

Spectral relaxations and branching strategies for global optimization of mixed-integer quadratic programs

Carlos J. Nohra* Arvind U. Raghunathan† Nikolaos V. Sahinidis‡

October 13, 2020

Abstract

We consider the global optimization of nonconvex quadratic programs and mixed-integer quadratic programs. We present a family of convex quadratic relaxations which are derived by convexifying nonconvex quadratic functions through perturbations of the quadratic matrix. We investigate the theoretical properties of these quadratic relaxations and show that they are equivalent to some particular semidefinite programs. We also introduce novel branching variable selection strategies which can be used in conjunction with the quadratic relaxations investigated in this paper. We integrate the proposed relaxation and branching techniques into the global optimization solver BARON, and test our implementation by conducting numerical experiments on a large collection of problems. Results demonstrate that the proposed implementation leads to very significant reductions in BARON’s computational times to solve the test problems.

1 Introduction

We address the global optimization of nonconvex quadratic programs (QPs) and mixed-integer quadratic programs (MIQPs) of the form:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & x^T Q x + q^T x \\ \text{s.t.} \quad & Ax = b \\ & Cx \leq d \\ & l \leq x \leq u \\ & x_i \in \mathbb{Z}, \quad \forall i \in J \subseteq \{1, \dots, n\} \end{aligned} \tag{1}$$

where $Q \in \mathbb{R}^{n \times n}$ is a symmetric matrix which may be indefinite, $q \in \mathbb{R}^n$, $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $C \in \mathbb{R}^{p \times n}$, and $d \in \mathbb{R}^p$. We assume that lower and upper bounds are finite, i.e., $-\infty < l_i < u_i < \infty$, $\forall i \in \{1, \dots, n\}$. For the sake of brevity, we use the notation $\mathcal{X} = \{x \in \mathbb{R}^n \mid Ax = b, Cx \leq d, l \leq x \leq u\}$ in the rest of the paper. Note also that even though we allow (1) to include constraints of the form $Cx \leq d$, we do not use information from these inequalities in order to convexify this problem.

QPs and MIQPs of the form (1) arise in a wide variety of applications including facility location and quadratic assignment [18], molecular conformation [23] and max-cut problems [14]. Given their practical importance, these problems have been studied extensively and are known to be very challenging to solve to global optimality.

*Mitsubishi Electric Research Laboratories (nohra@merl.com).

†Mitsubishi Electric Research Laboratories (raghunathan@merl.com)

‡Department of Chemical Engineering, Carnegie Mellon University (sahinidis@cmu.edu).

State-of-the-art global optimization solvers rely on spatial branch-and-bound algorithms to solve (1) to global optimality. The efficiency of these algorithms primarily depends on the quality of the relaxations utilized in the bounding step. Commonly used relaxations for non-convex QPs and MIQPs can be broadly classified in three groups. The first group consists of polyhedral relaxations typically derived via factorable programming methods [20, 27] and reformulation-linearization techniques (RLT) [25]. The second group is given by semidefinite programming (SDP) relaxations [6, 11, 26]. The third group involves convex quadratic relaxations derived through separable programming procedures [22], d.c. programming techniques [28], and quadratic convex reformulation methods [7, 8].

In this paper, we investigate a family of relaxations which falls under the third group. In particular, we consider convex quadratic relaxations derived by convexifying the objective function of (1) through uniform diagonal perturbations of Q . We revisit a very well-known technique which uses the smallest eigenvalue of Q to convexify the function $x^T Q x$. Through numerical experiments, we show that, despite its simplicity, this technique leads to convex quadratic relaxations which are often significantly tighter than the polyhedral relaxations typically used by state-of-the-art global optimization solvers. Motivated by these promising results, we refine this approach in several directions and make several theoretical and algorithmic contributions.

Our first contribution is a novel convex quadratic relaxation for problems of the form (1), derived by using information from both Q and the equality constraints $Ax = b$. Under this approach, the function $x^T Q x$ is convexified by constructing a perturbation of Q obtained by solving a generalized eigenvalue problem involving both Q and A . We show that the resulting relaxation is at least as tight as the relaxation constructed by using the smallest eigenvalue of Q .

In our second contribution, we consider another convex quadratic relaxation in which the function $x^T Q x$ is convexified by using the smallest eigenvalue of $Z^T Q Z$, where Z is a basis for the nullspace of A . We devise a simple procedure which allows us to approximate the bound given by this relaxation without having to compute Z . Moreover, we show that the relaxations obtained through this technique are at least as tight as the other two quadratic relaxations mentioned above. Unlike the polyhedral and SDP relaxations which are formulated in a higher dimensional space, the quadratic relaxations considered in this paper are constructed in the space of the original problem variables, which makes them very inexpensive to solve.

In our third contribution, we prove that the aforementioned quadratic relaxations are equivalent to some particular SDP relaxations. These results facilitate the theoretical comparisons with other relaxations that have been proposed in the literature. In particular, we show that the relaxation based on the smallest eigenvalue of $Z^T Q Z$ is the best among the class of relaxations considered in this paper.

Our fourth contribution is a method for improving the proposed quadratic relaxations with branching. We introduce a novel eigenvalue-based branching variable selection strategy for nonconvex binary quadratic programs. This strategy involves an effective approximation of the impact of branching decisions on relaxation quality.

In order to investigate the impact of the proposed techniques on the performance of branch-and-bound algorithms, we implement the quadratic relaxations and branching strategies considered in this paper in the state-of-the-art global optimization solver BARON [24]. The new quadratic relaxations are incorporated in BARON's portfolio of relaxations and are invoked according to a new dynamic relaxation selection rule which switches between different classes of relaxations based on their relative strength. We test our implementation

by conducting numerical experiments on a large collection of problems. Results demonstrate that the proposed implementation leads to a very significant improvement in the performance of BARON. Moreover, for many of the test problems, our implementation results in a new version of BARON which outperforms other state-of-the-art solvers including CPLEX and GUROBI.

The remainder of this paper is organized as follows. In §2 we review various relaxations which have been considered in the literature for bounding nonconvex QPs and MIQPs. Then, in §3 we present the convex quadratic relaxations considered in this paper and investigate their theoretical properties. In §4 we introduce novel eigenvalue-based branching strategies. This is followed by a description of our implementation in §5. In §6, we present the results of a computational study which includes a comparison between different classes of relaxations, an analysis of the impact of the proposed implementation on the performance of BARON, and a comparison between several global optimization solvers. Finally, §7 presents conclusions from this work.

Notation

We denote by \mathbb{Z} , \mathbb{R} and $\mathbb{R}_{\geq 0}$ the set integer, real and nonnegative real numbers, respectively. We use $\mathbf{1} \in \mathbb{R}^n$ to denote a vector of ones. The i -th element of $x \in \mathbb{R}^n$ is denoted by x_i . Given $d \in \mathbb{R}^n$, $\text{diag}(d)$ denotes the diagonal matrix whose diagonal entries are given by the elements of d . The i -th row of a matrix $A \in \mathbb{R}^{m \times n}$ is denoted by $A_{i\cdot}$, and its (i, j) -th entry by A_{ij} . Let \mathbb{S}^n denote the set of $n \times n$ real, symmetric matrices. Given $M \in \mathbb{S}^n$, we use λ_i to represent its i -th eigenvalue and v^i for the corresponding eigenvector. For $M \in \mathbb{S}^n$, the notation $M \succcurlyeq 0$ and $M \succ 0$, indicates that M is positive semidefinite and positive definite, respectively. We denote by I_n the $n \times n$ identity matrix. Let $M, N \in \mathbb{S}^n$ with $N \succ 0$. We use $\lambda_{\min}(M)$ to represent the smallest eigenvalue of M . Similarly, we denote by $\lambda_{\min}(M, N)$ the smallest generalized eigenvalue of the problem $Mv = \lambda Nv$, where $v \in \mathbb{R}^n$. The inner product between $M, P \in \mathbb{S}^n$ is denoted by $\langle M, P \rangle = \sum_{i=1}^n \sum_{j=1}^n M_{ij}P_{ij}$.

2 Current relaxations for nonconvex QPs and MIQPs

In this section, we review various types of relaxations that have been proposed for bounding (1).

2.1 Polyhedral relaxations

One of the simplest relaxations for (1) can be derived via factorable programming techniques [20, 27], leading to the linear program:

$$\min_{x \in \mathcal{X}, X} \sum_{i=1}^n \sum_{j=1}^n Q_{ij} X_{ij} + \sum_{i=1}^n q_i x_i \quad (2a)$$

$$\text{s.t. } X_{ij} \geq l_i x_j + l_j x_i - l_i l_j, \quad i = 1, \dots, n, j = i, \dots, n, \quad (2b)$$

$$X_{ij} \geq u_i x_j + u_j x_i - u_i u_j, \quad i = 1, \dots, n, j = i, \dots, n, \quad (2c)$$

$$X_{ij} \leq l_i x_j + u_j x_i - l_i u_j, \quad i = 1, \dots, n, j = i, \dots, n, \quad (2d)$$

$$X_{ij} \leq u_i x_j + l_j x_i - u_i l_j, \quad i = 1, \dots, n, j = i, \dots, n, \quad (2e)$$

$$X_{ij} = X_{ji}, \quad i = 1, \dots, n, j = (i + 1), \dots, n, \quad (2f)$$

where X is a symmetric matrix of introduced variables and (2b)–(2e) are the so-called *McCormick inequalities*. This relaxation is often referred to as the *McCormick relaxation* of (1). Even though this relaxation is simple to implement, it often leads to relatively weak bounds, and as a result, it is typically tightened by adding various classes of valid inequalities [5, 10, 21, 29].

Another polyhedral relaxation for (1) can be constructed through the reformulation linearization techniques (RLT), obtaining the linear program [25]:

$$\min_{x \in \mathcal{X}, X} \sum_{i=1}^n \sum_{j=1}^n Q_{ij} X_{ij} + \sum_{i=1}^n q_i x_i \quad (3a)$$

$$\text{s.t. Eqs. (2b) – (2f)} \quad (3b)$$

$$\sum_{i=1}^n A_{ki} X_{ij} = b_k x_j, \quad k = 1, \dots, m, \quad j = 1, \dots, n \quad (3c)$$

$$\sum_{i=1}^n C_{ki} X_{ij} - l_j \sum_{i=1}^n C_{ki} x_i - d_k x_j \leq -l_j d_k, \quad k = 1, \dots, p, \quad j = 1, \dots, n \quad (3d)$$

$$- \sum_{i=1}^n C_{ki} X_{ij} + u_j \sum_{i=1}^n C_{ki} x_i + d_k x_j \leq u_j d_k, \quad k = 1, \dots, p, \quad j = 1, \dots, n, \quad (3e)$$

$$- \sum_{i=1}^n \sum_{j=1}^n C_{ki} C_{lj} X_{ij} + \sum_{i=1}^n (d_l C_{ki} + d_k C_{li}) x_i \leq d_k d_l, \quad k, l = 1, \dots, p \quad (3f)$$

This relaxation is often referred to as the *first-level RLT relaxation* of (1).

2.2 SDP relaxations

The simplest SDP relaxation for (1) is given by:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (4a)$$

$$\text{s.t. } X - xx^T \succeq 0 \quad (4b)$$

This relaxation is often referred to as the *Shor relaxation* of (1) [26]. The Shor relaxation can be strengthened by including additional valid constraints. This can be achieved, for instance, by adding the McCormick inequalities corresponding to the diagonal elements of the matrix X , which results in the following SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (5a)$$

$$\text{s.t. } X - xx^T \succeq 0 \quad (5b)$$

$$X_{ii} \geq 2l_i x_i - l_i^2, \quad i = 1, \dots, n \quad (5c)$$

$$X_{ii} \geq 2u_i x_i - u_i^2, \quad i = 1, \dots, n \quad (5d)$$

$$X_{ii} \leq u_i x_i + l_i x_i - u_i l_i, \quad i = 1, \dots, n \quad (5e)$$

It is easy to verify that (5c) and (5d) are implied by $X - xx^T \succeq 0$ and hence redundant in this formulation. The SDP (5) can be further tightened by including constraints derived from $Ax = b$. For example, we can add the constraints obtained by lifting the valid equalities

$\sum_{j=1}^n A_{kj}x_i x_j = b_k x_i$, $k = 1, \dots, m$, $i = 1, \dots, n$, to the space of (x, X) . This leads to the following SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (6a)$$

$$\text{s.t. Eqs. (5b) - (5e)} \quad (6b)$$

$$\sum_{j=1}^n A_{kj} X_{ij} = b_k x_i, \quad k = 1, \dots, m, \quad i = 1, \dots, n \quad (6c)$$

Alternatively, we can add a single constraint derived by lifting the equality $(Ax - b)^T(Ax - b) = 0$ to the space of (x, X) , obtaining the following SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (7a)$$

$$\text{s.t. Eqs. (5b) - (5e)} \quad (7b)$$

$$\langle A^T A, X \rangle - 2(A^T b)^T x + b^T b = 0 \quad (7c)$$

The SDPs (6) and (7) are equivalent (see Proposition 5 in [13] for details).

2.3 Convex quadratic relaxations

In the following, we briefly review three important classes of convex quadratic relaxations for problems of the form (1).

2.3.1 Separable programming relaxation

This relaxation, which is constructed through the eigendecomposition of Q (see [22] for details), is given by:

$$\begin{aligned} \min_{x \in \mathcal{X}, y} \quad & \sum_{i:\lambda_i > 0} \lambda_i y_i^2 + \sum_{i:\lambda_i < 0} \lambda_i ((L_i + U_i)y_i - L_i U_i) + q^T x \\ \text{s.t.} \quad & y_i = v^{iT} x, \quad i = 1, \dots, n \\ & L_i \leq y_i \leq U_i, \quad i = 1, \dots, n \end{aligned} \quad (8)$$

where y_i are introduced variables and the bounds L_i and U_i are respectively determined by minimizing and maximizing the linear function $v^{iT} x$ over \mathcal{X} .

2.3.2 D.C. programming relaxations

The key idea behind this approach is to decompose the objective function of (1) as $f(x) = x^T Q x + q^T x = g(x) - h(x)$, where $g(x)$ and $h(x)$ are both convex quadratic functions [28]. In their most generic form, these relaxations can be expressed as:

$$\min_{x \in \mathcal{X}} \quad g(x) - \bar{h}_{\mathcal{X}}(x) \quad (9)$$

where $\bar{h}_{\mathcal{X}}(x)$ is a concave overestimator of $h(x)$ over \mathcal{X} . A particular d.c. programming relaxation is the classical α BB relaxation [3], which for (1) takes the form:

$$\min_{x \in \mathcal{X}} \quad x^T Q x + q^T x - \sum_{i=1}^n \alpha_i (x_i - l_i)(u_i - x_i) \quad (10)$$

where α_i , $i = 1, \dots, n$, are nonnegative parameters chosen such that the objective function of (10) is convex over \mathcal{X} . Other particular examples of d.c. programming relaxations are the relaxations constructed via undominated d.c. decompositions [9].

2.3.3 Relaxations based on quadratic convex reformulations

These relaxations are used in the Quadratic Convex Reformulation (QCR) methods. To illustrate these techniques, consider the binary quadratic program:

$$\min_{x \in \mathcal{X}_B} x^T Q x + q^T x \quad (11)$$

where $\mathcal{X}_B = \{x \in \{0, 1\}^n \mid Ax = b, Cx \leq d\}$. The QCR approaches involve two steps. The first step consists in reformulating (11) to an equivalent binary quadratic program whose continuous relaxation is convex. In the second step, the reformulated problem is solved using a branch-and-bound algorithm. At each node of the branch-and-bound tree, a lower bound is obtained by solving the continuous relaxation of the reformulated problem, which is a convex quadratic program.

One of the earliest references to these methods is found in a paper by Hammer and Rubin [16], in which the following reformulation for (11) is proposed:

$$\min_{x \in \mathcal{X}_B} x^T Q_\lambda x + q_\lambda^T x \quad (12)$$

where $Q_\lambda = Q - \min(0, \lambda_{\min}(Q))I_n$ and $q_\lambda = q + \min(0, \lambda_{\min}(Q))\mathbb{1}$. It is simple to check that $Q_\lambda \succcurlyeq 0$, and that the objective functions of (11) and (12) are equivalent $\forall x \in \{0, 1\}^n$. Another reformulation of (11) was proposed by Billionnet et al. [8]:

$$\min_{x \in \mathcal{X}_B} x^T Q_{d_q, \Theta_q} x + q_{d_q, \Theta_q}^T x \quad (13)$$

where $Q_{d_q, \Theta_q} = Q + \text{diag}(d_q) + \frac{1}{2}(\Theta_q^T A + A^T \Theta_q)$, $q_{d_q, \Theta_q} = q - d_q - \Theta_q^T b$, $d_q \in \mathbb{R}^n$ and $\Theta_q \in \mathbb{R}^{m \times n}$. The perturbation parameters d_q and Θ_q are chosen such that $Q_{d_q, \Theta_q} \succcurlyeq 0$ and the bound of the continuous relaxation of (13) is maximized. This is done by solving the SDP (6), and setting the entries of d_q and Θ_q to the optimal values of the dual variables associated with the constraints (5e) and (6c), respectively. Note that the continuous relaxation of (13) provides the same bound as the SDP (6).

3 Spectral relaxations for nonconvex QPs and MIQPs

In the following, we present a family of convex quadratic relaxations for problems of the form (1), and investigate their theoretical properties. Before providing a detailed derivation of these relaxations, we state two results which we will repeatedly use throughout this section. First, we recall that the minimum eigenvalue of a matrix M and the minimum generalized eigenvalue of a pair of matrices (M, N) , with $M, N \in \mathbb{S}^n, N \succ 0$, can be expressed in terms of the Rayleigh quotient as [15]:

$$\lambda_{\min}(M) = \min_{x \neq 0} \frac{x^T M x}{x^T x} \quad \text{and} \quad \lambda_{\min}(M, N) = \min_{x \neq 0} \frac{x^T M x}{x^T N x}. \quad (14)$$

Second, we provide a particularly useful formulation for the dual of a certain SDP.

Proposition 1. *Consider the following SDP*

$$\min_{x \in \mathcal{X}, X} \langle \hat{Q}, X \rangle + \hat{q}^T x \quad (15a)$$

$$\text{s.t. } X - x x^T \succcurlyeq 0 \quad (15b)$$

$$\langle \hat{C}_i, X \rangle + \hat{c}_i x + \hat{d}_i = 0, \quad i = 1, \dots, m_1 \quad (15c)$$

$$\langle \bar{C}_i, X \rangle + \bar{c}_i x + \bar{d}_i \leq 0, \quad i = 1, \dots, p_1 \quad (15d)$$

for some $\hat{Q}, \hat{C}_i, \bar{C}_i \in \mathbb{S}^n, \hat{q}, \hat{c}_i, \bar{c}_i \in \mathbb{R}^n$ and $\hat{d}_i, \bar{d}_i \in \mathbb{R}$. The dual of (15) is given by

$$\max_{\alpha \in \mathbb{R}^{m_1}, \beta \in \mathbb{R}_{\geq 0}^{p_1}: \hat{Q}_{\alpha, \beta} \succcurlyeq 0} \left\{ \min_{x \in \mathcal{X}} x^T \hat{Q}_{\alpha, \beta} x + \hat{q}_{\alpha, \beta}^T x + \hat{d}_{\alpha, \beta} \right\}.$$

where $\hat{Q}_{\alpha, \beta} = \hat{Q} + \sum_{i=1}^{m_1} \alpha_i \hat{C}_i + \sum_{i=1}^{p_1} \beta_i \bar{C}_i$, $\hat{q}_{\alpha, \beta} = \hat{q} + \sum_{i=1}^{m_1} \alpha_i \hat{c}_i + \sum_{i=1}^{p_1} \beta_i \bar{c}_i$, and $\hat{d}_{\alpha, \beta} = \sum_{i=1}^{m_1} \alpha_i \hat{d}_i + \sum_{i=1}^{p_1} \beta_i \bar{d}_i$.

Proof. By dualizing the constraints (15c) using the multipliers $\alpha_i \in \mathbb{R}$, for $i = 1, \dots, m_1$, and the constraints (15d) using the multipliers $\beta_i \in \mathbb{R}_{\geq 0}$, for $i = 1, \dots, p_1$, we have that the Lagrangian dual of the SDP (15) is given by:

$$\max_{\alpha \in \mathbb{R}^{m_1}, \beta \in \mathbb{R}_{\geq 0}^{p_1}} \left\{ \begin{array}{l} \min_{x \in \mathcal{X}} \langle \hat{Q}_{\alpha, \beta}, X \rangle + \hat{q}_{\alpha, \beta}^T x + \hat{d}_{\alpha, \beta} \\ \text{s.t.} \quad X - x x^T \succcurlyeq 0 \end{array} \right\} \quad (16)$$

Since the variables x have finite lower and upper bounds, the set \mathcal{X} is bounded. Hence, for the inner minimization to be bounded below, we need to choose α and β such that $\hat{Q}_{\alpha, \beta} \succcurlyeq 0$. This restriction on α and β implies that the optimal solution of the inner minimization problem satisfies $X = x x^T$. The claim follows after substituting $X = x x^T$. \square

3.1 Eigenvalue relaxation

We start by reformulating (1) as:

$$\begin{array}{ll} \min_{x \in \mathcal{X}} & x^T Q x + q^T x + \alpha_e \sum_{i=1}^n x_i^2 - \alpha_e \sum_{i=1}^n x_i^2 \\ \text{s.t.} & x_i \in \mathbb{Z}, \quad \forall i \in J \end{array} \quad (17)$$

where $\alpha_e \geq 0$. By dropping the integrality conditions from (17) and using the concave envelope of x_i^2 over $[l_i, u_i]$, we obtain the following relaxation:

$$\min_{x \in \mathcal{X}} x^T Q_{\alpha_e} x + q_{\alpha_e}^T x + k_{\alpha_e} \quad (18)$$

where $Q_{\alpha_e} = Q + \alpha_e I_n$, $q_{\alpha_e} = q - \alpha_e(l + u)$, and $k_{\alpha_e} = \alpha_e l^T u$.

To ensure that (18) is convex, it suffices to choose $\alpha_e \geq -\min(0, \lambda_{\min}(Q))$, since this renders Q_{α_e} positive semidefinite. Moreover, it is simple to check that $\alpha_e = -\min(0, \lambda_{\min}(Q))$ provides the tightest convex relaxation of the form (18) for which $Q_{\alpha_e} \succcurlyeq 0$. We refer to this relaxation as the *eigenvalue relaxation* of (1).

Two interesting observations can be made on (18) when $\alpha_e \geq -\min(0, \lambda_{\min}(Q))$. First, the derivation of (18) can be seen as an application of the d.c. programming technique reviewed in §2.3.2, whereby the objective function of (1) is expressed as the difference of the convex quadratic functions $g(x) = x^T Q x + q^T x + \alpha_e \sum_{i=1}^n x_i^2$ and $h(x) = \alpha_e \sum_{i=1}^n x_i^2$. Second, (18) is equivalent to the α BB relaxation discussed in §2.3.2 if we set $\alpha_i = \alpha_e$, $\forall i = 1, \dots, n$, in (10).

Note also that if all the variables in (1) are binary and $\alpha_e = -\min(0, \lambda_{\min}(Q))$, then (18) is equivalent to the continuous relaxation of the convex binary quadratic program (12), which was considered by Hammer and Rubin [16].

Even though the eigenvalue relaxation is relatively simple to construct, in many cases it can be significantly tighter than the polyhedral relaxations commonly used in state-of-the-art global optimization solvers (see §6.1). Motivated by this observation, we further investigate the theoretical properties of this relaxation. In particular, we show that the eigenvalue relaxation is equivalent to the following SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (19a)$$

$$\text{s.t. } X - xx^T \succcurlyeq 0 \quad (19b)$$

$$\langle I_n, X \rangle - (l + u)^T x + l^T u \leq 0 \quad (19c)$$

Proposition 2. *Suppose that the matrix Q is indefinite. Denote by $\mu_{\text{EIG}}(\alpha_e)$ and $\mu_{\text{SDP_EIG}}$ the optimal objective function values in (18) and (19), respectively. Then, the following hold:*

(i) *If $\alpha_e \geq -\lambda_{\min}(Q)$ in (18), then $\mu_{\text{EIG}}(\alpha_e) \leq \mu_{\text{SDP_EIG}}$.*

(ii) *If $\alpha_e = -\lambda_{\min}(Q)$ in (18), then $\mu_{\text{EIG}}(\alpha_e) = \mu_{\text{SDP_EIG}}$.*

Proof. The proof of this proposition relies on strong duality holding for (19). We start by showing that (19) admits a strictly feasible solution. Let $\bar{x} \in \mathbb{R}^n$ be a vector such that $A\bar{x} = b$, $C\bar{x} < d$, and $l < \bar{x} < u$. Recall that the concave envelope of x_i^2 over $[l_i, u_i]$ is given by $(l_i + u_i)x_i - l_i u_i$. Since $l_i < \bar{x}_i < u_i, \forall i = 1, \dots, n$, it follows that $(l_i + u_i)\bar{x}_i - l_i u_i - \bar{x}_i^2 > 0, \forall i = 1, \dots, n$. Define $\epsilon := (l_1 + u_1)\bar{x}_1 - l_1 u_1 - \bar{x}_1^2$. Clearly, there exists $\delta \in \mathbb{R}$ such that $0 < \delta < \epsilon$. Let $\bar{X} \in \mathbb{S}^n$ be the matrix satisfying:

$$\begin{aligned} \bar{X}_{11} &= (l_1 + u_1)\bar{x}_1 - l_1 u_1 - \delta \\ \bar{X}_{ii} &= (l_i + u_i)\bar{x}_i - l_i u_i, \quad i = 2, \dots, n \\ \bar{X}_{ij} &= \bar{X}_{ji} = \bar{x}_i \bar{x}_j, \quad i = 1, \dots, n, \quad j = (i + 1), \dots, n \end{aligned} \quad (20)$$

It is simple to check that (19c) is strictly satisfied by (\bar{x}, \bar{X}) . Define $\hat{X} := \bar{X} - \bar{x}\bar{x}^T$. Then, it follows that $\hat{X} \in \mathbb{S}^n$ is a diagonal matrix with entries given by:

$$\begin{aligned} \hat{X}_{11} &= (l_1 + u_1)\bar{x}_1 - l_1 u_1 - \bar{x}_1^2 - \delta \\ \hat{X}_{ii} &= (l_i + u_i)\bar{x}_i - l_i u_i - \bar{x}_i^2, \quad i = 2, \dots, n \end{aligned} \quad (21)$$

It is clear that $\hat{X}_{ii} > 0, \forall i = 1, \dots, n$. Hence, $\hat{X} \succ 0$ and (\bar{x}, \bar{X}) is strictly feasible in (19). Therefore, Slater's condition is satisfied by (19), which implies that strong duality holds and the optimal value of the dual problem is attained.

Now, we consider the dual of (19). By Proposition 1, this dual is given by:

$$\max_{\hat{\alpha}_e \in \mathbb{R}_{\geq 0}: Q_{\hat{\alpha}_e} \succcurlyeq 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\hat{\alpha}_e} x + q_{\hat{\alpha}_e}^T x + k_{\hat{\alpha}_e} \right\} \quad (22)$$

where $\hat{\alpha}_e$ is the multiplier for the constraint (19c), $Q_{\hat{\alpha}_e} = Q + \hat{\alpha}_e I_n$, $q_{\hat{\alpha}_e} = q - \hat{\alpha}_e(l + u)$, and $k_{\hat{\alpha}_e} = \hat{\alpha}_e l^T u$. Since Q is indefinite, it is clear that $Q_{\hat{\alpha}_e} \succcurlyeq 0$ for $\hat{\alpha}_e \geq -\lambda_{\min}(Q)$. For any such choice of $\hat{\alpha}_e$, the inner minimization in (22) takes the same form as the QP (18). Hence, the claim in (i) follows by weak duality. Let $\mu_{\text{DSDP_EIG}}$ be the optimal objective

function value in (22). Then, $\mu_{\text{DSDP_EIG}} = \max\{\mu_{\text{EIG}}(\hat{\alpha}_e) \mid \hat{\alpha}_e \in \mathbb{R}_{\geq 0} : Q_{\hat{\alpha}_e} \succcurlyeq 0\}$. It is easy to verify that $\mu_{\text{EIG}}(\hat{\alpha}_e)$ is monotonically decreasing in $\hat{\alpha}_e$. Hence, $\mu_{\text{DSDP_EIG}} = \mu_{\text{EIG}}(\hat{\alpha}_e)$ occurs when $\hat{\alpha}_e = -\lambda_{\min}(Q)$. Combining this with $\mu_{\text{DSDP_EIG}} = \mu_{\text{SDP_EIG}}$, which holds by strong duality, proves the claim in (ii). \square

3.2 Generalized eigenvalue relaxation

In this section, we consider a new quadratic relaxation which improves the bound of the eigenvalue relaxation by using information from the equality constraints. We start by reformulating (1) as:

$$\begin{aligned} \min_{x \in \mathcal{X}} \quad & x^T Q x + q^T x + \alpha_g \sum_{i=1}^n x_i^2 - \alpha_g \sum_{i=1}^n x_i^2 + \alpha_g \|Ax - b\|^2 \\ \text{s.t.} \quad & x_i \in \mathbb{Z}, \quad \forall i \in J \end{aligned} \quad (23)$$

where $\alpha_g \geq 0$. By dropping the integrality conditions from (23) and using the concave envelope of x_i^2 over $[l_i, u_i]$, we obtain the following quadratic relaxation:

$$\min_{x \in \mathcal{X}} \quad x^T Q_{\alpha_g} x + q_{\alpha_g}^T x + k_{\alpha_g} \quad (24)$$

where $Q_{\alpha_g} = Q + \alpha_g(I_n + A^T A)$, $q_{\alpha_g} = q - \alpha_g(l + u + 2A^T b)$, and $k_{\alpha_g} = \alpha_g(l^T u + b^T b)$.

In the following proposition, we provide a condition for choosing α_g which ensures that the above problem is a convex relaxation of (1).

Proposition 3. *Let $\alpha_g \geq -\min(0, \lambda_{\min}(Q, I_n + A^T A))$ in (24). Then, (24) is a convex quadratic program.*

Proof. To establish the convexity of (24), it suffices to verify that $Q_{\alpha_g} = Q + \alpha_g(I_n + A^T A)$ is positive semidefinite. From the definition of the Rayleigh quotient for the generalized eigenvalue pair $(Q, I_n + A^T A)$ in (14) we obtain:

$$\lambda_{\min}(Q, I_n + A^T A) \leq \frac{x^T Q x}{x^T (I_n + A^T A) x}, \quad \forall x \neq 0 \quad (25)$$

which using the positive definiteness of $(I_n + A^T A)$ can be equivalently written as:

$$x^T Q_{\alpha_g} x \geq (\alpha_g + \lambda_{\min}(Q, I_n + A^T A)) x^T (I_n + A^T A) x, \quad \forall x \neq 0. \quad (26)$$

It is readily verified that $Q_{\alpha_g} \succcurlyeq 0$ for $\alpha_g \geq -\min(0, \lambda_{\min}(Q, I_n + A^T A))$. \square

From Proposition 3, it follows that $\alpha_g = -\min(0, \lambda_{\min}(Q, I_n + A^T A))$ provides the tightest convex relaxation of the form (24) for which $Q_{\alpha_g} \succcurlyeq 0$. We refer to this convex relaxation as the *generalized eigenvalue relaxation* of (1). Next, we show that this relaxation is at least as tight as the eigenvalue relaxation.

Proposition 4. *Suppose that $\alpha_e = -\min(0, \lambda_{\min}(Q))$ in (18) and assume that $\alpha_g = -\min(0, \lambda_{\min}(Q, I_n + A^T A))$ in (24). Denote by μ_{EIG} and μ_{GEIG} the optimal objective function values in (18) and (24), respectively. Then, $\mu_{\text{GEIG}} \geq \mu_{\text{EIG}}$.*

Proof. To prove that $\mu_{\text{GEIG}} \geq \mu_{\text{EIG}}$, it suffices to show that $\alpha_g \leq \alpha_e$. We will use the definition of the Rayleigh quotient in (14). We consider the following cases:

- (a) $\lambda_{\min}(Q) \geq 0$. This implies that $x^T Q x \geq 0, \forall x \in \mathbb{R}^n \setminus \{0\}$. Moreover, it is clear that $x^T(I_n + A^T A)x > 0, \forall x \in \mathbb{R}^n \setminus \{0\}$. Then, from (14) it follows that $\lambda_{\min}(Q, I_n + A^T A) \geq 0$. Hence, $\alpha_e = \alpha_g = 0$, which implies that $\mu_{\text{GEIG}} = \mu_{\text{EIG}}$.
- (b) $\lambda_{\min}(Q) < 0$. This implies that $\exists x \in \mathbb{R}^n$ such that $x^T Q x < 0$. From (14), it follows that $\lambda_{\min}(Q, I_n + A^T A) < 0$. Then, it is clear that $\alpha_e = -\lambda_{\min}(Q)$ and $\alpha_g = -\lambda_{\min}(Q, I_n + A^T A)$. Define the set $D = \{x \in \mathbb{R}^n : x \neq 0, x^T Q x < 0\}$. Clearly, D is nonempty. It is easy to verify that the minimum in (14) occurs for $x \in D$. Combining this observation with $x^T(I + A^T A)x \geq x^T x$, we obtain

$$\frac{x^T Q x}{x^T (I_n + A^T A) x} \geq \frac{x^T Q x}{x^T x}, \quad \forall x \in D. \quad (27)$$

Hence, $\lambda_{\min}(Q, I_n + A^T A) \geq \lambda_{\min}(Q)$, which implies that $\alpha_g \leq \alpha_e$, and $\mu_{\text{GEIG}} \geq \mu_{\text{EIG}}$. □

The idea of using information from the equality constraints to convexify the objective function of (1) has been exploited before in the context of the QCR methods discussed in §2.3.3. Even though our approach also uses information from the equality constraints, it differs from the QCR techniques considered in [7, 8] in three ways. First, we do not seek the development of a reformulation of the original problem but instead the construction of cheap quadratic relaxations which can be incorporated in a branch-and-bound framework. Second, under our approach, at a given node of the branch-and-bound tree, we update the perturbation parameters used to construct these quadratic relaxations. This is done by solving the eigenvalue or generalized eigenvalue problems involving the submatrices of Q and $I_n + A^T A$ obtained after eliminating the rows and columns corresponding to the variables that have been fixed. This update results in tighter bounds, and as shown in §6.2–6.4, it can have a significant impact on the performance of branch-and-bound algorithms, especially in the binary case, in which our relaxations can be used in conjunction with the branching strategy introduced in §4. By contrast, in the QCR methods, the perturbation parameters used to convexify the problem are calculated only once, prior to the initialization of the branch-and-bound tree, and are not updated during the execution of the branch-and-bound algorithm. Third, in our method, the perturbation parameters can be obtained by solving an eigenvalue or generalized eigenvalue problem, which is often inexpensive. Under the QCR approaches, calculating the perturbation parameters involves the solution of an SDP, which is more computationally expensive.

Note also that, unlike our approach, the separable and d.c. programming techniques described in §2.3.1 and §2.3.2 do not use information from the equality constraints to improve the bound of the resulting relaxations.

We next show that the generalized eigenvalue relaxation is equivalent to the following SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (28a)$$

$$\text{s.t. } X - x x^T \succeq 0 \quad (28b)$$

$$\langle I_n, X \rangle - (l + u)^T x + l^T u + \langle A^T A, X \rangle - (2A^T b)^T x + b^T b \leq 0 \quad (28c)$$

Proposition 5. *Suppose that the matrix Q is indefinite. Denote by $\mu_{\text{GEIG}}(\alpha_g)$ and $\mu_{\text{SDP.GEIG}}$ the optimal objective function values in (24) and (28), respectively. Then, the following hold:*

(i) If $\alpha_g \geq -\lambda_{\min}(Q, I_n + A^T A)$ in (24), then $\mu_{\text{GEIG}}(\alpha_g) \leq \mu_{\text{SDP_GEIG}}$.

(ii) If $\alpha_g = -\lambda_{\min}(Q, I_n + A^T A)$ in (24), then $\mu_{\text{GEIG}}(\alpha_g) = \mu_{\text{SDP_GEIG}}$.

Proof. We will rely on strong duality holding for (28) and follow the same line of arguments used in the proof of Proposition of 2. We start by showing that (28) admits a strictly feasible solution. Let $\bar{x} \in \mathbb{R}^n$ be a vector such that $A\bar{x} = b$, $C\bar{x} < d$, and $l < \bar{x} < u$. Recall that the concave envelope of x_i^2 over $[l_i, u_i]$ is given by $(l_i + u_i)x_i - l_i u_i$. Since $l_i < \bar{x}_i < u_i, \forall i = 1, \dots, n$, it follows that $(l_i + u_i)\bar{x}_i - l_i u_i - \bar{x}_i^2 > 0, \forall i = 1, \dots, n$. Define $\epsilon := (l_1 + u_1)\bar{x}_1 - l_1 u_1 - \bar{x}_1^2$. Clearly, there exists $\delta \in \mathbb{R}$ such that $0 < \delta < \epsilon$. Let $\bar{X} \in \mathbb{S}^n$ be the matrix satisfying:

$$\begin{aligned}\bar{X}_{11} &= \frac{(l_1 + u_1)\bar{x}_1 - l_1 u_1 + \Phi_{11}\bar{x}_1^2 - \delta}{1 + \Phi_{11}} \\ \bar{X}_{ii} &= \frac{(l_i + u_i)\bar{x}_i - l_i u_i + \Phi_{ii}\bar{x}_i^2}{1 + \Phi_{ii}}, \quad i = 2, \dots, n, \\ \bar{X}_{ij} &= \bar{X}_{ji} = \bar{x}_i \bar{x}_j, \quad i = 1, \dots, n, \quad j = i + 1, \dots, n\end{aligned}\tag{29}$$

where Φ_{ii} is the (i, i) -th entry of $A^T A$. It is simple to check that (28c) is strictly satisfied by (\bar{x}, \bar{X}) . Define $\hat{X} := \bar{X} - \bar{x}\bar{x}^T$. It is clear that \hat{X} is diagonal with entries:

$$\begin{aligned}\hat{X}_{11} &= \frac{(l_1 + u_1)\bar{x}_1 - l_1 u_1 - \bar{x}_1^2 - \delta}{1 + \Phi_{11}} \\ \hat{X}_{ii} &= \frac{(l_i + u_i)\bar{x}_i - l_i u_i - \bar{x}_i^2}{1 + \Phi_{ii}}, \quad i = 2, \dots, n,\end{aligned}\tag{30}$$

Clearly, $\hat{X}_{ii} > 0, \forall i = 1, \dots, n$. Hence, $\hat{X} \succ$ and (\bar{x}, \bar{X}) is a strictly feasible solution to (28). Therefore, Slater's condition is satisfied by (28), which implies that strong duality holds and the optimal value of the dual problem is attained.

Now, we consider the dual of (28). By Proposition 1, this dual is given by:

$$\max_{\hat{\alpha}_g \in \mathbb{R}_{\geq 0}: Q_{\hat{\alpha}_g} \succ 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\hat{\alpha}_g} x + q_{\hat{\alpha}_g}^T x + k_{\hat{\alpha}_g} \right\}\tag{31}$$

where $\hat{\alpha}_g$ is the multiplier for the constraint (28c), $Q_{\hat{\alpha}_g} = Q + \hat{\alpha}_g(I_n + A^T A)$, $q_{\hat{\alpha}_g} = q - \hat{\alpha}_g(l + u + 2A^T b)$, and $k_{\hat{\alpha}_g} = \hat{\alpha}_g(l^T u + b^T b)$. Since Q is indefinite, Proposition 3 implies that $Q_{\hat{\alpha}_g} \succ 0$ for $\hat{\alpha}_g \geq -\lambda_{\min}(Q, I_n + A^T A)$. For any such choice of $\hat{\alpha}_g$, the inner minimization in (31) takes the same form as the QP (24). Hence, the claim in (i) follows by weak duality. Let $\mu_{\text{DSDP_GEIG}}$ be the optimal objective function value in (31). Then, $\mu_{\text{DSDP_GEIG}} = \max\{\mu_{\text{GEIG}}(\hat{\alpha}_g) \mid \hat{\alpha}_g \in \mathbb{R}_{\geq 0} : Q_{\hat{\alpha}_g} \succ 0\}$. Clearly, $\mu_{\text{GEIG}}(\hat{\alpha}_g)$ is monotonically decreasing in $\hat{\alpha}_g$. Hence, $\mu_{\text{DSDP_GEIG}} = \mu_{\text{GEIG}}(\hat{\alpha}_g)$ occurs when $\hat{\alpha}_g = -\lambda_{\min}(Q, I_n + A^T A)$. Combining this with $\mu_{\text{DSDP_GEIG}} = \mu_{\text{SDP_GEIG}}$, which holds by strong duality, proves the claim in (ii). \square

3.3 Eigenvalue relaxation in the nullspace of the equality constraints

In this section, we consider another convex quadratic relaxation of (1) which also incorporates information from the equality constraints in order to convexify the objective function. This relaxation can be formulated as:

$$\min_{x \in \mathcal{X}} x^T Q_{\alpha_z} x + q_{\alpha_z}^T x + k_{\alpha_z}\tag{32}$$

where $Q_{\alpha_z} = Q + \alpha_z I_n$, $q_{\alpha_z} = q - \alpha_z(l + u)$, $k_{\alpha_z} = \alpha_z l^T u$, and $\alpha_z \geq 0$. As discussed in §3.1, we must select a suitable α_z to ensure that (32) is a convex relaxation of (1). As indicated previously, one such α_z can be determined by using the smallest eigenvalue of the matrix Q . However, as we show in the next proposition, there exists another method for constructing such α_z which makes use of the nullspace of A .

Proposition 6. *Denote by Z an orthonormal basis for the nullspace of the matrix A . Let $\alpha_z \geq -\min(0, \lambda_{\min}(Z^T Q Z))$ in (32). Then, (32) is a convex quadratic program when restricted to the nullspace of the matrix A .*

Proof. Let $\mathcal{H} = \{x \in \mathbb{R}^n \mid Ax = b\}$, and denote by r the rank of A . It is clear than any point satisfying $Ax = b$ can be expressed as $x = x_h + Zx_z$, where $x_h \in \mathcal{H}$, $x_z \in \mathbb{R}^{n-r}$, and $Z \in \mathbb{R}^{n \times n-r}$. By using this transformation, we can write (32) as:

$$\begin{aligned} \min_{x_z} \quad & (x_h + Zx_z)^T Q_{\alpha_z} (x_h + Zx_z) + q_{\alpha_z}^T (x_h + Zx_z) + k_{\alpha_z} \\ \text{s.t.} \quad & C(x_h + Zx_z) \leq d \\ & l \leq (x_h + Zx_z) \leq u. \end{aligned} \quad (33)$$

It is easily verified that (33) is convex for all $\alpha_z \geq -\min(0, \lambda_{\min}(Z^T Q Z))$. \square

From Proposition 6, it follows that the tightest relaxation of the form (32) is obtained by setting $\alpha_z = -\min(0, \lambda_{\min}(Z^T Q Z))$. We refer to this convex relaxation of (1) as the *eigenvalue relaxation in the nullspace of A* . In the following proposition, we show that this relaxation is at least as tight as the generalized eigenvalue relaxation.

Proposition 7. *Assume that $\alpha_g = -\min(0, \lambda_{\min}(Q, I_n + A^T A))$ in (24) and $\alpha_z = -\min(0, \lambda_{\min}(Z^T Q Z))$ in (32). Let μ_{GEIG} and μ_{EIGZ} denote the optimal objective function values in (24) and (32), respectively. Then, $\mu_{\text{EIGZ}} \geq \mu_{\text{GEIG}}$.*

Proof. To prove that $\mu_{\text{EIGZ}} \geq \mu_{\text{GEIG}}$, it suffices to show that $\alpha_z \leq \alpha_g$. Similar to (14), the smallest eigenvalue of $Z^T Q Z$ can be expressed as:

$$\lambda_{\min}(Z^T Q Z) = \min_{x \neq 0, Ax=0} \frac{x^T Q x}{x^T x} = \min_{x \neq 0, Ax=0} \frac{x^T Q x}{x^T (I_n + A^T A) x} \quad (34)$$

where for the second equality we used the fact that the minimization is over vectors x that lie in the null space of A . The restriction of vectors x to the null space of A also implies that $\lambda_{\min}(Z^T Q Z) \geq \lambda_{\min}(Q, I_n + A^T A)$. This is easily seen by noting that the Rayleigh quotient expression for the generalized eigenvalue of the pair $(Q, I_n + A^T A)$ in (14) is over a larger domain. Hence, $\alpha_z \leq \alpha_g$, and $\mu_{\text{EIGZ}} \geq \mu_{\text{GEIG}}$. \square

From Proposition 7, it follows that the eigenvalue relaxation in the nullspace of A can be potentially tighter than the generalized eigenvalue relaxation. However, the computation of Z can be computationally expensive. Therefore, an important question is whether we can obtain a good approximation of $\lambda_{\min}(Z^T Q Z)$ without having to explicitly compute Z . This question is addressed by the following proposition.

Proposition 8. *Let δ be a real scalar. Then, the following hold:*

- (i) *If the matrix Q is indefinite, $\lambda_{\min}(Q, I_n + \delta A^T A)$ is a strictly increasing function of δ for $\delta \geq 1$.*

$$(ii) \lim_{\delta \rightarrow \infty} \lambda_{\min}(Q, I_n + \delta A^T A) = \min(0, \lambda_{\min}(Z^T Q Z)).$$

Proof. We start with the proof of (i). Let $\delta_1, \delta_2 \in \mathbb{R}$ be two scalars such that $\delta_2 > \delta_1 \geq 1$. Define the set $D = \{x \in \mathbb{R}^n : x \neq 0, x^T Q x < 0\}$. Since the matrix Q is indefinite by assumption, it is clear that $D \neq \emptyset$. From the definition of the set D , it is easy to check that the following inequality holds:

$$\frac{x^T Q x}{x^T (I_n + \delta_2 A^T A) x} > \frac{x^T Q x}{x^T (I_n + \delta_1 A^T A) x}, \quad \forall x \in D \quad (35)$$

Using the definition of the Rayleigh quotient in (14), $D \neq \emptyset$ and (35), it is simple to verify that $\lambda_{\min}(Q, I_n + \delta_2 A^T A) > \lambda_{\min}(Q, I_n + \delta_1 A^T A)$ which proves (i).

To prove (ii), consider the Rayleigh quotient definition in (14) for $(Q, I_n + \delta A^T A)$. Let $x = y + z$, where $y, z \in \mathbb{R}^n$ are orthogonal vectors which belong to the row space and nullspace of A , respectively. Then, by using this transformation in (14), we have:

$$\lim_{\delta \rightarrow \infty} \lambda_{\min}(Q, I_n + \delta A^T A) = \lim_{\delta \rightarrow \infty} \min_{(y+z) \neq 0} \frac{(y+z)^T Q (y+z)}{(y+z)^T (y+z) + \delta y^T A^T A y}. \quad (36)$$

To determine the limit in (36), we consider the following cases:

(i) $y \neq 0$. In this case, we obtain:

$$\min_{(y+z) \neq 0} \lim_{\delta \rightarrow \infty} \frac{(y+z)^T Q (y+z)}{(y+z)^T (y+z) + \delta y^T A^T A y} = 0. \quad (37)$$

(ii) $y = 0$. In this case, (36) reduces to:

$$\lim_{\delta \rightarrow \infty} \min_{z \neq 0} \frac{z^T Q z}{z^T z} = \lim_{\delta \rightarrow \infty} \min_{z \neq 0, Az=0} \frac{z^T Q z}{z^T z} = \lambda_{\min}(Z^T Q Z). \quad (38)$$

Then, it follows that $\lim_{\delta \rightarrow \infty} \lambda_{\min}(Q, I_n + \delta A^T A) = \min(0, \lambda_{\min}(Z^T Q Z))$. \square

Proposition 8 has very important consequences since it suggests we can approximate the bound given by the eigenvalue relaxation in the nullspace of A by solving the following quadratic program for a sufficiently large value of δ :

$$\min_{x \in \mathcal{X}} x^T Q x + q^T x + \alpha(\delta)(x^T x - (l+u)^T x + l^T u) + \alpha(\delta) \cdot \delta \cdot \|Ax - b\|^2 \quad (39)$$

where $\alpha(\delta) = -\lambda_{\min}(Q, I_n + \delta A^T A)$. Note that, for $\delta = 1$, (39) corresponds to the generalized eigenvalue relaxation introduced in §3.2.

Since $\lambda_{\min}(Q, I_n + \delta A^T A)$ is a strictly increasing function of δ for $\delta \geq 1$, Proposition 8 implies that as δ is increased, $\alpha(\delta)$ will converge to either 0 or $-\lambda_{\min}(Z^T Q Z)$. The case in which $\alpha(\delta)$ converges to 0 is particularly interesting since it indicates that $\lambda_{\min}(Z^T Q Z) \geq 0$, and the continuous relaxation of (1) is convex when restricted to the nullspace of A . Note that $\lambda_{\min}(Q) < 0$ does not necessarily imply that $\lambda_{\min}(Z^T Q Z) < 0$, and as a result, the continuous relaxation of (1) may be convex when restricted to the nullspace of A , even if it is nonconvex in the space of the original problem variables.

The quadratic term $\alpha(\delta) \cdot \delta \cdot \|Ax - b\|^2$ vanishes for any x feasible in (39). This term is included in the objective function of (39) to ensure that $Q + \alpha(\delta)(I_n + \delta A^T A)$ is positive

semidefinite. However, this term need not be included for (39) to be convex. Proposition 6 implies that (39) is convex $\alpha(\delta) \geq -\min(0, \lambda_{\min}(Z^T Q Z))$. The definition of $\alpha(\delta)$ and Proposition 8 imply that $\alpha(\delta) \geq -\min(0, \lambda_{\min}(Z^T Q Z))$ holds for any $\delta \geq 1$. As a result, the quadratic term $\alpha(\delta) \cdot \delta \cdot \|Ax - b\|^2$ can be dropped from the objective function of (39), which simplifies this relaxation to:

$$\min_{x \in \mathcal{X}} \quad x^T Q x + q^T x + \alpha(\delta)(x^T x - (l + u)^T x + l^T u) \quad (40)$$

This simplification has two significant advantages. First, it allows us to preserve the sparsity pattern defined by Q . Second, it prevents the relaxation from becoming ill-conditioned since δ does not figure in the objective function of (40) and is only used to determine $\alpha(\delta)$. We can use a simple iterative procedure to determine a value of δ which leads to a good approximation of the bound given by the eigenvalue relaxation in the nullspace of A . We detail such procedure in §5.

By considering a quadratic relaxation of the form (40), there is no need to project onto the nullspace of A . This is particularly advantageous in the context of the branching variable selection rules that we introduce in §4, since the branching decisions are easier to interpret in the space of the original problem variables.

We finish this section by showing that the eigenvalue relaxation in the null space of A is equivalent to the following SDP:

$$\min_{x \in \mathcal{X}, X} \quad \langle Q, X \rangle + q^T x \quad (41a)$$

$$\text{s.t. } X - xx^T \succeq 0 \quad (41b)$$

$$\langle I_n, X \rangle - (l + u)^T x + l^T u \leq 0 \quad (41c)$$

$$\langle A^T A, X \rangle - (2A^T b)^T x + b^T b = 0 \quad (41d)$$

Proposition 9. *Suppose that the matrix $Z^T Q Z$ is indefinite. Let $\mu_{\text{EIGZ}}(\alpha_z)$ and $\mu_{\text{SDP_EIGZ}}$ be the optimal objective function values in (32) and (41), respectively. Then, the following hold:*

(i) *If $\alpha_z \geq -\lambda_{\min}(Z^T Q Z)$ in (32), then $\mu_{\text{EIGZ}}(\alpha_z) \leq \mu_{\text{SDP_EIGZ}}$.*

(ii) *If $\alpha_z = -\lambda_{\min}(Z^T Q Z)$ in (32), then $\mu_{\text{EIGZ}}(\alpha_z) = \mu_{\text{SDP_EIGZ}}$.*

Proof. Unlike (19) and (28), (41) does not admit a strictly feasible solution. To illustrate this, note that, for any x satisfying $Ax = b$, (41d) can be written as:

$$\langle A^T A, X \rangle - (2A^T b)^T x + b^T b + \langle A^T A, xx^T \rangle - \langle A^T A, xx^T \rangle = 0 \quad (42a)$$

$$\implies \langle A^T A, X - xx^T \rangle = 0 \quad (42b)$$

which implies that $X - xx^T$ cannot be positive definite for the pairs (x, X) that are feasible in (41). It follows that we cannot apply the strong duality theorem to (41). As a result, the proof of this proposition relies on different arguments from those used in the proofs of Propositions 2 and 5. We proceed in two steps:

- (a) We show that the optimal objective function value of the dual problem of (41) provides an upper bound on (32) when $\alpha_z \geq -\lambda_{\min}(Z^T Q Z)$, and is equal to $\mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. By weak duality, this implies that $\mu_{\text{SDP_EIGZ}} \geq \mu_{\text{EIGZ}}(\alpha_z)$ for $\alpha_z \geq -\lambda_{\min}(Z^T Q Z)$, proving the claim in (i).

(b) We construct a feasible solution for (41) which attains the same objective function value as an optimal solution of (32) when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. This implies that $\mu_{\text{SDP_EIGZ}} \leq \mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. This observation combined with the result from (a) completes the proof of the claim in (ii).

To prove (a), we use Proposition 1 to write the dual of (41) as:

$$\max_{\hat{\alpha}_z \in \mathbb{R}_{\geq 0}, \hat{\beta}_z \in \mathbb{R}: Q_{\hat{\alpha}_z, \hat{\beta}_z} \succ 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\hat{\alpha}_z, \hat{\beta}_z} x + q_{\hat{\alpha}_z, \hat{\beta}_z}^T x + k_{\hat{\alpha}_z, \hat{\beta}_z} \right\} \quad (43)$$

where $\hat{\alpha}_z$ and $\hat{\beta}_z$ are multipliers for (41c) and (41d), respectively, $Q_{\hat{\alpha}_z, \hat{\beta}_z} = Q + \hat{\alpha}_z I_n + \hat{\beta}_z A^T A$, $q_{\hat{\alpha}_z, \hat{\beta}_z} = q - \hat{\alpha}_z(l+u) - 2\hat{\beta}_z A^T b$, and $k_{\hat{\alpha}_z, \hat{\beta}_z} = \hat{\alpha}_z l^T u + \hat{\beta}_z b^T b$. Let $\hat{\delta}_z = \hat{\beta}_z / \hat{\alpha}_z$. By substituting $\hat{\beta}_z = \hat{\delta}_z \hat{\alpha}_z$ in (43), the dual becomes:

$$\max_{\hat{\alpha}_z \in \mathbb{R}_{\geq 0}, \hat{\delta}_z \in \mathbb{R}: Q_{\hat{\alpha}_z, \hat{\delta}_z} \succ 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\hat{\alpha}_z, \hat{\delta}_z} x + q_{\hat{\alpha}_z, \hat{\delta}_z}^T x + k_{\hat{\alpha}_z, \hat{\delta}_z} \right\} \quad (44)$$

where $Q_{\hat{\alpha}_z, \hat{\delta}_z} = Q + \hat{\alpha}_z(I_n + \hat{\delta}_z A^T A)$, $q_{\hat{\alpha}_z, \hat{\delta}_z} = q - \hat{\alpha}_z(l+u + 2\hat{\delta}_z A^T b)$, and $k_{\hat{\alpha}_z, \hat{\delta}_z} = \hat{\alpha}_z(l^T u + \hat{\delta}_z b^T b)$. Note that the quadratic term $\hat{\alpha}_z \hat{\delta}_z \|Ax - b\|^2$ vanishes for any x feasible in the inner minimization problem. As a result, (44) can be posed as:

$$\max_{\hat{\alpha}_z \in \mathbb{R}_{\geq 0}, \hat{\delta}_z \in \mathbb{R}: Q_{\hat{\alpha}_z, \hat{\delta}_z} \succ 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\hat{\alpha}_z} x + q_{\hat{\alpha}_z}^T x + k_{\hat{\alpha}_z} \right\} \quad (45)$$

where $Q_{\hat{\alpha}_z} = Q + \hat{\alpha}_z I_n$, $q_{\hat{\alpha}_z} = q - \hat{\alpha}_z(l+u)$, and $k_{\hat{\alpha}_z} = \hat{\alpha}_z l^T u$. Since $Z^T Q Z$ is indefinite, Q is indefinite as well. Proposition 6 implies that, for a given value of $\hat{\delta}_z$, $Q_{\hat{\alpha}_z, \hat{\delta}_z} \succ 0$ when $\hat{\alpha}_z \geq -\lambda_{\min}(Q, I_n + \hat{\delta}_z A^T A)$. For all such $\hat{\alpha}_z, \hat{\delta}_z$ it is simple to check that the objective function of the inner minimization problem of (45) is monotonically decreasing in $\hat{\alpha}_z$. By using the fact that $\lambda_{\min}(Z^T Q Z) < 0$ and Proposition 8, it is easy to verify that the maximum of (45) is attained when $\hat{\alpha}_z = -\lim_{\hat{\delta}_z \rightarrow \infty} \lambda_{\min}(Q, I_n + \hat{\delta}_z A^T A) = -\lambda_{\min}(Z^T Q Z)$. Let $\mu_{\text{DSDP_EIGZ}}$ be the optimal objective function value in (45). Since the inner problem in (45) has the same form as the QP (32), it is simple to check that $\mu_{\text{DSDP_EIGZ}} \geq \mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z \geq -\lambda_{\min}(Z^T Q Z)$, and $\mu_{\text{DSDP_EIGZ}} = \mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. By weak duality, it follows that $\mu_{\text{SDP_EIGZ}} \geq \mu_{\text{EIGZ}}(\alpha_z)$ for $\alpha_z \geq -\lambda_{\min}(Z^T Q Z)$. This proves the claim in (i).

Next, we prove (b). Let \hat{x} denote an optimal solution of (32) when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. Define $\hat{X} = \hat{x}\hat{x}^T + \gamma Zv(Zv)^T$, where $\gamma = (l+u)^T \hat{x} - l^T u - \hat{x}^T \hat{x}$, and v denotes the eigenvector corresponding to the smallest eigenvalue of $Z^T Q Z$. We first show that (\hat{x}, \hat{X}) is feasible in (41). By definition, $\hat{x} \in \mathcal{X}$. Consider (41b). Recall that the concave envelope of x_i^2 over $[l_i, u_i]$ is given by $(l_i + u_i)x_i - l_i u_i$. As a result, it is clear that each term $(l_i + u_i)\hat{x}_i - l_i u_i - \hat{x}_i^2$ is nonnegative, which in turn implies that $\gamma \geq 0$. Moreover, since $Zv(Zv)^T \succ 0$, it follows that $\hat{X} - \hat{x}\hat{x}^T \succ 0$.

Consider (41c) and (41d). Substituting (\hat{x}, \hat{X}) in (41c), we obtain:

$$\begin{aligned} & \langle I_n, \hat{x}\hat{x}^T + \gamma Zv(Zv)^T \rangle - (l+u)^T \hat{x} + l^T u \\ &= \hat{x}^T \hat{x} + \gamma v^T Z^T Z v - (l+u)^T \hat{x} + l^T u = \hat{x}^T \hat{x} + \gamma - (l+u)^T \hat{x} + l^T u = 0. \end{aligned}$$

Similarly, substituting (\hat{x}, \hat{X}) in (41d) yields:

$$\begin{aligned} & \langle A^T A, \hat{x}\hat{x}^T + \gamma Zv(Zv)^T \rangle - (2A^T b)^T \hat{x} + b^T b \\ &= \hat{x}^T A^T A \hat{x} - (2A^T b)^T \hat{x} + b^T b + \gamma v^T Z^T A^T A Z v = (A\hat{x} - b)^T (A\hat{x} - b) = 0. \end{aligned}$$

Let $f(x, X)$ be the objective function of (41). The value of f at (\hat{x}, \hat{X}) is:

$$\begin{aligned} f(\hat{x}, \hat{X}) &= \langle Q, \hat{x}\hat{x}^T + \gamma Zv(Zv)^T \rangle + q^T \hat{x} \\ &= \hat{x}^T Q \hat{x} - \gamma \alpha_z + q^T \hat{x} \\ &= \hat{x}^T (Q + \alpha_z I_n) \hat{x} + (q - \alpha_z(l + u))^T \hat{x} + \alpha_z l^T u = \mu_{\text{EIGZ}} \end{aligned} \quad (46)$$

where we have used the fact that $v^T Z^T Q Z v = -\alpha_z = \lambda_{\min}(Z^T Q Z)$. Since (\hat{x}, \hat{X}) is feasible in (41), from (46) it follows that $\mu_{\text{SDP_EIGZ}} \leq \mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. This observation combined with the result from (a) proves the claim in (ii), by showing that $\mu_{\text{SDP_EIGZ}} = \mu_{\text{EIGZ}}(\alpha_z)$ when $\alpha_z = -\lambda_{\min}(Z^T Q Z)$. \square

3.4 Further insights into the proposed quadratic relaxations

The relaxations introduced in §3.1–3.3 can be derived through the following four-step recipe:

(R1) identify a (possibly empty) set \mathcal{J} of quadratic functions of the form $f_j(x) = x^T S_j x + s_j^T x + \eta_j$, where $S_j \in \mathbb{S}^n, s_j \in \mathbb{R}^n, \eta_j \in \mathbb{R}$, such that $f_j(x) = 0$ for $x \in \Omega := \{x \in \mathbb{R}^n \mid Ax = b\}$;

(R2) construct an initial relaxation for (1) as

$$\min_{x \in \mathcal{X}} \quad x^T Q x + q^T x + \alpha(x^T x - (l + u)^T x + l^T u) + \sum_{j \in \mathcal{J}} \beta_j f_j(x) \quad (47)$$

where $\alpha \in \mathbb{R}_{\geq 0}, \beta_j \in \mathbb{R}$, such that $Q + \alpha I_n + \sum_{j \in \mathcal{J}} \beta_j S_j \succcurlyeq 0$;

(R3) find α^*, β^* such that the bound given by the relaxation (47) is maximized

$$(\alpha^*, \beta^*) = \arg \max_{\alpha \in \mathbb{R}_{\geq 0}, \beta \in \mathbb{R}^{|\mathcal{J}|}: Q_{\alpha, \beta} \succcurlyeq 0} \left\{ \min_{x \in \mathcal{X}} x^T Q_{\alpha} x + q_{\alpha}^T x + k_{\alpha} \right\} \quad (48)$$

where $Q_{\alpha} = Q + \alpha I_n$, $Q_{\alpha, \beta} = Q_{\alpha} + \sum_{j \in \mathcal{J}} \beta_j S_j$, $q_{\alpha} = q - \alpha(l + u)$, $k_{\alpha} = \alpha l^T u$, and β is the $|\mathcal{J}|$ -dimensional vector whose entries are the parameters β_j ;

(R4) obtain the relaxation

$$\min_{x \in \mathcal{X}} \quad x^T Q_{\alpha^*} x + q_{\alpha^*}^T x + k_{\alpha^*}. \quad (49)$$

Observe that the parameters β_j are not present in the objective function of the inner minimization problem in (48) and the objective function in (49) since $f_j(x) = 0$ for $x \in \mathcal{X} \subset \Omega$ (due to (R1)). The three spectral relaxations presented in §3.1–3.3 can be identified with (48) by noting that:

- $\mathcal{J} = \emptyset$, $\alpha^* = -\min(0, \lambda_{\min}(Q))$ for the eigenvalue relaxation (18);
- $\mathcal{J} = \{1\}$, $f_1(x) = \sum_{i=1}^m (A_i x - b_i)^2$, $\alpha^* = -\min(0, \lambda_{\min}(Q, I_n + A^T A))$, $\beta_1^* = \alpha^*$ for the generalized eigenvalue relaxation (24). Note that in this case a further restriction that $\beta_1 = \alpha$ is imposed in (48); and
- $\mathcal{J} = \{1\}$, $f_1(x) = \sum_{i=1}^m (A_i x - b_i)^2$, $\alpha^* = -\min(0, \lambda_{\min}(Z^T Q Z))$ and $\beta_1^* = +\infty$ for the eigenvalue relaxation on the nullspace of A (32).

From Propositions 4 and 7 we know that the lower bound obtained from the eigenvalue relaxation in the nullspace of A (32) is at least as large as those provided by the other spectral relaxations. Further, the computation of α^* can be done efficiently.

The recipe (R1)-(R4) is preferable from a computational standpoint since the resulting relaxation is a quadratic program inheriting the sparsity of the problem. However, the step (R1) allows for other choice for the functions $f_j(x)$ that have been considered in the literature (see Faye and Roupin [13]). Some examples for the functions satisfying (R1) are [13]: $(x_j(A_i x - b_i))$, $((A_j x - b_j)(A_i x - b_i))$, $(x^T A_j^T A_i x - b_j b_i)$. This naturally raises the question: *Can we improve on the bound provided by (32) when restricted to the class of relaxations in (48)?* In the rest of the section, we show that we cannot improve on the bound provided by the eigenvalue relaxation on the nullspace of A (32). Thus, establishing that (32) is the best among the class of relaxations in (48).

We begin by recalling the properties of functions satisfying (R1).

Proposition 10. *Let $f(x) = x^T S x + s^T x + \eta$ be a quadratic function. Then, $f(x) = 0$ for all $x \in \Omega := \{x \in \mathbb{R}^n \mid Ax = b\}$ if and only if $S = A^T W^T + W A$, $s = A^T \nu - 2W b$, $\eta = -b^T \nu$ for some $W \in \mathbb{R}^{n \times m}$ and $\nu \in \mathbb{R}^m$.*

Proof. This follows from Theorem 1 in [13]. \square

Following Proposition 10, we assume without loss of generality that $S_j = A^T W_j^T + W_j A$ for some $W_j \in \mathbb{R}^{n \times m}$ in the rest of this section. We will compare the relaxations in the class (48) with the eigenvalue relaxation in the nullspace of A (32) through the respective SDP formulations. To this end, consider the SDP:

$$\min_{x \in \mathcal{X}, X} \langle Q, X \rangle + q^T x \quad (50a)$$

$$\text{s.t. } X - x x^T \succeq 0 \quad (50b)$$

$$\langle I_n, X \rangle - (l + u)^T x + l^T u \leq 0 \quad (50c)$$

$$\langle S_j, X \rangle + s_j^T x + \eta_j = 0, \quad j \in \mathcal{J}. \quad (50d)$$

The next proposition shows that SDP (50) is the dual of (48).

Proposition 11. *Let $\mathcal{J} \neq \emptyset$ be a set of quadratic functions satisfying (R1). The dual of the SDP (50) is given by (48).*

Proof. By dualizing the constraints (50c) and (50d) with the multipliers $\alpha \in \mathbb{R}_{\geq 0}$ and $\beta_j \in \mathbb{R}, j \in \mathcal{J}$, respectively, we can use Proposition 1 to obtain the claim. \square

The next result shows that the feasible set of the SDP (41) is in general a subset of the feasible set of the SDP (50). Further, we provide conditions on the choice of quadratic functions in \mathcal{J} so that equality holds.

Proposition 12. *Let $\mathcal{F}_{\text{SDP_EIGZ}}$ and $\mathcal{F}_{\text{SDP_EIGJ}}$ denote the feasible regions of the SDPs in (41) and (50), respectively. Then, the following holds:*

$$(i) \quad \mathcal{F}_{\text{SDP_EIGZ}} \subseteq \mathcal{F}_{\text{SDP_EIGJ}}.$$

$$(ii) \quad \text{If } \exists \omega_j, j \in \mathcal{J} \text{ such that } \sum_{j \in \mathcal{J}} \omega_j W_j = A^T \text{ then } \mathcal{F}_{\text{SDP_EIGZ}} = \mathcal{F}_{\text{SDP_EIGJ}}.$$

Proof. We start by proving (i). Recall from (42) that any $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGZ}}$ satisfies $\langle A^T A, \bar{X} - \bar{x}\bar{x}^T \rangle = 0$. Hence, \bar{X} takes the form $\bar{X} = \bar{x}\bar{x}^T + ZVZ^T$ for all $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGZ}}$, where $Z \in \mathbb{R}^{n \times n-r}$ is a basis for the null space of A and $V \in \mathbb{S}^{n-r}$. For any $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGZ}}$ it follows that for all $j \in \mathcal{J}$:

$$\langle S_j, \bar{X} \rangle + s_j^T \bar{x} + \eta_j \quad (51a)$$

$$= \langle S_j, \bar{X} - \bar{x}\bar{x}^T \rangle + \bar{x}^T S_j \bar{x} + s_j^T \bar{x} + \eta_j \quad (51b)$$

$$= \langle S_j, \bar{X} - \bar{x}\bar{x}^T \rangle = \langle A^T W_j^T + W_j A, ZVZ^T \rangle = 0 \quad (51c)$$

where (51b) follows from adding and subtracting $\bar{x}^T S_j \bar{x}$, the first equality in (51c) follows from (R1), the second equality in (51c) from Proposition 10 and the final equality due to Z being a basis for the nullspace of A . Thus $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGJ}}$ proving the claim in (i).

Consider the claim in (ii). Suppose that there exist $\omega_j, j \in \mathcal{J}$ such that the condition in (ii) holds. We perform a linear combination of the inequalities in (50d) using ω_j to obtain for any $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGJ}}$:

$$0 = \sum_{j \in \mathcal{J}} \omega_j (\langle S_j, \bar{X} \rangle + s_j^T \bar{x} + \eta_j) \quad (52a)$$

$$= \sum_{j \in \mathcal{J}} \omega_j (\langle S_j, \bar{X} - \bar{x}\bar{x}^T \rangle + \bar{x}^T S_j \bar{x} + s_j^T \bar{x} + \eta_j) \quad (52b)$$

$$= \sum_{j \in \mathcal{J}} \omega_j \langle S_j, \bar{X} - \bar{x}\bar{x}^T \rangle = 2 \langle A^T A, \bar{X} - \bar{x}\bar{x}^T \rangle \quad (52c)$$

where (52b) follows from adding and subtracting $\bar{x}^T S_j \bar{x}$, the first equality in (52c) follows from (R1), the second equality in (52c) from Proposition 10 and the condition in (ii). Thus $(\bar{x}, \bar{X}) \in \mathcal{F}_{\text{SDP_EIGZ}}$ proving the claim in (ii). \square

Faye and Roupin [13] proved the equivalence between the SDP (41) and a similar SDP where (41d) is replaced by the constraints derived by lifting the quadratic functions $x_j(A_i x - b_i) = 0$, $i = 1, \dots, m$, $j = 1, \dots, n$ to the space of (x, X) . Proposition 12 considerably expands the set of quadratic functions for which the feasible set of the resulting SDP is equal to $\mathcal{F}_{\text{SDP_EIGZ}}$ (claim in (ii)). It is easy to verify that all of the examples of quadratic functions satisfying (R1) described in [13] do satisfy the condition in (ii). Further, the claim in (i) shows that there exist no quadratic functions satisfying (R1) for which the resulting SDP can have a smaller feasible region than the SDP (41). This brings us to the main result on the claim that the relaxation (32) is indeed the best among the class of relaxations in (48).

Theorem 1. *Suppose that $Z^T Q Z$ is indefinite and that the set \mathcal{J} is chosen such that (R1) holds. Assume that $\alpha_z = -\lambda_{\min}(Z^T Q Z)$ in (32). Denote by μ_{EIGZ} and μ_{EIGJ} the optimal objective function values in (32) and (48), respectively. Then, $\mu_{\text{EIGJ}} \leq \mu_{\text{EIGZ}}$.*

Proof. Let $\mu_{\text{SDP_EIGZ}}$ and $\mu_{\text{SDP_EIGJ}}$ be the optimal objective values of (41) and (50), respectively. By Proposition 9(ii) we have that $\mu_{\text{EIGZ}} = \mu_{\text{SDP_EIGZ}}$. By Proposition (12)(i) we have that $\mu_{\text{SDP_EIGJ}} \leq \mu_{\text{SDP_EIGZ}}$. By Proposition 11 and weak duality we have that $\mu_{\text{EIGJ}} \leq \mu_{\text{SDP_EIGJ}}$. Hence, $\mu_{\text{EIGJ}} \leq \mu_{\text{EIGZ}}$. \square

We finish this section by providing a theoretical comparison between the spectral relaxations studied in §3.1–3.3 and some SDP relaxations described in §2.2.

Theorem 2. Assume that $Z^T Q Z$ is indefinite. Suppose that $\alpha_e = -\lambda_{\min}(Q)$ in (18), $\alpha_g = -\lambda_{\min}(Q, I_n + A^T A)$ in (24), and $\alpha_z = -\lambda_{\min}(Z^T Q Z)$ in (32). Denote by μ_{EIG} , μ_{GEIG} , μ_{EIGZ} , μ_{SDP_d} , $\mu_{\text{SDP}_{\text{dax}}}$, and $\mu_{\text{SDP}_{\text{da}}}$ the optimal objective function values of (18), (24), (32), (5), (6), and (7), respectively. Then, the following hold:

(i) $\mu_{\text{SDP}_d} \geq \mu_{\text{EIG}}$.

(ii) $\mu_{\text{SDP}_{\text{dax}}} = \mu_{\text{SDP}_{\text{da}}} \geq \mu_{\text{EIGZ}} \geq \mu_{\text{GEIG}} \geq \mu_{\text{EIG}}$.

Proof. We start by proving (i). Denote by $\mu_{\text{SDP}_{\text{EIG}}}$ the optimal objective function value in (19). By Proposition 2(ii), we have that $\mu_{\text{EIG}} = \mu_{\text{SDP}_{\text{EIG}}}$. Hence, we can prove (i) by comparing the SDPs (5) and (19). The constraints (5c) and (5d) are implied by (5b), and as a result, can be dropped from (5). Therefore, (5) and (19) only differ in the constraints (5e) and (19c). It is simple to verify that the inequality (19c) can be obtained by aggregating the McCormick inequalities (5e), which implies that $\mu_{\text{SDP}_d} \geq \mu_{\text{SDP}_{\text{EIG}}} = \mu_{\text{EIG}}$.

Now, we prove (ii). As stated in §2.3.1, the relationship $\mu_{\text{SDP}_{\text{dax}}} = \mu_{\text{SDP}_{\text{da}}}$ follows from a result given in [13]. To show that $\mu_{\text{SDP}_{\text{da}}} \geq \mu_{\text{EIGZ}}$, we follow the same line of arguments used for proving (i). Let $\mu_{\text{SDP}_{\text{EIGZ}}}$ be the optimal objective function value in (41). Proposition 9(ii) implies that $\mu_{\text{EIGZ}} = \mu_{\text{SDP}_{\text{EIGZ}}}$. Therefore, to prove (ii), we can simply compare the SDPs (7) and (41). The constraints (5c) and (5d) are redundant in (7), and can be dropped from this formulation. Similar to the previous case, (7) and (41) only differ in the constraints (5e) and (41c). As stated above, the inequality (41c) is implied by the inequalities (5e). Hence, $\mu_{\text{SDP}_{\text{da}}} \geq \mu_{\text{SDP}_{\text{EIGZ}}} = \mu_{\text{EIGZ}}$. The inequalities $\mu_{\text{GEIG}} \geq \mu_{\text{EIG}}$ and $\mu_{\text{EIGZ}} \geq \mu_{\text{GEIG}}$, follow from Propositions 4 and 7, respectively. \square

4 Spectral branching for nonconvex binary QPs

In this section, we introduce new eigenvalue-based branching variable selection strategies for nonconvex binary QPs. These strategies are inspired by the strong branching rule which was initially proposed for mixed-integer linear programs [4], and can be used along with the quadratic relaxations discussed in §3.1–3.3. For simplicity, we only describe our branching strategies for the eigenvalue relaxation, which rely on the smallest eigenvalue of Q and its associated eigenvector. The branching rules for the quadratic relaxations described in §3.2 and 3.3 are similar, but they make use of the smallest generalized eigenvalue of the pair $(Q, I + \delta A^T A)$ and its corresponding eigenvector.

We first introduce some notation. Let \mathcal{F} be the set of indices of the variables that are fixed at the current node. Denote by $\mathcal{B} = \{1, \dots, n\} \setminus \mathcal{F}$ the set of branching candidates. Let \bar{Q} be the $\mathbb{R}^{|\mathcal{B}| \times |\mathcal{B}|}$ sub-matrix of Q obtained by eliminating the rows and columns corresponding to the variables in \mathcal{F} . Define the bijection $\sigma : \mathcal{B} \rightarrow \{1, \dots, |\mathcal{B}|\}$, which maps $i \in \mathcal{B}$ to the $\sigma(i)$ -th row and $\sigma(i)$ -th column of \bar{Q} .

Assume that we branch on variable x_i , $i \in \mathcal{B}$ by creating two nodes, one where $x_i = 0$ and another where $x_i = 1$. At these descendant nodes, the eigenvalue relaxation is constructed by considering the smallest eigenvalue of the submatrix obtained by eliminating the $\sigma(i)$ -th row and $\sigma(i)$ -th column of \bar{Q} . We denote this submatrix by \hat{Q} . In this context, a potentially good branching rule may consist in branching on the variable which leads to the largest increase in the smallest eigenvalue of \hat{Q} . Note that, at a given node of the branch-and-bound tree, this rule requires the solution of $|\mathcal{B}|$ eigenvalue problems, each one involving a submatrix of \bar{Q} obtained by eliminating the row and column corresponding to a

particular index $i \in \mathcal{B}$. We call this rule *spectral branching with complete enumeration*. The index corresponding to this branching rule, denoted as $i_{\text{exact}} \in C$, can be mathematically expressed as:

$$i_{\text{exact}} = \arg \max_{i \in \mathcal{B}} \lambda_{\min} \left(P_{\sigma(i)} \bar{Q} P_{\sigma(i)}^T \right) \quad (53)$$

where $P_{\sigma(i)}$ is a $(|\mathcal{B}|-1) \times |\mathcal{B}|$ matrix obtained by removing the $\sigma(i)$ -th row from the $|\mathcal{B}| \times |\mathcal{B}|$ identity matrix. Note that $\hat{Q} = P_{\sigma(i)} \bar{Q} P_{\sigma(i)}^T$ results in a matrix where the $\sigma(i)$ -th row and $\sigma(i)$ -th column of \bar{Q} are removed.

The computational complexity of complete enumeration is $\Omega(|\mathcal{B}|^3)$. We are not aware of any efficient approach for obtaining i_{exact} that avoids complete enumeration. We instead rely on a lower bound for $\lambda_{\min}(\cdot)$ that will be obtained without computing an eigenvalue and is computationally inexpensive. Gershgorin's Circle Theorem (GCT) [15] provides such a lower bound estimate. The GCT states that: *every eigenvalue of a $t \times t$ matrix T lies in one of the circles $\mathcal{C}_k(T) = \{\lambda : |\lambda - T_{kk}| \leq \sum_{l \neq k} |T_{kl}|\}$ for $k = 1, \dots, t$* . A lower bound estimate for the smallest eigenvalue of the matrix T based on the GCT, denoted as $\lambda_{\min}^{\text{GCT}}(T)$ is:

$$\lambda_{\min}^{\text{GCT}}(T) = \min_{k \in \{1, \dots, t\}} \left(T_{kk} - \sum_{l \neq k} |T_{kl}| \right) \quad (54)$$

Using the GCT-based lower bound estimate we can then define a branching variable index as:

$$i_{\text{GCT}} = \arg \max_{i \in \mathcal{B}} \lambda_{\min}^{\text{GCT}} \left(P_{\sigma(i)} \bar{Q} P_{\sigma(i)}^T \right). \quad (55)$$

Note that the index i_{GCT} can be determined without having to compute the matrix $P_{\sigma(i)} \bar{Q} P_{\sigma(i)}^T$. This approach has a computational complexity of $O(|\mathcal{B}|^2)$ and is computationally inexpensive compared to complete enumeration.

The choice of i_{GCT} can be viewed as a pessimistic estimate since it is obtained by maximizing the worst-case bound for the smallest eigenvalue. Instead, we employ a different approach to determine the branching variable. Let v be the eigenvector corresponding to the smallest eigenvalue of \bar{Q} . Then, we select as a branching variable, denoted by i_{approx} , the one which corresponds to the entry of v with the largest absolute value, i.e.

$$i_{\text{approx}} = \arg \max_{i \in \mathcal{B}} |v_{\sigma(i)}| \quad (56)$$

where $v_{\sigma(i)}$ denotes the $\sigma(i)$ -th component of v . We call this rule *approximate spectral branching*. The computational complexity of this rule is $O(|\mathcal{B}|)$.

To appreciate the intuition behind this choice, we recall the proof for the GCT. From the definition of the eigenvalue, we have

$$\begin{aligned} \bar{Q}v &= \lambda_{\min}(\bar{Q})v \\ \implies \sum_{j \in \mathcal{B}} \bar{Q}_{\sigma(i_{\text{approx}})\sigma(j)} v_{\sigma(j)} &= \lambda_{\min}(\bar{Q})v_{\sigma(i_{\text{approx}})} \\ \implies \lambda_{\min}(\bar{Q}) - \bar{Q}_{\sigma(i_{\text{approx}})\sigma(i_{\text{approx}})} &= \sum_{\substack{j \in \mathcal{B}, \\ j \neq i_{\text{approx}}}} \bar{Q}_{\sigma(i_{\text{approx}})\sigma(j)} \frac{v_{\sigma(j)}}{v_{\sigma(i_{\text{approx}})}} \\ \implies |\lambda_{\min}(\bar{Q}) - \bar{Q}_{\sigma(i_{\text{approx}})\sigma(i_{\text{approx}})}| &\leq \sum_{\substack{j \in \mathcal{B}, \\ j \neq i_{\text{approx}}}} |\bar{Q}_{\sigma(i_{\text{approx}})\sigma(j)}| \end{aligned}$$

$$\implies \lambda_{\min}(\bar{Q}) \in \mathcal{C}_{\sigma(i_{\text{approx}})}(\bar{Q})$$

where the first implication follows from the $\sigma(i_{\text{approx}})$ -th row of the equality, the second implication is obtained by rearranging and dividing by $v_{\sigma(i_{\text{approx}})}$ and the inequality follows from $|v_{\sigma(j)}/v_{\sigma(i_{\text{approx}})}| \leq 1$ by definition of i_{approx} . In essence, i_{approx} identifies the particular Gershgorin circle that bounds the smallest eigenvalue $\lambda_{\min}(\bar{Q})$. Thus, the choice of i_{approx} as the branching variable can be interpreted as eliminating the particular Gershgorin circle to which $\lambda_{\min}(\bar{Q})$ belongs. In that sense, this can be viewed as an optimistic estimate.

To illustrate the effectiveness of i_{GCT} and i_{approx} in mimicking i_{exact} , we performed some numerical experiments. We generated matrices Q of sizes $n \in \{50, 100\}$ and densities $\rho \in \{0.25, 0.50, 1.00\}$, and computed i_{exact} by complete enumeration. Denote by i_{worst} the index corresponding to the worst choice of branching variable, i.e.:

$$i_{\text{worst}} = \arg \min_{i \in \mathcal{B}} \lambda_{\min} \left(P_{\sigma(i)} Q P_{\sigma(i)}^T \right) \quad (58)$$

Then, the effectiveness of i_x is measured using the metric:

$$\% \text{ gap} = \frac{\lambda_{\min} \left(P_{\sigma(i_x)} Q P_{\sigma(i_x)}^T \right) - \lambda_{\min} \left(P_{\sigma(i_{\text{exact}})} Q P_{\sigma(i_{\text{exact}})}^T \right)}{\lambda_{\min} \left(P_{\sigma(i_{\text{worst}})} Q P_{\sigma(i_{\text{worst}})}^T \right) - \lambda_{\min} \left(P_{\sigma(i_{\text{exact}})} Q P_{\sigma(i_{\text{exact}})}^T \right)} \times 100 \quad (59)$$

where $x \in \{\text{approx}, \text{GCT}\}$. A smaller value of % gap for i_x represents a better approximation of i_{exact} . To obtain a statistic of the effectiveness of these approaches, we generated 100 different instances of Q for each matrix size and density. Figure 1 shows cumulative plots of the percentage of instances for which the % gap is below a certain value. It is evident from the plots that the approximate spectral branching strategy is a better choice than the GCT-based branching rule.




Figure 1: Cumulative plots comparing the effectiveness of the approximate spectral branching and the GCT-based branching strategies.

5 Implementation of the proposed relaxation and branching strategies in BARON

By default, BARON's portfolio of relaxations consists of linear programming (LP), nonlinear programming (NLP) and mixed-integer linear programming (MILP) relaxations [17, 19, 27]. In our implementation, we have expanded this portfolio by adding a new class of convex

QP relaxations. These relaxations are constructed whenever the original model supplied to BARON is of the form (1). We take advantage of BARON’s convexity detector (see [17] for details) in order to determine the type of QP relaxation that will be constructed at a given node in the branch-and-bound tree. If the current node is convex, our QP relaxation is the continuous relaxation of (1) subject to the variable bounds of the current node. On the other hand, if the current node is nonconvex, we construct one of the QP relaxations introduced in §3.1–3.3. The relaxation (40) is selected by default if the original problem contains equality constraints. Otherwise, our QP relaxation constructor automatically switches to the eigenvalue relaxation (18).

To solve the eigenvalue and generalized eigenvalue problems that arise during the construction of the relaxations discussed in §3.1–3.3, we use the subroutines included in the linear algebra library LAPACK [2]. When constructing these quadratic relaxations, we only consider the variables that have not been fixed at the current node. We use CPLEX as a subsolver for the new QP relaxations. The relaxation solution returned by the QP subsolver is used at the current node only if it satisfies the KKT conditions. This KKT test is similar to the optimality checks that BARON performs on the solutions returned by the LP and NLP subsolvers (see [17] for details).

Another important component of our implementation is the approximate spectral branching rule described in §4. This strategy is activated whenever the original problem supplied to BARON is a nonconvex binary QP. When this strategy is disabled, BARON uses reliability branching [1] to select among binary branching variables.

Finding δ

When constructing the quadratic relaxation (40), we use a sufficiently large value of δ to obtain a good approximation of the bound given by the eigenvalue relaxation in the nullspace of A . We use an iterative procedure to determine such value of δ . We start by setting $\delta = 1$ and computing $\lambda_{\min}(Q, I_n + \delta A^T A)$. Then, in each iteration, we increase δ by a factor of σ and we use the resulting δ to compute a new value of $\lambda_{\min}(Q, I_n + \delta A^T A)$. The procedure terminates when either the relative change in $\lambda_{\min}(Q, I_n + \delta A^T A)$ is within a tolerance *relTol* or the number of iterations reaches *maxIter*. In our experiments, we set $\sigma = 10$, *maxIter* = 5, and *relTol* = 10^{-3} . This procedure is executed at the root node only, and the value of δ determined during its execution is used throughout the entire branch-and-bound tree.

Dynamic relaxation selection strategy

We have implemented a dynamic relaxation selection strategy which is used for problems of the form (1) and switches between polyhedral and quadratic relaxations based on their relative strength. This dynamic strategy is motivated by two key observations. First, the strength of a given relaxation may depend on particular characteristics of the problem under consideration. Second, a particular type of relaxation may become stronger than other classes of relaxations as we move down the branch-and-bound tree.

We dynamically adjust the frequencies at which we solve the different types of relaxations during the branch-and-bound search. Denote by $\omega_{lp} \in [1, \bar{\omega}_{lp}]$ and $\omega_{qp} \in [1, \bar{\omega}_{qp}]$ the frequencies with which we solve the LP and QP relaxations, respectively. Let f_{lp} and f_{qp} be the optimal objective function values of the LP and QP relaxations, respectively. At the beginning of the global search, we set $\omega_{lp} = 1$ and $\omega_{qp} = 1$, which indicates

that both the LP and QP relaxations will be solved at every node of the branch-and-bound tree. At nodes where both LP and QP relaxations are solved, we compare their corresponding objective function values. If $f_{qp} - f_{lp} \geq \text{absTol}$, we increase ω_{qp} and decrease ω_{lp} by setting $\omega_{qp} = \max(1, \omega_{qp}/\sigma_{qp})$ and $\omega_{lp} = \min(\bar{\omega}_{lp}, \omega_{lp} \cdot \sigma_{lp})$. Conversely, if $f_{qp} - f_{lp} < \text{absTol}$, we increase ω_{lp} by setting $\omega_{lp} = \max(1, \omega_{lp}/\sigma_{lp})$, and decrease ω_{qp} by setting $\omega_{qp} = \min(\bar{\omega}_{qp}, \omega_{qp} \cdot \sigma_{qp})$. In our experiments, we set $\sigma_{lp} = 10$, $\sigma_{qp} = 2$, $\bar{\omega}_{lp} = 1000$, $\bar{\omega}_{qp} = 10$, and $\text{absTol} = 10^{-3}$.

Although BARON also makes use of MILP relaxations, in our dynamic relaxation selection strategy, we only compare the bounds given by the LP and QP relaxations. As MILP relaxations can be expensive, BARON uses a heuristic to decide if an MILP relaxation will be solved at the current node [19]. In our implementation, this heuristic is invoked only if at the current node the QP relaxation is weaker than the LP relaxation. If the converse is true, the MILP relaxation is skipped altogether.

6 Computational results

In this section, we present the results of a computational study conducted to investigate the impact of the techniques proposed in this paper on the performance of branch-and-bound algorithms. We start in §6.1 with a numerical comparison between the spectral relaxations introduced in §3.1–3.3 and some relaxations reviewed in §2. In §6.2, we analyze the impact of the implementation described in §5 on the performance of the global optimization solver BARON. Then, in §6.3, we compare several state-of-the-art global optimization solvers. Finally, in §6.4, we compare BARON and the QCR approach discussed in §2.3.3.

Throughout this section, all experiments are conducted under GAMS 30.1.0 on a 64-bit Intel Xeon X5650 2.66GHz processor with a single-thread. We solve all problems in minimization form. For the experiments described in §6.1, the linear and convex quadratic programs are solved using CPLEX 12.10, whereas the SDPs are solved using MOSEK 9.1.9. For the experiments considered in §6.2–6.4, we consider the following global optimization solvers: ANTIGONE 1.1, BARON 19.12, COUENNE 0.5, CPLEX 12.10, GUROBI 9.0, LINDOGLOBAL 12.0 and SCIP 6.0. When dealing with nonconvex problems, we: (i) run all solvers with relative and absolute optimality tolerances of 10^{-6} and a time limit of 500 seconds, and (ii) set the CPLEX option `optimalitytarget` to 3 and the GUROBI option `nonconvex` to 2 to ensure that these two solvers search for a globally optimal solution. For other algorithmic parameters, we use default settings. The computational times reported in our experiments do not include the time required by GAMS to generate problems and interface with solvers; only times taken by the solvers are reported.

For our experiments, we use a large test set consisting of 960 Cardinality Binary Quadratic Programs (CBQPs), 30 Quadratic Semi-Assignment Problems (QSAPs), 246 Box-Constrained Quadratic Programs (BoxQPs), and 315 Equality Integer Quadratic Programs (EIQPs). These test libraries are described in detail in [12].

6.1 Comparison between relaxations

In this section, we provide a comparison between the spectral relaxations introduced in §3.1–3.3, the convex quadratic relaxation (8), the first-level RLT relaxation (3), and the SDP relaxations (5) and (7). We construct performance profiles based on the root-node

relaxation gap:

$$\text{GAP} = \left(\frac{f_{UBD} - f_{LBD}}{\max(|f_{LBD}|, 10^{-3})} \right) \times 100 \quad (60)$$

where f_{LBD} is the root-node relaxation lower bound, and f_{UBD} is the best upper bound available for a given instance. The following notation is used to refer to the different relaxations:

- EIG: Eigenvalue relaxation (18).
- GEIG: Generalized eigenvalue relaxation (24).
- EIGNS: Eigenvalue relaxation in the nullspace of A (32).
- EIGDC: Quadratic relaxation (8) based on the eigdecomposition of Q .
- RLT: First-level RLT relaxation (3).
- SDPd: SDP relaxation (5).
- SDPda: SDP relaxation (7).

Performance profiles are presented in Figures 2a–2d. These profiles show the percentage of models for which the gap defined in (60) is below a certain threshold. As seen in the figures, the SDP relaxations give the tighter bounds, followed by the spectral relaxations. For these instances, both the RLT relaxation (3) and the quadratic relaxation (8) provide relatively weak bounds.

We also compare these root-node relaxations in terms of their solution times. To that end, in Figures 3a–3d, we present the geometric means of the CPU times required to solve the different classes of relaxations. For the quadratic relaxation based on the eigdecomposition of Q , the CPU time includes the time required to solve the convex QP (8) and the time taken to solve the LPs used to determine the bounds on the y_i variables. We group the instances based on their size. As the figures indicate, the spectral relaxations are relatively inexpensive regardless of the characteristics of the problem. As the size of the problem increases, the RLT relaxations become more expensive to solve, and in some cases, these RLT relaxations are orders of magnitude more expensive than the other relaxations. Note that the separable programming procedure described in §2.3.1 does not only lead to relatively weak bounds, but it is also computationally expensive since it requires the solution of $2n$ linear programs. Even though for most of the problems considered in the experiments the SDP relaxations can be solved within 10 seconds, they are between one and two orders of magnitude more expensive than the spectral relaxations.

6.2 Impact of the implementation on BARON’s performance

In this section, we demonstrate the benefits the proposed relaxation and branching techniques on the performance of the global optimization solver BARON. We consider the following versions of BARON 19.12:

- BARONnoqp: BARON without the spectral relaxations and without the spectral branching rule.

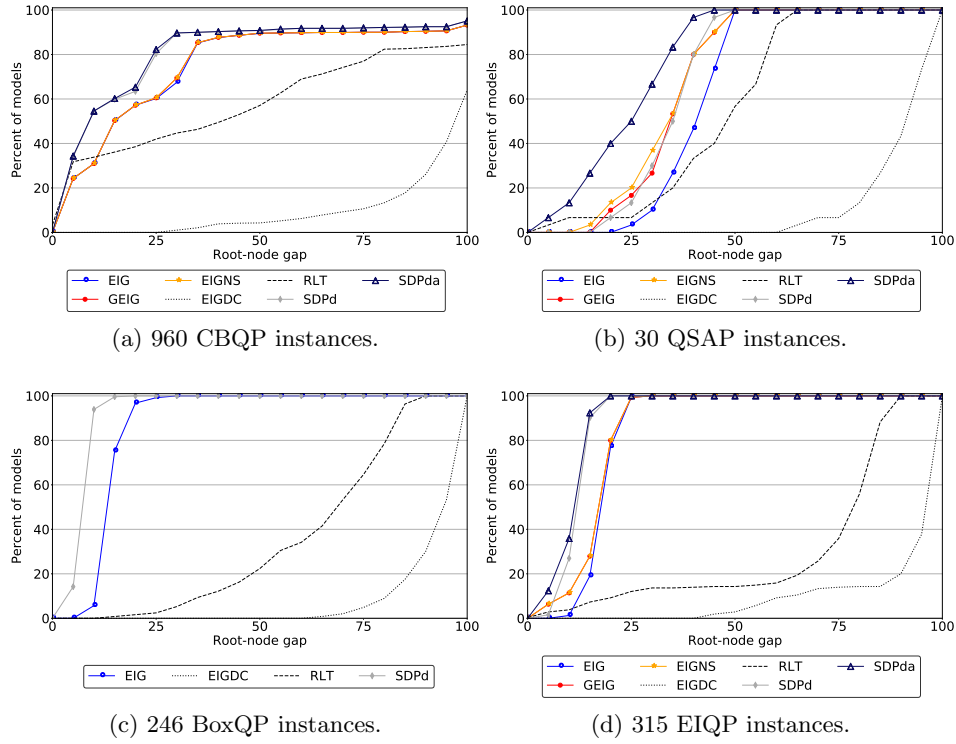


Figure 2: Comparison between the root-node relaxations gaps.

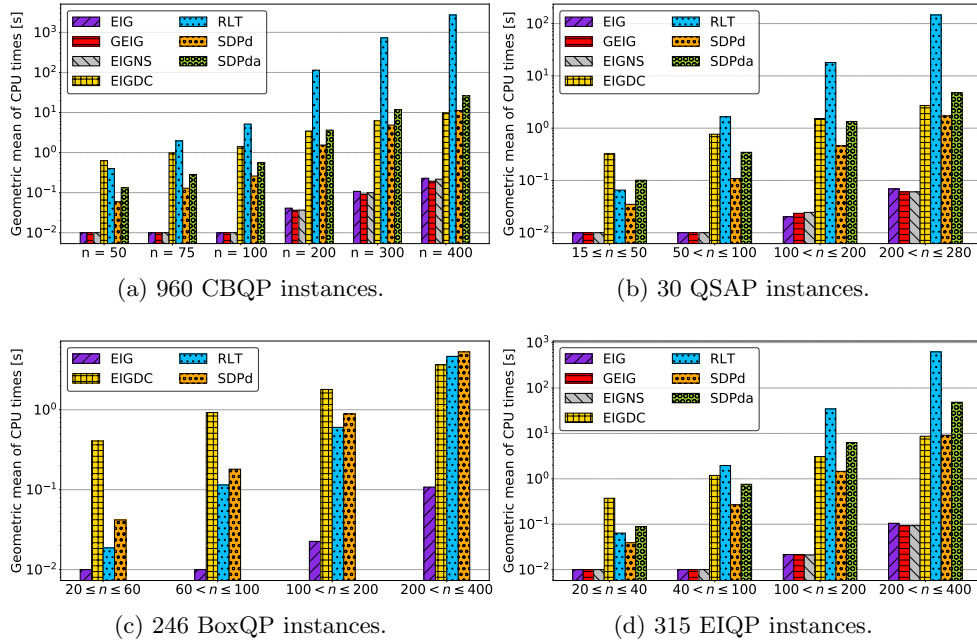


Figure 3: Geometric means of the CPU times required to solve the root-node relaxations.

- BARONnosb: BARON with the spectral relaxations but without the spectral branching rule.
- BARON: BARON with the spectral relaxations and the approximate spectral branching rule. This is the default version of this solver.

As mentioned previously, the spectral branching rule introduced in §4 is only used for the binary instances. In order to analyze the impact of our implementation, we start by comparing the different versions of BARON through performance profiles. For instances which can be solved to global optimality within the time limit of 500 seconds, we use performance profiles based on CPU times. In this case, for a given solver, we plot the percentage of models that can be solved within a certain amount of time. For problems for which global optimality cannot be proven within the time limit, we employ performance profiles based on the optimality gaps at termination. These gaps are determined according to (60) by using the best lower and upper bounds reported by the solver under consideration. In this case, for a given solver, we plot the percentage of models for which the remaining gap is below a given threshold.

The performance profiles are presented in Figures 4a–4d. As seen in the figures, our implementation leads to very significant improvements in the performance of BARON. Clearly, for the CBQP and QSAP instances, both the spectral relaxations and the spectral branching strategy result in a version of BARON which is able to solve many more problems to global optimality. In addition, for the four collections considered in this comparison, BARON terminates with much smaller relaxation gaps in cases in which global optimality cannot be proven within the time limit.




Figure 4: Comparison between the different versions of BARON.

Next, we provide a more detailed comparison between BARON and BARONnoqp. To

this end, we eliminate from the test set all the problems that can be solved trivially by both solvers (146 instances). A problem is regarded as trivial if it can be solved by both solvers in less than one second. After eliminating all of these problems from the original test set, we obtain a new test set consisting of 1405 instances.

We first consider the nontrivial problems that are solved to global optimality by at least one of the two the versions of the solver (412 instances). For this analysis, we compare the performance of the two solvers by considering their CPU times. In this comparison, we say that the two solvers perform similarly if their CPU times are within 10% of each other. The results are presented in Figure 5a. As the figure indicates, BARON is significantly faster than BARONnoqp. For nearly 50% of the problems considered in this comparison, BARON is at least one of magnitude faster than BARONnoqp

Now, we consider the nontrivial problems that neither of the two solvers are able to solve to global optimality within the time limit (993 instances). In this case, we analyze the performance of these solvers by comparing the gaps reported at termination. For the purposes of this comparison, we say that two solvers obtain similar gaps if their remaining gaps are within 10% of each other. The results are presented in Figure 5b. As seen in the figure, for more than 90% considered in this comparison, BARON reports significantly termination gaps than BARONnoqp.

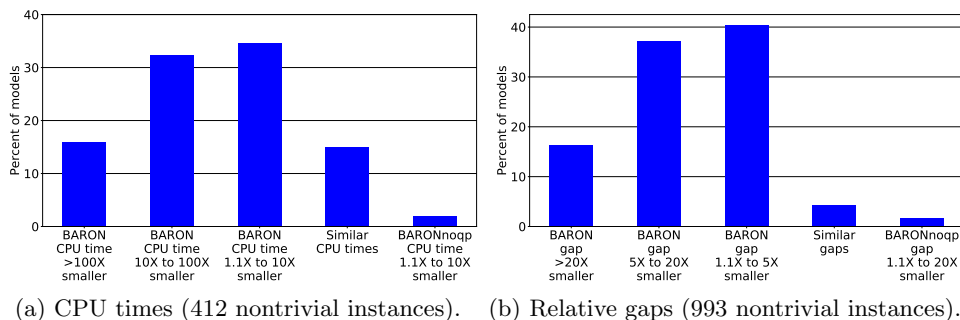


Figure 5: One-to-one comparison between BARON and BARONnoqp.

6.3 Comparison between global optimization solvers

In this section, we compare BARON with other global optimization solvers. We start by providing a comparison between different solvers using the same type of performance profiles considered in the previous section. These profiles are shown in Figures 6a–6d. As seen in these figures, BARON performs well in comparison to other solvers. For both the CBQP and QSAP instances, BARON is faster than the other solvers and solves many more problems to global optimality. For the QSAP and BoxQP instances, BARON terminates with smaller gaps than the other solvers in cases in which global optimality cannot be proven within the time limit. Many of the BoxQP and EIQP instances are very challenging and cannot be globally solved within the time limit by solvers considered in this analysis.

Next, we provide a detailed analysis involving BARON, CPLEX and GUROBI. For this analysis, we use the same type of bar plots employed in Figure 5. We start with a one-to-one comparison between BARON and CPLEX. First, we eliminate from the test set all the problems that can be solved trivially by both solvers (124 instances), obtaining a new test set with 1427 instances. In Figure 7a, we consider the nontrivial problems that are

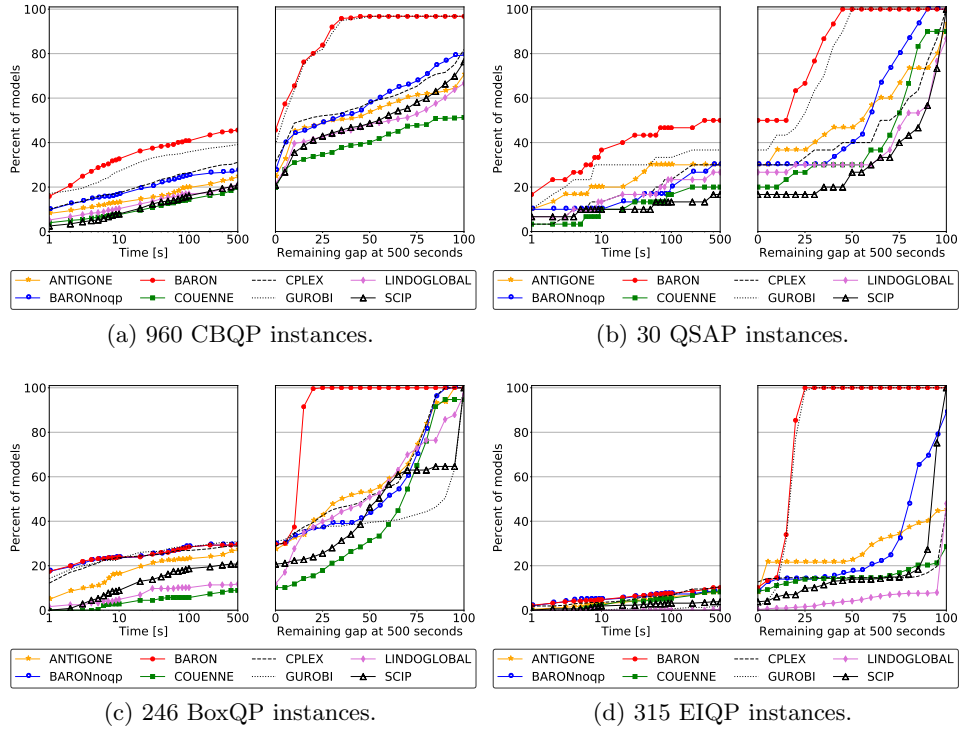


Figure 6: Comparison between global optimization solvers.

solved to global optimality by at least one of the two solvers (445 instances), whereas in Figure 7b, we consider nontrivial problems that neither of the two solvers are able to solve to global optimality within the time limit (982 instances). As both figures show, BARON performs significantly better than CPLEX. For 80% of the instances considered in Figure 7a, BARON is at least twice as fast as CPLEX. Similarly, for 90% of the instances considered in Figure 7b, BARON reports significantly smaller termination gaps than CPLEX.




Figure 7: One-to-one comparison between BARON and CPLEX.

Now, we present a similar one-to-one comparison between BARON and GUROBI. Once again, we eliminate from the test set all the problems that can be solved trivially by both solvers (185 instances), resulting in a new test set with 1366 instances. In Figure 8a, we consider the nontrivial problems that are solved to global optimality by at least one of the two solvers (380 instances), whereas in Figure 8b, we consider nontrivial problems that

neither of the two solvers are able to solve to global optimality within the time limit (986 instances). For more than 60% of the instances considered in Figure 8a, BARON is at least twice as fast as GUROBI, whereas for most of the problems considered in Figure 8b, the two solvers report similar termination gaps.

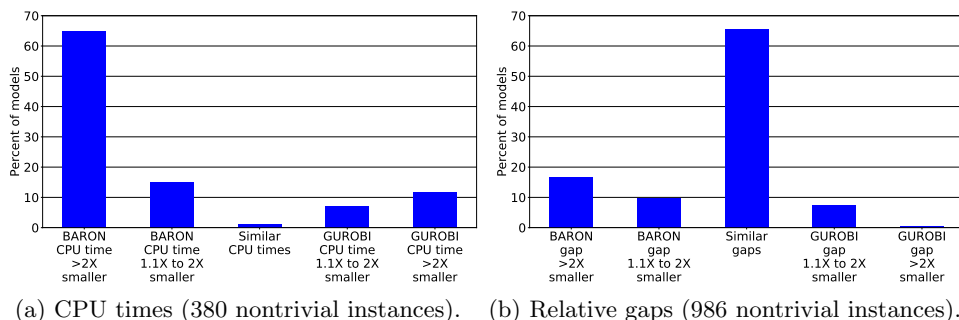


Figure 8: One-to-one comparison between BARON and GUROBI.

6.4 Comparison with the QCR method

In this section, we provide a numerical comparison between the relaxation and branching strategies proposed in this paper, and the QCR approach reviewed in §2.3.3. To this end, we consider the 990 binary CBQP and QSAP instances described in [12].

To apply the QCR method, we proceed in two steps. In the first step, for each test problem, we solve the SDP (6) with MOSEK and use its dual solution to construct a reformulated convex binary quadratic program of the form (13). In the second step, we solve the reformulated problems using a customized version of BARON, which denote by BARON_{qcr}. This version of BARON only differs from the default one in two aspects. First, BARON_{qcr} is devised in a way such that, at a given node of the branch-and-bound tree, the lower bound is obtained by solving the continuous relaxation of the reformulated problem, which is a convex QP. To this end, in the lower bounding routines of BARON_{qcr}, we have disabled the LP, NLP, MILP and recently introduced spectral relaxations. Second, in BARON_{qcr} we have also turned off the spectral branching rule and replaced it with the reliability branching strategy described in [19]. Recall that, under the QCR approach, the perturbation parameters used to derive the reformulated problem are not updated during the execution of the branch-and-bound algorithm. As a result, the QP relaxations constructed at different nodes of the branch-and-bound tree of BARON_{qcr} only differ from one another in the variables that are fixed. These convex QP relaxations are solved using CPLEX.

In our experiments, we run BARON_{qcr} with the same relative/absolute tolerances and time limit used for BARON. For all of the considered instances, the amount of time required to solve the SDP relaxation involved in the first step of the QCR method was much smaller than the CPU time corresponding to BARON_{qcr}. As a result, when comparing BARON and BARON_{qcr}, we ignore the time required to solve these SDPs.

We first compare BARON and BARON_{qcr} in terms of the lower bounds reported at the root-node. In this case, we say that a given solver obtains a better lower bound if the relative lower bound difference is greater than 10^{-3} . For cases in which the magnitude of the lower bound is below one, we use absolute differences. The results are presented in Table 1. In this table, for each test library, we provide the number, and in parentheses

the percentage, of problems for which a given solver reports better root-node lower bounds. As the results in this table indicate, BARONqcr obtains better root-node lower bounds for most of the instances considered in this comparison.

At the root node of the branch-and-bound tree, the lower bound obtained by BARONqcr is given by the continuous relaxation of (13), and as a result, it is equal to the bound provided by the SDP (6). On the other hand, for many of the problems considered in this comparison, the eigenvalue relaxation in the nullspace of A is tighter than the polyhedral relaxations implemented in BARON. Hence, in these cases, BARON relies on this quadratic relaxation to obtain lower bounds. As shown in Theorem 2, the SDP (6) is at least as tight as the eigenvalue relaxation in the nullspace of A . Therefore, it is not surprising that, for many of the problems considered in Table 1, BARONqcr provides tighter root-node bounds than BARON.

Table 1: Root-node lower bounds given by BARON and BARONqcr.

Test library	BARON better	BARONqcr better
CBQP	141 (15%)	819 (85%)
QSAP	2 (7%)	28 (93%)

Next, we compare BARON and BARONqcr using the same type of bar plots employed in previous sections. We start by eliminating from the test set all the problems that can be solved trivially by both solvers (11 instances), obtaining a new test set with 979 instances. In Figure 9a, we consider the nontrivial problems that are solved to global optimality by at least one of the two solvers (442 instances), whereas in Figure 9b, we consider nontrivial problems that neither of the two solvers are able to solve to global optimality within the time limit (537 instances). As both figures show, BARON performs significantly better than BARONqcr. For nearly 70% of the instances considered in Figure 9a, BARON is at least an order of magnitude faster than BARONqcr. Similarly, for more than 80% of the instances considered in Figure 9b, BARON reports smaller termination gaps than BARONqcr.

Even though BARONqcr reports tighter root-node lower bounds than BARON for most of the instances considered in this comparison, during the branch-and-bound search, the lower bounds obtained by BARON improve much more quickly than those provided by BARONqcr. This is due to the fact that, in BARON, we update the perturbation parameters used to construct the quadratic relaxations as we branch. In addition, BARON makes use of the approximate spectral rule introduced in §4, which as shown in §6.2, also has a significant impact on the performance of this solver.

7 Conclusions

We introduced a family of convex quadratic relaxations for nonconvex QPs and MIQPs. We studied the theoretical properties of these relaxations, showing that they are equivalent to certain SDPs. We devised a novel branching variable selection strategy which approximates the impact of the branching decisions on the quality of these relaxations. To assess the benefits of our approach, we incorporated the proposed relaxation and branching techniques in the global solver BARON, and tested our implementation on a large collection of problems. Results demonstrated that, our implementation leads to a significant improvement in the performance of BARON, enabling it to solve many more problems to global optimality.




Figure 9: One-to-one comparison between BARON and BARONqcr.

References

- [1] T. Achterberg, T. Koch, and A. Martin. Branching rules revisited. *Operations Research Letters*, 33:42–54, 2005.
- [2] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' guide*, volume 9. Siam, 1999.
- [3] I. P. Androulakis, C. D. Maranas, and C. A. Floudas. α BB: A global optimization method for general constrained nonconvex problems. *Journal of Global Optimization*, 7:337–363, 1995.
- [4] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. A recipe for semidefinite relaxations of (0,1) quadratic programming. Technical report, DIMACS, 1995.
- [5] X. Bao, A. Khajavirad, N. V. Sahinidis, and M. Tawarmalani. Global optimization of nonconvex problems with multilinear intermediates. *Mathematical Programming Computation*, 7:1–37, 2015.
- [6] X. Bao, N. V. Sahinidis, and M. Tawarmalani. Semidefinite relaxations for quadratically constrained quadratic programming: A review and comparisons. *Mathematical Programming*, 129:129–157, 2011.
- [7] A. Billionnet, S. Elloumi, and A. Lambert. An efficient compact quadratic convex reformulation for general integer quadratic programs. *Computational Optimization and Applications*, 54:141–162, 2013.
- [8] A. Billionnet, S. Elloumi, and M. C. Plateau. Improving the performance of standard solvers for quadratic 0-1 programs by a tight convex reformulation: The QCR method. *Discrete Applied Mathematics*, 157:1185–1197, 2009.
- [9] I. M. Bomze and M. Locatelli. Undominated d.c. decompositions of quadratic functions and applications to branch-and-bound approaches. *Computational Optimization and Applications*, 28:227–245, 2004.
- [10] P. Bonami, O. Günlük, and J. Linderoth. Globally solving nonconvex quadratic programming problems with box constraints via integer programming methods. *Mathematical Programming Computation*, pages 1–50, 2018.

- [11] C. Buchheim and A. Wiegele. Semidefinite relaxations for non-convex quadratic mixed-integer programming. *Mathematical Programming*, 141:435–452, 2013.
- [12] C. J. Nohra and A. U. Raghunathan and N. V. Sahinidis. A test set of quadratic, binary quadratic and integer quadratic programs. Available at <ftp://ftp.merl.com/pub/raghunathan/MIQP-TestSet/>.
- [13] A. Faye and F. Roupin. Partial lagrangian relaxation for general quadratic programming. *4OR*, 5:75–88, 2007.
- [14] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of ACM*, 42:1115–1145, 1995.
- [15] G. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, third edition, 1996.
- [16] P. L. Hammer and A. A. Rubin. Some remarks on quadratic programming with 0-1 variables. *RAIRO-Operations Research*, 4:67–79, 1970.
- [17] A. Khajavirad and N. V. Sahinidis. A hybrid LP/NLP paradigm for global optimization relaxations. *Mathematical Programming Computation*, 10:383–421, 2018.
- [18] E. M. Loiola, N. M. M. de Abreu, P.O. Boaventura-Netto, P. Hahn, and T. Querido. A survey for the quadratic assignment problem. *European Journal of Operational Research*, 176:657–690, 2007.
- [19] M. Kılınç and N. V. Sahinidis. Exploiting integrality in the global optimization of mixed-integer nonlinear programming problems in BARON. *Optimization Methods and Software*, 33:540–562, 2019.
- [20] G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems. *Mathematical Programming*, 10:147–175, 1976.
- [21] R. Misener, J. B. Smadbeck, and C. A. Floudas. Dynamically generated cutting planes for mixed-integer quadratically constrained quadratic programs and their incorporation into glomiqo 2. *Optimization Methods and Software*, 30:215–249, 2015.
- [22] P. M. Pardalos, J. H. Glick, and J. B. Rosen. Global minimization of indefinite quadratic problems. *Computing*, 539:281–291, 1987.
- [23] A.T. Phillips and J.B. Rosen. A quadratic assignment formulation of the molecular conformation problem. *Journal of Global Optimization*, 4:229–241, 1994.
- [24] N. V. Sahinidis. BARON: A general purpose global optimization software package. *Journal of Global Optimization*, 8:201–205, 1996.
- [25] H. D. Sherali and W. P. Adams. *A Reformulation-Linearization Technique for Solving Discrete and Continuous Nonconvex Problems*, volume 31 of *Nonconvex Optimization and its Applications*. Kluwer Academic Publishers, Dordrecht, 1999.
- [26] N.Z. Shor. Quadratic optimization problems. *Soviet Journal of Computer and Systems Sciences*, 25:1–11, 1987.

- [27] M. Tawarmalani and N. V. Sahinidis. Global optimization of mixed-integer nonlinear programs: A theoretical and computational study. *Mathematical Programming*, 99:563–591, 2004.
- [28] H. Tuy. DC optimization: Theory, methods and algorithms. In *Handbook of global optimization*, pages 149–216. Springer Verlag, 1995.
- [29] K. Zorn and N. V. Sahinidis. Global optimization of general nonconvex problems with intermediate bilinear substructures. *Optimization Methods and Software*, 29:442–462, 2013.