Tilburg University

Practical robust optimization techniques and improved inverse planning of HDR brachytherapy

Gorissen, B.L.

Publication date:
2014

Document Version
Publisher's PDF, also known as Version of record

Link to publication in Tilburg University Research Portal

Bram L. Gorissen

Practical robust optimization techniques and improved inverse planning of HDR brachytherapy

Practical robust optimization techniques and improved inverse planning of HDR brachytherapy

PROMOTOR: prof. dr. ir. D. den Hertog
COPROMOTOR: dr. ir. A. L. Hoffmann

OVERIGE LEDEN VAN DE PROMOTIECOMMISSIE:
            prof. dr. E. de Klerk
            prof. dr. D. J. Bertsimas
            prof. dr. B. J. M. Heijmen
            dr. R. C. M. Brekelmans
            dr. J. C. Vera Lizcano

Dankwoord

Ik ben alle "condiciones sine quibus non" veel dank verschuldigd. In het navolgende zal ik enkele personen apart bedanken.

Ik dank mijn promotor (Dick den Hertog), voor het aanbieden van een promotiepositie en de uitstekende begeleiding en discussies die daarop zijn gevolgd. Ik dank ook mijn copromotor (Aswin Hoffmann), voor het delen van zijn uitmuntende kennis van brachytherapie en zijn nieuwsgierigheid naar wiskundige optimalisatie. Ik dank de overige coauteurs bij mijn publicaties voor hun bijdragen, en de reviewers die hun commentaar op de verschillende hoofdstukken hebben geuit.

Enkelen die zich als heldere sterren hebben gemanifesteerd, en zodoende mijn pad naar dit proefschrift hebben gemarkeerd, bedank ik door ze nadrukkelijk te noemen. Het pad begon bij een familielid dat meende dat de nieuwste generatie nooit zo goed zou worden in hoofdrekenen als hij, een uitspraak waarin een zekere uitdaging schuilging. De stap van rekenen naar wiskunde bleek een kleine. Op een open dag van een middelbare school pakte ik een velletje met wiskundige puzzels mee. De enthousiaste wiskundedocent die deze puzzels heeft verspreid (dhr. Roos), heeft later zijn vrije tijd opgeofferd om mij onder meer de principes van de volledige inductie bij te brengen. Andere personen die mijn middelbareschoolperiode hebben getekend zijn mijn voormalige natuurkundedocent (dhr. Hoekstra) en informaticadocent (dhr. De Geest). Informatica, dat leek mij wel. 'Heb je wel eens naar econometrie of astronomie gekeken?', vroeg de lokatieleider van de middelbare school (dhr. Van der Vaart) mij, toen ik door baldadig gedrag met hem in aanraking kwam. Via de studievereniging der econometristen in Tilburg kon ik een dag meelopen, en een echt college analyse volgen. Krap vijf jaar later dacht ik klaar te zijn in Tilburg, maar wist mijn promotor mijn interesse met interessante projecten en warme contacten te wekken.

Tot slot wil ik familie, vrienden en oudcollegae[1] bedanken die mijn leven op positieve wijze hebben gekleurd. In het bijzonder bedank ik daarbij mijn ouders, voor hun goede en liefdevolle zorgen.

---

[1] Former colleagues: this is where I express my gratitude towards you.

ii

Contents

List of abbreviations

AARC	Affinely adjustable robust counterpart
ARC	Adjustable robust counterpart
ARO	Adjustable robust optimization
CQP	Conic quadratic program
DTMR	Dwell time modulation restriction
DVH	Dose-volume histogram
EBRT	External beam radiation therapy
FP	Fractional program
GLP	Generalized linear program
GRC	Globalized robust counterpart
GTV	Gross tumor volume
HDR-BT	High-dose-rate brachytherapy
IMRT	Intensity modulated radiation therapy
IPS	Inner approximation of the pareto set
LCP	Linear conic program
LP	Linear program
MILP	Mixed-integer linear program
MIP	Mixed-integer program
MIQP	Mixed-integer quadratic program
MOP	Multiobjective optimization problem
OAR	Organ at risk
OD	Optimistic dual
OPS	Outer approximation of the pareto set
PS	Pareto set
PTV	Planning target volume
QA	Quality Assurance
QP	Quadratic program
RC	Robust counterpart
RO	Robust optimization
SDP	Semidefinite program
SP	Stochastic programming
TRUS	Transrectal ultrasound

CHAPTER 1

Introduction

This dissertation comprises two topics, which are treated individually and get a separate introduction. The reader will be acquainted with Robust Optimization (RO) in Section 1.1, while inverse treatment planning of high-dose-rate brachytherapy (HDR-BT) for prostate cancer is introduced in Section 1.2.

While the topics seem unrelated, the majority of the chapters in this dissertation stems from the initial goal to apply RO to inverse treatment planning of HDR-BT. That goal turned out to be far-headed, since RO could not directly be applied to the optimization models used in HDR-BT, while those optimization models needed to be improved to relate more closely to the dosimetric goals of the treatment planner. The work in this thesis is valuable for RO and for inverse treatment planning separately, but will also make it easier to robustly optimize HDR-BT. Section 1.2.3 shows ideas for this.

1.1 Robust Optimization

Optimization problems are often affected by uncertainty. There are two leading generic methods that deal with this uncertainty. The first is Stochastic Programming (SP, see Prékopa, 1995; Ruszczyński and Shapiro, 2003), the second is RO (Ben-Tal et al., 2009a; Bertsimas et al., 2011). Section 1.1.2 highlights the differences between the two.

1.1.1 Robust Optimization: the paradigm

Robust Optimization (RO) is a paradigm for dealing with uncertain data in an optimization problem. The basic parts of RO originate from the seventies and eighties (Soyster, 1974; Thuente, 1980; Singh, 1982; Kaul et al., 1986), but most of the existing theory and applications followed after new results in the late nineties (Ben-Tal and Nemirovski, 1998; El Ghaoui and Lebret, 1997). An extensive overview of RO is given in (Ben-Tal et al., 2009a) and the survey (Bertsimas et al., 2011). The basic

idea of RO is that constraints have to hold for all parameter realizations in some given uncertainty region.

While RO can be applied to many optimization problems, let us demonstrate its use on a linear optimization problem. The "general" formulation of an uncertain linear optimization problem is as follows:

$$\max_{\boldsymbol{x} \geq \boldsymbol{0}} \left\{ \boldsymbol{c}^{\top} \boldsymbol{x} : \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b} \right\}_{(\boldsymbol{c}, \boldsymbol{A}, \boldsymbol{b}) \in \mathcal{U}}, \tag{1.1}$$

where $\boldsymbol{c} \in \mathbb{R}^n$, $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^m$ denote the uncertain coefficients, $\boldsymbol{x} \in \mathbb{R}^n$ is the optimization variable, and $\mathcal{U}$ denotes the user specified uncertainty set. The "basic" RO paradigm is based on the following three assumptions (Ben-Tal et al., 2009a, p. xii):

1. All decision variables $\boldsymbol{x}$ represent "here and now" decisions: they should get specific numerical values as a result of solving the problem before the actual data "reveals itself".

2. The decision maker is fully responsible for consequences of the decisions to be made when, and only when, the actual data is within the prespecified uncertainty set $\mathcal{U}$.

3. The constraints of the uncertain problem in question are "hard" - the decision maker cannot tolerate violations of constraints when the data is in the prespecified uncertainty set $\mathcal{U}$.

Without loss of generality, the objective coefficients ($\boldsymbol{c}$) and the right-hand side values ($\boldsymbol{b}$) can be assumed to be certain, since an uncertain objective may be reformulated as a constraint via an epigraph reformulation, and an uncertain right-hand side can be modeled as a column of the coefficient matrix $\boldsymbol{A}$ where the corresponding entry in $\boldsymbol{x}$ is -1. Often there is a small number of driving factors, called the primitive uncertainties $\boldsymbol{\zeta} \in \mathbb{R}^L$, such that the uncertain parameter $\boldsymbol{A}$ is an affine function of $\boldsymbol{\zeta}$:

$$\boldsymbol{A}(\boldsymbol{\zeta}) = \boldsymbol{A^0} + \sum_{\ell=1}^{L} \zeta_\ell \boldsymbol{A^\ell},$$

where $\boldsymbol{A^0}$ is the nominal value matrix, $\boldsymbol{A^\ell}$ are the shifting matrices, and $\mathcal{Z}$ is the user specified primitive uncertainty set. The robust reformulation of (1.1) that is generally referred to as the *robust counterpart* (RC), is then as follows:

$$\max_{\boldsymbol{x} \geq \boldsymbol{0}} \left\{ \boldsymbol{c}^{\top} \boldsymbol{x} \ : \ \boldsymbol{A}(\boldsymbol{\zeta})\boldsymbol{x} \leq \boldsymbol{b} \quad \forall \boldsymbol{\zeta} \in \mathcal{Z} \right\}.$$

This means that a solution $\boldsymbol{x}$ is robust feasible if it satisfies the uncertain constraints $[\boldsymbol{A}(\boldsymbol{\zeta})\boldsymbol{x} \leq \boldsymbol{b}]$ for all realizations of $\boldsymbol{\zeta}$ contained in the uncertainty set $\mathcal{Z}$. It can be

shown that the RC remains unchanged if the "for all" quantifier is applied per row (Ben-Tal et al., 2009a, p. 11):

$$\max_{\boldsymbol{x} \geq \boldsymbol{0}} \left\{ \boldsymbol{c}^\top \boldsymbol{x} \ : \ \left(\boldsymbol{a_i^0} + \boldsymbol{A_i}\boldsymbol{\zeta}\right)^\top \boldsymbol{x} \leq b_i \quad \forall \boldsymbol{\zeta} \in \mathcal{Z} \ \forall i \in I \right\},$$

where $\boldsymbol{a_i^0}$ is the transpose of the i^{th} row of $\boldsymbol{A^0}$, and the ℓ^{th} column of $\boldsymbol{A_i}$ is given by the i^{th} row of $\boldsymbol{A^\ell}$ ($\ell = 1, \ldots, L$).

Since the set $\mathcal{Z}$ may be uncountable, e.g. an ellipsoid or a polyhedron, the RC seems hard to solve at first sight as it has an infinite number of constrains. The power of RO is to reformulate the RC to an equivalent optimization problem with a finite number of variables and constraints. Two generic methods can be distinguished, while there are also specialized results based on the $\mathcal{S}$-lemma, sums of squares or enumeration of the vertices of $\mathcal{Z}$. The first generic method uses conic duality (e.g. used by Ben-Tal et al. (2009a)), while the second method uses Fenchel duality (Ben-Tal et al., 2014). Let us demonstrate these approaches on a single constraint with polyhedral uncertainty:

$$\left(\boldsymbol{a_i^0} + \boldsymbol{A_i}\boldsymbol{\zeta}\right)^\top \boldsymbol{x} \leq b_i \quad \forall i \quad \forall \boldsymbol{\zeta} : \boldsymbol{B_i}\boldsymbol{\zeta} \leq \boldsymbol{d_i}. \tag{1.2}$$

The conic duality approach is based on the following reformulation:

$$(\boldsymbol{a_i^0})^\top \boldsymbol{x} + \max_{\boldsymbol{\zeta} \in \mathbb{R}^L} \left\{ (\boldsymbol{A_i}\boldsymbol{\zeta})^\top \boldsymbol{x} : \boldsymbol{B_i}\boldsymbol{\zeta} \leq \boldsymbol{d_i} \right\} \leq b_i.$$

The left hand side contains a conic optimization problem (in fact, a linear program), whose value is equal to its dual. After replacing the optimization problem in the left hand side with its dual, the constraint becomes:

$$(\boldsymbol{a_i^0})^\top \boldsymbol{x} + \min_{\boldsymbol{w_i} \geq \boldsymbol{0}} \left\{ \boldsymbol{d_i}^\top \boldsymbol{w_i} : \boldsymbol{B_i}^\top \boldsymbol{w_i} = \boldsymbol{A_i}^\top \boldsymbol{x} \right\} \leq b_i. \tag{1.3}$$

Note that if this constraint holds for some $\boldsymbol{w_i}$, then it definitely holds for the minimum. The min operator may therefore be omitted. The final optimization problem becomes:

$$\max_{\boldsymbol{x} \geq \boldsymbol{0}, \boldsymbol{w_i} \geq \boldsymbol{0}} \left\{ \boldsymbol{c}^\top \boldsymbol{x} \ : \ (\boldsymbol{a_i^0})^\top \boldsymbol{x} + \boldsymbol{d_i}^\top \boldsymbol{w_i} \leq b_i, \quad \boldsymbol{B_i}^\top \boldsymbol{w_i} = \boldsymbol{A_i}^\top \boldsymbol{x} \quad \forall i \right\}.$$

The second approach uses Fenchel duality. Constraint (1.2) can be written as:

$$\max_{\boldsymbol{\zeta} \in \mathbb{R}^n} \left\{ f(\boldsymbol{\zeta}, \boldsymbol{x}) - g(\boldsymbol{\zeta}) \right\} \leq b_i,$$

with $f(\boldsymbol{\zeta}, \boldsymbol{x}) = \left(\boldsymbol{a_i^0} + \boldsymbol{A_i}\boldsymbol{\zeta}\right)^\top \boldsymbol{x}$, and $g(\boldsymbol{\zeta}) = \delta(\boldsymbol{\zeta}|\mathcal{Z})$, the indicator function on the uncertainty region taking the value 0 if $\boldsymbol{\zeta}$ is in $\mathcal{Z}$ and ∞ otherwise. By Fenchel's duality theorem, this constraint holds if and only if:

$$\min_{\boldsymbol{s_i} \in \mathbb{R}^n} \left\{ g^*(\boldsymbol{s_i}) - f_*(\boldsymbol{s_i}, \boldsymbol{x}) \right\} \leq b_i, \tag{1.4}$$

where the asterisk indicates the concave conjugate of f (with respect to the first argument) and the convex conjugate of g. These are computed as:

$$f_*(\boldsymbol{s_i}, \boldsymbol{x}) = \inf_{\boldsymbol{\zeta} \in \mathbb{R}^n} \left\{ \boldsymbol{\zeta}^\top \boldsymbol{s_i} - \left(\boldsymbol{a_i^0} + \boldsymbol{A_i} \boldsymbol{\zeta} \right)^\top \boldsymbol{x} \right\} = \begin{cases} -(\boldsymbol{a_i^0})^\top \boldsymbol{x} & \text{if } \boldsymbol{s_i} = \boldsymbol{A_i}^\top \boldsymbol{x} \\ -\infty & \text{otherwise} \end{cases}$$

$$\text{and } g^*(\boldsymbol{s_i}) = \sup_{\boldsymbol{\zeta} \in \mathbb{R}^n} \left\{ \boldsymbol{\zeta}^\top \boldsymbol{s_i} - \delta(\boldsymbol{\zeta} | \mathcal{Z}) \right\} = \sup_{\boldsymbol{\zeta} \in \mathbb{R}^n} \left\{ \boldsymbol{\zeta}^\top \boldsymbol{s_i} : \boldsymbol{B_i} \boldsymbol{\zeta} \leq \boldsymbol{d_i} \right\}$$

$$= \min_{\boldsymbol{w_i} \geq \boldsymbol{0}} \left\{ \boldsymbol{d_i}^\top \boldsymbol{w_i} : \boldsymbol{B_i}^\top \boldsymbol{w_i} = \boldsymbol{s_i} \right\}.$$

Plugging these into (1.4) yields:

$$\min_{\boldsymbol{s_i} \in \mathbb{R}^n} \left\{ \min_{\boldsymbol{w_i} \geq \boldsymbol{0}} \left\{ \boldsymbol{d_i}^\top \boldsymbol{w_i} : \boldsymbol{B_i}^\top \boldsymbol{w_i} = \boldsymbol{s_i} \right\} + (\boldsymbol{a_i^0})^\top \boldsymbol{x} : \boldsymbol{s_i} = \boldsymbol{A_i}^\top \boldsymbol{x} \right\} \leq b_i,$$

where the second min operator may be omitted. This constraint is equivalent to (1.3).

The second method captures all situations in which the first method can be used. Moreover, it is applicable when the constraints are nonlinear in the uncertain parameters or when the uncertainty set is not conic representable, as long as the concave conjugate has a closed-form expression or can be expressed as the minimum of a convex optimization problem.

Specific results exist for uncertain Conic Quadratic Programs or Semidefinite Programs (Ben-Tal et al., 2009a, Part II). These problems are not concave in the uncertain parameters, and can therefore not be solved with standard techniques.

There are extensions of RO that relax some of the three underlying assumptions from page 4. Adjustable RO (ARO) does not require the first assumption, i.e., some decisions may be taken after the uncertain data has revealed itself (Ben-Tal et al., 2004; Chen et al., 2008). The ARO solution provides a *decision rule* that dictates the value of the adjustable variables as a function of the uncertain parameters. For example, $x(\zeta) = 1 - \zeta$ (which is optimal for the example in Section 1.1.2). The Globalized RC (GRC) allows constraint violations for some realizations of the uncertain parameters, but the magnitude of the violation is still controlled (Ben-Tal et al., 2006). This relaxes the third assumption that the constraints should always be satisfied for the entire uncertainty set.

1.1.2 Numerical example and difference with Stochastic Programming

This section explains the difference between RO and SP. Both methods are demonstrated using a simple example of minimizing a convex real-valued function on the

Figure 1.1 – Graph of a univariate objective function

interval $[0, 5]$:

$$f(x) = \begin{cases} 6 - 5x & \text{if } 0 \leq x \leq 1 \\ 0.9 + 0.1x & \text{if } 1 < x \leq 5. \end{cases}$$

This function has been depicted in Figure 1.1. Without uncertainty, the optimal solution is $x = 1$, resulting in an objective value of 1. Now suppose there is an implementation error. Instead of x, $x + \zeta$ is implemented, resulting in an objective value of $f(x + \zeta)$, where ζ is a random perturbation taking values in $[-1, 1]$. By selecting $x = 1$, the objective value could become as high as 6. SP and RO will select a slightly larger x to avoid the possibility of implementing $x + \zeta = 0$, but do so in a very distinctive manner. To avoid infeasible solutions, x is restricted to the subinterval $[1, 4]$.

SP is founded on probability distributions. Given the probability distribution of ζ, a typical objective function in SP is the expected value:

$$\min_{x \in [1,4]} \mathbb{E}_\zeta \left[f(x + \zeta) \right],$$

i.e., determine x such that the expected objective value is minimized. For example, if the probability density function of ζ on the interval $[-1, 1]$ is given by $1 - \zeta$, the problem becomes:

$$\min_{x \in [1,4]} \int_{-1}^{1} (1 - \zeta) f(x + \zeta) d\zeta = \min_{x \in [1,4]} \begin{cases} \frac{17}{20}x^3 - 10x + \frac{46}{3} & \text{if } x \leq 2 \\ \frac{1}{5}x + \frac{26}{15} & \text{if } x > 2. \end{cases}$$

The optimal solution is $x = \sqrt{200/51}$.

RO on the other hand is founded on uncertainty regions. Given the uncertainty region $\mathcal{U}$, basic versions of RO solve the problem:

$$\min_{x\in[1,4]} \max_{\zeta\in\mathcal{Z}} f(x+\zeta),$$

i.e., determine x such that the worst case objective value is minimized. For example, if $\mathcal{Z} = [-1, 1]$, the problem becomes:

$$\min_{x\in[1,4]} \max_{\zeta\in[-1,1]} f(x+\zeta) = \min_{x\in[1,4]} \max\{f(x-1), f(x+1)\},$$

where the equality follows from the fact that a convex function takes its maximum over a closed set at the boundary of that set. It can be shown that the optimal solution is $x = 100/51$.

This dissertation does not consider SP, even though it may be a better approach for some problems. Nevertheless, RO has some clear advantages: it does not require knowledge about the probability distribution of the uncertain parameters, it keeps optimization problems computationally tractable, and it is the more intuitive choice for problems where the law of large numbers does not apply. For example, the expected outcome of a medical treatment is irrelevant if a patient undergoes the treatment only once.

1.1.3 Contributions

Hints for RO practitioners. RO has mainly been developed in the last fifteen years. The focus has been on the theoretical aspects, while applications are lagging. The practitioner may be reluctant to use RO because the theory is overwhelming or because the advantages are hard to measure. Chapter 3 treats many practical issues, as: (i) How to choose the uncertainty set? (ii) Should the decision rule be a function of the final or the primitive uncertain parameters? (iii) Should the objective be optimized for the worst case? (iv) How to deal with integer adjustable variables? (v) How to deal with equality constraints? (vi) What is the right interpretation of "RO optimizes for the worst case"? (vii) How to compare the robustness characteristics of two solutions? (viii) Is an uncertain objective different from an uncertain constraint?

Robust linear optimization can be solved for any convex uncertainty set. For many convex uncertainty sets, a tractable formulation of the RC of a linear program can not be obtained with the methods from Section 1.1 because the uncertainty set is not conic representable, the concave conjugate of the constraint w.r.t. the uncertain parameter does not have a closed-form expression and can not be expressed as the minimum of a convex optimization problem. Chapter 2 shows that the dual of the RC can be formulated as a convex problem with a finite number

of variables and constraints, and that the KKT vector of this convex problem gives the optimal solution of the original RC. The practical consequences are that robust linear optimization problems can now be solved for any convex uncertainty set, that the derivation of the RC becomes easier, and that the GRC - which is shown to be a special case of a normal RC - becomes tractable for any convex distance function.

Deriving the RC for constraints involving the sum of maxima of linear functions. Many optimization problems contain the sum of maxima of linear functions. Example are the sum of absolute values or the sum of $(\cdot)_+$ functions (which take the nonnegative part of their arguments). These are common in e.g. inventory problems with holding and backlogging costs. In many applications, such constraints are first reformulated as a set of linear constraints, and then the constraints are made robust. The resulting constraints are often more conservative than the RC of the original constraint. For solving the correct RC, approaches from literature are categorized in exact approaches and approximations, several new approaches are proposed, and their relations are explored. By testing all methods on three problems, general recommendations are given to reduce conservatism while keeping the problems computationally tractable.

Deriving the RC for Fractional Programs and observing that their benefit may be small. In a simulation study for three Fractional Programs (FPs), it is shown that the solutions that are optimal for the nominal data can deteriorate substantially due to uncertainty (Chapter 5). Standard RO techniques cannot be applied to mitigate the impact of uncertainty, since the objective of an FP is not concave in the uncertain parameters. RO is therefore extended to FPs. For some problems, the RC can be solved in a single step whereas for general FPs an iterative method should be deployed. For two problems, RO performs slightly better in the worst case than the nominal formulation. However, the RO solutions also deteriorate substantially due to uncertainty. Compared to this deterioration, the difference with the nominal solution is negligible. This shows that some FPs may not benefit from RO.

As a side result, two commonly used methods to solve *deterministic* FPs are shown to be dual to each other.

Smooth Pareto surfaces can be generated by solving a single problem. Multiobjective optimization problems have conflicting objectives for which the trade-off can be visualized with a Pareto surface. For generating this surface, previous methods compute one Pareto optimal solution per optimization run. The facets of the convex hull of Pareto optimal objective vectors provide a piecewise linear inner approximation of the Pareto surface. In Chapter 6, an alternative approach for multi-objective LPs based on ARO is presented, where the optimal solution of

minimizing one objective is a function of the bounds on the other objectives. The solution of a single Semidefinite Program (SDP) provides a polynomial inner approximation, where the degree of the polynomial can be selected by the user. The inner approximation thus has a simple functional description and is differentiable, which has advantages in numerous applications. A potentially useful application would be in treatment planning (see Section 1.2), where there is a trade-off between tumor coverage and sparing of the surrounding organs. Unfortunately, the optimization problems for generating a polynomial inner approximation are too large to be practically solvable.

1.2 HDR brachytherapy for prostate cancer

HDR-BT is a form of radiation therapy in which a high-activity radioactive sealed stepping source is brought inside or in close proximity to the tumor for a limited amount of time. Amongst others, it is used to treat localized prostate cancer. HDR-BT has shown to be effective for prostate cancer both as a boost to external beam radiation therapy (EBRT) and as a monotherapy (Ghilezan, 2012). Normal tissue gets less dose in comparison with EBRT (Georg et al., 2014). Yearly, thousands of patients are treated for prostate cancer.[1]

Figure 1.2 depicts the treatment set-up for HDR-BT. A transrectal ultrasound (TRUS) probe and a template are mounted on a stand. The TRUS probe is positioned inside the rectum and provides real-time data about the anatomy and the positions of the implanted source guides (i.e., catheters). Catheters have been inserted through the template and reach the base of the prostate. The urethra runs from the bladder, through the prostate, to the penis.

When a patient with prostate cancer is selected for HDR-BT, the treatment is fully tailored to his anatomical features. For each patient, the gross tumor volume (GTV) is defined as the entire prostate along with macroscopic spread of the tumor. This volume is typically extended isotropically with a margin of 3 mm to include microscopic spread. Some institutions add an extra margin to account for treatment delivery errors. The resulting volume is called the planning target volume (PTV, Hoskin et al., 2013).

A treatment plan tells how many catheters are used, how these are positioned such that the entire prostate can be adequately irradiated, and for how long the source should dwell at which positions inside the catheters. The goal of the treat-

---

[1]A precise estimate is not available, although some hospitals publish treatment statistics. Memorial Sloan Kettering Cancer Center treats 3,700 prostate cancer patients per year (http://www.mskcc.org/cancer-care/adult/prostate/diagnosis-and-treatment, May 2014).

Figure 1.2 – The treatment set-up. A TRUS probe and a template are mounted on a stand. The TRUS probe is positioned inside the rectum. Catheters have been inserted through the template and reach the base of the prostate. The urethra runs from the bladder, through the prostate, to the penis.

ment planner is to design a treatment plan that delivers a certain dose level to the PTV without simultaneously irradiating the surrounding organs at risk (OARs). Unfortunately, this is impossible since the dose range of the source extends beyond the PTV, though dose levels decrease approximately quadratically in the distance from the source. A treatment plan is therefore a trade-off between PTV coverage and OAR sparing.

The parameter space of a treatment plan is too large to explore by hand, even if the catheter positions and the dwell locations are given. Several mathematical optimization models have therefore been proposed to assist the treatment planner, some of which have found their way into commercial treatment planning software (De Boeck et al., 2014). These models determine a treatment plan based on specifications of the desired dose distribution, and are invaluable in the clinic.

1.2.1 Clinical workflow

To provide an impression of the actual treatment, the clinical workflow at a local hospital is outlined. The HDR-BT dose is delivered in two fractions of 8.5 Gy each within 24 hours as a boost to EBRT. A transperineal Martinez prostate template (Nucletron, an Elekta Company, Veenendaal, the Netherlands) is used to help positioning the catheters (Figure 1.3). The workflow can be divided into three phases: (a) in the pre-plan phase, an intended dose distribution is calculated based on vir-

(a)   **(b)**
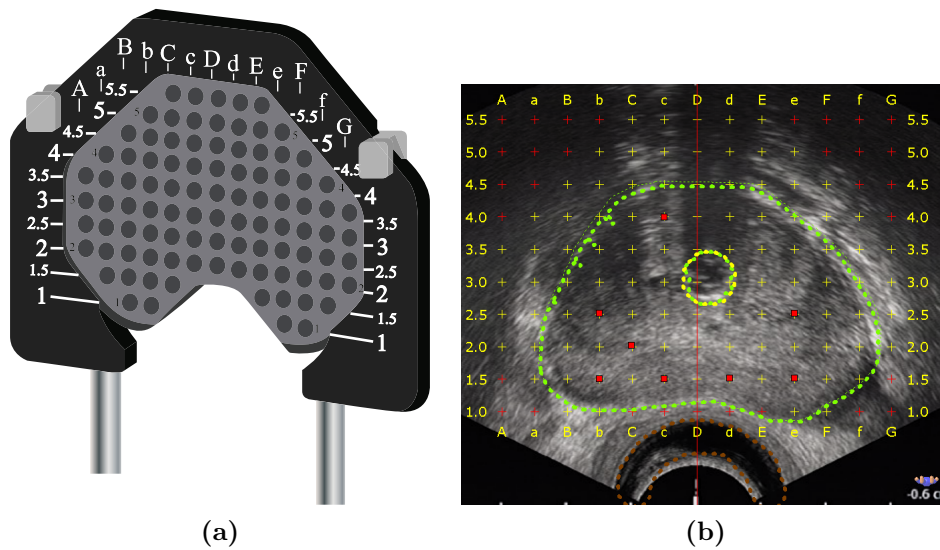
Figure 1.3 – (a) The Martinez prostate template and (b) a virtual model of the same template in the treatment planning system and a transverse view on an ultrasound image showing the superimposed PTV (green), urethra (yellow) and rectum (brown) contours.

tual catheter positons, after which (b) the catheters are implanted in the patient, and (c) a live-plan is created taking the true geometry of the catheters in relation to the patient's anatomy into account, and the optimized intended dose distribution is delivered to the patient. In more detail, the following steps can be distinguished:

1. The patient is positioned in dorsal lithotomy position and transveral images of the volume of interest are acquired using TRUS imaging.

2. The contours of the relevant organs (i.e., prostate, rectum and urethra) are delineated to turn real-life anatomy into digital data.

3. The treatment planning system (HDRplus, version 3.0, Eckert & Ziegler BEBIG GmbH, Berlin, Germany) provides orthogonal, transveral and saggital views of the delineated organs and a virtual model of the template.

4. The treatment planner decides through which template holes to insert catheters such that the catheters get distributed in the prostate volume.

5. The planning system activates dwell positions 3 mm apart within each catheter depending on whether they are inside the PTV or not.

6. A pre-plan is made by optimizing the dwell times at the activated dwell positions to ensure that the given catheter configuration produces a dose distribu-

tion that conforms to the PTV. Changes to the virtual catheter configuration are made if necessary.

7. Fixation catheters are implanted into the prostate under TRUS guidance to fix the position of the prostate relative to the template.

8. Catheters are implanted into the prostate under TRUS guidance according to the pre-plan.

9. After insertion, step 1 is repeated. Moreover, each catheter is reconstructed in the treatment planning software based on its actual position in the longitudinal live TRUS image.

10. Based on the actual catheter positions, the dwell times are re-optimized.

11. A Flexitron afterloader (Nucletron) delivers an iridium-192 radioactive source at the dwell positions for the planned dwell times.

12. The catheters remain inside the patient for delivering the second fraction.

13. After 24 hours, steps 1, 8, 9 and 10 are repeated and the second fraction is delivered.

14. The implant is removed.

This procedure differs between hospitals. Variations exist in the workflow, the treatment planning system, the dose fractionation scheme, the time between fractions and the use of fixation catheters and a template. Most of the results in this dissertation are valid regardless of the exact procedure. The first exception is when no template is used (i.e., for a free-hand implementation procedure), in which case the model for determining the optimal catheter configuration (Chapter 7) becomes inapplicable. The second exception is the simulation of the uncertainty (Chapter 9), which depends on the use of a template and rigid catheters.

1.2.2 Inverse treatment planning

Mathematical optimization can be used in the fourth, sixth and tenth step of the workflow outlined in the previous section. The optimization models described in Sections 1.2.2.1 and 1.2.2.2 are based on dose calculation points, which are artificial points inside the PTV and the OARs that are used to evaluate the dose distribution. Let I be the set of dose calculation points and let J be the set of dwell positions. The dose rate $\dot{d}_{ij}$ (with i in I and j in J) is a parameter indicating the dose per unit time, received by dose calculation point i from a source at dwell position j. The

Figure 1.4 – Dose-based optimization oftenly uses a linear penalty function.

total dose received by dose calculation point i can be expressed as a linear function of the dwell times t_j:

$$d_i = \sum_{j \in J} \dot{d}_{ij} t_j. \tag{1.5}$$

There are several types of dwell time optimization models. Dose-based models have been available for over a decade and have been implemented in commercial treatment planning software (Alterovitz et al., 2006; Karabis et al., 2005; Lessard and Pouliot, 2001). Recently, dose-volume based models have gained attention in the literature (Beliën et al., 2009; Lahanas et al., 2003b; Panchal, 2008; Siauw et al., 2011). These models are computationally more challenging than dose-based models, but also more interesting since they are based on clinically relevant dosimetric evaluation criteria. In this section, both types of models are explained. Other optimization models, such as those based on CVaR (Holm et al., 2013a), fall outside the scope of this introduction since they are not commonly used and not used in this dissertation.

1.2.2.1   Dose-based optimization

In dose-based optimization, the planner prescribes a dose interval $[L_i, U_i]$ for each dose calculation point, along with penalty weights α_i and β_i for underdosage and overdosage, respectively. So, if the dose in dose calculation point i fails to meet the lower bound L_i, the penalty contribution of that dose calculation point is $\alpha_i(L_i - d_i)$. The penalty contribution of a dose calculation point as a function of the dose received by that point is depicted in Figure 1.4. The objective is to minimize the total penalty, which is the sum of the penalty of the individual dose calculation points. The penalty for each voxel is multiplied with a weighting factor w_i. Often $w_i = 1/\{$number of dose

calculation points within the same structure}, to give each structure equal weight. The full objective function is given by:

$$\text{(LD)} \quad \min_{\boldsymbol{t} \geq \boldsymbol{0}} \quad \sum_{i \in I} w_i \max\{0,\, \alpha_i(L_i - \sum_{j \in J} \dot{d}_{ij} t_j),\, \beta_i(\sum_{j \in J} \dot{d}_{ij} t_j - U_i)\}.$$

Problem (LD) can be cast as a Linear Programming problem (LP, see Alterovitz et al., 2006):

$$\min \quad \sum_{i \in I} w_i x_i$$

$$\text{s.t.} \quad x_i \geq \alpha_i[L_i - \sum_{j \in J} \dot{d}_{ij} t_j] \qquad \forall i \in I \tag{1.6}$$

$$x_i \geq \beta_i[\sum_{j \in J} \dot{d}_{ij} t_j - U_i] \qquad \forall i \in I \tag{1.7}$$

$$x_i \geq 0 \qquad \forall i \in I$$

$$t_j \geq 0 \qquad \forall j \in J.$$

The penalty may also be quadratic in the deviation:

$$\min_{\boldsymbol{t} \geq \boldsymbol{0}} \quad \sum_{i \in I} w_i \max\{0,\, \alpha_i(L_i - \sum_{j \in J} \dot{d}_{ij} t_j)^2,\, \beta_i(\sum_{j \in J} \dot{d}_{ij} t_j - U_i)^2\}. \tag{1.8}$$

Sometimes the quadratic penalty formulation is formulated with a single prescribed dose:

$$\text{(QD)} \quad \min_{\boldsymbol{t} \geq \boldsymbol{0}} \quad \sum_{i \in I} w_i(\sum_{j \in J} \dot{d}_{ij} t_j - p_i)^2),$$

where p_i denotes the prescribed dose in dose calculation point i. This formulation has the advantage that the optimization time is independent of the cardinality of I (Lahanas et al., 2003a; Lahanas and Baltas, 2003).

1.2.2.2 Dose-volume based optimization

The (LD) model answers to the idea that a treatment plan is a trade-off between PTV coverage and OAR sparing. A few years after the development of this model in the early 2000s, the European Society for Radiotherapy & Oncology (ESTRO) recommended to report cumulative dose-volume histogram (DVH) statistics (Kovács et al., 2005). A cumulative DVH shows the fraction of the volume receiving at least a certain dose level. Initially, it was not clear for which dose levels or for which volumes these statistics were valuable. Nowadays, DVH statistics have become an important tool to assess the dosimetric quality of treatment plans (Martinez et al., 2001; D'Souza et al., 2004; Butler et al., 2009). The current ESTRO recommendations explicitly mention the constraints on the DVH statistics that should be satisfied

for a particular dose prescription. The correlation between the (LD) objective value and DVH statistics is weak (Holm (2011) and Chapter 7). Inverse treatment planning is therefore a trial and error process where the parameters α_i, β_i, L_i and U_i have to be adjusted until the intended dose distribution is clinically acceptable. Unfortunately, the effect of a change in these parameters on the DVH statistics is hardly predictable. The tuning process would be more intuitive for a small number of weights. For example, for an objective of the form $\gamma_1 f(x) + \gamma_2 g(x)$, the weights γ_1 and γ_2 can be interpreted as "f is improved by Δf if and only if g deteriorates by not more than $\gamma_2/\gamma_1 \cdot \Delta f$". Since (LD) has many weights α_i and β_i in addition to the parameters L_i and U_i, the interpretation of these is more complicated. The shortcomings of the dose-based optimization model have led to the development of dose-volume based models. In the latter, DVH statistics are directly used in the formulation of the optimization problem. In Chapter 7 the following stylistic formulation is used:

$$\text{(LDV)} \quad \max \quad V_{100\%}$$
$$\text{s.t.} \quad D_{10\%}(\text{rectum}) \leq 7.2 \text{ Gy}$$
$$D_{10\%}(\text{urethra}) \leq 10 \text{ Gy},$$

where $V_{100\%}$ denotes the relative volume of the PTV that receives at least the prescribed dose, and $D_{10\%}$ is the minimum dose in the hottest 10% of the OAR volume. (LDV) can be solved using Mixed Integer Programming (MIP, see Chapter 7), Simulated Annealing (SA, see Chapter 8), a combination of LP and SA (Beliën et al., 2009), or by a heuristic based on solving two LPs (Siauw et al., 2011).

Ultimately, a treatment plan should satisfy $D_{90\%} \geq 100\%$ and $V_{100\%} \geq 95\%$ (Hoskin et al., 2013). One may wonder why $V_{100\%}$ is maximized, instead of $D_{90\%}$. While it is possible to optimize on $D_{90\%}$ with similar techniques, the condition $V_{100\%} \geq 95\%$ implies $D_{95\%} \geq 100\%$, and is therefore a stronger condition than $D_{90\%} \geq 100\%$.

1.2.3 Robust Optimization for inverse treatment planning

HDR-BT is affected by many uncertainties (Kirisits et al., 2014). Among all plans that provide good tumor coverage and OAR sparing, an algorithm for automated treatment planning should pick the treatment plan that is least affected by uncertainty. In this section I outline my thoughts and expectations on how this can be developed.

Some uncertainties exist for all radiation therapy treatments, such as delineation errors or changes in the patient anatomy. I focus on uncertainty in the source position, which are specific to HDR-BT. Due to uncertainty in the catheter reconstruction and mechanical source positioning limitations of the afterloader (see Chapter 9), the

radioactive source may dwell at different positions than expected, which affects the delivered dose distribution. Although these errors only comprise a small part of the total error (Kirisits et al., 2014), there are two reasons that justify the focus on uncertainty in the source position. First, for a study on robustness, it is necessary to simulate the errors in order to measure their impact. For simulating changes in patient geometry, no quantitative data is currently available. Second, the error needs to be suitable for Robust Optimization. There need to be multiple near-optimal solutions, with varying sensitivity to uncertainty. For example, delineation errors do not fit this criterion, since it is ultimately the responsibility of the physician what volume to irradiate. The decision where to position the 100% isodose line should not be hidden in a model.

The model (LD) has a sum of maxima of linear functions in the objective function. Since this sum is not concave in the uncertain parameters (i.e., in the dose rates), standard RO techniques cannot be applied. The specialized techniques described in Chapter 4 can in theory be used to robustly optimize (LD). However, the exact methods have prohibitively large solution times and exploratory results show that the approximate methods yield solutions with insufficient target coverage.

Optimization methods that deal with uncertainty require a model that can be used to optimize a given scenario without intervention from the user. So, it should be able to automatically select reasonable treatment plans. This is another reason to disregard the (LD) model, since (LD) indirectly and inaccurately models the planners' preferences (Section 1.2.2.2). Instead, (LDV) is the model to consider. This model makes a clear trade-off between two conflicting goals, namely to maximize the target coverage as long as the OARs are not overly exposed. A clinical question is what would be desired when uncertainty enters the picture. One could keep the OAR sparing only for the nominal (expected) scenario while optimizing the coverage in a range of scenarios. Similarly, one could seek a plan that spares the OARs in the nominal scenario, or in each scenario. I think it is not necessary to have OAR sparing in each scenario for two reaons. First, the limits on the dose in the OARs are based on clinical experience. This experience is based on knowledge of the dose in the nominal scenario and the observations of side effects. So, a plan that satisfies the DVH criteria in the nominal scenario gives limited side effects. Secondly, sometimes the physician accepts a small overdose the the OARs to achieve good tumor coverage. This shows a preference for tumor control at the cost of a small extra risk on side effects.

So, the (LDV) needs to be modified to "optimize the coverage in a range of scenarios". Since for a given treatment plan the coverage is different in each scenario, optimizing a treatment plan is a multiobjective problem for which many methods

exist such as Pareto optimization, maximin optimization, lexicographic ordering and weighted sum optimization (Ehrgott, 2005). I will argue that maximin optimization (i.e., RO) is the most sensible method. Essentially, the physician wants to avoid scenarios in which the tumor survives. Suppose there are only two scenarios and the coverage in scenario 1 is 80% and the coverage in scenario 2 is 92%. The treatment is likely to fail in the first scenario, and unless that scenario is known to be unlikely, it is desired to increase the coverage even if that would reduce the coverage in the second scenario. RO answers to this desire by improving the worst scenario. Let S denote the set of scenarios under consideration and let s^0 be the nominal scenario. A stylistic formulation of the robust (LDV) model is given by:

$$
\begin{aligned}
\text{(R-LDV)} \quad \max \quad & y \\
\text{s.t.} \quad & y \leq V_{100\%}^s \quad \forall s \in S \\
& D_{10\%}^{s^0}(\text{rectum}) \leq 7.2 \text{ Gy} \\
& D_{10\%}^{s^0}(\text{urethra}) \leq 10 \text{ Gy},
\end{aligned}
$$

where the superscript of a DVH statistic indicates the scenario. The challenge in solving this problem is that the constraint $V_{100\%}$ has to be replicated for each scenario. Since the modelling of this constraint involves many dose rate matrices and many binary variables, the number of scenarios has to be low. However, neglecting the afterloader inaccuracy and considering only fifteen catheters and two scenarios per catheter, the number of scenarios is already $2^15 = 32,768$. Iterative methods can be used to solve (R-LDV) while keeping the problem size limited (see Chapter 4). An example of such a method is Algorithm 1. Basically, it starts with a smaller set $\tilde{S}$, which initially only includes s^0. Then in each iteration it solves (R-LDV), fixes the dwell times and determines the scenario with lowest $V_{100\%}$, and adds that scenario to $\tilde{S}$. The seventh step in Algorithm 1 can be solved as a MIP. Alternatively, if the number of scenarios is small, the scenario with lowest $V_{100\%}$ can be found by explicitly computing $V_{100\%}$ for each scenario. During execution, the algorithm provides lower and upper bounds on the optimal value. These bounds are valid only if the subproblems are solved to mathematical optimality, which is unlikely to happen. Nevertheless, the iterative method also works if the subproblems are not optimized to optimality, as long as the value of (R-LDV) is at least the value of LB in the previous iteration, and a scenario is found for which $V_{100\%} > UB$.

An idea to improve the performance of the method is by varying the density of the dose calculation points. The dose is likely to be high around a catheter, so there need to be less dose calculation points in this volume. This reduces the overall memory requirements and improves the solution time of the optimization models at the cost of more preprocessing.

Algorithm 1 Iterative method to solve (R-LDV)

1: $k := 0$

2: $\tilde{S} := \{s^0\}$ (the nominal value)

3: **repeat**

4:     $k := k + 1$

5:     Solve (R-LDV) for $S := \tilde{S}$

6:     Let $\boldsymbol{t^k}$ be the mathematically optimal dwell times and let UB be the value of (R-LDV)

7:     Determine the scenario with lowest $V_{100\%}$ for $\boldsymbol{t^k}$

8:     Let s^k denote this scenario and let $LB = V_{100\%}^{s^k}$

9:     $\tilde{S} := \tilde{S} \cup \{s^k\}$

10: **until** $UB - LB < \varepsilon$

1.2.4 Contributions

Simultaneous optimization of catheter positions and dwell times. Deciding which template holes to use for the catheters is considered computationally difficult due to the combinatorial structure of the problem. Karabis et al. (2009) have tried to solve this problem with a MIP solver and binary variables b_k indicating whether a catheter is inserted through template hole k. This approach was unsuccessful due to large computation times. In Chapter 7 this formulation is improved in two ways. First, constraints are added that forbid using neighboring catheter positions. This does not only speed up the optimization, but also avoids high-dose volumes around catheters to become connected, which is considered clinically desirable. Second, specific solver settings are used. With these changes, (LD) and (QD) can be solved to proven mathematical optimality and good solutions for (LDV) can be obtained.

This approach differs from the one in literature. Holm (2011) proposes a constraint that avoids using more than two catheters in any 2×2 subgrid of the template. This constraint is weaker than the aforementioned one, i.e., is less restrictive and therefore does not provide the same speedup. Holm (2013) proposes a custom branch & bound scheme to deal with problems that contain the binary variables b_k. She shows that for some patients this is better than standard branch & bound in CPLEX, whereas for other patients it is worse. In her comparison, exclusion constraints as proposed by me and by Holm (2011) were not used. Hence, it is not known how the custom branch & bound compares to the method in Section 7. Poulin et al. (2013) and Siauw et al. (2012) optimize catheter positions without simultaneously determining dwell times. Therefore, these algorithms may not choose the catheter positions that allow the lowest objective value for the dwell time optimization problem. Holm et al. (2013b) and Karabis et al. (2005) optimize the cathether positions and

dwell times heuristically, which may not result in a mathematically optimal solution.

Weak relation between (LD) objective and DVH criteria. Holm (2011) has shown that the (LD) objective correlates weakly with DVH statistics. This is confirmed in Chapter 7, where two treatment plans with similar DVH statistics have a factor 12 difference in (LD) objective value. The (LD) objective may therefore be unsuitable for determining reasonable treatment plans that satisfy a given list of constraints on DVH statistics. Further analysis shows that more than 90% of the factor 12 difference is explained by penalizing PTV overdosage, and that half the difference comes from less than 1% of the dose calculation points. These insights may lead to improvements of the (LD) objective.

Direct optimization on DVH criteria. The (LD) objective has two major disadvantages that make it hard to generate treatment plans that satisfy a given list of DVH criteria. First, plans that are mathematically optimal for (LD) may have suboptimal DVH statistics. Second, obtaining better plans by tuning the weights is not intuitive. To overcome these, several heuristics have been proposed to optimize dwell times that directly incorporate DVH statistics. The heuristics by Lahanas et al. (2003b) and Panchal (2008) use DVH statistics in the objective, while the heuristics by Beliën et al. (2009) and Siauw et al. (2011) can have constraints on DVH statistics. In Chapter 7 it is proposed to use a MIP solver with specific settings to directly optimize a model in which the objective and constraints are directly formulated in terms of DVH criteria. Similar to the heuristics, the solver comes up with a good solution in reasonable time but cannot prove that the solution is mathematically optimal (in reasonable time). An advantage of my method compared to all other heuristics is that it can also optimize catheter positions.

In Chapter 8, a new heuristic is presented based on Simulated Annealing as an alternative to existing heuristics for optimizing dwell times. It does not require an expensive solver, and is shown to have advantages in terms of speed and objective value.

Comparison between different optimization models. Previous papers on inverse treatment planning focus on solving a single optimization model (Alterovitz et al., 2006; Holm, 2013; Holm et al., 2013a,b; Karabis et al., 2005; Lessard and Pouliot, 2001; Karabis et al., 2009; Milickovic et al., 2002; Panchal, 2008; Ren et al., 2013; Siauw et al., 2011). The comparison between models is hindered because each paper uses a different data set and many papers employ heuristics whose suboptimality has not been quantified. In Chapter 7 the difference in dose statistics for treatment plans obtained with a commercially available treatment planning system that employs (LD) and good solutions of (LDV) for dwell time optimization are compared. This is also done for combined dwell time and catheter optimization, based

on the mathematically optimal solutions of (LD) and (QD), good solutions of (LDV), and heuristically obtained solutions of (1.8). The comparison shows large differences between optimization models, which induces the risk of treating a patient with a suboptimal plan. Since the live-plan is created intraoperatively, there is little time for Quality Assurance (QA). Recently, a fast QA procedure for optimization was proposed, based on the (QD) model using dose calculation points on the *surface* of each organ (Deufel and Furutani, 2014). Since the (QD) model has user-tunable weights and since cold spots in the PTV cannot be detected without dose calculation points inside the PTV, the (LDV) model would be better suitable for automated QA.

Heuristicical fast optimization of asymmetric quadratic penalty. The (QD) model has the advantage that the optimization time is independent of the number of dose calculation points, but forces the treatment planner to prescribe a single dose level for each voxel. In contrast, the model (1.8) allows the treatment planner to prescribe a lower and an upper bound on the dose for each voxel. This extra flexibility avoids that a penalty is incurred for reasonable dose levels, i.e., for dose levels between the lower and upper bound. Unfortunately, the solution time of (1.8) depends on the number of dose calculation points. Chapter 7 contains a heuristic that optimizes (1.8) by sequentially solving (QD) to combine the flexibility of (1.8) and the speed of (QD). While the heuristic does not always find the global optimum of (1.8), it generates treatment plans that are considered adequate by an expert treatment planner.

DTMR may not improve treatment plan quality and robustness. A dwell time modulation restriction (DTMR) limits the variation of the dwell times within a catheter, and was postulated to improve plan quality (Baltas et al., 2009). In Chapter 9 the effects of a DTMR on the reduction of selective intratumoral high-dose volumes and increase in robustness are quantified. For robustness, catheter reconstruction errors and treatment delivery uncertainties induced by the afterloader have been simulated. Neither was shown to improve by a DMTR. Moreover, a DTMR is a potential cause for underdosage for the (LD) model.

1.3   Overview of research papers

This thesis contains the following eight research papers.

Chapter 2   B. L. Gorissen, A. Ben-Tal, J. P. C. Blanc and D. den Hertog. Deriving robust and globalized robust solutions of uncertain linear programs with general convex uncertainty sets. *Operations Research*, 62(3), 672-679, 2014.

Chapter 3   B. L. Gorissen, İ. Yanıkoğlu and D. den Hertog. Hints for practical robust optimization. *CentER Discussion Paper*, 2013(065), 1–36, 2013.

Chapter 4   B. L. Gorissen and D. den Hertog. Robust counterparts of inequalities containing sums of maxima of linear functions. *European Journal of Operational Research*, 227(1), 30–43, 2013.

Chapter 5   B. L. Gorissen. Robust Fractional Programming. *Journal of Optimization Theory and Applications*, doi:10.1007/s10957-014-0633-4.

Chapter 6   B. L. Gorissen and D. den Hertog. Approximating the Pareto set of multiobjective linear programs via robust optimization. *Operations Research Letters*, 40(5), 319–324, 2012.

Chapter 7   B. L. Gorissen, D. den Hertog and A. L. Hoffmann. Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR brachytherapy. *Physics in Medicine and Biology*, 58(4), 1041–1058, 2013.

Chapter 8   T. M. Deist and B. L. Gorissen. HDR prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing. Submitted.

Chapter 9   M. Balvert, B. L. Gorissen, D. den Hertog and A. L. Hoffmann. Dwell time modulation restrictions do not necessarily improve treatment plan quality for prostate HDR brachytherapy implants. Submitted.

1.4   Disclosure

Each paper contains ideas and contributions from all its respective authors. Chapter 1 has been written by me, except Section 1.1.1 and the contribution of the second research paper, which are based on Sections 3.2 and 3.1, respectively. Chapters 2, 4, 5, 6 and 7 have been written by me, except the introductions of Chapters 6 and 7. In Chapter 3, Sections 3.3, 3.4, 3.10 and 3.11 have been written by me. All calculations and simulations were done by me, except in Sections 3.5 and 3.6 and in Chapter 8.

Deriving robust and globalized robust solutions of uncertain linear programs with general convex uncertainty sets

Abstract We propose a new way to derive tractable robust counterparts of a linear program by using the theory of Beck and Ben-Tal (2009) on the duality between the robust ("pessimistic") primal problem and its "optimistic" dual. First, we obtain a new *convex* reformulation of the dual problem of a robust linear program, and then show how to construct the primal robust solution from the dual optimal solution. Our result allows many new uncertainty regions to be considered. We give examples of tractable uncertainty regions that were previously intractable. The results are illustrated by solving a multi-item newsvendor problem. We also apply the new method to the globalized robust counterpart scheme and show its tractability.

2.1 Introduction

Robust Optimization (RO) is a paradigm for dealing with uncertain data in an optimization problem. Parts of RO originate from the seventies and eighties (Soyster, 1974; Thuente, 1980; Singh, 1982; Kaul et al., 1986), but most of the existing theory and applications followed after new results in the late nineties (Ben-Tal and Nemirovski, 1998; El Ghaoui and Lebret, 1997). An extensive overview of RO is given in (Ben-Tal et al., 2009a) and the survey (Bertsimas et al., 2011). The basic idea of RO is that constraints have to hold for all parameter realizations in some given uncertainty region.

Currently, two tractable methods to solve an RO problem can be distinguished. Both methods are applied constraint-wise, i.e. they reformulate individual constraints. The first method uses conic duality (e.g. used by Ben-Tal et al. (2009a)),

while the second method uses Fenchel duality (Ben-Tal et al., 2014). However, there are still some uncertainty sets for which both methods may not produce explicit tractable robust counterparts (RCs).

The purpose of this chapter is to present a new method for robust linear programs, based on the result "primal worst equals dual best" (Ben-Tal et al., 2009a), which gives tractable optimization problems for general convex uncertainty regions. In Section 2.3, we give examples of uncertainty sets where the new method generates explicit tractable robust counterparts, whereas the classical methods result in intractable RCs.

We also apply the new method to the *globalized robust counterpart* (GRC) model. In this model, there are two convex uncertainty regions, where the constraint must hold for uncertain events in the smaller uncertainty region, while the violation of the constraint for the events in the larger region is controlled via a convex distance function. We show that the GRC can be formulated as well as an ordinary robust linear program with a (different) convex uncertainty region, which implies that it can be solved efficiently with the method presented in this chapter.

2.2 A method for deriving a tractable dual of the Robust Counterpart

Consider the following Linear Conic Program (LCP):

$$\text{(LCP)} \quad \max_{\boldsymbol{x} \in \mathcal{K}} \{ \boldsymbol{c}^\top \boldsymbol{x} : \boldsymbol{a}_i^\top \boldsymbol{x} \le b_i \quad \forall i \in I \},$$

where $\mathcal{K}$ is a closed convex cone and I is a finite index set. If $\mathcal{K}$ is the nonnegative orthant $\mathbb{R}_+^n$, (LCP) is an LP in canonical form. Other common choices for $\mathcal{K}$ are the second-order cone and the semidefinite cone. This setting is more general than the title of this chapter indicates, but we will only allow uncertainty in the linear constraints and want to avoid confusion with robust conic optimization. We will use the prefix D for the dual, O for the optimistic counterpart, and R for the robust counterpart. The dual of (LCP) is given by:

$$\text{(D-LCP)} \quad \min_{\boldsymbol{y} \ge \boldsymbol{0}} \{ \boldsymbol{b}^\top \boldsymbol{y} : \sum_{i \in I} y_i \boldsymbol{a}_i - \boldsymbol{c} \in \mathcal{K}^* \},$$

where $\mathcal{K}^*$ is the dual cone of $\mathcal{K}$ and I is a finite index set. Assume that $\boldsymbol{a}_i$ are uncertain, but known to reside in some convex compact uncertainty region $\mathcal{U}_i = \{ \boldsymbol{a}_i : f_{ik}(\boldsymbol{a}_i) \le 0 \quad \forall k \in K \}$, where f_{ik} are given closed proper convex functions and

K is a finite index set. The robust counterpart of (LCP) is given by:

$$(\text{R-LCP}) \quad \max_{\boldsymbol{x} \in \mathcal{K}} \quad \boldsymbol{c}^\top \boldsymbol{x}$$

$$\text{s.t.} \quad \boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i \qquad \forall \boldsymbol{a}_i : f_{ik}(\boldsymbol{a}_i) \leq 0 \quad \forall i \in I \quad \forall k \in K, \qquad (2.1)$$

(R-LCP) is a semi-infinite convex optimization problem since it has a linear objective, $\boldsymbol{x}$ is finite dimensional, and the feasible region is the intersection of infinitely many half spaces. Since (R-LCP) is convex, a local optimum is also a global optimum. Numerical methods that find this optimum cannot be applied because of the semi-infinite representation. Current RO rewrites constraint (2.1) to a finite set of constraints, which works for a limited set of functions f_{ik}. The optimistic counterpart of (D-LCP) is given by:

$$(\text{OD-LCP}) \quad \min_{\boldsymbol{y} \geq \boldsymbol{0}} \{ \boldsymbol{b}^\top \boldsymbol{y} : \ \forall i \in I \ \exists \boldsymbol{a}_i, \ f_{ik}(\boldsymbol{a}_i) \leq 0, \ \forall k \in K, \ \sum_{i \in I} y_i \boldsymbol{a}_i - \boldsymbol{c} \in \mathcal{K}^* \}.$$

A result by Beck and Ben-Tal (2009) is that (OD-LCP), which is optimistic since it has to hold for a single $\boldsymbol{a}_i$, is a dual problem of (R-LCP), which has to hold for all $\boldsymbol{a}_i$. Moreover, the values of (R-LCP) and (OD-LCP) are equal if (OD-LCP) is bounded and satisfies the Slater condition. Less general but similar results can be found in (Falk, 1976; Römer, 2010; Soyster, 1974; Thuente, 1980). For $\mathcal{K} = \mathbb{R}^n_+$, (OD-LCP) is called a Generalized LP (GLP) (Dantzig, 1963, p. 434). It contains the product of variables $y_i \boldsymbol{a}_i$ and is in general nonconvex. However, it will be shown that (OD-LCP) can be converted to a convex program.

Dantzig mentions substituting $\boldsymbol{v}_i = y_i \boldsymbol{a}_i$ and multiplying the constraints containing f_{ik} with y_i as a solution approach to GLPs (Dantzig, 1963, p. 434), which has already been applied to the dual of LPs with polyhedral uncertainty by Römer (2010). When we apply this to (OD-LCP), we get the following convex optimization problem:

$$(\text{COD-LCP}) \quad \min_{\boldsymbol{y} \geq \boldsymbol{0}, \boldsymbol{v}_i} \quad \boldsymbol{b}^\top \boldsymbol{y}$$

$$\text{s.t.} \quad \sum_{i \in I} \boldsymbol{v}_i - \boldsymbol{c} \in \mathcal{K}^* \qquad (2.2)$$

$$y_i f_{ik} \left(\frac{\boldsymbol{v}_i}{y_i} \right) \leq 0, \qquad \forall i \in I \quad \forall k \in K, \qquad (2.3)$$

where $0 f_{ik}(\boldsymbol{v}_i/0) = \lim_{y_i \downarrow 0} y_i f_{ik}(\boldsymbol{v}_i/y_i)$ is the recession function of f. (COD-LCP) is indeed a convex problem, since the perspective function $g_{ik}(\boldsymbol{v}_i, y_i) := y_i f_{ik}(\boldsymbol{v}_i/y_i)$ is convex on $\mathbb{R}^n \times \mathbb{R}_+$. Here is a short proof of the convexity of $y_i f_{ik}(\boldsymbol{v}_i/y_i)$ on

$\mathbb{R}^n \times \mathbb{R}_+ \backslash \{0\}$ that uses convex analysis:

$$g_{ik}(\boldsymbol{v}_i, y_i) = y_i f_{ik}\left(\frac{\boldsymbol{v}_i}{y_i}\right) = y_i f_{ik}^{**}\left(\frac{\boldsymbol{v}_i}{y_i}\right) = y_i \sup_{\boldsymbol{x}}\left\{\frac{\boldsymbol{v}_i^\top}{y_i}\,\boldsymbol{x} - f_{ik}^*(\boldsymbol{x})\right\}$$
$$= \sup_{\boldsymbol{x}}\left\{\boldsymbol{v}_i^\top\,\boldsymbol{x} - y_i f_{ik}^*(\boldsymbol{x})\right\},$$

from which it follows that g_{ik} is jointly convex since it is the pointwise supremum of functions which are linear in $\boldsymbol{v}_i$ and y_i.

While (R-LCP) is difficult to solve because it has an infinite number of constraints, (COD-LCP) does not have "for all" constraints. For some popular choices of f_{ik} for which an exact reformulation of (R-LCP) is known, (COD-LCP) is at most as difficult to solve as (R-LCP). For instance, when the uncertainty region is polyhedral, (R-LCP) can be reformulated as an LP, and (COD-LCP) is also an LP. When the uncertainty region is an ellipsoid, (R-LCP) can be reformulated as a conic quadratic program, and (COD-LCP) is also a conic quadratic program.

Dantzig notes that (OD-LCP) and (COD-LCP) are equivalent only when $\boldsymbol{v}_i \neq \boldsymbol{0}$ is not possible if $y_i = 0$, since otherwise $\boldsymbol{v}_i = y_i \boldsymbol{a}_i$ does not hold. We call this the *substitution equivalence condition*. The following lemma shows that this condition is automatically satisfied:

Lemma 1 *Assume that the uncertainty region is bounded. Then* (2.3) *enforces the substitution equivalence condition.*

Proof. Let $i \in I$ and let $y_i = 0$. From the definition of $0 f_{ik}(\boldsymbol{v}_i/0)$, it is clear that $\boldsymbol{v}_i = \boldsymbol{0}$ is feasible for (2.3). It remains to show that a nonzero $\boldsymbol{v}_i^*$ is infeasible. Assume to the contrary that $\boldsymbol{v}_i^* \neq \boldsymbol{0}$ is feasible. Let us first construct two points where $g_{ik}(\boldsymbol{v}_i, y_i) \leq 0$.

The first point is $(c\boldsymbol{v}_i^*, 0)$, $c > 0$, and the second is $(2a_i, 2)$ where $a_i \in \mathcal{U}_i$. Indeed, for $k \in K$, $g_{ik}(c\boldsymbol{v}_i^*, 0) = \lim_{y_i \downarrow 0} y_i f_{ik}(c\boldsymbol{v}_i^*/y_i) = c \lim_{y_i \downarrow 0}(y_i/c) f_{ik}(\boldsymbol{v}_i^*/(y_i/c)) = c g_{ik}(\boldsymbol{v}_i^*, 0) \leq 0$ and clearly also $g_{ik}(2\boldsymbol{a}_i, 2) \leq 0$ for an arbitrary $\boldsymbol{a}_i \in \mathcal{U}_i$.

By convexity, we have $g_{ik}\left(\lambda\boldsymbol{v}_i^1 + (1-\lambda)\boldsymbol{v}_i^2, \lambda y_i^1 + (1-\lambda)y_i^2\right) \leq \lambda g_{ik}\left(\boldsymbol{v}_i^1, y_i^1\right) + (1-\lambda)g_{ik}\left(\boldsymbol{v}_i^2, y_i^2\right)$. In particular, for $\lambda = 0.5$ and the above two points, we get $g_{ik}((c\boldsymbol{v}_i^* + 2\boldsymbol{a}_i)/2, 1) \leq 0$ for all $k \in K$. This implies that $(c/2)\boldsymbol{v}_i^* + \boldsymbol{a}_i$ is in $\mathcal{U}_i$ for all $c > 0$. So, the uncertainty region recedes in the direction of $\boldsymbol{v}_i^*$, contradicting boundedness. ∎

In practice it is often necessary to have the primal robust solution $\boldsymbol{x}$ of (R-LCP), instead of the solution of (COD-LCP). The following theorem shows how $\boldsymbol{x}$ can be recovered from an optimal solution of (COD-LCP).

Theorem 1 *Assume that (COD-LCP) is bounded and satisfies the Slater condition. A KKT vector of constraint* (2.2) *corresponds to an optimal solution $\boldsymbol{x}$ of (R-LCP).*

Proof. First, we show that the dual variables associated with constraint (2.2) are the optimization variables of (R-LCP). The Lagrangian of (COD-LCP) is given by:

$$
L(\boldsymbol{y}, v, \boldsymbol{x}) = \begin{cases} \boldsymbol{b}^\top \boldsymbol{y} + \boldsymbol{x}^\top \left(\boldsymbol{c} - \sum_{i \in I} \boldsymbol{v}_i\right) & \text{if } y_i f_{ik}\left(\frac{\boldsymbol{v}_i}{y_i}\right) \leq 0, \ \forall i \in I \ \forall k \in K \\ \infty & \text{otherwise,} \end{cases}
$$

and hence, (R-LCP) is given by:

$$
\begin{aligned}
& \max_{\boldsymbol{x} \in \mathcal{K}} \quad \min_{\boldsymbol{y} \geq \boldsymbol{0}, v} L(\boldsymbol{y}, v, \boldsymbol{x}) \\
= \ & \max_{\boldsymbol{x} \in \mathcal{K}} \left\{ \boldsymbol{c}^\top \boldsymbol{x} + \min_{\boldsymbol{y} \geq \boldsymbol{0}, v} \left\{ \boldsymbol{b}^\top \boldsymbol{y} - \sum_{i \in I} \boldsymbol{v}_i^\top \boldsymbol{x} : y_i f_{ik}\left(\frac{\boldsymbol{v}_i}{y_i}\right) \leq 0, \quad \forall i \in I \ \forall k \in K \right\} \right\} \\
= \ & \max_{\boldsymbol{x} \in \mathcal{K}} \left\{ \boldsymbol{c}^\top \boldsymbol{x} + \min_{\boldsymbol{y} \geq \boldsymbol{0}, a} \left\{ \sum_{i \in I} y_i \left(b_i - \boldsymbol{a}_i^\top \boldsymbol{x}\right) : f_{ik}(\boldsymbol{a}_i) \leq 0 \quad \forall i \in I \ \forall k \in K \right\} \right\} \\
= \ & \max_{\boldsymbol{x} \in \mathcal{K}} \left\{ \boldsymbol{c}^\top \boldsymbol{x} : \boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i \quad \forall \boldsymbol{a}_i : f_{ik}(\boldsymbol{a}_i) \leq 0, \quad \forall i \in I \ \forall k \in K \right\},
\end{aligned}
$$

where in the second equality the substitution $\boldsymbol{a}_i = \boldsymbol{v}_i / y_i$ is made. If an optimal value is attained at $y_i = 0$ (and consequently at $\boldsymbol{v}_i = \boldsymbol{0}$) before the substitution, then any feasible $\boldsymbol{a}_i$ is optimal after the substitution. The problem in the last equality is indeed (R-LCP).

Since (COD-LCP) is bounded and satisfies the Slater condition, a KKT vector exists (Rockafellar, 1970, Theorem 28.2). An optimal $\boldsymbol{x}$ of (R-LCP) is equal to a KKT vector (Rockafellar, 1970, Corollary 28.4.1). ∎

This theorem is useful in practice because many solvers can output a KKT vector. There is also another way to obtain a solution of (R-LCP), similar to the method mentioned used by Soyster (1974). The idea is to use the "dual best $\boldsymbol{a}_i$" as the "primal worst $\boldsymbol{a}_i$": translate a solution of (COD-LCP) to a solution of (OD-LCP), then fix the variables $\boldsymbol{a}_i$, remove the constraints on $\boldsymbol{a}_i$, and dualize that problem with respect to y_i. The result is a problem similar to (LCP), where the vectors $\boldsymbol{a}_i$ have been replaced with "worst case" $\boldsymbol{a}_i$. We call this problem (M-LCP). This method only works if (COD-LCP) has a unique optimal $\boldsymbol{a}_i$, and if (M-LCP) has a unique optimal $\boldsymbol{x}$ (Soyster, 1974). Then, the value of (M-LCP) equals the value of (OD-LCP), and $\boldsymbol{x}$ is both feasible and optimal for (R-LCP).

The main advantage of our method is the fact that we use the perspectives of the original functions that define the uncertainty region, while existing methods in RO use the perspective of the conjugate functions to reformulate the "for all" constraints in (R-LCP) (Ben-Tal et al., 2014); the latter may not result in closed-form formulations for many uncertainty regions. We give examples of the above in Section 2.3.

There may be an additional computational advantage as the number of variables and constraints in (COD-LCP) can be smaller, compared to results obtained with existing RO techniques. For example, the latter may require an explicit conic representation (e.g. see (Ben-Tal et al., 2009a, Theorem 1.3.4)) which can significantly increase the number of variables and constraints.

Our method also has two disadvantages that is inherent to solving the *dual* problem. First, it cannot directly be applied to problems with integer variables (however, it can be used to solve LP relaxations such as those needed in cutting plane and branch & bound methods). Second, the primal solution has to be recovered from the KKT vector. This means that the dual problem has to be solved to high accuracy for otherwise the accuracy of the *primal* solution may suffer. However, the actual effect on the optimal primal objective function can be assessed by the duality gap. We show this in our numerical example in Section 2.5.

2.3   New tractable uncertainty regions

In this section we present three examples of uncertainty regions for which the robust counterpart cannot be obtained using the traditional approach:

1. The first example is given by problems in which several scenarios for the parameters can be distinguished, but the probabilities on these scenarios are not known. Suppose these unknown probabilities can be estimated based on historical data, and an optimization problem has a constraint involving these probabilities. An example of this is a constraint on expected value. For such problems, a wide class of uncertainty regions is given in terms of the distance between the real probability vector $\boldsymbol{p}$ and a historical estimate $\hat{\boldsymbol{p}}$, both indexed by the scenario s from a finite scenario set S:

$$\mathcal{U}_i = \left\{ \boldsymbol{p} : \boldsymbol{p} \geq \boldsymbol{0}, \sum_{s \in S} p^{(s)} = 1, d(\boldsymbol{p}, \hat{\boldsymbol{p}}) \leq \rho \right\}, \tag{2.4}$$

   where d is the distance measure and ρ is the level of uncertainty. Note that the constraint $\sum_{s \in S} p^{(s)} = 1$ is not necessary for the following results to hold, so $\boldsymbol{p}$ does not need to be a probability vector. We consider several classes of distance measures:

   (a) The first class of distance measures that contains previously intractable cases is ϕ–divergence, which for a convex function ϕ that satisfies $\phi(1) = 0$ is defined by $d(\boldsymbol{p}, \hat{\boldsymbol{p}}) = \sum_{s \in S} \hat{p}^{(s)} \phi \left(p^{(s)} / \hat{p}^{(s)} \right)$. Ben-Tal et al. (2013) show how to choose ρ in (2.4) based on historical observations, and give tractable

robust counterparts for several choices of ϕ. One example for which their method does not give a tractable reformulation is the Matusita distance (Matusita, 1967), where $\phi(t) = |1 - t^\alpha|^{1/\alpha}$ for given $\alpha \in (0, 1)$.

(b) The second class of distance measures is based on the Bregman distance, which is given by (Bregman, 1967):

$$d(\boldsymbol{p}, \hat{\boldsymbol{p}}) = g(\boldsymbol{p}) - g(\hat{\boldsymbol{p}}) - (\nabla g(\hat{\boldsymbol{p}}))^\top (\boldsymbol{p} - \hat{\boldsymbol{p}}),$$

where g is real-valued, continuously-differentiable and strictly convex on the set of probability vectors. The Bregman distance is convex in its first argument. Previously, uncertainty regions were intractable for many choices of g, while with our results any g gives a tractable optimistic counterpart.

(c) The third class of distance measures is the Rényi divergence (Rényi, 1961):

$$d(\boldsymbol{p}, \hat{\boldsymbol{p}}) = \frac{1}{\alpha - 1} \log \sum_{i \in I} \left(\hat{p}_i^{(s)} \right)^\alpha \left(p_i^{(s)} \right)^{1-\alpha},$$

where $\alpha > 0$ and $\alpha \neq 1$. After some rewriting, an uncertainty region based on this distance measure can also be reformulated using Fenchel duality (Ben-Tal et al., 2014). However, the rewriting is not always possible, e.g. when this divergence measure is clustered with other distance measures (Banerjee et al., 2005), while our result can then still be applied.

2. The second example of new tractable uncertainty regions is when the uncertainty region contains an uncertain parameter $\boldsymbol{B}_{ij}$:

$$\mathcal{U}_i = \left\{ (\boldsymbol{a}_i, \boldsymbol{\zeta}_i) : \exists \boldsymbol{B}_{ij}, \ \boldsymbol{\zeta}_i \geq \boldsymbol{0}, \ \boldsymbol{a}_i = \boldsymbol{a}_i^0 + \sum_j \zeta_{ij} \boldsymbol{B}_{ij}, \ g_{ijk}(\boldsymbol{B}_{ij}) \leq 0, \right.$$
$$\left. h_{ik}(\boldsymbol{\zeta}_i) \leq 0, \quad \forall j \in J, k \in K \right\},$$

where g_{ijk} and h_{ik} are convex functions and $\boldsymbol{B}_{ij}, \boldsymbol{a}_i^0$ are vectors. The same substitution we applied to (OD-CLP), can also be applied to this uncertainty region. Let $\boldsymbol{v}_{ij} = \zeta_{ij} \boldsymbol{B}_{ij}$. The uncertainty region $\mathcal{U}_i$ can be rewritten as:

$$\mathcal{U}_i = \left\{ (\boldsymbol{a}_i, \boldsymbol{\zeta}_i) : \exists \boldsymbol{v}_{ij}, \ \boldsymbol{\zeta}_i > \boldsymbol{0}, \ \boldsymbol{a}_i = \boldsymbol{a}_i^0 + \sum_j \boldsymbol{v}_{ij}, \right.$$
$$\left. \zeta_{ij} g_{ijk}(\boldsymbol{v}_{ij}/\zeta_{ij}) \leq 0, \ h_{ik}(\boldsymbol{\zeta}_i) \leq 0, \quad \forall j \in J, k \in K \right\},$$

which is convex, and hence, leads to a tractable optimistic counterpart. We mention three cases where this uncertainty region appears:

(a) First, it appears in factor models where the parameter $\boldsymbol{a}_i$ is estimated as $\boldsymbol{a}_i^0 + \sum_j \zeta_{ij} \boldsymbol{B}_{ij}$ with uncertainty in both the factors $\boldsymbol{\zeta}_i$ and the coefficients $\boldsymbol{B}_{ij}$. For example in finance, the return of an asset can be approximated by $\mu + V^\top f$, where f are the factors that drive the market (Goldfarb and Iyengar, 2003).

(b) Second, it appears in a constraint containing the steady-state distributions of a Markov chain, where the transition probabilities are uncertain. The uncertainty region then looks as follows:

$$\mathcal{U}_i = \left\{ \boldsymbol{\pi} \in \mathbb{R}_+^n : e^\top \boldsymbol{\pi} = 1, \sum_{j \in J} \pi_j \boldsymbol{B}_j = \boldsymbol{\pi}, \ g_{jk}(\boldsymbol{B}_j) \le 0, \quad \forall j \in J, k \in K \right\},$$

where $\boldsymbol{B}_j$ are the columns of the matrix with transition probabilities. Markov chains with column-wise uncertainty in the transition matrix were also considered by Blanc and den Hertog (2008).

(c) Third, it appears in a constraint on the next time period probability vector $\boldsymbol{p}_i$ of a Markov chain when there is uncertainty both in the transition matrix and in the current state:

$$\mathcal{U}_i = \left\{ \boldsymbol{p}_i \in \mathbb{R}_+^n : \boldsymbol{p}_i = \sum_{j \in J} (\boldsymbol{p}_i^0)_j \boldsymbol{B}_j, \ g_{jk}(\boldsymbol{B}_j) \le 0, \ h_k(\boldsymbol{p}_i^0) \le 0, \quad \forall j \in J, k \in K \right\},$$

where $\boldsymbol{B}_j$ are the columns of the transition matrix and $\boldsymbol{p}_i^0$ is the current probability vector.

3. The third example of new tractable uncertainty regions is illustrated by the following robust constraint:

$$\boldsymbol{a}_i^\top \boldsymbol{x} + \sum_{j \in J} h_{ij}(a_{ij}) x_j \le b_i \qquad \forall \boldsymbol{a}_i : ||\boldsymbol{a}_i||_\infty \le 1,$$

where the functions h_{ij} are convex. For many choices of h_{ij} this constraint is not tractable. To show that it can be solved with our method, we first move the nonlinearity to the uncertainty region:

$$\boldsymbol{a}_i^\top \boldsymbol{x} + \sum_{j \in J} d_{ij} x_j \le b_i \qquad \forall (\boldsymbol{a}_i, \boldsymbol{d}_i) : ||\boldsymbol{a}_i||_\infty \le 1, \quad h_{ij}(a_{ij}) = d_{ij} \quad \forall j \in J,$$

and then obtain an equivalent constraint by taking the convex hull of the uncertainty region:

$$\boldsymbol{a}_i^\top \boldsymbol{x} + \sum_{j \in J} d_{ij} x_j \le b_i \qquad \forall (\boldsymbol{a}_i, \boldsymbol{d}_i) : ||\boldsymbol{a}_i||_\infty \le 1, \quad h_{ij}(a_{ij}) \le d_{ij},$$

$$2 d_{ij} \le h_{ij}(1)(a_{ij} + 1) - h_{ij}(-1)(a_{ij} - 1) \quad \forall j \in J.$$

This transformation has also been applied by Ben-Tal et al. (2014, p. 20), but they require a closed form for the convex conjugate of h_{ij} to reformulate this constraint. With our method, this linear constraint with a convex uncertainty region is tractable for any convex h_{ij}.

2.4 Globalized Robust Counterpart

A robust constraint holds for all realizations of the uncertain parameters in the uncertainty region. If this region contains all "physically possible" events, the RC may result in a very pessimistic solution. Ben-Tal et al. (2006) proposes the *globalized robust counterpart* (GRC) model, which may be used to reduce the conservatism of the RC (also see (Ben-Tal et al., 2009a, Chapter 3)). Here we use a slightly modified version. Let $\mathcal{U}_i = \{\boldsymbol{a}_i : f_{ik}(\boldsymbol{a}_i) \leq 0, \quad \forall k \in K\}$ be the set of "physically possible" realizations, and let a smaller set $\mathcal{U}_i' = \{\boldsymbol{a}_i : g_{ik}(\boldsymbol{a}_i) \leq 0, \quad \forall k \in K\} \subset \mathcal{U}_i$ contain the "normal range" of realizations. We define the GRC as:

$$\boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i + \min_{\boldsymbol{a}_i' \in \mathcal{U}_i'}\{h_i(\boldsymbol{a}_i, \boldsymbol{a}_i')\} \qquad \forall \boldsymbol{a}_i \in \mathcal{U}_i, \tag{2.5}$$

where h_i is a nonnegative jointly convex distance-like function for which $h_i(\boldsymbol{a}_i', \boldsymbol{a}_i') = 0$ for all $\boldsymbol{a}_i'$ in $\mathcal{U}_i'$. Examples are norms and ϕ–divergence measures. The second term on the right-hand side of (2.5) expresses the allowable violation of the constraint and is equal to 0 if $\boldsymbol{a}_i$ is in the smaller set $\mathcal{U}_i'$.

We now show that (2.5) can be reformulated as a linear constraint with a convex uncertainty region. Constraint (2.5) is equivalent to:

$$\boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i + d_i \qquad \forall (\boldsymbol{a}_i, d_i) : f_{ik}(\boldsymbol{a}_i) \leq 0 \quad \forall k \in K, \quad d_i = \min_{\boldsymbol{a}_i' \in \mathcal{U}_i'}\{h_i(\boldsymbol{a}_i, \boldsymbol{a}_i')\},$$

which in turn is equivalent to:

$$\boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i + d_i \qquad \forall (\boldsymbol{a}_i, \boldsymbol{a}_i', d_i) : f_{ik}(\boldsymbol{a}_i) \leq 0 \quad g_{ik}(\boldsymbol{a}_i') \leq 0 \, \forall k \in K, \quad d_i \geq h_i(\boldsymbol{a}_i, \boldsymbol{a}_i').$$

This is indeed a linear constraint with a convex uncertainty region. We will now show how the GRC can be formulated. Consider the following Globalized Robust program:

$$\begin{aligned}
\text{(R-GRC)} \quad \max_{\boldsymbol{x} \in \mathcal{K}} \quad & \boldsymbol{c}^\top \boldsymbol{x} \\
\text{s.t.} \quad & \boldsymbol{a}_i^\top \boldsymbol{x} \leq b_i + d_i \quad \forall (\boldsymbol{a}_i, \boldsymbol{a}_i', d_i) : f_{ik}(\boldsymbol{a}_i) \leq 0, \quad \forall k \in K, \\
& \qquad\qquad g_{ik}(\boldsymbol{a}_i') \leq 0, \quad \forall k \in K \quad d_i \geq h_i(\boldsymbol{a}_i, \boldsymbol{a}_i') \quad \forall i \in I,
\end{aligned}$$

whose optimistic dual is:

$$(\text{OD-GRC}) \quad \min_{\boldsymbol{y} \geq \boldsymbol{0}, \boldsymbol{a}_i, \boldsymbol{a}'_i, \boldsymbol{d}} \quad \boldsymbol{b}^\top \boldsymbol{y} + \boldsymbol{d}^\top \boldsymbol{y}$$

$$\text{s.t.} \quad \sum_{i \in I} y_i \boldsymbol{a}_i - \boldsymbol{c} \in \mathcal{K}^*$$

$$f_{ik}(\boldsymbol{a}_i) \leq 0 \qquad \forall i \in I \quad \forall k \in K \tag{2.6}$$

$$g_{ik}(\boldsymbol{a}'_i) \leq 0 \qquad \forall i \in I \quad \forall k \in K \tag{2.7}$$

$$h_i(\boldsymbol{a}_i, \boldsymbol{a}'_i) \leq d_i \qquad \forall i \in I. \tag{2.8}$$

We substitute $\boldsymbol{v}_i = y_i \boldsymbol{a}_i$, $\boldsymbol{v}'_i = y_i \boldsymbol{a}'_i$ and $w_i = y_i d_i$ and multiply constraints (2.6)–(2.8) with y_i:

$$(\text{COD-GRC}) \quad \min_{\boldsymbol{y} \geq \boldsymbol{0}, \boldsymbol{v}_i, \boldsymbol{v}'_i, \boldsymbol{w}} \quad \boldsymbol{b}^\top \boldsymbol{y} + \sum_{i \in I} w_i$$

$$\text{s.t.} \quad \sum_{i \in I} \boldsymbol{v}_i - \boldsymbol{c} \in \mathcal{K}^*$$

$$y_i f_{ik}\left(\frac{\boldsymbol{v}_i}{y_i}\right) \leq 0 \qquad \forall i \in I \quad \forall k \in K$$

$$y_i g_{ik}\left(\frac{\boldsymbol{v}'_i}{y_i}\right) \leq 0 \qquad \forall i \in I \quad \forall k \in K$$

$$y_i h_i\left(\frac{\boldsymbol{v}_i}{y_i}, \frac{\boldsymbol{v}_i'}{y_i}\right) \leq w_i \qquad \forall i \in I. \tag{2.9}$$

Note that the product $y_i \boldsymbol{a}'_i$ does not appear in (OD-GRC), but that the substitution $\boldsymbol{v}'_i = y_i \boldsymbol{a}'_i$ is still necessary to make (COD-GRC) convex. From the objective and (2.9) it is clear that the substitution equivalence condition holds for w_i even though the uncertainty region is not bounded in $\boldsymbol{d}$. For the tractability of (COD-GRC), all results regarding the functions that define the uncertainty region $\mathcal{U}_i$ also apply to the functions that define $\mathcal{U}'_i$. For example, when h_i is a ϕ–divergence measure, constraint (2.8) is given by:

$$\sum_{j \in J} a_{ij} \phi\left(a'_{ij} / a_{ij}\right) \leq d_i.$$

Constraint (2.9) then contains the perspective of ϕ, which is just:

$$\sum_{j \in J} v_{ij} \phi\left(v'_{ij} / v_{ij}\right) \leq w_i.$$

When in (2.8), h_i is a norm, i.e., $||\boldsymbol{a}_i - \boldsymbol{a}'_i||$, constraint (2.9) contains the same norm: $||\boldsymbol{v}_i - \boldsymbol{v}'_i|| \leq w_i$.

2.5 Multi-item newsvendor example

We demonstrate our new method on a robust LP with a convex uncertainty region that currently cannot be solved with other methods. To obtain this problem, we take a slightly modified version of the multi-item newsvendor problem described by Ben-Tal et al. (2013). There are 12 items indexed by i, and each has its own ordering cost c_i, selling price v_i, salvage price r_i, and unsatisfied demand loss l_i. So, when the order quantity Q_i is less than the demand d_i, the profit equals $v_i Q_i + l_i (Q_i - d_i) - c_i Q_i$, and $v_i d_i + r_i (Q_i - d_i) - c_i Q_i$ otherwise. If $r_i \leq v_i + l_i$, this profit is concave piecewise linear in the decision variable Q_i. In practice the demand is not known, but for every item i we can define scenarios s in a scenario set S which occur with probability $p_i^{(s)}$, independently of other items, resulting in a demand of $d_i^{(s)}$. The goal is to determine Q_i such that the total ordering cost is minimized under the constraint that the expected profit is at least γ. This can be formulated as a robust LP as follows:

$$
\begin{aligned}
\text{(R-NV)} \quad \min_{\boldsymbol{Q} \geq \boldsymbol{0}, \boldsymbol{u}} \quad & \sum_{i \in \mathcal{I}} c_i Q_i \\
\text{s.t.} \quad & \sum_{i \in \mathcal{I}} \sum_{s \in S} p_i^{(s)} u_i^{(s)} \geq \gamma \\
& \forall \boldsymbol{p}_i \geq \boldsymbol{0} : \sum_{s \in S} p_i^{(s)} = 1, \sum_{s \in S} \left| \left(\hat{p}_i^{(s)} \right)^\alpha - \left(p_i^{(s)} \right)^\alpha \right|^{1/\alpha} \leq \rho, \quad \forall i \in \mathcal{I} \\
& u_i^{(s)} + (c_i - r_i) Q_i \leq d_i^{(s)} (v_i - r_i) \qquad \forall i \in \mathcal{I} \quad \forall s \in S \\
& u_i^{(s)} + (c_i - v_i - l_i) Q_i \leq -d_i^{(s)} l_i \qquad \forall i \in \mathcal{I} \quad \forall s \in S,
\end{aligned}
$$

(2.10)

where $u_i^{(s)}$ denotes the profit for item i in scenario s, and the uncertainty region is based on the Matusita distance with α in $(0, 1)$. (R-NV) cannot be solved with current methods, because the uncertainty region in (2.10) does not have a conic quadratic representation and (2.10) cannot be reformulated to finitely many constraints with closed form functions using Fenchel duality to the best of our knowledge (Ben-Tal

et al., 2014, Table 1). The optimistic dual of (R-NV) is given by:

$$(\text{OD-NV}) \quad \max_{x \geq 0, \boldsymbol{y} \leq \boldsymbol{0}, \boldsymbol{z} \leq \boldsymbol{0}, p \geq 0} \quad \gamma x + \sum_{i \in \mathcal{I}, s \in S} d_i^{(s)} \left((v_i - r_i) \, y_{is} - l_i z_{is} \right)$$

$$\text{s.t.} \quad p_i^{(s)} x + y_{is} + z_{is} = 0 \qquad \forall i \in \mathcal{I} \quad \forall s \in S$$

$$\sum_{s \in S} \left\{ (c_i - r_i) \, y_{is} + (c_i - v_i - l_i) \, z_{is} \right\} \leq c_i \qquad \forall i \in \mathcal{I}$$

$$\sum_{s \in S} p_i^{(s)} = 1 \qquad \forall i \in \mathcal{I} \tag{2.11}$$

$$\sum_{s \in S} \left| \left(\hat{p}_i^{(s)} \right)^\alpha - \left(p_i^{(s)} \right)^\alpha \right|^{1/\alpha} \leq \rho \qquad \forall i \in \mathcal{I}. \tag{2.12}$$

After substituting $w_{is} = p_i^{(s)} x$ and multiplying (2.11) and (2.12) with x, the convex reformulation becomes:

$$(\text{COD-NV}) \quad \max_{x \geq 0, \boldsymbol{y} \leq \boldsymbol{0}, \boldsymbol{z} \leq \boldsymbol{0}, w \geq 0} \quad \gamma x + \sum_{i \in \mathcal{I}, s \in S} d_i^{(s)} \left((v_i - r_i) \, y_{is} - l_i z_{is} \right)$$

$$\text{s.t.} \quad w_{is} + y_{is} + z_{is} = 0 \qquad \forall i \in \mathcal{I} \quad \forall s \in S \tag{2.13}$$

$$\sum_{s \in S} \left\{ (c_i - r_i) \, y_{is} + (c_i - v_i - l_i) \, z_{is} \right\} \leq c_i \qquad \forall i \in \mathcal{I}$$

$$\tag{2.14}$$

$$\sum_{s \in S} w_{is} = x \qquad \forall i \in \mathcal{I}$$

$$\sum_{s \in S} \left| \left(\hat{p}_i^{(s)} x \right)^\alpha - w_{is}^\alpha \right|^{1/\alpha} \leq \rho x \qquad \forall i \in \mathcal{I}. \tag{2.15}$$

We take $\alpha = 0.5$, $\gamma = 100$ and for all other parameters we take the same values as reported by Ben-Tal et al. (2013). This means that there are three scenarios per item, corresponding to low ($d_i^{(s)} = 4$), medium ($d_i^{(s)} = 8$) and high ($d_i^{(s)} = 10$) demand. The other parameters are listed in Table 2.1. We solve the problem for different values of ρ, varying between 0.000 and 0.030 in steps of 0.0001, where $\rho = 0$ corresponds to the nonrobust formulation where the estimates $\hat{p}_i^{(s)}$ are assumed to be exact, with AIMMS 3.11 and KNITRO 7.0. First, we check the conditions of Theorem 1. (COD-NV) is feasible for $\rho = 0$, and the solution is a Slater point of (COD-NV) for larger ρ. By (2.15) all w_{is} are forced to 0 if $x = 0$. Lastly, we have observed numerically that (COD-NV) is unbounded for ρ larger than 0.0306. This means that the primal problem is infeasible for ρ larger than 0.0306. The robust optimal $\boldsymbol{Q}$ and $\boldsymbol{u}$ are the elements of a KKT vector corresponding to constraints (2.13) and (2.14), respectively. As mentioned, we may have to verify the quality of the KKT vector. First, we check whether it is feasible to (R-NV). The constraint violation of (2.10) can be computed by maximizing a linear function over a convex

set, and was found to be at most $1.5 \cdot 10^{-5}$ among all solutions. Second, we need to verify optimality by comparing the objective values of (R-NV) and (COD-NV). We observe both positive and negative relative differences of at most $2.0 \cdot 10^{-4}$. So, the quality of the KKT vector reported by the solver is accurate for this particular problem. The optimal order quantities and corresponding ordering costs are listed in Table 2.2 for different values of ρ.

For every solution we have uniformly sampled 10,000 p matrices from the uncertainty region, and computed the corresponding expected profit. Because implementing the robust solution requires a larger investment, the following comparison is based on the expected return, which is obtained by dividing the expected profit by the total ordering costs. The mean value and the range of these expected returns are listed in Table 2.3. As can be seen from this picture, the mean value for the nonrobust solution is often worse than the worst case for the robust solution. We will explain why the robust solution performs much better using the expected return of a single item. In the same way as for all items together, we have computed the expected return for item 3 (Figure 2.1). The peak at $\rho = 0.022$ is not a simulation inaccuracy, but it is caused by buying a larger number of items than for slightly larger or smaller ρ. The largest increase in expected return is between $\rho = 0$ and $\rho = 0.005$, for which the order quantity increases from 4.00 to 5.87 (Table 2.2). The profits for item 3 in the three scenarios are $(12, -8, -18)$ for $Q_3 = 4.00$, and $(3.6, 7.0, -3.0)$ for $Q_3 = 5.87$. So, with a slightly larger investment, the variation of the profit becomes much smaller, and hence, deviations in the probabilities on the scenarios have a smaller impact on the expected profit. This reduces the range of the expected return of the the robust solution, which can be seen in Figure 2.1. In total, six items show this behaviour. One of these (item 10) is more robust only for ρ between 0.003 and 0.012. For five items the robust order quantity is the same as the nonrobust order quantity and hence they do not have better robust performance. One item (item 2) has slightly worse robust performance.

Acknowledgments

Table 2.1 – Parameter values for the multi-item newsvendor example

i	1	2	3	4	5	6	7	8	9	10	11	12
c_i	4	5	6	4	5	6	4	5	6	4	5	6
v_i	6	8	9	5	9	8	6	8	9	6.5	7	8
r_i	2	2.5	1.5	1.5	2.5	2	2.5	1.5	2	2	1.5	1
l_i	4	3	5	4	3.5	4.5	3.5	3	5	3.5	3	5
$\hat{p}_i^{(1)}$	0.375	0.250	0.375	0.127	0.958	0.158	0.485	0.142	0.679	0.392	0.171	0.046
$\hat{p}_i^{(2)}$	0.375	0.250	0.250	0.786	0.007	0.813	0.472	0.658	0.079	0.351	0.484	0.231
$\hat{p}_i^{(3)}$	0.250	0.500	0.375	0.087	0.035	0.029	0.043	0.200	0.242	0.257	0.345	0.723

Table 2.2 – Optimal ordering cost and quantities for the multi-item newsvendor problem

ρ	cost	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8	Q_9	Q_{10}	Q_{11}	Q_{12}
0.000	391	8.00	8.00	4.00	8.00	4.00	8.00	4.00	8.00	4.00	8.00	7.03	8.00
0.005	412	8.00	8.00	5.87	8.00	4.00	8.00	5.69	8.00	4.00	7.01	8.00	8.34
0.010	421	8.00	8.00	6.20	8.00	4.00	8.00	6.12	8.00	4.00	7.55	8.00	8.85
0.015	430	8.00	8.00	6.39	8.00	4.00	8.00	6.36	8.00	4.00	8.00	8.00	9.62
0.020	440	8.00	8.00	7.10	8.00	4.00	8.00	7.31	8.00	4.00	8.00	8.00	10.00
0.025	453	8.00	8.00	7.36	8.00	4.00	8.00	8.00	8.00	5.51	8.00	8.00	10.00
0.030	469	8.00	9.49	8.00	8.00	4.00	8.00	8.00	8.00	6.26	8.00	8.00	10.00

Table 2.3 – Simulation results of the expected return for the robust and nonrobust solutions for the multi-item newsvendor problem.

ρ	Robust solution			Nonrobust solution		
	min	mean	max	min	mean	max
0.000	0.2557	0.2557	0.2557	0.2557	0.2557	0.2557
0.005	0.2538	0.2695	0.2842	0.2369	0.2542	0.2737
0.010	0.2545	0.2753	0.2987	0.2269	0.2529	0.2795
0.015	0.2560	0.2806	0.3076	0.2169	0.2515	0.2855
0.020	0.2502	0.2855	0.3215	0.2121	0.2501	0.2903
0.025	0.2497	0.2815	0.3125	0.2088	0.2490	0.2955
0.030	0.2425	0.2809	0.3140	0.2026	0.2476	0.2900

Figure 2.1 – Simulation results of the expected return for item 3 in the robust and nonrobust solutions for the multi-item newsvendor problem.

Hints for practical Robust Optimization

Abstract Robust optimization (RO) is a young and active research field that has been mainly developed in the last 15 years. RO techniques are very useful for practice and not difficult to understand for practitioners. It is therefore remarkable that real-life applications of RO are still lagging behind; there is much more potential for real-life applications than has been exploited hitherto. The aim of this chapter is to help practitioners to successfully apply RO in practice. We pinpoint several important items that may be helpful for successfully applying RO. We use many small examples to illustrate our discussions.

3.1   Introduction

Real-life optimization problems often contain uncertain data. The reasons for data uncertainty could be measurement/estimation errors that come from the lack of knowledge of the parameters of the mathematical model (e.g., the uncertain demand in an inventory model) or could be implementation errors that come from the physical impossibility to exactly implement a computed solution in a real-life setting. There are two complementary approaches to deal with data uncertainty in optimization, namely robust and stochastic optimization. *Stochastic optimization* (SO) has an important assumption, i.e., the true probability distribution of uncertain data has to be known or estimated. If this condition is met and the deterministic counterpart of the uncertain optimization problem is computationally tractable, then SO is the methodology to solve the uncertain optimization problem at hand. For details on SO, we refer to Prékopa (1995); Birge and Louveaux (2011); Ruszczyński and Shapiro (2003), but the list of references can be easily extended. *Robust optimization* (RO), on the other hand, does not assume that probability distributions are known, but instead it assumes that the uncertain data resides in a so-called uncertainty set. Additionally, basic versions of RO assume "hard" constraints, i.e., constraint viola-

tion cannot be allowed for any realization of the data in the uncertainty set. RO is popular because of its computational tractability for many classes of uncertainty sets and problem types. For a detailed overview of the RO framework, we refer to Ben-Tal et al. (2009a); Ben-Tal and Nemirovski (2008), and Bertsimas et al. (2011).

RO is a relatively young and active research field, and has been mainly developed in the last 15 years. Especially in the most recent 5 years there have been many publications that show the value of RO in applications in many fields including finance (Lobo, 2000), management science (Ben-Tal and Nemirovski, 1998), supply chain (Bertsimas and Thiele, 2004), healthcare (Fredriksson et al., 2011), engineering (Ben-Tal and Nemirovski, 2002), etc. Indeed, the RO concepts and techniques are very useful for practice and not difficult to understand for practitioners. It is therefore remarkable that real-life applications are still lagging behind; there is much more potential for real-life applications than has been exploited hitherto. In this chapter we pinpoint several items that are important when applying RO and that are often not well understood or incorrectly applied by practitioners.

The aim of this chapter is to help practitioners to successfully apply RO in practice. Many practical issues are treated, as: (i) how to choose the uncertainty set? (ii) Should the decision rule be a function of the final or the primitive uncertain parameters? (iii) Should the objective also be optimized for the worst case? (iv) How to deal with integer adjustable variables? (v) How to deal with equality constraints? (vi) What is the right interpretation of "RO optimizes for the worst case"? (vii) How to compare the robustness characteristics of two solutions?

We also discuss several important insights and their consequences in applying RO. Examples are: (i) sometimes an uncertainty set is constructed such that it contains the true parameter with a prescribed probability. However, the actual constraint satisfaction probability is generally much larger than the prescribed value, since the constraint also holds for other uncertain parameters that are outside the uncertainty set. (ii) The robust reformulations of two equivalent deterministic optimization problems may not be equivalent. (iii) Comparing the robust objective value of the robust solution with the nominal objective value of the nominal solution is incorrect when the objective is uncertain. (iv) For several multi-stage problems the normal robust solution, or even the nominal solution, may outperform the adjustable solution both in the worst case and in the average performance when the solution is re-optimized in each stage.

The remainder of the chapter is organized as follows. Section 3.2 gives a recipe for applying RO. This recipe contains the important items in this chapter. Section 3.3 presents alternative ways of constructing uncertainty sets. Section 3.4 discusses how to model uncertainties in linear (or affine) decision rules. Section 3.5 proposes

a RO method to model adjustable integer variables. Section 3.6 shows that binary variables in big-M type constraints are automatically adjustable. Section 3.7 shows that robust counterparts of equivalent deterministic problems are not necessarily equivalent. Section 3.8 presents some ways to deal with equality constraints. Section 3.9 gives insights about maximin and minimax formulations in RO. Section 3.10 shows two tests to quantify the quality of a robust solution. Section 3.11 shows that static RO with folding horizon can take better decisions than linearly adjustable RO in multi-stage problems. Section 3.12 summarizes our conclusions, and indicates future research topics.

3.2   Recipe for robust optimization in practice

In this section we first give a brief introduction on RO, and then we give a recipe for applying RO in practice. Important items at each step of the recipe and the scopes of other sections that are related to these items are presented in this section.

For the sake of exposition, we use an uncertain linear optimization problem, but we point out that most of our discussions in this chapter can be generalized for other classes of uncertain optimization problems. The "general" formulation of the uncertain linear optimization problem is as follows:

$$\max_{\mathbf{x} \geq \mathbf{0}} \left\{ \mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d} \right\}_{(\mathbf{c},\mathbf{A},\mathbf{d}) \in \mathcal{U}}, \tag{3.1}$$

where $\mathbf{c}$, $\mathbf{A}$ and $\mathbf{d}$ denote the uncertain coefficients, and $\mathcal{U}$ denotes the user specified uncertainty set. The "basic" RO paradigm is based on the following three assumptions (Ben-Tal et al., 2009a, p. xii):

1. All decision variables $\mathbf{x}$ represent "here and now" decisions: they should get specific numerical values as a result of solving the problem before the actual data "reveals itself".

2. The decision maker is fully responsible for consequences of the decisions to be made when, and only when, the actual data is within the prespecified uncertainty set $\mathcal{U}$.

3. The constraints of the uncertain problem in question are "hard" - the decision maker cannot tolerate violations of constraints when the data is in the prespecified uncertainty set $\mathcal{U}$.

Without loss of generality, the objective coefficients ($\mathbf{c}$) and the right-hand side values ($\mathbf{d}$) can be assumed certain (see Section 1.1.1). Often there is a vector of primitive

uncertainties $\boldsymbol{\zeta}$ such that the uncertain parameter $\mathbf{A}$ is a linear function of $\boldsymbol{\zeta}$:

$$\mathbf{A}(\boldsymbol{\zeta}) = \mathbf{A^0} + \sum_{\ell=1}^{L} \zeta_\ell \mathbf{A}^\ell,$$

where $\mathbf{A^0}$ is the nominal value matrix, $\mathbf{A}^\ell$ are the shifting matrices, and $\mathcal{Z}$ is the user specified primitive uncertainty set. The robust reformulation of (3.1) that is generally referred to as the *robust counterpart* (RC) problem, is then as follows:

$$\min_{\mathbf{x} \geq \mathbf{0}} \{\mathbf{c}^\top \mathbf{x} : \quad \mathbf{A}(\boldsymbol{\zeta})\mathbf{x} \leq \mathbf{d} \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}\}.$$

A solution $\boldsymbol{x}$ is called robust feasible if it satisfies the uncertain constraints $[\mathbf{A}(\boldsymbol{\zeta})x \leq \mathbf{d}]$ for all realizations of $\boldsymbol{\zeta}$ in the uncertainty set $\mathcal{Z}$.

    In multistage optimization, the first assumption of the RO paradigm can be relaxed. For example, the amount a factory will produce next month is not a "here and now" decision, but a "wait and see" decision that will be taken based on the amount sold in the current month. Some decision variables can therefore be adjusted at a later moment in time according to a decision rule, which is a function of (some or all part of) the uncertain data. The *adjustable* RC (ARC) is given as follows:

$$\min_{\mathbf{x} \geq \mathbf{0}} \{\mathbf{c}^\top \mathbf{x} : \quad \mathbf{A}(\boldsymbol{\zeta})\mathbf{x} + \mathbf{B}\mathbf{y}(\boldsymbol{\zeta}) \leq \mathbf{d} \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}\},$$

where $\mathbf{B}$ denotes a certain coefficient matrix (i.e., fixed recourse), $\mathbf{x}$ is a vector of non-adjustable variables, and $\mathbf{y}(\boldsymbol{\zeta})$ is a vector of adjustable variables. Linear decision rules are commonly used in practice:

$$\mathbf{y}(\boldsymbol{\zeta}) := \mathbf{y^0} + \sum_{\ell=1}^{L} \mathbf{y}^\ell \zeta_\ell,$$

where $\mathbf{y^0}$ and $\mathbf{y}^\ell$ are the coefficients in the decision rule, which are to be optimized. Notice that we assume a fixed recourse situation for tractability. Another factor that affects the computational tractability of ARC is the type of the decision rule, but we shall focus on this issue later in Section 3.4. Now having introduced the general notation in RO and *adjustable* RO (ARO), we can give a recipe for applying RO in practice.

Practical RO Recipe

Step 0: Solve the nominal problem.

Step 1: **a)** Determine the uncertain parameters.

b) Determine the uncertainty set.

Step 2: Check robustness of the nominal solution.

IF the nominal solution is robust "enough" **THEN** stop.

Step 3: **a)** Determine the adjustable variables.

b) Determine the type of decision rules for the adjustable variables.

Step 4: Formulate the robust counterpart.

Step 5: **IF** an exact or approximate tractable reformulation of the (adjustable) robust counterpart can be derived **THEN** solve it.

ELSE use the adversarial approach.

Step 6: Check quality of the robust solution. **IF** the solution is too conservative **THEN** go to **Step 1b** or **Step 3**.

In the remainder of this section, we describe the most important items at each step of this algorithm. Several items need a more detailed description, and this is done in Sections 3–11.

Step 0 (*Solve the nominal problem*). First, we solve the problem with no uncertainty, i.e., the nominal problem.

Step 1a (*Determine uncertain parameters*). As already described above, in many cases the uncertain parameter is in fact a (linear) function of the *primitive uncertain parameter* ζ. Note that even though there could be many uncertain parameters in the problem at hand, the number of real or primitive sources of uncertainties is "generally" limited. An important example are the so-called *factor models* in finance, where the uncertain returns of different types of assets are linear functions of a limited number of economic factors. These economic factors are considered as the primitive

uncertain parameters. One of the most famous examples of that is the 3-factor model of Fama and French (1993).

Step 1b (*Determine uncertainty set*). We refer to Section 3.3 for a treatment on natural choices of uncertainty sets.

Step 2 (*Check robustness of nominal solution*). For several applications the nominal optimal solution may already be robust. However, in general using the nominal optimal solution often leads to "severe" infeasibilities. In this step we advocate to do a simulation study to analyze the quality of the nominal solution. Section 3.10 extensively describes how to do that. If the nominal solution is already robust "enough", then there is of course no need to apply RO.

In some applications the constraints are not that strict, and one is more interested in a good "average behavior". Note however that the RO methodology is primarily meant for protecting against the worst case scenario in an uncertainty set. However, often, as a byproduct, the robust solution shows good average behavior, but that is certainly not guaranteed.

If one is interested in a good average behavior, then one may try to use smaller uncertainty sets or use *globalized robust optimization* (GRO); for details on GRO we refer to Ben-Tal et al. (2009a, Chapters 3&11).

Step 3a & 3b (*Determine adjustable variables and decision rules*). We discuss several important issues with respect to Step 3, these are listed below.

Reducing extra number of variables. To obtain computationally tractable robust counterpart problems, one often has to use linear decision rules. However, when the number of uncertain parameters is high, the use of linear decision rules may lead to a big increase of the number of variables. Note that these extra variables are added to all constraints that contain adjustable variables. Moreover, when a constraint or the objective in the original problem does not contain uncertain parameters, but does contain adjustable variables, then after substituting the decision rule it will have uncertain parameters, and this will also lead to extra variables in the robust counterpart.

Sometimes, one can choose between a decision rule that is linear in the primitive uncertain parameter $\zeta \in \mathcal{Z}$ or linear in the "general" uncertain parameter $\mathbf{A} \in \mathcal{U}$. Often the number of primitive uncertain parameters is much smaller, and using them for the decision rule will lead to less variables. In Section 3.4 the advantages and disadvantages of both choices are treated. In many cases we have to restrict the linear decision rule to a subset of the uncertain vector ζ. This is especially the case in multi-period situations. In a production-inventory situation, for example, a linear decision rule in period t can only depend on the known demand of period 1 to $t-1$,

since the demand in periods $t, t+1$ and so on is not known yet. This also reduces the number of extra variables.

To further avoid a big increase in the number of variables because of the linear decision rule, one can try to use a subset of the uncertain vector $\boldsymbol{\zeta}$ that is called the "information base". In a production-inventory situation, for example, we may choose a linear decision rule in period t that depends on the known demand of, example given, the last two periods $t-1$ and $t-2$. This reduces the number of variables a lot, and numerical experiments have shown that often the resulting decision rule is almost as good as the full linear one; e.g., see Ben-Tal et al. (2009a). By comparing different information bases one could calculate the value of information.

Often an optimization problem contains *analysis* variables. As an example we give the inventory at time t in a production-inventory problem. For such analysis variables we can use a decision rule that depends on all the uncertain parameters, since we do not have to know the value of these analysis variables "here-and-now". The advantage of making analysis variables adjustable is that this may lead to better objective values. The disadvantage of this, however, is the increase of the number of extra variables.

Integer adjustable variables. A parametric decision rule, like the linear one, cannot be used for integer adjustable variables, since we have then to enforce that the decision rule is integer for all $\boldsymbol{\zeta} \in \mathcal{Z}$. In Section 3.5 we propose a new general way for dealing with adjustable integer variables. However, much more research is needed. In Section 3.6 we show that in some cases the integer variables are automatically adjustable.

Quadratic uncertainty. Suppose that we use a quadratic decision rule instead of a linear one. Then, the corresponding robust counterpart is still linear in all the optimization variables, but quadratic in the uncertain parameters. Hence, if the uncertainty set is ellipsoidal, we can use the results from Ben-Tal et al. (2009a) to obtain a tractable reformulation. In fact, the final constraint is then a semidefinite programming (SDP) constraint.

Suppose that the situation is not fixed recourse as assumed above, but that **B** is also uncertain and linear in ζ. Then using a linear decision rule for y results into quadratic uncertainties. Hence, if the uncertainty set is ellipsoidal, we can use the results from Ben-Tal et al. (2009a) to obtain a tractable reformulation. The resulting constraint is again an SDP.

Constraint-wise uncertainty in ARO. We emphasize that if an adjustable variable is used in multiple constraints, those constraints then contain the same set of uncertain parameters, since the adjustable variable is usually a function of all uncertain parameters; see Section 3.2. We have seen that, in RO, without loss of generality we can reformulate the robust problem such that we have constraint-wise uncertainty.

However, in ARO, we should first substitute the decision rules for adjustable variables, and then make the uncertainty constraint-wise; but not the other way around, since this may result in incorrect reformulations.

It can be shown that when the uncertainty in the original robust optimization problem is constraint-wise, then the objective values of ARC and RC are the same Ben-Tal et al. (2009a). Hence, in such cases using decision rules for adjustable variables does not lead to better objective values. However, there may still be value in using ARC since this may lead to (much) better average behavior; see the numerical example in Section 3.5.

Folding horizon. If one is allowed to reoptimize after each stage in a multi-stage problem, one can of course use adjustable robust optimization in each stage, using that part of the uncertain data that has been revealed. This is called a *folding horizon* (FH) strategy. To compare the ARC FH strategy with the nominal solution, one should also apply an FH strategy to the nominal optimization problem. One could also apply the RC approach in an FH. In many cases this is a good alternative for the ARC approach, e.g., when the ARC approach leads to too large problems. It even appeared that RC FH may lead to better solutions than ARC FH; see Section 3.11.

Step 4 (*Formulate robust counterpart*). RO has also to do with modeling. The modeling part is often overlooked in RO applications. An error often made in practice is that the robustness is added to the model after reformulation of the deterministic model. This often leads to solutions that are too conservative. Hence, an important warning is that the robust versions of two equivalent deterministic optimization problems may not be equivalent. We refer to Section 3.7 for a detailed treatment on these modeling issues.

We also observed that in several applications there are only one or a few uncertain parameters in each constraint, but the uncertainty set is a "joint" region (e.g., ellipsoidal region). Using the constraint-wise interpretation of the RO methodology may be too conservative for such problems, especially in the case where the constraint are not that strict.

It is very important to understand the basic RO concept. What does it mean that RO protects against the worst case scenario? Section 3.9 explains this in more detail.

Step 5 (*Solve RC via tractable reformulation*). If the constraints are linear in the uncertain parameters and in the optimization variables, then there are two ways to derive a tractable reformulation. The first way is the constraint-wise approach by Ben-Tal et al. (2014) that uses Fenchel duality; see Table 3.1 for a summary. The second way is to solve the dual problem of the robust counterpart problem. This

approach can handle all compact and convex uncertainty sets; see Chapter 2. If the constraints are nonlinear in the uncertain parameter and/or the variables, we refer to Ben-Tal et al. (2014) for deriving tractable robust counterparts. However, we emphasize that for many of such problems it might be not possible to derive tractable robust counterparts.

In Iancu and Trichakis (2014) it is observed that (A)RCs may have multiple optimal solutions. We advice to check whether this is the case, and to use a two-step procedure to find Pareto optimal solutions and to improve on the average behavior; for details see Section 3.5, Iancu and Trichakis (2014), and de Ruiter (2013).

Table 3.1 – Tractable reformulations for the uncertain constraint $[(\mathbf{a^0} + \mathbf{P}\boldsymbol{\zeta})^{\top}\mathbf{x} \leq d \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}]$,
and h_k^* is the convex conjugate of h_k

Uncertainty	$\mathcal{Z}$	Robust Counterpart	Tractability
Box	$\|\boldsymbol{\zeta}\|_{\infty} \leq \rho$	$(\mathbf{a^0})^{\top}\mathbf{x} + \rho\|\mathbf{P}^{\top}\mathbf{x}\|_1 \leq d$	LP
Ellipsoidal	$\|\boldsymbol{\zeta}\|_2 \leq \rho$	$(\mathbf{a^0})^{\top}\mathbf{x} + \rho\|\mathbf{P}^{\top}\mathbf{x}\|_2 \leq d$	CQP
Polyhedral	$\mathbf{D}\boldsymbol{\zeta} + \mathbf{d} \geq 0$	$\begin{cases} (\mathbf{a^0})^{\top}\mathbf{x} + \mathbf{d}^{\top}\mathbf{y} \leq d \\ \mathbf{D}^{\top}\mathbf{y} = -\mathbf{P}^{\top}\mathbf{x} \\ \mathbf{y} \geq \mathbf{0} \end{cases}$	LP
Convex cons.	$h_k(\boldsymbol{\zeta}) \leq 0 \quad \forall k$	$\begin{cases} (\mathbf{a^0})^{\top}\mathbf{x} + \sum_k u_k h_k^*\left(\frac{\mathbf{w^k}}{u_k}\right) \leq d \\ \sum_k \mathbf{w^k} = \mathbf{P}^{\top}\mathbf{x} \\ \mathbf{u} \geq \mathbf{0} \end{cases}$	Convex Opt.

Step 5 (*Solve RC via adversarial approach*). If the robust counterpart cannot be written as or approximated by a tractable reformulation, we advocate to perform the so-called *adversarial approach*. The adversarial approach starts with a finite set of scenarios $S_i \subset \mathcal{Z}_i$ for the uncertain parameter in constraint i. E.g., at the start, S_i only contains the nominal scenario. Then, the robust optimization problem, in which $\mathcal{Z}_i$ is replaced by S_i is solved. If the resulting solution is robust feasible, we have found the robust optimal solution. If that is not the case, we can find a scenario for the uncertain parameter that makes the last found solution infeasible. E.g., we can search for the scenario that maximizes the infeasibility. We add this scenario to S_i, and solve the resulting robust optimization problem, and so on. For a more detailed description, we refer to Bienstock and Özbay (2008). It appeared that this simple approach often converges to optimality in a few number of iterations. The advantage of this approach is that solving the robust optimization problem with S_i

instead of $\mathcal{Z}_i$ in each iteration, preserves the structure of the original optimization problem. Only constraints of the same type are added, since constraint i should hold for all scenarios in S_i.

We also note that in some cases it may happen that although a tractable reformulation of the robust counterpart can be derived, the size of the resulting problem becomes too big. For such cases the adversarial approach can also be used.

Step 6 (*Check quality of solution*). Since a robustness analysis is extremely important, and in practice one can easily draw wrong conclusions, we extensively describe in Section 3.10 how to perform such an analysis. Frequently stated criticism on RO is that it yields overly pessimistic solutions. Besides performing a wrong robustness analysis, there are several other possible reasons for such criticism. The first is that in the modeling phase one could easily introduce unnecessary pessimism when one does not realize that the robust counterpart of equivalent deterministic problems are not necessarily equivalent. For a detailed explanation on this issue, see Section 3.7. A second reason may be that the constraints that contain uncertain parameters are not that strict as e.g. safety restrictions for the design of a bridge or an airplane. In such cases violating the constraint for some scenarios of the uncertain parameters is not that serious. As it is explained in Step 2, for those cases one could use the GRO methodology or, alternatively, reduce the size of the uncertainty region. These alternatives can also be used when one is more interested in the average than the worst case behavior. Finally, there are also cases where indeed the nominal solution is already robust "enough", and where RO does not yield better and more robust solutions. We argue that in practice such a conclusion is already extremely valuable.

3.3   Choosing uncertainty set

In this section we describe different possible uncertainty sets and their advantages and disadvantages. Often one wants to make a trade-off between "full" robustness and the size of the uncertainty set: a box uncertainty set that contains the full range of realizations for each component of $\boldsymbol{\zeta}$ is the most robust choice and guarantees that the constraint is never violated, but on the other hand there is only a small chance that all uncertain parameters take their worst case values. This has led to the development of smaller uncertainty sets that still guarantee that the constraint is "almost never" violated. Such guarantees are inspired by chance constraints, which are constraints that have to hold with at least a certain probability. Often the underlying probability distribution is not known, and one seeks a distributionally robust solution. One application of RO is to provide a tractable safe approximation of the chance constraint in such cases, i.e. a tractable formulation that guarantees

that the chance constraint holds:

if x satisfies $\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} \leq d \quad \forall \boldsymbol{\zeta} \in \mathcal{U}_\varepsilon$, then x also satisfies $\mathbb{P}_\zeta(\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} \leq d) \geq 1 - \varepsilon$.

For $\varepsilon = 0$, a chance constraint is a traditional robust constraint. The challenge is to determine the set $\mathcal{U}_\varepsilon$ for other values of ε. We distinguish between uncertainty sets for uncertain parameters and for uncertain probability vectors.

For uncertain parameters, many results are given in (Ben-Tal et al., 2009a, Chapter 2). The simplest case is when the only knowledge about $\boldsymbol{\zeta}$ is that $||\boldsymbol{\zeta}||_\infty \leq 1$. For this case, the box uncertainty set is the only set that can provide a probability guarantee (of $\varepsilon = 0$). When more information becomes available, such as bounds on the mean or variance, or knowledge that the probability distribution is symmetric or unimodal, smaller uncertainty sets become available. Ben-Tal et al. (2009a, Table 2.3) list seven of these cases. Probability guarantees are only given when $||\boldsymbol{\zeta}||_\infty \leq 1$, $\mathbb{E}(\boldsymbol{\zeta}) = \mathbf{0}$ and the components of $\boldsymbol{\zeta}$ are independent. We mention the uncertainty sets that are used in practice when box uncertainty is found to be too pessimistic. The first is an ellipsoid (Ben-Tal et al., 2009a, Proposition 2.3.1), possibly intersected with a box (Ben-Tal et al., 2009a, Proposition 2.3.3):

$$\mathcal{U}_\varepsilon = \{\boldsymbol{\zeta} : ||\boldsymbol{\zeta}||_2 \leq \Omega \quad ||\boldsymbol{\zeta}||_\infty \leq 1\}, \tag{3.2}$$

where $\varepsilon = \exp(-\Omega^2/2)$. The second is a polyhedral set (Ben-Tal et al., 2009a, Proposition 2.3.4), called budgeted uncertainty set or the "Bertsimas and Sim" uncertainty set (Bertsimas and Sim, 2004):

$$\mathcal{U}_\varepsilon = \{\boldsymbol{\zeta} : ||\boldsymbol{\zeta}||_1 \leq \Gamma \quad ||\boldsymbol{\zeta}||_\infty \leq 1\}, \tag{3.3}$$

where $\varepsilon = \exp(-\Gamma^2/(2L))$. A stronger bound is provided in (Bertsimas and Sim, 2004). This set has the interpretation that (integer) Γ controls the number of elements of $\boldsymbol{\zeta}$ that may deviate from their nominal values. (3.2) leads to better objective values for a fixed ε compared to (3.3), but gives rise to a CQP for an uncertain LP while (3.3) results in an LP and is therefore from a computational point of view more tractable.

Bandi and Bertsimas (2012) propose uncertainty sets based on the central limit theorem. When the components of $\boldsymbol{\zeta}$ are independent and identically distributed with mean μ and variance σ^2, the uncertainty set is given by:

$$\mathcal{U}_\varepsilon = \left\{\boldsymbol{\zeta} : |\sum_{i=1}^{L} \zeta_i - L\mu| \leq \rho\sqrt{n}\sigma\right\},$$

where ρ controls the probability of constraint violation $1 - \varepsilon$. Bandi and Bertsimas also show variations on $\mathcal{U}_\varepsilon$ that incorporate correlations, heavy tails, or other distributional information. The advantage of this uncertainty set is its tractability, since

the robust counterpart of an LP with this uncertainty set is also LP. A disadvantage of this uncertainty set is that it is unbounded for $L > 1$, since one component of $\boldsymbol{\zeta}$ can be increased to an arbitrarily large number (while simultaneously decreasing a different component). This may lead to intractability of the robust counterpart or to trivial solutions. In order to avoid infeasibility, it is necessary to define separate uncertainty sets for each constraint, where the summation runs only over the elements of $\boldsymbol{\zeta}$ that appear in that constraint. Alternatively, it may help to take the intersection of $\mathcal{U}_\varepsilon$ with a box.

We now focus on uncertain probability vectors. These appear e.g. in a constraint on a risk measure expected value or variance. Ben-Tal et al. (2013) construct uncertainty sets based on ϕ-divergence. The ϕ-divergence between the vectors $\mathbf{p}$ and $\mathbf{q}$ is:

$$I_\phi(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^m q_i \phi\left(\frac{p_i}{q_i}\right),$$

where ϕ is the (convex) ϕ-divergence function. Let $\mathbf{p}$ denote a probability vector and let $\mathbf{q}$ be the vector with observed frequencies when N items are sampled according to $\mathbf{p}$. Under certain regularity conditions,

$$\frac{2N}{\phi''(1)} I_\phi(\mathbf{p}, \mathbf{q}) \xrightarrow{d} \chi^2_{m-1} \text{ as } N \to \infty.$$

This motivates the use of the following uncertainty set:

$$\mathcal{U}_\varepsilon = \{\mathbf{p} : \mathbf{p} \geq \mathbf{0}, \quad \mathbf{e}^\top \mathbf{p} = 1, \quad \frac{2N}{\phi''(1)} I_\phi(\mathbf{p}, \hat{\mathbf{p}}) \leq \chi^2_{m-1;1-\varepsilon}\},$$

where $\hat{\mathbf{p}}$ is an estimate of $\mathbf{p}$ based on N observations, and $\chi^2_{m-1;1-\varepsilon}$ is the $1 - \varepsilon$ percentile of the χ^2 distribution with $m - 1$ degrees of freedom. The uncertainty set contains the true $\mathbf{p}$ with (approximate) probability $1 - \varepsilon$. Ben-Tal et al. (2013) give many examples of ϕ-divergence functions that lead to tractable robust counterparts.

An alternative to ϕ-divergence is using the Anderson-Darling test to construct the uncertainty set (Ben-Tal et al., 2014, Ex. 15).

We conclude this section by pointing out a mistake that is sometimes made regarding the probability of violation. Sometimes an uncertainty set is constructed such that it contains the true parameter with high probability. Consequently, the constraint holds with the same high probability. However, the probability of constraint satisfaction is much larger than one expects, since the constraint also holds for the "good" realizations of the uncertain parameter outside the uncertainty set. We demonstrate this with a normally distributed $\boldsymbol{\zeta}$ of dimension $L = 10$, where the components are independent, and have mean 0 and variance 1. The singleton $\mathcal{Z}_\varepsilon = \{\mathbf{0}\}$

already guarantees that the uncertain constraint holds with probability 0.5. Let us now construct a set $\mathcal{Z}_\varepsilon$ that contains $\boldsymbol{\zeta}$ with probability 0.5. Since $\boldsymbol{\zeta}^\top \boldsymbol{\zeta} \sim \chi_L^2$, the set $\mathcal{Z}_\varepsilon = \{\boldsymbol{\zeta} : ||\boldsymbol{\zeta}||_2 \leq \sqrt{\chi_{L;1-\varepsilon}^2}\}$ contains $\boldsymbol{\zeta}$ with probability $1 - \varepsilon$. For $\varepsilon = 0.5$, $\mathcal{Z}_\varepsilon$ is a ball with radius 9.3 which is indeed much larger than the singleton. Consequently, it provides a much stronger probability guarantee. In order to compute this probability, we first write the explicit chance constraint. Since $(\mathbf{a^0} + \mathbf{P}\boldsymbol{\zeta})^\top \mathbf{x} \leq d$ is equivalent to $(\mathbf{a^0})^\top \mathbf{x} + (\mathbf{P}^\top \mathbf{x})^\top \boldsymbol{\zeta} \leq d$, and since the term $(\mathbf{P}^\top \mathbf{x})^\top \boldsymbol{\zeta}$ follows a normal distribution with mean 0 and standard deviation $\left|\left|\mathbf{P}^\top \mathbf{x}\right|\right|_2$, the chance constraint can explicitly be formulated as $(\mathbf{a^0})^\top \mathbf{x} + z_{1-\varepsilon} \left|\left|\mathbf{P}^\top \mathbf{x}\right|\right|_2 \leq d$, where $z_{1-\varepsilon}$ is the $1 - \varepsilon$ percentile of the normal distribution. This is the robust counterpart of the original linear constraint with ellipsoidal uncertainty and a radius of $z_{1-\varepsilon}$. The value $z_{1-\varepsilon} = 9.3$ coincides with $\varepsilon \approx 7.0 \cdot 10^{-21}$. So, while one thinks to construct a set that makes the constraint hold in 50% of the cases, the set actually makes the constraint hold in almost all cases. To make the chance constraint hold with probability $1 - \varepsilon$, the radius of the ellipsoidal uncertainty set is $z_{1-\varepsilon}$ instead of $\sqrt{\chi_{L;1-\varepsilon}^2}$. These only coincide for $L = 1$.

3.4 Linearly adjustable robust counterpart: linear in what?

Tractable examples of decision rules used in ARO are linear (or affine) decision rules (AARC) (Ben-Tal et al., 2009a, Chapter 14) or piecewise linear decision rules (Chen et al., 2008); see also Section 3.2. The AARC was introduced by Ben-Tal et al. (2004) as a computationally tractable method to handle adjustable variables. In the following constraint:

$$(\mathbf{a^0} + \mathbf{P}\boldsymbol{\zeta})^\top \mathbf{x} + \mathbf{b}^\top \mathbf{y} \leq d \qquad \forall \boldsymbol{\zeta} \in \mathcal{Z},$$

$\boldsymbol{y}$ is an adjustable variable whose value may depend on the realization of the uncertain $\boldsymbol{\zeta}$, while $\boldsymbol{b}$ does not depend on $\boldsymbol{\zeta}$ (fixed recourse). There are two different AARCs for this constraint:

AARC 1. $\mathbf{y}$ is linear in $\boldsymbol{\zeta}$ (e.g. see Ben-Tal et al. (2004) and Ben-Tal et al. (2009a, Chapter 14)), or

AARC 2. $\mathbf{y}$ is linear in $\mathbf{a^0} + \mathbf{P}\boldsymbol{\zeta}$ (e.g. see Roelofs and Bisschop (2012, Chapter 20.4)).

Note that AARC 2 is as least as conservative as AARC 1, since the linear transformation of $\boldsymbol{\zeta} \mapsto \mathbf{a^0} + \mathbf{P}\boldsymbol{\zeta}$ can only lead to loss of information, and that both methods are equivalent if the linear transformation is injective on $\mathcal{Z}$. The choice for a particular

method may be influenced by four factors: (i) the availability of information. An actual decision cannot depend on $\boldsymbol{\zeta}$ if $\boldsymbol{\zeta}$ has not been observed. (ii) The number of variables in the final problem. AARC 1 leads to $|\boldsymbol{\zeta}|$ extra variables compared to the RC, whereas AARC 2 leads to $|\mathbf{a^0}|$ extra variables. (iii) Simplicity for the user. Often the user observes model parameters instead of the primitive uncertainty vector. (iv) For analysis variables one should always use the least conservative method.

The practical issue raised in the first factor (availability of information) has been addressed with a information base matrix $\mathbf{P}$. Instead of being linear in $\boldsymbol{\zeta}$, $\mathbf{y}$ can be made linear in $\mathbf{P}\boldsymbol{\zeta}$. We give one example where uncertain demand is observed. Suppose there are two time periods and three possible scenarios for demand time period one and two, namely $(10, 10)^\top$, $(10, 11)^\top$ and $(11, 11)^\top$. So, the uncertainty set of the demand vector is the convex hull of these scenarios: $\{\boldsymbol{P}\boldsymbol{\zeta} : \boldsymbol{\zeta} \in \mathcal{Z}\}$ where $\boldsymbol{P}$ is the matrix with the scenarios as columns and $\mathcal{Z} = \Delta^2 = \{\boldsymbol{\zeta} \in \mathbb{R}^3 : \sum_{\ell=1}^3 \zeta_\ell = 1, \boldsymbol{\zeta} \geq 0\}$ is the standard simplex in $\mathbb{R}^3$. If the observed demand for time period one is 10, it is not possible to distinguish between $\boldsymbol{\zeta} = (1, 0, 0)^\top$ and $\boldsymbol{\zeta} = (0, 1, 0)^\top$. So, a decision for time period two can be modeled either as AARC 1 with $\mathbf{P} = (1, 1, 0)$ or as AARC 2. The latter leads to a decision rule that is easier to interpret, since it directly relates to previously observed demand.

3.5   Adjustable integer variables

Ben-Tal et al. (2009a, Chapter 14) use parametric decision rules for adjustable continuous variables. However, their novel techniques "generally" cannot be applied for adjustable integer variables. In the literature two alternative approaches have been proposed. Bertsimas and Georghiou (2013) introduced an iterative method to treat adjustable *binary* variables as piecewise constant functions. The approach by Bertsimas and Caramanis (2010) is different and is based on splitting the uncertainty region into smaller subsets, where each subset has its own binary decision variable (see also Vayanos et al. (2011)). In this section, we briefly show this last method to treat adjustable integer variables, and show how the average behavior can be improved. We use the following notation for the general RC problem:

$$(\text{RC1}) \quad \max_{\mathbf{x},\mathbf{y},\mathbf{z}} \ c(\mathbf{x}, \mathbf{y}, \mathbf{z})$$

$$\text{s.t.} \quad \mathbf{A}(\boldsymbol{\zeta})\,\mathbf{x} + \mathbf{B}(\boldsymbol{\zeta})\,\mathbf{y} + \mathbf{C}(\boldsymbol{\zeta})\,\mathbf{z} \leq \mathbf{d}, \quad \forall \boldsymbol{\zeta} \in \mathcal{Z},$$

where $\mathbf{x} \in \mathbb{R}^{n_1}$ and $\mathbf{y} \in \mathbb{Z}^{n_2}$ are "here and now" variables, i.e., decisions on them are made before the uncertain parameter $\boldsymbol{\zeta}$, contained in the uncertainty set $\mathcal{Z} \subseteq \mathbb{R}^L$, is revealed; $\mathbf{z} \in \mathbb{Z}^{n_3}$ is a "wait and see" variable, i.e., the decision on $\mathbf{z}$ is made after observing (part of) the value of the uncertain parameter. $\mathbf{A}(\boldsymbol{\zeta}) \in \mathbb{R}^{m_1 \times n_1}$ and

$\mathbf{B}(\boldsymbol{\zeta}) \in \mathbb{R}^{m_2 \times n_2}$ are the uncertain coefficient matrices of the "here and now" variables. Notice that the integer "wait and see" variable $\mathbf{z}$ has an uncertain coefficient matrix $\mathbf{C}(\boldsymbol{\zeta}) \in \mathbb{R}^{m_3 \times n_3}$. So, unlike the "classic" parametric method, this approach can handle uncertainties in the coefficients of the integer "wait and see" variables. For the sake of simplicity, we assume that the uncertain coefficient matrices to be linear in $\boldsymbol{\zeta}$ and, without loss of generality, $c(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is the certain linear objective function.

To model the *adjustable* RC (ARC) with integer variables, we first divide the given uncertainty set $\mathcal{Z}$ into m disjoint, excluding the boundaries, subsets ($\mathcal{Z}_i$, $i = 1, \dots, m$):

$$\mathcal{Z} = \bigcup_{i \in \{1, \dots, m\}} \mathcal{Z}_i,$$

and we introduce additional integer variables $\mathbf{z_i} \in \mathbb{Z}^{n_3}$ ($i = 1, \dots, m$) that model the decision in $\mathcal{Z}_i$. Then, we replicate the uncertain constraint and the objective function in (RC1) for each $\mathbf{z_i}$ and the uncertainty set $\mathcal{Z}_i$ as follows:

$$(\text{ARC1}) \quad \max_{\mathbf{x}, \mathbf{y}, \mathbf{Z}, t} \quad t$$
$$\text{s.t.} \quad c(\mathbf{x}, \mathbf{y}, \mathbf{z_i}) \geq t \qquad \forall i \in \{1, \dots, m\} \qquad (3.4)$$
$$\mathbf{A}(\boldsymbol{\zeta})\, \mathbf{x} + \mathbf{B}(\boldsymbol{\zeta})\, \mathbf{y} + \mathbf{C}(\boldsymbol{\zeta})\, \mathbf{z_i} \leq \mathbf{d} \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \dots, m\}.$$

Note that (ARC1) is more flexible than the non-adjustable RC (RC1) in selecting the values of integer variables, since it has a specific decision $\mathbf{z_i}$ for each subset $\mathcal{Z}_i$. Therefore, (ARC1) yields a robust optimal objective that is at least as good as (RC1).

Pareto efficiency. Iancu and Trichakis (2014) discovered that "the inherent focus of RO on optimizing performance only under worst case outcomes might leave decisions un-optimized in case a non worst case scenario materialized". Therefore, the "classical" RO framework might lead to Pareto inefficiencies; i.e., an alternative robust optimal solution may guarantee an improvement in the objective for (at least) a scenario without deteriorating it in other scenarios.

Pareto efficiency is also an issue in (ARC1) that coincides with the *worst case* objective value among m objective functions associated with the subsets. Henceforth, we must take into account the individual performance of the m subsets to have a better understanding of the general performance of (ARC1). To find Pareto efficient robust solutions, Iancu and Trichakis propose reoptimizing the slacks of "important" constraints, i.e., defined by a value vector, by fixing the robust optimal objective value of the classical RO problem that is initially optimized; for details on pareto efficiency in robust linear optimization we refer to Iancu and Trichakis (2014). Following a similar approach, we apply a reoptimization procedure to improve the average

performance of (ARC1). More precisely, we first solve (ARC1) and find the optimal objective t^*. Then, we solve the following problem:

$$(\text{re-opt}) \quad \max_{\mathbf{x}, \mathbf{y}, \mathbf{Z}, \mathbf{t}} \quad \sum_{i \in \{1, \ldots, m\}} t_i$$

$$\text{s.t.} \quad t_i \geq t^* \qquad\qquad\qquad \forall i \in \{1, \ldots, m\}$$

$$c(\mathbf{x}, \mathbf{y}, \mathbf{z_i}) \geq t_i \qquad\qquad \forall i \in \{1, \ldots, m\}$$

$$\mathbf{A}(\boldsymbol{\zeta}) \, \mathbf{x} + \mathbf{B}(\boldsymbol{\zeta}) \, \mathbf{y} + \mathbf{C}(\boldsymbol{\zeta}) \, \mathbf{z_i} \leq \mathbf{d} \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \ldots, m\},$$

that optimizes (i.e., maximizing) the slacks in (3.4), while the worst case objective value t^* remains the same. Note that t_i's are the additional variables associated with the objectives values of the subsets; (re-opt) mimics a multi-objective optimization problem that assigns equal weights to each objective, and finds Pareto efficient robust solutions.

Example

Here we compare the optimal objective values of (RC1), (ARC1), and (ARC1) with (re-opt) via a toy example. For the sake of exposition, we exclude continuous variables in this example. The non-adjustable RC is given as follows:

$$\max_{(w, \mathbf{z}) \in \mathbb{Z}_+^3} \quad 5w + 3z_1 + 4z_2$$

$$\text{s.t.} \quad (1 + \zeta_1 + 2\zeta_2)w + (1 - 2\zeta_1 + \zeta_2)z_1 + (2 + 2\zeta_1)z_2 \leq 18 \quad \forall \boldsymbol{\zeta} \in \text{Box}$$

$$(\zeta_1 + \zeta_2)w + (1 - 2\zeta_1)z_1 + (1 - 2\zeta_1 - \zeta_1)z_2 \leq 16 \qquad \forall \boldsymbol{\zeta} \in \text{Box},$$

$$(3.5)$$

where $\text{Box} = \{\zeta : -1 \leq \zeta_1 \leq 1, -1 \leq \zeta_2 \leq 1\}$ is the given uncertainty set, and w, z_1, and z_2 are nonnegative integer variables. In addition, we assume that z_1 and z_2 are adjustable on ζ_1; i.e., the decision on these variables is made after ζ_1 is being observed. Next, we divide the uncertainty set into two subsets:

$$\mathcal{Z}_1 = \{(\zeta_1, \zeta_2) : -1 \leq \zeta_1 \leq 0, -1 \leq \zeta_2 \leq 1\}$$

$$\mathcal{Z}_2 = \{(\zeta_1, \zeta_2) : 0 \leq \zeta_1 \leq 1, -1 \leq \zeta_2 \leq 1\}.$$

Then ARC of (3.5) is:

$$(\text{Ex:ARC}) \quad \max_{t,w,\mathbf{Z}} \quad t$$

$$\text{s.t.} \quad 5w + 3z_1^i + 4z_2^i \geq t \quad \forall i \in \{1, \ldots, m\}$$

$$(1 + \zeta_1 + 2\zeta_2)w + (1 - 2\zeta_1 + \zeta_2)z_1^i + (2 + 2\zeta_1)z_2^i \leq 18$$
$$\forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \ldots, m\}$$

$$(\zeta_1 + \zeta_2)w + (1 - 2\zeta_1)z_1^i + (1 - 2\zeta_1 - \zeta_1)z_2^i \leq 16$$
$$\forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \ldots, m\},$$

where $t \in \mathbb{R}$, $w \in \mathbb{Z}_+$, $\mathbf{Z} \in \mathbb{Z}_+^{2 \times m}$, and $m = 2$ since we have two subsets. Table 3.2 presents the optimal solutions of RC and ARC problems.

Table 3.2 – RC vs ARC

Method	Obj.	w	$\mathbf{z}$
RC	29	1	$(z_1, z_2) = (4, 3)$
ARC	31	0	$(z_1^1, z_2^1, z_1^2, z_2^2) = (0, 8, 5, 4)$

The numerical results show that using the adjustable reformulation we improve the objective value of the non-adjustable problem by 7%. On the other hand, if we assume that z_1 and z_2 are adjustable on ζ_2 (but not on ζ_1), and we modify the uncertainty subsets $\mathcal{Z}_1$ and $\mathcal{Z}_2$ accordingly, then RC and ARC yield the same objective 29. This shows that the value of information of ζ_1 is higher than that of ζ_2.

Next we compare the average performance of ARC and the second stage optimization problem (re-opt) that is given by:

$$\max_{\mathbf{t},w,\mathbf{Z}} \quad \sum_{i \in \{1, \ldots, m\}} t_i$$

$$\text{s.t.} \quad 5w + 3z_1^i + 4z_2^i \geq t_i, \quad t_i \geq t^* \quad \forall i \in \{1, \ldots, m\}$$

$$(1 + \zeta_1 + 2\zeta_2)w + (1 - 2\zeta_1 + \zeta_2)z_1^i + (2 + 2\zeta_1)z_2^i \leq 18$$
$$\forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \ldots, m\}$$

$$(\zeta_1 + \zeta_2)w + (1 - 2\zeta_1)z_1^i + (1 - 2\zeta_1 - \zeta_2)z_2^i \leq 16$$
$$\forall \boldsymbol{\zeta} \in \mathcal{Z}_i, \forall i \in \{1, \ldots, m\},$$

where $\mathbf{t} \in \mathbb{R}^m$. For changing the number of subsets, we again split the uncertainty sets ($\mathcal{Z}_i, i = 1, \ldots, m$) on ζ_1 but not on ζ_2. The numerical results are presented in Table 3.3.

The first column of the table presents the number of subsets used in ARC, and we assume that the domain of ζ_1 is divided into equally sized intervals (e.g., if the number

Table 3.3 – ARC vs re-opt for varying number of subsets

# Subsets	Worst Case Obj. Values per Subset		W.-C. Average	
	ARC	re-opt	ARC	re-opt
1	29	29	29	29.0
2	(32, 31*)	(34, 31*)	31.5	32.5
3	(33, 30*, 32)	(49, 30*, 35)	31.6	38.0
4	(33, 31*, 32, 32)	(64, 34, 31*, 54)	32	45.7
5	(33, 30*, 30*, 32, 32)	(80, 40, 30*, 33, 66)	31.4	49.8
8	(32, 32, 32, 34, 31*, 33, 33, 33)	(128, 64, 40, 34, 31* 36, 54, 108)	32.5	61.8
10	(32, 32, 32, 32, 34, 31*, 33, 33, 33, 33)	(160, 80, 52, 40, 34, 31*, 33, 45, 66, 135)	32.5	64.3

(∗) denotes the worst case (w.-c.) objective value over all subsets

of subsets is three, then the intervals are $-1 \leq \zeta_1 \leq -0.33, -0.33 \leq \zeta_1 \leq 0.33$, and $0.33 \leq \zeta_1 \leq 1$). The second column reports objective values of the subproblems associated with the subsets in ARC. The third column presents the objective values of the subproblems when we apply (re-opt). The fourth and fifth columns show the averages of the results in columns two and three. As anticipated, intra row comparisons show that ARC and (re-opt) yield the same worst case performance for a fixed number of subsets, and (re-opt) yields significantly better average performances than ARC. Moreover, the average performance improves with the number of subsets. Notice that the average performance of the RC solution is not reported in Table 3.3 because it has the same average performance, that is 29, for any given number of subsets. Nevertheless, it is important to see the significant average objective value improvement made by ARC with (re-opt) for the "fixed" performance of the RC. Last but not least, the optimal objective value 31, which is obtained when the number of subsets is two, four, eight and ten in Table 3.3, is the global optimal of the ARC; for details on optimality see the following section where this example will be revisited.

Optimality

To quantify how far is the optimal objective value (t^*) of (ARC1) from that of the best possible solution, we need to define an efficient lower bound (or an upper bound for a maximization problem) for the best objective. One way of finding such a bound is solving (ARC1) by defining an adjustable variable for each scenario, and scenarios are associated with a finite subset (denoted by $\hat{\mathcal{Z}}$) of the uncertainty set $\mathcal{Z}$ (Hadjiyiannis

et al., 2011; Bertsimas and Georghiou, 2013; Postek, 2013). The optimal objective value of such a formulation is always a lower bound for the best possible objective value, since we optimize adjustable variables for each unique scenario separately and use a finite subset $\hat{\mathcal{Z}}$ that is less conservative (or performs the same in the worst case) than the original uncertainty set $\mathcal{Z}$. More precisely, the lower bound problem is given as follows:

$$
\begin{aligned}
\text{(BRC)} \quad \min_{\mathbf{x}, \mathbf{y}, \mathbf{z}^{(\zeta)}, t^{\text{lb}}} \quad & t^{\text{lb}} \\
\text{s.t.} \quad & c(\mathbf{x}, \mathbf{y}, \mathbf{z}^{(\zeta)}) \leq t^{\text{lb}} && \forall \zeta \in \hat{\mathcal{Z}} \\
& \mathbf{A}(\boldsymbol{\zeta})\, \mathbf{x} + \mathbf{B}(\boldsymbol{\zeta})\, \mathbf{y} + \mathbf{C}(\boldsymbol{\zeta})\, \mathbf{z}^{(\zeta)} \leq \mathbf{d} && \forall \zeta \in \hat{\mathcal{Z}}
\end{aligned}
$$

where $\hat{\mathcal{Z}}$ is a finite subset of $\mathcal{Z}$, as explained above. Now the question that has to be answered is: how to construct $\hat{\mathcal{Z}}$ efficiently? Postek (2013) proposes to first find the optimal solution of (ARC1) for a given number of subsets, and then formulating the set of worst case uncertain parameters for the left-hand sides in active constraints to construct $\hat{\mathcal{Z}}$. For additional details on improving the lower bound we refer to Postek (2013, Chapter 4.2).

Example (Ex:ARC) revisited. The solution of (Ex:ARC) for two subsets (i.e., $m = 2$) is given in the second row of Table 3.2. The associated finite "worst case" subset for this solution is $\hat{\mathcal{Z}} = \{(0, 1), (0, -1)\}$, and the upper bound for the best possible worst case objective is $t^{\text{ub}} = 31$ (this is obtained by solving the upper bound reformulation of (BRC) for $\hat{\mathcal{Z}}$). Therefore, the optimal objective value of (Ex:ARC) is bounded above by 31 for any given number of subsets; but since we find the same objective value for $m = 2$ we can conclude that 31 is the global optimal value.

Tractability

It is important to point out that our adjustable reformulation and the "non-adjustable" RC have the same "general" mathematical complexity, but the adjustable reformulation increases the number of variables and constraints by a factor m (the number of subsets), so that if the number of integer variables is high (say a few hundreds) then the resulting adjustable RC may be intractable. Dividing the main uncertainty set $\mathcal{Z}$ into more subsets $\mathcal{Z}_i$ may improve the objective value by giving more freedom in making adjustable decisions, but the decision maker should make the tradeoff between optimality and computational complexity.

3.6 Binary variables in big-M-constraints are automatically adjustable

Often integer variables correspond to strategic here-and-now decisions, and then there is no need to make them adjustable. In this section we show that for an important class of $0, 1$ variables that are wait-and-see there is also no need to make them adjustable.

Suppose y is a $0, 1$ variable that is associated to a continuous variable x in such a way that:

$$y = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0. \end{cases}$$

Such $0, 1$ variables are often used, e.g., in supply chain models to model whether a facility is opened or not. Now suppose that both y and the continuous variable x are adjustable with respect to the uncertain parameter ζ. Let us use a linear decision rule for x:

$$x = u + \mathbf{v}^\top \boldsymbol{\zeta},$$

where $u \in \mathbb{R}$, and $\mathbf{v} \in \mathbb{R}^L$ are the coefficients of the linear decision rule, and we do not use a linear decision rule for y, although y is adjustable. For the optimal solution of the corresponding robust optimization problem, we either have (i) $u = 0$ and $\mathbf{v} = \mathbf{0}$ (i.e. $x = 0$), or (ii) $u \neq 0$ or $\mathbf{v} \neq \mathbf{0}$. In case (i) we get $y = 0$, and in case (ii) we get $y = 1$. Hence, the only problematic situation is when $u + \mathbf{v}^\top \boldsymbol{\zeta} = 0$ in case (ii), since then the optimal y should be 0 and not 1. Note however that the probability that $\boldsymbol{\zeta} \in \mathcal{U}$ is such that $u + \mathbf{v}^\top \boldsymbol{\zeta} = 0$ is zero, unless $\boldsymbol{\zeta}$ follows a discrete distribution or $u^* + (\mathbf{v}^*)^\top \boldsymbol{\zeta} = 0$, with u^* and $\mathbf{v}^*$ the optimal values for u and $\mathbf{v}$, is part of the description of $\mathcal{Z}$. Also observe that when $u = 0$ and $\mathbf{v} = \mathbf{0}$, we will automatically get $y = 0$, since the constraint $x \leq My$, with M a big number, is one of the constraints of the original problem and the objective is to minimize costs. A more efficient formulation would be:

$$u \leq My, \quad u \geq -My, \quad \mathbf{v} \leq My\mathbf{1}, \quad \mathbf{v} \geq -My\mathbf{1},$$

where $\mathbf{1}$ is the all one vector. We conclude that there is no need to make y adjustable, i.e. the final optimal linear decision rule,

$$y = \begin{cases} 1 & \text{if } u + \mathbf{v}^\top \boldsymbol{\zeta} > 0 \\ 0 & \text{if } u + \mathbf{v}^\top \boldsymbol{\zeta} = 0, \end{cases}$$

can be obtained by only using linear decision rules for x (and not for y). Note that this conclusion depends on the chosen class of decision rules. Suppose that we would have used piecewise linear decision rules, then we should also make y adjustable. One way to do that is to define a different y value for each interval of the piecewise linear decision rule.

Example. Let us consider the following problem. There are two possible production centers and together they have to produce at least ζ. Production costs per unit are 1 and 3, respectively for production center 1 and 2. Fixed costs for opening the centers are 40 and 10, respectively for center 1 and 2. The mathematical formulation is:

$$\min_{\mathbf{x},\mathbf{y}} \quad 40y_1 + 10y_2 + x_1 + 3x_2$$

$$\text{s.t.} \quad x_1 + x_2 \geq \zeta$$

$$x_1 \leq My_1$$

$$x_2 \leq My_2$$

$$x_1, x_2 \geq 0$$

$$y_1, y_2 \in \{0, 1\},$$

in which $M > 0$ is a big number. Now suppose that ζ is uncertain, with uncertainty interval $[10, 30]$, and both x and y are wait-and-see variables. Although y is adjustable we only use linear decision rules for x, and solve the following adjustable robust optimization problem:

$$\min_{\mathbf{x},\mathbf{y}} \max_{\zeta \in [10,30]} \quad 40y_1 + 10y_2 + u_1 + v_1\zeta + 3(u_2 + v_2\zeta)$$

$$\text{s.t.} \quad u_1 + v_1\zeta + u_2 + v_2\zeta \geq \zeta \qquad \forall \zeta \in [10, 30]$$

$$u_1 + v_1\zeta \leq My_1 \qquad \forall \zeta \in [10, 30]$$

$$u_2 + v_2\zeta \leq My_2 \qquad \forall \zeta \in [10, 30]$$

$$u_1 + v_1\zeta \geq 0 \qquad \forall \zeta \in [10, 30]$$

$$u_2 + v_2\zeta \geq 0 \qquad \forall \zeta \in [10, 30]$$

$$y_1, y_2 \in \{0, 1\}.$$

The optimal solution of this problem is:

$$x_1 = \zeta \quad (u_1 = 0, v_1 = 1)$$
$$x_2 = 0 \quad (u_2 = 0, v_2 = 0)$$
$$y_1 = 1$$
$$y_2 = 0,$$

which indeed can not be improved by using decision rules for y. Hence, indeed there is no need to make y adjustable. Now suppose the uncertainty interval is $[0, 30]$, then

the optimal linear decision rule is as above. However, now it can happen that $\zeta = 0$, in which case y_1 should be 0 instead of 1. Hence the final optimal decision rule is:

$$x_1 = \zeta$$
$$x_2 = 0$$
$$y_1 = \begin{cases} 1 & \text{if} \quad \zeta > 0 \\ 0 & \text{if} \quad \zeta = 0 \end{cases}$$
$$y_2 = 0.$$

3.7 Robust counterparts of equivalent deterministic problems are not necessarily equivalent

In this section we show that the robust counterparts of equivalent deterministic problems are not always equivalent. The message in this section is thus that one has to be careful with reformulating optimization problems, since the corresponding robust counterparts may not be the same.

Let us start with a few simple examples. The first one is similar to the example in Ben-Tal et al. (2009a, p. 13). Consider the following constraint:

$$(2 + \zeta)x_1 \leq 1,$$

where ζ is an (uncertain) parameter. This constraint is equivalent to:

$$\begin{cases} (2 + \zeta)x_1 + s = 1 \\ s \geq 0. \end{cases}$$

However, the robust counterparts of these two constraint formulations, i.e.

$$(2 + \zeta)x_1 \leq 1 \quad \forall \zeta \ : \ |\zeta| \leq 1, \tag{3.8}$$

and

$$\begin{cases} (2 + \zeta)x_1 + s = 1 \quad \forall \zeta \ : \ |\zeta| \leq 1 \\ s \geq 0, \end{cases} \tag{3.9}$$

in which the uncertainty set for ζ is the set $\{\zeta \ : \ |\zeta| \leq 1\}$, are not equivalent. It can easily be verified that the feasible set for robust constraint (3.8) is: $x_1 \leq 1/3$, while for the robust constraint (3.9) this is $x_1 = 0$. The reason why (3.8) and (3.9) are not equivalent is that by adding the slack variable, the inequality becomes an equality that has to be satisfied for all values of the uncertain parameter, which is

very restrictive. The general message is therefore: *do not introduce slack variables in uncertain constraints, unless they are adjustable like in Kuhn et al. (2011), and avoid uncertain equalities.*

Another example is the following constraint:

$$|x_1 - \zeta| + |x_2 - \zeta| \leq 2,$$

which is equivalent to:

$$\begin{cases} y_1 + y_2 \leq 2 \\ y_1 \geq x_1 - \zeta \\ y_1 \geq \zeta - x_1 \\ y_2 \geq x_2 - \zeta \\ y_2 \geq \zeta - x_2. \end{cases}$$

However, the robust versions of these two formulations, namely:

$$|x_1 - \zeta| + |x_2 - \zeta| \leq 2 \quad \forall \, \zeta \, : \, |\zeta| \leq 1, \tag{3.10}$$

and:

$$\begin{cases} y_1 + y_2 \leq 2 \\ y_1 \geq x_1 - \zeta & \forall \zeta \, : \, |\zeta| \leq 1 \\ y_1 \geq \zeta - x_1 & \forall \zeta \, : \, |\zeta| \leq 1 \\ y_2 \geq x_2 - \zeta & \forall \zeta \, : \, |\zeta| \leq 1 \\ y_2 \geq \zeta - x_2 & \forall \zeta \, : \, |\zeta| \leq 1, \end{cases} \tag{3.11}$$

are not equivalent. Indeed, it can easily be checked that the set of feasible solutions for (3.10) is $(\theta, -\theta)$, $-1 \leq \theta \leq 1$, but the only feasible solution for (3.11) is $\mathbf{x} = (0, 0)$. The reason for this is that in (3.11) the uncertainty is split over several constraints, and since the concept of RO is constraint-wise, this leads to different problems, and thus different solutions. The following linear optimization reformulation, however, is equivalent to (3.10):

$$\begin{cases} x_1 - \zeta + x_2 - \zeta \leq 2 & \forall \, \zeta \, : \, |\zeta| \leq 1 \\ x_1 - \zeta + \zeta - x_2 \leq 2 & \forall \, \zeta \, : \, |\zeta| \leq 1 \\ \zeta - x_1 + x_2 - \zeta \leq 2 & \forall \, \zeta \, : \, |\zeta| \leq 1 \\ \zeta - x_1 + \zeta - x_2 \leq 2 & \forall \, \zeta \, : \, |\zeta| \leq 1. \end{cases} \tag{3.12}$$

The general rule therefore is: *do not split the uncertainty in one constraint over more constraints, unless the uncertainty is disjoint.* In particular do not use "definition variables" if this leads to such a splitting of the uncertainty.

In the remainder we give a general treatment of some often used reformulation tricks to reformulate nonlinear problems into linear ones, and discuss whether the robust counterparts are equivalent or not.

- **Maximum function.** Consider the following constraint:

  $$\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + \max_k \mathbf{b_k}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) \quad \forall \boldsymbol{\zeta} \in \mathcal{Z},$$

  where $\boldsymbol{\zeta} \in \mathcal{Z}$ is the uncertain parameter, and $\mathbf{a}(\boldsymbol{\zeta})$, $\mathbf{b_k}(\boldsymbol{\zeta})$, and $d(\boldsymbol{\zeta})$ are parameters that depend linearly on $\boldsymbol{\zeta}$. The incorrect reformulation for this constraint is:

  $$\begin{cases} \mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + z \le d(\boldsymbol{\zeta}) & \forall \boldsymbol{\zeta} \in \mathcal{Z} \\ z \ge \mathbf{b_k}(\zeta)^\top \mathbf{x} & \forall k, \forall \boldsymbol{\zeta} \in \mathcal{Z}, \end{cases}$$

  since the uncertainty is split over more constraints. The correct reformulation is:

  $$\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + \mathbf{b_k}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) \quad \forall k, \forall \boldsymbol{\zeta} \in \mathcal{Z}.$$

  Note that in many cases we have "a sum of max":

  $$\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + \sum_i \max_k \mathbf{b_{ik}}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) \quad \forall \boldsymbol{\zeta} \in \mathcal{Z}.$$

  Important examples that contain such constraints are production-inventory problems. We refer to Chapter 4 for an elaborate treatment on exact and approximate reformulations of such constraints.

- **Absolute value function.** Note that $|x| = \max\{x, -x\}$, and hence this is a special case of the max function, treated above.

- **Linear fractional program.** Consider the following robust linear fractional problem:

  $$\begin{cases} \min\limits_{\mathbf{x}} & \max_{\boldsymbol{\zeta} \in \mathcal{Z}} \dfrac{\alpha(\boldsymbol{\zeta}) + \mathbf{c}(\boldsymbol{\zeta})^\top \mathbf{x}}{\beta(\boldsymbol{\zeta}) + \mathbf{d}(\boldsymbol{\zeta})^\top \mathbf{x}} \\ \text{s.t.} & \sum\limits_j a_{ij} x_j \ge b_i \quad \forall i \\ & \mathbf{x} \ge \mathbf{0}, \end{cases} \tag{3.13}$$

where $\boldsymbol{\alpha}(\boldsymbol{\zeta})$, $\mathbf{c}(\boldsymbol{\zeta})$, $\boldsymbol{\beta}(\boldsymbol{\zeta})$, and $\mathbf{d}(\boldsymbol{\zeta})$ are parameters that depend linearly on $\boldsymbol{\zeta}$. Moreover, we assume that $\beta(\boldsymbol{\zeta})+\mathbf{d}(\boldsymbol{\zeta})^\top\mathbf{x} > 0$, for all feasible $\mathbf{x}$ and for all $\boldsymbol{\zeta} \in \mathcal{Z}$. For the non-robust version one can use the Charnes-Cooper transformation that is proposed by Charnes and Cooper (1962) to obtain an equivalent linear optimization problem. However, if we apply this transformation to the robust version, we obtain:

$$
\begin{cases}
\min_{\mathbf{y},t} & \max_{\boldsymbol{\zeta}\in\mathcal{Z}} \alpha(\boldsymbol{\zeta})t + \mathbf{c}(\boldsymbol{\zeta})^\top\mathbf{y} \\
\text{s.t.} & \beta(\boldsymbol{\zeta})t + \mathbf{d}(\boldsymbol{\zeta})^\top\mathbf{y} = 1 \qquad \forall\boldsymbol{\zeta} \in \mathcal{Z} \\
& \sum_j a_{ij}y_j \geq b_i t \qquad\qquad \forall i \\
& \mathbf{y} \geq \mathbf{0}, \ t \geq 0,
\end{cases}
$$

which is not equivalent to (3.13) since the uncertainty in the original objective is now split over the objective and a constraint. A better way to deal with such problems is to solve the robust linear problem

$$
\begin{cases}
\min_{\mathbf{x}} & \max_{\boldsymbol{\zeta}\in\mathcal{Z}} \left[\alpha(\boldsymbol{\zeta}) + \mathbf{c}(\boldsymbol{\zeta})^\top\mathbf{x} - \lambda\left(\beta(\boldsymbol{\zeta}) + \mathbf{d}(\boldsymbol{\zeta})^\top\mathbf{x}\right)\right] \\
\text{s.t.} & \sum_j a_{ij}x_j \geq b_i \\
& \mathbf{x} \geq \mathbf{0},
\end{cases}
$$

for a fixed value of λ, and then find the minimal value of λ for which this optimization problem still has a non positive optimal value. One can use for example binary search on λ to do this. For a more detailed treatment of robust fractional problems we refer to Chapter 5.

- **Product of binary variables.** Suppose that a robust constraint contains a product of binary variables, say xy, with $x, y \in \{0, 1\}$. Then one can use the standard way to linearize this:

$$
\begin{cases}
z \leq x \\
z \leq y \\
z \geq x + y - 1 \\
z \geq 0,
\end{cases}
$$

and replace xy with z. One can use this reformulation since the added constraints do not contain uncertain parameters.

- **Product of binary and continuous variable.** A product of a binary and a continuous variable that occurs in a robust constraint can also be reformulated

in linear constraints, in a similar way as above. However, note that in the following robust constraint:

$$\mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + z\mathbf{b}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) \quad \forall \boldsymbol{\zeta} \in \mathcal{Z},$$

where $z \in \{0, 1\}$, one cannot use the standard trick:

$$\begin{cases} \mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + zy \le d(\boldsymbol{\zeta}) & \forall \boldsymbol{\zeta} \in \mathcal{Z} \\ y \ge \mathbf{b}(\boldsymbol{\zeta})^\top \mathbf{x} & \forall \boldsymbol{\zeta} \in \mathcal{Z}, \end{cases} \tag{3.14}$$

and then linearize zy. This is not possible since in (3.14) the uncertainty is split over different constraints. A correct reformulation is:

$$\begin{cases} \mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} + \mathbf{b}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) + M(1-z) & \boldsymbol{\zeta} \in \mathcal{Z} \\ \mathbf{a}(\boldsymbol{\zeta})^\top \mathbf{x} \le d(\boldsymbol{\zeta}) + Mz & \boldsymbol{\zeta} \in \mathcal{Z}. \end{cases} \tag{3.15}$$

- **K out of N constraints should be satisfied.** Suppose the restriction is that at least K out of the N robust constraints

  $$\mathbf{a_i}(\boldsymbol{\zeta})^\top \mathbf{x} \le d_i(\boldsymbol{\zeta}) \quad \forall \boldsymbol{\zeta} \in \mathcal{Z} \tag{3.16}$$

  should be satisfied, where $i \in \{1, \dots, N\}$. Then one can use the standard way

  $$\begin{cases} \mathbf{a_i}(\boldsymbol{\zeta})^\top \mathbf{x} \le d_i(\boldsymbol{\zeta}) + M(1-z_i) & \forall \boldsymbol{\zeta} \in \mathcal{Z}, \forall i \\ \sum_i z_i \ge K \\ z_i \in \{0, 1\} & \forall i, \end{cases}$$

  where M is a sufficiently big number. However, if the restriction is that $\forall \boldsymbol{\zeta} \in \mathcal{Z}$ at least K out of the N constraints should be satisfied (notice the difference with (3.16)), then the above constraint-wise formulation is not equivalent and is overly conservative. We do not see how to model such a constraint correctly. Maybe an adversarial approach could be used for such constraints.

- **If-then constraint.** Since an "if-then constraint" can be modeled as an at least 1 out of 2 constraints, the above remarks hold.

Up to now we only described linear optimization examples. Similar examples can be given for conic and nonlinear optimization. In Lobo et al. (1998) for example, many optimization problems are given that can be modeled as conic quadratic programming problems. However, for many of them it holds that the corresponding robust counterparts are not the same. This means that if an optimization problem is conic quadratic representable, then the robust counterparts are not automatically the same, and hence in such cases the robust optimization techniques for CQP cannot be used.

3.8 How to deal with equality constraints?

Equality constraints containing uncertain parameters should be avoided as much as possible, since often such constraints restrict the feasibility region drastically or even lead to infeasibility. Therefore, the advice is: *do not use slack variables unless they are adjustable, since using slack variables leads to equality constraints*; see Ben-Tal et al. (2009a, Chapter 2). However, equality constraints containing uncertain parameters cannot always be avoided. There are several ways to deal with such uncertain equality constraints:

- In some cases it might be possible to convert the equality constraints into inequality constraints. An illustrating example is the transportation problem: the demand constraints can either be formulated as equality constraints or as inequality constraints. The structure of the problem is such that at optimality these inequalities are tight.

- The equality constraints can be used to eliminate variables. This idea is mentioned in Ben-Tal et al. (2009a). However, several questions arise. First of all, after elimination of variables and after the resulting problem has been solved, it is unclear which values to take for the eliminated variables, since they also depend on the uncertain parameters. This is no problem if the eliminated variables are adjustable variables or analysis variables, since there is no need to know their optimal values. A good example is the production-inventory problem for which one can easily eliminate the analysis variables indicating the inventory in different time periods. See e.g. Ben-Tal et al. (2009a). Secondly, suppose the coefficients with respect to the variables that will be eliminated contain uncertain parameters. Eliminating such variables leads to problems that contain non-linear uncertainty, which are much more difficult to solve. To illustrate this, let us consider the following two constraints of an optimization problem:

$$\zeta_1 x_1 + x_2 + x_3 = 1, \quad x_1 + x_2 + \zeta_2 x_3 \le 5,$$

in which ζ_1 and ζ_2 are uncertain parameters. Suppose that x_1, x_2 and x_3 are all adjustable in ζ_1. Then there are three options for elimination:

1. *Elimination of x_1.* Let us assume that $\zeta_1 = 0$ is not in the uncertainty set. By substituting $x_1 = (1 - x_2 - x_3)/\zeta_1$ the inequality becomes:

$$\left(1 - \frac{1}{\zeta_1}\right) x_2 + \left(\zeta_2 - \frac{1}{\zeta_1}\right) x_3 \le 5 - \frac{1}{\zeta_1}.$$

The disadvantage of eliminating x_1 is thus that the uncertainty in the inequality becomes nonlinear.

2. *Elimination of x_2.* By substituting $x_2 = 1 - \zeta_1 x_1 - x_3$ the inequality becomes:

$$(1 - \zeta_1)x_1 + (\zeta_2 - 1)x_3 \leq 4,$$

which is linear in the uncertain parameters.

3. *Elimination of x_3.* By substituting $x_3 = 1 - \zeta_1 x_1 - x_2$ the inequality becomes:

$$(1 - \zeta_1 \zeta_2)x_1 + (1 - \zeta_2)x_2 \leq 5 - \zeta_2,$$

which is nonlinear in the uncertain parameters. We conclude that from a computational point of view it is more attractive to eliminate x_2.

It is important to note that different choices of variables to eliminate may lead to different optimization problems.

- If the constraint contains analysis variables one could make these variables adjustable and use decision rules, thereby introducing much more flexibility. One can easily prove that when the coefficients for such variables in the equality constraint do not contain uncertain parameters and the equality constraint is linear in the uncertain parameters, then using linear decision rules for such variables is equivalent to eliminating these variables. To be more precise: suppose the linear equality constraint is

$$\mathbf{q}(\boldsymbol{\zeta})^\top \mathbf{x} + y = r,$$

where $\mathbf{q}(\boldsymbol{\zeta})$ is linear in $\boldsymbol{\zeta}$, and y is an analysis variable (without loss of generality we assume the coefficient for y is 1). Then it can easily be proven that substituting $y = r - \mathbf{q}(\boldsymbol{\zeta})^\top \mathbf{x}$ everywhere in the problem is equivalent to using a linear decision rule for y. To reduce the number of extra variables, it is therefore better to eliminate such variables.

- Consider the following robust constraint:

$$(\mathbf{a} + \mathbf{P}\boldsymbol{\zeta})^\top \mathbf{x} = d \quad \forall \boldsymbol{\zeta} \in Z. \tag{3.17}$$

The equality constraint is satisfied for all $\boldsymbol{\zeta}$ in $\mathcal{Z}$ if $\mathbf{P}^\top \mathbf{x} = \mathbf{0}$. Hence, we could replace (3.17) by the stricter set of equations

$$\mathbf{a}^\top \mathbf{x} = d, \quad \mathbf{P}^\top \mathbf{x} = \mathbf{0}.$$

However, especially when L is large, this is much too restrictive.

- One could also drop the requirement that the constraints are hard, and make such constraints "soft", by adding, e.g., a quadratic penalty for the violations to the objective.

3.9   On maximin and minimax formulations of RC

In this section, we consider an uncertain LP of the following general form:

$$\max_{\mathbf{x} \geq \mathbf{0}} \{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d}\},$$

where without loss of generality $\mathbf{A}$ is the uncertain coefficient matrix that resides in the uncertainty set $\mathcal{U}$. So the general RC is given by

$$(\text{R-LP}) \quad \max_{\mathbf{x} \geq \mathbf{0}} \{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d} \ \ \forall \mathbf{A} \in \mathcal{U}\}.$$

Here we show that (R-LP) can be reformulated as:

$$(\text{RF}) \quad \min_{\mathbf{A} \in \mathcal{U}} \max_{\mathbf{x} \geq \mathbf{0}} \{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d}\},$$

if the uncertainty is constraint-wise; however if this condition is not met, then (RF) may not be equivalent to (R-LP).

Remark 1 *This shows that the statement "RO optimizes for the worst case* $\mathbf{A}$*" is too vague. Also the maximin reformulation:*

$$\max_{\mathbf{x} \geq \mathbf{0}} \min_{\mathbf{A} \in \mathcal{U}} \{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d}\},$$

is usually not equivalent to (R-LP). This is because we can almost always find an $\mathbf{x} \geq 0$ *such that no* $\mathbf{A} \in \mathcal{U}$ *exists for which* $\mathbf{A}\mathbf{x} \leq \mathbf{d}$*; therefore, we minimize over an empty set, and have* $+\infty$ *for the maximin objective. Also when* $\mathbf{x}$ *is selected such that at least one feasible* $\mathbf{A}$ *exists (e.g., see Falk (1973)), it is easy to find examples where both formulations are not equivalent.*

To show (R-LP)=(RF) when the uncertainty is constraint-wise, we first take the dual of the (inside) maximization problem of (RF) $[\max_{\mathbf{x} \geq \mathbf{0}} \mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{d}]$. Then, substituting the dual with the primal (maximization) problem in (RF) gives:

$$(\text{OC-LP}) \quad \min_{\mathbf{A} \in \mathcal{U}, \mathbf{y} \geq \mathbf{0}} \{\mathbf{d}^\top \mathbf{y} : \mathbf{A}^\mathbf{T} \mathbf{y} \geq \mathbf{c}\},$$

where val(RF)=val(OC-LP) at optimality. Note that the constraints of (RF) can be formulated as $[\mathbf{a_i^T}\mathbf{x} \leq d_i, \forall \mathbf{a_i} \in \mathcal{U}_i, i = 1, \ldots, m]$, if the uncertainty is constraint-wise. Beck and Ben-Tal (2009) show that (OC-LP)—which is the *optimistic counterpart of the dual problem*—is equivalent to the general robust counterpart (R-LP) for constraint-wise uncertainty and disjoint $\mathcal{U}_i$'s. However, if (some of) the constraints are dependent in (R-LP), then we may not sustain the associated equivalence. The following example shows such a situation.

Example

Consider the following toy RC example in which the uncertainty is not constraint-wise:

$$(\text{RC-Toy}) \quad \max_{\mathbf{y}} \ y_1 + y_2$$
$$\text{s.t. } a_1 y_1 \leq 1, a_2 y_2 \leq 1 \ \ \forall \mathbf{a} \in \mathbb{R}^2 : ||\mathbf{a}||_2 \leq 1,$$

where two constraints of the problem are dependent on each other via the ellipsoidal uncertainty set $[\mathbf{a} \in \mathbb{R}^2 : ||\mathbf{a}||_2 \leq 1]$. The robust reformulation of the (RC-Toy) is as follows:

$$(\text{RF-Toy}) \quad \min_{\mathbf{a}:||\mathbf{a}||_2 \leq 1} \max_{\mathbf{y}} \ y_1 + y_2$$
$$\text{s.t. } a_1 y_1 \leq 1, a_2 y_2 \leq 1,$$

and the optimistic counterpart (OC) of the problem is

$$(\text{OC-Toy}) \quad \min_{\mathbf{x} \geq \mathbf{0}, \ \mathbf{a}:||\mathbf{a}||_2 \leq 1} \ x_1 + x_2$$
$$\text{s.t. } a_1 x_1 = 1, a_2 x_2 = 1.$$

(RC-Toy) attains an optimal objective value of 2, whereas the (RF-Toy)'s optimal objective value is $2\sqrt{2}$. Therefore, the robust reformulation (RF-Toy) is not equivalent to the general RC problem (RC-Toy) in this situation. However, val(RF-Toy)=val(OC-Toy) from duality.

3.10   Quality of robust solution

In this section we describe how to assess the quality with respect to robustness of a solution based on a simulation study. We first identify four focus points for performing a Monte Carlo experiment, and conclude with two statistical tests that can be used to compare two solutions.

Choice of the uncertainty set. For a comparison between different solutions, it is necessary to define an uncertainty set $\mathcal{U}$ that is used for evaluation. This set should reflect the real-life situation. The uncertainty set that is used for optimization may be different than the set for evaluation. For example, an ellipsoidal set may be used to reduce the conservatism when the real-life uncertainty is a box, while still maintaining a large probability of constraint satisfaction (Ben-Tal et al., 2009a, p. 34).

Choice of the probability distribution. A simulation requires knowledge of the probability distribution on the uncertainty set. If this knowledge is ambiguous, it may be necessary to verify whether the simulation results are sensitive with respect to changes in this distribution. For example, Rozenblit (2010) performs different simulations, each based on a probability distribution with a different skewness level.

Choice of the sampling method. For univariate random variables it is computationally easy to draw a random sample from any given distribution. For multivariate random variables rejection sampling can be used, but it may be inefficient depending on the shape of the uncertainty set, e.g. for an uncertainty set with no volume. A more efficient method for sampling from an arbitrary continuous probability distribution is "hit and run" sampling (Bélisle et al., 1993). An R package for uniform hit and run sampling from a convex body is also available.

Choice of the performance characteristics. From a mathematical point of view there is no difference between uncertainty in the objective and uncertainty in the constraints since an uncertain objective can always be reformulated as a certain objective and an uncertain constraint. However, the distinction between an uncertain objective and an uncertain constraint is important for the interpretation of a solution. First, we look at the effects of adjustable RO and reformulations, then we present the performance characteristics.

Effect of adjustable RO. When one or more "wait and see" variables are modeled as adjustable variables, uncertain parameters may enter the objective function. In that case the performance characteristics for uncertainty in the objective become applicable.

Effect of reformulations. Reformulations are sometimes necessary to end up with a tractable model. The evaluation should be based on the original model, since reformulations introduce additional constraints whose violation is not necessarily a problem. Take for example an inventory model that has constraints on variables that indicate the cost at a certain time period (e.g. constraints (3.18) and (3.19)). These constraints have been introduced to model the costs in the objective function. A violation of these constraints does not render the solution infeasible but does affect the objective value (i.e. the costs of carrying out the solution).

Performance characteristics for uncertainty in the constraints. For an uncertain constraint $f(\mathbf{a}, \boldsymbol{\zeta}) \leq 0$ for all $\boldsymbol{\zeta}$ in $\mathcal{Z}$, the violation is $\max\{0, f(\mathbf{a}, \boldsymbol{\zeta})\}$. Meaningful

statistics are the probability on positive violation and the distribution of the violation (average, worst case, standard deviation) under the condition that the violation is positive. When multiple constraints are uncertain, these statistics can be computed per constraint. Additionally, the average number of violated constraints can be reported.

There is a clear trade-off between the objective value and constraint violations. The difference between the worst case objective value of the robust solution and the nominal objective value of the nominal solution is called the *price of robustness* (PoR) (Bertsimas and Sim, 2004). It is useful if the objective is certain, since in that case PoR is the amount that has to be paid for being robust against constraint violations. We observe that PoR is also used when the objective is uncertain. We discourage this, since it compares the nominal solution in case there is no uncertainty with the robust solution where the worst case occurs, so it compares two different scenarios.

Performance characteristics for uncertainty in the objective. Uncertainty in the objective affects the performance of a solution. For every simulated uncertainty vector, the actual objective value can be computed. One may be interested in the worst case, but also in the average value or the standard deviation. For a solution that is carried out many times, reporting the average performance is justified by the law of large numbers. The worst case may be more relevant when a solution is carried out only once or a few times, e.g. when optimizing a medical treatment plan for a single patient. These numbers show what objective value to expect, but they do not provide enough information about the quality of a solution since a high standard deviation is not necessarily undesirable. A robust solution is good when it is close to the *perfect hindsight* (PH) solution. The PH solution is the solution that is obtained by optimizing the decision variables for a specific uncertainty vector as if it is fully known beforehand. This has to be done for every simulated uncertainty vector, and yields an utopia solution. The PH solution may have a large variation, causing a high variation of good solutions as well.

Performance characteristics for any problem. Regardless of whether the uncertainty is in the objective or in the constraints, the mean and associated standard deviation of the difference between the actual performance of a solution and the PH solution are useful for quantifying the quality of a solution. The mean difference between the PH solution and a fully robust solution is defined as the *price of uncertainty* (PoU) by Ben-Tal et al. (2005). It is the maximum amount that a company should invest for reducing the level of uncertainty, e.g. by using more accurate forecasting techniques. It can also be interpreted as the regret of choosing a certain solution rather than the

PH solution. Alternative names for PoU are "cost of robustness" (Gregory et al., 2011) or "price of robustness" (Ben-Tal et al., 2004), which are less descriptive than "price of uncertainty" and may cause confusion with price of robustness from (Bertsimas and Sim, 2004). A low mean PoU and a low standard deviation characterize a good solution.

Subtracting the mean objective value of the nominal solution from the mean value of a robust solution yields the *actual price of robustness* (APoR) (Rozenblit, 2010). APoR can be interpreted as the expected price that has to be paid for using the robust solution rather than the nominal solution, which is negative if RO offers a solution that is better on average. PoR equals APoR when uncertainty only occurs in the constraints.

For multistage problems one may also follow a folding horizon (FH) approach. With FH in each stage where a part of the uncertain parameter is observed, that information is used to optimize for the remaining time periods. This is done by taking the original optimization problem, fixing the decision variables for previous stages, and fixing the elements of the uncertain parameter that have been observed. This allows a fair comparison between a dynamic solution (e.g. created by the AARC) and a static solution (e.g. the nominal solution) when in real-life the static solution is reoptimized in every stage.

Comparing two solutions. We provide several comparison criteria and provide the corresponding statistical test to verify whether one solution is better than another solution. The tests will be demonstrated in Section 3.11. We will assume that the data for the statistics test is available as n pairs (X_i, Y_i) $(i = 1, 2, \ldots, n)$, where X_i and Y_i are performance characteristics in the i'th simulation. For uncertainty in the objective, they can be objective values whereas for uncertainty in the constraints they can be the numbers of constraint violations or the sizes of the constraint violations. We assume that (X_i, Y_i) and (X_j, Y_j) are independent if $i \neq j$, and that smaller values are better. When a conjecture for a test is based on the outcome of a simulation study, the statistical test must be performed with newly generated data to avoid statistical bias. While for the statistical tests it is not necessary that X_i and Y_i are based on the same simulated uncertainty vector $\boldsymbol{\zeta}$, it increases the power of the test since X_i and Y_i will be positively correlated. This reduces the variance of the difference: $\mathrm{Var}(X_i - Y_i) = \mathrm{Var}(X_i) + \mathrm{Var}(Y_i) - 2\,\mathrm{Cov}(X_i, Y_i)$, which is used in the following tests:

- The sign test for the median validates H_0: $m_x = m_y$ against H_1: $m_x < m_y$ with confidence level α, where m_x and m_y are the medians of the distributions of X_i and Y_i, respectively. This tests the conjecture that the probability that

solution X outperforms solution Y is larger than 0.5. Let $n_=$ be the number of observations for which $X_i = Y_i$ and let Z be the number of negative signs of $X_i - Y_i$. Under the null hypothesis, Z follows a binomial distribution with parameters $n - n_=$ and 0.5. That means that the null hypothesis gets rejected if Z is larger than the $(1 - \alpha)$ percentile of the binomial distribution.

- The t-test for the mean validates H_0: $\mu_x = \mu_y$ against H_1: $\mu_x < \mu_y$ with confidence level α, where μ_x and μ_y are the means of the distributions of X_i and Y_i, respectively. This tests the conjecture that solution X outperforms solution Y in long run average behavior. This test assumes that $X_i - Y_i$ follows a normal distribution. Let $Z_i = X_i - Y_i$, $\bar{Z} = \sum_{i=1}^{n} Z_i/n$ and $s^2 = \sum_{i=1}^{n}(Z_i - \bar{Z})^2/(n-1)$, then $T = \sqrt{n}\sum_{i=1}^{n}(Z_i - \bar{Z})/s$ follows a t-distribution with $n - 1$ degrees of freedom under the null hypothesis. This means that H_0 gets rejected if T is smaller than the α percentile of the t-distribution with $n-1$ degrees of freedom.

3.11 RC may take better "here and now" decisions than AARC

A linear decision rule is a linear approximation of a more complicated decision rule. It dictates what to do at each stage as a linear function of observed uncertain parameters, but it is not guaranteed to be the optimal strategy. Every time a decision has to be made it is possible to either follow the linear decision rule, or to reoptimize the AARC for the remaining time periods based on everything that is observed up till then. We will refer to the latter as the AARC-FH, where FH stands for folding horizon. Ben-Tal et al. (2005) compare the AARC with the AARC-FH, and show that the latter produces better solutions on average. A comparison that involves AARC-FH assumes that there is time to reoptimize. It is therefore natural to also make a comparison with the RC-FH, where the RC is solved for the full time horizon and re-optimized for the remaining time period every time a part of the uncertain parameters is unveiled. On average, the RC-FH may outperform the AARC (Cohen et al., 2007; Rozenblit, 2010).

In the remainder of this section we will evaluate both the average and the worst case performance of the nominal solution with FH, the RC-FH and the AARC-FH. A comparison between RC-FH and AARC-FH is new, and shows which model takes the best "here and now" decisions.

We first give an example for the worst case performance. Consider a warehouse that transfers one good. The current inventory is $x_0 = 5$, the holding costs per time period are $h = 1$, the backlogging costs per time period are $b = 2$. In the first period,

any nonnegative (not necessarily integer) amount can be ordered while in the second period the maximum order quantity is $q_2^{max} = 3$. Let $T = \{1, 2\}$, let q_t be the order quantity in time period t, and let c_t denote the costs associated with time period t. The ending inventory can be returned to the supplier without penalty fee at time period three. The optimization problem can be formulated as:

$$\min \qquad \sum_{t \in T} c_t$$

$$\text{s.t.} \qquad c_t \geq (x_0 + \sum_{i=1}^{T} q_i - d_i)h \qquad \forall t \in T \qquad (3.18)$$

$$c_t \geq -(x_0 + \sum_{i=1}^{T} q_i - d_i)b \qquad \forall t \in T \qquad (3.19)$$

$$q_2 \leq q_2^{max}$$

$$q_t \in \mathbb{R}_+ \qquad \forall t \in T.$$

Suppose the demand vector $\mathbf{d}$ is uncertain but is known to reside in a ball around **5** with radius 5. We will use this uncertainty set both for optimization and for evaluation.

For this small example, it is possible to approximate the worst case costs for an FH approach as a function of the "here and now" decision q_1 as follows. For each q_1 in a range of values, we have randomly drawn 100 uncertain demand vectors from the boundary of the uncertainty set. For each demand vector we have computed the inventory level at the beginning of the second time period ($= x_0 + q_1 - d_1$). Based on this inventory level, we reoptimized the order quantity for the second time period, where d_2 was assumed to reside in the interval $[5 - r, 5 + r]$ with $r = \sqrt{25 - (d_1 - 5)^2}$ (so that the full $\mathbf{d}$ vector is in a ball around **5** with radius 5). Then we computed the total costs over both time periods. The maximum total costs over all 100 demand vectors approximates the worst case costs with the FH approach, and is depicted in Figure 3.1. From this picture it becomes clear that the optimal order quantity for the first time period is approximately 2.3, which has a worst case performance of 10.8.

We have solved the model for the full time horizon with the RC, with the AARC (where c_1 and c_2 are adjustable on the full $\mathbf{d}$, and q_2 is adjustable on d_1), and as a certain problem with $\mathbf{d} = \mathbf{5}$. The nominal solution gives $q_1 = 0$, the RC gives $q_1 \approx 4.4$, while the AARC yields $q_1 \approx 5.3$, leading to worst case costs of the FH approach of 17.8, 14.9 and 16.8, respectively. So, the RC takes the best "here and now" decision with respect to the worst case performance. It may be paradoxical that the AARC yields a worse solution than the RC, since the feasible region of the

Figure 3.1 – Approximation of the total worst case costs for an FH strategy as a function of the initial order quantity q_1.
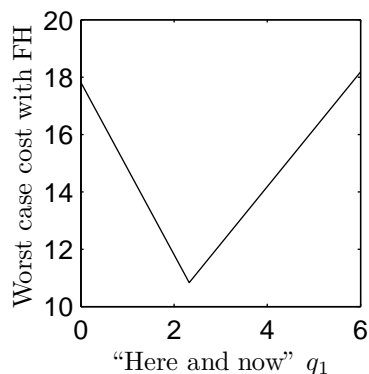
AARC includes the RC solution. However, neither of the two optimize the right objective function. Both approximate the objective value using (static or adjustable) auxiliary variables c_t. While AARC indeed has a better objective value than RC, the solution is not better for the original objective function.

We also perform a comparison on a more realistic problem, which is the retailer-supplier flexible commitment problem by Ben-Tal et al. (2005). At the starting time, the retailer commits himself to ordering certain quantities in later months. These commitments are flexible, i.e. deviations are allowed at a penalty cost. The objective is to minimize the total costs for the retailer, consisting of ordering costs (minus the salvage value), holding costs, backlogging costs, penalty costs associated with deviating from the commitments, and costs for variability in the commitments. The article provides two data sets for twelve time periods, A12 and D12, which we also use in our optimization and comparison.

In this problem the retailer faces an uncertain demand. Following Ben-Tal et al. (2005) we consider box uncertainty where the demand may deviate up to $\rho\%$ around the nominal value. For the simulation we draw demand vectors uniformly from this box region. For these demand vectors the nominal solution, RC and AARC are carried out completely. For the FH approach, the reoptimization is performed after each time period based on previously observed demand. 500 demand vectors were generated for each data set and each uncertainty level ρ, and the same demand vectors were used for all models. In addition, the PH solution was computed for each of these demand vectors.

The simulation results are listed in Tables 3.4 and 3.5. For data set A12, the nominal solution with FH results in the lowest average costs. This means that the nominal solution takes better "here and now" decisions than RC and AARC. Moreover, the RC-FH has lower average costs than the AARC-FH, so also the RC takes

better "here and now" decisions than AARC. The advantage of the nominal FH solution compared to RC-FH and AARC-FH increases when the uncertainty set becomes larger. For data set W12 the nominal solution is the best solution and FH leads to higher mean costs. For this data set, AARC takes significantly better "here and now" decisions than RC. When comparing Nominal-FH with AARC-FH in Table 3.4 it is not immediately clear which solution is better, since the lower mean value of Nominal-FH comes with a larger standard deviation. The Nominal-FH in Table 3.4 is an example of the statement in Section 3.10 that a high standard deviation is not necessarily bad. Its standard deviation is higher than that of AARC-FH but this is due to the high standard deviation of PH, as can be seen from Table 3.5. From this last table, it can be seen that Nominal-FH is strictly better than AARC-FH. For data set A12 and $\rho = 10\%$, it is not clear whether RC-FH outperforms AARC-FH. We now demonstrate the two statistical tests from Section 3.10 on this data set, each based on 100 newly generated uncertainty vectors, to test whether the RC-FH outperforms the AARC-FH. The null hypothesis that both solutions perform equal on average is rejected ($p = 6.1 \cdot 10^{-6}$), and also the null hypothesis that the medians of RC-FH and AARC-FH are equal is rejected ($p = 1.4 \cdot 10^{-10}$). These results show that the AARC is not necessarily better than the RC and support the statement in Section 3.10 that a simulation is required for comparing solutions. As mentioned in Section 3.5, RO may provide solutions that are not Pareto efficient when multiple optimal solutions exist. A different optimal solution to the RC or AARC may yield completely different simulation results, rendering our conclusions useless. This is not the case. We have verified this by solving the problem in two stages. In the first stage we solve the robust counterpart (RC or AARC). For the second stage we take the same problem as in the first stage, but we add the constraint that the robust objective value is not worse than the optimal value from the first stage, and we change the objective in minimizing the costs for the nominal demand trajectory. Thus, we find a solution that is robust optimal and that cannot be improved with respect to the nominal demand trajectory. Moreover, the resulting solution is Pareto optimal; see Iancu and Trichakis (2014). The second stage problem returns the same solution as the first stage problem, so our conclusions are unaffected.

Table 3.4 – Simulated mean (std) costs for the retailer-supplier flexible commitment problem.

	A12		W12	
	$\rho = 10\%$	$\rho = 50\%$	$\rho = 10\%$	$\rho = 50\%$
Nominal	688 (35)	848 (211)	12775 (708)	16163 (3899)
RC	731 (14)	1140 (70)	13656 (169)	20046 (1251)
AARC	702 (5)	1071 (53)	13314 (35)	18575 (192)
Nominal-FH	674 (14)	774 (87)	12869 (296)	16280 (1251)
RC-FH	699 (5)	979 (19)	13615 (125)	19260 (585)
AARC-FH	700 (5)	1027 (21)	13314 (35)	18572 (192)
PH	658 (11)	699 (50)	12194 (204)	12911 (1144)

Table 3.5 – Simulated mean (std) PoU for the retailer-supplier flexible commitment problem.

	A12		W12	
	$\rho = 10\%$	$\rho = 50\%$	$\rho = 10\%$	$\rho = 50\%$
Nominal	30 (29)	149 (178)	581 (575)	3252 (3186)
RC	73 (19)	441 (89)	1463 (308)	7135 (2266)
AARC	44 (9)	372 (74)	1120 (230)	5664 (1298)
Nominal-FH	16 (6)	75 (44)	675 (144)	3369 (717)
RC-FH	41 (9)	280 (49)	1421 (253)	6349 (1550)
AARC-FH	42 (9)	328 (51)	1120 (230)	5661 (1301)

3.12 Conclusion

In this chapter, we have presented a general recipe that shall be helpful for using RO in practice. Additionally, we give several practical insights and hints in applying RO. Examples of such practical insights are: the robust reformulations of equivalent deterministic optimization problems may not be equivalent; in multi-stage optimization problems, re-optimizing the given problem at each stage using static RO or nominal data may outperform solutions provided by ARO; and the actual probability guarantee of an uncertainty set is often higher than the probabilistic guarantee that is approximated by using a safe approximation technique. We also discuss many practical issues to apply RO in a successful and convincing way. Examples are: how to choose the uncertainty set; what is the right interpretation of "RO optimizes for the worst case"; and should the decision rule used in ARO be a function of the final or the primitive uncertainty? Moreover, we propose ideas on how to deal with equality constraints and integer adjustable variables, and on how to compare the robustness characteristics of two solutions. We have provided many numerical examples to illustrate our insights and discussions, and to demonstrate the effectiveness of the usefulness of our hints.

Robust counterparts of inequalities containing sums of maxima of linear functions

Abstract This chapter addresses the robust counterparts of optimization problems containing sums of maxima of linear functions. These problems include many practical problems, e.g. problems with sums of absolute values, and arise when taking the robust counterpart of a linear inequality that is affine in the decision variables, affine in a parameter with box uncertainty, and affine in a parameter with general uncertainty.

In the literature, often the reformulation is used that is exact when there is no uncertainty. However, in robust optimization this reformulation gives an inferior solution and provides a pessimistic view. We observe that in many papers this conservatism is not mentioned. Some papers have recognized this problem, but existing solutions are either conservative or their performance for different uncertainty regions is not known, a comparison between them is not available, and they are restricted to specific problems. We describe techniques for general problems and compare them with numerical examples in inventory management, regression and brachytherapy. Based on these examples, we give recommendations for reducing the conservatism.

4.1   Introduction

Robust Optimization (RO) first appeared in Soyster (1973), and after receiving very little attention in the subsequent decades it has been an active research area since Ben-Tal and Nemirovski (1999) and El Ghaoui and Lebret (1997) started publishing new results in the late nineties. In its basic form, it requires a solution of an optimization problem to be feasible for any realization of the uncertain parameters in

a given uncertainty set. For several choices of the uncertainty region this leads to tractable problems; for instance the robust counterpart (RC) of an LP with polyhedral uncertainty can be reformulated as an LP and the RC of an LP with ellipsoidal uncertainty is a conic quadratic program (CQP) (Ben-Tal et al., 2009a, p. 21).

In RO two distinct formulations that are equivalent in the nonrobust case may have different RCs. This is the case for optimization problems containing the sum of maxima of linear functions, of which the sum of absolute values is a special case ($|x| = \max\{x, -x\}$). These problems arise in inventory management, supply chain management, regression models, tumor treatment, and many other practical situations.

Because RO is applied constraint-wise w.l.o.g. and the objective function can always be formulated as a constraint (Ben-Tal et al., 2009a, p. 11), we focus on the following robust constraint:

$$\ell(\zeta, x) + \sum_{i \in I} \max_{j \in J}\{\ell_{ij}(\zeta, x)\} \leq d \qquad\qquad \forall \zeta \in \mathcal{Z}, \qquad\qquad (4.1)$$

where ℓ and ℓ_{ij} are biaffine functions in the uncertain parameter $\zeta \in \mathbb{R}^L$ and the decision variable $x \in \mathbb{R}^n$, $d \in \mathbb{R}$ is the right hand side, and $\mathcal{Z} \subset \mathbb{R}^L$ is a user-specified uncertainty region, e.g. a box or an ellipsoid. We can assume this uncertainty region to be closed and convex w.l.o.g., as the left hand side is convex in ζ and consequently the worst case is always located at an extreme point of the uncertainty region. In the remainder of this chapter our objective is always to minimize d, but all methods can still be applied when a different objective is used.

In the literature, often the following RC is used, obtained from a reformulation with analysis variables y_i that is exact when there is no uncertainty:

$$\text{(RC-R)} \qquad \ell(\zeta, x) + \sum_{i \in I} y_i \leq d \qquad\qquad \forall \zeta \in \mathcal{Z} \qquad\qquad (4.2)$$

$$y_i \geq \ell_{ij}(\zeta, x) \qquad \forall \zeta \in \mathcal{Z} \quad \forall i \in I \quad \forall j \in J. \qquad (4.3)$$

In this formulation, to which we refer as the RC-R (RC of the reformulation), $y_i \in \mathbb{R}$ is a fixed variable, taking the worst case value of the i^{th} term of the sum. In many cases the terms of the sum do not all reach their worst case value in the same realization of the uncertain parameter ζ, and therefore the RC-R is a conservative reformulation of constraint (4.1), i.e. a solution (x, d) that is feasible for RC-R is also feasible for constraint (4.1), but not necessarily vice versa. However, this reformulation is frequently used without mentioning its conservatism. It has the advantage that the constraints are linear, so that tractable reformulations exist for many uncertainty regions.

Bertsimas and Thiele (2006) use the RC-R for a robust inventory problem which includes the sum of holding and backlogging costs: $\sum_{i=1}^{n} \max\{c_h x_i, -c_b x_i\}$, where x_i is the inventory level at time period i and c_h and c_b are the holding and backlogging costs per time period respectively. Their uncertainty region is the intersection of $\{\zeta : ||\zeta_{1:i}||_1 \leq \Gamma_i\} \quad \forall i$ and $\{\zeta : ||\zeta||_\infty \leq 1\}$, where $\zeta_{1:i}$ is the vector consisting of the first i elements of ζ, and Γ_i is not necessarily integer (budgeted uncertainty). Their formulation allows that the values of different analysis variables take values based on different realizations of the uncertain parameter. The same uncertainty region is also used by Bertsimas and Thiele (2004) for a supply chain model, by Thiele (2004) for supply chains and revenue management, by Alem and Morabito (2012) for a production planning problem under uncertain demand, and by Wei et al. (2009) for a slightly more complicated inventory model in which items can be returned and remanufactured or disposed. With the exception of (Thiele, 2004, p.37), none of these papers mentions the conservatism of the formulation.

Ng et al. (2010) treat a lot allocation problem where each order is assigned to one or more production locations before the production capacity of each location is fully known. Every production location is assigned to at most one order. Their uncertainty region is an ellipsoid over all production locations, while an analysis variable is used for every order, so their formulation is conservative.

Kropat and Weber (2008) consider a robust linear cluster regression model using the sum of absolute values with polyhedral and ellipsoidal uncertainty regions. They introduce an analysis variable for every absolute value, which is conservative because every uncertain parameter affects two absolute values.

Ben-Tal et al. (2005) solve a multi-period inventory problem and use affine decision rules for the actual decisions (AARC). For every time period there are analysis variables indicating the costs in that time period. These costs are given by the maximum of holding costs and shortage costs, which both are linear functions of past demand. They note that an analysis variable should therefore be replaced with a function being the maximum of the two linear functions. Because that would lead to a very complicated robust counterpart, they replace the analysis variables with linear decision rules instead, which is conservative. Replacing analysis variables with linear decision rules was also done by Ben-Tal et al. (2009b) and Ben-Tal et al. (2011).

Bienstock and Özbay (2008) were the first to identify and eliminate the conservatism of the RC-R. The idea behind their solution is that it suffices to make constraint (4.1) hold for just the vertices of the uncertainty region. The constraint then also holds for all other elements in the (convex) uncertainty region, because the constraint is convex in ζ. They generalize it to a cutting plane method that only adds a subset of the vertices, which they successfully tested for computing basestock

levels under budgeted uncertainty. It is not yet known how well this cutting plane method performs on other problems or different uncertainty regions.

In this chapter, we make the following contributions. First, we identify that a conservative reformulation is often used in the literature without mentioning its conservatism. Second, we make a classification of solution approaches. Third, we present a new cutting plane method (Algorithm 3). Fourth, we show that a linear constraint with biaffine uncertainty is a special case. This provides a method to solve or approximate a conic quadratic constraint with box uncertainty, for which currently no efficient methods are available. Fifth, we demonstrate all approaches on several numerical examples. Such a comparison is currently not available. Sixth, based on our numerical examples, we give recommendations for reformulating inequalities containing sums of maxima of linear functions.

The structure of this chapter is as follows. Since it is not possible to compare the objective value of different formulations in RO, we introduce a new performance number which is independent of the reformulation in Section 4.2. In the examples we consider, this performance number is the slack in constraint (4.1). We provide an overview of exact (non-conservative) formulations, approximations, and cutting plane methods in Section 4.3. In that section we show similarities between different methods, and also show how several approaches can be combined. The application scope of this chapter is extended in Section 4.4, where we show that our methods can be applied to an uncertain conic quadratic constraint with box uncertainty. We evaluate the methods in Section 4.5 on some small toy problems and three larger problems. We give conclusions in Section 4.6, and show how the robust counterparts in the aforementioned papers could be improved.

4.2   The true robust value

This section explains how two solutions from different reformulations can be compared. For a minimization problem containing sums of maxima of linear functions, we define the true robust value of a solution to be the maximum over ζ in $\mathcal{Z}$ of the unreformulated problem with all decision variables fixed. We have assumed that in this chapter our objective is always to minimize d, an analysis variable at the right hand side of (4.1), so in our case the true robust value of a solution x is:

$$v_{true}(x) = \max_{\zeta \in \mathcal{Z}} \left\{ \ell(\zeta, x) + \sum_{i \in I} \max_{j \in J} \{\ell_{ij}(\zeta, x)\} \right\}. \tag{4.4}$$

Determining this value is a difficult problem, because it requires the maximization of a convex function over a convex set. The global maximum over $\zeta \in \mathcal{Z}$ of the sum of

maxima of linear functions does not necessarily coincide with a maximum over $\zeta \in \mathcal{Z}$ of one of those linear functions. One way to obtain the exact value is by solving the following optimization problem with integer variables for fixed x:

$$
\max \quad \ell(\zeta, x) + \sum_{i \in I} y_i \tag{4.5}
$$

$$
\text{s.t.} \quad y_i \leq \ell_{ij}(\zeta, x) + M(1 - z_{ij}) \qquad \forall i \in I \quad \forall j \in J
$$

$$
\sum_{j \in J} z_{ij} = 1 \qquad \forall i \in I
$$

$$
\zeta \in \mathcal{Z}, y \in \mathbb{R}^{|I|}, z \in \{0, 1\}^{|I| \times |J|},
$$

where M is a sufficiently large number. This problem is an MILP for polyhedral uncertainty, and an MIQCP for ellipsoidal uncertainty.

Another way of obtaining the exact value is by considering $|J|^{|I|}$ linear optimization problems, e.g. for $I = \{1, 2\}$ and $J = \{1, 2\}$ consider the following four optimization problems:

$$
\max_{\zeta \in \mathcal{Z}} \ell(\zeta, x) + \ell_{1,1}(\zeta, x) + \ell_{2,1}(\zeta, x), \tag{4.6}
$$

$$
\max_{\zeta \in \mathcal{Z}} \ell(\zeta, x) + \ell_{1,1}(\zeta, x) + \ell_{2,2}(\zeta, x),
$$

$$
\max_{\zeta \in \mathcal{Z}} \ell(\zeta, x) + \ell_{1,2}(\zeta, x) + \ell_{2,1}(\zeta, x), \text{ and}
$$

$$
\max_{\zeta \in \mathcal{Z}} \ell(\zeta, x) + \ell_{1,2}(\zeta, x) + \ell_{2,2}(\zeta, x).
$$

Each of these problems is easy, because the maximum of an affine function over a box or an ellipsoid can be computed in a few operations. The true robust value is the largest value of the computed maxima.

If determining $v_{true}(x)$ is intractable, bounds can still be obtained. Filling in any ζ from the set $\mathcal{Z}$ at the right hand side of equation (4.4), for instance the nominal value, gives a lower bound. If the dimension of ζ is small, the bound can be improved with global optimization techniques. Upper bounds can be obtained by fixing x in any conservative reformulation (such as those mentioned in Section 4.3.2) of constraint (4.1) and optimizing over the other variables.

4.3 Solution approaches

This section lists exact solution approaches and approximation methods, most of which can be applied to general RO problems containing the sum of maxima of linear functions. Many of these methods have been used before, but this full classification is new. This allows us to show the similary between some methods, and to show how approximations can be combined with exact methods.

4.3.1   Exact reformulations

4.3.1.1   Vertex enumeration

Vertex enumeration is an exact solution method first used for a RO problem containing the sum of maxima by Bienstock and Özbay (2008) that is powerful especially when the uncertainty region has a small number of vertices. Let V denote the finite set of vertices, and consider the following reformulation of constraint (4.1):

$$\text{(Vertex enumeration)} \qquad \ell(\zeta, x) + \sum_{i \in I} y_i^\zeta \leq d \qquad \forall \zeta \in V \qquad (4.7)$$

$$y_i^\zeta \geq \ell_{ij}(\zeta, x) \quad \forall i \in I \quad \forall j \in J \quad \forall \zeta \in V. \qquad (4.8)$$

Constraint (4.8) is no longer a semi-infinite constraint, because V is a finite set. This reformulation is exact because the left hand side of constraint (4.1) is convex in ζ, and a convex function takes its maximum at an extreme point of its domain.

4.3.1.2   Enumeration of robust linear constraints

The RC-R is inexact because it has analysis variables that may take values corresponding to different worst case scenarios. A constraint with a single max function does not suffer from this problem, because an equivalent set of linear constraints can be formulated without analysis variables. An exact reformulation of RC (4.1) can be obtained by first rewriting it as a constraint with a single $\max\{\cdot\}$ function by enumerating all combinations, and then applying RO to the reformulation:

$$\text{(EORLC)} \qquad \ell(\zeta, x) + \sum_{i \in I} \ell_{i,j(i)}(\zeta, x) \leq d \qquad \forall j(i) \in J \qquad \forall \zeta \in \mathcal{Z}.$$

We call this the enumeration of robust linear constraints (EORLC) formulation. It has the advantage that the constraints are linear, so that tractable reformulations exist for many uncertainty regions. For example, with $I = \{1, 2\}$ and $J = \{1, 2\}$, the EORLC formulation has the following constraints:

$$\ell(\zeta, x) + \ell_{1,1}(\zeta, x) + \ell_{2,1}(\zeta, x) \leq d \qquad \forall \zeta \in \mathcal{Z}$$
$$\ell(\zeta, x) + \ell_{1,1}(\zeta, x) + \ell_{2,2}(\zeta, x) \leq d \qquad \forall \zeta \in \mathcal{Z}$$
$$\ell(\zeta, x) + \ell_{1,2}(\zeta, x) + \ell_{2,1}(\zeta, x) \leq d \qquad \forall \zeta \in \mathcal{Z}$$
$$\ell(\zeta, x) + \ell_{1,2}(\zeta, x) + \ell_{2,2}(\zeta, x) \leq d \qquad \forall \zeta \in \mathcal{Z}.$$

A similar formulation is also given by Bienstock and Özbay (2008), where it was neglected for its exponential size. While the number of constraints $|J|^{|I|}$ indeed grows exponentially with the number of terms in the summation, it is effective for

small $|I|$. There are situations in which I is small indeed, for instance with a planning horizon of up to ten periods.

EORLC has a strong relation to vertex enumeration that was not observed before to the best of our knowledge. In fact, EORLC is vertex enumeration on a different set. Constraint (4.1) can be formulated as:

$$\ell(\zeta, x) + \sum_{i \in I} \sum_{j \in J} \lambda_{ij} \ell_{ij}(\zeta, x) \leq d \qquad \forall \lambda_i \in \Delta^{|J|-1} \quad \forall \zeta \in \mathcal{Z},$$

where $\Delta^{|J|-1} = \{\lambda_i \in \mathbb{R}^{|J|} : \sum_{j \in J} \lambda_{ij} = 1, \lambda_{ij} \geq 0\}$ (the standard simplex in $\mathbb{R}^{|J|}$). The vertices of the simplex are given by unit vectors. If λ_i is a unit vector, $\sum_{j \in J} \lambda_{ij} \ell_{ij}(\zeta, x)$ simplifies to a single $\ell_{ij}(\zeta, x)$. It follows that vertex enumeration on the $|I|$ simplices gives the EORLC reformulation.

EORLC can benefit from two preprocessing steps in order to reduce the final number of constraints. First, every $\max\{\cdot\}$ term should contain at most one function that does not depend on ζ. If it has more than one, those functions can all be replaced with a single analysis variable. Second, inequalities of the form $a \max\{0, \ell(\zeta, x)\} + b \max\{0, -\ell(\zeta, x)\} \leq d$ with $a, b > 0$, which are often used for holding and backlogging costs, should be reformulated as $\max\{a\ell(\zeta, x), -b\ell(\zeta, x)\} \leq d$.

4.3.1.3 Cases with special structure

There are several special cases of (4.1) that allow an exact reformulation. This section lists a few general cases. More specific cases can be found in (Ben-Tal et al., 2009a, Ch. 12.2) and in Xu et al. (2009).

The first case is when the uncertainty region $\mathcal{Z}$ is the direct product of sets $\mathcal{Z}_i$, where term i in the left hand side of constraint (4.1) is only affected by $\mathcal{Z}_i$. The RC-R is exact because all analysis variables can take their worst case values simultaneously.

The second case is when the inequality contains the sum of absolute values of linear functions of ζ_i and x:

$$\ell(\zeta, x) + \sum_{i \in I} |\alpha_i(x) + \beta_i(x)^\top \zeta| \leq d \qquad \forall \zeta \in \mathcal{Z}, \tag{4.9}$$

where $\alpha_i : \mathbb{R}^n \to \mathbb{R}$ and $\beta_i : \mathbb{R}^n \to \mathbb{R}$ are linear functions, the components of β_i that may be nonzero for one i are zero for all other i in I, and the uncertainty region $\mathcal{Z}$ is centrosymmetric around $\zeta = 0$, i.e. $\mathcal{Z}$ is closed under changing the sign of one or more vector elements (a box and an ellipsoid are examples of such sets). Note that the assumption that the symmetry is around 0 is made w.l.o.g. The following constraint is equivalent to (4.9):

$$\ell(\zeta, x) + \sum_{i \in I} \{|\alpha_i(x)| + \beta_i(x)\zeta_i\} \leq d \qquad \forall \zeta \in \mathcal{Z}. \tag{4.10}$$

Equivalence is readily checked by conditioning on the sign of $\alpha_i(x)$. The formulation for (4.10) where absolute values are replaced with analysis variables is equivalent to (4.10), since the analysis variables do not depend on ζ.

The third case is when, for a fixed $i \in I$, each of linear functions under the maximum is the sum of a common nonnegative linear function of ζ and a linear function of x:

$$\ell(\zeta, x) + \sum_{i \in I} \max_{j \in J}\{\alpha_i(\zeta) + \beta_i(\zeta)\ell_{ij}(x)\} \leq d \qquad \forall \zeta \in \mathcal{Z},$$

where $\alpha_i : \mathbb{R}^L \to \mathbb{R}$ and $\beta_i : \mathbb{R}^L \to \mathbb{R}_+$ are linear functions. The common functions of ζ can be placed outside the $\max\{\cdot\}$ expression:

$$\ell(\zeta, x) + \sum_{i \in I} \left[\alpha_i(\zeta) + \beta_i(\zeta) \max_{j \in J}\{\ell_{ij}(x)\}\right] \leq d \qquad \forall \zeta \in \mathcal{Z},$$

and the RC-R of this constraint is exact. If the range of β_i is $\mathbb{R}$ instead of $\mathbb{R}_+$, then $\max_{j \in J}\{\ell_{ij}(x)\}$ should be $\min_{j \in J}\{\ell_{ij}(x)\}$ when $\beta_i(\zeta) < 0$. This can be modeled as follows. Given a subset $I_+ \subseteq I$, we define a set which consists of those ζ for which $\beta_i(\zeta) \geq 0$ for $i \in I_+$, and $\beta_i(\zeta) \leq 0$ for $i \in I \backslash I_+$:

$$\mathcal{Z}(I_+) = \mathcal{Z} \cap \{\zeta : \beta_i(\zeta) \geq 0 \quad \forall i \in I_+, \quad \beta_i(\zeta) \leq 0 \quad \forall i \in I \backslash I_+\}.$$

Note that

$$\mathcal{Z} = \bigcup_{I_+ \subseteq I} \mathcal{Z}(I_+).$$

The constraint can now be written as:

$$\ell(\zeta, x) + \sum_{i \in I} \alpha_i(\zeta) + \sum_{i \in I_+} \beta_i(\zeta) \max_{j \in J}\{\ell_{ij}(x)\} + \sum_{i \in I \backslash I_+} \beta_i(\zeta) \min_{j \in J}\{\ell_{ij}(x)\} \leq d$$

$$\forall \zeta \in \mathcal{Z}(I_+) \; \forall I_+ \subseteq I.$$

The number of constraints is $2^{|I|}$ (one for each $I_+ \subseteq I$), which is less than the $|J|^{|I|}$ constraints obtained with EORLC. The max and min expressions do not depend on ζ, so the reformulation with analysis variables is exact. Each constraint is still convex despite the min expressions, because their coefficients $\beta_i(\zeta)$ are negative.

4.3.2   Conservative approximations

The RC-R (4.2)-(4.3) is a conservative approximation to (4.1). We discuss how the conservatism can be decreased, and also present a new method.

The variables y_i in the RC-R have been introduced for modeling the $\max\{\cdot\}$ expression. We do not need to know their values because they do not correspond to a "here and now" decision. Only x has to be known for implementing a solution. The values of y_i may be adjusted according to the realization of the uncertain parameter ζ as long as the constraints hold for every realization of ζ in the perturbation set (Adjustable RC). As first applied by Ben-Tal et al. (2005) for a multi-stage problem and later also done by Ben-Tal et al. (2009b) and Ben-Tal et al. (2011), we can make y_i an affine function of ζ, which leads to the Affinely Adjustable RC of the reformulation (AARC-R). After substituting $y_i = v_i + w_i^\top \zeta$ (with decision variables $v_i \in \mathbb{R}$ and $w_i \in \mathbb{R}^L$), constraints (4.2)-(4.3) become:

$$\text{(AARC-R)} \qquad \ell(\zeta, x) + \sum_{i \in I} \left(v_i + w_i^\top \zeta \right) \leq d \qquad \forall \zeta \in \mathcal{Z}$$

$$v_i + w_i^\top \zeta \geq \ell_{ij}(\zeta, x) \quad \forall \zeta \in \mathcal{Z} \quad \forall i \in I \quad \forall j \in J.$$

This substitution gives a less conservative reformulation, while the robust counterpart is often still tractable, because robust linear constraints are tractable for a wide class of uncertainty regions. The power of the AARC-R can be increased by lifting the uncertainty region to a higher dimension (Chen and Zhang, 2009). If ℓ_{ij} does not depend on one or more components of ζ for all j for some fixed i, the computational complexity can seemingly be reduced by making y_i a function of only those components of ζ that appear in term i for one or more j. However, it is easy to construct an example where this reduction introduces more conservatism.

There are two different approaches that also lead to the formulation AARC-R. The first approach is considering the Fenchel dual problem of maximizing the left hand side of constraint (4.1) over ζ in $\mathcal{Z}$. We give the full derivation and a proof of equivalence to the AARC-R in Appendix 4.A. The other approach is derived in Appendix 4.B.

An approach that is less conservative than an affine decision rule, is a quadratic decision rule:

$$y_i = v_i + w_i^\top \zeta + \zeta^\top W_i \zeta,$$

where $v_i \in \mathbb{R}$, $w_i \in \mathbb{R}^L$, and $W_i \in \mathbb{R}^{L \times L}$ are new analysis variables. This is called a Quadratically Adjustable RC of the reformulation (QARC-R) which is known to be tractable for ellipsoidal uncertainty and (under some restrictions) for box uncertainty. A deterministic reformulation of the QARC-R with ellipsoidal uncertainty is given in Appendix 4.C, resulting in $|I||J| + 1$ LMIs of size $|L + 1|$. For box uncertainty, when the quadratic terms are restricted such that each element of ζ is multiplied with itself and at most one other element of ζ, we can use the result by Yanıkoğlu et al. (2012) to write the RC as an SDP with $(|I||J| + 1)\lceil L/2 \rceil$ variable matrices of size three.

A new way to reduce the conservatism of the RC-R is by first combining several max expressions before reformulating. In order to do this, partition I into $|G|$ groups: $I = \bigcup_{g \in G} I_g$, where the groups are mutually disjoint: $I_{g_1} \cap I_{g_2} = \emptyset$ for $g_1, g_2 \in G, g_1 \neq g_2$, and partition in such a way that $|I_g|$ is small for all g. Now introduce analysis variables for all g:

$$\ell(\zeta, x) + \sum_{g \in G} y_g \leq d \qquad\qquad \forall \zeta \in \mathcal{Z}$$

$$y_g \geq \sum_{i \in I_g} \max\{\ell_{ij}(\zeta, x)\} \qquad\qquad \forall \zeta \in \mathcal{Z} \quad \forall g \in G.$$

Because the cardinality of I_g is small, the sum of max expression within each constraint can be transformed into a single max (EORLC). Each constraint in this reformulation is therefore tractable and can be solved exactly, but conservatism still comes from the analysis variables y_g. These analysis variables can also be written as a linear, quadratic or more general function of ζ.

4.3.3   Cutting plane methods

In this section we describe two cutting plane methods. The first one is based on vertex enumeration and was used in Bienstock and Özbay (2008), the other one is new and based on EORLC.

Vertex enumeration results in very large problems if there are many vertices. For box uncertainty, the number of vertices grows exponentially in the time horizon. Also for budgeted uncertainty, which is described in Section 4.1, the number of vertices quickly becomes very large. The cutting plane method outlined in Algorithm 2 adds only a subset of the vertices.

Algorithm 2 Cutting plane method based on vertex enumeration

Require: A linear program LP with constraints (4.7)
 1: $V := \{\zeta^{nom}\}$ (the nominal value)
 2: $k := 0$
 3: **repeat**
 4:      $k := k + 1$
 5:      Solve LP
 6:      Let x^* be the minimizer of LP, and LB be the value of d in LP
 7:      Let $UB := v_{true}(x^*)$, and ζ^k be its maximizer
 8:      $V := V \cup \{\zeta^k\}$
 9: **until** $UB - LB < \varepsilon$

While the algorithm is running, LB is a lower bound on the optimal value of d in (4.7) because it is the value of a relaxation of the original problem, and UB is an upper bound on the optimal d, because it is the maximum for some x and not for the optimal x. The difference $UB - LB$ indicates the current violation of the constraint. If ε is set to a larger value, the algorithm terminates quicker but does not give the optimal solution. If $\varepsilon = 0$ and the algorithm terminates, the final solution is robust feasible and robust optimal. It is possible that this algorithm still enumerates all constraints, but we have not encountered problems for which this is the case. Bienstock and Özbay (2008) find that for their problem the number of iterations does not increase when the time horizon is increased, and is around four on average even for 150 time periods. The number of vertices of their uncertainty region cannot be deducted from their report. In our numerical examples we often find a larger number of iterations. The same cutting plane method was used by Bohle et al. (2010), again for budgeted uncertainty, but solution times and the number of generated constraints are not reported. It is still an open question how well this method works on other problems, and for non-polyhedral uncertainty regions.

The most time consuming step is determining the true value, since this involves the maximization of a convex function. As pointed out by Bienstock and Özbay (2008), it is not necessary to find the optimal solution. It suffices to find any ζ for which the maximization problem has a larger objective value than LB, because it corresponds to a violated constraint in the relaxation, and adding that constraint strengthens the relaxation. Bienstock et al. report that also the problem LP does not need to be solved to optimality. Optimization can stop as soon as the objective value is less than UB, because at that point the current solution is already better than the solution in the previous run.

The second cutting plane method is new and based on EORLC. If $|I|$ becomes too large for EORLC, this cutting plane method adds only a subset of the robust linear constraints. It starts with the nominal problem and adds new robust constraints until the solution is robust feasible (Algorithm 3).

Similar to Algorithm 2, a lower and upper bound are reported and the stopping criterion can be adjusted. Also for this algorithm, we have not encountered numerical examples in which all constraints are enumerated. The main difference between this algorithm and Algorithm 2 is the constraint that is added in every iteration.

Just as with the cutting plane method based on vertex enumeration, determining the true value is the most time consuming step, and also here the optimization problem for determining the true value does not need to be solved to optimality as long as the value is larger than LB. Also LP does not need to be solved to optimality.

Both cutting plane methods can easily be combined so that two sets of constraints

---

Algorithm 3 Cutting plane method based on EORLC

Require: A linear program LP with constraints (4.7)
1: $V := \{\zeta^{nom}\}$ (the nominal value)
2: $k := 0$
3: **repeat**
4:      $k := k + 1$
5:      Solve LP
6:      Let x^* be the minimizer of LP, and LB be the value of d in LP
7:      Let $UB := v_{true}(x^*)$, and ζ^k be its maximizer
8:      Let j_i^k be the maximizer of $\max_{j \in J} \ell_{ij}(\zeta^k, x)$
9:      Add the following robust constraint to LP:

$$\ell(\zeta, x) + \sum_{i \in I} \ell_{i, j_i^k}(\zeta, x) \leq d \qquad \forall \zeta \in \mathcal{Z}$$

10: **until** $UB - LB < \varepsilon$

---

are added per iteration. This is done by inserting line 8 of Algorithm 2 in between lines 9 and 10 of Algorithm 3.

The stopping criterion in both algorithms is $UB - LB < \varepsilon$. It is also possible to use a dimensionless stopping criterion such as $2(UB - LB)/(1 + |UB + LB|) < \varepsilon$. This stopping criterion is based on the relative gap.

4.4   RC of a linear constraint with biaffine uncertainty

Constraint (4.1) may appear in RO itself when a constraint has biaffine uncertainty, and the uncertainty region of one parameter is a box. In general, such a constraint can be written as:

$$\tilde{\ell}(\zeta^{(1)}, \zeta^{(2)}, x) \leq d \qquad \forall \zeta^{(1)} : \left\| \zeta^{(1)} \right\|_\infty \leq \rho \quad \forall \zeta^{(2)} \in \mathcal{Z}, \tag{4.11}$$

where $\tilde{\ell} : \mathbb{R}^{|I|} \times \mathbb{R}^L \times \mathbb{R}^n \to \mathbb{R}$ is a triaffine function, ρ is the radius of the box, and $\mathcal{Z}$ is the uncertainty region of $\zeta^{(2)}$. Constraints with biaffine uncertainty have never been investigated before to the best of our knowledge. To show the equivalence to constraint (4.1), choose ℓ, ℓ_i and I in such a way that constraint (4.11) is equivalent to:

$$\ell(\zeta^{(2)}, x) + \sum_{i \in I} \zeta_i^{(1)} \ell_i(\zeta^{(2)}, x) \leq d \qquad \forall \zeta^{(1)} : \left\| \zeta^{(1)} \right\|_\infty \leq \rho \quad \forall \zeta^{(2)} \in \mathcal{Z},$$

and maximize the left hand side with respect to $\zeta^{(1)}$:

$$\ell(\zeta^{(2)}, x) + \sum_{i \in I} \rho |\ell_i(\zeta^{(2)}, x)| \leq d \qquad \forall \zeta^{(2)} \in \mathcal{Z}, \tag{4.12}$$

which is indeed in the form of constraint (4.1). This RC includes many practical problems, of which we give four examples.

The first example is a constraint $a^\top x \leq d$ where each element of a has been estimated using a regression model with the same regressors ζ for every element: $a_i = \beta_i^0 + (\beta^{(i)})^\top \zeta + \varepsilon_i$ ($\beta_i^0 \in \mathbb{R}, \beta^{(i)} \in \mathbb{R}^L, \varepsilon_i \in \mathbb{R}$), in that case the constraint is:

$$(\beta^0 + B\zeta + \varepsilon)^\top x \leq d.$$

We assume that the coefficients $\beta^0 \in \mathbb{R}^n$ and $B \in \mathbb{R}^{n \times L}$, the matrix having $\beta^{(i)}$ as rows, have been estimated with some error, and also that the regressors ζ are not fully known. The constraint can be written as constraint (4.11) if β^0, B, ζ and ε lie in some specified uncertainty region, and the uncertainty of either B or ζ is a box.

The second example appears in Adjustable RCs. Consider a robust constraint:

$$\ell(\zeta^{(1)}, \zeta^{(2)}, x) + b^\top y \leq d \qquad \forall \zeta^{(1)} : \left\| \zeta^{(1)} \right\|_\infty \leq \rho \quad \forall \zeta^{(2)} \in \mathcal{Z},$$

where $y \in \mathbb{R}^{m_1}$ is an adjustable variable that represents a "wait and see" decision that can be made after $\zeta^{(1)}$ and $\zeta^{(2)}$ are (partially) observed, and b is a vector of known coefficients (fixed recourse). Determining the true optimal policy for y as a function of $\zeta^{(1)}$ and $\zeta^{(2)}$ is often intractable, which is why a suboptimal y is often determined by a parameterization, i.e. y is written as a function of which the coefficients are decision variables. The first parameterization in RO was proposed by Ben-Tal et al. (2004) who proposed an affine decision rule, which results in a problem which is in the same class (LP, CQP or SDP, depending on the uncertainty region) as the problem with y as a "here and now" decision variable. This was later extended to a quadratic decision rule with an ellipsoidal uncertainty region, resulting in an SDP (Ben-Tal et al. (2009a)), and to a polynomial decision rule of arbitrarily large degree restricted to uncertainty regions described by polynomial inequalities, resulting in a conservative SDP (Bertsimas et al. (2012)). The latter includes the biaffine decision rule $y = \ell'(\zeta^{(1)}, \zeta^{(2)}, v)$, where ℓ' is a function $\mathbb{R}^{|I|} \times \mathbb{R}^L \times \mathbb{R}^{m_2} \to \mathbb{R}^{m_1}$, and v is a vector of coefficients to be determined by the model. This decision rule could be very useful if a problem is affected by several sources of uncertainty. The advantage over affine decision rules is that the former includes cross terms of $\zeta^{(1)}$ and $\zeta^{(2)}$. Applying the results of Bertsimas et al. (2012) gives an SDP which is not only conservative, but also has a potentially large instance size. E.g. if $|I|$ (the dimension of $\zeta^{(1)}$) is very small but $\zeta^{(2)}$ is in a box of (large) dimension L, the conservative SDP has at

least one variable matrix of size $(|I|+1)L + |I| + 1$ and over $2|L|$ smaller matrices, while our result is exact and gives a practically solvable LP. If instead $\zeta^{(2)}$ is in an ellipsoid of (large) dimension L, the SDP still has the large variable matrix of size $(|I|+1)L + |I| + 1$ while our exact result gives a CQP for which efficient solvers are available.

The third example is a constraint with unknown coefficients and implementation error. Consider the following constraint:

$$\ell(\zeta^{(1)})^\top x \le d \qquad \forall \zeta^{(1)} \in \mathcal{Z}_1, \tag{4.13}$$

where ℓ is a vector of n functions that are linear in the uncertain parameter $\zeta^{(1)}$. Now suppose that there is implementation error in x, i.e. instead of x we implement a vector of which component i is given by $x_i + \zeta_i^{(2)}$ (additive implementation error) or by $\zeta_i^{(2)} x_i$ (multiplicative implementation error). After substituting x, constraint (4.13) becomes:

$$\ell(\zeta^{(1)})^\top (x + \zeta^{(2)}) \le d \qquad \forall \zeta^{(1)} \in \mathcal{Z}_1 \quad \forall \zeta^{(2)} \in \mathcal{Z}_2,$$

in case of additive implementation error, and:

$$\sum_{i=1}^n \ell_i(\zeta^{(1)})\zeta_i^{(2)} x_i \le d \qquad \forall \zeta^{(1)} \in \mathcal{Z}_1 \quad \forall \zeta^{(2)} \in \mathcal{Z}_2,$$

in case of multiplicative implementation error. Both constraints are special cases of constraint (4.11) if either $\mathcal{Z}_1$ or $\mathcal{Z}_2$ is a box.

The fourth example is the following robust constraint with box uncertainty:

$$\ell(\zeta^{(1)}, x) + \sum_{k \in K} \left\| \ell_k(\zeta^{(1)}, x) \right\|_l \le d \qquad \forall \zeta^{(1)} : \left\| \zeta^{(1)} \right\|_\infty \le \rho, \tag{4.14}$$

where ℓ_k is a vector of L linear functions, and l equals 1, 2 or ∞. In order to see that this constraint is equivalent to constraint (4.11), note that constraint (4.14) is a reformulation of the following constraint:

$$\ell(\zeta^{(1)}, x) + \sum_{k \in K} (\zeta_k^{(2)})^\top \ell_k(\zeta^{(1)}, x) \le d \quad \forall \zeta^{(1)} : \left\| \zeta^{(1)} \right\|_\infty \le \rho$$

$$\forall \zeta^{(2)} : \left\| \zeta_k^{(2)} \right\|_l^* \le 1 \quad \forall k \in K,$$

where $\|\cdot\|_l^*$ is the dual norm, and $\zeta_k^{(2)}$ is a vector in $\mathbb{R}^L$ for all k in K. We will show how the results in this chapter can be used for solving the robust constraint (4.14) for different choices of l.

For $l = 2$ and $|K| = 1$, constraint (4.14) is a robust conic quadratic constraint, for which the RC is known only in special cases, one of which is when the vertices

of the uncertainty region can be enumerated (Ben-Tal et al., 2009a, p. 159). The case $|K| > 1$ has not been covered yet. A (conservative) reformulation with analysis variables for every $\left\|\ell_k(\zeta^{(1)}, x)\right\|_2$ reduces the problem to a problem with one linear constraint and $|K|$ robust conic quadratic constraints, each of which is of the form (4.14) with $|K| = 1$, so it can be reformulated using vertex enumeration. A different approach, that is not only exact but also results in a smaller problem than obtained by using analysis variables, is to apply vertex enumeration to constraint (4.14) directly. Vertex enumeration is exact because the constraint is convex in $\zeta^{(1)}$. If the dimension of the box is very large, (iterative) vertex enumeration is no longer tractable, so even the case $|K| = 1$ becomes unsolvable, and tractable conservative reformulations are not known. Our reformulation (4.12) allows the use of all approaches from Section 4.3.

For $l = 1$ constraint (4.14) is a special case of constraint (4.1), and we know from Section 4.3 how to solve it or how to find a conservative reformulation. The reformulation may be useful if it is solved with the RC-R, the AARC-R or the QARC-R, because the resulting formulation is very different. Vertex enumeration and EORLC are not useful on the reformulation (4.12), because they correspond with EORLC and vertex enumeration on the original constraint, respectively.

For $l = \infty$ constraint (4.14) is a special case of constraint (4.1) with $I = K$ and $J = \{1, 2, .., L\}$, and we know from Section 4.3 how to solve it or how to find a conservative reformulation. The reformulation (4.12) may be useful if it solved with the RC-R, the AARC-R or the QARC-R, because the resulting formulation is very different. The reformulation allows to do vertex enumeration on $\zeta^{(2)}$, which may be faster than vertex enumeration in the original constraint (which is done on the vertices of $\zeta^{(1)}$) if L and $|K|$ are small relative to $|I|$.

For all three cases of l, it holds that when EORLC is used on the reformulation (4.12), the resulting constraints are the same ones resulting from vertex enumeration on constraint (4.14). This implies that reformulating (4.12) is a redundant step if it is followed by EORLC.

4.5 Numerical examples

The LP, MILP, CQP and MICQP problems in this chapter have been solved with AIMMS 3.11 FR3 x32 with ILOG CPLEX 12.1 unless stated otherwise. SDP problems have been modeled with CVX (Grant and Boyd (2010)) and solved with SDPT3 (YALMIP (Löfberg (2012)) would nowadays be a better choice as it allows the user to specify a linear constraint with quadratic uncertainty with an ellipsoidal uncertainty region directly). Reported computing times have been obtained under Windows XP

SP3 on an Intel Core2 Duo E6400 (2.13 GHz) processor and 2 GB of RAM.

4.5.1   Computing the true robust value

We have run some experiments to determine how quickly the true robust value can be computed using the two different methods from Section 4.2. Using the problem with integer variables (4.5), the worst case ζ can be determined in less than a minute for $|I| = 50$ and $|J| = 3$ for box uncertainty of dimension 50, and in 7 seconds for $|I| = 20$ and $|J| = 2$ for ellipsoidal uncertainty of dimension 20. For testing the performance of maximizing the $|J|^{|I|}$ linear optimization problems (4.6), we have created a single threaded C++ application that creates an affine function with coefficients randomly taken from the interval $[-100, 100]$, maximizes that function, randomly selects new coefficients, et cetera. Only the running time of the maximization step is measured. It can maximize 1.5×10^8 affine functions $f : \mathbb{R}^{50} \to \mathbb{R}$ per minute over a box, which is 10^{23} times slower than the MILP. It can maximize 3×10^8 affine functions $g : \mathbb{R}^{20} \to \mathbb{R}$ per minute over an ellipsoid, which is 32 times faster than the MIQCP.

4.5.2   Illustrative small problems

Consider the following toy optimization problem:

$$
\begin{array}{llll}
\text{(TOY1)} & \min & d & \\
& \text{s.t.} & d \geq \max\{x, x + \zeta\} + \max\{x, x - \zeta\} & \forall \zeta \in [-1, 1] \\
& & d \in \mathbb{R}, x \in \mathbb{R}_+.
\end{array}
$$

The optimal value of this problem is 1 ($x = 0, d = 1$). If we model this problem as an LP and then apply RO, we get the following model:

$$
\begin{array}{llll}
\text{(RC-R)} & \min & y_1 + y_2 & \\
& \text{s.t.} & y_1 \geq x & \forall \zeta \in [-1, 1] \\
& & y_1 \geq x + \zeta & \forall \zeta \in [-1, 1] \\
& & y_2 \geq x & \forall \zeta \in [-1, 1] \\
& & y_2 \geq x - \zeta & \forall \zeta \in [-1, 1] \\
& & y_1, y_2 \in \mathbb{R}, x \in \mathbb{R}_+.
\end{array}
$$

The optimal value of this model is 2 ($x = 0, y_1 = y_2 = 1$), so the RC-R is not exact. If we substitute $y_i = v_i + w_i\zeta$, the model becomes:

$$
\begin{aligned}
\text{(AARC-R)} \quad &\min && y \\
&\text{s.t.} && y \geq v_1 + w_1\zeta + v_2 + w_2\zeta && \forall \zeta \in [-1, 1] \\
& && v_1 + w_1\zeta \geq x && \forall \zeta \in [-1, 1] \\
& && v_1 + w_1\zeta \geq x + \zeta && \forall \zeta \in [-1, 1] \\
& && v_2 + w_2\zeta \geq x && \forall \zeta \in [-1, 1] \\
& && v_2 + w_2\zeta \geq x - \zeta && \forall \zeta \in [-1, 1] \\
& && v_1, v_2, w_1, w_2, y \in \mathbb{R}, x \in \mathbb{R}_+.
\end{aligned}
$$

The optimal value of this problem is 1 ($v_1 = v_2 = \frac{1}{2}, w_1 = \frac{1}{2}, w_2 = -\frac{1}{2}, y = 1$), which is the same as the optimal value of the original problem. So, in this case the AARC-R closes the gap. This is not always the case as the following example shows:

$$
\begin{aligned}
\text{(TOY 2)} \quad &\min && d \\
&\text{s.t.} && d \geq \max\{x, x + \zeta_1 + \zeta_2\} + \max\{x, x + \zeta_1 - \zeta_2\} \\
& && \qquad + \max\{x, x - \zeta_1 + \zeta_2\} + \max\{x, x - \zeta_1 - \zeta_2\} \\
& && \qquad\qquad \forall \zeta \in [-1, 1]^2 \\
& && d \in \mathbb{R}, x \in \mathbb{R}_+.
\end{aligned}
$$

We obtain $v_{true} = 2$, $v_{RC-R} = 8$ and $v_{AARC-R} = 4$. The AARC-R is an improvement over the RC-R, but is not exact. The QARC-R, which can be derived by reparameterizing the uncertainty region to $[0, 1]^2$, applying the work of Yanıkoğlu et al. (2012) to obtain an LMI description of the uncertainty region, and formulating the RC to a deterministic program by dualization, gives $v_{QARC-R} = 2.83$. While this value is much closer to 2 than the value of the AARC-R, it is still inexact.

4.5.3 Least absolute deviations regression with errors-in-variables

An errors-in-variables regression model is a model $y_i = \beta_0 + \beta_1 x_i^* + \varepsilon_i$ ($\varepsilon_i \sim N(0, \sigma^2)$ i.i.d.) where x_i^* cannot be measured accurately. Only x_i and y_i with $x_i^* = (1+\zeta_i)x_i$ are observed, where ζ_i is an unknown measurement error. The least absolute deviations approach estimates β_0 and β_1 by minimizing $\sum_{i=1}^n |y_i - \beta_0 - \beta_1 x_i^*|$:

$$
\min_{\beta_0, \beta_1} \sum_{i=1}^n \max\{y_i - \beta_0 - \beta_1(1 + \zeta_i)x_i, \beta_0 + \beta_1(1 + \zeta_i)x_i - y_i\}.
$$

Because we do not know the values ζ_i, we can apply RO:

$$
\min_{\beta_0, \beta_1} \max_{\zeta \in \mathcal{Z}} \sum_{i=1}^n \max\{y_i - \beta_0 - \beta_1(1 + \zeta_i)x_i, \beta_0 + \beta_1(1 + \zeta_i)x_i - y_i\},
$$

where we choose the uncertainty region $\mathcal{Z}$ to be ellipsoidal: $\mathcal{Z} = \{\zeta \in \mathbb{R}^n : ||\zeta||_2 \leq \Omega\}$. Note that this is the second special cases in Section 4.3.1.3 because it is the sum of absolute values, ζ_i appears only in term i, and the uncertainty region is centrosymmetric around $\zeta = 0$. It follows that the RC can be written as:

$$\min_{\beta_0, \beta_1} \sum_{i=1}^{n} \max\{y_i - \beta_0 - \beta_1 x_i, \beta_0 + \beta_1 x_i - y_i\} + \Omega \, ||\beta_1 x||_2 \, , \tag{4.15}$$

which can be reformulated as a CQP. Even though there is this explicit and simple formulation of the exact RC, we also try the other approaches from Section 4.3 for a comparison. The RC-R is as follows:

$$
\begin{aligned}
\text{(RC-R)} \qquad \min \quad & \sum_{i=1}^{n} z_i \\
\text{s.t.} \quad & z_i \geq y_i - \beta_0 - \beta_1 x_i + \Omega|\beta_1 x_i| && \forall i \in I \\
& z_i \geq \beta_0 + \beta_1 x_i - y_i + \Omega|\beta_1 x_i| && \forall i \in I,
\end{aligned}
$$

which for comparison with formulation (4.15) can also be written as:

$$\min_{\beta_0, \beta_1} \sum_{i=1}^{n} \max\{y_i - \beta_0 - \beta_1 x_i, \beta_0 + \beta_1 x_i - y_i\} + \Omega \, ||\beta_1 x||_1 \, .$$

This is more pessimistic than constraint (4.15), since $||\cdot||_1 \geq ||\cdot||_2$. For the RC-R, the worst case realization in the ellipsoidal set has all but one elements equal to 0. Therefore, it may seem that the RC-R is the RC of a linear constraint with interval uncertainty, but this is not the case. The AARC-R is given by:

$$
\begin{aligned}
\min \quad & d \\
\text{s.t.} \quad & d \geq \sum_{i=1}^{n} v_i + w_i^\top \zeta && \forall \zeta \in \mathbb{R}^n : ||\zeta||_2 \leq \Omega \\
& v_i + w_i^\top \zeta \geq y_i - \beta_0 - \beta_1(1 + \zeta_i)x_i && \forall \zeta \in \mathbb{R}^n : ||\zeta||_2 \leq \Omega \quad \forall i \in I \\
& v_i + w_i^\top \zeta \geq \beta_0 + \beta_1(1 + \zeta_i)x_i - y_i && \forall \zeta \in \mathbb{R}^n : ||\zeta||_2 \leq \Omega \quad \forall i \in I,
\end{aligned}
$$

which after reformulation becomes:

$$
\begin{aligned}
\text{(AARC-R)} \qquad \min \quad & d \\
\text{s.t.} \quad & d \geq \Omega \, ||w||_2 + \sum_{i=1}^{n} v_i \\
& v_i \geq y_i - \beta_0 - \beta_1 x_i + \Omega \, ||\beta_1 x_i e_i + w_i||_2 && \forall i \in I \\
& v_i \geq \beta_0 + \beta_1 x_i - y_i + \Omega \, ||\beta_1 x_i e_i - w_i||_2 && \forall i \in I,
\end{aligned}
$$

where e_i is the i^{th} unit vector.

We have run these four models on $1,000$ cases. In each case, we fixed the parameters at $\beta_0 = 2$, $\beta_1 = 5$ and $\sigma^2 = 1$. A case consists of 15 observations of x_i, generated uniformly in $[0, 100]$ i.i.d., $\zeta \in \mathbb{R}^n : ||\zeta||_2 \leq \Omega$ uniformly and $\varepsilon_i \sim N(0, \sigma^2)$ i.i.d. Using these random draws, we have computed $y_i = \beta_0 + \beta_1(1 + \zeta_i)x_i + \varepsilon_i$. For the uncertainty region, we have picked $\Omega = 0.05$. For each of this cases, we solved the exact formulation (4.15), AARC-R, QARC-R, RC-R, and the nominal problem using CVX and SDPT3. All solution times are very low, and hence, not mentioned.

Next to comparing the usual way by means of v_{true}, we also compare solutions by looking at how well β_0 and β_1 are estimated. In all $1,000$ cases, the AARC-R and QARC-R give the same estimates of β_0 and β_1. Hence, they are considered equal. Histograms, the mean, and s^2 statistics for β_0, β_1 and v_{true} are shown in Figure 4.1. All models do well in estimating the parameters β_0 and β_1, except the AARC-R and the QARC-R for β_0.

The objective values (averaged over the $1,000$ cases) are listed in Table 4.1. The AARC-R and QARC-R objective values are 22.7% respectively 94.4% closer to the optimum than the RC-R objective value. However, for the true value we have $v_{true}(x_{RC-R}) = 92.095$ while $v_{true}(x_{AARC-R}) = v_{true}(x_{QARC-R}) = 99.322$. The RC-R outperforms the AARC-R and the QARC-R in 98% of the cases, and even the nominal solution outperforms the RC-R in 99% of the cases. So, even though the AARC-R and QARC-R provide a much better bound on the optimal true value, their true value is not necessarily better. This shows two things: it may be very misleading to compare robust solutions by their objective values, and using a better approximation of the true objective function might not improve the solution at all.

While we were able to solve the exact RC (4.15) directly, we have also looked at the efficiency of the cutting plane methods. This comparison is not based on the $1,000$ cases mentioned earlier, but on new cases where we varied the dimension of ζ (and consequently, the dimension of x). Each case was created in the same way as described before. Algorithms 2 and 3 can run very quickly because the worst case ζ can be computed efficiently, as this is a simple case. It is therefore interesting to look at the number of iterations the algorithms take. We have generated $1,000$ data sets in which we varied the number of observations. Algorithm 2 adds only 3 or 4 constraints, independent of the number of observations. This shows that the method is not only effective for polyhedral uncertainty regions. The number of iterations of Algorithm 3 is shown in Figure 4.2, where it seems to be a square root function of the number of observations. The regression model $numiter_i = \alpha\sqrt{numobs_i} + \varepsilon_i$ gives $\alpha = 1.258$ with a standard error of 0.007 and $R^2 = 0.966$. For 200 observations, the algorithm generates at most 18 constraints while full EORLC would result in a

Table 4.1 – A comparison of different solution methods for the regression example averaged over $1,000$ datasets.

Method	v_{method}	$v_{true}(x_{method})$
Nominal	36.326	91.994
RC-R	224.096	92.095
AARC-R	194.091	99.322
QARC-R	99.323	99.322
Exact	91.973	91.973

model with 2^{200} constraints. The algorithm needs more iterations than the vertex enumeration algorithm, but is still effective.

4.5.4   Brachytherapy

High dose rate brachytherapy (HDR-BT) is a form of radiation therapy where a highly radioactive sealed source is inserted into a tumor for short time periods via approximately fifteen till twenty catheters. When the catheter positions are fixed, a treatment plan specifies for how long the radioactive source has to stay at which position inside the catheters. A perfect treatment plan delivers a prescribed dose to the tumor while not delivering any dose to the surrounding organs at risk. Because this is physically impossible, the goal is to find a good trade-off between irradiating the tumor and saving the organs at risk. The quality of a plan can be measured by means of calculation points. These are artificial points where the received dose can be computed and compared to a prescribed lower and upper bound, that are distributed inside and around the tumor and organs at risk. The dose in a specific calculation point i can be computed as the sum of the individual contributions from every catheter k, which in turn is the sum of the dwell times of the individual dwell positions inside catheter k multiplied by given dose rate vectors d_{ik}:

$$\sum_{k \in K} d_{ik}^{\top} t_k.$$

If calculation point i fails to satisfy the lower bound L_i or the upper bound U_i, it contributes a linear penalty of α_i or β_i (respectively) per unit of violation to the objective function. This results in the following optimization problem:

$$\min_{t_k \in \mathbb{R}^n_+} \sum_{i \in I} \max\{0, \alpha_i(L_i - \sum_{k \in K} d_{ik}^{\top} t_k), \beta_i(\sum_{k \in K} d_{ik}^{\top} t_k - U_i)\}.$$
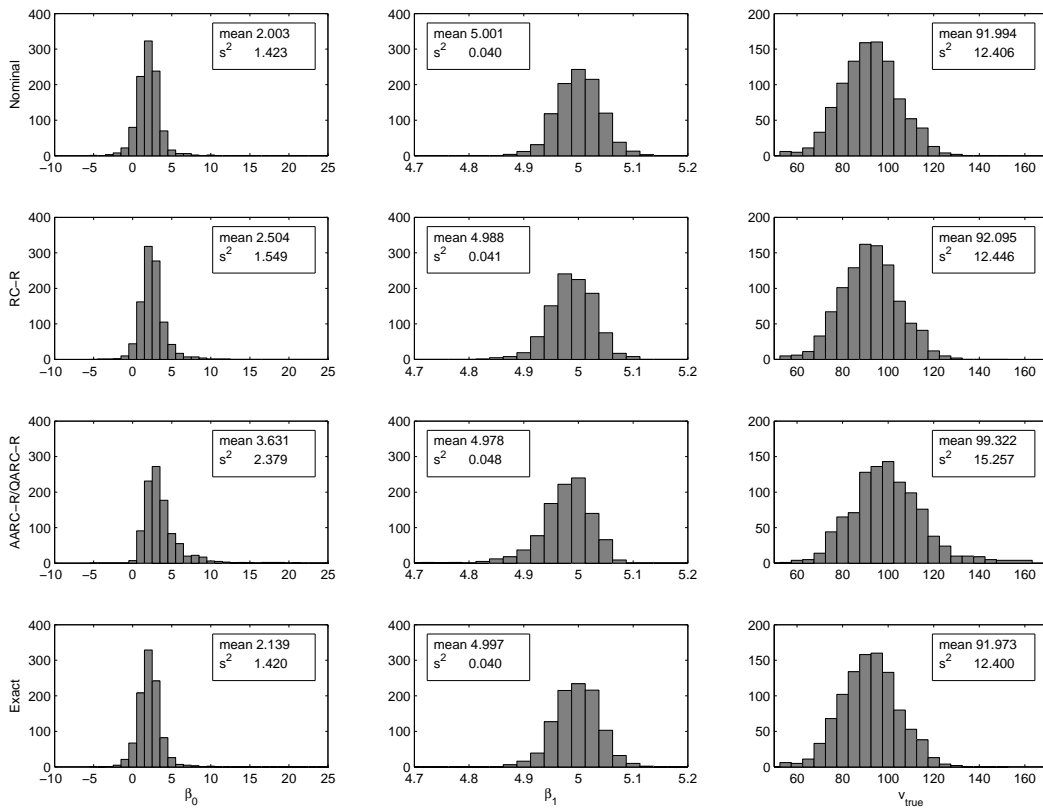
Figure 4.1 – Comparison of the nominal solution, the RC-R solution, the AARC-R/QARC-R solution and the exact solution for the regression model of Section 4.5.3.

Figure 4.2 – The number of iterations needed by Algorithm 3 to find an optimal solution versus the number of observations in the regression model of Section 4.5.3.

This convex piecewise linear objective function is commonly used for treatment planning (Alterovitz et al. (2006); Karabis et al. (2009); Lessard and Pouliot (2001)).

The parameters d_{ik} are computed based on the catheter positions. These positions cannot be measured accurately, hence the data in the optimization problem is uncertain. We assume that the true catheter position is within some cone around the measured position, which is justified because one side of the catheter is fixed at a known position. We replace the cone with a polyhedral cone with a 10-sided base to end up with a polyhedral set, and we make the simplification that the vector d_{ik} inside the cone is a convex combination of the vectors d_{ik} at the sides of the cone. Thus, we only need to know the vectors d_{ik} at the catheter positions corresponding to the 10 edges of the cone connecting the apex to the base. Using these vectors as the columns of a matrix B, we can write $d_{ik} = B_{ik}\zeta_k$, where $\zeta_k \in \Delta^{|S|-1}$ ($\Delta^{|S|-1} = \{\zeta \in \mathbb{R}^{|S|} : e^\top \zeta = 1, \zeta \geq 0\}$, the standard simplex in $\mathbb{R}^{|S|}$), and $|S|$ is the number of sides of the base of the polyhedral cone. This gives the following RO problem:

$$
\begin{aligned}
\min \quad & v \\
\text{s.t.} \quad & \sum_i \max\{0, \alpha_i(L_i - \sum_{k \in K}(B_{ik}\zeta_k)^\top t_k), \\
& \qquad \beta_i(\sum_{k \in K}(B_{ik}\zeta_k)^\top t_k - U_i)\} \leq v \quad \forall \zeta_k \in \Delta^{|S|-1}(k \in K) \\
& t_k \geq 0 \qquad\qquad\qquad\qquad\qquad\qquad \forall k \in K.
\end{aligned}
$$

The number of calculation points is usually in the order of magnitude of 5,000. We have data of one patient that has been exported from treatment planning software with a large number of calculation points. For the purpose of this chapter, we have reduced the number of calculation points to 40 by taking a random subset, because otherwise both the maximization step in Algorithm 3 and the AARC-R are intractable. The model does not lose its value from this reduction, because the objective could be a weighted average between the nominal objective based on all calculation point, and the robust objective based on a small fraction of calculation points that is well distributed inside the tumor.

The RC-R is based on the original constraint with the sum of 40 maxima of 3 functions. We have also solved the RC-R and AARC-R after splitting up the sum, as outlined in the last paragraph of Section 4.3.2. If we split the sum in groups of size 4, then the number of sums reduces to 10 while each term of sum is the maximum of 3^4 functions.

The results for the different methods are listed in Table 4.2. The first observation is that the nominal solution has a true value that is five times the optimal value, so the nominal solution is nonrobust. The AARC-R has value ≈ 143, which is much

Table 4.2 – A comparison of different solution methods for the brachytherapy example.

Method	Iterations	Sol. time (s)	v_{method}	$v_{true}(x_{method})$
Nominal	–	0	11.033	711.362
RC-R	–	1	229.986	157.277
AARC-R	–	314	143.104	140.827
RC-R (sum of 8 maxima of 243 functions)	–	51	186.219	157.568
RC-R (sum of 10 maxima of 81 functions)	–	17	203.100	157.030
RC-R (sum of 20 maxima 9 functions)	–	2	218.894	158.438
AARC-R (sum of 10 maxima of 81 functions)	–	136214	141.481	140.114
AARC-R (sum of 20 maxima 9 functions)	–	2774	142.447	139.957
Algorithm 2	25	1,173	137.628	137.676
Algorithm 3	36	1,065	137.629	137.662
Combined	20	894	137.628	137.650
Algorithm 2*	53	234	137.627	137.668
Algorithm 3*	> 400	> 50,000	–	–
Combined*	41	353	137.628	137.676

$\varepsilon = 0.05$

* The optimization problem for determining v_{true} is stopped as soon as the upper bound is at least 0.05 higher than the lower bound

closer to the true value (≈ 137) than to the RC-R value (≈ 229), so the AARC-R almost closes the gap. Algorithm 3 is outperformed by Algorithm 2. Its long running time when the optimization problem for determining v_{true} is not solved to optimality is not only due to the larger number of iterations, but more importantly, to the gradually increasing running time of solving LP in each iteration. This step becomes so memory consuming, that the algorithm cannot finish with CPLEX on 32 bit hardware. Combining both algorithms reduces the number of iterations, as more work per iteration is done, but also increases the solution time in case the optimization problem for determining v_{true} is not solved to optimality. Splitting up the sum does not give any benefit in terms of true value, though it gives a lower objective value when solving the RC-R.

4.5.5 Inventory planning

We consider a single item inventory model where backlogging is allowed to compare the AARC-R with the exact RC. At the beginning of each period, an order can be placed that is delivered instantly. At the end of each period, the holding and backlogging costs are c_h and c_b per unit, respectively. The objective is to minimize the costs:

$$
\min_{q,w\in\mathbb{R}_+^{|T|}} \max_{d\in\mathcal{Z}} \left(\sum_{t=1}^{|T|} \max\{c_h[x_0 + \sum_{i=1}^{t} q_i - d_i], c_b[x_0 + \sum_{i=1}^{t} q_i - d_i]\} \right),
$$

where x_0 is the starting inventory, q_i is the order quantity at time i, and d_i is the uncertain demand at time i. In all formulations we allow q_i to depend affinely on the demand in periods 1 up to $i-1$.

If the uncertainty region is a box, the AARC-R turns out to be exact for the intervals we have tried. This is in accordance with the numbers reported by Ben-Tal et al. (2005). We found that the AARC-R is no longer exact if the uncertainty region is an ellipsoid. Because demand is nonnegative, the ellipsoid is intersected with the nonnegative orthant:

$$
\mathcal{Z} = \{d \in \mathbb{R}_+^{|T|} : \left\|d - \bar{d}\right\|_2 \leq \Omega\}.
$$

In our numerical study we have looked at 12 time periods, with parameters $\Omega = 10$, $\bar{d} = 5$, $c_h = 1$, and $c_b = 2$.

Each determination of v_{true} in Algorithm 3 takes up to 2 minutes if continued to optimality using CPLEX, making it the most time consuming step in the algorithm. Because the dimension of the uncertain demand vector is 12, a global solver might be faster. We have tried LGO 1.0, whose accuracy can be adjusted with the parameters "maximal number of function evaluations" and "maximal number of stalled

evaluations", which initially both are 16,000. Every time the upper bound found by LGO is less than 0.1 larger than the lower bound, the parameters are increased by 25% until they exceed 1,000,000. This reflects the idea that it is still easy to find a violated constraint when the algorithm starts, but gradually becomes more difficult as the quality of the solution increases. Still, CPLEX finds better solutions in less time.

The RC-R is based on the original constraint with the sum of 12 maxima of 2 functions. We have also solved the RC-R and AARC-R after splitting up the sum, as outlined in the last paragraph of Section 4.3.2. In this example the uncertainty enters the model through time, so the most natural way of splitting up the summation is in consecutive time periods. We split the problem into two groups:

$$
\begin{aligned}
\min \quad & y_1 + y_2 \\
\text{s.t.} \quad & \sum_{t=1}^{6} \max\{c_h[x_0 + \sum_{i=1}^{t} q_i - d_i], c_b[x_0 + \sum_{i=1}^{t} q_i - d_i] \leq y_1 \qquad \forall d \in \mathcal{Z} \\
& \sum_{t=7}^{12} \max\{c_h[x_0 + \sum_{i=1}^{t} q_i - d_i], c_b[x_0 + \sum_{i=1}^{t} q_i - d_i] \leq y_2 \qquad \forall d \in \mathcal{Z},
\end{aligned}
$$

where again we allow q_i to depend affinely on the demand in periods 1 up to $i - 1$. This reformulation introduces more constraints with the sum of maxima, but each sum contains less terms, hopefully resulting in a shorter solution time. The worst case d may differ for the constraints that set y_1 and y_2, hence this reformulation is not exact.

The results are listed in Table 4.3. Again note that the order policies (q_i) are adjustable in all formulations, including the RC-R, so the differences in this table are caused only by different reformulations of the sum of maxima. The first observation is that the AARC-R gives a very small gain over the RC-R, but still has almost twice the optimal value. So, making the analysis variables adjustable does not significantly improve the solution. Algorithm 3 is the fastest cutting plane method, requiring approximately the same number of iterations when combined with Algorithm 2. The sum splitting method significantly reduces the computation time at the cost of nonoptimality. It performs much better than the AARC-R, both in optimal value and in true value. Using the AARC-R on the splitted sums gives a very small gain over the RC-R on the splitted sums, just as for the full problem, so it is not listed in the table. The large true value of the nominal solution comes from the fact that the order sizes are fixed in advance and are not adjusted to observed demand.

Table 4.3 – A comparison of different solution methods for the inventory example.

Method	Iterations	Sol. time (min)	v_{method}	$v_{true}(x_{method})$
Nominal	—	0	0.000	509.903
RC-R	—	0	120.000	94.641
AARC-R	—	0	120.000	93.315
EORLC	—	10	48.750	48.750
RC-R (sum of 2 maxima of 64 functions)	—	0	68.613	54.162
RC-R (sum of 3 maxima of 16 functions)	—	0	83.631	58.140
RC-R (sum of 4 maxima of 8 functions)	—	0	94.456	66.637
RC-R (sum of 6 maxima of 4 functions)	—	0	107.627	77.957
Algorithm 2	123	76	48.749	48.796
Algorithm 3	77	36	48.752	48.802
Combined	79	45	48.755	48.798
Algorithm 2*	309	180	48.752	48.827
Algorithm 3*	112	9	48.753	48.790
Combined *	103	29	48.755	48.825

$\varepsilon = 0.1$

* The optimization problem for determining v_{true} is stopped as soon as the upper bound is at least 0.1 higher than the lower bound

4.6 Conclusions

Because RO is applied constraint-wise, it is very important how constraints are formulated. In this chapter we list several approaches to an inequality constraint containing the sum of maxima of linear functions. The RC-R, often used in the literature, is the most pessimistic approach. It is obtained by first reformulating the deterministic constraint into linear constraints using analysis variables, and then applying RO. Its pessimism can be reduced by replacing the analysis variables with linear decision rules before applying RO, which gives the AARC-R. The AARC-R seems to work well for the practical problems we analyzed with polyhedral uncertainty regions, but we have constructed an example with polyhedral uncertainty where the value of the AARC-R is 100% higher than the true value. Nonlinear decision rules may give better results, but are computationally more challenging. The conservatism of the approximations can be reduced by combining several max expressions before reformulating. Especially for ellipsoidal uncertainty this method gives much better solutions at the cost of a slightly higher solution time.

In many cases it is not necessary to use an approximation because an exact reformulation can be practically solved. We identify four special cases in which an exact reformulation is often tractable. For the general case we give two exact general methods: vertex enumeration and EORLC. Both methods may result in very large optimization models, but cutting plane methods can be used to handle this. Vertex enumeration adds a set of constraints for every vertex of the uncertainty region, so this method is preferred if the uncertainty region has a low number of vertices. Surprisingly, its cutting plane version is also capable of solving problems where the uncertainty region has an infinite number of extreme points efficiently. EORLC is preferred if the number of terms with a maximum function is low. If it is not clear in advance which cutting plane method is faster, both methods have to be tried because our numerical examples do not show a clear preference. Both methods can be combined, but we have not found a situation in which it is beneficial to do so. EORLC can be combined with approximations, which is a new idea that can be used for large scale problems which shows promising numerical results.

The RC-R is often used in the literature while less conservative approaches could have been applied, mostly without explicitly mentioning that their approach is conservative. In the paper by Kropat and Weber (2008), the exact method EORLC would have increased the number of constraints by only a factor four while reducing the number of variables with almost a factor 2 and not changing the structure of the problem. The same authors applied vertex enumeration to an ellipsoidal constraint with polyhedral uncertainty (Özmen et al., 2011), which gives an exact reformulation. Bertsimas and Thiele (2006), Wei et al. (2009) and Alem and Morabito (2012)

apply the RC-R to a problem with polyhedral uncertainty. Our results, and also the numbers reported by Ben-Tal et al. (2005), show that the AARC-R often gives less conservative solutions while it has the advantage that the problem remains linear. For inventory problems in general, when the order quantities are made adjustable, then oftenly also the analysis variables are made adjustable. We show that the latter is not beneficial for ellipsoidal uncertainty, and that both the RC-R and the AARC-R give very bad solutions. For small planning horizons, an exact method has to be used, while for larger horizons splitting the sum in small groups and applying EORLC on the groups significantly improves the solution. Ng et al. (2010) solve a lot allocation problem with ellipsoidal uncertainty. Because the problem is computationally challenging, they solve a problem equivalent to our RC-R using Benders decomposition. Even though their problem is so challenging that even the simple RC-R cannot be solved within a day and their speed-ups are necessary to solve the problem efficiently, it is still interesting to know the conservatism of their approach. We have been able to get a suboptimal solution with cutting planes based on vertex enumeration, where both the minimization and the maximization step were stopped before optimality. The solution we got after four hours has a true value of 26.8, whereas the RC-R (which we tried to solve as an MILP) has a value between 28.7 and 47.2. So the objective value of the solution proposed by Ng et al. (2010) is at least 7-76% too pessimistic.

From our numerical examples it becomes clear that the RC-R is not necessarily better than the nominal problem. Neither of the two optimizes the true problem, so it cannot be determined a priori which one has a better true value. The same holds for the AARC-R and the RC-R: If the AARC-R gives a much lower value then at least it provides a guarantee on the worst case, but the RC-R may still outperform the AARC-R. When using an approximation, it is therefore crucial to measure its quality. This can be accomplished by comparing the true value of the solution of the approximation with the value of an exact formulation. If the problem is too large to be solved exactly, the comparison may be based on a smaller instance with similar structure.

Acknowledgments

4.A Derivation of AARC-R using Fenchel's duality

In this appendix we apply Fenchel's duality to robust constraints, a technique introduced in RO by Ben-Tal et al. (2014). First we will briefly mention the general theory, then we will apply it to constraints of general form, and finally we will apply the general results to constraint (4.1) and show that the result is the same as the AARC-R.

4.A.1 Fenchel's duality theorem

We start with some definitions that are necessary to formulate Fenchel's theorem:

Definition 1 *A function ϕ is proper convex if it is convex, its codomain is $\mathbb{R} \cup \{\infty\}$, and $\phi(x) < \infty$ for at least one x.*

Definition 2 *A function ψ is proper concave if it is concave, its codomain is $\mathbb{R} \cup \{-\infty\}$, and $\psi(x) > -\infty$ for at least one x.*

Theorem 2 *(Fenchel's duality (Rockafellar, 1970, p. 327)) Let ϕ be a proper convex function on $\mathbb{R}^n$, let ψ be a proper concave function on $\mathbb{R}^n$, and let either of the following conditions be satisifed:*

- *$ri(dom\ \phi) \cap ri(dom\ \psi) \neq \emptyset$*

- *ϕ and ψ are closed, and $ri(dom\ \phi^*) \cap ri(dom\ \psi_*) \neq \emptyset$,*

where ri is the relative interior, dom is the effective domain ($dom\ \phi = \{x : \phi(x) < \infty\}$), and ϕ^* and ψ_* are the convex and concave conjugate of ϕ and ψ, respectively. That is,

$$\phi^*(s) = \sup_x \{s^\top x - \phi(x)\}$$

$$\psi_*(s) = \inf_x \{s^\top x - \psi(x)\}.$$

Then the following equality holds

$$\inf_x \{\phi(x) - \psi(x)\} = \sup_s \{\psi_*(s) - \phi^*(s)\}.$$

4.A.1.1   Fenchel's duality applied to a robust constraint of general form.

We focus on the general robust constraint

$$g(\zeta, x) \le d \qquad \forall \zeta \in \mathcal{Z}, \tag{4.16}$$

where g is a proper concave function of ζ for any fixed value of x, and the condition for Fenchel's duality is satisfied with respect to the first argument for any fixed value of x. Because values of g are not of interest when $\zeta \notin \mathcal{Z}$, we assume that $g(\zeta, x) = -\infty$ for all $\zeta \notin \mathcal{Z}$. We also assume that $\mathcal{Z}$ is a compact set so that this constraint is equivalent to:

$$\max_{\zeta \in \mathbb{R}^L} \{g(\zeta, x) - \delta_{\mathcal{Z}}(\zeta)\} \le d, \tag{4.17}$$

with $\delta_{\mathcal{Z}}$ the indicator function ($\delta_{\mathcal{Z}}(\zeta) = 0$ if $\zeta \in \mathcal{Z}$, and ∞ otherwise). We can rewrite the left-hand side by applying Fenchel's duality:

$$\min_{s \in \mathbb{R}^L} \{\delta_{\mathcal{Z}}^*(s) - g_*(s, x)\} \le d,$$

which holds if and only if there exists some $s \in \mathbb{R}^L$ such that:

$$\delta_{\mathcal{Z}}^*(s) - g_*(s, x) \le d. \tag{4.18}$$

Note that this constraint is convex in s. Because every step is 'if and only if', constraint (4.16) is equivalent to constraint (4.18).

4.A.1.2 g is the sum of other functions.

If g can be written as the sum of several other functions, it might be impossible or very difficult to find a closed form solution for its conjugate function. Suppose we have a constraint of the form:

$$\sum_{i \in I} g_i(\zeta, x) \le d \qquad \forall \zeta \in \mathcal{Z}, \tag{4.19}$$

which is constraint (4.16) with $g = \sum_{i \in I} g_i$. If we want to formulate an equivalent constraint using Fenchel's duality, we need the concave conjugate of g. Under some mild assumptions on g_i, it turns out to be sufficient to have closed form solutions for the conjugates of g_i. The following lemma appears to be very useful (Rockafellar, 1970, p. 145):

Lemma 2 *Let ψ_i ($i \in I$) be proper concave functions on $\mathbb{R}^n$. If $\cap_{i \in I} ri(dom \ \psi_i) \ne \emptyset$ then*

$$\left(\sum_{i \in I} \psi_i\right)_*(s) = \sup_{\sum_{i \in I} s_i = s} \left\{\sum_{i \in I} (\psi_i)_*(s_i)\right\},$$

and the supremum is attained.

Applying this lemma, we can rewrite constraint (4.19) as:

$$\delta_{\mathcal{Z}}^*(s) - \max_{\sum_{i \in I} s_i = s} \{\sum_{i \in I} (g_i)_*(s_i, x)\} \leq d,$$

which is valid if and only if there exists $s_i \in \mathbb{R}^L (i \in I)$ such that

$$\delta_{\mathcal{Z}}^*(\sum_{i \in I} s_i) - \sum_{i \in I} (g_i)_*(s_i, x) \leq d. \tag{4.20}$$

So, if all conjugates have closed form solutions, the reformulated constraint also has a closed form. If all conjugates do not have a closed form, this reformulation is still useful because it allows computing each term separately. We will show this for piecewise linear convex functions later in this section. It should again be noted that this constraint is convex in s_i.

4.A.1.3 f is not convex or g is not concave.

Let us investigate the consequences to the reformulation if g is not concave. We assume g is finite on some non-empty set $\mathcal{Z}$, and ∞ elsewhere, so that its conjugate is non-trivial. Fenchel's inequality (Rockafellar, 1970, p. 105) states that:

$$\delta_{\mathcal{Z}}(\zeta) + \delta_{\mathcal{Z}}^*(s) = \delta_{\mathcal{Z}}(\zeta) + \sup_{\zeta' \in \mathcal{Z}} \{\zeta'^\top s - \delta_{\mathcal{Z}}(\zeta')\} \geq \delta_{\mathcal{Z}}(\zeta) + \zeta^\top s - \delta_{\mathcal{Z}}(\zeta) = \zeta^\top s$$

$$g(\zeta, x) + g_*(s, x) = g(\zeta, x) + \inf_{\zeta' \in \mathcal{Z}} \{\zeta'^\top s - g(\zeta', x)\} \leq g(\zeta, x) + \zeta^\top s - g(\zeta, x) = \zeta^\top s,$$

hence:

$$\delta_{\mathcal{Z}}(\zeta) + \delta_{\mathcal{Z}}^*(s) \geq g(\zeta, x) + g_*(s, x),$$

and consequently:

$$g(\zeta, x) - \delta_{\mathcal{Z}}(\zeta) \leq \delta_{\mathcal{Z}}^*(s) - g_*(s, x).$$

This implies that if constraint (4.18) is satisfied, then so is constraint (4.17), but the reverse implication is not necessarily true. Hence, constraint (4.18) is a conservative reformulation of constraint (4.17).

Also, Lemma 2 no longer holds with equality. We can rewrite it as an inequality:

$$(\sum_{i \in I} \psi_i)_*(s) = \inf_t \{s^\top t - \sum_{i \in I} \psi_i(t)\}$$

$$= \inf_t \{\sum_{i \in I} s_i^\top t - \psi_i(t)\},$$

for any $s_1, ..., s_m \in \mathbb{R}^L$ for which $\sum_{i \in I} s_i = s$. So in particular:

$$
\begin{aligned}
(\sum_{i \in I} \psi_i)_*(s) &= \sup_{\sum_{i \in I} s_i = s} \{\inf_t \{\sum_{i \in I} s_i^\top t - \psi_i(t)\}\} \\
&\geq \sup_{\sum_{i \in I} s_i = s} \{\sum_{i \in I} \inf_t \{s_i^\top t - \psi_i(t)\}\} \\
&= \sup_{\sum_{i \in I} s_i = s} \{\sum_{i \in I} (\psi_i)_*(s_i)\}.
\end{aligned}
$$

This implies that if constraint (4.20) is satisfied, then so is constraint (4.19), but the reverse implication is not necessarily true. Hence, constraint (4.20) can be seen as a conservative reformulation of constraint (4.19).

4.A.1.4 g is the sum of pointwise maxima of linear functions.

Let us first derive the conjugates of some functions before we arrive at the theorem:

$$
\delta_{\mathcal{Z}}^*(s) = \sup_{\zeta \in \mathbb{R}^L} \{s^\top \zeta - \delta_{\mathcal{Z}}(\zeta)\} = \sup_{\zeta \in \mathcal{Z}} \{s^\top \zeta\} = \max_{\zeta \in \mathcal{Z}} \{s^\top \zeta\},
$$

and

$$
\begin{aligned}
(\max_{j \in J} \{\ell_{ij}(\zeta, x)\})_*(s_i, x) &= \inf_{\zeta \in \mathcal{Z}} \{s_i^\top \zeta - \max_{j \in J} \{\ell_{ij}(\zeta, x)\}\} \\
&= \inf_{\zeta \in \mathcal{Z}} \{\min_{j \in J} \{s_i^\top \zeta - \ell_{ij}(\zeta, x)\}\} \\
&= \min_{j \in J} \{\inf_{\zeta \in \mathcal{Z}} \{s_i^\top \zeta - \ell_{ij}(\zeta, x)\}\}.
\end{aligned}
$$

Theorem 3 *Applying Fenchel's duality to:*

$$
\max_{\zeta \in \mathcal{Z}} \sum_{i \in I} \max_{j \in J} \{\ell_{ij}(\zeta, x)\} \leq d, \tag{4.21}
$$

gives a formulation that is equivalent to the AARC-R.

Proof. Constraint (4.21) is equivalent to constraint (4.19) where $g_i(\zeta, x)$ equals $\max_{j \in J} \{\ell_{ij}(\zeta, x)\}$ for ζ in $\mathcal{Z}$ and $-\infty$ otherwise. For for any fixed x, g_i is not concave in ζ so we will end up with a conservative instead of an equivalent reformulation. If we fill in the conjugate functions in constraint (4.20), the following conservative reformulation is obtained:

$$
\max_{\zeta \in \mathcal{Z}} \{\sum_{i \in I} s_i^\top \zeta\} - \sum_{i \in I} \min_{j \in J} \{\inf_{\zeta \in \mathcal{Z}} \{s_i^\top \zeta - \ell_{ij}(\zeta, x)\}\} \leq d.
$$

If we model the second terms as $\sum_{i \in I} z_i$, we can write this as:

$$\sum_{i \in I} s_i^\top \zeta - \sum_{i \in I} z_i \leq d \qquad\qquad \forall \zeta \in \mathcal{Z}$$

$$z_i \leq s_i^\top \zeta - \ell_{ij}(\zeta, x) \qquad\qquad \forall \zeta \in \mathcal{Z} \quad \forall i \in I \quad \forall j \in J,$$

and by rearranging the terms in each constraint we obtain:

$$\sum_{i \in I} [(-z_i) + s_i^\top \zeta] \leq d \qquad\qquad \forall \zeta \in \mathcal{Z}$$

$$(-z_i) + s_i^\top \zeta \geq \ell_{ij}(\zeta, x) \qquad\qquad \forall \zeta \in \mathcal{Z} \quad \forall i \in I \quad \forall j \in J,$$

which is the same as the AARC-R. ∎

4.B   Derivation of AARC-R by reformulating the nonrobust constraint

In this appendix we give a different derivation of the AARC-R of constraint (4.1) when both the biaffine functions and the uncertainty region are separable in the following way:

$$\sum_{i \in I} \max_{j \in J} \{ \sum_{k \in K} \ell_{ijk}(\zeta_k, x) \} \leq d \qquad \forall \zeta_k \in \mathcal{Z}_k \quad (k \in K), \qquad (4.22)$$

and $\mathcal{Z}_k$ is the convex hull of different scenarios ζ_k^s ($s \in S$). An example where this constraint is commonly used, is HDR brachytherapy optimization (Alterovitz et al. (2006); Karabis et al. (2009); Lessard and Pouliot (2001)). If the summation over k were outside the *max* expression, then an analysis variable could be used for every k without introducing any conservatism. Vertex enumeration can then be done on every $\mathcal{Z}_k$ separately. For problems not affected by uncertainty, we show that indeed there is an equivalent formulation where the summation over k is outside the *max* expression. Then we show equivalence to the AARC-R if there is uncertainty. First, we prove the following equality for fixed x and ζ_k:

Lemma 3

$$\sum_{i \in I} \max_{j \in J} \{ \sum_{k \in K} \ell_{ijk}(\zeta_k, x) \} = \min_{y \in \mathbb{R}^{|I||J||K|}: \sum_{k \in K} y_{ijk} = 0} \sum_{k \in K} \sum_{i \in I} \max_{j \in J} \{ y_{ijk} + \ell_{ijk}(\zeta_k, x) \}.$$

$$(4.23)$$

Proof. Note that:

$$\sum_{k \in K} \sum_{i \in I} \max_{j \in J}\{y_{ijk} + \ell_{ijk}(\zeta_k, x)\} \geq \sum_{k \in K} \sum_{i \in I} y_{ij(i)k} + \ell_{ij(i)k}(\zeta_k, x)$$
$$= \sum_{i \in I} \sum_{k \in K} \ell_{ij(i)k}(\zeta_k, x) \quad \forall j(i) \in J,$$

for any $j(i)$ in J, so in particular:

$$\sum_{k \in K} \sum_{i \in I} \max_{j \in J}\{y_{ijk} + \ell_{ijk}(\zeta_k, x)\} \geq \sum_{i \in I} \max_{j \in J}\{\sum_{k \in K} \ell_{ijk}(\zeta_k, x)\},$$

for any y, so in particular for the minimum. Hence, the right hand side of (4.23) is at least as large as the left hand side. On the other hand, given a feasible point for the left hand side of (4.23), we can always construct a feasible point for the right hand side with equal value by taking the same x, and $y_{ijk} = \frac{1}{|K|} \sum_{k' \in K} \ell_{ijk'}(\zeta_{k'}, x) - \ell_{ijk}(\zeta_k, x)$:

$$\sum_{k \in K} \sum_{i \in I} \max_{j \in J}\{y_{ijk} + \ell_{ijk}(\zeta_k, x)\} = \sum_{i \in I} \sum_{k \in K} \max_{j \in J}\{\frac{1}{|K|} \sum_{k' \in K} \ell_{ijk'}(\zeta_{k'}, x)\}$$
$$= \sum_{i \in I} \max_{j \in J}\{\sum_{k' \in K} \ell_{ijk'}(\zeta_{k'}, x)\}.$$

∎

Lemma 4 *A conservative reformulation of constraint* (4.22) *is given by:*

$$\sum_{k \in K} \sum_{i \in I} \max_{j \in J}\{y_{ijk} + \ell_{ijk}(\zeta_k, x)\} \leq d \qquad \forall \zeta_k \in \mathcal{Z}_k \quad (k \in K)$$
$$\sum_{k \in K} y_{ijk} = 0 \qquad \forall i \in I \quad \forall j \in J.$$

Proof. We use the proof of Lemma 3. The first part of the proof still holds, so the equality in constraint (4.23) becomes a "$\leq$". The construction of the feasible point in the second part of the proof depends on the value $\ell_{ijk}(\zeta_k, x)$. This value does not exist in the robust constraint (4.22), since there is no single ζ_k. Note that the left hand sides of (4.22) and (4.23) are the same, so by replacing the left hand side of (4.22) with the right hand side of (4.23), the conservative reformulation is obtained. ∎

The uncertainty is separable per k in the first constraint of the conservative reformulation. Hence, analysis variables per k can replace each term without changing the

solution. Vertex enumeration can then be done for every $\mathcal{Z}_k$ separately:

$$\sum_{k \in K} z_k \leq d \tag{4.24}$$

$$z_k \geq \sum_{i \in I} w_{iks} \qquad \forall s \in S$$

$$w_{iks} \geq y_{ijk} + \ell_{ijk}(\zeta_k^s, x) \qquad \forall i \in I \quad \forall s \in S \quad \forall k \in K$$

$$\sum_{k \in K} y_{ijk} = 0 \qquad \forall i \in I \quad \forall j \in J.$$

It remains to show that this formulation is equivalent to the AARC-R of constraint (4.22).

Theorem 4 *The conservative reformulation in Lemma 4 is the AARC-R in case of scenario generated uncertainty.*

Proof. The AARC-R of constraint (4.22) is given by:

$$\sum_{i \in I} \left(v_i + \sum_{k \in K} w_{ik}^\top \zeta_k \right) \leq d \qquad \forall \zeta_k \in \mathcal{Z}_k \quad (k \in K)$$

$$v_i + \sum_{k \in K} w_{ik}^\top \zeta_k \geq \sum_{k \in K} \ell_{ij}(\zeta_k, x) \quad \forall i \in I \quad \forall j \in J \quad \forall \zeta_k \in \mathcal{Z}_k \quad (k \in K).$$

$$\tag{4.25}$$

Constraint (4.25) can be reformulated as:

$$w_{ik}^\top \zeta_k \geq y_{ijk} + \ell_{ij}(\zeta_k, x) \qquad \forall i \in I \quad \forall j \in J \quad \forall \zeta_k \in \mathcal{Z}_k \quad \forall k \in K$$

$$\sum_{k \in K} y_{ijk} = -v_i \qquad \forall i \in I \quad \forall j \in J.$$

We have assumed scenario generated uncertainty, so our uncertainty region is $\mathcal{Z}_k = \Delta^{|S|-1}$, the standard simplex in $\mathbb{R}^{|S|}$. Suppose an optimal solution has $v_i \neq 0$, then an optimal solution with $v_i = 0$ can be obtained by increasing all elements of w_{ik} for a single random k with v_i because the elements of ζ_k sum to 1. Hence we fix $v_i = 0$. The AARC-R can then be formulated as:

$$\sum_{k \in K} z_k \leq d$$

$$z_k \geq \sum_{i \in I} w_{ik}^\top \zeta_k \qquad \forall \zeta_k \in \mathcal{Z}_k \quad (k \in K)$$

$$w_{ik}^\top \zeta_k \geq y_{ijk} + \ell_{ij}(\zeta_k, x) \qquad \forall i \in I \quad \forall j \in J \quad \forall \zeta_k \in \mathcal{Z}_k \quad \forall k \in K$$

$$\sum_{k \in K} y_{ijk} = 0 \qquad \forall i \in I \quad \forall j \in J.$$

Let w_{iks} denote the s^{th} component of w_{ik}. Equivalence to (4.24) now follows from vertex enumeration. ∎

4.C   Derivation of the QARC-R for an ellipsoidal uncertainty region

In this appendix the SDP reformulation of the QARC-R of (4.1) is derived for an ellipsoidal uncertainty region. For simplicity, we assume $\ell_{ij}(\zeta, x)$ and ℓ are bilinear functions in the parameters. Therefore, they can be expressed as $\ell(\zeta, x) = \zeta^\top L x$ and $\ell_{ij}(\zeta, x) = \zeta^\top L_{ij} x$, respectively, for some matrices L and L_{ij}. The QARC-R is given by:

$$\text{(QARC-R)} \qquad \zeta^\top L x + \sum_{i \in I} \left(v_i + w_i^\top \zeta + \zeta^\top W_i \zeta \right) \leq d \quad \forall \zeta \in \mathbb{R}^L : ||\zeta||_2 \leq \Omega$$

$$v_i + w_i^\top \zeta + \zeta^\top W_i \zeta \geq \zeta^\top L_{ij} x \quad \forall \zeta \in \mathbb{R}^L : ||\zeta||_2 \leq \Omega$$

$$\forall i \in I \quad \forall j \in J.$$

By application of the $\mathcal{S}$–lemma (Ben-Tal et al., 2009a, Lemma 6.5.3), this can be reformulated as the following LMIs:

$$\sum_{i \in I} \begin{pmatrix} W_i & \frac{1}{2} w_i \\ \frac{1}{2} w_i & v_i \end{pmatrix} + \begin{pmatrix} 0 & \frac{1}{2} L x \\ \frac{1}{2} L x & -d \end{pmatrix} \preceq \lambda_0 \begin{pmatrix} I & 0 \\ 0 & -\Omega^2 \end{pmatrix}$$

$$-\begin{pmatrix} W_i & \frac{1}{2} \left(w_i - L_{ij} x \right) \\ \frac{1}{2} \left(w_i - L_{ij} x \right) & v_i \end{pmatrix} \preceq \lambda_{ij} \begin{pmatrix} I & 0 \\ 0 & -\Omega^2 \end{pmatrix} \qquad \forall i \in I \quad \forall j \in J$$

$$\lambda_0 \geq 0, \lambda_{ij} \geq 0.$$

The QARC-R has $|I||J| + 1$ LMIs of size $|L + 1|$.

Robust fractional programming

Abstract We extend Robust Optimization to fractional programming, where both the objective and the constraints contain uncertain parameters. Earlier work did not consider uncertainty in both the objective and the constraints, or did not use Robust Optimization. Our contribution is threefold. First, we provide conditions to guarantee that either a globally optimal solution, or a sequence converging to the globally optimal solution, can be found by solving one or more convex optimization problems. Second, we identify two cases for which an exact solution can be obtained by solving a single optimization problem: (1) when uncertainty in the numerator is independent from the uncertainty in the denominator, and (2) when the denominator does not contain an optimization variable. Third, we show that the general problem can be solved with an (iterative) root finding method. The results are demonstrated on a return-on-investment maximization problem, data envelopment analysis, and mean-variance optimization. We find that the robust optimal solution is only slightly more robust than the nominal solution. As a side-result, we use Robust Optimization to show that two existing methods for solving fractional programs are dual to each other.

5.1   Introduction

A fractional program (FP) is an optimization problem, where the objective is a fraction of two functions. It can be used for an economical trade-off such as maximizing return/investment, maximizing return/risk, minimizing cost/time, and minimizing wasted/used material (Schaible and Ibaraki, 1983). A comprehensive overview of FP papers, containing over 550 references which include many applications, is given by Schaible (1982). More up-to-date references can be found in (Stancu-Minasian,

2013), which also refers to six preceding bibliographies by the same author.

Often, the parameters in an optimization problem are affected by uncertainty. Robust Optimization (RO) is about solving optimization problems with uncertain data in a computationally tractable way (see, e.g., Ben-Tal et al. (2009a); Bertsimas et al. (2011)). The key concept is that a solution has to be feasible for all realizations of the uncertain parameters, which are assumed to reside in convex uncertainty regions.

For a generalized FP, the objective is the maximum of finitely many fractions, and the feasible region is a convex set (see, e.g., Barros et al. (1996); Crouzeix and Ferland (1991) for solution methods). A generalized FP with infinitely many fractions in the objective was solved in (Lin and Sheu, 2005) using a cutting plane method, that uses a set of finitely many fractions that is extended in each step. They do not mention that their method can be used to deal with uncertain data, and do not use existing results from RO. Our work can be seen as an alternative approach, where we also deal with uncertainty in the constraints.

The Lagrange dual of a robust FP was studied by Jeyakumar et al. (2013), extending a result by Beck and Ben-Tal (2009). They assume that the uncertainty in the numerator of the objective is independent from the uncertainty in the denominator. The dual is tractable when the numerator, denominator and constraints are linear and the uncertainty regions are ellipsoidal or finite sets of scenarios. In this chapter, we focus on the primal problem. Nevertheless, in Section 5.4.2 we obtain and extend the list of tractable duals.

The aim of this chapter is to combine FP and RO, to provide a comprehensive overview of the solution methods, and to investigate the improvement of RO on numerical examples. First, we provide conditions that guarantee that a globally optimal solution, or a sequence that converges to the globally optimal solution can be found by solving one or more convex problems. The importance of these conditions is illustrated with a numerical example from literature. Second, we identify two cases for which an exact solution can be obtained by solving a single optimization problem. Third, we show that the general problem can be solved with an (iterative) root finding method.

In Section 5.2, we outline two existing solution methods for FPs, and present a new result showing that the two approaches are each others duals. In Section 5.3, we present existing results in RO, that will be used for FPs as well. Our main contribution is given in Section 5.4. The results are demonstrated on a return-on-investment maximization problem, data envelopment analysis, and mean-variance optimization in Section 5.5.

5.2   Solving nonrobust fractional programs

In this section, we present two existing methods to solve FPs, and show that these methods are dual to each other. To the best of our knowledge, this duality result is new. Consider the following general formulation of an FP:

$$(\text{FP}) \quad \min_{\boldsymbol{x} \in \mathbb{R}^n} \quad \frac{f(\boldsymbol{x})}{g(\boldsymbol{x})} \quad \text{s.t.} \quad h_i(\boldsymbol{x}) \leq 0, \ \forall i \in I.$$

We will assume that the constraint index set I is finite, that f is convex and non-negative and that g is concave and positive over the feasible region. When the functions f, g and h_i are affine, (FP) is a *linear* fractional program:

$$(\text{LFP}) \quad \min_{\boldsymbol{x} \in \mathbb{R}^n} \quad \frac{b_0 + \boldsymbol{b}^\top \boldsymbol{x}}{c_0 + \boldsymbol{c}^\top \boldsymbol{x}} \quad \text{s.t.} \quad d_{0i} + \boldsymbol{d_i}^\top \boldsymbol{x} \leq 0, \quad i \in I.$$

(Charnes and Cooper, 1962) show that (LFP) can be reformulated as an (equivalent) LP by substituting $\boldsymbol{y} = \boldsymbol{x}/(c_0 + \boldsymbol{c}^\top \boldsymbol{x})$ and $t = 1/(c_0 + \boldsymbol{c}^\top \boldsymbol{x})$:

$$(\text{CC-LFP}) \quad \min_{t \in \mathbb{R}_+, \boldsymbol{y} \in \mathbb{R}^n} \quad b_0 t + \boldsymbol{b}^\top \boldsymbol{y} \quad \text{s.t.} \quad d_{0i} t + \boldsymbol{d_i}^\top \boldsymbol{y} \leq 0, \ \forall i \in I,$$

$$c_0 t + \boldsymbol{c}^\top \boldsymbol{y} = 1.$$

An optimal solution of (LFP) is obtained from an optimal solution of (CC-LFP) by computing $\boldsymbol{x} = \boldsymbol{y}/t$. (Schaible, 1974) shows that a similar substitution ($\boldsymbol{y} = \boldsymbol{x}/g(\boldsymbol{x})$, $t = 1/g(\boldsymbol{x})$) transforms (FP), for which the numerator in the objective is positive on the entire feasible region, into an equivalent convex programming problem:

$$(\text{Schaible-FP}) \quad \min_{t \in \mathbb{R}_{++}, \boldsymbol{y} \in \mathbb{R}^n} \quad tf\left(\frac{\boldsymbol{y}}{t}\right) \quad \text{s.t.} \quad tg\left(\frac{\boldsymbol{y}}{t}\right) \geq 1, \quad th_i\left(\frac{\boldsymbol{y}}{t}\right) \leq 0, \ \forall i \in I.$$

(Schaible-FP) is indeed a convex problem, since the perspective function $p(\boldsymbol{y}, t) := tf(\boldsymbol{y}/t)$ is jointly convex on $\mathbb{R}^n \times \mathbb{R}_+$ when f is convex on $\mathbb{R}^n$. Furthermore, an optimal $\boldsymbol{x}$ is obtained from $\boldsymbol{x} = \boldsymbol{y}/t$. Schaible also shows that, if the constraint $tg(\boldsymbol{y}/t) \geq 1$ is formulated as an equality constraint: $tg(\boldsymbol{y}/t) = 1$, it is not necessary to require f to be positive (Schaible, 1974).

A different solution approach uses an auxiliary parameterized optimization problem. Let $F(\alpha) = \min_{\boldsymbol{x} \in \mathbb{R}^n} \{f(\boldsymbol{x}) - \alpha g(\boldsymbol{x}) : h_i(\boldsymbol{x}) \leq 0 \ \forall i \in I\}$. The objective value of (FP) is at least α if and only if $F(\alpha) \geq 0$ (Dinkelbach, 1967). So, the objective of (FP) equals the largest α such that $F(\alpha) \geq 0$:

$$(\text{P-FP}) \quad \max_{\alpha \in \mathbb{R}_+} \{\alpha : \min_{\boldsymbol{x} \in X} \{f(\boldsymbol{x}) - \alpha g(\boldsymbol{x})\} \geq 0\},$$

where X denotes the feasible region $\{\boldsymbol{x} \in \mathbb{R}^n : h_i(\boldsymbol{x}) \leq 0, \ \forall i \in I\}$. The usual way of solving this problem is by finding the root of F, since the corresponding $\boldsymbol{x}$ is optimal for (FP) (Dinkelbach, 1967). This is usually done with a Newton-like algorithm, where there is some freedom in choosing the next iteration point (Chen et al., 2009). The root of F is unique, since F is monotonically decreasing in α. The parameteric program $F(\alpha)$ is convex when g is affine or when f is non-negative on the feasible region (since then only $\alpha \in \mathbb{R}_+$ needs to be considered). For these cases, the Newton method to find the root of F was described by Dinkelbach (1967), which creates a monotonically decreasing sequence that converges superlinearly and often (local) quadratically to a root of F (Schaible, 1976).

We now show that these approaches are dual to each other when f is non-negative. The proof for affine g and possibly negative f is similar.

Theorem 5 *Assume that f is convex and non-negative on X, g is concave on X, X is closed and convex, and the optimal value of (FP) is attained. Then (Schaible-FP) and (P-FP) are dual to each other, and strong duality holds.*

Proof. First note that the following reformulation is equivalent to (P-FP):

$$(\text{RP-FP}) \quad \max_{\alpha \in \mathbb{R}_+}\{\alpha : f(\boldsymbol{x}) - \alpha g(\boldsymbol{x}) \geq 0, \ \forall \boldsymbol{x} \in X\}.$$

The remainder of the proof is based on the theory "primal worst is dual best" introduced by Beck and Ben-Tal (2009). Beck and Ben-Tal assume that X is compact and convex. Since we have assumed X to be closed and convex and since an optimal solution of (FP) is attained, compactness can be achieved by intersecting X with a box that includes the optimal solution. Additionally, Beck and Ben-Tal assume that the constraint in (RP-FP) is convex in α and concave in $\boldsymbol{x}$, which indeed holds. For fixed $\boldsymbol{x}$, (RP-FP) is an LP with the following dual:

$$(\text{D-LP}) \quad \min_{t \in \mathbb{R}_+}\{tf(\boldsymbol{x}) : tg(\boldsymbol{x}) \geq 1\}.$$

While (RP-FP) is robust since the constraint has to hold for all $\boldsymbol{x}$ in X, the constraint in the optimistic counterpart of (D-LP) has to hold for a single $\boldsymbol{x}$:

$$(\text{OD-LP}) \quad \min_{t \in \mathbb{R}_+, \boldsymbol{x} \in X}\{tf(\boldsymbol{x}) : tg(\boldsymbol{x}) \geq 1\}.$$

(RP-FP) and (OD-LP) are dual to each other (Beck and Ben-Tal, 2009). Strong duality holds since $(\boldsymbol{x}, t)$ is a Slater point for (OD-LP) for sufficiently large t. It is obvious that (OD-LP) and (Schaible-FP) are equivalent, since $t = 0$ is infeasible for (OD-LP). ∎

5.3   Robust Optimization

There are currently two generic methods that deal with an infinite number of constraints. The first method is applicable to both linear and nonlinear constraints, while the second method can only be applied to robust LPs.

The first ("constraint wise") approach uses conic duality (Ben-Tal et al., 2009a) or Fenchel duality (Ben-Tal et al., 2014). The vector $\boldsymbol{x}$ in $\mathbb{R}^n$ satisfies the following infinite number of constraints:

$$h_i(\boldsymbol{a_i}, \boldsymbol{x}) \le 0, \ \forall \boldsymbol{a_i} \in \mathbb{R}^L : \tau_{ij}(\boldsymbol{a_i}) \le 0, \ \forall j \in J,$$

if and only if there exists $u_{ij} \in \mathbb{R}_+$ and $\boldsymbol{v_{ij}} \in \mathbb{R}^L$ ($j \in J$, J being a finite set) such that $\boldsymbol{x}$ satisfies the following convex constraint:

$$\sum_{j \in J} u_{ij} \tau_{ij}^* \left(\frac{\boldsymbol{v_{ij}}}{u_{ij}} \right) - (h_i)_* \left(\sum_{j \in J} \boldsymbol{v_{ij}}, \boldsymbol{x} \right) \le 0, \tag{5.1}$$

where $\tau_{ij}^*(\boldsymbol{s}) = \sup_{\boldsymbol{a_i} \in \mathbb{R}^L} \{ \boldsymbol{s}^\top \boldsymbol{a_i} - \tau_{ij}(\boldsymbol{a_i}) \}$ and $(h_i)_*(\boldsymbol{s}, \boldsymbol{x}) = \inf_{\boldsymbol{a_i} \in \mathbb{R}^L} \{ \boldsymbol{s}^\top \boldsymbol{a_i} - h_i(\boldsymbol{a_i}, \boldsymbol{x}) \}$ are the convex and concave conjugates of τ_{ij} and h_i, respectively, and $\boldsymbol{v_{ij}}$ is a vector of the same dimension as $\boldsymbol{a_i}$. This approach requires the constraint to be concave in $\boldsymbol{a_i}$, the functions τ_{ij} to be convex, and $\mathrm{ri}(\mathrm{dom}\ h_i(\cdot, \boldsymbol{x})) \cap \mathrm{ri}(\mathcal{U}_i) \ne \emptyset$ for all $\boldsymbol{x} \in \mathbb{R}^n$, where $\mathcal{U}_i = \{ \boldsymbol{a_i} \in \mathbb{R}^L : \tau_{ij}(\boldsymbol{a_i}) \le 0, \ \forall j \in J \}$. This approach yields a tractable formulation for many constraints and many uncertainty sets (see Tables 1 and 2 in (Ben-Tal et al., 2014)), even if the conjugates do not have closed-form expressions. To give an impression of the broad applicability of this method, let us cite some examples, for which it provides a tractable reformulation. For uncertainty sets, one could have a norm-bounded (e.g., box or ball), polyhedral or conic representable set, or a generic set defined by (convex) power functions, exponential functions, negative logarithms, or any function for which the convex conjugate exists. Constraints could be linear or quadratic in the uncertain parameter. There are some operations on functions that preserve the availability of the conjugate. One is multiplication with a non-negative scalar: the concave conjugate of $tf(\boldsymbol{a_0}, \boldsymbol{y}/t)$ (with respect to $\boldsymbol{a_0}$) for $t \ge 0$, is the perspective of the concave conjugate of f: $tf_*(\boldsymbol{s}/t, \boldsymbol{y}/t)$. Another one is when h_i is the sum of two functions: $h_i = h_{i1} + h_{i2}$. Suppose closed-form conjugates exist for h_{i1} and h_{i2} separately; then:

$$(h_i)_*(\boldsymbol{s}, \boldsymbol{x}) = \max_{\boldsymbol{s_1} \in \mathbb{R}^L, \boldsymbol{s_2} \in \mathbb{R}^L} \{ (h_{i1})_*(\boldsymbol{s_1}, \boldsymbol{x}) + (h_{i2})_*(\boldsymbol{s_2}, \boldsymbol{x}) : \boldsymbol{s_1} + \boldsymbol{s_2} = \boldsymbol{s} \}. \tag{5.2}$$

When substituting (5.2) into (5.1), the max operator in (5.2) may be omitted, since if the resulting constraint holds for some $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$, then it surely holds for the

maximum. This example shows that a closed-form expression for the conjugate is indeed not always required.

The second method solves any robust LP with a convex uncertainty region via its Lagrange dual (Chapter 2). The transformation from the primal to the dual is a three step procedure. First, the dual of the nonrobust LP is formulated, where the uncertain parameters are assumed to be known. Second, instead of enforcing the constraints for all realizations of the uncertain parameters ("robust counterpart"), the constraints of the dual have to hold for a single realization of the uncertain parameters ("optimistic counterpart"). So, the uncertain parameters are added to the set of optimization variables. The last step is to reformulate the nonconvex optimistic counterpart to an equivalent convex optimization problem. The optimal solution of the resulting problem can be translated to an optimal solution of the original robust LP via the KKT vector.

In the remainder, we provide reformulations based on (5.1), but the reader should be aware that the other approaches may be useful when all functions involved are linear.

5.4   Solving robust fractional programs

In this section, we show how to solve (R-FP). It is our aim to obtain Robust Counterparts (RCs), that can be solved with existing Robust Optimization methods. First, we formulate conditions that give raise to convex optimization problems. Under these conditions, a globally optimal solution can be found by solving a single convex optimization problem (Sections 5.4.2 and 5.4.3). These conditions also guarantee that, in the general case (Section 5.4.4), a root finding method produces a sequence of convex optimization problems, whose solutions converge to a globally optimal solution. The results of this section are summarized in Tables 5.1-5.3.

5.4.1   Robust formulation and assumptions

The uncertain parameters, denoted by $\boldsymbol{a_i}$, are assumed to lie in sets $\mathcal{U}_i \subset \mathbb{R}^L$, which we define using functions $\tau_{ij} : \mathbb{R}^L \to \mathbb{R}$:

$$\mathcal{U}_i := \{\boldsymbol{a_i} \in \mathbb{R}^L : \tau_{ij}(\boldsymbol{a_i}) \leq 0, \ \forall j \in J\}, \tag{5.3}$$

where J is a finite set. In the RC of (FP), the constraints have to be satisfied by all realizations of the uncertain parameters:

$$(\text{R-FP}) \quad \min_{\alpha \in \mathbb{R}, \boldsymbol{x} \in \mathbb{R}^n} \quad \alpha \quad \text{s.t.} \quad \frac{f(\boldsymbol{a_0}, \boldsymbol{x})}{g(\boldsymbol{a_0}, \boldsymbol{x})} - \alpha \leq 0 \ \forall \boldsymbol{a_0} \in \mathcal{U}_0$$

$$h_i(\boldsymbol{a_i}, \boldsymbol{x}) \leq 0, \ \forall \boldsymbol{a_i} \in \mathcal{U}_i, \ \forall i \in I. \tag{5.4}$$

Note that the uncertainty is specified constraint-wise, which is possible even if the parameters in different constraints are correlated (Ben-Tal et al., 2009a, p. 12). We make the following assumptions:

1. τ_{ij} are convex and the sets $\mathcal{U}_i$ are convex and compact,

2. f and h_i are convex in $\boldsymbol{x}$ for every fixed value of $\boldsymbol{a_i}$ in $\mathcal{U}_i$,

3. g is concave in $\boldsymbol{x}$ for every fixed value of $\boldsymbol{a_i}$ in $\mathcal{U}_i$,

4. f and h_i are concave in $\boldsymbol{a_i}$ for every feasible $\boldsymbol{x}$,

5. g is convex in $\boldsymbol{a_i}$ for every feasible $\boldsymbol{x}$,

6. $g(\boldsymbol{a_0}, \boldsymbol{x}) > 0$ for every $\boldsymbol{a_0}$ in $\mathcal{U}_0$ and every feasible $\boldsymbol{x}$, and

7. $f(\boldsymbol{a_0}, \boldsymbol{x}) \geq 0$ for at least one $\boldsymbol{a_0} \in \mathcal{U}_0$ and for every feasible $\boldsymbol{x}$.

The last assumption is not necessary if g is biaffine, i.e., when g is affine in each parameter when the other parameter is fixed, of which we show the consequences in Section 5.4.5. In robust *linear* programming, the assumption that $\mathcal{U}_i$ is compact and convex is made without any loss of generality (Ben-Tal et al., 2009a, p. 12). For robust FP, compactness is not a restriction since the functions h_i are continuous and the constraints do not contain strict inequalities. So, the problem remains unchanged if the uncertainty region is replaced with its closure. However, requiring $\mathcal{U}_i$ to be convex is a restriction (unless f, g and h are affine in $\boldsymbol{a_i}$), that is necessary for using existing results in RO.

Assumptions 4 and 5 are made solely because they are required by generic methods to derive a tractable RC. There are some examples where the RC can be derived even though these conditions are not fulfilled, e.g., if the uncertainty region is the convex hull of a limited number of points and the constraint is convex in the uncertain parameter, for a conic quadratic program with implementation error, or when the $\mathcal{S}$-lemma or a sums of squares result can be applied (Ben-Tal et al., 2009a, 2014; Ben-Tal and den Hertog, 2014; Bertsimas et al., 2012). In these cases, assumptions 4 and 5 are not necessary.

In literature a problem is solved, that does not satisfy these requirements (Lin and Sheu, 2005). While the authors claim that this did not affect their computations and that they find the global optimum, in Appendix 5.A we show that their solution is suboptimal.

5.4.2   Special case: uncertainty in the numerator is independent of the uncertainty in the denominator

Suppose the uncertainty in the numerator of the objective is decoupled from the uncertainty in the denominator:

$$\text{(R-S1)} \quad \min_{\alpha \in \mathbb{R}_+, \boldsymbol{x} \in \mathbb{R}^n} \quad \alpha \quad \text{s.t.} \quad \frac{f(\boldsymbol{a_0}, \boldsymbol{x})}{g(\boldsymbol{a_0}', \boldsymbol{x})} - \alpha \leq 0, \ \forall \boldsymbol{a_0} \in \mathcal{U}_0, \ \forall \boldsymbol{a_0}' \in \mathcal{U}_0'$$

$$h_i(\boldsymbol{a_i}, \boldsymbol{x}) \leq 0, \ \forall \boldsymbol{a_i} \in \mathcal{U}_i, \ \forall i \in I.$$

We claim that (R-S1) is equivalent to the RC of the Schaible reformulation:

$$\text{(R-Schaible-S1)} \quad \min_{\alpha \in \mathbb{R}_+, t \in \mathbb{R}_{++}, \boldsymbol{y} \in \mathbb{R}^n} \quad \alpha \quad \text{s.t.} \quad tf\left(\boldsymbol{a_0}, \frac{\boldsymbol{y}}{t}\right) - \alpha \leq 0, \ \forall \boldsymbol{a_0} \in \mathcal{U}_0$$

$$tg\left(\boldsymbol{a_0}', \frac{\boldsymbol{y}}{t}\right) \geq 1, \ \forall \boldsymbol{a_0}' \in \mathcal{U}_0' \quad (5.5)$$

$$th_i\left(\boldsymbol{a_i}, \frac{\boldsymbol{y}}{t}\right) \leq 0, \ \forall \boldsymbol{a_i} \in \mathcal{U}_i \ \forall i \in I.$$

Clearly, an optimal solution of (R-Schaible-S1) exists, for which $t = \sup_{\boldsymbol{a_{01}} \in \mathcal{U}_0'} 1/g(\boldsymbol{a_0}', \boldsymbol{y}/t)$. Equivalence between (R-S1) and (R-Schaible-S1) readily follows from the substition $\boldsymbol{x} = \boldsymbol{y}/t$ that converts a feasible solution of one problem to a feasible solution of the other problem.

This result extends (Jeyakumar et al., 2013). They provide the dual of (R-S1), and show that strong duality holds. In case f, g and h_i are linear, they show that the dual of (R-S1) is tractable when the uncertainty region is ellipsoidal or consists of a finite set of scenarios. The resulting problems can also be obtained from our work by applying the solution method from Chapter 2 to (R-Schaible-S1). In addition to ellipsoids or scenarios, our method works with any convex uncertainty region, such as a polyhedral set or a conic quadratic representable set. A similar result is claimed by Kaul et al. (1986), but that result is wrong (see Appendix 5.B).

We provide a reformulation of (R-Schaible-S1) if the RO method using conjugates is used (eq. (5.1)). The resulting equivalent problem becomes:

$$\min \quad \alpha \quad \text{s.t.} \quad \sum_{j \in J} u_{0j} \tau_{0j}^* \left(\frac{\boldsymbol{v_{0j}}}{u_{0j}}\right) - tf_*\left(\frac{\sum_{j \in J} \boldsymbol{v_{0j}}}{t}, \frac{\boldsymbol{y}}{t}\right) - \alpha \leq 0$$

$$\sum_{j \in J} u_{0'j} \tau_{0'j}^* \left(\frac{\boldsymbol{v_{0'j}}}{u_{0'j}}\right) - tg_*\left(\frac{\sum_{j \in J} \boldsymbol{v_{0'j}}}{t}, \frac{\boldsymbol{y}}{t}\right) \geq 1$$

$$\sum_{j \in J} u_{ij} \tau_{ij}^* \left(\frac{\boldsymbol{v_{ij}}}{u_{ij}}\right) - t(h_i)_*\left(\frac{\sum_{j \in J} \boldsymbol{v_{ij}}}{t}, \frac{\boldsymbol{y}}{t}\right) \leq 0, \ \forall i \in I$$

$$\alpha \in \mathbb{R}_+, t \in \mathbb{R}_{++}, \boldsymbol{u} \in \mathbb{R}^{(|I|+2) \times |J|}, \boldsymbol{v} \in \mathbb{R}^{(|I|+2) \times |J| \times L}, \boldsymbol{y} \in \mathbb{R}^n.$$

5.4.3 Special case: the denominator does not depend on the optimization variable x

If the optimization variables do not appear in the denominator, (R-FP) is equivalent to (Ben-Tal et al., 2014, Ex. 30):

$$(\text{R-S2}) \quad \min_{\alpha \in \mathbb{R}_+, \boldsymbol{x} \in \mathbb{R}^n} \quad \alpha \quad \text{s.t.} \quad f(\boldsymbol{a_0}, \boldsymbol{x}) - \alpha g(\boldsymbol{a_0}) \leq 0, \ \forall \boldsymbol{a_0} \in \mathcal{U}_0$$

$$h_i(\boldsymbol{a_i}, \boldsymbol{x}) \leq 0, \ \forall \boldsymbol{a_i} \in \mathcal{U}_i, \ \forall i \in I.$$

Note that g indeed does not depend on $\boldsymbol{x}$. (R-S2) can be solved via the following equivalent convex reformulation using (5.1) and standard techniques for the conjugate of the sum of two functions:

$$\min \quad \alpha \quad \text{s.t.} \quad \sum_{j \in J} t_{0j} \tau_{0j}^* \left(\frac{\boldsymbol{v_{0j}}}{t_{0j}} \right) - f_* \left(\boldsymbol{s} + \sum_{j \in J} \boldsymbol{v_{0j}}, \boldsymbol{x} \right) + \alpha g^* \left(\frac{\boldsymbol{s}}{\alpha} \right) \leq 0$$

$$\sum_{j \in J} t_{ij} \tau_{ij}^* \left(\frac{\boldsymbol{v_{ij}}}{t_{ij}} \right) - (h_i)_* \left(\sum_{j \in J} \boldsymbol{v_{ij}}, \boldsymbol{x} \right) \leq 0, \ \forall i \in I$$

$$\alpha \in \mathbb{R}_+, \boldsymbol{t} \in \mathbb{R}_+^{(|I|+1) \times |J|}, \boldsymbol{s} \in \mathbb{R}^L, \boldsymbol{v} \in \mathbb{R}^{(|I|+1) \times |J| \times L}, \boldsymbol{x} \in \mathbb{R}^n.$$

5.4.4 General case

We now show how to solve the general problem (R-FP) using the following parametric problem:

$$F(\alpha) = \min_{\boldsymbol{x}, w} \quad w \quad \text{s.t.} \quad f(\boldsymbol{a_0}, \boldsymbol{x}) - \alpha g(\boldsymbol{a_0}, \boldsymbol{x}) \leq w, \ \forall \boldsymbol{a_0} \in \mathcal{U}_0 \tag{5.6}$$

$$h_i(\boldsymbol{a_i}, \boldsymbol{x}) \leq 0, \ \forall \boldsymbol{a_i} \in \mathcal{U}_i, \ \forall i \in I,$$

which is a convex optimization problem since we only have to solve it for $\alpha \in \mathbb{R}_+$. Let α^* be a root of F. Lin and Sheu show that an optimal solution of (R-FP) is the minimizer $\boldsymbol{x}$ of $F(\alpha^*)$ if the feasible region for $\boldsymbol{x}$ is compact (Lin and Sheu, 2005). We assume from now on that the constraint functions h_i define a compact feasible region. Moreover, they show that $F(\alpha) < 0$ if and only if $\alpha > \alpha^*$. Lin and Sheu do not use results from RO to arrive at the deterministic reformulation (5.7). Instead, they replace the set $\mathcal{U}_0$ with a finite set, approximate $F(\alpha)$ with entropic regularization method, and iteratively generate a sequence $\tilde{\alpha}_k$ that converges to α^*. The approximation becomes more accurate as the root of F is approached. The reason why they approximate $F(\alpha)$ is because they claim that computing its value is difficult. Our approach is to solve $F(\alpha)$ using RO, which inherently produces

tractable problems. The following convex reformulation using (5.1) is equivalent:

$$F(\alpha) = \min \ w \quad \text{s.t.} \quad \sum_{j \in J} t_{0j} \tau_{0j}^* \left(\frac{\boldsymbol{v_{0j}}}{t_{0j}} \right) - f_* \left(\boldsymbol{s} + \sum_{j \in J} \boldsymbol{v_{0j}}, \boldsymbol{x} \right) + \alpha g^* \left(\frac{\boldsymbol{s}}{\alpha}, \boldsymbol{x} \right) \leq w$$

(5.7)

$$\sum_{j \in J} t_{ij} \tau_{ij}^* \left(\frac{\boldsymbol{v_{ij}}}{t_{ij}} \right) - (h_i)_* \left(\sum_{j \in J} \boldsymbol{v_{ij}}, \boldsymbol{x} \right) \leq 0, \ \forall i \in I$$

$$\boldsymbol{t} \in \mathbb{R}_+^{(|I|+1) \times |J|}, \boldsymbol{s} \in \mathbb{R}^L, \boldsymbol{v} \in \mathbb{R}^{(|I|+1) \times |J| \times L}, w \in \mathbb{R}, \boldsymbol{x} \in \mathbb{R}^n.$$

Since F is monotonically decreasing in α, as for FPs and generalized FPs, existing root-finding methods can be used. We mention a few of these that produce a sequence $\{\alpha_k\}$ which converges to α^*:

1. The bisection method. Bounds on the interval that contain α^* are:

$$\alpha_{LB} = \min_{\boldsymbol{x} \in \mathbb{R}^n} \{ f(\boldsymbol{\bar{a}_0}, \boldsymbol{x})/g(\boldsymbol{\bar{a}_0}, \boldsymbol{x}) : h_i(\boldsymbol{a_i}, \boldsymbol{x}) \leq 0 : \forall \boldsymbol{a_i} \in \mathcal{U}_i \ \forall i \in I \} \qquad (5.8)$$

$$\alpha_{UB} = \sup_{\boldsymbol{a_0} \in \mathcal{U}_0} f(\boldsymbol{a_0}, \boldsymbol{x})/g(\boldsymbol{a_0}, \boldsymbol{x}), \qquad (5.9)$$

   where (5.8) is computed for a fixed $\boldsymbol{\bar{a}_0}$ from $\mathcal{U}_0$, and (5.9) is computed for some $\boldsymbol{x}$ that is (robust) feasible, i.e., for an $\boldsymbol{x}$ that satisfies (5.4). These bounds can be computed relatively easily using the Schaible reformulation. Alternatively, the lower bound (5.8) may be computed for fixed $\boldsymbol{\bar{a}_i}$ from $\mathcal{U}_i$. Since $F(\alpha_{LB}) \geq 0$ and $F(\alpha_{UB}) \leq 0$, and since F is clearly nonincreasing, α^* lies in $[\alpha_{LB}, \alpha_{UB}]$. The middle point of this interval is $\alpha_k := 0.5(\alpha_{LB} + \alpha_{UB})$. By evaluating $F(\alpha_k)$, the width of the interval that contains α^* can be halved: if $F(\alpha_k) > 0$, then set $\alpha_{LB} := \alpha_k$, otherwise set $\alpha_{UB} := \alpha_k$. By increasing k by 1 and repeating this procedure, the series $\{\alpha_k\}$ converges to α^*.

2. The Dinkelbach type algorithm by Crouzeix et al. (1985), adjusted for infinitely many ratios. The method starts with $k := 0$ and $\alpha_k := \sup_{\boldsymbol{a_0} \in \mathcal{U}_0} f(\boldsymbol{a_0}, \boldsymbol{x})/g(\boldsymbol{a_0}, \boldsymbol{x})$ for some feasible $\boldsymbol{x}$. Then $F(\alpha_k)$ is computed, with maximizer $\boldsymbol{x_k}$. If $F(\alpha_k) < 0$, then the next α is determined by $\alpha_{k+1} := \max_{\boldsymbol{a_0} \in \mathcal{U}_0} f(\boldsymbol{a_0}, \boldsymbol{x_k})/g(\boldsymbol{a_0}, \boldsymbol{x_k})$. Computing α_{k+1} requires solving an FP. The method proceeds by increasing k by 1 and again computing $F(\alpha_k)$. If the feasible region for $\boldsymbol{x}$ is compact, then the series $\{\alpha_k\}$ converges linearly to α_k.

3. The same as the second method, except that (5.6) is replaced with $f(\boldsymbol{a_0}, \boldsymbol{x}) - \alpha g(\boldsymbol{a_0}, \boldsymbol{x}) \leq w g(\boldsymbol{a_0}, \boldsymbol{x_k})$. Multiplying the right-hand side with $g(\boldsymbol{a_0}, \boldsymbol{x_k})$ may increase the speed of convergence for the same complexity of computation (Crouzeix et al., 1986).

4. The same as the second or third method, except that the $\boldsymbol{a_0}$ that maximizes $F(\alpha_k)$ is used to compute α_{k+1} instead of solving a new optimization problem. The worst case $\boldsymbol{a_0}$ in the final solution of computing $F(\alpha_k)$ can be recovered without much computational effort. Thus, $\alpha_{k+1} = f(\boldsymbol{a_0}, \boldsymbol{x_k})/g(\boldsymbol{a_0}, \boldsymbol{x_k})$ is computed more efficiently than in the second or third method. Additional work is required to ensure convergence to α^* (Crouzeix and Ferland, 1991, Section 5).

Let $\boldsymbol{x_k}$ be the maximizer of $F(\alpha_k)$. If a root finding method finds the root in a finite number of steps, then an exact solution of (R-FP) is found. Otherwise, Crouzeix et al. show that, if the sequence $\{\alpha_k\}$ converges to α^*, then any convergent subsequence of $\{\boldsymbol{x_k}\}$ converges to the optimal solution $\boldsymbol{x}^*$ of (R-FP) (Crouzeix et al., 1985, Theorem 4.1c).

5.4.5 Consequences when the denominator is biaffine

The assumption that the numerator is positive ensures that the objective value of (R-FP) is positive over the feasible region. Consequently, we could assume $\alpha \in \mathbb{R}_+$; this would produce a convex optimization problems. If g is biaffine, then the resulting problems are also convex for $\alpha < 0$. We shall discuss the results to each of the three aforementioned cases separately. For the *first special case* (Section 5.4.2), the restriction that the numerator is positive may be dropped only if the denominator does not contain an uncertain parameter. Then, (R-S1) and (R-Schaible-S1) are equivalent if $\alpha \in \mathbb{R}_+$ is replaced with $\alpha \in \mathbb{R}$ and (5.5) is stated as an equality (cf. (Schaible, 1974)). The reason why the denominator may not contain an uncertain parameter is because $t = 1/g(\boldsymbol{x}, \boldsymbol{a_{01}})$ is not possible for multiple $\boldsymbol{a_{01}}$.

For the *second special case* (Section 5.4.3), the denominator only depends on $\boldsymbol{a_0}$, so "biaffine" in the title of this subsection should be read as "affine". When (R-S2) is solved for $\alpha \in \mathbb{R}$, the restriction that the numerator is positive may be dropped.

For the *general case* (Section 5.4.4), no changes need to be made to drop the restriction that the numerator is positive.

Table 5.1 – Tractable cases when uncertainty in the numerator is independent of the uncertainty in the denominator. ℓ denotes an affine function.

f	g	h_i	$\mathrm{sgn}(f)$	(R-Schaible-S1)
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$g(\boldsymbol{a_0}', \boldsymbol{x})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	≥ 0	no modifications
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$\ell(\boldsymbol{x})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	any	$\alpha \in \mathbb{R}$, (5.5) as an equality

Table 5.2 – Tractable cases when the denominator does not depend on $\boldsymbol{x}$. ℓ denotes an affine function.

f	g	h_i	$\mathrm{sgn}(f)$	(R-S2)
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$g(\boldsymbol{a_0})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	≥ 0	no modifications
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$\ell(\boldsymbol{a_0})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	any	$\alpha \in \mathbb{R}$

Table 5.3 – Tractable cases for the general case. ℓ denotes a biaffine function.

f	g	h_i	$\mathrm{sgn}(f)$
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$g(\boldsymbol{a_0}, \boldsymbol{x})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	≥ 0
$f(\boldsymbol{a_0}, \boldsymbol{x})$	$\ell(\boldsymbol{a_0}, \boldsymbol{x})$	$h_i(\boldsymbol{a_i}, \boldsymbol{x})$	any

5.5 Numerical examples

In this section, we test our method on three examples: a multi-item newsvendor problem (Section 5.5.1), mean-variance optimization (Section 5.5.2) and data envelopment analysis (Section 5.5.3).

5.5.1 Multi-item newsvendor example

In Chapter 2, a multi-item newsvendor problem is solved by minimizing the investment cost under the condition that at least a certain expected profit is made.

We show how to directly optimize the expected return on investment for this example. Let us first recapitulate the problem. The newsvendor buys Q_i units of item i at the beginning of the day. Each item has its associated ordering cost c_i, selling price v_i, salvage price r_i, and unsatisfied demand loss l_i. We assume $r_i \leq v_i + l_i$. During the day the newsvendor faces a demand d_i, resulting in a profit of $\min\{v_i Q_i + l_i(Q_i - d_i) - c_i Q_i, v_i d_i + r_i(Q_i - d_i) - c_i Q_i\}$. The demand is not known in advance, but there are finitely many demand scenarios d_{is} (s in S) that occur with (uncertain) probability p_{is}, independently of other items.

The problem of maximizing expected return on investment can be formulated as:

$$
\text{(R-NV)} \quad \max_{\boldsymbol{Q} \in \mathbb{R}_+^{|I|}, \boldsymbol{u} \in \mathbb{R}^{|I| \times |S|}} \quad \min_{\boldsymbol{p_i} \in \mathcal{U}_i} \quad \frac{\sum_{i \in I} \sum_{s \in S} p_{is} u_{is}}{\sum_{i \in I} c_i Q_i}
$$
$$
\text{s.t.} \quad u_{is} + (c_i - r_i) Q_i \leq d_{is}(v_i - r_i), \ \forall i \in I, \ \forall s \in S
$$
$$
u_{is} + (c_i - v_i - l_i) Q_i \leq -d_{is} l_i, \ \forall i \in I, \ \forall s \in S,
$$

where u_{is} is the contribution to the profit of item i in scenario s, and the convex and compact uncertainty regions $\mathcal{U}_i$ are defined using the Matusita distance, which is a

ϕ–divergence measure (Ben-Tal et al., 2013):

$$\mathcal{U}_i = \left\{ \boldsymbol{p}_i \in \mathbb{R}_+^{|S|} : \sum_{s\in S} p_{is} = 1, \sum_{s\in S} |(\hat{p}_{is})^\alpha - (p_{is})^\alpha|^{1/\alpha} \leq \rho, \ \forall i \in I \right\}.$$

Note that the denominator is strictly positive if at least one item is bought, and also the other assumptions are fulfilled. (R-NV) can be classified under the first special case (Section 5.4.2). Since all functions are affine and the denominator is certain, the Schaible reformulation and the Charnes-Cooper reformulation are equivalent. (R-NV) is therefore equivalent to:

(R-CC-NV)

$$\max_{\boldsymbol{Q}\in\mathbb{R}_+^{|I|}, t\in\mathbb{R}_+, \boldsymbol{u}\in\mathbb{R}^{|I|\times|S|}} \quad \min_{\boldsymbol{p_i}\in\mathcal{U}_i} \quad \sum_{i\in I}\sum_{s\in S} p_{is}u_{is}$$

$$\text{s.t.} \quad u_{is} + (c_i - r_i)\,Q_i \leq d_{is}\,(v_i - r_i)\,t, \ \forall i \in I, \ \forall s \in S$$

$$u_{is} + (c_i - v_i - l_i)\,Q_i \leq -d_{is}l_i t, \ \forall i \in I, \ \forall s \in S$$

$$\sum_{i\in I} c_i Q_i = 1,$$

where $\boldsymbol{Q}$ and $\boldsymbol{u}$ in (R-CC-NV) have to be divided by t to obtain the $\boldsymbol{Q}$ and $\boldsymbol{u}$ in (R-NV). (R-CC-NV) is a linear program with a convex uncertainty region that we solve via its dual as outlined in the introduction. The last of the reformulation steps, a substitution, is not necessary since the uncertainty only appears in the objective. Let x_{is}, y_{is} and z be the dual variables of (CC-NV); then the optimistic dual (OD-CC-NV) is given by:

(OD-CC-NV) $\quad \min \quad z$

$$\text{s.t.} \quad x_{is} + y_{is} = p_{is}, \ \forall i \in I, \ \forall s \in S$$

$$\sum_{s\in S}\{(c_i - r_i)x_{is} + (c_i - v_i - l_i)y_{is}\} + c_i z \geq 0, \ \forall i \in I \tag{5.10}$$

$$\sum_{i\in I}\sum_{s\in S} -d_{is}(v_i - r_i)x_{is} + d_{is}l_i y_{is} \geq 0 \tag{5.11}$$

$$\sum_{s\in S} p_{is} = 1, \ \forall i \in I$$

$$\sum_{s\in S} |(\hat{p}_{is})^\alpha - (p_{is})^\alpha|^{1/\alpha} \leq \rho, \ \forall i \in I$$

$$\boldsymbol{p} \in \mathbb{R}_+^{|S|}, \boldsymbol{x} \in \mathbb{R}_+^{|I|\times|S|}, \boldsymbol{y} \in \mathbb{R}_+^{|I|\times|S|}, z \in \mathbb{R}. \tag{5.12}$$

The optimal value of (OD-CC-NV) is the robust expected return on investment. The corresponding optimal order quantities $\boldsymbol{Q}$ can be derived from the KKT vector of

(OD-CC-NV) by dividing its elements associated with (5.10) by the element corresponding to (5.11). This is the same as dividing $\boldsymbol{Q}$ by t in (R-CC-NV) to undo the Charnes-Cooper transformation.

We solve the problem for the same data as in Chapter 2 with AIMMS 3.11 (Paragon Decision Technology, the Netherlands) and KNITRO 7.0 (Zienna Optimization LLC, USA) with its default settings. Computation errors for negative p_{is} were avoided by using $|p_{is}|^{\alpha}$ instead of $(p_{is})^{\alpha}$. To obtain the nominal solution we take $\rho = 0$, while $\rho = 0.03$ for the robust solution. Solutions were obtained in less than 0.01 seconds. When the probabilities are as expected ($\rho = 0$), the expected return on investment of the nominal solution is 0.297, while for the robust solution it is 0.285. When $\rho = 0.03$ and the worst case probabilities occur for the nominal solution, i.e., the probabilities that minimize the expected return on investment for the nominal solution, the objective value drops to 0.211, while for the robust solution it drops to 0.214. So, the solution indeed becomes more robust, but the difference with the nominal solution is small. We verify if the decision maker could have done better, if he knows beforehand which probability vector realizes. This done by optimizing the nominal model ($\rho = 0$) while setting the probabilty estimates $\hat{p}_{is}$ equal to the worst case probabilities for the robust solution. This gives the so-called perfect hindsight solution. The objective value is as low as 0.214. So, even though the robust objective could deteriorate substantially, there is no other solution that performs better.

5.5.2 Mean-variance optimization

We are to present an example that involves a trade-off between mean and variance. This trade-off is commonly used in portfolio optimization, including the Modern Portfolio Theory (MPT) founded by Markowitz (1952), where the goal is to select the right mix of assets. In contrast to MPT, we do not impose that the expected returns on the assets and the covariance matrix are fully known. Instead, we assume that finitely many scenarios s (in S) for the future can be identified along with unknown probabilities of occurence p_s, which are estimated by $\hat{p}_s$. The return of asset i in scenario s is a constant r_{is}, so when x_i units of money are invested in asset i, the return in scenario s is given by $u_s = \sum_{i \in I} r_{is} x_i$ (possibly negative). The expected return and variance are given by:

$$\mathbb{E}(\text{return}) := \sum_{s \in S} p_s u_s$$

$$\text{Var}(\text{return}) := \sum_{s \in S} p_s \left(u_s - \mathbb{E}(\text{return})\right)^2 \tag{5.13}$$

$$= \sum_{s \in S} p_s u_s^2 - \left(\sum_{s \in S} p_s u_s\right)^2. \tag{5.14}$$

To remain in the minimization framework, the objective is to minimize the variance-to-mean ratio (or the dispersion index). The robust optimization problem is given by:

$$(\text{R-I}) \quad \min_{\boldsymbol{x} \in \mathbb{R}_+^{|I|}, \boldsymbol{u} \in \mathbb{R}^{|S|}} \quad \max_{\boldsymbol{p} \in \mathcal{U}} \quad \frac{\sum_{s \in S} p_s u_s^2 - \left(\sum_{s \in S} p_s u_s\right)^2}{\sum_{s \in S} p_s u_s}$$

$$\text{s.t.} \quad u_s = \sum_{i \in I} r_{is} x_i, \ \forall s \in S$$

$$\sum_{i \in I} x_i = C$$

$$\sum_{s \in S} p_s u_s > 0, \ \forall \boldsymbol{p} \in \mathcal{U}.$$

The last two constraints ensure that C units of money are invested, and that this model has a feasible solution only if the expected profit is positive. The numerator is convex in u_s (from (5.13)) and concave in p_s (from (5.14)). The denominator is clearly concave in p_s and convex in u_s. Moreover, the numerator is non-negative and the denominator is positive on the feasible region. For the uncertainty region we use the modified χ^2–distance as a ϕ–divergence measure, which can be justified by statistical theory (Ben-Tal et al., 2013):

$$\mathcal{U}_0 = \left\{ \boldsymbol{p} \in \mathbb{R}_+^{|S|} : \sum_{s \in S} p_s = 1, \sum_{s \in S} \frac{(p_s - \hat{p}_s)^2}{\hat{p}_s} \leq \rho \right\}.$$

(R-I) is not one of the special cases, so we solve this problem using the general method. In order to formulate (5.7) explicitly, we first derive some conjugate functions. The conjugate for f is from Ben-Tal et al. (2014, Ex. 25).

$$f(\boldsymbol{p}, \boldsymbol{u}) = \sum_{s \in S} p_s u_s^2 - \left(\sum_{s \in S} p_s u_s\right)^2 \qquad f_*(\boldsymbol{v}, \boldsymbol{u}) = \sup_z \left\{ -\frac{z^2}{4} : u_s^2 + u_s z = v_s, \ \forall s \in S \right\}$$

$$g(\boldsymbol{p}, \boldsymbol{u}) = \sum_{s \in S} p_s u_s \qquad g^*(\boldsymbol{v}, \boldsymbol{u}) = \begin{cases} 0, & \text{if } v_s = u_s, \ \forall s \in S \\ \infty, & \text{otherwise} \end{cases}$$

$$\tau_{i1}(\boldsymbol{p}) = \max_s \{-p_s\} \qquad \tau_{i1}^*(\boldsymbol{v}) = \begin{cases} 0, & \text{if } v_s \leq 0, \ \forall s \in S \text{ and } \sum_{s \in S} v_s \geq -1 \\ \infty, & \text{otherwise} \end{cases}$$

$$\tau_{i2}(\boldsymbol{p}) = \sum_{s \in S} p_s - 1 \qquad \tau_{i2}^*(\boldsymbol{v}) = \begin{cases} 1, & \text{if } v_s = 1, \ \forall s \in S \\ \infty, & \text{otherwise} \end{cases}$$

$$\tau_{i3}(\boldsymbol{p}) = 1 - \sum_{s \in S} p_s \qquad \tau_{i3}^*(\boldsymbol{v}) = \begin{cases} -1, & \text{if } v_s = -1, \ \forall s \in S \\ \infty, & \text{otherwise} \end{cases}$$

$$\tau_{i4}(\boldsymbol{p}) = \sum_{s \in S} \frac{(p_s - \hat{p}_s)^2}{\hat{p}_s} - \rho \qquad \tau_{i4}^*(\boldsymbol{v}) = \rho + \sum_{s \in S} \hat{p}_s \left(\frac{1}{4} v_s^2 + v_s \right).$$

Plugging in these formulas in (5.7) yields $F(\alpha) =$

$$\min \quad w \quad \text{s.t.} \quad t_{02} - t_{03} + \rho t_{04} + \sum_{s \in S} \hat{p}_s \left(\frac{(v_{04})_s^2}{4t_{04}} + (v_{04})_s \right) + \frac{z^2}{4} \leq w \qquad (5.15)$$

$$(v_{01})_s \leq 0 \quad (v_{11})_s \leq 0, \ \forall s \in S$$

$$\sum_{s \in S}(v_{01})_s \geq -t_{01} \quad \sum_{s \in S}(v_{11})_s \geq -t_{11}$$

$$u_s^2 + u_s z = u_s \alpha + (v_{01})_s + t_{02} - t_{03} + (v_{04})_s, \ \forall s \in S \qquad (5.16)$$

$$t_{12} - t_{13} + \rho t_{14} + \sum_{s \in S} \hat{p}_s \left(\frac{(v_{14})_s^2}{4t_{14}} + (v_{14})_s \right) < 0$$

$$u_s + (v_{11})_s + t_{12} - t_{13} + (v_{14})_s = 0, \ \forall s \in S$$

$$u_s = \sum_{i \in I} r_{is} x_i, \ \forall s \in S$$

$$\sum_{i \in I} x_i = C$$

$$\boldsymbol{t} \in \mathbb{R}_+^{2 \times 4}, \boldsymbol{u} \in \mathbb{R}^{|S|}, \boldsymbol{v} \in \mathbb{R}^{2 \times 4}, w \in \mathbb{R}, \boldsymbol{x} \in \mathbb{R}^{|I|}, z \in \mathbb{R}.$$

This problem is nonconvex because of the product $u_s z$ in (5.16). Similar to (Yanıkoğlu et al., 2013, Theorem 1) the problem can be made convex by replacing $t_{02} - t_{03}$ in (5.15) with $u_s^2 + u_s z - u_s \alpha - (v_{01})_s - (v_{04})_s$ and omitting (5.16) from the problem. Constraint (5.15) then becomes:

$$\left(u_s + \frac{z}{2} \right)^2 - u_s \alpha - (v_{01})_s - (v_{04})_s + \rho t_{04} + \sum_{s' \in S} \hat{p}_{s'} \left(\frac{(v_{04})_{s'}^2}{4t_{04}} + (v_{04})_{s'} \right) \leq w, \ \forall s \in S,$$
$$(5.17)$$

which is jointly convex in all variables. In order to improve the tractability of computing $F(\alpha)$, we cast it as a conic quadratic problem. The only complicating terms are $(v_{04})_s^2/(4t_{04})$, which can be reformulated using a standard trick. Constraint (5.17) is satisfied if and only if there exists auxiliary variables y_s such that the following inequalities are satisfied:

$$\left(u_s + \frac{z}{2} \right)^2 - u_s \alpha - (v_{01})_s - (v_{04})_s + \rho t_{04} + 2 \sum_{s' \in S} \hat{p}_{s'} \left(y_{s'} + (v_{04})_{s'} \right) \leq w, \ \forall s \in S$$

$$\left\| \begin{pmatrix} (v_{04})_s \\ \frac{y_s}{2} - 2t_{04} \end{pmatrix} \right\|_2 \leq \frac{y_s}{2} + 2t_{04}, \ \forall s \in S.$$

The problem (R-I) can now be solved by determining the root of F.

We perform a numerical analysis on 10 items and a generated data set of 50 scenarios. In order to incorporate correlations, we first construct a covariance matrix

$\boldsymbol{AA}^\top$, where $\boldsymbol{A}$ is a 10×10 matrix whose entries are uniformly and independently distributed on $[-0.5, 0.5]$. Then, to reflect the idea that a higher risk gives a higher expected return, a vector of expected returns $\boldsymbol{\mu}$ is constructed with a linear mapping on the variances of the items. The item with the smallest variance gets an expected return of 0.01, and the item with the largest variance gets an expected return of 0.20. Finally, the scenarios are drawn, each from a multivariate normal distribution with the constructed mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{AA}^\top$. We solve the model for $\rho = 1$, $\hat{p}_s = 0.02$ for all s, and $C = 100$ to obtain a robust solution using YALMIP (Löfberg, 2012) and MOSEK (Mosek ApS, Denmark) with their default settings. For this value of ρ the probabilities can vary, on average, between 0 and 0.04. Additionally, we solve the same problem for $\rho = 0$, i.e., when $p_s = \hat{p}_s$, to obtain a nonrobust solution. We use bisection search on the interval determined by (5.8) and (5.9) and stopped when the interval width was less than 10^{-10}. One step in bisection search takes around 2 seconds, of which around 7% is spent by MOSEK.

The convergence of the bisection method turns out to be adequate. Let $\boldsymbol{x_i}$ denote the solution in iteration i, and let $\boldsymbol{x}^*$ denote the final solution. The initial search interval is $[0.70, 20.09]$. The solution $\boldsymbol{x_1}$, obtained from solving $F((0.70 + 20.09)/2)$, is far from optimal: $||\boldsymbol{x_1} - \boldsymbol{x}^*||_\infty \approx 2.2$. In each three or four iterations, $\boldsymbol{x_i}$ gains one extra digit of accuracy. After 22 iterations, the accuracy has improved to $||\boldsymbol{x_{22}} - \boldsymbol{x}^*||_\infty \approx 4.1 \cdot 10^{-7}$. The algorithm terminates after 37 iterations, with no apparent improvement following the 22^{nd} iteration. Since $||\boldsymbol{x}^*||_2 \approx 44.9$, the error after 22 iterations is relatively small.

When $p_s = \hat{p}_s$ for all s, the mean-variance ratio of the nominal solution (which is 6.34) is indeed lower than that of the robust solution (which is 6.45). For both solutions we determined the worst case $\hat{\boldsymbol{p}}$ and the corresponding objective value. The objective of the robust solution (which is 18.62) is slightly better than that of the nominal solution (which is 18.98). This shows that uncertainty may cause a factor three deterioration of the objective value. Relative to this large difference, the difference between the two solutions is small. So, the nominal solution performs quite well for this example. For the worst case probabilities for the robust solution, we have computed the optimal portfolio as if these solutions were known beforehand (perfect hindsight solution). The objective value equals that of the robust solution. So, even though the robust objective could deteriorate substantially, there is no other solution that performs better.

5.5.3 Data envelopment analysis

Data Envelopment Analysis (DEA) is a tool to estimate the efficiency of different decision making units (DMUs) based on their inputs and outputs. DEA was orig-

inally introduced for not-for-profit companies, e.g., schools where inputs could be number of teacher hours and number of students per class, and outputs could be arithmetic scores and psychological tests of student attitudes, e.g., toward the community (Charnes et al., 1978). The applicability of DEA is not limited to nonprofit organizations. A reference list of more than 4,000 publications on DEA is given in (Emrouznejad et al., 2008).

Let n_i and n_o denote the number of in- and outputs, respectively. The efficiency of a DMU is defined as the largest fraction of weighted outputs divided by weighted inputs, given that the efficiency of the other DMUs is at most 1:

$$\text{(DEA)} \quad \max_{\boldsymbol{u} \in \mathbb{R}_+^{n_o}, \boldsymbol{v} \in \mathbb{R}_+^{n_i}} \quad \frac{\boldsymbol{u}^\top \boldsymbol{y_0}}{\boldsymbol{v}^\top \boldsymbol{x_0}} \quad \text{s.t.} \quad \frac{\boldsymbol{u}^\top \boldsymbol{y_i}}{\boldsymbol{v}^\top \boldsymbol{x_i}} \leq 1, \ \forall i \in I,$$

where $\boldsymbol{x_i}$ and $\boldsymbol{y_i}$ are the vectors of inputs and outputs of DMU i, and $\boldsymbol{u}$ and $\boldsymbol{v}$ are the non-negative weights.

The inputs and outputs are model parameters that have to be acquired from each DMU and are affected by measurement errors. Especially when a single DMU represents a group of smaller business units or is a pool of all activities in a certain region, and the inputs and outputs are aggregated, errors become practically inevitable. There have been many attempts to incorporate uncertainty in DEA. For an overview, e.g., see (Shokouhi et al., 2014). Since our focus is on RO, we only discuss the three papers that are relevant. The first only considers uncertain outputs (Sadjadi and Omrani, 2008). The second considers jointly uncertain inputs and outputs (Shokouhi et al., 2010). Unfortunately, the robust counterpart in this chapter is constructed in an ad-hoc manner that results in a nonconvex formulation, for which it is not clear whether globally optimal solutions were found. The third considers either uncertain inputs or uncertain outputs (Wang and Wei, 2010). In the last two papers, a simulation study is performed to quantify the improvement offered by the robust solution. For each randomly drawn set of inputs and outputs, they compute the relative efficiencies with the $\boldsymbol{u}$ and $\boldsymbol{v}$ obtained from the robust solution. However, in our view, when the inputs and outputs are fully known, the relative efficiencies can only be computed by optimizing DEA for those known inputs and outputs. Our results are therefore different in two ways. First, we consider both uncertain inputs and uncertain outputs and solve the correct problem. Second, we perform a valid simulation study to verify whether the robust solution is better than the nominal solution.

In this section, we take the data from Shokouhi et al. (2010), solve the correct robust counterpart with uncertain inputs and uncertain outputs, and perform a valid comparison between the robust and the nonrobust solution. In this data set there

are five DMUs, two inputs and two outputs. The in- and outputs are uncertain, but known to reside in given intervals (Table 5.4).

In order to get in the minimization framework, the objective of (DEA) is replaced with its reciprocal. The optimal solution of the robust counterpart of (DEA) then corresponds to the reciprocal of the root of:

$$F(\alpha) = \min_{\boldsymbol{u}\in\mathbb{R}_+^{n_o}, \boldsymbol{v}\in\mathbb{R}_+^{n_i}, w\in\mathbb{R}} \quad w \quad \text{s.t.} \quad \boldsymbol{v}^\top \boldsymbol{x_0} - \alpha \boldsymbol{u}^\top \boldsymbol{y_0} \le w, \ \forall(\boldsymbol{x},\boldsymbol{y})\in\mathcal{U},$$

$$\boldsymbol{u}^\top \boldsymbol{y_i} \le \boldsymbol{v}^\top \boldsymbol{x_i}, \ \forall(\boldsymbol{x},\boldsymbol{y})\in\mathcal{U}, \ \forall i\in I.$$

Following Shokouhi et al. (2010), we take the Bertsimas and Sim uncertainty region:

$$\mathcal{U} = \{(\boldsymbol{x},\boldsymbol{y})\in\mathbb{R}^{(|I|+1)\times n_i}\times\mathbb{R}^{(|I|+1)\times n_o} : \ x_{ij} = \bar{x}_{ij} + \zeta_{ij}^x\Delta\hat{x}_{ij},$$
$$y_{ij} = \bar{y}_{ij} + \zeta_{ij}^y\Delta\hat{y}_{ij}, \quad ||\text{vec}(\boldsymbol{\zeta^x},\boldsymbol{\zeta^y})||_\infty \le 1, \quad ||\text{vec}(\boldsymbol{\zeta^x},\boldsymbol{\zeta^y})||_1 \le \Gamma\},$$

where $\bar{x}_{ij}$ and $\bar{y}_{ij}$ are the midpoints, $\Delta\hat{x}_{ij}$ and $\Delta\hat{y}_{ij}$ are the half-widths of the uncertainty intervals, and the vec operator stacks the columns of the matrix arguments into a single vector. For robust *LP*, this set has the property that when Γ is integer, it controls the number of uncertain elements that can deviate from their nominal values (Bertsimas and Sim, 2004). This property also holds for a robust LFP, since $F(\alpha)$ is a robust LP.

The optimal weights $\boldsymbol{u}$ and $\boldsymbol{v}$ depend on the actual inputs and outputs. One may therefore be inclined to use Adjustable Robust Optimization (ARO) when Γ is larger than the dimensions of $\boldsymbol{x_i}$ and $\boldsymbol{y_i}$ added. Consequently, $\boldsymbol{u}$ and $\boldsymbol{v}$ are replaced by functions of the uncertain parameters. Unfortunately, even in the simple case of affine decision rules, this is often intractable. In the constraints of $F(\alpha)$, $\boldsymbol{u}$ and $\boldsymbol{v}$ are multiplied with uncertain parameters, which yields constraints that are quadratic in the uncertain parameters. These can currently only be solved efficiently for ellipsoidal uncertainty sets.

We use bisection search on the interval determined by (5.8) and (5.9) to determine the root of $F(\alpha)$, and stop when the interval width is less than 10^{-4} (which turns out to be accurate enough for ranking the DMUs). $F(\alpha)$ is computed using YALMIP and MOSEK with their default settings, and takes a few tenths of a second on a normal desktop computer where MOSEK accounts for around 10% of that time. The time it takes to compute $F(\alpha)$ turns out to be approximately constant, so independent of the remaining width of the interval and independent of the size of the uncertainty region Γ. The root of $F(\alpha)$ is determined in a few seconds.

We computed the robust efficiencies of the DMUs for Γ ranging between 0 and 4 in steps of 0.1, since each constraint has at most four uncertain parameters. For $\Gamma \le 0.2$, the list of DMUs ranked from most to least efficient, is 1, 2, 3, 5, 4. For

Table 5.4 – Data set for the DEA example of Section 5.5.3.

DMU$_i$	Input 1	Input 2	Output 1	Output 2
1	[14,15]	[0.06, 0.09]	[157, 161]	[28, 40]
2	[4,12]	[0.16, 0.35]	[157, 198]	[21, 29]
3	[10,17]	[0.10, 0.70]	[143, 159]	[28, 35]
4	[12,15]	[0.21, 0.48]	[138, 144]	[21, 22]
5	[19,22]	[0.12, 0.19]	[158, 181]	[21, 25]

$\Gamma \geq 0.3$, DMUs 3 and 5 switch positions. Hence, DMU 5 is more efficient than DMU 3 when $\Gamma \geq 0.3$. We have tried to verify this claim by running 100 simulations where in each simulation we uniformly drew inputs and outputs from the uncertainty region, solved (DEA) for each set of inputs and outputs, and ranked the DMUs based on efficiency. In 76 out of 100 simulations, DMU 3 was more efficient than DMU 5. This result advocates against the use of RO in DEA, since it shows that for $\Gamma = 0$ (i.e., the nonrobust solution), the ranking is better than for $\Gamma \geq 1$. We have also performed the simulation with more extreme data, by drawing the inputs and outputs only from the endpoints of their uncertainty intervals. This yielded similar results. Other experiments, where we used an ellipsoidal uncertainty region instead of the Bertsimas and Sim uncertainty region, or where we used the nominal objectives (based on $\bar{\boldsymbol{x}}$ and $\bar{\boldsymbol{y}}$) but kept the uncertain constraints, also yielded similar results.

5.6 Conclusions

We have shown how RO can be applied to FPs as a method to deal with uncertain data. The method has been tested on three problems. In all three examples we observe that the nominal solution, which is obtained by solving the deterministic problem, is severely affected by uncertainty. Surprisingly, this also holds for the robust solution, and in none of the examples the robust solution offers a significant improvement even when comparing worst case performance.

The first question that arises, is why the nominal solution performs so well. We try to answer this question for the mean-variance optimization problem, and note that the explanation for the multi-item newsvendor problem is similar. For a given solution, the worst case for the mean is when the probability vector is a unit vector, that assigns unit weight to the scenario with lowest return. For the variance, the worst case is when the scenarios with the lowest and highest return each occur with probability 0.5. For a robust solution w.r.t. the mean value, the scenario with lowest return should be optimized, whereas for the variance, the returns in the scenarios

with lowest and highest return should be close to each other. The nominal solution simultaneously maximizes the expected value and minimizes the variance. While not identical to the robust objective, it contains some aspects of it. For example, the mean is a weighted sum that contains the return for the scenario with lowest return.

The second question that arises, is why there is a realization of the uncertain parameters in Sections 5.5.1 and 5.5.2, for which no solution can outperform the robust solution; even if the former is optimized as if the realization of the uncertain parameters are known beforehand. This turns out to be due to Sion's minimax theorem (Sion, 1958). Our assumptions ensure that $f(\boldsymbol{a_0}, \boldsymbol{x})/g(\boldsymbol{a_0}, \boldsymbol{x})$ is quasi-convex in $\boldsymbol{x}$ (for fixed $\boldsymbol{a_0}$), quasi-concave in $\boldsymbol{a_0}$ (for fixed $\boldsymbol{x}$), and continuous, and that the uncertainty set is compact and convex, and the feasible set for $\boldsymbol{x}$, say X, is convex. Therefore, by Sion's minimax theorem, $\max_{\boldsymbol{a_0} \in \mathcal{U}} \min_{\boldsymbol{x} \in X} f(\boldsymbol{a_0}, \boldsymbol{x})/g(\boldsymbol{a_0}, \boldsymbol{x}) = \min_{\boldsymbol{x} \in X} \max_{\boldsymbol{a_0} \in \mathcal{U}} f(\boldsymbol{a_0}, \boldsymbol{x})/g(\boldsymbol{a_0}, \boldsymbol{x})$. This no longer holds when there is uncertainty in the constraints, since the feasible region X changes when the values for the uncertain parameters are known.

So, the robust solution is good in the sense that it cannot be improved in the worst case, even if the values of the uncertain parameters are known beforehand. On the other hand, the nominal solution performs well, at least in the examples studied. It shall be interesting to see the difference in real-life examples, especially with uncertainty in the constraints.

5.A   The importance of convexity conditions

We provide a short example to stress the importance of the convexity/concavity conditions on f and g. The second numerical example in (Lin and Sheu, 2005) is:

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \max_{a \in [0,1]} \quad \frac{a^2 x_1 x_2 + x_1^{2a} + a x_3^3}{5(a-1)^2 x_1^4 + 2x_2^2 + 4a x_3} \quad \text{s.t.} \quad 0.5 \le x_i \le 5, \quad i = 1, 2, 3.$$

This problem does not satisfy the convexity/concavity conditions from Section 5.4. Lin and Sheu claim that $\boldsymbol{x} = (0.5 \quad 1.5 \quad 0.5)$ and $a = 0$ is optimal with a value of 0.21 (reported as -0.21), but $\boldsymbol{x} = (0.5 \quad 5 \quad 0.5)$ and $a = 1$ is a better solution (maybe still not optimal) since the corresponding value is 0.06.

5.B   On the result by Kaul et al. (1986)

This appendix shows a mistake in the paper by Kaul et al. (1986). Essentially, they formulate the dual of:

$$\min_{\alpha\in\mathbb{R}_+,\boldsymbol{x}\in\mathbb{R}^n_+} \alpha \quad \text{s.t.} \quad \frac{b_0 + \boldsymbol{b}^\top\boldsymbol{x}}{c_0 + \boldsymbol{c}^\top\boldsymbol{x}} \leq \alpha, \ \forall(b_0,\boldsymbol{b})\in\mathcal{U}_1\times\mathcal{U}_2, \ \forall(c_0,\boldsymbol{c})\in\mathcal{U}_3\times\mathcal{U}_4,$$

$$\boldsymbol{Ax} \leq \boldsymbol{d}.$$

Note that $\boldsymbol{x}$ is non-negative. In their Lemma 2.1 they claim that the worst case $(c_0, \boldsymbol{c})$ does not depend on $\boldsymbol{x}$, and is given by $c_0^* = \min_{c_0\in\mathcal{U}_3}\{c_0\}$ and $\boldsymbol{c}^*$ with components $c_i^* = \min_{\boldsymbol{c}\in\mathcal{U}_4}\{c_i\}$. This implicitly assumes that $\boldsymbol{c}^*$ is a member of $\mathcal{U}_4$, which is not always true. The mistake becomes clear in their numerical example where they use $\boldsymbol{c}^* = [4;2]$, which is not in the uncertainty set. Consequently, the proposed approach gives the wrong dual problem and a suboptimal solution. Our results in Section 5.4.2 can provide the correct dual problem under milder conditions on the uncertainty sets.

Acknowledgments

Approximating the Pareto set of multiobjective linear programs via Robust Optimization

Abstract We consider problems with multiple linear objectives and linear constraints and use Adjustable Robust Optimization and Polynomial Optimization as tools to approximate the Pareto set with polynomials of arbitrarily large degree. The main difference with existing techniques is that we optimize a single (extended) optimization problem that provides a polynomial approximation whereas existing methods iteratively construct a piecewise linear approximation. One of the advantages of the proposed method is that it is more useful for visualizing the Pareto set.

6.1   Introduction

Multiobjective optimization problems (MOPs) have a vector valued objective function $\boldsymbol{f} = [f_i]_{1 \le i \le k}$, where each f_i is a separate objective. Often it is not possible to have optimal values for all f_i simultaneously, e.g. in portfolio optimization it is not possible to have minimum risk and maximum return at the same time. Another example is intensity-modulated radiation therapy, where tumour coverage is balanced with sparing of surrounding organs (Craft and Bortfeld, 2008; Rennen et al., 2011). Optimization of a vector valued function involves a trade-off between two or more objectives f_i $(1 \le i \le k)$.

A simple way to deal with multiple objectives is by assigning an importance factor $w_i > 0$ to each objective and optimizing $\sum_{i=1}^{k} w_i f_i$ (we make the assumption that all f_i should be minimized w.l.o.g.). If such importance factors are not known a priori, a Pareto set (PS) allows the decision maker to make the trade-off after optimization. The set PS consists of all objective vectors $\boldsymbol{f}$ in which one or more objectives can not be improved without deteriorating one or more other objectives. Overviews of

MOPs and approximation methods can be found in (Branke et al., 2008; Ehrgott, 2005; Miettinen, 1999).

In practice often approximations of PS are used, since the exact PS can often not be found. In literature many different approximation methods are proposed. It is desirable to approximate PS with as few optimization runs as possible (Rennen et al., 2011). A well-known class of such approximation methods is sandwich methods. Sandwich methods (Rennen et al., 2011) produce piecewise linear approximations in which between PS is located. In each iteration an optimization problem is solved, which leads to adding one or more facets to the approximations. All other methods (again see (Branke et al., 2008; Ehrgott, 2005; Miettinen, 1999)) are also sequential, i.e. in each iteration one has to solve an optimization problem which leads to an improvement of the approximation.

In this chapter, we focus on approximating PS for linear programming and propose a totally different way than those in the literature. The first difference is that our method is not sequential, but generates the approximation by solving one extended optimization problem. The second difference is that the final approximation is not piecewise linear but a polynomial. The way we construct this approximation is by using techniques from Adjustable Robust Optimization (ARO) (Ben-Tal et al., 2009a) and Polynomial Optimization (Laurent, 2009). We first explain the link to ARO. The Pareto set is seen as a function of the *uncertain parameters* $f_1, ..., f_{k-1}$. The area of interest, i.e. the domain for $f_1, ..., f_{k-1}$ for which we would like to approximate the Pareto set, is considered as the *uncertainty region*. All variables in the linear program are made *adjustable* in the parameters $f_1, ..., f_{k-1}$. We use polynomials for the decision rule, and use Polynomial Optimization theory to reformulate the resulting robust counterpart into a Semi-Definite Programming (SDP) problem. Since the number of uncertain parameters (i.e. $k - 1$) is often low, the sizes of the LMIs in the SDP are relatively small. Notice that in our approach ARO is merely used as a tool, uncertainty in the data is not considered.

The approach proposed in this chapter has the following advantages:

The first advantage is that the final approximation is more tractable for navigating through PS. The polynomial representation is useful for the user to visualize PS for selecting the final solution. This is the reason why in (Goel et al., 2007), afer determining points in the feasible region close to or on PS, polynomial regression is used to obtain a tractable representation of PS (a so-called response surface). Our method finds such a tractable representation directly, with the additional advantage that it is guaranteed to lie in the feasible region. Sometimes the decision maker needs a local approximation of PS around a given solution. Zhang et al. (2000) formulate and test a method that gives a local quadratic approximation of (not necessarily

convex) PS, but this approximation is neither an inner nor an outer approximation. Our method gives a polynomial of arbitrary degree and is guaranteed to give an inner approximation.

The second advantage is that our approach can be used to certify with a single optimization run that a given set V is dominated, i.e. that all elements of the set are dominated. If our method finds a feasible solution, then this solution is a certificate that the set V is dominated.

The third advantage is that the explicit polynomial approximation can be used in optimization problems with Pareto constraints. Such problems contain constraints that enforce that the solution should be (near) Pareto optimal for a certain multi-objective linear program. Examples of these problems can be found in e.g. (Hackman and Passy, 2002). A Pareto constraint can be replaced by the explicit polynomial approximation found by our method.

The fourth advantage is the possibility to quickly determine the shape of PS. (Craft and Bortfeld, 2008), e.g., show that in IMRT the set of feasible objective vectors is often "long and narrow" and therefore a linear approximation of PS suffices. This linear approximation can be easily obtained by our approach. Finally, after an initial approximation, the most interesting subregion can be selected, followed by one or more iterations of approximation and selection. An interesting subregion can also be used as input for another algorithm that explores it more carefully.

This method also has five disadvantages. First, the resulting problem is often an SDP while the original problem is LP. Only in case the approximation is linear and the region of interest is polyhedral, the resulting problem is LP. Note that this is still an interesting case; see (Craft and Bortfeld, 2008) that uses linear approximations in IMRT problems. Second, our method requires the region of interest to be known. Sandwich algorithms are capable of exploring the region of interest (Rennen et al., 2011). Third, it is difficult to approximate the Pareto set at its vertices, because polynomials are smooth functions. The further the vertex angle from 180°, the more difficult it is to approximate it. However, in some cases a smooth approximation is desirable, see e.g. (Mello et al., 2002). Fourth, the method can not be extended to nonlinear multiobjective problems with current ARO technology. Methods for approximating nonlinear MOPs can be found in (Luque et al., 2012; Utyuzhnikov et al., 2009). Fifth, while we show (Appendix 6.A) that the method can also produce outer approximations, this is practically impossible due to computational issues.

6.2 Notation

We use the notation from (Rennen et al., 2011) with some minor changes.

Throughout this chapter, we use the following orderings of vectors. Let $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$ with $n \geq 2$. With x_i, we denote the i^{th} element of the vector $\boldsymbol{x}$. To enumerate different vectors, we use superscripts. When ordering two vectors, we use:

- $\boldsymbol{x} < \boldsymbol{y} \Leftrightarrow x_i < y_i$ for all $i = 1, \ldots, n$.

- $\boldsymbol{x} \lneqq \boldsymbol{y} \Leftrightarrow x_i \leq y_i$ for all $i = 1, \ldots, n$ and $\boldsymbol{x} \neq \boldsymbol{y}$.

- $\boldsymbol{x} \leq \boldsymbol{y} \Leftrightarrow x_i \leq y_i$ for all $i = 1, \ldots, n$.

The symbols $>$, $\gneqq$ and $\geq$ are defined accordingly. We furthermore define the set $\mathbb{R}_-^n = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{x} \leq 0\}$. If $X \subseteq \mathbb{R}^n$, then we define $X + \mathbb{R}_-^n = \{\boldsymbol{x} + \boldsymbol{y} : \boldsymbol{x} \in X, \boldsymbol{y} \in \mathbb{R}_-^n\}$. The sets $\mathbb{R}_+^n$ and $X + \mathbb{R}_+^n$ are defined accordingly.

In this chapter, we consider the following multi-objective optimization problem:

$$\min \quad \boldsymbol{f}(\boldsymbol{x}) = [(\boldsymbol{c^1})^\top \boldsymbol{x}, \ldots, (\boldsymbol{c^k})^\top \boldsymbol{x}]^\top$$
$$A\boldsymbol{x} \leq \boldsymbol{b},$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is the optimization variable, $\boldsymbol{c^i} \in \mathbb{R}^n$ are the objective vectors, and $A \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^m$.

As it is generally impossible to find a feasible x that minimizes all objectives at the same time, our aim is to find a set of so-called Pareto optimal solutions.

Definition 3
An objective vector $\boldsymbol{f}(\boldsymbol{x})$, for $\boldsymbol{x}$ such that $A\boldsymbol{x} \leq \boldsymbol{b}$, is (strongly) dominated if there exists an $\widetilde{\boldsymbol{x}}$ such that $A\widetilde{\boldsymbol{x}} \leq \boldsymbol{b}$ and $\boldsymbol{f}(\widetilde{\boldsymbol{x}}) < \boldsymbol{f}(\boldsymbol{x})$. If no such $\widetilde{\boldsymbol{x}}$ exists, the objective vector $\boldsymbol{f}(\boldsymbol{x})$ is weakly Pareto optimal.
An objective vector $\boldsymbol{f}(\boldsymbol{x})$, for $\boldsymbol{x}$ such that $A\boldsymbol{x} \leq \boldsymbol{b}$, is weakly dominated if there exists an $\widetilde{\boldsymbol{x}}$ such that $A\widetilde{\boldsymbol{x}} \leq \boldsymbol{b}$ and $\boldsymbol{f}(\widetilde{\boldsymbol{x}}) \lneqq \boldsymbol{f}(\boldsymbol{x})$. If no such $\widetilde{\boldsymbol{x}}$ exists, the objective vector $\boldsymbol{f}(\boldsymbol{x})$ is (strongly) Pareto optimal.

The set of Pareto optimal solutions is denoted by PS. An inner and outer approximation of the Pareto set are defined as follows:

Definition 4 *A set $IPS \subseteq \mathbb{R}^k$ is an inner approximation of PS if it satisfies $IPS \subseteq PS + \mathbb{R}_+^k$.*

Definition 5 *A set $OPS \subseteq \mathbb{R}^k$ is an outer approximation of PS if it satisfies $PS \subseteq OPS + \mathbb{R}_+^k$.*

We will approximate PS with polynomials on multidimensional sets. The following definitions are used to define the degree of a polynomial.

Definition 6 *A monomial of degree d in $\boldsymbol{x} \in \mathbb{R}^n$ with powers $\boldsymbol{a} \in \mathbb{R}^n$ such that $\sum_{i=1}^{n} a_i = d$, is defined by $\prod_{i=1}^{n} x_i^{a_i}$.*

Definition 7 *A polynomial of degree d in $\boldsymbol{x} \in \mathbb{R}^n$ is defined as the sum of monomials in $\boldsymbol{x}$ of degree up to d. The degree of a polynomial f is denoted as $\deg(f)$.*

6.3   Inner approximation

Let $U \subseteq \mathbb{R}^{k-1}$ be the domain of interest for $((\boldsymbol{c^1})^\top \boldsymbol{x}, (\boldsymbol{c^2})^\top \boldsymbol{x}, \ldots, (\boldsymbol{c^{k-1}})^\top \boldsymbol{x})$. For a fixed $\boldsymbol{u}$ in U, the following optimization problem determines a single weakly Pareto optimal solution (Miettinen, 1999, Thm. 3.2.1):

$$\min_{x} \quad (\boldsymbol{c^k})^\top \boldsymbol{x}$$
$$(\boldsymbol{c^i})^\top \boldsymbol{x} \le u_i \qquad\qquad i = 1, 2, \ldots, k-1$$
$$A\boldsymbol{x} \le \boldsymbol{b}.$$

If the solution x is unique, it is (strongly) Pareto optimal (Miettinen, 1999, Thm. 3.2.4). For every $\boldsymbol{u}$, there will be a different optimal $\boldsymbol{x}$. So, we want to solve for $\boldsymbol{x}$ as a function of $\boldsymbol{u}$. The constraints should hold for all $\boldsymbol{x}(\boldsymbol{u})$ for which $\boldsymbol{u}$ is in U, and the goal is e.g. to minimize the average objective:

$$\min_{\boldsymbol{x}(\boldsymbol{u})} \quad \int_U (\boldsymbol{c^k})^\top \boldsymbol{x}(\boldsymbol{u}) du \tag{6.1a}$$
$$(\boldsymbol{c^i})^\top \boldsymbol{x}(\boldsymbol{u}) \le u_i \qquad \forall \boldsymbol{u} \in U, \quad i = 1, 2, \ldots, k-1 \tag{6.1b}$$
$$A\boldsymbol{x}(\boldsymbol{u}) \le \boldsymbol{b} \qquad \forall \boldsymbol{u} \in U. \tag{6.1c}$$

This is an ARO problem, where $\boldsymbol{u}$ is the uncertain parameter, U is the uncertainty region, and x is an adjustable variable Ben-Tal et al. (2009a). It is difficult to optimize over functions, therefore ARO uses parameterized functions for adjustable variables. The adjustable variables then become expressions that are linear in the parameters. For instance, if we take a linear parameterization $\boldsymbol{x}(\boldsymbol{u}) = \boldsymbol{\alpha^0} + \alpha^1 \boldsymbol{u}$, the parameters are $\boldsymbol{\alpha^0} \in \mathbb{R}^n$ and $\alpha^1 \in \mathbb{R}^{n \times (k-1)}$. After substituting $\boldsymbol{x}(\boldsymbol{u})$ in the problem (6.1a)-(6.1c), an ARO problem with constraints that are linear in $\boldsymbol{\alpha^0}$ and α^1 remains. In general, the tractability of (6.1a)-(6.1c) depends on the class of functions considered for x and the set U. Given a solution to this optimization problem, an inner approximation is given by
$\{((\boldsymbol{c^1})^\top \boldsymbol{x}(\boldsymbol{u}), (\boldsymbol{c^2})^\top \boldsymbol{x}(\boldsymbol{u}), \ldots, (\boldsymbol{c^k})^\top \boldsymbol{x}(\boldsymbol{u})) : \boldsymbol{u} \in U\}$. Constraint (6.1c) ensures that

$\boldsymbol{x}(\boldsymbol{u})$ is feasible. So, the resulting inner approximation indeed lies in $PS + \mathbb{R}_+^k$. The objective (6.1a) minimizes the volume under the inner approximation if (6.1b) is tight for all $\boldsymbol{u}$. Note that a constant function $\boldsymbol{x}(\boldsymbol{u})$ may be feasible for this optimization problem. The reason why this problem returns an interesting inner approximation and not just a constant function x is because from the objective it follows that a smaller $(\boldsymbol{c^k})^\top \boldsymbol{x}(\boldsymbol{u})$ is better, and a smaller $(\boldsymbol{c^k})^\top \boldsymbol{x}(\boldsymbol{u})$ can only be obtained by increasing $(\boldsymbol{c^i})^\top \boldsymbol{x}(\boldsymbol{u})$ for $i = 1, 2, \ldots, k-1$. This increase as a function of $\boldsymbol{u}$ is constrained by (6.1b).

A functional description of the inner approximation may be more tractable for a specific purpose. Therefore, we define:

$$IPS = \{(\boldsymbol{u}, (\boldsymbol{c^k})^\top \boldsymbol{x}(\boldsymbol{u})) : \boldsymbol{u} \in U\}.$$

If constraint (6.1b) is tight for all $\boldsymbol{u}$, IPS is the same as the inner approximation given before. Otherwise, the given inner approximation dominates IPS.

The question arises for which regions U and functions x this formulation is tractable. When U is polyhedral, for instance a box ($\{\boldsymbol{u} \in \mathbb{R}^{k-1} : ||\boldsymbol{u}||_\infty \leq 1\}$), and x is linear in $\boldsymbol{u}$, this problem can be formulated as an LP. The optimization problem minimizes the volume enclosed between the linear approximation, the Pareto curve and the boundary of U. In case of two objectives ($k = 2$), it will find the line connecting the point on the Pareto curve where $(\boldsymbol{c^1})^\top \boldsymbol{x} = -1$ with the point on the Pareto curve where $(\boldsymbol{c^1})^\top \boldsymbol{x} = 1$. Hence, it finds two Pareto optimal points. This extends to larger k, where it finds a plane going through $k-1$ Pareto optimal points. In case U is a box, the instance size also grows linearly in k. This result of linear growth is new, because determining a linear inner approximation over a box would require determining the Pareto optimal points at the exponentially growing number of extreme points of U.

The inner approximation becomes more interesting when x is nonlinear in $\boldsymbol{u}$. When $k = 2$, a tractable choice is given by polynomials: for polynomials of arbitrarily large degree, the problem can be formulated as an SDP (Ben-Tal et al., 2009a, Lemma 14.3.4), for which polynomial time solvers are available. When $k > 2$, the problem is tractable when U is ellipsoidal and x is quadratic in $\boldsymbol{u}$ (Ben-Tal et al., 2009a, Lemma 14.3.7). The resulting problem is an SDP with $m + k - 1$ variables matrices of size $k + 1$ and $m + k - 1$ constraints.

For $k > 2$, the problem can also be reformulated as an SDP when U is a semi-algebraic set ($U = \{\boldsymbol{u} : p_i(\boldsymbol{u}) \leq 0 \quad (i \in I)\}$, where p_i are polynomials of arbitrary degree), and $\boldsymbol{x}$ is a polynomial in $\boldsymbol{u}$, but the reformulation is not always equivalent. This means that the resulting optimal solution for the SDP reformulation may not be optimal for (6.1a)-(6.1c). However, the solution to the reformulation is always an inner approximation, and the numerical results in Section 6.4.2 are promising.

Table 6.1 – Tractability of the inner approximation.

k	U	$\boldsymbol{x(u)}$	Tractability	Exact
≥ 2	box	linear	LP	$\checkmark$
≥ 2	polyhedral	linear	LP	$\checkmark$
≥ 2	ball	linear	CQP	$\checkmark$
2	interval	polynomial	SDP	$\checkmark$
> 2	ellipsoidal	quadratic	SDP	$\checkmark$
≥ 2	semialgebraic	polynomial	SDP	$-$

The reformulation is based on polynomials that are sums of squares (SOS) of other polynomials. An example of an SOS polynomial is $5x^2 + 2x + 65$, because it can be decomposed as $(2x+4)^2 + (x-7)^2$. Testing whether a polynomial is SOS is equivalent to solving an SDP (Laurent, 2009, Lemma 3.8) with a matrix of size $\binom{k-1+d}{d}$, where $k-1$ is the number of variables and d is the degree of the polynomial. An SOS polynomial is obviously nonnegative. Let us focus on constraint (6.1b):

$$u_1 - (\boldsymbol{c^1})^\top \boldsymbol{x(u)} \geq 0 \qquad \forall \boldsymbol{u} : p_i(\boldsymbol{u}) \leq 0 \quad (i \in I).$$

By applying Putinar's Positivstellensatz (Laurent, 2009, Thm. 3.20), which has been done before in ARO Bertsimas et al. (2012), we can obtain a sufficient condition under which this constraint holds:

$$u_1 - (\boldsymbol{c^1})^\top \boldsymbol{x(u)} = \sigma_0(\boldsymbol{u}) + \sum_{i\in I} p_i(\boldsymbol{u})\sigma_i(\boldsymbol{u}), \tag{6.2}$$

where σ_0 and σ_i ($i \in I$) are SOS. Solving a problem with constraint (6.2) instead of (6.1b) is conservative for two reasons. First, (6.2) may not be a necessary condition when there are no σ_0, σ_i that are SOS for which the set $\{\boldsymbol{u} : \sigma_0(\boldsymbol{u}) + \sum_{i\in I} p_i(\boldsymbol{u})\sigma_i(\boldsymbol{u}) \geq 0\}$ is compact, or when (6.1b) is a not a strict inequality. Compactness can easily be guaranteed by including a constraint $\sum_{i=1}^{k-1} u_i^2 - R \leq 0$ to the description of U, which can be done without changing U because U is bounded (Laurent, 2009, p. 186). However, (6.1b) will in general not be a strict inequality. Second, solving a problem with constraint (6.2) as an SDP requires bounding the degree of σ_0 and σ_i.

We let the degree of x determine the complexity of the problem unless g is of higher degree, so we take the degree of σ_0 equal to $\max\{deg(x), \max_i\{deg(g_i)\}\}$, and the degree of σ_i equal to $\max\{0, deg(\sigma_0) - deg(g_i)\}$.

An overview of all tractable cases is given in Table 6.1.

For many uncertainty regions it may be difficult to reformulate the integral in the objective function as a simple linear function in the optimization variables. In that case the objective can be replaced with the average value of $(\boldsymbol{c}^k)^\top \boldsymbol{x}(\boldsymbol{u})$ at well distributed sampling points $\boldsymbol{u}$ in U. For efficient sampling from a polytope, see e.g. Kannan and Narayanan (2009).

The user has to specify the domain of interest U. If the specified region is too large, two things may occur. First, U may contain a vector with objective values that are too optimistic in the sense than they can not be met, in which case constraint (6.1b) is infeasible. Second, U may contain objective values that are not weakly dominated by $\boldsymbol{f}(\boldsymbol{x})$ for any feasible $\boldsymbol{x}$. In that case constraint (6.1b) will not be tight and also the objective is not fully related to the area of interest.

When the number of objectives is three or more, a weak parameterization of x can be another cause of infeasibility. If feasible solutions exist for all $\boldsymbol{u}$ in U, it is possible that these solutions can not be attained with the parameterization. An easy example is the case where U contains two different Pareto optimal solutions (projected on the first $k-1$ coordinates) while x is a constant function. When the optimization problem is infeasible, we do not see a possibility to detect whether the parameterization is too weak or U contains infeasible points.

It is known that PS is convex, and that it is nonincreasing. It may be the case that these properties do not hold for the inner approximation, which is problematic when the inner approximation is used in an algorithm that assumes these properties to hold. In case x is linear in $\boldsymbol{u}$, these conditions are automatically satisfied. In case $k = 2$, u is one-dimensional and convexity and nonincreasingness of $(\boldsymbol{c}^2)^\top \boldsymbol{x}(\boldsymbol{u})$ can be enforced by constraining the first and second derivative w.r.t. u. The first and second derivative of a polynomial is again a polynomial, so constraining these to be negative and positive, respectively, for polynomial x does not increase the complexity class of the problem. In case $k > 2$, U is a polynomial and x is quadratic in $\boldsymbol{u}$, say $x(\boldsymbol{u})_i = \alpha_i^0 + (\boldsymbol{\alpha_i^1})^\top \boldsymbol{u} + \boldsymbol{u}^\top \Gamma_i \boldsymbol{u}$ where $\alpha_i^0 \in \mathbb{R}, \boldsymbol{\alpha_i^1} \in \mathbb{R}^{k-1}$ and $\Gamma_i \in \mathbb{R}^{(k-1)\times(k-1)}$ are decision variables. Nonincreasingness can easily be enforced by adding the constraints $\boldsymbol{\alpha_i^1} + 2\Gamma_i \boldsymbol{u} \leq 0$ for all $\boldsymbol{u}$ in U, which is a set of $n(k-1)$ linear constraints with ellipsoidal uncertainty, each of which can be reformulated as a conic quadratic constraint. For convexity it is required that Γ_i is positive semidefinite, which is an SDP constraint.

6.4   Numerical examples

6.4.1   Two objectives

We construct a semi-random 150×170 matrix A, a 150-vector b, and two 170-vectors $\boldsymbol{c^1}$ and $\boldsymbol{c^2}$, such that the Pareto cuve is interesting on the interval $[0, 25]$. We compute

(a) degree 1

(b) degree 4

(c) degree 8

(d) degree 16
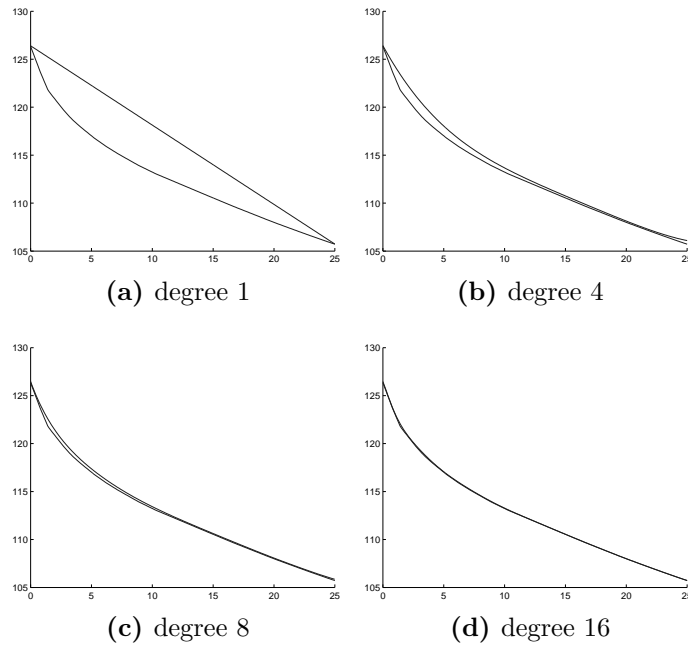
Figure 6.1 – Numerical example with two objectives indicating the quality of the inner approximation with a polynomial decision rule. The lowest curve represents the set PS which is usually not known.

a polynomial inner approximations of degree up to 16.

We solve linear programs with Matlab linprog. We enter linear constraints with LMI uncertainty into YALMIP Löfberg (2012). YALMIP reformulates this problem as an SDP. In Appendix 6.B we show how to do this reformulation by hand in case of a polynomial of degree 3. We let YALMIP export the resulting problem, then we reformulate free variables as the difference of two nonnegative variables using CSDP's convertf, and solve the problems with SDPA Yamashita et al. (2010) (SDPA-DD Nakata (2010) for the problem with a polynomial of degree 16). Figure 6.1 shows the resulting solutions. The solution time ranges from 21 seconds for the polynomial approximation of degree 4 to 4 minutes for degree 8 (with SDPA), and 45 minutes for degree 16 (with SDPA-DD).

6.4.2   Three objectives

We semi-randomly construct vectors $\boldsymbol{c^1} \in [0,1]^{10}$, $\boldsymbol{c^2} \in [0,1]^{10}$ and $\boldsymbol{c^3} \in [-1,0]^{10}$, and take $\mathbb{R}_+^{10}$ as the feasible region. Recall from Table 6.1 that we have an exact result for a quadratic inner approximation over an ellipsoidal set, and a conservative result for polynomial inner approximations of arbitrary degree over semialgebraic sets. Again

<div align="center">(a) degree 2   (b) degree 4</div>

Figure 6.2 – Numerical example with three objectives indicating the quality of the inner approximation with a polynomial decision rule. The lowest set represents the set PS which is usually not known.

we use YALMIP and SDPA to formulate the problem and solve the resulting SDPs. We use YALMIPs SOS module for constraining an expression to be SOS.

We take $\{\boldsymbol{u} : ||\boldsymbol{u} - \boldsymbol{5}||_2 \le 5\}$ as the area of interest for $(\boldsymbol{c^1})^\top \boldsymbol{x}$ and $(\boldsymbol{c^2})^\top \boldsymbol{x}$, and approximate the Pareto set with a polynomial of degree 2 and with a polynomial of degree 4. For degree 2, we solve the exact robust counterpart, while for degree 4 we solve the SOS approximation. We also solve the SOS approximation for degree 2, and notice that the inner approximation is the same as with the exact robust counterpart. Figure 6.2 shows that the polynomial of degree 4 gives an approximation that is closer to PS than the polynomial of degree 2. The solution time is around 1.6 seconds for all three approximations.

6.A   Outer approximation

In this appendix we show how ARO can be used to construct an outer approximation. While theoretically possible, it is practically not tractable to determine polynomial approximations of degree 2 or higher due to computational issues. For linear approximations, the results are trivial. We still mention the result because it also uses ARO to approximate PS.

Let $U = [a^1, b^1] \times [a^2, b^2] \times \cdots \times [a^{k-1}, b^{k-1}]$ with $[a^i, b^i]$ be the domain of interest for $(c^i)^\top \boldsymbol{x}$ $(i = 1, 2, \ldots, k-1)$. We construct the set OPS by creating a function function $\ell : \mathbb{R}^{k-1} \to \mathbb{R}$ for which $(u^1, u^2, \ldots, u^{k-1}, \ell(u^1, u^2, \ldots, u^{k-1}))$ is in $PS + \mathbb{R}^k_-$,

and optimizing over this function:

$$\max_{\ell} \quad \int_U \ell(u_1, u_2, \ldots, u_{k-1}) du \tag{6.3a}$$

$$\text{s.t. } \ell((\boldsymbol{c^1})^\top \boldsymbol{x}, (\boldsymbol{c^2})^\top \boldsymbol{x}, \ldots, (c^{k-1})^\top \boldsymbol{x}) \leq (\boldsymbol{c^k})^\top \boldsymbol{x}$$

$$\forall \boldsymbol{x} : A\boldsymbol{x} \leq \boldsymbol{b}, a^i \leq (\boldsymbol{c^i})^\top \boldsymbol{x} \leq b^i. \tag{6.3b}$$

Given a solution to this optimization problem, the outer approximation is $\{(u_1, u_2, \ldots, u_{k-1}, \ell(u_1, u_2, \ldots, u_{k-1})) : \boldsymbol{u} \in U\}$. The objective (6.3a) maximizes the volume under this approximation. Constraint (6.3b) ensures that the outer approximation as a function of $(\boldsymbol{c^1})^\top \boldsymbol{x}$ lies under $(\boldsymbol{c^2})^\top \boldsymbol{x}$ for every $\boldsymbol{x}$ in the domain of interest.

An optimal outer approximation is tangent to the Pareto curve at (at least) one point. This becomes clear from (6.3b): this constraint holds with equality for at least one x because otherwise we can add a constant to ℓ without losing feasibility, which contradicts optimality. Previous results force the decision maker to specify either this point of tangency or the derivative at this point a priori. Our formulation determines the point of tangency in such a way that the volume enclosed between this linear outer approximation and the Pareto curve over the set U, but in the linear case this turns out to give a trivial result.

When ℓ is linear, the problem (6.3b) can be reformulated as an LP using ARO. For the case $k = 2$ (two objectives) it can be shown that the optimal linear ℓ is a line tangent to PS at $\frac{a^1 + b^1}{2}$, i.e. halfway the interval of interest. We conjecture that in higher dimensions the point of tangency is the barycenter of U. This would imply that the formulation for the outer approximation is not interesting because it is already known how to obtain an outer approximation that is tangent at a given point.

For nonlinear ℓ the SOS framework used in Section 6.3 can be used to reformulate the problem as an SDP when ℓ is a polynomial of arbitrary degree d. This is a polynomial in the vector $\boldsymbol{x}$, so the number of terms is $\binom{n+d}{d}$, which is also the order of the matrix in the SDP. Even for a quadratic function, the size of the SDP is often too large to solve.

The user has to specify the domain of interest $[u^1, u^2]$. Specifying the wrong domain does not lead to infeasibility. However, if the interval is too large, i.e. $(\boldsymbol{c^1})^\top \boldsymbol{x}$ does not range through the full interval, part of the outer approximation is meaningless because PS is inexistent for some $\boldsymbol{u}$.

6.B   Derivation of the SDP formulation for a polynomial inner approximation with two objectives

We give a derivation of the SDP formulation of (6.1a)-(6.1c) in case $\boldsymbol{x(u)} = \boldsymbol{\alpha_0} + \boldsymbol{\alpha_1}u + \boldsymbol{\alpha_2}u^2 + \boldsymbol{\alpha_3}u^3$ (where $\boldsymbol{\alpha_i}$ in $\mathbb{R}^{170}$, i=1,2,3) for the numerical example of Section 6.4.1. Suppose the area of interest for $(\boldsymbol{c^1})^{\top}\boldsymbol{x}$ is $[0, 25]$, then u runs from 0 to 25. Because in the result by (Ben-Tal et al., 2009a, Lemma 14.3.4) u runs from -1 to 1, we use the following linear transformation $\boldsymbol{x} \rightarrow D\boldsymbol{x} + \boldsymbol{d}$ to transform $\mathcal{Z} = \{(u, u^2, u^3) : -1 \leq u \leq 1\}$ into $\{(u, u^2, u^3) : 0 \leq u \leq 25\}$:

$$D = \begin{pmatrix} 12.5 & 0 & 0 \\ 312.5 & 156.25 & 0 \\ 5859.375 & 5859.375 & 1953.125 \end{pmatrix} \qquad \boldsymbol{d} = \begin{pmatrix} 12.5 \\ 156.25 \\ 1953.125 \end{pmatrix}$$

The problem can now be written as follows:

$$\min \quad (\boldsymbol{c^2})^{\top} \left(25\boldsymbol{\alpha_0} + \frac{25^2}{2}\boldsymbol{\alpha_1} + \frac{25^3}{3}\boldsymbol{\alpha_2} + \frac{25^4}{4}\boldsymbol{\alpha_3} \right)$$

$$(\boldsymbol{c^1})^{\top} \left(\boldsymbol{\alpha_0} + [\boldsymbol{\alpha_1} \quad \boldsymbol{\alpha_2} \quad \boldsymbol{\alpha_3}](D\boldsymbol{\zeta} + \boldsymbol{d}) \right) \leq (D\boldsymbol{\zeta} + \boldsymbol{d})_1$$

$$\forall \boldsymbol{\zeta} \in \mathcal{Z} \tag{6.4a}$$

$$A \left(\boldsymbol{\alpha_0} + [\boldsymbol{\alpha_1} \quad \boldsymbol{\alpha_2} \quad \boldsymbol{\alpha_3}](D\boldsymbol{\zeta} + \boldsymbol{d}) \right) \leq \boldsymbol{b}$$

$$\forall \boldsymbol{\zeta} \in \mathcal{Z}, \tag{6.4b}$$

where:

$$\mathcal{Z} = \left\{ \boldsymbol{\zeta} \in \mathbb{R}^3 : \begin{pmatrix} 1 \\ \boldsymbol{\zeta} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 3 & 0 & 4 & 0 \\ 0 & 4 & 0 & 8 \\ 3 & 0 & 4 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}^{\top} \begin{pmatrix} \lambda_0 \\ \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \\ \lambda_5 \\ \lambda_6 \end{pmatrix}, \begin{pmatrix} \lambda_0 & \lambda_1 & \lambda_2 & \lambda_3 \\ \lambda_1 & \lambda_2 & \lambda_3 & \lambda_4 \\ \lambda_2 & \lambda_3 & \lambda_4 & \lambda_5 \\ \lambda_3 & \lambda_4 & \lambda_5 & \lambda_6 \end{pmatrix} \succeq 0 \right\}.$$

Constraints (6.4a) and (6.4b) are a total of 151 semi-infinite constraints that have to hold for an infinite number of $\boldsymbol{\zeta}$. Let $\boldsymbol{A_j}$ denote the j^{th} row of A. In order to allow for shorter notation, we define the linear function $\ell : \mathbb{R}^{170} \times \mathbb{R}^{170} \times \mathbb{R}^{170} \rightarrow \mathbb{R}^3$ vector $\boldsymbol{\ell(\alpha)} := A_j[\boldsymbol{\alpha_1} \quad \boldsymbol{\alpha_2} \quad \boldsymbol{\alpha_3}]D$. The j^{th} constraint of (6.4b) can be rearranged to:

$$\boldsymbol{A_j}\boldsymbol{\alpha_0} - \boldsymbol{b} + [\boldsymbol{\alpha_1} \quad \boldsymbol{\alpha_2} \quad \boldsymbol{\alpha_3}]\boldsymbol{d} + \boldsymbol{\ell(\alpha)}\boldsymbol{\zeta} \leq 0 \qquad\qquad \forall \boldsymbol{\zeta} \in \mathcal{Z},$$

which is equivalent to:

$$\boldsymbol{A_j}\boldsymbol{\alpha_0} - \boldsymbol{b} + \begin{bmatrix} \boldsymbol{\alpha_1} & \boldsymbol{\alpha_2} & \boldsymbol{\alpha_3} \end{bmatrix}\boldsymbol{d} + \max_{\boldsymbol{\zeta} \in \mathcal{Z}} \{\boldsymbol{\ell}(\boldsymbol{\alpha})\boldsymbol{\zeta}\} \leq 0. \tag{6.5}$$

The maximization problem is an SDP. Replacing this problem with its SDP dual and omitting the min operator is a well-known method to transform a semi-infinite constraint into a single constraint. We show how to do this. In the following SDP, we take the 4×4 matrix with λ_i's in the description of $\mathcal{Z}$ as our variable X. Let $\ell_i(\alpha)$ denote the i^{th} component of $\boldsymbol{\ell}(\boldsymbol{\alpha})$.

The optimization problem in (6.5) is:

$$\begin{aligned} \max \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A_i, X \rangle = b_i \quad (i = 1, 2, 3, 4) \\ & X \succeq 0, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ denotes the trace inner product, $b_1 = 1$, $b_2 = b_3 = b_4 = 0$, and:

$$C = \begin{pmatrix} \ell_3(\alpha) & \ell_2(\alpha) & g(\alpha) & \ell_2(\alpha) \\ \ell_2(\alpha) & g(\alpha) & \ell_2(\alpha) & g(\alpha) \\ g(\alpha) & \ell_2(\alpha) & g(\alpha) & \ell_2(\alpha) \\ \ell_2(\alpha) & g(\alpha) & \ell_2(\alpha) & \ell_3(\alpha) \end{pmatrix},$$

with $g(\alpha) = \dfrac{4}{3}\ell_1(\alpha) + \ell_3(\alpha)$, and:

$$A_1 = \begin{pmatrix} 0 & 0 & 0 & 2 \\ 0 & 0 & 2 & 0 \\ 0 & 2 & 0 & 0 \\ 2 & 0 & 0 & 0 \end{pmatrix}, \qquad A_2 = \begin{pmatrix} 0 & 0 & -\frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \\ -\frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$A_3 = \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \qquad A_4 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{2} \\ 0 & 0 & 1 & 0 \\ 0 & -\frac{1}{2} & 0 & 0 \end{pmatrix}.$$

By formulating the dual and putting this into constraint (6.5), we get the following robust counterpart:

$$\boldsymbol{A_j}\boldsymbol{\alpha_0} - \boldsymbol{b} + \begin{bmatrix} \boldsymbol{\alpha_1} & \boldsymbol{\alpha_2} & \boldsymbol{\alpha_3} \end{bmatrix}\boldsymbol{d} + y_1 \leq 0$$

$$\sum_{i=1}^{4} y_i A_i - C \succeq 0.$$

Constraint (6.4a) and the other constraints (6.4b) can be transformed in a smilar way.

Acknowledgments

CHAPTER 7

Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR-brachytherapy

Abstract Current inverse treatment planning methods that optimize both catheter positions and dwell times in prostate HDR brachytherapy use surrogate linear or quadratic objective functions that have no direct interpretation in terms of dose-volume histogram (DVH) criteria, do not result in an optimum or have long solution times.

We decrease the solution time of existing linear and quadratic dose-based programming models (LP and QP, respectively) to allow optimizing over potential catheter positions using mixed integer programming. An additional average speed-up of 75% can be obtained by stopping the solver at an early stage, without deterioration of the plan quality. For a fixed catheter configuration, the dwell time optimization model LP solves to optimality in less than 15 seconds, which confirms earlier results. We propose an iterative procedure for QP that allows to prescribe the target dose as an interval, while retaining independence between the solution time and the number of dose calculation points. This iterative procedure is comparable in speed to the LP model, and produces better plans than the non-iterative QP.

We formulate a new dose-volume based model that maximizes $V_{100\%}$ while satisfying pre-set DVH-criteria. This model optimizes both catheter positions and dwell times within a few minutes depending on prostate volume and number of catheters, optimizes dwell times within 35 seconds, and gives better DVH statistics than dose-based models. The solutions suggest that the correlation between objective value and clinical plan quality is weak in existing dose-based models.

7.1   Introduction

7.1.1   HDR brachytherapy optimization

Interstitial high-dose-rate (HDR) brachytherapy is a form of internal radiation therapy where a high activity iridium-192 stepping source is temporarily placed into the tumour volume or its proximity through the use of implanted catheters. This type of radiotherapy has shown to be an excellent option for the definitive treatment of localized prostate cancer in any risk category (Yamada et al., 2012). Clinical outcome data shows high tumour control and low toxicity rates because of the precision and control with which highly conformal optimized HDR treatment can be delivered.

Treatment planning is one of the steps in the process to deliver the prescribed dose, and entails two design problems. Firstly, the number and spatial configuration of the catheters to be implanted has to be determined. Secondly, the spatio-temporal source stepping pattern within the implanted catheters needs to be calculated. The number and configuration of catheters depend on the prostate shape, volume, and regional anatomy. Often a template is used for the transperineal implantation of needles. As each catheter offers a range of potential stopping points (dwell positions) where the source can stay for a predefined time (dwell time), the design problem has many degrees of freedom. Hence, the problem formulation to design an HDR brachytherapy plan comes down to: 1) determine where to insert the catheters in the template, and 2) determine where and for how long to stop the source inside the catheters to achieve adequate coverage of the planning target volume (PTV) while sufficiently limiting dose to surrounding organs at risk (OAR).

Computerized techniques for anatomy-based inverse treatment planning of HDR bra-chytherapy enable solving these problems by calculating catheter configurations and source dwell time distributions, based on mathematical optimization models. These models require the above problem formulation to be translated into a mathematical framework that describes how the dose distribution depends on the decision variables. Furthermore, it is mandatory to establish several levels of abstraction to assess the dosimetric quality of the resulting treatment plan. In this chapter, we discriminate between three levels of abstraction. At the highest level, a dose penalty function is used to assign a penalty measure to a treatment plan. At the intermediate level, dose-volume-histogram (DVH) statistics are used to evaluate the dose and dose-volume characteristics of a given dose distribution. At the lowest level, the opinion of a human expert forms a subjective judgement of the three dimensional dose distribution. Ideally, lower values at the highest level correspond to better plans when evaluated at the lowest level. We briefly discuss these three levels, as they are relevant for the remainder of this chapter.

Dose penalty functions. As it is impossible to calculate dose deposited to every single cell, the most relevant tissue volumes (i.e., prostate, urethra, rectum) are discretized into finite sets of dose calculation points that each represents an adequately small subvolume where the dose is considered to be uniform. At each point i, the delivered dose, d_i, is compared with a prescribed lower bound L_i and upper bound U_i. If the delivered dose is not between these bounds, the violation is penalized. We use the linear and quadratic penalty function to develop our linear and quadratic dose-based optimization models, respectively (Figure 7.1). For a linear penalty function, costs of α_i or β_i are incurred per unit dose (Gy) violation of the lower or upper bound, respectively, whereas for a quadratic penalty function a dose violation is penalized to the second degree. The total penalty per tissue structure is the summed penalty over all calculation points within the structure.

The linear penalty function has been used in two well-known algorithms for anatomy-based inverse treatment planning: Inverse Planning by Simulated Annealing (IPSA) (Lessard and Pouliot, 2001; Alterovitz et al., 2006) and Hybrid Inverse treatment Planning and Optimization (HIPO) (Karabis et al., 2005). The quadratic penalty function has also been discussed in the literature (Milickovic et al., 2002; Lahanas et al., 2003a; Lahanas and Baltas, 2003).

Figure 7.1 – The penalty in calculation point i for not satisfying the lower or upper bound is either linear (a) or quadratic (b) in the violation.

DVH statistics. A typical example of clinically relevant dose-volume criteria for an HDR brachytherapy prostate plan with 2 fractions of 8.5 Gy following external beam radiation treatment with 13 fractions of 2.75 Gy is listed in Table 7.1. The criterion $D_{90\%} \geq 90\%$ requires the hottest 90% of the PTV to receive at least 90% of the prescribed dose, and $V_{150\%} \leq 55\%$ requires the relative volume that is exposed to more than 150% of the prescribed dose to be less than 55%. For the urethra, $D_{0.1cc} \leq 10\,\mathrm{Gy}$ means that the minimum dose delivered to the hottest 0.1 cc does not exceed 10 Gy.

Table 7.1 – Local protocol based on (Hoskin et al., 2007) of DVH criteria for a prescribed dose of 8.5 Gy per fraction.

PTV	Rectum	Urethra
$D_{90\%} \geq 90\%$	$D_{10\%} \leq 7.2\,\mathrm{Gy}$	$D_{10\%} \leq 10\,\mathrm{Gy}$
$V_{100\%} \geq 90\%$	$D_{2cc} \leq 6.7\,\mathrm{Gy}$	$D_{0.1cc} \leq 10\,\mathrm{Gy}$
$V_{150\%} \leq 55\%$	$V_{94\%} = 0\,\mathrm{cc}$	$V_{125\%} = 0\,\mathrm{cc}$
$V_{200\%} \leq 20\%$		

Expert opinion. Clinically established DVH criteria are often used as quality indicators of a treatment plan. Computer optimized treatment plans often do not satisfy pre-set DVH criteria, and hence require *a posteriori* adjustment of the dwell time distributions. This is often accomplished using so-called 'graphical optimization', see e.g. (Morton et al., 2008), or manual adaptation of individual dwell times. The quality of an optimization method can therefore also be expressed by the time spent to post-process the plan. It is obvious that the perceived plan quality strongly depends on the level of experience of the treatment planner.

Mixed integer programming. There is a matured field of research that deals with optimization problems where (some of the) variables have to be integer, which are inherently difficult to solve to optimality. When applied to problems in which variables are restricted to 0 or 1, the method starts by solving a relaxed problem where the binary variables are allowed to take any value between 0 and 1. If in the optimal solution there is at least one binary variable which is not integer, the method procedes by either (1) adding a constraint that does not exclude the optimal binary solution, but excludes the current optimal solution, or (2) splitting the problem into two subproblems, one with a selected binary variable fixed to 0 and one with the same binary variable fixed to 1. These steps may be combined, and are repeated until the optimal solution is found. This way, the method can report an upper and a lower bound on the objective value, where one corresponds to the current best solution and the other to the solution of a relaxed problem. We refer the interested reader to (Nemhauser and Wolsey, 1999).

7.1.2 Our contributions

Currently, only HIPO can solve the problem of catheter placement for prostate HDR brachytherapy. However, because HIPO is partially based on heuristical optimization, mathematical optimality of the solution cannot be guaranteed. This means that there may be solutions with lower objective values, that better satisfy the pre-

scribed dose. We apply modifications to existing dose-based optimization models, so that a state-of-the-art mixed-integer linear programming (MILP) or mixed-integer quadratic programming (MIQP) solver can solve them to optimality.

Conventional quadratic dose-based optimization models have the advantage that the complexity of the optimization problem is independent of the number of dose calculation points if the target dose is specified as a single value rather than the interval $[L_i, U_i]$ (Lahanas et al., 2003a; Lahanas and Baltas, 2003). This implies that underdosage and overdosage are penalized equally, even though the latter is allowed up to U_i. It is our second contribution to present an iterative procedure for quadratic penalty functions that retains the advantage of the conventional quadratic dose-based models.

As a third contribution, we present a new model that maximizes PTV coverage while constraining DVH parameters on the OAR(s). This model has a more direct clinical interpretation than linear penalty functions, solves in clinically acceptable time and is expected to produce better treatment plans than dose-based models. Using this model, we show that dose-based objective functions are bad surrogates for dose-volume based plan evaluation criteria.

The structure of this chapter is as follows. We start by introducing the mathematical notations and optimization models in Section 7.2. Our methods are numerically evaluated on three clinical sample cases in Section 7.3. The relation to other work is discussed in Section 7.4.

7.2 Methods

7.2.1 Mathematical notation

We use the sets, parameters and variables listed in Tables 7.2-7.4. The dose delivered to calculation point i is a linear function of dwell time and dose rate:

$$(\dot{\mathbf{D}}\mathbf{t})_i = \sum_{j \in J} \dot{d}_{ij} t_j. \tag{7.1}$$

Table 7.2 – Sets used for inverse treatment planning.

set	description
S	Tissue structures {PTV, R(ectum), (U)rethra}
I_s	Calculation points in structure $s \in S$
I	All calculation points, $I_{PTV} \cup I_R \cup I_U$
K	Catheters
J_k	Dwell positions in catheter $k \in K$
J	All dwell positions, $\bigcup_{k \in K} J_k$
$\Gamma(j)$	Dwell positions adjacent to dwell position $j \in J_k$ within catheter k

Table 7.3 – Parameters used for inverse treatment planning.

parameter	unit	description
N	1	Upper bound on the number of catheters allowed
$t_{\max}$	s	An upper bound on the dwell time for a single dwell position
$\dot{d}_{ij}$	Gy	The dose rate delivered to calculation point $i \in I$ by a source at dwell position $j \in J$ per unit of time
$\dot{\mathbf{D}}$	Gy	The first order dose kernel matrix, i.e. the matrix with elements $\dot{d}_{ij}$
L_i	Gy	Prescribed lower bound on the dose for calculation point $i \in I$
U_i	Gy	Prescribed upper bound on the dose for calculation point $i \in I$
p_i	Gy	Prescribed dose for calculation point $i \in I$ (in case $L_i = U_i$)
$\mathbf{p}$	Gy	The prescribed dose vector with elements p_i
τ_s	1	Percentage of calculation points receiving a dose less than L_i in structure $s \in S$
α_i	Gy^{-1}	Penalty per Gy below the lower bound L_i for calculation point $i \in I$
β_i	Gy^{-1}	Penalty per Gy exceeding the upper bound U_i for calculation point $i \in I$
γ	%	The maximum allowable relative difference in dwell times between two adjacent dwell positions
w_i	1	Relative weight of calculation point i, $w_i = 1/|I_s|$ where s is the structure containing i
$\mathbf{W}$	1	Weight matrix, with w_i on the diagonal and 0 at other positions

<div align="center">**Table 7.4** – Variables used for inverse treatment planning.</div>

variable	unit	description
t_j	s	Dwell time at dwell position $j \in J$
$\mathbf{t}$	s	The dwell time vector with elements $t_j, j \in J$
v_i	1	Binary variable indicating whether calculation point i receives at most (or at least) its prescribed dose
b_k	1	Binary variable indicating whether catheter position $k \in K$ in the treatment template is used
x_i	1	Penalty for calculation point $i \in I$

7.2.2   Optimization models

7.2.2.1   Linear Dose-based (LD) model.

Using the notation introduced in Section 7.2.1, the linear dose-based objective function in Figure 7.1a can be written as:

$$\min \quad \sum_{i \in I} \max\{0, \alpha_i(L_i - \sum_{j \in J} \dot{d}_{ij} t_j), \beta_i(\sum_{j \in J} \dot{d}_{ij} t_j - U_i)\}. \tag{7.2}$$

In order to transform this objective function into a model with a linear objective function and linear constraints, we introduce a variable x_i replacing the argument of the max operator in (7.2). Additional constraints limit the number of catheters used and the relative difference in dwell time between adjacent positions. The latter implements the dwell time modulation restriction (DTMR), that is often used to prevent hot spots (Baltas et al., 2009). The full model then becomes:

$$
\begin{align}
(LD) \quad \min \quad & \sum_{i \in I} w_i x_i \tag{7.3a} \\
\text{s.t.} \quad & x_i \geq \alpha_i[L_i - \sum_{j \in J} \dot{d}_{ij} t_j] & \forall i \in I \tag{7.3b} \\
& x_i \geq \beta_i[\sum_{j \in J} \dot{d}_{ij} t_j - U_i] & \forall i \in I \tag{7.3c} \\
& t_j \leq b_k t_{\max} & \forall k \in K \quad \forall j \in J_k \tag{7.3d} \\
& t_{j_1} \leq (1 + 100\gamma)t_{j_2} & \forall j_1 \in J \quad \forall j_2 \in \Gamma(j_1) \tag{7.3e} \\
& \sum_{k \in K} b_k \leq N \tag{7.3f} \\
& b_k \in \{0, 1\} & \forall k \in K \tag{7.3g} \\
& x_i \geq 0 & \forall i \in I \tag{7.3h} \\
& t_j \geq 0 & \forall j \in J. \tag{7.3i}
\end{align}
$$

The penalty function for calculation point i is a convex piecewise linear function in the dose $\sum_{j \in J} \dot{d}_{ij} t_j$. Constraints (7.3b), (7.3c) and (7.3h) together with objective (7.3a) make x_i equal to the pointwise maximum. If catheter k is not used then $b_k = 0$.

Constraint (7.3d) sets the dwell times within that catheter to 0 seconds. Constraints (7.3e) and (7.3f) implement the DMTR and enforce no more than N catheters to be selected, respectively.

If the DTMR would be dropped, the model is equivalent to the one described by Karabis et al. (2009). If additionally the values of b_k are fixed, i.e., if all catheter positions are fixed, the model corresponds to the one by Alterovitz et al. (2006).

When we solve it as an MILP, the solution times are very high and clinically unacceptable. This is in line with the results from Karabis et al. (2009), where a similar model could not be solved in less than 5 hours when the number of catheter positions was more than 25–30.

We have improved the solution time by making two improvements. The first is to specify constraint (7.3d) as an indicator constraint, which is an option offered by our solver that helps treating this constraint more efficiently. Only if $b_k = 0$, the constraint $t_j = 0$ becomes visible to the solver. The second improvement is to make two adjacent catheter positions mutually exclusive by adding an exclusion restriction $b_{k_1} + b_{k_2} \leq 1$ for any two catheters k_1 and k_2 that are adjacent in the template. The rationale for this is that two adjacent catheters are likely to cause high-dose subregions to become connected and form undesirable hot spots.

7.2.2.2   Quadratic Dose-based (QD) model.

As an alternative to the (LD) model, we propose a convex quadratic model. If we use a quadratic objective function, the number of calculation points no longer plays a role, thus greatly reduces complexity. By using identity (7.1), $(\dot{\mathbf{D}}\mathbf{t})_i$ it is evident that $\sum_i (w_i(\mathbf{D}\mathbf{t})_i - p_i)^2$ measures the deviation from the prescribed dose p_i in calculation point i. This can also be written as the squared 2-norm $||\mathbf{W}(\dot{\mathbf{D}}\mathbf{t} - \mathbf{p})||_2^2$, which is convex in $\mathbf{t}$. Consider the following constrained least-squares approximation model:

$$
\begin{align*}
(QD) \quad \min \quad & ||\mathbf{W}(\dot{\mathbf{D}}\mathbf{t} - \mathbf{p})||_2^2 & & \text{(7.4a)} \\
\text{s.t.} \quad & t_j \leq b_k t_{\max} & \forall k \in K \quad \forall j \in J_k & \text{(7.4b)} \\
& t_{j_1} \leq (1 + 100\gamma)t_{j_2} & \forall j_1 \in J \quad \forall j_2 \in \Gamma(j_1) & \text{(7.4c)} \\
& \textstyle\sum_{k \in K} b_k \leq N & & \text{(7.4d)} \\
& b_k \in \{0, 1\} & \forall k \in K & \text{(7.4e)} \\
& t_j \geq 0. & & \text{(7.4f)}
\end{align*}
$$

The objective can be rewritten as: $(\mathbf{W}(\dot{\mathbf{D}}\mathbf{t}-\mathbf{p}))^\top(\mathbf{W}(\dot{\mathbf{D}}\mathbf{t}-\mathbf{p})) = \mathbf{t}^\top\dot{\mathbf{D}}^\top\mathbf{W}^\top\mathbf{W}\dot{\mathbf{D}}\mathbf{t} - 2\mathbf{p}^\top\mathbf{W}^\top\mathbf{W}\dot{\mathbf{D}}\mathbf{t}+\mathbf{p}^\top\mathbf{W}^\top\mathbf{W}\mathbf{p}$. Instead of the full $|I| \times |J|$ matrix $\dot{\mathbf{D}}$ it suffices to specify the $|J| \times |J|$ matrix $\dot{\mathbf{D}}^\top\mathbf{W}^\top\mathbf{W}\dot{\mathbf{D}}$ and the $|J| \times |1|$ vector $\dot{\mathbf{D}}^\top\mathbf{W}^\top\mathbf{W}\mathbf{p}$, whose sizes do not increase with the number of calculation points. The latter has been observed before by Lahanas et al. (2003a), but has not been used for formulating a convex

quadratic programming model. Instead, these authors formulate nonconvex models, for which optimality cannot be guaranteed. However, the (QD) model can be solved to optimality. The solution time is greatly reduced by adding the exclusion restriction.

In this model, the prescribed dose p_i for calculation point i has to be specified. It is not immediately clear which value should be taken. For points outside the PTV, a value of 0 is reasonable. For points inside the PTV it is difficult not to penalize very reasonable values. Finding a good value for p_i only gives a target value that is good for the average calculation point. All calculation points will still contribute some amount to the objective function even though they receive a dose between L_i and U_i. We can alleviate this disadvantage by solving the problem iteratively. The algorithm starts by initializing each p_i at $(L_i + U_i)/2$. For each iteration, first the problem is solved, then p_i gets adjusted to a value in $[L_i, U_i]$ closest to the dose received in the current optimal solution. The algorithm stops when the improvement in objective value is sufficiently small. A very precise description is given in Algorithm 4 in 7.B. This procedure in general does not necessarily converge to the global optimum that could have been obtained by minimizing simultaneously over $\mathbf{t}$, $\dot{d}_{ij}$ and b_k, which is proven with a small example in 7.C.

7.2.2.3 Linear Dose-Volume based (LDV) model.

We propose a new model that maximizes the fraction of the PTV receiving the prescribed dose, while constraining DVH parameters for OARs. For the rectum and urethra we enforce $D_{10\%} \leq 7.2\,\text{Gy}$ and $D_{10\%} \leq 10\,\text{Gy}$, respectively. In accordance with Table 7.1, we do not allow a dose higher than $8\,\text{Gy}$ in the rectum, and $10.6\,\text{Gy}$ in the urethra. This formulation has the advantage that a feasible solution exists (e.g. take all dwell times equal to 0), and that the solution shows the best target coverage that satisfies clinically derived DVH constraints. We formulate the model

as:

$$(LDV) \quad \max \quad \frac{1}{|I_{PTV}|} \sum_{i \in I_{PTV}} v_i \tag{7.5a}$$

$$\text{s.t.} \quad \sum_{j \in J} \dot{d}_{ij} t_j \geq L_i v_i \quad \forall i \in I_{PTV} \tag{7.5b}$$

$$\sum_{j \in J} \dot{d}_{ij} t_j \leq L_i + (U_i - L_i)(1 - v_i) \quad \forall i \in I_R \cup I_U \tag{7.5c}$$

$$\sum_{i \in I_s} v_i \geq \tau_s |I_s| \quad \forall s \in \{R, U\} \tag{7.5d}$$

$$t_j \leq b_k t_{\max} \quad \forall k \in K \quad \forall j \in J_k \tag{7.5e}$$

$$t_{j_1} \leq (1 + 100\gamma) t_{j_2} \quad \forall j_1 \in J \quad \forall j_2 \in \Gamma(j_1) \tag{7.5f}$$

$$\sum_{k \in K} b_k \leq N \tag{7.5g}$$

$$b_k \in \{0, 1\} \quad \forall k \in K \tag{7.5h}$$

$$v_i \in \{0, 1\} \quad \forall i \in I \tag{7.5i}$$

$$t_j \geq 0 \quad \forall j \in J. \tag{7.5j}$$

Here L_i and U_i have a slightly different interpretation than in the dose-based models. We select $L_i = 8.5 \, \text{Gy}$ for the PTV, $(L_i, U_i) = (7.2, 8) \, \text{Gy}$ for the rectum, $(L_i, U_i) = (10, 10.6) \, \text{Gy}$ for the urethra and $\tau_R = \tau_U = 0.9$.

Constraint (7.5b) allows v_i for $i \in I_{PTV}$ to be 1 only if the dose exceeds L_i. Hence objective function (7.5a) maximizes the number of points inside the PTV that receive the prescribed dose. Constraint (7.5c) allows v_i for $i \in I_R$ to be 1 only if the dose is less than $7.2 \, \text{Gy}$, and never allows a dose higher than $8 \, \text{Gy}$. Constraint (7.5d) then enforces 90% of the v_i to be 1. Similarly, the same constraints enforce that 90% of the urethra receives a dose less than $10 \, \text{Gy}$, and no part in the urethra receives a dose higher than 125%. Constraints (7.5e)–(7.5g) are the same as constraints (7.3d)–(7.3f) in the (LD) model.

By fixing the parameters b_k, (LDV) becomes a dwell time optimization model. This enables a comparison with IPSA, which is not able to optimize catheter positions.

One of the advantages of directly optimizing on clinically relevant criteria is the possibility to extend the model in a clinically interpretable way. Suppose for instance that the maximum number of catheters is not fixed *a priori*, but the planner wants to insert an extra catheter only if it leads to an improvement of $V_{100\%}$ of at least 5%. This can be incorporated into (LDV) by changing the objective to $1/|I_{PTV}| \sum_{i \in I_{PTV}} v_i + 0.05N$ and by treating N as a variable rather than as a parameter.

7.3 Numerical evaluation

7.3.1 Patient data

Clinical data from three different patients have been obtained from the treatment planning system (HDRplus, version 3.0, Eckert & Ziegler BEBIG GmbH, Berlin, Germany). Characteristics are summarized in Table 7.5. Approximately 2500 calculation points have been hexagonally distributed over the PTV, rectum and urethra. The number of potential catheter positions in the template is 40, 49 and 43 for patient 1, 2 and 3, having a prostate volume of 50 cc, 75 cc and 81 cc, respectively. According to our clinical protocol, the PTV had been extended with a 2 mm margin, and dwell positions were activated with a separation of 3 mm. A transperineal needle template with a hole resolution of 5 mm was used (Martinez Prostate Template, Nucletron BV, Veenendaal, the Netherlands). The dose rates have been calculated using the TG-43 formalism with the source parameters according to (Granero et al., 2006).

Table 7.5 – Characteristics of the patient data.

| Structure | α_i | L_i | β_i | U_i | Number of calculation points | | |
					Patient 1	Patient 2	Patient 3
PTV	8	8.5	3	25	1732	1834	1791
Rectum	0	0	10	8	246	234	240
Urethra	0	0	10	10	489	473	495

7.3.2 Inverse planning simulated annealing (IPSA)

We compare our results with the IPSA implementation in HDRplus, which exploits the linear penalty function (7.2) and was configured as follows. The composite objective function did not include the total dwell time. A maximum weight was used for the DTMR. The trade-off between speed and quality was set to its default value. After three consecutive runs, the plan with the lowest objective value was selected.

7.3.3 Our optimization models

For the model parameter values we set $t_{\max}$ at 5 seconds for an apparent source activity of 370 GBq. The DTMR parameter γ was set at 10%, and the maximum allowed number of catheters N was varied between 15 and 20.

All computing times reported have been obtained with the optimization software AIMMS 3.10 x64 using ILOG CPLEX 12.1 as solver running on Windows 7 x64 on

an Intel Core i5 660 (3.33 GHz) processor with 8 GB of RAM.

The exclusion restriction reduces the number of allowed catheter configurations with a factor 10^3–10^8, depending on the prostate volume and number of catheters.

The solution time and objective values of all models are listed in Tables 7.6–7.8. The first line is read as follows: for patient 1, the IPSA model, which is optimized for 16 preselected catheters, returns an optimized plan within 0.8 minutes. If the (LD) model, (QD) model, Algorithm 7.B or the (LDV) model had chosen the same dwell times as IPSA, they would have given objective values of 0.9280, 189, 1.1229 or 86.1, respectively. The other columns show the dosimetric plan performance.

7.3.3.1   Linear dose-based optimization

The solution times for the (LD) model are 5, 364 and 3 minutes for patients 1, 2 and 3, respectively when the allowed number of catheters is 20. The high solution time for patient 2 is probably due to the large number of feasible catheter configurations. We have not been able to obtain a solution within 24 hours without the exclusion restriction. Specifying constraint (7.3d) as an indicator constraint decreases calculation time by about 10%.

The convergence rate of the lower and upper bound of $V_{100\%}$ during the optimization process is depicted in Figure 7.2. We observe that most time is spent on obtaining a better lower bound. If we would terminate the solver when the upper bound is at most twice the lower bound, we would get a solution four times quicker at the cost of a 5% higher objective value. A slightly higher objective value does not translate into a significantly lower plan quality when evaluated at a lower level (Alterovitz et al., 2006; Karabis et al., 2009). Hence, stopping early will on average not result in clinically worse treatment plans.

When the catheter positions are fixed, the (LD) model reduces to an LP. CPLEX finds the optimal solution in 2.4 seconds. This is comparable to the solution time reported for a similar LP (Alterovitz et al., 2006; Karabis et al., 2009).

7.3.3.2   Quadratic dose-based optimization

First, we searched for a good value of p_i in the PTV. By repeatedly solving the model for different target doses and evaluating the dose distribution, we found that 21 Gy gave an acceptable treatment plan for patient 1. We used the same target value for patient 2 and 3.

The (QD) model solved in between 18 seconds and 1.5 minutes for $N = 20$ catheter positions. Again, we observed that the solution time decreased with a larger allowed number of catheters. Comparing the solution times to those of the (LD) model, we found that the (QD) model was at least ten times faster.

For the iterative procedure (Algorithm 4), we observed a drop in objective value

Figure 7.2 – Convergence of the lower and upper bound during the optimization of the (LD) model.

to 5% of the initial value after the first iteration. In the subsequent 15 steps, the additional decrement was 40–50%. After 14–24 iterations in total, the decrements in objective value were smaller than 10^{-3}, which is when the procedure was halted.

When the catheter positions are fixed, and only the dwell times are to be optimized, the (QD) model solves in 0.14, 0.9 and 0.9 seconds for patients 1, 2 and 3 respectively, while Algorithm 4 solves in 8.5, 13.3, and 11.8 seconds, respectively.

7.3.3.3 Dose-volume based optimization

The solver requires more than 24 hours to solve the (LDV) model to optimality. This is not problematic, because during execution the solver reports both $V_{100\%}$ of the best solution found so far, and an upper bound on $V_{100\%}$ that gradually gets lower. Hence, in a clinical setting, the treatment planner can stop the solver as soon as the value of $V_{100\%}$ is satisfactory. Here, we stopped as soon as $V_{100\%} \geq 95\%$ or after 15 minutes. Stopping the solver before optimality is reached is not a novel idea. It has been applied before with low-dose-rate brachytherapy (Gallagher and Lee, 1997; Lee and Zaider, 2003). For dwell time optimization, the solution time is 8.5 minutes for patient 2. By changing solver parameters, we have reduced the solution time for dwell time optimization to 14, 30 and 35 seconds for patients 1, 2 and 3, respectively. The new parameters make the solver first consider solutions in

which binary variables that are close to 0 in the LP relaxation, are fixed at 0. At least 90% of the binary variables must be 1. After fixing 10% of those variables to 0, the other variables must take the value 1 in the LP relaxation (Pryor and Chinneck, 2011). We have also tried Benders' decomposition (Benders, 1962), but it did not provide a speed-up.

7.3.4   Plan performance evaluation

All models generate plans with good OAR sparing. The DVH bounds on the rectum and the urethra are (almost) satisfied in all cases. The $V_{100\%}$ for the PTV is sufficiently high for all models except the (QD) model. $V_{150\%}$ and $V_{200\%}$ are sufficiently low for all models, except for the (QD) model.

The plan performance as assessed from the three-dimensional (3D) dose distribution by an experienced planner (ALH) is listed in the last column of Tables 7.6-7.8. The experienced planner paid most attention to conformality of the dose distribution and to whether high dose subvolumes (150% and 200%) around dwell positions were connected. Indeed a high correlation was found between the plan performance scored by the expert and the COIN value.

The (QD) model produces inacceptable plans due to a low $V_{100\%}$, especially for patient 2. All other models perform very well, the (LDV) model for patient 1 being an exception. The latter is due to the activation of dwell positions outside the PTV, giving rise to significant dose contributions outside the PTV. This can probably be avoided by activating dwell positions only inside the PTV or by adding constraints on auxiliary avoidance structures that limit dose outside the PTV.

The relation between the linear penalty function value and the expert opinion is very weak. This becomes clear from the objective values of the solution of the (LDV) model: for patient 3 with 18 catheters the plan of the (LDV) model is still preferred by the expert even though the linear objective value is 4.6 times higher (0.1038 vs. 0.4816) than the optimal plan of the (LD) model. Also the relation between the linear objective value and DVH statistics is weak. For patient 1 with 16 catheters, the plans of the (LD) and (LDV) models have similar DVH statistics, but their linear penalty value differs by a factor 12 (0.1338 vs. 1.6512).

The effect on the clinical evaluation criteria of extra catheters above 15 is small for all three patients. In most cases, the models could slightly improve $V_{100\%}$ by increasing the number of catheters, the largest improvement being 2.1%. The expert opinion is almost constant for a specific patient and model, with most variability related to patient 3. For this patient, the expert opinion always becomes more positive when more catheters were inserted.

7.4 Discussion

With existing dose-based models for inverse treatment planning of HDR brachytherapy, it is often a trial-and-error process to obtain adequate treatment plans that satisfy pre-set DVH criteria. In a clinical setting, this is undesirable because of the time burden and the required degree of user experience. Often, the aim is to design a plan with maximum achievable target coverage under fixed OAR dose-volume constraints. This type of optimization problem can be formulated and solved using mixed integer programming. This improves comprehensibility and plan quality compared to traditional dose-based inverse optimization. We investigated enhancements of existing linear and quadratic programming models for dose-based optimization and showed that the solution time could be decreased substantially.

We have limited our analysis to three representative clinical cases that cover a range of prostate sizes. The limited number of patients has allowed us to perform an extensive analysis of the effects of different algorithms on the allowed number of catheters. We realize that it is necessary to include more patients to further strengthen our conclusions.

In this dosimetric study, perturbations unavoidable in a clinical implementation were not taken into account. There is still a lack of validated data on the uncertainties, which makes it hard to assess the impact of these uncertainties on the optimality of a treatment plan. In future work the uncertainties need to be identified and quantified, and the effect on treatment plans needs to be evaluated. Current optimization techniques that deal with uncertainties such as "stochastic programming" or "robust optimization" require a model that finds good treatment plans when there is no uncertainty (Ben-Tal et al., 2009a; Kall and Wallace, 1994). This means that our work is also relevant for a future study on finding robust treatment plans.

For dwell time optimization, all models except (QD) can produce clinically good plans. This confirms current practice where the (LD) model is used.

Choosing catheter positions is still a difficult problem. Despite the exclusion restriction, only the (LDV) model produces plans in clinically acceptable time for all patients. There is one other article that uses an exclusion restriction, allowing a maximum of two catheters in any 2×2 square of template holes (Holm, 2011). This is weaker than our restriction. The HIPO algorithm is widely used and can optimize catheter positions in clinically acceptable time. HIPO inherently differs from our mixed integer programming approach, making it difficult to clarify any discrepancies.

DVH-based optimization has been applied to HDR brachytherapy by others (Lahanas et al., 2003b; Panchal, 2008; Siauw et al., 2011). All restrict to dwell time optimization and use heuristics for which mathematical optimality cannot be guaranteed.

Of these, only Siauw et al. (2011) provide a fast solution to a MILP formulation, but cannot provide a good upper bound for $V_{100\%}$.

Weak correlation between the (LD) objective value and DVH parameters was also observed by Holm (2011). Holm claims that the (LD) objective can make a rough division between good and bad plans. We confirm that there could be an order of magnitude difference in (LD) objective values among good plans. This implies there is a gap between the objective function in dose-based models and clinically desired properties of a dose distribution. The (LDV) model has the potential to close this gap, and to give the planner a better tool to steer the optimization.

7.5 Conclusion

With the proposed extensions, existing dose-based optimization models that simultaneously optimize catheter positions and dwell times can be solved more quickly to proven optimality with mixed integer programming techniques. Our dose-volume based model relates more closely with clinical parameters compared to dose-based models, and is faster to solve.

Acknowledgments

7.A Numerical evaluation

Table 7.6 – Treatment plan performance indicators for patient 1.

Model	N	Sol. time (min)	(LD) value	(QD) value	Alg. 7.B value	COIN	PTV (%)				Rectum (Gy)		Urethra (Gy)		Expert opinion
							$D_{90\%}$ ≥ 90%	$V_{100\%}$ ≥ 90%	$V_{150\%}$ ≤ 55%	$V_{200\%}$ ≤ 20%	$D_{10\%}$ ≤ 7.2 Gy	D_{2cc} ≤ 6.7 Gy	$D_{10\%}$ ≤ 10 Gy	$D_{0.1cc}$ ≤ 10 Gy	
IPSA*	16	2.1	0.109	166	0.029	0.666	110.8	97.2	46.6	17.2	6.5	5.5	9.6	9.7	+
(LD)	**15**	**11.0**	**0.158**	**163**	**0.123**	**0.729**	**111.3**	**96.6**	**44.8**	**13.4**	**6.0**	**4.9**	**9.7**	**9.9**	++
(LD)	**16**	**7.0**	**0.134**	**162**	**0.069**	**0.727**	**111.8**	**97.1**	**45.6**	**13.6**	**6.0**	**4.9**	**9.7**	**9.9**	++
(LD)	**17**	**4.9**	**0.131**	**162**	**0.064**	**0.729**	**111.7**	**97.1**	**44.8**	**13.1**	**5.9**	**4.8**	**9.7**	**9.9**	++
(LD)	**18**	**4.9**	**0.123**	**162**	**0.031**	**0.738**	**111.6**	**97.4**	**44.0**	**12.9**	**5.7**	**4.7**	**9.7**	**9.9**	++
(LD)	**19**	**5.3**	**0.121**	**162**	**0.029**	**0.741**	**111.6**	**97.1**	**42.6**	**12.2**	**5.6**	**4.7**	**9.7**	**9.9**	++
(LD)	**20**	**4.9**	**0.118**	**162**	**0.028**	**0.739**	**112.4**	**97.2**	**43.1**	**12.0**	**5.6**	**4.7**	**9.8**	**9.9**	++
(QD)	15	1.0	1.296	151	1.417	0.625	99.3	89.5	50.0	22.5	5.7	4.7	8.8	8.9	-
(QD)	16	0.5	1.238	150	1.101	0.641	98.7	88.8	50.1	25.2	5.5	4.6	8.9	9.1	-
(QD)	17	0.3	1.167	149	1.015	0.645	99.0	89.3	52.0	24.5	5.5	4.6	9.1	9.3	-
(QD)	18	0.3	1.167	149	1.015	0.645	99.0	89.3	52.0	24.5	5.5	4.6	9.1	9.3	-
(QD)	19	0.3	1.167	149	1.015	0.645	99.0	89.3	52.0	24.5	5.5	4.6	9.1	9.3	-
(QD)	20	0.3	1.167	149	1.015	0.645	99.0	89.3	52.0	24.5	5.5	4.6	9.1	9.3	-
Alg. 7.B	15	4.9	0.341	159	0.055	0.700	106.4	94.2	46.5	14.7	5.8	4.7	9.4	9.6	+
Alg. 7.B	16	4.0	0.226	160	0.031	0.713	108.9	95.6	45.0	13.7	5.9	4.8	9.5	9.7	+
Alg. 7.B	17	3.6	0.243	158	0.032	0.710	106.9	94.7	45.2	15.1	5.7	4.7	9.4	9.5	+
Alg. 7.B	**18**	**3.2**	**0.218**	**158**	**0.025**	**0.715**	**109.7**	**96.1**	**47.0**	**14.8**	**5.7**	**4.7**	**9.7**	**9.9**	++
Alg. 7.B	19	2.8	0.250	158	0.025	0.714	110.1	96.2	49.2	14.7	5.7	4.7	9.8	10.0	+
Alg. 7.B	20	2.6	0.245	158	0.023	0.712	110.5	96.3	51.1	14.3	5.8	4.8	9.8	10.0	+
(LDV)	12	0.3	3.002	188	22.249	0.613	113.6	98.9	52.3	27.0	6.2	5.3	9.8	9.9	- -
(LDV)	16	0.3	1.651	176	8.558	0.687	113.4	98.9	43.2	19.0	6.0	5.0	9.9	10.0	- -
(LDV)*	**16**	**0.1**	**1.623**	**173**	**8.813**	**0.703**	**113.1**	**98.6**	**48.2**	**22.7**	**6.0**	**5.1**	**9.9**	**10.0**	++

* Using the catheter configuration chosen by an expert planner (ALH).

Abbreviations: N = allowed number of catheters, Sol. time = solution time, COIN = conformality index (Baltas et al., 1998).

Plans rated ++ are displayed in bold.

Table 7.7 – Treatment plan performance indicators for patient 2.

Model	N	Sol. time (min)	(LD) value	(QD) value	Alg. 7.B value	COIN	PTV (%)				Rectum (Gy)		Urethra (Gy)		Expert opinion
							$D_{90\%}$ ≥90%	$V_{100\%}$ ≥90%	$V_{150\%}$ ≤55%	$V_{200\%}$ ≤20%	$D_{10\%}$ ≤7.2 Gy	D_{2cc} ≤6.7 Gy	$D_{10\%}$ ≤10 Gy	$D_{0.1cc}$ ≤10 Gy	
IPSA*	16	2.2	0.107	203	0.019	0.709	111.3	97.6	37.8	11.2	7.6	6.9	9.6	9.7	+
(LD)	15	1986.6	0.175	198	0.033	0.748	108.3	95.9	36.7	12.2	7.3	6.6	9.8	9.9	+
(LD)	16	593.8	0.156	194	0.030	0.724	110.2	96.1	42.9	13.9	7.3	6.6	9.8	9.9	+
(LD)	17	344.6	0.146	192	0.028	0.716	111.0	96.6	46.3	15.4	7.3	6.7	9.8	9.9	+
(LD)	18	520.0	0.144	191	0.028	0.717	111.1	96.7	47.1	15.6	7.4	6.7	9.8	9.9	+
(LD)	19	1276.3	0.137	193	0.028	0.718	110.6	96.6	44.9	15.5	7.4	6.7	9.8	9.9	+
(LD)	20	364.4	0.134	193	0.029	0.723	110.4	96.7	43.7	13.8	7.4	6.7	9.8	9.9	+
(QD)	15	22.8	3.242	174	1.127	0.557	80.6	74.3	42.4	23.2	6.0	5.4	8.3	8.5	– –
(QD)	16	9.0	3.186	173	1.105	0.561	80.6	74.9	43.4	23.4	6.1	5.5	8.3	8.5	– –
(QD)	17	5.3	3.171	172	1.065	0.560	80.4	74.8	43.4	24.4	6.1	5.5	8.3	8.5	– –
(QD)	18	2.3	2.879	172	1.000	0.569	82.2	75.8	42.8	23.4	6.1	5.5	8.4	8.5	– –
(QD)	19	1.8	2.762	171	0.964	0.574	82.8	76.4	42.4	23.3	6.1	5.6	8.3	8.5	– –
(QD)	20	1.5	2.747	171	0.906	0.576	82.6	76.6	44.0	23.4	6.1	5.6	8.3	8.5	– –
Alg. 7.B	15	134.0	0.393	186	0.048	0.725	103.7	92.6	45.3	14.5	6.8	6.2	9.7	9.8	+
Alg. 7.B	16	70.6	0.362	185	0.043	0.719	104.0	92.9	47.5	14.8	6.8	6.2	9.7	9.8	+
Alg. 7.B	17	41.0	0.359	183	0.043	0.700	104.0	92.8	49.0	15.2	6.8	6.2	9.7	9.8	+
Alg. 7.B	18	27.3	0.326	184	0.041	0.709	105.9	93.6	48.8	13.4	7.0	6.4	9.7	9.8	+
Alg. 7.B	19	17.1	0.321	184	0.040	0.709	106.0	93.5	50.4	13.5	7.0	6.4	9.7	9.9	+
Alg. 7.B	20	24.9	0.321	183	0.040	0.713	105.7	93.5	50.6	14.2	7.0	6.4	9.7	9.9	+
(LDV)	15	13	1.228	202	4.579	0.753	108.6	95.9	39.9	14.1	7.0	6.4	10.0	10.1	+
(LDV)	17	15	0.929	205	2.378	0.731	105.5	93.8	31.5	10.8	7.2	6.5	9.9	10.1	+
(LDV)	19	2.7	1.875	211	16.519	0.736	106.5	94.4	36.5	13.8	6.8	6.2	9.7	9.8	+
(LDV)	20	2.3	1.228	205	6.033	0.739	107.7	95.5	35.4	12.2	7.1	6.4	9.8	9.9	+
(LDV)*	16	0.3	1.143	203	6.016	0.754	107.5	95.6	39.8	11.6	7.0	6.4	9.9	10.0	+

* Using the catheter configuration chosen by an expert planner (ALH).

Abbreviations: N = allowed number of catheters, Sol. time = solution time, COIN = conformality index (Baltas et al., 1998).

Table 7.8 – Treatment plan performance indicators for patient 3.

Model	N	Sol. time (min)	(LD) value	(QD) value	Alg. 7.B value	COIN	PTV (%)				Rectum (Gy)		Urethra (Gy)		Expert opinion
							$D_{90\%}$ $\geq 90\%$	$V_{100\%}$ $\geq 90\%$	$V_{150\%}$ $\leq 55\%$	$V_{200\%}$ $\leq 20\%$	$D_{10\%}$ $\leq 7.2\,\mathrm{Gy}$	D_{2cc} $\leq 6.7\,\mathrm{Gy}$	$D_{10\%}$ $\leq 10\,\mathrm{Gy}$	$D_{0.1cc}$ $\leq 10\,\mathrm{Gy}$	
IPSA*	14	2.0	0.347	196	0.391	0.706	108.2	95.1	36.1	12.7	7.1	6.6	9.6	9.7	+
(LD)	15	10.1	0.120	203	0.011	0.723	112.0	97.2	32.0	10.2	7.3	6.8	10.1	10.1	+
(LD)	16	7.9	0.112	204	0.009	0.737	113.0	97.5	32.0	9.9	7.3	6.8	10.1	10.2	+
(LD)	17	5.9	0.106	203	0.009	0.735	113.3	97.5	32.4	9.7	7.3	6.8	10.1	10.2	+
(LD)	18	3.5	0.104	204	0.009	0.743	113.0	97.5	31.6	9.2	7.3	6.8	10.1	10.2	+
(LD)	**19**	**2.7**	**0.102**	**205**	**0.009**	**0.748**	**112.3**	**97.5**	**30.6**	**9.3**	**7.4**	**6.8**	**10.1**	**10.1**	++
(LD)	**20**	**2.9**	**0.102**	**205**	**0.009**	**0.748**	**112.3**	**97.5**	**30.6**	**9.3**	**7.4**	**6.8**	**10.1**	**10.1**	++
(QD)	15	1.3	2.355	176	0.713	0.613	87.9	81.6	44.8	21.9	6.7	6.1	9.2	9.4	- -
(QD)	16	1.2	2.081	175	0.619	0.623	89.8	83.0	44.8	21.3	6.7	6.1	9.3	9.5	-
(QD)	17	0.9	2.008	175	0.587	0.629	90.4	83.4	45.0	21.0	6.7	6.1	9.3	9.5	-
(QD)	18	0.9	1.998	175	0.580	0.630	90.4	83.4	45.1	20.8	6.7	6.1	9.3	9.5	-
(QD)	19	0.9	1.998	175	0.580	0.630	90.4	83.4	45.1	20.8	6.7	6.1	9.3	9.5	-
(QD)	20	0.8	1.998	175	0.580	0.630	90.4	83.4	45.1	20.8	6.7	6.1	9.3	9.5	-
Alg. 7.B	15	22.2	0.502	183	0.057	0.689	107.2	94.8	50.1	15.7	7.2	6.6	10.1	10.3	+
Alg. 7.B	16	19.1	0.542	184	0.058	0.683	107.7	94.8	52.0	13.9	7.3	6.7	10.1	10.3	+
Alg. 7.B	**17**	**16.6**	**0.504**	**184**	**0.055**	**0.684**	**108.4**	**95.3**	**52.2**	**13.2**	**7.3**	**6.8**	**10.1**	**10.3**	++
Alg. 7.B	**18**	**14.9**	**0.508**	**184**	**0.056**	**0.683**	**108.4**	**95.3**	**52.3**	**14.6**	**7.3**	**6.8**	**10.1**	**10.3**	++
Alg. 7.B	**19**	**14.6**	**0.529**	**184**	**0.055**	**0.686**	**108.4**	**95.3**	**53.0**	**14.2**	**7.3**	**6.7**	**10.2**	**10.3**	++
Alg. 7.B	**20**	**16.1**	**0.543**	**184**	**0.053**	**0.688**	**108.5**	**95.3**	**53.5**	**13.4**	**7.3**	**6.7**	**10.2**	**10.4**	++
(LDV)	15	15	0.814	202	3.153	0.707	108.4	94.7	35.9	12.3	6.7	6.2	10.0	10.1	o
(LDV)	16	1.4	0.426	204	1.835	0.732	111.6	96.9	31.2	9.9	7.1	6.6	10.1	10.2	+
(LDV)	17	1.6	0.483	202	1.377	0.717	109.7	95.9	37.1	10.4	7.0	6.4	10.1	10.1	+
(LDV)	**18**	**1.8**	**0.482**	**204**	**1.565**	**0.733**	**109.6**	**95.7**	**33.4**	**9.9**	**6.8**	**6.3**	**10.1**	**10.2**	++
(LDV)*	14	8.5	0.963	200	3.718	0.710	108.7	95.5	39.1	14.2	7.2	6.6	10.1	10.1	+

* Using the catheter configuration chosen by an expert planner (ALH).
Abbreviations: N = allowed number of catheters, Sol. time = solution time, COIN = conformality index (Baltas et al., 1998).
Plans rated ++ are displayed in bold.

7.B   Iterative procedure for adjusting the target dose in the (QD) model

---

Algorithm 4 Iterative procedure for adjusting target dose in (QD)

---

    for $i \in I$ **do**
        $p_i := (L_i + U_i)/2$
    end for
    CURVAL $:= \infty$
    repeat
        solve the (QD) model
        OLDVAL $:=$ CURVAL
        CURVAL $:=$ (QD).value
        for $i \in I$ **do**
            if $(\dot{\mathbf{D}}\mathbf{t})_i < L_i$ **then**
                $p_i := L_i$
            else if $(\dot{\mathbf{D}}\mathbf{t})_i > U_i$ **then**
                $p_i := U_i$
            else
                $p_i := (\dot{\mathbf{D}}\mathbf{t})_i$
            end if
        end for
    until OLDVAL $-$ CURVAL $\leq \epsilon$

7.C   Counter example for global convergence of the iterative procedure

We show that Algorithm 4 does not always converge to the global optimum that we would have obtained by optimizing the (QD) model simultaneously over $\mathbf{t}$, $\mathbf{p}$ and b_k.

Consider two mutually exclusive catheter positions, each offering one dwell position. There are three calculation points; the first one has $(L_1, U_1) = (8, 10)$ and the other two have $(L_2, U_2) = (L_3, U_3) = (10, 15)$. All dose rates are 1 except $d_{22} = d_{23} = 2$. In the first iteration we start with prescribed dose vector $\mathbf{p} = (9\ 12.5\ 12.5)$, in which case catheter position 2 is optimal with a dwell time of $118/18$ s. The new prescribed dose vector becomes $\mathbf{p} = (8\ 13.1\ 13.1)$. In every subsequent iteration, the dose vector gets closer to $(8\ 15\ 15)$, resulting in a total penalty of $2/9$. But the optimal dose distribution is obtained by selecting catheter position 1, which can deliver a dose of 10 Gy to all three calculation points, resulting in zero penalty.

HDR prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing

Abstract High-dose-rate brachytherapy is a tumor treatment method where a highly radioactive source is brought in close proximity to the tumor. In this chapter we develop a simulated annealing (SA) algorithm to optimize the dwell times at preselected dwell positions to maximize tumor coverage under dose-volume constraints on the organs at risk. SA is extended to cope with hard constraints, and changes have been made to ensure good performance. Compared to existing algorithms, our algorithm has advantages in terms of speed and objective value and does not require an expensive general purpose solver. Its success mainly depends on exploiting the efficiency of matrix multiplication and a careful selection of the neighboring states. In this chapter we outline its details and make an in-depth comparison with existing methods using real patient data.

8.1 Introduction

High-dose-rate brachytherapy is a form of radiation therapy in which the tumor is temporarily exposed to a highly radioactive source which dwells at different positions in or around the Planning Target Volume (PTV). For prostate tumors, the dwell positions are offered by a temporary transperineal implant of catheters which run through the prostate. A remote afterloader, which is connected to all the catheters, then sends a radioactive source through the catheters one by one, stopping at several dwell positions inside the PTV. We assume that the catheter locations and the dwell positions are known. The goal of the treatment planner is to irradiate the PTV while sparing the surrounding Organs at Risk (OARs) by optimizing the dwell time at each dwell position. Traditionally, this has been done by forward planning,

Figure 8.1 – Dose-volume histogram for a prescribed dose of 8.5 Gy.

a trial-and-error procedure in which dwell times are adjusted until the dose distribution is satisfactory. In the last fifteen years, inverse planning has been developed as a computerized technique where the dwell times are optimized according to the treatment planner's preferences.

In the remainder of this introduction, we explain how dose distributions are evaluated, we provide a literature review on the ongoing progress of inverse planning, and we show how this work contributes to these developments.

DVH Criteria. The dose distribution within the PTV or an OAR is described in a dose-volume histogram (DVH). Such a histogram displays the percentage of the structure receiving at least a certain dose. An example histogram is displayed in Figure 8.1. Points on this histogram are denoted by D_x or V_y, e.g. $D_{90\%}$ is the minimum dose received in the hottest 90% of the structure under consideration and $V_{150\%}$ represents the percentage of the volume receiving 150% of the prescribed dose. These statistics are currently the most important quantitative evaluation criteria (Hoskin et al., 2013). Besides these, the treatment planner inspects the isodose lines in order to avoid undesired hot spots and cold spots.

Table 8.1 displays clinically used DVH criteria from a local hospital. In order to achieve tumor control, $V_{100\%}$ of the PTV needs to be at least 90%, i.e., at least 90% of the PTV's volume need to receive at least 100% of the prescribed dose.

PTV	Rectum	Urethra
$V_{100\%} \geq 90\%$	$D_{10\%} \leq 7.2\,\mathrm{Gy}$	$D_{10\%} \leq 10\,\mathrm{Gy}$
	$D_{\max} \leq 8\,\mathrm{Gy}$	$D_{\max} \leq 10.6\,\mathrm{Gy}$

Table 8.1 – DVH protocol for a prescribed dose of 8.5 Gy for HDR prostate brachytherapy based on Hoskin et al. (2007).

The limitation of complications is expressed in DVH constraints on the rectum and urethra. $D_{10\%} \leq 7.2\,\mathrm{Gy}$ for the rectum imposes that 90% of the rectum's tissue may not receive more than 7.2 Gy. A hard upper bound for tissue in the rectum is set by $D_{\max} \leq 8\,\mathrm{Gy}$. For the urethra, 10 Gy is the upper bound for 90% of the tissue. The remaining 10% are allowed to reach radiation levels of at most 10.6 Gy.

Dose measurement. In order to compute DVH statistics, the PTV and the OARs are discretized into a grid of calculation points. The dose rates d_{ij} describe the amount of radiation emitted from dwell position $j \in J$ towards calculation point $i \in I$ per second. The dose at a calculation point i is equal to the sum of radiation received from each dwell position j:

$$(Dt)_i = \sum_{j \in J} d_{ij} t_j, \tag{8.1}$$

where t_j is the dwell time at dwell position j. Let D be the $|I| \times |J|$ matrix with all dose rates d_{ij}. The matrix-vector product Dt yields a vector that denotes the dose per calculation point.

Literature review. Traditionally, inverse planning is dose-based: the treatment planner prescribes a lower bound L_i and and an upper bound U_i on the desired dose for each calculation point. Typically these bounds are the same for all calculation points within the same structure. At each calculation point i, the received dose $(Dt)_i$ is compared to its respective prescribed lower and upper bounds. If the dose lies outside the interval $[L_i, U_i]$, a penalty is imposed that is linear or quadratic in the deviation. A treatment plan is optimal if it minimizes the sum of penalties over all calculation points. Linear penalty functions have been successfully implemented in the commercially available algorithms Inverse Planning by Simulated Annealing (IPSA) (Lessard and Pouliot, 2001) and Hybrid Inverse Treatment Planning and Optimization (HIPO) (Karabis et al., 2005). Quadratic penalty functions are used by Lahanas et al. (2003a); Lahanas and Baltas (2003); Milickovic et al. (2002).

The disadvantage of dose-based penalty functions is that the resulting treatment plans may need a posteriori adjustments as they do not necessarily adhere to the DVH criteria. More recently, dose-volume based optimization methods have been

developed that directly optimize on DVH criteria. Panchal (2008) has formulated a Harmony Search heuristic. Siauw et al. (2011) have presented a Mixed Integer Linear Program (MILP) and suggested a heuristic (IPIP) that determines a solution that is feasible and near-optimal by solving two LPs. In Chapter 7, we have devised to directly optimize an MILP using specific solver settings. Beliën et al. (2009) have introduced a hybrid approach based on simulated annealing and linear programming. Their objective is dose-based, while the delivered radiation to OARs is now also subject to dose-volume constraints. For the constraint $D_{10\%} \leq 7.2$ Gy, the linear constraint $(Dt)_i \leq 7.2$ needs to hold for 90% of the calculation points. The authors utilize simulated annealing to determine which calculation points are in the group of 90%. According to Beliën et al. (2009), the algorithm yields results superior to running the integer program alone based on ten problem instances with 30 minutes computation time each.

Our contribution. This chapter follows up on Beliën et al. (2009)'s suggestion to develop a pure simulated annealing (SA) heuristic to optimize on DVH statistics. We present DOPSA (DVH Optimization by Pure Simulated Annealing), a novel algorithm that solves the same MILP model as Siauw et al. (2011) and as in Chapter 7. In this model the objective is to maximize $V_{100\%}$ for the PTV, i.e., the volume of the PTV receiving the prescribed dose, while meeting DVH constraints on the OARs. Our method exploits the speed of matrix multiplication, and has been tuned by conducting many trial-and-error experiments. The parameters have been selected using a metamodel.

The advantage of SA over the existing dose-volume based models lies in its simple implementation independent of costly LP or MILP solvers. Since optimization is often offered as a separate module for treatment planning systems, clinics may choose cheaper dose-based optimization modules or keep using them if the price difference is too large, leaving the potential of dose-volume based methods unused. Furthermore, the results in Section 8.4 show that our heuristic has clear advantages over existing methods in terms of speed and objective value.

The significance of this result becomes apparent when investigating the structure of the MILP formulation. The objective is to maximize the number of calculation points that receive the prescribed dose of 8.5 Gy. Let us first define the vector x, indexed by i in I_{PTV} as follows:

$$x_i = \max\{0, \frac{8.5 - (Dt)_i}{8.5}\}. \tag{8.2}$$

Since x_i is zero if calculation point i receives at least the prescribed dose, minimizing the number of non-zero elements in x yields a treatment plan with maximal $V_{100\%}$. The number of nonzero elements is $||x||_0$ (recall that the l_0-norm counts the number

of non-zero elements in a vector). Attaining sparse vectors by minimizing the l_0-norm has been shown to be NP-hard (Natarajan, 1995).

An MILP solver first solves the LP relaxation:

$$\max_{x \in [0,1], t \geq 0} \Big\{ \sum_{i \in I_{PTV}} x_i : (Dt)_i \geq 8.5 x_i, (x,t) \in \mathcal{X} \Big\}, \tag{8.3}$$

where $\mathcal{X}$ is a polyhedron. Substituting $y_i = 1 - x_i$ yields the following equivalent program:

$$\min_{y \in [0,1], t \geq 0} \Big\{ \sum_{i \in I_{PTV}} |y_i| : y_i \geq \tfrac{8.5 - (Dt)_i}{8.5}, (1 - y, t) \in \mathcal{X} \Big\}. \tag{8.4}$$

So, the LP relaxation is equivalent to l_1-norm minimization. Literature, predominately in the field of image processing, suggests that solutions found for the LP relaxation (l_1-norm minimization) provide good approximations of the solutions for the MILP (l_0-norm minimization).

Donoho (2006) shows that, in most cases, the l_1-approximation is also the sparsest solution for underdetermined systems of equations. Li (2010) discusses necessary and sufficient conditions for the equivalence of l_0-norm and l_1-norm solutions in general. These conditions do not hold in our case. Nevertheless, Candes et al. (2005) provide numerical evidence for the power of these l_1-approximations and also the results in Chapter 7 show that the LP relaxation is strong.

This may explain the result in Chapter 7 that an MILP solver quickly determines good treatment plans. However, our results in Sections 8.4 show that the DOPSA still outperforms an MILP solver in finding good solutions in a short time-span. Bearing in mind the observations from literature, one can therefore conclude that this SA heuristic is a powerful optimization tool for inverse treatment planning.

In the next section, the concept of SA will be outlined. The implementation details of DOPSA will be explained in Section 8.3, followed by a discussion of the results and a comparison with existing algorithms in Section 8.4.

8.2 Simulated annealing

SA is an iterative metaheuristic introduced by Kirkpatrick et al. (1983) in analogy to an annealing process in metallurgy. The heuristic is initiated in a feasible state t^{start} and proceeds through the feasible region searching the global maximizer t^{opt} of an objective function f. The distinguishing characteristic of SA is that it avoids premature termination in a local optimum by accepting non-improving states with a given probability. In every iteration, a new state t^{new} is selected from the neighborhood

of the current state t^{cur}. This new state is accepted with an acceptance probability, also named the metropolis criterion:

$$\mathbb{P}(\text{accept } t^{new}) = \begin{cases} 1 & \text{if } f(t^{new}) \geq f(t^{cur}) \\ \exp \frac{f(t^{new}) - f(t^{cur})}{c_k} & \text{otherwise,} \end{cases} \tag{8.5}$$

where t^{cur} is the state with the current objective value $f(t^{cur})$, and $c_k \in (0, \infty)$ is the variable control parameter representing the temperature. Hence, any new state t^{new} is accepted with probability 1 if the objective is at least as high as in the current state t^{cur}. States with lower objective value are accepted with a probability decreasing in the loss in objective value $f(t^{new}) - f(t^{cur})$ and increasing in temperature c. Therefore, the search path is increasing in objective values except for occasional drops, which facilitate break-outs from local maxima. Kirkpatrick et al. (1983) have proven, given certain conditions, that objective values increase along the search path and converge to the global optimum.

Cooling schedules govern the decrease in temperature c_k from its initial value c_0, which in turn decreases the probability to accept states with lower objective values. Essentially, SA steadily reduces to a hill-climbing procedure as the temperature approaches zero. The Exponential Schedule (Salamon et al. 2002, p. 91) sets the temperature at iteration k to

$$c_k = c_0 \alpha^{\left\lfloor \frac{k}{m} \right\rfloor}, \tag{8.6}$$

where $m \in (0, \infty)$ is a fixed number of iterations after which the temperature is rescaled by $\alpha \in (0, 1)$.

8.3 Implementation

In this section we provide in-depth information about DOPSA. Its final implementation is the result of a long trial-and-error process where many ideas have been tested. Here, we restrict ourselves to a description of the final algorithm. For the selection procedure and parameter tuning based on trial-and-error and a metamodel we refer the reader to Deist (2013).

First, we describe the optimization model (Section 8.3.1) and the main steps in the algorithm (Section 8.3.2). Then we focus on the two most important steps, which are the generation of neighborhood states (Section 8.3.3) and checking feasibility (Section 8.3.4).

8.3.1 Optimization model

We use the same model as (Siauw et al., 2011) and Chapter 7 to allow a comparison with their results. This model is based on the DVH protocol from Table 8.1. The objective is to maximize $V_{100\%}$ under $D_{10\%}$ and $D_{\max}$ constraints on the OARs. Additionally, to avoid high doses outside the PTV in regions where there is no OAR to limit the dose in that region, we also consider an artificial structure that surrounds the prostate at 2 mm distance. This collection of tissue is henceforth denoted as the 'shell', and a constraint is added to ensure that the maximum dose in this structure is below $8.5\,\mathrm{Gy}$. Additionally, a dwell time modulation restriction (DTMR) is added to avoid excessively high dwell times. The DTMR prohibits the dwell times of two adjacent dwell positions within the same catheter to differ more than a factor two.

8.3.2 Algorithm

DOPSA has been implemented in MATLAB according to Algorithm 5. DOPSA's code is designed to generate and evaluate multiple states simultaneously. All time consuming procedures within the heuristic thus become matrix-operations which decreases the required CPU time per state. Essentially, it is faster to check the feasibility of n states in one operation, than to check the feasibility of the n states individually, since computing DT, where T is a $|J| \times n$ matrix with the n states as columns, takes less than n times longer than computing Dt for a single state t.

Algorithm 5 Simulated Annealing Code

---
    set the initial state to $t = 0$
    set the temperature to c_0
    while running time does not exceed limit **do**
        every 15000 iterations: return to current optimum
        generate neighborhood states
        discard infeasible states
        choose t^{new} from neighborhood according to highest $V_{100\%}$
        decide acceptance
        lower the temperature by scaling with α
    end while

DOPSA is initialized with the dwell times set to zero because it should start in the feasible region and, a priori, any other feasible state can only be determined at relatively high computational effort. Then, one state is chosen from the set of feasible candidates in the neighborhood and is evaluated by the metropolis criterion as described in Section 8.2. Subsequently, the temperature is decreased in every itera-

tion according to an exponential cooling schedule (see Section 8.2) with parameters $\alpha = 0.99$, $m = 1$, and $c_0 = 1.5$. The search continues in the chosen state, if accepted, or restarts in the current state. After each 15,000 iterations, the state with the highest $V_{100\%}$ is selected and the search is restarted in the neighborhood of this state. This is to recover from an unsuccessful attempt to escape from a local optimum. Returning to the state with the highest $V_{100\%}$ focuses the search to areas with high objective function values and avoids unnecessary search efforts in low-potential regions.

The running time of DOPSA is restricted to a fixed limit to allow a comparison with existing methods and to test whether DOPSA can provide satisfactory results within a short time-span. In practice, the treatment planner would run the heuristic until satisfactory results are returned, i.e., when the coverage is sufficiently high.

The performance of the algorithm revealed to be insensitive to changes in the parameter values of α and c_0. This suggests that hill-climbing methods might be similarly successful.

8.3.3 Generating neighborhood states

DOPSA generates a neighborhood solution t by perturbing the dwell times of the current state t^{cur}:

$$t = t^{cur} + \lfloor r \rceil,$$

where r is a vector of the same size as t. The vector r is chosen randomly according to $r_j \sim N(0, 0.05)$ i.i.d. for all j in J_c, the set of dwell positions that are changed, and $r_j = 0$ for all j in $J \backslash J_c$. The operator $\lfloor \rceil$ rounds the elements of r to one decimal place, which is the input precision of the Flexitron remote afterloader (Nucletron BV, Veenendaal, the Netherlands). Negative dwell times are subsequently set to 0.

For the DTMR, the dwell times in each catheter are adjusted one by one to ensure the relative difference with the previous dwell time is less than a factor 2. This may increase the number of changed entries in the dwell time vector.

Number of changed entries. We observed that controlling the number of dwell time changes applied to a new state is fundamental to the performance of DOPSA. A lower number of changes is more likely to yield a feasible new state. However, lowering this number decreases the potential maximal increase in $V_{100\%}$ in the new state, which is most notable in early iterations.

For this reason, a negative correlation between the number of changed entries and the improvement over the past 200 iterations is introduced. In the first 200 iterations, all entries are subject to changes before the dynamic adjustment begins.

The delay of 200 iterations only affects the initial phase of DOPSA since the total number of iterations is approximately 50000.

Let J_c denote the index set of entries in t^{cur} that are varied to construct a new vector t. The relation between the cardinality of J_c and the improvement in $V_{100\%}$ is defined by

$$|J_c| = \lceil \frac{V_{100\%}^k - V_{100\%}^{k-200,max}}{V_{100\%}^{k-200,max}} |J| \rceil, \tag{8.7}$$

where $V_{100\%}^k$ is the $V_{100\%}$ attained in the current state t^{cur}. $V_{100\%}^{k-200,max}$ is the overall highest $V_{100\%}$ attained until iteration $k - 200$.

Note that the states chosen during the search can also exhibit $V_{100\%}$ lower than in preceding iterations due to the metropolis criterion, resulting in a negative $|J_c|$. For that reason and to maintain a natural upper bound, $|J_c|$ is restricted to the interval $[\lceil |J|/50 \rceil, |J|]$. The lower bound is chosen to be higher than 1 (assuming $|J| > 50$) since improvements in $V_{100\%}$ might only be attained with increases in dwell times for one dwell position while simultaneously decreasing the dwell time of another position. Taking 2% of all entries as the lower bound ensures sufficient dwell time changes in generated states.

Empirical results indicate that a newly generated state has a six times higher probability of being feasible by using J_c compared to perturbing all dwell times. The number of states that is dropped due to infeasibility decreases from 95% to 70%.

Dwell position selection. Once $|J_c|$, the number of dwell positions for which the dwell times are perturbed, is chosen, the next step is to select that many dwell positions. The intention is to alter dwell times such that the objective value is improved. This is achieved by changing the dwell times of the dwell positions close to the calculation points in the PTV that do not yet receive the prescribed dose. Since the distance between dwell position j and calculation point i has an inverse squared relation to the dose rate d_{ij}, the probability that dwell position j is chosen in a single draw is therefore set to:

$$P(t_j \in J_c) = \frac{\sum\limits_{i \in I_{PTV}:(Dt)_i < 8.5} d_{ij}}{\sum\limits_{j \in J} \sum\limits_{i \in I_{PTV}:(Dt)_i < 8.5} d_{ij}}, \tag{8.8}$$

where I_{PTV} is the set of calculation points in the PTV. The denominator normalizes the term to attain a probability distribution. Using probabilities rather than simply picking the closest dwell positions bears the advantage that the set of states that can be generated from the current state does not become too small. Occasionally also dwell positions further away are selected which diversifies the search.

Using this methods requires a random sample of dwell positions in each iteration of the search. Randomly sampling dwell positions without replacement in MATLAB is a computationally expensive procedure. Our experiments have shown that sampling $|J_c|$ items with replacement and removing duplicates yields better results, due to a larger number of iterations in the same running time. In a direct comparison, the average PTV coverage after a running time of 3 minutes over 50 replications was 0.12% higher when using sampling without replacement.

Number of generated states. In the early stages of the search process, improving states can be determined with few generated states per iteration since the improvement potential in the neighborhood of t^{cur} is high. This saves computational effort. Later on, an intensification of the search in one of those areas is most important. Moreover, intensification becomes necessary only when the search approaches the boundary, where one expects to find optimal solutions. Therefore, DOPSA generates more states when the boundary is reached (as suggested by Hedar and Fukushima (2006)). This can be controlled with an opposite relation between the number of states generated in a single iteration g and $V_{100\%}$:

$$g = \lceil (1 - \frac{V_{100\%}^{k} - V_{100\%}^{k-200,max}}{V_{100\%}^{k-200,max}}) \rceil. \tag{8.9}$$

g is limited to the interval $[1, 40]$ because trial runs indicated that excessive state generation does not further improve the DOPSA's performance. In the first 200 iterations, $g = 1$ is used.

8.3.4 Checking feasibility

The feasibility of each generated state needs to be verified, which is one of the computationally most expensive procedures next to computing $V_{100\%}$. Testing each constraint discussed in Section 8.1 requires a row subset of the matrix D corresponding to a specific OAR to be multiplied with the dwell time vector t. The computation time of checking a constraint increases with the number of calculation points in the corresponding OAR.

The sequence in which the constraints on the DVHs for the OARs are checked is chosen to minimize the required computation time: the most restrictive constraints with lowest computational cost are evaluated first. States that do not satisfy a constraint are immediately excluded in order to reduce the computation time for checking other constraints.

The calculation of constraint values and $V_{100\%}$ for newly generated states can be sped up by using the values for the current state t^{cur} that have been computed in the preceding iteration. Let T^{cur} denote the $|J| \times g$ matrix containing the current

state as columns, and let T^{new} denote the $|J| \times g$ matrix with the neighboring states as columns.

The dose computation for each generated state DT^{new} can be reformulated as follows:

$$DT^{new} = DT^{cur} + D(T^{new} - T^{cur}). \tag{8.10}$$

The matrix product DT^{cur} is known from the previous iteration, so only $D(T^{new} - T^{cur})$ needs to be computed. $T^{new} - T^{cur}$ is a sparse matrix with at most $|J_c|$ non-zero entries per column, since states generated in the neighborhood of a current state only differ in $|J_c|$-many entries. Moreover, the sparsity increases over the search process. Using MATLAB's toolpack for multiplication of sparse matrices, an increase of the number of iterations over the total running time of 5% up to 21% could be observed, depending on the number of dwell positions.

8.4 Results

The performance of DOPSA is assessed by a direct comparison with two existing dose-volume-based methods: the MILP formulation from Chapter 7 and the IPIP algorithm by Siauw et al. (2011). The running time has been fixed to 3 minutes on an Intel Xeon E5620 2.4 GHz with 6 GB of memory. Both MILP and IPIP are solved using the state-of-the-art solver CPLEX 12.4 (IBM).

8.4.1 Patient data

Tables 8.2 and 8.3 provide an overview of the three patient data sets. Each data set uses the same number of calculation points in the PTV and the OARs. The PTV contains the biggest share of calculation points since it is the largest organ. All dwell positions inside the PTV have been activated based on a step size of 2.5 mm, resulting in 115 to 236 dwell positions per patient. The number of inserted catheters is similar among patients with 17 for Patient 2 and 16 for the other two patients. The dose rates have been calculated using the TG-43 formalism.

8.4.2 Running time and objective value

DOPSA has been tested on three patients. Since it is a random local search procedure, the result can be subject to random effects. Therefore, the heuristic has been run 250 times for each patient in order to provide a detailed analysis of the average treatment plan quality and the variation in quality. If a low variation can be ensured

Table 8.2 – The number of calculation points per structure.

Structure	Calculation Points
PTV	1750
Urethra	500
Rectum	250
Shell	250

Table 8.3 – The number of dwell positions and catheters for each patient.

Patient	Dwell Positions	Catheters
1	236	16
2	115	17
3	182	16

for this optimization algorithm, only a single optimization run per patient will be necessary in a clinical setting to obtain treatment plans with consistent quality.

Figures 8.2-8.4 show the results for all three patient data sets. The average and standard deviation in PTV coverage over 250 replications are displayed over the running time of 180 seconds. The positioning of catheters determines the maximally achievable level of $V_{100\%}$ in the PTV, which explains the difference in objective values between patients.

In all three cases, the PTV coverage increases rapidly within the first 30 seconds and further improvements on the treatment plan quality are achieved over the remaining running time. Every treatment plan determined by DOPSA is feasible, i.e., it always satisfies the DVH criteria imposed on the OARs (see Table 8.1).

The absolute standard deviations in PTV coverage after a running time of 3 minutes for each of the 250 replications are 0.09%, 0.33%, and 0.23%, respectively. Therefore, DOPSA delivers treatment plans with consistent quality across replications.

Figure 8.2 – Average and standard deviation in PTV coverage for Patient 1 (250 replications).

Figure 8.3 – Average and standard deviation in PTV coverage for Patient 2 (250 replications).

Figure 8.4 – Average and standard deviation in PTV coverage for Patient 3 (250 replications).

8.4.3 Comparison with existing DVH-based optimizers

For the comparison with existing DVH-based optimizers, we have run an MILP solver with the solver settings from Chapter 7 and we have implemented IPIP (Siauw et al., 2011). These have been run on the same three patient data sets and on the same model. It is important to recognize that IPIP can only provide a single solution per data set, whereas MILP and DOPSA continue improving the solution. This is observable in Figures 8.5-8.7, where the line for IPIP remains flat after a single increase. In theory, MILP could stop when it proves that the solution is optimal, but this did not happen for any of the patients.

Figures 8.5 to 8.7 show the average PTV coverage obtained by DOPSA and the coverage obtained by the existing methods MILP and IPIP. The running time for Patient 3 has been extended to 5 minutes since MILP requires more than 3 minutes to find nonzero feasible points.

IPIP attains adequate PTV coverage within less than 6 seconds for each patient. All those solutions are initially superior to treatment plans found by DOPSA. Over the running time of 3 minutes, however, DOPSA consistently outperforms IPIP for all patient data sets. In each of the 250 replications for all three patients, the PTV coverage after 3 minutes is higher for DOPSA than for IPIP. The average PTV coverage attained by DOPSA exceeds the results by IPIP after only 18, 6, and 12

seconds, respectively.

MILP requires substantially more time to determine good treatment plans. For 2 out of 3 patients, MILP requires more than 1 minute to determine a treatment plan with nonzero PTV coverage. However, after 3 minutes for Patients 1-2 and 5 minutes for Patient 3 it has found the best plan among the three algorithms.

Unfortunately, the optimal objective values are not known. There is still a large gap between the best solution found by any of the algorithms and the upper bound provided by the MILP solver, and this gap does not decrease substantially after running the solver for a week. DOPSA differs greatly from the existing dose-volume based methods while it finds similar objective values. This could be an indication that much better solutions do not exist.

Figure 8.5 – Comparison of PTV coverage for DVH-based optimizers for Patient 1.

Figure 8.6 – Comparison of PTV coverage for DVH-based optimizers for Patient 2.

Figure 8.7 – Comparison of PTV coverage for DVH-based optimizers for Patient 3.

8.5 Conclusions

The existing dose-volume based optimization methods MILP and IPIP have not been compared before. Our results show that each method has its own advantages. IPIP is faster while MILP determines a better solution.

DOPSA, our new algorithm, is a viable alternative to both. Its main advantage is that it does not require a costly solver. The time it takes to determine a good solution is less than MILP but more than IPIP. Since IPIP returns a single solution quickly but does not keep improving the solution, DOPSA eventually determines better treatment plans. This leads to two observations. First, advances in computing power will reduce the absolute time difference between both methods, making IPIP lose its edge. Second, it is possible to combine the strength of two methods and use the IPIP solution as a starting point for DOPSA.

An interesting observation is that the currently used method IPSA, which was developed over a decade ago, uses SA to optimize a dose-based objective function as a surrogate for dose-volume optimization. Recently it has been shown that the correlation between dose-based objective functions and dose-volume criteria is weak (Holm (2011) and Chapter 7). Consequently, a posteriori manual tuning is required for plans that are optimized with a dose-based objective function to obtain plans that are good with respect to dose-volume criteria. Our results show that SA is capable of directly optimizing on dose-volume criteria. Dose-based methods may not have become the prevalent method in treatment planning systems, if IPSA had been designed to optimize dose-volume criteria.

We have disclosed all information necessary to implement DOPSA. We hope it receives more testing by the community, and will eventually become available to treatment planners at a relatively low price.

Acknowledgments

CHAPTER 9

Dwell time modulation restrictions do not necessarily improve treatment plan quality for prostate HDR brachytherapy implants.

Abstract Inverse planning algorithms for dwell time optimisation in interstitial high-dose-rate (HDR) brachytherapy may produce solutions with large dwell time variations within catheters, which may result in undesirable selective high-dose subvolumes. Extending the dwell time optimisation model with a dwell time modulation restriction (DTMR) that limits dwell time differences between neighboring dwell positions has been suggested to eliminate this problem. DTMRs may additionally reduce the sensitivity for uncertainties in dwell positions that inevitably result from catheter reconstruction errors and afterloader source positioning inaccuracies. This study quantifies the reduction of high-dose subvolumes and the robustness against these uncertainties by applying a DTMR to template-based prostate HDR brachytherapy implants.

Three different DTMRs were consecutively applied to a linear dose-based penalty model (LD) and a dose-volume based model (LDV), both obtained from literature. The models were solved with DTMR levels ranging from no restriction to uniform dwell times within catheters in discrete steps. Uncertainties were simulated on clinical cases using in-house developed software, and dose-volume metrics were calculated in each simulation. For the assessment of high-dose subvolumes, the dose homogeneity index (DHI) and the contiguous dose volume histogram were analysed. Robustness was measured by the improvement of the lowest $D_{90\%}$ of the planning target volume (PTV) observed in the simulations.

For (LD), a DTMR yields an increase in DHI of approximately 30% and reduces the size of the largest high-dose volume by 2 to 5 cc. However, this comes at a cost of a reduction in $D_{90\%}$ of the PTV of 10%, which often implies that it drops below the desired minimum of 100%. For (LDV), none of the DTMRs were able to improve high-dose volume measures. DTMRs were not capable of improving robustness of PTV $D_{90\%}$ against uncertainty in dwell positions for both models.

9.1   Introduction

Interstitial high-dose-rate (HDR) brachytherapy with a high activity ^{192}Ir stepping source using remote afterloading has shown to be an excellent treatment option for localised prostate cancer in any risk category. Its high tumour control and low toxicity rates result from the precision and control with which this highly conformal treatment can be delivered Yamada et al. (2012).

In modern treatment planning systems, automated techniques for anatomy-based inverse treatment planning enable a fast adjustment of the source dwell time distribution within implanted catheters. The fundament of these automated techniques is a mathematical optimisation model that uses dose penalty functions for the planning target volume (PTV) and all relevant organs at risk (OARs) to achieve pre-set dose requirements. Several optimisation algorithms, like the inverse planning simulated annealing (IPSA, Lessard and Pouliot, 2001) and hybrid inverse planning optimisation (HIPO, Karabis et al., 2005) algorithms have been described in the literature to solve this task.

Often these algorithms produce solutions with large dwell time variations within catheters Holm et al. (2012). This may give rise to the following problems related to treatment plan robustness against uncertainties in dose delivery. Firstly, large dwell times may produce selective high-dose subvolumes around dominant dwell positions. As stated by Baltas et al. (2009), these high-dose subvolumes should be avoided unless inhomogeneities in the GTV structure ask for an inhomogeneous dose distribution. Secondly, catheters with large dwell time variations are expected to yield a heterogeneous dose distribution. Since a longitudinal displacement of the catheter implies that a dwell position may shift to the location of its neighbor, heterogeneous dose distributions are expected to be more susceptible to inter- and intra-fraction catheter displacement. Thirdly, it is generally believed that irregular dwell time distributions may be sensitive to uncertainties in dwell positions due to catheter reconstruction errors and (mechanical) source positioning inaccuracies of the afterloader.

9.1.1 Clinical procedure and workflow

This study is based on a clinical procedure for HDR brachytherapy delivered in two fractions of 8.5 Gy with a week interfraction interval. The brachytherapy fractions are a boost to external radiotherapy delivered in 13 fractions of 2.75 Gy. Dose prescriptions are according to Hoskin et al. (2007) and Hoskin et al. (2013). Prior to treatment, the patient is placed in dorsal lithotomy position, and is not moved until the treatment fraction has been completed. First, TRUS images are made and relevant tissue structures are contoured. Second, a pre-plan is made, where the catheter configuration as well dwell times are optimised via inverse planning. For each fraction, a new implant is made using rigid steel catheters. In the third step, catheters are inserted under transrectal ultrasound (TRUS) guidance, using a transperineal template with a hole spacing of 5 mm. Fourth, new TRUS images are made, on which catheters are reconstructed and structure delineations are adapted when necessary. Since we make use of rigid catheters which can be assumed to be unbent, catheter reconstruction is performed by localising two points. The first point is the location of the catheter at the template, which is known, and the second point is the catheter tip, which is identified on the scan. In the fifth phase, a treatment plan is developed based on the delineations and catheter reconstructions using inverse planning. This work only considers the dwell time optimisation in this fifth step, where the catheter configuration is fixed. Finally, the treatment can be delivered.

9.1.2 High-dose subvolumes

Small high-dose subvolumes inevitably occur around each of the dwell positions, where the size of this region depends on the dwell time. When large dwell times occur, the high-dose subvolumes around two neighboring dwell positions may get connected, resulting in a large high-dose volume. Without (radio)biological information about the intra-tumoural heterogeneity, such large high-dose subvolumes may cause irreversible damage to the stromal tissue causing necrosis, and are thus considered an undesirable property of a dose distribution. Therefore, it is reasonable to avoid the formation of such high-dose subvolumes.

The size of the largest high-dose volume for a range of dose levels is reflected in the contiguous dose-volume histogram (DVH^c Thomas et al., 2007).

Besides DVH^c, which is a local measure, we consider the global measure DHI, which is calculated as Wu et al. (1988):

$$\text{DHI} = \frac{V_{100\%} - V_{150\%}}{V_{100\%}}.$$

DHI can be interpreted as the fraction of calculation points receiving a dose between

100% and 150% of the prescribed dose relative to the points receiving at least 100% of the prescribed dose. DHI attains a value between 0 and 1, where a value close to 1 corresponds to a homogeneous dose distribution.

While DVH^c shows the largest volume receiving at least a certain dose, interest may also be in the second or third largest volume. We only take these into account indirectly via the DHI for a dose of 150% of the prescribed dose. In order to see this, note that $1-$ DHI $= V_{150\%}/V_{100\%}$, which is the total volume of all regions receiving at least 150% of the prescribed dose, as a fraction of the total volume receiving at least the prescribed dose.

9.1.3   Uncertainty in dwell locations

A generally acknowledged problem encountered in HDR brachytherapy is the uncertainty in the location of the dwell positions, resulting in a difference between the intended dose distribution and the one actually delivered. Here, we identify three causes. The first cause is a catheter positioning error, which occurs when the catheter is not exactly positioned during implantation as it was planned in the pre-plan Abolhassani et al. (2007). We do not consider this type of error here, since we focus on cases where the catheter positions are given. The second cause is a reconstruction error: it is difficult to determine the exact location of the catheter tip from TRUS imaging data Siebert et al. (2009). This uncertainty translates to an uncertainty in the location of each of the dwell positions within that catheter. The third cause is a source positioning error caused by mechanical inaccuracies of the afterloader, resulting from the resolution of the afterloader as well as uncertainty in the path of the source guide transit tube. The latter may be slightly bent, causing an error in the distance from the afterloader to the source. These inaccuracies result in all dwell positions within one catheter to shift in the same direction with the same magnitude.

When the expected locations of the dwell positions differ from the true locations, the delivered dose distribution differs from the expected dose distribution. Treatment plans should be robust in the sense that the influence of these uncertainties on the dose distribution is small.

9.1.4   Dwell time modulation restrictions

In the literature it has been proposed to regularise the dwell time distribution per catheter by adding a dwell time modulation restriction (DTMR) to the optimisation model Baltas et al. (2009). This is a constraint that puts a restriction on the difference between dwell times of adjacent dwell positions. In Chapter 7, a DTMR restricting the relative dwell time difference is used. Van der Laarse and Prins (1994) considered

the sum of squared differences between dwell times of adjacent dwell locations.

DTMRs are available in commercial treatment planning systems, as for example in the real-time intra-operative planning system Oncentra Prostate (Nucletron B.V., Veenendaal, The Netherlands) that employs the HIPO Karabis et al. (2005) and the IPSA algorithm Lessard and Pouliot (2001); Alterovitz et al. (2006). Here, a user-selectable level between 0 and 1 allows for different degrees of dwell time modulation. Unfortunately, no mathematical definition of the DTMR is given for these particular dwell-time optimisation algorithms, and the interpretation of the various restriction levels remains unclear. Baltas et al. (2009) and Mavroidis et al. (2010) have studied the effects of a DTMR in HIPO, and concluded that including a DTMR results in treatment plans with fewer high-dose subvolumes and lower total dwell time.

9.1.5   Aim of the paper

The aim of this study is to quantify the assumed improvement in treatment plan quality caused by DTMRs in HDR brachytherapy of the prostate. We measure the reduction in high-dose subvolumes caused by three different DTMRs in existing dose- or dose-volume based inverse treatment planning models and are the first to do so in a quantitative manner. Furthermore, this study is the first to assess robustness against uncertainties in dwell locations as a result of these DTMRs. It is our aim to investigate the trend caused by an increasing DTMR level, not to consider treatment quality for a single DTMR level.

9.2   Methods and materials

9.2.1   Dwell time optimisation models

For our analysis, we used two different dwell time optimisation models. The first model is the linear dose (LD) model by Alterovitz et al. (2006), which forms the basis for the IPSA algorithm. This model assigns a penalty to each dose calculation point receiving a dose below the preset lower bound or above the preset upper bound, where the penalty is linear in the difference between the dose and the corresponding bound. The objective is to minimise the total penalty. Recently, new optimisation models have been developed that directly optimise dose-volume histogram (DVH) parameters (Siauw et al. (2011) and Chapter 7). The second model we use is the linear dose volume (LDV) model decribed in Chapter 7. This model maximises the number of dose calculation points in the PTV that receive at least the prescribed dose, while restricting the dose received by the hottest ten percent of the rectum and urethra, denoted by $D_{10\%}(\text{rectum})$ and $D_{10\%}(\text{urethra})$, respectively.

9.2.2 Modulation restrictions

Both the (LD) and (LDV) models contain the dwell time variable t_j, which denotes the dwell time of dwell position j. A restriction can be placed on the dwell time gradient of neighboring dwell positions using these dwell time variables. DTMR-R was introduced in Chapter 7:

$$t_{j_1} \leq (1 + \gamma)t_{j_2} \qquad \forall j_1 \in J, \ \forall j_2 \in \Gamma(j_1), \tag{9.1}$$

where γ denotes the pre-set maximum *relative* difference between two adjacent dwell positions, J denotes the set of dwell positions and $\Gamma(j)$ denotes the set of all dwell positions adjacent to dwell position j. Note that constraint (9.1) also implies the reversed constraint, where the dwell time of dwell position j_2 cannot exceed $(1+\gamma)t_{j_1}$.

Instead of restricting the relative differences as in constraint (9.1) we can restrict the absolute difference between dwell times of two adjacent dwell positions. We introduce DTMR-A, formulated as:

$$t_{j_1} - t_{j_2} \leq \theta \qquad \forall j_1 \in J, \ \forall j_2 \in \Gamma(j_1), \tag{9.2}$$

where θ is the pre-set maximum *absolute* difference between two adjacent dwell positions. Just as constraint (9.1), constraint (9.2) works two ways.

Finally, DTMR-Q is introduced as a modification of the quadratic penalty first described by Van der Laarse and Prins (1994):

$$\frac{1}{2} \sum_{j_1 \in J} \sum_{j_2 \in \Gamma(j_1)} (t_{j_1} - t_{j_2})^2 \leq \rho,$$

where ρ is some pre-set maximum on the sum of squared differences between dwell times of adjacent dwell positions.

9.2.3 Patient data and simulations

In order to investigate the effects of the different DTMRs on the quality of the dose distribution, the three restrictions are applied to the (LD) and (LDV) dwell time optimisation models.

For the numerical evaluation clinical data from three prostate cancer patients were used, where the rectum and urethra are the delineated OARs. These three patients cover various prostate sizes: 32, 55 and 48 cc, respectively. The PTV is defined as the clinical target volume extended with a 2 mm margin. The catheter

positions used have been chosen by an experienced treatment planner (ALH). For patients 1 and 2, 16 catheters were used, and 14 were used for patient 3.

Patient data were obtained from the treatment planning system (HDRplus, version 3.0, Eckert and Ziegler BEBIG GmbH, Berlin, Germany), comprising the sets of dwell positions, catheter positions and dose calculation points, the parameters necessary for optimisation, and the dose rate matrix. We used the same data as in Chapter 7. Dose calculation points had been hexagonally distributed over the delineated structures. For treatment planning, a resolution of approximately 40, 30 and 180 calculation points per cc was used for the PTV, rectum and urethra, respectively. Approximately four times as many dose calculation points were used for evaluation of the resulting dose distribution, in order to increase accuracy. The PTV consists of the prostate extended with a 2 mm margin, according to our clinical protocol. The dose rates were determined according to the TG-43 formalism Nath et al. (1995), with parameters according to Granero et al (2006).

For all three patients different treatment plans were obtained by solving the (LD) and (LDV) models extended with each DTMR for different values for γ, θ and ρ. In the models extended with DTMR-R, $1 + \gamma$ ranges from 1 to 4.6 with an incremental step size of 10%, implying that the relative difference between dwell times of adjacent dwell positions is restricted to be 0 up to 360%. Note that $1 + \gamma = 4.6$ is closest to the unrestricted case, and for $1 + \gamma = 1$ all dwell times within the same catheter are forced to be equal. In the models extended with DTMR-A, θ varies from 0 to 5 in steps of 0.05, where $\theta = 5$ approaches the unrestricted case. In the models extended with DTMR-Q, ρ varies from 0 to the value implying free modulation, taking 50 steps. The optimisation models extended with DTMR-R or DMTR-A were solved using CPLEX 12.2 Optimiser (IBM Corporation, Somers, USA). For (LD) extended with DTMR-Q we obtained more accurate results using the MOSEK 6.0 solver (Mosek ApS, Copenhagen, Denmark) due to its strongly developed interior point method. For model (LDV) we stopped the solver as soon as 95% of the PTV received at least the prescribed dose, or after 30 minutes (on a computer with an Intel Q8400 processor). The (LDV) model extended with DTMR-Q could not be solved within reasonable time, and is thus not included in the analysis.

After the treatment plans had been generated, the actual locations of the catheters were perturbed by means of simulation, resulting in different locations of the dwell positions. Due to the large number of possible scenarios, at least 10,000 simulations were calculated per patient. Simulated locations are based on deviations from the nominal (measured) scenario. The accuracies used for simulation are consistent with values reported in the literature (Pantelis et al., 2004, page 62). For each simulation, the location of the catheter tip was uniformly sampled from a sphere of 2 mm around

the measured position, resulting in a change in angle between the catheter and the template and in catheter depth. Consequently, all dwell positions in that catheter were moved with the catheter. A single longitudinal shift was uniformly sampled on the interval [-1,1] mm, which applies to all dwell positions within the catheter simultaneously. Dislocations were sampled for each catheter separately. For each treatment plan obtained, the resulting objective values and DVH evaluation criteria were calculated for every simulation. In order to assess whether the models provide good and robust treatment plans, the objective value and the DVH criteria were compared for different DTMR parameter values by simulation.

9.2.4   Plan quality indicators

We consider the following performance indicators to assess treatment plan quality. In Table 9.1 the DVH criteria are shown, based on the 8.5 Gy per fraction boost according to the treatment protocol by Hoskin et al. (2007). Here, $V_{100\%}$ (PTV) is referred to as the fraction of the PTV receiving at least 100% of the prescribed dose, while $D_{90\%}$ reflects the minimum dose received by the hottest 90% of the PTV as a fraction of the prescribed dose.

Table 9.1 – Dose-volume criteria, based on the protocol by Hoskin et al. (2007).

PTV	Rectum	Urethra
$D_{90\%} \geq 100\%$	$D_{10\%} \leq 7.2$ Gy	$D_{10\%} \leq 10$ Gy
$V_{100\%} \geq 95\%$	$D_{2\text{cc}} \leq 6.7$ Gy	$D_{0.1\text{cc}} \leq 10$ Gy
$V_{150\%} \leq 55\%$		
$V_{200\%} \leq 20\%$		

For the assessment of robustness against uncertainty in catheter positions, the objective value is a useful indicator. Ideally, a robust model does not deteriorate the objective value, while it improves the worst case value in the simulations. Furthermore, it decreases the standard deviation among simulations.

The objective of (LD) is only a surrogate for the actual goal, which is to satisfy the preset DVH criteria as well as possible (Holm (2011) and Chapter 7). Therefore, we also consider the results for $D_{90\%}$ of the PTV. For the (LDV) model, we consider the objective ($V_{100\%}$) as well as $D_{90\%}$.

9.3   Results

The objective values, DVH metrics and high-dose volume measures resulting from the different treatment plans applied to the simulations are summarised in graphs. Since the number of graphs is huge, we only discuss graphs for patient 1 for (LD) extended with DTMR-R as an illustration. The remaining graphs have been included as supplementary data.

In all graphs except for those concerning DVH^c, the modulation parameter is shown on the horizontal axis. The smallest value implies a strong DTMR, while the largest value implies free or almost free modulation. Consequently, the two extreme observations are the same for all modulation restrictions. For example, it does not matter whether DTMR-A, DTMR-R or DTMR-Q is used when the best plan is the one with the strongest DTMR, because for all DTMRs a parameter value exists that forces all dwell times to be equal within the same catheter. In contrast, a specific DTMR is favourable if the best plan is found for non-extreme DTMR parameter values, *i.e.*, in the middle of the graph. Note that, since (LDV) could not always be solved to optimality, the results for the extreme modulation restrictions are not exactly the same.

In the figures displaying dose-volume parameters, the three solid lines represent the average, the maximum and the minimum value over all simulations, and the dotted curve represents the value when all locations of the dwell positions are as derived from the imaging data (*i.e.*, nominal value). The grey area denotes the values within a distance of one standard deviation from the mean.

By way of example, we briefly describe the interpretation of Figure 9.1b. From this figure we see that for the strongest modulation restriction, *i.e.*, when $1 + \gamma = 1$, the minimum, mean and maximum $D_{90\%}(\text{PTV})$ over all simulations are approximately 99.9%, 105% and 109%, respectively. The $D_{90\%}(\text{PTV})$ is 106% if the dwell locations are exactly as in the nominal scenario, which can be seen from the dotted line representing the nominal value. When the treatment plan developed with (LD) and DTMR-R with $1 + \gamma = 1$ is used, the values for $D_{90\%}(\text{PTV})$ that deviate at most one standard deviation from the mean lie between 104% and 106%. When moving along the x-axis, the $D_{90\%}(\text{PTV})$ for less strict DTMRs can be found, until we reach the values for the almost unrestricted case at $1 + \gamma = 4.6$.

As opposed to DHI and dose-volume metrics, DVH^c is not a single value, but a complete graph for every DTMR level. Therefore, a 3D graph shows the DVH^c for various DTMR levels.

The most prominent aspect of all graphs for (LDV) is the large variation of each curve. The main cause is that we did not solve (LDV) to optimality. Therefore, we should consider the trend, rather than the individual values.

9.3.1   Relative dwell time difference restricted in (LD)

The results obtained from (LD) extended with the relative dwell time difference constraint (9.1) are discussed in this subsection. First, the results concerning robustness are discussed, followed by high-dose volume measures. Note that we are not interested in the absolute values of the quality indicators, but merely in the improvements as a result of applying a DTMR. Hence, we will only consider deviations in parameter values caused by applying a stronger restriction.

From Figure 9.1a, it is evident that the range of penalties among the simulations becomes smaller when the DTMR gets stronger. Furthermore, the penalty for the worst case scenario becomes smaller. From these two observations, we can conclude that adding a strong relative dwell time modulation restriction to the (LD) model yields solutions with robust penalty values. However, the results in Figure 9.1b show that a strong modulation restriction yields a decrease in the $D_{90\%}(\text{PTV})$. The decrease in the $D_{90\%}(\text{PTV})$ for patient 1 does not result in insufficient target coverage, but for patients 2 and 3 it drops below the minimum desired level of 100% (Figures 9.7b and 9.7c of the supplementary data). The standard deviation decreases slightly, but the $D_{90\%}(\text{PTV})$ for worst case scenario strongly decreases. We can thus conclude that DTMR-R does not yield robustness of the $D_{90\%}$ for the PTV, and leads to compromised target coverage.

Figure 9.1 – Graphs for patient 1 showing: (a) penalty and (b) $D_{90\%}(\text{PTV})$ generated by (LD) extended with DTMR-R. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

The results for the remaining DVH parameters for PTV as well as those for the OARs can be found in 9.A of the supplementary data. From these figures, we conclude that a strong restriction on the relative dwell time differences results in a

large reduction of $V_{100\%}$, and a slight reduction of $V_{150\%}$ and $V_{200\%}$. Also rectum D_{2cc} and $D_{10\%}$, whereas urethra $D_{0.1cc}$ and $D_{10\%}$ stay at the same level.

High-dose volumes are the second feature under consideration. Figure 9.2a shows that for patient 1, DHI is optimal for the weakest as well as strongest DTMR. For patients 2 and 3, the highest DHI is obtained for a strong modulation restriction (Figures 9.4b and 9.4c). From Figure 9.2b, one can see that the strongest DTMR yields the best DVH^c. The same conclusion can be drawn for patients 2 and 3 (Figures 9.5b and 9.5c). Note that this is a trivial consequence from the decrease in $D_{90\%}$.

(a)   **(b)**

Figure 9.2 – Graphs for patient 1 showing: high-dose volume performance indicators (a) DHI and (b) DVH^c generated by (LD) extended with DTMR-R. The solid lines in graph (a) represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

From the above we conclude that including DMTR-R in (LD) can result in undesirable treatment plans when considering the $D_{90\%}$ for the PTV. Furthermore, it does result in robustness of the penalty, but not in robustness of the $D_{90\%}$ for the PTV. A strong DMTR-R slightly improves high-dose subvolumes according to the DHI and DVH^c, but this is a trivial consequence of an undesirably low $D_{90\%}$.

9.3.2 General results

For (LD) extended with DTMR-A or DTMR-Q, the results are similar to those obtained from (LD) extended with DTMR-R and do not need a separate discussion. For the (LDV) model extended with DTMR-R and DTMR-A the DTMRs did not

provide treatment plans with improved high-dose volume indicator values, and sometimes even worse high-dose volume indicator values were found. Furthermore, our DTMRs often yield lower values for the $D_{90\%}$ of the PTV. Detailed results can be found in the supplementary data.

9.4   Discussion

The effects of DTMRs were assessed in several other papers as well. In their work, Baltas et al. (2009) and Mavroidis et al. (2010) used data from 12 clinical implants in combination with the HIPO algorithm. The dosimetric quality of these implants was assessed for plans optimised with and without DTMR, such that DVH parameters of their protocol were completely fulfilled. This resulted in almost equivalent DVHs for the prostate and a more pronounced sparing of the OARs, especially urethra and bladder. However, our work does not support this conclusion. We observed a decrease in target coverage when using a DTMR.

Baltas et al. (2009) and Mavroidis et al. (2010) have found a lower mean dwell time per implant and a mean total dwell time reduction of 1.4%, that both were proven to be statistically significant. In our opinion, part of the reduction in total dwell time for the plan with DTMR can be explained by the average observed reduction in $D_{90\%}$ for the prostate. Using (LD) extended with one of the three DTMRs, we observed a mean total dwell time reduction of 3-5%. For DTMR-Q the reduction takes place only when dwell times are forced to be equal within the same catheter, while a more gradual reduction was observed for the other two DTMRs as the restriction became stronger. For (LDV) there were no differences in mean total dwell time between the plans with and without DTMR.

Both Baltas et al. (2009) and Mavroidis et al. (2010) use radiobiological models to assess the expected treatment outcome in terms of the probabilities of benefit, injury and uncomplicated tumour control. The results of this radiobiological evaluation supported to a large extend the conclusions derived from the dosimetric comparisons between plans with and without DTMRs. They concluded that optimisation with a DTMR can introduce a minor improvement in the effectiveness of the dose distribution obtained compared to the optimisation without DTMR. We question whether the parameter values of the radiobiological models, which have been derived from 3D dose distributions delivered with external beam radiotherapy, can be applied to highly heterogeneous dose distributions delivered with HDR brachytherapy, and thus whether they give usable results. Furthermore, the parameters of the radiobiological models are uncertain, and the effect of this uncertainty is not assessed in these papers. Moreover, in our opinion there is a lack of statistical evidence in assessing the

significance of the difference in radiobiological indices, and it is unclear for how many patients the DTMR gives an improvement in any of the radiobiological indices.

The results of the current study show that the DTMR is not suitable for the development of robust treatment plans. Therefore, these uncertainties need to be included in the optimisation model in a different way. Examples of methods to take uncertainties into account are robust optimisation Ben-Tal et al. (2009a) and stochastic programming Kall and Wallace (1994).

High-dose volume parameters do not show improvements as a result of the DTMR for the (LDV) model. For the (LD) model, the DTMR can only marginally improve high-dose volume measures while compromising target coverage. A trade-off between these two goals thus needs to be made. For patient 1, the decrease in coverage is not problematic, but for patients 2 and 3 it drops below the desired minimum. The effects of a DTMR do not only differ among patients, but also among models, which indicates that the different effects may not depend on the implant quality and patient characteristics.

A method to improve the homogeneity of the dose distribution is to include the DHI in the objective, as proposed by Holm et al. (2013a). When comparing the newly developed model to the original linear penalty model by Alterovitz et al. (2006), Holm et al. (2013a) found no significant difference in treatment plan quality when comparing the homogeneity index.

For the assessment of a DTMR's capabilities of improving robustness against uncertainties in dwell locations, we took catheter reconstruction errors and the source positioning inaccuracy of the afterloader into account. According to Pantelis et al. (2004), such errors hardly affect the dose-volume histogram and corresponding DVH metrics. However, our results show a variation in $D_{90\%}$ and $V_{100\%}$ that can be up to 5%, which corresponds to the magnitude of variation reported by Kirisits et al. (2014), Table 5. Pantelis et al. (2004) derive DVH metrics from the average dose to a single calculation point over all sampled catheter locations instead of deriving DVH metrics for a single sample, and as such ignore the risks occurring in individual scenarios.

Additional uncertainties arise in various stages of the treatment process. In the pre-planning stage, where catheter locations are chosen, uncertainties caused by inter- and intra observer variability, catheter deflection during insertion and organ deformations caused by catheter insertion arise. In addition, uncertainties arise in dose calculations and after catheters have been inserted (Kirisits et al., 2014, page 207 and Table 5). According to Kirisits et al. (2014), the magnitude of the effects on dose distribution of these uncertainties are of a similar magnitude as the magnitude of the errors considered in this work. The effects of uncertainties other than those

considered in this article, such as inter- and intra observer variabilities and inter fractional organ deformations, as well as methods to overcome them are interesting topics for future research, especially for different catheter configurations.

In our study, we only considered the case where rigid catheters are used. However, flexible catheters are used in clinics and hospitals where the implant remains in situ between fractions. When performing a similar study with flexible catheters, indicating two points on a catheter is insufficient for identification of the complete catheter location, and uncertainties should be simulated in a different way.

Since we have considered only three patients, we were able to perform the detailed study that was necessary to address the effects of a DTMR on high-dose subvolumes and robustness. The results for all three patients were negative, and DTMRs sometimes even showed to result in deterioration of the plan quality. Including more patients in this study would not change the negative results for these three patients. Therefore, despite the small number of patients, we can conclude that DTMRs do not in general improve treatment plan quality, as they were expected to. This does not imply that a DTMR has negative effects for every patient: patients for whom a DTMR would have positive effects may exist. In clinical practice, it is thus desirable to conduct tests to investigate whether a plan generated with a DTMR would yield fewer high-dose volumes and more robustness than a plan generated without a DTMR for a specific patient. However, such a lengthy study would take up too much time during the treatment.

9.5   Conclusion

Robustness in the penalty of the (LD) model is obtained from all three DTMRs, however the goal is to achieve robustness in the $D_{90\%}$ for the PTV. No improvement in robustness of the $D_{90\%}$ against uncertainty in dwell position measurement and afterloader precision was obtained by applying any of the three DTMRs for the (LD) and (LDV) models. Furthermore, these DTMRs do not reduce the high-dose subvolumes without simultaneously deteriorating $D_{90\%}$. Finally, hardly any sparing of the OARs is achieved, unless the dose delivered at the PTV is decreased as well. These conclusions hold for the patients from our institution, though different results may be obtained in other institutions. Therefore, we recommend other institutions to perform a similar study in order to see how their treatment plans are influenced by the use of DTMRs.

Acknowledgment

9.A   Relative dwell time difference restricted

9.A.1   (LD) model

Figure 9.3 – Objective function value for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.4 – DHI for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.
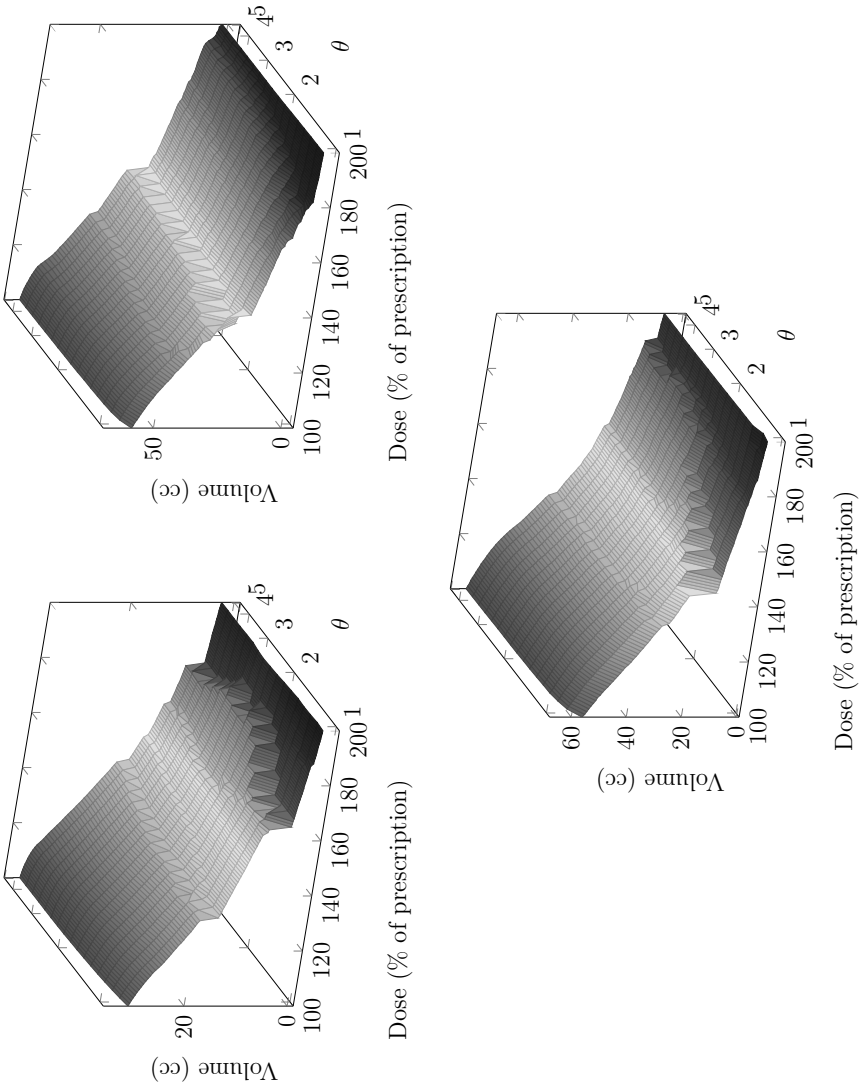
Figure 9.5 – PTV DVH^c for all patients.

Figure 9.6 – $D_{90\%}$(PTV) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.7 – $V_{100\%}$(PTV) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.8 – $V_{150\%}(\text{PTV})$ for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.9 – $V_{200\%}(\text{PTV})$ for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.
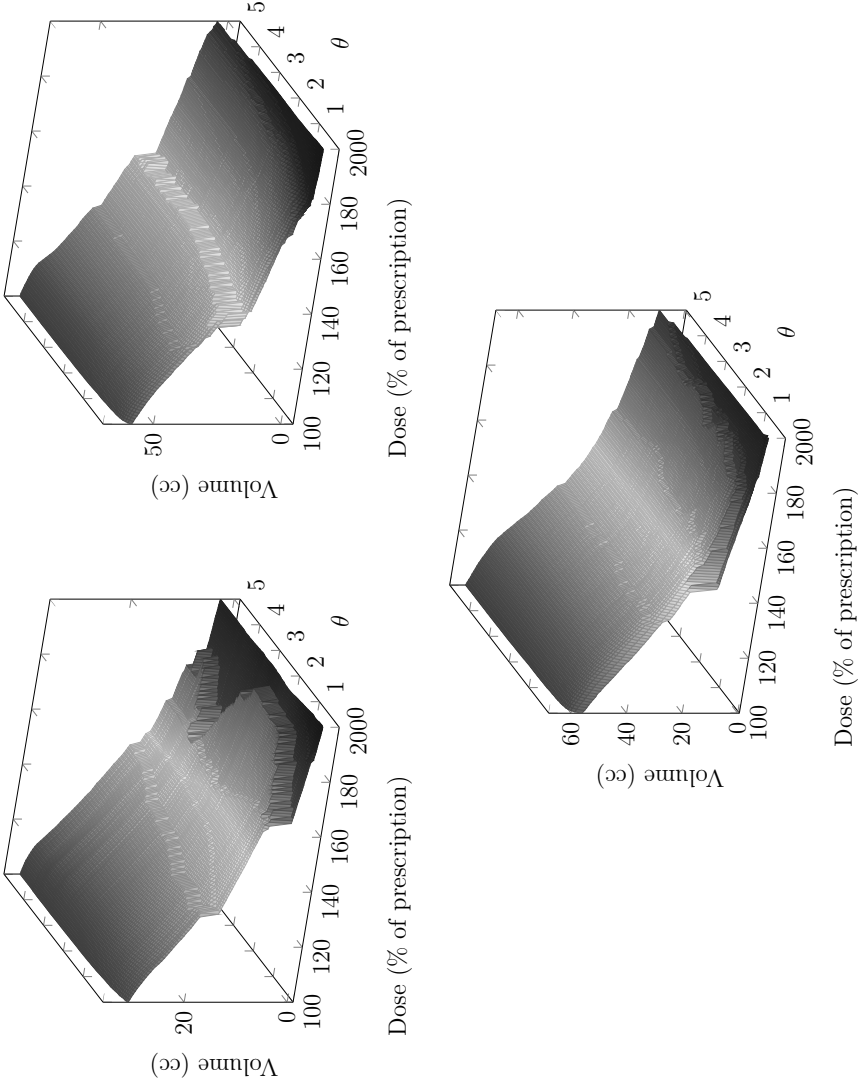
Figure 9.10 – $D_{10\%}$(rectum) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.11 – D_{2cc}(rectum) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.12 – $D_{10\%}$(urethra) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.13 – $D_{0.1cc}$(urethra) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

9.A.2 (LDV) model

Figure 9.14 – $V_{100\%}$(PTV) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.15 – DHI for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.
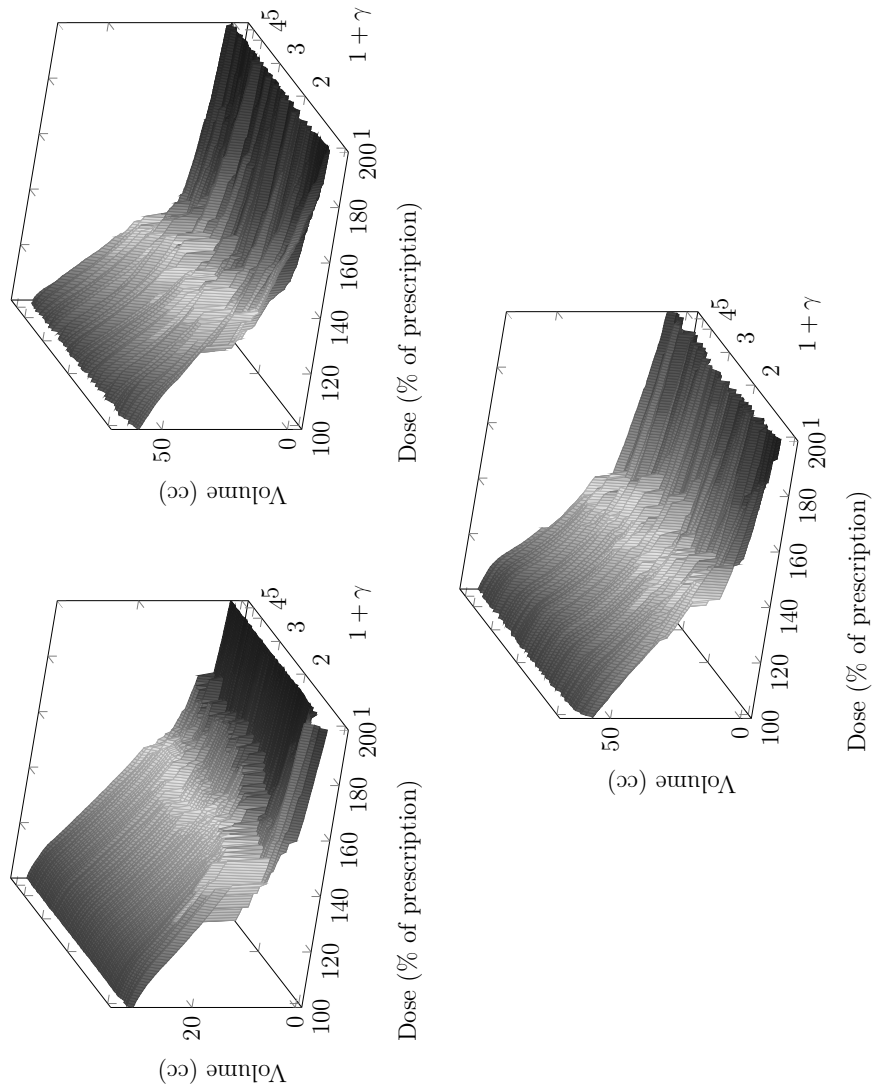
Figure 9.16 – PTV DVH^c for all patients.

Figure 9.17 – $D_{90\%}$(PTV) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.18 – $V_{150\%}$(PTV) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.19 – $V_{200\%}(\text{PTV})$ for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.20 – $D_{10\%}(\text{rectum})$ for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.21 – D_{2cc}(rectum) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.22 – $D_{10\%}$(urethra) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.23 – $D_{0.1cc}$(urethra) for all patients as function of the maximum relative dwell time difference γ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

9.B Absolute dwell time difference restricted

9.B.1 (LD) model

Figure 9.24 – Objective function value for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.25 – DHI for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.
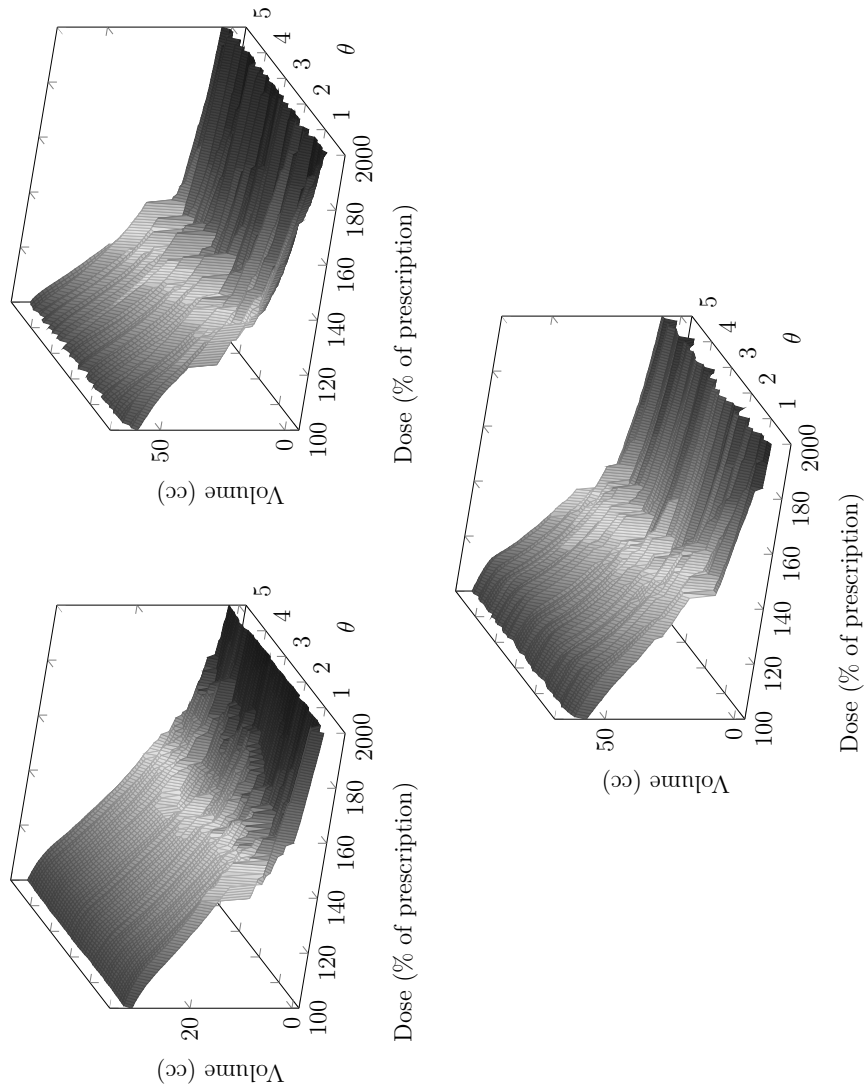
Figure 9.26 – PTV DVH^c for all patients.

Figure 9.27 – $D_{90\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.28 – $V_{100\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.29 – $V_{150\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.30 – $V_{200\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.31 – $D_{10\%}$(rectum) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.32 – D_{2cc}(rectum) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.33 – $D_{10\%}$(urethra) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.34 – $D_{0.1cc}$(urethra) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

9.B.2 (LDV) model

Figure 9.35 – $V_{100\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.36 – DHI for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.37 – PTV DVH^c for all patients.

Figure 9.38 – $D_{90\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.39 – $V_{150\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.40 – $V_{200\%}$(PTV) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.41 – $D_{10\%}$(rectum) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.42 – D_{2cc}(rectum) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.43 – $D_{10\%}$(urethra) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.44 – $D_{0.1cc}$(urethra) for all patients as function of the absolute dwell time difference θ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

9.C   Quadratic dwell time difference restricted

9.C.1   (LD) model

Figure 9.45 – Objective function value for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.46 – DHI for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.
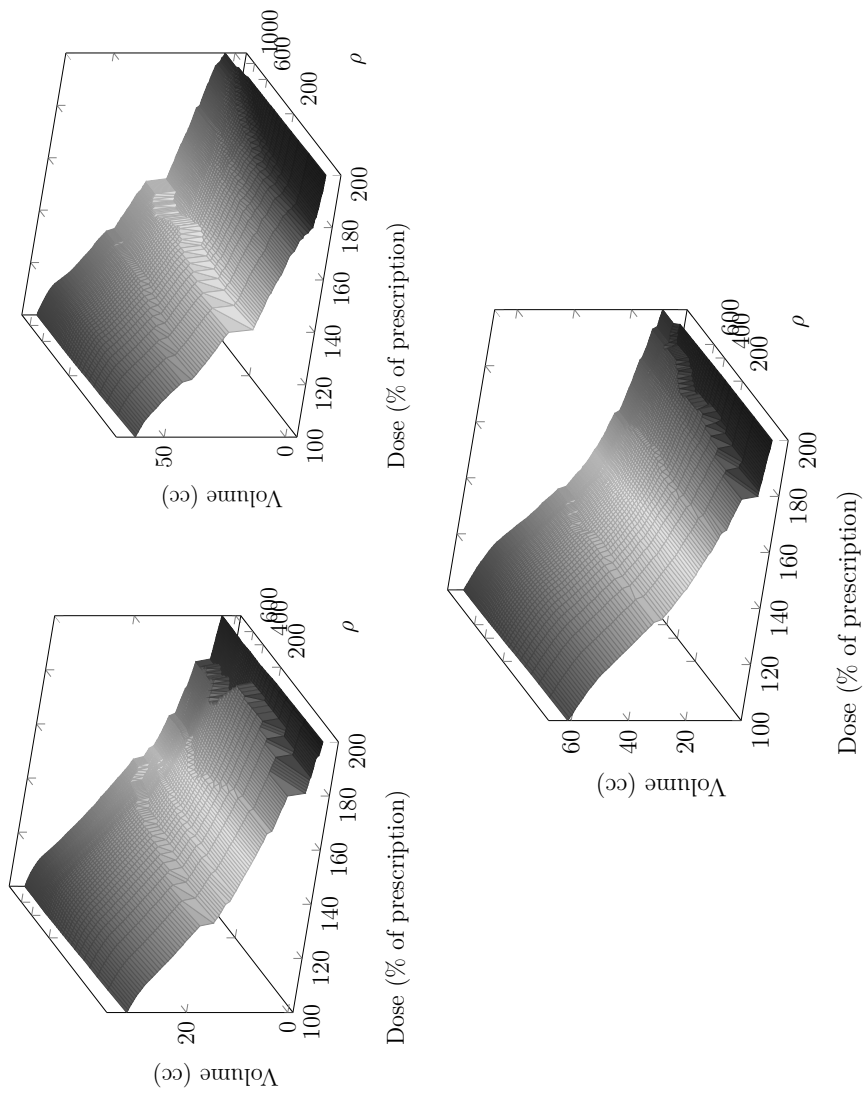
Figure 9.47 – PTV DVH^c for all patients.

Figure 9.48 – $D_{90\%}$(PTV) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.


Figure 9.49 – $V_{100\%}$(PTV) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.50 – $V_{150\%}$(PTV) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.51 – $V_{200\%}$(PTV) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.52 – $D_{10\%}$(rectum) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.53 – D_{2cc}(rectum) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.54 – $D_{10\%}$(urethra) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Figure 9.55 – $D_{0.1cc}$(urethra) for all patients as function of the maximum sum of squared differences ρ. The solid lines represent minimum, mean and maximum values, and the dotted line is the pre-plan value. The grey area denotes values at most one standard deviation from the mean.

Bibliography

N. Abolhassani, R. Patel, and M. Moallem. Needle insertion into soft tissue: a survey. *Medical Engineering and Physics*, 29(4):413–431, 2007.

D. J. Alem and R. Morabito. Production planning in furniture settings via robust optimization. *Computers & Operations Research*, 39(2):139–150, 2012.

R. Alterovitz, E. Lessard, J. Pouliot, I. J. Hsu, J. F. O'Brien, and K. Goldberg. Optimization of HDR brachytherapy dose distributions using linear programming with penalty costs. *Medical Physics*, 33(11):4012–4019, 2006.

D. Baltas, C. Kolotas, K. Geramani, R. F. Mould, G. Ioannidis, M. Kekchidi, and N. Zamboglou. A conformal index (COIN) to evaluate implant quality and dose specification in brachytherapy. *International Journal of Radiation Oncology * Biology * Physics*, 40(2):515–524, 1998.

D. Baltas, Z. Katsilieri, V. Kefala, S. Papaioannou, A. Karabis, P. Mavroidis, and N. Zamboglou. Influence of modulation restriction in inverse optimization with HIPO of prostate implants on plan quality: Analysis using dosimetric and radiobiological indices. *IFMBE Proceedings*, 25(1):283–286, 2009.

C. Bandi and D. Bertsimas. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming*, 134(1):23–70, 2012.

A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh. Clustering with Bregman divergences. *The Journal of Machine Learning Research*, 6:1705–1749, 2005.

A. I. Barros, J. B. G. Frenk, S. Schaible, and S. Zhang. A new algorithm for generalized fractional programs. *Mathematical Programming*, 72(2):147–175, 1996.

A. Beck and A. Ben-Tal. Duality in robust optimization: primal worst equals dual best. *Operations Research Letters*, 37(1):1–6, 2009.

J. Beliën, J. Colpaert, and L. de Boeck. A hybrid simulated annealing linear programming approach for treatment planning in HDR brachytherapy with dose volume constraints. In *Proceedings of the 35 International Conference on Operational Research Applied to Health Services (ORAHS)*, 2009.

C. J. P. Bélisle, H. E. Romeijn, and R. L. Smith. Hit-and-run algorithms for generating multivariate distributions. *Mathematics of Operations Research*, 18(2):255–266, 1993.

A. Ben-Tal and D. den Hertog. Hidden conic quadratic representation of some non-convex quadratic optimization problems. *Mathematical Programming*, 143(1-2): 1–29, 2014.

A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–13, 1999.

A. Ben-Tal and A. Nemirovski. Robust optimization–methodology and applications. *Mathematical Programming*, 92(3):453–480, 2002.

A. Ben-Tal and A. Nemirovski. Selected topics in robust convex optimization. *Mathematical Programming*, 112(1):125–158, 2008.

A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

A. Ben-Tal, B. Golany, A. Nemirovski, and J.-Ph. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.

A. Ben-Tal, S. Boyd, and A. Nemirovski. Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Mathematical Programming*, 107(1):63–89, 2006.

A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, 2009a.

A. Ben-Tal, B. Golany, and S. Shtern. Robust multi-echelon multi-period inventory control. *European Journal of Operations Research*, 199(3):922–935, 2009b.

A. Ben-Tal, B. D. Chung, S. R. Mandala, and T. Yao. Robust optimization for emergency logistics planning: Risk mitigation in humanitarian relief supply chains. *Transportation Research Part B: Methodological*, 45(8):1177–1189, 2011.

A. Ben-Tal, D. den Hertog, A. M. B. De Waegenaere, B. Melenberg, and G. Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

A. Ben-Tal, D. den Hertog, and J.-Ph. Vial. Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming*, Online First, 2014.

J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962.

D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control*, 55(12):2751–2766, 2010.

D. Bertsimas and A. Georghiou. Design of near optimal decision rules in multistage

adaptive mixed-integer optimization. *Optimization Online*, Sept, 2013.

D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

D. Bertsimas and A. Thiele. A robust optimization approach to supply chain management. In D. Bienstock and G. Nemhauser, editors, *Integer Programming and Combinatorial Optimization*, volume 3064 of *Lecture Notes in Computer Science*, 145–156. Springer Berlin / Heidelberg, 2004.

D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, 2006.

D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

D. Bertsimas, D. A. Iancu, and P. A. Parrilo. A hierarchy of near-optimal policies for multi-stage adaptive optimization. *IEEE Automatic Control*, 56(12):2809–2824, 2012.

D. Bienstock and N. Özbay. Computing robust basestock levels. *Discrete Optimization*, 5(2):389–414, 2008.

J. R. Birge and F. V. Louveaux. *Introduction to Stochastic Programming*. Springer, 2011.

H. Blanc and D. den Hertog. On Markov chains with uncertain data. *CentER Discussion Paper*, 2008(050), 2008.

C. Bohle, S. Maturana, and J. Vera. A robust optimization approach to wine grape harvesting scheduling. *European Journal of Operations Research*, 200(1):245–252, 2010.

J. Branke, K. Deb, K. Miettinen, and R. Slowinski. *Multiobjective Optimization: Interactive and evolutionary approaches*, volume 5252 of *Lecture Notes in Computer Science*. Springer, 2008.

L. M. Bregman. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.

W. M. Butler, R. R. Stewart, and G. S. Merrick. A detailed radiobiological and dosimetric analysis of biochemical outcomes in a case-control study of permanent prostate brachytherapy patients. *Medical Physics*, 36(3):776–787, 2009.

E. Candes, M. Rudelson, T. Tao, and R. Vershynin. Error correction via linear programming. In *Proceedings of the 46th Annual IEEE Symposium on FOCS*, 295–308, 2005.

A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3–4):181–186, 1962.

A. Charnes, W. W. Cooper, and E. Rhodes. Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2(6):429–444, 1978.

H. J. Chen, S. Schaible, and R. L. Sheu. Generic algorithm for generalized fractional programming. *Journal of Optimization Theory and Applications*, 141(1):93–105, 2009.

X. Chen and Y. Zhang. Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research*, 57(6):1469–1482, 2009.

X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357, 2008.

I. Cohen, B. Golany, and A. Shtub. The stochastic time-cost tradeoff problem: A robust optimization approach. *Networks*, 49(2):175–188, 2007.

D. Craft and T. Bortfeld. How many plans are needed in an imrt multi-objective plan database? *Physics in Medicine and Biology*, 53(11):2785–2796, 2008.

J.-P. Crouzeix and J. A. Ferland. Algorithms for generalized fractional programming. *Mathematical Programming*, 52(1–3):191–207, 1991.

J.-P. Crouzeix, J. A. Ferland, and S. Schaible. An algorithm for generalized fractional programs. *Journal of Optimization Theory and Applications*, 47(1):35–49, 1985.

J.-P. Crouzeix, J. A. Ferland, and S. Schaible. A note on an algorithm for generalized fractional programs. *Journal of Optimization Theory and Applications*, 50(1):183–187, 1986.

G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ, 1963.

L. De Boeck, J. Beliën, and W. Egyed. Dose optimization in high-dose-rate brachytherapy: A literature review of quantitative models from 1990 to 2010. *Operations Research for Health Care*, 3(2):80–90, 2014.

F. de Ruiter. Adjustable robust optimization based on inexact information. Master's thesis, Tilburg University, the Netherlands, 2013.

T. M. Deist. Fast inverse planning for HDR prostate brachytherapy. Master's thesis, Tilburg University, Tilburg, The Netherlands, 2013.

C. L. Deufel and K. M. Furutani. Quality assurance for high dose rate brachytherapy treatment planning optimization: using a simple optimization to verify a complex optimization. *Physics in Medicine and Biology*, 59(3):525–540, 2014.

W. Dinkelbach. On nonlinear fractional programming. *Management Science*, 13(7):492–498, 1967.

D. Donoho. For most large underdetermined systems of equations, the minimal l(1)-norm near-solution approximates the sparsest near-solution. *Communications on Pure and Applied Mathematics*, 59(7):907–934, 2006.

W. D. D'Souza, H. D. Thames, and D. A. Kuban. Dose–volume conundrum for response of prostate cancer to brachytherapy: summary dosimetric measures and their relationship to tumor control probability. *International Journal of Radiation Oncology * Biology * Physics*, 58(5):1540–1548, 2004.

M. Ehrgott. *Multicriteria Optimization*. Lecture Notes in Economics and Mathematical Systems. Springer, 2005.

L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

A. Emrouznejad, B. R. Parker, and G. Tavares. Evaluation of research in efficiency and productivity: A survey and analysis of the first 30 years of scholarly literature in DEA. *Socio-Economic Planning Sciences*, 42(3):151–157, 2008.

J. E. Falk. A linear max–min problem. *Mathematical Programming*, 5(1):169–188, 1973.

J. E. Falk. Exact solutions of inexact linear programs. *Operations Research*, 24(4):783–787, 1976.

E. F. Fama and K. R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.

A. Fredriksson, A. Forsgren, and B. Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Medical Physics*, 38(3):1672–1684, 2011.

R. J. Gallagher and E. K. Lee. Mixed integer programming optimization models for brachytherapy treatment planning. *Proceedings of the AMIA Annual Fall Symposium*, 278–282, 1997.

D. Georg, J. Hopfgartner, J. Gòra, P. Kuess, G. Kragl, D. Berger, N. Hegazy, G. Goldner, and P. Georg. Dosimetric considerations to determine the optimal technique for localized prostate cancer among externalăphoton, proton, or carbon-ion therapy and high-dose-rate or low-dose-rate brachytherapy. *International Journal of Radiation Oncology * Biology * Physics*, 88(3):715–722, 2014.

M. Ghilezan. Role of high dose rate brachytherapy in the treatment of prostate cancer. *Cancer/Radiothérapie*, 16(5–6):418–422, 2012.

T. Goel, R. Vaidyanathan, R. Haftka, W. Shyy, N. Queipo, and K. Tucker. Response surface approximation of Pareto optimal front in multi-objective optimization. *Computer Methods in Applied Mechanics and Engineering*, 196(4-6):879–893, 2007.

D. Goldfarb and G. Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38, 2003.

D. Granero, J. Perez-Calatayud, E. Casal, F. Ballester, and J. Venselaar. A dosimetric study on the ir-192 high dose rate flexisource. *Medical Physics*, 33(12): 4578–4582, 2006.

M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. http://cvxr.com/cvx, Aug. 2010.

C. Gregory, K. Darby-Dowman, and G. Mitra. Robust optimization and portfolio selection: The cost of robustness. *European Journal of Operational Research*, 212 (2):417–428, 2011.

S. T. Hackman and U. Passy. Maximizing a linear fractional function on a pareto efficient frontier. *Journal of Optimization Theory and Applications*, 113(1):83–103, 2002.

M. J. Hadjiyiannis, P. J. Goulart, and D. Kuhn. A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. In *2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 7386–7391. IEEE, 2011.

A.-R. Hedar and M. Fukushima. Derivative-free filter simulated annealing method for constrained continuous global optimization. *Journal of Global Optimization*, 35(4):521–549, 2006.

Å. Holm. *Dose plan optimization in HDR brachytherapy using penalties properties and extensions*. PhD thesis, Linköping University, Linköping, Sweden, 2011.

Å. Holm. A tailored branch-and-bound method for optimizing the dwelling time pattern and catheter positioning in HDR brachytherapy. *Preprint*, 1–10, 2013.

Å. Holm, T. Larsson, and Å. Tedgren. Impact of using linear optimization models in dose planning for HDR brachytherapy. *Medical Physics*, 39(2):1021–1028, 2012.

Å. Holm, T. Larsson, and Å. C. Tedgren. A linear programming model for optimizing HDR brachytherapy dose distributions with respect to mean dose in the DVH-tail. *Medical Physics*, 40(8):081705, 2013a.

Å. Holm, Å. C. Tedgren, and T. Larsson. Heuristics for integrated optimization of catheter positioning and dwell time distribution in prostate HDR brachytherapy. *Annals of Operations Research*, 208(1):1–21, 2013b.

P. J. Hoskin, K. Motohashi, P. Bownes, L. Bryant, and P. Ostler. High dose rate brachytherapy in combination with external beam radiotherapy in the radical treatment of prostate cancer: initial results of a randomised phase three trial. *Radiotherapy and Oncology*, 84(2):114–120, 2007.

P. J. Hoskin, A. Colombo, A. Henry, P. Niehoff, T. P. Hellebust, F.-A. Siebert, and G. Kovács. GEC/ESTRO recommendations on high dose rate afterloading brachytherapy for localised prostate cancer: An update. *Radiotherapy and Oncol-*

ogy, 107(3):325–332, 2013.

D. Iancu and N. Trichakis. Pareto efficiency in robust optimization. *Management Science*, 60(1):130–147, 2014.

V. Jeyakumar, G. Y. Li, and S. Srisatkunarajah. Strong duality for robust minimax fractional programming problems. *European Journal of Operational Research*, 228 (2):331–336, 2013.

P. Kall and S. W. Wallace. *Stochastic programming*. Wiley Interscience Series in Systems and Optimization. Wiley, 1994.

R. Kannan and H. Narayanan. Random walks on polytopes and an affine interior point method for linear programming. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC, 561–570, 2009.

A. Karabis, S. Giannouli, and D. Baltas. HIPO: A hybrid inverse treatment planning optimization algorithm in HDR brachytherapy. *Radiotherapy and Oncolology*, 76: S29, 2005.

A. Karabis, P. Belotti, and D. Baltas. Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming. In *World Congress on Medical Physics and Biomedical Engineering*, volume 25 of *Germany IFMBE Proceedings*, 612–615, 2009.

R. N. Kaul, S. Kaur, and V. Lyall. Duality in inexact fractional programming with set-inclusive constraints. *Journal of Optimization Theory and Applications*, 50(2): 279–288, 1986.

C. Kirisits, M. J. Rivard, D. Baltas, F. Ballester, M. D. Brabandere, R. van der Laarse, Y. Niatsetski, P. Papagiannis, T. P. Hellebust, J. Perez-Calatayud, K. Tanderup, J. L. Venselaar, and F.-A. Siebert. Review of clinical brachytherapy uncertainties: Analysis guidelines of GEC-ESTRO and the AAPM. *Radiotherapy and Oncology*, 110(1):199–212, 2014.

S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

G. Kovács, R. Pötter, T. Loch, J. Hammer, I.-K. Kolkman-Deurloo, J. J. M. C. H. De la Rosette, and H. Bertermann. GEC/ESTRO-EAU recommendations on temporary brachytherapy using stepping sources for localised prostate cancer. *Radiotherapy and Oncology*, 74(2):137–148, 2005.

E. Kropat and G.-W. Weber. Robust regression analysis for gene-environment and eco-finance networks. 2008. Unpublished results.

D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

M. Lahanas and D. Baltas. Are dose calculations during dose optimization in brachytherapy necessary? *Medical Physics*, 30(9):2368–2375, 2003.

M. Lahanas, D. Baltas, and S. Giannouli. Global convergence analysis of fast multiobjective gradient-based dose optimization algorithms for high-dose-rate brachytherapy. *Physics in Medicine and Biology*, 48(5):599–617, 2003a.

M. Lahanas, D. Baltas, and N. Zamboglou. A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy. *Physics in Medicine and Biology*, 48(3):399–415, 2003b.

M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In M. Putinar and S. Sullivant, editors, *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, 157–270. Springer New York, 2009.

E. K. Lee and M. Zaider. Mixed integer programming approaches to treatment planning for brachytherapy - application to permanent prostate implants. *Annals of Operations Research*, 119(1-4):147–163, 2003.

E. Lessard and J. Pouliot. Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Medical Physics*, 28(5):773–779, 2001.

Y. Li. Two conditions for equivalence of 0-norm solution and 1-norm solution in sparse representation. *Neural Networks*, 21(7):1189–1196, 2010.

J. Lin and R. Sheu. Modified Dinkelbach-type algorithm for generalized fractional programs with infinitely many ratios. *Journal of Optimization Theory and Applications*, 126(2):323–343, 2005.

M. S. Lobo. *Robust and convex optimization with applications in finance*. PhD thesis, Stanford University, Stanford CA, United States, 2000.

M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284(1):193–228, 1998.

J. Löfberg. Automatic robust convex programming. *Optimization Methods and Software*, 27(1):115–129, 2012.

M. Luque, K. Miettinen, A. B. Ruiz, and F. Ruiz. A two-slope achievement scalarizing function for interactive multiobjective optimization. *Computers & Operations Research*, 39(7):1673–1681, 2012.

H. M. Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.

A. A. Martinez, I. Pataki, G. Edmundson, E. Sebastian, D. Brabbins, and G. Gustafson. Phase II prospective study of the use of conformal high-dose-rate brachytherapy as monotherapy for the treatment of favorable stage prostate cancer: a feasibility report. *International Journal of Radiation Oncology * Biology **

Physics, 49(1):61–69, 2001.

K. Matusita. On the notion of affinity of several distributions and some of its applications. *Annals of the Institute of Statistical Mathematics*, 19(1):181–192, 1967.

P. Mavroidis, Z. Katsilieri, V. Kefala, N. Milickovic, N. Papanikolaou, A. Karabis, N. Zamboglou, and D. Baltas. Radiobiological evaluation of the influence of dwell time modulation restriction in HIPO optimized HDR prostate brachytherapy implants. *Journal of Contemporary Brachytherapy*, 2(3):117–128, 2010.

J. C. C. B. S. D. Mello, M. P. E. Lins, and E. G. Gomes. Construction of a smoothed dea frontier. *Pesquisa Operacional*, 22(2):183–201, 2002.

K. Miettinen. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, 1999.

N. Milickovic, M. Lahanas, M. Papagiannopoulou, N. Zamboglou, and D. Baltas. Multiobjective anatomy-based dose optimization for HDR-brachytherapy with constraint free deterministic algorithms. *Physics in Medicine and Biology*, 47(13): 2263–2280, 2002.

G. C. Morton, R. Sankreacha, P. Halina, and A. Loblaw. A comparison of anatomy-based inverse planning with simulated annealing and graphical optimization for high-dose-rate prostate brachytherapy. *Brachytherapy*, 7(1):12–16, 2008.

M. Nakata. A numerical evaluation of highly accurate multiple-precision arithmetic version of semidefinite programming solver: SDPA-GMP, -QD and -DD. In *2010 IEEE Multi-Conference on Systems and Control*, 29–34, 2010.

B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.

R. Nath, L. Anderson, G. Luxton, K. Weaver, J. Williamson, and A. Meigooni. Dosimetry of interstitial brachytherapy sources: recommendations of the AAPM radiation therapy committee task group no. 43. *Medical Physics*, 22(2):209–334, 1995.

G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1st edition, 1999.

T. S. Ng, Y. Sun, and J. Fowler. Semiconductor lot allocation using robust optimization. *European Journal of Operations Research*, 205(3):557–570, 2010.

A. Özmen, G.-W. Weber, I. Batmaz, and E. Kropat. RCMARS: Robustification of CMARS with different scenarios under polyhedral uncertainty set. *Communications in Nonlinear Science and Numerical Simulation*, 16(12):4780–4787, 2011.

A. Panchal. *Harmony search optimization for HDR prostate brachytherapy*. PhD thesis, Rosalind Franklin University of Medicine and Science, North Chicago, United

States, 2008.

E. Pantelis, P. Papagiannis, G. Anagnostopoulos, D. Baltas, P. Karaiskos, P. Sandilos, and L. Sakelliou. Evaluation of a TG-43 compliant analytical dosimetry model in clinical ^{192}Ir HDR brachytherapy treatment planning and assessment of the significance of source position and catheter reconstruction uncertainties. *Physics in Medicine and Biology*, 49(55):55–67, 2004.

K. Postek. Adjustable robust optimization by splitting the uncertainty set. Master's thesis, Tilburg University, the Netherlands, 2013.

E. Poulin, C.-A. C. Fekete, M. Létourneau, A. Fenster, J. Pouliot, and L. Beaulieu. Adaptation of the CVT algorithm for catheter optimization in high dose rate brachytherapy. *Medical Physics*, 40(11):111724, 2013.

A. Prékopa. *Stochastic Programming*. Kluwer Academic Publishers, 1995.

J. Pryor and J. W. Chinneck. Faster integer-feasibility in mixed-integer linear programs by branching to force change. *Computers & Operations Research*, 38(8): 1143–1152, 2011.

J. Ren, G. Menon, and R. Sloboda. Comparative evaluation of two dose optimization methods for image-guided, highly-conformal, tandem and ovoids cervix brachytherapy planning. *Physics in Medicine and Biology*, 58(7):2045–2058, 2013.

G. Rennen, E. R. van Dam, and D. den Hertog. Enhancement of sandwich algorithms for approximating higher-dimensional convex pareto sets. *INFORMS Journal on Computing*, 23(4):493–517, 2011.

A. Rényi. On measures of entropy and information. In *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability*, volume I, 547–561. University of California Press, 1961.

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

M. Roelofs and J. Bisschop. AIMMS 3.12: The Language Reference, Paragon Decision Technology BV, 2012.

M. Römer. Von der Komplexmethode zur robusten Optimierung und zurück. In *Beiträge zum Festkolloquium "Angewandte Optimierung" anlässlich des 65. Geburtstages von Prof. Dr. Rolf Rogge*, 37–61. Universität Halle–Wittenberg, 2010.

M. Rozenblit. Robust optimization of a multi-period production planning problem under uncertainty. Master's thesis, Technion – Israel Institute of Technology, 2010.

A. P. Ruszczyński and A. Shapiro. *Stochastic programming (handbooks in operations research and management science)*. Elsevier, 2003.

S. J. Sadjadi and H. Omrani. Data envelopment analysis with uncertain data: an application for Iranian electricity distribution companies. *Energy Policy*, 36(11): 4247–4254, 2008.

P. Salamon, P. Sibani, and R. Frost. *Facts, Conjectures, and Improvements for Simulated Annealing.* SIAM, Philadelphia, USA, 2002.

S. Schaible. Parameter-free convex equivalent and dual programs of fractional programming problems. *Zeitschrift für Operations Research*, 18(5):187–196, 1974.

S. Schaible. Fractional programming. II, On Dinkelbach's algorithm. *Management Science*, 22(8):868–873, 1976.

S. Schaible. Bibliography in fractional programming. *Zeitschrift für Operations Research*, 26(1):211–241, 1982.

S. Schaible and T. Ibaraki. Fractional programming. *European Journal of Operational Research*, 12(4):325–338, 1983.

A. H. Shokouhi, A. Hatami-Marbini, M. Tavana, and S. Saati. A robust optimization approach for imprecise data envelopment analysis. *Computers & Industrial Engineering*, 59(3):387–397, 2010.

A. H. Shokouhi, H. Shahriari, P. Agrell, and A. Hatami-Marbini. Consistent and robust ranking in imprecise data envelopment analysis under perturbations of random subsets of data. *OR Spectrum*, 36(1):133–160, 2014.

T. Siauw, A. Cunha, A. Atamtürk, I.-C. Hsu, J. Pouliot, and K. Goldberg. IPIP: A new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices. *Medical Physics*, 38(7):4045–4051, 2011.

T. Siauw, A. Cunha, D. Berenson, A. Atamtürk, I.-C. Hsu, K. Goldberg, and J. Pouliot. NPIP: A skew line needle configuration optimization system for HDR brachytherapy. *Medical Physics*, 39(7):4339–4346, 2012.

F. Siebert, M. Hirt, P. Niehoff, and G. Kovács. Imaging of implant needles for real-time HDR-brachytherapy prostate treatment using biplane ultrasound transducers. *Medical Physics*, 39(8):3406–3412, 2009.

C. Singh. Convex programming with set-inclusive constraints and its applications to generalized linear and fractional programming. *Journal of Optimization Theory and Applications*, 38(1):33–42, 1982.

M. Sion. On general minimax theorems. *Pacific Jouronal of Mathematics*, 8(1):171–176, 1958.

A. L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21(5):1154–1157, 1973.

A. L. Soyster. A duality theory for convex programming with set-inclusive constraints. *Operations Research*, 22(4):892–898, 1974.

I. M. Stancu-Minasian. A seventh bibliography of fractional programming. *Advanced Modeling and Optimization*, 15(2):309–386, 2013.

A. Thiele. *A robust optimization approach to supply chains and revenue management.*

PhD thesis, Massachusetts Institute of Technology, Cambridge MA, United States, 2004.

C. Thomas, A. Kruk, C. McGahan, I. Spadinger, and W. Morris. Prostate brachytherapy post-implant dosimetry: a comparison between higher and lower source density. *Radiotherapy & Oncology*, 83(1):18–24, 2007.

D. J. Thuente. Duality theory for generalized linear programs with computational methods. *Operations Research*, 28(4):1005–1011, 1980.

S. V. Utyuzhnikov, P. Fantini, and M. D. Guenov. A method for generating a well-distributed pareto set in nonlinear multiobjective optimization. *Journal of Computational and Applied Mathematics*, 223(2):820–841, 2009.

R. Van der Laarse and T. Prins. *Introduction to HDR brachytherapy optimization. In: Brachytherapy from Radium to Optimization, R. Mould, J. Battermann, A. Martinez and B. Speiser (Eds)*. Nucletron International B.V., 1994.

P. Vayanos, D. Kuhn, and B. Rustem. Decision rules for information discovery in multi-stage stochastic programming. In *2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 7368–7373. IEEE, 2011.

K. Wang and F. Wei. Robust data envelopment analysis based MCDM with the consideration of uncertain data. *Journal of Systems Engineering and Electronics*, 21(6):981–989, 2010.

C. Wei, Y. Li, and X. Cai. Robust optimal policies of production and inventory with uncertain returns and demand. *International Journal of Production Economics*, 134(2):357–367, 2009.

A. Wu, K. Ulin, and E. Sternick. A dose homogeneity index for evaluating ^{192}Ir interstitial breast implants. *Medical Physics*, 15(1):104–107, 1988.

H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *The Journal of Machine Learning Research*, 10:1485–1510, 2009.

Y. Yamada, L. Rogers, D. J. Demanes, G. Morton, B. R. Prestidge, J. Pouliot, G. N. Cohen, M. Zaider, M. Ghilezan, and I.-C. Hsu. American Brachytherapy Society consensus guidelines for high-dose-rate prostate brachytherapy. *Brachytherapy*, 11(1):20–32, 2012.

M. Yamashita, K. Fujisawa, K. Nakata, M. Nakata, M. Fukuda, K. Kobayashi, and K. Goto. A high-performance software package for semidefinite programs: SDPA 7. Technical report, Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, Tokyo, Japan, 2010.

İ. Yanıkoğlu, A. Ben-Tal, D. den Hertog, and M. Laurent. New nonlinear decision rules for adjustable robust counterparts. 2012. Unpublished results.

İ. Yanıkoğlu, D. den Hertog, and J. P. C. Kleijnen. Adjustable robust parameter design with unknown distributions. *CentER Discussion Paper*, 2013(022), 2013.

J. Zhang, M. M. Wiecek, and W. Chen. Local approximation of the efficient frontier in robust design. *Journal of Mechanical Design*, 122(2):232–236, 2000.