Risk and Robust Optimization

by

David Benjamin Brown

B.S., Electrical Engineering, Stanford University (2000)
M.S., Electrical Engineering, Stanford University (2001)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

Author ....
Department of Electrical Engineering and Computer Science
May 3, 2006

Certified by ....
Dimitris J. Bertsimas
Boeing Professor of Operations Research
Sloan School of Management
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students
Department of Electrical Engineering and Computer Science

Risk and Robust Optimization

by

David Benjamin Brown

B.S., Electrical Engineering, Stanford University (2000)

M.S., Electrical Engineering, Stanford University (2001)

Submitted to the Department of Electrical Engineering and Computer Science
on May 3, 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

This thesis develops and explores the connections between risk theory and robust optimization. Specifically, we show that there is a one-to-one correspondence between a class of risk measures known as coherent risk measures and uncertainty sets in robust optimization. An important consequence of this is that one may construct uncertainty sets, which are the critical primitives of robust optimization, using decision-maker risk preferences. In addition, we show some results on the geometry of such uncertainty sets. We also consider a more general class of risk measures known as convex risk measures, and show that these risk measures lead to a more flexible approach to robust optimization. In particular, these models allow one to specify not only the values of the uncertain parameters for which feasibility should be ensured, but also the *degree* of feasibility. We show that traditional, robust optimization models are a special case of this framework. As a result, this framework implies a family of probability guarantees on infeasibility at different levels, as opposed to standard, robust approaches which generally imply a single guarantee. Furthermore, we illustrate the performance of these risk measures on a real-world portfolio optimization application and show promising results that our methodology can, in some cases, yield significant improvements in downside risk protection at little or no expense in expected performance over traditional methods.

While we develop this framework for the case of linear optimization under uncertainty, we show how to extend the results to optimization over more general cones. Moreover, our methodology is scenario-based, and we prove a new rate of convergence result on a specific class of convex risk measures. Finally, we consider a multi-stage problem under uncertainty, specifically optimization of quadratic functions over uncertain linear systems. Although the theory of risk measures is still undeveloped with respect to dynamic optimization problems, we show that a set-based model of uncertainty yields a tractable approach to this problem in the presence of constraints.

Moreover, we are able to derive a near-closed form solution for this approach and prove new probability guarantees on its resulting performance.

Thesis Supervisor: Dimitris J. Bertsimas
Title: Boeing Professor of Operations Research
Sloan School of Management

Acknowledgments

When I first came to MIT in 2001, I did not have an advisor. I only knew I had research interests, roughly speaking, in the area of optimization. Consequently, I was taking a course on optimization from Dimitris Bertsimas, and I asked him one day after class if he would be willing to meet to discuss research. We met in his office a few days later, and, after a brief discussion about my background and interests, he told me he "had a good feeling" about me and was willing to take me on as a student.

Dimitris had on the order of ten students under his supervision at that point already, and it would have been very easy for him to tell me that he was unable at that time to accommodate another one. Instead, however, he relied on his intuition and gave me an opportunity for which I am very grateful. It is indeed difficult to imagine what trajectory my life would have taken if Dimitris had felt differently that day.

Working with Dimitris since then has been a true learning experience. Just like his approach to taking me as a student, he has shown me, more generally, the power of intuition in research. His message about maintaining a grounded perspective about the practicality of any theory has resonated with me, and I will carry it with me throughout my career. He has taught me the value of creativity the role of the proverbial "bigger picture" in research. It has been an honor, a privilege, and a pleasure to work with Dimitris Bertsimas.

I would also like to acknowledge the members of my thesis committee, Asu Ozdaglar, John Tsitsiklis, and Rob Freund, for their time and energy as thesis readers. I am particularly grateful to Asu for her extensive efforts in reading drafts of the thesis. I am also fortunate to have done teaching work with all three of them, as well as Dimitris and Dimitri Bertsekas. In addition to always being available for advice, John also helped me with funding one semester. Pablo Parrilo has shown a genuine interest in my work. and I am appreciative to have had some research discussions with him. I also would like to thank Sanjoy Mitter for providing me with some computing resources and research references during my time in LIDS. In terms of funding, this thesis work was also partially supported by Singapore-MIT Alliance and DARPA under grant number N666001-05-1-6030.

Chapter 4 is joint work between me, Dimitris, and Ronny Ben-Tal, who has been a visiting professor at MIT for the past year. Though I have only worked with Ronny for a few months, I have already learned quite a bit from him. It has been a joy working with him, and I feel very fortunate to have had this opportunity.

I would also like to recognize Alan Gous, Miguel Lobo, and Arman Maghbouleh, for it was my time working with them at Panopticon during the end of my undergraduate tenure at Stanford that really got me interested in optimization. Miguel, in particular, has been a mentor and role model from that point on and throughout my time as a graduate student. He has gone out of his way to help me on many occasions.

There are a number of support staff members at MIT who have helped me tremendously, particularly over the past year. Especially noteworthy are Rachel Cohen, Lynne Dell, and Christine Liberty, all of whom helped me with reference letters, thesis issues, and a whole host of other matters. I have constantly leaned on them for assistance.

I also have a number of friends at MIT who have supported me throughout the past five years. Constantine Caramanis and Georgios Kotsalis have not only been notable research peers but also great friends. I have had innumerable discussions about research with Constantine, and the past year, which was filled with the stresses of graduation and the academic job market, would have been very difficult without having him as a colleague. I also want to acknowledge Jeff and Monica Hixon for being extremely kind, generous, and loyal friends. I am additionally very grateful to Bora and Liba Mikic, housemasters at Next House, where I have served as a Graduate Resident Tutor for the past four years.

Finally, I want to thank my parents. They have never once failed to support me when I needed it. This thesis represents a culmination of many of the tremendous opportunities I have been blessed with throughout my life, and the foundation for all of this has been their unconditional love.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The fundamental problem of decision-making under uncertainty permeates nearly all fields of science, engineering, and industry. Many real-world decision problems can be modelled mathematically and cast as optimization problems, but the underlying parameters are rarely, if ever, known exactly. As Rockafellar [100] aptly puts it:

> *Decisions must often be taken in the face of the unknown. Actions decided upon in the present will have consequences that can't be fully determined until a later stage... the uncertainties in a problem have to be represented in such a manner that their effects on present decision-making can properly be taken into account. This is an interesting and challenging subject.*

The focus of this thesis will, in fact, be the problem of how to model uncertainty within an optimization framework. From the broadest viewpoint, we are interested in the following questions:

1. How do we model uncertainty? What primitives should we utilize in these models? Can we develop models which have both interesting, theoretical implications as well as practical merit?

2. How do we balance the tradeoff between models that are *flexible*, i.e., can accommodate a wide range of decision maker preferences, and those that are *prescriptive*, i.e., offer a clear methodology for modelling uncertainty?

3. What is the structure of these models, as well as their associated solutions in the context of optimization? Can we gain theoretical insights into how these solutions perform? What is their empirical performance in applications of interest?

4. How do we create models that are computationally tractable? This is paramount if we want our approach to have utility in practice.

As suggested by these questions, a running theme throughout this thesis will be an approach to modelling uncertainty which not only is theoretically motivated and justified, but also has value from a practical perspective as well.

1.1 Problem statement and approaches

We can state the problem of *certain* optimization in its most generic, mathematical form as

$$\begin{array}{ll} \text{minimize} & f_0(\boldsymbol{x}, \boldsymbol{\omega}) \\ \text{subject to} & f_i(\boldsymbol{x}, \boldsymbol{\omega}) \leq 0, \qquad i = 1, \ldots, m, \end{array} \tag{1.1}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is a decision vector, $\boldsymbol{\omega} \in \Omega$ is a parameter vector, and the functions $f_i : \mathbb{R}^n \times \Omega \to \mathbb{R}$ are the objective (for $i = 0$) and constraints of the problem, respectively.

Problem (1.1), as stated, is a family of problems parameterized by the specific value of $\boldsymbol{\omega}$. One computes an optimal solution $\boldsymbol{x}^*(\boldsymbol{\omega})$ as a function of the uncertain parameter. Clearly this can only be done if the decision-maker has perfect foresight on $\boldsymbol{\omega}$, in which case, uncertainty plays no role in the decision-making process. Obviously, this is not very interesting for our purposes.

What is more interesting, however, is a way of computing optimal solutions *before* $\boldsymbol{\omega}$ is revealed. Implicit in any such approach is not only a model for the uncertainty but also a metric for measuring a solution's performance in handling uncertainty. Presumably, the decision-maker would like the solutions to retain a low cost value and feasibility with some degree of reliability against the inherent uncertainty.

(1.1) does not account for the inherent uncertainty in $\boldsymbol{\omega}$ which is prevalent in most real-world problems. What is needed is some sort of model for this uncertainty embedded within the optimization problem. The literature on optimization under uncertainty bifurcates into two, essentially distinct approaches for modelling uncertainty. We now briefly discuss these methodologies.

1.1.1 Stochastic optimization

The approach of *stochastic optimization*, detailed in, for instance, Birge and Louveaux [36], Prekopa [94], and Ruszczynski and Shapiro [103], utilizes an underlying probability model to handle uncertainty. Specifically, one assumes the existence of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ for $\boldsymbol{\omega}$. From here, there are many variants in the problem statement, but they all depend critically on this underlying probability model. One possibility is the following:

$$\begin{aligned} \text{minimize} \quad & \mathbb{E}_{\mathbb{P}}\left[f_0(\boldsymbol{x}, \boldsymbol{\omega})\right] \\ \text{subject to} \quad & f_i(\boldsymbol{x}, \boldsymbol{\omega}) \leq 0 \qquad \text{almost surely}, \ i = 1, \ldots, m. \end{aligned} \tag{1.2}$$

Obviously, depending on the description of Ω, this problem could be quite conservative. More generally, one has the problem

$$\begin{aligned} \text{minimize} \quad & \mathbb{E}_{\mathbb{P}}\left[f_0(\boldsymbol{x}, \boldsymbol{\omega})\right] \\ \text{subject to} \quad & \mathbb{P}\left\{f_i(\boldsymbol{x}, \boldsymbol{\omega}) > 0\right\} \leq \epsilon_i, \qquad i = 1, \ldots, m, \end{aligned} \tag{1.3}$$

where $\epsilon_i \geq 0$. This is a problem with so-called *chance constraints*. Clearly, (1.2) is a special case of (1.3) in the limit when $\epsilon_i \to 0$.

The difficulties with stochastic optimization are twofold. First, such problems generally require detailed, if not full, distributional information for $\boldsymbol{\omega}$. In practice, this information is rarely known, and, in fact, it is questionable whether such a distribution even exists. More troublesome, however, is the fact that, even with total distributional information, stochastic optimization problems are computationally intractable (Shapiro and Nemirovski, [106]) for all but a few very select and special problem instances.

There are two common remedies to the intractability issue, and they have received considerable attention recently, particularly in the context of chance constraints. Given that chance constraints of the form (1.3) are, generally, highly non-convex in the decision variable $\boldsymbol{x}$, one approach is to find a convex approximation to the chance constraint. In particular, the goal of this approach is to find a convex function $g_i(\cdot)$ such that

$$g_i(\boldsymbol{x}) \leq 0 \quad \Rightarrow \quad \mathbb{P}\left\{f_i(\boldsymbol{x}, \boldsymbol{\omega}) > 0\right\} \leq \epsilon_i.$$

For more details on convex approximations to chance constraints, see Nemirovski and Shapiro [90].

The second approach, which is more relevant to this thesis, is scenario-based, i.e., using a samples of the uncertain parameter $\boldsymbol{\omega}$ (Calafiore and Campi, [41], [42], Nemirovski and Shapiro, [89]). This methodology approximates the exact chance constraint $\mathbb{P}\left\{f(\boldsymbol{x}, \boldsymbol{\omega}) > 0\right\} \leq \epsilon$ with N constraints of the form $f(\boldsymbol{x}, \boldsymbol{\omega}_i) \leq 0$, where the $\boldsymbol{\omega}_j$ are N independent samples of $\boldsymbol{\omega}$, then computes optimal solutions $\boldsymbol{x}_N^*$ to this scenario-based approximation. Specifically, one approximates (1.3) with the problem

$$\begin{array}{lll} \text{minimize} & \displaystyle\frac{1}{N}\sum_{j=1}^{N} f_0(\boldsymbol{x}, \boldsymbol{\omega}_j) & \\ \text{subject to} & f_i(\boldsymbol{x}, \boldsymbol{\omega}_j) \leq 0, & i = 1, \ldots, m, \;\; j = 1, \ldots, N. \end{array} \tag{1.4}$$

Notice that the probability of infeasibility of $\boldsymbol{x}_N^*$ for such problems stems not only from the uncertainty in $\boldsymbol{\omega}$, but also from the possibility that the random samples used to compute $\boldsymbol{x}_N^*$ provide a "misleading" description of the feasible set. In other words, we are interested not only in the probability of infeasibility due to the inherent randomness in $\boldsymbol{\omega}$, but also the *reliability* of this approach. There are a number of results related to this tradeoff. For instance, we have the following, due to Calafiore and Campi [41].

Theorem 1.1.1. *Consider a chance-constrained problem of the form*

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\boldsymbol{x} \\ \text{subject to} \quad & \mathbb{P}\{f(\boldsymbol{x},\boldsymbol{\omega}) > 0\} \leq \epsilon, \\ & \boldsymbol{x} \in X, \end{aligned} \tag{1.5}$$

where $\boldsymbol{\omega}$ is a random variable with support Ω, f is convex in $\boldsymbol{x}$ for all $\boldsymbol{\omega} \in \Omega$, and X is convex. Consider the problem:

$$\begin{aligned} \text{minimize} \quad & \mathbf{c}'\boldsymbol{x} \\ \text{subject to} \quad & f(\boldsymbol{x},\boldsymbol{\omega}_j) \leq 0, \qquad j = 1,\dots,N, \\ & \boldsymbol{x} \in X, \end{aligned} \tag{1.6}$$

where $\boldsymbol{\omega}_j$ are N independent samples of ω. Then we have the implication

$$N \geq \frac{2n}{\epsilon}\ln\left(\frac{12}{\epsilon}\right) + \frac{2}{\epsilon}\ln\left(\frac{2}{\delta}\right) + 2n \;\Rightarrow\; \mathbb{P}\{\boldsymbol{x}^* \in \mathcal{X}(\epsilon)\} \geq 1-\delta, \tag{1.7}$$

where $\boldsymbol{x}^*$ is an optimal solution to (1.6) and $\mathcal{X}(\epsilon) = \{\boldsymbol{x} \in X \mid \mathbb{P}\{f(\boldsymbol{x},\boldsymbol{\omega}) > 0\} \leq \epsilon\}$.

The idea of using scenarios will be a central theme for much of this thesis. The motivation is to directly utilize *data*, which, in many applications, is the only information we have regarding our uncertain parameters. Notice that, for fixed N, the scenario-based approximation used in the context of chance constraints (e.g., Theorem 1.1.1) allows only a *single* problem: namely, replacing the chance constraint with N constraints based on the sampled values of $\boldsymbol{\omega}$. When N is large, while the implied reliability could be high, the optimal solutions could be very conservative.

A key idea in this thesis will be a more general method of utilizing data directly within optimization that is based on decision maker *risk preferences*. Scenario approximations such as those in Theorem 1.1.1 can be viewed as limiting cases of the framework in this thesis in which the risk preferences are the most conservative possible for a fixed number of samples N. More critically, our approach will have connections to the uncertainty model used in robust optimization, a model which we now discuss.

1.1.2 Robust optimization

In contrast to the probabilistic nature of stochastic optimization, the approach of *robust optimization* does not rely on an underlying probability model. In this paradigm, one ensures feasibility for all realizations of the uncertain parameters within some prescribed *uncertainty set* $\mathcal{U}$. A robust optimization constraint, then, has the form

$$f_i(\boldsymbol{x}, \boldsymbol{\omega}) \leq 0, \qquad \forall\, \boldsymbol{\omega} \in \mathcal{U}.$$

The drawbacks of robust optimization, like stochastic optimization, are twofold. First, robust optimization is, even for convex functions and convex uncertainty sets, generally intractable. Unlike stochastic optimization, however, there are very large and practically relevant classes of problems for which we can efficiently solve their robust counterparts. This will become evident in Chapter 2 when we present a survey of the main tractability results in the robust optimization literature.

The second primary difficulty with robust optimization is that it is unclear what $\mathcal{U}$ should be. Obviously, this uncertainty set is the key primitive element of the problem, and its structure has a critical impact on the resulting solution. The uncertainty set structures suggested in the literature are primarily *ad hoc* ones which yield tractable problems. Ideally, we would like to retain these tractability results while proposing uncertainty sets which depend on the key primitives of *data* and *risk preferences*. This will be a central theme of this thesis.

1.2 Contributions and thesis outline

We summarize the contributions of this thesis as follows:

I. *Providing a tractable approach constructing uncertainty sets in robust optimization.*

We will present a methodology for uncertainty set construction in the optimization framework which relies on two primitive elements: first, a *data set* of realizations of the uncertain parameter; second, a *risk measure* reflecting the preferences of the decision maker. We will show, when this risk measure falls into a certain axiomatic class, how to explicitly construct uncertainty sets in the context of robust optimization. Moreover, we will prove tractability results for the resulting robust optimization problems. Our data-driven methodology is a significant generalization of other, scenario-based methods for stochastic optimization problems, and the links to robust optimization were previously unexplored.

At a higher level, we are connecting risk theory and robust optimization. In particular, we show that there is a one-to-one correspondence between risk measures of a particular class and uncertainty sets; specifically, the correspondence is between a class of risk measures called *coherent risk measures* and convex uncertainty sets. The risk measure and the uncertainty set descriptions are effectively dual representations of one another. In addition, we study classes of coherent risk measure which give rise to uncertainty sets of particular structure (e.g., polyhedral, conic, etc.) and provide explicit and tractable descriptions of the corresponding robust optimization problems.

II. *Developing a more flexible approach to robust optimization.*

We extend these results to the case when the risk measure belongs to the more general family of *convex risk measures.* We show that such risk measures yield a richer framework for robustness which allows one to control the degree of feasibility over the range of the uncertain parameters. We consider four, primary types of convex risk measures, each a variant of a special *certainty equivalent* measure. Each of these risk measures has different implications in terms of feasibility protection for the corresponding robust optimization problem. Furthermore, we illustrate the efficacy of this approach on a real-world, portfolio optimization problem.

III. *Providing geometric insights.*

We explore the geometric structure of the resulting uncertainty sets, and provide an explicit description of the class of measures which yield centrally symmetric, polyhedral uncertainty sets.

IV. *Extending to the multi-stage case.*

The discussion thus far has centered on single stage uncertainty, i.e., a single solution is computed, then the uncertain parameter is revealed entirely. Obviously, many problems, especially those involving sequential decisions over time, have an interleaving of stages in which decisions are implemented, then uncertainty is revealed. Loosely speaking, we refer to such problems as *dynamic*. Although the extension of risk measures to dynamic problems is an interesting problem which has received much recent attention, the literature on this subject is still very much incomplete. Nonetheless, we are able to develop a tractable model based on uncertainty sets for the important class of dynamic problems involving linear systems, quadratic costs, and constraints. We show that, in the unconstrained case, our approach has a near closed-form solution analogous to the Riccati equation from dynamic programming. For constrained problems, our approach maintains tractability, and this is in stark contrast to dynamic programming.

V. *Proving new probability guarantees.*

There are a number of new probability results in the thesis. First, for the case of convex risk measures, we are able to prove a family of probability guarantees for the robust solutions for various levels of infeasibility. Next, we prove an extension of the well-known Hoeffding inequality [70] for a conditional expectation risk measure which will be important throughout the thesis. Although this inequality does not directly apply to the case of data-driven optimization, it can be used as an *a posteriori* tool for analyzing the risk properties of resulting solutions. Finally, for our robust approach in the multi-stage case, we are able to prove probability guarantees on the distribution of the optimal cost function under normally distributed disturbances. This is important not only because the normality assumption is so common within the literature on control theory, but also because it illustrates that our approach, which is based on bounded uncertainty sets, performs well even when the underlying uncertainty model is based on unbounded random variables.

The remainder of the thesis is organized as follows. Chapter 2 provides background to both robust optimization and risk theory. Chapter 3 considers the case when the underlying problem is linear and the risk measure is coherent. Chapter 4 provides a more general framework based on convex risk measures; in addition, we also demonstrate this approach within the context of a real-world portfolio optimization application in this chapter. Chapter 5 shows how to generalize the approach to conic optimization models and contains some rate of convergence results for estimating risk measures from data. Finally, Chapter 6 considers the multi-stage problem with linear systems and quadratic costs, and Chapter 7 concludes the thesis and presents a summary of the results as well as important, open questions for future research.

Some notes on notation: We will use the following notation at various points throughout the thesis. For a function f, we denote the conjugate function by

$$f^*(\boldsymbol{y}) = \sup_{\boldsymbol{x}}\{\boldsymbol{y}'\boldsymbol{x} - f(\boldsymbol{x})\}.$$

For any norm $\|\cdot\|$, we denote the dual norm by $\|\boldsymbol{y}\|^* = \sup_{\|\boldsymbol{x}\|\leq 1}\{\boldsymbol{y}'\boldsymbol{x}\}$. Finally, the probability simplex in N dimensions will be denoted by Δ^N.

Chapter 2

Background

In this chapter, we present a short review of the literature on robust optimization and risk theory as preliminaries for the remainder of the thesis. Our goal is not only to provide an overview of the two subjects, but also to present some of the main results that we will later use in this thesis.

2.1 Robust optimization

A static, robust optimization problem (i.e., one without recourse variables) can be rather generically stated as

$$\begin{array}{ll} \text{minimize} & f_0(\boldsymbol{x}) \\ \text{subject to} & f_i(\boldsymbol{x}, \boldsymbol{u}_i) \leq 0, \qquad \forall\, \boldsymbol{u}_i \in \mathcal{U}_i,\ i = 1, \ldots, m, \end{array} \tag{2.1}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is a vector of decision variables, $f_0 : \mathbb{R}^n \to \mathbb{R}$ is an objective (cost) function, $f_i : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ are m constraint functions, $\boldsymbol{u}_i \in \mathbb{R}^m$ are *disturbance vectors* or *parameter uncertainties*, and $\mathcal{U}_i \subseteq \mathbb{R}^m$ are *uncertainty sets*, which, for our purposes, will always be closed. The goal of (2.1) is to compute minimum cost solutions $\boldsymbol{x}^*$ among all those solutions which are feasible for *all* realizations of the disturbances $\boldsymbol{u}_i$ within $\mathcal{U}_i$. Thus, if some of the $\mathcal{U}_i$ are continuous sets, (2.1), as stated, has an infinite number of constraints. Intuitively, this problem offers some measure of feasibility protection for optimization problems containing parameters which are not known exactly.

It is worthwhile to notice the following, straightforward facts about the problem statement of (2.1):

1. The fact that the objective function is unaffected by parameter uncertainty is without loss of generality; indeed, if there is parameter uncertainty in the objective, we may always introduce an auxiliary variable, call it t, and minimize t subject to the additional constraint $\max_{\boldsymbol{u}_0 \in \mathcal{U}_0} f_0(\boldsymbol{x}, \boldsymbol{u}_0) \leq t$. This, of course, assumes the objective of minimizing the worst or largest possible realization of the function $f_0(\boldsymbol{x}, \boldsymbol{u}_0)$ aligns with the goals of the decision-maker.

2. It is also without loss of generality to assume that the parameters $\boldsymbol{u}_i$ belong to distinct uncertainty sets $\mathcal{U}_i$. Indeed, if we have a single uncertainty set $\mathcal{U}$ for which we require $(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_m) \in \mathcal{U}$, then the constraint-wise feasibility requirement implies an equivalent problem is (2.1) with the $\mathcal{U}_i$ taken as the projection of $\mathcal{U}$ along the corresponding dimensions (see Ben-Tal and Nemirovski, [14] for more on this).

3. Exact constraints are also captured in this framework by assuming the corresponding $\mathcal{U}_i$ to be singletons.

4. Problem (2.1) also contains the instances when the decision or disturbance vectors are contained in more general vector spaces than $\mathbb{R}^n$ or $\mathbb{R}^m$ (e.g., $\mathbb{S}^n$ in the case of semidefinite optimization) with the definitions modified accordingly.

We emphasize that robust optimization is distinctly different than the field of *sensitivity analysis*, which is typically applied as a post-optimization tool for quantifying the change in cost for small perturbations in the underlying problem data. Here, our goal is to *compute* solutions with *a priori* ensured feasibility when the problem parameters vary within the prescribed uncertainty set. We refer the reader to some of the standard optimization literature (e.g., Bertsimas and Tsitsiklis, [35], Boyd and

Vandenberghe, [40]) and works on perturbation theory (e.g., Freund, [61], Renegar, [99]) for more on sensitivity analysis.

2.1.1 Prior work

In the early 1970s, Soyster [107] was one of the first researchers to investigate explicit approaches to robust optimization. This short note focused on robust linear optimization in the case where the column vectors of the constraint matrix were constrained to belong to ellipsoidal uncertainty sets; Falk [55] followed this a few years later with more work on "inexact linear programs." The optimization community, however, was relatively quiet on the issue of robustness until the work of Ben-Tal and Nemirovski (e.g., [13], [14], [16]) and El Ghaoui et al. [66] in the late 1990s. This work, coupled with advances in computing technology and the development of fast, interior point methods for convex optimization, particularly for semidefinite optimization (e.g., Boyd and Vandenberghe, [39]) sparked a massive flurry of interest in the field or robust optimization.

It is not at all clear when (2.1) is efficiently solvable. One might imagine that the addition of robustness to a general optimization problem comes at the expense of significantly increased computational complexity. Although this is indeed generally true, there are many robust problems which may be handled in a tractable manner, and much of the literature since the modern resurgence has focused on specifying classes of functions f_i, coupled with the types of uncertainty sets $\mathcal{U}_i$, that yield tractable problems. For the most part, this is tantamount to the feasible set

$$X(\mathcal{U}) \quad = \quad \{\boldsymbol{x} \mid f_i(\boldsymbol{x}, \boldsymbol{u}_i) \leq 0 \ \forall \ \boldsymbol{u}_i \in \mathcal{U}_i, \ i = 1, \ldots, m\}, \tag{2.2}$$

being convex in $\boldsymbol{x}$, where $\mathcal{U} = \mathcal{U}_1 \times \ldots \times \mathcal{U}_m$. We now present an abridged taxonomy of some of the main results related to this issue. We note that the background we present here will be for robust optimization in static optimization problems, although we will apply some of these results in a dynamic setting in Chapter 6. For more on robustness in an adaptive setting, the reader may also consult, for instance, Ben-Tal et al. [10] and Caramanis [43].

2.1.2 Complexity and tractability results

Robust linear optimization

The robust counterpart of a linear optimization problem is written, without loss of generality, as

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}, \qquad \forall\, \boldsymbol{A} \in \mathcal{U}, \end{aligned} \tag{2.3}$$

where $\mathcal{U} \subseteq \mathbb{R}^{m \times n}$. Ben-Tal and Nemirovski [14] show the following:

Theorem 2.1.1. *(Ben-Tal and Nemirovski, [14]) When $\mathcal{U}$ is "ellipsoidal," i.e., $\mathcal{U}$ satisfies:*

1. $\mathcal{U} = \bigcap_{l=0}^{k} U(\Pi_l, \boldsymbol{Q}_l)$, *where*

$$U(\Pi, \boldsymbol{Q}) \;=\; \{\Pi(\boldsymbol{u}) \mid \|\boldsymbol{Q}\boldsymbol{u}\| \leq 1\},$$

where $\boldsymbol{u} \rightarrow \Pi(\boldsymbol{u})$ is an affine embedding of $\mathbb{R}^L$ into $\mathbb{R}^{m \times n}$ and $\boldsymbol{Q} \in \mathbb{R}^{M \times L}$; and

2. *$\mathcal{U}$ is bounded with nonempty interior,*

then Problem (2.3) is equivalent to a second-order cone program (SOCP), i.e., (2.3) may be converted to a problem with the constraints

$$\|\boldsymbol{B}_i \boldsymbol{x} + \boldsymbol{b}_i\| \leq \boldsymbol{a}_i' \boldsymbol{x} + \alpha_i, \; i = 1, \ldots, K,$$

for some number K of appropriate matrices $\boldsymbol{B}_i$, vectors $\boldsymbol{a}_i$ and $\boldsymbol{b}_i$, and reals α_i.

Additionally, the case of ellipsoidal uncertainty covers the case of polyhedral uncertainty, as the authors show [14]. In fact, when $\mathcal{U}$ is polyhedral, the robust counterpart is equivalent to a linear optimization problem. Furthermore, the size of such problems grows polynomially in the size of the nominal problem and the dimensions of the uncertainty set.

Bertsimas et al. [30] show that robust linear optimization problems with uncertainty sets described by more general norms lead to convex problems with constraints related to the dual norm. Here we use the notation vec($\boldsymbol{A}$) to denote the vector formed by concatenating all the rows of the matrix $\boldsymbol{A}$.

Theorem 2.1.2. *(Bertsimas et al., [30]) With the uncertainty set*

$$\mathcal{U} = \{\boldsymbol{A} \mid \|\boldsymbol{M}(\text{vec}(\boldsymbol{A}) - \text{vec}(\bar{\boldsymbol{A}}))\| \leq \Delta\},$$

where $\boldsymbol{M}$ is an invertible matrix, $\bar{\boldsymbol{A}}$ is any constant matrix, and $\|\cdot\|$ is any norm, Problem (2.3) is equivalent to the problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \bar{\boldsymbol{a}}_i'\boldsymbol{x} + \Delta\|(\boldsymbol{M}')^{-1}\boldsymbol{x}_i\|^* \leq b_i, \qquad i = 1, \ldots, m, \end{array}$$

where $\boldsymbol{x}_i \in \mathbb{R}^{(m \cdot n) \times 1}$ is a vector that contains $\boldsymbol{x} \in \mathbb{R}^n$ in entries $(i-1) \cdot n + 1$ through $i \cdot n$ and 0 everywhere else, and $\|\cdot\|^*$ is the corresponding dual norm of $\|\cdot\|$.

Thus the norm-based model shown in Theorem 2.1.2 yields an equivalent problem with corresponding dual norm constraints. For the most part, then, robust linear optimization problems over uncertainty sets described by norms are tractable; in particular, the l_1 and l_∞ norms result in linear optimization problems, and the l_2 norm results in a second-order cone problem.

In short, for many choices of the uncertainty set, robust linear optimization problems are tractable.

Robust quadratic optimization

For $f_i(\boldsymbol{x}, \boldsymbol{u}_i)$ of the form

$$\|\boldsymbol{A}_i\boldsymbol{x}\|^2 + \boldsymbol{b}_i'\boldsymbol{x} + c_i \leq 0.$$

i.e., *quadratically constrained quadratic programs* (QCQP), where $\boldsymbol{u}_i = (\boldsymbol{A}_i, \boldsymbol{b}_i, c_i)$, the robust counterpart is a semidefinite optimization problem if $\mathcal{U}$ is a single ellipsoid, and NP-hard if $\mathcal{U}$ is polyhedral (Ben-Tal and Nemirovski, [13], [15]).

For robust SOCPs, the $f_i(\boldsymbol{x}, \boldsymbol{u}_i)$ are of the form

$$||\boldsymbol{A}_i\boldsymbol{x} + \boldsymbol{b}_i|| \leq \boldsymbol{c}_i'\boldsymbol{x} + d_i.$$

If $(\boldsymbol{A}_i, \boldsymbol{b}_i)$ and $(\boldsymbol{c}_i, d_i)$ each belong to a set described by a single ellipsoid, then the robust counterpart is a semidefinite optimization problem; if $(\boldsymbol{A}_i, \boldsymbol{b}_i, \boldsymbol{c}_i, d_i)$ varies within a shared ellipsoidal set, however, the robust problem is NP-hard (Ben-Tal et al., [20], Bertsimas and Sim, [33]).

Robust semidefinite optimization

With ellipsoidal uncertainty sets, robust counterparts of semindefinite optimization problems are NP-hard (Ben-Tal and Nemirovski, [13], Ben-Tal et al. [8]). Similar negative results hold even in the case of polyhedral uncertainty sets (Nemirovski, [88]).

Computing approximate solutions, i.e., solutions that are robust *feasible* but not robust *optimal* to robust semidefinite optimization problems has, as a consequence, received considerable attention (e.g., El Ghaoui et al., [66], Ben-Tal and Nemirovski, [19], [18], and Bertsimas and Sim, [33]). Some tightness results have been obtained. For instance, Ben-Tal and Nemirovski show [19] that an appropriately defined "level of conservativeness" of a particular approximation to a robust SDP with box uncertainty grows no faster than $\pi\sqrt{\mu}/2$, where μ is the maximum rank of the matrices describing $\mathcal{U}$. Here, they define the level of conservativeness as

$$\rho(\mathrm{AR} : \mathrm{R}) \quad = \quad \inf\left\{\rho \geq 1 \mid X(\mathrm{AR}) \supseteq X(\mathcal{U}(\rho))\right\},$$

where $X(\mathrm{AR})$ is the feasible set of the approximate robust problem and $X(\mathcal{U}(\rho))$ is the feasible set of the original robust SDP with the uncertainty set "inflated" by a factor of ρ.

Bertsimas and Sim [33] develop an approach to which is flexible enough to approximate robust optimization problems over general, convex cones.

Robust geometric programming

A *geometric program* (GP) is a convex optimization problem of the form

$$\begin{array}{lll} \text{minimize} & \boldsymbol{c}'\boldsymbol{y} & \\ \text{subject to} & g(\boldsymbol{A}_i\boldsymbol{y} + \boldsymbol{b}_i) \leq 0, & i = 1, \ldots, m, \\ & \boldsymbol{G}\boldsymbol{y} + \boldsymbol{h} = \boldsymbol{0}, & \end{array}$$

where $g : \mathbb{R}^k \rightarrow \mathbb{R}$ is the *log-sum-exp* function, i.e.,

$$g(\boldsymbol{x}) = \log \left(\sum_{i=1}^{k} e^{x_i} \right),$$

and the matrices and vectors $\boldsymbol{A}_i$, $\boldsymbol{G}$, $\boldsymbol{b}_i$, and $\boldsymbol{h}$ are of appropriate dimension. For many engineering, design, and statistical applications of GP, see Boyd and Vandenberghe [40]. Hsiung et al. [71] study a robust version of GP with the constraints

$$g(\tilde{\boldsymbol{A}}_i(\boldsymbol{u})\boldsymbol{v} + \tilde{\boldsymbol{b}}_i(\boldsymbol{u})) \leq 0 \qquad \forall\, \boldsymbol{u} \in \mathcal{U},$$

where $(\tilde{\boldsymbol{A}}_i(\boldsymbol{u}), \tilde{\boldsymbol{b}}_i(\boldsymbol{u}))$ are affinely dependent on the uncertainty $\boldsymbol{u}$, and $\mathcal{U}$ is an ellipsoid or a polyhedron. The complexity of this problem is unknown; the approach in [71] is to use a piecewise linear approximation to get upper and lower bounds to the robust GP.

Robust discrete optimization

Kouvelis and Yu [77] study robust models for some discrete optimization problems, although their approach yields robust counterparts to a number of polynomially solvable combinatorial problems which are in turn NP-hard. For instance, the problem minimizing the maximum shortest path on a graph over only two scenarios for the cost vector can be shown to be an NP-hard problem [77].

Bertsimas and Sim [31], however, present a model for cost uncertainty in which each coefficient c_j is allowed to vary within the interval $[\bar{c}_j, \bar{c}_j + d_j]$, but no more than $\Gamma \geq 0$ coefficients may vary. They then apply this model to a number of combinatorial problems, i.e., they attempt to solve

$$\begin{array}{ll} \text{minimize} & \bar{\boldsymbol{c}}'\boldsymbol{x} + \max\limits_{\{S \,\mid\, S \subseteq N,\ |S| \leq \Gamma\}} \sum\limits_{j \in S} d_j x_j \\ \text{subject to} & \boldsymbol{x} \in X, \end{array}$$

where $N = \{1, \ldots, n\}$ and X is a fixed set. They show that under this model for uncertainty, the robust version of a combinatorial problem may be solved by solving no more than $n + 1$ versions of the underlying, nominal problem. They also show that this result extends to approximation algorithms for combinatorial problems. For network flow problems, they show that the above model can be applied and the robust solution can be computed by solving a logarithmic number of nominal, network flow problems.

Atamtürk [5] shows that, under an appropriate uncertainty model for the cost vector in a mixed 0-1 integer program, there is a tight, linear programming formulation of the robust mixed 0-1 problem with size polynomial in the size of a tight linear programming formulation for the nominal mixed 0-1 problem.

2.1.3 Some probability guarantees

In addition to tractability, a central question in the robust optimization literature has been probability guarantees on feasibility under particular distributional assumptions for the disturbance vectors. Specifically, what does robust feasibility imply about probability of feasibility, i.e., what is the smallest ϵ we can find such that

$$\boldsymbol{x} \in X(\mathcal{U}) \Rightarrow \mathbb{P}\left\{f_i(\boldsymbol{x}, \boldsymbol{u}_i) > 0\right\} \leq \epsilon$$

under (generally mild) assumptions on a distribution for $\boldsymbol{u}_i$? In this section, we briefly survey some of the results in this vein.

For linear optimization, Ben-Tal and Nemirovski [16] propose a robust model based on ellipsoids of radius Ω. Under this model, if the uncertain coefficients have bounded, symmetric support, they show that the corresponding robust feasible solutions are feasible with probability $e^{-\Omega^2/2}$. In a similar spirit, Bertsimas and Sim [32] propose an uncertainty set of the form

$$\mathcal{U}_\Gamma = \left\{ \bar{\boldsymbol{a}} + \sum_{j \in J} z_j \hat{\boldsymbol{a}}_j \;\middle|\; \|\boldsymbol{z}\|_\infty \leq 1, \; \sum_{j \in J} \mathbf{1}(z_j) \leq \Gamma \right\}, \tag{2.4}$$

for the coefficients $\boldsymbol{a}$ of an uncertain, linear constraint. Here, $\mathbf{1} : \mathbb{R} \to \mathbb{R}$ denotes the indicator function of y, i.e., $\mathbf{1}(y) = 0$ if and only if $y = 0$, $\bar{\boldsymbol{a}}$ is a vector of "nominal" values, $J \subseteq \{1, \ldots, n\}$ is an index set of uncertain coefficients, and $\Gamma \leq |J|$ is an integer[1] reflecting the number of coefficients which are allowed to deviate from their nominal values. The following then holds.

Theorem 2.1.3. *(Bertsimas and Sim [32]) Let $\boldsymbol{x}^*$ satisfy the constraint*

$$\max_{\boldsymbol{a} \in \mathcal{U}_\Gamma} \boldsymbol{a}'\boldsymbol{x} \leq b,$$

where $\mathcal{U}_\Gamma$ is as in (2.4). If the random vector $\tilde{\boldsymbol{a}}$ has independent components with a_j distributed symmetrically on $[\bar{a}_j - \hat{a}_j, \bar{a}_j + \hat{a}_j]$ if $j \in J$ and $a_j = \bar{a}_j$ otherwise, then

$$\mathbb{P}\left\{\tilde{\boldsymbol{a}}'\boldsymbol{x}^* > b\right\} \leq e^{-\frac{\Gamma^2}{2|J|}}.$$

In the case of linear optimization with partial moment information (specifically, known mean and covariance), Bertsimas et al. [30] prove guarantees for the general norm uncertainty model used in Theorem 2.1.2. For instance, when $\|\cdot\|$ is the Euclidean norm, $\boldsymbol{x}^*$ feasible to the robust problem in Theorem 2.1.2 can be shown

[1] The authors also consider Γ non-integer, but we omit this straightforward extension for notational convenience.

[30] to imply the guarantee

$$\mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x}^* > b\} \leq \frac{1}{1+\Delta^2},$$

where Δ is the radius of the uncertainty set, and the mean and covariance are used for $\bar{\boldsymbol{A}}$ and $\boldsymbol{M}$, respectively.

For more general, robust conic optimization problems, there are less results on probability guarantees. Bertsimas and Sim [33] are able to prove probability guarantees for their approximate robust solutions. They use the following model for data uncertainty:

$$\tilde{\boldsymbol{D}} = \boldsymbol{D}^0 + \sum_{j \in N} \Delta \boldsymbol{D}^j \tilde{z}_j,$$

where $\boldsymbol{D}^0$ is the nominal data value and $\Delta \boldsymbol{D}^j$ are data perturbations. The $\tilde{z}_j$ are random variables with mean zero and independent, identical distributions. The robust problem in this case is

$$\max_{\tilde{\boldsymbol{D}} \in \mathcal{U}_\Omega} f(\boldsymbol{x}, \tilde{\boldsymbol{D}}) \leq 0, \tag{2.5}$$

where

$$\mathcal{U}_\Omega = \left\{ \boldsymbol{D}^0 + \sum_{j \in N} \Delta \boldsymbol{D}^j u_j \;\middle|\; \|\boldsymbol{u}\| \leq \Omega \right\}, \tag{2.6}$$

and f satisfies some convexity assumptions. They then prove the following probability bound under normal distributions.

Theorem 2.1.4. *(Bertsimas and Sim, [33]) Let $\boldsymbol{x}^*$ be robust feasible to the robust conic optimization problem with the model of uncertainty (2.6). When we use the l_2-norm in (2.6), and under the assumption that $\boldsymbol{u} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$, we have the probability* *bound*

$$\mathbb{P}\{f(\boldsymbol{x}, \boldsymbol{D}) > 0\} \leq \frac{\sqrt{e}\Omega}{\alpha} \exp\left(-\frac{\Omega^2}{2\alpha^2}\right).$$

where $\alpha = 1$ *for LPs,* $\alpha = \sqrt{2}$ *for SOCPs, and* $\alpha = \sqrt{m}$ *for SDPs (*m *is the dimension of the matrix in the SDP).*

2.1.4 Applications of robust optimization

Techniques from robust optimization have been applied across a vast array of problems; here we mention and cite only a few of the many applications and corresponding studies which have utilized these ideas. See the references within the following for much more on robust optimization in applications.

1. *Communications*: Ben-Tal and Nemirovski [17], Hsiung et al. [72], [81].

2. *Control theory*: Zhou et al. [119], El Ghaoui et al. [66], El Ghaoui and Calafiore [63], Bertsimas and Brown [27], Ben-Tal et al. [7].

3. *Network problems*: Bertsimas and Sim, [31].

4. *Estimation/classification*: El Ghaoui and Lebret [64], Lanckriet et al. [78], Eldar et al. [53], Kim et al. [76].

5. *Markov decision processes*: El Ghaoui and Nilim [65].

6. *Engineering design*: Ben-Tal and Nemirovski [17], [12], Boyd et al. [38], Xu et al. [116].

7. *Finance*: Ben-Tal et al. [11], Goldfarb and Iyengar [68], Lobo and Boyd [80], and Lobo [79].

8. *Supply chain problems*: Bertsimas and Thiele [34], Ben-Tal et al. [9].

2.2 Risk theory

2.2.1 Preliminaries

The field of risk theory focuses on axiomatic approaches to quantifying random, future outcomes, and its development has been led by both economists and mathematicians alike. The generic problem setup is to assume the existence of a fixed, underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We will occasionally refer to $\mathbb{P}$ as the "reference measure" for reasons that will become clear shortly. Furthermore, for the purposes of this thesis, the case of Ω finite will be sufficient, although almost all of the fundamental results discussed do extend to more general probability spaces.

This probability space is known but beyond our control. What is potentially within our control, however, is the selection of a real-valued random variable $X : \Omega \to \mathbb{R}$. Let $\mathcal{X}$ be the space of all random variables on Ω.

To provide a concrete illustration relevant to this discussion, we can consider the case of linear optimization under uncertainty. In this context, an uncertain constraint vector $\tilde{\boldsymbol{a}}$ could be random with an event space $\Omega \subseteq \mathbb{R}^n$. Our decision vector $\boldsymbol{x}$ must be chosen to belong to some feasible set $X \subseteq \mathbb{R}^n$. Through $\boldsymbol{x}$, we control the distribution of the random variable $\tilde{\boldsymbol{a}}'\boldsymbol{x}$, and the corresponding space of random variables is

$$\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \mid \boldsymbol{x} \in X\} \subseteq \mathcal{X}.$$

What the decision-maker desires, then, is a compact way of comparing such random variables; this motivates the following definition.

Definition 2.2.1. A **risk measure** is a function $\mu : \mathcal{X} \to \mathbb{R}$.

Throughout this thesis, we will use the convention that *smaller* values of μ are preferred, i.e., a random variable X is preferred over a random variable Y under μ if and only if $\mu(X) \leq \mu(Y)$. Thus, the risk measure effectively measures *loss*, as opposed to gain. This convention runs somewhat counter to some of the following literature on risk theory, but will be more convenient for our purposes.

2.2.2 Foundations

The concept of risk measures can be traced back to the work of Bernoulli [25], who offered a proposal to resolve the so-called "St. Petersburg Paradox" [93]. This work served as the genesis for *utility theory*, starting most notably with the celebrated work of von Neumann and Morgenstern [111]. One of their biggest contributions was to show the equivalence between an axiomatized preference relation $\succeq$ on random variables and the existence of a utility function $u : \mathbb{R} \to \mathbb{R}$ in the sense that, for any $X,\ Y \in \mathcal{X}$, $X \succeq Y$ if and only if $\mathbb{E}\left[u(X)\right] \geq \mathbb{E}\left[u(Y)\right]$. For a more recent treatment of utility theory, see, for instance, Fishburn [57].

The notion of variance as a risk measure also became highly popularized due to the seminal work on portfolio optimization by Markowitz [83]. More recent works, motivated in part by prospect theory (e.g., Kahneman and Tversky [75], who cite inconsistencies between utility theory and observed human behavior) such as Quiggin [95] and Yaari [117], have attemped to extend the notion of utility theory to risk measures satisfying different axioms.

2.2.3 Coherent risk measures

At first glance, it may seem that the definition of a risk measure above is quite restrictive in that risks are only described by a single, real number. Indeed, one may imagine that is a significant loss of information. In a seminal paper, Artzner et al. [3] address this by attempting to axiomatize risk measures. Their starting point is the notion of an "acceptance set," of viable random variables, and they show a correspondence between acceptance sets satisfying particular axioms and risk measures satisfying related axioms. In this context, the notion of describing risk by a single number is quite natural, as a random variable is either acceptable or it is not.

Specifically, the authors do in fact consider the case when Ω is finite, and denote by $\mathcal{A} \subseteq \mathcal{X}$ the "acceptance set," i.e., the set of random variables on $(\Omega, \mathcal{F}, \mathbb{P})$ which are acceptable. They argue that any reasonable acceptance set should satisfy the following axioms:

A1. $\{X \in \mathcal{X} \mid X(\omega) \leq 0 \; \forall \, \omega \in \Omega\} \subseteq \mathcal{A}$.

A2. $\{X \in \mathcal{X} \mid X(\omega) > 0 \; \forall \, \omega \in \Omega\} \cap \mathcal{A} = \emptyset$.

A3. $\mathcal{A}$ is convex.

A4. $\mathcal{A}$ is a positively homogeneous cone.

A1 simply states that any random variable which never results in a positive loss should always be accepted; A2, conversely, argues that any random variable which results in a positive loss in every scenario should never be accepted. A3 essentially reflects risk aversion on the part of the decision maker. A4 does not have a particularly meaningful interpretation in terms of acceptance sets, but will have an interpretation in terms of the corresponding risk measure.

These acceptance sets, in fact, naturally induce a risk measure and vice versa.

Definition 2.2.2. For an acceptance set $\mathcal{A} \subseteq \mathcal{X}$, the **induced risk measure** is

$$\mu_{\mathcal{A}}(X) \;\; = \;\; \inf_{m \in \mathbb{R}} \{m \mid X - m \in \mathcal{A}\}.$$

Conversely, if μ is a risk measure, then the **induced acceptance set** is

$$\mathcal{A}_{\mu} \;\; = \;\; \{X \in \mathcal{X} \mid \mu(X) \leq 0\}.$$

The interpretation of $\mu_{\mathcal{A}}(X)$ is the minimum sure payment the decision maker would need to receive to find X acceptable. The authors also introduced the following, now famous, axiomatized definition for risk measures.

Definition 2.2.3. A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is **coherent** if it satisfies the following four properties:

(a) (*positive homogeneity*)	$\mu(aX) = a\mu(X)$,	$\forall \, X \in \mathcal{X}, \; a \geq 0$.
(b) (*translation invariance*)	$\mu(X + b) = \mu(X) + b$,	$\forall \, X \in \mathcal{X}, \; b \in \mathbb{R}$.
(c) (*monotonicity*)	$\mathbb{P}\{X \leq Y\} = 1 \Rightarrow \mu(X) \leq \mu(Y)$,	$\forall \, X, Y \in \mathcal{X}$.
(d) (*subadditivity*)	$\mu(X + Y) \leq \mu(X) + \mu(Y)$,	$\forall \, X, Y \in \mathcal{X}$.

The positive homogeneity axiom implies that risk scales linearly with the size of a random variable. Translation invariance implies that there is an absolute scale for measuring risks; note that it immediately rules out the commonly used risk measures of standard deviation or variance. The monotonicity axiom simply states the reasonable requirement that the risk of a random variable should always be less than the corresponding risk of a random variable which always results in a greater loss (this also rules out standard deviation). Finally, subadditivity ensures that risk cannot increase from a merger; it is a crucial property related to convexity and rules out quantile as a risk measure.

Artzner et al. [3] show the correspondence between their axiomatized acceptance sets and coherent risk measures.

Proposition 2.2.1. (Artzner et al. [3]) For a closed acceptance set satisfying axioms A1-A4, the induced risk measure $\mu_{\mathcal{A}}$ is coherent. Conversely, for a coherent risk measure μ, the induced acceptance set $\mathcal{A}_\mu$ is closed and satisfies axioms A1-A4.

The intuition behind Proposition 2.2.1 is that coherent risk measures correspond exactly to a set of acceptable random variables, and this acceptance set satisfies the axioms A1-A4. A critical result shown in [3] is a representation theorem for coherent risk measures, which we now state.

Theorem 2.2.1. *(Artzner et al., [3]) A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is coherent if and only if there exists a family of probability measures $\mathcal{Q}$ that*

$$\mu(X) = \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q [X], \qquad \forall\, X \in \mathcal{X}, \tag{2.7}$$

where $\mathbb{E}_Q [X]$ denotes the expectation of the random variable X under the measure Q (as opposed to the reference measure $\mathbb{P}$).

Theorem 2.2.1 was known before the emergence of coherent risk measures; see, Huber [73], for example, for a proof. The essence of the proof is the separating hyperplane theorem for convex sets. Indeed, Theorem 2.2.1 is a duality result, and it says that any coherent risk measure has a dual description in terms of expectations over a family $\mathcal{Q}$ of "generating measures."

This theorem will be central in making the connections between these risk measures and uncertainty sets in robust optimization. The interpretation is that all coherent risk measures may be written as the supremum of the expected value over a set of "generalized scenarios" [3] representing different probability measures.

At first glance of Theorem 2.2.1, it may seem that the reference measure $\mathbb{P}$, which is the underlying probability measure for X, is irrelevant in evaluating a coherent risk measure. Implicit in this definition, however, is the fact that all measures in $\mathcal{Q}$ are absolutely continuous with respect to $\mathbb{P}$. In the finite case, this is straightforward, as the measures are finite dimensional vectors.

Here are a few examples of generating families and the corresponding description of the risk measure for the case when X is a discrete random variable with support of cardinality N:

1. When $\mathcal{Q} = \{\mathbb{P}\}$, we obtain $\mu(X) = \mathbb{E}[X]$.

2. When $\mathcal{Q} = \operatorname{conv}(\{\boldsymbol{q}_1, \ldots, \boldsymbol{q}_m\})$, where $\boldsymbol{q}_i \in \Delta^N$ then $\mu(X) = \max\limits_{i=1,\ldots,m} \mathbb{E}_{\boldsymbol{q}_i}[X]$.

3. When $\mathcal{Q} = \left\{\boldsymbol{q} \in \Delta^N \mid q_i \leq p_i/\alpha,\ i = 1, \ldots, N\right\}$ for some $\alpha \in (0,1]$, then it can be shown (Rockafellar and Uryasev [102]) that

$$\mu(X) = \inf_{\nu} \left\{ \nu + \frac{1}{\alpha} \mathbb{E}\left[(X-\nu)^+\right] \right\}. \tag{2.8}$$

   For random variables with a continuous distribution function, it can be further shown (e.g., Acerbi and Tasche [1]) that the risk measure in (2.8) is equivalent to $\mathbb{E}[X \mid X \geq \mathrm{VaR}_\alpha(X)]$, where

$$\mathrm{VaR}_\alpha(X) = \sup\{x \mid \mathbb{P}[X \geq x] \geq \alpha\}.$$

We will discuss more examples of coherent risk measures in Chapter 3.

The issue regarding absolute continuity with respect to the reference measure becomes more explicit in the case of continuous probability spaces. Indeed, with more subtle mathematics and appropriate technical conditions, Delbaen [49] extends the main results for coherent risk measures to more general probability spaces.

2.2.4 Convex risk measures

One primary criticism of coherent risk measures is their linear growth with the size of a position. In a financial setting, for instance, one can imagine that liquidity risk may grow nonlinearly with respect to the size of a transaction. To account for this issue with coherent risk measures, Föllmer and Schied [58] and Fritelli and Gianin [62] proposed relaxing the positive homogeneity and subadditivity axioms for coherence, and replacing them with a convexity axiom.

Föllmer and Schied, like Artzner et al., also use acceptance sets as a starting point. They propose the following axioms for acceptance sets:

B1. If $X \in \mathcal{A}$ and $Y \in \mathcal{X}$ satisfies $Y \leq X$, then $Y \in \mathcal{A}$ as well.

B2. If $X \in \mathcal{A}$ and $Y \in \mathcal{X}$, then

$$\{\lambda \in [0,1] \mid \lambda X + (1-\lambda)Y \in \mathcal{A}\}$$

is closed in $[0,1]$.

B3. $\mathcal{A}$ is convex and non-empty.

They then propose the following axioms for risk measures:

Definition 2.2.4. A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is **convex** if it satisfies the following three properties:

(a) (*translation invariance*) $\mu(X+b) = \mu(X) + b, \quad \forall\, b \in \mathbb{R}.$

(b) (*monotonicity*) $\mathbb{P}\{X \leq Y\} = 1 \Rightarrow \mu(X) \leq \mu(Y).$

(c) (*convexity*) $\mu(\lambda X + (1-\lambda)Y) \leq \lambda\mu(X) + (1-\lambda)\mu(Y), \quad \forall\, \lambda \in [0,1].$

The properties of convex risk measures are intuitively appealing. The convexity property can be interpreted as the fact that hedging reduces risk. The monotonicity property enforces the natural requirement that, if one random variable is always less than another one (i.e., incurs less loss), then it is more favorable from a risk standpoint. The translation invariance property says that if we are going to pay a certain fixed penalty b in addition to the risky project X, we are indifferent as to whether we pay before or after X is realized.

Note that, unlike coherent risk measures, convex risk measures need not satisfy $\mu(0) = 0$. As convex risk measures are translation invariant, however, we may always normalize them to satisfy this condition. This will be the case throughout this thesis, and it allows us, again, to interpret a convex risk measure $\mu(X)$ as the minimum additional payment the decision maker requires in order to find the risk X acceptable.

Similar to [3], Föllmer and Schied are able to connect the acceptance sets satisfying B1-B3 to convex risk measures.

Proposition 2.2.2. (Föllmer and Schied, [58]) If $\mu : \mathcal{X} \to \mathbb{R}$ is a convex risk measure, then the induced acceptance set $\mu_{\mathcal{A}}$ satisfies axioms B1-B3, and $\mu_{\mathcal{A}_\mu} = \mu$. Conversely, if $\mathcal{A} \subseteq \mathcal{X}$ is an acceptance set satisfying B1-B3, and the induced risk measure $\mu_{\mathcal{A}}$ satisfies $\mu_{\mathcal{A}}(0) < \infty$, then $\mu_{\mathcal{A}}$ is convex and $\mathcal{A}_{\mu_{\mathcal{A}}} = \mathcal{A}$.

The intuition behind Proposition 2.2.2, like Proposition 2.2.1, is that convex risk measures correspond one-to-one with acceptance sets of random variables. Unlike the coherent risk measures, which correspond to acceptance sets that are positively homogeneous cones, however, the acceptance sets for convex risk measures need only be convex.

Just like with coherent risk measures, convex risk measures have a representation theorem related to expectations under other probability measures. The following result is also shown independently by Heath and Ku [69].

Theorem 2.2.2. *(Föllmer and Schied [58]) Denote the set of all probability measures over Ω by $\mathcal{Q}$. A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is convex if and only if there exists a closed,* *convex function* $\alpha : \mathcal{Q} \to (-\infty, \infty]$ *such that*

$$\mu(X) \;=\; \sup_{Q \in \mathcal{Q}} \left\{ \mathbb{E}_Q [X] - \alpha(Q) \right\}, \qquad \forall\, X \in \mathcal{X}, \tag{2.9}$$

where $\mathbb{E}_Q [X]$ *denotes the expectation of the random variable* X *under the measure* Q.

The interpretation for the representation theorem for convex risk measures is similar to the coherent case, except that, here, we consider *all* measures on Ω, but we "penalize" by the function α. This function will typically be some sort of quasi-distance function from the reference measure $\mathbb{P}$. As we will see, this richer framework has implications for robust optimization models.

It is easy to see [58] that the class of coherent risk measures is a special case of convex risk measures. Indeed, a convex risk measure μ will satisfy positive homogeneity if and only if the penalty function α is an indicator function on set of measures $\mathcal{Q}$, i.e., μ is coherent if and only if

$$\alpha(Q) \;=\; \begin{cases} 0 & \text{if } Q \in \mathcal{Q}, \\ +\infty & \text{otherwise.} \end{cases}$$

In this case, it is easy to see that we obtain

$$\mu(X) \;=\; \sup_{Q \in \mathcal{Q}} \left\{ \mathbb{E}_Q [X] \right\},$$

which, according to Theorem 2.2.1, is a coherent risk measure. We will discuss more examples of penalty functions $\alpha(\cdot)$ and their corresponding convex risk measures in Chapter 4.

Although all of the results presented here assumed knowledge of the reference measure $\mathbb{P}$, work has been done in utility theory by Gilboa and Schmeidler [67] and convex risk measures by Föllmer and Schied [59] in extending these ideas to the case when the underlying measure $\mathbb{P}$ is not known explicitly.

Chapter 3

Robust linear optimization and coherent risk measures

In this chapter, we focus on robust linear optimization problems, which have the form

$$\min\{\boldsymbol{c}'\boldsymbol{x} \mid \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}, \qquad \forall\ \boldsymbol{A} \in \mathcal{U}\}, \tag{3.1}$$

where $\mathcal{U}$ is the corresponding uncertainty set for the uncertain constraint matrix $\boldsymbol{A}$. As we have discussed, the idea of this approach is to compute optimal solutions which retain feasibility for all possible realizations of $\boldsymbol{A}$ within this prescribed uncertainty set $\mathcal{U}$.

The theory of robust optimization, however, is essentially silent on the question of *how* to construct uncertainty sets. As we have seen in Chapter 2, ellipsoidal uncertainty sets, as well as other norms are common in many treatments (e.g., Bertsimas et al. [30]). Although these approaches are often rooted in simple statistical considerations and some probability guarantees have been proven, no theoretical justification is given for such choices.

Here we provide an axiomatic methodology for constructing uncertainty sets within a robust optimization framework for linear optimization problems with uncertain data. We accomplish this by taking as primitive a risk measure on the outcome of an uncertain constraint as well as a finite number of observations of the uncertain data.

For the class of coherent risk measures, we show that this approach leads directly to a robust optimization problem with an explicit, convex uncertainty set whose structure depends on the specific form of the coherent risk measure chosen by the decision maker.

As we will discuss in greater detail shortly, a particularly interesting class of coherent risk measures is one which can be represented as the expected value of a random variable under a distortion of the probability distribution. Such risk measures have been studied widely in actuarial settings; see, e.g., Wang [114], [113]. A result by Schmeidler [105] shows that this class is equivalent to the class of risk measures which are additive under sums of comonotonic random variables. Comonotonicity is an interesting and useful property because it can be used to provide bounds on sums of random variables with arbitrary dependencies (see Dhaene et al. [52], Kaas et al. [74]).

The key results of this chapter are the following:

- Given a coherent risk measure as a primitive, as well as realizations of the uncertain data in the problem, we construct a corresponding convex uncertainty set in a robust optimization framework. This is important as the uncertainty set becomes a *consequence* of the particular risk measure the decision maker selects. In other words, we argue that any robust approach which is to protect against uncertainty should depend intimately on the decision maker's *attitude* towards this uncertainty. When this attitude can be cast as a coherent risk measure, convex uncertainty sets of an explicit construction arise.

- For the important class of coherent risk measures which satisfy comonotonic additivity, we obtain a special and interesting class of polyhedral uncertainty sets. We study this class in detail and show that the entire space of such polyhedral uncertainty sets is, under an appropriate generating mechanism, finitely generated by the class of **conditional tail expectation** risk measures. We also study the sub-class of these polyhedral uncertainty sets which are centrally symmetric and show that they are also finitely generated by a specific set of coherent risk measures. These uncertainty sets are useful because they naturally induce a norm.

- We study a particular class of coherent risk measures based on higher-order tail moments, which are studied by Fischer [56]. These measures lead to l_p-norm uncertainty sets of a particular form and the resulting robust problems are conic optimization problems.

- We consider a converse problem; namely, starting with a convex uncertainty set as primitive, an obvious corollary of the representation theorem is that there is a corresponding coherent risk measure. In addition we illustrate how to compute approximations to arbitrary norm-bounded and polyhedral uncertainty sets from a particular subclass of coherent risk measures satisfying comonotonic additivity.

With regards our data-driven approach, we feel it has the following benefits:

1. It is *tractable*. Indeed, with N observations, our uncertainty sets lead directly to optimization over $\mathbb{R}^N$. As we will see, we can solve these problems efficiently for large N. Moreover, we can make meaningful theoretical statements without relying on involved assumptions and complex results from measure theory.

2. It is *parsimonious*. We are making the bare minimum of distributional assumptions.

3. It is *practical*. In reality we do not have, or there does not exist, a distribution for the uncertain quantities of interest. In practice we usually have *data* and must operate without any other information.

Risk measures and data-driven approaches have received considerable attention from the optimization community recently. Recent work by Ruszczynski and Shapiro [104] considers optimization problems involving risk measures. In many ways, the framework in this chapter can be thought of as a significant generalization of scenario-based approximations to chance constraints (Calafiore and Campi, [41], Nemirovski and Shapiro, [89]). There has been some work on *distributional robustness* in the chance constraint literature (e.g., Erdogan and Iyengar, [54]), but, to the best of our knowledge, the explicit connection between coherent risk measures and convex uncertainty sets in a robust optimization framework has not yet been made.

The outline of the chapter is as follows. Section 3.1 briefly describes some of the relevant classes of coherent risk measures. Section 3.2 considers general coherent risk measures. Section 3.3 considers coherent risk measures which satisfy comonotonic additivity and studies the associated polyhedral uncertainty sets in detail. Section 3.4 deals with the coherent measures based on higher-order tail moments. Section 3.5 considers converse constructions and approximations, and Section 3.6 concludes with a computational example demonstrating our methodology.

3.1 Classes of coherent risk measures

We recall the representation theorem for coherent risk measures, stated below for convenience. It states that we can describe *any* coherent risk measure equivalently in terms of expectations over a family of distributions. The result is largely a consequence the separation theorem for convex sets.

Theorem 3.1.1. *A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is coherent if and only if there exists a family of probability measures $\mathcal{Q}$ that*

$$\mu(X) = \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q [X], \qquad \forall\, X \in \mathcal{X}, \tag{3.2}$$

where $\mathbb{E}_Q [X]$ denotes the expectation of the random variable X under the measure Q (as opposed to the measure of X itself).

Again, the representation theorem says that all coherent risk measures may be represented as the worst-case expected value over a family of "generalized scenarios." This will be the crucial idea as we attempt to construct uncertainty sets in a robust optimization framework from a given coherent risk measure.

3.1.1 Distorted probability measures and the Choquet integral

In this section, we will examine a particularly interesting class of risk measures known as **Choquet integrals**. These measures are expectations of a random variable under an appropriate distortion of the original distribution. It turns out that the Choquet integral is a fairly general risk measure with quite a bit of modelling freedom; in fact, many commonly used risk measures may be cast as a Choquet integral. The Choquet integral has been used extensively for pricing insurance premia (see, e.g., Wang [114], [113]). The name follows from the work on the *theory of capacities* developed by Choquet [45], which would lead in part to a large body of work on belief functions in decision-making (see, e.g., Ngyuen and Ngyuen [91]).

Here we will denote the **de-cumulative distribution function (ddf)** of the random variable $X \in \mathcal{X}$ as $S(x)$; i.e., we have $S(x) = \mathbb{P}\{X \geq x\}$. We give the following definition.

Definition 3.1.1. A **distortion function** g is any non-decreasing function on $[0,1]$ such that $g(0) = 0$ and $g(1) = 1$. The **distorted probability distribution** for a random variable $X \in \mathcal{X}$ with ddf $S(x)$ is the unique[1] distribution defined by the distorted ddf $S^*(x) = g(S(x))$.

We use these distorted distributions to define the Choquet integral; see Denneberg [51] for a more formal definition.

Definition 3.1.2. The **Choquet integral** of a random variable $X \in \mathcal{X}$ with respect to the distortion function g is defined as

$$\mu_g(X) = \int_0^\infty S^*(x)dx + \int_{-\infty}^0 [S^*(x) - 1]\,dx. \tag{3.3}$$

Note that we may write $\mathbb{E}[X] = \int_0^\infty S(x)dx + \int_{-\infty}^0 [S(x) - 1]\,dx$, so the Choquet integral

[1] If g is has discontinuities, we assume the ddf relation holds only at continuity points and thus the distribution is uniquely defined.

is indeed the expected value under the distorted distribution. For any distortion function g, the Choquet integral satisfies translation invariance, monotonicity, and positive homogeneity. We need to say something more about g, however, in order to ensure subadditivity; see Reesor et al. [98] for one proof of the following.

Proposition 3.1.1. The Choquet integral μ_g with a distortion function in equation (3.3) satisfies monotonicity, translation invariance, and positive homogeneity. In addition, μ_g satisfies subadditivity if and only if g is concave. Thus, μ_g is coherent if and only if g is concave.

Examples

In this section, we illustrate the modelling flexibility of the Choquet integral by showing that common risk measures are Choquet integrals of distorted measures. Reesor et al. [98] provide a much more general and extensive list of examples. For simplicity we assume that the random variable X is non-negative, but the examples all extend to the more general case.

Example 3.1.1. *Value-at-risk.*
For some $\alpha \in [0, 1]$ we define

$$g(u) = \begin{cases} 0, & \text{if } u < \alpha, \\ 1, & \text{otherwise.} \end{cases}$$

Then we have

$$\mu_g(X) = \int_0^\infty g(S(x))dx = \int_0^{S^{-1}(\alpha)} dx = S^{-1}(\alpha) = \sup\{x \mid \mathbb{P}[X \geq x] \geq \alpha\},$$

which is commonly referred to as the **value-at-risk at level** α, or $\mathrm{VaR}_\alpha(X)$, in mathematical finance literature. Note that g is not concave, which implies, by Proposition 3.1.1, that value-at-risk is *not* a coherent risk measure.

Example 3.1.2. *Conditional tail expectation (CTE).*

Let us define the distortion function $g(u) = \min(u/\alpha, 1)$ for some $\alpha \in [0, 1]$. Then we have

$$\begin{aligned}
\mu_g(X) &= \int_0^\infty g(S(x))dx \\
&= \frac{1}{\alpha} \int_{S^{-1}(\alpha)}^{\infty} S(x)dx + \int_0^{S^{-1}(\alpha)} dx \\
&= \frac{x}{\alpha} S(x)\Big|_{S^{-1}(\alpha)}^{\infty} - \frac{1}{\alpha} \int_{S^{-1}(\alpha)}^{\infty} x dS(x) + S^{-1}(\alpha) \\
&= \mathbb{E}[X | X \geq S^{-1}(\alpha)] \\
&= \mathbb{E}[X | X \geq \mathrm{VaR}_\alpha(X)].
\end{aligned}$$

Note that this is a coherent risk measure since g is concave. This coherent risk measure is called the **conditional tail expectation of X at level α** and denoted $\mathrm{CTE}_\alpha(X)$. This risk measure is of central importance and has a variety of interesting properties which we will examine and discuss in Section 3.3.

Example 3.1.3. *Proportional hazards distortion.*

Consider the distortion function $g(u) = u^\alpha$. When $\alpha \in [0, 1]$, the distorted risk measure is coherent. The *hazard rate* of t is a function $\theta(t)$ such that

$$e^{\theta(t)} = \mathbb{P}[X \geq t] = S(t).$$

It is easy to see that the hazard function θ^* associated with the distorted distribution satisfies $\theta^*(t) = \alpha\theta(t)$ and hence the term **proportional hazards distortion**. This distorted risk measure has been applied in many insurance applications (see, e.g., Wang [114]). Also note that the limiting case $\alpha = 1$ gives us $\mu_g(X) = \mathbb{E}[X]$, which is clearly a coherent risk measure (albeit not a very conservative one).

3.1.2 Comonotone risk measures

It is not obvious whether or not all coherent risk measures can be represented as a Choquet integral under a concave distortion function. Despite the modelling power of the Choquet framework, the answer is no. We complete our road map of the risk measure world by examining this question in this section. It turns out the key property is a risk measure's behavior under sums of comonotonic random variables. We begin with some definitions.

Definition 3.1.3. The set $A \subseteq \mathbb{R}^n$ is a **comonotonic set** if for all $\boldsymbol{x} \in A$, $\boldsymbol{y} \in A$, we have either $\boldsymbol{x} \leq \boldsymbol{y}$ or $\boldsymbol{y} \leq \boldsymbol{x}$.

Clearly any one-dimensional set is comonotonic. It is also not difficult to see that any set in $\mathbb{R}^n$ with $n > 1$ cannot be full-dimensional and also be comonotonic. We can extend this idea naturally to random variables.

Definition 3.1.4. A random variable $X = (X_1, \ldots, X_n)$ is **comonotonic** if its support $A \subseteq \mathbb{R}^n$ is a comonotonic set.

A simple example of a comonotonic random variable is the joint payoff of a stock and a call option on the stock. Indeed, let S be the stock value at the exercise time, C be the call value, and K be the strike price. Then $C = \max(0, S - K)$. It is easy to see that any pair of payoffs (S_1, C_1), (S_2, C_2) satisfy either $S_1 \geq S_2$ and $C_1 \geq C_2$ or $S_1 \leq S_2$ and $C_1 \leq C_2$, and hence the support of the random variable (S, C) is comonotonic.

Comonotonicity is an important property when considering sums of random variables with arbitrary dependencies. Comonotonic random variables have "worst-case" summation properties among all dependence structures, and, as such, have been used by Dhaene et al. [52] to compute upper bounds on sums of random variables. In our case, comonotonicity has a very specific relationship to coherent risk measures and Choquet integrals. In particular, we are interested in the following property.

Definition 3.1.5. A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ is **comonotonically additive** if, for any comonotonic random variables X and Y, we have

$$\mu(X+Y) \;=\; \mu(X) + \mu(Y).$$

Moreover, we also introduce the following (shorthand) terminology.

Definition 3.1.6. If a coherent risk measure is comonotonically additive, we say the risk measure is **comonotone**.

The following result is due to Schmeidler [105].

Theorem 3.1.2. *A risk measure $\mu : \mathcal{X} \to \mathbb{R}$ can be represented as the Choquet integral with a concave distortion function if and only if μ is comonotone.*

The subadditivity property says we can do no worse by aggregating risk when dealing with a coherent risk measure; Schmeidler's result implies that we will not benefit from diversifying risk when our risk measure is a Choquet integral and the underlying random variables are comonotonic. The theorem also allows us to answer the question of whether all coherent risk measure can be represented in the Choquet integral form. The answer is no, and using Schmeidler's theorem, this can be done by constructing a coherent risk measure which violates comonotonic additivity. See Delbaen [50] for an example of this. In Figure 3-1 we illustrate the landscape of the risk measure universe.

3.2 From coherent risk measures to convex uncertainty sets

In this section, we show how the concepts from risk theory, in particular, coherent risk measures, allow us to construct a meaningful robust counterpart to a linear optimization problem with uncertain data. We will focus on a single constraint of the form $\tilde{\boldsymbol{a}}'\boldsymbol{x} \leq b$, but the results all extend in straightforward fashion to the case with $m > 1$ such constraints.

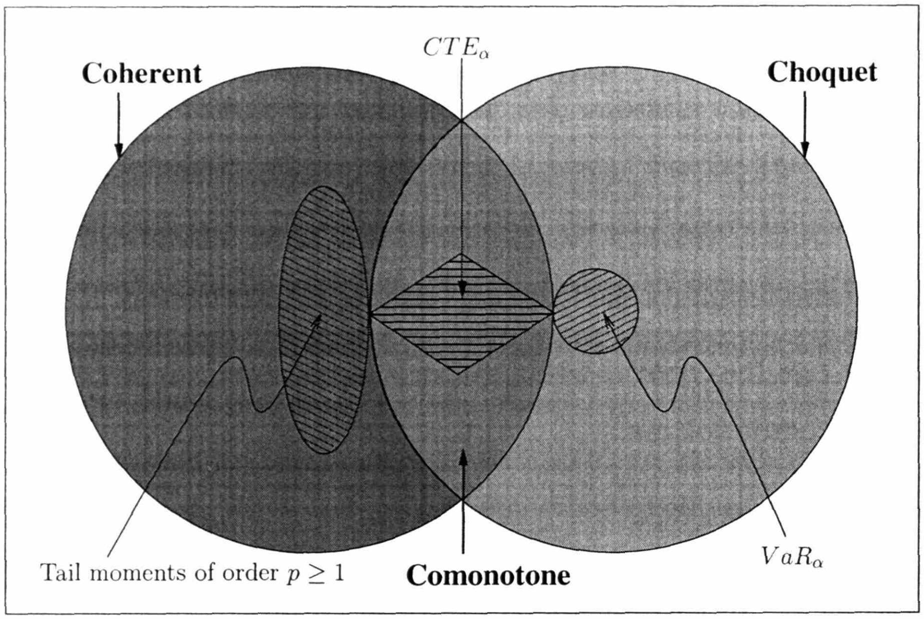

Figure 3-1: *Venn diagram of the risk measure universe. The box represents all functions $\mu : \mathcal{X} \to \mathbb{R}$. In bold are the three main classes of risk measures: coherent, distorted, and their intersection, comonotone. Also illustrated are the specific subclasses CTE_α, VaR_α, and tail moments of higher order (discussed in Section 3.4). Note that these subclasses intersect at limiting values of the various parameters.*

We note the following issues in a practical context.

(1) We generally do not know the distribution of $\tilde{\boldsymbol{a}}$. In fact, we usually only have some finite number N of *observations* of the uncertain vector $\tilde{\boldsymbol{a}}$.

(2) Even equipped with a perfect description of the distribution of $\tilde{\boldsymbol{a}}$, it is not clear how we should construct an uncertainty set $\mathcal{U}$ with some specific, desirable properties.

To address the first issue, we will make the following assumption.

Assumption 3.2.1. The uncertain vector $\tilde{\boldsymbol{a}}$ is a random variable in $\mathbb{R}^n$ satisfying $|\text{supp}\,(\tilde{\boldsymbol{a}})\,| = N$.

Remark 3.2.1. We will typically refer to $\mathcal{A}$ as the **data** or **data set** of the problem. In some cases, it will also be convenient to use the matrix form $\boldsymbol{A} = [\boldsymbol{a}_1 \cdots \boldsymbol{a}_N]$. Also, where convenient, we denote $\mathcal{N} = \{1, \ldots, N\}$.

Thus, we assume that the sample space is confined to $\{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N\}$, and $\tilde{\boldsymbol{a}}$ is distributed across these N values. Although this may seem restrictive, it is in many ways sensible as the data is the *only information* we have about the distribution of $\tilde{\boldsymbol{a}}$.

For the second issue, we take as primitive a coherent risk measure. The choice of this risk measure clearly depends on the preferences of the decision maker. Given a constraint based on this coherent risk measure and distribution defined as above, we will show that there exists an equivalent robust optimization problem with a unique convex uncertainty set. We first define more formally the problem in question.

Definition 3.2.1. For a linear optimization problem with uncertain data $\tilde{\boldsymbol{a}}$ and scalar b, along with a risk measure μ, we define the **risk averse problem** to be

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b. \end{aligned} \tag{3.4}$$

Note that when μ satisfies translation invariance (as all coherent risk measures must), the constraint in (3.4) is equivalent to the constraint $\mu(\tilde{\boldsymbol{a}}'\boldsymbol{x} - b) \leq 0$. We have the first result, which stems directly from the representation theorem (Theorem 3.1.1).

Theorem 3.2.1. *If the risk measure μ is coherent and $\tilde{\boldsymbol{a}}$ is distributed as in Assumption 3.2.1, then the risk averse problem (3.4) is equivalent to the robust optimization problem*

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{a}'\boldsymbol{x} \leq b, \qquad \forall\, \boldsymbol{a} \in \mathcal{U}, \end{aligned} \tag{3.5}$$

where

$$\mathcal{U} = \text{conv}\left(\left\{ \sum_{i=1}^{N} q_i \boldsymbol{a}_i \;\middle|\; \boldsymbol{q} \in \mathcal{Q} \right\}\right) \subseteq \text{conv}\left(\mathcal{A}\right),$$

and $\mathcal{Q}$ is the set of generating measures for μ from Theorem 3.1.1.

Proof. μ is coherent, so, by Theorem 3.1.1 and the fact that $\tilde{\boldsymbol{a}}$ is distributed uniformly on $\mathcal{A}$, we have

$$\mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) = \sup_{\boldsymbol{q}\in\mathcal{Q}} \mathbb{E}_{\boldsymbol{q}}\left[\tilde{\boldsymbol{a}}'\boldsymbol{x}\right] = \sup_{\boldsymbol{q}\in\mathcal{Q}} \sum_{i=1}^{N} (\boldsymbol{a}_i'\boldsymbol{x}) q_i = \sup_{\boldsymbol{q}\in\mathcal{Q}} \left(\sum_{i=1}^{N} q_i \boldsymbol{a}_i\right)' \boldsymbol{x} = \sup_{\boldsymbol{a}\in\tilde{\mathcal{U}}} \boldsymbol{a}'\boldsymbol{x} = \sup_{\boldsymbol{a}\in\mathcal{U}} \boldsymbol{a}'\boldsymbol{x},$$

where $\tilde{\mathcal{U}} = \{\sum_{i=1}^{N} q_i \boldsymbol{a}_i \mid \boldsymbol{q} \in \mathcal{Q}\}$ and $\mathcal{U}$ is as in the statement of the theorem, i.e., $\mathcal{U} = \text{conv}(\tilde{\mathcal{U}})$. Note that the last line follows from the simple observation that the supremum of a linear function over any bounded set is equal to the supremem of that function over the convex hull of that set. Hence the risk averse problem (3.4) is equivalent to the robust optimization problem (3.5) with uncertainty set given as in the statement of the theorem. □

Theorem 3.2.1 provides a methodology for constructing robust optimization problems with uncertainty sets possessing a direct, physical meaning. The decision maker has some risk measure μ which depends on their preferences. If μ is coherent, there is an *explicit* uncertainty set that should be used in the robust optimization framework. This uncertainty set is convex and its structure depends on the generating family $\mathcal{Q}$ for μ and the data $\mathcal{A}$. We emphasize that the reference measure $\mathbb{P}$ need not play any explicit role in the description of $\mathcal{Q}$. Indeed, as discussed in Chapter 2, $\mathcal{Q}$ may simply be a finite set of measures. Of course, as also discussed in Chapter 2, many risk measures (such as CVaR) have a generating family $\mathcal{Q}$ which is parameterized by the reference measure $\mathbb{P}$. So, whether or not the structure of the uncertainty set in the robust optimization framework depends on $\mathbb{P}$ will be determined by the underlying risk measure that the decision maker uses.

The remainder of this paper primarily focuses on classes of coherent risk measures which give rise to uncertainty sets with special structure.

3.3 From comonotone measures to polyhedral uncertainty sets

For an arbitrary coherent risk measure μ, the uncertainty set in the corresponding robust optimization problem depends crucially on the generator $\mathcal{Q}$ of μ given by Theorem 3.1.1. In general it is not obvious how to compute this generator. We now show how to construct $\mathcal{Q}$ in the important case where μ is a comonotone risk measure (i.e., the Choquet integral of a concave distortion function) and find that the resulting uncertainty set is a polyhedron. We first note the following simple observation.

Lemma 3.3.1. *If $\mathcal{Q}$ is a finite set, then we have the following:*

$$\sup_{Q \in \mathcal{Q}} \mathbb{E}_Q[X] = \max_{Q \in \mathcal{Q}} \mathbb{E}_Q[X] = \max_{Q \in \mathrm{conv}(\mathcal{Q})} \mathbb{E}_Q[X], \tag{3.6}$$

where $\mathrm{conv}(\mathcal{Q})$ *denotes the convex hull of $\mathcal{Q}$.*

Proof. The first equality is obvious, since $\mathcal{Q}$ is a closed set. The second equality follows from the fact that the maximum of a linear function over a polytope has an optimal solution at a vertex, and since the vertices of $\mathrm{conv}(\mathcal{Q})$ are all elements of $\mathcal{Q}$, the result follows. □

For simplicity, in what follows, we will rely on a stronger assumption regarding the distribution of $\tilde{\boldsymbol{a}}$.

Assumption 3.3.1. In addition to Assumption 3.2.1, the random variable $\tilde{\boldsymbol{a}}$ satisfies $\mathbb{P}\{\tilde{\boldsymbol{a}} = \boldsymbol{a}_i\} = 1/N$ for $i = 1, \ldots, N$.

Assumption 3.3.1 simply allows for a cleaner description of the generating families $\mathcal{Q}$ for the coherent risk measures we will discuss. The results may be generalized for an arbitrary, discrete distribution vector $\boldsymbol{p} \in \Delta^N$, but we omit this extension.
We first show how to calculate μ_g under Assumption 3.3.1.

Proposition 3.3.1. For a comonotone risk measure with distortion function g on a random variable Y with support $\{y_1, \ldots, y_N\}$ such that $\mathbb{P}[Y = y_i] = 1/N$, we have

$$\mu_g(Y) = \sum_{i=1}^{N} q_i y_{(i)}, \tag{3.7}$$

where $y_{(i)}$ is the ith order statistic of Y, i.e., $y_{(1)} \leq \ldots \leq y_{(N)}$, and

$$q_i = g\left(\frac{N+1-i}{N}\right) - g\left(\frac{N-i}{N}\right). \tag{3.8}$$

Proof. We assume without loss of generality that $y_i = y_{(i)}$ for all $i \in \{1, \ldots, N\}$. We show the result for the case $y_1 \geq 0$, and the general case follows by straightforward extension. We first note that we can write the de-cumulative distribution of Y as

$$S_Y(y) = \begin{cases} 1, & \text{if } y < y_1, \\ \dfrac{N-i}{N}, & \text{if } y_i \leq y < y_{i+1},\ i = 1, \ldots, N-1, \\ 0, & \text{otherwise.} \end{cases}$$

Now, applying Equation (3.3), we have

$$\begin{aligned} \mu_g(Y) &= \int_0^\infty g\left(S_Y(y)\right) dy \\ &= \int_0^{y_1} g(1) dy + \sum_{i=1}^{N-1} \int_{y_i}^{y_{i+1}} g\left(\frac{N-i}{N}\right) dy \\ &= g(1) y_1 + \sum_{i=1}^{N-1} g\left(\frac{N-i}{N}\right)(y_{i+1} - y_i) \\ &= \sum_{i=1}^{N}\left(g\left(\frac{N-i+1}{N}\right) - g\left(\frac{N-i}{N}\right)\right) y_i \\ &= \sum_{i=1}^{N} q_i y_i \\ &= \sum_{i=1}^{N} q_i y_{(i)}. \end{aligned}$$

□

Remark 3.3.1. Note that, for q_i defined in equation (3.8), $q_1 \leq \cdots \leq q_N$. This is easy to see from the fact that g is nondecreasing and concave.

Thus, in order to compute μ_g for a comonotone risk measure, we need an order statistic on the N possible values the random variable Y (in the context we are considering, we have $y_i = \boldsymbol{a}_i'\boldsymbol{x}$). We now argue that this ordering can be implicitly done via solution of a linear optimization problem.

Proposition 3.3.2. For a concave distortion function g and with q_i defined as in Equation (3.8), we have

$$\sum_{i=1}^{N} q_i y_{(i)} = z^*_{\boldsymbol{q},\boldsymbol{y}}, \tag{3.9}$$

where $z^*_{\boldsymbol{q},\boldsymbol{y}}$ is the optimal value of the linear programming problem

$$\begin{array}{ll} \text{maximize} & \sum_{i=1}^{N} q_i \sum_{j=1}^{N} w_{ij} y_j \\ \text{subject to} & \boldsymbol{w} \in W(N), \end{array} \tag{3.10}$$

in variables w_{ij}, and where $W(N)$ is the assignment polytope in N dimensions, i.e.,

$$W(N) = \left\{ \boldsymbol{w} \in \mathbb{R}_+^{N^2} \;\middle|\; \sum_{i=1}^{N} w_{ij} = 1 \; \forall j \in \{1, \ldots, N\}, \; \sum_{j=1}^{N} w_{ij} = 1 \; \forall i \in \{1, \ldots, N\} \right\}.$$

Proof. As indicated in Remark 3.3.1, we have $q_i \leq q_{i+1}$ for all $i \in \{1, \ldots, N-1\}$, so the sum in Equation (3.9) is sorting the y_i in increasing order and assigning the largest y_i to q_N, the second largest to q_{N-1}, and so on.

We assume without loss of generality that the y_i are already ordered, i.e., $y_{(i)} = y_i$ for all $i \in \mathcal{N}$. Then the solution

$$w_{ij} = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{otherwise,} \end{cases}$$

is feasible to problem (3.10) with the same value as the sum. Thus, $\sum_{i=1}^{N} q_i y_{(i)} \leq z^*_{\boldsymbol{q},\boldsymbol{y}}$. Conversely, given any feasible $\boldsymbol{w}$ to (3.10), and noting that $q_i \leq q_{i+1}$, we have

$$\sum_{i=1}^{N} q_i \sum_{j=1}^{N} w_{ij} y_j \leq \sum_{i=1}^{N} q_i y_i = \sum_{i=1}^{N} q_i y_{(i)},$$

which shows that $z^*_{\boldsymbol{q},\boldsymbol{y}} \leq \sum_{i=1}^{N} q_i y_{(i)}$ and we are done. □

We see that the formulation in Proposition 3.3.2 explicitly defines a generator $\mathcal{Q}$ for μ_g. Indeed, we may write the generator here as

$$\mathcal{Q} = \left\{ \sum_{j=1}^{N} w_{ij} q_i \;\middle|\; \boldsymbol{w} \in W(N) \right\},$$

or, alternatively,

$$\mathcal{Q} = \{\boldsymbol{p} \in \mathbb{R}^N \mid \exists\, \sigma \in \mathrm{S}_N \;:\; p_i = q_{\sigma(i)},\ \forall\, i \in \mathcal{N}\}, \tag{3.11}$$

where S_N is the symmetric group on N elements. This second description (3.11) is valid by noting the well-known result that $W(N)$ has integral extreme points and Lemma 3.3.1.

Remark 3.3.2. Although an arbitrary coherent risk measure requires a family $\mathcal{Q}$ of measures to generate it, we see that a comonotone risk measure is effectively generated by a *single* measure $\boldsymbol{q}$ as given by Equation (3.8). When speaking of comonotone risk measures, we will refer to $\boldsymbol{q}$ as the **generator of** μ_g, with the understanding that $\mathcal{Q}$ described above in equation (3.11) is the actual generating family.

So comonotone risk measures lead to a very special class of polyhedral uncertainty sets: the convex hull of all $N!$ convex combinations of $\mathcal{A}$ induced by all permutations of the generator $\boldsymbol{q}$ (from Equation (3.8)). This class of polytopes has a very special structure, so we define it formally. From here on out we denote the vector of ones by

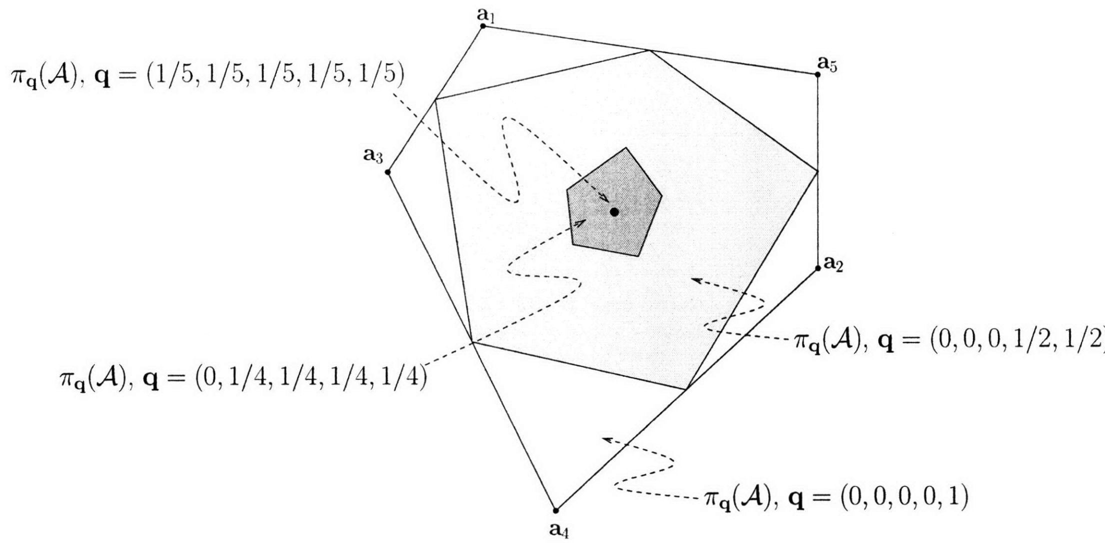

Figure 3-2: $\pi_{\boldsymbol{q}}(\mathcal{A})$ *for various* $\boldsymbol{q}$ *for an example with* $N = 5$.

$\boldsymbol{e}$ and the N-dimensional probability simplex by Δ^N, i.e.,

$$\Delta^N = \{\boldsymbol{p} \in \mathbb{R}^N_+ \mid \boldsymbol{e}'\boldsymbol{p} = 1\}.$$

Definition 3.3.1. For some measure $\boldsymbol{q} \in \Delta^N$ and discrete set $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N\}$ with $\boldsymbol{x}_i \in \mathbb{R}^n$ for all $i \in \{1, \ldots, N\}$, we define the **$\boldsymbol{q}$-permutohull** of X by

$$\pi_{\boldsymbol{q}}(X) \quad = \quad \operatorname{conv}\left(\left\{\sum_{i=1}^{N} q_{\sigma(i)}\boldsymbol{x}_i \;\middle|\; \sigma \in \mathrm{S}_N\right\}\right). \tag{3.12}$$

If $\boldsymbol{e}_i$ is any unit vector in $\mathbb{R}^N$, then $\pi_{\boldsymbol{e}_i}(X) = \operatorname{conv}(X)$. We also see that $\pi_{\boldsymbol{e}/N}(X) = \{1/N \sum_{i=1}^N \boldsymbol{x}_i\}$, i.e., the sample mean of X, where $\boldsymbol{e}$ is the vector of ones in $\mathbb{R}^N$. Note also the difference between a $\boldsymbol{q}$-permutohull and a **permutohedron** (see Ziegler [120]), which is the convex hull of all permutations of the vector $[1 \ \cdots \ N]$.

A simple illustration of $\pi_{\boldsymbol{q}}(\mathcal{A})$ is given in Figure 3-2. We are ready for the main result of this section.

Theorem 3.3.1. *For a risk averse problem with comonotone risk measure μ_g generated by a measure $\boldsymbol{q} \in \Delta^N$ and uncertain vector $\tilde{\boldsymbol{a}}$ distributed as given by Assumption 3.3.1, problem (3.4) is equivalent to a robust optimization problem with uncertainty* *set given by the polyhedron* $\pi_{\boldsymbol{q}}(\mathcal{A})$*, i.e., (3.4) is equivalent to*

$$\begin{array}{lll} \text{minimize} & \boldsymbol{c}'\boldsymbol{x} & \\ \text{subject to} & \boldsymbol{a}'\boldsymbol{x} \leq b, & \forall\, \boldsymbol{a} \in \pi_{\boldsymbol{q}}(\mathcal{A}). \end{array} \tag{3.13}$$

Moreover, (3.13) is equivalent to the linear optimization problem

$$\begin{array}{lll} \text{minimize} & \boldsymbol{c}'\boldsymbol{x} & \\ \text{subject to} & \boldsymbol{e}'\boldsymbol{y}_1 + \boldsymbol{e}'\boldsymbol{y}_2 \leq b, & \\ & \boldsymbol{e}_i'\boldsymbol{y}_1 + \boldsymbol{e}_j'\boldsymbol{y}_2 \geq q_i \cdot \boldsymbol{a}_j'\boldsymbol{x} & \forall\, (i,j) \in \mathcal{N} \times \mathcal{N}, \end{array} \tag{3.14}$$

in decision variables $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{y}_1 \in \mathbb{R}^N$, $\boldsymbol{y}_2 \in \mathbb{R}^N$.

Proof. Equivalence to (3.4) follows by applying Theorem 3.1.1 and Proposition 3.3.2 (as well as the remarks following it) to the risk averse problem (3.4). Since $\pi_{\boldsymbol{q}}(\mathcal{A})$ is a polyhedron, we know, by standard robust optimization results, that (3.4) is a linear optimization problem. To get the specific form of this problem, consider the problem, for any $\boldsymbol{x}$, $\max\limits_{\boldsymbol{a} \in \pi_{\boldsymbol{q}}(\mathcal{A})} \boldsymbol{x}'\boldsymbol{a}$. This is equivalent to the problem

$$\begin{array}{ll} \text{maximize} & \sum\limits_{i,j} q_i \cdot (\boldsymbol{a}_j'\boldsymbol{x}) \cdot w_{ij} \\ \text{subject to} & \boldsymbol{w} \in W(N), \end{array} \tag{3.15}$$

in variables $w_{ij} \in \mathbb{R}^{N^2}$. The dual problem is

$$\begin{array}{ll} \text{minimize} & \boldsymbol{e}'\boldsymbol{y}_1 + \boldsymbol{e}'\boldsymbol{y}_2 \\ \text{subject to} & \boldsymbol{e}_i'\boldsymbol{y}_1 + \boldsymbol{e}_j'\boldsymbol{y}_2 \geq q_i \cdot \boldsymbol{a}_j'\boldsymbol{x}, \end{array} \tag{3.16}$$

in variables $\boldsymbol{y}_1 \in \mathbb{R}^N$, $\boldsymbol{y}_2 \in \mathbb{R}^N$. Since strong duality holds between (3.15) and (3.16) (as (3.15) has a nonempty, bounded feasible set), we may replace left-hand side of the constraint in (3.13) with the objective from (3.16) and add in the dual constraints as well, leaving us with the desired result. □

We remark that although $\pi_{\boldsymbol{q}}(\mathcal{A})$ has as many as $N!$ extreme points, the complexity of using it as an uncertainty set is polynomial in N. Specifically, the equivalent problem (3.14) has only $2N$ extra variables and N^2 extra constraints.

3.3.1 Structure of the comonotone family

It is interesting and relevant to explore the properties of the family of uncertainty sets $\pi_{\boldsymbol{q}}(\mathcal{A})$ over appropriate generating measures $\boldsymbol{q}$. We will find that, under an appropriate generating mechanism, the set of all comonotone risk measures is finitely generated by the class of conditional tail expectation measures ($\mathrm{CTE}_{i/N}(\cdot)$). *In a very direct sense, then, the* $\mathrm{CTE}_{i/N}(\cdot)$ *measures are basis measures for the entire space of comonotone risk measures on random variables with a discrete state space of cardinality* N*.*

We first notice that, as a comonotone risk measure produces generators with $q_i \leq q_{i+1}$ for all $i \in \{1, \ldots, N-1\}$, the space of possible generators $\boldsymbol{q}$ is a strict subset of the N-dimensional unit simplex.

Definition 3.3.2. The **restricted simplex in N-dimensions** is denoted by $\tilde{\Delta}^N$ and defined as

$$\tilde{\Delta}^N = \left\{ \boldsymbol{q} \in \Delta^N \mid q_1 \leq \cdots \leq q_N \right\}. \tag{3.17}$$

It is easy to see that there is a one-to-one correspondence between $\tilde{\Delta}^N$ and the space of comonotone risk measures on random variables distributed on N values, as we now show.

Proposition 3.3.3. There exists a bijection between $\tilde{\Delta}^N$ and the space of comonotone risk measures on random variables with a finite sample space of cardinality N.

Proof. Clearly, any such comonotone risk measure defines a $\boldsymbol{q} \in \tilde{\Delta}^N$ via Equation (3.8). Conversely, given any $\boldsymbol{q} \in \tilde{\Delta}^N$, we may define a distortion function (on N

values) as

$$
\begin{aligned}
g(0) &= 0, \\
g\left(\frac{i}{N}\right) &= \sum_{j=1}^{i} q_{N-j+1}, \qquad i = 1, \ldots, N.
\end{aligned}
$$

One can easily verify that such a g satisfies:

$$
\begin{aligned}
g(1) &= 1, \\
g(i/N) &\geq g((i-1)/N) \qquad \forall\, i \in \mathcal{N}, \\
g(i/N) - g((i-1)/N) &\leq g((i-1)/N) - g((i-2)/N) \qquad \forall\, i \in \{2, \ldots, N\},
\end{aligned}
$$

so g is a valid distortion function corresponding to a comonotone risk measure. □

The restricted simplex and the space of comonotone risk measures, then, are isomorphic. We now argue that a very special class of risk measures generates the entire space of comonotone risk measures.

Theorem 3.3.2. *The restricted simplex $\tilde{\Delta}^N$ is generated by the N-member family*

$$
\mathcal{G}_N = \{\boldsymbol{q} \in \tilde{\Delta}^N \mid \exists\, k \in \mathcal{N} :\ q_i = 0\ \forall\, i \leq N-k,\ q_i = 1/k\ \forall\, i > N-k\}, \quad (3.18)
$$

i.e., $\operatorname{conv}(\mathcal{G}_N) = \tilde{\Delta}^N$. *Moreover, each $\hat{\boldsymbol{q}} \in \mathcal{G}_N$ corresponds to the measure* $\mathrm{CTE}_{i/N}(\cdot)$ *for some $i \in \mathcal{N}$.*

Proof. Clearly $\operatorname{conv}(\mathcal{G}_N) \subseteq \tilde{\Delta}^N$, as $\tilde{\Delta}^N$ is convex and any $\hat{\boldsymbol{q}} \in \mathcal{G}_N$ also satisfies $\hat{\boldsymbol{q}} \in \tilde{\Delta}^N$ by construction. Now consider any $\boldsymbol{q} \in \tilde{\Delta}^N$. We write the matrix with columns the members of $\mathcal{G}_N$ as

$$
\boldsymbol{G}_N = \begin{bmatrix}
\frac{1}{N} & 0 & 0 & \cdots & 0 \\
\frac{1}{N} & \frac{1}{N-1} & 0 & \cdots & 0 \\
\frac{1}{N} & \frac{1}{N-1} & \frac{1}{N-2} & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
\frac{1}{N} & \frac{1}{N-1} & \frac{1}{N-2} & \cdots & 1
\end{bmatrix},
$$

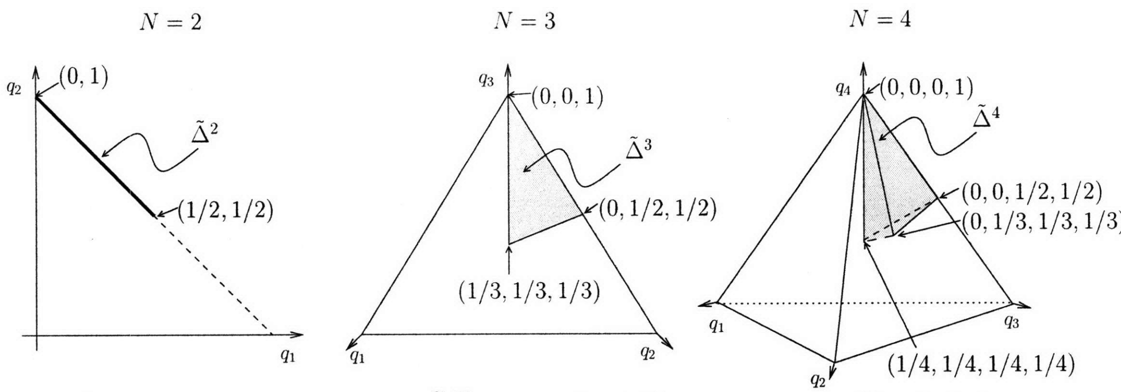

Figure 3-3: *Generation of* $\tilde{\Delta}^N$ *by the* CTE $(\cdot)$ *measures for* $N = 2, 3, 4$.

and define the vector $\boldsymbol{\lambda} \in \mathbb{R}^N$ as $\lambda_1 = Nq_1$, $\lambda_i = (N - i + 1)(q_i - q_{i-1})$, for all $i \in \{2, \ldots, N\}$. We have $\boldsymbol{q} \in \tilde{\Delta}^N$, so $q_{i+1} \geq q_i$ and thus $\boldsymbol{\lambda} \geq \mathbf{0}$. In addition,

$$\sum_{i=1}^{N} \lambda_i = Nq_1 + \sum_{i=2}^{N} (N - i + 1)(q_i - q_{i-1}) = \sum_{i=1}^{N} q_i = 1.$$

Finally, we compute the vector $\boldsymbol{p} = \boldsymbol{G}_N \boldsymbol{\lambda}$ and see that

$$p_i = q_1 + \sum_{j=2}^{i} \frac{\lambda_j}{N - j + 1} = q_1 + \sum_{j=2}^{i} \frac{(N - j + 1)}{(N - j + 1)} (q_j - q_{j-1}) = q_i,$$

so $\boldsymbol{q} \in \operatorname{conv}(\mathcal{G}_N)$, implying that $\tilde{\Delta}^N \subseteq \operatorname{conv}(\mathcal{G}_N)$.

To see that the N members of $\mathcal{G}_N$ indeed correspond to a risk measure $\mathrm{CTE}_{i/N}(\cdot)$ for some i, note that $\boldsymbol{q} \in \mathcal{G}_N$ implies $\boldsymbol{q} = \{0, \ldots, 0, 1/k, \ldots, 1/k\}$ for some $k \in \mathcal{N}$. We construct the corresponding distortion function and find $g(i/N) = \min(i/k, 1)$ for all $i \in \{0, \ldots, N\}$. Now, by the simple calculations in Example 3.1.2, we see that this function corresponds to $\mathrm{CTE}_{k/N}(\cdot)$. □

Figure 3-3 highlights the result of Theorem 3.3.2. The $\mathrm{CTE}_\alpha(\cdot)$ measures are, in a sense, fundamental; this is consistent with other work, such as Delbaen [49], who shows that $\mathrm{CTE}_\alpha(\cdot)$ is the smallest distribution invariant coherent risk measure greater than $\mathrm{VaR}_\alpha(\cdot)$. Nemirovski and Shapiro [90] also illustrate this idea in the context of convex approximations to chance-constrained problems. $\mathrm{CTE}_\alpha(\cdot)$ is also related closely to the concept of **shortfall**, which is just the mean plus the $\mathrm{CTE}_\alpha(\cdot)$.

Shortfall has been studied extensively in the mathematical finance community, see, e.g., Bertsimas et al. [28] for a treatment exploring a number of its special properties.

3.3.2 Comonotone measures with centrally symmetric uncertainty sets

In this section, we study a more restricted class of generating measures $\boldsymbol{q}$; specifically, we study those that lead to polyhedral uncertainty sets obeying a specific symmetry property. These structures are important because they naturally induce norm spaces which we will use in the next section to approximate arbitrary polyhedral uncertainty sets by those corresponding to comonotone risk measures.

Definition 3.3.3. A set P is **centrally symmetric around $\boldsymbol{x}_0 \in P$** if $\boldsymbol{x}_0 + \boldsymbol{x} \in P$ implies $\boldsymbol{x}_0 - \boldsymbol{x} \in P$.

Here we will be interested in uncertainty sets which are symmetric around the sample mean of the data. In the space of discrete measures of dimension N, this is equivalent to symmetry around the measure $\boldsymbol{q} = \boldsymbol{e}/N$.

Definition 3.3.4. The **symmetric restricted simplex in N-dimensions** is denoted by $\tilde{\Delta}^N_{\text{sym}}$ and defined by

$$\tilde{\Delta}^N_{\text{sym}} = \left\{ \boldsymbol{q} \in \tilde{\Delta}^N \mid \pi_{\boldsymbol{q}}\left(\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\}\right) \text{ centrally symmetric around } \boldsymbol{e}/N \right\}, \quad (3.19)$$

where the $\boldsymbol{e}_i$ are the unit basis vectors, $\boldsymbol{e}$ is the vector of ones, and $\pi_{\boldsymbol{q}}(\cdot)$ is defined in (3.12).

We begin with the following simple observation.

Proposition 3.3.4. For a vector $\boldsymbol{q} \in \tilde{\Delta}^N$, we have $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$ if and only if $\tilde{q}_i = q_{\sigma(i)}$, where σ is a permutation such that $\sigma(i) = N - i + 1$ and $\tilde{\boldsymbol{q}}$ is the vector

$$\tilde{\boldsymbol{q}} = \frac{2}{N}\boldsymbol{e} - \boldsymbol{q}. \quad (3.20)$$

Proof. We first note that clearly $\boldsymbol{e}/N \in \pi_{\boldsymbol{q}}(\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\})$. Indeed, taking all $N!$ permutations of $\boldsymbol{q}$ and summing their average, we see by symmetry that the result is $\boldsymbol{e}/N$. Now $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$ if and only if every extreme point of $\pi_{\boldsymbol{q}}(\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\})$ satisfies the symmetry property. Furthermore, it is sufficient to check the extreme point $\boldsymbol{q} \in \tilde{\Delta}^N$, as all other extreme points of $\pi_{\boldsymbol{q}}(\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\})$ are just permutations of this and the arguments follow through by symmetry. We can write $\boldsymbol{q} = \boldsymbol{e}/N + \boldsymbol{p}$ for some vector $\boldsymbol{p} \in \mathbb{R}^N$. We need to check that $\tilde{\boldsymbol{q}} = \boldsymbol{e}/N - \boldsymbol{p} \in \pi_{\boldsymbol{q}}(\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_N\})$ as well. Substituting for $\boldsymbol{p}$ we have $\tilde{\boldsymbol{q}} = 2/N\boldsymbol{e} - \boldsymbol{q}$. But since $\boldsymbol{q}$ is an extreme point, we must have $\tilde{\boldsymbol{q}}$ as an extreme point as well, which means that $\tilde{\boldsymbol{q}}$ must be permutation of $\boldsymbol{q}$. Moreover, as $\boldsymbol{q} \in \tilde{\Delta}^N$, its components are non-decreasing, which implies that the components of $\tilde{\boldsymbol{q}}$ are non-increasing. It is then easy to see that $\tilde{\boldsymbol{q}}$ must be the specified permutation σ of $\boldsymbol{q}$. □

We now prove that, like the restricted simplex, $\tilde{\Delta}^N_{\text{sym}}$ is generated by a finite family of comonotone risk measures.

Theorem 3.3.3. *The symmetric restricted simplex $\tilde{\Delta}^N_{\text{sym}}$ is generated by the $\hat{N} = (\lfloor N/2 \rfloor + 1)$-member family*

$$\mathcal{S}_N = \{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{\hat{N}}\}, \tag{3.21}$$

i.e., $\operatorname{conv}(\mathcal{S}_N) = \tilde{\Delta}^N_{\text{sym}}$*, where the vectors $\boldsymbol{s}_i$ satisfy*

$$s_{i,j} = \begin{cases} 0, & j < i, \\ \dfrac{1}{N}, & i \le j \le N - i + 1, \\ \dfrac{2}{N}, & j > N - i + 1. \end{cases} \tag{3.22}$$

Proof. We assume N is even; the proof for N odd is nearly identical with some small changes on summation limits. We define the matrix $\boldsymbol{S}_N \in \mathbb{R}^{N \times \hat{N}}$ by

$$\boldsymbol{S}_N = \begin{bmatrix} \boldsymbol{s}_1 & \cdots & \boldsymbol{s}_{\hat{N}} \end{bmatrix}.$$

and consider an arbitrary convex combination $\boldsymbol{\lambda} \in \mathbb{R}^{\hat{N}}$ of the columns of $\boldsymbol{S}_N$, i.e., a vector $\boldsymbol{q} \in \mathbb{R}^N$ such that $\boldsymbol{q} = \boldsymbol{S}_N \boldsymbol{\lambda}$. We find that

$$q_i = \begin{cases} \dfrac{1}{N} \displaystyle\sum_{j=1}^{i} \lambda_j, & i < \hat{N}, \\ \dfrac{1}{N} + \dfrac{1}{N} \displaystyle\sum_{j=N+2-i}^{\hat{N}} \lambda_j, & i \geq \hat{N}. \end{cases}$$

We see that we can rearrange this to find that

$$\begin{aligned} q_1 &= \lambda_1 / N, \\ q_i &= q_{i-1} + \lambda_i / N \qquad i < \hat{N}, \\ q_{\hat{N}} &= 1/N + \lambda_{\hat{N}} / N, \\ q_i &= q_{i-1} + \lambda_{N+2-i} / N \qquad i > \hat{N}, \end{aligned} \tag{3.23}$$

from which it follows that $\boldsymbol{q} \geq \boldsymbol{0}$ and $q_i \leq q_{i+1}$. Combining this with the observation that the rows of $\boldsymbol{S}_N$ each sum to one, we see that $\sum_{i=1}^{N} q_i = \sum_{i=1}^{\hat{N}} \lambda_i = 1$, so $\boldsymbol{q} \in \tilde{\Delta}^N$. We now check the symmetry condition from Proposition 3.3.4. We have

$$\begin{aligned} \tilde{\boldsymbol{q}} &= \frac{2}{N} \boldsymbol{e} - \boldsymbol{q} \\ &= \begin{cases} \dfrac{2}{N} - \dfrac{1}{N} \displaystyle\sum_{j=1}^{i} \lambda_j, & i < \hat{N}, \\ \dfrac{1}{N} - \dfrac{1}{N} \displaystyle\sum_{j=N+2-i}^{\hat{N}} \lambda_j, & i \geq \hat{N}, \end{cases} \\ &= \begin{cases} \dfrac{1}{N} + \dfrac{1}{N} \displaystyle\sum_{j=i+1}^{\hat{N}} \lambda_j, & i < \hat{N}, \\ \dfrac{1}{N} \displaystyle\sum_{j=1}^{N+1-i} \lambda_j, & i \geq \hat{N}, \end{cases} \end{aligned}$$

and we see that $\tilde{q}_i = q_{\sigma(i)}$ under the permutation $\sigma(i) = N - i + 1$ for all $i \in \mathcal{N}$. This shows that $\operatorname{conv}(\mathcal{S}_N) \subseteq \tilde{\Delta}^N_{\text{sym}}$.

For the reverse inclusion, consider any $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$. From the permutation in Proposition 3.3.4, we see that we must have $q_{\hat{N}-1} + q_{\hat{N}} = 2/N$, and since $\boldsymbol{q} \in \tilde{\Delta}^N$, we have $q_{\hat{N}} \geq q_{\hat{N}-1}$. Taken together with the previous remark, this implies that $q_{\hat{N}} \geq 1/N$. We now construct a $\boldsymbol{\lambda} \in \mathbb{R}^{\hat{N}}$ by reversing the construction above in (3.23). This leads us to

$$\begin{aligned} \lambda_1 &= Nq_1, \\ \lambda_i &= N(q_i - q_{i-1}) \qquad \forall i \in \{2, \ldots, \hat{N} - 1\}, \\ \lambda_{\hat{N}} &= Nq_{\hat{N}} - 1. \end{aligned}$$

From the fact that the q_i are componentwise nondecreasing, and the above argument that $q_{\hat{N}} \geq 1/N$, we see that $\boldsymbol{\lambda} \geq \mathbf{0}$. In addition, we find that

$$\sum_{i=1}^{\hat{N}} \lambda_i = N(q_{\hat{N}-1} + q_{\hat{N}}) - 1 = N(2/N) - 1 = 1,$$

so $\boldsymbol{q} \in \operatorname{conv}(\mathcal{S}_N)$, and we are done. □

Figure 3-4 depicts a simple two-dimensional case with $N = 6$ data points. Shown are the basis generators for all of $\pi_{\boldsymbol{q}}(\mathcal{A})$ as well as those for the subclass of centrally symmetric $\pi_{\boldsymbol{q}}(\mathcal{A})$.

3.4 From one-sided moments to norm-bounded uncertainty

In this section, we examine a class of coherent risk measures which depend on higher-order moments. We will see that they induce uncertainty sets which are norm-bounded. As the representation theorem for coherent risk measures (Theorem 2.2.1) is essentially a duality result, it is not surprising that these norms are the dual norms of the moments in the original risk measure description.

We now present a result, shown in Delbaen [50] and Fischer [56], and provide an alternate proof using Lagrange duality. We use the notation $x^+ = \max(0, x)$.

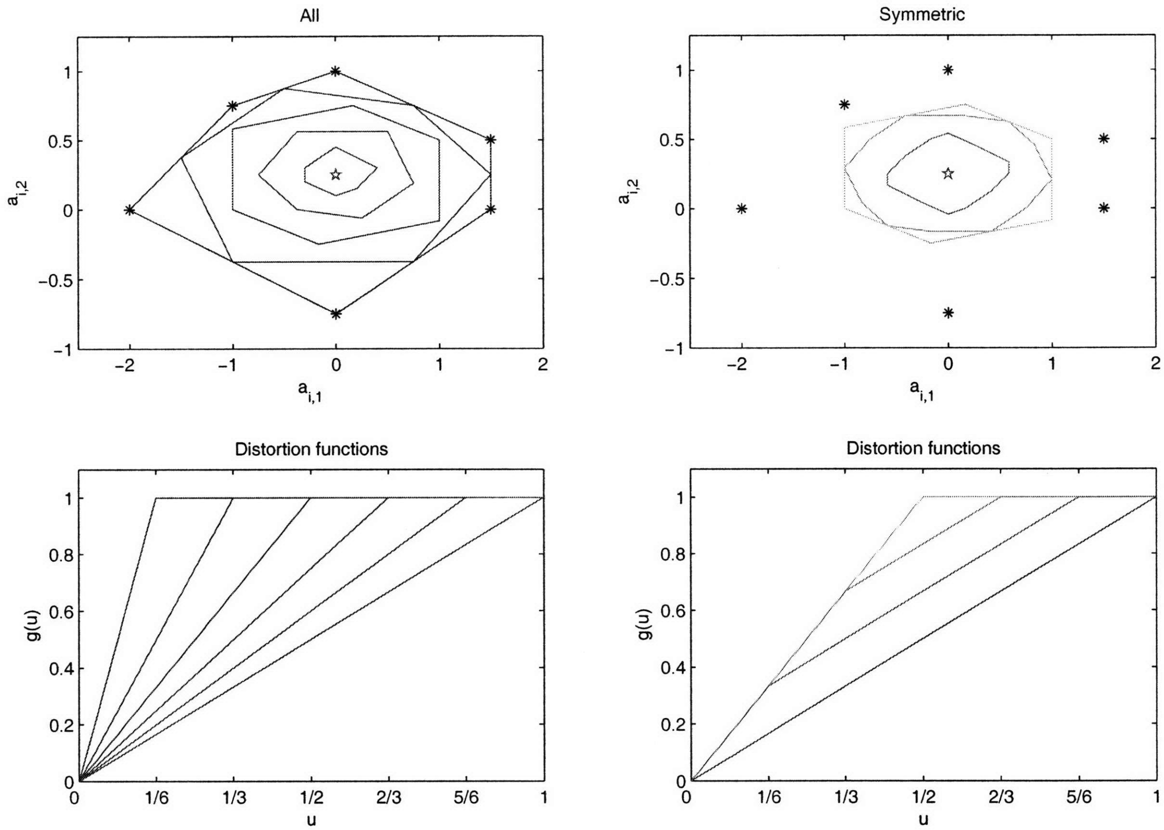

Figure 3-4: *For a case with $N = 6$ data points (denoted by $*$): the 6 basis generators for $\pi_{\boldsymbol{q}}(\mathcal{A})$ (upper left) and the corresponding distortion ($\mathrm{CTE}_{i/N}(\cdot)$) functions (lower left); the 4 basis generators for the centrally symmetric subclass of $\pi_{\boldsymbol{q}}(\mathcal{A})$ (upper right) and the corresponding distortion functions (lower right).*

Theorem 3.4.1. *If X is a random variable distributed as in Assumption 3.3.1, and for any $\alpha \in [0,1]$ and $p \geq 1$, the risk measure*

$$\mu_{p,\alpha}(X) \quad = \quad \mathbb{E}[X] + \alpha\sigma_{p,+}(X), \tag{3.24}$$

where

$$\sigma_{p,+}(X) \quad = \quad \left[\mathbb{E}\left[\left((X - \mathbb{E}[X])^+\right)^p\right]\right]^{1/p},$$

is a coherent risk measure. Moreover, it is representable by the family of measures

$$\mathcal{Q}_{q,\alpha} \quad = \quad \left\{\frac{1}{N}\left(\boldsymbol{e} + \alpha\left(\boldsymbol{g} - \hat{g}\boldsymbol{e}\right)\right) \mid \boldsymbol{g} \geq 0,\ \|\boldsymbol{g}\|_q \leq 1\right\}, \tag{3.25}$$

where $q = p/(p-1)$ *and* $\hat{g} = 1/N \sum_{i=1}^{N} g_i$.

Proof. The proof of coherence is in Fischer [56]. For the representability, it is easy to see that any $\boldsymbol{q} \in \mathcal{Q}$ sums to 1. For nonnegativity, using the fact that $\|\boldsymbol{g}\|_1 \leq N\|\boldsymbol{g}\|_q$ for all $q \geq 1$, and since $\boldsymbol{g} \geq 0$, $\alpha \in [0,1]$, we have

$$\boldsymbol{e} + \alpha(\boldsymbol{g} - \hat{g}\boldsymbol{e}) \;\geq\; \boldsymbol{e} - \frac{\alpha}{N}\|\boldsymbol{g}\|_1 \boldsymbol{e} \;\geq\; \boldsymbol{e}\,(1 - \|\boldsymbol{g}\|_q) \;\geq\; \boldsymbol{0},$$

since $\|\boldsymbol{g}\|_q \leq 1$. Now, denoting by x_i the values of the random variable X on the discrete sample space, and also denoting $\tilde{x}_i = x_i - \mathbb{E}[X]$, we consider the optimization problem

$$\begin{array}{ll} \text{maximize} & \tilde{\boldsymbol{x}}'\boldsymbol{g} \\ \text{subject to} & \|\boldsymbol{g}\|_q \leq 1, \\ & \boldsymbol{g} \geq \boldsymbol{0}. \end{array}$$

We form the Lagrangian

$$\mathcal{L}(\boldsymbol{g}, \lambda) \;=\; \tilde{\boldsymbol{x}}'\boldsymbol{g} + \lambda(1 - \|\boldsymbol{g}\|_q),$$

and note that $\mathcal{L}(\boldsymbol{g}, \lambda) \geq \tilde{\boldsymbol{x}}'\boldsymbol{g}$ for all $\boldsymbol{g} \geq \boldsymbol{0}$, $\lambda \geq 0$. Consider setting λ such that $\lambda^{1/q-1} = (\|\tilde{\boldsymbol{x}}^+\|_p)^{p-1}$ and then set

$$\boldsymbol{g}^* \;=\; \left(\frac{\tilde{\boldsymbol{x}}^+}{\lambda}\right)^{\frac{1}{q-1}}.$$

It is easy to see that $\boldsymbol{g}^* \geq \boldsymbol{0}$ and $\|\boldsymbol{g}^*\|_q = 1$, which leads us to the conclusion that $\boldsymbol{g}^*$ is optimal to the original optimization problem (by applying strong duality). Finally, we have

$$
\begin{aligned}
\mathbb{E}_{\boldsymbol{g}^*}[X] &= \mathbb{E}[X] + \frac{\alpha}{N}[\boldsymbol{x}'\boldsymbol{g}^* - \hat{x}(\boldsymbol{e}'\boldsymbol{g}^*)] \\
&= \mathbb{E}[X] + \frac{\alpha}{N}[(\boldsymbol{x} - \mathbb{E}[X]\boldsymbol{e})'\boldsymbol{g}^*] \\
&= \mathbb{E}[X] + \frac{\alpha}{N}\frac{\sum_{i=1}^{N}(\tilde{x}_i^+)^{\frac{1}{q-1}+1}}{(\|\tilde{\boldsymbol{x}}^+\|_p)^{p-1}} \\
&= \mathbb{E}[X] + \frac{\alpha}{N}\|\tilde{\boldsymbol{x}}^+\|_p \\
&= \mathbb{E}[X] + \alpha\sigma_{p,+}(X),
\end{aligned}
$$

which proves the result. □

This leads us directly to the following result.

Theorem 3.4.2. *The risk averse optimization problem (3.4) with risk measure* $\mu = \mu_{p,\alpha}$ *for some* $p \geq 1$, $\alpha \in [0,1]$ *defined in Equation (3.24), and uncertain vector* $\tilde{\boldsymbol{a}}$ *distributed according to Assumption 3.3.1, is equivalent to the robust optimization problem*

$$
\begin{array}{lll}
\text{minimize} & \boldsymbol{c}'\boldsymbol{x} & \\
\text{subject to} & \boldsymbol{a}'\boldsymbol{x} \leq b, & \forall \boldsymbol{a} \in \mathrm{P}_{q,\alpha}(\mathcal{A}),
\end{array} \tag{3.26}
$$

with uncertainty set

$$
\mathrm{P}_{q,\alpha}(\mathcal{A}) = \left\{ \sum_{i \in \mathcal{N}} p_i \boldsymbol{a}_i \;\middle|\; \boldsymbol{p} \in \mathcal{Q}_{q,\alpha} \right\}, \tag{3.27}
$$

where $\mathcal{Q}_{q,\alpha}$ *is as in (3.25) and* $q = p/(p-1)$. *Moreover, (3.26) is equivalent to the* *convex optimization problem*

$$
\begin{aligned}
\text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\
\text{subject to} \quad & \hat{\boldsymbol{a}}'\boldsymbol{x} + (\alpha/N)\hat{\lambda} \le b \\
& \lambda_i \ge (\boldsymbol{a}_i - \hat{\boldsymbol{a}})'\boldsymbol{x} \qquad \forall\, i \in \mathcal{N}, \\
& \|\boldsymbol{\lambda}\|_p \le \hat{\lambda},
\end{aligned} \tag{3.28}
$$

in variables $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{\lambda} \in \mathbb{R}^N$, $\hat{\lambda} \in \mathbb{R}$.

Proof. The fact that (3.4) with risk measure $\mu_{p,\alpha}$ is equivalent to (3.26) follows in similar fashion to the proof of Theorem 3.2.1 (simply apply Theorems 3.1.1 and 3.4.1, as well as Assumption 3.3.1 to $\mu_{p,\alpha}$). For the equivalence to (3.28), we have

$$
\begin{aligned}
\mu_{p,\alpha}(\boldsymbol{x}'\tilde{\boldsymbol{a}}) \le b \quad & \Leftrightarrow \quad \sup_{\boldsymbol{q} \in \mathcal{Q}_{q,\alpha}} \mathbb{E}_{\boldsymbol{q}}\left[\boldsymbol{x}'\tilde{\boldsymbol{a}}\right] \le b \\
& \Leftrightarrow \quad \boldsymbol{e}'\boldsymbol{A}\boldsymbol{x} + \alpha \max_{\boldsymbol{g} \ge \boldsymbol{0}, \|\boldsymbol{g}\|_q \le 1} \{(\boldsymbol{x}'\boldsymbol{A} - 1/N\,(\boldsymbol{x}'\boldsymbol{A}\boldsymbol{e})\,\boldsymbol{e}')\,\boldsymbol{g}\} \le Nb.
\end{aligned}
$$

Denoting $y_i = \boldsymbol{x}'\boldsymbol{a}_i$, $\hat{y} = 1/N\boldsymbol{e}'\boldsymbol{y}$, and $\tilde{y}_i = y_i - \hat{y}$, we need to then consider the optimization problem

$$
\begin{aligned}
\text{maximize} \quad & \tilde{\boldsymbol{y}}'\boldsymbol{g} \\
\text{subject to} \quad & \|\boldsymbol{g}\|_q \le 1, \\
& \boldsymbol{g} \ge \boldsymbol{0}.
\end{aligned}
$$

The dual of this problem is the convex conic problem

$$
\begin{aligned}
\text{minimize} \quad & \hat{\lambda} \\
\text{subject to} \quad & \lambda_i \ge \tilde{y}_i, \qquad \forall\, i \in \mathcal{N}, \\
& \|\boldsymbol{\lambda}\|_p \le \hat{\lambda},
\end{aligned}
$$

in variables $\boldsymbol{\lambda} \in \mathbb{R}^N$, $\hat{\lambda} \in \mathbb{R}$. Now noting that strong duality holds between the two problems (both are convex with strict interiors), we may insert the dual problem into the original problem and we get the desired result. □

Thus optimization over the coherent risk measure (3.24) based on higher-order tail moments is equivalent to a robust optimization problem with a norm-bounded uncertainty set. This robust problem can be solved efficiently as it is equivalent to a convex optimization problem with $N+1$ additional variables and $N+2$ total constraints ($N+1$ of which are linear). We conclude this section by noting that (3.27) naturally induces families of nested uncertainty sets.

Proposition 3.4.1.
(a) For $1 \leq q_1 \leq q_2$ and $\alpha \in [0,1]$, we have $\mathrm{P}_{q_1,\alpha}(\mathcal{A}) \subseteq \mathrm{P}_{q_2,\alpha}(\mathcal{A})$.
(b) For $q \geq 1$ and $\alpha_1,\ \alpha_2 \in [0,1]$ with $\alpha_1 \leq \alpha_2$ we have $\mathrm{P}_{q,\alpha_1}(\mathcal{A}) \subseteq \mathrm{P}_{q,\alpha_2}(\mathcal{A})$.

Proof.
(a) Follows form the definitions and the fact that $\|\boldsymbol{x}\|_{q_1} \leq 1 \Rightarrow \|\boldsymbol{x}\|_{q_2} \leq 1$ if $q_1 \leq q_2$.
(b) If $\boldsymbol{a} \in \mathrm{P}_{q,\alpha_1}(\mathcal{A})$ then $\boldsymbol{a} = 1/N(\boldsymbol{e} + \alpha_1(\boldsymbol{g}_1 - \hat{g}_1\boldsymbol{e}))$ for some $\boldsymbol{g}_1 \geq \mathbf{0}$, $\|\boldsymbol{g}_1\|_q \leq 1$. Set $\boldsymbol{g}_2 = (\alpha_1/\alpha_2)\boldsymbol{g}_1$ (so $\boldsymbol{g}_2 \geq \mathbf{0}$ and $\|\boldsymbol{g}_2\|_q \leq 1$) and we can also express $\boldsymbol{a}$ as $1/N(\boldsymbol{e} + \alpha_2(\boldsymbol{g}_2 - \hat{g}_2\boldsymbol{e}))$, implying that $\boldsymbol{a} \in \mathrm{P}_{q,\alpha_2}(\mathcal{A})$. □

3.5 From uncertainty sets to coherent risk measures

Up to this point we have focused on the situation where some coherent risk measure is the primitive element. We have seen that this primitive leads to an uncertainty set with some structure depending on the risk measure of choice. We now consider the converse construction; namely, we take the uncertainty set as the primitive. It is easy to see that the reverse construction holds; that is, a robust problem with a convex uncertainty set (contained within the convex hull of the data) corresponds to a risk averse problem with a coherent risk measure.

Proposition 3.5.1. The robust optimization problem (2.3) with convex uncertainty set $\mathcal{U} \subseteq \mathrm{conv}(\mathcal{A})$ for a single, uncertain constraint vector $\tilde{\boldsymbol{a}}$ distributed as in Assumption 3.3.1 is equivalent to a risk averse problem (3.4) with coherent risk measure generated (via Theorem (3.1.1)) by the following families of measures $\mathcal{Q}$:

(a) For arbitrary, convex $\mathcal{U} \subseteq \text{conv}(\mathcal{A})$,

$$\mathcal{Q} = \left\{ \boldsymbol{q} \in \Delta^N \mid \exists \boldsymbol{a} \in \mathcal{U} \; : \; \boldsymbol{a} = \boldsymbol{A}\boldsymbol{q} \right\}. \tag{3.29}$$

(b) For finitely generated, polyhedral $\mathcal{U} \subseteq \text{conv}(\mathcal{A})$, i.e., $\mathcal{U} = \text{conv}(\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_K\})$, we have

$$\mathcal{Q} = \left\{ \boldsymbol{q} \in \Delta^N \mid \exists i \in \{1, \ldots, K\} \; : \; \boldsymbol{u}_i = \boldsymbol{A}\boldsymbol{q} \right\}. \tag{3.30}$$

Proof. We only need to prove **(a)**, as **(b)** is merely a special case of it. Since $\mathcal{U}$ is convex and contained in $\text{conv}(\mathcal{A})$, for any $\boldsymbol{a} \in \mathcal{U}$, there exists a vector $\boldsymbol{q} \in \mathbb{R}^N$ such that $\boldsymbol{q} \geq \boldsymbol{0}$, $\boldsymbol{e}'\boldsymbol{q} = 1$, and $\boldsymbol{a} = \boldsymbol{A}\boldsymbol{q}$. This then implies that there exists a set $\mathcal{Q}$ such that

$$\max_{\boldsymbol{a} \in \mathcal{U}} \boldsymbol{a}'\boldsymbol{x} = \max_{\boldsymbol{q} \in \mathcal{Q}} \boldsymbol{q}'\boldsymbol{A}'\boldsymbol{x},$$

which is the supremum of the expected value of the random variable $\tilde{\boldsymbol{a}}'\boldsymbol{x}$ over a family of measures $\mathcal{Q}$. From Theorem 3.1.1 we know such a risk measure is coherent. □

3.5.1 Comonotone approximations

As we have noted, an arbitrary convex uncertainty set corresponds to a coherent risk measure. In general, however, there may not be an efficient representation for the corresponding measure. For instance, if the uncertainty set is a polytope with a facet description, then, from Proposition 3.5.1(b), the risk measure can be written as the supremum over a number of scenarios, where each scenario corresponds to an extreme point of the uncertainty set. In general, of course, this set may be very large and therefore difficult to describe.

It is natural, then, to consider efficiently representable approximations to these risk measures. In particular, we consider comonotone approximations to the risk measures derived from convex uncertainty sets. From a complexity standpoint, this is very desirable because, as we have noted, comonotone risk measures are uniquely defined by a single generating measure. In addition, optimizing over comonotone measures may be done efficiently (Theorem 3.3.1). It turns out that the class of centrally symmetric $\pi_{\boldsymbol{q}}(\mathcal{A})$ from Section 3.3.2 naturally induces a norm, and thus provides a convenient way of approximating more general measure structures, as we now illustrate.

The norm $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$

We now describe the norm induced by the sets $\pi_{\boldsymbol{q}}(\mathcal{A})$, where $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$.

Proposition 3.5.2. For a comonotone generator $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$ and any N distinct points in $\mathbb{R}^n$, $\mathcal{A} = \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N\}$, with $\hat{\boldsymbol{a}} = 1/N \sum_{i=1}^N \boldsymbol{a}_i$, the function

$$\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \;=\; \inf\left\{\alpha > 0 \;\middle|\; \frac{\boldsymbol{a} - \hat{\boldsymbol{a}}}{\alpha} \in \tilde{\pi}_{\boldsymbol{q}}(\mathcal{A})\right\}, \tag{3.31}$$

where $\tilde{\pi}_{\boldsymbol{q}}(\mathcal{A})$ is $\pi_{\boldsymbol{q}}(\mathcal{A})$ shifted by $-\hat{\boldsymbol{a}}$, is a norm.

Proof. $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$ is a form of a Minkowski function, and it is well-known that this function is convex whenever the underlying set in question (here $\tilde{\pi}_{\boldsymbol{q}}(\mathcal{A})$) is closed and convex, which is the case in this construction. Without loss of generality we assume $\hat{\boldsymbol{a}} = \boldsymbol{0}$ in the remainder of this proof.

$\|\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}} = 0$ implies that $\boldsymbol{a}/\epsilon \in \pi_{\boldsymbol{q}}(\mathcal{A})$ for all $\epsilon > 0$. As $\pi_{\boldsymbol{q}}(\mathcal{A})$ is a bounded set, this can only be the case if $\boldsymbol{a} = \boldsymbol{0}$.

If $\beta > 0$, it is easy to see that $\|\beta\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}} = \beta\|\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}}$ by a simple scaling argument. If $\beta < 0$, we have $\beta\boldsymbol{a} \in \pi_{\boldsymbol{q}}(\mathcal{A})$ if and only if $\pi_{\boldsymbol{q}}(\mathcal{A})$ is centrally symmetric about zero: this is the case, however, since $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$. Combining all this, we see that $\|\beta\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}} = |\beta|\|\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}}$.

Finally, noting that this function is convex, we have, for all $\boldsymbol{a}_1, \boldsymbol{a}_2 \in \mathbb{R}^n$,

$$\begin{aligned} \|\boldsymbol{a}_1 + \boldsymbol{a}_2\|_{\boldsymbol{q},\mathcal{A}} &= \|2(1/2\boldsymbol{a}_1 + 1/2\boldsymbol{a}_2)\|_{\boldsymbol{q},\mathcal{A}} \\ &= 2\|1/2\boldsymbol{a}_1 + 1/2\boldsymbol{a}_2\|_{\boldsymbol{q},\mathcal{A}} \\ &\leq 2\left[\|1/2\boldsymbol{a}_1\|_{\boldsymbol{q},\mathcal{A}} + \|1/2\boldsymbol{a}_2\|_{\boldsymbol{q},\mathcal{A}}\right] \\ &= \|\boldsymbol{a}_1\|_{\boldsymbol{q},\mathcal{A}} + \|\boldsymbol{a}_2\|_{\boldsymbol{q},\mathcal{A}}, \end{aligned}$$

which completes the proof that it is a norm. □

The norm $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$ has one particular property which is of interest.

Lemma 3.5.1. *Let $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$ be any centrally symmetric generator and $\lambda \in \mathbb{R}$. Then the vector $\tilde{\boldsymbol{q}} = \lambda\boldsymbol{q} + (1-\lambda)\boldsymbol{e}/N$ satisfies*

$$\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\tilde{\boldsymbol{q}},\mathcal{A}} = \frac{1}{|\lambda|}\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \tag{3.32}$$

for all $\boldsymbol{a} \in \mathbb{R}^n$.

Proof. The proof follows by noting that $\pi_{\tilde{\boldsymbol{q}}}(\mathcal{A})$ is a scaled version of $\pi_{\tilde{\boldsymbol{q}}}(\mathcal{A})$ around $\hat{\boldsymbol{a}}$ by a factor of λ. Indeed, it is easy to see that the extreme points of $\pi_{\tilde{\boldsymbol{q}}}(\mathcal{A})$ affine combinations of the extreme points of $\pi_{\tilde{\boldsymbol{q}}}(\mathcal{A})$ and $\boldsymbol{e}/N$ (which is permutation-invariant). (Note that because of this we may as well assume $\lambda \geq 0$, as $\lambda < 0$ reflects the set through $\hat{\boldsymbol{a}}$, and it is centrally symmetric through this point by construction.) The result then follows from the definition of $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$. □

Approximating norm-bounded uncertainty sets

In this section, we consider finding inner and outer approximations to uncertainty sets described by

$$\mathcal{U} = \{\boldsymbol{a} \mid \|\boldsymbol{a} - \hat{\boldsymbol{a}}\| \leq \Delta\}$$

for some norm $\|\cdot\|$ in $\mathbb{R}^n$. We may generate approximations from the class of centrally symmetric $\pi_{\boldsymbol{q}}(\mathcal{A})$ using the fact that the associated norm $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$ is equivalent to any other norm $\|\cdot\|$.

Proposition 3.5.3. Let $\|\cdot\|$ be any norm in $\mathbb{R}^n$ and $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$ be the norm induced by a centrally symmetric generator $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$, and denote $\hat{\boldsymbol{a}} = 1/N \sum_{i=1}^N \boldsymbol{a}_i$. Assuming there exists a basis $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n\}$ of $\mathbb{R}^n$ such that $\|\boldsymbol{e}_i\| = 1$ for all $i \in \{1, \ldots, n\}$, we have the following:

(a) There exists a constant $\underline{c} \geq 1/\sum_{i=1}^N q_i y_{(i)}$, where $y_{(i)}$ are the order statistics of $\|\boldsymbol{a}_i\|$, such that

$$\|\boldsymbol{x} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \geq \underline{c}\|\boldsymbol{x} - \hat{\boldsymbol{a}}\|. \tag{3.33}$$

(b) There exists a constant $0 \leq \bar{c} \leq \sum_{i=1}^n \|\boldsymbol{e}_i\|_{\boldsymbol{q},\mathcal{A}}$ such that

$$\|\boldsymbol{x} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \leq \bar{c}\|\boldsymbol{x} - \hat{\boldsymbol{a}}\|. \tag{3.34}$$

Proof. We assume without loss of generality that $\hat{\boldsymbol{a}} = \boldsymbol{0}$.

(a) Consider the problem $\max\{\|\boldsymbol{a}\| \mid \|\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}} \leq 1\}$. Denote the optimal value of this problem $1/\underline{c}$. We have

$$\begin{aligned} 1/\underline{c} &\leq \Big\|\sum_{i=1}^N q_{\sigma(i)}\boldsymbol{a}_i\Big\| \qquad \text{for some } \sigma \in \mathrm{S}_N \\ &\leq \sum_{i=1}^N q_{\sigma(i)}\|\boldsymbol{a}_i\| \\ &\leq \sum_{i=1}^N q_i y_{(i)}, \end{aligned}$$

which shows (3.33).

(b) Consider the problem $\max\{\|\boldsymbol{a}\|_{\boldsymbol{q},\mathcal{A}} \mid \|\boldsymbol{a}\| \leq 1\}$, and denote its optimal value by $\bar{c}$. Clearly, we must have $\bar{c} \geq 0$, for otherwise the optimal solution of this problem would be $\boldsymbol{a} = \boldsymbol{0}$, which cannot be the case (as it is a maximization of a nonnegative function over a nonempty compact set). We have

$$\begin{aligned}\bar{c} &\leq \|\sum_{i=1}^{n}\lambda_i \boldsymbol{e}_i\|_{\boldsymbol{q},\mathcal{A}} \qquad \text{for some } \boldsymbol{\lambda}\in\mathbb{R}^n : \sum_{i=1}^{n}|\lambda_i|\leq 1\\ &\leq \sum_{i=1}^{n}|\lambda_i|\|\boldsymbol{e}_i\|_{\boldsymbol{q},\mathcal{A}}\\ &\leq \sum_{i=1}^{n}\|\boldsymbol{e}_i\|_{\boldsymbol{q},\mathcal{A}},\end{aligned}$$

which shows (3.34) and we are done. □

Remark 3.5.1. For specific $\boldsymbol{q}\in\tilde{\Delta}^N_{\text{sym}}$ and data $\mathcal{A}$, computing the constants $\underline{c}$, $\bar{c}$ exactly seems to be a difficult combinatorial problem. The bounds derived in Proposition 3.5.3, however, are efficiently computable.

We now see that Lemma 3.5.1 and Proposition 3.5.3 provide a simple way of computing inner and outer approximations to uncertainty sets $\mathcal{U}$ described as norms centered around $\hat{\boldsymbol{a}}$.

Theorem 3.5.1. *Consider any centrally symmetric generating measure $\boldsymbol{q}\in\tilde{\Delta}^N_{\text{sym}}$. The uncertainty set described by*

$$\mathcal{U} \;=\; \{\boldsymbol{a} \mid \|\boldsymbol{a}-\hat{\boldsymbol{a}}\|\leq\Delta\},$$

where $\|\cdot\|$ is any norm and $\Delta\geq 0$, satisfies the inclusion $\pi_{\underline{\boldsymbol{q}}}(\mathcal{A})\subseteq\mathcal{U}\subseteq\pi_{\overline{\boldsymbol{q}}}(\mathcal{A})$, where

$$\underline{\boldsymbol{q}} \;=\; (\alpha\Delta)\boldsymbol{q}+(1-\alpha\Delta)\boldsymbol{e}/N, \tag{3.35}$$

$$\overline{\boldsymbol{q}} \;=\; (\beta\Delta)\boldsymbol{q}+(1-\beta\Delta)\boldsymbol{e}/N, \tag{3.36}$$

and α and β satisfy $0\leq\alpha\leq\underline{c}$, $\beta\geq\bar{c}$ from Proposition 3.5.3. Moreover, $\pi_{\underline{\boldsymbol{q}}}(\mathcal{A})$ and $\pi_{\overline{\boldsymbol{q}}}(\mathcal{A})$ correspond to comonotone risk measures on $\tilde{\boldsymbol{a}}$ under Assumption 3.3.1 if and *only if*

$$\alpha \leq \frac{1}{\Delta(1 - Nq_{min})}, \tag{3.37}$$

$$\beta \leq \frac{1}{\Delta(1 - Nq_{min})}, \tag{3.38}$$

respectively, where $q_{min} = \min_i q_i$.

Proof. We prove the inclusion $\pi_{\underline{q}}(\mathcal{A}) \subseteq \mathcal{U}$; the other inclusion follows in similar fashion. Since $0 \leq \alpha \leq \underline{c}$ and using Proposition 3.5.3, we have

$$\begin{aligned} \|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \leq \alpha\Delta \Rightarrow \|\boldsymbol{a} - \hat{\boldsymbol{a}}\| &\leq \frac{1}{\underline{c}}\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \\ &\leq \frac{1}{\alpha}\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \\ &\leq \Delta. \end{aligned}$$

Now, by Lemma 3.5.1, it is easy to see that

$$\begin{aligned} \{\boldsymbol{a} \mid \|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\boldsymbol{q},\mathcal{A}} \leq \alpha\Delta\} &= \{\boldsymbol{a} \mid \|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_{\underline{\boldsymbol{q}},\mathcal{A}} \leq 1\} \\ &= \pi_{\underline{\boldsymbol{q}}}(\mathcal{A}), \end{aligned}$$

and the result follows. (3.37) follows by noticing that $\underline{\boldsymbol{q}}$ must be a valid measure on $\mathbb{R}^N$ in order for it to be associated with a comonotone risk measure, i.e., all its components must be nonnegative. □

An example is illustrated in Figure 3-5.

Tight inner approximations to uncertainty polytopes

Here we consider the case when $\mathcal{U}$ is a polytope. Given a robust optimization problem over some uncertainty polytope, is there a corresponding comonotone risk measure? Not surprisingly, the answer is, in general, no. We first state an obvious case for which this converse does hold; we omit the proof, as it follows trivially from the preceding results.

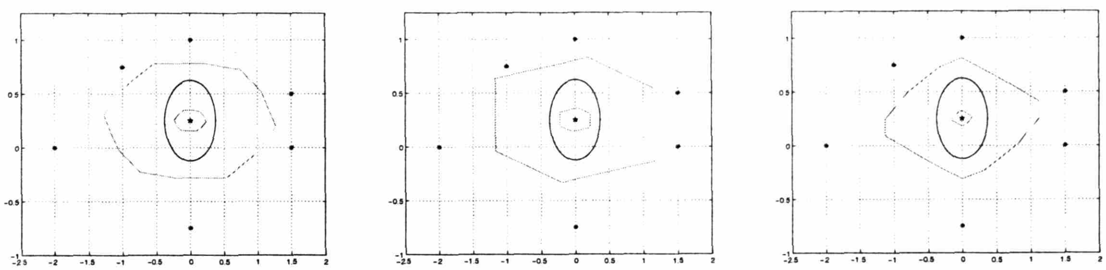

Figure 3-5: *A series of centrally symmetric for the data points from Figure 3-4 and the uncertainty set* $\mathcal{U} = \{\boldsymbol{a} \mid \|\boldsymbol{a} - \hat{\boldsymbol{a}}\|_2 \leq 3/8\}$ *using the result of Theorem 3.5.1. The inner approximations here all correspond to comonotone risk measures; the outer approximations here do not.*

Theorem 3.5.2. *A robust linear optimization problem (3.13) over an uncertainty set of the form* $\pi_{\boldsymbol{q}}(\mathcal{A})$ *for some* $\boldsymbol{q} \in \tilde{\Delta}^N$ *where the uncertain vector* $\tilde{\boldsymbol{a}}$ *is distributed according to Assumption 3.3.1 is equivalent to a risk averse problem (3.4) with a comonotone risk measure* μ_g.

In general, however, a robust optimization problem may certainly be defined over a polyhedral set which is not a $\boldsymbol{q}$-permutohull. In many cases the uncertainty set is an arbitrary polyhedron $\mathcal{U}$ described in constraint-wise form:

$$\mathcal{U} \quad = \quad \{\boldsymbol{a} \in \mathbb{R}^n \mid \boldsymbol{u}_i'\boldsymbol{a} \leq v_i, \ \forall\, i \in \{1, \ldots, m\}\}. \tag{3.39}$$

We assume that the sample mean $\hat{\boldsymbol{a}} \in \mathcal{U}$. In such cases, we can always find a comonotone risk measure which leads to an inner approximation (lower bound) on the robust problem with $\mathcal{U}$. We will find an inner approximation which is centrally symmetric about $\hat{\boldsymbol{a}}$ and use the norm $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$ derived in Proposition 3.5.2 to measure the quality of the approximation. We begin with a simple fact about $\|\cdot\|_{\boldsymbol{q},\mathcal{A}}$.

Lemma 3.5.1 immediately suggests a method for finding an inner approximation to an arbitrary uncertainty polytope $\mathcal{U}$: begin with a centrally symmetric generator $\boldsymbol{q} \in \tilde{\Delta}^N_{\text{sym}}$ and "mix" it with as little of the generator $\boldsymbol{e}/N$ such that the result is contained in $\mathcal{U}$. From Lemma 3.5.1 the resulting $\pi_{\boldsymbol{q}^*}(\mathcal{A})$ will be the largest in the

$\|\cdot\|_{q,\mathcal{A}}$ sense among all such mixtures contained in $\mathcal{U}$. We now show how to compute this algorithmically via linear optimization.

Theorem 3.5.3. *Given a centrally symmetric generator $\hat{\boldsymbol{q}} \in \tilde{\Delta}^N_{\text{sym}}$ and an arbitrary polytope $\mathcal{U} \subseteq \mathbb{R}^n$ described by (3.39) such that $\hat{\boldsymbol{a}} \in \mathcal{U}$, the centrally symmetric $\pi_{\boldsymbol{q}}(\mathcal{A})$ which is largest in the $\|\cdot\|_{\hat{q},\mathcal{A}}$ sense among all such $\pi_{\boldsymbol{q}}(\mathcal{A}) \subseteq \mathcal{U}$ is given by the solution to the linear optimization problem*

$$\begin{aligned} \text{maximize} \quad & \lambda \\ \text{subject to} \quad & \boldsymbol{q} = \lambda \hat{\boldsymbol{q}} + (1-\lambda)\boldsymbol{e}/N, \\ & \boldsymbol{e}'(\boldsymbol{s}_k + \boldsymbol{t}_k) \le v_k \ \forall\ k \in \{1,\dots,m\}, \\ & s_{k,i} + t_{k,j} \ge (\boldsymbol{u}_k'\boldsymbol{a}_j)q_i \ \forall\ i,j \in \mathcal{N}\times\mathcal{N},\ \forall\ k \in \{1,\dots,m\}, \end{aligned} \tag{3.40}$$

in variables $\boldsymbol{s}_k \in \mathbb{R}^N$, $k \in \{1,\dots,m\}$, $\boldsymbol{t}_k \in \mathbb{R}^N$, $k \in \{1,\dots,m\}$, $\boldsymbol{q} \in \mathbb{R}^N$, and $\lambda \in \mathbb{R}$. Moreover, the resulting approximating uncertainty set corresponds to a comonotone risk measure if and only if the optimal value λ^* of (3.40) satisfies

$$\lambda^* \le \frac{1}{1 - Nq_{min}}. \tag{3.41}$$

Proof. Consider a single inequality constraint $\boldsymbol{u}'\boldsymbol{a} \le v$. We have $\boldsymbol{u}'\boldsymbol{a} \le v$ for all $\boldsymbol{a} \in \pi_{\boldsymbol{q}}(\mathcal{A})$ if and only if the optimal value of the problem

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^N q_i \sum_{j=1}^N (\boldsymbol{u}'\boldsymbol{a}_j) y_{ij} \\ \text{subject to} \quad & \sum_{i=1}^N y_{ij} = 1 \qquad \forall\ j \in \mathcal{N}, \\ & \sum_{j=1}^N y_{ij} = 1 \qquad \forall\ i \in \mathcal{N}, \\ & y_{ij} \ge 0, \qquad \forall\ i,j \in \mathcal{N}\times\mathcal{N}, \end{aligned}$$

is no greater than v. We note that this is optimization over a bounded, nonempty polyhedron and thus by strong duality the optimal value of this problem equals the

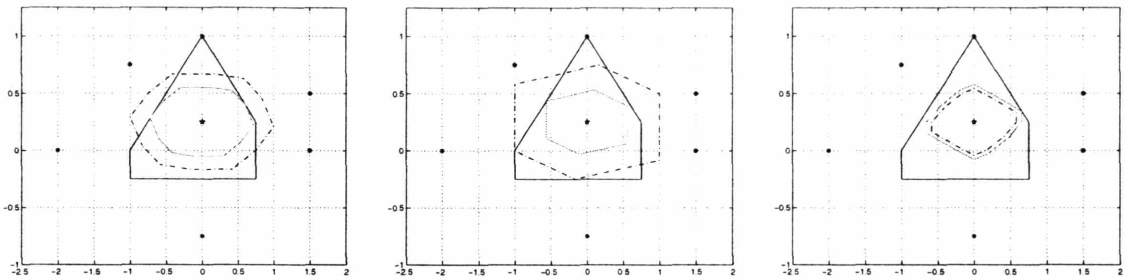

Figure 3-6: *Optimal inner approximation for a class of centrally symmetric generators for the example from Figure 3-4 and an arbitrary polyhedral uncertainty set. The dashed line indicates the non-scaled version of* $\pi_{\boldsymbol{q}}(\mathcal{A})$ *in each case, and in dark gray is the tightest inner approximation. In the first two cases, the approximations are "shrunken" and thus correspond to comonotone risk measures. In the the last case, the optimal approximation is actually larger than* $\pi_{\boldsymbol{q}}(\mathcal{A})$.

optimal value of its dual

$$\begin{array}{ll} \text{minimize} & \boldsymbol{e}'(\boldsymbol{s}+\boldsymbol{t}) \\ \text{subject to} & s_i + t_j \geq (\boldsymbol{u}'\boldsymbol{a}_j)q_i \qquad \forall\, i,j \in \mathcal{N}\times\mathcal{N}. \end{array}$$

It is then easy to see that $\pi_{\boldsymbol{q}}(\mathcal{A}) \subseteq \mathcal{U}$ if and only if there exist $\boldsymbol{s}_1,\ldots,\boldsymbol{s}_m,\boldsymbol{t}_1,\ldots,\boldsymbol{t}_m$ such that

$$\begin{array}{rcll} \boldsymbol{e}'(\boldsymbol{s}_k+\boldsymbol{t}_k) & \leq & v_k & \forall\, k\in\{1,\ldots,m\}, \\ s_{k,i}+t_{k,j} & \geq & (\boldsymbol{u}_k'\boldsymbol{a}_j)q_i & \forall\, i,j\in\mathcal{N}\times\mathcal{N},\ \forall\, k\in\{1,\ldots,m\}. \end{array}$$

Now, to find the largest such $\pi_{\boldsymbol{q}}(\mathcal{A}) \subseteq \mathcal{U}$ in the $\|\cdot\|_{\hat{\boldsymbol{q}},\mathcal{A}}$ sense, we can, by Lemma 3.5.1, set $\boldsymbol{q} = \lambda\hat{\boldsymbol{q}} + (1-\lambda)\boldsymbol{e}/N$ and maximize λ, which leads us to the desired linear program. The bound (3.41) follows as in the proof of the bound (3.37) in Theorem 3.5.1. □

Figure 3-6 shows an example of an inner approximations to an arbitrary polyhedral uncertainty set.

3.6 Computational example

In this section, we provide a computational example to illustrate the connection between robust optimization and coherent risk measures. The goal here is to compare the distribution, structure, and performance for robust solutions generated according to some specific coherent risk measures to robust solutions generated from some more *ad hoc* uncertainty sets. Specifically, we consider optimal policies for some risk averse problems using uncertainty sets from $\mathrm{CTE}_\alpha(\cdot)$ risk measures as well as optimal policies for some ellipsoidal uncertainty sets. Our major findings are the following:

(1) The $\mathrm{CTE}_\alpha(\cdot)$ policies are much easier to understand and interpret in terms of their structure and distribution than their more *ad hoc* ellipsoidal counterparts.

(2) Even using a small number of observations (e.g., $N = 100$), the $\mathrm{CTE}_\alpha(\cdot)$ policies do in fact perform the best in terms of actual $\mathrm{CTE}_\alpha(\cdot)$ (measured on a large number of *new* observations) out of all the policies we study for the specific value of α for which they are computed.

(3) The $\mathrm{CTE}_\alpha(\cdot)$ policies outperform the ellipsoidal policies (in terms of $\mathrm{CTE}_\alpha(\cdot)$) over a wide range of α levels (i.e., not just for the specific value of α for which they are optimized). In some cases, they outperform for all $\alpha \in [0, 1]$.

3.6.1 Problem description and data generating process

In particular, we consider the problem an uncertain objective vector $\tilde{\boldsymbol{c}}$. The structure of the problem we examine is

$$\begin{array}{ll} \text{minimize} & \max_{c \in \mathcal{U}} \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \boldsymbol{e}'\boldsymbol{x} = 1, \\ & \boldsymbol{x} \geq \boldsymbol{0}. \end{array} \tag{3.42}$$

One can think of this problem in terms of portfolio theory; here we are attempting to allocate optimal fractions of our resources into various risky assets to minimize some risk function subject to no short sales. We reiterate that we are minimizing here; thus lower values of the risk function are preferred (which goes in hand with our framework that a higher risk is associated with a greater loss).

For the specific problem, we consider a problem with $\boldsymbol{x} \in \mathbb{R}^{10}$ and assume we have $N = 100$ observations of the uncertain vector $\tilde{\boldsymbol{c}}$. We generate these observations randomly according to uniform distributions for each component. Specifically, we take, for all $i = 1, \ldots, N$, $j = 1, \ldots n$, $c_{i,j} \sim \mathrm{U}[j/2, 3j/2]$ and concatenate the results as $\mathcal{C} = \{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_N\}$. The idea behind this distribution process is to generate various "assets" such that lower mean (and thus preferable in a minimization framework) assets have higher variance.[2]

3.6.2 Robust formulations

We then solved Problem (3.42) for the following two families of uncertainty sets,

$$\begin{aligned} \mathcal{U}_\alpha &= \pi_{\boldsymbol{q}_\alpha}(\mathcal{C}) \\ \mathcal{U}_\rho &= \{\boldsymbol{c} \in \mathbb{R}^n \mid \|\hat{\boldsymbol{\Sigma}}^{-1}(\boldsymbol{c} - \hat{\boldsymbol{c}})\|_2 \leq \rho\}, \end{aligned}$$

with $\boldsymbol{q}_\alpha$ set to the corresponding value to achieve the associated $\mathrm{CTE}_\alpha(\cdot)$ risk measures, and where $\hat{\boldsymbol{\Sigma}}^{-1}$ is the sample covariance matrix of $\mathcal{C}$, and $\hat{\boldsymbol{c}}$ is the sample mean of $\mathcal{C}$.

We solved the problem with the values $\alpha = .01$, $\alpha = .5$, and $\alpha = 1$. Note that $\mathrm{CTE}_1(X) = \mathbb{E}[X]$. For the $\mathcal{U}_\rho$, we set $\rho = \sqrt{p/(1-p)}$ for probability guarantees $p = .99$, $p = .95$, and $p = 0.9$ (leading to $\rho = 9.9$, $\rho = 4.4$, and $\rho = 3$, respectively). These imply, under any distribution for $\tilde{\boldsymbol{c}}$ with mean $\hat{\boldsymbol{c}}$ and covariance $\hat{\boldsymbol{\Sigma}}$, the guarantee $\mathbb{P}\left\{\tilde{\boldsymbol{c}}'\boldsymbol{x}^* \geq \max_{\boldsymbol{c} \in \mathcal{U}_\rho} \boldsymbol{c}'\boldsymbol{x}^*\right\} \leq 1 - p$ (Bertsimas et al. [30]).

[2] The uniform choice for the distribution may seem curious. The reason we chose this as opposed to a normal distribution was to yield a discernable variety of distributions after taking linear combinations of $\tilde{\boldsymbol{c}}$ with our decision vector $\boldsymbol{x}$. If $\tilde{\boldsymbol{c}}$ were normal, the resulting random variables $\tilde{\boldsymbol{c}}'\boldsymbol{x}$ would also be normal and their resulting distributions would not be as easily distinguished.

3.6.3 Performance comparison

After computing the optimal solutions, we then generated 100,000 new samples of the $\boldsymbol{c}_i$ and plotted the distributions of $\tilde{\boldsymbol{c}}'\boldsymbol{x}^*$. The results are shown in Figure 3-8.

The distributions on the left correspond to $\mathcal{U} = \pi_{\boldsymbol{q}_\alpha}(\mathcal{C})$ uncertainty sets and they match intuition; as we increase α in minimizing $\text{CTE}_\alpha(\tilde{\boldsymbol{c}}'\boldsymbol{x})$, we improve the mean but fatten the tails of the resulting distribution. We gain more in terms of mean performance, and correspondingly less in terms of protection against for "bad" events. On the right are the distributions for the uncertainty sets $\mathcal{U}_\rho$. Clearly larger ρ leads to a more conservative performance, as expected.

In the left part of Figure 3-8 we see the performance of the various optimal policies in terms of $\text{CTE}_\alpha(\tilde{\boldsymbol{c}}'\boldsymbol{x}^*)$ for all $\alpha \in [.01, 1]$. Careful inspection at the values $\alpha = 1$, $\alpha = .5$, and $\alpha = .01$ shows that the optimal policies which are to minimize these quantities do in fact perform the best in this example. *Moreover, for all* $\alpha \in [.01, 1]$, *one or more of the optimal* CTE $(\cdot)$ *policies outperforms all of the optimal Euclidean norm robust policies.*

In the right part of Figure 3-8 we see the optimal policies in terms of allocation for the 10 variables for the various uncertainty sets. $\alpha = 1$ simply maximizes expected value; smaller α results in a more even distribution of resources. The optimal policies for the $\mathcal{U}_\rho$ do not appear to have an intuitive interpretation.

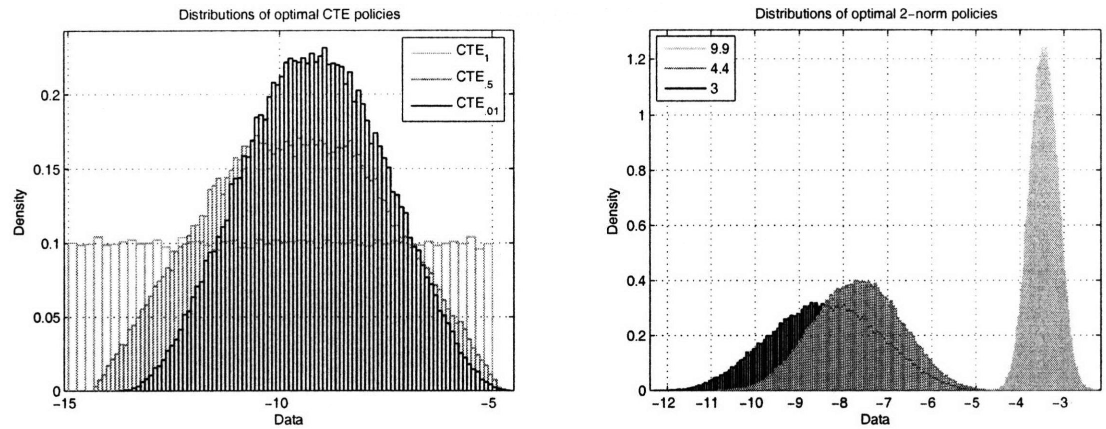

Figure 3-7: *Distributions of the optimal polices for the numerical example of Section 3.6.*

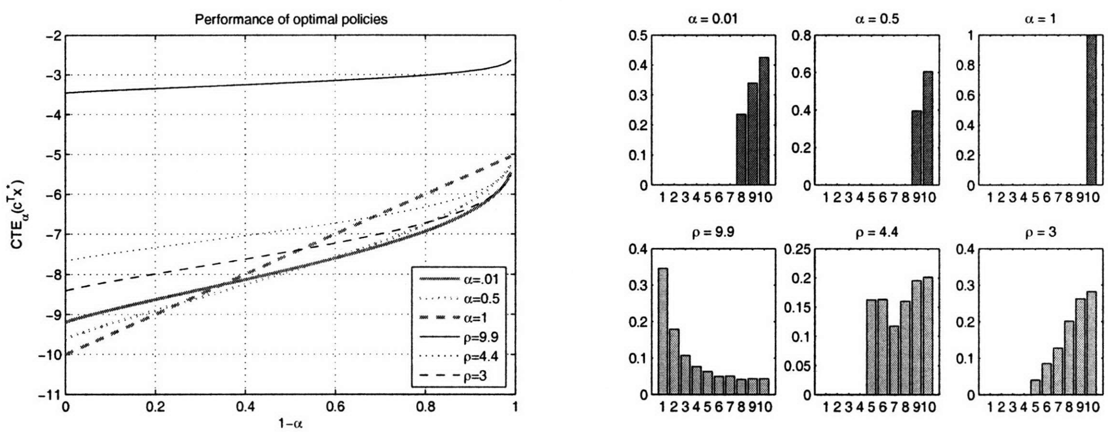

Figure 3-8: *Performance (left) and optimal allocations (right) for the numerical example of Section 3.6.*

Chapter 4

A flexible approach to robust optimization via convex risk measures

In this chapter, we attempt to extend the methodology of robust optimization to uncertain linear programs by employing a more flexible approach based on more general decision-maker risk preferences. This approach allows one to specify the tolerance to various degrees of infeasibility. To be specific, consider the constraint on a vector $\boldsymbol{x} \in \mathbb{R}^n$:

$$\tilde{\boldsymbol{a}}'\boldsymbol{x} \leq b, \tag{4.1}$$

where the constraint vector $\tilde{\boldsymbol{a}}$ is uncertain. As before, the model of uncertainty in our setting here is "data-driven," i.e., the only information on the uncertain vector $\tilde{\boldsymbol{a}}$ at our disposal is a finite set of sampled vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N$.

In the robust optimization approach, the uncertain constraint (4.1) is represented by its so-called robust counterpart (RC):

$$\boldsymbol{a}'\boldsymbol{x} \leq b \qquad \forall\, \boldsymbol{a} \in \mathcal{U}. \tag{4.2}$$

where $\mathcal{U}$ is a user-specified uncertainty set which is closed, bounded, and convex (in our uncertainty model here, we have $\mathcal{U} \subseteq \text{conv}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N)$). The RC can be equivalently written as

$$\boldsymbol{a}'\boldsymbol{x} \leq b + \beta(\boldsymbol{a}) \qquad \forall\, \boldsymbol{a} \in \mathbb{R}^n, \tag{4.3}$$

where $\beta(\boldsymbol{a}) = 0$ if $\boldsymbol{a} \in \mathcal{U}$ and $+\infty$ otherwise.

One can then view the RC as ensuring feasibility by employing a particular "extreme" penalty function - the indicator function of $\mathcal{U}$. In this chapter, we consider the possibility of using "milder" penalties which arise naturally by considering the decision-maker's risk preferences. Specifically, the risk preference criteria we use here are based on the theory of *convex risk measures* as developed by Föllmer and Schied [58]; in particular, we focus on the class of convex risk measures related to the Optimized Certainty Equivalent (OCE) measure due to Ben-Tal and Teboulle ([22], [23], [24]).

In Section 4.1, we revisit in more detail the class of convex risk measures introduced in Chapter 2 and show that for this class, a risk constraint on an uncertain linear function is equivalent to the inequality (4.3) with a specific function $\beta(\boldsymbol{a})$ arising from the Föllmer-Schied representation theory of convex risk measures [58]. We further obtain a dual formulation of (4.3) which is particularly amenable to tractable computation. In Section 4.2, we restrict our attention to the subclass of OCE-related risk measures introduced by Ben-Tal and Teboulle and show that the penalty functional is given in terms of a (generalized) *relative entropy* functional (see Csiszár [47]). In particular, we suggest four variants of the representation (4.3), each offering a different level of protection against infeasibility of the original constraint. In Section 4.3. we derive for each variant a probabilistic guarantee on the level of infeasibility.

A motivating example for our approach is a portfolio optimization problem of the form

$$\begin{aligned} \text{maximize} \quad & \hat{\boldsymbol{r}}'\boldsymbol{x} \\ \text{subject to} \quad & \tilde{\boldsymbol{r}}'\boldsymbol{x} \geq \gamma \\ & \boldsymbol{x} \in X, \end{aligned} \tag{4.4}$$

where $\tilde{\boldsymbol{r}}$ is an uncertain return vector, X is a deterministic feasible set, and $\hat{\boldsymbol{r}}$ is the expected return. Here, the constraint $\tilde{\boldsymbol{r}}'\boldsymbol{x} \geq \gamma$ is an uncertain constraint which, presumably, must be modelled in a way to reflect the risk preferences of the investor. This problem is in the spirit of the classical work of Markowitz [84], in which the risk measure of variance is used. In Section 4.4, we apply our proposed approach to a real-world asset allocation problem (4.4) using historical financial data. The results suggest that our milder-RC approach offers a tradeoff between expected return and downside risk protection that is, in many cases, more favorable than the pure-RC approach or the one based on conditional-value-at-risk (CVaR).

A direct connection between robust optimization and uncertainty sets associated with risk measures was been explored in Chapter 3 and has also been investigated by Natarajan et al., [87], but for the smaller class of *coherent* risk measures (Arztner et al., [3]). Our work here can be thought of as an extension of these papers, where here we connect these risk measures with more general and milder notions of robustness.

4.1 The general approach: convex risk measures

Recall the following class of risk measures:

Definition 4.1.1. Let $\tilde{\mathcal{M}}$ be the class of risk measures $\mu : \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\mu(X) \quad = \quad \sup_{\boldsymbol{q} \in \Delta^N} \left\{ \mathbb{E}_{\boldsymbol{q}}\left[X\right] - \alpha(\boldsymbol{q}) \right\} \tag{4.5}$$

where $\alpha : \Delta^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is a closed, convex function.

As discussed in Chapter 2, this is in fact the class of *convex risk measures* (Föllmer and Schied [58]) over $(\Omega, \mathcal{F}, \mathbb{P})$.

For convenience, we will work within a mildly restricted class of normalized convex risk measures. In particular, we will be interested in the following class:

Definition 4.1.2. Let $\mathcal{M}$ denote the set of $\mu \in \tilde{\mathcal{M}}$ which satisfy $\inf_{\boldsymbol{q} \in \Delta^N} \alpha(\boldsymbol{q}) = \alpha(\mathbb{P}) = 0$ and $\operatorname{dom} \alpha \subseteq \mathbb{R}_+^N$.

Note that the condition $\inf_{\boldsymbol{q} \in \Delta^N} \alpha(\boldsymbol{q}) = 0$ implies that $\mu(0) = 0$, which, in conjunction with the translation invariance property, means that $\mu(c) = c$ for every constant c. It also allows us to interpret μ as a minimum capital requirement to make a random variable X acceptable.

The fact that $\mathbb{P}$ attains the minimum, implies that $\mu(X) \geq \mathbb{E}_{\mathbb{P}}[X]$ for all $X \in \mathcal{X}$, i.e., expected value is the *least* conservative risk measure in this class.

The condition on the domain of α can be imposed without loss of generality since we are always penalizing over nonnegative measures.

4.1.1 Tractability of convex risk measures within linear optimization

We now illustrate how we will use these convex risk measures in the context of linear optimization under uncertainty. Consider the uncertain, linear constraint $\tilde{\boldsymbol{a}}'\boldsymbol{x} \leq b$, where $\boldsymbol{x} \in \mathbb{R}^n$ is a decision vector, $\tilde{\boldsymbol{a}}$ is an uncertain constraint vector, and $b \in \mathbb{R}$ is a known constant. For simplicity, we are only considering a single, uncertain constraint; clearly the discussion can be carried out in constraint-wise fashion for multiple, uncertain constraints.

In this chapter, as before, we assume a crude knowledge of the uncertainty associated with $\tilde{\boldsymbol{a}}$, namely:

Assumption 4.1.1. The uncertain vector $\tilde{\boldsymbol{a}}$ has support $\mathcal{A} = \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_N\}$ and a probability measure $\mathbb{P} \in \Delta^N$, i.e., $\mathbb{P}\{\tilde{\boldsymbol{a}} = \boldsymbol{a}_i\} = p_i$, where $p_i > 0$ and $\sum_{i=1}^{N} p_i = 1$.

Assumption 4.1.1 captures a prevailing situation in many practical problems where one has at his disposal merely N samples of the uncertain vector $\tilde{\boldsymbol{a}}$ (presumably, obtained from historical data). Thus, in this case, we have $\Omega \subseteq \mathbb{R}^n$ with $|\Omega| = N$, so the space of measures which are absolutely continuous with respect to $\mathbb{P}$ is just Δ^N.

Definition 4.1.3. Given an uncertain constraint vector $\tilde{\boldsymbol{a}}$, a known constant $b \in \mathbb{R}$, and a risk measure $\mu : \mathcal{X} \to \mathbb{R}$, we say that $\boldsymbol{x} \in \mathbb{R}^n$ satisfies the **risk-averse constraint** under μ if

$$\mu\left(\tilde{\boldsymbol{a}}'\boldsymbol{x}\right) \leq b. \tag{4.6}$$

Note that when μ satisfies translation invariance (as all convex risk measures must), (4.6) is equivalent to the constraint $\mu(\tilde{\boldsymbol{a}}'\boldsymbol{x} - b) \leq 0$.

It is not immediately obvious that (4.6) is tractable, even when μ is convex. In fact, the general principle for finding a tractable approach to a robust optimization problem is to utilize duality principles and replace the supremum with an infimum of a convex function over a potentially lifted space of variables. For instance, consider the case of a usual robust linear constraint of the form

$$\sup_{\boldsymbol{a} \in \mathcal{U}} \boldsymbol{a}'\boldsymbol{x} \leq b.$$

If the description of $\mathcal{U}$ is such that we can find a constraint of the form

$$\inf_{\boldsymbol{x} \in X,\ \boldsymbol{y} \in \mathcal{Y}} f(\boldsymbol{x}, \boldsymbol{y}) \leq b$$

which is equivalent to the original, robust constraint, and f is convex in $(\boldsymbol{x}, \boldsymbol{y})$, with $\mathcal{Y}$ convex, then we have a tractable representation (provided we have appropriate oracles for evaluating f, etc.).

We show next that we can accomplish similar results in our framework.

Theorem 4.1.1. *Let $\mu \in \mathcal{M}$ and let $\tilde{\boldsymbol{a}}$ satisfy Assumption 4.1.1. Then the following* *relations are equivalent:*

$$
\begin{aligned}
&\text{(A)} \quad \mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b \\
&\text{(B)} \quad \boldsymbol{a}'\boldsymbol{x} \le b + \beta(\boldsymbol{a}), \qquad \forall\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A}) \\
&\text{(C)} \quad \inf_{\nu\in\mathbb{R}} \{\nu + \alpha^*(\boldsymbol{A}\boldsymbol{x} - \nu\boldsymbol{e})\} \le b
\end{aligned}
$$

where $\boldsymbol{A}$ *is the* $N \times n$ *matrix:*

$$
\boldsymbol{A} \;=\; \left[\boldsymbol{a}_1 \cdots \boldsymbol{a}_N\right]'
$$

and

$$
\beta(\boldsymbol{a}) \;=\; \inf\left\{\alpha(\boldsymbol{q}) \mid \boldsymbol{q} \in \Delta^N, \boldsymbol{A}'\boldsymbol{q} = \boldsymbol{a}\right\}.
$$

Proof. We prove that (A)⇔(B) and (A)⇔(C). The key is the representation result of Föllmer and Schied. Let $\Delta^N(\boldsymbol{a}) = \{\boldsymbol{q} \in \Delta^N \mid \boldsymbol{A}'\boldsymbol{q} = \boldsymbol{a}\}$. To show equivalence of (A) to (B), we have

$$
\begin{aligned}
\mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b \;&\Leftrightarrow\; \boldsymbol{q}'\boldsymbol{A}\boldsymbol{x} - \alpha(\boldsymbol{q}) \le b, \;\; \forall\ \boldsymbol{q} \in \Delta^N \\
&\Leftrightarrow\; \boldsymbol{a}'\boldsymbol{x} \le b + \alpha(\boldsymbol{q}), \qquad \forall\ \boldsymbol{q} \in \Delta^N(\boldsymbol{a}),\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A}) \\
&\Leftrightarrow\; \boldsymbol{a}'\boldsymbol{x} \le b + \beta(\boldsymbol{a}), \qquad \forall\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A}).
\end{aligned}
$$

To show equivalence of (A) to (C), we have

$$
\begin{aligned}
\mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b \;&\Leftrightarrow\; \boldsymbol{q}'\boldsymbol{A}\boldsymbol{x} - \alpha(\boldsymbol{q}) \le b, \qquad \forall\ \boldsymbol{q} \in \Delta^N \\
&\Leftrightarrow\; \sup_{\boldsymbol{q}\in\Delta^N} \{\boldsymbol{x}'\boldsymbol{A}'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \le b \\
&\Leftrightarrow\; \inf_{\nu\in\mathbb{R}} \sup_{\boldsymbol{q}\ge\boldsymbol{0}} \{\boldsymbol{x}'\boldsymbol{A}'\boldsymbol{q} - \alpha(\boldsymbol{q}) + \nu(1 - \boldsymbol{e}'\boldsymbol{q})\} \le b \\
&\Leftrightarrow\; \inf_{\nu\in\mathbb{R}} \nu + \sup_{\boldsymbol{q}\ge\boldsymbol{0}} \{(\boldsymbol{A}\boldsymbol{x} - \nu\boldsymbol{e})'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \le b \\
&\Leftrightarrow\; \inf_{\nu\in\mathbb{R}} \nu + \sup_{\boldsymbol{q}\in\operatorname{dom}(\alpha)} \{(\boldsymbol{A}\boldsymbol{x} - \nu\boldsymbol{e})'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \le b \\
&\Leftrightarrow\; \inf_{\nu\in\mathbb{R}} \nu + \alpha^*(\boldsymbol{A}\boldsymbol{x} - \nu\boldsymbol{e}) \le b.
\end{aligned}
$$

where the second equivalence follows by standard duality arguments. This completes the proof. □

Formulation (B) in Theorem 4.1.1 gives an interpretation of the risk-averse problem in terms of "generalized robustness." Indeed, instead of strictly enforcing feasibility for all realizations of $\boldsymbol{a}$ within $\operatorname{conv}(\mathcal{A})$, (B) shows that we require a milder $\beta(\boldsymbol{a})$-feasibility. In other words, we not only control *where* we wish to be feasible, but we also control *how* feasible we are for a particular realization of $\tilde{\boldsymbol{a}}$. Thus, by appropriately choosing μ (and hence the penalty function α), one can balance the tradeoff between feasibility and conservatism in a much more flexible way than traditional robust models can.

It is also easy to see that the conditions requiring $\mu \in \mathcal{M}$ imply that the "penalty term" $\beta(\boldsymbol{a})$ is nonnegative. Thus, we have a family of formulations in Theorem 4.1.1 which are *less* conservative than the robust constraint

$$\boldsymbol{a}'\boldsymbol{x} \leq b, \qquad \forall\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A})\,.$$

Formulation (C), on the other hand, essentially shows that the risk-averse problem, when μ is convex, is tractable. Indeed, Theorem 4.1.1 implies the equivalence of the two problems:

$$\left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x}} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \mu(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b\end{array}\right\} \Leftrightarrow \left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x},\nu} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \nu + \alpha^*(\boldsymbol{A}\boldsymbol{x} - \nu) \leq b\end{array}\right\},$$

where the latter problem is convex from the well-known facts that the conjugate function of any function is convex (e.g., Rockafellar, [101]), and that convexity is preserved by affine transformations. Therefore, provided we have an oracle for efficiently evaluating α^* (and, possibly, its subgradients), then the risk-averse problem is tractable. In fact, for many choices of the penalty function α, the conjugate function can be computed analytically.

Example 4.1.1. (*Indicator functions*). Let the penalty function α in the definition of μ be an indicator function on a convex set, i.e.,

$$\alpha(\boldsymbol{q}) = \begin{cases} 0, & \text{if } \boldsymbol{q} \in \mathcal{Q}, \\ +\infty, & \text{otherwise,} \end{cases}$$

where $\mathcal{Q} \subseteq \Delta^N$ is a nonempty, closed, convex set. In this case $\alpha^*(\boldsymbol{y}) = \sup_{\boldsymbol{q}\in\mathcal{Q}}\{\boldsymbol{y}'\boldsymbol{q}\}$, so the risk-averse constraint in Theorem 4.1.1 is

$$\begin{aligned} \inf_{\nu\in\mathbb{R}}\{\nu + \sup_{\boldsymbol{q}\in\mathcal{Q}}\{(\boldsymbol{A}\boldsymbol{x} - \nu\boldsymbol{e})'\boldsymbol{q}\}\} \le b &\Leftrightarrow \sup_{\boldsymbol{q}\in\mathcal{Q}}\{(\boldsymbol{A}\boldsymbol{x})'\boldsymbol{q}\} \le b \\ &\Leftrightarrow \sup_{\boldsymbol{a}\in\mathcal{U}}\{\boldsymbol{a}'\boldsymbol{x}\} \le b, \end{aligned}$$

where $\mathcal{U} = \{\boldsymbol{A}'\boldsymbol{q} \mid \boldsymbol{q} \in \mathcal{Q}\}$.

This choice of α, then, leads to robust optimization in the traditional sense. At the same time the choice of α as an indicator function of a convex set yields the class of *coherent* risk measures [3], and we explored the connection between these risk measures and robust optimization in detail in Chapter 3.

4.2 Convex certainty equivalents and robustness

In this section, we consider various choices for the penalty function $\alpha(\cdot)$ which correspond to different notions of robustness, then demonstrate their connection to convex risk measures originating from *certainty equivalent measures.* In particular, we will connect the following four notions of robustness to convex risk measures:

1. Feasibility within an amount dependent on a distance measure, for all realizations of $\boldsymbol{a}$.

2. Feasibility for all realizations of $\boldsymbol{a}$ within a convex, compact set (this is the standard notion of robustness).

3. Feasibility for all realizations of $\boldsymbol{a}$ within a convex, compact set, and, in addition, feasibility within an amount dependent on a distance measure for all realizations

of $\boldsymbol{a}$ outside this set ("comprehensive robustness").

4. Feasibility within an amount dependent on a distance measure, for all realizations of $\boldsymbol{a}$ within a convex, compact set ("soft robustness").

In addition to showing the corresponding risk measures for these penalty functions, in this section, we will also explicitly construct the corresponding robust optimization problems, provide intuition behind the feasibility guarantees which are offered, and, finally, prove tractability of the approach with these types of penalty terms.

4.2.1 Optimized certainty equivalents

Our starting point will be penalty functions which depend on divergence measures from the underlying probability distribution $\mathbb{P}$. We need a bit of terminology before developing these ideas further and relating them to robust optimization problems.

We first recall a class of *certainty equivalents* introduced by Ben-Tal and Teboulle [22] and further developed in [23], [24].

Definition 4.2.1. Let $u : \mathbb{R} \to [-\infty, \infty)$ be a closed, concave, and nondecreasing utility function with nonempty domain. The **optimized certainty equivalent (OCE)** of a random variable $X \in \mathcal{X}$ under u is

$$S_u(X) \;=\; \sup_{\nu \in \mathbb{R}} \left\{ \nu + \mathbb{E}_{\mathbb{P}} \left[u(X - \nu) \right] \right\}. \tag{4.7}$$

The OCE can be interpreted as the value obtained by an optimal allocation between receiving a sure amount ν out of the future uncertain amount X now, and the remaining, uncertain amount $X - \nu$ later, where the utility function u effectively captures the "present value" of this uncertain quantity. It turns out that OCE measures have a dual description in terms of a convex risk measure with a penalty function described by a type of generalized relative entropy function called the *ϕ-divergence* (see Csiszár, [47]).

Definition 4.2.2. Let Φ be the class of all functions $\phi : \mathbb{R} \to \mathbb{R} \cup \{+\infty\}$ which are closed, convex, have a minimum value of 0 attained at 1, and satisfy $\operatorname{dom} \phi \subseteq \mathbb{R}_+$.

Specifically, we have the following, known result.

Theorem 4.2.1. *(Ben-Tal and Teboulle, [24]) With $u(t) = -\phi^*(-t)$, where $\phi \in \Phi$, we have*

$$S_u(X) = \inf_{\boldsymbol{q} \in \Delta^N} \left\{ \mathbb{E}_{\boldsymbol{q}}[X] + \sum_{i=1}^{N} p_i \phi(q_i/p_i) \right\}. \tag{4.8}$$

The term $\sum_{i=1}^{N} p_i \phi(q_i/p_i)$ is called the *ϕ-divergence of $\boldsymbol{q}$ with respect to $\mathbb{P}$*. It is a distance-like measure from $\boldsymbol{q}$ to $\mathbb{P}$; indeed, noting that $\phi \in \Phi$, by Jensen's inequality, we have

$$\sum_{i=1}^{N} p_i \phi(q_i/p_i) \geq \phi\left(\sum_{i=1}^{N} p_i \cdot (q_i/p_i)\right) = \phi(1) = 0,$$

with equality holding if $\boldsymbol{q} = \boldsymbol{p}$.

The framework defining the OCE in terms of a concave utility funcstion is derived in the context of random variables representing gains, whereas our concern is with random variables representing losses. To capture this difference, we will use the risk measure $\mu_\phi(X) := -S_u(-X)$, where $u(t) = -\phi^*(-t)$. Note that, in this case, we have

$$\begin{aligned} \mu_\phi(X) &= -S_u(-X) \\ &= -\sup_{\nu} \{\nu + \mathbb{E}_{\mathbb{P}}[u(-X - \nu)]\} \\ &= \inf_{\nu} \{-\nu - \mathbb{E}_{\mathbb{P}}[u(-X - \nu)]\} \\ &= \inf_{\nu} \{\nu - \mathbb{E}_{\mathbb{P}}[u(\nu - X)]\} \\ &= \inf_{\nu} \{\nu + \mathbb{E}_{\mathbb{P}}[\phi^*(X - \nu)]\}. \end{aligned}$$

In view of Theorem 4.2.1, it is also easy to see that

$$\begin{aligned}\mu_\phi(X) &= -S_u(-X) \\ &= -\inf_{\boldsymbol{q}\in\mathcal{Q}}\left\{\mathbb{E}_{\boldsymbol{q}}\left[-X\right]+\sum_{i=1}^N p_i\phi(q_i/p_i)\right\} \\ &= \sup_{\boldsymbol{q}\in\mathcal{Q}}\left\{\mathbb{E}_{\boldsymbol{q}}\left[X\right]-\sum_{i=1}^N p_i\phi(q_i/p_i)\right\},\end{aligned}$$

and therefore, the penalty function $\alpha_{\mathbb{P}}(\boldsymbol{q})$ in our setting is just the ϕ-divergence of $\boldsymbol{q}$ with respect to the reference measure $\mathbb{P}$, i.e., $\alpha_{\mathbb{P}}(\boldsymbol{q}) = \sum_{i=1}^N p_i\phi(q_i/p_i)$. Since $\alpha_{\mathbb{P}}(\boldsymbol{p}) = 0$, we have that $\mu_\phi \in \mathcal{M}$.

The interpretation of the risk measure μ_ϕ is analogous to that of the OCE, but in terms of losses: the decision-maker can pay off a sure amount ν of the uncertain debt X immediately and thereby leave the remaining uncertain amount $X-\nu$ remaining to be paid back at a later time. The function ϕ^* plays the role of a loss function, and $\mathbb{E}_{\mathbb{P}}\left[\phi^*(X-\nu)\right]$ is the current expected value of the future remaining debt. The risk measure $\mu_\phi(X)$ then reflects the optimal payoff allocation between these two time periods.

In the context of linear optimization under uncertainty, the risk measure μ_ϕ also has a very clear interpretation in terms of robustness, as we now make explicit.

Proposition 4.2.1. With $\phi \in \Phi$, and under the conditions of Theorem 4.1.1, the risk-averse constraint $\mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b$ is equivalent to each of the following:

$$\begin{aligned}&\textit{(Robustness)} \quad \boldsymbol{a}'\boldsymbol{x} \le b + \beta_\phi(\boldsymbol{a}), \qquad \forall\, \boldsymbol{a} \in \operatorname{conv}(\mathcal{A});\\ &\textit{(Risk aversion)} \quad \mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b,\end{aligned}$$

where

$$\beta_\phi(\boldsymbol{a}) = \inf_{\boldsymbol{q}\in\Delta^N}\left\{\sum_{i=1}^N p_i\phi(q_i/p_i) \;\middle|\; \boldsymbol{A}'\boldsymbol{q} = \boldsymbol{a}\right\}. \tag{4.9}$$

Proof. Both formulations follow directly from Theorem 4.1.1. □

Within the context of optimization, Proposition 4.2.1 implies the following equivalence:

$$\left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x}} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b\end{array}\right\} \Leftrightarrow \left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x},\nu} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \nu + \sum_{i=1}^{N} p_i \phi^*(\boldsymbol{a}_i'\boldsymbol{x} - \nu) \le b\end{array}\right\}.$$

Similar to the discussion after Theorem 4.1.1, it is clear that the problem on the right is convex in $(\boldsymbol{x}, \nu)$, and therefore modelling uncertainty using the OCE measure yields a tractable problem in this case.

Example 4.2.1. *(Exponential utility).* Consider $\phi(t) = t\log(t) - t + 1$. The corresponding dual function is $\phi^*(y) = e^y - 1$. The OCE in this case, then, is

$$\begin{aligned}\mu_\phi(X) &= \inf_{\nu \in \mathbb{R}} \left\{\nu + \mathbb{E}_{\mathbb{P}}\left[e^{X-\nu} - 1\right]\right\} \\ &= \log \mathbb{E}_{\mathbb{P}}\left[e^X\right].\end{aligned}$$

The corresponding, risk-averse constraint in the context of linear optimization under uncertainty, then, is

$$\begin{aligned}\mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b \;&\Leftrightarrow\; \log \mathbb{E}_{\mathbb{P}}\left[e^{\tilde{\boldsymbol{a}}'\boldsymbol{x}}\right] \le b \\ &\Leftrightarrow\; \log\left(e^{-b}\sum_{i=1}^{N} p_i e^{\boldsymbol{a}_i'\boldsymbol{x}}\right) \le 0.\end{aligned}$$

Such a constraint is convex in $\boldsymbol{x}$ and is, in fact, the convex form of a posynomial function associated with Geometric Programming (Boyd and Vandenberghe, [40]). Note that the penalty function $\alpha_{\mathbb{P}}(\boldsymbol{q})$ in this case is

$$\begin{aligned}\alpha_{\mathbb{P}}(\boldsymbol{q}) &= \sum_{i=1}^{N} p_i \phi(q_i/p_i) \\ &= \sum_{i=1}^{N} p_i((q_i/p_i)\log(q_i/p_i) - (q_i/p_i) + 1) \\ &= \sum_{i=1}^{N} q_i \log(q_i/p_i).\end{aligned}$$

known commonly in information theory as the *relative entropy* (or *Küllback-Liebler distance*) from $\boldsymbol{q}$ to $\mathbb{P}$ (see, for instance, Cover and Thomas, [46]). The robust problem, in this case, then, corresponds to one which enforces feasibility within $\beta_\phi(\boldsymbol{a})$ for all realizations of $\boldsymbol{a} \in \operatorname{conv}(\mathcal{A})$, where $\beta_\phi(\boldsymbol{a})$ measures the minimum relative entropy (from $\mathbb{P}$) among all measures which generate $\boldsymbol{a}$ in expectation.

Example 4.2.2. *(γ-divergence).* Consider the family of functions

$$\phi_\gamma(t) \;=\; \frac{1}{\gamma}t + \frac{1}{1-\gamma} - \frac{1}{\gamma(1-\gamma)}t^{1-\gamma},$$

parameterized by $\gamma \in [0, 1]$. The corresponding conjugate function is given by

$$\phi_\gamma^*(y) \;=\; \begin{cases} \frac{1}{1-\gamma}\left[(1-\gamma y)^{(\gamma-1)/\gamma} - 1\right], & \text{if } y < 1/\gamma, \\ +\infty, & \text{otherwise.} \end{cases}$$

The corresponding, risk-averse constraint is

$$\inf_{\nu \geq \max_i \boldsymbol{a}_i'\boldsymbol{x} + 1/\gamma} \left\{ \nu + \frac{1}{1-\gamma}\mathbb{E}_{\mathbb{P}}\left[(1-\gamma(\tilde{\boldsymbol{a}}'\boldsymbol{x} - \nu))^{(\gamma-1)/\gamma} - 1\right] \right\} \leq b.$$

For any $\gamma \in [0, 1]$, this is a convex constraint on $\boldsymbol{x}$ and $\nu \geq \max_i \boldsymbol{a}_i'\boldsymbol{x} + 1/\gamma$. For $\gamma = 1/2$, the associated divergence function is $\phi_{1/2}(t) = 2(\sqrt{t}-1)^2$, which is the divergence function generating the *Hellinger distance* $2\sum_{i=1}^{N}(\sqrt{q_i} - \sqrt{p_i})^2$ between measures $\boldsymbol{q}$ and $\boldsymbol{p}$. In the limiting case when $\gamma = 0$, we recover the relative entropy; indeed, this is easy to see from the fact that

$$\begin{aligned} \lim_{\gamma \to 0} \phi_\gamma^*(y) &= \lim_{\gamma \to 0} \frac{1}{1-\gamma}\left[(1-\gamma y)^{\frac{\gamma-1}{\gamma}} - 1\right] \\ &= \lim_{\gamma \to 0}(1-\gamma y)^{-1/\gamma} - 1 \\ &= e^y - 1. \end{aligned}$$

Thus, in this case, the robust problem corresponds to one which enforces feasibility within $\beta(\boldsymbol{a})$ for all realizations of $\boldsymbol{a} \in \operatorname{conv}(\mathcal{A})$, where $\beta(\boldsymbol{a})$ measures the minimum ϕ_γ-divergence (from $\mathbb{P}$) among all measures which generate $\boldsymbol{a}$ in expectation. The corresponding risk constraint is an OCE with the utility function given by ϕ_γ^*.

See Amari [2] for a more detailed discussion of these types of divergence measures.

Example 4.2.3. *(Piecewise linear utility).* Consider the divergence function

$$\phi(t) = \begin{cases} 0, & \text{if } t \in [\gamma_1, \gamma_2], \\ +\infty, & \text{otherwise,} \end{cases}$$

where $0 \le \gamma_1 < 1 < \gamma_2$. Clearly, $\phi \in \Phi$. A simple calculation shows that $\phi^*(y) = \max(\gamma_1 y, \gamma_2 y)$. With this choice of ϕ, then, we have

$$\begin{aligned} \mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b \;&\Leftrightarrow\; \exists \nu \in \mathbb{R} :\; \nu + \sum_{i=1}^N \max(\gamma_1(\boldsymbol{a}_i'\boldsymbol{x} - \nu), \gamma_2(\boldsymbol{a}_i'\boldsymbol{x} - \nu)) \le b \\ &\Leftrightarrow\; \exists (\nu, \boldsymbol{t}) \in \mathbb{R}^{N+1} : \begin{Bmatrix} \nu + \sum_{i=1}^N p_i t_i & \le & b \\ \gamma_1(\boldsymbol{a}_i'\boldsymbol{x} - \nu) & \le & t_i \\ \gamma_2(\boldsymbol{a}_i'\boldsymbol{x} - \nu) & \le & t_i \end{Bmatrix}, \; i = 1, \ldots, N, \end{aligned}$$

which is clearly convex (in fact, linear) in $(\boldsymbol{x}, \nu, \boldsymbol{t})$. Ben-Tal and Teboulle [24] show that the OCE is coherent if and only if ϕ has this form. Note also that when $\gamma_1 = 0$ and $\gamma_2 = 1/\alpha$ for some $\alpha \in (0, 1)$, we obtain the risk measure

$$\mu_\phi(X) = \inf_{\nu \in \mathbb{R}} \left\{ \nu + \frac{1}{\alpha} \mathbb{E}_{\mathbb{P}} \left[(X - \nu)^+ \right] \right\},$$

which is a representation for CVaR (e.g., Rockafellar and Uryasev, [102]).

Variants of penalty functions α related to the ϕ-divergence, and their connection to other robust models, will now be our focus.

4.2.2 Standard robustness

We have already discussed the connection of indicator functions of convex sets to coherent risk measures and, in turn, to the usual notion of robustness. Here we make this connection concrete in terms of a convex set related closely to ϕ-divergence.

Proposition 4.2.2. Consider the following penalty function

$$\alpha(\boldsymbol{q}) = \begin{cases} 0, & \text{if } \sum\limits_{i=1}^{N} p_i \phi(q_i/p_i) \le \rho, \\ +\infty, & \text{otherwise,} \end{cases}$$

where $\rho > 0$, $\phi \in \Phi$. Denote the corresponding risk measure by $R_{\phi,\rho}(X)$. Under the conditions of Theorem 4.1.1, the risk-averse constraint $R_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \le b$ is equivalent to each of the following:

$$\begin{array}{rl} \textit{(Robustness)} & \boldsymbol{a}'\boldsymbol{x} \le b, \qquad \forall\, \boldsymbol{a} \in \mathcal{U}_\phi(\rho); \\ \textit{(Risk aversion)} & \inf\limits_{\lambda>0} \rho\lambda + \lambda\mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}/\lambda) \le b, \end{array}$$

where

$$\mathcal{U}_\phi(\rho) = \left\{ \boldsymbol{A}'\boldsymbol{q} \,\middle|\, \boldsymbol{q} \in \Delta^N, \; \sum_{i=1}^{N} p_i \phi(q_i/p_i) \le \rho \right\}. \tag{4.10}$$

Proof. Since $\phi \in \Phi$, we can apply Theorem 4.1.1 once we have the conjugate function for α. We have

$$\begin{aligned} \alpha^*(\boldsymbol{y}) &= \sup_{\boldsymbol{q}} \{\boldsymbol{y}'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \\ &= \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} \,\middle|\, \sum_{i=1}^{N} p_i \phi(q_i/p_i) \le \rho \right\} \\ &= \inf_{\lambda \ge 0} \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} + \lambda \left(\rho - \sum_{i=1}^{N} p_i \phi(q_i/p_i) \right) \right\} \\ &= \inf_{\lambda>0} \rho\lambda + \lambda \sup_{\boldsymbol{q}} \left\{ (\boldsymbol{y}/\lambda)'\boldsymbol{q} - \sum_{i=1}^{N} p_i \phi(q_i/p_i) \right\} \\ &= \inf_{\lambda>0} \rho\lambda + \lambda \sum_{i=1}^{N} p_i \sup_{q_i} \left\{ \frac{y_i}{\lambda} \cdot \frac{q_i}{p_i} - \phi(q_i/p_i) \right\} \\ &= \inf_{\lambda>0} \rho\lambda + \lambda \sum_{i=1}^{N} p_i \sup_{q_i \in \operatorname{dom} \phi} \left\{ \frac{y_i}{\lambda} \cdot q_i - \phi(q_i) \right\} \\ &= \inf_{\lambda>0} \rho\lambda + \lambda \mathbb{E}_{\mathbb{P}} \left[\phi^*(\boldsymbol{y}/\lambda) \right], \end{aligned}$$

where the third equality follows by standard, convex duality arguments (the Slater condition holds, since $\phi \in \Phi$ implies that $\sum_{i=1}^{N} p_i \phi(p_i/p_i) = \phi(1) = 0 < \rho$). Now applying Theorem 4.1.1 to the random variable $\tilde{\boldsymbol{a}}'\boldsymbol{x}$, we have

$$\begin{aligned} R_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) &= \inf_{\nu \in \mathbb{R}} \inf_{\lambda > 0} \{\nu + \rho\lambda + \lambda \mathbb{E}_{\mathbb{P}}\left[\phi^*((\tilde{\boldsymbol{a}}'\boldsymbol{x} - \nu)/\lambda)\right]\} \\ &= \inf_{\lambda > 0} \lambda \inf_{\nu \in \mathbb{R}} \{\nu + \rho + \mathbb{E}_{\mathbb{P}}\left[\phi^*(\tilde{\boldsymbol{a}}'\boldsymbol{x}/\lambda - \nu)\right]\} \\ &= \inf_{\lambda > 0} \rho\lambda + \lambda \mu_{\phi}(\tilde{\boldsymbol{a}}'\boldsymbol{x}/\lambda), \end{aligned}$$

and we are done. □

Remark 4.2.1. The function $\lambda \mu_{\phi}(\tilde{\boldsymbol{a}}'\boldsymbol{x}/\lambda)$ which appears in Proposition 4.2.2 will be a recurring theme in a number of the models we consider. It is easy to show that this function, in addition to being nonincreasing in λ for any fixed $\boldsymbol{x}$, is also convex in $(\lambda, \boldsymbol{x})$, and therefore such problems may be efficiently solved with convex optimization techniques. One possible approach to such problems, provided one can efficiently solve problems with constraints of the form $\mu_{\phi}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b$, is to then bisect over $\lambda \in (0, U]$, where U is a predetermined upper bound on an optimal value of λ, and solve a polynomial number of nominal problems.

For a concrete example of Proposition 4.2.2 within the context of optimization, consider the case when $\phi(t) = t\log(t) - t + 1$. We then have the equivalence

$$\left\{\begin{array}{ll} \text{minimize}_{\boldsymbol{x}} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & R_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b \end{array}\right\} \Leftrightarrow \left\{\begin{array}{ll} \text{minimize}_{\boldsymbol{x},\lambda>0} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \rho\lambda + \lambda \log\left(\sum_{i=1}^{N} p_i e^{\boldsymbol{a}_i'\boldsymbol{x}/\lambda}\right) \leq b \end{array}\right\},$$

which is a convex optimization problem in $(\boldsymbol{x}, \lambda)$ over $\mathbb{R}^n \times \mathbb{R}_+$.

4.2.3 Comprehensive robustness

With $\alpha(\boldsymbol{q}) = \sum_{i=1}^{N} p_i \phi(q_i/p_i)$ for $\phi \in \Phi$, the interpretation in terms of generalized robustness is that $\boldsymbol{x}$ must satisfy $\boldsymbol{a}'\boldsymbol{x} \leq b + \beta(\boldsymbol{a})$ for all $\boldsymbol{a} \in \text{conv}(\mathcal{A})$, where $\beta(\boldsymbol{a})$ is the minimum ϕ-divergence with respect to $\mathbb{P}$ over all measures $\boldsymbol{q}$ such that $\mathbb{E}_{\boldsymbol{q}}[\tilde{\boldsymbol{a}}] = \boldsymbol{a}$.

Notice, however, that feasibility of $\boldsymbol{x}$ is only guaranteed for realizations $\boldsymbol{a}$ such that $\beta(\boldsymbol{a}) = 0$, and this set can be as small as a singleton.[1] On the flip side, a penalty function which is the indicator function of a convex set $\mathcal{Q} \subseteq \Delta^N$ (and thus the risk measure is coherent) offers *no* guarantees for realizations of $\tilde{\boldsymbol{a}}$ outside of the set $\mathcal{U} = \{\boldsymbol{A}'\boldsymbol{q} \mid \boldsymbol{q} = \Delta^N\}$. This motivates the following, related type of convex risk measure, which also has a certainty equivalent interpretation.

Proposition 4.2.3. Consider the following penalty function

$$\alpha(\boldsymbol{q}) = \max\left(0, \sum_{i=1}^{N} p_i \phi(q_i/p_i) - \rho\right),$$

where $\rho \geq 0$ and $\phi \in \Phi$, and denote the corresponding risk measure by $C_{\phi,\rho}(X)$. Under the conditions of Theorem 4.1.1, the risk-averse constraint $C_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b$ is equivalent to each of the following:

$$\begin{aligned} &\textit{(Robustness)} && \boldsymbol{a}'\boldsymbol{x} \leq b, \qquad \forall\ \boldsymbol{a} \in \mathcal{U}_\phi(\rho) \\ & && \boldsymbol{a}'\boldsymbol{x} \leq b + (\beta_\phi(\boldsymbol{a}) - \rho) \qquad \forall\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A}) \setminus \mathcal{U}_\phi(\rho); \\ &\textit{(Risk aversion)} && \inf_{\lambda \in (0,1]} \rho\lambda + \lambda\mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}/\lambda) \leq b, \end{aligned}$$

where $\beta_\phi(\boldsymbol{a})$ is as in Equation (4.9).

Proof. We need to evaluate the conjugate function of α, then use Theorem 4.1.1. We

[1]This, in fact, is the case when ϕ is strictly convex. In this case, the only guarantee is on the expected value of $\tilde{\boldsymbol{a}}'\boldsymbol{x}$, i.e., $\hat{\boldsymbol{a}}'\boldsymbol{x} \leq b$, where $\hat{\boldsymbol{a}} = \mathbb{E}_{\mathbb{P}}[\boldsymbol{a}]$.

have

$$
\begin{aligned}
\alpha^*(\boldsymbol{y}) &= \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} - \max\left(0, \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) - \rho\right) \right\} \\
&= \sup_{\boldsymbol{q}} \left\{ \min\left(\boldsymbol{y}'\boldsymbol{q}, \boldsymbol{y}'\boldsymbol{q} - \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) + \rho\right) \right\} \\
&= \sup_{\boldsymbol{q},\ t \leq 0} \left\{ \boldsymbol{y}'\boldsymbol{q} + t \;\middle|\; t \leq \rho - \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) \right\} \\
&= \inf_{(\lambda_1,\lambda_2)\in\mathbb{R}_+^2} \sup_{(\boldsymbol{q},t)} \left\{ \boldsymbol{y}'\boldsymbol{q} + t - \lambda_1 t - \lambda_2 \left(t + \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) - \rho \right) \right\} \\
&= \inf_{(\lambda_1,\lambda_2)\in\mathbb{R}_+^2} \sup_{(\boldsymbol{q},t)} \left\{ \boldsymbol{y}'\boldsymbol{q} + (1 - \lambda_1 - \lambda_2) t - \lambda_2 \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) + \rho\lambda_2 \right\}.
\end{aligned}
$$

Clearly, the supremum is finite if and only if $\lambda_1 + \lambda_2 = 1$, and therefore, it suffices to consider $\lambda_2 \in [0, 1]$. Continuing, we have

$$
\begin{aligned}
\alpha^*(\boldsymbol{y}) &= \inf_{\lambda\in[0,1]} \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} - \lambda \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) + \rho\lambda \right\} \\
&= \inf_{\lambda\in(0,1]} \rho\lambda + \lambda \left(\sup_{\boldsymbol{q}} \left\{ (1/\lambda)\boldsymbol{y}'\boldsymbol{q} - \sum_{i=1}^{N} p_i \phi\left(q_i/p_i\right) \right\} \right) \\
&= \inf_{\lambda\in(0,1]} \rho\lambda + \lambda \sum_{i=1}^{N} p_i \sup_{q_i} \left\{ \frac{y_i}{\lambda} \cdot \frac{q_i}{p_i} - \phi\left(\frac{q_i}{p_i}\right) \right\} \\
&= \inf_{\lambda\in(0,1]} \rho\lambda + \lambda \sum_{i=1}^{N} p_i \phi^*(y_i/\lambda) \\
&= \inf_{\lambda\in(0,1]} \rho\lambda + \lambda \mathbb{E}_{\mathbb{P}} \left[\phi^*(\boldsymbol{y}/\lambda)\right].
\end{aligned}
$$

Now, since $\phi \in \Phi$, $\operatorname{dom} \phi \subseteq \mathbb{R}_+$, and, therefore, we can apply Theorem 4.1.1 in similar fashion to the proof of Proposition 4.2.2 to obtain the desired result. □

Within the context of optimization, Proposition 4.2.3 implies the equivalence:

$$\left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x}} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & C_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b\end{array}\right\} \Leftrightarrow$$

$$\left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x},\nu,\lambda\in(0,1]} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \rho\lambda + \lambda\left(\nu + \sum_{i=1}^{N} p_i \phi^*(\boldsymbol{a}_i'\boldsymbol{x}/\lambda - \nu)\right) \leq b\end{array}\right\},$$

which is not convex in $(\boldsymbol{x}, \nu, \lambda)$ because of the term $\lambda\nu$, but by the change of variable $\eta = \lambda\nu$, it converts to an equivalent convex problem in variables $(\boldsymbol{x}, \eta, \lambda)$:

$$\left\{\begin{array}{ll}\text{minimize}_{\boldsymbol{x},\eta,\lambda\in(0,1]} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \rho\lambda + \eta + \lambda\sum_{i=1}^{N} p_i \phi^*\left(\frac{\boldsymbol{a}_i'\boldsymbol{x}-\eta}{\lambda}\right)\end{array}\right\}.$$

Remark 4.2.2. *(Certainty equivalent interpretation of $C_{\phi,\rho}$)* In light of Proposition 4.2.3, the risk measure $C_{\phi,\rho}$ has a very natural interpretation. Indeed, it is somewhat similar to the usual OCE in Proposition 4.2.1, with two differences. First, as before, we can pay a certain amount ν as an immediate credit towards the uncertain debt $\tilde{\boldsymbol{a}}'\boldsymbol{x}$, the remainder of which, $\tilde{\boldsymbol{a}}'\boldsymbol{x} - \nu$ is revealed at a later time period. In addition to this allocation choice, we also have control over $\lambda \in (0, 1]$, which can be interpreted as a *discount rate* describing the value over time. The uncertain amount $\tilde{\boldsymbol{a}}'\boldsymbol{x} - \nu$ accrues "interest" to a future-value loss of $(1/\lambda)(\tilde{\boldsymbol{a}}'\boldsymbol{x} - \nu)$, ϕ^* reflects its utility, and λ then scales the expected utility back to present value units. Finally, we pay a certain penalty of ρ at the later time period, and the present cost of this is thus $\lambda\rho$. This sure penalty represents the fact that we are protecting completely against all realizations of $\tilde{\boldsymbol{a}}$ within the set $\mathcal{U}_\phi(\rho)$.

Remark 4.2.3. *(Robust interpretation of $C_{\phi,\rho}$)* The robustness interpretation of Proposition 4.2.3 is straightforward: we require $\boldsymbol{a}'\boldsymbol{x} \leq b$ for all $\boldsymbol{a} \in \mathcal{U}_\phi(\rho)$, and enforce $\beta(\boldsymbol{a})$ feasibility for all $\boldsymbol{a} \in \mathcal{A} \setminus \mathcal{U}_\phi(\rho)$, where $\beta(\boldsymbol{a})$ represents the minimum ϕ-divergence among all measures $\boldsymbol{q} \in \Delta^N$ such that $\mathbb{E}_{\boldsymbol{q}}[\tilde{\boldsymbol{a}}] = \boldsymbol{a}$. This is similar in spirit to the work of Ben-Tal et al. [7], who explore models for "comprehensive robust optimization."

These models enforce feasibility over a pre-specified set $\mathcal{U}$ which represents the "normal range" of values that can be realized by the uncertain vector $\tilde{\boldsymbol{a}}$ (here, this is the set $\mathcal{U}_\phi(\rho)$). For values of $\tilde{\boldsymbol{a}}$ outside the normal range, the infeasibility is bounded by a term proportional to the distance of $\tilde{\boldsymbol{a}}$ to $\mathcal{U}$ (here, this distance is captured instead by $\beta(\boldsymbol{a})$).

4.2.4 Soft robustness

The penalty function in Proposition 4.2.3, then, is more conservative than a penalty function of the form

$$\alpha(\boldsymbol{q}) \;=\; \begin{cases} 0, & \text{if } \sum_{i=1}^{N} p_i\phi(q_i/p_i) \leq \rho, \\ +\infty, & \text{otherwise,} \end{cases}$$

as it offers protection for corresponding realizations of $\tilde{\boldsymbol{a}}$ *outside* the set $\mathcal{U}_\phi(\rho)$. We may also wish to consider a model for robustness which is less conservative than the standard approach. This motivates the following.

Proposition 4.2.4. Consider the following penalty function

$$\alpha(\boldsymbol{q}) = \begin{cases} \sum_{i=1}^{N} p_i\phi(q_i/p_i), & \text{if } \sum_{i=1}^{N} p_i\phi(q_i/p_i) \leq \rho, \\ +\infty, & \text{otherwise,} \end{cases}$$

where $\rho > 0$, $\phi \in \Phi$. Denote the corresponding risk measure by $S_{\phi,\rho}(X)$. Under the conditions of Theorem 4.1.1, the risk-averse constraint $S_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b$ is equivalent to each of the following:

$$\begin{aligned} \textit{(Robustness)} \quad & \boldsymbol{a}'\boldsymbol{x} \leq b + \beta_\phi(\boldsymbol{a}), \qquad \forall\ \boldsymbol{a} \in \mathcal{U}_\phi(\rho); \\ \textit{(Risk aversion)} \quad & \inf_{\lambda \geq 0} \{\rho\lambda + (\lambda+1)\mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}/(\lambda+1))\} \leq b. \end{aligned}$$

Proof. As in Proposition 4.2.3, the proof is largely an exercise in Lagrangian duality to find the conjugate function α^*, then a direct application of Theorem 4.1.1 by virtue of the fact that $\phi \in \Phi$. We have

$$\begin{aligned}
\alpha^*(\boldsymbol{y}) &= \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} - \sum_{i=1}^{N} p_i \phi(q_i/p_i) \;\middle|\; \sum_{i=1}^{N} p_i \phi(q_i/p_i) \leq \rho \right\} \\
&= \inf_{\lambda \geq 0} \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} - \sum_{i=1}^{N} p_i \phi(q_i/p_i) + \lambda \left(\rho - \sum_{i=1}^{N} p_i \phi(q_i/p_i) \right) \right\} \\
&= \inf_{\lambda \geq 0} \lambda\rho + \sup_{\boldsymbol{q}} \left\{ \boldsymbol{y}'\boldsymbol{q} - (\lambda + 1) \sum_{i=1}^{N} p_i \phi(q_i/p_i) \right\} \\
&= \inf_{\lambda \geq 0} \lambda\rho + (\lambda + 1)\mathbb{E}_{\mathbb{P}} \left[\phi^* \left(\frac{\boldsymbol{y}}{\lambda + 1} \right) \right],
\end{aligned}$$

where the second equality follows by convex duality. From here, we apply Theorem 4.1.1 to obtain the desired result. □

4.3 Probability guarantees

In this section, we derive probability guarantees for the various robust formulations discussed in Section 4.2. The key will be the relationship of the corresponding risk measures to the following coherent risk measure.

Definition 4.3.1. For a random variable $X \in \mathcal{X}$ with $|\Omega| = N$, and $\alpha \in (0, 1]$, the **conditional value-at-risk at level** α, $\mathrm{CVaR}_\alpha(X)$, is

$$\mathrm{CVaR}_\alpha(X) = \sup_{\boldsymbol{q} \in \mathcal{P}_\alpha} \mathbb{E}_{\boldsymbol{q}}[X], \tag{4.11}$$

where $\mathcal{P}_\alpha = \{\boldsymbol{q} \in \Delta^N \mid q_i \leq p_i/\alpha,\ i = 1, \ldots, N\}$.

For continuous distributions, we have $\mathrm{CVaR}_\alpha(X) = \mathbb{E}[X \mid X \geq \mathrm{VaR}_\alpha(X)]$, where

$$\mathrm{VaR}_\alpha(X) = \sup\{x \mid \mathbb{P}\{X \geq x\} \geq \alpha\},$$

called the **value-at-risk at level** α. For detailed treatments on CVaR and its properties, see Rockafellar and Uryasev [102] and Delbaen [49]. It is not hard to see, for any $\alpha \in (0,1]$, we have that $\mathrm{CVaR}_\alpha(X) \geq \mathrm{VaR}_\alpha(X)$, which means

$$\mathrm{CVaR}_\alpha(X) \leq \gamma \Rightarrow \mathbb{P}\{X \geq \gamma\} \leq \alpha.$$

Note that the use of CVaR to bound VaR is not arbitrary. In fact, CVaR is known to be the smallest *law-invariant* (i.e., dependent only on the distribution of the random variable) convex risk measure which upper bounds VaR (e.g., Föllmer and Schied [60], chapter 4.5).

This will be the key fact that we use to prove probability guarantees on the various robust formulations corresponding to different types of convex certainty equivalents. We need the following fact.

Lemma 4.3.1. *Let* $\mu_\rho(X) = \sup_{\boldsymbol{q} \in \mathcal{Q}_\rho} \mathbb{E}_{\boldsymbol{q}}[X]$, *where*

$$\mathcal{Q}_\rho = \left\{ \boldsymbol{q} \in \Delta^N \;\middle|\; \sum_{i=1}^{N} p_i \phi(q_i/p_i) \leq \rho \right\}, \tag{4.12}$$

with $\phi \in \Phi$ *and* $\rho \geq 0$. *Then, for any* $X \in \mathcal{X}$, *with* $|\Omega| = N$ *and* $p_i = 1/N$ *for all* i, *we have* $\mu_\rho(X) \geq \mathrm{CVaR}_\alpha(X)$ *for all* $\alpha \geq \hat{\alpha}(\rho) := \max(1/N, \bar{\alpha}(\rho))$, *where*

$$\bar{\alpha}(\rho) = \inf_{\alpha > 0} \{\alpha \mid \alpha\phi(1/\alpha) + (1-\alpha)\phi(0) + (1/N)\max(\phi(1/\alpha), \phi(0)) \leq \rho\} \tag{4.13}$$

Additionally, we have, in the following important cases:

$$\phi(t) = t\log(t) - t + 1 \;\Rightarrow\; \bar{\alpha}(\rho) \geq e^{-\rho},$$

$$\phi(t) = 2(\sqrt{t} - 1)^2 \;\Rightarrow\; \bar{\alpha}(\rho) \geq \begin{cases} \left(1 - \sqrt{\rho/2 - 1/N}\right)^2, & \text{if } \rho \geq 2/N, \\ 1, & \text{otherwise.} \end{cases}$$

Proof. For general $\phi \in \Phi$, our goal is to find the smallest $\alpha > 0$ such that $\mathcal{P}_\alpha \subseteq \mathcal{Q}_\rho$. Clearly, we can do no better than $1/N$ as $\alpha = 1/N$ implies that $\mathcal{P}_\alpha = \Delta^N$, which proves the first part of the bound.

For given $\alpha \geq 1/N$, $\rho \geq 0$, the desired containment is true if and only if

$\sum_{i=1}^{N} p_i \phi(q_i/p_i) \le \rho$ for all $\boldsymbol{q} \in \mathcal{P}_\alpha$, which is equivalent to this statement being true for all $\boldsymbol{q}$ that are extreme points of $\mathcal{P}_\alpha$, since this set is polyhedral. There is a bijection between the extreme points of $\mathcal{P}_\alpha$ and the family of index sets

$$\mathcal{I} \;=\; \left\{ I \subseteq \{1,\ldots,N\} \;\middle|\; q_i = 1/(N\alpha)\;\forall\, i \in I,\; q_i = 0\;\forall\, i \notin I, \sum_{i\in I} q_i = 1 \right\}.$$

For any $I \in \mathcal{I}$, we have the corresponding extreme point $\boldsymbol{q}$ with $q_i = 1/(N\alpha)$ for $i \in I$, $q_j = 1 - \lfloor N\alpha \rfloor/(N\alpha) \le 1/(N\alpha)$ for some $j \notin I$, and $q_i = 0$ otherwise. Note that $0 \le Nq_j \le 1/\alpha$ and convexity of ϕ together imply that $\phi(Nq_j) \le \max(\phi(1/\alpha), \phi(0))$. Using this fact, for such a $\boldsymbol{q}$, we have

$$\begin{aligned}
\sum_{i=1}^{N} p_i \phi(q_i/p_i) &= \frac{1}{N}\left[\left(\sum_{i\in I}\phi(1/\alpha)\right) + \left(\sum_{i\notin I, i\neq j}\phi(0)\right) + \phi(Nq_j)\right] \\
&\le \frac{1}{N}\left[\left(\sum_{i\in I}\phi(1/\alpha)\right) + \left(\sum_{i\notin I, i\neq j}\phi(0)\right) + \max(\phi(1/\alpha),\phi(0))\right] \\
&= \frac{1}{N}\left[\lfloor N\alpha\rfloor\phi(1/\alpha) + (N - \lfloor N\alpha \rfloor - 1)\phi(0) + \max(\phi(1/\alpha),\phi(0))\right] \\
&\le \frac{1}{N}\left[N\alpha\phi(1/\alpha) + N(1-\alpha)\phi(0) + \max(\phi(1/\alpha),\phi(0))\right] \\
&= \alpha\phi(1/\alpha) + (1-\alpha)\phi(0) + \frac{1}{N}\max\left(\phi(1/\alpha),\phi(0)\right).
\end{aligned}$$

Now finding the smallest such α gives us the desired result in the general case.

For the case of relative entropy, we require $\sum_{i=1}^{N} q_i \log(q_i/p_i) \le \rho$ for all extreme points $\boldsymbol{q}$ of $\mathcal{P}_\alpha$. Using a similar argument to above, we have, for any $I \in \mathcal{I}$, a corresponding $\boldsymbol{q}$ satisfying

$$\begin{aligned}
\sum_{i=1}^{N} q_i \log(Nq_i) &= \frac{1}{N}\left[\sum_{i\in I}\phi(1/\alpha) + \left(1 - \frac{\lfloor N\alpha\rfloor}{N\alpha}\right)\log(Nq_j)\right] \\
&\le \frac{\lfloor N\alpha\rfloor}{N\alpha}\log(1/\alpha) + \left(1 - \frac{\lfloor N\alpha\rfloor}{N\alpha}\right)\log(1/\alpha) \\
&= \log(1/\alpha).
\end{aligned}$$

where the inequality follows from the fact that $Nq_j \le 1/\alpha$ and log being an increasing function.

For the case of Hellinger distance, we require $2\sum_{i=1}(\sqrt{(q_i/p_i)} - 1)^2 \leq \rho$ for all extreme points $\boldsymbol{q}$ of $\mathcal{P}_\alpha$. As above, we have for any $I \in \mathcal{I}$, a corresponding $\boldsymbol{q}$ satisfying

$$
\begin{aligned}
\frac{2}{N}\sum_{i=1}^{N}\left(\sqrt{Nq_i}-1\right)^2 &= \frac{2}{N}\left[\sum_{i\in\mathcal{I}}\left(\sqrt{\frac{1}{\alpha}}-1\right)^2 + \left(\sqrt{Nq_j}-1\right)^2\right] \\
&\leq \frac{2}{N}\left[\lfloor N\alpha\rfloor\left(\sqrt{\frac{1}{\alpha}}-1\right)^2 + 1\right] \\
&\leq 2\alpha\left(\sqrt{\frac{1}{\alpha}}-1\right)^2 + \frac{2}{N},
\end{aligned}
$$

and, when $\rho \geq 2/N$, the first case holds. Otherwise, the bound $\hat{\alpha}(\rho) \geq 1$ holds vacuously. □

With Lemma 4.3.1, we can now state our main result regarding probability guarantees for various classes of convex risk measures.

Theorem 4.3.1. *For any $\phi \in \Phi$ and $\rho \geq 0$, when $\tilde{\boldsymbol{a}}$ satisfies Assumption 4.1.1 with $p_i = 1/N$ for all i, we have the following implications for all $\epsilon \geq 0$:*

$$
\begin{array}{llcll}
\textit{(OCE)} & \mu_\phi(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b & \Rightarrow & \mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \geq b+\epsilon\} & \leq \hat{\alpha}(\epsilon), \\
\textit{(robust)} & R_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b & \Rightarrow & \mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \geq b+\epsilon\} & \leq \hat{\alpha}(\rho), \\
\textit{(comprehensive robust)} & C_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b & \Rightarrow & \mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \geq b+\epsilon\} & \leq \hat{\alpha}(\rho+\epsilon), \\
\textit{(soft robust)} & S_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b & \Rightarrow & \mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \geq b+\epsilon\} & \leq \hat{\alpha}(\min(\epsilon,\rho)).
\end{array}
$$

Proof. For the case of $R_{\phi,\rho}$, the result holds directly as a result of Lemma 4.3.1. We prove the implication for the case of $C_{\phi,\rho}$, and the rest follow in similar fashion. For any $\epsilon \geq 0$, consider the penalty function

$$
\alpha_\rho^\epsilon(\boldsymbol{q}) = \begin{cases} 0, & \text{if } \sum_{i=1}^{N} p_i\phi(q_i/p_i) \leq \rho+\epsilon, \\ +\infty, & \text{otherwise.} \end{cases}
$$

Name	Risk measure	Robust guarantees*	$\mathbb{P}\{\tilde{\boldsymbol{a}}'\boldsymbol{x} \geq b+\epsilon\}$
μ_ϕ	$\inf_{\nu\in\mathbb{R}} \{\nu + \mathbb{E}_{\mathbb{P}}[\phi^*(X-\nu)]\}$	$\beta_\phi(\boldsymbol{a})\ \forall\ \boldsymbol{a} \in \operatorname{conv}(\mathcal{A})$	$\hat{\alpha}(\epsilon)$
$R_{\phi,\rho}$	$\inf_{\lambda>0} \{\rho\lambda + \lambda\mu_\phi(X/\lambda)\}$	$\begin{cases} 0 \text{ if } \boldsymbol{a} \in \mathcal{U}_\phi(\rho), \\ +\infty \text{ otherwise.} \end{cases}$	$\hat{\alpha}(\rho)$
$C_{\phi,\rho}$	$\inf_{\lambda\in(0,1]} \{\rho\lambda + \lambda\mu_\phi(X/\lambda)\}$	$\begin{cases} 0 \text{ if } \boldsymbol{a} \in \mathcal{U}_\phi(\rho), \\ \beta_\phi(\boldsymbol{a}) - \rho \text{ otherwise.} \end{cases}$	$\hat{\alpha}(\rho+\epsilon)$
$S_{\phi,\rho}$	$\inf_{\lambda>0} \{\rho\lambda + (\lambda+1)\mu_\phi(X/(\lambda+1))\}$	$\begin{cases} \beta_\phi(\boldsymbol{a}) \text{ if } \boldsymbol{a} \in \mathcal{U}_\phi(\rho), \\ +\infty \text{ otherwise.} \end{cases}$	$\hat{\alpha}(\min(\epsilon,\rho))$

Table 4.1: *Summary of the properties of the various risk measures discussed. The robust guarantees (*) represent the smallest value $\beta(\boldsymbol{a})$ such that $\boldsymbol{a}'\boldsymbol{x} \leq b + \beta(\boldsymbol{a})$ for each $\boldsymbol{a} \in \operatorname{conv}(\mathcal{A})$.*

For any $\boldsymbol{q} \in \Delta^N$, we have $\max(\sum_{i=1}^N p_i\phi(q_i/p_i) - \rho, 0) \leq \alpha_\rho^\epsilon(\boldsymbol{q}) + \epsilon$, implying

$$\begin{aligned} C_{\phi,\rho}(X) &= \sup_{\boldsymbol{q}\in\Delta^N} \left\{ \mathbb{E}_{\boldsymbol{q}}[X] - \max\left(\sum_{i=1}^N p_i\phi(q_i/p_i) - \rho, 0 \right) \right\} \\ &\geq \sup_{\boldsymbol{q}\in\Delta^N} \left\{ \mathbb{E}_{\boldsymbol{q}}[X] - \alpha_\rho^\epsilon(\boldsymbol{q}) \right\} - \epsilon \\ &= \sup_{\boldsymbol{q}\in\mathcal{Q}_{\rho+\epsilon}} \{\mathbb{E}_{\boldsymbol{q}}[X]\} - \epsilon \\ &= R_{\phi,(\rho+\epsilon)}(X) - \epsilon. \end{aligned}$$

Therefore, $C_{\phi,\rho}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b \Rightarrow R_{\phi,\rho+\epsilon}(\tilde{\boldsymbol{a}}'\boldsymbol{x}) \leq b + \epsilon$, and, invoking Lemma 4.3.1, the result follows. □

Table 4.1 summarizes the risk measures we have discussed.

4.4 Application to portfolio optimization

Here we apply some of the models discussed in the previous sections to an asset allocation problem using real-world financial data. Our focus is exploring and comparing the empirical performance of various classes of convex risk measures in terms of the observed distribution of returns. For theoretical insights and structural results using general convex risk measures, see Lüthi and Doege [82]; for a theoretical treatment of OCEs applied to portfolio problems, see Ben-Tal and Teboulle [21].

4.4.1 Problem set-up

We consider an investor with wealth level $w > 0$ who wishes to allocate his wealth among n assets. The decision vector $\boldsymbol{x} \in \mathbb{R}^n$ denotes the vector of weights the investor allocates to each asset for the current time period. We require $\boldsymbol{x} \in X := \{\boldsymbol{x} \in \mathbb{R}^n \mid \boldsymbol{e}'\boldsymbol{x} = 1,\ \boldsymbol{x} \geq \boldsymbol{0}\}$. In a true real-world setting, of course, there will be many others constraints in addition to a no short-sales restriction; as we are largely examining the relative performance of the various risk measures, however, and we expect similar relative results with the addition of these constraints.

The n assets have an associated, random return vector $\tilde{\boldsymbol{r}}$ over the time period, with $\mathbb{P}\{\tilde{\boldsymbol{r}} \geq \boldsymbol{0}\} = 1$. The final wealth after a single period is therefore just $w \cdot \tilde{\boldsymbol{r}}'\boldsymbol{x}$. Denote $\mathbb{E}[\tilde{\boldsymbol{r}}]$ by $\hat{\boldsymbol{r}}$. At each time period, the investor solves the problem

$$\begin{array}{ll} \text{maximize} & \hat{\boldsymbol{r}}'\boldsymbol{x} \\ \text{subject to} & -\frac{1}{w}\mu(-w \cdot \tilde{\boldsymbol{r}}'\boldsymbol{x}) \geq \gamma, \\ & \boldsymbol{x} \in X, \end{array} \tag{4.14}$$

where $\mu \in \mathcal{M}$ is a convex risk measure specified by the investor, and $\gamma \geq 0$ is a parameter reflecting the investor's target risk level.

Remark 4.4.1. The presence of the minus signs in the risk constraint simply accounts for the fact that we defined convex risk measures to measure *losses* in Definition 2.2.4. For instance, we have

$$\begin{aligned} -\mathrm{VaR}_\alpha(-X) &= -\sup_x \{-x \mid \mathbb{P}\{-X \geq -x\} \geq \alpha\} \\ &= \inf_x \{x \mid \mathbb{P}\{X \leq x\} \geq \alpha\} \\ &= \inf_x \{x \mid \mathbb{P}\{X \geq x\} \leq 1-\alpha\} \\ &= \mathrm{VaR}_{1-\alpha}(X). \end{aligned}$$

where the last equality holds provided the distribution is continuous. This, in turn, implies we have

$$\begin{aligned} -\mathrm{CVaR}_\alpha(-X) &= -\mathbb{E}\left[-X \mid -X \geq \mathrm{VaR}_\alpha(-X)\right] \\ &= \mathbb{E}\left[X \mid X \leq -\mathrm{VaR}_\alpha(-X)\right] \\ &= \mathbb{E}\left[X \mid X \leq \mathrm{VaR}_{1-\alpha}(X)\right]. \end{aligned}$$

If $\mu = \mathrm{CVaR}_1(\cdot)$, then $-\mu(-X) = \mathbb{E}[X]$. Similarly, $-\mathrm{CVaR}_0(-X) = x_{\min}$.

Thus, if μ is $\mathrm{CVaR}_\alpha(\cdot)$, the constraint in (4.14) simply states that the expected value of the portfolio return given that the return is less than its $1-\alpha$ VaR, is no smaller than the level γ. In this setting, and, in fact, generally, one would expect an investor to have $\gamma \approx 1$.

It is not obvious that (4.14) is feasible for a particular value of $\gamma > 0$. Note, however, that if we assume the presence of a risk-free asset with constant return r_f, the problem is feasible for all $\gamma \leq r_f$. Indeed, if the investor invests all of his wealth in the risk free asset, we have

$$-\frac{1}{w}\mu(-w \cdot r_f) = -\frac{1}{w}\left[\mu(0) - wr_f\right] = r_f,$$

provided that $\mu(0) = 0$, which is in fact the case since $\mu \in \mathcal{M}$.

Note that when μ is coherent, by the positive homogeneity property, the wealth level w is irrelevant. For an arbitrary, convex risk measure, however, the wealth level does impact the risk constraint in (4.14). Specifically, we have the following, straightforward fact.

Proposition 4.4.1. Let $\mu_w(X) = (1/w)\mu(wX)$, where $w \geq 0$ and $\mu : \mathcal{X} \to \mathbb{R}$ satisfies $\mu \in \mathcal{M}$. Then we have the following:

(a) $\mu_w(X)$ is increasing in w.

(b) $\lim_{w \to \infty} \mu_w(X) = x_{\max}$, where $x_{\max} = \max_{\omega \in \Omega} X(\omega)$.

(c) $\lim_{w \to 0} \mu_w(X) = \mathbb{E}[X]$, provided that α is strictly convex.

Proof. The proof in each case follows easily by noting, from the representation theorem for convex risk measures, that

$$\begin{aligned}\mu_w(X) &= \frac{1}{w}\sup_{\boldsymbol{q}\in\Delta^N}\left\{\mathbb{E}_{\boldsymbol{q}}\left[wX\right]-\alpha(\boldsymbol{q})\right\}\\ &= \sup_{\boldsymbol{q}\in\Delta^N}\left\{\mathbb{E}_{\boldsymbol{q}}\left[X\right]-\frac{1}{w}\alpha(\boldsymbol{q})\right\}.\end{aligned}$$

From here, (a) clearly holds, since $\mu \in \mathcal{M}$ implies that $\alpha(\boldsymbol{q}) \geq 0$ for all $\boldsymbol{q} \in \Delta^N$. (b) is clear. Finally, for (c), we see that, as $w \to 0$, the second term in the supremum dominates, and therefore, in the limit, any $\boldsymbol{q}$ which achieves the sup must satisfy $\alpha(\boldsymbol{q}) = 0$. Since $\mu \in \mathcal{M}$, however, $\mathbb{P}$ is one such measure, and if α is strictly convex, it is the only such measure, which gives the desired result. □

Proposition 4.4.1 means, specifically, that an investor whose risk preferences do not change over time will become more conservative as their wealth level grows. Ben-Tal and Teboulle [21] prove an analogous result for more general probability spaces, but over the more restrictive class of OCE risk measures.

4.4.2 Empirical data

For our empirical study, we use monthly historical returns for $n = 11$ publicly traded asset classes over the period from April, 1981 through February, 2006. The asset classes are listed in Table 4.2. In Table 4.3, we list the realized CVaR and VaR for the various assets based on the data from this time period. Note again that $-\text{CVaR}_\alpha(-R)$ can be interpreted as that the expected value of the asset's return, given that the return is in the lower α-tail of its distribution; in particular, $-\text{CVaR}_1(-R)$ is the expected return of the asset.

4.4.3 Experiment and results

Using the data described in the previous section, we solved (4.14) for several different risk measures and compared the results. In this setup, we used a sliding window of the

Asset name	*Symbol*	*Category*
S&P 500 Index	SP500	U.S. equity
Russell Mid-cap Index	RMidC	U.S. equity
Russell 2000 Index	R2000	U.S. equity
MSCI EAFE Index	MSCIEAFE	International equity
MSCI Emerging Market Index	MSCIEmer	International equity
NAREIT Index	NAREIT	Real estate
Lehman Brothers' U.S. Aggregate Index	LBUS	U.S. bond
Lehman Brothers' U.S. Corporate High Yield	LBHY	U.S. corporate bond
Global Governments Bond	GlobGovBnd	International bond
Emerging Markets Bond	EmerMktBnd	International bond
3-month LIBOR	LIBOR	Cash

Table 4.2: *Descriptions for the various asset classes used in the experiment.*

	α			
$-\text{CVaR}_\alpha(-R)$ / $-\text{VaR}_\alpha(-R)$	1.0	.50	.10	.05
SP500	1.141	1.013	0.802	0.752
	1.480	1.156	0.871	0.776
RMidC	1.154	1.020	0.846	0.785
	1.508	1.141	0.880	0.851
R2000	1.133	0.976	0.805	0.730
	1.638	1.131	0.847	0.811
MSCIEAFE	1.139	0.965	0.794	0.743
	1.858	1.120	0.869	0.764
MSCIEmer	1.169	0.900	0.731	0.682
	1.796	1.033	0.795	0.717
NAREIT	1.134	0.998	0.885	0.775
	1.546	1.107	0.941	0.899
LBUS	1.098	1.045	1.018	1.012
	1.319	1.087	1.024	1.017
LBHY	1.115	1.040	0.989	0.980
	1.509	1.082	1.004	0.981
GlobGovBnd	1.085	0.989	0.924	0.912
	1.347	1.067	0.945	0.923
EmerMktBnd	1.139	1.037	0.863	0.863
	1.388	1.132	0.880	0.866
LIBOR	1.062	1.040	1.017	1.012
	1.144	1.059	1.021	1.016

Table 4.3: *CVaR and VaR of annualized returns for the 11 asset classes from April, 1981 through February, 2006.*

past three years of returns as the sample data for solving (4.14) (therefore, $N = 36$, as we have monthly return data), implemented the optimal portfolio over the following year's worth of data, then re-balanced.[2] We repeated this process over each year within the entire data range and tabulated performance statistics (described below) for the various risk measures. The parameters used were $\gamma = 1$ and and initial wealth level of $w = 50$.

Specifically, we examined the performance of this approach under the following risk measures:

(a) CVaR_α, for $\alpha = 0.1, 0.2, \ldots, 1$;

(b) $R_{\phi,\rho}$ (standard robustness);

(c) $C_{\phi,\rho}$ (comprehensive robustness);

(d) $S_{\phi,\rho}$ (soft robustness).

For all trials, we used $\phi(t) = t\log(t) - t + 1$, i.e., an exponential utility function $u(t) = 1 - e^{-t}$, or $\phi^*(t) = e^t - 1$. Guided by Theorem 4.3.1, we therefore chose $\rho = \log(1/\alpha)$ in all cases.

In Table 6.3, we see a comparison of the realized performance in terms of CVaR_α and expected return for the various risk measures. Table 4.5 shows the probability that the realized return (annualized) drops below a pre-specified threshold for the different risk measures. Finally, Figure 4-1 shows the cumulative return over the historical time period for the risk measures under various choices of α (and hence ρ).

Some key observations from these empirical results are the following:

1. The standard robust ($R_{\phi,\rho}$) and comprehensive robust measures generally did the the best in terms of closeness to achieving a realized CVaR_α of at least 1. This matches intuition, as they are the most conservative risk measures of the four.

[2] We did not account for transactions costs in our results. The turnover levels for the various risk measures, however, were similar, so the relative performance should be similar with such costs included.

2. Somewhat surprisingly, the soft robust measure significantly outperformed CVaR in terms of realized risk (measured in terms of CVaR_α for $\alpha = 0.3, \ldots, 0.8$). This risk reduction was not always offset by a decrease in rate of return (e.g., for $\alpha = 0.8$ and $\alpha = 0.9$, the soft robust measure significantly outperformed CVaR; this is clearly demonstrated in Figure 4-1).

3. The comprehensive robust solutions were quite similar to the standard robust solutions for $\alpha \leq 0.5$, as was the corresponding performance and risk. For $\alpha > 0.5$, however, $R_{\phi,\rho}$ offered an average of +3.68% expected return over $C_{\phi,\rho}$. The probability of bad performance, however, was significantly higher for large α for $R_{\phi,\rho}$ (e.g., 25.0% vs. 2.0% of the (annualized) monthly return dropping below 0.8 for the case $\alpha = 1$).

4. Over the 10 values of α, the average benefit of the soft robust measure over $R_{\phi,\rho}$ was +0.38% of expected return. This was traded off at a cost of an average of 1.97% of realized CVaR_α. The probability of bad performance, shown in Table 4.5, was similar for $\alpha > 0.5$ for both $R_{\phi,\rho}$ and $S_{\phi,\rho}$.

5. Although CVaR had the highest expected performance in many of the cases, this was not always so (see point 2 above), and, in every case listed in Table 4.5, the investment strategies using CVaR had the highest probability of bad performance. Figure 4-1 emphasizes this graphically; note the large dips in cumulative return for the CVaR investment strategy, particularly for the case $\alpha = 0.8$.

	$-\mathrm{CVaR}_\alpha(-R)$ (annualized)				$\mathbb{E}[R]$ (annualized)			
α	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$
.1	0.941	0.954	0.879	0.954	1.062	1.062	1.069	1.062
.2	0.943	0.975	0.931	0.975	1.065	1.063	1.069	1.063
.3	0.943	0.984	0.955	0.987	1.072	1.064	1.070	1.064
.4	0.915	0.990	0.970	0.998	1.079	1.066	1.071	1.065
.5	0.906	0.992	0.978	1.008	1.107	1.069	1.074	1.066
.6	0.896	0.993	0.985	1.018	1.128	1.075	1.079	1.067
.7	0.899	0.999	0.996	1.029	1.120	1.088	1.093	1.067
.8	0.945	1.007	1.005	1.040	1.117	1.118	1.120	1.068
.9	1.020	1.039	1.038	1.052	1.118	1.125	1.125	1.068

Table 4.4: *Realized conditional-value-at-risk (left) and expected return (right) for the experiment run with the 4 types of risk measures at different levels of* α. $\rho = \log(1/\alpha)$ *and* $\phi(t) = t\log(t) - t + 1$ *in all cases. Recall that* $-\mathrm{CVaR}_\alpha(-R) = \mathbb{E}[R \mid R \leq \mathrm{VaR}_{1-\alpha}(R)]$.

	$\mathbb{P}\{\text{Return} < 1\}$				$\mathbb{P}\{\text{Return} < 0.9\}$			
α	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$
0.1	0.127	0.095	0.190	0.091	0.020	0.012	0.040	0.020
0.2	0.187	0.119	0.198	0.119	0.032	0.020	0.040	0.020
0.3	0.218	0.151	0.210	0.139	0.052	0.024	0.044	0.024
0.4	0.266	0.190	0.222	0.151	0.115	0.028	0.048	0.028
0.5	0.282	0.210	0.238	0.167	0.171	0.040	0.071	0.032
0.6	0.310	0.250	0.258	0.175	0.230	0.083	0.099	0.032
0.7	0.321	0.266	0.274	0.190	0.266	0.107	0.119	0.032
0.8	0.357	0.294	0.310	0.190	0.290	0.194	0.210	0.032
0.9	0.373	0.329	0.329	0.190	0.302	0.266	0.266	0.040
1.0	0.373	0.365	0.369	0.194	0.302	0.302	0.302	0.040
	$\mathbb{P}\{\text{Return} < 0.8\}$				$\mathbb{P}\{\text{Return} < 0.7\}$			
α	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$	CVaR_α	$R_{\phi,\rho}$	$S_{\phi,\rho}$	$C_{\phi,\rho}$
0.1	0.004	0.000	0.020	0.000	0.0040	0	0.0079	0
0.2	0.008	0.004	0.016	0.004	0.0079	0	0.0079	0
0.3	0.016	0.008	0.016	0.004	0.0119	0.0040	0.0079	0
0.4	0.052	0.012	0.016	0.008	0.0198	0.0040	0.0119	0.0040
0.5	0.091	0.016	0.020	0.012	0.0516	0.0119	0.0119	0.0040
0.6	0.155	0.028	0.036	0.012	0.0992	0.0159	0.0159	0.0040
0.7	0.206	0.063	0.079	0.016	0.1627	0.0278	0.0317	0.0040
0.8	0.238	0.099	0.111	0.020	0.1905	0.0635	0.0675	0.0079
0.9	0.262	0.198	0.198	0.020	0.2103	0.1468	0.1508	0.0079
1.0	0.262	0.250	0.262	0.020	0.2103	0.1984	0.2024	0.0079

Table 4.5: *Probabilities that the realized, monthly returns (annualized) are less than 1 (upper left), 0.9 (upper right), 0.8 (lower left), and 0.7 (lower right) for the various risk measures in the experiment.*

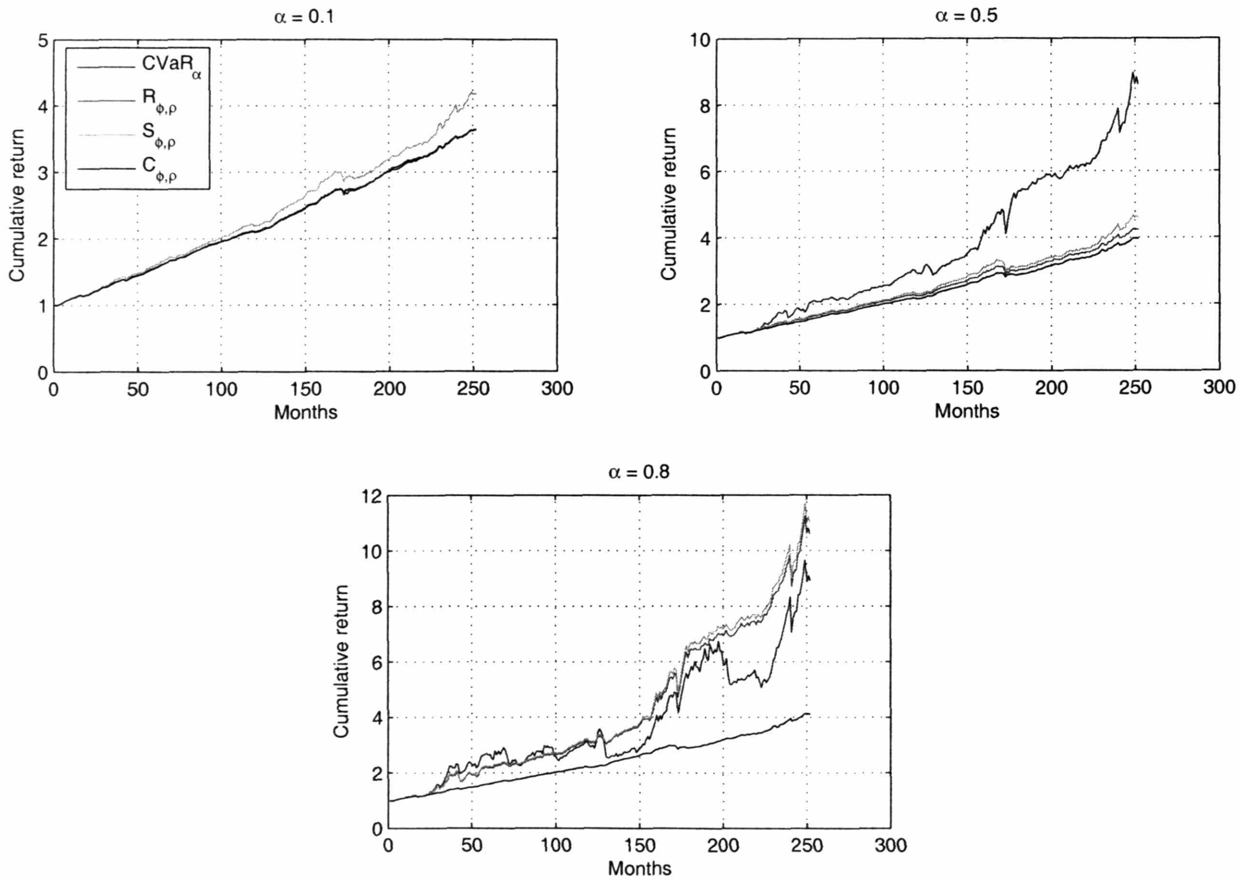

Figure 4-1: *Realized cumulative return under the 4 risk measures for various* α.

Chapter 5

Extensions and probability guarantees

Chapters 3 and 4 developed our approach with coherent and convex measures, respectively, to the case of linear optimization under uncertainty. While linear optimization certainly captures a wide range of applications, there are nonetheless many problems with inherent nonlinearities. In this chapter, we briefly extend the theory of the previous two chapters. We also consider the issue of implied probability guarantees by solving scenario approximations of the risk-averse problems.

5.1 Risk measures and conic optimization

In this section, we explain the problem of interest and provide an overview of our approach to it. We begin by defining the nominal problem.

5.1.1 Problem setup

Definition 5.1.1. A **conic optimization problem** is a problem of the form

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \boldsymbol{A}\boldsymbol{x} - \boldsymbol{b} \preceq_{\mathbb{K}} \boldsymbol{0}, \end{aligned} \tag{5.1}$$

where $\mathbb{K} \subseteq \mathcal{V}$ is a closed, convex cone contained in a finite-dimensional vector space $\mathcal{V}$.

For full generality we consider cones in arbitrary, finite-dimensional vector spaces $\mathcal{V}$, but we note that we are essentially always interested in the case when $\mathcal{V} = \mathbb{R}^m$, $\mathcal{V} = \mathbb{S}^m$ (i.e., the space of symmetric matrices of dimension m), or $\mathcal{V} = \mathbb{R}^{m_1} \times \mathbb{S}^{m_2}$. Throughout our work, we will operate under the following, mild assumption on $\mathbb{K}$, which will allows us to exploit strong duality results to the fullest.

Assumption 5.1.1. The cone $\mathbb{K}$ in Problem (5.1), in addition to being closed and convex, is also pointed and has a nonempty interior. We will say such cones $\mathbb{K}$ are *regular*.

Some practically relevant cones are the following:

1. The nonnegative orthant, $\mathbb{R}^m_+$.
2. The second-order (or *Lorentz*) cone, i.e.,

$$\mathbb{L}^m \quad = \quad \left\{ \boldsymbol{x} \in \mathbb{R}^m \;\middle|\; \sqrt{\sum_{i=1}^{m-1} x_i^2} \leq x_m \right\}.$$

3. The positive semidefinite cone $\mathbb{S}^m_+$, i.e., the cone of all symmetric, $m \times m$ matrices which are positive semidefinite.

Another relevant cone is the epigraph of certain classes of convex functions.

Lemma 5.1.1. *Let $f : \mathcal{V} \to \mathbb{R}$ be a convex, lower semi-continuous function satisfying the following properties:*
(a) $|f(\boldsymbol{x})| < \infty \;\forall\; \boldsymbol{x} \in \mathcal{V}$,
(b) $f(k\boldsymbol{x}) \leq kf(\boldsymbol{x}) \;\forall\; k \geq 0,\; \boldsymbol{x} \in \mathcal{V}$.
Then the set

$$\mathbb{K}_f \quad = \quad \{(\boldsymbol{x}, t) \in \mathcal{V} \times \mathbb{R} \mid f(\boldsymbol{x}) - t \leq 0\} \tag{5.2}$$

is a closed, convex cone with a nonempty interior.

Proof. Closedness follows from standard results in convex analysis relating to convex, lower semi-continuous functions (e.g., Rockafellar [101]). That $\mathbb{K}_f$ is nonempty follows by Property **(a)**; for any $\boldsymbol{x} \in \mathcal{V}$, the point $(\boldsymbol{x}, t)$ where $f(\boldsymbol{x}) < t < \infty$ is in the interior of $\mathbb{K}_f$. To show that $\mathbb{K}_f$ is a convex cone, we need to show that it is closed under addition and nonnegative scaling. The latter follows directly by Property **(b)**. For closure under addition, consider $(\boldsymbol{x}_1, t_1)$, $(\boldsymbol{x}_2, t_2) \in \mathbb{K}_f$. Then we have

$$\begin{aligned} f(\boldsymbol{x}_1 + \boldsymbol{x}_2) &= f(2(1/2\boldsymbol{x}_1 + 1/2\boldsymbol{x}_2)) \\ &\leq 2f(1/2\boldsymbol{x}_1 + 1/2\boldsymbol{x}_2) \qquad \text{(Property } \mathbf{(b)}) \\ &\leq f(\boldsymbol{x}_1) + f(\boldsymbol{x}_2) \qquad \text{(convexity of } f) \\ &\leq t_1 + t_2, \end{aligned}$$

so $(\boldsymbol{x}_1 + \boldsymbol{x}_2, t_1 + t_2) \in \mathbb{K}_f$. □

Note that any norm in $\mathcal{V}$ satisfies the requirements of Lemma 5.1.1. For instance, the second-order cone $\mathbb{L}^m$ is an example of $\mathbb{K}_f$ with the function f being the Euclidean norm.

The fact that the objective is linear in Problem (5.1) is not particulary limiting. In fact, given a conic optimization problem similar to (5.1) but with a nonlinear objective function $f(\boldsymbol{x})$ satisfying the requirements of Lemma 5.1.1, we may transform the problem to an equivalent one in the form of Problem (5.1). This may be done by introducing an extra variable t and taking the product of the original cone in the constraints with the epigraph cone $\mathbb{K}_f$. We then proceed to minimize t in the objective.

We are interested in the case when there is uncertainty in the problem data $(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$ associated with Problem (5.1). As a consequence of the above commentary, any uncertainty in $\boldsymbol{c}$ is irrelevant; we may always introduce epigraph form if $\boldsymbol{c}$ is unknown, then minimize the certain quantity t. Similarly, uncertainty in $\boldsymbol{b}$ may be ignored, as we may always add a extra decision variable x_{n+1}, aggregate the uncertain vector $\boldsymbol{b}$ with the matrix $\boldsymbol{A}$, and constrain $x_{n+1} = 1$. The transformations in the case of uncertainty in $\boldsymbol{b}$ or $\boldsymbol{c}$ are easy enough to see that we do not state them rigorously.

Without loss of generality, then, we are interested in the following problem.

Definition 5.1.2. A conic optimization problem with **data uncertainty** is a problem of the form

$$\begin{array}{ll} \text{minimize} & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} & \tilde{\boldsymbol{A}}\boldsymbol{x} - \boldsymbol{b} \preceq_{\mathbb{K}} \boldsymbol{0}, \end{array} \tag{5.3}$$

where $\tilde{\boldsymbol{A}}$ is an uncertain matrix, and $\mathbb{K}$ is regular.

As stated, Problem (5.3) is ill-posed; the cone membership constraint is meaningless in the absence of an uncertainty model and a metric for measuring the quality of feasibility. As before, we will use data for the uncertain matrix and a risk measure on the uncertain constraint to accomplish this. There are three key elements to our approach:

1. A *data set of size* N for the uncertain matrix $\tilde{\boldsymbol{A}}$. We denote the data set by $\mathcal{A}_N$; the ith element of $\mathcal{A}_N$ is denoted by $\boldsymbol{A}_i$.

2. A convex function $g : \mathcal{V} \to \mathbb{R}_+$ satisfying the following:

$$g(\boldsymbol{y}) \leq 0 \quad \Leftrightarrow \quad \boldsymbol{y} \preceq_{\mathbb{K}} \boldsymbol{0}. \tag{5.4}$$

3. A coherent risk measure $\mu : \mathcal{X} \to \mathbb{R}$.

g is essentially a type of convex indicator function for membership in $\mathbb{K}$. Some examples of possible choices of g are:

- $g(\boldsymbol{y}) = \sum_{i=1}^{m} \max(0, y_i)$ and $\mathbb{K} = \mathbb{R}^m_+$.
- $g(\boldsymbol{y}) = \|\boldsymbol{y}^+\|$, where $\|\cdot\|$ is any norm in $\mathbb{R}^m$ and $\mathbb{K} = \mathbb{R}^m_+$.
- $g(\boldsymbol{Y}) = \|\Lambda^+(\boldsymbol{Y})\|$, where $\|\cdot\|$ is any norm in $\mathbb{R}^m$, $\Lambda(\boldsymbol{Y})$ denotes the vector of m eigenvalues of a symmetric $m \times m$ matrix $\boldsymbol{Y} \in \mathbb{S}^m$, and $\mathbb{K} = \mathbb{S}^m_+$.

- $g(\boldsymbol{y}, \boldsymbol{Y}) = \| \left(\boldsymbol{y}^+, \boldsymbol{\Lambda}^+(\boldsymbol{Y})\right) \|$, where $\mathbb{K} = \mathbb{R}_+^{m_1} \times \mathbb{S}_+^{m_2}$.

- $g(\boldsymbol{y}) = \max(0, \|\boldsymbol{y}_{m-1}\|_2 - y_m)$, where $\mathbb{K} = \mathbb{L}^m$ and $\boldsymbol{y}_{m-1}$ is the vector of the first $m-1$ components of $\boldsymbol{y}$.

We now describe formally the problem we solve.

Definition 5.1.3. The (μ, g) **risk averse counterpart** of Problem (5.3) at level $\gamma \geq 0$ is the problem

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \mu\left(g\left(\tilde{\boldsymbol{A}}\boldsymbol{x} - \boldsymbol{b}\right)\right) \leq \gamma, \end{aligned} \tag{5.5}$$

in decision variables $\boldsymbol{x} \in \mathbb{R}^n$.

In short, Problem (5.5) ensures that any feasible solution $\boldsymbol{x}$ is no worse than γ-infeasible (measured by the function g) to $\tilde{\boldsymbol{A}}\boldsymbol{x} - \boldsymbol{b} \leq_{\mathbb{K}} \boldsymbol{0}$ for all scenarios captured by the risk measure μ. Although Definition 5.1.3 makes no mention of an uncertainty model for $\tilde{\boldsymbol{A}}$, clearly we must have a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ for $\tilde{\boldsymbol{A}}$ in order to evaluate μ. As in the linear case, we will utilize the following probability space, which works quite naturally in a data-driven context.

Assumption 5.1.2. The uncertain matrix $\tilde{\boldsymbol{A}}$ has support $\mathcal{A}_N = \{\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N\}$ and a corresponding probability measure $\mathbb{P} \in \Delta^N$.

5.1.2 Tractability of the approach

We now show that (5.5) is tractable for various classes of convex risk measures. We have all the necessary machinery to immediately state the main result.

Theorem 5.1.1. *Let $\mu \in \mathcal{M}$ be a convex risk measure with associated penalty function α. Under Assumption 5.1.2, Problem (5.5) is equivalent to the convex optimization* *problem*

$$\begin{array}{lll} \textit{minimize} & \boldsymbol{c}'\boldsymbol{x} & \\ \textit{subject to} & \nu + \alpha^*(\boldsymbol{y} - \nu\boldsymbol{e}) \leq \gamma, & \\ & g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i, \qquad 1 = 1, \ldots, N, & \end{array} \tag{5.6}$$

in decision variables $(\boldsymbol{x}, \boldsymbol{y}, \nu) \in \mathbb{R}^{n+N+1}$.

Proof. Invoking the representation theorem (Theorem 3.1.1) for convex risk measures, and using Assumption 5.1.2, we have

$$\begin{aligned} \mu(g(\tilde{\boldsymbol{A}}\boldsymbol{x} - \boldsymbol{b})) &= \sup_{\boldsymbol{q} \in \Delta^N} \left\{ \sum_{i=1}^N q_i \underbrace{g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b})}_{z_i} - \alpha(\boldsymbol{q}) \right\} \\ &= \inf_{\nu \in \mathbb{R}} \sup_{\boldsymbol{q} \in \text{dom } \alpha} \{\boldsymbol{z}'\boldsymbol{q} - \alpha(\boldsymbol{q}) + \nu(1 - \boldsymbol{e}'\boldsymbol{q})\} \\ &= \inf_{\nu \in \mathbb{R}} \left\{ \nu + \sup_{\boldsymbol{q} \in \text{dom } \alpha} \{(\boldsymbol{z} - \nu\boldsymbol{e})'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \right\} \\ &= \inf_{\nu \in \mathbb{R}} \{\nu + \alpha^*(\boldsymbol{z} - \nu\boldsymbol{e})\}, \end{aligned}$$

where we used the shorthand notation $z_i = g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b})$ for convenience. Now recall that $\mu \in \mathcal{M}$ requires dom $\alpha \subseteq \mathbb{R}^N_+$, which, in turn, implies that α^* is nondecreasing. Indeed, consider $\boldsymbol{z}_1 \in \mathbb{R}^N$, $\boldsymbol{z}_2 \in \mathbb{R}^N$, with $\boldsymbol{z}_2 \geq \boldsymbol{z}_1$. Then

$$\begin{aligned} \alpha^*(\boldsymbol{z}_1) &= \sup_{\boldsymbol{q} \in \text{dom } \alpha} \{\boldsymbol{z}_1'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \\ &= \sup_{\boldsymbol{q} \in \text{dom } \alpha} \{\boldsymbol{z}_2'\boldsymbol{q} - \alpha(\boldsymbol{q}) + (\boldsymbol{z}_1 - \boldsymbol{z}_2)'\boldsymbol{q}\} \\ &\leq \sup_{\boldsymbol{q} \in \text{dom } \alpha} \{\boldsymbol{z}_2'\boldsymbol{q} - \alpha(\boldsymbol{q})\} \\ &= \alpha^*(\boldsymbol{z}_2), \end{aligned}$$

where the inequality follows from the fact that any $\boldsymbol{q}$ in the sup must be nonnegative, and $\boldsymbol{z}_2 \geq \boldsymbol{z}_1$. As a consequence of α^* being nondecreasing, then, we may replace the $g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b})$ within the argument of α^* with auxiliary variables y_i with $g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i$ to achieve the desired result. □

Thus, provided we have an oracle for efficiently evaluating the function α^*, Problem (5.5) is efficiently solvable using standard convex optimization techniques. We can also explicitly state (5.5) for the types of risk measures discussed in the previous two chapters.

Corollary 5.1.1. Under the conditions of Theorem 5.1.1, the risk constraint in (5.5), for the following choices of $\mu \in \mathcal{M}$, is equivalent to the sets of constraints:

$$\mu \text{ coherent} \quad \Rightarrow \quad \begin{cases} \sup_{\boldsymbol{q} \in \mathcal{Q}} \boldsymbol{q}'\boldsymbol{y} \leq \gamma, \\ g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i, \qquad 1 = 1, \ldots, N. \end{cases}$$

$$\mu = \mu_\phi \quad \Rightarrow \quad \begin{cases} \nu + \sum_{i=1}^{N} \phi^*(y_i - \nu) \leq \gamma, \\ g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i, \qquad 1 = 1, \ldots, N. \end{cases}$$

$$\mu = S_{\phi,\rho} \quad \Rightarrow \quad \begin{cases} \nu + \rho\lambda + (\lambda + 1) \sum_{i=1}^{N} p_i \phi^* \left(\frac{y_i - \nu}{\lambda + 1}\right), \\ g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i, \qquad 1 = 1, \ldots, N, \\ \lambda \geq 0. \end{cases}$$

$$\mu = C_{\phi,\rho} \quad \Rightarrow \quad \begin{cases} \nu + \rho\lambda + \lambda \sum_{i=1}^{N} p_i \phi^* \left(\frac{y_i - \nu}{\lambda}\right), \\ g(\boldsymbol{A}_i\boldsymbol{x} - \boldsymbol{b}) \leq y_i, \qquad 1 = 1, \ldots, N, \\ \lambda \in (0, 1]. \end{cases}$$

Here, $\mathcal{Q}$ denotes the generating family for μ when μ is coherent.

Proof. Each case follows directly from Theorem 5.1.1 and the results developed in Chapters 3 and 4. □

We emphasize that the sets of constraints for each of the cases in Corollary 5.1.1 are convex in their corresponding variables.

5.1.3 Connection to robust optimization

The emphasis of Chapters 3 and 4 was the connection of constraints based on risk measures to robustness constraints. In particular, we essentially showed that when the underlying problem was linear, the correspondence between robustness and risk measures was one-to-one. In the case of conic optimization, however, we can only show that constraints with risk measures *imply* robustness, but not vice versa. In other words, solutions based on risk constraints, which we have shown to be tractable when the risk measure is convex, are *robust feasible* but, in general, not *robust optimal*, where we define "robustness" in a particular way. We now make this explicit.

Proposition 5.1.1. Let $\boldsymbol{x}^*$ be a feasible solution to (5.5) under the conditions of Theorem 5.1.1. Then $\boldsymbol{x}^*$ is feasible to the robust constraint

$$g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \leq \gamma + \beta(\boldsymbol{A}), \qquad \forall\ \boldsymbol{A} \in \operatorname{conv}(\mathcal{A}_N), \tag{5.7}$$

where

$$\beta(\boldsymbol{A}) \;=\; \inf_{\boldsymbol{q}\in\Delta^N} \left\{ \alpha(\boldsymbol{q}) \;\middle|\; \sum_{i=1}^{N} q_i \boldsymbol{A}_i = \boldsymbol{A} \right\}.$$

Proof. Consider any $\boldsymbol{A} \in \operatorname{conv}(\mathcal{A}_N)$, and let $\mathcal{Q}(\boldsymbol{A}) = \left\{ \boldsymbol{q} \in \Delta^N \;\middle|\; \sum_{i=1}^{N} q_i \boldsymbol{A}_i = \boldsymbol{A} \right\}$. We then have, for any $\boldsymbol{q} \in \mathcal{Q}(\boldsymbol{A})$,

$$\begin{aligned} g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) &= g\left(\left(\sum_{i=1}^{N} q_i \boldsymbol{A}_i \right) \boldsymbol{x}^* - \boldsymbol{b} \right) \\ &\leq \sum_{i=1}^{N} q_i g(\boldsymbol{A}_i \boldsymbol{x}^* - \boldsymbol{b}) \\ &\leq \gamma + \alpha(\boldsymbol{q}), \end{aligned}$$

where the first inequality follows by convexity of g, and the second by the fact that $\boldsymbol{x}^*$ is feasible to (5.5) and Theorem 3.1.1, since $\mu \in \mathcal{M}$. Taking the minimum over all $\boldsymbol{q} \in \mathcal{Q}(\boldsymbol{A})$ gives us the desired result. $\square$

Proposition 5.1.1, then, shows that solutions to (5.5), are feasible to a robust optimization problem, although the converse does not necessarily hold. Thus, one way to view the risk-constrained problem is as an *inner approximation* to robust problems of the form (5.7). Quantifying the tightness of this inner approximation is an open question of interest.

5.2 Probability guarantees

In this section, we consider probability guarantees on optimal solutions to problems with convex risk constraints. Recall that all of the development up to this point has centered on a data-driven approach. Our goal here is to provide a probability guarantee on the feasibility of optimal solutions, given that we used N independent observations of the uncertain matrix $\tilde{\boldsymbol{A}}$ to compute our solution. Obviously, then, a central question will be the *rate of convergence* of optimal solutions in our scenario-based methodology as we increase N.

This issue, in the context of risk measures, is a nontrivial one, and the results we present serve only as a starting point. Specifically, we present some convergence results for fixed solutions. Extending these results to solutions to solutions chosen as a function of the realizations, which is clearly what we would ultimately want, is a challenging problem. Again, we emphasize that the results here are incomplete and only a starting point.

To our knowledge, the only paper which considers this issue explicitly in the context of optimization of risk measures is the work of Takeda and Kanamori [109], who derive convergence bounds of optimal solutions to CVaR problems using learning theory (i.e., VC dimension). It is quite likely that such bounds are very loose, as this theory, while powerful, is also quite conservative. The results here could be extended to arbitrary, convex risk measures, but it is likely that the VC dimension for such classes would be so large that they would be useless from a practical perspective.

Convergence results in the context of chance constraints on convex functions are considered in a number of studies. de Farias and Van Roy [48] prove probability of containment of the approximate feasible set within the actual feasible set for the case of linear constraints. Calafiore and Campi [41] prove similar bounds for optimal solutions, but for the case of a general, convex function. Erdogan and Iyengar [54] extend these results when the underlying distribution is not exactly known, i.e., the chance constraints are *ambiguous.* Nemirovski and Shapiro [89] derive bounds based on moment generating information for the underlying distribution.

Notice that development thus far has relied on the assumption of finite support $\mathcal{A}_N$ (Assumption 5.1.2) for the uncertain matrix $\tilde{\boldsymbol{A}}$. Now we are taking a different viewpoint: let us assume that $\tilde{\boldsymbol{A}}$ has an underlying distribution, and that $\mathcal{A}_N$ represents N independent samples from this distribution.

We would like to quantify explicitly how well a sample-based optimization approximates the actual risk constraint for some N. In practice, we usually know very little about the underlying distribution, so this motivates us to make as few distributional assumptions as possible. We now define the problem formally, with the definition of randomness explicitly stated. Note that all of our results will focus on the convex OCE risk measure μ_ϕ as discussed in Chapter 4.

Definition 5.2.1. The risk-averse problem of interest is

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \mu_\phi(g(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})) \leq \gamma \\ & \boldsymbol{x} \in X, \end{aligned} \tag{5.8}$$

and where the matrix $\boldsymbol{A}$ is a random variable with $\operatorname{supp}(\boldsymbol{A}) = \mathcal{A} \subseteq \mathbb{R}^{m \times n}$, and $\mu_\phi \in \mathcal{M}$ is a convex risk measure with $\phi \in \Phi$.

As we will solve this problem approximately with N independent trials for $\boldsymbol{A}$, we can define the following. This is the problem we have been solving in the previous two chapters.

Definition 5.2.2. Let $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N$ be N independent samples of the random variable $\boldsymbol{A}$, and let μ be the convex risk measure $\mu_\phi \in \mathcal{M}$, where $\phi \in \Phi$. The **N-sampled**

counterpart of (5.8) is the problem

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{c}'\boldsymbol{x} \\ \text{subject to} \quad & \hat{\mu}_{\phi,\boldsymbol{x}}(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N) \leq \gamma \\ & \boldsymbol{x} \in X, \end{aligned} \tag{5.9}$$

where

$$\hat{\mu}_{\phi,\boldsymbol{x}}(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N) = \inf_{\nu} \left\{ \nu + \frac{1}{N} \sum_{i=1}^{N} \phi^* \left(g(\boldsymbol{A}_i \boldsymbol{x} - \boldsymbol{b}) - \nu \right) \right\}. \tag{5.10}$$

The only distributional assumption we will make is bounded support for the uncertain quantity of interest.

Assumption 5.2.1. There exists a finite $U \geq 0$ such that the random variable $g(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})$ satisfies $\operatorname{supp}(g(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})) \subseteq [0, U]$ for all $\boldsymbol{x} \in X$.

Note that the function g was taken to be nonnegative, so the lower bound in the assumption must hold by construction.[1] In many applications, the upper bound U is also reasonable. For instance, in the portfolio optimization example in Chapter 4, we had $X \subseteq \Delta^n$, and the uncertain return vector $\boldsymbol{r}$ played the role of $\boldsymbol{A}$. Assuming we believe there is a limit to the size of asset returns, we can readily derive an upper bound U as above. Of course, the quality of the resulting probability guarantees will depend intimately on the quality of our bound for U, as we will see shortly.

The probability of infeasibility of a solution involves a tradeoff between "bias" from the choice of risk measure and learning error due to the fact that we are approximating the original problem. Recall from Chapter 4 the sets

$$\begin{aligned} \mathcal{Q}_\alpha &= \{Q \ll \mathbb{P} \mid Q \leq \mathbb{P}/\alpha\} \\ \mathcal{Q}_\phi(\epsilon) &= \left\{ Q \ll \mathbb{P} \;\middle|\; \mathbb{E}_{\mathbb{P}} \left[\phi \left(\frac{dQ}{d\mathbb{P}} \right) \right] \leq \epsilon \right\}, \end{aligned}$$

[1] In the case of single constraints in linear optimization, where there is no function g, we can readily modify Assumption 5.2.1 to be the bounded, but not necessarily nonnegative, interval $[a, b]$, and all of the results still hold with appropriate modifications.

and note that we also have

$$\begin{aligned}\mathrm{CVaR}_\alpha(X) &= \sup_{Q\in\mathcal{Q}_\alpha} \mathbb{E}_Q[X] \\ \mu_\phi(X) &= \inf_\nu \{\nu + \mathbb{E}_{\mathbb{P}}[\phi^*(X-\nu)]\},\end{aligned}$$

for any $\phi \in \Phi$. The tradeoff in these two error terms is straightforward to show, as we now illustrate.

Proposition 5.2.1. Let $\boldsymbol{x}^* = \boldsymbol{x}^*(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N)$ be a feasible solution to (5.9). Then, for any $\epsilon \geq 0$, $\boldsymbol{x}^*$ satisfies

$$\mathbb{P}\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon\} \leq \inf_{\lambda\in[0,1]} \{\hat{\alpha}((1-\lambda)\epsilon) + \mathbb{P}\{\mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) > \gamma + \lambda\epsilon\}\},$$

where $\hat{\alpha}(\epsilon) = \inf\{\alpha \in [0,1] \mid \mathcal{Q}_\alpha \subseteq \mathcal{Q}_\phi(\epsilon)\}$.

Proof. We have, for any $\epsilon \geq$, for all $\lambda \geq 0$, the relation

$$\begin{aligned}&\mathbb{P}\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon\} \\ =\ &\mathbb{P}\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) \leq \gamma + \lambda\epsilon\} \underbrace{\mathbb{P}\{\mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) \leq \gamma + \lambda\epsilon\}}_{\leq 1} \\ &+ \underbrace{\mathbb{P}\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) > \gamma + \lambda\epsilon\}}_{\leq 1} \mathbb{P}\{\mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) > \gamma + \lambda\epsilon\} \\ \leq\ &\mathbb{P}\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) \leq \gamma + \lambda\epsilon\} \\ &+ \mathbb{P}\{\mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) > \gamma + \lambda\epsilon\}.\end{aligned}$$

For the first term, we proceed similarly to the proof of the bound for μ_ϕ in Theorem 4.3.1. Note that, for any $\alpha \geq \hat{\alpha}((1-\lambda)\epsilon)$, we have, by definition, $\mathcal{Q}_\alpha \subseteq \mathcal{Q}_\phi((1-\lambda)\epsilon)$. Similar to the proof of Theorem 4.3.1, we can compare penalty functions from the representation theorem for convex risk measures to see that, for such an α, we have the relation

$$\mathrm{CVaR}_\alpha(X) - (1-\lambda)\epsilon \geq \mu_\phi(X).$$

This in turn implies

$$\begin{aligned}
&\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) \leq \gamma + \lambda\epsilon\right\} \\
\leq\; &\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mathrm{CVaR}_\alpha\left(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})\right) \leq \gamma + \lambda\epsilon + (1-\lambda)\epsilon\right\} \\
=\; &\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mathrm{CVaR}_\alpha\left(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})\right) \leq \gamma + \epsilon\right\} \\
\leq\; &\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon \mid \mathrm{CVaR}_{\hat{\alpha}((1-\lambda)\epsilon)}\left(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})\right) \leq \gamma + \epsilon\right\} \\
\leq\; &\hat{\alpha}((1-\lambda)\epsilon),
\end{aligned}$$

where the second to last inequality follows by the definition that $\hat{\alpha}((1-\lambda)\epsilon)$ is the smallest such α such that $\mathcal{Q}_\alpha \subseteq \mathcal{Q}_\phi((1-\lambda)\epsilon)$, and the last inequality follows from the fact that

$$\mathrm{CVaR}_\alpha\left(X\right) \leq x \Rightarrow \mathbb{P}\left\{X \geq x\right\} \leq \alpha.$$

Noting that this is true for any $\lambda \in [0,1]$, we can take the minimum over all λ to obtain the result. □

Of course, to evaluate the bound in Proposition 5.2.1, we need both the function (or an upper bound on it) $\hat{\alpha}$, as well as a bound on the error term

$$\mathbb{P}\left\{\mu_\phi(g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b})) > \gamma + \lambda\epsilon\right\}.$$

For $\hat{\alpha}$, the case of discrete random variables is handled by Theorem 4.3.1. Similar results should hold over continuous spaces of measures.

A more challenging issue is bounding the learning error term. The primary difficulty is that the solution $\boldsymbol{x}^*$ is itself a random variable, which means, among other things, that the samples $\boldsymbol{A}_i\boldsymbol{x}^*$ are not independent. Additionally, even if we can describe the distribution for $\boldsymbol{A}_i\boldsymbol{x}^*$, we are ultimately interested in the distribution of $\boldsymbol{A}\boldsymbol{x}^*$, where $\boldsymbol{A}$ is a *new* realization of the random matrix $\boldsymbol{A}$.

5.2.1 Bounds for fixed solutions

Obviously, one way around this issue is to assume the solution $\boldsymbol{x}^*$ is chosen deterministically. In this case, bounding the learning error term simply becomes an estimation issue, and we can leverage results from the theory of concentration inequalities to obtain some bounds. Of course, in practice, all of our framework depends on the fact that we choose $\boldsymbol{x}^*$ as a function of the realizations $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N$; the following results are meant only to provide some qualitative insights into the rate of convergence of risk functions.

The main machinery for proving a rate of convergence for fixed solutions and convex OCE measures will be the following, powerful result from McDiarmid [86].

Theorem 5.2.1. *(McDiarmid [86]) Consider a function $f : S^n \to \mathbb{R}$ which satisfies*

$$\sup_{x_1,\ldots,x_n,x_i' \in S} |f(x_1,\ldots,x_n) - f(x_1,\ldots,x_i',\ldots,x_n)| \;\leq\; c_i, \tag{5.11}$$

for all $i = 1,\ldots,n$. Let $X_1,\ldots,X_n$ be independent random variables taking values in S. Then

$$\mathbb{P}\left\{\mathbb{E}\left[f(X_1,\ldots,X_n) - f(X_1,\ldots,X_n)\right] \geq \epsilon\right\} \;\leq\; e^{-2\epsilon^2/\sum\limits_{i=1}^{n} c_i^2}. \tag{5.12}$$

We now apply McDiarmid's inequality to derive a rate of convergence on (5.9) for fixed solutions $\boldsymbol{x}^*$.

Theorem 5.2.2. *When $\boldsymbol{x}^*$ is a feasible solution to (5.9) chosen independently of $(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)$, under Assumption 5.2.1, we have*

$$\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon\right\} \;\leq\; \inf_{\lambda \in [0,1]} \left\{\hat{\alpha}((1-\lambda)\epsilon) + e^{-2N(\lambda\epsilon/\phi^*(U))^2}\right\}.$$

Proof. The main task is to find the bounded differences, then apply McDiarmid's inequality and Proposition 5.2.1. Without loss of generality, we may assume the left term (using $\boldsymbol{A}_1$, not $\boldsymbol{A}_1'$, is the larger of the two). Let ν_R be a value of ν which achieves the infimum for the righthand term below (such a value exists and attains the infimum, since, by assumption, $g(\boldsymbol{Ax}-\boldsymbol{b}) \in [0,U]$). We have

$$\begin{aligned}
& \sup_{(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N,\boldsymbol{A}_1')\in\mathcal{A}^{N+1}} \hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N) - \hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1',\ldots,\boldsymbol{A}_N) \\
= & \sup_{(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N,\boldsymbol{A}_1')\in\mathcal{A}^{N+1}} \inf_{\nu}\left\{\nu + \frac{1}{N}\sum_{i=1}^{N}\phi^*\left(g(\boldsymbol{A}_i\boldsymbol{x}^*-\boldsymbol{b})-\nu\right)\right\} \\
& -\inf_{\nu}\left\{\nu + \frac{1}{N}\phi^*\left(g(\boldsymbol{A}_1'\boldsymbol{x}^*-\boldsymbol{b})-\nu\right)\frac{1}{N}\sum_{i=2}^{N}\phi^*\left(g(\boldsymbol{A}_i\boldsymbol{x}^*-\boldsymbol{b})-\nu\right)\right\} \\
\leq & \sup_{(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N,\boldsymbol{A}_1')\in\mathcal{A}^{N+1}} \nu_R + \frac{1}{N}\sum_{i=1}^{N}\phi^*\left(g(\boldsymbol{A}_i\boldsymbol{x}^*-\boldsymbol{b})-\nu_R\right) \\
& -\nu_R - \frac{1}{N}\phi^*\left(g(\boldsymbol{A}_1'\boldsymbol{x}^*-\boldsymbol{b})-\nu_R\right) - \frac{1}{N}\sum_{i=2}^{N}\phi^*\left(g(\boldsymbol{A}_i\boldsymbol{x}^*-\boldsymbol{b})-\nu_R\right) \\
= & \frac{1}{N}\sup_{(\boldsymbol{A}_1,\boldsymbol{A}_1')\in\mathcal{A}^2} \phi^*(g(\boldsymbol{A}_1\boldsymbol{x}^*-\boldsymbol{b})-\nu_R) - \phi^*(g(\boldsymbol{A}_1'\boldsymbol{x}^*-\boldsymbol{b})-\nu_R) \\
\leq & \frac{1}{N}\phi^*(U).
\end{aligned}$$

Now, note that the function $\hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)$ is the infimum over a family of convex functions in $(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)$, and, therefore, we have, by Jensen's inequality,

$$\begin{aligned}
\mathbb{E}\left[\hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)\right] &= \mathbb{E}\left[\inf_{\nu}\left\{\nu + \frac{1}{N}\sum_{i=1}^{N}\phi^*(g(\boldsymbol{A}_i\boldsymbol{x}-\boldsymbol{b})-\nu)\right\}\right] \\
&\geq \inf_{\nu}\left\{\nu + \mathbb{E}\left[\phi^*(g(\boldsymbol{A}_i\boldsymbol{x}-\boldsymbol{b})-\nu)\right]\right\} \\
&= \mu_\phi(g(\boldsymbol{Ax}-\boldsymbol{b})).
\end{aligned}$$

We then have

$$\begin{aligned}
\mathbb{P}\left\{\mu_\phi(g(\boldsymbol{Ax}^*-\boldsymbol{b})) > \gamma+\epsilon\right\} &\leq \mathbb{P}\left\{\mu_\phi(g(\boldsymbol{Ax}^*-\boldsymbol{b})) \geq \gamma+\epsilon\right\} \\
&\leq \mathbb{P}\left\{\mu_\phi(g(\boldsymbol{Ax}^*-\boldsymbol{b})) \geq \hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)+\epsilon\right\} \\
&\leq \mathbb{P}\left\{\mathbb{E}\left[\hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)\right] \geq \hat{\mu}_{\phi,\boldsymbol{x}^*}(\boldsymbol{A}_1,\ldots,\boldsymbol{A}_N)+\epsilon\right\} \\
&\leq e^{-2N(\epsilon/\phi^*(U))^2}.
\end{aligned}$$

Now putting this bound into Proposition 5.2.1, we obtain the stated result. □

5.2.2 Bounds for fixed solutions and CVaR

When the convex risk measure μ_ϕ corresponds to CVaR_α, it is well-known (e.g., Ben-Tal and Teboulle, [24]), that the corresponding "utility" function is $\phi^*(t) = 1/\alpha \max(0,t)$. Directly applying Theorem 5.2.2, we find, in this case, with a fixed solution $\boldsymbol{x}^*$,

$$\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \geq \gamma + \epsilon\right\} \quad \leq \quad \inf_{\lambda \in [0,1]} \left\{\hat{\alpha}((1-\lambda)\epsilon) + e^{-2N\alpha^2(\lambda\epsilon/U)^2}\right\}. \qquad (5.13)$$

Intuitively, this bound seems quite loose, due to the following result, which is actually a special case of McDiarmid's inequality, but discovered independently much earlier by Hoeffding [70].

Theorem 5.2.3. *(Hoeffding, [70]) Let $X_1, \ldots, X_N$ be i.i.d. random variables with* $\text{supp}\,(X_1) \subseteq [0, U]$*, where $U \geq 0$. Then, for any $\epsilon \geq 0$, we have*

$$\mathbb{P}\left\{\frac{1}{N}\sum_{i=1}^{N} X_i \leq \mathbb{E}\left[X\right] - \epsilon\right\} \leq e^{-2(\epsilon/U)^2 N}.$$

Hoeffding's inequality says we need $N_H = \mathcal{O}((U/\epsilon)^2 \log(1/\delta))$ samples to estimate the sample mean within a precision of ϵ with probability at least $1 - \delta$. We would expect for CVaR_α to need $\mathcal{O}(1/\alpha \cdot N_H)$ samples, as CVaR is essentially a conditional expectation of the α-tail of the distribution, and $\mathcal{O}(1/\alpha)$ samples fall in the α-tail. The bound in (5.13), however, suggests $\mathcal{O}(1/\alpha^2 \cdot N_H)$ samples are needed, and this seems too conservative. It turns out, by exploiting the structure of CVaR in more detail, we can reduce this to $\mathcal{O}(1/\alpha \cdot N_H)$ to match intuition.

To show this, we first need the following, straightforward fact.

Proposition 5.2.2. Let $X_1, \ldots, X_N$ be i.i.d. random variables with $\text{supp}\,(X_i) \subseteq \mathbb{R}_+$, and let $\hat{C}_\alpha(X_1, \ldots, X_N)$ be the sample estimate of $\text{CVaR}_\alpha\,(X_i)$, i.e.,

$$\hat{C}_\alpha(X_1, \ldots, X_N) \quad = \quad \inf_{\nu} \left\{\nu + \frac{1}{N\alpha}\sum_{i=1}^{N}(X_i - \nu)^+\right\}.$$

Then we have

$$\hat{C}_\alpha(X_1,\ldots,X_N) \;\geq\; \frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)}, \tag{5.14}$$

where $X_{(i)}$ are the decreasing order statistics of X_i, i.e., $X_{(1)} \geq X_{(2)} \geq \ldots \geq X_{(N)}$.

Proof. The proof follows by simply carrying out the minimization of the piecewise linear, convex function $\nu + \frac{1}{N\alpha}\sum_{i=1}^{N}(X_i - \nu)^+$, which has $N+1$ pieces. Quick inspection shows that the slope of this function changes sign from positive to negative at $\nu^* = X_{(\lceil N\alpha \rceil)}$, which means $\nu = \nu^*$. We then have

$$\begin{aligned}
\hat{C}_\alpha(X_1,\ldots,X_N) &= \nu^* + \frac{1}{N\alpha}\sum_{i=1}^{N}(X_i - \nu^*)^+ \\
&= X_{(\lceil N\alpha \rceil)} + \frac{1}{N\alpha}\sum_{i=1}^{\lceil N\alpha \rceil}\left(X_{(i)} - X_{(\lceil N\alpha \rceil)}\right) \\
&= X_{(\lceil N\alpha \rceil)} + \frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha \rfloor}\left(X_{(i)} - X_{(\lceil N\alpha \rceil)}\right) \\
&= \left(1 - \frac{\lfloor N\alpha \rfloor}{N\alpha}\right) X_{(\lceil N\alpha \rceil)} + \frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)} \\
&\geq \frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)},
\end{aligned}$$

where the inequality follows from the fact that $X_i \geq 0$. □

We can now generalize the Hoeffding Inequality to apply to the sample estimates of $\mathrm{CVaR}_\alpha(X)$ for random variables with bounded support. Our result applies to underlying random variables with a continuous distribution function.

Lemma 5.2.1. *Let $X_1,\ldots,X_N$ be i.i.d. random variables with a continuous distribution function and* $\mathrm{supp}(X_1) \subseteq [0,U]$, *where $U \geq 0$. When $N \geq 1/\alpha$, for any $\epsilon \geq 0$, we have*

$$\mathbb{P}\left\{\hat{C}_\alpha(X_1,\ldots,X_N) \leq \mathrm{CVaR}_\alpha(X) - \epsilon\right\} \leq 3e^{-c\alpha(\epsilon/U)^2 N},$$

where $c = \hat{\beta}^2/2 \approx 0.2123$, *and* $\hat{\beta}$ *is the unique root in* $[0,1]$ *to the cubic equation* $4\beta^3 - 16\beta^2 + 21\beta - 8 = 0$.

Proof. We have, by Proposition 5.2.2,

$$\mathbb{P}\left\{\hat{C}_\alpha(X_1,\ldots,X_N) \leq \mathrm{CVaR}_\alpha(X) - \epsilon\right\} \leq \mathbb{P}\left\{\frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha\rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon\right\}.$$

It is well-known (e.g., Acerbi and Tasche, [1]), that when the distribution function for X has no discontinuities, that $\mathrm{CVaR}_\alpha(X) = \mathbb{E}[X \mid X \geq \mathrm{VaR}_\alpha(X)]$. Thus, from here, we are basically trying to bound the error in estimating conditional expectation. The key behind to doing this is to condition on the random variable $K_{N,\alpha}$, defined as

$$K_{N,\alpha} = \max\{i \mid X_{(i)} \in [\mathrm{VaR}_\alpha(X)), U]\}.$$

(Note that $K_{N,\alpha}$ is a function of $(X_1,\ldots,X_N)$, but we omit this dependence for brevity). Clearly $K_{N,\alpha}$ is binomially distributed with parameters N and α. From here, if we condition on $K_{N,\alpha} = k$, one can see that $1/k\sum_{i=1}^k X_{(i)}$ is equal in distribution to $1/k\sum_{i=1}^k \tilde{X}_i$, where $\tilde{X}_i$ are i.i.d. and $\tilde{X}_1$ is equal in distribution to $\{X_1 \mid X_1 \in [\mathrm{VaR}_\alpha(X), U]\}$. We then have

$$\begin{aligned}
&\mathbb{P}\left\{\tfrac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha\rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon\right\} \\
=\ &\sum_{k=0}^{N} \mathbb{P}\{K_{N,\alpha} = k\}\,\mathbb{P}\left\{\tfrac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha\rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \,\middle|\, K_{N,\alpha} = k\right\} \\
=\ &\underbrace{\sum_{k=0}^{\lfloor N\alpha\rfloor} \mathbb{P}\{K_{N,\alpha} = k\}\,\mathbb{P}\left\{\frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha\rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \,\middle|\, K_{N,\alpha} = k\right\}}_{I_2} \\
&\underbrace{+ \sum_{k=\lceil N\alpha\rceil}^{N} \mathbb{P}\{K_{N,\alpha} = k\}\,\mathbb{P}\left\{\frac{1}{N\alpha}\sum_{i=1}^{\lfloor N\alpha\rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \,\middle|\, K_{N,\alpha} = k\right\}}_{I_1}.
\end{aligned}$$

We now distinguish two cases.

Case 1: $k \geq \lceil N\alpha \rceil$. Then

$$\begin{aligned}
& \mathbb{P}\left\{ \frac{1}{N\alpha} \sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \;\middle|\; K_{N,\alpha} = k \right\} \\
\leq \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \;\middle|\; K_{N,\alpha} = k \right\} && (k \geq \lceil N\alpha \rceil \geq N\alpha) \\
= \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} \tilde{X}_i \leq \mathrm{CVaR}_\alpha(X) - \epsilon \;\middle|\; K_{N,\alpha} = k \right\} \\
\leq \;& e^{-2(\epsilon/U)^2 k}. && (\text{Thm. } 5.2.3)
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
I_1 &\leq \sum_{k=\lceil N\alpha \rceil}^{N} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} e^{-2(\epsilon/U)^2 k} \\
&\leq e^{-2(\epsilon/U)^2 N\alpha} \sum_{k=\lceil N\alpha \rceil}^{N} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} \\
&\leq e^{-2(\epsilon/U)^2 \lceil N\alpha \rceil} \\
&\leq e^{-2\alpha(\epsilon/U)^2 N}.
\end{aligned}$$

Case 2: $k \leq \lfloor N\alpha \rfloor$. Then

$$\begin{aligned}
& \mathbb{P}\left\{ \frac{1}{N\alpha} \sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \;\middle|\; K_{N,\alpha} = k \right\} \\
\leq \;& \mathbb{P}\left\{ \frac{1}{N\alpha} \sum_{i=1}^{k} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \;\middle|\; K_{N,\alpha} = k \right\} && (X_i \geq 0) \\
= \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \frac{N\alpha}{k} \left(\mathrm{CVaR}_\alpha(X) - \epsilon\right) \;\middle|\; K_{N,\alpha} = k \right\} \\
\leq \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \frac{\lfloor N\alpha \rfloor}{k} \left(\mathrm{CVaR}_\alpha(X) - \epsilon\right) \;\middle|\; K_{N,\alpha} = k \right\} \\
\leq \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \frac{\lfloor N\alpha \rfloor}{k}\epsilon + \left(\frac{\lfloor N\alpha \rfloor - k}{k}\right) U \;\middle|\; K_{N,\alpha} = k \right\} && (\mathrm{CVaR}_\alpha(X_i) \leq U) \\
\leq \;& \mathbb{P}\left\{ \frac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon'(k) \;\middle|\; K_{N,\alpha} = k \right\},
\end{aligned}$$

where $\epsilon'(k) = \frac{\lfloor N\alpha \rfloor}{k}(\epsilon - U) + U$. Note that $\epsilon'(k) \geq 0$ if and only if $k \geq (1-\epsilon/U)\lfloor N\alpha \rfloor$. Let $\gamma = 1 - \epsilon/U$ and let $k^* = \lceil \gamma \lfloor N\alpha \rfloor \rceil$. Finally, for some $\beta \in [0,1]$, let $k^*_\beta = \lceil (\beta\gamma + (1-\beta)) \lfloor N\alpha \rfloor \rceil = \lceil (1 - \beta(\epsilon/U)) \lfloor N\alpha \rfloor \rceil \in [k^*, \lfloor N\alpha \rfloor]$. We furthermore note by a simple concavity and Taylor series argument that, for any $k \in [k^*, \lfloor N\alpha \rfloor]$, we have

$$
\begin{aligned}
(\epsilon'(k)/U)^2 &= \left(1 - \frac{\gamma \lfloor N\alpha \rfloor}{k}\right)^2 \\
&\geq \left(\frac{k}{\lfloor N\alpha \rfloor} - \gamma\right)^2 .
\end{aligned} \tag{5.15}
$$

Continuing, we have

$$
\begin{aligned}
&\mathbb{P}\left\{ \tfrac{1}{k} \sum_{i=1}^{k} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon'(k) \,\middle|\, K_{N,\alpha} = k \right\} \\
\leq\ & \begin{cases} e^{-2(\epsilon'(k)/U)^2 k}, & k^* \leq k \leq \lfloor N\alpha \rfloor, \\ 1, & k < k^*, \end{cases} \qquad \text{(Thm. 5.2.3)} \\
\leq\ & \begin{cases} e^{-2k(\frac{k}{\lfloor N\alpha \rfloor} - \gamma)^2} & k^* \leq k \leq \lfloor N\alpha \rfloor, \\ 1, & k < k^*, \end{cases} \qquad \text{(Eq. 5.15)} \\
\leq\ & \begin{cases} e^{-2k(\frac{k}{\lfloor N\alpha \rfloor} - \gamma)^2} & k^*_\beta \leq k \leq \lfloor N\alpha \rfloor, \\ 1, & k < k^*_\beta, \end{cases} \\
\leq\ & \begin{cases} e^{-2k(\frac{k^*_\beta}{\lfloor N\alpha \rfloor} - \gamma)^2} & k^*_\beta \leq k \leq \lfloor N\alpha \rfloor, \\ 1, & k < k^*_\beta, \end{cases} \\
\leq\ & \begin{cases} e^{-2k((1-\beta)\epsilon/U)^2} & k^*_\beta \leq k \leq \lfloor N\alpha \rfloor, \\ 1, & k < k^*_\beta. \end{cases}
\end{aligned}
$$

Going back to our original expansion, we have

$$
\begin{aligned}
I_2 &= \sum_{k=0}^{\lfloor N\alpha \rfloor} \mathbb{P}\{K_{N,\alpha} = k\} \, \mathbb{P}\left\{ \tfrac{1}{N\alpha} \sum_{i=1}^{\lfloor N\alpha \rfloor} X_{(i)} \leq \mathrm{CVaR}_\alpha(X) - \epsilon \,\middle|\, K_{N,\alpha} = k \right\} \\
&\leq \underbrace{\left[\sum_{k=0}^{k^*_\beta - 1} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} \right]}_{I_{2a}} + \underbrace{\left[\sum_{k=k^*_\beta}^{\lfloor N\alpha \rfloor} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} e^{-2k((1-\beta)\epsilon/U)^2} \right]}_{I_{2b}} .
\end{aligned}
$$

For the first term, we have

$$\begin{aligned}
I_{2a} &= \mathbb{P}\left\{K_{N,\alpha} \le k_\beta^* - 1\right\} \\
&\le \mathbb{P}\left\{K_{N,\alpha} \le (1-\beta(\epsilon/U)) \lfloor N\alpha \rfloor\right\} \\
&\le \mathbb{P}\left\{K_{N,\alpha} \le (1-\beta(\epsilon/U)) (N\alpha)\right\} \\
&\le e^{-(\alpha/2)(\beta(\epsilon/U))^2 N}. && \text{(Chernoff bound)}
\end{aligned}$$

For the second term, we have

$$\begin{aligned}
I_{2b} &= \sum_{k=k_\beta^*}^{\lfloor N\alpha \rfloor} \binom{N}{k} \alpha^k (1-\alpha)^{N-k} e^{-2k((1-\beta)\epsilon/U)^2} \\
&= \sum_{k=k_\beta^*}^{\lfloor N\alpha \rfloor} \binom{N}{k} \left(\alpha \cdot e^{-2((1-\beta)\epsilon/U)^2}\right)^k (1-\alpha)^{N-k} \\
&\le \left(1 - \left(1 - e^{-2((1-\beta)\epsilon/U)^2}\right)\alpha\right)^N && \text{(Binomial expansion)} \\
&\le e^{-\left(1-\exp\left(-2((1-\beta)\epsilon/U)^2\right)\alpha N\right)} && (1-\Delta \le e^{-\Delta} \text{ for } \Delta \ge 0) \\
&\le e^{-\alpha(2\beta(2-\beta)(1-\beta)^2)(\epsilon/U)^2 N},
\end{aligned}$$

where the last line follows from the fact that that, for any $\rho \in [0,1]$ and $x \in [0,1]$, we have

$$\begin{aligned}
1 - e^{-2\rho^2 x^2} &= 1 - (1 - 2\rho^2x^2 + 2\rho^4x^4 - \cdots) \\
&= 2\rho^2x^2 - 2\rho^4x^4 + \cdots \\
&\ge 2\rho^2x^2 - 2\rho^4x^4 \\
&\ge ax^2,
\end{aligned}$$

provided that $a \le 2\rho^2(1-\rho^2)$. Using $\rho = 1-\beta$, $x = \epsilon/U$, and setting such an a to its highest value gives the above bound. To combine the three terms into a single, exponential bound with the highest decay coefficient, we want to choose $\beta \in [0,1]$ such that the $f(\beta) = \min\left(\beta^2/2, 2\beta(2-\beta)(1-\beta)^2\right)$ is maximum. which occurs when these two terms in the minimization are equal, and hence the cubic equation in the result. Putting everything together, we have

$$
\begin{aligned}
& \mathbb{P}\left\{\hat{C}_\alpha(X_1,\ldots,X_N) \le \mathrm{CVaR}_\alpha(X) - \epsilon\right\} \\
\le\ & I_1 + I_{2a} + I_{2b} \\
\le\ & e^{-2\alpha(\epsilon/U)^2 N} + e^{-\alpha(\beta^2/2)(\epsilon/U)^2 N} + e^{-\alpha(2\beta(2-\beta)(1-\beta)^2)(\epsilon/U)^2 N} \\
\le\ & 3e^{-\alpha \min(2,\beta^2/2,2\beta(2-\beta)(1-\beta)^2)(\epsilon/U)^2 N} \\
\le\ & 3e^{-c\alpha(\epsilon/U)^2 N},
\end{aligned}
$$

where c is as stated above. □

Lemma 5.2.1 leads us directly to the main result for a probability guarantee with fixed solutions and CVaR.

Theorem 5.2.4. *Let $\mu_\phi = CVaR_\alpha$ for some $\alpha \in [1/N, 1]$. When $\boldsymbol{x}^*$ is a feasible solution to (5.9) chosen independently of $(\boldsymbol{A}_1, \ldots, \boldsymbol{A}_N)$, under Assumption 5.2.1, we have*

$$
\mathbb{P}\left\{g(\boldsymbol{A}\boldsymbol{x}^* - \boldsymbol{b}) \ge \gamma + \epsilon\right\} \le \alpha + 3e^{-c\alpha(\epsilon/U)^2 N},
$$

where c is as in Lemma 5.2.1.

Proof. The proof follows directly from Proposition 5.2.1 and Lemma 5.2.1. □

Chapter 6

A tractable approach to constrained multi-stage LQC

The previous chapters focused on static optimization problems in which there is a single decision to be made. In this chapter, we consider a dynamic problem with multiple decision stages. Although some attempts at extending the risk theoretic ideas we have discussed have been made (e.g., Artzner et al., [4]), the theory is still undeveloped from an algorithmic perspective. As a result, our starting point here will not be risk measures; instead, we will start directly with an uncertainty set-based model of uncertainty from robust optimization, and show how to tractably approach optimization under uncertainty of linear systems with quadratic costs. Risk in this setting will come to play in the form of probability guarantees on the performance of the resulting solutions.

The standard approach to multi-stage optimization is the theory of dynamic programming. This theory, while conceptually elegant, is computationally impractical for all but a few special cases of system dynamics and cost functions. One of the notable triumphs of dynamic programming is its success with stochastic linear systems and quadratic cost functions (stochastic linear-quadratic control, SLQC) . It is easily shown (e.g., Bertsekas [26]) in this case that the cost-to-go functions are quadratic in the state, and therefore the resulting optimal controls are linear in the current state. As a result, solving Bellman's equation in this case is tantamount to finding appropriate gain matrices, and these gain matrices are described by the well-known Riccati equation (Willems [115]).

This success, however, has some limitations. In particular, Bellman's equation in the SLQC cannot handle even the simplest of constraints on either the control or state vectors. It is not difficult to find applications that demand constraints on the controls or state. Bertsimas and Lo [29] describe the dynamics of an optimal share-purchasing policy for stockholders. The unconstrained policy based on the Riccati equation requires the investor to purchase and *sell* shares, which is clearly absurd. This can be mitigated by a nonnegativity constraint on the control, which causes the cost-to-go function to become piecewise quadratic with an exponential number of pieces. Thus, a very simple constraint destroys the tractability of this approach. Much of the current literature (e.g., Verriest and Panunen [110], Voulgaris [112]) derives necessary conditions for optimality for simple control constraints, but does not explicitly describe solution methods.

A further drawback of the Riccati approach from dynamic programming is that it only deals with the expected value of the resulting cost. In many cases, we may wish to know more information about the distribution of the cost function (e.g., cases in which we want to provide a probabilistic level of protection guaranteeing some system performance).

In this chapter, we take an entirely different approach to the SLQC problem. Rather than attempting to solve Bellman's equation, we exploit relatively new results from robust optimization to propose an alternative solution technique for SLQC. Our approach has the following advantages over the traditional dynamic programming approach:

1. It inherits much greater modelling freedom by being able to tractably handle a wide variety of constraints on both the control and state vectors.

2. It admits a probabilistic description of the resulting cost, allowing us to understand and, in some cases, *control* the system cost distribution. We can capture risk in this setting, then, via related probability guarantees on the distribution

of the cost.

3. In the unconstrained case, its complexity is not much more than the complexity of linear feedback (i.e., the Riccati approach). In particular, optimal policies in this case may be computed by optimizing a convex function over a scalar, then multiplying the initial state by appropriate matrices.

Our approach is based on convex conic optimization. Although the use of these optimization techniques is widespread in the control literature (see, e.g., Boyd et al. [37], El Ghaoui and Lebret [64], Rami and El Ghaoui [96], Rami and Zhou[97]), we believe this is a new approach. Chen and Zhou [44] provide an elegant solution to the SLQC problem with conic control constraints, but their solution is limited to a scalar-valued state variable and homogeneous system dynamics. Our approach here is more general. We emphasize that we are *not* proposing a solution for robust control (see e.g., Zhou et al. [119] for a start to the vast literature on the subject); *rather, we are proposing an approach to the SLQC with the conceptual framework of robust optimization as a guide.*

The structure of the chapter is as follows. In Section 6.1, we present a description of the SQLC problem, as well as the currently known results from dynamic programming and a conceptual description of our methodology. In addition, we provide background for the robust optimization results we will later use. In Section 6.2, we develop our approach for the unconstrained SLQC problem. This approach is based on semi-definite programming (SDP) and robust quadratic programming results from Ben-Tal and Nemirovski [17]. We further show that this SDP has a very special structure which allows us to derive a closed-loop control law suitable for real-time applications. Unfortunately, in the presence of constraints, this simplification no longer applies, and the complexity of solving the SDP is impractical in a closed-loop setting. This motivates us to simplify the SDP, which we do in Section 6.3. Here we use recent results from robust conic optimization developed by Bertsimas and Sim [33] to develop a tight SOCP approximation which is far easier to solve. We then show in Section 6.4 how this approach admits various constraints and performance guarantees. These constraints may be deterministic constraints on the control, or probabilistic guarantees on the state and objective function. In Section 6.5, we show that a particular model for imperfect state information fits into the framework already developed, and in Section 6.6, we provide computational results. Section 6.7 concludes the paper.

6.1 Problem statement and preliminaries

Throughout this paper we will work with discrete-time stochastic linear systems of the form:

$$\boldsymbol{x}_{k+1} \;=\; \boldsymbol{A}_k\boldsymbol{x}_k + \boldsymbol{B}_k\boldsymbol{u}_k + \boldsymbol{C}_k\boldsymbol{w}_k, \quad k = 0,\ldots,N-1, \tag{6.1}$$

where $\boldsymbol{x}_k \in \mathbb{R}^{n_x}$ is a state vector, $\boldsymbol{u}_k \in \mathbb{R}^{n_u}$ is a control vector, and $\boldsymbol{w}_k \in \mathbb{R}^{n_w}$ is a disturbance vector (an unknown quantity). We assume throughout that the matrices $\boldsymbol{A}_k \in \mathbb{R}^{n_x \times n_x}$, $\boldsymbol{B}_k \in \mathbb{R}^{n_x \times n_u}$, and $\boldsymbol{C}_k \in \mathbb{R}^{n_x \times n_w}$ are known exactly.

It is desired to control the system in question in a way that keeps the cost function

$$J(\boldsymbol{x}_0, \boldsymbol{u}, \boldsymbol{w}) = \sum_{k=1}^{N} \left(\boldsymbol{x}_k^T\boldsymbol{Q}_k\boldsymbol{x}_k + 2\boldsymbol{q}_k^T\boldsymbol{x}_k\right) + \sum_{k=0}^{N-1} \left(\boldsymbol{u}_k^T\boldsymbol{R}_k\boldsymbol{u}_k + 2\boldsymbol{r}_k^T\boldsymbol{u}_k\right), \tag{6.2}$$

as small as possible. Here we will assume $\boldsymbol{Q}_k \succeq \boldsymbol{0}$, $\boldsymbol{R}_k \succ \boldsymbol{0}$, and, again, that the data $\boldsymbol{Q}_k$, $\boldsymbol{q}_k$, $\boldsymbol{R}_k$, and $\boldsymbol{r}_k$, are known exactly. We are also using the shorthand $\boldsymbol{u}$ and $\boldsymbol{w}$ to denote the entire vector of controls and disturbances, i.e.,

$$\boldsymbol{u}^T \;=\; \begin{bmatrix}\boldsymbol{u}_0^T & \boldsymbol{u}_1^T & \cdots & \boldsymbol{u}_{N-1}^T\end{bmatrix}, \tag{6.3}$$

$$\boldsymbol{w}^T \;=\; \begin{bmatrix}\boldsymbol{w}_0^T & \boldsymbol{w}_1^T & \cdots & \boldsymbol{w}_{N-1}^T\end{bmatrix}. \tag{6.4}$$

Finally, our convention will be for the system to be in some initial state $\boldsymbol{x}_0$. Unless otherwise stated, we assume this initial state is also known exactly.

Note that (6.2) is an uncertain quantity, as it depends on the realization of $\boldsymbol{w}$, which is unknown. Most approaches assume $\boldsymbol{w}$ is a random variable possessing some distributional properties and proceed to minimize (6.2) in an expected value sense. We now survey the traditional approach to this problem.

6.1.1 The traditional approach: Bellman's recursion

The dynamic programming approach requires a few distributional assumptions on the disturbance vectors. Typically, it is assumed that the $\boldsymbol{w}_k$ are independent, and independent of both $\boldsymbol{x}_k$ and $\boldsymbol{u}_k$. Moreover, we have $\mathbb{E}[\boldsymbol{w}_k] = \boldsymbol{0}$, and $\boldsymbol{w}_k$ has finite second moment. For this derivation, we will assume $\boldsymbol{q}_k = \boldsymbol{0}$, $\boldsymbol{r}_k = \boldsymbol{0}$, and $\boldsymbol{C}_k = \boldsymbol{I}$ for ease of notation, but the result holds more generally after some simple manipulations. Modifications of some of the distributional assumptions (such as nonzero mean, correlations) are also possible, but we do not detail them here.

The literature on this subject is vast, and the problem is well-understood. The main result is that the expected cost-to-go functions, $J_k(\boldsymbol{x}_k)$, defined by

$$
\begin{aligned}
J_N(\boldsymbol{x}_N) &= \boldsymbol{x}_N^T \boldsymbol{Q}_N \boldsymbol{x}_N, \\
J_k(\boldsymbol{x}_k) &= \boldsymbol{x}_k^T \boldsymbol{Q}_k \boldsymbol{x}_k + \min_{\boldsymbol{u}_k} \mathbb{E}\left[\boldsymbol{u}_k^T \boldsymbol{R}_k \boldsymbol{u}_k + J_{k+1}(\boldsymbol{A}_k \boldsymbol{x}_k + \boldsymbol{B}_k \boldsymbol{u}_k + \boldsymbol{w}_k)\right],
\end{aligned} \tag{6.5}
$$

are quadratic in the state $\boldsymbol{x}_k$. Thus, it follows that the optimal policy is *linear* in the current state. In particular, one can show (see, e.g., Bertsekas [26]) that the optimal control $\boldsymbol{u}_k^\star$ is given by

$$
\boldsymbol{u}_k^\star = \boldsymbol{L}_k \boldsymbol{x}_k,
$$

where $\boldsymbol{L}_k = -(\boldsymbol{B}_k^T \boldsymbol{K}_{k+1} \boldsymbol{B}_k + \boldsymbol{R}_k)^{-1} \boldsymbol{B}_k^T \boldsymbol{K}_{k+1} \boldsymbol{A}_k$, and $\boldsymbol{K}_k$ are symmetric, positive semi-definite matrices given recursively by

$$
\begin{aligned}
\boldsymbol{K}_N &= \boldsymbol{Q}_N \\
\boldsymbol{K}_k &= \boldsymbol{A}_k^T \left(\boldsymbol{K}_{k+1} - \boldsymbol{K}_{k+1} \boldsymbol{B}_k \left(\boldsymbol{B}_k^T \boldsymbol{K}_{k+1} \boldsymbol{B}_k + \boldsymbol{R}_k\right)^{-1} \boldsymbol{B}_k^T \boldsymbol{K}_{k+1}\right) \boldsymbol{A}_k + \boldsymbol{Q}_k.
\end{aligned} \tag{6.6}
$$

The fact that the recursion given in (6.5) works so well (from a complexity standpoint) is quite particular to the case of linear systems and quadratic costs. For more arbitrary systems or cost functions such an approach is, in general, intractable.

A more troubling difficulty, however, is that even with the same system and cost function, this approach explodes computationally with ostensibly simple constraints, such as $\boldsymbol{u}_k \geq \mathbf{0}$. For instance, the cost-to-go function (6.5) in this case becomes piecewise quadratic with an exponential (in N) number of pieces (Bertsimas and Lo, [29]).

Thus, the traditional, dynamic programming approach can be solved very rapidly with linear feedback in the unconstrained case, but becomes, for all practical purposes, impossible when we add constraints to the problem. This is a very unfavorable property of the DP approach, and it is in direct contrast to the field of convex optimization, whose problem instances are quite robust (in terms of complexity) to perturbations in the constraint structure. Our approach, which we now detail, will leverage this useful property of convex optimization.

6.1.2 A tractable approach: overview

The traditional approach above is not amenable to problem changes such as the addition of constraints for two primary reasons:

1. *Complexity of distributional calculations.* Computing the expectation in (6.5), except for very special cases, is cumbersome computationally.

2. *Intractability of Bellman's recursion.* The recursion in (6.5) requires us, when computing the current control, to have advance knowledge of *all future controls for all possible future states*, even states that are extraordinarily improbable. While this recursion is an elegant idea conceptually, it is not well-suited to computation because the number of possible future states grows so rapidly with problem size.

We propose the following approach, which circumvents these difficulties:

(a) Given our current state $\boldsymbol{x}_0$ and problem data, we consider the entire control and disturbance vectors $\boldsymbol{u} \in \mathbb{R}^{N \cdot n_u}$, $\boldsymbol{w} \in \mathbb{R}^{N \cdot n_w}$, respectively, as in (6.3)-(6.4).

(b) We do not assume a particular distribution for $\boldsymbol{w}$. Assume only that $\boldsymbol{w}$ belongs within some "reasonable" uncertainty set. In particular, assume $\boldsymbol{w}$ belongs to some norm-bounded set

$$\mathcal{W}_\gamma \;=\; \{\boldsymbol{w} \mid \|\boldsymbol{w}\|_2 \leq \gamma\}, \tag{6.7}$$

parameterized by $\gamma \geq 0$.[1]

(c) Discard the notion of Bellman's recursion. Instead, do the best we can for all possible disturbances within $\mathcal{W}_\gamma$. That is, rather than computing controls for *every* possible state realization, we simply choose a control vector for the remaining stages which performs best for the most pessimistic disturbance within this "reasonable" uncertainty set. Specifically, we search for an optimal control $\boldsymbol{u}^*$ to the problem

$$\min_{\boldsymbol{u} \in \mathbb{R}^{N \cdot n_u}} \max_{\boldsymbol{w} \in \mathcal{W}_\gamma} J(\boldsymbol{x}_0, \boldsymbol{u}, \boldsymbol{w}). \tag{6.8}$$

Of course, this brings up the issue of open-loop versus closed-loop control. At first glance, this approach appears to be an open-loop method only. We can, however, compute a solution $\boldsymbol{u}^*$ to (6.8), take the first n_u components, and apply this as the current control. After a new state observation, we can repeat the calculation in (6.8) with the updated problem data (most of this updating can be done off-line). The only issue is that the routine for solving (6.8) be computationally simple enough for the application at hand. The complexity of these solution procedures will indeed be a central issue for much of the remaining discussion.

This approach, as will be shown, has the following properties:

[1]If we wish instead to have $\boldsymbol{w} \in \{\boldsymbol{w} \mid \boldsymbol{w}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{w} \leq \gamma^2\}$, where $\boldsymbol{\Sigma} \succ \boldsymbol{0}$, then we may re-scale coordinates and obtain a problem of the same form. Note that the statistical appropriateness of ellipsoids and their explicit construction is not the subject of this chapter, but the interested reader may see Paganini [92] for uncertainty set modelling for the case of white noise.

1. It is tractable, even in the presence of control and state space constraints.

2. It admits greater insight and control into the stochastic behavior of the cost function under appropriate distributional assumptions. The traditional approach, on the other hand, only minimizes the expected value.

3. In the unconstrained case, it yields an efficient control law which is linear in the current state after a simple, scalar optimization procedure. In addition, for $\gamma = 0$, we recover the traditional (Riccati) solution, whereas, for $\gamma > 0$, we have a family of increasingly conservative approaches.

Thus, our framework is a methodology for SLQC which pays essentially nothing in complexity in the nominal (unconstrained) case but is vastly more amenable to the addition of constraints and performance guarantees. We do this by discarding the cumbersome calculations over distributions as well as the total adaptability demanded by Bellman's recursion. Of course, given a known distribution for the $\boldsymbol{w}_k$ our approach will be a suboptimal solution to the problem of minimizing the expected cost-to-go. The quality of this approximate solution under normally distributed disturbances will be the subject of Sections 6.4 and 6.6.

In terms of modern control theory literature, our approach can be thought of as H^∞, model predictive control (MPC) for a discrete time, finite horizon problem. The H^∞ aspect is due to the sup over all disturbances within $\mathcal{W}_\gamma$ (for more on H^∞, particularly as it relates to the problem set-up here, see Başar and Bernhard [6]). The MPC portion of our approach is related to the fact that we are not enforcing Bellman's equation, but rather recomputing a solution which does not assume future optimality at each stage. A good starting point for the vast literature on MPC is Mayne et al. [85].

Although this Chapter uses these standard ideas from control theory, we believe a number of the results, including the solution structure shown in Section 6.2.1, the SOCP approximation in Section 6.3, and the probability guarantees in 6.4.3, are new.

To solve (6.8) we will utilize a number of results from robust optimization.

6.1.3 Results from robust quadratic optimization over ellipsoids

We will leverage some robust quadratic programming results popularized by Ben-Tal and Nemirovski [17]. In particular, they consider the conic quadratic constraint

$$\|\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}\|_2 \leq \boldsymbol{c}^T\boldsymbol{x} + d$$

when the data $(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}, d)$ are uncertain and known only to belong to some bounded uncertainty set $\mathcal{U}$. The goal of robust quadratic programming is to optimize over the set of all $\boldsymbol{x}$ such that the constraint holds for all possible values of the data within the set $\mathcal{U}$. In other words, we desire to find $\boldsymbol{x}$ such that

$$\|\boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}\|_2 \leq \boldsymbol{c}^T\boldsymbol{x} + d \qquad \forall\ (\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}, d) \in \mathcal{U}.$$

Ben-Tal and Nemirovski show that in the case of an ellipsoidal uncertainty set, the problem of optimizing over an uncertain conic quadratic inequality may be solved tractably using semi-definite programming (SDP). This turns out to also be the case for (6.8) above. To this end, we will need the following two classical results, proofs of which may be found in [17], among others. First, we have the *Schur complement lemma.*

Lemma 6.1.1. *Let*

$$\boldsymbol{A} \;=\; \begin{bmatrix} \boldsymbol{B} & \boldsymbol{C}^T \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix},$$

where $\boldsymbol{B} \succ \boldsymbol{0}$. *Then* $\boldsymbol{A}$ *is positive (semi-) definite if and only if the matrix* $\boldsymbol{D} - \boldsymbol{C}\boldsymbol{B}^{-1}\boldsymbol{C}^T$ *is positive (semi-) definite.*

In addition, we have the $\mathcal{S}$*-lemma.*

Lemma 6.1.2. *Let* $\boldsymbol{A}$, $\boldsymbol{B}$ *be symmetric* $n \times n$ *matrices, and assume that the quadratic* *inequality*

$$\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} \geq 0$$

is strictly feasible. Then the minimum value of the problem

$$\begin{array}{ll} \text{minimize} & \boldsymbol{x}^T \boldsymbol{B} \boldsymbol{x} \\ \text{subject to} & \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} \geq 0 \end{array}$$

is nonnegative if and only if there exists a $\lambda \geq 0$ *such that* $\boldsymbol{B} - \lambda \boldsymbol{A} \succeq \boldsymbol{0}$.

6.1.4 Results from robust conic optimization over norm-bounded sets

To improve the complexity of solving (6.8) when we have constraints, we will utilize recent results from robust conic optimization results due to Bertsimas and Sim [33]. Most of these results are presented in Chapter 2, but we repeat them here in more detail for convenience.

The approach is a relaxation of the exact min-max approach, but is computationally less complex and leads to a unified probability bound across a variety of conic optimization problems. We survey the main ideas and developments here.

Bertsimas and Sim use the following model for data uncertainty:

$$\tilde{\boldsymbol{D}} = \boldsymbol{D}^0 + \sum_{j \in N} \Delta \boldsymbol{D}^j \tilde{z}_j,$$

where $\boldsymbol{D}^0$ is the nominal data value and $\Delta \boldsymbol{D}^j$ are data perturbations. The $\tilde{z}_j$ are random variables with mean zero and independent, identical distributions. The goal is to find a policy $\boldsymbol{x}$ such that a given constraint is "robust feasible," i.e.,

$$\max_{\tilde{\boldsymbol{D}} \in \mathcal{U}_\Omega} f(\boldsymbol{x}, \tilde{\boldsymbol{D}}) \leq 0, \tag{6.9}$$

where

$$\mathcal{U}_\Omega = \left\{ \boldsymbol{D}^0 + \sum_{j \in N} \boldsymbol{\Delta D}^j u_j \;\middle|\; \|\boldsymbol{u}\| \leq \Omega \right\}. \tag{6.10}$$

For our purposes we typically use the Euclidean norm on $\boldsymbol{u}$, as it is self-dual, but many other choices for the norm may be tractably used [33]. We operate under some restrictions on the function $f(\boldsymbol{x}, \boldsymbol{D})$.[2]

Assumption 6.1.1. The function $f(\boldsymbol{x}, \boldsymbol{D})$ satisfies:

(a) $f(\boldsymbol{x}, \boldsymbol{D})$ is convex in $\boldsymbol{D}$ for all $\boldsymbol{x} \in \mathbb{R}^n$.

(b) $f(\boldsymbol{x}, k\boldsymbol{D}) = kf(\boldsymbol{x}, \boldsymbol{D})$ for all $k \geq 0$, $\boldsymbol{D}$, $\boldsymbol{x} \in \mathbb{R}^n$.

One of the central ideas of [33] is to linearize the model of robustness as follows:

$$\max_{(\boldsymbol{u},\boldsymbol{v}) \in \mathcal{V}_\Omega} f(\boldsymbol{x}, \boldsymbol{D}^0) + \sum_{j \in N} \left\{ f(\boldsymbol{x}, \boldsymbol{\Delta D^j}) u_j + f(\boldsymbol{x}, -\boldsymbol{\Delta D^j}) v_j \right\} \leq 0, \tag{6.11}$$

where

$$\mathcal{V}_\Omega = \left\{ (\boldsymbol{u}, \boldsymbol{v}) \in \mathbb{R}_+^{2|N|} \;\middle|\; \|\boldsymbol{u} + \boldsymbol{v}\| \leq \Omega \right\}.$$

In the framework developed thus far, Eq. (6.11) turns out to be a relaxation of Eq. (6.9), i.e., we have the following:

Proposition 6.1.1. *(Bertsimas-Sim)*

(a) If $f(\boldsymbol{x}, \boldsymbol{A} + \boldsymbol{B}) = f(\boldsymbol{x}, \boldsymbol{A}) + f(\boldsymbol{x}, \boldsymbol{B})$, then $\boldsymbol{x}$ satisfies (6.11) if and only if $\boldsymbol{x}$ satisfies (6.9).

(b) Under Assumption 6.1.1, if $\boldsymbol{x}$ is feasible in (6.11), then $\boldsymbol{x}$ is feasible in (6.9).

Finally, Eq. (6.11) is tractable due to the following.

[2] In [33], the authors assume the function is concave in the data. For our purposes, convexity is more convenient. All results follow up to sign changes and we report them accordingly.

Theorem 6.1.1. (Bertsimas-Sim) *Under Assumption 6.1.1, we have:*

(a) *Constraint (6.11) is equivalent to*

$$f(\boldsymbol{x}, \boldsymbol{D}^0) + \Omega \|\boldsymbol{s}\|^* \leq 0, \tag{6.12}$$

where

$$s_j = \max\{f(\boldsymbol{x}, \boldsymbol{\Delta D}^j), f(\boldsymbol{x}, -\boldsymbol{\Delta D}^j)\}.$$

(b) *Eq. (6.12) can be written as:*

$$\begin{aligned} \exists \quad & y, \boldsymbol{t} \in \mathbb{R}^{|N|} : \\ & f(\boldsymbol{x}, \boldsymbol{D}^0) \leq -\Omega y \\ & t_j \geq f(\boldsymbol{x}, \boldsymbol{\Delta D}^j) \qquad \forall j \in N \\ & t_j \geq f(\boldsymbol{x}, -\boldsymbol{\Delta D}^j) \qquad \forall j \in N \\ & \|\boldsymbol{t}\|^* \leq y. \end{aligned} \tag{6.13}$$

Finally, Bertsimas and Sim derive a probability of constraint violation.

Theorem 6.1.2. (Bertsimas-Sim) *In the model of uncertainty in Equation (6.10), when we use the l_2-norm, i.e., $\|\boldsymbol{s}\|^* = \|\boldsymbol{s}\|_2$, and under the assumption that $\boldsymbol{u} \sim \mathcal{N}(\boldsymbol{0}, I)$, we have the probability bound*

$$\mathbb{P}\left\{f(\boldsymbol{x}, \boldsymbol{D}) > 0\right\} \leq \frac{\sqrt{e}\Omega}{\alpha} \exp\left(-\frac{\Omega^2}{2\alpha^2}\right),$$

where $\alpha = 1$ for LPs, $\alpha = \sqrt{2}$ for SOCPs, and $\alpha = \sqrt{m}$ for SDPs (m is the dimension of the matrix in the SDP).

6.2 An exact approach using SDP

In this section, we apply the robust quadratic optimization results to formulate (6.8) as an SDP. We then show that we can compute optimal solutions to this SDP with a very simple control law.

First, exploiting the linearity of the system, we have the following straightforward result.

Proposition 6.2.1. The cost function (6.2) for system (6.1) can be written in the form

$$\begin{aligned} J(\boldsymbol{x}_0, \boldsymbol{u}, \boldsymbol{w}) &= 2\boldsymbol{a}^T\boldsymbol{x}_0 + \boldsymbol{x}_0^T\boldsymbol{A}\boldsymbol{x}_0 + 2\boldsymbol{b}^T\boldsymbol{u} + \boldsymbol{u}^T\boldsymbol{B}\boldsymbol{u} \\ &+ 2\boldsymbol{c}^T\boldsymbol{w} + \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} + 2\boldsymbol{u}^T\boldsymbol{D}\boldsymbol{w}, \end{aligned} \tag{6.14}$$

for appropriate vectors $\boldsymbol{a} \in \mathbb{R}^{(N\cdot n_x)\times 1}$, $\boldsymbol{b} \in \mathbb{R}^{(N\cdot n_u)\times 1}$, $\boldsymbol{c} \in \mathbb{R}^{(N\cdot n_w)\times 1}$ and matrices $\boldsymbol{A} \in \mathbb{R}^{(N\cdot n_x)\times(N\cdot n_x)}$, $\boldsymbol{B} \in \mathbb{R}^{(N\cdot n_u)\times(N\cdot n_u)}$, $\boldsymbol{C} \in \mathbb{R}^{(N\cdot n_w)\times(N\cdot n_w)}$, $\boldsymbol{D} \in \mathbb{R}^{(N\cdot n_u)\times(N\cdot n_w)}$, and where $\boldsymbol{B} \succ \boldsymbol{0}$, $\boldsymbol{C} \succeq \boldsymbol{0}$.

Proof. Since the system is linear, we can write the state at any instant k as

$$\boldsymbol{x}_k = \tilde{\boldsymbol{A}}_{k-1}\boldsymbol{x}_0 + \tilde{\boldsymbol{B}}_{k-1}\boldsymbol{u} + \tilde{\boldsymbol{C}}_{k-1}\boldsymbol{w},$$

where

$$\begin{aligned} \tilde{\boldsymbol{A}}_{k-1} &= \prod_{i=0}^{k-1} \boldsymbol{A}_i \\ \tilde{\boldsymbol{B}}_{k-1} &= \begin{bmatrix} (\prod_{j=1}^{k-1} \boldsymbol{A}_j)\boldsymbol{B}_0 & \cdots & \boldsymbol{B}_{k-1} & \boldsymbol{0}_{n_x\times(N-k)\cdot n_u} \end{bmatrix} \\ \tilde{\boldsymbol{C}}_{k-1} &= \begin{bmatrix} (\prod_{j=1}^{k-1} \boldsymbol{A}_j)\boldsymbol{C}_0 & \cdots & \boldsymbol{C}_{k-1} & \boldsymbol{0}_{n_x\times(N-k)\cdot n_w} \end{bmatrix}. \end{aligned}$$

Now the cost of any state term is written

$$\begin{aligned}
\boldsymbol{x}_k^T \boldsymbol{Q}_k \boldsymbol{x}_k + 2\boldsymbol{q}_k^T \boldsymbol{x}_k &= \boldsymbol{x}_0^T \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} \boldsymbol{x}_0 + 2\boldsymbol{x}_0^T \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{Q}_k \left(\tilde{\boldsymbol{B}}_{k-1} \boldsymbol{u} + \tilde{\boldsymbol{C}}_{k-1} \boldsymbol{w} \right) \\
&+ \boldsymbol{u}^T \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{B}}_{k-1} \boldsymbol{u} + \boldsymbol{w}^T \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \boldsymbol{w} + 2\boldsymbol{u}^T \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \boldsymbol{w} \\
&+ 2\boldsymbol{q}_k^T \left(\tilde{\boldsymbol{A}}_{k-1} \boldsymbol{x}_0 + \tilde{\boldsymbol{B}}_{k-1} \boldsymbol{u} + \tilde{\boldsymbol{C}}_{k-1} \boldsymbol{w} \right).
\end{aligned}$$

Thus, the overall cost is clearly written in the form stated above, with

$$\begin{aligned}
\boldsymbol{a} &= \sum_{k=1}^{N} \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{q}_k \\
\boldsymbol{A} &= \sum_{k=1}^{N} \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} \\
\boldsymbol{b} &= \hat{\boldsymbol{r}} + \left(\sum_{k=1}^{N} \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} \right) \boldsymbol{x}_0 + \sum_{k=1}^{N} \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{q}_k \\
\boldsymbol{B} &= \hat{\boldsymbol{R}} + \sum_{k=1}^{N} \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{B}}_{k-1} \\
\boldsymbol{c} &= \left(\sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} \right) \boldsymbol{x}_0 + \sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{q}_k \\
\boldsymbol{C} &= \sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \\
\boldsymbol{D} &= \sum_{k=1}^{N} \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1},
\end{aligned}$$

where

$$\begin{aligned}
\hat{\boldsymbol{r}} &= \left[\boldsymbol{r}_0^T \cdots \boldsymbol{r}_{N-1}^T \right]^T, \\
\hat{\boldsymbol{R}} &= \operatorname{diag}\left(\boldsymbol{R}_0, \ldots, \boldsymbol{R}_{N-1} \right).
\end{aligned}$$

Finally, positive (semi-) definiteness of $\boldsymbol{B}$ and $\boldsymbol{C}$ follow from positive (semi-) definiteness of $\boldsymbol{R}_k$ and $\boldsymbol{Q}_k$. □

Next, for ease of notation, we will transform the coordinates of the control space.

Proposition 6.2.2. To minimize the cost function $J(\boldsymbol{x}_0, \boldsymbol{u}, \boldsymbol{w})$ in Proposition 6.2.1 over all $\boldsymbol{u} \in \mathbb{R}^{N \cdot n_u}$, it is sufficient instead to optimize over all $\boldsymbol{y} \in \mathbb{R}^{N \cdot n_u}$ the cost function

$$\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) = \boldsymbol{y}^T\boldsymbol{y} + 2\boldsymbol{h}^T\boldsymbol{w} + 2\boldsymbol{y}^T\boldsymbol{F}\boldsymbol{w} + \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} \tag{6.15}$$

with $\boldsymbol{h} = \boldsymbol{c} - \boldsymbol{D}^T\boldsymbol{B}^{-1}\boldsymbol{b}$, $\boldsymbol{F} = \boldsymbol{B}^{-1/2}\boldsymbol{D}$.

Proof. The proof is immediate from the fact that $\boldsymbol{B}^{-1}$ exists since $\boldsymbol{B} \succ \boldsymbol{0}$, and then using the transformation $\boldsymbol{u} = \boldsymbol{B}^{-1/2}\boldsymbol{y} - \boldsymbol{B}^{-1}\boldsymbol{b}$. □

By Proposition 6.2.2, then, (6.8) is equivalent to the problem

$$\min_{\boldsymbol{y} \in \mathbb{R}^{N \cdot n_u}} \max_{\boldsymbol{w} \in \mathcal{W}_\gamma} \tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}). \tag{6.16}$$

This problem may be solved using SDP, as we now show.

Theorem 6.2.1. *Problem (6.16) may be solved by the following SDP:*

$$\begin{aligned} \text{minimize} \quad & z \\ \text{subject to} \quad & \begin{bmatrix} \boldsymbol{I} & \boldsymbol{y} & \boldsymbol{F} \\ \boldsymbol{y}^T & z - \gamma^2\lambda & -\boldsymbol{h}^T \\ \boldsymbol{F}^T & -\boldsymbol{h} & \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0}, \\ & \lambda \geq 0, \end{aligned} \tag{6.17}$$

in decision variables $\boldsymbol{y}$, z, and λ.

Proof. We first rewrite the problem as

$$\begin{aligned} \text{minimize} \quad & z \\ \text{subject to} \quad & z - \boldsymbol{y}^T\boldsymbol{y} - 2\boldsymbol{h}^T\boldsymbol{w} - 2\boldsymbol{y}^T\boldsymbol{F}\boldsymbol{w} - \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} \geq 0, \ \forall\ \boldsymbol{w} : \boldsymbol{w}^T\boldsymbol{w} \leq \gamma^2. \end{aligned} \tag{6.18}$$

We may homogenize the system and rewrite this equivalently as

$$\begin{aligned} &\text{minimize} \quad z \qquad\qquad (6.19) \\ &\text{subject to} \quad t^2(z - \boldsymbol{y}^T\boldsymbol{y}) - 2t\boldsymbol{h}^T\boldsymbol{w} - 2t\boldsymbol{y}^T\boldsymbol{F}\boldsymbol{w} - \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} \geq 0, \ \forall\ \boldsymbol{w}, t : \boldsymbol{w}^T\boldsymbol{w} \leq \gamma^2 t^2. \end{aligned}$$

Clearly, feasibility of $(z, \boldsymbol{y})$ in (6.19) implies feasibility of $(z, \boldsymbol{y})$ in (6.18) (by setting $t = 1$). For the other direction, assume $(z, \boldsymbol{y})$ is feasible in (6.18) and set $\tilde{\boldsymbol{w}} = t\boldsymbol{w}$, where $\boldsymbol{w}^T\boldsymbol{w} \leq \gamma^2$. This implies $\tilde{\boldsymbol{w}}^T\tilde{\boldsymbol{w}} \leq \gamma^2 t^2$, and

$$\begin{aligned} 2t\boldsymbol{h}^T\tilde{\boldsymbol{w}} + 2t\boldsymbol{y}^T\boldsymbol{F}\tilde{\boldsymbol{w}} - \tilde{\boldsymbol{w}}^T\boldsymbol{C}\tilde{\boldsymbol{w}} &= t^2\left(2\boldsymbol{h}^T\boldsymbol{w} + 2\boldsymbol{y}^T\boldsymbol{F}\boldsymbol{w} - \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w}\right) \\ &\leq t^2(z - \boldsymbol{y}^T\boldsymbol{y}), \end{aligned}$$

where the inequality follows by (6.18). Thus, the claim is true.

But now we wish to check whether a homogenous quadratic form in $(t, \boldsymbol{w})$ is nonnegative over all $(t, \boldsymbol{w})$ satisfying another quadratic form. Invoking Lemma 6.1.2, we know the constraint holds if and only if there exists a $\lambda \geq 0$ such that

$$\begin{bmatrix} z - \gamma^2\lambda - \boldsymbol{y}^T\boldsymbol{y} & -\boldsymbol{h}^T - \boldsymbol{y}^T\boldsymbol{F} \\ -\boldsymbol{h} - \boldsymbol{F}^T\boldsymbol{y} & \lambda\boldsymbol{I} - \boldsymbol{C} \end{bmatrix} \succeq \boldsymbol{0}$$

$$\Updownarrow$$

$$\begin{bmatrix} z - \gamma^2\lambda & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} - \begin{bmatrix} \boldsymbol{y}^T \\ \boldsymbol{F}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y} & \boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0}.$$

Finally, utilizing Lemma 6.1.1 with $\boldsymbol{B} = \boldsymbol{I}$ we see that this is equivalent to

$$\begin{bmatrix} \boldsymbol{I} & \boldsymbol{y} & \boldsymbol{F} \\ \boldsymbol{y}^T & z - \gamma^2\lambda & -\boldsymbol{h}^T \\ \boldsymbol{F}^T & -\boldsymbol{h} & \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0}.$$

Thus we arrive at the desired SDP. □

There is a tie between the standard DP approach (the Riccati equation) and this SDP, and the connection is not difficult to see.

Corollary 6.2.1. With $\gamma = 0$, the optimal solution to SDP (6.17) solves the Riccati equation, i.e., minimizes the cost-to-go in an expected value sense.

Proof. As argued in proposition 6.2.1, the total cost can be written in the form

$$J_{\boldsymbol{u}}(\boldsymbol{x}_0) \;=\; 2\boldsymbol{a}^T\boldsymbol{x}_0 + \boldsymbol{x}_0^T\boldsymbol{A}\boldsymbol{x}_0 + 2\boldsymbol{b}^T\boldsymbol{u} + \boldsymbol{u}^T\boldsymbol{B}\boldsymbol{u} + 2\boldsymbol{c}^T\boldsymbol{w} + \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} + 2\boldsymbol{u}^T\boldsymbol{D}\boldsymbol{w}.$$

With $\gamma = 0$, we require $\boldsymbol{w} = \boldsymbol{0} = \mathbb{E}[\boldsymbol{w}]$. Hence, we have

$$\max_{\boldsymbol{w}=\boldsymbol{0}} J_{\boldsymbol{u}}(\boldsymbol{x}_0) \;=\; \mathbb{E}[J(\boldsymbol{x}_0)] + c,$$

where $c = -\mathbb{E}\left[\boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w}\right]$. Since the goal of the dynamic programming approach is to minimize the expected cost, the equivalence of the two approaches in this case follows. □

In summary, Theorem 6.2.1 provides an exact SDP approach towards solving problem (6.16), and in the limiting case $\gamma = 0$, this SDP yields the same solution as the Riccati equation. It is not surprising that the complexity of the problem is that of solving an SDP; in fact, Yao et al. [118] have shown that solving the Riccati equation can be cast as an SDP.

6.2.1 Simplifying the SDP for closed-loop control

Although Theorem 6.2.1 ostensibly provides us with an open-loop policy, we can certainly run this approach in closed-loop. We would do this by solving (6.17) and applying the first n_u components of the solution as the current control. Then, with a new state observation, we update the data to Problem (6.17) (in fact, only $\boldsymbol{h}$ depends on the current state, so all other data for the problem can be computed off-line) and solve (6.17) again.

For large problem sizes and applications demanding rapid feedback, however, this approach is clearly impractical. In particular, solving large SDPs of the form (6.17) is expensive for large problem sizes, and this a serious drawback in real-time control settings. We would like a simplification which allows us to compute solutions much faster.

In this section, we show that Problem (6.17) has a very special structure which allows us to dramatically reduce the problem complexity. We will show how to compute optimal policies with only linear operations (i.e., matrix multiplication) in the current state plus a very simple optimization of a convex function over a scalar variable. In short, we will derive for our approach a control law which is an analog of the linear control law (Riccati) in the distributional framework. This control law, while nonlinear in the current state, can nonetheless be computed extremely efficiently and thus may be used very efficiently in closed-loop control.

To begin, we need the following simple observation.

Lemma 6.2.1. *Consider matrices $\boldsymbol{G}$, $\boldsymbol{U}$, and $\boldsymbol{V}$, of appropriate size, and let $\tilde{\boldsymbol{U}}$ and $\tilde{\boldsymbol{V}}$ be the orthogonal complements of $\boldsymbol{U}$ and $\boldsymbol{V}$, respectively (i.e., full-rank matrices such that $\tilde{\boldsymbol{U}}^T\boldsymbol{U} = 0$, $\tilde{\boldsymbol{V}}^T\boldsymbol{V} = 0$). If there exists a matrix $\boldsymbol{X}$ such that*

$$\boldsymbol{G} + \boldsymbol{U}\boldsymbol{X}\boldsymbol{V}^T + \boldsymbol{V}\boldsymbol{X}^T\boldsymbol{U}^T \succeq \boldsymbol{0},$$

then the following hold

$$\tilde{\boldsymbol{U}}^T\boldsymbol{G}\tilde{\boldsymbol{U}} \succeq \boldsymbol{0}, \tag{6.20}$$

$$\tilde{\boldsymbol{V}}^T\boldsymbol{G}\tilde{\boldsymbol{V}} \succeq \boldsymbol{0}. \tag{6.21}$$

Proof. Clear, by multiplying inequalities (6.20) and (6.21) by $(\tilde{\boldsymbol{U}}^T, \tilde{\boldsymbol{U}})$ and $(\tilde{\boldsymbol{V}}^T, \tilde{\boldsymbol{V}})$, respectively. □

Lemma 6.2.1 is actually a special case of an "elimination lemma." The statement is in fact true in both directions when the inequalities are all made strict (Boyd et al., [37]). We can now apply Lemma 6.2.1 to simplify Problem (6.17). From here on out we use the notation $||\boldsymbol{P}||_2$ as the spectral norm of a positive semi-definite matrix

$\boldsymbol{P} \in \mathbb{R}^{n\times n}$, i.e.,

$$\|\boldsymbol{P}\|_2 = \sqrt{\max_{i=1,\ldots,n} \lambda_i(\boldsymbol{P})},$$

where λ_i are the eigenvalues of $\boldsymbol{P}$.

Proposition 6.2.3. Let z^* be the optimal value of Problem (6.17). Then $z^* \geq z_R^*$, where z_R^* is the optimal value of the problem

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & \begin{bmatrix} z - \gamma^2\lambda & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \lambda \boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0}, \\ & \lambda \geq \|\boldsymbol{C}\|^2, \end{array} \tag{6.22}$$

in decision variables (z, λ).

Proof. Applying Lemma 6.2.1 to the constraint in (6.17), we see that we can write it in the form

$$\boldsymbol{G}(z,\lambda) + \boldsymbol{U}\boldsymbol{y}\boldsymbol{V}^T + \boldsymbol{V}\boldsymbol{y}^T\boldsymbol{U}^T \succeq \boldsymbol{0},$$

where

$$\begin{aligned} \boldsymbol{G}(z,\lambda) &= \begin{bmatrix} \boldsymbol{I} & \boldsymbol{0} & \boldsymbol{F} \\ \boldsymbol{0}^T & z - \gamma^2\lambda & -\boldsymbol{h}^T \\ \boldsymbol{F}^T & -\boldsymbol{h} & \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix}, \\ \boldsymbol{U} &= \begin{bmatrix} \boldsymbol{I}_{N\cdot n_u} \\ \boldsymbol{0} \end{bmatrix}, \\ \boldsymbol{V} &= \boldsymbol{e}_{(N\cdot n_u)+1}. \end{aligned}$$

We now invoke Lemma 6.2.1 with

$$\tilde{\boldsymbol{U}} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{I}_{N \cdot n_w + 1} \end{bmatrix},$$

$$\tilde{\boldsymbol{V}} = \begin{bmatrix} \boldsymbol{I}_{N \cdot n_u} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{I}_{N \cdot n_w} \end{bmatrix},$$

where all zero matrices are sized appropriately. With these choices, if there exists a $\boldsymbol{y}$ such that (6.17) is feasible for a given (z, λ), then the following inequalities also hold for such a (z, λ):

$$\begin{bmatrix} z - \gamma^2 \lambda & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \lambda \boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T \boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0},$$

$$\begin{bmatrix} \boldsymbol{I} & \boldsymbol{F} \\ \boldsymbol{F}^T & \lambda \boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T \boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0},$$

and, by Schur complements, the latter inequality is equivalent to $\lambda \boldsymbol{I} - \boldsymbol{C} \succeq \boldsymbol{0} \Rightarrow \lambda \geq \|\boldsymbol{C}\|^2$, since $\boldsymbol{C} \succeq \boldsymbol{0}$. It follows that any (z, λ) which is feasible to (6.17) is feasible to (6.22), and hence $z^* \geq z_R^*$. □

Before analyzing the structure of Problem (6.22) in more detail, we note the following definition, which we will employ for notational convenience.

Definition 6.2.1. Consider a matrix $\boldsymbol{X}$ such that $\boldsymbol{X} \succeq \boldsymbol{0}$ with eigenvalue decomposition written as $\boldsymbol{X} = \boldsymbol{Q} \Lambda \boldsymbol{Q}^T$. Denote by $\hat{\boldsymbol{X}}$ the **unit full-rank version of $\boldsymbol{X}$**, with

$$\hat{\boldsymbol{X}} = \boldsymbol{Q} \hat{\Lambda} \boldsymbol{Q}^T, \tag{6.23}$$

where $\hat{\mathbf{\Lambda}}$ is a diagonal matrix such that

$$[\hat{\mathbf{\Lambda}}]_{ii} = \begin{cases} [\mathbf{\Lambda}]_{ii}, & \text{if } [\mathbf{\Lambda}]_{ii} > 0, \\ 1, & \text{otherwise.} \end{cases}$$

Note $\hat{\boldsymbol{X}} \succ \mathbf{0}$ always, so $\hat{\boldsymbol{X}}^{-1}$ always exists, and that if $\boldsymbol{X} \succ \mathbf{0}$, then $\hat{\boldsymbol{X}} = \boldsymbol{X}$.

We now show that Problem (6.22) can be reduced to a simpler optimization problem involving no semidefinite constraints and just the variable λ. In what follows we will denote the eigenvalues and eigenvectors of $\boldsymbol{F}^T\boldsymbol{F}-\boldsymbol{C}$ by λ_i and $\boldsymbol{q}_i$, $i = 1, \ldots, N{\cdot}n_w$, respectively.

Proposition 6.2.4. Problem (6.22) is equivalent to the convex optimization problem (in single variable λ)

$$\begin{array}{ll} \text{minimize} & f(\lambda) \\ \text{subject to} & \lambda \geq \|\boldsymbol{C}\|^2, \end{array} \tag{6.24}$$

where

$$f(\lambda) = \begin{cases} \gamma^2\lambda + \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda)\boldsymbol{h}, & \text{if } \boldsymbol{q}_i^T\boldsymbol{h} = 0 \quad \forall\, i \,:\, \lambda + \lambda_i = 0, \\ +\infty & \text{otherwise,} \end{cases} \tag{6.25}$$

and $\boldsymbol{H}(\lambda) = \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F}$.

Proof. By Schur complements, a pair (z, λ) is feasible in (6.22) if and only if

$$\begin{aligned} \boldsymbol{H}(\lambda) - \left(\frac{1}{z - \gamma^2\lambda}\right)\boldsymbol{h}\boldsymbol{h}^T \succeq \mathbf{0} \;&\Leftrightarrow\; (z - \gamma^2\lambda)\boldsymbol{x}^T\boldsymbol{H}(\lambda)\boldsymbol{x} \geq (\boldsymbol{h}^T\boldsymbol{x})^2 \;\forall\, \boldsymbol{x} \\ &\Leftrightarrow\; z - \gamma^2\lambda \geq v^2, \end{aligned}$$

where v is the optimal value of the problem

$$\begin{array}{ll} \text{maximize} & \boldsymbol{h}^T \boldsymbol{x} \\ \text{subject to} & \boldsymbol{x}^T \boldsymbol{H}(\lambda) \boldsymbol{x} \leq 1. \end{array}$$

Note that feasibility of λ implies that $\lambda + \lambda_i \geq 0$ for all i. Let $I_+ = \{i \mid \lambda + \lambda_i > 0\}$. Then carrying out the above optimization problem, we find

$$\begin{aligned} v &= \begin{cases} \sqrt{\sum_{i \in I_+} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{\lambda + \lambda_i}}, & \text{if } \boldsymbol{h}^T \boldsymbol{q}_j = 0 \; \forall \; j \notin I_+, \\ +\infty, & \text{otherwise,} \end{cases} \\ &= \begin{cases} \sqrt{\boldsymbol{h} \hat{\boldsymbol{H}}^{-1}(\lambda) \boldsymbol{h}}, & \text{if } \boldsymbol{h}^T \boldsymbol{q}_j = 0 \; \forall \; j \notin I_+, \\ +\infty, & \text{otherwise.} \end{cases} \end{aligned}$$

From the above equivalences, then, we have (z, λ) feasible to (6.22) if and only if

$$z \geq \gamma^2 \lambda + \boldsymbol{h} \hat{\boldsymbol{H}}^{-1}(\lambda) \boldsymbol{h}.$$

Since we wish to minimize z, for a fixed λ, we should set z equal to $\gamma^2 \lambda + \boldsymbol{h} \hat{\boldsymbol{H}}^{-1}(\lambda) \boldsymbol{h}$. Thus (6.22) is equivalent to minimization of $f(\lambda)$ over all $\lambda \geq \|\boldsymbol{C}\|^2$. □

We now argue that optimization of $f(\lambda)$ may be done efficiently.

Proposition 6.2.5. The function $f(\lambda)$ in (6.25) satisfies the following:

(a) $f(\lambda)$ is convex on $\lambda \geq \|\boldsymbol{C}\|^2$.

(b) If $\gamma \geq \gamma_{\text{thresh}}$, where

$$\gamma_{\text{thresh}} = \begin{cases} \|\hat{\boldsymbol{H}}^{-1}(\|\boldsymbol{C}\|^2) \boldsymbol{h}\|, & \text{if } \boldsymbol{q}_i^T \boldsymbol{h} = 0 \; \forall \; i \; : \; \|\boldsymbol{C}\|^2 + \lambda_i = 0, \\ +\infty, & \text{otherwise.} \end{cases} \tag{6.26}$$

then $\lambda^* = \|\boldsymbol{C}\|^2$ minimizes $f(\lambda)$ over all $\lambda \geq \|\boldsymbol{C}\|^2$.

(c) If $\gamma < \gamma_{\text{thresh}}$, then the minimizer λ^* of $f(\lambda)$ over all $\lambda \geq \|\boldsymbol{C}\|^2$ may be found (within tolerance ϵ) in time $\mathcal{O}(\kappa)$, where

$$\kappa = \log_2 \left\lceil \frac{\sqrt{m}\|\boldsymbol{Q}^T\boldsymbol{h}\|_\infty - \gamma(\|\boldsymbol{C}\|^2 + \min_{i=1,\ldots,N\cdot n_w} \lambda_i)}{\gamma\epsilon} \right\rceil, \tag{6.27}$$

the columns of $\boldsymbol{Q}$ are the eigenvectors of $\boldsymbol{F}^T\boldsymbol{F} - \boldsymbol{C}$, and $m \leq N \cdot n_w$ is the number of eigenvectors such that $\boldsymbol{q}_i^T\boldsymbol{h} \neq 0$.

Proof.

(a) We may write $f(\lambda)$ as

$$f(\lambda) = \gamma^2\lambda + \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T\boldsymbol{h})^2}{\lambda + \lambda_i},$$

which is clearly a convex function in λ over $\lambda \geq \|\boldsymbol{C}\|^2 \Rightarrow \lambda + \lambda_i \geq 0$.

(b) If $\gamma \geq \gamma_{\text{thresh}}$, we have, for all $\lambda \geq \|\boldsymbol{C}\|^2$,

$$\begin{aligned} f'(\lambda) &= \gamma^2 - \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T\boldsymbol{h})^2}{(\lambda + \lambda_i)^2} \\ &\geq \gamma_{\text{thresh}}^2 - \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T\boldsymbol{h})^2}{(\lambda + \lambda_i)^2} \\ &= \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T\boldsymbol{h})^2}{(\|\boldsymbol{C}\|^2 + \lambda_i)^2} - \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T\boldsymbol{h})^2}{(\lambda + \lambda_i)^2} \\ &\geq 0. \end{aligned}$$

so $f(\lambda)$ is nondecreasing over $\lambda \geq \|\boldsymbol{C}\|^2$.

(c) If $\gamma < \gamma_{\text{thresh}}$, then we must search for $\lambda^* > \|\boldsymbol{C}\|^2$ such that $f'(\lambda^*) = 0$ (since, by **(a)**, $f(\lambda)$ is convex). This is the same as finding a root $\lambda^* > \|\boldsymbol{C}\|^2$ of the nonlinear equation

$$\sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{(\lambda + \lambda_i)^2} \;=\; \gamma^2. \tag{6.28}$$

We claim $\lambda^* \leq \bar{\lambda}$, where

$$\bar{\lambda} \;=\; \frac{\sqrt{m}\|\boldsymbol{Q}^T \boldsymbol{h}\|_\infty}{\gamma} - \min_{i=1,\ldots,N\cdot n_w} \lambda_i.$$

Indeed, assume $\lambda^* > \bar{\lambda}$. Then

$$\begin{aligned}
f(\lambda^*) &= \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{(\lambda^* + \lambda_i)^2} \\
&< \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{(\bar{\lambda} + \lambda_i)^2} \\
&\leq \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{m\|\boldsymbol{Q}^T \boldsymbol{h}\|_\infty^2/\gamma^2} \\
&= \frac{\gamma^2}{m} \sum_{i=1}^{N\cdot n_w} \frac{(\boldsymbol{q}_i^T \boldsymbol{h})^2}{\|\boldsymbol{Q}^T \boldsymbol{h}\|_\infty^2} \\
&\leq \gamma^2,
\end{aligned}$$

which implies that λ^* cannot be a solution of (6.28). The result then follows by applying bisection (with tolerance ϵ) on the interval $[\|\boldsymbol{C}\|^2, \bar{\lambda}]$. □

We are finally ready for the main result of this section. Given an optimal solution λ^* to (6.24), we can then compute our closed-loop optimal simply by performing matrix multiplications.

Theorem 6.2.2. *Let λ^* be the optimal solution to Problem (6.24). Then the solution $(\lambda^*, z^*, \boldsymbol{y}^*)$, where*

$$\begin{aligned}
z^* &= f(\lambda^*), \\
\boldsymbol{y}^* &= -\boldsymbol{F}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h}.
\end{aligned} \tag{6.29}$$

is an optimal solution to the SDP (6.17).

Proof. By Proposition 6.2.3 we know that $z^* \geq z_R^*$, and, by Proposition 6.2.4 we have $z_R^* = f(\lambda^*)$, so if we can just show that the solution $(\lambda^*, z^*, \boldsymbol{y}^*)$ in (6.29) is feasible to (6.17), then we know it must be optimal. First, since $\lambda^* \geq \|\boldsymbol{C}\|^2 \geq 0$, we have $\lambda^* \geq 0$, as required. For the semi-definite constraint in (6.17), by Schur complements, we require

$$\begin{bmatrix} z^* - \gamma^2\lambda^* & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \lambda^*\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} - \begin{bmatrix} \boldsymbol{y}^{*T} \\ \boldsymbol{F}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y}^* & \boldsymbol{F} \end{bmatrix} \succeq \boldsymbol{0}.$$

As before, let $\boldsymbol{H}(\lambda^*) = \lambda^*\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F}$; note that feasibility of λ^* in (6.24) implies $\boldsymbol{H}(\lambda^*) - \boldsymbol{F}^T\boldsymbol{F} = \lambda^*\boldsymbol{I} - \boldsymbol{C} \succeq \boldsymbol{0}$, so $\lambda^*\boldsymbol{I} - \boldsymbol{C} = \boldsymbol{G}^T\boldsymbol{G}$ for some matrix $\boldsymbol{G}$. Finally, λ^* minimizes $f(\lambda^*)$ over all $\lambda \geq \|\boldsymbol{C}\|^2$, so $f(\lambda^*)$ must be finite and thus $z^* - \gamma^2\lambda^* = \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h}$. Putting all of this together, we have

$$\begin{aligned}
& \begin{bmatrix} z^* - \gamma^2\lambda^* & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \lambda^*\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} - \begin{bmatrix} \boldsymbol{y}^{*T} \\ \boldsymbol{F}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y}^* & \boldsymbol{F} \end{bmatrix} \\
= & \begin{bmatrix} \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & -\boldsymbol{h}^T \\ -\boldsymbol{h} & \hat{\boldsymbol{H}}(\lambda^*) \end{bmatrix} - \begin{bmatrix} \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{F}^T\boldsymbol{F}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & -\boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{F}^T\boldsymbol{F} \\ -\boldsymbol{F}^T\boldsymbol{F}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & \boldsymbol{F}^T\boldsymbol{F} \end{bmatrix} \\
= & \begin{bmatrix} \boldsymbol{h}\left(\hat{\boldsymbol{H}}^{-1}(\lambda^*) - \hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{F}^T\boldsymbol{F}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\right)\boldsymbol{h} & -\boldsymbol{h}^T\left(I - \hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{F}^T\boldsymbol{F}\right) \\ -\left(\boldsymbol{I} - \boldsymbol{F}^T\boldsymbol{F}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\right)\boldsymbol{h} & \lambda^*\boldsymbol{I} - \boldsymbol{C} \end{bmatrix} \\
= & \begin{bmatrix} \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\left(\hat{\boldsymbol{H}}(\lambda^*) - \boldsymbol{F}^T\boldsymbol{F}\right)\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & -\boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\left(\hat{\boldsymbol{H}}(\lambda^*) - \boldsymbol{F}^T\boldsymbol{F}\right) \\ \left(\hat{\boldsymbol{H}}(\lambda^*) - \boldsymbol{F}^T\boldsymbol{F}\right)\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & \lambda^*\boldsymbol{I} - \boldsymbol{C} \end{bmatrix} \\
= & \begin{bmatrix} \boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\left(\lambda^*\boldsymbol{I} - \boldsymbol{C}\right)\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & -\boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\left(\lambda^*\boldsymbol{I} - \boldsymbol{C}\right) \\ -\left(\lambda^*\boldsymbol{I} - \boldsymbol{C}\right)\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & \lambda^*\boldsymbol{I} - \boldsymbol{C} \end{bmatrix} \\
= & \begin{bmatrix} -\boldsymbol{h}^T\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{G}^T \\ \boldsymbol{G}^T \end{bmatrix} \begin{bmatrix} -\boldsymbol{G}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\boldsymbol{h} & \boldsymbol{G} \end{bmatrix} \\
\succeq & \ 0.
\end{aligned}$$

which completes the proof. □

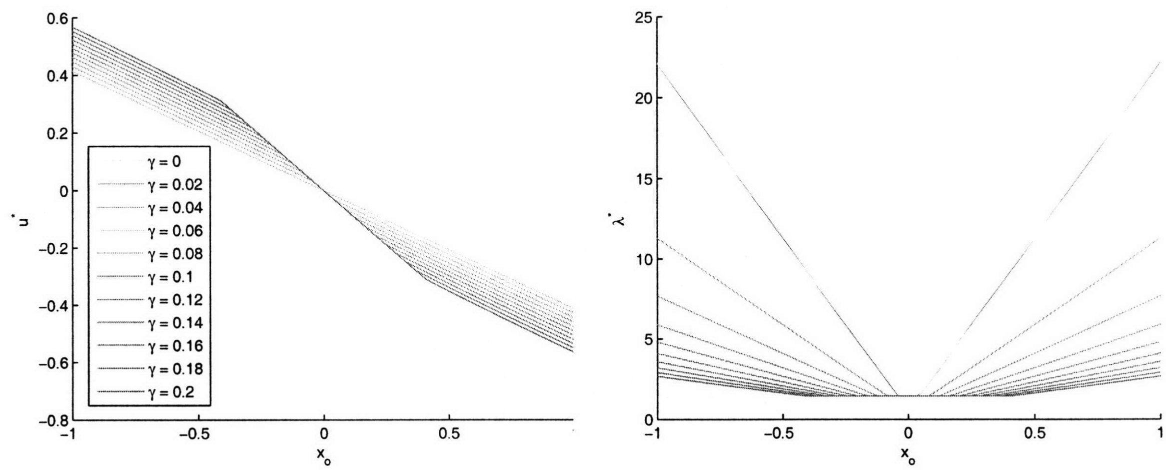

Figure 6-1: *Optimal control law (left) versus initial state; optimal value of λ in optimal control law (right) versus initial state. Both plotted for various values of γ.*

We reiterate that Theorem 6.2.2 provides us with a control vector $\boldsymbol{y} \in \mathbb{R}^{N \cdot n_u}$ for *all* remaining stages. When we run this in closed-loop, however, we would just take the first n_u components and apply that as the control to the current stage.

Figure 6-1 displays a simple example illustrating the optimal control law (as well as optimal value of λ) calculated via Theorem 6.2.2 versus the initial state for various values of γ. This is the first control for a 10-stage problem with time-invariant state, control, and disturbance matrices $\boldsymbol{A} = 1$, $\boldsymbol{B} = 1$, and $\boldsymbol{C} = 1$. The cost function is given by $\boldsymbol{Q}_k = (1/2)^k$, $\boldsymbol{R}_k = (1/2)^k$, and $\boldsymbol{q}_k = \boldsymbol{r}_k = 0$ for all $k \in \{0, \dots, 10\}$. In this case, the optimal control is *piecewise* linear in the initial state. Note that the case $\gamma = 0$ yields a linear policy (i.e., the Riccati approach), which we know must be the case from Corollary 6.2.1.

Thus, we see that the optimal, closed-loop control law for the approach given by Problem (6.16) may be computed in the following way: first, by computing the optimal value λ^* (which may be done by bisection, via Proposition 6.2.5), then by matrix multiplication. The only data in the SDP (6.17) which depends on the current state $\boldsymbol{x}_0$ is the vector $\boldsymbol{h}$. Thus, much of the work may be done offline. We now quantify explicitly the online computational burden.

Corollary 6.2.2. Consider the problem setup from (6.17), with N the number of stages, n_x, n_u, and n_w the sizes of the state, control, and disturbance vectors, respectively. Then the optimal closed-loop policy for a single period may be computed in $\mathcal{O}\left(Nn_w\left(\kappa + n_x + n_u\right)\right)$ time, where κ is given in Equation (6.27).

Proof. Since we only care about the current control, and $\boldsymbol{h}$ is linear in the current state $\boldsymbol{x}_0$, we may write re-express the closed-loop control for the current optimal control $\boldsymbol{u}^*_{\text{curr}}$ as

$$\boldsymbol{u}^*_{\text{curr}} \quad = \quad \bar{\boldsymbol{u}} + \tilde{\boldsymbol{F}}\hat{\boldsymbol{H}}^{-1}(\lambda^*)\tilde{\boldsymbol{G}}\boldsymbol{x}_0,$$

where $\tilde{\boldsymbol{F}} \in \mathbb{R}^{n_u \times Nn_w}$, $\tilde{\boldsymbol{G}} \in \mathbb{R}^{Nn_w \times n_x}$, and $\hat{\boldsymbol{H}}^{-1}(\lambda^*) \in \mathbb{R}^{Nn_w \times Nn_w}$. Of course, since $H(\lambda) = \lambda\boldsymbol{I} - \boldsymbol{C} + \boldsymbol{F}^T\boldsymbol{F}$, we may compute an eigenvalue decomposition for $H(0) = \boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}^T$ offline, and computing the inverse of $\hat{\boldsymbol{H}}(\lambda^*)$ may be done simply by adding λ^* to the diagonal elements of $\boldsymbol{\Lambda}$. Our total computational burden breaks down as follows

+ $\mathcal{O}(Nn_wn_x)$ iterations for setting up the optimization of $f(\lambda)$ (i.e., computing $\boldsymbol{Q}^T\boldsymbol{h}$).

+ $\mathcal{O}(\kappa)$ bisection calls, each requiring a sum over Nn_w terms, for a total effort of $\mathcal{O}(\kappa \cdot Nn_w)$ searching for λ^*.

+ Right-matrix multiplication of the current state: $\mathcal{O}(Nn_wn_x)$.

+ Scaling the resulting vector componentwise by $\lambda^* + \lambda_i$: $\mathcal{O}(Nn_w)$.

+ Left-matrix multiplication to obtain optimal control: $\mathcal{O}(Nn_wn_u)$.

The total effort required is thus

$$\mathcal{O}\left(Nn_w\left(1 + \kappa + n_x + 1 + n_u\right)\right) \quad = \quad \mathcal{O}\left(Nn_w\left(\kappa + n_x + n_u\right)\right).$$

□

	n_w		
N	1	10	100
1	1e-8	1e-6	1e-5
10	1e-7	1e-5	1e-4
100	1e-6	1e-4	1e-3
1000	1e-5	1e-3	1e-2
10000	1e-4	1e-2	1e-1

Table 6.1: *Computational effort (sec.) for various problem sizes on a 1GHz machine. Here,* $\kappa \approx 10$ *and* $n_x \approx n_u \approx n_w$ *are assumed.*

We note that the computational effort, for all other inputs fixed, grows *linearly* with the number of stages N. Of course, by computing the matrices offline, we are reducing the computational effort by increasing storage requirements. Our total memory usage is to store N matrices of sizes $n_u \cdot Nn_w$ and $Nn_w \cdot n_x$, for a total memory requirement of $\mathcal{O}\left(N^2 n_w \left(n_x + n_u\right)\right)$. Hence, memory increases quadratically with the time horizon N.

Table 6.1 illustrates order-of-magnitude estimates for the computational time for various problem sizes.

6.3 An inner approximation using SOCP

In the presence of constraints, we cannot use the simplification results of Section 6.2.1. This means for constrained control with feedback we would need to solve a problem of the same form as (6.17) with constraints at each stage. For large problems and applications demanding fast decisions, this will not be feasible. This motivates us to find a simplification of the exact SDP in (6.17).

Here we will develop an inner approximation using SOCP (i.e., any feasible solution to the SOCP will be feasible to SDP (6.17)). We will exploit the robust conic optimization results highlighted in Section 6.1.4.

Recall that our cost-to-go can be written as

$$\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) \quad = \quad \boldsymbol{y}^T\boldsymbol{y} + 2(\boldsymbol{h} + \boldsymbol{F}^T\boldsymbol{y})^T\boldsymbol{w} + \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w},$$

As before, we would like to find the policy $\boldsymbol{y}$ which minimizes the maximum value of $\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w})$ over all $\boldsymbol{w} \in \mathbb{R}^{N \cdot n_w}$ in some ellipsoidal uncertainty set. We may write this problem as

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & f(\boldsymbol{y}, \boldsymbol{w}) \leq z - \|\boldsymbol{y}\|_2^2 \qquad \forall\, \boldsymbol{w} \in \mathcal{W}_\Omega, \end{array} \tag{6.30}$$

where $f(\boldsymbol{y}, \boldsymbol{w}) = 2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{w} + \boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w}$, our uncertainty model is

$$\mathcal{W}_\Omega = \left\{ \boldsymbol{w} \;\middle|\; \exists \boldsymbol{u} \in \mathbb{R}^{N \cdot n_u} : \; \boldsymbol{w} = \sum_{j=1}^{N \cdot n_w} u_j \boldsymbol{e}^j, \; \|\boldsymbol{u}\|_2 \leq \Omega \right\}, \tag{6.31}$$

and $\{\boldsymbol{e}^1, \ldots, \boldsymbol{e}^{N \cdot n_w}\}$ is any orthonormal basis of $\mathbb{R}^{N \cdot n_w}$. Note that $\mathcal{W}_\Omega$ is precisely the same uncertainty set utilized in Section 6.2, with Ω assuming an analogous role as γ. In the framework of Section 6.1.4, we are using the assignments $\boldsymbol{D} = \boldsymbol{w}$, $\boldsymbol{D}^0 = \boldsymbol{0}$, and $\boldsymbol{D}^j = \boldsymbol{e}^j$.

From here we would like to directly apply the results from Section 6.1.4. The difficulty, however, is that the quadratic term $\boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w}$ causes $f(\boldsymbol{y}, \boldsymbol{w})$ to violate Assumption 6.1.1(b). We remove this difficulty with a slight relaxation of problem (6.30).

Proposition 6.3.1. Consider the problem

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & \tilde{f}(\boldsymbol{y}, \boldsymbol{w}) \leq z - \|\boldsymbol{y}\|_2^2 - \Omega^2 \|\boldsymbol{C}\|_2^2 \qquad \forall\, \boldsymbol{w} \in \mathcal{W}_\Omega, \end{array} \tag{6.32}$$

where $\tilde{f}(\boldsymbol{y}, \boldsymbol{w}) = 2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{w}$, $\mathcal{W}_\Omega$ is as described in Eq. (6.31). We then have the following:

(a) If a solution $(z, \boldsymbol{y})$ is feasible in problem (6.32), then it is also feasible in problem (6.30).

(b) The function $\tilde{f}(\boldsymbol{y}, \boldsymbol{w})$ satisfies Assumption 6.1.1.

Proof.

(a) Let $(z, \boldsymbol{y})$ be feasible in (6.32), and consider any $\boldsymbol{w} \in \mathcal{W}_\Omega$. Then we have

$$\begin{aligned} f(\boldsymbol{y}, \boldsymbol{w}) &= 2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{w} + \boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w} \\ &\leq 2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{w} + \Omega^2 \|\boldsymbol{C}\|_2^2 \qquad (6.33) \\ &\leq z - \|\boldsymbol{y}\|_2^2, \qquad (6.34) \end{aligned}$$

where (6.33) follows from the fact that

$$\max_{\|\boldsymbol{x}\| \leq \Omega} \boldsymbol{x}^T \boldsymbol{C} \boldsymbol{x} = \Omega^2 \|\boldsymbol{C}\|_2^2,$$

and (6.34) follows from feasibility of $(z, \boldsymbol{y})$ in (6.32). Thus, $(z, \boldsymbol{y})$ is also feasible in the original formulation in (6.30).

(b) This follows trivially, since $\tilde{f}(\boldsymbol{y}, \boldsymbol{w})$ is linear in $\boldsymbol{w}$. □

Now that we have cast the problem in the framework of Section 6.1.4, we may apply the corresponding results. This leads us to our formulation of problem (6.30) as a second-order cone problem, as we now illustrate.

Theorem 6.3.1. *Consider the SOCP*

$$\begin{aligned} \text{minimize} \quad & z \\ \text{subject to} \quad & \|\boldsymbol{y}\|_2^2 \leq z - \Omega \tilde{y} - \Omega^2 \|\boldsymbol{C}\|_2^2 \\ & |2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{e}^j| \leq t_j, \qquad j = 1, \ldots, N \cdot n_w, \qquad (6.35) \\ & \|\boldsymbol{t}\|_2 \leq \tilde{y}, \end{aligned}$$

in decision variables $(z, \boldsymbol{y}, \tilde{y}, \boldsymbol{t})$. If $(z, \boldsymbol{y})$ are part of a feasible solution to (6.35), then $(z, \boldsymbol{y})$ are also feasible in problem (6.30).

Proof. From Theorem 6.1.1, problem (6.35) is equivalent to the problem

$$\begin{array}{ll} \text{minimize} & z \\ \text{subject to} & \tilde{f}(\boldsymbol{y},\boldsymbol{0}) + \max\limits_{(\boldsymbol{u},\boldsymbol{v})\in\mathcal{V}_\Omega} \sum\limits_{j=1}^{N\cdot n_w} \left\{ \tilde{f}(\boldsymbol{y},\boldsymbol{e}^j)u_j + \tilde{f}(\boldsymbol{y},-\boldsymbol{e}^j)v_j \right\} \leq z - \|\boldsymbol{y}\|_2^2 - \Omega^2\|\boldsymbol{C}\|_2^2, \end{array}$$

where

$$\mathcal{V}_\Omega \;=\; \left\{ (\boldsymbol{u},\boldsymbol{v}) \in \mathbb{R}_+^{2(N\cdot n_w)} \;\middle|\; \|\boldsymbol{u}+\boldsymbol{v}\|_2 \leq \Omega \right\}.$$

Now, since $\tilde{f}(\boldsymbol{y},\boldsymbol{w})$ is linear in $\boldsymbol{w}$, this problem is equivalent to problem (6.32), by Proposition 6.1.1(a). Finally, invoking Proposition 6.3.1(a), we have that feasibility of $(z,\boldsymbol{y})$ in (6.32) implies feasibility in (6.30), and thus we are done. □

Theorem 6.3.1 thus gives us an inner approximation to the exact problem given in (6.30). This is in contrast to Theorem 6.2.1, which solves the problem exactly using SDP (and, in the unconstrained case, can be simplified via the results in Section 6.2.1). Theorem 6.3.1 gives us an SOCP formulation, which is a significant reduction in complexity from the SDP. In addition, we expect this approximation to be quite tight, as the *only* inequality we have exploited is $\boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} \leq \Omega^2\|\boldsymbol{C}\|_2^2$, which holds for all $\boldsymbol{w} \in \mathcal{W}_\Omega$. We now quantify this difference.

Corollary 6.3.1. Let z^*_{SDP} and z^*_{SOCP} be the optimal values of the SDP and SOCP (with $\Omega = \gamma$) given in Theorems 6.2.1 and 6.3.1, respectively. Then we have the relation

$$z^*_{\text{SOCP}} - z^*_{\text{SDP}} \;\leq\; 2\gamma\|\boldsymbol{h}\|_2. \tag{6.36}$$

Proof. Note that we have

$$\begin{aligned} z^*_{\text{SOCP}} &= \gamma^2\|\boldsymbol{C}\|_2^2 + \min_{\boldsymbol{y}} \left[\|\boldsymbol{y}\|_2^2 + 2 \max_{\boldsymbol{w}:\|\boldsymbol{w}\|_2\leq\gamma} \left(\boldsymbol{h} + \boldsymbol{F}^T\boldsymbol{y}\right)^T \boldsymbol{w} \right] \\ &\leq \gamma^2\|\boldsymbol{C}\|_2^2 + 2 \max_{\boldsymbol{w}:\|\boldsymbol{w}\|_2\leq\gamma} \boldsymbol{h}^T\boldsymbol{w} \\ &= \gamma^2\|\boldsymbol{C}\|_2^2 + 2\gamma\|\boldsymbol{h}\|_2, \end{aligned}$$

where the inequality follows from feasibility of $\boldsymbol{y} = \mathbf{0}$. On the other hand, positive semi-definiteness of the matrix in Theorem 6.2.1 requires

$$\begin{aligned} z^*_{\text{SDP}} &\geq \|\boldsymbol{y}\|_2^2 + \gamma^2\lambda \\ &\geq \gamma^2\lambda \\ &\geq \gamma^2\|\boldsymbol{C}\|_2^2, \end{aligned}$$

where the last line follows from Lemma 6.1.2. The result in Equation (6.36) now follows. □

6.4 Constraints and performance guarantees

We now demonstrate the modelling power of the approaches developed in Theorems 6.2.1 and 6.3.1. In particular, we show that both approaches readily lend themselves towards handling a wide variety of constraints. These constraints fit into three categories: control constraints, probabilistic guarantees on the state, and probabilistic guarantees on the cost function. For the probabilistic guarantees, we will assume the disturbances are independently and normally distributed.[3]

We present the results here for both the SDP and SOCP frameworks. Since the presence of constraints destroys the simple control law for the SDP from Section 6.2.1, however, the SOCP is more viable in a constrained, closed-loop control setting (in fact, we reiterate that this was the primary motivation for the development of the SOCP approach).

Throughout this section we will make claims about the "complexity type" of the problem being unchanged. By this we mean the SDP remains an SDP and the SOCP remains an SOCP. We implicitly appeal to the fact that the class of SDP problems includes the class of SOCP problems, and thus we may add second-order

[3]The fact that we are now making distributional assumptions whereas the approaches developed in Sections 6.2 and 6.3 are independent of such assumptions should not be viewed as contradictory. Our point is that these approaches are general in that they are derived in the absence of a probabilistic framework, yet flexible and powerful in that they can be augmented with probabilistic guarantees on both the state and cost function under the common assumption of normally distributed disturbances.

cone constraints to an SDP without increasing its complexity type.

We turn first to the simplest case of control constraints.

6.4.1 Control constraints

We will show that both approaches may handle any convex quadratic constraint on the control vector. In this and the following section we temporarily revert to the traditional notation $\boldsymbol{u}$ for the controls and note that the simple affine transformation listed in Proposition 6.2.2 allows us to implement these constraints in our $\boldsymbol{y}$ control space. We first need the following well-known result.

Proposition 6.4.1. The quadratic constraint $\boldsymbol{x}^T\boldsymbol{x} \leq t$ is equivalent to the second-order cone constraint

$$\left\| \begin{pmatrix} \boldsymbol{x} \\ \frac{t-1}{2} \end{pmatrix} \right\|_2 \leq \frac{t+1}{2}.$$

Proof. This is a standard result (see, e.g. Ben-Tal and Nemirovski [17]) which follows by noting that $t = \frac{(t+1)^2}{4} - \frac{(t-1)^2}{4}$. □

We now have the rather straightforward result of this section.

Theorem 6.4.1. *Any control constraints of the form*

$$\|\boldsymbol{G}\boldsymbol{u}\|_2^2 + 2\boldsymbol{g}^T\boldsymbol{u} + \hat{g} \leq 0, \tag{6.37}$$

where $\boldsymbol{G} \in \mathbb{R}^{(N \cdot n_u)\times(N \cdot n_u)}$, $\boldsymbol{g} \in \mathbb{R}^{(N \cdot n_u)\times 1}$, and $\hat{g} \in \mathbb{R}$, may be suitably added to problems (6.17) and (6.35) without increasing their respective complexity types.

Proof. By Proposition 6.4.1 we may write (6.37) as

$$\left\| \begin{pmatrix} \boldsymbol{G}\boldsymbol{u} \\ -\boldsymbol{g}^T\boldsymbol{u} - \frac{(\hat{g}+1)}{2} \end{pmatrix} \right\|_2 \leq -\boldsymbol{g}^T\boldsymbol{u} - \frac{(\hat{g}-1)}{2}.$$

which is a second-order cone constraint and hence may be added to either problem without raising the complexity type. □

Note that Theorem 6.4.1 implies we can tractably deal with any polyhedral or ellipsoidal constraints on the control.

6.4.2 Probabilistic state guarantees

Since the state $\boldsymbol{x}$ of the system is not exactly known, any constraints on $\boldsymbol{x}$ can only be enforced in a probabilistic sense. To ensure probabilistic guarantees in what follows, we will operate under the typical assumption that the disturbances $\boldsymbol{w}$ are independently and normally distributed.[4]

Assumption 6.4.1. The disturbances $\boldsymbol{w}$ are independently and normally distributed with zero mean, i.e.,

$$\boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}).$$

Note that, if instead we have $\boldsymbol{w} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} \succ \boldsymbol{0}$, we may simply rotate coordinates and multiply the $\boldsymbol{C}_k$ matrices in the dynamics in Equation (6.1) accordingly.

We now show how to explicitly ensure that linear constraints on the state will hold with a desirably high probability. The notation Φ and χ_n stand for the cumulative distribution functions of standard normal and n-degree chi-squared variables, respectively, and the notation $\|\cdot\|_{\boldsymbol{A}}$ will represent the the Euclidean norm induced under the matrix $\boldsymbol{A} \succeq \boldsymbol{0}$, i.e., $\|\boldsymbol{x}\|_{\boldsymbol{A}} = \sqrt{\boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x}}$.

Theorem 6.4.2. *Consider a linear system described by Eq. (6.1), with the state written as*

$$\boldsymbol{x} = \bar{\boldsymbol{A}}\boldsymbol{x}_0 + \bar{\boldsymbol{B}}\boldsymbol{u} + \bar{\boldsymbol{C}}\boldsymbol{w},$$

for appropriate matrices $\bar{\boldsymbol{A}}$, $\bar{\boldsymbol{B}}$, and $\bar{\boldsymbol{C}}$. Then under Assumption 6.4.1, we have the following:

[4]Of course, other distributional assumptions may be made: we present the normality assumption primarily because (a) it provides the cleanest analytical results and (b) it is the most common assumption in the literature.

(a) *The constraint*

$$\boldsymbol{p}^T \boldsymbol{u} \leq q, \tag{6.38}$$

where $\boldsymbol{p} = \bar{\boldsymbol{B}}^T \boldsymbol{g}$, $q = \hat{g} - \boldsymbol{g}^T \bar{\boldsymbol{A}} \boldsymbol{x}_0 - \|\bar{\boldsymbol{C}}^T \boldsymbol{g}\|_2 \Phi^{-1}(1-\epsilon)$ *implies the following guarantee:*

$$\mathbb{P}\left\{\boldsymbol{g}^T \boldsymbol{x} > \hat{g}\right\} \leq \epsilon. \tag{6.39}$$

(b) *The constraint*

$$\|\bar{\boldsymbol{A}} \boldsymbol{x}_0 + \bar{\boldsymbol{B}} \boldsymbol{u}\|_{\boldsymbol{H}} \leq r, \tag{6.40}$$

where $r = 1 - \|\bar{\boldsymbol{C}}^T \boldsymbol{H} \bar{\boldsymbol{C}}\|_2 \chi^{-1}_{Nk}(1-\epsilon)$ *implies the following guarantee:*

$$\mathbb{P}\left\{\boldsymbol{x}^T \boldsymbol{H} \boldsymbol{x} > 1\right\} \leq \epsilon, \tag{6.41}$$

where $\boldsymbol{H} \succeq \boldsymbol{0}$.

Proof.

(a) We have

$$\begin{aligned} \mathbb{P}\left\{\boldsymbol{g}^T \boldsymbol{x} > \hat{g}\right\} &= \mathbb{P}\left\{\boldsymbol{g}^T \bar{\boldsymbol{C}} \boldsymbol{w} > \hat{g} - \boldsymbol{g}^T \bar{\boldsymbol{A}} \boldsymbol{x}_0 - \boldsymbol{g}^T \bar{\boldsymbol{B}} \boldsymbol{u}\right\} \\ &= 1 - \Phi\left(\frac{\hat{g} - \boldsymbol{g}^T \bar{\boldsymbol{A}} \boldsymbol{x}_0 - \boldsymbol{g}^T \bar{\boldsymbol{B}} \boldsymbol{u}}{\|\bar{\boldsymbol{C}}^T \boldsymbol{g}\|_2}\right), \end{aligned}$$

and the result follows by setting this less than or equal to ϵ and inverting Φ.

(b) We have

$$\begin{aligned}
\mathbb{P}\left\{\boldsymbol{x}^T\boldsymbol{H}\boldsymbol{x} > 1\right\} &= \mathbb{P}\left\{\|\bar{\boldsymbol{A}}\boldsymbol{x}_0 + \bar{\boldsymbol{B}}\boldsymbol{u} + \bar{\boldsymbol{C}}\boldsymbol{w}\|_{\boldsymbol{H}} > 1\right\} \\
&\leq \mathbb{P}\left\{\|\bar{\boldsymbol{A}}\boldsymbol{x}_0 + \bar{\boldsymbol{B}}\boldsymbol{u}\|_{\boldsymbol{H}} + \|\bar{\boldsymbol{C}}\boldsymbol{w}\|_{\boldsymbol{H}} > 1\right\} \\
&\leq \mathbb{P}\left\{\|\bar{\boldsymbol{C}}^T\boldsymbol{H}\bar{\boldsymbol{C}}\|_2\|\boldsymbol{w}\|_2 > 1 - \|\bar{\boldsymbol{A}}\boldsymbol{x}_0 + \bar{\boldsymbol{B}}\boldsymbol{u}\|_{\boldsymbol{H}}\right\} \\
&= 1 - \chi_{Nk}\left(\frac{1 - \|\bar{\boldsymbol{A}}\boldsymbol{x}_0 + \bar{\boldsymbol{B}}\boldsymbol{u}\|_{\boldsymbol{H}}}{\|\bar{\boldsymbol{C}}^T\boldsymbol{H}\bar{\boldsymbol{C}}\|_2}\right),
\end{aligned}$$

and, again, the result follows by setting this less than or equal to ϵ and inverting χ_{Nk}. □

Both parts of Theorem 6.4.2 are constraints which can be added to either the SDP or SOCP approaches without increasing their respective complexity types. Note that the constraint in part (a) of the theorem is exact, while the constraint in (b) is somewhat conservative. Care must be taken to ensure that $\boldsymbol{H}$ and ϵ do not result in a constraint which forces the problem to be infeasible.

6.4.3 Probabilistic performance guarantees

In this section, we analyze the probability distribution of the cost-to-go function. We first derive a bound on the performance distribution of the cost-to-go function for a given control policy $\boldsymbol{y}$ under Assumption 6.4.1, then describe the protection guarantees and expected losses for both problems (6.17) and (6.35). Finally, we show how to probabilistically ensure certain levels of performance.

We emphasize that the results proven here in terms of performance guarantees are for *open-loop* control. In general, analyzing our approach in a feedback context seems difficult. Instead, we will study the closed-loop performance computationally in the Section 6.6.

For a given policy $\boldsymbol{y}$, the cost function is a random variable (a function of the random disturbances):

$$\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) = \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} + 2(\boldsymbol{h} + \boldsymbol{F}^T\boldsymbol{y})^T\boldsymbol{w} + \boldsymbol{y}^T\boldsymbol{y}.$$

For our results in this section, we need a slightly stronger assumption.

Assumption 6.4.2. In addition to Assumption 6.4.1, we have $\boldsymbol{C} \succ \mathbf{0}$.

We see that under Assumption 6.4.2, we have

$$\mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w})\right] = \|\boldsymbol{y}\|_2^2 + \mathrm{Tr}(\boldsymbol{C}),$$

with $\mathrm{Tr}(\boldsymbol{C}) = \sum_{i=1}^{N \cdot n_w} \lambda_i$, where λ_i are the eigenvalues of C. We now derive a key result; the proof is quite similar to a proof from Bertsimas and Sim [33].

Proposition 6.4.2. Under Assumption 6.4.2, we have

$$\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) > z\right\} \leq c_\rho \exp\left(-\frac{1}{2\mathrm{Tr}(\boldsymbol{C})}\left[z - \|\boldsymbol{y}\|_2^2 - \frac{1}{\rho - 1}\|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}}^2\right]\right), \quad (6.42)$$

where

$$c_\rho = \left(\frac{\rho}{\rho - 1}\right)^{\rho/2}, \quad (6.43)$$

$$\rho = \frac{\mathrm{Tr}(\boldsymbol{C})}{\max\limits_{i=1,\ldots,N \cdot n_w} \lambda_i} > 1, \quad (6.44)$$

$$\boldsymbol{g} = \boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y}. \quad (6.45)$$

Proof. Let $\lambda_i > 0$, $i = 1, \ldots, N \cdot n_w$, be the eigenvalues of $\boldsymbol{C} \succ \mathbf{0}$, and $\boldsymbol{Q}^T \boldsymbol{\Lambda} \boldsymbol{Q}$ be its eigenvalue decomposition. We then have

$$\begin{aligned}
\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) > z\right\} &= \mathbb{P}\left\{\boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w} + 2(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})^T \boldsymbol{w} > z - \|\boldsymbol{y}\|_2^2\right\} \\
&= \mathbb{P}\left\{\sum_{i=1}^{N \cdot n_w} \lambda_i \left(v_i + f_i/\lambda_i\right)^2 > z - \|\boldsymbol{y}\|_2^2 + \|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}}^2\right\} \\
&= \mathbb{P}\left\{\sum_{i=1}^{N \cdot n_w} \lambda_i u_i^2 > \tilde{z}\right\},
\end{aligned}$$

where we employ the transformations $\boldsymbol{v} = \boldsymbol{Q}\boldsymbol{w}$, $\boldsymbol{f} = \boldsymbol{Q}(\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y})$, and we have $u_i \sim \mathcal{N}(f_i/\lambda_i, 1)$, independent. Finally, $\tilde{z} = z - \|\boldsymbol{y}\|_2^2 + \|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}}^2$ for notational convenience.

Continuing, we have

$$\begin{aligned}
\mathbb{P}\left\{\sum_{i=1}^{N\cdot n_w} \lambda_i u_i^2 > \tilde{z}\right\} &\leq \frac{\mathbb{E}\left[\exp\left(\theta \sum_{i=1}^{N\cdot n_w} \lambda_i u_i^2\right)\right]}{\exp(\theta\tilde{z})} \\
&= \frac{\prod_{i=1}^{N\cdot n_w} \mathbb{E}\left[\exp\left(\theta\lambda_i u_i^2\right)\right]}{\exp(\theta\tilde{z})} \\
&= \frac{\prod_{i=1}^{N\cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)^{\beta\theta\lambda_i}\right]}{\exp(\theta\tilde{z})} \\
&\leq \frac{\prod_{i=1}^{N\cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)\right]^{\beta\theta\lambda_i}}{\exp(\theta\tilde{z})},
\end{aligned}$$

where we require $\theta > 0$, $\beta\theta\lambda_i \leq 1$, and $\beta > 2$. The first line above follows from the Markov inequality, the second follows from independence of the u_i, and the last line follows from Jensen's inequality. Now noting that under $u_i \sim \mathcal{N}(f_i/\lambda_i, 1)$, we have

$$\mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)\right] = \sqrt{\frac{\beta}{\beta-2}}\exp\left(\frac{1}{\beta-2}\left(f_i/\lambda_i\right)^2\right),$$

we have, after some rearranging with $\theta = 1/(\beta \max_{i=1,\ldots,N\cdot n_w} \lambda_i)$,

$$\begin{aligned}
\frac{\prod_{i=1}^{N\cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)\right]^{\beta\theta\lambda_i}}{\exp(\theta\tilde{z})} &= \frac{\prod_{i=1}^{N\cdot n_w}\left(\sqrt{\frac{\beta}{\beta-2}}\exp\left(\frac{1}{\beta-2}\left(f_i/\lambda_i\right)^2\right)\right)^{\beta\theta\lambda_i}}{\exp(\theta\tilde{z})} \\
&= \frac{\left(\frac{\beta}{\beta-2}\right)^{\rho/2}\exp\left(\frac{2\rho}{\beta(\beta-2)}\frac{\|\boldsymbol{g}\|^2_{\boldsymbol{C}^{-1}}}{\mathrm{Tr}(\boldsymbol{C})}\right)}{\exp\left(\frac{\rho}{\beta}\left[\frac{z-\|\boldsymbol{y}\|_2^2}{\mathrm{Tr}(\boldsymbol{C})}\right]\right)}.
\end{aligned}$$

Finally, we set $\beta = 2\rho > 2$ and the result follows. □

Note that the bound in Equation (6.42) is for *any* policy $\boldsymbol{y}$, and we have in no way imposed the structure of either of the approaches developed in Sections 6.2 or 6.3. Now utilizing the structure of the SDP and the SOCP, we may quantify more precisely what we gain in terms of performance protection with both approaches.

This protection comes at the price of some degradation of expected performance, and we also quantify this decrease.

Theorem 6.4.3. *Under Assumption 6.4.2, and with* $\mathrm{E}[\tilde{J}(\boldsymbol{x}_0, \mathbf{0}, \boldsymbol{w})]$ *the expected value of the Riccati approach, we have the following:*

(a) *If* $(\boldsymbol{y}^*_{\mathrm{SDP}}, z^*_{\mathrm{SDP}})$ *are feasible in SDP (6.17), then the expected performance loss is bounded as*

$$\mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SDP}}, \boldsymbol{w})\right] - \mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \mathbf{0}, \boldsymbol{w})\right] \;\leq\; 2\|\boldsymbol{h}\|_2\gamma, \tag{6.46}$$

while we gain the following level of probabilistic protection:

$$\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SDP}}, \boldsymbol{w}) > z^*_{\mathrm{SDP}}\right\} \;\leq\; \alpha \exp\left(-\frac{\rho}{2}\gamma^2\right), \tag{6.47}$$

where

$$\begin{aligned} \alpha &= c_\rho \exp\left(\frac{1}{2(\rho-1)\mathrm{Tr}(\boldsymbol{C})}\delta^2\right) \\ \delta &= \|\boldsymbol{C}^{-1}\|_2\left[\|\boldsymbol{h}\|_2 + \|\boldsymbol{F}\boldsymbol{F}^T\|_2\sqrt{2\|\boldsymbol{h}\|_2\gamma}\right]. \end{aligned}$$

(b) *If* $(\boldsymbol{y}^*_{\mathrm{SOCP}}, z^*_{\mathrm{SOCP}})$ *are feasible in SOCP (6.35), then the expected performance loss is bounded as*

$$\mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SOCP}}, \boldsymbol{w})\right] - \mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \mathbf{0}, \boldsymbol{w})\right] \;\leq\; 4\|\boldsymbol{h}\|_2\Omega, \tag{6.48}$$

while we gain the following level of probabilistic protection:

$$\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SOCP}}, \boldsymbol{w}) > z^*_{\mathrm{SOCP}}\right\} \;\leq\; \sqrt{\frac{e}{2}}\Omega \exp\left(-\frac{\Omega^2}{4}\right). \tag{6.49}$$

Proof.

(a) For the expected loss from the Riccati approach, note that the optimal Riccati solution, from Corollary 6.2.1, is $\boldsymbol{y} = \mathbf{0}$, and that the expected performance of *any* policy is just given by

$$\mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w})\right] = \|\boldsymbol{y}\|_2^2 + \mathrm{Tr}(\boldsymbol{C}).$$

Therefore, noting that feasibility of $(\boldsymbol{y}^*_{\mathrm{SDP}}, z^*_{\mathrm{SDP}})$ in SDP (6.17) requires

$$\|\boldsymbol{y}^*_{\mathrm{SDP}}\|_2^2 \leq z^*_{\mathrm{SDP}} - \|\boldsymbol{C}\|_2^2\gamma^2,$$

we have

$$\begin{aligned}
\mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SDP}}, \boldsymbol{w})\right] - \mathbb{E}\left[\tilde{J}(\boldsymbol{x}_0, \boldsymbol{0}, \boldsymbol{w})\right] &= \|\boldsymbol{y}^*_{\mathrm{SDP}}\|_2^2 \\
&\leq z^*_{\mathrm{SDP}} - \|\boldsymbol{C}\|_2^2\gamma^2 \\
&= \min_{\boldsymbol{y}}\left[\|\boldsymbol{y}\|_2^2 + \max_{\boldsymbol{w}\in\mathcal{W}_\gamma} \boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} + 2(\boldsymbol{h} + \boldsymbol{F}^T\boldsymbol{y})^T\boldsymbol{w}\right] \\
&\quad -\|\boldsymbol{C}\|_2^2\gamma^2 \\
&\leq \max_{\boldsymbol{w}\in\mathcal{W}_\gamma}\left[\boldsymbol{w}^T\boldsymbol{C}\boldsymbol{w} + 2(\boldsymbol{h} + \boldsymbol{F}^T\boldsymbol{y})^T\boldsymbol{w}\right] - \|\boldsymbol{C}\|_2^2\gamma^2 \\
&\leq \|\boldsymbol{C}\|_2^2\gamma^2 + 2\|\boldsymbol{h}\|_2\gamma - \|\boldsymbol{C}\|_2^2\gamma^2 \\
&= 2\|\boldsymbol{h}\|_2\gamma.
\end{aligned}$$

The first line follows from feasibility of $(\boldsymbol{y}^*_{\mathrm{SDP}}, z^*_{\mathrm{SDP}})$ as stated above, the next line follows from the definition of z^*_{SDP}, the next line follows by noting that $\boldsymbol{y} = \boldsymbol{0}$ is feasible in the SDP, and the second-to-last line follows by bounding the maximum value of the given function over all $\boldsymbol{w} \in \mathcal{W}_\gamma$.

For the probabilistic guarantee of Eq. (6.47), we note, from Proposition 6.4.2 and feasibility of $(\boldsymbol{y}^*_{\mathrm{SDP}}, z^*_{\mathrm{SDP}})$ that

$$\begin{aligned}
\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}^*_{\mathrm{SDP}}, \boldsymbol{w}) > z^*_{\mathrm{SDP}}\right\} &\leq c_\rho \exp\left(-\frac{1}{2\mathrm{Tr}(\boldsymbol{C})}\left[z^*_{\mathrm{SDP}} - \|\boldsymbol{y}^*_{\mathrm{SDP}}\|_2^2 - \frac{1}{\rho-1}\|\boldsymbol{g}\|^2_{\boldsymbol{C}^{-1}}\right]\right) \\
&\leq c_\rho \exp\left(-\frac{1}{2\mathrm{Tr}(\boldsymbol{C})}\left[\|\boldsymbol{C}\|_2^2\gamma^2 - \frac{1}{\rho-1}\|\boldsymbol{g}\|^2_{\boldsymbol{C}^{-1}}\right]\right) \\
&= c_\rho \exp\left(\frac{1}{2(\rho-1)\mathrm{Tr}(\boldsymbol{C})}\|\boldsymbol{g}\|^2_{\boldsymbol{C}^{-1}}\right)\exp\left(-\frac{\rho}{2}\gamma^2\right).
\end{aligned}$$

Now, we have

$$
\begin{aligned}
\|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}} &= \|\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y}^*_{\text{SDP}}\|_{\boldsymbol{C}^{-1}} \\
&\leq \|\boldsymbol{C}^{-1}\|_2 \|\boldsymbol{h} + \boldsymbol{F}^T \boldsymbol{y}^*_{\text{SDP}}\|_2 \\
&\leq \|\boldsymbol{C}^{-1}\|_2 \left[\|\boldsymbol{h}\|_2 + \|\boldsymbol{F}^T \boldsymbol{y}^*_{\text{SDP}}\|_2\right] \\
&\leq \|\boldsymbol{C}^{-1}\|_2 \left[\|\boldsymbol{h}\|_2 + \|\boldsymbol{F}\boldsymbol{F}^T\|_2 \|\boldsymbol{y}^*_{\text{SDP}}\|_2\right] \\
&\leq \|\boldsymbol{C}^{-1}\|_2 \left[\|\boldsymbol{h}\|_2 + \|\boldsymbol{F}\boldsymbol{F}^T\|_2 \sqrt{2\|\boldsymbol{h}\|_2 \gamma}\right],
\end{aligned}
$$

where we have repeatedly used matrix norm bounds, the Schwartz inequality, and, in the last line, utilized the bound for $\|\boldsymbol{y}^*_{\text{SDP}}\|_2^2$ derived in the expected loss above. The result now follows.

(b) The expected loss for the SOCP follows exactly analogously to the proof for the expected loss for the SDP in part (a) by noting that feasibility of $(\boldsymbol{y}^*_{\text{SOCP}}, z^*_{\text{SOCP}})$ in the SOCP implies feasibility in the SDP (Proposition 6.3.1), replacing γ with Ω, and then noting the result of Corollary 6.3.1 (namely, $z^*_{\text{SOCP}} - z^*_{\text{SDP}} \leq 2\Omega\|\boldsymbol{h}\|_2$). The probabilistic bound follows by directly applying Theorem 6.1.2. □

Theorem 6.4.3 quantifies the expected loss and a probabilistic protection level for both approaches. Note that the expected loss from the Riccati equation can be bounded by a quantity linear in the size of the uncertainty set (γ or Ω). Moreover, the protection level bounds are both of a similar nature ($\mathcal{O}\left(\gamma \exp(-\gamma^2)\right)$ and $\mathcal{O}\left(\Omega \exp(-\Omega^2)\right)$).

In addition to simply describing performance, we may also want to explicitly protect against certain threshold performance levels. We now show how to do this.

Theorem 6.4.4. *Under Assumption 6.4.2, the convex quadratic constraint*

$$\boldsymbol{y}^T \boldsymbol{P} \boldsymbol{y} + 2\boldsymbol{q}^T \boldsymbol{y} \leq r, \tag{6.50}$$

where

$$\begin{aligned}
\boldsymbol{P} &= \boldsymbol{I} + \frac{1}{\rho - 1}\boldsymbol{F}\boldsymbol{C}^{-1}\boldsymbol{F}^T \succ \boldsymbol{0}, &(6.51)\\
\boldsymbol{q} &= \frac{1}{\rho - 1}\boldsymbol{F}\boldsymbol{C}^{-1}\boldsymbol{h}, &(6.52)\\
r &= z + 2\mathrm{Tr}(\boldsymbol{C})\ln(\epsilon/c_\rho) - \frac{1}{\rho - 1}\boldsymbol{h}^T\boldsymbol{C}^{-1}\boldsymbol{h}, &(6.53)
\end{aligned}$$

implies the following guarantee:

$$\mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) > z\right\} \leq \epsilon. \tag{6.54}$$

Proof. The proof follows directly from Proposition 6.4.2 in straightforward fashion, noting the implications

$$\begin{aligned}
\boldsymbol{y}^T\boldsymbol{P}\boldsymbol{y} + 2\boldsymbol{q}^T\boldsymbol{y} \leq r \;&\Rightarrow\; \|\boldsymbol{y}\|_2^2 + \frac{1}{\rho - 1}\|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}}^2 \leq z + 2\mathrm{Tr}(\boldsymbol{C})\ln(\epsilon/c_\rho)\\
&\Rightarrow\; c_\rho \exp\left(-\frac{1}{2\mathrm{Tr}(\boldsymbol{C})}\left[z - \|\boldsymbol{y}\|_2^2 - \frac{1}{\rho - 1}\|\boldsymbol{g}\|_{\boldsymbol{C}^{-1}}^2\right]\right) \leq \epsilon\\
&\Rightarrow\; \mathbb{P}\left\{\tilde{J}(\boldsymbol{x}_0, \boldsymbol{y}, \boldsymbol{w}) > z\right\} \leq \epsilon,
\end{aligned}$$

where $\boldsymbol{P}$, $\boldsymbol{q}$, and r are defined in Eqs. (6.51)-(6.53). Positive definiteness of P follows from positive definiteness of $\boldsymbol{I}$ and $\boldsymbol{C}$, and thus the constraint is a convex quadratic constraint. □

We see that (6.50) is a convex quadratic constraint and hence may be added (in the same manner as in the proof of Theorem 6.4.1) to either approach without increasing their respective complexity types. Note that we may only ensure against appreciably high levels of cost. In fact, a simple necessary (but not sufficient) condition to retain feasibility of the problem is the requirement

$$z \;\geq\; -2\mathrm{Tr}(\boldsymbol{C})\ln(\epsilon/c_\rho).$$

6.5 Imperfect state information

In some cases, we may not know the state $\boldsymbol{x}_0$ of the system exactly. Rather, we may only have an estimate $\hat{\boldsymbol{x}}_0$ of the current state. In standard dynamic programming texts (e.g., Bertsekas [26]), it is shown that in the case where noise-corrupted state observations of the form

$$\boldsymbol{v}_k \;=\; \boldsymbol{H}_k \boldsymbol{x}_k + \boldsymbol{\delta}_k$$

are available, with $\boldsymbol{H}_k$ known matrices and $\boldsymbol{\delta}_k$ additive noise with finite second moment, then the resulting optimal policy is a modified Riccati equation. Here, we assume the following model for the state estimate $\hat{\boldsymbol{x}}_0$:

$$\hat{\boldsymbol{x}}_0 \;=\; \boldsymbol{x}_0 + \boldsymbol{\eta}, \tag{6.55}$$

where $\boldsymbol{\eta}$ is a noise term with some distribution. We will show that the form of the cost-to-go function is unchanged by the added uncertainty in the state by now viewing the disturbances as $\hat{\boldsymbol{w}} = \begin{bmatrix}\boldsymbol{\eta}^T & \boldsymbol{w}^T\end{bmatrix}^T$. As a consequence, we can apply either of the robust approaches to the problem with imperfect state information of the form given in Eq. (6.55).

Proposition 6.5.1. With noisy estimates of the state given by Eq. (6.55), the cost-to-go can be written in the form

$$\hat{J}_{\boldsymbol{u}}(\hat{\boldsymbol{x}}_0) = \hat{\boldsymbol{x}}_0^T \hat{\boldsymbol{A}} \hat{\boldsymbol{x}}_0 + 2\hat{\boldsymbol{a}}^T \hat{\boldsymbol{x}}_0 + 2\hat{\boldsymbol{b}}^T \boldsymbol{u} + \boldsymbol{u}^T \boldsymbol{B} \boldsymbol{u} + 2\hat{\boldsymbol{w}}^T \hat{\boldsymbol{c}} + \hat{\boldsymbol{w}}^T \hat{\boldsymbol{C}} \hat{\boldsymbol{w}} + 2\boldsymbol{u}^T \hat{\boldsymbol{D}} \hat{\boldsymbol{w}}, \tag{6.56}$$

where $\hat{\boldsymbol{A}}$, $\hat{\boldsymbol{a}}$, and $\boldsymbol{B}$ are as in Proposition 6.2.1. $\hat{\boldsymbol{b}}$ is as in Proposition 6.2.1 with $\hat{\boldsymbol{x}}_0$

replacing $\boldsymbol{x}_0$, and

$$\begin{aligned}
\hat{\boldsymbol{c}} &= \begin{bmatrix} -\hat{\boldsymbol{a}} - \hat{\boldsymbol{A}}\hat{\boldsymbol{x}}_0 \\ \left(\sum_{k=0}^{N-1} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1}\right) \hat{\boldsymbol{x}}_0 + \sum_{k=0}^{N-1} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{q}_k \end{bmatrix} \\
\hat{\boldsymbol{C}} &= \begin{bmatrix} \hat{\boldsymbol{A}} & -\sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \\ -\sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} & \boldsymbol{C} \end{bmatrix} \\
\hat{\boldsymbol{D}} &= \begin{bmatrix} -\sum_{k=1}^{N} \left(\boldsymbol{R}_{k+1} + \tilde{\boldsymbol{B}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{B}}_{k-1}\right) \\ D \end{bmatrix}
\end{aligned}$$

Furthermore, the matrix $\hat{\boldsymbol{C}}$ is positive semi-definite.

Proof. The proof follows by recalling the original, perfect state cost form of

$$J_{\boldsymbol{u}}(\boldsymbol{x}_0) = 2\hat{\boldsymbol{a}}^T \boldsymbol{x}_0 + \boldsymbol{x}_0^T \hat{\boldsymbol{A}} \boldsymbol{x}_0 + 2\boldsymbol{b}^T \boldsymbol{u} + \boldsymbol{u}^T \boldsymbol{B} \boldsymbol{u} + 2\boldsymbol{c}^T \boldsymbol{w} + \boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w} + 2\boldsymbol{u}^T \boldsymbol{D} \boldsymbol{w}$$

(Eq. (6.14)). Simply substituting $\boldsymbol{x}_0 = \hat{\boldsymbol{x}}_0 - \boldsymbol{\eta}$ and collecting terms, Eq. (6.56) follows. To see that $\hat{\boldsymbol{C}} \succeq \boldsymbol{0}$, note, using the original definitions of $\hat{\boldsymbol{A}}$ and $\boldsymbol{C}$, we have

$$\begin{aligned}
\hat{\boldsymbol{C}} &= \begin{bmatrix} \sum_{k=1}^{N} \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} & -\sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \\ -\sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} & \sum_{k=1}^{N} \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \end{bmatrix} \\
&= \sum_{k=1}^{N} \begin{bmatrix} \tilde{\boldsymbol{A}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{A}}_{k-1} & -\tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \\ -\tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} & \tilde{\boldsymbol{C}}_{k-1}^T \boldsymbol{Q}_k \tilde{\boldsymbol{C}}_{k-1} \end{bmatrix} \\
&= \sum_{k=1}^{N} \left[\tilde{\boldsymbol{A}}_{k-1} - \tilde{\boldsymbol{C}}_{k-1}\right]^T \boldsymbol{Q}_k \left[\tilde{\boldsymbol{A}}_{k-1} - \tilde{\boldsymbol{C}}_{k-1}\right] \\
&\succeq \boldsymbol{0},
\end{aligned}$$

where the last line follows, since it is a sum of similarity transformations of positive definite matrices ($\boldsymbol{Q}_k \succ \boldsymbol{0}$). □

6.6 Computational results

We have written routines for solving the control law of Section 6.2.1 as well as SOCP (6.35) in Section 6.3. The routines have been implemented in a Matlab environment and the SeDuMi (Sturm, [108]) package has been used for the underlying optimization problems. In this section, we explore computationally the performance of our approach in closed-loop control in a variety of ways.

6.6.1 Performance in the unconstrained case

Here we compare the performance of optimal policy in Theorem 6.2.2 to that given by the Riccati equation for a problem without constraints. We considered a simple, 10-stage problem with time-invariant state, control, and disturbance matrices $\boldsymbol{A} = 1$, $\boldsymbol{B} = 1$, and $\boldsymbol{C} = 1$, initial state $\boldsymbol{x}_0 = -1$. The cost function was given by $\boldsymbol{Q}_k = (1/2)^k$, $\boldsymbol{R}_k = (1/2)^k$, and $\boldsymbol{q}_k = \boldsymbol{r}_k = 0$ for all $k \in \{0, \dots, 10\}$.

We ran 1000 trials of the closed-loop policies for the Riccati approach and the control law of Theorem 6.2.2 and tabulated the average percentage increase in cost (over Riccati) for various values of γ. Disturbances vectors were generated at each iteration by $\mathcal{N}(0, \sigma^2)$, where σ is a parameter we varied. Table 6.2 lists the results, which are also illustrated graphically in Figure 6-2.

We observe the following from these computational results:

1. For γ small, this approach does not result in a marked increase in expected cost.

2. For γ beyond a certain value, the expected cost increase does not change. This is not surprising, since for γ large enough, we have $\gamma > \gamma_{\text{thresh}}$ (Proposition 6.2.5, (b)) at each iteration. In this case, there will be no change in the policy given by Theorem 6.2.2 for further increases in γ, so the performance, on average, will not change.

3. The policies given by Theorem 6.2.2 are more conservative than the traditional, Riccati approach. As such, the distribution in the cost is more stable for larger γ. Here, we measure stability in terms of the standard deviation of the first

Average relative cost increase (%)						
γ		.001	.01	.1	1	10
σ	.01	.00	.05	4.0	51	51
	.1	.00	.13	4.8	54	54
	1	.04	.45	5.0	48	49
	10	.00	.02	.20	2.8	19
Stability increase (%)						
γ		.001	.01	.1	1	10
σ	.01	.16	1.5	17	110	120
	.10	.13	1.5	15	80	86
	1	.06	.57	5.1	15	19
	10	.00	.01	.13	.41	7.7

Table 6.2: *Average relative cost increase (top) and stability increase (bottom) for our approach versus Riccati for various γ and disturbance distributions. All numbers in %. Stability is measured here by the standard deviation of the first upper tail moment.*

upper tail moment. Specifically, if we denote the cost distribution under a control policy π by the random variable J_π, then the stability results reported are

$$\text{Stability}_\pi = \sigma\left(\mathbb{E}\left[\max(J_\pi - \mathbb{E}\left[J_\pi\right], 0)\right]\right).$$

We chose the upper tail moment because downside variability in the cost is potentially beneficial. The policies given by Theorem 6.2.2 give, for the most part, significantly more stable policies than the Riccati policy. Thus, by varying γ, there is a tradeoff between expected value of the cost and *variability* of the cost.

4. Although we only report results for a $N = 10$ stage problem here, the results are similar for other dimensions and other problem instances.

6.6.2 SOCP performance in the unconstrained case

The SOCP approach from Section 6.3 has been developed for use in the constrained case (i.e., when we cannot use the control law given in Theorem 6.2.2). As it is

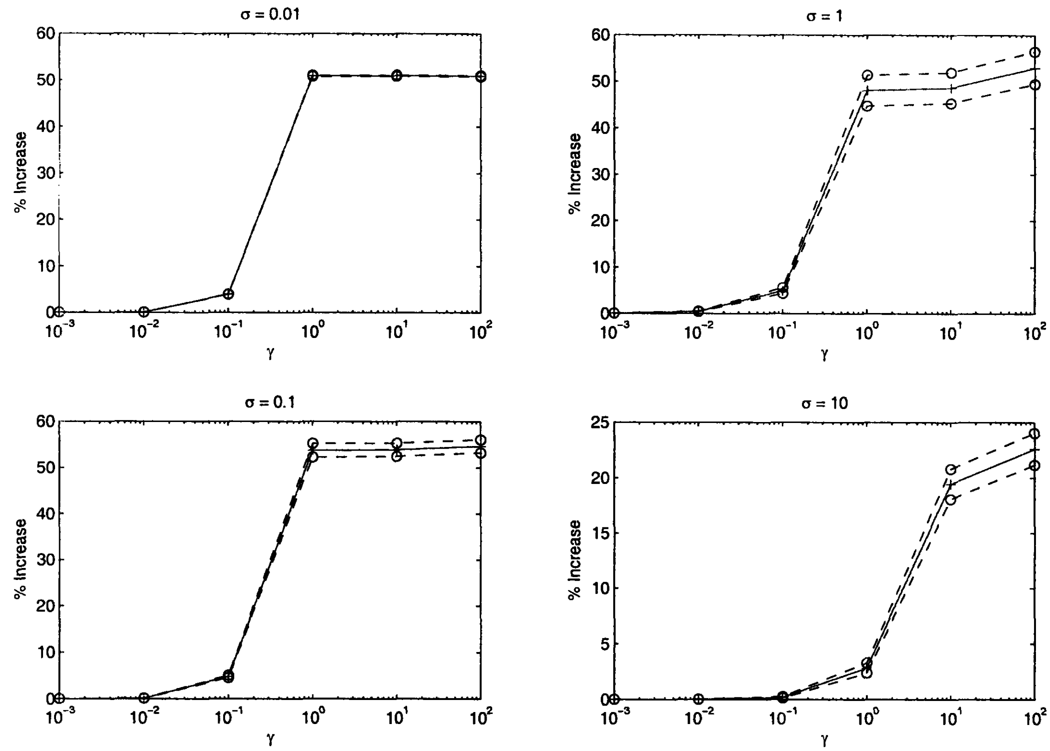

Figure 6-2: *Expected % increase in total cost from Riccati under various distributions. The dashed lines are 95% confidence intervals for the expected % increase.*

an approximation to SDP (6.17), however, it is relevant to examine how good this approximation is when we have no constraints. We ran 100 trials of the problem instance from the previous section and compared the performance of SOCP (6.35) to the control law in Theorem (6.2.2), with both approaches run closed-loop. The results are shown in Figure 6-3, again versus γ and under various disturbance distributions.

Note that the shape of the expected cost increase for SOCP (6.35) versus γ is essentially the same as that for the control law of Theorem 6.2.2, just with a higher asymptote for large γ. For $\gamma \ll 1$, however, the performance of the SOCP approximation is essentially indistinguishable from that of the control law in Theorem 6.2.2.

6.6.3 Effect of constraints on runtime

Here we present resulting run times for various values of N for a problem with $n_x = n_u = n_w = 1$, $\boldsymbol{A}_k = \boldsymbol{B}_k = 1$, $\boldsymbol{C}_k = 1/(2N)$, $\boldsymbol{Q}_k = 1$, $\boldsymbol{q}_k = 0$, $\boldsymbol{R}_k = 0$, $\boldsymbol{r}_k = 0$, and

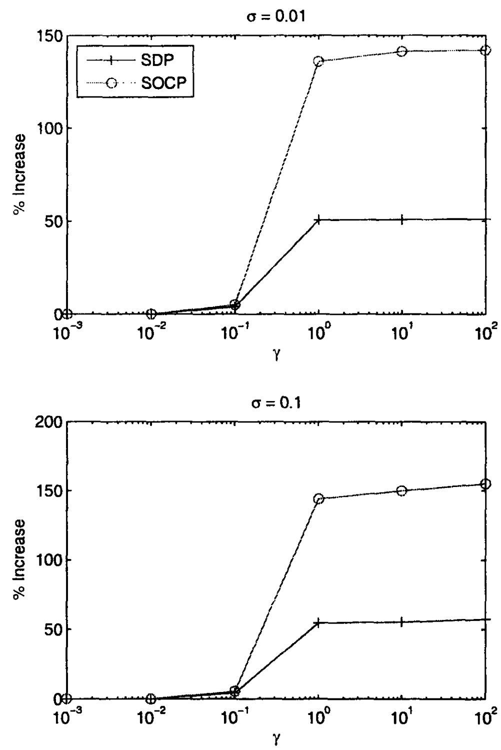

Figure 6-3: *Expected % increase in total cost from Riccati for both the control law of Theorem 6.2.2 (+) and SOCP (6.35) (o).*

$\boldsymbol{x}_0 = 1$. For all trials we use $\Omega = 0$. The machine is running Linux on a 2.2GHz processor with 1.0GB RAM. The results are listed in Table 6.3. Note that in this case we are solving the problem in open-loop fashion, i.e., we are computing a single control vector in each case and calculating the runtime to do so.

We note the following from the computational results:

1. The presence of any of the constraints listed does not result in marked increases in run time for fixed N.

2. For the objective guarantee, we use $\epsilon = 0.4$ for $N = 2$, $\epsilon = 0.2$ for $N = 5, 10, 20$, and $\epsilon = 0.1$ for all other N. In the unconstrained case, we do not have a very good guarantee; in fact, we only know

$$\mathbb{P}\left\{J_{\boldsymbol{u}_0}(\boldsymbol{x}_0) > 2\mathrm{E}[J_{\boldsymbol{u}_0}(\boldsymbol{x}_0)]\right\} \leq 1,$$

Run Time (sec.)						
N	2	5	10	20	50	100
No constraints	0.06	0.05	0.06	0.07	0.12	0.47
$\mathbb{P}\{J_{\boldsymbol{u}_0}(\boldsymbol{x}_0) > 2\mathrm{E}[J_{\boldsymbol{u}_0}(\boldsymbol{x}_0)]\} \leq \epsilon$	0.13	0.13	0.16	0.20	0.44	1.87
$\mathbb{P}\{x_i < 0\} \leq .01$	0.09	0.11	0.14	0.14	0.40	1.54
$\boldsymbol{u} \geq 0$	0.08	0.08	0.11	0.15	0.38	1.57
$\mathbb{P}\{x_i < 0\} \leq .01$ & $\boldsymbol{u} \geq 0$	0.07	0.09	0.11	0.16	0.38	1.55
Increase from Unconstrained, Expected Cost Per Stage (%)						
N	2	5	10	20	50	100
$\mathbb{P}\{J_{\boldsymbol{u}_0}(\boldsymbol{x}_0) > 2\mathrm{E}[J_{\boldsymbol{u}_0}(\boldsymbol{x}_0)]\} \leq \epsilon$	16.96	2.47	0.79	0.27	0.67	1.35
$\mathbb{P}\{x_i < 0\} \leq .01$	61.00	78.58	86.24	88.53	89.93	91.89
$\boldsymbol{u} \geq 0$	293.93	1909.38	7261.24	27993.9	1.71E5	6.82E5
$\mathbb{P}\{x_i < 0\} \leq .01$ & $\boldsymbol{u} \geq 0$	293.93	1909.38	7261.24	27993.9	1.71E5	6.82E5

Table 6.3: *Run time in seconds and cost increase from unconstrained for the SOCP approach with various constraints.*

where $\boldsymbol{u}_0$ is the optimal policy in that case. As we can see, however, without significant increase in expected cost we are able to ensure that this bad event occurs with probability no greater than ϵ.

3. The constraints $\boldsymbol{u} \geq 0$ result in very large increases in expected cost. This is merely due to the fact that it is a very restrictive restraint and not a drawback of the proposed approaches.

4. Although we do not report the run-times here, we ran this simulation for the SDP with the listed constraints as well. Typically this runs $\mathcal{O}(N)$ longer than the SOCP, solidifying our assertion that the SOCP is much more suitable to efficient, closed-loop control.

6.6.4 Performance on a problem with constraints

Here we compared the performance of SOCP (6.35) versus the optimal policy for a 5-stage problem with the constraints $u_k \geq 0$ for $k \in \{0, \ldots, 4\}$. The problem data were $\boldsymbol{A}_k = 1$, $\boldsymbol{B}_k = 1$, and $\boldsymbol{C}_k = 1$ for $k = \{0, \ldots, 4\}$, initial state $\boldsymbol{x}_0 = -.6$. The cost function was given by $\boldsymbol{Q}_k = \beta^k$, $\boldsymbol{R}_k = \beta^k$, and $\boldsymbol{q}_k = 0$, $\boldsymbol{r}_k = -\beta^k$ for all $k \in \{0, \ldots, 4\}$.

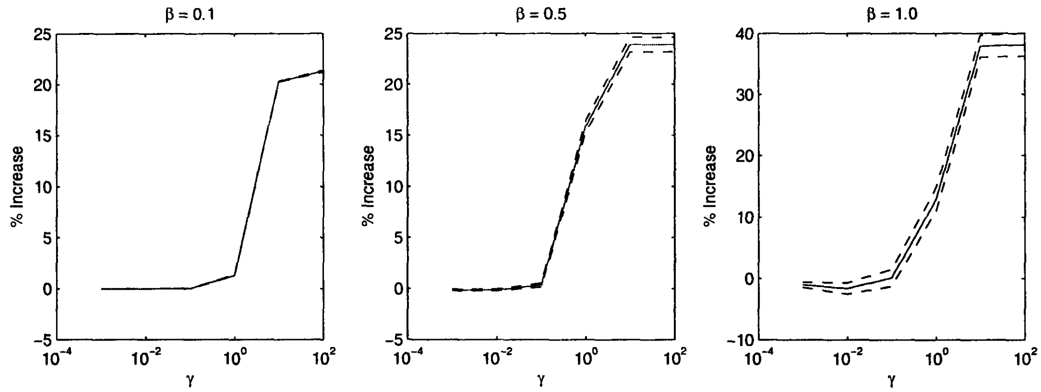

Figure 6-4: *Expected % increase in total cost from optimal policy versus γ for a problem with nonnegativity constraints on the control for various discount factors β. The dashed lines are 95% confidence intervals on the expected % increase.*

Here, β is a discounting factor which we varied in the simulations. Note that smaller β implies that the cost associated with later stages are less important.

The optimal policy was computed by enumeration to solve Bellman's recursion (using a discrete grid for the control and state spaces), under the assumption that the disturbances are generated independently $w_k \sim \mathcal{N}(0, 0.1)$. The results are illustrated graphically in Figure 6-4.

We make the following observations from these results:

1. The form of the increase in expected relative cost is very similar to the unconstrained case. In particular, for $\gamma \ll 1$, there is very little degradation in performance from the optimal policy. For large enough γ, the performance approaches an upper limit (in the range $20 - 40\%$ in this case) and does not go beyond that.

2. The performance loss is better for smaller β. This makes intuitive sense, as our approach does not exactly capture the tradeoff between current costs and future costs. When β is smaller, then, future costs are less relevant and SOCP (6.35), which is myopic, will perform better.

3. For $\beta = 1$, we see that our approach appears to outperform the optimal policy for small γ. Obviously, this cannot be the case. This can be attributed due to

estimation error, as well as small error in computing the optimal policy due to the discrete approximation of the state and control spaces.

We emphasize that it is *not* obvious that SOCP (6.35) will perform as well as the optimal policy for small γ (in fact, in general, it will not be the case). This is true in the unconstrained case via Corollaries 6.2.1 and 6.3.1. In this example, however, we can in fact do just as well as the optimal policy (which, in general, is very difficult to compute), by solving (6.35) (with constraints) closed-loop and keeping γ small.

6.7 Conclusions

A primary open question of interest is how to simplify this approach even further in the presence of constraints. The best we have here is to solve an SOCP at each step. From Section 6.6, we see how the complexity of this problem grows faster than linearly with the size of the problem. For large enough problems, solving this in closed-loop will become overly burdensome. It remains to be seen if the constrained SOCP has a structure which can be simplified for more efficient computations. Also of interest is extending many of the performance guarantees in Section 6.4.3 to our framework run in closed-loop fashion. Finally, an open question is if there are classes of constrained LQC problems for which a certainty equivalence principle holds. For such problems, it would likely be the case that controlling with the SOCP here with $\gamma = 0$ would result in an optimal policy.

Chapter 7

Reflections and future directions

This thesis has focused on a set-based model for uncertainty within the context of optimization and its relationship to various notions of risk. For static optimization problems, we were able to derive explicit connections between robust approaches and the theory of coherent and, more generally, convex risk measures. This framework, originally developed in the context of linear optimization, is applicable for optimization problems over convex cones. In the context of dynamic optimization, this thesis also presented a set-based model of uncertainty for optimization of quadratic costs over linear systems.

There are a number of open problems related to these ideas which demand further research efforts. We now briefly discuss these directions.

7.1 Open theoretical questions

The following problems represent some of the critical open problems related to the theory developed in this thesis.

7.1.1 Extensions to general distributions

The connections between risk theory and robust optimization explored in Chapters 3 and 4 have emphasized a data-driven approach in which the decision maker is equipped with scenarios for the uncertain parameters. While there is a clear, practical motivation for this approach, it is interesting from a theoretical perspective to see how the methodologies extend under, for instance, the underlying parameters possess a continuous distribution function. It would be interesting to see what types of robust problems arise in this setting.

There is reason to believe extending this theory in this way is more involved than simply utilizing similar, but more rigorous, duality results. For instance, from a tractability perspective, it is well-known that even evaluating CVaR, except for very special distributions, involves multi-dimensional integration, which can be computationally expensive (e.g., Nemirovski and Shapiro, [90]). In other words, it may be possible to derive equivalent robust problems, but such problems may be, in general, intractable. One possibility would be to develop approximation techniques to such problems, as the authors do in [90].

7.1.2 Probability guarantees for nondeterministic solutions

Chapter 5 derives implied probability guarantees for fixed solutions to risk-averse problems. Obviously, a more interesting question is what types of bounds can be shown for solutions which are chosen as functions of the realized data. The discussion in Chapter 5 highlighted some of the challenges in obtaining such results.

As we have mentioned, one way around many of these difficulties is to apply Vapnik-Chervonenkis theory, similar to Takeda and Kanamori [109]. This seems like a reasonable approach for linear problems and CVaR, but, for more general problems and more complex risk measures, the conservatism may result in unenlightening bounds.

The thrust of Chapter 5 was to assume only bounded support for the underlying random elements. In some situations, one may possess more distributional information. It would be interesting to see how any derived probability guarantees could be strengthened as one utilizes more distributional information, such as a hierarchy of increasing moments.

7.1.3 Tightness of approximation results

In Chapters 3 and 4, we showed a one-to-one correspondence between robust optimization problems and problems based on risk measures for the case of linear optimization under uncertainty. Chapter 5 extended these ideas to optimization of more general cones, and showed that these problems based on risk measures were an inner approximation to robust problems.

As we have discussed in Chapter 2, equivalent robust optimization problems to general, conic optimization problems generally has a very unfavorable increase in computational complexity. This has motivated approximation methods to these problems, such as the work of Bertsimas and Sim [33]. Our work may actually be viewed as a generalization of this approach. Indeed, Bertsimas and Sim approximate the following type of robust, conic constraint

$$f(\boldsymbol{x}, \boldsymbol{D}) \;\leq\; 0, \qquad \forall\, \boldsymbol{D} \in \mathcal{U} = \left\{ \boldsymbol{D}_0 + \sum_{j=1}^{N} z_j \boldsymbol{D}_j \;\middle|\; \|\boldsymbol{z}\| \leq \Omega \right\} \tag{7.1}$$

with the constraint

$$f(\boldsymbol{x}, \boldsymbol{D}_0) + \sum_{j=1}^{N} z_j f(\boldsymbol{x}, \boldsymbol{D}_j) \;\leq\; 0, \qquad \forall\, \boldsymbol{z} :\; \|\boldsymbol{z}\| \leq \Omega. \tag{7.2}$$

It is not hard to see that this type of constraint would also arise from a constraint based on a risk measure which depends on a Euclidean norm. Specifically, a second-order tail moment risk measure, as in Section 3.4, would induce a constraint like (7.2).

In any event, Chapter 5 does not quantify how tight the inner approximation is for the corresponding robust problem. It would be interesting to see what kinds of approximation results, similar in spirit to the work of Ben-Tal and Nemirovski [19], could be derived, and what classes of risk measures yield better approximations.

7.1.4 Dynamic risk measures

The starting point in Chapter 6 was from uncertainty sets, not risk measures. A primary reason for this is that the theory of dynamic risk measures, from an algorithmic perspective at least, is certainly far from developed.

An interesting, open problem is developing an axiomatized description of risk measures in a dynamic setting (e.g., Artzner et al. [4]) which result in practically solvable problem structures. Ideally, such a theory would not suffer the brittleness that plagues many dynamic optimization approaches; for instance, tractability would be preserved under small changes in the underlying constraint structure.

7.2 Application directions

While most of our focus here has been on questions related to the theory of these approaches, one can argue the ultimate metric of success for a theory of optimization under uncertainty would be its performance and utility for compelling applications. Given that the theory of risk was largely derived from an economic perspective, an obvious application domain of interest here is problems related to finance.

Portfolio optimization is a canonical problem in this category. CVaR and, more generally, convex risk measures, have certainly been applied to this problem (e.g., Rockafellar and Uryasev, [102], Lüthi and Doege, [82], Ben-Tal and Teboulle, [24]). In addition to understanding the structural properties of optimal portfolios using these ideas, more studies on the empirical performance of these approaches are needed. The example in Chapter 4 was merely a first step in this direction.

Risk measures have been widely used in the actuarial sciences as a pricing mechanism (e.g., Wang [114], [113]), and it would be interesting to see if the relationship of this theory to set-based uncertainty models would yield tractable methods for challenging valuation problems, such as options pricing. Clearly, this would rely on further developments of the theory within a dynamic setting. Other types of dynamic problems, such as the optimal timing of large market transactions (e.g., Bertsimas and Lo, [29]) could potentially benefit from the ideas presented here.

Outside of financial applications, there are also a number of important operations research problems which have been extended to corresponding, robust versions (inventory control, supply chain problems, revenue management, etc.). It would be interesting to utilize some of the connections here to see the implied benefits, from a risk perspective, of such extensions.

Bibliography

[1] C. Acerbi and D. Tasche. On the coherence of expected shortfall. *Journal of Banking and Finance*, 26(7), 2002.

[2] S. Amari. *Differential-geometrical methods in statistics*. Springer-Verlag, 1985.

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

[4] P. Artzner, F. Delbaen, J-M. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellman's principle. Working paper, 2005.

[5] A. Atamtürk. Strong formulations of robust mixed 0-1 programming. Forthcoming in *Mathematical Programming*, 2005.

[6] T. Başar and P. Bernhard. *H^∞-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Berkhäuser, Boston, 1995.

[7] A. Ben-Tal, S. Boyd, and A. Nemirovski. Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Mathematical Programming Series B*, 107, 2006.

[8] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. Robust semidefinite programming. in Saigal, R., Vandenberghe, L., Wolkowicz, H., eds., *Semidefinite programming and applications*. Kluwer Academic Publishers, 2000.

[9] A. Ben-Tal, B. Golany, A. Nemirovski, and J.P. Vial. Supplier-retailer flexible commitments contracts: A robust optimization approach. Submitted to *Manufacturing and Service Operations Management*. 2003.

[10] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Math. Programming*, 99:351–376, 2003.

[11] A. Ben-Tal, T. Margalit, and A. Nemirovski. Robust modeling of multi-stage portfolio problems. *High Performance Optimization.*

[12] A. Ben-Tal and A. Nemirovski. Robust truss topology design via semindefinite programming. *SIAM Journal on Optimization*, 7, 1997.

[13] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. Oper. Res.*, 23:769–805, 1998.

[14] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–13, 1999.

[15] A. Ben-Tal and A. Nemirovski. Robust solutions to uncertain programs. *Oper. Res. Let.*, 25, 1999.

[16] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming*, 88:411–421, 2000.

[17] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization.* MPS-Siam Series on Optimization, 2001.

[18] A. Ben-Tal and A. Nemirovski. On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty. *SIAM Journal on Optimization*, 12, 2002.

[19] A. Ben-Tal and A. Nemirovski. On approximate robust counterparts of uncertain semidefinite and conic quadratic programs. *Proceedings of 20th IFIP TC7 Conference on System Modelling and Optimization*, July 23-27, 2001, Trier, Germany.

[20] A. Ben-Tal, A. Nemirovski, and C. Roos. Robust solutions of uncertain quadratic and conic-quadratic problems. *SIAM Journal on Optimization*, 13, 2002.

[21] A. Ben-Tal and M. Teboulle. Portfolio theory for the recourse certainty equivalent maximizing investor. *Annals of Operations Research*, 31.

[22] A. Ben-Tal and M. Teboulle. Expected utility, penalty functions and duality in stochastic nonlinear programming. *Management Science*, 32, 1986.

[23] A. Ben-Tal and M. Teboulle. Penalty functions and duality in stochastic progamming via ϕ-divergence functionals. *Mathematics of Operations Research*, 12, 1987.

[24] A. Ben-Tal and M. Teboulle. An old-new concept of convex risk measures: the optimized certainty equivalent. To appear in *J. Math. Finance*, 2006.

[25] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22, 1954 [1738].

[26] D.P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, Belmont, Mass., 1995.

[27] D. Bertsimas and D.B. Brown. Constrained Stochastic LQC: A Tractable Approach. Technical Report LIDS 2658, M.I.T., August 2005.

[28] D. Bertsimas, G.J. Lauprete, and A. Samarov. Shortfall as a risk measure: properties, optimization and applications. *Journal of Economic Dynamics and Control*, 28(7):1353–1381, 2004.

[29] D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, pages 1–50, 1998.

[30] D. Bertsimas, D. Pachamanova, and M. Sim. Robust linear optimization under general norms. *Operations Research Letters*, 32:510–516, 2004.

[31] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming Series B*, 98:49–71, 2003.

[32] D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

[33] D. Bertsimas and M. Sim. Tractable approximations to robust conic optimization problems. To appear, Mathematical Programming, 2005.

[34] D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. Forthcoming in *Operations Research*, 2005.

[35] D. Bertsimas and J. Tsitsiklis. *Introduction to Linear Optimization.* Athena Scientific, 1997.

[36] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer-Verlag, 1997.

[37] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory.* SIAM Studies in Applied Mathematics, 1994.

[38] S. Boyd, S.-J. Kim, D. Patil, and M. Horowitz. Digital circuit sizing via geometric programming. *Operations Research*, 53(6), 2005.

[39] S. Boyd and L. Vandenberghe. Semidefinite programming. *SIAM Review*, 38(1), 1996.

[40] S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge University Press, 2004.

[41] G. Calafiore and M.C. Campi. Uncertain convex programs: randomized solutions and confidence levels. Working paper, 2003.

[42] G. Calafiore and M.C. Campi. Decision making in an uncertain environment: the scenario-based optimization approach. Working paper, 2004.

[43] C. Caramanis. *Adaptable Optimization: Theory and Algorithms.* Phd dissertation, Massachusetts Institute of Technology, 2006.

[44] X. Chen and X. Zhou. Stochastic LQ control with conic control constraints on an infinite time horizon. *SIAM Journal on Control and Optimization*, 43:1120–1150, 2004.

[45] G. Choquet. Theory of capacities. *Ann. Inst. Fourier*, pages 131–295, 1954.

[46] T.M. Cover and J.A. Thomas. *Elements of Information Theory.* Wiley, New York, 1991.

[47] I. Csiszár. Information-type measures of difference probability distributions and indirect observations. *Studia Sci. Math. Hungarica*, 2, 1967.

[48] D.P. de Farias and B. Van Roy. On constraint sampling in the linear programming approach to approximate dynamic programming. To appear in *Math. Oper. Res.*, 2001.

[49] F. Delbaen. Coherent risk measures on general probability spaces. Working paper, 2000.

[50] F. Delbaen. Pisa lecture notes: Coherent risk measures. 2000.

[51] D. Denneberg. *Non-Additive Measure and Integral.* Kluwer Academic Publishers, 1994.

[52] J. Dhaene, M. Denuit, M.J. Goovaerts, R. Kaas, and D. Vyncke. The concept of comonotonicity in actuarial science and finance: Theory. Working paper, 2001.

[53] Y. Eldar, A. Ben-Tal, and A. Nemirovski. Robust mean-squared error estimation in the presence of model uncertainties. To appear in *IEEE Trans. on Signal Processing*, 2005.

[54] E. Erdogan and G. Iyengar. Ambiguous chance constrained problems and robust optimization. To appear, Mathematical Programming, 2004.

[55] J.E. Falk. Exact solutions of inexact linear programs. *Operations Research*, 24, 1976.

[56] T. Fischer. Examples of coherent risk measures depending on one-sided moments. Working paper, 2001.

[57] P.C. Fishburn. *The Foundations of Expected Utility.* Reidel Publishing Company, Dordrecht, 1982.

[58] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6, 2002.

[59] H. Föllmer and A. Schied. Robust preferences and convex risk measures. In K. Sandmann and Ph. Schönbucher, editors, *Advances in Finance and Stochastics: Essays in Honor of Dieter Sondermann.* Springer, New York, Berlin, Heidelberg, 2002.

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

[61] R. Freund. Postoptimal analysis of a linear program under simultaneous changes in matrix coefficients. *Math. Prog. Study*, 24, 1985.

[62] M. Fritelli and E.R. Gianin. Putting order in risk measures. *Journal of Banking and Finance*, 26, 2002.

[63] L. El Ghaoui and G. Calafiore. Worst-case state prediction under structured uncertainty. *American Control Conference*, June, 1999.

[64] L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal Matrix Analysis and Applications*, 1997.

[65] L. El Ghaoui and A. Nilim. Robust markov decision processes with uncertain transition matrices. Forthcoming in *Operations Research*, 2005.

[66] L. El Ghaoui, F. Oustry, and H. Lebret. Robust solutions to uncertain semidefinite programs. *Siam J. Optimization*, 9(1), 1998.

[67] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18, 1989.

[68] D. Goldfarb and G. Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 1, 2003.

[69] D. Heath and H. Ku. Pareto equilibria with coherent measures of risk. *Mathematical Finance*, 14, 2004.

[70] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58, 1963.

[71] K.-L. Hsiung, S.-J. Kim, and S. Boyd. Tractable approximate robust geometric programming. Submitted to *Mathematical Programming*, 2005.

[72] K.-L. Hsiung, S.-J. Kim, and S. Boyd. Power control in lognormal fading wireless channels with uptime probability specifications via robust geometric programming. *Proceedings American Control Conference*, 6, June, 2005.

[73] P. Huber. *Robust Statistics.* Wiley, New York, 1981.

[74] R. Kaas, J. Dhaene, D. Vyncke, M.J. Goovaerts, and M. Denuit. A simple geometric proof that comonotonic risks have the convex-largest sum. *ASTIN Bulletin*, 32, 2002.

[75] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 1979.

[76] S.-J. Kim, A. Magnani, and S. Boyd. Robust fisher discriminant analysis. Accepted to *Neural Information Processing Systems*, 2005.

[77] P. Kouvelis and G. Yu. *Robust discrete optimization and its applications.* Kluwer Academic Publishers, Norwell, MA, 1997.

[78] G.R.G. Lanckriet, L. El Ghaoui, C. Bhattacharyya, and M.I. Jordan. A robust minimax approach to classification. *Journal of Machine Learning Research*, 3, 2002.

[79] M. Lobo. *Robust and Convex Optimization with Application in Finance.* PhD thesis, Stanford University, March 2000.

[80] M.S. Lobo and S. Boyd. The worst-case risk of a portfolio and robust portfolio optimization. Working paper. 2000.

[81] R. Lorenz and S. Boyd. Robust minimum variance beamforming. *IEEE Transactions on Signal Processing*, 53(5), 2005.

[82] H.-J. Lüthi and J. Doege. Convex risk measures for portfolio optimization and concepts of flexibility. *Mathematical Programming Series B*, 104(2-3), 2005.

[83] H.M. Markowitz. Portfolio selection. *Journal of Finance*, 7, 1952.

[84] H.M. Markowitz. Portfolio selection. *Journal of Finance*, 7(1), 1952.

[85] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 2000.

[86] C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics 1989*. Cambridge University Press, Cambridge, 1989.

[87] K. Natarajan, D. Pachamanova, and M. Sim. Constructing risk measures from uncertainty sets. Working paper, 2005.

[88] A. Nemirovski. Several np-hard problems arising in robust stability analysis. *Math. Control Signals Systems*, 6, 1993.

[89] A. Nemirovski and A. Shapiro. Scenario approximations of chance constraints. Available online at http://www.optimization-online.org, 2004.

[90] A. Nemirovski and A. Shapiro. Convex approximations of chance constrained programs. Available online at http://www.optimization-online.org, 2005.

[91] H.T. Nguyen and N.T. Nguyen. Random sets in decision-making. *Random Sets: Theory and Applications, IMA*, 97:297–320, 1997.

[92] F. Paganini. A set-based methodology for white noise modeling. Working paper, 2003.

[93] The St. Petersburg Paradox. `http://plato.stanford.edu/archives/win1999/entries/paradox-stpetersburg/`.

[94] A. Prékopa. *Stochastic Programming.* Kluwer, 1995.

[95] J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3, 1982.

[96] M.A. Rami and L. El Ghaoui. LMI optimizaiton for nonstandard Riccati equations arising in stochastic control. *IEEE Trans. Aut. Control*, 1997.

[97] M.A. Rami and X.Y. Zhou. Linear matrix inequalities, Riccati equations, and indefinite stochastic linear controls. *IEEE Trans. Aut. Control*, 45:1131–1143, 2000.

[98] M.R. Reesor and D.L. McLeish. Risk, entropy, and the transformation of distributions. Available online at http://www.bankofcanada.ca/en/res/wp02-11.htm, 2002.

[99] J. Renegar. Some perturbation theory for linear programming. *Mathematical Programming*, 65, 1994.

[100] R.T. Rockafellar. Optimization under uncertainty. Lecture notes, University of Washington.

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

[102] R.T. Rockafellar and S.P. Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2:21–41, 2000.

[103] A. Ruszczynski and A. Shapiro. *Stochastic Programming.* North-Holland Publishing Company, 2003.

[104] A. Ruszczynski and A. Shapiro. Optimization of risk measures. Available online at http://www.optimization-online.org, 2004.

[105] D. Schmeidler. Integral representation without additivity. In *Proceedings of the American Math Society*, volume 97, pages 255–261, 1986.

[106] A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. Available online at http://www.optimization-online.org, 2005.

[107] A.L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21:1154–1157, 1973.

[108] J.F. Sturm. Sedumi matlab optimization toolbox. Available online at http://fewcal.kub.nl/sturm/software/sedumi.html.

[109] A. Takeda and T. Kanamori. A robust optimization approach based on conditional value-at-risk measure and its applications to statistical learning problems. Working paper, 2005.

[110] E.I. Verriest and G.A. Pajunen. Quadratically saturated regulator for constrained linear systems. *IEEE Trans. Aut. Control*, 41:992–995, 1996.

[111] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, 1947.

[112] P.G. Voulgaris. Optimal H_2/l_1 control via duality theory. *IEEE Trans. Aut. Control*, 40:1881–1888, 1995.

[113] S. Wang. A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1):15–36.

[114] S. Wang. Insurance pricing and increase limits ratemaking by proportional hazards transforms. *Insurance: Mathematics and Economics*, 17:43–54.

[115] J.C. Willems. Least squares stationary optimal control and the algebraic Riccati equation. *IEEE Trans. Aut. Control*, 16(6):621–634, 1971.

[116] Y. Xu, K.-L. Hsiung, X. Li, I. Nausieda, S. Boyd, and L. Pileggi. Opera: Optimization with ellipsoidal uncertainty for robust analog ic design. *Proceedings of the ACM/IEEE Design Automation Conference*, June, 2005.

[117] M.E. Yaari. The dual theory of choice under risk. *Econometrica*, 55, 1987.

[118] D.D. Yao, S. Zhang, and X.Y. Zhou. Stochastic linear-quadratic control via semidefinite programming. *Siam Journal on Control and Optimization*, 40(3):801–823.

[119] K. Zhou, J.C. Doyle, and K. Glover. *Robust Optimal Control.* Prentice Hall, 1996.

[120] G.M. Ziegler. *Lectures on Polytopes.* Springer-Verlag, 1994.