JEMS

Manuel Bodirsky · Mario Kummer · Andreas Thom

Spectrahedral shadows and completely positive maps on real closed fields

Abstract. In this article we develop new methods for exhibiting convex semialgebraic sets that are not spectrahedral shadows. We characterize when the set of nonnegative polynomials with a given support is a spectrahedral shadow in terms of sums of squares. As an application we prove that the cone of copositive matrices of size $n \geq 5$ is not a spectrahedral shadow, answering a question of Scheiderer. Our arguments are based on the model-theoretic observation that any formula defining a spectrahedral shadow must be preserved by every unital $\mathbb{R}$-linear completely positive map $R \to R$ on a real closed field extension R of $\mathbb{R}$.

Keywords: spectrahedral shadow, real closed field, completely positive map, relational structure.

1. Introduction

A set $S \subseteq \mathbb{R}^n$ is called a *spectrahedron* if

$$S = \{(x_1, \ldots, x_n) \mid A_0 + A_1 x_1 + \cdots + A_n x_n \succeq 0\}$$

for some $k \times k$ real symmetric matrices $A_0, A_1, \ldots, A_n$. Here $A \succeq 0$ says that the real symmetric matrix A is positive semidefinite, and the expression $A_0 + A_1 x_1 + \cdots + A_n x_n \succeq 0$ is called a *linear matrix inequality (LMI) of size k*. Note that spectrahedra are convex and *semialgebraic*, i.e., definable over $\mathbb{R}$ by a first-order formula in the language of real closed fields (i.e., the language of ordered rings).

There has been ample interest in representing convex semialgebraic sets $K \subseteq \mathbb{R}^n$ as a *spectrahedral shadow*. This amounts to finding real symmetric matrices $A_0, \ldots, A_n$,

Manuel Bodirsky: Institut für Algebra, Technische Universität Dresden, 01062 Dresden, Germany;
manuel.bodirsky@tu-dresden.de

Mario Kummer: Institut für Geometrie, Technische Universität Dresden, 01062 Dresden, Germany; mario.kummer@tu-dresden.de

Andreas Thom: Institut für Geometrie, Technische Universität Dresden, 01062 Dresden, Germany; andreas.thom@tu-dresden.de

$B_1, \ldots, B_m$ of the same size such that

$$K = \{x \in \mathbb{R}^n \mid \exists y \in \mathbb{R}^m \colon A_0 + x_1 A_1 + \cdots + x_n A_n + y_1 B_1 + \cdots + y_m B_m \succeq 0\}.$$

The interest stems primarily from the fact that there are efficient methods for approximately optimizing a linear function over such a set. It was shown by Helton and Nie [11] that every convex semialgebraic set satisfying certain smoothness and curvature conditions is a spectrahedral shadow. Moreover, Scheiderer [30] proved that every convex semialgebraic set in the plane is a spectrahedral shadow. On the other hand, not every convex semialgebraic set is a spectrahedral shadow. The first counterexamples were constructed by Scheiderer [31]. Further examples were found in [4, 8], both using techniques similar to Scheiderer's in his seminal work.

Recent efforts also focus on finding a representation as spectrahedral shadow with several LMIs of a bounded size. For instance Averkov [2] showed that the cone of $k \times k$ real symmetric positive semidefinite matrices has no semidefinite extended representation with LMIs of size less than k (for a formal definition, see Section 2). A generalization of this result in terms of convex geometry was given by Saunderson [29]. On the positive side, Scheiderer [32] proved that for every closed convex semialgebraic set $K \subseteq \mathbb{R}^2$ LMIs of size 2 are sufficient.

Our approach to these questions is based on an easy consequence of the Tarski principle for real closed fields. The question whether a given semialgebraic set $K \subseteq \mathbb{R}^n$ is a spectrahedral shadow is equivalent to the question whether K has a primitive positive definition over the structure $\mathfrak{A}$ with domain $\mathbb{R}$ and with a relation $S \subseteq \mathbb{R}^n$ for every spectrahedron S. By the Tarski principle for real closed fields, we may equivalently study this question over real closed field extensions R of $\mathbb{R}$ instead of $\mathbb{R}$. More precisely, if ϕ is a formula in the language of real closed fields that defines K, let K^* be the set defined by ϕ over R. Likewise, let $\mathfrak{A}^*$ denote the structure with domain R which contains, for every formula ψ that defines a spectrahedron over $\mathbb{R}$, the relation defined by ψ over R. Then K has a primitive positive definition in $\mathfrak{A}$ if and only if K^* has a primitive positive definition in $\mathfrak{A}^*$. A map $h \colon R \to R$ *preserves* a relation $K \subseteq R^n$ if $(h(a_1), \ldots, h(a_n)) \in K$ whenever $(a_1, \ldots, a_n) \in K$. We use the easy fact that K^* does not have an existential positive definition in $\mathfrak{A}^*$ (and hence in particular no primitive positive definition) if there exists a map $h \colon R \to R$ that preserves all the relations of $\mathfrak{A}^*$ but does not preserve K^*. In Appendix B we discuss this approach in more detail and prove a converse of the last statement.

We show that a map $L \colon R \to R$ preserves $\mathfrak{A}^*$ if and only if it is unital, $\mathbb{R}$-linear, and completely positive in the sense of operator theory. Thus, in order to prove that K is not a spectrahedral shadow, it suffices to exhibit a unital, $\mathbb{R}$-linear, and completely positive map $L \colon R \to R$ on a real closed field extension $R \supseteq \mathbb{R}$ that does not preserve the set defined by ϕ. In this way we prove that the set of copositive matrices of size n is not a spectrahedral shadow whenever $n \geq 5$, a family of convex semialgebraic sets that resisted Scheiderer's efforts [31, Question 5.2]. A real symmetric matrix A is *copositive* if $x^\top A x \geq 0$ for every nonnegative real vector x. Note that for $n \leq 4$ it follows from the main result in [7] that this set is a spectrahedral shadow.

Note that convexity is preserved not only by conjunction (i.e., intersection) and existential quantification (i.e., projections), but also by universal quantification. In the light of the results mentioned above one may ask whether every convex semialgebraic set has a definition using a formula from *positive first-order logic* over the structure formed by all spectrahedra. Positive first-order logic is the extension of existential positive logic by universal quantification (so only negation from first-order logic is prohibited). Our proof shows that the convex semialgebraic sets under consideration are not definable from spectrahedra even in this more expressive logic. The reason is that the map $L: R \to R$ that we construct is *surjective*. Hence, we may combine the Tarski principle with the easy direction of Lyndon's preservation theorem (see, e.g., [13, Corollary 10.3.5]) to conclude that for example the set of copositive matrices does not have a definition from spectrahedra in positive first-order logic.

Apart from the Tarski principle, the only ingredient in our proof is a separation theorem (Theorem 3.7) that follows from the results in [21]. We will give a self-contained proof in Appendix A. In particular, we do not make use of the theory developed in [31]. Our result on copositive matrices is a consequence of the following more general result:

Theorem (see Theorem 3.15 below). *Let $S \subset \mathbb{Z}_{\geq 0}^n$ be a finite set and let $\mathrm{P}_+(S)$ be the set of all nonnegative polynomials with coefficients from $\mathbb{R}$ whose support is contained in S. Then the following are equivalent:*

(1) $\mathrm{P}_+(S)$ *is a spectrahedral shadow.*

(2) *There exists an integer $d > 0$ such that $p(x_1^d, \ldots, x_n^d)$ is a sum of squares for all $p \in \mathrm{P}_+(S)$.*

Finally, in Section 4 we reprove Averkov's above mentioned result in a similar fashion. Here we use results about positive definite maps [15].

We now give a brief outline of our proof strategy. We first observe in Section 2 that a map $L: R \to R$ for some real closed field extension R of $\mathbb{R}$ preserves $\mathfrak{A}^*$ if and only if L is unital, $\mathbb{R}$-linear, and completely positive in the sense that applying L entrywise to a positive semidefinite symmetric matrix over R gives a positive semidefinite matrix. When $R = R'[[\varepsilon^\Lambda]]$ is the field of Hahn series with value group Λ and coefficients in another real closed field R', we consider maps $L_f: R \to R$ of the form $L_f(\sum_{a \in \Lambda} c_a \varepsilon^a) = \sum_{a \in \Lambda} f(a) c_a \varepsilon^a$ for functions $f: \Lambda \to R'$. On the other hand, extending f linearly, we obtain a linear functional $T_f: R'[\Lambda] \to R'$ on the group ring of Λ. The crucial insight of Section 2 is that L_f is completely positive whenever T_f is positive on every nonzero square in $R'[\Lambda]$. In Section 3 we focus on the case $\Lambda = \mathbb{Q}^n$. Elements of the group ring $R'[\mathbb{Q}^n]$ can be regarded as functions on $(R'_{>0})^n$. If such a function takes only nonnegative values but is not a sum of squares in $R'[\mathbb{Q}^n]$, then it can be separated by a linear functional T_f from the cone of sums of squares after a base change to a larger real closed field by Theorem 3.7. The corresponding map L_f is then completely positive but does not preserve a certain convex cone of nonnegative polynomials. Applying this construction to the so-called Horn polynomial proves that the cone of copositive matrices of size $n \geq 5$ is not a spectrahedral shadow.

2. k-positive maps

In this section let R always be a real closed field extension of $\mathbb{R}$. Let Γ_k be the set of all first-order formulas in the language of real closed fields with parameters from $\mathbb{R}$ that express a linear matrix inequality of size k with coefficients from $\mathbb{R}$, and let $\Gamma = \bigcup_{k=1}^\infty \Gamma_k$. Let $\mathfrak{A}_k$ be the relational structure with domain $\mathbb{R}$ whose relations consist of the sets defined by the formulas in Γ_k. We follow standard terminology in model theory (see, e.g., [13, 33]). For the convenience of the reader, we have collected the relevant concepts in Appendix B.1. The signature of $\mathfrak{A}_k$ should not be confused with the signature of real closed fields. If a set K has a primitive positive definition over $\mathfrak{A}_k$ then following Averkov [2] we say that K has a *semidefinite extended formulation with LMIs of size k*.

We say that a map $h\colon R \to R$ *preserves* a formula in the language of real closed fields if it preserves the relation defined by the formula over R. The map h is called *unital* if it preserves the formula $x = 1$, and for $R' \subseteq R$ it is called R'-*linear* if it preserves the formula $z = \lambda \cdot x + \mu \cdot y$ for all $\lambda, \mu \in R'$.

Definition 2.1. Let $k \in \mathbb{Z}_{>0}$. A map $L\colon R \to R$ is called k-*positive* if for every $k \times k$ symmetric positive semidefinite matrix $A = (a_{ij})_{i,j}$ with $a_{ij} \in R$ the matrix $L(A) := (L(a_{ij}))_{i,j}$ is positive semidefinite as well. For technical reasons, we declare every map $L\colon R \to R$ to be 0-positive. If L is k-positive for every k, then, following the nomenclature of operator theory (see, e.g., [6, Section 1.5]), we say that L is *completely positive*.

Remark 2.2. Note that if $L\colon R \to R$ is k-positive, then it is also k'-positive for all $k' \in \{0, \dots, k\}$ since principal submatrices of positive semidefinite matrices are again positive semidefinite.

Lemma 2.3. *Let $k \in \mathbb{Z}_{>0}$. A map $L\colon R \to R$ preserves all formulas in Γ_k (respectively Γ) if and only if L is unital, $\mathbb{R}$-linear, and k-positive (respectively completely positive).*

Proof. Assume that L preserves all formulas in Γ_k. The formula $x = 1$ is equivalent to

$$(x - 1) \succeq 0 \wedge (1 - x) \succeq 0.$$

Thus $L(1) = 1$. A similar argument shows that $z = \lambda \cdot x + \mu \cdot y$ is in Γ_k for all $\lambda, \mu \in \mathbb{R}$. To prove that L is k-positive, note that a symmetric matrix $A = (a_{ij})_{1 \le i,j \le k}$ is positive semidefinite if and only if $(a_{11}, a_{12}, \dots, a_{1k}, a_{22}, a_{23}, \dots, a_{kk})$ satisfies the linear matrix inequality $\sum_{1 \le i \le j \le k} x_{ij} E_{ij} \succeq 0$ where E_{ij} has entry 1 at position i, j and at position j, i and entry 0 at all other positions. Since L preserves this formula, we see that $L(A) \succeq 0$ whenever $A \succeq 0$.

Conversely, assume that L is unital, $\mathbb{R}$-linear, and k-positive. Consider $k \times k$ real symmetric matrices $A_0, A_1, \dots, A_n$. Let

$$K = \{x \in R^n \mid A_0 + x_1 A_1 + \cdots + x_n A_n \succeq 0\}$$

and let $x = (x_1, \dots, x_n) \in K$. Then $L(x) = (L(x_1), \dots, L(x_n))$ is also in K. Indeed, by

the properties of L we have

$$A_0 + L(x_1)A_1 + \cdots + L(x_n)A_n = L(A_0 + x_1 A_1 + \cdots + x_n A_n) \succeq 0. \qquad \blacksquare$$

Lemma 2.4. *Let $R \subset R'$ be a field extension and $L: R' \to R'$ be an R-linear map. Let A be an $m \times n$ matrix with entries in R'. Let M and N be a $\times m$ and $n \times b$ matrices with entries in R. Then $L(MAN) = ML(A)N$.*

Proof. This follows directly from R-linearity. $\qquad \blacksquare$

Now let Λ a divisible ordered abelian group. We consider $H = R[[\varepsilon^\Lambda]]$, the field of Hahn series over R with value group Λ and the valuation $v: H \to \Lambda \cup \{\infty\}$ given by $v(\sum_{e \in \Lambda} c_e \varepsilon^e) = \min \{e \in \Lambda \mid c_e \neq 0\}$. Since H is real closed [1, p. 218, (2)], it is equipped with an order $\leq$ defined by $x \leq y \Leftrightarrow \exists z \, (z^2 + x = y)$. Further, let $\mathfrak{o}$ be the valuation ring of H, that is, all elements of H with nonnegative valuation. Then $\mathfrak{o}$ is a local ring with residue field R, the maximal ideal consisting of all elements with positive valuation. We denote by $\pi: \mathfrak{o} \to R$ the natural projection.

Lemma 2.5. *Let $L: H \to H$ be R-linear and $(k-1)$-positive and suppose that for all $a_1, \ldots, a_k \in H_{>0}$ with $v(a_1) < \cdots < v(a_k)$ the matrix $L(a \cdot a^\top)$, where $a = (a_1, \ldots, a_k)^\top$, is positive semidefinite. Then L is k-positive.*

Proof. Let B be a $k \times k$ symmetric positive semidefinite matrix over H. Then

$$B = \sum_{i=1}^k b_i \cdot b_i^\top$$

for some column vectors $b_1, \ldots, b_k \in H^k$. Since L is R-linear, we have

$$L(B) = \sum_{j=1}^k L(b_i \cdot b_i^\top).$$

Since the sum of positive semidefinite matrices is positive semidefinite, it suffices to show that $L(a \cdot a^\top)$ is positive semidefinite for all $a = (a_1, \ldots, a_k)^\top$ with $a_1, \ldots, a_k \in H$. Since L is $(k-1)$-positive, we may assume that $a_i \neq 0$ for all i. Let $M \in \mathbf{GL}_n(R)$. The matrix $L(a \cdot a^\top)$ is positive semidefinite if and only if $ML(a \cdot a^\top)M^\top$ is positive semidefinite, which equals $L((Ma) \cdot (Ma)^\top)$ by Lemma 2.4. Thus, after choosing M to be a suitably signed permutation matrix, we may assume that $a_1 \geq \cdots \geq a_k > 0$. This implies that $v(a_1) \leq \cdots \leq v(a_k)$. Now, after choosing M to be a product of suitable elementary matrices, we can make each inequality strict. $\qquad \blacksquare$

Lemma 2.6. *Let A be a symmetric $n \times n$ matrix with entries in $\mathfrak{o}$.*

(1) *If $A = (a_{ij})_{1 \leq i,j \leq n}$ is positive semidefinite, then $\pi(A) = (\pi(a_{ij}))_{1 \leq i,j \leq n}$ is positive semidefinite.*

(2) *If $\pi(A)$ is positive definite, then A is positive definite.*

Proof. We use the well-known fact that a symmetric matrix is positive semidefinite (resp. positive definite) if and only if all its principal minors are nonnegative (resp. positive); see, e.g., [14, Theorem 7.2.5]. Let $x \in \mathfrak{o}$ with $x \geq 0$. Since every nonnegative element in $\mathfrak{o}$ has a square root in $\mathfrak{o}$, we find that $\pi(x) \in R$ has a square root as well and hence $\pi(x) \geq 0$. If B is a principal submatrix of A, then $\det(\pi(B)) = \pi(\det(B))$. Therefore, if every principal minor of A is nonnegative, the same is true for every principal minor of $\pi(A)$. This shows (1).

For (2), assume that A is not positive definite. Then there is a principal minor of A which is nonpositive. The same minor of $\pi(A)$ will also be nonpositive, showing that $\pi(A)$ is not positive definite. ∎

Remark 2.7. Note that the converse of neither (1) nor (2) in Lemma 2.6 is true, as witnessed by the 1×1 matrices $(-\varepsilon)$ and (ε).

For any function $f \colon \Lambda \to R$ we define the function $L_f \colon H \to H$ via

$$L_f\Big(\sum_{e \in \Lambda} c_e \varepsilon^e\Big) = \sum_{e \in \Lambda} f(e) c_e \varepsilon^e.$$

We want to derive a criterion on f for L_f to be k-positive. We start with the following refinement of Lemma 2.5.

Lemma 2.8. *Let $L \colon H \to H$ be R-linear and $(k-1)$-positive and suppose that for all $a_1, \ldots, a_k \in H_{>0}$ with $v(a_1) < \cdots < v(a_k)$ and $\pi(a_i / \varepsilon^{v(a_i)}) = 1$ the matrix $L(a \cdot a^\top)$ is positive semidefinite, where $a = (a_1, \ldots, a_k)^\top$. Then L is k-positive.*

Proof. This follows in the same way as Lemma 2.5 by another multiplication with a suitable diagonal matrix. ∎

Corollary 2.9. *Let $L \colon H \to H$ be R-linear and suppose that for all $a_1, \ldots, a_k \in H_{>0}$ with $v(a_1) < \cdots < v(a_k)$ and $\pi(a_i / \varepsilon^{v(a_i)}) = 1$ the matrix $L(a \cdot a^\top)$ is positive semidefinite, where $a = (a_1, \ldots, a_k)^\top$. Then L is k-positive.*

Proof. We prove by induction on i that L is i-positive for every $i \in \{0, \ldots, k\}$. The case $i = 0$ is trivial. Let $i \in \{1, \ldots, k\}$ and assume that L is $(i-1)$-positive. Let $a_1, \ldots, a_k \in H_{>0}$ with $v(a_1) < \cdots < v(a_k)$ and $\pi(a_i / \varepsilon^{v(a_i)}) = 1$. By assumption $L(a \cdot a^\top)$ is positive semidefinite. This implies in particular that the top left $i \times i$ submatrix of $L(a \cdot a^\top)$ is positive semidefinite as well. Thus, Lemma 2.8 proves that L is i-positive. ∎

For later reference we also note the following condition for L_f being a bijection.

Lemma 2.10. *Let $f \colon \Lambda \to R$ be such that $L_f \colon H \to H$ is unital and 2-positive. Then L_f is bijective.*

Proof. Let $a \in \Lambda$ and consider the positive semidefinite matrix

$$A = \begin{pmatrix} \varepsilon^a & 1 \\ 1 & \varepsilon^{-a} \end{pmatrix}.$$

Since L_f is unital, we have

$$L_f(A) = \begin{pmatrix} f(a)\varepsilon^a & 1 \\ 1 & f(-a)\varepsilon^{-a} \end{pmatrix}.$$

Since L_f is 2-positive, this shows that $f(a)f(-a) \geq 1$, so in particular $f(a) \neq 0$. Now one can define

$$g : \Lambda \to R, \quad a \mapsto 1/f(a),$$

and one has $L_f \circ L_g = L_g \circ L_f = \mathrm{id}_H$. ∎

Definition 2.11. A function $f \colon \Lambda \to R$ is called *k-positive semidefinite* if for all $x_1, \ldots, x_k \in \Lambda$ the matrix $(f(x_i + x_j))_{i,j}$ is positive semidefinite, and f is *k-positive definite* if for all pairwise different $x_1, \ldots, x_k \in \Lambda$ this matrix is positive definite. Finally, $f \colon \Lambda \to R$ is *positive semidefinite* (resp. *positive definite*[1]) if f is k-positive semidefinite (resp. k-positive definite) for all $k \in \mathbb{Z}_{\geq 0}$.

Remark 2.12. If $f_1, f_2 \colon \Lambda \to R$ are k-positive (semi-)definite and $\lambda, \mu \in R_{>0}$, then λf_1, μf_2 and $\lambda f_1 + \mu f_2$ are k-positive (semi-)definite as well. Indeed, this follows from the fact that the cone of positive (semi-)definite matrices is convex.

The motivation for considering Definition 2.11 is the following result.

Theorem 2.13. *Let* $f \colon \Lambda \to R$ *and consider* $L_f \colon H \to H$ *as defined above.*

(1) *If* f *is k-positive definite, then* L_f *is k-positive.*

(2) *If* L_f *is k-positive, then* f *is k-positive semidefinite.*

Proof. In order to prove (2) let $v_1, \ldots, v_k \in \Lambda$. We have to prove that $(f(v_i + v_j))_{i,j}$ is positive semidefinite. Let $a = (\varepsilon^{v_1}, \ldots, \varepsilon^{v_k})^\top$. By assumption $L_f(a \cdot a^\top)$ is positive semidefinite. Let $b = (\varepsilon^{-v_1}, \ldots, \varepsilon^{-v_k})^\top$ and let $B = b \cdot b^\top$, which is positive semidefinite. Then the Hadamard product C of $L_f(a \cdot a^\top)$ with B is positive semidefinite as well by the Schur product theorem. Hence, Lemma 2.6 (1) implies that $\pi(C)$ is positive semidefinite. One calculates that $\pi(C)$ is the matrix $(f(v_i + v_j))_{1 \leq i,j \leq k}$, which proves (2).

We will prove (1) by applying Corollary 2.9. To that end let $a_1, \ldots, a_k \in H_{>0}$ be such that $v(a_1) < \cdots < v(a_k)$ and $\pi(a_i/\varepsilon^{v(a_i)}) = 1$. We have to show that the matrix $L_f(a \cdot a^\top)$ is positive semidefinite, where $a = (a_1, \ldots, a_k)^\top$. Let $b = (\varepsilon^{-v(a_1)}, \ldots, \varepsilon^{-v(a_k)})^\top$ and $B = b \cdot b^\top$. Since B is positive semidefinite of rank 1 with nonzero entries, the matrix $L_f(a \cdot a^\top)$ is positive semidefinite if and only if its Hadamard product C with B is positive semidefinite. As above one calculates that $\pi(C)$ is the matrix $(f(v(a_i) + v(a_j)))_{1 \leq i,j \leq k}$, which is positive definite since f is k-positive definite and $v(a_1), \ldots, v(a_k)$ are pairwise different. Thus, C is positive definite by Lemma 2.6 (2). Therefore, $L_f(a \cdot a^\top)$ is positive semidefinite and (1) follows from Corollary 2.9. ∎

---

[1]In [3, p. 88], positive definite functions are called *positive definite in the semigroup sense*; we use the shorter form since no confusion can arise in our context.

Remark 2.14. Another point of view on positive semidefinite functions can be taken by looking at $R[\Lambda]$. The group ring $R[\Lambda]$ can be identified with the subring of $H = R[[\varepsilon^{\Lambda}]]$ consisting of series with finite support. Now any function $f\colon \Lambda \to R$ can be extended linearly to an R-linear functional $T_f\colon R[\Lambda] \to R$. Any element of $R[\Lambda]$ is of the form $a = \sum_{i=1}^{k} c_i \varepsilon^{a_i}$ for some $k \in \mathbb{N}$, $(c_1, \ldots, c_k) \in R^k$, and $(a_1, \ldots, a_k) \in \Lambda^k$. Then we have

$$T_f(a^2) = T_f\Big(\sum_{i,j=1}^{k} c_i c_j \varepsilon^{a_i + a_j} \Big) = \sum_{i,j}^{k} c_i c_j f(a_i + a_j)$$

$$= \begin{pmatrix} c_1 & \ldots & c_k \end{pmatrix} \cdot (f(a_i + a_j))_{1 \le i,j \le k} \cdot \begin{pmatrix} c_1 \\ \vdots \\ c_k \end{pmatrix}.$$

From this we see that f is positive semidefinite (resp. positive definite) if and only if $T_f(a^2) \ge 0$ (resp. $T_f(a^2) > 0$) for all $a \in R[\Lambda] \setminus \{0\}$.

For the rest of this section we focus on the case $\Lambda = \mathbb{R}$. Let $\mathcal{H}(\mathbb{C})$ be the ring of holomorphic functions on $\mathbb{C}$. We have a homomorphism of $\mathbb{R}$-algebras $\Psi\colon \mathbb{R}[\Lambda] \to \mathcal{H}(\mathbb{C})$ defined as follows. For $a \in \mathbb{R}$, set $\Psi(\varepsilon^a)$ to be the holomorphic function $z \mapsto \exp(az)$ and extend linearly to all of $\mathbb{R}[\Lambda]$.

Lemma 2.15. *Let $g \in \mathbb{R}[\Lambda]$. If $\Psi(g)$ vanishes on $[a, b]$ for $a < b$, then $g = 0$ as element of $\mathbb{R}[\Lambda]$. In particular, the homomorphism Ψ is injective.*

Proof. Let $g = \sum_{i=1}^{n} c_i \varepsilon^{a_i}$ for some $c_i, a_i \in \mathbb{R}$, $i = 1, \ldots, n$, with $a_1 < \cdots < a_n$. The identity theorem from complex analysis implies that $\Psi(g)$ vanishes on all of $\mathbb{C}$. Thus we have $\sum_{i=1}^{n} c_i \exp(a_i z) = 0$ for all $z \in \mathbb{C}$. Linear independence of characters now implies that $c_i = 0$ for $i \in \{1, \ldots, n\}$ because the functions $\mathbb{R} \to \mathbb{R}$, $t \mapsto \exp(a_i t)$, for $i \in \{1, \ldots, n\}$ are pairwise different characters. ∎

This enables us to construct positive definite functions $\mathbb{R} \to \mathbb{R}$.

Example 2.16. Let $\Lambda = \mathbb{R}$ and μ a measure on $\mathbb{R}$ with $\int \exp(ax)\, d\mu(x) < \infty$ for all $a \in \mathbb{R}$. Then the functional

$$T\colon \mathbb{R}[\Lambda] \to \mathbb{R}, \quad g \mapsto \int \Psi(g)(x)\, d\mu(x),$$

clearly satisfies $T(g^2) \ge 0$ for all $g \in \mathbb{R}[\Lambda]$. If we choose for instance μ to be the push-forward to $\mathbb{R}$ of the Lebesgue measure on some compact interval $[a, b]$ with $a < b$, then we even have $T(g^2) > 0$ for every $g \in \mathbb{R}[\Lambda] \setminus \{0\}$. Indeed, for any $g \in \mathbb{R}[\Lambda]$ with $T(g^2) = 0$ the function $\Psi(g)$ must vanish on $[a, b]$, which implies $g = 0$ by Lemma 2.15. Let $f\colon \mathbb{R} \to \mathbb{R}$ be defined by $f(a) = T(\varepsilon^a)$ for $a \in \mathbb{R}$. Since T is $\mathbb{R}$-linear, we have $T = T_f$ and thus f is positive definite by Remark 2.14.

Corollary 2.17. *Let Λ be a divisible ordered abelian group whose dimension as a $\mathbb{Q}$-vector space is at most $2^{\aleph_0}$. Further, let R be a real closed field extension of $\mathbb{R}$. Then there exists an R-linear map $T\colon R[\Lambda] \to R$ which satisfies $T(g^2) > 0$ for every $g \in R[\Lambda] \setminus \{0\}$.*

Proof. By Example 2.16 there exists a positive definite map $f : \mathbb{R} \to \mathbb{R}$. By our assumptions on Λ there is an injective group homomorphism $\iota_1 : \Lambda \to \mathbb{R}$. Let $\iota_2 : \mathbb{R} \to R$ be the inclusion map. Then the map $\bar{f} = \iota_2 \circ f \circ \iota_1 : \Lambda \to R$ is also positive definite. By Remark 2.14 the map $T_{\bar{f}} : R[\Lambda] \to R$ has the desired properties. ∎

3. The group ring $R[\mathbb{Q}^n]$

From now on let R be a real closed field (not necessarily an extension of $\mathbb{R}$). We fix a total order on $\mathbb{Q}^n$ that is translation-invariant, for instance the lexicographical ordering. As in Remark 2.14, we consider the group ring $R[\mathbb{Q}^n]$ as the subring of $R[[\varepsilon^{\mathbb{Q}^n}]]$ of series with finite support. For $a = (a_1, \dots, a_n) \in \mathbb{Q}^n$ we write

$$\varepsilon^a = \varepsilon_1^{a_1} \cdots \varepsilon_n^{a_n}.$$

By plugging in positive elements from R for ε_i, we can regard elements from $R[\mathbb{Q}^n]$ as functions on the positive orthant $R_{>0}^n$. Clearly, every element of $R[\mathbb{Q}^n]$ that can be written as a sum of squares of elements of $R[\mathbb{Q}^n]$ is nonnegative on $R_{>0}^n$. We will see in Example 3.4 that the converse is not true. We will consider the polynomial ring $R[\underline{x}] = R[x_1, \dots, x_n]$ as a subring of $R[\mathbb{Q}^n]$ by replacing x_i by ε_i.

Remark 3.1. Evaluating elements $f \in R[\mathbb{Q}^n]$ at negative numbers is in general not well defined since the exponents can have even denominator. Similarly, a polynomial can be a sum of squares of elements from the group ring $R[\mathbb{Q}^n]$ without being nonnegative on all of R^n. For example, $\varepsilon \in R[\varepsilon]$ takes negative values on R, but equals $(\varepsilon^{1/2})^2$.

Remark 3.2. A polynomial $p \in R[\underline{x}]$ is a sum of squares of elements in $R[\mathbb{Q}^n]$ if and only if $p(x_1^m, \dots, x_n^m)$ is a sum of squares of polynomials from $R[\underline{x}]$ for some $m \in \mathbb{Z}_{>0}$. Indeed, negative exponents cannot appear in the sum of squares decomposition since lowest degree terms cannot cancel out.

Example 3.3. Consider the *Motzkin polynomial*

$$M = x_3^6 - 3x_1^2 x_2^2 x_3^2 + x_1^2 x_2^4 + x_1^4 x_2^2 \in R[x_1, x_2, x_3].$$

It follows from the inequality of arithmetic and geometric means that M is nonnegative on R^3 (since $\frac{1}{3}x_1^4 x_2^2 + \frac{1}{3}x_1^2 x_2^4 + \frac{1}{3}x_3^6 \geq x_1^{4/3} x_2^{2/3} x_1^{2/3} x_2^{4/3} x_3^{6/3}$) but it is not a sum of squares of polynomials from $R[\underline{x}]$ [18]. However, $M(x_1^2, x_2^2, x_3^2)$ is a sum of squares of polynomials [28, Example 5.3]:

$$M(x_1^2, x_2^2, x_3^2) = 2 \cdot (x_1^3 x_2^3 - x_1 x_2 x_3^4)^2 + (x_1^4 x_2^2 - x_1^2 x_2^4)^2 + (x_3^6 - x_1^2 x_2^2 x_3^2)^2.$$

This shows that M is a sum of squares of elements of $R[\mathbb{Q}^3]$.

Example 3.4. Consider the *Horn polynomial*

$$h = (x_1 + x_2 + x_3 + x_4 + x_5)^2 - 4(x_1 x_2 + x_2 x_3 + x_3 x_4 + x_4 x_5 + x_5 x_1).$$

It was proven in [10, p. 335] that h is nonnegative on the positive orthant. However, for any integer $k > 0$ the polynomial $h(x_1^k, \ldots, x_5^k)$ is not a sum of squares of elements of $R[\underline{x}]$. This was shown in [26, Section 4] by adapting the proof from [10] for the case $k = 2$. This shows that h is not a sum of squares of elements of $R[\mathbb{Q}^5]$.

For the sake of completeness, we repeat the argument here. Assume h is a sum of squares in $R[\mathbb{Q}^5]$. Then for some integer $d > 0$ the polynomial

$$H_d = h(x_1^d, \ldots, x_5^d) = g_1^2 + \cdots + g_m^2$$

is a sum of squares of homogeneous polynomials g_j of degree d (since highest and lowest degree terms cannot cancel out). Restricting H_d to $x_i = x_{i+1} = 0$ we obtain

$$(x_{i+2}^d - x_{i+3}^d + x_{i+4}^d)^2$$

for $i \in \{1, \ldots, 5\}$, where the indices are taken modulo 5. Thus, for $j \in \{1, \ldots, m\}$ the restriction of g_j to $x_i = x_{i+1} = 0$ must be a scalar multiple of $x_{i+2}^d - x_{i+3}^d + x_{i+4}^d$. Indeed, if

$$(x_{i+2}^d - x_{i+3}^d + x_{i+4}^d)^2 = g_1'^2 + \cdots + g_m'^2$$

for $g_j' \in \mathbb{R}[x_{i+2}, x_{i+3}, x_{i+4}]$, then each g_j' must vanish on the real zero set of $x_{i+2}^d - x_{i+3}^d + x_{i+4}^d$. The latter is Zariski dense in the complex zero set, which implies by Hilbert's Nullstellensatz that g_j' is divisible by the irreducible polynomial $x_{i+2}^d - x_{i+3}^d + x_{i+4}^d$. Letting c_i be the coefficient of x_i^d in g_j, this shows that $c_i = -c_{i+1}$. Since 5 is an odd number, this actually implies that $c_1 = \cdots = c_5 = 0$. In particular, $g_j(1, 0, \ldots, 0) = 0$ for $j \in \{1, \ldots, m\}$. However, since $H_d(1, 0, \ldots, 0) \neq 0$, this is a contradiction.

Example 3.5. The Horn polynomial considered in Example 3.4 is a special case of the following more general construction. Recall that a *stable set* of a graph $G = (V, E)$ is a subset of elements of V that are pairwise nonadjacent to each other. The *stability number* $\alpha(G)$ of G is the largest size of a stable set of G. By a theorem of Motzkin–Straus [19] the quadratic form associated to the symmetric matrix

$$Q(G) = \alpha(G)(I + A) - J$$

is nonnegative on the positive orthant [23, Example 3.163], i.e., the matrix $Q(G)$ is *copositive*. Here, A is the adjacency matrix of G and J is the all-ones matrix. The Horn polynomial corresponds to the matrix $Q(C_5)$ where C_5 is the five-cycle. The argument from Example 3.4 can in fact be applied to any odd cycle of size at least 5.

We also note that nonnegative elements in $\mathbb{R}[\mathbb{Q}^n]$ cannot be distinguished from sums of squares by an $\mathbb{R}$-linear functional $\mathbb{R}[\mathbb{Q}^n] \to \mathbb{R}$.

Proposition 3.6. *Let* $T \colon \mathbb{R}[\mathbb{Q}^n] \to \mathbb{R}$ *be an $\mathbb{R}$-linear functional with $T(a^2) \geq 0$ for all* $a \in \mathbb{R}[\mathbb{Q}^n]$ *and let $p \in \mathbb{R}[\mathbb{Q}^n]$ be nonnegative on $R_{>0}^n$. Then $T(p) \geq 0$.*

Proof. In the case $n = 1$ the statement of [3, Proposition 5.10, p. 211] says that for every such T there is a measure μ on $[0, \infty)$ such that $T(g) = \int g \, d\mu$. This shows in particular that $T(p) \geq 0$. The case $n > 1$ then follows from [3, Theorem 5.4, p. 204]. ∎

However, if we replace the target field by some real closed extension field, then separation by a linear functional is always possible. The following theorem follows from the results of [21]. We will give self-contained proofs in Appendix A for the results of [21] that we use in the proof of Theorem 3.7.

Theorem 3.7. *Let $p \in R[\mathbb{Q}^n]$ not be a sum of squares of elements of $R[\mathbb{Q}^n]$. Then there exists a real closed field extension $R' \supset R$ and an R'-linear functional $T \colon R'[\mathbb{Q}^n] \to R'$ such that for every $a \in R'[\mathbb{Q}^n] \setminus \{0\}$ we have $T(a^2) > 0$ and $T(p) < 0$.*

Proof. By Corollary A.3 there exists a real closed field extension $R' \supset R$ and an R'-linear functional $T' \colon R'[\mathbb{Q}^n] \to R'$ such that for every $a \in R'[\mathbb{Q}^n]$ we have $T'(a^2) \geq 0$ and $T'(p) < 0$. Further, by Corollary 2.17 there exists an R'-linear map $T_0 \colon R'[\mathbb{Q}^n] \to R'$ that satisfies $T_0(a^2) > 0$ for all $0 \neq a \in R'[\mathbb{Q}^n]$. Let $\lambda \in R'$ be such that $\lambda > 0$ and $\lambda |T_0(p)| < |T'(p)|$. Then $T = \lambda T_0 + T'$ has the desired properties. ∎

Now we are ready to formulate and prove our main result.

Definition 3.8. Let $S \subset \mathbb{Z}_{\geq 0}^n$ be a finite set and R a real closed field. We denote by $\mathrm{P}(S, R)$ the (finite-dimensional) R-vector space of all polynomials in n variables and coefficients in R whose support is contained in S. We further denote by $\mathrm{P}_+(S, R)$ the convex cone of all $p \in \mathrm{P}(S, R)$ such that $p(x) \geq 0$ for all $x \in R^n$. Finally, we denote by $\Sigma(S, R)$ the convex cone of all $P \in \mathrm{P}(S, R)$ that are sums of squares of polynomials (with coefficients in R). If we do not specify the real closed field, then we work over the real numbers, i.e., we write $\mathrm{P}(S) = \mathrm{P}(S, \mathbb{R})$, $\mathrm{P}_+(S) = \mathrm{P}_+(S, \mathbb{R})$, and $\Sigma(S) = \Sigma(S, \mathbb{R})$.

Remark 3.9. Note that $\mathrm{P}_+(S, R)$ is the set of all tuples $(c_s)_{s \in S} \subset R^S$ such that

$$\forall x_1 \cdots \forall x_n \colon \quad \sum_{s \in S} c_s \cdot \prod_{i=1}^n x_i^{s_i} \geq 0$$

holds true over R. This is a formula in the language of real closed fields without parameters. Such a formula also exists for $\Sigma(S, R)$. Indeed, let $T_1 = \mathbb{Z}_{\geq 0}^n \cap \mathrm{conv}(S)$ and $T_2 = \mathbb{Z}_{\geq 0}^n \cap (\frac{1}{2} \mathrm{conv}(S))$. Every element of $\Sigma(S, R)$ is a sum of squares of elements of $\mathrm{P}(T_2, R)$, i.e., lies in the conical hull of the set

$$\{q^2 \mid q \in \mathrm{P}(T_2, R)\} \subset \mathrm{P}(T_1, R).$$

Then by Carathéodory's theorem for the conical hull every such element can be written as a sum of $|T_1| = \dim(\mathrm{P}(T_1, R))$ squares. Therefore, $\Sigma(S, R)$ is the set of all $p \in \mathrm{P}(S, R)$ for which there exist $c_{i,t} \in R$, for $i \in \{1, \ldots, |T_1|\}$ and $t \in T_2$, such that

$$p = \sum_{i=1}^{|T_1|} \left(\sum_{t \in T_2} c_{i,t} \prod_{j=1}^n x_j^{t_j} \right)^2.$$

This shows that $\Sigma(S, R)$ is defined by a formula in the language of real closed fields without parameters.

Remark 3.10. For all finite sets S the set $\Sigma(S)$ is a spectrahedral shadow [22, Section 6.3.1, p. 263].

Proposition 3.11. *Let $S \subset \mathbb{Z}_{\geq 0}^n$ be a finite set. Assume that there exists a real closed field R and $p \in P_+(S, R)$ such that p is not a sum of squares of elements from $R[\mathbb{Q}^n]$. Then $P_+(S)$ is not a spectrahedral shadow.*

Proof. Without loss of generality we may assume that R contains $\mathbb{R}$. By Theorem 3.7 there exists a real closed field extension $R' \supset R$ and an R'-linear functional $T : R'[\mathbb{Q}^n] \to R'$ such that $T(a^2) > 0$ for all $a \in R'[\mathbb{Q}^n] \setminus \{0\}$ and $T(p) < 0$. Let $f : \mathbb{Q}^n \to R'$ be defined by $f(w) = T(\varepsilon_1^{w_1} \cdots \varepsilon_n^{w_n}) \in R'$ for $w \in \mathbb{Q}^n$. By the R'-linearity of T we then have $T = T_f$ as in Remark 2.14. As $T(a^2) > 0$ we know that f is positive definite (see Remark 2.14). After replacing T by $\frac{1}{T(1)} \cdot T$, we can further assume that $f(0) = 1$. Therefore, $L_f : R'[[\varepsilon^{\mathbb{Q}^n}]] \to R'[[\varepsilon^{\mathbb{Q}^n}]]$ is completely positive by Theorem 2.13, $\mathbb{R}$-linear by construction, and unital because $f(0) = 1$. The polynomial $q = p(\varepsilon_1 x_1, \ldots, \varepsilon_n x_n)$ with coefficients in $R[[\varepsilon^{\mathbb{Q}^n}]]$ is nonnegative on $R'[[\varepsilon^{\mathbb{Q}^n}]]^n$. By construction we have

$$L_f(q)(\varepsilon_1^{-1}, \ldots, \varepsilon_n^{-1}) = T_f(p) < 0.$$

Thus, L_f preserves all formulas in Γ by Lemma 2.3 but not the formula defining the cone $P_+(S, R'[[\varepsilon^{\mathbb{Q}^n}]])$. This shows that $P_+(S)$ is not a spectrahedral shadow. ∎

A first-order formula is called *positive* if it is built from atomic formulas by existential and universal quantification, conjunction, and disjunction. That is, only negation from first-order logic is forbidden.

Proposition 3.12. *Let $S \subset \mathbb{Z}_{\geq 0}^n$ and $p \in P_+(S, R)$ be as in Proposition 3.11. Then $P_+(S)$ does not have a positive definition over the relational structure whose relations consist of all spectrahedra and the inequality relation $\neq$.*

Proof. Note that by Lemma 2.10 the map $L_f : R'[[\varepsilon^{\mathbb{Q}^n}]] \to R'[[\varepsilon^{\mathbb{Q}^n}]]$ constructed in the proof of Proposition 3.11 is surjective. It follows from the easy direction of the Lyndon preservation theorem for positive first-order logic (see, e.g., [13, Corollary 10.3.5]) and the Tarski principle for real closed fields (as outlined for existential positive logic in the introduction) that the convex semialgebraic set $P_+(S)$ does not have a definition by a positive formula over spectrahedra. In fact, the map L_f is also injective and hence preserves $\neq$, and so $P_+(S)$ is inexpressible even if we additionally allow the use of $\neq$. ∎

Remark 3.13. If full first-order logic is permitted, then the graph of multiplication can already be defined from fairly restricted sets of relations over the reals [17, 24] (and in combination with the graph of addition we then obtain all semialgebraic relations). In particular, it is easy to see that all semialgebraic relations can be defined with first-order formulas and parameters from the graph of addition and the spectrahedron defined by $y \geq x^2$.

Let R be a real closed field, $d > 0$ an integer, $S \subset \mathbb{Z}_{\geq 0}^n$ a finite set and $s \in \{\pm 1\}^n$. We consider the following isomorphisms of R-vector spaces:

$$\psi_d \colon P(S, R) \to P(d \cdot S, R), \quad p(x_1, \ldots, x_n) \mapsto p(x_1^d, \ldots, x_n^d),$$
$$\varphi_s \colon P(S, R) \to P(S, R), \quad p(x_1, \ldots, x_n) \mapsto p(s_1 x_1, \ldots, s_n x_n).$$

Further, we define

$$\Sigma_d(S, R) = \bigcap_{s \in \{\pm 1\}^n} \varphi_s(\psi_d^{-1}(\Sigma(d \cdot S, R))).$$

Again we write $\Sigma_d(S) = \Sigma_d(S, \mathbb{R})$.

Lemma 3.14. *The following hold true for all integers $d > 0$:*

(1) $\Sigma_d(S, R)$ *can be defined by a formula in the language of real closed fields without parameters.*

(2) $\Sigma_d(S, R) \subseteq \Sigma_{i \cdot d}(S, R)$ *for all integers $i > 0$.*

(3) $\Sigma_d(S, R) \subseteq P_+(S, R)$.

(4) $\Sigma_d(S)$ *is a spectrahedral shadow.*

Proof. Part (1) follows from Remark 3.9 and the fact that ψ_d and φ_s are first-order definable without parameters in the language of real closed fields. If a polynomial p is a sum of squares, then $p(x_1^i, \ldots, x_n^i)$ is also a sum of squares. This shows (2). Let $p \in \Sigma_d(S, R)$ and $x_1, \ldots, x_n \in R$. To show (3), we have to show that $p(x_1, \ldots, x_n) \geq 0$. For $i \in \{1, \ldots, n\}$, let $s_i \in \{\pm 1\}$ be such that $s_i x_i \geq 0$. Then we let $y_i = \sqrt[d]{s_i x_i}$. We have $q = \psi_d(\varphi_s^{-1}(p)) \in \Sigma(d \cdot S, R) \subseteq P_+(d \cdot S, R)$. Therefore,

$$0 \leq q(y_1, \ldots, y_n) = (\varphi_s^{-1}(p))(s_1 x_1, \ldots, s_n x_n) = p(x_1, \ldots, x_n).$$

This shows (3). Finally, $\Sigma_d(S)$ is a spectrahedral shadow as the intersection of finitely many spectrahedral shadows. $\blacksquare$

Theorem 3.15. *Let $S \subset \mathbb{Z}_{\geq 0}^n$ be a finite set. The following are equivalent:*

(1) $P_+(S)$ *is a spectrahedral shadow.*

(2) *For every real closed field R and every $p \in P_+(S, R)$, there exists $d > 0$ such that $p(x_1^d, \ldots, x_n^d)$ is a sum of squares.*

(3) *There exists $d > 0$ such that for every real closed field R and every polynomial $p \in P_+(S, R)$ the polynomial $p(x_1^d, \ldots, x_n^d)$ is a sum of squares.*

(4) *There exists $d > 0$ such that for every polynomial $p \in P_+(S)$ the $p(x_1^d, \ldots, x_n^d)$ is a sum of squares.*

(5) $P_+(S)$ *has a positive definition over the relational structure whose relations consist of all spectrahedra and the relation $\neq$.*

Proof. The implication (1)$\Rightarrow$(2) follows from Proposition 3.11 and Remark 3.2. Now assume (2). We first claim that

$$P_+(S, R) = \bigcup_{k=1}^{\infty} \Sigma_k(S, R).$$

The inclusion $\supseteq$ is Lemma 3.14 (3). Let $p \in P_+(S, R)$. By assumption, for all $s \in \{\pm 1\}^n$ there is $k_s > 0$ such that $\psi_{k_s}(\varphi_s(p)) \in \Sigma(k_s \cdot S, R)$. Therefore, letting $k = \prod_{s \in \{\pm 1\}^n} k_s$ we have $\psi_k(\varphi_s(p)) \in \Sigma(k \cdot S, R)$ and thus $p \in \Sigma_k(S, R)$.

Both $P_+(S, R)$ and each $\Sigma_k(S, R)$ can be defined by a formula without parameters. Thus, letting $\mathbb{R}^*$ be an $\aleph_1$-saturation of $\mathbb{R}$, we have $P_+(S, \mathbb{R}^*) = \Sigma_d(S, \mathbb{R}^*)$ for some $d > 0$ by [27, Theorem 2.2.11]. By the Tarski principle we thus have $P_+(S, R) = \Sigma_d(S, R)$ for every real closed field R. This shows (3).

The direction (3)$\Rightarrow$(4) is trivial. Finally, assume (4) which implies $P_+(S) = \Sigma_d(S)$. This shows (1) by Lemma 3.14 (4) and we see that (1)–(4) are equivalent.

It is clear that (1) implies (5), and Proposition 3.12 shows that (5) implies (2). ∎

Remark 3.16. Theorem 3.15 also gives a characterization of when the set of all $p \in P(S)$ which are nonnegative on the nonnegative orthant is a spectrahedral shadow. Indeed, a polynomial $p \in P(S)$ is nonnegative on the nonnegative orthant if and only if $p(x_1^2, \ldots, x_n^2) \in P_+(2S)$.

Remark 3.17. A closed convex cone $K \subseteq \mathbb{R}^n$ is a spectrahedral shadow if and only if its dual cone is a spectrahedral shadow [22, Exercise 6.23]. The dual cone of $P_+(S)$ is the closed conical hull of the image of $\mathbb{R}^n$ under the monomial map given by the monomials from S.

The following answers [31, Question 5.2]. For $s \in \mathbb{Z}_{\geq 0}^n$, we write $|s|$ for $s_1 + \cdots + s_n$.

Corollary 3.18. *The set of copositive matrices of size $n \geq 5$ is not a spectrahedral shadow. Moreover, it does not have a positive definition in the relational structure over the reals whose relations consist of all spectrahedra and the inequality relation.*

Proof. By Remark 3.16 the set of copositive matrices of size n can be identified with the set of all polynomials $q \in \mathbb{R}[x_1, \ldots, x_n]$ that are homogeneous of degree two such that $q(x_1^2, \ldots, x_n^2)$ is nonnegative on $\mathbb{R}^n$. Thus, the set of copositive matrices of size n is identified with $P_+(2S)$ where $S = \{s \in \mathbb{Z}_{\geq 0}^n \mid |s| = 2\}$. We have seen in Example 3.4 that the Horn polynomial h satisfies $p = h(x_1^2, \ldots, x_n^2) \in P_+(2S)$ but $p(x_1^d, \ldots, x_n^d)$ is not a sum of squares for every $d > 0$. Thus Theorem 3.15 implies the claim. ∎

For the rest of this section we summarize the current knowledge on whether $P_+(S)$ is a spectrahedral shadow for a given finite set $S \subset \mathbb{Z}_{\geq 0}^n$.

Example 3.19 (Hilbert). Let $S \subset \mathbb{Z}_{\geq 0}^n$ a finite set and $\Delta_{\leq d}$ the set of all $s \in \mathbb{Z}_{\geq 0}^n$ such that $|s| \leq d$. Hilbert [12] proved that every element of $P_+(S)$ is a sum of squares if we are in one of the following cases:

(1) $S \subseteq \Delta_{\leq 2}$;

(2) $n = 1$;

(3) $n = 2$ and $S \subseteq \Delta_{\leq 4}$.

Thus, in all these cases $P_+(S)$ is a spectrahedral shadow.

Example 3.20 (Fawzi, Scheiderer). Let $S \subset \mathbb{Z}_{\geq 0}^n$ be a finite set which is *downward closed* in the sense that if $s' \leq s$ (coordinatewise) and $s \in S$, then $s' \in S$. Then by [8, Theorem 3] the set $P_+(S)$ is a spectrahedral shadow if and only if $P_+(S) = \Sigma(S)$. Thus, for every nonnegative polynomial which is not a sum of squares, letting S be the downward closure of its support, we obtain an instance where $P_+(S)$ is not a spectrahedral shadow. Scheiderer [31] used a related criterion in several examples. A generalization of [8, Theorem 3] to finite-dimensional spaces of polynomials not necessarily spanned by monomials is given in [4, Theorem 2.14].

Example 3.21 (SONC). A finite subset $A \subset \mathbb{Z}_{\geq 0}^n$ is called a *circuit set* if its convex hull $\mathrm{conv}(A) \subset \mathbb{R}^n$ is a k-dimensional simplex whose vertices $v_0, \dots, v_k$ are in $2\mathbb{Z}_{\geq 0}^n$ and $A \setminus \{v_0, \dots, v_k\} \subseteq \{w\}$ for some w in the relative interior of $\mathrm{conv}(A)$. For any finite $S \subset \mathbb{Z}_{\geq 0}^n$ the cone of *sums of nonnegative circuit polynomials* is defined as

$$\mathrm{SONC}(S) = \sum_{A \subseteq S \text{ a circuit set}} P_+(A).$$

Averkov [2, Section 6] showed that $\mathrm{SONC}(S)$ has a semidefinite extended representation with LMIs of size at most 2. Without the bound on the LMI sizes this also follows from a result by Reznick [26, 28] which implies that for every $p \in \mathrm{SONC}(S)$ the polynomial $p(x_1^d, \dots, x_n^d)$ is a sum of squares for $d = \max(2, n - 1)$. In particular, $P_+(S)$ is a spectrahedral shadow whenever $P_+(S) = \mathrm{SONC}(S)$. By definition this is the case if S is a circuit set. Equality is also straightforward if every element of S is a vertex of $\mathrm{conv}(S)$. More cases of equality can be deduced from [20, Theorem 11] and [9, Theorem 8.1].

The proof of Corollary 3.18 shows the existence of a set $S \subset \mathbb{Z}_{\geq 0}^n$ of size $\frac{1}{2}n(n + 1)$, for $n \geq 5$, such that $P_+(S)$ is not a spectrahedral shadow. For $n = 5$ this means $|S| = 15$. Moreover, Scheiderer [31, Remark 4.21] constructs such an $S \subset \mathbb{Z}_{\geq 0}^2$ with $|S| = 14$. We do not know the smallest size of a set $S \subset \mathbb{Z}_{\geq 0}^n$ such that $P_+(S)$ is not a spectrahedral shadow. The following shows that any such set must satisfy $|S| > 5$.

Corollary 3.22. *Let $S \subset \mathbb{Z}_{\geq 0}^n$ with $|S| \leq 5$. Then $P_+(S)$ is a spectrahedral shadow.*

Proof. Consider the convex hull $K = \mathrm{conv}(S) \subset \mathbb{R}^n$ and let N be the number of vertices of K. First suppose that $N \leq 2$. Then there is a finite subset $T \subset \mathbb{Z}_{\geq 0}$ and $v \in \mathbb{Z}_{\geq 0}^n$, $w \in \mathbb{Z}^n \setminus \{0\}$ such that $S = \{v + k \cdot w \mid k \in T\}$ and the entries of w do not have a nontrivial common divisor. If $P_+(S) \neq \{0\}$ we can further assume that $v \in 2\mathbb{Z}_{\geq 0}^n$. Then the vector space isomorphism

$$P(T) \to P(S), \quad p \mapsto \underline{x}^v \cdot p(\underline{x}^w),$$

identifies $P_+(T)$ with $P_+(S)$. Since every univariate nonnegative polynomial is a sum of squares, the set $P_+(T)$, and hence $P_+(S)$, is a spectrahedral shadow.

In the case $N \geq 3$ we claim that $P_+(S) = \text{SONC}(S)$. This implies the claim since $\text{SONC}(S)$ is a spectrahedral shadow by [2, Section 6]; see also Example 3.21. First assume that K is a simplex. Since $|S| \leq 5$, there are at most two elements of S that are not vertices of K. Hence $P_+(S) = \text{SONC}(S)$ follows from [20, Theorem 11 (1)] (applied to the trivial partition of S). Now assume that K is not a simplex. We first note that $N \geq 3$ implies that $\dim(K) \geq 2$. Therefore, since K is not a simplex, we even have $N \geq 4$. Thus there is at most one element of S that is not a vertex of K and we can apply [20, Theorem 11 (2)] (again to the trivial partition). ∎

4. Separating Γ_k from Γ_{k+1}

The following criteria for a function $f \colon \mathbb{R} \to \mathbb{R}$ to be k-positive definite or semidefinite are due to Karlin [15].

Lemma 4.1. *For $k \in \mathbb{Z}_{>0}$ let $f \in C^{2k-2}(\mathbb{R})$ and for all $x \in \mathbb{R}$ consider the Hankel matrix $H_f(x) = (f^{(i+j-2)}(x))_{1 \leq i,j \leq k}$ of order k.*

(1) *If f is k-positive semidefinite, then $H_f(x)$ is positive semidefinite for every $x \in \mathbb{R}$.*

(2) *If $H_f(x)$ is positive definite for every $x \in \mathbb{R}$, then f is k-positive definite.*

Proof. Part (1) follows from applying suitable transformations to the matrix

$$(f(x_i + x_j))_{i,j}$$

for $x_1 < \cdots < x_k$ and then letting all x_i converge to the same value x, as is done in [15, Section 2.1]. Part (2) follows from [15, Theorem 2.6] applied to the kernel $K(x, y) = f(x + y)$. ∎

This criterion allows us to exhibit a function $\mathbb{R} \to \mathbb{R}$ that is k-positive definite, but not $k + 1$-positive semidefinite. First we make some observations on the Hankel matrices $H_f(x)$.

Remark 4.2. By definition the Hankel matrix $H_f(x)$ is linear in f.

For the rest of the section let $0 < a_1 < \cdots < a_k$ and $g(x) = \sum_{m=1}^{k} \exp(a_m x)$. Let $\bar{g}(x) = g(-x)$. For any $\varepsilon > 0$ consider the function $f_\varepsilon \colon \mathbb{R} \to \mathbb{R}$ given by $f_\varepsilon(x) = g(x) + \bar{g}(x) - \varepsilon$.

Lemma 4.3. *For any $\varepsilon > 0$ the function f_ε is not $(k + 1)$-positive semidefinite.*

Proof. Let $p_m(x) = \exp(a_m x)$ for $m \in \{1, \ldots, k\}$. The Hankel matrix $H_{p_m}(x)$ of order $k + 1$ equals the rank 1 matrix

$$\exp(a_m x) \cdot v_m^\top \cdot v_m$$

where $v_m = (1, a_m, \ldots, a_m^k)$ for $m \in \{1, \ldots, k\}$. Let $w \in \mathbb{R}^{k+1}$ be a nonzero vector in

the orthogonal complement of the subspace spanned by $v_1, \ldots, v_k$. We claim that $w^\top$ is in the kernel of $H_g(x)$ for all $x \in \mathbb{R}$. Indeed, for all $x \in \mathbb{R}$ we have

$$H_g(x) \cdot w^\top = \sum_{m=1}^{k} H_{p_m}(x) \cdot w^\top = \sum_{m=1}^{k} \exp(a_m x) \cdot v_m^\top \cdot \underbrace{v_m \cdot w^\top}_{=0} = 0.$$

Further, we claim that the first component of w is nonzero. Indeed, assume otherwise and let $w' \in \mathbb{R}^k$ be such that $w = (0, w')$. Then w' is in the orthogonal complement of the vectors

$$(a_1, \ldots, a_1^k), \ldots, (a_k, \ldots, a_k^k).$$

However, these vectors are linearly independent because $a_1, \ldots, a_k$ are pairwise different and nonzero. This contradicts our assumption that w and thus w' is nonzero. Thus, we have shown that the first component of w is nonzero. Without loss of generality, we may assume that it equals 1. For $m \in \{1, \ldots, k\}$ let $b_m = -a_m$ and $u_m = (1, b_m, \ldots, b_m^k)$. Then, as above, we compute

$$wH_{\tilde{g}}(x)w^\top = \sum_{m=1}^{k} \exp(-a_m x) \cdot \langle u_m, w \rangle^2,$$

which tends to zero as $x \to \infty$ because $a_1, \ldots, a_k$ are positive. Let $x_0 \in \mathbb{R}$ be such that $wH_{\tilde{g}}(x)w^\top < \varepsilon$ for all $x > x_0$. Then for all $x > x_0$ we have

$$wH_{f_\varepsilon}w^\top = wH_{\tilde{g}}w^\top - \varepsilon < 0.$$

By Lemma 4.1 this proves the claim. ∎

Lemma 4.4. *For sufficiently small $\varepsilon > 0$ the function f_ε is k-positive definite.*

Proof. Let $H_g(x)$ and $H_{\tilde{g}}(x)$ be Hankel matrices of order k. They are positive linear combinations of positive semidefinite rank 1 matrices so they are positive semidefinite for all $x \in \mathbb{R}$. We can write $H_g(x) = A + B(x)$ where $A = \sum_{m=1}^{k} v_m^\top \cdot v_m$, $B(x) = \sum_{m=1}^{k} (\exp(a_m x) - 1) \cdot v_m^\top \cdot v_m$ and $v_m = (1, a_m, \ldots, a_m^{k-1})$. Similarly, we can write $H_{\tilde{g}}(x) = A' + B'(-x)$ where $A' = \sum_{m=1}^{k} u_m^\top \cdot u_m$, $B'(x) = \sum_{m=1}^{k} (\exp(-a_m x) - 1) \cdot u_m^\top \cdot u_m$ and $u_m = (1, -a_m, \ldots, (-a_m)^{k-1})$. We note that if $x \geq 0$ then $B(x)$ is positive semidefinite because $\exp(a_m x) \geq 1$ for $x \geq 0$, and similarly $B'(-x)$ is positive semidefinite when $x \leq 0$. In any case, A and A' are positive definite since the vectors $v_1, \ldots, v_k$ resp. $u_1, \ldots, u_k$ are linearly independent. Thus, it suffices to choose $\varepsilon > 0$ in such a way that both $A - \varepsilon \cdot \delta_1 \cdot \delta_1^\top$ and $A' - \varepsilon \cdot \delta_1 \cdot \delta_1^\top$ are positive definite, where δ_1 is the first unit vector, since $\varepsilon \cdot \delta_1 \cdot \delta_1^\top$ is the Hankel matrix of the constant function ε. Indeed, for $x \geq 0$ the matrix

$$H_{f_\varepsilon}(x) = (A - \varepsilon \cdot \delta_1 \cdot \delta_1^\top) + B(x) + H_{\tilde{g}}(x)$$

is positive semidefinite. Similarly, for $x \leq 0$, the matrix

$$H_{f_\varepsilon}(x) = (A' - \varepsilon \cdot \delta_1 \cdot \delta_1^\top) + B'(-x) + H_g(x)$$

is positive semidefinite. ∎

Corollary 4.5. *For every $k \in \mathbb{Z}_{\geq 0}$ there is a function $\mathbb{R} \to \mathbb{R}$ that is k-positive definite but not $k + 1$-positive semidefinite.*

Proof. For sufficiently small $\varepsilon > 0$ the function f_ε has the desired properties. ∎

Example 4.6. Let $k = 2$ and $0 < a_1 < a_2$. Then

$$f_\varepsilon(x) = \exp(a_1 x) + \exp(a_2 x) + \exp(-a_1 x) + \exp(-a_2 x) - \varepsilon.$$

For $x \geq 0$ resp. $x \leq 0$ the Hankel matrix $H_f(x)$ is of the form

$$\underbrace{\begin{pmatrix} 2 & \pm(a_1 + a_2) \\ \pm(a_1 + a_2) & a_1^2 + a_2^2 \end{pmatrix}}_{=A \text{ resp. } A' \text{ in the proof of } 4.4} + \begin{pmatrix} -\varepsilon & 0 \\ 0 & 0 \end{pmatrix} + C(x),$$

where $C(x)$ is positive semidefinite. For instance, letting $a_1 = \ln(2)$, $a_2 = \ln(3)$, and $\varepsilon = 1/11$ makes $H_{f_{1/11}}(x)$ positive definite for all $x \in \mathbb{R}$. Indeed, in this case we have

$$\det \begin{pmatrix} 2 - \varepsilon & \pm(a_1 + a_2) \\ \pm(a_1 + a_2) & a_1^2 + a_2^2 \end{pmatrix} = \frac{21}{11} \cdot (\ln(2)^2 + \ln(3)^2) - (\ln(2) + \ln(3))^2 \approx 0.011 > 0.$$

Thus, $f_{1/11}$ is 2-positive definite. However, the matrix $(f_{1/11}(x_i + x_j))_{1 \leq i, j \leq 3}$ for $x_1 = 5, x_2 = 6, x_3 = 7$ equals

$$\begin{pmatrix} \dfrac{39956170693955}{665127936} & \dfrac{715125242123209}{3990767616} & \dfrac{12823220129727323}{23944605696} \\[2ex] \dfrac{715125242123209}{3990767616} & \dfrac{12823220129727323}{23944605696} & \dfrac{230229525738486289}{143667634176} \\[2ex] \dfrac{12823220129727323}{23944605696} & \dfrac{230229525738486289}{143667634176} & \dfrac{4137070068201557555}{862005805056} \end{pmatrix}.$$

Its determinant is

$$-\frac{2277541160576348197}{107750725632} < 0$$

and thus the matrix is not positive semidefinite. Therefore, $f_{1/11}$ is not 3-positive semidefinite.

In order to apply Theorem 2.13 consider the real closed field $H = \mathbb{R}[[\varepsilon^{\mathbb{R}}]]$ of Hahn series over $\mathbb{R}$ in the (infinitesimal) indeterminate ε and with value group $\mathbb{R}$.

Corollary 4.7 ([2, Corollary 2.5]). *The cone K of $k \times k$ real symmetric positive semidefinite matrices has no semidefinite extended formulation with LMIs of size less than k.*

Proof. By Corollary 4.5 there exists a function $f \colon \mathbb{R} \to \mathbb{R}$ that is $(k-1)$-positive definite but not k-positive semidefinite. By Remark 2.12 we can assume without loss of generality that $f(0) = 1$. Then $L_f \colon H \to H$ does not preserve K by Theorem 2.13 (2). But L_f preserves all spectrahedra of size $k - 1$ by Lemma 2.3, and hence it preserves all sets that have a semidefinite extended formulation with LMIs of size at most $k - 1$. This shows that K does not have such a formulation. ∎

In the same way as in the proof of Proposition 3.12 we even obtain the following stronger statement.

Corollary 4.8. *The cone K of $k \times k$ real symmetric positive semidefinite matrices does not have a positive definition over the relational structure whose relations consist of all sets that can be described by an LMI of size less than k and the inequality relation $\neq$.*

Appendix A. Real closed separation

We follow the line of arguments from [21]. We start with a well-known lemma which we include for completeness.

Lemma A.1. *Let R be a real closed field and I some finite index set. Let V_1, V_2 be finite-dimensional R-vector spaces. Let $\alpha_i \colon V_1 \to V_2$, $i \in I$, be some homogeneous polynomial maps with the property that no nontrivial positive linear combination of images of the α_i is trivial, meaning that*

$$\sum_{j=1}^{r} \lambda_j \alpha_{i_j}(v_j) = 0$$

for $r \in \mathbb{N}$, $i_j \in I$ and $\lambda_j > 0$ implies that all $v_j \in V_{i_j}$ are zero. Then

$$C := \operatorname{conv}\left(\bigcup_{i \in I} \alpha_i(V_1) \right)$$

is a closed convex cone. Moreover, if $b \notin C$, then there exists an R-linear map $\phi \colon V_2 \to R$ such that $\phi(c) \geq 0$ for $c \in C$ and $\phi(b) < 0$.

Proof. After choosing bases we can assume that $V_1 = R^d$ and $V_2 = R^n$. Let

$$S = \{ v \in R^d \mid \|v\|_2 = 1 \}$$

be the unit sphere with respect to the Euclidean norm. Since each α_i is homogeneous, C is the conical hull of $\bigcup_{i \in I} \alpha_i(S)$. In particular, it is a convex cone. Thus it remains to prove that C is closed. Let $m = |I|$. We consider the map

$$\Phi \colon S^{mn} \times \Delta_{mn} \to R^n, \quad (v_{11}, \dots, v_{mn}, \lambda_{11}, \dots, \lambda_{mn}) \mapsto \sum_{j=1}^{n} \sum_{i=1}^{m} \lambda_{ij} \alpha_i(v_{ij}),$$

where $\Delta_{mn} \subset R^{mn}$ is the standard simplex. The image $B \subset R^n$ of Φ is semialgebraically compact as the image of a semialgebraically compact set under a continuous semialgebraic map. We claim that

$$\Sigma = \{ (v, b) \in R^n \times B \mid \exists \lambda \geq 0 : v = \lambda \cdot b \}$$

is a closed subset of $R^n \times B$. Indeed, it can be written as

$$\Sigma = \{ (v, b) \in R^n \times B \mid \forall i, j : (v_i b_j = v_j b_i) \wedge (v_i b_i \geq 0) \}$$

since $0 \notin B$ by assumption. Since B is semialgebraically compact, the projection of Σ onto the first coordinate, which equals C by Carathéodory's theorem for the conical hull, is closed as well. Now the conical separation theorem implies the additional statement. ∎

Let R be a real closed field and A a commutative R-algebra. For $g_1, \dots, g_r \in A$ we denote

$$\mathrm{U}_A(g_1, \dots, g_r) = \{\alpha \in \mathrm{Sper}(A) \mid g_1(\alpha) > 0, \dots, g_r(\alpha) > 0\},$$

where $\mathrm{Sper}(A)$ is the *real spectrum* of A (see for example [5, Section 7]). We further consider the *quadratic module* $\mathrm{QM}_A(g_1, \dots, g_r)$, i.e., the set of all

$$s_0 + s_1 g_1 + \dots + s_r g_r \in A,$$

where the $s_i \in A$ are sums of squares. We say that A is *real reduced* if $\sum_{i=1}^k a_i^2 = 0$ implies $a_1 = \dots = a_k = 0$ for all $k \in \mathbb{N}$.

Theorem A.2. *Let R be a real closed field and let A be a real reduced commutative R-algebra. Let $b, g_1, \dots, g_r \in A$ be such that*

(1) $b \notin \mathrm{QM}_A(g_1, \dots, g_r)$,

(2) $\mathrm{U}_A(g_1, \dots, g_r)$ is Zariski dense in $\mathrm{Sper}(A)$.

Then there exists a real closed field R' containing R and an R'-linear map

$$\phi \colon A' = R' \otimes_R A \to R'$$

such that $\phi(1 \otimes b) < 0$ and $\phi(c) \geq 0$ for all $c \in \mathrm{QM}_{A'}(g_1, \dots, g_r)$.

Proof. Since A is real reduced, we know that $\mathrm{Sper}(A) \neq \emptyset$. If $g_i = 0$ for some $i \in \{1, \dots, r\}$, then $\mathrm{U}_A(g_1, \dots, g_r) = \emptyset$, contradicting our second assumption. Therefore, all the g_i are nonzero. In order to simplify notation we further let $g_0 = 1$. Let H be a finite-dimensional R-subspace of A containing $b, g_0, \dots, g_r$. Consider the set

$$C(H) := \left\{ \sum_{i=0}^r g_i \cdot \sum_{j=1}^{k_i} a_{ij}^2 \;\middle|\; k_i \in \mathbb{N},\, a_{ij} \in H \right\}.$$

Note that $C(H)$ is a convex cone inside the finite-dimensional subspace $H^{(2)} := \mathrm{span}_R\{g_i a_1 a_2 \in A \mid a_1, a_2 \in H,\, 0 \leq i \leq r\}$ and $b \in H^{(2)}$ but $b \notin C(H)$. Since A is real reduced and by condition (2), we can apply Lemma A.1 to the maps

$$\alpha_i : H \to H^{(2)}, \quad a \mapsto a^2 \cdot g_i.$$

Therefore, there exists an R-linear map $\phi_H \colon H^{(2)} \to R$ such that $\phi_H(c) \geq 0$ for all $c \in C(H)$ and $\phi_H(b) < 0$. Let $\phi'_H \colon A \to R$ be an arbitrary linear extension of ϕ_H to A. Consider an ultrafilter $\mathcal{U}$ on the set of finite-dimensional subspaces of A, containing the filter generated by sets of the form $S_H := \{H' \mid H \subset H'\}$ for H a finite-dimensional subspace, and consider the ultraproduct $R' := \prod^{\mathcal{U}} R$, which is again a real closed field.

There is a natural R-linear map $\phi' \colon A \to R'$ induced by all ϕ'_H for H as above. Consider the R'-linear map

$$\phi := 1_{R'} \otimes_R \phi' \colon R' \otimes_R A \to R'.$$

It is clear by construction that $\phi(1 \otimes b) < 0$. We claim that ϕ is nonnegative on $\mathrm{QM}_{A'}(g_1, \ldots, g_r)$. It suffices to show that ϕ is nonnegative on elements of the form $a^2 g_k$ for $a \in R' \otimes_R A$ and $0 \le k \le r$. We write $a = \sum_{j=1}^{l} x_j \otimes y_j$ with $x_j \in R'$ and $y_j \in A$. Let H be a finite-dimensional subspace of A containing $b, g_0, \ldots, g_r$ and $y_1, \ldots, y_l$. Consider the matrix

$$(\phi(1 \otimes g_k y_i y_j))_{i,j=1}^{l} = (\phi'(g_k y_i y_j))_{i,j=1}^{l} \in M_l(R').$$

For $x = (x_1, \ldots, x_l) \in R'^l$ and $y := (g_k y_i y_j)_{i,j=1}^{l}$ we get $\phi(a^2 g_k) = x \phi'(y) x^\top$, so that it is sufficient to show that $\phi_H(y) \in M_l(R)$ is positive semidefinite for all such H. However, this is clear since for $x' \in R^l$, we have

$$x' \phi_H(y) x'^\top = \phi_H((x'_1 y_1 + \cdots + x'_l y_l)^2 \cdot g_k) \ge 0$$

by construction. This finishes the proof. ∎

Corollary A.3. *Let R be a real closed field and let A a real reduced commutative R-algebra. Suppose that $b \in A$ is not a sum of squares. Then there exists a real closed field R' containing R and an R'-linear map*

$$\phi \colon R' \otimes_R A \to R'$$

which is nonnegative on sums of squares such that $\phi(1 \otimes b) < 0$.

Remark A.4. Let $X \subset R^n$ be some algebraic set and $A = R[X]$ its coordinate ring, i.e., the ring of all polynomial functions $X \to R$. Then A is real reduced and the second condition of Theorem A.2 is equivalent to the geometric condition that

$$\{a \in X \mid g_1(a) > 0, \ldots, g_r(a) > 0\}$$

is Zariski dense in X.

Appendix B. Preservation theorems

The approach to study definable sets via structure-preserving maps can be used to obtain exact characterizations of existential positive, positive, and first-order definability. There are various formulations of such preservation theorems. Our formulation in Theorem B.1 (in Section B.2) is closest to the setting of this article and we have been unable to find it in the literature, but it should be considered to be known.

B.1. Basic terminology

A *signature* τ is a set of *function symbols* and *relation symbols*. Each symbol is equipped with an arity $k \in \mathbb{N}$. A τ-*structure* $\mathfrak{A}$ consists of a set A (the *domain*) and

- for each function symbol $f \in \tau$ of arity k, an operation $f^{\mathfrak{A}} \colon A^k \to A$,
- for each relation symbol $R \in \tau$ of arity k, a relation $R \subseteq A^k$.

Note that the arity might in particular be 0; function symbols of arity 0 are also called *constants*. We denote structures by $\mathfrak{A}$, $\mathfrak{B}$, $\mathfrak{C}$, and so on, and the corresponding underlying sets by the corresponding Roman letters A, B, C, and so on. If $\mathfrak{A}$ and $\mathfrak{B}$ are structures with the same domain and $\mathfrak{A}$ is obtained from $\mathfrak{B}$ by dropping some of the relations and operations, then $\mathfrak{A}$ is called a *reduct* of $\mathfrak{B}$, and $\mathfrak{B}$ is called an *expansion* of $\mathfrak{A}$.

In the usual way we define τ-*terms* inductively; constant symbols from τ are τ-terms, variables are τ-terms, and a function symbol of arity k applied to a sequence of τ-terms of length k is a τ-term. If t is a τ-term, we write $t(x_1, \ldots, x_n)$ if all the variables that appear in t are from $x_1, \ldots, x_n$. An *atomic τ-formula* is an expression of the form

- $\bot$ (which stands for *false*), or
- $t_1 = t_2$ where t_1 and t_2 are τ-terms, or
- $R(t_1, \ldots, t_k)$ where $R \in \tau$ is a relation symbol of arity k and $t_1, \ldots, t_k$ are τ-terms.

A *first-order τ-formula* is built from atomic τ-formulas in the usual inductive manner, using conjunction $\wedge$, disjunction $\vee$, negation $\neg$, universal quantification $\forall$, and existential quantification $\exists$. If ϕ is a first-order formula, we write $\phi(x_1, \ldots, x_n)$ if all the variables that appear *free* in ϕ (i.e., are not bound by a universal or existential quantifier; for formal details, see [13]) are from $x_1, \ldots, x_n$. Formulas without free variables are called *sentences*. If $\phi(x_1, \ldots, x_k)$ is a τ-formula and $a_1, \ldots, a_k \in A$ we define $\mathfrak{A} \models \phi(a_1, \ldots, a_n)$ if $(a_1, \ldots, a_k)$ *satisfies* ϕ in $\mathfrak{A}$ (see [13] for a precise definition). A τ-formula $\phi(x_1, \ldots, x_k)$ defines over $\mathfrak{A}$ a relation $R \subseteq A^k$, namely

$$\{(a_1, \ldots, a_k) \in A^k \mid \mathfrak{A} \models \phi(a_1, \ldots, a_k)\}.$$

A τ-formula $\phi(x_1, \ldots, x_k)$ defines over $\mathfrak{A}$ an operation $f \colon A^{k-1} \to A$ if ϕ defines over $\mathfrak{A}$ the *graph* of f, i.e., the relation

$$\{(a_1, \ldots, a_k) \in A^k \mid f(a_1, \ldots, a_{k-1}) = a_k\}.$$

We say that $\mathfrak{A}$ is a *first-order reduct* of $\mathfrak{B}$ if $\mathfrak{A}$ is a reduct of an expansion of $\mathfrak{B}$ by relations and operations that are first-order definable in $\mathfrak{B}$.

B.2. Preservation theorems

A first-order formula is called *existential positive* if it is built without negation and universal quantification. Note that every existential positive formula is equivalent to a disjunction of *primitive positive formulas*, i.e., formulas of the form

$$\exists x_1, \ldots, x_k \, (\psi_1 \wedge \cdots \wedge \psi_m)$$

where $\psi_1, \ldots, \psi_m$ are atomic formulas.

Theorem B.1. *Let $\mathfrak{A}$ be a structure. Then $\mathfrak{A}$ has an elementary extension $\mathfrak{A}^*$ such that the following holds. If $\mathfrak{B}$ is a first-order reduct of $\mathfrak{A}$ and $K \subseteq A^k$ is first-order definable over $\mathfrak{A}$, then*

- *K has an existential positive definition in $\mathfrak{B}$ if and only if K^* is preserved by all endomorphisms of $\mathfrak{B}^*$;*

- *K has a positive definition in $\mathfrak{B}$ if and only if K^* is preserved by all surjective endomorphisms of $\mathfrak{B}^*$;*

- *K has a first-order definition in $\mathfrak{B}$ if and only if K^* is preserved by all automorphisms of $\mathfrak{B}^*$.*

Here, $\mathfrak{B}^*$ is the first-order reduct of $\mathfrak{A}^*$ defined by the same formulas as the first-order reduct $\mathfrak{B}$ of $\mathfrak{A}$, and K^* is the relation defined over $\mathfrak{A}^*$ by the same first-order formula that defines K over $\mathfrak{A}$.

We prove Theorem B.1 in Section B.4.

Remark B.2. In our setting, the real closed field $\mathbb{R}$ plays the role of $\mathfrak{A}$. For the results in Section 3 we then work with the first-order reduct $\mathfrak{B}$ of $\mathfrak{A}$ that contains all spectrahedra as relations, to show that specific relations K with a first-order definition in $\mathfrak{A}$ do not have a positive definition in $\mathfrak{B}$. In Section 4 we work with the first-order reduct $\mathfrak{B}$ of $\mathfrak{A}$ that contains all relations that are defined by LMIs of size k to show that some specific relations K with a first-order definition in $\mathfrak{A}$ do not have a positive definition in $\mathfrak{B}$. In both settings, Lemma 2.3 is a statement about the endomorphisms of $\mathfrak{B}^*$.

The elementary extension $\mathfrak{A}^*$ of $\mathfrak{A}$ in Theorem B.1 will be a so-called *special model* (see, e.g., [13]). In the context of preservation theorems, *saturated structures* are often used. Special structures are more technical to define than saturated structures. However, the existence of saturated structures requires the Generalized Continuum Hypothesis which we would like to avoid.

B.3. Saturated and special structures

For $n \geq 0$, an *n-type* of a theory T is a set p of formulas with free variables $x_1, \ldots, x_n$ such that $p \cup T$ is satisfiable. An *n-type over a structure* $\mathfrak{A}$ is an n-type of the first-order theory of $\mathfrak{A}$; note that we do not admit parameters in these formulas (i.e., we allow constants to appear in the formulas of the type only if they belong to the signature). An n-type p of $\mathfrak{A}$ is *realized* in $\mathfrak{A}$ if there exist $a_1, \ldots, a_n \in A$ such that $\mathfrak{A} \models \phi(a_1, \ldots, a_n)$ for each $\phi \in p$. For an infinite cardinal κ, a structure $\mathfrak{A}$ is *κ-saturated* if, for all $\beta < \kappa$ and expansions $\mathfrak{A}'$ of $\mathfrak{A}$ by at most β constants, every 1-type of $\mathfrak{A}'$ is realized in $\mathfrak{A}'$. We say that an infinite structure $\mathfrak{A}$ is *saturated* if it is $|A|$-saturated. If λ is a cardinal, then λ^+ denotes the successor cardinal of λ.

Theorem B.3 ([16, Theorem 4.3.12]). *Let τ be a signature and $\lambda \geq |\tau|$. Then every τ-structure $\mathfrak{B}$ has a λ^+-saturated elementary extension of cardinality $\leq |B|^\lambda$.*

So if we assume the Generalized Continuum Hypothesis, then there are saturated models of size κ^+ for all cardinals κ [13, Corollary 10.2.4].

Definition B.4. A structure $\mathfrak{A}$ of cardinality $\kappa \geq \omega$ is *special* if $\mathfrak{A}$ is the union of an elementary chain $(\mathfrak{B}_\lambda)_{\lambda < \kappa}$ and each $\mathfrak{B}_\lambda$ is λ^+-saturated.

Saturated structures are special, because we can write a saturated structure $\mathfrak{A}$ of cardinality κ as $(\mathfrak{B}_\lambda)_{\lambda < \kappa}$ where $\mathfrak{B}_\lambda = \mathfrak{A}$; since $\mathfrak{A}$ is λ^+-saturated for all $\lambda < \kappa$, this proves that $\mathfrak{A}$ is special. The converse is not true; see [13, Exercises 6 to 8, p. 514]. The advantage of special structures is that they always exist, independently from set-theoretic assumptions.

Theorem B.5 ([13, Theorem 10.4.2]). *Every infinite structure has a special elementary extension.*

B.4. Proof of the preservation theorems

The following is at the heart of the proof of Theorem B.1.

Theorem B.6. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be special structures of the same signature and cardinality.*

- *If every first-order sentence that is true in $\mathfrak{A}$ is true in $\mathfrak{B}$, then there exists an isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$.*
- *If every positive sentence that is true in $\mathfrak{A}$ is true in $\mathfrak{B}$, then there exists a surjective homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.*
- *If every existential positive sentence that is true in $\mathfrak{A}$ is true in $\mathfrak{B}$, then there exists a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.*

Proof. The first statement can be shown by a back and forth argument, and appears in [13, Theorem 10.4.4]. The other two statements can be shown analogously. ∎

Lemma B.7 (see [13, Theorem 10.4.5]). *First-order reducts of special structures are special.*

Proof of Theorem B.1. In each of the three statements, the forward direction is immediate, since homomorphisms preserve existential positive formulas, surjective homomorphisms preserve positive formulas, and automorphisms preserve first-order formulas. Moreover, if K has a certain definition in $\mathfrak{B}$, then K^* has the corresponding definition in $\mathfrak{B}^*$ since all the structures and sets involved are first-order definable in $\mathfrak{A}$ and $\mathfrak{A}^*$, respectively, and $\mathfrak{A}$ and $\mathfrak{A}^*$ have the same first-order theory.

The three statements are well-known if $\mathfrak{A}$ is finite (see, e.g., [25]) so let us assume that $\mathfrak{A}$ is infinite. Let $\mathfrak{A}^*$ be a special elementary extension of $\mathfrak{A}$ which exists by Theorem B.5, and let S be the first-order theory of $\mathfrak{A}$ (and $\mathfrak{A}^*$). Let $\mathfrak{B}$ be a first-order reduct of $\mathfrak{A}$ with signature τ and let $K \subseteq A^k$ be first-order definable over $\mathfrak{A}$. Let $\mathfrak{B}^*$ be the reduct of $\mathfrak{A}^*$ and $K^* \subseteq (A^*)^k$ be as in the statement of the theorem. By Lemma B.7, $\mathfrak{B}^*$ is special as well.

To show the converse of the first statement, suppose that K^* is preserved by all endomorphisms of $\mathfrak{B}^*$. Let Ψ be the set of all existential positive τ-formulas $\psi(x_1, \ldots, x_k)$ that hold on all tuples from K^* in $\mathfrak{B}^*$. Let $\underline{a}$ be a tuple satisfying Ψ in $\mathfrak{B}^*$. Note that if there is no such tuple, then ϕ must be unsatisfiable in $\mathfrak{B}$. In this case, ϕ is equivalent to the existential positive formula $\bot$ and we are done.

Let U be the set of all primitive positive τ-formulas $\theta(\underline{x})$ such that $\mathfrak{B}^* \models \neg\theta(\bar{a})$. For any first-order τ-formula η let $T(\eta)$ be the first-order formula obtained from η by replacing each atomic subformula of η by its first-order definition over $\mathfrak{A}^*$. We claim that

$$S \cup \{T(\neg\theta) \mid \theta \in U\} \cup \{\phi\}$$

is satisfiable. Otherwise, by the compactness theorem of first-order logic there would be a finite subset V of U such that $S \cup \{T(\neg\theta) \mid \theta \in V\} \cup \{\phi\}$ is unsatisfiable. Since $\neg T(\theta) = T(\neg\theta)$ we find that $S \cup \{\phi\}$ implies $\bigvee_{\theta \in V} T(\theta)$. Thus, $\psi := \bigvee_{\theta \in V} \theta$ is an existential positive formula such that $S \cup \{\phi\} \models T(\psi)$. But then $\psi \in \Psi$ and $\mathfrak{B}^* \models \psi(\underline{a})$. This is in contradiction to the assumption that $\mathfrak{B}^* \models \neg\theta(\underline{a})$ for all $\theta \in U$.

Since $\mathfrak{A}^*$ is a special model of S, the n-type $\{T(\neg\theta(\bar{x})) \mid \theta \in U\} \cup \{\phi(\bar{x})\}$ is realized by some $b \in (A^*)^k$ in $\mathfrak{A}^*$. Let $c_1, \ldots, c_k$ be new constant symbols and let $\mathfrak{C}$ be the $\tau \cup \{c_1, \ldots, c_k\}$-expansion of $\mathfrak{B}^*$ where $\underline{c}^{\mathfrak{C}} = \underline{b}$. Let $\mathfrak{D}$ be the $\tau \cup \{c_1, \ldots, c_k\}$-expansion of $\mathfrak{B}^*$ where $\underline{c}^{\mathfrak{C}} = \underline{a}$. Clearly, $\mathfrak{C}$ and $\mathfrak{D}$ are special as well. Every primitive positive $(\tau \cup \{c_1, \ldots, c_k\})$-sentence θ that is true in $\mathfrak{D}$ is also true in $\mathfrak{C}$: otherwise, if θ were false in $\mathfrak{C}$, then write θ as $\theta(\underline{c})$ for $\theta \in U$, in contradiction to the assumption that $\mathfrak{D} \models \{\neg\theta(\underline{c}) \mid \theta \in U\}$. Hence, since $\mathfrak{C}$ is special and by Theorem B.6, there exists a homomorphism from $\mathfrak{D}$ to $\mathfrak{C}$. This homomorphism maps $\underline{b}$ to $\underline{a}$. Since $\underline{b}$ satisfies ϕ in $\mathfrak{A}^*$ we deduce that $\underline{b} \in K^*$, and since K^* is preserved by endomorphisms of $\mathfrak{B}^*$, we have $\underline{a} \in K^*$. We conclude that $S \cup \{T(\theta) \mid \theta \in \Psi\} \cup \{\neg\phi\}$ is unsatisfiable, and again by compactness there exists a finite subset Ψ' of Ψ such that $S \cup \{T(\theta) \mid \theta \in \Psi'\} \cup \{\neg\phi\}$ is unsatisfiable. Note that $\psi := \bigwedge \Psi'$ is an existential positive formula and

$$S \models \forall \underline{x}\, (T(\psi) \Leftrightarrow \phi),$$

which is what we wanted to show.

The converse directions for the second and third statements of the theorem can be shown analogously. ∎

References

[1] Alling, N. L.: Foundations of analysis over surreal number fields. North-Holland Math. Stud. 141, North-Holland, Amsterdam (1987) Zbl 0621.12001 MR 0886475

[2] Averkov, G.: Optimal size of linear matrix inequalities in semidefinite approaches to polynomial optimization. SIAM J. Appl. Algebra Geom. **3**, 128–151 (2019) Zbl 1420.90043 MR 3922906

[3] Berg, C., Christensen, J. P. R., Ressel, P.: Harmonic analysis on semigroups. Grad. Texts in Math. 100, Springer, New York (1984) Zbl 0619.43001 MR 0747302

[4] Bettiol, R. G., Kummer, M., Mendes, R. A. E.: Convex algebraic geometry of curvature operators. SIAM J. Appl. Algebra Geom. **5**, 200–228 (2021) Zbl 1475.14112 MR 4252070

[5] Bochnak, J., Coste, M., Roy, M.-F.: Real algebraic geometry. Ergeb. Math. Grenzgeb. (3) 36, Springer, Berlin (1998) Zbl 0912.14023 MR 1659509

[6] Brown, N. P., Ozawa, N.: C^*-algebras and finite-dimensional approximations. Grad. Stud. Math. 88, Amer. Math. Soc., Providence, RI (2008) Zbl 1160.46001 MR 2391387

[7] Diananda, P. H.: On non-negative forms in real variables some or all of which are non-negative. Proc. Cambridge Philos. Soc. **58**, 17–25 (1962) Zbl 0108.04803 MR 0137686

[8] Fawzi, H.: The set of separable states has no finite semidefinite representation except in dimension 3×2. Comm. Math. Phys. **386**, 1319–1335 (2021) Zbl 1472.14066 MR 4299123

[9] Forsgård, J., de Wolff, T.: The algebraic boundary of the sonc-cone. SIAM J. Appl. Algebra Geom. **6**, 468–502 (2022) Zbl 1506.14107 MR 4477805

[10] Hall, M., Jr., Newman, M.: Copositive and completely positive quadratic forms. Proc. Cambridge Philos. Soc. **59**, 329–339 (1963) Zbl 0124.25302 MR 0147484

[11] Helton, J. W., Nie, J.: Sufficient and necessary conditions for semidefinite representability of convex hulls and sets. SIAM J. Optim. **20**, 759–791 (2009) Zbl 1190.14058 MR 2515796

[12] Hilbert, D.: Ueber die Darstellung definiter Formen als Summe von Formenquadraten. Math. Ann. **32**, 342–350 (1888) JFM 20.0198.02 MR 1510517

[13] Hodges, W.: Model theory. Encyclopedia Math. Appl. 42, Cambridge University Press, Cambridge (1993) Zbl 0789.03031 MR 1221741

[14] Horn, R. A., Johnson, C. R.: Matrix analysis. Cambridge University Press, Cambridge (1985) Zbl 0576.15001 MR 0832183

[15] Karlin, S.: Total positivity. Vol. I. Stanford University Press, Stanford, CA (1968) Zbl 0219.47030 MR 0230102

[16] Marker, D.: Model theory. Grad. Texts in Math. 217, Springer, New York (2002) Zbl 1003.03034 MR 1924282

[17] Marker, D., Peterzil, Y., Pillay, A.: Additive reducts of real closed fields. J. Symbolic Logic **57**, 109–117 (1992) Zbl 0774.03020 MR 1150928

[18] Motzkin, T. S.: The arithmetic-geometric inequality. In: Inequalities (Proc. Sympos. Wright-Patterson Air Force Base, Ohio, 1965), Academic Press, New York, 205–224 (1967) MR 0223521

[19] Motzkin, T. S., Straus, E. G.: Maxima for graphs and a new proof of a theorem of Turán. Canad. J. Math. **17**, 533–540 (1965) Zbl 0129.39902 MR 0175813

[20] Murray, R., Chandrasekaran, V., Wierman, A.: Newton polytopes and relative entropy optimization. Found. Comput. Math. **21**, 1703–1737 (2021) Zbl 1483.90121 MR 4343021

[21] Netzer, T., Thom, A.: Real closed separation theorems and applications to group algebras. Pacific J. Math. **263**, 435–452 (2013) Zbl 1318.14004 MR 3068552

[22] Nie, J.: Semidefinite representability. In: Semidefinite optimization and convex algebraic geometry, MOS-SIAM Ser. Optim. 13, SIAM, Philadelphia, PA, 251–291 (2013) Zbl 1360.90192 MR 3050245

[23] Parrilo, P. A.: Polynomial optimization, sums of squares, and applications. In: Semidefinite optimization and convex algebraic geometry, MOS-SIAM Ser. Optim. 13, SIAM, Philadelphia, PA, 47–157 (2013) Zbl 1360.90194 MR 3050242

[24] Peterzil, Y.: Reducts of some structures over the reals. J. Symbolic Logic **58**, 955–966 (1993) Zbl 0790.03040 MR 1242049

[25] Pöschel, R., Kalužnin, L. A.: Funktionen- und Relationenalgebren. Mathematische Monographien 15, Deutscher Verlag der Wissenschaften, Berlin (1979) Zbl 0421.03049 MR 0543839

[26] Powers, V., Reznick, B.: A note on mediated simplices. J. Pure Appl. Algebra **225**, article no. 106608 (2021) Zbl 1464.11045 MR 4217929

[27] Prestel, A., Delzell, C. N.: Positive polynomials. Springer Monogr. Math., Springer, Berlin (2001) Zbl 0987.13016 MR 1829790

[28] Reznick, B.: Forms derived from the arithmetic-geometric inequality. Math. Ann. **283**, 431–464 (1989) Zbl 0637.10015 MR 0985241

[29] Saunderson, J.: Limitations on the expressive power of convex cones without long chains of faces. SIAM J. Optim. **30**, 1033–1047 (2020) Zbl 1436.90108 MR 4080391

[30] Scheiderer, C.: Semidefinite representation for convex hulls of real algebraic curves. SIAM J. Appl. Algebra Geom. **2**, 1–25 (2018) Zbl 1390.14173 MR 3755651

[31] Scheiderer, C.: Spectrahedral shadows. SIAM J. Appl. Algebra Geom. **2**, 26–44 (2018) Zbl 1391.90462 MR 3755652

[32] Scheiderer, C.: Second-order cone representation for convex sets in the plane. SIAM J. Appl. Algebra Geom. **5**, 114–139 (2021) Zbl 1462.90085 MR 4234144

[33] Tent, K., Ziegler, M.: A course in model theory. Lect. Notes Logic 40, Association for Symbolic Logic, La Jolla, CA; Cambridge University Press, Cambridge (2012) Zbl 1245.03002 MR 2908005