

LEAST MULTIVARIATE CHEBYSHEV POLYNOMIALS ON DIAGONALLY-DETERMINED SETS

M. DRESSLER, S. FOUCART, M. JOLDES, E. DE KLERK, J. B. LASSERRE, AND Y. XU

ABSTRACT. We consider a new multivariate generalization of the classical monic (univariate) Chebyshev polynomial that minimizes the uniform norm on the interval $[-1, 1]$. Let Π_n^* be the subset of polynomials of degree at most n in d variables, whose homogeneous part of degree n has coefficients summing up to 1. The problem is determining a polynomial in Π_n^* with the smallest uniform norm on a set $\Omega \subset \mathbb{R}^d$, which we call a least Chebyshev polynomial (associated with Ω). Our main result solves the problem for Ω belonging to a non-trivial class of sets that we call *diagonally-determined*, and establishes the remarkable result that a least Chebyshev polynomial can be given via the classical, univariate, Chebyshev polynomial. In particular, the solution can be independent of the dimension. Diagonally-determined sets include centered balls in \mathbb{R}^d in any norm, but can be non-convex and even non-simply connected. We also introduce a computational procedure, based on semidefinite programming hierarchies, to detect if a given semi-algebraic set is diagonally-determined.

1. INTRODUCTION

Among its numerous properties, the Chebyshev polynomial $T_n(x) = \cos(n \arccos x)$ provides a solution for the best approximation to the monomial x^n on the interval $[-1, 1]$ in the uniform norm. More precisely, the polynomial

$$q_n^*(x) = x^n - 2^{1-n}T_n(x)$$

of degree $n - 1$ ($n \in \mathbb{N}$) is the best polynomial of approximation to x^n on $[-1, 1]$; that is

$$(1.1) \quad q_n^* = \arg \min_{q \in \Pi_{n-1}} \sup_{x \in [-1, 1]} |x^n - q(x)|,$$

with Π_{n-1} being the vector space of univariate polynomials of degree at most $n - 1$. In other words, the monic Chebyshev polynomial $x^n - q_n^*$ is the *least polynomial* in the sense that it has the least uniform norm among all monic polynomials of degree n .

There have been multiple extensions of Chebyshev polynomials to multivariate settings from different angles. From the point of view of approximation, an immediate generalization is finding the best approximation to monomials [1, 2, 5, 3, 9, 7, 11, 15, 16, 17]. To this end we replace the interval $[-1, 1]$ by a general set $\Omega \subset \mathbb{R}^d$, for $d > 1$. Letting $\alpha \in \mathbb{N}^d$ and $|\alpha| := \sum_{i=1}^d \alpha_i = n$, we consider the problem

$$(1.2) \quad \min_{q \in \Pi_{n-1}^d} \sup_{\mathbf{x} \in \Omega} |\mathbf{x}^\alpha - q(\mathbf{x})|,$$

2020 *Mathematics Subject Classification.* Primary: 41A10, 41A63, 90C22.

Key words and phrases. Chebyshev polynomial, Chebyshev approximation, extremal signature, semidefinite programming, Lasserre hierarchy.

where Π_{n-1}^d denotes the space of real polynomials of total degree at most $n-1$ in the variables $\mathbf{x} = (x_1, \dots, x_d)$, and where we define the monomial $\mathbf{x}^\alpha := x_1^{\alpha_1} \dots x_d^{\alpha_d}$. Since we will deal with monomials of degree n throughout the paper, we introduce the notation

$$\mathbb{N}_n^d := \{\alpha \in (\mathbb{N}_0)^d : |\alpha| = n\},$$

where $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ denotes the set of nonnegative integers. Thus \mathbb{N}_n^d indexes all monomials of degree *exactly* n in d variables.

Problem (1.2) can be regarded as a natural multivariate generalization of (1.1), where $\Omega \subset \mathbb{R}^d$ is a subset of \mathbb{R}^d . While the interval $[-1, 1]$ is a prototype of a compact connected set of the real line, there is no ‘prototype’ in higher dimensions for such a set. In the literature, this problem has been studied primarily on a few special regular domains. In two variables, the problem (1.2) is solved for the square, the disk, and the isosceles right triangle [5, 3, 7, 11, 15]. While the solution on the square can be extended to the cube for $d > 2$, the problem is solved only for a few cases, mostly monomials of lower degrees, on the ball and the simplex [1, 2, 16]. Moreover, the existing examples indicate an increasing complexity, so much so that it does not appear possible to find an analytic solution even for these regular domains.

Recently, in [4] we have proposed to investigate (1.2) for various choices of $\Omega \subset \mathbb{R}^d$ and $\alpha \in \mathbb{N}_n^d$ by combining analytical tools with numerical tools from optimization (and notably the so-called moment-SOS hierarchy [6]). During this study, we have encountered an optimization problem that has initiated a change of view: Namely, instead of studying the best polynomial of approximation to monomials, we can study the least polynomial instead. While the two concepts are identical in one variable, they can be quite different in higher dimensions, as seen from the definition below.

We emphasize that – even though our paper deals with the characterization of least polynomials on special sets – it is intrinsically connected to constructive approximation and numerical optimization due to the close links with the work in [4] and [6] respectively.

Definition 1.1. Let Π_n^d denote the space of polynomials of total degree at most n in d variables, and Π_n^* the subset of Π_n^d that consists of polynomials of the form

$$\mathbf{x} \mapsto P(\mathbf{x}) := \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha \mathbf{x}^\alpha + Q(\mathbf{x}) \quad \text{with} \quad \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha = 1 \quad \text{and} \quad Q \in \Pi_{n-1}^d.$$

Let Ω be a set in \mathbb{R}^d . We consider the optimization problem

$$(1.3) \quad \inf_{P \in \Pi_n^*} \|P\|_\Omega, \quad \text{where} \quad \|P\|_\Omega := \sup_{\mathbf{x} \in \Omega} |P(\mathbf{x})|.$$

If it exists, we call a minimizer $P^* \in \Pi_n^*$ of (1.3) a least Chebyshev polynomial of degree n on the set Ω .

For $d = 1$, there is only one monomial of degree n . In the case $d > 1$, every element of Π_n^* is ‘monic’ and the monomial \mathbf{x}^α in (1.2) is only one among many possible choices in Π_n^* . However, rather than approximating a fixed monomial by polynomials of lower degree, the problem (1.3) requires finding a polynomial that has the least norm among all polynomials in Π_n^* . As far as we are aware, this problem has not been considered in the literature.

Contribution. The main purpose of this paper is to report our findings on the optimization problem (1.3). It turns out, much to our surprise, that the problem (1.3) can be solved analytically for a fairly general family of sets Ω in \mathbb{R}^d for all $d \geq 2$. This family of sets will be referred to as *diagonally-determined*.

Organization of this paper. The paper is organized as follows. We start by defining and describing several examples of diagonally-determined sets in Section 2. Our main results are presented in Section 3, where we describe the least polynomial for a diagonally-determined set; see Theorem 3.1. We discuss the dual problem of (1.3) in Section 4 and show that it too has a closed-form solution in the case of diagonally-determined sets; see Theorem 4.8. We also rephrase this result for the dual problem in the framework of *extremal signatures*. Finally, we consider the problem of deciding whether a given semi-algebraic set is diagonally determined in Section 5, and propose an algorithmic approach based on semidefinite programming hierarchies. We conclude the paper with a discussion of related problems and open questions in Section 6.

2. DIAGONALLY-DETERMINED SETS

We define the *diagonal* of a set $\Omega \subset \mathbb{R}^d$ as the set $\text{diag}(\Omega) := \{t \in \mathbb{R} : t\mathbf{1} \in \Omega\}$, where $\mathbf{1}$ denotes the all-ones vector in \mathbb{R}^d . For two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, we use $\langle \mathbf{x}, \mathbf{y} \rangle$ for their standard inner product.

Definition 2.1. We call a set $\Omega \subset \mathbb{R}^d$ diagonally-determined if the following two conditions hold:

- (1) The diagonal of Ω is an interval, say $\text{diag}(\Omega) = [a, b]$, and
- (2) there exists $\mathbf{v} \in \mathbb{R}^d$ such that $\langle \mathbf{v}, \mathbf{1} \rangle = 1$ and $\langle \mathbf{v}, \mathbf{x} \rangle \in [a, b]$ for all $\mathbf{x} \in \Omega$.

Importantly, Definition 2.1 covers sets that could be non-convex, non-compact, and not simply connected. Figure 1 gives an example of a non-convex diagonally-determined set in \mathbb{R}^2 .




FIGURE 1. Example of a diagonally-determined, non-convex set. Here, $\text{diag}(\Omega) = [1, 3]$, and $\mathbf{v} = (1, 0)$.

As a first observation, a dilation of a diagonally-determined set is again diagonally-determined.

Lemma 2.2. *If $\Omega \subset \mathbb{R}^d$ is diagonally-determined with $\text{diag}(\Omega) = [a, b]$ and vector \mathbf{v} , then, for any $r > 0$, the dilation*

$$r\Omega := \{r\mathbf{x} : \mathbf{x} \in \Omega\}$$

is also diagonally-determined with $\text{diag}(r\Omega) = [ra, rb]$ and vector \mathbf{v} .

Proof. The proof is an immediate consequence of Definition 2.1. \square

We will now show that balls in \mathbb{R}^d (in any norm), centered at the origin, are examples of diagonally-determined sets. To this end, we first need to recall some definitions.

Definition 2.3. *For a given norm $\|\cdot\|$ on \mathbb{R}^d , we use $\|\cdot\|_*$ to denote its dual norm, namely*

$$\|\mathbf{u}\|_* = \sup\{\langle \mathbf{u}, \mathbf{x} \rangle : \|\mathbf{x}\| = 1\}, \quad \mathbf{u} \in \mathbb{R}^d.$$

We will also consider the special case of monotone norms.

Definition 2.4. *We call a norm $\|\cdot\|$ on \mathbb{R}^d monotone, if*

$$\|\mathbf{x}\| \leq \|\mathbf{x}\| \quad \forall \mathbf{x} \in \mathbb{R}^d,$$

where $|\mathbf{x}| = (|x_1|, \dots, |x_d|)$.

We may now show that any centered ball is diagonally-determined.

Proposition 2.5. *Consider a ball of radius r in \mathbb{R}^d centered at the origin:*

$$\Omega := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq r\},$$

where $\|\cdot\|$ denotes any norm on \mathbb{R}^d . Then Ω is a diagonally-determined set. If, in addition, the norm $\|\cdot\|$ is monotone, then the vector \mathbf{v} may be assumed to be entrywise nonnegative, denoted by $\mathbf{v} \geq \mathbf{0}$.

Proof. We prove the statement for unit balls; the required result then follows from Lemma 2.2. For a centered unit ball, one has $\text{diag}(\Omega) = [a, b]$ with $a = -1/\|\mathbf{1}\|$ and $b = 1/\|\mathbf{1}\|$.

Now let $\mathbf{v} = \mathbf{u}/\|\mathbf{1}\|$, where $\mathbf{u} \in \mathbb{R}^d$ is ‘dual’ to $\mathbf{1}$, in the sense that $\|\mathbf{u}\|_* = 1$ and $\langle \mathbf{u}, \mathbf{1} \rangle = \|\mathbf{1}\|$. We then have $\langle \mathbf{v}, \mathbf{1} \rangle = \langle \mathbf{u}/\|\mathbf{1}\|, \mathbf{1} \rangle = 1$, as required. Moreover we have, for any $\mathbf{x} \in \Omega$,

$$|\langle \mathbf{v}, \mathbf{x} \rangle| = |\langle \mathbf{u}/\|\mathbf{1}\|, \mathbf{x} \rangle| \leq (1/\|\mathbf{1}\|)\|\mathbf{u}\|_*\|\mathbf{x}\| \leq 1/\|\mathbf{1}\|.$$

Thus we see that $\langle \mathbf{v}, \mathbf{x} \rangle \in [a, b]$ for all $\mathbf{x} \in \Omega$, proving that Ω is diagonally-determined.

Assume now that the norm $\|\cdot\|$ is monotone. We have

$$\langle \mathbf{u}, \mathbf{1} \rangle = \|\mathbf{1}\| \geq \|\text{sgn}(\mathbf{u})\| \geq \langle \mathbf{u}, \text{sgn}(\mathbf{u}) \rangle,$$

where, for $i \in \{1, \dots, d\}$,

$$\text{sgn}(\mathbf{u})_i = \begin{cases} 1 & \text{if } u_i \geq 0 \\ -1 & \text{else.} \end{cases}$$

This means that $\sum_{i=1}^d u_i \geq \sum_{i=1}^d |u_i|$, which holds if and only if $\mathbf{u} \geq \mathbf{0}$. Subsequently, since $\mathbf{v} = \mathbf{u}/\|\mathbf{1}\|$, we have $\mathbf{v} \geq \mathbf{0}$. \square

By Definition 2.1, we have the following immediate corollary.

Corollary 2.6. *Assume $\Omega \subset \mathbb{R}^d$ is a subset of a ball in \mathbb{R}^d in any norm, centered at the origin, and that the diagonal of Ω coincides with the diagonal of the ball. Then Ω is a diagonally-determined set.*

To illustrate this corollary, the example in Figure 2 shows a non-convex subset of a unit (Euclidean) ball, that has the same diagonal as the ball.




FIGURE 2. Example of a diagonally-determined, non-convex set.

Here, $\text{diag}(\Omega) = \left[-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right]$, and $\mathbf{v} = \left(\frac{1}{2}, \frac{1}{2}\right)$.

Example 2.7. *It is insightful to consider the specific example of the unit ball of the owl norm (ordered weighted ℓ_1 -norm) defined relative to weights $w_1 \geq w_2 \geq \dots \geq w_d \geq 0$ by*

$$\|\mathbf{x}\|_{\text{owl}} = \sum_{i=1}^d w_i x_i^*,$$

with (x_1^*, \dots, x_d^*) being the nonincreasing rearrangement of $(|x_1|, \dots, |x_d|)$. Note that $\|\mathbf{1}\|_{\text{owl}} = \sum_{i=1}^d w_i =: W$, so that $b = 1/W$. Letting $\mathbf{w} = (w_1, \dots, w_d)$, we set $\mathbf{v} = \mathbf{w}/W$, since $\langle \mathbf{w}/W, \mathbf{1} \rangle = (\sum_{i=1}^d w_i)/W = 1$ and, for $\|\mathbf{x}\|_{\text{owl}} \leq 1$,

$$\begin{aligned} |\langle \mathbf{w}/W, \mathbf{x} \rangle| &= \left| \sum_{i=1}^n w_i x_i \right| / W \\ &\leq \left(\sum_{i=1}^n w_i |x_i| \right) / W \\ &\leq \left(\sum_{i=1}^n w_i x_i^* \right) / W \\ &\leq 1/W = b. \end{aligned}$$

The intersection of certain balls in \mathbb{R}^d with the nonnegative orthant in \mathbb{R}^d is also diagonally-determined, as the next result shows.

Proposition 2.8. *Assume $\|\cdot\|$ is a norm on \mathbb{R}^d with dual norm $\|\cdot\|_*$, and that the following holds: there is an entrywise-nonnegative vector $\mathbf{u} \geq \mathbf{0}$ such that $\|\mathbf{u}\|_* = 1$ and $\langle \mathbf{u}, \mathbf{1} \rangle = \|\mathbf{1}\|$.*

Further assume $\Omega \subset \mathbb{R}^d$ has diagonal $\text{diag}(\Omega) = [0, r/\|\mathbf{1}\|]$ for some $r > 0$, and, for all $\mathbf{x} \in \Omega$, $\mathbf{x} \geq \mathbf{0}$ and $\|\mathbf{x}\| \leq r$. Then Ω is a diagonally-determined set with vector $\mathbf{v} = \mathbf{u}/\|\mathbf{1}\|$.

Proof. We again prove the statement for the case $r = 1$; the result for general $r > 0$ then follows from Lemma 2.2. Setting $\mathbf{v} = \mathbf{u}/\|\mathbf{1}\|$, one has $\langle \mathbf{v}, \mathbf{1} \rangle = 1$. Moreover, for $\mathbf{x} \in \Omega$, $\langle \mathbf{v}, \mathbf{x} \rangle \geq 0$, since $\mathbf{x}, \mathbf{v} \geq 0$, and

$$\langle \mathbf{v}, \mathbf{x} \rangle = \frac{1}{\|\mathbf{1}\|} \langle \mathbf{u}, \mathbf{x} \rangle \leq \frac{1}{\|\mathbf{1}\|} \|\mathbf{u}\|_* \|\mathbf{x}\| \leq \frac{1}{\|\mathbf{1}\|},$$

so that $\langle \mathbf{v}, \mathbf{x} \rangle \in \text{diag}(\Omega)$ for all $\mathbf{x} \in \Omega$. \square

We note again that the assumption on the norm in Proposition 2.8 is met by, for example, all *monotone norms*, as shown in the proof of Proposition 2.5.

As a simple corollary of the proposition, the simplex

$$\Omega = \left\{ \mathbf{x} \in \mathbb{R}^d : \mathbf{x} \geq \mathbf{0}, \sum_{i=1}^d x_i \leq 1 \right\},$$

is diagonally-determined, since it is the intersection of the unit ℓ^1 -ball with the non-negative orthant. See Figure 3 for an illustrative example of Proposition 2.8.




FIGURE 3. Example to illustrate Proposition 2.8. Here, $\text{diag}(\Omega) = [0, b]$, with $b \approx 1.35$. Moreover, Ω is contained in the ℓ^1 -ball of radius $r = bd \approx 2.7$, intersected with the nonnegative quadrant.

Finally, the property of being diagonally-determined is not invariant under translation in general. The exception is a translation by a constant multiple of the all-ones vector. We state this observation as a lemma for later use. The proof again follows immediately from Definition 2.1 and is omitted.

Lemma 2.9. *Fix $\alpha \in \mathbb{R}$. Then a given $\Omega \subset \mathbb{R}^d$ is diagonally determined with $\text{diag}(\Omega) = [a, b]$ and vector \mathbf{v} if and only if its translation $\Omega + \alpha \mathbf{1}$ is diagonally determined with diagonal $[a + \alpha, b + \alpha]$ and vector \mathbf{v} .*

3. LEAST CHEBYSHEV POLYNOMIAL FOR A DIAGONALLY-DETERMINED SET

In this section, and with Ω being a set in \mathbb{R}^d , we consider the problem of finding a least polynomial $P^* \in \Pi_n^*$ such that

$$\|P^*\|_\Omega = \inf \{ \|P\|_\Omega : P \in \Pi_n^* \},$$

for the special case where Ω is diagonally-determined.

Theorem 3.1. *Let Ω be diagonally-determined with vector \mathbf{v} and $\text{diag}(\Omega) = [a, b]$. Then*

$$(3.1) \quad \inf \{ \|P\|_{\Omega} : P \in \Pi_n^* \} = \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}},$$

and the infimum is attained by the polynomial

$$(3.2) \quad \mathbf{x} \mapsto P^*(\mathbf{x}) := \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}} T_n \left(-1 + 2 \frac{\langle \mathbf{v}, \mathbf{x} \rangle - a}{b-a} \right).$$

Proof. Let $P \in \Pi_n^*$, say

$$P(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}_n^d} a_{\alpha} \mathbf{x}^{\alpha} + Q(\mathbf{x}) \quad \text{with} \quad \sum_{\alpha \in \mathbb{N}_n^d} a_{\alpha} = 1 \text{ and } Q \in \Pi_{n-1}^d.$$

Then $P(t, \dots, t) = t^n + q_{n-1}(t)$, where $q_{n-1}(t) = Q(t, \dots, t)$ is a univariate polynomial of degree at most $n-1$ in the single variable t . Hence,

$$(3.3) \quad \|P\|_{\Omega} \geq \max_{a \leq t \leq b} |t^n + q_{n-1}(t)| \geq \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}},$$

where we have used a classical result in one variable (e.g. [12, Theorem 2.1]), and equality in the last inequality is attained by choosing q_{n-1} such that $t^n + q_{n-1}(t)$ is the rescaled Chebyshev polynomial in the right-hand side of (3.2), when restricted to $\mathbf{x} = (t, \dots, t)$ ($t \in \mathbb{R}$).

Moreover, as Ω is diagonally-determined, then for all $\mathbf{x} \in \Omega$, one has $\langle \mathbf{v}, \mathbf{x} \rangle \in [a, b]$, and therefore

$$\|P^*\|_{\Omega} = \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}} \max_{\mathbf{x} \in \Omega} \left| T_n \left(-1 + 2 \frac{\langle \mathbf{v}, \mathbf{x} \rangle - a}{b-a} \right) \right| \leq \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}}.$$

Finally, we need to show that $P^* \in \Pi_n^*$. First note that $t \mapsto P^*(t\mathbf{1})$ is a monic univariate polynomial in t , since $\langle \mathbf{v}, \mathbf{1} \rangle = 1$. On the other hand, we may write P^* in the form

$$P^*(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}_n^d} a_{\alpha}^* \mathbf{x}^{\alpha} + Q^*(\mathbf{x}) \quad \text{with } Q^* \in \Pi_{n-1}^d,$$

so that

$$P^*(t\mathbf{1}) = \sum_{\alpha \in \mathbb{N}_n^d} a_{\alpha}^* t^n + Q^*(t\mathbf{1}).$$

Thus $\sum_{\alpha \in \mathbb{N}_n^d} a_{\alpha}^* = 1$, and therefore $P^* \in \Pi_n^*$, as required. \square

The minimal value in (3.1) does not depend explicitly on the dimension d , but the diagonal $[a, b]$ may. For example, if Ω is the unit Euclidean ball in \mathbb{R}^d , then $[a, b] = [-1/\sqrt{d}, 1/\sqrt{d}]$.

Example 3.2. *For the cube $\Omega = [-1, 1]^d$, Theorem 3.1 shows that $\|P^*\|_{\Omega} = 2^{-n+1}$. In contrast, for every $\alpha \in \mathbb{N}^d$ with $|\alpha| = n$, we have*

$$\inf_{P \in \Pi_{n-1}^d} \|\mathbf{x}^{\alpha} - P\|_{[-1, 1]^d} = 2^{-n+d},$$

as shown in [18] (see also [15]), which depends on the dimension d .

Since an optimal solution of problem (1.3) is given in terms of the univariate Chebyshev polynomial T_n in Theorem 3.1, the reader may wonder if this optimal solution is in fact unique. This turns out to be not the case, as the next example shows.

Example 3.3. The simplex $\Delta^d = \{\mathbf{x} \in \mathbb{R}^d : 0 \leq x_1 \leq x_2 \leq \dots \leq x_d \leq 1\}$ is diagonally-determined with $\mathbf{v} = (1, 0, \dots, 0)$, and $\text{diag}(\Delta^d) = [0, 1]$. Hence, by Theorem 3.1,

$$\|P^*\|_{\Delta^d} = \inf \{\|P\|_{\Delta^d} : P \in \Pi_n^*\} = \left(\frac{1}{2}\right)^n \frac{1}{2^{n-1}} = 2^{-2n+1}$$

is independent of the dimension d . Moreover, one has $P^*(\mathbf{x}) = 2^{-2n+1}T_n(2x_1 - 1)$, which is univariate. Another valid choice for \mathbf{v} is $\mathbf{v} = (0, 1, 0, \dots, 0)$ that leads to a different least Chebyshev polynomial, namely $\mathbf{x} \mapsto 2^{-2n+1}T_n(2x_2 - 1)$.

Example 3.4. Returning to Example 2.7, where Ω is the unit ball for the owl norm, we have $[a, b] = [-1/W, 1/W]$ and $\mathbf{v} = \mathbf{w}/W$, so that, by Theorem 3.1, the least Chebyshev polynomial of degree n becomes

$$\mathbf{x} \mapsto P^*(\mathbf{x}) = \left(\frac{1}{W}\right)^n \frac{1}{2^{n-1}} T_n(\langle \mathbf{w}, \mathbf{x} \rangle).$$

We point out that the requirement $\sum_{\alpha \in \mathbb{N}_n^d} a_\alpha = 1$, is susceptible to an affine change of variables other than a translation. For instance, the triangle $\hat{\Delta}^2 = \{\mathbf{x} : x_1 \geq 0, x_2 \geq 0, x_1 + x_2 \leq 1\}$, is a mirror image of Δ^2 under $x_2 \mapsto 1 - x_2$. However, under this affine change of variables, the leading monomial $a_0x_1^2 + a_1x_1x_2 + a_2x_2^2$ of a polynomial in Π_2^* becomes, $a_0x_1^2 - a_1x_1x_2 + a_2x_2^2$, which is, in general, no longer an element of Π_2^* .

Moreover, the extremal problem (1.3) in Definition 1.1 is not rotationally invariant. To see this consider the following three sets in \mathbb{R}^2 , that are rotations of the same line segment:

$$\begin{aligned} \Omega_1 &:= \{(x, x) \mid x \in [-1, 1]\} \\ \Omega_2 &:= \{(0, x) \mid x \in [-\sqrt{2}, \sqrt{2}]\} \\ \Omega_3 &:= \{(x, -x) \mid x \in [-1, 1]\}. \end{aligned}$$

Of these three sets, only Ω_1 is diagonally-determined. Even though the three sets are rotations of each other, the solution to the extremal problem (1.3) is starkly different in each case:

- For Ω_1 , Theorem 3.1 provides the solution $(x_1, x_2) \mapsto \frac{1}{2^{n-1}}T_n(x_1)$ as a least Chebyshev polynomial, and thus $\inf_{P \in \Pi_n^*} \|P\|_{\Omega_1} = \frac{1}{2^{n-1}}$;
- For Ω_2 , a least Chebyshev polynomial is clearly given by $(x_1, x_2) \mapsto x_1^n$, so that $\inf_{P \in \Pi_n^*} \|P\|_{\Omega_2} = 0$;
- For Ω_3 , a least Chebyshev polynomial is $(x_1, x_2) \mapsto \frac{1}{2^n}(x_1 + x_2)^n$, so that $\inf_{P \in \Pi_n^*} \|P\|_{\Omega_3} = 0$.

The example shows that the diagonal of Ω plays a key role in determining the least Chebyshev polynomial, and that it is not only an artefact of the analysis.

4. THE DUAL FRAMEWORK AND SIGNATURES

In this section we consider the dual problem of (1.3). Our goal is to show that, in the special case of diagonally-determined sets, the dual problem has a closed-form solution. Moreover, we will construct a dual solution that is atomic (discrete), and supported on $n + 1$ points.

4.1. The dual problem. With Ω being a set in \mathbb{R}^d , let $C(\Omega)$ be the space of continuous functions on Ω , and let $C(\Omega)^*$ be the dual space of $C(\Omega)$. For compact Ω , one has the following, strong duality result for problem (1.3).

Theorem 4.1. *Let $\Omega \subset \mathbb{R}^d$ be compact, and consider the dual problem of (1.3), namely (4.1)*

$$\gamma^* := \sup_{\substack{\gamma \in \mathbb{R} \\ L \in C(\Omega)^*}} \left\{ \gamma : L|_{\Pi_{n-1}^d} = 0, \|L\|_{C(\Omega)^*} = 1, \text{ and } L(\mathbf{x} \mapsto \mathbf{x}^\alpha) = \gamma \text{ for all } \alpha \in \mathbb{N}_n^d \right\},$$

where the norm of L is defined as

$$\|L\|_{C(\Omega)^*} = \sup_{h \in C(\Omega)} \{|Lh| : \|h\|_\Omega \leq 1\}.$$

One has $\gamma^* \leq \inf \{\|P\|_\Omega : P \in \Pi_n^*\}$ (weak duality), and $\gamma^* = \inf \{\|P\|_\Omega : P \in \Pi_n^*\}$ if Ω is compact (strong duality).

One may derive this result directly through conic linear programming duality theory, e.g. [14, Proposition 2.9], and we will omit the proof here.

It is insightful, though, to make a link with classical Chebyshev approximation. To this end, we denote an optimal solution of problem (1.3) by

$$(4.2) \quad \mathbf{x} \mapsto P^*(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha + Q^*(\mathbf{x}) \quad \text{such that} \quad \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* = 1, \quad Q^* \in \Pi_{n-1}^d.$$

Note that $P^* \in \Pi_n^*$. Clearly,

$$\|P^*\|_\Omega = \min_{Q \in \Pi_{n-1}^d} \left\| \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha + Q(\mathbf{x}) \right\|_\Omega = \min_{Q \in \Pi_{n-1}^d} \left\| \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha - Q(\mathbf{x}) \right\|_\Omega,$$

while the latter problem is the classical Chebyshev problem of approximating the homogeneous polynomial $\mathbf{x} \mapsto \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha$ from Π_{n-1}^d .

Recall the strong duality result for the classical Chebyshev approximation problem as given in the paper by Rivlin and Shapiro [13, Corollary 2].¹

Theorem 4.2. *For any $f \in C(\Omega)$, with Ω compact, one has the following*

$$(4.3) \quad \min_{Q \in \Pi_{n-1}^d} \|f - Q\|_\Omega = \max_{L \in C(\Omega)^*} \left\{ L(f) : L|_{\Pi_{n-1}^d} = 0, \|L\|_{C(\Omega)^*} = 1 \right\}.$$

From this, we can immediately deduce the following relation between the two dual problems.

Lemma 4.3. *For $f(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha$ with a_α^* as defined in (4.2), and Ω compact, any optimal solution of problem (4.1) is also optimal for the dual problem in (4.3).*

Proof. Let L^* denote an optimal solution of problem (4.1). Then L^* is feasible for the dual (maximization) problem in (4.3). Using that $\sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* = 1$, one has

$$L^*(f) = L^* \left(\sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha \right) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* L^*(\mathbf{x}^\alpha) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \gamma^* = \gamma^* = \|P^*\|_\Omega,$$

¹In [13, Corollary 2], the objective is in fact given as $\sup |L(f)|$, but the absolute value may be omitted without loss of generality, since L is feasible for the dual problem if and only if $-L$ is feasible.

where the last equality is the strong duality relation in Theorem 4.1. Since the optimal value of the primal problem (4.3) is also $\|P^*\|_\Omega$ when $f(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* \mathbf{x}^\alpha$, it follows that L^* is an optimal solution for the dual problem in (4.3), as required. \square

In the next subsection, we review the fact that the optimal dual solutions may be assumed to be atomic (discrete) without loss of generality.

4.2. Atomic dual solutions and signatures. We first recall a classical interpolation formula for linear functionals, as given in [13, Corollary 3].

Theorem 4.4. *Let L be any linear functional on a finite dimensional subspace V of $C(\Omega)$, with Ω compact. Then there exist points $\omega_1, \dots, \omega_r \in \Omega$ with $r \leq \dim(V)$, and non-zero scalars τ_1, \dots, τ_r , such that, defining the point evaluation functionals*

$$L_{\omega_i}(f) = f(\omega_i) \quad i = 1, \dots, r,$$

one has

$$(4.4) \quad L = \sum_{i=1}^r \tau_i L_{\omega_i}, \quad \|L\|_{C(\Omega)^*} = \sum_{i=1}^r |\tau_i|.$$

As a consequence, there exists a discrete (atomic) solution to problem (4.1). We give a proof below only for the sake of completeness and later reference — the type of argument we use is classical.

Corollary 4.5. *There exist points $\omega_1, \dots, \omega_r \in \Omega$ with $r \leq \dim(\Pi_n^d)$, and non-zero scalars τ_1, \dots, τ_r with $\sum_{i=1}^r |\tau_i| = 1$, such that an optimal solution of problem (4.1) is given by (4.4). Moreover, the points $\omega_1, \dots, \omega_r$ are extremal points of any optimal solution P^* to problem (1.3), i.e.,*

$$|P^*(\omega_i)| = \|P^*\|_\Omega \quad \text{for all } i = 1, \dots, r.$$

Proof. Let L^* denote an optimal solution of problem (4.1). The first statement in the corollary now follows immediately from Theorem 4.4, applied to L^* , and using $V = \Pi_n^d$, so that

$$(4.5) \quad L^* = \sum_{i=1}^r \tau_i L_{\omega_i}, \quad \sum_{i=1}^r |\tau_i| = 1.$$

It remains to show that the points $\omega_1, \dots, \omega_r$ are extremal points of any optimal solution P^* to problem (1.3).

By Lemma 4.3, L^* is also an optimal solution for the dual problem in (4.3). Thus, for any optimal solution P^* of problem (1.3), one has

$$\|P^*\|_\Omega = L^*(P^*) = \sum_{i=1}^r \tau_i P^*(\omega_i).$$

Since $\sum_{i=1}^r |\tau_i| = 1$, $\|P^*\|_\Omega$ is a weighted average of the values $P^*(\omega_i)$ with $\tau_i > 0$ and $-P^*(\omega_i)$ with $\tau_i < 0$. As a consequence, we have

$$\|P^*\|_\Omega = \begin{cases} P^*(\omega_i) & \text{if } \tau_i > 0, \\ -P^*(\omega_i) & \text{if } \tau_i < 0, \end{cases}$$

completing the proof. \square

The extremal points in Corollary 4.5 are usually called the *support of an extremal signature*, defined along the line of [12, Section 2.2] as follows.

Definition 4.6. A signature with finite support $S \subset \Omega$ is simply a (partition) function from S to $\{\pm 1\}$. A signature σ with support S is said to be extremal for a subspace $V \subset C(\Omega)$ if there exist weights $\lambda_\omega > 0$, $\omega \in S$, such that $\sum_{\omega \in S} \lambda_\omega \sigma(\omega) v(\omega) = 0$ for all $v \in V$. A signature σ with support S is said to be associated with a function $g \in C(\Omega)$ if S is included in the set $\{\omega \in \Omega : |g(\omega)| = \|g\|_\Omega\}$ of extremal points of g and if $\sigma(\omega) = \text{sgn}(g(\omega))$ for all $\omega \in S$.

Thus, the result of Corollary 4.5 may be restated as the existence of a signature with support $S = \{\omega_1, \dots, \omega_r\}$ where $r \leq \dim(\Pi_n^d)$, so that this signature is extremal for Π_{n-1}^d , and associated with every optimal solution of problem (1.3). Therefore, we will simply refer to an *optimal signature* for problem (1.3). Formally we have the following result.

Proposition 4.7. Any atomic solution of the dual problem (4.1), say L^* of the form (4.5), gives rise to an optimal signature for problem (1.3), by setting:

$$S = \{\omega_1, \dots, \omega_r\}, \quad \sigma(\omega_i) = \text{sgn}(\tau_i), \quad \lambda_i = |\tau_i|, \quad i = 1, \dots, r.$$

This is essentially a reformulation of the well-known characterization of best approximation by polynomials in terms of the extremal signature ([12, Theorem 2.6]) for the dual problem.

4.3. Signatures for diagonally-determined sets. For the special case where Ω is a diagonally-determined set, we may now infer information about an optimal signature from Theorem 3.1 and Proposition 4.7. The idea of the proof is to construct an atomic optimal solution for the dual problem (4.1).

Theorem 4.8. Let Ω be a diagonally-determined set with $\text{diag}(\Omega) = [a, b]$ and vector \mathbf{v} . Then, an optimal signature for problem (1.3) is defined by a set S , of $n+1$ points

$$(4.6) \quad \omega_j = a\mathbf{1} + \frac{b-a}{2} \left(1 + \cos \left(\frac{(j-1)\pi}{n} \right) \right) \mathbf{1} \quad \text{for } j = 1, \dots, n+1,$$

where $\mathbf{1}$ again denotes the all-ones vector in \mathbb{R}^d , together with the partition function $\sigma : S \rightarrow \{-1, 1\}$, given by

$$\sigma(\omega_j) = \begin{cases} 1 & \text{if } j \text{ is odd} \\ -1 & \text{if } j \text{ is even,} \end{cases}$$

as well as $\lambda_1 = \lambda_{n+1} = \frac{1}{2}$, and $\lambda_i = 1$ if $i = 2, \dots, n$.

Proof. By Theorem 3.1, an optimal solution to problem (1.3) is given by the polynomial P^* in (3.2), namely

$$\mathbf{x} \mapsto P^*(\mathbf{x}) = \left(\frac{b-a}{2} \right)^n \frac{1}{2^{n-1}} T_n \left(-1 + 2 \frac{\langle \mathbf{v}, \mathbf{x} \rangle - a}{b-a} \right).$$

Using the well-known fact that the extremal points of T_n are given by the Gauss-Lobatto-Chebyshev points $\xi_j := \cos((j-1)\pi/n)$ for $j = 1, \dots, n+1$, we have that the points listed in (4.6) are extremal points of P^* . Indeed, using $\langle \mathbf{v}, \mathbf{1} \rangle = 1$, one has

$$\langle \mathbf{v}, \omega_j \rangle = a + \frac{b-a}{2} \left(1 + \cos \left(\frac{(j-1)\pi}{n} \right) \right) \quad \text{for } j = 1, \dots, n+1,$$

which is the same as

$$-1 + 2 \frac{\langle \mathbf{v}, \omega_j \rangle - a}{b-a} = \xi_j \quad \text{for } j = 1, \dots, n+1.$$

Next, we will construct an atomic solution of the dual problem (4.1), which will lead to the required optimal signature, by Proposition 4.7.

To this end, define the linear operator

$$L^* = \frac{1}{n} \sum_{i=1}^{n+1} (-1)^{i+1} \lambda_i L_{\omega_i},$$

so that $\|L^*\|_{C(\Omega)^*} = \frac{1}{n} \sum_{i=1}^{n+1} |(-1)^{i+1} \lambda_i| = 1$, by construction.

We claim that L^* is an optimal solution of the dual problem (4.1). We first verify feasibility. We will show that L^* vanishes on Π_{n-1}^d , as required, by using the Gauss-Lobatto-Chebyshev quadrature formula (see e.g. [12, Exercise 1.5.29]), which implies, for $k < n$:

$$0 = \int_{-1}^1 T_n(t) T_k(t) \frac{1}{\sqrt{1-t^2}} dt = \frac{\pi}{n} \sum_{j=1}^{n+1} \lambda_j T_n(\xi_j) T_k(\xi_j),$$

due to the orthogonality of Chebyshev polynomials. Using $T_n(\xi_j) = (-1)^{1+j}$ yields

$$(4.7) \quad L^*(T_k) = \frac{1}{n} \sum_{j=1}^{n+1} \lambda_j (-1)^{1+j} T_k(\xi_j) = 0 \quad \text{if } k < n.$$

Thus L^* vanishes on univariate (and separable) polynomials in Π_{n-1}^d . For the general case, it suffices to consider $d = 2$ and the monomial $p(x_1, x_2) = T_{k_1}(x_1) T_{k_2}(x_2)$, where $k_1 + k_2 < n$, and show that $L^*(p) = 0$. To this end, note that

$$\begin{aligned} L^*(p) &= \frac{1}{n} \sum_{j=1}^{n+1} \lambda_j (-1)^{1+j} T_{k_1}(\xi_j) T_{k_2}(\xi_j) \\ &= \frac{1}{n} \sum_{j=1}^{n+1} \lambda_j (-1)^{1+j} \frac{1}{2} (T_{|k_1-k_2|}(\xi_j) + T_{k_1+k_2}(\xi_j)) \\ &= \frac{1}{2n} \sum_{j=1}^{n+1} \lambda_j (-1)^{1+j} T_{|k_1-k_2|}(\xi_j) + \frac{1}{2n} \sum_{j=1}^{n+1} \lambda_j (-1)^{1+j} T_{k_1+k_2}(\xi_j) = 0, \end{aligned}$$

using (4.7). The result for general d now follows by induction.

Secondly, we verify that $L^*(\mathbf{x}^\alpha)$ is independent of α , when $|\alpha| = n$, say $L^*(\mathbf{x}^\alpha) = \hat{\gamma}$ whenever $\alpha \in \mathbb{N}_n^d$. This follows immediately from the construction of L^* , since all coordinates of the vectors ω_i are equal for all $i = 1, \dots, n+1$. Finally, to prove optimality, we note that, for P^* written in the form (4.2),

$$\|P^*\|_\Omega = L^*(P^*) = \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* L^*(\mathbf{x}^\alpha) = \hat{\gamma} \sum_{\alpha \in \mathbb{N}_n^d} a_\alpha^* = \hat{\gamma},$$

so, by the weak duality theorem, L^* is indeed optimal. The required result now follows from Proposition 4.7. \square

Remark 4.1. It is interesting to note that — for diagonally-determined sets — there exists an optimal signature for problem (1.3) with support size $n+1$. One can compare this to the general result for compact Ω of a signature support size of $\dim(\Pi_n^d) = \binom{n+d}{d}$. Remarkably the minimum support size is independent of d in the diagonally-determined case. Moreover, we did not use compactness of Ω in the proof of Theorem 4.8, whereas

compactness is needed for the general bound. Finally, note that the optimal signature in Theorem 4.8 depends on Ω only through its diagonal $[a, b]$, and is independent of \mathbf{v} .

5. DETECTING THE DIAGONALLY-DETERMINED PROPERTY

The goal of this section is to provide a numerical scheme to detect whether a given compact basic semi-algebraic set

$$(5.1) \quad \Omega := \{ \mathbf{x} \in \mathbb{R}^d : g_j(\mathbf{x}) \geq 0, \quad j = 1, \dots, m \}$$

for some given polynomials g_j , is diagonally-determined.

The *quadratic module* of g_1, \dots, g_m , denoted by $\mathcal{Q}(g_1, \dots, g_m)$, is defined as the set of all polynomials of the form

$$(5.2) \quad p = \sigma_0 + \sigma_1 g_1 + \dots + \sigma_m g_m$$

for some sum-of-squares polynomials $\sigma_0, \sigma_1, \dots, \sigma_m$. Clearly, if $p \in \mathcal{Q}(g_1, \dots, g_m)$, then $p(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \Omega$.

Assumption 5.1. *We assume that $\mathcal{Q}(g_1, \dots, g_m)$ satisfies the Archimedean condition, meaning that there exists $p \in \mathcal{Q}(g_1, \dots, g_m)$ such that*

$$\{ \mathbf{x} \in \mathbb{R}^d : p(\mathbf{x}) \geq 0 \} \text{ is compact.}$$

Under this assumption, if a polynomial is (strictly) positive on Ω , then it belongs to $\mathcal{Q}(g_1, \dots, g_m)$. We state this result — known as Putinar’s Positivstellensatz — as a theorem for later use.

Theorem 5.2 (Putinar [10]). *Assume Ω is of the form (5.1) and that Assumption 5.1 holds. If $p \in \Pi_n^d$ satisfies $p(\mathbf{x}) > 0$ for all $\mathbf{x} \in \Omega$, then $p \in \mathcal{Q}(g_1, \dots, g_m)$.*

We also introduce the concept of the *truncated quadratic module* of degree $k \in \mathbb{N}_0$, denoted by $\mathcal{Q}(g_1, \dots, g_m)_k$, as the subset of $\mathcal{Q}(g_1, \dots, g_m)$ where each term on the right-hand-side of (5.2) has degree at most k , i.e. $\deg(\sigma_0) \leq k$, and $\deg(\sigma_j g_j) \leq k$ for all $j = 1, \dots, m$. The membership problem for $\mathcal{Q}(g_1, \dots, g_m)_k$ may be reformulated as a semidefinite programming feasibility problem; see Lasserre [6], and the references therein.

Finding the diagonal $[a, b]$. We first note that one may find the diagonal of Ω by finding the real roots of the univariate polynomials

$$t \mapsto p_j(t) := g_j(t\mathbf{1}), \quad j = 1, \dots, m.$$

Indeed, for each p_j we may find the set of all (possibly unbounded) intervals where it is nonnegative, and subsequently take the intersection of all these sets for $j = 1, \dots, m$. If the result is a single interval $[a, b]$, then $\text{diag}(\Omega) = [a, b]$, and we may proceed with the next step, since the first property holds in Definition 2.1.

Once $[a, b]$ is known. We assume that we now know $[a, b] = \text{diag}(\Omega)$, and that $a < 0 < b$. The latter assumption is without loss of generality, due to Lemma 2.9. Then, to detect whether the second property holds in Definition 2.1, consider the optimization problem:

$$(5.3) \quad \rho = \min_{\mathbf{v}'} \{ |1 - \langle \mathbf{v}', \mathbf{1} \rangle| : a \leq \langle \mathbf{v}', \mathbf{x} \rangle \leq b \quad \forall \mathbf{x} \in \Omega \}.$$

Then Ω is diagonally-determined if and only if $\rho = 0$. Note that one may in fact omit the absolute value in the objective of (5.3), since $\mathbf{x} = b\mathbf{1} \in \Omega$, and $\langle \mathbf{v}', \mathbf{x} \rangle \leq b$ therefore implies $\langle \mathbf{v}', \mathbf{1} \rangle \leq 1$ for all \mathbf{v}' that are feasible for (5.3).

We now relax this problem to a sequence of semidefinite programs, using the approach introduced by Lasserre [6]. The key idea is to replace the nonnegativity conditions in (5.3) by the sufficient condition of membership of a truncated quadratic module. To this end, we introduce the two affine functions, for given $\mathbf{v}' \in \mathbb{R}^d$:

$$\mathbf{x} \mapsto b - \langle \mathbf{v}', \mathbf{x} \rangle = \ell_{\mathbf{v}', b}(\mathbf{x}), \quad \mathbf{x} \mapsto \langle \mathbf{v}', \mathbf{x} \rangle - a = \ell_{\mathbf{v}', a}(\mathbf{x}),$$

to obtain the reformulation of problem (5.3) as

$$\rho = \min_{\mathbf{v}'} \{ 1 - \langle \mathbf{v}', \mathbf{1} \rangle : \ell_{\mathbf{v}', a}(\mathbf{x}), \ell_{\mathbf{v}', b}(\mathbf{x}) \geq 0 \quad \forall \mathbf{x} \in \Omega \}.$$

For every $k \in \mathbb{N}_0$, we now define:

$$(5.4) \quad \rho_k = \min_{\mathbf{v}'} \{ 1 - \langle \mathbf{v}', \mathbf{1} \rangle : \ell_{\mathbf{v}', a}, \ell_{\mathbf{v}', b} \in \mathcal{Q}(g_1, \dots, g_m)_k \}.$$

Importantly, and as already mentioned, we emphasize that the optimization problem (5.4) is a (convex) semidefinite program that can be solved efficiently with arbitrary precision (fixed in advance), under mild assumptions, using, e.g., interior point methods; see Lasserre [6].

Proposition 5.3. *Assume Ω is of the form (5.1), that Assumption 5.1 holds, and that $\text{diag}(\Omega) = [a, b]$ where $a < 0 < b$. Then one has $\rho_k \in [0, 1]$ for all k . Moreover, if Ω is diagonally-determined, then the non-increasing sequence $\{\rho_k\}$ converges to zero. If $\rho_k = 0$ for some finite value of k , then the optimal $\mathbf{v}' \in \mathbb{R}^d$ of (5.4) certifies that Ω is diagonally-determined with vector $\mathbf{v} = \mathbf{v}'$ and diagonal $[a, b]$.*

Proof. For $k = 0$, a feasible solution of (5.4) is given by $\mathbf{v}' = \mathbf{0}$, since $a < 0 < b$, so that $\rho_0 \leq 1$. Consequently, $\rho_k \in [0, 1]$ for all k , since $\mathcal{Q}(g_1, \dots, g_m)_k \subseteq \mathcal{Q}(g_1, \dots, g_m)_{k+1}$ for all $k \in \mathbb{N}_0$. Assume now that Ω is diagonally-determined with vector \mathbf{v} and diagonal $[a, b]$. Then, for any $\varepsilon > 0$ and $\mathbf{x} \in \Omega$, one has

$$\langle (1 - \varepsilon)\mathbf{v}, \mathbf{x} \rangle \in [(1 - \varepsilon)a, (1 - \varepsilon)b] \subset (a, b),$$

where we again use that $a < 0 < b$. Thus, $\mathbf{x} \mapsto \langle (1 - \varepsilon)\mathbf{v}, \mathbf{x} \rangle - a \in \mathcal{Q}(g_1, \dots, g_m)$ and $\mathbf{x} \mapsto b - \langle (1 - \varepsilon)\mathbf{v}, \mathbf{x} \rangle \in \mathcal{Q}(g_1, \dots, g_m)$, by Theorem 5.2. This implies that, for sufficiently large $k^* \in \mathbb{N}_0$, $\mathbf{v}' = (1 - \varepsilon)\mathbf{v}$ is feasible for (5.4) with $k = k^*$. Thus $\rho_{k^*} \leq 1 - \langle (1 - \varepsilon)\mathbf{v}, \mathbf{1} \rangle = \varepsilon$, and therefore the sequence $\{\rho_k\}$ is non-increasing, and converges to zero. \square

Example 5.4. *It is insightful to see how the above procedure works in the case of the unit ball in the ∞ -norm, i.e.*

$$\Omega = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_\infty \leq 1\}.$$

By Proposition 2.5, and its proof, we know that Ω is diagonally-determined, with diagonal $[a, b] = [-1, 1]$ and vector $\mathbf{v} = \frac{1}{d}\mathbf{1}$. We may describe Ω as a semi-algebraic set in the form (5.1) by using the polynomials $g_j(\mathbf{x}) = 1 - x_j^2$ for $j = 1, \dots, d$, i.e. $m = d$. The next step is to find the diagonal of Ω by finding the roots of the univariate polynomials $p_j(t) := g_j(t\mathbf{1}) = 1 - t^2$ for $j = 1, \dots, d$. These polynomials are the same for all j and nonnegative on the interval $[-1, 1]$, so that we obtain $[a, b] = [-1, 1]$. The next step is to solve (5.4) for increasing values of $k \in \mathbb{N}$. Using the identity

$$1 \pm x = \frac{1}{2} ((1 \pm x)^2 + 1 - x^2), \quad x \in \mathbb{R},$$

one obtains that

$$b - \langle \mathbf{v}, \mathbf{x} \rangle = 1 - \frac{1}{d} \sum_{j=1}^d x_j = \frac{1}{d} \sum_{j=1}^d (1 - x_j) = \frac{1}{2d} \sum_{j=1}^d ((1 - x_j)^2 + g_j(\mathbf{x})).$$

Similarly one has

$$\langle \mathbf{v}, \mathbf{x} \rangle - a = \frac{1}{2d} \sum_{j=1}^d ((1 + x_j)^2 + g_j(\mathbf{x})) \in \mathcal{Q}(g_1, \dots, g_d)_2.$$

Thus, for $k = 2$, we find that $\mathbf{v}' = \mathbf{v}$ is feasible in (5.4), so that $\rho_2 = 0$.

In the example we have finite convergence of the sequence $\{\rho_k\}$, but it is an open question when this occurs. There are known sufficient conditions for the finite convergence of Lasserre-type semidefinite programming hierarchies, e.g. [8], but these are not satisfied for our problem in general.

6. CONCLUSION AND DISCUSSION

The extremal problem (1.3) we considered in this paper relies on a specific generalization of the concept of a monic univariate polynomial to the multivariate case. In particular, in the multivariate case we require that the leading coefficients (in the standard monomial basis) sum to one, leading to the introduction of the set Π_n^* in Definition 1.1. There are other natural choices, like also requiring the leading coefficients to be nonnegative. Having said that, the least Chebyshev polynomial in Theorem 3.1 already has nonnegative leading coefficients for many diagonally-determined sets, e.g. the unit ball in a monotone norm $\|\cdot\|$. To see this, let Ω be such a unit ball, centered at the origin. By Proposition 2.5, Ω is then diagonally-determined with diagonal $[a, b] = [-1/\|\mathbf{1}\|, 1/\|\mathbf{1}\|]$ and allows an entrywise-nonnegative vector \mathbf{v} . By Theorem 3.1, the least Chebyshev polynomial in this case is given by

$$\mathbf{x} \mapsto P^*(\mathbf{x}) = \left(\frac{1}{\|\mathbf{1}\|} \right)^n \frac{1}{2^{n-1}} T_n(\|\mathbf{1}\| \langle \mathbf{v}, \mathbf{x} \rangle).$$

Since $\frac{1}{2^{n-1}} T_n$ is a monic univariate polynomial of degree n , the leading coefficients of the least Chebyshev polynomial P^* are therefore obtained from the multinomial expansion of $\mathbf{x} \mapsto \langle \mathbf{v}, \mathbf{x} \rangle^n$. Since \mathbf{v} is entrywise-nonnegative by assumption, it follows that these leading coefficients are also nonnegative. In this sense our generalization of monic polynomials to the multivariate case is a natural one, but it remains an interesting question which diagonally-determined sets allow a entrywise-nonnegative vector \mathbf{v} .

A second point to mention, is that we have not explored the geometric and invariance properties of diagonally-determined sets to any great extent in this paper, and this deserves to be investigated further. The same may be said for the decision problem of whether a given set is diagonally-determined. If Ω is a compact semi-algebraic set, one may formulate an algorithmic procedure to decide this question, as we did in Section 5. An interesting special case is the (complexity) question: may one decide in polynomial time if a given polyhedron is diagonally-determined?

Finally, we note that our main results in Theorems 3.1 and 4.8 are relatively straightforward to prove, once the correct formulations have been discovered. We ‘discovered’ these results by doing numerical experiments on least Chebyshev polynomials for the Euclidean unit ball, using the tools described in our earlier paper [4]. Thus we observed

the structure of the atomic dual solution in Theorem 4.8, and we could infer the main results from there. In the immortal words of Bernhard Riemann:²

If only I had the theorems! Then I should find the proofs easily enough.

Acknowledgments. This work is the result of a collaboration made possible by the SQuaRE program at the American Institute of Mathematics (AIM). We are truly grateful to AIM for the supportive and mathematically rich environment they provided. In addition, we owe thanks to several funding agencies, as M. D. is supported by the Australian Research Council Discovery Early Career Award DE240100674, S. F. is partially supported by grants from the National Science Foundation (DMS-2053172) and from the Office of Naval Research (N00014-20-1-2787), E. de K. is supported by grants from the Dutch Research Council (NWO) (OCENW.M.23.050 and OCENW.GROOT.2019.015), J. B. L. is supported by the AI Interdisciplinary Institute through the French program “Investing for the Future PI3A” (ANR-19-PI3A-0004) and by the National Research Foundation, Singapore, through the DesCartes and Campus for Research Excellence and Technological Enterprise (CREATE) programs, and Y. X. is partially supported by the Simons Foundation (grant #849676).

REFERENCES

- [1] N. N. Andreev and V. A. Yudin, Polynomials of least deviation from zero, and Chebyshev-type cubature formulas, *Tr. Mat. Inst. Steklova* **232** (2001), 45–57.
- [2] N. N. Andreev and V. A. Yudin, Best approximation of polynomials on the sphere and on the ball, in *Recent progress in multivariate approximation*, 23–30; Internat. Ser. Numer. Math., **137**, Birkhäuser, Basel, 2001.
- [3] B. D. Bojanov, W. Haussmann, and G. P. Nikolov, Bivariate polynomials of least deviation from zero, *Canad. J. Math.* **53** (2001), 489–505.
- [4] M. Dressler, S. Foucart, E. de Klerk, M. Joldes, J. B. Lasserre, Y. Xu, *Optimization-Aided Construction of Multivariate Chebyshev Polynomials*, *J. Approx. Theory*, **305** (2025), 106116.
- [5] W. B. Gearhart, Some Chebyshev approximations by polynomials in two variables, *J. Approx. Theory*, **8** (1973), 195–209.
- [6] J.B. Lasserre. *Global optimization with polynomials and the problem of moments*. SIAM Journal on Optimization **11** (2001), 796–817.
- [7] D. J. Newman and Y. Xu, Tchebycheff polynomials on a triangular region, *Constr. Approx.* **9** (1993), no. 4, 543–546.
- [8] J. Nie. Optimality conditions and finite convergence of Lasserre’s hierarchy. *Math. Program.* **146** (2014), 97–121.
- [9] I. Moale and F. Peherstorfer, An explicit class of min–max polynomials on the ball and on the sphere, *J. Approx. Theory* **163** (2011), no. 6, 724–737.
- [10] M. Putinar. *Positive polynomials on compact semi-algebraic sets*. Indiana University Mathematics Journal, **42/3** (1993), 969–984.
- [11] M. Reimer, On multivariate polynomials of least deviation from zero on the unit ball, *Math. Z.* **153** (1977), no. 1, 51–58.
- [12] T. J. Rivlin. Chebyshev Polynomials. Second edition. Courier Dover Publications, 2020.
- [13] T. J. Rivlin and H. S. Shapiro, A unified approach to certain problems of approximation and minimization, *J. Soc. Indust. Appl. Math.* **9** (1961), 670–699.
- [14] A. Shapiro. On duality theory of conic linear problems, in *Semi-Infinite Programming: Recent Advances* (M.Á. Goberna and M.A. López, eds.), Springer, 135–165 (2001).
- [15] J. M. Sloss, Chebyshev approximation to zero, *Pacific J. Math.* **15** (1965), 305–313.
- [16] Y. Xu, On polynomials of least deviation from zero in several variables, *Experiment. Math.* **13** (2004), no. 1, 103–112.

²As quoted in I. Lakatos, *Proofs and Refutations* (Cambridge University Press, 1976).

- [17] Y. Xu, Best approximation of monomials in several variables, *Rend. Cir. Mat. Palermo (2) Suppl.* (2005), no. 76, 129–155.
- [18] V. A. Yudin, Best approximation to monomials on a cube, *Matematicheskii Sbornik* **199**, (2008), 149–160.

MAREIKE DRESSLER, SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NEW SOUTH WALES, SYDNEY, NSW 2052, AUSTRALIA.

Email address: `m.dressler@unsw.edu.au`

SIMON FOUCART, DEPARTMENT OF MATHEMATICS, TEXAS A&M UNIVERSITY, UNITED STATES.

Email address: `foucart@tamu.edu`

MIOARA JOLDES, LAAS-CNRS, FRANCE.

Email address: `joldes@laas.fr`

ETIENNE DE KLERK, DEPARTMENT OF ECONOMETRICS AND OPERATIONS RESEARCH, TILBURG UNIVERSITY, THE NETHERLANDS.

Email address: `E.deKlerk@tilburguniversity.edu`

JEAN-BERNARD LASSERRE, LAAS-CNRS AND TOULOUSE SCHOOL OF ECONOMICS, FRANCE.

Email address: `lasserre@laas.fr`

YUAN XU, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF OREGON, UNITED STATES.

Email address: `yuan@uoregon.edu`