

Worst-Case-Optimal Estimation of a Quadratic Form by Quadratic Functionals of its Linear Observations

Simon Foucart^{a,1}

This manuscript was compiled on October 10, 2026

In this article, it is established that, in order for a quadratic form of an unknown object to be estimated optimally, in a worst-case sense, from the *a priori* knowledge that the object belongs to an hyperellipsoid and the *a posteriori* availability of linear observations, a quadratic functional of these observations can typically be used. Working in finite dimensions, a computational construction of an optimal quadratic functional is uncovered. It is accompanied by a certificate guaranteeing that the returned quadratic functional is indeed optimal among all estimation functionals. The certificate appears to always be validated on generic numerical examples. Importantly, the results are further extended to take into account some observation errors modeled deterministically.

optimal recovery | semidefinite programming | S-lemma | S-procedure

1. Framing of the Problem

The adage that “optimal estimation of linear functionals can be achieved with linear functionals” is a cornerstone of Optimal Recovery (1), and in turn of Information-Based Complexity (2). The result is credited to Smolyak (3), although it was published later by his doctoral advisor Bakhvalov (4). The purpose of this article is to establish a similar result where “quadratic” replaces “linear”. One starts by formalizing precisely what is meant through a brief presentation of the general Optimal Recovery problem, to be narrowed down later.

The goal here is to estimate a quantity $\Gamma(f)$ depending on an element f from a vector space F —often, but not necessarily, a space of functions. This element f is only partially known through:

- some *a priori* information, expressed as

$$f \in \mathcal{K} \quad \text{for some model set } \mathcal{K} \subseteq F;$$

- some *a posteriori* information resulting from linear observations $y_i = \lambda_i(f)$, $i = 1, \dots, m$, summarized as

$$y = \Lambda f \quad \text{for some linear map } \Lambda : F \rightarrow \mathbb{R}^m.$$

The estimation of the quantity of interest $\Gamma : F \rightarrow G$, which is assumed to be known and accessible, is encapsulated by a map $\Delta : \mathbb{R}^m \rightarrow G$, whose performance is measured in a worst-case framework

- either via $\text{lwce}_y(\Delta(y))$, when $y \in \mathbb{R}^m$ is considered fixed and

$$\text{lwce}_y(g) := \sup\{\|\Gamma(f) - g\| : f \in \mathcal{K}, \Lambda f = y\};$$

- or, when $y \in \mathbb{R}^m$ is unfixed, via

$$\text{gwce}(\Delta) := \sup\{\|\Gamma(f) - \Delta(\Lambda f)\| : f \in \mathcal{K}\}.$$

Significance

In Learning Theory, one typically faces the task of estimating a quantity depending on an object which was acquired via linear observations. The estimation procedure can be assessed from a worst-case perspective provided some prior knowledge about the object is available as well. One could then pursue the best estimation procedures, with the hope that they shall be simple when the quantity to be estimated is itself simple. In the field of optimal recovery, it has been known for decades that the hope can be realized when simple refers to linear. But much less is known in a nonlinear setting. The importance of this article is the finding that the hope can also be realized when simple refers to quadratic.

Author affiliations: ^aDepartment of Mathematics, Texas A&M University, College Station, TX 77843

¹To whom correspondence should be addressed. E-mail: foucart@tamu.edu

The grail of Optimal Recovery is to find estimation maps Δ^{opti} that minimize these notions of worst-case error, either locally, i.e., $\text{lwce}_y(\Delta^{\text{opti}}(y)) \leq \text{lwce}_y(g)$ for all $g \in G$, or globally, i.e., $\text{gwce}(\Delta^{\text{opti}}) \leq \text{gwce}(\Delta)$ for all $\Delta : \mathbb{R}^m \rightarrow G$. The adage from the beginning refers to this global setting. It conveys that, when the model set $\mathcal{K}$ is convex and symmetric about the origin, if $\Gamma = \gamma : F \rightarrow \mathbb{R}$ is a linear functional—lowercase Greek letters are used for functionals—then it is possible to select an optimal estimation map $\delta^{\text{opti}} : \mathbb{R}^m \rightarrow \mathbb{R}$ as a linear functional. In fact, the requirement that $\mathcal{K}$ is symmetric about the origin can be dropped if one allows for estimation functionals that are affine instead of linear, as originally shown by Sukharev (5).

Much less is known when the functional γ to be estimated is not linear. Assuming again that $\mathcal{K}$ is a convex set, it was recently established in (6) that, if $\gamma(f)$ is the largest of several linear functionals of f (or even the k th largest), then it is possible to choose a globally optimal estimation functional δ^{opti} as a maximum of affine functionals (or as the maximum of minima of affine functionals). The present article focuses on the case where γ is a quadratic functional, which has previously been examined in the Statistics literature, see e.g. (7–10). The Statistical Estimation problem considered there is similar in many aspects to the Optimal Recovery problem considered here, and correspondence between the two is discussed in (11). The main difference lies in the assessment of the performance of a recovery map $\delta : \mathbb{R}^m \rightarrow \mathbb{R}$ —called estimator in the Statistics literature—by way of a worst-case mean-squared error $\text{wcms}(\delta) := \sup\{\mathbb{E}[|\gamma(f) - \delta(\Lambda f + e)|^2] : f \in \mathcal{K}\}$, where $e \sim \mathcal{N}(0, \sigma^2 I_m)$ represents stochastic observation errors. The results obtained here parallel the ones obtained in the article (7), say, which targeted the estimation of the quadratic form $\int_{\mathbb{T}} [f^{(k)}]^2$ of a periodic function f , based on the *a priori* information that $\int_{\mathbb{T}} [f^{(\ell)}]^2 \leq 1$ with $\ell > k$ and the *a posteriori* information consisting of noisy evaluations $y_i = f(t_i) + e_i$ of f at m equidistant points on the trigonometric circle $\mathbb{T}$. This article uncovered a particular estimator δ_0 such that $\text{wcms}(\delta_0) \leq C \inf_{\delta} \text{wcms}(\delta)$, with a near-optimality constant $C \rightarrow 1$ as $m \rightarrow \infty$. It relied on a reduction to the case where γ is a linear functional, which had been resolved in (11) (see also (12) for a fresh take). The present article will avoid such a reduction, while dealing with no observation errors or with deterministic rather than stochastic ones. It will also handle less restrictive quantities of interest, model sets, and observation maps and will produce genuinely optimal recovery maps (near-optimality factor $C = 1$), albeit not in closed form. The perspective compares to that of (10), where the authors advocate an “operational” approach: enable more general situations by sacrificing explicit solutions for computed ones.

To be more specific, for an arbitrary quadratic form $\gamma(f) = \langle Af, f \rangle$ defined by a self-adjoint operator A on the space F —from now on considered a Hilbert space and thus denoted by H —and for an ellipsoidal model set $\mathcal{K}$, it will be shown in Section 3 that it is possible to select a globally optimal recovery functional $\delta^{\text{opti}} : \mathbb{R}^m \rightarrow \mathbb{R}$ expressed as $\delta^{\text{opti}}(y) = \langle By, y \rangle + \beta$ for some symmetric matrix $B \in \mathbb{R}^{m \times m}$ and some scalar $\beta \in \mathbb{R}$. The result (Theorem 4) is in fact conditional on the

verification of property of the optimal quadratic estimation functional. In Section 4, a computational construction of this estimation functional is presented (Theorem 5) and complemented by a streamlining of the verification step (Theorem 6). The main ingredient is Yakubovich S-lemma, which the survey (13) discusses, along with related results to be introduced in due time. For instance, in Section 5, Polyak S-procedure will allow for the computational construction tied to an extension of the main results (Theorems 7 and 8) covering the presence of deterministic observation errors. In Section 6, Finsler lemma will also be used to highlight the computation, redundant in retrospect, of a key lower bound (Proposition 9). All the content just described shall be preceded by Section 2, which quickly settles the local setting.

2. Locally Optimal Recovery Functionals

In the scenario considered here, the quantity of interest γ can be a quadratic functional on a Hilbert space H , instead of merely a quadratic form, i.e.,

$$\gamma(f) = \langle Af, f \rangle + 2\langle a, f \rangle + \alpha, \quad f \in H, \quad [1]$$

for some self-adjoint operator A on H , some vector $a \in H$, and some scalar $\alpha \in \mathbb{R}$. The *a posteriori* information, considered fixed in this section, is given by $y = \Lambda(f)$ for some unrestricted linear map $\Lambda : H \rightarrow \mathbb{R}^m$, while the model set $\mathcal{K}$ describing the *a priori* information is specifically given as

$$\mathcal{K} = \{f \in H : \|Tf\| \leq 1\} \quad [2]$$

for some linear operator $T : H \rightarrow H$. The goal is to determine the so-called Chebyshev center of $\gamma(\mathcal{K}_y)$, where the set $\mathcal{K}_y := \{f \in \mathcal{K} : \Lambda(f) = y\}$ collects all the model- and data-consistent elements f . In other words, one aims at determining the minimizer g^{cheb} of the local worst-case error $\text{lwce}_y(g) = \sup\{|\gamma(f) - g| : \|Tf\| \leq 1, \Lambda(f) = y\}$. Since the set $\mathcal{K}_y$ is connected, the set $\gamma(\mathcal{K}_y)$ is a connected subset of the real line, i.e., an interval. It is then clear, and very much folklore in Optimal Recovery and Information-Based Complexity, that the Chebyshev center and Chebyshev radius are obtained as

$$g^{\text{cheb}} = \frac{\sigma + \iota}{2}, \quad \text{lwce}_y(g^{\text{cheb}}) = \frac{\sigma - \iota}{2}, \quad [3]$$

where

$$\sigma := \sup_{f \in H} \{\gamma(f) : \|Tf\| \leq 1, \Lambda f = y\}, \quad [4]$$

$$\iota := \inf_{f \in H} \{\gamma(f) : \|Tf\| \leq 1, \Lambda f = y\}. \quad [5]$$

For the problem to make sense, i.e., for $\sigma < +\infty$ and $\iota > -\infty$ independently of γ , it is assumed that

$$\ker(\Lambda) \cap \ker(T) = \{0\}, \quad [6]$$

otherwise a nonzero $h \in \ker(\Lambda) \cap \ker(T)$ and a fixed $f_0 \in \mathcal{K}_y$ would generate the unbounded family $f_t = f_0 + th$, $t \in \mathbb{R}$, of elements belonging to $\mathcal{K}_y$. Elementary as the above considerations are, computing the quantities [4] and [5] is not completely straightforward. This is where the S-lemma comes into play. Here is its statement for

$$\begin{aligned} \underset{c \geq 0}{\text{minimize } d} \quad & \text{subject to } \left[\begin{array}{c|c} P(cT^*T - A)P & -P(Af_y + a) \\ \hline -(Af_y + a)^*P & d - c(1 - \|Tf_y\|^2) - \langle Af_y, f_y \rangle - 2\langle a, f_y \rangle - \alpha \end{array} \right] \succeq 0, \quad [7] \\ \underset{c \geq 0}{\text{maximize } d} \quad & \text{subject to } \left[\begin{array}{c|c} P(cT^*T + A)P & P(Af_y + a) \\ \hline (Af_y + a)^*P & -d - c(1 - \|Tf_y\|^2) + \langle Af_y, f_y \rangle + 2\langle a, f_y \rangle + \alpha \end{array} \right] \succeq 0, \quad [8] \end{aligned}$$

where $P = \text{Id}_H - \Lambda^*(\Lambda\Lambda^*)^{-1}\Lambda$ and $f_y = \text{argmin}\{\|Tf\|^2 : \Lambda(f) = y\}$.

Fig. 1. The semidefinite programs to construct the locally optimal estimation map (see Theorem 1) in the absence of observation errors.

a finite-dimensional Hilbert space H —note that, since an optimization program eventually needs to be solved, assuming finite-dimensionality for the rest of this article is not considered a restriction.

Yakubovich S-lemma: given two quadratic functionals $q_i(h) = \langle R_i h, h \rangle + 2\langle r_i, f \rangle + \rho_i$, $i = 0, 1$, defined on a finite-dimensional Hilbert space H , the equivalence

$$[q_0(h) \leq 0 \text{ whenever } q_1(h) \leq 0] \iff [\exists c \geq 0 : q_0 \leq c q_1]$$

holds provided $q_1(h) < 0$ for some $h \in H$.

In the present scenario, this tool allows one to uncover a computational recipe for the construction of the locally optimal estimation functional—also called central algorithm in Optimal Recovery and Information-Based Complexity.

Theorem 1. For a finite-dimensional Hilbert space H , a linear observation map $\Lambda : H \rightarrow \mathbb{R}^m$, and the model set $\mathcal{K}$ of [2] obeying the assumption [6], the locally optimal recovery map for the estimation of the quadratic functional [1] is given by $\delta^{\text{cheb}} : y \in \mathbb{R}^m \mapsto (d' + d'')/2 \in \mathbb{R}$, where d' and d'' are, respectively, the optimal values of the semidefinite programs [7] and [8] displayed in Fig. 1

Proof. As mentioned above, the locally optimal estimation functional δ^{opti} associates to each $y \in \mathbb{R}^m$ the Chebyshev center g^{cheb} of $\gamma(\mathcal{K}_y)$ as given in [3]. Accordingly, the task at hand consists in showing that σ coincides with d' and that ι coincides with d'' . Focusing first on $\sigma = \sup\{\gamma(f) : \|Tf\| \leq 1, \Lambda(f) = y\}$, note that it can be expressed as

$$\sigma = \inf_d d \quad \text{subject to } \langle Af, f \rangle + 2\langle a, f \rangle + \alpha \leq d \quad \text{whenever } \|Tf\|^2 \leq 1, \Lambda f = y.$$

Writing $f = f_y + h$ with $f_y = \text{argmin}\{\|Tf\|^2 : \Lambda(f) = y\}$ and $h \in \ker(\Lambda)$, one views the above constraint as the nonpositivity of a quadratic functional of h occurring as a consequence of the nonpositivity of another quadratic functional of h . Thus, the S-lemma will apply provided the latter quadratic functional can take a strictly negative value. And indeed, the strict inequality $\|Tf_y\|^2 < 1$ holds: select some $f \neq f_y$ with $\|Tf\|^2 \leq 1$ and $\Lambda f = y$, so that $f - f_y \in \ker(\Lambda)$, hence $T(f - f_y) \neq 0$ by [6], and then the condition $T(f_y) \perp T(\ker(\Lambda))$ characterizing the optimality of f_y yields $\|Tf_y\|^2 = \|Tf\|^2 - \|T(f - f_y)\|^2 < 1$. Therefore, the S-lemma ensures that the above constraint is equivalent to the existence of $c \geq 0$ such that

$$\langle A(f_y + h), f_y + h \rangle + 2\langle a, f_y + h \rangle + \alpha - d \leq c(\|T(f_y + h)\|^2 - 1)$$

for all $h \in \ker(\Lambda)$. Rearranging and also making use of the orthogonality condition $\langle Tf_y, Th \rangle = 0$, this reads

$$\langle (cT^*T - A)h, h \rangle - 2\langle Af_y + a, h \rangle + d - c(1 - \|Tf_y\|^2) - \langle Af_y, f_y \rangle - 2\langle a, f_y \rangle - \alpha \geq 0$$

for all $h \in \ker(\Lambda)$. More succinctly, with P representing the orthogonal projection onto $\ker(\Lambda)$ (so that $\text{Id}_H - P$ is the orthogonal projection onto $\ker(\Lambda)^\perp = \text{ran}(\Lambda^*)$, easily leading to the expression of P announced in Fig. 1), this $c \geq 0$ must satisfy **constraint** _{A, a, α} (c, d):

$$\left[\begin{array}{c|c} P(cT^*T - A)P & -P(Af_y + a) \\ \hline -(Af_y + a)^*P & d - c(1 - \|Tf_y\|^2) - \gamma(f_y) \end{array} \right] \succeq 0.$$

All in all, incorporating $c \geq 0$ as an optimization variable, one arrives at

$$\sigma = \inf_{c \geq 0, d} d \quad \text{subject to } \text{constraint}_{A, a, \alpha}(c, d),$$

so that σ is indeed equal to the optimal value d' of the semidefinite program [7]. As for ι , one notices that $-\iota = \sup\{(-\gamma)(f) : \|Tf\| \leq 1, \Lambda(f) = y\}$, so it admits an expression similar to that of σ , up to the change $(A, a, \alpha) \leftrightarrow (-A, -a, -\alpha)$, namely $-\iota = \inf\{d : c \geq 0, d \text{ satisfy } \text{constraint}_{-A, -a, -\alpha}(c, d)\}$. One obtains $\iota = \sup\{-d : c \geq 0, d \text{ satisfy } \text{constraint}_{-A, -a, -\alpha}(c, d)\}$. Changing d to $-d$ shows that ι is equal to the optimal value d'' of the semidefinite program [8]. $\square$

Remark. The construction of Theorem 1 is implemented in the reproducible MATLAB file accompanying this article*. There, it is exploited to show that the locally optimal functional is not quadratic, even when $a = 0$. This indicates that, for quadratic functionals in the local setting, the adage fails. A conceptual difference between the quadratic and linear cases is thereby revealed, since it is known that, when the model set is ellipsoidal, the central algorithm for the estimation of a linear functional γ (more generally a linear map Γ) actually depends linearly on the observations.

3. Conditional Globally Optimal Recovery Functionals

Generally speaking, a drawback of the local setting is that an expensive optimization program needs to be solved anew for each data vector $y \in \mathbb{R}^m$ —except in favorable circumstances, e.g. when Γ is a linear map between Hilbert spaces and $\mathcal{K}$ is an ellipsoid, in which case the Chebyshev

* Available on the author's webpage and the repository <https://github.com/foucart/COR>.

center depends linearly on y . This aspect partly explains the predominance of the global setting, considered in this section. Here, the quantity of interest γ is restricted to be a quadratic form, i.e.,

$$\gamma(f) = \langle Af, f \rangle, \quad f \in H, \quad [9]$$

for a self-adjoint operator A on a Hilbert space H . As for the restrictions on *a posteriori* and *a priori* information, they are minimal: the observation map $\Lambda : H \rightarrow \mathbb{R}^m$ is linear and the model set $\mathcal{K}$ is compact, convex, and symmetric about the origin—it will be particularized in the next sections. The main result of this section is stated after the two preparatory lemmas below. The first one recalls a well known lower bound for the global worst-case error of any estimation functional. To use the terminology of Information-Based Complexity, this lower bound is half the diameter of information.

Lemma 2. With $\Lambda : F \rightarrow \mathbb{R}^m$, $\mathcal{K} \subseteq F$, and $\gamma : F \rightarrow \mathbb{R}$ all being arbitrary, any estimation functional $\delta : \mathbb{R}^m \rightarrow \mathbb{R}$ satisfies

$$\text{gwce}(\delta) \geq \text{lb} := \sup_{\substack{f', f'' \in \mathcal{K} \\ \Lambda f' = \Lambda f''}} \frac{\gamma(f') - \gamma(f'')}{2}. \quad [10]$$

The second lemma concerns the global worst-case error of estimation functionals taken specifically to be quadratic.

Lemma 3. With an arbitrary linear map Λ from a Hilbert space H to $\mathbb{R}^m$, if $\mathcal{K} \subseteq H$ is convex and symmetric about the origin and if the functional $\gamma : H \rightarrow \mathbb{R}$ obeys $\gamma(-f) = \gamma(f)$ for all $f \in H$, then any quadratic functional of the form $\delta_{B,b,\beta}(y) = \langle By, y \rangle + 2\langle b, y \rangle + \beta$, $y \in \mathbb{R}^m$, satisfies

$$\text{gwce}(\delta_{B,b,\beta}) \geq \text{gwce}(\delta_B) = \frac{\sigma_B - \iota_B}{2},$$

where δ_B is defined by $\delta_B(y) = \langle By, y \rangle + \frac{\sigma_B + \iota_B}{2}$, $y \in \mathbb{R}^m$, and where

$$\sigma_B := \sup_{f \in \mathcal{K}} (\gamma(f) - \langle B(\Lambda f), \Lambda f \rangle), \quad [11]$$

$$\iota_B := \inf_{f \in \mathcal{K}} (\gamma(f) - \langle B(\Lambda f), \Lambda f \rangle). \quad [12]$$

The main result now establishes the global optimality of quadratic estimation functionals, conditionally on a property to be assessed post computation.

Theorem 4. Given a Hilbert space H , a linear observation map $\Lambda : H \rightarrow \mathbb{R}^m$, and a model set $\mathcal{K} \subseteq H$ which is compact, convex, and symmetric about the origin, let a symmetric matrix $B_0 \in \mathbb{R}^{m \times m}$ be a minimizer of $\text{gwce}(\delta_B) = (\sigma_B - \iota_B)/2$. If the extremizers of the expressions [11] and [12] for σ_{B_0} and ι_{B_0} are unique (up to a sign), then the quadratic functional δ_{B_0} is globally optimal for the estimation of the quadratic form [9], i.e.,

$$\text{gwce}(\delta_{B_0}) \leq \text{gwce}(\delta) \quad \text{for all } \delta : \mathbb{R}^m \rightarrow \mathbb{R}.$$

It is now time to justify the above statements, starting with the lemmas.

Proof of Lemma 2. For $f', f'' \in \mathcal{K}$ with $\Lambda f' = \Lambda f''$, let y denote this common value. One has

$$\begin{aligned} \frac{\gamma(f') - \gamma(f'')}{2} &= \frac{1}{2}(\gamma(f') - \delta(y)) + \frac{1}{2}(\delta(y) - \gamma(f'')) \\ &\leq \frac{1}{2} \sup_{f \in \mathcal{K}} |\gamma(f) - \delta(\Lambda f)| + \frac{1}{2} \sup_{f \in \mathcal{K}} |\gamma(f) - \delta(\Lambda f)| \\ &= \text{gwce}(\delta). \end{aligned}$$

The announced result follows by taking the supremum over these f' and f'' . $\square$

Proof of Lemma 3. As a first step, one observes that the global worst-case error of $\delta_{B,b,\beta}$ is minimized when $b = 0$. To justify this claim, notice that, for any $y \in \mathbb{R}^m$, $\delta_{B,0,\beta}(y) = (\delta_{B,b,\beta}(y) + \delta_{B,b,\beta}(-y))/2$ and that, for $f \in \mathcal{K}$, and hence $-f \in \mathcal{K}$, one has

$$\begin{aligned} |\gamma(f) - \delta_{B,0,\beta}(\Lambda f)| &= \frac{1}{2} |[\gamma(f) - \delta_{B,b,\beta}(\Lambda f)] + [\gamma(-f) - \delta_{B,b,\beta}(\Lambda(-f))]| \\ &\leq \frac{1}{2} |\gamma(f) - \delta_{B,b,\beta}(\Lambda f)| + \frac{1}{2} |\gamma(-f) - \delta_{B,b,\beta}(\Lambda(-f))| \\ &\leq \frac{1}{2} \text{gwce}(\delta_{B,b,\beta}) + \frac{1}{2} \text{gwce}(\delta_{B,b,\beta}) = \text{gwce}(\delta_{B,b,\beta}). \end{aligned}$$

Taking the supremum over $f \in \mathcal{K}$ yields the desired inequality $\text{gwce}(\delta_{B,0,\beta}) \leq \text{gwce}(\delta_{B,b,\beta})$.

As a second step, one establishes that $(\sigma_B + \iota_B)/2$ is the optimal value of the parameter β . Indeed, considering f^+, f^- achieving (or coming arbitrarily close to) the supremum and infimum defining σ_B and ι_B , one has

$$\begin{aligned} \text{gwce}(\delta_{B,0,\beta}) &= \sup_{f \in \mathcal{K}} |\gamma(f) - \langle B(\Lambda f), \Lambda f \rangle - \beta| \\ &\geq \begin{cases} \gamma(f^+) - \langle B(\Lambda f^+), \Lambda f^+ \rangle - \beta = \sigma_B - \beta, \\ -\gamma(f^-) + \langle B(\Lambda f^-), \Lambda f^- \rangle + \beta = -\iota_B + \beta, \end{cases} \end{aligned}$$

so the lower estimate $\text{gwce}(\delta_{B,0,\beta}) \geq (\sigma_B - \iota_B)/2$ follows by averaging these two values. This lower estimate is also achieved at δ_B , i.e., with $\beta = \beta^{\text{opti}} := (\sigma_B + \iota_B)/2$, since then one has, for any $f \in \mathcal{K}$,

$$\begin{aligned} \gamma(f) - \langle B(\Lambda f), \Lambda f \rangle - \beta^{\text{opti}} &= \begin{cases} \leq \sigma_B - \frac{\sigma_B + \iota_B}{2} = +\frac{\sigma_B - \iota_B}{2}, \\ \geq \iota_B - \frac{\sigma_B + \iota_B}{2} = -\frac{\sigma_B - \iota_B}{2}. \end{cases} \end{aligned}$$

Taken together, these two inequalities mean that $\text{gwce}(\delta_B) = \sup_{f \in \mathcal{K}} |\gamma(f) - \langle B(\Lambda f), \Lambda f \rangle - \beta^{\text{opti}}|$ does not exceed $(\sigma_B - \iota_B)/2$, as desired. $\square$

Lemmas 2 and 3 being settled, one moves to the proof of the main result.

Proof of Theorem 4. The objective is to justify that $\text{gwce}(\delta_{B_0}) \leq \text{lb}$, where the global worst-case error of a quadratic functional δ_B defined in Lemma 3 is expressed, with $A_B := A - \Lambda^* B \Lambda$, as

$$\text{gwce}(\delta_B) = \sup_{f', f'' \in \mathcal{K}} \frac{\langle A_B f', f' \rangle - \langle A_B f'', f'' \rangle}{2},$$

and the lower bound lb from Lemma 2 is expressed as

$$\text{lb} = \sup_{f', f'' \in \mathcal{K}} \frac{\langle Af', f' \rangle - \langle Af'', f'' \rangle}{2} \quad \text{s.t. } \Lambda f' = \Lambda f''.$$

The argument is now divided in two parts:

(i) with $\mathcal{K}'_0$ and $\mathcal{K}''_0$ denoting the extremizer sets for σ_{B_0} and ι_{B_0} , i.e., $\mathcal{K}'_0 := \{f' \in \mathcal{K} : \langle A_{B_0} f', f' \rangle = \sigma_{B_0}\}$ and $\mathcal{K}''_0 := \{f'' \in \mathcal{K} : \langle A_{B_0} f'', f'' \rangle = \iota_{B_0}\}$,

$$\text{gwce}(B_0) = \inf_B \sup_{\substack{f' \in \mathcal{K}'_0 \\ f'' \in \mathcal{K}''_0}} \frac{\langle A_B f', f' \rangle - \langle A_B f'', f'' \rangle}{2}. \quad [13]$$

(ii) when $\mathcal{K}'_0 = \{\pm f'_0\}$ and $\mathcal{K}''_0 = \{\pm f''_0\}$, if $\Lambda f'_0 \neq \pm \Lambda f''_0$, then there exists B such that

$$\sup_{\substack{f' \in \mathcal{K}'_0 \\ f'' \in \mathcal{K}''_0}} \frac{\langle A_B f', f' \rangle - \langle A_B f'', f'' \rangle}{2} < \text{gwce}(B_0). \quad [14]$$

As a result, because the two conclusions are incompatible and given the uniqueness assumption of the extremizers f'_0 and f''_0 , one must have $\Lambda f'_0 = \pm \Lambda f''_0$. It follows that

$$\begin{aligned} \text{gwce}(B_0) &= \frac{\langle (A - \Lambda^* B \Lambda) f'_0, f'_0 \rangle - \langle (A - \Lambda^* B \Lambda) f''_0, f''_0 \rangle}{2} \\ &= \frac{\langle A f'_0, f'_0 \rangle - \langle A f''_0, f''_0 \rangle}{2} \leq \text{lb}, \end{aligned}$$

which is the desired inequality.

Thus, it remains to establish (i) and (ii). Starting with the latter, notice that the left-hand side of [14] with $B = B_0 + C$ is

$$\text{gwce}(B_0) - \frac{1}{2} \text{tr}(C \Lambda(f'_0 f'^*_0 - f''_0 f''^*_0) \Lambda^*),$$

which can be made smaller than $\text{gwce}(B_0)$ by taking $C = \Lambda(f'_0 f'^*_0 - f''_0 f''^*_0) \Lambda^*$, provided it is nonzero. But if $\Lambda f'_0 f'^*_0 \Lambda^* = \Lambda f''_0 f''^*_0 \Lambda^*$, then $\Lambda f'_0$ and $\Lambda f''_0$ would be equal in direction by looking at the range and equal in norm by looking at the trace, contradicting $\Lambda f'_0 \neq \pm \Lambda f''_0$.

As for (ii), it is a consequence of a general Kolmogorov-type criterion from Approximation Theory. Since this is a little-known result, an argument tailored to the current situation is provided. By way of contradiction, assume that [13] fails, i.e., that there exists B_1 such that

$$\mu := \sup_{\substack{f' \in \mathcal{K}'_0 \\ f'' \in \mathcal{K}''_0}} \frac{\langle A_{B_1} f', f' \rangle - \langle A_{B_1} f'', f'' \rangle}{2} < \text{gwce}(B_0) =: \nu.$$

For $\tau \in [0, 1]$, define $B_\tau := (1 - \tau)B_0 + \tau B_1$, and additionally, for $(f', f'') \in \mathcal{K} \times \mathcal{K}$,

$$\Theta((f', f''), \tau) := \frac{\langle A_{B_\tau} f', f' \rangle - \langle A_{B_\tau} f'', f'' \rangle}{2}.$$

Note that, if $(f', f'') \notin \mathcal{K}'_0 \times \mathcal{K}''_0$, then $\Theta((f', f''), 0) < \nu$, while if $(f', f'') \in \mathcal{K}'_0 \times \mathcal{K}''_0$, then $\Theta((f', f''), 0) = \nu$ and $\Theta((f', f''), 1) \leq \mu$. Picking some $\bar{\mu} \in (\mu, \nu)$, the set $\mathcal{O} := \{(f', f'') \in \mathcal{K} \times \mathcal{K} : \Theta((f', f''), 1) < \bar{\mu}\}$ is an open neighborhood of $\mathcal{K}'_0 \times \mathcal{K}''_0$. By compactness of the complement of $\mathcal{O}$ in $\mathcal{K} \times \mathcal{K}$, one observes that

$\underline{\nu} := \max\{\Theta((f', f''), 0) : (f', f'') \in (\mathcal{K} \times \mathcal{K}) \setminus \mathcal{O}\}$ satisfies $\underline{\nu} < \nu$. Next, considering some $\tau > 0$ with $\tau < (\nu - \underline{\nu})/\kappa$, $\kappa := \|B_1 - B_0\| \max\{\|A f\|^2 : f \in \mathcal{K}\}$, one has

• if $(f', f'') \notin \mathcal{O}$, then

$$\begin{aligned} \Theta((f', f''), \tau) &= \Theta((f', f''), 0) \\ &\quad + \tau \frac{\langle (B_1 - B_0) \Lambda f', \Lambda f' \rangle - \langle (B_1 - B_0) \Lambda f'', \Lambda f'' \rangle}{2} \\ &\leq \underline{\nu} + \tau \kappa < \nu; \end{aligned}$$

• if $(f', f'') \in \mathcal{O}$, then

$$\begin{aligned} \Theta((f', f''), \tau) &= (1 - \tau) \Theta((f', f''), 0) + \tau \Theta((f', f''), 1) \\ &\leq (1 - \tau) \nu + \tau \bar{\mu} < \nu. \end{aligned}$$

All in all, one deduces that the above choice of τ leads to $\max\{\Theta((f', f''), \tau) : (f', f'') \in \mathcal{K} \times \mathcal{K}\} < \nu$, which reads $\text{gwce}(B_\tau) < \text{gwce}(B_0)$, thus producing the required contradiction. The proof is now complete. $\square$

4. Computation and Validation

In this section, the compact, convex, and symmetric about the origin model set is specified to be the ellipsoid

$$\mathcal{K} = \{f \in H : \|Tf\| \leq 1\}, \quad T \text{ isomorphism}. \quad [15]$$

The goal here is twofold: firstly, to provide a practical construction of a quadratic estimation functional with minimal global worst-case error; secondly, to certify that this quadratic functional is optimal among all possible estimation functionals. The construction involves semidefinite programming again.

Theorem 5. In a finite-dimensional Hilbert space H , for the estimation of the quadratic form [9] using a linear observation map $\Lambda : H \rightarrow \mathbb{R}^m$ and the model set $\mathcal{K}$ of [15], a globally optimal quadratic recovery functional is given as

$$\delta_{B_0}(y) = \langle B_0 y, y \rangle + \beta_0, \quad y \in \mathbb{R}^m,$$

where the symmetric matrix $B_0 \in \mathbb{R}^{m \times m}$ is obtained as solution to the semidefinite program [16] displayed in Fig. 2. The solutions $\bar{c}', \bar{c}'' \in \mathbb{R}$ yield the optimal parameter $\beta_0 = (\bar{c}' - \bar{c}'')/2$, as well as the minimal global worst-case estimation error $(\bar{c}' + \bar{c}'')/2$.

Proof. By Lemma 3, a globally optimal quadratic estimation functional can be selected of the form $\delta_B(y) = \langle B y, y \rangle + \beta$, $y \in \mathbb{R}^m$. The optimal β and the global worst-case error depend on the two quantities σ_B and ι_B via

$$\beta = \frac{\sigma_B + \iota_B}{2} \quad \text{and} \quad \text{gwce}(\delta_B) = \frac{\sigma_B - \iota_B}{2}.$$

The quantity σ_B , for instance, can be converted with the help of the S-lemma as follows (the strict feasibility

$$\begin{aligned} & \underset{\substack{c', c'' \geq 0 \\ B \text{ symmetric}}}{\text{minimize}} \quad \frac{c' + c''}{2} \quad \text{subject to} \quad \begin{cases} c'T^*T - A + \Lambda^*B\Lambda \succeq 0, \\ c''T^*T + A - \Lambda^*B\Lambda \succeq 0. \end{cases} \end{aligned} \quad [16]$$

Fig. 2. The semidefinite program to construct a globally optimal quadratic estimation map (see Theorem 5) in the absence of observation errors.

condition being easily met by $f = 0$):

$$\begin{aligned} \sigma_B &:= \sup_{\|Tf\| \leq 1} \langle (A - \Lambda^*B\Lambda)f, f \rangle \\ &= \inf_d d : \langle (A - \Lambda^*B\Lambda)f, f \rangle \leq d \text{ whenever } \|Tf\|^2 \leq 1 \\ &= \inf_d d : \exists c \geq 0 \text{ such that, for all } f \in H, \\ &\quad \langle (A - \Lambda^*B\Lambda)f, f \rangle - d \leq c(\|Tf\|^2 - 1) \\ &= \inf_{c \geq 0} d : \text{for all } f \in H, \\ &\quad \langle (cT^*T - (A - \Lambda^*B\Lambda))f, f \rangle + d - c \geq 0. \end{aligned}$$

The constraint decouples as $\langle (cT^*T - A + \Lambda^*B\Lambda)f, f \rangle \geq 0$ for all $f \in H$ and $d - c \geq 0$, i.e., as $cT^*T - A + \Lambda^*B\Lambda \succeq 0$ and $d \geq c$, so the minimal admissible value for d is c , which implies that

$$\sigma_B = \inf_{c \geq 0} c : cT^*T - A + \Lambda^*B\Lambda \succeq 0.$$

The quantity $\iota_B = \inf\{\langle (A - \Lambda^*B\Lambda)f, f \rangle : \|Tf\| \leq 1\}$ —or rather $-\iota_B$ —is likewise converted to

$$-\iota_B = \inf_{c \geq 0} c : cT^*T + A - \Lambda^*B\Lambda \succeq 0.$$

Combining these two expressions with the fact that $\text{gwce}(\delta_B) = (\sigma_B - \iota_B)/2$ justifies that, for a fixed B ,

$$\begin{aligned} \text{gwce}(\delta_B) &= \inf_{c', c'' \geq 0} \frac{c' + c''}{2} \\ \text{s.t. to } &\begin{cases} c'T^*T - A + \Lambda^*B\Lambda \succeq 0, \\ c''T^*T + A - \Lambda^*B\Lambda \succeq 0. \end{cases} \end{aligned} \quad [17]$$

Further minimizing over symmetric matrices $B \in \mathbb{R}^{m \times m}$ —which is computationally feasible thanks to the linear dependence on B of the constraints—gives rise to the semidefinite program announced in [16]. Note that the solutions $\bar{c}', \bar{c}'' \geq 0$ correspond to σ_B and $-\iota_B$ for the optimal B , justifying the final part of the statement. $\square$

After having produced B_0 , the next step is to make sure that δ_{B_0} is optimal not only among quadratic, but among all estimation functionals. This shall be done via an eigencomputation involving B_0 .

Theorem 6. Under the setting of Theorem 5, once B_0 has been computed, if the largest and smallest eigenvalues of $(T^*T)^{-1/2}(A - \Lambda^*B_0\Lambda)(T^*T)^{-1/2}$ are strictly positive and strictly negative, respectively, and both simple, then

$$\text{gwce}(\delta_{B_0}) \leq \text{gwce}(\delta) \quad \text{for all } \delta : \mathbb{R}^m \rightarrow \mathbb{R}.$$

Proof. Replacing T by $(T^*T)^{1/2}$ if necessary, it is assumed in the proof that T is self-adjoint. By Theorem 4, it

suffices to establish that the extremizers for σ_{B_0} and ι_{B_0} are unique (up to a sign), where, say,

$$\sigma_{B_0} = \sup_{\|Tf\| \leq 1} \langle (A - \Lambda^*B_0\Lambda)f, f \rangle \quad [18]$$

$$= \sup_{\|g\| \leq 1} \langle T^{-1}(A - \Lambda^*B_0\Lambda)Tg, g \rangle. \quad [19]$$

Notice that $\sigma_{B_0} = 0$ if $T^{-1}(A - \Lambda^*B_0\Lambda)T^{-1} \preceq 0$ and that $\sigma_{B_0} = \text{eig}_{\max}(T^{-1}(A - \Lambda^*B_0\Lambda)T^{-1}) > 0$ otherwise, the latter being in effect here. Based on the one-to-one correspondence $g = Tf$ between extremizers of [18] and extremizers of [19], which are exactly the unit-norm eigenvectors of $T^{-1}(A - \Lambda^*B_0\Lambda)T^{-1}$ for its largest eigenvalue, the assumption that this eigenvalue has multiplicity one establishes the uniqueness claim for σ_{B_0} . The argument for ι_{B_0} follows in a similar fashion by considering $\Lambda^*B_0\Lambda - A$ instead of $A - \Lambda^*B_0\Lambda$. $\square$

Remark. The construction from Theorem 5 is implemented in the MATLAB file accompanying this article, which reveals that the “plug-in estimate is not optimal”. Precisely, by plug-in estimate, one means the functional δ^{plug} equal to $\gamma \circ \Delta^{\text{opti}}$, plus an optimally chosen scalar β , where Δ^{opti} is the locally (hence globally) optimal recovery map for $\Gamma = \text{Id}_H$. This map associates to $y \in \mathbb{R}^m$ the Chebyshev center of $\mathcal{K}_y = \{f \in H : \|Tf\| \leq 1, \Lambda f = y\}$ and is known (see (14, Lemma 6) for a bit more) to be $y \in \mathbb{R}^m \mapsto \text{argmin}\{\|Tf\| : \Lambda f = y\}$, i.e., the linear map $\Delta^{\text{opti}} = T^{-1}(T^{-1})^*\Lambda^*(\Lambda T^{-1}(T^{-1})^*\Lambda^*)^{-1}$. Thus, the plug-in functional δ^{plug} is a quadratic functional whose global-worst case error can be computed according to [17]. Numerical tests reveal that in general it is strictly larger than the minimal global worst-case error as computed according to [16]—this experimental conclusion is not quite apparent for a random T , but it is for $T = \text{Id}_H$.

5. Incorporating Observation Errors

In practical situations, the vector $y \in \mathbb{R}^m$ is likely to be contaminated by observation errors, so that $y = \Lambda(f) + e$ for some nonzero $e \in \mathbb{R}^m$. This error vector can be modeled stochastically, as in (7–10), or deterministically, as in this article, by way of an uncertainty set

$$\mathcal{E} = \{e \in \mathbb{R}^m : \|Ue\|_2 \leq 1\}, \quad [20]$$

where U is an isomorphism on $\mathbb{R}^m$. The case of error vectors of small magnitude $\eta > 0$ is covered by this formalism via the choice $U = \eta^{-1}\text{Id}_{\mathbb{R}^m}$. It is well known that such an “inaccurate scenario” can be reduced to the “accurate scenario”. Precisely, the adjusted global-worst

case error takes the form

$$\begin{aligned}\widetilde{\text{gwce}}(\delta) &:= \sup_{\substack{f \in \mathcal{K} \\ e \in \mathcal{E}}} |\gamma(f) - \delta(\Lambda f + e)| \\ &= \sup_{\tilde{f} \in \tilde{\mathcal{K}}} |\tilde{\gamma}(\tilde{f}) - \delta(\tilde{\Lambda}\tilde{f})|,\end{aligned}$$

with “tilde” objects given by $\tilde{\mathcal{K}} := \mathcal{K} \times \mathcal{E} \subseteq \tilde{H} := H \times \mathbb{R}^m$,

$$\begin{aligned}\tilde{f} &:= \begin{bmatrix} f \\ e \end{bmatrix} \in \tilde{H}, \quad \tilde{\Lambda}\tilde{f} := \Lambda f + e = \begin{bmatrix} \Lambda & \text{Id}_{\mathbb{R}^m} \end{bmatrix} \begin{bmatrix} f \\ e \end{bmatrix}, \\ \tilde{\gamma}(\tilde{f}) &:= \gamma(f) = \left\langle \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} f \\ e \end{bmatrix}, \begin{bmatrix} f \\ e \end{bmatrix} \right\rangle =: \langle \tilde{A}\tilde{f}, \tilde{f} \rangle.\end{aligned}$$

This is a formal reduction and does not automatically allow one to call upon earlier results that relied on the specificity of the model set $\mathcal{K}$, especially when computational issues are concerned. For instance, the original model set $\mathcal{K}$ was an ellipsoid but the new model set $\tilde{\mathcal{K}}$ is, in the adjusted space $\tilde{H}$, the intersection of two ellipsoids, so the proof of Theorem 1, say, would not carry over, since the use of the S-lemma was contingent on $\mathcal{K}$ being described by the lone quadratic inequality $\|Tf\| \leq 1$. Fortunately, there is an analog of the S-lemma covering the case of one quadratic inequality being a consequence of two—not one—quadratic inequalities. Discussed in (13), this result derived by Polyak in (15) has already bear fruits in Optimal Recovery: it was invoked in (16) to show how the optimal parameter of a Tikhonov-style regularization (which is a globally optimal map for the estimation of the identity) can actually be computed.

Polyak S-procedure: given three quadratic functionals of the specific form $q_i(h) = \langle R_i h, h \rangle + \rho_i$, $i = 0, 1, 2$, defined on a Hilbert space H of finite dimension at least 3, the equivalence

$$\begin{aligned}q_0(h) \leq 0 \quad \text{whenever } q_1(h) \leq 0 \text{ and } q_2(h) \leq 0 \\ \iff [\exists c_1, c_2 \geq 0 : q_0 \leq c_1 q_1 + c_2 q_2]\end{aligned}$$

holds provided $q_1(h) < 0$ and $q_2(h) < 0$ for some $h \in H$ and $d_1 R_1 + d_2 R_2 \succ 0$ for some $d_1, d_2 \in \mathbb{R}$.

Although of no assistance in the local setting, this S-procedure is quite fitted to the global setting. Indeed, the main results for the “accurate scenario” admit the following extensions.

Theorem 7. In a finite-dimensional Hilbert space H , for the estimation of the quadratic form [9] using a linear observation map $\Lambda : H \rightarrow \mathbb{R}^m$, the model set $\mathcal{K}$ of [15], and the uncertainty set $\mathcal{E}$ of [20], a globally optimal quadratic recovery functional is given as

$$\delta_{B_0}(y) = \langle B_0 y, y \rangle + \beta_0, \quad y \in \mathbb{R}^m,$$

where the symmetric matrix $B_0 \in \mathbb{R}^{m \times m}$ is obtained as solution to the semidefinite program [21] displayed in Fig. 3. The solutions $\tilde{c}'_1, \tilde{c}'_2, \tilde{c}''_1, \tilde{c}''_2 \geq 0$ yield the values of the optimal parameter $\beta_0 = ((\tilde{c}'_1 - \tilde{c}''_1) + (\tilde{c}'_2 - \tilde{c}''_2))/2$, as well as of the minimal global worst-case estimation error, equal to $((\tilde{c}'_1 + \tilde{c}''_1) + (\tilde{c}'_2 + \tilde{c}''_2))/2$.

Theorem 8. Under the setting of Theorem 7, once B_0 has been computed, if the optimal values $\tilde{c}'_1, \tilde{c}'_2, \tilde{c}''_1, \tilde{c}''_2$ are all strictly positive, and if the matrix

$$\left[\begin{array}{c|c} A - \Lambda^* B_0 \Lambda & -\Lambda^* B_0 \\ \hline -B_0 \Lambda & -B_0 \end{array} \right]$$

has a largest eigenvalue which is simple when left- and right-multiplied by $\text{diag}[(\tilde{c}'_1 T^* T)^{-1/2}; (\tilde{c}'_2 U^* U)^{-1/2}]$ and a smallest eigenvalue which is simple when left- and right-multiplied by $\text{diag}[(\tilde{c}''_1 T^* T)^{-1/2}; (\tilde{c}''_2 U^* U)^{-1/2}]$, then

$$\widetilde{\text{gwce}}(\delta_{B_0}) \leq \widetilde{\text{gwce}}(\delta) \quad \text{for all } \delta : \mathbb{R}^m \rightarrow \mathbb{R}.$$

Proof of Theorem 7. The considerations from Section 3 carry over, *mutatis mutandis*, using the “tilde” objects. In particular, with the adjusted model set $\tilde{\mathcal{K}}$ taking the form $\tilde{\mathcal{K}} = \{\tilde{f} \in \tilde{H} : \|\tilde{T}\tilde{f}\| \leq 1, \|\tilde{U}\tilde{f}\| \leq 1\}$ involving operators defined for $\tilde{f} = [f; e]$ by $\tilde{T}\tilde{f} = [Tf; 0]$ and $\tilde{U}\tilde{f} = [0; Ue]$, Lemma 3 ensures that the task at hand consists in minimizing $(\tilde{\sigma}_B - \tilde{\iota}_B)/2$, with obvious notation. Concentrating on the first quantity, it is

$$\begin{aligned}\tilde{\sigma}_B &:= \sup_{\tilde{f} \in \tilde{\mathcal{K}}} \langle (\tilde{A} - \tilde{\Lambda}^* B \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle \\ &= \inf_d d : \langle (\tilde{A} - \tilde{\Lambda}^* B \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle \leq d \\ &\quad \text{whenever } \|\tilde{T}\tilde{f}\|^2 \leq 1 \text{ and } \|\tilde{U}\tilde{f}\|^2 \leq 1.\end{aligned}$$

From the S-procedure, which applies since the strict feasibility constraint is met with $\tilde{f} = 0$ and with any $d_1, d_2 > 0$ (recall that T and U are invertible), the above constraint is equivalent to the existence of $c_1, c_2 \geq 0$ such that

$$\langle (\tilde{A} - \tilde{\Lambda}^* B \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle - d \leq c_1 (\|\tilde{T}\tilde{f}\|^2 - 1) + c_2 (\|\tilde{U}\tilde{f}\|^2 - 1)$$

for all $\tilde{f} \in \tilde{H}$, which decouples as $d \geq c_1 + c_2$ and $\langle (c_1 \tilde{T}^* \tilde{T} + c_2 \tilde{U}^* \tilde{U} - \tilde{A} + \tilde{\Lambda}^* B \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle \geq 0$ for all $\tilde{f} \in \tilde{H}$. Incorporating $c_1, c_2 \geq 0$ as optimization variables, taking d as its minimal admissible value, and spelling out the form of $\tilde{A}$, $\tilde{\Lambda}$, $\tilde{T}$, and $\tilde{U}$ leads to

$$\begin{aligned}\tilde{\sigma}_B &= \inf_{c_1, c_2 \geq 0} c_1 + c_2 \\ \text{s.to } &\left[\begin{array}{c|c} c_1 T^* T - A + \Lambda^* B \Lambda & \Lambda^* B \\ \hline B \Lambda & c_2 U^* U + B \end{array} \right] \succeq 0.\end{aligned}$$

In a very similar way, one derives

$$\begin{aligned}-\tilde{\iota}_B &:= \sup_{\tilde{f} \in \tilde{\mathcal{K}}} \langle (-\tilde{A} + \tilde{\Lambda}^* B \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle \\ &= \inf_{c_1, c_2 \geq 0} c_1 + c_2 \\ \text{s.to } &\left[\begin{array}{c|c} c_1 T^* T + A - \Lambda^* B \Lambda & -\Lambda^* B \\ \hline -B \Lambda & c_2 U^* U - B \end{array} \right] \succeq 0.\end{aligned}$$

Putting these two infima together in $(\tilde{\sigma}_B - \tilde{\iota}_B)/2$ and further minimizing over symmetric matrices $B \in \mathbb{R}^{m \times m}$ yields the semidefinite program [21]. Finally, keeping in mind that $\tilde{\sigma}_B$ corresponds to $\tilde{c}'_1 + \tilde{c}'_2$ and that $-\tilde{\iota}_B$ corresponds to $\tilde{c}''_1 + \tilde{c}''_2$, the values of the optimal parameter β_0 and of the minimal global worst-case error follow. $\square$

$$\begin{aligned} & \text{minimize} \quad \frac{c'_1 + c''_1 + c'_2 + c''_2}{2} \quad \text{subject to} \quad \begin{cases} \left[\begin{array}{c|c} c'_1 T^* T - A + \Lambda^* B \Lambda & \Lambda^* B \\ \hline B \Lambda & c'_2 U^* U + B \end{array} \right] \succeq 0, \\ \left[\begin{array}{c|c} c''_1 T^* T + A - \Lambda^* B \Lambda & -\Lambda^* B \\ \hline -B \Lambda & c''_2 U^* U - B \end{array} \right] \succeq 0. \end{cases} \\ & \begin{matrix} c'_1, c'_2 \geq 0 \\ c''_1, c''_2 \geq 0 \\ B \text{ symmetric} \end{matrix} \end{aligned} \quad [21]$$

Fig. 3. The semidefinite program to construct a globally optimal quadratic estimation map (see Theorem 7) in the presence of observation errors.

Proof of Theorem 8. By Theorem 4, it suffices to establish that the extremizers for $\tilde{\sigma}_{B_0}$ and $\tilde{\iota}_{B_0}$ are unique (up to a sign). For instance, let $\tilde{f}^\sharp$ denote a maximizer of

$$\begin{aligned} \tilde{\sigma}_{B_0} &= \sup_{\tilde{f} \in \tilde{H}} \langle (\tilde{A} - \tilde{\Lambda}^* B_0 \tilde{\Lambda}) \tilde{f}, \tilde{f} \rangle \quad \text{s.t.} \quad \|\tilde{T} \tilde{f}\| \leq 1, \|\tilde{U} \tilde{f}\| \leq 1 \\ &= \inf_{c_1, c_2 \geq 0} c_1 + c_2 \quad \text{s.t.} \quad c_1 \tilde{T}^* \tilde{T} + c_2 \tilde{U}^* \tilde{U} - \tilde{A} + \tilde{\Lambda}^* B_0 \tilde{\Lambda} \succeq 0, \end{aligned}$$

where the latter expression was obtained in the previous proof. Writing $\bar{c}_1, \bar{c}_2$ for its minimizers, the dual feasibility condition implies that

$$\langle (\bar{c}_1 \tilde{T}^* \tilde{T} + \bar{c}_2 \tilde{U}^* \tilde{U} - \tilde{A} + \tilde{\Lambda}^* B_0 \tilde{\Lambda}) \tilde{f}^\sharp, \tilde{f}^\sharp \rangle \geq 0,$$

which, together with the primal feasibility conditions, further implies that

$$\begin{aligned} \tilde{\sigma}_{B_0} &= \langle (\tilde{A} - \tilde{\Lambda}^* B_0 \tilde{\Lambda}) \tilde{f}^\sharp, \tilde{f}^\sharp \rangle \\ &\leq \bar{c}_1 \|\tilde{T} \tilde{f}^\sharp\|^2 + \bar{c}_2 \|\tilde{U} \tilde{f}^\sharp\|^2 \quad [22] \\ &\leq \bar{c}_1 + \bar{c}_2 = \tilde{\sigma}_{B_0}. \quad [23] \end{aligned}$$

Therefore, equality must occur all the way through. In particular, in view of $\bar{c}_1 > 0$ and $\bar{c}_2 > 0$, equality in [23] forces $\|\tilde{T} \tilde{f}^\sharp\| = 1$ and $\|\tilde{U} \tilde{f}^\sharp\| = 1$. Moreover, in view of $\tilde{A} - \tilde{\Lambda}^* B_0 \tilde{\Lambda} \preceq \tilde{M} := \bar{c}_1 \tilde{T}^* \tilde{T} + \bar{c}_2 \tilde{U}^* \tilde{U}$, equality in [22] means that $\tilde{g}^\sharp := \tilde{M}^{1/2} \tilde{f}^\sharp$ is an eigenvector of $\tilde{M}^{-1/2} (\tilde{A} - \tilde{\Lambda}^* B_0 \tilde{\Lambda}) \tilde{M}^{-1/2}$ for its largest eigenvalue. Since this eigenvalue has multiplicity one, the eigenvector $\tilde{g}^\sharp$ —normalized via $\|\tilde{g}^\sharp\|^2 = \bar{c}_1 + \bar{c}_2$ —is uniquely determined (up to a sign), and then so is the maximizer $\tilde{f}^\sharp$ for $\tilde{\sigma}_{B_0}$. The maximizer for $-\tilde{\iota}_B$, i.e., the minimizer for $\tilde{\iota}_B$, is proved to be unique (up to a sign) in a similar fashion. $\square$

Remark. The compute-first-verify-later paradigm is more subtle in the presence of observation error than in its absence. Indeed, without error, the computed minimal worst-case error over functionals δ_B could be followed by a comparison with the computed lower bound in order to certify the optimality of quadratic functionals. As revealed next in Proposition 9, this is realizable, although not numerically reliable. With error, however, computing the lower bound along the lines of Proposition 9 is doomed, as a version of the S-procedure where a quadratic inequality is a consequence of more than two quadratic inequalities would be required. This is known to fail in general, hence underlining the importance of Theorem 8.

6. Postlude on Computational Discovery

I wish to thank Chunyang Liao for asking about the optimal estimation of quadratic functionals at the end of a presentation of the results from (6). It prompted me to investigate this question, numerically at first because a way to compute the minimal global worst-case error over quadratic functionals was relatively easy to come by (see Theorem 5), so an answer could be speculated if there was also a way to compute the well-known lower bound on the global worst-case error of any estimation functional (see Lemma 2). And there was a way—it led me to the empirical conviction that “optimal estimation of quadratic forms can be achieved with quadratic functionals”. This discovery owes a lot to the availability of computational tools such as CVX (17). Later came the formal justification presented in this article, whose unfortunate consequence is to erase a decisive part of the process. Still, I believe that it is well worth mentioning it here. The main tool, also discussed in (13), is again related to the S-lemma.

Finsler lemma: given two self-adjoint operators R_0, R_1 defined on a finite-dimensional Hilbert space H , the following equivalence holds:

$$\begin{aligned} & \langle R_0 f, f \rangle > 0 \text{ whenever } \langle R_1 f, f \rangle = 0 \text{ and } f \neq 0 \\ & \iff [\exists d \in \mathbb{R} : R_0 + d R_1 \succ 0]. \end{aligned}$$

Proposition 9. For a finite-dimensional Hilbert space H , a linear observation map $\Lambda : H \rightarrow \mathbb{R}^m$, and the model set $\mathcal{K}$ of [15], the lower bound [10] on the global worst-case error of any functional for the estimation of the quadratic form [9] can be computed as the optimal value of the semidefinite program [24] displayed in Fig. 4

Proof. With the notation $\hat{f} := [f'; f''] \in \hat{H} := H \times H$, $\hat{A}(\hat{f}) := [A f'; -A f'']$, $\hat{T}'(\hat{f}) := T(f')$, $\hat{T}''(\hat{f}) := T(f'')$, $\hat{\Lambda}(\hat{f}) := \Lambda(f') - \Lambda(f'')$, the lower bound takes the form

$$\begin{aligned} \text{lb} &= \sup_{f', f''} \left\{ \frac{\langle A f', f' \rangle - \langle A f'', f'' \rangle}{2} : \right. \\ & \quad \left. \|T f'\|^2 \leq 1, \|T f''\|^2 \leq 1, \Lambda(f') = \Lambda(f'') \right\} \\ &= \sup_{\hat{f}} \left\{ \frac{\langle \hat{A} \hat{f}, \hat{f} \rangle}{2} : \|\hat{T}' \hat{f}\|^2 \leq 1, \|\hat{T}'' \hat{f}\|^2 \leq 1, \hat{f} \in \ker(\hat{\Lambda}) \right\} \\ &= \inf_d d : \text{for all } \hat{f} \in \ker(\hat{\Lambda}), \\ & \quad \frac{\langle \hat{A} \hat{f}, \hat{f} \rangle}{2} \leq d \text{ whenever } \|\hat{T}' \hat{f}\|^2 \leq 1, \|\hat{T}'' \hat{f}\|^2 \leq 1. \end{aligned}$$

$$\underset{c_1, c_2 \geq 0}{\text{minimize}} \frac{c_1 + c_2}{2} \quad \text{subject to} \quad \left[\begin{array}{c|c} c_1 T^* T - A + d \Lambda^* \Lambda & -d \Lambda^* \Lambda \\ \hline -d \Lambda^* \Lambda & c_2 T^* T + A + d \Lambda^* \Lambda \end{array} \right] \succeq 0. \quad [24]$$

Fig. 4. The semidefinite program to determine the lower bound on the global worst-case error (see Proposition 9) in the absence of observation errors.

First, relying on Polyak S-procedure in the space $\ker(\hat{\Lambda})$, one obtains (the strict feasibility condition being easily verified)

$$\begin{aligned} \text{lb} &= \inf_d d : \exists c_1, c_2 \geq 0 \text{ such that, for all } \hat{f} \in \ker(\hat{\Lambda}), \\ &\langle \hat{A} \hat{f}, \hat{f} \rangle - 2d \leq c_1 (\|\hat{T}' \hat{f}\|^2 - 1) + c_2 (\|\hat{T}'' \hat{f}\|^2 - 1) \\ &= \inf_{c_1, c_2 \geq 0} \frac{c_1 + c_2}{2} \\ &\text{s.to } \langle (c_1 \hat{T}'^* \hat{T}' + c_2 \hat{T}''^* \hat{T}'' - \hat{A}) \hat{f}, \hat{f} \rangle \geq 0 \\ &\text{whenever } \langle \hat{\Lambda}^* \hat{\Lambda} \hat{f}, \hat{f} \rangle = 0. \end{aligned}$$

Note that this coincides with

$$\begin{aligned} \text{lb}' &:= \inf_{c_1, c_2 \geq 0} \frac{c_1 + c_2}{2} \\ &\text{s.to } \langle (c_1 \hat{T}'^* \hat{T}' + c_2 \hat{T}''^* \hat{T}'' - \hat{A}) \hat{f}, \hat{f} \rangle > 0 \\ &\text{whenever } \langle \hat{\Lambda}^* \hat{\Lambda} \hat{f}, \hat{f} \rangle = 0 \text{ and } \hat{f} \neq 0. \end{aligned}$$

Indeed, on the one hand, if $c_1, c_2 \geq 0$ satisfy the constraint of lb' , then they satisfy the constraint of lb , yielding $(c_1 + c_2)/2 \geq \text{lb}$, and in turn $\text{lb}' \geq \text{lb}$; on the other hand,

given $\varepsilon > 0$, if $c_1, c_2 \geq 0$ satisfy the constraint of lb , then $c_1 + \varepsilon, c_2 + \varepsilon \geq 0$ satisfy the constraint of lb' (recall that T is invertible), yielding $(c_1 + \varepsilon + c_2 + \varepsilon)/2 \geq \text{lb}'$, and in turn $\text{lb} \geq \text{lb}' - \varepsilon$. Therefore, Finsler lemma can be utilized to arrive at

$$\begin{aligned} \text{lb} &= \inf_{c_1, c_2 \geq 0} \frac{c_1 + c_2}{2} \\ &\text{s.to } \exists d \in \mathbb{R} : c_1 \hat{T}'^* \hat{T}' + c_2 \hat{T}''^* \hat{T}'' - \hat{A} + d \hat{\Lambda}^* \hat{\Lambda} \succ 0 \\ &= \inf_{c_1, c_2 \geq 0} \frac{c_1 + c_2}{2} \\ &\text{s.to } c_1 \hat{T}'^* \hat{T}' + c_2 \hat{T}''^* \hat{T}'' - \hat{A} + d \hat{\Lambda}^* \hat{\Lambda} \succ 0. \end{aligned}$$

As before, this infimum can be seen to equal the one with positive semidefiniteness replacing positive definiteness—although the straightforward inequality would be enough, as one seeks a lower bound. It now remains to spell out the form of $\hat{A}$, $\hat{T}'$, $\hat{T}''$, and $\hat{\Lambda}$ in order to arrive at the semidefinite program [24]. $\square$

ACKNOWLEDGMENTS. The author is partially supported by a grant from the National Science Foundation (DMS-2505204).

1. CA Micchelli, TJ Rivlin, Lectures on optimal recovery in *Numerical Analysis Lancaster 1984*, pp. 21–93 (1985).
2. JF Traub, GW Wasilkowski, H Woźniakowski, *Information-based complexity*. (Academic Press), (1988).
3. SA Smolyak, “On optimal restoration of functions and functionals of them,” PhD thesis, Moscow State University (1965).
4. NS Bakhvalov, On the optimality of linear methods for operator approximation in convex classes of functions. *USSR Comput. Math. Math. Phys.* **11**, 244–249 (1971).
5. AG Sukharev, On the existence of optimal affine methods for approximating linear functionals. *J. Complex.* **2**, 317–322 (1986).
6. S Foucart, Optimal algorithms for nonlinear estimation with convex models. *arXiv preprint arXiv:2512.20826* (2025).
7. DL Donoho, M Nussbaum, Minimax quadratic estimation of a quadratic functional. *J. Complex.* **6**, 290–323 (1990).
8. J Fan, On the estimation of quadratic functionals. *The Annals Stat.* **19**, 1273–1294 (1991).
9. TT Cai, MG Low, Nonquadratic estimators of a quadratic functional. *The Annals Stat.* **33**, 2930–2956 (2005).
10. A Juditsky, A Nemirovski, Estimating linear and quadratic forms via indirect observations. *Bernoulli* **26**, 2639–2669 (2020).
11. DL Donoho, Statistical estimation and optimal recovery. *The Annals Stat.* **22**, 238–270 (1994).
12. S Foucart, G Paouris, Near-optimal estimation of linear functionals with log-concave observation errors. *Inf. Inference* **12**, 2546–2561 (2023).
13. I Pólik, T Terlaky, A survey of the S-lemma. *SIAM Rev.* **49**, 371–418 (2007).
14. S Foucart, N Hengartner, Worst-case learning under a multifidelity model. *SIAM/ASA J. on Uncertain. Quantification* **13**, 171–194 (2025).
15. BT Polyak, Convexity of quadratic transformations and its use in control and optimization. *J. Optim. Theory Appl.* **99**, 553–583 (1998).
16. S Foucart, C Liao, Optimal recovery from inaccurate data in Hilbert spaces: regularize, but what of the parameter? *Constr. Approx.* **57**, 489–520 (2023).
17. M Grant, S Boyd, CVX: Matlab software for disciplined convex programming, version 2.1 (<https://cvxr.com/cvx>) (2014).