Explicit minimisation of a convex quadratic under a general quadratic constraint

Casper J. Albers[1]

Department of Psychometrics and Statistics,
University of Groningen,
Grote Kruisstraat 2/1,
9712TS Groningen,
The Netherlands.

Frank Critchley, John C. Gower

Department of Mathematics and Statistics,
The Open University,
Walton Hall,
Milton Keynes MK7 6AA,
United Kingdom.

Abstract

A complete, explicit, analytical resolution is given of a very widely occurring problem: minimisation of a convex quadratic under a general quadratic, equality or inequality, constraint. Completeness comes via identification of a set of mutually exclusive and exhaustive special cases. Explicitness, via algebraic expressions for each solution set, depending at most upon a specified real root of a known polynomial. Overall, the analysis presented gives insight into the diverse forms taken by the problem and its solution set, showing how both may be intrinsically unstable, with implications for algorithm performance and design. Underlying geometry illuminates algebraic development throughout. In particular, affine equivalence is exploited to re-express the same problem in simpler coordinate systems. Points of contact with simultaneous diagonalisation results are noted.

Keywords: constrained optimisation, quadratic programming, simultaneous diagonalisation
2000 MSC: 15A18, 90C20, 90C30

1. Introduction

1.1. Background

The problem of minimising a convex quadratic under a general quadratic, equality or inequality, constraint occurs very widely. It may arise either by itself, or as a component of a larger problem as, for example, when iteratively minimising a smooth convex objective function under a smooth constraint.

Statistical instances of this problem occur, for example, in minimum distance estimation, Bayesian decision theory, generalised linear models with smooth constraints, canonical variate analysis with fewer samples than variables, various forms of oblique Procrustes analysis, Fisher/Guttmann estimation of optimal scores, spline fitting, estimation of Hardy-Weinberg equilibrium, and size and location constraints in iterative missing value procedures in Procrustes analysis. Relevant literature is reviewed, united and extended in Albers et al. [1]. Some of these instances are elaborated in Albers et al. [2, 3].

Email addresses: c.j.albers@rug.nl (Casper J. Albers), f.critchley@open.ac.uk (Frank Critchley), john.gower@open.ac.uk (John C. Gower)
URL: http://www.casperalbers.nl (Casper J. Albers)
[1]Corresponding author

Preprint submitted to Linear Algebra and its Applications *February 17, 2015*

The problem considered here has also received considerable attention in the numerical analysis and optimisation literatures. In particular, Gander [4] characterises properties of the solution to a special case of our problem, related work appearing in Eldén [5]. Algorithms to handle the extreme forms of ill-conditioning which can occur are presented, for example, in Eldén [6, 7]. Moré [8] develops an algorithm to find an approximate solution to the more general problem in which the convexity requirement for the objective function is dropped.

More widely, the whole area of quadratic optimisation under quadratic constraints is under active development in global optimisation. See, for example, Tuy and Hoai-Phuong [9] and the references therein. These authors point out the importance of providing algorithms that are stable under small constraint perturbations. Unfortunately, as we show, this is not always possible, some problems being intrinsically unstable, while non-strict convexity of the objective function is, itself, intrinsically unstable.

Concerning the present analytical approach, Critchley [10] made a preliminary study of the strictly convex case, while Gower and Dijksterhuis [11] addressed the problem in the context of Procrustes analysis and gave a preliminary algorithm, further worked out in Albers and Gower [12].

We provide a complete, explicit resolution of the problem described. Completeness comes via identification of a set of mutually exclusive and exhaustive special cases, in each of which the solution set is obtained explicitly, depending at most on a specified real root of a polynomial. The treatment given affords analytical insight into the diverse nature both of the special cases identified and of their solution sets, shedding light on algorithm performance and design.

The underlying geometry is emphasised at several points, illuminating the development. In particular, nonsingular affine transformations are used to re-express the *same* problem in simpler coordinate systems. We call such problems *affinely equivalent*.

There are points of contact with simultaneous diagonalisation results (cf. Newcomb [13]).

The following notation and conventions are used. Terms involving arrays of vanishing order are absent. R^n is endowed with the standard Euclidean inner product, inherited by each of its subspaces. Its zero member is denoted by 0_n, and the span of its first r unit coordinate vectors S_r. M_n denotes the set of all $n \times n$ real symmetric matrices, with zero member O_n. Subscripts denoting the order of arrays may be omitted when no confusion is possible. Positive (semi-)definiteness of a matrix A is denoted by $A > O$ (respectively, $A \geq O$), the latter terminology here implying that A is singular. The Moore-Penrose inverse of A is denoted by A^-. Finally, $\mathrm{diag}(\cdot, \ldots, \cdot)$ denotes a (block)diagonal matrix with the diagonal entries listed, while $\subset$ denotes strict inclusion.

For brevity, straightforward proofs are omitted.

1.2. The general problem $\mathbb{P}_\omega$

The general equality-constrained problem is as follows.

Definition 1.1. For A, B in M_n, t, b in R^n and k in R, $A \neq O_n$, writing $\omega = (A, B, t, b, k)$, the n-variable problem $\mathbb{P}_\omega$ is:

$$\begin{aligned} \text{find } \underline{L}_\omega &:= \inf_{x \in R^n} L_\omega(x), \text{ with } L_\omega(x) := (x-t)'A(x-t), \\ \text{and } \widehat{X}_\omega &:= \{\widehat{x} \in X_\omega : L_\omega(\widehat{x}) = \underline{L}_\omega\} \\ \text{subject to } Q_\omega(x) &:= x'Bx + 2b'x - k = 0, \end{aligned}$$

where the objective (loss) function $L_\omega(\cdot)$ is convex and the feasible set $X_\omega := \{x \in R^n : Q_\omega(x) = 0\}$ nonempty; when the solution set $\widehat{X}_\omega$ is nonempty, $\underline{L}_\omega = \min\{L_\omega(x) : x \in X_\omega\}$ may be written as $\widehat{L}_\omega$. The set of all such ω is denoted by Ω_n. Where no confusion is possible, we may omit the subscript ω. We say that $\mathbb{P}$ is (a) *centred* if the target $t = 0_n$, (b) *a (partitioned) least-squares problem* if A has the form $\mathrm{diag}(1, ...1, 0, ..., 0)$, and (c) *a full least-squares problem* if $A = I$. For any least-squares problem, we partition

$$x = \begin{pmatrix} x_1 \\ x_0 \end{pmatrix}, \; B = \begin{pmatrix} B_{11} & B_{10} \\ B_{01} & B_{00} \end{pmatrix}, \; t = \begin{pmatrix} t_1 \\ t_0 \end{pmatrix} \text{ and } b = \begin{pmatrix} b_1 \\ b_0 \end{pmatrix} \tag{1}$$

conformably with A, subscripts 1 and 0 connoting its range and null spaces, so that $L_\omega(x) = \|x_1 - t_1\|^2$. In this way, terms with 0 subscripts are absent when $A = I_n$, so that $x = x_1, B = B_{11}, t = t_1$ and $b = b_1$.

Remark 1.1. Denoting the rank of A by $0 < r \leq n$, $A > O$ and $L(\cdot)$ is strictly convex when $r = n$, while otherwise $A \geq O$, with $L(x) = 0$ if and only if $(x - t)$ lies in the $(n - r)$-dimensional null space of A. Symmetry of A and B is assumed without loss, while taking A nonzero avoids one obvious triviality. $X \neq \emptyset$ avoids another, entailing restrictions on (B, b, k) characterised (by negation) in Lemma 1.1 below. The overall sign of the constraint coefficients can be reversed at will, $\mathbb{P}_\omega$ being unchanged under $(B, b, k) \to (-B, -b, -k)$. Additional consistent affine constraints do not require separate treatment: they can be substituted out to arrive at an equivalent instance of $\mathbb{P}_\omega$ in fewer variables.

Lemma 1.1. *Let $Q(x) := x'Bx + 2b'x - k$, where $B \in M_n$, $b \in R^n$ and $k \in R$. Let b have decomposition $b = Bx_b + b_\perp$, $x_b := B^- b$, along the range and null spaces of B, so that $b'_\perp x_b = 0$, and let $k_+ := k + b'B^- b$. Then:*

$$Q(x) \equiv (x + x_b)'B(x + x_b) + 2b'_\perp(x + x_b) - k_+. \tag{2}$$

Accordingly, $Q(x) = 0$ does <u>not</u> have a solution if and only if

$$\begin{array}{rlll} \textit{(i):} & [B = O, & b_\perp = 0, & k_+ \neq 0], \\ \textit{or (ii):} & [B \textit{ is nonsingular}, & (-k_+ B) > O], & \\ \textit{or (iii):} & [B \neq O \textit{ is singular}, & b_\perp = 0, & k_+ \neq 0, \\ & k_+ B \textit{ has no positive eigenvalues}] & & \end{array}$$

1.3. Overview of paper

The solution to $\mathbb{P}_\omega$ when $B = O$ is well-known. For all other $\omega \in \Omega$, the infimal value $\underline{L}_\omega$ and solution set $\widehat{X}_\omega$ are given explicitly. The variant of $\mathbb{P}_\omega$ in which the constraint is relaxed to a weak inequality is also completely resolved. The organisation and principal subsidiary results of the paper are as follows.

Sections 2 and 3: Theorem 2.1 describes *affine equivalence* on the set of problems $\{\mathbb{P}_\omega : \omega \in \Omega_n\}$, Theorem 3.1 showing that every $\mathbb{P}_\omega$ is affinely equivalent to a centred least-squares problem.

Section 4 deals with the inequality constrained case. We say that two minimisation problems with the same loss function are *effectively equivalent* if they have the same infimal value attained on the same solution set. A simple continuity argument shows that, for every $\mathbb{P}_\omega$, its weak inequality constraint variant either has a trivial solution or is effectively equivalent to $\mathbb{P}_\omega$ itself, Theorem 4.1 specifying exactly when each occurs.

Section 5 discusses the $A \geq O$ case, Lemma 5.2 showing that every centred least-squares problem is now affinely equivalent to a simplified form. Theorem 5.1 establishes that there are at most three possibilities for such a form, specifying exactly when each occurs. In two of them, distinguished by whether or not $\widehat{X}$ is empty, $\underline{L} = 0$. In the third, (a) $\underline{L} > 0$ while (b) a *reduced form* of $\mathbb{P}_\omega$ is effectively equivalent to an induced (centred) full least-squares problem in $r < n$ variables. Remark 5.2 establishes direct points of contact with simultaneous diagonalisation.

Section 6: given the above results, there is no loss in assuming now that A is positive definite, in which case, $\mathbb{P}_\omega$ always has a solution (Theorem 6.1), Theorem 6.2 establishing that it is affinely – indeed, linearly – equivalent to a *canonical form* $\mathbb{P}_{\omega^*}$, governed by the spectral decomposition of B.

Section 7: exploiting orthogonal indeterminacy within eigenspaces, Theorem 7.1 establishes that solutions to $\mathbb{P}_{\omega^*}$ can be characterised in terms of those of a dimension-reduced canonical form $\mathbb{P}_{\omega^{**}}$. There is no loss in restricting attention to *regular* such forms, denoted by $\mathbb{P}_{\overline{\omega}^{**}}$ (Remark 7.2).

Section 8 presents a series of auxiliary results culminating in Theorem 8.1, specifying the minimised objective function and solution set for any $\mathbb{P}_{\overline{\omega}^{**}}$.

Section 9 discusses intrinsic instability of solution sets and problem forms, while Section 10 concludes the paper with two worked examples.

2. Affine equivalence

Recalling that the nonsingular affine transformations G_n on R^n form a group under composition, whose general member we denote by $g : x \to x_g$, the *same* instance of $\mathbb{P}_\omega$ can be re-expressed in different, affinely equivalent, coordinate systems.

Definition 2.1. For any vector a and for any nonsingular T, $g = g_{(T,a)}$ denotes the map $x \to x_g := T^{-1}(x-a)$, inducing $\omega \to \omega_g$ via:

$$\begin{array}{lcll} t & \to & t_g & := & T^{-1}(t-a) \\ A & \to & A_g & := & T'AT, \\ B & \to & B_g & := & T'BT, \\ b & \to & b_g & := & T'(b+Ba) \text{ and} \\ k & \to & k_g & := & k - a'(2b+Ba). \end{array}$$

For linear maps ($a = 0_n$), we abbreviate g as g_T.

Remark 2.1. Note that: (a) the rank and signature of A and B are maintained in those of A_g and B_g respectively, (b) A and B are unchanged under translation ($T = I$), and (c) k is unchanged under linear maps ($a = 0_n$).

Theorem 2.1. *For all $\omega \in \Omega_n$, $x \in R^n$ and $g \in G_n$: $L_\omega(x) = L_{\omega_g}(x_g)$ and $Q_\omega(x) = Q_{\omega_g}(x_g)$, so that $x \in X_\omega \Leftrightarrow x_g \in X_{\omega_g}$, $\underline{L}_\omega = \underline{L}_{\omega_g}$, $\widehat{x} \in \widehat{X}_\omega \Leftrightarrow \widehat{x}_g \in \widehat{X}_{\omega_g}$ and $\widehat{X}_\omega \neq \emptyset \Leftrightarrow \widehat{X}_{\omega_g} \neq \emptyset$, in which case $\widehat{L}_\omega = \widehat{L}_{\omega_g}$.*

In view of Theorem 2.1, we say that $\mathbb{P}_\omega$ and $\mathbb{P}_{\omega_g}$ are *affinely equivalent*, writing $\omega \sim \omega_g$. Again, if $g = g_{(T,a)}$ with T orthogonal, we call $\mathbb{P}_\omega$ and $\mathbb{P}_{\omega_g}$ *Euclideanly* equivalent.

Recall that $\omega \to \omega_- := (A, -B, t, -b, -k)$ also leaves $\mathbb{P}_\omega$ unchanged. For later use (Section 6), we note here

Lemma 2.1. *$\mathbb{P}_\omega \to \mathbb{P}_{\omega_-}$ commutes with $\mathbb{P}_\omega \to \mathbb{P}_{\omega_g}$, $g \in G$.*

3. Centred least-squares form

We characterise here the set of centred least-squares problems to which a given problem $\mathbb{P}_\omega$ is affinely equivalent. Partitioning T conformably with B, as in (1), let $\mathcal{T}$ denote $\{T : T_{11}$ is orthogonal, $T_{10} = O$ and T_{00} is nonsingular$\}$, noting that $\mathcal{T}$ forms a group under multiplication.

Theorem 3.1. *Let $\omega \in \Omega_n$ and $T_A := U_A \operatorname{diag}(D_A^{-1}, I_{n-r})$ where A has spectral decomposition $A = U_A \operatorname{diag}(D_A^2, O_{n-r})U_A'$ with U_A orthogonal and D_A diagonal, positive definite. Then:*
(i) $\mathbb{P}_{\omega_{g_0}}$, $g_0 = g_{(T_A,t)}$, is a centred least-squares problem;
(ii) $\mathbb{P}_{\omega_{g \circ g_0}}$ is also such a problem if and only if $g = g_T$ for some $T \in \mathcal{T}$.

In view of Theorems 2.1 and 3.1, there is no loss in restricting attention to centred least-squares problems $\mathbb{P}_\omega$, in which case $L_\omega(x) = \|x_1\|^2$.

4. Inequality contrained variant

We denote by $\mathbb{P}_\leq$ the inequality constrained variant of $\mathbb{P}$ in which the feasible set required to be nonempty $X := \{x \in R^n : Q(x) = 0\} \neq \emptyset$ is replaced by $X_\leq := \{x \in R^n : Q(x) \leq 0\} \neq \emptyset$, its infimal value and solution set being denoted by $\underline{L}_\leq$ and $\widehat{X}_\leq$ respectively. The reverse inequality is accommodated by changing the overall sign of (B, b, k).

Affine equivalence generalises at once to the inequality constrained case, as does Theorem 3.1. Accordingly, in discussing $\mathbb{P}_\leq$, there is again no loss in restricting attention to the centred least-squares case, when $L(x) = \|x_1\|^2$.

Theorem 4.1. *Let $\mathbb{P}_\leq$ be an inequality constrained, centred least-squares problem.*

(i) When $A > O$:
if $Q(0_n) \leq 0$, $\underline{L}_\leq = 0$ and $\widehat{X}_\leq = \{0_n\}$;
otherwise, $X \neq \emptyset$ and $\mathbb{P}_\leq$ is effectively equivalent to $\mathbb{P}$.
(ii) When $A \geq O$, putting $X_{0,\leq} := \{x_0 \in R^{n-r} : Q(0_r', x_0')' \leq 0\}$:
if $X_{0,\leq} \neq \emptyset$, $\underline{L}_\leq = 0$ and $\widehat{X}_\leq = \{(0_r', x_0')' : x_0 \in X_{0,\leq}\}$;
otherwise, $X \neq \emptyset$ and $\mathbb{P}_\leq$ is effectively equivalent to $\mathbb{P}$.

Proof. The proof is similar in both cases.

(i) If $Q(0_n) \leq 0$, the result is immediate. Else, $Q(0_n) > 0$, continuity of $Q(\cdot)$ ensuring that, $\forall\, x \in X_{\leq}$, $Q(x) < 0 \Rightarrow \exists\, 0 < \kappa < 1$ with $Q(\kappa x) = 0$.

(ii) If $X_{0,\leq} \neq \emptyset$, the result is immediate. Else, $Q(0_r', x_0')' > 0\ \forall\, x_0 \in R^{n-r}$ while, $\forall\, x \in X_{\leq}$, $Q(x) < 0 \Rightarrow \exists\, 0 < \kappa < 1$ with $Q(\kappa x_1', x_0')' = 0$.

□

5. Solving $\mathbb{P}_\omega$ when A is positive semi-definite

In this section, we take A positive semi-definite ($r < n$), so that x_0 occurs in the constraint but not in the objective function. Accordingly, any centred least-squares problem takes an associated reduced form, an immediate lemma providing geometric insight.

Denoting orthogonal projection of R^n onto S_r by $P : x \to x_1$, we have:

Definition 5.1. For any centred least-squares problem $\mathbb{P}_\omega$ with $r < n$, its *reduced form* is:

$$\text{find } \underline{L}_1 := \inf \{L_1(x_1) : x_1 \in X_1\}, L_1(x_1) := \|x_1\|^2, X_1 := P(X). \tag{3}$$

Remark 5.1. Note that X_1 is (a) nonempty, since X is nonempty, and (b) given by

$$X_1 = \{x_1 \in R^r : X_0(x_1) \neq \emptyset\}$$

where $X_0(x_1) := \{x_0 \in R^{n-r} : Q(x_1', x_0')' = 0\}$.

Lemma 5.1. *Let $\mathbb{P}_\omega$ be a centred least-squares problem with $r < n$. Then $\forall x \in R^n$, $L(x) = L_1(x_1)$ while $x \in X \Leftrightarrow [x_1 \in X_1,\ x_0 \in X_0(x_1)]$, so that $\underline{L}_1 = \underline{L}$, while $(x_1', x_0')'$ solves $\mathbb{P}_\omega \Leftrightarrow [x_1$ solves (3) and $x_0 \in X_0(x_1)]$.*

Geometrically, the reduced form seeks the infimal (squared) distance from the origin in S_r to the orthogonal projection of the conic X onto that subspace, its solution set $\widehat{X}_1$ being the orthogonal projection of $\widehat{X}$ onto S_r.

To help solve the reduced form (3), we introduce a simplifying linear transformation, via a decomposition of the null space of A according to its intersections with the range and null spaces of B_{00}.

Definition 5.2. Let $\mathbb{P}_\omega$ be a centred least-squares problem with $r < n$. Then, $\mathbb{P}_\omega$ is said to be *in simplified form* if, for some $0 \leq s_0 \leq n - r$ and for some nonsingular diagonal Γ_0 of order s_0, B has the partitioned form:

$$B = \begin{pmatrix} B_{11} & C_{10} & O \\ C_{10}' & O_{n-r-s_0} & O \\ O & O & \Gamma_0 \end{pmatrix}. \tag{4}$$

Accordingly, any term involving Γ_0 is absent if and only if $B_{00} = O$, and any involving C_{10} if and only if B_{00} is nonsingular – in particular, if B is (positive or negative) definite.

Lemma 5.2. *Let $\mathbb{P}_\omega$ be a centred least-squares problem with $r < n$. Then, $\exists\, T \in \mathcal{T}$ with $\mathbb{P}_{\omega_{gT}}$ in simplified form.*

Proof. Let B_{00} have rank $0 \leq s_0 \leq n - r$ and spectral decomposition $U_0 \operatorname{diag}(O_{n-r-s_0}, \Gamma_0)U_0'$, with U_0 orthogonal and Γ_0 nonsingular diagonal. Define C_{10} and D_{10} implicitly via

$$\operatorname{diag}(I_r, U_0)' B \operatorname{diag}(I_r, U_0) = \begin{pmatrix} B_{11} & C_{10} & D_{10} \\ C_{10}' & O_{n-r-s_0} & O \\ D_{10}' & O & \Gamma_0 \end{pmatrix}$$

and put

$$T := \operatorname{diag}(I_r, U_0) \begin{pmatrix} I_r & O & O \\ O & I_{n-r-s_0} & O \\ -\Gamma_0^{-1} D_{10}' & O & I_{s_0} \end{pmatrix}. \tag{5}$$

Then $T \in \mathcal{T}$ and so, by Theorem 3.1, $\mathbb{P}_{\omega_{gT}}$ is a centred least-squares problem in simplified form. □

We note in passing that, while preserving simplified form, a further linear transformation establishes direct points of contact with simultaneous diagonalisation.

Remark 5.2. If $\mathbb{P}_\omega$ is in simplified form, while B_{11} in (4) has spectral decomposition:

$$B_{11} = U_1 \operatorname{diag}(O_{r-s_1}, \Gamma_1)U_1', \quad \Gamma_1 \text{ nonsingular},$$

the further transformation $x \to T^{-1}x$ with $T := \operatorname{diag}(U_1, I_{n-r}) \in \mathcal{T}$ induces:

$$B \to \begin{pmatrix} O_{r-s_1} & O & E_{10} & O \\ O & \Gamma_1 & F_{10} & O \\ E_{10}' & F_{10}' & O_{n-r-s_0} & O \\ O & O & O & \Gamma_0 \end{pmatrix} \text{ in which } \begin{pmatrix} E_{10} \\ F_{10} \end{pmatrix} = U_1' C_{10} \tag{6}$$

so that, using again Theorem 3.1, $\mathbb{P}_{\omega_g}$, $g = g_T$, remains in its simplified form. This can be seen as extending Newcomb [13] who showed that any two symmetric matrices, neither of which is indefinite, can be simultaneously diagonalised. For, there is no loss in restricting attention to matrices that, like A, are either positive definite or positive semi-definite, which, if true of B, entails that E_{10} and F_{10} in (6) are absent or zero respectively.

Returning to the mainstream, the following additional terms are used.

Definition 5.3. For any centred least-squares problem $\mathbb{P}_\omega$ in simplified form, we sub-partition x_0 and b_0 so that

$$x_0 = \begin{pmatrix} y_0 \\ z_0 \end{pmatrix}, \; b_0 = \begin{pmatrix} c_0 \\ d_0 \end{pmatrix} \text{ and } B_{00} = \begin{pmatrix} O_{n-r-s_0} & O \\ O & \Gamma_0 \end{pmatrix}$$

conform. Accordingly, any term involving y_0 or c_0 is absent if and only if B_{00} is nonsingular; and any involving z_0 or d_0 if and only if $B_{00} = O$. For given x_1, $Q(\cdot)$ depends quadratically on z_0, but only linearly on y_0, since

$$Q(x) \equiv (z_0 + \Gamma_0^{-1}d_0)'\Gamma_0(z_0 + \Gamma_0^{-1}d_0) + 2(C_{10}'x_1 + c_0)'y_0 + Q_1(x_1) \tag{7}$$

in which $Q_1(x_1) := x_1'B_{11}x_1 + 2b_1'x_1 - k_1$, with $k_1 = k_1(\omega) := k + d_0'\Gamma_0^{-1}d_0$.
If $B_{00} \neq O$, $Z_0(\alpha) := \{z_0 \in R^{s_0} : (z_0 + \Gamma_0^{-1}d_0)'\Gamma_0(z_0 + \Gamma_0^{-1}d_0) = \alpha\}$, $\alpha \in R$, so that $\alpha(\Gamma_0) := \{\alpha : Z_0(\alpha) \neq \emptyset\}$ is R, $[0, \infty)$ or $(-\infty, 0]$ according as Γ_0 is indefinite, positive definite or negative definite, while $Z_0(0) = \{-\Gamma_0^{-1}d_0\}$ if Γ_0 is definite.
If B_{00} is singular and $c_0 \neq 0$, $\alpha(y_0) := k_1 - 2c_0'y_0$, while $y_0(c_0) := k_1c_0/(2\,\|c_0\|^2)$, so that $\alpha(y_0(c_0)) = 0$.

In view of previous results, when $r < n$, it suffices to solve $\mathbb{P}_\omega$ – or, giving also each $X_0(x_1)$ ($x_1 \in X_1$), its reduced form (3) – for centred least-squares problems in simplified form.

Definition 5.4. Let $\mathbb{P}_\omega$ with $r < n$ be a centred least-squares problem in simplified form. Then, $\omega_1 := (I_r, B_{11}, 0_r, b_1, k_1)$ is called the *projected form* of $\omega = (\operatorname{diag}(I_r, O), B, 0_n, b, k)$, and we say that $\mathbb{P}_\omega$ admits:
(a) *a perfect solution* if $L(x) = 0$ for some $x \in X$;
(b) *an essentially perfect solution* if $L(x) > 0$ ($x \in X$), but $\underline{L} = 0$;
(c) *a projected, yet imperfect, reduced form* if (i) $X_{\omega_1} \neq \emptyset$, so that $\omega_1 \in \Omega_r$, and (ii) its reduced form is effectively equivalent to $\mathbb{P}_{\omega_1}$:

$$\text{find } \inf\{L_1(x_1) : x_1 \in X_{\omega_1}\}, \; X_{\omega_1} = \{x_1 \in R^r : Q_1(x_1) = 0\},$$

in which $k_1 \neq 0$.

Example 5.1. Examples of these three possibilities are the problems, with $r = 1$ and $n = 2$, of finding the infimum of x_1^2 over all $(x_1, x_0)'$ satisfying, respectively, (a) $x_1^2 - x_0^2 = -1$, (b) $x_1x_0 = 1$ and (c) $x_1^2 - x_0^2 = +1$, the corresponding sets X_1 being R, $R\backslash\{0\}$ and $\{x_1 : x_1^2 \geq 1\}$.

Indeed, there are no other possibilities. Denoting the closure and boundary of X_1 by $cl(X_1)$ and ∂X_1, we have:

Theorem 5.1. *Let $\mathbb{P}_\omega$ with $r < n$ be a centred least-squares problem in simplified form. There are at most three, mutually exclusive, possibilities:*

(a) $\mathbb{P}_\omega$ *admits a perfect solution,*
(b) $\mathbb{P}_\omega$ *admits an essentially perfect solution,*
(c) $\mathbb{P}_\omega$ *admits a projected, yet imperfect, reduced form,*

whose occurrences are characterised thus:

(a) occurs $\Leftrightarrow 0_r \in X_1 \Leftrightarrow [\underline{L} = 0, \widehat{X} \neq \emptyset]$,
(b) occurs $\Leftrightarrow 0_r \in \partial X_1 \Leftrightarrow [\underline{L} = 0, \widehat{X} = \emptyset]$,
(c) occurs $\Leftrightarrow 0_r \notin cl(X_1) \Leftrightarrow \underline{L} > 0$,

these arising as follows:

(i) If $s_0 = 0$*:*

$$\textit{(a) occurs} \Leftrightarrow [k_1 = 0,\ c_0 = 0] \textit{ or } c_0 \neq 0,$$

$X_0(0)$ *being* R^{n-r} *or* $\{y_0(c_0)\}$ *respectively. Otherwise,* $[k_1 \neq 0,\ c_0 = 0]$*, with (c) or (b) occurring according as* C_{10} *does, or does not, vanish. When (c) occurs,* $X_1 = X_{\omega_1}$ *with each* $X_0(x_1) = R^{n-r}$*.*

(ii) If $s_0 = n - r$*:*

$$\textit{(a) occurs} \Leftrightarrow k_1 = 0 \textit{ or } k_1\Gamma_0 \textit{ has a positive eigenvalue,}$$

$X_0(0)$ *being* $Z_0(0)$ *or* $Z_0(k_1)$ *respectively, while (b) does* *<u>not</u>* *occur. Thus, (c) occurs* $\Leftrightarrow (-k_1\Gamma_0) \succ O$*, when* $X_1 = \{x_1 \in R^r : k_1 Q_1(x_1) \geq 0\}$ *with each* $X_0(x_1) = \{x_0 : z_0 \in Z_0(-Q_1(x_1))\}$*.*

(iii) If $0 < s_0 < n - r$*:*

$$\textit{(a) occurs} \Leftrightarrow \quad [k_1 = 0,\ c_0 = 0] \textit{ or } [k_1\Gamma_0 \textit{ has a positive eigenvalue, } c_0 = 0] \textit{ or } c_0 \neq 0,$$

$X_0(0)$ *comprising all* x_0 *with* z_0 *in* $Z_0(0)$*,* $Z_0(k_1)$ *or* $\cup_{\alpha(y_0)\in\alpha(\Gamma_0)} Z_0(\alpha(y_0))$*, respectively. Otherwise,* $[(-k_1\Gamma_0) \succ O,\ c_0 = 0]$*, with (c) or (b) occurring according as* C_{10} *does, or does not vanish. When (c) occurs,* X_1 *and each* $X_0(x_1)$ *are as in (ii).*

Proof. Using Lemma 5.1, the characterisations of the occurrence of (a) and of (b) are immediate. Since

$$0_r \notin cl(X_1) \Leftrightarrow [\exists\, \eta > 0 \text{ such that } \|x_1\| < \eta \Rightarrow x_1 \notin X_1] \tag{8}$$

so, too, is the fact that (c) occurs $\Rightarrow 0_r \notin cl(X_1) \Leftrightarrow \underline{L} > 0$. Recall that

$$x_1 \notin X_1 \Leftrightarrow Q(x_1', x_0')' \neq 0 \text{ for all } x_0 \in R^{n-r}. \tag{9}$$

Using the $x_1 = 0$ instances of (7) and (9), and identifying the former as an instance of the identity (2), Lemma 1.1 gives:

$$0_r \notin X_1 \Leftrightarrow \left\{ \begin{array}{llll} \text{(i):} & [s_0 = 0, & c_0 = 0, & k_1 \neq 0], \\ \text{or (ii):} & [s_0 = n - r, & (-k_1\Gamma_0) \succ O], & \\ \text{or (iii):} & [0 < s_0 < n - r, & (-k_1\Gamma_0) \succ O, & c_0 = 0] \end{array} \right\}.$$

Similarly, using also (8), we have:

$$0_r \notin cl(X_1) \Leftrightarrow [0_r \notin X_1 \text{ and } C_{10} \text{ is either absent } (s_0 = n - r) \text{ or zero}].$$

Suppose $0_r \notin cl(X_1)$. If also $s_0 = 0$, $Q(x) = Q_1(x_1)$, so that $X_{\omega_1} = X_1 \neq \emptyset$ and the reduced form (3) of $\mathbb{P}_\omega$ *is* $\mathbb{P}_{\omega_1}$. Thus, (c) occurs, while each $X_0(x_1) = R^{n-r}$. If, instead, we also have $s_0 > 0$, $(-k_1\Gamma_0) \succ O$ and

$$k_1 Q(x) = k_1 Q_1(x_1) - (z_0 + \Gamma_0^{-1} d_0)'(-k_1\Gamma_0)(z_0 + \Gamma_0^{-1} d_0),$$

so that $X_1 = \{x_1 : k_1 Q_1(x_1) \geq 0\}$. Since $X_1 \neq \emptyset$ while $k_1 \neq 0$, Theorem 4.1 establishes that (c) again occurs, while each $X_0(x_1)$ has the form stated. The forms taken by $X_0(0)$ when (a) occurs are immediate. □

Corollary 5.1. *Under the hypotheses of Theorem 5.1, and adopting its terminology, it is necessary for (b) to occur that B be indefinite.*

Proof. If B is not indefinite, C_{10} is either absent or zero. □

Geometrically, a hyperbolic quadratic constraint is necessary for $\mathbb{P}$ to admit an essentially perfect solution.

As Theorem 5.1 shows – and as the example of (b) given immediately before it illustrates – it is possible that $\mathbb{P}_\omega$ has no solution when $r < n$, reflecting the fact that, although X is closed, its projection $X_1 = P(X)$ may not be. In contrast, a simple compactness argument, used in proving Theorem 6.1 below, shows that $\mathbb{P}_\omega$ *always* has a solution when $r = n$.

6. Canonical form

Given the above results, it suffices to solve the centred, full least-squares form $A = I$, $t = 0$ of problem $\mathbb{P}_\omega$. The affinely constrained case being well-known, there is no essential loss in also requiring $B \neq O$. Reversing if necessary the overall sign of (B, b, k), and recalling Lemma 2.1, there is no loss in further assuming that B has a positive eigenvalue.

Such a problem always has a solution.

Theorem 6.1. $\widehat{X}_\omega \neq \emptyset$ *for any full least-squares problem* $\mathbb{P}_\omega$.

Proof. For any $\widetilde{x} \in X_\omega$, $\mathbb{P}_\omega$ is essentially equivalent to minimising $\|x - t\|^2$ over the compact set $X_\omega \cap \{x \in R^n : \|x - t\| \leq \|\widetilde{x} - t\|\}$. □

Geometrically, there is always a shortest normal from t to the conic X.

We introduce next the notion of a canonical form, a full least-squares form of $\mathbb{P}_\omega$ in which the constraint has been simplified according to the spectral decomposition of B. In it, $u_m := (1, 0'_{m-1})'$ denotes the first unit coordinate vector in R^m.

Definition 6.1. Let $\mathbb{P}_\omega$ be a centred, full least-squares problem in which B has a positive eigenvalue. Let B have rank s and distinct nonzero eigenvalues $\gamma_1 > ... > \gamma_q$ $(1 \leq q \leq s)$, so that $\gamma_1 > 0$. For each $i = 1, ..., q$, let γ_i have multiplicity $m_i \geq 1$ and let the orthogonal projection of b onto the corresponding eigenspace have length $l_i \geq 0$. Let $m_0 := n - s \geq 0$ denote the dimension of the null space of B, and $l_0 \geq 0$ the length of the orthogonal projection of b onto this subspace.

With $\varepsilon := l_0$ and $\delta = (\delta_i) \in R^q$, $\delta_i := |\gamma_i|^{-1} l_i$, let $\omega^* := (A^*, B^*, t^*, b^*, k^*)$ in which $A^* = I$, $B^* = \text{diag}(O_{m_0}, \Gamma)$ where $\Gamma := \text{diag}(\gamma_1 I_{m_1}, ..., \gamma_q I_{m_q})$,

$$t^* = \begin{pmatrix} 0_{m_0} \\ \delta_1 u_{m_1} \\ \vdots \\ \delta_q u_{m_q} \end{pmatrix}, b^* = \begin{pmatrix} \varepsilon u_{m_0} \\ 0_{m_1} \\ \vdots \\ 0_{m_q} \end{pmatrix}$$

and $k^* = k + \sum_{i=1}^{q} \gamma_i^{-1} l_i^2$. Then, we call $\mathbb{P}_{\omega^*}$ the *canonical form* of $\mathbb{P}_\omega$.

The reason for this terminology is made clear by the following result.

Theorem 6.2. $\mathbb{P}_\omega$ *is Euclideanly equivalent to* $\mathbb{P}_{\omega^*}$.

Proof. By the spectral decomposition of B, there is an orthogonal matrix T_B with $T'_B B T_B = \text{diag}(O_{m_0}, \Gamma)$, where T_B is unique up to postmultiplication by $U \in \mathcal{U} := \{\text{diag}(U_0, U_1, ..., U_q) : U'_i U_i = I_{m_i}\ (0 \leq i \leq q)\}$.

Let $T'_B b = (c', d'_1, ..., d'_q)'$ and $T'_B B T_B$ conform, so that $\|c\| = l_0$ while $\|d_i\| = l_i$ $(1 \leq i \leq q)$, and let now $U = U_B \in \mathcal{U}$ be such that the first column of U_0 is $c/\|c\|$ if $c \neq 0$ and of U_i $(1 \leq i \leq q)$ is $d_i/\|d_i\|$ if $d_i \neq 0$.

Finally, let $g : x \to x_g = U'_B\{T'_B x + (0', \gamma_1^{-1} d'_1, ..., \gamma_q^{-1} d'_q)'\}$. Then, using Lemma 1.1, $\mathbb{P}_\omega$ is Euclideanly equivalent to $\mathbb{P}_{\omega_g} = \mathbb{P}_{\omega^*}$. □

Remark 6.1. Viewed geometrically, the isometry inducing $\omega \to \omega^*$ so translates, rotates and reflects coordinate axes that:

(a) B^*, determining the quadratic part of the constraint, becomes diagonal,

(b) b^*, determining its linear part, vanishes when B^* is nonsingular and, otherwise, has at most one nonzero element – its coordinate $\varepsilon \geq 0$ on the first dimension of the null space of B^* – and:

(c) the target t^* is *orthogonal* to this subspace, and has at most one nonzero element associated with each nonzero eigenvalue of B^* – its coordinate $\delta_i \geq 0$ on the first dimension of the corresponding eigenspace.

As is explicit in the proof of Theorem 6.2, the simple structure enjoyed by a canonical form within each eigenspace of B^* is made possible by a well-known orthogonal indeterminacy there. This simple structure invites a natural dimension reduction, further exploiting this same indeterminacy.

7. Dimension-reduced canonical form

Solutions to $\mathbb{P}_{\omega^*}$ can be characterised in terms of those of a dimension-reduced canonical form $\mathbb{P}_{\omega^{**}}$, having just one variable for each distinct eigenvalue of B^*. Variables $z = (z_i) \in R^q$ enter the constraint *quadratically*. When B^* is singular, an additional scalar variable y associated with its null space enters the constraint *linearly*.

A unified account is provided in which we denote these variables by $w \in R^{\underline{n}}$, whether B^* is nonsingular ($m_0 = 0$) or not ($m_0 > 0$). In the former case, $\underline{n} = q$ and $w = z$. In the latter, $\underline{n} = q + 1$ and $w = (y, z')'$.

Definition 7.1. Adopting the assumptions and notation of Definition 6.1, let $\Gamma^* := \text{diag}(\gamma_1, ..., \gamma_q)$, so that Γ^* is either positive definite ($\gamma_q > 0$) or nonsingular indefinite ($\gamma_1 > 0 > \gamma_q$, requiring $q > 1$).
If B^* is singular, let $w_0 := (0, \delta')'$ and $\Delta := \text{diag}(0, \Gamma^*)$ conform with $w := (y, z')'$, while $d := \varepsilon u_{q+1}$. Otherwise, put $w := z$, $w_0 := \delta$, $\Delta := \Gamma^*$ and $d := 0_q$. Then, the *dimension-reduced canonical form* of $\mathbb{P}_{\omega^*}$ is $\mathbb{P}_{\omega^{**}}$, $\omega^{**} := (A^{**}, B^{**}, t^{**}, b^{**}, k^{**})$, with $A^{**} = I_{\underline{n}}$, $B^{**} = \Delta$, $t^{**} = w_0$, $b^{**} = d$ and $k^{**} = k^*$. That is, the problem $\mathbb{P}_{\omega^{**}}$ is to:

$$\begin{aligned} \text{find } \underline{L}^* &:= \inf_{w \in R^{\underline{n}}} L^*(w),\ L^*(w) := \|w - w_0\|^2, \\ \text{and } \widehat{W} &:= \{\widehat{w} \in W : L^*(w) = \underline{L}^*\} \\ \text{subject to } Q^*(w) &:= w'\Delta w + 2d'w - k^* = 0, \end{aligned} \tag{10}$$

where the feasible and solution sets $W := \{w \in R^{\underline{n}} : Q^*(w) = 0\}$ and $\widehat{W}$ are nonempty, as X_ω and $\widehat{X}_\omega$ are so. We may write $\widehat{w} \in \widehat{W}$ as $\widehat{z} = (\widehat{z_i})$ ($m_0 = 0$), or $(\widehat{y}, \widehat{z}')'$ ($m_0 > 0$). Reflecting the form of the conic, when $d = 0$, we call the constraint in $\mathbb{P}_{\omega^{**}}$ *elliptic* or *hyperbolic* according as $\gamma_q > 0$ or $\gamma_q < 0$. When $d \neq 0$, we call it *parabolic-elliptic* or *parabolic-hyperbolic* in the same two cases.

Remark 7.1. When each eigenvalue of B – equivalently, of B^* – is simple, no such dimension reduction is possible, and $\mathbb{P}_{\omega^{**}}$ coincides with $\mathbb{P}_{\omega^*}$. That is,

$$[s = q \text{ and } m_0 \in \{0, 1\}] \Rightarrow [\underline{n} = n \text{ and } \omega^{**} = \omega^*].$$

Theorem 7.1. *Let $\widehat{x} := (\widehat{x}_0', \widehat{x}_1', ..., \widehat{x}_q')'$, where $\widehat{x}_i \in R^{m_i}$ $(0 \leq i \leq q)$. Then, $\widehat{x}$ solves $\mathbb{P}_{\omega^*}$ if and only if*

$$\widehat{x}_0 = \widehat{y} u_{m_0} \text{ and, for } i = 1, ..., q, \left\{ \begin{array}{ll} \widehat{x}_i = \widehat{z}_i u_{m_i} & \text{if } \delta_i > 0 \\ \left\|\widehat{x}_i\right\|^2 = \widehat{z}_i^2 & \text{if } \delta_i = 0 \end{array} \right\},$$

where $\widehat{w} = (\widehat{y}, \widehat{z}')'$ solves $\mathbb{P}_{\omega^{**}}$. In particular, $\underline{L}^* = \underline{L}$.

Clearly, $\delta_i > 0 \Rightarrow \widehat{z}_i \geq 0$. Moreover, Theorem 7.1 has the immediate

Corollary 7.1. *$\widehat{w}$ solving $\mathbb{P}_{\omega^{**}}$ determines more than one solution to $\mathbb{P}_{\omega^*}$ $\Leftrightarrow$ $[\delta_i = 0$ and $\widehat{z}_i \neq 0]$ for some i, in which case the sign of $\widehat{z}_i$ is indeterminate, while $\widehat{x}_i$ is orthogonally indeterminate.*

Example 7.1. A clear example of Corollary 7.1 is when the problem $\mathbb{P}_{\omega^*}$ is to minimise the (squared) distance to the unit sphere in R^n from its centre, the origin 0_n. The solution set $\widehat{X}_{\omega^*}$ is, of course, the unit sphere, while $\mathbb{P}_{\omega^{**}}$ is one-dimensional with $\widehat{X}_{\omega^{**}} = \{\pm 1\}$.

Remark 7.2. If B^* is singular, but b^* has zero component in its null space – that is, if $m_0 > 0$, but $\varepsilon = 0$ – the linear term in the constraint vanishes so that the optimal $\widehat{y} = 0$, reducing $\mathbb{P}_{\omega^{**}}$ to the corresponding $m_0 = 0$ case. It suffices, then, to solve dimension-reduced canonical forms which are regular in the sense defined next.

Definition 7.2. A dimension-reduced canonical form is called *regular* if either $m_0 = 0$ or $[m_0 > 0$ and $\varepsilon > 0]$. We denote such forms by $\mathbb{P}_{\overline{\omega}^{**}}$.

Since, trivially, $m_0 = 0 \Rightarrow \varepsilon = 0$, we have at once

Remark 7.3. For any regular dimension-reduced canonical form $\mathbb{P}_{\overline{\omega}^{**}}$:

$$\varepsilon > 0 \Leftrightarrow m_0 > 0 \Leftrightarrow d \neq 0 \Leftrightarrow \text{ the linear term in } Q^*(\cdot) \text{ does NOT vanish,}$$

in which case, substituting out the constraint, $\mathbb{P}_{\overline{\omega}^{**}}$ can be rephrased as the unrestricted minimisation of a quartic in z.

8. Solving $\mathbb{P}_{\overline{\omega}^{**}}$

We solve any regular dimension-reduced canonical form via a series of simple, insightful, auxiliary results. The first shows that a Lagrangian approach establishes sufficient conditions for this.

8.1. Sufficient conditions

The Lagrangian $\mathcal{L} := L^*(w) - \lambda Q^*(w)$ for $\mathbb{P}_{\overline{\omega}^{**}}$ has normal equations

$$[I - \lambda\Delta]w = (w_0 + \lambda d) \tag{11}$$

and Hessian $H := 2[I - \lambda\Delta]$.

Definition 8.1. We refer to $\Lambda := \{\lambda : H$ has no negative eigenvalues$\}$ as the *admissible region* – that is:

$$\Lambda := (-\infty, \gamma_1^{-1}] \text{ if } \gamma_q > 0, \text{ while } \Lambda := [\gamma_q^{-1}, \gamma_1^{-1}] \text{ if } \gamma_q < 0 \tag{12}$$

– its interior, denoted by Λ°, being where $H \succ O$. For each $\lambda \in \Lambda$,

$$W_N(\lambda) := \{w \in R^{\underline{n}} : (11) \text{ holds}\},$$

so that $W(\lambda) := W \cap W_N(\lambda)$ denotes the set of feasible vectors, if any, obeying the normal equations.

Lemma 8.1. *For any $\lambda \in \Lambda$, $W(\lambda) \neq \emptyset \Rightarrow W(\lambda) = \widehat{W}$. Indeed, if $w_N \in W(\lambda)$,*

$$\underline{L}^* = \|w_N - w_0\|^2 \text{ and } \widehat{W} = \{w \in W : H(w - w_N) = 0\} = W(\lambda).$$

In particular, if $\lambda \in \Lambda^\circ$, w_N ***uniquely*** *solves $\mathbb{P}_{\overline{\omega}^{**}}$.*

Proof. It suffices to note that, by the Cosine Law:

$$\forall w \in W,\ L^*(w) - L^*(w_N) = (w - w_N)'[I - \lambda\Delta](w - w_N) \geq 0.$$

□

In view of Lemma 8.1, $\mathbb{P}_{\overline{\omega}^{**}}$ is completely resolved if we can find, in explicit form, a nonempty set $W(\lambda)$ for any $\lambda \in \Lambda$.

8.2. Feasible solutions to the normal equations

Characterising when feasible solutions to the normal equations exist requires extensions of terminology established in Section 7. Continuing its unified presentation of $\underline{n} = q$ and $\underline{n} = q + 1$, terms involving any of y, $\widehat{y}$ or, introduced below, $\widehat{y}_1$ or $\widehat{y}_q$ are absent by convention if $m_0 = 0$.

We distinguish between interior and boundary points of Λ. Accordingly, recalling (12), there are either two or three cases to consider according as $\gamma_q > 0$ or $\gamma_q < 0$. We call these Cases A ($\lambda \in \Lambda^\circ$), B_1 ($\lambda = \gamma_1^{-1}$) and B_q ($\lambda = \gamma_q^{-1}$), this last arising when and *only* when $\gamma_q < 0$. By convention this last condition – that the constraint (10) be, at least in part, hyperbolic – is implicit whenever any of the terms defined only in Case B_q is mentioned. In particular, this applies to $\widehat{y}_q$, $\widehat{z}_{(q)}$, $\Gamma^*_{(q)}$ and f_q, defined next.

Definition 8.2. Let $\mathbb{P}_{\overline{\omega}^{**}}$ be a regular dimension-reduced canonical form and let $\lambda \in \Lambda$.

Case A ($\lambda \in \Lambda^\circ$) : $w(\lambda) := [I - \lambda\Delta]^{-1}(w_0 + \lambda d)$ denotes the unique solution to the normal equations, while the function $f : \Lambda^\circ \to R$ is defined by $f(\lambda) := Q^*(w(\lambda))$. That is, for all $\lambda \in \Lambda^\circ$:

$$f(\lambda) := \textstyle\sum_{i=1}^{q}(1 - \lambda\gamma_i)^{-2}\gamma_i\delta_i^2 + 2\varepsilon^2\lambda - k^*. \tag{13}$$

$\underline{f}$ and $\overline{f}$ respectively denote the infimum and supremum of $\{f(\lambda) : \lambda \in \Lambda^\circ\}$.

Case B_1 ($\lambda = \gamma_1^{-1}$) : $\widehat{y}_1 := \varepsilon\gamma_1^{-1}$ ($m_0 > 0$), $\widehat{z}_{(1)} \in R^{q-1}$ has general element $[1 - (\gamma_i/\gamma_1)]^{-1}\delta_i$ ($i = 2, ..., q$), while $f_1 := \widehat{z}'_{(1)}\Gamma^*_{(1)}\widehat{z}_{(1)} + 2\varepsilon\widehat{y}_1 - k^*$ where $\Gamma^* \equiv \mathrm{diag}(\gamma_1, \Gamma^*_{(1)})$.

Case B_q ($\lambda = \gamma_q^{-1}$) : when and *only* when $\gamma_q < 0$: $\widehat{y}_q := \varepsilon\gamma_q^{-1}$ ($m_0 > 0$), $\widehat{z}_{(q)} \in R^{q-1}$ has general element $[1 - (\gamma_i/\gamma_q)]^{-1}\delta_i$ ($i = 1, ..., q-1$), while $f_q := \widehat{z}'_{(q)}\Gamma^*_{(q)}\widehat{z}_{(q)} + 2\varepsilon\widehat{y}_q - k^*$ where $\Gamma^* \equiv \mathrm{diag}(\Gamma^*_{(q)}, \gamma_q)$.

Remark 8.1. The constraint $f(\lambda) = 0$ is *polynomial* in λ.

Key properties of $f : \Lambda^\circ \to R$ are summarised in

Proposition 8.1. *For any* $\mathbb{P}_{\overline{\omega}^{**}}$:

if $\delta = 0_q$ *and* $\varepsilon = 0$, $f(\lambda) = -k^*$ *for all* $\lambda \in \Lambda^\circ$, *so that* $\underline{f} = -k^* = \overline{f}$;

else, if either δ *or* ε *does not vanish,* f *is continuous and strictly increasing,* $\overline{f} > \underline{f}$ *being given by*:

$$\begin{aligned}
\overline{f} &= f_1 \quad \text{or } +\infty \text{ according as } \delta_1 = 0 \text{ or } \delta_1 > 0, \text{ while}\\
\underline{f} &= f_q \quad \text{or } -\infty \text{ according as } \delta_q = 0 \text{ or } \delta_q > 0, \ (\gamma_q < 0),\\
\underline{f} &= -k^* \text{ or } -\infty \text{ according as } \varepsilon = 0 \text{ or } \varepsilon > 0, \ (\gamma_q > 0).
\end{aligned}$$

Consider now Case A. Regularity of $\mathbb{P}_{\overline{\omega}^{**}}$ and Proposition 8.1 give at once

Lemma 8.2. *$w(\lambda)$ does not depend upon λ* $\Leftrightarrow$ [$\delta = 0_q$ *and* $\varepsilon = 0$] $\Leftrightarrow$ $w(\lambda) = 0_q$ *for every* $\lambda \in \Lambda^\circ \Leftrightarrow \underline{f} = -k^* = \overline{f}$.

Again, using Lemma 8.1, we have at once

Lemma 8.3. *Let* $\lambda \in \Lambda^\circ$. *Then,*

$$W(\lambda) = \begin{cases} \{w(\lambda)\} & \text{if } f(\lambda) = 0 \\ \emptyset & \text{if } f(\lambda) \neq 0 \end{cases}$$

so that: $f(\lambda_1) = f(\lambda_2) = 0 \Rightarrow W(\lambda_1) = W(\lambda_2) = \widehat{W} \Rightarrow w(\lambda_1) = w(\lambda_2)$.

In view of Lemma 8.3, we may define $\widehat{w}_\circ$ as the common value of $w(\lambda)$ among all solutions to $f(\lambda) = 0$, when at least one such exists. Putting $W_\circ := \cup_{\lambda\in\Lambda^\circ}W(\lambda)$, Lemmas 8.1 and 8.3 now give

Lemma 8.4. $W_\circ = \begin{cases} \widehat{W} = \{\widehat{w}_\circ\} & \text{if } \exists\, \lambda \text{ with } f(\lambda) = 0 \\ \emptyset & \text{else.} \end{cases}$

Combining the above auxiliary results makes $\widehat{w}_\circ$ explicit and, with it, the following summary of Case A.

Proposition 8.2. *For any $\mathbb{P}_{\overline{\omega}^{**}}$, there are three possibilities*:

(a) If $\underline{f} < 0 < \overline{f}$, $f(\lambda) = 0$ has a unique solution $\widehat{\lambda}$ and $\widehat{W} = W_\circ = \{w(\widehat{\lambda})\}$.

(b) If $\underline{f} = 0 = \overline{f}$, $f(\lambda) = 0$ for every $\lambda \in \Lambda^\circ$ and $\widehat{W} = W_\circ = \{0_q\}$.

(c) In all other cases, $f(\lambda) = 0$ has no solutions and $W_\circ = \emptyset$.

Remark 8.2. When it exists, computing $\widehat{\lambda}$ is straightforward (see Albers et al. [1, 2]).

In view of Proposition 8.2, we refer to $W_\circ$ as a *potential Lagrangian solution set*, this potential being realised if and only if it is non-empty. Turning now to the boundary cases, there are two other potential Lagrangian solution sets $W_1 := W(\gamma_1^{-1})$ and $W_q := W(\gamma_q^{-1})$, this second definition being made when and *only* when $\gamma_q < 0$. The hessian H being singular here, the normal equations (11) have either no solution, or a unique solution for all but one member of w, which they leave unconstrained. We have

Proposition 8.3. *For any $\mathbb{P}_{\overline{\omega}^{**}}$:*

(1) Case B_1 $(\lambda = \gamma_1^{-1})$: $W_1 \neq \emptyset \Leftrightarrow [\delta_1 = 0$ and $f_1 \leq 0]$, in which case:

$$\widehat{W} = W_1 = \{(\widehat{y}_1, z_1, \widehat{z}'_{(1)})' : z_1^2 = (-f_1)/\gamma_1\}.$$

(2) Case B_q $(\lambda = \gamma_q^{-1} < 0)$: $W_q \neq \emptyset \Leftrightarrow [\delta_q = 0$ and $f_q \geq 0]$, in which case:

$$\widehat{W} = W_q = \{(\widehat{y}_q, \widehat{z}'_{(q)}, z_q)' : z_q^2 = f_q/(-\gamma_q)\}.$$

Proof. $W_N(\gamma_1^{-1}) \neq \emptyset \Leftrightarrow \delta_1 = 0$ in which case:

$$w \in W_N(\gamma_1^{-1}) \Leftrightarrow [y = \widehat{y}_1 \text{ and } z_{(1)} = \widehat{z}_{(1)}], \text{ where } z \equiv (z_1, z'_{(1)})',$$

z_1 being unconstrained. Thus,

$$W_1 \neq \emptyset \Leftrightarrow [\delta_1 = 0 \text{ and } \exists\, z_1 \text{ with } \gamma_1 z_1^2 + f_1 = 0] \Leftrightarrow [\delta_1 = 0 \text{ and } f_1 \leq 0].$$

(1) now follows from Lemma 8.1. The proof of (2) is entirely similar. □

In summary, $\mathbb{P}_{\overline{\omega}^{**}}$ has either two $(\gamma_q > 0)$ or three $(\gamma_q < 0)$ types of potential Lagrangian solution set – $W_\circ$, W_1 and W_q – Propositions 8.1 to 8.3 together establishing precisely when these potentials are realised and, in each case, what the corresponding solution set $\widehat{W}$ then is.

8.3. The minimised objective function and the solution set

We are now ready to solve any regular dimension-reduced canonical form $\mathbb{P}_{\overline{\omega}^{**}}$ in which, by definition, either $m_0 = 0$ or $[m_0 > 0$ and $\varepsilon > 0]$.

To aid geometric interpretation, recall that, under the assumptions and notation of Definition 6.1:
(a) $\varepsilon := l_0 \geq 0$ denotes the length of the orthogonal projection of b onto the null space of B, whose dimension is $m_0 \geq 0$;
(b) for each distinct nonzero eigenvalue $\gamma_1 > ... > \gamma_q$ of B, $\delta \in R^q$ has general element $\delta_i := l_i / |\gamma_i|$ in which $l_i \geq 0$ denotes the length of the orthogonal projection of b onto the corresponding eigenspace of B, whose dimension is $m_i \geq 1$. Accordingly:

the linear part of the constraint vanishes	$\Leftrightarrow$	$m_0 = 0$	$\Leftrightarrow \varepsilon = 0$,
the origin is feasible	$\Leftrightarrow$	$k^* = 0$,	
the origin is the target	$\Leftrightarrow$	$\delta = 0$,	and:
the constraint is, at least in part, elliptic	$\Leftrightarrow$	$\gamma_q > 0$.	

We distinguish three mutually exclusive and exhaustive types of regular dimension-reduced canonical form $\mathbb{P}_{\overline{\omega}^{**}}$. The first two are trivial.

We call $\mathbb{P}_{\overline{\omega}^{**}}$ *non-Lagrangian* if $[m_0 = 0,\ k^* = 0,\ \delta \neq 0_q$ and $\gamma_q > 0]$, (so that, in particular, W_q is undefined). Geometrically, if the feasible set is the origin, the target being a positive distance away. As Theorem 8.1 establishes, in this case, both $W_\circ$ and W_1 are empty. That is, the constraint and normal equations are inconsistent, so that $\mathbb{P}_{\overline{\omega}^{**}}$ is *not* amenable to Lagrangian solution.

We call $\mathbb{P}_{\overline{\omega}^{**}}$ *multiply-Lagrangian* if $\underline{f} = 0 = \overline{f}$. That is, if $[m_0 = 0,\ k^* = 0$ and $\delta = 0_q]$. Geometrically, if the target is the origin, through which the conic defining the constraint passes. As Theorem 8.1 establishes, in this case, $\widehat{W}$ coincides with *each* of the nonempty sets $W_\circ$, W_1 and, when defined, W_q.

Finally, we call $\mathbb{P}_{\overline{\omega}^{**}}$ *singly-Lagrangian* if it is neither non-Lagrangian nor multiply-Lagrangian. Algebraically, if either $[m_0 = 0,\ k^* = 0,\ \delta \neq 0_q$ and $\gamma_q < 0]$, or $[m_0 = 0$ and $k^* \neq 0]$, or $[m_0 > 0$ and $\varepsilon > 0]$. As Theorem 8.1 establishes, in this case, exactly *one* of $W_\circ$, W_1 and W_q is nonempty, and so provides the solution set required.

Theorem 8.1. *The minimised objective function $\underline{L}^*$ and solution set $\widehat{W}$ for a regular dimension-reduced canonical form $\mathbb{P}_{\overline{\omega}^{**}}$ are as follows.*

(a) *If $\mathbb{P}_{\overline{\omega}^{**}}$ is non-Lagrangian, $W_\circ = W_1 = \emptyset$, while*:

$$\underline{L}^* = \textstyle\sum_{i=1}^q \gamma_i^{-2} l_i^2 > 0 \text{ attained on } \widehat{W} = W = \{0_q\}.$$

(b) *If $\mathbb{P}_{\overline{\omega}^{**}}$ is multiply-Lagrangian*:

$$\underline{L}^* = 0 \text{ attained on } \widehat{W} = W_\circ = W_1 = W_q = \{0_q\}.$$

(c) *Otherwise, if $\mathbb{P}_{\overline{\omega}^{**}}$ is singly-Lagrangian, $\widehat{W}$ is the unique nonempty member of $\{W_\circ, W_1, W_q\}$. Specifically, $\underline{f}$ and $\overline{f}$ being as in Proposition 8.1*:

if $\gamma_q > 0$, $\underline{f} < 0$, while $\widehat{W}$ is $W_\circ$ or W_1 according as $\overline{f} > 0$ or $\overline{f} \leq 0$;
if $\gamma_q < 0$, $\widehat{W}$ is $W_\circ$, W_1 or W_q according as $\underline{f} < 0 < \overline{f}$, $\overline{f} \leq 0$ or $\underline{f} \geq 0$.

When $\widehat{W} = W_\circ$,

$$\underline{L}^* = \widehat{\lambda}^2\{l_0^2 + \textstyle\sum_{i=1}^q (1 - \widehat{\lambda}\gamma_i)^{-2} l_i^2\},$$

attained at $w = w(\widehat{\lambda})$, where $\widehat{\lambda}$ unique solves $f(\lambda) = 0$.

When $\widehat{W} = W_1$,

$$\underline{L}^* = (\gamma_1^{-1})^2\{l_0^2 + \textstyle\sum_{i=2}^q (1 - \gamma_1^{-1}\gamma_i)^{-2} l_i^2\} + \zeta_1^2,$$

attained at $w = (\widehat{y}_1,\ \pm\zeta_1,\ \widehat{z}'_{(1)})'$, where $\zeta_1 := \sqrt{(-f_1)/\gamma_1} \geq 0$.

When $\widehat{W} = W_q$,

$$\underline{L}^* = (\gamma_q^{-1})^2\{l_0^2 + \textstyle\sum_{i=1}^{q-1} (1 - \gamma_q^{-1}\gamma_i)^{-2} l_i^2\} + \zeta_q^2,$$

attained at $w = (\widehat{y}_q,\ \widehat{z}'_{(q)},\ \pm\zeta_q)'$, where $\zeta_q := \sqrt{f_q/(-\gamma_q)} \geq 0$.

Proof. The result follows from detailed, straightforward application of the auxiliary results of Section 8.2, noting the following:

(a) Here, $W_\circ = W_1 = \emptyset$ as $\underline{f} = 0$ and $f_1 > 0$ while, by definition, $W = \{0_q\}$.
(b) Here, $w_0 = \widehat{w}_\circ = 0_q$ and $\widehat{z}_{(1)} = \widehat{z}_{(q)} = 0_{q-1}$, while $f_1 = f_q = 0$.
(c) Here, if $\gamma_q > 0$, either $\varepsilon > 0$ so that $\underline{f} = -\infty$, or $m_0 = 0 \neq k^*$ so that $\underline{f} = -k^*$, consistency of the constraint implying $k^* \geq 0$. In either case, $\underline{f} < 0$. Again, $\underline{f} = 0 = \overline{f}$ is impossible, so that $W_\circ \neq \emptyset \Leftrightarrow \underline{f} < 0 < \overline{f}$ when, f being continuous and strictly increasing, $f(\lambda) = 0$ has a unique solution.

□

Corollary 8.1. *Let $\mathbb{P}_{\omega^*}$ be a canonical form whose dimension-reduced form $\mathbb{P}_{\overline{\omega}^{**}}$ is regular. Then:*

(1) *$\mathbb{P}_{\omega^*}$ has a unique solution whenever $\mathbb{P}_{\overline{\omega}^{**}}$ does.*

*(2) $\mathbb{P}_{\overline{\omega}^{**}}$ does **not** have a unique solution if, and only if, it is singly-Lagrangian and either:*

(a) $[f_1 < 0, \delta_1 = 0]$, when its solutions are unique up to the sign of $\widehat{z_1} \neq 0$; or:

(b) $[f_q > 0, \delta_q = 0]$, when its solutions are unique up to the sign of $\widehat{z_q} \neq 0$,

in which cases solutions to $\mathbb{P}_{\omega^*}$ are unique up to orthogonal indeterminacy of (a) $\widehat{x_1}$ or (b) $\widehat{x_q}$ respectively.

Proof. The characterisation of non-uniqueness of solution to $\mathbb{P}_{\overline{\omega}^{**}}$ is immediate from Theorem 8.1 and Proposition 8.1, the rest of (2) then following from Corollary 7.1. Inspection of Theorem 8.1 also establishes that, whenever $\mathbb{P}_{\overline{\omega}^{**}}$ has a unique solution, $\delta_i = 0 \Rightarrow \widehat{z_i} = 0$. Part (1) now follows, using again Corollary 7.1. □

Theorem 8.1 and its Corollary 8.1 complete our primary objective, providing the minimised objective function and solution set of any regular dimension-reduced $\mathbb{P}_{\overline{\omega}^{**}}$ and so, via Theorem 7.1, of any initial canonical form $\mathbb{P}_{\omega^*}$.

9. Intrinsic instability

The above analysis of any, equality or inequality constrained, problem $\mathbb{P}_\omega$ rests on several partitions of possibilities. Passage between different members of a partition can involve movement between equality and inequality of the same two reals. Or, again, between weak (≤) and strict (<) versions of an inequality. As a result, both the form of the solution set and, indeed, of the problem itself can be intrinsically unstable under arbitrarily small perturbations of problem parameters.

Noteworthy in itself, this has direct implications for algorithm performance and design, finite machine precision preventing wholly accurate detection of which partition member any given problem belongs to. One strategy here is to implement the above analysis *as if* calculations were exact, while flagging up (and, potentially, pursuing) when other members of the same partition may be appropriate. Whether or not the solution set $\widehat{X_\omega}$ varies continuously with ω across such partition boundaries is implicit in the above analysis. In particular, in its summative Theorems 4.1, 5.1, 7.1 and 8.1, and their key Corollaries 7.1 and 8.1. The discontinuities involved can be dramatic, as the following instances illustrate.

9.1. Instability of the form of the solution set

We begin with a marked instance of the passage from non-unique to unique solutions. As in Example 7.1, consider minimisation of the distance to a sphere from its centre, whose solution set is, of course, the sphere itself. Changing this to minimising the same distance from *any* other point, the solution suddenly becomes unique. This instability corresponds to the passage from $\delta_1 = 0$ to $\delta_1 > 0$ in Corollary 8.1. Overall, a careful algorithmic implementation of the above analysis will flag up the possibility of non-unique solutions, indicating what they are.

In the hyperbolic case, there can be extreme directional instability, due to the orthogonality of the eigenspaces of γ_1 and γ_q. For instance, consider minimisation of the distance from the origin to the hyperbola $z_1^2 - z_2^2 = k^*$, both of whose eigenvalues, ±1, are simple. When $k^* = 0$, this comprises the lines $z_2 = \pm z_1$ which meet at the (repeated) solution $(0, 0)'$, while $\mathbb{P}_{\overline{\omega}^{**}}$ is multiply-Lagrangian. Otherwise, it comprises two branches with these lines as asymptotes, $\mathbb{P}_{\overline{\omega}^{**}}$ being singly-Lagrangian. For $k^* > 0$, these branches meet the z_1–axis at the twin solutions $\pm(\sqrt{k^*}, 0)'$ while, for $k^* < 0$, they meet the z_2–axis at the twin solutions $\pm(0, \sqrt{-k^*})'$. Thus, no matter how small we take $k^*_+ > 0 > k^*_-$, the twin solutions for k^*_+ are always orthogonal to those for k^*_-. In terms of Theorem 8.1, the solution set $\widehat{W}$ changes from W_1 to W_q via $W_\circ$ as k^* goes from positive to negative via zero. Again, a careful algorithmic implementation of the above analysis will flag these three possibilities up.

9.2. Instability of the form of the problem

Such an implementation will also flag up when more than one form of the *problem* may be appropriate. This type of instability itself comes in a variety of forms, as we illustrate.

As a first instance of intrinsic instability of the form of the problem, consider the variant of Example 5.1(b) with constraint $(x_1 + \xi_1)x_0 = 1$ for given $\xi_1 \neq 0$. Whereas the perfect solution here has $\widehat{x_1} = 0$ for every ξ_1, it has $\widehat{x_0} = 1/\xi_1$ which tends to $+\infty$ as $\xi_1 \to 0_+$ and to $-\infty$ as $\xi_1 \to 0_-$. Such instability holds generally. Indeed, Theorem 5.1 shows that, if $\mathbb{P}_\omega$ admits an essentially perfect solution, it is intrinsically unstable under small constraint perturbations. And, relatedly, that the occurrence of essentially perfect solutions to $\mathbb{P}_\omega$ is itself intrinsically unstable, depending as it does on the vanishing of the constraint parameter c_0. Replacing c_0 by C_{10}, the same is true of reducibility of $\mathbb{P}_\omega$ when B_{00}

is singular. That is, when A is positive semi-definite, there is intrinsic instability at the boundaries between the three possibilities – $\mathbb{P}_\omega$ admits a perfect solution, $\mathbb{P}_\omega$ admits an essentially perfect solution, or $\mathbb{P}_\omega$ admits a projected, yet imperfect, reduced form – their occurrences being characterised in Theorem 5.1.

As a second instance, any positive semi-definite A is arbitrarily close to a positive definite matrix $A(\kappa)$, $\kappa > 0$ – for example, $A(\kappa) = A + \kappa I$ – analysis of the problem changing from that of Section 5 to that discussed in the sequel. Alternative forms of $A(\kappa)$ are of course possible, as the worked examples of Section 10 illustrate.

Similar intrinsic instabilities in problem form occur at the boundaries between members of the following partitions made for positive definite A:

(i) the four possible forms of constraint – elliptic, hyperbolic, or the partly parabolic variant of either – as detailed in Definition 7.1;

(ii) in the dimension-reduced canonical form $\mathbb{P}_{\omega^{**}}$, whether it is regular or not – as detailed in Definition 7.2;

(iii) the three possibilities – non-Lagrangian, multiply Lagrangian and singly Lagrangian – as detailed in Section 8.3.

10. Worked examples and conclusion

We end with two worked examples and a short conclusion. The examples differ only by a single element of A, yet involve completely different approaches to their solution. As such, they illustrate one of the intrinsic instabilities noted in Section 9. Namely, that at the boundary between $A \geq O$ where Section 5 applies, and $A > O$ where Sections 6 to 8 pertain.

Whereas this generic type of instability can of course arise whatever the scale of the problem, to make things explicit, we work here with $n = 3$ and with values of ω for which exact calculations are fairly straightforward. Specifically, we take $k = 1$, $b = 0_3$, $t = (1, 1, 1)'$, $B = I_3$ and

$$A = A(\kappa) := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -1 & 1+\kappa \end{pmatrix}, \ \kappa \geq 0,$$

so that $A(0) \geq O$, while $A(\kappa) > O$ for all $\kappa > 0$. While reporting exact results, we also indicate where there is – or is not – numerical uncertainty, especially concerning which member of a partition of possibilities the problem on hand belongs to.

Examples 10.1 and 10.2 concern $\kappa = 0$ and $\kappa = 3/2$ respectively. In the first case, one of the computed eigenvalues of A will be within machine accuracy of zero; in the second, all will be clearly positive. As $\kappa \to 0+$, both problem forms will be flagged up. Numerically, B is clearly positive definite.

Example 10.1. Here, $A = A(0)$ clearly has rank $r = 2$, and we follow the approach of Section 5.

Analysis begins with a 3-stage transformation: first, to centred least-squares form, as in part (i) of Theorem 3.1; second, to simplified form, via (5) of Lemma 5.2; and third, to simultaneous diagonal form, via the further linear transformation of Remark 5.2. Overall, this transformation is $x \to x_g = T^{-1}(x - t)$ where

$$T = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -\frac{1}{2} & -\frac{1}{\sqrt{2}} \\ 0 & \frac{1}{2} & -\frac{1}{\sqrt{2}} \end{pmatrix}$$

inducing $\omega \to \omega_g$ with $t_g = 0_3$ and $A_g = T'AT = \text{diag}(1, 1, 0)$. These two components of ω_g are known (and so do not require calculation), as is the diagonal nature of B. To within machine accuracy, $B_g = T'BT = \text{diag}(1, \frac{1}{2}, 1)$, $b_g = (1, 0, -\sqrt{2})'$ and $k_g = -2$.

Next, we apply Theorem 5.1 to $\mathbb{P}_{\omega_g}$, dropping the subscript g where notationally convenient. Identifying terms in Definitions 5.2 and 5.3 for this transformed member of Ω_3, and noting that the third diagonal element B_g is clearly positive, we have $s_0 = 1$, so that $B_{00} = \Gamma_0 = (1)$, while $b_0 = d_0 = (-\sqrt{2})$. In particular, $s_0 = n - r$, so that part (ii) of Theorem 5.1 applies. Moreover, to within machine accuracy, $k_1 = 0$ and $Z_0(0) = \{\sqrt{2}\}$, so that $\mathbb{P}_{\omega_g}$ admits perfect solution $\widehat{x}_g = (0_2', \sqrt{2})'$, $\widehat{L}_{\omega_g} = 0$.

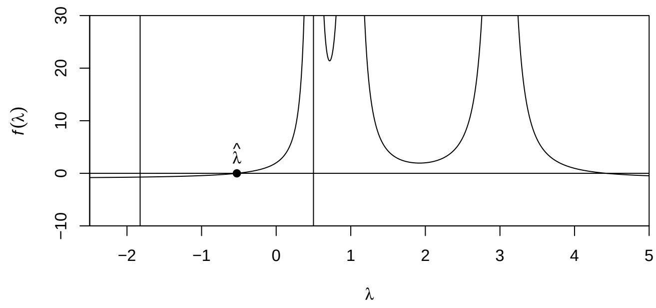

Figure 1: Display of the functional form of $f(\lambda)$. The vertical lines are at $\lambda = l^{\circ}$ and $\lambda = u^{\circ}$, defined in the text.

Finally, invoking Theorem 2.1, we transform back to conclude that $\mathbb{P}_{\omega}$ has perfect solution $\widehat{x} = t + T\widehat{x_g} = (1, 0, 0)'$, $\widehat{L_{\omega}} = 0$.

Example 10.2. Here, $A = A(3/2)$ clearly has full rank, and we follow Sections 6–8.

Analysis begins with a 2-stage transformation: first, to centred – here, full – least-squares form, using again Theorem 3.1(i); second, to canonical form, via the Euclidean transformation of Theorem 6.2. Overall, apart from a constant, this transformation is $x \to x_g = T^{-1}x$ where

$$T = \begin{pmatrix} 0 & -1 & 0 \\ -\sqrt{\frac{8}{5}} & 0 & \frac{1}{\sqrt{15}} \\ -\sqrt{\frac{2}{5}} & 0 & -\frac{2}{\sqrt{15}} \end{pmatrix}$$

inducing $\omega \to \omega_g$ in which $A_g = T'AT$ is I_3 and $B_g = T'BT$ is diagonal. To within machine accuracy, $B_g = \mathrm{diag}(2, 1, \frac{1}{3})$, which clearly has distinct, positive, eigenvalues. Accordingly, $\mathbb{P}_{\omega_g}$ is, itself, dimension-reduced and regular. Finally, in the notation of Definition 6.1, we compute that $\delta = (3/\sqrt{10}, 1, \sqrt{3/5})'$ and $k^* = 1$.

We solve the regular dimension-reduced canonical form $\mathbb{P}_{\omega_g}$ via the steps described in Section 8. Since, clearly, $m_0 = 0$ and $k^* \neq 0$, $\mathbb{P}_{\omega_g}$ is singly-Lagrangian. Here, $\Lambda^{\circ} = (-\infty, \gamma_1^{-1}) = (-\infty, \frac{1}{2})$, the function $f : \Lambda^{\circ} \to R$ of (13) taking the form:

$$\frac{9}{5(1-2\lambda)^2} + \frac{1}{(1-\lambda)^2} + \frac{9}{5(3-\lambda)^2} - 1. \tag{14}$$

Again, it is numerically clear that $\gamma_3 > 0$, $\varepsilon = 0$ and $\delta_1 > 0$ so that, by Proposition 8.1, $\underline{f} = -k^* = -1 < 0$ and $\overline{f} = +\infty$. Accordingly (Proposition 8.2), $\widehat{W} = W_{\circ} = \{w(\widehat{\lambda})\}$ where $\widehat{\lambda}$ is the unique root of $f(\lambda) = 0$. Albers et al [1] show how to reduce Λ° to a finite interval (l°, u°), $u^{\circ} := \gamma_1^{-1}$, containing $\widehat{\lambda}$. Here, $l^{\circ} = \frac{1}{2} - 3\sqrt{3/5}$. This is visualised in Figure 1, where the functional form (14) is plotted over an interval containing (l°, u°), its two vertical lines being at $\lambda = l^{\circ}$ and $\lambda = u^{\circ}$, while there are vertical asymptotes at $\lambda = \gamma_i^{-1}$, $(i = 1, 2, 3)$. Numerical solution over (l°, u°) – for example, by the bisection method – is straightforward, yielding here the approximate value $\widehat{\lambda} \approx -0.527$.

Finally, substituting in the relevant part of Theorem 8.1(c), and back transforming $\widehat{x_g} \to \widehat{x}$, yields the final solution $\widehat{x} \approx (0.655, 0.414, 0.632)'$, $\widehat{L_{\omega}} \approx 0.370$.

Conclusion. To summarise, we start from the general problem $\mathbb{P}_{\omega}$, $\omega \in \Omega$, of Definition 1.1, or its inequality constrained variant. After a number of steps to reduce the problem to a simpler form, we provide a list of mutually exclusive and exhaustive categories into which it must fall, giving the solution for each.

References

[1] C. J. Albers, F. Critchley, J. C. Gower, Quadratic minimisation problems in statistics, Journal of Multivariate Analysis 102 (3) (2011) 698 – 713. doi:10.1016/j.jmva.2009.12.018.

[2] C. J. Albers, F. Critchley, J. C. Gower, Applications of quadratic minimisation problems in statistics, Journal of Multivariate Analysis 102 (3) (2011) 714 – 722. doi:10.1016/j.jmva.2010.11.009.

[3] C. J. Albers, J. C. Gower, Canonical analysis: Ranks, ratios and fits, Journal of Classification 31 (1) (2014) 2–27. doi:10.1007/s00357-014-9146-y.

[4] W. Gander, Least squares with a quadratic constraint, Numerische Mathematik 36 (3) (1980) 291–307. doi:10.1007/BF01396656.

[5] L. Eldén, A weighted pseudoinverse, generalized singular values, and constrained least squares problems, BIT Numerical Mathematics 22 (4) (1982) 487–502. doi:10.1007/BF01934412.

[6] L. Eldén, Algorithms for the computation of functionals defined on the solution of a discrete ill-posed problem, BIT Numerical Mathematics 30 (3) (1990) 466–483. doi:10.1007/BF01931661.

[7] L. Eldén, Solving quadratically constrained least squares problems using a differential-geometric approach, BIT Numerical Mathematics 42 (2) (2002) 323–335. doi:10.1023/A:1021998908811.

[8] J. J. Moré, Generalizations of the trust region problem, Optimization Methods and Software 2 (3-4) (1993) 189–209. doi:10.1080/10556789308805542.

[9] H. Tuy, N. T. Hoai-Phuong, A robust algorithm for quadratic optimization under quadratic constraints, Journal of Global Optimization 37 (4) (2007) 557–569. doi:10.1007/s10898-006-9063-7.

[10] F. Critchley, On the minimisation of a positive definite quadratic form under a quadratic constraint: analytical solution and statistical applications, Statistics Research Report 165, Department of Statistics, University of Warwick (1989).

[11] J. C. Gower, G. B. Dijksterhuis, Procrustes Problems, no. 30 in Oxford Statistical Science Series, Oxford University Press, 2004.

[12] C. J. Albers, J. C. Gower, A general approach to handling missing values in procrustes analysis, Advances in Data Analysis and Classification 4 (4) (2010) 223–237. doi:10.1007/s11634-010-0077-0.

[13] R. W. Newcomb, On the simultaneous diagonalization of two semidefinite matrices, Quarterly Journal of Applied Mathematics 19 (1960) 144–146.