



Newton Polytopes and Relative Entropy Optimization

Riley Murray¹ · Venkat Chandrasekaran¹ · Adam Wierman¹

Received: 8 October 2018 / Revised: 12 May 2020 / Accepted: 20 December 2020 /
Published online: 5 March 2021
© SFoCM 2021

Abstract

Certifying function nonnegativity is a ubiquitous problem in computational mathematics, with especially notable applications in optimization. We study the question of certifying nonnegativity of signomials based on the recently proposed approach of Sums-of-AM/GM-Exponentials (SAGE) decomposition due to the second author and Shah. The existence of a SAGE decomposition is a sufficient condition for nonnegativity of a signomial, and it can be verified by solving a tractable convex relative entropy program. We present new structural properties of SAGE certificates such as a characterization of the extreme rays of the cones associated to these decompositions as well as an appealing form of sparsity preservation. These lead to a number of important consequences such as conditions under which signomial nonnegativity is equivalent to the existence of a SAGE decomposition; our results represent the broadest-known class of nonconvex signomial optimization problems that can be solved efficiently via convex relaxation. The analysis in this paper proceeds by leveraging the interaction between the convex duality underlying SAGE certificates and the face structure of Newton polytopes. After proving our main signomial results, we direct our machinery toward the topic of globally nonnegative polynomials. This leads to (among other things) efficient methods for certifying polynomial nonnegativity, with complexity independent of the degree of a polynomial.

Keywords Arithmetic-geometric-mean inequality · Certifying nonnegativity · Fewnomials · SAGE · Signomials · Sparse polynomials

Mathematics Subject Classification 52A40 · 90C30 · 14P15

Communicated by James Renegar.

✉ Riley Murray
rmurray@caltech.edu

Venkat Chandrasekaran
venkatc@caltech.edu

Adam Wierman
adamw@caltech.edu

¹ California Institute of Technology, Pasadena, CA 91125, USA

1 Introduction

The problem of certifying function nonnegativity is broadly applicable in the mathematical sciences. Optimization provides an especially notable example: since $f^* = \inf_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$ can be expressed as $f^* = \sup\{\gamma : f(\mathbf{x}) - \gamma \geq 0 \forall \mathbf{x} \in \mathbb{R}^n\}$, any certificate that $f - \gamma$ is nonnegative on \mathbb{R}^n is a proof that γ is a lower bound on f^* . Other tasks that can be framed as nonnegativity certification problems include deciding stability of a dynamical system, bounding Lipschitz constants, or determining the consistency of nonlinear equations. It is computationally intractable to decide function nonnegativity in general, but we may yet seek *sufficient conditions* for nonnegativity which are tractable to certify. This article is concerned with the Sums-of-AM/GM-Exponentials or “SAGE” certificates of signomial nonnegativity [7].

Signomials are functions of the form $\mathbf{x} \mapsto \sum_{i=1}^m c_i \exp(\mathbf{a}_i^\top \mathbf{x})$, where $\mathbf{c} \in \mathbb{R}^m$ is the coefficient vector and the \mathbf{a}_i are called *exponent vectors*. Signomial nonnegativity is NP-Hard to decide by reduction to matrix copositivity [32]. Global signomial optimization is tractable in the special case of convex *geometric programming* [6], however high-fidelity signomial models often do not fit into this framework [5,8,15,19,23,26]. Our study of SAGE certificates is motivated by a desire to better understand signomials as a fundamental class of functions, which evidently stands to improve our understanding of nonconvex signomial programming. What is perhaps less evident, but just as important, is that studying signomials forces us to completely set aside the notion of *degree* that is often taken for granted in the study of polynomials. This leads us to develop computational tools that apply to polynomials when $\{\mathbf{a}_i\}_{i=1}^m \subset \mathbb{N}^n$, but in a way that is agnostic to polynomial degree.

There are two perspectives with which one can look at SAGE certificates. On the one hand, there is an interpretation where the summands are certified as nonnegative by a weighted AM/GM inequality. Equivalently, the summands may be certified as nonnegative by the principle of strong duality as applied to a suitable convex program. The connection to convex duality is essential, as it provides a way to tractably represent the SAGE cone via the *relative entropy function*

$$D(\mathbf{v}, \boldsymbol{\lambda}) = \sum_{i=1}^m v_i \ln(v_i/\lambda_i), \quad \mathbf{v}, \boldsymbol{\lambda} \in \mathbb{R}_+^m.$$

Over the course of this article, we employ convex analysis and convex geometry in an effort to better understand fundamental aspects of SAGE certificates. The convex-analytic arguments are used to establish structural results concerning the SAGE cone itself—for example, the nature of its extreme rays, and sparsity-preservation properties of SAGE decompositions. The convex-geometric considerations come into play when we ask precisely how the SAGE cone relates to the nonnegativity cone. Here we use *Newton polytopes*, which are convex hulls of the exponents $\{\mathbf{a}_i\}_{i=1}^m$, often jointly with consideration to the sign pattern of a signomial’s coefficient vector.

Analysis by Newton polytopes has a longer history in study of polynomials, and in particular in the study of sparse polynomials. Prominent examples in this area include Khovanskii’s fewnomials [21,22], Reznick’s agiforms [41,42], and Bajbar and Stein’s work on polynomial coercivity [4]. Many such works “signomialize” polynomials via a substitution $x_j \leftarrow \exp y_j$ in certain intermediate proofs. We adopt a different

perspective, where signomials are the first-class object. This perspective allows us to properly study signomials as a class of functions distinct from polynomials, while also providing a transparent mechanism to obtain polynomial results down the line.

1.1 Our Contributions

Section 2 briefly reviews relevant background on SAGE certificates for signomial nonnegativity.

In Sect. 3, we prove a number of new structural properties of SAGE certificates. Theorem 3 proves an important sparsity-preserving property of SAGE functions. If a signomial f is SAGE, then there exists a decomposition $f = \sum_k f_k$ that certifies this property (i.e., each f_k is an “AGE function”; certifiably nonnegative by a relative entropy inequality, as described in Sect. 2), such that each f_k consists only of those terms that appear in f . Furthermore, the process of summing the f_k to obtain f results in no cancelation of coefficients on basis functions $\mathbf{x} \mapsto \exp(\mathbf{a}_i^T \mathbf{x})$. Theorem 5 goes on to provide a characterization of the extreme rays of the cone of SAGE functions; in particular, all nontrivial extreme rays are given by AGE functions that are supported on simplicial Newton polytopes.

Section 4 leverages the understanding from Sect. 3 to derive a collection of structural results which describe when nonnegative signomials are SAGE, with the Newton polytope being the primary subject of these theorems’ hypotheses. Theorem 9 is concerned with cases where the Newton polytope is simplicial, while Theorems 10 and 11 concern when it “decomposes” in an appropriate sense. Each of these theorems exhibits invariance under nonsingular affine transformations of the exponent vectors. Corollaries 13 and 14 show how Theorem 9 applies to signomial optimization problems. We conclude the section with a result on conditions under which SAGE can recognize signomials that are bounded below (Theorem 15).

In Sect. 5, we specialize our results on signomials to polynomials, by defining a suitable “signomial representative” of a polynomial and requiring that the signomial admit a SAGE decomposition. The resulting class of “SAGE polynomials” inherits a tractable representation from the cone of SAGE signomials (Theorem 16) as well as other structural properties on sparsity preservation and extreme rays (Corollaries 20 and 21). Moving from a polynomial to a signomial representative is simple but somewhat delicate, yielding both stronger (Corollary 17) and weaker (Corollary 18) results than in the signomial case. We then situate our results on polynomial SAGE certificates in the broader literature, with specific emphasis on “Sums-of-Nonnegative-Circuits” (SONC, see Sects. 5.3 and 5.4) and “Sums-of-Squares” (SOS, see Sect. 5.5). The section is concluded with discussion on how our results provide the basis for a sparsity-preserving hierarchy of convex relaxations for polynomial optimization problems.

Section 6 demonstrates that there are meaningful senses in which our results from Sect. 4 cannot be improved upon. Through Theorem 31, we provide a novel dual characterization of conditions under which the SAGE cone and the cone of nonnegative signomials coincide.

1.2 Related Work

1.2.1 Signomials and Signomial Optimization

The literature on signomials is quite fragmented, owing to a wide range of conventions used for this class of functions across fields and over time. A recurring theme is to consider signomials as functions $\mathbf{y} \mapsto \sum_{i=1}^m c_i \mathbf{y}^{a_i}$, where monomials $\mathbf{y}^{a_i} \doteq \prod_{j=1}^n y_j^{a_{ij}}$ remain well-defined for real a_{ij} so long as \mathbf{y} is restricted to the positive orthant. In analysis of biochemical reaction networks, signomials are often called *generalized polynomials* [30,31], or simply “polynomials over the positive orthant” [37]. In amoeba theory, one usually calls signomials *exponential sums* [12,51]. Much of the earlier optimization literature referred to signomial programming as “generalized geometric programming,” but this term now means something quite different [6].

In the taxonomy of optimization problems, geometric programming is to signomial optimization what convex quadratic programming is to polynomial optimization. Current approaches to global signomial optimization use successive linear or geometric programming approximations together with branch-and-bound [28,48,49]. Equality constrained signomial programs are often treated by penalty or augmented Lagrangian methods [43], and are notoriously difficult to solve [35,57]. The SAGE approach to signomial optimization does not involve branch-and-bound and does not entail added complexity when considering signomial equations instead of inequalities.

1.2.2 Sums-of-Squares Certificates and Polynomial Optimization

There is a large body of work on SOS certificates for polynomial nonnegativity, and the resulting convex relaxations for polynomial optimization problems [27,39,50]. Over the course of this article, we make two contributions which have direct parallels in the SOS literature.

Our results in Sect. 3 are along the lines of David Hilbert’s 1888 classification of the number of variables “ n ” and the degrees “ $2d$ ” for which SOS-representability coincides with polynomial nonnegativity [16]. The granularity with which we seek such a classification is distinct from that in the SOS literature, as there is no canonical method to take finite-dimensional subspaces of the infinite-dimensional space of signomials.

A principal drawback of the SOS method is that its canonical formulation requires a semidefinite matrix variable of order $\binom{n+d}{d}$ —and the size of this matrix is exponential in the degree d . In Sect. 5, we use SAGE signomials to certify polynomial nonnegativity in a way which is unaffected by the polynomial’s degree. Section 5.5 compares our proposed method to SOS, as well as refinements and variations of SOS which have appeared in the literature: [2,24,34,54].

1.2.3 Certifying Function Nonnegativity via the AM/GM Inequality

The “AGE functions” in a SAGE decomposition may be proven nonnegative in either of two ways. The first approach is to certify a particular relative entropy inequality over a signomial’s coefficients, which is described in Sect. 2. This approach is known to be computationally tractable, and it provides a convenient tool for proving structural

results for the set of SAGE certificates. The second method is to find weights for an appropriate AM/GM inequality over a signomial’s coefficients; this latter method directly connects SAGE to a larger literature on certifying function nonnegativity via the AM/GM inequality, which we summarize in the sequel.

The earliest systematic theoretical studies in this area were undertaken by Reznick [41,42] in the late 1970s and 1980s. The first developments of any computational flavor came from Pébay, Rojas and Thompson in 2009 [40], via their study of polynomial maximization. Pébay et al. used tropical geometry and \mathcal{A} -discriminants to develop an AM/GM type certificate for boundedness of functions supported on matroid-theoretic circuits. In this context, a function is supported on a circuit if the monomial exponents $\{\mathbf{a}_i : c_i \neq 0\}$ form a minimal affinely dependent set. In 2011, Ghasemi et al. pioneered the use of geometric programming to recognize functions which were certifiably nonnegative by the AM/GM-inequality and a sums-of-binomial-squares representation [13,14]. In 2012, Pantea, Koepl, and Craciun derived an AM/GM condition to certify \mathbb{R}_+^n -nonnegativity of polynomials supported on circuits [37, Theorem 3.6]. Follow-up work by August, Craciun, and Koepl used [37, Theorem 3.6] to determine invariant sets for biological dynamical systems [3]. A short while later, Iliman and de Wolff suggested taking sums of globally nonnegative circuit polynomials [17]; the resulting *SONC polynomials* have since become an established topic in the literature [11,18,45,55,56].

We continue to make connections to the AM/GM-certificate literature throughout this article; [11,17,37,45] are revisited in Sect. 5.3, and [17,55,56] are addressed in Sect. 5.4.

1.3 Notation and Conventions

Special sets include $[\ell] \doteq \{1, \dots, \ell\} \subset \mathbb{N}$, the probability simplex $\Delta_\ell \subset \mathbb{R}^\ell$, the nonnegative reals \mathbb{R}_+ , and the positive reals \mathbb{R}_{++} . Vectors and matrices always appear in boldface. For indexing reasons, we find it useful to define certain vectors by writing $\mathbf{v} = (v_i)_{i \in [\ell] \setminus k}$; such a vector is said to belong to “ $\mathbb{R}^{[\ell] \setminus k}$.” If we wish to drop a the k^{th} component from a vector $\mathbf{v} \in \mathbb{R}^\ell$, we write $\mathbf{v}_{\setminus k} \doteq (v_i)_{i \in [\ell] \setminus k}$. The *support* of a vector $\mathbf{v} \in \mathbb{R}^\ell$ is $\text{supp } \mathbf{v} \doteq \{i \in [\ell] : v_i \neq 0\}$. We use \mathbf{e}_i to denote the i^{th} standard basis vector in \mathbb{R}^ℓ with ℓ inferred from context, and set $\mathbf{1} = \sum_{i=1}^\ell \mathbf{e}_i$. The operator \oplus is used for vector concatenation. We often call finite a point-set $\{\mathbf{u}_i\}_{i=1}^\ell$ *simplicial* if it is affinely independent.

The operators “cl” and “conv” return a set’s closure and convex hull, respectively; “ext” returns the extreme points of a compact convex set. We use $U + V$ to denote the Minkowski sum of sets U and V within \mathbb{R}^ℓ . For any convex cone K contained in \mathbb{R}^ℓ there is an associated *dual cone* $K^\dagger \doteq \{\mathbf{y} : \mathbf{y}^\top \mathbf{x} \geq 0 \text{ for all } \mathbf{x} \text{ in } K\}$.

We reserve \mathbf{A} for a real $n \times m$ matrix with distinct columns $\{\mathbf{a}_i\}_{i=1}^m$, and \mathbf{c} for a vector in \mathbb{R}^m . Writing $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ means that f takes values $f(\mathbf{x}) = \sum_{i=1}^m c_i \exp(\mathbf{a}_i^\top \mathbf{x})$. For a fixed signomial $\text{Sig}(\mathbf{A}, \mathbf{c})$, we often refer to $\mathbf{a}_i \in \mathbb{R}^n$ as an *exponent vector*, and use $\mathcal{P}(\mathbf{A})$ to denote its Newton polytope. The cone of coefficients for nonnegative signomials over exponents \mathbf{A} is

$$C_{\text{NNS}}(\mathbf{A}) = \{ \mathbf{c} : \text{Sig}(\mathbf{A}, \mathbf{c})(\mathbf{x}) \geq 0 \text{ for all } \mathbf{x} \text{ in } \mathbb{R}^n \}.$$

We sometimes overload terminology and refer to $C_{\text{NNS}}(\mathbf{A})$ as a cone of signomials, rather than a cone of coefficients. We routinely work with matrices “ $\mathbf{A}_{\setminus k}$ ” formed by deleting some k^{th} column from \mathbf{A} . The matrix $\mathbf{A}_{\setminus k}$ is applied to a vector $\mathbf{v} \in \mathbb{R}^{[m]\setminus k}$ by $\mathbf{A}_{\setminus k}\mathbf{v} = \sum_{i \in [m]\setminus \{k\}} \mathbf{a}_i v_i$.

2 Background Theory on SAGE Functions

In their debut, SAGE functions were used as a building block for a hierarchy of convex relaxations to challenging nonconvex signomial optimization problems [7]. Underlying this entire hierarchy were the simple facts that SAGE functions are globally nonnegative and efficiently recognizable. The purpose of this section is to review the theory of SAGE functions to the extent that it is needed for subsequent development.

Section 2.1 introduces the idea of an *AGE function*, which serve as the building blocks of SAGE functions. The way in which AGE functions extend to SAGE functions is given in Sect. 2.2. Section 2.3 describes the connection between nonnegativity and optimization in the context of SAGE relaxations. These three sections are essential for understanding the present article.

2.1 AM/GM Exponentials

We need some additional structure to make it easier to verify membership in the nonnegativity cone $C_{\text{NNS}}(\mathbf{A})$. The structure used by Chandrasekaran and Shah [7] was that the coefficient vector \mathbf{c} contained at most one negative entry c_k ; if such a function was globally nonnegative, they called it an *AM/GM Exponential*, or an *AGE function*. To facilitate the study of such functions, [7] defines the k^{th} AGE cone

$$C_{\text{AGE}}(\mathbf{A}, k) = \{ \mathbf{c} : \mathbf{c}_{\setminus k} \geq \mathbf{0} \text{ and } \mathbf{c} \text{ belongs to } C_{\text{NNS}}(\mathbf{A}) \}. \tag{1}$$

Elements of $C_{\text{AGE}}(\mathbf{A}, k)$ are sometimes called *AGE vectors*.

It is evident that $C_{\text{AGE}}(\mathbf{A}, i)$ is a proper convex cone which contains the nonnegative orthant. By using a convex duality argument, one may show that a vector \mathbf{c} with $\mathbf{c}_{\setminus k} \geq \mathbf{0}$ belongs to $C_{\text{AGE}}(\mathbf{A}, k)$ if and only if

$$\text{some } \mathbf{v} \in \mathbb{R}_+^{[m]\setminus k} \text{ satisfies } [\mathbf{A}_{\setminus k} - \mathbf{a}_k \mathbf{1}^T] \mathbf{v} = \mathbf{0} \text{ and } D(\mathbf{v}, e\mathbf{c}_{\setminus k}) \leq c_k. \tag{2}$$

It is crucial that the representation in (2) is jointly convex in \mathbf{c} and the auxiliary variable \mathbf{v} , and moreover that no assumption is made on the sign of c_k . Using the representation (2), one may derive the following expression for the dual of the k^{th} AGE cone

$$C_{\text{AGE}}(\mathbf{A}, k)^\dagger = \text{cl} \{ \mathbf{v} : \mathbf{v} > 0, \text{ and for some } \boldsymbol{\mu}_k \text{ in } \mathbb{R}^n \text{ we have } v_k \ln(v_k/v_i) \leq (\mathbf{a}_k - \mathbf{a}_i)^T \boldsymbol{\mu}_k \text{ for } i \text{ in } [m] \}. \tag{3}$$

The “size” of a primal or dual AGE cone refers to the number of variables plus the number of constraints in the above representations, which is $O(m)$ assuming $n \leq m$.

2.2 From “AGE” to “SAGE”

We call a signomial *SAGE* if it can be written as a sum of AGE functions. SAGE functions are globally nonnegative by construction; cones of coefficients for SAGE signomials are denoted by

$$C_{SAGE}(A) \doteq \sum_{k=1}^m C_{AGE}(A, k). \tag{4}$$

Standard calculations in conic duality yield the following expression for a dual SAGE cone

$$C_{SAGE}(A)^\dagger = \cap_{k=1}^m C_{AGE}(A, k)^\dagger. \tag{5}$$

Equations 4 and 5 provide natural definitions, but they also contain redundancies.

Proposition 1 [7, Section 2.4] *If $a_k \in \text{ext } \mathcal{P}(A)$ and $f = \text{Sig}(A, c)$ is nonnegative, then $c_k \geq 0$. Consequently, if $a_k \in \text{ext } \mathcal{P}(A)$ then $C_{AGE}(A, k) = \mathbb{R}_+^m$.*

Proposition 1 is the most basic way Newton polytopes appear in the analysis of non-negative signomials. In our context, it means that so long as $\text{ext } \mathcal{P}(A) \subsetneq \{a_i\}_{i=1}^m$, we can take $C_{SAGE}(A)$ as the Minkowski sum of AGE cones $C_{AGE}(A, k)$ where a_k are nonextremal in $\mathcal{P}(A)$.

2.3 From Nonnegativity to Optimization

A means for certifying nonnegativity provides a natural method for computing lower bounds for minimization problems: given a function f , find the largest constant γ where $f - \gamma$ can be certified as globally nonnegative. We now formalize this procedure for signomials and SAGE certificates.

Given a signomial $f = \text{Sig}(A, c)$ and a constant γ in \mathbb{R} , we want to check if $f - \gamma$ is SAGE. To do this, we need an unambiguous representation of $f - \gamma$ in terms of $\text{Sig}(\cdot, \cdot)$ notation. Toward this end, we suppose A has $a_1 = \mathbf{0}$, and we make a point of allowing any entry of c to be zero. Under these assumptions, the function $f - \gamma$ can be written as $\text{Sig}(A, c - \gamma e_1)$. Thus the optimization problem

$$f_{SAGE} \doteq \sup\{ \gamma : c - \gamma e_1 \text{ in } C_{SAGE}(A) \} \tag{6}$$

is well defined, and its optimal value satisfies $f_{SAGE} \leq f^*$. We also analyze the dual problem to (6), and for reference, we obtain via conic duality that

$$f_{SAGE} = \inf\{ c^T v : e_1^T v = 1, v \text{ in } C_{SAGE}(A)^\dagger \}. \tag{7}$$

Although not done in [7], it can be shown that strong duality holds in the primal-dual pair (6)-(7). This fact is important for our later theorems, and so we make a point to state it clearly in the following proposition:

Proposition 2 *Strong duality always holds in the computation of f_{SAGE} .*

The proposition is proven in Appendix 7.2 using convex analysis.

3 Structural Results for SAGE Certificates

This section presents two new geometric results and analytical characterizations on the SAGE cone. These results have applications to polynomial nonnegativity, as discussed later in Sect. 5. Statements of the theorems are provided below along with remarks on the theorems' significance. Proofs are deferred to later subsections.

3.1 Summary of Structural Results

Our first theorem shows that when checking if \mathbf{c} belongs to $\mathcal{C}_{SAGE}(\mathbf{A})$ we can restrict the search space of SAGE decompositions to those exhibiting a very particular structure. It highlights the sparsity-preserving property of SAGE, and in so doing has significant implications for both the practicality of solving SAGE relaxations, and Sect. 5's development of SAGE polynomials.

Theorem 3 *If \mathbf{c} is a vector in $\mathcal{C}_{SAGE}(\mathbf{A})$ with nonempty index set $N \doteq \{i : c_i < 0\}$, then there exist vectors $\{\mathbf{c}^{(i)} \in \mathcal{C}_{AGE}(\mathbf{A}, i)\}_{i \in N}$ satisfying $\mathbf{c} = \sum_{i \in N} \mathbf{c}^{(i)}$ and $c_j^{(i)} = 0$ for all distinct $i, j \in N$.*

We can use Theorem 3 to define some parameterized AGE cones that will be of use to us in Section 4. Specifically, for an index set J contained within $[m]$, and an index i in $[m]$, define

$$\mathcal{C}_{AGE}(\mathbf{A}, i, J) = \{\mathbf{c} : \mathbf{c} \text{ in } \mathcal{C}_{AGE}(\mathbf{A}, i), c_j = 0 \text{ for all } j \text{ in } J \setminus \{i\}\}.$$

In terms of such sets we have the following corollary of Theorem 3.

Corollary 4 *A signomial $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ with $\mathbf{a}_1 = \mathbf{0}$ has*

$$f_{SAGE} = \sup\{\gamma : \mathbf{c} - \gamma \mathbf{e}_1 \text{ in } \sum_{i \in N \cup \{1\}} \mathcal{C}_{AGE}(\mathbf{A}, i, N)\}$$

for both $N = \{i : c_i < 0\}$ and $N = \{i : c_i \leq 0\}$.

This corollary has two implications concerning practical algorithms for signomial optimization. First, it shows that for $k = |\{i : c_i < 0\}|$, computing f_{SAGE} can easily be accomplished with a relative entropy program of size $O(km)$; this is a dramatic improvement over the naïve implementation for computing f_{SAGE} , which involves a relative entropy program of size $O(m^2)$. Second—the improved conditioning resulting from restricting the search space in this way often makes the difference in whether existing solvers can handle SAGE relaxations of moderate size. This point is highlighted in recent experimental demonstrations of relative entropy relaxations; the authors of [20] discuss various preprocessing strategies to more quickly solve such optimization problems.

Our next theorem characterizes the extreme rays of the SAGE cone. To describe these extreme rays, we use a notion from matroid theory [36]: a set of points $X = \{x_i\}_{i=1}^\ell$ is called a *circuit* if it is affinely dependent, but any proper subset $\{x_i\}_{i \neq k}$ is affinely independent. If the convex hull of a circuit with ℓ elements contains $\ell - 1$ extreme points, then we say the circuit is *simplicial*.

Theorem 5 *If $c \in \mathbb{R}^m$ generates an extreme ray of $C_{SAGE}(A)$, then $\{a_i : i \in [m], c_i \neq 0\}$ is either a singleton or a simplicial circuit.*

Theorem 5 can be viewed as a signomial generalization of a result by Reznick concerning *agiforms* [42, Theorem 7.1]. The theorem admits a partial converse: if $X = \{a_j\}_{j \in J \cup \{i\}}$ is a simplicial circuit with nonextremal term a_i , then there is an extreme ray of $C_{AGE}(A, i)$ supported on $J \cup \{i\}$. When specialized to the context of polynomials, this result gives us an equivalence between SAGE polynomials (suitably defined in Sect. 5) and the previously defined SONC polynomials [17], thus providing an efficient description of the latter set which was not known to be tractable.

3.2 Proof of the Restriction Theorem for SAGE Decompositions (Theorem 3)

Our proof requires two lemmas. The first such lemma indicates the claim of the theorem applies far more broadly than for SAGE functions alone.

Lemma 6 *Let $K \subset \mathbb{R}^m$ be a convex cone containing the nonnegative orthant. For an index $i \in [m]$, define $C_i = \{c \in K : c_{\setminus i} \geq \mathbf{0}\}$, and sum these to $C = \sum_{i=1}^m C_i$. We claim that a vector c with at least one negative entry belongs to C if and only if*

$$c \in \sum_{i:c_i < 0} C_i.$$

Proof Suppose $c \in C$ has a decomposition $c = \sum_{i \in N} c^{(i)}$ where each $c^{(i)}$ belongs to C_i . If $N = \{i : c_i < 0\}$, then there is nothing to prove, so suppose there is some $k \in N$ with $c_k \geq 0$. We construct an alternative decomposition of c using only cones C_i with i in $N \setminus \{k\}$.

The construction depends on the sign of $c_k^{(k)}$. If $c_k^{(k)}$ is nonnegative then the problem of removing dependence on C_k is simple: for i in $N \setminus \{k\}$, the vectors

$$\tilde{c}^{(i)} = c^{(i)} + c^{(k)} / (|N| - 1)$$

belong to C_i (since $C_i \supset \mathbb{R}_+^m$), and sum to c . If instead $c_k^{(k)} < 0$, then there exists some index $i \neq k$ in N with $c_k^{(i)}$ positive. This allows us to define the distribution λ with $\lambda_i = c_k^{(i)} / \sum_{j \in N \setminus \{k\}} c_k^{(j)}$ for $i \neq k$ in N . With λ we construct the $|N| - 1$ vectors

$$\tilde{c}^{(i)} = c^{(i)} + \lambda_i c^{(k)}.$$

The vectors $\tilde{c}^{(i)}$ belong to K because they are a conic combination of vectors in K ($c^{(i)}$ and $c^{(k)}$). We claim that for every $i \neq k$ in N , the coordinate $\tilde{c}_k^{(i)}$ is nonnegative.

This is certainly true when $\lambda_i = 0$, but more importantly, $\lambda_i > 0$ implies

$$\frac{1}{\lambda_i} \tilde{c}_k^{(i)} = \frac{1}{\lambda_i} \left(c_k^{(i)} + \lambda_i c_k^{(k)} \right) = \left[\sum_{j \in N \setminus \{k\}} c_k^{(j)} \right] + c_k^{(k)} = c_k \geq 0.$$

Hence \mathbf{c} can be expressed as the sum of vectors $\{\tilde{\mathbf{c}}^{(i)}\}_{i \in N \setminus \{k\}}$ where each vector $\tilde{\mathbf{c}}^{(i)}$ belongs to C_i .

From here, update $N \leftarrow N \setminus \{k\}$. If N contains another index k' with $c_{k'} \geq 0$, then repeat the above procedure to remove the unnecessary cone $C_{k'}$. Naturally, this process continues until $N = \{i : c_i < 0\}$. □

Lemma 7 *Let \mathbf{w}, \mathbf{v} be vectors in \mathbb{R}^m with distinguished indices $i \neq j$ so that*

$$\mathbf{w}_{\setminus i}, \mathbf{v}_{\setminus j} \geq \mathbf{0} \quad \text{and} \quad w_k + v_k < 0 \text{ for } k \text{ in } \{i, j\}.$$

Then there exist vectors $\hat{\mathbf{w}}, \hat{\mathbf{v}}$ in the conic hull of $\{\mathbf{w}, \mathbf{v}\}$ which satisfy

$$\hat{\mathbf{w}} + \hat{\mathbf{v}} = \mathbf{w} + \mathbf{v} \quad \text{and} \quad \hat{w}_j = \hat{v}_i = 0.$$

Proof By reindexing, take $i = 1$ and $j = 2$. Such $\hat{\mathbf{w}}, \hat{\mathbf{v}}$ exist if and only if some λ in \mathbb{R}_+^4 solves

$$\begin{bmatrix} w_1 & 0 & v_1 & 0 \\ 0 & w_2 & 0 & v_2 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}. \tag{8}$$

The determinant of the matrix above is $d = w_1 v_2 - v_1 w_2$. If w_2 or $v_1 = 0$, then $d > 0$. If $w_2, v_1 \neq 0$, then $d > 0 \Leftrightarrow |v_2/w_2| \cdot |w_1/v_1| > 1$. In this case, we use the assumptions on \mathbf{w}, \mathbf{v} to establish the slightly stronger condition that $|v_2/w_2| > 1$ and $|w_1/v_1| > 1$. In both cases we have a nonzero determinant, so there exists a unique λ in \mathbb{R}^4 satisfying system (8). Now we need only prove that this λ is nonnegative.

One may verify that the symbolic solution to (8) is

$$\begin{aligned} \lambda_1 &= -(w_2 + v_2)v_1/d, & \lambda_2 &= (w_1 + v_1)v_2/d, \\ \lambda_3 &= w_1(w_2 + v_2)/d, & \lambda_4 &= -(w_1 + v_1)w_2/d, \end{aligned}$$

and furthermore that all numerators and denominators are nonnegative. □

Proof (Theorem 3) Let \mathbf{c}^* be a vector in $C_{\text{SAGE}}(\mathbf{A})$ with k negative entries c_1^*, \dots, c_k^* . It is clear that the AGE cones $C_{\text{AGE}}(\mathbf{A}, i)$ satisfy the hypothesis of Lemma 6, with $K = C_{\text{NNS}}(\mathbf{A})$. Therefore there exists a k -by- m matrix \mathbf{C} with i^{th} row $\mathbf{c}_i \in C_{\text{AGE}}(\mathbf{A}, i)$, and $\mathbf{c}^* = \sum_{i=1}^k \mathbf{c}_i$. We prove the result by transforming \mathbf{C} into a matrix with rows \mathbf{c}_i satisfying the required properties, using only row-sum preserving conic combinations from Lemma 7.

It is clear that for any pair of distinct i, j , the vectors $\mathbf{c}_i, \mathbf{c}_j$ satisfy the hypothesis of Lemma 7, thus there exist $\hat{\mathbf{c}}_i, \hat{\mathbf{c}}_j$ in the conic hull of $\mathbf{c}_i, \mathbf{c}_j$ where $\hat{c}_{ij} = \hat{c}_{ji} = 0$

and $\hat{c}_i + \hat{c}_j = c_i + c_j$. Furthermore, this remains true if we modify C by replacing $(c_i, c_j) \leftarrow (\hat{c}_i, \hat{c}_j)$.

We proceed algorithmically: apply Lemma 7 to rows $(1, 2)$, then $(1, 3)$, and continuing to rows $(1, k)$. At each step of this process, we eliminate $c_{j1} = 0$ for $j > 1$ and maintain $c_{ji} \geq 0$ for off-diagonal c_{ji} . We then apply the procedure to the second column of C , beginning with rows 2 and 3. Since $c_{j1} = 0$ for $j > 1$, none of the row operations introduce an additional nonzero in the first column of C , and so the first column remains zero below c_{11} , and the second column becomes zero below c_{22} . Following this pattern, we reduce C to have zeros on the strictly lower-triangular block in the first k columns, in particular terminating with $c_{kk} = c_k^* < 0$.

The next phase is akin to back-substitution. Apply Lemma 7 to rows $(k, k - 1)$, then $(k, k - 2)$, and continue until rows $(k, 1)$. This process zeros out the k^{th} column of C above c_{kk} . The same procedure applies with rows $(k - 1, k - 2)$, then $(k - 1, k - 3)$, through $(k - 1, 1)$, to zero the $(k - 1)^{\text{st}}$ column of C except for the single entry $c_{[k-1][k-1]} = c_{k-1}^* < 0$. The end result of this process is that the first k columns of C comprise a diagonal matrix with entries $(c_1^*, \dots, c_k^*) < \mathbf{0}$.

The resulting matrix C satisfies the claimed sparsity conditions. Since all row operations involved conic combinations, each row of the resulting matrix C defines a nonnegative signomial. The theorem follows since row i of the resulting matrix has a single negative component $c_{ii} = c_i^* < 0$. □

3.3 Proof of Extreme Ray Characterization of the SAGE Cone (Theorem 5)

Because every ray in the SAGE cone (extreme or otherwise) can be written as a sum of rays in AGE cones, it suffices to characterize the extreme rays of AGE cones. For the duration of this section, we discuss the AGE cone $C_{\text{AGE}}(\mathbf{A}, k)$, where \mathbf{a}_k is nonextremal in $\mathcal{P}(\mathbf{A})$.

It can easily be shown that for any index i in $[m]$, the ray $\{r\mathbf{e}_i : r \geq 0\}$ is extremal in $C_{\text{AGE}}(\mathbf{A}, k)$. We call these rays (those supported on a single coordinate) the *trivial* extreme rays of the AGE cone. The work in showing Theorem 5 is to prove that all nontrivial extreme rays of the AGE cone are supported on simplicial circuits. Our proof will appeal to the following basic fact concerning polyhedral geometry, which we establish in the appendix.

Lemma 8 Fix $\mathbf{B} \in \mathbb{R}^{n \times d}$, $\mathbf{h} \in \mathbb{R}^n$, and $\Lambda = \{\boldsymbol{\lambda} \in \Delta_d : \mathbf{B}\boldsymbol{\lambda} = \mathbf{h}\}$. For any $\boldsymbol{\lambda} \in \Lambda$, there exist $\{\boldsymbol{\lambda}^{(i)}\}_{i=1}^\ell \subset \Lambda$ and $\boldsymbol{\theta} \in \Delta_\ell$ for which $\{\mathbf{b}_j : \lambda_j^{(i)} > 0\}$ are affinely independent, and $\boldsymbol{\lambda} = \sum_{i=1}^\ell \theta_i \boldsymbol{\lambda}^{(i)}$.

Proof (Theorem 5) We seek an $\ell \in \mathbb{N}$ where we can decompose $\mathbf{c} \in C_{\text{AGE}}(\mathbf{A}, k)$ as a sum of $\ell + 1$ AGE vectors $\{\mathbf{c}^{(i)}\}_{i=1}^{\ell+1} \subset C_{\text{AGE}}(\mathbf{A}, k)$, where $\{\mathbf{a}_j : c_j^{(i)} \neq 0\}$ are simplicial circuits for $i \in [\ell]$ and $\mathbf{c}^{(\ell+1)} \geq \mathbf{0}$. Since \mathbf{c} is an AGE vector, there is an associated $\mathbf{v} \in \mathbb{R}_+^{[m] \setminus k}$ for which $[\mathbf{A}_{\setminus k} - \mathbf{a}_k \mathbf{1}^\top] \mathbf{v} = \mathbf{0}$ and $D(\mathbf{v}, \mathbf{e}\mathbf{c}_{\setminus k}) \leq c_k$. If \mathbf{v} is zero, then $D(\mathbf{v}, \mathbf{e}\mathbf{c}_{\setminus k}) = 0 \leq c_k$, so $\ell = 0$ and $\mathbf{c}^{(\ell+1)} = \mathbf{c}$ provides the required decomposition. The interesting case, of course, is when \mathbf{v} is nonzero. Our proof proceeds by providing a mechanism to decompose \mathbf{v} into a convex combination of certain vectors $\{\mathbf{v}^{(i)}\}_{i=1}^\ell$, and from there obtain suitable AGE vectors $\mathbf{c}^{(i)}$ from each $\mathbf{v}^{(i)}$.

Given $\mathbf{v} \neq \mathbf{0}$, the vector $\boldsymbol{\lambda} \doteq \mathbf{v}/\mathbf{v}^\top \mathbf{1}$ belongs to the probability simplex $\Delta_{[m]\setminus k} \subset \mathbb{R}^{[m]\setminus k}$. We introduce this $\boldsymbol{\lambda}$ because $[\mathbf{A}_{\setminus k} - \mathbf{a}_k \mathbf{1}^\top] \mathbf{v} = \mathbf{0}$ is equivalent to $\mathbf{A}_{\setminus k} \boldsymbol{\lambda} = \mathbf{a}_k$, and the latter form is amenable to Lemma 8. Apply Lemma 8 to decompose $\boldsymbol{\lambda}$ into a convex combination of vectors $\{\boldsymbol{\lambda}^{(i)}\}_{i=1}^\ell \subset \Delta_{[m]\setminus k}$ for which $\{\mathbf{a}_j : j \in \text{supp } \boldsymbol{\lambda}^{(i)}\}$ are simplicial and $\boldsymbol{\lambda}^{(i)}$ satisfy $\mathbf{A}_{\setminus k} \boldsymbol{\lambda}^{(i)} = \mathbf{a}_k$; let $\boldsymbol{\theta} \in \Delta_\ell$ denote the vector of convex combination coefficients for this decomposition of $\boldsymbol{\lambda}$. For each $\boldsymbol{\lambda}^{(i)}$, define $\mathbf{v}^{(i)} = \boldsymbol{\lambda}^{(i)} (\mathbf{v}^\top \mathbf{1})$. These values for $\mathbf{v}^{(i)}$ evidently satisfy $[\mathbf{A}_{\setminus k} - \mathbf{a}_k \mathbf{1}^\top] \mathbf{v}^{(i)} = \mathbf{0}$ and $\sum_{i=1}^\ell \theta_i \mathbf{v}^{(i)} = \mathbf{v}$. From these $\mathbf{v}^{(i)}$ we construct $\mathbf{c}^{(i)} \in \mathbb{R}^m$ by

$$c_j^{(i)} = \begin{cases} (c_j/v_j)v_j^{(i)} & \text{if } v_j > 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } j \neq k,$$

and for $j = k$ we take $c_k^{(i)} = D(\mathbf{v}^{(i)}, e\mathbf{c}_{\setminus k}^{(i)})$.

By construction these $\mathbf{c}^{(i)}$ belong to $\mathcal{C}_{\text{AGE}}(\mathbf{A}, k)$, and $\{\mathbf{a}_j : c_j^{(i)} \neq 0\}$ comprise simplicial circuits. We now take a componentwise approach to showing $\sum_{i=1}^\ell \theta_i \mathbf{c}^{(i)} \leq \mathbf{c}$. For indices $j \neq k$ with $v_j > 0$, the inequality actually holds with equality $\sum_{i=1}^\ell \theta_i c_j^{(i)} = (c_j/v_j)(\sum_{i=1}^\ell \theta_i v_j^{(i)}) = c_j$. Now we turn to showing $\sum_{i=1}^\ell \theta_i c_k^{(i)} \leq c_k$; we specifically claim that

$$\sum_{i=1}^\ell \theta_i c_k^{(i)} = \sum_{i=1}^\ell \theta_i D(\mathbf{v}^{(i)}, e\mathbf{c}_{\setminus k}^{(i)}) = D(\mathbf{v}, e\mathbf{c}_{\setminus k}) \leq c_k. \tag{9}$$

For the three relations in display (9), the first holds from the definitions of $c_k^{(i)}$, and the last holds from our assumptions on (\mathbf{c}, \mathbf{v}) , so only the second equality needs explaining. For this we use the fact that definitions of $c_j^{(i)}$ relative to $v_j^{(i)}$ preserve ratios with c_j relative to v_j , i.e.,

$$D(\mathbf{v}^{(i)}, e\mathbf{c}_{\setminus k}^{(i)}) = \sum_{j \neq k} v_j^{(i)} \ln \left(\frac{v_j^{(i)}}{e c_j^{(i)}} \right) = \sum_{j \neq k} v_j^{(i)} \ln \left(\frac{v_j}{e c_j} \right). \tag{10}$$

One may then prove the middle equality in display (9) by summing $\theta_i D(\mathbf{v}^{(i)}, \mathbf{c}_{\setminus k})$ over i , applying the identity in equation (10), and then interchanging the sums over i and j . Formally,

$$\sum_{i=1}^\ell \theta_i D(\mathbf{v}^{(i)}, e\mathbf{c}_{\setminus k}^{(i)}) = \sum_{j \neq k} \log \left(\frac{v_j}{e c_j} \right) \underbrace{\left(\sum_{i=1}^\ell \theta_i v_j^{(i)} \right)}_{=v_j} = \sum_{j \neq k} v_j \log \left(\frac{v_j}{e c_j} \right) = D(\mathbf{v}, e\mathbf{c}_{\setminus k}).$$

We have effectively established the claim of the theorem. To find a decomposition of the form desired at the beginning of this proof, one rescales $\mathbf{c}^{(i)} \leftarrow \theta_i \mathbf{c}^{(i)}$ and sets $\mathbf{c}^{(\ell+1)} = \mathbf{c} - \sum_{i=1}^\ell \mathbf{c}^{(i)}$. □

4 The Role of Newton Polytopes in SAGE Signomials

This section begins by introducing two theorems (Theorems 9 and 10) concerning SAGE-representability versus signomial nonnegativity. These theorems are then combined to obtain a third theorem (Theorem 11), which provides the most general yet-known conditions for when the SAGE and nonnegativity cones coincide. The proofs of Theorems 9 and 10 are contained in Sects. 4.2 and 4.3. Applications of Theorem 9 are given in Sect. 4.4. Section 4.5 uses a distinct proof strategy (nevertheless Newton-polytope based) to determine a condition on when SAGE can recognize signomials which are bounded below.

4.1 When SAGE Recovers the Nonnegativity Cone

The following theorem is the first instance beyond AGE functions when SAGE representability is known to be equivalent to nonnegativity.

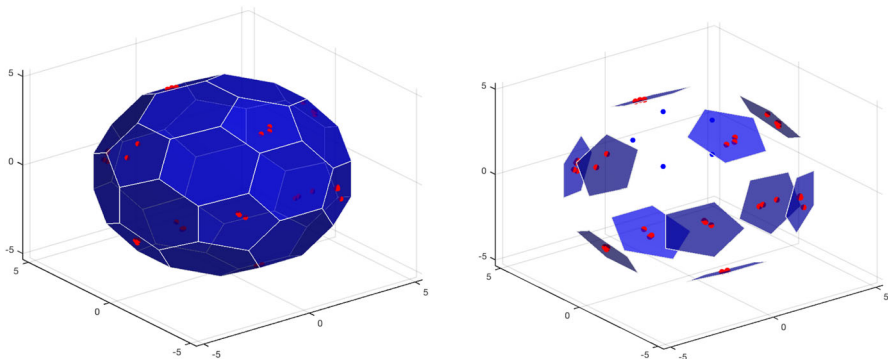
Theorem 9 *Suppose $\text{ext } \mathcal{P}(\mathbf{A})$ is simplicial, and that \mathbf{c} has $c_i \leq 0$ whenever \mathbf{a}_i is nonextremal. Then \mathbf{c} belongs to $C_{\text{SAGE}}(\mathbf{A})$ if and only if \mathbf{c} belongs to $C_{\text{NNS}}(\mathbf{A})$.*

Our proof of the theorem (Sect. 4.2) uses convex duality in a central way, and provides intuition for why the theorem’s assumptions are needed. Section 6 provides counter-examples to relaxations of Theorem 9 obtained through weaker hypothesis.

This section’s next theorem (proven in Sect. 4.3) concerns conditions on \mathbf{A} for when the SAGE and nonnegativity cones can be expressed as a Cartesian product of simpler sets. To aid in exposition, we introduce a definition: a matrix \mathbf{A} can be *partitioned into k faces* if by a permutation of its columns it can be written as a concatenation $\mathbf{A} = [\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(k)}]$, where $\mathbf{A}^{(i)}$ are submatrices of \mathbf{A} and $\{\mathcal{P}(\mathbf{A}^{(i)})\}_{i=1}^k$ are mutually disjoint faces of $\mathcal{P}(\mathbf{A})$.

Theorem 10 *If $\{\mathbf{A}^{(i)}\}_{i=1}^k$ are matrices partitioning the block matrix $\mathbf{A} = [\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(k)}]$, then $C_{\text{NNS}}(\mathbf{A}) = \bigoplus_{i=1}^k C_{\text{NNS}}(\mathbf{A}^{(i)})$ and $C_{\text{SAGE}}(\mathbf{A}) = \bigoplus_{i=1}^k C_{\text{SAGE}}(\mathbf{A}^{(i)})$.*

The following figure illustrates partitioning a matrix \mathbf{A} where $\text{ext } \mathcal{P}(\mathbf{A})$ are the vertices of the truncated icosahedron, and nonextremal terms (marked in red) lay in the relative interiors of certain pentagonal faces.



Note that every matrix A admits the trivial partition with $k = 1$. In fact, a natural regularity condition (one that we consider in Section 6) would be that A *only* admits the trivial partition. Regularity conditions aside, Theorems 9 and 10 can be combined with known properties of AGE functions to establish new conditions for when the SAGE and nonnegativity cones coincide.

Theorem 11 *Suppose A can be partitioned into faces where (1) simplicial faces contain at most two nonextremal exponents, and (2) all other faces contain at most one nonextremal exponent. Then $C_{SAGE}(A) = C_{NNS}(A)$.*

Proof Let A satisfy the assumptions of Theorem 11 with associated faces $\{F_i\}_{i=1}^k$ and column blocks $A^{(i)}$, and fix c in $C_{NNS}(A)$. For i in $[k]$, define the vector $c^{(i)}$ so that $c = \bigoplus_{i=1}^k c^{(i)}$. By Theorem 10, the condition $C_{SAGE}(A) = C_{NNS}(A)$ holds if and only if $C_{SAGE}(A^{(i)}) = C_{NNS}(A^{(i)})$ for all i in $[k]$. Because we assumed that c belongs to $C_{NNS}(A)$ it suffices to show that each $c^{(i)}$ belongs to $C_{SAGE}(A^{(i)})$.

Per Proposition 1, any vector $c^{(i)} \in C_{NNS}(A^{(i)})$ cannot have a negative entry $c_j^{(i)}$ when $a_j^{(i)}$ is extremal in $\mathcal{P}(A^{(i)})$. By assumption, $A^{(i)}$ has at most two nonextremal terms, and so $c^{(i)} \in C_{NNS}(A^{(i)})$ can have at most two negative entries. If $c^{(i)}$ has at most one negative entry, then $c^{(i)}$ is an AGE vector. If on the other hand $c^{(i)}$ has two negative entries $c_j^{(i)}$, then both of these entries must correspond to nonextremal a_j , and F_i must be simplicial. This allows us to invoke Theorem 9 on $c^{(i)}$ to conclude $c^{(i)} \in C_{SAGE}(A^{(i)})$. The result follows. \square

4.2 Simplicial Sign Patterns for SAGE Versus Nonnegativity (Theorem 9)

The proof of Theorem 9 begins by exploiting two key facts about signomials and SAGE relaxations: (1) that $C_{SAGE}(A)$ and $C_{NNS}(A)$ are invariant under translation of the exponent set A , and (2) that strong duality *always* holds when computing f_{SAGE} . These properties allow us to reduce the problem of checking SAGE decomposability to the problem of exactness of a convex relaxation for a signomial optimization problem.

Proof (Theorem 9) Begin by translating A in $\mathbb{R}^{n \times m}$ to $A \leftarrow A - a_j \mathbf{1}^T$ where a_j is an arbitrary extremal element of $\mathcal{P}(A)$. Next, permute the columns of A so that $a_1 = \mathbf{0}$. Fix c in $C_{NNS}(A)$ and define $f = \text{Sig}(A, c)$ so that $f^* \geq 0$. We show that $f_{SAGE} = f^*$, thereby establishing $c \in C_{SAGE}(A)$.

Let $N = \{i : c_i \leq 0\}$ and $E = [m] \setminus N$; apply Corollary 4 with Proposition 2 to obtain

$$f_{SAGE} = \inf c^T v$$

s.t. v in \mathbb{R}_{++}^m has $v_1 = 1$, and there exist $\{\mu_i\}_{i \in N \cup \{1\}} \subset \mathbb{R}^n$ with

$$v_i \ln(v_i/v_j) \leq (a_i - a_j)^T \mu_i \text{ for } j \text{ in } E \text{ and } i \text{ in } N \cup \{1\}. \quad (11)$$

In order to show $f_{SAGE} = f^*$, we reformulate (11) as the problem of computing f^* by appropriate changes of variables and constraints.

We begin with a change of constraints. By the assumption that $c_i \leq 0$ for all nonextremal \mathbf{a}_i , the set E satisfies $\{\mathbf{a}_i\}_{i \in E} \subset \text{ext } \mathcal{P}(\mathbf{A})$. Combine this with extremality of $\mathbf{0} = \mathbf{a}_1$ and the assumption that $\text{ext } \mathcal{P}(\mathbf{A})$ is simplicial to conclude that $\{\mathbf{a}_i : i \in E \setminus \{1\}\}$ are linearly independent. The linear independence of these vectors ensures that for fixed \mathbf{v} we can always choose $\boldsymbol{\mu}_1$ to satisfy the following constraints *with equality*

$$v_1 \ln(v_1/v_j) \leq (\mathbf{a}_1 - \mathbf{a}_j)^\top \boldsymbol{\mu}_1 \text{ for all } j \text{ in } E.$$

Therefore we can equivalently reformulate f_{SAGE} as

$$\begin{aligned} f_{\text{SAGE}} &= \inf \mathbf{c}^\top \mathbf{v} \\ \text{s.t. } \mathbf{v} &\text{ in } \mathbb{R}_{++}^m \text{ has } v_1 = 1, \text{ and there exist } \{\boldsymbol{\mu}_i\}_{i \in N \cup \{1\}} \subset \mathbb{R}^n \\ &\text{ with } \ln(v_j) = \mathbf{a}_j^\top \boldsymbol{\mu}_1 \text{ for all } j \text{ in } E, \text{ and} \\ &v_i \ln(v_i/v_j) \leq (\mathbf{a}_i - \mathbf{a}_j)^\top \boldsymbol{\mu}_i \text{ for } j \text{ in } E, i \text{ in } N. \end{aligned}$$

Next we rewrite the constraint $v_i \ln(v_i/v_j) \leq (\mathbf{a}_i - \mathbf{a}_j)^\top \boldsymbol{\mu}_i$ as $\ln(v_i) - \ln(v_j) \leq (\mathbf{a}_i - \mathbf{a}_j)^\top \boldsymbol{\mu}_i$ by absorbing v_i into $\boldsymbol{\mu}_i$. If we also substitute the expression for $\ln(v_j)$ given by the equality constraints, then the inequality constraints become

$$\ln(v_i) \leq \mathbf{a}_i^\top \boldsymbol{\mu}_i + \mathbf{a}_j^\top (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_i) \text{ for all } j \text{ in } E, i \text{ in } N. \tag{12}$$

We now show that for every i in N , the choice $\boldsymbol{\mu}_i = \boldsymbol{\mu}_1$ makes these inequality constraints as loose as possible.

Toward this end, define $\psi_i(\mathbf{x}) = \mathbf{a}_i^\top \mathbf{x} + \min_{j \in E} \{\mathbf{a}_j^\top (\boldsymbol{\mu}_1 - \mathbf{x})\}$; note that for fixed i and $\boldsymbol{\mu}_i$, the number $\psi_i(\boldsymbol{\mu}_i)$ is the minimum over all $|E|$ right hand sides in (12). It is easy to verify that ψ_i is concave, and because of this we know that ψ_i is maximized at \mathbf{x}^* if and only if $\mathbf{0} \in (\partial \psi_i)(\mathbf{x}^*)$. Standard subgradient calculus tells us that $(\partial \psi_i)(\mathbf{x})$ is precisely the convex hull of vectors $\mathbf{a}_i - \mathbf{a}_k$ where k is an index at which the minimum (over $j \in E$) is obtained. Therefore $(\partial \psi_i)(\boldsymbol{\mu}_1) = \text{conv}\{\mathbf{a}_i - \mathbf{a}_j : j \in E\}$, and this set must contain the zero vector (unless perhaps $c_i = 0$, in which case the constraints on v_i are inconsequential). Hence $\max_{\mathbf{x} \in \mathbb{R}^n} \{\psi_i(\mathbf{x})\} = \mathbf{a}_i^\top \boldsymbol{\mu}_1$, and so inequality constraints (12) reduce to

$$\ln(v_i) \leq \mathbf{a}_i^\top \boldsymbol{\mu}_1 \text{ for all } i \text{ in } N. \tag{13}$$

Since the objective $\mathbf{c}^\top \mathbf{v}$ is decreasing in v_i for i in N , we can actually take the constraints in (13) to be binding. We established much earlier that $v_i = \exp \mathbf{a}_i^\top \boldsymbol{\mu}_1$ for i in E . Taking these together we see $v_i = \exp \mathbf{a}_i^\top \boldsymbol{\mu}_1$ for all i , and so

$$f_{\text{SAGE}} = \inf \{ \sum_{i=1}^m c_i \exp \mathbf{a}_i^\top \boldsymbol{\mu}_1 : \boldsymbol{\mu}_1 \text{ in } \mathbb{R}^n \} = f^* \tag{14}$$

as required. □

Let us now recap how the assumptions of Theorem 9 were used at various stages in the proof. For one thing, all discussion up to and including the statement of Problem (11) was fully general; the expression for f_{SAGE} used none of the assumptions of the theorem. The next step was to use linear independence of nonzero extreme points to

allow us to satisfy $v_1 \ln(v_1/v_j) \leq (\mathbf{a}_1 - \mathbf{a}_j)^\top \boldsymbol{\mu}_1$ with equality. The reader can verify that if we did not have linear independence but we were told that those constraints were binding at the optimal \mathbf{v}^* , then we would still have $f_{\text{SAGE}} = f^*$ under the stated sign pattern assumption on \mathbf{c} . Note how the sign pattern assumption on \mathbf{c} was only really used to replace $\ln(v_i) \leq \mathbf{a}_i^\top \boldsymbol{\mu}_1$ from (13) by $\ln(v_i) = \mathbf{a}_i^\top \boldsymbol{\mu}_1$.

4.3 Proof of the Partitioning Theorem (Theorem 10)

The following lemma adapts claim (iv) from Theorem 3.6 of Reznick [42] to signomials. Because the lemma is important for our subsequent theorems, the appendix contains a more complete proof than can be found in Reznick’s [42]. As a matter of notation: for any $F \subset \{\mathbf{a}_i\}_{i=1}^m$, write $\text{Sig}_F(\mathbf{A}, \mathbf{c})$ to mean the signomial with exponents \mathbf{a}_i in F and corresponding coefficients c_i .

Lemma 12 *If F is a face of $\mathcal{P}(\mathbf{A})$ then $\text{Sig}_F(\mathbf{A}, \mathbf{c})^* < 0$ implies $\text{Sig}(\mathbf{A}, \mathbf{c})^* < 0$.*

Proof (Theorem 10) Let \mathbf{A} have partition $\mathbf{A} = [\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(k)}]$, where the submatrices $\mathbf{A}^{(i)}$ have sizes $n \times m_i$ and $\sum_{i=1}^k m_i = m$. It is clear from the definition of the SAGE cone that $\mathbf{C}_{\text{SAGE}}(\mathbf{A}) = \bigoplus_{i=1}^k \mathbf{C}_{\text{SAGE}}(\mathbf{A}^{(i)})$. The bulk of this proof is to show that $\mathbf{C}_{\text{NNS}}(\mathbf{A})$ admits the same decomposition.

Let $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ for some \mathbf{c} in \mathbb{R}^m . The vector \mathbf{c} is naturally decomposed into $\mathbf{c} = \bigoplus_{i=1}^k \mathbf{c}^{(i)}$ where $\mathbf{c}^{(i)} \in \mathbb{R}^{m_i}$ align with $\mathbf{A}^{(i)}$. For each i in $[k]$ define $f_i = \text{Sig}(\mathbf{A}^{(i)}, \mathbf{c}^{(i)})$ so that $f = \sum_{i=1}^k f_i$. If any f_i^* is negative, then Lemma 12 tells us that f^* must also be negative. Meanwhile if all f_i^* are nonnegative, then the same must be true of $f^* \geq \sum_{i=1}^k f_i^*$. The result follows. □

4.4 Corollaries for Signomial Programming

Signomial minimization is naturally related via duality to checking signomial non-negativity. Thus we build on groundwork laid in Sects. 3 and 4 to obtain consequences for signomial minimization.

Corollary 13 *Assume $\mathcal{P}(\mathbf{A})$ is simplicial with $\mathbf{a}_1 = \mathbf{0}$, and that nonzero nonextremal \mathbf{a}_i have $c_i \leq 0$. Then either $f_{\text{SAGE}} = f^*$, or $f^* \in (f_{\text{SAGE}}, c_1)$.*

Proof It suffices to show that $f_{\text{SAGE}} < f^*$ implies $f^* < c_1$. This follows as the contrapositive of the following statement: “If $f^* \geq c_1$, then by Theorem 9 the nonnegative signomial $f - f^*$ is SAGE, which in turn ensures $f_{\text{SAGE}} = f^*$.” □

Now we consider constrained signomial programs. Starting with problem data (f, g) where $f = \text{Sig}(\mathbf{A}, \mathbf{c})$, $g_j = \text{Sig}(\mathbf{A}, \mathbf{g}_j)$ for j in $[k]$, $\mathbf{a}_1 = \mathbf{0}$, consider the problem of computing

$$(f, g)^* \doteq \inf\{f(\mathbf{x}) : \mathbf{x} \text{ in } \mathbb{R}^n \text{ satisfies } g(\mathbf{x}) \geq \mathbf{0}\}. \tag{15}$$

It is evident¹ that we can relax the problem to that of

$$(f, g)_{\text{SAGE}} \doteq \inf \{ \mathbf{c}^\top \mathbf{v} : \mathbf{v} \text{ in } \mathcal{C}_{\text{SAGE}}(\mathbf{A})^\dagger \text{ satisfies } v_1 = 1 \text{ and } \mathbf{G}^\top \mathbf{v} \geq \mathbf{0} \} \leq (f, g)^*$$

where \mathbf{G} is the $m \times k$ matrix whose columns are the \mathbf{g}_j .

Corollary 14 *Suppose $\mathcal{P}(\mathbf{A})$ is simplicial with vertex $\mathbf{a}_1 = \mathbf{0}$, and that when \mathbf{a}_i is nonextremal we have (i) $\mathbf{c}^\top \mathbf{v}$ is decreasing in v_i , and (ii) each $\mathbf{g}_j^\top \mathbf{v}$ is increasing in v_i . Then $(f, g)_{\text{SAGE}} = (f, g)^*$.*

Proof (sketch) The claim that $(f, g)_{\text{SAGE}} = (f, g)^*$ can be established by a change-of-variables and change-of-constraints argument of the same kind used in the proof of Theorem 9.

Suffice it to say that rather than using Corollary 4 to justify removing constraints from the dual without loss of generality, one can simply throw out those constraints to obtain some $(f, g)'$ with $(f, g)' \leq (f, g)_{\text{SAGE}}$. One then shows $(f, g)' = (f, g)^*$ to sandwich $(f, g)^* \leq (f, g)' \leq (f, g)_{\text{SAGE}} \leq (f, g)^*$. \square

4.5 Finite Error in SAGE Relaxations

This section’s final theorem directly considers SAGE as a relaxation scheme for signomial minimization. It exploits the primal formulation for f_{SAGE} to establish sufficient conditions under which SAGE relaxations can only exhibit finite error.

Theorem 15 *Suppose $\mathbf{0} \in \mathcal{P}(\mathbf{A})$ and there exists an $\epsilon > 0$ so that $(1 + \epsilon)\mathbf{a}_j$ belongs to $\mathcal{P}(\mathbf{A})$ for all nonextremal \mathbf{a}_j . Then $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ is bounded below if and only if f_{SAGE} is finite.*

The requirements Theorem 15 imposes on the Newton polytope are significantly weaker than those found elsewhere in this work. Theorem 15 is especially notable as we do not know of analogous theorems in the literature on SOS relaxations for polynomial optimization.

Proof (Theorem 15) Let $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ have $\mathbf{a}_1 = \mathbf{0}$ and $f^* > -\infty$. We may assume without loss of generality that $c_1 = 0$. Use $E = \{i : \mathbf{a}_i \text{ nonzero, extremal}\}$ to denote indices of extremal exponents of f , excluding the possibly extremal exponent $\mathbf{a}_1 = \mathbf{0}$. The desired claim holds if there exists a positive constant γ so that the translate $f_\gamma = f + \gamma$ is SAGE.

Define $\hat{\mathbf{c}} = \mathbf{c} + \gamma \mathbf{e}_1$ as the coefficient vector of f_γ . Because $f^* > -\infty$ we have $c_i = \hat{c}_i \geq 0$ for every $i \in E$ (Proposition 1). Let N denote the set of indices i for which $\hat{c}_i < 0$. For each such index $i \in N$ we define the vector $\hat{\mathbf{c}}^{(i)}$ in \mathbb{R}^m by

$$\hat{c}_j^{(i)} = \begin{cases} \hat{c}_i & \text{if } j = i \\ \hat{c}_j / |N| & \text{if } j \in [m] \setminus N \\ 0 & \text{if } j \in N \setminus \{i\} \end{cases}$$

¹ See Section 3.4 of [7].

Certainly, $\sum_{i \in N} \hat{c}^{(i)} = \hat{c}$ and $\hat{c}_{\setminus i}^{(i)}$ is nonnegative (in particular $\hat{c}_1^{(i)} = \gamma/|N|$ is positive).

Now we build the vectors $\mathbf{v}^{(i)}$ for i in N . Because N is contained within $[m] \setminus E$, we have that each i in N satisfies $(1 + \epsilon)\mathbf{a}_i$ in $\mathcal{P}(\mathbf{A})$ for some positive ϵ . Therefore for i in N , the vector \mathbf{a}_i is expressible as a convex combination of extremal exponents and the zero vector. Let $(\lambda_j^{(i)})_{j \in E \cup \{1\}}$ be positive convex combination coefficients so that $\mathbf{a}_i = \sum_{j \in E \cup \{1\}} \lambda_j^{(i)} \mathbf{a}_j$.

Now define the vector $\mathbf{v}^{(i)}$ in $\mathbb{R}^{[m] \setminus i}$ by $v_j^{(i)} = \lambda_j^{(i)}$ for j in $E \cup \{1\}$, and $v_j^{(i)} = 0$ for all remaining indices. Each $\mathbf{v}^{(i)}$ is nonnegative, satisfies $v_1^{(i)} > 0$, and belongs to the kernel of $[\mathbf{A}_{\setminus i} - \mathbf{a}_i \mathbf{1}^\top]$. Because $v_1^{(i)}$ is positive, the quantity $D(\mathbf{v}^{(i)}, \hat{c}_{\setminus i}^{(i)})$ can be made to diverge to $-\infty$ by sending γ to ∞ . It follows that there exists a sufficiently large M so that $\gamma \geq M$ implies

$$D(\mathbf{v}^{(i)}, e\hat{c}_{\setminus i}^{(i)}) - \hat{c}_i \leq 0 \quad \text{for all } i \text{ in } N. \tag{16}$$

Hence for sufficiently large γ , we have $\hat{c}^{(i)}$ in $C_{\text{AGE}}(\mathbf{A}, i)$ for all i in N —and the result follows. □

5 Certifying Global Nonnegativity of Polynomials

Throughout this section we write $p = \text{Pol}(\mathbf{A}, \mathbf{c})$ to mean that p takes values $p(\mathbf{x}) = \sum_{i=1}^m c_i \mathbf{x}^{\mathbf{a}_i}$. We refer to polynomials in this way to reflect our interest in *sparse polynomials*. Vectors \mathbf{a}_i are sometimes called *terms*, where a term is *even* if \mathbf{a}_i belongs to $(2\mathbb{N})^n$. To an n -by- m matrix of nonnegative integers \mathbf{A} , we associate the sparse nonnegativity cone

$$C_{\text{NNP}}(\mathbf{A}) \doteq \{\mathbf{c} : \text{Pol}(\mathbf{A}, \mathbf{c})(\mathbf{x}) \geq 0 \text{ for all } \mathbf{x} \text{ in } \mathbb{R}^n\}.$$

Beginning with Sect. 5.1 we introduce *polynomial SAGE certificates*. We shall see that polynomial SAGE certificates offer a tractable avenue for optimizing over a subset of $C_{\text{NNP}}(\mathbf{A})$, where the complexity depends on \mathbf{A} exclusively through the dimensions n and m .

Section 5.2 demonstrates how our study of SAGE signomials yields several corollaries in this new polynomial setting. Perhaps most prominently, Sect. 5.2 implies that a polynomial admits a SAGE certificate if and only if it admits a SONC certificate. The qualitative relationship between SAGE and SONC as proof systems is explained in Sects. 5.3, and 5.4 addresses how some of our corollaries compare to earlier results in the SONC literature.

In Sect. 5.5, we compare polynomial SAGE certificates to the widely studied Sums-of-Squares certificates. We conclude with Sect. 5.6, which outlines how to use SAGE polynomials to obtain a hierarchy for constrained polynomial optimization.

5.1 Signomial Representatives and Polynomial SAGE Certificates

To a polynomial $p = \text{Pol}(A, c)$ we associate the *signomial representative* $q = \text{Sig}(A, \hat{c})$ with

$$\hat{c}_i = \begin{cases} c_i & \text{if } a_i \text{ is even} \\ -|c_i| & \text{otherwise} \end{cases} \tag{17}$$

By a termwise argument, we have that if the signomial q is nonnegative on \mathbb{R}^n , then the polynomial p must also be nonnegative on \mathbb{R}^n . Moving from a polynomial to its signomial representative often entails some loss of generality. For example, the univariate polynomial $p(x) = 1 + x - x^3 + x^4$ never has both “ $+x < 0$ ” and “ $-x^3 < 0$,” and yet the inner terms appearing in the signomial representative $q(y) = 1 - \exp(y) - \exp(3y) + \exp(4y)$ are both negative.

There is a natural condition A and the sign pattern of c where passing to the signomial representative is at no loss of generality. Specifically, if there exists a point $x_0 \in (\mathbb{R} \setminus \{0\})^n$ where $c_i x_0^{a_i} \leq 0$ for all $a_i \notin (2\mathbb{N})^n$, then $\text{Pol}(A, c)$ is nonnegative if and only if its signomial representative is nonnegative. We call such polynomials *orthant-dominated*. Checking if a polynomial is orthant-dominated is a simple task. Given A and c , define b by $b_i = 0$ if $c_i \leq 0$ or a_i is even, and $b_i = 1$ if otherwise. Then assuming every $c_i \neq 0$, the polynomial $\text{Pol}(A, c)$ is orthant-dominated if and only if the system $A^T s = b \pmod{2}$ has a solution over $s \in \mathbb{F}_2^n$.

In what should feel natural, we call $p = \text{Pol}(A, c)$ a *SAGE polynomial* if its signomial representative $q = \text{Sig}(A, \hat{c})$ is a SAGE signomial. Subsequently, we define a *polynomial SAGE certificate* for $p = \text{Pol}(A, c)$ as a set of signomial AGE certificates $\{(\hat{c}^{(i)}, v^{(i)})\}_{i=1}^m$ where $\hat{c} \doteq \sum_{i=1}^m \hat{c}^{(i)}$ defines the signomial representative for p . Because the signomial SAGE cone contains the nonnegative orthant, the cone of coefficients for SAGE polynomials admits the representation

$$\begin{aligned} C_{\text{SAGE}}^{\text{POLY}}(A) = \{c : \text{there exists } \hat{c} \in C_{\text{SAGE}}(A) \text{ where } \hat{c} \leq c \\ \text{and } \hat{c}_i \leq -c_i \text{ for all } i \text{ with } a_i \text{ not in } (2\mathbb{N})^n\}. \end{aligned} \tag{18}$$

We use this representation to obtain the following theorem.

Theorem 16 *Let $L : \mathbb{R}^\ell \rightarrow \mathbb{R}^m$ be an injective affine map, $A \in \mathbb{N}^{n \times m}$ be a matrix of exponents ($n \leq m$), and h be a vector in \mathbb{R}^ℓ . An ϵ -approximate solution to*

$$\inf_{z \in \mathbb{R}^\ell} \{h^T z : L(z) \in C_{\text{SAGE}}^{\text{POLY}}(A)\} \tag{19}$$

can be computed in time $O(p(m) \log(1/\epsilon))$ for a polynomial p .

Proof We appeal to standard results on interior point methods (IPMs) for conic programming. The task is to show that $C_{\text{SAGE}}^{\text{POLY}}(A)$ can be expressed as a projection of a convex cone “ K ,” which possesses a tractable self-concordant barrier with a complexity parameter ϑ bounded by a polynomial in m . From there, the meaning of “ ϵ -approximate” and its relationship to the polynomial “ p ” depends highly on the details of a given IPM; relevant sources for general conic IPMs include [33, §4] and

[53, §5]. In particular we rely on algorithms for optimizing over the *exponential cone* $K_{\text{exp}} = \text{cl}\{(u, v, w) : v \exp(u/v) \leq w, v > 0\}$, and defer to [38,47,52] for formal meanings of “ ϵ -approximate” in our context.

For each $i \in [m]$, let M_i denote a matrix with “ m_i ” columns spanning $\ker(A_{\setminus i} - a_i \mathbf{1}^\top) \subset \mathbb{R}^{[m] \setminus i}$ and define the cone $K_i = \{(u, v, t) : u, v \in \mathbb{R}_+^{[m] \setminus i}, D(u, ev) \leq t\}$. In terms of M_i and K_i we can reformulate the i^{th} signomial AGE cone as

$$\{c^{(i)} : \text{some } w^{(i)} \in \mathbb{R}^{m_i} \text{ satisfies } (M_i w^{(i)}, \hat{c}_{\setminus i}^{(i)}, \hat{c}_i^{(i)}) \in K_i\}.$$

Since K_i can be represented with $m - 1$ copies of K_{exp} and one linear inequality over $m - 1$ additional scalar variables, the preceding display tells us that $C_{\text{SAGE}}(\mathbf{A})$ can be represented with $m(m - 1)$ copies of K_{exp} , m linear inequalities, and $O(m^2)$ scalar auxiliary variables. Combine this with the representation (18) to find that the feasible set for (19) can be described with $O(m^2)$ exponential cone constraints, $O(m)$ linear inequalities, and $O(\ell + m^2) \in O(m^2)$ scalar variables. As the exponential cone has a tractable self-concordant barrier with complexity parameter $\vartheta_{\text{exp}} = 3$, $C_{\text{SAGE}}^{\text{POLY}}(\mathbf{A})$ has a tractable self-concordant barrier with complexity parameter $O(m^2)$. \square

5.2 Simple Consequences of Our Signomial Results

Section 5.1 suggested that the signomial SAGE cone is more fundamental than the polynomial SAGE cone. This section serves to emphasize that idea, by showing how our study of the signomial SAGE cone quickly produces results in the polynomial setting. The following corollaries are obtained by viewing Theorems 9 and 11 through the lens of orthant-dominance.

Corollary 17 *If exponent vectors \mathbf{A} induce a simplicial Newton polytope $\mathcal{P}(\mathbf{A})$, and nonextremal exponents are linearly independent mod 2, then $C_{\text{SAGE}}^{\text{POLY}}(\mathbf{A}) = C_{\text{NNP}}(\mathbf{A})$.*

Corollary 18 *Suppose \mathbf{A} belonging to $p = \text{Pol}(\mathbf{A}, \mathbf{c})$ can be partitioned into faces where (1) each simplicial face induces an orthant-dominated polynomial with at most two nonextremal terms, and (2) all other faces have at most one nonextremal term. Then p is nonnegative if and only if it is SAGE.*

Unfortunately, it is not possible to reduce the dependence of Corollary 18 on the coefficient vector \mathbf{c} of the polynomial p . The obstruction is that taking a signomial representative is not without loss of generality, as the case $\mathbf{A} = [0, 1, 3, 4]$ shows.

To more deeply understand the polynomial SAGE cone it is necessary to study its extreme rays, as well as its sparsity preservation properties. We now show how this can be done by leveraging Theorems 3 and 5 from Sect. 3.

Theorem 19 *Defining the cone of “AGE polynomials” for exponents \mathbf{A} and index k as*

$$C_{\text{AGE}}^{\text{POLY}}(\mathbf{A}, k) \doteq \{c : \text{Pol}(\mathbf{A}, \mathbf{c}) \text{ is globally nonnegative, and } c_{\setminus k} \geq \mathbf{0}, c_i = 0 \text{ for all } i \neq k \text{ with } \mathbf{a}_i \notin (2\mathbb{N})^n\}, \quad (20)$$

we have $\sum_{k=1}^m C_{AGE}^{POLY}(A, k) = C_{SAGE}^{POLY}(A)$.

Proof The inclusion $C_{AGE}^{POLY}(A, k) \subset C_{SAGE}^{POLY}(A)$ is obvious, since polynomials satisfying (20) have AGE signomial representatives. We must show the reverse inclusion $C_{SAGE}^{POLY}(A) \subset \sum_{k=1}^m C_{AGE}^{POLY}(A, k)$.

Given a polynomial $p = \text{Pol}(A, c)$, testing if c belongs to $C_{SAGE}^{POLY}(A)$ will reduce to testing if \hat{c} (given by Equation (17)) belongs to $C_{SAGE}(A)$. Henceforth let $\hat{c} \in C_{SAGE}(A)$ be fixed and set $N = \{i : \hat{c}_i < 0\}$. By Theorem 3, there exist vectors $\{\hat{c}^{(i)} \in C_{AGE}(A, i)\}_{i \in N}$ where $\hat{c}_i^{(i)} = \hat{c}_i < 0$ for each i and $\hat{c}_j^{(i)} = 0$ for all $j \in N \setminus \{i\}$. The sign patterns here are important: $\hat{c}^{(i)}$ is supported on the index set $\{i\} \cup ([m] \setminus N)$, and $\hat{c}_j^{(i)} \geq 0$ for all j in $[m] \setminus N$. By construction of \hat{c} , any index j in $[m] \setminus N$ corresponds to an exponent vector a_j in $(2\mathbb{N})^n$. Therefore the carefully chosen vectors $\{\hat{c}^{(i)}\}_{i \in N}$ define not only AGE signomials, but also AGE polynomials $\hat{p}_i = \text{Pol}(A, \hat{c}^{(i)})$. Lastly, for each index $i \in N$ set $c^{(i)}$ by $c_{\setminus i}^{(i)} = \hat{c}_{\setminus i}^{(i)}$, and $c_i^{(i)} = -1 \cdot \text{sign}(c_i) \cdot \hat{c}_i^{(i)}$. The resulting polynomials $p_i = \text{Pol}(A, c^{(i)})$ inherit the AGE property from \hat{p}_i , and sum to p . As we have decomposed our SAGE polynomial into an appropriate sum of “AGE polynomials,” the proof is complete. \square

Corollary 20 Any SAGE polynomial can be decomposed into a sum of AGE polynomials in a manner that is cancelation-free.

Proof The cancelation-free decomposition is given constructively in the proof of Theorem 19. \square

Corollary 21 If $c \in \mathbb{R}^m$ generates an extreme ray of $C_{SAGE}^{POLY}(A)$, then $\{a_i : i \in [m], c_i \neq 0\}$ is either a singleton or a simplicial circuit.

Proof In view of Theorem 19, it suffices to show that for fixed k the extreme rays of $C_{AGE}^{POLY}(A, k)$ are supported on single coordinates, or simplicial circuits. This follows from Theorem 5, since vectors in $C_{AGE}^{POLY}(A, k)$ are—up to a sign change on their k^{th} component—in 1-to-1 correspondence with vectors in $C_{AGE}(\hat{A}, k)$, where \hat{A} is obtained by dropping suitable columns from A . \square

5.3 AM/GM Proofs of Nonnegativity, Circuits, and SAGE

In 1989, Reznick defined an *agiform* as any positive multiple of a homogeneous polynomial $f = \text{Pol}([A, \beta], [\lambda, -1]^T)$, where $A \in (2\mathbb{N})^{n \times m}$ and $\beta = A\lambda$ for a weighting vector $\lambda \in \Delta_m$ [42]. Agiforms have AGE signomial representatives, which follows by plugging $v = \lambda$ into (2). Reznick’s investigation concerned extremality in the cone of nonnegative polynomials and identified a specific subset of simplicial agiforms which met the extremality criterion [42, Theorem 7.1].

Agiform-like functions were later studied by Pantea, Koepl, and Craciun for analysis of biochemical reaction networks [37, Proposition 3]. Pantea et al. spoke in terms of *posynomials* $f(x) = \sum_{i=1}^{k+1} c_i x^{a_i}$ where all $c_i \geq 0$; a posynomial f was said to *dominate* the monomial x^β if $x \mapsto f(x) - x^\beta$ was nonnegative on \mathbb{R}_+^n . If we adopt


the notation where $\Theta(\mathbf{c}, \boldsymbol{\lambda}) = \prod_{i=1}^{k+1} (c_i/\lambda_i)^{\lambda_i}$, [37, Theorem 3.6] says that in the case of a simplicial Newton polytope (i.e., $k = n$), monomial domination is equivalent to $1 \leq \Theta(\mathbf{c}, \boldsymbol{\lambda})$ where $\boldsymbol{\lambda}$ gives the barycentric coordinates for $\boldsymbol{\beta} \in \mathcal{P}(\mathbf{A})$.

A few years following Pantea et al., Iliman and de Wolff suggested taking sums of nonnegative circuit polynomials, which are globally nonnegative polynomials $f = \text{Pol}([\mathbf{A}, \boldsymbol{\beta}], [\mathbf{c}, b]^\top)$ where $\{\mathbf{a}_i\}_{i=1}^{n+1} \cup \{\boldsymbol{\beta}\}$ form a simplicial circuit [17]. Iliman and de Wolff’s Theorem 1.1 states that if all \mathbf{a}_i are even, f is a circuit polynomial, and $\boldsymbol{\beta} \in \mathcal{P}(\mathbf{A})$ has barycentric coordinates $\boldsymbol{\lambda} \in \Delta_{n+1}$, then f nonnegative if and only if

$$\text{either } |b| \leq \Theta(\mathbf{c}, \boldsymbol{\lambda}) \text{ and } \boldsymbol{\beta} \notin (2\mathbb{N})^n \text{ or } -b \leq \Theta(\mathbf{c}, \boldsymbol{\lambda}) \text{ and } \boldsymbol{\beta} \in (2\mathbb{N})^n. \quad (21)$$

It is clear that [17, Theorem 1.1] extends [37, Theorem 3.6], to account for sign changes of $b \cdot \mathbf{x}^\beta$ and to impose no scaling on $|b|$.

The approach of taking sums of nonnegative circuit polynomials is now broadly known as ‘‘SONC.’’ Prior formulations for the SONC cone work by enumerating every simplicial circuit which could possibly be of use in a SONC decomposition (see [11, §5.2], and subsequently [55,56]). The circuit enumeration approach is extremely inefficient, as Example 22 shows an m -term polynomial can contain as many as $2^{(m-1)/2}$ simplicial circuits.



Example 22 Let d be divisible by 2 and n . Construct an $n \times 2n$ matrix \mathbf{A} by setting \mathbf{a}_{2i-1} and \mathbf{a}_{2i} to distinct points in $\mathbb{N}^n \cap d\Delta_n$ adjacent to $d\mathbf{e}_i$. Then for large enough d , $\boldsymbol{\beta} = d\mathbf{1}/n$ will be contained in exactly 2^n simplices. Above: $(d, n) = (12, 3)$, and a projection of $(d, n) = (16, 4)$.

Circuit enumeration is not merely a theoretical issue. When using the heuristic circuit-selection technique from [45], Seidler and de Wolff’s POEM software package fails to certify nonnegativity of the AGE polynomial $f(x, y) = (x - y)^2 + x^2y^2$ and moreover only returns a bound $f^* \geq -1$ [46].

Of course- Corollary 21 tells us that a polynomial admits a SAGE certificate if and only if it admits a SONC certificate. This is good news, since Theorem 16 says we can optimize over this set in time depending polynomially on m . In particular, we may avoid SONC’s severe problems of circuit enumeration and circuit selection. The qualitative distinction here is that while Pantea et al. and Iliman and de Wolff consider the weights $\boldsymbol{\lambda}$ as fixed (given by barycentric coordinates), the analogous quantity \mathbf{v} in the SAGE approach is an optimization variable. At a technical level,

the relative entropy formulation (2) affords a joint convexity whereby SAGE can search simultaneously over coefficients $\mathbf{c}^{(i)}$ and weighting vectors $\mathbf{v}^{(i)}$. As our proof of Theorem 16 points out, we can be certain that $[A_{\setminus i} - \mathbf{a}_i \mathbf{1}^\top] \mathbf{v}^{(i)} = \mathbf{0}$ holds in exact arithmetic simply by defining $\mathbf{v}^{(i)} \leftarrow \mathbf{M}_i \mathbf{w}^{(i)}$ for the indicated matrix \mathbf{M}_i .

5.4 Comparison to Existing Results in the SONC Literature

Due to the equivalence of the class of nonnegative polynomials induced by the SAGE and the SONC approaches, some of our results have parallels in the SONC literature.

Corollary 17 is not stated in the literature, though it may be deduced from [17, Corollary 7.5]. Iliman and de Wolff prove [17, Corollary 7.5] by signomializing $g(\mathbf{x}) = f(\exp \mathbf{x})$ and introducing an additional regularity condition so that $\nabla g(\mathbf{x}) = \mathbf{0}$ at exactly one $\mathbf{x} \in \mathbb{R}^n$. Our proof of Corollary 17 stems from Theorem 9, which employs a convex duality argument applicable to constrained signomial optimization problems in the manner of Corollary 14.

Wang showed that nonnegative polynomials in which at most one term $c_i \mathbf{x}^{\mathbf{a}_i}$ takes on a negative value at some $\mathbf{x} \in \mathbb{R}^n$ (either $c_i < 0$ or $\mathbf{a}_i \notin (2\mathbb{N})^n$) are SONC polynomials [55, Theorem 3.9]. This result can be combined with the definition of AGE polynomial given in Theorem 19 in order to prove a weaker form of Corollary 21, where all \mathbf{a}_i belong to $\text{ext } \mathcal{P}(A)$ or $\text{int } \mathcal{P}(A)$. We emphasize that Corollary 21 is not responsible for the major efficiency gains of SAGE from Theorem 16; the SONC formulation in [55, §5] uses $2^{(m-1)/2}$ circuits for the m -term polynomials from Example 22.

Finally, in a result that was announced contemporaneously to the original submission of the present paper, Wang showed that summands in a SONC decomposition of a polynomial $f = \text{Pol}(A, \mathbf{c})$ may have supports restricted to A without loss of generality [56, Theorem 4.2]. In light of the equivalence between the class of SAGE polynomials and of SONC polynomials, this result may be viewed as a weaker analog of our Corollary 20; specifically, [56, Theorem 4.2] shows SONC certificates are sparsity-preserving but it does not provide a cancellation-free decomposition.

The distinctions between our polynomial Corollaries 17 and 18 versus our signomial Theorems 9 and 11 make clear that polynomial results should not be conflated with signomial results. With the exception of Sect. 5, our setup and results in this paper pertain to the class of signomials, which in general can have $\mathbf{a}_i \in (\mathbb{R} \setminus \mathbb{Q})^n$. The developments in the SONC literature only consider polynomials and employ analysis techniques of an algebraic nature which rely on integrality of exponents in fundamental ways (c.f. [56, Theorem 4.2]). In contrast, our techniques are rooted in convex duality and are applicable to the broader question of certifying signomial nonnegativity.

5.5 SAGE and SOS

The SOS approach to polynomial nonnegativity considers polynomials f in n variables of degree $2d$, and attempts to express $f(\mathbf{x}) = L(\mathbf{x})^\top \mathbf{P} L(\mathbf{x})$ where \mathbf{P} is a PSD matrix and $L : \mathbb{R}^n \rightarrow \mathbb{R}^{\binom{n+d}{d}}$ is a lifting which maps \mathbf{x} to all monomials of degree at most d evaluated at \mathbf{x} [27,39,50]. The identity $f(\mathbf{x}) = L(\mathbf{x})^\top \mathbf{P} L(\mathbf{x})$ can be enforced

with linear equations on the coefficients of f and the entries of \mathbf{P} , so deciding SOS-representability reduces to a semidefinite program.

Because it is extremely challenging to solve semidefinite programs at scale, several modifications to SOS have been proposed to offer reduced complexity. Kojima et. al built on earlier work of Reznick [41] to replace the lifting “ L ” appearing in the original SOS formulation with a smaller map using fewer monomials [24]. Their techniques had meaningful use-cases, but could fail to perform any reduction in some very simple situations [24, Proposition 5.1]. Subsequently, Waki et. al introduced the *correlative sparsity* heuristic to induce structured sparsity in the matrix variable \mathbf{P} [54]. Shortly thereafter Nie and Demmel suggested replacing the standard lifting by a collection of smaller $\{L_i\}_i$, so as to express $f(\mathbf{x}) = \sum_i L_i(\mathbf{x})^\top \mathbf{P}_i L_i(\mathbf{x})$ with order $\binom{k+d}{d}$ PSD matrices \mathbf{P}_i for some $k \ll n$ [34]. Very recently, Ahmadi and Majumdar suggested one use the standard lifting together with a *scaled diagonally dominant* matrix \mathbf{P} of order $\binom{n+d}{d}$; these “SDSOS polynomials” are precisely those polynomials admitting a decomposition as a sum of binomial squares [2].

Each of these SOS-derived works suffers from a drawback that SOS decompositions may require cancelation on coefficients of summands $f_i = g_i^2$ as one recovers $f = \sum_i f_i$. As a concrete example, consider $f(x, y) = 1 - 2x^2y^2 + x^8/2 + y^8/2$; this polynomial is nonnegative (in fact, AGE) and admits a decomposition as a sum of binomial squares. The trouble is that to decompose f as a sum of binomial squares, the summands $f_i = g_i^2$ require additional terms $+x^4y^4$ and $-x^4y^4$. By contrast, SAGE certificates need only involve the original monomials in f , and one may take summand AGE polynomials to be cancelation-free with no loss of generality (Corollary 20). The SAGE approach also has the benefit of being formulated with a relative entropy program of size $O(m^2)$ (Theorem 16), while SOS-derived works have complexity scaling exponentially with a polynomial’s degree d .

We make two remarks in closing. First, it is easy to verify that every binomial square is an AGE polynomial, and so SAGE can certify nonnegativity of all SDSOS polynomials. Second, it is well known that proof systems leveraging the AM/GM inequality (SAGE among them) can certify nonnegativity of some polynomials which are *not* SOS. A prominent example here is the Motzkin form $f(x, y, z) = x^2y^4 + x^4y^2 + z^6 - 3x^2y^2z^2$.

5.6 Extending SAGE Polynomials to a Hierarchy

We conclude this section by discussing how to obtain hierarchies for constrained polynomial optimization problems, in a manner which is degree-independent and sparsity preserving. Adopt the standard form (15) for minimizing a polynomial f subject to constraint polynomials $\{g_i\}_{i=1}^k$. Here, all polynomials are over a common set of exponents $\mathbf{A} \in \mathbb{N}^{n \times m}$, with $\mathbf{a}_1 = \mathbf{0}$ and $n \leq m$. Our development is based on a hierarchy for signomials that is described in [7, §3.3].

² A symmetric matrix \mathbf{P} is scaled diagonally dominant if there exists diagonal $\mathbf{D} > \mathbf{0}$ so that $\mathbf{D}\mathbf{P}\mathbf{D}^\top$ is diagonally dominant. Such matrices can be represented as a sum of 2×2 PSD matrices with appropriate zero padding.

Consider operators \mathcal{A} and \mathcal{C} taking values $\mathcal{A}(\text{Pol}(\mathbf{A}, \mathbf{c})) = \mathbf{A}$ and $\mathcal{C}(\text{Pol}(\mathbf{A}, \mathbf{c})) = \mathbf{c}$ respectively. We shall say our SAGE polynomial hierarchy is indexed by two parameters: p and q . The parameter p controls the complexity of Lagrange multipliers; when $p = 0$, the Lagrange multipliers are simply $\lambda_i \geq 0$. For general p , the Lagrange multipliers are SAGE polynomials over exponents $\mathbf{A}' \doteq \mathcal{A}(\text{Pol}(\mathbf{A}, \mathbf{1})^p)$. The parameter q controls the number of constraints in the nonconvex primal problem: $H = \{h_i\}_{i=1}^{k^q}$ are obtained by taking all q -fold products of the g_i . Once the Lagrangian $L = f - \gamma - \sum_{h \in H} h \cdot s_h$ is formed, it will be a polynomial over exponents $\mathbf{A}'' \doteq \mathcal{A}(\text{Pol}(\mathbf{A}, \mathbf{1})^{p+q})$. By the minimax inequality we have

$$(f, g)^{(p,q)} \doteq \sup_{\gamma, \{s_h\}_{h \in H}} \{ \gamma : \mathcal{C}(L) \in \mathbb{C}_{\text{SAGE}}^{\text{POLY}}(\mathbf{A}'') \text{ and } \mathcal{C}(s_h) \in \mathbb{C}_{\text{SAGE}}^{\text{POLY}}(\mathbf{A}') \forall h \in H \} \leq (f, g)^*$$

Following Theorem 16, the above can be solved in time polynomial in m, k for each fixed p, q . As p and q increase, we obtain improved bounds at the expense of an increase in computation. Mirroring [7], one can appeal to representation theorems from the real algebraic geometry literature [25,29,44] to prove that this hierarchy can provide arbitrarily accurate lower bounds for sparse polynomial optimization problems in which the constraint set is Archimedean (for example, if all variables have explicit finite upper and lower bounds).

Our broader message here—beyond results on convergence to the optimal value of specific hierarchies—is that the above construction qualitatively differs from other hierarchies in the literature, because the optimization problems encountered at every level of our construction *depend only on the nonnegative lattice generated by the original exponent vectors \mathbf{A}* . The theoretical underpinnings of this sparsity-preserving hierarchy trace back to the decomposition result given by Theorem 3. Thus, it is possible to obtain entire families of relative entropy relaxations that are sparsity-preserving, which reinforces our message about the utility of SAGE-based relative entropy optimization for sparse polynomial problems.

6 Toward Necessary and Sufficient Conditions for SAGE versus Nonnegativity

We conclude this paper with a discussion on the extent to which our results tightly characterize the distinction between SAGE and nonnegativity for signomials. This section is split into three parts. In the first part, we describe a process for identifying cases where $\mathbb{C}_{\text{SAGE}}(\mathbf{A}) \subsetneq \mathbb{C}_{\text{NNS}}(\mathbf{A})$. This process is illustrated with several examples which suggest that our results from Sect. 4 are essentially tight. Section 6.2 presents a formal conjecture regarding the ways in which our results might be improved, and Sect. 6.3 provides a novel dual formulation for when $\mathbb{C}_{\text{SAGE}}(\mathbf{A}) = \mathbb{C}_{\text{NNS}}(\mathbf{A})$.

6.1 Constructing Examples of Non-Equality

Given a matrix of exponent vectors A , we are interested in finding a coefficient vector c so that $f = \text{Sig}(A, c)$ satisfies $f_{\text{SAGE}} < f^*$. If such c exists, then it is evident that $C_{\text{SAGE}}(A) \neq C_{\text{NNS}}(A)$.

The naïve approach to this process would be to carefully construct signomials where the infimum f^* is known by inspection, to compute f_{SAGE} , and then to test if the measured value $|f_{\text{SAGE}} - f^*|$ is larger than would be possible from rounding errors alone. A serious drawback of this approach is that it can be quite difficult to construct A and c where f^* is apparent, and yet $\{a_i : c_i \neq 0\}$ satisfy the properties for the conjecture under test.

To address this challenge, we appeal to the idea alluded to in Sect. 2 that SAGE provides a means of computing a *sequence* of lower bounds $(f_{\text{SAGE}}^{(\ell)})_{\ell \in \mathbb{N}}$. For details on this “unconstrained SAGE hierarchy,” we refer the reader to [7]. For our purposes, suffice it to say that

$$f_{\text{SAGE}}^{(\ell)} \doteq \sup\{\gamma : \text{Sig}(A, \mathbf{1})^\ell(f - \gamma) \text{ is SAGE}\}$$

defines a non-decreasing sequence bounded above by f^* . Thus, while we cannot readily check if $|f_{\text{SAGE}} - f^*| \gg 0$, we *can* compute a few values of $f_{\text{SAGE}}^{(\ell)}$ for $\ell > 0$, and check if $|f_{\text{SAGE}}^{(0)} - f_{\text{SAGE}}^{(\ell)}| \gg 0$.

The remainder of this section goes through case studies in which we probe the sensitivity our earlier theorems’ conclusions to their stated assumptions. All computation was performed with a late 2013 MacBook Pro with a 2.4GHz i5 processor, using CVXPY [1,9] as an interface to the conic solver ECOS [10,47].³ Numerical precision is reported to the farthest decimal point where the primal and dual methods for computing $f_{\text{SAGE}}^{(\ell)}$ agree.

Example 23 We test here whether it is possible to relax the assumption of simplicial Newton polytope in Theorem 9. Since every Newton polytope in \mathbb{R} is trivially simplicial, the simplest signomials available to us are over \mathbb{R}^2 . With that in mind, consider

$$A = \begin{bmatrix} 0 & 2 & 1 & 0 & 0 & 2 \\ 0 & 0 & 0 & 2 & 1 & 2 \end{bmatrix}.$$

This choice of A is particularly nice, because were it not for the last column $a_6 = [2, 2]^T$, we would very clearly have $C_{\text{SAGE}}(A) = C_{\text{NNS}}(A)$. We test tested a few values for c before finding

$$c = [0, 3, -4, 2, -2, 1]^T,$$

which resulted in $f_{\text{SAGE}}^{(0)} \approx -1.83333$, and $f_{\text{SAGE}}^{(1)} \approx -1.746505595 = f^*$. Because the absolute deviation $|f_{\text{SAGE}} - f^*| \approx 0.08682$ is much larger than the precision to

³ See github.com/rileyjmmurray/sigpy or data.caltech.edu/records/1427 for code.

which we solved these relaxations, we conclude that $C_{SAGE}(A) \neq C_{NNS}(A)$ for this choice of A .

Example 24 Let us reinforce the conclusion from Example 23. Applying a 180 degree rotation about the point (1,1) to the columns of A , we obtain

$$A = \begin{bmatrix} 0 & 2 & 0 & 2 & 1 & 2 \\ 0 & 0 & 2 & 2 & 2 & 1 \end{bmatrix}.$$

We then choose the coefficients in a manner informed by the theory developed in Sect. 6.3

$$c = [0, 1, 1, 1.9, -2, -2]^T$$

which subsequently defines $f = \text{Sig}(A, c)$. In this case the primal formulation for f_{SAGE} is infeasible, and so $f_{SAGE}^{(0)} = -\infty$. Meanwhile, the second level of the unconstrained hierarchy produces $f_{SAGE}^{(1)} \approx -0.122211863 = f^*$. Thus in a very literal sense, the gap $|f_{SAGE} - f^*|$ could not be larger.

We know from Theorem 11 that any signomial with at most four terms is nonnegative if and only if it is SAGE. It is natural to wonder if in some very restricted setting (e.g., univariate signomials) the SAGE and nonnegativity cones would coincide for signomials with five or more terms; Example 25 shows this is not true in general.

Example 25 For $f = \text{Sig}(A, c)$ with $A = [0, 1, 2, 3, 4]$ and $c = [1, -4, 7, -4, 1]^T$, we have $f_{SAGE}^{(0)} \approx -0.3333333$ and $f_{SAGE}^{(1)} \approx 0.2857720944$. Per the affine-invariance invariance properties of the SAGE and nonnegativity cones, this examples shows $C_{SAGE}(A)$ is a strict subset of $C_{NNS}(A)$ for every 1×5 matrix A with equispaced values.

Together, Examples 23 through 25 demonstrate there are meaningful senses in which Theorems 9 through 11 cannot be improved upon.

6.2 A Conjecture, Under Mild Regularity Conditions

Despite the conclusion in the previous subsection, there *are* settings when we can prove $C_{SAGE}(A) = C_{NNS}(A)$ in spite of A not satisfying the assumptions of Theorem 11. For example, one case in which SAGE equals nonnegativity is when $A = [0, I, D]$ where D is a diagonal matrix with diagonal entries in $(0, 1)$. Here one proves equality as follows: for each possible sign pattern of $c \in C_{NNS}(A)$, there exists a lower dimensional simplicial face F of $\mathcal{P}(A)$ upon which we invoke Theorem 9, and for which the remaining exponents (those outside of F) have positive coefficients. We know that the signomial induced by the exponents outside of F is trivially SAGE, and so by Theorem 10 we conclude $c \in C_{SAGE}(A)$. As this holds for all possible sign patterns on c in $C_{NNS}(A)$, we have $C_{SAGE}(A) = C_{NNS}(A)$. However, this case is somewhat degenerate, and we wish to exclude it in our discussion via some form of regularity on A .

The most natural regularity condition on A would be that it admits only the trivial partition, and indeed we focus on the case when every a_i belongs to either $\text{ext } \mathcal{P}(A)$ or $\text{int } \mathcal{P}(A)$. In this setting, we have the following corollary of Theorem 11.

Corollary 26 *If $\mathcal{P}(A)$ is full dimensional with either*

1. *at most one interior exponent, or*
2. *$n + 1$ extreme points and at most two interior exponents*

then $C_{\text{SAGE}}(A) = C_{\text{NNS}}(A)$.

Along with this corollary, we present a conjecture for the reader’s consideration.

Conjecture 27 *If $\mathcal{P}(A)$ has every a_i in either $\text{ext } \mathcal{P}(A)$ or $\text{int } \mathcal{P}(A)$, but A does not satisfy the hypothesis of Corollary 26, then $C_{\text{SAGE}}(A) \neq C_{\text{NNS}}(A)$.*

Note that when A satisfies the stated assumptions and further has some $a_i = \mathbf{0}$ in the interior, Theorem 15 ensures that $f = \text{Sig}(A, c)$ can have f_{SAGE} deviate from f^* only by a finite amount. To overcome a potential obstacle posed by this result in the resolution of Conjecture 27, one can also consider modifying the hypotheses of the conjecture to require that all a_i lie in the relative interior of the Newton polytope.

To finish discussion on Conjecture 27, we provide empirical support with the following examples.

Example 28 Let f be a signomial in two variables with

$$\begin{bmatrix} A \\ c^T \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0.30 & 0.21 & 0.16 \\ 0 & 0 & 1 & 0.58 & 0.08 & 0.54 \\ \hline 33.94 & 67.29 & 1 & 38.28 & -57.75 & -40.37 \end{bmatrix}.$$

Then $f_{\text{SAGE}} = -24.054866 < f_{\text{SAGE}}^{(1)} = -21.31651$. This example provides the minimum number of interior exponents needed to be relevant to Conjecture 27 in the simplicial case.

Example 29 Let f be a signomial in two variables with

$$\begin{bmatrix} A \\ c^T \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 2 & 0.52 & 1.30 \\ 0 & 0 & 1 & 2 & 0.15 & 1.38 \\ \hline 0.31 & 0.85 & 2.55 & 0.65 & -1.48 & -1.73 \end{bmatrix},$$

then $f_{\text{SAGE}} = 0.00354263 < f_{\text{SAGE}}^{(1)} = 0.13793126$. This signomial has the minimum number of interior exponents needed to be relevant to Conjecture 27 in the nonsimplicial case.

6.3 A Dual Characterization of SAGE Versus Nonnegativity

In this section, we provide a general necessary and sufficient dual characterization in terms of certain moment-type mappings for the question of $C_{\text{SAGE}}(A) = C_{\text{NNS}}(A)$. To

establish this dual characterization we use some new notation. Given two vectors \mathbf{u} , \mathbf{v} the Hadamard product $\mathbf{w} = \mathbf{u} \circ \mathbf{v}$ has entries $w_i = u_i v_i$; this is extended to allow sets in either argument in the same manner as the Minkowski sum. The operator $\mathcal{R}(\cdot)$ returns the range of a matrix.

We begin with the following proposition (proven in the appendix).

Proposition 30 *If \mathbf{A} has $\mathbf{a}_1 = \mathbf{0}$, then the following are equivalent:*

1. *For every vector \mathbf{c} , the function $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ satisfies $f^* = f_{\text{SAGE}}$.*
2. $C_{\text{NNS}}(\mathbf{A}) = C_{\text{SAGE}}(\mathbf{A})$.
3. $\{\mathbf{v} : v_1 = 1, \mathbf{v} \text{ in } C_{\text{SAGE}}(\mathbf{A})^\dagger\} \subset \text{cl conv exp } \mathcal{R}(\mathbf{A}^\top)$.

Our dual characterization consists of two new sets, both parameterized by \mathbf{A} . The first of these sets relates naturally to the third condition in Proposition 30. Formally, the *moment preimage* of some exponent vectors \mathbf{A} is the set

$$T(\mathbf{A}) \doteq \log \text{cl conv exp } \mathcal{R}(\mathbf{A}^\top).$$

Here, we extend the logarithm to include $\log 0 = -\infty$ in the natural way. The second set appearing in our dual characterization is defined less explicitly. For a given \mathbf{A} , we say that $S(\mathbf{A})$ is a set of *SAGE-feasible slacks* if $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ has

$$f_{\text{SAGE}} = \inf\{\mathbf{c}^\top \exp \mathbf{y} : \mathbf{y} \text{ in } \mathcal{R}(\mathbf{A}^\top) + S(\mathbf{A})\}$$

for every \mathbf{c} in \mathbb{R}^m .

Theorem 31 *Let \mathbf{A} have $\mathbf{a}_1 = \mathbf{0}$, and let $S(\mathbf{A})$ be any set of SAGE-feasible slacks over exponents \mathbf{A} . Then $C_{\text{SAGE}}(\mathbf{A}) = C_{\text{NNS}}(\mathbf{A})$ if and only if $S(\mathbf{A}) \subset T(\mathbf{A})$.*

Proof (Theorem 31) To keep notation compact write $U = \mathcal{R}(\mathbf{A}^\top)$ and $S = S(\mathbf{A})$. Also, introduce $W = \{\mathbf{v} : v_1 = 1, \mathbf{v} \text{ in } C_{\text{SAGE}}(\mathbf{A})^\dagger\}$ to describe the feasible set to the dual formulation for f_{SAGE} . By the supporting-hyperplane characterizations of convex sets, the definitions of S and W ensure

$$W = \text{cl conv exp}(U + S).$$

Thus by the equivalence of 1 and 3 in Proposition 30, it follows that all SAGE relaxations will be exact if and only if $\text{exp}(U + S) \subset \text{cl conv exp } U$. We apply a pointwise logarithm to write the latter condition as $U + S \subset \log \text{cl conv exp } U$.

Now we prove that $T \doteq \log \text{cl conv exp } U$ is invariant under translation by vectors in U . It suffices to show that $\text{exp}(\mathbf{v} + T) = \text{exp } T$ for all vectors \mathbf{v} in U . Fixing \mathbf{v} in U we have

$$\begin{aligned} \text{exp}(\mathbf{v} + T) &= \text{exp}(\mathbf{v}) \circ \text{exp}(T) \\ &= \text{exp}(\mathbf{v}) \circ \text{cl conv exp}(U) \\ &= \text{cl conv exp}(\mathbf{v} + U) \\ &= \text{cl conv exp}(U) = \text{exp}(T) \end{aligned}$$

as claimed. This translation invariance establishes that $U + S \subset \log \text{cl conv exp } U$ is equivalent to $S \subset \log \text{cl conv exp } U$, and in turn that condition 1 of Proposition 30 holds if and only if $S \subset \log \text{cl conv exp } U$. The claim now follows by the equivalence of 1 and 2 in Proposition 30. \square

It is the authors' hope that Theorem 31 may help future efforts to resolve Conjecture 27. A starting point in understanding the moment preimage could be to use cumulant generating functions from probability theory. For constructing sets of SAGE-feasible slacks, one might use a change-of-variables argument similar to that seen in the proof of Theorem 9.

Acknowledgements The authors are thankful for the detailed suggestions of anonymous referees, which have led to a much-improved revision of our original manuscript. V.C. would like to acknowledge helpful conversations with Parikshit Shah, particularly on the connections between SAGE and SDSOS polynomials. R.M. was supported in part by NSF grant CCF-1637598 and by an NSF Graduate Research Fellowship. V.C. was supported in part by NSF grants CCF-1350590 and CCF-1637598, AFOSR grant FA9550-16-1-0210, and a Sloan Research Fellowship. A.W. was supported in part by NSF grant CCF-1637598.

7 Appendix

7.1 Proof of Lemma 8

Denote $\text{supp}_{\mathbf{B}}(\lambda) = \{b_j : j \in [d], \lambda_j \neq 0\}$. The proof is constructive, where there is nothing to prove when $\text{supp}_{\mathbf{B}} \lambda$ is simplicial. Suppose then that $\lambda \in \Lambda$ has *nonsimplicial* $\text{supp}_{\mathbf{B}}(\lambda)$. We show that it is possible to decompose $\lambda = z\lambda^{(1)} + (1 - z)\lambda^{(2)}$ for some $z \in (0, 1)$ and $\lambda^{(i)} \in \Lambda$ where $\text{supp } \lambda^{(i)} \subsetneq \text{supp } \lambda$. It should be clear that if this is possible, then the process may be continued in a recursive way if either $\text{supp}_{\mathbf{B}}(\lambda^{(i)})$ are nonsimplicial, and so the claim would follow.

The statement " $\lambda \in \Lambda$ " means that \mathbf{h} may be expressed as a convex combination of vectors in $\text{supp}_{\mathbf{B}}(\lambda)$, and so by Minkowski-Carathéodory, there exists at least one $\lambda^{(1)} \in \Lambda_x$ with $\text{supp } \lambda^{(1)} \subsetneq \text{supp } \lambda$ and simplicial $\text{supp}_{\mathbf{B}}(\lambda^{(1)})$. We will use λ and $\lambda^{(1)}$ to construct the desired $\lambda^{(2)}$ and z .

For each real t , consider $\lambda'_t \doteq \lambda^{(1)} + t(\lambda - \lambda^{(1)})$. It is easy to see that for all t the vector λ'_t belongs to the affine subspace $\{\mathbf{w} : \mathbf{h} = \mathbf{B}\mathbf{w}, \mathbf{1}^\top \mathbf{w} = 1\}$, and furthermore, the support of λ'_t is contained within the support of λ . Now define $T = \max\{t : \lambda'_t \text{ in } \Delta_d\}$; we claim that $T > 1$ and that the support of λ'_T is a proper subset of the support of λ . The latter claim is more or less immediate. To establish the former claim consider how λ'_t (as an affine combination of $\lambda^{(1)}, \lambda$) belongs to Δ_d if and only if it is elementwise nonnegative. This lets us write $T = \max\{t : \lambda'_t \geq \mathbf{0}\}$. Next, use our knowledge about the support of λ'_t to rewrite the constraint " $\lambda'_t \geq \mathbf{0}$ " as " $\lambda_i^{(1)} + t(\lambda_i - \lambda_i^{(1)}) \geq 0$ for all i in $\text{supp } \lambda$." Once written in this form, we see that for $t = 1$ all constraints are satisfied strictly. It follows that $T > 1$ at optimality, and furthermore that the support of λ'_T is distinct from (read: *a proper subset of*) that of λ .

We complete the proof by setting $\lambda^{(2)} = \lambda'_T$ and $z = 1 - 1/T$.

7.2 Proof of Lemma 12

Denote $f = \text{Sig}(A, c)$ and $g = \text{Sig}_F(A, c)$. For brevity write $P = \mathcal{P}(A)$; we may assume without loss of generality that P contains the origin. If $F = P$ then $g = f$ and the claim is trivial. If otherwise, the affine hull of F must have some positive codimension ℓ , and there exist supporting hyperplanes $\{S_i\}_{i=1}^\ell$ such that $F = [\bigcap_{i=1}^\ell S_i] \cap P$. We can express S_i as $\{x : s_i^\top x = r_i\}$ for a vector s_i and a scalar r_i . Because P is convex we know that it is contained in one of the half spaces $\{x : s_i^\top x \leq r_i\}$ or $\{x : s_i^\top x \geq r_i\}$. By possibly replacing (s_i, r_i) by $(-s_i, -r_i)$, we can assume that P is contained in $\{x : s_i^\top x \leq r_i\}$. In addition, the assumption that $\mathbf{0}$ belongs to P ensures that each r_i is nonnegative. Now define $s = \sum_{i=1}^\ell s_i$ and $r = \sum_{i=1}^\ell r_i \geq 0$. The pair (s, r) is constructed to satisfy the following properties:

- For every a_j in F , we have $a_j^\top s = r$.
- For every a_j not in F , we have $a_j^\top s < r$.

Finally, define $h = \text{Sig}_{P \setminus F}(A, c)$ so $f = g + h$. The remainder of the proof is case analysis on r .

If $r = 0$ then we must have $r_i = 0$ for all i . The condition that $r_i = 0$ for all i implies that F is contained in a linear subspace U which is orthogonal to s , and so nonnegativity of g over \mathbb{R}^n reduces to nonnegativity of g over U . Suppose then that there exists some \hat{x} in U where $g(\hat{x})$ is negative. For any vector y in the orthogonal complement of U we have $g(\hat{x} + y) = g(\hat{x})$. Meanwhile no matter the value of \hat{x} we know that $\lim_{t \rightarrow \infty} h(\hat{x} + ts) = 0$. Using $f^* \leq \inf\{f(\hat{x} + ts) : t \in \mathbb{R}\} \leq g(\hat{x})$, we have the desired result for $r = 0$: $g^* < 0$ implies $f^* < 0$.

Now consider the case when r is positive. Define the vector $\hat{s} = rs/\|s\|^2$; we produce an upper bound on f^* by searching over all hyperplanes $\{x : \hat{s}^\top x = t\}$ for t in \mathbb{R} . Specifically, for any x in \mathbb{R}^n there exists a scalar t and a vector y such that $x = t\hat{s} + y$ and $\hat{s}^\top y = 0$. In these terms we have

$$g(t\hat{s} + y) = \exp(t\|\hat{s}\|^2) \sum_{a_i \in F} c_i \underbrace{\exp(t[a_i - \hat{s}]^\top \hat{s})}_{=1 \text{ for all } t} \exp(a_i^\top y). \tag{22}$$

Hence assuming $g^* < 0$ means $\sum_{a_i \in F} c_i \exp(a_i^\top \hat{y}) < 0$ for some \hat{y} in $\text{Span}(s)^\perp$. Using this \hat{y} , one may verify that

$$\lim_{t \rightarrow \infty} f(t\hat{s} + \hat{y}) = -\infty, \tag{23}$$

and so when r is positive, $g^* < 0$ implies $f^* < 0$.

7.3 Proof of Proposition 2

In Sect. 2, we asserted strong duality held between (6) and (7). Here we prove a more general result with two lemmas. In what follows, “co” is an operator that computes a set’s conic hull.

Lemma 32 Fix a closed convex cone K in \mathbb{R}^n . If \mathbf{a} in K^\dagger is such that

$$X \doteq \{\mathbf{x} : \mathbf{a}^\top \mathbf{x} = 1, \mathbf{x} \text{ in } K\}$$

is nonempty, then $\text{cl co } X = K$.

Proof Certainly the conic hull of X is contained within K , and the same is true of its closure. The task is to show that every \mathbf{x} in K also belongs to $\text{cl co } X$; we do this by case analysis on $b \doteq \mathbf{a}^\top \mathbf{x}$.

By the assumptions $\mathbf{a} \in K^\dagger$ and $\mathbf{x} \in K$, we must have $b \geq 0$. If b is positive then the scaling $\tilde{\mathbf{x}} \doteq \mathbf{x}/b$ belongs to K and satisfies $\mathbf{a}^\top \tilde{\mathbf{x}} = 1$. That is, $b > 0$ gives us $\tilde{\mathbf{x}}$ in X . Simply undo this scaling to recover \mathbf{x} and conclude $\mathbf{x} \in \text{co } X$. Now suppose $b = 0$. Here we consider the sequence of points $\mathbf{y}_n \doteq \mathbf{x}_0 + n\mathbf{x}$, where \mathbf{x}_0 is a fixed but otherwise arbitrary element of X . Each point \mathbf{y}_n belongs to K and has $\mathbf{a}^\top \mathbf{y}_n = 1$, hence the \mathbf{y}_n are contained in X . It follows that the scaled points \mathbf{y}_n/n are contained in $\text{cl co } X$, and the same must be true of their limit $\lim_{n \rightarrow \infty} \mathbf{y}_n/n = \mathbf{x}$.

Since \mathbf{x} in K was arbitrary, we have $\text{cl co } X = K$. □

Lemma 33 Let C be a closed and pointed convex cone, and fix $\mathbf{a} \in C \setminus \{\mathbf{0}\}$. Then the primal dual pair

$$f_p = \sup\{\gamma : \mathbf{c} - \gamma \mathbf{a} \text{ in } C\} \quad \text{and} \quad f_d = \inf\{\mathbf{c}^\top \mathbf{v} : \mathbf{a}^\top \mathbf{v} = 1, \mathbf{v} \text{ in } C^\dagger\}$$

exhibits strong duality.

Proof Because C is pointed, C^\dagger is full-dimensional, so there exists no nonzero vector $\tilde{\mathbf{a}} \in \mathbb{R}^n$ where $\tilde{\mathbf{a}}^\top \mathbf{v} = 0$ for all $\mathbf{v} \in C^\dagger$. Consider this fact with $\mathbf{a} \in C \setminus \{\mathbf{0}\}$ to see that there exists a $\mathbf{v} \in C^\dagger$ with $\mathbf{a}^\top \mathbf{v} = 1$. This tells us that the dual feasible set $\{\mathbf{v} : \mathbf{a}^\top \mathbf{v} = 1, \mathbf{v} \text{ in } C^\dagger\}$ is nonempty. Since the dual problem is feasible, a proof that $f_d = f_p$ can be divided into the cases $f_d = -\infty$, and f_d in \mathbb{R} . The proof in former case is trivial; weak duality combined with $f_p \geq -\infty$ gives $f_d = f_p$. In the latter case we prove $f_p \geq f_d$ by showing that $\mathbf{c}^* \doteq \mathbf{c} - f_d \mathbf{a}$ belongs to C .

To prove $\mathbf{c}^* \in C$, we will appeal to Lemma 32 with $K \doteq C^\dagger$. Clearly the set $X = \{\mathbf{v} : \mathbf{a}^\top \mathbf{v} = 1, \mathbf{v} \text{ in } K\}$ is precisely the [nonempty] feasible set for computing f_d , and so from the definition of f_d we have $\mathbf{c}^{*\top} \mathbf{v} \geq 0$ for all \mathbf{v} in X . The inequality also applies to any \mathbf{v} in $\text{cl co } X$, which by Lemma 32 is equal to K^\dagger . Therefore the definition of f_d ensures \mathbf{c}^* is in K^\dagger . Using $K^\dagger \equiv C$, we have the desired result. □

Strong duality in computation of f_{SAGE} for $f = \text{Sig}(\mathbf{A}, \mathbf{c})$ readily follows from Lemma 33. Letting $N = \{i : c_i < 0\}$, simply take $C = \sum_{i \in N \cup \{1\}} C_{\text{AGE}}(\mathbf{A}, i, N)$, and use $\mathbf{a} = \mathbf{e}_1$. What’s more, with appropriate bookkeeping one can use Lemma 33 to prove strong duality in computation of $f_{\text{SAGE}}^{(p)}$ for any nonnegative integer p !

7.4 Proof of Proposition 30

The cases (2) \Rightarrow (1) and (3) \Rightarrow (1) are easy.

$\neg(2) \Rightarrow \neg(1)$. Because $C_{\text{NNS}}(\mathbf{A})$ and $C_{\text{SAGE}}(\mathbf{A})$ are full-dimensional closed convex sets, the condition $C_{\text{SAGE}}(\mathbf{A}) \neq C_{\text{NNS}}(\mathbf{A})$ implies that $C_{\text{NNS}}(\mathbf{A}) \setminus C_{\text{SAGE}}(\mathbf{A})$ has nonempty interior. Assuming this condition, fix a vector \tilde{c} and a radius r such that $B(\tilde{c}, r) \subset C_{\text{NNS}}(\mathbf{A}) \setminus C_{\text{SAGE}}(\mathbf{A})$.⁴ This allows us to strictly separate \tilde{c} from $C_{\text{SAGE}}(\mathbf{A})$, which establishes $f^* \geq f_{\text{SAGE}} + r > f_{\text{SAGE}}$.

(1) \Rightarrow (3). Now suppose that $f^* = f_{\text{SAGE}}$ for all relevant f . In this case, the function $c \mapsto \inf\{c^T x : x \in \Omega\}$ is the same for $\Omega = \text{cl conv exp } \mathcal{R}(\mathbf{A}^T)$ or $\Omega = \{v : v_1 = 1 \text{ and } v \text{ in } C_{\text{SAGE}}(\mathbf{A})^*\}$. This function completely determines the set of all half spaces containing Ω . Since Ω is closed and convex, it is precisely equal to the intersection of all half spaces containing it; the result follows.

References

1. Agrawal, A., Verschuere, R., Diamond, S., Boyd, S.: A rewriting system for convex optimization problems. *Journal of Control and Decision* **5**(1), 42–60 (2018)
2. Ahmadi, A., Majumdar, A.: DSOS and SDSOS optimization: LP and SOCP-based alternatives to sum of squares optimization. In: 2014 48th Annual Conference on Information Sciences and Systems (CISS), pp. 1–5 (2014)
3. August, E., Craciun, G., Koeppl, H.: Finding invariant sets for biological systems using monomial domination. In: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC). IEEE (2012). <https://doi.org/10.1109/cdc.2012.6426491>
4. Bajbar, T., Stein, O.: Coercive polynomials and their newton polytopes. *SIAM Journal on Optimization* **25**(3), 1542–1570 (2015)
5. Björk, K.M., Westerlund, T.: Global optimization of heat exchanger network synthesis problems with and without the isothermal mixing assumption. *Computers & Chemical Engineering* **26**(11), 1581–1593 (2002)
6. Boyd, S., Kim, S.J., Vandenberghe, L., Hassibi, A.: A tutorial on geometric programming. *Optimization and Engineering* **8**(1), 67–127 (2007)
7. Chandrasekaran, V., Shah, P.: Relative entropy relaxations for signomial optimization. *SIAM Journal on Optimization* **26**(2), 1147–1173 (2016)
8. Dembo, R., Avriel, M.: Optimal design of a membrane separation process using signomial programming. *Mathematical Programming* **15**(1), 12–25 (1978)
9. Diamond, S., Boyd, S.: CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* **17**(83), 1–5 (2016)
10. Domahidi, A., Chu, E., Boyd, S.: ECOS: An SOCP solver for embedded systems. In: European Control Conference (ECC), pp. 3071–3076 (2013)
11. Dressler, M., Ilman, S., de Wolff, T.: A positivstellensatz for sums of nonnegative circuit polynomials. *SIAM Journal on Applied Algebra and Geometry* **1**(1), 536–555 (2017)
12. Forsgård, J., de Wolff, T.: The lattice of amoebas (2017)
13. Ghasemi, M., Lasserre, J.B., Marshall, M.: Lower bounds on the global minimum of a polynomial. *Computational Optimization and Applications* **57**(2), 387–402 (2013)
14. Ghasemi, M., Marshall, M.: Lower bounds for polynomials using geometric programming. *SIAM Journal on Optimization* **22**(2), 460–473 (2012)
15. Hall, D., Dowdle, A., Gonzalez, J., Trollinger, L., Thalheimer, W.: Assessment of a boundary layer ingesting turboelectric aircraft configuration using signomial programming. In: 2018 Aviation Technology, Integration, and Operations Conference. American Institute of Aeronautics and Astronautics (2018). <https://doi.org/10.2514/6.2018-3973>
16. Hilbert, D.: Ueber die darstellung definiter formen als summe von formenquadraten. *Mathematische Annalen* **32**(3), 342–350 (1888)
17. Ilman, S., de Wolff, T.: Amoebas, nonnegative polynomials and sums of squares supported on circuits. *Research in the Mathematical Sciences* **3**(1), 9 (2016). <https://doi.org/10.1186/s40687-016-0052-2>

⁴ $B(x, d)$ is ℓ_2 ball centered at x of radius d .

18. Ilman, S., de Wolff, T.: Lower bounds for polynomials with simplex newton polytopes based on geometric programming. *SIAM Journal on Optimization* **26**(2), 1128–1146 (2016)
19. Jabr, R.: Inductor design using signomial programming. *COMPEL - The international journal for computation and mathematics in electrical and electronic engineering* **26**(2), 461–475 (2007)
20. Karaca, O., Darivianakis, G., Beuchat, P., Georghiou, A., Lygeros, J.: The REPOP toolbox: Tackling polynomial optimization using relative entropy relaxations. *IFAC-PapersOnLine* **50**(1), 11652–11657 (2017). 20th IFAC World Congress
21. Khovanskii, A.: Fewnomials and pfaff manifolds. In: *Proceedings of the International Congress of Mathematicians*, vol. 1 (1983)
22. Khovanskii, A.: *Fewnomials*, vol. 88. American Mathematical Soc. (1991)
23. Kirschen, P., York, M., Ozturk, B., Hoburg, W.: Application of signomial programming to aircraft design. *Journal of Aircraft* **55**(3), 965–987 (2018)
24. Kojima, M., Kim, S., Waki, H.: Sparsity in sums of squares of polynomials. *Mathematical Programming* **103**(1), 45–62 (2004)
25. Krivine, J.L.: Anneaux préordonnés. *Journal d’analyse mathématique* **12**, p. 307–326 (1964)
26. Kwan, H., Ho, Y., Wong, N., Ho, K.: Designing globally optimal delta-sigma modulator topologies via signomial programming. In: 2007 International Symposium on VLSI Design, Automation and Test (VLSI-DAT), pp. 1–4 (2007)
27. Lasserre, J.B.: Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization* **11**(3), 796–817 (2001)
28. Maranas, C., Floudas, C.: Global optimization in generalized geometric programming. *Computers & Chemical Engineering* **21**(4), 351–369 (1997)
29. Marshall, M.: *Positive Polynomials and Sums of Squares*. American Mathematical Society (2008). <https://doi.org/10.1090/surv/146>
30. Müller, S., Feliu, E., Regensburger, G., Conradi, C., Shiu, A., Dickenstein, A.: Sign conditions for injectivity of generalized polynomial maps with applications to chemical reaction networks and real algebraic geometry. *Foundations of Computational Mathematics* **16**(1), 69–97 (2015)
31. Müller, S., Hofbauer, J., Regensburger, G.: On the bijectivity of families of exponential/generalized polynomial maps (2018)
32. Murty, K., Kabadi, S.: Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming* **39**(2), 117–129 (1987)
33. Nesterov, Y., Nemirovskii, A.: *Interior-point polynomial algorithms in convex programming*. Society for Industrial and Applied Mathematics, Philadelphia (1994)
34. Nie, J., Demmel, J.: Sparse SOS relaxations for minimizing functions that are summations of small polynomials. *SIAM Journal on Optimization* **19**(4), 1534–1558 (2009)
35. Opgeenoord, M., Cohen, B., Hoburg, W.: Comparison of algorithms for including equality constraints in signomial programming. Tech. Rep. ACDL TR-2017-1, MIT (2017)
36. Oxley, J.: What is a matroid? <https://www.math.lsu.edu/~oxley/survey4.pdf> (2014)
37. Pantea, C., Koepl, H., Craciun, G.: Global injectivity and multiple equilibria in uni- and bi-molecular reaction networks. *Discrete and Continuous Dynamical Systems - Series B* **17**(6), 2153–2170 (2012)
38. Papp, D., Yilliz, S.: On “a homogeneous interior-point algorithm for non-symmetric convex conic optimization” (2017)
39. Parillo, P.: *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. Ph.D. thesis, California Institute of Technology, Pasadena, CA (2000)
40. Pébay, P., Rojas, J.M., Thompson, D.: Optimization and NP_r-completeness of certain fewnomials. In: *Proceedings of the 2009 conference on Symbolic numeric computation*. ACM Press (2009). <https://doi.org/10.1145/1577190.1577212>
41. Reznick, B.: Extremal PSD forms with few terms. *Duke Mathematical Journal* **45**(2), 363–374 (1978)
42. Reznick, B.: Forms derived from the arithmetic-geometric inequality. *Mathematische Annalen* **283**(3), 431–464 (1989)
43. Rountree, D., Rigler, A.: A penalty treatment of equality constraints in generalized geometric programming. *Journal of Optimization Theory and Applications* **38**(2), 169–178 (1982)
44. Schweighofer, M.: An algorithmic approach to Schmüdgen’s Positivstellensatz. *Journal of Pure and Applied Algebra* **166**(3), 307–319 (2002)
45. Seidler, H., de Wolff, T.: An experimental comparison of SONC and SOS certificates for unconstrained optimization (2018)

46. Seidler, H., de Wolff, T.: POEM: Effective methods in polynomial optimization, version 0.2.1.0(a). <http://www.iaa.tu-bs.de/AppliedAlgebra/POEM/index.html> (2019)
47. Serrano, S.: Algorithms for unsymmetric cone optimization and an implementation for problems with the exponential cone. Ph.D. thesis, Stanford University, Palo Alto, CA (2015)
48. Shen, P.: Linearization method of global optimization for generalized geometric programming. *Applied Mathematics and Computation* **162**(1), 353–370 (2005)
49. Serali, H.: Global optimization of nonconvex polynomial programming problems having rational exponents. *Journal of Global Optimization* **12**(3), 267–283 (1998)
50. Shor, N.Z.: Class of global minimum bounds of polynomial functions. *Cybernetics* **23**(6), 731–734 (1988)
51. Silipo, J.: The Ronkin number of an exponential sum. *Mathematische Nachrichten* **285**(8-9), 1117–1129 (2012)
52. Skajaa, A., Ye, Y.: A homogeneous interior-point algorithm for nonsymmetric convex conic optimization. *Mathematical Programming* **150**(2), 391–422 (2014)
53. Tunçel, L.: Generalization of primal—dual interior-point methods to convex optimization problems in conic form. *Foundations of Computational Mathematics* **1**(3), 229–254 (2001)
54. Waki, H., Kim, S., Kojima, M., Muramatsu, M.: Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM Journal on Optimization* **17**(1), 218–242 (2006)
55. Wang, J.: Nonnegative polynomials and circuit polynomials. [arXiv:1804.09455v1](https://arxiv.org/abs/1804.09455v1) (2018)
56. Wang, J.: On supports of sums of nonnegative circuit polynomials. [arXiv:1809.10608v1](https://arxiv.org/abs/1809.10608v1) (2018)
57. Yan, J.: Signomial programs with equality constraints: numerical solution and applications. Ph.D. thesis, University of British Columbia (1976)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.