RECOGNIZING SIGNOMIAL CONVEXITY IS HARD

Anonymous authors
Paper under double-blind review

ABSTRACT

Signomials—finite sums of generalized monomials with real exponents—are a natural machine learning model, combining parsimonious representations of power laws, inverse relationships, and multiplicative interactions with universal approximation and interpretable parameters. In fact, 45 of the 100 equations in the AI Feynman benchmark admit signomial representations. Moreover, signomial optimization is widely used in engineering design, communications, economics, and machine learning, but is computationally intractable in general. Convexity, the gold standard for efficient optimization, offers three routes to tractability for signomials: a signomial may be convex in its original variables, become convex after a logarithmic change of variables, or, when positive, become convex after additionally taking the logarithm of its value—the structure underlying disciplined geometric programming. We show that recognizing each form is strongly NP-hard, both on compact domains and globally. Our proofs start from gap-promise variants of polynomial convexity and develop polynomial-to-signomial reductions that preserve curvature gaps under the relevant logarithmic transformations.

1 INTRODUCTION

Signomials are finite linear combinations of generalized monomials with real exponents. Formally, a signomial in n variables is a function

$$f(x) = \sum_{\alpha \in A} c_\alpha \prod_{j=1}^{n} x_j^{\alpha_j}, \qquad x_1, \ldots, x_n > 0, \tag{1}$$

where $A \subset \mathbb{R}^n$ is finite and $c_\alpha \in \mathbb{R}$. In particular, signomials substantially generalize polynomial models by allowing fractional and negative powers: expressions such as $\sqrt{x_1}$, $x_2^{-1} = 1/x_2$, and $x_1^{3/2} x_2^{-1/3}$ are all generalized monomials. Because arbitrary real powers need not be real-valued at nonpositive arguments, the natural domain of a signomial is the positive orthant $\mathbb{R}^n_{++}$.

Signomials are a natural model class for scientific equation discovery. The AI Feynman benchmark contains 100 equations drawn from the three volumes of *The Feynman Lectures on Physics* (Udrescu & Tegmark, 2020), 45 of which admit signomial representations (Lumadjeng et al., 2026). This prevalence reflects the ability of signomials to represent power laws, inverse relationships, roots, and multiplicative interactions exactly. For machine-learning applications, signomials combine parsimony, interpretability, and expressive power (Karcher, 2023; Lee et al., 2014; Hwang et al., 2015; Lumadjeng et al., 2026). First, relationships such as $\sqrt{x}$ and $1/x$ require a single signomial term, whereas polynomials and neural networks can approximate them but do not explicitly encode their power-law structure. Second, the model parameters (the exponents) have direct interpretations as elasticities. Third, signomials inherit the universal-approximation property of polynomials on compact subsets of $\mathbb{R}^n_{++}$.

The problem of minimizing a signomial $f(x)$ over $x \in \mathbb{R}^n_{++}$, subject to signomial constraints $f_i(x) \leq 1$ and $g_i(x) = 1$, is called a *signomial optimization problem* (Boyd et al., 2007, Section 9.1). Such problems arise in machine learning, wireless communications, inventory management, biotechnology, aerospace engineering, and chemical engineering (Karcher, 2023; Févotte & Idier, 2011; Zimmert & Seldin, 2021; Chiang et al., 2007; Chiu et al., 2007; San-José et al., 2022; Marin-Sanguino et al., 2007; Kirschen et al., 2018; Björk & Westerlund, 2002; Dressler & Murray, 2022; Maranas & Floudas, 1997; Lee et al., 2014; Hwang et al., 2015; Lumadjeng et al., 2026).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

This expressiveness, however, comes with a computational cost: signomial optimization is NP-hard in general (see, e.g., Chandrasekaran & Shah, 2016). This follows already from matrix copositivity: deciding whether $x^\top Qx \geq 0$ for every $x \geq 0$ is coNP-complete (Murty & Kabadi, 1987). Consequently, deciding whether the quadratic signomial $x^\top Qx$ attains a negative value is also hard.

The hardness of signomial optimization motivates the search for tractable subclasses, foremost among them convex optimization problems. For signomials, however, convexity can arise in three distinct forms. First, f may be convex in its original positive variables x. Second, after the logarithmic reparameterization $x = e^y$, the function $y \mapsto f(e^y)$ may be convex on $\mathbb{R}^n$; we call this *log-coordinate convexity*. Third, when $f > 0$, the function $y \mapsto \log f(e^y)$ may be convex; this property is known as *log-log convexity*. Figure 1 summarizes these three notions and their relationships.

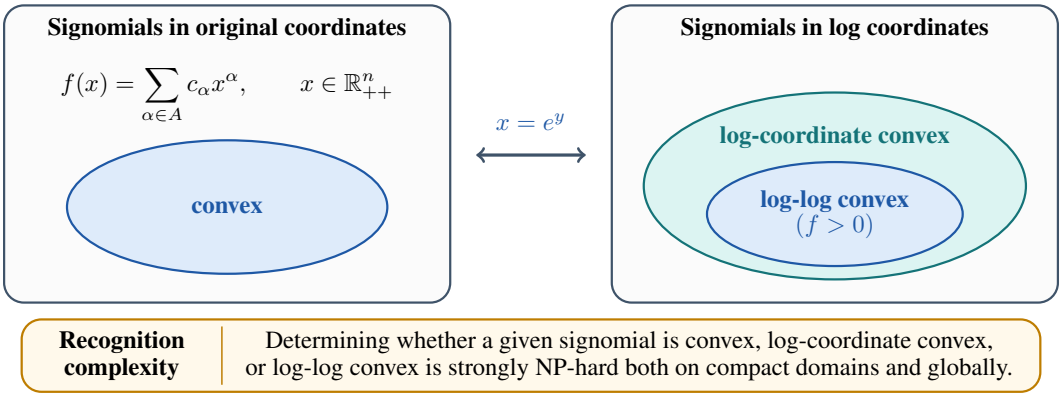


Figure 1: The three notions of convexity for signomials and their recognition complexity.

Log-log convexity plays a central role in disciplined geometric programming (DGP) (Agrawal et al., 2019), much as ordinary convexity does in disciplined convex programming (Grant et al., 2006). DGP begins with a library of atoms whose log-log curvature—affine, convex, or concave—is known in advance. It then uses curvature-preserving composition rules, guaranteeing that every compliant problem becomes convex after a logarithmic change of variables. Classical geometric programming is a canonical special case: applying the logarithmic transformation to a posynomial—a signomial with positive coefficients—produces a log-sum-exp function (Boyd et al., 2007).

However, signomial models learned from data need not satisfy the curvature conditions enforced by these composition rules. Regression and equation-learning methods estimate their coefficients and exponents directly from observations (Karcher, 2023; Lee et al., 2014; Lumadjeng et al., 2026), so the resulting models carry no a priori guarantee of convexity in any of the three senses above. This leads to the central question of this paper: *Can we efficiently determine whether a given signomial is convex, log-coordinate convex, or log-log convex?*

Contributions. We give a strongly negative answer to our central question: recognizing each of the three forms of convexity is strongly NP-hard, both on compact domains and globally; see Figure 1. These statements correspond to six distinct recognition problems. In the original variables x, compact-domain hardness concerns a box contained in $\mathbb{R}^n_{++}$, whereas global hardness concerns the entire positive orthant. For log-coordinate and log-log convexity, the transformed variable $y = \log x$ belongs to $\mathbb{R}^n$, so the corresponding domains are a compact box in $\mathbb{R}^n$ and all of $\mathbb{R}^n$, respectively. Importantly, compact-domain hardness does not immediately imply global hardness: a function that is convex on a box need not remain convex outside it, so the positive instances of the compact problem are not necessarily positive instances of the global problem.

Our reductions are from strongly NP-hard gap-promise problems for polynomial convexity derived from the results of Ahmadi et al. (2013); Ahmadi & Hall (2020). Hardness in the original variables follows easily from the polynomial case: an affine change of variables maps the relevant polynomial instances to signomials over the positive orthant (Theorem 2.1). However, hardness for log-coordinate and log-log convexity is more delicate: exponential reparameterization alters curvature, while the additional logarithm in log-log convexity introduces a gradient-dependent term. Our main

technical contribution is a family of polynomial-to-signomial reductions that preserves the curvature gap under these transformations.

Section 2 develops the gap-preserving reduction framework used to control these curvature changes. Section 3 applies it to prove strong NP-hardness of recognizing log-coordinate convexity (Theorem 3.1) and log-log convexity (Theorem 3.2) on the unit box. Finally, Section 4 establishes the corresponding global results over $\mathbb{R}^n$ (Theorems 4.2 and 4.3).

Related work. To the best of our knowledge, no prior work has studied the computational complexity of recognizing convexity, log-coordinate convexity, or log-log convexity of signomials. The relevant hardness results concern polynomial convexity. Ahmadi et al. (2013) proved that deciding whether a quartic polynomial is globally convex is strongly NP-hard, while Ahmadi & Hall (2020) established strong NP-hardness for cubic polynomials over a box. These results underlie the global quartic and cubic box gap-convexity problems used in our reductions and yield direct reductions to signomial convexity in the original variables. They do not, however, settle the complexity of log-coordinate or log-log convexity: exponential reparameterization changes curvature, and the additional outer logarithm in the log-log setting introduces a gradient-dependent term.

Related work on signomial optimization has largely focused on identifying tractable subclasses or constructing global relaxations. Disciplined geometric programming provides a compositional framework for constructing problems whose convexity after a logarithmic transformation is certified in advance (Boyd et al., 2007; Agrawal et al., 2019). A complementary line of work develops global optimization methods based on SAGE certificates and relative-entropy relaxations (Chandrasekaran & Shah, 2016; Murray et al., 2021a;b; Wang et al., 2020; Dressler & Murray, 2022). These methods certify signomial nonnegativity and construct hierarchies of convex relaxations for approximating global optima, but do not address recognition of the three convexity notions studied here.

2 Convexity Problems and Reduction Framework

2.1 Gap Convexity: definition and hardness

Let $\Omega \subseteq \mathbb{R}^n$ be a convex set given by a finite description, and let $\mathcal{F}(\Omega)$ be a finitely encoded class of twice continuously differentiable functions $p : \Omega \to \mathbb{R}$. For $p \in \mathcal{F}(\Omega)$, define

$$\kappa_\Omega(p) = \inf_{\xi \in \Omega} \lambda_{\min}\big(\nabla^2 p(\xi)\big).$$

Problem 1 (Gap-Convexity)**.** *The input consists of a function $p \in \mathcal{F}(\Omega)$ and two positive rational thresholds γ_+ and γ_-. It is promised that exactly one of the following alternatives holds:*

$$\begin{aligned} \textsc{Yes}: &\quad \kappa_\Omega(p) \geq \gamma_+, \\ \textsc{No}: &\quad \kappa_\Omega(p) \leq -\gamma_-. \end{aligned}$$

The task is to decide which alternative holds.

The positive alternative implies that p is uniformly γ_+-strongly convex on Ω, whereas the negative alternative guarantees negative curvature and hence that p is not convex on Ω.

Given two gap-convexity problems, a reduction from the first to the second is a polynomial-time algorithm that maps an instance p of the first problem to an instance q of the second, so that solving the latter determines which alternative holds for p. For such a reduction, suppose that

$$\big|\kappa_\Omega(q) - \kappa_\Omega(p)\big| \leq E.$$

If p satisfies the positive alternative $\kappa_\Omega(p) \geq \gamma_+$, then $\kappa_\Omega(q) \geq \gamma_+ - E$. Similarly, if p satisfies the negative alternative $\kappa_\Omega(p) \leq -\gamma_-$, then $\kappa_\Omega(q) \leq -\gamma_- + E$. Figure 2 illustrates this mechanism. To preserve the sign of the curvature in both alternatives, we require $E < \min\{\gamma_+, \gamma_-\}$. In particular, choosing $E \leq \frac{1}{2}\min\{\gamma_+, \gamma_-\}$ retains curvature margins of at least $\gamma_+/2$ and $\gamma_-/2$, respectively. The factor $1/2$ is only for convenience; any fixed fraction below one would suffice.

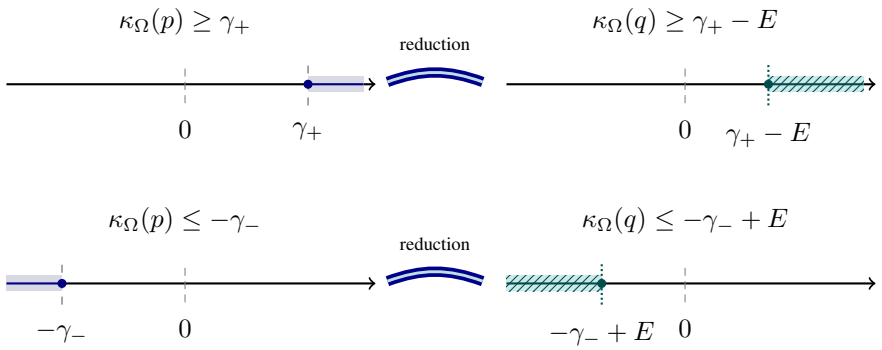


Figure 2: In the positive alternative, the lower curvature bound decreases by at most E. In the negative alternative, the negative-curvature threshold increases by at most E.

We base our reductions on the following two restrictions of GAP-CONVEXITY:

(CBGC) *Cubic box gap convexity:* $\Omega = [-1, 1]^n$ and $p : \mathbb{R}^n \to \mathbb{R}$ is a cubic polynomial.

(GQGC) *Global quartic gap convexity:* $\Omega = \mathbb{R}^n$ and $p : \mathbb{R}^n \to \mathbb{R}$ is a quartic polynomial.

Both CBGC and GQGC are strongly NP-hard. Ahmadi & Hall (2020) established hardness of CBGC with inverse-polynomial curvature margins: $\gamma_+, \gamma_- \geq \frac{1}{\mathrm{poly}(n)}$. For GQGC, Ahmadi et al. (2013) established hardness with the one-sided gap $\gamma_+ = 0, \gamma_- \geq \frac{1}{\mathrm{poly}(n)}$. In Appendix A.1, we modify their construction to obtain $\gamma_+, \gamma_- \geq \frac{1}{\mathrm{poly}(n)}$, as required by our gap-preserving reductions.

In both cases, the hardness persists even when every coefficient a/b of p, written in lowest terms with $b > 0$, satisfies $|a|, b \leq \mathrm{poly}(n)$. This is the meaning of the "strong" qualification in our hardness statements: the coefficient numerators and denominators are polynomially bounded, and the curvature margins are inverse-polynomial in n. Thus, the hardness does not rely on exponentially large coefficients or exponentially small curvature gaps.

2.2 SIGNOMIAL CONVEXITY-RECOGNITION PROBLEMS

We study three signomial convexity-recognition problems, each with a constrained and an unconstrained variant. Throughout, f is a signomial as in Equation 1, with rational coefficients and rational exponents encoded in binary. Whenever $f(e^y)$ appears, exponentiation is applied componentwise. In the constrained variants, the input additionally includes a finitely described compact convex set $X \subseteq \mathbb{R}^n_{++}$ or $Y \subseteq \mathbb{R}^n$, as appropriate.

Problem 2 (SIGNOMIAL-CONVEXITY)**.** *Given a signomial f, consider the following two variants:*

(i) Constrained: *Decide whether f is convex on X.*

(ii) Unconstrained: *Decide whether f is convex on $\mathbb{R}^n_{++}$.*

Problem 3 (LOG-COORDINATE-CONVEXITY)**.** *Given a signomial f, consider the following two variants:*

(i) Constrained: *Decide whether $y \mapsto f(e^y)$ is convex on Y.*

(ii) Unconstrained: *Decide whether $y \mapsto f(e^y)$ is convex on $\mathbb{R}^n$.*

Problem 4 (LOG-LOG-CONVEXITY)**.** *Given a signomial f, under the promise that $f > 0$ throughout the relevant domain, consider the following two variants:*

(i) Constrained: *Decide whether $y \mapsto \log f(e^y)$ is convex on Y.*

(ii) Unconstrained: *Decide whether $y \mapsto \log f(e^y)$ is convex on $\mathbb{R}^n$.*

Define $g(y) = f(e^y)$. Accordingly, LOG-COORDINATE-CONVEXITY asks whether g is convex, whereas, under the promise that $g > 0$, LOG-LOG-CONVEXITY asks whether $\log g$ is convex. Thus, f is log-log convex precisely when its exponential form g is log-convex.

The two SIGNOMIAL-CONVEXITY problems admit direct reductions from CBGC and GQGC. The constrained reduction uses an affine bijection from a positive box onto the unit box, while the unconstrained reduction uses the linear map $(u, v) \mapsto u - v$ from $\mathbb{R}^{2n}_{++}$ onto $\mathbb{R}^n$. Both substitutions produce ordinary polynomials, and hence signomials on the positive orthant, so no approximation construction is needed.

Theorem 2.1. *Both the constrained and unconstrained* SIGNOMIAL-CONVEXITY *problems are strongly NP-hard.*

Proof. For the constrained problem, let $p : \mathbb{R}^n \to \mathbb{R}$ be an instance of CBGC and define $X = [1, 2]^n$, $T(x) = 2x - 3\mathbf{1}$, $f(x) = p(T(x))$. The affine map T sends X bijectively onto $[-1, 1]^n$, and $\nabla^2 f(x) = 4\nabla^2 p(T(x))$. Consequently, f is convex on X if and only if p is convex on $[-1, 1]^n$. Moreover, f is a polynomial and hence a signomial on $\mathbb{R}^n_{++}$.

For the unconstrained problem, let $p : \mathbb{R}^n \to \mathbb{R}$ be an instance of GQGC and define $f(u, v) = p(u - v)$, $(u, v) \in \mathbb{R}^{2n}_{++}$. If p is convex on $\mathbb{R}^n$, then f is convex as its composition with a linear map. Conversely, suppose that f is convex on $\mathbb{R}^{2n}_{++}$. For any $x, d \in \mathbb{R}^n$, choose $u, v \in \mathbb{R}^n_{++}$ such that $u - v = x$. Let $z = (d^\top, 0)^\top$, then $z^\top \nabla^2 f(u, v) z = d^\top \nabla^2 p(x) d \geq 0$. Since x and d are arbitrary, p is convex on $\mathbb{R}^n$. Thus, f is convex on $\mathbb{R}^{2n}_{++}$ if and only if p is convex on $\mathbb{R}^n$.

Both transformations are computable in polynomial time, preserve fixed degree, and produce coefficients of polynomial encoding length. Strong NP-hardness therefore follows. □

Unlike SIGNOMIAL-CONVEXITY, neither transformed problem inherits hardness directly from polynomial convexity. The exponential change of variables does not preserve curvature; for log-log convexity, taking the logarithm further changes the curvature condition and requires positivity.

The preceding discussion in Section 2.1 motivates a *gap-preserving reduction*. Given an instance p of a known hard gap-convexity problem, we seek to construct in polynomial time a function q for the target recognition problem such that $|\kappa_Y(q) - \kappa_Y(p)| \leq E$, $E \leq \frac{1}{2}\min\{\gamma_+, \gamma_-\}$. This bound preserves positive curvature in one alternative and negative curvature in the other. Thus, an algorithm for the target recognition problem would solve the known hard gap-convexity problem. If the construction also preserves the required encoding bounds, strong NP-hardness transfers as well. By Weyl's inequality (Bhatia, 2013), it suffices to establish

$$\sup_{\xi \in Y} \|\nabla^2 q(\xi) - \nabla^2 p(\xi)\|_2 \leq E. \tag{2}$$

2.3 A GAP-PRESERVING REDUCTION FRAMEWORK

We now develop a gap-preserving reduction framework for the two transformed problems. Given a polynomial p, we construct a family q_h that is the exponential form of a signomial and satisfies the Hessian perturbation bound Equation 2 for sufficiently small $h > 0$. The construction applies a coordinatewise deformation ϕ_h that approaches the identity as $h \to 0$, so that the Hessian of q_h uniformly approximates that of p on the unit box. This approximation preserves the curvature gap.

Curvature requirements. Let $\phi_h : \mathbb{R} \to \mathbb{R}$ be twice continuously differentiable and depend on $h > 0$. Apply it coordinatewise as $\Phi_h(y) = (\phi_h(y_1), \ldots, \phi_h(y_n))^\top$, and set $q_h(y) = p(\Phi_h(y))$. Thus q_h is obtained by evaluating p at a coordinatewise deformation of y.

To understand how this deformation changes curvature, write $J_h(y) = \mathrm{Diag}(\phi_h'(y_1), \ldots, \phi_h'(y_n))$. The Hessian chain rule gives $\nabla^2 q_h = J_h^\top \nabla^2 p(\Phi_h) J_h + \mathrm{Diag}(\phi_h''(y_i) \partial_i p(\Phi_h(y)))_{i=1}^n$. Adding and subtracting $\nabla^2 p(\Phi_h(y))$ yields

$$\begin{aligned} \nabla^2 q_h(y) - \nabla^2 p(y) = &\underbrace{J_h(y)^\top \nabla^2 p(\Phi_h(y)) J_h(y) - \nabla^2 p(\Phi_h(y))}_{\text{Jacobian error}} \\ &+ \underbrace{\nabla^2 p(\Phi_h(y)) - \nabla^2 p(y)}_{\text{argument error}} + \underbrace{\mathrm{Diag}(\phi_h''(y_i)\, \partial_i p(\Phi_h(y)))_{i=1}^n}_{\text{second-derivative error}}. \end{aligned} \tag{3}$$

Equation 3 identifies three sources of error. Uniformly on $[-1, 1]$, the argument error reflects the fact that the Hessian of p is evaluated at $\Phi_h(y)$ rather than at y; it therefore becomes small when $\phi_h(t) \to t$. The Jacobian error is caused by the local rescaling $J_h(y)$ introduced by the coordinate map, and becomes small when $\phi_h'(t) \to 1$. The second-derivative error is generated by the nonlinearity of ϕ_h, and becomes small when $\phi_h''(t) \to 0$.

To quantify the three errors in Equation 3, we measure how closely ϕ_h follows the identity map on $[-1, 1]$. Let $E_0(h) = \sup_{|t|\le 1} |\phi_h(t) - t|$, which measures how far the transformed argument can move from the original one; $E_1(h) = \sup_{|t|\le 1} |\phi_h'(t) - 1|$, which measures how far the local scaling deviates from the identity; and $E_2(h) = \sup_{|t|\le 1} |\phi_h''(t)|$, which measures the nonlinearity of the coordinate map. The first quantity also determines where the transformed argument can lie. If $E_0(h) \le 1$, then $\Phi_h([-1, 1]^n) \subseteq [-2, 2]^n$, since $|\phi_h(t)| \le |t| + |\phi_h(t) - t| \le 2$. It therefore suffices to control the derivatives of p on the fixed enlarged box $[-2, 2]^n$.

We next introduce the derivative bounds needed to control the three errors. Let G_p be a rational upper bound on the gradient of p: $G_p \ge \sup_{\xi\in[-2,2]^n} \|\nabla p(\xi)\|_\infty$, and let H_p be a rational upper bound on its Hessian: $H_p \ge \sup_{\xi\in[-2,2]^n} \|\nabla^2 p(\xi)\|_2$. Since p is a polynomial, $\nabla^2 p$ is Lipschitz on the compact box $[-2, 2]^n$. Let L_p be a rational Lipschitz bound satisfying $\|\nabla^2 p(\xi) - \nabla^2 p(\zeta)\|_2 \le L_p\|\xi - \zeta\|_\infty$ for $\xi, \zeta \in [-2, 2]^n$. Thus L_p controls how rapidly the Hessian can vary across the box.

These quantities control the three errors directly. The argument error is at most $L_pE_0(h)$, because $\Phi_h(y)$ lies within $E_0(h)$ of y. For the Jacobian error, writing $J_h = I_n + \Delta_h$ gives $\|\Delta_h\|_2 \le E_1(h)$, so its norm is bounded by $H_p(2E_1(h) + E_1(h)^2)$. Finally, the second-derivative error is diagonal, and its spectral norm is at most $G_pE_2(h)$. Combining the three bounds gives

$$\sup_{y\in[-1,1]^n} \|\nabla^2 q_h(y) - \nabla^2 p(y)\|_2 \le L_pE_0(h) + H_p\big(2E_1(h) + E_1(h)^2\big) + G_pE_2(h). \tag{4}$$

Hence, if $E_0(h)$, $E_1(h)$, and $E_2(h)$ are all $O(h)$, then the entire Hessian perturbation is $O(h)$.

Canonical exponential deformation. The error $E_0(h)$ measures how closely ϕ_h approximates the identity, so we require $\phi_h(t) \to t$ as $h \to 0$. At the same time, ϕ_h should be built from exponentials so that $p \circ \Phi_h$ is the exponential form of a signomial. The simplest starting point is e^{ht}. For small h, Taylor's first-order expansion gives $e^{ht} = 1 + ht + O(h^2)$. This suggests the choice

$$\phi_h(t) = \frac{e^{ht} - 1}{h}. \tag{5}$$

We then have $(e^{ht} - 1)/h = t + O(h)$, so $\phi_h(t) \to t$ as $h \to 0$. It remains to verify the derivative requirements. Direct differentiation gives $\phi_h'(t) = e^{ht}$ and $\phi_h''(t) = he^{ht}$, so $\phi_h'(t) \to 1$ and $\phi_h''(t) \to 0$ as $h \to 0$. More precisely, for $|t| \le 1$ and $0 < h \le 1/2$, Taylor's theorem yields $E_0(h) \le h$, $E_1(h) \le 2h$, and $E_2(h) \le 2h$. Thus, all three errors are $O(h)$, and substituting these bounds into Equation 4 shows that the Hessian perturbation is also $O(h)$.

Moreover, substituting $\phi_h(y_i) = (e^{hy_i} - 1)/h$ into the finitely many monomials of p and expanding their integer powers expresses $q_h = p \circ \Phi_h$ as a finite sum of terms $c_\ell e^{h\ell^\top y}$. Thus, for rational $h > 0$, q_h is the exponential form of a signomial with rational coefficients and rational exponents.

The preceding construction shows that the Hessian perturbation vanishes as $h \to 0$. To apply the gap-preservation principle, we need a quantitative bound that determines how small h must be for Equation 2 to hold. The following lemma provides this bound.

Lemma 2.2 (Gap-preserving curvature approximation)**.** *Let $p : \mathbb{R}^n \to \mathbb{R}$ be a polynomial, and let $q_h = p \circ \Phi_h$. Then q_h is the exponential form of a signomial and, for every rational $0 < h \le 1/2$,*

$$\sup_{y\in[-1,1]^n} \big\|\nabla^2 q_h(y) - \nabla^2 p(y)\big\|_2 \le C_{\mathrm{err}}(p)h,$$

where $C_{\mathrm{err}}(p) = L_p + 6H_p + 2G_p$.

Thus, the Hessian perturbation condition Equation 2 holds with $E = C_{\mathrm{err}}(p)h$. Let $\gamma = \frac{1}{2}\min\{\gamma_+, \gamma_-\}$ and $\overline{C}_{\mathrm{err}}(p) = \max\{1, C_{\mathrm{err}}(p)\}$, and choose $h = \min\{\frac{1}{2}, \gamma/\overline{C}_{\mathrm{err}}(p)\}$. Since $C_{\mathrm{err}}(p) \le \overline{C}_{\mathrm{err}}(p)$, this choice gives $E \le \overline{C}_{\mathrm{err}}(p)h \le \gamma = \frac{1}{2}\min\{\gamma_+, \gamma_-\}$. Hence the two alternatives retain curvature margins of at least $\gamma_+/2$ and $\gamma_-/2$, respectively. In our setting, $\overline{C}_{\mathrm{err}}(p)$

is polynomially bounded, while γ_+ and γ_- have inverse-polynomial magnitude. Therefore, h can be chosen rational with inverse-polynomial magnitude and polynomial encoding length. The resulting q_h can be constructed in polynomial time and also has polynomial encoding length. The details are given in Appendix A.2.

3 Constrained Hardness for Transformed Convexity

We now establish strong NP-hardness for the constrained versions of the two transformed convexity problems. Both reductions apply the gap-preserving exponential approximation of Section 2.3 on a compact domain, with the approximation parameter chosen so that the Hessian perturbation remains smaller than the curvature margins of CBGC. We begin with log-coordinate convexity.

Theorem 3.1. *The constrained* Log-Coordinate-Convexity *problem is strongly NP-hard.*

Proof. Let p be an instance of CBGC on $Y = [-1, 1]^n$, with curvature margins γ_+ and γ_-. By Lemma 2.2, we can choose a positive rational h, with polynomial encoding length, and construct a signomial $q_h(y) = f_h(e^y)$ such that $\sup_{y\in Y} \|\nabla^2 q_h(y) - \nabla^2 p(y)\|_2 \le E$, $E \le \frac{1}{2}\min\{\gamma_+, \gamma_-\}$. By the gap-preserving argument of Section 2.1, q_h is convex on Y in the positive case and nonconvex on Y in the negative case. Since $q_h(y) = f_h(e^y)$, the signomial f_h is log-coordinate convex on Y exactly in the positive case.

The construction of f_h is computable in polynomial time and preserves the required coefficient and exponent encoding bounds. Thus, an algorithm for the constrained version of Log-Coordinate-Convexity would solve CBGC. Since CBGC is strongly NP-hard and the reduction preserves the required encoding bounds, the result follows. □

We next consider log-log convexity. For any positive twice continuously differentiable function g, $\nabla^2 \log g(y) = (g(y)\nabla^2 g(y) - \nabla g(y)\nabla g(y)^\top)/g(y)^2$. Since $g(y)^2 > 0$, the function $\log g$ is convex exactly when the numerator is positive semidefinite throughout the domain. We therefore define the log-curvature matrix by

$$M_g(y) = g(y)\nabla^2 g(y) - \nabla g(y)\nabla g(y)^\top. \tag{6}$$

Our reduction adds a positive constant to the exponential approximation and scales its remaining terms so that g stays positive and, in the positive alternative, $g\nabla^2 g$ dominates $\nabla g \nabla g^\top$.

Theorem 3.2. *The constrained* Log-Log-Convexity *problem is strongly NP-hard.*

Proof. Let p be an instance of CBGC on $Y = [-1, 1]^n$. As in the proof of Theorem 3.1, choose a positive rational h and construct $q_h(y) = f_h(e^y)$. The gap-preserving approximation retains inverse-polynomial margins $\gamma_+, \gamma_- > 0$, so $\kappa_Y(q_h) \ge \gamma_+$ in the positive case, $\kappa_Y(q_h) \le -\gamma_-$ in the negative case.

Choose rational bounds $B_0, B_1 \ge 1$ such that $|q_h(y)| \le B_0$ and $\|\nabla q_h(y)\|_2 \le B_1$ for every $y \in Y$, and take a positive rational $\varepsilon \le \min\{1/(2B_0), \gamma_+/(4B_1^2)\}$. Define $g_h(y) = 1 + \varepsilon q_h(y)$. Then $g_h \ge 1/2$ on Y, and $M_{g_h} = \varepsilon g_h \nabla^2 q_h - \varepsilon^2 \nabla q_h \nabla q_h^\top$.

In the positive case, Weyl's inequality gives, for every $y \in Y$,

$$\lambda_{\min}(M_{g_h}(y)) \ge \varepsilon g_h(y)\lambda_{\min}\big(\nabla^2 q_h(y)\big) - \varepsilon^2 \|\nabla q_h(y)\|_2^2 \ge \varepsilon\left(\frac{\gamma_+}{2} - \varepsilon B_1^2\right) \ge \frac{\varepsilon\gamma_+}{4}.$$

Hence $M_{g_h}(y) \succeq 0$ throughout Y, and therefore $\log g_h$ is convex on Y.

In the negative case, there exists $\hat{y} \in Y$ such that $\lambda_{\min}(\nabla^2 q_h(\hat{y})) \le -\gamma_-$. Since the rank-one term in M_{g_h} is negative semidefinite,

$$\lambda_{\min}(M_{g_h}(\hat{y})) \le \varepsilon g_h(\hat{y})\lambda_{\min}\big(\nabla^2 q_h(\hat{y})\big) \le -\varepsilon g_h(\hat{y})\gamma_- < 0.$$

Thus, $\log g_h$ is not convex on Y.

The exponential-approximation bounds allow B_0, B_1, and ε to be chosen with polynomial encoding length. Since q_h is the exponential form of a signomial, so is $g_h = 1 + \varepsilon q_h$. The corresponding signomial can therefore be constructed in polynomial time while preserving the required encoding bounds. This gives a reduction from CBGC, and strong NP-hardness follows. □

4 UNCONSTRAINED HARDNESS FOR TRANSFORMED CONVEXITY

We now establish strong NP-hardness for the unconstrained versions of the two transformed convexity problems. Compact-domain hardness need not extend to $\mathbb{R}^n$: an interior negative-curvature witness remains valid globally, but convexity on the unit box gives no control outside it. To bridge this gap, both reductions augment the same hard exponential form q with an auxiliary construction that extends the curvature separation beyond the unit box while preserving the interior witness.

Let p be an instance of GQGC, choose a rational h according to the gap-preservation condition of Section 2.3, and write $q = q_h = p \circ \Phi_h$. Since p has degree four and $h \le \frac{1}{2}$, $D = \max_{\alpha \in \mathrm{supp}(q)} \|\alpha\|_1 \le 2$. Let $\gamma = \frac{1}{2}\min\{\gamma_+, \gamma_-\}$. By Lemma 2.2 and the interior-witness property proved in Appendix A.1, it is strongly NP-hard to distinguish $\kappa_{[-1,1]^n}(q) \ge \gamma$ from $\kappa_{[-\frac{1}{2},\frac{1}{2}]^n}(q) \le -\gamma$.

To extend this curvature separation to all of $\mathbb{R}^n$, we augment q with an auxiliary exponential form:

$$Q_L(y) = q(y) + \delta_L \Psi_L(y). \tag{7}$$

Here, Ψ_L is designed to provide sufficiently large positive curvature outside the unit box to dominate any negative curvature of q. The coefficient $\delta_L > 0$ will be chosen small enough that Q_L retains the interior negative curvature of q, while L controls how rapidly the curvature supplied by Ψ_L grows outside the unit box.

A natural first choice for Ψ_L is $\sum_i (e^{Ly_i} + e^{-Ly_i})$, since, whenever $|y_i|$ is large, one of e^{Ly_i} and e^{-Ly_i} grows as $e^{L|y_i|}$. However, its Hessian is diagonal, so the rapidly growing curvature generated by y_i is confined to the e_i-direction and may not dominate negative curvature of q in other directions. To distribute this growth across every coordinate direction, we couple $e^{Ly_i} + e^{-Ly_i}$ with $e^{y_j} + e^{-y_j}$ for every i, j and define

$$\Psi_L(y) = \sum_{i,j=1}^{n} \left(e^{Ly_i} + e^{-Ly_i}\right)\left(e^{y_j} + e^{-y_j}\right) = \sum_{i,j=1}^{n} \sum_{s,t \in -1,1} e^{sLy_i + ty_j}, \qquad L > 1. \tag{8}$$

In the factorized representation, $e^{Ly_i} + e^{-Ly_i}$ provides rapid growth when $|y_i|$ is large, while $e^{y_j} + e^{-y_j}$ transfers this growth into curvature in the e_j-direction. Summing over all i, j ensures that, regardless of which coordinate of y is large, Ψ_L provides growing curvature in every direction.

The properties used in both reductions are summarized below.

Lemma 4.1 (Globalizing auxiliary)**.** *The function Ψ_L is positive, globally convex, and globally log-convex. Writing $R = \|y\|_\infty$, and letting M_{Ψ_L} denote the log-curvature matrix of Ψ_L, the smallest eigenvalues of its Hessian and log-curvature matrix satisfy $\lambda_{\min}(\nabla^2 \Psi_L(y)) \ge 2e^{(L-1)R}$, $\lambda_{\min}(M_{\Psi_L}(y)) \ge 4e^{2LR}$. Moreover, there exist polynomials $P_1(n, L)$ and $P_2(n, L)$ such that $\Psi_L(y)$, $\|\nabla \Psi_L(y)\|_2$, and $\|\nabla^2 \Psi_L(y)\|_2$ are bounded above by $P_1(n, L)e^{(L+1)R}$, while $\|M_{\Psi_L}(y)\|_2 \le P_2(n, L)e^{2(L+1)R}$.*

The lower bounds allow the auxiliary term to dominate the curvature of q outside the unit box, while the upper bounds allow its coefficient to be chosen small enough to preserve negative curvature inside the inner box. The proof and explicit expressions for P_1 and P_2 are given in Appendix B.1.

We first apply the globalization construction to log-coordinate convexity. We give the main argument below and defer the complete quantitative proof to Appendix B.2.

Theorem 4.2. *The unconstrained* LOG-COORDINATE-CONVEXITY *problem is strongly NP-hard.*

Proof sketch. Consider $Q_L = q + \delta_L \Psi_L$. In the negative alternative, $\kappa_{[-\frac{1}{2},\frac{1}{2}]^n}(q) \le -\gamma$. On $[-\frac{1}{2}, \frac{1}{2}]^n$, Lemma 4.1 gives $\|\nabla^2 \Psi_L\|_2 \le P_1(n, L)e^{(L+1)/2}$. We therefore choose $\delta_L > 0$ such that $\delta_L P_1(n, L)e^{(L+1)/2} \le \frac{\gamma}{2}$. By the curvature perturbation bound, $\kappa_{[-\frac{1}{2},\frac{1}{2}]^n}(Q_L) \le -\frac{\gamma}{2}$, so Q_L remains nonconvex.

In the positive alternative, $\kappa_{[-1,1]^n}(q) \ge \gamma$. Since Ψ_L is convex, Q_L is convex on the unit box. For $R = \|y\|_\infty \ge 1$, the exponent bound $D \le 2$ gives $\|\nabla^2 q(y)\|_2 \le C_2 e^{DR}$ for a polynomially bounded constant C_2. Lemma 4.1 and Weyl's inequality then give

$$\lambda_{\min}\left(\nabla^2 Q_L(y)\right) \ge 2\delta_L e^{(L-1)R} - C_2 e^{DR}.$$

The ratio of the auxiliary curvature term to the adverse curvature term is $2\delta_L e^{(L-1)R}/(C_2 e^{DR}) = (2\delta_L/C_2)e^{(L-1-D)R}$. When $L-1 > D$, this ratio increases with R. Hence it suffices to make the preceding lower bound nonnegative at $R = 1$. The quantitative bounds show that an integer $L = O(\log n)$ and a positive rational δ_L of inverse-polynomial magnitude can simultaneously preserve the negative alternative and ensure $\lambda_{\min}\big(\nabla^2 Q_L(y)\big) \geq 0$ whenever $\|y\|_\infty \geq 1$.

Thus, Q_L is convex on $\mathbb{R}^n$ in the positive alternative and remains nonconvex in the negative alternative. Since Q_L is the exponential form of a signomial and L and δ_L have polynomial encoding length, this gives a polynomial-time reduction from GQGC. Strong NP-hardness follows. □

We next adapt the globalization construction to log-log convexity, where curvature is governed by the log-curvature matrix M rather than the Hessian. We give the main argument below and defer the complete quantitative proof to Appendix B.3.

Theorem 4.3. *The unconstrained* LOG-LOG-CONVEXITY *problem is strongly NP-hard.*

Proof sketch. Starting from the same exponential polynomial q, choose $\varepsilon > 0$ as in the compact log-log reduction and define $g_0 = 1 + \varepsilon q$. The choice of ε ensures that $g_0 \geq \frac{1}{2}$ on the unit box in both alternatives. In the positive alternative, g_0 is log-convex there; in the negative alternative, it retains a strict M_{g_0}-curvature witness in $[-\frac{1}{2}, \frac{1}{2}]^n$. We augment g_0 by setting $G_L = g_0 + \delta_L \Psi_L$.

Unlike the Hessian, $M_g = g\nabla^2 g - \nabla g \nabla g^\top$ is nonlinear under addition. For twice continuously differentiable functions u and v, define $\mathcal{C}(u, v) = u\nabla^2 v + v\nabla^2 u - \nabla u \nabla v^\top - \nabla v \nabla u^\top$. A direct expansion gives $M_{G_L} = M_{g_0} + \delta_L \mathcal{C}(g_0, \Psi_L) + \delta_L^2 M_{\Psi_L}$. For a fixed L, we first choose $\delta_L > 0$ sufficiently small that the last two terms preserve the negative eigenvalue of M_{g_0} at the interior witness. Hence, G_L remains non-log-convex in the negative alternative. In the positive alternative, both g_0 and Ψ_L are positive and log-convex on the unit box. Since positive sums of log-convex functions are log-convex, G_L is log-convex there.

Now let $R = \|y\|_\infty \geq 1$. The growth bounds for q, together with Lemma 4.1, give polynomially bounded constants A_0 and A_1 such that $\|M_{g_0}(y)\|_2 \leq A_0 e^{2DR}$, $\|\mathcal{C}(g_0, \Psi_L)(y)\|_2 \leq A_1 e^{(D+L+1)R}$. Using the lower bound on M_{Ψ_L}, we obtain

$$\lambda_{\min}(M_{G_L}(y)) \geq 4\delta_L^2 e^{2LR} - A_0 e^{2DR} - \delta_L A_1 e^{(D+L+1)R}.$$

When $L > D + 1$, the positive term grows faster than both potentially negative terms. It therefore suffices to make this lower bound nonnegative at $R = 1$. The quantitative estimates show that an integer $L = O(\log n)$ and a positive rational δ_L of inverse-polynomial magnitude can be chosen to satisfy this condition while preserving the interior negative-curvature witness.

It remains to establish global positivity in both alternatives. For some polynomially bounded constant C_0, we have $G_L(y) \geq 1 - \varepsilon C_0 e^{DR} + 2\delta_L e^{LR}$. Because $L > D$, the auxiliary term grows faster than the potentially negative contribution from q, so positivity outside the unit box can likewise be enforced at $R = 1$. The same L and δ_L chosen above also satisfy this condition. Together with $g_0 \geq \frac{1}{2}$ and $\Psi_L > 0$ on the unit box, this proves that $G_L > 0$ on $\mathbb{R}^n$ in both alternatives.

Consequently, G_L is globally log-convex in the positive alternative and remains positive but non-log-convex in the negative alternative. Since G_L is the exponential form of a signomial and L, δ_L have polynomial encoding length, this gives a polynomial-time reduction from GQGC, proving strong NP-hardness. □

5 CONCLUSION

We established that recognizing convexity, log-coordinate convexity, and log-log convexity of signomials is strongly NP-hard, both on compact domains and globally. Our reductions transfer inverse-polynomial curvature gaps from low-degree polynomials through exponential approximations and use auxiliary exponential terms to extend this separation beyond a compact region without destroying an interior nonconvexity witness. Thus, although these properties can enable tractable optimization once certified, no polynomial-time method can recognize them for arbitrary signomials unless P = NP. These worst-case results leave room for efficient methods on structured subclasses and motivate sound, efficiently verifiable certificates and systematically improvable approximation hierarchies for recognizing convex structure in signomial models.

AI USE STATEMENT

In this work, we used generative AI tools to support literature discovery and improve the presentation of the manuscript. Specifically, these tools were used to polish the writing and organization, suggest clearer wording for author-developed proof explanations, and flag potential issues in algebra, notation, and cross-references. They were not used to formulate the research questions, devise the reductions, or establish the mathematical results. The authors developed the proofs and independently verified all final definitions, derivations, claims, and cited sources. The authors take full responsibility for the content of the paper.

REFERENCES

Akshay Agrawal, Steven Diamond, and Stephen Boyd. Disciplined geometric programming. *Optimization Letters*, 13(5):961–976, 2019.

Amir Ali Ahmadi and Georgina Hall. On the complexity of detecting convexity over a box. *Mathematical Programming*, 182(1):429–443, 2020.

Amir Ali Ahmadi, Alex Olshevsky, Pablo A Parrilo, and John N Tsitsiklis. NP-hardness of deciding convexity of quartic polynomials and related problems. *Mathematical Programming*, 137(1): 453–476, 2013.

Rajendra Bhatia. *Matrix analysis*. Springer Science & Business Media, 2013.

Kaj-Mikael Björk and Tapio Westerlund. Global optimization of heat exchanger network synthesis problems with and without the isothermal mixing assumption. *Computers & Chemical Engineering*, 26(11):1581–1593, 2002.

Stephen Boyd, Seung-Jean Kim, Lieven Vandenberghe, and Arash Hassibi. A tutorial on geometric programming. *Optimization and Engineering*, 8(1):67–127, 2007.

Venkat Chandrasekaran and Parikshit Shah. Relative entropy relaxations for signomial optimization. *SIAM Journal on Optimization*, 26(2):1147–1173, 2016.

Mung Chiang, Chee Wei Tan, Daniel P Palomar, Daniel O'Neill, and David Julian. Power control by geometric programming. *IEEE Transactions on Wireless Communications*, 6(7):2640–2651, 2007.

Singa Wang Chiu, Chia-Kuan Ting, and Yuan-Shyi Peter Chiu. Optimal production lot sizing with rework, scrap rate, and service level constraint. *Mathematical and Computer Modelling*, 46(3-4): 535–549, 2007.

Mareike Dressler and Riley Murray. Algebraic perspectives on signomial optimization. *SIAM Journal on Applied Algebra and Geometry*, 6(4):650–684, 2022.

Cédric Févotte and Jérôme Idier. Algorithms for nonnegative matrix factorization with the β-divergence. *Neural Computation*, 23(9):2421–2456, 2011.

Michael Grant, Stephen Boyd, and Yinyu Ye. Disciplined convex programming. In *Global optimization: From theory to implementation*, pp. 155–210. Springer, 2006.

Kyoungmi Hwang, Kyungsik Lee, Chungmok Lee, and Sungsoo Park. Multi-class classification using a signomial function. *Journal of the Operational Research Society*, 66(3):434–449, 2015.

Cody J Karcher. Data fitting with signomial programming compatible difference of convex functions. *Optimization and Engineering*, 24(2):973–987, 2023.

Philippe G Kirschen, Martin A York, Berk Ozturk, and Warren W Hoburg. Application of signomial programming to aircraft design. *Journal of Aircraft*, 55(3):965–987, 2018.

Kyungsik Lee, Norman Kim, and Myong K Jeong. The sparse signomial classification and regression model. *Annals of Operations Research*, 216(1):257–286, 2014.

Chen Ling, Jiawang Nie, Liqun Qi, and Yinyu Ye. Biquadratic optimization over unit spheres and semidefinite programming relaxations. *SIAM Journal on Optimization*, 20(3):1286–1310, 2010.

Adia Lumadjeng, Ilker Birbil, and Erman Acar. ECSEL: Explainable classification via signomial equation learning. *arXiv preprint arXiv:2601.21789*, 2026.

Costas D Maranas and Christodoulos A Floudas. Global optimization in generalized geometric programming. *Computers & Chemical Engineering*, 21(4):351–369, 1997.

Alberto Marin-Sanguino, Eberhard O Voit, Carlos Gonzalez-Alcon, and Nestor V Torres. Optimization of biotechnological systems through geometric programming. *Theoretical Biology and Medical Modelling*, 4(1):38, 2007.

Theodore S Motzkin and Ernst G Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.

Riley Murray, Venkat Chandrasekaran, and Adam Wierman. Newton polytopes and relative entropy optimization. *Foundations of Computational Mathematics*, 21(6):1703–1737, 2021a.

Riley Murray, Venkat Chandrasekaran, and Adam Wierman. Signomial and polynomial optimization via relative entropy and partial dualization. *Mathematical Programming Computation*, 13(2): 257–295, 2021b.

Katta G Murty and Santosh N Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

Luis A San-José, Manuel González-De-la Rosa, Joaquín Sicilia, and Jaime Febles-Acosta. An inventory model for multiple items assuming time-varying demands and limited storage. *Optimization Letters*, 16(6):1935–1961, 2022.

Silviu-Marian Udrescu and Max Tegmark. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020.

Allen Houze Wang, Priyank Jaini, Yaoliang Yu, and Pascal Poupart. A Positivstellensatz for conditional SAGE signomials. *arXiv preprint arXiv:2003.03731*, 2020.

Julian Zimmert and Yevgeny Seldin. Tsallis-INF: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22(28):1–49, 2021.

A Proofs for Section 2

A.1 Global Quartic Gap-Convexity (GQGC) Hardness

Proof. We combine the biquadratic hardness construction of Ling et al. (2010) with the quartic convexity reduction of Ahmadi et al. (2013, Theorem 2.3).

Let (G, K) be an instance of Independent Set, where G has m vertices and $2 \leq K \leq m$, and set $\tau = K - 1$. The construction below has ambient dimension $n = 2m$. Let $\bar{G}$ denote the complement of G. Following Ling et al. (2010), define

$$b_{\bar{G}}(x; y) = -2 \sum_{(i,j) \in E(\bar{G})} x_i x_j y_i y_j.$$

Let $\mathrm{ind}(G)$ denote the independence number of G. By the Motzkin–Straus identity (Motzkin & Straus, 1965), $\min_{\|x\|_2 = \|y\|_2 = 1} b_{\bar{G}}(x; y) = -1 + 1/\mathrm{ind}(G)$.

Now set

$$b(x; y) = \tau b_{\bar{G}}(x; y) + (\tau - 1)\|x\|_2^2 \|y\|_2^2.$$

On the product of unit spheres, $\min b = \tau/\mathrm{ind}(G) - 1$. Hence, if $\mathrm{ind}(G) < K$, then $\mathrm{ind}(G) \leq \tau$, so $b \geq 0$ on the product of unit spheres and therefore globally, by biquadratic homogeneity. If $\mathrm{ind}(G) \geq K = \tau + 1$, there instead exist unit vectors $\bar{x}, \bar{y}$ such that $b(\bar{x}; \bar{y}) \leq -1/K$.

We next apply the quartic construction of Ahmadi et al. (2013, Theorem 2.3). Let $C_b(x, y)$ be the matrix of mixed second derivatives of b, and let Λ_b be the largest absolute coefficient appearing in its entries. Define

$$P_b(x, y) = b(x; y) + \frac{m^2 \Lambda_b}{2} \left(\sum_{i=1}^{m} x_i^4 + \sum_{i=1}^{m} y_i^4 + \sum_{i<j} x_i^2 x_j^2 + \sum_{i<j} y_i^2 y_j^2 \right).$$

Their construction shows that b is globally nonnegative if and only if P_b is globally convex.

Consider first the case $\mathrm{ind}(G) \geq K$. Let $w_0 = (\bar{x}, 0)$ and $d = (0, \bar{y})$, so that $\|d\|_2 = 1$. The quartic construction gives

$$d^\top \nabla^2 P_b(w_0) d = 2b(\bar{x}; \bar{y}) \leq -\frac{2}{K}.$$

Since P_b is homogeneous of degree four, its Hessian is homogeneous of degree two. Thus, with $\widehat{w} = w_0/2 \in [-1/2, 1/2]^n$,

$$d^\top \nabla^2 P_b(\widehat{w}) d \leq -\frac{1}{2K}.$$

To create a uniform positive margin in the other alternative, define

$$p(w) = P_b(w) + \frac{1}{8K} \|w\|_2^2.$$

If $\mathrm{ind}(G) < K$, then P_b is globally convex, and therefore

$$\nabla^2 p(w) \succeq \frac{1}{4K} I_n \qquad \text{for every } w \in \mathbb{R}^n.$$

If $\mathrm{ind}(G) \geq K$, then at the interior witness $\widehat{w}$,

$$d^\top \nabla^2 p(\widehat{w}) d \leq -\frac{1}{2K} + \frac{1}{4K} = -\frac{1}{4K}.$$

By Rayleigh–Ritz, $\lambda_{\min}(\nabla^2 p(\widehat{w})) \leq -1/(4K)$.

Thus the two alternatives have the common margin $\gamma_+ = \gamma_- = 1/(4K)$, which has inverse-polynomial magnitude because $K \leq m$. Moreover, p has degree at most four, contains only polynomially many variables and monomials, and has rational coefficients whose numerators and denominators are polynomially bounded in the input size. Hence p has polynomial encoding length and can be constructed in polynomial time. It follows that GQGC is strongly NP-hard. □

A.2 Proof of Lemma 2.2

Proof. Write

$$p(\xi) = \sum_{|\alpha| \leq d} p_\alpha \xi^\alpha, \qquad S_p = \sum_{|\alpha| \leq d} |p_\alpha|.$$

We first verify the exponential representation. For each multi-index α,

$$\Phi_h(y)^\alpha = h^{-|\alpha|} \prod_{i=1}^{n} (e^{hy_i} - 1)^{\alpha_i} = h^{-|\alpha|} \sum_{0 \leq \ell \leq \alpha} \binom{\alpha}{\ell} (-1)^{|\alpha| - |\ell|} e^{h\ell^\top y}.$$

Since $|\alpha| \leq d$, summing over the monomials of p gives $q_h(y) = \sum_{|\ell| \leq d} c_\ell e^{h\ell^\top y}$ for rational coefficients c_ℓ.

We next bound the Hessian perturbation. On $[-2, 2]^n$, termwise differentiation gives $\|\nabla p(\xi)\|_\infty \leq G_p$ and $\|\nabla^2 p(\xi)\|_2 \leq H_p$, where $G_p = d 2^d S_p$ and $H_p = n d^2 2^d S_p$. Moreover, every third partial derivative is bounded in absolute value by $d^3 2^d S_p$. Applying the fundamental theorem of calculus along the line segment joining $\xi, \zeta \in [-2, 2]^n$, together with $\|A\|_2 \leq \|A\|_F \leq n \max_{i,j} |A_{ij}|$, gives

$$\|\nabla^2 p(\xi) - \nabla^2 p(\zeta)\|_2 \leq L_p \|\xi - \zeta\|_\infty, \qquad L_p = n^2 d^3 2^d S_p.$$

For $|t| \leq 1$ and $0 < h \leq \frac{1}{2}$, $|\phi_h(t) - t| \leq h$, $|\phi_h'(t) - 1| \leq 2h$, and $|\phi_h''(t)| \leq 2h$. Hence, for $y \in [-1, 1]^n$, $\|\Phi_h(y) - y\|_\infty \leq h$, $\Phi_h(y) \in [-2, 2]^n$, and, with $J_h(y) = \mathrm{Diag}(e^{hy_1}, \ldots, e^{hy_n})$, $\|J_h(y) - I_n\|_2 \leq 2h$ and $\|J_h(y)\|_2 \leq 2$.

Using the decomposition in Equation 3, the Jacobian error satisfies

$$\begin{aligned}&\left\|J_h^\top \nabla^2 p(\Phi_h) J_h - \nabla^2 p(\Phi_h)\right\|_2 \\ &\quad \le \|J_h - I_n\|_2 \|\nabla^2 p(\Phi_h)\|_2 \|J_h\|_2 + \|\nabla^2 p(\Phi_h)\|_2 \|J_h - I_n\|_2 \le 6H_p h.\end{aligned}$$

The argument error is at most $L_p\|\Phi_h(y) - y\|_\infty \le L_p h$. Since the remaining term in Equation 3 is diagonal, its spectral norm is at most $2G_p h$. Therefore

$$\|\nabla^2 q_h(y) - \nabla^2 p(y)\|_2 \le (6H_p + L_p + 2G_p)h$$

uniformly on $[-1,1]^n$. Thus we may take $C_{\mathrm{err}}(p) = 6H_p + L_p + 2G_p$, which is rational, computable in polynomial time from p, and polynomially bounded in the input size.

It remains to verify the encoding claim. Since d is fixed, the expansion above contains only $O(n^d)$ exponent vectors. Each coefficient c_ℓ is a finite sum of terms involving p_α, binomial coefficients, and powers $h^{-|\alpha|}$. Hence, if h has polynomially bounded rational data, then so do the coefficients c_ℓ and exponents $h\ell$. Thus q_h satisfies the required encoding bounds. □

B Proofs for Section 4

B.1 Proof of Lemma 4.1

Proof. Positivity is immediate. Each summand of Ψ_L is the exponential of an affine function and is therefore convex, so Ψ_L is globally convex. Moreover, $\log \Psi_L$ is a log-sum-exp of affine functions, hence Ψ_L is globally log-convex.

Fix $y \in \mathbb{R}^n$ and write $R = \|y\|_\infty$. Choose an index i and a sign $s \in \{-1, 1\}$ such that $sy_i = R$. For arbitrary $z \in \mathbb{R}^n$ and any j, pairing the two terms corresponding to $t = \pm 1$ gives

$$\begin{aligned}&e^{sLy_i + y_j}(sLz_i + z_j)^2 + e^{sLy_i - y_j}(sLz_i - z_j)^2 \\ &\quad \ge e^{LR - |y_j|}\left[(sLz_i + z_j)^2 + (sLz_i - z_j)^2\right] \\ &\quad = 2e^{LR-|y_j|}\left(L^2 z_i^2 + z_j^2\right) \\ &\quad \ge 2e^{(L-1)R}\left(L^2 z_i^2 + z_j^2\right),\end{aligned}$$

where the last inequality uses $|y_j| \le R$. Summing over j and discarding the nonnegative terms involving $L^2 z_i^2$ yields

$$z^\top \nabla^2 \Psi_L(y) z \ge 2e^{(L-1)R}\|z\|_2^2.$$

Since it holds for every z, we obtain $\nabla^2 \Psi_L(y) \succeq 2e^{(L-1)R} I_n$, thus $\lambda_{\min}\left(\nabla^2 \Psi_L(y)\right) \ge 2e^{(L-1)R}$.

We next record the required upper bounds. Every exponent vector appearing in Ψ_L has ℓ_1-norm at most $L+1$, and there are $4n^2$ summands. Since the absolute value of every exponent is at most $(L+1)R$, termwise differentiation gives $\Psi_L(y) \le 4n^2 e^{(L+1)R}$, $\|\nabla \Psi_L(y)\|_2 \le 4n^2(L+1)e^{(L+1)R}$, and $\|\nabla^2 \Psi_L(y)\|_2 \le 4n^2(L+1)^2 e^{(L+1)R}$. Thus the first group of upper bounds in the lemma holds, for example, with $P_1(n, L) = 4n^2(L+1)^2$.

It remains to bound the log-convexity matrix M_{Ψ_L}. For $g(y) = \sum_a c_a e^{a^\top y}$, direct expansion of $M_g = g\nabla^2 g - \nabla g \nabla g^\top$ gives

$$M_g(y) = \frac{1}{2}\sum_{a,b} c_a c_b e^{(a+b)^\top y}(a-b)(a-b)^\top.$$

For each j, the support of Ψ_L contains the two exponent vectors $sLe_i + e_j$ and $sLe_i - e_j$. Their difference is $2e_j$, while their sum has inner product $2LR$ with y. The two ordered pairs formed from these vectors therefore contribute $4e^{2LR} e_j e_j^\top$ to $M_{\Psi_L}(y)$. Summing over j gives

$$M_{\Psi_L}(y) \succeq 4e^{2LR} I_n,$$

and hence $\lambda_{\min}\left(M_{\Psi_L}(y)\right) \ge 4e^{2LR}$.

Finally, using $\|M_{\Psi_L}\|_2 \le \Psi_L\|\nabla^2 \Psi_L\|_2 + \|\nabla \Psi_L\|_2^2$ together with the bounds above gives $\|M_{\Psi_L}(y)\|_2 \le 32n^4(L+1)^2 e^{2(L+1)R}$. Hence the second upper bound holds with $P_2(n, L) = 32n^4(L+1)^2$. □

B.2 Proof of Theorem 4.2

Proof. Let $q(y) = \sum_{|\ell|\le 4} c_\ell e^{h\ell^\top y}$ be the hard exponential form with curvature margin $\gamma > 0$. Recall that $D = \max_{\alpha\in\mathrm{supp}(q)} \|\alpha\|_1 \le 2$. Let $C_2 = \sum_{|\ell|\le 4} |c_\ell| \|h\ell\|_2^2$. Then, writing $R = \|y\|_\infty$, $\|\nabla^2 q(y)\|_2 \le C_2 e^{DR}$.

We construct $Q_L(y) = q(y) + \delta_L \Psi_L(y)$, where Ψ_L is the auxiliary function from Lemma 4.1. We first choose δ_L to preserve the interior negative-curvature witness, and then choose L so that the auxiliary curvature dominates outside the unit box.

On $[-\frac{1}{2}, \frac{1}{2}]^n$, Lemma 4.1 and the explicit bound from its proof give

$$\|\nabla^2 \Psi_L(y)\|_2 \le 4n^2(L+1)^2 e^{(L+1)/2}.$$

Using $\sqrt{e} < 5/3$, let

$$U_L = 4n^2(L+1)^2 \left(\frac{5}{3}\right)^{L+1}, \qquad \delta_L = \frac{\gamma}{2U_L}.$$

Then $\delta_L \|\nabla^2 \Psi_L(y)\|_2 \le \gamma/2$ throughout the witness box.

In the negative alternative, let $\widehat{y} \in [-\frac{1}{2}, \frac{1}{2}]^n$ and a unit vector z satisfy $z^\top \nabla^2 q(\widehat{y}) z \le -\gamma$. The preceding choice of δ_L gives

$$z^\top \nabla^2 Q_L(\widehat{y}) z = z^\top \nabla^2 q(\widehat{y}) z + \delta_L z^\top \nabla^2 \Psi_L(\widehat{y}) z \le -\gamma + \delta_L \|\nabla^2 \Psi_L(\widehat{y})\|_2 \le -\frac{\gamma}{2}.$$

Thus the strict negative-curvature witness is preserved.

Now consider the positive alternative. On $[-1, 1]^n$, $\nabla^2 q(y) \succeq \gamma I_n$, and Ψ_L is globally convex, so Q_L is convex on the unit box. For $R \ge 1$, $\|\nabla^2 q(y)\|_2 \le C_2 e^{DR}$ implies $\nabla^2 q(y) \succeq -C_2 e^{DR} I_n$, while Lemma 4.1 gives $\nabla^2 \Psi_L(y) \succeq 2e^{(L-1)R} I_n$. Hence

$$\nabla^2 Q_L(y) \succeq \left(2\delta_L e^{(L-1)R} - C_2 e^{DR}\right) I_n.$$

We choose L so that the auxiliary term already dominates at $R = 1$. Let $L \ge 4$ be the smallest integer satisfying

$$\left(\frac{8}{5}\right)^L \ge \frac{160 n^2 (L+1)^2 C_2}{\gamma}.$$

Such an L exists because the left-hand side grows exponentially, whereas the dependence on L on the right is polynomial. For the source family here, C_2/γ is polynomially bounded in n, so $L = O(\log n)$.

Indeed, substituting the choices of U_L and δ_L, and using $8/3 < e < 3$, gives

$$2\delta_L e^{L-1} = \frac{\gamma e^{L-1}}{4n^2(L+1)^2(5/3)^{L+1}} = \frac{3\gamma}{20e\, n^2(L+1)^2}\left(\frac{3e}{5}\right)^L > \frac{\gamma}{20n^2(L+1)^2}\left(\frac{8}{5}\right)^L \ge 8C_2.$$

Since $D \le 2$ and $e^2 < 8$, this implies $2\delta_L e^{L-1} > C_2 e^D$. Moreover, $L \ge 4$ and $D \le 2$ give $L - 1 - D > 0$, so the ratio $2\delta_L e^{(L-1)R}/(C_2 e^{DR})$ increases with R. As it exceeds one at $R = 1$, we obtain $\nabla^2 Q_L(y) \succeq 0$ for every $R \ge 1$.

Consequently, Q_L is globally convex in the positive alternative and remains nonconvex in the negative alternative. Since Q_L is the exponential form of a signomial, these alternatives are exactly global log-coordinate convexity and nonconvexity.

Finally, $L = O(\log n)$, so the auxiliary exponents and the $4n^2$ added terms satisfy the encoding bounds. Moreover, $\delta_L = \gamma/(2U_L)$ has polynomially bounded rational data. Hence the reduction preserves strong NP-hardness. □

B.3 Proof of Theorem 4.3

Proof. Let $q(y) = \sum_{|\ell|\le 4} c_\ell e^{h\ell^\top y}$ be the hard exponential form with curvature margin $\gamma > 0$. Recall that $D = \max_{\alpha\in\mathrm{supp}(q)} \|\alpha\|_1 \le 2$. We first apply the same scaling and offset used in the

compact log-log reduction. Let $C_q = \sum_{|\ell|\le 4} |c_\ell|$. Since $h \le \frac{1}{2}$, $|\ell| \le 4$, and $\|y\|_\infty \le 1$, we have $|q(y)| \le 8C_q$ and $\|\nabla q(y)\|_2 \le 16C_q$ on the unit box. Let $B_0 = 1 + 8C_q$, $B_1 = 1 + 16C_q$, and choose $\varepsilon = \min\{1/(2B_0), \gamma/(4B_1^2)\}$. Set $g_0(y) = 1 + \varepsilon q(y)$. Then $g_0 \ge \frac{1}{2}$ on the unit box and

$$M_{g_0} = \varepsilon g_0 \nabla^2 q - \varepsilon^2 \nabla q \nabla q^\top.$$

In the positive alternative, $M_{g_0}(y) \succeq (\varepsilon\gamma/4) I_n$ throughout the unit box. In the negative alternative, at the witness $\widehat{y}$,

$$z^\top M_{g_0}(\widehat{y}) z \le -\frac{\varepsilon\gamma}{2}.$$

Let $\nu = \min\{1, \varepsilon\gamma/2\}$. Thus g_0 is positive and log-convex on the unit box in the positive alternative, while in the negative alternative it has a strict M_{g_0}-curvature witness of margin at least ν. Since ε and γ have inverse-polynomial magnitude, so does ν.

Write $g_0(y) = \sum_{|\ell|\le 4} b_\ell e^{h\ell^\top y}$, absorbing the constant term into b_0. Let $K_0 = \sum |b_\ell|$, $K_1 = \sum |b_\ell| \|h\ell\|_1$, and $K_2 = \sum |b_\ell| \|h\ell\|_1^2$. For $R = \|y\|_\infty$, $|g_0(y)| \le K_0 e^{DR}$, $\|\nabla g_0(y)\|_2 \le K_1 e^{DR}$, and $\|\nabla^2 g_0(y)\|_2 \le K_2 e^{DR}$. Consequently, $\|M_{g_0}(y)\|_2 \le K_M e^{2DR}$, where $K_M = K_0 K_2 + K_1^2$.

Now define $G_L(y) = g_0(y) + \delta_L \Psi_L(y)$. For twice differentiable functions u, v, let $\mathcal{C}(u, v) = u\nabla^2 v + v\nabla^2 u - \nabla u \nabla v^\top - \nabla v \nabla u^\top$. Then

$$M_{G_L} = M_{g_0} + \delta_L \mathcal{C}(g_0, \Psi_L) + \delta_L^2 M_{\Psi_L}.$$

Using the bounds from Lemma 4.1 gives

$$\|\mathcal{C}(g_0, \Psi_L)(y)\|_2 \le K_L e^{(D+L+1)R},$$

where $K_L = 4n^2\big(K_0(L+1)^2 + K_2 + 2K_1(L+1)\big)$.

We first choose δ_L to preserve the interior witness. On $[-\frac{1}{2}, \frac{1}{2}]^n$, the bound $D \le 2$ gives $\|\mathcal{C}(g_0, \Psi_L)\|_2 \le K_L e^{(L+3)/2}$. Using $\sqrt{e} < 5/3$ and $e^{3/2} < 5$, let $V_L = 5K_L(5/3)^L$, so that the cross term is bounded by V_L throughout the witness box. On the same box, Lemma 4.1 gives $\|M_{\Psi_L}\|_2 \le 32n^4(L+1)^2 e^{L+1}$. Let $S_L = 10n^2(L+1)(5/3)^L$; then $\|M_{\Psi_L}\|_2 \le S_L^2$.

Set $T_L = \max\{1, V_L, S_L\}$ and $\delta_L = \nu/(8T_L)$. Then $\delta_L V_L \le \nu/8$ and $\delta_L^2 S_L^2 \le \nu^2/64 \le \nu/64$. At the negative witness,

$$z^\top M_{G_L}(\widehat{y}) z \le -\nu + \delta_L V_L + \delta_L^2 S_L^2 \le -\nu + \frac{\nu}{8} + \frac{\nu}{64} < -\frac{\nu}{2}.$$

Moreover, $G_L(\widehat{y}) > 0$ because $g_0(\widehat{y}) \ge \frac{1}{2}$ and $\Psi_L > 0$. Hence $z^\top \nabla^2 \log G_L(\widehat{y}) z < 0$, so the negative alternative is preserved.

It remains to choose L so that G_L is globally positive and, in the positive alternative, globally log-convex. Let $P_L = 1 + 5K_L + 10n^2(L+1)$. Since $T_L \le P_L(5/3)^L$, we have $\delta_L \ge \frac{\nu}{8P_L}(3/5)^L$.

We choose $L \ge 4$ to satisfy three requirements. The first makes the auxiliary term dominate g_0 outside the unit box, the second makes $\delta_L^2 M_{\Psi_L}$ dominate the cross term, and the third makes it dominate M_{g_0}. Concretely, choose L as the smallest integer satisfying

$$\left(\frac{8}{5}\right)^L \ge \frac{128 P_L K_0}{\nu} \qquad \text{(global positivity)},$$

$$\left(\frac{8}{5}\right)^L \ge \frac{84 P_L K_L}{\nu} \qquad \text{(cross-term domination)},$$

$$\left(\frac{64}{25}\right)^L \ge \frac{1760 P_L^2 K_M}{\nu^2} \qquad \text{(source-curvature domination)}.$$

Such an L exists because P_L and K_L grow only polynomially in L and n, whereas the left-hand sides grow exponentially. Hence $L = O(\log n)$.

We first verify global positivity. For $R \ge 1$, choose i, s with $sy_i = R$, and choose t so that $ty_j \ge 0$ for any fixed j. The corresponding summand of Ψ_L is at least e^{LR}, so $\Psi_L(y) \ge e^{LR}$. On the other

hand, $|g_0(y)| \le K_0 e^{DR}$. At $R = 1$, the lower bound on δ_L, the global-positivity condition above, and $3e/5 > 8/5$ give

$$\delta_L e^L \ge \frac{\nu}{8P_L}\left(\frac{3e}{5}\right)^L > \frac{\nu}{8P_L}\left(\frac{8}{5}\right)^L \ge 16K_0.$$

Since $D \le 2$ and $e^2 < 8$, we have $16K_0 > K_0 e^D$. Thus $\delta_L e^L > K_0 e^D$ at $R = 1$. Because $L > D$, the ratio $\delta_L e^{LR}/(K_0 e^{DR})$ increases with R, so $\delta_L e^{LR} > K_0 e^{DR}$ for every $R \ge 1$. Hence $G_L(y) > 0$ outside the unit box. Together with $g_0 \ge \frac{1}{2}$ and $\Psi_L > 0$ on the unit box, this proves that G_L is globally positive in both alternatives.

Now suppose the positive alternative holds. On the unit box, g_0 and Ψ_L are positive and log-convex, so their positive sum $G_L = g_0 + \delta_L \Psi_L$ is log-convex there.

For $R \ge 1$, the bounds above and Lemma 4.1 give

$$M_{G_L}(y) \succeq \left(4\delta_L^2 e^{2LR} - \delta_L K_L e^{(D+L+1)R} - K_M e^{2DR}\right) I_n.$$

We first compare the positive auxiliary term with the cross term. At $R = 1$,

$$\frac{4\delta_L^2 e^{2L}}{\delta_L K_L e^{D+L+1}} = \frac{4\delta_L e^{L-D-1}}{K_L} \ge \frac{\nu}{2P_L K_L e^{D+1}}\left(\frac{3e}{5}\right)^L > \frac{\nu}{42P_L K_L}\left(\frac{8}{5}\right)^L \ge 2,$$

where $D \le 2$, $e^3 < 21$, and the last step uses the cross-term domination condition.

Similarly, comparing the positive auxiliary term with M_{g_0} gives, at $R = 1$,

$$\frac{4\delta_L^2 e^{2L}}{K_M e^{2D}} \ge \frac{\nu^2}{16P_L^2 K_M e^{2D}}\left(\frac{3e}{5}\right)^{2L} > \frac{\nu^2}{880P_L^2 K_M}\left(\frac{64}{25}\right)^L \ge 2,$$

where $D \le 2$, $e^4 < 55$, and the last step uses the source-curvature domination condition.

Thus, at $R = 1$, the positive auxiliary term is at least twice each negative contribution. Moreover, $L \ge 4$ and $D \le 2$ imply $2L > D + L + 1$ and $2L > 2D$, so both ratios increase with R. Therefore, for every $R \ge 1$,

$$4\delta_L^2 e^{2LR} \ge 2\delta_L K_L e^{(D+L+1)R}, \qquad 4\delta_L^2 e^{2LR} \ge 2K_M e^{2DR}.$$

It follows that $M_{G_L}(y) \succeq 0$ throughout the exterior region. Combined with log-convexity on the unit box and the global positivity proved above, G_L is globally log-convex in the positive alternative.

Thus, G_L is positive on $\mathbb{R}^n$ in both alternatives. In the positive alternative, $\log G_L$ is convex on $\mathbb{R}^n$; in the negative alternative, the preserved interior witness shows that $\log G_L$ is not convex. Since G_L is the exponential form of a signomial, the corresponding signomial satisfies the positivity promise in both alternatives and distinguishes the two cases of the unconstrained LOG-LOG-CONVEXITY problem.

Finally, $L = O(\log n)$, while ε, ν, and δ_L have polynomially bounded rational data, and only $4n^2$ auxiliary terms are introduced. Hence the construction preserves the encoding bounds and strong NP-hardness. □