i^{th} polynomial by a_{i-1}, and use it as the i^{th} portion of the spline. To acheive better scaling, we define m polynomials

$$u^{(i)}(x) = \sum_{\ell=0}^{3} u_\ell^{(i)} x^\ell \qquad\qquad i = 1, \ldots, m,$$

where $u_\ell^{(i)} \in \mathbb{R}$, $\ell = 0, \ldots, 3$, $i = 1, \ldots, m$, and translate and scale the $[0, 1]$ portion of $u^{(i)}$ to obtain the i^{th} segment of the spline. Thus, whenever $t \in [a_{i-1}, a_i]$, we have

$$\lambda(t) = u^{(i)}\!\left(\frac{t - a_{i-1}}{a_i - a_{i-1}}\right) = u^{(i)}\!\left(\frac{t - a_{i-1}}{d_i}\right) = \sum_{\ell=0}^{3} u_\ell^{(i)} \left(\frac{t - a_{i-1}}{d_i}\right)^{\ell}.$$

For $\lambda(t)$ to be everywhere nonnegative, we need each polynomial $u^{(i)}$ to be nonnegative on $[0, 1]$. Applying Theorem 1 with $d_i = 1$ and eliminating the variable $s_i = u_i^{(0)}$ from (2.26), we obtain the equivalent conditions

$$\begin{aligned}
u_1^{(i)} &= x_i - s_i + 2v_i & \qquad\qquad x_i z_i &\geq y_i^2 \\
u_2^{(i)} &= 2y_i - 2v_i + w_i & s_i w_i &\geq v_i^2 \\
u_3^{(i)} &= z_i - w_i & u_0^{(i)}, x_i, z_i, w_i &\geq 0.
\end{aligned}$$

After some routine manipulations, we obtain the formulation in Figure 2.3.3, equivalent to (2.34). Note that in the case of equidistant knots, we have $d_{i+1}/d_i = 1$ for all $i = 1, \ldots, m$, along with $d_1/d_m = 1$, so the d_i do not appear in the constraints. For aggregated data, we similarly obtain a model with the objective function

$$\max \quad \sum_{j=1}^{\nu} n_j \ln\left(\sum_{\ell=0}^{3} u_\ell^{(i(j))} \frac{(q_j - q_{j-1})^{\ell+1}}{(\ell+1)\, d_i^\ell}\right) - \, c \sum_{i=1}^{m} d_i \sum_{\ell=0}^{3} \frac{u_\ell^{(i)}}{\ell+1}, \tag{2.36}$$

and the same constraints as in Figure 2.3.3. Unlike (2.34) and (2.35), such scaled models were solved to optimality by all the nonlinear programming solvers we attempted to use, including IPOPT [WB06], LOQO [BVS02], and KNITRO [BHN99, NW03]. Furthermore, although KNITRO could also solve the unscaled models, it converged more rapidly on the scaled models. Nonperiodic variants of the scaled models may be obtained by deleting model elements in an analogous manner to the unscaled versions.

$$
\begin{array}{llcll}
\max & \displaystyle\sum_{j=1}^{n} \ln\left[\sum_{\ell=0}^{3} u_\ell^{(i_j)}\left(\frac{t_j\,(\mathrm{Mod}\,T) - a_{i_j-1}}{d_i}\right)^{\ell}\right] - \; c\sum_{i=1}^{m} d_i \sum_{\ell=0}^{3}\frac{u_\ell^{(i)}}{\ell+1} & & & \\
\text{s.t.} & u_0^{(i+1)} & = & u_0^{(i)} + u_1^{(i)} + u_2^{(i)} + u_3^{(i)} & i = 1,\ldots,m-1 \\
 & (d_{i+1}/d_i)\,u_1^{(i+1)} & = & u_1^{(i)} + 2u_2^{(i)} + 3u_3^{(i)} & i = 1,\ldots,m-1 \\
 & 2(d_{i+1}/d_i)^2\,u_2^{(i+1)} & = & 2u_2^{(i)} + 6u_3^{(i)} & i = 1,\ldots,m-1 \\
 & u_0^{(1)} & = & u_0^{(m)} + u_1^{(m)} + u_2^{(m)} + u_3^{(m)} & \\
 & (d_1/d_m)\,u_1^{(1)} & = & u_1^{(m)} + 2u_2^{(m)} + 3u_3^{(m)} & \\
 & 2(d_1/d_m)^2\,u_2^{(1)} & = & 2u_2^{(m)} + 6u_3^{(m)} & \\
 & u_1^{(i)} & = & x_i - s_i + 2v_i & i = 1,\ldots,m \\
 & u_2^{(i)} & = & 2y_i - 2v_i + w_i & i = 1,\ldots,m \\
 & u_3^{(i)} & = & z_i - w_i & i = 1,\ldots,m \\
 & x_i z_i & \geq & y_i^2 & i = 1,\ldots,m \\
 & s_i w_i & \geq & v_i^2 & i = 1,\ldots,m \\
 & u_0^{(i)}, x_i, z_i, w_i & \geq & 0 & i = 1,\ldots,m
\end{array}
$$

Figure 2.1: Scaled version of full periodic optimization model with exact arrival times.

2.3.4 Cross-validation approaches for determining m

We now consider the problem of choosing the number of spline knots m. As discussed in §2.1, our method should only be considered nonparametric if the choice of m corresponds in a reasonable way to the input data $\mathbf{t}$ or $(\mathbf{n}, \mathbf{q})$. Qualitatively, too small a value of m does not fully exploit the available data, whereas too large a value results in overfitting: rapidly fluctuating rate estimates overly tailored to individual arrival events in the dataset. In fact, for exact arrival time data $\mathbf{t}$, allowing m to grow indefinitely results in splines that resemble linear combinations of Dirac δ distributions: such arrival rate estimates in effect assign positive probability to the observed arrival times t_j and zero probability to other times. A similar effect can occur for aggregated input data if it is of sufficiently high resolution.

Here, we propose using a cross-validation method to determine m from the input data. Roughly speaking, cross validation proceeds as follows: we start with a relatively small value of m, say m_1. Next, we choose a random subset of arrival times $\mathcal{T}$ and set them aside. We solve the optimization problem described above for the remaining set of arrivals, with $m = m_1$. Once the estimated spline function is determined, we calculate its likelihood function value for the arrival times in the set $\mathcal{T}$. We may repeat this process R times, each for a different random subset $\mathcal{T}$, and calculate the average value $\overline{L}(m_1)$ of the likelihood function over this sample. Next, we increase the number of knots to some $m_2 > m_1$, and repeat the above process, obtaining an average likelihood $\overline{L}(m_2)$. After similarly testing a sequence of values $m_1, m_2, \ldots, m_M$, we select the number of knots which maximizes $\overline{L}(m_i)$. In our numerical experiments, we apply a variant of this technique known as K-*folding*; details are described in §2.4.1.

In general, as in the case of nonparametric regression and nonparametric density estimation, cross-validation procedures are most effective when the sample size is large. It should be noted that each time a step of the cross-validation process is applied, a different instance of the optimization problem described in this section must be solved. Thus, having fast methods to solve the nonnegative maximum likelihood spline arrival rate problem is essential.

2.4 Numerical Experiments

2.4.1 An e-mail arrival dataset

In this section, we present results obtained for a dataset of approximately 10,000 e-mail arrivals recorded over a period of 446 days. The recorded arrival times are rounded to the nearest integer second, resulting in aggregated data with one-second intervals. Thus, we used model (2.35), or in scaled form, the model of Figure 2.3.3 and (2.36). Examination of the data revealed a clear weekly periodicity: see Figures 2.4.1-2.4.1 below, where the jagged lines show a standard piecewise-constant approximation to $\lambda(\cdot)$ using 64 intervals. The optimization models are expressed in the AMPL modeling language [FGK02a], and solved by KNITRO [BHN99, NW03] on the NEOS servers [CMM98, Dol01, GM97].

Determining the number of knots.

The specific variant of cross validation we use to determine the number of knots is K-folding, also known as the leave-one-out cross-validation method [Sha93].

In our implementation of K-folding, we randomly divide the observation period into K regions $D_1, \dots, D_K$ of equal size. Then, for each $i = 1, \dots, K$, we estimate the arrival rate, omitting D_i from the model input. Next, by evaluating the appropriate log-likelihood function given by (2.6), we examine how well the estimated arrival intensity function $\lambda(\cdot)$ describes the behavior of the process during D_i. Specifically, we evaluate $L_d\big(\mathbf{n}(i), \mathbf{q}(i), \lambda^i\big)$, where λ^i is the arrival intensity function estimate derived from all the subsets *except* D_i, and $\mathbf{n}(i)$ and $\mathbf{q}(i)$ are derived *only* from D_i.

For each choice under consideration for the number of spline knots m, we perform the above procedure R times, randomly selecting $D_1, \dots, D_K$ differently each time. Thus, for each possible value of m, we obtain RK different values of the likelihood $L_d\big(\mathbf{n}(i), \mathbf{q}(i), \lambda^i\big)$, whose average is $\overline{L}(m)$. Among the values tried for m, we then select the one which maximizes $\overline{L}(m)$. The entire procedure requires RKM solutions of our optimization model, where M is the number of different values considered for m.

For the e-mail dataset, Figure 2.4.1 shows the average likelihood level results for $K = 5$ and $R = 10$; thus, each point in the figure is the average of $RK = 50$ observations. The tested

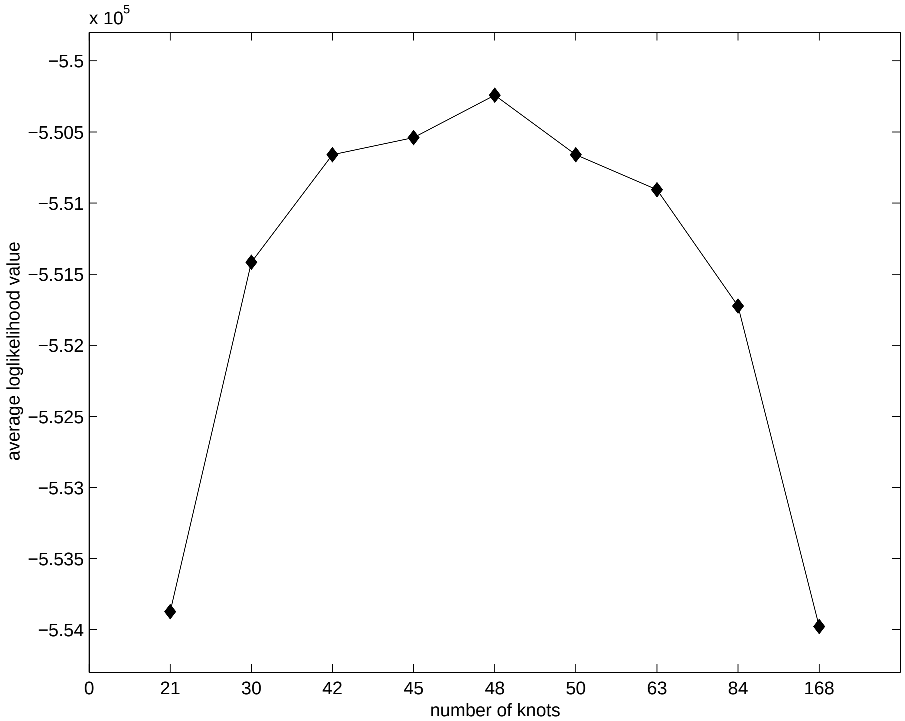

Figure 2.2: Average likelihoods for 10 different 5-foldings, large e-mail dataset.

values of m are 21, 30, 42, 45, 48, 50, 63, 84, and 168. Based on these results, we use 48 knots. Figure 2.4.1 shows the estimate using 14 knots, which does not provide sufficient detail to describe the arrival process, and Figure 2.4.1 shows the results for 48 knots. Figure 2.4.1 shows the estimate using 336 knots, which appears to overfit the data. In Figures 2.4.1-2.4.1, as well as in Figures 2.4.1, 2.4.2, and 2.4.2, the time axis is measured in weeks.

Sensitivity to data aggregation.

We also consider the sensitivity of the results to the degree of aggregation in the input data. In the results of Figures 2.4.1-2.4.1, we set $q_j - q_{j-1} = 1$ second, the recorded precision of the dataset, for all j. Starting with this data representation, we successively merge adjacent intervals, aggregating their arrival information. Setting $q_j - q_{j-1} = s$ for all j, we evaluate various aggregation intervals s ranging from the original one second to 24 hours. We write $\mathbf{n}(s)$ and $\mathbf{q}(s)$ for vectors resulting from data aggregation with interval length s. For each value of s, we recompute the spline estimate of $\lambda(\cdot)$ with 48 knots, denoted $\lambda^{[s]}(\cdot)$, and record the value of the log likelihood function $L_d\big(\mathbf{n}(s), \mathbf{q}(s), \lambda^{[s]}\big)$.

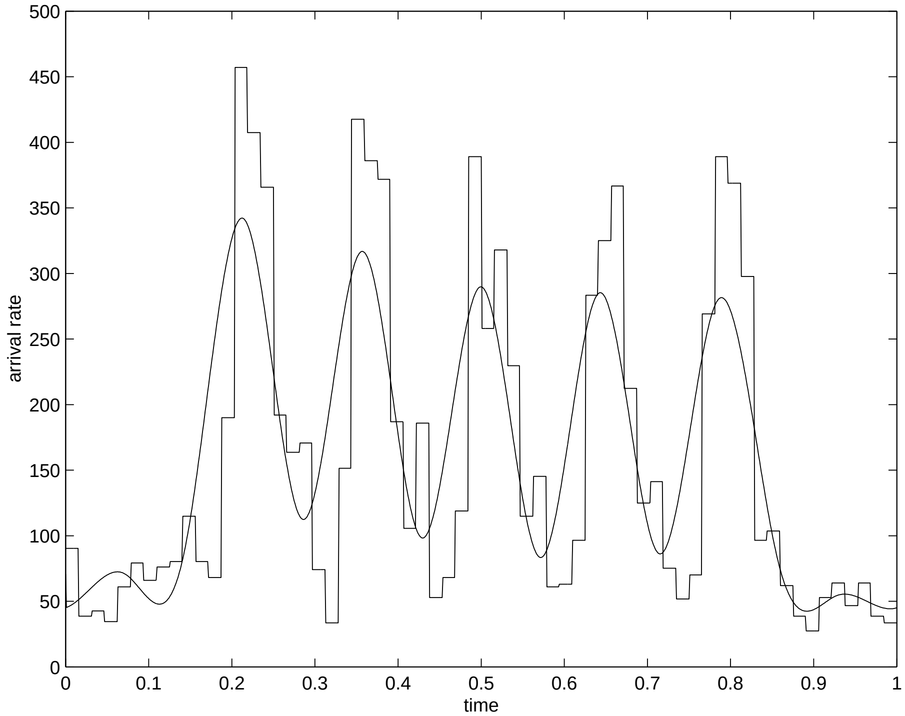

Figure 2.3: 14-knot approximation for the large e-mail dataset.

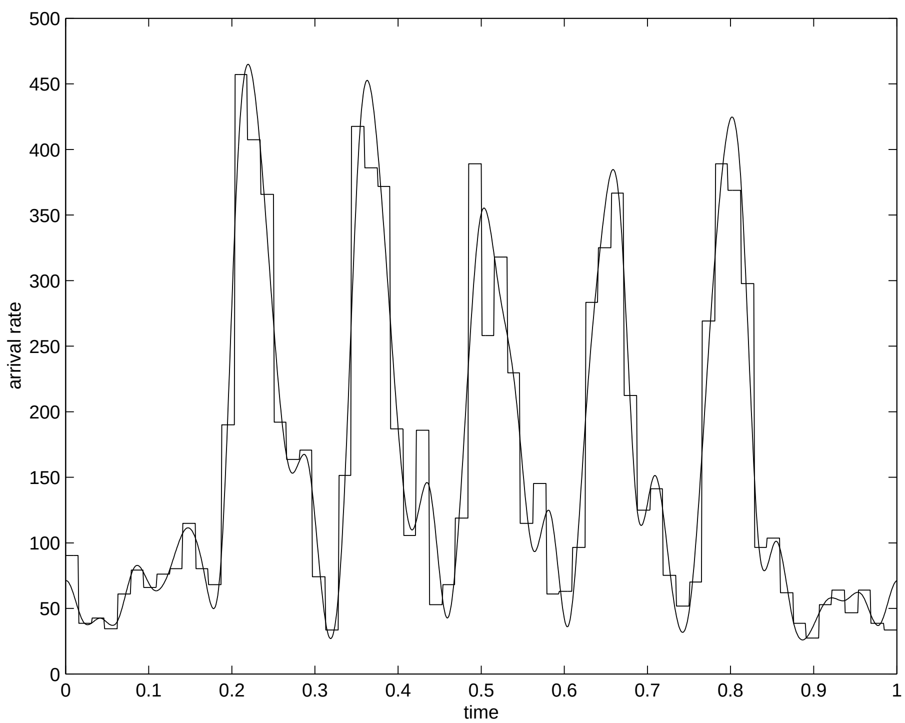

Figure 2.4: 48-knot approximation for the large e-mail dataset.

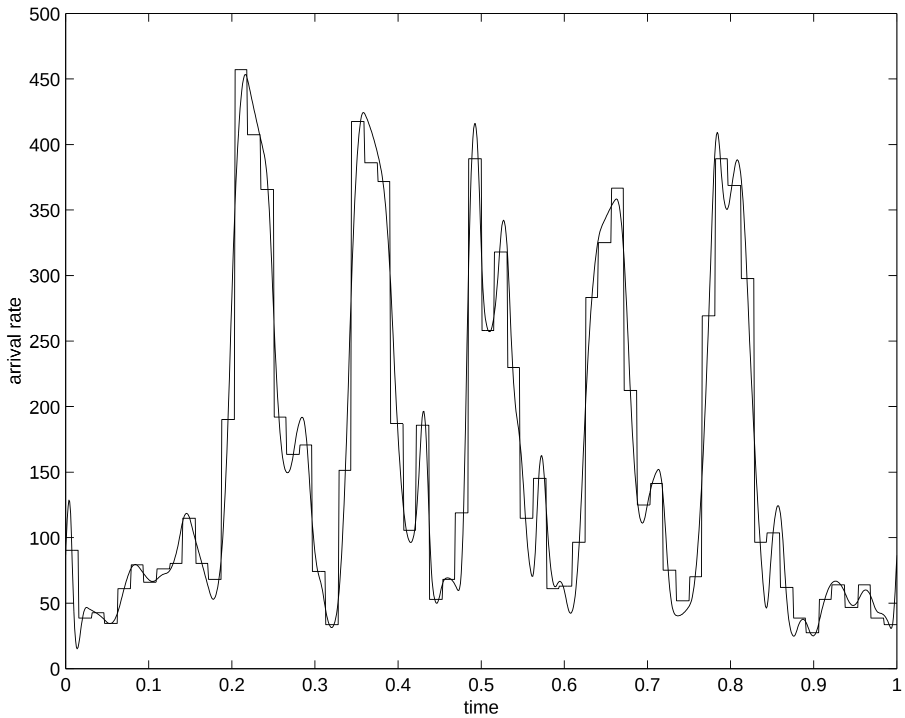

Figure 2.5: 336-knot approximation for the large e-mail dataset.

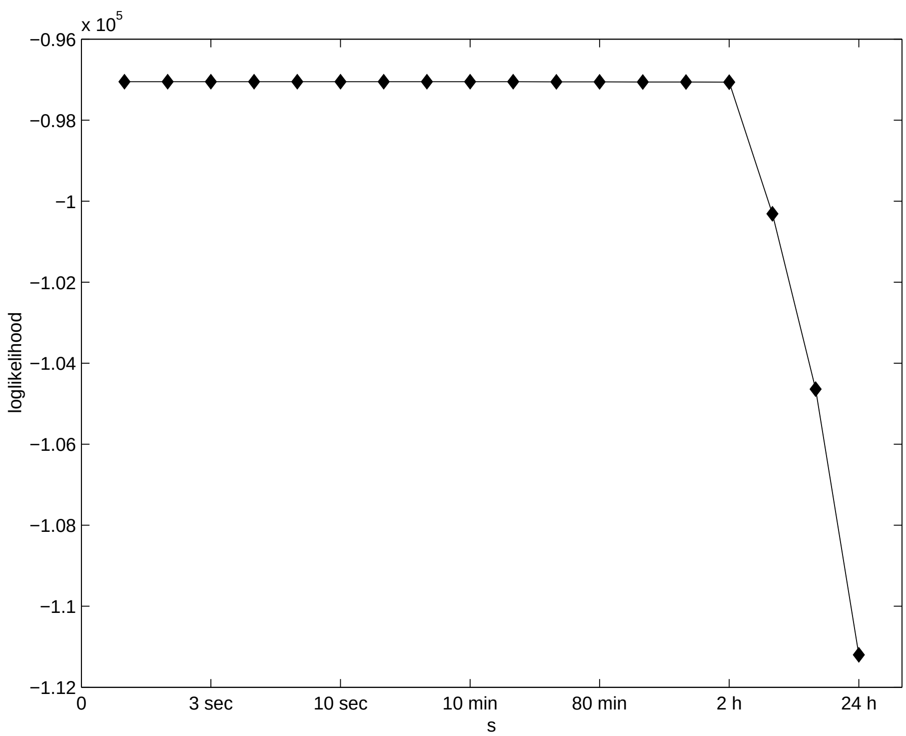

Figure 2.6: Optimal log likelihood as a function of the aggregation interval s.

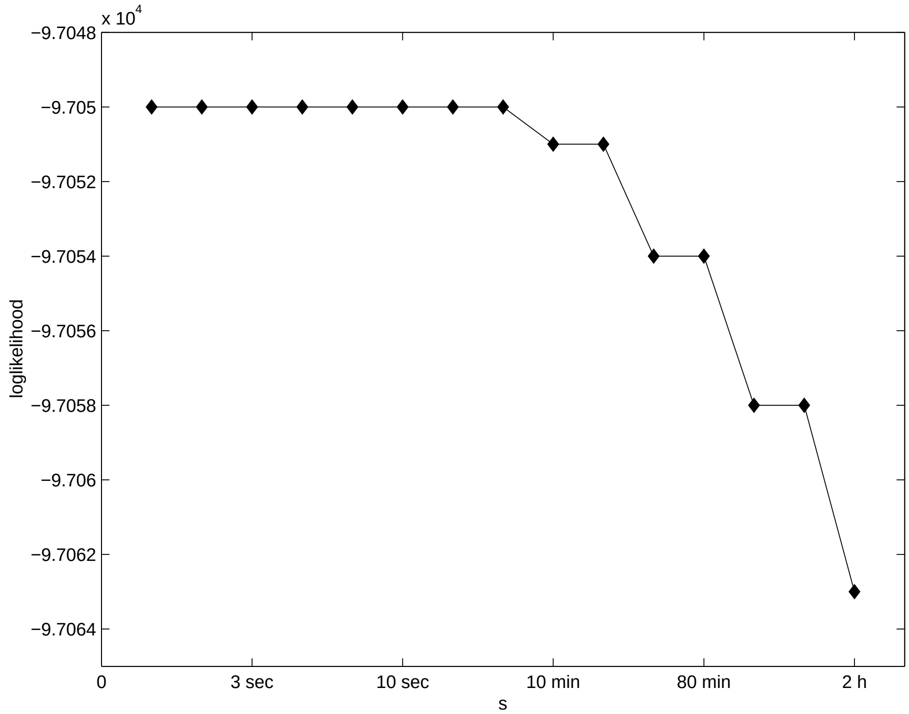

Figure 2.7: Magnified view of the critical portion of Figure 2.4.1.

Figure 2.4.1 shows the results. A sharp drop in the objective value occurs when the estimating spline reaches zero and the inequalities (2.30)-(2.32) in the nonnegativity constraints become binding. Figure 2.4.1 shows the section of the graph around this drop in more detail, while Figure 2.4.1 shows the estimating spline computed just after the drop. These results indicate that aggregating the arrival data with periods of up to 5 minutes does not significantly affect the fit.

2.4.2 Illustrating the importance of the nonnegativity constraints

In most of the experiments of §2.4.1, the spline nonnegativity constraints were not binding at the optimal solution, that is, the influence of the logarithmic terms in the objective function was sufficient to make the spline nonnegative without the help of the variables x_i, y_i, z_i, v_i, w_i, s_i and the associated constraints. In general, however, we cannot rely on such automatic satisfaction of nonnegativity. For smaller datasets or arrival rate functions that sometimes approach zero, nonnegativity constraints are essential. Further, as observed in §2.3, the estimation problem may even be unbounded and thus ill-posed if the nonnegativity constraints are omitted.

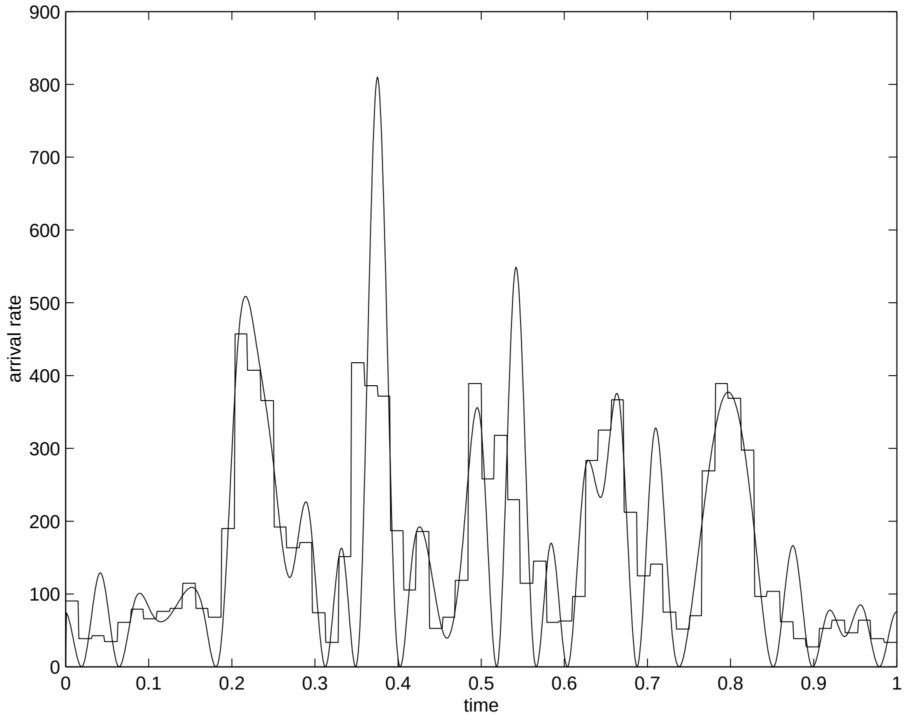

Figure 2.8: Approximating spline with excessively aggregated data.

To illustrate practical importance of the nonnegativity constraints, we consider a small dataset of 550 e-mail arrival times. Figure 2.4.2 shows the 42-knot spline estimate of the arrival rate, while Figure 2.4.2 shows the results when the nonnegativity constraints are omitted. Note from the piecewise-constant approximation that there is a significant time period (corresponding to Friday nights) when there are no arrivals in the dataset. Without nonnegativity constraints, the maximum-likelihood spline takes negative values during this period.

2.4.3 Numerical comparison to other approaches

We now present a comparison of our numerical results with those reported by [KB00] using nonnegative wavelets, and (in one instance) by [KDW98], using a parametric model and a least squares estimation method.

For testing purposes, we simulate arrival data using the same arrival rate functions $\lambda_1(\cdot)$, $\lambda_2(\cdot)$, and $\lambda_3(\cdot)$ as [KB00]. All three of these functions have the general form (2.2); for $\lambda_1(\cdot)$, the "trend" component $\sum p_i t^i$ of the function is constant, and the the trigonometric part consists of a single sine function. For $\lambda_2(\cdot)$, the trend $\sum p_i t^i$ is a quadratic polynomial, and the

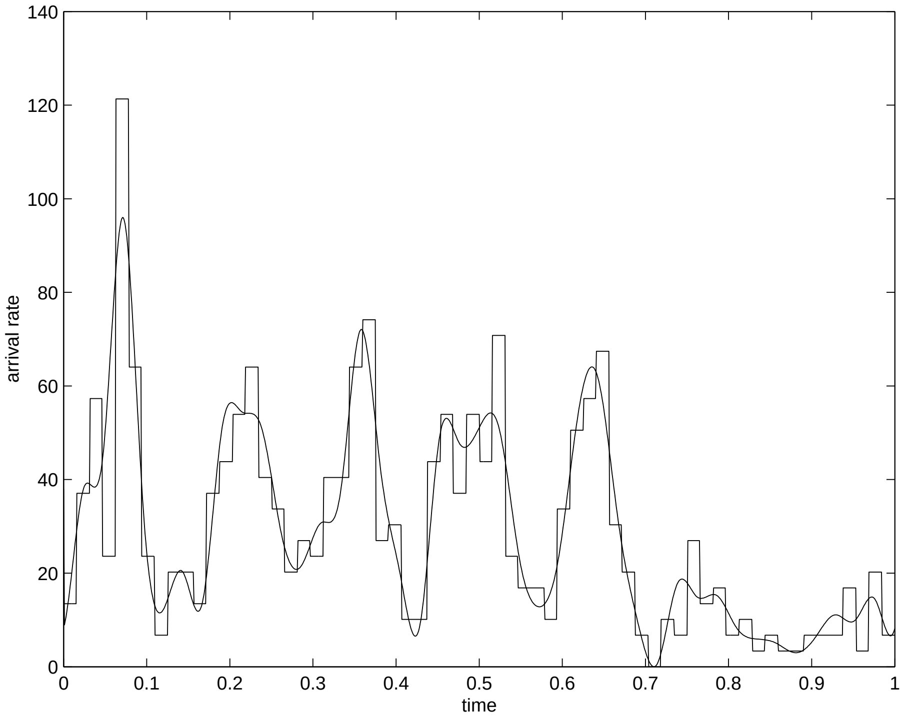

Figure 2.9: Arrival rate estimate for the small e-mail dataset.

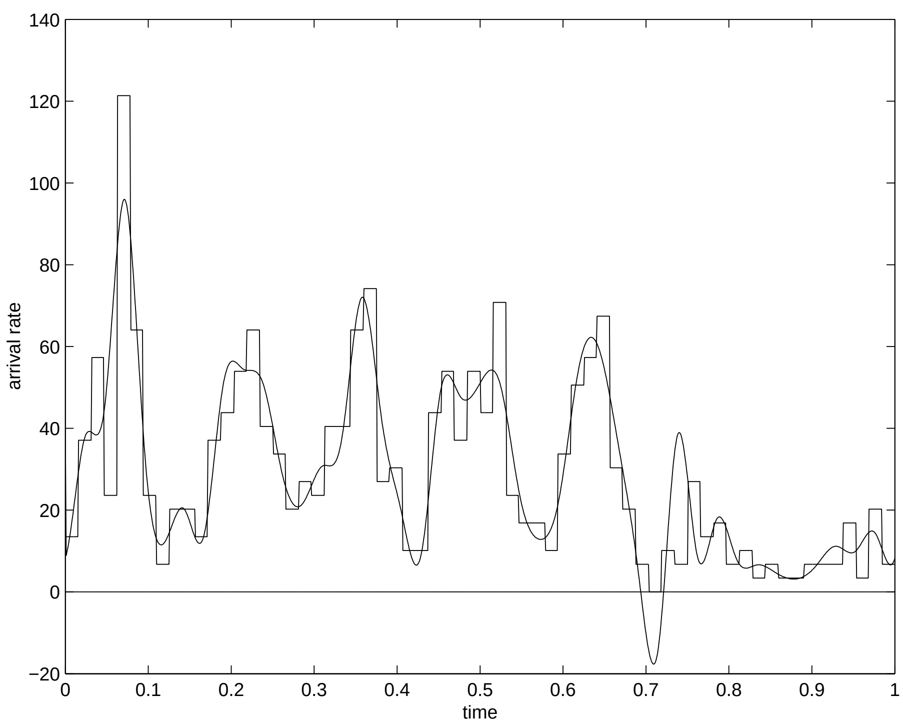

Figure 2.10: Estimate for the small e-mail dataset without spline nonnegativity constraints.

periodic part is again a single sine function. Finally, for $\lambda_3(\cdot)$, the polynomial $\sum p_i t^i$ is again constant, but the periodic part consists of two sine terms.

Since we are comparing our results to those of [KDW98], we have tried to closely follow their testing procedure. For each function $\lambda_i(\cdot)$, $i = 1, 2, 3$, we simulate arrivals in the interval [0,7] using the procedure described in [LS79]. We repeat this process twenty times for each of the three functions. Next, for each of the twenty simulations, we use our nonnegative spline method to construct an estimate $\hat{\lambda}_{i,j}$ of the arrival rate, where $i = 1, 2, 3$ and $j = 1, \ldots, 20$. For each $t \in [0, 7]$, we next form an empirical 90% tolerance interval $[\underline{\lambda}_i(t), \overline{\lambda}_i(t)]$ for $\lambda(t)$ by discarding the lowest and highest of these values, that is, letting $\underline{\lambda}_i(t) = \hat{\lambda}_{i,(2)}(t)$ and $\overline{\lambda}_i(t) = \hat{\lambda}_{i,(19)}(t)$, where $\hat{\lambda}_{i,(r)}(t)$ denotes the r^{th} smallest element of the set $\{\hat{\lambda}_{i,1}(t), \ldots, \hat{\lambda}_{i,20}(t)\}$.

Figures 2.4.3-2.4.3 display $\underline{\lambda}_i(t)$, $\lambda_i(t)$, and $\overline{\lambda}_i(t)$ as functions of t for each value of i. For $i = 1$, we make two sets of calculations: in Figure 2.4.3, we take into account the known periodicity of the function, and in Figure 2.4.3 we do not.

We also calculate the quality of these results with respect to the measures described in [KB00] and [KDW98]. In particular,

- δ is the average error from original arrival rate
- δ^* the maximum error from original arrival rate
- Δ is the average error from original mean value function
- Δ^* is the maximum error from original mean value function
- D is the average discrepancy between the estimated mean value function and the *empirical* arrival count function $N(t)$.
- D^* is the maximum discrepancy between the estimated mean value function and the *empirical* arrival count function $N(t)$.

In the above, the averages and maxima are taken over time, that is, over $t \in [0, 7]$; all the reported statistics are also averaged over the 20 replications for each i. For each statistic X, we also define a normalized version Q_X: for example, Q_δ is the average over t of $|\lambda_i(t) - \hat{\lambda}_{i,j}(t)|/\lambda_i(t)$; the other statistics are normalized similarly. This information is reported in Table 2.4.3. In the table, the notation $V[X]$ denotes the coefficient of variation of the

Table 2.1: Comparison of quality of estimated arrival rate functions computed by nonnegative splines, wavelet, and parametric least squares methods.

	$\lambda_1(t)$				$\lambda_2(t)$		$\lambda_3(t)$	
	Wavelet	Spline	Periodic Spline	Least Squares	Wavelet	Spline	Wavelet	Spline
δ	10.28	**9.36**	**4.40**	10.00	14.37	**12.37**	12.50	**12.08**
Q_δ	0.21	**0.19**	**0.09**	0.21	0.18	**0.16**	0.20	**0.19**
$V[\delta]$	0.14	0.07	0.27	0.65	0.15	0.08	0.15	0.11
δ^*	33.07	31.57	**11.59**	23.70	49.01	**47.98**	56.25	**50.63**
Q_{δ^*}	0.67	0.64	**0.24**	0.48	0.63	**0.61**	0.88	**0.80**
$V[\delta^*]$	0.29	0.17	0.32	0.60	0.33	0.03	0.24	0.15
Δ	8.49	**7.23**	**7.03**	12.4	7.12	**7.04**	10.81	**7.46**
Q_Δ	0.048	**0.040**	**0.040**	0.043	0.034	**0.030**	0.048	**0.030**
$V[\Delta]$	0.43	0.43	0.75	0.84	0.4	0.25	0.61	0.42
Δ^*	19.56	**17.11**	**14.10**	25.10	23.52	**17.59**	21.85	**17.20**
$\mathcal{Q}_{\Delta^*}$	(*) 0.49	**0.10**	**0.08**	0.09	(*) 0.54	**0.08**	0.49	**0.07**
$\mathcal{V}[\Delta^*]$	0.11	0.30	0.72	0.74	0.1	0.28	0.09	0.39
D	1.70	1.82	6.01	N/A	2.32	2.86	2.20	3.40
Q_D	0.01	0.01	0.06	N/A	0.01	0.01	0.01	0.01
D^*	5.80	6.96	14.78	N/A	8.97	9.65	7.63	10.77
Q_{D^*}	(*) 0.03	0.04	0.15	N/A	(*) 0.06	0.05	0.05	0.05

"N/A" indicates data not available from [KB00] or [KDW98]. We suspect the values marked (*) are incorrect and may have resulted from typographical errors in [KDW98].

statistic X over the 20-element sample. In our view, D and D^* are the least relevant statistics: they may not be good indicators of performance, since they measure the error from the sample path arrival count function, as opposed to the original mean value function. Such measures cannot detect overfitting.

In the case of the periodic estimator for $\lambda_1(\cdot)$, for example, our estimator stays quite close to the actual mean value function, but deviates from the empirical arrival count. Instances in which the spline estimators result in lower error statistics than reported in [KB00] and [KDW98] are marked in boldface; we have not boldfaced any of the coefficients of variation. Note that in nearly all cases, except for D and D^*, the spline estimators outperform those of [KB00] and [KDW98], even though they employ a parametric model of the same form used to generate the data. Our technique thus compares favorably with such prior work, especially considering that our approach is nonparametric.

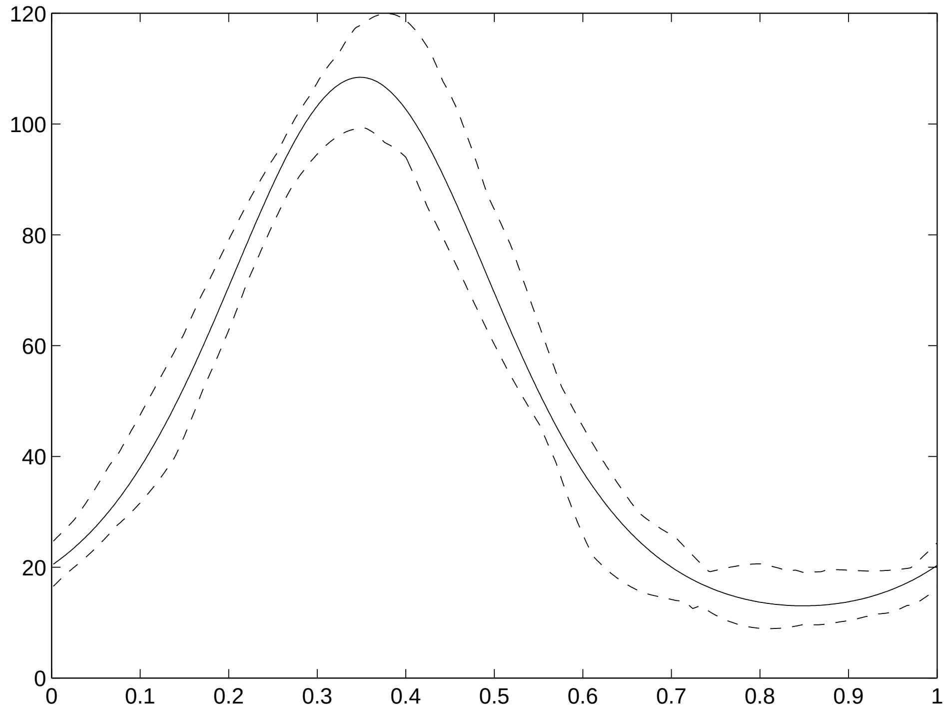

Figure 2.11: The function $\lambda_1(t)$ and computed empirical 90% tolerance interval, with explicit periodicity included in the spline model.

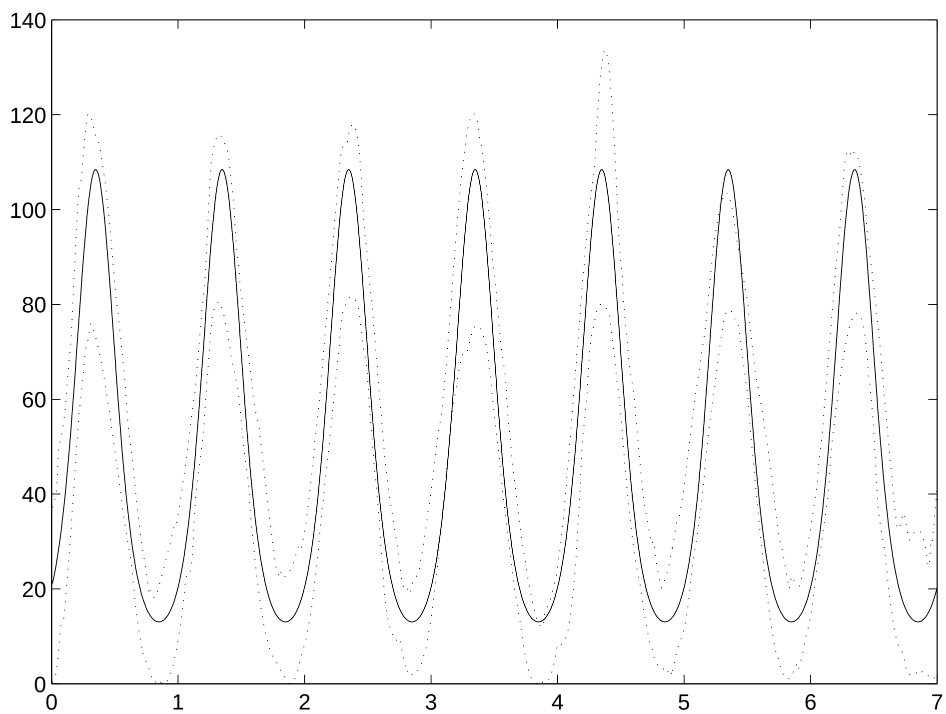

Figure 2.12: The function $\lambda_1(t)$ and computed empirical 90% tolerance interval, with periodicity *not* included in the spline model.

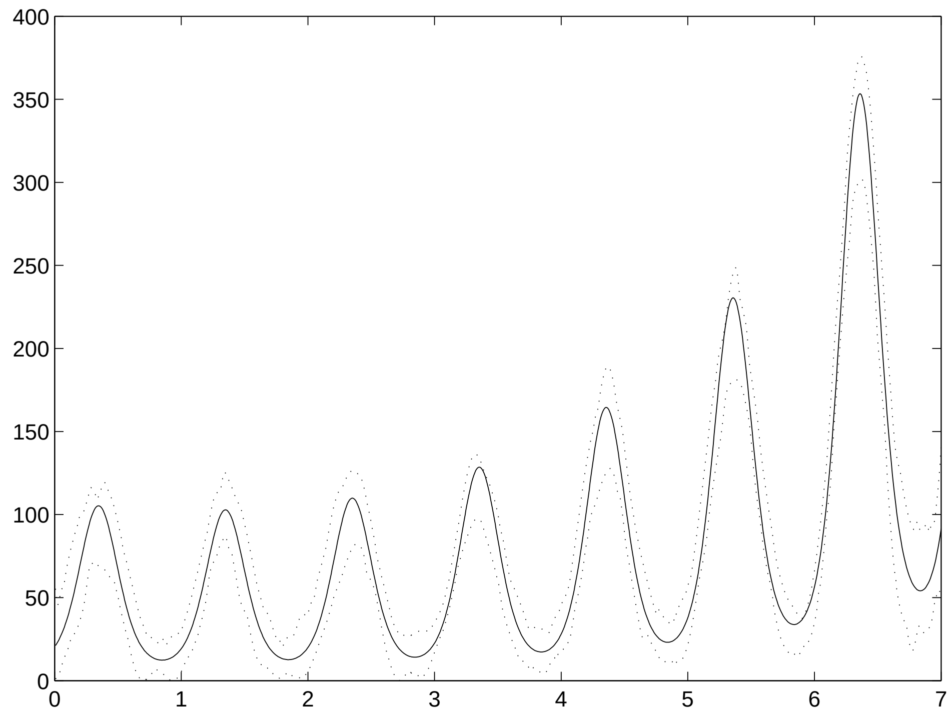

Figure 2.13: $\lambda_2(t)$ and its associated empirical 90% tolerance interval.

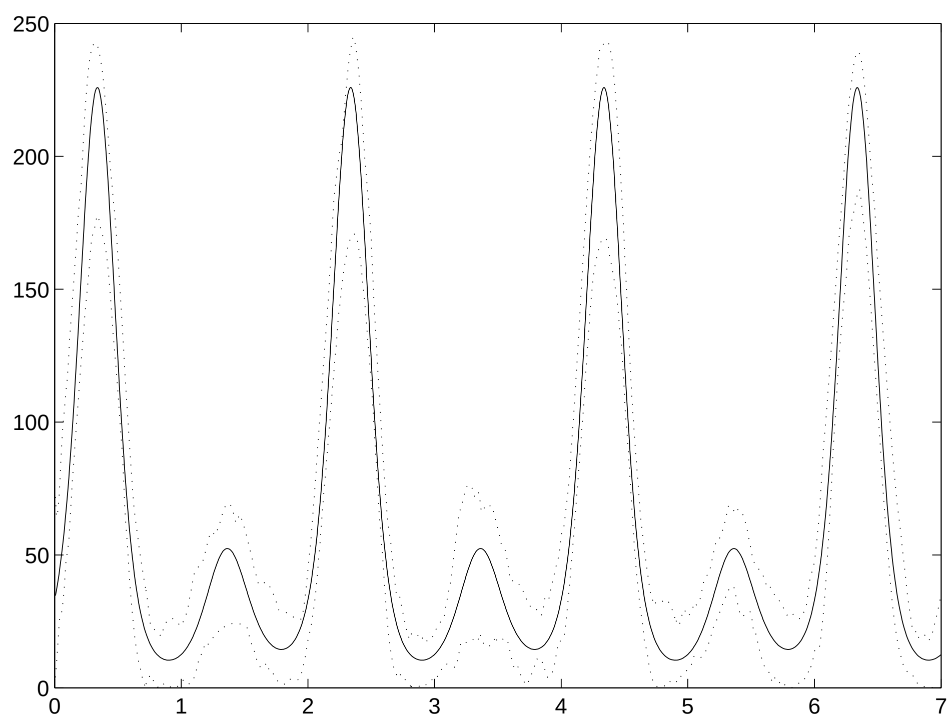

Figure 2.14: The function $\lambda_3(t)$ and its associated empirical 90% tolerance interval.

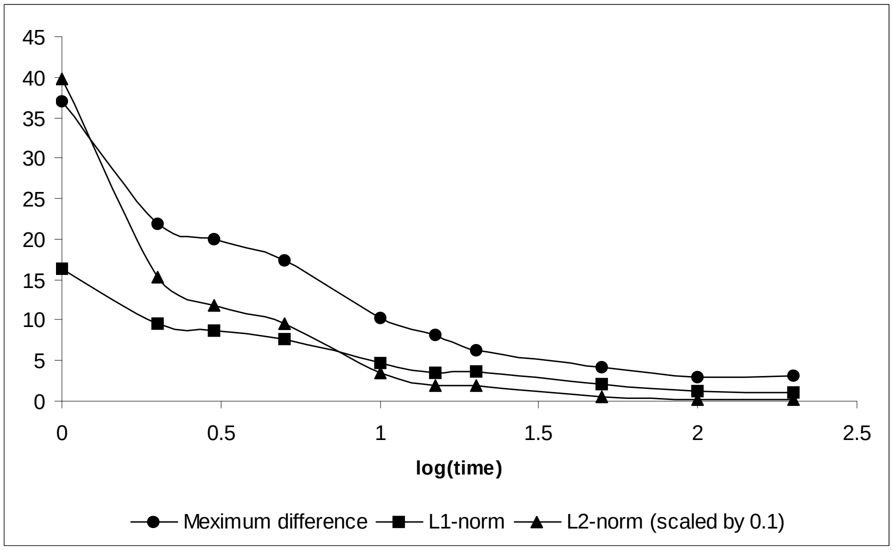

Figure 2.15: Convergence of estimates to $\lambda_4(\cdot)$.

2.4.4 Datasets generated from a known arrival rate function

We also used our own methodology to evaluate our approach's effectiveness on data generated from known arrival rate functions. In these experiments, we randomly generate datasets for time periods of different lengths, using periodic arrival rate functions $\lambda_4(\cdot)$ and $\lambda_5(\cdot)$ with a period of 1; $\lambda_4(\cdot)$ is a cubic spline with 6 knots, while $\lambda_5(t) = 100(\sin(2\pi t) + 1)$.

We estimated both of these functions with our method, using 6-knot splines. To measure the accuracy of the resulting estimate $\lambda^*(\cdot)$, we computed the $\mathcal{L}_1$-norm, $\mathcal{L}_2$-norm and the maximum of the absolute value of the difference $\lambda_i(\cdot) - \lambda^*(\cdot)$, that is, for $i = 4, 5$,

$$\mathcal{L}_{1,i} = \int_0^1 \left|\lambda_i(t) - \lambda^*(t)\right| dt \quad \mathcal{L}_{2,i} = \left(\int_0^1 \left(\lambda_i(t) - \lambda^*(t)\right)^2 dt\right)^{1/2} \quad \mathcal{L}_{\infty,i} = \max_{t\in[0,1]} \left|\lambda_i(t) - \lambda^*(t)\right|.$$

Figure 2.4.4 shows the results of these experiments for $\lambda_4(\cdot)$, while Figure 2.4.4 displays the results for $\lambda_5(\cdot)$. In both figures, the horizontal axis shows the logarithm of the length of time for which data were generated.

2.5 Conclusion and future work

We have demonstrated that it is computationally feasible to use recorded arrival data to compute nonnegative cubic spline estimates for the arrival rates of nonhomogeneous Poisson processes.

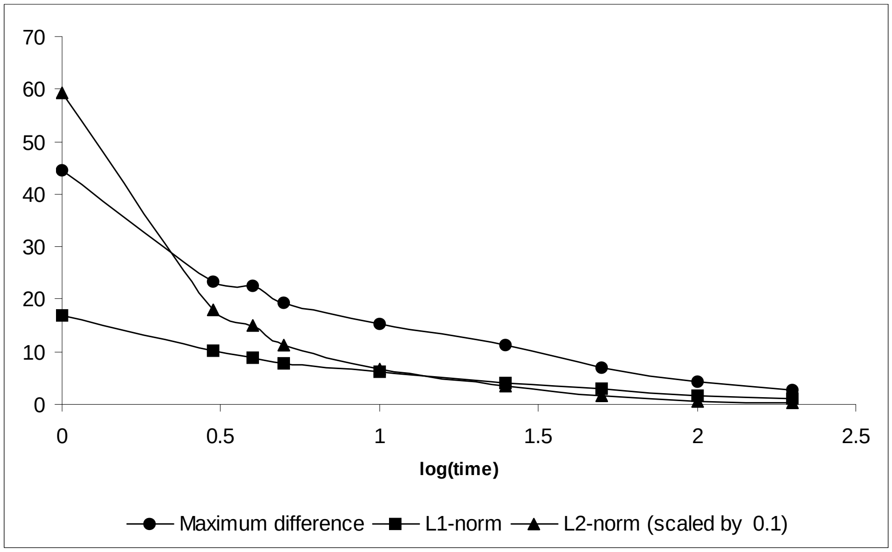

Figure 2.16: Convergence of estimates to $\lambda_5(\cdot)$.

The problem reduces to a convex nonlinear programming problem which is consistently solved to global optimality by the KNITRO software package, and also by a variety of other packages if proper scaling is employed, as suggested in §2.3.3

There are a number of possible extensions which bear further investigation. One extension is to model situations in which an arbitrary nonparametric cyclic intensity pattern modeled by splines is combined with a long term growth or decay trend analogous to the $\sum_i p_i t^i$ component of (2.2). An additive trend is easily combined with our model, but a more desirable multiplicative trend would require some modifications to our approach. It may be interesting to try to identify techniques that would guarantee a global likelihood maximum in such settings.

Another issue is determining whether there is any significant practical benefit in departing from the equidistant spline knot spacing we have used. Turning the spline knot positions $\{a_i\}$ into additional decision variables in the nonlinear programming formulation results in a nonlinear integer model that is most likely intractable. However, if our cross-validation procedure considered knot positions as well as just the number of knots m, there might be some benefit.

It is also interesting to consider whether our techniques may be extended to non-Poisson models. In particular, in order to validate our approximate solutions statistically, we applied the Kolmogorov-Smirnov goodness-of-fit test between the observed interarrival intervals and the nonhomogeneous Poisson distribution function implied by our estimates of $\lambda(\cdot)$; see [GE01] and [GES03] for a more detailed discussion. For our e-mail dataset, the results of these tests are poor for both the spline and traditional piecewise-constant estimates. The reason, we believe, is that the e-mail arrival process displays some autocorrelation or "burstiness". Specifically, some e-mail messages provoke other messages to appear shortly afterwards. Such autocorrelation occurs, for example, during "flame wars", or when there is an ongoing discussion. An interesting challenge is to devise a model which can take into account such effects in combination with natural cyclic time variations, and estimate these time variations in a manner similar to the techniques of this chapter.

We also note that our model can be extended to more general settings than polynomial splines. The basic characterization of nonnegative functions that we have employed applies in any functional space generated by a Tchebychev system $\{u_1(\cdot), \ldots, u_m(\cdot)\}$; see [KS66] for details. The log likelihood functional in this situation is still concave, and the cone of nonnegative functions remains convex. In many important special cases, this cone may be SD-representable, as in the case of the Tchebychev system generated by trigonometric polynomials of the form $\sum_i a_i \sin(it) + b_i \cos(it)$ (in this case, the trigonometric moment cones are characterized by positive semidefinite Töplitz matrices). For general Tchebychev systems, the cone of positive functions may not be SD-representable, but there are still efficient barrier functions from which practical interior point methods may be constructed; see [Fay02].

Finally, it is possible to consider extending our work to estimating multivariate probability distributions and multivariate nonhomogeneous Poisson processes. For instance, we may wish to estimate multivariate probability densities, or two- or three-dimensional arrival rates in spatial Poisson processes. Most of the theory described in this chapter carries over to such situations, with the notable exception of the nonnegativity requirement. It is in fact $\mathcal{NP}$-hard to decide whether a multivariate polynomial is nonnegative over a given region. However, it may be possible to estimate nonnegative functions by polynomials that are weighted sums of squares using the construction in [Nes00]. Thus, it could prove worthwhile to study thin-plate multivariate polynomial splines made up of such sum-of-squares polynomials.

Chapter 3

Bilinear Optimality Constraints for the Cone of Positive Polynomials

3.1 Introduction

In this chapter we examine the complementarity conditions for convex cones. In particular, we are interested in those cones where complementarity can be expressed using bilinear relations. Our main result is that the complementarity conditions for the cone of positive polynomials and its dual, the closure of the moment cone over the real line, *cannot* be represented by bilinear relations alone.

The cone of positive polynomials is a non-symmetric cone with many practical applications such as shape-constrained regression and the approximation of nonnegative functions (see for example [AENR08, PA08]).

It is well-known that positive polynomials over the real line are precisely those polynomials that can be written as the sum of squares of other polynomials. This property directly leads to the expression of the cone of positive polynomials as a linear image of the cone of positive semidefinite matrices, see for example [Nes00]. For instance, optimization over the cone of positive polynomials of degree $2n$ can be expressed as the dual of a semidefinite program over $n \times n$ Hankel matrices [DS97]. However, this approach may significantly increase the size of the problem and introduce degeneracy. This motivates us to look for solution methods and optimality conditions which directly apply to the cone of positive polynomials.

As a first step we wish to find as simple complementary slackness conditions as is possible for the positive polynomials and the moment cones. For instance, in linear programming complementary slackness conditions are given by $x_i s_i = 0$ where x_i are the primal variables and s_i are the dual slack variables. In semidefinite programming (SDP) the complementary slackness theorem is given by $XS + SX = 0$, where, again, X is the primal matrix variable and S is the dual slack matrix. Finally for second order cone programming (SOCP) we have $\langle \mathbf{x}, \mathbf{s} \rangle = 0$ and $s_0 x_i + s_i x_0 = 0$ (see the next section for more details). All of these relations are bilinear in the primal and dual slack variables. This property turns out to be essential in the design of primal-dual interior point algorithms. Furthermore, these bilinear forms make the machinery of certain algebraic structures available to help the understanding and improvement of such algorithms; this is especially true for SDP and SOCP.

According to a result of Güler, for every closed, pointed, convex cone $\mathcal{K}$ and its dual cone $\mathcal{K}^*$, the complementarity set $C(\mathcal{K})$, that is, the set of vector pairs $(\mathbf{x}, \mathbf{s}) \in \mathbb{R}^{2n}$, where $\mathbf{x} \in \mathcal{K}$, $\mathbf{s} \in \mathcal{K}^*$ and $\langle \mathbf{x}, \mathbf{s} \rangle = 0$, is an n-dimensional manifold. In many cases, this fact translates to a computationally tractable set of n equations $f_i(\mathbf{x}, \mathbf{s}) = 0$ $(i = 1, \ldots, n)$, which form the basis of complementary slackness theorems in optimization problems. Thus, it is an interesting endeavor to seek the simplest and most natural expressions for such relations. In fact, if it is at all possible to represent complementarity relations with bilinear forms, then that would be ideal, because potentially primal-dual interior point algorithms can be designed for such cones. Furthermore, bilinear relations induce algebras, and properties of these algebras may shed light on the properties of these cones and optimization problems over them [**?**].

In this chapter we develop some techniques for proving that for certain cones, bilinear relations are not sufficient to express complementary slackness. The method we apply relies on results allowing the parametrization of the boundaries of these cones based on the theory of Chebyshev systems [KS66].

The chapter is structured as follows: in Section 3.2 we present some fundamental concepts and results related to complementarity for proper cones, and introduce the notion of algebraic cones. In Section 3.3 we present a simple proof template for showing that cones are not algebraic. In the process we show a few simple cones that are not algebraic. We review necessary background information about the cone of positive polynomials $\mathcal{P}_{2n+1}$ and its dual, the closure of the moment cone $\mathcal{M}_{2n+1}$ in Section 3.4. Section 3.5 contains our main results concerning bilinear optimality constraints where we show that for the cone positive polynomials there are exactly four linearly independent bilinear complementarity relations.

3.1.1 Notation

For a polynomial represented by the vector of its coefficients $\mathbf{p} = (p_0, \ldots, p_n)$ the corresponding polynomial function is denoted by $p(t) = p_0 + p_1 t + p_2 t^2 + \cdots + p_n t^n$. For a real $t \in \mathbb{R}$ and nonnegative integer n, $\mathbf{c}_{n+1}(t)$ denotes the moment vector $(1, t, \ldots, t^n)^\top$.

Throughout the chapter we adopt the following convention: if for a range of indices the lower bound is greater than the upper bound, the range is considered to be empty.

The convex hull of a set $S \subset \mathbb{R}^n$ is denoted by $\mathrm{conv}(S)$, the closure of S is denoted by $\bar{S}$.

The linear space spanned by vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ is denoted by $\mathrm{span}(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k)$.

The inner product of vectors $\mathbf{x}$ and $\mathbf{s}$ is denoted by $\langle \mathbf{x}, \mathbf{s} \rangle = \mathbf{x}^\mathsf{T} \mathbf{s}$.

The parity of an integer m is denoted by $m \pmod 2 = \begin{cases} 0 \text{ if } m \text{ is even} \\ 1 \text{ if } m \text{ is odd} \end{cases}$.

For a matrix $A = (a_{ij})_{m \times n}$, $\mathrm{vec}(A) \stackrel{\text{def}}{=} (a_{11}, \ldots, a_{n1}, a_{12}, \ldots, a_{n2}, \ldots, a_{mn})^\top$. For two column vectors $\mathbf{u}$ and $\boldsymbol{v}$, their Kronecker product is defined to be $\mathbf{u} \otimes \boldsymbol{v} \stackrel{\text{def}}{=} \mathrm{vec}(\mathbf{u}\boldsymbol{v}^\top)$.

3.2 Algebraic Cones

Let $\mathcal{K}$ be a proper cone in $\mathbb{R}^n$ (that is, a closed, pointed, and convex cone with nonempty interior in $\mathbb{R}^n$), and let

$$\mathcal{K}^* = \{\mathbf{z} \mid \langle \mathbf{x}, \mathbf{z} \rangle \geq 0, \quad \forall \mathbf{x} \in \mathcal{K}\}$$

be its *dual cone*. A pair of vectors $(\mathbf{x}, \mathbf{s})$, $\mathbf{x} \in \mathcal{K}$, $\mathbf{s} \in \mathcal{K}^*$ is said to satisfy the *complementary slackness conditions* with respect to $\mathcal{K}$ if $\langle \mathbf{x}, \mathbf{s} \rangle = 0$. We are interested in the following set:

Definition 1 *Let $\mathcal{K}$ be a proper cone, and $\mathcal{K}^*$ its dual. Then the set*

$$C(\mathcal{K}) = \{(\mathbf{x}, \mathbf{s}) \mid \mathbf{x} \in \mathcal{K}, \mathbf{s} \in \mathcal{K}^*, \langle \mathbf{x}, \mathbf{s} \rangle = 0\}$$

is called the complementarity set *of* $\mathcal{K}$.

Since for every proper cone $(\mathcal{K}^*)^* = \mathcal{K}$, it is immediate from the definition that $C(\mathcal{K})$ and $C(\mathcal{K}^*)$ are congruent: one can be obtained from the other by exchanging the first and last n coordinates.

Let us now recall the following basic geometric fact, which will be necessary to prove an important result underlying the complementary slackness theorems for all convex optimization problems.

Proposition 1 *Let* $S \subseteq R^n$ *be a closed convex set and* $a \in \mathbb{R}^n$*. Then there is a unique point* $x = \Pi_S(a)$ *in* S*, called the* projection of a to S*, which is closest to* a*, i.e., there is a unique point* $x \in S$ *such that* $x = \operatorname{argmin}_{y \in S} \|a - y\|$*. Furthermore, if* S *is a closed convex cone, then* $\langle x, x - a \rangle = 0$*.*

The proof of the following theorem is due to O. Güler [Gül97].

Theorem 2 *For each proper cone* $\mathcal{K}$ *in* $\mathbb{R}^n$*,* $C(\mathcal{K})$ *is an* n*-dimensional manifold homeomorphic to* $\mathbb{R}^n$*.*

Proof: We need to show a continuous bijection between the complementarity set $C(\mathcal{K})$ of $\mathcal{K}$ and $\mathbb{R}^n$, whose inverse is also continuous.

Let $\varphi\colon \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^n$ be defined by $\varphi(a) = (x, s)$, where $x = \Pi_{\mathcal{K}}(a)$ and $s = x - a$. Clearly φ is continuous; we first show that $\varphi(a) \in C(\mathcal{K})$ for every a. By definition $\Pi_{\mathcal{K}}(a) \in \mathcal{K}$, and by the above proposition $\langle x, s \rangle = 0$. It remains to show that $s \in \mathcal{K}^*$.

For an arbitrary $u \in \mathcal{K} \setminus \{x\}$, define the convex combination $u_\alpha = \alpha u + (1 - \alpha)x$ where $0 \le \alpha \le 1$, and let $\zeta(\alpha) = \|a - u_\alpha\|^2$. Then ζ is a differentiable function on the interval $[0, 1]$, and $\min_{0 \le \alpha \le 1} \zeta(\alpha)$ is attained at $\alpha = 0$. Hence $\frac{d\zeta}{d\alpha}\big|_{\alpha=0} \ge 0$.

Now, using $\langle x, s \rangle = 0$, we have

$$\left.\frac{d\zeta}{d\alpha}\right|_{\alpha=0} = 2\langle s, u - x \rangle = 2\langle s, u \rangle \ge 0$$

for every $u \in \mathcal{K} \setminus \{x\}$. Note that the inequality $\langle s, u \rangle \ge 0$ also holds for $u = x$, implying $\langle s, u \rangle \ge 0$ for every $u \in \mathcal{K}$. Therefore $s \in \mathcal{K}^*$.

Consider now the continuous function $\bar{\varphi}\colon C(\mathcal{K}) \to \mathbb{R}^n$ defined by $\bar{\varphi}(x, s) = x - s$. To conclude the proof we show that $\bar{\varphi} \circ \varphi = \iota_{\mathbb{R}^n}$ and $\varphi \circ \bar{\varphi} = \iota_{C(\mathcal{K})}$, where ι_S denotes the identity function of the set S. The first one is easy:

$$(\bar{\varphi} \circ \varphi)(a) = \bar{\varphi}\left(\Pi_{\mathcal{K}}(a), \Pi_{\mathcal{K}}(a) - a\right) = a.$$

To show $\varphi \circ \bar{\varphi} = \iota_{C(\mathcal{K})}$, it suffices to prove that $\Pi_{\mathcal{K}}(x - s) = x$ for every $(x, s) \in C(\mathcal{K})$.

Suppose on the contrary that there is a point $u \in \mathcal{K}$ such that $\|a - u\| < \|a - x\|$, where $a = x - s$. Then, again using $\langle x, s \rangle = 0$,

$$0 > \langle a-u, a-u\rangle - \langle a-x, a-x\rangle = \langle x-s-u, x-s-u\rangle - \langle s,s\rangle = \|x-u\|^2 + 2\langle s,u\rangle,$$

in contradiction with $\langle s, u\rangle \geq 0$, which completes the proof. ■

To see the implications of this result for optimization problems over affine images or pre-images of proper cones, consider the following pair of dual *cone-LP* problems:

$$\begin{array}{llll} & \textbf{Primal} & & \textbf{Dual} \\ \inf & \langle c, x\rangle & \sup & \langle y, b\rangle \\ \text{s.t.} & Ax = b & \text{s.t.} & A^\top y + s = c \\ & x \in \mathcal{K} & & s \in \mathcal{K}^* \end{array} \tag{3.1}$$

It is easy to see that for any feasible solution x of the **Primal** problem and any feasible solution (y, s) of the **Dual** problem the quantities $\langle c, x\rangle - \langle y, b\rangle$ and $\langle x, s\rangle$ are equal and nonnegative. The *strong duality theorem* for cone-LP problems states the following: Under certain regularity conditions, if both the **Primal** and **Dual** problems are feasible, then inf and sup can be replaced by min and max. Moreover, the optimal objective values are equal, i.e., $\langle c, x^*\rangle - \langle y^*, b\rangle = \langle x^*, s^*\rangle = 0$. It follows that at the optimum we have $(x^*, s^*) \in C(\mathcal{K})$. Since $C(\mathcal{K}) \in \mathbb{R}^2 n$ is n-dimensional, it is often possible to obtain a square system of equations

$$\begin{aligned} &Ax = b \\ &A^\top y + s = c \\ &f_i(x, s) = 0 \quad \text{for } i = 1, \ldots, n, \end{aligned} \tag{3.2}$$

where $f_i(x, s) = 0$ are the complementarity equations. Many primal-dual algorithms for linear, second order and semidefinite programming problems, are based on strategies for solving this system of equations.

Let us examine some familiar examples.

Example 1 (Nonnegative orthant) When $\mathcal{K}$ is the nonnegative orthant, $\mathcal{K}^* = \mathcal{K}$. In this case if x and s contain only nonnegative components, and $\langle x, s\rangle = 0$, then we must have $x_i s_i = 0$ for $i = 1, \ldots, n$. This is the basis of the familiar complementary slackness theorem in linear programming. ■

Example 2 (Positive semidefinite cone) If $\mathcal{K}$ is the cone of real, symmetric positive semidefinite matrices, then $\mathcal{K}^* = \mathcal{K}$. If both X and S are real symmetric positive semidefinite matrices, and $\langle X, S\rangle = \sum_{ij} X_{ij}S_{ij} = 0$, then it is easy to show that the matrix product $XS = 0$, or equivalently $XS + SX = 0$. This is the basis of the complementary slackness theorem in semidefinite programming. ■

Example 3 (Second order cones) Let $\mathcal{K} \in \mathbb{R}^{n+1}$ be the cone defined by all vectors $\boldsymbol{x}$ such that $x_0 \geq \|\overline{\boldsymbol{x}}\|$, where $\boldsymbol{x} = (x_0, x_1, \dots, x_n)$, $\overline{\boldsymbol{x}} = (x_1, \dots, x_n)$, and $\|\cdot\|$ is the Euclidean norm. This cone is also self-dual. Now if $\boldsymbol{x}, \boldsymbol{s} \in \mathcal{K}$ and $\langle \boldsymbol{x}, \boldsymbol{s}\rangle = 0$, then from Cauchy-Schwarz-Bunyakovsky inequality it follows that $x_0 s_i + x_i s_0 = 0$ for $i = 1, \dots, n$. These relations along with $\langle \boldsymbol{x}, \boldsymbol{s}\rangle = 0$ are the basis of the complementary slackness theorem for the *second order cone programming* problem. ■

Example 4 (L_p cones) Generalizing the previous example, suppose instead the cone $\mathcal{K}_p$ consists of vectors $\boldsymbol{x}$ such that $x_0 \geq \|\overline{\boldsymbol{x}}\|_p$, where $\|\cdot\|_p$ is the L_p norm for some real number $p > 1$. Then it is known that the dual cone is $\mathcal{K}_q$ where $\frac{1}{p} + \frac{1}{q} = 1$. In this case one can deduce from Hölder's inequality that if $\boldsymbol{x} \in \mathcal{K}_p$ and $\boldsymbol{s} \in \mathcal{K}_q$ and $\langle \boldsymbol{x}, \boldsymbol{s}\rangle = 0$, then $s_0^q\, |x_i|^p - x_0^p\, |s_i|^q = 0$ for $i = 1, \dots, n$. ■

Example 5 (L_1 and L_∞ cones) A limiting case of the previous example is when $p = 1$ (and thus $q = \infty$). Here $\mathcal{K}_1$ consists of vectors $\boldsymbol{x}$ such that $x_0 \geq |x_1| + \cdots + |x_n|$, and $\mathcal{K}_\infty$ consists of vectors $\boldsymbol{s}$ where $s_0 \geq \max_i |s_i|$. In this case, if $\boldsymbol{x} \in \mathcal{K}_1$, $\boldsymbol{s} \in \mathcal{K}_\infty$ and $\langle \boldsymbol{x}, \boldsymbol{s}\rangle = 0$, then $x_i(s_0 - |s_i|) = 0$ for $i = 1, \dots, n$. ■

Recall that an *algebra* is a linear space with an additional multiplication operation: $\boldsymbol{x} \cdot \boldsymbol{y} = \boldsymbol{z}$ defined on its vectors. The main requirement is that the components of $\boldsymbol{z}$ be expressed as bilinear functions of $\boldsymbol{x}$, and $\boldsymbol{y}$; in algebraic terms this multiplication must satisfy the distributive law. Therefore, there are matrices Q_i such that $z_i = \boldsymbol{x}^\top Q_i \boldsymbol{y}$. If for a cone the complementarity relations can be exclusively expressed by bilinear forms, then, since these bilinear forms also define an algebra with multiplication, say "$\cdot$", the complementarity relations may be characterized by $\boldsymbol{x} \cdot \boldsymbol{s} = \boldsymbol{0}$. The machinery of this algebra may be useful in studying optimization problems over these cones. This motivates the following definitions.

Definition 2 *Let* $\mathcal{K} \in \mathbb{R}^n$ *be a proper cone. The* $n \times n$ *matrix* Q *is* a bilinear optimality condition for $\mathcal{K}$ *if every* $(\mathbf{x}, \mathbf{s}) \in C(K)$ *satisfies* $\mathbf{x}^\top Q\mathbf{s} = 0$*.*

Note that the set of all bilinear optimality conditions for $\mathcal{K}$, denoted by $\mathcal{Q}(\mathcal{K})$, is a linear subspace of $\mathbb{R}^{n\times n}$.

Definition 3 *A proper cone* $\mathcal{K} \subseteq \mathbb{R}^n$ *is called* algebraic *if there exist at least* n *linearly independent bilinear optimality conditions for* $\mathcal{K}$*.*

Remark 1 *An algebraic cone* $\mathcal{K} \subseteq \mathbb{R}^n$ *may have more than* n *bilinear optimality conditions, as the following example shows. Let* $\mathcal{K}$ *be the three-dimensional second order cone (see Example 3), and let*

$$Q_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, Q_2 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, Q_3 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, Q_4 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}.$$

Then every $(x, s) \in C(K)$ *satisfies* $x^\top Q_i s = 0$*,* $i = 1, 2, 3, 4$*. These four equations are linearly independent.*

Since $C(\mathcal{K})$ and $C(\mathcal{K}^*)$ are congruent, the cone $\mathcal{K}^*$ is algebraic if and only if $\mathcal{K}$ is.

From the examples above we observe that the cones in Examples 1, 2, and 3 are algebraic. Note that in Example 5, even though $\mathcal{K}_1$ and $\mathcal{K}_\infty$ are polyhedral, the complementarity relations are not completely bilinear due to the absolute values. In Theorem 3 we show that $\mathcal{K}_1$ and $\mathcal{K}_\infty$ do not have any non-trivial bilinear complementarity relations.

The largest class of cones known to be algebraic are the *symmetric cones*. These are cones that are self-dual and homogeneous (that is, for any two points in the interior of the cone, there is a linear automorphism of the cone mapping the first point to the second one [FK94]). The cones in Examples 1, 2, and 3 are all symmetric. In addition, the cones of positive semidefinite complex Hermitian and quaternion Hermitian matrices are also symmetric. The second order cone, and the cones of positive semidefinite symmetric, complex Hermitian and quaternion Hermitian matrices, along with an exceptional 27 dimensional cone, are essentially the only symmetric cones; any other symmetric cones can be decomposed into direct sums of these five classes of cones.

Symmetric cones are intimately related to *Euclidean Jordan algebras*, see [FK94] and [Koe99]. In such algebras the binary operation "$\circ$" is the abstraction of the operation $X \circ S = \frac{XS+SX}{2}$ in matrices. The properties of these algebras have played a major role in all aspects of optimization over such cones. In particular, design and analysis of interior point algorithms, duality, complementarity, and design of numerically efficient algorithms have been greatly simplified using the machinery of Jordan algebras. This is particularly true in the design of *primal-dual* interior point algorithms [Fay97], [AS00].

There is an easy way to manufacture algebraic cones from other algebraic cones.

Definition 4 *The proper cones $\mathcal{K}$ and $\mathcal{L}$ are* algebraically equivalent *if there is a nonsingular (one-to-one and onto) linear transformation A such that $A\mathcal{K} = \mathcal{L}$.*

If two cones are algebraically equivalent, then one is algebraic if and only if the other one is. In fact, in the next section we introduce the concept of *bilinearity rank* of a cone and prove that this rank is invariant among all algebraically equivalent cones.

In the next two sections we develop techniques to prove certain cones are not algebraic.

3.3 A simple approach for proving cones are *not* algebraic

Recall that $\mathcal{Q}(K)$ denotes the linear space of all bilinear optimality conditions for $\mathcal{K}$, and consider the linear space

$$L(\mathcal{K}) \stackrel{\text{def}}{=} \text{span}\{\mathbf{s}\mathbf{x}^\top \mid (\mathbf{x}, \mathbf{s}) \in C(\mathcal{K})\}.$$

Proposition 2 *For every proper cone $\mathcal{K}$ we have*

$$\dim(\mathcal{Q}(\mathcal{K})) = \text{co-dim}(L(\mathcal{K})).$$

Proof: Follows immediately from the identity $\mathbf{x}^\top Q \mathbf{s} = \langle \mathbf{s}\mathbf{x}^\top, Q^\top \rangle$. ■

Since by definition $X \in L(\mathcal{K})$ implies $\operatorname{trace} X = \langle X, I \rangle = 0$, the co-dimension of $L(\mathcal{K})$ as a subspace of $\mathbb{R}^{n\times n}$ is at least 1. Now if there are m linearly independent bilinear forms Q_i such that $\langle X, Q_i \rangle = 0$ for all $X \in L(\mathcal{K})$, then co-dim$(L(\mathcal{K})) \geq m$. Therefore, if we show $n^2 - k$ linearly independent matrices $X \in L(\mathcal{K})$, then this proves that there can be at most k bilinear forms in any characterization of $C(\mathcal{K})$. In particular, $\mathcal{K}$ is algebraic if and only if co-dim$(L(\mathcal{K})) \geq n$. Note that, as Remark 1 shows, it is possible that co-dim$(L(\mathcal{K})) > n$ for an algebraic cone $\mathcal{K}$.

Definition 5 *The quantity* $\dim(Q(\mathcal{K})) =$ co-dim$(L(\mathcal{K}))$ *is called the* bilinearity rank *of* K *and is denoted by* $\beta(\mathcal{K})$.

The manifolds $C(\mathcal{K})$ and $C(\mathcal{K}^*)$ are congruent for every proper cone $\mathcal{K}$, implying $\beta(\mathcal{K}) = \beta(\mathcal{K}^*)$. Furthermore, we have:

Lemma 1 *If* $\mathcal{K}$ *and* $\mathcal{L}$ *are algebraically equivalent proper cones then* $\beta(\mathcal{K}) = \beta(\mathcal{L})$.

Proof: Let A be a nonsingular linear transformation such that $A\mathcal{K} = \mathcal{L}$. Then the dual cone of $A\mathcal{K}$ is the cone $A^{-\top}\mathcal{K}^*$. Furthermore, Q_i $(i = 1, \ldots, m)$ define linearly independent bilinear complementarity conditions for $\mathcal{K}$ if and only if $A^{-\top}Q_iA^{\top}$ $(i = 1, \ldots, m)$ define linearly independent bilinear complementarity conditions for $A\mathcal{K}$. ∎

To derive our main results, we use the following simple fact.

Proposition 3 *If there are* k *pairs of vectors* $(\mathbf{x}_i, \mathbf{s}_i) \in C(\mathcal{K})$ *for* $i = 1, \ldots, k$*, such that the matrices* $\mathbf{s}_i\mathbf{x}_i^{\top}$ *are linearly independent, then* $\beta(\mathcal{K}) \leq n^2 - k$*. In particular, if* $k > n^2 - n$*, then* $\mathcal{K}$ *is not algebraic.*

These results lead to the following template for proving certain cones are not algebraic: Suppose $\mathcal{K}$ is a proper cone in $\mathbb{R}^n$.

Step 1 Select a finite set S of orthogonal pairs of vectors $(\mathbf{x}, \mathbf{s})$, where $\mathbf{x}$ is a boundary vector of $\mathcal{K}$ and $\mathbf{s}$ is a boundary vector of $\mathcal{K}^*$.

Step 2 Form the matrix T whose rows are $\mathbf{x} \otimes \mathbf{s} = \mathrm{vec}(\mathbf{s}\mathbf{x}^{\top})$, $(\mathbf{x}, \mathbf{s}) \in S$.

Step 3 If rank $T > n^2 - n$, then $\mathcal{K}$ is not algebraic. More generally, $\beta(\mathcal{K}) \leq n^2 - \mathrm{rank}\, T$.

To see how this template works let us show that the dual cones $\mathcal{K}_1, \mathcal{K}_\infty \subseteq \mathbb{R}^{n+1}$ from Example 5 are not algebraic for $n \geq 2$.

Theorem 3 $\beta(\mathcal{K}_1) = \beta(\mathcal{K}_\infty) = 1$.

Proof: As before, we assume that vectors are indexed from zero. We begin by introducing the following notation:

- $\boldsymbol{e}_i = (0,\dots,0,1,0\dots,0)^\top \in \mathbb{R}^{n+1}$, with the single nonzero element in the ith position ($i = 0,\dots,n$),
- $\mathbf{f} = (1,\dots,1) \in \mathbb{R}^{n+1}$,
- $\mathbf{f}_i = (1,\dots,1,-1,1,\dots,1) \in \mathbb{R}^{n+1}$, with all entries equal to 1 except in the ith position ($i = 0,\dots,n$),
- $\mathbf{f}_{ij} = (1,\dots,1,-1,1,\dots,1,-1,,1,\dots1) \in \mathbb{R}^{n+1}$ with all entries equal to 1 except in the ith and jth positions ($i,j = 0,\dots,n$).

The extreme rays of $\mathcal{K}_1$ are the $2n$ vectors $\boldsymbol{e}_0 \pm \boldsymbol{e}_i$ ($i = 1,\dots,n$), while the extreme rays of $\mathcal{K}_\infty$ are the 2^n vectors of the form $(1,\pm1,\pm1,\dots,\pm1)^\top$. Specifically, for every $i,j = 1,\dots,n$, the vectors $\mathbf{f}$, $\mathbf{f}_i$, and $\mathbf{f}_{ij}$ are among the extreme vectors of $\mathcal{K}_\infty$.

Let the set S (as described in Step 1 of the previous template) consist of the following orthogonal pairs $(\boldsymbol{x},\boldsymbol{s})$ from $C(\mathcal{K}_1)$:

- $(\boldsymbol{e}_0 + \boldsymbol{e}_i, \mathbf{f}_i), \quad i = 1,\dots,n,$
- $(\boldsymbol{e}_0 - \boldsymbol{e}_i, \mathbf{f}), \quad i = 1,\dots,n,$
- $(\boldsymbol{e}_0 + \boldsymbol{e}_i, \mathbf{f}_{ij}), \quad i,j = 1,\dots,n, \;\; i \neq j,$
- $(\boldsymbol{e}_0 - \boldsymbol{e}_i, \mathbf{f}_j), \quad i,j = 1,\dots,n, \;\; i \neq j,$

and let the matrix T be constructed as in Step 2. The following vectors can be obtained as linear combinations of the rows of T.

$$
\begin{aligned}
r_{0j} &= \boldsymbol{e}_0 \otimes \boldsymbol{e}_j = \frac{1}{4}\big((\boldsymbol{e}_0+\boldsymbol{e}_1)\otimes\mathbf{f}_1 - (\boldsymbol{e}_0+\boldsymbol{e}_1)\otimes\mathbf{f}_{1j} + (\boldsymbol{e}_0-\boldsymbol{e}_1)\otimes\mathbf{f} - (\boldsymbol{e}_0-\boldsymbol{e}_1)\otimes\mathbf{f}_j\big) \quad j = 1,\dots,n,\\
r_{ij} &= \boldsymbol{e}_i \otimes \boldsymbol{e}_j = \frac{1}{4}\big((\boldsymbol{e}_0+\boldsymbol{e}_i)\otimes\mathbf{f}_i - (\boldsymbol{e}_0+\boldsymbol{e}_i)\otimes\mathbf{f}_{ij} - (\boldsymbol{e}_0-\boldsymbol{e}_i)\otimes\mathbf{f} + (\boldsymbol{e}_0-\boldsymbol{e}_i)\otimes\mathbf{f}_j\big) \quad i,j = 1,\dots,n,\; i \neq\\
r_{ii} &= -\boldsymbol{e}_0\otimes\boldsymbol{e}_0 + \boldsymbol{e}_i\otimes\boldsymbol{e}_i = \sum_{\substack{1\leq j\leq n\\ j\neq i}} r_{0j} - \frac{1}{2}\big((\boldsymbol{e}_0+\boldsymbol{e}_i)\otimes\mathbf{f}_i + (\boldsymbol{e}_0-\boldsymbol{e}_i)\otimes\mathbf{f}\big), \quad i = 1,\dots,n\\
r_{i0} &= \boldsymbol{e}_i\otimes\boldsymbol{e}_0 = -(\boldsymbol{e}_0+\boldsymbol{e}_i)\otimes\mathbf{f} + \sum_{j=1}^{n} r_{0j} - \sum_{\substack{1\leq j\leq n\\ j\neq i}} r_{ij} - r_{ii}, \quad i = 1,\dots,n.
\end{aligned}
$$

Let $R \in \mathbb{R}^{[(n+1)^2-1]\times(n+1)^2}$ denote the matrix consisting of rows $\mathbf{r}_{01}, \mathbf{r}_{02}, \ldots, \mathbf{r}_{0n}, \mathbf{r}_{10}, \ldots, \mathbf{r}_{nn}$. Notice that by deleting the first column of R we obtain the identity matrix $I_{(n+1)^2-1}$. The rows of R were obtained as linear combinations of the rows of T, which in turn implies $\operatorname{rank} T \geq (n+1)^2 - 1$. In accordance with Step 3 of the previous template this completes the proof. ■

The template we used to prove Theorem 3 is a special case of the following, formally more general, framework:

Step 1 Select a set S of orthogonal pairs of vectors $(\mathbf{x}, \mathbf{s})$, where $\mathbf{x}$ is a boundary vector of $\mathcal{K}$ and $\mathbf{s}$ is a boundary vector of $\mathcal{K}^*$.

Step 2 Consider the set $\mathcal{T} = \{\mathbf{x} \otimes \mathbf{s} \mid (\mathbf{x}, \mathbf{s}) \in S\}$.

Step 3 If $\dim(\operatorname{span}(\mathcal{T})) > n^2 - n$, then $\mathcal{K}$ is not algebraic. More generally, $\beta(\mathcal{K}) \leq n^2 - \dim(\operatorname{span}(\mathcal{T}))$.

After presenting some necessary structural results in Section 3.4, we shall use these steps to prove our main results in Section 3.5.

3.4 Positive Polynomials and Moment Cones

Let us first introduce the cones of positive polynomials and moment cones:

Definition 6 *The* cone of positive polynomials *(also referred to as* cone of nonnegative polynomials*) of degree* $2n$

$$\mathcal{P}_{2n+1} \stackrel{\text{def}}{=} \left\{(p_0, \ldots, p_{2n}) \in \mathbb{R}^{2n+1} \;\mid\; p(t) = p_0 + p_1 t + p_2 t^2 + \cdots + p_{2n} t^{2n} \geq 0 \quad \forall t \in \mathbb{R}\right\}$$

consists of the coefficient vectors of nonnegative polynomials of degree $2n$*. The* moment cone *of dimension* $2n+1$ *is defined as*

$$\mathcal{M}_{2n+1} \stackrel{\text{def}}{=} \operatorname{conv}\left(\{\mathbf{c}_{2n+1}(t) \;\mid\; t \in \mathbb{R}\}\right), \textit{ where } \mathbf{c}_{2n+1}(t) \stackrel{\text{def}}{=} (1, t, t^2, \ldots, t^{2n})^\top.$$

Remark 2 *This is not the traditional definition of the moment cone. See [KS66] (Ch.VI) for the original definition and proof of its equivalence with the one given above.*

The cone of positive polynomials and the moment cone are closely related [KS66]:

Proposition 4 $\mathcal{P}^*_{2n+1} = \bar{\mathcal{M}}_{2n+1}$.

We will repeatedly use the following simple observation.

Proposition 5 *If* $\mathbf{p} \in \mathbb{R}^{n+1}$ *is the coefficient vector of a polynomial* p*, and* t *is real number, then* $p(t) = \langle \mathbf{p}, \mathbf{c}_{n+1}(t) \rangle$*. In particular,* $p(t) = 0$ *if and only if* $\langle \mathbf{p}, \mathbf{c}_{n+1}(t) \rangle = 0$*.*

In order to use the templates presented in Section 3.3 and prove that a cone $\mathcal{K}$ is not algebraic, it is useful to know the boundary or extreme rays of the cones $\mathcal{K}$ and $\mathcal{K}^*$. The extreme rays of $\mathcal{M}_{2n+1}$ are well known:

Proposition 6 ([KS66])

The extreme vectors of $\mathcal{M}_{2n+1}$ *are the vectors* $\alpha\mathbf{c}(t)$ *for every* $\alpha > 0$ *and* $t \in \mathbb{R}$*, and the vectors* $(0, \dots, 0, \alpha)^\top$ *for every* $\alpha \geq 0$*.*

Finally, in the subsequent sections we will also use the following observation:

Proposition 7 ([KS66])

Every root of a nonnegative polynomial in $\mathcal{P}_{2n+1}$ *is a multiple root with even multiplicity.*

3.5 Main Results

In this section we show our main result, namely that the cone of positive polynomials over the real line is not algebraic. Moreover, we give the exact bilinearity rank of this cones.

To prove our main results we need the following elementary fact from linear algebra.

Lemma 2 *Let* k *be a positive integer and let* $\mathcal{B} = \{b_1, \dots, b_k\}$ *be a set of linearly independent vectors in a real vector space. For a set* $\{m_1, \dots, m_k\} \subset \text{span}(\mathcal{B})$ *consider the coordinates* $\alpha_{i,j} \in \mathbb{R}$ $(i, j = 1, \dots, k)$ *uniquely defined by the representations* $m_i = \sum_{j=1}^{k} \alpha_{i,j} b_j$*. (We refer to this as the* $\mathcal{B}$-representation *of* m_i*.) If the conditions*

$$\begin{aligned} \alpha_{i,i} &\neq 0 \quad \textit{for all} \quad 1 \leq i \leq k, \\ \alpha_{i,j} &= 0 \quad \textit{for all} \quad 1 \leq i < j \leq k \end{aligned}$$

hold, then the set $\{m_1, \dots m_k\}$ *is also linearly independent.*

Proof: The claim follows immediately from the observation that the matrix $(\alpha_{i,j})_{k\times k}$ is lower triangular with a nonzero diagonal, and hence non-singular. ■

We are going to use the following, formally more general version of the above lemma:

Corollary 1 *Let* $\mathcal{B} \subset \mathbb{R}[x_1, \dots, x_n]$ *be a finite set of linearly independent polynomials and consider a set* $\mathcal{M} \subset \operatorname{span}(\mathcal{B})$ *with coordinates* $\alpha_{m,b}$ $(m \in \mathcal{M},\ b \in \mathcal{B})$ *defined by the representations* $m = \sum_{b\in\mathcal{B}} \alpha_{m,b} b$. *Assume that there exists an injection* $\varphi : \mathcal{B} \to \mathcal{M}$ *and a linear order* $\prec$ *on* $\varphi(\mathcal{B})$ *such that*

$$\alpha_{\varphi(b),b} \neq 0 \quad \textit{for all} \quad b \in \mathcal{B},$$

$$\alpha_{\varphi(b),d} = 0 \quad \textit{for all} \quad b, d \in \mathcal{B} \textit{ satisfying } \varphi(b) \prec \varphi(d).$$

Then $\dim(\operatorname{span}(\mathcal{M}(\mathbb{R}^n))) = |\mathcal{B}|$, *where* $\mathcal{M}(\mathbb{R}^n) \stackrel{\text{def}}{=} \{(m(\mathbf{x}))_{m\in\mathcal{M}} \,|\, \mathbf{x} \in \mathbb{R}^n\}$.

Proof: Let $k = |\mathcal{B}|$. It is well known that for a vector $P = (p_1, \dots, p_k) \in (\mathbb{R}[x_1, \dots, x_n])^k$ consisting of linearly independent polynomials we have $\dim(\operatorname{span}(P(\mathbb{R}^n))) = k$, therefore it suffices to find a k-element linearly independent subset of $\mathcal{M}$. As φ is injective, there exists an indexing $\mathcal{B} = \{b_1, \dots, b_k\}$ such that $\varphi(b_1) \prec \cdots \prec \varphi(b_k)$. Let $m_i = \varphi(b_i) \in \mathcal{M}$ (for all $i = 1, \dots, k$). It is easy to verify that the sets $\{b_1, \dots, b_k\}$ and $\{m_1, \dots, m_k\}$ satisfy the conditions of Lemma 2. Consequently the set $\{m_1, \dots, m_k\} \subset \mathcal{M}$ is linearly independent, which implies our claim. ■

3.5.1 Positive polynomials over the real line

Theorem 4 *The cone* $\mathcal{P}_{2n+1}$ *is not algebraic, unless* $n = 1$. *More specifically, for every* n, $\beta(\mathcal{P}_{2n+1}) \leq 4$.

The second claim immediately implies the first. Note that when $n = 1$, we do have an algebraic cone algebraically equivalent to the cone of 2×2 positive semidefinite matrices.

Proof: Consider the matrix valued functions $M\colon \mathbb{R}^n \mapsto \mathbb{R}^{(2n+1)\times(2n+1)}$ defined as

$$M(t_1, \dots, t_n) = \mathbf{c}\mathbf{p}^\top,$$

where $\mathbf{p} \in \mathcal{P}_{2n+1}$ is the coefficient vector of the polynomial $p(x) = \prod_{k=1}^n (x - t_k)^2$, and $\mathbf{c} = \mathbf{c}_{2n+1}(t_1) = (1, t_1, \dots, t_1^{2n})$ is the moment vector corresponding to the first root of $\mathbf{p}$. It is easy to verify that the entries of $M = (m_{i,j})_{i,j=0}^{2n}$ satisfy the polynomial equation

$$\sum_{j=0}^{2n} m_{i,j} x^j \equiv t_1^i \prod_{k=1}^{n} (x - t_k)^2. \tag{3.3}$$

The polynomial $p(x)$ is clearly nonnegative everywhere, and $\mathbf{c}$ is a moment vector, furthermore, by Proposition 5, $\langle \mathbf{p}, \mathbf{c} \rangle = 0$. Therefore, following the general template of Section 3.3 (with $\mathbf{p}$ and $\mathbf{c}$ playing the roles of $\mathbf{x}$ and $\mathbf{s}$, and $M(\mathbb{R}^n)$ playing the role of $\mathcal{T}$), the theorem follows if $\dim(\operatorname{span}(M(\mathbb{R}^n))) = (2n+1)^2 - 4$. We show this equality using the sufficient condition presented in Corollary 1, with the set $\{m_{i,j}\}$ playing the role of set $\mathcal{M}$.

Let us define the n-variate polynomials $\Pi(k, \ell)$ by

$$\Pi(k, \ell)(t_1, \ldots, t_n) \stackrel{\text{def}}{=} \sum_{\substack{0 \le \alpha_2, \ldots, \alpha_n \le 2 \\ \alpha_2 + \cdots + \alpha_n = \ell}} t_1^k \prod_{j=2}^{n} 2^{(\alpha_j \bmod 2)} t_j^{\alpha_j}, \tag{3.4}$$

whenever $0 \le k \le 2n+2$ and $0 \le \ell \le 2n-2$; for values of k and ℓ outside these ranges let us define $\Pi(k, \ell)$ to be the zero polynomial. Let $\mathcal{B}$ denote the set $\{\Pi(k, \ell) \mid 0 \le k \le 2n+2, 0 \le \ell \le 2n-2\}$. It follows from the definition that $|\mathcal{B}| = (2n+1)^2 - 4$, and that $\mathcal{B}$ is linearly independent, because no two polynomials share a common monomial. It remains to show that $\mathcal{M}$ is indeed a subset of $\operatorname{span}(\mathcal{B})$, and exhibit the injection φ and the linear order $\prec$ of Corollary 1.

The coefficient of $x^{2n-k-\ell}$ in the polynomial $\prod_{j=1}^{n} (x - t_j)^2$ is $\sum_{k=0}^{2} \Pi(k, \ell)$. From this observation it follows immediately that $\operatorname{span}(\mathcal{B})$ contains the entries of M; more specifically, for every $0 \le i, j \le 2n$,

$$m_{i,j} = \Pi(i, 2n-j) + \Pi(i+1, 2n-1-j) + \Pi(i+2, 2n-2-j). \tag{3.5}$$

We now introduce an injection $\varphi \colon \mathcal{B} \mapsto \mathcal{M}$ by defining its inverse (where it exists): let $m_{i,j}$ be the image of the polynomial

$$\varphi^{-1}(m_{i,j}) = q_{i,j} \stackrel{\text{def}}{=} \begin{cases} \Pi(i, 2n-j) & j \ge \max\{2, i\} \\ \Pi(i+2, 2n-2-j) & j \le \min\{i-1, 2n-2\} \\ \text{not defined} & \text{otherwise} \end{cases}. \tag{3.6}$$

In particular, we assign a polynomial to each entry $m_{i,j}$ of $\mathcal{M}$ except for $m_{0,0}$, $m_{0,1}$, $m_{1,1}$, and $m_{2n,2n-1}$, and we assign different polynomials to different entries of M, because if $q_{i_1,j_1} =$

q_{i_2,j_2} for some $(i_1, j_1) \neq (i_2, j_2)$ and $i_1 \leq j_1$, then $j_1 \geq i_1$, $i_2 - 1 \geq j_2$, $i_1 = i_2 + 2$, and $2n - j_1 = 2n - 2 - j_2$, a contradiction, as the sum of these inequalities reduces to $-1 \geq 0$. Consequently, each $\Pi(k, \ell)$ is equal to $q_{i,j}$ for precisely one pair (i, j), therefore φ is indeed an injection.

Equation (3.5) shows that the coefficient of $q_{i,j}$ in the $\mathcal{B}$-representation of $m_{i,j}$ is 1, so using the notation of Corollary 1, $\alpha_{\varphi(\Pi(k,\ell)),\Pi(k,\ell)} = 1$ for all $\Pi(k, \ell) \in \mathcal{B}$.

Let us define a linear order $\succ$ on $\varphi(\mathcal{B})$ in the following way: $m_{i_1,j_1} \succ m_{i_2,j_2}$ precisely when one of the following three conditions holds:

1. $i_1 - j_1 \geq 1 > i_2 - j_2$;

2. $i_1 - j_1 \geq 1$, $i_2 - j_2 \geq 1$, and either $i_1 > i_2$, or $i_1 = i_2$ but $j_1 < j_2$;

3. $i_1 - j_1 < 1$, $i_2 - j_2 < 1$, and either $j_1 < j_2$, or $j_1 = j_2$ but $i_1 > i_2$.

An easy case analysis using Equations (3.5) and (3.6) shows that if $m_{i_1,j_1} \succ m_{i_2,j_2}$, then the coefficient of q_{i_1,j_1} in the $\mathcal{B}$-representation of m_{i_2,j_2} is zero:

1. If $i_1 - j_1 \geq 1 > i_2 - j_2$, then Equations (3.5) and (3.6) show that the three terms of m_{i_1,j_1} have higher degree than those of m_{i_2,j_2}, so in particular $\Pi(i_1 + 2, 2n - 2 - j_1)$ does not appear in the $\mathcal{B}$-representation of m_{i_2,j_2}.

2. If both $i_1 - j_1, i_2 - j_2 \geq 1$, then $i_1 + 2 > i_2 + 2$ or $2n - 2 - j_1 > 2n - 2 - j_2$, and by Equation (3.5) $\Pi(i_1 + 2, 2n - 2 - j_1)$ does not appear in the $\mathcal{B}$-representation of m_{i_2,j_2}.

3. If both $i_1 - j_1, i_2 - j_2 \leq 0$, then $i_1 > i_2$ or $2n - j_1 > 2n - j_2$, and by Equation (3.5), $\Pi(i_1, 2n - j_1)$ does not appear in the $\mathcal{B}$-representation of m_{i_2,j_2}.

The injection $m_{i,j} \mapsto q_{i,j}$ and the linear order $\succ$ satisfy the conditions of Corollary 1, therefore, by Equation (3.6),

$$\dim(\text{span}(M(\mathbb{R}^n))) = \left|\mathcal{B}\right| = (2n + 1)^2 - 4,$$

which completes the proof. ■

3.5.2 Lower bounds

To simplify the proof of the validity of bilinear optimality conditions, we will use the following lemma.

Lemma 3 *The bilinear optimality condition* $\mathbf{x}^\top Q\mathbf{s} = 0$ *is satisfied by every* $(\mathbf{x}, \mathbf{s}) \in C(\mathcal{K})$ *if and only if it is satisfied by every* $(\mathbf{x}, \mathbf{s}) \in C(\mathcal{K})$ *such that* $\mathbf{x}$ *is an extreme vector of* $\mathcal{K}$ *and* $\mathbf{s}$ *is an extreme vector of* $\mathcal{K}^*$.

Proof: The *only if* direction is obvious. To show the converse implication, observe that every $\mathbf{x} \in \mathcal{K}$ and $\mathbf{s} \in \mathcal{K}^*$ can be expressed as a sum of finitely many extreme vectors of $\mathcal{K}$ and $\mathcal{K}^*$, respectively. Furthermore, if $\mathbf{x} = \sum_{i=1}^{k} \mathbf{x}_i$ and $\mathbf{s} = \sum_{j=1}^{\ell} \mathbf{s}_j$, then $\langle \mathbf{x}, \mathbf{s} \rangle = 0$ if and only if $\langle \mathbf{x}_i, \mathbf{s}_j \rangle = 0$ for every $1 \le i \le k$, $1 \le j \le \ell$. Therefore, if $\langle \mathbf{x}, \mathbf{s} \rangle = 0$, and the optimality condition is satisfied by every orthogonal pair of extreme vectors, then $\langle \mathbf{x}_i, \mathbf{s}_j \rangle = 0$ for every $1 \le i \le k$, $1 \le j \le \ell$, and

$$\mathbf{x}^\top Q\mathbf{s} = \left(\sum_{i=1}^{k} \mathbf{x}_i\right)^\top Q \left(\sum_{j=1}^{\ell} \mathbf{s}_j\right) = \sum_{i=1}^{k}\sum_{j=1}^{\ell} \mathbf{x}_i^\top Q\mathbf{s}_j = 0.$$

■

We are now ready to show that the upper bound on the number of linearly independent bilinear optimality conditions given in Theorem 4 is sharp.

Theorem 5 *For every integer* $n \ge 1$, $\beta(\mathcal{P}_{2n+1}) = 4$.

Proof: We have already proven $\beta(\mathcal{P}_{2n+1}) \le 4$. Now we prove that the following bilinear optimality conditions satisfied by every $(\mathbf{p}, \mathbf{c}) \in C(\mathcal{P}_{2n+1})$:

$$\sum_{i=0}^{2n} p_i c_i = 0, \tag{3.7a}$$

$$\sum_{i=1}^{2n} i p_i c_{i-1} = 0, \tag{3.7b}$$

$$\sum_{i=0}^{2n-1} (2n-i) p_i c_i = 0, \tag{3.7c}$$

$$\sum_{i=0}^{2n-1} (2n-i) p_i c_{i+1} = 0. \tag{3.7d}$$

It is easy to see that these conditions are indeed linearly independent. By Lemma 3 it is enough to show that the conditions are satisfied for pairs of vectors $(\mathbf{p},\mathbf{c}) \in C(\mathcal{P}_{2n+1})$ where $\mathbf{c}$ is an extreme vector of $\bar{\mathcal{M}}_{2n+1}$.

If $\mathbf{c} = c_{2n}\mathbf{e}_{2n} = (0,\dots,0,c_{2n})$ with some $c_{2n} > 0$ and $\langle \mathbf{p}, \mathbf{e}_{2n}\rangle = 0$, then (3.7a), (3.7b), and (3.7c) trivially hold, since all the terms on the left-hand sides these equations are zeros. Furthermore, the left-hand side of (3.7d) simplifies to $p_{2n-1}c_{2n}$, which must be zero, because otherwise $p_{2n-1} \neq 0$, $p_{2n} = 0$, and p would be a polynomial of odd degree, which cannot be nonnegative over the entire real line.

If $\mathbf{c}$ is an extreme vector of $\mathcal{M}_{2n+1}$, then, by Proposition 5, $\mathbf{c} = \mathbf{c}(t_0)$ for some $t_0 \in \mathbb{R}$, and $\mathbf{c}$ is orthogonal to $\mathbf{p}$ if and only if $p(t_0) = 0$. But this equation is equivalent to (3.7a), since

$$p(t_0) = \sum_{i=0}^{2n} p_i t_0^i = \sum_{i=0}^{2n} p_i c_i.$$

By Proposition 7, every root of p has even multiplicity, therefore $(\mathbf{p},\mathbf{c}) \in C(\mathcal{K})$ implies $p'(t_0) = 0$, which is equivalent to (3.7b), as

$$p'(t_0) = \sum_{i=1}^{2n} p_i i t_0^{i-1} = \sum_{i=1}^{2n} i p_i c_{i-1}.$$

Furthermore, if $p(t_0) = p'(t_0) = 0$, then $2np(t_0) - t_0 p'(t_0) = 0$, which translates to (3.7c), since

$$2np(t_0) - t_0 p'(t_0) = \sum_{i=0}^{2n} 2np_i t_0^i - \sum_{i=1}^{2n} p_i i t_0^i = \sum_{i=0}^{2n} 2np_i c_i - \sum_{i=1}^{2n} i p_i c_i = \sum_{i=0}^{2n} (2n-i) p_i c_i.$$

Finally, $p(t_0) = p'(t_0) = 0$ also implies $2nt_0 p(t_0) - t_0^2 p'(t_0) = 0$, which is equivalent to (3.7d):

$$2nt_0p(t_0) - t_0^2 p'(t_0) = \sum_{i=0}^{2n} 2np_i t_0^{i+1} - \sum_{i=1}^{2n} p_i i t_0^{i+1} = \sum_{i=0}^{2n} 2np_i c_{i+1} - \sum_{i=1}^{2n} i p_i c_{i+1} = \sum_{i=0}^{2n-1} (2n-i) p_i c_{i+1}.$$

■

3.6 Conclusion and recent results

Our main motivation for this research came from our work on solving statistical nonparametric estimation problems using polynomials and polynomial splines where the estimated functions themselves required to be nonnegative, [AENR08] and [PA08]. Our goal was to see if there is an easier way than formulating these problems as semidefinite programs. In particular are there efficient algorithms for cone-LP problems over positive polynomials? This questions led us to consider the simplest form of complementarity relations for positive polynomials, and we have found that bilinear complementarity relations alone are not sufficient.

As mentioned before, the results of this chapter originally appeared in [NRA05]. The revised framework presented here can also be used to prove similar results for positive polynomials over finite intervals and, using algebraic equivalence, several other cones of functions. Below we provide a quick overview of these recent developments, more details can be found in [NRAP].

Definition 7 *For real numbers* $a < b$*, the cone of* positive polynomials *(or* nonnegative polynomials*)* over the interval $[a, b]$ *of degree* n *is the cone*

$$\mathcal{P}_{n+1}^{[a,b]} \stackrel{\text{def}}{=} \Big\{(p_0, \ldots, p_n) \in \mathbb{R}^{n+1} \ \mid \ p(t) = p_0 + p_1 t + p_2 t^2 + \cdots + p_n t^n \geq 0 \quad \forall t \in [a, b]\Big\}.$$

The proof of the following result is similar to that of Theorems 4 and 5, although the computations are more complicated as case analysis becomes necessary.

Theorem 6 *The cone* $\mathcal{P}_{n+1}^{[a,b]}$ *is not algebraic. More specifically, for every* n*,* $\beta(\mathcal{P}_{n+1}^{[a,b]}) = 2$.

Using Lemma 1 we can extend our results to include cones which are algebraically equivalent to cones with a known bilinearity rank.

Theorem 7 *The cone* $\mathcal{P}_{n+1}^{[0,\infty]}$ *of polynomials non-negative over the half-line is algebraically equivalent to* $\mathcal{P}_{n+1}^{[0,1]}$*. Therefore* $\beta\big(\mathcal{P}^{[0.\infty]}\big) = 2$ *for every* $n \in \mathbb{N}$.

Theorem 8 *The cone of positive trigonometric polynomials of degree* $2n$

$$\mathcal{P}_{2n+1}^{\mathit{trig}} = \Big\{\mathbf{r} \in \mathbb{R}^{2n+1} \mid r_0 + \sum_{k=1}^{n}\big(r_{2k-1}\cos(kt) + r_{2k}\sin(kt)\big) \geq 0 \textit{ for all } t \in \mathbb{R}\Big\}$$

is algebraically equivalent to $\mathcal{P}_{2n+1}$*. Therefore,* $\beta\big(\mathcal{P}_{2n+1}^{\mathit{trig}}\big) = 4$ *for every* $n \in \mathbb{N}$.

The central question remaining open is whether there are algebraic cones other than symmetric cones and their algebraic equivalents?

Another direction is to investigate more sets of cones and estimate their bilinearity rank. For example one can examine all cones of positive functions over Chebyshev systems, and cones of functions of several variables which can be expressed as sums of squares of functions over a given finite set of functions.

Chapter 4

Optimization with Second Order Stochastic Dominance Constraints

4.1 Introduction

Our objective is to develop new approaches to stochastic optimization problems with a constraint in the form of the second order stochastic dominance relation. Such problems, introduced and analyzed in [DR03, DR04a], are new models of risk-averse optimization, in which risk aversion is expressed by the stochastic dominance constraint. Due to its specific structure, the constraint poses new theoretical and computational challenges.

The relation of *stochastic dominance* (introduced in statistics in [Leh55, MW47] and in economics in [HR69, QS62]) is defined as follows. Let X and Y be random variables on a probability space $(\Omega, \mathcal{F}, P)$ with distribution functions F_X and F_Y, respectively. We say that X *dominates* Y *in the first order* if $F_X(\eta) \le F_Y(\eta)$ for all $\eta \in \mathbb{R}$, and we denote this relation by $X \succeq_{(1)} Y$. An equivalent condition is that for every nondecreasing function $u(\cdot)$ one has

$$\mathbb{E}[u(X)] \ge \mathbb{E}[u(Y)], \tag{4.1}$$

provided the expected values above are finite.

For two integrable random variables X and Y, we say that X *dominates* Y *in the second order* if $\int_{-\infty}^{\eta} F_X(t)\,dt \le \int_{-\infty}^{\eta} F_Y(t)\,dt$ for all $\eta \in \mathbb{R}$, and we denote this relation by $X \succeq_{(2)} Y$. An equivalent condition is that for every concave nondecreasing function $u(\cdot)$ condition (4.1) holds true, provided that the expected values on both sides are finite.

We refer the readers to the monographs [MS02, SS94] for a modern view on the stochastic dominance relations and other comparison methods for random outcomes.

More generally, for an interval $I \subset \mathbb{R}$ let $X \succeq_{(2,I)} Y$ denote the relation

$$\int_{-\infty}^{\eta} F_X(t)\,dt \le \int_{-\infty}^{\eta} F_Y(t)\,dt, \quad \forall\, \eta \in I.$$

It is a relaxation of the second order dominance relation. If the interval I is compact, then this relaxation allows us to overcome technical difficulties in dealing with the second order dominance relation, as discussed in [DR03, DR04a]. If the interval I is reduced to one point, then the relation $X \succeq_{(2,I)} Y$ becomes the *integrated chance constraint* of [KH86].

An alternative representation of the second order dominance relation can be derived by using the *shortfall* of a random variable X from a target $\eta \in \mathbb{R}$, defined as $\max(0, \eta - X)$ (and written compactly $[\eta - X]_+$). By changing the order of integration one can easily verify that the expected value of the shortfall is given by the formula $\mathbb{E}([\eta - X]_+) = \int_{-\infty}^{\eta} F_X(t)\,dt$. Therefore we can rewrite the relation $X \succeq_{(2,I)} Y$ in the following form:

$$\mathbb{E}([\eta - X]_+) \le \mathbb{E}([\eta - Y]_+), \quad \forall\, \eta \in I. \tag{4.2}$$

Consider a stochastic model in which our decisions $z \in Z$ affect a random outcome $X = G(z)$. We assume that $z \in Z \subset \mathcal{Z}$, where $\mathcal{Z}$ is a Banach space and Z is a convex closed set. The mapping $G : \mathcal{Z} \to \mathcal{L}_1(\Omega, \mathcal{F}, P)$ is assumed to be continuous and concave in the sense that for P-almost all $\omega \in \Omega$ the function $z \mapsto [G(z)](\omega)$ is concave. Finally, let $f : \mathcal{Z} \to \mathbb{R}$ be a concave objective functional (for example, $f(z) = \mathbb{E}G(z)$). We are interested in the following problem

$$\begin{aligned} \underset{z}{\text{maximize}} \quad & f(z) \\ \text{subject to} \quad & G(z) \succeq_{(2,I)} Y, \\ & z \in Z. \end{aligned} \tag{4.3}$$

Here $Y \in \mathcal{L}_1(\Omega, \mathcal{F}, P)$ is a benchmark random outcome and I is an interval in $\mathbb{R}$.

As the second order dominance relation carries over to expectations of concave nondecreasing utility functions, no risk averse decision maker will prefer random outcome Y over random outcome $G(z)$ (if $I = \mathbb{R}$). Therefore, if the benchmark outcome Y represents an "acceptable" risk exposure, the risk exposure of $G(z)$ is even "more acceptable." Furthermore, suppose that the objective functional is monotone (consistent) with respect to the second order stochastic dominance relation, as defined in [OR99, OR01, OR02]: $G(z') \succeq_{(2)} G(z) \Rightarrow f(z') \ge f(z)$. For example, we may use $f(z) = \mathbb{E}[G(z)]$ or $f(z)$ being a negative of a coherent measure of risk. If the solution of problem (4.3) is unique, then no other feasible outcome $G(z')$ can strictly dominate the solution $G(z)$ (see [OR99, OR01, OR02]). The essence of the approach via stochastic dominance constraints is that the distribution of the outcome $G(z)$ is indirectly shaped by the distribution of the benchmark Y, which may also be an artificially constructed random variable. Reference [NR08] illustrates this modeling flexibility on an example of a portfolio problem.

The papers [DR03,DR04a] provide optimality and duality theory for problem (4.3) in which Lagrange multipliers associated with the dominance constraints are identified with concave nondecreasing utility functions. In [DR06a] an equivalent inverse form of the second order stochastic dominance constraint was analyzed and it was shown that it is equivalent to a continuum of Conditional (Average) Value at Risk constraints [RU02]. Moreover, Lagrange multipliers associated with the inverse form of stochastic dominance constraints were identified in [DR06a] with concave rank dependent utility functions of the dual utility theory [Yaa87]. In this way, model (4.3) is related to several classical models of risk averse decision making.

However, efficient solution of problem (4.3), even in the finite dimensional linear case, remains a challenge.

In what follows we focus on the stochastic dominance constraint $G(z) \succeq_{(2,I)} Y$ as the novel element in model (4.3), leaving aside considerations about possible objective functionals. We also remark that setting the problem in a Banach space $\mathcal{Z}$ does not lead to any significant technical difficulties, as compared to the finite dimensional case $\mathcal{Z} = \mathbb{R}^n$. Moreover, we hope to apply our formulation to multistage stochastic optimization problems, with $\mathcal{Z}$ representing the space of policies, which is usually modeled as a subspace of the space of integrable functions (see [RS03]).

Using (4.2) we obtain a more explicit formulation of (4.3):

$$\begin{aligned} \underset{z}{\text{maximize}} \quad & f(z) \\ \text{subject to} \quad & \mathbb{E}([\eta - G(z)]_+) \le \mathbb{E}([\eta - Y]_+), \quad \forall\, \eta \in I. \\ & z \in Z. \end{aligned} \tag{4.4}$$

When the functions $f(\cdot)$ and $G(\cdot)$ are affine and the set Z is a convex closed polyhedron, in §4.2 we develop a linear programming formulation of problem (4.4). But even in the finite dimensional case, this problem is difficult to solve, because its size grows quadratically with the number of the elementary events considered.

Another approach to (4.4) is the dual method of [DR04a]. It is a specialized nonsmooth optimization algorithm applied to the dual problem, in the space of concave nondecreasing functions playing the role of Lagrange multipliers associated with the dominance constraint. While efficient for some problems, especially portfolio problems of [DR06b], the dual method is rather complicated.

Our objective is to develop new efficient linear programming formulations, which exploit the specific structure of the stochastic dominance constraint in cut generation schemes. This results in a significant increase of the size of computationally tractable problems, as well as in a speedup in the solution of smaller instances. Furthermore, for problems with first order stochastic dominance constraints $\mathsf{G}(z) \succeq_{(1)} \mathsf{Y}$, which are typically much more difficult, due to the potential nonconvexity of the feasible region, model (4.3) serves as a powerful convex relaxation (see [DR04b,NRR06,NR08]). Thus, the speedup also benefits some advanced iterative methods of [NR08] for problems with first order constraints.

In §4.2 and §4.3 we present a primal cutting plane method based on formulation (4.4). In §4.4 we develop a new version of the duality theory for an extended reformulation of problem (4.4). In §4.5 we show how a reduction of the number of variables in the dual problem can be achieved by employing Strassen's theorem about the existence of measures on product spaces with given marginals. This leads to a dual cutting plane method of §4.7. Finally, in §4.8 we present numerical results, along with performance comparisons of the various methods, for portfolio optimization problems based on real data.

We remark that work to improve cutting plane methods for problems with second order stochastic constraints is still ongoing; see [DR] for some recent advances.

4.2 A linear representation of the second order stochastic dominance constraint

In order to solve (4.3) it is necessary to represent the SSD relation $\succeq_{(2,\mathrm{I})}$ in a tractable form. The usual approach to achieve this is to introduce *shortfall functions*. In the finite dimensional case they correspond to slack variables, but in the infinite dimensional case we need to introduce an appropriate space of the shortfall functions.

Denote by ℓ the Lebesgue measure on I, and let $\mathcal{B}$ be the σ-algebra of Borel subsets of I. We denote the Banach space of continuous functions on I by $\mathcal{C}(I)$. Let $\mathcal{S}$ be the vector space of all real-valued measurable functions s on $(I \times \Omega, \mathcal{B} \times \mathcal{F}, P \times \ell)$ satisfying the following conditions:

(i) for every $\eta \in I$ the function $s(\eta, \cdot)$ is an element of $\mathcal{L}_1(\Omega, \mathcal{F}, P)$;

(ii) for P-almost all $\omega \in \Omega$ the function $s(\cdot, \omega)$ is an element of $\mathcal{C}(I)$;

(iii) the function $\omega \mapsto \max_{\eta \in I} |s(\eta, \omega)|$ is an element of $\mathcal{L}_1(\Omega, \mathcal{F}, P)$.

Owing to the Lebesgue theorem, the function $w(\eta) = \int_\Omega s(\eta, \omega)\, dP$ is an element of $\mathcal{C}(I)$. It can be verified directly from the definition that $\mathcal{S}$ a Banach space with the norm

$$\|s\| = \int_\Omega \max_{\eta \in I} |s(\eta, \omega)|\, dP.$$

Immediately from (4.2) we obtain the following observation.

Lemma 4 *Assume that* $X, Y \in \mathcal{L}_1(\Omega, \mathcal{F}, P)$. *Then* $X \succeq_{(2,I)} Y$ *if and only if there exists a nonnegative function* $s \in \mathcal{S}$ *such that*

$$\begin{aligned} & s(\eta, \omega) \geq \eta - X(\omega), \quad \forall\, \eta \in I, \quad \forall\, \omega \in \Omega, \\ & \int_\Omega s(\eta, \omega)\, dP \leq \int_\Omega [\eta - Y(\omega)]_+\, dP, \quad \forall\, \eta \in I. \end{aligned}$$

Let us introduce the notation $v(\eta) = \mathbb{E}([\eta - Y]_+) = \int_\Omega [\eta - Y(\omega)]_+\, dP$ for the shortfalls of the benchmark variable. Applying Lemma 4, we can formulate another optimization problem which is equivalent to (4.3):

$$\begin{aligned} \underset{z,s}{\text{maximize}} \quad & f(z) \\ \text{subject to} \quad & \int_\Omega s(\eta, \omega)\, dP \leq v(\eta), \quad \forall\, \eta \in I, \\ & [G(z)](\omega) + s(\eta, \omega) \geq \eta, \quad \forall \eta \in I, \quad \forall\, \omega \in \Omega, \\ & s \geq 0, \\ & z \in Z, \ s \in \mathcal{S}. \end{aligned} \tag{4.5}$$

If the functional $f(\cdot)$ and the mapping $G(\cdot)$ are affine and the set Z is polyhedral, then problem (4.5) becomes a linear programming problem in Banach spaces. When the distribution of the benchmark outcome is discrete, one can restrict the range of η in (4.5) to the realizations of the benchmark Y.

A potential drawback of the above approach is the introduction of the auxiliary variables $s(\cdot,\cdot)$ indexed by the set $I \times \Omega$. As we shall see in §4.8, even in the finite dimensional case, when $\mathcal{Z} = \mathbb{R}^n$ and the probability space Ω is finite, formulation (4.5) may be impractical to solve.

We now present an alternative representation which does not require additional variables $s(\cdot,\cdot)$. It is an extension of the representation developed in [KHVDV06] for integrated chance constraints.

Theorem 9 *Assume that* $X, Y \in \mathcal{L}_1(\Omega, \mathcal{F}, P)$. *Then* $X \succeq_{(2,I)} Y$ *if and only if for all* $\eta \in I$ *and all events* $A \in \mathcal{F}$,

$$\mathbb{E}\left((\eta - X)\,\mathbf{1}_A\right) \leq \nu(\eta).$$

Proof: For every event $A \in \mathcal{F}$

$$\mathbb{E}\left((\eta - X)\,\mathbf{1}_A\right) \leq \mathbb{E}([\eta - X]_+).$$

This inequality becomes an equation if $A = \{X < \eta\}$, and thus

$$\max_{A \in \mathcal{F}} \mathbb{E}\left((\eta - X)\,\mathbf{1}_A\right) = \mathbb{E}([\eta - X]_+).$$

The theorem immediately follows from (4.2). ■

Using this result we obtain another equivalent formulation of the optimization problem (4.2):

$$\begin{aligned} \underset{z}{\text{maximize}} \quad & f(z) \\ \text{subject to} \quad & \int_A \big(\eta - G(z)\big)\, dP \leq \nu(\eta), \quad \forall\, \eta \in I, \quad \forall\, A \in \mathcal{F}, \\ & z \in Z. \end{aligned} \tag{4.6}$$

Although the auxiliary variables are no longer present, we have introduced an infinite family of constraints indexed by the set $I \times \mathcal{F}$. However, we shall show that the new family of constraints can be efficiently dealt with by a cut generation method.

4.3 A primal cutting plane method

In this section we assume that $I = \mathbb{R}$, the probability space Ω is finite, with elementary events $\omega_1, \ldots, \omega_N$ and corresponding probabilities $p_1, \ldots, p_N$. The realizations of the benchmark outcome Y are denoted by $y_1, \ldots, y_D$ and the corresponding benchmark shortfalls are $v_j = \mathbb{E}([y_j - Y]_+)$. We also write $G_i(z)$ for $[G(z)](\omega_i)$.

It is known from [DR03, DR04a] that in the case of a discrete benchmark, the second order dominance condition $G(z) \succeq_{(2)} Y$ is equivalent to finitely many inequalities:

$$\mathbb{E}([y_j - G(z)]_+) \le v_j, \quad j = 1, \ldots, D. \tag{4.7}$$

We can thus rewrite problem (4.6) as follows

$$\begin{aligned} &\underset{z}{\text{maximize}} \quad f(z) \\ &\text{subject to} \quad \sum_{i \in A} p_i\big(y_j - G_i(z)\big) \le v_j, \quad \forall\, j = 1 \ldots, D, \quad \forall\, A \subset \{1, \ldots, N\}, \\ &\qquad\qquad\quad z \in Z. \end{aligned} \tag{4.8}$$

The last formulation allows for construction of a cutting plane method. At iteration k we have a collection of subsets (events) $A_1, \ldots, A_{k-1}$ of $\{1, \ldots, N\}$. We solve a relaxation of (4.8):

$$\begin{aligned} &\underset{z}{\text{maximize}} \quad f(z) \\ &\text{subject to} \quad \sum_{i \in A_m} p_i\big(y_j - G_i(z)\big) \le v_j, \quad j = 1 \ldots, D, \quad m = 1, \ldots, k-1, \\ &\qquad\qquad\quad z \in Z. \end{aligned} \tag{4.9}$$

If the solution z^k of this problem (which is assumed to exist) satisfies all constraints (4.7), then we stop. Otherwise we find j^* for which (4.7) is violated and we define

$$A_k = \{1 \le i \le N : y_{j^*} > G_i(z^k)\}.$$

The iteration index k is increased by one, and we solve (4.9) again.

Since (4.7) is violated,

$$\sum_{i \in A_k} p_i\big(y_{j^*} - G_i(z)\big) > v_{j^*},$$

and thus A_k is different than A_m, $m = 1, \ldots, k-1$, used in problem (4.9). As the possible number of sets that can be added is finite, the method must stop at an optimal solution of (4.8).

Examples in §4.8 suggest that in practice a small number of sets A_k need to be generated in order to find the optimal solution.

4.4 Lagrangian duality

In this section we derive duality relations for the extended formulation (4.5). Our derivation uses ideas and techniques developed in [DR04a]. The main difference is that we develop duality relations for the formulation (4.5) involving explicit shortfall variables, in contrast to the duality theory of [DR04a], where we focused on the dominance constraint in the nonsmooth formulation (4.4).

The difficulty with formulation (4.5) is that no Slater condition can be formulated for the inequality constraint on the shortfall variables:

$$[G(z)](\omega) + s(\eta, \omega) \geq \eta, \quad \forall\, \eta \in I, \quad \forall\, \omega \in \Omega,$$

because the nonnegative cone in the space $\mathcal{S}$ has no interior. Because of that, we cannot simply apply general duality schemes from [Roc74] or [KH86]. We need to exploit the special structure of problem (4.5).

At first, we introduce several relevant topological vector spaces. We denote by $\mathbf{rca}(I)$ the space of finite signed measures on I and by $\mathcal{L}_\infty(\Omega, \mathcal{F}, P)$ the space of essentially bounded measurable real functions on $(\Omega, \mathcal{F}, P)$. Let $\mathcal{M}$ denote the vector space of signed measures on $(I \times \Omega, \mathcal{B} \times \mathcal{F})$, such that for every measure $\lambda \in \mathcal{M}$ the *marginal measures* λ_I and λ_Ω, defined by the equations

$$\lambda_I(B) = \lambda(B \times \Omega), \quad B \in \mathcal{B},$$
$$\lambda_\Omega(A) = \lambda(I \times A), \quad A \in \mathcal{F},$$

satisfy the conditions:

$$\lambda_I \in \mathbf{rca}(I), \qquad \frac{d\lambda_\Omega}{dP} \in \mathcal{L}_\infty(\Omega, \mathcal{F}, P). \tag{4.10}$$

Here we implicitly assume that λ_Ω is absolutely continuous with respect to P.

Theorem 10 *The space $\mathcal{M}$ is the topological dual space to the space $\mathcal{S}$, that is, ℓ is a continuous linear functional on $\mathcal{S}$ if and only if there exists $\lambda \in \mathcal{M}$ such that for all $s \in \mathcal{S}$*

$$\ell(s) = \iint\limits_{I\times\Omega} s(\eta,\omega)\, d\lambda. \tag{4.11}$$

Proof: Fix any $\lambda \in \mathcal{M}$, $\lambda \geq 0$, and consider the linear functional (4.11). Its value can be bounded as follows:

$$\begin{aligned} |\ell(s)| &\leq \iint\limits_{I\times\Omega} \max_{\eta\in I} |s(\eta,\omega)|\, d\lambda = \int\limits_{\Omega} \max_{\eta\in I} |s(\eta,\omega)|\, d\lambda_\Omega \\ &= \int\limits_{\Omega} \max_{\eta\in I} |s(\eta,\omega)| \frac{d\lambda_\Omega}{dP}(\omega)\, dP \leq \left\| \frac{d\lambda_\Omega}{dP} \right\|_{\mathcal{L}_\infty} \|s\|. \end{aligned}$$

For a general signed measure $\lambda \in \mathcal{M}$ we use its Jordan decomposition into a difference of two nonnegative measures: $\lambda = \lambda^+ - \lambda^-$, and we define Ω^+ and Ω^- to be the support sets of λ^+ and λ^-, respectively. Using the last displayed inequality we obtain the estimate

$$\begin{aligned} |\ell(s)| &\leq \Bigg| \iint\limits_{I\times\Omega^+} s(\eta,\omega)\, d\lambda_+ \Bigg| + \Bigg| \iint\limits_{I\times\Omega^-} s(\eta,\omega)\, d\lambda_- \Bigg| \\ &\leq \left\| \frac{d\lambda^+_\Omega}{dP} \right\|_{\mathcal{L}_\infty} \int\limits_{\Omega^+} \max_{\eta\in I} |s(\eta,\omega)|\; dP + \left\| \frac{d\lambda^-_\Omega}{dP} \right\|_{\mathcal{L}_\infty} \int\limits_{\Omega^-} \max_{\eta\in I} |s(\eta,\omega)|\; dP \\ &\leq \max\left(\left\| \frac{d\lambda^+_\Omega}{dP} \right\|_{\mathcal{L}_\infty}, \left\| \frac{d\lambda^-_\Omega}{dP} \right\|_{\mathcal{L}_\infty} \right) \int\limits_{\Omega} \max_{\eta\in I} |s(\eta,\omega)|\; dP = \left\| \frac{d\lambda_\Omega}{dP} \right\|_{\mathcal{L}_\infty} \|s\|, \end{aligned}$$

and we conclude that the linear functional (4.11) is continuous. Thus $\mathcal{S}^* \supset \mathcal{M}$.

To prove the converse inclusion, consider the linear subspace of S:

$$\mathcal{S}_0 = \big\{ s \in \mathcal{S} : s = \varphi\xi,\ \varphi \in \mathcal{C}(I),\ \xi \in \mathcal{L}_1(\Omega, \mathcal{F}, P) \big\}.$$

Let $\ell \in \mathcal{S}_0^*$. Fix $A \in \mathcal{F}$ and consider the functional $\varphi \mapsto \ell(\varphi \mathbf{1}_A)$. It is continuous on $\mathcal{C}(I)$. By the Riesz representation theorem, there exists a measure $\mu^\ell_A \in \mathbf{rca}(I)$ such that

$$\ell(\varphi \mathbf{1}_A) = \int\limits_I \varphi(\eta)\, d\mu^\ell_A, \quad \forall\, \varphi \in \mathcal{C}(I).$$

Define the measure λ^ℓ on $(I \times \Omega, \mathcal{B} \times \mathcal{F})$ by the formula:

$$\lambda^\ell(B \times A) = \mu^\ell_A(B), \quad \forall\, B \in \mathcal{B}, \quad \forall\, A \in \mathcal{F}.$$

Then

$$\ell(\varphi \mathbf{1}_A) = \iint\limits_{I\times\Omega} \varphi(\eta)\mathbf{1}_A(\omega)\, d\lambda^\ell.$$

It follows that for every $s = \varphi\xi$ such that $\varphi \in \mathcal{C}(I)$ and ξ is a step function, i.e., $\xi = \sum_{k=1}^{K} \alpha_k \mathbf{1}_{A_k}$ with some $\alpha_k \in \mathbb{R}$ and $A_k \in \mathcal{F}$, $k = 1, \dots, K$, the functional ℓ has form (4.11) with $\lambda = \lambda^\ell$. As the step functions are dense in $\mathcal{L}_1(\Omega, \mathcal{F}, P)$, for every $\xi \in \mathcal{L}_1(\Omega, \mathcal{F}, P)$ we can find a sequence of step functions $\xi^j \to \xi$, $j \to \infty$. Since ℓ is continuous, we obtain

$$\ell(\varphi\xi) = \lim_{j\to\infty} \ell(\varphi\xi^j) = \lim_{j\to\infty} \iint\limits_{I\times\Omega} \varphi(\eta)\xi_j(\omega)\, d\lambda^\ell = \iint\limits_{I\times\Omega} \varphi(\eta)\xi(\omega)\, d\lambda^\ell,$$

and thus the functional ℓ has form (4.11) on $\mathcal{S}_0$. Moreover, the marginal measure λ_I^ℓ satisfies the first part of condition (4.10):

$$\lambda_I^\ell(B) = \lambda^\ell(B \times \Omega) = \mu_\Omega^\ell \in \mathbf{rca}(I).$$

Consider now functions $s(\eta, \omega) = \xi(\omega)$ with $\xi \in \mathcal{L}_1(\Omega, \mathcal{F}, P)$. As ℓ is continuous, the functional $\xi \mapsto \ell(\mathbf{1}\xi)$ must be continuous on $\mathcal{L}_1(\Omega, \mathcal{F}, P)$. Since

$$\ell(\mathbf{1}\xi) = \iint\limits_{I\times\Omega} \xi(\omega)\, d\lambda^\ell = \int\limits_\Omega \xi(\omega)\, d\lambda_\Omega^\ell,$$

it is necessary that also the second part of (4.10) is satisfied by λ_Ω^ℓ. Thus $\mathcal{S}^* \subset \mathcal{S}_0^* \subset \mathcal{M}$. ■

We can now formulate the *Lagrangian* $L : \mathcal{Z} \times \mathcal{S} \times \mathcal{M} \times \mathbf{rca}(I) \to \mathbb{R}$ of the optimization problem (4.5) as follows:

$$L(z, s, \lambda, \mu) = f(z) - \iint\limits_{I\times\Omega} \Big(\eta - [G(z)](\omega) - s(\eta, \omega)\Big)\, d\lambda + \int\limits_I \Big(\nu(\eta) - \int\limits_\Omega s(\eta, \omega)\, d\mathbb{P}\Big)\, d\mu.$$

The corresponding *Lagrangian dual function* $L_D : \mathcal{M} \times \mathbf{rca}(I) \to \mathbb{R}$ is given by

$$\begin{aligned} L_D(\lambda, \mu) &= \sup_{s\ge 0,\ z\in Z} L(z, s, \lambda, \mu) \\ &= \sup_{s\ge 0,\ z\in Z} \Big\{ f(z) + \iint\limits_{I\times\Omega} s(\eta, \omega)\, d(\lambda - \mu \times P) - \int\limits_I \eta\, d\lambda_I + \int\limits_\Omega [G(z)](\omega)\, d\lambda_\Omega + \int\limits_I \nu(\omega)\, d\mu \Big\}. \end{aligned}$$

By examining the second term of this expression we obtain

$$L_D(\lambda, \mu) = \begin{cases} -\int\limits_I \eta\, d\lambda_I + \int\limits_I \nu(\eta)\, d\mu + \sup\limits_{z\in Z} \Big\{ f(z) + \int\limits_\Omega [G(z)](\omega)\, d\lambda_\Omega \Big\} & \text{if } \lambda \le \mu \times P, \\ +\infty & \text{otherwise.} \end{cases} \tag{4.12}$$

This leads to the following dual problem:

$$\begin{aligned} &\underset{\lambda,\mu}{\text{minimize}} \; -\int_I \eta \, d\lambda_I + \int_I \nu(\omega)\, d\mu + \sup_{z\in Z}\Big\{ f(z) + \int_\Omega [G(z)](\omega)\, d\lambda_\Omega \Big\} \\ &\text{subject to } \lambda \le \mu \times P, \\ &\qquad\qquad \lambda \in \mathcal{M}^+, \quad \mu \in \mathbf{rca}^+(I). \end{aligned} \tag{4.13}$$

Here we use $\mathcal{M}^+$ and $\mathbf{rca}^+(I)$ to denote the sets of all nonnegative measures in $\mathcal{M}$ and $\mathbf{rca}(I)$, respectively.

Theorem 11 (Weak Duality) *Let c_* and c_L denote the optimum values of the original problem* (4.5) *and the Lagrangian dual* (4.13)*, respectively. Then $c_* \le c_L$.*

Proof: Let (z,s) be feasible for (4.5). Then for every $(\lambda,\mu) \in \mathcal{M}^+ \times \mathbf{rca}^+(I)$ we have the inequalities

$$L_D(\lambda,\mu) \ge L(z,s,\lambda,\mu) \ge f(z).$$

Taking the infimum of the left hand side with respect to (λ,μ) and the supremum of the right hand side with respect to feasible (z,s) we obtain the assertion. ■

In order to prove the strong duality relation we need a constraint qualification condition, introduced in [DR03, DR04a].

Definition 8 Problem (4.5) satisfies the *uniform dominance condition* if there exists $\tilde{z} \in Z$ such that

$$\max_{\eta\in I}\Big\{\mathbb{E}\big[(\eta - G(\tilde{z}))_+\big] - \nu(\eta)\Big\} < 0.$$

Theorem 12 (Strong Duality) *Assume that problem* (4.5) *satisfies the uniform dominance condition and that it has an optimal solution. Then the dual problem* (4.13) *has an optimal solution and $c_* = c_L$.*

Proof: Due to Theorem 11, it is sufficient to find $(\hat{\lambda},\hat{\mu}) \in \mathcal{M}^+ \times \mathbf{rca}^+(I)$ such that $L_D(\hat{\lambda},\hat{\mu}) = c_*$.

Let $(\hat{z},\hat{s})$ be an optimal solution of (4.5). Consider the equivalent problem formulation

(4.4). Following [DR03] we can rewrite in the abstract form:

$$\begin{aligned}
&\underset{z}{\text{maximize}}\ f(z)\\
&\text{subject to } \Gamma(z) \in K,\\
&\qquad\qquad z \in Z,
\end{aligned}$$

where $\Gamma : \mathcal{Z} \to \mathcal{C}(I)$ is a continuous operator defined as

$$[\Gamma(z)](\eta) = \nu(\eta) - \mathbb{E}([\eta - G(z)]_+), \quad \eta \in I.$$

The set K is the cone of nonnegative functions in $\mathcal{C}(I)$. Observe that the function $z \mapsto \eta - G(z)$ is convex, for almost all $\omega \in \Omega$, and the function $x \mapsto (x)_+$ is convex and nondecreasing. Therefore, the composition $\mathbb{E}[(\eta - G(z))_+]$ is a convex function of z. It follows that the operator Γ is concave with respect to the cone K, that is, for any z_1, z_2 in Z and all $\lambda \in [0, 1]$,

$$\Gamma(\lambda z_1 + (1 - \lambda)z_2) - [\lambda\Gamma(z_1) + (1 - \lambda)\Gamma(z_2)] \in K.$$

As the topological dual space to $\mathcal{C}(I)$ is $\mathbf{rca}(I)$, we can introduce the Lagrangian $\Lambda : \mathcal{Z} \times \mathbf{rca}(I) \to \mathbb{R}$,

$$\Lambda(z, \mu) = f(z) + \int_I [\Gamma(z)](\eta)\, d\mu. \tag{4.14}$$

Let us observe that the uniform dominance condition implies that the following generalized Slater condition is satisfied: there exists a point $\tilde{z} \in Z$ such that $\Gamma(\tilde{z}) \in \operatorname{int} K$. Therefore we can use the necessary conditions of optimality in Banach spaces (see, e.g., [BS00, Theorem 3.4]). We conclude that there exists a measure $\hat{\mu} \in \mathbf{rca}^+(I)$ such that

$$\Lambda(\hat{z}, \hat{\mu}) = \max_{z \in Z} \Lambda(z, \hat{\mu}) \tag{4.15}$$

and

$$\int_I \big(\nu(\eta) - \mathbb{E}([\eta - G(\hat{z})]_+)\big)\, d\hat{\mu} = 0. \tag{4.16}$$

This means that $c_* = f(\hat{z}) = \Lambda(\hat{z}, \hat{\mu})$.

Define the set $U = \big\{(\beta, X, z) \in \mathbb{R} \times \mathcal{L}_1(\Omega, \mathcal{F}, P) \times Z : \beta \le f(z),\ X \le G(z)\big\}$. It follows from (4.15) that $\hat{\beta} = f(\hat{z})$, $\hat{X} = G(\hat{z})$ and $\hat{z}$ are the solution of the convex optimization problem

$$\underset{(\beta, X, z) \in U}{\text{maximize}}\ \beta - \iint_{I \times \Omega} [\eta - X]_+\, dP\, d\hat{\mu}. \tag{4.17}$$

Indeed, the best value of β is $f(z)$, and, due to the monotonicity of the function $x \mapsto -[\eta - x]_+$, the best value of X is $G(z)$. By carrying out the partial maximization with respect to (β, X) we reduce (4.17) to the right hand side of (4.15).

Consider the function $\varphi : \mathcal{S} \to \mathbb{R}$ defined by

$$\varphi(s) = \iint\limits_{I\times\Omega} \big[s(\eta,\omega)\big]_+ \, d\mathbb{P}\, d\hat{\mu},$$

at the point $\hat{s}(\eta,\omega) = \eta - [G(\hat{z})](\omega)$, $\eta \in I$, $\omega \in \Omega$. By virtue of the necessary and sufficient condition of optimality for problem (4.17), there exists a subgradient $\gamma \in \partial\varphi(\hat{s})$ such that $(\hat{\beta}, \hat{X}, \hat{z})$ is also a solution of the problem

$$\underset{(\beta,X,z)\in U}{\text{maximize}} \;\; \beta + \iint\limits_{I\times\Omega} \gamma(\eta,\omega) X \, d\mathbb{P}\, d\hat{\mu}. \tag{4.18}$$

By Strassen's disintegration theorem [Str65, Theorem 1],

$$\gamma(\eta,\omega) \in \partial\big(\eta - \hat{X}(\omega)\big)_+ = \begin{cases} \{1\} & \text{if } \eta > \hat{X}(\omega), \\ \{0\} & \text{if } \eta < \hat{X}(\omega), \\ [0,1] & \text{if } \eta = \hat{X}(\omega). \end{cases}$$

From the definition of the set U and from the fact that $\gamma(\eta,\omega) \geq 0$, for every value of z the best values of β and X in (4.18) are $f(z)$ and $G(z)$, respectively. It follows that $\hat{z}$ is an optimal solution of the problem

$$\underset{z\in Z}{\text{maximize}} \left\{ f(z) + \iint\limits_{I\times\Omega} \gamma(\eta,\omega) G(z) \, d\mathbb{P}\, d\hat{\mu} \right\}.$$

Define the measure $\hat{\lambda}$ as absolutely continuous with respect to $\hat{\mu} \times P$ with the Radon-Nikodym derivative

$$\frac{d\hat{\lambda}}{d(\hat{\mu} \times P)} = \gamma.$$

Since $0 \leq \gamma \leq 1$, we have $0 \leq \hat{\lambda} \leq \hat{\mu} \times P$. From (4.12) we obtain

$$\begin{aligned} L_D(\hat{\lambda}, \hat{\mu}) &= -\int\limits_I \eta \, d\hat{\lambda}_I(\eta) + \int\limits_I \nu(\eta)\, d\hat{\mu} + \sup_{z\in Z} \left\{ f(z) + \int\limits_\Omega [G(z)](\omega)\, d\hat{\lambda}_\Omega \right\} \\ &= -\iint\limits_{I\times\Omega} \eta\gamma(\eta,\omega)\, d\mathbb{P}\, d\hat{\mu} + \int\limits_I \nu(\eta)\, d\hat{\mu} + f(\hat{z}) + \iint\limits_{I\times\Omega} [G(\hat{z})](\omega)\gamma(\eta,\omega)\, d\mathbb{P}\, d\hat{\mu}. \end{aligned}$$

It follows from the definition of γ that the first and the last term in this expression can be written as

$$-\iint_{I\times\Omega} \eta\gamma(\eta,\omega)\, dP\, d\hat{\mu} + \iint_{I\times\Omega} [G(\hat{z})](\omega)\gamma(\eta,\omega)\, dP\, d\hat{\mu}$$
$$= -\iint_{I\times\Omega} \big(\eta - [G(\hat{z})](\omega)\big)_+\, dP\, d\hat{\mu} = -\int_I \mathbb{E}\big(\eta - [G(\hat{z})](\omega)\big)_+\, d\hat{\mu}.$$

Substituting into the last formula for $L_D(\hat{\lambda},\hat{\mu})$ we conclude that

$$L_D(\hat{\lambda},\hat{\mu}) = f(\hat{z}) + \int_I \Big[\nu(\eta) - \mathbb{E}\big(\eta - [G(\hat{z})](\omega)\big)_+\Big]\, d\hat{\mu} = f(\hat{z}) = c_*.$$

In the last equation we have used the complementarity condition (4.16). ■

Finally, let us observe that the condition $\lambda \le \mu \times P$ appearing in the dual problem implies that $\lambda_I \le \mu$. This is of importance for the solution method we describe later in §4.7.

4.5 Reducing the space of Lagrange multipliers

Notice that apart from the condition $\lambda \le \mu \times p$ the measure $\lambda \in \mathcal{M}$ on the product space $I \times \Omega$ appears in the dual optimization problem (4.13) via its marginal measures λ_I and λ_Ω. We can exploit this fact to achieve a reduction of the space of variables similar to that seen in the case of the primal problem. The main tool for this reduction is Strassen's theorem on the existence of measures with given marginals [Str65, Theorem 6]. We present here its version in the setting suitable for direct application to our problem.

Theorem 13 *Let $\kappa \in \mathcal{M}^+$, $\beta \in \mathbf{rca}(I)$, and let α be a measure on $(\Omega,\mathcal{F})$. There exists a measure $\lambda \in \mathcal{M}^+$ having marginal measures $\lambda_I = \beta$ and $\lambda_\Omega = \alpha$ and such that $\lambda \le \kappa$, if and only if*

$$\beta(B) + \alpha(A) \le \psi + \kappa(B \times A), \quad \forall\, B \in \mathcal{B}, \quad \forall\, A \in \mathcal{F},$$

where $\psi = \beta(I) = \alpha(\Omega)$.

Observe that setting $B = I$ we obtain $\alpha(A) \le \kappa(B \times A)$ for all $A \in \mathcal{F}$. Employing the definition of $\mathcal{M}^+$ we conclude that it is necessary that $d\alpha/dP \in \mathcal{L}_\infty(\Omega,\mathcal{F},P)$.

Applying Theorem 13 to the dual problem (4.13) with $\kappa = \mu \times P$, we obtain the following equivalent formulation of the dual problem:

$$\begin{aligned} \underset{\alpha,\beta,\mu,\psi}{\text{minimize}} \quad & -\int_I \eta \, d\beta + \int_I \nu \, d\mu + \sup_{z \in Z} \left(f(z) + \int_\Omega [G(z)](\omega) \, d\alpha \right) \\ \text{subject to} \quad & \beta(I) = \psi, \qquad \alpha(\Omega) = \psi, \\ & \beta(B) + \alpha(A) \le \psi + \mu(B)P(A), \quad \forall\, B \in \mathcal{B}, \quad \forall\, A \in \mathcal{F}, \\ & \alpha \ge 0, \quad \frac{d\alpha}{dP} \in \mathcal{L}_\infty(\Omega, \mathcal{F}, P), \quad \beta, \mu \in \mathbf{rca}^+(I). \end{aligned} \tag{4.19}$$

Note that the measure λ on the product space $I \times \Omega$ is eliminated from this formulation, at the cost of introducing new constraints indexed by the family $\mathcal{B} \times \mathcal{F}$. The merits of this tradeoff become apparent for problems with discrete distributions, where we propose a column generation method.

4.6 An implied transportation problem

We now focus again on the finite probability space $\Omega = \{\omega_1, \ldots, \omega_N\}$ with corresponding probabilities $p_1, \ldots, p_N$. The realizations of the benchmark outcome Y are denoted by $y_1, \ldots, y_D$ and the corresponding benchmark shortfalls are $\nu_j = \mathbb{E}([y_j - Y]_+)$.

We recall for convenience the dual problem (4.13) in this case. The measure λ becomes an array λ_{ij}, $i = 1, \ldots, N$, $j = 1, \ldots, D$. The marginal measures are its row and column sums, respectively. We obtain the formulation:

$$\begin{aligned} \underset{\lambda,\mu}{\text{minimize}} \quad & -\sum_{j=1}^{D} \sum_{i=1}^{N} \lambda_{ij} y_j + \sum_{j=1}^{D} \mu_j \nu_j + \sup_{z \in Z} \left\{ f(z) + \sum_{i=1}^{N} \sum_{j=1}^{D} \lambda_{ij} G_i(z) \right\} \\ \text{subject to} \quad & \lambda_{ij} \le p_i \mu_j, \quad i = 1, \ldots, N, \quad j = 1, \ldots, D, \\ & \lambda \ge 0, \quad \mu \ge 0. \end{aligned} \tag{4.20}$$

Consider the marginal sums

$$\alpha_i = \sum_{j=1}^{D} \lambda_{ij}, \quad i = 1, \ldots, N,$$

$$\beta_j = \sum_{i=1}^{N} \lambda_{ij}, \quad j = 1, \ldots, D.$$

Vectors $\alpha \geq 0$ and $\beta \geq 0$ are marginal sums of a feasible dual variable λ if and only if the following conditions are satisfied:

(i) for some $\psi \geq 0$ we have $\sum_{i=1}^{N} \alpha_i = \sum_{j=1}^{D} \beta_j = \psi$;

(ii) there exists a transportation flow of value ψ in the network having N source nodes with supplies α, D destination nodes with demands β, and with arc capacities equal to $p_i\mu_j$ for every arc (i,j), $i = 1, \ldots, N$, $j = 1, \ldots, D$.

In the discrete case Strassen's Theorem takes on the form of the *Maximum Flow–Minimum Cut Theorem* [FF56, Gal57] specialized for the above network: for all $A \subset \{1, \ldots, N\}$ and all $B \subset \{1, \ldots, D\}$

$$\sum_{i\in A} \alpha_i + \sum_{j\in B} \beta_j - \sum_{i\in A} p_i \sum_{j\in B} \mu_j \leq \psi. \tag{4.21}$$

The dual formulation based on this fact and corresponding to (4.19) is

$$\begin{aligned}
\underset{\alpha,\beta,\mu,\psi}{\text{minimize}} \quad & -\sum_{j=1}^{D} y_j\beta_j + \sum_{j=1}^{D} \nu_j\mu_j + \sup_{z\in Z} \left\{ f(z) + \sum_{i=1}^{N} \alpha_i G_i(z) \right\} \qquad (4.22)\\
\text{subject to} \quad & \sum_{i=1}^{N} \alpha_i = \psi, \qquad \sum_{j=1}^{D} \beta_j = \psi,\\
& \sum_{i\in A} \alpha_i + \sum_{j\in B} \beta_j - \sum_{i\in A} p_i \sum_{j\in B} \mu_j \leq \psi, \quad \forall\, A \subset \{1, \ldots, N\}, \quad \forall\, B \subset \{1, \ldots, D\},\\
& \alpha \geq 0, \quad \beta \geq 0, \quad \mu \geq 0, \quad \psi \geq 0.
\end{aligned}$$

In this formulation ND dual variables of (4.20) are replaced by $N + D$ marginal sums, at the cost of introducing 2^{N+D} new constraints, indexed by all possible sets A and B.

If the functions $f(\cdot)$ and $G_i(\cdot)$ are affine and the set Z is defined by linear equations and inequalities, then problem (4.22) becomes a linear programming problem. In a standard way, the term $\sup_{z\in Z} \left\{ f(z) + \sum_{i=1}^{N} \alpha_i G_i(z) \right\}$ can be replaced by an affine function of α and linear inequalities involving α. All these manipulations are the same as in linear programming duality theory. We illustrate them for a portfolio example in §4.8.

The main difficulty associated with problem (4.22) is the large number of constraints. We show in §4.7 a way to overcome this difficulty, by generating only a subset of relevant constraints.

4.7 A dual column generation method

Formulation (4.22) suggests a cutting plane method of the following form. At iteration k we have pairs of sets $A_m \subset \{1,\dots,N\}$ and $B_m \subset \{1,\dots,D\}$, $m = 1,\dots,k-1$. We solve a relaxation of problem (4.22):

$$\begin{aligned}
\underset{\alpha,\beta,\mu,\psi}{\text{minimize}} \quad & -\sum_{j=1}^{D} y_j\beta_j + \sum_{j=1}^{D} \nu_j\mu_j + \sup_{z\in Z}\Big\{ f(z) + \sum_{i=1}^{N} \alpha_i G_i(z) \Big\} \\
\text{subject to} \quad & \sum_{i=1}^{N} \alpha_i = \psi, \qquad \sum_{j=1}^{D} \beta_j = \psi, \\
& \sum_{i\in A_m} \alpha_i + \sum_{j\in B_m} \beta_j - \sum_{i\in A_m} p_i \sum_{j\in B_m} \mu_j \le \psi, \quad m = 1,\dots,k-1, \\
& \alpha \ge 0, \quad \beta \ge 0, \quad \mu \ge 0, \quad \psi \ge 0.
\end{aligned} \tag{4.23}$$

The next step is to verify inequalities (4.21) for all possible sets A and B, at the optimal solution $(\alpha^k, \beta^k, \mu^k, \psi^k)$. To this end, we find a pair A_k, B_k, which solves the problem

$$\underset{\substack{A\subset\{1,\dots,N\}\\ B\subset\{1,\dots,D\}}}{\text{maximize}} \quad -\psi^k + \sum_{i\in A} \alpha_i^k + \sum_{j\in B} \beta_j^k - \sum_{i\in A} p_i \sum_{j\in B} \mu_j^k. \tag{4.24}$$

Defining the complement event $A^c = \{1,\dots,N\} \setminus A$ we observe that the first three terms in (4.24) describe the required inflow to the set of nodes $A^c \cup B$. The last term in (4.24) is the total capacity of the arcs leading to this set, that is, the arcs starting in A and ending in B. It follows that problem (4.24) is a problem of finding a minimal cut in a bipartite graph. It can be solved in a very efficient way by special network algorithms, as described in [AMO93]. One method, which is closely related to our transformation, is the following. We formulate the maximum flow problem:

$$\begin{aligned}
\underset{\lambda}{\text{maximize}} \quad & \sum_{i=1}^{N}\sum_{j=1}^{D} \lambda_{ij} \\
\text{subject to} \quad & \sum_{j=1}^{D} \lambda_{ij} \le \alpha_i^k, \quad i = 1,\dots,N, && (4.25) \\
& \sum_{i=1}^{N} \lambda_{ij} \le \beta_j^k, \quad j = 1,\dots,D, && (4.26) \\
& 0 \le \lambda_{ij} \le p_i\mu_j^k, \quad i = 1,\dots,N, \quad j = 1,\dots,D.
\end{aligned}$$

If the flow equals ψ^k, then the optimal solution of (4.23) is also optimal for (4.22). Otherwise, we denote the Lagrange multipliers associated with (4.25) by $\zeta_i, = 1, \dots, N$, and the Lagrange multipliers associated with (4.26) by ξ_j, $j = 1, \dots, D$ (they all are equal either 0 or 1). We set

$$A_k = \{i : \zeta_i = 0\}, \qquad B_k = \{j : \xi_j = 0\},$$

we add the pair (A_k, B_k) to the pairs of sets included in (4.23), increase k by 1, and continue.

Observe that if the maximum in (4.24) is positive (and thus the maximum flow in the last displayed problem is smaller than ψ^k), the new cut is different from the cuts already included in problem (4.23). As the number of possible cuts is finite, the method must eventually stop at an optimal solution. In that case the flow in the network gives us the optimal values of the multipliers λ in the dual problem (4.20).

4.8 Numerical illustration

Let $R_1, \dots, R_n$ be random return rates of assets $1, \dots, n$. We denote the fractions of the initial capital invested in these assets by $z_1, \dots, z_n$. Clearly, the portfolio return rate equals

$$G(z) = R_1 z_1 + \cdots + R_n z_n.$$

The set of possible asset allocations is the simplex

$$Z = \{z \in \mathbb{R}^n : z_1 + \cdots + z_n = 1,\ z_k \geq 0,\ k = 1, \dots, n\},$$

but the approach outlined here easily extends to more general polyhedral sets Z. Finally, let a benchmark random return rate Y be given; for example, Y may represent the return rate of an index or the return rate of the current portfolio. The dominance-constrained portfolio optimization problem takes on the form

$$\begin{aligned} &\underset{z}{\text{maximize}}\ \mathbb{E}\big[R_1 z_1 + \cdots + R_n z_n\big] \\ &\text{subject to } R_1 z_1 + \cdots + R_n z_n \succeq_{(2)} Y \\ &\qquad\qquad z \in Z. \end{aligned}$$

This model was introduced as an example in [DR03] and analyzed in [DR06b].

As discussed in the introduction, no risk averse decision maker will prefer a portfolio with return rate Y over a portfolio with return rate $R_1z_1+\cdots+R_nz_n$. Therefore, the risk exposure of the portfolio return rate is "more acceptable" than that of Y. In our model we use the expected value of the portfolio return rate as the objective functional, and thus the entire burden of controlling risk it carried by the stochastic dominance constraint. As the distribution of the returns at the solution is indirectly shaped by the distribution of the benchmark Y, it is essential that Y be "acceptable," for example, the return rate of a broad market index. However, it is easy to additionally incorporate risk-averse objective functionals to our model, such as coherent measures of risk (see [ADEH99, FS02, FS04, MR08, RS06] and the references therein).

In the discrete distribution case, we denote the return of asset k in event i by r_{ik}, $i = 1,\dots,N$, $k = 1,\dots,n$, the probabilities of the elementary events by p_i, $i = 1,\dots,N$, the realizations of the benchmark returns by y_j, and the benchmark shortfalls by

$$\nu_j = \sum_{i=1}^{N} p_i(y_j - y_i)_+.$$

We obtain the problem:

$$\begin{aligned}
\underset{z}{\text{maximize}} \quad & \sum_{i=1}^{N}\sum_{k=1}^{n} p_i r_{ik} z_k \\
\text{subject to} \quad & \sum_{i=1}^{N} p_i\Big(y_j - \sum_{k=1}^{n} r_{ik}z_k\Big)_+ \le \nu_j, \quad j = 1,\dots,N, \\
& z \in Z.
\end{aligned}$$

The piecewise linear constraint is dealt with by the primal cutting plane method.

The dual problem (4.22) takes on the form:

$$\begin{aligned}
\underset{\alpha,\beta,\mu,\psi,\zeta}{\text{minimize}} \quad & -\sum_{j=1}^{N} y_j\beta_j + \sum_{j=1}^{N} \nu_j\mu_j + \zeta \\
\text{subject to} \quad & \sum_{i=1}^{N} \alpha_i = \psi, \qquad \sum_{j=1}^{N} \beta_j = \psi, \\
& \sum_{i=1}^{N} (p_i + \alpha_i) r_{ik} \le \zeta, \quad k = 1,\dots,n, \\
& \sum_{i\in A} \alpha_i + \sum_{j\in B} \beta_j - \sum_{i\in A} p_i \sum_{j\in B} \mu_j \le \psi, \ \ \forall\, A \subset \{1,\dots,N\}, \ \ \forall\, B \subset \{1,\dots,N\}, \\
& \alpha \ge 0, \quad \beta \ge 0, \quad \mu \ge 0, \quad \psi \ge 0.
\end{aligned}$$

The variable ζ in the objective function of this problem represents the term

$$\sup_{z\in Z}\Big\{f(z)+\sum_{i=1}^{N}\alpha_i G_i(z)\Big\} = \sup_{z\in Z}\Big\{\sum_{i=1}^{N}\sum_{k=1}^{n} p_i r_{ik} z_k + \sum_{i=1}^{N}\sum_{k=1}^{n}\alpha_i r_{ik} z_k\Big\}$$
$$= \max_{1\le k\le n}\sum_{i=1}^{N}(p_i+\alpha_i) r_{ik}.$$

The constraints involving the sets A and B are dealt with by the dual cutting plane method.

We considered several problem instances of different sizes, obtained from historical data on realizations of joint daily returns of $n = 500$ assets in N days, for seven different values of N, ranging from 50 to 1000. We used the returns in each day as equally probable realizations of the n-dimensional random vector R. The benchmark outcome Y was the return rate of the S&P 500 index. All calculations were carried out on a 2.00 GHz Pentium 4 PC with 1.00 GB of RAM, by using the AMPL modeling language [FGK02b] and with version 9.1 of the CPLEX solver [ILO05].

Table 4.1: Dimensions of the three formulations.

Scenarios	Linear Programming		Primal Formulation		Dual Formulation	
	Variables	Constraints	Variables	Constraints	Variables	Constraints
50	3000	2551	500	1	152	502
100	10500	10101	500	1	302	502
150	23000	22651	500	1	452	502
200	40500	40201	500	1	602	502
500	250500	250001	500	1	1502	502
750	563000	562501	500	1	2252	502
1000	1000500	1000001	500	1	3002	502

Table 4.2: Performance of the three approaches.

Scenarios	Linear Programming		Primal Method			Dual Method		
	CPU	Iterations	CPU	Cuts	Iterations	CPU	Cuts	Iterations
50	3.44	570	0.55	9	9	13.81	68	6883
100	20.23	3161	2.03	33	75	407.26	259	156304
150	372.52	7272	3.49	53	267	9144.25	552	1155166
200	373.63	16666	3.90	61	180	-	-	-
500	-	-	6.59	88	924	-	-	-
750	-	-	9.74	123	477	-	-	-
1000	-	-	10.23	117	530	-	-	-

Table 4.1 compares the sizes of the three formulations: the straightforward linear programming model (4.5), the primal cutting plane formulation (4.6), and the dual cutting plane formulation (4.19). In the last two cases we report the initial numbers of constraints only, without the cuts indexed by the sets $A \in \mathcal{F}$ and $B \in \mathcal{B}$. The numbers of cuts, which were actually

generated in the course of the solution, are reported in Table 4.2. This table provides also the CPU times of the simplex solver in the three cases, and the total numbers of simplex iterations performed.

It can be seen from these results that the primal cut generation method is quite efficient and it dramatically outperforms the direct linear programming approach. This is consistent with the results of [KHVDV06] for integrated chance constraints. In fact, the direct linear programming model was too large for our computer for 500 scenarios and more. The dual method is much slower for this problem class, mainly due to minimal differences between many cuts and severe numerical difficulties associated with that. For problems with $N = 200$ and more scenarios, we interrupted the calculation, because of excessive time. Apparently, the number of Strassen cuts is too large. However, we still believe that the dual formulation is interesting in its own right and that one day it may find its application.

Finally, Figure 4.8 compares the cumulative distribution functions of the return rates of the benchmark portfolio (the S&P 500 index) and of the solution to the dominance constrained problem for the case of 1000 scenarios. The optimal portfolio contains only 11 assets, but we can see that they are sufficient to shape the distribution function in a favorable way. Close inspection reveals that the optimal distribution function is not entirely below the benchmark (this would mean first order stochastic dominance); in the range between -0.02 and -0.015 it is slightly above. However, the expected shortfall (4.2) is always smaller at the solution than at the benchmark. This is in line with the results of [NR08], where similar examples are presented.

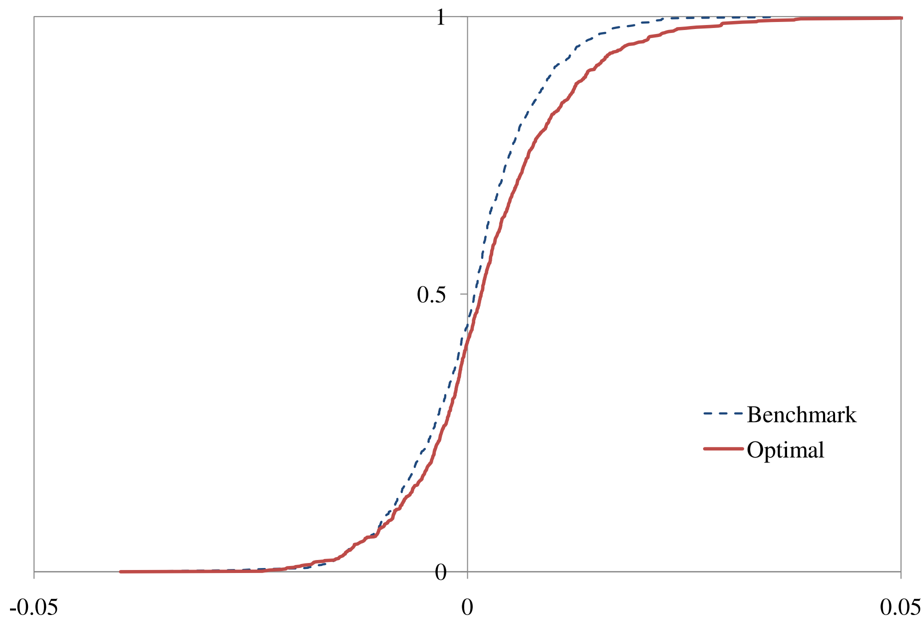

Figure 4.1: Cumulative distribution functions of the return rates of the benchmark and optimal portfolios in the 1000 scenario example.

Chapter 5

Strassen's Theorem Revisited

In this chapter we present a new proof of Strassen's Theorem about the existence of measures with given marginals. As we saw in Section 4.6, when applied to a discrete space, Strassen's Theorem yields a special case of the Maximum Flow–Minimum Cut Theorem. Our goal here is to go in the reverse direction and use the discrete result to prove the general case.

In order to proceed, we first establish a strong duality theorem; the exposition below closely follows the framework laid out in [Roc74].

5.1 Abstract Duality

Let $\mathcal{X}$ and $\mathcal{V}$ be paired locally convex topological vector spaces. That is, $\mathcal{X}$ and $\mathcal{V}$ are equipped with respective topologies which make them locally convex topological vector spaces and these topologies are compatible with the real-valued bilinear form $\langle x, v\rangle$. The latter condition means that every linear continuous functional on $\mathcal{X}$ can be represented in the form $\langle \cdot, v\rangle$ for some $v \in \mathcal{V}$, and every linear continuous functional on $\mathcal{V}$ can be represented in the form $\langle x, \cdot\rangle$ for some $x \in \mathcal{X}$. In particular, we can equip each space $\mathcal{X}$ and $\mathcal{V}$ with its weak topology induced by its paired space. This will make $\mathcal{X}$ and $\mathcal{V}$ paired locally convex topological vector spaces provided that for any $x \in \mathcal{X} \setminus \{0\}$ there exists $v \in \mathcal{V}$ such that $\langle x, v\rangle \neq 0$, and for any $v \in \mathcal{V} \setminus \{0\}$ there exists $x \in \mathcal{X}$ such that $\langle x, v\rangle \neq 0$. Further, let $\mathcal{U}$ and $\mathcal{Y}$ be paired locally convex topological vector spaces, with the corresponding bilinear form $\langle u, y\rangle$ (we use the same notation, because it will never lead to misunderstanding).

Suppose $A : \mathcal{X} \to \mathcal{U}$ is a continuous linear operator, $c \in \mathcal{V}$, Q is a closed convex cone in

$\mathcal{U}$, K is a closed convex cone in $\mathcal{X}$. Consider the problem

$$\begin{aligned} &\min \langle c, x\rangle \\ &\text{s.t. } Ax + u \in Q, \\ &x \in K. \end{aligned} \tag{5.1}$$

We consider u as a parameter of this problem and we are interested in the optimality conditions for (5.1) at $u = u_0$.

Define $\varphi : \mathcal{U} \to \overline{\mathbb{R}}$ as the infimum value of problem (5.1), as a function of $u \in \mathcal{U}$; if the problem is infeasible, we set $\varphi(u) = +\infty$.

Define the polar cone to K as

$$K^\circ = \{v \in \mathcal{V} : \langle x, v\rangle \leq 0,\ \forall x \in \mathcal{X}\}.$$

Similarly, we define the polar cone $Q^\circ \subset Y$.

Together with problem (5.1) we define the dual problem:

$$\begin{aligned} &\max \langle u, y\rangle \\ &\text{s.t. } A^* y + c \in -K^\circ, \\ &y \in Q^\circ. \end{aligned} \tag{5.2}$$

A^* is the adjoint of A.

The following theorem is a corollary of [Roc74, Thm.15].

Theorem 14 *Suppose that the optimal value function is lower semicontinuous at u_0. Then*

$$\inf (5.1) = \sup (5.2),$$

where in both problems we set $u = u_0$.

5.2 The Capacitated Mass Transportation Problem

Suppose Ω_1 and Ω_2 are measurable spaces, with σ-algebras $\mathcal{F}_1$ and $\mathcal{F}_2$, respectively, and let $\Omega = \Omega_1 \times \Omega_2$, $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2$. Suppose that on the spaces Ω_1 and Ω_2 regular nonnegative measures P_1 and P_2 are defined, and let P be their product measure. Consider $\mathcal{X}$ as the space of measures which are absolutely continuous with respect to P, with a square-integrable density.

Then we can identify $\mathcal{X}$ with the space $\mathcal{L}_2(\Omega, \mathcal{F}, P)$. In an analogous way we define the spaces $\mathcal{X}_1 = \mathcal{L}_2(\Omega_1, \mathcal{F}_1, P_1)$ and $\mathcal{X}_2 = \mathcal{L}_2(\Omega_2, \mathcal{F}_2, P_2)$. For a set $A \in \mathcal{F}$ define $\mathbb{1}_A(\omega) = 1$, if $\omega \in A$, and $\mathbb{1}_A(\omega) = 0$, if $\omega \notin A$. We write $\mathbb{1}$ for the constant 1 function of the corresponding space.

We define the linear operators $\Pi_1 : \mathcal{X} \to \mathcal{X}_1$ and $\Pi_2 : \mathcal{X} \to \mathcal{X}_2$ as follows:

$$\langle \Pi_1 x, \mathbb{1}_A \rangle = \langle x, \mathbb{1}_{A \times \Omega_2} \rangle, \quad \forall A \in \mathcal{F}_1,$$
$$\langle \Pi_2 x, \mathbb{1}_B \rangle = \langle x, \mathbb{1}_{\Omega_1 \times B} \rangle, \quad \forall B \in \mathcal{F}_2.$$

In other words, $\Pi_1(x)$ and $\Pi_2(x)$ are the marginal measures of x on Ω_1 and Ω_2, respectively.

In this setting Strassen's Theorem can be formulated as follows:

Theorem 15 *Consider two elements $u_1 \in \mathcal{X}_1$ and $u_2 \in \mathcal{X}_2$ (the marginal measures), and an element $\bar{x} \in \mathcal{X}$ (the upper bound). Assume that $\langle \mathbb{1}, u_1 \rangle = \langle \mathbb{1}, u_2 \rangle$, and denote this common value by $\varkappa$. Then there exists an element $x \in \mathcal{X}$ such that*

$$\begin{aligned} \Pi_1 x &= u_1, \\ \Pi_2 x &= u_2, \\ 0 \leq x &\leq \bar{x}. \end{aligned} \tag{5.3}$$

if and only if the condition

$$\int\limits_A u_1(\omega_1)\, dP_1 + \int\limits_B u_2(\omega_2)\, dP_2 + \iint\limits_{(\Omega_1 \setminus A) \times (\Omega_2 \setminus B)} \bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 \geq \varkappa.$$

holds for all $A \in \mathcal{F}_1$ and $B \in \mathcal{F}_2$.

The remainder of this chapter is devoted to the proof of this theorem. First, in order to be able to apply the duality results from the previous section, we convert (5.3) to a linear optimization problem.

Theorem 16 *System* (5.3) *has a solution if and only if problem* (5.4) *below has an optimal solution and its optimal value is at least $\varkappa$.*

$$\begin{aligned} \max\ & \langle \mathbb{1}, x\rangle \\ \textit{s.t.}\ & \Pi_1 x \le u_1, \\ & \Pi_2 x \le u_2, \\ & x \le \bar{x}, \\ & x \ge 0. \end{aligned} \tag{5.4}$$

Proof: If system (5.3) has a solution x^* then it is feasible for problem (5.4). By the definition of Π_1

$$\langle \mathbb{1}_{\Omega_1\times\Omega_2}, x^*\rangle = \langle \mathbb{1}_{\Omega_1}, \Pi_1 x^*\rangle = \langle \mathbb{1}_{\Omega_1}, u_1\rangle = \varkappa.$$

Thus the optimal value is at least $\varkappa$. Notice that it cannot be higher, because for every feasible solution x, by the above relations and the first constraint,

$$\langle \mathbb{1}_{\Omega_1\times\Omega_2}, x\rangle \le \langle \mathbb{1}_{\Omega_1}, u_1\rangle = \varkappa.$$

Now suppose that problem (5.4) has an optimal solution x^* with objective value $\varkappa$. Then

$$\varkappa = \langle \mathbb{1}_{\Omega_1\times\Omega_2}, x^*\rangle = \langle \mathbb{1}_{\Omega_1}, \Pi_1 x^*\rangle \le \langle \mathbb{1}_{\Omega_1}, u_1\rangle = \varkappa.$$

For this to hold true, we must in fact have $\Pi_1 x^* = u_1$. Similarly, $\Pi_2 x^* = u_2$, and thus x^* solves system (5.3). ■

Problem (5.4), after changing the sign of the objective, is of the form (5.1) with

$$c = -\mathbb{1}, \quad A = \begin{bmatrix} \Pi_1 \\ \Pi_2 \\ I \end{bmatrix}, \quad u = \begin{bmatrix} -u_1 \\ -u_2 \\ -\bar{x} \end{bmatrix},$$

K being the non-negative cone in $\mathcal{X}$, and $-Q$ the nonnegative cone in $\mathcal{X}_1 \times \mathcal{X}_2 \times \mathcal{X}$.

In this case the optimal value function $\varphi(u)$ is indeed lower semicontinuous. Consider a sequence $u^k \to u$, and the corresponding sequence of ε_k-optimal solutions x^k, with $\varepsilon_k \downarrow 0$. To verify lower semicontinuity, we only need to consider the case when $\varphi(u^k) < \infty$, and thus all x^k are indeed feasible for their problems. This implies that for every $\delta > 0$

$$0 \le x^k \le \bar{x} + \delta \mathbb{1},$$

for all sufficiently large k. This is a closed convex set, and in $\mathcal{L}_2$ it is weakly compact. Thus x^k have a weak limit x^* in this set. It is obvious that $0 \le x^* \le \bar{x}$, because $\delta > 0$ was arbitrary. Consider an arbitrary set $A \in \mathcal{F}_1$. As x^k is feasible for the kth problem, we have

$$\langle u_1^k, \mathbb{1}_A \rangle \ge \langle x^k, \mathbb{1}_{A\times\Omega_2} \rangle \to \langle x^*, \mathbb{1}_{A\times\Omega_2} \rangle.$$

The last relation follows from the weak convergence of x^k to x^*. We conclude that for every $A \in \mathcal{F}_1$

$$\langle x^*, \mathbb{1}_{A\times\Omega_2} \rangle \le \langle u_1, \mathbb{1}_A \rangle.$$

This means that $\Pi_1 x^* \le u_1$. Similarly, $\Pi_2 x^* \le u_2$. Consequently, x^* is feasible for the primal problem, and

$$\varphi(u) \le -\langle \mathbb{1}, x^* \rangle = -\lim \langle \mathbb{1}, x^k \rangle \le \liminf \left[\varphi(u_k) + \varepsilon_k \right] = \liminf \varphi(u^k).$$

We can thus use the duality theorem derived in Section 5.1.

The next step is to formulate the dual problem:

$$\begin{aligned} &\max\ -\langle u_1, y_1 \rangle - \langle u_2, y_2 \rangle - \langle \bar{x}, z \rangle \\ &\text{s.t. } \Pi_1^* y_1 + \Pi_2^* y_2 + z \ge \mathbb{1}, \\ &\quad y_1 \ge 0,\ y_2 \ge 0,\ z \ge 0, \end{aligned}$$

with $y_1 \in \mathcal{Y}_1 = \mathcal{L}_2(\Omega_1, \mathcal{F}_1, P_1)$, $y_2 \in \mathcal{Y}_2 = \mathcal{L}_2(\Omega_2, \mathcal{F}_2, P_2)$, $z \in \mathcal{Z} = \mathcal{L}_2(\Omega, \mathcal{F}, P)$.

This is equivalent (after changing the sign of the objective function) to the problem

$$\begin{aligned} &\min\ \langle u_1, y_1 \rangle + \langle u_2, y_2 \rangle + \langle \bar{x}, z \rangle \\ &\text{s.t. } \Pi_1^* y_1 + \Pi_2^* y_2 + z \ge \mathbb{1}, \\ &\quad y_1 \ge 0,\ y_2 \ge 0,\ z \ge 0. \end{aligned}$$

More explicitly, the dual problem can be reformulated as follows:

$$\min \int_{\Omega_1} y_1(\omega_1) u_1(\omega_1)\, dP_1 + \int_{\Omega_2} y_2(\omega_2) u_2(\omega_2)\, dP_2 + \iint_{\Omega_1\times\Omega_2} z(\omega_1,\omega_2)\bar{x}(\omega_1,\omega_2)\, dP_1\, dP_2 \tag{5.5}$$

$$\text{s.t. } y_1(\omega_1) + y_2(\omega_2) + z(\omega_1,\omega_2) \ge 1, \quad P\text{-a.s.}, \tag{5.6}$$

$$y_1(\omega_1) \ge 0,\ y_2(\omega_2) \ge 0,\ z(\omega_1,\omega_2) \ge 0, \quad P\text{-a.s.} \tag{5.7}$$

By Theorem 14 if the primal problem has optimal value $\varkappa$, the dual problem also has optimal value $\varkappa$. It follows that:

$$\int_{\Omega_1} y_1(\omega_1)u_1(\omega_1)\,dP_1 + \int_{\Omega_2} y_2(\omega_2)u_2(\omega_2)\,dP_2 + \iint_{\Omega_1\times\Omega_2} z(\omega_1,\omega_2)\bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2 \geq \varkappa,$$

for all functions y_1, y_2, and z such that

$$y_1(\omega_1) + y_2(\omega_2) + z(\omega_1,\omega_2) \geq 1, \quad \text{P-a.s.},$$

$$y_1(\omega_1) \geq 0,\ y_2(\omega_2) \geq 0,\ z(\omega_1,\omega_2) \geq 0, \quad \text{P-a.s.}$$

Setting $y_1 = \mathbb{1}_A$, $y_2 = \mathbb{1}_B$ we see that the worst value of z is

$$z = \mathbb{1}_{(\Omega_1\setminus A)\times(\Omega_2\setminus B)}.$$

We conclude that it is necessary that for all $A \in \mathcal{F}_1$ and $B \in \mathcal{F}_2$

$$\int_A u_1(\omega_1)\,dP_1 + \int_B u_2(\omega_2)\,dP_2 + \iint_{(\Omega_1\setminus A)\times(\Omega_2\setminus B)} \bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2 \geq \varkappa. \tag{5.8}$$

In the next section we show that this condition is also sufficient.

5.3 A discretization method

Assume that the common optimum of (5.4) and (5.5) is strictly less than $\varkappa$. Then there exists $\epsilon > 0$ and a feasible solution y_1^*, y_2^*, z^* of (5.5) such that

$$\int_{\Omega_1} y_1^*(\omega_1)u_1(\omega_1)\,dP_1 + \int_{\Omega_2} y_2^*(\omega_2)u_2(\omega_2)\,dP_2 + \iint_{\Omega_1\times\Omega_2} z^*(\omega_1,\omega_2)\bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2 < \varkappa - \epsilon.$$

Notice that, since the replacement of y_1^*, y_2^* and z^* by $\min(y_1^*, 1)$, $\min(y_2^*, 1)$ and $\min(z^*, 1)$, respectively, does not increase the objective value and preserves feasibility, we can assume $y_1^*, y_2^*, z^* \leq 1$ without loss of generality. Let us introduce the step functions

$$y_1^N(\omega_1) = \frac{\lceil N y_1^*(\omega_1)\rceil}{N}$$
$$y_2^N(\omega_2) = \frac{\lceil N y_2^*(\omega_2)\rceil}{N}$$
$$z^N(\omega_1,\omega_2) = \Big(1 - y_1^N(\omega_1) - y_2^N(\omega_2)\Big)_+$$

for some $N \geq \frac{2\varkappa}{\epsilon}$, where $(\cdot)_+$ is the *positive part* operation . It is readily verified that y_1^N, y_2^N, z^N is also a feasible solution of (5.5), and the relations $y_1^N \leq y_1^* + \frac{1}{N}$, $y_2^N \leq y_2^* + \frac{1}{N}$, $z^N \leq z^*$ and $y_1^N, y_2^N, z^N \leq 1$ hold.

Let $A_j = \left\{\omega_1 \in \Omega_1 \mid y_1^N(\omega_1) = \frac{j}{N}\right\}$ and $B_j = \left\{\omega_2 \in \Omega_2 \mid y_2^N(\omega_2) = \frac{j}{N}\right\}$ denote the level sets of y_1^N and y_2^N, respectively, for $j = 0, \dots, N$. Furthermore, let $\alpha_i = \int\limits_{A_i} u_1(\omega_1)\, dP_1$, $\beta_j = \int\limits_{B_j} u_2(\omega_2)\, dP_2$, $\gamma_{ij} = \iint\limits_{A_i \times B_j} \bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2$.

Lemma 5 *The optimum of the following linear programming problem with variables* $\tilde{\mathbf{y}}_1 \in \mathbb{R}^{N+1}$, $\tilde{\mathbf{y}}_2 \in \mathbb{R}^{N+1}$, $\tilde{\mathbf{z}} \in \mathbb{R}^{(N+1)\times(N+1)}$ *is strictly less than* $\varkappa$.

$$
\begin{array}{lll}
\min & \sum\limits_{i=0}^{N} \alpha_i \tilde{y}_1^i + \sum\limits_{j=0}^{N} \beta_j \tilde{y}_2^j + \sum\limits_{i=0}^{N}\sum\limits_{j=0}^{N} \gamma_{ij}\tilde{z}^{ij} & \\
\text{s.t.} & \tilde{y}_1^i + \tilde{y}_2^j + \tilde{z}^{ij} \geq 1 & i, j = 0, \dots, N \\
 & 0 \leq \tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \tilde{\mathbf{z}} \leq 1 &
\end{array}
\tag{5.9}
$$

Proof: The solution $\tilde{y}_1^i = \frac{i}{N}$, $\tilde{y}_2^j = \frac{j}{N}$, $\tilde{z}^{ij} = (1 - \frac{i}{N} - \frac{j}{N})_+$ is obviously feasible. Using $\int\limits_{\Omega_1} u_1(\omega_1)\, dP_1 = \int\limits_{\Omega_2} u_2(\omega_2)\, dP_2 = \varkappa$ and $N \geq \frac{2\varkappa}{\epsilon}$ the corresponding objective value is:

$$
\begin{aligned}
& \sum_{i=0}^{N} \alpha_i \tilde{y}_1^i + \sum_{j=0}^{N} \beta_j \tilde{y}_2^j + \sum_{i=0}^{N}\sum_{j=0}^{N} \gamma_{ij}\tilde{z}^{ij} \\
&= \sum_{i=0}^{N} \left(\int\limits_{A_i} u_1(\omega_1)\, dP_1 \right) \frac{i}{N} + \sum_{j=0}^{N} \left(\int\limits_{B_j} u_2(\omega_2)\, dP_2 \right) \frac{j}{N} + \sum_{i=0}^{N}\sum_{j=0}^{N} \left(\iint\limits_{A_i \times B_j} \bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 \right) (1 - \frac{i}{N} - \frac{j}{N})_+ \\
&= \sum_{i=0}^{N} \int\limits_{\mathcal{A}_i} y_1^N(\omega_1) u_1(\omega_1)\, dP_1 + \sum_{j=0}^{N} \int\limits_{\mathcal{B}_j} y_2^N(\omega_2) u_2(\omega_2)\, dP_2 + \sum_{i=0}^{N}\sum_{j=0}^{N} \iint\limits_{A_i \times B_j} z^N(\omega_1, \omega_2)\bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 \\
&= \int\limits_{\Omega_1} y_1^N(\omega_1) u_1(\omega_1)\, dP_1 + \int\limits_{\Omega_2} y_2^N(\omega_2) u_2(\omega_2)\, dP_2 + \iint\limits_{\Omega_1 \times \Omega_2} z^N(\omega_1, \omega_2)\bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 \\
&\leq \int\limits_{\Omega_1} (y_1^*(\omega_1) + \frac{1}{N}) u_1(\omega_1)\, dP_1 + \int\limits_{\Omega_2} (y_2^*(\omega_2) + \frac{1}{N}) u_2(\omega_2)\, dP_2 + \iint\limits_{\Omega_1 \times \Omega_2} z^*(\omega_1, \omega_2)\bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 \\
&= \int\limits_{\Omega_1} y_1^*(\omega_1) u_1(\omega_1)\, dP_1 + \int\limits_{\Omega_2} y_2^*(\omega_2) u_2(\omega_2)\, dP_2 + \iint\limits_{\Omega_1 \times \Omega_2} z^*(\omega_1, \omega_2)\bar{x}(\omega_1, \omega_2)\, dP_1\, dP_2 + 2\frac{\varkappa}{N} \\
&< \varkappa - \epsilon + \frac{2\varkappa}{N} \leq \varkappa.
\end{aligned}
$$

■

Now we can apply the integrality of the discretized system:

Lemma 6 *The matrix of the LP given in (5.9) is totally unimodular.*

Proof: The columns corresponding to the variable $\tilde{z}$ constitute an identity matrix, while the remaining columns form the incidence matrix of the complete bipartite graph $K_{N+1,N+1}$. ■

It follows from Lemma 6 that (5.9) has a $0-1$ valued optimal solution $\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \hat{\mathbf{z}}$. Let $A = \bigcup_{i:\hat{y}_1^i=1} A_i$ and $B = \bigcup_{j:\hat{y}_2^j=1} B_j$. Notice that (by the feasibility of this solution) $\hat{y}_1^i = \hat{y}_2^j = 0$ implies $z^{ij} = 1$. Using this fact and Lemma 5 we conclude that

$$\int_A u_1(\omega_1)\,dP_1 + \int_B u_2(\omega_2)\,dP_2 + \iint_{(\Omega_1\setminus A)\times(\Omega_2\setminus B)} \bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2$$

$$= \sum_{i:\hat{y}_1^i=1} \int_{A_i} u_1(\omega_1)\,dP_1 + \sum_{j:\hat{y}_2^j=1} \int_{B_j} u_2(\omega_2)\,dP_2 + \sum_{i:\hat{y}_1^i=0} \sum_{j:\hat{y}_1^j=1} \iint_{A_i\times B_j)} \bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2$$

$$\le \sum_{i=0}^{N} \left(\int_{A_i} u_1(\omega_1)\,dP_1 \right) \hat{y}_1^i + \sum_{j=0}^{N} \left(\int_{B_j} u_2(\omega_2)\,dP_2 \right) \hat{y}_2^j + \sum_{i:\hat{y}_1^i=0}^{N} \sum_{j=0}^{N} \left(\iint_{A_i\times B_j)} \bar{x}(\omega_1,\omega_2)\,dP_1\,dP_2 \right) \hat{z}^{ij} =$$

$$\sum_{i=0}^{N} \alpha_i \hat{y}_1^i + \sum_{j=0}^{N} \beta_j \hat{y}_2^j + \sum_{i=0}^{N} \sum_{j=0}^{N} \gamma_{ij} \hat{z}^{ij} < \varkappa.$$

Therefore the sets A and B violate condition (5.8), which completes the proof.

References

[ADEH99] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.

[AENR08] F. Alizadeh, J. Eckstein, N. Noyan, and G. Rudolf. Arrival rate approximation by nonnegative cubic splines. *Operations Research*, 56:140–156, 2008.

[AG03] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Math. Program.*, 95(1):3–51, 2003.

[AMO93] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows. Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[AS00] F. Alizadeh and S.H. Schmieta. Symmetric Cones, Potential Reduction Methods and Word-By-Word Extensions. In R. Saigal, L. Vandenberghe, and H. Wolkowicz, editors, *Handbook of Semidefinite Programming, Theory, Algorithms and Applications*, pages 195–233. Kluwer Academic Publishers, 2000.

[BHN99] R. M. Byrd, M. E. Hribar, and J. Nocedal. An interior point method for large scale nonlinear programming. *SIAM J. Optim.*, 9(4):877–900, 1999.

[BS00] J.F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer-Verlag, 2000.

[BVS02] H. Y. Benson, R. J. Vanderbei, and D. F. Shanno. Interior-point methods for nonconvex nonlinear programming: filter methods and merit functions. *Comput. Optim. Appl.*, 23(2):257–272, 2002.

[CMM98] J. Czyzyk, M. Mesnier, and J. Moré. The NEOS server. *IEEE J. Comput. Sci. Engrg.*, 5(3):68–75, 1998.

[dB78] C. de Boor. *A Practical Guide to Splines*. Springer-Verlag, New York, 1978.

[Dol01] E. Dolan. The NEOS server 4.0 administrative guide. Technical Report ANL/MCS-TM-250, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, 2001.

[DR] D. Dentcheva and A. Ruszczyński. Inverse cutting plane methods for optimization problems with second order stochastic dominance constraints. *Optimization*. To appear.

[DR03] D. Dentcheva and A. Ruszczyński. Optimization with stochastic dominance constraints. *SIAM Journal on Optimization*, 14:548–566, 2003.

[DR04a] D. Dentcheva and A. Ruszczyński. Optimality and duality theory for stochastic optimization problems with nonlinear dominance constraints. *Mathematical Programming*, 99:329–350, 2004.

[DR04b] D. Dentcheva and A. Ruszczyński. Semi-infinite probabilistic optimization: first order stochastic dominance constraints. *Optimization*, 53:583–601, 2004.

[DR06a] D. Dentcheva and A. Ruszczyński. Inverse stochastic dominance constraints and rank dependent expected utility theory. *Mathematical Programming*, 108:297–311, 2006.

[DR06b] D. Dentcheva and A. Ruszczyński. Portfolio optimization with stochastic dominance constraints. *Journal of Banking and Finance*, 30, 2006.

[DS97] H. Dette and W. J. Studden. *The Theory of Canonical Moments with Applications in Statistics, Probability, and Analysis*. Wiley, New York, 1997.

[Fay97] L. Faybusovich. Euclidean Jordan algebras and interior-point algorithms. *Positivity*, 1(4):331–357, 1997.

[Fay02] L. Faybusovich. Self-concordant barriers for cones generated by Chebychev systems. *SIAM J. Optim.*, 12(3):770–781, 2002.

[FF56] L.R. Ford and D.R. Fulkerson. Maximal flow through a network. *Canad. J. Math.*, 8:399–404, 1956.

[FGK02a] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming, 2nd edition*. Boyd and Fraser, Danvers, MA, 2002.

[FGK02b] R. Fourer, D.M. Gay, and B.W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. The Duxbury Press, 2002.

[FK94] J. Faraut and A. Korányi. *Analysis on Symmetric Cones*. Oxford University Press, Oxford, UK, 1994.

[FS02] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6:429–447, 2002.

[FS04] H. Föllmer and A. Schied. *Stochastic Finance. An Introduction in Discrete Time*. de Gruyter, 2004.

[Gal57] D. Gale. A theorem on flows in networks. *Pacific J. Math.*, 7:1073–1082, 1957.

[GE01] A. Gal and J. Eckstein. Managing periodically updated data in relational databases: a stochastic modeling approach. *J. ACM*, 46(6):1141–1183, 2001.

[GES03] A. Gal, J. Eckstein, and Z. Stoumbos. Scheduling of data transcription in periodically connected databases. *Stoch. Anal. Applic.*, 21(5):1021–1058, 2003.

[GM97] W. Gropp and J. Moré. Optimization environments and the NEOS server. In M. D. Buhman and A. Iserles, editors, *Approximation Theory and Optimization*, pages 167–182. Cambridge University Press, Cambridge, UK, 1997.

[Gül97] O. Güler. Personal communication. 1997.

[HR69] J. Hadar and W. Russell. Rules for ordering uncertain prospects. *The American Economic Review*, 59:25–34, 1969.

[ILO05] ILOG CPLEX. *CPLEX 9.0 Users Manual and Reference Manual*. Incline Village, 2005.

[KB00] M. E. Kuhl and P. S. Bhairgond. New frontiers in input modeling: nonparametric estimation of nonhomogeneous Poisson process using wavelets. In *Proc. 2000 Winter Simulation Conf.*, pages 562–571, Orlando, FL, 2000.

[KDW98] M. E. Kuhl, H. Damerdji, and J. R. Wilson. Least squares estimation of nonhomogeneous Poisson processes. In *Proc. 1998 Winter Simulation Conf.*, pages 637–646, Washington, DC, 1998.

[KH86] W.K. Klein Haneveld. *Duality in Stochastic Linear and Dynamic Programming*, volume 274 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, 1986.

[KHVDV06] W.K. Klein Haneveld and M.H. Van Der Vlerk. Integrated chance constraints: reduced forms and an algorithm. *Computational Management Science*, 3:245–269, 2006.

[Koe99] Koecher, Max. *The Minnesota Notes on Jordan Algebras and Their Applications*. Springer-Verlag, 1999. Edited by Kreig, Aloys and Walcher, Sebastian based on Lectures given in The University of Minnesota, 1960.

[KS66] S. Karlin and W. J. Studden. *Tchebycheff Systems, with Applications in Analysis and Statistics*. Wiley, New York, 1966.

[KSW06] M. E. Kuhl, S. G. Sumant, and J. R. Wilson. An automated multiresolution procedure for modeling complex arrival processes. *INFORMS J. Comput.*, 18(1):3–18, 2006.

[KW01] M. E. Kuhl and J. R. Wilson. Modeling and simulating Poisson processes having trends or nontrigonometric cyclic effects. *Europ. J. Oper. Res.*, 133(3):566–582, 2001.

[KWJ97] M. E. Kuhl, J. R. Wilson, and M. A. Johnson. Estimating and simulating Poisson processes having trends or multiple periodicities. *IIE Trans.*, 29(3):201–211, 1997.

[Leh55] E. Lehmann. Ordered families of distributions. *Annals of Mathematical Statistics*, 26:399–419, 1955.

[LS79] P. A. W. Lewis and G. S. Shedler. Simulation of nonhomogeneous Poisson processes by thinning. *Naval Res. Logist. Quart.*, 26(3):403–413, 1979.

[MR08] N. Miller and A. Ruszczyński. Risk-adjusted probability measures in portfolio optimization with coherent measures of risk. *European Journal of Operational Research*, 191:193–206, 2008.

[MS02] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, 2002.

[MW47] H.B. Mann and D.R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, 18:50–60, 1947.

[Nem99] A. Nemirovski. What can be expressed via conic quadratic and semidefinite programming?, August 1999. Talk presented at RUTCOR weekly seminar, Rutgers University.

[Nes00] Y. Nesterov. Squared functional systems and optimization problems. In J. B. G. Frenk, C. Roos, T. Terlaky, and S. Zahng, editors, *High Performance Optimization*, pages 405–440. Kluwer, Dordrecht, The Netherlands, 2000.

[NN94] Y. Nesterov and A. Nemirovski. *Interior Point Polynomial Methods in Convex Programming: Theory and Applications*. SIAM, Philadelphia, PA, 1994.

[NR08] N. Noyan and A. Ruszczyński. Valid inequalities and restrictions for stochastic programming problems with first order stochastic dominance constraints. *Mathematical Programming*, 114:249–275, 2008.

[NRA05] N. Noyan, G. Rudolf, and F. Alizadeh. Optimality Constraints for the Cone of Positive Polynomials. Technical Report 5-2005, Rutgers Center for Operations Research, Rutgers University, Piscataway, NJ, 2005.

[NRAP] N. Noyan, G. Rudolf, F. Alizadeh, and D. Papp. Bilinear optimality constraints for the cone of positive polynomials and related cones. *submitted for publication*.

[NRR06] N. Noyan, A. Ruszczyński, and G. Rudolf. Relaxations of linear programming problems with first order stochastic dominance constraints. *Operations Research Letters*, 34:653–659, 2006.

[NW03] J. Nocedal and R. A. Waltz. KNITRO user's manual. Technical Report OTC 2003/05, Optimization Technology Center, Northwestern University, Evanston, IL, 2003.

[OR99] W. Ogryczak and A. Ruszczyński. From stochastic dominance to mean–risk models: semideviations as risk measures. *European Journal of Operational Research*, 116:33–50, 1999.

[OR01] W. Ogryczak and A. Ruszczyński. On consistency of stochastic dominance and mean–semideviation models. *Mathematical Programming*, 89:217–232, 2001.

[OR02] W. Ogryczak and A. Ruszczyński. Dual stochastic dominance and related mean–risk models. *SIAM Journal on Optimization*, 13:60–78, 2002.

[PA08] D. Papp and F. Alizadeh. Linear and Second Order Cone Programming Approaches to Statistical Estimation Problems. Technical Report RRR 13-2008, RUTCOR, Rutgers Center For Operations Research, 2008. Submitted.

[QS62] J.P. Quirk and R. Saposnik. Admissibility and measurable utility functions. *Review of Economic Studies*, 29:140–146, 1962.

[Roc70] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.

[Roc74] R.T. Rockafellar. *Conjugate Duality in Optimization*. SIAM, 1974.

[RR98] R. Reemtsen and J.-J. Rückmann, editors. *Semi-Infinite Programming*. Kluwer, 1998.

[RR08] G. Rudolf and A. Ruszczyński. Optimization problems with second order stochastic dominance constraints: duality, compact formulations, and cut generation methods. *SIAM Journal on Optimization*, 19(3):1326–1343, 2008.

[RS03] A. Ruszczyński and A. Shapiro, editors. *Stochastic Programming*. Elsevier Science, 2003.

[RS06] A. Ruszczyński and A. Shapiro. Optimization of convex risk functions. *Mathematics of Operations Research*, 31:433–452, 2006.

[RU02] R.T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26:1443–1471, 2002.

[Sch81] L. L. Schumaker. *Spline Functions: Basic Theory*. Wiley, New York, 1981.

[Sco76] D. W. Scott. *Nonparametric probability density estimation by optimization theoretic techniques*. PhD thesis, Rice University, Houston, TX, 1976.

[Sha93] J. Shao. Linear model selection by cross-validation. *J. Amer. Statist. Assoc.*, 88:486–494, 1993.

[SS94] M. Shaked and J.G. Shanthikumar. *Stochastic Orders and Their Applications*. Academic Press, Boston, 1994.

[Str65] V. Strassen. The existence of probability measures with given marginals. *Annals of Mathematical Statistics*, 38:423–439, 1965.

[STT80] D. W. Scott, R. A. Tapia, and J. R. Thompson. Nonparametric probability density estimation by discrete maximum penalized-likelihood criteria. *Ann. Statist.*, 8(4):820–832, 1980.

[Stu01] Jos F. Sturm. *Using SeDuMi 1.02, a Matlab Toolbox for Optimization over Symmetric Cones (Updated for Version 1.05)*, 2001. Available from `http://sedumi.mcmaster.ca`.

[TT90] J. R. Thompson and R. A. Tapia. *Nonparametric Function Estimation, Modeling, and Simulation*. SIAM, Philadelphia, PA, 1990.

[Wah90] G. Wahba. *Spline Models for Observational Data*. SIAM, Philadelphia, PA, 1990.

[WB06] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, 2006.

[WS98] G. G. Walter and X. Shen. Positive estimation with wavelets. In *Wavelets, Multiwavelets, and their Applications*, volume 216 of *Contemporary Mathematics*, pages 63–79. American Mathematical Society, Providence, RI, 1998.

[WSV00] H. Wolkowicz, R. Saigal, and L. Vandenberghe, editors. *Handbook of Semidefinite Programming: Theory, Algorithms and Applications*. Kluwer, Dordrech, The Netherlands, 2000.

[Yaa87] M.E. Yaari. The dual theory of choice under risk. *Econometrica*, 55:95–115, 1987.

Vita

Gábor Rudolf

2002-2009	Rutgers University, New Brunswick, NJ, Ph.D., Operations Research, October 2009
1997-2002	ELTE University, Budapest, Hungary M.S., Mathematics, June 2002

Publications

Rudolf, G., A. Ruszczyński, Optimization Problems With Second Order Stochastic Dominance Constraints: Duality, Compact Formulations, and Cut Generation Methods, SIAM Journal on Optimization 19 (2008), No. 3

Khachiyan, L., E. Boros, K. Borys, K. Elbassioni, V. Gurvich, G. Rudolf, J. Zhao, On Short Paths Interdiction problems: Total and Node-wise Limited Interdiction, Theory of Computing Systems, 43(2)2008.

Lozin, V., G. Rudolf, Minimal Universal Bipartite Graphs, Ars Combinatoria, 84(2007).

Boros, E., K. Borys, K. Elbassioni, V. Gurvich, K. Makino, G. Rudolf, Generating Minimal k-Vertex Connected Spanning Subgraphs. Computing and Combinatorics, 13th Annual International Conference, COCOON Proceedings (2007).

Boros, E., K. Borys, V. Gurvich, G. Rudolf, Generating 3-vertex connected spanning subgraphs, Discrete Mathematics Vol. 308(2008) Issue 24

Noyan, N., G. Rudolf, A. Ruszczyński, Relaxations of Linear Programming Problems with First Order Stochastic Dominance Constraints, Operations Research Letters 34(2006).

Rudolf, G., On the Sharpness and Generalizations on the Fine-Wilf Periodicity Theorem, Pure Mathematics and Applications, Vol.13, 1-2 (2003).