European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Invited Review

Optimization under uncertainty and risk: Quadratic and copositive approaches

Immanuel M. Bomze[a,*], Markus Gabl[b]

[a] ISOR/VCOR/ds:UniVie, University of Vienna, Oskar-Morgenstern-Platz 1, Wien 1090, Austria
[b] IOR, Karlsruhe Institute of Technology, Kaiserstraße 12, Karlsruhe 76131, Germany

ARTICLE INFO

Article history:
Received 9 February 2022
Accepted 11 November 2022
Available online xxx

Keywords:
Conic programming and interior point methods
Quadratically constrained quadratic problems
Two-stage stochastic standard quadratic problems
Adjustable robust optimization
Distributionally robust optimization

ABSTRACT

Robust optimization and stochastic optimization are the two main paradigms for dealing with the uncertainty inherent in almost all real-world optimization problems. The core principle of robust optimization is the introduction of parameterized families of constraints. Sometimes, these complicated semi-infinite constraints can be reduced to finitely many convex constraints, so that the resulting optimization problem can be solved using standard procedures. Hence flexibility of robust optimization is limited by certain convexity requirements on various objects. However, a recent strain of literature has sought to expand applicability of robust optimization by lifting variables to a properly chosen matrix space. Doing so allows to handle situations where convexity requirements are not met immediately, but rather intermediately.

In the domain of (possibly nonconvex) quadratic optimization, the principles of copositive optimization act as a bridge leading to recovery of the desired convex structures. Copositive optimization has established itself as a powerful paradigm for tackling a wide range of quadratically constrained quadratic optimization problems, reformulating them into linear convex-conic optimization problems involving only linear constraints and objective, plus constraints forcing membership to some matrix cones, which can be thought of as generalizations of the positive-semidefinite matrix cone. These reformulations enable application of powerful optimization techniques, most notably convex duality, to problems which, in their original form, are highly nonconvex.

In this text we want to offer readers an introduction and tutorial on these principles of copositive optimization, and to provide a review and outlook of the literature that applies these to optimization problems involving uncertainty.

1. Introduction

Robust optimization is one of the two main paradigms, next to *stochastic optimization*, for dealing with the uncertainty inherent in almost all real-world optimization problems. The core principle of robust optimization is the introduction of parameterized families of constraints, enforced for all realizations of the *uncertainty parameters* belonging to a so-called *uncertainty set*. Using tools from convex optimization theory, these complicated semi-infinite constraints can often be reformulated into finitely many convex constraints, so that the resulting optimization problem can be solved using standard convex optimization procedures. Based on this simple idea, the framework of robust optimization allows for a unified treatment of a vast array of approaches to optimization under uncertainty, spawning countless generalizations such as *adjustable robust optimization (ARO)* and *distributionally robust optimization (DRO)*, yielding elaborate models, which in essence boil down to robust optimization problems. This lends further credence to the versatility of this framework.

Unfortunately, the flexibility of robust optimization is limited by certain convexity requirements on various objects, such as the parameterized constraints as well as the uncertainty set. However, a recent strain of literature has sought to expand applicability of robust optimization by lifting variables to a properly chosen matrix space. Doing so allows to handle situations where convexity requirements are not met immediately, but rather intermediately.

As stated above, robust counterparts can often be reformulated into a tractable, finite convex optimization problems. At the core of the machinery enabling these reformulations lies the observation

* Corresponding author.
E-mail addresses: immanuel.bomze@univie.ac.at (I.M. Bomze), markus@gabl.com (M. Gabl).

https://doi.org/10.1016/j.ejor.2022.11.020

that a robust constraint of the form

$$f(\mathbf{x},\mathbf{u}) \geq 0 \quad \text{for all } \mathbf{u} \in \mathcal{U} \quad \text{is equivalent to} \quad \inf_{\mathbf{u}\in\mathcal{U}} \{ f(\mathbf{x},\mathbf{u}) \} \geq 0\,.$$

In case the infimum problem admits a dual (involving dual variables $\boldsymbol{\lambda}$, a dual feasible set $\mathcal{D}$ (which typically also involves $\mathbf{x}$ in some tractable manner) and an appropriate dual objective function $\tilde{f}(\mathbf{x},\boldsymbol{\lambda})$, say) attaining its optimal value with zero duality gap, the robust constraint can further be reformulated into

$$\sup_{\boldsymbol{\lambda}\in\mathcal{D}} \left\{ \tilde{f}(\mathbf{x},\boldsymbol{\lambda}) \right\} \geq 0\,,$$

where finally the supremum operator can be dropped, since any nonnegative feasible value certifies that the supremum is nonnegative as well, so that the robust constraint is fulfilled.

The desired strong duality property is readily available in case the infimum problem is a convex optimization problem: here only mild additional regularity conditions, such as Slater's condition, need to be satisfied. Outside the domain of convex optimization, such strong duality results are much more scarce.

In the domain of (possibly nonconvex) quadratic optimization, the principles of *copositive optimization* act as a bridge leading to recovery of the desired convex structures. Copositive optimization has established itself as a powerful paradigm for tackling a wide range of *quadratically constrained quadratic optimization problems (QCQPs)*. It aims at reformulating QCQPs into linear convex-conic optimization problems involving only linear constraints and objective, plus constraints forcing membership to so-called *set-copositive* matrix cones, which can be thought of as generalizations of the positive-semidefinite matrix cone. These reformulations allow for the application of the powerful tools of convex optimization, most notably convex duality, to problems which, in their original form, are highly nonconvex.

In this text we want to offer readers an introduction and tutorial on these principles of copositive optimization, and to provide a review and outlook of the literature that applies these to robust optimization problems. We hope that the reader will acquire the following benefits:

- gaining an overview on existing copositive optimization approaches to robust optimization as well as open questions in this field;
- understanding basic principles of convexifying nonconvex QCQPs in the style of copositive optimization with a focus to practice-oriented applications;
- being exposed to open problems and interesting research directions, which hopefully inspire the pursuit of new research in this area.

Regarding the final point we will discuss open problems throughout the text. However, for the readers' convenience we will attach a dedicated "**section with open problems**" at the end of each topic, where we will summarize interesting research directions point by point.

In the sequel, we will not delve into much much detail on robust optimization theory, since there are great tutorials available, providing excellent introductions to the field and its various sub-genres, for example Bertsimas, Brown, & Caramanis (2011); Gorissen, Yanikoglu, & den Hertog (2015); Rahimian & Mehrotra (2019); Wiesemann, Kuhn, & Sim (2014); Yanikoglu, Gorissen, & den Hertog (2019). In the interest of a focused and concise presentation, we will also omit discussions on another strain of literature dealing with convexifications of QCQPs by means of the so-called *S-Lemma* and its many variants. However, let us highlight that this topic has strong ties with copositive optimization as well as robust optimization. Most notably, copositive optimization is sometimes referred to as an alternative to the S-Lemma in the context of robust optimization. While we will comment on this circumstance sporadically throughout the text, our discussion on the topic will be limited. The interested reader may refer to Ben-Tal, Goryashko, Guslitzer, & Nemirovski (2004); Bomze & Gabl (2021); Jeyakumar, Li, & Woolnough (2021); Pólik & Terlaky (2007); Woolnough, Jeyakumar, & Li (2021).

The rest of this article is organized as follows: in Section 2 we will give a detailed but by no means exhaustive account of copositive optimization theory and related topics, concluding with a guide through surrounding literature. After briefly introducing basic concepts of robust optimization and some of its variants in Section 3, we will discuss in greater detail the various ways copositive optimization has been applied in robust optimization contexts. A core technique in this regard is the reformulation of semi-infinite constraints with quadratic index, which we will discuss extensively in Section 4. Some of the adjustable robust models discussed there can be tackled by an alternative approach which seeks to reformulate the entire problem rather than individual constraints and is discussed in Section 5. We then review robust versions and a two-stage stochastic version of the so-called *Standard Quadratic Optimization Problem* in Sections 6 and 7, respectively. A copositive approach to mixed-binary linear optimization under objective uncertainty, that sits conceptually in-between stochastic optimization and distributionally robust optimization, is presented in Section 8. Finally, we discuss a conic approach to two-stage distributionally robust optimization in Section 9.

1.1. Notation

Throughout the paper, matrices are denoted with sans-serif capital letters, e.g., E is the matrix of all ones, I is the identity matrix and O the matrix of all zeros (the matrix order will depend on the context). Vectors will be given as boldface lower case letters, for instance the vector of all ones (a column of E) is $\mathbf{e}$, the vector of zeros is $\mathbf{o}$ and the vector $\mathbf{e}_i$ is the ith column of I. By $^\top$ we denote transpose. For a square matrix M, $\operatorname{diag}\mathsf{M}$ extracts its diagonal as a column vector while $\operatorname{Diag}\mathbf{x}$ produces a diagonal matrix with diagonal $\mathbf{x}$. For any $\mathbf{x} = [x_i]_i \in \mathbb{R}^n$ we denote by $\mathbf{x}\circ\mathbf{x} = [x_i^2]_i \in \mathbb{R}^n$ its Hadamard square. We will also use the shorthand

$$\mathsf{Y}(\mathbf{x},\mathsf{X}) := \begin{bmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathsf{X} \end{bmatrix}.$$

Sets will mostly be indicated using letters or acronyms in capital calligraphic font. Most importantly: $\mathcal{S}^n$ is the space of symmetric $n\times n$ matrices, $\mathcal{N}^n \subset \mathcal{S}^n$ those of them with no negative entries and $\mathcal{S}^n_+$ those of them with no negative eigenvalues, i.e., positive-semidefinite (psd) symmetric matrices of order n (sometimes the cone $\mathcal{S}^n_+$ is referred to as the psd-cone in short), $\mathcal{SOC}^n = \left\{ (x_0,\mathbf{x}^\top)^\top \in \mathbb{R}^n : \|\mathbf{x}\| \leq x_0 \right\}$ is the second-order cone.

There are occasional exceptions, e.g., the n-dimensional Euclidean space $\mathbb{R}^n$, its nonnegative orthant $\mathbb{R}^n_+$, or the index set $[i\colon j] = \{i, i+1, \ldots, j-1, j\}$, where $i<j$ are integer numbers. For a set $\mathcal{A}$ we denote $\operatorname{cl}(\mathcal{A}), \operatorname{int}(\mathcal{A}), \operatorname{conv}(\mathcal{A})$ its closure, interior, and convex hull, respectively, and for a convex set $\mathcal{A}$ we denote by $\operatorname{relint}(\mathcal{A})$ its relative interior, as well by $\operatorname{ext}(\mathcal{A})$ the set of its extremal points. For a cone $\mathcal{K} \in \mathbb{R}^n$ we denote the dual cone as $\mathcal{K}^* := \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{y}^\top\mathbf{x} \geq 0 \text{ for all } \mathbf{y} \in \mathcal{K} \right\}$. For any optimization problem (P), we denote by $\operatorname{val}(P)$ its optimal value, regardless whether it is attained or not.

2. Convexifying QCQPs via set-copositive optimization

2.1. Basic lifting strategies and their core ingredients

A QCQP consists of minimizing a quadratic function subject to quadratic constraints, formally given by

$$\inf_{\mathbf{x}\in\mathcal{K}} \left\{ \mathbf{x}^\top\mathsf{Q}_0\mathbf{x} + 2\mathbf{q}_0^\top\mathbf{x} - \omega_0 : \mathbf{x}^\top\mathsf{Q}_i\mathbf{x} + 2\mathbf{q}_i^\top\mathbf{x} \leq \omega_i,\ i \in [1:m] \right\} \qquad (1)$$

where $\{\mathsf{Q}_i : i \in [0{:}m]\} \subset \mathcal{S}^n$, $\{\mathbf{q}_i : i \in [1{:}m]\} \subset \mathbb{R}^n$ and ω_i are real numbers. $\mathcal{K} \subseteq \mathbb{R}^n$ is a cone which one could choose to be any cone representable by (finitely many) linear or quadratic inequality constraints without leaving the domain of QCQPs, for instance $\mathcal{K} \in \left\{\mathbb{R}^n, \mathbb{R}^n_+, \mathcal{SOC}^n\right\}$. Note that neither the objective nor the feasible set need be convex, the latter may even be disconnected. Indeed, general QCQPs are NP-hard as they contain many NP-hard problems as special cases (see e.g. Pardalos & Vavasis, 1991).

In our discussion we want to familiarize the reader with a specific type of convexification of QCQPs, that is simple, yet ultimately very powerful. To convince even readers who are unfamiliar with the subject of the simplicity of the approach, we will now discuss some simple examples that nonetheless exhibit all the ingredients that are necessary for understanding the machinery.

Example 1. Consider the following optimization problem:

$$\min_{\mathbf{x}\in\mathbb{R}^n} \left\{\mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} + 2\mathbf{q}^\mathsf{T}\mathbf{x} : \mathbf{x} \in \{\mathbf{a}, \mathbf{b}\} \subset \mathbb{R}^n\right\}. \tag{2}$$

Clearly, the problem is easily solved by just evaluating the objective at both feasible points and then choosing the minimizer. Still, we have a (possibly) nonconvex objective that is optimized over a nonconvex feasible set, so that the problem belongs to a class of actually difficult problems and it is in fact a nice take-off point for thinking about how to convexify more general problems in this class. Firstly, observe the following equivalence: $\mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} = \mathrm{Tr}(\mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x}) = \mathrm{Tr}(\mathsf{Q}\mathbf{x}\mathbf{x}^\mathsf{T}) = \mathsf{Q} \bullet \mathbf{x}\mathbf{x}^\mathsf{T}$, which holds since the trace of a number is the identity function and the trace-operator is invariant under cyclic permutation of matrix products. Note that the Frobenius product is bilinear, so that we can achieve a linearization of the problem via the following modifications:

$$\min_{\mathbf{x}\in\mathbb{R}^n, \mathsf{X}\in\mathcal{S}^n} \left\{\mathsf{Q} \bullet \mathsf{X} + 2\mathbf{q}^\mathsf{T}\mathbf{x} : \mathsf{X} = \mathbf{x}\mathbf{x}^\mathsf{T},\ \ \mathbf{x} \in \{\mathbf{a}, \mathbf{b}\}\right\} \tag{3}$$

Further, we can eliminate the explicit relation between X and $\mathbf{x}$ by pushing it into the description of the feasible set in order to obtain

$$\min_{\mathbf{x}\in\mathbb{R}^n, \mathsf{X}\in\mathcal{S}^n} \left\{\mathsf{Q} \bullet \mathsf{X} + 2\mathbf{q}^\mathsf{T}\mathbf{x} : (\mathbf{x}, \mathsf{X}) \in \left\{\left(\mathbf{a}, \mathbf{a}\mathbf{a}^\mathsf{T}\right), \left(\mathbf{b}, \mathbf{b}\mathbf{b}^\mathsf{T}\right)\right\}\right\}. \tag{4}$$

A convexification is now easily obtained by replacing the feasible set with its convex hull, since the linear constraint will attain its optimum at an extreme point of the so obtained convex feasible set. In our case, the latter is a line segment connecting the two points in the feasible set of (4), which also are the extreme points of this line segment. Rather than expressing this convexification in the space of tuples of the form $(\mathbf{x}, \mathsf{X})$, it is instructive to represent it entirely in the space $\mathcal{S}^{n+1}$ in the following manner:

$$\min_{\mathbf{x}\in\mathbb{R}^n, \mathsf{X}\in\mathcal{S}^n} \left\{\mathsf{Q} \bullet \mathsf{X} + 2\mathbf{q}^\mathsf{T}\mathbf{x} : \mathsf{Y}(\mathbf{x}, \mathsf{X}) \in \mathrm{conv}\left\{\mathsf{Y}(\mathbf{a}, \mathbf{a}\mathbf{a}^\mathsf{T}), \mathsf{Y}(\mathbf{b}, \mathbf{b}\mathbf{b}^\mathsf{T})\right\}\right\}. \tag{5}$$

Note that

$$\mathsf{Y}\left(\mathbf{a}, \mathbf{a}\mathbf{a}^\mathsf{T}\right) = \begin{bmatrix} 1 & \mathbf{a}^\mathsf{T} \\ \mathbf{a} & \mathbf{a}\mathbf{a}^\mathsf{T} \end{bmatrix} = \begin{bmatrix} 1 \\ \mathbf{a} \end{bmatrix}\begin{bmatrix} 1 \\ \mathbf{a} \end{bmatrix}^\mathsf{T}$$

so that in fact, the feasible set is the convex hull of symmetric dyadic matrices (i.e., matrices of the form $\mathbf{x}\mathbf{x}^\mathsf{T}$) where the last n components of the factors $\mathbf{x}$ are feasible solutions to the original optimization problem. In addition, these dyadic matrices form the extreme points of the feasible set and the linear function will attain its minimum at one of these points. Finally, at these dyadic extreme points the linear objective will evaluate identically to the quadratic function at the respective feasible points, so that the convexification enjoys zero gap.

The hitherto exemplified construction of the feasible set of the convex reformulation is critical for the understanding of the convexification strategy we want to convey to the reader. What we demonstrated in the above example for the case where the original feasible set contained just two points can be generalized to the case where the feasible set, say $\mathcal{F}$, is arbitrary. In this case a general convexification can by achieved via a lifted set given by

$$\mathcal{G}(\mathcal{F}) := \mathrm{clconv}\left\{\begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix}\begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix}^\mathsf{T} : \mathbf{x} \in \mathcal{F}\right\}.$$

Characterizing $\mathcal{G}(\mathcal{F})$ for a given set $\mathcal{F}$ is challenging, and we will spend a considerable part of this text discussing known strategies, and highlighting open questions in this regard. However, irrespective of the characterization, optimizing a linear function over this set will always yield optimal points that are dyadic matrices whose factors contain $\mathbf{x} \in \mathcal{F}$.

It is however noteworthy that not all optimal solutions to problems of the type (5) and its generalization have this quality. But in general optimal solutions are always in the convex hull of the optimal dyadic solutions.

We also want to highlight the fact that all dyadic matrices are positive-semidefinite. In fact, the psd-cone is the convex hull of all symmetric dyadic matrices, which are also the generators of the extreme rays of that cone. This foreshadows the fact that, in practice, many characterizations of $\mathcal{G}(\mathcal{F})$ are achieved via conic intersections involving suitable sub-cones of the psd-cone, namely the so-called *set-completely positive cones* whose extreme rays are generated by the dyadic matrices the factors of which are elements of certain sets. We will discuss these objects in more detail later in the text.

At this point we want to further the intuition regarding our convexification strategy by repeating a neat example originally given in Burer (2015), which we will discuss in extensively in order to highlight its connection to the rest of our exposition.

Example 2. The next example is an extended take on an example discussed in Burer (2015). Consider the following optimization problem:

$$\min_{x\in\mathbb{R}} \left\{Qx^2 + 2qx : 1 \geq x \geq -1\right\}. \tag{6}$$

Depending on the sign of the coefficient Q this can be a nonconvex quadratic optimization problem, which we will now conexify in the style discussed in this section. In some simple steps we can obtain

$$\begin{aligned} &\min_{x\in\mathbb{R}} \left\{Qx^2 + 2qx : 1 \geq x \geq -1\right\} \\ =\ &\min_{(x,X)\in\mathbb{R}^2} \left\{QX + 2qx : X = x^2,\ \ 1 \geq x \geq -1\right\} \\ =\ &\min_{(x,X)\in\mathbb{R}^2} \left\{QX + 2qx : 1 \geq X = x^2\right\} \\ =\ &\min_{(x,X)\in\mathbb{R}^2} \left\{QX + 2qx : 1 \geq X \geq x^2\right\} \end{aligned}$$

where only the last equality merits justification. The feasible set of the final optimization problem is the convex hull of the parabola where $X = x^2$ that is truncated at height equal to one. Since all extreme points of this set correspond to points at the parabola, the relaxation is tight.

We will now give some more insight on how this geometry relates to the discussion so far. Consider the fact that by Schur complementation we have

$$X \geq x^2 \Leftrightarrow \begin{bmatrix} 1 & x \\ x & X \end{bmatrix} \in \mathcal{S}^2_+.$$

We can again write the optimization problem in the lifted space of 2×2 symmetric matrices as to obtain

$$\min_{(x_0,x,X)\in\mathbb{R}^3} \begin{bmatrix}0 & q\\ q & Q\end{bmatrix} \bullet \begin{bmatrix}x_0 & x\\ x & X\end{bmatrix}$$
$$\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix} \bullet \begin{bmatrix}x_0 & x\\ x & X\end{bmatrix} = 1\,,$$
$$\begin{bmatrix}0 & 0\\ 0 & 1\end{bmatrix} \bullet \begin{bmatrix}x_0 & x\\ x & X\end{bmatrix} \le 1\,,$$
$$\begin{bmatrix}x_0 & x\\ x & X\end{bmatrix} \in \mathcal{S}^2_+\,.$$

Here we explicitly write down the de-homogenizing equation $x_0 = 1$ in order to make the geometry of the feasible set as transparent as possible. We see that the set of feasible matrices is again a subset of the psd-cone. More importantly, the extreme points of the feasible set are all boundary points of the psd-cone, which, in case of 2×2 matrices, are all dyadic matrices (in higher dimensions the psd-cone has non-dyadic boundary points). Hence, the above optimization problem will attain its optimal value at a point where

$$\begin{bmatrix}x_0 & x\\ x & X\end{bmatrix} = \begin{bmatrix}1 & x\\ x & x^2\end{bmatrix} = \begin{bmatrix}1\\ x\end{bmatrix}\begin{bmatrix}1\\ x\end{bmatrix}^{\mathsf T},\ \text{with } -1\le x\le 1\,,$$

or in other words, the set of feasible matrices of the relaxation is precisely $\mathcal{G}(\mathcal{F})$, where $\mathcal{F}$ is the original feasible set.

In the previous example, consider the case where $Q=-1$ and $q=0$, so that the original quadratic problem is a nonconvex problem with optimal value given by -1, which is attained at $x\in\{-1,1\}$. The convex reformulation gives the same optimal value and indeed the points $(x,X)\in\{(-1,1),(1,1)\}$ are optimal solutions. But so are all the points $(x,X)=\lambda(-1,1)+(1-\lambda)(1,1)$, $\lambda\in[0,1]$, or, expressed in the lifted space

$$\begin{bmatrix}x_0 & x\\ x & X\end{bmatrix} = \lambda\begin{bmatrix}1\\ -1\end{bmatrix}\begin{bmatrix}1\\ -1\end{bmatrix}^{\mathsf T} + (1-\lambda)\begin{bmatrix}1\\ 1\end{bmatrix}\begin{bmatrix}1\\ 1\end{bmatrix}^{\mathsf T},\quad \lambda\in[0,1]\,,$$

which illustrates that the optimal solutions of the relaxation are in the convex hull of its dyadic solutions. Since the latter correspond to optimal solutions of the original problem, the x component of the optimal solution to our relaxation are always in the convex hull of optimal solutions to the original problem. Hence, unless the original feasible set is already convex, the $\mathbf{x}$ components of a solution to the reformulation are not necessarily feasible to the original problem.

One must however not confuse convex combinations in the original space of variables with convex hulls in the lifted space! It is vital to understand that $\mathcal{G}(\mathrm{conv}(\mathcal{F}))$ is always a strictly larger set than $\mathcal{G}(\mathcal{F})$, unless $\mathcal{F}$ is a singleton. Said differently: convex combinations in the original space do not correspond to convex combinations in the lifted space. To illustrate this point, let us revisit the problem in Example 1 for the special case $n=1$, $\mathbf{a}=-1$ and $\mathbf{b}=1$. As we can see, the feasible set of the problem in Example 2 is just the convex hull of these points. However, the feasible set of the latter problems convexification is not just the convex hull of the two lifted extreme points, but the convex hull of an entire curve of points, each of which represents a lifting of a convex combination of the points $\{1,-1\}$. Merely considering the convex hull of the lifted extreme points of the interval yields $\mathcal{G}(\{1,-1\})$, i.e., the feasible set of the convexification problem in Example 1, which is a much smaller lifted set. In fact, no dyadic matrix can be expressed as the convex combination of two dyadic matrices which are not just re-scalings of that matrix, i.e., $\mathbf{x}\mathbf{x}^{\mathsf T}=\lambda\mathbf{y}_1\mathbf{y}_1^{\mathsf T}+(1-\lambda)\mathbf{y}_2\mathbf{y}_2^{\mathsf T}$ implies $\mathbf{y}_i\mathbf{y}_i^{\mathsf T}=\mu_i\mathbf{x}\mathbf{x}^{\mathsf T}$, $\mu_i\ge 0$, $i\in[1\!:\!2]$, as we prove later (see the proof of Proposition 11 in the appendix).

With the preceding discussion in mind, the following theorem, which is at the heart of all convexifications of QCQPs we will discuss in this text, should be easily accessible to the reader.

Theorem 1. *Let* $\mathcal{F} := \left\{\mathbf{x}\in\mathcal{K} : \mathbf{x}^{\mathsf T}\mathsf{Q}_i\mathbf{x}+2\mathbf{q}_i^{\mathsf T}\mathbf{x}\le\omega_i,\ i\in[1\!:\!m]\right\}\subseteq\mathbb{R}^n$ *be a feasible set of a QCQP where* $\mathcal{K}$ *is a closed cone, and denote by*

$$\mathcal{G}(\mathcal{F}) = \mathrm{clconv}\left\{\begin{bmatrix}1\\ \mathbf{x}\end{bmatrix}\begin{bmatrix}1\\ \mathbf{x}\end{bmatrix}^{\mathsf T} : \mathbf{x}\in\mathcal{F}\right\},$$

where $\mathrm{clconv}(\mathcal{A})$ *stands for the closure of the convex hull of a set* $\mathcal{A}$. *Then for any* $\mathsf{Q}_0\in\mathcal{S}^n$ *and* $\mathbf{q}_0\in\mathbb{R}^n$ *we have*

$$\begin{aligned}\mathrm{val(P)} &:= \inf_{\mathbf{x}\in\mathcal{F}}\left(\mathbf{x}^{\mathsf T}\mathsf{Q}_0\mathbf{x}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0\right)\\ &= \inf_{\mathsf{Y}(\mathbf{x},\mathsf{X})\in\mathcal{G}(\mathcal{F})}\left(\mathsf{Q}_0\bullet\mathsf{X}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0\right) =: \mathrm{val(R)}\,.\end{aligned}$$

Proof. See, e.g. Burer & Anstreicher (2013); Eichfelder & Povh (2013). For the readers' convenience we repeat the argument here. We refer to the QCQP as (P) and to the reformulation as (R). Let $\mathbf{x}$ be feasible for (P), then $\left(\mathbf{x},\mathbf{x}\mathbf{x}^{\mathsf T}\right)$ is feasible for (R) with identical objective function value given by $\mathsf{Q}_0\bullet\mathbf{x}\mathbf{x}^{\mathsf T}+\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0=\mathbf{x}^{\mathsf T}\mathsf{Q}_0\mathbf{x}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0$. Thus $\mathrm{val(R)}\le\mathrm{val(P)}$. For the converse, let $(\mathbf{x},\mathsf{X})$ be ε-optimal for (R), i.e., $\mathsf{Q}_0\bullet\mathsf{X}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0\le\mathrm{val(R)}+\varepsilon$. (We need an arbitrarily small $\varepsilon>0$ in case that is val(R) not attained.) Then by definition of $\mathcal{G}(\mathcal{F})$ as the closure we have likewise $d\Big((\mathbf{x},\mathsf{X}),\sum_{i=1}^k\lambda_i\big(\mathbf{x}_i,\mathbf{x}_i\mathbf{x}_i^{\mathsf T}\big)\Big)<\delta$ with $\mathbf{x}_i\in\mathcal{F}$, $\sum_{i=1}^k\lambda_i=1$ and $\lambda_i\ge 0$, $i\in[1\!:\!k]$, and $\delta>0$ so small that, by continuity,

$$|\mathsf{Q}_0\bullet\mathsf{X}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\sum_{i=1}^k\lambda_i\left[\mathbf{x}_i^{\mathsf T}\mathsf{Q}_0\mathbf{x}_i+2\mathbf{q}_0^{\mathsf T}\mathbf{x}_i\right]|<\varepsilon\,.$$

So, on one hand, $\mathbf{x}_i^{\mathsf T}\mathsf{Q}_0\mathbf{x}_i+2\mathbf{q}_0^{\mathsf T}\mathbf{x}_i-\omega_0\ge\mathrm{val(P)}$ for all $i\in[1\!:\!k]$ and on the other hand,

$$\begin{aligned}\mathrm{val(R)}+\varepsilon &\ge \mathsf{Q}_0\bullet\mathsf{X}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\omega_0\\ &= \mathsf{Q}_0\bullet\mathsf{X}+2\mathbf{q}_0^{\mathsf T}\mathbf{x}-\ \omega_0-\sum_{i=1}^k\lambda_i\left[\mathbf{x}_i^{\mathsf T}\mathsf{Q}_0\mathbf{x}_i+2\mathbf{q}_0^{\mathsf T}\mathbf{x}_i-\omega_0\right]\\ &\quad+\sum_{i=1}^k\lambda_i\left[\mathbf{x}_i^{\mathsf T}\mathsf{Q}_0\mathbf{x}_i+2\mathbf{q}_0^{\mathsf T}\mathbf{x}_i-\omega_0\right]\\ &\ge -\varepsilon+\sum_{i=1}^k\lambda_i\left[\mathbf{x}_i^{\mathsf T}\mathsf{Q}_0\mathbf{x}_i+2\mathbf{q}_0^{\mathsf T}\mathbf{x}_i-\omega_0\right]\\ &\ge -\varepsilon+\sum_{i=1}^k\lambda_i\,\mathrm{val(P)}=\mathrm{val(P)}-\varepsilon\,,\end{aligned}$$

which shows $\mathrm{val(R)}+2\varepsilon\ge\mathrm{val(P)}$. As ε was arbitrarily small, we arrive at $\mathrm{val(R)}\ge\mathrm{val(P)}$. □

Despite the simplicity of the theorem we want to take a moment and reconsider the core ingredients that enable its validity. The first one is a linearization by lifting to matrix variables: from a quadratic form $\mathbf{x}^{\mathsf T}\mathsf{Q}\mathbf{x}=\mathrm{Tr}(\mathbf{x}^{\mathsf T}\mathsf{Q}\mathbf{x})=\mathrm{Tr}(\mathsf{Q}\mathbf{x}\mathbf{x}^{\mathsf T})=\mathsf{Q}\bullet\mathbf{x}\mathbf{x}^{\mathsf T}$ we pass on to a linear form $\mathsf{Q}\bullet\mathsf{X}$, in substituting X_{ij} for x_ix_j. The second ingredient is the set $\mathcal{G}(\mathcal{F})$. Merely requiring that $(\mathbf{x},\mathsf{X})\in\left\{(\mathbf{x},\mathbf{x}\mathbf{x}^{\mathsf T}):\mathbf{x}\in\mathcal{F}\right\}$ would obviously render the linearization to be exact. But linear optimization is invariant to taking the convex hull of the feasible set, a fact often exploited in, for example, mixed integer linear optimization, where one seeks to find the convex hull of integer points.

The characterization of $\mathcal{G}(\mathcal{F})$ is the major challenge when employing the reformulation strategy depicted in Theorem 1 and a general workable description of $\mathcal{G}(\mathcal{F})$ is not known. There are, however, characterizations for specific instances of $\mathcal{F}$.

References to important examples of such reformulations in literature will be given in the sequel and will be summarized in Section 2.4.

2.1.1. Lower bounds by Shor relaxation: exactness and strengthening results

A natural starting point for the construction of $\mathcal{G}(\mathcal{F})$ is based upon the so-called *Shor relaxation* introduced in Shor (1987). A central role here is played by the *set-completely positive matrix cone* defined as

$$\mathcal{CPP}(\mathcal{K}) := \text{conv}\left\{\mathbf{x}\mathbf{x}^\mathsf{T} : \mathbf{x} \in \mathcal{K}\right\},$$

for a cone $\mathcal{K} \subseteq \mathbb{R}^n$. The matrix cone $\mathcal{CPP}(\mathcal{K})$ is a closed cone whenever $\mathcal{K}$ is closed, and with nonempty interior whenever $\mathcal{K}$ has nonempty interior (see e.g., Mittal & Hanasusanto, 2021, Lemma 4 or Tuncel & Wolkowicz, 2012, Theorem 5.1). It is the convex hull of extreme rays spanned by dyadic matrices. These are precisely the positive-semidefinite matrices of rank equal to 1, except for the zero matrix $\mathsf{O} = \mathbf{o}\mathbf{o}^\mathsf{T}$, which has rank equal to zero. In general, $\mathcal{CPP}(\mathcal{K})$ is an intractable cone in that membership of a given matrix is hard to decide (Dickinson & Gijben, 2014). Thus, when working with this object, one is bound to use either approximations or clever tools to check membership. Since these tools are essential when working with $\mathcal{CPP}(\mathcal{K})$, we will devote an entire section to this matter (see Section 2.3.1). In the present section, we will merely focus on its relation to the Shor relaxation, which can be best explained by looking at a homogeneous QCQP:

$$\begin{aligned}
&\min_{\mathbf{x}\in\mathcal{K}} \left\{\mathbf{x}^\mathsf{T}\mathsf{Q}_0\mathbf{x} : \mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} \le \omega_i,\ i \in [1\!:\!m]\right\},\\
&= \min_{\mathsf{X}} \left\{\mathsf{Q}_0 \bullet \mathsf{X} : \mathsf{Q}_i \bullet \mathsf{X} \le \omega_i,\ i \in [1\!:\!m],\ \mathsf{X} \in \left\{\mathbf{x}\mathbf{x}^\mathsf{T} : \mathbf{x} \in \mathcal{K}\right\}\right\},\\
&\ge \min_{\mathsf{X}} \left\{\mathsf{Q}_0 \bullet \mathsf{X} : \mathsf{Q}_i \bullet \mathsf{X} \le \omega_i,\ i \in [1\!:\!m],\ \mathsf{X} \in \mathcal{CPP}(\mathcal{K})\right\},
\end{aligned}$$

where we added an intermediate step for the sake of transparency. In case additional linear terms $\mathbf{q}_i^\mathsf{T}\mathbf{x}$, $i \in [0\!:\!m]$, are present, one can always recover the homogeneous case by enriching $\mathbf{x}$ by an additional coordinate $x_0 \in \mathbb{R}_+$, so that $\left(x_0, \mathbf{x}^\mathsf{T}\right)^\mathsf{T} \in \mathbb{R}_+ \times \mathcal{K}$, and adding one de-homogenizing constraint $x_0^2 = 1$. Note that it is indeed important to have $x_0 \ge 0$ in order to secure the implication $x_0^2 = 1 \Rightarrow x_0 = 1$. The Shor relaxation then is given by adding a row and column to the data and to the matrix variables:

$$\begin{aligned}
&\min_{\mathsf{X},\mathbf{x}}\{\widehat{\mathsf{Q}}_0 \bullet \mathsf{Y}(\mathbf{x},\mathsf{X}) : \widehat{\mathsf{Q}}_i \bullet \mathsf{Y}(\mathbf{x},\mathsf{X})\\
&\quad \le 0\,,\ i \in [1\!:\!m]\,,\ \mathsf{Y}(\mathbf{x},\mathsf{X}) \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})\},
\end{aligned}$$

with

$$\mathsf{Y}(\mathbf{x},\mathsf{X}) := \begin{bmatrix} 1 & \mathbf{x}^\mathsf{T} \\ \mathbf{x} & \mathsf{X} \end{bmatrix}, \quad \widehat{\mathsf{Q}}_i := \begin{bmatrix} -\omega_i & \mathbf{q}_i^\mathsf{T} \\ \mathbf{q}_i & \mathsf{Q}_i \end{bmatrix},$$

so that

$$\widehat{\mathsf{Q}}_i \bullet \mathsf{Y}(\mathbf{x},\mathsf{X}) = \mathsf{Q}_i \bullet \mathsf{X} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} - \omega_i\,,\ \text{all } i \in [0\!:\!m]\,.$$

While $\mathcal{G}(\mathcal{F})$ is the convex hull of intersections of the halfspaces induced by $\mathsf{Q}_i \bullet \mathsf{X} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} \le \omega_i$ and $\left\{(\mathbf{x}, \mathbf{x}\mathbf{x}^\mathsf{T}) : \mathbf{x} \in \mathcal{K}\right\}$, the feasible set of the Shor relaxation $\mathcal{F}_{Shor}$ is the intersection of said half spaces with $\text{conv}\left\{(\mathbf{x}, \mathbf{x}\mathbf{x}^\mathsf{T}) : \mathbf{x} \in \mathcal{K}\right\}$, hence the latter must be the bigger set in general. Another way to see that in general $\mathcal{F}_{Shor} \supseteq \mathcal{G}(\mathcal{F})$ is the simple fact that not all matrices in $\mathcal{CPP}(\mathcal{K})$ are of the form $\mathbf{x}\mathbf{x}^\mathsf{T}$. Only the extreme matrices of $\mathcal{CPP}(\mathcal{K})$ have this property. The following examples illustrate one case where $\mathcal{G}(\mathcal{F})$ and $\mathcal{F}_{Shor}$ coincide and another one where they differ.

Example 3. Consider the following quadratic program and its Shor relaxation

$$\left.\begin{array}{l}
\inf_{\mathbf{x}\in\mathbb{R}^2} q_{11}x_1^2 + 2q_{12}x_1x_2 + q_{22}x_2^2\\
\quad \text{s.t.}: 2x_1^2 + x_2^2 \le 12,\\
\qquad\quad x_1^2 + 2x_2^2 \le 12\\
\qquad\quad 4x_1^2 + x_2^2 \ge 4\\
\qquad\quad x_1^2 + 4x_2^2 \ge 4
\end{array}\right\} \text{ and}$$

$$\left\{\begin{array}{l}
\inf_{\mathsf{X}\in\mathcal{S}_+^2} q_{11}X_{11} + 2q_{12}X_{12} + q_{22}X_{22}\\
\quad \text{s.t.}: 2X_{11} + X_{22} \le 12,\\
\qquad\quad X_{11} + 2X_{22} \le 12\\
\qquad\quad 4X_{11} + X_{22} \ge 4\\
\qquad\quad X_{11} + 4X_{22} \ge 4\,.
\end{array}\right.$$

The feasible sets of these problems are depicted in Fig. 1. Since all matrices at the boundary of $\mathcal{S}_+^2$ are dyadic matrices, we see that the extreme points of the lifted feasible set are also dyadic. Therefore the relaxation has optimal solutions of the form $\mathbf{x}\mathbf{x}^\mathsf{T}$ and $\mathbf{x}$ is feasible for the original QCQP, hence the relaxation is exact. Of course, there are more potentially optimal solutions to the Shor relaxation (depending on the objective function), but these are convex combinations of dyadic optimal solutions. An example can be seen in Fig. 1 as the line connecting the two lower vertices in the lifted feasible set.

Algorithm 1: Solving copositive optimization problems.

```
Result: v*
set k = 1;
construct outer approximation C_k ⊇ COP(K) ;
repeat
    generate a feasible point for v(C_k) to obtain (S_k, y_k);
    check S_k ∈ COP(K);
    if S_k ∉ COP(K) then
        obtain certificate x_k ∈ K
    else
        x_k = o
    set C_{k+1} = C_k ∩ {(S, y) : x_k^T S x_k ≥ 0} ∩ C'_k    (using
    additional cuts via C'_k, see below)
until some stopping criterion is met;
```

Example 4. Consider the following QCQP and its Shor relaxation:

$$\left.\begin{array}{l}
\min_{\mathbf{x}\in\mathbb{R}^2} q_{11}x_1^2 + 2q_{12}x_1x_2 + q_{22}x_2^2\\
\quad \text{s.t.}: 3x_1^2 + 3x_2^2 + 2x_1x_2 \le 6\\
\qquad\quad 3x_1^2 + 3x_2^2 - 2x_1x_2 \le 6\\
\qquad\quad 2x_1^2 + x_2^2 \le 3
\end{array}\right\} \text{ and}$$

$$\left\{\begin{array}{l}
\min_{\mathsf{X}\in\mathcal{S}_+^2} q_{11}X_{11} + 2q_{12}X_{12} + q_{22}X_{22}\\
\quad \text{s.t.}: 3X_{11} + 3X_{22} + 2X_{12} \le 6\\
\qquad\quad 3X_{11} + 3X_{22} - 2X_{12} \le 6\\
\qquad\quad 2X_{11} + X_{22} \le 3\,.
\end{array}\right.$$

This example gives an instance where the Shor relaxation fails to be tight. We can see this by examining the extreme points of its feasible set. Ignoring the psd-constraint for a moment, intersections of three halfspaces can yield an extreme points only where the three associated hyperplanes meet. A simple calculation shows that it is the point $X_{11} = X_{22} = 1,\ X_{12} = 0$, hence $\mathsf{X} = \mathsf{I}$, the identity matrix. Clearly, this matrix is the interior of the psd-cone, so that it is a feasible solution for the Shor relaxation and indeed one of its extreme points. However, it is not a dyadic matrix as I always has full rank. We therefore can get an optimal solution for the Shor relaxation that has no corresponding solution in the original QCQP. More formally, we cannot find $(x_1, x_2) \in \mathbb{R}^2$ where $x_1^2 = x_2^2 = 1$ and $x_1x_2 = 0$.

Indeed, if we set the objective function coefficients $q_{11} = -8,\ q_{22} = -7$ and $q_{12} = 0$, the optimal value of the Shor relaxation is -15 attained at $\mathsf{X} = \mathsf{I}$, while the original QCQP attains its optimal value of -14 at $(x_1, x_2) = (0, \sqrt{2})$.

Nonetheless, there are choices for the objective function coefficients where the two problems give identical optimal values. To see this, let us identify the rest of the extreme points of $\mathcal{F}_{Shor}$. Note

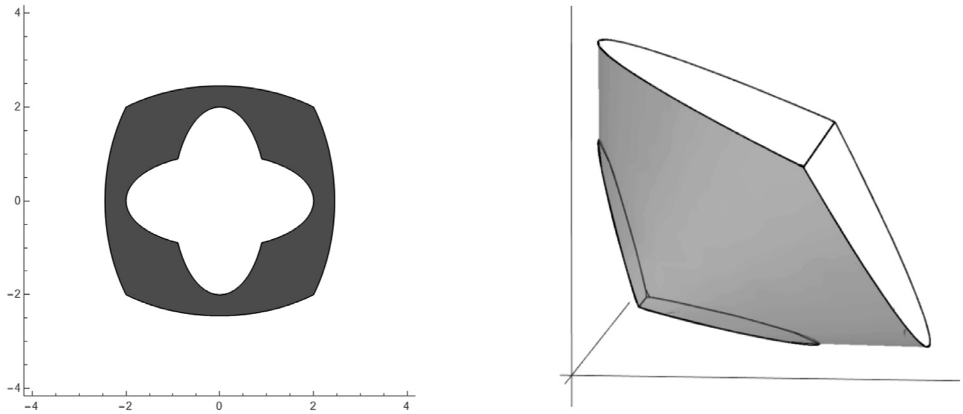

a) b)

Fig. 1. (a) The feasible set $\mathcal{F}$ of the QCQP in Example 3. (b) The feasible set $\mathcal{F}_{Shor}$ of its Shor relaxation. As a consequence of Theorem 2 below, $\mathcal{F}_{Shor}$ coincides with $\mathcal{G}(\mathcal{F})$; we show a projection of $\mathcal{G}(\mathcal{F})$, given by the map $(\mathbf{x}, \mathsf{X}) \mapsto (X_{11}, \sqrt{2}X_{12}, X_{22})^\mathsf{T}$ from $\mathbb{R}^2 \times \mathcal{S}^2$ to $\mathbb{R}^3$, illustrating the intersection of four half-spaces and the psd-cone.

that the matrices in

$$\mathcal{M} := \left\{ \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix}, \begin{bmatrix} 3 & -1 \\ -1 & 3 \end{bmatrix}, \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \right\}$$

are all positive-definite, so that the set $\mathcal{F}_\mathsf{M} := \left\{ \mathsf{X} \in \mathcal{S}^n_+ : \mathsf{M} \bullet \mathsf{X} = \alpha \right\}$ is nonempty, compact, conic intersections whenever $\mathsf{M} \in \mathcal{M}$ and $\alpha \geq 0$. The three linear inequalities are precisely of the form $\mathsf{M} \bullet \mathsf{X} \leq \alpha$ where M is one of the matrices in $\mathcal{M}$.

Now let us examine those extreme points of the $\mathcal{F}_{Shor}$ where none of these linear inequalities are binding. Hence, we ask for the extreme points of the psd-cone and the only one there is the zero matrix $\mathsf{O} = \mathbf{oo}^\mathsf{T}$. If only one linear constraint is active, the extreme points are those extreme points of $\mathcal{F}_\mathsf{M}$, with $\mathsf{M} \in \mathcal{M}$, which fulfill the other two inequalities in $\mathcal{F}_{Shor}$ strictly. But $\mathcal{F}_\mathsf{M}$ is a compact conic intersection, so that its extreme points are points in the intersection of the hyperplane with extreme rays of $\mathcal{S}^2_+$, i.e. rays spanned by dyadic matrices. Therefore they are themselves dyadic matrices. Finally, let us examine the extreme points that fulfill exactly two of the linear inequalities. The points that fulfill two of the inequalities must form either a line, a half line or a line segment that is a subset of $\mathcal{F}_\mathsf{M}$ for $\mathsf{M} \in \mathcal{M}$, but these are compact sets, so that they form a line segment, given by the intersection of $\mathcal{S}^2_+$ and a line. The extreme points of these sets are therefore the two points where the respective lines intersect the boundary of $\mathcal{S}^2_+$, which is entirely comprised of dyadic matrices. (Note that this is the case for the psd-cone $\mathcal{S}^2_+$ only, for $\mathcal{S}^n_+$ with $n > 2$ there are boundary points that are not dyadic. However, we will later see in Theorem 2 that the Shor relaxation is exact whenever only two inequality constraints are present, so that the argument would in fact stay valid if $n > 2$.)

In total, we see that all extreme points of $\mathcal{F}_{Shor}$ are dyadic except for the one we have identified as the identity matrix I. Thus, if we choose the objective function coefficients such that the optimal solution of the Shor relaxation is attained at a point other than I, then the relaxation will be tight. As an example for the latter case, let us consider $q_{11} = q_{22} = -1$ and $q_{21} = 0$. In this case, the optimal value of the QCQP is given by -2 attained at $(x_1, x_2) = (0, \sqrt{2})$. The Shor relaxation attains the same optimal value of -2 at $\mathsf{X} = \mathsf{I}$, but clearly this is not the only optimal point since the dyadic matrix formed from the optimal solution of the QCQP gives a feasible solution with the same optimal value, that is:

$$\mathsf{X} = \begin{bmatrix} 0 \\ \sqrt{2} \end{bmatrix} \begin{bmatrix} 0 \\ \sqrt{2} \end{bmatrix}^\mathsf{T} = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix},$$

which is feasible and yields the optimal value of -2. Thus, there is an optimal dyadic solution to the Shor relaxation, which is enough to eliminate the relaxation gap.

Let us summarize our observations. The dyadic matrices are at the boundary of $\mathcal{S}^n_+$ and for $n = 2$, this boundary is entirely comprised of dyadic matrices so that it is actually ext $\mathcal{CPP}(\mathbb{R}^2)$. However, since we consider the convex hull of the latter, i.e., $\mathcal{S}^2_+$, we produced an extreme point in the interior of $\mathcal{S}^2_+$, which thus is of rank greater than one. For certain choices of the objective function coefficients, there will therefore be a gap between the two optimization problems. On the bright side, we also see that even if we are far from describing $\mathcal{G}(\mathcal{F})$, the Shor relaxation can be exact for some choices of the objective function coefficients.

The above discussion makes it apparent that the Shor relaxation is not necessarily tight, since its feasible set $\mathcal{F}_{Shor}$ can have extreme points that are not dyadic matrices. Under additional assumptions, one can close the gap at least for the homogeneous case. To this end, we introduce the following geometric condition.

Condition 1. For a collection of matrices $\mathsf{Q}_i \in \mathcal{S}^n$ and real numbers b_i, $i \in [1\!:\!m]$ we say that Condition 1 holds if for any $\mathsf{X} \in \mathcal{S}^n_+$ with $\mathsf{Q}_i \bullet \mathsf{X} \leq \omega_i$ for all $i \in [1\!:\!m]$,

$$\mathsf{Q}_k \bullet \mathsf{X} < \omega_k \quad \text{for all } k \in [1\!:\!m] \setminus \{i, j\} \text{ whenever}$$
$$\mathsf{Q}_i \bullet \mathsf{X} = \omega_i \text{ and } \mathsf{Q}_j \bullet \mathsf{X} = \omega_j \text{ for } i \neq j\,.$$

The condition requires that for any feasible $\mathsf{X} \in \mathcal{S}^n$ at most two constraints can be binding at the same time. If $\mathcal{F}_{Shor} := \left\{ \mathsf{X} \in \mathcal{S}^n_+ : \mathsf{Q}_i \bullet \mathsf{X} \leq \omega_i, i \in [1\!:\!m] \right\}$ is bounded (as assumed in Theorem 2), one can check Condition 1 by solving $(m^3 - 3m^2 + 2m)/6$ semidefinite optimization problems of the form $\sup_{\mathsf{X} \in \mathcal{F}_{Shor}} \left\{ \mathsf{Q}_k \bullet \mathsf{X} - \omega_k : \mathsf{Q}_i \bullet \mathsf{X} = \omega_i,\ \mathsf{Q}_j \bullet \mathsf{X} = \omega_j \right\}$. For Condition 1 to hold, all the optimal values must be strictly smaller than 0. Note that $\mathcal{K} = \mathbb{R}^n$ here.

Theorem 2. *Suppose that Condition 1 holds for the matrices* $\mathsf{Q}_i \in \mathcal{S}^n$ *and real numbers* $\omega_i \in \mathbb{R}$, $i \in [1\!:\!m]$. *Further, suppose that the set* $\mathcal{F}_{Shor} := \left\{ \mathsf{X} \in \mathcal{S}^n_+ : \mathsf{Q}_i \bullet \mathsf{X} \leq \omega_i,\ i \in [1\!:\!m] \right\}$ *is bounded. Then*

$$\inf_{\mathbf{x} \in \mathbb{R}^n} \left\{ \mathbf{x}^\mathsf{T} \mathsf{Q}_0 \mathbf{x} : \mathbf{x}^\mathsf{T} \mathsf{Q}_i \mathbf{x} \leq \omega_i,\ i \in [1\!:\!m] \right\}$$
$$= \inf_{\mathsf{X} \in \mathcal{S}^n_+} \{ \mathsf{Q}_0 \bullet \mathsf{X} : \mathsf{Q}_i \bullet \mathsf{X} \leq \omega_i,\ i \in [1\!:\!m] \}.$$

Proof. See Bomze & Gabl (2021). □

While Theorem 1 and the results above clarify the role of $\mathcal{G}(\mathcal{F})$ for optimization problems, explicit characterizations of the set $\mathcal{G}(\mathcal{F})$ have been given in terms of $\mathcal{F}_{Shor}$ and additional cuts. The respective results are summarized in the following theorem:

Theorem 3. *Consider the following feasible sets of QCQPs:*

- $\mathcal{F}_1 := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| \le 1, \mathsf{A}\mathbf{x} \le \mathbf{b}\}$, *with* $\mathsf{A} \in \mathbb{R}^{m\times n}$, *where the m hyperplanes described by* $\mathsf{A}\mathbf{x} = \mathbf{b}$ *do not intersect inside the unit ball.*
- $\mathcal{F}_2 := \left\{\mathbf{x} \in \mathbb{R}^2 : \mathsf{A}\mathbf{x} \le \mathbf{b}\right\}$ *with* $\mathsf{A} \in \mathbb{R}^{3\times 2}$, $\mathbf{b} \in \mathbb{R}^3$ *such that* $\mathcal{F}_2$ *is a nondegenerate planar triangle.*
- $\mathcal{F}_3 := \left\{\mathbf{x} \in \mathbb{R}^2 : \mathsf{A}\mathbf{x} \le \mathbf{b}\right\}$ *with* $\mathsf{A} \in \mathbb{R}^{4\times 2}$, $\mathbf{b} \in \mathbb{R}^4$ *such that* $\mathcal{F}_3$ *is a nondegenerate planar quadrangle.*

Let $\mathbf{a}_i$ *be the ith row of* A. *Then*

- $\mathcal{G}(\mathcal{F}_1) = \left\{\mathsf{Y}(\mathbf{x},\mathsf{X}) \in \mathcal{S}^{n+1}_+ : \mathrm{trace}(\mathsf{X}) \le 1, \begin{array}{l} \|b_i\mathbf{x} - \mathsf{X}\mathbf{a}_i\| \le b_i - \mathbf{a}_i^\mathsf{T}\mathbf{x}, \ i \in [1\!:\!m], \\ b_i\mathbf{a}_j^\mathsf{T}\mathbf{x} + b_j\mathbf{a}_i^\mathsf{T}\mathbf{x} - \mathbf{a}_i^\mathsf{T}\mathsf{X}\mathbf{a}_j \le b_ib_j, \ (i,j) \in [1\!:\!m]^2 \end{array}\right\}$,
- $\mathcal{G}(\mathcal{F}_2) = \left\{\mathsf{Y}(\mathbf{x},\mathsf{X}) \in \mathcal{S}^3_+ : b_i\mathbf{a}_j^\mathsf{T}\mathbf{x} + b_j\mathbf{a}_i^\mathsf{T}\mathbf{x} - \mathbf{a}_i^\mathsf{T}\mathsf{X}\mathbf{a}_j \le b_ib_j, \ (i,j) \in [1\!:\!3]^2\right\}$
- $\mathcal{G}(\mathcal{F}_3) = \left\{\mathsf{Y}(\mathbf{x},\mathsf{X}) \in \mathcal{S}^3_+ : b_i\mathbf{a}_j^\mathsf{T}\mathbf{x} + b_j\mathbf{a}_i^\mathsf{T}\mathbf{x} - \mathbf{a}_i^\mathsf{T}\mathsf{X}\mathbf{a}_j \le b_ib_j, \ (i,j) \in [1\!:\!4]^2\right\}$.

Proof. The characterizations are due to Anstreicher & Burer (2010); Burer & Anstreicher (2013) respectively, the characterizations of $\mathcal{G}(\mathcal{F}_1)$ with no, or just a single linear inequality, go back to Sturm & Zhang (2003); Yakubovich (1971). □

So far we only considered examples where $\mathcal{K} = \mathbb{R}^n$, so that the Shor relaxation took the form of a positive-semidefinite optimization problem. For other choices of $\mathcal{K}$, one leaves this familiar territory and is confronted with optimizing over $\mathcal{CPP}(\mathcal{K})$, a potentially much harder task. However, conceptually much of the intuition we garnered so far stays intact: unless the Shor relaxation produces extreme points that are not dyadic matrices, the relaxation gap vanishes. Such higher-rank extreme points may either arise from the interaction of the linear constraints with each other inside the interior of the respective matrix cone, or from the interaction of these constraints with the boundary of the said cone. The following example demonstrates an instance of the latter.

Example 5. Now consider another pair of QCQP and its Shor relaxation:

$$\left.\begin{array}{l} \min_{\mathbf{x}\in\mathbb{R}^2_+} q_{11}x_1^2 + 2q_{12}x_1x_2 + q_{22}x_2^2 \\ \quad \text{s.t.}: 3x_1^2 + x_2^2 = 6 \\ \qquad\quad x_1^2 + 3x_2^2 = 6 \end{array}\right\} \quad \text{and}$$

$$\left\{\begin{array}{l} \min_{\mathsf{X}\in\mathcal{CPP}(\mathbb{R}^2_+)} q_{11}X_{11} + 2q_{12}X_{12} + q_{22}X_{22} \\ \quad \text{s.t.}: 3X_{11} + X_{22} = 6 \\ \qquad\quad X_{11} + 3X_{22} = 6\,. \end{array}\right.$$

The system of quadratic equations has exactly one non-negative solution which is $\mathbf{x} := \left[\sqrt{\frac{3}{2}}, \sqrt{\frac{3}{2}}\right]^\mathsf{T}$. In fact the point $\mathbf{x}\mathbf{x}^\mathsf{T}$ is an extreme point of the feasible set of the Shor relaxation. But unfortunately there is another one. Consider the fact that $\mathcal{CPP}\left(\mathbb{R}^2_+\right) = \mathcal{S}^2_+ \cap \mathcal{N}_2 = \left\{\mathsf{X} \in \mathcal{S}^n_+ : X_{12} \ge 0\right\}$. The matrix that fulfills the equality constraints and $X_{12} = 0$, i.e. $\sqrt{\frac{3}{2}}\,\mathsf{I}$, is also an extreme point, and the feasible set is in fact the convex hull of those two extreme points.

Again, it is the convexification that gives rise to this spurious, extreme point. Indeed, had we only used $\mathrm{ext}\left(\mathcal{CPP}(\mathbb{R}^2_+)\right)$, the additional extreme point would not have appeared since the identity matrix I is not dyadic, and therefore does not span a ray in $\mathrm{ext}\left(\mathcal{CPP}(\mathbb{R}^2_+)\right)$. But by shifting to the convex hull of the latter cone, i.e. $\mathcal{CPP}(\mathbb{R}^2_+)$, we gave rise to an extreme point in the feasible set of the relaxation of rank 2, and therefore has no match in the feasible set of the original quadratic problem. Note, that compared to the previous example, this time it is not the constellation of the linearized inequalities, but the geometry of the convex hull of $\mathrm{ext}\left(\mathcal{CPP}(\mathbb{R}^2_+)\right)$, namely $\mathcal{CPP}\left(\mathbb{R}^2_+\right) = \mathcal{S}^2_+ \cap \mathcal{N}_2$ that generated the problem.

The above example demonstrates that the complex geometry of $\mathcal{CPP}(\mathcal{K})$ may present a formidable challenge if one seeks to close the relaxation gap. In the following section we introduce a powerful machinery that meets this challenge by exploiting this very geometry in an elegant way.

2.1.2. Burer's convex reformulation of a large class of QCQPs

One of the most celebrated examples of an application of Theorem 1 is Burer's completely positive reformulation of a quite large class of QCQPs:

Theorem 4. *Let* $\mathcal{K} \subseteq \mathbb{R}^n$ *be a closed, convex cone and let* $\mathcal{L} := \{\mathbf{x} \in \mathcal{K} : \mathsf{A}\mathbf{x} = \mathbf{b}\}$ *be nonempty so that* $\mathcal{L}_\infty = \{\mathbf{x} \in \mathcal{K} : \mathsf{A}\mathbf{x} = \mathbf{o}\}$ *is its recession cone. Further let* $\mathsf{Q}_i \in \mathcal{S}^n$, $\mathbf{q}_i \in \mathbb{R}^n$, $i \in [1\!:\!l]$, *and define* $\mathcal{B} := \left\{j : \mathsf{Q}_i\mathbf{e}_j \ne \mathbf{o} \text{ or } \mathbf{q}_i^\mathsf{T}\mathbf{e}_j \ne 0 \text{ for some } i \in [1\!:\!l]\right\}$. *Assume that*

(a) $\mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} \ge \omega_i$ for all $\mathbf{x} \in \mathcal{L}$ and $i \in [1\!:\!l]$, *and*
(b) $\mathbf{d} \in \mathcal{L}_\infty \Rightarrow d_j = 0$ *for all* $j \in \mathcal{B}$.

Then, any feasible QCQP of the form

$$\min_{x\in\mathcal{K}} \left\{\mathbf{x}^\mathsf{T}\mathsf{Q}_0\mathbf{x} + 2\mathbf{q}_0^\mathsf{T}\mathbf{x} : \mathsf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} + \mathbf{q}_i^\mathsf{T}\mathbf{x} = \omega_i\,,\ i \in [1\!:\!l]\right\}$$

is equivalent to

$$\begin{aligned} &\min_{\mathbf{x},\mathsf{X}} \mathsf{Q} \bullet \mathsf{X} + \mathbf{q}^\mathsf{T}\mathbf{x} \\ &\quad \text{s.t.}: \mathsf{A}\mathbf{x} = \mathbf{b}\,, \\ &\mathrm{Diag}\left(\mathsf{A}\mathsf{X}\mathsf{A}^\mathsf{T}\right) = \mathbf{b} \circ \mathbf{b}, \\ &\mathsf{Q}_i \bullet \mathsf{X} + \mathbf{q}_i^\mathsf{T}\mathbf{x} = \omega_i\,,\ i \in [1\!:\!l]\,, \\ &\begin{bmatrix} 1 & \mathbf{x}^\mathsf{T} \\ \mathbf{x} & \mathsf{X} \end{bmatrix} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}). \end{aligned}$$

Proof. See Burer (2009) for the original proof for $\mathcal{K} = \mathbb{R}^n_+$, see Eichfelder & Povh (2013) for the proof under the assumption that $\mathcal{K}$ is a norm-cone. Considering the results in Kim, Kojima, & Toh (2020), which we will discuss shortly, no such assumption is needed for the theorem to hold. □

Together, assumptions a) and b) are colloquially referred to as the *key assumption*. Note, that it is met, for example, by the constraint $x_i - x_i^2 = 0$ (hence x_i is a binary variable) if $\mathbf{x} \in \mathcal{L} \Rightarrow x_i \in [0,1]$, which can always be achieved by adding $x_i + s = 1,\ s, x_i \in \mathbb{R}_+$ to the description of the feasible set, where s acts as a slack variable. In case $x_i, x_j \in \mathbb{R}_+$ the complementarity constraint $x_ix_j = 0$ also fulfills the key assumption if both variables are bounded over $\mathcal{L}$.

The description of the linear portion of the completely positive reformulation can be modified without changing the feasible set. The following proposition summarizes the most important reformulations, all of which will appear later in our discussion.

Proposition 5. *Suppose* $\mathsf{Y} = \mathsf{Y}(\mathbf{x},\mathsf{X}) \in \mathcal{S}^{n+1}_+$ *and define* $\mathsf{M} := [\mathbf{b}, -\mathsf{A}]$ *to be the matrix containing* $\mathbf{b} \in \mathbb{R}^n$ *and the columns of* $-\mathsf{A} \in \mathbb{R}^{m\times n}$ *concatenated. Then the following are equivalent:*

(i) $\mathsf{A}\mathbf{x} = \mathbf{b}$ *and* $\operatorname{diag}\left(\mathsf{A}\mathsf{X}\mathsf{A}^\mathsf{T}\right) = \mathbf{b} \circ \mathbf{b}$;
(ii) $\mathsf{M}\mathsf{Y}\mathsf{M}^\mathsf{T} = \mathsf{O}$;
(iii) $\mathsf{M}\mathsf{Y} = \mathsf{O}$.

Proof. See Burer (2012, Proposition 3). □

The original proof of the theorem is quite algebraic and seems somewhat removed from the simple, geometric motivation of Theorem 1. Fortunately (Kim et al., 2020) recently provided a geometrical perspective on the subject. The concepts they introduce are quite versatile and allow proofs for generalizations of Theorem 4 as well as exactness proofs for relaxations of polynomial optimization problems. The theorems presented in the remainder of this section are simplified (and thus less powerful) versions of results in Kim et al. (2020) for presentational reasons. Also, they will be strong enough to prove a weaker version of Theorem 4, under the additional assumption that $\mathcal{L}$ is bounded.

We start out by investigating a more general question. Let $\mathbb{V}$ be a vector space of dimension n. For a (possibly nonconvex) cone $\mathbb{K} \subseteq \mathbb{V}$, and vectors $\mathsf{Q}, \mathsf{H}_0 \in \mathbb{V}$ and a convex set $\mathbb{J} \subseteq \operatorname{conv}(\mathbb{K})$, we want to know which conditions establish the equality:

$$\min_{\mathsf{X}\in\mathbb{V}} \{\langle \mathsf{Q}, \mathsf{X}\rangle : \mathsf{X} \in \mathbb{K}\cap\mathbb{J},\ \langle \mathsf{H}_0, \mathsf{X}\rangle = 1\} = \min_{\mathsf{X}\in\mathbb{V}} \{\langle \mathsf{Q}, \mathsf{X}\rangle : \mathsf{X} \in \mathbb{J},\ \langle \mathsf{H}_0, \mathsf{X}\rangle = 1\}.$$

Defining $\mathbb{H} := \{\mathsf{X} : \langle \mathsf{H}_0, \mathsf{X}\rangle = 1\} \subseteq \mathbb{V}$, we can equivalently ask for conditions for the equality

$$\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}) = \mathbb{H}\cap\mathbb{J}.$$

The following theorem gives an answer based on convex geometry.

Theorem 6. *For* $\mathbb{H}, \mathbb{K}, \mathbb{J}$ *as above, assume that* $\mathbb{H}\cap\mathbb{J} \neq \emptyset$ *is bounded and that* $\mathbb{J}$ *is a face of* $\operatorname{conv}(\mathbb{K})$*. Then* $\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}) = \mathbb{H}\cap\mathbb{J}$.

Proof. See Kim et al. (2020). □

This theorem motivates the search for a condition that lets us identify faces of convex cones, which are provided in the following theorem.

Theorem 7. *Assume that* $\mathbb{J} = \{\mathsf{X} \in \operatorname{conv}(\mathbb{K}) : \langle \mathsf{Q}_i, \mathsf{X}\rangle = 0,\ i \in [0\!:\!m]\}$ *and define*

$$\mathbb{J}_p := \{\mathsf{X} \in \operatorname{conv}(\mathbb{K}) : \langle \mathsf{Q}_i, \mathsf{X}\rangle = 0,\ i \in [0\!:\!p]\},$$

so that $\mathbb{J}_m = \mathbb{J}$ *and* $\mathbb{J}_{-1} = \operatorname{conv}(\mathbb{K})$*. If* $\mathsf{Q}_p \in \mathbb{J}_{p-1}^*$ for all $p \in [0\!:\!m]$ *then* $\mathbb{J}$ *is a face of* $\operatorname{conv}(\mathbb{K})$.

Proof. See Kim et al. (2020). □

Before we apply this machinery to convexify QCQPs, we will supply a small example for illustrating above theorems. The example itself is not immediately connected to QCQPs, but the geometric intuition it seeks to convey may further the understanding of the convexification strategy as a whole.

Example 6. Consider the nonconvex cone

$$\mathbb{K} := \left\{\lambda\begin{bmatrix}1\\0\\1\end{bmatrix} : \lambda \geq 0\right\} \cup \left\{\lambda\begin{bmatrix}-1\\0\\1\end{bmatrix} : \lambda \geq 0\right\} \cup \left\{\lambda\begin{bmatrix}0\\1\\0\end{bmatrix} : \lambda \geq 0\right\} \subset \mathbb{R}^3,$$

which is the union of three half-rays emanating from the origin in three different directions, two of which form a "V" in the xz-plane and the other one covers half of the y-axis. The intersection of $\mathbb{K}$ with the hyperplane

$$\mathbb{H} := \left\{\mathbf{x} \in \mathbb{R}^3 : x_3 = 1\right\},$$

which is a plane parallel to the xy-plane at height 1, are the points in

$$\mathbb{K}\cap\mathbb{H} = \left\{\begin{bmatrix}1\\0\\1\end{bmatrix}, \begin{bmatrix}-1\\0\\1\end{bmatrix}\right\},$$

which is a nonconvex set. We now want to find a convex set $\mathbb{J} \subseteq \operatorname{conv}(\mathbb{K})$ so that $\mathbb{H}\cap\mathbb{J} = \operatorname{conv}(\mathbb{K}\cap\mathbb{H})$. We claim that desired set is

$$\mathbb{J} := \{\mathbf{x} \in \operatorname{conv}(\mathbb{K}) : x_2 = 0\} = \left\{\lambda_1\begin{bmatrix}1\\0\\1\end{bmatrix} + \lambda_2\begin{bmatrix}-1\\0\\1\end{bmatrix} : \lambda_1, \lambda_2 \geq 0\right\}. \tag{7}$$

To see this, let us first check that $\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}) = \operatorname{conv}(\mathbb{H}\cap\mathbb{K})$, which follows by merely showing that $\mathbb{H}\cap\mathbb{K}\cap\mathbb{J} = \mathbb{H}\cap\mathbb{K}$. Clearly $\mathbb{H}\cap\mathbb{K}\cap\mathbb{J} \subseteq \mathbb{H}\cap\mathbb{K}$, but also

$$\mathbb{J}\cap\mathbb{K} = \left\{\lambda\begin{bmatrix}1\\0\\1\end{bmatrix} : \lambda \geq 0\right\} \cup \left\{\lambda\begin{bmatrix}-1\\0\\1\end{bmatrix} : \lambda \geq 0\right\},$$

which contains $\mathbb{H}\cap\mathbb{K}$, so that the desired equivalence is obvious. Now we can use Theorem 6 to establish $\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}) = \mathbb{H}\cap\mathbb{J}$. We see that $\mathbb{H}\cap\mathbb{J}$ is bounded since

$$\left(\lambda_1\begin{bmatrix}1\\0\\1\end{bmatrix} + \lambda_2\begin{bmatrix}-1\\0\\1\end{bmatrix}\right)_3 = \lambda_1 + \lambda_2 = 1\,,\ \ \lambda_1, \lambda_2 \geq 0\ \Rightarrow \lambda_i \in [0,1]\,,\ \ i = 1, 2\,, \tag{8}$$

so that all that is left to show is that $\mathbb{J}$ is a face of $\operatorname{conv}(\mathbb{K})$. We have that $\mathbf{x} \in \operatorname{conv}(\mathcal{K})$ implies that $x_2 \geq 0$ so that $\mathbb{J}$ is such a face by Theorem 7. Geometrically, it is the convex hull of the two "legs" of $\mathbb{K}$ that point the z-direction. It is also an exposed face of $\operatorname{conv}(\mathbb{K})$, where the exposing hyperplane is described by $x_2 = 0$.

Let us convince ourselves that the conclusion of the procedure is actually true. It is immediate that

$$\operatorname{conv}(\mathbb{K}\cap\mathbb{H}) = \left\{\lambda\begin{bmatrix}1\\0\\1\end{bmatrix} + (1-\lambda)\begin{bmatrix}-1\\0\\1\end{bmatrix} : \lambda \in [0,1]\right\},$$

on the other hand, in this simple example, (8) already tells us that $\mathbb{H}\cap\mathbb{J}$ is the same set.

We can use this simple setup to test the conditions of Theorem 6. First let us study a failure of boundedness of $\mathbb{H}\cap\mathbb{J}$, which we can construct by choosing $\mathbb{J} = \operatorname{conv}(\mathbb{K})$. In this case $\mathbb{J}$ is still a (trivial) face of $\operatorname{conv}(\mathbb{K})$ but

$$\mathbb{H}\cap\mathbb{J} = \operatorname{conv}(\mathbb{K}\cap\mathbb{H}) + \left\{\lambda\begin{bmatrix}0\\1\\0\end{bmatrix} : \lambda \geq 0\right\} \supset \operatorname{conv}(\mathbb{H}\cap\mathbb{K}),$$

hence, we get a strictly bigger set than the desired convex hull. Now, let us consider a slightly enlarged version of the $\mathbb{J}$ defined in (7) given by

$$\mathbb{J} := \left\{\lambda_1\begin{bmatrix}1\\0\\1\end{bmatrix} + \lambda_2\begin{bmatrix}-1\\0\\1\end{bmatrix} + \lambda_3\begin{bmatrix}0\\1\\1\end{bmatrix} : \lambda_1, \lambda_2, \lambda_3 \geq 0\right\},$$

for which we can easily check both $\mathbb{J} \subseteq \operatorname{conv}(\mathbb{K})$ and $\mathbb{H}\cap\mathbb{K}\cap\mathbb{J} = \mathbb{H}\cap\mathbb{K}$. Also, boundedness of $\mathbb{H}\cap\mathbb{J}$ is immediate from an argument analogous to (8). However, $\mathbb{J}$ is no longer a face of $\operatorname{conv}(\mathbb{K})$ and in fact

$$\mathbb{H}\cap\mathbb{J} = \left\{\lambda_1\begin{bmatrix}1\\0\\1\end{bmatrix} + \lambda_2\begin{bmatrix}-1\\0\\1\end{bmatrix} + \lambda_3\begin{bmatrix}0\\1\\1\end{bmatrix} : \lambda_1 + \lambda_2 + \lambda_3 = 1,\ \lambda_1, \lambda_2, \lambda_3 \geq 0\right\},$$

so that, again, the conclusion of the theorem is not sustained.

Finally we would like to point out that the present example is not entirely unrelated to QCQPs. Consider again Example 1 with $\mathbf{a} = [1,1]^\mathsf{T}$, $\mathbf{b} = [-1,1]^\mathsf{T}$. Then the feasible set of (3) can be described as a conic intersection given by

$$\left\{\begin{bmatrix}x_1\\x_2\end{bmatrix} \in \mathbb{R}^2 : x_2 = 1\,, \mathbf{x} \in \mathcal{K} := \left\{\lambda\begin{bmatrix}1\\1\end{bmatrix},\ \lambda \geq 0\right\} \cup \left\{\lambda\begin{bmatrix}-1\\1\end{bmatrix},\ \lambda \geq 0\right\}\right\}.$$

The Shor relaxation is then given by

$$\left\{\begin{bmatrix} x_0 & x \\ x & X \end{bmatrix} \in \mathcal{S}^2 : x_0 = 1,\ \mathsf{Y}(x,X) \in \mathcal{CPP}(\mathcal{K})\right\}, \tag{9}$$

which can certainly be made exact by replacing $\mathcal{CPP}(\mathcal{K})$ with $\operatorname{ext}(\mathcal{CPP}(\mathcal{K}))$. However, the preceeding discussion allows us to conclude that the Shor relaxation is tight anyways. We merely have to consider isomorphism $\pi : \mathcal{S}^2 \to \mathbb{R}^3$ given by

$$\pi\left(\begin{bmatrix} x_0 & x \\ x & X \end{bmatrix}\right) \mapsto \begin{bmatrix} X \\ x \\ x_0 \end{bmatrix}, \tag{10}$$

to see that $\pi(\operatorname{ext}(\mathcal{CPP}(\mathcal{K})))$ is essentially $\mathbb{K}$ where the third leg, which was spurious for the derivation of the convexification, got removed. Also, the hyperplane spanned by $x_0 = 1$ corresponds to $\pi^{-1}(\mathbb{H})$. Finally, removing the constraint $x_0 = 1$ from the set in (9) leaves us with $\pi^{-1}(\mathbb{J})$, so that the set itself is the inverse image $\pi^{-1}(\mathbb{H} \cap \mathbb{J})$ and therefore represents the exact convexification of the feasible set of our underlying QCQP.

To see how this is relevant for convex reformulations of QCQPs, consider the following simple reformulation:

$$\min_{\mathbf{x}\in\mathbb{R}^n}\left\{\mathbf{x}^\mathsf{T}\mathsf{Q}_0\mathbf{x} + 2\mathbf{q}_0^\mathsf{T}\mathbf{x} - \omega_0 : \mathsf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \in \mathcal{K},\ \mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} + \mathbf{q}_i^\mathsf{T}\mathbf{x} - \omega_i = 0,\ i \in [1\!:\!m]\right\}$$
$$= \min_{\mathsf{Y}\in\mathcal{S}^{n_1}}\left\{\bar{\mathsf{Q}}_0 \bullet \mathsf{Y} : \mathsf{H}_0 \bullet \mathsf{Y} = 1,\ \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m],\ \mathsf{Y} \in \left\{\mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K}\right\}\right\}$$

where as in Proposition 5, we have $\mathsf{M} = [\mathbf{b}, -\mathsf{A}]$ and

$$\bar{\mathsf{Q}}_0 := \begin{bmatrix} -\omega_0 & \mathbf{q}_0^\mathsf{T} \\ \mathbf{q}_0 & \mathsf{Q}_0 \end{bmatrix}, \quad \widehat{\mathsf{Q}}_0 := \mathsf{M}^\mathsf{T}\mathsf{M}$$

$$\widehat{\mathsf{Q}}_i := \begin{bmatrix} -\omega_i & \mathbf{q}_i^\mathsf{T} \\ \mathbf{q}_i & \mathsf{Q}_i \end{bmatrix}, i \in [1\!:\!m], \quad \mathsf{H}_0 := \begin{bmatrix} 1 & \mathbf{o}^\mathsf{T} \\ \mathbf{o} & \mathsf{O} \end{bmatrix}.$$

The final result has the desired form with $\mathbb{K} = \left\{\mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K}\right\}$, $\mathbb{H} = \left\{\mathsf{Y} \in \mathcal{S}^{n+1} : \mathsf{H}_0 \bullet \mathsf{Y} = 1\right\}$. We actually have $\left\{\mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K}\right\} = \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}))$, as explained in the proof of Proposition 11 in the appendix. Hence, by Theorem 6 we can show the equivalence

$$\min_{\mathsf{Y}}\left\{\bar{\mathsf{Q}}_0 \bullet \mathsf{Y} : \mathsf{H}_0 \bullet \mathsf{Y} = 1,\ \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m],\ \mathsf{Y} \in \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}))\right\}$$
$$= \min_{\mathsf{Y}}\left\{\bar{\mathsf{Q}}_0 \bullet \mathsf{Y} : \mathsf{H}_0 \bullet \mathsf{Y} = 1,\ \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m],\ \mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})\right\},$$

if we can show that

$$\mathbb{J} = \left\{\mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}) : \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m]\right\}$$

is a face of $\operatorname{conv}\mathbb{K} = \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})$ and that $\mathbb{J} \cap \mathbb{H} \neq \emptyset$ is bounded. Also, from $\operatorname{conv}(\mathbb{H} \cap \mathbb{K} \cap \mathbb{J}) = \mathbb{H} \cap \mathbb{J}$ we see that $\mathcal{G}(\mathcal{F}) = \mathbb{J} \cap \mathbb{H}$ where $\mathcal{F}$ is the feasible set of the underlying QCQP.

Thus, we can describe a recipe for the characterization of $\mathcal{G}(\mathcal{F})$ by the following steps:

Step 1: Given a feasible set $\mathcal{F} = \left\{\mathsf{Y} : \mathsf{H}_0 \bullet \mathsf{Y} = 1, \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m],\ \mathsf{Y} \in \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}))\right\}$, set

$$\mathbb{K} = \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})),$$
$$\mathbb{J} = \left\{\mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}) : \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m]\right\} \text{ and}$$
$$\mathbb{H} = \{\mathsf{Y} \in \mathcal{S}^n : \mathsf{H}_0 \bullet \mathsf{Y} = 1\}.$$

Step 2: Show that $\mathbb{J} \cap \mathbb{H}$ is bounded.

Step 3: Show that, perhaps after a reordering, we have for all $p \in [0\!:\!m]$

$$\widehat{\mathsf{Q}}_p \bullet \mathsf{Y} \geq 0 \quad \text{for all } \mathsf{Y} \in \mathbb{J}_{p-1} \text{ (i.e., } \widehat{\mathsf{Q}}_p \in \mathbb{J}_{p-1}^*)$$

with $\mathbb{J}_p := \left\{\mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}) : \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!p],\ \right\}$, $p \in [0\!:\!m]$ and $\mathbb{J}_{-1} = \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})$, and apply Theorem 7 to conclude that $\mathbb{J}$ is a face of $\operatorname{conv}\mathbb{K}$.

Step 4: Conclude that $\mathcal{G}(\mathcal{F}) = \mathbb{J} \cap \mathbb{H} = \left\{\mathsf{Y} : \mathsf{H}_0 \bullet \mathsf{Y} = 1, \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} = 0,\ i \in [0\!:\!m],\ \mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})\right\}$, by Theorem 6.

As a reference and illustration we will prove a special case of Theorem 4, where the feasible set is bounded, using this recipe in the appendix.

2.1.3. Unions of feasible sets and subtractions of ellipsoids

Given a workable description of $\mathcal{G}(\mathcal{F}_i),\ i \in [1\!:\!k]$ it is always possible to derive characterizations of $\mathcal{G}(\cup_{i=1}^k \mathcal{F}_i)$ and it is also possible to give a characterization of $\mathcal{G}(\mathcal{F}_1 \setminus \cup_{i=2}^k \operatorname{int}\mathcal{F}_i)$ in case $\mathcal{F}_i,\ i \in [2\!:\!k]$ are ellipsoids that fulfill certain regularity conditions. We summarize the respective procedures in the following two theorems.

Theorem 8. *Let $\mathcal{F}_i,\ i \in [1\!:\!k]$ be feasible sets of QCQPs and such that*

$$\mathcal{G}(\mathcal{F}_i) = \{\mathsf{X} \in \mathcal{S}^n : \mathsf{H} \bullet \mathsf{X} = 1,\ \mathscr{A}_i(\mathsf{X}) = \mathbf{o},\ \mathsf{X} \in \mathcal{C}_i\},\ i \in [1\!:\!k],$$

where for all $i \in [1\!:\!m]$, $\mathscr{A}_i : \mathcal{S}_n \to \mathbb{R}^m$ are appropriate linear operators and $\mathcal{C}_i$ are appropriate convex matrix cones. Further, assume $\mathsf{H} \bullet \mathsf{X} > 0$ whenever, for at least one $i \in [1\!:\!k]$, we have $\mathsf{X} \in \mathcal{C}_i$ and $\mathscr{A}_i(\mathsf{X}) = \mathbf{o}$. Then

$$\mathcal{G}(\cup_{i=1}^k \mathcal{F}_i) = \left\{\mathsf{X} = \sum_{i=1}^k \mathsf{X}_i : \mathsf{H} \bullet \left(\sum_{i=1}^k \mathsf{X}_i\right) = 1,\ \mathscr{A}_i(\mathsf{X}_i) = \mathbf{o},\ \mathsf{X}_i \in \mathcal{C}_i,\ i \in [1\!:\!k]\right\}.$$

Proof. The statement can be derived by leveraging results from disjunctive programming (Balas, 1979), but we give a short proof in the appendix. □

Theorem 9. *Let $\mathcal{F}_1$ be a feasible set of a QCQP set and let $\mathcal{F}_i = \left\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} \leq \omega_i\right\}$ be such that the inequalities $\mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} \geq \omega_i,\ i \in [1\!:\!m]$, induce non-intersecting holes into $\mathcal{F}_1$. Then $\mathcal{G}(\mathcal{F}_1 \setminus \cup_{i=2}^k \operatorname{int}\mathcal{F}_i) = \left\{(\mathbf{x}, \mathsf{X}) \in \mathcal{G}(\mathcal{F}_1) : \mathsf{Q}_i \bullet \mathsf{X} + 2\mathbf{q}_i^\mathsf{T}\mathbf{x} \geq \omega_i,\ i \in [1:m]\right\}$.*

Proof. See Yang, Anstreicher, & Burer (2016). □

These techniques have so far not been utilized for the sake of robust optimization, but they are simple and might be relevant for future research, which we want to foster with this article.

2.1.4. Open problems

The two-trustregion-subproblem (TTRS): The TTRS is the problem of minimizing a nonconvex quadratic function over a feasible set, say $\mathcal{F}$, described by two convex quadratic constraints. It is known to be solvable in polynomial time by an algorithm described in Bienstock (2016), which is, unfortunately, very impractical in the same way the ellipsoid method is impractical for solving linear problems. Recent work by Anstreicher (2022) provides more practical ways of solving it. However, despite substantial effort by the community (see e.g. Bomze & Overton, 2015) no description of $\mathcal{G}(\mathcal{F})$ is known. A discussion on the difficulties in this endeavor can be found in Yang & Burer (2013), partial results can be found in Bomze, Jeyakumar, & Li (2018); Jeyakumar et al. (2021). A description of $\mathcal{G}(\mathcal{F})$ would be highly appreciated by the community.

The key assumption in Burer's reformulation: The most limiting requirement in Theorem 4 is without a doubt the so-called key assumption. In light of the discussion in Section 2.1.2, specifically regarding the results in Kim et al. (2020), we already introduced some tools to relax part *a*) of the key assumption. It is only necessary that there is an order in which one can add quadratic equations to $\mathcal{L}$ so that every new quadratic function is non-negative over the set $\mathcal{L}$ intersected with the already introduced quadratic equations.

There are also recent results by Bomze & Peng (2022) on the relaxation of part *b*). In Theorem 9 we discussed the introduction of holes via quadratic inequalities. However, the question when a quadratic constraint can be added to the description of the feasible set without losing tightness of the set-completely positive reformulation remains, in general, an open one.

2.2. Duality of linear optimization over $\mathcal{CPP}(\mathcal{K})$

One of the decisive advantages of convex reformulations of QCQPs is that the resulting optimization problems enjoy the rich duality theory that convex optimization offers. General results on convex optimization duality, such as strong duality under Slater's condition, can be immediately applied to optimization problems involving $\mathcal{CPP}(\mathcal{K})$. For the readers' convenience we formulate a general linear completely positive optimization problem and its dual here, to review the conditions for full strong conic duality in the sequel.

So let

$$\begin{aligned}&\inf_{\mathsf{X}\in\mathcal{S}^n} \mathsf{Q}_0\bullet\mathsf{X}\\ &\text{s.t.}: \mathsf{Q}_i\bullet\mathsf{X}\le b_i,\quad i\in[1\!:\!m],\\ &\qquad \mathsf{X}\in\mathcal{CPP}(\mathcal{K}),\end{aligned}\tag{11}$$

then its dual is given by

$$\begin{aligned}&\sup_{\boldsymbol{\lambda}\in\mathbb{R}^m_+} -\sum_{i=1}^m b_i\lambda_i\\ &\text{s.t.}: \mathsf{Q}_0+\sum_{i=1}^m\lambda_i\mathsf{Q}_i\in\mathcal{COP}(\mathcal{K}).\end{aligned}\tag{12}$$

Here we use the definition

$$\begin{aligned}\mathcal{COP}(\mathcal{K}) := \mathcal{CPP}(\mathcal{K})^* &= \{\mathsf{M}\in\mathcal{S}^n : \mathsf{M}\bullet\mathsf{X}\ge 0 \quad\text{for all } \mathsf{X}\in\mathcal{CPP}(\mathcal{K})\}\\ &= \left\{\mathsf{M}\in\mathcal{S}^n : \mathbf{x}^\mathsf{T}\mathsf{M}\mathbf{x}\ge 0\quad\text{for all } \mathbf{x}\in\mathcal{K}\right\},\end{aligned}$$

where the second equality is valid since all the extreme rays of $\mathcal{CPP}(\mathcal{K})$ are of the form $\mathbf{x}\mathbf{x}^\mathsf{T}$ with $\mathbf{x}\in\mathcal{K}$. The cone $\mathcal{COP}(\mathcal{K})$ is called the set-copositive matrix cone, and can be thought of as a generalization of the positive-semidefinite matrix cone. It is a central object in our discussion and we provide a more thorough treatment of this subject in Section 2.3.1. We now state a well known theorem on strong duality between the two optimization problems.

Theorem 10. *For (11) and (12) we always have that* $\operatorname{val}(11)\ge\operatorname{val}(12)$. *Further,*

- *if (11) has a feasible point* $\mathsf{X}\in\operatorname{relint}\mathcal{CPP}(\mathcal{K})$ *then* $\operatorname{val}(11)=\operatorname{val}(12)$ *and (12) attains its optimal value,*
- *if (12) has a feasible point* $\boldsymbol{\lambda}\in\mathbb{R}^m_+$ *such that* $\mathsf{Q}_0+\sum_{i=1}^m\lambda_i\mathsf{Q}_i\in\operatorname{relint}\mathcal{COP}(\mathcal{K})$, *then* $\operatorname{val}(11)=\operatorname{val}(12)$, *and (11) attains its optimal value.*

An immediate consequence of the above theorem is that, without any assumptions, (12) offers a rigorous lower bound of any QCQP (1) whose Shor-relaxation is transformed into (11). This is of particular importance in situations where primal values are offered which are claimed to be nearly optimal.

2.3. Solving copositive optimization problems

The conic reformulations discussed so far introduce many of the comforts of convex optimization, most notably convex duality theory, to an area that is, in general, highly nonconvex. However, they do not alleviate the core difficulty of these problems in most cases: set-copositive and set-completely positive optimization problems are still NP-hard in general. But this does not mean that the convexification approach has no merit for solving the problems. We will now discuss two major routes by which the convexifications can be exploited in order to either solve the problem exactly or to give very good bounds.

2.3.1. Characterizations and inner/outer approximations of $\mathcal{CPP}(\mathcal{K})$ and $\mathcal{COP}(\mathcal{K})$

As stated before, in general, certifying membership in either $\mathcal{CPP}(\mathcal{K})$ or $\mathcal{COP}(\mathcal{K})$ is intractable save for some particular instances of $\mathcal{K}$. One justification for reformulating QCQPs into copositive optimization problems anyway is the fact that there are powerful approximations of these cones and in some cases even tractable characterizations. We will now discuss some of the more prominent and easily explained approximations and give some interesting references to more involved theory on the matter. Before we start this discussion, we want to provide some general and useful properties of the two cones: Most of them seem to be common knowledge within the community, so attributing historically correct credits is difficult. However, we believe the concise compilation may be of some use here, and for completeness we will provide a proof in the appendix.

Proposition 11. *For any cones* $\mathcal{K},\mathcal{K}_1,\mathcal{K}_2\subseteq\mathbb{R}^n$ *we have the following relations:*

1. $\mathcal{COP}(\mathcal{K})=\mathcal{COP}(-\mathcal{K})=\mathcal{COP}(\mathcal{K}\cup-\mathcal{K})$, *which also holds if* $\mathcal{COP}$ *is replaced with* $\mathcal{CPP}$,
2. *If* $\mathcal{K}_1\subseteq\mathcal{K}_2$, *then* $\mathcal{CPP}(\mathcal{K}_1)\subseteq\mathcal{CPP}(\mathcal{K}_2)$ *with equality if and only if* $\mathcal{K}_2\subseteq\mathcal{K}_1\cup-\mathcal{K}_1$,
3. *If* $\mathcal{K}_1\subseteq\mathcal{K}_2$, *then* $\mathcal{COP}(\mathcal{K}_1)\supseteq\mathcal{COP}(\mathcal{K}_2)$; *if in addition we assume* $\operatorname{int}\mathcal{K}_2^*\neq\emptyset$, *we have* $\mathcal{COP}(\mathcal{K}_1)=\mathcal{COP}(\mathcal{K}_2)$ *if and only if*[1] $\mathcal{K}_2\subseteq\operatorname{cl}\mathcal{K}_1$.
4. $\mathcal{CPP}(\mathcal{K})\subseteq\mathcal{S}^n_+\subseteq\mathcal{COP}(\mathcal{K})$; *all three sets are equal if and only if* $\mathcal{K}\cup-\mathcal{K}=\mathbb{R}^n$, *in particular*
5. $\mathcal{COP}(\mathbb{R}_+\times\mathbb{R}^m)=\mathcal{CPP}(\mathbb{R}_+\times\mathbb{R}^m)=\mathcal{S}^{m+1}_+$, *more generally,*
6. $\mathcal{CPP}(\mathcal{K}\times\mathbb{R}^m)=\left\{\begin{bmatrix}\mathsf{M}_{11} & \mathsf{M}_{21}^\mathsf{T}\\ \mathsf{M}_{21} & \mathsf{M}_{22}\end{bmatrix}\in\mathcal{S}^{m+n}_+ : \mathsf{M}_{11}\in\mathcal{CPP}(\mathcal{K})\right\}$ *if* $\mathbf{o}\in\mathcal{K}$,
7. $\mathcal{COP}(\mathcal{K}_1\cup\mathcal{K}_2)=\mathcal{COP}(\mathcal{K}_1)\cap\mathcal{COP}(\mathcal{K}_2)$,
8. $\mathcal{CPP}(\mathcal{K}_1\cup\mathcal{K}_2)=\mathcal{CPP}(\mathcal{K}_1)+\mathcal{CPP}(\mathcal{K}_2)$,
9. $\mathcal{CPP}(\operatorname{conv}\mathcal{K})\supseteq\mathcal{CPP}(\mathcal{K})$ *with equality if* $\mathcal{K}$ *is convex,*
10. $\mathcal{COP}(\operatorname{conv}\mathcal{K})\subseteq\mathcal{COP}(\mathcal{K})$ *with equality if* $\mathcal{K}$ *is convex,*
11. $\operatorname{int}\mathcal{CPP}(\mathcal{K})=\left\{\sum_{i=1}^k\mathbf{x}_i\mathbf{x}_i^\mathsf{T} : \mathbf{x}_i\in\operatorname{int}\mathcal{K},\ \operatorname{span}\{\mathbf{x}_1,\dots,\mathbf{x}_k\}=\mathbb{R}^n\right\}$ *if* $\mathcal{K}$ *is closed, convex and* $\operatorname{int}\mathcal{K}\neq\emptyset$,
12. $\mathcal{COP}(\mathcal{K})=\operatorname{cl}\mathcal{COP}(\mathcal{K})=\mathcal{COP}(\operatorname{cl}\mathcal{K})$ *while* $\mathcal{CPP}(\operatorname{cl}\mathcal{K})=\operatorname{cl}\mathcal{CPP}(\mathcal{K})$;
13. $\mathcal{COP}(\mathcal{K})=\mathcal{COP}(\operatorname{relint}\mathcal{K})$, *if* $\mathcal{K}$ *is convex,*
14. $\operatorname{int}\mathcal{COP}(\mathcal{K})=\left\{\mathsf{Q}\in\mathcal{S}^n : \mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x}>0\quad\text{for all } \mathbf{x}\in\mathcal{K}\setminus\{\mathbf{o}\}\right\}$.

Proof. See appendix. □

For the case of $\mathcal{K}=\mathbb{R}^n_+$ we have the following chain of inclusions

$$\mathcal{CPP}(\mathbb{R}^n_+)\subseteq\mathcal{S}^n_+\cap\mathcal{N}\subseteq\mathcal{S}^n_++\mathcal{N}\subseteq\mathcal{COP}(\mathbb{R}^n_+)\tag{13}$$

where $\mathcal{N}$ is the orthant of nonnegative matrices. The cone $\mathcal{S}^n_+\cap\mathcal{N}$ is often call the *doubly nonnegative matrix cone* $\mathcal{DNN}^n$, and $\mathcal{S}^n_++\mathcal{N}$ is often called the *nonnegative-decomposable matrix cone* $\mathcal{NND}^n$. Despite their conceptual simplicity, these cones often turn out to be quite powerful in practice. We will also discuss some impressive theoretical guarantees that involve these simple approximations later in this and other sections (see Theorem 13 and the succeeding discussion, but also Section 5.2).

[1] note that $\mathcal{K}_1\subseteq\mathcal{K}_2\subseteq\operatorname{cl}(\mathcal{K}_1\cup-\mathcal{K}_1)$ and $\operatorname{int}\mathcal{K}_2^*\neq\emptyset$ already implies $\mathcal{K}_2\subseteq\operatorname{cl}\mathcal{K}_1$, so that this criterion coincides with the criterion of 2. up to closure

For polyhedral cones $\mathcal{K} := \{\mathbf{x} \in \mathbb{R}^n : \mathsf{A}\mathbf{x} \leq \mathbf{o}\}$, with $\mathsf{A} \in \mathbb{R}^{m \times n}$, there are simple polyhedral approximations given by

$$\mathcal{PI}(\mathcal{K}) := \left\{\mathsf{A}^\mathsf{T}\mathsf{N}\mathsf{A} : \mathsf{N} \in \mathcal{N}^m\right\} \subseteq \mathcal{COP}(\mathcal{K})$$

$$\mathcal{PO}(\mathcal{K}) := \left\{\mathsf{M} \in \mathcal{S}^n : \mathsf{A}\mathsf{M}\mathsf{A}^\mathsf{T} \in \mathcal{N}^m\right\} \supseteq \mathcal{CPP}(\mathcal{K})$$

It is always possible to use Theorem 4 in order to characterize $\mathcal{CPP}(\bar{\mathcal{K}})$, where $\bar{\mathcal{K}} := \left\{(\mathbf{x}, \mathbf{s}) \in \mathbb{R}^n \times \mathbb{R}^m_+ : \mathsf{A}\mathbf{x} + \mathbf{s} = 0\right\}$ in order to derive $\mathcal{G}(\bar{\mathcal{K}}) = \mathcal{CPP}(\bar{\mathcal{K}})$ so that $\mathcal{CPP}(\mathcal{K})$ is the projection on the north-west $n \times n$ entries. The cone $\mathcal{CPP}(\bar{\mathcal{K}})$ is thereby described via linear constraints and a conic constraint involving $\mathcal{CPP}(\mathbb{R}^n \times \mathbb{R}^m_+)$. The latter constraint can then be reformulated via Proposition 11 point 6. where any approximation for $\mathcal{CPP}(\mathbb{R}^m_+)$ can be inserted in order to obtain inner and outer approximations of $\mathcal{CPP}(\mathcal{K})$. For the *second-order cone* case $\mathcal{K} = \mathcal{SOC}^n$, the celebrated S-Lemma (Yakubovich, 1971) allows for an exact characterization of both, the set-completely positive and the set-copositive matrix cone in terms of psd-constraints, namely

$$\mathcal{CPP}(\mathcal{SOC}^n) = \left\{\mathsf{M} \in \mathcal{S}^n_+ : \mathsf{M} \bullet \mathsf{J} \leq 0\right\},$$

$$\mathcal{COP}(\mathcal{SOC}^n) = \left\{\mathsf{M} \in \mathcal{S}^n : \mathsf{M} + \lambda\mathsf{J} \in \mathcal{S}^n_+,\ \lambda \geq 0\right\},$$

where J is the identity matrix up to the first entry in the first row, which is flipped to -1. Due to Proposition 11 point 2. we have $\mathcal{CPP}(\mathcal{SOC}^n) = \mathcal{CPP}(\mathcal{K})$ with $\mathcal{K} := \left\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\mathsf{T}\mathsf{J}\mathbf{x} \leq 0\right\}$, hence a cone described by a homogeneous quadratic inequality. For the case where multiple such inequalities are present, only limited results are available. For example, Bomze & Gabl (2021) proved the following theorem:

Theorem 12. *Let $\mathcal{K} := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\mathsf{T}\mathsf{Q}_i\mathbf{x} \leq 0,\ i \in [1\!:\!m]\}$ with $\mathsf{Q}_i \in \mathcal{S}^n$. Assume that there is some $\mathbf{x}_0$ with $\mathbf{x}_0^\mathsf{T}\mathsf{Q}_i\mathbf{x}_0 < 0$ for all $i \in [1\!:\!m]$. Further, suppose that for all $i \in [1\!:\!m]$*

$$\mathsf{X} \in \mathcal{S}^n_+ \setminus \{\mathsf{O}\} \quad \text{and} \quad \mathsf{Q}_i \bullet \mathsf{X} = 0 \Rightarrow \mathsf{Q}_j \bullet \mathsf{X} < 0 \quad \text{for all } j \in [1\!:\!m] \setminus \{i\}. \tag{14}$$

Then

$$\mathcal{COP}(\mathcal{K}) = \left\{\mathsf{M} : \mathsf{M} + \sum_{i=1}^m \lambda_i\mathsf{Q}_i \in \mathcal{S}^n_+ \text{ for some } \boldsymbol{\lambda} \in \mathbb{R}^m_+\right\},$$

$$\mathcal{CPP}(\mathcal{K}) = \left\{\mathsf{M} \in \mathcal{S}^n_+ : \mathsf{M} \bullet \mathsf{Q}_i \leq 0,\ i \in [1\!:\!m]\right\}.$$

The theorem does not cover the case where $\mathcal{K}$ is the intersection of (perhaps linearly transformed) second-order cones. A respective characterization of set-copositivity/set-completely positivity would provide a long desired convex reformulation of the multi-trustregion subproblem. So far, this remains an open problem, despite substantial effort by the community. Still, one may study (Yang & Burer, 2013) to find inspirations for approximations for instances of $\mathcal{K}$ that involve two second-order cone constraints.

In case $\mathcal{K} := \{\mathbf{x} \in \mathcal{SOC}^n : \mathsf{A}\mathbf{x} \leq \mathbf{o}\}$ where the hyperplanes encoded by the linear inequalities do not intersect within the second-order cone, one may use a homogeneous version of Theorem 3 (regarding $\mathcal{F}_1$) in order do derive a tractable characterization of $\mathcal{CPP}(\mathcal{K})$ and $\mathcal{COP}(\mathcal{K})$. However, Xu & Hanasusanto (2018) found an elegant way to neatly summarize approximations and exactness results for a slightly more general instance of $\mathcal{K}$.

Theorem 13. *Consider*

$$\mathcal{K} := \{\mathbf{x} \in \mathbb{R}^n \times \mathbb{R}_+ : \mathsf{A}\mathbf{x} \geq \mathbf{o},\ \ \mathsf{B}\mathbf{x} \in \mathcal{SOC}^r\}$$

where $\mathsf{A} \in \mathbb{R}^{p \times (n+1)}$ and $\mathsf{B} \in \mathbb{R}^{r \times (n+1)}$ and define

$$\mathcal{SI}(\mathcal{K}) := \left\{\mathsf{M} \in \mathcal{S}^{n+1} : \begin{array}{l} \mathsf{W} \in \mathcal{S}^{n+1}_+,\ \mathsf{U} \in \mathcal{N}^p, \\ \mathsf{V} \in \mathcal{S}^{n+1},\ \mathsf{T} \in \mathbb{R}^{p \times r},\ \lambda \in \mathbb{R}_+ \\ \mathsf{M} = \mathsf{W} + \lambda\mathsf{S} + \mathsf{A}^\mathsf{T}\mathsf{U}\mathsf{A} + \mathsf{V} \\ \mathsf{V} = \frac{1}{2}\left(\mathsf{A}^\mathsf{T}\mathsf{T}\mathsf{B} + \mathsf{B}^\mathsf{T}\mathsf{T}^\mathsf{T}\mathsf{A}\right),\ \mathrm{Rows}(\mathsf{T}) \in \mathcal{SOC}^r \end{array}\right\},$$

where $\mathsf{S} := \mathsf{B}^\mathsf{T}\mathbf{e}_1\mathbf{e}_1^\mathsf{T}\mathsf{B} - \sum_{i=2}^r \mathsf{B}^\mathsf{T}\mathbf{e}_i\mathbf{e}_i^\mathsf{T}\mathsf{B}$. Then $\mathcal{SI}(\mathcal{K}) \subseteq \mathcal{COP}(\mathcal{K})$. Further, equality holds under one of the following conditions:

- $\mathsf{A} = \mathsf{O}$, *hence no linear inequalities are present.*
- *If $\mathbf{x} \in \mathbb{R}^{n+1}$ satisfies $\mathsf{B}\mathbf{x} \in \mathcal{SOC}^r$ and $\mathbf{a}_i^\mathsf{T}\mathbf{x} = 0$ for some $i \in [1\!:\!p]$, then $\mathbf{x} \in \mathcal{K}$.*

Clearly the dual of $\mathcal{SI}(\mathcal{K})$ is an outer approximation of $\mathcal{CPP}(\mathcal{K})$, but we will not go through the effort of deriving it here. Instead, we want to comment on the philosophy behind its construction. Note that for any two convex cones $\mathcal{K}_1$ and $\mathcal{K}_2$ containing the origin we have

$$\mathcal{K}_1 + \mathcal{K}_2 = \mathrm{conv}(\mathcal{K}_1 \cup \mathcal{K}_2)\,. \tag{15}$$

Now, $\mathcal{SI}(\mathcal{K})$ is such a sum where the components consist of $\mathcal{S}^{n+1}_+$, an instance of $\mathcal{PI}(\mathcal{K})$, a single ray $\{\lambda\mathsf{S} : \lambda \geq 0\}$ and the fourth cone described in terms of V and T which differs from any of the previous inner approximations, but whose containment in $\mathcal{COP}(\mathcal{K})$ can be easily checked. Hence, whenever a new inner approximation is identified, one can combine it with all other inner approximations to obtain a potentially much stronger inner approximation. We want to highlight that due to (15), even adding a single ray may increase the size of the inner approximation substantially.

In addition, this inner approximation improves on another popular construction discussed in Ben-Tal, El Ghaoui, & Nemirovski (2009, Theorem B.3.1) where the authors propose the so-called *approximate S-Lemma,* which can be used to derive an alternative inner approximation of $\mathcal{COP}(\mathcal{K})$, with $\mathcal{K}$ as defined in Theorem 13. However, in Xu & Hanasusanto (2018, Proposition 3) it is demonstrated that $\mathcal{SI}(\mathcal{K})$ gives a superset of the approximations based on the approximate S-Lemma.

2.3.2. Algorithmic approaches via copositivity detection

Recently Badenbroek & de Klerk (2022) and Anstreicher & Gabl (2022) proposed algorithmic approaches to solve a copositive optimization problem where the ground cone is either $\mathbb{R}^n_+$ or a polyhedral cone, but it seems plausible that similar approaches are feasible for other ground cones $\mathcal{K} \subseteq \mathbb{R}^n$. We will give a high-level abstraction of their approaches here.

We consider a general set-copositive optimization problem given by

$$v^* = \sup_{\mathbf{y},\mathsf{S}} \left\{\mathbf{b}^\mathsf{T}\mathbf{y} : \mathsf{C} - \sum_{i=1}^m y_i\mathsf{A}_i = \mathsf{S},\ \ \mathsf{S} \in \mathcal{COP}(\mathcal{K})\right\}.$$

The algorithms are based on relaxed problems:

$$v(\mathcal{C}) := \sup_{\mathbf{y},\mathsf{S}} \left\{\mathbf{b}^\mathsf{T}\mathbf{y} : \mathsf{C} - \sum_{i=1}^m y_i\mathsf{A}_i = \mathsf{S},\ \ (\mathsf{S}, \mathbf{y}) \in \mathcal{C}\right\}.$$

where $\mathcal{C}$ is a convex set such that its projection on the S-coordinate contains $\mathcal{COP}(\mathcal{K})$ and over which we can optimize efficiently. If $v(\mathcal{C})$ attains its optimum at a point $(\mathsf{S}, \mathbf{y})$ such that $\mathsf{S} \in \mathcal{COP}(\mathcal{K})$ then $v(\mathcal{C}) = v^*$, and we solved the problem. If $\mathsf{S} \notin \mathcal{COP}(\mathcal{K})$ then there is a certificate $\mathbf{x} \in \mathcal{K}$ such that $\mathbf{x}^\mathsf{T}\mathsf{S}\mathbf{x} < 0$. We assume that we have an oracle that is capable of testing set-copostivity and produces a certificate in case of negative answer. The algorithm proceeds as follows:

The two papers employ different variations of this algorithm. Both have in common that in each iteration, set-copositivity of the iterate S_k is tested and the approximations $\mathcal{C}_k$ are updated via the cut generated by the certificate $\mathbf{x}_k$, in case the test result is negative. The algorithms differ in the generation of the feasible points S_k and $\mathbf{y}_k$, in the method by which copositivity is checked and in a set of additional cuts $\mathcal{C}_k'$, which we did not discuss so far.

In (Badenbroek & de Klerk, 2022) the authors deal with the case where $\mathcal{K} = \mathbb{R}^n_+$. The feasible points are generated by finding the analytic center of the feasible set of $v(\mathcal{C}_k)$ and at every iteration where in case $\mathcal{S}_k \in \mathcal{COP}(\mathcal{K})$ they implement an optimality cut $\mathcal{C}'_k = \left\{(\mathsf{S}, \mathbf{y}) : \mathbf{b}^\mathsf{T}\mathbf{y} \geq \mathbf{b}^\mathsf{T}\mathbf{y}_k\right\}$. In addition, at each iteration one either obtains a lower bound on the problem in case $\mathsf{S}_k \in \mathcal{COP}(\mathcal{K})$ or an upper bound otherwise. The algorithm stops if the relative gap between the best lower and the best upper bound shrinks below predetermined threshold. Also, the copositivity check is performed by solving the standard quadratic optimization problem parameterized by S_k, given by $\min_{\mathbf{x}\in\mathbb{R}^n_+} \left\{\mathbf{x}^\mathsf{T}\mathsf{S}_k\mathbf{x} : \mathbf{e}^\mathsf{T}\mathbf{x} = 1\right\}$ via a mixed-integer programming approach outlined in Xia, Vera, & Zuluaga (2020).

In contrast, (Anstreicher & Gabl, 2022) solve $\nu(\mathcal{C}_k)$ to optimality at every iteration. As long as $\mathcal{S}_k \notin \mathcal{COP}(\mathcal{K})$ they generate a cut based on the certificate $\mathbf{x}_k$. In addition they employ various second-order cone cuts (which would take the role $\mathcal{C}'_k$ in our present notation). The algorithm stops as soon as the copositivity test is positive. In addition the authors provide their own mixed integer optimization based approach to set-copositivity testing, which is able to deal with cases where $\mathcal{K}$ is a polyhedral cone described by intersection of the non-negative orthant and arbitrarily many hyperplanes. Their approach is of particular interest to this text since they apply their algorithm to copositive reformulations of robust optimization problems (of the kind discussed in Section 4 below), and show that it can be used in conjunction with the approximation-based approaches discussed in the previous section, in order to test the quality of the latter approximations.

2.4. Concise guide: convex reformulations, Shor lifting and copositivity

In what follows we will provide the reader with a roadmap through the literature which may assist in understanding and further developing the theory around convex reformulations and copositive optimization. This is by no means an exhaustive list, nor does it imply any judgements on articles not mentioned here. More complete accounts of the respective literature may be found in Bomze, Schachinger, & Uchida (2012); Dür & Rendl (2021).

Historically, the idea of copositive matrices, hence matrices in $\mathcal{COP}(\mathbb{R}^n_+)$ goes back to Motzkin (1818), where the term and the concept were introduced originally. The dual term of complete positivity can be found in the early paper (Hall & Newman, 1963). However, the standard reference, as far as linear algebra is concerned, is the classic book (Berman & Shaked-Monderer, 2003), which mostly deals with $\mathcal{CPP}(\mathbb{R}^n_+)$. Further developments on the analysis of $\mathcal{COP}(\mathbb{R}^n_+)$ and $\mathcal{CPP}(\mathbb{R}^n_+)$ can be found in Dickinson (2010, 2013); Dür & Still (2008), which present interesting geometrical and topological insights on the two cones. For many of these results it is still an open question, whether they can be generalized to cases where the ground cones differ from the non-negative orthant. Some results for a general closed, convex ground cones can be found in Sturm & Zhang (2003). More extensive surveys on copositive and completely positive matrices are (Bomze, 2012; Bomze et al., 2012; Dür, 2010).

The classical Shor relaxation where $\mathcal{K} = \mathbb{R}^n$ was introduced in Shor (1987). Exactness proofs of this relaxation are regularly achieved via the results on the rank of extreme matrices of feasible sets of SDPs given in Pataki (1998), see for example Bomze & Gabl (2021); Burer & Anstreicher (2013). The first exactness result for a convex reformulation where $\mathcal{K} = \mathbb{R}^n_+$ is given in Bomze et al. (2000), where a convex reformulation for the standard quadratic optimization was derived. The core papers that introduce the methodology based on $\mathcal{G}(\mathcal{F})$ are (Anstreicher & Burer, 2010; Burer, 2009; 2012; Burer & Anstreicher, 2013; Eichfelder & Povh, 2013; Yang et al., 2016). An earlier contribution is however given in Sturm & Zhang (2003), who laid out many fundamental ideas of that machinery. Still, for the purposes of introduction we rather recommend (Burer, 2015), which will prepare the reader to deal with the more involved texts cited here. For a very recent survey see Dür & Rendl (2021). Finally, accounts of the strengths of convex relaxations of this style can be found in Anstreicher (2009, 2012); Anstreicher & Burer (2005); Bomze (2015), in which the reader may find theoretical guarantees as well as numerical studies.

Many more approximations have been proposed in literature, often in the form of hierarchies approximate the cones $\mathcal{COP}(\mathcal{K})$ or $\mathcal{CPP}(\mathcal{K})$ to arbitrary good accuracy, at the cost of introducing an exponentially increasing number of additional constraints. The interested reader may be referred to Bomze & de Klerk (2002); Bundfuss & Dür (2008, 2009); Dickinson & Povh (2013); de Klerk & Pasechnik (2002); Lasserre (2001); Parrilo (2000a,b); Peña, Vera, & Zuluaga (2007); Sponsel, Bundfuss, & Dür (2012); Yıldırım (2012). Due to significantly higher computational cost however, these approximations have not featured prominently in the literature on optimization under uncertainty yet, which is why we do not go into detail here.

3. A brief account on robust optimization and some variants

As mentioned above, we trust that most readers are familiar with the core concepts of robust optimization. Therefore, the following exposition is just exhaustive enough to make the subsequent discussion understandable.

In theory there are many types of optimization problems that can be solved efficiently to any desired accuracy, provided the structure of the problem, including the relevant data, is known. However in practice the latter is often not the case and one is confronted with an *uncertain optimization problem*:

$$\inf_{\mathbf{x}\in\mathbb{R}^n} \{f_0(\mathbf{x},\mathbf{u}) : f_i(\mathbf{x},\mathbf{u}) \geq 0,\ i\in[1\!:\!m]\} \text{ where } \mathbf{u}\in\mathcal{U}. \tag{16}$$

The parameters of the functions $f_i,\ i\in[0\!:\!m]$ are uncertain and governed by the *uncertainty parameter vector* $\mathbf{u}$ that lives in an *uncertainty set* $\mathcal{U}\subseteq\mathbb{R}^q$. This set encompasses all realizations of $\mathbf{u}$, for which the decision maker takes responsibility. Examples for designing appropriate uncertainty sets can be found in Ben-Tal et al. (2009); Bertsimas & Brown (2009); Bertsimas, Gupta, & Kallus (2018); Gorissen et al. (2015).

Under the robust optimization paradigm, one seeks to select a decision with the best worst-case performance among all decisions that are feasible for any realization of the uncertain data (see Ben-Tal et al., 2009; Gorissen et al., 2015 and references therein). The mathematical model encompassing this philosophy, the so-called *robust counterpart* of an uncertain optimization problem, is given by

$$\inf_{\mathbf{x}\in\mathbb{R}^n} \left\{ \sup_{\mathbf{u}\in\mathcal{U}} \{f_0(\mathbf{x},\mathbf{u})\} : f_i(\mathbf{x},\mathbf{u}) \geq 0,\ i\in[1\!:\!m] \text{ for all } \mathbf{u}\in\mathcal{U} \right\}. \tag{17}$$

In the rest of the text we will be mainly concerned with cases where $f_i,\ i\in[0\!:\!m]$ are quadratic functions in $\mathbf{u}$ and affine or concave quadratic in $\mathbf{x}$. For many specifications of f_i and $\mathcal{U}$, the robust counterpart can be reformulated into a tractable optimization problem, solvable via standard solutions strategies. The downside of this framework is that it is inherently conservative due to its pessimistic perspective on the eventual outcome of the uncertain process.

Many approaches have been proposed to remedy this shortcoming of conservativeness. One such approach is called *adjustable robust optimization (ARO)*. The domain of this approach are situations where parts of the decision can be delayed until uncertainty is revealed. These adjustable decisions are modeled as function-valued decision variables, hence one looks for the optimal policy which, conditional on the outcome of the uncertain process, will yield a good feasible solution of the optimization problem. Adjustable robust optimization was first introduced in

Ben-Tal et al. (2004), for a detailed survey see Yanikoglu et al. (2019). The *adjustable robust counterpart* can be written as

$$\inf_{\mathbf{x}\in\mathbb{R}^{n_1},\mathbf{y}(\mathbf{u})} \left\{ \sup_{\mathbf{u}\in\mathcal{U}} \{f_0(\mathbf{x},\mathbf{y}(\mathbf{u}),\mathbf{u})\} : f_i(\mathbf{x},\mathbf{y}(\mathbf{u}),\mathbf{u}) \geq 0 \quad \text{for all } \mathbf{u}\in\mathcal{U}\,,\ \ i\in[1\!:\!m] \right\}. \tag{18}$$

Compared to a robust optimization problem, the decision vector is split into two parts: the *first-stage decision vector* $\mathbf{x}\in\mathbb{R}^{n_1}$ and the *second-stage decision vector* $\mathbf{y}(\mathbf{u}) : \mathcal{U}\to\mathbb{R}^{n_2}$, where $\mathbf{y}(\mathbf{u})$ is allowed to adapt to the uncertainty and is thus a function of $\mathbf{u}$. Since the space of all functions is intractable, so is (18), and thus it is much harder to solve in practice than (17). However, there are many powerful approaches to (approximately) solve it (see Yanikoglu et al., 2019 and references therein), for example contracting the search space to the space of affine-linear functions, isomorphic to the tractable Euclidean vector space.

The final concept we will be interested in, as far as this text is concerned, is *distributionally robust optimization (DRO)*. It operates under the assumption that the uncertainty parameter vector is a random vector $\tilde{\mathbf{u}}$ governed by a probability distribution that is not known entirely, but assumed to reside in a set of distributions called the *ambiguity set*. The aim is now to optimize expected values of uncertain functions, under the assumption that the worst-case distribution will materialize for the chosen solution. The mathematical model that captures this paradigm is the *distributionally robust counterpart* given by:

$$\inf_{\mathbf{x}\in\mathbb{R}^{n_1}} \left\{ \sup_{P\in\mathcal{P}} \{\mathbb{E}_P[f_0(\mathbf{x},\tilde{\mathbf{u}})]\} : \mathbb{E}_P[f_i(\mathbf{x},\tilde{\mathbf{u}})] \geq 0 \quad \text{for all } P\in\mathcal{P},\ \ i\in[1\!:\!m] \right\}. \tag{19}$$

The inner supremum is taken over all expected values of the (random) objective function, w.r.t. all distributions $P\in\mathcal{P}$. The same way, we may rephrase the constraints as

$$\inf_{P\in\mathcal{P}} \mathbb{E}_P[f_i(\mathbf{x},\tilde{\mathbf{u}})] \geq 0 \quad \text{for all} \quad i\in[1\!:\!m]\,.$$

Note that if all members of the ambiguity set $\mathcal{P}$ are distributions with one-point support, we recover the robust counterpart. Hence, distributionally robust optimization is a generalization of robust optimization. However, the practicability of this approach stems from the fact that in many interesting cases the distributionally robust counterpart can be reduced to a robust counterpart, which can be tackled with all the instruments known from robust optimization.

3.1. Wasserstein ambiguity sets

An important way of constructing an ambiguity set involves the so-called Wasserstein balls, which are sets of probability measures with a Wasserstein distance to a certain reference distribution upper bounded by a constant $\varepsilon > 0$. The Wasserstein distance between two probability measures is a metric that can loosely be interpreted as the cost of transporting the mass of one distribution to the other. It enjoys rich theoretical background, most importantly, it gives a natural framework for data-driven optimization: if the empirical distribution of the uncertain data is used as reference distribution, ε can be chosen large enough to include the true distribution of the data-generating process. In fact (Esfahani & Kuhn, 2018; Zhao & Guan, 2018) give an explicit, closed-form description for $\varepsilon(\beta)$ which guarantees a $(1-\beta)$-confidence that the true distribution is contained in a Wasserstein ball of radius $\varepsilon(\beta)$ around the empirical distribution.

Formally we have the following

Definition 1. For any $r > 1$, let $\mathcal{M}^r(\mathcal{U})$ be the set of probability distributions P supported on $\mathcal{U}$ that satisfy $\mathbb{E}_P[d(\mathbf{u},\mathbf{u}_0)] = \int_{\mathcal{U}} d(\mathbf{u},\mathbf{u}_0)P(\mathrm{d}\mathbf{u}) < \infty$, where $\mathbf{u}_0\in\mathcal{U}$ is some reference point and $d(\mathbf{u},\mathbf{u}_0)$ is a continuous reference metric on $\mathcal{U}$. The r-Wasserstein distance between two distributions $P_1, P_2\in\mathcal{M}^r(\mathcal{U})$ is defined as

$$W^r(P_1,P_2) = \inf \left\{ \left[\int_{\mathcal{U}^2} d(\mathbf{u}_1,\mathbf{u}_2)^r\, Q(\mathrm{d}\mathbf{u}_1,\mathrm{d}\mathbf{u}_2)\right]^{\frac{1}{r}} : \begin{array}{c} Q \text{ is any joint distribution of } (\mathbf{u}_1,\mathbf{u}_2) \\ \text{with marginals } P_1 \text{ and } P_2 \end{array} \right\}.$$

Based on this notion, the ambiguity sets are often modelled as a ball induced by W^r, centered around an empirical distribution:

$$\mathcal{B}^r_\varepsilon(\hat{P}_I) := \left\{P\in\mathcal{M}^r(\mathcal{U}) : W^r(P,\hat{P}_I) \leq \varepsilon\right\}, \tag{20}$$

where $\hat{P}_I$ is the empirical probability measure based upon a sample $\left\{\hat{\mathbf{u}}_1,\dots,\hat{\mathbf{u}}_I\right\}$, i.e., $\hat{P}_I := \frac{1}{I}\sum_{i\in[1:I]}\delta_{\hat{\mathbf{u}}_i}$ where $\delta_{\hat{\mathbf{u}}}$ is the Dirac measure, which centers all its probability mass at $\hat{\mathbf{u}}\in\mathbb{R}^k$.

4. Robust constraints with quadratic index

In the case of quadratic optimization, such results are often obtained by invoking the so-called S-Lemma or some variants of it. However, copositive optimization theory opens an alternative path, which we will now review in great depth.

4.1. General strategy for finite reformulation

Assume that we are confronted with a robust constraint of the form:

$$f(\mathbf{x},\mathbf{u}) = \mathbf{u}^\mathsf{T}\mathsf{Q}(\mathbf{x})\mathbf{u} + 2\mathbf{q}(\mathbf{x})^\mathsf{T}\mathbf{u} + \omega(\mathbf{x}) \geq 0 \quad \text{for all } \mathbf{u}\in\mathcal{U}\,, \tag{21}$$

where $\mathsf{Q}(\mathbf{x}), \mathbf{q}(\mathbf{x}), \omega(\mathbf{x})$ are appropriate matrix-, vector- and scalar-valued functions of the decision vector $\mathbf{x}$. For ease of notation we suppress the dependence on $\mathbf{x}$. The reasoning to achieve a finite reformulation of (21) follows these steps:

- We again observe that
$$\mathbf{u}^\mathsf{T}\mathsf{Q}\mathbf{u}+2\mathbf{q}^\mathsf{T}\mathbf{u}+\omega \geq 0 \quad \text{for all } \mathbf{u}\in\mathcal{U} \quad \Longleftrightarrow \quad \inf_{\mathbf{u}\in\mathcal{U}}\left[\mathbf{u}^\mathsf{T}\mathsf{Q}\mathbf{u}+2\mathbf{q}^\mathsf{T}\mathbf{u}+\omega\right] \geq 0\,. \tag{22}$$
We will regularly refer to the optimization problem as the *inner* or *implied QCQP*.
- Next we need a convex reformulation of the inner infimum-problem, e.g., based on the ideas outlined in Section 2. For the sake of presentation we assume that
$$\inf_{\mathbf{u}\in\mathcal{U}}\left[\mathbf{u}^\mathsf{T}\mathsf{Q}\mathbf{u}+2\mathbf{q}^\mathsf{T}\mathbf{u}+\omega\right] = \inf_{\mathsf{Y}\in\mathcal{C}}\left\{\begin{bmatrix}\omega & \mathbf{q}^\mathsf{T}\\ \mathbf{q} & \mathsf{Q}\end{bmatrix}\bullet\mathsf{Y} : \mathsf{G}_i\bullet\mathsf{Y}\leq g_i\,,\ i\in[1\!:\!m]\right\}, \tag{23}$$
using an appropriate, convex matrix cone $\mathcal{C}$ and appropriate matrices $\mathsf{G}_i\in\mathcal{S}^{n+1}$, real numbers $b_i\in\mathbb{R}$, $i\in[1\!:\!m]$.
- If for the convex reformulation we can establish full strong duality, i.e., zero duality gap and dual attainability, we further have
$$\inf_{\mathsf{Y}\in\mathcal{C}}\left\{\begin{bmatrix}\omega & \mathbf{q}^\mathsf{T}\\ \mathbf{q} & \mathsf{Q}\end{bmatrix}\bullet\mathsf{Y} : \mathsf{G}_i\bullet\mathsf{Y}\leq g_i\,,\ i\in[1\!:\!m]\right\},$$
$$= \sup_{\boldsymbol{\lambda}\in\mathbb{R}^m_+}\left\{-\mathbf{g}^\mathsf{T}\boldsymbol{\lambda} : \begin{bmatrix}\omega & \mathbf{q}^\mathsf{T}\\ \mathbf{q} & \mathsf{Q}\end{bmatrix} + \sum_{i=1}^m \lambda_i\mathsf{G}_i \in\mathcal{C}^*\right\}.$$
where $\mathbf{g} := [g_1,\dots,g_m]^\mathsf{T}$ and $\boldsymbol{\lambda} := [\lambda_1,\dots,\lambda_m]^\mathsf{T}$.
- Since dual attainability guarantees the existence of the dual maximizers, we can enforce the semi-infinite constraint in (22) by demanding that
$$\mathbf{g}^\mathsf{T}\boldsymbol{\lambda}\leq 0 \quad \text{and} \quad \begin{bmatrix}\omega & \mathbf{q}^\mathsf{T}\\ \mathbf{q} & \mathsf{Q}\end{bmatrix} + \sum_{i=1}^m \lambda_i\mathsf{G}_i \in\mathcal{C}^* \quad \text{for some } \boldsymbol{\lambda}\in\mathbb{R}^m_+\,. \tag{24}$$

Readers experienced with robust optimization may have noticed that the most of this general strategy is part of the standard repertoire of techniques used in this field. Indeed, if $f(\mathbf{x},\mathbf{u})$ were linear in $\mathbf{u}$, the lifting and convexification step could be skipped and the remaining steps would be the familiar way of reformulating semi-infinite constraints via linear convex duality theory. Of course the difficult part is the convexification step, which is one of the main reasons why the techniques introduced in Section 2 are so vital for robust optimization. Once the hurdle of providing a convex reformulation of the inner QCQP is taken successfully, one may once again tread on familiar territory.

We also want to highlight that if $\mathcal{C}^*$ in (24) is replaced by an inner approximation $\mathcal{C}^*_{inner} \subseteq \mathcal{C}^*$, then (24) and hence (21) is still implied, so that we obtain a conservative approximation of the latter constraint. This is important since we will mostly work with cases of $\mathcal{C}$ that involve $\mathcal{CPP}(\mathcal{K})$ (so that $\mathcal{C}^*$ involves $\mathcal{COP}(\mathcal{K})$) in some capacity, and the latter cone is intractable, so that approximations are necessary, which is the major motivation behind the detailed discussion in Sections 2.3.1 and 2.3.2.

At this point we also like to comment on a common modelling choice, to construct the uncertainty set as a conic intersection $\mathcal{U} := \left\{\mathbf{u} : (1,\mathbf{u}^\mathsf{T})^\mathsf{T} \in \mathcal{K} \subseteq \mathbb{R}^{q+1}\right\}$. This is in fact a generic way to construct convex sets, as discussed in Rockafellar (2015, Section 8). The motivation behind this construction is a practical one: most studies that apply the general strategy do so in conjunction with Theorem 4 as workhorse which delivers the convexification step, and this theorem talks about feasible sets that are modelled as conic intersections. Hence, constructing $\mathcal{U}$ in this manner makes the application of the theorem more straightforward.

Finally, before reviewing literature where this general strategy has come to pass, we want to discuss the critical ingredients of the above strategy. We already discussed extensively how to close the relaxation gap in Section 2. The duality gap is usually easy to close since $\mathcal{U}$ is a bounded set so that the conic reformulation will also have a bounded feasible set, which is enough to guarantee a dual Slater point and thus a zero duality gap, albeit without dual attainability. The boundedness of $\mathcal{U}$ is in fact a generic property of an adequate uncertainty set. If it were unbounded, then the feasible set could be empty in case there is no $\mathbf{x}$ such that the constraint function is unbounded in $\mathbf{u}$ over $\mathcal{U}$. However, if there is a feasible $\mathbf{x}$ then the infinitely many constraints that are associated with $\mathbf{u}$ from the directions of recession of $\mathcal{U}$ are redundant. Hence, it does not make sense to consider unbounded uncertainty sets and in fact, to the best of our knowledge, uncertainty sets are generally assumed to be compact (see Ben-Tal et al., 2009; Gorissen et al., 2015; Yanikoglu et al., 2019). As a consequence, eliminating the duality gap is of little concern in most cases.

Howeve, dual attainability is the more elusive quality. For the conic reformulations we discussed, a Slater point in the primal problem, hence a feasible point in $\operatorname{int}\mathcal{CPP}(\mathcal{K})$, guarantees dual attainability. While a simple generalization of the results in Tuncel (2001) shows that $\mathcal{G}(\mathcal{F})$ has interior whenever $\mathcal{F}$ has interior, for reformulations based on Theorem 4, the most important type of reformulations, it is well known that the feasible set never has interior. However, the requirement of dual attainability can be loosened quite a bit. As shown in Bomze & Gabl (2021), one loses merely boundary points of the feasible set described by (21) when applying our general reformulation strategy without guaranteeing dual attainability. Hence, if the feasible set described by a collection of robust constraints is not connected merely by boundary points, e.g., if the sets described by the individual robust constraints have a common point in their respective relative interiors, one does not not need dual attainability.

4.2. Various applications of the general strategy for robust optimization

We will now discuss different instances of robust optimization that have appeared in the literature, where (21) takes a particular form, and where a reformulation into (24) is possible, given that the requirements of our general strategy are fulfilled. We will briefly describe the models, specify the values for $(\mathsf{Q},\mathbf{q},\omega)$ in the respective reformulation and discuss some features of their applications as stated in the original literature.

4.2.1. Linear ARO under uncertain recourse and affine decision rule

The generic linear ARO problem is given by

$$\min_{\mathbf{x}\in\mathcal{X},\mathbf{y}(\mathbf{u})} \left\{\mathbf{c}^\mathsf{T}\mathbf{x} : (\mathsf{A}_i\mathbf{u}+\mathbf{a}_i)\mathbf{x} + (\mathsf{B}_i\mathbf{u}+\mathbf{b}_i)\mathbf{y}(\mathbf{u}) + (\mathbf{u}^\mathsf{T}\mathsf{D}_i\mathbf{u}+\mathbf{d}_i^\mathsf{T}\mathbf{u}+d_i) \geq 0 \right. \\ \left. \text{for all } \mathbf{u}\in\mathcal{U},\ i\in[1\!:\!k]\right\}, \tag{25}$$

hence we have a linear optimization problem with uncertain coefficients, which we model as affine functions and quadratic functions in $\mathbf{u}$. More specifically, we model the coefficients of the first-stage decision $\mathbf{x}$ in the i-th constraint as affine functions involving the matrices $\mathsf{A}_i \in \mathbb{R}^{n_1\times q}$ and vectors $\mathbf{a}_i \in \mathbb{R}^{n_1}$, and the respective coefficients of the second-stage decisions $\mathbf{y}(\mathbf{u})$ as affine functions involving matrices $\mathsf{B}_i \in \mathbb{R}^{n_2\times q}$ and vectors $\mathbf{b}_i := (b_1,\dots,b_{n_2}) \in \mathbb{R}^{n_2}$. Finally, the offsets independent of $\mathbf{x},\mathbf{y}$ are modeled as quadratic functions involving matrices $\mathsf{D}_i \in \mathcal{S}^q$, vectors $\mathbf{d}_i \in \mathbb{R}^q$ and numbers d_i.

If the matrices B_i and D_i, $i\in[1\!:\!k]$ were zero, then the above model would coincide with the one studied in Ben-Tal et al. (2004), the seminal paper on ARO. In that case, if one applies an affine decision rule by specifying $\mathbf{y}(\mathbf{u}) = \mathsf{Y}\mathbf{u}+\mathbf{y}_0$, where the coefficients $\mathsf{Y}\in\mathbb{R}^{n_2\times m}$ and $\mathbf{y}_0\in\mathbb{R}^{n_2}$ take the role of the decision vector, then linear, convex duality is readily applicable, modulo some regularity conditions on $\mathcal{U}$, in order to obtain a finite reformulation of the robust constraints. The complication arises if one considers uncertain recourse, i.e., when B_i are not zero. Then, bilinear terms in $\mathbf{u}$ arise and duality of the implied, inner infimum is no longer guaranteed. However, the general strategy allows us to proceed anyway. Focusing on a single constraint of the above model, we are concerned with

$$(\mathsf{A}\mathbf{u}+\mathbf{a})^\mathsf{T}\mathbf{x} + (\mathsf{B}\mathbf{u}+\mathbf{b})^\mathsf{T}(\mathsf{Y}\mathbf{u}+\mathbf{y}_0) + \mathbf{u}^\mathsf{T}\mathsf{D}\mathbf{u} + \mathbf{d}^\mathsf{T}\mathbf{u} + d \geq 0 \quad \text{for all } \mathbf{u}\in\mathcal{U}, \tag{26}$$

where an affine decision rule has already been put into place. We omit an index indicating which of the k constraints we are concerned with, since they are all structurally identical. Also, letting $\mathsf{D}\neq\mathsf{O}$ does not hinder the application of our techniques, which gives some additional modelling power aside from uncertain recourse. Applying the general strategy in a straightforward manner allows us to achieve the following result.

Theorem 14. *Assume that (26) has an exact conic reformulation of the form (23) enjoying full strong duality. Then problem (26) is equivalent to*

$$\mathbf{g}^\mathsf{T}\boldsymbol{\lambda} \leq 0,$$
$$\begin{bmatrix} \mathbf{a}^\mathsf{T}\mathbf{x}+\mathbf{b}^\mathsf{T}\mathbf{y}_0+d & \frac{1}{2}\left(\mathsf{A}^\mathsf{T}\mathbf{x}+\mathsf{Y}^\mathsf{T}\mathbf{b}+\mathsf{B}^\mathsf{T}\mathbf{y}_0+\mathbf{d}\right)^\mathsf{T} \\ \frac{1}{2}\left(\mathsf{A}^\mathsf{T}\mathbf{x}+\mathsf{Y}^\mathsf{T}\mathbf{b}+\mathsf{B}^\mathsf{T}\mathbf{y}_0+\mathbf{d}\right) & \mathsf{D}+\frac{1}{2}\left(\mathsf{B}^\mathsf{T}\mathsf{Y}+\mathsf{Y}^\mathsf{T}\mathsf{B}\right)\end{bmatrix} + \sum_{i=1}^m \lambda_i\mathsf{G}_i \in \mathcal{C}^*,$$
$$\boldsymbol{\lambda}\in\mathbb{R}^m_+.$$

Proof. The theorem follows immediately from our general strategy. Note that in order to symmetrize the quadratic term we use $\mathbf{u}^\mathsf{T}\mathsf{B}^\mathsf{T}\mathsf{Y}\mathbf{u} = \frac{1}{2}\mathbf{u}^\mathsf{T}\left(\mathsf{B}^\mathsf{T}\mathsf{Y}+\mathsf{Y}^\mathsf{T}\mathsf{B}\right)\mathbf{u}$. □

Already in Ben-Tal et al. (2004), the authors provided finite convex reformulations of ARO under uncertain recourse and affine decision rule, in case the uncertainty set is an ellipsoid, where the

S-Lemma provided the necessary convexification. Apart from this special case, the authors also provided a conservative approximation based on an approximate S-Lemma. In contrast, our general strategy in conjunction with the results discussed in Section 2 allows for a wider range of uncertainty sets to be utilized. The first paper to apply this machinery was (Xu & Hanasusanto, 2018), where the convexification was achieved by means of Theorem 4. Of course the involvement of $\mathcal{CPP}(\mathcal{K})$ may again necessitate the use of approximations, but the authors provide such approximations for interesting choices of $\mathcal{U}$ and prove that these perform at least as good as the approximations based on the approximate S-Lemma (see Xu & Hanasusanto, 2018, Proposition 3).

4.2.2. Linear ARO under fixed recourse and quadratic decision rules

We again consider (25) with the slight modification that $\mathsf{B}_i,\ i \in [1\!:\!k]$ are set to zero, hence, we have fixed recourse. In this case the introduction of an affine decision rule does not lead to bilinear terms in $\mathbf{u}$, and standard reformulation procedures can be applied. However, we can do better than that. Utilizing our general strategy allows us to expand the search space for the second-stage decision from the space of affine functions to the space of quadratic functions. Thus, we specify

$$\mathbf{y}(\mathbf{u}) = \begin{bmatrix} \mathbf{u}^\mathsf{T}\mathsf{Y}_1\mathbf{u} + \mathbf{y}_1^\mathsf{T}\mathbf{u} + y_1 \\ \dots \\ \mathbf{u}^\mathsf{T}\mathsf{Y}_{n_2}\mathbf{u} + \mathbf{y}_{n_2}^\mathsf{T}\mathbf{u} + y_{n_2} \end{bmatrix},$$

so that the robust constraint can be written as

$$\mathbf{u}^\mathsf{T}\left[\sum_{j=1}^{n_2} b_j\mathsf{Y}_j + \mathsf{D}\right]\mathbf{u} + \left[\sum_{j=1}^{n_2} b_j\mathbf{y}_j + \mathsf{A}^\mathsf{T}\mathbf{x} + \mathbf{d}\right]^\mathsf{T}\mathbf{u} + \mathbf{a}^\mathsf{T}\mathbf{x} + \sum_{j=1}^{n_2} b_j y_j + d \geq 0 \quad \text{for all } \mathbf{u} \in \mathcal{U}. \tag{27}$$

Note, that under fixed recourse the coefficients of $\mathbf{y}(\mathbf{u})$ reduce to the vector $\mathbf{b} \in \mathbb{R}^{n_2}$, and we again suppressed the row index.

Theorem 15. *Assume that (27) has an exact conic reformulation of the form (23) enjoying full strong duality. Then (27) is equivalent to*

$$\mathbf{g}^\mathsf{T}\boldsymbol{\lambda} \leq 0\,,$$
$$\begin{bmatrix} \mathbf{a}^\mathsf{T}\mathbf{x} + \sum_{j=1}^{n_2} b_j y_j + d & \frac{1}{2}\left(\sum_{j=1}^{n_2} b_j\mathbf{y}_j + \mathsf{A}^\mathsf{T}\mathbf{x} + \mathbf{d}\right)^\mathsf{T} \\ \frac{1}{2}\left(\sum_{j=1}^{n_2} b_j\mathbf{y}_j + \mathsf{A}^\mathsf{T}\mathbf{x} + \mathbf{d}\right) & \mathsf{D} + \sum_{j=1}^{n_2} b_j\mathsf{Y}_j \end{bmatrix} + \sum_{i=1}^{m}\lambda_i\mathsf{G}_i \in \mathcal{C}^*\,,$$
$$\boldsymbol{\lambda} \in \mathbb{R}^m_+\,.$$

Proof. The theorem follows immediately from our general strategy. □

Quadratic decision rules have been applied in various articles, usually under some restrictions regarding the uncertainty set or the structure of the quadratic forms in $\mathbf{y}(\mathbf{u})$. For example, in case the uncertainty set is ellipsoidal, the S-Lemma allows for a finite convex reformulation of (27), and an exhaustive list of similar approaches can be found in Yanikoglu et al. (2019, Table 3). The approaches often restrict the form of the quadratic decision rule, for example to separable quadratic functions, where no bilinear terms are present. However, as first shown in Xu & Hanasusanto (2018) and again presented here, the quadratic decision rule is much more generally applicable if one uses the general strategy in conjunction with Theorem 4.

4.2.3. Convex quadratic robust optimization

The model of interest here is

$$\min_{\mathbf{x}\in\mathcal{X}}\left\{\mathbf{c}^\mathsf{T}\mathbf{x} : -\|\mathsf{A}_i(\mathbf{x})\mathbf{u}\|^2 + (\mathbf{a}_i(\mathbf{x}))^\mathsf{T}\mathbf{u} + a_i(\mathbf{x}) + \mathbf{u}^\mathsf{T}\mathsf{D}_i\mathbf{u} + \mathbf{d}_i^\mathsf{T}\mathbf{u} + d_i \geq 0 \quad \text{for all } \mathbf{u} \in \mathcal{U}\,,\ i \in [1\!:\!k]\right\}, \tag{28}$$

where $\mathsf{A}_i(\mathbf{x}) : \mathbb{R}^{n_1} \to \mathbb{R}^{k\times q}$ and $\mathbf{a}_i(\mathbf{x}) : \mathbb{R}^{n_1} \to \mathbb{R}^q$ are affine matrix and vector pencils, respectively, and $a_i(\mathbf{x}) : \mathbb{R}^{n_1} \to \mathbb{R}$ is a real-valued, affine function of $\mathbf{x}$. This case was recently addressed by Mittal, Gökalp, & Hanasusanto (2019), in a way similar to the approach presented here. We will slightly generalize their result, again focusing on an arbitrary constraint in (28) given by

$$-\|\mathsf{A}(\mathbf{x})\mathbf{u}\|^2 + (\mathbf{a}(\mathbf{x}))^\mathsf{T}\mathbf{u} + a(\mathbf{x}) + \mathbf{u}^\mathsf{T}\mathsf{D}\mathbf{u} + \mathbf{d}^\mathsf{T}\mathbf{u} + d \geq 0 \quad \text{for all } \mathbf{u} \in \mathcal{U}\,. \tag{29}$$

It is clear from the general strategy that we can reformulate (29) as

$$\mathbf{b}^\mathsf{T}\boldsymbol{\lambda} \leq 0,$$
$$\begin{bmatrix} d + a(\mathbf{x}) & \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d})^\mathsf{T} \\ \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d}) & \mathsf{D} - \mathsf{A}(\mathbf{x})^\mathsf{T}\mathsf{A}(\mathbf{x}) \end{bmatrix} + \sum_{i=1}^{m}\lambda_i\mathsf{G}_i \in \mathcal{C}^*\,. \tag{30}$$

The entries of the south-east diagonal block of the constraints matrix in (30) are now quadratic functions. In case $\mathcal{C}^* = \mathcal{COP}(\mathcal{K})$ for some closed, convex cone $\mathcal{K} \subseteq \mathbb{R}^{q+1}$ (which is the case for all the conic reformulations of QCQPs discussed in this text), we can linearize the constraints by employing the following lemma, which is a straightforward generalization of Mittal et al. (2019, Lemma 4).

Lemma 16. *Assume $\mathcal{C}^* = \mathcal{COP}(\mathcal{K})$ for some cone $\mathcal{K} \subseteq \mathbb{R}^{q+1}$. Then a vector $\mathbf{x} \in \mathbb{R}^n$ fulfills the conic constraint in (30) if and only if there exists a matrix $\mathsf{H} \in \mathcal{S}^q$ such that*

$$\begin{bmatrix} d + a(\mathbf{x}) & \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d})^\mathsf{T} \\ \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d}) & \mathsf{D} - \mathsf{A}(\mathbf{x})^\mathsf{T}\mathsf{A}(\mathbf{x}) \end{bmatrix} + \sum_{i=1}^{m}\lambda_i\mathsf{G}_i \in \mathcal{COP}(\mathcal{K}) \quad \text{and} \quad \begin{bmatrix} \mathsf{H} & \mathsf{A}(\mathbf{x})^\mathsf{T} \\ \mathsf{A}(\mathbf{x}) & \mathsf{I} \end{bmatrix} \in \mathcal{S}^{q+k}_+\,.$$

Using this lemma we can derive the following theorem

Theorem 17. *Assume that (29) has an exact conic reformulation of the form (23), with $\mathcal{C} = \mathcal{CPP}(\mathcal{K})$ for some appropriate cone $\mathcal{K}$, enjoying full strong duality. Then (29) is equivalent to*

$$\mathbf{g}^\mathsf{T}\boldsymbol{\lambda} \leq 0\,, \quad \boldsymbol{\lambda} \in \mathbb{R}^m_+\,,$$
$$\begin{bmatrix} d + a(\mathbf{x}) & \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d})^\mathsf{T} \\ \frac{1}{2}(\mathbf{a}(\mathbf{x}) + \mathbf{d}) & \mathsf{D} - \mathsf{A}(\mathbf{x})^\mathsf{T}\mathsf{A}(\mathbf{x}) \end{bmatrix} + \sum_{i=1}^{m}\lambda_i\mathsf{G}_i \in \mathcal{COP}(\mathcal{K})$$
$$\text{and} \quad \begin{bmatrix} \mathsf{H} & \mathsf{A}(\mathbf{x})^\mathsf{T} \\ \mathsf{A}(\mathbf{x}) & \mathsf{I} \end{bmatrix} \in \mathcal{S}^{q+k}_+\,.$$

Proof. The theorem follows immediately from our general strategy. □

The setting can be transferred to the ARO case in a straightforward manner, using the tools discussed in this and the previous section. The second-stage variables may enter linearly with fixed or uncertain recourse, in which case the all the strategies that we discussed apply immediately. In case the second-stage enters in a convex quadratic manner, analogous to the vector $\mathbf{x}$ in this section, one can apply an affine policy and use Lemma 16 in order to obtain a convex conic formulation. At this point, for the sake of brevity we leave the details to the reader and skip the respective presentation.

4.2.4. Distributionally robust, and two-stage distributionally robust, optimization

Two recent papers exploit reformulations of distributionally robust optimization problems into semi-infinite optimization problems in order to arrive at representations of these problems where constraints are amenable to the general strategy. We will briefly discuss their approach in the following paragraphs.

The first paper in this regard (Jiang, Ryu, & Xu, 2019), deals with appointment scheduling under data ambiguity, where the ambiguity set is constructed using Wasserstein balls, the construction of which was discussed in Section 3.1. The authors investigate the model

$$\inf_{\mathbf{x}\in\mathcal{X}} \left\{ \sup_{P\in\mathcal{B}^r_\varepsilon(\hat{P}_I)} \{\mathbb{E}_P[f(\mathbf{x},\tilde{\mathbf{u}})]\} \right\}, \tag{31}$$

where $\mathcal{X}\subseteq\mathbb{R}^n$ is a feasible set, not affected by uncertainty, and f is an objective function. In addition, the metric in the definition of the Wasserstein distance is chosen to be the p-norm with $p=r$ so that

$$W^r(P_1,P_2) = \inf \left\{ \left[\int_{\mathcal{U}^2} \|\mathbf{u}_1,\mathbf{u}_2\|_r^r\, Q(\mathrm{d}\mathbf{u}_1,\mathrm{d}\mathbf{u}_2)\right]^{\frac{1}{r}} : \begin{array}{c} Q \text{ is any joint distribution of } (\mathbf{u}_1,\mathbf{u}_2) \\ \text{with marginals } P_1 \text{ and } P_2 \end{array} \right\}.$$

For this model the authors derive the following semi-infinite representation

$$\begin{aligned} &\inf_{\mathbf{x}\in\mathcal{X},\rho,\theta} \varepsilon^r\rho + \frac{1}{I}\sum_{j=1}^{I}\theta_j \\ &\text{s.t.}: f(\mathbf{x},\mathbf{u}) - \rho\|\mathbf{u}-\hat{\mathbf{u}}_j\|_r^r \le \theta_j \ \text{ for all } \mathbf{u}\in\mathcal{U},\ \text{ all } j\in[1:I] \\ &\qquad \rho\ge 0,\ \ \theta\in\mathbb{R}^I. \end{aligned} \tag{32}$$

In case $r\in\{1,2\}$ the second term in the semi-infinite constraint is linear or quadratic in $\mathbf{u}$ respectively. If in addition

$$f(\mathbf{x},\mathbf{u}) := \sup_{\mathbf{w}\in\mathcal{W}} q(\mathbf{x},\mathbf{u},\mathbf{w}), \quad \text{where } q(\mathbf{u},\mathbf{w}) := \begin{bmatrix}\mathbf{u}\\ \mathbf{w}\end{bmatrix}^\top \mathsf{Q}(\mathbf{x}) \begin{bmatrix}\mathbf{u}\\ \mathbf{w}\end{bmatrix}, \tag{33}$$

and hence, a pointwise maximum of quadratic functions involving some matrix valued function $\mathsf{Q}:\mathbb{R}^n\to\mathcal{S}^{k+n}$ and an index set $\mathcal{W}\subseteq\mathbb{R}^k$, we can reformulate the semi-infinite constraint in (33) as

$$q(\mathbf{x},\mathbf{u},\mathbf{w}) - \|\mathbf{u}-\hat{\mathbf{u}}_j\|_r^r \le \theta_j \ \text{ for all } (\mathbf{u}^\top,\mathbf{w}^\top)^\top\in\mathcal{U}\times\mathcal{W}, \quad \text{all } j\in[1:I].$$

Since we have produced semi-infinite constraints with quadratic index, we can apply the general strategy in order to obtain a convex reformulation. Note, that as long as the dependence of Q on $\mathbf{x}$ is linear or convex quadratic, we can use the strategy directly or consecutively invoke Lemma 16 in order to obtain a problem with only linear terms in $\mathbf{x}$. The authors apply this methodology to robust appointment scheduling, in which case f is a certain pointwise maximum of linear functions linear in $\mathbf{u}$, so that q is bilinear in $(\mathbf{u}^\top,\mathbf{w}^\top)^\top$ and $\mathcal{W}$ is and appropriate polyhedron.

The second paper (Fan & Hanasusanto, 2021) deals with risk-averse two-stage distributionally robust optimization under the conditional value at risk (CVaR) as risk measure. The respective model is given by

$$\inf_{\mathbf{x}\in\mathcal{X}} \mathbf{c}^\top\mathbf{x} + \sup_{P\in\mathcal{P}} \mathrm{CVaR}_\delta^P(Z(\mathbf{x},\mathbf{u})), \tag{34}$$

where $\mathrm{CVaR}_\delta^P(.)$ is the conditional value at risk at level δ of a risky position whose distribution is P, $\mathbf{u}$ is the uncertain parameter, $\mathcal{P}$ is a set of plausible distributions supported on a conic intersection $\mathcal{U} := \left\{\mathbf{u} : (1,\mathbf{u}^\top)^\top\in\mathcal{K}\subseteq\mathbb{R}^{q+1}\right\}$, $\mathcal{X}\subseteq\mathbb{R}^{n_2}$ is a feasible set not affected by uncertainty, and $Z(\mathbf{x},\mathbf{u})$ is the recourse problem given by

$$Z(\mathbf{x},\mathbf{u}) := \inf_{\mathbf{y}\in\mathbb{R}^{n_2}} \left\{\mathbf{u}^\top\mathsf{D}^\top\mathbf{y} : \mathsf{T}_l(\mathbf{x})^\top\mathbf{u} \le \mathbf{u}^\top\mathsf{W}_l^\top\mathbf{y} \ \text{ for all } l\in[1:L]\right\},$$

with appropriate matrices $\mathsf{D},\mathsf{W}_l,\ l\in[1:L]$ and matrix valued functions $\mathsf{T}_l(\mathbf{x}),\ l\in[1:L]$. Hence, we look for a first stage decision $\mathbf{x}$ so that the worst case CVaR of our second stage response to an uncertain parameter is optimized. The authors show that (34) can be reformulated as

$$\begin{aligned} &\inf_{\mathbf{x}\in\mathcal{X},\theta,\tau(\mathbf{u}),\mathbf{y}(\mathbf{u})} \mathbf{c}^\top\mathbf{x} + \theta + \frac{1}{\delta}\sup_{P\in\mathcal{P}}\mathbb{E}_P(\tau(\mathbf{u})) \\ &\quad \left.\begin{array}{r} \text{s.t.}: \tau(\mathbf{u})\ge 0\,, \\ \tau(\mathbf{u}) \ge \mathbf{u}^\top\mathsf{D}^\top\mathbf{y}(\mathbf{u}) - \theta\,, \\ \mathsf{T}_l(\mathbf{x})^\top\mathbf{u} \le \mathbf{u}^\top\mathsf{W}_l^\top\mathbf{y}(\mathbf{u}) \ \text{ for all } l\in[1:L], \end{array}\right\} \text{ for all } \ \mathbf{u}\in\mathcal{U}. \end{aligned} \tag{35}$$

This reformulation should make it tangible for the reader that the second stage decision $(\tau(\mathbf{u}),\mathbf{y}(\mathbf{u}))$ can be subjected to linear and quadratic policies, so that the semi-infinite constraints can be tackled via the general strategy. However, the supremum term in the objective still needs to be taken care of first, which would require detailing the intricate construction of the ambiguity set used in Fan & Hanasusanto (2021) and some extensive massaging of that term depicted therein. But this lies beyond the scope of this text, and we refer the reader to the original source for these details. Nonetheless, the general strategy is a core ingredient of the authors' derivations, the results of which are eventually applied to network inventory allocation and the multi-item newsvendor problem. We do, however, like to mention the fact, that said construction of the ambiguity sets necessitates the introduction of additional semi-infinite constraints, which are duplications of the ones present in (35) corresponding to certain subsets of the support $\mathcal{U}$. The authors tackle the computational challenge of the potentially large number of matrix blocks that arise from the general strategy via a Bender's decomposition approach, which allows for a parallelization of the solution of the copositive sub-problems

4.3. Viable uncertainty sets

So far we have demonstrated how convex reformulations expand the modeling capabilities with respect to the functional form of the robust constraints. However, the theorems that enable these reformulations put some requirements on the feasible sets of the inner QCQP and therefore on the uncertainty sets, while at the same time they are allowing new modeling choices there as well. We will now provide an overview over the uncertainty sets that can be managed with the machinery outlined above, and discuss their benefits and limitations.

4.3.1. Primitive uncertainty sets

A number of uncertainty sets are regularly cited as being standard or classic, among them ellipsoidal and polyhedral uncertainty sets. We will briefly discuss how they are handled in context of our general strategy.

Ellipsoidal uncertainty sets are easily tackled by the general strategy via Theorem 3 (regarding $\mathcal{F}_1$ with no linear constraints), which in essence boils down to a roundabout way of using the S-Lemma since the respective characterization of $\mathcal{G}(\mathcal{F})$ is based on that result. However, the S-Lemma can be employed directly to the infimum problem in (22) in order to obtain a dual supremum problem and thus a finite reformulation. While our framework does not offer anything new in this respect, it is neither restrictive as well.

Polyhedral uncertainty sets can be tackled using Theorem 4. However, there is some ambiguity to which we like to draw some attention. One way to generally represent polyhedra is $\mathcal{P}_1 := \left\{\mathbf{x}\in\mathbb{R}^n_+ : \mathsf{A}\mathbf{x}=\mathbf{b}\right\}$ in which case Theorem 4 readily provides a description of $\mathcal{G}(\mathcal{P}_1)$ involving $\mathcal{CPP}(\mathbb{R}^n_+)$. However, another generic description is given by $\mathcal{P}_2 := \{\mathbf{x}\in\mathbb{R}^n : \mathsf{A}\mathbf{x}\le\mathbf{b}\}$ in which case Theorem 4 can be applied after introducing slack variables $\mathbf{s}\in\mathbb{R}^m$, where m the number of inequality constraints in the description of $\mathcal{P}_2$. The resulting characterization of $\mathcal{G}(\mathcal{P}_2)$ would involve $\mathcal{CPP}(\mathbb{R}^n\times\mathbb{R}^m_+)$ which by Proposition 11 point 6. can be expressed using only $\mathcal{S}_+^{n+m}$ and $\mathcal{CPP}(\mathbb{R}^m)$. Exploiting this ambiguity, one might choose the description that yields the smaller completely positive constraint, which may reduce complexity.

In (Xu & Hanasusanto, 2018) the authors study combinations of these of ellipsoidal and polyhedral uncertainty sets, where the facets of the polyhedron do not meet inside the ellipsoid. In that case Theorem 13 provides an exact representation of the conic constraints present in $\mathcal{G}(\mathcal{U})$.

4.3.2. Mixed-integer uncertainty sets

In (Mittal et al., 2019) the authors use Theorem 4 in order to introduce uncertainty sets with mixed-integer components, namely

$$\mathcal{U} := \left\{ \mathbf{u} \in \mathbb{R}_+^k : \mathsf{A}\mathbf{u} = \mathbf{b},\ u_l \in \mathbb{Z} \quad \text{for all } l \in \mathcal{L} \right\} \tag{36}$$

where $\mathcal{L} \subseteq [1\!:\!k]$. One can assume without loss of generality that $\mathcal{L} := [1\!:\!L]$ for some $L \le k$. Under the additional assumption that $\mathcal{U}$ is bounded we can always express any integer component of a member of $\mathcal{U}$ by binary expansion as $u_l = \sum_{i=1}^{Q} 2^{i-1} v_{il} = \mathbf{q}^\mathsf{T}\mathbf{v}_l$ for some integer Q. Hence the set

$$\mathcal{U}' := \left\{ (\mathbf{u}, \mathsf{V}, \mathsf{S}) \in \mathbb{R}_+^k \times \{0,1\}^{Q\times L} \times \mathbb{R}_+^{Q\times L} : \mathsf{A}\mathbf{u} = \mathbf{b},\ u_l = \mathbf{q}^\mathsf{T}\mathbf{v}_l,\ \mathbf{v}_l + \mathbf{s}_l = \mathbf{e},\ l \in [1\!:\!L] \right\}$$

has $\mathcal{U}$ as its projection on the $\mathbf{u}$-coordinates. Note that next to the variables in V we also had to introduce additional constraints and slack variables. This is done in order to meet the requirements of Theorem 4. Hence, any robust constraint with quadratic index in an uncertainty set $\mathcal{U}$ can be cast as a robust constraints over $\mathcal{U}'$, which can then be reformulated using the general strategy in conjunction with Theorem 4. The resulting copositive constraint will involve $\mathcal{COP}(\mathbb{R}_+^{k+2QL})$, however, the authors of Mittal et al. (2019) prove that even the simplest inner approximation based on $\mathcal{NND}^{k+2QL}$ outperforms the classical approach based on the approximate S-Lemma introduced in (Ben-Tal et al., 2004).

Note that the convex formulation based on Theorem 4 scales quadratically in the dimension of the original quadratic problem. Hence, the introduction of the additional variables may come at a potentially high cost of optimizing over a large set-copositive constraint. Providing reformulation strategies that do not require the excessive lifting when changing from $\mathcal{U}$ to $\mathcal{U}'$ is therefore a desirable achievement to be pursued in future research.

4.3.3. Adapting the uncertainty set to piecewise affine/quadratic decision rules

In (Xu & Hanasusanto, 2018) the authors skillfully exploited the modeling capabilities offered by Theorem 4 in order to enable piecewise linear and quadratic decision rules. Given that the uncertainty set is defined as a compact, convex, conic intersection given by:

$$\mathcal{U} := \left\{ \mathbf{u} : (1, \mathbf{u}^\mathsf{T})^\mathsf{T} \in \mathcal{K} \subseteq \mathbb{R}^{q+1} \right\},$$

one can lift the uncertainty set to obtain

$$\mathcal{U}' := \left\{ (\mathbf{u}, \mathbf{w}) \subseteq \mathcal{U} \times \mathbb{R}^L : w_l = \max\left\{0, \mathbf{g}_l^\mathsf{T}\mathbf{u} - h_l\right\},\ l \in [1\!:\!L] \right\}.$$

Here $\mathbf{g}_l$ is interpreted as the folding direction of the lth piece of the piecewise policy and h_l is its breakpoint. Clearly, a general adjustable robust constraint in (25) is equivalent to

$$(\mathsf{A}\mathbf{u} + \mathbf{a})^\mathsf{T}\mathbf{x} + (\mathsf{B}\mathbf{u} + \mathbf{b})^\mathsf{T}\mathbf{y}'(\mathbf{u}, \mathbf{w}) + \mathbf{u}^\mathsf{T}\mathsf{D}\mathbf{u} + \mathbf{d}^\mathsf{T}\mathbf{u} + d \ge 0 \quad \text{for all } (\mathbf{u}, \mathbf{w}) \in \mathcal{U}', \tag{37}$$

since $\mathbf{y}(\mathbf{u}) := \mathbf{y}'\left(\mathbf{u}, \max\left\{0, \mathbf{g}_1^\mathsf{T}\mathbf{u} - h_1\right\}, \ldots, \max\left\{0, \mathbf{g}_L^\mathsf{T}\mathbf{u} - h_L\right\}\right)$ is a function that maps $\mathcal{U}$ into $\mathbb{R}^{n_2}$, and vice versa any function of $\mathbf{u}$ can be generated from functions of $(\mathbf{u}, \mathbf{w})$, with $\mathbf{w}$ defined as in $\mathcal{U}'$. However, if we restrict $\mathbf{y}'$ to be affine or quadratic in its arguments, then $\mathbf{y}(\mathbf{u})$ is a piecewise affine/quadratic function in $\mathbf{u}$. Hence, we can easily introduce piecewise policies by merely updating the uncertainty set accordingly, albeit at the price of a having to work with a nonconvex uncertainty set. A simple argument shows that we have

$$\mathcal{U}' = \left\{ (\mathbf{u}, \mathbf{w}) \subseteq \mathcal{U} \times \mathbb{R}^L : \begin{array}{ll} \mathbf{o} \le \mathbf{w} \le \bar{\mathbf{w}}, & \\ w_l \ge \mathbf{g}_l^\mathsf{T}\mathbf{u} - h_l, & l \in [1\!:\!L], \\ w_l(w_l - \mathbf{g}_l + h_l) = 0, & l \in [1\!:\!L] \end{array} \right\}$$
$$= \left\{ (\mathbf{u}, \mathbf{w}) : (1, \mathbf{u}, \mathbf{w}) \in \mathcal{K}',\ w_l(w_l - \mathbf{g}_l + h_l) = 0,\ l \in [1\!:\!L] \right\},$$

where

$$\mathcal{K}' := \left\{ (u_0, \mathbf{u}, \mathbf{w}) \in \mathcal{K} \times \mathbb{R}_+^L : \begin{array}{ll} \mathbf{w} \le u_0\bar{\mathbf{w}}, & \\ w_l \ge \mathbf{g}_l^\mathsf{T}\mathbf{u} - h_l, & l \in [1\!:\!L] \end{array} \right\}.$$

Note that $\mathcal{U}'$ is a bounded set and for all $(u_0, \mathbf{u}, \mathbf{w}) \in \mathcal{K}'$ we have $w_l(w_l - \mathbf{g}_l^\mathsf{T}\mathbf{u} + h_l) \ge 0,\ l \in [1\!:\!L]$, so that the key condition in Theorem 4 is satisfied for any quadratic optimization problem over $\mathcal{U}'$. Hence, after replacing $\mathbf{y}'$ by an affine or an quadratic policy (in case of fixed recourse) in $(\mathbf{u}, \mathbf{w})$ we can use the general strategy in conjunction with Theorem 4 to obtain a finite reformulation of (38) under a piecewise affine/quadratic policy. The final result involves the cone $\mathcal{COP}(\mathcal{K}')$ for which the authors of Xu & Hanasusanto (2018) find tractable outer, hence conservative, approximations based on $\mathcal{SI}(\mathcal{K}')$ from Theorem 13.

4.4. Application: disjoint convex-convex quadratic optimization

Following the core idea of Zhen, Marandi, de Moor, den Hertog, & Vandenberghe (2022), the authors of Bomze & Gabl (2021) proposed a convex lower bound of special type of QCQP based on a reformulation as an ARO problem that can be approximated, using the general strategy and the results presented in the preceding sections. The following theorem presents the QCQP and its adjustable robust reformulation:

Theorem 18. *Let $\mathsf{Q}_x \in \mathcal{S}^{n_1}$, $\mathsf{Q}_{xy} \in \mathbb{R}^{n_1 \times n_2}$, $\mathsf{F} \in \mathbb{R}^{k \times n_2}$ and $\mathsf{G} \in \mathbb{R}^{r \times n_2}$. Further, assume $\mathcal{X} \subseteq \mathbb{R}^{n_1}$ is a compact set and $\mathcal{Y} := \{\mathbf{y} \in \mathbb{R}_+^{n_2} : \mathsf{F}\mathbf{y} = \mathbf{d}\} \subseteq \mathbb{R}^{n_2}$ has a Slater point and let $\mathcal{Z}(\mathbf{x}) := \{(\mathbf{z}, \mathbf{w}) : \mathsf{F}^\mathsf{T}\mathbf{z} + \mathsf{G}^\mathsf{T}\mathbf{w} \le \mathsf{Q}_{xy}^\mathsf{T}\mathbf{x}\}$. Then*

$$\inf_{\mathbf{x}\in\mathcal{X}, \mathbf{y}\in\mathcal{Y}} \mathbf{x}^\mathsf{T}\mathsf{Q}_x\mathbf{x} + \mathbf{x}^\mathsf{T}\mathsf{Q}_{xy}\mathbf{y} + \|\mathsf{G}\mathbf{y}\|^2 \tag{38}$$

$$= \sup_\tau \left\{ \tau : \forall\, \mathbf{x} \in \mathcal{X}\ \exists\ (\mathbf{z}(\mathbf{x}), \mathbf{w}(\mathbf{x})) \in \mathcal{Z}(\mathbf{x}) \text{ with } \mathbf{x}^\mathsf{T}\mathsf{Q}_x\mathbf{x} + \mathbf{d}^\mathsf{T}\mathbf{z}(\mathbf{x}) - \tfrac{1}{4}\|\mathbf{w}(\mathbf{x})\|^2 \ge \tau \right\}, \tag{39}$$

where the decision variables $\mathbf{z} : \mathbb{R}^{n_1} \to \mathbb{R}^k$ and $\mathbf{w} : \mathbb{R}^{n_1} \to \mathbb{R}^r$ are functions.

In the ARO problem the variables $\mathbf{z}(\mathbf{x})$ and $\mathbf{w}(\mathbf{x})$ take the role of the second-stage variables, the decision vector $\mathbf{x}$ takes the role of the uncertainty parameter vector and its former feasible set $\mathcal{X}$ becomes the uncertainty set. If the adjustable variables are restricted to a quadratic and affine policy respectively, i.e., $(\mathbf{z}(\mathbf{x}))_j = \mathbf{x}^\mathsf{T}\mathsf{Z}_j\mathbf{x} + \mathbf{x}^\mathsf{T}\mathbf{z}_j + z_j,\ j \in [1\!:\!k]$, $\mathbf{w}(\mathbf{x}) = \mathsf{W}\mathbf{x} + \mathbf{w}$, then all the semi-infinite constraints become quadratic in $\mathbf{x}$ and are thus amenable to a reformulation based on the general strategy. Since the application of the policies contracts the feasible set of the supremum problem, we generate a lower bound.

The authors test the resulting lower bound against lower bounds based on relaxation of the completely positive reformulation from Theorem 4 on random instances with $\mathcal{X} := \{\mathbf{x} \in \mathcal{K} : \mathsf{B}\mathbf{x} = \mathbf{c}\}$ given by a compact conic intersection. The results are mixed, but it is noted that in case the number of constraints in $\mathcal{Y}$ is much bigger than the number of linear equality constraints in $\mathcal{X}$, the ARO lower bound has computational advantages. Currently, a direct real-world application of this model is not in sight, but we are confident that future research will reveal relevant areas where we can profit from the strength of the ARO lower bound, and also ways to exploit the special structure of the lower bound. The interested reader may inspect these structural details in Bomze & Gabl (2021).

4.5. Outlook on new research direction: robust convex optimization

Recently (Bertsimas, den Hertog, Pauphilet, & Zhen, 2022) introduced a reformulation of a general robust convex constraint into a robust bilinear constraint. The argument rests on the characterization of a closed, convex function as the bi-dual conjugate, hence the conjugate function of its conjugate function (see Rockafellar, 2015, Section 12). For the readers' convenience we repeat their derivation here. So consider a robust constraint

$$h(\mathsf{A}(\mathbf{x})\mathbf{u}+\mathbf{b}(\mathbf{x})) \leq 0\,, \quad \mathbf{u}\in\mathcal{U}, \tag{40}$$

for a convex function $h:\mathbb{R}^m\to[-\infty,\infty]$ and appropriate affine functions $\mathsf{A}:\mathbb{R}^n\to\mathbb{R}^{m\times q}$ and $\mathbf{b}:\mathbb{R}^n\to\mathbb{R}^m$. Checking whether the constraints hold for an $\mathbf{x}\in\mathbb{R}^n$ is hard since it is equivalent to finding the supremum of a convex function. We now use the fact that any closed, convex function can be characterized as the bi-dual convex conjugate, hence

$$h(\mathbf{x}) = h^{**} = \sup_{\mathbf{w}\in\mathrm{dom}h^*}\left\{\mathbf{x}^\mathsf{T}\mathbf{w}-h^*(\mathbf{w})\right\}$$

where $h^*(\mathbf{y}) := \sup_{\mathbf{w}\in\mathrm{dom}h}\left\{\mathbf{y}^\mathsf{T}\mathbf{w}-h(\mathbf{w})\right\}$ is the convex conjugate of h (see Rockafellar, 2015 for a detailed explanation). We can now reformulate the implied supremum problem in (40) into a bilinear problem:

$$\begin{aligned}\sup_{\mathbf{u}\in\mathcal{U}}\{h(\mathsf{A}(\mathbf{x})\mathbf{u}+\mathbf{b}(\mathbf{x}))\} &= \sup_{\mathbf{u}\in\mathcal{U}}\sup_{\mathbf{w}\in\mathrm{dom}h^*}\left\{\mathbf{u}^\mathsf{T}\mathsf{A}(\mathbf{x})^\mathsf{T}\mathbf{w}+\mathbf{b}(\mathbf{x})^\mathsf{T}\mathbf{w}-h^*(\mathbf{w})\right\}\\ &= \sup_{(\mathbf{w},w_0,\mathbf{u})\in\mathcal{U}'}\mathsf{A}(\mathbf{x})\bullet\mathbf{u}\mathbf{w}^\mathsf{T}+\mathbf{b}(\mathbf{w})^\mathsf{T}\mathbf{w}-w_0\end{aligned}$$

with $\mathcal{U}' := \{(\mathbf{w},w_0,\mathbf{u}) : \mathbf{u}\in\mathcal{U},\, \mathbf{w}\in\mathrm{dom}h^*,\, h^*(\mathbf{w})\leq w_0\}$. The authors of Bertsimas et al. (2022) proceed to further reformulate using an exact lifting, i.e.,

$$\sup_{(\mathbf{w},w_0,\mathbf{u},\mathsf{V})}\Big\{\mathsf{A}(\mathbf{x})\bullet\mathsf{V}+\mathbf{b}(\mathbf{x})^\mathsf{T}\mathbf{w}-w_0 : \mathsf{V}=\mathbf{u}\mathbf{w}^\mathsf{T},\ \mathbf{u}\in\mathcal{U},\ \mathbf{w}\in\mathrm{dom}h^*,\, h^*(\mathbf{w})\leq w_0\Big\}$$

and define Θ to be the (nonconvex) feasible set of the latter supremum problem. By providing convex supersets $\bar{\Theta}\supseteq\Theta$ they achieve safe approximations of the original robust constraint that take the form

$$\mathsf{A}(\mathbf{x})\bullet\mathsf{V}+\mathbf{b}(\mathbf{x})^\mathsf{T}\mathbf{w}-w_0\leq 0 \quad \text{for all } (\mathbf{w},w_0,\mathbf{u},\mathsf{V})\in\bar{\Theta}$$

and can be reformulated using standard techniques based on convex duality.

However, here we want to demonstrate another application of our approach, focusing on the bilinear reformulation. It is clear that this formulation is amenable to the general strategy if we can find characterizations of $\mathcal{G}(\mathcal{U}')$. A lot of interesting questions are worth investigation with regard to such an approach.

The obvious problem is finding $\mathcal{G}(\mathcal{U}')$, which may be an unattainable goal for most instances of $h(\cdot)$, but for some instances is, perhaps, just a matter of cleverly exploiting known results. In complicated cases, one may resort to simpler approximations of $\mathcal{U}'_{inner}\subseteq\mathcal{U}'\subseteq\mathcal{U}'_{outer}$ in order to obtain manageable approximations $\mathcal{G}(\mathcal{U}'_{inner})\subseteq\mathcal{G}(\mathcal{U}')\subseteq\mathcal{G}(\mathcal{U}'_{outer})$. Ideally, one can look for performance guarantees of such approximations. Since the copositive approach can manage discontinuous quadratic optimization problems, it is also reasonable to ask if the above approach can be extended to discontinuous compositions of convex functions under a copositive optimization paradigm. Also, in the light of Lemmas 16 and Theorem 17 the copositive approach can be expected to be able manage cases where the argument of $h(\cdot)$ is a convex quadratic function.

Further, it would be interesting to study the relationship between said approximations and the approach from Bertsimas et al. (2022) mentioned above. Specifically, their approach might inspire approximations of $\mathcal{G}(\mathcal{U}')$ that can be used in other contexts. We would be interested to cooperate towards this goal, as it has significant overlap with our research agenda.

4.6. Open problems

Robust convex optimization: The entirety of the discussion of Section 4.5 is preliminary and hopefully inspires some readers to take up the questions we outlined there.

Quality of the quadratic policy: Conditions under which affine decision rules are optimal are well understood (see Bertsimas & Goyal, 2012; Bertsimas, Iancu, & Parrilo, 2010; Iancu, Sharma, & Sviridenko, 2013). Further, it was shown in Bemporad, Borrelli, & Morari (2003); Ben-Tal, El Housni, & Goyal (2020) that the optimal set of an adjustable robust optimization problems contains piece-wise affine policies in many interesting cases. Clearly, it follows that the same results hold for the quadratic and piecewise quadratic decision rules, since they contain affine policies as a special case. However, it would be interesting to give conditions under which the quadratic policy performs provably better than the affine ones, or, on top of that, if there are cases where a quadratic policy is provably optimal among all (also non-quadratic) policies while all affine ones are provably suboptimal.

Breakpoints and folding directions of the piecewise quadratic policy: In Section 4.3.3 we presented results that allow for the implementation of a piecewise decision rule. This is a significant technique since the set of optimal decision rules of an adjustable robust optimization problem is known to contain piecewise linear decision rules. However, in order to implement the technique in the aforementioned section, one needs to fix the number of pieces as well as the folding directions and breakpoints beforehand. We currently do not know whether there is a preferable way of making that choice, or whether there is a way to update a such a choice once the solution under that choice is known.

Probabilistic guarantees under structured uncertainty sets: The discussion in the previous sections shows that copositive optimization techniques allow for the application of uncertainty sets that are not applicable under the standard convex duality-based paradigm. In Sections 4.3.2, 4.3.3 this modeling capabilities were used in order to implement discrete uncertainty sets as well as uncertainty sets that allow for the application of piecewise decision rules. The motivation for the lattter is clear, the former is motivated by the fact that in many applications, the outcome of the uncertain process is most naturally described by a discrete set of alternatives. However, in literature we often find that specific uncertainty sets are motivated by the desire to give certain probabilistic guarantees that the robust solution does not violate an uncertain constraint, perhaps under some broad assumption on the distribution of the uncertain process. So far, we do not know an approach where the additional modeling power granted by the copositive approach was used in order to tighten such guarantees or give such guarantees under novel sets of assumptions.

5. ARO with uncertain right-hand side: an alternative copositive approach

In (Xu & Burer, 2018) the authors proposed a copositive reformulation of a certain class of linear ARO problems based on Theorem 4 obtained by means very different from the general strategy we outlined above, and consequently results in a distinctive type of copositive reformulation. The derivation is simple and elegant, and we will give a condensed account of their methodology in the sequel, extending their model by introducing additional uncertainty, also on the left-hand side and in the objective.

5.1. Copositive reformulation à la Xu and Burer

The class of ARO models we consider here is given by

$$\begin{aligned}&\min_{\mathbf{x}\in\mathcal{X},\mathbf{y}(\mathbf{u})}\left\{\max_{\mathbf{u}\in\mathcal{U}}\mathbf{c}^\mathsf{T}\mathbf{x}+\mathbf{d}(\mathbf{u})^\mathsf{T}\mathbf{y}(\mathbf{u}) : \mathbf{a}(\mathbf{x},\mathbf{u})+\mathsf{B}\mathbf{y}(\mathbf{u})\geq\mathbf{f}(\mathbf{u})\ \text{ for all } \mathbf{u}\in\mathcal{U}\right\}\\ &=\min_{\mathbf{x}\in\mathcal{X}}\mathbf{c}^\mathsf{T}\mathbf{x}+\max_{\mathbf{u}\in\mathcal{U}}\min_{\mathbf{y}}\left\{\mathbf{d}(\mathbf{u})^\mathsf{T}\mathbf{y} : \mathbf{a}(\mathbf{x},\mathbf{u})+\mathsf{B}\mathbf{y}\geq\mathbf{f}(\mathbf{u})\right\}\end{aligned}\tag{41}$$

where the latter reformulation is proved by using standard arguments from optimization theory. Also, $\mathbf{d}(\mathbf{u})=\mathbf{d}_0+\mathsf{D}\mathbf{u}$, $\mathbf{f}(\mathbf{u})=\mathbf{f}_0+\mathsf{F}\mathbf{u}$ for appropriate matrices and vectors, $\mathbf{a}(\mathbf{x},\mathbf{u}):=\mathbf{a}_0(\mathbf{x})+\mathsf{A}(\mathbf{x})\mathbf{u}$ for appropriate vector-valued, affine mappings, and $\mathcal{X}\subseteq\mathbb{R}^n$ is a feasible set of the first-stage decision not affected by uncertainty. Again, the uncertainty set is modeled as compact, conic intersection:

$$\mathcal{U}:=\left\{\mathbf{u} : (1,\mathbf{u}^\mathsf{T})^\mathsf{T}\in\mathcal{K}\right\},$$

for some closed, convex cone $\mathcal{K}\subseteq\mathbb{R}^{q+1}$ The reformulation strategy we are about to lay out rests on the following assumptions:

Assumption 1. For problem (41) the following statements hold:

(a) it is feasible with finite optimal value;
(b) it possesses relatively complete recourse, i.e., for all $\mathbf{x}\in\mathcal{X}$ and $\mathbf{u}\in\mathcal{U}$ the innermost LP (in the min-max-min reformulation) is feasible.

The innermost minimization problem can be dualized to obtain

$$\begin{aligned}&\min_{\mathbf{y}}\left\{\mathbf{d}(\mathbf{u})^\mathsf{T}\mathbf{y}(\mathbf{u}) : \mathsf{B}\mathbf{y}\geq\mathbf{f}(\mathbf{u})-\mathbf{a}(\mathbf{x},\mathbf{u})\right\}\\ &=\max_{\mathbf{w}\in\mathbb{R}^m_+}\left\{\mathbf{w}^\mathsf{T}[]\mathbf{f}(\mathbf{u})-\mathbf{a}(\mathbf{x},\mathbf{u})] : \mathsf{B}^\mathsf{T}\mathbf{w}=\mathbf{d}(\mathbf{u})\right\}.\end{aligned}$$

We can now plug in the dual and the definitions of the functions representing the uncertain data, to obtain a bilinear optimization problem that can be reformulated into a set-completely positive optimization problem:

$$\begin{aligned}&\max_{(\mathbf{u},\mathbf{w})\in\mathcal{U}\times\mathbb{R}^m_+}\left\{\mathbf{w}^\mathsf{T}[\mathsf{F}+\mathsf{A}(\mathbf{x})]\mathbf{u}+[\mathbf{f}_0-\mathbf{a}_0(\mathbf{x})]^\mathsf{T}\mathbf{w} : \mathsf{B}^\mathsf{T}\mathbf{w}-\mathsf{D}\mathbf{u}=\mathbf{d}_0\right\}\\ &\max_{(u_0,\mathbf{u},\mathbf{w})\in\mathcal{K}\times\mathbb{R}^m_+}\left\{\mathbf{w}^\mathsf{T}[\mathsf{F}+\mathsf{A}(\mathbf{x})]\mathbf{u}+u_0[\mathbf{f}_0-\mathbf{a}_0(\mathbf{x})]^\mathsf{T}\mathbf{w} : \mathsf{D}\mathbf{u}\right.\\ &\qquad\left.+\mathbf{d}_0u_0-\mathsf{B}^\mathsf{T}\mathbf{w}=\mathbf{o}, u_0=1\right\}\\ &=\max_{\mathbf{z}\in\mathcal{K}\times\mathbb{R}^m_+}\left\{\mathbf{z}^\mathsf{T}\mathsf{Q}(\mathbf{x})\mathbf{z} : z_0=1,\ \mathsf{E}\mathbf{z}=\mathbf{o}\right\}\\ &=\max_{\mathsf{Z},\mathbf{z}}\left\{\mathsf{Q}(\mathbf{x})\bullet\mathsf{Z} : (\mathsf{Z})_{00}=1,\ \mathsf{Z}\mathsf{E}^\mathsf{T}=\mathsf{O},\ \mathsf{Z}\in\mathcal{CPP}\left(\mathcal{K}\times\mathbb{R}^m_+\right)\right\}\end{aligned}$$

with

$$\mathsf{Q}(\mathbf{x}):=\begin{bmatrix}0 & \mathbf{o}^\mathsf{T} & \frac{1}{2}[\mathbf{f}_0-\mathbf{a}_0(\mathbf{x})]^\mathsf{T}\\ \mathbf{o} & \mathsf{O} & \frac{1}{2}[\mathsf{F}+\mathsf{A}(\mathbf{x})]^\mathsf{T}\\ \frac{1}{2}[\mathbf{f}_0-\mathbf{a}_0(\mathbf{x})] & \frac{1}{2}[\mathsf{F}+\mathsf{A}(\mathbf{x})] & \mathsf{O}\end{bmatrix},\ \mathsf{E}:=\left[\mathbf{d}_0,\mathsf{D},-\mathsf{B}^\mathsf{T}\right],\ \mathbf{z}:=\begin{bmatrix}u_0\\ \mathbf{u}\\ \mathbf{w}\end{bmatrix}.$$

Under Assumption 1a), the quadratic problem is feasible, since any $\mathbf{x}\in\mathcal{X}$ that would render it infeasible would be optimal for (41) with minus infinity as optimal value. Thus, the convex relaxation is exact by Theorem 4. Further, under the Assumption 1b) the dual of the innermost LP is feasible with finite value regardless of $\mathbf{u}$, hence attaining its optimal value on an extreme point of its feasible set. The latter set is polyhedral so that its extreme points can be contained in a ball of sufficient size, rendering $\mathbf{w}^\mathsf{T}\mathbf{w}\leq r_w$ redundant for the bilinear problem given that $r_w>0$ is large enough. Also, since $\mathcal{U}$ is bounded, the constraint $\mathbf{u}^\mathsf{T}\mathbf{u}\leq r_u$ is redundant for large enough $r_u>0$. It follows that we can always introduce the constraint $\mathbf{z}^\mathsf{T}\mathbf{z}\leq r$ with sufficiently large $r\geq 0$ to the bilinear optimization problem without changing the optimal value, hence $\mathsf{Z}\bullet\mathsf{I}\leq r$ is redundant for the conic optimization problem. After doing so, the dual of the conic problem is given by

$$\begin{aligned}&\min_{\lambda,\Lambda,\rho}\lambda+r\rho\\ &\text{s.t.}:\ \mathsf{Q}(\mathbf{x})+\lambda\mathbf{e}_1\mathbf{e}_1^\mathsf{T}+\frac{1}{2}\left(\Lambda\mathsf{E}+\mathsf{E}^\mathsf{T}\Lambda^\mathsf{T}\right)+\rho\mathsf{I}\in\mathcal{COP}\left(\mathbb{R}^m_+\times\mathcal{K}\right)\end{aligned}$$

and since for the identity matrix I we have $\mathsf{I}\in\operatorname{int}\mathcal{COP}(\mathcal{K})$ for any cone $\mathcal{K}$, the latter problem has a Slater point so that the duality gap is zero. Thus, the original optimization problem can be equivalently reformulated as

$$\begin{aligned}&\min_{\mathbf{x},\lambda,\Lambda,\rho}\ \mathbf{c}^\mathsf{T}\mathbf{x}+\lambda+r\rho\\ &\text{s.t.}:\ \mathsf{Q}(\mathbf{x})+\lambda\mathbf{e}_1\mathbf{e}_1^\mathsf{T}+\frac{1}{2}\left(\Lambda\mathsf{E}+\mathsf{E}^\mathsf{T}\Lambda^\mathsf{T}\right)+\rho\mathsf{I}\in\mathcal{COP}\left(\mathbb{R}^m_+\times\mathcal{K}\right),\ \mathbf{x}\in\mathcal{X}.\end{aligned}\tag{42}$$

The reformulation is exact but the cone $\mathcal{COP}(\mathbb{R}^m_+\times\mathcal{K})$ is intractable even if $\mathcal{COP}(\mathcal{K})$ is tractable. Hence one has to resort again to inner, hence conservative, approximations.

5.2. Improving the affine policy

This raises the question whether any benefit can be incurred by this strategy when compared to other conservative approximations such as the ones based on affine decision rules. The authors of Xu & Burer (2018) find an elegant answer to this question, at least for the case where $\mathbf{d}(\mathbf{u})$ is constant. We summarize their findings in the following theorem:

Theorem 19. *For (41) assume that $\mathbf{d}(\mathbf{u})$ is a constant. Further denote by v_{Aff} the optimal value of (41) under an affine policy and with v_{IA} the optimal value of (42) after replacing $\mathcal{COP}\left(\mathbb{R}^m_+\times\mathcal{K}\right)$ with*

$$\mathcal{IA}\left(\mathcal{K}\times\mathbb{R}^m_+\right):=\left\{\mathsf{S}=\begin{bmatrix}\mathsf{S}_{11} & \mathsf{S}_{21}^\mathsf{T}\\ \mathsf{S}_{21} & \mathsf{S}_{22}\end{bmatrix} : \begin{array}{l}\mathsf{S}_{11}=\mathbf{e}_1\mathbf{g}^\mathsf{T}+\mathbf{g}\mathbf{e}_1^\mathsf{T},\ \mathbf{g}\in\mathcal{K},\\ \text{Rows}(\mathsf{S}_{21})\in\mathcal{K}^*,\ \mathsf{S}_{22}\geq 0\end{array}\right\}.$$

Then we have $v_{\text{IA}}\leq v_{\text{Aff}}$. Further, denoting by v_{IB} the optimal value of (42) where $\mathcal{COP}\left(\mathbb{R}^m_+\times\mathcal{K}\right)$ is replaced with any cone $\mathcal{IB}\left(\mathcal{K}\times\mathbb{R}^m_+\right)$ for which

$$\mathcal{COP}\left(\mathbb{R}^m_+\times\mathcal{K}\right)\supseteq\mathcal{IB}\left(\mathbb{R}^m_+\times\mathcal{K}\right)\supseteq\mathcal{IA}\left(\mathbb{R}^m_+\times\mathcal{K}\right)\tag{43}$$

holds, we have $\text{val}(42)\leq v_{\text{IB}}\leq v_{\text{Aff}}$, where the first inequality holds even if $\mathbf{d}(\mathbf{u})$ is not constant.

The authors of Xu & Burer (2018) propose the following candidate for $\mathcal{IB}\left(\mathbb{R}^m_+\times\mathcal{K}\right)$ with the desired property, namely

$$\mathcal{IB}\left(\mathbb{R}^m_+\times\mathcal{K}\right):=\left\{\mathsf{S}+\mathsf{M}+\mathsf{R} : \begin{array}{l}\mathsf{S}\in\mathcal{IA}\left(\mathbb{R}^m_+\times\mathcal{K}\right),\ \mathsf{M}\in\mathcal{S}^{m+q}_+,\\ \mathsf{R}_{11}\in\mathcal{C},\ \mathsf{R}_{21}=\mathsf{O},\ \mathsf{R}_{22}=\mathsf{O}\end{array}\right\}\tag{44}$$

where $\mathcal{C}\subseteq\mathcal{COP}(\mathcal{K})$, can be replaced by inner approximations such as the ones discussed in Section 2.3.1. The authors also show in numerical experiments, that the optimal value v_{IB} obtained by using the above cone can be strictly smaller than v_{Aff}.

5.3. Open problems

Improving the affine policy under nonconstant $\mathbf{d}(\mathbf{u})$: In the original article (Xu & Burer, 2018), both $\mathbf{d}(\mathbf{u})$ as well as $\mathsf{A}(\mathbf{u})$ were assumed constant. While establishing the above theorem for the case where the latter function is not constant is simply a matter of carrying along some additional terms through the discussion presented in Xu & Burer (2018), the same is not true for non-constant $\mathbf{d}(\mathbf{u})$. The reason for this lies in the proof strategy that achieves $v_{\mathrm{IA}} \le v_{\mathrm{Aff}}$. It is based on first deriving the finite reformulation of (41) under an affine policy, and then showing that one can turn any feasible solution of that reformulation into a feasible solution of (42), under the required specifications. The finite reformulation under the affine policy is thereby achieved using standard linear convex duality, which does not apply if $\mathbf{d}(\mathbf{u})$ is not constant.
It is however possible to give a finite reformulation based on the general strategy as discussed in Section 4.2.1. It remains to clarify how the resulting reformulation can be projected into the feasible set of (42). Answering this question, one may be able to find a modification of $\mathcal{IA}\left(\mathbb{R}^m_+ \times \mathcal{K}\right)$ that allows for similar performance guarantees.

Characterizing implied policies: As noted in Xu & Burer (2018), (42) is powerful enough to represent the original ARO problem, and by the discussion in Section 5.2 we see that the affine policy can be mapped into the solution space of the reformulation. However, there is no similar analysis regarding other types of policies.

Improving the quadratic policy: On a related note, it is not clear whether the conic constraint in (42) can be replaced by an inner approximation that performs at least as good as the quadratic policy. Such a result seems tangible since we know from the discussion in Section 4.2.1 that (41) under a quadratic policy does have a conic reformulation, where each row of the constraints is reformulated individually, resulting in a collection of conic constraints. However, there seems to be no straightforward way in which the feasible set of such a reformulation can be projected into the feasible set of (42).

The case of uncertain recourse: The reformulation presented in the above section assumes that the matrix B is not affected by uncertainty. If this were the case, we would have to deal with quadratic constraints, which would jeopardize the application of Theorem 4 at the penultimate step of the reformulation strategy. Specifically, instead of the linear constraints $\mathsf{D}\mathbf{u} + \mathbf{d}_0 u_0 - \mathsf{B}^\mathsf{T}\mathbf{w} = \mathbf{o}$ we would have to deal with the constraint $\mathsf{D}\mathbf{u} + \mathbf{d}_0 u_0 - (\mathsf{B}(\mathbf{u}))^\mathsf{T}\mathbf{w} = \mathbf{o}$ which is bilinear in case $\mathsf{B}(\mathbf{u})$ is linear in $\mathbf{u}$. Theorem 4 does not place any restrictions on linear constraints, but quadratic ones have to respect the key condition, in order for the relaxation to be exact. Hence, the case of uncertain recourse could be tackled if the problem data is such that the key condition is either satisfied or can be relaxed, e.g., as in Bomze & Peng (2022). However, we do not know whether either of these strategies are feasible for interesting instances of (41) with uncertain recourse.

6. Robust standard quadratic optimization

Standard quadratic optimization deals with minimizing an indefinite quadratic form over the standard simplex (also known as the probability simplex) $\Delta^n := \left\{\mathbf{x} \in \mathbb{R}^n_+ : \mathbf{e}^\mathsf{T}\mathbf{x} = 1\right\}$, i.e.,

$$\min_{\mathbf{x}\in\mathbb{R}^n_+} \left\{\mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} : \mathbf{e}^\mathsf{T}\mathbf{x} = 1\right\} = \min_{\mathbf{x}\in\Delta^n} \mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} \tag{45}$$

Despite its simplicity, this *Standard Quadratic Problem (StQP)* features prominently in diverse application areas such as game theory, graph theory, machine learning and copositivity detection. It was the first problem for which an exact copositive reformulation was presented in Bomze et al. (2000):

$$\min_{\mathbf{x}\in\Delta^n} \mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} = \min_{\mathsf{X}\in\mathcal{CPP}(\mathbb{R}^n_+)} \{\mathsf{Q}\bullet\mathsf{X} : \mathsf{E}\bullet\mathsf{X} = 1\},$$

or in other words it holds that $\mathcal{G}(\Delta^n) = \{\mathsf{X}\in\mathcal{CPP}(\mathbb{R}^n_+) : \mathsf{E}\bullet\mathsf{X} = 1\}$. While the original proof is straightforward, this by now classical result can also easily be derived via the methodology discussed in Section 2.1.2. Specifically, one can apply Theorem 6 where the $\mathbb{J}$ is chosen to be all of $\mathcal{CPP}(\mathbb{R}^n_+)$ and $\mathbb{H}$ is the hyperplane associated with $\mathsf{E}\bullet\mathsf{X} = 1$; the details are left to the reader.

In (Bomze, Kahr, & Leitner, 2021b) the authors investigate robust counterparts of this problem, which are generically given by

$$\min_{\mathbf{x}\in\Delta^n} \max_{\mathsf{U}\in\mathcal{U}} \mathbf{x}^\mathsf{T}(\mathsf{Q}+\mathsf{U})\mathbf{x}\,.$$

Since the constraints are a structural aspect of the problem (e.g., probabilities are are always positive and sum up to one), only the objective function is affected by uncertainty. An immediate question is whether the completely positive relaxation given by

$$\min_{\mathsf{X}\in\mathcal{G}(\Delta^n)} \max_{\mathsf{U}\in\mathcal{U}} (\mathsf{Q}+\mathsf{U})\bullet\mathsf{X},$$

is again tight. While the answer is negative in general, the authors prove that the relaxation gap is exactly the min-max gap.

Theorem 20. *Consider the robust Standard Quadratic Problem with uncertainty set $\mathcal{U}$.*

(a) For general $\mathcal{U}$ we have

$$\min_{\mathbf{x}\in\Delta^n} \max_{\mathsf{U}\in\mathcal{U}} \mathbf{x}^\mathsf{T}(\mathsf{Q}+\mathsf{U})\mathbf{x} \le \min_{\mathsf{X}\in\mathcal{G}(\Delta^n)} \max_{\mathsf{U}\in\mathcal{U}} (\mathsf{Q}+\mathsf{U})\bullet\mathsf{X}\,.$$

(b) Suppose $\mathcal{U}$ is closed and convex. Then

$$\min_{\mathsf{X}\in\mathcal{G}(\Delta^n)} \max_{\mathsf{U}\in\mathcal{U}} (\mathsf{Q}+\mathsf{U})\bullet\mathsf{X} = \max_{\mathsf{U}\in\mathcal{U}} \min_{\mathbf{x}\in\Delta^n} \mathbf{x}^\mathsf{T}(\mathsf{Q}+\mathsf{U})\mathbf{x}\,,$$

so that the completely positive relaxation gap is exactly the gap in the min-max inequality.

Proof. See Bomze et al. (2021b, Theorem 1). □

However, there are instances in which the inner maximization problem can be evaluated independently of $\mathbf{x}$. In these cases the robust counterpart reduces to a deterministic standard quadratic problem so that the exactness of the relaxation stays intact. The first set of instances for which this is the case are those where the uncertainty set is constructed via suitable cones.

Theorem 21. *Let $\mathcal{C}\subseteq\mathcal{COP}(\mathbb{R}^n)$ be a sub-cone of the cone of copositive matrices and $\mathsf{L},\mathsf{B}\in\mathcal{S}^n$ be given matrices. Assume that $\mathcal{U} = \{\mathsf{U} : \mathsf{U}-\mathsf{L}\in\mathcal{C},\ \mathsf{B}-\mathsf{U}\in\mathcal{C}\}$. Then the completely positive relaxation is an exact reformulation and the robust counterpart reduces to a standard quadratic problem with data $\mathsf{Q}+\mathsf{L}$.*

Proof. See Bomze et al. (2021b, Corollary 1). □

Choices for $\mathcal{C}$ that fulfill the assumption of the above theorem are

$$\mathcal{C}\in\{\mathcal{CPP}(\mathcal{K}_1),\ \mathcal{N}^n,\ \mathcal{DNN}^n,\ \mathcal{NND}^n,\ \mathcal{COP}(\mathcal{K}_2)\}, \tag{46}$$

where $\mathcal{K}_1\subseteq\mathbb{R}^n$ and $\mathbb{R}^n_+\subseteq\mathcal{K}_2$. Besides this, there is another interesting case of a different kind for which a similar result can be derived, namely the case of ellipsoidal uncertainty.

Theorem 22. *Let* $\mathsf{C} \in \mathbb{R}^{n\times n}$ *be a nonsingular matrix and define, for some scalar* $\rho > 0$, *the uncertainty set* $\mathcal{U} := \left\{\mathsf{U} \in \mathcal{S}^n : \|\mathsf{C}^\top \mathsf{U}\mathsf{C}\|_F \le \rho\right\}$. *Then*

$$\min_{\mathsf{X}\in\mathcal{G}(\Delta^n)} \max_{\mathsf{U}\in\mathcal{U}} (\mathsf{Q}+\mathsf{U}) \bullet \mathsf{X} = \min_{\mathsf{X}\in\mathcal{G}(\Delta^n)} \left(\mathsf{Q} - \rho\left(\mathsf{C}\mathsf{C}^\top\right)^{-1}\right) \bullet \mathsf{X}, \tag{47}$$

i.e., the robust counterpart reduces to a single standard quadratic optimization problem.

6.1. Open problems

Robust properties of StQP solutions: Many problems in graph theory such as the stable-set problem and the maximum-clique problem have a reformulation as an StQP, and one can infer the solution to these problems from the optimal solution of the respective StQP. It is however, unclear if these properties also hold for the robustified StQP. Conversely it is not known whether robust versions of the stable-set problem or the maximum-clique problem can be modelled as a robust StQP or, perhaps, a generalization of the latter model.

Convexified robust StQP: While the above discussion presents cases where the robust StQP reduces to an instance of StQP which can be tackled via standard convexification approaches, a general convexification approach applicable outside of these special cases is not known. The complication arises from the fact that the pointwise maximum of linear function is itself not linear but convex, and convex functions may attain their optimum at points which are not extreme. Hence, $\mathcal{G}(\Delta^n)$ fails to deliver the effectiveness we enjoy in the deterministic case.

7. Two-stage stochastic optimization for StQPs

In (Bomze, Gabl, Maggioni, & Pflug, 2021a) the authors considered a two-stage stochastic version of the StQP.

Stochastic optimization deals with optimizing expected outcomes of uncertain optimization problems, i.e.,

$$\min_{\mathbf{x}\in\mathcal{X}} \{\mathbb{E}_{\tilde{\mathbf{u}}}(f(\mathbf{x},\tilde{\mathbf{u}}))\}. \tag{48}$$

where the expected value is taken with respect to the random vector $\tilde{\mathbf{u}}$, which is defined by a known probability space $(\Xi, \mathcal{A}, P)$ with support Ξ, probability distribution P and σ-field $\mathcal{A}$. Analogously to the adjustable robust setting, in two stage stochastic optimization one seeks a decision on the first-stage variables and on a second-stage decision rule that adapts to the realization of the random event. Thus, we are considering

$$\min_{\mathbf{x}\in\mathcal{X}} \left\{ f_1(\mathbf{x}) + \mathbb{E}_{\tilde{\mathbf{u}}}\left[\min_{\mathbf{y}\in\mathcal{Y}(\mathbf{x},\tilde{\mathbf{u}})} f_2(\mathbf{x},\mathbf{y},\tilde{\mathbf{u}})\right]\right\}. \tag{49}$$

Here we make a choice on the first-stage variables $\mathbf{x}$ and second-stage policies $\mathbf{y}$ so that we optimize the sum of a deterministic first-stage outcome and the expected value of the optimal second-stage choice. Note, that the innermost minimization problem depends on the random vector $\tilde{\mathbf{u}}$ so that the decision vector $\mathbf{y}$ is implicitly a function of $\tilde{\mathbf{u}}$. Hence, the setting is indeed analogous to the adjustable robust setting. However, for our purposes it will not be necessary to model $\mathbf{y}$ explicitly as a function as it is done in ARO. Also, in our case the constraints linking $\mathbf{y}$ to $\mathbf{x}$ are not uncertain but deterministic: $\mathbf{y} \in \mathcal{Y}(\mathbf{x})$.

Here we are dealing with the special case of the (typically non-convex) StQP of the form

$$\min_{\mathbf{z}\in\Delta^n} \mathbf{z}^\top \tilde{\mathsf{Q}}\mathbf{z}\,,$$

where uncertainty is considered only in the objective function. Suppose a (possibly) small $n_1 \times n_1$ principal submatrix A of $\tilde{\mathsf{Q}}$ is known (more or less) exactly whereas the rest of the problem data are subject to uncertainty with known probability distribution:

$$\tilde{\mathsf{Q}} = \begin{bmatrix} \mathsf{A} & \tilde{\mathsf{B}}^\top \\ \tilde{\mathsf{B}} & \tilde{\mathsf{C}} \end{bmatrix}. \tag{50}$$

Here $\tilde{\mathbf{u}} = [\tilde{\mathsf{B}}, \tilde{\mathsf{C}}]$ are the uncertain data. Such a situation may for example arise in portfolio optimization, when the relevant statistics on younger securities can be assessed less accurately due to lack of historical data.

Decomposing $\mathbf{z} \in \mathbb{R}^n$ via $\mathbf{z}^\top = [\mathbf{x}^\top, \mathbf{y}^\top]$ with $(\mathbf{x},\mathbf{y}) \in \mathbb{R}^{n_1}_+ \times \mathbb{R}^{n_2}_+$ and $n_1 + n_2 = n$ we arrive at the following problem reformulation of the (random) objective function as

$$q(\mathbf{z},\tilde{\mathbf{u}}) = \mathbf{z}^\top \tilde{\mathsf{Q}}\mathbf{z} = \mathbf{x}^\top \mathsf{A}\mathbf{x} + 2\mathbf{x}^\top \tilde{\mathsf{B}}^\top \mathbf{y} + \mathbf{y}^\top \tilde{\mathsf{C}}\mathbf{y}\,.$$

Taking the expectation with respect to the probability distribution of ξ, we obtain the so-called *recourse function*

$$r(\mathbf{x}) := \mathbb{E}_{\tilde{\mathbf{u}}}\left\{\min_{\mathbf{y}\in\mathbb{R}^{n2}_+}\left[2\mathbf{x}^\top \tilde{\mathsf{B}}^\top \mathbf{y} + \mathbf{y}^\top \tilde{\mathsf{C}}\mathbf{y} : \mathbf{e}^\top \mathbf{y} = 1 - \mathbf{e}^\top \mathbf{x}\right]\right\}$$

and the two-stage stochastic StQP can be formulated as follows:

$$\min_{\mathbf{x}\in T^{n_1}} \left\{s(\mathbf{x}) := \mathbf{x}^\top \mathsf{A}\mathbf{x} + r(\mathbf{x})\right\},$$

with $T^{n_1} = \operatorname{conv}\left\{\mathbf{o}, \mathbf{e}^i : i \in [1:n_1]\right\} = \operatorname{conv}(\Delta^{n_1} \cup \{\mathbf{o}\})$.

In most cases, a two-stage stochastic problem cannot be solved directly, since merely evaluating the expected value can be intractable. Thus, in practice one resorts to approximating the true uncertainty measure by a finite discretization. This gives rise to the so-called scenario problem, which in our case is given by:

$$\begin{aligned} \min_{\mathbf{x},\mathbf{y}_1,\dots,\mathbf{y}_S} \ & \mathbf{x}^\top \mathsf{A}\mathbf{x} + \sum_{s=1}^{S} p_s(2\mathbf{x}^\top \tilde{\mathsf{B}}_s^\top \mathbf{y}_s + \mathbf{y}_s^\top \tilde{\mathsf{C}}_s \mathbf{y}_s) \\ & \mathbf{e}^\top \mathbf{x} + \mathbf{e}^\top \mathbf{y}_s = 1, \quad s \in [1:S], \\ & \mathbf{y}_s \ge \mathbf{o}, \quad s \in [1:S], \\ & \mathbf{x} \ge \mathbf{o}. \end{aligned} \tag{51}$$

As we can see, the discretization is achieved by condensing the true probability measures to a set of S scenarios with associated probabilities p_s $s \in [1:S]$. There are many schemes on how to obtain these discretizations, and it would be beyond the scope of this text to dicuss them here; the interested reader may consult the references given in Bomze et al. (2021a, Section 2). Other techniques are preoccupied with reducing the size of an existing discretization, in order to obtain a more manageable problem size. For example Bomze et al. (2021a) employed a dissection technique to the discretized probability measure. In essence, scenarios are grouped together into m groups. Then the smaller scenario problems, that only involve scenarios from one group at time, are solved using probabilities conditional on the respective group. The so obtained solutions are averaged, with weights given by the probability of the respective group, in order to obtain a lower bound on the scenario problem. By varying the size of the groups one can trade-off accuracy against the benefit of having to solve smaller problems.

Since (51) describes a class of non-convex QCQPs, which contains the StQP as a special case, it is NP-hard. However, it clearly is amenable to a convex reformulation based on Theorem 4. Such a reformulation would involve the cone $\mathcal{CPP}(\mathbb{R}^{n_1+Sn_2+1}_+)$, hence a lifting in a space that scales quadratically with S. This is problematic as the quality of the approximations yielded by the scenario problem depends on the number of scenarios considered. As a consequence, the classical convex reformulation becomes impractical for those very cases where the scenario problem is relevant, namely when S is large. The authors of Bomze et al. (2021a) therefore propose an alternative, albeit weaker, relaxation that scales linearly with S:

Theorem 23. *Consider the problem*

$$\min_{\mathbf{x},\mathbf{y}_1,\dots,\mathbf{y}_S} \mathbf{x}^\mathsf{T}\mathsf{A}\mathbf{x} + \mathbf{a}^\mathsf{T}\mathbf{x} + \sum_{s=1}^{S} \mathbf{x}^\mathsf{T}\mathsf{B}_s\mathbf{y}_s + \mathbf{y}_s^\mathsf{T}\mathsf{C}\mathbf{y}_s + \mathbf{c}_s^\mathsf{T}\mathbf{y}_s$$
$$\text{s.t.}: \mathbf{e}^\mathsf{T}\mathbf{x} + \mathbf{e}^\mathsf{T}\mathbf{y}_s = 1\,, \quad s \in [1\!:\!S]\,,$$
$$\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_S \geq \mathbf{o}\,. \tag{52}$$

The following conic optimization problem gives a lower bound, and if $\mathbf{c}_s = \alpha_s\mathbf{e},\ \mathsf{B}_s = \mathbf{b}_s\mathbf{e}^\mathsf{T},\ \alpha_s \in \mathbb{R},\ \mathbf{b}_s \in \mathbb{R}^{n_1},\ i \in [1\!:\!S]$ *the bound is actually tight:*

$$\min_{\mathsf{X},\mathsf{Y}_s,\mathsf{Z}_s,\mathbf{y}_s} \mathsf{A}\bullet\mathsf{X} + \mathbf{a}^\mathsf{T}\mathbf{x} + \sum_{s=1}^{S} \mathsf{B}_s\bullet\mathsf{Z}_s + \mathsf{C}_s\bullet\mathsf{Y}_s + \mathbf{c}_s^\mathsf{T}\mathbf{y}_s$$
$$\text{s.t.}: \quad \mathbf{e}^\mathsf{T}\mathbf{x} + \mathbf{e}^\mathsf{T}\mathbf{y}_s = 1\,, \quad s \in [1\!:\!S]\,,$$
$$\mathsf{E}\bullet\mathsf{X} + 2\,\mathsf{E}\bullet\mathsf{Z}_s + \mathsf{E}\bullet\mathsf{Y}_s = 1\,, \quad s \in [1\!:\!S]\,,$$
$$\begin{bmatrix} 1 & \mathbf{x}^\mathsf{T} & \mathbf{y}_s^\mathsf{T} \\ \mathbf{x} & \mathsf{X} & \mathsf{Z}_s^\mathsf{T} \\ \mathbf{y}_s & \mathsf{Z}_s & \mathsf{Y}_s \end{bmatrix} \in \mathcal{CPP}(\mathbb{R}_+^{n_1+n_2+1}), \quad s \in [1\!:\!S]\,. \tag{53}$$

Compared to the classical reformulation which involves $\mathcal{CPP}(\mathbb{R}_+^{n_1+Sn_2+1})$, the above relaxation merely exhibits S conic constraints involving $\mathcal{CPP}(\mathbb{R}_+^{n_1+n_2+1})$, hence growing linearly with S. This advantage comes at the cost of losing the exactness, so that outside of the special cases mentioned in the theorem, the conic problem only provides a lower bound. However, numerical experiments conducted in Bomze et al. (2021a), comparing the bounds obtained by solving the $\mathcal{DNN}$-relaxation of both the classical reformulation and (53), suggest that the gap between the two tends to be very small. In fact the gap is so small that the authors hinted at the possibility that it is merely a numerical artefact. The reduction of computational effort on the other hand is substantially in favor of the lower-dimensional bound.

We also would like to stress that the proof of Theorem 23 relies heavily on the theory laid out in Section 2. By replacing the $\mathcal{CPP}$ constraint with a more complicated conic constraint in a follow-up paper (Gabl, 2022), it is even possible to close the relaxation gap between (52) and (53) entirely. Among the two proofs of this result, one follows the strategy described in Section 2.1.2. The conic constraint used in Gabl (2022) is a structured generalization of $\mathcal{CPP}$-type cones and can be approximated via similar means.

Another interesting feature of the methodology proposed in Bomze et al. (2021a) was their combination of upper bounds obtained by relaxations, first-order methods and global optimization solvers. As it turns out, (53) preserves the original space of variables and thus yields not only a lower but also an upper bound. This feasible solution can be used as starting point for local algorithms such as the pairwise Frank–Wolfe algorithm, or for global solvers such as `Gurobi`. The quality of these refined upper bounds can then be assessed relative to the lower bound obtained by the relaxation. As numerical experiments suggest, optimality gaps can be reduced substantially and with reasonable computational effort, and moreover the combination of procedures yields better results than each method would produce on their own.

7.1. Open problems

Efficacy of the sparse model: As stated before, the bounds produced by applying the $\mathcal{DNN}$-relaxation to (53) are almost identical to the ones obtained from the $\mathcal{DNN}$-relaxation of the classical model based on Theorem 4. Based on the experiments in Bomze et al. (2021a), we cannot rule out the possibility that the sparse relaxation is in fact tight. However, despite some effort in Gabl (2022), no such result was found so far, nor were the authors able to produce a counterexample.

8. Mixed-binary linear optimization problems under objective uncertainty

In (Natarajan, Teo, & Zheng, 2011) the authors considered the following optimization problem

$$Z(\tilde{\mathbf{c}}) := \max_{\mathbf{x}\in\mathbb{R}_+^n} \left\{ \tilde{\mathbf{c}}^\mathsf{T}\mathbf{x} : \mathsf{A}\mathbf{x} = \mathbf{b},\ x_j \in \{0,1\},\ j \in \mathcal{B} \right\}$$

where $\tilde{\mathbf{c}}$ are uncertain objective function coefficients whose true probability distribution P is assumed to have support in $\mathbb{R}_+^n$ and apart from that is ambiguous up to its first two moments, the mean $\boldsymbol{\mu} := \mathbb{E}(\tilde{\mathbf{c}})$ and covariance matrix $\Sigma := \mathbb{E}(\tilde{\mathbf{c}}\tilde{\mathbf{c}}^\mathsf{T})$. The authors aim to give an upper bound on $\mathbb{E}_P[Z(\tilde{\mathbf{c}})]$ by considering

$$\sup_{\tilde{\mathbf{c}}\sim(\boldsymbol{\mu},\Sigma)^+} \mathbb{E}[Z(\tilde{\mathbf{c}})] \geq \mathbb{E}_P[Z(\tilde{\mathbf{c}})]$$

where $(\boldsymbol{\mu},\Sigma)^+$ is the set of all distributions with nonnegative support, mean $\boldsymbol{\mu}$ and covariance matrix Σ. While the approach seems related to the two-stage distributionally robust paradigm, since the decision variables are allowed to adjust to the outcome of the uncertainty the same way it would in a recourse problem, it is different in that we do not consider the worst-case distribution, but rather the best-case distribution. However, the worst-case interpretation remains valid if the underlying optimization problem already is a worst-case estimation, such as for the longest path problem. Another way to interpret this model is to see it as the second stage of a two-stage distributionally robust optimization problem where $Z(\tilde{\mathbf{c}})$ is the dual of recourse problem problem with uncertain right-hand sides (which is a valid interpretation if $\mathcal{B} = \emptyset$). Indeed both interpretations have featured in literature following up (Natarajan et al., 2011), which we will briefly discuss at the end of this section.

The authors approach this bound by providing a copositive reformulation of

$$\sup_{\tilde{\mathbf{c}}\sim(\boldsymbol{\mu},\Sigma)^+} \mathbb{E}\left[\max_{\mathbf{x}\in\mathbb{R}_+^n} \left\{ \tilde{\mathbf{c}}^\mathsf{T}\mathbf{x} : \mathsf{A}\mathbf{x} = \mathbf{b},\ x_j \in \{0,1\} \text{ for all } j \in \mathcal{B} \right\} \right] \tag{54}$$

which necessitates the following set of assumptions:

Assumption 2. The following statements on $Z(\tilde{\mathbf{c}})$ hold :

1. The set $(\boldsymbol{\mu},\Sigma)^+$ is nonempty.
2. $\mathbf{x} \in \mathbb{R}_+^n : \mathsf{A}\mathbf{x} = \mathbf{b}$ implies $x_j \leq 1,\ j \in \mathcal{B}$.
3. The feasible region of the inner maximization problem is nonempty and bounded.

Note that the first assumption holds exactly if

$$\begin{bmatrix} 1 & \boldsymbol{\mu}^\mathsf{T} \\ \boldsymbol{\mu} & \Sigma \end{bmatrix} \in \mathcal{CPP}(\mathbb{R}_+^{n+1}),$$

which is, of course, an NP-hard task unless $n+1 \leq 4$. The other two assumptions are checked easily, the second one can even be enforced generically by introducing additional constraints and slack variables (see our discussion succeeding Theorem 4).

The reformulation rests on a particular mixed-moment lifting. More precisely, let $\mathbf{x}(\mathbf{c})$ denote the optimal solution (or in case of non-uniqueness, a measurable selection from the set of optimal solutions) to $Z(\mathbf{c})$ where $\mathbf{c}$ is a realization of $\tilde{\mathbf{c}}$. Then $\mathbf{x}(\tilde{\mathbf{c}})$ is a random vector, and we define the random vector

$$\mathbf{y}(\tilde{\mathbf{c}}) := \begin{bmatrix} 1 \\ \tilde{\mathbf{c}} \\ \mathbf{x}(\tilde{\mathbf{c}}) \end{bmatrix} \in \mathbb{R}_+^{2n+1}$$

so that

$$\begin{bmatrix} 1 & \boldsymbol{\mu}^\top & \mathbf{p}^\top \\ \boldsymbol{\mu} & \Sigma & \mathsf{Y}^\top \\ \mathbf{p} & \mathsf{Y} & \mathsf{X} \end{bmatrix} := \mathbb{E}\left[\mathbf{y}(\tilde{\mathbf{c}})\mathbf{y}(\tilde{\mathbf{c}})^\top\right] = \begin{bmatrix} 1 & \mathbb{E}\left[\tilde{\mathbf{c}}^\top\right] & \mathbb{E}\left[\mathbf{x}(\tilde{\mathbf{c}})^\top\right] \\ \mathbb{E}[\tilde{\mathbf{c}}] & \mathbb{E}\left[\tilde{\mathbf{c}}\tilde{\mathbf{c}}^\top\right] & \mathbb{E}\left[\tilde{\mathbf{c}}\mathbf{x}(\tilde{\mathbf{c}})^\top\right] \\ \mathbb{E}[\mathbf{x}(\tilde{\mathbf{c}})] & \mathbb{E}\left[\mathbf{x}(\tilde{\mathbf{c}})\tilde{\mathbf{c}}^\top\right] & \mathbb{E}\left[\mathbf{x}(\tilde{\mathbf{c}})\mathbf{x}(\tilde{\mathbf{c}})^\top\right] \end{bmatrix}$$

Since $\mathbf{y}(\tilde{\mathbf{c}}) \in \mathbb{R}^{2n+1}_+$ almost surely, and since $\mathcal{CPP}(\mathbb{R}^d_+)$ is closed, by definition of an integral as limit of finite positively weighted sums, above matrix is clearly completely positive. Also, by construction the constraints

$$\mathsf{A}\mathbf{p} = b_i,\ \operatorname{diag}\left(\mathsf{A}\mathsf{X}\mathsf{A}^\top\right) = \mathbf{b}\circ\mathbf{b},\ X_{jj} = p_j,\ j \in \mathcal{B},$$

are clearly valid for the so constructed variables, and the objective can be restated as

$$\mathbb{E}[Z(\tilde{\mathbf{c}})] = \mathbb{E}\left[\tilde{\mathbf{c}}^\top \mathbf{x}(\tilde{\mathbf{c}})\right] = \mathsf{I} \bullet \mathsf{Y},$$

where the max-operator was dropped since $\mathbf{x}(\mathbf{c})$ is an optimal solution to $Z(\mathbf{c})$. Based on this lifting, the authors were able to prove the following theorem.

Theorem 24. *Under Assumption 2, Problem (54) is equivalent to*

$$\begin{aligned} \max_{\mathbf{p},\mathsf{X},\mathsf{Y}} \ & \mathsf{I}\bullet\mathsf{Y} \\ & \mathsf{A}\mathbf{p} = b_i, \\ \operatorname{diag}\left(\mathsf{A}\mathsf{X}\mathsf{A}^\top\right) &= \mathbf{b}\circ\mathbf{b}, \\ & X_{jj} = p_j,\ j\in\mathcal{B}, \\ \begin{bmatrix} 1 & \boldsymbol{\mu}^\top & \mathbf{p}^\top \\ \boldsymbol{\mu} & \Sigma & \mathsf{Y}^\top \\ \mathbf{p} & \mathsf{Y} & \mathsf{X} \end{bmatrix} & \in \mathcal{CPP}(\mathbb{R}^{2n+1}_+), \end{aligned}$$

in the sense that their optimal value is the same, and that for an optimal solution $(\mathbf{p}^*, \mathsf{Y}^*, \mathsf{X}^*)$ *there exists a sequence of non-negative random objective coefficient vectors* $\tilde{\mathbf{c}}^*_\varepsilon$ *and feasible solutions* $\mathbf{x}^*(\tilde{\mathbf{c}}^*_\varepsilon)$ *converging in moments to this optimal solution, i.e.,*

$$\lim_{\varepsilon\searrow 0}\mathbb{E}\left(\begin{bmatrix} 1 \\ \tilde{\mathbf{c}}_\varepsilon \\ \mathbf{x}(\tilde{\mathbf{c}}_\varepsilon)\end{bmatrix}\begin{bmatrix} 1 \\ \tilde{\mathbf{c}}_\varepsilon \\ \mathbf{x}(\tilde{\mathbf{c}}_\varepsilon)\end{bmatrix}^\top\right) = \begin{bmatrix} 1 & \boldsymbol{\mu}^\top & (\mathbf{p}^*)^\top \\ \boldsymbol{\mu} & \Sigma & (\mathsf{Y}^*)^\top \\ \mathbf{p}^* & \mathsf{Y}^* & \mathsf{X}^* \end{bmatrix}.$$

An interesting feature of this reformulation is that for binary variables x_j, $j\in\mathcal{B}$, the optimal solutions p^*_j have an interpretation as success probabilities (for $x_j(\tilde{\mathbf{c}}) = 1$):

$$p^*_j = \mathbb{E}\left[x_j(\tilde{\mathbf{c}})\right] = 1 * \mathbb{P}(x_j(\tilde{\mathbf{c}}) = 1) + 0 * \mathbb{P}(x_j(\tilde{\mathbf{c}}) = 0) = \mathbb{P}(x_j(\tilde{\mathbf{c}}) = 1)$$

under the limiting distribution of $\tilde{\mathbf{c}}$.

The authors further extend their approach to different cases where $(\boldsymbol{\mu}, \Sigma)$ are not exactly known but also to instances where the support of the objective function coefficients is all of $\mathbb{R}^n$, in which case the conic constraint merely needs to enforce membership in $\mathcal{CPP}(\mathbb{R}_+ \times \mathbb{R}^n \times \mathbb{R}^n_+)$. The later case is further investigated in the follow-up paper (Natarajan & Teo, 2017), where the authors apply Proposition 11.6., in order to derive models with reduced computational complexity and some applications. In (Padmanabhan, Natarajan, & Murthy, 2021) it is shown that in the above model one can exploit the structure of uncertainty in Σ in order to obtain reformulations with conic constraints of smaller dimension. For the sake of conciseness we will not present these models here. Beyond that, the approach has sparked some promising developments that involve interesting generalizations and applications of the approach described in this section. In (Kong, Lee, Teo, & Zheng, 2013) the authors investigated scheduling of arrivals to a stochastic service delivery system for which they modeled the second-stage problem in a fashion similar to (54), where additionally the uncertain objective $\tilde{\mathbf{c}}$ also depends on a first-stage decision. Their approach is also interesting since they derive a model similar to that in Theorem 24 with a different angle, based on moment decompositions. In (Kong, Li, Liu, Teo, & Yan, 2015) the authors investigated appointment scheduling under schedule-dependent patient no-show behaviour, where a similar model was employed that also incorporates uncertainty in the right-hand side $\mathbf{b}$. The discussion on the design of structures in operations in Yan, Gao, & Teo (2018) is an example where $Z(\tilde{\mathbf{c}})$ is interpreted as the dual of a linear problem with uncertain right-hand sides. In addition, they avoid the computational cost of introducing slack variables, which would increase the copositive matrix block, by replacing the respective constraints by a single bilinear constraint. Thus, their exposition also expands the applicability of the model in this regard.

8.1. Open problems

- **Geometrical analysis of Theorem 24**: In the original paper (Natarajan et al., 2011), the authors prove their main result with a methodology that has a striking resemblance with the proof strategy that was first used in Burer (2009) to prove Theorem 4 (for the special case $\mathcal{K} = \mathbb{R}^n_+$), where the added complication comes from the fact that the analysis has to proceed in the space of probability measures. In our discussion in Section 2.1.2 we presented an alternative strategy to the classical approach to Theorem 4 that rests on the geometrical analysis provided in Kim et al. (2020). However, it is not clear that a similar geometrical proof can be achieved for Theorem 24, since for Theorem 6 we assumed the vector space to be of finite dimension. It would be interesting to investigate whether a geometric approach similar to Kim et al. (2020) can be extended to the case of infinite-dimensional vector spaces in order to prove results such as Theorem 24.
- **General data uncertainty and problem structure:** An immediate question is whether the approach can be generalized to cases where not only the objective function coefficients, but also other parts of the problem data are uncertain. Also, in analogy to Theorem 4, it would be interesting to study generalizations of (54) where $\mathbf{x}\in\mathcal{K}$ for choices of $\mathcal{K}$ other than the positive orthant, and where binarity constraints are generalized to other types of quadratic constraints.

9. Two-stage distributionally robust optimization: conic formulation

In (Hanasusanto & Kuhn, 2018) the authors applied a distributionally robust framework to two-stage robust optimization. Analogously to the adjustable robust framework, the second-stage variables are allowed to adapt to the uncertainty, and the goal is to optimize expected performance under distributional ambiguity. The authors introduce copositive reformulations and relaxations for this model, and show that the resulting approximations outperform state-of-the-art approaches.

In the following exposition, ambiguity sets are modelled as a Wasserstein balls, centered around an empirical distribution:

$$\mathcal{B}^r_\varepsilon(\hat{P}_I) := \left\{P\in\mathcal{M}^r(\mathcal{U}) : W^r(P,\hat{P}_I) \le \varepsilon\right\}, \tag{55}$$

where, from now on, $\mathcal{U} := \left\{\mathbf{u}\in\mathbb{R}^k_+ : \mathsf{S}\mathbf{u}\le\mathbf{t}\right\}$ is a nonempty, polyhedral support, with $\mathsf{S}\in\mathbb{R}^{j\times k}$ and $\mathbf{t}\in\mathbb{R}^j$, and $\hat{P}_I$ is the empirical probability measure based as defined in Section 3.1.

9.1. Problem formulation

The model under consideration is

$$\min_{\mathbf{x}\in\mathcal{X}} \left[\mathbf{c}^\top\mathbf{x} + R(\mathbf{x})\right], \tag{56}$$

where $\mathcal{X} \subseteq \mathbb{R}^{n_1}$ is a feasible set not stricken with uncertainty, and $R(\mathbf{x})$ is the distributionally robust analog to the recourse function given by

$$R(\mathbf{x}) := \sup_{P\in\mathcal{P}} \mathbb{E}_P[Z(\mathbf{x},\mathbf{u})], \tag{57}$$

where $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^q$ is a random vector, and $\mathcal{P}$ is an ambiguity set of the possible distributions of $\mathbf{u}$. The recourse problem is given by

$$Z(\mathbf{x},\mathbf{u}) := \inf_{\mathbf{y}\in\mathbb{R}^{n_2}} \left\{(\mathsf{Q}\mathbf{u}+\mathbf{q})^\top\mathbf{y} : \mathsf{T}(\mathbf{x})\mathbf{u} + \mathbf{h}(\mathbf{x}) \le \mathsf{W}\mathbf{y}\right\}$$
$$\ge Z_d(\mathbf{x},\mathbf{u}) := \sup_{\mathbf{p}\in\mathbb{R}^m_+} \left\{[\mathsf{T}(\mathbf{x})\mathbf{u} + \mathbf{h}(\mathbf{x})]^\top\mathbf{p} : \mathsf{Q}\mathbf{u}+\mathbf{q} = \mathsf{W}^\top\mathbf{p}\right\}$$

where $\mathsf{T}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$ are appropriate matrix- and vector-valued functions and $Z_d(\mathbf{x},\mathbf{u})$ denotes the dual of $Z(\mathbf{x},\mathbf{u})$. The following two assumptions on the recourse problem have critical influence on the behaviour copositive bounds we are about to derive:

Assumption 3. Regarding the recourse problem we consider the following qualities:

(a) *Complete recourse:* There exists $\mathbf{y}^+ \in \mathbb{R}^{n_2}$ so that $\mathsf{W}\mathbf{y}^+ > \mathbf{o}$.
(b) *Suffiently expensive recourse:* For any $\mathbf{u} \in \mathcal{U}$ the dual problem Z_d is bounded.

Complete recourse ensures that $Z(\mathbf{x},\mathbf{u})$ is always finite while sufficiently expensive recourse ensures $Z(\mathbf{x},\mathbf{u})$ is always feasible. If the recourse problem exhibits either of these qualities then $Z(\mathbf{x},\mathbf{u}) = Z_d(\mathbf{x},\mathbf{u})$. Assumption 3b) will be maintained throughout the discussion, while different results will be presented no matter whether Assumption 3a) is satisfied or not.

9.2. Conic reformulation

The following theorem is at the heart of the derivation of the copositive bounds:

Theorem 25. *If $\mathcal{P} = \mathcal{B}^r_\varepsilon(\hat{P}_I)$ the worst-case expectation (57) coincides with the optimal value of the generalized moment problem*

$$R(\mathbf{x}) = \sup_{P_i\in\mathcal{M}^r(\mathcal{U})} \left\{\tfrac{1}{I}\sum_{i=1}^I \int_{\mathcal{U}} Z(\mathbf{x},\mathbf{u})P_i(\mathrm{d}\mathbf{u}) : \tfrac{1}{I}\sum_{i=1}^I \int_{\mathcal{U}} [d(\mathbf{u},\hat{\mathbf{u}}_i)]^r P_i(\mathrm{d}\mathbf{u}) \le \varepsilon^r\right\}. \tag{58}$$

Furthermore, for $\varepsilon > 0$ this problem admits the strong dual robust optimization problem

$$R(\mathbf{x}) = \inf_{\lambda\in\mathbb{R}_+} \left[\varepsilon^r\lambda + \tfrac{1}{I}\sum_{i=1}^I \sup_{\mathbf{u}\in\mathcal{U}} Z(\mathbf{x},\mathbf{u}) - \lambda\, d(\mathbf{u},\hat{\mathbf{u}})^r\right]. \tag{59}$$

The first formulation can be related to a completely positive optimization problem, the second one to a copositive optimization problem. These conic problems can be shown to be duals of each other. Assumption 3a) can then be used in order to close the duality gap, so that one achieves a conic reformulation of $R(\mathbf{x})$ that enjoys strong duality. The conic optimization problems in question are:

$$\bar{R}(\mathbf{x}) := \inf_{\lambda\in\mathbb{R}_+, s_i\in\mathbb{R}, \boldsymbol{\psi}_i,\boldsymbol{\phi}_i\in\mathbb{R}^{n_2+j}} \left\{\varepsilon^2 + \tfrac{1}{I}\sum_{i=1}^I \left[s_i + \bar{\mathbf{q}}^\top\boldsymbol{\psi}_i - \lambda\|\hat{\mathbf{u}}_i\|^2 + \sum_{k=1}^{n_2+j}\phi_i(\bar{\mathbf{q}})_k\right]\right\}$$

$$\text{s.t.}: \begin{bmatrix} \lambda\mathsf{I} + \bar{\mathsf{Q}}^\top\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{Q}} & -\tfrac{1}{2}\bar{\mathsf{T}}(\mathbf{x})^\top - \bar{\mathsf{Q}}^\top\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{W}}^\top & -\lambda\hat{\mathbf{u}}_i - \tfrac{1}{2}\bar{\mathsf{Q}}^\top\boldsymbol{\psi}_i \\ -\tfrac{1}{2}\bar{\mathsf{T}}(\mathbf{x}) - \bar{\mathsf{W}}\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{Q}} & \bar{\mathsf{W}}\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{W}}^\top & \tfrac{1}{2}\left[\bar{\mathsf{W}}\boldsymbol{\psi}_i - \bar{\mathbf{h}}(\mathbf{x})\right] \\ \left[-\lambda\hat{\mathbf{u}}_i - \tfrac{1}{2}\bar{\mathsf{Q}}^\top\boldsymbol{\psi}_i\right]^\top & \tfrac{1}{2}\left[\bar{\mathsf{W}}\boldsymbol{\psi}_i - \bar{\mathbf{h}}(\mathbf{x})\right]^\top & s_i \end{bmatrix} \in \mathcal{COP}(\mathbb{R}^{k+m+j+1}_+),$$

$$i \in [1\!:\!I],$$

where

$$\bar{\mathsf{Q}} := \begin{bmatrix}\mathsf{Q}\\ \mathsf{S}\end{bmatrix},\ \bar{\mathbf{q}} := \begin{bmatrix}\mathbf{q}\\ -\mathbf{t}\end{bmatrix},\ \bar{\mathsf{T}}(\mathbf{x}) := \begin{bmatrix}\mathsf{T}(\mathbf{x})\\ \mathsf{O}\end{bmatrix},$$
$$\bar{\mathbf{h}}(\mathbf{x}) := \begin{bmatrix}\mathbf{h}(\mathbf{x})\\ \mathbf{o}\end{bmatrix},\ \bar{\mathsf{W}} := \begin{bmatrix}\mathsf{W} & \mathsf{O}\\ \mathsf{O} & \mathsf{I}\end{bmatrix},$$

while

$$\underline{R}(\mathbf{x}) := \sup \tfrac{1}{I}\sum_{i=1}^I \left[\mathrm{Tr}(\bar{\mathsf{T}}(\mathbf{x})\mathsf{Y}_i) + \bar{\mathbf{h}}(\mathbf{x})^\top\boldsymbol{\gamma}_i\right]$$
$$\text{s.t.}: \bar{\mathsf{Q}}\boldsymbol{\mu}_i + \bar{\mathbf{q}} = \bar{\mathsf{W}}^\top\boldsymbol{\gamma}_i, \quad i \in [1\!:\!I],$$
$$\mathrm{diag}\left(\begin{bmatrix}\bar{\mathsf{Q}}\\ -\bar{\mathsf{W}}\end{bmatrix}\begin{bmatrix}\Omega_i & \mathsf{Y}_i\\ \mathsf{Y}_i^\top & \Gamma_i\end{bmatrix}\begin{bmatrix}\bar{\mathsf{Q}}\\ -\bar{\mathsf{W}}\end{bmatrix}^\top\right) = \bar{\mathbf{q}}\circ\bar{\mathbf{q}}, \quad i \in [1\!:\!I],$$
$$\tfrac{1}{I}\sum_{i=1}^I \left[\mathrm{Tr}(\Omega_i) - 2\hat{\mathbf{u}}_i^\top\boldsymbol{\mu}_i + \hat{\mathbf{u}}_i^\top\hat{\mathbf{u}}_i\right] \le \varepsilon^2,$$
$$\begin{bmatrix}\Omega_i & \mathsf{Y}_i & \boldsymbol{\mu}_i\\ \mathsf{Y}_i^\top & \Gamma_i & \boldsymbol{\gamma}_i\\ \boldsymbol{\mu}_i^\top & \boldsymbol{\gamma}_i^\top & 1\end{bmatrix} \in \mathcal{CPP}\left(\mathbb{R}^{k+m+j}\right), \quad i \in [1\!:\!I].$$

In a nutshell, the argument proceeds as follows:

- Show that $R(\mathbf{x}) \le \bar{R}(\mathbf{x})$ by replacing $Z(\mathbf{x},\mathbf{u})$ by $Z_d(\mathbf{x},\mathbf{u})$ so that the inner supremum problems become quadratic optimization problems that, after squaring their linear constraints, are upper bounded by their completely positive relaxations. Combining the suprema we obtain a conic problem that is itself upper bounded by its dual that is in fact given by $\bar{R}(\mathbf{x})$.
- Show that $R(\mathbf{x}) \ge \underline{R}(\mathbf{x})$ by an argument resembling the proof of Theorem 4, with the added difficulty that the decomposition of feasible solutions of $\underline{R}(\mathbf{x})$ have to be translated into approriately constructed, discrete probability measures.
- The gaps can then be closed by showing that strong duality holds between the two conic optimization problems. While weak duality is immediate from the standard derivation of conic duals, closing the gap involves a generalization of the sufficiency part of the Shur-complementation criterion for positive definiteness to copositive matrices. For the so obtained relaxation, a Slater point can be constructed under Assumption 3a), which by sufficiency yields a Slater point for $\bar{R}(\mathbf{x})$.

Hence, a conic reformulation of the recourse problem that enjoys strong duality is obtainable given that the problem has complete recourse. Replacing $R(\mathbf{x})$ by $\bar{R}(\mathbf{x})$ in the description of (56) yields a finite, conic reformulation that can be approximated with standard techniques.

In case the latter assumption fails, the authors prove approximation results that use a slight modification of the conic problems. Define $\bar{R}_\delta(\mathbf{x})$ the same way as $\bar{R}(\mathbf{x})$ with $\bar{\mathsf{W}}\,\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{W}}^\top$ replaced by $\bar{\mathsf{W}}\,\mathrm{Diag}(\boldsymbol{\phi}_i)\bar{\mathsf{W}}^\top + \delta\mathsf{I}$, $i \in [1\!:\!I]$, with $\delta > 0$ a constant, and consider

$$\min_{\mathbf{x}\in\mathcal{X}} \left[\mathbf{c}^\mathsf{T}\mathbf{x} + \bar{R}_\delta(\mathbf{x})\right]. \tag{60}$$

Again we have a finite, conic optimization problem for which the following theorem can be proved.

Theorem 26. *The following statements hold:*

1. *If $\delta = 0$ and (56) has complete recourse, then it is equivalent to (60).*
2. *If $\delta = 0$ and (56) fails to have complete recourse, then it is upper bounded by (60).*
3. *If $\delta > 0$, then (56) is lower bounded by (60).*
4. *If $\mathcal{X}$ is compact, then the optimal value of (60) converges to that of (56) as $\delta \searrow 0$. Moreover, every cluster point $\mathbf{x}^*$ of a sequence $\left\{\mathbf{x}^*_\delta\right\}_{\delta\searrow 0}$ of minimizers to (60) is a minimizer to (56).*

9.3. Open problems

Multi-stage distributionally robust optimization: Results in two-stage robust optimization are often not that easy to generalize to the multi-stage setting. It would be an interesting challenge to investigate whether the approach above is useful for tackling the multi-stage distributionally robust optimization problem under Wasserstein ambiguity.

The case of uncertain recourse: Similar to the discussion in Section 5.3, the case of uncertain recourse is an open problem for the approach. The reason is basically the same: quadratic terms in the constraints of the recourse problem limit the application of copositive techniques. Hence, more research is needed in this regard.

Unified analysis of the two-stage setting: There is a strong relation between some of the approaches we have discussed so far. In fact the authors of both (Hanasusanto & Kuhn, 2018; Xu & Burer, 2018) point out that the exactness of (43) from Section 5 is equivalent to Theorem 26. In addition, the recourse function $R(\mathbf{x})$ bears striking resemblance to (55) discussed in Section 8, once the infimum problem $Z(\mathbf{x},\mathbf{u})$ is replaced replaced by the supremum problem $Z_d(\mathbf{x},\mathbf{u})$, and both models are eventually reformulated into a copositve optimization problem. However, all three approaches use very different methods to derive their results and even more, results similar to those in Natarajan et al. (2011) are derived in Kong et al. (2013) with a different proof strategy. It is therefore plausible that there is a unifying lens under which all these approaches can be understood, perhaps based on the geometrical analysis hinted at in Section 8.1, but so far such an approach is absent from the literature.

10. Conclusions

Copositive optimization tools for quadratic optimization are an active area of research that yields a plethora of possible applications in optimization under risk and uncertainty. Research in recent history has shown that the interplay between these two fields sparks powerful approaches that are competitive especially when it comes to modelling and improving bounds on difficult problems. Still many open questions and potential new research areas remain, some of which we pointed out in our discussion. At this point we also like to point out a major area where, hopefully, improvements can be made in the future and that is the question of how to make copositive optimization more practical. The major drawback of the this technology remains the substantial computational burden that come along with solving even simple approximations. We believe there are many untapped sources for improvement on that front. We hope that this text is able to encourage readers to engage in these challenging questions and to expand on these ideas in future research.

Appendix A. Longer proofs

Proof of Theorem 4. We will prove only a weaker version of the theorem, where we assume $\mathcal{L}$ to be bounded. As in Proposition 5 we denote by $\mathsf{M} = [\mathbf{b}, -\mathsf{A}]$ and define

$$\mathbb{K} := \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K})),$$
$$\mathbb{J} := \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[0\colon m]\right\},$$

and let $\widehat{\mathsf{Q}}_i$ be defined as in Section 2.1.2. We can rewrite the QCQP as

$$\begin{aligned}&\min_{\mathsf{Y}}\left\{\bar{\mathsf{Q}}_0\bullet\mathsf{Y} : \mathsf{H}_0\bullet\mathsf{Y} = 1,\ \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[0\colon m],\right.\\ &\qquad\left.\mathsf{Y}\in\operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}))\right\}\\ &=\min_{\mathsf{Y}}\left\{\langle\bar{\mathsf{Q}}_0,\mathsf{Y}\rangle : \mathsf{Y}\in\mathbb{K}\cap\mathbb{J},\ \langle\mathsf{H}_0,\mathsf{Y}\rangle = 1\right\}\end{aligned}$$

while the convex reformulation can be rewritten as

$$\begin{aligned}&\min_{\mathsf{Y}}\left\{\bar{\mathsf{Q}}_0\bullet\mathsf{Y} : \mathsf{H}_0\bullet\mathsf{Y} = 1,\ \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[0\colon m],\right.\\ &\qquad\left.\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K})\right\}\\ &=\min_{\mathsf{Y}}\left\{\langle\bar{\mathsf{Q}}_0,\mathsf{Y}\rangle : \mathsf{Y}\in\mathbb{J},\ \langle\mathsf{H}_0,\mathsf{Y}\rangle = 1\right\}\end{aligned}$$

If we want to use Theorem 6 to show that the two problems are equivalent we need to show that $\mathbb{H}\cap\mathbb{K}\neq\emptyset$ is bounded and that $\mathbb{J}$ is a face of $\operatorname{conv}(\mathbb{K})$. To show that boundedness consider that

$$\begin{aligned}\mathbb{H}\cap\mathbb{J} &:= \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[0\colon m],\ \mathsf{H}_0\bullet\mathsf{Y} = 1\right\}\\ &\subseteq\left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \widehat{\mathsf{Q}}_0\bullet\mathsf{Y} = 0,\ \mathsf{H}_0\bullet\mathsf{Y} = 1\right\} =: \mathbb{J}_0.\end{aligned}$$

We will prove that $\mathbb{J}_0$ is bounded. The recession cone of $\mathbb{J}_1$ is given by

$$0^+\mathbb{J}_0 = \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \widehat{\mathsf{Q}}_0\bullet\mathsf{Y} = 0,\ \mathsf{H}_0\bullet\mathsf{Y} = 0\right\}.$$

So assume $\mathsf{O}\neq\mathsf{Y}\in 0^+\mathbb{J}_0$. We have

$$\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) \Rightarrow \mathsf{Y} = \sum_{i=1}^k \mathbf{y}_i\mathbf{y}_i^\mathsf{T} \text{ with } \mathbf{y}_i\in\mathbb{R}_+\times\mathcal{K}$$

$$\widehat{\mathsf{Q}}_0\bullet\mathsf{Y} = \sum_{i=1}^k \mathbf{y}_i^\mathsf{T}\mathsf{M}^\mathsf{T}\mathsf{M}\mathbf{y}_i = 0 \Rightarrow \mathsf{M}\mathbf{y}_i = \mathbf{o}$$

$$\mathsf{H}_0\bullet\mathsf{Y} = \sum_{i=1}^k (\mathbf{y}_i)^2_{n+1} = 0 \Rightarrow (\mathbf{y}_i)_{n+1} = 0$$

So $\mathbf{y}_i = \left[\ \mathbf{x}_i\ \mathbf{o}\ \right]$ for some $\mathbf{x}_i\in\mathcal{K}$ with $\mathsf{A}\mathbf{x}_i = \mathbf{o}$, but then $\mathcal{L} := \{\mathbf{x}\in\mathcal{K} : \mathsf{A}\mathbf{x} = \mathbf{b}\}$ is not bounded contrary to our assumption. Hence $0^+\mathbb{J}_1$ contains only the origin so that $\mathbb{J}_1$ is bounded.

Let's unpack

$$\begin{aligned}\mathbb{J} &= \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[0\colon m]\right\}\\ &= \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \mathsf{M}^\mathsf{T}\mathsf{M}\bullet\mathsf{Y} = 0, \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[1\colon m]\right\},\end{aligned}$$

and define $\mathbb{J}_{-1} = \operatorname{conv}(\mathbb{K}) = \mathcal{CPP}(\mathbb{R}_+\times\mathcal{K})$ and

$$\mathbb{J}_p := \left\{\mathsf{Y}\in\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K}) : \mathsf{M}^\mathsf{T}\mathsf{M}\bullet\mathsf{Y} = 0, \widehat{\mathsf{Q}}_i\bullet\mathsf{Y} = 0,\ i\in[1\colon p]\right\}.$$

First, note that $\mathsf{M}^\mathsf{T}\mathsf{M}\in\mathcal{S}^{n+1}_+$ so $\mathsf{M}^\mathsf{T}\mathsf{M}\bullet\mathbf{x}\mathbf{x}^\mathsf{T}\geq 0$ for all $\mathbf{x}\in\mathcal{K}$, which implies that $\mathsf{M}^\mathsf{T}\mathsf{M}\bullet\mathsf{X}\geq 0$ for all $\mathsf{X}\in\operatorname{conv}(\mathbb{K})$ so that $\mathsf{M}^\mathsf{T}\mathsf{M}\in\operatorname{conv}(\mathbb{K})^*$. As a side product we get that $\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}_1) = \mathbb{H}\cap\mathbb{J}_1$ by Theorem 6 since $\mathbb{J}_1$ is a face of $\mathbb{K}$ by Theorem 7. Thus we have

$$\begin{aligned}\operatorname{conv}(\mathbb{H}\cap\mathbb{K}\cap\mathbb{J}_1) &= \operatorname{conv}\left\{\mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y}\in\mathbb{R}_+\times\mathcal{K} : \mathbf{y}^\mathsf{T}\mathsf{M}^\mathsf{T}\mathsf{M}\mathbf{y} = 0,\ \mathbf{y}^2_{n+1} = 1\right\}\\ &= \operatorname{conv}\left\{\mathsf{Y}\in\operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+\times\mathcal{K})) : \widehat{\mathsf{Q}}_0\bullet\mathsf{Y} = 0,\ \mathsf{Y}_{n+1} = 1\right\}\end{aligned}$$

$$= \left\{ \mathsf{Y} \in \mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K}) : \widehat{\mathsf{Q}}_0 \bullet \mathsf{Y} = 0,\ \mathsf{Y}_{n+1} = 1 \right\} = \mathbb{H} \cap \mathbb{J}_1$$

Now by the key assumption for any $i \in [1\!:\!m]$ we have

$$\widehat{\mathsf{Q}}_i \bullet \mathbf{y}\mathbf{y}^\mathsf{T} \geq 0 \quad \text{for all } \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K} : \mathsf{M}\mathbf{y} = \mathbf{o},\ \mathbf{y}_{n+1} = 1,$$

so that

$$\begin{aligned} \widehat{\mathsf{Q}}_i \bullet \mathsf{Y} \geq 0 \quad \text{for all} \quad \mathsf{Y} &\in \operatorname{conv}\left\{ \mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K} : \mathsf{M}\mathbf{y} = \mathbf{o},\ \mathbf{y}_{n+1} = 1 \right\} \\ &= \operatorname{conv}\left\{ \mathbf{y}\mathbf{y}^\mathsf{T} : \mathbf{y} \in \mathbb{R}_+ \times \mathcal{K} : \mathbf{y}^\mathsf{T}\mathsf{M}^\mathsf{T}\mathsf{M}\mathbf{y} = 0,\ \mathbf{y}_{n+1}^2 = 1 \right\} \\ &= \operatorname{conv}\left\{ \mathsf{Y} \in \operatorname{ext}(\mathcal{CPP}(\mathbb{R}_+ \times \mathcal{K})) : \widehat{\mathsf{Q}}_0 \bullet \mathsf{Y} = 0,\ \mathsf{Y}_{n+1} = 1 \right\} \\ &= \operatorname{conv}(\mathbb{H} \cap \mathbb{K} \cap \mathbb{J}_1) = \mathbb{H} \cap \mathbb{J}_1 \end{aligned}$$

Thus $\widehat{\mathsf{Q}}_i \in \mathbb{J}_1^*$ for any $i \in [1\!:\!m]$. But $\mathbb{J}_p \subseteq \mathbb{J}_1$ for any $p \in [2\!:\!m]$ which implies $\mathbb{J}_1^* \subseteq \mathbb{J}_p^*$ for any $p \in [2\!:\!m]$. Consequently, $\widehat{\mathsf{Q}}_i \in \mathbb{J}_p^*$ for any $i \in [1\!:\!m],\ p \in [2\!:\!m]$. We have argued that $\widehat{\mathsf{Q}}_0 := \mathsf{M}^\mathsf{T}\mathsf{M} \in \operatorname{conv}(\mathbb{K})^*$ and that $\widehat{\mathsf{Q}}_p \in \mathbb{J}_{p-1}^*$. So $\mathbb{J}$ is a face of $\operatorname{conv}(\mathbb{K})$. □

Proof of Theorem 8. An elementary argument shows that $\mathcal{G}(\cup_{i=1}^k \mathcal{F}_i) = \operatorname{conv} \cup_{i=1}^k \mathcal{G}(\mathcal{F}_i)$. Clearly, for a convex combination $\mathsf{X} = \sum_{i=1}^k \lambda_i \mathsf{X}_i$ with $\mathsf{X}_i \in \mathcal{G}(\mathcal{F}_i),\ i \in [1\!:\!k]$ we have $\hat{\mathsf{X}}_i := \lambda_i \mathsf{X}_i \in \mathcal{C}_i$, with $\mathscr{A}(\hat{\mathsf{X}}_i) = \mathbf{o}$ so that $\mathsf{H} \bullet (\sum_{i=1}^k \hat{\mathsf{X}}_i) = \sum_{i=1}^k \lambda_i \mathsf{H} \bullet \mathsf{X}_i = \sum_{i=1}^k \lambda_i = 1$. Conversely, for an $\mathsf{X} := \sum_{i=1}^k \mathsf{X}_i$ with $\mathsf{X}_i \in \mathcal{C}_i$, with $\mathscr{A}(\hat{\mathsf{X}}_i) = \mathbf{o}$ and $\mathsf{H} \bullet \mathsf{X} = 1$, we can write $\mathsf{X} = \sum_{i=1}^k (\mathsf{H} \bullet \mathsf{X}_i) \frac{\mathsf{X}_i}{\mathsf{H} \bullet \mathsf{X}_i}$ since $\mathsf{H} \bullet \mathsf{X}_i > 0$ by assumption and we have $\frac{\mathsf{X}_i}{\mathsf{H} \bullet \mathsf{X}_i} \in \mathcal{C}_i$, $\mathscr{A}(\frac{\mathsf{X}_i}{\mathsf{H} \bullet \mathsf{X}_i}) = \mathbf{o}$ and $\sum_{i=1}^k (\mathsf{H} \bullet \mathsf{X}_i) = \mathsf{H} \bullet \mathsf{X} = 1$ as desired. □

Proof of Proposition 11. We start observing extremality of dyadic matrices: suppose $\mathbf{x}\mathbf{x}^\mathsf{T} = \mathsf{A} + \mathsf{B}$ with $\{\mathsf{A}, \mathsf{B}\} \subset \mathcal{S}_+^n$. Then

$$(\mathbf{x}^\mathsf{T}\mathbf{u})^2 = \mathbf{u}^\mathsf{T}\mathsf{A}\mathbf{u} + \mathbf{u}^\mathsf{T}\mathsf{B}\mathbf{u} \quad \text{for all } \mathbf{u} \in \mathbb{R}^n$$

which implies $\mathbf{x}^\perp \subseteq \ker \mathsf{A} \cap \ker \mathsf{B}$. If $\mathbf{x} = \mathbf{o}$, this already yields $\mathsf{A} = \mathsf{B} = \mathsf{O}$. If $\mathbf{x} \neq \mathbf{o}$, spectral decomposition of A and B yields by above $\mathsf{A} = \alpha\mathbf{x}\mathbf{x}^\mathsf{T}$ and $\mathsf{B} = \beta\mathbf{x}\mathbf{x}^\mathsf{T}$ for some $\{\alpha, \beta\} \subset \mathbb{R}_+$ with $\alpha + \beta = 1$, which shows extremality of $\mathbf{x}\mathbf{x}^\mathsf{T}$ in $\mathcal{S}_+^n$ and also in all $\mathcal{CPP}(\mathcal{K})$ for any cone $\mathcal{K}$. Next we observe uniqueness up to reflection for vectors building dyadic matrices:

$$\mathbf{x}\mathbf{x}^\mathsf{T} = \mathbf{y}\mathbf{y}^\mathsf{T} \quad \Longrightarrow \quad \mathbf{x} \in \{-\mathbf{y}, \mathbf{y}\}\,. \tag{61}$$

Indeed, considering

$$\|\mathbf{x}\|^4 = \mathbf{x}^\mathsf{T}(\mathbf{x}\mathbf{x}^\mathsf{T})\mathbf{x} = \mathbf{x}^\mathsf{T}(\mathbf{y}\mathbf{y}^\mathsf{T})\mathbf{x} = (\mathbf{x}^\mathsf{T}\mathbf{y})^2 = \mathbf{y}^\mathsf{T}(\mathbf{x}\mathbf{x}^\mathsf{T})\mathbf{y} = \|\mathbf{y}\|^4$$

we have the equality case of the Cauchy–Schwarz inequality which yields either $\mathbf{y} = \mathbf{o}$ or else $\mathbf{x} = \alpha\mathbf{y}$ with $\alpha^4 = 1$, so in any case $\mathbf{x} = -\mathbf{y}$ or $\mathbf{x} = \mathbf{y}$.

Now let us consider each point individually:

1. $\mathcal{COP}(\mathcal{K}) = \mathcal{COP}(-\mathcal{K}) = \mathcal{COP}(\mathcal{K} \cup -\mathcal{K})$, which also holds if $\mathcal{COP}$ is replaced with $\mathcal{CPP}$.
   We can simply appeal to $\mathbf{y}^\mathsf{T}\mathsf{X}\mathbf{y} = (-\mathbf{y})^\mathsf{T}\mathsf{X}(-\mathbf{y})$ and $\mathbf{y}\mathbf{y}^\mathsf{T} = (-\mathbf{y})(-\mathbf{y})^\mathsf{T}$.
2. If $\mathcal{K}_1 \subseteq \mathcal{K}_2$, then $\mathcal{CPP}(\mathcal{K}_1) \subseteq \mathcal{CPP}(\mathcal{K}_2)$ with equality if and only if $\mathcal{K}_2 \subseteq \mathcal{K}_1 \cup -\mathcal{K}_1$.
   The inclusions are obvious, as is sufficiency for the equalities, using 1. Now assume $\mathbf{x} \in \mathcal{K}_2$ and identity of the $\mathcal{CPP}$ cones. Then, by extremality of dyadic matrices $\mathbf{x}\mathbf{x}^\mathsf{T} = \mathbf{y}\mathbf{y}^\mathsf{T}$ for some $\mathbf{y} \in \mathcal{K}_1$, so $\mathbf{x} = \pm\mathbf{y} \in \mathcal{K}_1 \cup -\mathcal{K}_1$ by (61).
3. If $\mathcal{K}_1 \subseteq \mathcal{K}_2$, then $\mathcal{COP}(\mathcal{K}_1) \supseteq \mathcal{COP}(\mathcal{K}_2)$; if in addition we assume $\operatorname{int} \mathcal{K}_2^* \neq \emptyset$, we have $\mathcal{COP}(\mathcal{K}_1) = \mathcal{COP}(\mathcal{K}_2)$ if and only if[2] $\mathcal{K}_2 \subseteq \operatorname{cl} \mathcal{K}_1$.
   Again, the inclusion statements of the $\mathcal{COP}$ cones is obvious. Assume now they are identical and suppose, arguing by contradiction, the existence of an $\mathbf{x} \in \mathcal{K}_2 \setminus \operatorname{cl}\mathcal{K}_1$. Then $\mathbf{x} \neq \mathbf{o}$ and moreover, there is a $\mathbf{c} \in \operatorname{int}\mathcal{K}_2^* \setminus \{\mathbf{o}\}$ such that $\mathbf{c}^\mathsf{T}\mathbf{x} > 0$. Since $\mathbf{c} \in \operatorname{int}\mathcal{K}_2^* \subseteq \operatorname{int}\mathcal{K}_1^*$, the set

   $$B := \{\mathbf{z} \in \operatorname{cl}\mathcal{K}_1 : \mathbf{c}^\mathsf{T}\mathbf{z} = \mathbf{c}^\mathsf{T}\mathbf{x}\}$$

   is a compact base of $\operatorname{cl}\mathcal{K}_1$ and $\rho := \operatorname{dist}(\mathbf{x}, \operatorname{cl}\mathcal{K}_1) > 0$. It follows that the quadratic form

   $$\mathbf{y}^\mathsf{T}\mathsf{Q}\mathbf{y} = q(\mathbf{y}) := \|\mathbf{y} - \frac{\mathbf{c}^\mathsf{T}\mathbf{y}}{\mathbf{c}^\mathsf{T}\mathbf{x}}\mathbf{x}\|^2 - \rho^2\left(\frac{\mathbf{c}^\mathsf{T}\mathbf{y}}{\mathbf{c}^\mathsf{T}\mathbf{x}}\right)^2$$

   satisfies $q(\mathbf{x}) = -\rho^2 < 0$ while on the other hand, we have for any $\mathbf{y} \in \operatorname{cl}\mathcal{K}_1 \setminus \{\mathbf{o}\}$ that $\bar{\mathbf{y}} := \frac{\mathbf{c}^\mathsf{T}\mathbf{y}}{\mathbf{c}^\mathsf{T}\mathbf{x}}\mathbf{y} \in B$, so

   $$q(\mathbf{y}) = \left(\frac{\mathbf{c}^\mathsf{T}\mathbf{y}}{\mathbf{c}^\mathsf{T}\mathbf{x}}\right)^2 q(\bar{\mathbf{y}}) \geq 0\,,$$

   since $q(\bar{\mathbf{y}}) = \|\bar{\mathbf{y}} - \mathbf{x}\|^2 - \rho^2 \geq 0$ by definition of ρ. In particular $\mathsf{Q} \in \mathcal{COP}(\mathcal{K}_1) = \mathcal{COP}(\mathcal{K}_2)$ which is absurd in view of the relations $\mathbf{x} \in \mathcal{K}_2$ and $q(\mathbf{x}) < 0$. Hence $\mathcal{K}_2 \subseteq \operatorname{cl}\mathcal{K}_1$. This inclusion implies, conversely, the already observed $\mathcal{COP}$ inclusions (for the leftmost identity see 12. below)

   $$\mathcal{COP}(\mathcal{K}_1) = \mathcal{COP}(\operatorname{cl}\mathcal{K}_1) \subseteq \mathcal{COP}(\mathcal{K}_2) \subseteq \mathcal{COP}(\mathcal{K}_1)\,.$$

   Note that for the $\mathcal{CPP}$ cones, we have a different implication: if $\mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \operatorname{cl}(\mathcal{K}_1 \cup -\mathcal{K}_1)$ then not necessarily $\mathcal{CPP}(\mathcal{K}_1) = \mathcal{CPP}(\mathcal{K}_2)$ but $\operatorname{cl}\mathcal{CPP}(\mathcal{K}_1) = \mathcal{CPP}(\mathcal{K}_2)$, cf., again, 12. below. For the footnote [0] in this point, note that, by virtue of $\mathcal{K}_1 \subseteq \mathcal{K}_2$, all $\mathbf{z} \in \mathcal{K}_2 \setminus \operatorname{cl}\mathcal{K}_1 \subseteq \mathcal{K}_2 \cap (-\operatorname{cl}\mathcal{K}_1) \subseteq \mathcal{K}_2 \cap (-\operatorname{cl}\mathcal{K}_2)$ satisfy $\mathbf{c}^\mathsf{T}\mathbf{z} = 0$ for any $\mathbf{c} \in \operatorname{int}[\mathcal{K}_2]^*$, implying $\mathbf{z} = \mathbf{o} \in \operatorname{cl}\mathcal{K}_1$, which is absurd. Thus the assumptions yield $\mathcal{K}_2 \subseteq \operatorname{cl}\mathcal{K}_1$.
4. $\mathcal{CPP}(\mathcal{K}) \subseteq \mathcal{S}_+^n \subseteq \mathcal{COP}(\mathcal{K})$; all three sets are equal if and only if $\mathcal{K} \cup -\mathcal{K} = \mathbb{R}^n$.
   We obviously have $\mathcal{CPP}(\mathbb{R}^n) = \mathcal{S}_+^n = \mathcal{COP}(\mathbb{R}^n)$. Now specialize $\mathcal{K}_1 = \mathcal{K} \cup -\mathcal{K}$ and $\mathcal{K}_2 = \mathbb{R}^n$ in 2. and 3., to arrive at the claim, using 1.
5. $\mathcal{COP}(\mathbb{R}_+ \times \mathbb{R}^m) = \mathcal{CPP}(\mathbb{R}_+ \times \mathbb{R}^m) = \mathcal{S}_+^{m+1}$.
   The statement follows from 4. since $\mathcal{K} = \mathbb{R}_+ \times \mathbb{R}^m$ satisfies $\mathcal{K} \cup -\mathcal{K} = \mathbb{R}^{m+1}$.
6. $\mathcal{CPP}(\mathcal{K} \times \mathbb{R}^m) = \left\{ \begin{pmatrix} \mathsf{M}_{11} & \mathsf{M}_{21}^\mathsf{T} \\ \mathsf{M}_{21} & \mathsf{M}_{22} \end{pmatrix} \in \mathcal{S}_+^{m+n} : \mathsf{M}_{11} \in \mathcal{CPP}(\mathcal{K}) \right\}$ if $\mathbf{o} \in \mathcal{K}$.
   The statement was proved first in Dickinson (2013) and then independently in Natarajan & Teo (2017), both for the case that $\mathcal{K}$ is closed. We present an alternative proof that merely requires $\mathbf{o} \in \mathcal{K}$. Sufficiency is clear, since any set-completely positive matrix cone is a subset of the positive semidefinite matrix cone and the north-west block of any matrix in $\mathcal{CPP}(\mathcal{K} \times \mathbb{R}^m)$ is in $\mathcal{CPP}(\mathcal{K})$. So consider an element M of the right-hand set. Since it is a psd matrix we have a decomposition

   $$\mathsf{M} = \begin{bmatrix} \mathsf{M}_{11} & \mathsf{M}_{21}^\mathsf{T} \\ \mathsf{M}_{21} & \mathsf{M}_{22} \end{bmatrix} = \begin{bmatrix} \mathsf{X} \\ \mathsf{Y} \end{bmatrix}\begin{bmatrix} \mathsf{X} \\ \mathsf{Y} \end{bmatrix}^\mathsf{T} \text{ with } \mathsf{X} \in \mathbb{R}^{n \times r},\ \mathsf{Y} \in \mathbb{R}^{m \times r},$$

   so that $\mathsf{M}_{11} = \mathsf{X}\mathsf{X}^\mathsf{T}$. But we also have $\mathsf{M}_{11} \in \mathcal{CPP}(\mathcal{K})$, so that in fact $\mathsf{M}_{11} = \mathsf{Z}\mathsf{Z}^\mathsf{T}$ for some $n \times s$ matrix Z with $\mathbf{z}_i \in \mathcal{K}$ for all $i \in [1\!:\!s]$. If $s \neq r$ we can always append columns of zeroes to the smaller matrix without changing the relation $\mathsf{X}\mathsf{X}^\mathsf{T} = \mathsf{Z}\mathsf{Z}^\mathsf{T}$, so that w.l.o.g. we can assume $s = r$. From Groetzner & Dür (2020, Lem. 2.6) we then have that $\mathsf{X}\mathsf{X}^\mathsf{T} = \mathsf{Z}\mathsf{Z}^\mathsf{T}$ is the case exactly if $\mathsf{Z} = \mathsf{X}\mathsf{Q}$ for some $\mathsf{Q} \in \mathbb{R}^{r \times r}$ with $\mathsf{Q}\mathsf{Q}^\mathsf{T} = \mathsf{I}$. Set $\bar{\mathsf{Y}} = \mathsf{Y}\mathsf{Q}$. Then the decomposition

   $$\begin{bmatrix} \mathsf{Z} \\ \bar{\mathsf{Y}} \end{bmatrix}\begin{bmatrix} \mathsf{Z} \\ \bar{\mathsf{Y}} \end{bmatrix}^\mathsf{T} = \begin{bmatrix} \mathsf{X}\mathsf{Q} \\ \mathsf{Y}\mathsf{Q} \end{bmatrix}\begin{bmatrix} \mathsf{X}\mathsf{Q} \\ \mathsf{Y}\mathsf{Q} \end{bmatrix}^\mathsf{T} = \begin{bmatrix} \mathsf{X}\mathsf{Q}\mathsf{Q}^\mathsf{T}\mathsf{X}^\mathsf{T} & \mathsf{X}\mathsf{Q}\mathsf{Q}^\mathsf{T}\mathsf{Y}^\mathsf{T} \\ \mathsf{Y}\mathsf{Q}\mathsf{Q}^\mathsf{T}\mathsf{X}^\mathsf{T} & \mathsf{Y}\mathsf{Q}\mathsf{Q}^\mathsf{T}\mathsf{Y}^\mathsf{T} \end{bmatrix}$$

[2] note that $\mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \operatorname{cl}\,(\mathcal{K}_1 \cup -\mathcal{K}_1)$ and $\operatorname{int}\,\mathcal{K}_2^* \neq \emptyset$ already implies $\mathcal{K}_2 \subseteq \operatorname{cl}\,\mathcal{K}_1$, so that this criterion coincides with the criterion of 2. up to closure

$$= \begin{bmatrix} \mathsf{X} \\ \mathsf{Y} \end{bmatrix} \begin{bmatrix} \mathsf{X} \\ \mathsf{Y} \end{bmatrix}^\mathsf{T} = \mathsf{M} \tag{62}$$

has $\bar{\mathsf{Y}} \in \mathbb{R}^{m\times r}$ and all columns of Z in $\mathcal{K}$, since by assumption also $\mathbf{o} \in \mathcal{K}$ (which holds automatically if $\mathcal{K}$ is closed), so that it certifies the membership of M in $\mathcal{CPP}(\mathcal{K}\times\mathbb{R}^m)$.

7. $\mathcal{COP}(\mathcal{K}_1 \cup \mathcal{K}_2) = \mathcal{COP}(\mathcal{K}_1) \cap \mathcal{COP}(\mathcal{K}_2)$.
Here we use the fact that

$$\mathbf{x}^\mathsf{T}\mathsf{M}\mathbf{x} \geq 0 \quad \text{for all } \mathbf{x}\in\mathcal{K}_1\cup\mathcal{K}_2 \iff$$
$$\iff \mathbf{x}^\mathsf{T}\mathsf{M}\mathbf{x} \geq 0 \quad \text{for all } \mathbf{x}\in\mathcal{K}_1 \text{ and } \mathbf{x}^\mathsf{T}\mathsf{M}\mathbf{x} \geq 0 \quad \text{for all } \mathbf{x}\in\mathcal{K}_2$$

8. $\mathcal{CPP}(\mathcal{K}_1 \cup \mathcal{K}_2) = \mathcal{CPP}(\mathcal{K}_1) + \mathcal{CPP}(\mathcal{K}_2)$.
For $\mathbf{x}_i \in \mathcal{K}_1\cup\mathcal{K}_2, i\in[1\!:\!k]$, we can divide the vectors into two groups $\mathbf{y}_j \in \mathcal{K}_1$, $\mathbf{z}_r \in \mathcal{K}_2$ so that $\sum_i \mathbf{x}_i\mathbf{x}_i^\mathsf{T} = \sum_j \mathbf{y}_j\mathbf{y}_j^\mathsf{T} + \sum_r \mathbf{z}_r\mathbf{z}_r^\mathsf{T}$.
9. $\mathcal{CPP}(\mathrm{conv}\mathcal{K}) \supseteq \mathcal{CPP}(\mathcal{K})$ with equality if $\mathcal{K}$ is convex.
Follows by 2., since conv$\mathcal{K}$ contains $\mathcal{K}$.
10. $\mathcal{COP}(\mathrm{conv}\mathcal{K}) \subseteq \mathcal{COP}(\mathcal{K})$ with equality if $\mathcal{K}$ is convex.
Again the set inclusion is obvious by 3.
11. $\mathcal{CPP}(\mathcal{K}) = \left\{ \sum_{i=1}^k \mathbf{x}_i\mathbf{x}_i^\mathsf{T} : \mathbf{x}_i \in \mathrm{int}\mathcal{K},\ \mathrm{span}\{\mathbf{x}_1,\dots,\mathbf{x}_k\} = \mathbb{R}^n \right\}$ if $\mathcal{K}$ is closed, convex and int$\mathcal{K} \neq \emptyset$.
The proof for $\mathcal{K} = \mathbb{R}^n_+$ presented in Dickinson (2010) can be extended to any closed convex cone $\mathcal{K}$ with nonempty interior.
12. $\mathcal{COP}(\mathcal{K}) = \mathrm{cl}\mathcal{COP}(\mathcal{K}) = \mathcal{COP}(\mathrm{cl}\mathcal{K})$ while $\mathcal{CPP}(\mathrm{cl}\mathcal{K}) = \mathrm{cl}\mathcal{CPP}(\mathcal{K})$.
The first two equalities follow from the continuity of quadratic functions. The last one is obtained as follows: it is clear from continuity that $\mathcal{CPP}(\mathrm{cl}\,\mathcal{K}) \subseteq \mathrm{cl}[\mathcal{CPP}(\mathcal{K})]$. To establish the reverse inclusion, consider a sequence with $\boldsymbol{\lambda}^{(\nu)} \in \Delta^k$ and $\mathbf{x}_i^{(\nu)} \in \mathcal{K}$, all $i \in [1\!:\!k]$:

$$\mathsf{X}^{(\nu)} = \sum_{i=1}^k \lambda_i^{(\nu)}\mathbf{x}_i^{(\nu)}[\mathbf{x}_i^{(\nu)}]^\mathsf{T} \to \mathsf{A} \in \mathcal{S}^n \quad \text{as } \nu\to\infty\,.$$

Because

$$\|\sqrt{\lambda_i^{(\nu)}}\mathbf{x}_i^{(\nu)}\|^2 \leq \sum_{j=1}^k \lambda_j^{(\nu)}\|\mathbf{x}_j^{(\nu)}\|^2 = \mathrm{Tr}(\mathsf{X}^{(\nu)}) \to \mathrm{Tr}(\mathsf{A})$$

remains bounded for all $i\in[1\!:\!k]$, we may select a subsequence along which $\sqrt{\lambda_i^{(\nu)}}\mathbf{x}_i^{(\nu)} \to \mathbf{y}_i \in \mathrm{cl}\,\mathcal{K}$ as $\nu\to\infty$, for all $i\in[1\!:\!k]$. Again by continuity, we infer

$$\mathsf{A} = \sum_{i=1}^k \tfrac{1}{k}\,[\sqrt{k}\mathbf{y}_i][\sqrt{k}\mathbf{y}_i]^\mathsf{T} \in \mathcal{CPP}(\mathrm{cl}\,\mathcal{K})\,,$$

which shows $\mathrm{cl}[\mathcal{CPP}(\mathcal{K})] \subseteq \mathcal{CPP}(\mathrm{cl}\,\mathcal{K})$.
13. $\mathcal{COP}(\mathcal{K}) = \mathcal{COP}(\mathrm{relint}\mathcal{K})$, if $\mathcal{K}$ is convex.
follows from 12. and the fact that for all convex sets $\mathcal{K}$, we have $\mathrm{cl}\,\mathcal{K} = \mathrm{cl}(\mathrm{relint}\mathcal{K})$.
14. $\mathrm{int}\mathcal{COP}(\mathcal{K}) = \left\{\mathsf{Q}\in\mathcal{S}^n : \mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} > 0 \text{ for all } \mathbf{x}\in\mathcal{K}\setminus\{\mathbf{o}\}\right\}$.
The statement follows from

$$\begin{aligned}\mathrm{int}\mathcal{COP}(\mathcal{K}) &= \{\mathsf{Q} : \mathsf{Q}\bullet\mathsf{X} > 0 \text{ for all } \mathsf{X}\in\mathcal{CPP}(\mathcal{K})\setminus\{\mathsf{O}\}\}\\ &= \{\mathsf{Q} : \mathsf{Q}\bullet\mathsf{X} > 0 \text{ for all } \mathsf{X}\in\mathcal{CPP}(\mathcal{K})\setminus\{\mathsf{O}\}\}\\ &= \{\mathsf{Q} : \mathsf{Q}\bullet\mathsf{X} > 0 \text{ for all } \mathsf{X}\in\mathrm{ext}(\mathcal{CPP}(\mathcal{K}))\setminus\{\mathsf{O}\}\}\\ &= \left\{\mathsf{Q} : \mathbf{x}^\mathsf{T}\mathsf{Q}\mathbf{x} > 0 \text{ for all } \mathbf{x}\in\mathcal{K}\setminus\{0\}\right\}. \quad\square\end{aligned}$$

References

Anstreicher, K. M. (2009). Semidefinite programming versus the reformulation-linearization technique for nonconvex quadratically constrained quadratic programming. *Journal of Global Optimization, 43*(2–3), 471–484.
Anstreicher, K. M. (2012). On convex relaxations for quadratically constrained quadratic programming. *Mathematical Programming, 136*(2), 233–251.
Anstreicher, K. M. (2022). Solving two-trust-region subproblems using semidefinite optimization with eigenvector branching. *Journal of Optimization Theory and Applications*, 1–17.
Anstreicher, K. M., & Burer, S. (2005). DC versus copositive bounds for standard QP. *Journal of Global Optimization, 33*(2), 299–312.
Anstreicher, K. M., & Burer, S. (2010). Computable representations for convex hulls of low-dimensional quadratic forms. *Mathematical Programming, 124*(1-2), 33–43.
Anstreicher, K. M., & Gabl, M. (2022). Solving generalized copositive optimization problems using outer approximations, cutting surfaces and set-copositivity detection. Preprint, in preparation.
Badenbroek, R., & de Klerk, E. (2022). An analytic center cutting plane method to determine complete positivity of a matrix. *INFORMS Journal on Computing, 34*(2), 1115–1125.
Balas, E. (1979). Disjunctive programming. *Annals of Discrete Mathematics, 5*, 3–51.
Bemporad, A., Borrelli, F., & Morari, M. (2003). Min-max control of constrained uncertain discrete-time linear systems. *IEEE Transactions on Automatic Control, 48*(9), 1600–1606.
Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization*. Princeton and Oxford: Princeton University Press.
Ben-Tal, A., El Housni, O., & Goyal, V. (2020). A tractable approach for designing piecewise affine policies in dynamic robust optimization. *Mathematical Programming, 182*, 57–102.
Ben-Tal, A., Goryashko, A., Guslitzer, E., & Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming, 99*(2), 351–376.
Berman, A., & Shaked-Monderer, N. (2003). *Completely positive matrices*. World Scientific.
Bertsimas, D., & Brown, D. B. (2009). Constructing uncertainty sets for robust linear optimization. *Operations Research, 57*(6), 1483–1495.
Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review, 53*(3), 464–501.
Bertsimas, D., & Goyal, V. (2012). On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical Programming, 134*(2), 491–531.
Bertsimas, D., Gupta, V., & Kallus, N. (2018). Data-driven robust optimization. *Mathematical Programming, 167*(2), 235–292.
Bertsimas, D., den Hertog, D., Pauphilet, J., & Zhen, J. (2022). Robust convex optimization: A new perspective that unifies and extends. *Mathematical Programming*. Online 12 Sept.2022
Bertsimas, D., Iancu, D. A., & Parrilo, P. A. (2010). Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research, 35*(2), 363–394.
Bienstock, D. (2016). A note on polynomial solvability of the CDT problem. *SIAM Journal on Optimization, 26*(1), 488–498.
Bomze, I. M. (2012). Copositive optimization-recent developments and applications. *European Journal of Operational Research, 216*(3), 509–520.
Bomze, I. M. (2015). Copositive relaxation beats Lagrangian dual bounds in quadratically and linearly constrained quadratic optimization problems. *SIAM Journal on Optimization, 25*(3), 1249–1275.
Bomze, I. M., Dür, M., de Klerk, E., Roos, C., Quist, A., & Terlaky, T. (2000). On copositive programming and standard quadratic optimization problems. *Journal of Global Optimization, 18*(4), 301–320.
Bomze, I. M., & Gabl, M. (2021). Interplay of non-convex quadratically constrained problems with adjustable robust optimization. *Mathematical Methods of Operations Research, 93*, 115–151.
Bomze, I. M., Gabl, M., Maggioni, F., & Pflug, G. C. (2021a). Two-stage stochastic standard quadratic optimization. *European Journal of Operational Research, 299*(1), 21–34.
Bomze, I. M., Jeyakumar, V., & Li, G. (2018). Extended trust-region problems with one or two balls: exact copositive and Lagrangian relaxations. *Journal of Global Optimization, 71*(3), 551–569.
Bomze, I. M., Kahr, M., & Leitner, M. (2021b). Trust your data or not - StQP remains StQP: Community detection via robust standard quadratic optimization. *Mathematics of Operations Research, 46*(1), 301–316.
Bomze, I. M., & de Klerk, E. (2002). Solving standard quadratic optimization problems via linear, semidefinite and copositive programming. *Journal of Global Optimization, 24*(2), 163–185.
Bomze, I. M., & Overton, M. L. (2015). Narrowing the difficulty gap for the Celis–Dennis–Tapia problem. *Mathematical Programming, 151*(2), 459–476.
Bomze, I. M., & Peng, B. (2022). Conic formulation of QPCCs applied to truly sparse QPs. *Computational Optimization and Application, online 13 Dec 2022*. https://doi.org/10.1007/s10589-022-00440-5.
Bomze, I. M., Schachinger, W., & Uchida, G. (2012). Think co(mpletely)positive! Matrix properties, examples and a clustered bibliography on copositive optimization. *Journal of Global Optimization, 52*(3), 423–445.
Bundfuss, S., & Dür, M. (2008). Algorithmic copositivity detection by simplicial partition. *Linear Algebra and its Applications, 428*(7), 1511–1523.
Bundfuss, S., & Dür, M. (2009). An adaptive linear approximation algorithm for copositive programs. *SIAM Journal on Optimization, 20*(1), 30–53.
Burer, S. (2009). On the copositive representation of binary and continuous nonconvex quadratic programs. *Mathematical Programming, 120*(2), 479–495.
Burer, S. (2012). Copositive programming. In M. F. Anjos, & J. B. Lasserre (Eds.), *Handbook of semidefinite, cone and polynomial optimization: Theory, algorithms, software and applications*. In *International series in operations research and management science* (pp. 201–218). New York: Springer.
Burer, S. (2015). A gentle, geometric introduction to copositive optimization. *Mathematical Programming, 151*(1), 89–116.

Burer, S., & Anstreicher, K. M. (2013). Second-order-cone constraints for extended trust-region subproblems. *SIAM Journal on Optimization, 23*(1), 432–451.
Dickinson, P. (2010). An improved characterisation of the interior of the completely positive cone. *The Electronic Journal of Linear Algebra, 20*, 723–729.
Dickinson, P. J. (2013). *The copositive cone, the completely positive cone and their generalisations*. The Netherlands: University of Groningen Ph.D. dissertation.
Dickinson, P. J., & Povh, J. (2013). Moment approximations for set-semidefinite polynomials. *Journal of Optimization Theory and Applications, 159*(1), 57–68.
Dickinson, P. J. C., & Gijben, L. (2014). On the computational complexity of membership problems for the completely positive cone and its dual. *Computational Optimization and Applications, 57*(2), 403–415.
Dür, M. (2010). Copositive programming-a survey. In M. Diehl, F. Glineur, E. Jarlebring, & W. Michiels (Eds.), *Recent advances in optimization and its applications in engineering* (pp. 3–20). Berlin Heidelberg: Springer.
Dür, M., & Rendl, F. (2021). Conic optimization: A survey with special focus on copositive optimization and binary quadratic problems. *EURO Journal of Computational Optimization, 9*, 100021.
Dür, M., & Still, G. (2008). Interior points of the completely positive cone. *The Electronic Journal of Linear Algebra, 17*, 48–53.
Eichfelder, G., & Povh, J. (2013). On the set-semidefinite representation of nonconvex quadratic programs over arbitrary feasible sets. *Optimization Letters, 7*(6), 1373–1386.
Esfahani, P. M., & Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming, 171*(1), 115–166.
Fan, X., & Hanasusanto, G. A. (2021). A decision rule approach for two-stage data-driven distributionally robust optimization problems with random recourse. arXiv preprint arXiv:2110.00088
Gabl, M. (2022). Sparse conic reformulation of structured QCQPs based on copositive optimization with applications in stochastic optimization. arXiv preprint arXiv:2101.06219
Gorissen, B. L., Yanikoglu, I., & den Hertog, D. (2015). A practical guide to robust optimization. *Omega, 53*, 124–137.
Groetzner, P., & Dür, M. (2020). A factorization method for completely positive matrices. *Linear Algebra and its Applications, 591*, 1–24.
Hall, M., Jr., & Newman, M. (1963). Copositive and completely positive quadratic forms. *Proceedings of the Cambridge Philosophical Society, 59*, 329–339.
Hanasusanto, G. A., & Kuhn, D. (2018). Conic programming reformulations of two-stage distributionally robust linear programs over Wasserstein balls. *Operations Research, 66*(3), 849–869.
Iancu, D. A., Sharma, M., & Sviridenko, M. (2013). Supermodularity and affine policies in dynamic robust optimization. *Operations Research, 61*(4), 941–956.
Jeyakumar, V., Li, G., & Woolnough, D. (2021). Quadratically adjustable robust linear optimization with inexact data via generalized S-lemma: Exact second-order cone program reformulations. *EURO Journal on Computational Optimization, 9*, 100019.
Jiang, R., Ryu, M., & Xu, G. (2019). Data-driven distributionally robust appointment scheduling over Wasserstein balls. arXiv preprint arXiv:1907.03219
Kim, S., Kojima, M., & Toh, K.-C. (2020). A geometrical analysis on convex conic reformulations of quadratic and polynomial optimization problems. *SIAM Journal on Optimization, 30*(2), 1251–1273.
de Klerk, E., & Pasechnik, D. V. (2002). Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization, 12*(4), 875–892.
Kong, Q., Lee, C.-Y., Teo, C.-P., & Zheng, Z. (2013). Scheduling arrivals to a stochastic service delivery system using copositive cones. *Operations Research, 61*(3), 711–726.
Kong, Q., Li, S., Liu, N., Teo, C.-P., & Yan, Z. (2015). Appointment scheduling under schedule-dependent patient no-show behavior. *Management Science, 68*(8), 3480–3500.
Lasserre, J. B. (2001). Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization, 11*(3), 796–817.
Mittal, A., Gökalp, C., & Hanasusanto, G. A. (2019). Robust quadratic programming with mixed-integer uncertainty. *INFORMS Journal on Computing, 32*(2), 201–218.
Mittal, A., & Hanasusanto, G. A. (2021). Finding minimum volume circumscribing ellipsoids using generalized copositive programming. *Operations Research, 70*(5), 2867–2882.
Motzkin, T. (1818). Copositive quadratic forms. *National Bureau of Standards Report, 1952*, 11–22.
Natarajan, K., & Teo, C.-P. (2017). On reduced semidefinite programs for second order moment bounds with applications. *Mathematical Programming, 161*(1), 487–518.
Natarajan, K., Teo, C.-P., & Zheng, Z. (2011). Mixed 0–1 linear programs under objective uncertainty: A completely positive representation. *Operations Research, 59*(3), 713–728.
Padmanabhan, D., Natarajan, K., & Murthy, K. (2021). Exploiting partial correlations in distributionally robust optimization. *Mathematical Programming, 186*(1), 209–255.
Pardalos, P. M., & Vavasis, S. A. (1991). Quadratic programming with one negative eigenvalue is NP-hard. *Journal of Global Optimization, 1*(1), 15–22.
Parrilo, P. A. (2000a). Semidefinite programming based tests for matrix copositivity. In *Proceedings of the 39th IEEE conference on decision and control (Cat. No. 00CH37187): vol. 5* (pp. 4624–4629). IEEE.
Parrilo, P. A. (2000b). *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology Ph.D. Thesis.
Pataki, G. (1998). On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of Operations Research, 23*(2), 177–203.
Peña, J., Vera, J., & Zuluaga, L. F. (2007). Computing the stability number of a graph via linear and semidefinite programming. *SIAM Journal on Optimization, 18*(1), 87–105.
Pólik, I., & Terlaky, T. (2007). A survey of the S-lemma. *SIAM Review, 49*(3), 371–418.
Rahimian, H., & Mehrotra, S. (2019). Distributionally robust optimization: A review. ArXiv preprint arXiv:1908.05659
Rockafellar, R. T. (2015). *Convex analysis*. Princeton and Oxford: Princeton University Press.
Shor, N. Z. (1987). Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences, 25*(6), 1–11.
Sponsel, J., Bundfuss, S., & Dür, M. (2012). An improved algorithm to test copositivity. *Journal of Global Optimization, 52*(3), 537–551.
Sturm, J. F., & Zhang, S. (2003). On cones of nonnegative quadratic functions. *Mathematics of Operations Research, 28*(2), 246–267.
Tuncel, L. (2001). On the Slater condition for the SDP relaxations of nonconvex sets. *Operations Research Letters, 29*(4), 181–186.
Tuncel, L., & Wolkowicz, H. (2012). Strong duality and minimal representations for cone optimization. *Computational Optimization and Applications, 53*(2), 619–648.
Wiesemann, W., Kuhn, D., & Sim, M. (2014). Distributionally robust convex optimization. *Operations Research, 62*(6), 1358–1376.
Woolnough, D., Jeyakumar, V., & Li, G. (2021). Exact conic programming reformulations of two-stage adjustable robust linear programs with new quadratic decision rules. *Optimization Letters, 15*, 25–44.
Xia, W., Vera, J. C., & Zuluaga, L. F. (2020). Globally solving nonconvex quadratic programs via linear integer programming techniques. *INFORMS Journal on Computing, 32*(1), 40–56.
Xu, G., & Burer, S. (2018). A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Computational Optimization and Applications, 70*(1), 33–59.
Xu, G., & Hanasusanto, G. A. (2018). Improved decision rule approximations for multi-stage robust optimization via copositive programming. Preprint, Available at http://www.optimization-online.org/DB_HTML/2018/08/6776.html.
Yakubovich, V. A. (1971). S-procedure in nonlinear control theory. *Vestnik Leningrad University, 1*, 62–77.
Yan, Z., Gao, S. Y., & Teo, C.-P. (2018). On the design of sparse but efficient structures in operations. *Management Science, 64*(7), 3421–3445.
Yang, B., Anstreicher, K., & Burer, S. (2016). Quadratic programs with hollows. *Mathematical Programming, 170*(2), 541–553.
Yang, B., & Burer, S. (2013). A two-variable analysis of the two-trust-region subproblem. *SIAM Journal on Optimization, 26*(1), 661–680.
Yanikoglu, I., Gorissen, B., & den Hertog, D. (2019). A survey of adjustable robust optimization. *European Journal of Operational Research, 277*, 799–813.
Yıldırım, E. A. (2012). On the accuracy of uniform polyhedral approximations of the copositive cone. *Optimization Methods and Software, 27*(1), 155–173.
Zhao, C., & Guan, Y. (2018). Data-driven risk-averse stochastic optimization with Wasserstein metric. *Operations Research Letters, 46*(2), 262–267.
Zhen, J., Marandi, A., de Moor, D., den Hertog, D., & Vandenberghe, L. (2022). Disjoint bilinear optimization: A two-stage robust optimization perspective. *INFORMS Journal on Computing, 34*(5), 2410–2427.